Sound Design for a System of 1000 Distributed Independent Audio-Visual Devices

Sam Ferguson
Creativity & Cognition Studios
Faculty of Engineering & IT
University of Technology
Sydney
samuel.ferguson@uts.edu.au

Anthony Rowe
Squidsoup.org
ant@squidsoup.org

Oliver Bown
Interactive Media Laboratory
UNSW Art and Design
o.bown@unsw.edu.au

Liam Birtles
Bournemouth University
lbirtles@bournemouth.ac.uk

Chris Bennewith
School of Art, Design & Architecture
Plymouth University, UK
chris.bennewith@plymouth.ac.uk

ABSTRACT

This paper describes the sound design for *Bloom*, a light and sound installation made up of 1000 distributed independent audio-visual pixel devices, each with RGB LEDs, Wifi, Accelerometer, GPS sensor, and sound hardware. These types of systems have been explored previously, but only a few systems have exceeded 30-50 devices and very few have included sound capability, and therefore the sound design possibilities for large systems of distributed audio devices are not well understood. In this article we describe the hardware and software implementation of sound synthesis for this system, and the implications for design of media for this context.

Author Keywords

distributed, internet of things, embedded audio

1. INTRODUCTION

This paper examines the use of sound in display systems that are made of a large number of independent distributed pixel devices. Research interest in distributed pixel systems is increasing but their implementation and use are still at an early stage, and much of the capability that has been shown is both low resolution and demonstrative in nature. Few, if any, systems have demonstrated the ability to produce sound and so research into sound design in this context has been limited.

The accelerating decline in cost of the devices that make up these systems have meant that the feasibility of the systems is no longer prohibitive, and this lower cost has also meant that hardware for other modalities can be included. As the 'Internet of Things' becomes more dominant [11], opportunities for deploying systems of this nature will become more frequent. Including other modalities, apart from the visual modality, will be a critical part of the evolution of these systems both for sensing capability and display purposes.

In this paper we report on the creation of a system made up of nearly 1000 distributed pixels all of which included audio capability and hardware, and which were connected together on a wireless network. Additionally, each device used a GPS sensor to be able to sense its physical location, and in accelerometer reported information about the motion. The system was installed at Kew Gardens in London for a period of six weeks, and was experienced by the public as part of the *Christmas at Kew* exhibition (Figure 1). This system is one of the first of its kind (that we know of) and therefore we encountered many challenges in designing for such a platform, especially as concerns sound design for the platform.

Figure 1: *Bloom* installed at Kew Gardens, London.

This system poses several opportunities, but also some new design problems of some significance for the NIME community. It mimics several theoretical/virtual approaches to musical composition and performance in a physical artefact that necessitates programming for a distributed mode of experience, rather than as a musical interface that maps musical intent to sound output [14,

NIME'17, May 15-19, 2017, Aalborg University Copenhagen, Denmark.

31]. This paper will focus on developing the design principles and approaches and will seek to contrast this design context to typical sound design frameworks, focusing on opportunities for artistic and musical expression.

1.1 Existing research

Many authors have reported on projects that developed distributed pixel systems and we will discuss a few examples of systems of independent devices which are coordinated to communicate media of some type. A system for integrating a set of independent devices into a display system capable of displaying images was described by Sato et al. [25, 26] – it used a camera to calibrate the position of each device and some careful choice of how to represent the figure. Similarly, but on a larger scale, Seitinger et al. [29] developed a set of 50 'urban pixels', which were distributed and independent of both power and network cables and were designed to be positioned on buildings. They were controlled by either and infrared flashlights or individual SMS control messages, allowing interaction with their audience. Barker et al. [1] and Hauesler et al. [12] focused on the development process and capabilities of firstly 'Janus' (a face-shaped facade of addressable pixels used to present facial expressions), and then 'polymedia pixels', a system of networked pixels. The work of SquidSoup, using 3-dimensional arrays of thousands of individually addressable LEDs, has shown the potential for volumetric displays to transform the experience of spaces [24, 23], but has also incorporated programmatic media design for novel installations. The *PushPins* project [7] was made up of a set of small layered computing devices that could be pushed into a 'power substrate' made up of two layers of conducting material – a type of conducting corkboard. They then formed a decentralised network to become a distributed display system, without a central controller. Paradiso et al. [20] went on to discuss *Tribble*, a system of tiled computers that formed a sensate sphere, and *Z-Tiles* which made up a sensate flooring system. Similarly, *Junkyard Jumbotron* [3] showed how everyday display systems (mobile phones, TVs, tablets) could be co-ordinated ad-hoc to become a composite display capable of displaying images on a larger scale. While *Junkyard Jumbotron* used fiducial markers to encode location, orientation and size on each display, *Phone as a Pixel* [28] achieved similar using colour transitions.

