The Dissertation Committee for Yago Parra Moutinho Stucky de Quay Certifies that this is the approved version of the following dissertation:

The Development and Experience of Gesture and Brainwave Interaction in Audiovisual Performances

Committee:

Sharon Strover, Supervisor

Bruce Pennycook, Co-Supervisor

Shanti Kumar

Russell Pinkston

Yacov Sharir

Rui Penha

António Coelho

The Development and Experience of Gesture and Brainwave Interaction in Audiovisual Performances

by

Yago Parra Moutinho Stucky de Quay, B.Music.; M.Multimedia

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 2016

Dedication

To my parents, Isabel and José

Acknowledgements

I want to thank Bruce Pennycook, Yacov Sharir, and Sharon Strover for their significant intellectual and artistic contributions to this work. I also want to thank A. R. Rahman for inviting me to join his fantastic project with Intel, and Lakshman Krishnamurthy for his inspiring leadership in that project. Finally, I am grateful to have had Reema Bounajem supporting me emotionally throughout the writing process.

The Development and Experience of Gesture and Brainwave Interaction in Audiovisual Performances

Yago Parra Moutinho Stucky de Quay, PhD

The University of Texas at Austin, 2016

Supervisor: Sharon Strover

Co-Supervisor: Bruce Pennycook

Software and hardware developments in the 21st century have greatly expanded the realm of possibilities for interactive media. These developments have fueled an increased interest in using unconventional methods to translate a performer's intentions to music and visuals. And at the heart of every computer mediated performance is the question of what human action causes what audiovisual effect. This research explores this question and others by examining two performances that I created; performances featuring audiovisual media that respond to a performer's gestures and brainwaves. The aim is to improve our understanding of what factors influence interaction design within and across the performing arts. I argue that technology, collaborators, and transparency have a decisive impact on a performer's actions and interactive media. Furthermore, my findings suggest that synchronicity, elegance and visual feedback are important characteristics that can enhance the bond between the performer and audiovisual content. Lastly, I explain how universities and companies can improve research and development by together creating a live show.

Table of Contents

Chapter 1: Introduction

Background

Digital art is reinventing itself at an accelerating rate due to the exponential growth of technology. Furthermore, artists are routinely using unconventional methods to connect live performers with interactive media. These include open-source software, wearable computers, motion tracking, brain-computer interfaces, signal processing algorithms, and advances in processing power. At the heart of this multidisciplinary landscape is a question asked in the course of developing every new interactive work: what human action causes what audiovisual effect?

There are various theories exploring how to answer this question. Many come from human-computer interaction and cognitive sciences because of their successful usability models. However, many multimedia artists resist limiting their vision by basing their works solely on intuitive modes of interaction. Moreover, funding agents, collaborators, available technology, time, and aesthetic trends have a big impact on the trajectory of an artwork. We need a better understanding of the wider context in which audiovisual performances are created so that we can devise more flexible and effective guidelines for interactive media.

It is from this perspective that I explore how technology, collaborators, and transparency affect the relationship between a performer's actions and interactive media. This research revolves around two multimedia concerts that I participated in: *Curie* and *Ad Mortuos*. The common thread for these two shows is the development of specific technologies that enable real-time gesture or mental control of music and visuals. Using a variety of sensors and multimedia computers, the performers could control visual elements on a display and manipulate sound effects in front of a live audience. By creating and analyzing these two concerts, my aim was to understand what factors influence the design of action-media relationships, and under what circumstances these factors arise.

Main Question

- What factors influence the relationship between the performer and interactive media?

Sub-Questions

- How does access to resources affect the collaboration?
- How does the show impact the organization sponsoring it?
- How does the creative team cope with technical limitations?
- What are the technological challenges to real-time mappings?
- How and to whom is it important to make the mapping transparent?
- What technological considerations are important for interactive works?

METHODS

This dissertation analyzes two interactive audiovisual performances that I created during my doctoral studies—*Curie* and *Ad Mortuos*. Both performances involve a live show, a collaborative team, and mapping user data to musical and visual elements.

Inspired by ethnography, this investigation used a standard set of questions that examined the relationship between the performative structure, process, experience, and artists. Data was collected from self-reflection and interviews with collaborators. This information was used to create a detailed report of each case study and identify themes within each case. Following that is a comparative thematic analysis across the cases and an interpretation of the results.

Although this dissertation draws from a rich variety of literature and concludes with new theory, its foundations are practice-based and exploratory. My research started out as a creative endeavor, free of any scientific hypothesis that needed to be tested. Similarly, those participating in either of the two shows reviewed here were mainly motivated by aesthetical, practical, and intellectual considerations. At the end of the day, we measured success by our ability to make something novel, move audiences, overcome challenges, and adopt new technologies. This practice-based strategy to scientific inquiry is common within the arts (Ansari, Ansari, & Jaffri, 2014; Dallow, 2003; Edmonds, 2010). The reason for this has less to do with the attitudes of practitioners and more to do with the nature of artistic practice. Artworks are situated in a complex interplay of culture, collaborators, resources, and audiences. Adding theoretical constraints can undermine an artist's ability to capitalize on these factors to provoke significant change in tastes, meaning, identity, and the economy using the primary resource of creativity.

Therefore, this study integrates a mixed approach of inductive and deductive reasoning, where practice shapes theory and vice versa. The process of applying and building theory is illustrated in *Figure 1.1.* Looking at the bottom of this figure, we find

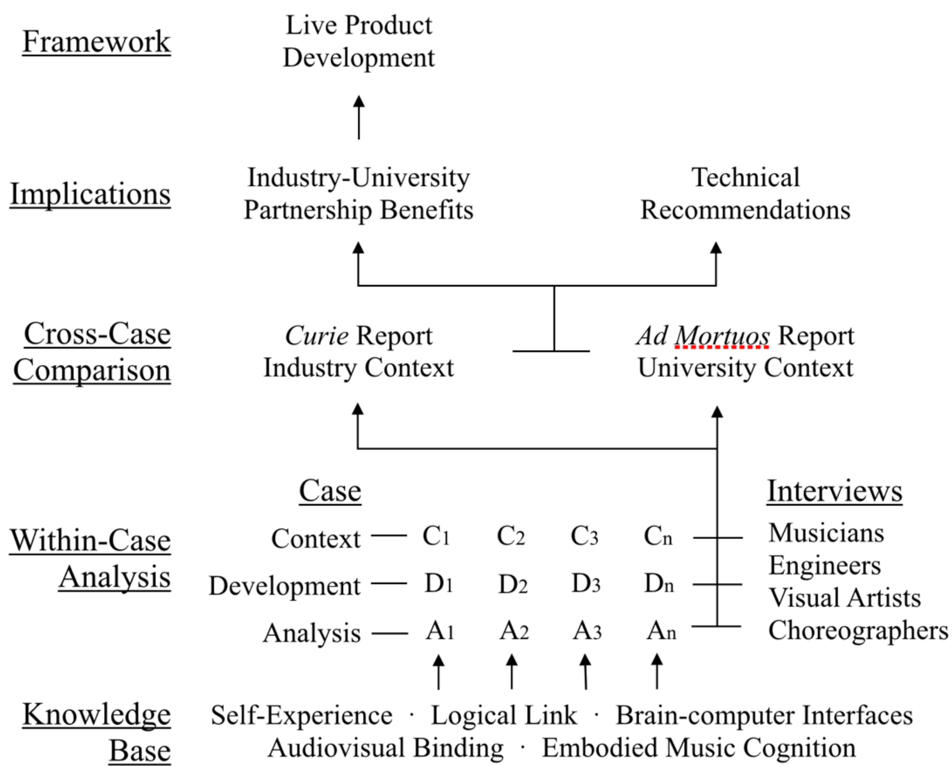

Figure 1.1: Conceptual framework illustrating the relationships between practice and theory.

the theoretical lenses through which the case studies are examined: audiovisual binding, embodied music cognition, and brain-computer interfaces, as well as first-hand experience. Next, each case study is divided into three stages that loosely resemble the chronological sequences of its design: context, development, and analysis. The intersection of theory and practice breaks up these three stages into research variables, which in turn are presented as questions in interviews with the artists and engineers involved in the projects. The ensuing case study report provides a comprehensive description of each project in an industrial or academic context. From this report, themes are identified and used for a cross-case comparative analysis (Baxter & Jack, 2008;

Maxwell, 2005; Miles & Huberman, 1994). The next level aggregates the technical and organizational findings and lays the foundations for a new industry-university partnership framework for the performing arts, which I have called Live Product Development.

Because of the case studies' artistic nature—practice-based research that springs from intuition and collaboration—there was not a set of *a priori* propositions that were tested systematically (Sharir, 2013). Instead, this dissertation explored the two performances in depth by looking at the common variables, which are illustrated in *Figure 1.2.* Both works were selected because they involved considerable artistic, technical, and collaborative development. Together, they showed different perspectives on the same issue. The analysis of these variables informed the development of *a posteriori* propositions regarding the relationships between the settings, performers, engineers, events, and processes within and across the cases (Miles & Huberman, 1994). By developing a broad theoretical understanding of the interactive shows, grounded in self-experience, this study remains framed by empirical data and aims to generalize its insights to a wider field of interactive art and business practice (Maréchal, 2010).

The analytical strategy of this research borrowed from the pattern-based method for case studies advocated by Creswell (2007) and Miles and Huberman (1994). Designed for qualitative studies, this kind of analysis is intended to condense information, elicit meanings, and compare themes. Collected data compiled during the study were assigned descriptive or interpretive labels. Labels were then consolidated and grouped to produce themes. Themes are summarized at the end of every case study and possible relationships made explicit (Creswell, 2016).

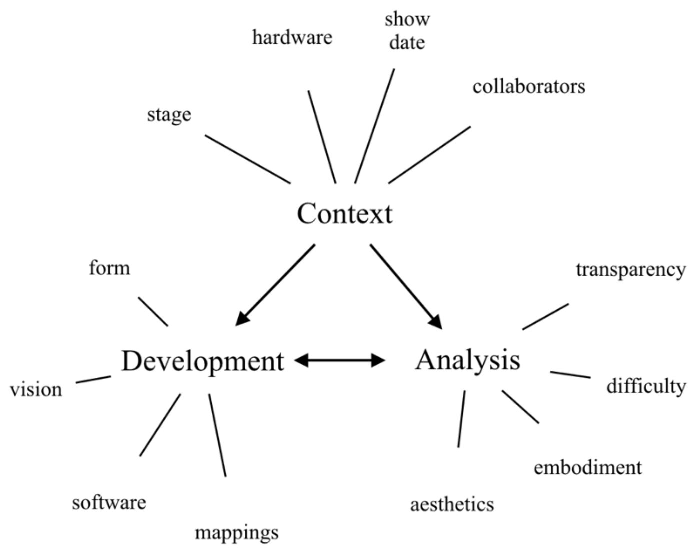

Figure 1.2: Examples of case study variables.

DATA COLLECTION

Data collection draws from self-reflection and interviews with collaborators. My own experiences and those of the interviewees are presented in the case study as a key source for understanding the performative world observed, yielding both self- and social knowledge. The interview procedure can be viewed in *Figure 1.3*. Using video footage of the performances as an aid, self-reflection and interviews answer various questions related to the variables within the context, development, and analysis of each case study. All collaborators involved in the decision-making process were interviewed; their names are listed in *Table 1.1*.

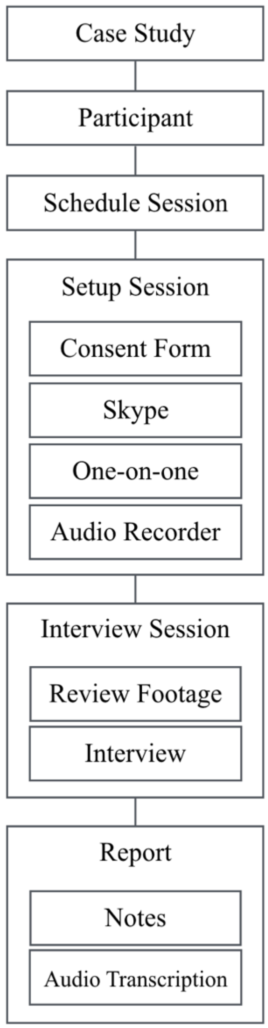

Figure 1.3: Interview procedure.

Case Study	Interviewees
Ad Mortuos	Rodrigo Carvalho (visual artist)
	Bruce Pennycook (composer)
	Yacov Sharir (choreographer)
Curie	Peter Sistrom (visual artist)
	Adam Nurre (musician)
	Mitchell Vandenburg (musician)
	Reema Bounajem (choreographer)
	Jean-Sebastien Rousseau (programmer)
	Chris Legaspi (programmer)
	Jason Blanchard (programmer)
	Tim DeBaillie (engineer)
	Swarnendu Kar (engineer)
	Lakshman Krishnamurthy (project manager)

Table 1.1: List of interviewees.

Dissertation Outline

Chapter 2 reviews theoretical and practical models for brain-driven and body-driven multimedia performances. It starts by talking about embodied music cognition models that describe the relationship between a human subject and a musical environment. The next section goes beyond movement-sound pairings into movement-sound-image interaction and presents cognitive mechanisms that help us bind what we see and what we hear into a single multisensory object. In Brain-Computer Interfaces, I investigate the ways in which brain signals can be decoded and transformed into computer commands. The last two sections offer examples of interactive performances and enabling software toolboxes. Altogether, this chapter is of use to researchers intending to create live shows that incorporate body movement, brainwaves, and digital media.

Chapter 3 describes at length the development of two live multimedia shows, *Curie* and *Ad Mortuos*, respectively created in commercial and academic environments.[1] These two reports aim to thoroughly examine the dynamic relationship between machines, aesthetics, and people. In the context section, I describe the production equipment, infrastructure, and origins of each show. Through the voice of participants and my first-hand experience, the development section charts the evolution of the technology, mappings, and artwork as a whole. The technology portion describes the hardware and software, while the mapping section reveals how the performers were connected to interactive digital media. Visuals, choreography, and music are explained by their respective artists. The information given by the interviewees and my own

[1] The videos of the shows for each case studies can be found at the following links: *Curie*: https://vimeo.com/170378294; *Ad Mortuos*: https://vimeo.com/123801965.

experience are analyzed in the end—the large themes identified here revolve around technological, performative, and collaborative concerns.

Chapter 4 puts forth a set of technical recommendations for multimedia shows and discusses their implications for businesses and academia. In Technical Recommendations, I cover latency, transparency, and feedback. After that, I outline the implications for companies and universities, which include accelerated development; improved product quality, usability, and exposure; expanded team competency and skills; and an increase in project resources and funding. I wrap up this section by introducing the Live Product Development framework, a type of industry-university partnership that uses technology-driven shows to accelerate a company's new product development process and to augment a university's research resources.

Chapter 2: Literature Review

The literature review looks at how humans cognitively and perceptually bind brainwaves, movement, sound, and images. I start by describing embodied music cognition theory and how its intuitive metaphors can strengthen movement-sound relationships. The following section adds the image component to the previous interactive pair. Here, audiovisual binding looks at an alternative way to connect visual and sonic elements by exploiting our innate tendency to identify a common cause to what we see and hear. In doing so, audiovisual binding frees artists from relying on abstract metaphors for cohesion, opting instead for semantic, temporal, or spatial characteristics as unifiers for all interactive elements. Ultimately, these techniques help audiences and performers fuse the various elements they see and hear in a show. Expanding interactivity to incorporate visuals can undermine a performer presence on stage—this tension is the topic of the third section. Following that, I reassert the role of the live performer through interactive transparency and technology concealment.

The second part of this literature review sheds some light on what is still a very experimental field—brain-computer interfaces for musical expression. It describes relevant neural signals, what features to look for, how to extract them with electroencephalography, and typical interaction models. Their application centers around controlling external software with deliberate mental actions, mental states and facial expressions. The last two sections list examples of performances using body motion and brainwaves.

Embodied Music Cognition

Virtual instruments enable a computer-mediated conversation between a performer's movements and digital sounds. At the heart of this relationship rests a question asked in the course of developing every new interactive work: what action causes what audiovisual effect? In answering this question, embodied music cognition has contributed much to our understanding of how humans and instruments interact. This theory provides us with intuitive models that associate movement with artificial sounds that are not a result of a body's interaction with acoustic objects.

Jensenius (2007) argues that our understanding of music is informed by cognitive capacities and limitations of the body in relation to the environment. Motor areas of the brain become activated by thinking of an action, watching someone perform an action, and hearing the sound of an action. Ecological knowledge of how actions and objects interact dictates the perceived relationship between actions and sounds and their mechanical and acoustical coupling.

The author states that while action-sound couplings seem natural because they emerge from interaction with the physical world, artificial action-sound relationships have severed their mechanical and acoustic coupling by having the sound generated electronically through speakers. This relationship is perceived as unstable because the same action on an interface might produce a different, unexpected sound. Artificial action-sound relationships range from weak to strong depending on how natural they feel.

Visi, Schramm, and Miranda (2014) describe what is a natural gesture-to-sound relationship by arguing that the body can be understood as a mediator between physical sound and emotional musical experience. As a mediator, the "body will build up a repertoire of gestures and gesture/action consequences" (Visi et al., 2014, p. 2) that allows the musician and the audience to attribute feeling and intention to a gesture and predict its musical outcome.

Paine (2009) interviewed skilled acoustic musicians to examine fundamental control parameters related to their instruments, and applies those parameters to computer music. Adding relevance to the importance of explicit mappings, the author states that gestures embody the quality of a sound and engage an audience. The instrument should support these functions by encouraging performability (i.e. complexity and uniqueness) and communication of energy and intent.

The language used to describe relationships between body movement and sound plays a major role in interactive systems. Antle, Corness, and Droumeva (2009) carried out an experiment to evaluate children's learning with a musical system based on body movements. Results showed that a particular set of movement metaphors increased the learning rate. The authors state that "embodied metaphors conceptually extend embodied schemata through the linking of a source domain that is an embodied schema and a target domain that is an abstract concept" (Antle et al., 2009, p. 68).

Embodied schemata were described by Johnson (1987) as metaphors that unconsciously arise from the body's interaction with objects, orientation in space, and experiential gestalts. For example, he stated that a metaphor that connects proximity to pitch is "near is high; far is low"; a metaphor that connects activity to volume is "more is loud; less is quiet." Lakoff and Johnson (1980) describe these abstract concepts as objects with spatial and physical attributes. These metaphors, which link abstract concepts to actual body movements, are the basis for embodied cognition.

Metaphors ground abstract concepts to universal sensory-motor abilities and our associated low-level inference capabilities. Based on an experiment using whole-body interaction with the tonal harmony visual interface *Song Walker*, Holland, Wilkie, Bouwer, Dalgleish, and Mulholland (2011) claim that users can develop various musical skills such as rhythm, composition, and memorization if motor actions are implicated in the manipulation of a musical task. Their results provide evidence of kinesthetic learning (vs. cognitive learning), successful collaboration between two users, and a quicker learning curve compared to desktop interaction. When designing for whole-body interaction, the authors recommend crafting metaphors that physically ground abstract concepts and explore cues that stimulate physical manipulation.

Eckel (2012) developed an interactive genetic algorithm that aims to provide insight into the relationship between bodily and musical expression. In this system, the faster the dancer's movement, the lower the delay of the sonic consequence, and the higher its pitch modulation. The delay time (minimum 0 seconds and maximum 4 seconds) hints at an imagined agent mediating the interaction, and enables the performer to transition between an instrumental and structural role. The author concludes that rationally, it is hard to manipulate the system, while kinesthetic knowledge becomes an intuitive way to play with it. In this case, interactive consistency and coherence is apparent to the body, not the mind.

Yet movement-sound metaphors are only concerned with musical interaction. The reason why embodied cognition is at the center of many discussions regarding interactive instruments is because these instruments lack visual feedback. This act of compensation is apparent from the visual language employed in these metaphors, such as “fast,” “near,” “smooth,” “slow,” and “high” (Antle et al., 2009, p. 60). If we move further from movement-sound relationships and look more generally at how humans combine images and sound, we start to see how our perception looks for cues that help organize these two senses into coherent chunks. Audiovisual cues can overcome embodied music cognition’s restrictions on the interactive possibilities, as well as limitations on sound-only interaction. The following literature explains how causality can be exploited in multimedia artworks to bind visual and auditory experience.

AUDIOVISUAL BINDING

Metaphors can be thought more broadly to incorporate scenarios where the performer’s intentions are expressed visually as well as sonically. Interactive images provide powerful cues that help combine movement, sound, and image. Callear (2012) and Coulter (2010) agree that culturally-learned associations are partially responsible for the perceived connection between audio and visual materials. Media pairs that clearly reference the same object via metaphoric or linguistic conventions are more easily integrated than abstract materials that require further explanation. However, consistent and synchronous repetition of abstract media pairs may in time bunch together into a coherent whole. Coulter adds that interactivity, through feedback loops, helps maintain morphological correspondences.

Cooke (2010) describes how different audiovisual environments affect how we make sense of what we see and hear. An individual in an immersive environment becomes a member of a collective experience and a participant in the image-sound nexus. Visual representations of musical ideas tend to explore the relation of color to sound while stimulating audio and visual senses to create a kind of synesthetic experience. Chance or algorithmic art takes the process of creating meaning away from the artists, allowing order and coherency to arise spontaneously from the system.

Whitelaw (2008) claims that through analogy, transcoded audiovisuals attempt to create a synesthetic experience. Biologically, synthetic perceptions are automatic, involuntary, and caused by cross-connections in the neural channels of the perceptual system. This network of sensory channels is not fixed, but reconfigurable over time, and perhaps barely even distinct from each other. Therefore, sensory channels can learn to substitute and take over nonexistent perceptual channels, such as sound over vision in the blind. When discussing the analogy between audiovisual works and synesthesia, Whitelaw states:

> Audiovisual works are artifacts; objects of perception, not perceptions. To put it bluntly, synesthesia, by definition, occurs in the perceptual system of a synesthete, not in the crossed connections of a video synth. […] we can use the gap as a provocation, rather than an obstacle. (p. 267)

Audiovisual works do not provoke sensory substitution because this process requires long-term learning and both sensory channels are stimulated. Instead, Whitelaw offers an alternative approach to the aesthetics of fused audiovisuals: binding. When our mind perceives correlations between objects, it binds them into groups that often correspond to objects in our physical reality. Binding is persistent, sticky, and ultimately pleasurable. Our limbic system seems to reward us when we discover correlations in our environment, suggesting that it is part of our survival mechanism. This also suggests that artists, maybe unknowingly, seek to produce the pleasure of sensory binding by offering clues to potential objects in our environment. Furthermore, cross-sensory binding can alter our interpretation of an event, since the correlated stimuli can point to different causes to explain their coherence.

Coulter (2010) agrees that integrating visual and auditory elements serves the purpose of increasing perceptual clarity. We have a natural tendency to integrate audiovisual objects, but the ability to separate them allows a switch between homogenous and heterogeneous experiences. Our innate binding mechanism highlights homogenous objects and masks those that are heterogeneous. Artists can manipulate shared characteristics such as duration, intensity, and proximity to catalyze media integration. However, perceptual bandwidth limits the ability to integrate a large number of homogenous objects into a coherent whole.

Yet image and sound do not need to be semantically or metaphorically similar to be bound together. Moody, Fells, and Bailey (2006) and Callear (2012) identify temporal synchronicity in audio and visual objects as the most effective form of creating audiovisual congruency. Even unrelated objects are perceptually bound together when they happen at the same time. This binding is subconscious, attracts attention, and makes an image more real or hyper-real. Synchronous materials help clarify correspondences and facilitate audience attention during a complex audiovisual scene. Callear concludes by stating that deliberate manipulation of the temporal congruency is a mechanism for creating suspense and expectation.

Most audiovisual works, in some way or another, try to make perceptual or conceptual sense. Cooke (2010) argues that audiences often expect there to be a narrative somewhere, and look to bind the objects they experience into narrative fragments. Modern audiences are accustomed to witnessing, and are ready to apply, narrative strategies in entertainment, economy, and culture. The problem with narratives is that they orient us inevitably towards an end, which when reached transforms the original experiences leading to it. He suggests that instead of creating meaning, the real potential of live audiovisual performance rests in its ability to explore new relationships between media, data, and the human body:

> Rather than asking what it is, we might take an approach from Spinoza, who in asking what a body is capable of—what can a body do? —shifts the definitional emphasis and burden of proof, orienting the mode of questioning towards context, capacity, and possibility rather than meaning and ontological or phenomenological certainty. (p. 205)

The ability to capture any source of input, modulate it, and output it through various media can be understood as a broader movement related to data capture and transformation. In some cases, it is more useful to ask "What can the body do?" rather than "What is it about?" New media performances are not concerned with what is, but what can be done and experienced.

MEDIATIZATION ON STAGE

Despite the interactive possibilities afforded by audiovisual binding, some authors argue that digital media diminishes the role of the live performer. Digital media and electric amplification can mediatize live performances to such an extent that the live event itself becomes a product of media technologies.

Kelly (2007) and Auslander (2008) claim that massive video screens produce an effect of intimacy and immediacy at the cost of liveness. In fact, intimacy and immediacy are precisely the factors originally attributed to television that enabled live performance to be displaced. Both authors use Madonna's performances as an example. Some of Madonna's tour staff have freely admitted that the performances emulate her music videos as closely as possible, as they are the real standard the audience expects. Assuming the audience is familiar with the video clips, regardless of live mediatization, the feeling of proximity and intimacy associated with video can be reproduced live by the faithful replication of the video clips. At the same time, the music videos gain a certain realness from the live performance:

> The kind of proximity and intimacy we can experience with television, which has become our model for close-up perception, but that is traditionally absent from these performances, can be reintroduced only by means of their "videation." […] Because we

> are already intimately familiar with the images from our televisual and filmic experience of them, we see them as proximate, irrespective of how far away they may be in physical distance. (Auslander, 2008, p. 39)

According to Auslander, the pervasive, reproducible, and mediatized character of performance has led to its own depreciation, and can only be compensated for by further mediatization. The conventions of mediatized experiences shape the way we think of live events and possibly turn them into a degraded version of the former. This leads to his famous equation, Dance + Virtual = Virtual. The presence of a body and digital media does not result in fusion of realms; instead, the dominant media incorporates the live element as its raw material.

Klich and Scheer (2012) on the other hand, claim that the above equation is only valid when the virtual (that is, the media) is self-sufficient and exists independently of the performer. They propose instead the notion of intermedial theater that unites both the live and media elements in a symbiotic, non-hierarchical communication system within the frame of the performance. These elements derive their meaning from their relationships, and as such cannot be understood as discrete. According to Klich and Scheer, through the process of remediation, new media can refashion old media, a concept characterized by two logics. On one hand, transparent mediatization presents only the object of mediation and is positively correlated with immersion. On the other hand, conspicuous technologies attract participants' attention to the medium as a discrete element with intrinsic mediation qualities.

Interaction Transparency

Concealment of media technologies can redirect attention back to the performer. Mullis (2013) states that the technological device that enables audiovisual communication between the performer and media can be characterized by the visibility of the physical machinery. For example, a projector is normally hidden from view, exposing only its result, the projection. Since these technologies are necessary for performance but purposely concealed, they exist as entities without a context. Their functional, interactive, and concealed nature makes them interchangeable with any other technology so long as the desired end result is achieved.

This concept is similarly described by Klich and Scheer (2012) when they talk about medium transparency. Kelly (2007) also observes that obscure technologies promote an illusory, immediate, and immersive experience, while visible technologies acknowledge a constructive performative system.

Furthermore, Mullis (2013) claims that technology concealment and interactivity guide attention towards the intersection of the performer and media and disconnect the artist behind the media from the experience. This develops a tension between two creative forces, the non-corporeal digital artist and the unique embodied performer. Interestingly, the interactive technology has no affinity with the performer, since it will capture anyone who interacts with it. The performer's body then is simply a "focal thing" (p. 112) that presents the fusion of means and ends and can be interchanged with someone else with the same skills. As this tense relationship unfolds artistically, the audience gains understanding of how much training is needed to develop an expressive interaction. However, the dialog between the technology and the performer is always strained, since the former can only respond to the latter through a few interactive parameters.

The interaction can also be made explicit or obscure. When discussing intermedia mapping strategies, Callear (2012) describes transparency as the perceptual binding clarity of input and output. Parameter mapping in media-driven audiovisual compositions translates data from one medium to another by exchanging morphological structures. On the other hand, data-driven compositional morphological structures are derived from an external data source, or its salient properties. The perceived transparency between these mappings is very personal and is affected by preexisting knowledge and cultural background. As described earlier, scene complexity and audiovisual synchrony also affect transparency. Manipulating transparency can negate the fundamental predictability of strongly mapped audiovisual material.

However, obscuring the interactive rules may reduce binding. Whitelaw (2008) recalls that the ecological function of binding is to identify a common cause. Alternative causes can be inferred from an interactive system's input signal. This signal forms an underlying structure of sequential similarities that are played out in the auditory and visual field. Better understood as a procedure than an object, this specific but abstract underlying structure is where the mapping unfolds. Ultimately, mapping can produce noesis or revelation by making apparent the common cause and thus allowing audiovisual binding.

INTERIM SUMMARY

Combining sounds with a performer's movement is not a straightforward task—audiovisual works need to pay attention to the cognitive processes governing how humans combine what they see with what they hear. Embodied music cognition explains that music is at least partially experienced and understood through low-level sensory motor areas. Metaphors based on how the body experiences the world inform our perception of musical elements. Artificial mappings that disobey our intuitive mechanical and acoustical heuristics are perceived as weak action-sound relationships.

Nevertheless, embodied music cognition does not take into account reactive visuals in its interactive framework. Responsive images add another layer of information about the relationship between the performer and digital media. Furthermore, all dynamic audio and visual elements can incorporate semantic, temporal, and spatial characteristics, which cue our perceptual system to bind them together into one coherent phenomenon. The advantage of visual and auditory cues is that they free artists to pursue more interactive choices than those afforded by embodied music cognition metaphors alone.

In these interactive audiovisual environments, concealed technologies promote an illusory, immediate, and immersive experience, while visible technologies become part of a constructive performative system. The same applies to media and mediums: those that are embedded transparently are more immersive, while those that are obvious attract attention. Transparent technology directs the attention away from the machinery and towards the intersection of the performer and media, while disconnecting the artist behind the media from the experience. In multimedia works, the performer can be absorbed by the dominant media or bound into a symbiotic relationship.

Brain-Computer Interfaces

Making music with brainwaves is still very much an experimental endeavor. The evolution of our knowledge, technology, and learning algorithms in the field of brain-computer interfaces continues to provide new and more effective strategies for translating mental intentions into computer commands. This section explores ways in which computers can be controlled by physiological bio-signals originating from the brain. Techniques for brain-computer interaction are described in the first part, followed by methods for extracting reliable signals and relevant features. The last part presents four examples of brain-driven musical systems.

Wolpaw and Wolpaw (2011) state that the role of the central nervous system (CNS), which includes the brain, is to respond to events in the environment by producing outputs that serve a person's needs. All natural outputs are either neuromuscular or hormonal, but brain-computer interfaces (BCI) add a new output that is neither. Instead, "a BCI is a system that measures CNS activity and converts it into artificial output that replaces, restores, enhances, supplements, or improves natural CNS output and thereby changes the ongoing interactions between the CNS and its external or internal environment" (p. 3).

Drawing from the field of neuroprosthetics, He, Gao, Yuan, and Wolpaw (2013) describe five kinds of applications for BCI: 1) to replace a natural output that has been lost to injury or disease; 2) to restore the functionality of a lost organ or limb; 3) to enhance a natural ability to perform a particular task; 4) to supplement a natural output with an artificial one; or 5) to improve and develop a natural output.

Royer, Rose, and He (2011) identify two types of BCI commands: process control and goal selection. Process control enables the user to control the details of the system that accomplishes the intended task with little or no software automation. Goal selection allows the user only to communicate a goal to the software, with the intention that the software accomplishes that goal. In their study, the process control scenario had users mentally drag a mouse cursor left or right by actively producing the necessary brain signals; the goal selection scenario had users mentally ordering the mouse to move fully right or left.

Jackson and Mappus (2010) split goal selection into either binary or n-ary. Binary selection tasks evoke brain activity in two separate brain regions, while n-ary selection tasks discretize a brain signal with a number of progressive thresholds. In goal selection, the software adjusts its actions according to the feedback of the current state in relation to the declared goal. In process control, the user interprets the feedback and chooses the subsequent actions needed to fulfill an intent (Wolpaw & Wolpaw, 2011). BCI can combine both interaction models, so that like ordinary motor activity, control over an output would include automatic reactions as well as voluntary actions (Curran & Stokes, 2003).

Effective use of a BCI requires training the user to produce reliable, voluntary brain signals that the BCI can decode to fulfill an intent (Tan & Nijholt, 2010). Users are trained for two reasons: first to introduce them to the control paradigm, and second to produce training data for the software's categorization algorithm. It is worth noting that brain activity is influenced by concentration, distraction, frustration, fatigue, motivation, emotions, and intentions. And, unlike motor control, it is hard to gauge, since the user can neither identify nor discern internal activity aside from the software’s feedback (Jackson & Mappus, 2010).

Vallabhaneni, Wang, and He (2005) describe how two methods—cognitive tasks and operant conditioning—enable the user to learn how to control a BCI’s output. The cognitive task method requires the user to perform mental actions such as arithmetic, visual counting, geometric figure rotation, composing letters, and imagining moving an arm. This system also requires a contrasting relaxed baseline reading. Cognitive tasks can be further divided into either endogenous (not requiring any external stimulus) or exogenous (where the user responds positively to external stimuli).

Operant conditioning, on the other hand, expects the user to modify (condition) his/her (operant) behavior as a result of the consequence. While both methods are voluntary, cognitive tasks require mental engagement, while operant conditioning becomes automatic. One of the biggest challenges in BCI research is how to configure the two-way learning process—the software adapts to the brain signals, while the user modifies his/her behavior in response to the consequences (Wolpaw & Wolpaw, 2011). *Figure 2.1* Illustrates the three methods for operating a BCI described here.

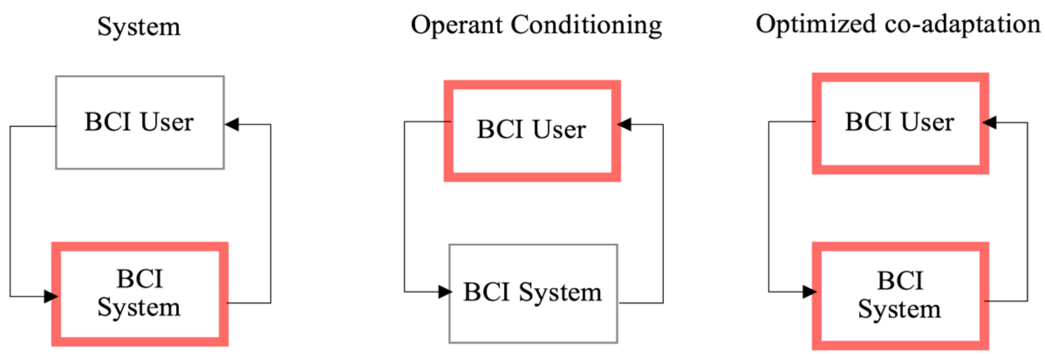

Figure 2.1: Three methods for BCI operation. The red boxes around the system or user identifies the task of learning how to improve and maintain the correlation between intent and output.

Brain activity consists of electric, chemical, and metabolic phenomena such as neuronal action potentials, synaptic potentials, neurotransmitter releases, and oxygen consumption. This activity can be measured by monitoring "electric or magnetic fields, hemoglobin oxygenation, or other parameters using sensors on the scalp, on the surface of the brain, or within the brain" (He et al., 2013).

He et al. (2013); Otto, Ludwig, and Kipke (2011); and Srinivasan (2011) investigate the main technologies used for brain signal acquisition. The suitability of the technology depends on the kind of operating environment and resolution required for proper translation. These technologies can be categorized into either invasive or non-invasive. Invasive acquisition of brain signals is accomplished by electrophysiological readings from electrodes that are surgically implanted inside or on the surface of the brain. For example, miniaturized microelectrode arrays can be implanted into the cerebral cortex to record neuronal action potentials or local field potentials from highly localized sets of neurons and synapses. Surface activity is acquired with electrocorticography (ECoG), where an electrode array is implanted subdurally over the cortex (on the outermost layer of the meninges that surround the brain). ECoG takes advantages of the fact that most large cortical neurons are oriented perpendicular to the cortical surface and their synchronous activity can lead to a detectable signal.

The above authors state that the two principal non-invasive signal acquisition techniques are magnetoencephalography (MEG) and electroencephalography (EEG). Brain-current sources generate a small magnetic field that can be measured through MEG—more precisely, a superconducting quantum interference device (SQUID) magnetometer placed 1-2 cm above the scalp surface. To maintain superconductivity, SQUID coils must operate within a few degrees of absolute zero, cooled with a liquid helium-containing Dewar insulator. Being extremely sensitive, MEG has to be shielded from ambient magnetic field variations by placing the user inside a high-permeability mu-metal chamber. EEG recording involves placing electrodes on the scalp: at least, a ground electrode and two recording electrodes. EEG is by far the most commonly used technology for BCI.

