Enactive Sound Machines: Theatrical Strategies for Sonic Interaction Design

Fiona Keenan

Doctor of Philosophy

University of York
Theatre, Film and Television

November 2018

For my parents Clare Forbes and Dave Keenan, who gave me two life-defining gifts - my love of words and my love of noises.

Abstract

Embodied interaction with digital sound has been subject to much prior research, but a method of coupling simple and intuitive hand actions to the vast potential of digital soundmaking in a perceptually meaningful way remains elusive. At the same time, artistic practices centred on performative soundmaking with objects remain overlooked by researchers. This thesis explores the design and performance of theatre sound effects in Europe and the U.S. in the late nineteenth and early twentieth century in order to converge the embodied knowledge of soundmaking at the heart of this historical practice with present-day design and evaluation strategies from Sonic Interaction Design and Digital Musical Instrument design.

An acoustic theatre wind machine is remade and explored as an interactive sounding object facilitating a continuous sonic interaction with a wind-like sound. Its main soundmaking components are digitally modelled in Max/MSP. A prototype digital wind machine is created by fitting the acoustic wind machine with a rotary encoder to activate the digital wind-like sound in performance. Both wind machines are then evaluated in an experiment with participants. The results show that the timbral qualities of the wind-like sounds are the most important factor in how they are rated for similarity, that the rotational speed of both wind machines is not clearly perceivable from their sounds, and that the enactive properties of the acoustic wind machine have not yet been fully captured in the digital prototype. The wind machine's flywheel mechanism is also found to be influential in guiding participants in their performances. The findings confirm the acoustic wind machine's ability to facilitate enactive learning, and a more complete picture of its soundmaking components emerges. The work presented in this thesis opens up the potential of mechanisms to couple simple hand actions to complex soundmaking, whether acoustic or digital, in an intuitive way.

Contents

List of Figures

List of Tables

Acknowledgements

Firstly I would like to thank my supervisor, Dr Sandra Pauletto, for her support, guidance and expertise, and in particular for helping me to develop as a researcher throughout my PhD. Thanks to the staff at the Department of Theatre, Film and Television, who welcomed me so warmly and made a space for my work to grow. In particular I would like to thank my postgraduate comrades Romana, Hasmik and Max for their friendship, advice and support.

This research would not have been possible without the support of the AHRC through WRoCAH, which I am eternally grateful for. Thank you also to Caryn Douglas and Clare Meadley at WRoCAH for their guidance, and for supporting our cohort of researchers with such valuable training and development opportunities, which made this experience so enriching and positive.

Thank you to Prof. Jentery Sayers, Teddie Brock, Tiffany Chan and Katherine Goertz for welcoming me to MLab at the University of Victoria and the wonderful world of making and prototyping, which has opened up so many new avenues for my work.

Thank you to all the people who assisted with my sound effect construction ambitions, especially Kevin Lordan at Yorkshire Hut Company for his woodwork mentoring, Erik Olafsen at TFTV for his nautical sewing expertise, and Dr James Hilder at York Robotics Laboratory for 3D printing and laser cutting advice and assistance. Thanks to Stefano Delle Monache for his advice on the SDT.

Thanks to all of my experiment participants for giving their time and opinions, which have been so useful to my research. Thanks also to Andy Fravigar at TFTV for his assistance with my experiment equipment setup.

This research involved a lot of digging for old and overlooked publications. Thank you to the archivists at the Theatre Collection, University of Bristol and the librarians at the British Library, Boston Spa and the University of York for their advice, patience and expertise in tracking down some hidden gems and supporting the very early stages of this research. Thank you to the publishers who assisted with copyright clearance to allow me to include historical images in this thesis. Thanks also to all of those who gave permission for images of their own work to be included.

This research, and thesis, would not have been possible without the freely available open source tools and resources I was able to make use of. Thank you to the data scientists, audio researchers and graphic artists who made their expertise and skills freely available. Their work has been credited throughout this thesis. Thanks also to Jonathan Tate for the LaTex typesetting template you are reading right now.

Thank you to my ever-supportive family, who were ready with advice and encouragement at every stage of this journey. Finally but by no means least, my immesurable gratitude to my partner and fellow traveller John McAreavey, whose love and support kept me moving forward through each challenge, even when my gears were grinding to a halt. This thesis would not be here without him.

This work was supported by the Arts and Humanities Research Council (grant number AH/L503848/1) through the White Rose College of the Arts and Humanities.

Declaration

I declare that the research described in this thesis is my own work, which I undertook at the University of York during 2014 - 2018. All source material has been referenced. This work has not previously been presented for an award at this, or any other, University.

Some parts of this thesis have been previously published in the following conference proceedings and journals. For those items that were published jointly, this thesis includes only the material that I have contributed. For each published item the primary author is the first listed author.

- Keenan, Fiona. (2016). A theatre wind machine as interactive sounding object. In *International Conference on Live Interfaces (ICLI) Doctoral Colloquium*, University of Sussex. Available online [Accessed November 2018].

- Keenan, Fiona and Pauletto, Sandra. (2016). An acoustic wind machine and its digital counterpart. In *Interactive Audio Systems Symposium*, University of York. Available online [Accessed November 2018].

- Keenan, Fiona and Pauletto, Sandra. (2017). Listening back: exploring the sonic interactions at the heart of historical sound effects performance. *The New Soundtrack*, Volume 7 Issue 1, pp. 15-30. Available online [Accessed November 2018].

- Keenan, Fiona and Pauletto, Sandra. (2017). Design and evaulation of a digital theatre wind machine. In *International Conference on New Interfaces for Musical Expression (NIME17)*, Copenhagen. Available online [Accessed November 2018].

- Keenan, Fiona and Pauletto, Sandra. (2017). A mechanical mapping model for real-time control of a complex physical modelling synthesis engine with a simple gesture. In *International Conference on Digital Audio Effects (DAFx-17)*, Edinburgh. Available online [Accessed November 2018].

Chapter 1

Introduction

Our digital experience is filled with sound. As we navigate software menus, every tap or swipe of our fingers is rewarded with a digital *swish*, *beep* or *click* to let us know that our instruction has been received or acted on. These digital sounds display the workings of a device like a smartphone to us in a responsive way. In the virtual space of gaming, digital sound gives life to objects that respond to the movement of our thumbsticks and trigger presses with *clunks*, *zips* or *thwacks*, and selects atmospheric backdrops and responsive soundtracks as we move through locations and levels. Despite the constant presence of responsive sound in these digital environments, we rarely have the opportunity to make sound the object of a digital interaction. We do not hold, shape or manipulate sound directly through the use of touchscreens or game controllers. The potential of digital sound as an interactive and expressive material has been explored almost exclusively from the perspective of music making. Performers with prior experience of musical instrumentation and production, or audio programmers and designers, possess the technicality required to innovate within this area. There has been a wealth of interesting developments in digital music performance systems, including instruments fully realised within Virtual Reality (VR) environments (Serafin et al. 2016). Everyday human perceptual experience of sound is a much broader phenomenon however, and digital systems have yet to offer a simple gateway to exploiting this more general understanding in the facilitation of creative bodily-guided soundmaking.

This thesis begins work to bridge the gap between simple hand actions and the vast potential of digital soundmaking through an exploration of a historical context of sound design and performance; the creation of sound effects in theatres in the late nineteenth and early twentieth century. As interactive objects that produced performable sounds mimicking everyday events such as rain, wind and thunder, theatre sound effects are simultaneously a simple action, an interactive object affording that action, and a performable complex sonic illusion. They are also, crucially, "enactive sound designs" (Franinović 2013, p.21) in that they offer the potential to develop bodily skill in sound performance through a simple process of exploration and rehearsal while listening to self-produced sound. They have much to teach us about how to design an object to facilitate simple and meaningful interactions with digital sound and how bodily skill in creative expression with digital sound could be accumulated enactively. They may also hold the key to how prior bodily knowledge of soundmaking by manipulating real-world objects, such as that accumulated by Foley practitioners creating sounds for film soundtracks (Pauletto 2017, p.343), can be transferred into digital soundmaking.

When performing with a theatre sound effect, the sound itself feels like it is the material being shaped and manipulated. The sound effect is designed to place the sound directly into the hand of the performer. This research focuses on the experience of sound performance from a performer's perspective. It explores the potential of historical theatre sound effects to inform the design and evaluation of simple continuous interactions with digital sounds, and works to understand and transfer the real-world perceptual experience of performance with a sound effect to a digital sonic interaction. This chapter discusses the scope of this research and presents the research questions. The structure of this thesis and its contribution to knowledge are also outlined.

1.1 Research Scope

This research works at the point of convergence of two main areas of sound and music computing research, that of Sonic Interaction Design (SID) and Digital Musical Instrument (DMI) design. SID is an interdisciplinary design approach that focuses on the use value of sound in interactions with electronic devices or digital systems - how sound can make meaning by giving information, guiding behaviour, or influencing how a user feels about an object (Franinović & Serafin 2013, p. vii). The design potential of sound in this context is informed by research into everyday human sensory experiences of touch, movement and audition. How ordinary real-world sound events, such as impacts or scrapes, are perceived is also important to SID in exploring how sound might convey meaning to a user. DMI design focuses on ways to connect the human body to digital music making systems to facilitate music creation and performance. Of particular interest to this research is the area of DMI design focused on musical *expression*, where designers look to create meaning for the performer through the physical properties of the musical tools that they use, adding tangible materials or resistances to traditionally light and effortless computer interfaces like game controllers, for example (Miranda & Wanderley 2006). Sound in this context responds to a performer's intention as they move and gesture in the activation and modulation of complex and even abstract sound events and textures. By converging these two fields, this thesis presents a cross-pollination of SID and DMI design that helps to expand the scope of each. In particular, this research explores the influence of the material qualities of a handheld interactive object on the performer's experience of a sonic interaction, something not previously considered in SID research. It also brings a rigorous approach to the way that new systems for expressive digital sound performance can be evaluated in an experiment with participants, an area still in development within DMI design.

In addition to this convergence of SID and DMI design research, this research engages with a historical design practice, and examines the creation of sound effects for theatre performances in the late nineteenth and early twentieth century. Theatre sound effects have not been previously researched within the context of a practical and technical enquiry into sonic interactivity. As such, this work integrates knowledge and methods from a variety of other disciplines to pursue a design-led enquiry into the interactive potential of this era of soundmaking. It uncovers historical knowledge of sound creation through human performance by researching the history of theatre sound effects design, and then applies that historical understanding to inform contemporary methods of digitally designing sounds. It uses current understanding of the human perceptual experience of sound to contextualise and probe further into historical descriptions of sounds heard and performed. It incorporates practical methods of remaking (Elliott et al. 2012) to extend the historical enquiry further and create a working theatre sound effect, and then a digital prototype to mimic its sound. Finally, it combines qualitative and quantitative evaluation methods to examine the sonic interaction this sound effect, and its digital prototype, afford in an experimental setting. This integration across disciplinary approaches makes this research project interdisciplinary in nature, but it also moves towards transdisciplinarity in its attempt to unify these approaches even further (Stember (1991) and Jensenius (2012)) into an overarching strategy for sound-and-enactive-object design and evaluation. A design practice based on creating and evaluating enactive sound machines - simple, learnable interactive objects for complex soundmaking through hand movement - is what draws these different areas of research together (Figure 1.1).

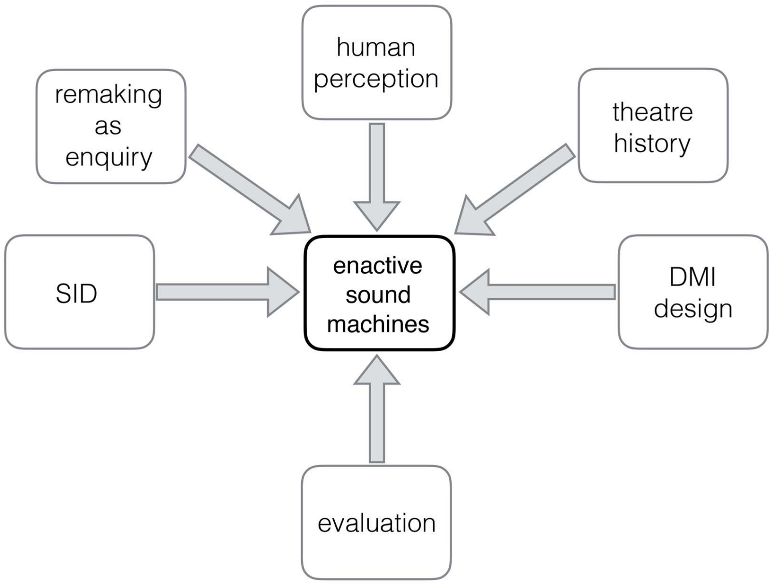

Figure 1.1: The research areas that converge within this thesis.

1.1.1 Sonic Interaction Design Approach

This research explores how simple hand actions might be connected to complex digital sounds in a more immediately meaningful and playable way in order to open up digital sound to creative practices based on building skill in bodily movement. It centres on the particular design problem at the heart of SID research. This is embodied interaction (Dourish 2001), or how to make use of everyday bodily understanding of action, movement and manipulation in the creation of an encounter with a digital system. SID research into the creation of new sonic interactions has so far focused on how potentially meaningful connections between actions and sounds can be discovered and then captured as particular designs (Hug 2008, p.15). The idea of meaning in relation to action and sound metaphors is complex; this is dependent on the bodily experiences that a person has already accumulated, and how the sonic interaction upholds or confounds these. Does a *twist* make sense when it results in a *creak* to only one individual, or is this a much broader cross-culturally shared understanding, for example? Is a *twist-creak* configuration easier to learn through rehearsal than a *pull-creak* configuration?

To discover these potential meanings, SID design work has drawn inspiration from everyday experiences not directly related to soundmaking. Action-sound configurations are examined and unpacked to become the subject of a new design. A good example of this is Franinović's work with simple mechanical kitchen tools, like a garlic crusher or salad spinner (2013, p.15). Through observation and reflection exercises, Franinović worked to closely observe her own use of these tools, producing a palette of action primitives that she then used to inform design concepts for sonically interactive objects. These interactive objects were realised with various sensors housed in 3D printed plastic shells, shaped for intuitive manipulation by a user in real-time. One of the designs that closely informs this research, known as Spinotron (Lemaitre et al. 2009, p.978), will be discussed in further detail in Chapter 2.

While an interactive action (such as *twist*) might be meaningfully configured to activate and modulate a digital sound (such as *creak*), the way that action is materially facilitated also impacts on how it can be understood in light of previous bodily experiences. It is this issue of material that connects the work of this research with previous research in DMI design. DMI design has focused much more closely on how the transition between an interaction with an acoustic musical instrument to one with a digital system for sound performance implies a loss of sensory information, and therefore meaning, for the performer. This information is the constant resistance to our bodily action that we experience in daily life. Pushing a finger against a rubber pad to produce a musical note from software is very different to pushing a finger against a tensioned string, for example. The complexity of the tensioned string's texture under the fingertip, and its resistance to pressure and movement, is missing from the rubber pad.

We can unpack this transition further by returning to the example of a metal kitchen garlic crusher. While exploring what sound might be meaningfully configured with the action of opening and closing this tool was the focus of Franinović's research (2013, p.15), this research examines a further potential dimension of this kind of interaction. Interacting with the garlic crusher in everyday use while it is empty is very different to when it is packed with cloves of garlic. While empty, there is much less resistance to the user's manipulations. If we move on to design an interactive object in 3D printed plastic to mimic the garlic crusher exactly, the action of opening and closing remains the same, but the tactile and resistive qualities of the metal version have been lost. While the garlic cloves have been removed, properties such as the metal's smoothness, weight, initially cool temperature, and tight metal hinge are also missing.

This research engages directly with the difficulty in noticing, unpacking and capturing the complexity of the transition from acoustic to digital soundmaking. It does this by applying the expertise and design focus of SID research in a different direction. Rather than trying to generate new meaningful connections between actions and sounds, it works to examine the already realised, tested and refined collection of real-world action-sound configurations at the heart of historical theatre sound effects. These configurations are not just metaphorical links between actions and sounds, but are materially and mechanically facilitated. This requires the performer to continuously navigate force, friction and weight in a simple action to make sound just as they would when performing with a more complex acoustic musical instrument. One particular historical sound effect design is remade, observed and reflected on in detail to examine how the qualities of the sonic interaction it offers a performer can be transferred to the digital domain. Working to resolve this design problem necessitates an exploration of how the sonic interaction is perceptually meaningful to the performer, how the material resistances of the design guide the performer in performative soundmaking, and how these resistances might be captured and realised digitally. This work also brings a historical perspective to bear on present-day methods for interacting with digital sound, revealing more about their inherent limitations, challenges and opportunities.

1.1.2 Making Use of Designerly Histories

This research makes use of history by establishing a collaborative exchange between the expertise of historical theatre sound effects designers in the late nineteenth and early twentieth century and present-day research into the potential of embodied digital soundmaking. This strategy is informed by previous SID research. The need for an artistic skillset in materially situated soundmaking to help realise new digital sonic interactions has been highlighted by Franinović (2013, p.20) and Pauletto (2017, p.346). They both look to present-day Foley practitioners, who use objects and materials to create sounds for film, as a potentially rich source of knowledge about meaningful configurations of actions and sounds.

SID research has also engaged directly with historical methods of soundmaking. Serafin et al. (2006) created a performable digital version of a family of historical musical instruments, Luigi Russolo's intonarumori, in what they have termed an "enactive preservation" (Serafin & De Götzen 2009) of their action and sound properties. Russolo's family of noise intoners were created in the early twentieth century to realise the ideas of his Futurist manifesto "The Art of Noises" (Russolo 1913). Russolo proposed that noises such as *howls*, *crashes* or *scrapes* should be used to create a new kind of music. Each intonarumori consisted of a wooden box housing a crank handle mechanism and rotating disc, which turned against a tight catgut string threaded through a drumskin. A flared cone around the drumskin amplified the soundmaking apparatus inside the box. As a digital controller, the enactive recreation of the intonarumori incorporates few of the material resistances to bodily movement that a fully acoustic version would have presented (Fig.1.2). In the context of this research, the proposed enactive preservation method is framed as another important step in exploring the transition between acoustic and digital soundmaking.

Figure 1.2: Comparing (L) a fully acoustic reconstruction of Russolo's intonarumori (author's own) with (R) the digital enactive reconstruction (Serafin et al. 2006) (images by Stefania Serafin, used here with her permission).

This research extends Serafin and De Götzen's approach even further by making connections between the field of SID and new methods of enquiry into history, material culture and technological practices. The "maker turn" (Staley 2018) in humanities research has seen hands-on activities like making, prototyping and fabrication pursued to generate new knowledge and engage more fully and critically with digital media (Sayers et al. 2015). Elliott et al.'s prototyping work (2012), which recreated visual illusions from nineteenth century manuals of stage magic, is of particular relevance to this research. As part of this enquiry, a fully working acoustic theatre sound effect is constructed, which facilitates a deeper exploration of the sonic illusion it creates and the perceptual experience it offers to a performer. Its design is supported by an enquiry into a variety of historical sources reporting on traces of what must have been a practice that imagined sonic possibilities only in terms of action, mechanism and material. Informed by a research through design (RTD) approach (Frayling (1993) and Durrant et al. (2017)), the practical and technical aspects of the work undertaken to realise this working sound effect are a critical part of the enquiry itself, as the author's own implicit knowledge of interface design for sound performance becomes part of a dialogue with the expertise of historic designers through an enactment of their instructions. Then, in an investigation of the sonic interactivity of the working sound effect, the author's embodied knowledge of sound performance also converges with the descriptions of historic performers to explore the experience of soundmaking it affords.

This tactile and experiential work also draws on previous research into embodied skill acquisition and making practice from the field of anthropology (Ingold (2006, 2013)). By bringing these fields of enquiry around making practice and technologies together to examine a historical context of soundmaking, this approach foregrounds the importance of the sonic interaction designer in engaging with and interpreting the embodied knowledge that might be secreted away within an artistic soundmaking practice like present-day Foley. It also positions the sonic interaction designer as the practitioner best placed to connect traces of information on soundmaking from different historical sources together, and to observe, unpack and reflect in light of their own bodily knowledge. This connects the fields of SID and DMI design, and the technical work of sound design and audio programming, to broader areas of scholarship in historical sound studies, design practice, anthropologies of skill and making, and critical approaches to technology.

It is proposed that this approach is vital to a better understanding of early histories of the creative use of sound. As will be argued in Chapter 3, histories of performance-based soundmaking practices have so far been traced from Russolo's avant-garde "Art of Noises" manifesto (Russolo (1913) and Kahn (1999, p.10)), and so more popular entertainment forms like theatre have not been considered interesting either to sound studies (Ovadija 2013, p.9) or to present-day SID or DMI designers. This state of affairs fails to recognise that the technology at the heart of Russolo's manifesto - the intonarumori themselves - may have firm roots in a long history of sound creation for theatre. Intonarumori were sometimes fabricated and repaired in theatre workshops, for example (Davies 1994).

This research proposes that soundmaking practitioners should be embraced, no matter their site of performance, as designers working to solve the problem of sonic interaction. Theatre sound designers in the late nineteenth and early twentieth century shared the same struggles as avant-garde composers in trying to realise their ambitions for total control over sound with a universal sound machine (Hopkins 1920, p.80) for example. This approach has the potential to open up a new and interesting field of enquiry into sound and the performing body through history. It will also aid further reflection on the present-day design conventions around performing sound as electrical signal (e.g. analog synthesizers) and sound as data (e.g. audio software) in light of another era, when sound was only designed and performed through action, material and mechanism. Present-day methods of soundmaking have established workflows that may overly influence how we perceive the potential of extending a performers intent into digital systems. This contextualisation of the conditions in which the designer works, or a "dialectic of tools" (Hug & Kemper 2014), runs throughout this thesis, and brings a critical design philosophy that interrogates naturalness (Norman 2010, p.6) to bear on the creation of interactive sounding objects.

This section has outlined the transdisciplinary scope of this thesis, highlighting its approach to sonic interaction design and engagement with historical sources. The research questions that direct this enquiry will now be presented, along with the structure of this thesis and its contribution to knowledge.

1.2 Research Questions

This thesis works to answer the following research questions:

1. How can historical theatre sound effects be fully investigated as interactive sounding objects?

2. How can the enactive properties of historical theatre sound effects be captured in the design of a digital sonic interaction?

From this, there follows several sub-questions:

- What theories, methods and frameworks are available for studying historical theatre sound effects as interactive sounding objects?

- Which historical theatre sound effect designs offer the most interesting action-sound configuration to the performer?

- How might a historical theatre sound effect be made digital?
- How can a historical theatre sound effect best be evaluated?

1.3 Thesis Structure

Chapter 2 presents the main research areas and theoretical concepts that inform the work presented in thesis. Section 2.1 defines what a continuous sonic interaction is, and explores research into human perceptual experience of action and sound. Section 2.2 discusses available methods for designing a meaningful continuous interaction with digital sound from SID and DMI design research. Section 2.3 explores the challenges of evaluating continuous sonic interactions, and presents potential solutions from SID and DMI design research. Section 2.4 introduces theatre sound effect design practice in the late nineteenth and early twentieth century, and outlines its potential to solve the design and evaluation challenges highlighted in this chapter. Section 2.5 presents a summary of the chapter.

Chapter 3 presents a design history of the practice that created theatre sound effects in the late nineteenth and early twentieth century, and introduces a framework for the creation of digital sounding objects based on some particular designs. Section 3.1 establishes the approach and scope of the historical enquiry, and explores some potential reasons for theatre's soundmaking expertise being so far overlooked by researchers. Section 3.2 explores the expertise historical practitioners used to create sound effects in theatres during this era using the framework of Cross' "five aspects of designerly ways of knowing" (2006, p.29). Section 3.3 explores the most widely used theatre sound effects used at this time in more detail, and presents a framework for the realisation of new correspondences between actions and sounds based on these designs and the Sound Design Toolkit (SDT) digital synthesis taxonomy (Baldan et al. 2017, p.258). Section 3.4 examines one of the historical effects presented in this chapter, a wind machine, in more detail as a particularly interesting design to use as the focus of this research. Section 3.5 summarises this chapter.

Chapter 4 introduces the methodology that this research pursues in order to answer the research questions. Section 4.1 presents the design-led approach. The method of remaking to explore the designerly knowing of historical practitioners is outlined. The programming and interactive sounding object design approach is also presented. Section 4.2 presents the acoustical and experimental evaluation methods used to evaluate the design work undertaken and indicate directions for future work. Section 4.3 summarises the research methodology presented in this chapter.

Chapter 5 outlines work to remake a working acoustic wind machine according to historical design instructions and investigate the sonic interaction it offers the performer. Section 5.1 outlines work to synthesise several examples of design instructions from historical texts to create a design for the wind machine. Section 5.2 presents the work undertaken to construct this design and make it work as it should. Section 5.3 explores the perceptual experience the wind machine offers the performer from the perspective of the author. Section 5.4 presents an overview of how the wind machine creates sound in performance in order to inform the programming work outlined in the next chapter. Section 5.5 outlines a summary of the work presented in this chapter.

Chapter 6 builds on the case study in Chapter 5, outlining work to create a prototype digital wind machine based on the Sound Design Toolkit's (SDT) physical model of friction in the Max/MSP graphical programming environment. Section 6.1 examines how the movement of the acoustic wind machine was captured as data by coupling a rotary encoder to its cylinder, allowing it to drive the progress of its digital counterpart. Section 6.2 presents the work undertaken to model the main soundmaking elements of the acoustic wind machine digitally. Section 6.3 outlines the mapping approach to connect the data from the rotary encoder to the model in Max/MSP. Section 6.4 presents the acoustical evaluation work undertaken to compare the acoustic and digital wind-like sounds in performance in order to validate the prototype digital wind machine before proceeding to the experiment outlined in Chapter 7. Section 6.5 outlines the stages of work undertaken in order to create a digital crank interface to drive the model in Max/MSP without using the acoustic wind machine. Section 6.6 summarises this chapter.

Chapter 7 reports on the procedure and results of an experiment with participants to evaluate the wind-like sounds and sonic interactivity of the acoustic and prototype digital wind machines. Section 7.1 presents the aims and hypotheses and reports on a pilot study undertaken to test the procedure. Section 7.2 reports on the procedure and results of the first listening step, which evaluated the similarity of the acoustic and digital wind-like sounds. Section 7.3 reports on the procedure and results of the second listening step, which evaluated participants' understanding of the metaphor of rotational speed in relation to the wind-like sounds. Section 7.4 reports on the final step of the experiment, which asked participants to perform with the acoustic and prototype digital wind machines and describe their experiences. Section 7.5 summarises and discusses the results of each step of the experimental evaluation. Section 7.6 summarises the work presented in this chapter.

Chapter 8 presents a general discussion of the research undertaken and the findings presented in the previous chapters, and draws conclusions in relation to the research question presented. Recommendations for future work are also given. Section 8.1 re-states the research questions and aims of the research. Section 8.2 summarises the methodology used to answer the research questions. Section 8.3 discusses the work undertaken in light of the theoretical background and prior research presented in Chapter 2. Section 8.4 summarises the findings of this research in relation to the research questions. Section 8.5 explores some potentially fruitful areas for further work arising from this thesis. Section 8.6 summarises this chapter and concludes this thesis.

1.4 Contribution to Knowledge

This research contributes to knowledge through the establishment of the historical practice of theatre sound effects design and performance in the late nineteenth and early twentieth century as a site for the generation of new approaches to facilitating continuous interactions with digital sounds. It takes the first steps in applying the embodied knowledge and soundmaking expertise of theatre practitioners in the late nineteenth and early twentieth century to the present-day design problem of embodied interaction through the construction and detailed study of one particular sound effect design, and the creation of a performable prototype digital model of its soundmaking components. Through a design-led exploration, acoustical analysis and experimental evaluation of the acoustic and digital effects, the importance of the dynamics of the sound effect's mechanism in guiding the performer to perform a complex continuous sound with a simple gesture is evidenced. Engaging with a very different context of creative soundmaking builds new connections between very simple real-world interactions and more complex encounters with digital sounds, creating a coherent phenomenological framework of sound performance across a variety of materials, mechanisms and digital systems. This research foregrounds the importance of a transdisciplinary approach in forging fruitful connections between SID and DMI design research and the embodied expertise of performative soundmaking practitioners. It positions SID and DMI design research methods as potentially vital contributions to wider fields of enquiry encompassing the cultural history of technology, making, performance and the senses. It also opens up sonic interactivity, and the wicked problems of facilitating acoustic and digital sound performance, to potential contributions from design researchers working beyond the conventions of present-day audio and sound practices.

1.5 A Note on Terminology

This research aims to connect the practice of late nineteenth and early twentieth century theatre practitioners with present-day research into the creation of perceptually meaningful interactions with digital sounds. Exploring creative soundmaking techniques from an age before sound as signal or data necessitates engagement with a language of sound that is unfamiliar to present-day practitioners in SID or DMI design. Particular attention has been paid to how historical terminology has been deployed throughout this thesis to ensure that this connection between design practices in disparate contexts is engaging and fruitful. Some definitions are stated here to aid the understanding of the reader.

1.5.1 Theatre Sound Effects

The terminology used to describe sound, and sound performance, in late nineteenth and early twentieth century theatres originated long before the present-day term *sound effects* (or SFX) implied the use or manipulation of recorded sounds, or a collection of audio files (.wav, .mp3 etc.) ready to insert into a film or game soundtrack.[1] Theatre sound effects were instead designed purely through action, mechanism and material at this time, and the naming conventions around this practice were much more loosely defined without the process of recording to capture a sound event and commodify it as a piece of fixed media. Thus, the term *sound effect* means something very different when deployed in historical sources.

[1]Indeed, present-day SFX used in theatres are largely a collection of audio files played back through loudspeakers.

In his instructional manual "Equipment for Stage Production", Krows (1928, p.114) offers a useful definition, saying that an effect is "any impression produced upon an audience by mechanical means". Although an effect is merely an *impression*, the term is regularly deployed as if referring to an object, such as "working an effect" (Green 1958, p.15) for example. Freed from the process of recording, these sound effects are more ethereal than our present-day understanding of recorded sound. They are born in the space between the performer, the interactive sounding object, the audience, the dramaturgy, and the theatre space. To engage more fully with historical designers, their terminology is retained as much as possible throughout this thesis to allow their understandings of the design and performance of sound to be clearly heard. Inadequate terms such as *apparatus*, *device*, *machine* or *prop* will not be used to qualify the term *effect* to avoid overburdening the reader and also to keep the meaning clear.[2]

Chapter 2 presents a new term, *interactive sounding object*, which is used in this thesis to capture these various methods and connect them to digital soundmaking with objects. The word *theatre* is used throughout to qualify the term *sound effects* and draw the reader's attention to the historical context of soundmaking. When a specific method is being discussed the interactive sounding object in question (e.g. a rain machine) will be clearly defined. Otherwise, the term *theatre sound effects* will refer to the general collection of late nineteenth and early twentieth century backstage soundmaking devices and the sounds and impressions they produce, just as it did when they were in common use.

1.5.2 Designer, Performer, Practitioner

Finding the best term to describe who created and operated sound effects within theatres in the late nineteenth and early twentieth century is similarly challenging. The historical research presented in Chapter 3 evidences that sound effects were not designed by a dedicated group of sound specialists such as we would recognise in present-day theatre, film or game design. The activity of designing the sound effect was somewhat separate from constructing it, and in turn separate from performing with it. Expertise in soundmaking is therefore diffused throughout different backstage and production roles during this era. Any member of a production - stage manager, director or patron - can be assigned credit

[2]For example, *machine* implies mechanical power and movement, but not all effects were produced in this way.

for a particularly interesting design in historical sources.

It is the prop-maker, or property master, that has been singled out by writers as the person who actually constructed sound effects (Vincent (1904, p.418) and Somerfield (1934, p.77)), but the director has ultimate control over the "intensity and duration" of the performance (Jackson & Wilson 1976, p.132). Similarly, a successful performance of effects is attributed to directors and stage managers rather than to the stagehands that performed them. There is a division of labour implied in the way this historical practice is described. As will be explored further in this thesis, the finer details of mechanical principles and carpentry techniques, and how designs were specifically refined and made to work, have not been very well preserved. This research works to reveal some of the implied expertise in carpentry and performance that is not explicitly described in historical sources.

Throughout this thesis, two simple terms are used to distinguish between the making and performance stages of a theatre sound effect. *Designer* refers to all those who worked to design and construct a sound effect, and *performer* refers to all those who performed and directed the performance of the sound effects. To encapsulate both aspects of the practice, the term *practitioner* is used.

1.6 Summary

This chapter has presented the context and motivations of this research project and outlined the research questions that guide the discussion in the following chapters. An outline of the thesis and its contribution to knowledge was also presented. The issue of terminology when engaging with a historical context of soundmaking was highlighted, and the terms used throughout this thesis when discussing theatre sound effects practice in the late nineteenth and early twentieth century were clearly defined for the reader. Next, Chapter 2 explores the theoretical background to this research enquiry.

Chapter 2

Interactive Sounding Objects

This research explores the potential of historical theatre sound effects to inform the design and evaluation of simple continuous interactions with digital sound. It is concerned not only with how digital sound can make meaning in an interaction, but also how digital sound can be performed through interaction. As outlined in the previous chapter, this approach converges prior research in the fields of Sonic Interaction Design (SID) and Digital Musical Instrument (DMI) design. Both in their own way tackle the design problem of embodied interaction (Dourish 2001), or how to build on everyday human sensory modalities and perceptual experience to design meaningful encounters with digital technology. In the case of SID and DMI design, these encounters specifically produce sound, whether to enhance an interaction or to facilitate music creation. In SID, sound is used specifically to influence a user by displaying information or guiding sensorimotor activity in the completion of a task. The aim is to use sound in a design to connect a user's touch, movement and audition together in a natural multimodal way (Franinović & Serafin 2013, p.x) in order to mimic real-world perceptual experience. When evaluating new sonic interaction designs, users are not generally asked to focus on soundmaking for its own sake, but are rather helped to focus on something else (e.g. a visual task) by a sound. SID research attempts to address two main questions: how to study the process of meaning creation with sound, and how to create and evaluate sonic interactions (Pauletto 2013, p.2).

DMI design research is concerned with facilitating music creation using digital sounds, interfaces and systems. Scholarship in this field has centred on the International Conference on New Interfaces for Musical Expression (NIME), which takes an increasingly broad view of how technology can be used for musical aims (Jensenius & Lyons 2017, p.vii). This research draws on prior DMI design research specifically focused on creating a continuum of possibility for "musical expression" for performers across acoustic and digital instrumentation.[1] Strategies include acoustic instruments extended with sensors or a control interface to integrate digital sound capabilities (known as hyper or augmented instruments), and fully digital configurations of performance interfaces and software. In the use of a DMI, the focus of an interaction with an instrument is the resulting sound, its musical potential, and the level of expression it affords the performer. The quality of the resulting sound guides the performer in the use of the instrument. DMI design is informed by research into the qualities of acoustic musical instruments and the expertise of instrumentalists and musicians. Connecting human bodily action to a computer necessitates the use of digital interfaces such as keyboards, game controllers and touchscreens. These devices are developed as part of Human-Computer Interaction (HCI) research.

SID and DMI design both borrow from HCI's expertise in that they make use of interfaces such as game controllers, or use HCI components like sensors to construct new interfaces of their own. They also look to HCI for ways to evaluate users and performers when they are interacting with sound or playing digital music. Given their different perspectives, there is much that SID and DMI researchers can learn from each other, and indeed the two fields are already beginning to converge (Heinrichs 2018, p.19). SID's expertise in empirical evaluation methods and the rigorous deployment of perceptual research in the design of objects to link action and sound together would benefit DMI designers. DMI design, in striving to unpack, examine and recreate some of the complexity of the feel of the acoustic instrumental experience, has much to offer SID in its understanding of nuanced physical movement guided by material qualities, and expressive intent in the performance of sound.

This chapter provides a theoretical context to the work of this research and introduces the relevant areas of prior research underpinning this thesis. The term *interactive sounding object* is used throughout to guide the subsequent discussion that connects and converges prior research from both SID and DMI design. This term is formulated from "digital sounding objects," which is used in SID. These are digital models of real-world sounds that respond to human interaction (Rocchesso et al. 2003, p.2). In the context of this research,

[1] Musical *expression* is the subject of much debate in NIME research, with technology also framed as offering the potential to interrogate more traditional notions of human virtuosity or even performance itself (Tanaka (2000), Gurevich & Treviño (2007)).

the interactive sounding object is defined as a real-world object designed specifically for manipulation by a performer, in order to produce sound. That sound can be acoustic or digital depending on how the particular object has been designed. This connects everyday experiences of soundmaking, creative acoustic soundmaking with objects, and techniques for digital soundmaking together (Figure 2.1).

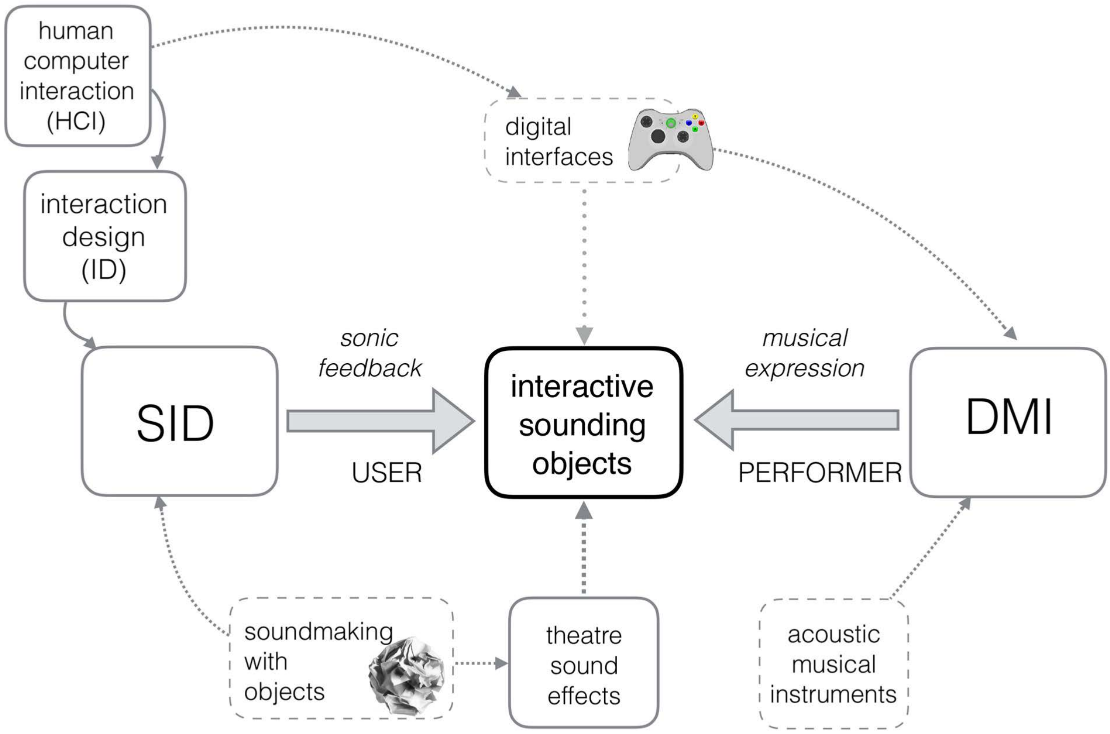

Figure 2.1: Defining an interactive sounding object.

Section 2.1 defines the term *continuous sonic interaction*, and explores research into human perceptual experience of sound and action. It also explores how sound and action are intimately linked and how this connection influences embodied learning and performative intention. Section 2.2 introduces the design problem at the heart of this research - how the rich and intuitive nature of real-world sonic interactions can be made digital. It outlines how research in both SID and DMI design has approached the problem of connecting bodily movement with continuous digital sound. Section 2.3 discusses the challenges inherent in evaluating continuous digital sonic interactions. It focuses on some potential solutions suggested by both SID and DMI research that are explored as part of this research. Section 2.4 presents the case for a closer examination of the practice of theatre sound effects design, arguing that it has the potential to open up a new area of design and enquiry into simple and rich interactions with digital sound. Section 2.5

presents a summary of the discussion in this chapter.

2.1 Continuous Sonic Interaction

In an interaction with a digital sound, users and performers have a direct effect upon an interactive sounding object through manipulating, shaping, configuring or acting on it in some way. The degree to which this digital sonic interaction feels natural and fluid, or *embodied*, is determined by how well it facilitates the natural real-world feedback loop of human perception and sensorimotor activity. We are constantly engaged in an embodied dialogue with our everyday surroundings, as input from visual (sight), auditory (listening), tactile (touch) and kinaesthetic (movement) sensory modalities influence our ongoing cognition or action. This continuous orientation of the body in the world is always *multimodal*; these sensory inputs and movement adjustments are not distinct from each other, but form a whole experience that is challenging for designers to transfer to the digital domain.

In a sonic interaction, a user controls a sound through movement, while simultaneously getting sonic, tactile and kinaesthetic information back about the result of their own movements, in turn helping to inform further action and adjustment. Figure 2.2 illustrates this process in the case of a very simple sonic interaction with a ball of crumpled paper and a wooden surface. By holding the crumpled paper and scraping it across the surface, the performer experiences tactile feedback from the feel of the texture of the paper and the hardness of the wood beneath it. They also feel the resistance to movement between the paper and the surface caused by friction. The pushing action produces a scraping sound, and the performer can continually adjust their grip on the paper or the trajectory of their arm to influence the quality of the sound they are making. The interaction is utterly *continuous*; action reliably, directly and constantly produces and modulates the quality of the sound.

By activating and experiencing this sonic interaction, the performer can develop embodied knowledge about how the scraping sound can be performed. This knowledge will build on any previous understanding they have of the link between their bodily actions and the potential they have for soundmaking. They can also learn to focus their intention according to the sonic feedback during the interaction to perform a particular kind of scrape. SID engages with theories of human perception and cognition to describe and understand embodied knowledge, conceptual links between action and sound, and the accumulation of bodily skill in soundmaking. These theories will now be considered in more detail with reference to the example of a continuous sonic interaction with crumpled paper shown in Figure 2.2.

2.1.1 Embodied Knowledge of Soundmaking

The embodied interaction approach stems from phenomenology, a school of philosophy founded by Edmund Husserl at the turn of the twentieth century, which focuses on the minutiae of daily human experience as a source of insight into consciousness and behaviour. This close examination of daily experience often centres on encounters with objects and technologies (Ihde 2016, p. xiii). The theories of philosophers Heidegger and Merleau-Ponty, who drew on Husserl's philosophical method, are of particular interest to researchers exploring the potential of embodied interaction (Svanæs 2013). Heidegger's fomulation of human experience of tool use is central to this enquiry, and will be explored further in Section 2.1.4.

Merleau-Ponty proposed that human perception is a continually active process, and that what he termed the "lived body" does not view the world in an objective way. Instead, humans are "situated" beings; constantly influenced by the environment around them and the skills and experience they have already accumulated (Merleau-Ponty 2002, p.90). This idea established the theory of enactive cognition within psychology, which frames human beings not as passive receivers of information from their environment, but instead constantly engaged in sensorimotor activity in order to perceive the world around them (Varela et al. 1992, pp.149-150). Our senses are all "touch-like" (Noë 2004, p.1), and require us to move, adjust and direct our focus in order to clearly perceive. Therefore how we act in the world, or how we know how to act, determines what we see, hear, smell,

Figure 2.2: A continuous sonic interaction with a ball of crumpled paper to produce a scraping sound.

taste and touch.

In the case of the continuous sonic interaction in Figure 2.2, moving the crumpled paper will help the performer to understand its potential for scraping. Looking, listening, feeling and moving in a conscious way during their interaction will allow them to continually monitor the effect of their action and adjustment on the resulting scraping sound. This adds to the performer's repository of embodied knowledge gained through acting to perceive. Embodied knowledge is tacit (Polanyi 1967, p.29) and expands the human body into the environment around it, allowing it to reach out and interact confidently and effectively. Tacit knowing is gained through experiences of actions, such as tool use. Following a continuous sonic interaction with the ball of crumpled paper, the performer will be able to approach another similar interaction with more understanding of how their actions can be directed to make sound. If the performer is actually a Foley artist creating sounds for a film soundtrack, they may be able to accumulate embodied knowledge of a highly nuanced soundmaking practice focused on different qualities of scraping sounds. In both cases, the knowledge will be personal and situated. It will be difficult to pass on to someone else in a set of explicit instructions. For example, it may be possible to write a set of instructions or show the Figure 2.2 diagram to another performer and ask them to perform this scraping sound, but the fine experiential details around the trajectory of the hand and how feel guides its movement is difficult to explain in the abstract. It is only through an individual rehearsal process with the crumpled paper that the performer can learn exactly how aspects like the material of the paper and surface, the hand's grip pressure and the arm's speed of movement are linked to the resulting sound.

This kind of embodied learning is built on the constancy of the physical properties of our environment. Materials and objects usually respond to our manipulations in familiar and expected ways. When performing the scraping sound, the wooden surface beneath the crumpled paper maintains its hardness, and does not become elastic and springy in response to the performer's movement. The crumpled paper ball on the other hand will change its size or shape in response to grip pressure, but this happens within expected parameters. Depending on the stiffness of the paper, it will spring back slightly as the ball opens up in response to the performer loosening their grip. It is this repeatability and reliability that facilitates embodied learning.

The reliability of the scraping sound itself is another of the physical properties that guides the performer to acquire embodied knowledge in this interaction. Listening to self-produced sound is something that Chion calls "ergoaudition" (2015, p.671). Through their ergoaudition of the scraping sound as they move, the performer in Figure 2.2 is guided to continually adjust the trajectory and pressure of their hand according to the kind of scrape that they intend to perform. If the performer has no prior embodied knowledge of acting specifically for the purpose of soundmaking (if they are making a Foley sound for the first time, for example), this new learning experience will still be scaffolded with other sources of knowing. They may be very familiar with the tactile feel and sound of paper pages from notebooks for example, which will prove useful here. They may understand conceptually that the sound of scraping is somehow metaphorically connected to a sliding or pushing motion.

It is this enmeshing of sensory modalities and prior bodily knowledge that gives the performer the ability to continuously develop their tacit knowledge of soundmaking with the ball of crumpled paper. If an interactive sounding object is truly enactive, it should give a performer the same kind of opportunity for embodied knowledge accumulation through exploration that this ball of crumpled paper gives our performer in Figure 2.2. Designing such an encounter also requires a consideration of the potential meaning of a sound event, and how a performer might be invited to begin the process of interaction and exploration. We will now consider how the meaning of a sound is intimately tied to action.

2.1.2 Action-Listening

Touch, movement and audition are bound together in a multimodal way when our performer in Figure 2.2 acts to make a scraping sound. The scraping sound produced by the interaction has meaning for the performer, but also for any potential listener who merely hears the scraping without simultaneously seeing the action creating it. This is because sounds themselves are important sources of information for human actors within their environment. The material quality, size and particular movement of an object all have the potential to be understood from the sound it makes. This is because we directly perceive sounds as events rather than as acoustic phenomena (Vanderveer 1980). Gaver defines this default mode of listening as "everyday listening" (1993*a*, p.286).

In the case of the scraping sound, a listener will perceive it as a scraping *event*, rather than immediately analysing individual acoustical aspects of the sound such as how loud or how high-pitched the scrape is. In his exploration of the potential of designing digital sounds to help convey information to users, Gaver produced a map of these everyday sound events. He breaks them down into components and describes them in terms of a hierarchy of complexity; basic level events like impacts combine to produce more complex temporally patterned sounds like breaking, for example (Gaver 1993*b*, pp.22-24)(Figure 2.3).

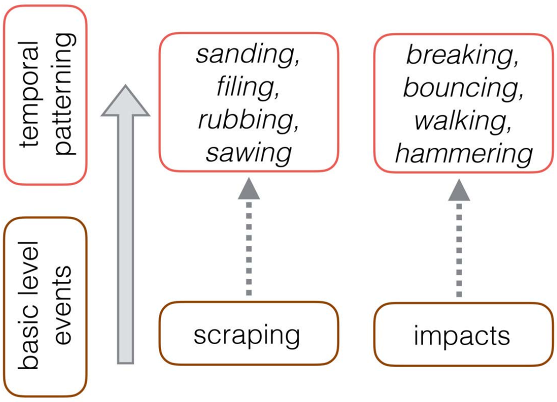

Figure 2.3: Basic level scraping and impact events with temporal patterning, adapted from Gaver (1993*b*, pp.22-24).

There is much discussion around other modes of listening that might focus on acoustical aspects of sounds, which is outside the scope of this research project. It is worth noting here though that everyday sound events can be scrutinised differently in the context of an artistic practice. As described previously, a Foley artist interacting with the crumpled paper ball in Figure 2.2 will direct their movements in a nuanced way to produce exactly the kind of scraping sound they feel is right for a scene. The listening mode here, which examines acoustical aspects of the scrape while perhaps also syncing the sound with onscreen action, is closer to what Gaver terms "musical listening" (1993*a*, p.286). Rather than just perceiving the sound as an event, the Foley artist judges its acoustical qualities in more detail in order to create the sound that they require.

Everyday listening also helps us to acquire embodied knowledge of how sound is connected to events such as breaking, bouncing and scraping. Through daily experience, we build up a lexicon of sound events that facilitates recognition and helps us to contextualise new sonic experiences in light of old ones (Adams & Bigand 1993, p.6). This lexicon is imperfect, however. We are much better at determining what action caused a sound rather than what materials were involved in that action. This was confirmed in a study by Lemaitre & Heller (2012), who asked participants to listen to recordings of scraping, rolling, hitting and bouncing events generated from cylinders made of glass, plastic, wood and metal. Participants were much faster and more accurate in identifying the actions rather than the materials involved. This chimes with other research into human perception of environmental sounds. When asked to compare a wide range of different sounds in a listening test, participants tend to group sounds together into categories based on the events rather than the materials that created them - discrete impacts are perceived as very distinct from continuous sounds like tearing, for example (Gygi et al. (2007) and Houix et al. (2012)).

Human auditory perception is therefore intimately linked with action. This clear connection between sound and action has been evidenced by research into mirror neurons in macaque monkeys. In a study by Keysers et al. (2003), these nerve cells in the brain were shown to fire both when a macaque monkey is performing an action (e.g. tearing a piece of paper), and also when they hear a recording of the sound produced by that action (e.g. a piece of paper being torn). Mirror neurons were later observed in humans, confirming the existence of an action-listening network for the more complex experience of musical performance. Lahav et al. (2007) trained non-musicians to play a simple piece of music, and observed their mirror neurons firing when they heard it played. Playing an unfamiliar piece of music did not result in the same response, suggesting that sensorimotor skills (or embodied knowledge) are crucial to the ability of sound to evoke action. This idea of distinct event-based groupings, and an underlying perceptual structure that determines the meaning of sounds heard, is explored further as part of this research.

We have so far considered how embodied knowledge can be accumulated through action and how human perception of sound is intimately linked to action. We will now move on to explore how objects in the environment, and even a sound itself, might directly suggest a particular action to a performer.

2.1.3 Sound Affords Action

When presented with the crumpled paper ball in Figure 2.2, our performer might understand instantly, based on tacit knowledge, just how they can hold it in their hand and manipulate it. The crumpled paper presents them with a visual opportunity for action. Gibson (1977, 1979) first proposed this theory, suggesting that the environment around us offers specific opportunities to act in order to gain tacit knowledge. He named these opportunities *affordances*. Gibson was an ecological psychologist, a field also influenced by the philosophical school of phenomenology. He theorised that animals, and therefore humans, were skilled in perceiving the palette of possible actions offered to them by their surroundings. For example, a forest environment might afford swinging to a chimpanzee just through the presence of its tree branches and vines. Returning to the crumpled paper example, this object affords holding to the performer through its shape. Once held, the stiff texture and lightweight composition of the paper might afford squeezing and sliding across a surface, or the opportunity for something else - throwing the paper into the air, for example.

The theory of affordances has long been subject to discussion in interaction design. Norman (1999) makes some useful distinctions between affordances, perceived affordances and conventions. He defines perceived affordances as those that designers actually make use of. For example, a desktop computer screen has always naturally afforded lifting, tilting and touching, but this was not part of how a user usually interacted with software functions during this era of computing. Instead, icons placed on the desktop offered visual cues to the user for targeted clicking actions, defining a set of perceived affordances. Conventions, such as software toolbars or the size of a screen, were established in this era of desktop computer systems software. Users were constrained by these conventions, and did not perceive lifting, tilting or touching as affordances in computing. Once design conventions establish themselves, moving on to new ones can cause confusion. Touchscreen and mobile tablet systems are now ubiquitous, and lifting, tilting and touching are all part of computing. Norman reminds us that swiping or dragging are not natural affordances either, but a new set of conventions that need to be carefully deployed by designers (2010, p.10).

In the case of an interactive sounding object, a consideration of its perceived affordances must incorporate both its visual appearance and the sound it produces. Given that sound is so closely tied to action in human perception, it also has the potential to afford action. Previous research has shed some light on this. It has been shown in listening experiments that sounds, when heard without seeing the action that caused them, can be linked directly to bodily actions rather than just sound events. Participants might describe a *thunk* as *hitting* for example. This is again linked to prior embodied knowledge of sound. When sounds are not familiar to participants, they describe them more abstractly, by tracing what they perceive as the shape of the sound in the air (Godøy et al. (2006*b*), Caramiaux et al. (2011) and Caramiaux et al. (2014)). Altavilla et al. (2013) have experimented with the idea that sound has innate sonic affordances that can help design sonic interactions. In a study with non-musicians, they provided three different sounds for the participants to interact with using a wireless sensor attached to one of their hands. This allowed for free hand movements in the exploration of percussion, a rattle, and a continuous oscillating sound. Participants explored each sound in turn and were able to understand how to play them with the sensor. The researchers highlighted that again, familiarity with the sound source was a factor in how participants decided to interact with the sounds. Their choice of a wireless sensor was an attempt to remove any cultural constraints inherent in an interactive sounding object's visual appearance that might have influenced participants' experiences.

Interactive sound's ability to bind itself to sensory modalities, conceptual understandings of events and human actions, and even objects themselves, makes it a challenging material to use in designing a perceived affordance. It is clear though that when interacting with sound, a performer should feel some correspondence between the action that they are performing and the sound it is producing. The next part of this section explores how the experience of this correspondence within a continuous sonic interaction can be understood conceptually from the point of view of the performer directly engaged in soundmaking.

2.1.4 Sounds-In-Hand

In a continuous sonic interaction with crumpled paper (Figure 2.2), the performer's experience of touch, movement and audition corresponds with the knowledge they have previously embodied about actions, materials and sounds through acting to perceive. Familiar modes of feedback in the encounter guide experiential learning in others; touch directing the enactive exploration of movement or movement directing listening, for example. In SID research, the use of sound has so far been investigated as a method of guiding user activity that is directed at a visual target. In this formulation, sound is merely part of the multimodal feedback that guides a user's action towards a visual outcome on a screen. In the case of the crumpled paper example, a participant in a SID experiment might be asked to try to produce a scraping sound of a particular loudness in order to influence a responsive graphical display of the sounds amplitude. The participant would be targeting their interaction to the graphical display. This targeting shifts the focus of the participant away from the tool (the paper) and towards the target (the display).