With independence and portability comes the possibility of interaction, with physical game based pieces being a particular area of motivation. Fischer et al. [10] designed independent devices with unique shapes that altered the way participants could interact with them, and in *Urban Musical Game* by Rasamimanana et al. [21] a ball was instrumented with an accelerometer, which drove sonification algorithms in Max/MSP. Similarly, Distributed Interactive Audio Devices [6] – distributed systems for interaction with audio synthesis capabilities based on Raspberry Pi devices – were designed to be physically handled and passed between participants that are part of an audience, and were also used to develop simple musical games employing physical interaction [4]. The musical potential of mobile phones has also been a focus of research, with the creation of several mobile phone ensembles – for instance Schiemer's [27] *pocket gamelan* work, and the work of Wang et al. [33, 19] in establishing the MoPho mobile phone orchestra. These research areas have been associated with interaction [8] and localisation research [15] to investigate better sensing capabilities that can be employed in such compositions.

Motility of displays is another area of investigation that is rapidly evolving. For instance Kuznetsov and co-authors developed *Wallbots*, a system of robots that could traverse metal walls and surfaces with their magnetic wheels, but also acted as display elements and have senses for interaction [16]. The *Spaxels* system of Ars Electronica Future Lab [13] has expanded to 100 quadrotor drones outfitted with RGB LEDs, which are able to co-ordinate movements to produce visual images and choreography in a three-dimensional public arena, with 'light-painting' being a particular capability they are able to demonstrated. On a smaller scale, the *Pixelbots* system of small coordinated mobile pixels that has been developed by Digurmati et al. [9], is an impressive display of how animation of pixel-based images could be implemented. Merrill's *Siftables* are computers within tiles that implemented interaction systems employing physical rearrangement interaction procedures as input to [17]. While not independently mobile, they used careful sensing of their physical position as a prime source of interaction input in their applications. Each of these examples show how distributed pixel systems, when each device has more output controls apart from simply a visual outputs, gain many more capabilities and configuration possibilities.

Many artists, such as Zimoun, Rafael Lozano-Hemmer, Nils Völker or Anthony Gormley have created artworks that employ the multiple as a central concern. For instance, Zimoun's sculptures, made up of cardboard boxes, DC motors and usually wire beaters that together produce an imposing ambient noise texture rely on the large number of elements in each work, and indeed many of the work's titles reference the number of elements involved. Völker's works, made up of groups of inflatable elements (usually plastic bags, but also Hoberman spheres) draw attention to the relationship between each of the elements, as they often seem to breathe together. Anonymous has proposed 'media multiplicities', as a term to describe these artworks and systems, which rely on the multiple acting in coordination [5]. Generally these systems seem to vary structurally along three main dimensions: a) heterogeneity vs. homogeneity, b) object vs. substrate, and c) composed vs. self-organised. Understanding how these systems have autonomy to rapidly shift along these axes helps decompose the possibilities that exist for design and analysis of works of this nature.

Ambient visualisation devices [18] are closely related to the rise of distributed pixel systems, as the devices performing the visualisation are extremely similar in nature to the pixels that make up a distributed pixel system. There has been a longer history of investigation of ambient visualisation, and Tomitsch et al. [30] summarised a proportion of this history in a taxonomy of different ways in which ambient devices communicated data. Tellingly, of the 18 systems that they investigated, only one of them used sound as a display output. This is symptomatic of the larger deficit in research on audio modality in new interfaces, with the visual modality dominating for the majority. However, where ambient computing devices have a range of interface possibilities, a distributed system of thousands of similar devices has many extra capabilities made possible by the co-operation of each of these devices.