> EEG is the most prevalent method of signal acquisition for BCIs. EEG recording has high temporal resolution: it is capable of measuring changes in brain activity that occur within a few milliseconds. The spatial resolution of EEG is not as good as that of implanted methods, but signals from up to 256 electrode sites can be measured at the same time. EEG is easy to set up, portable, inexpensive, and has a rich literature of past performance. The practicality of EEG in the laboratory and the real-world setting is unsurpassed. EEG recording equipment are portable and the electrodes can be easily placed on the subject's

> scalp by simply donning a cap. In addition, since EEG systems have been widely used in numerous fields since their inception more than 80 years ago, the methods and technology of signal acquisition with this modality have been standardized. Finally, and most important, the method is noninvasive. (He et al., 2013, p. 15)

EEG recordings are bipolar and require electrode pairs to measure scalp potentials by having the current pass through a measuring circuit (Srinivasan, 2011). The placement of the electrodes follows the International 10/20 system, so that the distance between adjacent electrodes is 10% or 20% of the total front-back or right-left distance of the skull (He et al., 2013).

EXTRACTING RELEVANT NEURAL SIGNALS

> *Oscillatory EEG activity* is caused by a complex network of neurons that create feedback loops. The synchronized firing of the neurons in these feedback loops generates observable oscillations. The frequency of oscillations decreases as the number of synchronized neuronal bodies increases (Vallabhaneni et al., 2005, p. 96).

Schomer (2007) describes the commonly used frequency bands extracted from a power spectrum analysis of the oscillations. Delta band (1-4 Hz) is associated with deep sleep. Mild during wake states, it tends to exhibit a non-oscillatory nature. The same can be said about theta bands (4-8 Hz), which are found during states of drowsiness, meditation, and deep dreaming. Alpha (8-12 Hz) emerges when the eyes are closed and during relaxation. It resembles a sinusoidal wave within a narrow frequency range at 10

Hz. Beta (12-30 Hz) consists mostly of broadband noise, does not have single frequency peaks, shows bilateral symmetrical distributions with more frontal presence, and is connected to motor behavior. The final band mentioned here, gamma (>30 Hz), is noticeable during complex cognitive and motor tasks, but shows very low power and reliability except at 40Hz (Schomer, 2007).

Event-related potentials (ERPs) are described by He et al. (2013); Sellers, Arbel, & Donchin (2011); and Vallabhaneni et al. (2005) as unique patterns of voltage deflections that are time-locked to a stimulus. The initial 150 ms of an ERP following a stimulus, referred as exogenous ERP, tends to reflect the unconscious activity of the primary sensory system, and is dependent on the stimulus modality. Later components, called endogenous ERP, normally represent internal information processing and cognition, and are more dependent on the significance of the stimulus. P300 is an endogenous ERP component that occurs when the subject is presented with a rare stimulus, making it useful for selection tasks. Lastly, the frequency of continuously flickering lights or pulsating tones can evoke a similar frequency (up to 40 Hz) in the brain. These signals are called Steady-state evoked potentials (SSEPs) (Birbaumer, 2011).

Allison, Faller, and Neuper (2011) reveal the relationship between actual movement or imagined movement and slow voltaic changes, called slow cortical potentials (SCP), over the sensorimotor cortex. SCPs are time-locked and phase-locked to sensorimotor events. The contingency negative variation (CNV) is a negative SCP that begins 200-500 ms after a stimulus, in anticipation to another imperative stimulus (i.e.

demanding a behavior response). Readiness potential is another negative SCP that begins 500-1000 ms after a self-initiated movement.

Vallabhaneni et al. (2005) provide a concise description of how to determine what brain activity features are useful as command components for BCI:

> The goal of all processing and extraction techniques is to characterize an item by discernible measures whose values are very similar for those in the same category but very different for items in another category. Such characterization is done by choosing relevant features from the numerous choices available. This selection process is necessary, because unrelated features can cause the translation algorithms to have poor generalization, increase the complexity of calculations, and require more training samples to attain a specific level of accuracy. (p. 16)

Krusienski, Mcfarland, and Principe (2011) argue that brain activity features should possess the following attributes: 1) spatial, temporal, spectral characteristics, and dynamics that can be precisely distinguished for a user or group of users; 2) the ability to be modulated and combined to convey intent; and 3) long-term feedback stability and correlation with a user's intent. The first step in feature extraction is signal conditioning and noise reduction through filtering, data reduction, normalization, and elimination of environmental noise and artifacts. Furthermore, Wolpaw et al. (2000) state that maximizing the signal-to-noise ratio is crucial for successful feature extraction. Problems identifying noise arise when features such as frequency, amplitude, and time are similar to those of the desired signal. Two of the main sources of noise are 50/60 Hz power-line interference and facial muscles, particularly eye movement.

Mcfarland and Krusienski (2011) explain how extracted features represent indirect measurements of a user's intent—they must be translated by algorithms into appropriate commands that convey that intent. Crucial to the translation algorithm is the mathematical model that accepts feature vectors as its input to output commands for a device. At its core, the model describes the relationship between feature vectors and the user's intent in a way that is simpler than the measured data. The model's parameters are often trained using data consisting of feature vectors and intended outputs. Using supervised learning, parameters are continuously adjusted until the model translates feature vectors into output commands as accurately as possible. Furthermore, the model must function well with new data (i.e. adapt) due to the real-time demands of BCI.

INTERIM SUMMARY

Brain signals are used in interfaces to produce continuous or discrete commands that are subject to conscious or unconscious control. Electroencephalography (EEG) is the most popular technology used for measuring brain activity because of its high temporal resolution and convenience. Power spectral analysis is the most used method for identifying particular mind states; it works by measuring brain frequency bands, categorized as delta, theta, alpha, beta, and gamma.

Examples of Interactive Performances

We can trace the roots of modern-day interactive performances to works created in the 1960s, 70s, and 80s by pioneering electronic music artists. Alvin Lucier is famously known for being one of the first composers to use brainwaves as a musical instrument (Straebel & Thoben, 2014). In his piece *Music for Solo Performer*, which debuted in 1965, Lucier amplified the alpha signals produced by his brain and routed them to loudspeakers coupled with percussive instruments. He controlled his alpha waves through relaxation. Between 1976-79, David Rosenboom built on this idea by incorporating more brain signals and developing a generative music system. His non-linear interactive system *On Being Invisible*, for example, used event-related potentials from the brain to generate musical structures in real time (Rosenboom, 1990).

Michael Waisvisz was a performer and inventor who created groundbreaking instruments that converted sensor data into musical information. Only one year after MIDI standards were introduced, in 1984 he built *The Hands*, a pair of hand controllers that featured switches, an ultrasonic transmitter, wheel potentiometers, toggle buttons, tilting senors, and finger motion detection (Waisvisz, 1985). (MIDI is a communication protocol for virtual instruments.)

In the following works, we will turn our attention toward more recent multimedia performances that use data captured directly from the performer's body or brain to manipulate audio or visual parameters in real time. Although the interaction paradigm has not changed drastically since the early days of interactive music, the evolution of sensors and computers has greatly expanded our visual and musical capabilities.

Brainwaves

Thies (2012) reviews three models that use emotions as controls for computer processes. Intuitive but challenging to apply, the "Basic emotions" model groups emotions into categories of happiness, fear, disgust, sadness, anger, and surprise. The "Dimensional" model measures emotions in a range of active to passive and positive to negative (activation and valence). This is described as being less intuitive, requiring training sessions. "Appraisal-approach," the most recent model, is the least intuitive and most challenging of all, but provides a helpful way to distinguish between primary and secondary emotions.

Using power spectrum analysis, Makeig et al. (2011) trained a program to recognize five feelings so that they can be applied for note recall in live performance. During training, subjects were played five notes while imagining each one to be a human emotion and paying attention to their own somatic sensations. *Figure 2.2* maps some of the components that made up the emotion classifier. During the performance, subjects successfully produced a five-note musical sequence (84% correct performance) by re-experiencing feelings they had associated with notes during the training sessions. While the subject produced these notes, a violinist, flautist, and cellist performed an original composition over them.

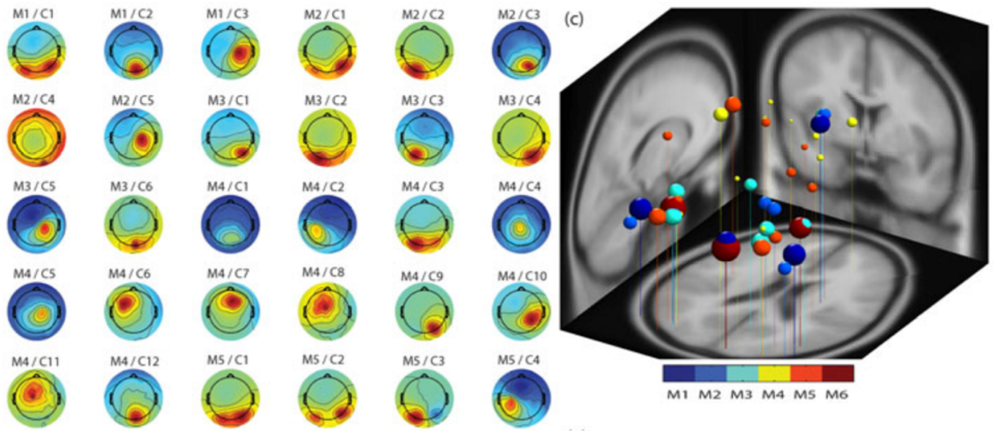

Figure 2.2: Activity map and location of some of the independent components in the brain used to identify feelings.[2]

Folgieri and Zichella (2012) performed a study where subjects' beta brainwave patterns became associated with problem-solving to evoke notes after a short training session. The authors use Neurosky Mindwave's EEG because of its comfort, wireless capabilities, and single dry sensor positioned on the area of interest, the premotor frontal cortex. During the training session, subjects were played a single note, shown an image, and asked to perform a gesture. The reinforcement stimuli, a color and gesture in this case, reduced training by about 45% by increasing concentration when compared to just being played the note.

---

[2] Adapted from "Demonstration of a Musical Emotion BCI," by S. Makeig, G. Leslie, T. Mullen, D. Sarma, N. Bigdely-Shamlo, & C. Kothe. (2011). *Affective Computing and Intelligent Interaction*, 6975, p. 494. Copyright 2011 by Springer-Verlag Berlin Heidelberg.

By evoking the associated sound, color, and gesture, subjects were able to reproduce three to seven notes after one training session. Results show that subjects have individual cerebral patterns that the BCI needs to learn during a training phase. The experiment demonstrates that there is a correlation between the execution of an action and the will to execute it.

Grierson (2008) experimented with matching musical notes to specific electrical signals produced when subjects looked at a visual cue. The author measured P300 event-related potentials produced by the brain in response to an external stimulus. In the first part of the experiment, subjects were displayed a sequence of flashing note-names while the P300 ERP for each note was recorded with an eight-channel EEG. Each P300 ERP was thus associated with a note. In the second part, subjects were asked to look at a sequence of notes on the display and the system attempted to reproduce the notes. Results show that the system correctly identified 75% of the notes after seven trials for most subjects. Two issues identified were the large latency of more than seven seconds between note selection and performance, and the body movement of participants that often contaminated the EEG data.

The above musical BCIs generate music indirectly from the brain. That is, the user is not actually thinking of a musical sound but attempting to re-experience a feeling that the software has learned to sonify. Klonowski, Duch, Perovic and Jovanovic (2009), however, discuss their results in extracting tones directly from the subject's brain while they are in the process of mentally imagining a tone. Their goal is to use these signals as a musical command language for BCI and to test its musical abilities.

> Just like an imagination of hand or finger movement is related to changes in activity of the brain somehow resembling those connected with the real movement, so the process of mental hearing and comprehending music is related to changes in brain activity somehow resembling those occurring in the brain when listening to real physical sounds of music. Such a cognitive process of auditory imagery, of singing in the mind, is also called audiation. (Klonowski et al., 2009, p. 1)

During the experiment, subjects were asked to listen to a calibration tone and then imagine the same tone while brain activity was recorded for 5 to 10 seconds. Brain data was captured using an eight-channel EEG and analyzed in real-time with the Fourier method. Detection was complicated by the high amount of noise present in the signal and the low number of harmonics in the imagined tone. Nevertheless, 123 out of 147 experiments correctly detected imagined tones through their spectral lines. Subjects lacking musical abilities produced negative results.

Canibal, Moura, and Pedro (2013) equipped a performer with an EEG, Electrocardiogram (ECG), and a depth camera. Their one-man, 45-minute-long piece, *Câmara Neuronal*, blends performance art with interactive sounds and visuals to create one coherent narrative. Captured signals include the heartbeat, depth sensor body scan, levels of anxiety, stress, concentration, and focus. These inputs ensure that the performance always has an element of unpredictability.

Motion Tracking

Tekh, Hammond, and Burns (2014) present several pieces that use a depth sensor to manipulate musical parameters and generate a virtual representation of the performer. The sound control parameters vary as the piece progresses, allowing gestures and sounds to change dynamically. Chris Vik, the performer, states that in order to make the system appear responsive to the audience, movement-to-sound cues are kept simple.

Heap (2011) performs a musical piece using her voice, a depth sensor, and a pair of custom-made gloves called mi.mu. The gloves track hand orientation, posture, direction, finger flex, and inertia, while the depth sensor camera tracks the location of the performer. These enable precise and diverse control over sounds. However, the engineers behind the gloves could not accommodate all the gestures because of sensor conflicts. Simple audio gestures were regarded as effective and engaging.

Siegel (2009) describes a series of experimental dance performances called *The Pandora Project* that used a combination of computer vision and wearable accelerometers to sonify the movements of dancers. The authors found that filter and volume manipulation produced the most intuitive and expressive relation between sound and movement. The accelerometer offered dancers precise control over sound, but the computer vision system provided a more obvious relationship. Sound triggers based on thresholds became tedious and limited the dancers' freedom to move.

James et al. (2006) opted to use a marker-based motion capture system to give a dance trio control over musical and visual parameters on stage. Rather than focusing on individual gestures, their piece *Lucidity* measured five group dynamics: proximity, translational correlation, grouping, formations, and body motion correlation. The piece was divided into three sections: In the first, proximity would thicken the musical texture, while projected lines would swirl faster. In the second, position and activity would affect the fitness function of a genetic algorithm while producing a 3D curved surface sweep. In the last section, activity produced rhythmic bell-like sounds while the visuals created a history of the dancer's position with nodes and connections.

The authors state that conventional fixed choreography is unsuitable for interactive systems. Instead, works need to contain structured improvisation or free movement to allow dancers the opportunity to contribute their own uniqueness to the work as well as manipulate the audiovisual environment.

Clay et al. (2012) also explored the issue of intuitive and expressive relations between sound, image, and movement. They developed an augmented live dance performance that combined inertial motion tracking, body emotion recognition, interactive music, and an emotionally responsive virtual robot. In the show, the dancer performed various scenes highlighting different technical and artistic possibilities of the interactive system.

Feedback was collected from the choreographer, dancer, and audience in the form of a questionnaire with rated questions. For the choreographer, technological constraints did not restrain creativity, but oriented it. However, controlling music was not considered in the choreography. This goes hand in hand with the dancer's impression of low control over the responsive music. Nonetheless, even if the relationship between movement and sound was obscure, the resulting music and visuals had an important impact on improvisation. In the end, the audience favored direct, straightforward interactions. Innovative and magical interpretations by the audience are related to the staging of the show, while the artistic focus should be on direct interactions.

In their piece *Hakanaï*, Mondot & Bardainne (2013) expanded the application of virtual reality by immersing a dancer inside a four-walled cube composed of a translucent mesh that displayed interactive projections. Visuals included numbers, letters, lines, and grids. With the spectators sitting around the box, the dancer's proximity to the visuals created obvious effects like deformation and object translation, while the music reacted in subtler ways. The dancer was given the freedom to improvise within the structure of the piece.

Biosignals

Novello (2013) attempted to sonify and visualize the thought process of a performer as he tries to navigate an avatar through a maze projected on the stage floor. His piece *Fragmentation* aims to communicate to the audience the performer's mental stress, stimulation, and saturation through graphs, an oscilloscope, and glitch audiovisuals. Using an EEG, brain signals are sonically sampled, looped, and compared to three brain signal patterns that serve as commands for the avatar. The performer must remain calm and focused despite the glitchy sounds and flickering visuals.

In *Fragmentation*, the most stable brain patterns were thoughts related to imagined kinetic movements. Increasing the number of thought patterns that the system could identify beyond three made it very unstable. The author concludes by stating that direct manipulation of the avatar produced higher involvement from the audience.

Donnarumma (2012) describes his piece *Hypo Chrysus* as located at the intersection of bioscience, music technology, performance art, and human affect. Microphones placed on the arm capture blood flow, muscle contraction, and bone crackles as the performer tries to drag two concrete blocks around the stage for 20 minutes. The sounds are amplified, distorted, and reproduced through eight speakers. Sounds also affect the visuals projected on a screen. Donnarumma concludes that using muscle tension to control musical parameters requires great concentration, since there is no tactile feedback as with a traditional instrument. Furthermore, practice increased proprioceptive sensitivity and flexibility.

Chapter 3: *Curie*

Date: January 5 and February 1, 2016

Venue 1: Consumer Electronics Show, The Venetian, Las Vegas

Venue 2: Anaheim Convention Center, Anaheim, California

Video: https://vimeo.com/170378294

Curie opened Intel's keynote at the Consumer Electronics Show in 2016. The four-minute show featured electronic music with gesture-controlled instruments, a live band, two dancers, and quasi-holographic projections. The interactive audiovisual elements were controlled using Intel's Curie-based wristbands and RealSense camera technology. Aside from the performers, audience members were also given wristbands to contribute musically and visually. The project took three months to develop and encompassed a wide range of technological and artistic elements including sensors, gesture algorithms, radio transmission, music, visuals, and dance.

This case study will start by contextualizing *Curie*, followed by an in-depth description of the artistic and technical development. Because of its complexity, *Curie* offered a rich source of information about the factors that influence the interactive design. I conclude this chapter with an analysis of the technical, performative, and collaborative concerns that emerged from the project.

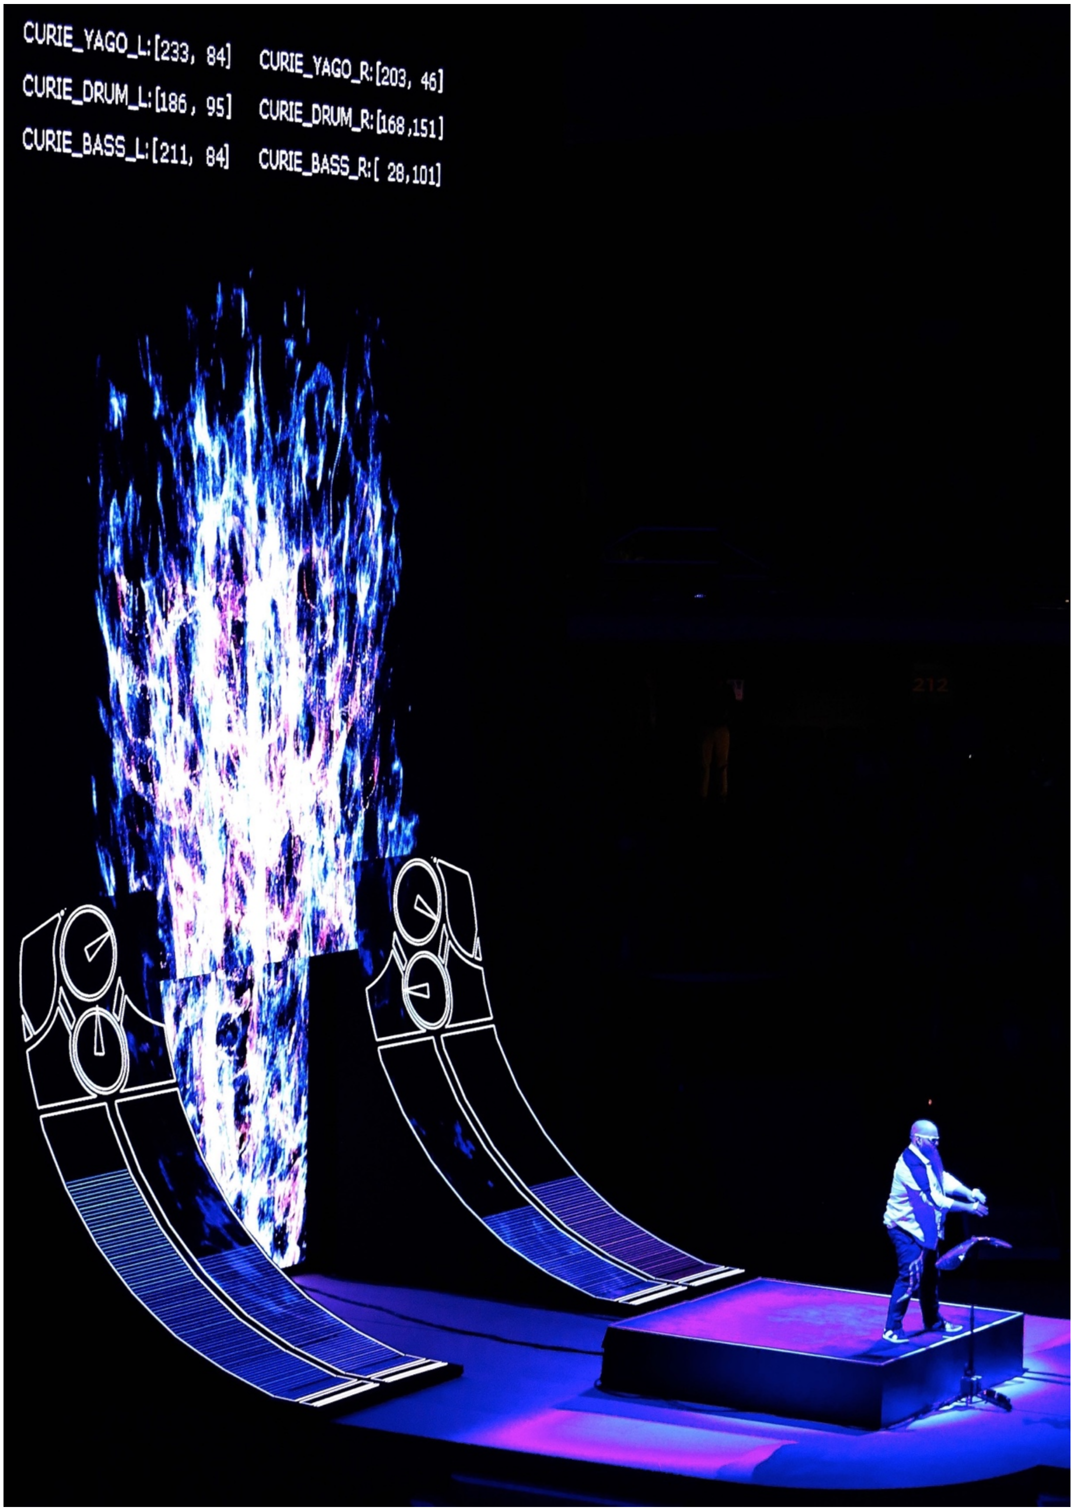
CURIE_YAGO_L:[233, 84]
CURIE_YAGO_R:[203, 46]
CURIE_DRUM_L:[186, 95]
CURIE_DRUM_R:[168,151]
CURIE_BASS_L:[211, 84]
CURIE_BASS_R:[28,101]

Team

	Talent
Visual Artist	Peter Sistrom
Composer	Zackery Wilson
Performer	Yago de Quay
Choreographer	Reema Bounajem
Costume Designer	Kelsey Vidic
Dancers	Reema Bounajem Shoko Fujita
Bassist	Mitchell Vandenburg
Drummer	Adam Nurre
	Marketing
Executive Producer	Bill Welter
	Engineers
Logistics	Lakshman Krishnamurthy
Programmers	Stephen Xing Manan Goel Matt Pinner Jason Blanchard Tal Ein-Gar
Sensors	Saurin Shah
Algorithms	Swarnendu Kar
900 MHz Radio	Phillip Sitbon
UWB Radio	Tim DeBaillie
	Operators
Audio Rig	Chris Legaspi Fred Carlton
Video Rig	Evan Pierre

Table 3.1: The team responsible for *Curie*.

Context

Origin

Intel invited me to contribute to the development of a motion-tracking wristband and perform at the Consumer Electronics Show. Their goal was to showcase the new Curie Module by expanding on a gesture-based instrument created in collaboration with A. R. Rahman, a famous Indian composer and performer. We were also asked to use the RealSense camera for hand tracking.

Due to the complexity of the performance, various experts were invited to address specific areas. I was allowed to invite Reema, the choreographer, and Peter, the visual artist, to aid me in creating the artistic content for the show. The technical team consisted of operators hired by me as well as Intel engineers. The contracted rig operators had the important role of ensuring that all the software booted and performed correctly. Intel's software engineers created the gesture detection algorithm, the Curie Module firmware and Mac software that interpreted the gesture data from the wristbands. Ciholas built the ultra-wideband radio array and the server racks that managed the data packages. The goal of the show was to create an immersive audiovisual performance that responded to gestures.

Stage. The performance was presented on two very different stages. The Bellagio ballroom at the Consumer Electronics Show in Las Vegas had a proscenium stage with a scrim as a projection surface. All the performers on stage were close to the scrim and facing the audience. The 360-degree stage at the Anaheim Convention Center had a circular center stage connected by two bridges to two semicircular stages on either side. The audience sat on both sides of the stage. The projection surfaces for this venue included the floor and two LED walls on the sides. This case study will focus mostly on the show for the Consumer Electronics Show, since most of the development was geared for it.

The Consumer Electronics Show stage was 24 meters wide and six meters deep. The floor and walls were mostly white so they could diffuse projected light. A black Rose Brand shark-tooth scrim covered the width of the stage and was four meters high. There was a 2.4 m^2 ultra-wideband radio platform (used for motion tracking) in the center flanked by two 2.4 m-by-3 m instrument risers. The 900 MHz radio receivers (used to communicate with the wristband) outlined the front, back, and sides of the stage. Cameras above the audience captured the wristband activity of clusters of 200 users. *Figure 3.1* illustrates the placement of the sensors, scrim, performers, and audience.

Anaheim Convention Center has a long 360-degree stage and housed around 5,000 attendees. I performed in the center, the musicians on the bridges, and the dancers next to the outer LED walls. The center stage was 10 meters in diameter, the bridge was eight meters long, and the semi-circular stages were nine meters wide. The dancers and I stood over an ultra-wideband (UWB) radio platform that tracked the position of our hands. The projection area covered the circular center stage as well as the floor around it, as depicted in blue in *Figure 3.2*. The show on this stage did not have an overhead camera system monitoring the audience's activity.

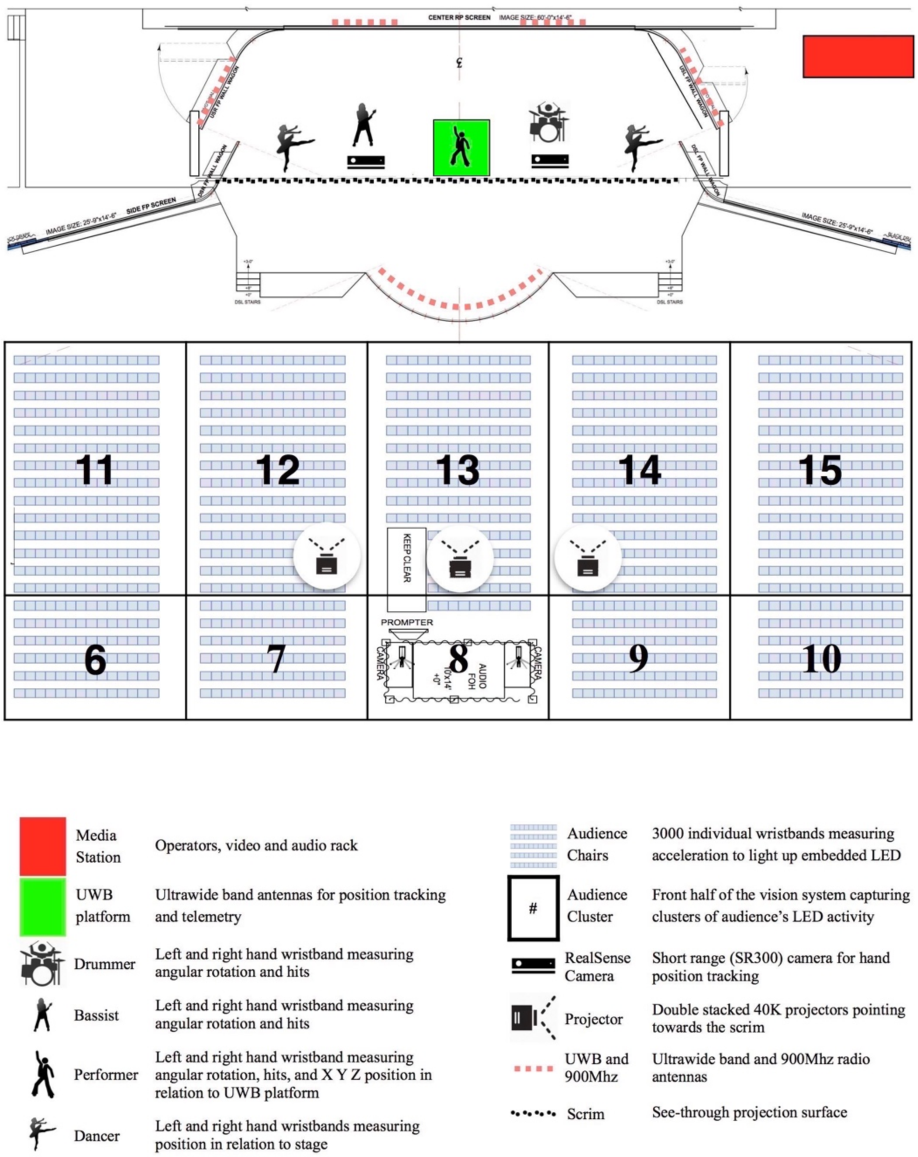

Figure 3.1: CES stage and audience setup.

Rehearsal space. Before the performance, *Curie* was rehearsed in three other locations. The first space in Evansville, Indiana, was used mainly for "blocking" and testing the ultra-wideband and 900 MHz sensors, since the company that developed our ultra-wideband technology was located there. The second rehearsal space was a studio in Las Vegas, Nevada, that had a similar setup to the main Consumer Electronics Show stage. For the Anaheim concert, we rehearsed at an NBC studio in Studio City, California.

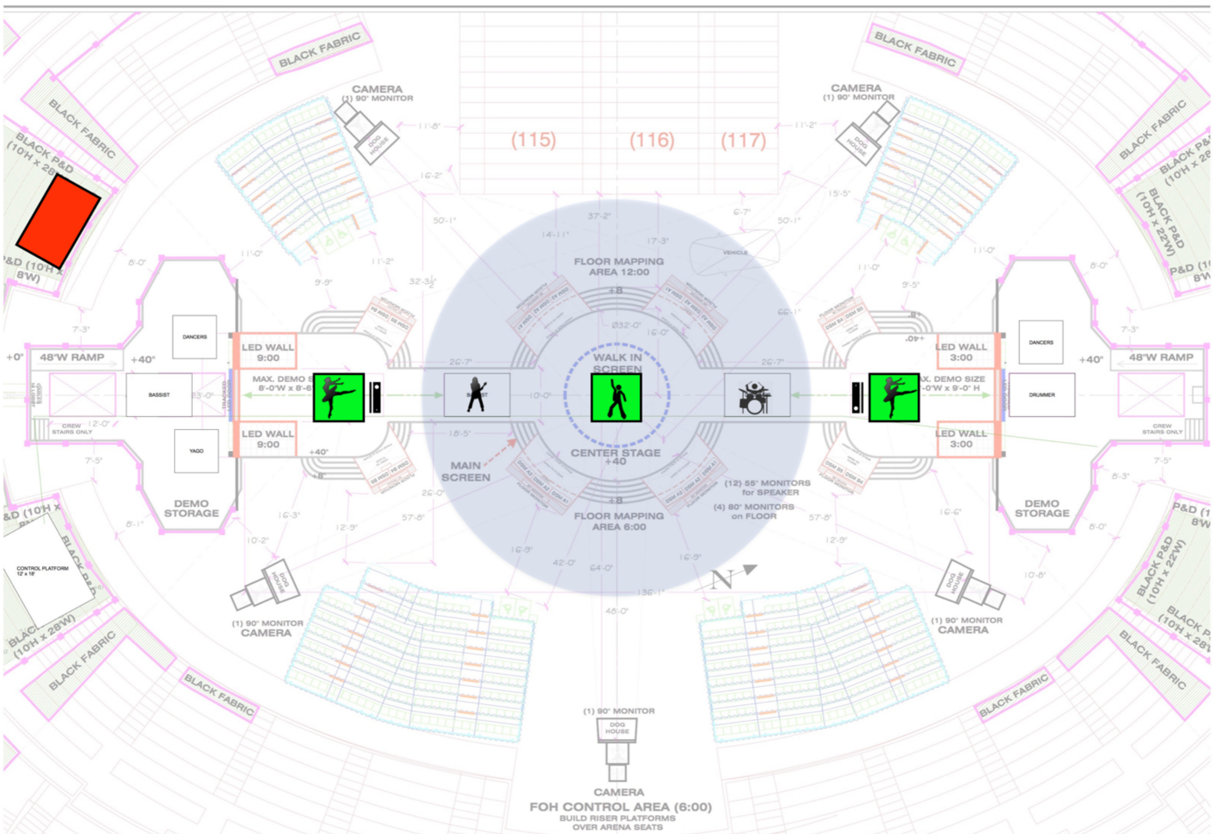

Figure 3.2: 360-degree stage diagram including performer positions (off-stage labels refer to platform storage).

Media Technology

Media Station. The media station backstage housed the operators and the computers generating the music and visuals. These computers mapped multimedia content to performer and audience hand positions and gestures. The ultra-wideband antennas communicated with, and triangulated the location of, the performers' sensors. The 900 MHz antennas served as a backup communication channel. The RealSense cameras were placed on a jig in front of the musicians so they could walk up to it during the performance.

Sound. The venue had a speaker array on both sides of the stage. The mix was mostly monophonic so that the whole audience would listen to all the instruments regardless of their seats. The musicians and I had an in-ear monitor mix where they could listen to the click track and fixed musical layer.

Projection. Three projector nodes—each stacking two 30,000-lumen projectors—were rigged in the ceiling and converged their images on the scrim in plain view of the performers and audience members. Backstage, everyone in the media station saw the performance through a TV monitor.

Lighting. The production team placed dark blue lights on the floor and ceiling next to the scrim. These lights did not interfere with the visuals on the scrim because it did not diffuse or reflect dark blue light.

Audio rack. The audio rack was a crucial component of the setup, consisting of computers, audio interfaces, and data management software. The audio interfaces also had an A/B (primary/backup) configuration. Mac Pro A and a Mac Mini C—illustrated in *Figure 3.3*—were the primary computers, while a Mac Pro B and a Mac Mini D served as backups. The KVM switch switched between the A and B computers. The MacBook Pro hosted Syncphony, The State Machine, The Portal, Ableton Live, and data management and networking software. The Mac Mini monitored remote configurations. *Figure 3.4* displays how the different components were interconnected.

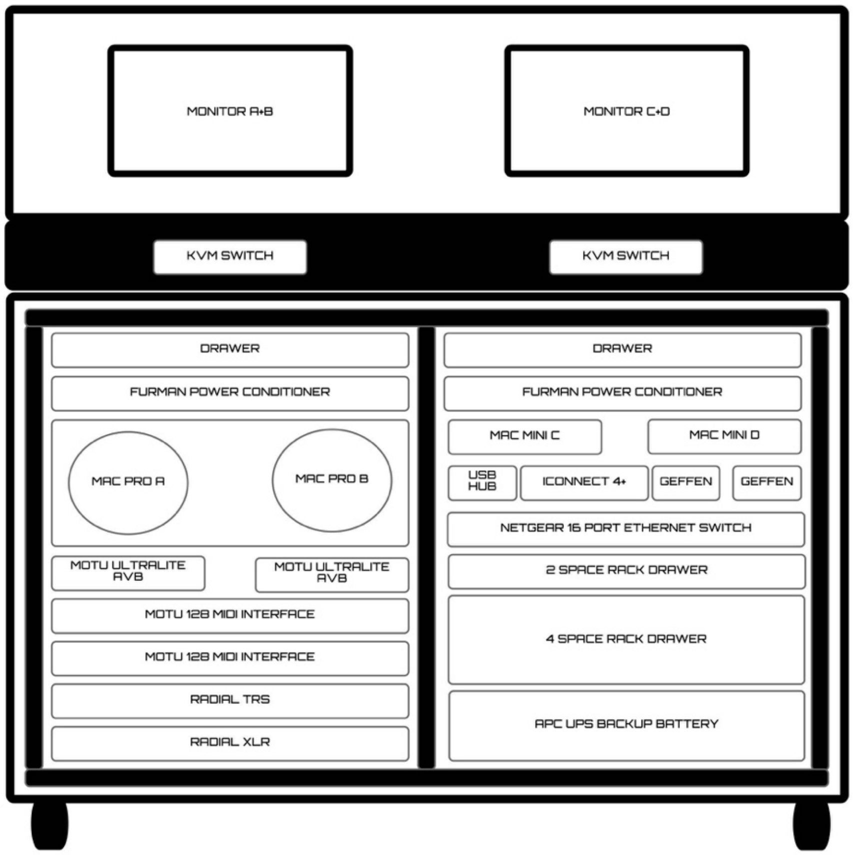

Figure 3.3: Audio rack hardware placement diagram.