A focus shift like this is something the philosophy of phenomenology has considered in depth.

In his phenomenological examination of the embodied experience of tool use, Heidegger distinguishes between items that are "present-at-hand" and "ready-to-hand" (Heidegger 1962, p.59). Tools *present-at-hand* are the focus of objective theorising - something is broken or requires a close examination of its workings. Tools *ready-to-hand* are available for use, offering a simple and reliable affordance. Heidegger gives the example of a hammer, which extends the capability of the human body during use, and seems to disappear from the user's perception, as the target of the activity (the head of the nail) becomes the only focus. Intent is focused from the body, through the tool, and towards the target. Returning to the example of the SID experiment, the crumpled paper is *ready-to-hand* during the task as long as it is reliable. If the material qualities of the paper suddenly change mid-gesture (it becomes rubbery in texture, for example), the paper will appear *at-hand* and the user will have to focus on it to understand the new sensory input before being able to return their focus to the task. It is the *ready-to-hand* quality of interactive sounding objects that facilitate enactive learning. This is a very challenging kind of object to design, something that will be considered in detail in the next section of this chapter.

This research examines sonic interactivity specifically in relation to sound performance, and so removes the influence of a visual task on the experience of a continuous sonic interaction. If sound is the only focus of an interaction, then this is where the performer's intention must be directed. It is useful then to extend Heidegger's philosophy to the case of soundmaking. Let us consider the example of the crumpled paper again, but this time it is no longer the means by which a user completes a task focused on a visual outcome. Instead, a Foley artist is using it to rehearse a performance of scraping that will later be used for a scene. In this continuous sonic interaction, the crumpled paper disappears as the object of focus and is instead *ready-to-hand*, extending the performer's intent towards soundmaking. As the body's capabilities are extended into expressive soundmaking, the scraping sound itself could be said to be *in-hand* as the continuously responsive and reliable outcome of each physical adjustment and movement. This research is concerned with this possibility, and examines the potential of designing an interactive sounding object that, through its ready-to-hand qualities, can actually give a performer the feeling that the sound they are making is *in-hand*.

This section has outlined how a continuous sonic interaction is defined in the context of this research. It has explored the potential of enactive learning and the acquisition of embodied knowledge, the perceptual weld between sounds, events and actions, and the potential of sound to afford action and feel in-hand when it is the object of an interaction. The next section will move on to consider the design of interactive sounding objects in more detail. It will also highlight potential issues in trying to transfer qualities of simple, natural and rich experiences of real-world sonic interactions into the digital domain.

2.2 Designing Interactions with Digital Sound

In a continuous sonic interaction, sound has the potential to create meaning for the human user or performer. This meaning is dependent on an array of factors. These are centred on the accumulated embodied knowledge of the user or performer - how the sound and action correspond according to their prior experience, for example - and also the qualities of the object or material being interacted with. The ability of a material like crumpled paper to facilitate an interaction with a scraping sound-in-hand rests on its reliable physical properties and understandable soundmaking parameters. The perceptual experience of *continuousness* the crumpled paper might provide is a huge challenge to recreate in the digital domain. This is due to the fact that digital systems decouple natural connections between action, material properties, and sound that exist in the real world. While this offers huge freedom to the designer, it means that a digital sound must be specifically designed to feel intuitively in-hand during a continuous interaction. This is not a simple proposition. Franinović has formulated an overarching SID approach that directly engages with the potential of ergoaudition to facilitate learning, and designs continuous digital sound purely with action in mind. This research explicitly engages with her proposed method of “enactive sound design,” a design strategy in which sound:

- Affects the sensorimotor activity of the user by building on previously accumulated tacit knowledge, guiding them to learn something new.
- Engages the user’s willed action through ergoaudition with an object that only produces sound in response to a user’s own movement.
- Enhances multisensory experience by not separating audition from the other senses, such as touch.

- Responds directly and continuously to the user's movements. (2013, p.21)

As explored in the previous section, the enactive potential of an interactive sounding object is determined by its ability to be ready-at-hand, or perceivably integrated into bodily movement, during an interaction. Successfully realising this kind of design requires careful consideration of how the performer's hand can feel like it is extending seamlessly into the digital sound. Design strategies must focus on the material properties of the interactive object itself, how that object converts human action into data, how that data activates and modulates sound, and what kind of sound should result from the interaction (Medeiros et al. 2014).

Technical limitations can have a profound effect on the performer's experience. The presence of latency, a delay in the onset of a sound after the performer has moved to activate it (McPherson et al. 2016), can make the sonic interaction feel less continuous, redirecting the performer's intent away from the sound-in-hand and back to the interactive object itself as it reappears at-hand. If technical limitations can be overcome, the choice of sound and activating action presents another interesting challenge to the designer. Replicating familiar configurations of actions and sounds to facilitate intuitive explorations of digital sound may unnecessarily limit the potential of this design space, while a very open approach that does not build on prior embodied knowledge may provoke too much confusion for the user or performer. Somewhere between these two extremes lies the sweet spot - a rich, intuitive continuous sonic interaction with a complex digital sound that facilitates enactive learning.

This section shifts the discussion to how designers in SID and DMI research create continuous interactions with digital sound. Figure 2.4 illustrates a generic setup to facilitate this kind of continuous sonic interaction in order to focus the subsequent discussion. An overview of the challenges inherent in designing a digital interactive sounding object, and potential solutions, will now be presented with reference to relevant research from both SID and DMI design.

2.2.1 Gestural Affordance

When holding a ball of crumpled paper, the performer can hold it and perform a gesture, a movement event with an identifiable duration that communicates energy to create sound (Choi 2000, p.147). We have already considered how the visual appearance of a designed object might suggest a particular kind of action, and how prior experience of design conventions might confuse what kind of affordance is actually perceivable (Norman (1999, 2013)). It is therefore useful to examine how gestural affordances can be defined and described in the design process. This helps to determine how they might correspond meaningfully with a resulting sound.

As outlined in the previous chapter, SID research has focused on gathering potentially meaningful correspondences between actions and sounds and connecting them together in new designs. Work such as Franinović's decomposition of the gestures used to activate mechanical kitchen tools gives many useful descriptors for observing, unpacking and examining what kind of gestures an interactive sounding object might afford. In particular, she defines composed gestures as those comprised of more than one hand action at a time, *hold* and *tilt* together, for example (2009, p.15). Similarly, research into the role of human action and gesture in music performance has established useful descriptors for human movement in interaction with an instrument that has helped to inform the design of new DMIs. Instrumental gesture, specifically targeted at sound creation (Cadoz & Wanderley 2000), has been classified into fundamental movements called primitives (Choi (2000), Cadoz (1988) and Godøy et al. (2006*b*)). These can be force-based (lean, push, pull etc.), and involve excitation to activate sound (hitting, stroking) or modulation to continuously change it (shaking, deforming). While many sound-producing gestures can have an excitation phase with a perceivable beginning and ending (Jensenius & Wanderley 2010, p.11), this research focuses on one continuous gesture of rotation that continually activates and modulates sound. This is discussed in further detail in Chapter 5.

To interact with digital sound, the potential gestures of the hand must be converted into data. This is a challenge for both SID and DMI designers, as available digital interfaces

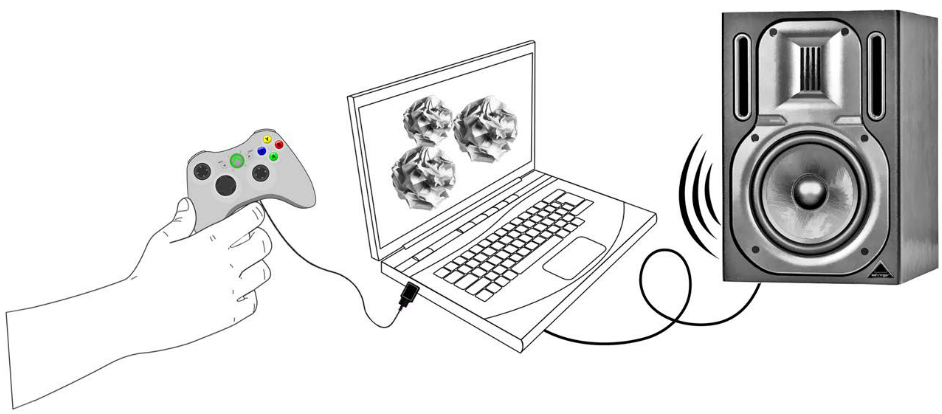

Figure 2.4: A generic setup for a digital interaction with a scraping sound.

from HCI or audio production rarely facilitate truly continuous hand movements. They instead impose design conventions that often centre on precise and restricted movement of the fingertips, and may afford no grasping or holding action in the hand at all (Figure 2.5).

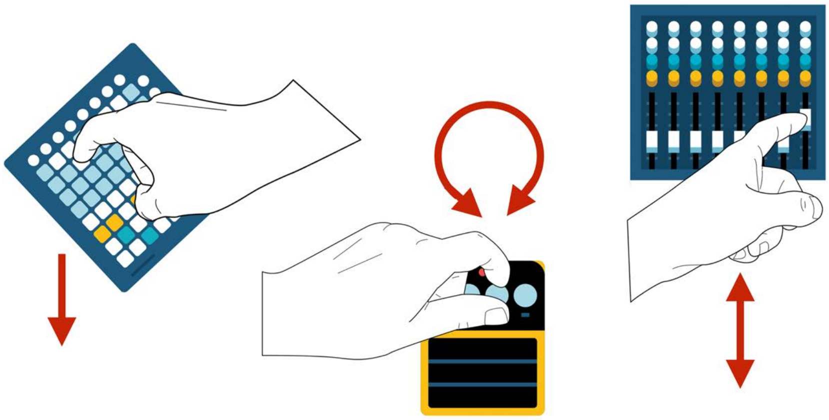

Figure 2.5: Fingertip-based control conventions in audio production (L-R): pushing a pad, turning an encoder and moving a fader.

The search for more gesturally rich and interesting interfaces is an ongoing part of DMI design research, centred on the International Conference on New Interfaces for Musical Expression (NIME) (Jensenius & Lyons 2017), which is too broad a field to be summarised here. However, this research implements some interface design methods from DMI research by configuring a very simple crank handle and optical rotary encoder to capture a rotational gesture as data. The rotary encoder component would usually be activated with small precise turning movements of the fingertips (Figure 2.5), but its gestural affordance is changed when connected to a larger handle. Other researchers have also taken this approach. Ward & Torre (2014) created The Twister, an optical encoder based interface that performers could hold between two hands and twist. Heinrichs (2018, p.136) augmented a potentiometer with a corkscrew handle to create an interface that afforded twisting back and forth in one hand. Gelineck & Serafin (2009) conducted research to compare knobs (rotary potentiometers), sliders (linear potentiometers) and an optical encoder augmented with a simple crank handle in a performance experiment with participants. Participants found the continuous crank handle movement particularly satisfying to perform with. The potential of a crank handle paradigm for the creation of an interactive sounding object is explored further as part of this research.

We have so far considered how the kind of gesture an interactive sounding object affords will impact on its potential to place a digital sound intuitively in-hand for a performer. Next, the impact of touch will be explored.

2.2.2 Touch

The surfaces of naturally occurring materials, such as wood, are texturally complex to the hand and fingertips. Natural materials have been used when constructing DMIs to give the performer more texture to explore (Crevoisier & Polotti 2005). Designers have worked to connect touch, movement and sound together in the creation of new continuous sonic interactions. Essl & O'Modhrain (2006) created two very simple interfaces based on texturally complex materials for the performance of a complex digital sound. Pebblebox affords shuffling, stroking or dropping movements through a box of small smooth stones to interact with a granular synthesizer. Crumblebag, a soft bag filled with a choice of crushable materials, affords grabbing and squeezing when performing the same granular synthesizer.

When performing with an acoustic musical instrument, the player experiences vibrotactile feedback, or vibratory movement that receptors in the hands and body are particularly sensitive to (Rovan & Hayward 2000, p.298). This is another important element contributing to the tactility of the instrument. These complex sensations of texture and vibrotactile feedback must be explicitly designed into a digital interactive sounding object. Haptic technology can add vibrotactile sensations to new interfaces. Small vibration motors have been shown to improve a performer's engagement with a new DMI (Marshall & Wanderley 2011). Commercially available game controllers that provide vibrotactile feedback have also been used as DMIs in and of themselves, such as Hayes' (2011) musical performances with the Novint Falcon. If haptic technology is used effectively, it can make encounters with the more technical materials of plastic and glass more texturally alive under the hand, helping to bind sound to the other sensory modalities in a continuous sonic interaction.

This research makes use of natural materials in the design of an interactive sounding object, and also considers how the perceptual experience of these materials could be recreated in an interface comprised of manufactured parts.

2.2.3 Effort and Resistance

Gestural affordance and touch have so far been explored as critical considerations in the design of a digital interactive sounding object. Closely related to these is the performer's experience of bodily effort and resistance to their movements. A piece of crumpled paper will resist a sliding motion by a performer according to the friction between it and a surface. An acoustic musical instrument resists manipulation by a performer in an even more profound way. Properties like stiffness, string tension, weight or size and shape must be overcome in some way to develop expressive mastery of a material or instrument's soundmaking capabilities. Acoustic and even electroacoustic instruments (such as the electric guitar) also afford "non-linear adjacencies" to the performer, with accurate pitches and unwanted noises separated from each other only by very slight physical movements (Collins 2015, p.27), so the performer must also be highly accurate in their manipulation and articulation.

The multifarious effort required on the part of a performer has been highlighted as crucial to a satisfying and meaningful musical performance (Ryan 1991, p.6). The experience of this can also be found in present-day soundmaking practices that use objects. The design of a scraping sound being developed by an experienced Foley artist may require a very specific grip pressure and movement trajectory, requiring careful rehearsal with the crumpled paper. While not requiring a large physical force, this performance nonetheless requires effort to specifically tailor the movement to produce the desired quality of sound.

The quality of material resistance an interactive sounding object offers to a performer is critical in helping to guide their enactive learning. Designers and musicians have proposed several strategies to foreground effort in DMI design and bring a heightened sensory experience of resistance and more possibilities for articulation to the digital musician. The resistive potential of mechanisms has been investigated, an area of particular interest to this research. Sinyor and Wanderley's Gyrotyre (2005) explores the experience of rotational dynamics in the performance of a digital sound with a spinning bicycle wheel interface. As the wheel rotates on its axis, it gives a sensation of vibrotactile feedback to the performer through gyroscopic precession, which is the phenomenon of a spinning object wobbling back and forth counter to its axis of rotation. The performer is also required to expend effort to move against the rotational inertia of the wheel, which is its ability to store movement energy and resist changes in rotational speed once it is moving. Bennett and O'Modhrian's Damper (2007) uses a fluid brake interface that resists movement as the performer pulls its two handles back and forth. This interface allows different levels of stiffness to be calibrated by applying a current to the magnetic brake. Both of these mechanism-based interfaces can be seen in Figure 2.6. This research explores the potential of another mechanism, a rotating crank-and-axle configuration, as an effortful performance interface.

We have so far considered how the material qualities of an interactive sounding object - its gestural affordance, tactile quality and material resistances - can be configured to ensure that it feels ready-to-hand during performance. Next, strategies for mapping the gestural data output to the activation of digital sound will be outlined.

Figure 2.6: Effortful DMIs Gyrotyre (Sinyor & Wanderley 2005) and Damper (Bennett et al. 2007). Images used here with permission of IDMIL and Pete Bennett respectively.

2.2.4 Action-Sound Couplings and Virtual Mechanisms

The way a user's action, translated to data by an interface, is mapped to the activation of a digital sound can have a significant effect on the experience of continuous sonic interaction. A perceivable "continuum of energy" from gesture to sound (Cadoz 2009, p.218) will make the sonic interaction more immediately meaningful in light of previously embodied knowledge of actions and sounds. As Franinović (2013, p.29) suggests, if the interactive object responds predictably and continuously with sound the performer or user will be able to enactively learn how it works.

The difference between the perceptual experience of soundmaking with real-world materials and digital systems has been explored within DMI research. Jensenius defines the distinction between them as the difference between an action-sound relationship, any kind of connection between action and sound, and an action-sound coupling, which mechanically couples action directly to an acoustic sound in a stable and predictable way (2007, p.22). The experience of pushing a button to activate a doorbell is very different to knocking directly on a door with your knuckles, for example. How far digital action-sound relationships might be configured to replicate the stability of action-sound couplings is an interesting challenge for designers. Different mapping strategies to connect control data with meaningful digital sound creation has long been a focus of DMI research (Hunt et al. 2003). Control data plays a pivotal role in the excitation and modulation of real-time digital sound, and can be programmed to mimic real-world dynamics such as energy loss and accumulation (Menzies 2002). These models can also be focused on controlling the perceptual features of the sound in a physically-inspired way (Heinrichs et al. 2014). Strategies like these might allow even very simple actions to activate and modulate digital sounds in a perceptually meaningful way.

SID research offers an interesting paradigm for designing couplings between actions and digital sounds that informs the work of this research project. The potential of a virtual mechanism has been explored to facilitate a meaningful interaction with a digital sound. The Spinotron affords a vertical pumping gesture to the user, but this movement is mapped to a rotating ratchet mechanism sound model (Figure 2.7). As the user pumps the Spinotron up and down under their hand, they hear the clicking sound of the ratchet turning in response; the two movements are mechanically coupled together metaphorically in a virtual space (Lemaitre et al. 2009, p.977). Another sound model using a rolling ball proved too ambiguous for users to interact with using the Spinotron, suggesting that some virtual mechanical connections between action and sound are innately more intuitive than others. The potential of virtual mechanisms is explored further as part of this research as a way to understand the sonic affordance at the heart of a historical theatre sound effect.

2.2.5 Digital Sound

The material qualities of an interactive sounding object, and the way that data is mapped to activate its sound, are both important factors to consider in the design process. The digital sound to place in-hand as part of the interaction must also be selected with care. This is a particularly interesting challenge for the designer. Digital signal processing offers enormous potential for sound creation, activation and modulation. This vast palette of sonic possibilities encompasses many methods of synthesizing and processing sound beyond the scope of this thesis. What they all share, as digital methods, is a decoupling from real-world events such as *breaking*. There are no innate couplings between human action and digital sound, and this is a design space based on conventions that have evolved from working with sound as signal. As such, designers working to develop perceptually meaningful sonic interactions have looked for ways to synthesize digital sound that models real-world physics.

This research explores the potential of this method of synthesis as a way to design a continuum of energy between action and sound. Heinrichs (2018, p.30) encapsulates the two main methods in this area of digital synthesis with the useful term “computationally generated audio”. These approaches - physical modelling synthesis, and what he terms “practical synthesis” - both create sound as the result of behaviour or process. Physical modelling synthesis creates sound using mathematical expressions of real-world physical phenomena. This method was first used extensively to model the components of acoustic musical instruments, such as a plucked string (Karjalainen et al. 1998) or the dispersion of that string sound through the material of an instrument body such as a violin (Smith 2010, pp.40-41). These models were originally computationally expensive, but have become increasingly available to designers and musicians for prototyping in accessible audio programming environments. Ronan’s (2010) work to program a complex physical model of a sitar (an Indian string instrument) entirely in the graphical programming environment Max/MSP is a good example of this. This research implements some of these more classical methods of modelling musical instruments as part of a prototype of a digital theatre sound effect.

Practical synthesis methods take a much broader view of modelling, focusing more on the perceptual aspects of the sound being created or the real-time efficiency of the model itself (Heinrichs 2018, p.30). These models may not use exact mathematical expressions, but try to produce the same effect as the sound event they are trying to model. In this way, this digital approach is close to the practice of creating theatre sound effects themselves, as will be explored further in Chapter 3. One particularly useful method of practical synthesis that directly informs this research is Farnell’s (2010) “procedural audio” approach, which programs sound from the point of view of events and processes. This method observes, unpacks and describes complex systems to allow their individual sound-producing components to be programmed.

This research implements a procedural approach to sound modelling, and also makes use of a practical synthesis toolkit developed specifically for SID research. The Sound Design Toolkit (SDT) is a suite of objects created for the graphical programming environment Max/MSP (Baldan et al. 2017). This toolkit was produced by the Sounding Object (SOb) project (Rocchesso 2014, p.160), which developed a combined approach for digital sound synthesis that incorporates aspects of physical modelling synthesis and practical synthesis together with a focus on Gaver's basic level components of everyday sounds (1993*b*, pp.22-24). This allows large categories of sounds to be synthesized and controlled in a simple way - a friction model informs higher-level events such as creaking, squeaking, rubbing and scratching, for example (Figure 2.8). The underlying physical models are simplified, with important features of the resulting sounds exaggerated to ensure that they are clearly perceived and understood in an interactive context (Rocchesso et al. 2003, p.2). The work undertaken to create a prototype digital model of a theatre sound effect as part of this research is explored in Chapter 6.

This section has explored how real-world continuous sonic interactions might be made digital, highlighting some potential design problems and the solutions so far proposed by SID and DMI researchers. Creating an interactive sounding object that offers a performer a seamless and expressive experience of soundmaking requires an understanding of how each aspect - gestural affordance, touch, effort and resistance, virtual mechanism and resulting sound - influences how the design might build on previously embodied knowledge and facilitate enactive learning. The challenge inherent in this process, and explored further within this thesis, is how to connect the digital sound to the intuitive movement of the hand in a perceptually meaningful way. The next section considers some of the challenges of evaluating a continuous sonic interaction.

2.3 Evaluating Continuous Sonic Interactions

Evaluating a continuous interaction with a digital sound in an experiment with participants is a complex undertaking. While the design of the interactive sounding object may be based on a certain perceived affordance, with the aim of creating a certain meaning for a performer, new meanings emerge according to the individual embodied knowledge each performer brings to the encounter. Features the designer intended to be important may be less so in another's experience, for example. Technical features of the design of the interactive object and its production of digital sound (such as those previously outlined in Section 2.2) may overly influence the experience, and cause the interactive object

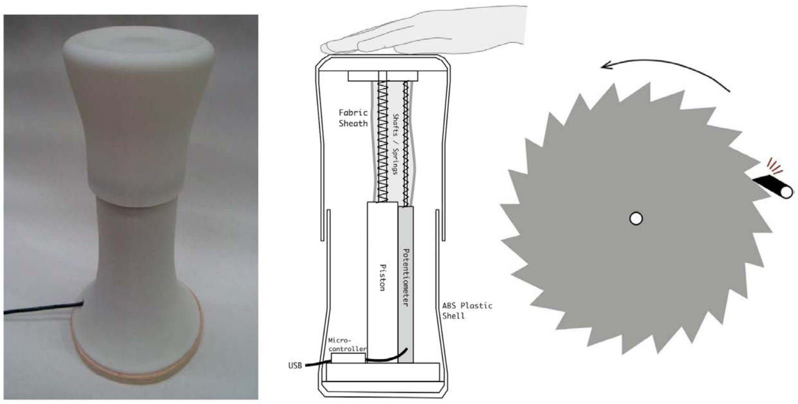

Figure 2.7: (L-R)The Spinotron, its internal mechanism and the virtual mechanism of its sound model. Images used here with permission of Karmen Franinović.

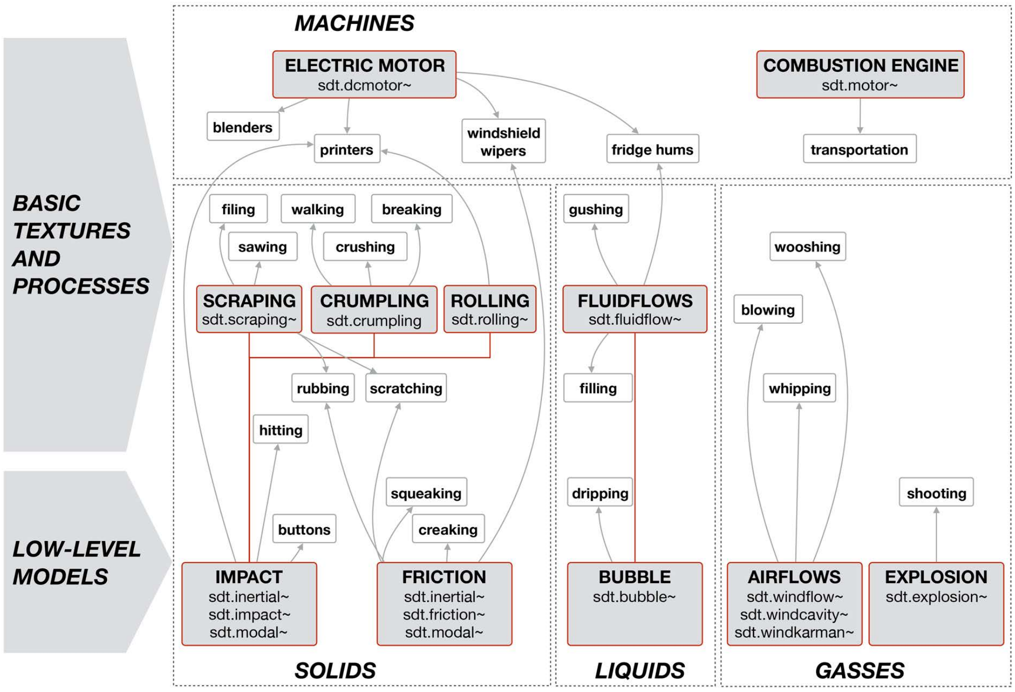

Figure 2.8: The SDT taxonomy (Baldan et al. 2017). Included here with permission of Stefano Delle Monache.

to be present-at-hand (Heidegger 1962, p.59), potentially confusing or disrupting the experimental task. Because of these issues, evaluation is a very important part of the design process. It can generate useful information to refine and improve the interactive sounding object, and can also create knowledge about the perceptual role of sound in interactions. However, the experiment must be designed in way that elicits useful information from participants, while also engaging them in the process of interacting with sound.

SID and DMI design approach evaluation in different and developing ways. SID focuses on the use of sound to produce a target behaviour or specific perceptual experience in the user (Giordano et al. 2013, p.151). DMI design works to examine the more elusive metric of musical expression in performance with interactive digital systems. How an evaluation is best directed presents another kind of design problem. Empirical methods of evaluation use objective and replicable measurements, producing numerical data that can be statistically analysed and verified. Qualitative methods focus heavily on participant perspectives, and use methods such as interviews to obtain rich and more personally situated data. In the case of an interactive sounding object, a highly targeted experiment designed to empirically evaluate it may be too narrowly focused, while a very open and exploratory setup may miss the opportunity to establish useful empirical findings.

SID's evaluation approach attempts to mitigate this by bringing more objective methods of data gathering such as acoustical analysis and motion capture together with perceptual information from participants obtained through rating tasks, all of which can be statistically evaluated to produce findings applicable to further design work. DMI design evaluation is usually much more qualitative, focused on participants with a musical skillset in the examination of specific designs in detail, with less emphasis on the generalisation of its findings. There is in fact little agreement on the criteria for evaluation among DMI designers (Barbosa et al. 2015). There is similarly no single clearly defined procedure within SID for evaluating a newly designed interactive sounding object. A synthesis of formal methods must be applied in an exploratory way. This can help to establish a useful approach to evaluation, even in the case where the findings are not entirely conclusive.

Evaluation is an important part of this research, and the experiment conducted to investigate the sonic interactivity of a theatre sound effect is fully reported in Chapter 7. The design of the experiment has been closely informed by SID research, but also borrows from appropriate DMI evaluation methods. This section considers the main issues in this kind of evaluation with reference to relevant prior research from both fields.

2.3.1 Participants

Given that embodied knowledge is such an important factor in the experience of a continuous sonic interaction, the selection of participants for an evaluation is an important influence on the potential of the process to produce knowledge that can be put to use elsewhere. SID evaluation methods are focused on the generalisation of findings, and so incorporate some simple approaches centred on everyday sonic interactions. These include asking participants to produce sounds with everyday tools, or performing simple scraping actions much like the continuous sonic interaction introduced in Figure 2.2 (Giordano et al. 2013, p.153). These methods, which produce useful findings about human perceptual experience of sound in an interaction, are easily extended to participants with non-specialist experience of soundmaking.

By contrast, evaluation of DMI designs from the perspective of musical expression is often focused on participants with significant musical ability. Although beyond the scope of this research, it is noted here for the reader that musical expression in the performance of an instrument has been defined as comprising of the technical skill, motor precision and style of the performer, as well as the acoustic parameters and instrument-specific characteristics available to them (Juslin 2003, p.273). As we have explored earlier in this chapter, Foley artists also acquire significant embodied knowledge in sound performance with objects. There is also a focus on expression in Foley practice, as onscreen events and characters are interpreted and brought to life through sound performance (Pauletto 2017, p.345). These expressive skills can take many years to acquire. This kind of participant can prove very useful for a designer in determining how to improve the finer details of the quality of the performance experience afforded by a continuous sonic interaction. Heinrichs' (2018, pp.140-173) study with Foley artists asked them to perform digital sounds for an animation, for example.

Although experienced participants can bring a lot of useful embodied knowledge to the evaluation tasks, an interactive sounding object design can still confound their expectations if it does not build on their specific experience. It is also challenging to recruit enough suitably experienced participants for a study if the aim is to produce statistically significant results. The experiment conducted as part of this research aims to establish some statistically significant findings, and as such focuses on a larger group of participants with a broader range of experiences of music and soundmaking. As Chapter 4 discusses, the experiment is designed to allow inexperienced participants to engage in soundmaking and explore expression in the performance of a sound effect. This approach is informed by the historical sources on late nineteenth and early twentieth century theatre sound effects discussed in Chapter 3, which confirm that sound performance in theatres was often delegated to stagehands without a specific soundmaking practice of their own. This suggests that the interactive sounding objects constructed to produce sound through performance at this time were inherently learnable through rehearsal, and therefore enactive. The potential of an interactive sounding object to enactively teach expressive soundmaking is something that this research explores.

2.3.2 Controlling Experimental Conditions

Evaluation in an experiment ordinarily requires careful control of any effects that might influence the outcome being measured. If participants are asked to repeatedly perform a digital scraping sound with the setup in Figure 2.4 for example, and their movement is being measured for speed, then it is crucial that the experimenter controls the setting so as not to unnecessarily influence participants' speed of movement. This might involve asking participants to each sit at the same chair and the same desk for testing, and use the same digital interface to control the sound, for example. Controlling experimental conditions is particularly challenging in the case of a continuous sonic interaction. Different modes of everyday sensorimotor feedback are difficult to disentangle from each other. Determining the influence of one mode in particular (for example, a change in sound only) requires careful experimental design and an understanding of the potential issues caused by the interactive sounding object in use. Making digital sonic interactions may overlook important factors in human perception, particularly when these factors conflict with established audio equipment and workflows. Complex experimental setups can be required.

A SID study by Giordano et al. (2012) explicitly engaged with these issues. It examined whether participants could identify the kinds of materials they were walking on under different multisensory feedback conditions. These conditions were strictly controlled. Participants were first blindfolded to eliminate visual feedback. The specific influence of vibrotactile feedback when walking was measured. This was provided to participants through motors worn in the soles of a pair of specially designed shoes. As these motors made audible noise when they were vibrating, participants had to wear headphones playing white noise during the task to block out the sound of the motors. In another step, the vibrotactile feedback was removed. Researchers ensured that participants wore exactly the same shoes, with the motors switched off, for this part of the experiment so that the experience of wearing the shoes themselves would not influence the results. In another step of the experiment, participants listened to recordings of footsteps. These recordings had been made with a microphone positioned close to the feet however, and so not from the perspective of a walker listening to their own footsteps. Researchers highlighted that this small oversight may have influenced the results.

The interactive sounding object designed for this research specifically allows one mode of sensorimotor feedback - the sound - to be isolated from the others (touch, vision and movement). This allows the experiment design to focus on evaluating only the influence of a change in sonic feedback on participants' experiences. The interactive sounding object design approach, and the experiment design, is discussed in more detail in Chapter 4.

2.3.3 Task Design

The question of what constitutes an appropriate task with which to observe, somehow capture and then evaluate a continuous sonic interaction is complex. Task design in SID research focuses on evaluating the potential of sound in an interaction, and so is often based around a learning scenario that involves the use of sound rather than a direct evaluation of the interactive sounding object design itself.

The SID approach is highly influenced by HCI methods. For example, a participant's focus is often directed towards a visual task, while sound only informs the interaction as a part of the multisensory feedback. Rath and Rocchesso's Ballancer (2005) presented participants in their study with a screen display showing a virtual marble on a virtual track. A digital interface was created which replicated the virtual track as a controller. The track controller was activated with a tilting action. When the controller was tilted, the virtual track onscreen also tilted, which moved the virtual marble along it. This allowed participants to interact with the virtual marble and track on the screen. The study showed that sonic feedback provided by a physical model of a rolling ball allowed participants to understand and learn how to tilt the track interface to keep the virtual marble in position on the screen. Similarly, the Spinotron was evaluated through a task where participants pumped the interface at speed in order to control a visual display on a screen. Participants who received sonic feedback as part of the interaction were found to have performed better in learning to control the visual display (Lemaitre et al. 2009).

By contrast, DMI research has examined both the usability or controllability of a design, and the experience of expression and potential for virtuosity it offers (Kiefer et al. (2008), Dobrian & Koppelman (2006)). Simple tasks informed by evaluation methods from HCI have been used to evaluate DMIs (Young & Murphy 2015), but the possibility of engaging participants in the expressive potential of these interactions is still considered to be important (Stowell et al. 2009). Poepel (2005) offers an interesting approach through his "operationalizing" of musical expression. His study evaluated three viola-based DMI designs with expert musical performers. The performers were given musical tasks to perform with the instruments based on a list of cue-groups he devised from Juslin's (2003) definition of musical expression. The cue-groups focused on the accuracy of the DMI in relation to timing, pitch, dynamics, articulation and timbre. Participants performed musical tasks devised from these cue-groups and also rated the DMIs according to these criteria. Participants were also asked some general questions about each instrument, including specifically whether the action-sound relationship was adequate and whether they had managed to play the sound they had intended to. As will be explored in Chapter 4, this approach has been adapted for the experiment conducted as part of this research in order to direct participants to perform with an interactive sounding object and elicit their views on their performance.

This section has discussed some of the inherent challenges in evaluating of digital interactive sounding objects in experiments, and presented some possible solutions from the fields of both SID and DMI research. The skillset of participants, the control of experimental conditions and the kind of task that directs the evaluation can all contribute to the usefulness of experimental findings. This research approaches evaluation through a synthesis of quantitative and qualitative approaches, which are outlined in Chapter 4.

2.4 Theatrical Sounding Objects

This chapter has so far examined human perceptual experience of real-world continuous sonic interactions, and how actions that produce sound facilitate enactive learning and the accumulation of embodied knowledge. The perceptual connection between action and sound has been explored, including the potential of a sound to evoke action and even a specific gestural affordance. The main challenges of building on this perceptual connection in the design of interactive objects that produce digital sound were also outlined, along with design solutions from SID and DMI research that might help to make a digital sound feel in-hand for a performer. These strategies include control over the qualities of the interactive object itself (gestural affordance, touch and resistance), the mapping of data produced by the performer's action, and the design of a meaningful and responsive digital sound through modelling synthesis approaches. Issues in designing a formal evaluation of an interactive sounding object were also presented, again with potential solutions offered by SID and DMI research.

As outlined in Chapter 1, this research engages with the historical practice of theatre sound effects design and performance from the late nineteenth and early twentieth century in order to expand on current approaches to designing interactive sounding objects and evaluating continuous sonic interactions. This section introduces the historical practice of theatre sound effects design and performance in the late nineteenth and early twentieth century, and argues that it has the potential to make a meaningful contribution to the design and evaluation of interactive sounding objects. The reasons why this historical context of soundmaking is so interesting are explored. The practice that created theatre sound effects, which will be examined more fully in Chapter 3, is explicitly connected to the theoretical background and prior research so far discussed in this chapter.

Theatre sound effects were once created exclusively with actions, mechanisms, and materials long before sound could be created or shaped as signal or data. Present-day Foley artists also work exclusively with the manipulation of objects to create sounds for film. Through the performance of movements with everyday items such as clothing or cutlery, Foley artists build up sonic layers within a film soundtrack that focus on heightening the dramatic texture of the narrative (Wright 2014, p.205). The embodied knowledge of soundmaking with objects inherent in Foley practice (Pauletto 2017) has been suggested as worthy of investigation as a potential resource for SID research (Franinović 2013, p.19). This work has already begun with Heinrichs's (2018, pp.140-173) study to evaluate different configurations of performance hardware and software in the creation of digital sounds for an animation. Foley artists were used as expert participants in this study and invited to perform with various digital interfaces to interact with a digitally synthesised model of a creaking door. Given that present-day Foley artists possess such significant and valuable experience in the performance of sound with objects, and could potentially participate in this research, it is useful to outline here why this enquiry has instead focused on a historical context of soundmaking. The term Foley originates from Jack Donovan Foley, a sound effects operator for film who was the first to add the sound of his footsteps to a soundtrack as part of the film Showboat (1929) (Ament 2009, p.7). Although centred on sound creation for film, the term Foley has come to denote the use of objects to create sounds through performance in any context in common parlance. This reflects that fact that performative soundmaking is really no longer part of our popular live entertainment forms. Foley, however, is by no means the only word on this practice. While the work of present-day Foley artists has already been connected with that of sound effects performers for silent cinema at the turn of the twentieth century (Wright 2014, p.206), this research connects these cinematic practices to an earlier history of sound effects performance in theatres.

The creative use of sonic effects in theatre dates back to 1 BCE (Brown 2010, p.158), and its methods evolved over centuries through various different sites of live performance, only latterly becoming part of media such as radio and film at the turn of the twentieth century. Traces of theatre's influence and methods can sometimes be seen in contemporary Foley practice, particularly when it engages with contraptions or mechanisms to create sound. As Chapter 3 of this thesis will explore, theatre practitioners in the late nineteenth and early twentieth century shared some of the central tenets of present-day Foley practice, such as an understanding of the importance of synchronisation with the onstage action and a concern with the dramatic potential of sound. As theatre sound effects were first designed before the age of sound as signal, all sounds were required to be performed and activated live, along with the onstage action. However, these sounds were not always performed by artists with the kind of embodied soundmaking expertise that Foley artists accumulate today. Instead, the effects could be constructed and refined by a stage carpenter before being passed to other stage workers, who were instructed in the use of the effect, and rehearsed the required performance to enactively learn their part until the director was satisfied with the result. As outlined in Chapter 1, the embodied expertise of sound effects design and performance was diffused throughout several roles in the theatre. As such, it was the interactive sounding objects themselves that extended the potential of the human body into complex soundmaking through mechanical methods. These objects designed specifically to perform sounds were created to be learned through rehearsal. It is this enactive quality of theatre sound effects that is so useful to this research.

By taking a step away from the direct manipulation of materials, and towards the design of interactive sounding objects, these historical practitioners were already grappling with a central design problem of SID and DMI design research; how the interactive sounding object can give a perceptually meaningful experience of performance to the performer. As will be explored in Chapter 3, historical theatre practitioners continuously highlight the ability of particular sound effects designs to directly teach the skill of performative soundmaking. The embodied knowledge of soundmaking that resulted in these designs is therefore still embedded in the interactive sounding objects. In the case of present-day Foley, the soundmaking knowledge is instead embodied within the individual artist creating the sound. This is accumulated over years of selecting and configuring materials and actions, and then performing sound in sync to picture. There is much to explore in the case of the sonic interactivity of Foley practice. Potentially interesting investigations into the evolution of sound performance with actions, materials and objects could also be undertaken in collaboration with Foley artists, given the origins of their practice. By focusing on the inherent sonic interactivity of late nineteenth and early twentieth century theatre sound effects practice, this research begins that work. Figure 2.9 traces theatre's potential historical influence on other interesting sites of performative soundmaking with objects. The design and performance practice that created of theatre sound effects in the late nineteenth and early twentieth century is explored in more detail in Chapter 3.

Historical theatre sound effects present a body of design work for sound performance that is bound to real-world materials and human actions. The potential of this approach to inform the design and evaluation of continuous sonic interactions with digital sound will now be considered. Figure 2.10 revisits the continuous sonic interactions first presented in Figure 2.2 and Figure 2.4, this time to compare them with a generic theatre sound effect configuration that mechanically activates sound.

Theatre sound effects represent the solution to a particular design problem - that of a complex continuous sonic interaction in the acoustic domain. Rather than the performer's hand grasping the soundmaking material (in this case, our crumpled paper) to directly manipulate it, a mechanism couples a grasping action performed on a handle, which is then used to perform a gesture (push, pull, twist, etc.). This gesture is coupled to a mechanism that translates the movement to another moving component. This moving component is what directly engages the soundmaking material to produce sound. As evidenced in this thesis, this configuration of action, mechanism and soundmaking material meets the criteria of an "enactive sound design" (Franinović 2013, p.21). The performer's movement directly and continuously creates sound, sound is only produced in response to

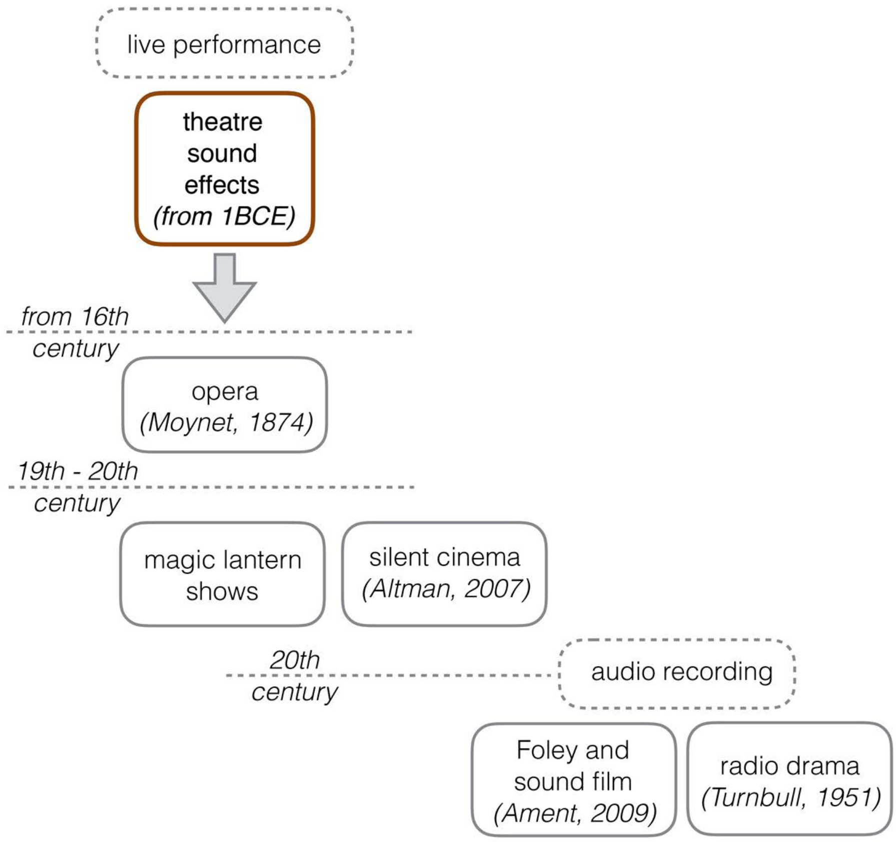

Figure 2.9: Making connections between historical sites of performative soundmaking with objects.

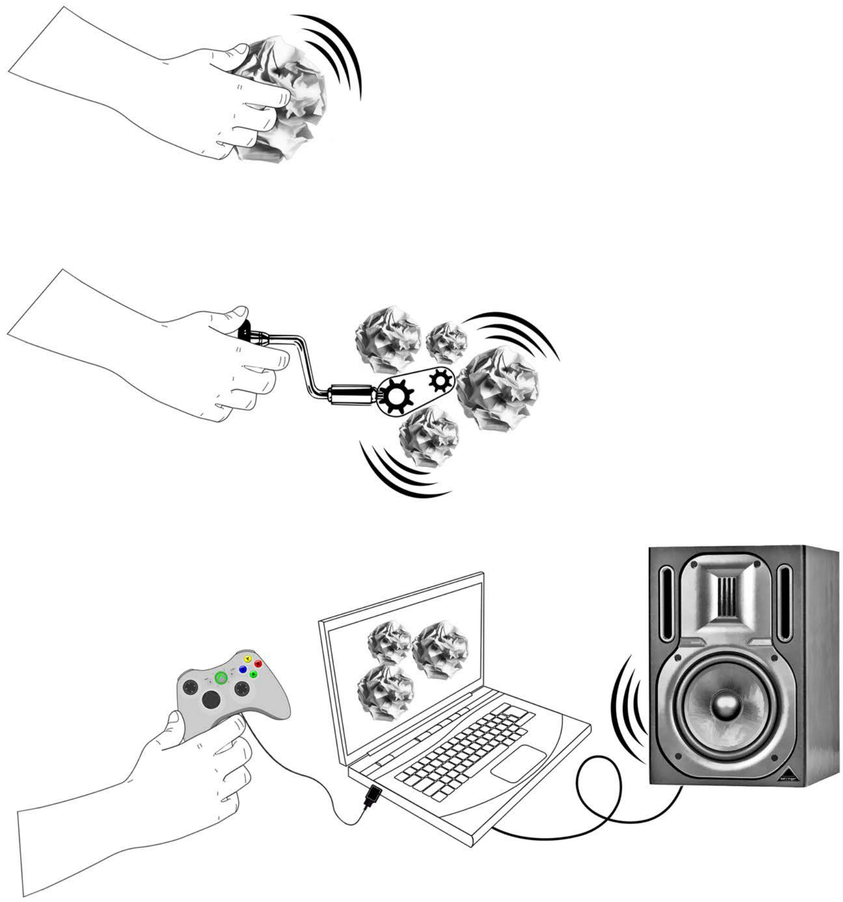

Figure 2.10: Comparing sonic interactions with simple materials (top), mechanisms and materials (middle) and digital interfaces (bottom).

the performer's movement, sound is not separated from other modalities such as touch, and reliable sonic feedback guides the performer in learning how to use the mechanism. A real-world continuous sonic interaction is therefore mediated by a handle and mechanism, potentially without any loss of multisensory experience, but with an increased potential for complex soundmaking. This approach, which increases the potential complexity of the resulting sound while simultaneously simplifying the action required to continuously interact with it, is why these historical designs are so interesting. Unpacking this process with a view to designing continuous interactions with digital sound may offer the key to a *sound-as-action* approach to the configuration of hardware, mapping and digital synthesis, and to the way encounters with these objects can best be evaluated. These aspects will now be explored in more detail.

2.4.1 Designing from Sonic Affordances

Historical theatre sound effects were created through an analysis of real-world sound events and an understanding of simple mechanisms. Designers prototyping these must have in some way been able to imagine sounds as the potential result of these mechanisms activated by human actions. Sound must have conceptually been inseparable from movement or human action in the design process. This is very different to contemporary sound design methods that work across areas of sound performance, recording, signal processing and programming. How the sounds directly suggested actions to the designers through their "sonic affordances" (Altavilla et al. 2013) is very interesting, as exploring this may offer more insight into how sounds are perceptually linked to human actions. It may also be possible to understand how certain sounds could have directly informed the design of objects to afford interactions with them.

Engaging with this historical practice opens up a collection of already realised configurations between actions and sounds; a ready-made design space for the exploration of continuous sonic interactions with objects. Examining how these interactive objects link action to sound might reveal some particularly meaningful couplings. Perhaps a *tilting* gesture makes more sense to a performer when activating *scraping* or *creaking* rather than *crushing*, for example. If perceived sonic affordances directly inspired the designs of particular historical mechanisms to mimic sounds, perhaps digital sounds could also eventually help to design interactive objects. This work could proceed from design exercises with very simple sounds like *crushing*, but eventually uncover more broad methods of making meaning in digital sonic interactions, allowing the designer to build skills in how to offer affordances for skill development in encounters with more abstract digital sounds and textures. This could produce some robust digital interactive sounding objects that afford expression with digital sound in a meaningful way by building on previous embodied knowledge of continuous sonic interactions. This may make digital sound more accessible as an interactive medium for non-specialists. It could also allow musicians skilled in acoustic instrumentation, and Foley artists, to explore the potential of digital sound more fully as part of their own embodied musical or soundmaking practice.

2.4.2 Acoustic Interfaces and Mechanisms

These historical methods of soundmaking made use of real-world materials such as wood, but use them to construct a control apparatus that mechanically activates and modulates sound in response to a simple action performed upon a handle. As acoustic controllers, these devices offer the tactility, vibratory response, and mechanical resistance to movement that acoustic musical instruments do, but as a simple handle-based interface. They are therefore a move away from the complexity of acoustic musical instrumentation, and simultaneously a move towards the simplicity of HCI controllers that ordinarily capture human actions as data. Despite this simplicity, they preserve a direct coupling between action and resulting sound.

This presents an interesting area of potential investigation that is explored further in this thesis. In particular, the simplicity of the activating actions means that these historic acoustic interfaces can be recreated as digital "enactive preservations" (Serafin & De Götzen 2009) based on contemporary electronic components like potentiometers. DMI design strategies, such as those discussed in Section 2.2, could be adapted to this work. With a historical design to work from, different versions of the same interface could be made to explore whether a particular kind of feel of object was more satisfying when performing the same gesture to create the same sound.

These mechanical configurations are not the result of very broad design possibilities, like those afforded by digital systems, but are instead realised according to the limitations of real-world materials. Action and sound are directly and continuously coupled together via a simple mechanism in these designs. As such, historical theatre sound effects meet the definition of action-sound couplings (Jensenius 2007, p.22). If the mechanisms that facilitate these couplings can be examined in further detail, they may offer a way to program mappings for digital sound performance as models of kinematic processes, taking account of movement, force and inertia in the control of digital sound. This paradigm has the potential to expand the potential of virtual mechanisms (Lemaitre et al. 2009, p.977) and couple simple actions to complex digital sounds in a meaningful way.

2.4.3 Effects-in-Hand

Historical theatre sound methods are not just concerned with the production of sounds using materials and mechanisms; they produce *effects* that model sounds produced by real-world events. A crumpled piece of paper affords a simple and direct experience of ergoaudition when interacting with a scraping sound. By contrast, theatre sound effects invite the performer to interact directly and continuously with a much more complex sound produced by a configuration of different materials - rain or wind, for example. These methods build on the perceptual experience of sounds as events, but achieve a sonic complexity beyond basic level sound events like impacts or friction (Gaver 1993*b*, pp.22-24) in their final designs.

This presents two main interesting design questions. The first is that attempting to modelling these effects digitally using a combination of low-level models might bring an added layer of complexity to the synthesis methods offered by programming tools such as the Sound Design Toolkit (Baldan et al. 2017). How this could be realised, and what potential this method has, is a question that this research engages with. As will be explored in Chapter 4, the programming work undertaken to model a theatre sound effect digitally is approached in a creative and exploratory way informed by a practical synthesis (Heinrichs 2018, p.30) and procedural audio (Farnell 2010) approach with the aim of creating an effective sonic illusion activated by human action.

Secondly, if the crumpled paper can place the scraping sound *in-hand* during a continuous sonic interaction, does a more complex sound such as a rain effect similarly feel like it is in-hand during the use of a historical design? If the interactive sounding object perceivably disappears during the course of an interaction and focuses the performer's intent beyond the handle towards the sound, does it also focus the performer beyond the soundmaking material (e.g. dried peas and metal) towards the effect (e.g. rain)? In other words, do these effects work on the performer as well as on the audience during performance? If so, this may also help to develop a method of creating interactions with very abstract digital sounds and textures that feel very much in-hand, and are perceived as the direct result of a performer's action.

2.4.4 Evaluation

Theatre sound effects are simpler in operation than musical instruments, but afford continuous sonic interactions with complex sounds. Performance with these interactive sounding objects was learned only through rehearsal, making them enactive sound designs. They are very simple to operate, but allow the performer to develop some skill in soundmaking through bodily action. This combination of simplicity and expression opens up the possibility of using these designs to study continuous sonic interactions with digital sound using participants without prior musical experience. A design could be evaluated to discover how it helps an inexperienced performer to learn how to express with digital sound, for example. This could also help to discover more about whether digital sound can be placed in-hand for performers with little experience of soundmaking, revealing more about how embodied knowledge of soundmaking is developed. These designs could also help create interactive sounding objects that make use of Foley artists' embodied knowledge of soundmaking, making digital sound a part of their practice.

Theatre sound effects may also help to understand how experiment design can directly engage with the task of soundmaking in the evaluation of continuous sonic interactions, rather than giving participants a visual target for their interaction. As a body of work that links actions, mechanisms and sound together, these interactive sounding objects may facilitate a more rigorous examination of how meaning is created for the performer during the interaction as they extend their intent beyond an everyday sound like a scrape, and instead place their hands into the activation of more complex events like rain. The simplicity of these designs and their functionality as interactive sounding objects allows the perceptual space between acoustic and digital soundmaking to be more clearly mapped and examined. Removing particular components of the action-sound coupling, such as shifting from an acoustic to a digital sound during an evaluation, will enable a closer examination of what factors might contribute to the sensation of sounds-in-hand.

This section has presented an overview of the potential of historical theatre sound effects design to facilitate further research into the design and evaluation of interactive sounding objects. The reasons for investigating this historical practice of soundmaking rather than present-day Foley methods were presented. Some particularly interesting areas that are explored further in the course of this thesis were outlined. This chapter will now be summarised in full.

2.5 Summary

This chapter has introduced the theoretical background to this research project, outlining the main factors in human perceptual experience of real-world continuous sonic interactions, and how SID and DMI designers have explored the most effective ways to facilitate similarly rich experiences in interaction with objects that make digital sound. An overview of the inherent challenges in evaluating interactive sounding objects was presented. The historical practice of theatre sound effects design was also introduced, and its potential contributions to this design and evaluation space were outlined.

In Chapter 3, the historical context and practice of theatre sound effects design in the late nineteenth and early twentieth century is explored in more detail. A design framework for the creation of new interactive sounding objects based on historical theatre sound effects designs is presented. The theatre sound effect design that forms the basis of the main design and evaluation work is also explored in detail.

Chapter 3

A Hidden History of Sound and Action

Strategies for the creation and manipulation of sound existed long before the advent of sound recording or digital media (Curtin 2014, p.xii). These strategies, which extended the capabilities of the human body into complex soundmaking with actions, mechanisms and materials, can be traced back to an origin in theatre performance. This chapter presents an overview of the historical practice of theatre sound effects design in the late nineteenth and early twentieth century, and evidences how designers created interactive sounding objects that could be learned through rehearsal in order to build an embodied knowledge of soundmaking. The most widely used historical sound effect methods are also examined to discover their action-sound correspondences, which are then applied to a framework for the creation of new digital interactive sounding objects. One particularly interesting theatre sound effect, a wind machine, is then explored in detail.