Similarly the *Internet of Things* [11], made up of devices interacting with wireless sensor networks, is still rapidly developing, and many of the research areas it is made up of are under detailed investigation. The precipitous miniaturisation of power sources and of computer hardware that it has both benefited from and propelled, are a crucial contributor to the possibility for distributed pixel systems. However, in a typical Internet-of-Things-type scenario, where household appliances and devices coordinate and communicate their actions, it's hard to see how a typical visual display made up of a grid of visual pixels, might be configured and indeed, the visual modality relies on a high degree of attention. By comparison, the use of audio devices as part of a distributed audio display system would be quite feasible, and so it seems useful to consider sound design for these devices.

1.2 Computer Music & New Interfaces

Design considerations for distributed pixel systems echo a lot of the ideas of computer music composition and new interfaces for musical expression. Digital musical instruments and mapping techniques, spatial sound systems, granular synthesis, and wave field synthesis are some of the critical areas relevant to sound design for this system.

Mapping is a central concern for designers of musical instruments, and authors have discussed methods for conceptualising and

Figure 2: The devices and their cabling were required to be waterproof and durable, given the installation location and conditions and the 6-week installation period. The stands were manufactured from flexible plastic, and carried the DC 12V power supply that powers each device. Image Copyright 2016

designing mappings [32] as well as methods for assessing their effectiveness [14].

Spatial sound – the ability to position sounds in space and to move that sound's location through perceived space has been an artistic technique used by many sound designers and musicians, and much audio research has focused on standards and technical methods for doing so with realism and flexibility. Most of these methods situate a small set of loudspeakers in a circular arrangement around a seated listener, and imply a 'sweet spot'. By contrast, Wave Field Synthesis [2] is a technique whereby closely located arrays of loudspeakers are temporally synchronized so as to be able to physically recreate the acoustic waves that would be created by a physical acoustic source. Using this approach to mimic travelling sound waves emanating from particular virtual locations, the system can move the apparent source location around the room both behind the loudspeaker system, but also in front of the system in the physical space of the audience, and audience members do not need to sit in only one specific sweet spot. As we will discuss later, the difference between sound elements placed on the periphery of the space, and placed within the space is significant.

Granular synthesis is a process of developing sounds from thousands of small fragments, or grains, of sound. Each of these fragments is overlaid on top of each other in ways that can be parametrically controlled to achieve particular textures [22]. The grains themselves can be drawn from particular audio recordings, and the characteristics determine a great proportion of the type of sound that will be produced. In the example of this system, given the huge number of elements, and their co-located visual and auditory reproduction capability, the statistical characteristics and parameterisation shown in granular synthesis are invoked, in that tone pulses (rather than fragments of audio recordings) would be produced in such numbers that they could coalesce to produce a consistent textural sound.

1.3 Aims and Overview

In this paper we aim to:

- describe the development of a system of 1000 distributed pixels that includes the capability for audio output.
- investigate and describe the design considerations that come into play with large numbers of audio elements fused to visual pixels
- describe a systematic set of approaches to developing the software and media architecture for designing audio content in distributed pixel systems.

In the following section, we will carefully describe the distributed pixel system hardware developed for this installation, alongside the software and network design. We will also outline the sensor capability of the devices, and the approach we used for the network communication. Following this, we will focus on the Sound design and specifically the design possibilities that are enabled by specially distributed set of audio elements. Finally we'll discuss the Implications of this research, and outline future research possibilities.

2. IMPLEMENTATION

The installation was commissioned for a space within Kew Gardens, a garden in south east London. It was to be installed for 6 weeks in November and December, meaning cold weather and rain were environmental factors of importance (Figure 2). The devices were placed approximately 1.2 m from the ground on flexible stands, with 12V DC power cables run across the ground to each device. The space used was approximately 80 m x 25 m, bisected by a path along which the audience could pass through the space (and across which no cables were run).

The device was designed to be lightweight, low-power, network enabled and to create sound and light (Figure 3). It included several important capabilities:

Onboard Processing: The device used an ESP3266 type module (a Digistump Acorn) which included an over-the-air programmer and various other advanced capabilities.