Signal network. Our network had four gesture input streams and three media outputs. The musicians and dancers communicated digitally over 900 MHz radio, ultra-wideband radio, and RealSense, while the audience used the vision system. The three audiovisual outputs consisted of projectors, audio speakers, and lights on audience wristbands.

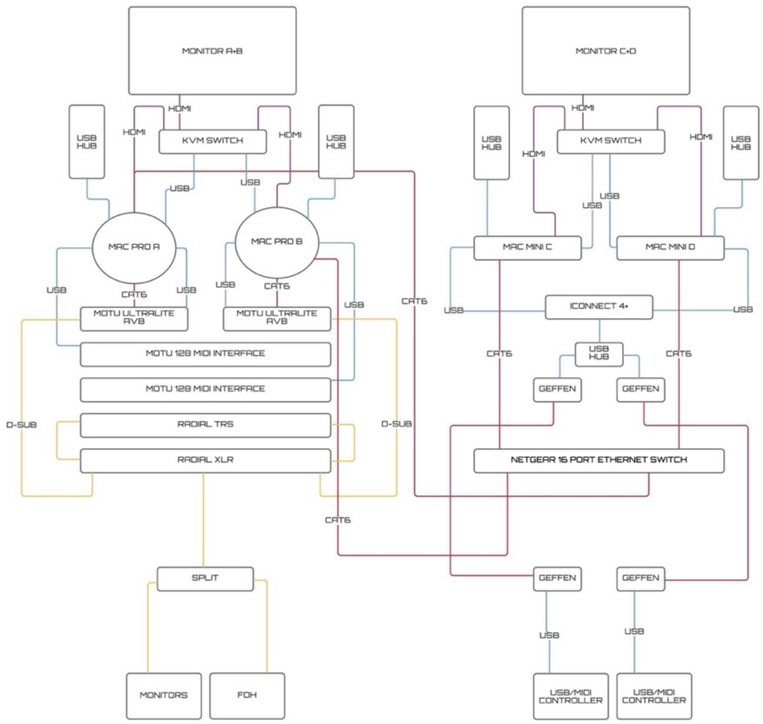

Figure 3.4: Audio rack data, audio, and MIDI connections.

Figure 3.5 delineates the different interactive input/output components. Those on the left side were responsible for capturing gestures, on the middle processing the data, and on the right outputting the media. Data from the performers' wristbands were transmitted over ultra-wideband and 900 MHz to radio receivers positioned around the stage. The ultra-wideband receivers consolidated the data into a server before going to the Ethernet hub. The 900 MHz radio receivers were connected by USB to an Intel Next Unix of Computing (NUC). The RealSense was connected to two Intel NUCs, one for each sensor. The camera system on the ceiling monitoring the audience's wristbands was wired to 20 Intel NUCs and then processed by a master server. All the sensor traffic converged into one Ethernet hub on the audio rack before being distributed to other computers.

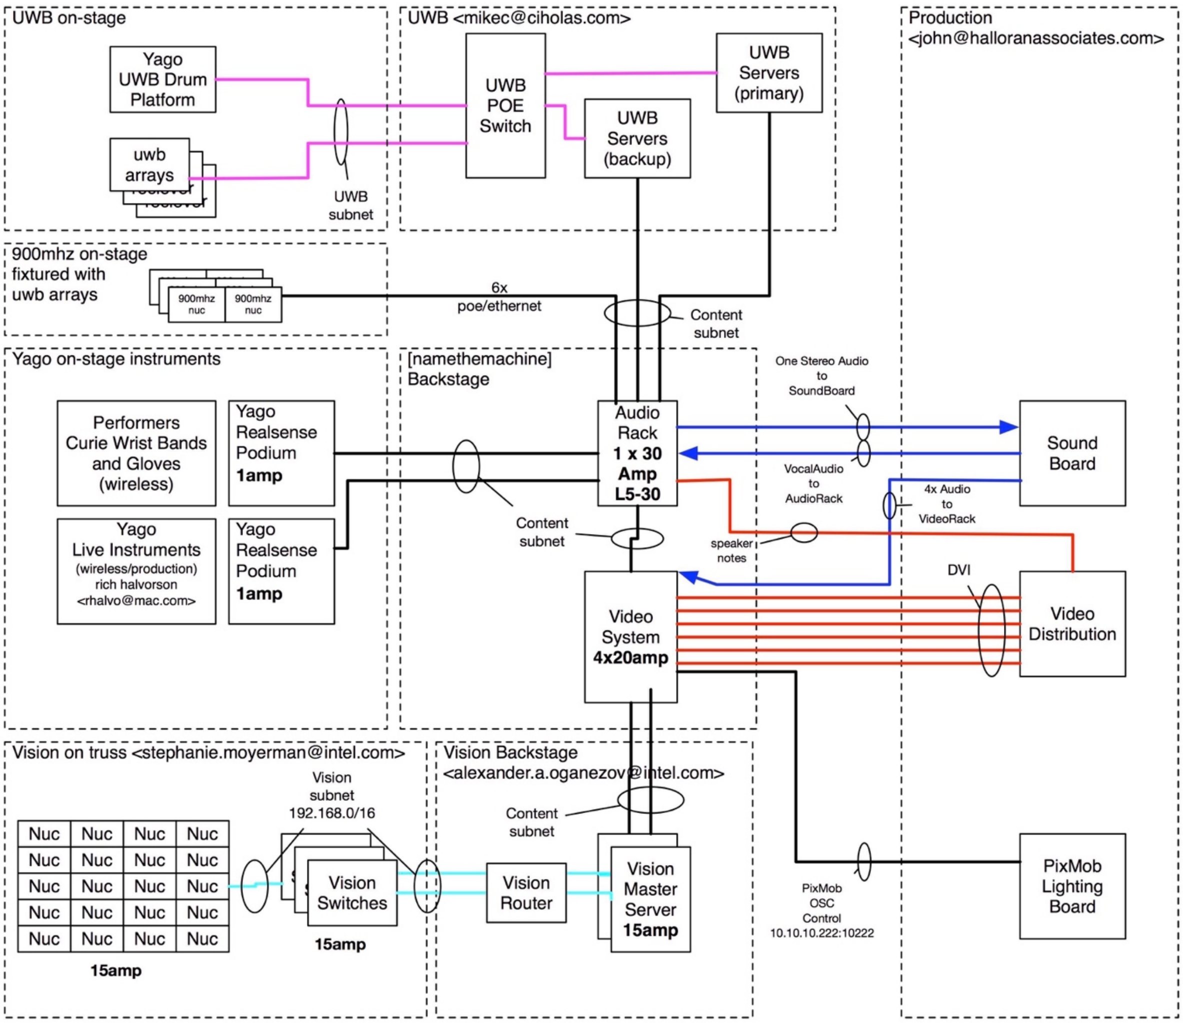

Figure 3.5: Sensor, signal and media network. On the top left, the ultra-wideband and 900 MHz radio systems connect with the musician and dancer wristbands. The RealSense and wristbands tracking the gestures of the performers on are located on center left. The bottom section comprises of the audience motion tracking system. The center and right sections illustrate the media stations and the multimedia output paths.

The audio rack, video rack, and PixMob computer pulled the sensor data from the Ethernet hub and applied it to audio and visual effects. Audio was output to the sound board, while the video signals were sent to a video distribution rig. The PixMob computer used the sensor data to communicate with the audience's wristbands via infrared blasters.

Sensors

Intel developed three technologies to capture and transmit hand gestures. The Curie Module was a small processor and sensor unit that was embedded into a wristband and programmed to identify and report gestures. Intel's RealSense camera was an infrared sensor that, together with a small computer, could determine the location of a user's hand at close proximity. The ultra-wideband radio array located the position of the wristbands while transmitting their data.

Intel Curie Module. The Curie module is a tiny, low-power chip with a 32-bit Intel Quark microcontroller and a six-axis accelerometer and gyroscope. To track movement, the chip was either placed on a 3D-printed elastic wristband or fitted into costumes. *Curie* employed three motion-tracking gadgets based on the Curie Module: musician wristbands, audience wristbands, and dancer tags. Audience wristbands communicated via infrared and visible light. The rest communicated wirelessly through 900 MHz and ultra-wideband.

The chip's on-board algorithms detected gestures, managed data packages, and emitted signals for triangulation. Musician wristbands reported hand rotation and hit movements. Audience wristbands reported hits, while dancer tags served only for location tracking. Musician wristbands would turn red with a down hit, green with a side hit, and blue with rotation. Audience wristbands temporarily turned red with a fist pump. The musicians and I had two Curie Modules on each wristband for redundancy; if one module failed to register or transmit a gesture, the other would take over.

Figure 3.6: A recently assembled pack of audience wristbands presented by Lakshman Krishnamurthy.

Intel RealSense R200. The Intel RealSense R200 is an optical sensor combining an infrared camera and laser source to locate a musician's hand position. The sensor was placed on a stand under the hands of the performer and connected by USB to an Intel NUC which used the 3D-point cloud data to calculate the vertical, horizontal, and depth position of the hands. Hand position was tracked between 10-30 cm above the camera using an algorithm developed by Intel. Each stand had two cameras facing up and slightly out. A translucent plastic covered the sensor and housed a reactive LED strip that would rotate a color wheel as the hand moved horizontally. The final design of the controller is shown in *Figure 3.7*.

Figure 3.7: RealSense camera stand for hand position tracking.

Ultra-wideband radio. The ultra-wideband radio was a low-latency positioning and network technology. The ultra-wideband system consisted of an array of radio anchors strategically placed around the stage that transmitted gesture data from the performers' wristbands while also determining their positions. A dedicated computing server determined the location of the wristbands by measuring tiny discrepancies in the data's time of flight to the different anchors. This data was then routed to the audio and visual rigs. The ultra-wideband anchors were mounted on wooden structures surrounding the stage to reduce occlusion, as shown in *Figure 3.8*.

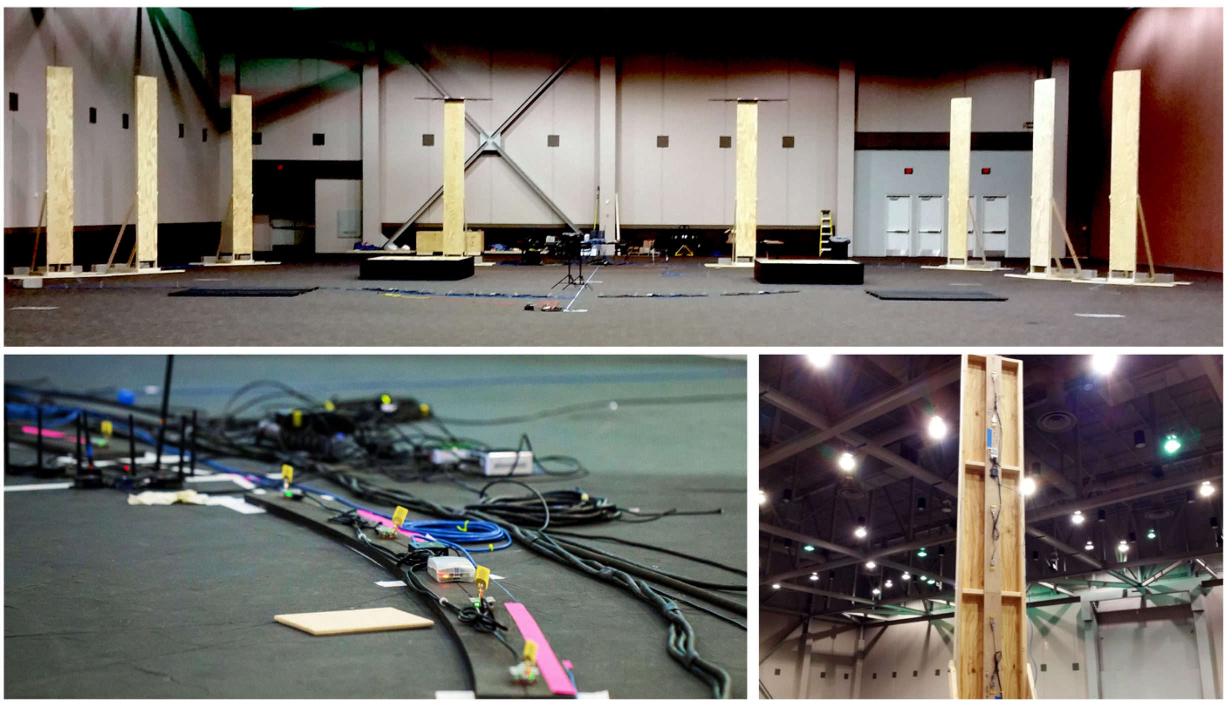

Figure 3.8: Wooden structure surrounding the stage that held the UWB radio anchors.

DEVELOPMENT

Curie required extensive technological development. The first section will look at the creation of the musician and audience wristbands and the gesture-detection algorithm. Following that, we will examine how Intel's depth camera was hastily incorporated into the show. The extensive yet invisible communication and local positioning radio system will be discussed at length. The software that handled the incoming sensor data will be explained afterwards.

The show's mapping scheme went through four major revisions before arriving at its final form documented in the list of mappings. The last mapping section narrates the evolution of the mappings in response to technological and performative challenges. Peter and Reema give a detailed description of the visuals and choreography that includes their vision, progress, and challenges encountered. I describe how the music was created for the show. By charting the evolution of these components, we get a better idea of what factors are implicated in the design of an interactive system.

Repurposing Intel's technology to enable motion tracking required developing sensors, building a wireless infrastructure, creating gesture algorithms, and coding software. The next sections will explore each of the technological developments in detail.

Intel had to build wristbands that encased a circuit board fitted with a Curie Module and a radio antenna. The PixMob company expanded on this technology by building 3,000 units for the audience. Swarnendu Kar worked together with the performers to create the gesture-detection algorithm for the wristbands. Hand tracking was also achieved with the incorporation of a depth-sensing camera. Ciholas built the wireless radio system that tracked the position of the wristbands. The last sections describe software that handled the sensor data.

Musician Wristbands

Intel made wristbands for the show that had the capacity to detect fast hit movements and hand rotation. The former was useful for percussive instruments, the latter for continuous control of audiovisual effects. Once the data was transmitted through a wireless protocol, it was interpreted and recast by an application called Syncphony. The development of the wristband's hardware focused on four areas: radio transmission, processing unit, firmware, and form factor. There were two stages in the development of Intel's wristbands.

My initial mapping experimentations were done using Myo's inertial sensor, as seen in *Figure 3.9*. The gyroscope and accelerometer information enabled me to derive three axes of rotation and downward hits. Myo communicated over Bluetooth to a USB adapter, which resulted in unreliable connectivity and high latency. Getting the gesture data from Myo over Open Sound Control (OSC) required running the software Myo Connect and Myo-OSC. (OSC is a communication protocol for computers and music software.)

Figure 3.9: Myo armband and graphs showing muscle and gyroscopic data.

Intel's first version of the wristbands, shown on both of my hands in *Figure 3.10*, consisted of a single black box containing a Piccolo microcontroller that detected downward hand hits and rotation (pitch and roll). The resolution of the pitch and roll data had to be specified on the firmware and was intended to report only if the user was within an angular range. This wristband transmitted data over Bluetooth or 900 MHz radio (the white cable coming out of the box was the antenna). The data was encoded in a MIDI format and could only be accessed in Syncphony. Changes to the algorithm or MIDI format were done by uploading profiles via Syncphony. The 900 MHz implementation was inconsistent and the Bluetooth connection suffered from high latency.

The second version of the wristbands, depicted in *Figure 3.11*, improved connectivity, refined the visual look, and incorporated the Curie Module—the technology

Intel was eager to demonstrate. The top box communicated over 900 MHz and the bottom over ultra-wideband, since many of the latter's radio receivers were on the floor. This duplication ensured exceptional network reliability, because if a data package was lost on one radio the other would cover it.

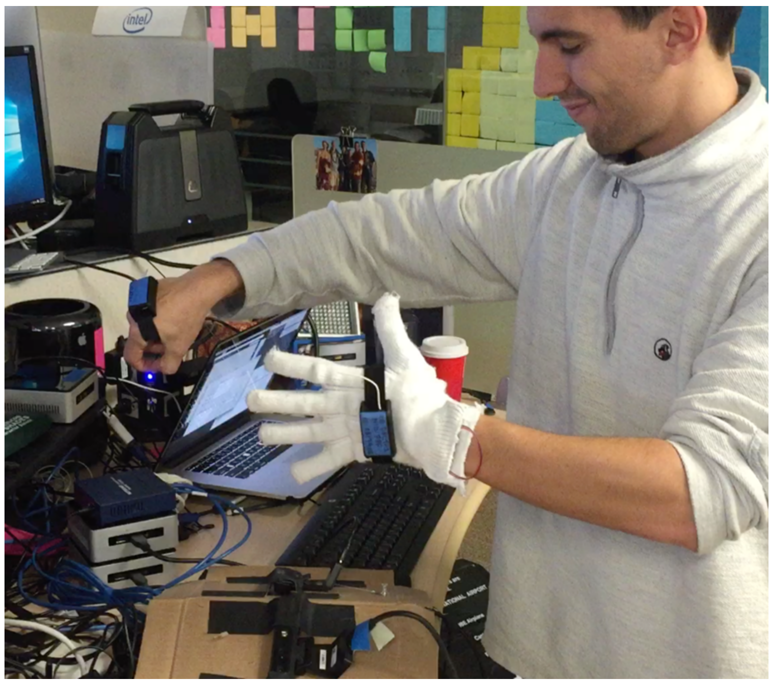

Figure 3.10: First version of the gesture capture wristbands on both my hands. I am also wearing a white glove that reports finger flexion.

The radio network produced much less latency than Bluetooth and did not require authentication. The Syncphony application interpreted the data packages and broadcast them using OSC. The wristband could be tightened with a Velcro strap on one side. Intel added LED lights that responded to hits and rotation. The features of this wristband included:

- Pitch and roll rotation reported every 5 degrees
- Blue light every time rotation was reported
- Detection of fast side and down hits
- Red light for down hits and green for side hits
- Dual-processor and dual-radio
- 900 MHz on top and ultra-wideband on the bottom
- Open Sound Control protocol

Figure 3.11: A close look at the last wristband prototype.
It hosted an Intel Curie Module that could detect fast hits and measure rotation. The top (image right) unit communicated over 900 MHz while the bottom (image left) used ultra-wideband radio.

Gesture Algorithm

The gesture algorithm developed by Swarnendu Kar detected and described specific hand motions using Intel's wristbands. These gestures were performed by the musicians and mapped to percussive, harmonic, or melodic instruments. The algorithm was looking at four features of a gesture: 1) when a gesture was executed; 2) how long it lasted; 3) intensity; and 4) direction. There were two gestures: down tap or side-swipe.

The algorithm used mainly gyroscopic data to track hand movements. The wristband's acceleration data proved unreliable for the musicians because it drifted. Swarnendu explained:

> As you move your hand more and more within a particular motion sequence, the direction of the gravity as perceived by the algorithm deviates from the true direction of the gravity.… So that was the feedback we got from the performers.… They were performing several of those gestures in a short amount of time, and because the device was losing track of gravity, it was either not reporting the trigger or the velocity was being estimated inaccurately (personal communication, March 10, 2016).

To prevent drift, Swarnendu swapped the accelerometer with the gyroscope data, slightly altering the gesture in the process:

> So then we had to think of using a different sensing mode and ultimately figured out that gyroscopes are the kind of sensors that are not dependent on the direction of gravity. There was a little bit of tradeoff.… When we switched between the sensors we had to fix the axis in which we are looking at, so now our axis became pretty much tied to the orientation of the device, not the hand. Each of those devices had to be programmed differently based on how a performer wore it and which location of the hand or the leg. (personal communication, March 10, 2016).

The performers also had to adapt to the change: the tap or swipe motion now required a slight rotation of the wrist. Although it increased the complexity of the gesture, swapping linear for angular acceleration proved very reliable. "It increased the complexity a little bit compared to the acceleration-based [version]. But overall it resulted in better accuracy so far as the performer's requirements were concerned," he noted. The accelerometer was "used as a sanity check that had a veto power over the final decision. But was mostly used to weed out false positives."

Figure 3.12 illustrates how gestures were detected and converted to MIDI velocity values. Swarnendu described what the figure means:

> I am showing 3-axis gyro data of 2 seconds, during which there are four taps – two in the X-direction (down hit), one in the Y direction (wrist rotation), and one in the Z direction (side hit). The tap-detection timing (black dots) is based on peak detection, event segmentation (bold overlays) is based on backtracking following peak detection. [MIDI] velocity (numbers specifying theta) is based on integrating the angle over the event segment. (personal communication, March 10, 2016).

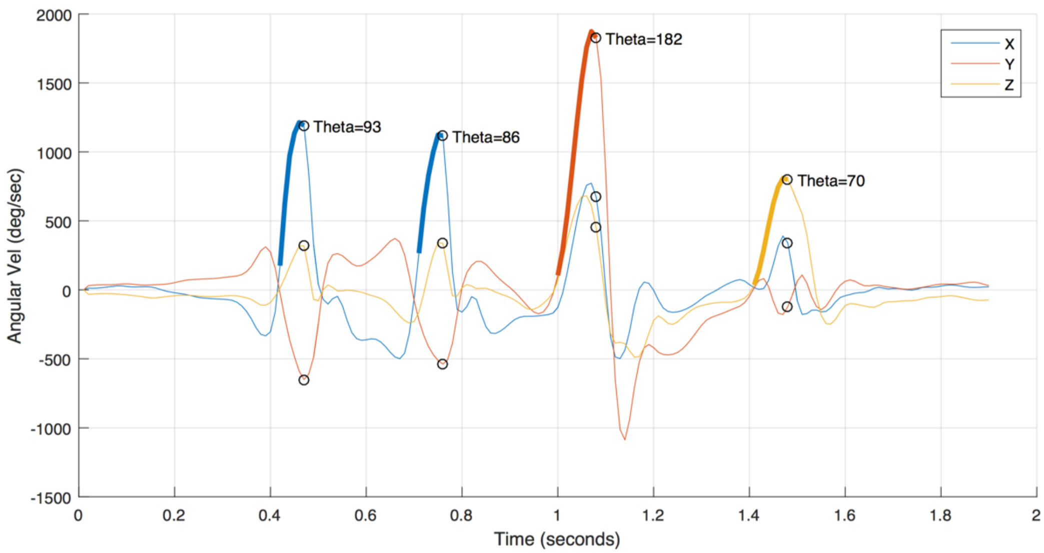

Figure 3.12: Gyroscopic data of four hits over two seconds with peak detection (black dots) and velocity measurement (bold overlays). [3]

By working with the performers, Swarnendu discovered that playing directional gestures was more intuitive than zoned drums. He originally attempted to use a magnetometer to identify unique locations around the performer with the goal of creating zoned drums. This would have allowed the performer to play different sounds depending on where he was hitting. However, the magnetometer proved unreliable because of usability issues, magnetic interference, and the impracticality of using the Earth's magnetic field. Intel realized that using directional gestures would be more effective after looking at how the musicians and I were using the wristband. Swarnendu explained:

---

[3] Used with permission by Swarnendu Kar.

> We started with just one trigger—the down direction—and then thought of adding another kind of trigger. We started trying to achieve positional drums, but we couldn't quite achieve that.… Even with only two positions the magnetometer was giving us inaccuracies in terms of finding out in which position the trigger was happening.… So we had to look at the performers and we figured out that maybe there was something else that we could do. What we found was that in order to make the transition from [zone] 1 to [zone] 2, the performer actually ended up doing a sideways movement—it is difficult to move your hand sideways and then go down. We found out that maybe having one downward hit and one sideways hit was something that the performer would be able to do naturally. It also fit well with our transition to a gesture-based algorithm where we could see the sideways hit and actually translate it to a rotation in a different axis. (personal communication, March 10, 2016).

Latency was reduced by making the network system more efficient and incorporating the gyroscope in the gesture algorithm. Each network upgrade enabled a faster transmission of the signal, starting with Bluetooth, then 900 MHz radio, and finally ultra-wideband radio. The incorporation of the gyroscope serendipitously sped up detection by 20 ms because the angular acceleration peak happens earlier than the linear acceleration—halfway through the gesture rather than at the end. "The latency improvement was almost a welcomed byproduct of the algorithm switch from accelerometer to gyroscope," Swarnendu said.

All in all, Swarnendu said he was happy with the results. In future versions, he hopes to increase the amount of control given to the performer by incorporating note duration, more gestures, and continuous control.

Optical Motion Tracking

Intel introduced us to three different RealSense (depth sensor) cameras for motion tracking, the R200, F200, and SR300. Each camera was engineered differently and their features varied widely. Neither of them had robust skeleton tracking, but the SR300 was able to track hand position reliably at close range. The teams responsible for each camera's development insisted on us using them and offered help with coding and construction. *Table 3.2* specifies a few of the features of each camera available at the time of writing.

Name	Image	Skeleton Tracking	Face Tracking	Hand Tracking	Range (m)
R200	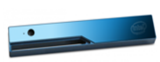	X	X		3-4
SR300			X	X	0.2-1.5
F200			X	X	0.2-1.2

Table 3.2: Tracking capabilities of three of Intel's RealSense cameras.
We ended up using the SR300 because of its robust hand tracking.

Because the RealSense cameras were relatively new, most of the software features were immature and unreliable. Few of the software development kits (SDKs) had been compiled into standalone programs that would work automatically when booted. The SDKs also lacked network capabilities. We proposed using Microsoft's Kinect camera because it was effective at extracting skeletal data, but Intel was against using a different company's depth-sensing camera technology. We experimented with the F200 camera, but hand tracking was inconsistent and jittery. One Intel employee complained that F200's "range is really small—two feet—and it is unreliable" (personal communication).

A month into the project we discovered a hand tracking prototype at Intel that used the SR300 camera (an upgrade of the F200), which accurately and reliably located hand position. The software's default settings required only a few modifications and network integration. To further develop the prototype, Matt Pinner and I flew to Intel's headquarters in Santa Clara, California, and worked directly with the team responsible for the hand tracking software. This hands-on involvement hastened the development and ensured that our specifications regarding the hardware, network, and use case scenario were adhered to.

Intel's design division worked independently to enclose the SR300 in a wide metal and acrylic stand which proved to be impractical (*Figure 3.13*). It had an LED color wheel on the top that responded to hand movement of any connected camera, which meant that any local hand movement was overridden by other performers' hand movements. Aside from the color wheel's limited ability to represent the X-axis boundaries, there was no way to know the boundaries of the Z and Y axes. Because the hands were out of bounds, the algorithm frequently lost track of them. I tried to specify beforehand that the stand should be simple and easy to travel with, but I was not involved in its creation. The complexity and bulkiness of the stand made experimentation difficult.

Figure 3.13: Adam playing the RealSense instrument that uses Intel's SR300 camera to track the location of the user's hand.

Radio Motion Tracking

Commissioned by Intel and built by the engineering company Ciholas, the ultra-wideband (UWB) radio array was a technology that doubled as a wireless network and real-time location system. With it we could communicate with the wristbands and determine their location. Tim DeBaillie, who supervised the development of the technology, explained:

> Our enabling piece of technology is a DW1000 DecaWave chip on our main board. It's an ultra-wideband radio frequency chip. It has a very precise clock that can measure

> when a reception comes in down to 15 picoseconds of resolution. That allows you to tell distances based on the time-of-flight or other methods…. You can schedule a transmit at a specific time and then expect it to be sent precisely at that time and on the other end you can measure what you actually got. The further step that we did that takes it to the next realm is that we have a unique algorithm that actually takes that [time-of-flight] data and calculates a very precise position. (personal communication, March 11, 2016).

One major benefit of UWB is that it operates on frequency bands centered around 6.5, 4.5, and 3.5 GHz that are mostly free from interference. "It was really nice to have UWB because no one is really operating in that band to this date," Tim said. "900 MHz band is very widely used for handheld radios. 2.4 and 5.4 GHz are used very heavily for Wi-Fi. So in a case like at [the Consumer Electronics Show] where you have 2,000-plus people in the audience and they all have a cellphone that's trying to get a Wi-Fi signal, it can be very polluted." Furthermore, packages sent over UWB are spread across different frequency bands, reducing the odds of complete interference or corruption.

Ciholas combined wired and wireless connections to reduce latency. The wristbands communicated wirelessly with the radio receivers, which communicated through Ethernet and USB to the main server.

"When Intel came to us, they told us what they wanted to do with the [zone drums]. To meet their needs, we used Ethernet and USB as our data link between our [radio receivers] and the host server so that we could get more airtime," Tim said. Airtime is the time a device has to transmit a package wirelessly; when there are many devices transmitting, they compete for smaller and less frequent airtime slices. The radio receivers can only receive one transmission at a time; otherwise there is interference.

Limiting the wireless transmissions to only the wristbands reduced calculations and latency. With the wristbands sending packages at 200 Hz, by the end latency was down to 5 ms.

The show utilized two UWB networks: one built around the stage for the dancers and musicians, and one inside a platform for me. The dancer and musician network consisted of an array of 42 radio anchors. The anchors were placed on a curved platform downstage and on wooden boards screwed to the walls of the stage. The UWB team took 4 hours to install the arrays and 3 hours to survey their position for accurate spatial modeling. My radio array was physically hardwired to the platform, so it required no installation and surveying time. *Figure 3.14* illustrates how the network was assembled on stage.

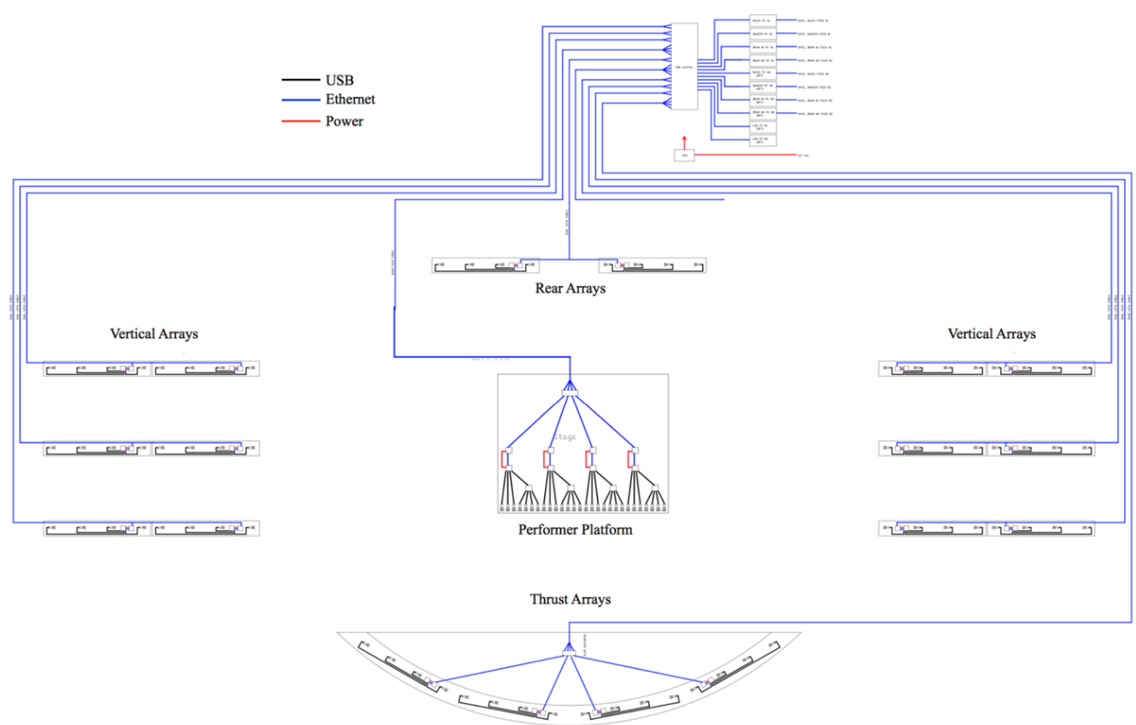

Figure 3.14: Diagram of UWB stage and platform network with arrays and wire paths.

The motion-tracking system had limitations due to occlusion and the placement of the anchors. It suffered signal loss when the line of sight between the anchor and wristband was blocked by a human body. The chances of this happening increased when my hands were close to the scrim and on the edge of the platform. Tim proposed placing all the anchors on the ceiling so they could have a clearer line of sight to the wristbands, but the ceiling was entirely reserved for stage lights and other equipment.

The system was also imprecise—the position data jittered around the wristband's actual location. The height axis of the platform network was particularly unreliable because the anchors were spread across the horizontal plane. Averaging the samples of the position data made it smoother but increased latency by 5 ms for each averaged sample. Furthermore, the rapid output frequency tended to overwhelm the audio and visual software. Once during rehearsals, the whole network broke down because it could not handle all the packages. We reduced the amount of data entering our interactive software by limiting the output from the wristbands: we removed my down hit (unused), reduced the musician's rotation resolution, and reduced the number of sensors on the dancers.

Regarding the collaboration, Tim commented how "the [show] went pretty smooth as far as that was concerned. You guys would come to us and say 'we think we want to do this.' We'd go back to the drawing board and say, is that possible?"

Ciholas expanded the UWB system to accommodate the increasing number of connected performers. Tim pointed out how this enhancement was a major challenge:

“The big challenge that comes to mind is that originally it was supposed to be one drum stage, and then it was two drum stages, and then it was dancers, and then it was musicians. It just kind of grew. We had to adapt to allow more than one network to be run in one space.”

In a conversation with me after the show, Mike Ciholas, owner of the eponymous company, confessed that if he had known earlier and more precisely how the UWB system was going to be used, he would have designed it differently. In the future, Tim said, he hopes to improve anchor placement and the occlusion mitigation algorithm that detects whether anchors have bad line of sight.

Finger Tracking Glove

I experimented with a glove created by Intel and the company StretchSense that could measure finger flexing, but ended up not using it because it was introduced too late. The glove used flex sensors along all the fingers that changed capacitance when bent (*Figure 3.15*). Data was processed in a Curie Module inside a black box on the back of the hand and communicated to a computer via Bluetooth Low Energy. The glove required calibration every time it was put on. Gestures were stored as finger patterns that could be used to trigger events such as note selection. The final gloves, pictured in *Figure 3.15*, were delivered a few weeks before the show. Since all the mappings had already been created for other technologies, I found no use for the device. Furthermore, adding this glove to the system would have increased its complexity.

Figure 3.15: Glove created by Intel and StretchSense that can measure finger flexion. It ended up not being used due to time constrains.

Audience Wristbands

PixMob is a company that specializes in building light-emitting wearables for audience participation. In *Curie*, their system consisted of 3,000 interactive wristbands distributed among audience members, which lit up in response to either audience or musician gestures. An audience wristband served two purposes: to detect gestures, and to contribute to the light show.

During the breakdown section of the show, audience members were invited to contribute both musically and visually by fist pumping. Shaking their hand activated an LED that instructed the interactive system—through light-sensitive overhanging cameras—to play a sound and generate visuals.

Figure 3.16 depicts an audience member wearing an active wristband during the show. During the other sections of the show, gestures performed by the musicians lit the audience's wristbands. Jean-Sebastien Rousseau was responsible for making the middleware that listened to data from the wristbands and output light instructions.

Figure 3.16: Audience member with an active PixMob wristband.

Jean-Sebastien commented that he worked mostly on creating the circuit board for the audience wristbands:

> For us it was more of an engineering challenge to design the circuit board. Intel manufactured the board. We just provided them with the parts we used normally and they added their parts. Just to do that for us and add the firmware was the main challenge. (personal communication, March 24, 2016).

Intel and PixMob developed various solutions to adapt the original PixMob wristbands to the CES show:

- Reducing latency down to 300 ms
- Creating a reliable gesture algorithm
- Increasing the battery life from 15 minutes to a few hours
- Integrating the Curie Module with the PixMob chip board
- Creating a backup plan in case the audience did not participate

Aside from engineering the wristband, PixMob addressed several issues related to integration with the interactive system. Mainly, these issues were related to the network data handling.

> Our challenges [with the integration] were the stability of the data stream, figuring out who was going to send what to whom, if there was consistency, latency, and if we added more wristbands. And the problem for PixMob was that we would never know until the show what it is going to look like because we never have that many wristbands during rehearsals. (personal communication, March 24, 2016).

He concluded by admitting that "the light show was the easy part for us…. The wristbands were working, everything was working, it was just a question of integration."

PixMob managed the input to the wristbands while the vision system captured the light from the wristbands and relayed that information to the audio rig. *Figure 3.17* illustrates the data flow of the whole audience interaction system. A Curie Module detected gestures while a light sensor decoded infrared (IR) messages. Gesture or IR cues were transmitted to the PixMob chip to produce colored light patterns. The IR emitter worked by sending a message through DMX, a digital communication standard for stage lighting, to infrared emitters that transmitted 9-byte data packages at 20 Hz to the wristbands. The vision system consisted of an array of cameras on the ceiling of the venue that measured luminosity from the crowd and relayed that information by OSC to the audio rig. Luminosity was divided into 15 zones: five columns and three rows (refer to *Figure 3.1*).