Section 3.1 explores some of the reasons why the interesting qualities of theatre sound effects have been overlooked in previous scholarship. The scope of this historical enquiry is established, and the frameworks used to explore historical sources and build a coherent picture of the practice are introduced. Section 3.2 presents an overview of how theatre sound effects were designed and performed in the late nineteenth and early twentieth century. Cross' "designerly ways of knowing" (2006, p.29) are used as a framework to draw out what practitioners knew about sound and perception, as well as examining some of the embodied knowledge and skills implied by the accounts in historical sources. Section 3.3 explores some specific designs for the most widely used sound effects in the late nineteenth and early twentieth century, as corroborated by several historical sources. Descriptions of these effects are examined with reference to their activating actions, mechanisms, materials and sounds. The Sound Design Toolkit (SDT) (Baldan et al. 2017) taxonomy of low-level sound events and basic sound textures is expanded using this information. This opens up a new design framework of action-sound configurations, harnessing the digital synthesis algorithms of the SDT to realise enactive recreations of historical theatre sound effects. Section 3.4 describes the workings of one of these sound effects, a wind machine, in detail as a particularly interesting design to realise digitally. As will be outlined in Chapter 4, the wind machine design is used as the basis for further work undertaken to answer the research questions. Section 3.5 summarises the discussion presented in this chapter.

3.0.1 *Effects* Terminology

As outlined in Chapter 1, exploring the history of soundmaking in theatre requires us to engage with texts discussing sound and performance from an era before sound recording was ubiquitous. The approach taken to historical terminology is briefly restated here to aid the reader. The historical term *effect* refers to an impression created in an audience (Krows 1928, p.114), but is deployed by historical practitioners as if they were referring to an object to construct and perform with. The term used in this chapter is *theatre sound effects*. As the design and performance work of sound effects is often attributed to many different roles within the theatre, the general term *designer* is used to refer to the activities around their design and construction, and *performer* is used to specifically refer to a sound effects performance. The general term *practitioner* encapsulates both aspects of the practice.

3.1 Historical Method

This research positions theatre sound effects design practice from the late nineteenth and early twentieth century as a useful source of knowledge for the design of interactions with digital sound, and connects this practice to other potentially interesting sites of performative soundmaking such as silent cinema and present-day Foley. Investigating the history of theatre, tied as it is to live performance and sensory experience, has long proved a challenge for historians (Nagler 1952, p.ix). More recent interdisciplinary enquiries have helped to shed new light on performance practices that have not been well documented in texts. Previous research approaching theatre and performance history from the perspective of sound studies and audio has explored how the environments of performance spaces were a critical part of audiences' sonic experiences, for example (Smith (1999), Lopez (2015) and Murphy et al. (2017)). This research extends this work by converging SID and DMI design research and methods with theatre history.

However, examining theatre sound effects as a single body of work through these present-day design perspectives is challenging. Historical sources do not accord sound a particularly special status in theatre productions. Sound was merely one part of a range of effects (Krows 1928, p.114), and historical theatre practitioners approached sound design and performance work with this sensibility. Within the context of theatre, the design problems of sound performance are resolved in between other priorities in production, and soundmaking expertise is diffused throughout several backstage and production roles. Important considerations for this research, such as how action-sound configurations were decided upon, or how an interactive sounding object was refined and made to work using carpentry techniques, were not well documented historically. Theatre also found its soundmaking expertise marginalised in the early twentieth century by a new aesthetics of modernism, and by technological advances in sound recording and reproduction. As a result, theatrical sound practice from the late nineteenth and early twentieth century has not been widely researched, making historical sources of information from the period challenging to compile and corroborate.

This section explores the issues inherent in investigating this historical design practice, and outlines the approach and scope of the enquiry. The historical context of theatre sound effects, and some possible reasons for their hitherto minor place in sound studies and theatre scholarship are explored. The historical parameters of the enquiry are clearly defined. Finally, Cross' "designerly ways of knowing" (2006, p.29) and the Sound Design Toolkit (SDT) (Baldan et al. 2017) digital synthesis taxonomy are presented as two useful frameworks with which to explore the embodied knowledge at the heart of this historical soundmaking practice.

3.1.1 Theatre's Place in the History of Soundmaking

Theatre has probably always incorporated the creative use of sound. Writings on the acoustics of the theatre space, with methods for sound reinforcement, date back to 1 BCE (Brown 2010, p.158). Despite its status as the original site of sound effects design and performance, theatre's methods have so far received little attention from scholarship in sound studies (Ovadija 2013, p.9). Work within theatre history has revealed some interesting information on sound effects methods, but there are so far no detailed studies that explore what practitioners may have understood about sound and perception, and how specific designs were realised. This research opens up this area of enquiry. Some more recent examinations of the history of theatre sound effects have proved useful for this work, however. Shirley's (1963) study of how sound effects were used in Shakespearean plays, Culver's (1981) thesis that presents a historical timeline of theatre sound effects methods to 1927, and Collison's (2008) publications on his career as a theatre sound designer in the UK from the 1950s have all helped to uncover more sources from the late nineteenth and early twentieth century and corroborate their accounts. The potential reasons for theatre's history of soundmaking being under-researched merit further consideration, as they may suggest the means to discover and examine other unexplored areas of soundmaking from history.

The late nineteenth and early twentieth century brought great changes to the creative possibilities of sound that eventually overtook the mechanical methods long in use in theatres. This also influenced the trajectory of subsequent scholarship. This era has been characterised as one of "sonic modernity" (Curtin 2014, p.12), or "Ensoniment" (Sterne 2003, p.2).[1] Sound became separated from music into its own particular area of thought and practice, resulting in an explosion of new ideas about sound, and new technologies to create and reproduce it. A new modernist philosophy arose that rejected popular culture, and brought technical experimentation and a desire for new forms of media to the forefront in artistic fields like music (Born 1995, p.40). Composers began to reject traditional methods and take an interest in the use of sound, rather than musical instruments, in their compositions.

This approach grew new soundmaking practices that were outside of the very functional role of sound effects in theatres. For example, many of the sounds Luigi Russolo defined in his proposed "six families of noises" - rumbles, roars, creaks, scrapes and hisses - were already being created as sound effects in theatres. As part of his Futurist manifesto "The Art of Noises" (Russolo 1913) however, these sounds take on a radical new musical and aesthetic meaning. Audiences found Russolo's first noise performances with his mechanical intonarumori instruments controversial (*Noise Makers for the Futurist Concert of Noises*

[1] "Sonic modernity" (Curtin 2014, p.12) occurs between 1890 and 1925, while "Ensoniment" (Sterne 2003, p.2) covers the period between 1750 and 1925.

1914), but his approach influenced the composers that followed him. However, these composers became fascinated not with the intonarumori themselves, but with the potential of new media - electronics and sound recording - to realise these ideas (Manning (2013), Chadabe (1997)):

> Static between the stations. Rain. We want to capture and control these sounds, to use them not as sound effects but as musical instruments . . . Given four phonographs, we can compose and perform a quartet for explosive motor, wind, heartbeat and landslide. (Cage 1937, p.26)

Here, in his own manifesto "The Future of Music: Credo," the avant-garde composer John Cage makes an interesting distinction between sound effects and musical instruments. Although it is the same sound recording on his phonograph (e.g. rain), he suggests that using it for the purposes of an *effect* is a less artistic (and therefore less valuable) pursuit than using it to compose music. Cage captures here the two main challenges that theatre sound effects faced in the era of *Ensoniment* that would change their status in the twentieth century. As functional, entertaining and popular methods of creating and using sound, they did not fit with the new philosophy of modernist music. Rooted in materials and mechanisms, their methods did not fit with the new and exciting signal-based technologies of sound recording and synthesis.[2] Therefore despite broader cultural shifts in sound technology and thought, theatre sound effects did not manage to find a new home beyond narrative arts in the twentieth century. They instead moved into the areas of silent cinema, radio and film, and developed according to the conventions of those sites of performance.

As historical enquiry within the field of sound studies has so far focused on twentieth century artistic uses of sound (Kahn 1999), or the expansion of sound technologies such as radio or magnetic tape recording (Sterne (2003) and Thompson (2004)), it is understandable that theatre sound effects are overlooked. They stand very much apart from avant-garde musical innovation and the development of new electronic media for sound. This has recently begun to shift, with more detailed research in sound studies that reconsiders the history of early avant-garde sound performances with interactive sounding objects and voice in theatres (Ovadija (2013) and Curtin (2014)). This research expands

[2]Interestingly, the technology of the phonograph has already removed bodily performance from soundmaking in Cage's example (e.g. a recording of rain rather than a theatre rain machine).

on these previous enquiries in sound studies through its examination of the design practice that created theatre sound effects in the late nineteenth and early twentieth century.

The status of sound in theatres has also contributed to the lack of detailed scholarship on sound effects methods. Sound effects methods did not develop as a distinct specialism, and they also adapted to the circumstances of theatre performance over their long history. As outlined in Chapter 1, devices for soundmaking were a mere subclass of the "properties" (stage props) and only part of an entire group of effects that encompassed scenic, lighting and other elements (Krows 1928, p.114). Sound was not considered equal to other theatrical elements, such as costume or scenography. There were no established notation methods for sound performances, and techniques were probably only passed on through apprenticeship, as evidenced in the later practice that created sounds for silent cinema (Curtin 2011). Theatre sound performers were largely unskilled stagehands with little control over the creative aspects of an overall performance. Those with more creative control, such as stage directors, were multidisciplinary practitioners with priorities across sound, lighting and scenography.[3] Even in the nineteenth century, critics only tended to write about unique or particularly complex sound effects (Culver 1981, p.12).

This apparent disinterest around sound production methods is found in other areas of performative soundmaking, such as silent cinema (Altman 2007, p.77). These factors led to a failure to document the practice that created theatre sound effects to the same level of detail that other aspects of staging and production have enjoyed. Information on theatre sound is usually to be found in between instructions for lighting, visual or scenic effects in historical sources. Sound effects also responded to changing genres in theatre throughout its long history, and were not consistently perceived as important to dramaturgy. The spectacular visual and sonic illusions in nineteenth century productions did not persist in the face of modernism, for example. In particular, theatre in the mid-twentieth century aimed to deconstruct "theatre's mechanics" and bring realism to the stage (Counsell 2013, p.95). When sound effects were required in twentieth century productions, practitioners increasingly turned to sound recordings, which eventually replaced mechanical effects.[4] By

[3]A good example is Frank Napier, Stage Director at the Old Vic Theatre, London 1931-1934. While his "Noises Off" (1962) has been a compelling and useful text in the context of this study, he also published another book about the visual side of theatre, "Curtains for Stage Settings", in 1937.

[4]Culver proposes that as early as 1927, theatres could achieve a full thunderstorm sound effect using electrical methods (1981, p.133).

the 1960s, theatre was adopting many of the conventions around music production and live performance - recordings of sounds, mixing desks and loudspeakers (Collison 2010).

Theatre scholarship has more recently begun to connect its historical legacy of sound creation to present-day approaches in a continuum of sonic practice (Brown 2010), to engage more critically with the dramaturgy of sound (Roesner & Kendrick 2011) and to examine the unique sonic potential, or "aurality," of the theatre space (Kendrick 2017). This may lead to new and interesting soundmaking innovations and performance practices.

We have so far considered the historical context of theatrical soundmaking, and how external and internal influences have impacted its sound effects methods and their perceived importance in scholarship. Next, the scope of this historical enquiry will be outlined.

3.1.2 Historical Scope

The long history of soundmaking for theatrical performance is far too broad in scope to examine fully as part of this research. Instead, this historical enquiry focuses only on theatre memoirs, technical manuals, newspaper and magazine articles that detail the practice during the late nineteenth and early twentieth century in Europe and the U.S. This era offers the most interesting, and immediately accessible, information on the design practice that created theatre sound effects. This allows the embodied knowledge of the practitioners who created sound to be examined more closely. Some particular designs can also be examined in greater detail for their potential to inform new digital interactive sounding objects. The advantages of examining this particular era are presented here.

As this era straddles "sonic modernity" (Curtin 2014, p.12) and "Ensoniment" (Sterne 2003, p.2), examining theatre sound effects practice at this time places it in context with some of the wider cultural changes in sound thought and technology that have been previously explored in this chapter. This allows an examination of how these changes might have influenced the approach of theatre sound practitioners, particularly in terms of how they compared their established methods of soundmaking to the emerging use of sound recordings.

Theatre practitioners do seem to have been influenced by cultural shifts in soundmaking practices. Nineteenth century sources focus heavily on the technical prowess and inventiveness of theatre sound effects practitioners. By the early twentieth century, designers with a particular interest in sound were producing publications about their own particular methods, even giving some detail on their design ethos or devising process (Rose (1928) and Napier (1962)). This may have been in response to an increased demand for information on sound effects technique from practitioners in emerging areas such as silent cinema (Bottomore 2001), or perhaps to defend theatre sound effects as methods for recording and playing back sounds became more popular (Brown 2010, p.24). Looking at this era allows theatre sound effects methods to be more directly connected to the soundmaking practices for silent cinema and magic lantern shows (Bottomore (1999), Altman (2007) and Curtin (2011)), radio (Turnbull (1951) and Mott (1990)) and early sound film (Ament 2009). It seems that there was some cross-pollination between theatre and these emerging popular entertainment forms. For example, a New York Times profile of the playwright Lincoln J. Carter discusses both his visual and sound effects for theatre as well as a visual effect he achieved for an early film (1913, p.SM12).

By the late nineteenth and early twentieth century, some sound effects methods were very well established, forming what Brown has called a "culturally defining repertoire" or an "acoustemology of theatre" (2010, p.9).[5] With this shared group of theatre sound effects available to survey, comparisons can be made between the approaches of designers and performers across different productions tackling the same design problem. This provides a broad overview of the common themes and design problems inherent in the practice, and allows the often-brief descriptions of sound effects methods in historical sources to be gathered together to build more detail on each design.

Theatre's acoustemology is particularly interesting at this time. Practitioners seem to have been very interested in theatre's historical productions and methods during this era, particularly those used for Elizabethan plays. This is evidenced in sources such as Dutton Cook's "A Book of the Play," which discusses sixteenth-century sources of information on theatrical effects alongside more recent nineteenth century performances of sounds for Shakespearean plays (1876, p.246). Napier's 1962 publication also references his own work to provide sound effects for Shakespearean plays in the 1930s. Leverton's 1936 publication

[5] Acoustemology, or "acoustic knowing" (Feld 2015) is understood as culturally situated in the context of this research. Techniques to create performable weather, battle and animal sounds within European and U.S. theatres during this era may not have been shared by other parts of the world, and so there may be new methods to discover. Further research may be able to explore other geographical contexts and cultures to expand on the specific "acoustemology" of late nineteenth and early twentieth century theatre explored in this chapter.

presents a syllabus for teaching and staging an authentic recreation of nineteenth-century U.S. dramas, and includes instructions for sound effects as he suggests they would have been created. This awareness of historical methods is also to be found in texts where writers reference the origins of the particular sound effect that they are describing. Though not always accurate, this has proved useful in cross-referencing approaches as part of this research.

Theatre's acoustemology is also expanding in the late nineteenth and early twentieth century to encompass new and more complex sounds. There is evidence of new inventions to produce the sound of machines and vehicles. There are also some detailed descriptions of how sound effects producing large thunderstorms or moving trains were devised and performed, as will be explored in Section 3.2. These accounts help to reveal more of the embodied knowledge of soundmaking, and the implicit knowledge about the sonic possibilities of materials, that informed this practice.

Having established the historical scope of this research and surveyed some of the available sources of information on theatre sound effects in the late nineteenth and early twentieth century, we will now move on to consider the frameworks used to study this historical context of soundmaking and apply the embodied knowledge of theatre practitioners to the design of digital interactive sounding objects.

3.1.3 Connecting with a Historical Design Practice

This research converges a historical context of soundmaking with a present-day design problem, that of creating a continuous sonic interaction with a digital sound. It engages with historical sources on theatre sound effects from the late nineteenth and early twentieth century to understand how practitioners made everyday sounds performable, in order to apply their expertise to the design of a digital interactive sounding object. As explored in Chapter 2, SID research has already highlighted the need for this kind of artistic knowledge to help design new digital sonic interactions (Franinović (2013, p.20) and Pauletto (2017, p.346)), but how an artistic and embodied soundmaking practice can be best uncovered and applied to a new digital design has not yet been established. This is a central concern of this research.

Exploring a historical context of soundmaking documented only in written sources presents particular challenges. Clearly, practitioners working in the context of late nineteenth and early twentieth century theatre were not concerned with embodied interaction as we define it in the present day. Historical sources therefore do not record all of the information and perspective that is specifically interesting for this research. For example, there is little information on how the performers themselves experienced continuous sonic interactions with the interactive sounding objects they used; almost all sound effects are discussed only in terms of the audience experience. This historical enquiry must therefore work to make connections between the concerns of this research and the concerns of the past in its examination of sound effects practice and its survey of specific interactive sounding objects.[6] Two distinct frameworks are used to help establish a designerly dialogue with available sources on the soundmaking practice that created theatre sound effects, and unpack some of the embodied knowledge that informed their design in a way that is relevant to the design of digital interactive sounding objects. This connects the historical enquiry to previous research through design, which makes use of frameworks and prototyping to guide and structure its explorations (Stappers et al. 2014). This approach allows the historical sources investigated here to establish a focus for the design methodology of this research, which will be examined in Chapter 4. The two frameworks will now be explored further.

Although theatre sound effects were realised through the work of many practitioners (carpenter, property master, operator, stage director), this research frames these interactive sounding objects and their use in performance as the result of a unified chain of “designerly ways of knowing”, as defined by Cross (2006, pp.22-29). Cross’ five aspects of designerly knowing are extended beyond the design of objects to also encompass the design and performance of sound effects. These five aspects are used as a framework to examine various historical sources of information connect them together, allowing the insights and solutions of multiple theatre practitioners to be unified into a single design practice:

1. *Designers tackle “ill-defined” problems:* Ill-defined or “wicked” (Rittel & Webber 1973) problems are those that can only be understood through a series of attempts to solve them. Each new solution enables a further understanding of the problem. Theatre sound effects practice tackled a problem that was particularly ill-defined in an age before sound recording became widespread; how to produce sound effects that

[6]This research gathers information from historical sources on specific sound effects methods, a method previously used to discover more about the history of sound performance in silent cinema (Bottomore 1999).

are not just realistic imitations of real-world sound events, but that also communicate that intended imitation successfully to an audience.

2. *Their mode of problem-solving is "solution-focused":* This kind of problem solving works to create a timely and *good enough* solution rather than a perfect outcome after consideration of all possible approaches. Theatre sound effects were produced to serve a specific function within a theatre production, and were made to work according to the requirements of the stage director. Those involved in their construction and performance would not have had time to explore every potential iteration of a design to find the optimal solution. Instead, they worked on good enough solutions to perform each sound effect in a way that accommodated the timeframe of pre-production, rehearsal, and performance run.

3. *Their mode of thinking is "constructive":* Constructive thinking is based not in verbal, literary or numerical modes, but instead sketches and prototypes to develop, communicate and realise ideas. Theatre sound effects were not designed from standardised technical diagrams or specific materials and actions. Instead, practitioners constructively explored the sonic properties of materials and prototyped configurations of objects and mechanisms. Rather than describing the sound effect verbally or in writing, they performed it for the stage director's approval, and adjusted the interactive sounding object or performance to any additional requirements.

4. *They use "codes" that translate abstract requirements into concrete objects:* Designers translate abstract requests into physical artefacts. A feeling of connection for example could be added to a user's experience of using a smartphone app by adding subtle vibration under the touchscreen. Similarly, theatre sound effects were created not from specific requests, such as for a machine of specific length and height. Instead, dramas called for sound in an abstract and often elaborate way, such as Shakespeare's call for "a tempestuous noise of thunder and lightning heard" in Act I Scene 1 of The Tempest (Shirley 1963, p.42). Practitioners had to translate instructions like this into interactive sounding objects and performances through a synthesis of codes knowledge of the possible actions of an operator, the potential of mechanisms, the sonic range of materials, the acoustics of the theatre space, and the dramatic approach of the director.

5. *They use these codes to both "read and write" in "object languages":* Designers develop a fluency in their process that enables them to extend the knowledge they have already accumulated to the creation of new designs. As the acoustemology of theatre sound effects consolidated in the late nineteenth and early twentieth century, the codes of actions, mechanisms, materials and the theatre space enabled practitioners to explore the design of specific new mechanisms or complex new sounds, such as train effects.

Using these various *ways of knowing* as a framework draws out common themes in the creation of late nineteenth and early twentieth century sound effects, giving a more complete picture of the thought and practical work that created them than what is implied in the often-brief descriptions in individual historical sources. The distributed expertise that realised historical sound effects can also be brought closer to present-day design conventions, where performative digital soundmaking is often the work of one designer.[7]

Once this unified overview of the design practice has been established, the historical enquiry moves on to examine the specific action-sound configurations of the most widely used sound effect designs in the acoustemology of late nineteenth and early twentieth century theatre, as corroborated by many historical sources. This approach is also informed by Cross, who states:

> Objects are a form of knowledge about how to satisfy certain requirements, about how to perform certain tasks. (2006, p.26)

In this formulation, the embodied knowledge of the designer is embedded into the object they create. Cross' framing of the designed object as the repository of designerly knowledge has already been explored within the context of a research through design approach to interaction and HCI (Zimmerman et al. 2007, p.498). In the case of theatre sound effects, a rain machine is a designer's answer to the question *how can we perform a rain effect?* To connect this knowledge to present-day research into meaningful couplings between actions and digital sounds, the Sound Design Toolkit (SDT) (Baldan et al. 2017) digital synthesis taxonomy of basic sound events and textures is used as a framework to map out the action-sound configuration of each of these interactive sounding objects. This

[7]DMI designers possess skills across sketching, sourcing electronic components, constructing interfaces, programming audio, mapping, testing, performing and sound recording, for example.

produces an expanded design framework of digital sounds and actions that could be used to enactively recreate these historical theatre sound effects as digital interfaces controlling physical modelling synthesis algorithms. This work also reveals more about how historical designers connected simple sound events and actions together to realise more complex sonic illusions. One of these sound effects within this acoustemology, a wind machine, is chosen as a particularly interesting example to examine in more detail and use as the focus of the work undertaken in order to answer the remaining research questions.

This section has examined the way in which the historical enquiry of this research is pursued. The potential reasons for why theatre's historical expertise in soundmaking has been so far overlooked by researchers were outlined, and the scope of this historical enquiry was established. Finally, the two frameworks used to engage with available historical sources were presented. The next section examines the historical practice of theatre sound effects design and performance using the framework of Cross' five aspects of "designerly ways of knowing" (2006, p.29).

3.2 Designerly Knowledge of Soundmaking

Theatrical soundmaking in the late nineteenth and early twentieth century has so far been considered only as a chronology of specific sound effects methods. This research aims to make connections between the approaches of sound practitioners in various roles across theatres and productions in order to build a more complete picture of the practice. This helps to reveal more of the embodied knowledge of soundmaking that informed the creation of sound effects in theatres. This section evidences the designerly way in which sound effects were created using actions, mechanisms and materials through a consideration of historical sources on theatre from the late nineteenth to early twentieth century. These sources are examined using the framework of Cross' "designerly ways of knowing" (2006, p.29), and are also considered in light of the theoretical background of research previously presented in Chapter 2.

3.2.1 The Problem of Sound Effects

Theatre practitioners understood that the success of a sound effect lay in its ability to create meaning for an audience. There is a sense of the fragility of that potential success in historical writings. The interactive sounding objects did not just have to produce a particular sound, but also successfully give the audience an impression (Krows 1928, p.114) of particular real-world sound events. This was a very challenging design problem to solve, as several factors could potentially impact a sound effect's ability to create meaning. Historical practitioners are not always united on what these factors are, but they do offer some general insights into what should be considered when undertaking this work. These insights have been gained through a process of trying new solutions and adjusting established ones to fit each production. As Frank Napier notes in his theatrical sound manual "Noises Off":

> Whereas the design and execution of period costume is guided by rules, which are to be ascertained from the study of historical data, there are no rules, but only a few principles, involved in making noises off, for the problems differ with every play and every theatre. (1962, p.1)

A central concern of designers and performers was that audiences understood which sounds the designers and performers intended them to hear. Theatrical manuals advise that a sound effect should be as realistic as possible (Rose 1928, p.2), and also that it is instantly recognisable when heard (Napier 1962, p.42). The design process involved listening to an original sound, or recalling one from memory (Rose 1928, p.1), and then trying to imitate it with available materials (Napier 1962, p.13). The acoustemology of theatre in the late nineteenth and early twentieth century included effects to mimic familiar everyday sounds such as rain, as we will explore further in Section 3.3.

It is likely that regular theatregoers would also have become familiar with certain kinds of complex sound effects, such as thunderstorms, which were often a part of productions (Scientific American (1913, p.378) and Leverton (1936, p.50)). Weather sounds such as thunder and rain also had a dramatic and symbolic meaning for audiences, foreshadowing conflict or cosmic imbalance (Brown 2010, p.36). This would have made certain sound effects more immediately meaningful. The perceptual lexicon of real-world sound events that audiences brought to their experience of performances was also recognised as having the potential to reduce the efficacy of certain sound effects. In a chapter entitled "The Problem of Sound Effects", Philip Barber highlights the difficulty in recreating the sound of a bomber flying overhead after U.S. audiences' recent experience of hearing wartime air traffic [8] - the previous illusion, a loudspeaker playing a phonograph recording (presumably

[8]Barber refers to "events of the last few years" (1953, p.744), which suggests he is alluding to WWII.

from a height), no longer works as it once did. He links this kind of phenomenon to a growing reluctance among practitioners to include any sound effects in productions (Gassner & Barber 1953, p.744).

It does seem that practitioners were very cautious, and would rather exclude a sound effect altogether rather than use one they thought was imperfect (Napier (1962, p.3) and Peterson (1934, p.245)). Besides a consideration of whether the audience would recognise, understand and not be overly critical of a sound, it was also crucial to ensure that each effect fitted with the other sonic elements of the play. A breaking glass could not be louder than a thunderstorm, and that thunderstorm could not be loud enough to mask ongoing dialogue, for example (Peterson 1934, p.133). The size of the theatre space was a significant influence on the intelligibility and loudness of sounds (Napier 1962, p.5). To disperse sound effects to the full audience, devices were placed on each side of the stage (Scientific American 1913, p.378), under the stage (Collison 2008, p.44) and in the ceiling of the theatre (Southern 1944). Timing was also crucial as to how smoothly sound effects fitted with the performance, particularly if they were to cue a part of the drama. Green (1958, p.13) gives the examples of a ringing phone or explosion which, if delayed, will leave the actors waiting and potentially cause the atmosphere of the play to break down, drawing unintended laughter from the audience. In his introduction to A. E. Peterson's "Stage Effects and Noises Off", Arnold Ridley[9] goes further, saying that synchronisation is key to the success of a sound effect, and that at the time of his writing this is a more pressing concern than for previous generations of theatre practitioners (Peterson 1934, p.30).

However, creating a realistic, familiar and accurately timed sound was no guarantee of success. Barber links the ability of a sound effect to stir the imagination directly to the qualities of the play's text, set, lighting and actors' performances (Gassner & Barber 1953, p.744). Ridley similarly suggests that audiences' recognition of a sound is contingent upon the onstage action. He gives the example of a scene being introduced with the sound of rain. If a character appears onstage dripping wet, the audience immediately understand that they are hearing rain. If the same character enters the scene with no such visual clue, the audience might be confused by the same sound (Peterson 1934, p.29). This

[9]It is noted here for the reader that as well as being a playwright and producer of effects, Arnold Ridley was also an actor, and starred as Private Godfrey in the BBC television sitcom "Dad's Army"(1968 - 1977).

appreciation of the role of context in making a sound effect meaningful may reflect the fact that practitioners understood that certain sounds were perceptually close to each other. Some descriptions of specific methods add weight to this theory. For example, in his manual of sound effects, Green describes devices specifically for avalanches, crashes and rumbling sounds that are reported elsewhere as being suitable for making thunder (1958, p.43). It is clear that theatre practitioners considered carefully how the previous perceptual experiences of their audience, as well as the technical prowess they themselves could bring to their practice, would impact the ability of their sound effects to convey meaning. As Ridley explains:

> An audience must be *with* the off-stage effect. They must never be allowed to say: "I wonder what that is supposed to be?" They must *know* what it is supposed to be in advance. Then it is *they themselves* who provide the effect and not the workers behind the scenes. (1934, p.30)

We have so far explored the particularly wicked problems that practitioners were faced with in the design and performance of theatre sound effects. Next, the approach practitioners took to solving these problems will be outlined.

3.2.2 Solutions for Stage Directions

It is clear from historical sources that sound effects in the late nineteenth and early twentieth century were recast for each theatre building or performance run. This led to several distinct designs being produced to perform different qualities of the same popular sound. There are several ways to perform thunder for example, as will be explored in Section 3.3. It seems that practitioners were caught in a constant design cycle of iteration and reconfiguration, and there little agreement on specifically which devices should be used for particular effects (Krows 1928, p.114). This constant search for design solutions must have caused frustration:

> There has been little improvement in them in the last 50 years, and it really seems as though there is a field for some inventor to get up a universal sound machine the component parts of which could be operated from one central point. (Hopkins 1920, p.80)

How particular sound effect designs were first arrived at is not well documented, leading to inaccurate claims of inventiveness or confusion around the origin of particular methods

(Culver 1981, p.228). There is certainly an air of mystery, invention and individual skill around accounts of stage effects, particularly in the nineteenth century. The story of composer Giacomo Meyerbeer's thunder effect is a good example of this. In 1859, his opera "Le pardon de Ploërmel" (later known as "Dinorah") was due to be performed in Paris, but Meyerbeer was apparently holding up the whole production because of his dissatisfaction with the imperfect thunder effect he was being offered by the stage carpenter (Logan 1874, p.637):

> Passing under the scaffolding outside the Louvre, which at that time was being restored, he noticed the builders discharging plaster rubble from upstairs windows by letting it slide down a wooden chute. The dull rumbling produced by the fall of plaster rubble consisting of irregular shapes and weights gave him an idea, which he carried out on the stage of the Opera-Comique. A huge chimney or hopper was built out of thick pine planks running from the grid to the stage. Angled cross-pieces were fitted inside and a trap-door closed in the upper end of this huge pipe. When the thunder effect was called for, a load of quarry rubble - pebbles and foundry clinker - was placed upon the trap-door which was swung open on cue. Everything fell into the chimney, bouncing against the obstacles, hitting the sides and falling with a deafening noise upon the floor. (Moynet (1874), translated by Baugh & Wilmore (2015, pp.165-166))

This is an elegant solution that physically models the process Meyerbeer witnessed in the street, but offers the potential for further control over the thunder sound through adjusting the quality of the materials emptied into the chimney. The story is interesting in that Meyerbeer is both the sound effect designer and the composer with the authority to decide on the resulting sound. It was more usual that theatre sound practitioners did not set their own design briefs. The requests for sounds came from the play's text, or from the director and stage manager, and had to be responded to with timely solutions. Historical sources emphasise the difficulty and specificity of the challenges put to theatre sound effects practitioners. This description of a request for a more dramatic gale is typical:

> The gale I want, he says, must have teeth, or how is the hero to be rescued in them? I want roaring and pelting and smashing and clashing. None of

your land zephyrs for me. I must hear the dolphin howl and the shark shriek! (Vincent 1904, p.418)

Barber outlines a methodical approach to soundmaking that could give us a solution for even this exaggerated request, which is summarised here:

1. Listen to the actual sound you seek to imitate if you can, and analyse it.
2. Select the conventional sound machines you need or devise new ones.
3. Imitate the sound as closely as possible.
4. When satisfied, rehearse the sound for the director.
5. Adjust according to the director's changes.
6. Note down exactly how the sound effect was produced so that it is ready for dress rehearsal. (1953, pp.745-746)

This suggests that although the conventional machines of theatre's acoustemology are available for use, they may still need to be improved upon or adjusted somehow to accommodate a very specific request for sound. Theatrical manuals confirm this theory, as they tend to present several methods of making a sound to their readers, advising that they test and choose according to the needs of their own production (Napier 1962, p.62). Methods are described as ranging from very simple, but sufficient, to more complicated but better-sounding (Leverton 1936, p.55). Manuals offer instructions for larger objects to facilitate louder sounds (Rose 1928, p.14), and scale down these larger designs usually installed within theatres to make use of more readily available and less expensive materials (Peterson 1934, p.134). This suggests that the practice was solution-focused. Any of the factors previously outlined that might impact on a sound effect's ability to create meaning for an audience could be heightened in the circumstances of a particular production (for example, a change of theatre building or a request for a particularly complicated or abstract effect). This would necessitate some adjustment and development on previously satisfactory methods.

Performance was also a crucial part of how sound effects were developed and improved upon. Cue systems were installed within the theatre to ensure that sound effects were accurately timed and added to the drama of the play rather than causing disruption. These were originally achieved by tying string to the performers who were out of the line of sight, and tugging on it to activate their performance (Napier 1962, p.23). A coding system of coloured electric lights was later implemented (Green 1958, p.16). Not all solutions offered by theatre sound effects designers were successful, however. Nineteenth century theatre writers in particular take delight in reporting the mishaps that occurred in theatrical performances. These stories give a sense of how practitioners developed their embodied knowledge of soundmaking as they improved on previous failures.

Edward Dutton Cook reports an attempt to replicate an Elizabethan method for making thunder in an Edinburgh theatre, which was achieved by pushing a wheelbarrow full of cannonballs over an uneven surface behind the stage. This works well until the operator trips up, turning over the wheelbarrow and sending cannonballs rolling all over the stage, into the orchestra pit and even towards the audience (1876, pp.246-247). Similarly, Percy Fitzgerald recounts a last minute attempt by an effects operator to produce the sound of a rain shower in a country theatre production in the absence of proper equipment.[10] Someone is sent to buy a large sheet of brown paper, and the operator grabs it in both hands, desperately rubbing it against a wall behind the stage. Fitzgerald reports that this sounded much more like hissing air rather than rain, however (1881, pp.62-63). Perhaps this effect was improved on for the next performance.

Having established some of the solutions theatre sound effects practitioners developed in response to requests for particular sounds, we will move on to consider some examples of the constructive thinking that practitioners used to create interactive sounding objects and choreograph sound performances will be outlined.

3.2.3 Soundmaking Thinking

Despite Barber's advice to take written notes on exactly how a sound effect was produced (Gassner & Barber 1953, p.746), there is little evidence of a culture of archiving specific mechanical diagrams or of notating performances as part of theatre sound practice in the late nineteenth and early twentieth century in the sources accessed for this research. Some practitioners are also very clear that this was not part of their work. For example, Napier advises his readers that unmusical sound cannot be drawn or notated in any way (1962, p.1). There is evidence that this practice was centred on constructive thinking however, as listening, configuring and performing led the design process and communicated ideas

[10]Fitzgerald actually suggests that real water was the missing material. However, it is likely that actual water would not have been considered as a source of a rain shower sound, but rather used as a visual effect.

to others.

In his manual “Noises Off”, Frank Napier advises his readers to train the auditory observation, and analyse the components of a sound. He gives the example of a door slam effect, which is usually achieved by dropping one wooden board against another one laid flat on the floor. He complains that this only replicates one part of the sound, which should also include the rattle of a letter box or the squeak of a hinge (1962, pp.13-15). This is a simple example. An everyday sound, a wooden door with a letterbox being slammed, is replicated with a similar action and configuration of materials - wooden boards, and a metal rattle, being slammed to the floor. Further evidence of the deconstruction of a real-world sound, and its reconstruction as an action and interactive sounding object can be found in designs that specifically played with resonance, when mechanisms were enclosed within wooden boxes (Napier 1962, p.34), objects swirled on the surface of drums (Collison 2008, p.44), or sheets of metal struck to give echoes to a gunfire effect (Napier 1962, p.76).

The translation of a stage direction or sound heard into action, mechanism and material could be much more complex, however. In the foreword to “Noises Off,” director Tyrone Guthrie quips that Napier could easily sound the direction “the nightbirds exult in the tree-tops” armed with only an empty matchbox and three bent pins (1962, p.v). At first glance this seems overly complimentary, but within his own text Napier does discuss in detail his prototyping of the sound of a steam-ship engine with his fists, a signet ring, a writing desk, a typewriter on a rubber sheet, and some rattling desk drawers:

> By pounding alternately with my closed fists I find that I get a surprisingly good representation, for the table booms, the drawers make one rattle, the typewriter another, and the ring striking through the rubber sheet gives exactly the right knocking sound. I doubt whether one could keep this up for any great length of time, however, but the work of each hand could be given to a separate person, thereby halving the labour. (1962, p.45).

Napier goes on to discuss how this prototyped sound effect could be then made loud enough for a theatre setting through the use of different objects. This kind of embodied skill and constructive thinking for soundmaking must have been at the heart of other practitioners' approaches. There are clues to this in the texts that single out particular practitioners for their ability to imitate any possible sound (Vincent 1904, p.418), and in some of the legends told about their skill:

> There has been current in the theatre for years the story of a veteran property-man who saw Niagara Falls for the first time. When asked what he thought of that grand sight he replied contemptuously that he could make a much better effect with a wheel-barrow and a piece of tin. The amazing part of the story is that he probably could as far as the theatre is concerned. (Krows 1928, p.115)

As well as having enough implicit knowledge to be able to reach for and configure just the right materials to make a specific sound, practitioners also developed embodied knowledge of performance techniques. The interactive sounding objects that were purpose-built to perform specific effects offered some flexibility in their sonic potential, as will be outlined in Section 3.3. A good example of this is the wind machine, which Napier (1962, p.50) highlights as particularly performable. Rehearsal was considered crucial to the success of a sound effect, particularly if it required a complex performance (Bax 1936, p.56). Napier suggests that this soundmaking skill can be developed within the body through an enactive process of repetition of movement and an assessment of the resulting sonic feedback:

> The *feel* of it, transmitted to our brains through the sense of touch seated in the joints, muscles and skin. This *feel* is retained by what I have called the tactile memory. Let the man rehearse his gunfire, therefore, often enough to impress the correct *feel* on his tactile memory. (1962, p.6)

It is evident that listening, configuring materials and rehearsing performances were all a crucial part of the constructive thinking designers used to develop sound effects. We will now consider some of the codes of this practice as implied by the available historical sources.

3.2.4 The Codes of Soundmaking

Theatre sound practitioners used codes to translate the stage director's requirements into appropriate sound effects. As previously explored within this chapter, descriptions of sounds could be highly complex and abstract. These had to be realised through the application of an embodied knowledge of sound, the sonic possibilities of materials, an understanding of mechanism, and a familiarity with theatre performance practices. One of the codes of theatrical soundmaking that is explored in more detail as part of this research relates to the craft of carpentry for machine design. Backstage in the late nineteenth and early twentieth century theatre was a truly mechanical space, with "innumerable drums and windlasses" (Logan 1874, p.629) used to move scenery, curtains and visual illusions. These mechanisms were hand-operated. The skill of the "carpenter-machinists" (Jackson & Wilson 1976, p.109) who made them was highly specialised:

> Stage carpentry is a separate branch of the craft, and most first-class stage carpenters have served their time with one of the scenic contractors. In this way the stage carpenter learns all sorts of tricks of the trade that an ordinary man never thinks about. (Bax 1936, p.53)

It is suggested here that these carpenters learned their trade through apprenticeship. There is little detail on the specific mechanics of particular sound effects designs, or the carpentry techniques used to realise them, in the available historical sources that discuss theatre sound. As mechanical diagrams had long been part of engineering by the late nineteenth century (Lefèvre 2004), their absence in theatrical sound manuals does suggest that stage carpentry was an embodied practice that was not documented in written or diagrammatic form. In his exploration of the craft of medieval carpentry, Ingold offers some useful insights that can be extended to the ill-documented practice of stage carpentry for theatre sound effects. He suggests that buildings were constructed in the Middle Ages through a "tactile and practical geometry" (2013, p.51). Carpenters would have used their own bodies and perspectives to calculate distances and angles to determine, for example, the eventual height of a tower. Medieval buildings were therefore constructed through an iterative process of solving design problems as they arose, rather than sketching and planning the final structure in advance (2013, p.53). The lack of mechanical diagrams in historical sources on theatre sound effects practice suggests that a similar approach was taken to their design and construction. The lack of explicit reference to the acoustics of specific sound effects designs in historical sources, despite advances in acoustical science in the nineteenth century (Beyer & Raichel 1999), also lends weight to this.

Within the late nineteenth and early twentieth century theatre, practitioners understood the relationship between the performance space and the potential loudness, duration and position of a sound. This is evidenced by the design of what are termed *primed* sound effects in the context of this research. These effects are prepared in advance, and merely activated by a performer rather than continuously interacted with. We have already examined one of these within this section, Meyerbeer's *trémie* design, where masonry and other materials were emptied down a specially designed chute. This design was used in several theatres (Vincent (1904, p.421) and Napier (1962, p.55)).

Another design, known as a *rabbit hutch*, consisted of stacked cannonballs inside a wooden cupboard with sloping shelves (Rose 1928, p.7) built high up into the theatre above the stage. To perform the effect, the doors of the cupboard were simply opened, sending the cannonballs rolling down from the shelves and through some wooden chutes designed to let them roll down the back wall behind the stage (Peterson 1934, p.133), or perhaps across the flies above the stage.

Theatres also had channels called *thunder runs* built into their ceilings (Southern 1944). Performers rolled individual cannonballs through these chutes to give the effect of thunder rumbling above the audience (Southern 1952, p.101). The *venetian blind* design to produce a thunderclap consisted of large planks of wood strung up like a blind and hung high up backstage. To perform the effect, the rope was pulled and released to send the planks crashing down to the floor (Logan (1874, p.637) and Rose (1928, p.17)).

These four methods the *trémie, rabbit hutch, thunder run* and *venetian blind* all share a common approach, in that they do not afford a continuous interaction with the thunder sound. Instead, in each case the performer primes an apparatus and releases something to trigger the sound. The importance of the theatre space is evident in these four designs, and must have been a critical part of how thunder was made loud enough and of long enough duration for performances. These devices too must have been the result of an embodied practice of tactile geometry, as designs were adjusted to produce just the right quality of sound. This embodied skill of stage carpentry must have realised interactive sounding objects in theatres that allowed a sound to be performed, but also allowed that sound to be varied by the performer.

As recorded sound began to be used in theatres, practitioners struggled with a new way of working that did not conform to the old codes around sound creation with actions, mechanisms and materials. This can be glimpsed in some sources. For example, Napier (1962, p.50)) outlines his prejudices against a gramophone recording of wind, explaining that its limited duration and inability to respond to events onstage results in obvious repetition, and ultimately an alienating experience for the audience when compared with a performed mechanical wind machine. The recorded sound has violated previous codes around the way a wind sound was made performable, creating a new set of design problems to solve. Part of the success of a performance involved what Napier calls “stage sense” (1962, p.9), or an understanding of the importance of dramatic timing. We have already explored how practitioners stressed the importance of timing to the success of an effect, and how cue systems were used to help time sound effects performance. Dramatic timing, and knowledge of what is required of the play in production, is another design code implied in historical sources.

Here, we have considered some of the codes of theatrical soundmaking practice that are implied in historical sources from the late nineteenth and early twentieth century. Finally, we move to explore how the object language of soundmaking with actions, mechanisms and materials was used to read and write in a designerly way, through an examination of some more inventive sound effects designs and complex performances.

3.2.5 The Object Language of Sound Effects

Theatre’s sound design problems and solutions, the constructive way in which practitioners devised their sound effects, and the codes that designers might have made use of in their practice have already been considered. Further evidence of theatre sound effects practitioners’ ability to read and write in their designerly language of action, mechanism and material for sound creation can be found in examples of particularly inventive designs or ambitious performance methods. Stage directions for sound became much more specific towards the end of the nineteenth century (Culver 1981, p.121), and therefore more demanding on the skills of designers and performers. With an acoustemology of theatre sound already established, practitioners worked to create new sounds that were not already a regular part of performances. There is evidence of some unique interactive sounding objects in historical sources. For example, in his manual “Stage Effects: How to Make and Work Them”, Rose presents his readers with a design to imitate the sounds of an aeroplane with a rotating toothed wheel and crank handle (Figure 3.1):

> Now if a banjo or guitar string is made loose or tightened up by turning the peg to which the string is attached, a moaning and variation of sound will be the result, as given by the propeller of an aeroplane; but when a number of strings are worked in unison, the volume of sound is greater, and may be so manipulated as to give a very good imitation of an aeroplane speeding through the clouds (1928, p.24)

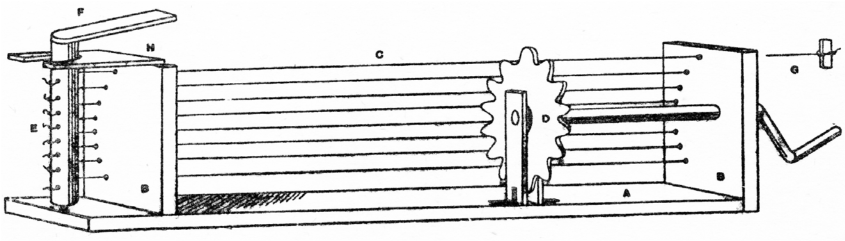

Figure 3.1: An aeroplane or propeller effect (Rose 1928, p.23). Image included here courtesy of Taylor and Francis.

Without a textual description like this of a particularly unique design, it can be unclear as to what the original purpose of a sound effect was. Collison describes a machine preserved from the Theatre Royal in Drury Lane, London,[11] that makes a very unique sound when rotated, which at first seems to mimic a carriage wheel. Without any written record, it is difficult to establish the specific effect that it has been created to produce:

> This machine has a revolving wheel made of iron with a heavy metal weight resting on it. The scraping sound could resemble the mechanics of carriage wheels or, perhaps, some kind of lathe - were it not for the piece of wood fixed to the rim of the wheel. This lifts the weight on each rotation, briefly interrupting the scraping noise, then lets it fall with a metallic clank. Could it be for a ship's engine or a factory machine or a comic car? (2008, p.61)

Unique examples like this give weight to the myths about the skills of particular practitioners. It is clear that some designers and performers were able to develop a sophisticated practice that could accommodate requests for very specific sound effects. As well as inventing new devices, practitioners also devised elaborate performances of evolving sound effects for larger set pieces, particularly thunderstorms or trains. This description of a thunder effect used in Century Theatre in New York 1913 evidences a configuration of different objects and carefully timed performances (Figure 3.2):

> Thunder is produced in two parts of the stage. The crash takes place in one part of the fly galleries, while the reverberation or roll is produced on the thunder drum, which is beaten by a hand near the backdrop. Then the fly

[11]This device was part of the Theatre Museum collection in London when Collison documented it in 2002, and must now be in the care of the Victoria and Albert Museum.

> man releases a bolt, which allows a pivoted box to dump about eight hundred pounds of stones and assorted junk down a chute, striking a large iron place at the bottom. This is immediately followed by three hundred pounds of chain, which is also released and allowed to fall on an iron plate. Then the roll is begun by dropping five or six cannon balls down an incline or chute, which is technically called a "rabbit hutch." The finale of the effect is produced by deforming a large sheet of iron, which is twisted. This is used in connection with the thunder drum. (Scientific American, 1913, p.874)

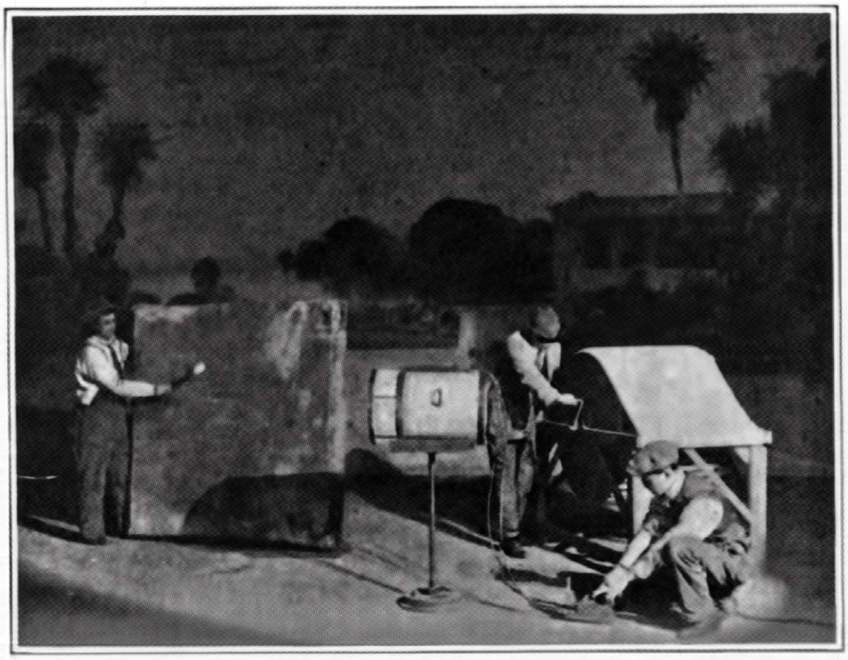

Figure 3.2: Some of the theatre sound effects apparatus from Century Theatre. A square thunder drum can be seen on the left, and a wind machine on the right. The original caption reads "Thunder drum for rain and reverberating thunder, lightning flasher and wind machine" (Scientific American, 1913, p.378). Image included here courtesy of Scientific American.

Train sounds were also produced through a combination of objects, such as wire brushes, drums and metal tubing (Krows 1916, p.228), and were considered particularly complex and important effects to create. In his interview with the New York Times, Lincoln J. Carter claims that he first devised his visual and sound effects for his 1889 play "The Fast Mail," and then wrote his script around them, for example (1913, p.SM12). One particularly famous train effect from "The Ghost Train", a play by Arnold Ridley which first ran in London in 1925, is cited by several sources as an impressive example of the potential of sound effects (Somerfield (1934, p.79) and Napier (1962, p.39)). Ridley gives his own brief overview of the effect:

> The effect itself was perfectly simple, being merely a judicious mixture of thunder sheets, compressed air, a garden roller, and pieces of sandpaper and wire brushes combined with a kettle-drum. (Peterson 1934, p.29)

In fact, the play included three distinct train effects, and gave audiences the sensation of a train moving through the theatre. A complete breakdown of each train sound, along with detailed stage directions, can be found in Collison's "Stage Sound" (1976). Each train was produced through a combination of a thunder sheet, various whistles, a garden roller, an iron tank, wire brushes and some electrical motors. The stage directions cue each specific element very carefully in order to create three distinct train effects (1976, pp.89-90). The performances must have been very carefully rehearsed before production (Figure 3.3).

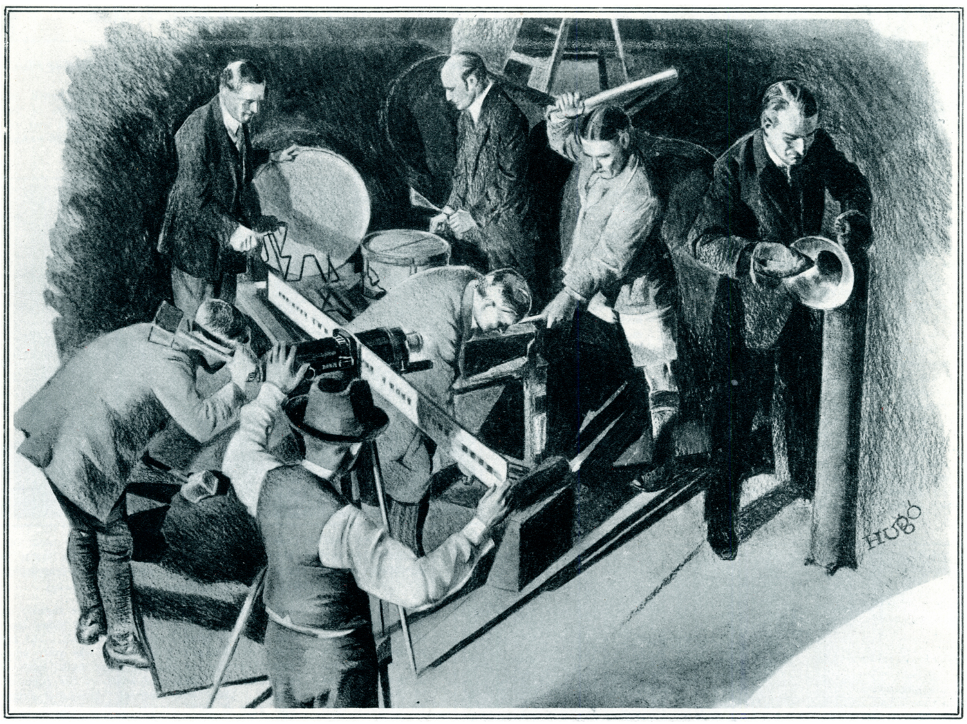

Figure 3.3: An illustration of some of the train effect performers for Arnold Ridley's "The Ghost Train" in a Munich Production. The original caption reads "Realism in Sound Behind the Scenes: Railway Noises "Off."" (Grein 1928, p.517). Image used here with permission of Mary Evans Picture Library.

As theatre sound effects practice developed in the late nineteenth and early twentieth century, practitioners played with their established acoustemology to encompass new requests for sounds and produce performances to rival those of other popular entertainments. This required them to develop a fluency in the object language of theatre sound effects.

This section has explored historical sources on theatre sound effects from the late nineteenth and early twentieth century using the framework of Cross' five aspects of "designerly ways of knowing" (2006, p.29) to uncover more about the design and performance practice that created them. The next section examines the specific designs that placed theatre's acoustemology into the hands of performers, and uses the SDT (Baldan et al. 2017) taxonomy as a framework to classify them in terms of their actions and sounds.

3.3 An Acoustemology of Actions, Mechanisms and Materials

This chapter has so far examined how designers and performers worked to produce sound effects in theatres within Europe and the U.S. in the late nineteenth and early twentieth century. Sources from this period also contain descriptions of varying detail that explain how specific sound effects were made. These have been collected and classified as part of this enquiry in order to examine the action and sound configuration of each interactive sounding object. Classifying these descriptions in this way reveals how particular theatre sound effects could be realised as digital interactive sounding objects. The SDT (Baldan et al. 2017) taxonomy introduced in Chapter 2 is used here as a framework to map out a design space for new enactive recreations of these sound effects using its digital synthesis algorithms.

In the interests of a robust enquiry, this research engages directly with those designs that make up theatre's acoustemology in the late nineteenth and early twentieth century. These are the most widely used methods at this time, as corroborated by the sources gathered for this historical enquiry. Sound effects that have been used across different productions and performance settings have had the principles of their design tried and tested by many practitioners, even if the interactive sounding objects themselves were usually constructed and adjusted for each production. This suggests that there is some agreement on the efficacy of these methods, and perhaps also the meaning created for the performer when using the interactive sounding objects.