Onboard WiFi Networking: One of the capabilities provided by the Acorn module was onboard wireless networking.

GPS: a compact GPS aerial and chipset was soldered onto the circuitboard and connected to the Acorn module directly.

Accelerometer: An ADXL345 accelerometer was included on the device, and connected to the Acorn module by an I^2C interface.

Audio amplifier: a simple audio signal amplifier was included whose gain was digitally controllable with high temporal resolution, through an I^2C interface.

Transducer: a contact-type transducer was glued to the device, so that sound produced resonated through the device housing.

LEDs: Two NeoPixel type RGB LEDs were positioned on either side of the device.

Open Sound Control was used to communicate with the devices, with each mode of operation being a separate command with control parameters specified as arguments. Global messages were sent to the devices, and each device interpreted the messages based on their local sensed information. The devices were connected by a WiFi network, and 4 separate Xirrus Networks access points were used to allow the 1000 devices to connect simultaneously, and to connect to the central server. The outdoor park location made for low levels of interference from other networks, and the space was not obscured by any buildings either. Each of the devices had independent computation capability, with a full WiFi networking stack and the capability of being reprogrammed over the network. The ease by which they could be programmed and reprogrammed meant that the way they responded to the network messages could be rapidly adjusted and redesigned in response to changing requirements. The devices were provided with an Arduino module to allow them to be programmed using the Arduino IDE, but with a server deployment scheme so as to allow a simple over-the-air update process. Iteration was quick enough that an update could be deployed in approximately 60-120 seconds.

The devices included sound capability, but it was of a rudimentary basis. In fact, the main capability was the use of the `tone()` command (a simple square-wave tone generator), alongside the

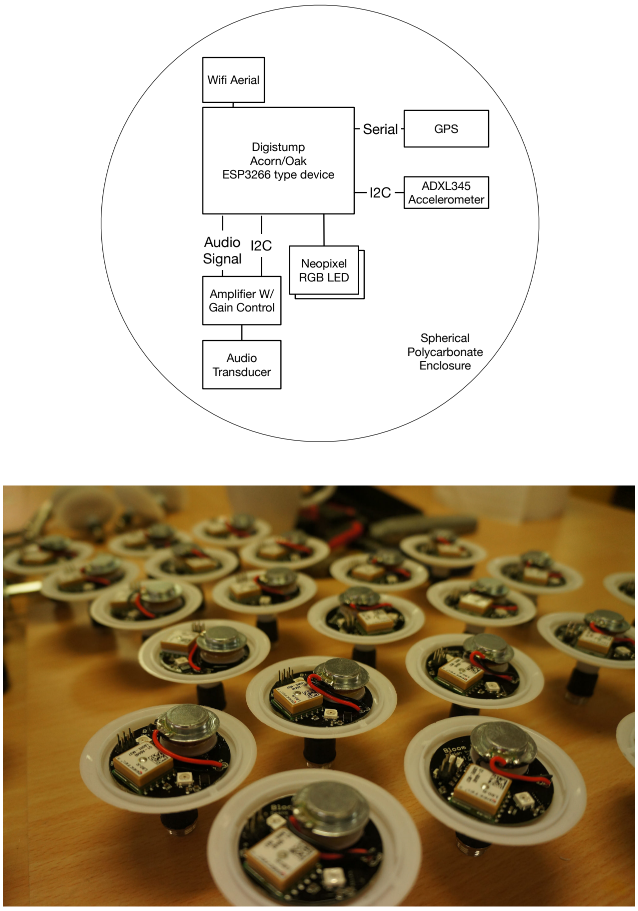

Figure 3: Diagram, and photograph of components of distributed pixel system with spherical enclosure covers removed. The GPS unit, the 2 NeoPixel LEDs and the audio transducer are clearly shown, while the Digistump Acorn unit was soldered to the underside of the circuitboard. Because of the size of the elements, the circuitboards were able to be mounted on the lower-most portion of the sphere, meaning that the rest of the sphere was not in shadow.

addition of fine temporal control of the amplifier gain. The audio transducer produced sound by transmitting vibrations to the circuitboard and through to the spherical cover, which was a more space efficient approach than trying to utilise a loudspeaker and associated mounting, and also neatly avoids shadowing. This is a highly unorthodox combination of synthesis (a digital tone generator combined with analogue amplifier with digital amplitude control), but for this context the limitations it provided still allowed for many design possibilities, which we discuss in section 3 below. Having fine-grained control over the audio amplitude allowed simple amplitude envelopes as well as longer-term ramping effects, meaning tones that are sometimes annoying when they are constant, could be modulated both over the duration of a note (100-1000 ms) and over the term of a *crescendo* or phrase (10-60 s).