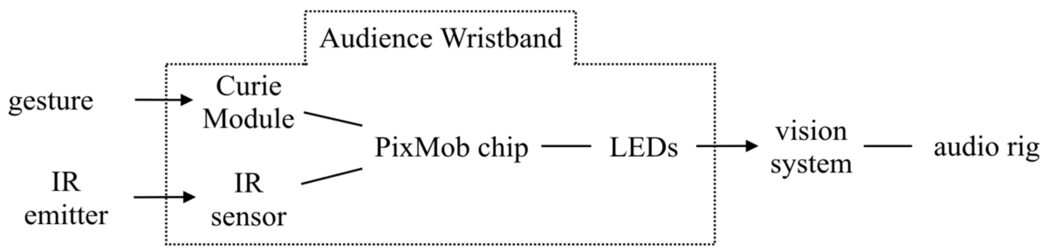

Figure 3.17: Data flow of the PixMob and Vision system.

A simple fist pump was the most responsive gesture to trigger lights and sounds. Initially, Intel suggested teaching the audience to play seven gestures and then splitting the audience into groups that played unique instruments, in a manner similar to an orchestra. Jean-Sebastien explained:

> From Intel's perspective, it was really about having an Intel [Curie] chip in the PixMob wristband and be able to say, 'look at what the Intel chip can do.' We did that. However, the way the gestures were detected on the Intel side was working, but reaction time was slow. We used a simpler technique to trigger things rather than their gesture algorithms. (personal communication, March 24, 2016).

One idea that was never fulfilled because of time constraints was to render shapes on the audience like a massive LED screen. To do this, each wristband would have had to be assigned a location in the audience so that it could act like a pixel. Jean-Sebastien said:

> We wanted to be able to figure out where every single person was in the audience and with the infrared system be able to, let's say, draw a line or draw a particular shape. But we never got to that point; if we had more time we could have done it. (personal communication, March 24, 2016).

Jean-Sebastien said the most interesting moments were the preparations, rehearsals, and the "wow" moment when, "I saw everyone was [participating] just on that row; they actually listed to what you said." He didn't expect everyone in the audience to participate to the extent they did.

Software

Syncphony. Syncphony is software that interprets gesture and position data transmitted wirelessly from the performers' wristbands and sends it as Open Sound Control (OSC) packages to audio and visual software. Each wristband had two active Curie Modules for redundancy—the top one transmitted over 900 MHz and the bottom over ultra-wideband—in case a package was lost over either of the radio frequencies. If both packages arrived, Syncphony detected which packaged arrived first and ignored the other. An image of the Syncphony interface can be seen in *Figure 3.18*.

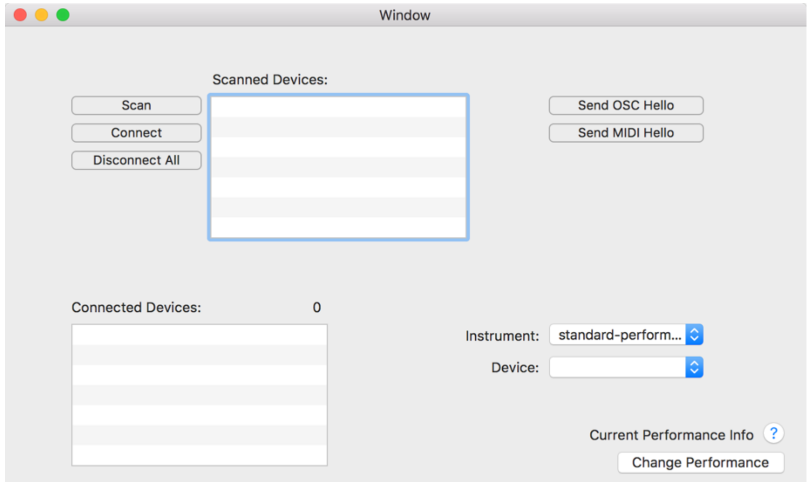

Figure 3.18: Syncphony interpreted data from the Curie wristbands and routed it to other software.

Jason Blanchard, who coded much of the Syncphony application, observed that "[Syncphony] actually went through quite a bit of change." He explained that "the usability requirements were really originally driven from the engineer's perspective … but once we started bringing you to the table, it became more collaborative. In the sense that we tried to move ourselves along in a way that made sense for the artists."

The first iteration of Syncphony received gestures from a single wristband over Bluetooth Low Energy and converted it into MIDI output to play sounds in a digital audio workstation. Then, as Jason explained,

> We moved on to using 900 MHz radio communication as well as the ultra-wideband radio communications—that was one big change. Another was actually in response to your requests, which was to be able to output OSC rather than just MIDI. We also added

> the ability for multiple gesture devices to correspond to the same instrument so there was redundancy handling. (personal communication, March 11, 2016).

Open Sound Control facilitated communication between software; 900 MHz and ultra-wideband increased range and lowered latency.

Latency was a major concern throughout the development phase—"we ended up doing a lot of things to try to limit latency," Jason said. There were improvements in three areas: wireless networking, software streamlining, and gesture recognition. Switching the wireless technology from Bluetooth to 900 MHz and ultra-wideband considerably reduced latency. "We worked a lot on [Syncphony] on trying to streamline the processing … that being redundancy removal and the mapping of gestures to the proper behaviors," Jason said. Lastly, the gesture-recognition algorithm was tweaked to reduce computation and speed up detection.

Two issues arose during the development phase: the network's complexity and Syncphony's stability. Regarding the first, Jason explained:

> We would often have a hard time determining where things went wrong within the system. For example: was it the Syncphony app that was not sending out the sounds properly, or was it further down the line, like the ultra-wideband not working? Those kinds of things were always an issue. (personal communication, March 11, 2016).

Monitoring scripts were inserted along the network path. but troubleshooting still proved time-consuming. As the *Curie* show evolved, new features were added to Syncphony that compromised its stability. This was a big problem, "because when the

Syncphony app crashes, essentially the show's dead, because there is nothing coming through. Making sure that we cleaned up the app so that it didn't crash was a big issue that we had along the way," Jason said.

Due to time constraints, two features were not completed: Syncphony's interface for changing performer configurations and a wireless method to update the Curie Module's firmware. The performer configuration mapped a set of interactive gestures to a performer for them to play the appropriate instrument. Jason explains that they "did not have enough time to make [the interface] as easy as we would have liked." The firmware on the Curie Module configured parameters of the gesture algorithms. The plan was to be able to quickly change the firmware by updating it wirelessly. As Jason recalled:

> Unfortunately, we did not have time for that. That was actually kind of a big thing. It would have made it easier on everyone on making sure that the bands were acting properly. Instead we ended up having to flash new firmware onto the bands when anything needed to change." (personal communication, March 11, 2016).

In the end, flashing the devices required taking them off and manually connecting them to a computer.

In the end, "I think we were pretty happy," Jason remarked. "It ended up being fairly stable and it got the job done. But there are definitely improvements [to be made]." He hopes to revise the user interface and re-code it in C or C++ to increase speed and compatibility.

*The State Machine***.** The State Machine software, developed by me and shown in *Figure 3.19*, mapped gesture and position data from wristbands and RealSense cameras to musical parameters. Each mapping configuration was saved in a state module. A collection of states for a song was called a project. The user could activate a state by clicking on the interface or sending a MIDI note corresponding to the index of a state; only one state could be enabled at any time.

As shown in *Figure 3.20*, The State Machine accepted OSC and MIDI information and routed it to a state containing a particular mapping instruction. These instructions converted input data into control information that was sent to Ableton Live for audio effects and TouchDesigner for visual effects. The State Machine's modularity provided a host of benefits:

- an easy platform to create and prototype mappings,
- no mapping conflicts between states,
- mappings could be copied/pasted within or across projects,
- projects were ready to go when launched,
- worked well with a digital audio workstation's timeline, and
- it had a clean interface.

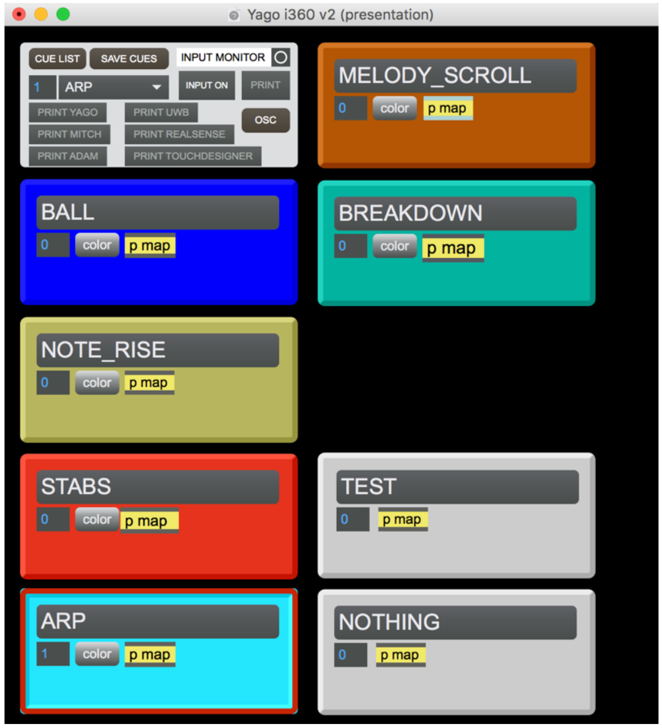

Figure 3.19: The State Machine.

Each colored box was a state containing a unique mapping configuration.

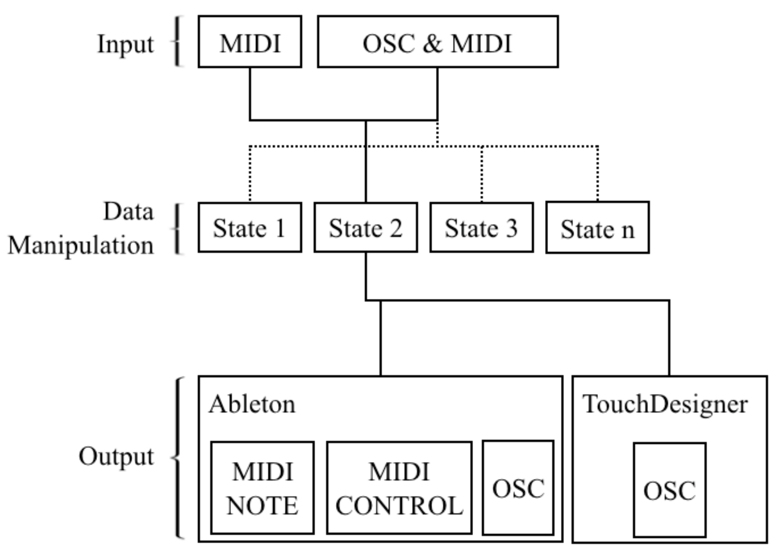

Figure 3.20: Data flow in The State Machine.

Figure 3.21 exposes the mapping configuration inside the STABS state. All OSC and MIDI messages were gated so they only passed data if the state was active (active and inactive states were commanded by 1/0 message to the gate objects). In this example, if we imagine the data cascading downwards after the top gate object, the first thing we do is define what kind of data will be accepted with the route object. The STABS state accepts only right-hand and left-hand side hits. If the acceleration of the right- or left-hand hit is higher than 30, then a 4 (right hand) or 1 (left hand) is sent in MIDI to Ableton. Ableton remaps the 4 and 1 MIDI triggers into chords. These chords are fed back into the STABS state through the “notein” object on the lower left and sent to TouchDesigner.

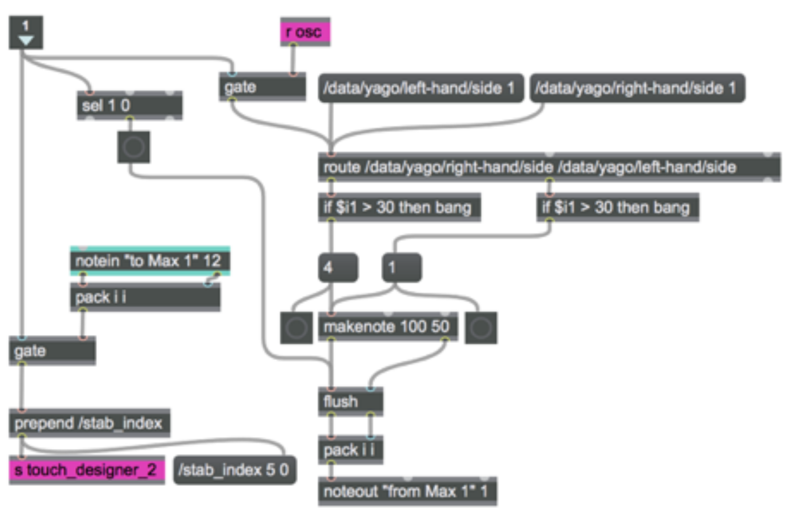

Figure 3.21: Mapping configuration inside the “STABS” state connecting right- and left-hand side hits to MIDI notes.

The Portal. The Portal is a Max/MSP program created by Chris Legaspi that parses data from Syncphony and distributes it to various computers. Its purpose was to route messages from Syncphony, RealSense cameras, and PixMob to The State Machine and Peter's multiple computers. As illustrated in *Figure 3.22,* the interface was designed to be user-friendly: the operator could quickly change the IPs, ports, and presets as well as monitor and test OSC messages.

One of the problems we came across was that Max/MSP could not unbundle multicast messages. Multicasting efficiently distributes OSC messages to different computers across a network. Our solution was to create a terminal application that converted multicast to unicast and sent individual OSC messages to each IP, which increased the load on the network. Furthermore, the various software applications running in the main computer conflicted over the allocation of UDP ports. Lastly, The Portal did not always load a preset correctly. Nevertheless, we were able to circumvent these problems for the show.

Jean-Sebastien criticized The Portal when he said that, with regard to data handling, "the main failure point was [The Portal]…. It should have been done in something else. I felt like it was slowing down when there was too much stuff going on."

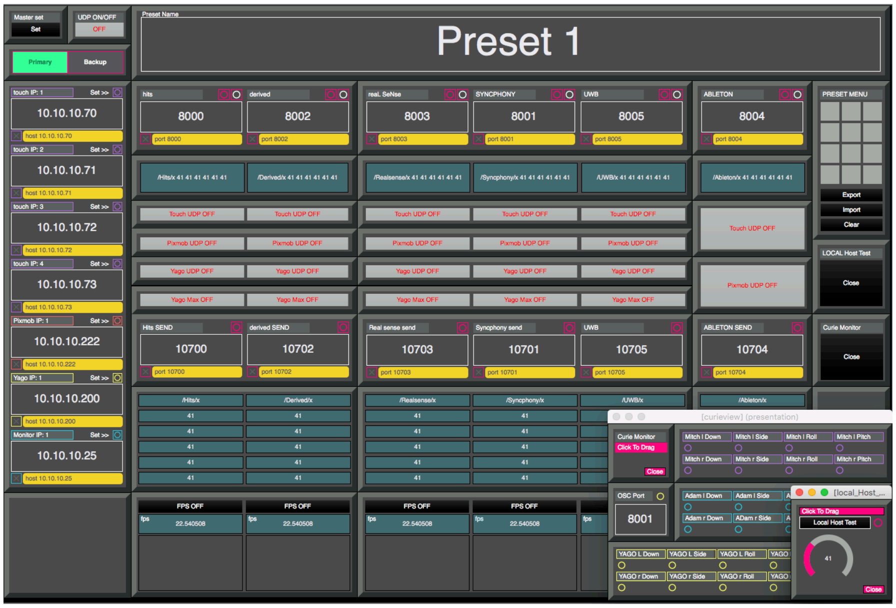

Figure 3.22: The Portal monitors and distributes OSC messages from the sensors to audio and visual software.

Mappings

All the visuals were manipulated on the scrim in front of the performers. Most of the audiovisual elements were controlled by me throughout the piece. Musicians and audience members only had interactive control during the breakdown sections. The dancers who appeared at the third verse only controlled visuals, not sound. Mappings between the gesture and position data and the audiovisual content were coded on The State Machine or TouchDesigner. Depending on the section, the State Machine and TouchDesigner would automatically enable or disable a set of predetermined mappings; while audio tracks and visuals were muted or un-muted. All the states were cued in the timeline of the music software.

The gesture-to-multimedia mapping went through four major revisions, named here after where they happened: Los Angeles, Santa Clara, Evansville, and Las Vegas. This section will detail how music, image, and technology co-evolved to create the mappings. The song's instrumentation introduced the majority of the mappings. Towards the end, the need to simplify the interactive system required many of them to be dropped.

Los Angeles. I presented a demonstration of an interactive audiovisual song in Los Angeles. The visuals were pulled from a previous project and the music was selected from Zackery's list of compositions. His song was picked apart, three tracks were extracted and made interactive. The technology, none of which was used in the final show, consisted of Myo's armband inertial sensor and Microsoft's Kinect for skeleton tracking. The song ran for 1:30, featured three audiovisual mappings, an upper body visualization, and plans to incorporate a *Guitar Hero*-like instrument.

Santa Clara. When I was asked to present a demonstration at Intel's headquarters in California, I received the first version of the gesture capture wristbands. No changes were done to the song structure, but the wristbands substituted for the Myo sensor.

Evansville. Zackery finished composing the song for the show at the beginning of this phase. Around seven mappings were created to accommodate the new instruments. Three out of four of the mappings from the demo song were re-applied to similar sounds in the new composition. The song's instrumentation was still the major force behind the creation of the new mappings, with the exception of Mitch's virtual instrument, which was added to the composition.

Las Vegas. More mappings were deleted than added in this phase. Those that were added were not based on the music composition; rather, they connected Peter's visuals to the performer or the PixMob audience participation to audiovisual effects. I'm not entirely sure how the choreography and visualization evolved together, but I believe Peter already had ideas in adding particles or lines to the limbs of the performers. The deleted mappings were most likely a result of aesthetic and practical decisions. The RealSense and PixMob technology required the composition to adapt to more instruments and modify the breakdown section.

List of Mappings. In the following page is a list of mappings for the performers and their reactive audiovisual elements, organized chronologically. The middle column indicates who performed what gesture, while the left and right columns describe the accompanying visualization and sound, respectively. (~) indicates that the sensor's signal consisted of a range of values, (bool) was an on/off operation, and (trig) a single event trigger.

Visualization	*Performer* Gesture	Sound
	Intro	
	Yago	
White stripes opacity	Right hand vertical angle (~)	Noise volume
Blue circle diameter	Right hand clockwise roll (~)	Pad volume
Left green meter	Left hand vertical angle (~)	Synth Pitch
Right green meter	Right hand vertical angle (~)	
	Pre-Verse 1	
	Yago	
Panels blink and spectrum level meter	Right or left hand side hit (trig)	Chords
	Verse 1	
	Yago	
Circle over each hand	Hands horizontal and vertical position (~)	
	Right hand vertical angle (~)	Low-pass filter on synth
Hand display rotation	Right hand clockwise roll (~)	Arpeggiator speed
	Chorus 1	
	Yago	
Color boxes under scrolling notes	Horizontal hand position (bool)	Un-/mute melody track
	Pre-Verse 2	
	Yago	
Panels blink	Right or left hand side hit (trig)	Chords
	Verse 2	
	Yago	
Blue trail	Horizontal and vertical hand position (~)	None
	Breakdown 1/2	
	Audience	
LED lights on wristbands and particles on the scrim	Quantity of motion of five different sections (~)	Volume of five different frequency bands of pad
	Yago	
Selection of five vertical columns	Hands horizontal and vertical position (~)	

Table 3.3: Mappings used in *Curie* listed
by order of appearance in the song.

	Breakdown 2/2	
	Yago	
Circle over each hands	Hands horizontal and vertical position (~)	
Green meter	Right hand vertical angle (~)	Low-pass filter on synth
Hand display rotation	Right hand rotation (~)	Arpeggiator speed
While stripes opacity	Right hand vertical angle (~)	Noise volume
Blue circle diameter	Right hand roll (~)	Pad volume
	Mitch	
	Right hand down hit (trig)	Play a note
	Right hand vertical angle (~)	Synth release length
	Left hand vertical angle (~)	Synth volume
Dial	Left hand roll (~)	Synth pitch
	RealSense right hand depth position (~)	Filter morphing
	RealSense left hand horizontal position (~)	Reverb volume
	Adam	
	Right hand down hit (trig)	Crash snare sound
	Left hand down hit (trig)	Kick sound
	Right hand side hit (trig)	High hat sound
	Left hand side hit (trig)	Wood snare sound
	Left hand vertical angle (~)	Kick release length
	Right hand vertical angle (~)	Crash snare release length
Dial	Left hand rotation (~)	Saturator & kick frequency
Dial	Right hand rotation (~)	Crash snare low-pass cut
	RealSense right hand depth position (~)	Crash snare tremolo effect
	RealSense left hand horizontal position (~)	Reverb and delay
	Verse 3	
	Yago	
Circles over hands	Hands horizontal and vertical position (~)	
	Right hand vertical angle (~)	Low-pass filter on synth
Hand display rotation	Right hand clockwise roll (~)	Synth arpeggiator speed
	Dancers	
Origin of particle generators	Right and left hand position (~)	
	Chorus 2	
	Yago	
Color boxes under scrolling notes	Horizontal hand position under scrolling notes (bool)	Un-/mute melody track

Table 3.3 cont.

Chorus 3		
	Yago	
Blue trail	Horizontal and vertical hand position (~)	None
Coda		
	Dance Jockey	
Panels blink	Right or left hand side hit (trig)	Chords
Opacity of blue mist	Right hand vertical angle (~)	Low-pass cut of noise

(~) continuous signal (trig) discrete event (bool) true or false

Table 3.3 cont.

Mapping Development

Intro. The musicians moved into place while the stage was dark. I started the piece with my hands down and slowly rose them to increase the volume of a noisy sample and the opacity of a series of white stripes. Then my hands rotated and pushed out to increase the size of the blue circle. Raising my arms increased the volume of a pad. After about a dozen seconds, I pointed left and right to introduce the drummer and the bassist with a short solo, while visual structures around the performers were assembled and the lighting designer lit the blue stage lights. After the musicians were introduced, I lifted my hands to increase the frequency of a square waveform sound.

The primary objective of the *intro* was to introduce the interactive concept to the audience by making explicit the connection between the gestures, visuals, and sound. Controlling both the volume of the noise and pad with the right hand made it hard to control each of the sounds independently. I could have switched one of the sounds to the other hand, but there wasn't enough time to make this change and rehearse it. My original plan was to increase the volume of the pad by expanding my arms because that would have been visually consistent with the interactive blue circle that grew in size. But the positional data was not reliable enough for that approach.

Pre-Verse 1. Three buttons were arranged radially on my left and right side. By extending my arms outward, I could press the button and play a chord. Only one button on both sides was responsive at any time. The responsive button changed at particular beats in accordance to the harmony.

Early in development, we tried to make hand rotation or hand height select the button. but it was difficult to play—my arms were moving too fast to be detected. So although I seemed to be hitting different buttons in the video, their selection was actually predetermined. I accomplished this by automating the small device that pitch-shifted the incoming trigger MIDI. The right hand always triggered 1 and left hand 4. The pitch shift would then change the pitch to somewhere between 1 through 6 to play one of six chord samples in a sampler.

The hit detection algorithm produced false positives, and changing the code in the Curie chip was a hassle. Changes to the threshold of the hit gesture had to be flashed to the firmware of the Curie chip. We could only do this by changing the code, connecting the Curie Module through USB, and uploading the new code. Early in the project, we had the ability to wirelessly change some of the settings of the gesture algorithm with Syncphony and Bluetooth. Luckily, the hit values from the wristband indicated the degree of acceleration on a scale from 0-100, so I was able to filter out low values in The State Machine, since most accidental hits had a low acceleration.

Verse 1. The buttons disappeared to give way to dials over my hands that controlled a couple of parameters in an arpeggiator. Rotating my right hand clockwise increased the number of notes being played. Pointing my right hand down or up opened or closed a low pass filter.

I tried to get as close as possible to the scrim when playing the arpeggiator instrument to better superimpose the projected circle over my hands and reduce the parallax for those watching from the sides. However, moving the wristbands away from the ultra-wideband platform's center decreased tracking precision and increased jitter. Due to this, my hand movements had to be slow, and I had to stand one foot away from the scrim.

Controlling both interactive parameters with my right hand was not very effective; rotating the hand to change the arpeggiator's speed was particularly hard. Like in the *intro*, I had planned to control the filter by measuring the distance between my hands, but the positional data was not reliable enough.

Chorus 1. Melodic notes were arranged horizontally like a piano and scrolled from top to bottom on the scrim in front of me. I had to place either of my hands exactly where the scrolling notes landed to keep the melody track on. I had to adjust the range that defined what was over and what was not for each show.

At the Anaheim Convention Center show, the notes were arranged radially in 360 degrees on the floor. This show required increasing the range because the notes were coming from 360 degrees and were extremely hard to hit. Nevertheless, I was hesitant in facilitating this section because it could become less engaging. In practice, it was hard to know why some notes were not being played correctly, so learning was a result of trial and error during rehearsals.

It was impossible to practice playing the melody without the ultra-wideband positioning system and the projections. I could memorize roughly the position of the notes as they scrolled towards me without the projections, but the interactive system required the spatial location of my hands to check if they were correctly positioned. The ultra-wideband system was not very convenient, since it required a lot of hardware and at least one programmer to start, manage, and debug the software. This limited the extent to which I could experiment and rehearse with the mappings.

I wanted to add a layer of interaction that would require me to perform a down hit in addition to placing my hands over the note. I believed that that gesture would make the section more engaging and legible to the audience. However, I was not able to include the down hit because it was disabled to reduce network traffic, and the person that could re-enable it was unavailable. The lack of flexibility in changing the Curie Module's gesture capabilities limited the mapping possibilities. Nevertheless, I gestured a down hit over some of the notes to impart to the audience that I was playing them.

Pre-Verse 2. Same as *pre-verse 1*.

Verse 2. Blue trails followed the vertical and horizontal position of both my hands as if I was painting on a canvas. The tail of the streak slowly faded away after some time. My movements followed a simple choreography. There was no sound interaction.

During this section's development, I experimented with two different techniques for controlling both the melody and visuals, but ended up with only the latter because of technical difficulties and time constraints. The first approach using IRCAM's Gesture Follower gave my hands the ability to trace a melody. The second approach compared the current gesture data with the stored, time-based gesture data. Both were matching algorithms that compared a previously recorded hand choreography with what I was performing in real time.

The ultra-wideband positional data was inappropriate as an input to both algorithms because the data was too jittery and I had limited practice time on the ultra-wideband platforms. Fortunately, the Curie Module's inertial data served as an alternative.

IRCAM's Gesture Follower is a powerful tool that records time-stamped gesture data that can be recreated by a performer in real-time. In *Curie*, the timeline of the recorded gestures held MIDI notes that played a synthesizer. Gesture Follower worked independently of the DAW's timeline, meaning that rhythmical conflicts would arise if the gestures were not being followed correctly. I decided not to use Gesture Follower because it did not always sync up with the very precise rhythms of our song.

I created a similar and simpler program called Dance Match that measured the degree of similarity between real-time gesture data and pre-recorded data that was time-locked to the Digital Audio Workstation's timeline. It worked by converting the gesture data into MIDI control data and saving it on a MIDI track. As the song went through this section, the recorded gesture MIDI control data was sent back to The State Machine and compared to my real-time gesture data to produce a correlation value. Depending on whether I was performing the choreography correctly or not, the correlation value raised or lowered the melody track's volume. What was particularly pragmatic about this approach is that new data could be recorded instantly if the choreography was modified.

Dance Match was not implemented in the final show for three reasons. First, I did not have enough rehearsal time to learn a new choreography for this section. Secondly, the mapping landscape for the entire piece was already getting too complex. Lastly, because of the former reasons, I did not have a chance to work with Peter on how to express this matching system visually. That led me to decide that if the audience was not going to see the relationship between my dance and the melody track, then it was not worth including it. Ultimately, I dropped the musical interaction in favor of just having visual manipulation.

Breakdown 1/2. During the first half, the breakdown audience members were invited to contribute musically and visually to the performance by performing a fist pump. The crowd was divided into five columns that were spread evenly across the scrim. Sonically, each column represented a band pass filter applied to a synthesizer. My arms could highlight a column on the scrim and I would ask the crowd to help me activate that column. When a column of audience members performed a fist pump three things happened: 1) Their wristbands lit up; 2) their respective column on the scrim exploded with particles; and 3) the band pass of that column opened up. The audio rig operation looped this section until I introduced the musicians for the following section.

Intel's investment in the wristbands bolstered their resolve to include audience participation despite the creative team's initial skepticism. Intel had produced 3,000 wristbands with onboard gesture algorithms and LEDs. They were also working together with PixMob to develop a visual feedback system using overhanging cameras and the LEDs. One of the reasons the creative team was concerned with the outcomes of audience engagement was that the system was never fully functional until showtime. We did not have an audience to test on, and Intel installed the feedback camera system only on the main stage, where we had very little rehearsal time. The other challenge was a creative one: the visual artist and the composer had to design an interaction for a hypothetical crowd.

The engineers at Intel were excited by the idea of using multiple gestures and dividing the audience into different groups, similarly to an orchestra. However, the final version allowed only very simple control of light, visuals, and sound. How to engage the audience was debated throughout the development process. It was inconclusive for the most part because the technology was not mature and it lacked user testing. The final solution was driven by four technical and creative constraints:

- The vision system was able to detect if the wristband was on or off
- The easiest gesture to capture was a fist pump
- Audience interaction was going to be unpredictable

- There was no way to fully test the audience interaction system
- The audience interaction section had to be around 45 seconds long

These five constraints suggested a simple interaction model: By shaking her hand, an audience member would light up the wristbands' LEDs, create particles on the scrim before her, and increase the volume of a synthesizer. Although we did not use it, in case the audience did not shake their hands, PixMob could signal the wristbands to light up regardless.

Breakdown 2/2. In the second part of the breakdown, the bassist and drummer left their acoustic instruments and grabbed a pair of interactive wristbands. The drummer played an electronic drum kit while the bassist played melodic notes with a long sustain. Joined by my arpeggiator, this section built to a crescendo and then contracted for a few seconds while I played a noisy pad and the musicians got back to their instruments.

Aside from the wristbands, Adam and Mitchell also could interact with the RealSense cameras. Adam's right hand moved forward and backwards to manipulate a tremolo effect, while his left hand moved out sideways to add a reverb and delay effect. Mitchell's right hand moved forward and backward to control a filter while his left hand moved out sideways to add a reverb effect.

I designed a mapping scheme that made it easier for Mitchell to playfully hit a few pitch classes over many octaves rather than carefully select every note and risk hitting wrong ones. His instrument could only play a set of harmonically correct notes that updated at every chord change. Rotating his hand produced no useful visual or physical feedback, so knowing what notes were selected before they were triggered was impossible. Initially, Mitchell's right-hand vertical angle was inversely mapped to note duration; hand up produced short notes, while hand down produced long notes. During a rehearsal, I accidentally inverted the mapping, resulting in much longer notes that fit the song better. During the show, Mitchell chose not to use the RealSense camera.

Adam played a percussive instrument and manipulated sound effects. During rehearsals, we explored a number of samples to find the best ones for his four hit gestures. Adam would occasionally miss a hit gesture, either because the gesture was performed incorrectly or the network lost a package. These two problems were eliminated after Adam trained with Swarnendu (developer of the gesture algorithm) and we performed in a better network on the main stage. Adam mostly used the right RealSense camera (depth axis) because it produced a noticeable sound effect. Latency was low enough that he could play percussion with precision.

Verse 3. The second half of the show introduced two dancers whose hands generated purple and yellow smoky particles on the scrim. The horizontal and vertical position of their hands determined the location of a particle generator; they did not produce any sounds. My hands controlled the same arpeggiator as in *verse 1*.

Chorus 2. The second chorus was very similar to the first, except that the last passage of the melody contained more notes.

Chorus 3. The interactive blue trails from *verse 2* were used here again. My movements were improvised. The dancers continued to generate particles with their hands.

Coda. The coda reintroduced the buttons from *pre-verse 1* with a different set of chords. The dancers stopped generating particles. The section concluded with me rising and lowering my hands to control the opacity of the white smoky particles covering the scrim.

Visuals

Peter Sistrom and his team Name the Machine worked directly with me and Intel to address the artistic and technical requirements for the show. They designed the visuals, managed the interactive data and built the computer rigs.

The role of the visuals was to a large extent functional in that they presented information about the sensor status to the performers. Peter intended the visuals to look like a heads-up display—a transparent interface that presents data within the user's field of vision—acting as a "visual analog to the instrument we were creating."

Adam, Mitchell, and I were encased in a heads-up display represented as a "frame" or "architectural structure" that provided a platform for the interactive portions. On the other hand, the dancers' visuals were more about "aesthetic and reactivity." They did not require an interface; rather their movements were amplified.

Peter said the overall theme had a very “techy” vibe: “One thing that these real-time systems often need is simplicity in their aesthetic because of the performance needs. So a lot of times, that actually informs the aesthetic.” The scale of the architectural structures surrounding the performers was a function of performer and stage size.

Peter referred to the process of creating the interactive visuals as sculpting:

> The act of setting up these data pipelines and all these logic switches [by] massaging these pathways and changing them and rerouting them to make it all come alive when it is inputted with real-time interactivity, it was always a living entity … all the way through the performance. (personal communication, March 30, 2016).

He admits that this workflow is taxing on the system because all the content is generated in real time. He keeps an eye on the graphics card performance meter so that the frame rate stays at either 30 or 60 Hz, otherwise it will lag behind the performers. The process of developing mappings is comprised of “tinkering," creating robust connections between the software and hardware; and “practicing,” exploring and rehearsing said connections to find the “cool things you can do with it.” Peter admits that he sometimes had to fight the urge to continuously tweak the mappings.

Peter describes a successful mapping as being immediate, nuanced, and evocative. First and foremost, the interface designer has to ensure that sound and visuals respond instantaneously to the actions of the performer. Reducing latency ensures that the media can carry out a proper functional role in the interaction. According to Peter, nuanced mappings provide a rich variety of possible outcomes from a limited amount of inputs:

> We take most advantage out of the input technologies we are given and allow a performer like yourself to play with it, to jam with it, to find the spaces that this mapping allows…. Maybe the most successful mappings is the one that strengthens or informs the next mapping, whether it's an evolution of itself or the setting up of a new connection. (personal communication, March 30, 2016).

This process is a feedback loop where old mappings are re-examined in light of new discoveries. Finally, he said that the mapping has to "produce evocative audio and visual results."

Name The Machine used the TouchDesigner software to generate the visuals because it was quick to program, versatile and effective for real-time rendering. Peter, the visual designer, described the software's interface as "high-level," "modular," and "open ended." Furthermore, it "allows me and other people to leverage components that have already been made by the community of TouchDesigner users." And it accepts "a significant amount of inputs from other protocols that other artists and creatives use to transmit data." When designing the interactive software, it is necessary to keep the programming environment simple and clean; a messy environment can slow down the prototyping phase and obscure the code.

Choreography

Reema Bounajem was the choreographer for the show; she and Shoko Fujita were the dancers. Reema explains that she started the project by choreographing and sharing some ideas before having the visuals and music. She explained:

> We had two weeks to record ourselves doing some rough material to submit to the rest of the team just to let them see this is the kind of movement I'm doing, how it will work with the visuals…. At this point we didn't know what any of the visuals would look like. We didn't even know what the music was going to be, it was sort of a shot in the dark of how we imagined how the show could look like. (personal communication, March 18, 2016).

Before having the technology ready to produce the visuals, she had to "essentially imagine if I had attachments to my arm, how would I draw in the air." Reema took into consideration four major elements:

- What kind of movement, how to share the stage with the other performers,
- What kind of music, when to perform during the piece, how to develop the form;
- What kind of visuals, when were they interactive, how to control them appropriately; and
- What parts of the bodies were being tracked, how would they constrain movement.

The music and visuals called for conflicting choreographic styles. Reema explained:

> Once I heard the music there was a little bit of a conflict in terms of aesthetics…. Because the music was very quick, and techno, and sharp; you couldn't have counts to the music, it was very subtle and nuanced…. Whereas with the visuals you had to make it flowy: if you are fast, the visuals can't keep up with you. (personal communication, March 18, 2016).

Furthermore, the short length of dance meant that the choreography "didn't have time for an introduction, rising, breakdown, climax, and then the cool down into an ending." The final choreography struck a balance between aesthetics, energy, and enhancing the visuals.

Sensor position increased the expressivity of the limbs and limited the gestures. Reema recalls thinking, "Is the movement going to be impaired by the placement of the sensors? how are the sensors going to work with the visuals? There was a lot to take into consideration."

The sensors on the wrists and legs inspired her to "choreograph things with long kicks and larger leg movement to utilize the sensors on our ankles" and put "emphasis on the upper body." Meanwhile, the sensors also limited the movement: they "were these little bulky rectangles, so we had to be careful of their placement on our body. We did not want to roll on them and hurt ourselves, like on the outsides of our ankles or on our backs."

Reducing the number sensors on the dancer's body was an unexpected change to which the dancers did not have enough time to adapt. Intel removed the ankle and lumbar sensors and left only the wrist sensors because the network was getting saturated. When the network overloaded, the sensor's position would get stuck. The leg movement became "no longer necessary for the sensors; it was just an aesthetic attachment."