As this research centres on hand-operated interactive objects that continuously produce sound, some categories of sound effects that were part of theatre's late nineteenth and early twentieth century acoustemology are excluded from this survey of specific methods. These include those effects that were produced merely through simple object manipulations, such as twisting one clay pot inside another to produce creaking sounds (Leverton 1936, p.56), *primed effects* as described in Section 3.2, and sound produced without hand-manipulated materials; human vocal mimicking, whistles, bells and theatrical firearms (Culver 1981). These methods were also widely used, but have less to offer in terms of strategies for designing a purpose-built object to facilitate continuous interaction with a hand-operated sound. The specific interactive sounding objects that were part of theatre's acoustemology during this era will now be considered, classified by the sound effect they produced.

3.3.1 Crashes

Theatre practitioners used two dedicated methods to produce crashes for productions. These were both based on an action to produce rotation and a sound created through impacts, but implemented in slightly different ways. Both were known as "crash machines" (Leverton (1936, p.55) and Culver (1981, p.169)). The first machine consisted of a wooden drum enclosure that housed material like cannonballs, masonry or pieces of wood. This was mounted on a frame and axle, and turned either with a crank handle or by holding the drum directly in both hands and pulling or pushing it to facilitate rotation. The material inside of the drum hit against the sides and against itself, sliding and impacting to create sound. Green suggests that this crash machine design produces a rumbling effect (1958, p.45), but Napier describes the sound as a "splendid crashing, rumbling grinding noise" (1962, p.60). Interestingly, Napier's version has a barrel made of sheet iron and drum sides made of wood, whereas Green does not specify the material of the drum. It is likely that the use of sheet iron had a significant impact on the quality of the sound (Figure 3.4).

The second crash machine design was based on a ratchet mechanism, or a "hand-operated, magnified version of the watchman's wooden rattle" (Electric 1947, p.12). This was based on a wooden cylinder mounted on a frame and turned with a crank handle. The slats were arranged to make a toothed cylinder, which turned against some flexible

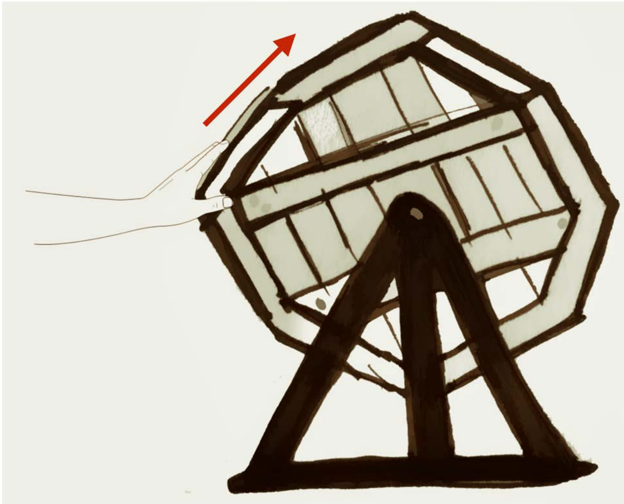

Figure 3.4: Crash Machine Method 1.

wooden slats mounted to the frame, causing them to bend, snap back and clatter against the cylinder (Rose (1928, p.15) and Peterson (1934, p.565)). This machine could be rotated at different speeds, and was used to perform gunfire sound effects as well as more generic crash sounds (Leverton (1936, p.55) and Collison (2008, p.50)). Peterson confirms that this machine can produce continuous crashes, and that larger constructions will produce very loud sounds (1934, p.566). As explored in Section 3.2, the apparatus of thunder sound effects was also used to also perform crash sounds (Green 1958, p.40), suggesting that these sounds were perceived as perceptually close to each other (Figure 3.5).

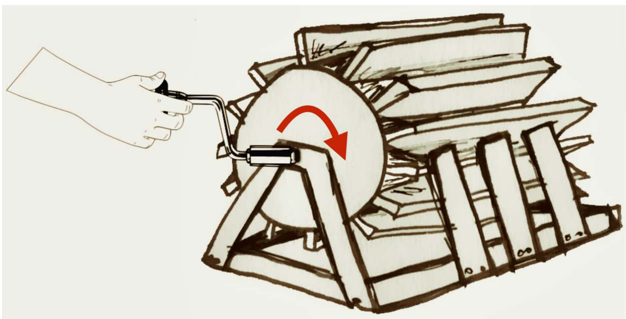

Figure 3.5: Crash Machine Method 2.

3.3.2 Creaking and Roaring

The sound of creaking, or animal roars (a lion's roar is sometimes specified by practitioners (Vincent 1904, p.419)), was created with a barrel and string configuration. This device is named as a *bull roarer* in historical sources,[12] but sometimes it is given other names - Napier calls it a “creak-tub” for example (1962, p.89). The two ends of the barrel are removed, and one of them is replaced with a drumskin. A string, usually made from cat gut, is threaded through this drumhead and fixed in place:

> The operator, ordinarily the stage manager, places the drum between his knees. Then, hand covered by a glove coated with resin, he rubs more or less vigorously on the cat gut cord, and he thus produces some prolonged growlings, muffled or ringing. (Jackson & Wilson 1976, p.134)

This method does seem to have had great potential, and could produce several effects. Green suggests it can produce “a variety of creaks and squeaks” as well as a lion’s roar (1958, p.46). Adjustment of the tension of the string, the amount of resin or rosin on the cloth, and the grip of the operator could all help to articulate different sounds. Substituting the wooden barrel for a drum (Jackson & Wilson 1976, p.88), or even a large tin can (Gassner & Barber 1953, p.748) would also have changed the quality of the sound effect. This is a particularly interesting design (Figure 3.6), as it appears to be an early iteration of Luigi Russolo’s intonarumori musical instruments (Russolo 1913). As outlined in Chapter 1, his designs were centred on a string threaded through a drumhead, although this was enclosed inside a wooden box rather than a barrel, and excited with a mechanism rather than by hand.

3.3.3 Rain and Surf

Rain, sea and surf sounds were created with three main devices. Again, as these effects methods are interchangeable, these sounds must have been understood as perceptually close to each other. All of the designs listed here are described as being able to perform rain or sea-related sounds across the sources examined as part of this enquiry. The first and simplest of these was the rain box or rain tray. This was a wooden box, tray or wire sieve, which was held in two hands, and tilted back and forth or rotated with a slow

[12] This is potentially confusing - what we would term a *bullroarer* today is a prehistoric noisemaker that creates a wind-like or whistling sound. A piece of string is attached to a specially carved piece of wood, and whirled around in the air to create a whistling sound (Fletcher et al. 2002). The barrel and string method does seem to be very old though, and perhaps dates back to Elizabethan times. It seems that the *bull roarer* name for the barrel method may come from it being known for imitating a bull’s roar.

Figure 3.6: Bull Roarer.

movement. This action rolled and swirled the material in the tray, which is specified as dried peas, metal shot or marbles (Logan (1871, p.119), Rose (1928, p.14) and Peterson (1934, p.230)). Altering the tray by studding it with metal nails (Rose 1928, p.42), or using a wooden hoop with a paper base (Leverton 1936, p.45) changed the quality of the sound. Krows also suggests dropping the dried peas or shot directly onto the surface of the tray during performance to produce a slightly different rain sound (1928, p.121) (Figure 3.7).

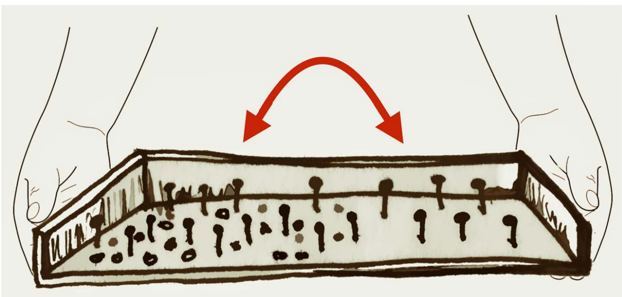

Figure 3.7: Rain Tray.

The second method of producing a rain effect was known as a rain box (Napier 1962, p.53). This was a long oblong enclosed wooden structure mounted on an axle much like a seesaw. Performers would hold the box in both hands and tilt it, which moved the shot, peas and other material to the other side, producing a rain or surf-like sound. Ledges and obstructions within the box ensured that the material fell slowly enough to be effective (Vincent 1904, p.418) (Figure 3.8).

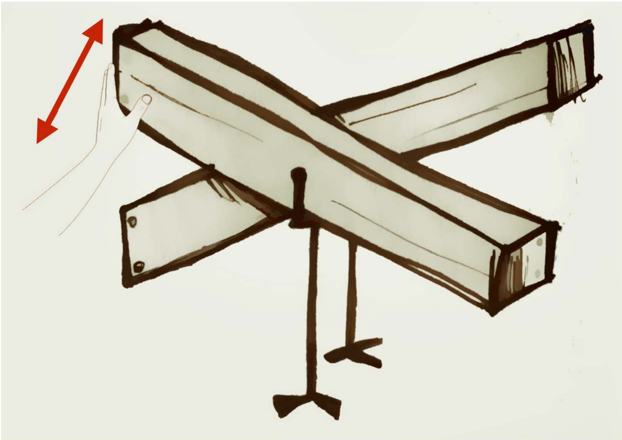

Figure 3.8: Rain Box.

The final method of performing rain and sea sounds mechanised the sound producing process even further. Like the crash machine design, an enclosed cylinder was constructed and mounted on a rotating axle. A crank handle allowed the cylinder to be rotated. As the handle was turned, the material inside moved, sliding, rolling and hitting the interior. Again, an irregularly shaped channel inside slowed the progress of the dried peas or shot. Practitioners adjusted this design to facilitate different kinds of rain sounds:

> If beans are used, the effect is of rain on a roof. Shot instead of peas gives the effect of rain on a tin roof. Large shot gives the effect of hail. (Gassner & Barber 1953, p.750)

There are also reports of seeds and sea shells being used in these rotating machines (Cook 1876, p.254), which must have produced a different quality of sound. Rose adapts this design with a tin cylinder and steel shot to produce the sound of a steam engine (1928, p.30) (Figure 3.9).

3.3.4 Thunder

Thunder sound effects were produced in several different ways in theatre productions in the late nineteenth and early twentieth century. Some of the primed effects that did not afford a continuous sonic interaction with the thunder sound - the trémie, thunder run, rabbit hutch and venetian blind - have been previously examined in Section 3.2. Two further

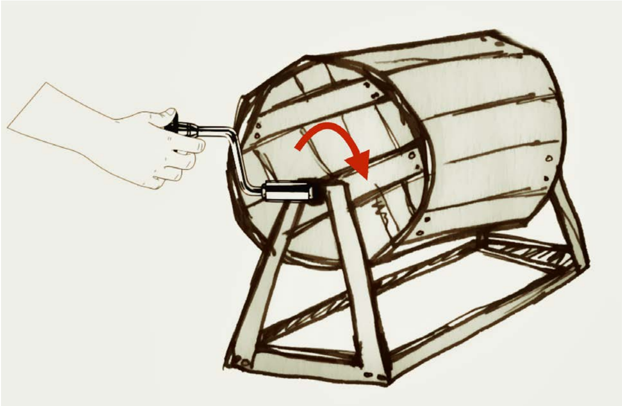

Figure 3.9: Rain Machine.

designs do produce a thunder sound through a continuous sonic interaction. One of these is the thunder cart or wagon, which is constructed from wood. It has uneven wheels to make it wobble, and is pushed forward behind or above the stage (Kranich (1929, p.244) and Vincent (1904, p.421)). The rolling and impacts produced by the wheels created a rumbling thunder effect (Figure 3.10):

> The cart was weighted by being filled with old iron, bricks, and rounded stones or cobbles. The volume of sound that was produced by it was enormous. (Peterson 1934, p.133)

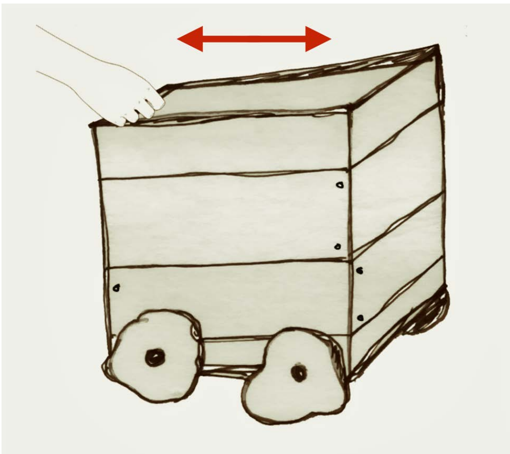

Figure 3.10: Thunder Cart.

Another method that produced thunder sounds for performance is the thunder sheet. This was a long sheet of metal hung up behind the stage that was equipped with a handle at the bottom end (Somerfield 1934, p.77). Performers held the sheet and shook it continuously, creating rumbling and peals of thunder (Green 1958, p.39) (Figure 3.11). Barber gives some specific measurements for the thunder sheet, and links the thickness of the metal to the realism of the sound effect it will produce:

> The thunder sheet is conventionally a piece of approximately 20 gauge sheet metal, 3' wide and 6' to 12' long, suspended lengthways so that the bottom edge is about three feet from the stage floor. When the bottom edge is shaken, a somewhat metallic reverberation results. A sheet of 3/16" profile similarly operated gives a deeper, more realistic tone.(Gassner & Barber 1953, p.752)

Bax also suggests that some plywood can be used as a temporary measure if there is no sheet iron available (1936, p.149).

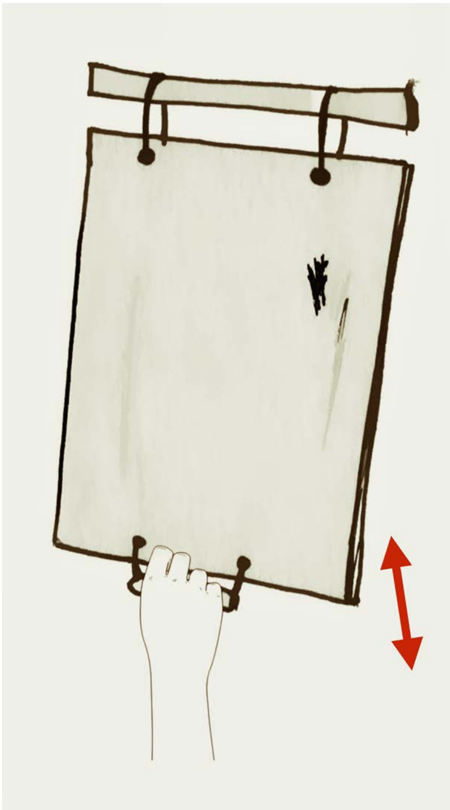

Figure 3.11: Thunder Sheet.

3.3.5 Wind

A wind effect was produced in theatres using another cylinder and axle machine design. A slatted cylinder was mounted on a frame, and turned with a crank handle. The cylinder was covered by a piece of canvas, and turning the handle scraped the wooden slats against the cloth to produce the sound of wind (Green 1958, p.41). Some sources identify the cloth as silk rather than canvas (Scientific American, 1913, p.378). By varying the speed of the handle's rotation, performers could change the force of the gale (Krows 1916, p.219) (Figure 3.12). This method seems to have been very popular among practitioners, who debated different methods of adjusting and improving the sound:

> The noise is absolutely indistinguishable from the real thing. Some people favour wetting the canvas, but I have never been able to notice any improvement. (Somerfield 1934, p.79)

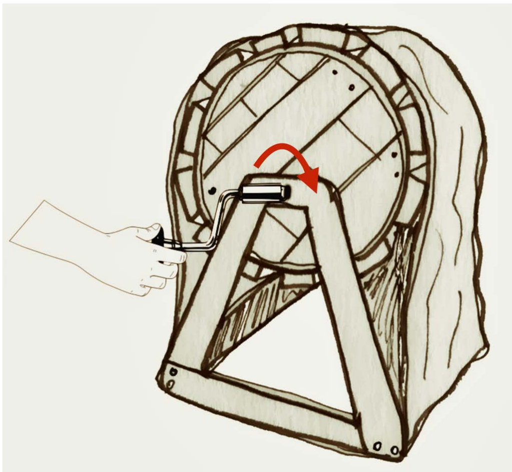

Figure 3.12: Wind Machine.

The wind machine will be considered in further detail in Section 3.4 as a particularly interesting theatre sound effect to use as a focus for this research.

3.3.6 Action and Sound Configurations

Examining these historical methods in this way allows some general conclusions to be drawn about these theatre sound effects methods:

- These sound effects were performed through continuous sonic interaction. The performer held something in their hands, and then they activated and continuously modulated the sound directly through their own movement.

- Gestures were not tied to specific sound effects or to particular categories of sound. Rotation, for example, is used to perform wind, rain and crash sounds.

<table>
<tr><th>Sound Effect</th><th>Action</th><th>Basic Sound Event</th></tr>
<tr><td>Crash</td><td>Rotate</td><td>Impacts</td></tr>
<tr><td>Creaking/Roaring</td><td>Squeeze and Pull</td><td>Scraping</td></tr>
<tr><td rowspan="3">Rain</td><td>Tilt and Turn</td><td>Rolling, Sliding</td></tr>
<tr><td>Pull and Push</td><td rowspan="2">Rolling, Sliding, Impacts</td></tr>
<tr><td>Rotate</td></tr>
<tr><td rowspan="2">Thunder</td><td>Pull and Push</td><td>Rolling, Impacts</td></tr>
<tr><td>Shake</td><td>Deformation, Crumpling</td></tr>
<tr><td>Wind</td><td>Rotate</td><td>Scraping</td></tr>
</table>

Table 3.1: Classification of Theatre Sound Effects by Basic Sound Event and Action.

- Adjustment and experimentation with materials was an important part of the design process. Practitioners were aware, for example, that the use of metal rather than wood would change the quality of a sound effect.

- Some of the interactive sounding objects produce several different effects. The crash machine can produce crashes, thunder or gunfire for example. Purpose-built methods for producing rain can also create the sound of the sea. This suggests that certain sounds were perceived as perceptually close to others. New sounds produced by already established effects methods must have been discovered during the rehearsal process.

- Sound effects methods are based on certain principles rather than on specific interactive sounding objects. This is evident in the designs for producing rain and surf sounds, which encompass a range of methods of varying complexity. The simplest is a tray of dried peas rolling and sliding to produce sound. The rain box and rain machine develop this method in increasingly mechanised ways, but still produce sound through the activation and modulation of the same rolling and sliding.

Having examined these core sound effects designs from theatre's acoustemology in the late nineteenth and early twentieth century, their action-sound configurations can now be classified more clearly in terms of the basic sound events of the SDT (Baldan et al. 2017) (Table 3.1).

3.3.7 A Framework of Enactive Recreations

Having established the action-sound configurations at the heart of late nineteenth and early twentieth century theatre's acoustemology, this information can now be added to the SDT digital synthesis taxonomy. This maps out how enactive recreations of these historical designs might be realised digitally (Figure 3.13).

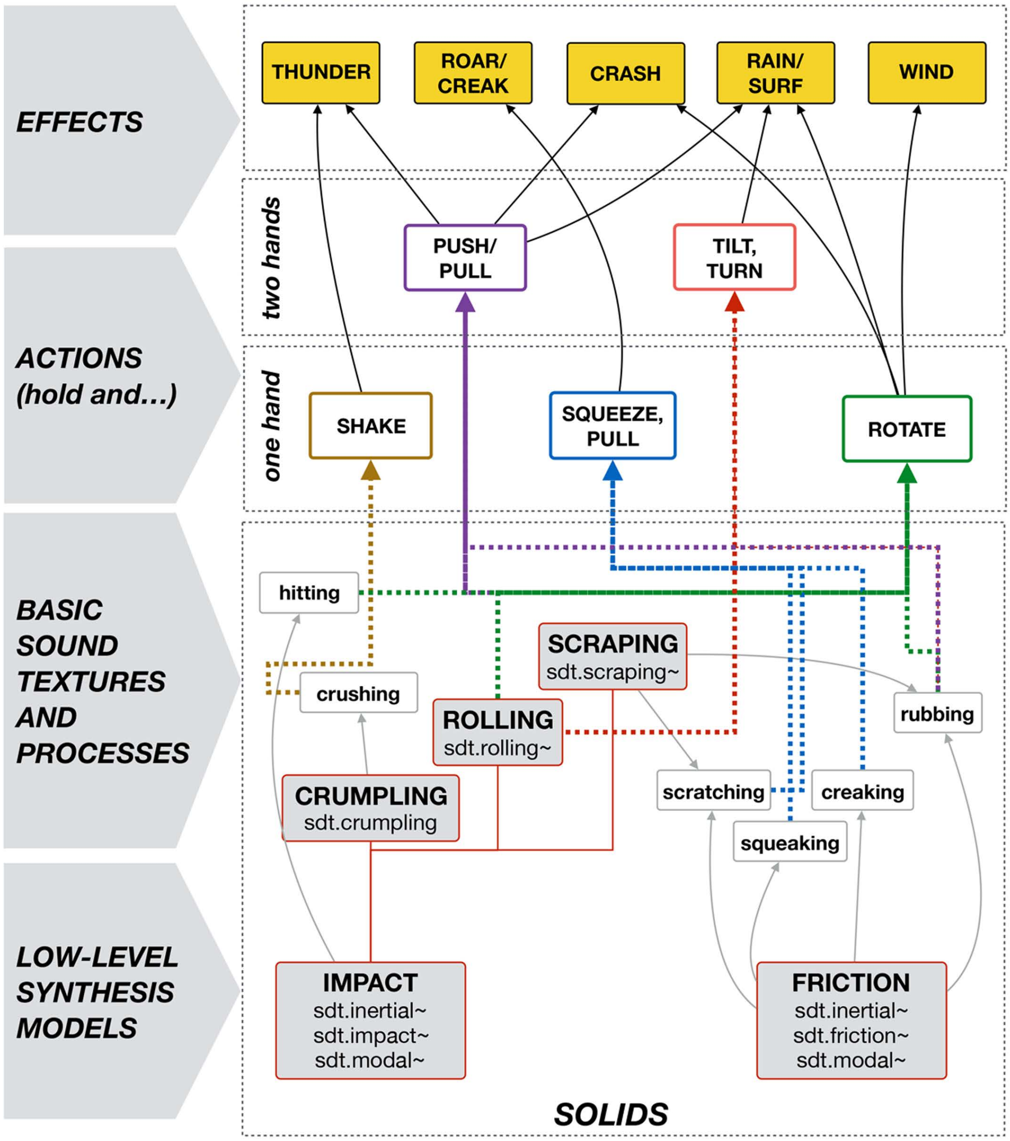

Figure 3.13: Expanding the Sound Design Toolkit (SDT) (Baldan et al. 2017) with the action-sound configurations of theatre's acoustemology in the late nineteenth and early twentieth century.

This new framework connects the SDT digital synthesis algorithms directly to potentially meaningful actions. It also shows how a low level sound model, like crumpling, could be used to make a perceptually meaningful effect like thunder if the model is activated with a shaking action. These are so far very simple couplings between action and sound, however. Some actions, particularly rotation, imply that there must be a virtual mechanism to connect the simple sound model (e.g. scraping, rubbing) to the performance of a more complex and yet meaningful sound (e.g. wind). In order to explore the full potential of this new proposed design space, this research focuses on one particular theatre sound effect.

The next section examines this effect, the wind machine, more closely, and considers its origins and design in detail. Its particularly interesting qualities as an interactive sounding object affording continuous control over a continuous sound are highlighted.

3.4 A Machine to Imitate Wind

This chapter has so far outlined the historical method that informed this part of the research enquiry, explored aspects of the designerly knowledge of theatre sound practitioners and examined the specific action-sound configurations of some of the main theatre sound effects methods in use in the late nineteenth and early twentieth century for their potential to create new interactions with digital sounds.

This section now moves on to present a more detailed examination of one theatre sound effect design that is used as the basis for the remaining work undertaken to answer the research questions. As outlined briefly in Section 3.3, a wind machine consists of a wooden slatted cylinder mounted to rotate freely on an axle fixed to a frame. The performer rotates the cylinder by means of a crank handle. As the cylinder rotates, its wooden slats rub against a piece of cloth or canvas draped over it, which produces a wind-like sound. The wind machine is explored in detail as part of this research as a particularly interesting example of a theatre sound effect. Here, we examine why this particular effect is worthy of further investigation. This informs the research methodology that will be introduced in Chapter 4.

3.4.1 An Enduring Design

The wind machine is presented as a particularly successful theatre sound effect method in historical sources:

> A wind machine is an instrument upon which the sensitive effectsman can play any required tune, and, moreover, "accompany" a scene in the most telling manner. (Napier 1962, p.50)

It is most likely that this sound effect method was first established in the nineteenth century. The earliest reference to the cylinder-based design can be found in Jean-Pierre Moynet's 1874 "L'Envers du Théâtre":

> There are several kinds of machine in use to imitate the sound of wind, but this is the most usual: a solid framework supports the axle of a cylinder resting upon two trunnions. The cylinder is made up of individual sections each of which in cross-section looks like the tooth of a cog wheel, and form projections on the surface of the cylinder. Generally, there are between fifteen and twenty of these ridges fastened to the turning cylinder. Strong silk fabric on top of the frame runs over the cylinder and small bolts, which can be tightened as required, allow it to be tensioned. When the handle of the cylinder is turned, the friction of the silk over the ridges produces a continuous sound exactly like that of wind whistling in chimneys or corridors. (Baugh & Wilmore 2015, p.169)

Practitioners, sometimes with reference to Moynet himself or the fact that it was originally a "French" method, consistently report this method of creating a wind sound in historical sources. Once in use in theatres, the rotating cylinder wind machine was not developed into different iterations as we have seen with rain or thunder effects, but rather remained very consistent across productions and theatres. The persistence of this design over so many years, even when recorded wind sounds were available, suggests that this method was particularly successful as an effect for the audience, and potentially also for the performer. As previously highlighted in Section 3.2, the wind machine was even compared favourably to recordings of real-world wind (Napier 1962, p.50). How realistic and performable this sound effect is warrants further investigation.

3.4.2 Synthesis Method

The synthesis principle at the heart of the wind machine, the *rubbing* or *scraping* of the wood against the cloth, was actually established before the nineteenth century:

> A sound of rushing, whistling wind, as successfully produced by De Loutherbourg for the theatre of David Garrick at the close of the eighteenth century, was made by rubbing two silk-covered discs together; and this remains perhaps the simplest of the good ways of giving the effect. If one will run his

> nails with a rotary movement over a tightly stretched cloth of any kind, he will illustrate the principle. (Krows 1928, p.116)

Here, Krows connects a scraping sound produced by running his fingernails over silk in a rotary motion with an early mechanical design that rubbed two silk discs together. Practitioners must have understood this connection. Rose, for example, shows a simpler iteration of the cylinder wind machine design in his theatrical manual - a single slat of wood scraped across a stretched piece of silk (1928, p.9) (Figure 3.14).

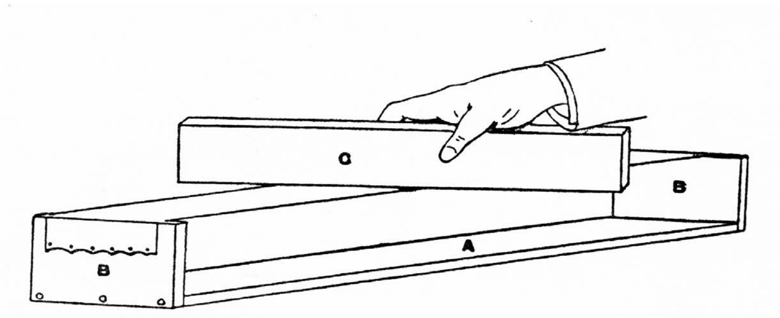

Figure 3.14: A simple single-slat wind effect (Rose 1928, p.9). Image included here courtesy of Taylor and Francis.

The cylinder based wind machine may have been developed from an investigation of these simpler methods of facilitating scraping. As will be explored in Chapter 5, practitioners do not make connections between the aeroacoustics of real-world wind sounds and this design, despite their being studied in the nineteenth century (Beyer & Raichel 1999, p.95). If the original design was not informed by a knowledge of acoustics, it may have been the sound of natural wind itself, perhaps through a sonic affordance suggesting a gesture of rotation to practitioners. How these simple scrapes could become a complex wind sound, and how continuous interaction with a wind machine creates a perceivably wind-like sound through scraping is worthy of further investigation.

3.4.3 Continuous Action, Continuous Sound

The wind machine couples a continuous action (a gesture of rotation) to a continuous sound (wind-like, produced through friction). Cylindrical structures became ubiquitous in the nineteenth century with developments in mechanical design (Müller-Sievers 2012, p.4). However, the technical motion of a cylinder rotating on an axle is very different to the natural movements of the human body. How a simple and smooth gesture of rotation has the potential to produce a continuously varying sound in an interaction is interesting. Historical sources do link the speed of rotation to a variation in the sound. Krows suggests that "the force of the gale (depends) upon the rapidity with which it is turned" (1916, p.219). Barber is more specific, and links speed to a variation in the pitch and volume of the wind (1953, p.753).

Changes in rotational speed may be linked to a complex process of sound excitation and modulation. How the speed of a hand rotating a crank handle is linked to the creation of a more forceful gale, and how this linkage might facilitate enactive learning through rehearsal, is not fully explained in historical sources. A rotational gesture is also used for other historical sound effect designs, and so any interactive sounding object produced as part of this research could be adapted to control different digital synthesis models as required. These other designs include rain machines and crash machines, as outlined in Section 3.3. Unlike these other designs, the wind machine's sound is produced through continuous scraping rather than a sequence of distinct impacts.

3.4.4 Expression and Rotation

The wind machine is operated by rotating a crank handle. This control gesture offers a somewhat familiar perceived affordance to users of audio control surfaces for music or sound performance. A rotary potentiometer, which is usually turned between fingers and thumb between its minimum and maximum range, will increase or decrease the amount of an already moving signal or sound. The crank handle, by contrast, requires the whole hand and indeed arm to be engaged in a gesture of rotation. Besides requiring a larger movement on the part of the operator, the crank has the potential to be the actual source of that moving signal or sound, given that it can be rotated continuously. Previous research has suggested that the experience of performing a sound with a crank handle is particularly satisfying when compared with a rotary or linear potentiometer (Gelineck & Serafin 2009). The wind machine design gives the opportunity to examine the perceptual experience of performing a continuous sound with a crank handle more closely.

This section has examined a wind machine as a particularly interesting example of a theatre sound effect to create digitally. Some of the peculiarities of the design, and potential areas of investigation, have been highlighted. A summary of this chapter now follows.

3.5 Summary

This chapter has presented an overview of theatre sound effects design and performance practice in late nineteenth and early twentieth century Europe and the U.S., and explored its potential to inform the design of new continuous interactions with digital sounds. Some potential reasons for the neglect of theatre's soundmaking expertise in previous scholarship were presented. The historical method that informed this part of the research was outlined.

Cross' "five aspects of designerly ways of knowing" (2006, p.29) were used as a framework to explore historical practitioners embodied knowledge of sound and perception, the sonic possibilities of materials, and the potential of sound to create meaning in a theatre performance. The Sound Design Toolkit (SDT) (Baldan et al. 2017) taxonomy of basic sound events was used as a framework to examine some specific methods used for producing theatre sound effects in this era, and classify them by their actions and sounds. This information was then added to the SDT taxonomy to extend its potential and map some potentially meaningful connections between its digital synthesis algorithms and performance actions to enactively recreate these theatre sound effects.

Finally, one specific effect, a theatre wind machine, was explored in more depth as a particularly interesting example with which to answer the remaining research questions of this thesis.

Next, Chapter 4 outlines the research methodology used to fully investigate the wind machine as an interactive sounding object, realise it digitally, and evaluate the continuous sonic interaction it affords.

Chapter 4

Design and Evaluation Methodology

This thesis has so far presented a theoretical background exploring human perception of real-world continuous sonic interactions and prior research into the possibilities of creating similarly rich multisensory experiences of digital soundmaking. An overview of the history of theatre sound effects design and performance practice, using the frameworks of "designerly ways of knowing" (Cross 2006, p.29) and the Sound Design Toolkit (SDT) digital synthesis taxonomy (Baldan et al. 2017) was also presented. This positioned theatre sound effects as late nineteenth and early twentieth century interactive sounding objects, with direct potential for the design of meaningful interactions with digital models of everyday basic level sound events. The focus then shifted to one specific sound effect method, a theatre wind machine, which was introduced as a particularly interesting example to realise digitally. This chapter now outlines the research methodology used to investigate the sonic interactivity of the theatre wind machine design and realise it digitally.

As discussed in Chapter 1, the main aim of this research is to examine the contribution historical theatre sound effects from the late nineteenth and early twentieth century can make to methods for designing and evaluating continuous sonic interactions with digital sound. This research methodology therefore concentrates on two main areas; how a digital interactive sounding object might be created based on historical information about the wind machine design, and how that digital interactive sounding object can then be best evaluated.

As theatre sound effects have not previously been subjected to this kind of study, this endeavour involves a synthesis of methods from distinct fields, and draws theatre history, making in humanities research, sound design, audio programming, acoustical evaluation, experiment design and statistical analysis together. The first part of the methodology used is design-led, and the second part focuses on evaluation methods. A research through design approach (Zimmerman & Forlizzi (2014), Stappers & Giaccardi (2017)) threads throughout this work. Practical and perceptual investigation and making inform the design of objective evaluations, and technical work such as programming is conducted in an exploratory way with the aim of examining and revealing more *wicked problems* in the transition from acoustic to digital sound performance. As will be outlined in this chapter, this methodology encompasses not only the creation of a digital model of an acoustic wind machine, but also a full examination of the sonic properties and interactive potential of the acoustic wind machine design. This helps to ground the programming and evaluation work in a clear comparison between the soundmaking properties of the acoustic and digital systems, while also allowing for an investigation into the meaning created for a performer when interacting with a continuous everyday sound like wind.

Section 4.1 explores how a digital model of a wind machine can be created based on historical information from the late nineteenth and early twentieth century. The method of remaking that informs the work to create an acoustic wind machine presented in Chapter 5 is discussed. The physical modelling synthesis and interactive sounding object design approach presented in Chapter 6 is also introduced.

Section 4.2 presents the acoustical evaluation methods used to examine the sonic properties of the acoustic wind machine and calibrate the digital model of its workings. The background to the experiment design used to evaluate the continuous sonic interaction afforded by both wind machines is also presented, along with an explanation of the statistical analysis procedures used. The experiment procedure and results are fully reported in Chapter 7.

Section 4.3 summarises the methodology presented in this chapter.

4.1 Digital Design from Historic Interactions

As we explored in Chapter 3, theatre sound effects design and performance practice in the late nineteenth and early twentieth century was rooted in an embodied knowledge of the soundmaking possibilities of actions, materials, mechanisms, and an understanding of the acoustical properties of the theatre space. While some of this knowledge can be glimpsed through a survey of historical sources from this era, there is much embodied expertise - the "tactile and practical geometry" (Ingold 2013, p.51) that created specific interactive sounding objects for example - that has not been carefully documented. It is therefore difficult to model a theatre wind machine digitally using only historical texts as sources of information on its workings. As such, the design-led methodology first works to create a working example of a theatre wind machine, and focuses the programming work on modelling this specific example. The resulting digital model of the wind machine is framed as a "prototype" (Stappers 2013) - a carrier for reflection, examination and objective evaluation of the potential for theatre sound effects to inform new satisfying interactions with digital sound. This section outlines the stages of the design-led methodology.

First, the available historical sources of information on the wind machine design are briefly revisited to examine the variations in descriptions and design approaches that challenge the possibility of modelling a definitive digital version. Remaking is then examined as a solution to this variation in historical accounts, and as a way to engage more fully with the implied expertise and embodied knowledge embedded in historical design instructions. The programming approach used to create a prototype digital model of the wind machine in Max/MSP using the Sound Design Toolkit (SDT) (Baldan et al. 2017) is then presented. Finally, all of the potential methods of performing a continuous acoustic or digital wind sound explored so far within this thesis are mapped out in detail. This helps to examine the multisensory experience offered by each system a little more, connects real-world materials to digital soundmaking, introduces how the digital model of the wind machine is made performable, and foregrounds the aims of the experiment design that will be introduced in Section 4.2.

4.1.1 Describing a Wind Machine

Chapter 3 explored the solution-focused way in which practitioners in the late nineteenth and early twentieth century responded to requests for particular theatrical sound effects. Methods were often adjusted or redeveloped for each performance run, and the wind machine was no exception to this despite the popularity of the design. While there seems to have been agreement that the crank-and-cylinder design for performing wind was the most effective available to them, each theatre practitioner presents their own adjusted version of that design in historical sources. These adjustments might be discussed in detail within a text, or simply shown in a drawing presented for the reader. Unlike the sonic qualities of an acoustic musical instrument, like a violin or guitar, that vary depending on very small adjustments to the way in which it is constructed (the choice of woods, for example), the wind machine is realised in often quite distinct ways. These may have offered quite a different experience of performance when interacting with the wind machine, and also seem to have had the potential to produce many different qualities of wind sound.

Practitioners sometimes highlight the uniqueness of their approach to the core principles of the wind machine (the crank, cylinder and cloth) as crucial to the success of the resulting effect. This is evidenced in historical sources, such as the detailed discussion offered by theatre practitioner Ernest Maurice Laumann, in which he describes a long process of adjustment and testing to create his perfect wind effect that improved upon the "high-pitched whistling" found in Parisian theatres at the time (1897, pp.136-137). First, the wooden slats of the wind machine are covered in strips of what he terms "metallic canvas", presumably to produce a particular quality of friction between them and the encompassing silk cloth. He does not stop there, however:

> After some trial and error, we arrived at this. The silk was fixed only at one end, the other end to a wooden bracket for use as a foot pedal. We had, moreover, stretched violin strings, and other silk ribbons rubbed with rosin, over the cylinder. By actuating the cylinder, we produced a noise analogous to that which was produced in the other theaters: but by pulling and releasing the silk ribbon held by the foot, we obtained modulations, and harsh or deep sounds, such as those produced by real wind. (Laumann (1897, p.137), translation by the author and Culver (1981, p.110))

These continual processes of adaptation means that there are in fact as many different implementations of the principles of the wind machine design as there are historical examples. Despite these variations, most historical sources do not provide the kind of detail in Laumann's own account of his work. Instead, some very general guidelines are offered as to how the choice of material or performance gesture impacts the resulting wind sound, for example:

> The tension of the canvas must be adjusted for the required sound. The faster the drum is turned, the higher the note of the wind. (Green 1958, p.41)

> Vary volume and pitch of sound by speed of revolution. Heavy silk instead of canvas gives an interesting variation. (Gassner & Barber 1953, p.753)

While the mechanical design and material qualities of the wind machine are variable and ill documented, there is also very little information to be found in historical sources on how it felt to perform with. As previously highlighted in Chapter 3, practitioners at this time are not interested in the present-day concerns of this research, such as whether any factors like weight, tactile feedback or effort improved the experience of performance. Historical descriptions of the wind sounds this interactive sounding object is capable of producing are also inconsistent. The wind machine is described as capable of "shrieking" (Napier 1962, p.50) and of both "moaning" and "shrieking" (Bax 1936, p.147), for example. The translation of texts into English is also a factor here. For example, Moynet's wind machine design (1874, p.169) is described as producing "swishing" (Jackson & Wilson 1976, p.135) but this is more accurately translated as "whistling" (Baugh & Wilmore 2015, p.169). Leverton (1936, p.50) compiles descriptions of wind machines from various nineteenth century writers and notes adjectives such as "howling, roaring, sighing" and "whistling."

Given all of these descriptors, it is difficult to determine whether a wind machine should *whistle* or *roar*, or be capable of both. In addition, the writers of these historical descriptions of wind sounds were of course bound to their own context and sensory experiences (Jay 2011, p.310), which must have influenced their perceptual lexicon of sounds. The historical meaning of a descriptor like roaring may not therefore be shared in the present day. These variations in the construction and sonic descriptions of the wind machine, as well as the lack of clarity on performers' experiences of its use make it difficult to create a definitive digital model of this sound effect based on the information in historical sources alone. This research employs a hands-on approach to historical enquiry and works to remake a wind machine from historical design instructions. We will now move on to consider how this approach reveals more of the embodied and implicit knowledge behind the wind machine design, which helps to model it digitally.

4.1.2 Remaking to Enhance Text

As outlined in Chapter 1, this research connects with the "maker turn" (Staley 2018) in humanities research, and uses remaking as a method to investigate the wind machine design in more detail. This extends the historical enquiry presented in Chapter 3 by allowing a deeper exploration of some of the embodied knowledge of soundmaking with actions, materials and mechanisms implied in historical sources. It facilitates a full and detailed examination of an acoustic wind machine as an interactive sounding object, which informs the creation and calibration of a prototype digital synthesis engine to model its main soundmaking components. It also draws the author's own soundmaking practice together with the expertise of historical practitioners, allowing the design process to generate knowledge and avenues of enquiry (Durrant et al. 2017). Making, fabricating and prototyping have all been used as research methods to explore the creation of historical objects and technologies. This facilitates a tactile and critical approach to material histories. Prototyping, or building simple models of historical technologies allows the design of even one-off, fictitious or no longer existing objects to be explored (Sayers 2015, p.158). Engaging in making and fabrication also connects historical design and manufacturing practices to present-day fabrication methods such as 3D printing (Turkel & Elliott 2014, p.179).

As part of this research, remaking allows historical mechanisms, designerly knowing and even sensory experiences to become part of a design-led enquiry. Elliott et al.'s (2012) work to prototype stage magic illusions from the nineteenth century is particularly relevant to this research. They created dollhouse-sized figurines and accompanying props to enhance textual descriptions of particular methods, such as the levitation effect shown in Figure 4.1. This informed an enquiry into the designerly knowledge implied by historical texts, and also a new understanding of the perceptual aspects of the performance and audience experience - sightlines, lighting and misdirection, for example. Elliot et al. highlight that the sensory experiences, and in particular the visual perception, of the practitioners who created these stage effects may have been very different to their own in the present day (2012, p.123). Crucially, they stress that this technique of remaking allows them to experience the intended effect and link their own perceptual experiences with those of historical practitioners:

> Working with actual physical recreations, we can experiment with viewpoint, lighting, colours, patterns, posture, misdirection, sightlines, material properties, and so on, discovering which set of factors makes an effect seem more or less convincing to our senses and comparing our own subjective experiences with those recorded a century ago. (2012, p.124)

In the same way, it is possible that theatre practitioners in the late nineteenth and early twentieth century may not have heard the world as we do today. Connecting present-day sonic experiences to those reported in the past helps to apply the embodied expertise of historical practitioners more directly to the design problem of making the wind machine design digital. This research therefore foregrounds remaking as a particularly powerful method of investigating how historical theatre sound effects might inform the design of digital interactive sounding objects. It extends the enactive recreation approach previously explored in SID (Serafin & De Götzen 2009), and instead engages with the design problem presented to historical theatre practitioners - how a fully acoustic wind sound effect can be made performable. Remaking the continuous sonic interaction afforded by the historical wind machine design allows a more detailed study of how it creates sound and what it feels like to perform - both crucial considerations in the design of a digital sounding object.

While this work aims to examine what has been omitted from or implied within historical sources, and is informed by a re-enactment of historical design instructions, it is also acknowledged that the possibility of really "being there" in a fully shared understanding with historical practitioners is not attainable (Sayers 2015, p.173). Instead, information from historical sources is examined and implemented through the author's own embodied knowledge of acoustic and digital soundmaking, and in the context of the theoretical background of research outlined in Chapter 2. The process to remake the acoustic wind machine is fully detailed in Chapter 5. It is proposed that remaking has particular potential as a method to research historical theatre sound effects, and may be fruitfully applied to some of the other designs detailed in Chapter 3. This approach enables a full investigation of the rigor of historical design instructions by using them as a guide for construction. It also allows a more full understanding of the meaning created by the sonic and ergonomic limits of the wind machine in performance. Materials can be touched and manipulated, and the sounds described in texts can finally be heard.

Next, the discussion shifts to an examination of the programming approach used to realise a digital model of the remade acoustic wind machine.

4.1.3 Programming Digital Wind

As outlined in Chapter 3, historical descriptions of the wind machine design explain that its wind-like sound is produced by a *rubbing* or *scraping* action. This rubbing and scraping occurs between the wooden slats of the cylinder as they are rotated with the crank handle and the cloth, which resists their movement. The wood and the cloth are both materials with many microscopic surface irregularities, which creates this resistance to movement, or frictional force, between them. As friction is the principle on which the wind machine's sound is based, the prototype digital wind synthesis engine produced as part of this research is based on the Sound Design Toolkit (SDT) physical model of friction (Baldan et al. 2017). The programming work is informed by a *practical synthesis* approach, with a focus on the perceptual and interactive characteristics of the digital sound created. This allows for a detailed exploration of the experience of the digital wind-like sound in performance, as well as an objective analysis of its sonic properties. The acoustic wind machine is first studied as an interactive sounding object following its construction in order to discover how best to model its soundmaking process digitally. An entity-action model is then created to deconstruct how the movement of the crank handle creates sound (Farnell 2010, p.36). Modelling the workings of the acoustic wind machine in turn begins to reveal more about the underlying physics inherent in its process of sound production. This is a method long in use in musical instrument acoustics research (Woodhouse 1992, p.43). This modelling approach is unique, however, in that it is modelling an acoustic sound effect that is itself a model of something else.

The acoustic wind machine design is an attempt to model the everyday sound event of wind using action, mechanism and material. It is therefore useful to briefly consider how the sound of wind has been digitally synthesised using other methods. Naturally occurring wind does not innately produce sound. The sound of wind is instead the result of moving air meeting an obstruction, which produces disturbances known as *vortices* in the airflow that oscillate back and forth and can produce audible sound. This occurs particularly when moving air interacts with cylindrical objects, such as telephone wires strung on telegraph poles (Farnell 2010, p.474). Moving air produces a variety of aeroacoustic sounds, including the Aeolian tone, an audible pitch produced by wind through this process of flowing against and around obstacles. As wind speed increases, this Aeolian tone rises. The Aeolian tone has been successfully modelled as a digital sound using a mathematical description of airflow against a cylinder (Selfridge et al. 2016). A convincing digital wind sound can also be produced through a signal-based method, where a noise source is passed through a moving band-pass filter (Farnell 2010, p.475). Given that the prototype digital wind synthesis engine produced as part of this research is based on the SDT digital model of friction, it implements a new approach when compared with these other synthesis methods (Figure 4.2).

Figure 4.1: Elliott et al.'s (2012, p.124) model of a levitation effect. Photograph by William J. Turkel and used with his permission.

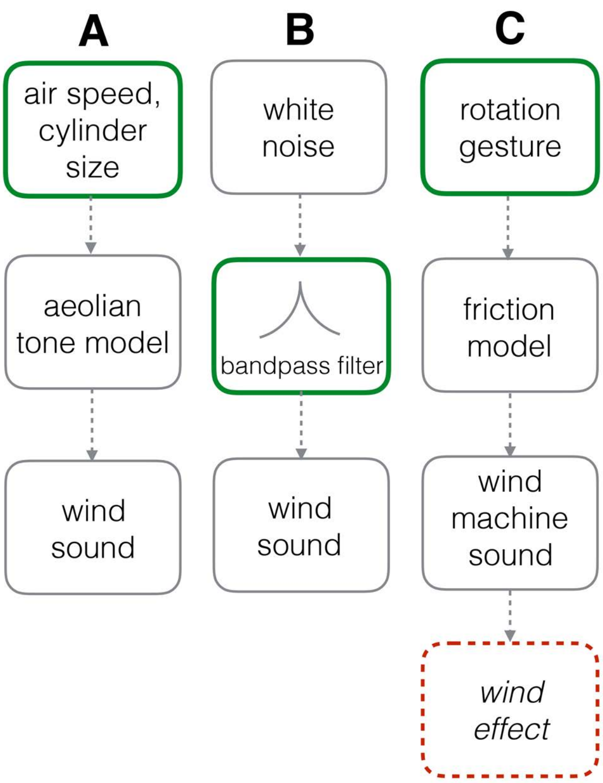

Figure 4.2: A comparison of approaches to digitally synthesizing a wind sound from A) Selfridge et al. (2016), B) Farnell (2010), and C) this research. The green outline denotes which parameters performance data might be mapped to.

This programming approach may not be as computationally efficient or even as sonically rich as its aeroacoustic or signal-based counterparts, but crucially it produces the sound of wind only through a continuous sonic interaction. Expressive human action is intimately linked to the activation and modulation of the wind-like sound in performance. Although beyond the scope of this research, it is tempting to speculate whether developments in the science of aeroacoustics in the nineteenth century directly influenced the wind machine design (Beyer & Raichel 1999, p.95). As will be explored in Chapter 5, the motion of the cylinder and cloth of the acoustic wind machine are not explicitly described as an attempt to imitate the interaction between real-world moving air and cylindrical obstacles.

The SDT suite of physical modelling algorithms is used within the software program Max/MSP.[1] Max/MSP is a graphical programming environment, with data flow facilitated by drawing connections between objects with patchcords. Programming in this visual way allows for an intuitive workflow and fast iterations, and processes can be quickly reconfigured to tailor the response of the prototype digital wind machine in performance. As will be explored further in Chapter 6, this prototyping work engages explicitly with design problems and tensions inherent in the programming process, and makes space for further questions to be generated as the digital model takes shape. This is facilitated with a close examination of iterations of the design work, framed by a consideration of how the sound is placed *in-hand* (Section 4.1.4). As will be explored as part of the discussion on potential future work arising from this research in Chapter 8, the prototype in Max/MSP could potentially be translated to another programming language, such as C++, as part of future development work.

Having established the programming approach that informs the design of a prototype digital synthesis engine to imitate the sound of the acoustic wind machine, we will now move on to examine how that synthesis engine is activated and modulated in real-time performance.

4.1.4 Performing Acoustic and Digital Wind-Like Sounds

This chapter has so far outlined the methodology pursued to remake a working example of an acoustic wind machine and then program a prototype digital synthesis engine that

[1] https://cycling74.com/

models its wind-like sound. We will now consider how the digital synthesis engine is made performable in real time. As previously discussed in Chapter 2, performing a digital sound with more conventional interfaces often implies a loss of tactile feedback or sensation of physical effort when compared with the experience of manipulating an acoustic musical instrument, something which prior research in DMI design has focused on addressing. The choice of interactive object that places the digital wind-like sound in-hand in performance is therefore critical to the meaning that might be created during a continuous sonic interaction.

However, with a working acoustic wind machine to perform with and explore, it is possible to begin to close the gap between the fully acoustic and fully digital experience of soundmaking so far explored in this thesis. Chapter 5 details an investigation by the author into the sonic interactivity of the acoustic wind machine, and the tactile and kinaesthetic experience it offers a performer. This close examination is informed by previous research in the field of anthropology, which considers the process of embodied skill acquisition. In "Walking the Plank" (2006), Ingold describes the process of sawing wood to create a bookcase, revealing the fine details of the material resistances and continual bodily adjustments that are inherently part of the process of woodworking. A similar process of observation and description is employed within this research to reveal more about the influence of the acoustic wind machine's crank handle and mechanism in making its wind-like sound feel truly in-hand.

Several potential methods of making an everyday wind sound performable have been touched on so far in the preceding chapters. These are gathered together here to examine the detail of the multisensory experience each method affords a performer. This allows the potential ways to create a wind sound through human action to be mapped as a succession of stages, beginning with the simplest, most physical and tactile interaction with handheld acoustic materials, and ending with the most virtual and digital interaction, when open hand movements generate data with worn sensors to produce a digital sound, and there is no experience of physical manipulation of an object or digital interface. The transitions between the most limited and acoustic, and the most free and digital methods can now be examined more closely (Figure 4.3).

1. A very simple wind-like sound is created by scraping paper directly over a surface (Fitzgerald (1881, p.63) and Dodge (1912, p.16)). The experience of soundmaking is simple and direct, with continuous multimodal feedback in performance.

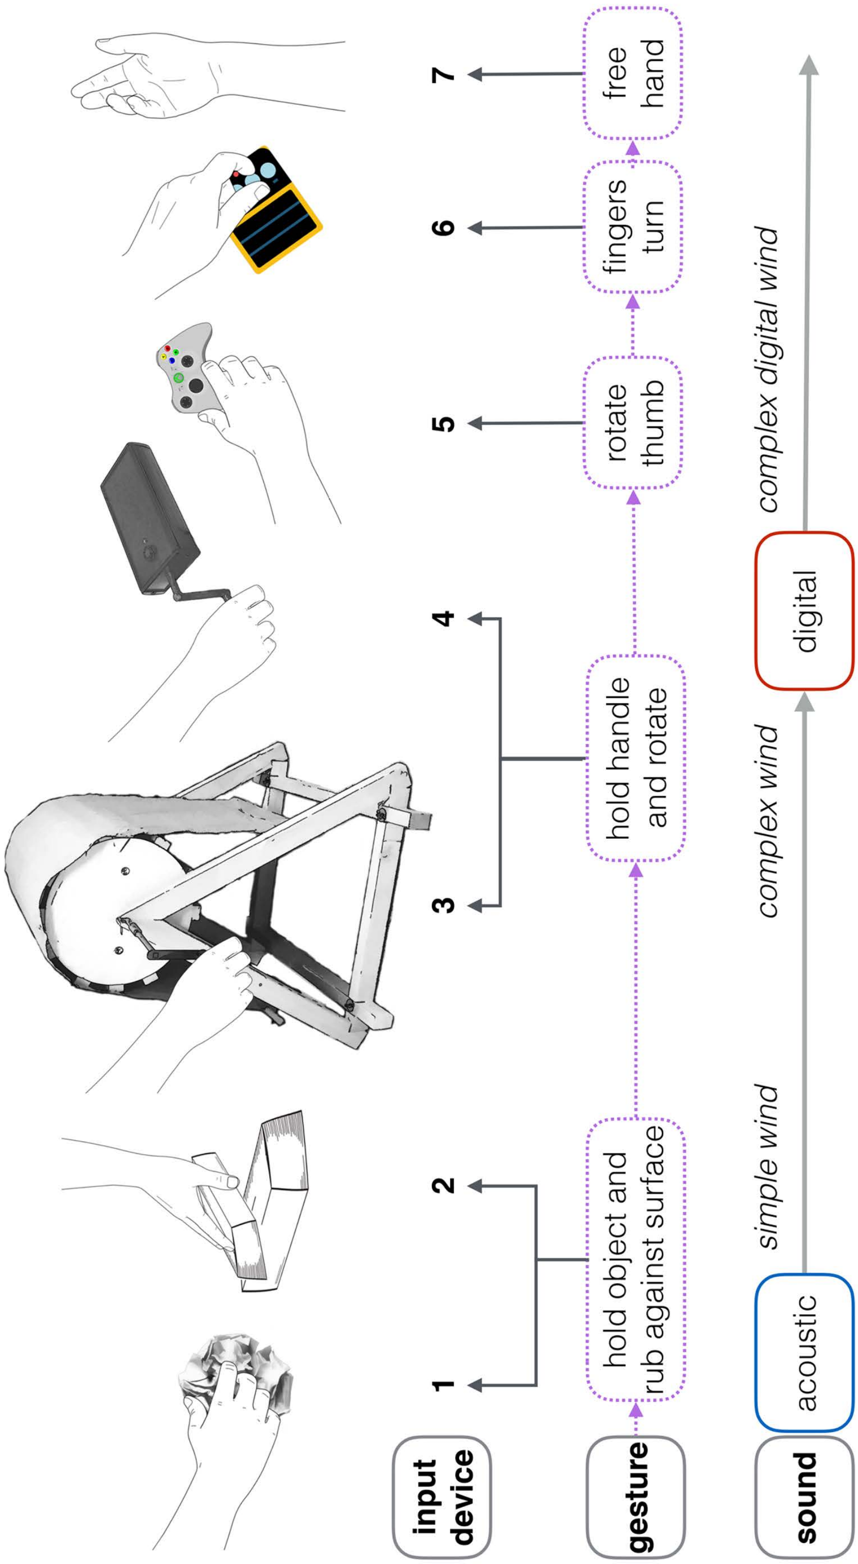

Figure 4.3: A framework for *sounds-in-hand*: Putting the control of a wind-like sound into the hands of a performer.