The device included two types of sensors, an accelerometer (an ADXL345) for sensing orientation, and a GPS chipset for sensing location. Each sensor was set up to store data about the device location and motion locally, rather than broadcasting it to a central host. Each procedure, therefore, used this local information to alter its display output in response.

2.1 Design for *Bloom*

For Bloom, the sound was designed with a tightly synchronised approach – meaning that each visual pulse of brightness, was accompanied by an auditory pulse of the same temporal envelope. This tight temporal synchronization meant that the only degree of freedom between sound and visual mode was in the colours and pitches chosen as these were not linked. In design terms, however, this temporal synchronization meant that the visual design was easily paired to appropriate sound design, and the design process was greatly simplified.

The design used a sequenced series of global OSC network commands to trigger particular behaviours, which each ran for between 10 and 120 seconds. Each of these commands was associated with a set of parameters, which were then interpreted by the functions executed on each of the devices. These functions altered the behaviour based on either random selection, pre-assigned grouping values (group 1, 2, ...), or by the GPS location of the device. In this way, the tight temporal control required for audio processing and for rapid visual content was controlled by each pixel, but the group's was both synchronised and programmed by a central control mechanism. While high level control was the central controller's, much of the finer temporal and dynamic behaviour was retained by the pixels themselves.

3. SYNTHESIS AND SOUND DESIGN

For the purpose of sound design for 1000 or more elements, some concepts and terms used in typical spatial sound design may need to be rethought and reconsidered.

3.1 Spatial Sound Design

Spatial sound reproduction has focused on developing systems of reproduction designed for sound fields made up of moving and stationary sources at arbitrary locations. Sound formats (e.g. stereo, surround, ambisonic, wave field synthesis) etc. are engineered to record and encode a sound field that can be transmitted as audio signals to a reproduction system for a listener to experience that same sound field with as much fidelity as possible, as mediated by the human perceptual system. A significant user of these spatial sound formats have been designers of film sound, where the relationship between the visual content and audio can be considered either diegetic (fused to visual images on the screen) and non-diegetic (sound not fused to any visual content - for instance background music). In the context of multiplicitous independent distributed visual pixels such spatial sound concepts may need adjustment to take account of the changed nature of the context. Reproduction of sound fields usually pre-assumes that a small number (for instance 2 to 16 at most) of loudspeakers are positioned at the periphery of the space, facing inwards and implying a 'sweet spot' for a small number of stationary listeners positioned within. The system we have described positions the listener amongst a much larger number of sound reproduction devices within the space, positioned in a scatterable, dispersed fashion across the inhabited space of the installation, and the listener is able to reposition themselves arbitrarily.

3.2 Fusing Sound and Image in Space

Similarly, where a typical sound reproduction system may be partnered to a visual screen located close or far from the reproducing elements, in this context the distributed pixel is both sound and visual element, as the spherical physical shell performs both the light and sound diffusion. Although the devices are figurative in nature, and are not attempting to communicate an image in a typical filmic way, sound can be temporally fused to their visual stimulus, and listeners can focus on individual pixels, groups of pixels, or change (eg. colour, brightness) within a pixel. Whether or not this is diegetic sound in the traditional sense ('diegetic' as a term has been defined within the context of film), the design possibilities are analagous.

Furthermore, sound designers have rarely been able to employ movable, but co-located visual and auditory images. Wave field synthesis systems can move auditory images of sounds both behind or between loudspeakers (as can most systems) as well as in front of the sound reproduction elements and within and around the listener space. However, these systems cannot move the *visual* image, except within the considerably narrow space of typical visual display. They certainly cannot make a fused audio-visual image exist in arbitrary spatial locations within the space of the listener, even with overhead projection. But perceptually locating and tracking auditory images, without a concomitant visual image, is not easy and prone to breakdown of a subtle image. The system we have described, with its 1000 source locations for both visual and auditory stimuli, can show fused images given appropriate sound design.