Reema was excited by the idea of making visuals in real time and intended to make it clear to the audience how the dancers were controlling the visuals:

> I didn't want to make it difficult for the audience to see the interaction. I wanted to have moments where it was pretty obvious. I was always thinking what would make the dancing blend the best with the sensors, like utilize it the most, so that we would make it really obvious that the dancers aren't just moving, they are actually part of the visuals. They actually have sensors in their clothes that create movement on the scrim as well. (personal communication, March 18, 2016).

The choreography was co-dependent on the visual particles; they relied on each other for the performance.

Reema regarded the visuals on the scrim at CES as more engaging and legible when compared to the Anaheim Convention Center's rear LED panels:

> When we walked right on to the stage at CES, we could see the visuals, and we could really have fun playing with it; we could see what we were affecting. That affected our choices of making the movement as big as possible, as sharp, to draw as much in the air. At [Anaheim], we couldn't see it. So we just had to dance very specifically in the way we had imagined it. (personal communication, March 18, 2016).

Seeing the visuals helped reconnect the dancers with the visuals: "There was one moment in [Anaheim] where I even included in the dance that we would turn around to glance back to face the visuals to at least establish some sort of connection," Reema said.

Reema worked closely with the engineering team on the development of the sensors. She also recorded various motion-tracking clips for the visual artist to design some interaction mockups. Together, they explored artistic possibilities, investigated problems, and discovered some limitations like the low precision of horizontal and diagonal movements. “It was an ongoing process,” she said, “making something, listening, gaining new information, and thinking about how to move forward, what can I bring with me and what is no longer necessary for the choreography.”

Although Reema was satisfied with the final choreography, she wished she had more time to adapt to two main changes: the shortening of the song and the deactivation of the interactive particles during the *coda* section. She remembered:

> It was getting so close to the show that I wouldn’t want to change something when I know that everyone else is trying to finalize their own thing. I just wanted to keep it as stable as possible. But if we had more time … I would have loved to know when visuals were no longer being connected to dance, so that I would have more freedom to choreograph instead of thinking the entire time [about] how to create shapes and draw. (personal communication, March 18, 2016).

Music

The music was composed by Dr. Zackery Wilson and was inspired by funk and video game music. The song’s instrumentation was upbeat, fast, and intricate; the synthesizers used waveforms (i.e. sound fonts) based on the *Final Fantasy* video game series. The final version of Zackery’s composition featured a fixed soundtrack, live drums and bass, and interactive instruments. The duration and compositional structure of the live show adhered to a specified timeline. The musicians and I were cued to the sections using a click track.

During our first collaborative meeting, Zackery and I broke down one of his songs to make it interactive; the process set a precedent on how we would work together in creating the final track. Early on, while Zackery was still composing the song, I used a placeholder paired with my own visuals to demonstrate the interactive concept to Bill Welter, the show's executive producer. The placeholder song—"Suck 'R Punch [Kirby's Dream Land]" by Zackery—was selected after listening to many of his previous compositions. It offered a variety of sounds, both percussive and melodic, that could be played in a number of ways. We experimented with an arpeggiator, chord stabs, and generating a melody. The visuals were pulled from my previous project *Be Real*.

When Zackery finished the first draft of the final song, we listened to the individual tracks and I identified potential interactive instruments. I had to reconstruct the interactive tracks in Ableton Live because Zackery composed on Fruity Loops. Setting up the interactive tracks consisted of:

- Loading samplers with chord or percussive samples,
- Loading wave samples into synthesizers, and
- Placing audio clips and creating sound effects on tracks.

The State Machine sent all the MIDI and OSC commands to the interactive tracks. Zackery created most of the audio clips; we collaborated on the sound effects. Aside from the interactive and backing tracks, we had utility tracks that served various functions.

Fred Carlton, the audio rig operator, recommended not consolidating all the backing tracks so that each sound could be mixed individually after hearing them in the venue. The functional tracks included:

- A full song track for reference during rehearsals,
- Cue tracks to synchronize The State Machine and TouchDesigner,
- Safety tracks for the interactive tracks in case of sensor failure, and
- A count track to cue the musicians.

During the 2015 Christmas break, before the last set of rehearsals, producer Bill Welter had Zackery trim the song. We shortened the show from 5:30 to 4:30 to accommodate for the rest of the keynote material. I was going to sing in the verses and choruses but Intel decided against that since they wanted to focus on the interactive instruments. These unexpected and late changes forced the artistic team to adjust their works in a rushed manner. Except for the song length, its structure did not change much during the development phase.

ANALYSIS

Reports from the participants indicated that concerns were addressed in a satisfactory manner, but the approaches to coping with artistic and technical issues varied. In this section, we will see how the engineers attempted to resolve technological limitations for the performance team, the consequences of the ones left unresolved, as well as the conditions that cause them to be unresolved. In the transparency section, most of the creative team was united under a common set of concerns related to how they wanted the interactive instruments to be perceived for themselves and by the audience. The last section identifies crucial themes that emerged from the collaborative efforts.

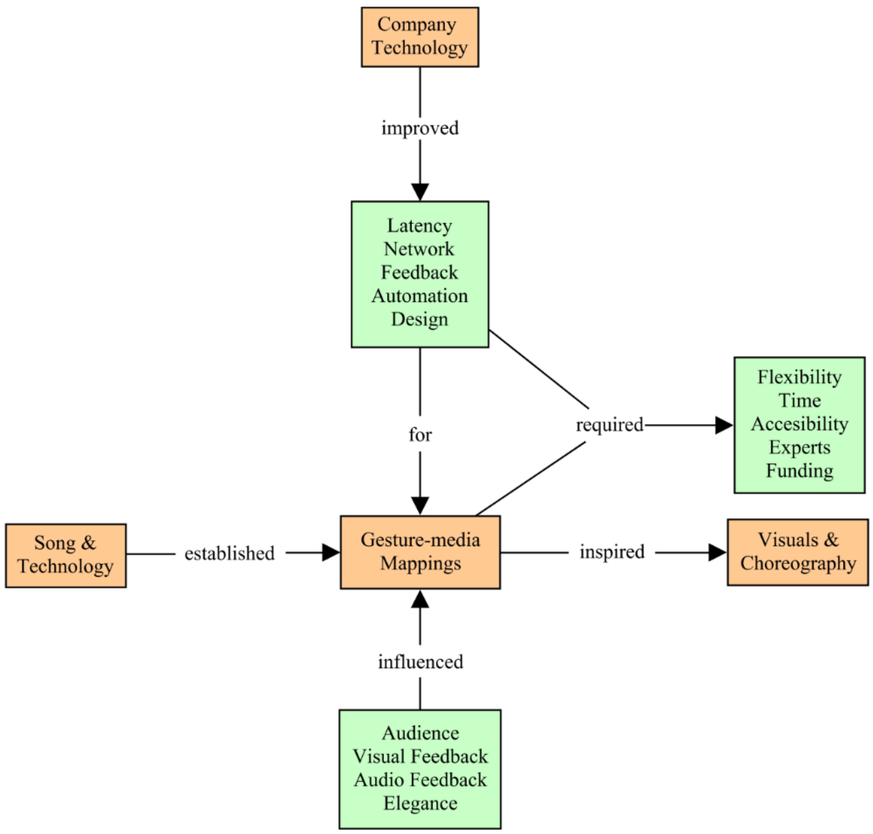

Figure 3.23: Technological, performative, and collaborative areas that influenced or were influenced by the gesture-media mappings.

Figure 3.23 illustrates the relationship between the various development areas that influenced or were influenced by the gesture-media mapping. The orange boxes depict information given in the preceding sections, while in green we see the technological, performative, and collaborative elements that will be analyzed here. The final mapping scheme for *Curie* was mostly a product of aesthetic considerations that dictated the development path of the technology. The magnitude of these developments and the geographic constraints presented unique aspects pertaining to the management of resources.

Technology

The wide range of technologies developed for *Curie* provided useful information that led to insightful discoveries with respect to common problems across all the engineering challenges.

Latency is a problem inherent in all electronic devices due to the time it takes to transport and process signals. Different approaches were taken to reduce it and coping mechanisms were devised by the performers. Network management became more and more important as the complexity of the project grew. System operation feedback in some ways was not taken seriously enough, but its significance became more evident as the digital network expanded. Automation reduced the need to oversee and manage the technologies and it also streamlined the boot-up process. Ensuring that the technology had an attractive design enhanced its aesthetic value.

Latency. The impact of latency was felt in every element of the production and was a major area of focus. Latency originated from the sensors, network, and motion-capture algorithms. The most successful solutions—like the improvements done to the musicians' wristbands—had to address all of these elements. When latency was sufficiently low, it supported the execution of fast gestures; otherwise, they had to be slowed down. Yet the perception of a delayed reaction depended on the task. Trigger gestures required faster systems than those responsible for continuous control. Reema, for example, said that "the visuals were not fast and snappy," but she was comfortable with them because fast mappings were "not the goal." Adam, on the other hand, stated in his interview that successful mappings should have low latency. He felt it was crucial for the virtual instrument to respond quickly like an acoustic instrument.

> You have to go through the technical hemming and hawing to make sure that the sync is as good as possible before you go anywhere. - Peter Sistrom

Most of the technology was optimized to process and transmit the sensors' information as quickly as possible. Jason Blanchard of Intel shared a recurrent theme across the development teams: "We ended up doing a lot of things that tried to limit latency."

The wristband's network, processing, and gesture algorithm were improved to deliver the time-sensitive gesture hits with the most amount of consistency and least amount of delay. Running the gesture algorithm locally on the Curie Module was more efficient than on the computer, and it also reduced the amount of data being transmitted wirelessly. Moving the wristbands from the Bluetooth protocol to the dual-radio solution provided a major reduction in wireless transition time. Jason Blanchard streamlined the processes in Syncphony. Swarnendu sped up the hit-detection algorithm when he switched the accelerometer for the gyroscope. The transmission speed of the ultra-wideband system was increased when the internal network was converted to USB and Ethernet. Latency was not a priority for the RealSense because its tracking algorithm was adequate for the gestures we were performing.

Despite all the effort towards reducing latency, there were cases when it was still inadequate. The low precision of ultra-wideband required an averaging function that considerably slowed the tracking speed. This limitation forced the dancers and me to slow down our movements, or the visualizations—which were based on the ultra-wideband position data—would be constantly late in following our movements. The PixMob infrastructure, including the vision system and wristbands, had a considerable structural delay of 300 ms. The original gesture recognition algorithm designed by Intel added over a second to the process, because it required repetitive execution of a gesture to recognize it. Our faster solution was to switch the gesture algorithm for a simple acceleration trigger in whatever direction was indicated.

Network. Connecting all the technologies to the video and audio racks posed one of the major challenges in concretizing the show. Tasked with connecting the user to the audiovisual output, the network was the backbone of the interactive system. Its structure was comprised of transmission (optical, wireless, wired), consolidation (Syncphony, The Portal, ultra-wideband and vision servers) and communication (MIDI and OSC). The stage network had to accommodate 20 sensors, including gesture and positional data from five performers. The audience network managed the activity of 3,000 audience members. Gesture data from the audience's wristbands, the musicians' wristbands, and RealSense traveled over optical, radio, and Ethernet networks, respectively.

> The wristbands were working, everything was working. It was just a question of integration. - Jean-Sebastien Rousseau

The concerns facing the engineers revolved around lowering latency, increasing bandwidth, and upping reliability. The three wireless updates to the musician's wristband (from Bluetooth, 900 MHz, to ultra-wideband) produced marked improvements. 900 MHz and ultra-wideband allowed a much higher level of control than Bluetooth standards. Latency was reduced and bandwidth was increased by optimizing the wireless airtime scheduling service and streamlining the server's code. Some of the ultra-wideband system's internal components were moved over to a wired connection to further reduce latency.

Network bandwidth became a problem when we started to integrate all of the elements at Las Vegas. The ultra-wideband system's high sample rate occasionally swamped our video and audio rigs. It also had issues allocating enough airtime to all the data packages from the wristbands. These issues were solved by reducing the total number of sensors and gestures to only those necessary.

Redundancy ensured that the musician's hits were transmitted reliably to the audio and visual rig. Each musician wore two sensors that communicated via two different radio systems with multiple anchors and channels. Data was duplicated at the sensor level and captured by multiple receivers before being de-duped (the process of removing duplicates). These receivers were strategically placed around the stage to cover different lines of sight in case one of them suffered from occlusion.

Feedback. Adding a visual indicator to the sensors helped with debugging, provided feedback on how they were responding, and communicated the gesture to the audience. We would know that if a gesture was not going through but the wristband lit up, then the data was being lost over the network. Adding more visible watch points along the network path would have facilitated troubleshooting for dropped data packages.

> We would often have a hard time determining where things went wrong within the system. For example: was it the Syncphony app that was not sending out the sounds properly, or was it further down the line like the UWB not working? Those kinds of things were always an issue. - Jason Blanchard

The main purpose of the LED lights on the RealSense stand was to provide control feedback for the performers. An issue that contributed to failures in capturing some gestures was that the lights did not inform the user if his hand was moving out of the capture area. The ultra-wideband position tracking system did not provide any local feedback, so Peter's visuals had to take on that responsibility.

Automation. The motion capture and media systems had to run dozens of applications and scripts that loaded many media and configuration files. Setting up all these components manually would take too much time, so we tried to make them load automatically once the system booted. This saved time and labor and reduced errors. Automation consisted of making sure applications, configuration files, presets, and network settings were loaded correctly. The musician's wristbands, RealSense, and media sets achieved this to a large extent, while the radio networks, Syncphony, and The Portal required considerable supervision.

> Luckily [during the show] we didn't have to do too much. We set [the system] up in a way that everything that needed to be done on the fly was being done by a triggering system. - Peter Sistrom

The wristbands were programmed to automatically send data to the ultra-wideband and 900 MHz radios when turned on. Its firmware loaded a custom-made gesture algorithm to each performer. Intel's engineers also added a "watchdog" feature that reloaded the motion capture application if gestures were not being registered.

Both the ultra-wideband and 900 MHz systems required someone to load the scripts and oversee the quality of the tracking. Syncphony and The Portal needed someone to configure the software. Requiring someone with unique know-how to run these systems conditioned our access to the technology to their presence.

Design. Intel designed the wristbands and RealSense to look aesthetically pleasing onstage. The wristband's designs went through several iterations that took into account how the device was being used and how to make it look eye-catching. Intel did not consult with us when designing the RealSense acrylic cover, which resulted in a bulky and inconveniently large podium. Overall, Intel's consideration on how to present the technology contributed to the show's aesthetic.

Transparency

Transparency refers to how clear and evident the mapping is to the spectators or the performers. The performers and Intel were very interested in making the mapping legible for the audience, but for different reasons that had to do with the response they wanted to evoke. Peter, the visual artist, on the other hand thought excessive transparency was detrimental to the art form. Yet he agreed with the musicians that the visual interface had a functional role which demanded clarity. In this section, we will cover how functional audio feedback has its limitation due to the aesthetic role that music has in these kinds of shows. We will also see why aesthetic elegance was deemed to be a crucial element for increasing transparency.

Audience. Intel and the performers showed a big concern for communicating to the audience that they were controlling the audiovisual parameters with their gestures. Intel wanted to demonstrate to the audience what their devices were capable of, and that required a clear cause-and-effect relationship between the gestures and the audiovisual content. The performers had their own unique reasons for communicating the interaction to the audience.

> I think that if the technology is in the spotlight, and it is not obviously explained, then it can get confusing for the audience. - Mitchell Vandenburg

The dancers and I were interested in using expressive gestures that communicated to the audience that we were controlling the interactive digital content. Reema said that she was always thinking of how to connect the dancers with the interactive visuals. For that, she emphasized the upper body where the two sensors were placed. Most of my mappings were designed to be played using big and expressive gestures, so that the audience could see what I was doing to control the music and visuals. The blue trail and the chorus melody were examples of mappings that required big movements. Some instruments, particularly the chord stabs, could have been played more effectively if I used more conservative hand movement, but I preferred making them very visible and responsive to the visuals.

I also performed expressive non-interactive gestures together with interactive gestures for clarity. For example, interactive gestures, such as hand pitch and rotation, were normally accompanied by non-interactive gestures like rising or lowering and contracting or extending my arms. Although these movements integrated with the interactive media, they could not be tracked with the technology we had.

Both Adam and Mitch commented that communicating the causal relationship between gesture and sound to the audience could increase their appreciation of the performance. They didn't indicate that they wanted the audience to know "how" but just "what" was being controlled. "I think a basic understanding [for the audience] that the gestures are triggering the sounds would be important. Otherwise it could be perceived as something prerecorded as opposed to a live performance," Adam said. For Mitch, transparency is only relevant when the technology "is a focal point of the show. If it is background, then not as much."

I also shared their sentiment. Knowing what instruments the musicians were controlling added more meaning to our movements, particularly mine, since I used a lot of gestures not normally associated with generating sounds, such as raising my arm up and down or closing and opening my arms.

Clarifying the interaction to the audience helped Intel demonstrate their technology. Peter commented that corporate shows tend to "imply a more straightforward connection." Intel wanted to impress the audience and promote the RealSense camera and Curie Module. We achieved this not only by making the gesture-media interaction clear, but also by verbally informing the audience how we were manipulating the music and visuals.

Peter was the only member of the team to whom communicating to the audience that the performance was interactive was "not that big of a concern. There is always this sort of tech demo interest by certain people. They really just care about beating people over the head about the new tech and what is allowing these things to happen," he said. Personally, he said he would be happy with creating a performance with mappings that eluded even the most tech-savvy spectators. His interest laid in integrating interactive media into rigorous performative art forms to produce results that could not have been achieved with traditional rendering software.

Visual feedback. My visual interface provided feedback on how I was using the virtual instruments. When commenting about the inspiration for the visuals, Peter said:

> This aesthetic, pretty much from the get-go, was intended to be primarily sort of a heads-up display, a visual analog to the instrument that we were creating that actually was functional. The simplicity, cause-and-effect, and that sort of thing is extremely important because you are actually providing a performer with a tool that they need to be able to rely on.

The other two forms of feedback consisted of the audio and muscle memory. I mainly used these two forms of feedback for most of my instruments. The melody in *chorus 1 & 2,* on the other hand, required me to watch the scrolling notes and slider to calculate where to move next.

> I think that it's invaluable to have some sort of heads up display so you're not just swinging in the wind and hoping you're producing the right tones. - Mitchell Vandenburg

Peter would design the interfaces to be more legible and functional if the performer relied on them for feedback. For example, Adam could play with the wristbands entirely by ear and the dancers could perform without looking at the particles. On the other hand, I needed to see the interface of some of my instruments to play them correctly. "Everything surrounding [Yago] was quite functional. As we got to the other two band members and the dancers, we stepped away from function and was more about aesthetic and reactivity," Peter said. Each performer had a unique relationship with the visuals, and the degree to which their images were legible depended on the interactive task they were performing.

Visual feedback was particularly important for continuous gestures—those that control a range of values as opposed to on/off triggers—because it provided the performer with information about where they were in the possible range of motion. Adam's percussive instrument consisted mostly of hit gestures that did not require feedback because they were instantaneous triggers. Controlling the RealSense, however, presented some problems, as Adam explained:

> Actually getting that one particular forward and backward thing dialed in, that was probably the most, in terms of consistency, difficult one in the whole set of sounds that I had. Because it was based on crossing this imaginary line that sometimes was where I thought it was going to be, and sometimes it was a little different.

In fact, the RealSense did have reactive LED lights, but they were not designed to be informative. Mitch also stated that he would have liked to see what notes he was selecting: "There was no way for me to control individual tones. It would have been cool to have a visual aspect to it, so that by the time I articulate with my right hand that it's time to play a note, I know that note is going to come out," he said. Adam and Mitch could have had more precise control over their instrument's continuous parameters if they could somehow see a visualization of how they were affecting it.

Although the dancers did not have to see the visual particles to perform, they still saw it as an indispensable part of the choreography and an inspiration for their movements during the show. Reema explained:

> I couldn't make the dance independent of the visuals because they were so intertwined. I wouldn't perform this dance by itself…. When we walked right onto the stage of [the Consumer Electronics Show], we could see the visuals, we could really have fun playing with it and dancing—we could see what we were affecting. And that affected our choices of making the movement as big as possible and as sharp to draw as much in the air. For 360 we couldn't see it so we just had to dance very specifically in the way we had imagined it.

I felt the same way with most of my instruments. Even though I could play the instruments by ear, I also looked at the projection surfaces as a source for improvisational and playful interaction. Not all interactions require audiovisual feedback, but its presence inspired interaction and improvisation.

For the most part, my interactive gestures always had a visual response on the interface. When it became clear that the Dance Match mapping for *verse 2* would not have a reactive visual element, I decided to drop it. The blue trail created for that part did not reflect the matching algorithm.

Audio Feedback. Auditory feedback was limited to sounds that fit the harmony and rhythm of the music; they could never serve a purely functional role. Interactions that featured boundaries and selection zones were particularly unsuitable for audio response because all sounds had to fit the musical structure. For example, during *breakdown 2/2*, Mitch selected a pitch with his left hand and triggered it with the right hand. If the selection gesture with his left hand continuously played sounds as a form of feedback, it would easily inundate the music. The same can be said about my note selection instrument in *chorus 1 & 2*, where my hand glided underneath the falling notes to play them. A continuous, gliding sound would not fit aesthetically with music Zackery had composed. Furthermore, all the interactive audio tracks and the backing tracks mixing in together made it hard to hear each gesture individually. Contrary to visuals, the way sounds are mixed in concert arenas limits our ability to spatialize and segregate the sources. Mitchell hesitated using some of his interactive gestures because "there was a lot of noise going on during [*breakdown*]. I was more or less trying to stay out of the way," he said. Visual feedback was particularly useful when the interaction could not be represented by a sound that fit the harmony and rhythm of the music.

Elegance. Mappings comprised of a single gesture controlling a single audiovisual effect established a clear cause-and-effect relationship. When it came to designing the engagement between the performers and interactive media, Intel and the musicians showed a preference for elegant solutions that embraced simple one-to-one mappings. Intel wanted the audience to understand that their technology was controlling the audiovisual effects. "Obviously, in a more corporate setting trying to imply a little more straightforward connection, there is a mandate to keep it simple so that people can see the cause and effect," Peter said. The musicians wanted to keep the instruments uncomplicated to avoid technical challenges. Adam explained:

> I feel like for the performer, simpler is always better. I had a very simple set of parameters. The more you add on, the more of a learning curve and cross triggering. I personally kind of prefer having that limited set of four or five things you can choose from and that's all you got. Because that's really all you need.

Peter and I strived to make the interaction elegant and legible by limiting the number of parameters a performer could control and connecting a gesture with only one effects parameter.

> One thing that these real-time systems often need is a simplicity in
> their aesthetic because of the performance needs. - Peter Sistrom

An elegant approach to interaction was also reflected in the instrumentation of the piece. By occasionally limiting the number non-interactive audiovisual elements, we ensured that interactive gestures were highlighted throughout the show. The *intro* section, for example, featured only interactive media, to communicate right from the beginning that the show was interactive. During *breakdown 2/2*, the musicians introduced their instruments one at a time. "We singled out individual sounds and individual people in spot lights to help [the audience] make that connection," Mitchell said. We tried to minimize the overlapping of the interactive media with the fixed media. Furthermore, my instruments, which boiled down to five different mappings, were reused when a particular section repeated.

Peter, on the other hand, was more interested in complex, nuanced mappings that were not straightforward. “Personally,” he said, “my interest is not so much in the bone simple thing, but really the melding of the two into something that is so amorphous you can’t discern what is coming from where.” Nevertheless, his visuals had to adhere to a simpler aesthetic because of the performer and client needs.

Collaboration

The collaborative aspects for *Curie* were multifaceted. Most of those interviewed demonstrated a preoccupation with how resources were managed, time was spent, and how to respond to change. The melding of the client’s marketing goals, the artists’ creative ambitions, and performer needs indicated more than anything the need for flexibility in attitude as well a guiding principle for system design. Time alone and time together was crucial to the positive development of the artworks and technology. Overall we had enough time to address the challenges, but we will also see consequences of when it was inadequate. The geographic dispersion of the team and the bulkiness of the technology presented challenges that will be discussed here. We will also explore how the music, visuals and choreography co-evolved and the unequal influence each had on the other. The last two items reveal the crucial role that experts and funding have on large projects.

Flexibility. The technology had to adapt to performance needs as they arose. Ciholas changed the configuration of the ultra-wideband network in response to the number of performers and sensors. The gesture algorithm for the musician and audience wristbands could be programmed to apply different calculations and inertial sensors. Meanwhile, the network configuration changes in Syncphony and The Portal greatly facilitated the flow of information. Moving part of the sensor data processing away from the wristbands into The State Machine meant that I could more easily manipulate the data. Changing the processes in the wristband's firmware required writing code in a specific programming language and re-flashing the Curie Module by USB. A lot of reliability, usability and communications improvements to the technology were unplanned, arising instead form the interplay between the performers, digital media, and sensing devices.

> It was an on-going process. Making something, listening, gaining new information, and thinking about how to move forward, what can I bring with me and what is no longer necessary for the choreography. - Reema Bounajem

Although most of the mappings were laid out early in the development process, some had to be created to accommodate technologies introduced by Intel. The audience interaction section of the piece was finalized very close to the show date. We had to create new audiovisual material for the audience and incorporate PixMob's data flow into the network. We had to work around the hand position jitter of ultra-wideband and figure out alternative sources of movement data. Similar to the approach discovered by

Swarnendu during the programming of the wristband's gesture algorithm, we figured out that having alternative sources of movement data—accelerometer, gyroscope, radio motion tracking, optical motion tracking—helped overcome issues in any particular source.

The dynamic environment for this project demanded a flexible mapping organization scheme. The technique Peter and I pursued encapsulated gesture-media mappings into modules assigned to sections of the song. Changing the mapping assignments in one module did not affect the other modules. "It made more sense to make more modules that had different interactions so that we could more easily change things up," Peter said. The creative team's ability to adapt to technological changes was evident when Intel rejected our suggestion to use Microsoft's Kinect camera, prompting a reconfiguration in the artwork and mapping scheme.

The ability to quickly map the performer's gestures to different sounds was regarded as a very useful feature. "I think that is the most impressive part is being able to customize each individual sound and tailor it to what you want," Mitchell said. Adam echoed this feeling when he said:

> The freedom that we had in programming the gestures we wanted was really cool. Meaning I wasn't really tied down into doing some unnatural stuff that felt weird. We were able to kind of pick and choose [different combinations of gestures and sounds] until it all felt very natural and organic.

This level of flexibility could only be achieved by being able to access the mapping in real time in The State Machine and Ableton Live. Having a programing environment that allowed changing the mappings and sound assignments on the fly enabled the creation of usable and interesting instruments.

The music, visuals, and choreography required different degrees of flexibility. The song defined the timeline, structure, pace, and instrumentation. After the song was composed, the only two changes we had to do were incorporating the audience interaction in breakdown and trimming the verses to shorten the song's show length. Structural changes to the song required modifying the mappings and other artworks. The interactive visuals were created in response to the interactive instrument and the type of sensor data we were using. Peter designed the visuals in TouchDesigner, which he described as high-level, modular, and open-ended. Furthermore, he said that it accepted "a significant amount of inputs from other protocols that other artists and creatives use to transmit data." The choreography experienced a lot of changes throughout the development due to the placement of the sensors and shortening of *verse 3*. All the art forms had to adapt to each other's development and the technological demands; the music changed the least and the choreography the most.

Some of the technological and artistic elements of the show became less flexible as they developed. Reema and Ciholas, for example, would have liked their roles to have been specified more clearly at an early stage. Ciholas, the owner of the company that developed the ultra-wideband system, admitted that he would have designed it differently if he had a better idea of how it was going to be used. Reema wanted to keep things as stable as possible towards the end of the development phase.

Time. Having time to collaborate, experiment and practice with the art forms was instrumental to their success. Peter pointed out the importance not only of spending time tinkering with the sensor data but also to “rehearse and practice with them and find out what you can do with those possible things…. And you find all the cool things you can do with it.” By working together with Reema, he developed a responsive particle system, while she choreographed movements that enhanced his visuals. Last-minute changes to the song and visuals close to the show day did not give her enough time to change her choreography. I too had to spend time with Zackery to pick apart his song and incorporate the sounds into the interactive system. Since we lived in different cities, Peter and I shared our ideas by providing diagrams and talked over the phone. Then we would each go our own way. It would have been beneficial to spend more time collaborating on the audiovisual material. The creative team tried to meet as often as possible to share and respond to each other’s ideas.

> If we had more time we could have done it. - Jean-Sebastien Rousseau

Behind the scenes, a lot of time went into building, coding, and testing the technology. Major changes to Syncphony, the ultra-wideband system, and the wristbands came about after the artists spent time with the developers. When talking about

Syncphony, Jason stated that “the usability requirements were really originally driven from the engineer’s perspective … but once we started bringing you to the table it became more collaborative.” Tim echoed this when he said that “you guys would come to us and say ‘we think we want to do this.’ We’d go back to the drawing board and say is that possible?” The time Adam spent with Swarnendu testing the wristbands led to the creation of the side hit gesture that became crucial for my instruments. Reema also spent time with Intel figuring out where the ultra-wideband sensor pods would be placed on the dancer’s body to enable motion capture but not hamper movement. I spent a whole weekend with the RealSense team to discover the best camera and repurpose their software for our ends.

On the other hand, the dragonfly design for the RealSense was made by Intel with no input from us and the final product turned out bulky and difficult to play. But most of the technologies were developed in collaboration with its users.

Full-stage rehearsals with all the technology, art, and performers were crucial for practicing the instruments and debugging the technology. The instruments Adam, Mitchell, and I played required some time to learn. “It just took a little time to like ‘OK, this is how this instrument works,’” Adam said. Mitchell noted that he did not play with the RealSense because he did not have enough time to be completely comfortable with his instrument: “Given more time with the instrument would have led to further exploration of the effect field.”

Errors were also spotted during rehearsals. Network issues related to the ultra-wideband and 900 MHz systems surfaced after repeated use of the wristbands.

Computing and communication systems were streamlined to make the technology more autonomous and reliable. Both the performers and technical team were able to improve their contributions during rehearsals.

Resources. The team's geographic disparity made it hard to access the tools and talent needed to develop, test, and practice the show's technical and artistic components. *Curie's* technicians and artists were located in five cities: Los Angeles, Austin, New York, Santa Clara, and Portland. Providing a fully equipped and staffed space to rehearse at Evansville and Las Vegas was crucial to the success of the show. Otherwise we could not have incorporated the PixMob, optimize the wristbands, practice the interactive instruments, refine the audiovisual media, and debug the technology.

For most of the development process, the sensing technologies—RealSense camera, 900 MHz, ultra-wideband, and PixMob—could only be operated in the rehearsal spaces with the network infrastructure and the technical staff. The UWB system required a large sensor arrays, a server rack, and someone to run the server. Theoretically, the 900 MHz could have run on a laptop, but the USB port configuration was never configured properly. The RealSense rig was only completed close to show day. The PixMob system required a large number of bracelets, IR emitters, and a camera array. The dependency on network infrastructure and technical staff meant that the musicians, dancers, and I could not practice interactive gestures outside of rehearsal times.

Imagination, specifications, and recorded data were used when we could not access the talent and technology. Outside of the rehearsals in Evansville and Las Vegas, Peter, Reema, and I had to use our imagination to contextualize our work around the other art forms. Reema said that at the initial stage of the development, she had to imagine how the show would look like. Artistic decisions about the visuals, performer placement, and movement on stage had to be mentally juxtaposed with the 'big picture.' Specifications about how the sensors functioned and sensor-recorded data served as a substitute for the real thing during development. In fact, most of the mappings were coded without the sensors. We had an idea of the sensors' value ranges and mapped the data we expected to receive to the audiovisual parameters. Imagination and "fake" sensor data helped overcome the lack of accessibility to the talent and technology.

Artwork. The song outlined the structure and defined the mappings which in turn inspired the interactive visuals. The choreography responded to both the song and visuals. Music was considered the fundamental art form for the show and the main content to be manipulated by the technology. The audio mappings emerged from a combination of the song's instrumentation, my experiential knowledge of motion tracking, and what I imagined would be the accompanying visual interface. These audio mappings and suggestions for the visual interface informed part of Peter's designs. He stated that the interface "represented the instrument." Their visuals did not influence the creation of new interactive music instruments but it did alter some of my gestures. Peter spent time recording Reema's movements so that he could see how they would affect his particle system. Reema's choreography took into consideration the particle system and song structure. The music influenced the visuals, and together they informed the choreography.

> The language that new media artists speak a lot more is not so much audio or visual but the two together. - Peter Sistrom

Gestures were conceptualized together with the interactive sound and visuals. If this gesture was not captured effectively by the sensor, then one of four things happened: 1) a suitable input gesture was added to the original gesture; 2) the original gesture was kept and the software was changed to accommodate it; 3) the original gesture was modified to fit the sensor's requirements; or 4) the mapping was rejected.

The artists suggested that we strive for simplicity among such a complex collaboration. There were different reasons to keep things simple. Jean-Sebastien explained that giving simply one gesture to the audience interaction was important because there was not time to explain more complex interactions: "We don't want to go too complicated in this kind of event. It went by so quick, we didn't have time to explain anything."

Both Adam and Mitchell noted their preference for keeping the instruments simple to improve reliability and usability. "We were finding the more simple it was the less problems you would run into," Mitchell said. Finally, Peter observed that, "it's probably important to stay simple so that you don't lose yourself building these systems…. But certainly from a programmer or sculpture standpoint in these systems, simple is necessary." Keeping the mappings simple among such a complex technological system ensured the intuitiveness and fidelity of the interactive system.

The interactive visuals and music had backing tracks in case of sensor or network failure. There was a real danger that the technology would not work during the show. Because we were pressured to impress the audience regardless of system failure, we had to create a backup plan. "Because of the nature of those particular shows, you have the finger on the 'oh shit' switch, which attempts to slide in some pre-done recording," Peter said. At any point during the show, we could switch between the interactive tracks and the prerecorded ones.

Experts. Curie required a wide variety of experts to create artistic content and make or modify new technologies. We had to produce an original work and improve sensors that had not been originally developed for our particular purposes. The three artistic disciplines—music, visuals, and dance—were each managed by an expert with intimate knowledge of the field. Peter and I could delegate the task of creating content to others so that we could focus on the mappings. Intel's mission to retrofit their Curie Module into wristbands for the audience and musicians required the knowledge of various experts from fabricators to software, circuit, and electronic engineers. The algorithm to detect the gestures was the sole focus of one of the programmers at Intel.

The ultra-wideband system required mobilizing a whole team to build the support beams, calibrate the radio anchors, code the motion tracking algorithm, and transport it to the stage. The venue also designated specific roles to the lighting designer, production manager, stage manager, etc. To create new content and technologies it was necessary to invite experts from different fields to design the art work or provide technical solutions that would enable the realization of the artistic vision.

Funding. Enabling access to space, equipment, talent, and experts required considerable funding. Curie was particularly expensive because of the geographical dispersion of its team, the size of the show, and the development of new technologies. We had to hire talent and experts in multiple fields to help develop the content and technology for the show. These people had to be flown around the country to experiment and rehearse with the sensors. Meanwhile, the technology also had to be transported to the stages in Evansville and Las Vegas. The pressure to create an extravagant show and the need to repurpose Intel's technology demanded considerable resources.

Chapter 4: *Ad Mortuos*

Date: March 6-14, 2015

Location: Oscar G. Brockett Theatre, Austin

Video: https://vimeo.com/123801965

Ad Mortuos was a dance piece featuring interactive music and visuals that responded to brain signals, facial expressions, and body movement. I wore various technologies that allowed the broadcast of brain and body bio-signals through Wi-Fi connection. Elevated seating surrounded the stage and visuals were projected on the floor. The piece was created for the MOVE! concert series in the Department of Theater and Dancer at the University of Texas at Austin.

This case study is divided into three chapters. The first describes the origins of the show and the media and venue specifications. The development chapter starts by detailing the improvements made to the EEG sensor and how we made it mobile. It also describes two approaches to motion-tracking and software programs developed for practice purposes. The next section provides information about the creation of the mapping scheme. The last three sections in the development describe how the visuals, choreography, and music were created. The last chapter analyzes technological, transparency, and collaborative issues that arose during the development and production of the piece.

Team	
Poet	Stephanie Pope
Spoken Word	LaQuetta Carpenter
Visual Artists	João Beira
	Rodrigo Carvalho
Choreographer	Yacov Sharir
Composer	Bruce Pennycook
Dancers	Becca Bagley
	Gianina Casale
	Katie McCarn
	Allyson Morales
	Summer Fiaschetti
	Emily Snouffer
Singer and EEG Performer	Yago de Quay
Costuming	Kelsey Vidic

Table 4.1: The team responsible for *Ad Mortuos*.

CONTEXT

Origin

Ad Mortuos was a multimedia dance piece created for the MOVE! Concert series that tracked brainwaves, facial expressions, and movement. The black box venue had audience members sitting on elevated seats on three sides of the stage, giving them a good view of the visuals projected on the floor. Using the venue and technology as a starting point, the team met regularly for five months prior to the show to brainstorm ideas. Aside from creating the artwork, considerable work was invested in improving the reliability of the sensors and making the interactive system portable. Conflicts with the production team had an impact on the placement of the projections and speakers.