2. A single slat is held in the hand and scraped over stretched silk (Rose 1928, p.9). Again, the experience of soundmaking is simple and continuous. This time, the material being interacted with to make sound is not directly under the hand like the crumpled paper. Instead, the tactile feedback from the stretched cloth is slightly mediated by the wooden slat in the hand.

3. The stage 2 method is mechanised as an acoustic wind machine, an early performance interface, which produces sound through a gesture of rotation. This time, the scraping action is facilitated between the wooden slatted cylinder and a silk or canvas cloth. The experience of sound creation is intuitive and direct, again with multimodal feedback in performance. The performer experiences a direct coupling between action and sound despite the introduction of a handle into the process of scraping, and the resulting wind sound in-hand is more complex than at stages 1 and 2.

4. The acoustic wind machine is enactively recreated as a small digital crank handle interface driving a rotary encoder, which plays a digital synthesis engine based on a physical model of friction to create a wind-like sound. The experience of sound creation has the potential to be intuitive, but not as rich as the acoustic wind machine due to the relatively smooth rotation of the small crank handle, which meets no resistance from the rotary encoder and requires little effort from the performer.

5. The same digital synthesis engine is controlled with a digital interface fabricated from smooth and light plastic parts. Continuous sound is activated using a thumbstick. Although the movement is continuous, this engages only one thumb rather than the whole hand grasping and turning a handle. The gesture is smaller and more precise, and the full rotation is completed more quickly than with the digital crank handle.

6. The digital synthesis engine is now activated by turning an individual rotary encoder turned between two fingertips. The gesture of rotation has become very small and subtle, and the possibilities for complex sound creation may be limited due to the range of movement this hand position is capable of, in comparison with the experience of continuously rotating a crank handle or moving a thumbstick. The sound may be less continuous and stop or start as a result of the fingers re-adjusting their grip on the encoder as it turns.

7. Free hand gestures, captured as data from worn sensors or motion tracking, control the digital synthesis engine. The experience of sound creation does not offer any tactile feedback or material resistance, as the performer is unable to rely on the limitations inherent in the manipulation of a physical object. The performer can activate the digital sound with continuous motion in performance.

Mapping each of these methods of performing a wind sound reveals an obvious break between the experience of performing with the acoustic wind machine and that of the digital crank interface, which enactively recreates it. As will be explored further in Chapter 6, the digital crank replicates the same rotational gesture, but affords none of the tactile feedback or material resistance of the acoustic wind machine. It also imitates the way the acoustic wind machine produces sound, but produces a fully digital wind-like sound. These two changes present an interesting opportunity that directs the focus of this research.

Rather than exclusively creating a digital crank interface to perform the digital wind engine, and working to transfer as much of the richness of the acoustic wind machine's multisensory feedback to this system, the acoustic wind machine itself is used as the interface for the performance of the digital wind sound. This is achieved by coupling a rotary encoder to the acoustic wind machine's mechanism. The rotation data generated by the encoder is then read by an Arduino[2] prototyping board and routed to Max/MSP. What this approach achieves is the ability to isolate and change only one mode of multisensory feedback - the wind-like sound - and preserve all of the tactile feedback and material resistance of the acoustic wind machine when performing the digital sound. As outlined in Chapter 6, this allows for more precise perceptual and objective comparisons between the acoustic and digital sounds to inform the programming of the prototype digital wind engine. Section 4.2 examines how this configuration enables a single modality - sound - to be isolated as the subject of study in the design of an experiment to evaluate both of the wind-like sounds in performance. This addresses some of the issues encountered in previous research in separating modes of feedback from each other in an interaction to examine their specific effects (Giordano et al. 2012).

[2]Arduino is a physical computing environment that facilitates fast prototyping.

4.1.5 Recording Simultaneous Performances of Wind-Like Sounds

With the acoustic wind machine also activating its digital prototype in performance, a single gesture can simultaneously activate both wind-like sounds, which in turn can be simultaneously recorded. Capturing wind performances as recordings to facilitate reflection, acoustical analysis or to create experiment stimuli is an important part of the work discussed in chapters 5, 6 and 7. To aid the understanding of the reader, the recording setups discussed in the following chapters will be briefly outlined here:

- Chapter 5 describes work undertaken to compare the acoustic wind machine in performance with a field recordings of natural real-world wind sounds. To facilitate this exploration, field recordings are delivered over loudspeakers while the author performs and rehearses listening to both the natural wind sounds and the acoustic sound of the wind machine in the studio. Later, the field recordings are delivered to the author over headphones from Pro Tools[3] software while the acoustic wind machine is simultaneously captured with a microphone and recorded to the same Pro Tools session. The headphone feed is split, with the natural wind delivered to the author in one ear and the acoustic wind machine delivered to the author in the other ear.

- Chapter 6 describes work undertaken to acoustically evaluate and compare the acoustic wind machine's sound with that of its developing digital prototype in Max/MSP. To facilitate this, the same performance gestures are simultaneously captured from both the acoustic wind machine and the program in Max/MSP. The acoustic wind machine is recorded with a microphone into Pro Tools, and the Max/MSP program running on a second computer is routed via the line outs from an additional audio interface into the same Pro Tools session. These recordings are then acoustically analysed.

- For the evaluation with participants discussed in Chapter 7, the simultaneous recording setup used to generate audio for acoustical analysis is also used to create stimuli of wind-like sounds for participants. This includes recordings of the same performance gesture with the acoustic and digital wind-like sounds in response to a field recording of natural real-world wind. When participants perform with both

[3] Digital audio workstation software from Avid (https://www.avid.com/pro-tools).

wind-like sounds during the experiment, their performances are recorded in the same way to produce a corpus of audio examples which can be used for acoustical analysis.

This section has outlined the design-led methodology that is pursued as part of this research. The sonic interactivity of the theatre wind machine design will be examined through:

- Remaking a working example of an acoustic wind machine based on the instructions of historical practitioners, exploring the multisensory experience it affords in performance, and examining how exactly it produces sound in response to a performer's action.
- Using this information to inform the design of a prototype digital synthesis engine, programmed in Max/MSP using the SDT model of friction, to imitate the wind-like sound of this specific acoustic wind machine example.
- Fitting the acoustic wind machine with a rotary encoder and Arduino to capture its movement as data, allowing the acoustic and digital wind-like sounds to be simultaneously performed and recorded.

These stages of design work are discussed in detail in Chapter 5 and 6. Having established how the design-led methodology of this research is pursued, the next section moves on to consider how the acoustic and prototype digital wind machines are evaluated, both as sound effects that imitate wind and as continuously interactive sounding objects.

4.2 Evaluating Wind Effects

As discussed in Chapter 2, evaluation is an important part of the process of developing the design of a continuous sonic interaction with a digital sound. Given that the designer will have created an interactive sounding object with the aim of making meaning for the user or performer in an interaction, its ability to do so should be verified. How others perceive the sound model will need further confirmation. The meaning created in the interaction should also be investigated further in an experiment with participants.

Two main avenues of evaluating the design work to create the acoustic and prototype digital wind machines are pursued as part of this research. First, during the process of programming the prototype digital wind machine, the main characteristics of its sound are compared with those of the acoustic wind machine in order to refine its configuration and calibration. Secondly, the two wind machines are subjected to an experimental study with participants. This examines the perceptual qualities of their wind-like sounds more closely, and also seeks feedback from participants on their experience of performing both sounds. These two approaches will now be discussed in more detail.

4.2.1 Acoustical Evaluation

This research makes use of available analysis methods from musical acoustics in order to examine and describe the sound of the acoustic and digital wind machines more precisely. Theatre sound effects have not yet been subjected to an acoustical study, and so this method of analysis helps to establish how the sound of the acoustic wind machine is activated and modulated by the rotation of the crank handle. This informs the configuration and calibration of the prototype digital wind machine as it is programmed. With a rotary encoder and Arduino coupled to the acoustic wind machine, the same performance gesture can activate both the acoustic and digital wind-like sounds, which are then simultaneously recorded to facilitate their comparison. The approach taken to this simultaneous recording has been outlined in Section 4.1.5. This helps to examine not just how closely the digital sound resembles its acoustic counterpart in performance, but also whether the digital wind-like sound activates and modulates efficiently enough in time when compared to its acoustic counterpart. This simultaneous performance setup is also used as part of the experiment to capture and compare participants' own performances of the wind-like sounds.

Acoustical measurements allow the features of digital sound recordings to be analysed and expressed graphically and numerically. This allows an objective examination of the properties of a sound. The MIR Toolbox audio analysis environment (Lartillot & Toiviainen 2007) in MATLAB [4] is used in this research to analyse the recorded acoustic wind machine and its digital counterpart, reveal their acoustic features and compare them to each other. It is also used to compare the progress of the acoustic and digital wind-like sounds with recordings of real-world wind, in order to examine how the effect of wind is created through human performance.

As explored in Chapter 3, the theatre space was itself an important part of the design of

[4]Software for analysis, programming and modelling by MathWorks.

sound effects in the late nineteenth and early twentieth century. This research is primarily concerned with the experience of the performer in a continuous sonic interaction with a wind machine rather than the audience listening to that performance. As such, the potential influence of the acoustical properties of a performance space on the resulting wind sound is not included in this evaluation. This may be an interesting avenue for future research, however. To facilitate as accurate a comparison as possible between the acoustic and digital wind sounds, the acoustic wind machine is recorded in a dry acoustically treated studio space for this part of the evaluation. The results of the acoustical analysis and comparison are reported in Chapter 6. The main audio descriptors used to analyse the acoustic and digital wind-like sounds will now be briefly outlined.

4.2.1.1 Spectral Measurements

A spectrum (Figure 4.4) displays which combinations of audible frequencies make up a specific sound. A spectrogram (Figure 4.5) displays how the frequencies of the sound vary with time.

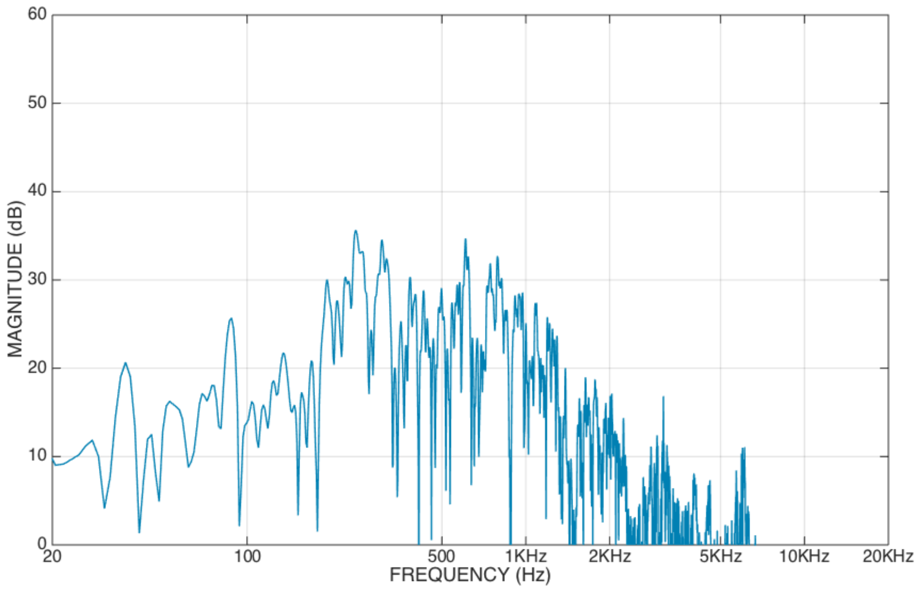

Figure 4.4: A spectral analysis of the sound of a hammer hitting a nail, produced with MIR Toolbox in MATLAB.

Some further measures that numerically describe characteristics of the spectrum are also used in this research to express its features. These are:

- *Brightness:* This measures how much high frequency energy the spectrum has above 1500Hz, expressed as a number between 0 and 1.

- *Inharmonicity:* This is a simple measure to show how noisy a sound is relative to its fundamental, or lowest frequency. In the case of the wind effects, their fundamental frequencies will be their modelled Aeolian tones. This is expressed as a number between 0 and 1.

- *Spectral Centroid:* The statistical centre of mass of the sound, or the frequency where the most energy of the sound is located. In the case of a wind effect, this may be where the Aeolian tone is to be found. This is expressed as a frequency in Hz.

- *Spread:* This is the statistical measure of the standard deviation, or the square root of the variance, of the spectrum. It shows how spread out the spectrum of the sound is, or how wide a range of frequencies it contains, relative to the spectral centroid. It is expressed as a value in Hz.

- *Skewness:* This shows whether the spectrum of the sound is clustered towards the low or towards the high frequencies. It is another statistical measure expressed as a value between -3.0 and 3.0.

4.2.1.2 Amplitude Envelope Measurements

The amplitude envelope (Figure 4.6) shows how the loudness of a sound evolves in time.

One further measure is used to give a numerical value to the envelope of the wind sounds. This is *event density* - a measure of the average number of onsets (the beginning of a sound) in a sound's amplitude envelope per second.

Having established which acoustical measurements are used to examine the acoustic and digital winds more precisely, we will now explore the experiment design and statistical analysis part of the evaluation work.

4.2.2 Experiment Design

This research incorporates an experiment to evaluate both the acoustic and prototype digital wind machines in performance with a number of participants. The experiment is designed to explore how participants perceive the acoustic and digital wind-like sounds, as well as how they experience a continuous sonic interaction with each of them. Real-world wind sounds have not yet been subjected to a dedicated study to examine their underlying perceptual structure. Theatre sound effects have also not yet been studied for their perceptual qualities. While this leaves the way open to design a bespoke experimental

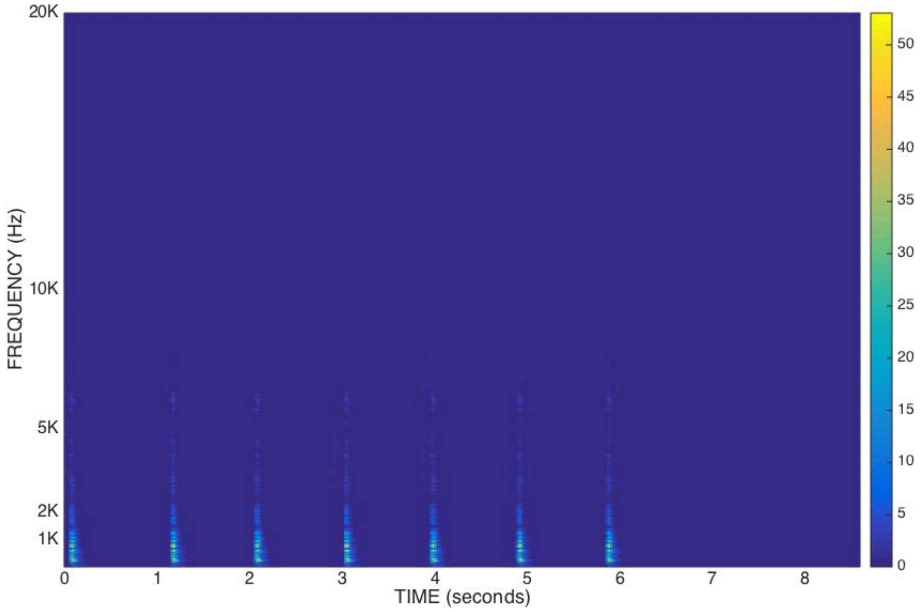

Figure 4.5: A spectrogram of the same sound of a hammer hitting a nail, produced with MIR Toolbox in MATLAB.

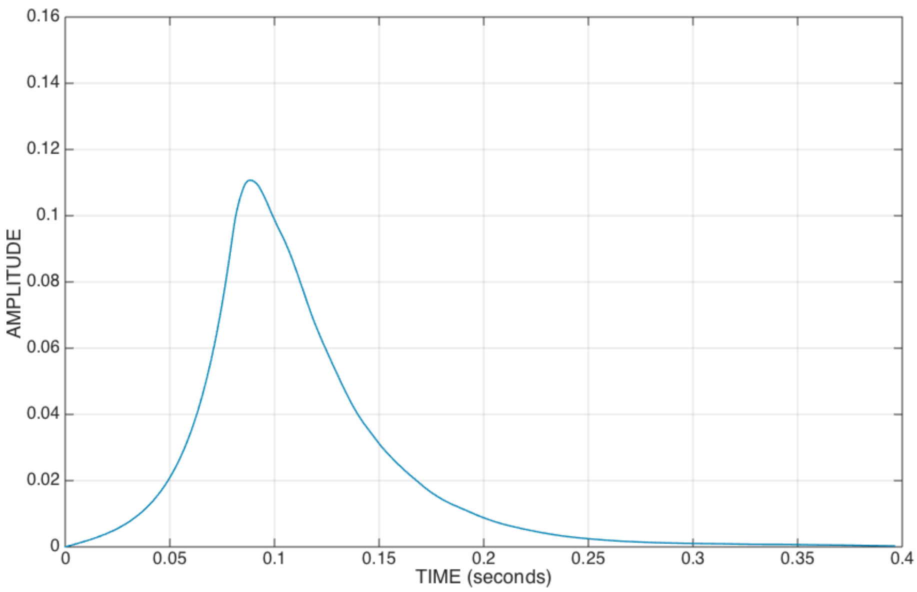

Figure 4.6: An amplitude envelope of the same sound of a hammer hitting a nail, produced with MIR Toolbox in MATLAB.

procedure, it is also important to ground this part of the evaluation in prior research so that any results can be fully understood.

The experiment design therefore builds on procedures from prior research into the perception of environmental sounds, SID, and the evaluation of DMIs. The design-led methodology of remaking, modelling and simultaneously performing the acoustic and prototype digital wind machines detailed in Section 4.1 offers some unique opportunities in terms of evaluating a continuous sonic interaction. Given the apparent simplicity of the wind machine design and the everyday sound it produces, non-specialist participants with no musical experience can be recruited for this study. This allows for a greater number of participants and the use of robust statistical testing examining the gathered data in addition to qualitative measures inviting free descriptions of their experiences. As the acoustic wind machine will be able to perform both the acoustic and digital winds, only one mode of feedback in the interaction - the wind-like sound - can be changed to examine the perceptual result. This facilitates a more precise examination of how the prototype digital wind machine compares with its acoustic counterpart in performance.

Guidelines by Bonebright et al. (2005, p.2) recommend that the task time for a perceptual test should not exceed 30 minutes if possible, as this kind of testing is onerous on participants. Given that sound must be carefully listened to or performed and then described, fatigue or loss of motivation can impact the data gathered. Mindful of this, the experiment undertaken as part of this research is designed around three distinct steps. First, two listening tests are presented for participants to complete. Then, following a short break, participants take part in a performance test. The full experiment takes approximately 40 minutes to run.

Before the research background that informs each step of the experiment is considered here in more detail, the statistical analysis concepts that inform the experiment design will be briefly outlined for the understanding of the reader, using definitions found in Field & Hole (2002) and Field et al. (2012).

4.2.2.1 General Statistical Analysis Concepts

Experiment design involves the control, manipulation and measurement of variables. This facilitates an analysis of whether an independent variable (e.g. the amplitude of an audio stimulus played to a participant) has an effect on a dependent variable (e.g. the perceived loudness of that audio stimulus, which will be scored on a rating scale by a participant).

Statistical testing can help to understand if such an effect has indeed occurred. Robust statistical testing is useful particularly in this evaluation of the acoustic and prototype digital wind machines as there is no prior research specifically investigating a performable wind-like sound on which any assumptions can be based. The statistical significance of a statistical test, also known as the *p-value* or α, is used in this research to confirm whether an observed effect in the gathered data is systematic, or whether it may also have occurred by chance. The minimum acceptable level for α is taken as 0.05.

To ensure that the statistically significant result is measuring an effect of meaningful magnitude, the effect size is also measured and reported. This allows findings to be compared across different studies. The minimum acceptable level is a large effect, or $r = 0.5$. Finally, a calculation of statistical power is used to ensure that enough participants are included in the study to be able to detect the expected effect accurately. As a general guideline, a minimum of 28 participants is required to be able to detect a large effect size ($r = 0.5$) at $\alpha = 0.05$ with a statistical power of 0.8. In order to ensure the integrity of the measurements in this underexplored area of evaluation and facilitate several different statistical measures, a total of 54 participants were recruited for this experiment. As will be reported in Chapter 7, some participants are excluded from particular steps of the experiment due to issues with data collection during the procedure. All of these exclusions, manipulations and measures are fully reported, along with the results of the statistical testing.

The background to each step of the experiment, and the particular methods of statistical analysis used to evaluate the results, will now be outlined.

4.2.2.2 The Perceptual Space of Wind Sounds

The acoustic and digital wind-like sounds are first compared to each other in a listening test to understand how similar participants perceive them to be. Similarity ratings like this are useful for revealing the underlying perceptual structure of a group of sounds (Bonebright et al. 2005, p.5). This procedure builds on previous research examining the perceptual structure of a large group of everyday sounds. Gygi et al. (2007) used similarity ratings tasks to show that participants categorise everyday sound events into three distinct clusters (impacts, continuous sounds and harmonic sounds), and that this holds true both for sounds that participants have listened to and sounds that they have been asked to imagine. In the first step of the experiment conducted as part of this research, participants

1	2	3	4	5	6	7
Not similar at all						As similar as they can possibly be

Table 4.1: The similarity scale used in this research, based on the one implemented in a study of environmental sounds by Gygi et al. (2007).

are asked to compare sounds in pairs and rate how similar they perceive them to be to each other on a similarity scale (Table 4.1). This kind of scale is known as a Likert scale.

Each sound is compared with each other sound in the group of stimuli. Further analysis of the resulting data can reveal more about which acoustical features of the sounds informed the decisions they made. The procedure used within this research closely follows that of Gygi et al. (2007). The similarity ratings are used to produce a matrix of similarity values averaged across the participants. Multidimensional scaling (MDS) analysis is then used to produce a visual representation of the perceptual structure. MDS is a method of exploratory factor analysis. Factor analysis attempts to measure "latent variables," or those that cannot be measured directly (Field et al. 2012, p.750).

In the case of a group of sound stimuli, MDS visually represents the perceived similarity of the sounds to each other in terms of their distance from each other (Bartholomew et al. 2008, p.55). The analysis arranges the sound stimuli as individual points in a 2-dimensional space. Sounds that are perceived to be similar to each other will be placed closer together, and those that have been rated as very different from each other will be placed further apart. The resulting coordinates of each sound can then be tested against the acoustic features of the sound stimuli themselves, suggesting which factors have influenced the similarity ratings. Research by Vanderveer (1980) concluded that temporal patterning (the evolution of the sound over time) is the dominant feature in sound similarity, for example.

Rather than focus this listening test exclusively on the acoustic and prototype digital wind-like sounds, the group of stimuli incorporates some other sounds in order to contextualise the perceived similarity of the wind effects and how participants ratings compare with previous research into the perceptual structure of a broader range of everyday sounds. First, natural wind recordings are used to guide performances of the acoustic and digital wind-like sounds, which are simultaneously recorded as per the setup described in Section 4.1.5, producing a corpus of wind sounds with a natural, acoustic and digital component. This allows participants to compare a natural recording of wind to both an acoustic and digital imitation of it, offering further insight into the sonic affordance of a natural wind sound. Two other sounds are added to the group of stimuli to help contextualise the perceived similarity of the continuous wind sounds. A harmonic sound (a horn) and an impact sound (wood being dropped) are both included in this test. Although these sounds are from different perceptual categories, they share a similar amplitude envelope in that the audio clips include a repetition. This helps to further clarify the temporal patterning factor when analysing the final similarity ratings.

We will now move on to consider the second listening step of the experiment.

4.2.2.3 Rotational Speed of Wind

A wind machine's crank handle affords a gesture of rotation to the performer, and variations in its rotational speed produce changes in the resulting wind sound. This step of the experiment aims to understand two distinct aspects of this process. First, whether the connection between rotational movement of a certain speed and a continuous sound like wind is perceptually meaningful to participants. Secondly, whether the acoustic and digital wind-like sounds can equally communicate variations in the rotational speed of their machine interface to participants. It is important to investigate this as the final step of the experiment guides participants to perform with the acoustic and digital wind-like sounds by asking them to first listen to a recording of a wind performance. They must perceive a connection between the wind sound they listen to and a possible gesture with the crank handle if they are to successfully imitate what they have heard.

This method follows part of the procedure used to evaluate the interactive sounding object Spinotron (Lemaitre et al. 2009) introduced in Chapter 2. Before using Spinotron in an interactive task, researchers first evaluated participants' ability to perceive variations in the speed of its virtual ratchet sound model. Participants estimated the rotational speed of recordings of the ratchet sound model moving at various RPM (rotations per minute). Researchers were able to establish that participants could estimate the rotational speed of the model just by listening to its sound. In the case of the wind machines, it is not clear exactly how or if their sounds can communicate rotational speed to a listener, so this test will also help to establish if this is the case. This is also another way of comparing the acoustic and prototype digital wind machines to see if the digital model is adequately recreating the soundmaking components of its acoustic counterpart.

To produce the stimuli for this test, the acoustic wind machine is performed at different

1	2	3	4
Slowest			Fastest

Table 4.2: The speed ranking scale used for the second listening step of the experiment.

rotational speeds, and the acoustic and digital wind-like sounds simultaneously recorded as per the description of the setup in Section 4.1.5. During the production of these stimuli, the author listens to a split headphone setup, with the acoustic wind-like sound delivered in one ear and the digital wind-like sound delivered in the other. Given the size and weight of the acoustic wind machine and the fact that it is hand-operated, it is not be possible to record many different sounds at precise RPMs. Instead, four distinct rotational speeds are captured as recorded wind-like sounds. In order to link these sounds to the speed of the crank handle, participants are told that the sounds have been produced by a handle being turned at different speeds. Participants are then asked to rank the acoustic and digital wind-like sounds separately, from slowest to fastest, on a 4-point rating scale (Table 4.2).

This facilitates an exploration of whether participants can imagine a virtual rotating mechanism producing the wind sounds they hear. The results from this test are compared and summarised visually using a boxplot. Statistical test is also used to analyse the results. A Friedman rank-sum test (Field et al. 2012, p.688) is performed on participants' rankings to examine whether there is a statistically significant difference between the ranks given to each of the wind-like sounds. A further Wilcoxon signed-rank test (Field et al. 2012, p.668) is then used to establish which pairs of wind-like sounds have been ranked significantly differently from each other.

Next, the design of the performance step of the experiment will be outlined.

4.2.2.4 Performing Wind Sounds

Historical sources make many claims as to the simplicity of operation, the expressive potential and the efficacy of the wind machine as a performable sound effect. This step of the experiment attempts to verify some of these claims while also comparing the acoustic wind machine to its digital counterpart in a performance task.

The conditions of the experiment must also be controlled as much as possible in order to reduce any unwanted influences on participants perceptual experiences. The experimental setup makes the multimodal feedback as consistent as possible across the acoustic and prototype digital wind machine performances. The acoustic wind machine will constantly produce sound for example, even during a performance of the digital wind-like sound. To control this as far as is practical both wind-like sounds are delivered to participants in headphones as they perform. A microphone captures the sound of the acoustic wind machine, and the prototype digital wind machine sound comes from the program in Max/MSP. The headphones are closed, and so reduce the level of any external sounds during use. There will be a very slight spill of sound from the acoustic wind machine during each part of the performance, but this will be equal for both the acoustic and digital wind-like sounds.

Visual feedback is also controlled in this experimental step in order to focus the interaction purely on the gestural performance of wind afforded by the crank handle. Any sight of the acoustic wind machine's rotating wooden cylinder or moving cloth is removed. The machine is concealed within a cardboard structure with only its crank handle protruding. This should also eliminate the possibility that the sight of the unusually large wooden structure does not cause participants to make quick assumptions about the weight of the wind machine or the effort required to play it before they have a chance to try it out for themselves (Figure 4.7).

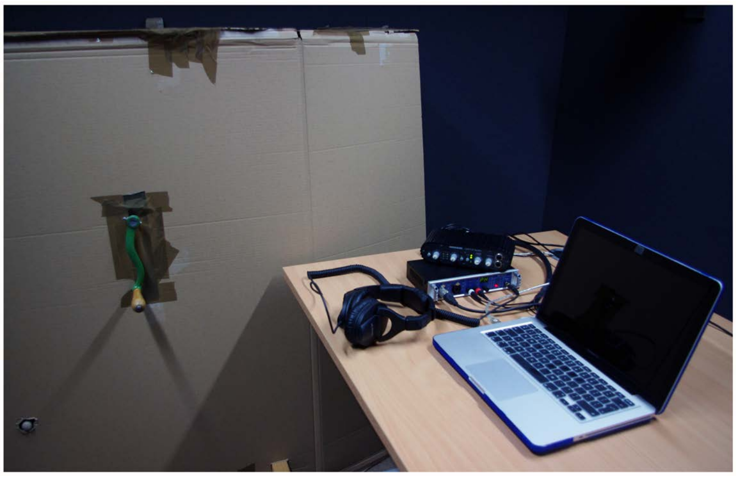

Figure 4.7: Concealing the acoustic wind machine to remove any visual feedback for the experiment.

This step of the experiment is based on a task that "operationalizes" (Poepel 2005) the expressive potential of performance with a wind machine. The task is designed for participants of varying expertise to complete, while making space for them to enactively learn how to perform wind in the restricted timeframe of the experimental setup. To focus participants clearly on the wind-like sound and how to perform it, recordings of two simple gestures performed both with the acoustic and prototype digital wind machines are used as stimuli to elicit performances. Participants are asked to listen to these sounds, and then imitate what they have heard by using the crank handle. The sounds are not identified as either *acoustic* or *digital* for participants, to ensure that the sounds can be experienced as they are presented without any prior judgement as to the quality of their origin.

To encourage participants to fully explore the expressive range of the acoustic and digital wind-like sounds in performance, they are asked to imitate both the matched source of the sound (e.g. acoustic wind machine imitating an acoustic wind recording) and its counterpart (e.g. acoustic wind machine imitating a digital wind recording). These performances are themselves recorded for later acoustical analysis. Participants are asked to evaluate and describe their own experiences of their performances throughout this step of the experiment, and rate their performances for similarity and easiness on 7-point Likert scales. They are also asked to describe each wind sound that they perform by selecting from a list of descriptive words. This produces qualitative data in the form of free descriptions and categorisations of the wind-like sounds, as well as numerical data to be analysed.

Participants' ratings are described and summarised visually using bar graphs. Statistical testing is then performed to analyse the results. To ensure that the order in which the stimuli are presented does not influence how the acoustic or prototype digital wind machine are rated in terms of similarity to the stimuli or ease of play, participants are split into four groups, each with a slightly adjusted experiment order. A Kruskal-Wallis (Field et al. 2012, p.675) test is then performed to compare the ratings given by each group of participants to confirm that the order of presentation of the stimuli did not significantly influence their ratings. A further Wilcoxon signed-rank test (Field et al. 2012, p.668) is also performed to compare the ratings for similarity to the stimuli and ease of play given to the acoustic and prototype digital wind machines to establish if any difference between them is statistically significant.

Participants' descriptions of the wind-like sounds and their free descriptions are visually summarised and described. The recordings of participants' performances with the acoustic and prototype digital wind machines are then collated and categorised before being acoustically analysed using numerical descriptors from the MIR Toolbox in MATLAB (Lartillot & Toiviainen 2007). A Wilcoxon signed-rank test (Field et al. 2012, p.668) is then performed on the resulting data to examine whether the choice of stimulus (acoustic or digital) had an influence on how participants performed with the wind machines.

This section has presented an overview of the design of the experimental evaluation with participants undertaken as part of this research. It consists of:

- A listening step where participants compare pairs of sounds and rate them in terms of their perceived similarity.
- A further listening step, where participants rank the acoustic and digital wind-like sounds in terms of the rotational speed that they perceive has produced them.
- A performance step, where participants perform wind-like sounds in response to stimuli, rate their experience in terms of the similarity of their performance to what they heard and how easy they found the sound to play, and also describe the wind-like sounds that they have performed.

The full experimental procedure and results are presented and discussed in Chapter 7. Having established the evaluation-led methodology pursued as part of this research, this chapter will now be summarised.

4.3 Summary

This chapter has outlined the design-led and evaluative research methodology pursued to remake the wind machine design, prototype a digital model imitating its sound, and then use the acoustic wind machine's interface to activate and modulate the digital wind-like sound in performance.

The wind machine design is first *remade* from a synthesis of instructions found in historical sources, allowing the approach of several practitioners to be combined with the author's own embodied knowledge of soundmaking. This work is discussed in Chapter 5.

The working acoustic wind machine is then used to inform the programming approach in Max/MSP to create a digital model of its soundmaking components. This model is based on the SDT model of friction. The programming work, and an acoustical evaluation to compare the acoustic and prototype digital wind machine in performance, is reported in Chapter 6.

These two new working wind machines are then used in an experiment with participants in order to explore:

- The underlying perceptual structure of natural wind, along with acoustic and digital sound effects designed to imitate it.

- How a performative human action like rotation might be perceptually linked to a continuous everyday sound like wind.
- Whether the prototype digital wind machine is perceptually similar to the acoustic wind machine in performance.

The next chapter outlines the stages of work to remake a working example of an acoustic wind machine and explore the continuous sonic interaction it affords in performance.

Chapter 5

An Acoustic Wind Machine as Interactive Sounding Object

This chapter presents the first part of the design-led methodology discussed in Chapter 4. The continuous sonic interaction afforded by an acoustic wind machine, a theatre sound effect originating in the late nineteenth century, is explored in detail. Using a synthesis of several historical design instructions, a working example is constructed. This process examines the embodied knowledge of historical theatre practitioners in more depth, and expands on the descriptions of the wind machine design offered in historical texts. The working acoustic wind machine is then explored as an interactive sounding object. This exploration reveals further details on the multisensory feedback the acoustic wind machine offers in performance, and allows its process of sound production to be unpacked further to inform an entity-action model describing its workings. This work informs the approach taken to programming a digital synthesis engine to mimic the soundmaking components of the acoustic wind machine, which will be presented in Chapter 6.

Section 5.1 details work to produce a design for the wind machine according to a synthesis of disparate design instructions contained in historical texts. Section 5.2 discusses the process of constructing the wind machine and refining it to make it work. Section 5.3 explores the multisensory experience of performing wind with this sound effect, and examines the dynamics of its mechanism. Section 5.4 closely examines how the wind machine creates sound in an interaction. An entity-action model is then produced to describe the process of sound production in detail. Section 5.5 presents a summary of this chapter.

5.1 Designing an Acoustic Wind Machine

As outlined previously in Chapter 4, bringing together historical accounts describing the wind machine design shows that each practitioner adjusted the core principles of the sound effect to their own requirements, resulting in many different versions. Remaking is used as part of this research as a method of addressing these variations in construction approaches, as well as the diverse descriptions of the quality of wind sound this design is capable of producing. With a working example to explore, the wind machine's qualities as an interactive sounding object can be investigated more thoroughly to inform the programming of a digital synthesis engine to imitate its workings.

This section presents the work undertaken to design a working wind machine from historical design instructions. Chapter 3 examined the often short and simple nature of the descriptions of wind machines found in historical sources from the late nineteenth and early twentieth century, even in more instructional theatrical manuals. The brevity of these sources may point to some of the assumptions writers are making about the skills of their readers, as will be explored in this chapter. Similarly, little information is given on the origins of this cylinder-and-axle design. While the developing science of aeroacoustics did connect the diameter of a string or wire to the frequency of an Aeolian tone in the nineteenth century (Beyer & Raichel 1999, p.95), there is no reference to this in sources on the wind machine design that have been uncovered as part of this research. As outlined in Chapter 3, some practitioners do connect the simple action that produces a *scraping* sound to this design that produces a wind-like sound. The wind machine and the action that activates it are also not explicitly connected to an innate "sonic affordance" (Altavilla et al. 2013) of real-world wind either; rather, the design is described in terms of the materials used to create it.

In light of the scarcity of detailed information on the technique of constructing a working wind machine, four distinct historical sources were chosen to inform a new sketch for a design based on a synthesis of their observations and implementations. These sources provide some of the most detailed accounts of the wind machine's construction encountered in this research, and are drawn from across the late nineteenth and early twentieth century so as to include a range of theatrical circumstances and writing styles.

The four practitioners and their designs will now be introduced and their particular approaches examined more closely. The full texts of their descriptions can be found in Appendix A.

5.1.1 Jean-Pierre Moynet (1874)

Moynet was a trained architect, scene painter and stage designer in nineteenth century Paris. This wind machine is described in his 1874 publication on theatre production "L'Envers du Théâtre" (Baugh & Wilmore 2015, p.169) (Figure 5.1). This is the earliest historical source detailing the design of a wind machine that has been found in the course of this research. While not providing any measurements for his readers to describe the exact dimensions of the machine, Moynet does provide some useful details, and has drawn the illustration of the effect himself.

The cylinder's axle (the rod passing through its centre) is "resting upon two trunnions." A trunnion is a cylindrical protrusion that can be used as a pivot. As is suggested by the drawing, this means that the axle is extruded from the edges of the cylinder to also mount the cylinder on the frame. With a square frame like this, it may be that there is a groove cut in the wood for the trunnion to rest on as it rotates in response to the crank handle movement. Moynet also advises that the cylinder is assembled from individual wooden slats, and that there should be fifteen to twenty of these. The cloth in this case is "strong silk fabric," and bolts at each end can adjust the tension of the cloth. The cloth bolted at each end probably helped to fix the cylinder to the frame. Without the cloth in place, it looks like the cylinder could be detached from its frame, which might have helped with maintenance and adjustment.

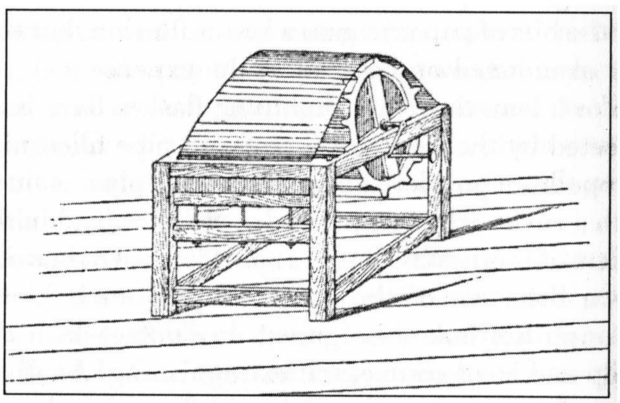

Figure 5.1: Wind machine design by Moynet, from his "L'Envers du Théâtre" (Baugh & Wilmore 2015, p.168). Image used here with permission of Christopher Baugh.

5.1.2 Van Dyke Browne (1913)

There is little biographical information to be found on Browne, who may have been a theatre practitioner in London in the late nineteenth century.[1] This wind machine design comes from the fifth edition of his theatrical manual “Secrets of Scene Painting and Stage Effects,” published in 1913 (Browne 1913, p.70) (Figure 5.2). Again, it is specified here that the cylinder (“drum” or “wheel” as Browne says) is assembled from individual wooden slats, but this time Browne specifies a distance of 3 inches between each slat.

The frame of this wind machine is made from “two uprights” that look like solid pieces of wood supported by a base and a rod across the top. Browne states that the crank handle must turn “easily,” and suggests that the cloth is fixed to the top rod and draped down freely over the cylinder. Fixing the cloth on only one side of the cylinder will have lessened the friction between the cloth and the slats on the loose side, and given the performer a different experience of resistance to the rotation of the crank when compared with Moynet’s wind machine.

Browne specifies moiré silk as the cloth, which is a particularly shiny silk with a watery appearance due to a unique finishing process. This must have been a particularly smooth but stiff fabric to use. The choice of fabric and the unique fixing of the cloth will have changed the quality of wind sound produced.

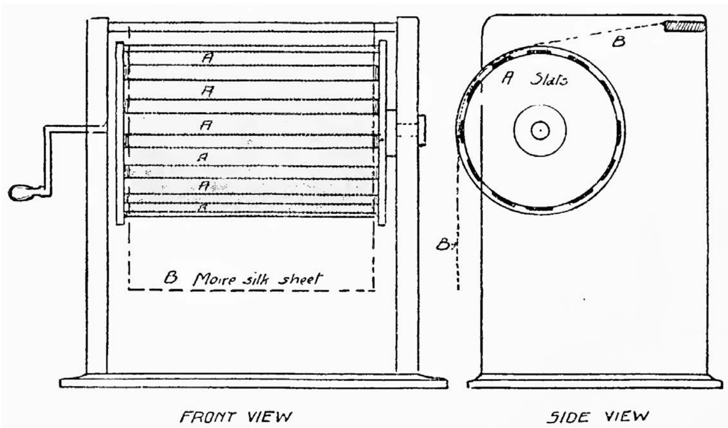

Figure 5.2: Wind machine design by Browne, from his “Secrets of Scene Painting and Stage Effects” (1913, p.70). Image included here courtesy of Taylor and Francis.

[1] *Vandyke brown* is also a term used for a particular shade of black-brown paint as early as the sixteenth century (Eastaugh et al. 2008, p.388), which suggests that this name may also have been a pseudonym for a scene painter.

5.1.3 Garrett H. Leverton (1936)

Leverton was a theatre producer, scholar and educator in the U.S. in the first half of the twentieth century. This wind machine design is from the first edition of his manual "The Production of Later Nineteenth Century Drama: A Basis for Teaching" (1936, p.50) (Figure 5.3). As a professor of theatre as well as a practitioner, Leverton himself surveys other historical sources to compliment his own production experience to inform this design. He even cites Browne's wind machine description as previously discussed in Section 5.1.2 (1913, p.70).

Looking at the included illustration, this wind machine design seems to consist of a solid cylinder, but Leverton does refer to it as a "cylindrical frame" in the text. What is clear is the similarity between this arrangement of the slats and that illustrated by Moynet (Baugh & Wilmore 2015, p.168). Leverton specifies the dimensions of the slats themselves, referring to them as "two-inch by one-inch batten." This specifies not the length, but the width and depth of the slats.

The cloth seems to be fixed to both sides of the frame in this example, and the frame is a triangular A-frame rather than the more square varieties we have so far encountered. The A-frame is a particularly useful choice, as the angled pieces of wood directly support the pivot point where the axle turns, making the wind machine very stable in performance. Leverton highlights that with a little practice this machine can produce "winds of various degrees of intensity and pitch." He also advises that different kinds of cloth will vary the sound, but only specifies silk as one of these.

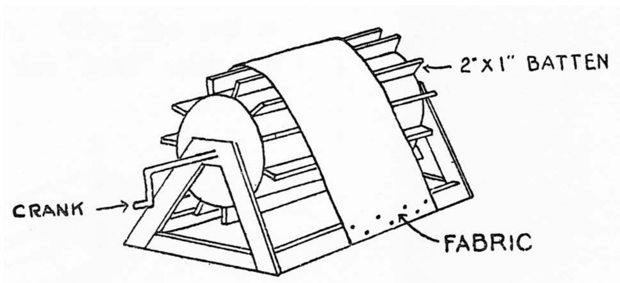

Figure 5.3: Wind machine design by Leverton, from his "The Production of Later Nineteenth Century Drama: A Basis for Teaching" (1936, p.50). Image included here courtesy of the Leverton Theatre Collection at Lake Forest College.

5.1.4 Frank Napier (1936)

Napier was an actor and stage director, notably at the Old Vic Theatre in London in the 1930s. This wind machine design comes from the fourth edition of his theatrical manual "Noises Off," which was first published in 1936 (1962, p.51) (Figure 5.4). Napier gives a little more detail in his description of how the wooden cylinder is constructed, with "two circles of wood connected by strips." He also says that these wooden strips have been "chamfered." This is a carpentry technique that involves paring down or shaving off a right angle on a piece of wood. This suggests that there is some smoothing involved in getting the cylinder to rotate properly.

The cloth, specified as "sail-canvas," which will be heavier than the silks suggested by the other practitioners, is again fixed to only one side of the machine's frame. The illustration shows the canvas weighted on one side with a batten threaded through its end, and it is referred to as being "stretched," which suggests that the canvas is always under tension. Napier draws his readers' attention to the way the overall machine is constructed, as it apparently will produce unwanted sounds like "bumps and squeaks" if the fixings of the frame and cylinder "work loose" through the course of many performances (1962, p.52). He suggests the use of bolts, so that they can be tightened if too much unwanted movement of the frame or axle occurs.

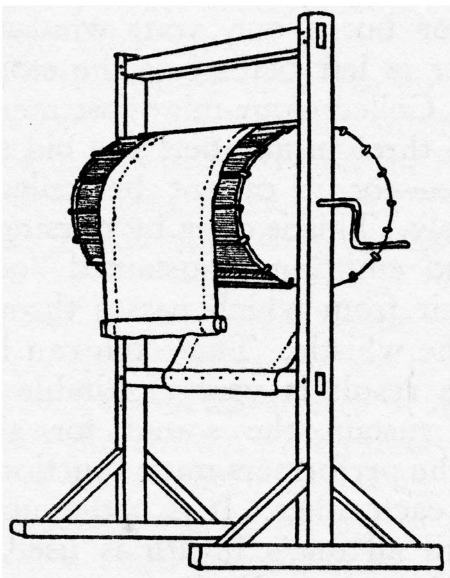

Figure 5.4: Wind machine design by Napier from his "Noises Off" (1962, p.51). Image used here with permission of Cressrelles Publishing Co. Ltd.

5.1.5 Specifications for a New Design

Comparing these historical sources reveals the distinct choices made about the kind of material used for the wind machine's cloth, the way its cylinder and frame are constructed, and how the motion of the slats against the cloth is facilitated. There are, however, some notable omissions. The lack of detail on how the crank handle is made, how it connects to the axle, and how that axle is fixed to the cylinder in order to move it in rotation shows that readers of these texts were presumed to have a certain level of mechanical knowledge, or perhaps to have encountered other crank-and-axle designs previously. The writers of these theatrical manuals may not have been experienced carpenter-machinists themselves, and so concentrated on the principles of the sound creation method rather than how the construction of the final mechanism ensures that the right sound effect is created.

The ability to make a large and heavy wooden machine responsive to variation in performance, and reliable across many performances, will have been crucial to the success of these wind machine designs. Only Napier hints at the carpentry technique and skill involved in his discussion of chamfering and the maintenance required to keep the wind machine playable (1962, p.52). The initial design specifications for the remade acoustic wind machine attempted to address some of these omissions as well as accommodating some additional requirements. The aim was to produce a working example of an acoustic wind machine with readily available materials that was sturdy enough to be later subjected to a performance experiment:

- The tension of the cloth would be adjustable to allow the resulting wind sound to be refined further if necessary.
- The wooden cylinder would be mounted on an A-frame as in Leverton's (1936, p.50) example, in order to support its weight in rotation as efficiently as possible.
- The cylinder and frame would be constructed to allow for easy disassembly. This would make any necessary modifications easier following construction, and also accommodate the installation of any sensors needed to capture the movement of the wind machine as data.
- For additional portability, the acoustic wind machine's frame would be fitted with wheels and brakes to allow it to be manoeuvred easily while assembled, and then secured in place for performance.

With the four historical sources considered and the design brief for the acoustic wind machine now established, the discussion moves to consider the process of constructing this particular example and refining it to make it work.

5.2 Remaking and Construction

Bringing four distinct historical sources on the wind machine design together informed a design to guide the construction of a working example. Some further considerations also informed this process.

First, without any clear guidelines on the dimensions of the machine, the final size was approximated from the few examples photographed in historical sources such as Vincent (1904, p.417) and Scientific American (1913, p.378) and from the author's own height and ideal playing position.[2] The acoustic wind machine would be freestanding on its own frame and large enough to produce an audible wind-like sound when heard acoustically. The wooden cylinder would be approximately 0.8m long, and the frame would be just under 1m high.

Secondly, with no details available on the construction of the crank handle and axle, a Victorian table winder and a steel rod were sourced. The steel rod could be heated and hammered into a square shape to fit the socket of the table winder handle (Figure 5.5). This ensured that the crank handle would already be shaped to be very comfortable to hold and not negatively impact the overall experience of performing with the wind machine, as a purpose-built one might. The steel rod would be able to take the weight of the wooden cylinder and protrude far enough to hold it onto a frame, as Moynet's trunnions did (Baugh & Wilmore 2015, p.168).

Finally, the historical sources considered previously give no indication as to which kind of wood should be used to construct the cylinder and frame. The materials available for this research were more modern woods commonly used in the building trade, like white pine lumber (in 2" x 4" planks) and 18mm plywood. These relatively inexpensive softwoods are easy to cut and shape using tools but are not usually used in the creation of musical instruments, which are generally constructed from hardwoods.[3] To ensure that

[2]The author measures 5' 3" in height.

[3]Hardwoods come from deciduous trees like walnut or maple. These woods are used to create instruments like violins or acoustic guitars, which need thin resonating soundboards.

this wind machine would produce a loud enough sound, a more rough cloth would be used to increase the friction produced between it and the slats. If this did not prove sufficient, it would be possible to fix rougher material onto the outside of the slats, as in the design suggested by Laumann (1897, p.137) explored in Chapter 4.

The initial stage of work to construct the wind machine took place at the York Arts and Crafts workshop.[4] After reviewing the available workshop materials and tools, a cutting list, or a list of the parts with their dimensions, was drawn up (Figure 5.6). This, along with the initial design and the illustrations of the chosen designers to refer to, formed the basis of the construction work. Construction proceeded according to an iterative process to allow for any adjustments that might be required - the final size of the cylinder would be able to influence the size of the frame, for example. The wood was cut to length, planed[5] and sanded using hand tools and the available workshop machines.[6] Assembly and finishing was completed by hand.

The wooden cylinder was constructed first. Each of its faces was cut from 18mm plywood, and the outside edges sanded smooth. The two resulting plywood discs were then clamped together, and a central hole drilled through them to fit the steel rod axle. A large block of wood was then glued and bolted to the inside of each disc, and they were also then drilled through. Besides strengthening the axle hole through the plywood so it could fully bear the cylinder's weight, these blocks would also allow the axle to be fixed to the cylinder so that it could be rotated.

The slats were then cut from 2" x 4" lumber, which was split into 1" x 2" battens. Twelve of the battens were made to allow them to be placed more easily around the plywood discs in a simple clock face formation by hand. These ended up being more irregularly placed however, although within tolerances, to create a smooth rotating cylinder. The slats were arranged in four groups of three, with a slightly larger gap between each group (Figure 5.7). To ensure that the flat slats fitted snugly to the circular

[4]This was part of a 6-week Introduction to Woodworking course run in June 2015, run by Kevin Lordan at Yorkshire Hut Company.

[5]This involves applying a hand tool called a plane to the surface of the wood to shave it smooth with a sharp blade.

[6]The workshop's electrical bench sander made it very easy to shape the circular discs for the wind machines cylinder, for example. Tools like this helped to make the woodworking more accurate and efficient in the absence of years of carpentry skill for the author to draw on, but of course would not have been available to historical practitioners.

Figure 5.5: Fitting a Victorian table winder handle to the steel rod axle.

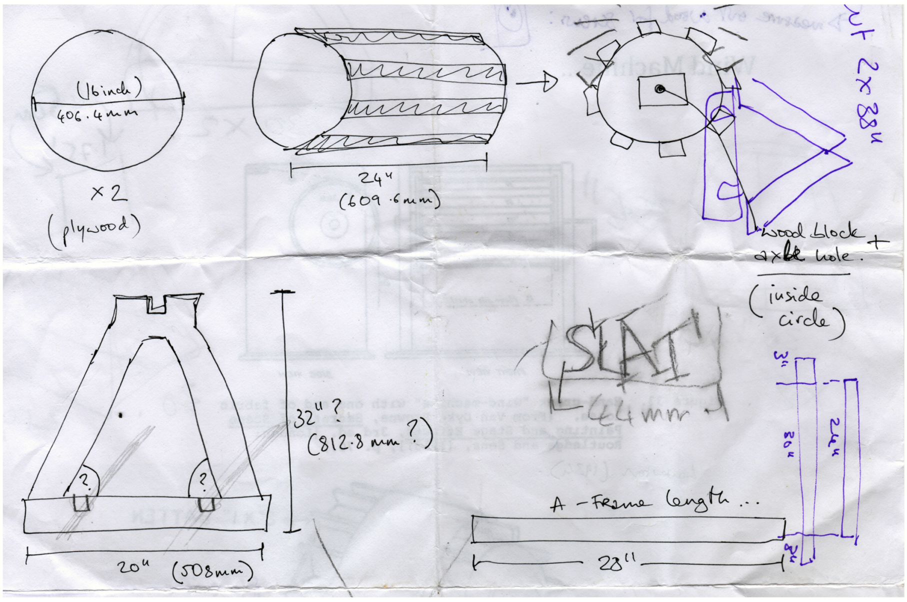

Figure 5.6: An extract of the cutting list produced for the acoustic wind machine, with layers of notes added during the construction process as the parts took shape.

shape, the disc was sanded flatter where it contacted the underside of each slat. Each slat and contact point on the plywood disc was marked with a code so that they could be matched to their original position in the event of disassembly. The slats were then fixed in place with screws. Each of the screws was countersunk so that it sat below the top of each slat so as not to catch the cloth during rotation.