3.3 Location Sensing and GPS

Few loudspeaker, or spatial format systems have included the capability of location sensing for the reproduction elements (although it is not completely novel). Access to this information for scatterable systems of sound reproduction radically expands spatial sound design possibilities. While GPS is clearly limited to the outdoors, as localization technologies of many more types become widespread, sound design that exploits reproduction element location will become easy tom implement in many spaces. Having pixel location information means that sound can not only be fused to a particular (but unknown) visual location, this source can be accurately located in space in a large number of places. In practice, the accuracy of location data is critical, as groupings and formations of pixels are only perceptible when they can be specified within appropriate limits.

3.4 Statistical Approaches to Sound

Given this system increases the number of sound reproducing elements by such a significant amount, one approach to sound design for the system is to treat the devices in a statistical manner, in such a way as granular synthesis does. Rather than sounding identical tones at identical times, small variations in timing, pitch or duration, when applied to 1000 devices simultaneously, produce a granular texture whose characteristics are altered by parameters such as fragment length, fragment frequency and tone pitch randomness. While, of course, the grains in this conception are not extracted from a source sound, as they are in typical granular synthesis, apart from this much of the design process is similar to the procedure of granular synthesis, except that the addition of tones takes place in the physical sound field, rather than in the digital realm.

4. CONCLUSIONS & FUTURE RESEARCH

This paper has presented an implementation of a distributed pixel system, which for the first time we are aware of, incorporated both sound capability and individual localization hardware. This system was deployed in a garden space and covered an area of approximately 18 m by 25 m, which was bisected by a pedestrian path. Given that there were 1000 devices across the space, and each of them could reproduce sound, new approaches to spatial sound design naturally followed.

After analysing the process of sound design for this system, it is clear that several elements of typical sound design need to be rethought: a) many of the fundamental assumptions of spatial sound design do not hold when sound reproduction devices are scattered across the space which the listener is expected to traverse rather than being restricted to the borders or corners b) the fusion of sound and image means that new approaches to designing sound, that aren't possible where auditory images and visual images are separated, are now possible c) when the numbers of devices are so high, a straightforward solution to a sound design is to take a statistical view of sound reproduction, leading to a granular type sound texture.

5. ACKNOWLEDGEMENTS

We acknowledge the design assistance of Digistump LLC, who designed the electronics and coordinated the manufacture of the devices. Networking support was generously provided by Xirrus Networks. *Bloom* was a Squidsoup project sponsored by the Royal Botanical Gardens, Kew, and curated by Culture Creative.

6. REFERENCES

[1] T. Barker, M. H. Haeusler, F. Maguire, and J. McDermott. Investigating political and demographic factors in crowd based interfaces. *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group on Design: Open 24/7 - OZCHI '09*, page 413, 2009.

[2] A. J. Berkhout. Acoustic control by wave field synthesis. *The Journal of the Acoustical Society of America*, 93(5):2764, 1993.

[3] R. Borovoy and B. Knep. Junkyard Jumbotron, 2011.

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

[5] O. Bown and S. Ferguson. Media Technology + the Internet of Things = Media Multiplicities. *Leonardo Transactions*, submitted, 2017.

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

[7] W. Butera. Text display and graphics control on a paintable computer. In *First International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2007*, pages 45–54, 2007.

[8] L. Dahl and G. Wang. Sound Bounce: Physical Metaphors in Designing Mobile Music Performance. *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 178–181, 2010.

[9] T. Digumarti, J. Alonso-Mora, R. Siegwart, and P. Beardsley. Pixelbots 2014. *Leonardo*, 49(4):366–367, 2016.

[10] P. T. Fischer, F. Gerlach, J. G. Acuna, D. Pollack, I. Schäfer, J. Trautmann, and E. Hornecker. Movable, Kick-/Flickable Light Fragments Eliciting Ad-hoc Interaction in Public Space. *Proceedings of The International Symposium on Pervasive Displays - PerDis '14*, pages 50–55, 2014.