Yacov Sharir, the choreographer, was interested in "exploring the not-knowing zone" regarding the use of brainwaves in dance pieces, prompting him to assemble a team of experts to create a piece. Bruce adds that he and Yacov wanted to produce a "new major work … a piece that was long enough to be a complete idea." They invited me to perform and develop the EEG technology. Later, we invited Rodrigo to do the visuals.

***Stage*.** The Oscar G. Brockett is a black box theater that provides an intimate setting for 255 elevated seats surrounding the stage. *Ad Mortuos* was performed there eight times during two consecutive weekends. The venue has four rows of elevated seating on three sides of the stage (*Figure 4.1*). With the exception of the marley floor, the whole venue is painted in black. The projectors and speakers were placed above the stage. The projection area is represented by the blue rectangle in *Figure 4.1*. The speakers were facing outwards towards each of the audience seating areas.

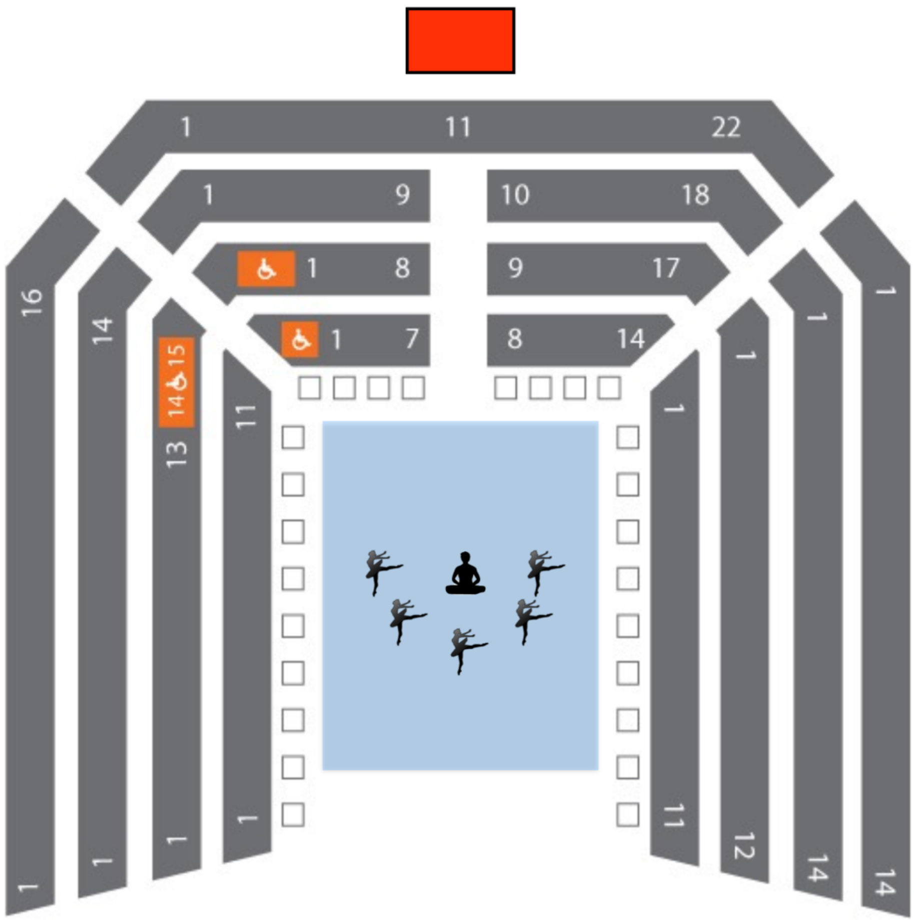

Figure 4.1: Oscar G. Brockett Theater layout. Five dancers surround the EEG performer. The blue rectangle delineates the projection area, the red rectangle the media booth.

Rehearsals. We had nine rehearsals in the space in the weeks leading up to show day. The rehearsals almost always had an audio engineer and lighting designer there to help us. The first few days in the space were dedicated to setting up the projector and computer system. Before that we ran tests in smaller studios.

Media Technology

Media Station. The graphics and music computers, together with their operators—Rodrigo and Bruce, respectively—were housed in a booth on the second floor that overlooked the stage. The tablet worn by me communicated with a Wi-Fi router in the booth.

Sound. The speaker arrays were rigged above the stage facing out towards the audience. The mixing table was located next to the audio computer in the booth. I wore a discreet headset microphone that was routed to the audio computer through the mixer. Except for the voice, all the audio was produced from the software in Bruce's computer.

Projection. The Texas Performing Arts organization lent us two 5,000-lumen projectors that were set up side by side on the ceiling pointing down to the stage. We had to hire our own rigger to install the projectors. The signal from the laptops was sent to a VGA matrix, split in half, and routed to each projector.

Lighting. All the lights were rigged on the ceiling. We used mainly the red hue. The venue provided us with a lighting designer and a stage manager. The projector and lights sometimes conflicted with each other because they were both pointing down on the same area.

Signal Network. All the input data was aggregated in a Windows tablet before being sent to the media computers. The EEG sensor was connected wirelessly to the tablet via a USB dongle. The connection suffered from interruptions when they were more than a few feet apart so I kept the dongle hidden in the hood next to the EEG. Brainwave data was processed in Emotiv Control Panel then shared by UDP through

MindYourOSC. An iPhone reported the gyroscopic information via the Wi-Fi network to the tablet. The brainwave and gyroscope data were aggregated in a Max/MSP patch before being distributed by OSC to the music and visuals computers (*Figure 4.2)*. The music was generated in a Max/MSP patch and the visuals in Quartz and VDMX.

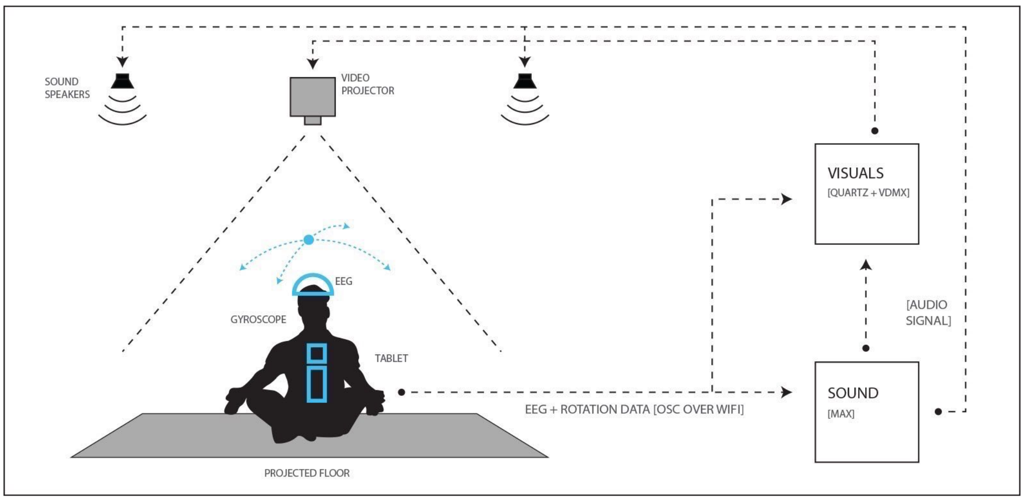

Figure 4.2: Signal network path.[4]

DEVELOPMENT

This part describes *Ad Mortuos'* technological, mapping, and creative developments. Despite some changes to the hardware, most of the technical work was invested in training the EEG algorithms, coding the network protocol, and creating

---

[4] Image designed by Rodrigo Carvalho.

software for the team to experiment with. The mapping section explains how the brainwave, facial, and motion data was connected to the interactive media. Some issues related to the quality of the data are described here. The last sections provide insight into how the visuals, choreography, and music were created based on interviews with Rodrigo, Yacov, and Bruce, respectively.

The following sections are dedicated to the technological developments for *Ad Mortuos.* They comprised mostly of implementing existing devices in a performative context. The first part describes how Emotiv's EEG headset was adapted to make it portable and reliable. I had to make the headset more secure and connected it to a portable computing unit so that the EEG data could be analyzed on the user. The next section delves into two motion-tracking strategies pursued by the team. Tracking multiple performers on stage was dropped because it conflicted with the venue's production resources, so we opted for simply tracking me on stage. The last section briefly describes two software applications created by Rodrigo and me that enabled the team to experiment with prerecorded EEG data and media.

EEG Headset

Emotiv EPOC is a sensor-and-software package that includes a 14-channel electroencephalogram (EEG) and software with algorithms that interpret the data. The EEG relies on wet felt pad electrodes to capture signals. The electrodes attach to a plastic headset that the user can put on himself without the need of a technician (*Figure 4.3*). The headset communicates by a proprietary wireless protocol through a USB dongle to a computer hosting the application.

Figure 4.3: Emotiv EPOC electroencephalogram headset.

Emotiv's software program extracts facial expressions, affective levels, and cognitive commands. *Figure 4.4* displays the four panels inside the software for each of the detection modes. Facial expressions consist of electrical impulses on the forehead caused by smiling and raising eyebrows. Affective levels display the intensity of mind states consisting of excitement, engagement, frustration, and meditation. These levels are measured against benchmarks created by the Emotiv research team, although they are, to some extent, adaptable to the user. Cognitive command is a pattern matching algorithm that detects particular mental tasks.

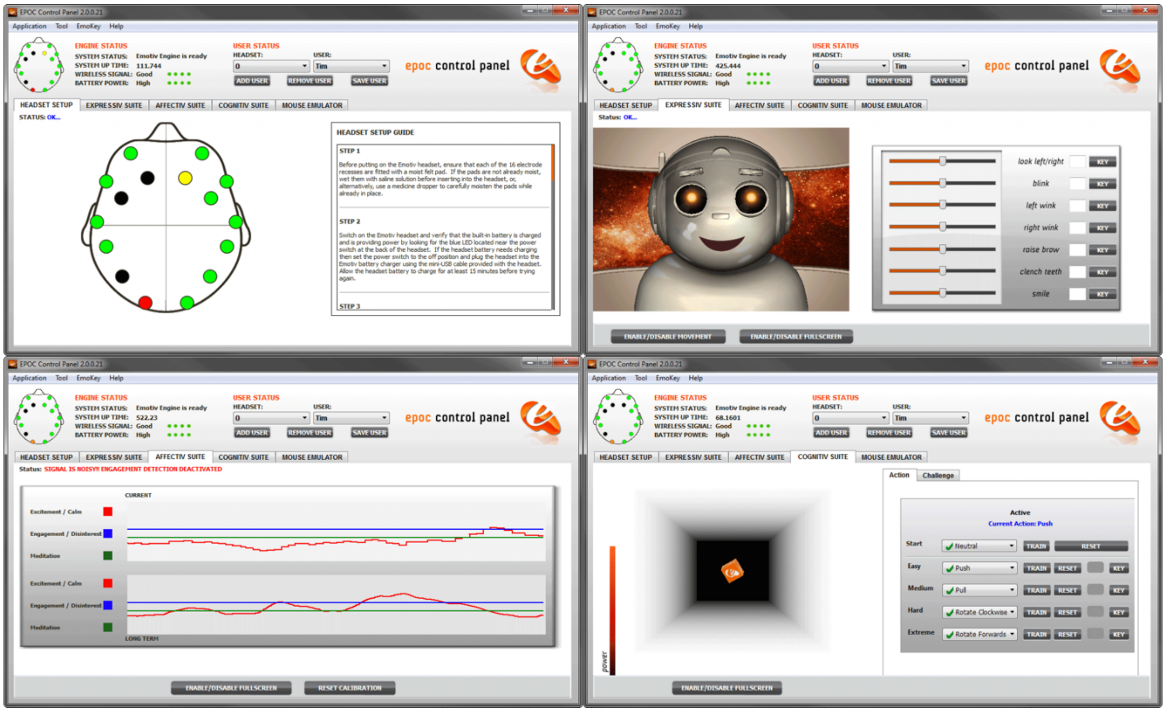

Figure 4.4: The four different information displays in Emotiv's Control Panel. From top left, clockwise: Sensor quality, facial expression, affective levels, cognitive commands.[5]

The process of creating a cognitive command starts by having the user perform a mental task—most likely imagining displacing an object—while recording the brainwave pattern. The user can apply an action on a 3D cube in the control panel (bottom right panel in *Figure 4.4*) to help mentally visualize the movement. When in detection mode, the software compares this recording (and other recordings) with the real-time pattern of

[5] Used with permission from http://www.hksinc.com.

activation. By repeating a mental task, the user can command the interactive software to execute a function, or in our case, modify an audiovisual parameter.

The headset had to be reinforced to increase sensor quality, because the curved plastic beam that wrapped around the head was fragile and did not put enough pressure on the scalp. The fragility of the beam increased the chances that the beam would break during the show and reduced the quality of the electrode's connection. In fact, since the headset was not custom-made for my head, some of the electrodes did not touch my scalp at all. I fastened a curved metal strip to strengthen the plastic beam and tighten the curvature of the headset. The resulting structure increased sensor connectivity but also its weight.

The felt pads required daily maintenance to retain a good connection. I purchased new felt pads shortly before the show because the old ones had lost some of their conductivity. Emotiv recommends damping the felt pads with saline fluid, but I noticed a better performance when they were completely soaked. High levels of saline reacted with the silver-based coating that interfaced between the wet felt pad and the lead wire, creating a green grunge that had to be scrubbed and cleaned occasionally.

Tablet Computer

Issues with Emotiv's wireless connectivity was solved by using a wearable tablet computer. The EEG headset could connect via a proprietary 2.4 GHz protocol or Bluetooth (4.0 Low Energy). Both protocols were prone to interference from the surrounding equipment, walls, audience cellphones, and human bodies. I found out that the connection was only reliable when the headset was very close to the computer's wireless receiver. There was no optimal place in the venue to install the wireless receiver since I needed to activate the EEG headset in two spaces separated by two concrete walls: the stage and dressing room. Furthermore, I had to watch the Control Panel while putting on the EEG to check the status of the electrodes, so the computer had to be visible at all times.

The tablet computer enabled the EEG and software to be fully functional and wireless while I moved around the stage and dressing room. This freedom of movement allowed Yacov to choreograph my movement on stage unencumbered by the technology. We chose the Dell Venue 8 Pro 5000 tablet because it could run Windows 8 (required by Emotiv) and had a USB port to connect the wireless EEG dongle. The battery lasted up to four hours. *Figure 4.5* shows the tablet running the Control Panel and connected to the EEG. Kelsey Vidic, the costume designer, created a transparent, touch-sensitive pouch on the chest area of my costume so that I could carry, view, and interact with the tablet. A USB cable went around my back and into my hood, placing the wireless dongle less than two feet away from the headset.

I created a simple visualizer for the tablet with the intention of illustrating the brainwave data to the audience and myself. It exhibited levels of meditation, excitement and facial expressions. However, the Max/MSP patch generating the visuals was crashing when the real-time data started pouring in. The final visualizers were randomly generated in the app to keep its aesthetic purpose.

The Dell tablet had a manufacturing bug that occasionally jammed the Wi-Fi network adapted. The solution required disabling and re-enabling the adaptor in Windows' device settings.

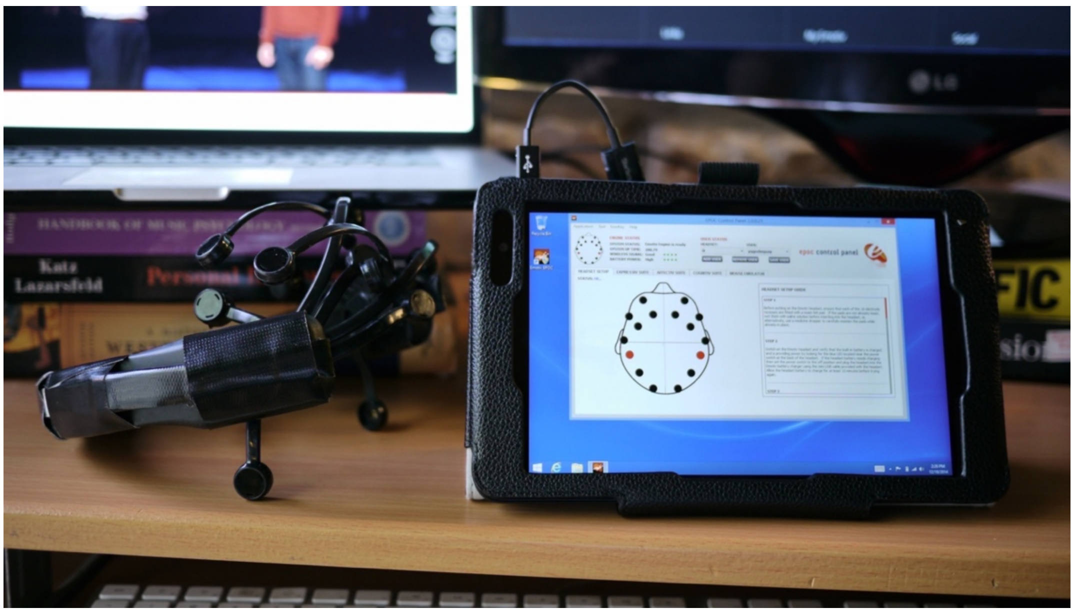

Figure 4.5: Emotiv's EEG headset and the Dell tablet.

LED Lights

We added lights to the headset to illuminate my face and highlight the technology. Kelsey constructed a metal wire bridge on the front of the headset that held small button LEDs to illuminate my face. I required illumination because we could not use the venue's lights on me since they would interfere with the floor projections. I tried to make the headset more visible by highlighting it with a battery-powered electroluminescent strip.

However, the strip's electrical current interfered with the EEG readings so it had to be discarded.

Figure 4.6: LED buttons inside the frontal bridge illuminated my face. The discarded electroluminescent strip attempted to highlight the EEG.

Motion Tracking

The iPhone's gyroscope sensor was added because we could not extract the same data from the tablet or the EEG headset. I worked with three programmers on Emotiv's SDK but we were not able to get this data to stream correctly. As an alternative to these sources, I used the iPhone app GyrOSC to transmit inertial information over OSC protocol. GyrOSC could only send data to one client, so I routed it to the tablet, where it was redistributed to each of the media computers. The iPhone was connected to the same Wi-Fi as the tablet and media computers, but the Wi-Fi range did not reach the dressing room, so I could not verify if GyrOSC was working until I went onstage. The orientation needed to be calibrated before each performance. The iPhone was placed inside a small touch-sensitive pouch made by Kelsey in a shirt underneath my costume.

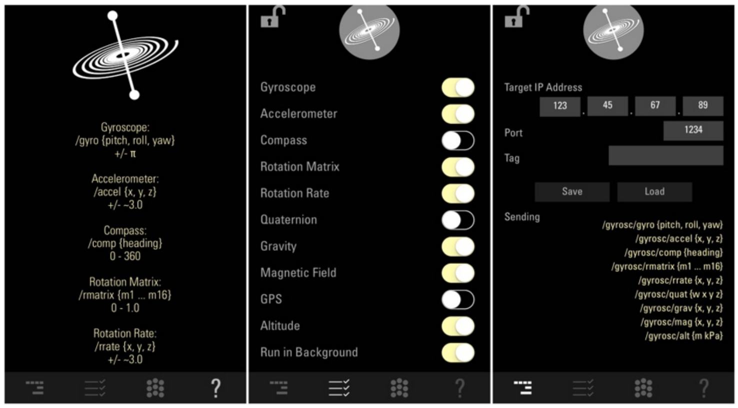

Figure 4.7: Screen shot of the GyrOSC iPhone app.[6]

Position Tracking

Rodrigo and Bruce considered tracking the performers with a ceiling-mounted Kinect camera, but the idea was dropped because of technical and time limitations. The first problem was that the production team was not available to help us mount the camera at the appropriate height. The second was that the camera interfered with the front projection to the cyclorama.

---

[6] Used with permission from www.bitshape.com.

“We wanted to use an overhead camera but didn’t. because the ceiling was too high for a Kinect to work and using a DIY camera system would have been too complicated,” Rodrigo said. “We had little time in the space. That’s one of the reasons we did not use a camera system.”

Bruce had initially thought about mapping movement on the stage to localized sound, but the venue was not acoustically suitable for surround sound. The loudspeakers were very close to the audience and emphasizing some loudspeakers more than others would have been acoustically unpleasant. He also rejected the idea of having soundscapes mapped to the performer’s position in the room. He explained why he ultimately decided against that: “For me, the piece was a more linear dramatic narrative than the abstractness that any kind of generative system would have brought into play.”

Practice Software

Prompted by the lack of a practice space to collaborate, Rodrigo and I each developed an application that simulated interaction. Rodrigo’s application, compiled in Processing, replicated the visuals that were going to be projected on the floor, as shown in *Figure 4.8*. It accepted OSC data and responded with all the visual transformations described in the mapping list. My EEG player, presented in *Figure 4.9*, streamed stored facial and cognitive data for Bruce and Rodrigo to experiment with. It allowed them to play with the input data without me being there. When we started to practice in the rehearsal space, our individual interactive components were prepared to work together.

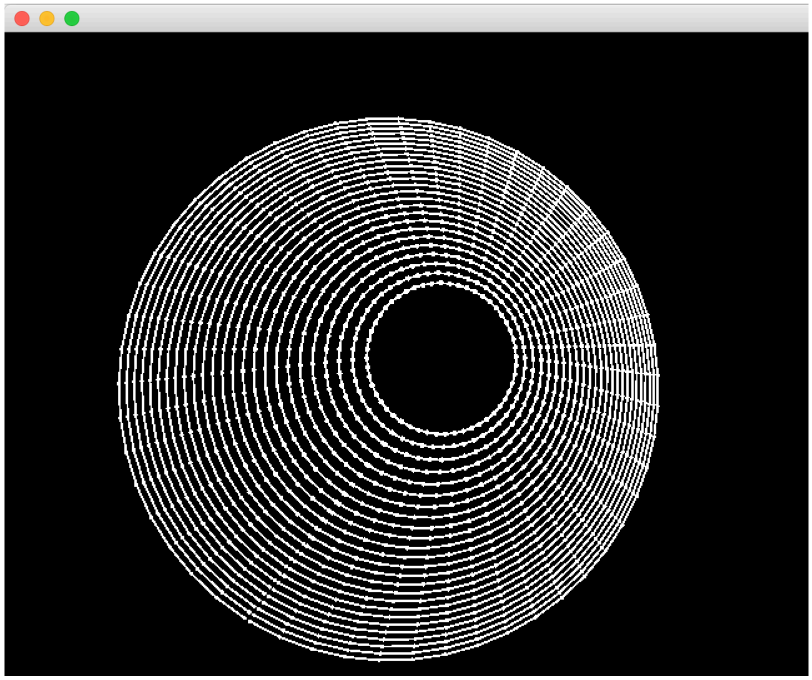

Figure 4.8: Gyromove app created by Rodrigo Carvalho for testing purposes.

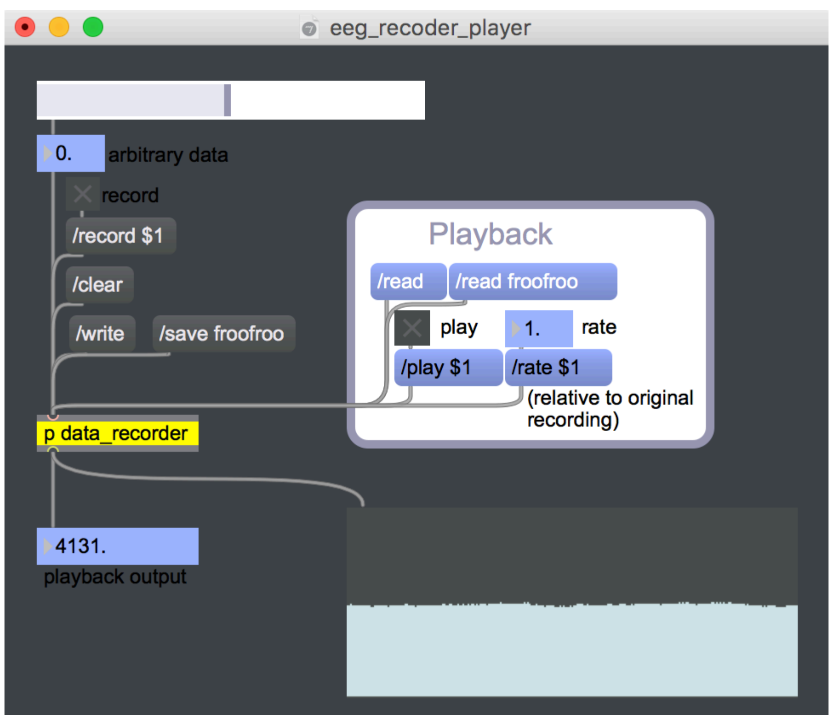

Figure 4.9: EEG recorder and player for experimenting with pre-recorded brainwave data.

Mappings

This section explores how brainwaves, facial expressions, and body movement were applied to media. Most interactions featured both an audio and a visual effect. EEG monitoring of brain and facial activity produced modest reliability results. Even after training the software to better recognize my activity, we noticed considerable false-positive triggers. The full potential of the interactive system was not explored because my movement on stage reduced the quality of the EEG data.

Floor Projections	*Performer* Brain/Face/Body Signal	Sound
	EEG Section	
	Yago de Quay	
Radius of concentric rings	Brainwave pattern	Frequency shift
	Eye blink	Percussive sound
	Smile	Time-based delay
Distortion of circles	Furrow brow	Time-based delay
Slant of concentric circles	Torso	Repeat-rate and pitch

Table 4.2: List of mappings used in *Ad Mortuos.*

Mapping Development

Brainwave Patterns. Closing my eyes commanded an increase in the frequency of a frequency shifter on a track and the expansion of the concentric rings. I raised my hands during this command to help establish a connection between me and the interactive media.

I trained Emotiv to recognize when my eyes were closed because it was the most reliable brainwave pattern. The system analyzed the brain in real time and matched the current brainwave pattern with other recorded patterns of activation to produce the most likely match. Emotiv recognized up to eight brainwave patterns that normally consist of imagining moving an object in space in eight different ways. Each movement produced a particular pattern of neural activation that was understood as a unique computer command. Recreating that imagined movement with accuracy is very difficult, particularly on stage when there are other stimuli present.

Technical difficulties, illegibility, and the limited number of interactive parameters provoked me not to incorporate more than one brainwave command. The problem with increasing the number of patterns are twofold: It reduces the tolerance for identifying a pattern which can decrease recognition, and demands a more precise and clean mental thought that is harder to reproduce. Furthermore, the performer's process of creating different brainwave patterns is not picked up by the audience. I thought it unnecessary to complicate the performance if the mental action was not apparent to them. Lastly, the array of mental, facial, and body action data streams were enough to address all the available mappings.

I tried to develop my own machine learning algorithm, but the real-time raw EEG values did not produce useful data, and performing effective spectral analysis of brain signals for brain-computer interaction was difficult. Furthermore, MindYourOSC did not transmit raw EEG—I had to rely a separate Javascript project included in the SDK that, together with MindYourOSC, was too complex and CPU-intensive to run on a tablet. I trained Emotiv's detection algorithms to yield better results, yet many of the mental and facial actions were dropped because they were unreliable or illegible.

Facial Expressions. Eye blinks produced a variety of short percussive sounds. Borth smiles and furrowed brows activated and deactivated different settings of a time-based delay effect and distorted the geometry of the circles.

Facial expressions were visible and easy to extract from the software, but with the exception of eye blinks were unreliable. I could improve the detection by training the Emotiv's facial recognition algorithm and fine-tuning the threshold. Smile and eyebrow actions could be sustained and expressed as a Boolean value, while eye blinks acted like a trigger. The software failed to detect a sustained smile or furrowed brow without producing false negatives. This was acceptable because it did not produce many false positives. Placing my hand on my head when I performed facial expressions served to attract attention to my face and cued the dancers to mimic my facial expressions.

Torso. The angle of the torso was mapped to the slant of the cone below me to create the illusion that the floor was tilting. The music also responded by changing the repeat rate and pitch of prerecorded vocalizations.

Initially, we planned on using the EEG's gyroscope, but were unsuccessful in retrieving that data. The SDK provided by the company was incomplete and the documentation poor. An alternative had to be found because dropping the torso interaction would deprive me of my most visible interaction. The Dell tablet had inertial sensors too, but I was not able to extract that data. The only solution was to add another device, the iPhone, to the costume and transmit its gyroscopic information. The iPhone app for the task, GyrOSC, sent the data though Wi-Fi to the tablet before being routed to the audio and visual computers.

Visuals

Rodrigo, the visual artist, was invited to join the team in January, a few months before the show, while on his visiting trip as a researcher at University of Texas at Austin.

The space played a major role in shaping the visual aesthetic. Rodrigo described the black box as “an immersive environment” because the audience was surrounding it and looking from above. This angle of view inspired him to design a “tunnel vision” that created the illusion of depth using the geometry of the circles underneath me. Rodrigo also tried to match the aesthetics to that of the rest of the piece, which was designed by João Beira. His style was very minimal, consisting of black and white lines and circles.

The team was not able to fully explore the mapping possibilities together because of our limited time in the rehearsal space. The time we had on stage together was spent rehearsing and troubleshooting problems like the projectors, network, and technology, rather than making new content. More time would have allowed us to “explore more with the content,” Rodrigo said, “and the interactions between the performer and the media.” He noted that the time constraints reduced serendipity and the emergence of interesting mappings. Because of this, much of the piece was speculative until we started rehearsals.

Lack of time in the venue also contributed to the decision to drop the overhead camera system.

Rodrigo complained that overall, the interactive system was too complicated and inefficient. The complexity of the system—EEG, iPhone, and tablet—meant that there were many things to manage. One of the problems with the hardware was, “the fact that [Yago] had to have a bunch of things with [him], the iPhone, the tablet, plus the EEG. It would have been much better if we had a unified solution,” he said.

Rodrigo commented that in the end he was happy with the results. He particularly like the rotation mapping: “it was very effective on stage,” he said. If he could change something, he would have involved the dancers in the interactive system.

Figure 4.10: Rodrigo’s project in VDMX.

Choreography

Yacov Sharir was the choreographer behind *Ad Mortuos*. Aside from his role as co-creator of the overall structure, he also directed my movement as well as those of the dancers. Yacov was interested in merging dance and brainwaves. He said that once a technology has been used successfully in a performative environment, his desire is to "move on and try something new—otherwise I am rehashing something that we know already." Part of that exploration consists of listening and collaborating with other artists that have their own interests. His goal with this piece to "devise a system that the dancers will, in a coherent and obvious way, drive the technology." One of the consequences of this philosophy, he admits, is that we don't "explore all the possibilities."

Ample time during the development phase allowed the maturation of the technology, convergence of the different visions pursued by the creative team, and the formulation of the show's structure. Regular meetings for six months prior to the show were critical to its success. Yacov noted that it took a "whole semester to create a coherent and clear idea that was accepted by everyone…. We had enough time to develop the how, what and where…. So it's like we had a score." Furthermore, developing the "technology requires as much time and attention as the physical body needs in order to create a work," he said.

Specifying beforehand the limitations of the technology shaped the choreography and made rehearsals more effective, Yacov said:

> Once we figured out the content and the meaning of the work, the technology went along really well—the reason being that we had discussed during the first semester, before we made the choreography, all the possibilities that the technology would allow us to do. (personal communication, May 20, 2016).

He added that recruiting experts that can tackle the technical challenges brought by new technologies reduces the complexity of the system.

> If you have a complex system, you need a person that activates this system; and it's not complex because of that person. If you have an idea for a work that uses specific technologies, you assemble a team that has the expertise. (personal communication, May 20, 2016).

Anecdotally, he recalled how "In the 80s, we used to go to the studio and make the choreography. Three days before the show, we brought in the technology and we didn't know what we were doing."

The interactive EEG section of the piece focused on full-body motion "with specific and immediate attention to the gestures [Yago] was creating by activating the system that was attached to [his] brain," Yacov said. My movements consisted of tilt, circular, and hand motions. "Two movement systems were co-existing, the full body movement by the dancers and the gestures [Yago] was utilizing to activate the system." There were moments when the dancers mimicked me to "visually accentuate" my movements and others when they performed unique, juxtaposed solos. I suggested that the piece start with one-to-one mappings and mimicry to help the audience connect the gestures with the visuals and music.

Yacov did not regard brain signals as a physically expressive medium and was not concerned with their legibility. Brainwaves were only articulated through the visuals. He explained: “You’re not performing the brainwaves. The brainwaves merely activate the visuals that you select. Do we see your brain? No. We take the activity of the brain and give it visual interpretation.” Yacov was not worried about whether the audience understood that the brain was affecting the visuals.

Yacov choreographed my movement in the space. Before my interactive section, I sat on a chair upstage while the other performers communicated the poem in American Sign Language and danced with the visuals on the floor. Around four minutes into the piece, I stood up and sang. A few minutes later, I walked up to the center of the stage and sat cross-legged. Only my hands, torso, and face moved during the interactive section, which lasted around 3 1/2 minutes. I stood on the side of the stage vocalizing words from the poem for the rest of the piece until the finale, when I walked to the center again.

Overall, Yacov said he was very satisfied with the piece. “The preparation for the piece was very adequate. We all agreed there was no conflict,” he said.

Music

Stephanie Pope’s poem informed the show’s design, length, theme, and aura. Bruce explained how “the central glue of the whole piece was the poem … Things really didn’t start to move forward until this poem became the focus.” The final form of the piece came about after several revisions that took into consideration the art forms, technology, and creative visions. The poem and the spoken voice were commissioned by Bruce Pennycook.

The piece went through several theoretical iterations. Elements that were decided early—space, performers, EEG technology, spoken voice, singing, and sign language—were bundled into chunks and shuffled around with consideration to the desired emotional arch. As the music, choreography, and visuals in each chunk crystallized, so did its duration. The result was a three-part piece consisting of an "introduction song, the séance, and then dance," Bruce said. He added: "it's really hard to manage lengthy pieces without some formal structure."

The music for the show consisted of vocals, fixed tracks, interactive tracks, and sound effects. With the exception of the eye blinks that generated percussive sounds, all the mappings controlled sound effects parameters. *Figure 4.11* presents the user interface of the music software developed by Bruce Pennycook in Max/MSP. The input trims adjust balances of all the processes and effects. The output faders control the processes' output mix. The matrix allowed Bruce to instantly assign processes (represented vertically) to the left and right loudspeakers as well as the reverb (represented horizontally). The EEG levels panel monitored the signals produced by me.

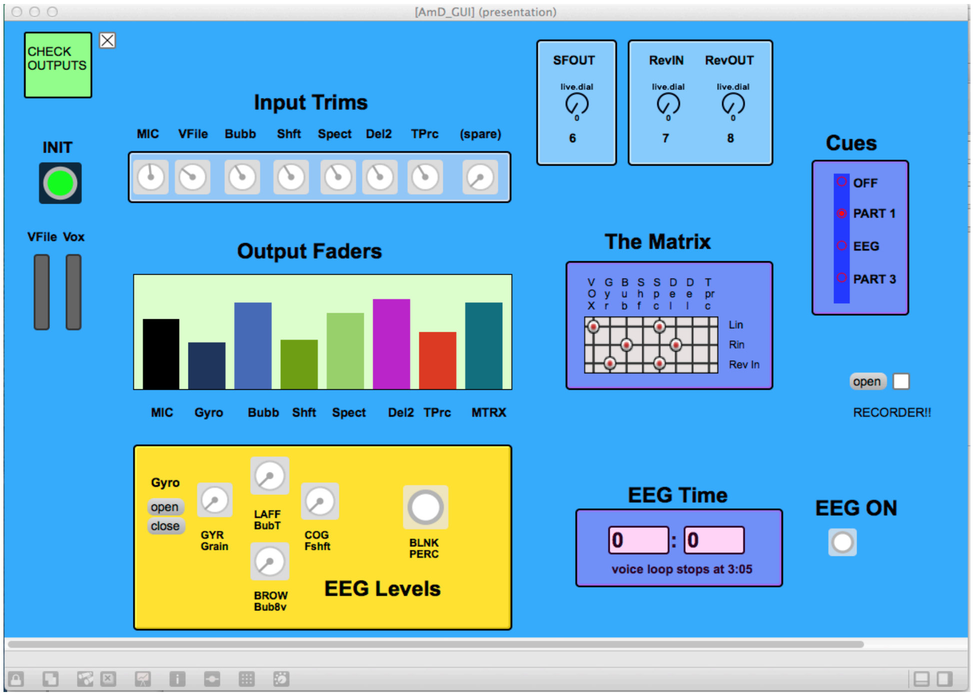

Figure 4.11: Interface for the music software developed by Bruce Pennycook.

The show wasn't fully automated; Bruce was "tweaking" the music during the performance. He exercised control over the mix and sound effects using an eight-channel motorized fader unit on which he could change the levels of sonic elements like voice, baseline tracks, interactive tracks, and reverb. He also manually triggered the different sections. Bruce wanted to respond to what he was "seeing on the floor, what the motions were, the sound in the hall, the audience … There were some interactive dynamics."

The size and circular look of the stage as well as the color of the lights inspired the composition. The former led to the repetitive, cyclical idea of the interactive sound loop and the mixing of the tracks. The latter inspired Bruce to play around with effects that influenced the audio's frequency spectrum. Furthermore, Bruce looked to the gestures that I performed for rhythm and timbre. He described the sonic structure of the interactive sounds as “fairly onomatopoeic.”