The axle was threaded through the central hole of the cylinder and the slat screws were tightened up together to balance the tension of the cylinder evenly. Following the cylinder construction, the A-frame was built from 2" x 4" lumber. Two triangular shapes of equal height were constructed, with some parallel pieces of wood bolted between them to create a secure frame. This would allow the frame itself to be disassembled if needed. Two deep grooves were cut in the top of the frame to mount the wooden cylinder via its steel axle, enabling it to be rotated freely. A roughly woven heavy cloth was chosen to complete the design.[7]

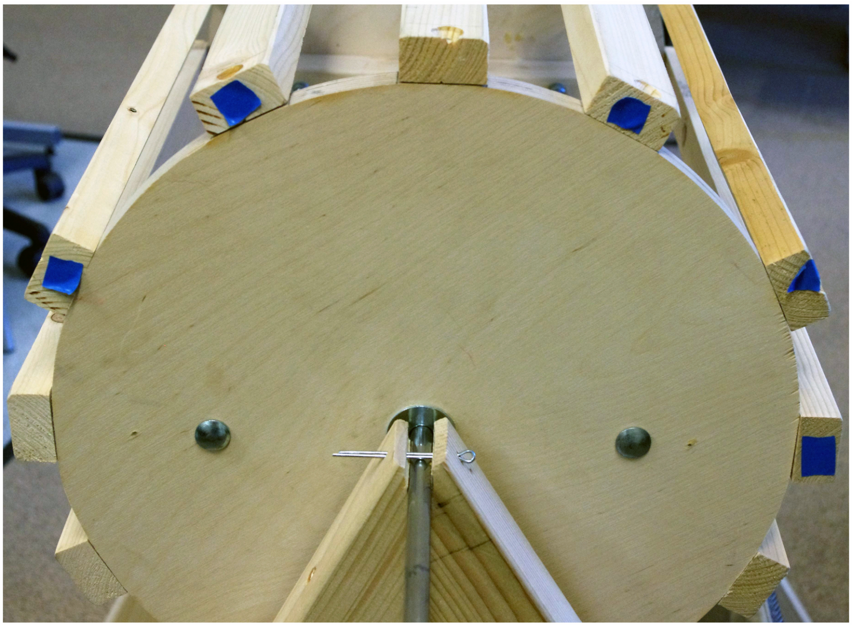

Figure 5.7: The final placement of the slats on the acoustic wind machine.

Once the cylinder and frame had been built and the cloth had been chosen, the parts were transferred to the Department of Theatre, Film and Television at the University of York for the next stage of construction. First, corresponding holes were drilled through the wooden blocks inside the cylinder and through the steel rod axle itself. The axle was

[7]Fittingly, this cloth had once served as a backdrop for a local theatre.

pinned tightly in place with a split pin at each end, threaded through the blocks of wood and the holes in the axle (Figure 5.8). The axle was hammered into a square shape at one end to fit the socket of the crank handle, allowing the cylinder to be rotated by hand once mounted on the frame.

Then, the cloth was cut to length and pockets were sewn in at each end. A wooden pole was then threaded through one end of the cloth and fixed to one side of the frame with another pair of split pins. The other side of the cloth was weighted with a piece of threaded steel bar slipped into its pocket. A split pin was also used at either end of the axle to fix the cylinder to the frame during rotation. Finally, castor wheels with brakes were fitted to the base of the frame to allow the acoustic wind machine to be moved and positioned easily when assembled.

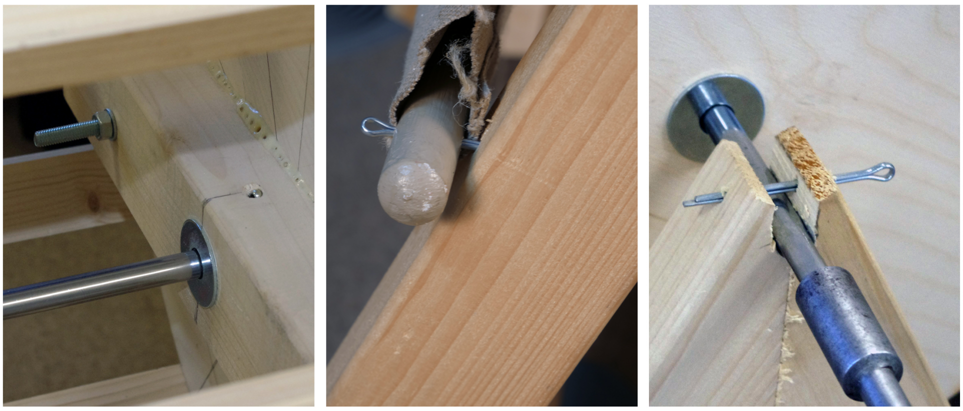

Figure 5.8: Split pins used to (L-R) fix the steel rod axle to the cylinder, fix the cloth to the frame, and fix the rotating axle onto the frame.

Although the acoustic wind machine was complete, some additional adjustments were required to refine the build and make it work. Firstly, the wooden slats of the cylinder did not immediately scrape against the cloth when the crank handle was rotated. They instead caught in the cloth, slowing and stopping rotation. To resolve this, the slats were sanded with fine sandpaper until their flat surfaces were very smooth. The exterior right angles of the slats were also chamfered as per Napier's advice (1962, p.51). This made the cylinder a more successful *rubber* against the cloth, allowing the handle to be rotated (Figure 5.9).

Secondly, the crank handle initially required too much effort to rotate with any meaningful variability in speed. The handle felt very heavy to turn and the cloth seemed to be very resistive to the movement of the slats, even after sanding and chamfering. This did not chime with the historical descriptions discussed in Section 5.1, which highlighted that the cylinder should move easily when the crank handle was rotated. This issue was addressed by firstly removing the steel rod weighing down the freely hanging side of the cloth. This reduced the downward force on that side of the cylinder and allowed it to move more easily.

After a period of performance practice with the wind machine, it was discovered that the crank handle was positioned too far from the pivot point of the axle on the A-frame. This was increasing the effort required to turn the crank handle unnecessarily, and also causing the axle itself to twist between the handle and the pivot point in response to the rotational movement. To address this, a new axle and crank handle were produced to replace the first iteration. The new crank handle came from an iron kitchen meat grinder, and could be positioned much closer to the pivot point. This lightened the load in the hand, enabling the crank to be turned more easily and at higher speeds, and also made the acoustic wind machine's axle more robust for use by experiment participants (Figure 5.10).

Finally, as the frame lacked additional bracing towards the top of its two triangles, the wind machine was initially prone to longitudinal movement along its axle (rocking away from and towards the crank handle) during rotation. This made continuous rotation more tiring, as the hand and arm had to push against this unwanted movement in performance. To resolve this, the frame was tightened longitudinally and laterally with an adjustable ratchet strap (Figure 5.11). This could be removed quickly if the frame needed to be disassembled. Some further maintenance was also required to ensure playability and eliminate some creaking that occasionally developed, including an application of WD40 to grease both the pivot point of the axle against the frame and the movement of the wooden part of the crank handle against its shaft. In addition, the bolts of the frame needed to be kept tight as they did tend to "work loose" as Napier advised (1962, p.52).

Remaking the acoustic wind machine from historical design instructions has revealed more about the embodied knowledge of theatre practitioners in the late nineteenth and early twentieth century. While all of the sources described in Section 5.1 highlight their particular approach to constructing the wooden cylinder and facilitating its motion against the cloth, there is scant detail included on the issues that might arise in the construction of the mechanism that will negatively impact the wind machine's performance potential.

The construction of a wind machine in the late nineteenth and early twentieth century

Figure 5.9: A slat after chamfering and sanding. The countersunk screw can also be seen fixing the slat to the plywood disc.

Figure 5.10: (L-R) The first iteration of the wind machine's crank handle showing the development of the twist in the axle, working to fabricate the new meat grinder handle and axle, and fitting the new crank handle and axle.

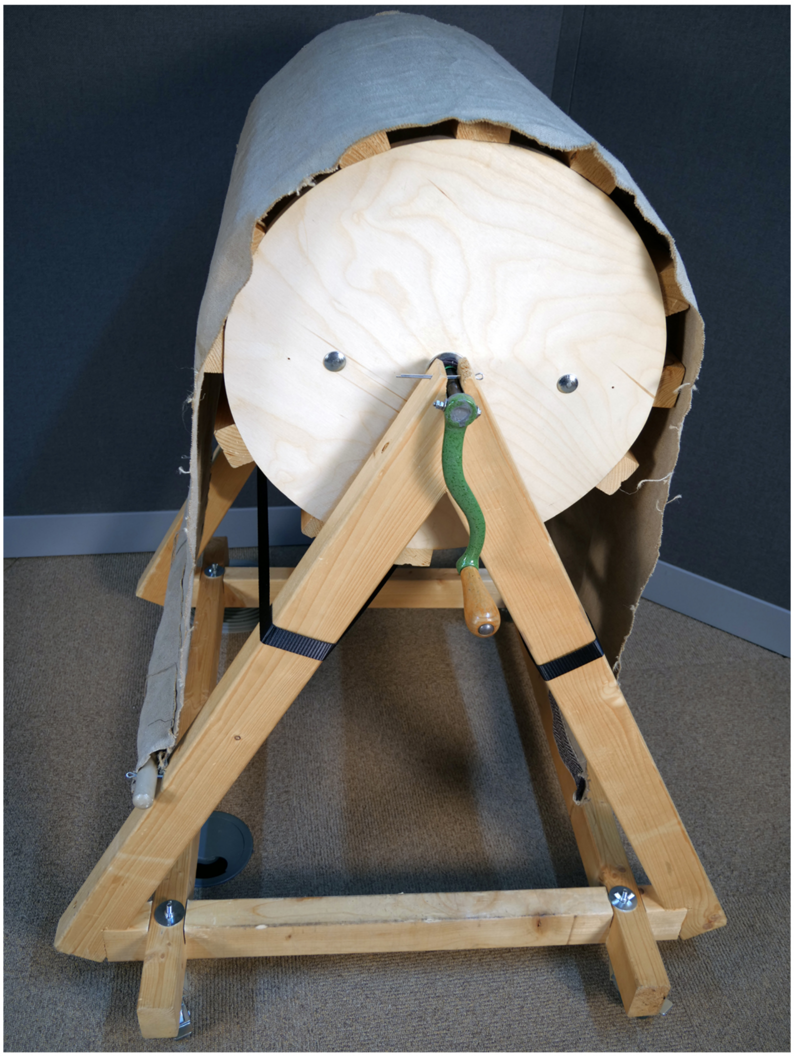

Figure 5.11: The completed acoustic wind machine with ratchet strap and wheels.

required the embodied knowledge of a carpenter-machinist. The success of the design depended on an understanding of how to fix its components tightly to each other, and how to ensure that the rotational movement was facilitated efficiently both for the benefit of the performer and for the machine itself. This was a demanding process. The fitting and adjustment of the component parts, for example the tight fixing required between the cylinder and axle, would have been undertaken using only hand tools. Any unwanted motion, for example longitudinal rocking or torsion in the axle as it twists in response to the rotation, would have overly tired the performer or shortened the life of the wind machine as it buckled under the stress of many performances. Even a slight change to the weight of the cloth has a significant impact on the ease of rotation.

This section has evidenced the precision and attention to detail required in the construction of this sound effect, and revealed some of the language of sound machinery that must have been employed to invent and refine these acoustic interfaces in the late nineteenth and early twentieth century. The next section considers the perceptual experience of a continuous sonic interaction with the remade acoustic wind machine in more detail.

5.3 Wind-In-Hand: Continuous Sonic Interaction with an Acoustic Wind Machine

With a working example of an acoustic wind machine available to explore, the continuous sonic interaction it affords a performer can be examined in more depth. This initial exploration is informed by the author's own embodied experience of soundmaking with objects. As outlined in Chapter 4, more empirical methods of acoustical and experimental evaluation are also used within this research to examine and validate the work of the design-led methodology.

This section describes the perceptual experience of performing with the acoustic wind machine in a continuous sonic interaction, and in particular how its handle and mechanism make very simple gesture feel complex and expressive to the performer. A process of playing and documenting to notice and unpack reveals yet more about this design to expand further upon the historical sources introduced in Section 5.1. This work also helps to inform the design of the tasks of the experiment that will be presented in Chapter 7, which evaluates both this acoustic wind machine and its digital model in a study with participants.

To focus the discussion in this section, we will first return to the simple conceptual model of a theatre sound effect first introduced in Chapter 2. Here, a handle and mechanism extend the performer's hand into complex acoustic soundmaking (Figure 5.12). Using the acoustic wind machine as an example, the perceptual experience of a handle and mechanism can be explored further, and the potential dynamics of such a system can be explored in more depth.

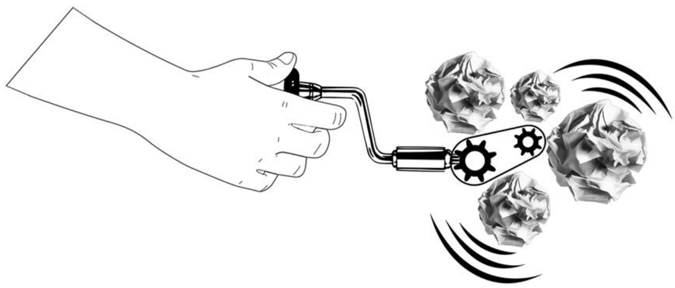

Figure 5.12: The simple conceptual model of a theatre sound effect introduced in Chapter 2.

The acoustic wind machine translates a composed gesture of embracing pressure (*hold handle*) and rotation (*turn handle*) into the continuous friction of wooden slats against a rough cloth. The crank handle offers one degree of freedom (DOF) to the performer, in that it is constrained to rotate in one axis around the wind machine's central axle. Furthermore, it can only be rotated in a clockwise direction owing to the way that the cloth has been fixed to the wind machine's frame on one side.

Rotating the crank handle anticlockwise will eventually pull the cloth off the wooden cylinder. Turning the crank handle in a clockwise direction immediately activates the acoustic wind machine to produce sound. Although the size of the machine gives the impression of it being very heavy to move, the wooden cylinder turns more easily on its axle than perhaps initially expected when the crank handle is rotated. The performer can face the machine to hold and rotate the handle in a circular motion in front of them,[8] or they can turn to one side and push the handle away from and then towards themselves to turn it.

[8]This is the position favoured by the author. The clockwise rotation favours right-handedness, an aspect of the acoustic wind machine's design that is somewhat adjusted for the experiment with participants reported in Chapter 7.

The machine enforces a somewhat static standing position on the performer during play. It is challenging to adjust the angle of the body while still maintaining the rotational movement. The performer has a choice in how they initially grasp the crank handle depending on the bodily position they have chosen. When facing the crank handle, it can be held in a neutral (palm to the side, like a handshake) or supine (palm facing upwards) grasp. The wooden part of the crank is smooth to the touch, satisfyingly tactile, and shaped to fit the hand well. It also spins freely on its own metal shaft. This means that as the crank is rotated, the wooden part of the handle also turns on its own shaft, changing the position of the hand and engaging the wrist and the whole arm in the movement (Figure 5.13). The freely moving wooden handle also responds with a very slight side-to-side motion on its metal shaft as it is turned. This gives the feeling of movement under the hand.

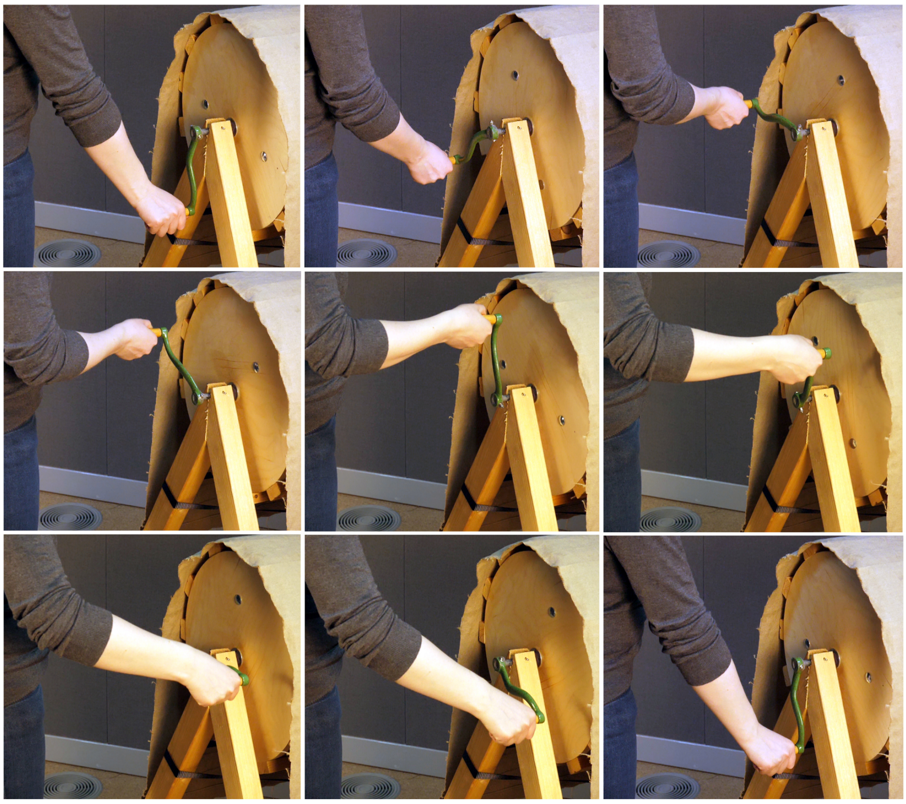

Figure 5.13: Video stills showing the author's performance of one rotation with the acoustic wind machine.

The rotational gesture required to activate the wind machine's sound is far from smooth, and is in fact comprised of two distinct parts. The first half of the rotation (from bottom to top) requires more effort on the part of the performer, as the movement pushes upward against the side of the cloth that is tightly coupled to the frame. The second half of the rotation (from top to bottom) requires less effort as the performer moves against the freely hanging loose side of the cloth, which resists the downward motion with a smaller frictional force. Transitioning between more and less resistance to the movement of the crank handle creates a feeling of more and then less weight being lifted by the wrist and arm as it adjusts position back and forth during a continuous movement. This shifting sensation of perceived weight and effort couples to the slight side-to-side motion of the wooden handle, connecting the movement of the arm, wrist and hand directly to the ongoing motion of the crank handle and cylinder.

While the perceived weight and effort required to rotate the acoustic wind machine's crank handle continually shifts in the process of a single rotation, it also changes throughout the course of a performance. Some clear effort is required to first set the cylinder in motion, but once it is continually moving there is a perceptible reduction in the feeling of *weight* through the crank handle held in the hand, and so the wind machine begins to feel perceptibly lighter and easier to rotate. This quality is particularly evident at higher rotational speeds, when the axle can spin so fast that the performer's interaction with the crank handle takes on the role of slowing rather than driving the progress of the wind machine. The simple cylinder and axle design is therefore storing the rotational energy input to it via the crank handle.

This means that the mechanism of the wind machine is working much like a flywheel, the dynamics of which warrant further examination. A flywheel is a mechanical device used to store rotational energy and therefore smooth changes in rotational speed, the principle of which can be seen in early human-powered potter's wheels or steam engines, as well as present-day automobile engines and toy fidget spinners. Setting a flywheel in motion takes a lot of force, but once it is moving it stays in motion and requires considerable force to stop. The mass of a flywheel will influence how much of this rotational energy it can store; larger, heavier examples can store more. The energy stored within a flywheel can be calculated with the equation:

$$E = \frac{1}{2}(I\omega)^2$$

where I is the rotational inertia of the flywheel and ω is its angular velocity.

By removing the cloth from the wind machine, its flywheel properties are accentuated, and they can be observed more clearly. Without the friction of the cloth to resist its movement, the wooden cylinder can be set in motion by rotating the crank handle at an accelerating speed to wind it up, and then letting the handle go. The cylinder will then continue to move on its own, only eventually slowing and coming to a stop due to the friction between the steel axle and wooden frame at each pivot point.[9] This means that the cylinder has a considerable angular momentum (motion in a direction on its axis) relative to this very small friction force. The mass of the wooden slats and their distance from the central axle gives the cylinder a high rotational inertia, which is the equivalent of mass for a spinning object. This rotational inertia combined with a high speed enables the cylinder to store a lot of rotational energy. Like a potter's wheel or toy fidget spinner, the bare wooden cylinder can be kept continually in motion with an occasional twist of the crank handle or slats to help it overcome the friction at the axle points.

Replacing the cloth reduces this flywheel effect, and the performer must use much more energy to get the wooden cylinder moving. The cloth is again acting as a brake to the moving flywheel of the cylinder and axle. However, as more and more energy is input into the machine and the speed of rotation is increased, more and more rotational energy is stored. This helps the cylinder to resist the friction of the cloth. At particularly high speeds, the cloth can be seen lifting very slightly from the slats as the cylinder turns, reducing the friction between them even further. This is due to some moving air produced by the cylinder in rotation (Figure 5.14).

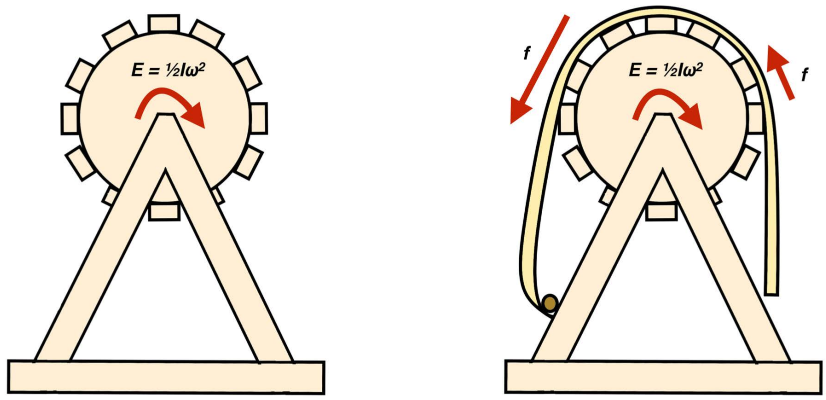

Figure 5.14: (L-R) The acoustic wind machine without its cloth, free to accumulate rotational energy, and the direction of the frictional force exerted by the cloth on the cylinder's motion which shape the performance gesture into two perceivable parts.

[9]As described in the previous section, these points were greased with WD40 to ease the cylinder's movement, which will of course reduce this friction.

The experience of rotating the crank handle is therefore far from simple. The wind machine is a complex dynamic performance mechanism that continually accumulates and responds to the energy it receives. It requires more effort to play at slow and regular speeds. Adding extra energy to the rotational gesture causes the feeling of weight in the hand and arm to perceptibly *lift*. As the movement of the wind machine's cylinder accelerates, it feels like it is moving ahead of the hand's rotational motion rather than being directly driven by it. At points the performer must push the wind machine forward, and then at other points resist that motion to slow it down. This dynamic system of shifting movement and effort produced by the orientation of the hand, wrist and arm in interaction with the wooden handle, the varying tension of the cloth and the continual resistance of the wooden cylinder to acceleration and deceleration gives a sensation of continually stirring a very viscous fluid in an almost figure-8 motion. This stirring motion produces a wind-like sound.

The sound is tightly coupled to the performer's action, and begins when the rotation of the crank handle begins. The sound source is positioned close to the performer and propagates from them, rather than towards them from a loudspeaker. Although the performer is very close to the wind machine when playing it (rather than being seated in a theatre audience), its sound is experienced as wind *effect* rather than a complex scraping sound produced by the slats and cloth. Playing the wind machine at a very regular speed of rotation can make the resulting sound more regular and *machine-like*, and therefore less convincing as a wind effect.

To perform the wind-like sound, the key is to vary the speed of the crank handle's rotation. Continually increasing and then disrupting the speed of the wind machine subtly varies the quality of the wind effect and makes it more perceptually natural. The resistance of the wind machine to starting and stopping its motion produces a particular envelope in the wind-like sound that contributes to the realism of the effect. It is this aspect of the design that affords the performer the potential to develop skill, and even expression, in playing the sound. Learning to subtly shift the speed of the crank handle's rotation is crucial to the success of the resulting wind effect. At higher and increasing rotational speeds, a characteristic whistling can be heard. During performance, the wind-like sound binds to the complex sensation of movement produced by the wind machine's mechanism, giving the impression of directly *stirring* or *tracing* a wind sound of a particular trajectory or *shape*. Sound seems to vibrate through the handle and the resulting wind effect feels like it is in-hand during performance. As the crank handle is turned, the wind machine also provides visual feedback to the performer as the slats rotate and the cloth can be seen moving in response. It is clear that historical practitioners had the opportunity to develop an embodied knowledge of the performance of wind, rather than a complex scraping sound, with this particular design.

5.3.1 Imitating Recordings of Natural Wind

As outlined previously in Chapter 4, perceptual and objective comparisons between the sound of the acoustic wind machine, its digital counterpart and recordings of real-world wind sounds form a critical part of the design and evaluation approach of this research. This work began with an exploration of the perceptual connection between the acoustic wind-like sound and natural real-world wind, with field recordings of wind from the BBC Sound Effects Library (BBC 1988) used as stimuli to guide the author's performances. The technical setup for this work will be briefly restated here to aid the understanding of the reader. The field recordings were first delivered through stereo loudspeakers at sufficient level to facilitate the author's rehearsal with the acoustic wind machine. The acoustic wind machine was first played acoustically without any sound reinforcement, and later a microphone captured its sound and delivered this to the same loudspeakers mixed with the field recording of natural real-world wind. Then, to allow the performances of the acoustic wind machine to be recorded simultaneously along with the recordings of natural wind, an AKG microphone and an RME Fireface 400 interface were then used to capture the acoustic sound using Pro Tools and deliver it to the performer in headphones. The Pro Tools session also delivered the natural wind sounds to the performer through the same headphones. The headphone channels were split, with the natural wind sound being delivered in one ear and the acoustic wind-like sound delivered into the other ear.

Performing the acoustic wind machine to imitate the various recordings of natural wind produced some interesting results. When rehearsing with the loudspeakers, it becomes challenging to differentiate clearly between the two sources. Both the acoustic and natural winds tend to complement each other and sometimes seem to blend together as one unified source of a wind-like sound. The performance limitations of the acoustic wind machine become more apparent when it is played to specifically imitate a recording of natural wind. For a slow and continuous steady breeze, the crank handle must be rotated very slowly to avoid producing the more usual gentle gusts in the amplitude envelope of the acoustic wind sound. More high pitched and whistling natural wind recordings are best imitated with very irregular gestures of rotation. By deliberately shaking or wobbling the crank handle as it is rotated, the cylinder will shudder in response. This buffets the cloth against the slats and produces a series of audible short scrapes, which can result in more small and frantic gusts in the acoustic wind-like sound.

When performing to imitate these more frenetic natural winds, the crank handle can feel too slow and cumbersome to rotate fast enough. When performing using headphones and tracking the acoustic wind machine simultaneously with a particular natural wind recording, turning the crank handle begins to feel less like a constrained rotation in one direction, and instead more like drawing or tracing out a perceivable *shape* of the natural wind sound in the air using the handle. There is a sensation of alternately flinging the crank handle forward and pulling it backwards in order to create more shifts in the acoustic wind machine's envelope. With practice, it is possible to track the amplitude envelope of the acoustic wind machine's sound quite closely with that of the natural wind sound (Figure 5.15).

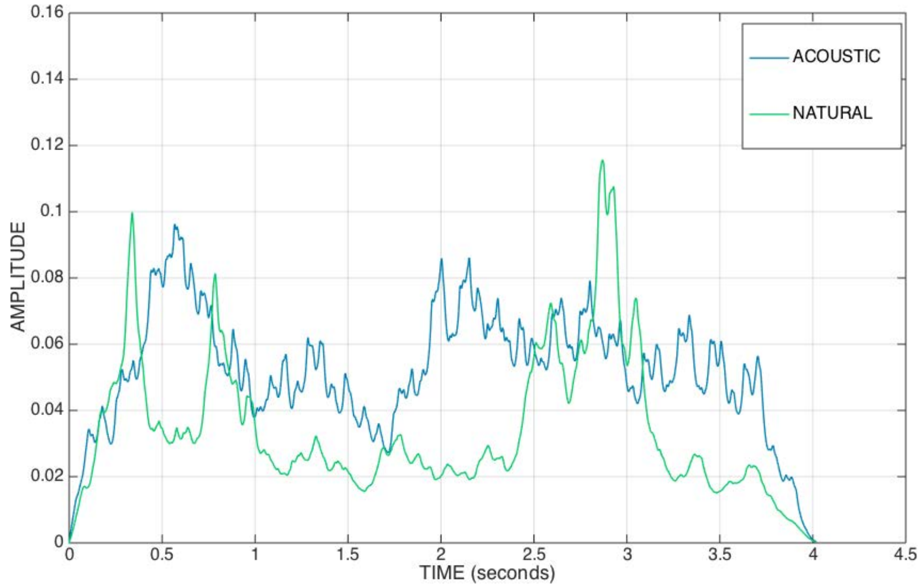

Figure 5.15: A comparison of the amplitude envelopes of a particularly frenetic and high-pitched whistling natural wind sound (green) and the acoustic wind machine (blue) imitating that sound.

While the acoustic wind machine produces a sonically variable and convincing wind effect in performance, directly imitating a natural wind sound requires a more focused approach to how its wind-like sound is activated and modulated. This highlights the inherent limitations of the crank and axle mechanism, which at times cannot be moved slowly enough and at other times is too cumbersome to be moved at speed in order to easily produce an accurate imitation. However, performing the acoustic wind machine to imitate specific recordings of natural wind sounds has revealed a further dimension to the potential of the gesture of rotation that activates its wind-like sound. The continuous resistance of the flywheel to changes in rotational speed requires increased effort from the performer to articulate just the right kind of sound. The crank handle becomes more compelling, as it traces out a perceptual shape or gesture of the wind being imitated.

We have so far considered the rich multimodal experience the acoustic wind machine gives a performer when continuously interacting with it to produce sound, and examined some of the dynamics of the mechanism that create this experience for the performer. The acoustic wind machine's progress in performance was also compared to a recording of a natural wind sound. Next, the soundmaking components behind the production of the wind-like sound will be considered in more detail.

5.4 Soundmaking Components

Examining a continuous sonic interaction with the acoustic wind machine has revealed more about the complexity of the perceptual experience it affords in performance. The process behind its method of sound production is also more complex than it initially appears. This section begins to examine how exactly the acoustic wind machine produces its wind-like sound, and draws the information together to produce an entity-action model (Farnell 2010, p.36) describing its workings. This work informs the programming approach taken to create the prototype digital model of this acoustic wind machine that will be presented in Chapter 6. In turn, programming the prototype digital wind machine will facilitate a deeper understanding of how the acoustic wind machine produces sound.

As discussed in Section 5.3, the wind machine produces a wind-like sound when the crank handle is rotated. While at first glance this sound appears to be produced by the friction between two objects - the wooden cylinder and cloth - the sound is in fact produced by a number of individual slats in contact with the cloth at once. In the case of this particular wind machine, only seven of the twelve slats touch the cloth at any one time. This can be seen clearly when the wind machine is at rest (Figure 5.16).

As the crank handle is turned to rotate the cylinder in performance, each slat moves to come into contact with the cloth, rub against it and then move out of range. Therefore,

Figure 5.16: The acoustic wind machine at rest. The arrows mark each of the seven slats in contact with the cloth.

rather than one source of the sound (i.e. one rubbing cylinder) there are in fact twelve rubbers continuously interacting with the cloth, each with their own circular trajectory. The circular path of the slats can be clearly seen by marking them with electrical tape as a reference, and then filming a continuous rotation (Figure 5.17).

This means that this acoustic wind machine extends the performer's single rotational movement with the crank handle to control twelve individual sources of friction. The brief mentions of the number of slats to consider using, or the particular space to put between them, in the historical sources outlined in Section 5.1 takes on a new significance. Clearly the number of slats and their particular placement is also critical to the kind of wind effect the machine is capable of producing. This is in fact a one-to-many mapping approach (Hunt et al. 2003) implemented in the design of an acoustic interface that can be applied directly in the digital domain. The prototype digital wind machine will therefore be based on twelve instances of a physical model of friction, each activated according to the placement of the slats on the acoustic wind machine.

The cloth of the acoustic wind machine is its main resonator. It vibrates in response to the friction of the slats against it, and must help to blend the individual sources of friction together to create a complex wind sound. Because of the tight coupling of the slats to the plywood discs, the fixing of the cloth to the frame and the connection between the axle and the frame at the pivot points, the sound produced through friction will also transfer through the wooden structure and even through the handle to the performer. Given that the cylinder and frame are made of softwoods and are not particularly resonant, much of this sympathetic vibration will be inaudible. The cloth is quite a complex resonator. As discussed in Section 5.3, it is tight on one side of the cylinder and loose on the other, which will give it slightly different resonant properties on each side.

In addition, the cloth responds to the continuous motion of the cylinders slats with a sympathetic and irregular ripple, and seems to both tighten and loosen at points during rotation (Figure 5.18). It can also pull taut and lift slightly from the slats when the cylinder is moving at very high speeds. Although this kind of cloth behaviour has not been physically modelled for the purposes of soundmaking thus far, there are useful methods from physical modelling synthesis for musical instruments that are available for use as part of the digital prototype produced in this research. These will be discussed in more detail in Chapter 6.

With the main soundmaking components of the acoustic wind machine described in

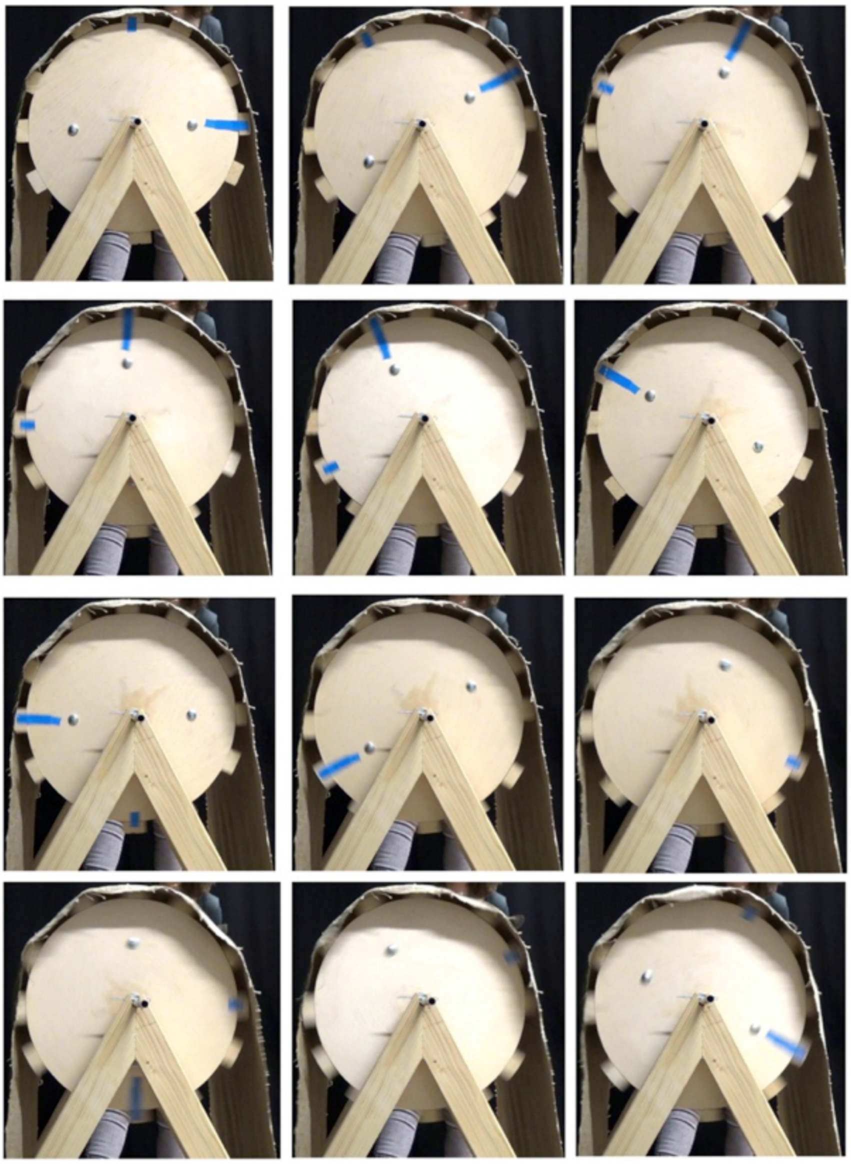

Figure 5.17: Video stills showing the slat movement through one rotation of the wind machine's handle. Rotation is anticlockwise when viewed from this side of the machine.

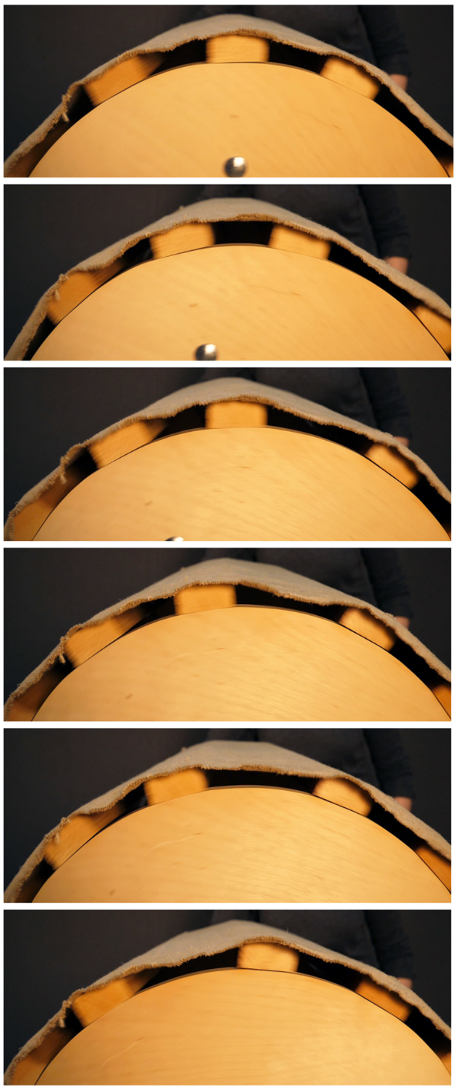

Figure 5.18: Video stills showing the shifting motion of the cloth against the wooden slats during rotation.

more detail, an entity-action model (Farnell 2010, p.36) can now be created to describe them (Figure 5.19). The two main features of the model, the individual slats and cloth resonator, will form the core of the prototype digital wind engine described in Chapter 6.

Some additional features of the acoustic wind machine's sound need to be considered to inform the prototype digital wind engine. These will be used to refine the mapping of data produced by a rotary encoder coupled to the acoustic wind machine to perform the digital wind sound. As will be explored in Chapter 6, these features are important in replicating the mechanical behaviour of the acoustic wind machine in terms of how it produces sound. From the exploration of the continuous sonic interaction with this sound effect described in Section 5.3, it is evident that:

- There is a link between the increased friction produced in the first half of the rotational gesture due to the tight side of the cloth and the loudness of the wind-like sound.

- The wind-like sound can seem repetitive and *machine-like* if the performer does not vary the speed of the crank handle during play.

- There is a link between the speed of the crank handle's rotation and the pitch of the characteristic whistling sound. As the rotation increases in speed, so does the pitch. As the rotation slows, the pitch decreases.

This section has examined the unique way in which this acoustic wind machine produces its wind-like sound in performance, and highlighted some particular aspects that will inform the programming approach detailed in Chapter 6. The next section summarises the discussion in this chapter.

5.5 Summary

This chapter has described the first stage of the design-led methodology outlined in Chapter 4, which investigates the historical acoustic wind machine design as an interactive sounding object. Four distinct wind machine designs from the late nineteenth and early twentieth century were chosen from the historical sources available for this research, and explored for their particular approach to realising this sound effect. This work informed a new design for a working example of a wind machine. This design was constructed and

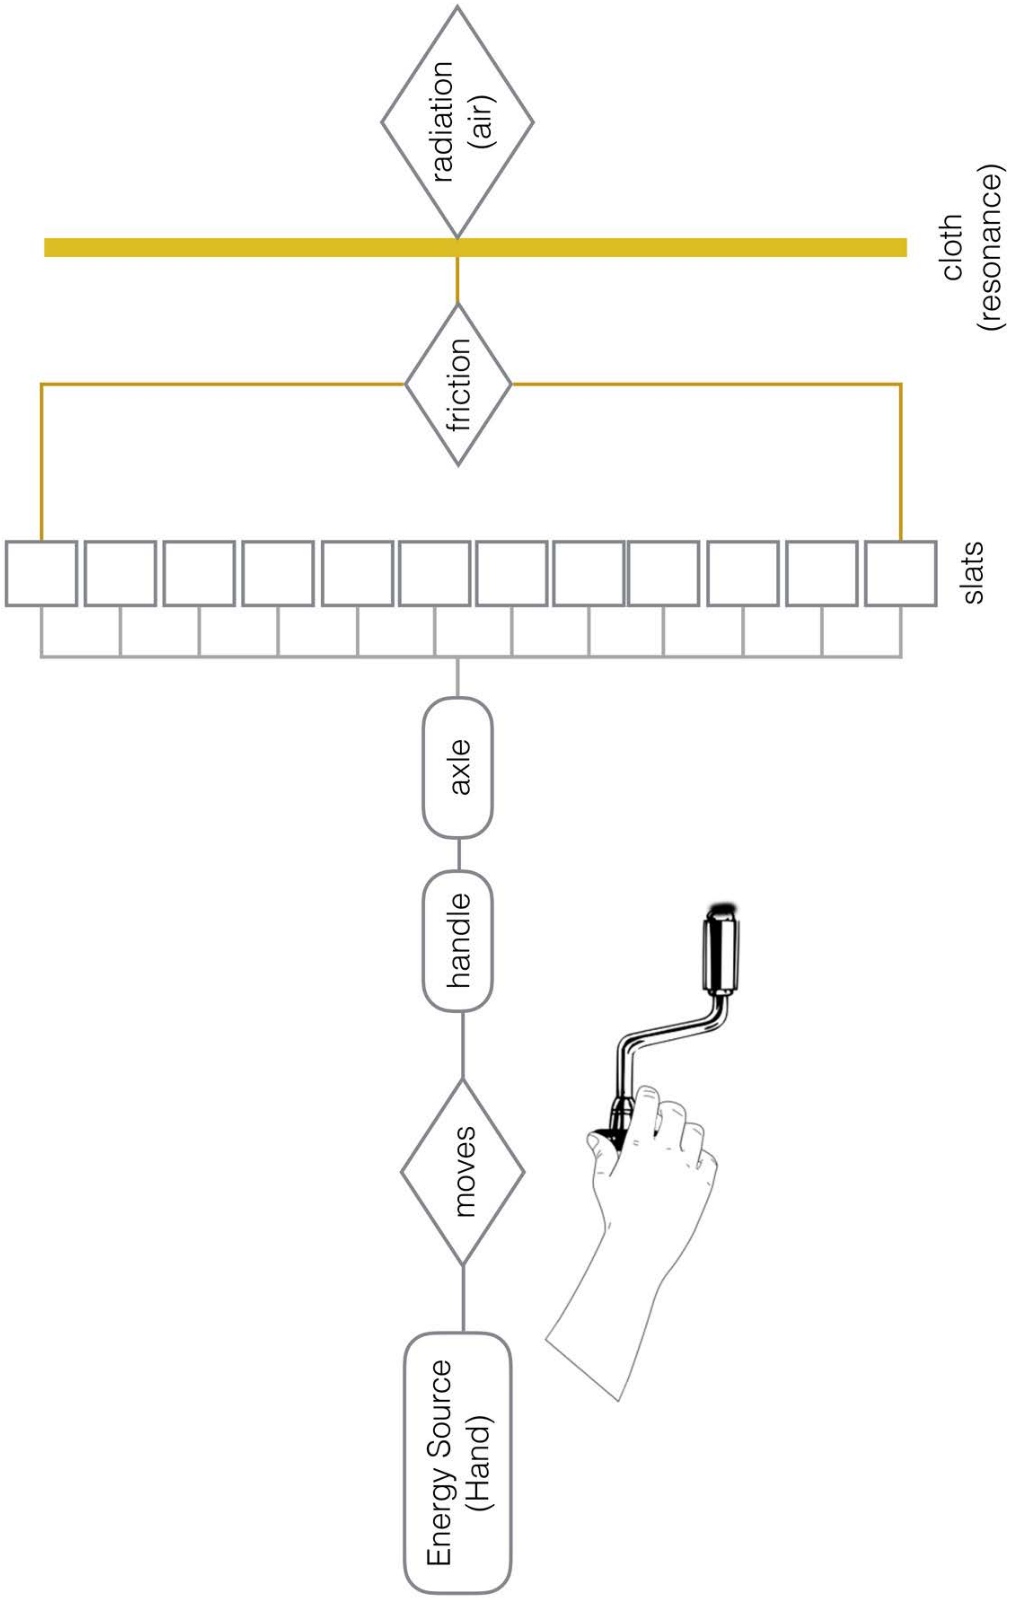

Figure 5.19: An entity-action model describing the workings of the acoustic wind machine.

refined to produce a working acoustic wind machine robust enough to be subjected to a full investigation and an experiment with participants.

The acoustic wind machine was then interacted with and documented in order to unpack the perceptual experience of its use in performance. This process also allowed the acoustic wind machine to be described in more detail in terms of its sound-producing components, informing the creation of an entity-action model of its workings and highlighting some particular features of the wind-like sound it produces.

Chapter 6 discusses the installation of a rotary encoder to capture the movement of this acoustic wind machine as data, and the programming of a prototype digital wind engine to model its wind-like sound.

Chapter 6

Protoyping a Digital Wind Machine

This chapter outlines the second stage of the design-led methodology of this research, which centres on the programming of a digital synthesis engine in Max/MSP to model the workings of the acoustic wind machine presented in Chapter 5. As outlined previously in Chapter 4, this work is undertaken in a creative and exploratory way, producing a practically synthesised *prototype* in order to generate further reflection (Stappers 2013). Technical aspects of both the programming approach and the objective analysis of both wind-like sounds are therefore given space to provoke further examination of the digital potential of historical theatre sound effect designs. In the spirit of this design-led method, several iterations of the sensor configuration, audio programming, acoustical evaluation and interface design work are reported within this chapter to highlight the *wicked* nature of the problems that emerged throughout and draw the reader's attention to the potential implications for the *sounds-in-hand* framework introduced in Chapter 4. This discussion will be developed further in Chapter 8.

The work described here begins with the aim of using the acoustic wind machine itself as the interface to activate and modulate the digital wind-like sound in order to preserve the complex experience of performance with the cylinder and axle mechanism across the two different wind effects. The acoustic and digital wind-like sounds are also acoustically evaluated to compare the responses of both systems across three distinct performance gestures in order to validate the prototype digital wind machine before proceeding to the experimental evaluation with participants reported in Chapter 7. A fully digital crank interface is also created to activate the digital wind-like sound in performance in order to facilitate an exploration of the perceptual difference between the experience it affords and that of the acoustic wind machine it enactively recreates.

As previously outlined in Chapter 4, the programming work to create a digital model of the acoustic wind machine's sound centres on the graphical Max/MSP environment and the Sound Design Toolkit's (SDT) (Baldan et al. 2017) physical model of friction. Figure 6.1 provides an overview of the main components of the prototype digital wind machine and introduces the stages of work reported within this chapter. A guide to the full Max/MSP program can also be found in Appendix B.

This chapter proceeds from the working acoustic wind machine presented in Chapter 5, and examines the work undertaken to capture its motion, deconstruct its soundmaking properties and extend them into the digital domain. As such, Section 6.1 begins with a discussion of the work undertaken to couple a rotary encoder to the acoustic wind machine and translate its rotational movement into data. Section 6.2 then outlines the creation of the main soundmaking components of the prototype digital wind machine by taking a practical synthesis approach to the use of the Sound Design Toolkit (SDT) in Max/MSP (Baldan et al. 2017). Section 6.3 explores the mapping that connects the rotational data from the rotary encoder to the parameters of the digital synthesis engine, allowing the prototype digital wind machine to be activated and modulated in performance. Section 6.4 presents an acoustical evaluation and comparison of the acoustic and prototype digital wind machines performing the same rotational gestures in order to validate the program in Max/MSP before proceeding to the experiment with participants presented in Chapter 7. Section 6.5 outlines a parallel piece of work undertaken to create a fully digital crank interface for the program in Max/MSP, and examines the perceptual experience of using it to perform the digital wind-like sound. This is another iteration of the framework for *sounds-in-hand* presented in Chapter 4. Section 6.6 presents a summary of this chapter.

6.1 Digital Action

As we explored in Chapter 5, the very simple rotational gesture afforded by the acoustic wind machine's crank handle gives a highly complex sensation of tactile feedback, continuous movement and energy exchange to a performer continuously interacting with its wind-like sound. To preserve this experience in the activation and modulation of the prototype digital wind machine, the acoustic wind machine itself was chosen as its main performance interface. This section outlines work undertaken to capture the rotational

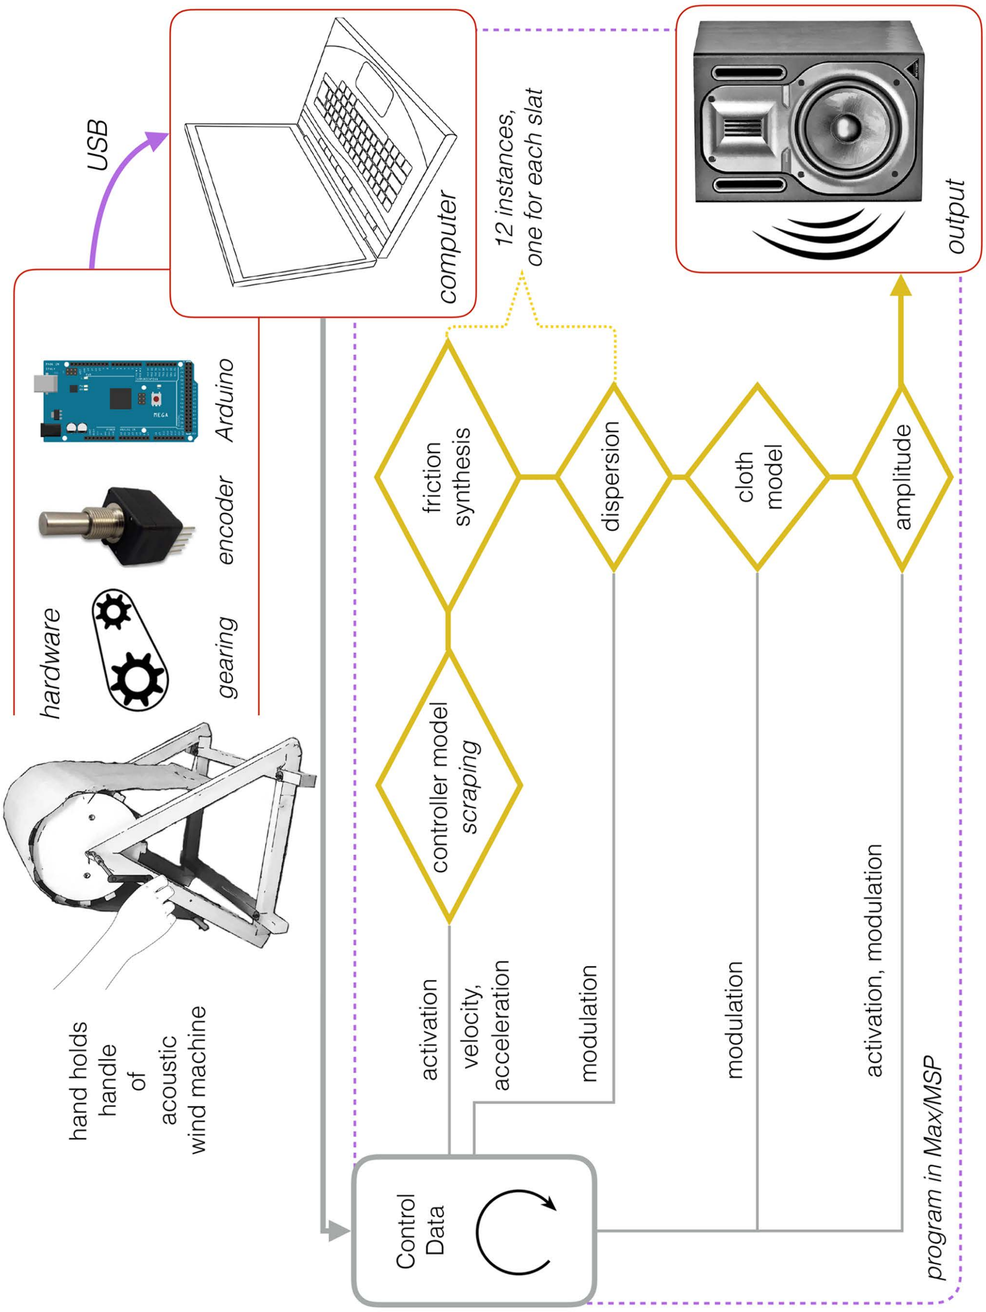

Figure 6.1: The main components of the prototype digital wind machine presented in this chapter.

motion of the acoustic wind machine as data.

6.1.1 Intertial Measurement Unit (IMU) Sensor

Given the dynamic response of the acoustic wind machine's flywheel mechanism to the movement of the crank handle, the first iteration of this work explored the potential of an Adafruit 10-DOF Inertial Measurement Unit (IMU) sensor (Townsend 2013). This sensor tracks motion in 3D space using a combination of an accelerometer (acceleration), gyroscope (speed), magnetometer (magnetic field) and barometer (atmospheric pressure). Sinyor and Wanderley's DMI Gyrotyre (2005) discussed in Chapter 2, which used the dynamic properties of a spinning bicycle wheel to give a sensation of performance effort and vibrotactile feedback, also incorporated an accelerometer and gyroscope into its design. The IMU was mounted onto the acoustic wind machine to track its movement, and the resulting data read with an Arduino prototyping board.[1] The data was transmitted wirelessly to the computer using XBee hardware (Sparkfun 2015) so as not to impede the normal working of the acoustic wind machine with unnecessary cable connections (Figure 6.2).

The IMU sensor captured much of the complexity of the acoustic wind machine's motion, but its sensitivity produced a rich data stream that needed to be smoothed and simplified by both the Arduino and the Max/MSP program in order to activate the prototype digital wind machine efficiently. An initial analysis and comparison between the acoustic wind machine and the first iteration of the prototype digital wind machine driven by this sensor configuration found that this choice of sensor hardware was potentially causing an unnecessary delay in the progress of the digital wind-like sound during performance (Keenan & Pauletto 2016). In particular, wireless transmission of performance data using the XBee has previously been found to introduce a minimum latency of 57ms (McPherson et al. 2016, p.5). This delay would have an impact on the results of a perceptual experiment to compare the two wind machines. As such, the next stage of work proceeded with the aim of simplifying and increasing the efficiency of the way the acoustic wind machine's motion was captured as a stream of data to transmit to the computer.

[1] Arduino is an open-source platform for working with hardware inputs (sensors, switches) and outputs (motors, LEDs) to create interactive electronics: https://www.arduino.cc/en/Guide/Introduction

6.1.2 Rotary Encoder

The wooden cylinder of the acoustic wind machine moves continuously in rotation, and its plywood discs are mounted close to the A-frame, which remains static. The rear side of the wind machine, having no crank handle attached to the axle, provides an ideal mounting point for a mechanical coupling between the motion of the cylinder and the static frame to capture the movement as data with an incremental rotary encoder,[2] which can rotate continuously. Some prototype 3D printed gearing was intitally produced for this purpose, in addition to a mounting bracket. Following testing and calibration, the 3D printed gearing was replaced with a more robust version laser cut from acrylic.