[11] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami. Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7):1645–1660, 2013.

[12] H. Hauesler, T. Barker, and K. Beilharz. 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*, pages 1–7. Marmara University, 2010.

[13] H. Hörtner, M. Gardiner, R. Haring, C. Lindinger, and F. Berger. Spaxels, Pixels in Space. In *Proceedings of the International Conference on Signal Processing and Multimedia Applications and Wireless Information Networks and Systems*, pages 19–24, 2012.

[14] A. Hunt and M. Wanderley. Mapping performer parameters to synthesis engines. *Organised Sound*, 2002.

[15] H. S. Kim, J. Herrera, and G. Wang. Ping-Pong: Musically Discovering Locations. *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 273–276, 2014.

[16] S. Kuznetsov, E. Paulos, and M. D. Gross. WallBots: Interactive Wall-Crawling Robots In the Hands of Public Artists and Political Activists. In *Proceedings of the 8th ACM*

Conference on Designing Interactive Systems - DIS '10, page 208, 2010.

[17] D. Merrill, J. Kalanithi, and P. Maes. Siftables: Towards sensor network user interfaces. *Proceedings of the 1st international conference on Tangible and embedded interaction - TEI '07*, page 75, 2007.

[18] a. V. Moere. Towards Designing Persuasive Ambient Visualization. In *Issues in the Design & Evaluation of Ambient Information Systems Workshop.*, page 48, 2007.

[19] J. Oh, J. Herrera, N. J. Bryan, L. Dahl, and G. Wang. Evolving The Mobile Phone Orchestra. *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 82–87, 2010.

[20] J. A. Paradiso, J. Lifton, and M. Broxton. Sensate media - Multimodal electronic skins as dense sensor networks. *BT Technology Journal*, 22(4):32–44, 2004.

[21] N. Rasamimanana, F. Bevilacqua, J. Bloit, N. Schnell, E. Fléty, A. Cera, U. Petrevski, and J.-L. Frechin. The Urban Musical Game: Using Sport Balls As Musical Interfaces. *Proc. CHI EA'12*, pages 1027–1030, 2012.

[22] C. Roads. *Microsound*. MIT Press, Cambridge, 2001.

[23] A. Rowe. Within an Ocean of Light: Creating Volumetric Lightscapes. *Leonardo*, 45(4):358–365, 2012.

[24] A. Rowe, G. Bushell, L. Birtles, C. Bennewith, and O. Bown. Submergence 2013. *Leonardo*, 49(4):356–357, 2016.

[25] M. Sato. 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*, pages 3771–3776, 2008.

[26] M. Sato, A. Hiyama, T. Tanikawa, and M. Hirose. Particle Display System - Virtually Perceivable Pixels with Randomly Distributed Physical Pixels. *Journal of Information Processing*, 17:280–291, 2009.

[27] G. Schiemer and M. Havryliv. Pocket Gamelan: Tuneable Trajectories for Flying Sources in Mandala 3 and Mandala 4. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression*, NIME '06, pages 37–42, Paris, France, 2006. IRCAM — Centre Pompidou.

[28] J. Schwarz, D. Klionsky, and C. Harrison. Phone as a pixel: enabling ad-hoc, large-scale displays using mobile devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3–6, 2012.

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

[30] M. Tomitsch, K. Kappel, A. Lehner, and T. Grechenig. Towards a Taxonomy for Ambient Information Systems. In *Pervasive '07 Workshop: W9 - Ambient Information Systems.*, 2007.

[31] V. Verfaille, M. M. Wanderley, and P. Depalle. Mapping Strategies for Gestural and Adaptive Control of Digital Audio Effects. *Journal Of New Music Research*, 35(1):71–93, 2006.

[32] M. M. Wanderley and N. Orio. Evaluation of Input Devices for Musical Expression: Borrowing Tools from HCI. *Computer Music Journal*, 26(3):62–76, 2002.

[33] G. Wang, G. Essl, and H. Penttinen. Do mobile phones dream of electric orchestras. *Proceedings of the International Computer Music Conference*, 16(10):1252–61, 2008.