Having a semester of deliberation was beneficial to the form and quality of the piece. “The success of the piece is a huge function of the weekly meetings and [artistic] offerings,” Bruce said. Together, the team was able to distill the piece to its essence by figuring out what mattered and what did not—a crucial consideration in any kind of production. Time also enhanced the collaboration by allowing each artist to review and incorporate the other’s work. Bruce explained:

> It wasn’t like I went and created a giant electroacoustic piece that you never heard before and said, ‘Alright, here is all 12 minutes of it, go’. We did it in sections with quite a lot of care. The idea of the poem being narrated led to the decision to have the American Sign Language—which was stunningly beautiful—that I don’t think would have happened without just thinking about it. (personal communication, May 24, 2016).

Despite benefiting from ample development time, we did not have a practice room, which meant that we could not collaboratively experiment with mapping the technology to audio and visuals. “We desperately need a place to create these kind of pieces,” he said. Not having a rehearsal space “made us all very careful. We brought into the theater what we knew would work.” The solution to not being able to experiment together was to share recordings of the interactive data. Each collaborator had to explore the mapping individually using prerecorded EEG data, and then combine the multimedia elements during rehearsals. When we finally got together to rehearse on stage, the trial data proved to be very close to the live data.

Bruce wanted to create a surround sound experience, but the venue proved technically inadequate. The major limitations included the lack of directionality of the speakers, the small size of the venue, and lack of technical support. In an email sent to the creative team during the development phase, Bruce explained his reasons for not using the surround sound: "After some consideration, I have decided to drop the five-channel audio and go with the stereo tracks you have heard. For one, there is no direct artistic motivation at this point—I'm not really doing any localized sound processing. For another, it is very complicated to mount that whole system in the Brockett with nearly zero help from [venue] staff." Furthermore, Bruce "had enough to do with the I/O, processes, level mixing." During the performance, the stereo system "worked just fine."

Overall, Bruce said he was happy with the mappings. He thought the "gesture-to-sound and EEG-to-sound connections worked as well as could be expected." His only constructive comment was "we could have refined it a bit more."

ANALYSIS

This chapter describes how technology, transparency, and collaboration influenced the mapping scheme and production of *Ad Mortuos*. The first section explores issues with the network and reliability as well as concerns about how the technology was presented. The demand for freedom of movement pitted against the limitations of the sensors, the networking protocols, and our technical skills gave rise to a set of problems that were tackled in different ways. The middle section regards transparency—we will see different opinions among the creative team that range from complete transparency to none at all. Those interested in enhancing understanding were found using elegance, congruence, and accentuation. The last section focuses on how the collaboration between the creative and the production team unfolded with a critical look at the role of time. It highlights the importance of specifying early on what the technology can do and what is needed on stage to concretize the show.

Figure 4.12 provides an overview of how the different technical and aesthetic elements influenced each other. The orange boxes were addressed in the previous chapter; the green boxes will be analyzed in this chapter. Technological capabilities were specified early on during the conceptualization of the project. The freedom of movement required for the show narrowed the quality and quantity of information that we could extract from the brain. The brain/movement data was used to generate audiovisual content and informed the choreography for the dancers. These artistic elements were also impacted by our concern over how to communicate the interactive nature of the show to the audience. Technical and performative needs, namely time, resources, and production, were addressed with some challenges when the venue was engaged.

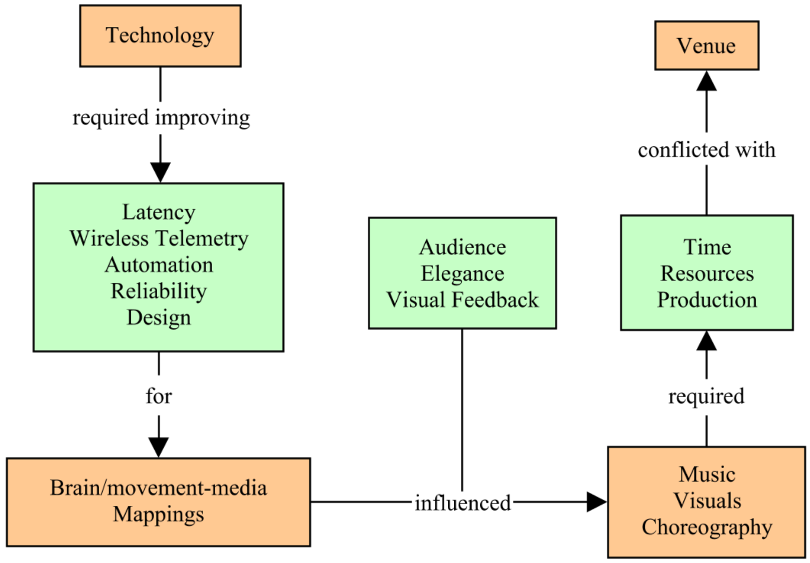

Figure 4.12: Technological, performative, and collaborative elements and their relationship to the mapping scheme.

Technology

The technological development of *Ad Mortuos* was mainly focused on enabling the transfer of information among the network devices and the improvement of the EEG's reliability. The nature of the hardware and software used for brainwave analysis, together with an inefficient network scheme, produced unusually high latency. However, the interaction designed for this show and the low priority placed on overt transparency made this delay inconsequential. This section will also highlight some reliability issues that were experienced due to electrodes in the EEG and network structure.

Latency. The latency resulting from the technology did not have an impact on the show because the audiovisual content did not need to respond instantly to my movement and brain data. All three collaborators agreed that the delay was not noticeable.

> The movements were slow and soft, so even if there was latency, it wouldn't have been noticeable. - Rodrigo Carvalho

The asynchronous relationship between the interactive elements was very convenient, since the time it took to for my input data to travel through the system probably was too slow for fast interactions. For example, the iPhone's gyroscopic data had to travel to the router, to the tablet, back to the router, and then to the media computers. Furthermore, my brain's electrical activity required several seconds to change when I closed my eyes. So although there was a lot of latency in the input due to biological or technical reasons, the interaction Bruce and Rodrigo designed was slow enough to accommodate that delay.

Wireless Telemetry. Freedom of movement was a priority since I was required to move around the stage during the show. All the technologies were battery-powered and capable of wireless communication. The iPhone sent the gyroscopic data, the EEG brainwave activity, and the tablet collected all the data and sent it to the media computers. The iPhone and tablet were embedded into the costume to make them portable; pockets held both the devices in place while I moved.

Neither Max/MSP, GyrOSC, nor MindYourOSC could send the data through multicasting, which meant that the tablet had to collect and resend all the data. This increased latency and the load on the network, because each message was sent twice from Max/MSP. This was particularly problematic for GyrOSC, since it required a router to connect to the tablet and there was none in the dressing room. I could only get the system running on stage. Despite my efforts to make the software network more efficient, I was not able to overcome the need to redistribute the sensor data though the tablet.

Automation. Activating the devices was a laborious task. Before each performance, I had to set-up the EEG hardware and software. After each performance, I had to charge the devices. In chronological order, my responsibilities were:

In the dressing room

- Getting into costume
- Applying saline solution on the EEG's electrodes to increase conductivity
- Putting on the EEG, tablet and iPhone
- Launching Emotiv, MindYourOSC, the UDP broadcaster, and data visualization on the tablet
- Launching GyrOSC on the iPhone
- Checking if all the electrodes were applied correctly on the scalp

In the back stage

- Connecting to the Wi-Fi network
- Checking if all the data was transmitting on the UDP broadcaster

Table 4.3: The author's hardware task list for *Ad Mortuos*.

Ideally, the applications would have launched automatically when the iPhone or tablet were booted and the wireless network would have connected instantly. It would have also been advantageous to have someone aside from me to place the devices in the costume. I was responsible for activating all the interactive technology before the show, and that concern distracted me from the show.

The mappings between the performer and visuals were completely automatic, while the interactive music was curated in real time by Bruce. He said he would mix in the interactive parts when he saw me perform a gesture:

> I had an interface that was reading the dataset from your EEG transmissions according to the parameters we had chosen—I was receiving those in real time and monitoring those levels and using thresholds of those levels to initiate other sounds from the Max patch. (personal communication, May 24, 2016).

Rodrigo had simply to fade in his visuals at the beginning and fade out at the end. Both Bruce and Rodrigo had the ability to manually override my input. This came in handy when we had a network failure during a show and they had to manually control the data as if my sensors were doing it.

Reliability. Wi-Fi connection was a problem in all the interactive technologies used for *Ad Mortuos*. The Wi-Fi's authentication procedure had the potential to produce drops or errors in the connection. Moreover, both the iPhone and tablet required signing in to the private network each time I went to the backstage area from the dressing room. In some cases, the tablet's wireless adapter would fail, requiring a reboot. The complexity of transferring data over the Wi-Fi network reduced reliability; a simpler radio broadcast network would have been better.

> "You want a computer that always detects that movement or interaction without having doubts, jumps, noise, or false positives." - Rodrigo Carvalho

Other issues were a result of the sensors. The EEG's electrodes routinely suffered from poor connection. In fact, a few interactive modes were dropped early on due to this. Aside from facial expression and brainwave pattern matching, Emotiv could also detect affective information about the user, such as meditation, excitement, and concentration levels. The measurement of affective levels required an excellent electrode connection, because otherwise the software would stop reporting the values. This quality could not be ensured on stage, because my movement displaced the electrodes' position, and over time the wet felt sensors would dry up and lose conductivity. The electrode connectivity problem produced false positive eye blink triggers that forced Rodrigo to abandon using them in his mapping schema.

Design. All the technology that was visible during the show was presented in a way that added to the ambience of the show. The tablet on my chest displayed an animation that simulated the data from the EEG. It featured a scrolling black-and-white histogram to signal that the performer had some connection with digital data. Kelsey Vidic designed the costume to look serious and sinister.

Transparency

As the participants made clear, it was challenging to demonstrate the EEG data to the audience. Rodrigo said, “because the data comes from an EEG, there isn’t something physical that the audience can see happening. If [a piece] uses movement, the audience can see the performer moving and creating something.”

Each member decided to portray the EEG data uniquely. Their approach was linked to the degree of willingness to make the interaction apparent to the audience. Coherence and elegance were pointed out as important techniques with an influence on transparency.

Audience. Bruce, Rodrigo, and I were more concerned with interaction transparency than Yacov. Rodrigo, and I were interested in complete transparence; we wanted the audience to know as much as possible what input was connected to what output. Bruce was satisfied with giving the audience a general feeling of reactivity. Both Bruce and Rodrigo remarked a preference for performances with an obvious interaction.

“I hate going to a piece where if you haven’t read the program notes you have no clue what’s going on,” Bruce said. Rodrigo echoed this sentiment when he said: “I’ve seen many EEG shows were you are always in doubt on whether the [audiovisuals] were prerendered or seriously happening in real-time. Transparency was an important concern for us." With the exception of Yacov, most of the team showed interest in having the audience understand that the piece was interactive.

Yacov was markedly more interested in the aesthetic outcome of the piece and not on making it clear to the audience the technological means to achieve that outcome. He explained

> The piece has to stand alone with the technology. Everything you present on stage is one whole coherent venture. The dancers, technology, wearer, activators, and computers are one. It's really a new way to making work…. If you use technology and there is no raison d'être for the work, then technology is not important for the people that are watching. It's a successful work by the way it comes together.

Coherence was an important aspect to the action-sound mapping. Bruce described the interactive sound from the EEG as "almost onomatopoeic." Bruce was careful to choose sound effects that matched naturally with the actions and data created by me.

> There was no point having the EEG system without it having some direct, audible, visual impact on the piece. In other words, I could have made the [eye] blink do anything, but I chose to have the blink do something more or less like the gesture: I called that sound 'boink.'

To me, the main goal of interactive transparency is to give more meaning to the performers' actions on stage. Hand movements, or meditative moments, are seen in a different light when the audience knows they are responsible for altering the multimedia content. Control of an audiovisual effect requires talent, and it is in the interest of the audience and performer for this talent to be made clear. I suggested starting out the piece with a quiet backing track and limited movement by the dancers to help focus the audience's attention to me playing with the audiovisual content.

Despite our concern with transparency, the audience probably did not completely understand how and what was interactive. However, Bruce stated that we were successful at giving a "sense that [Yago] was there for a reason. Did they know [Yago] was wearing an EEG? Probably not. Did they know that the sound and light was doing something relative to what [he] was doing? Yes, they were."

Elegance. Bruce, Yacov, and I acknowledged that an elegant interactive dialog that was simple, parsimonious, and direct would be the most effective way to communicate cause and effect to an audience.

Bruce wanted the beginning of *Ad Mortuos* to highlight the interaction. So he composed a slow, uneventful backing track to emphasize the interactive sounds. The music in the interactive section was lean, simple, and undramatic; it was meant to "just float, so that what happened visually with the performers would be immediately apparent as an acoustical event and not lost in some giant thing," he said. One day, during rehearsals, Bruce said:

> We are very excited about the EEG part – the direct influence on the audio and visuals will be made very clear to the audience. The base track will be much quieter than in the rehearsal, so that your EEG signals are evident. Let's make sure that the first 90 minutes

> or so of the EEG make really evident you are sending data and influencing the audiovisuals, and then add some brief solos.

The main strategies Rodrigo employs to increase transparency are elegance and consistency. "I'm concerned with making the interaction simple and stable, so that the audience clearly understands from the beginning what is the rule," he said. Elegance and consistency relates to both the mapping scheme and the gesture: the interaction-to-audiovisuals mapping should be clear enough to be identifiable, and the tracking algorithm uncomplicated and reliable. It cannot have "noise, jitter, or false positive [triggers]." Moreover, it's important to have a big and interesting audiovisual output. Rodrigo used as an example gyroscope mapping:

> The rotation of the body is a simple and effective thing to detect and it always works. There are no errors with that—you're never going to not have the visuals rotate with the body. And it has quite a big, effective audiovisual effect that the audience can immediately understand.

Visual Feedback. Yacov and I came up with gestures to accompany the facial and brainwave actions to strengthen the connection between me and the interactive media. The gestures consisted of raising my hands during the brainwave control and placing my hand over my head when doing the facial expressions. Rodrigo recalled that "we thought it important to give physical feedback, a physical action, to each interaction to help the public understand the interaction."

Most of the time these actions were performed one at a time to clarify the input-to-output correlation. Bruce mentioned how "the audience, even if they couldn't decode it technically, could see that a gesture did something and it wasn't buried in a complex mix of other stuff…. There was almost a 1-to-1 mapping of gesture to sound design."

Yacov's choreography emphasized the interactivity by restraining the activity of the dancers and having them mimic my actions. The first half of the piece focused on me so I could clearly expose all the mappings to the audience. During this moment, the dancer's role was to mimic my moves for emphasis. Only later in the piece did they start dancing again.

Collaboration

Ad Mortuos benefited from a long conceptualization phase, but suffered from too few rehearsals and a serious conflict with the production team. This section unveils the crucial role that technical specifications have on the production side.

The nature of these interactive shows is sometimes at odds with the venue's production timeline, leading to misallocation of resources. We will see how some of the reliability issues mentioned earlier could have been detected and solved if we had more experts and more time in the rehearsal space. A shared creative vision and test data helped attenuate these issues.

Production. The projections were initially planned to cover the cyclorama as well as the floor, but serious disagreements and animosity between the creative and production teams prevented us from sharing the front projectors. In fact, we could have lost all the projectors had we not hired an external rigger to set them up. As shown in *Figure 4.13*, these alternative surfaces could have possibly allowed more interactive possibilities. The conflict arose due to incompatible procedures and poor communication.

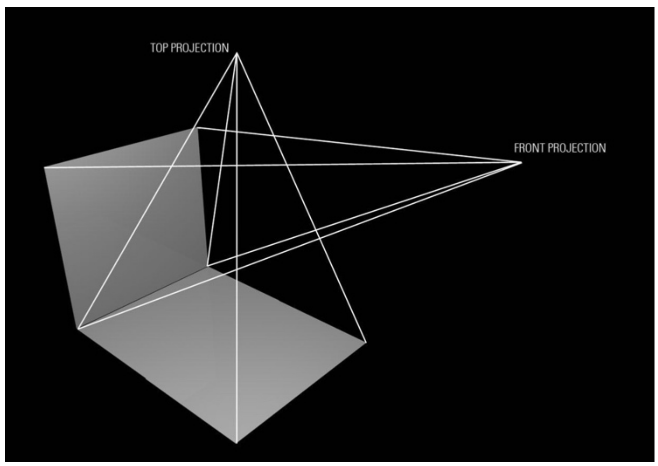

Figure 4.13: Front and top projection plan.

The timeline for delivering the technical specifications was incompatible because of the difference in each party's procedures. It is the nature of interactive works to tread through uncertain paths and work with experimental technology; technical requirements emerge during the development process and are only mature enough to specify after considerable time. Production teams, on the other hand, have a rigid timeline bound to the needs of the staff and talent who require information at an early stage to move forward, or conflicts may arise with the allocation of equipment, space, and technicians.

The production team complained that our piece suffered from "insufficient planning and failure to meet deadlines, respect for a timeline that was created to insure the success of your show." On the other hand, João Beira, who designed the non-interactive visuals for *Ad Mortuos* and helped set up the equipment, described in a message to the production team our view on the issue:

> There are a few elements missing at this point to have a final exact technical outline for the projection design, due to the experimental process and workflow that we are following for this project. There are variables that are difficult to calculate before we start rehearsals, and we can only test specific elements regarding motion-tracking strategies with video projections on the floor.

The second aspect to this problem resulted from poor communication between us and the production team regarding the level of detail for the projection specifications. A member of the production team said that our "system design looks very vague, and I can't tell whether the drawing is to scale." We did not see it as our responsibility to produce a detailed layout, nor did we have the skills to do it. This misunderstanding of responsibilities led to a breakdown in communication. In the end, we were forbidden from using the front projection, unable to set-up the Kinect system and had to hire an independent rigger for the top projector.

Yacov stated that the issues with securing the equipment was a "hindrance" to the process. Nevertheless, he said that it did not diminish the quality of the piece.

Time. The meetings held every Sunday for four months prior to the show played a crucial role in defining the piece's structure and artistic vision.

> "We had a lot of time to think…. The weekly meetings impacted enormously the form, reasoning, and strength of the piece." - Bruce Pennycook

Both Bruce and Yacov, who attended all the meetings, linked the success of the show to the collective brainstorming sessions. "The reason the piece was so successful was because we had a full semester to talk about the concept and develop the ideas that everyone was invested in," Yacov said.

The structure had to take into account all the creative and technical elements pursued by each member. For example, these elements included the poem, spoken word, number of dancers, movement in space, EEG, and projections. In the meetings, we would figure out if, when, and how these elements would work together. Aside from helping formulate the form of the piece, the meetings also helped crystalize the ambience and mood that we were trying to achieve. As Bruce said: "All the design ideas had time to percolate and be reconsidered weekly until we were quite comfortable with them."

Not having a space to permanently set up and play with the interactive system and the performers reduced experimentation. Bruce, Rodrigo, and I experienced this throughout the development process. Yacov, on the other hand, did not because he was not as involved with the technology. The lack of time invested in experimentation felt more pronounced on the creative end than on the technical challenges. Individually, we could work on the development of the sensors, software, and network, but only together could we have discovered “new things and see what comes out of it,” Rodrigo said.

Even in the two weeks we had of rehearsals before the show, we could not leave all of the equipment there or use it without the venue production staff. The limited time we had to collaborate together on stage made us more cautious, as we only brought into the interactive system what we knew would work.

The negative consequences of the lack of time collaborating on stage were attenuated by the creative meetings and the use of prerecorded data. The lengthy meetings each week leading up to the show helped create new ideas and solidify accepted ones, so that although we did not work together in creating content, we had a similar vision to realize. Yacov emphasized this point when he said how in the meetings “we had the chance to develop the how, what and where. That’s great ammunition to arrive on stage.”

A technical solution Rodrigo and I came up with to help the collaboration was to create a small app that simulated the EEG or visuals so that we could individually experiment with it. Both Bruce and Rodrigo could use my EEG player and feed prerecorded brainwave and facial expression data to their audio or visual application. Rodrigo’s visual app Gyromove provided a simulation of the projections in the computer.

Resources. Specifying early on the capabilities of the interactive technologies facilitated the process of creating the artwork. Bruce and Rodrigo knew in advance what data was going to be fed into their music and visual software. Yacov also was informed of the constraints that the EEG placed on my movements. Laying out the opportunities and constraints afforded by the technology enabled the team to move forward without having to backtrack in response to technical difficulties.

> You should be very clear as to what it is that you want to achieve and associate yourself with artists that have these expertise. And you should know that you can never make your work on your own. - Yacov Sharir

Some of the issues related to reliability could have been solved if we had more experts working on the project. Incorporating the iPhone into the show added complexity that could have been avoided if we had someone dedicated to coding the network data transfer system. Unfortunately, the budget for the show did not allow for hiring a programmer. Further network issues were experienced with the iPhone and tablet's wireless adapter. Lastly, the wireless microphone lost connection during one of the shows. Having someone to troubleshoot problems during rehearsals would have prevented these failures.

Chapter 5: Discussion

This chapter discusses the technical and organization implications of the two case studies. The first section presents a series of network and usability recommendations for researchers aiming to create interactive multimedia performances.

Firstly, I propose that user experience is tied to the specifications of four areas that control information flow—sensors, network, mappings, and interaction modality. Secondly, interaction transparency is shown to be result of our perceptions of synchronicity and elegance, and dependent on the feedback needs of the user. Lastly, these same feedback needs point to importance of visual displays as a means to improve instrument playability and variety. Optimizing these areas improves the functionality of virtual instruments to performer and audiences.

The last section considers the impact of technology-driven shows on technology-based companies and universities. I outline how co-developed shows grant companies benefits such as the discovery new product applications, improved usability and stability, increased productivity, and promotion of a market-oriented vision. Universities, on the other hand, have an opportunity to commercialize their projects in return for skills, technology, funding, and distribution from the private sector. Ultimately, technology, research, and the arts are processes intimately bound up in the imagining and generation of new ideas, products, and ways of interpreting the world. Live shows offer a convergence point for these processes that in turn create wealth and knowledge.

TECHNICAL RECOMMENDATIONS

Network Architecture

The wide range of technical and interactive approaches featured in *Curie* and *Ad Mortuos* provided a wealth of information with respect to their implications on network architecture. The main goal of a network is moving data reliably and quickly. Wired connections and broadcast protocols are more reliable than wireless authentication protocols. The perception of fast or slow audiovisual response time was shaped by the physical properties of the interactive system and user expectations. All in all, experienced latency resulted from three technical and one performative factor: starting with the sensor, then the network structure, leading to the mapping algorithms, and ending with the performer's interaction modality. *Figure 5.1* illustrates the four areas that contribute to latency and signal flow.

A sensor's sample rate, algorithms, and the nature of physical event it is monitoring determines detection latency. The first delay is contained in the execution of an action such as hand movement, eye blinking, or brainwaves. The sample of rate of the measuring sensor specifies how often this action is polled. Further processing of the data to extract high-level information will require more time. Inertial and electrical sensors output data at high sample rates below 7 ms while off-the-shelf camera-based sensors tend to take at least 17 ms. After the raw data is captured detection algorithms that use machine learning (ex. skeleton tracking and brainwave pattern) require more processing time than those based on thresholds (ex. hit gestures). The physical events being captured can be measured in different ways. For example, the two inertial sensors in the *Curie* wristbands had practically the same sample rate, yet peaks in rotation were detected earlier with the gyroscope, due to the nature of the gesture, then peak changes in speed with the accelerometer. As another example, Emotiv's EEG had a high resolution rate but detection was slow due to the pattern-matching algorithm and the time it takes to generate a particular brainwave pattern.

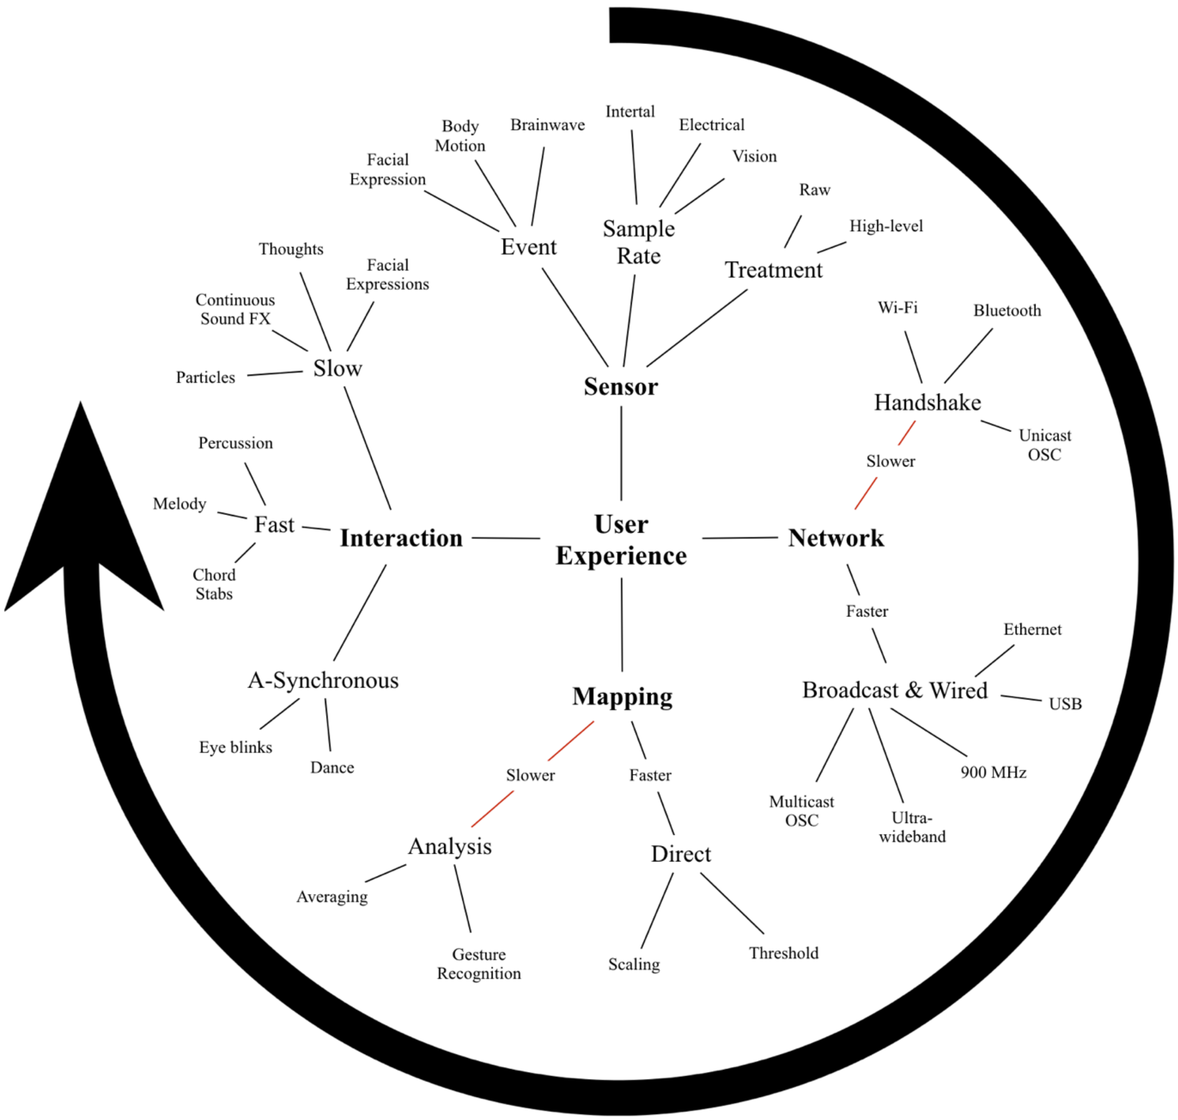

Figure 5.1: The four areas that contribute to a user's perceived interaction delay. The arrow indicates the signal flow through the interactive system.

Detection time in *Curie* was reduced by increasing the sensor's sampling rate, substituting the sensor or optimizing the gesture. For instance, the sample rate of the ultra-wideband position tracking was increased from 50 Hz to 200 Hz. The wristband's accelerometer was substituted for the gyroscope while the hit gesture was changed to optimize the latter's response time. Some devices had sample rate ceilings, such as the shutter speed of a camera or the ultra-wideband airtime, and some actions like an eye blink required multiple samples to detect.

After capturing the action, sensors need to transmit it across some kind of medium to the visual and audio software. Findings indicate that information travels fastest through wired connections and that broadcast and multicast were more efficient than handshake[7] connections. Communications requiring a handshake experience slower transmission speeds and demanded a time-consuming pairing procedure. For example, switching from the Bluetooth protocol to 900 MHz and ultra-wideband produced a marked decrease in latency.

The same improvements would have probably been noticed in *Ad Mortuos* if we had switched from Wi-Fi to a broadcast technology. However, at that point in time, we

---

[7] I use here the term broadcast in reference to any communication technology that does not require an agreement between a sender and receiver, and handshake to those requiring identification or authentication. Using this definition, 900 MHz and ultra-wideband were broadcast, while Wi-Fi and Bluetooth were handshake.

did not know that these broadcast solutions existed nor did we have a network expert on the team.

Higher data transfer speeds are achieved with USB and Ethernet—the former connected the devices to computers, while the latter interconnected the computer. Lastly, sending OSC messages containing sensor data over UDP unicast were distributed less efficiently than multicast.

Mapping schemes that require averaging or pattern matching have a slower response than direct action-to-media. This is because averaging (as with the ultra-wideband positioning system) uses multiple samples and pattern matching (as with the glove gesture identification) requires more calculations. Removing the analysis reduced latency. For example, Intel's idea of having the wristband identify a set of gestures was too slow and was reduced to one gesture.

User expectations regarding the responsive media have an impact on perceived latency. These expectations are tied to the attack and release time of an action or media. Actions with slow attacks like facial expression and thoughts can cope with slow media response. Fast gestures like percussive hand hits will most likely require a fast response.

Lastly, some interactions are not expected to synchronize with external stimuli and thus have loose requirements for latency. The eye blinks in *Ad Mortuos* and the dancer's movements in *Curie* fell under this category. Users expect immediate response when they perform an action with a fast attack that meant to start a media event with another fast attack. Conversely, slow actions or slow media tolerate higher levels of latency.

Interaction Transparency

The interactive contract between a performer and media helps establish a repertoire and define expectations. And audiences build action-sound-image associations after repeated exposure. If this contract is broken, the credibility of the interaction is compromised: when Peter unexpectedly disconnected the particle interaction from the dancers, Reema complained that her movements lost some of their significance because they were conceived with the visuals in mind. When a performer switches to a new instrument, a change in the visual display can indicate that something new is being controlled. In this case, the performer-audio-visual associations are clearly contained within a multi-sensorial instrument. Consistency is also dependent on reliability; data redundancy and network upgrades in *Curie* improved the cohesion of the gesture to media pairings.

Synchronicity is perhaps the most important contribution to transparency. If an action and a media element happen at the same time, a "spontaneous and irresistible weld" will be produced between them (Chion, Gorbman, & Murch, 1994, p. 63). This weld allows apparently unrelated pairs of audiovisual material to come together. Latency was dramatically reduced in *Curie* to allow exact synchronicity for the hit gestures.

When engaging with interactive media, designers and performers said that elegant solutions proved to be the most transparent. This meant carefully limiting live gestures, visuals, and sounds to those that were interactive. Simple interactive environments draw attention to the connection between the performer and interactive media. In particular, interactive sounds need to have a very pronounced role in the composition if they are going to be noticed.

Elegance is particularly important during the introduction of an interactive instrument. When multiple performers are involved, introducing them one by one ensures that the audience will associate the sound to the performer correctly. Elegant mappings (i.e. one action produces one audiovisual effect) are more easily understood than complex or abstract ones. These results are supported by literature: audiovisual complexity can make it hard for participants to perceptually bind the interactive elements (Callear, 2012; Coulter, 2010); while simple stimuli pairs inevitable lead to a connection (Vatakis & Spence, 2007; Welch, 1999).

The case studies identified individuals in the team who were not concerned with making the mapping evident. Peter and Yacov saw it as a continuum that they explored for the sake of the artwork—for which the results justify the transparency. Preference for transparency may be related to the excitement some of the performers felt for the technology and interaction. They felt they were doing something unique and wanted to share it with the audience. Moreover, transparency highlights the capabilities of a technology. Musicians interested in exhibiting virtuosity require strong and direct associations between their actions and the audiovisual output.

Aside from the addressing the needs of the user, the physical design of the instrument can add to the ambience of the piece and draw attention to the technology. Aesthetic design considerations follow the practical shape of the device. If the design does not correctly address usability features, like the poor visual feedback of the RealSense instrument in *Curie*, it risks low user involvement.

Visual Feedback

If the performer cannot see the reactive visuals, the interaction stops. Reema Bounajem noticed this when she danced with the visuals behind her: “[On the 360-degree stage], we couldn’t see [the visuals] so we just had to dance very specifically in the way we had imagined it.” I felt the same way when I performed in *Ad Mortuos* and could not see the projections on the floor.

Without the visuals, the performer’s actions become scripted and the active dialog with the media is lost. Not all venues and visual configurations permit a display that can be visible for both the audience and performers. Scrims can be cost-effective solutions but require a powerful projector.

Visual feedback is important for instruments involving continuous control or event selection. Adam, Mitch, and I found it useful to see where we were in the range of motion. One of the advantages of visual feedback is that it can be placed next to the performer without overlapping other images. Furthermore, visual feedback for instrument control does not need to be visible to the audience. Instruments performed at the New Interfaces for Musical Expression conference often use a laptop screen for feedback. However, most digital instrument makers miss out on the opportunity of exhibiting the parameter space of their instrument to the audience. Displays add another dimension to a show and enhance transparency.

Auditory feedback is mandatory for musical interaction, but it may be insufficient as a means of communication. Because audio signals overlap with other instruments, they tend to get muddled. Another issue is that functional auditory feedback may not suit the composition. Despite the advantages afforded by the audiovisual feedback, few digital instrument makers design for both modalities. This is most likely due to the technical and material challenges of building a visual display. The advantages are not only performative but also perceptual—audiovisual instruments are more easily understood by audience members.

IMPLICATIONS TO INDUSTRY-UNIVERSITY PARTNERSHIPS

One of the surprising implications of this research came from the informal industry-university partnership that was created in *Curie*. When compared with the academic piece *Ad Mortuos*, we see a stark difference in the way technology, aesthetics, and collaboration unfolded. My findings contribute to a growing body of work that supports the claim that industry-university partnerships can provide numerous benefits to the collaborators. Beyond supporting literature, these case studies present new evidence that artistic fields can participate in research and development, a role normally reserved to science, technology, engineering and mathematics disciplines.

Industry-university partnerships have become a powerful transformative force in the world today. Over the last decades, a growing number of companies and universities have been working together to replace insulated R&D environments with distributed platforms that promote rich transfers of knowledge, resources, and technology

(Chesbrough & Brunswicker, 2013; Science | Business Innovation Board AISBL, 2012; Mora Valentín, 2000). With business being increasingly pressed to reinvent themselves and evolve their strategies before they are disrupted from the outside or by their competitors, multidisciplinary co-development partnerships have become a key enabler of the changes they need to make (Shekar, 2014).

Quite a lot has been written on how these industry-university partnerships can fuel innovation, but the particular contributions of artistic disciplines remain unstudied. Innovation ecosystems, the thinking goes, are reserved for science, technology, engineering, and mathematics (STEM) disciplines (Public and Corporate Economic Consultants, 2012; Abreu, Grinevich, Hughes, & Kitson, 2011). This restrictive view has prompted universities to overlook performing arts' true value to the private sector, and has been an important factor in reducing policy attention and funding for commercialization (Benneworth & Jongbloed, 2010). Artists and tech companies, the thinking goes, have different, incompatible priorities that prevent them from working together in towards a common goal (England, 2013).

This research presents new evidence that challenges this view. Technology-driven industry-university partnerships can encompass the performing arts, with the goal of improving new product development for companies and increasing skill and technological resources for universities. It turns out that with regard to products and art, both tech companies and tech-driven arts hope to maximize usability, innovation, and impact. Drawing from my experience and analysis of two collaborative performing arts projects within the academic and corporate environment, I demonstrate that a strategic partnership between tech companies and artistic research communities makes economic, technological, and intellectual sense. More specifically, they lead to accelerated development, improved product quality, usability and exposure, expanded team competency and skills, and an increase in project resources and funding.

In the following sections I point out the collaborative benefits to the industry and academic partners. The last section offers a framework—Live Product Development—for the successful collaboration between industry and university within the performing arts. This framework is aimed primarily at companies developing technologies that enable people to do more things, and at researchers creating technology-driven interactive performances. What is at stake is a means to empower students and businesses with competitive qualities that will allow them to contribute and lead transformative changes in a world increasingly characterized by technological disruption (the emergence of technologies that transform life) and evolving customer needs.

Impact on Industry

As strategic discussions routinely focus on how business can evolve and capitalize on new innovations, it is important to recognize the enhanced role that artistic communities should play in the creative use of disruptive technologies. The finding reviewed in their research suggest that artist and researchers can promote innovation and foster creativity during the new product development phase. This section highlights various ways in which creative, scientific, and commercial partnerships can contribute to a company's economic and cultural potential by increasing a product's competitiveness. More precisely, we will see how implementing prototype technologies in live events can lead to the discovery new product applications, improve usability and stability, increase productivity, and promote a market-oriented vision.