The first gear was threaded onto the acoustic wind machine's axle, and fixed with some glue to the centre of the cylinder's rear face. It was also tightened around the axle with a set screw.[3] The rotary encoder was then mounted to the acoustic wind machine's A-frame with the bracket, and the corresponding gear fitted to its shaft. The gears mesh together in a ratio of 1:1, meaning that each complete rotation of the wooden cylinder also

[2]The encoder used was Bourns type ENA1J-B28-L00128L.

[3]A set screw works much like the principle of the split pins used in the construction of the acoustic wind machine. By tightening the screw, the gear can be fixed onto the central shaft of the axle.

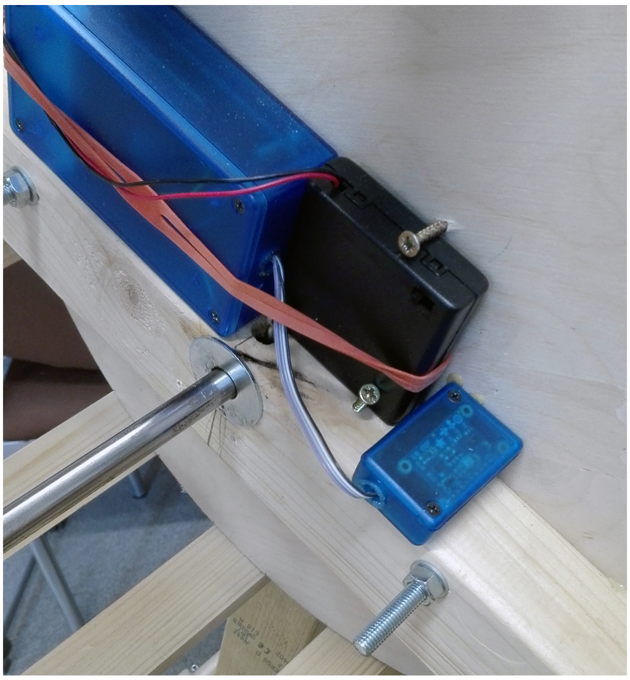

Figure 6.2: The initial configuration of (L-R) an Arduino and transmitter, battery pack and IMU sensor installed on the acoustic wind machine.

completes a full rotation of the rotary encoder. To ensure that the gears meshed together properly while moving, a spacer was added to the shaft of the rotary encoder and also to the acoustic wind machine's axle behind the crank handle.

As discussed in Chapter 5, the acoustic wind machine was initially prone to some longitudinal movement in rotation, which was resolved by tightening its frame with a ratchet strap. Due to the length of the acoustic wind machine's axle, the cylinder could also slide forwards and backwards on the frame if pushed with the crank handle. While the spacer on the axle resolved most of this movement, the ratchet strap tension had also to be carefully calibrated to keep the cylinder positioned within tolerances, as if it slid too far towards the rear side of the frame the rotary encoder's shaft could grind against the plywood (Figure 6.3).

Figure 6.3: (L-R) Fitting the first gearing prototype with the rotary encoder's spacer, adding a spacer between the cylinder and axle, and the final configuration of laser cut gearing, rotary encoder and Arduino on the acoustic wind machine, with grind marks evident on the plywood.

With the rotary encoder installed and moving in response to the acoustic wind machine's rotation, the Arduino could be configured to read its data. When connected to the Arduino to receive power, the rotary encoder outputs an electrical signal on two of its pins when it is turned. The two signals are square wave pulses running 90° out of phase with each other, known as quadrature outputs. By counting the pulses produced by the rotary encoder as it is moving, the Arduino can determine the encoder's angular position, and by monitoring both of the square wave signals together, it can determine the direction of the encoder's rotation (Figure 6.4). The Arduino was loaded with a script instructing it to decode the output of the rotary encoder and produce one incremental pulse for use within Max/MSP. As the Arduino was now mounted on the stationary A-frame, it could be connected directly to the computer with a USB cable to receive power and send its data into Max/MSP.

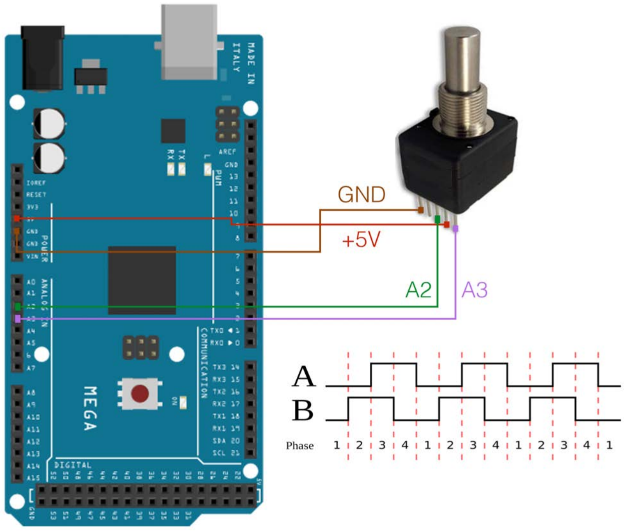

Figure 6.4: Wiring diagram showing the pin connections between the rotary encoder and Arduino prototyping board (left), and the encoder A/B input that it monitors to produce incremental degree values (right). The Arduino sends 5 volts of power to the encoder, and reads the A and B quadrature outputs as it turns.

6.1.3 Rotational Data to Max/MSP

A patch was then programmed in Max/MSP to read the encoder data from the Arduino. When connected to the acoustic wind machine's gearing, the resolution of the rotary encoder was 261 PPR, or pulses per rotation. As the encoder had no absolute value, it would output a continuously increasing value when rotated continuously with the crank handle. The data input to Max/MSP was therefore constrained and mapped to a repeating circular rotation between the values of 0° and 360°, expressed as floating-point numbers.[4] The rest position of the acoustic wind machine's crank handle was set as the 0° reference point, and a reset button added to the patch to return the rotary encoder to this value at the end of each performance.

Passing through the reference point in the course of a continuous rotation initially

[4]Converting the encoder's steps into numbers with several decimal places ensured a smoother transition between values.

caused unwanted spikes in the degree values, as the rotary encoder values could not be easily interpolated in the transition between 360.0° and 0.0°. The [*dot.unwrap*] abstraction from the Digital Orchestra Toolbox suite in Max/MSP (Malloch et al. 2007) was used to smooth out this transition and eliminate any sudden spike in values. With a smooth 360° rotation calculated from the acoustic wind machine's crank handle movement, a continuous dial was added to the patch in Max/MSP to give a visual display of its position. The patch was then programmed to calculate the number of rotations, rotations per minute (RPM), velocity and acceleration from this continuous stream of data (Figure 6.5).

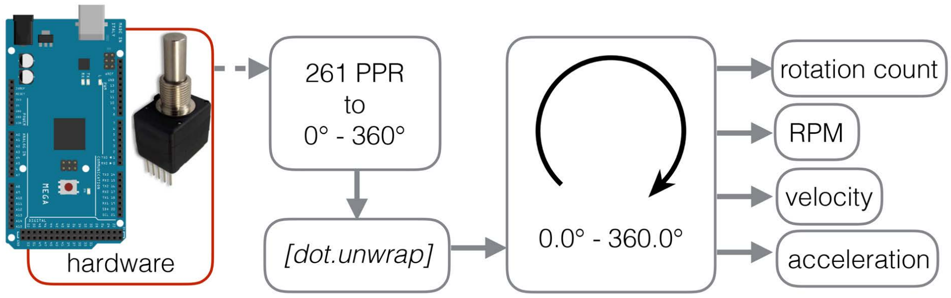

Figure 6.5: Diagram of the first part of the Max/MSP program, which captures the acoustic wind machine's movement as rotational data, smoothing it to a range of 0.0°-360.0°, and then calculates some additional data from this continuous rotation.

6.1.4 Slat Data Streams

As highlighted in Chapter 5, the acoustic wind machine's sound is produced not just by the motion of one moving object (the wooden cylinder) rubbing against the rough cloth, but by a combination of each individual slat's movement against the cloth during rotation. The prototype digital wind machine was therefore based on twelve distinct instances of the SDT physical model of friction.

To allow each of these friction models to be individually controlled, the rotary encoder's movement was parsed into a continuous data stream for each individual slat on the acoustic wind machine, allowing their progress to be individually tracked in real time. The first 360° rotational data stream into Max/MSP was designated as the trajectory of the crank handle, and also that of the bottom slat of the acoustic machine which tracks the same rotational path. The degree position of each slat on the acoustic wind machine was measured relative to this bottom slat in order to capture their irregular clock-face configuration around the plywood disc (Figure 6.6).

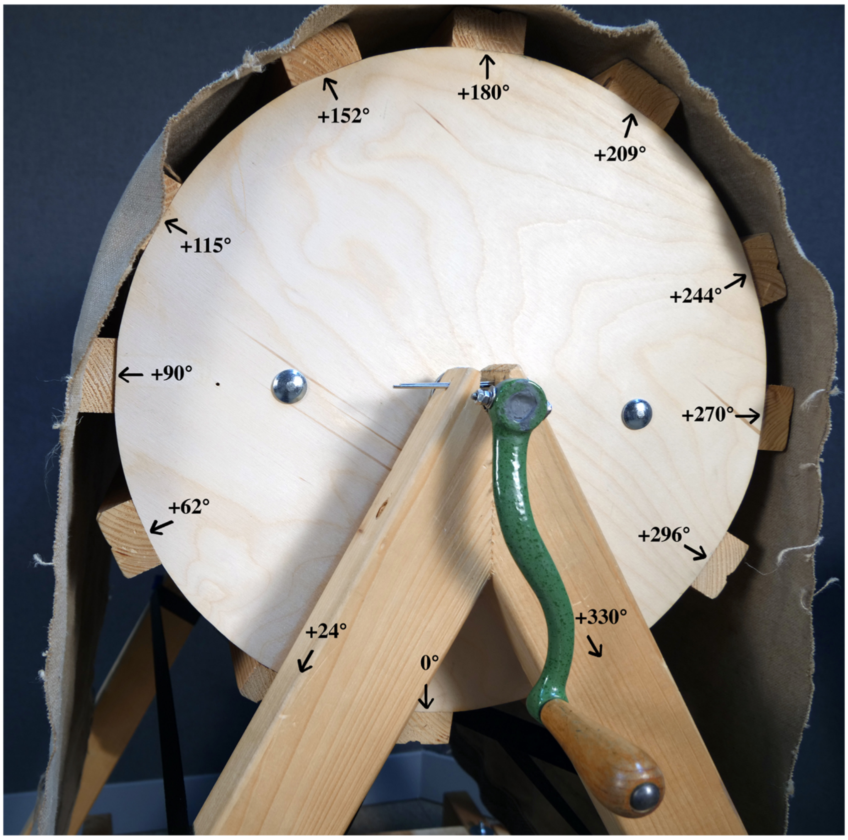

Figure 6.6: Side view of the acoustic wind machine with its slat positions marked in a 360° clockwise rotation from their position of origin. The degree values used to parse each slat position from the main data stream in Max/MSP and replicate those positions digitally are also shown.

This produced a set of degree values for use in Max/MSP, enabling the main 360° rotational data stream to be placed out of phase a further eleven times to produce an individual stream of rotational data for each of the twelve slats. Rotating the acoustic wind machine's crank handle could now drive a multi-pronged data stream that modelled the movement of individual digital slats ahead of or behind the position of the handle, mirroring how the wooden slats move on the acoustic wind machine. With the first part of the program in Max/MSP established, the next section considers work to model the main soundmaking components of the acoustic wind machine.

6.2 Digital Sound

Chapter 5 outlined work to investigate and deconstruct the main soundmaking components of the working acoustic wind machine. It was determined that the twelve individual slats each produce sound through their scraping action against the rough cloth as they are rotated, and the cloth itself may act as a resonator, continuously vibrating in response to amplify and blend the scraping sounds and give the wind machine its characteristic whistling. This section presents the programming work in Max/MSP to prototype these soundmaking components for real-time performance. The entity-action model (Farnell 2010, p.36) of the acoustic wind machine previously presented in Chapter 5 forms the basis for this work, which proceeded with the aim of practically synthesising a digital model to be activated in performance by data mapped from the rotary encoder described in Section 6.1. As such, the digital model developed over several stages of programming, each adding further complexity to the sonic potential of the prototype digital wind machine. The stages of this work were as follows:

1. Use the Sound Design Toolkit's (SDT) physical model of friction (Baldan et al. 2017) to create a digital model of the sound produced by one slat of the acoustic wind machine as it rubs against the cloth.

2. Combine several instances of this *slat model* to create the sound of twelve slats moving at different trajectories depending on their position on the acoustic wind machine's cylinder.

3. Create a model of the acoustic wind machine's cloth to blend the sound produced by a combination of twelve slat models being activated in performance, and offer a way to develop some of its characteristic whistling Aeolian tone digitally.

These stages of work will now be described more fully. The data mapping that activates and modulates the sound of these components in performance will then be discussed in detail in Section 6.3.

6.2.1 Slat Model

The interaction between each slat of the acoustic wind machine's cylinder and the encompassing cloth is modelled with its own instance of the Sound Design Toolkit's (SDT) physical model of friction (Baldan et al. 2017). Physical models of friction can produce a wide variety of sounds, and have previously been used as the basis for interactions with a virtual bowed string (Young & Serafin 2003), a digital model of Russolo's intonarumori (Serafin & Nordahl 2005), and creaking or squeaking sounds for animations (Avanzini et al. 2005). Friction models can be static, mathematically describing the behaviour of one object moving against another stationary object, or they can be dynamic, accounting for more subtle changes in sliding velocity and therefore friction force when motion occurs.

The SDT friction model is based on a dynamic elastoplastic friction algorithm (Dupont et al. 2002). The model interprets friction as being caused by multiple bristles making micro-contacts between surfaces, and it is the behaviour of these bristles that is mathematically modelled. In the case of the acoustic wind machine, it may be the interaction between the *bristles* of the wooden surface of each slat and those of the encompassing cloth that are responsible for much of the resulting wind-like sound. Figure 6.7 illustrates how this interaction might look. While the acoustic wind machine's cloth is described here as the stationary surface of the friction interaction, it does in fact move in response to the motion of the slats in performance. As will be explored in Section 6.3, the data mapping approach to activate the digital wind-like sound from the rotary encoder does attempt to account for the potential sonic results of this motion.

The SDT environment is implemented across several Max/MSP objects, allowing models to be configured to specific requirements. Each interaction algorithm is made up of two solid resonating objects connected together by an interactor, which in this case is friction. The interactor can be manipulated further in real time by imposing a controller model onto it, as will be explored further in Section 6.3. The interaction between each individual slat of the acoustic wind machine and the cloth was modelled with one instance of this configuration of SDT objects (Figure 6.8):

- *[sdt.inertial~]:* This is the first solid resonating object, a physical model of an inertial mass which accelerates according to Newton's second law of motion, where Force = mass x acceleration. As the mass of the moving object in the interaction, this was configured to represent the mass of the rotating slat.

- *[sdt.modal∼]:* This is the second solid resonating object, a modal resonator to physically model three distinct resonant frequencies. Modal resonance is produced with models of mechanical oscillation, which mathematically describes the motion of a spring weighed down with a mass that loses energy through damping (Cook 2002, p.40). As the second stationary object of the friction interaction, this could be said to represent the resonance of the acoustic wind machine's cloth, but as previously highlighted this is not a simple case of one moving and one stationary object. Parameters to *[sdt.modal∼]* were chosen to create as convincing a wind *effect* as possible in performance when compared with the acoustic sound, and some of the sound will be produced by the resonance of the wooden slats themselves, or the air moving between them as they rotate.

- *[sdt.friction∼]:* This is the interactor, an elastoplastic dynamic friction model that couples [*sdt.inertial∼*] and [*sdt.modal∼*] together and outputs sound. It models the dynamic friction between the cloth and the slat in rotation.

- *[sdt.scraping∼]:* This is the controller model, which simulates a probe sliding on a surface. It modulates the quality of the friction produced by [*sdt.friction∼*] to model the scraping motion of the slat against the cloth.

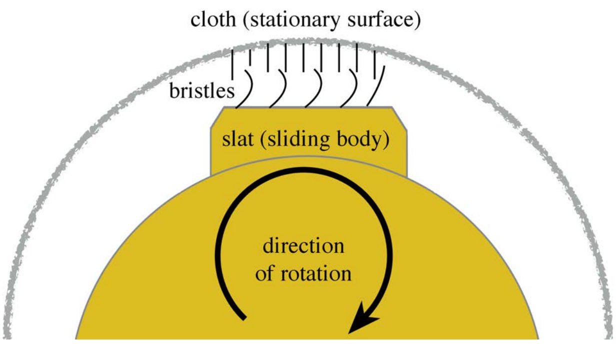

Figure 6.7: Diagram of the bristle interpretation of friction used within the SDT model and how it might be applied to the interaction between each wooden slat and the rough cloth (based on Fontana & Rocchesso (2003, p.152))

6.2.2 Digital Slat Configuration

Following the creation of a single slat model to produce the sound of the interaction between one slat and the encompassing cloth, twelve instances of this were programmed in Max/MSP. This approach was taken to account for the particular trajectory taken by each slat on the acoustic wind machine as it comes into range of the cloth, rubs and scrapes as it rotates, and then falls out of range again. This means that each slat is silent for a particular part of a full rotation, which must have some effect on the overall resulting sound of the acoustic wind machine. As the computer was required to run these simultaneously, each one was implemented within its own [*poly~*] object in Max/MSP. [*poly~*] is usually used for polyphonic voice allocation. It allows digital signal processing (DSP) to be temporarily muted or downsampled for any object encapsulated within it, allowing the overall program to work more efficiently. These functions were used as part of the mapping between the rotary encoder's data and the activation and modulation of the digital slats, helping the computer's CPU to work more efficiently when running the Max/MSP program, resulting in a much smoother operation of the SDT object configurations and an improved resulting sound. The mapping approach for the prototype digital wind machine will be explored in detail in Section 6.3.

Each of the SDT objects used to create the digital slat model offers a number of parameters that can be calibrated to modify the quality of the resulting sound produced at the output of [*sdt.friction~*]. The final parameters for the slat models were chosen through a process of rehearsal and testing by the author to approximate the sound of the acoustic wind machine. The slat models were activated through a performance with the crank handle, with monitoring of the program in Max/MSP over headphones. The

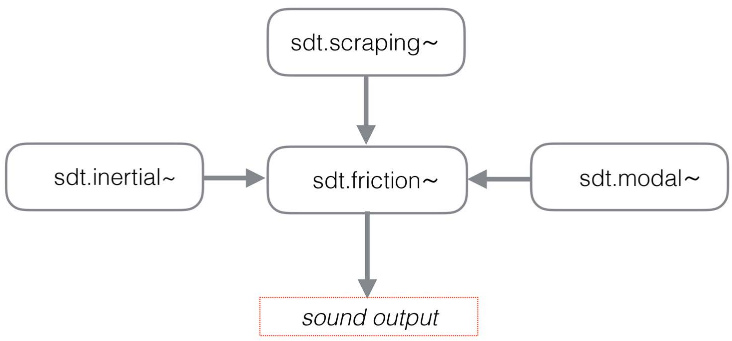

Figure 6.8: The configuration of SDT objects used to model the sound produced by each of the acoustic wind machine's slats in interaction with the cloth.

author could then remove the headphones and compare the sound to that of the acoustic wind machine, making adjustments to the program as necessary. This was an interesting process. Due to the nonlinearity of the friction model and its ability to synthesise many different sounds, a small adjustment of the parameters might transform the sound into something closer to the squeaking of a rubbery material, or the sound of a glass being rubbed. Similarly, small adjustments to the mapping connecting the rotary encoder's data to the real-time activation and modulation of the slat models could also bring considerable variation to the sound.

The design process focused on modelling as much of the roughness of the acoustic wind machine sound as possible, and was informed by acoustical analysis work undertaken to compare the acoustic and digital wind machine sounds in performance. Some of the stages of this work have been published elsewhere (Keenan & Pauletto (2016) and 2017*a*), and will be discussed in further detail in Section 6.4. The main parameters to each SDT object in the slat model are detailed in Table 6.1. Some of these are continuously modulated in performance by the motion of the crank handle. This mapping will be discussed in more detail in Section 6.3.

6.2.3 Cloth Model

The combination of several slat models activated together in performance produced quite a promising imitation of the wind-like sound of the acoustic wind machine. However, as will be explored further in Section 6.4, early iterations of the prototype digital wind machine produced some audible stepping during rotations of the crank handle due to the muting function implemented within each [*poly∼*] object, and also failed to recreate the responsive whistling sound of the acoustic wind machine. To improve both of these issues, the acoustic wind machine's cloth was investigated further as a way to model the blending of the sounds of each individual slat together as well as account for more of the whistling Aeolian tone in its wind-like sound. A simple digital model of the cloth was developed to augment the previous configuration of SDT friction models and add these features.

When the acoustic wind machine's cylinder is in motion, the friction produced between each slat and the encompassing cloth propagates through the cloth as vibrations, which become audible as sound (Farnell 2010, p.19). The material of the acoustic wind machine's cloth must have some role in dispersing the sound produced by each of its slats, as it will allow certain frequencies in the spectrum to pass through it quickly, and remove or delay

Parameters to [sdt.scraping∼]	
Surface profile (a signal)	pink∼
Probe width or grain (density of micro-impacts)	*0.03, modulated by encoder data*
Velocity (m/s)	*0.63 initial value, then activated by encoder data*
Force	*0.06, modulated by encoder data*
Parameters to [sdt.inertial∼]	
Mass of inertial object (kg)	0.01
Fragment size (to simulate crumpling)	1
Parameters to [sdt.friction∼]	
External rubbing force	output from [sdt.scraping∼]
Bristle stiffness	513.26
Bristle dissipation	24
Viscosity (speed of timbre evolution and pitch)	1.20
Amount of sliding noise (perceived surface roughness)	0.41
Dynamic friction coefficient (high values reduce sound bandwith)	0.16
Static friction coefficient (smoothness of sound attack)	0.5
Breakaway coefficient (transient of elasto-plastic state)	0.17
Stribeck velocity (smoothness of sound attacks)	0.17
Parameters to [sdt.modal∼]	
Active modes	3
Frequency Factor	1.
Frequency of each mode [Hz]	1000, 10100, 380
Decay factor	0.005
Decay of each mode (s)	0.8, 0.45, 0.02

Table 6.1: Main parameters to the SDT objects for each slat model.

others. This is evident in the characteristic whistling produced during performance. To model this aspect of the cloth, some dispersion was added to the output of each slat model within its [*poly~*] object. This dispersion was modelled using the simplest method of a digital delay line in series with an all-pass filter (Smith 2010, pp.40-41). Adding this dispersion to each digital slat produced very slight spectral shifts in the rough wind-like sound. This was achieved by mapping control data to modulate the delay time according to the position of the slat in the circular rotation, which will be discussed further in Section 6.3. This had the effect of making the combined sound of several slats together perceptually more complex.

Next, a further global model of the cloth was added to the combined output of all of the digital slats to blend their sounds together and smooth their amplitude envelopes to reduce the audible stepping. This would add more of the reverberant characteristics of the cloth to the digital model, which had not been fully accounted for by [*sdt.modal~*]. As outlined in Chapter 5, the cloth has a slightly different tension on each side of the wooden cylinder due to the way it is coupled to the frame. When the acoustic wind machine is rotated, the friction sound produced by the rubbing slats propagates through the cloth, vibrating it slightly differently on either side of the cylinder. These differences were modelled within Max/MSP. The tight side of the cloth coupled to the wooden pole on the acoustic wind machines frame is much like a string coupled to the bridge of a violin being activated with the motion of a bow in performance. As such, a method of modelling string vibration from Karplus-Strong synthesis (Karjalainen et al. 1998) was adapted to model this digitally. A digital waveguide (a bidirectional delay line) in series with a low-pass filter and an all-pass filter was used to simulate dispersion of the friction sound through the tight side of the cloth, allowing for some damping (loss of energy, and therefore a loss of frequencies) due to its coupling to the acoustic wind machine's bridge. This damping was modelled with a feedback path from the low pass filter back into the delay.

The freely hanging side of the cloth does not incorporate the same tension and damping as the tight side and so another simple model of dispersion was added to account for this, with another digital delay line in series with an all-pass filter (Figure 6.9). Implementing the added dispersion and simple cloth model within the prototype digital wind machine improved the sonic response of the program to changes in the rotational gesture and also reduced the audible stepping produced by the model. The addition of the delay lines to the model also offered the potential of further modulation of the frequency of the digital wind-like sound in performance to mimic the characteristic whistling of the acoustic wind machine.

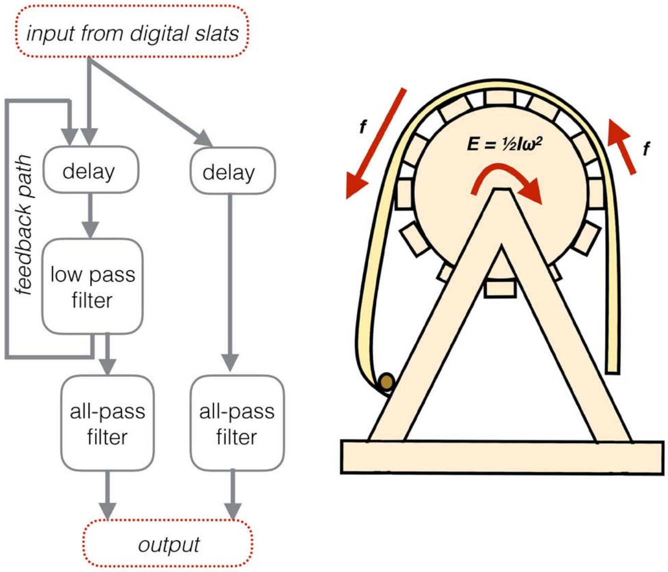

Figure 6.9: Diagram of the cloth model in Max/MSP.

This section has outlined the main soundmaking components of the prototype digital wind machine programmed in Max/MSP and the stages of work undertaken to develop them. The next section describes the mapping approach taken to facilitate the activation and modulation of the digital wind-like sound in performance.

6.3 Action-to-Sound Mapping

This chapter has so far presented the work undertaken to couple a rotary encoder to the acoustic wind machine in order to capture its movement as data, and the creation of a program in Max/MSP to map that motion to a continuous rotation between 0.0° and 360.0°. From this rotational data, the individual degree positions of each slat in rotation were parsed to create twelve distinct data streams, and global measures of velocity, acceleration, RPM and a count of rotations were also calculated in real time. The main soundmaking components of the acoustic wind machine were prototyped within Max/MSP as twelve individual SDT friction models with a small amount of dispersion, and a global model of the propagation of their combined sound through the cloth.

This section outlines the mapping approach taken to connect the data streams to the soundmaking components in Max/MSP and capture the dynamic mechanical activation and modulation of the acoustic wind machine's wind-like sound digitally, in order to produce as much of the same behaviour in the digital wind-like sound as possible during performance. Some of this work has previously been published elsewhere (Keenan & Pauletto 2017*b*).

To begin and end the digital wind sound appropriately, changes in the rotary encoder's data stream were monitored in real time to trigger a [*line∼*] object to smoothly increase or decrease the output amplitude of the program in Max/MSP and mimic the overall envelope of the acoustic wind machine's sound. As explored in Chapter 5, high speeds of rotation produce a louder wind-like sound from the acoustic wind machine. To replicate this behaviour, the velocity calculated from the rotary encoder's main data stream was used to scale the output amplitude of the prototype digital wind machine, ensuring overall lower amplitudes at slower speeds and allowing for higher amplitudes when the crank handle was rotated more quickly. Three further features of the acoustic wind machine's sound were then used to focus the rest of the mapping work:

- The amplitude envelope of each slat in rotation is complex. The rotational gesture with the crank handle produces an amplitude envelope in the acoustic wind-like sound that rises in the first half of each rotation, and then decreases in the second half.

- As the cloth moves in response to the motion of the cylinder, it alternates between losing contact with the slats or pulling tight against them. This continually shifts the quality of the friction between the slats and the cloth.

- At higher and increasing rotational speeds with the crank handle, the friction between the slats and cloth creates a characteristic *whistling* wind-like sound (an audible pitch) that seems to ascend during the first part of each rotation and descends during the second half. As the overall speed of rotation slows, so the pitch of the wind-like sound decreases.

The mapping work undertaken to recreate these features will now be considered in more detail.

6.3.1 Slat Scrape Design

As outlined in Section 6.2, the SDT object [*sdt.scraping∼*] simulates a probe sliding on a surface, and controls the output of [*sdt.friction∼*], activating the sound produced by the modelled friction interaction. There are three parameters to [*sdt.scraping~*] that can be activated in real time:

- velocity
- probe width
- external force on the probe

Each of these is allocated a [*multislider*] object that can be adjusted in real time within a range of 0.0-100.0. These simple ranges are then mapped to more precise values for [*sdt.scraping~*].

The velocity slider is what activates [*sdt.scraping~*], and therefore [*sdt.friction~*], to produce sound. At high velocities, the amplitude of the scrape is high. Manipulating the velocity slider in real time gives the resulting scraping sound a perceivable gestural shape. As the SDT environment is configured for fast prototyping of sounds, it offers some velocity profile presets with the [*function*] object to continually move the velocity slider up and down in real time and shape the quality of the resulting scrape. Just as with the crumpled ball of paper example discussed in Chapter 2, profiles such as a fast initial increase in scraping velocity followed by a slower decrease will produce a perceivable gesture in the scrape (Figure 6.10).

It would have been possible to design a velocity profile with the graphical function editor in Max/MSP and have each stream of rotational data move through the predetermined shape in order to create a scraping sound. However, without prior research establishing what the acoustic wind machine's velocity profile might be, this might overly limit the potential of this digital prototype to model its workings. Preliminary testing using this approach produced too uniform a response in the digital wind-like sound, particularly at high rotational speeds. The mapping work therefore moved to introducing more of the dynamic behaviour of the acoustic wind machine into the model. The velocity profile was instead replaced with a real-time data stream from the rotary encoder's movement. The way this data stream was shaped and filtered to dynamically activate each digital slat model will now be outlined.

The process of the acoustic wind machine's sound production, where seven individual slats rub the cloth at any one time, was used to shape the activation of each digital slat in Max/MSP. Through observing the acoustic wind machine in rotation, it was calculated that each digital slat should be silent as it moved clockwise between 290.0° and 65.0°.

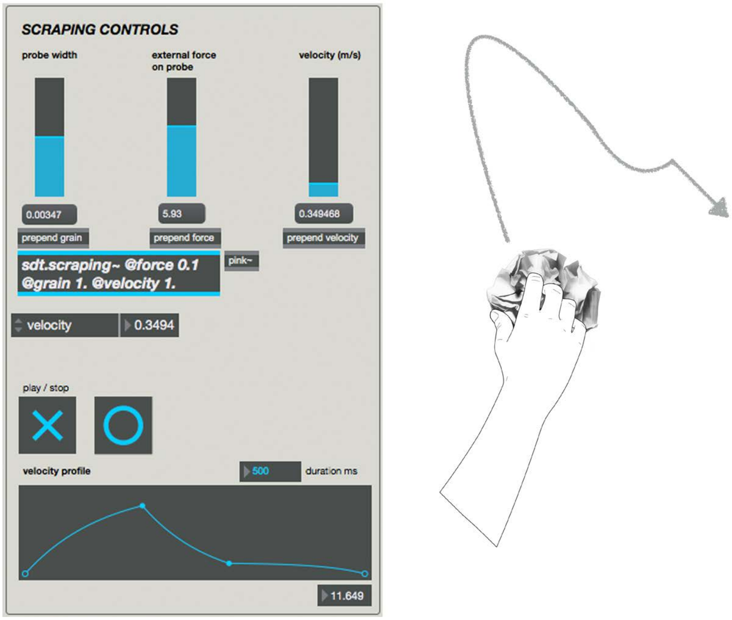

Figure 6.10: (L) The controls to [*sdt.scraping~*] as offered by its Max/MSP help file. The velocity profile on the bottom automatically triggers the scrape by moving the velocity slider up steadily, and then down more slowly. (R) How that velocity profile might perceivably suggest the shape of a real-world scraping action with a crumpled ball of paper to a listener.

The 360.0° rotational data stream parsed for each individual digital slat was tracked to determine when the slat should be producing sound in contact with the cloth, and when it should be silent. This ensured that only seven of the digital slats produced sound at any one time, mimicking the mechanical process at the heart of the acoustic wind machine's sound production. Once it passed out of activation range, the DSP processing for each slat's [*poly~*] object was muted to silence its sound.[5] The continuous rotation data was then converted to a linear up-down motion and scaled to fit the 0.0-100.0 range of the velocity slider to [*sdt.scraping~*] (Figure 6.11).

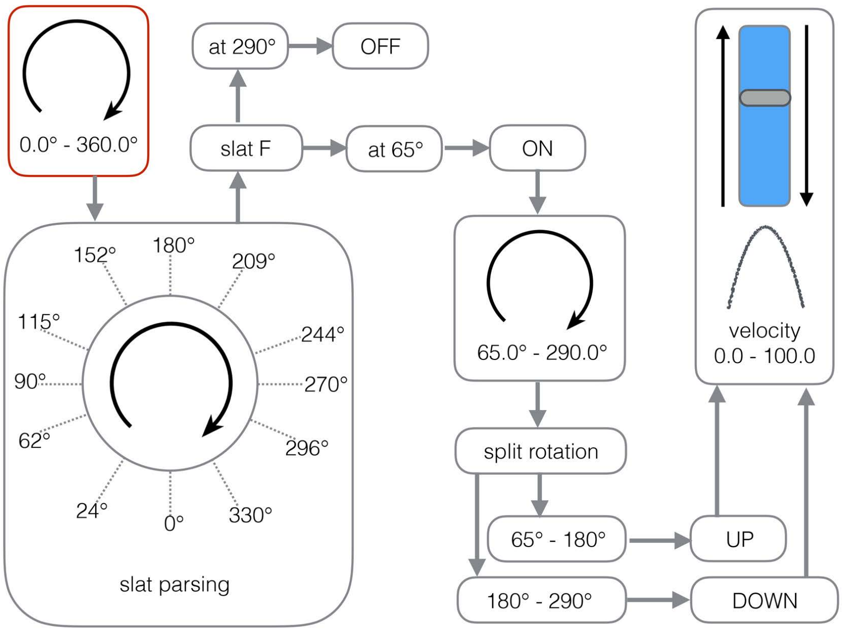

Figure 6.11: The main rotational data from the rotary encoder is parsed to create twelve individual streams, which are then tracked to facilitate DSP muting. The rotational movement of each digital slat is also converted to a linear movement to move the velocity slider up and down, producing sound from [*sdt.scraping~*].

As explored in Chapter 5, a study of the acoustic wind machine revealed that the second half of each rotation was of slightly less amplitude than the first. To replicate this behaviour in the amplitude of each digital slat, a slight offset was added to the second half of the linear movement as the velocity slider descended. The overall linear movement of the

[5]As discussed previously in Section 6.2, while the choice to temporarily mute the sound of each digital slat improved the overall response of the prototype digital wind machine in terms of the CPU demands on the computer, this necessitated the further development of a cloth model to reduce the audible stepping that was produced in its wind-like sound by this muting.

velocity slider was then also scaled according to the main rotational velocity calculated from the rotary encoder's movement, ensuring that each individual digital slat would produce a sound of lower amplitude when the crank handle was rotated more slowly. To model the increased effort required in the upstroke of the rotational gesture, the data was then logarithmically filtered through a simple signal inertia model [6] to ensure that the scraping gesture would not activate in a smooth, linear way, but would require some energy to produce a maximum value (Figure 6.12).

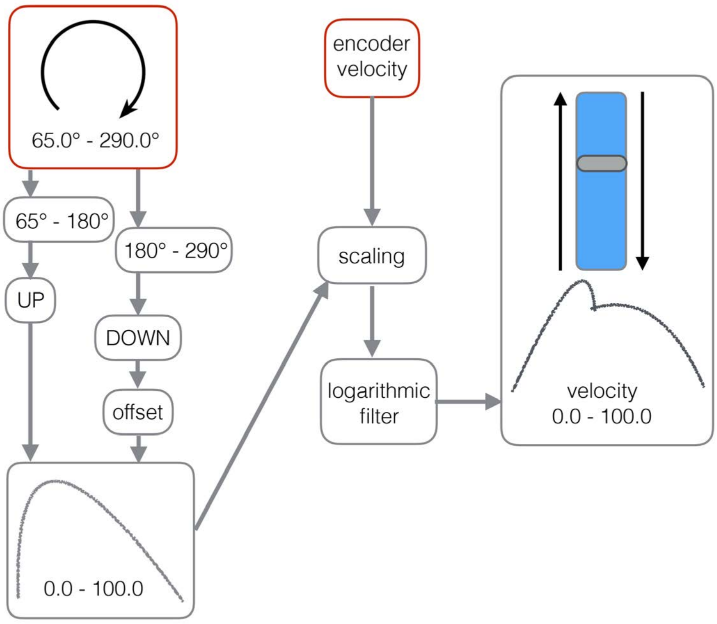

Figure 6.12: Adding complexity to the velocity slider data for each digital slat, with the addition of an offset, scaling according to the main rotational velocity of the encoder, and a logarithmic filter.

In this way, complexity was added to the simple linear movement of the velocity slider for each [*sdt.scraping~*], producing a dynamic velocity profile for real-time performance. With the design of the scrape of each digital slat established, the way the movement of the cloth was modelled will now be considered.

6.3.2 Cloth Movement

As explored in Chapter 5, the acoustic wind machine's cloth also moves in response to the movement of the wooden cylinder as it is rotated with the crank handle. The

[6] The logarithmic filtering was developed from code by jvkr on the Cycling '74 forum.

surface area of each slat in contact with the cloth, and the friction force between them, is continually shifting depending on the flywheel response of the wooden cylinder to the rotational motion. The way the slats are placed around the cylinder also has a role in this, as even at slow rotational speeds the cloth can be seen tightening and loosening as the cylinder rotates. These mechanical processes increase the complexity of the friction sound produced by the interaction of the slats and cloth.

The lifting of the cloth observed at particularly high rotational speeds, which pulls it taut against the slats it does make contact with, may also have some role in the characteristic whistling sound of the acoustic wind machine. To model the changing surface area of the slat and force of the cloth, the probe width and force parameters to [*sdt.scraping~*] were modulated slightly around their initial value in real time. The probe width parameter was programmed to vary with the degree position of the digital slat, with the highest value corresponding to the top of the acoustic wind machine at 180°, where the slat is most fully in contact with the cloth. The probe width slider to each [*sdt.scraping~*] moved up and down within a restricted range of values during the digital slat's rotation.

The acceleration data calculated from the movement of the rotary encoder was used to modulate the force parameter to each [*sdt.scraping~*], again with the highest potential value corresponding to the top of the acoustic wind machine at 180°, where the slat is flattest against the cloth. This moved the force slider to each [*sdt.scraping~*] up and down in response to the acceleration of the crank handle. Rather than have these two parameters increase and decrease in a smooth and regular way during rotation, a very small amount of phase shifting was added to randomly shift the phase of the functions driving the probe width and force changes in real-time. Each digital slat's probe width and force sliders would be at a slightly different stage in their movement in real-time, mimicking the highly responsive and sometimes chaotic motion of the cloth (Figure 6.13).

Next, the addition of the characteristic whistling behaviour to the prototype digital wind machine will be outlined.

6.3.3 Responsive Whistling

Although the acoustic wind machine produces a perceivably noisy and rough sound during rotation, it responds to increased speed and energy from the crank handle with an audible whistling pitch. The complex modelling and mapping so far described in this chapter did

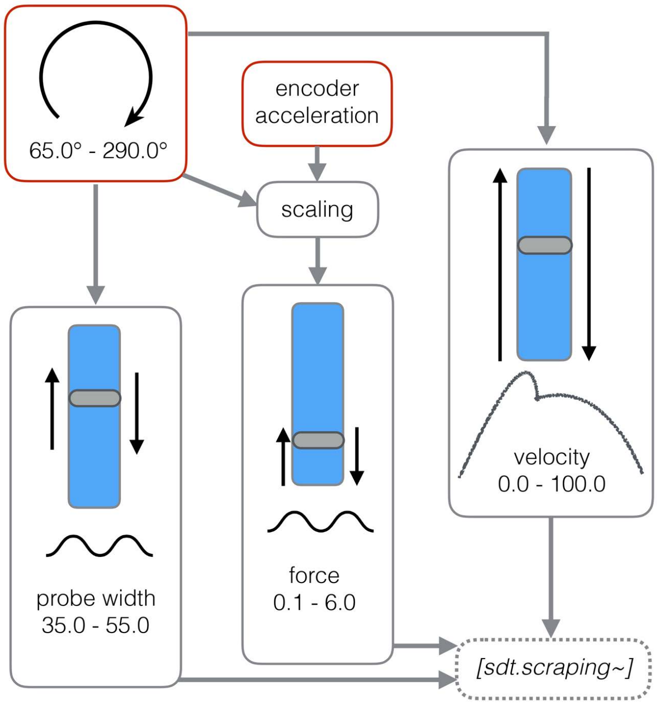

Figure 6.13: The real-time modulation of the probe width and force sliders for each [*sdt.scraping~*].

not produce enough of this behaviour in the prototype digital wind machine to make a similarly perceivable pitch. This made the early iterations of the digital wind-like sound perceivably less responsive to changes in rotational speed during performance, something that was highlighted by participants during a pilot of the experimental procedure outlined in Chapter 7. Some additional spectral complexity was first added by modulating the delay time to the dispersion model within each digital slat according to its degree position. To add some of the characteristic whistling of the acoustic wind machine's wind-like sound, the delay time to each side of the digital cloth model was programmed to modulate slightly according to the degree position of the main rotation data (and therefore the crank handle). This produced a slightly ascending pitch in the first half of the rotational gesture that descended during the second half. The variation in pitch was also scaled according to the velocity of the rotary encoder to ensure a lower pitch at slower speeds (Figure 6.14).

With the full mechanical mapping implemented in Max/MSP, a simple gesture recorder was programmed using a [*coll*] object to display the main data stream of the slat tracking with the crank handle and rotary encoder and record its motion during performance.

This section has described the mapping approach taken to recreate the complexity of the acoustic wind machine's sonic response in performance as a prototype digital wind

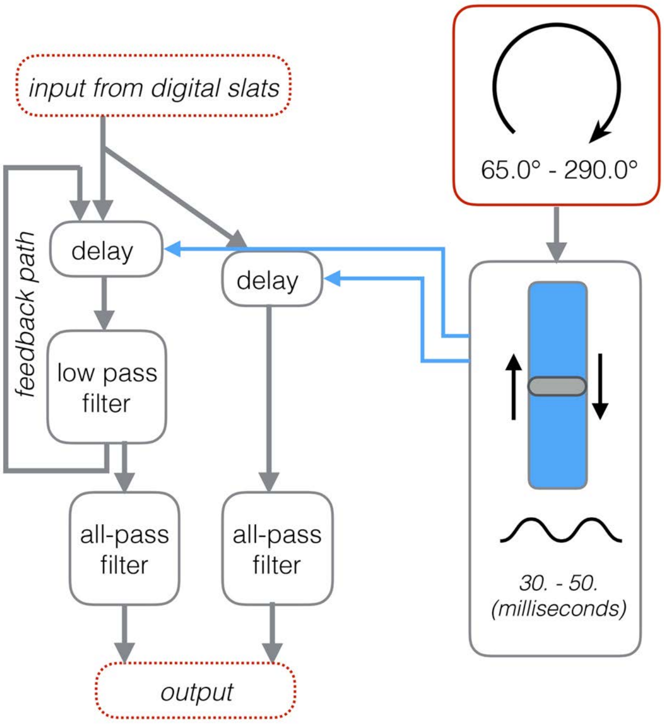

Figure 6.14: A diagram of the modulation of the cloth pitch.

machine in Max/MSP. By adding complexity to the velocity slider for each [*sdt.scraping*∼] object, the scraping action of each slat on the acoustic wind machine was modelled more convincingly. The probe width and force sliders to each instance of [*sdt.scraping*∼] allowed the responsive motion of the cloth to be added to the digital model. Finally, by modulating the delay time to the global cloth model in real time, some of the audible whistling pitch produced by the acoustic wind machine was added to its digital counterpart. Throughout the process of refining and adjusting the prototype digital wind machine, acoustical analysis was used to determine how closely it was managing to mimic the sonic behaviour of the acoustic wind machine when performing the same rotational gesture.

The next section outlines this analysis work, and presents the final acoustical evaluation and comparison of the acoustic and digital wind-like sounds that were used for the experiment with participants reported in Chapter 7.

6.4 Acoustical Evaluation

The wind-like sounds of the acoustic and prototype digital wind machines were acoustically evaluated and compared during this part of the design-led methodology in order to inform the choice of sensor to activate the digital wind-like sound, as well as the programming approach in Max/MSP. Some of the stages of this work have previously been published elsewhere (Keenan & Pauletto (2016) and 2017*a*). A summary of the procedure used and the issues that arose in the evaluation of the wind machines acoustically now follows. The final acoustical comparison undertaken to compare the sonic responses of the two wind machines in performance before proceeding with the experiment outlined in Chapter 7 is also presented.

6.4.1 Gestures for Analysis

To observe the acoustic wind machine's dynamic response to variations in the rotation of its crank handle and examine how well this had been modelled in Max/MSP, three distinct rotational gestures were chosen for the simultaneous recording of the acoustic and digital wind-like sounds in performance:

- A single steady rotation.
- Five steady rotations to produce a continuous sound.
- Ten rotations that start at speed, but then diminish in energy.

As outlined previously in Section 4.1.5, the simultaneous recording setup captured both the acoustic and digital wind-like sounds into the same Pro Tools session. This approach also allowed the effectiveness of the rotary encoder hardware and Max/MSP program to be evaluated, as any significant delay in the progress of the digital wind-like sound could be easily observed. As this recording process was being used to determine how well the digital wind-like sound tracked with its acoustic counterpart during performance, the author listened to the acoustic wind machine in headphones while both wind-like sounds were recorded to Pro Tools. Many examples of the distinct gestures were recorded, and representative audio clips were chosen for analysis using the MIR Toolbox in Matlab (Lartillot & Toiviainen 2007). As discussed previously in Chapter 4, the amplitude envelope and spectrum of each of the sounds was examined, as well as some further numerical measures of these features like event density and spectral centroid.

6.4.2 Stages of Acoustical Evaluation

Acoustical evaluation proved vital to the development of the final prototype digital wind machine. The wind-like sound produced by the acoustic wind machine is highly complex.

Perceivably noisy and rough, it shifts from a machine-like repetitive sound to a convincing wind effect depending on the energy exchange between the performer grasping the crank handle and the flywheel mechanism of the cylinder and A-frame. Early iterations of the prototype digital wind machine focused on replicating the perceivable roughness in the acoustic wind effect produced by the cloth. This produced promising results in the prototype digital wind machine in terms of how its amplitude envelope tracked with its acoustic counterpart in performance, but the resulting sound possessed none of the perceivable shifts in frequency to recreate the responsive whistling from the multiple sources of friction (Figure 6.15). This produced a digital wind-like sound that seemed to track well with the rotation of the acoustic wind machine's crank handle, but did not seem to react to changes in the energy of the performer's rotational gesture. Audible stepping could also be heard as each slat model was muted when its acoustic counterpart moved out of range of the cloth. This was highlighted by pilot study participants, as will be reported in Chapter 7.

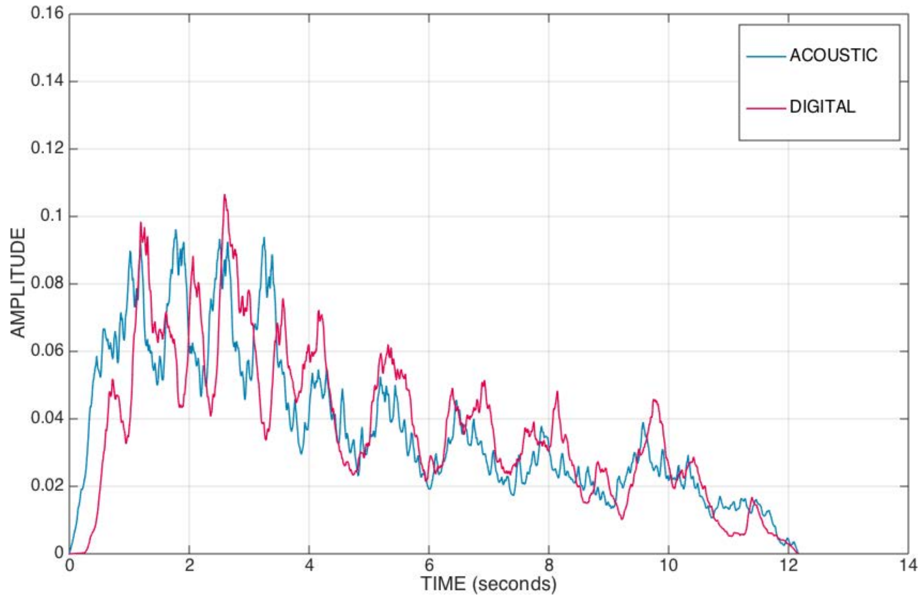

Figure 6.15: An amplitude envelope for 10 rotations starting at speed and diminishing in energy performed with both the acoustic wind machine (shown in blue), and an early iteration of the prototype digital wind machine (shown in red) before some responsive whistling behaviour was added to its program (First published in Keenan & Pauletto (2017*a*)).

Further work was then undertaken to introduce more modulation of the parameters to the prototype digital wind machine in performance, which successfully added this responsive whistling in performance, as discussed in Section 6.3. However, this added responsiveness made the prototype digital wind machine perceivably less rough, and slightly more wind-like than its acoustic counterpart in performance. The added modulation in Max/MSP changed the robust amplitude envelope that had previously tracked more closely with that of the acoustic wind machine. This made changes in the energy of the rotational gesture sometimes more difficult to perceive in the resulting wind-like sound. Further rehearsal and calibration tightened the range of the final modulation values in order to keep as much of the gestural responsiveness in the prototype digital wind machine as possible, while reducing the delay in the amplitude envelope.

The final acoustical analysis of the wind machines used for the experiment with participants discussed in Chapter 7 will now be presented.

6.4.3 Amplitude Envelope

The amplitude envelope of the acoustic and digital wind machines was examined to determine how responsive the prototype digital wind machine was in comparison with its acoustic counterpart, and also how the combination of the twelve digital slats compared with the combination of the twelve acoustic slats. In the following graphs, the acoustic sound is shown in blue and the digital sound is shown in red.

6.4.3.1 Single Rotation

For a single steady rotation of the crank handle, the acoustic and digital wind-like sounds sounds share a similar envelope shape, showing that the combination of twelve digital slats is being triggered in rotation in a way that accurately models the amplitude envelope of the acoustic wind machine. However, the increase in amplitude for the first part of the rotational gesture is more pronounced in the acoustic wind machine. There is a small amount of latency in the start of the digital sound, and it ends more abruptly at the close of the gesture (Figure 6.16).

6.4.3.2 Five Rotations

For five steady rotations of the crank handle, the acoustic and digital wind sounds again share a similar amplitude envelope shape. There is some latency in the progress of the prototype digital wind machine, although the distinct rotations can be seen in its envelope. The digital wind sound appears to be moving slightly ahead of its acoustic counterpart

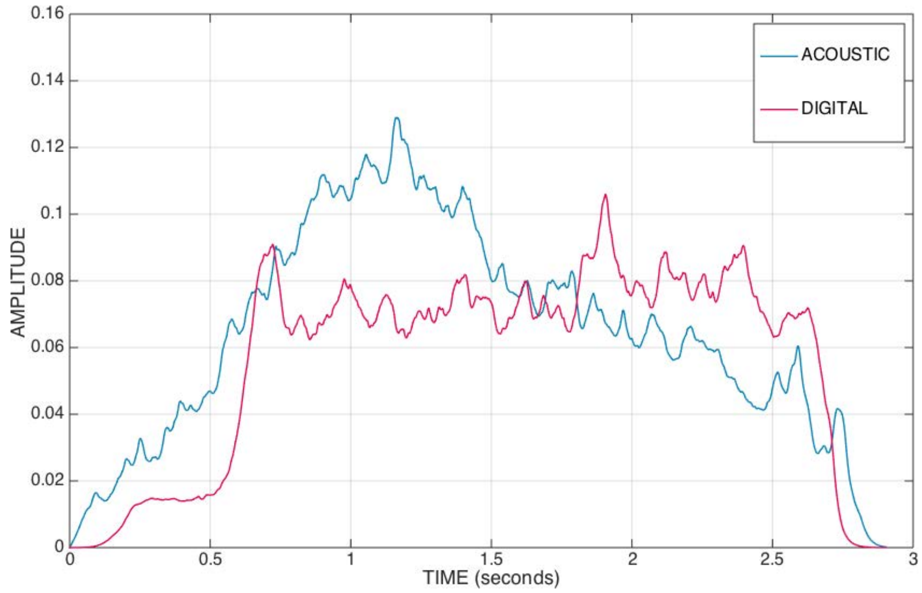

Figure 6.16: A comparison of amplitude envelopes for the same single rotation of the acoustic wind machine and the prototype digital wind machine programmed in Max/MSP.

during the fifth rotation. This may be due to the accumulated energy of the logarithmic filter producing signal inertia in Max/MSP, or it may be due to the hardware of the rotary encoder and gearing moving slightly ahead of the crank handle. There is also an anomaly in the acoustic wind machine's envelope during the fifth rotation as it responds to the energy in the rotational gesture (Figure 6.17).