As the analysis sections have shown, experts play a vital role in enabling the fulfillment of the creative vision. Consider the challenges presented in each case study—solving them required knowledge in machine learning, network protocols, signal processing, radio transmission, computer vision, and programming languages. Although companies employ experts in various fields, they may not have all the skills needed to address the wide range of challenges that arise during the development and testing of a technology. By connecting with music and visualization research communities—which includes the fields of performing art, fine arts, computer graphics, audio engineering, signal processing, HCI, and computer science to name a few (Marquez-Borbon & Stapleton, 2015)—companies can transfer scientific, technical, and creative knowledge into their process. In our case, we saw Intel absorbing various methods and technologies such as the Open Sound Control protocol, ultra-wideband radio, gesture repertoire, and data visualizer. Due to the multidisciplinary nature of product development, people with different backgrounds contribute to its success, some of which can be found in the scientific community.

Another advantage of collaborating with artists, according to Lakshman Krishnamurthy, a fellow at Intel's New Devices Group, is that they can help "discover new applications of technology" because "artists push the limits of technology and use it in unexpected ways." Initial observations suggest that there may be a link between creativity and the discovery of new ways that technology enables processes to be done differently or better—possibly leading business to explore a completely new direction. For instance, our work with Intel's RealSense camera showed that tracking the hands and limbs of a user could enable gestural control of user interface elements. Likewise, the collaborative efforts between performers, artists, and engineers led to the musical application of Intel's Curie chip as a wearable instrument. This device is now being shared among students in at the Center for Arts and Entertainment Technologies in Austin, Texas. In general, therefore, is seems that the use cases envisioned by the artists can suggest new markets and ways for consumers to use a product.

Equally important was the feedback the artists gave to the engineers regarding the usability of a product. It appears that no matter how well thought out a product is, without significant testing, users will find applications that are not fully compatible with the device's features. Within the performing arts context, the most productive way to optimize these features is to give artists significant creative freedom while also requiring them to learn how the necessary devices work. Such a process guides artists and engineers with surprising effectiveness towards use-oriented development.

Nowhere was this more evident in than the fruitful collaboration between the performers and the Curie wristbands. Here we saw user feedback lead to more responsive software and hardware, represented by the improvements in the gesture-detection algorithms and the ergonomics of the instrument. Understanding how artists use a device provides supporting knowledge on both the machine and the human side in connection with the design, evaluation, and implementation of interactive systems.

Yet our most noticeable contributions to Intel was to their product development cycle. According to Lakshman, the benefits were twofold: the collaboration "accelerated Curie's hardware and software development" and attracted "lots of resources to the project."

The first claim is supported by Stevens et al. (1999): an analysis of 267 early-stage product development projects indicated that the creativity of those evaluating products had a significant impact on speed and productivity. I would argue that acceleration is most likely due to the added emphasis on four actions implicated in the collaboration: setting targets, exploring new techniques, creating prototypes, and giving feedback. Within this context, these actions happen at much faster iterations than in the traditional linear model of product development, and have a combined effect of focusing engineering efforts towards practical concerns.

This can be illustrated briefly with the Curie Module's RealSense Cursor Mode and skeleton-tracking features, which were developed with the assistance of the creative team and made tremendous headway in both the quality of the tracking and improvement of the software. Furthermore, after noticing the importance of the technological and promotional contributions to the show, Intel invested financial and human resources to strengthen the engineering and marketing outcomes. In sum, products get to enjoy faster development cycles when they are tested with artists, due to a boost in resources and theoretical and practical knowledge.

In addition to accelerating development, live shows also provide an opportunity to improve the quality of a product. Lakshman stated that because of the development needs of the show, his team "understood a lot better what works and what doesn't, where should we focus on, how many things we should do…. It was trial by fire."

This is due, in part, to the fact that technologies developed for concert venues must overcome multiple sources of interference: cell phones, Wi-Fi signals, ambient noise, ambient light, electrical equipment, and so on. Aside from the high performance targets, live concerts insist on streamlining the functionality of a device so that it can fit into the complex performative environment, which tends to favor fast and efficient execution. For instance, the performance needs placed on the Curie-based wristbands motivated the construction of an extremely fast wireless network, precise localization, fail-safe data transmission, and an instant system boot sequence. It may well be that Curie would not have enjoyed such a widespread success as compared to other Intel projects had it not gone through such rigorous development and testing ahead of the Consumer Electronics Show.

Lastly, concert-oriented development can improve the sales of a product by linking it to marketing goals. The integration of marketing and R&D is widely recognized as a key success factor in new product development in both theoretical literature and the *Curie* case study (Brown & Eisenhardt, 1995; Griffin & Hauser, 1996). However, a recent study has shown that engineering and marketing have different goals and time frames which can lead to conflict and hinder integration (Homburg & Jensen, 2007). Live shows, due to their technical and performative characteristics, enhance the cooperation between these two departments by providing a range of cross-departmental activities.

Two examples from *Curie* illustrate this point. The first was the active involvement of the marketing team in the conceptualization of the dance, visuals, music, and design aspects of the show. Secondly, in recognition to the well-received live demonstration of their technology, Intel was awarded "Best Production of an Event" by the Consumer Electronics Show[8]. Given these points, we start to see how live concerts can bridge the R&D-marketing gap with the goal of identifying new, promising ideas and refining them into product concepts.

Impact on Universities

The emergence of every new technology creates abundant opportunities for research communities eager to innovate their practices. Yet some valuable student ideas are not developed due to lack of a university's resources. In order to discover and implement technologies while they are still novel, research communities need to rapidly gain the right skills and equipment by collaborating with original creators. Rather than wait for the solutions they need to appear through traditional sources, they can take an active approach by making continuous contact with developers in technology companies a core competency. Survey data presented by Lee (1996) suggest universities would do

[8] http://www.eventmarketer.com/article/2016-ex-awards-winners/

right to initiate or expand university-industry partnerships as long as the purpose is to accelerate the flow of expertise, skills, knowledge, and technology.

Similar kinds of collaboration, such as Open Innovation (OI), have become "a top priority with world‐class universities for the benefits it has been shown to provide" (Shekar, 2014, p. 2). Industry-entertainment ecosystems are evolving to enable co-created solutions to serve student needs in ways that are beyond the means of any single institution. As shown below, there is huge potential to make a lasting, transformative impact on students by allowing them to commercialize their projects in return for skills, technology, funding, and distribution from the private sector.

This research found overwhelming evidence for the notion that the discovery of new methods and solutions to problems are strongly related to the skill set of the creative team. Successful shows require different combinations of core artists, freelancers, and devices for each new challenge. Driven in large part by the fast pace of technological advancements, technical competencies of the multidisciplinary teams are constantly changing. To make better use of the artistic tools they have at their disposal, universities should form alliances with developers to help overcome challenges by tapping into resources that they don't have to own.

It is important to realize that very frequently, the sensors implicated in the case studies required major modifications to make them useful within the framework of the project. While the *Ad Mortuos* skill set was limited to the four core creative team members, *Curie*, on the other hand, was able draw skills from various sources within and outside of the corporate and academic environment and, as a result, was able to achieve a successful adoption. Had *Ad Mortuos* partnered with Emotiv, the makers of the EEG headset, perhaps we would have been able to overcome the headset's sensor and network limitations.

However, artistic researchers in the digital age do much more than tick off a checklist of technical capabilities. They know their ability to drive change and disrupt the status quo hinges on the technological tools on their disposal. Many aspects of a live show make good use of new hardware and software to conceive of and accelerate the process of producing multimedia content.

In this regard, *Ad Mortuos* and *Curie* differed in how the hardware was discovered and acquired. The *Ad Mortuos* team "sought" the technology, while in *Curie* we were "exposed" to it. Sought technologies, like the EEG headset, iPhone, and computer tablet, were confined to established commercial products, while those presented to us by Intel, such as the RealSense camera and Curie Module, introduced the next generation of sensing devices. The embedded design of the wristband's microprocessor, for instance, wirelessly detected the type and location of a performer's gesture; in effect, the tech industry provided solutions that we did not imagine existed.

Furthermore, allowing artists to keep the technology can promote sustainable innovation that goes beyond the original collaborative assignment. Intel's sponsorship of donated equipment to the Center for Arts and Entertainment Technologies following *Curie* illustrates this point clearly. In brief, it is almost certain that technology transfers from industry to artistic research groups can assist the economic development and commercialization of digitally-enhanced performances.

Besides technology, production resources also play a crucial role in enabling the execution of the creative vision. This includes rehearsal spaces, stage rental and design, audiovisual equipment, lighting, and the technical staff that operates it. Universities normally have spaces dedicated to live performances, but equipment shortages and time constraints happen frequently due to competing acts. For example, the production conflicts in *Ad Mortuos*, which happened within the academic context, prevented the implementation of one of the sensors and limited projection capabilities.

On the other hands, once the vision has been approved by the client, live corporate events can offer production resources tailored to satisfy the needs of the project. For instance, leading up to *Curie*, Intel's commitment to providing pre-show production support for the both the Consumer Electronics Show and Anaheim events allowed the artists to tap into vast pools of production resources across different domains.

More important, though, is acknowledging the tremendous impact that funding has on a project's human and technological resources. As the literature on interactive performances indicated, the trend towards media-performance integration has been steadily accelerating over the years. It is driven by increasingly sophisticated embedded systems and visual displays that are supported by complex interactions across multiple domains of knowledge in the performing arts, media design, and human-computer interaction. Students, however, are often left with few means to create complex multi-disciplinary projects other than the small grants provided by the department for equipment and labor. In addition to poor funding, they are often faced with the inability to accommodate innovative techniques that do not fit established modes of production.

Nowhere was this more evident than the limitations pointed out in the *Ad Mortuos* case study regarding the conflicts in the production timeline and implementing the technology at the venue.

Given that technology and skilled labor are universal ingredients that intersect both interactive digital content and the performing arts, we should be asking ourselves if research communities are providing sufficient funding for students to catalyze significant disruptive changes while at the same time pursuing commercial ambitions. Forming strategic partnerships with tech companies in support of their products may increase investment in current artistic projects as well as future university-industry knowledge and technology transfers. This funding will likely be most beneficial when used to initiate conversations about the performative and multimedia resources needed to potentially reshape artistic processes or fundamentally recast audience experiences.

The last point I want to mention is that corporate events also present an additional performance outlet to the academic theater. In general, conventions and trade shows provide an opportunity to raise awareness of an artist's work, capture leads from new buyers and prospects, and create or strengthen industry relationships (Sashi & Perretty, 1992). The performances at the Consumer Electronics Show and in Anaheim drew over 5,000 attendees. But exposure is not limited to live audiences; because *Curie* was featured in the Intel-sponsored show "America's Greatest Makers," it was able to expand its online footprint.

Live Product Development

Over the last decades, various mechanisms have been designed to promote, forge, and manage complex relationships between the private technology-based sectors and science, technology, engineering, and mathematics (STEM) disciplines. These collaborative strategies have to be expanded to embrace an innovation ecosystem driven by artistic research. Having outlined the major economic benefits of merging performing arts and science with a company's product development process, there is an opportunity for a specialized type of industry-university alliance. Institutions looking to catalyze these partnerships must create leadership positions to broker connections, expand online presence to increase visibility, and engage with intermediary agents responsible for producing corporate events. Nurturing these areas will build the foundations of an effective collaboration policy that matches business needs to those universities that best meet their requirements within the appropriate domain.

At the same time, businesses should scout for academic partners through a disciplined multi-step filtering process to help improve the odds of a profitable, lasting match that successfully captures the vibrancy and rigor of creative inquiry and applies it to research and development. The success of these communication pipelines will be measured by their ability to discover opportunities and may articulate how, in essence, an art and product co-development partnership can offer a competitive and innovative means for wealth and knowledge creation.

In this study, I propose one form in which these joint ventures might occur: Live Product Development. Conceptualized as a collaborative framework, Live Product

Development is a joint product-and-artwork development methodology that derives economic and scientific value through the co-creation of live entertainment. Its goal is to empower companies and researchers with the knowledge, technology, and financial means to constantly adapt, learn, create new solutions, foment change, and disrupt the status quo. *Figure 5.2* illustrates the convergence of these seemingly disparate activities into a common goal and outlines the benefits discovered in this study.

Unlocking the potential of Live Product Development involves expanding industry-university partnerships to encompass creative disciplines, creating channels for knowledge, talent, and technology transfer, and strengthening resources with which to imagine diverse new futures.

Initiating a Live Product Development entails creating a show that highlights the technical capabilities of a company's technology using live performers and interactive digital media. The reason for implementing real-time interactivity in the show is to ensure that the technology is pushed to the limits and to affirm its role as the enabling nexus in human-computer interaction. To this end, the company and creative team must work together through a series of engineering and artistic activities leading up the live event—while keeping in mind that the ultimate goal is it to increase the economic success of a product and expand research resources.

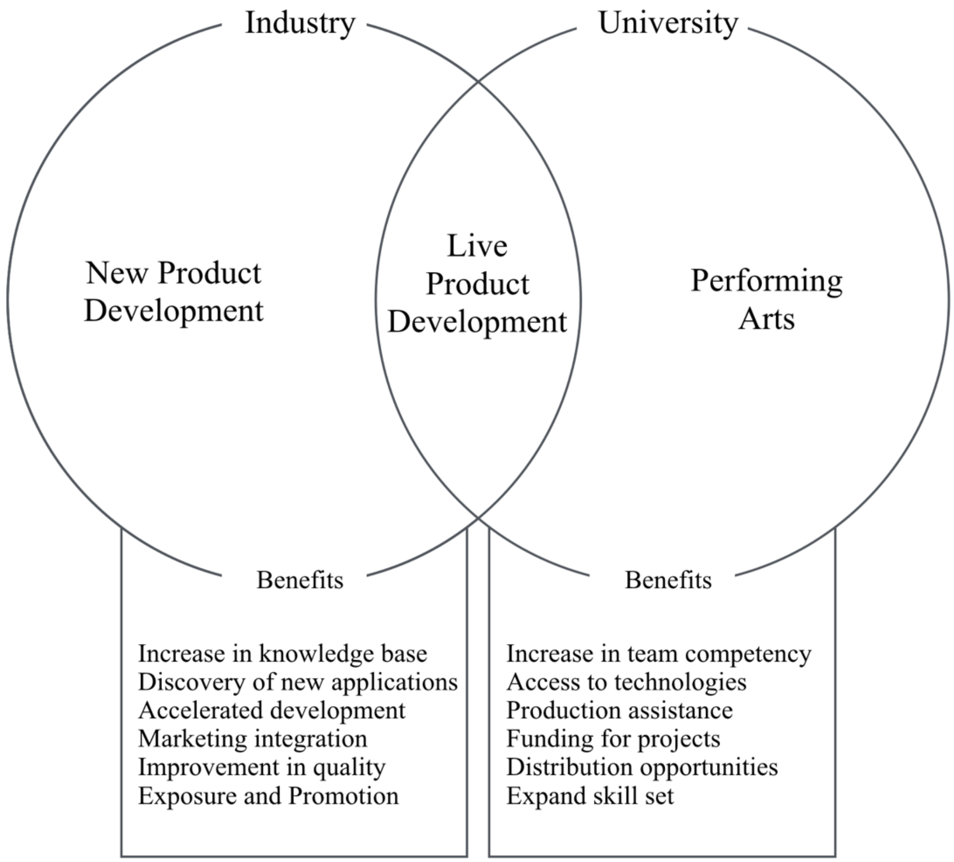

Figure 5.2: Located at the intersection of new product development and the performing arts, Live Product Development is a technology-driven alliance that foments innovation through collaboration.

Looking more closely at the two case studies, it is clear that a successful implementation of Live Product Development takes into consideration three collaborative aspects: resources, activities, and organizational structure. Firstly, the engine of innovation in this collaborative ecosystem is powered by the contributions of its partners: artistic research groups provide creativity, scientific knowledge, user testing, and artistic content; companies supply funding, technology, and production resources. Secondly, there is not a standard protocol for joint art and product development; instead, deliverables initially emerge in response to spontaneous business and performative demands before crystalizing into an action plan. Lastly, when forming the alliance, those in leadership positions need to understand that successful projects require different combinations of self-organizing, cross-functional teams consisting of creative, technical, and consultative experts who can address artistic, engineering, and production challenges.

An effective way to frame the collaborative nuances in Live Product Development is to envision it as a sophisticated product-placement strategy. The artistic team has to be prepared to cleverly embed the technology into the narrative of the artwork by integrating it as a core element of the piece, fully capturing its transformative potential in a captivating show. Businesses, on the other hand, need to be prepared to provide extensive support throughout the process while embracing the artistic initiatives that underpin the generation of new ideas, products, or ways of interpreting the world.

The next industry trade show, board meeting, or product launch date are just a few of the reasons to initiate a Live Product Development partnership. Prior to these events, this study suggests that three months of collaboration including weekly meetings and deliverables was associated with increased reliability and maturation of the technology and artwork.

It is important to note that while Live Product Development is an especially effective framework for speeding up R&D and prioritizing specific applications, it also requires an additional set of activities that can take months to complete. These include ideation, content creation, training, exploration, hardware and software building, testing, and rehearsing. Fortunately, these event-centered activities can be executed in tandem with the usual product development process and are fully compatible with agile cycles of development and experimentation. Irrespective of events, companies and universities are recommended to pursue a collaborative agreement in the early stages of new product development—when there is room for flexibility and adaptation—so that the insights and experience gathered from the show can have a chance to inform the design and functionality of the technology.

Conclusion

In this dissertation, I investigated some of the factors that influence the design of gesture and brainwave interaction in audiovisual performances. Defining these relationships is very often a top priority when creating interactive artwork. Interaction models and guidelines based on embodied music cognition theory are often too restrictive or inadequate. As evident in my two case studies, *Curie* and *Ad Mortuos*, funding agents, aesthetic trends, available technology, time, and collaborators play a much larger role in the overall direction of the piece. In this environment, there is a need for a more flexible technique to enhance the connection between the performer and interactive media.

This study has shown that perceptual cues offer an effective means to explicitly connect humans, sounds, and images on stage. This was found to be the case in both body-driven and brain-driven interactive works. The interviews in my case studies report that designers and performers paid most attention to synchronicity, elegance, and visual feedback when creating or playing with interactive content. This argument is supported by cognitive studies in audiovisual binding. Manipulating these variables can strengthen or weaken the apparent relationship between the performer and the interactive content. Ultimately, they affect the audience's overall impression of the show and the performer's behavior toward the interactive interface.

Furthermore, the case studies exposed the underlying architecture that drives user experience in human-sound-image interaction. Sensors, network, mapping, and interaction modality were found to be the four configurable pillars of an interactive system. While elegance and visual feedback are often seen as aesthetic decisions, enhancing synchronicity means dedicating a considerable amount of time to reducing latency. Raw inertial data, broadcast communication protocols, and direct mappings were found to be the fastest means to connect a performer's actions to interactive audiovisual content. In the interaction component, however, it appears that the speed of a performer's action and the tempo of an audiovisual reaction were the deciding factors in how much latency was tolerated by performers and audience alike.

In my opinion and that of my collaborators, *Curie* and *Ad Mortuos* successfully met our artistic expectations and managed the limitations imposed by the technical and collaborative environment. They are a reminder that complex large-scale artworks involve as much creative deliberation as negotiations between different stakeholders. Negotiations serve to develop a unifying vision and bring together the resources needed to execute it. There is no doubt that interactive technologies add another layer of complexity to what already an elaborate stage production. But as these new modes of audiovisual productions, collaborative structures, and network systems become commonplace, we will probably see them streamlined and incorporated into our daily practice.

Lastly, one unexpected finding that came out of the cross-case analysis was the impact of live performance on academic research, as well as the sponsoring company, Intel. By working together with artists to create a technology-driven show, Intel enjoyed accelerated development, improved quality, increased exposure, and new applications for its products. Meanwhile, I as a university student benefited from expanded resources, commercialization of research, funding for projects, and access to new technology.

Interestingly, these benefits are thought to arise only from STEM-related industry-university partnerships, and are not normally associated with the arts. This study provides new evidence that creating interactive audiovisual performances can in fact improve research and development in industry-university partnerships. What is at stake is a means to empower students and businesses with competitive qualities that will allow them to contribute and effect changes in a world increasingly characterized by technological disruption and evolving customer needs. To address this gap, I introduced Live Product Development as a preliminary framework for initiating this type of partnership.

Bibliography

Abreu, M., Grinevich, V., Hughes, A., & Kitson, M. (2011). Knowledge Exchange between Academics and the Busines, Public and Third Sectors. Swindon: Centre for Business Research, University of Cambridge and the Arts and Humanities Research Council.

Allison, B. Z., Faller, J., & Neuper, C. (2011). BCIs That Use Steady-State Visual Evoked Potentials or Slow Cortical Potentials. In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 241-249). Oxford: Oxford University Press, USA.

Ansari, K., Ansari, S., & Jaffri, S. (2014). The Atypical Creative Arts Research Methodology(s): Integrating Practice with Performance. *International Research Journal of Art & Humanities, 42*(42), 141-165.

Antle, A. N., Corness, G., & Droumeva, M. (2009). What the body knows: Exploring the benefits of embodied metaphors in hybrid physical digital environments. *Interacting with Computers, 21*(1-2), 66-75.

Auslander, P. (2008). *Liveness: Performance in a Mediatized Culture* (2 ed.). New York, NY: Taylor and Francis.

Baxter, P., & Jack, S. (2008). Qualitative Case Study Methodology: Study Design and Implementation for Novice Researchers. *The qualitative report, 13*(4), 544-559.

Benneworth, P., & Jongbloed, B. W. (2010). Who matters to universities? A stakeholder perspective on humanities, arts and social sciences valorisation. *Higher Education, 59*(5), 567-588.

Bevilacqua, F., Müller, R., & Schnell, N. (2005). *MnM: a Max/MSP mapping toolbox.* Paper presented at the New Interfaces for Musical Expression, Vancouver, BC, Canada.

Bevilacqua, F., Zamborlin, B., Sypniewski, A., Schnell, N., Guedy, F., & Rasamimanana, N. (2010). Continuous Realtime Gesture Following and Recognition. *Gesture in Embodied Communication and Human-Computer Interaction, 5934*, 73-84.

Birbaumer, N. (2011). Steady State Evoked Potentials. http://gocognitive.net/interviews/steady-state-evoked-potentials: goCognitive.

Brown, S. L., & Eisenhardt, K. M. (1995). Product Development: Past Research, Present Findings, and Future Directions. *The Academy of Management Review, 20*(2), 343-378.

Callear, S. (2012). *Audiovisual particles: parameter mapping as a framework for audiovisual composition.* (Ph.D.), Bath Spa University.

Canibal, A. L., Moura, J. M., & Pedro, M. (2013). *Câmara Neuronal: a Neuro/Visual/Audio Performance*. Paper presented at the Computation Communication Aesthetics and X, Bergamo, Italy.

Chesbrough, H., & Brunswicker, S. (2013). Managing open innovation in large firms. Fraunhofer Verlag.

Chion, M., Gorbman, C., & Murch, W. (1994). *Audio-Vision: Sound on Screen*. New York: Columbia University Press.

Clay, A., Couture, N., Nigay, L., de la Riviere, J.-B., Martin, J.-C., Courgeon, M., . . . Domengero, G. (2012, 5-8 Nov. 2012). *Interactions and Systems for Augmenting* *a Live Dance Performance.* Paper presented at the 2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-AMH).

Cooke, G. (2010). Start making sense: Live audio-visual media performance. *International Journal of Performance Arts and Digital Media, 6*(2), 193-208.

Coulter, J. (2010). Electroacoustic Music with Moving Images: the art of media pairing. *Organised Sound, 15*(1), 26-34.

Creswell, J. W. (2007). *Qualitative Inquiry & Research Design: Choosing Among Five Approaches* (Vol. 2nd). Thousand Oaks: Sage Publications.

Creswell, J. W. (2016). *30 Essential Skills for the Qualitative Researcher*. Thousand Oaks, California: SAGE.

Curran, E. A., & Stokes, M. J. (2003). Learning to control brain activity: A review of the production and control of EEG components for driving brain-computer interface (BCI) systems. *Brain and Cognition, 51*(3), 326-336.

Dallow, P. (2003). Representing creativeness: practice-based approaches to research in creative arts. *Art, Design & Communication in Higher Education, 2*(1/2), 49-66.

Donnarumma, M. (2012). Performing Proprioception and Effort in “Hypo Chrysos”, an Action Art Piece for the Xth Sense. *Biotechnological Performance Practice - Canadian Electroacoustic Community.*

Eckel, G. (2012). Embodied Generative Music. In D. Peters, G. Eckel, & A. Dorschel (Eds.), *Bodily Expression in Electronic Music: Perspectives on Reclaiming Performativity*. Hoboken: Taylor and Francis.

Arts and Humanities Research Council and Creative England. (2013). connecting and growing creative businesses through engagement with higher education institutions.

Ernest Edmonds, L. C. (2010). Relating Theory, Practice and Evaluation in Practitioner Research. *Leonardo, 43*(5).

EY. (2015). Cultural times: The first global map of cultural and creative industries.

Folgieri, R., & Zichella, M. (2012). A BCI-based application in music: Conscious playing of single notes by brainwaves. *Computer Entertainment, 10*(3), 1-10.

Gelineck, S., & Böttcher, N. (2012). *6to6Mappr: an educational tool for fast and easy mapping of input devices to musical parameters*. Paper presented at the Proceedings of the 7th Audio Mostly Conference: A Conference on Interaction with Sound, Corfu, Greece.

Grierson, M. (2008). *Composing With Brainwaves: Minimal Trial P300 Recognition As An Indication Of Subjective Preference For The Control Of A Musical Instrument.* Paper presented at the International Computer Music Conference, Belfast.

Griffin, A., & Hauser, J. R. (1996). Integrating R&D and Marketing: A Review and Analysis of the Literature. *Journal of Product Innovation Management, 13*(3), 191-215.

Hagel, J., Brown, J. S., Maar, A. d., & Wooll, M. (2016). Connect peers: Leveraging the evolving digital infrastructure to empower the edge.

He, B., Gao, S., Yuan, H., & Wolpaw, J. (2013). Brain–Computer Interfaces. In B. He (Ed.), *Neural Engineering* (pp. 87-151): Springer US.

Heap, I. (2011). Me The Machine.

Holland, S., Wilkie, K., Bouwer, A., Dalgleish, M., & Mulholland, P. (2011). Whole Body Interaction in Abstract Domains. In D. England (Ed.), *Whole Body Interaction* (pp. 19-34): Springer London.

Homburg, C., & Jensen, O. (2007). The Thought Worlds of Marketing and Sales: Which Differences Make a Difference? *Journal of Marketing, 71*(3), 124-142.

I-CubeX. (2011). Biosensing in Max. 2014, from http://cycling74.com/toolbox/biosensing-in-max/

Im, S., & Workman, J. P. (2004). Market Orientation, Creativity, and New Product Performance in High-Technology Firms. *Journal of Marketing, 68*(2), 114-132.

Jackson, M. M., & Mappus, R. (2010). Neural Control Interfaces. In D. S. Tan & A. Nijholt (Eds.), *Brain-Computer Interfaces* (pp. 21-33): Springer London.

James, J., Ingalls, T., Qian, G., Olsen, L., Whiteley, D., Wong, S., & Rikakis, T. (2006). *Movement-based interactive dance performance.* Paper presented at the 14th annual ACM international conference on Multimedia, Santa Barbara, CA, USA.

Jensenius, A. R. (2007). *Action-sound: developing methods and tools to study music-related body movement.* (Ph.D.), University of Oslo.

Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago: The University of Chicago Press.

Kelly, J. (2007). Pop Music, Multimedia and Live Performance. In J. Sexton (Ed.), *Music, Sound and Multimedia: From the Live to the Virtual*. Edinburgh: Edinburgh University Press.

Klich, R., & Scheer, E. (2012). *Multimedia performance*. Houndmills, Basingstoke, Hampshire; New York, NY: Palgrave Macmillan.

Klonowski, W., Duch, W., Perovic, A., & Jovanovic, A. (2009). Some Computational Aspects of the Brain Computer Interfaces Based on Inner Music. *Computational Intelligence and Neuroscience*.

Krusienski, D. J., Mcfarland, D. J., & Principe, J. C. (2011). BCI Signal Processing: Feature Extraction. In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 123-145). Oxford: Oxford University Press, USA.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.

Lee, Y. S. (1996). 'Technology transfer' and the research university: a search for the boundaries of university-industry collaboration. *Research Policy, 25*(6), 843-863.

Lewald, J., & Guski, R. (2003). Cross-modal perceptual integration of spatially and temporally disparate auditory and visual stimuli. *Cognitive Brain Research, 16*(3), 468-478.

Makeig, S., Leslie, G., Mullen, T., Sarma, D., Bigdely-Shamlo, N., & Kothe, C. (2011). First Demonstration of a Musical Emotion BCI. *Affective Computing and Intelligent Interaction, 6975*, 487-496.

Malloch, J., Sinclair, S., & Schumacher, M. (2014). Digital Orchestra Toolbox for MaxMSP. Retrieved August, 2014, from http://idmil.org/software/digital_orchestra_toolbox

Maréchal, G. (2010). Autoethnography *Encyclopedia of Case Study Research* (pp. 44-47). Thousand Oaks, CA: SAGE Publications, Inc.

Marquez-Borbon, A., & Stapleton, P. (2015). *Fourteen Years of NIME: The Value and Meaning of 'Community' in Interactive Music Research.* Paper presented at the New Interfaces for Musical Expression.

Maxwell, J. A. (2005). *Qualitative Research Design: an Interactive Approach* (Vol. 41). Thousand Oaks, CA: Sage Publications.

Mcfarland, D. J., & Krusienski, D. J. (2011). BCI Signal Processing: Feature Translation. In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 147-165). Oxford: Oxford University Press, USA.

Miles, M. B., & Huberman, A. M. (1994). *Qualitative Data Analysis: an Expanded Sourcebook* (Vol. 2nd). Thousand Oaks: Sage Publications.

Monache, S. D., Polotti, P., & Rocchesso, D. (2010). *A Toolkit for Explorations in Sonic Interaction Design.* Paper presented at the 5th Audio Mostly Conference: A Conference on Interaction with Sound, Piteå, Sweden.

Mondot, A., & Bardainne, C. (2013). Hakanaï.

Moody, N., Fells, N., & Bailey, N. (2006). *Motion as the Connection Between Audio and Visual and how it may inform the design of new audiovisual instruments.* Paper presented at the International Computer Music Conference.

Mora Valentín, E. M. (2000). University–industry cooperation: a framework of benefitts and obstacles. Industry and Higher Education, 14(3), 165-172.

Moreton, S. (2016). Rethinking ‘knowledge exchange’: new approaches to collaborative work in the arts and humanities. *International Journal of Cultural Policy, 22*(1), 100-115.

Mullis, E. (2013). Dance, Interactive Technology, and the Device Paradigm. *Dance Research Journal, 45*(3), 111-123.

Nerone, M. A., Junior, O. C., & Liao, Y. (2014). Classification of the Open Innovation Practices: the Creativity LevelMoving Integrated Product Development to Service Clouds in the Global Economy (Vol. 1, pp. 871 - 879).

Novello, A. (2013). *Exposing the Brain Activity in an EEG Performance: the Case of Fragmentation - a Brain-Controlled Performance.* Paper presented at the ESSEM@AI*IA'13.

Nymoen, K., & Jensenius, A. R. (2011). *A Toolbox for Storing and Streaming Music-related Data.* Paper presented at the 8th Sound and Music Computing Conference, Anchorage, Alaska.

Otto, K. J., Ludwig, K. A., & Kipke, D. R. (2011). Acquiring Brain Signals from Within the Brain In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 81-103). Oxford: Oxford University Press, USA.

Paine, G. (2009). Towards unified design guidelines for new interfaces for musical expression. *Organised Sound, 14*(2), 142-155.

Public and Corporate Economic Consultants. (2012). Strengthening the Contribution of English Higher Education Institutions to the Innovation System: Knowledge Exchange and HEIF Funding.

Rajesh Sethi, Daniel C. Smith, & Park, C. W. (2001). Cross-Functional Product Development Teams, Creativity, and the Innovativeness of New Consumer Products. *Journal of Marketing Research, 38*(1), 73-85.

Rosenboom, D. (1990). The Performing Brain*.* *Computer Music Journal, 14*(1), 48-66.

Royer, A. S., Rose, M. L., & He, B. (2011). Goal Selection vs. Process Control while Learning to Use a Brain-Computer Interface. *Journal of Neural Engineering, 8*(3).

Sashi, C. M., & Perretty, J. (1992). Do trade shows provide value? *Industrial Marketing Management, 21*(3), 249-255.

Schomer, D. L. (2007). The Normal EEG in an Adult. In A. S. Blum & S. B. Rutkove (Eds.), *The Clinical Neurophysiology Primer* (pp. x, 526 p.). Totowa, N.J.: Humana Press.

Science | Business Innovation Board AISBL. (2012). MAKING INDUSTRY-UNIVERSITY PARTNERSHIPS WORK: Lessons from successful collaborations.

Sellers, E. W., Arbel, Y., & Donchin, E. (2011). BCIs That Use P300 Event-related Potentials. In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 215-226). Oxford: Oxford University Press, USA.

Sharir, Y. (2013). *Beyond the Electronic Connection: The Technologically Manufactured Cyber-Human and Its Physical Human Counterpart in Performance: A Theory Related to Convergence Identities.* (Doctor of Philosophy), University of Plymouth.

Shekar, A. (2014). Open Innovation: A Framework for Collaborative Product Development between Industry and Universities *Open Innovation* (pp. 225-255): John Wiley & Sons, Inc.

Siegel, W. (2009). Dancing the Music: Interactive Dance and Music. In R. T. Dean (Ed.), *The Oxford Handbook of Computer Music*. New York: Oxford University Press.

Smith, B. D., & Garnett, G. E. (2011). *Unsupervised Play: Machine Learning Toolkit for Max.* Paper presented at the New Interfaces for Musical Expression, Ann Arbor, Michigan.

Srinivasan, R. (2011). Acquiring Brain Signals from Outside the Brain. In J. Wolpaw & E. W. Wolpaw (Eds.), Brain-Computer Interfaces: Principles and Practice (1 ed., pp. 105-122). Oxford: Oxford University Press, USA.

Stevens, G., Burley, J., & Divine, R. (1999). Creativity + Business Discipline = Higher Profits Faster from New Product Development. *Journal of Product Innovation Management, 16*(5), 455-468.

Straebel, V., & Thoben, W. (2014). Alvin Lucier's Music for Solo Performer: Experimental music beyond sonification. *Organised Sound, 19*(1), 17-29.

Tan, D., & Nijholt, A. (2010). Brain-Computer Interfaces and Human-Computer Interaction. In D. S. Tan & A. Nijholt (Eds.), *Brain-Computer Interfaces* (pp. 3-19): Springer London.

Teder-Sälejärvi, W. A., Russo, F. D., McDonald, J. J., & Hillyard, S. A. (2005). Effects of Spatial Congruity on Audio-Visual Multimodal Integration. *Journal of Cognitive Neuroscience, 17*(9), 1396-1409.

Tekh, E., Hammond, B., & Burns, S. (2014). Ethno Tekh 2014, from http://www.ethnotekh.com/

Thies, M. J. (2012). *Controlling Game Music in Real Time with Biosignals.* (B.M.), The University of Texas at Austin.

United Nations Development Programme. (2013). Creative Economy Report. New York: UNDP/UNESCO.

Vallabhaneni, A., Wang, T., & He, B. (2005). Brain—Computer Interface. In B. He (Ed.), *Neural Engineering* (pp. 85-121): Springer US.

van der Linden, M., van Turennout, M., & Fernandez, G. (2011). Category Training Induces Cross-modal Object Representations in the Adult Human Brain. *Journal of Cognitive Neuroscience, 23*(6), 1315-1331.

Vatakis, A., & Spence, C. (2007). Crossmodal binding: Evaluating the "unity assumption" using audiovisual speech stimuli. *Perception & Psychophysics, 69*(5), 744-756.

Visi, F., Schramm, R., & Miranda, E. (2014). *Gesture in performance with traditional musical instruments and electronics: Use of embodied music cognition and* *multimodal motion capture to design gestural mapping strategies.* Paper presented at the International Workshop on Movement and Computing, Paris, France.

Waisvisz, M. (1985). *THE HANDS, a set of remote MIDI-controllers*. Paper presented at the International Computer Music Conference, Ann Arbor, MI.

Welch, R. B. (1999). Meaning, Attention, and the "Unity Assumption" in the Intersensory Bias of Spatial and Temporal Perceptions. In G. Aschersleben, T. Bakhman, & J. Müsseler (Eds.), *Cognitive Contributions to the Perception of Spatial and Temporal Events*. Amsterdam: Elsevier.

Whitelaw, M. (2008). Synesthesia and Cross-Modality in Contemporary Audiovisuals. *Senses & Society, 3*(3).

Wolpaw, J., & Wolpaw, E. W. (2011). *Brain-Computer Interfaces: Principles and Practice* (1 ed.). Oxford: Oxford University Press, USA.

Wolpaw, J. R., Birbaumer, N., Heetderks, W. J., McFarland, D. J., Peckham, P. H., Schalk, G., . . . Vaughan, T. M. (2000). Brain-computer interface technology: a review of the first international meeting. *IEEE Transactions on Rehabilitation Engineering, 8*(2), 164-173.