6.4.3.3 Ten Rotations

For ten rotations of the crank handle, starting at speed and then diminishing in energy, the acoustic and digital wind-like sounds do not share as similar an envelope as the other two test gestures. The prototype digital wind machine is unable to increase its amplitude as quickly as its acoustic counterpart does at the start of the gesture, although when it does increase in amplitude it appears to track several fast rotations. When the acoustic wind machine falls in amplitude as the gesture slows to a stop, it takes the prototype digital wind machine longer to show a similar drop (Figure 6.18). This may be due to the logarithmic filtering of the data to each digital slat keeping the values to the velocity slider high for too long, or the resonance of the global cloth model delaying the fall in amplitude of the combined sources of friction. It may also be due to the hardware configuration capturing the acoustic wind machine's motion as data - in particular, using gearing with

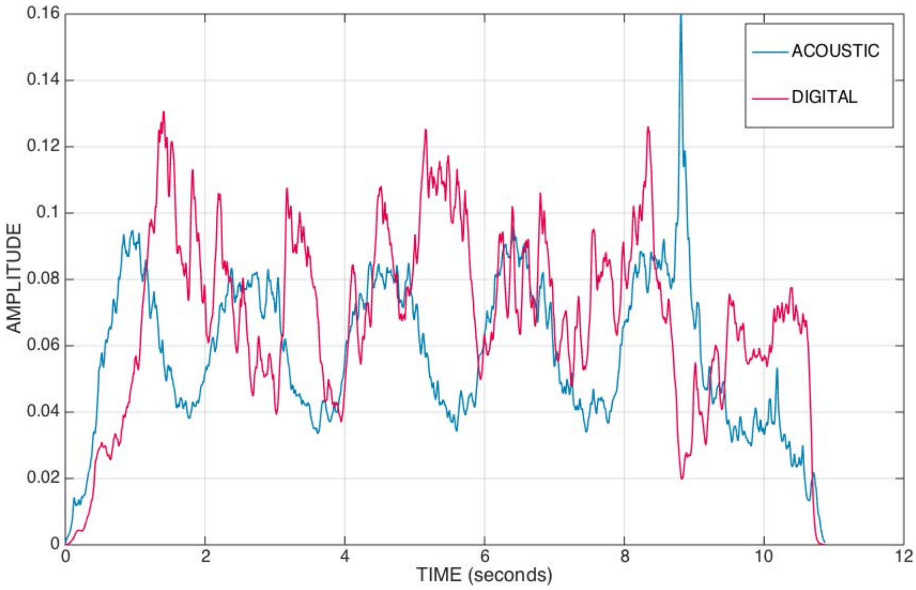

Figure 6.17: A comparison of amplitude envelopes for the same five rotations of the acoustic wind machine and the prototype digital wind machine programmed in Max/MSP.

finer teeth may allow the rotary encoder to track more smoothly with the acoustic wind machine's crank handle.

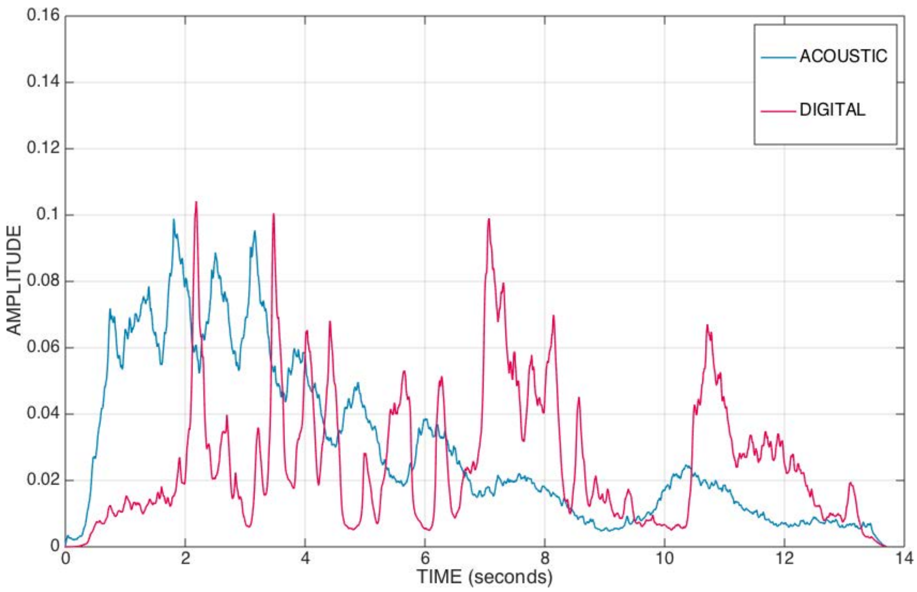

Figure 6.18: A comparison of amplitude envelopes for the same ten rotations (starting at speed and then diminishing in energy) of the acoustic wind machine and the prototype digital wind machine in Max/MSP.

6.4.4 Event Density

A further measure of the amplitude envelope, event density, was calculated for each of the gestures performed with the acoustic and digital wind machines in order to compare the number of onsets in each per second. The results are outlined in Table 6.2.

Gesture	Acoustic Wind Machine	Prototype Digital Wind Machine
1 rotation	3.20	4.62
5 rotations	0.55	3.25
10 rotations	0.59	1.03

Table 6.2: Event density measures for the three gestures simultaneously performed with both wind machines.

The event density results for both of the wind-like sounds are close to each other, although it seems that for five regular rotations the prototype digital wind machine is not producing as smooth an amplitude envelope as its digital counterpart. This may be due to the influence of the cloth in blending the combined sources of friction not being fully accounted for as part of the Max/MSP program.

6.4.5 Spectra

The spectra of the acoustic and digital wind-like sounds were examined to see how closely the prototype digital wind machine was modelling the complex sound of its acoustic counterpart.

6.4.5.1 Acoustic Wind Machine

The spectrum of the acoustic wind machine is very noisy, but has some peaks in the low and mid frequencies and a significant amount of high frequency energy (Figure 6.19).

6.4.5.2 Prototype Digital Wind Machine

The spectrum of the prototype digital wind machine has replicated some of the noisiness of its acoustic counterpart, but has a distinct peak in its spectrum at around 1KHz. It also lacks the same energy at high frequencies (Figure 6.20). A lack of power at high frequencies may be a common pitfall when creating digital models to imitate acoustic sounds (Ronan 2010, p.49). In this case, the choice to downsample within each [*poly∼*] object may have produced the high frequency rolloff evident here.

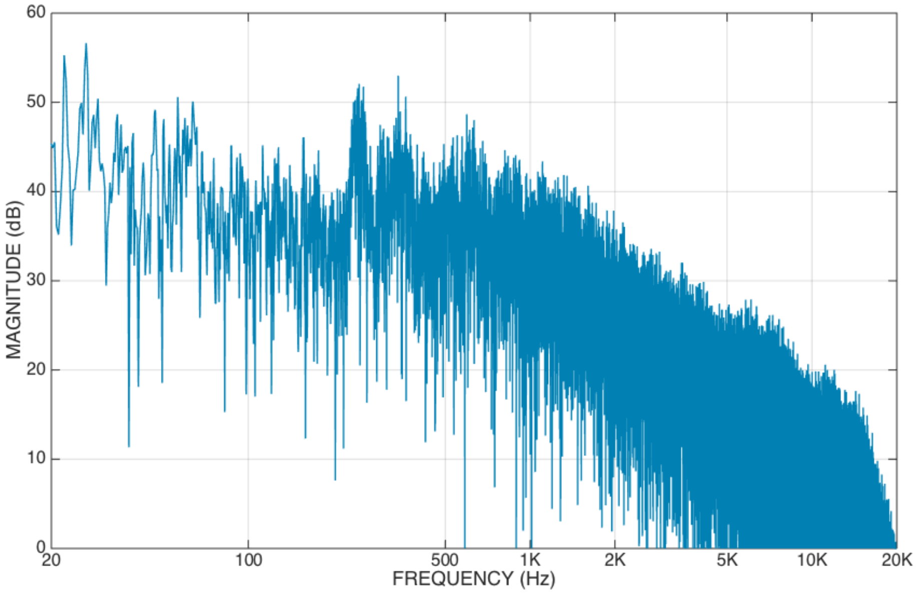

Figure 6.19: Spectrum of the Acoustic Wind Machine (1 rotation)

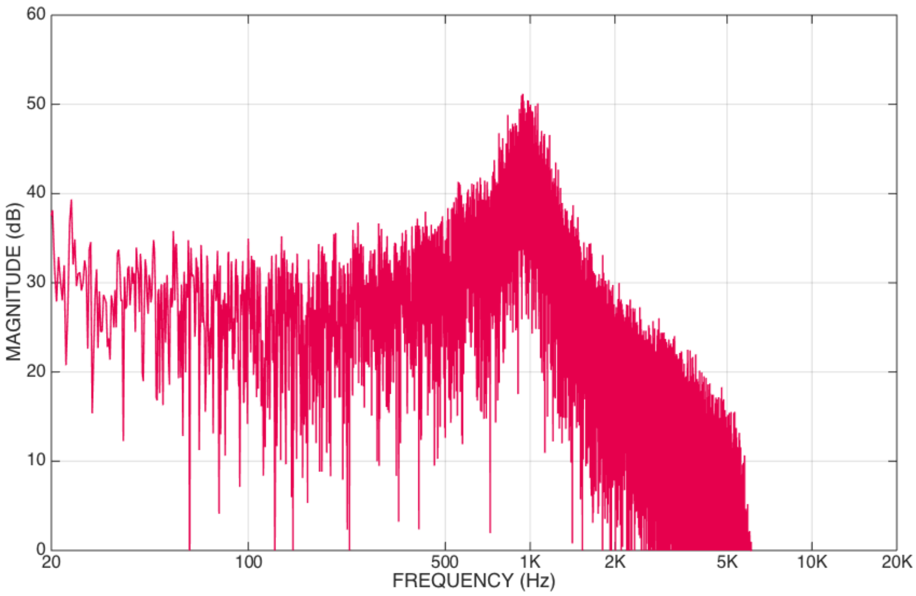

Figure 6.20: Spectrum of Digital Wind Machine (1 rotation)

6.4.5.3 Spectrograms

Spectrograms of the acoustic and digital wind-like sounds show how their frequency responses change in time, and with variation in the rotational gesture. Figure 6.21 shows the gesture of ten rotations starting at speed and diminishing in energy performed with the acoustic wind machine. Each rotation can be clearly seen, and the high frequency energy diminishes as the rotational gesture slows down.

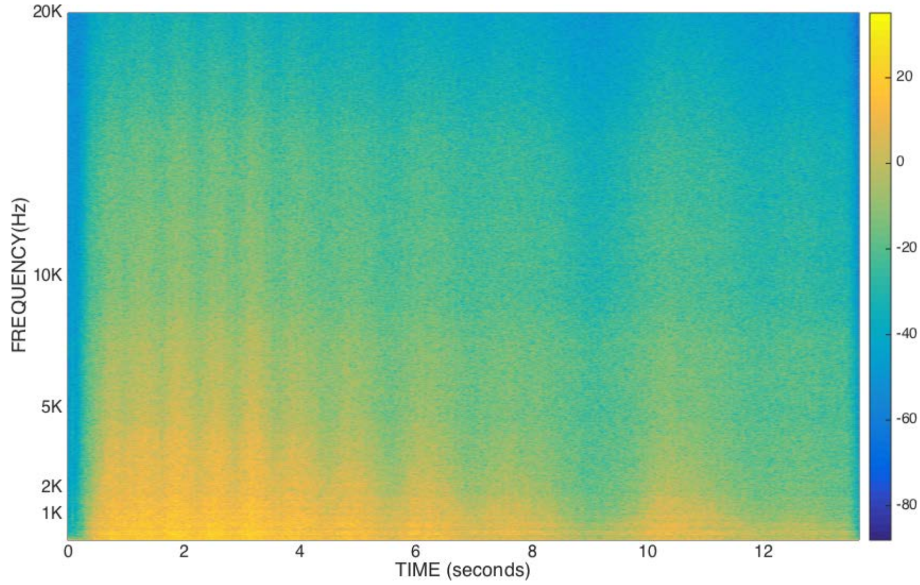

Figure 6.21: Spectrogram of 10 rotations with the acoustic wind machine.

The spectrogram of the prototype digital wind machine does not show the same distinctive rotations, and it is clear that this wind-like sound has a lot of energy at 1KHz throughout variations in the speed of rotation (Figure 6.22).

6.4.6 Numerical Measures of the Spectra

As outlined previously in Chapter 4, several numerical measures of the spectra of both wind machines were also used to evaluate and compare their sounds. The results are presented in Table 6.3.

These measures confirm that the acoustic wind machine is significantly brighter than the digital prototype, although both wind-like sounds maintain their brightness across the three different performance gestures. Both wind-like sounds share quite a similar harmonic-to-noise ratio, as evidenced by the simple inharmonicity measure. Again, this measure does not change significantly with the performance gesture across both sounds. The statistical measures of the spectra (centroid, spread and skewness) confirm that the

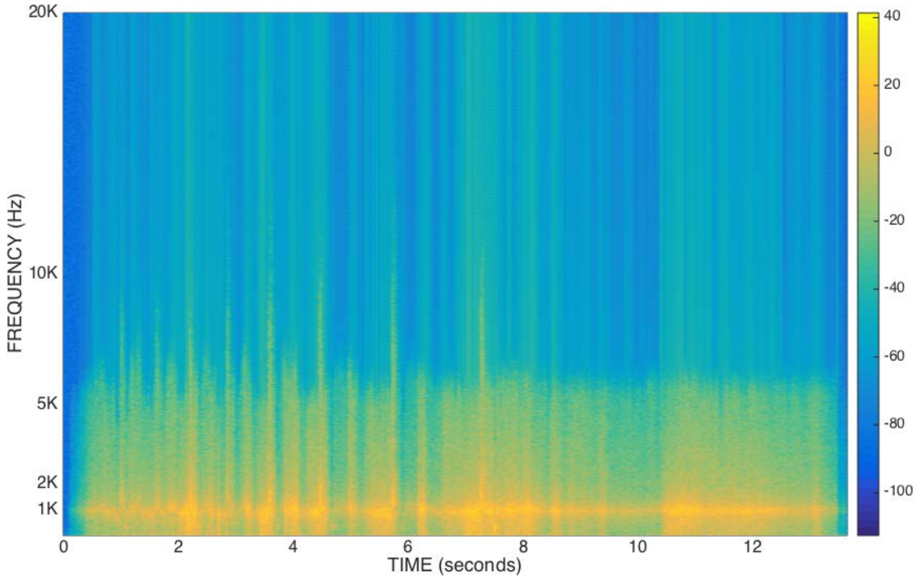

Figure 6.22: Spectrogram of 10 rotations with the prototype digital wind machine.

Spectral Measure	**Gesture**	**Acoustic Wind Machine**	**Prototype Digital Wind Machine**
Brightness (0. to 1.)	1 rotation	0.49	0.22
	5 rotations	0.49	0.16
	10 rotations	0.54	0.21
Inharmonicity (0. to 1.)	1 rotation	0.49	0.41
	5 rotations	0.48	0.40
	10 rotations	0.48	0.42
Spectral Centroid (Hz)	1 rotation	3293.16	1301.98
	5 rotations	3276.48	1119.19
	10 rotations	3145.02	1228.28
Spectral Spread (Hz)	1 rotation	4122.05	1098.46
	5 rotations	4086.12	926.68
	10 rotations	3669.32	1235.88
Spectral Skewness (-3.0 to 3.0)	1 rotation	1.79	3.12
	5 rotations	1.82	2.78
	10 rotations	1.96	4.46

Table 6.3: Numerical measures of the spectra of the acoustic wind machine and prototype digital wind machine across the three performance gestures.

acoustic wind machine has a significant amount of very high frequency energy while also being spread out across the spectrum. The prototype digital wind machine, on the other hand, lacks that energy at very high frequencies and does not have spectral energy across the spectrum. Instead, it is highly concentrated at the single peak around 1KHz.

6.4.7 Comparing the Two Wind-Like Sounds

This acoustical evaluation has established that the wind-like sound produced by the prototype digital wind machine is quite similar to that of its acoustic counterpart, although it does require further development to increase its spectral energy at high frequencies. The digital model is responsive in performance, and is capable of some sonic variation depending on the rotational gesture performed with the acoustic wind machine's crank handle.

The progress of the amplitude envelope of the digital wind-like sound is quite close to that of its acoustic counterpart at regular speeds, although it does not track as closely with greater variation in rotational speed, as evidenced by the analysis of the gesture of ten rotations. This confirms that the programming approach presented in this chapter, which centred on twelve individual physical models of friction and a mechanical mapping model, has begun to accurately replicate the soundmaking components of the acoustic wind machine.

In performance, the rotations of the crank handle can be made audible in the resulting sound of the prototype digital wind machine when it is activated with a repetitive gesture, much like the acoustic machine. However, at less regular speeds the prototype digital wind machine can sound perceivably more wind-like than its acoustic counterpart. With robust mapping and an efficient Max/MSP program, it may be that the prototype digital wind machine is satisfying to interact with but will not be able to imitate the sound of the acoustic wind machine fully. It should therefore be established whether the digital wind-like sound is perceivably similar to its acoustic counterpart in performance. As outlined in Chapter 4, this is investigated by undertaking an experimental evaluation of both wind machines with participants. Given the prototype digital wind machine's response to variability in rotational speed, the performance step of the experiment will ask participants to play simple gestures. The full experimental procedure is reported in Chapter 7.

The next section outlines a parallel piece of work undertaken as part of this research to create a digital crank interface to activate the program in Max/MSP and examine how it might compare perceptually in performance to the acoustic wind machine.

6.5 Digital Crank Interface

This chapter has so far examined the activation and modulation of the prototype digital wind machine using its acoustic counterpart as a performance interface. A parallel piece of work was also undertaken to construct a fully digital crank interface to allow the program in Max/MSP to be performed independently of the acoustic wind machine. Although the experiment reported in Chapter 7 does not evaluate this performance interface, this work is presented here in order to examine the perceptual experience it affords when performing the digital wind-like sound, and to explore how its tactile and kinaesthetic feedback compares with that of the acoustic wind machine. As outlined in Chapter 4, a digital interface that affords a rotational gesture is another potential method of placing a wind-like sound in-hand so that it can be continuously interacted with.

6.5.1 First Iteration

The first iteration of the digital crank design was based on an Arduino and digital rotary encoder. The crank handle to drive the rotary encoder was created by welding pieces of 8mm steel rod together (Figure 6.23). The digital rotary encoder rotated continuously, but clicked with each increment as it turned. The shaft of the crank handle was fixed to the shaft of the rotary encoder with metal epoxy, and the two were pressed together. The Arduino and rotary encoder were housed in a metal enclosure, allowing the crank handle to protrude from the front and offer an obvious perceived affordance to a prospective performer.

With the data from this digital encoder mapped to the Max/MSP program described in this chapter, this crank handle activates and modulates the digital wind-like sound, but offers quite a different sensory experience to the performer. The crank is light, and given the short distance from its handle to its shaft it completes a full rotation much more quickly than the acoustic wind machine's crank handle. The performer does not need to stand to use this interface, but can engage the hand, arm and shoulder in the rotational gesture if the digital crank is placed on a desktop and operated from a standing position.

This interface offers reliable sonic feedback, as sound is tightly coupled to the

Figure 6.23: The first iteration of the digital crank handle and the digital rotary encoder it turned.

performer's action, and begins when the rotation begins. During performance, sound does not vibrate through the crank handle. The continuous clicking of the digital rotary encoder is sometimes felt in the hand when the handle is rotated. There is no sensation of continuously shifting weight or motion due to the absence of the acoustic wind machine's flywheel mechanism. The smooth metal handle offers no movement or sensation of turning under the hand during performance, and the rotational movement feels much smaller and more precise than that of the acoustic wind machine.

6.5.2 Second Iteration

To develop this design further, a second iteration of the digital crank was created. This time, the same kind of smoothly rotating rotary encoder that was coupled to the acoustic wind machine was used.[7] The rotary encoder and Arduino were housed in the same metal enclosure as the first iteration of the interface. A new crank handle was also fabricated from a repurposed table winder handle. The original shaft of the table winder was drilled out and replaced with a piece of 10mm steel rod. The steel rod was coupled to the table winder handle with a large bolt held in place with metal epoxy. The rotary encoder's shaft was then covered in a layer of metal epoxy, and the steel rod was pushed onto it to couple the two together (Figure 6.24).

Once tested, the metal enclosure was sealed (Figure 6.25). This new digital crank interface was then used to perform the program in Max/MSP.

While the experience of performance with this new version of the digital crank was

[7]This was a Bourns type ENA1J-B28-L00128L.

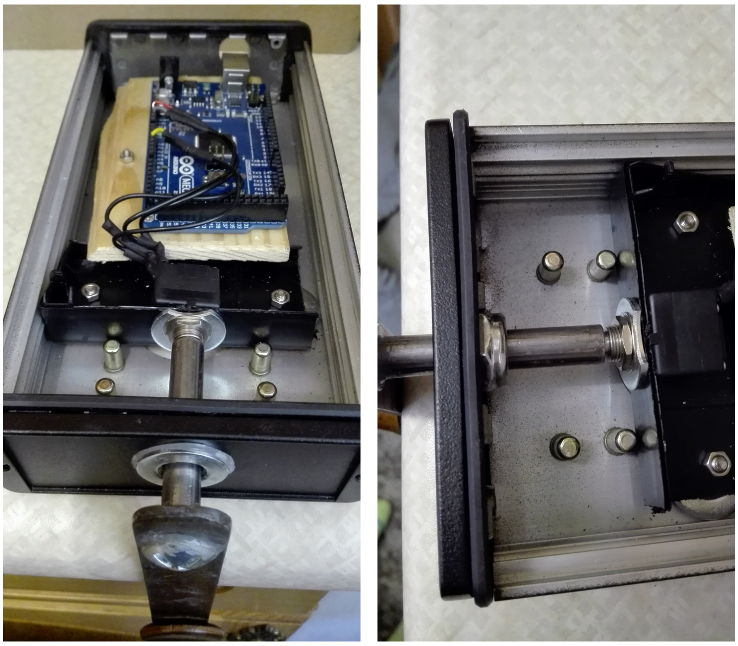

Figure 6.24: (L): The Arduino, rotary encoder and crank configuration in the metal enclosure. (R): Fixing the shaft of the crank handle to that of the rotary encoder.

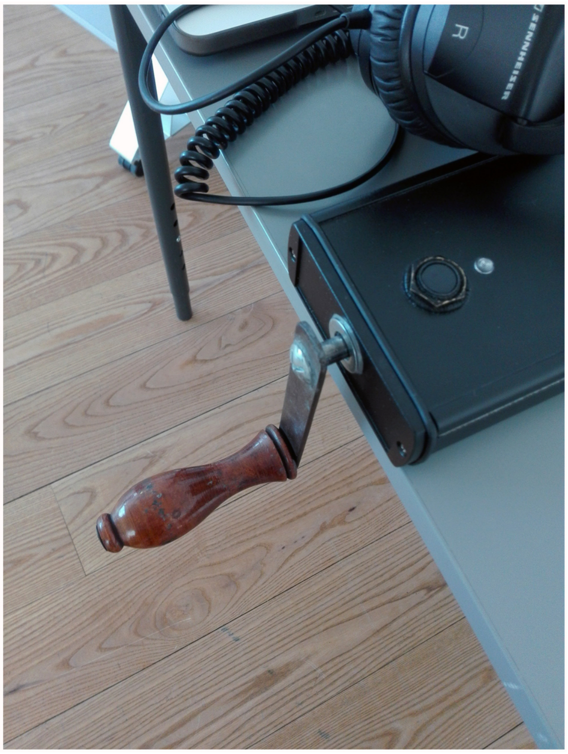

Figure 6.25: The second iteration of the digital crank interface.

quite similar to the first iteration of the design, some qualities of this interface proved interesting. The increased distance of the wooden handle from the central shaft makes the rotational gesture feel perceivably larger in size and less precise than the first handle. This sense of largeness and inefficiency gives a more natural feeling of connection to the resulting digital wind sound. The freely spinning wooden handle of the crank allows the hand to continuously reorient itself as the rotation progresses, making the gesture more compelling and bringing the sensation perceivably closer to that of the acoustic wind machine's handle. While there is still no sensation of shifting weight from the flywheel mechanism of the acoustic wind machine's cylinder and axle, the wooden handle does move freely back and forth during rotation. This crank is heavier than the first iteration of the design, and the wooden handle gives a more satisfying sensation of tactile feedback than the smooth metal handle.

6.5.3 Comparisons to the Acoustic Wind Machine

When performing the digital wind-like sound with the digital crank interface, the rotational gesture is freed from the mass of the acoustic wind machine's cylinder and requires none of the same effort to begin, adjust and end. The rotation feels more uniform under the hand without the tightness of the cloth on one side, and so the perceivable shape the hand makes is more circular. Given the uniformity of the rotary encoder's movement in this interface, the response of the program in Max/MSP is itself more uniform and steady.

Performance with the digital crank requires more concentration on how the action modulates the wind sound, particularly as this gesture is much faster and more efficient than that afforded by the acoustic wind machine. There is a sensation of trying to keep the movement slow in performance to bring the resulting sound closer to that produced when activated by the acoustic wind machine interface, and a feeling of holding and guiding the sound rather than it being intuitively in-hand.

This section has briefly examined the stages of parallel work undertaken to create and explore the potential of a digital crank interface for performing the digital wind model in Max/MSP as one of the stages of the transition between the acoustic and digital wind-like sound presented in Chapter 4. The next section summarises the work presented in this chapter.

6.6 Summary

This chapter has outlined the second stage of the design-led methodology to program a prototype digital wind machine using the SDT physical model of friction in Max/MSP:

- A rotary encoder was installed on the acoustic wind machine with laser cut gearing to send rotational data to Max/MSP.
- The entity-action model of the acoustic wind machine presented in Chapter 5 was used to inform the programming of a digital model of each slats process of sound production using the Sound Designer's Toolkit (SDT) physical model of friction.
- A mechanical mapping model was used to recreate the real-time performance of a compound sound created by twelve sources of friction in the digital domain.
- Performance gestures with the acoustic and digital wind machines were simultaneously recorded to facilitate an acoustical comparison between the results of both systems, allowing for further adjustment and mapping to the sound model parameters.
- A parallel piece of work was undertaken to create a fully digital crank interface to facilitate performance of the program in Max/MSP without the acoustic wind machine. The experience of performance with this interface was compared to a performance using the acoustic wind machine to activate the digital wind-like sound.

The next chapter reports on an experiment designed to evaluate the acoustic and digital wind-like sounds, and the sonic interactivity of the two wind machines, with participants.

Chapter 7

Experimental Evaluation and Analysis

This chapter reports on the evaluation of the acoustic wind machine and its prototype digital counterpart in an experiment with participants. As outlined in Chapter 4, this experiment was designed to encompass a range of procedures in order to explore both the perceptual qualities of the wind-like sounds and the continuous sonic interaction afforded by the two wind machines as fully as possible without overly burdening participants with a lengthy procedure. Participants were presented with three distinct experimental steps. The first step asked them to rate a corpus of sounds, including the wind-like sounds produced by both wind machines, in terms of their perceived similarity to each other. The second asked them to rank the wind-like sounds in terms of the perceived rotational speed that produced them. The final step of the experiment asked participants to perform with the acoustic and prototype digital wind machines and describe their experience of the continuous sonic interactions.

Section 7.1 outlines the aims and hypotheses of this evaluation, and reports briefly on the pilot study conducted prior to the full experiment. Section 7.2 presents the procedure and results of the first listening step of the experiment, where participants rated a corpus of sounds in terms of their similarity. Section 7.3 presents the procedure and results of the second listening step of the experiment, where participants ranked the acoustic and digital wind-like sounds in terms of the rotational speed of the mechanism that may have produced them. Section 7.4 presents the procedure and results of the final step of the experiment, where participants performed with both of the wind machines to imitate some recordings of acoustic and digital wind-like sounds, and also describe their experiences of the interactions. Section 7.5 summarises and discusses the results reported from each step of the experiment.

7.1 Experiment Aims and Hypotheses

As explored in Chapter 4, real-world wind sounds have been subject to little perceptual study, and theatre sound effects have not yet been evaluated as sonically interactive objects. As such, this experiment design builds on prior research across the fields of environmental sound perception, SID and DMI evaluation to examine how a wind effect compares perceptually to a real-world wind sound, and how a wind-like sound might suggest a performative action to a listener. The compelling continuous sonic interaction afforded by the acoustic wind machine explored in Chapter 5 is also compared to that of the prototype digital wind machine, to examine how much of the enactive quality of the original sonic interaction has been transferred to the encounter with a digital sound.

The aims of this experiment were:

- To examine the perceptual space of a wind-like sound through a comparison of acoustic, natural and digital winds in order to explore how a continuous sound like wind could be perceived as the result of a gesture, and therefore be performable.
- To evaluate whether the prototype digital wind machine produced a sound that was perceivably similar to its acoustic counterpart in order to verify the programming approach discussed in Chapter 6.
- To understand how evaluation procedures from SID examining the link between sound and perceived action could be applied to the comparison of the acoustic and prototype digital wind machines.
- To evaluate whether the continuous sonic interaction afforded by the prototype digital wind machine was perceivably similar enough to its acoustic counterpart in operation in order to be used as a substitute in a future experiment.
- To discover more about the experience of performing an everyday sound like wind with a rotational gesture, and develop ideas for future work.

The main hypotheses developed for the experiment were as follows:

- Hypothesis 1 (H1): There is perceived similarity between the acoustic wind machine sound and the prototype digital wind machine sound.

- Hypothesis 2a (H2a): Participants can perceive different rotational speeds in a continuous wind-like sound produced through acoustic or digital means.

- Hypothesis 2b (H2b): The perception of rotational speed is equally accurate when the continuous wind-like sound is produced through acoustic or digital means

- Hypothesis 3 (H3): There is perceived similarity between the experience of performing with the acoustic wind machine and that of performing with the prototype digital wind machine.

The experiment was therefore designed to include three distinct steps:

- The first step aimed to evaluate the perceptual space of the acoustic and digital wind sounds through a listening test in which participants rated pairs of sounds for perceived similarity. This step addresses H1.

- The second step attempted to evaluate whether the speed of rotation of the acoustic and prototype digital wind machines could be communicated equally to a listener through their resulting wind-like sounds, by asking participants to rank these sounds in order of their perceived rotational speed. This step addresses H2a and H2b.

- The third step focused on the acoustic and prototype digital wind machines in performance. Participants first listened to a wind-like sound produced by acoustic or digital means, and were then invited to imitate what they had heard by playing the acoustic or prototype digital wind machine with the crank handle. They were then asked to evaluate their own performance, offer some descriptors for the wind sound they had played, and comment on what their experience was like when performing. This step addresses H3.

As outlined previously in Section 4.1.5, each step of the experiment makes use of a simultaneous recording setup to capture both the acoustic and digital wind-like sounds activated by the same gesture. This is used both to create stimuli for participants to listen to and to capture participants' own performances. The particular approach taken for each step of the experiment will be clearly stated for the reader throughout this chapter.

7.1.1 Pilot Study

The experimental procedure discussed in this chapter was first piloted with 6 participants. Results were gathered from each of the steps of the experiment, and participants were asked for their verbal feedback on their experience of the experimental procedure. This work helped to inform some small improvements to both the experiment design and the program in Max/MSP before proceeding to the full experiment. In the first step of the pilot study, participants rated pairs of acoustic, digital and natural wind sounds for similarity. The sounds were presented without being identified as *acoustic* or *digital* to ensure that participants were not influenced by these terms. While this part of the evaluation proceeded successfully, it was determined that some further sounds should be added to the group of stimuli for this step of the full experiment. This would help to contextualise the ratings comparing the wind-like sounds to each other in terms of their similarity to very different sounds, and also facilitate a comparison of the results with other larger studies of the factors involved in environmental sound perception (Gygi et al. 2007). For the second step of the experiment, pilot study participants ranked four audio clips of each wind machine in terms of their perceived rotational speed. This step proceeded successfully.

During the final step of the pilot study participants were asked to perform with both the acoustic and prototype digital wind machines in response to a video clip showing a windy scene, and then describe their experience of each wind-like sound. Again, each sound was presented to participants to interact with without being labelled as *acoustic* or *digital*. The results of this step showed that the visual stimulus seemed to be preoccupying participants with trying to sync their performance to the video clip, and this in turn coloured their perceptual evaluation of the sounds they were performing. It was determined that the full experiment should focus participants more closely on the wind-like sounds through the use of audio rather than video stimuli for this step. In discussion with pilot study participants, it was established that they found the digital wind-like sound less responsive to changes in the speed of their rotational gesture with the crank handle. In response to this feedback, some further subtle modulation of the parameters to the digital model in Max/MSP was added to improve the responsive whistling of the digital sound in performance. This final mapping between the rotary encoder data and the modulation of the synthesis parameters was presented in Chapter 6.

7.1.2 Participants and Statistical Power

In order to ensure that enough data was available for robust statistical testing, a total of 54 participants were included in the full experiment. As outlined in Chapter 4, the total of 54 participants ensures a statistical power of 0.8 when detecting a large effect size ($r = 0.5$) at $\alpha \leq 0.05$. Each step of the experiment would therefore have enough statistical power to detect a large effect. However, due to issues with data gathering during testing that would have affected the final analysis, some participants were excluded from each step of the experiment. The exact number of participants retained for each step is reported in the following sections. All manipulations of the experimental data, and the power and significance of each statistical test is also fully reported within this chapter.

This section has introduced the aims and hypotheses of the experimental evaluation, briefly summarised a pilot study carried out to confirm the robustness of the experiment design, and reported on the number of participants recruited for the full experiment. The next sections present a summary of the procedure and results for each step of the full experiment.

7.2 The Perceptual Space of Acoustic, Digital and Natural Winds

The first step of the experiment invited participants to listen to pairs of sounds and rate them in terms of their perceived similarity. As outlined in Chapter 4, studies using similarity ratings to compare everyday sounds have shown that there is an underlying perceptual structure that clusters sounds into distinct groups of continuous, harmonic and discrete impact sounds (Gygi et al. 2007). This procedure facilitates an examination of how a wind *effect* might compare perceptually with a natural wind sound, and whether the prototype digital wind machine produces a wind-like sound perceivably similar to its acoustic counterpart.

7.2.1 Stimuli

Field recordings of real-world wind were chosen from the BBC Sound Effects Library (BBC 1988) to serve as natural wind stimuli for this listening test. As per the recording setup previously described in Section 4.1.5, these field recordings were first loaded into a Pro Tools session. A microphone captured the sound of the acoustic wind machine via an audio interface into the same session, and the prototype digital wind machine sound was also captured through the same interface, routed as line level audio from the separate computer and audio interface running its program in Max/MSP. The acoustic wind machine and prototype digital wind machine were then simultaneously performed by the author to imitate the field recordings of natural wind sounds, and the results were recorded. During the recording process, the author listened to a split headphone feed, with the natural wind sounds routed to one ear and the acoustic and digital wind-like sound routed to the other ear. This produced a corpus of wind gestures, each with an acoustic, digital and natural component. Two distinct wind gestures were chosen as stimuli for this listening test. These were designated as *steady* (a consistent, sustained wind sound) (Figure 7.1), and *gusty* (a wind with exaggerated changes in intensity) (Figure 7.2), respectively.

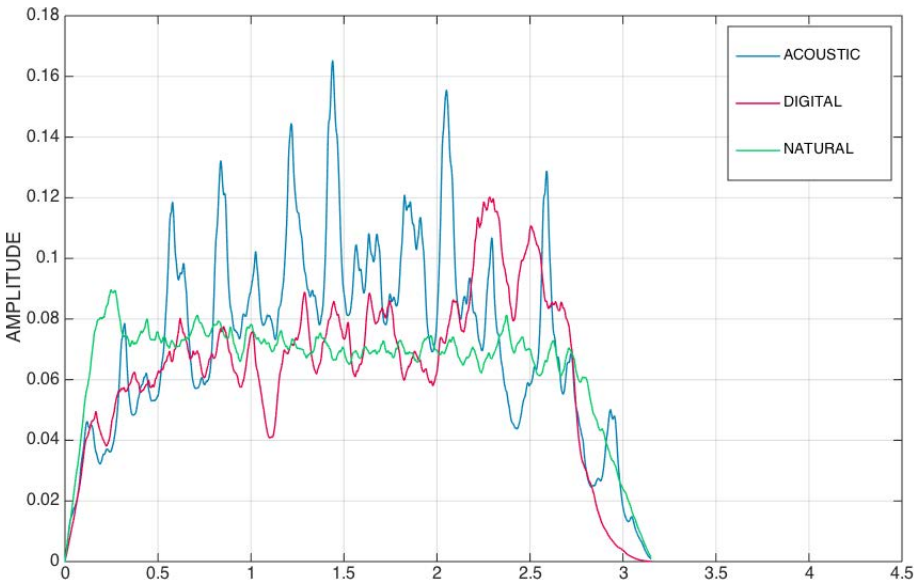

Figure 7.1: A comparison of amplitude envelopes for the steady Acoustic (blue), Digital (red) and Natural (green) wind sounds.

Previous research by (Gygi et al. 2007) found that everyday environmental sounds were clustered into three distinct categories. The cluster of continuous sounds included thunder, pouring water and a car engine. The cluster of harmonic sounds included pitched sounds like bells, animal vocalisations and honking horns. The final cluster of discrete impact sounds included footsteps, a typewriter and breaking glass.

To contextualise the six continuous wind sounds within this broader perceptual

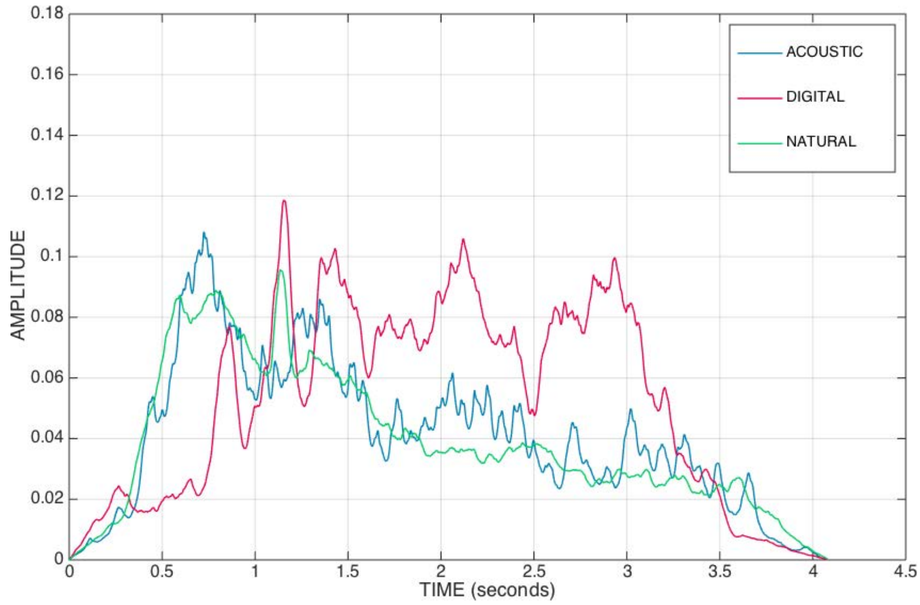

Figure 7.2: A comparison of amplitude envelopes for the gusty Acoustic (blue), Digital (red) and Natural (green) wind sounds.

structure of everyday sounds, a harmonic sound (a hand-operated horn) (BBC 1986) and a discrete impact sound (a piece of wood dropping onto a surface) (Edge 1991) were added to the corpus for this listening step. Although the horn and wood sound are from distinct clusters of everyday sounds according to Gygi et al. (2007), they share a similarly repetitive amplitude envelope, in order to help establish the importance of temporal patterning in the similarity ratings of the sounds. The stimuli for this listening test then consisted of a total of eight distinct sounds. To facilitate the creation of a matrix from the final similarity ratings, each sound was compared to itself and to each other sound, giving a total of 36 pairs of sounds for participants to rate as part of this task.

7.2.2 Apparatus

The experiment took place in an acoustically treated room at the Department of Theatre, Film and Television at the University of York. A MacBook Pro running the Python-based Open Sesame experiment platform (Mathôt et al. 2012) presented the audio stimuli and collected questionnaire data from participants. The 24bit/48KHz audio clips were played back through an RME Fireface 400 audio interface, and participants listened through a closed-back pair of Sennheiser HD280 Pro headphones.

7.2.3 Subjects

Data from a total of 51 participants was retained from this step of the experiment for the purposes of analysis. All of these participants reported normal hearing. 33 participants identified themselves as female, and 17 identified themselves as male. 40 participants identified themselves as 18-24 years of age, 9 as 25-34, and 2 as 45-54. 38 of the participants reported that they had experience of playing a musical instrument, with 16 designating themselves as beginners, 13 intermediate, and 9 with advanced ability. All participants were paid for their participation.

7.2.4 Procedure

Participants were asked to listen to each of the 36 sound pairs and rate the similarity of the two sounds to each other on a Likert scale from 1 (*not similar at all*) to 7 (*as similar as they can possibly be*), the same scale used by Gygi et al. (2007) in their study of environmental sounds. Participants were first presented with a practice step, which asked them to rate two sound pairs from the corpus, one comparing a sound to itself, and another pair of two distinct sounds, on the similarity scale. During the test step, the order of the presentation of each of the 36 sound pairs was randomised by the Open Sesame program. Each sound pair was presented only once, apart from the practice pairs, which were presented again during the main test step. The sounds were not labelled when presented to participants to avoid the use of any terms such as *acoustic*, *digital*, *continuous* or *impact*, which may have influenced their ratings. The full trial lasted for approximately 20 minutes.

7.2.5 Results and Analysis

The similarity ratings were scored according to their place on the scale with values from 1 to 7. The overall mean similarity score was 3.77, with a standard deviation (SD) of 0.42. The similarity ratings were not normally distributed. This was expected due to the high ratings given to the pairs that compared sounds to themselves. An examination of these scores showed that, while the median similarity of pairs comparing a sound to itself was consistently 7, participants did not always rate these sounds *as similar as they could possibly be* to each other. Both of the digital wind sounds achieved slightly lower mean similarity scores, but the horn sound achieved the lowest mean similarity score when compared to itself. The scores for these pairs are presented in Table 7.1 in ascending order

Sound compared to itself	**Mean similarity**	**Standard deviation (SD)**	**Median**
Horn	6.57	1.04	7
Gusty Digital Wind	6.82	0.43	7
Steady Digital Wind	6.84	0.61	7
Gusty Natural Wind	6.84	0.36	7
Steady Acoustic Wind	6.88	0.33	7
Gusty Acoustic Wind	6.90	0.3	7
Steady Natural Wind	6.90	0.27	7
Wood Drop	6.98	0.14	7

Table 7.1: Similarity ratings of sounds compared to themselves, in order of mean similarity.

of mean similarity.

To ensure that the next stage of the analysis of the similarity ratings would produce some robust conclusions for this small group of sound stimuli, participants' individual ratings were adjusted to account for intersubject differences in mean ratings. This process produces a mean similarity rating for each pair of sounds based on an *agreed* similarity rating scale across the participants, rather than trying to incorporate all of the individual differences in how the decisions on similarity were made. An adjustment factor was calculated for each participant by subtracting the mean similarity of their individual ratings for each pair of sounds from the overall mean similarity score (3.77). Each participant's ratings were then adjusted by adding their adjustment factor to the rating they had assigned to each pair (Field et al. 2012, p.364). The scores for each same-sound pairs were also adjusted to the maximum value of 7. A similarity matrix was then created from the resulting mean similarity scores for each sound pair (Table 7.2).

7.2.5.1 Multidimensional Scaling Analysis

A multidimensional scaling (MDS) analysis was performed to express the similarity matrix data as a series of points (one for each sound in the corpus of stimuli) in a two dimensional space. The similarity matrix was first transformed with the *dist()* function in R into a distance matrix, which expressed the mean similarities between the sounds as Euclidean

	Steady N	Steady A	Steady D	Gusty N	Gusty A	Gusty D	Horn	Wood
Steady N	*7.0*							
Steady A	5.06	*7.0*						
Steady D	3.98	3.78	*7.0*					
Gusty N	5.06	4.74	3.51	*7.0*				
Gusty A	4.19	5.19	4.04	5.0	*7.0*			
Gusty D	3.21	3.45	5.37	3.51	4.27	*7.0*		
Horn	1.39	1.39	1.74	1.33	1.41	1.15	*7.0*	
Wood	1.09	1.11	1.19	1.15	1.13	1.17	1.35	*7.0*

Table 7.2: Matrix of mean similarity scores produced from participants' similarity ratings. Wind sounds are designated by gesture descriptor and then by source: natural (N), acoustic (A), digital (D).

distances (Table 7.3). A Euclidean distance is the distance between two points in the geometry of Euclidean space. It is calculated as

$$\sqrt{\sum_{i=1}^{N}(A_i - B_i)^2}$$

where N is the number of dimensions in the space and A and B are the two points (Giordano et al. 2013, p.186).

MDS analysis can take several dimensions into account, but requires that the number of dimensions be specified before proceeding. Each dimension represents a factor influencing the perceived similarity ratings, which in this case is an acoustical quality of the sound stimuli. To discover how many dimensions were required for the MDS procedure, the factoextra package in R (Kassambara & Mundt 2016) was used to calculate eigenvalues[1] from the distance matrix and express how much variance each factor (or dimension) might account for as a percentage of the total variance (Figure 7.3). The resulting scree plot showed that most of the variance in the distance matrix (95.3% in total) could be accounted for with the first three factors, meaning that three dimensions would be appropriate for

[1] Eigenvalues are part of the mathematics of matrices and linear algebra. In this case, they express the underlying characteristics of the distance matrix, or how much influence each factor may have had on the similarity ratings that produced the distances.

	Steady N	**Steady A**	**Steady D**	**Gusty N**	**Gusty A**	**Gusty D**	**Horn**
Steady A	2.95						
Steady D	5.21	5.35					
Gusty N	2.93	3.21	5.57				
Gusty A	4.12	2.85	4.81	3.13			
Gusty D	5.98	5.79	2.53	5.79	4.79		
Horn	10.36	10.58	9.86	10.52	10.65	10.12	
Wood	11.05	11.24	10.72	11.10	11.30	10.53	8.03

Table 7.3: Matrix of Euclidean distances between the sound stimuli produced by transforming the matrix of mean similarity ratings with the *dist()* function in R. Wind sounds are designated by gesture descriptor and then by source: natural (N), acoustic (A), digital (D).

the final MDS analysis. This result is confirmed by Gygi et al.'s 2007 study, which proposed a three dimensional solution as most appropriate for evaluating a wider variety of environmental sounds.

An MDS analysis using the simplest CLASCAL method was then performed on the distance matrix using three dimensions and the cmdscale function from the core statistics package in R. This produced coordinates for the sound stimuli along each of the three dimensions. MDS solutions are rotationally invariant, and so can be rotated or reflected while retaining the distances between the points for the purpose of analysis (Bartholomew et al. 2008, p.58). Dimensions 1 and 2 were rotated before producing the final two-dimensional plots to position the points in a manner consistent with the MDS analysis produced in Gygi et al. (2007), which positioned continuous sounds in a distinct cluster in the centre and towards the top of the y-axis (Figure 7.4).

The resulting plot of Dimensions 1 and 2 cluster the sound stimuli into three distinct groups - continuous sounds (winds), harmonic sounds (horn) and discrete impacts (wood). Within the continuous sound cluster, the digital winds sounds are closer to each other than to the acoustic and natural winds. The acoustic wind sounds are very close to the natural wind sounds. Dimension 1 was also plotted against Dimension 3 (Figure 7.5).

The plot of Dimension 1 against Dimension 3 shows the horn and wood sounds further apart from each other, with the continuous wind sounds again clustered together. The acoustic wind sound is closer to the natural wind sound here, particularly its steady

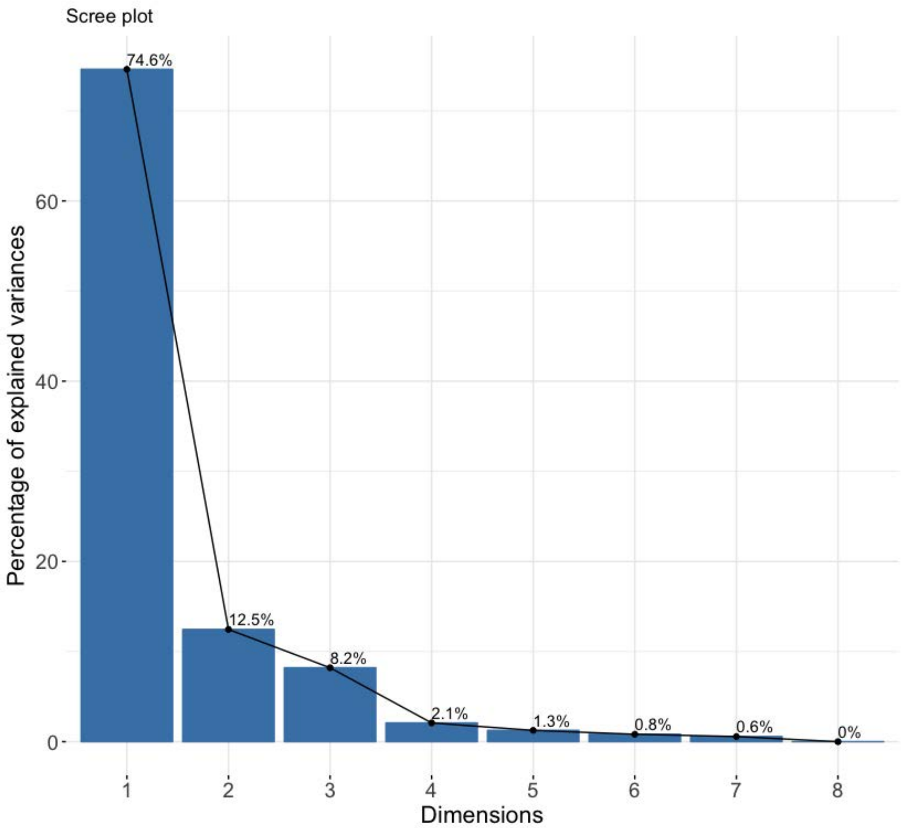

Figure 7.3: An analysis of variance in the distance matrix produced by the factoextra package in R, showing that 95.3% of the variance in the data can be accounted for by the first three dimensions.

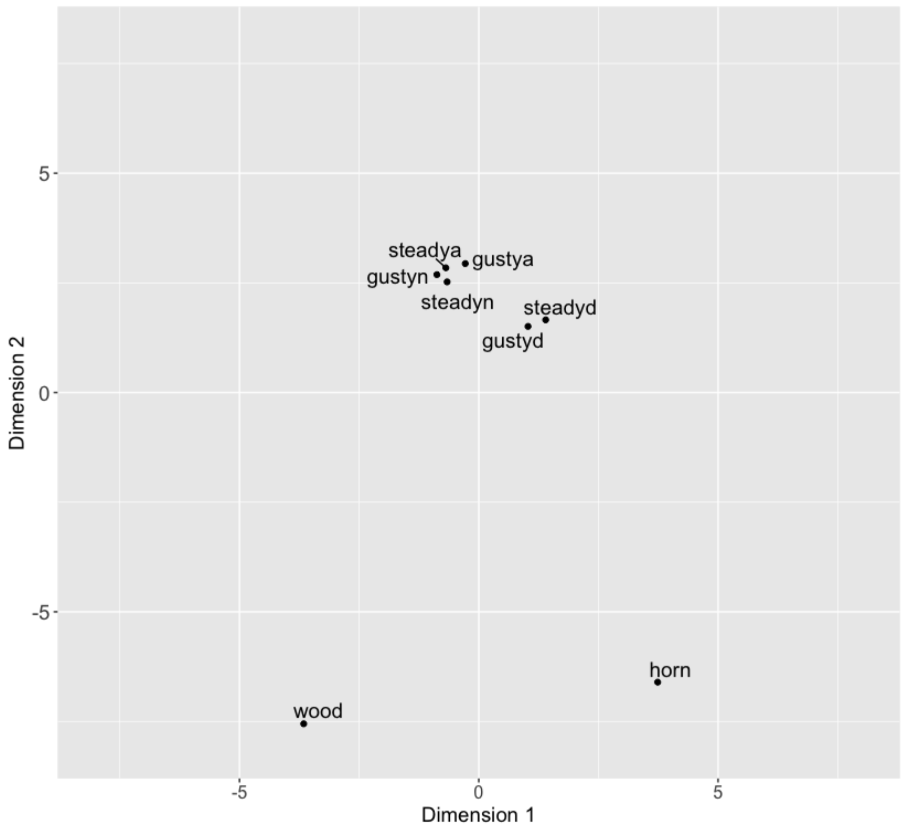

Figure 7.4: Two-dimensional plot of the MDS solution for Dimension 1 against Dimension 2.

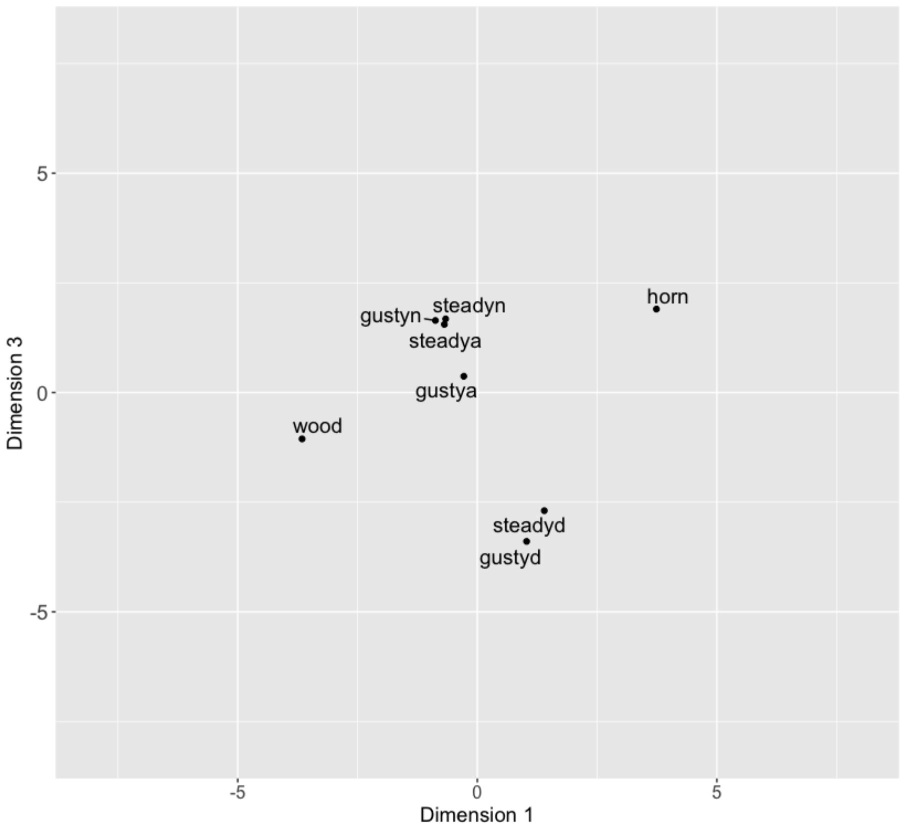

Figure 7.5: Two-dimensional plot of the MDS solution for Dimension 1 against Dimension 3.

gesture. The digital wind sounds are much further removed from the other wind sounds in this configuration. A hierarchical cluster dendrogram (Figure 7.6) was also produced from the distance matrix to make the groupings more explicit. This shows that the horn and wood sounds are very distinct from the cluster of continuous sounds. Within the continuous sounds, each source of the wind (natural, acoustic or digital) is responsible for the grouping rather than the gestures themselves (steady or gusty), but the acoustic and natural wind sounds are grouped much more closely together while the digital wind sounds are in their own distinct part of the continuous cluster.

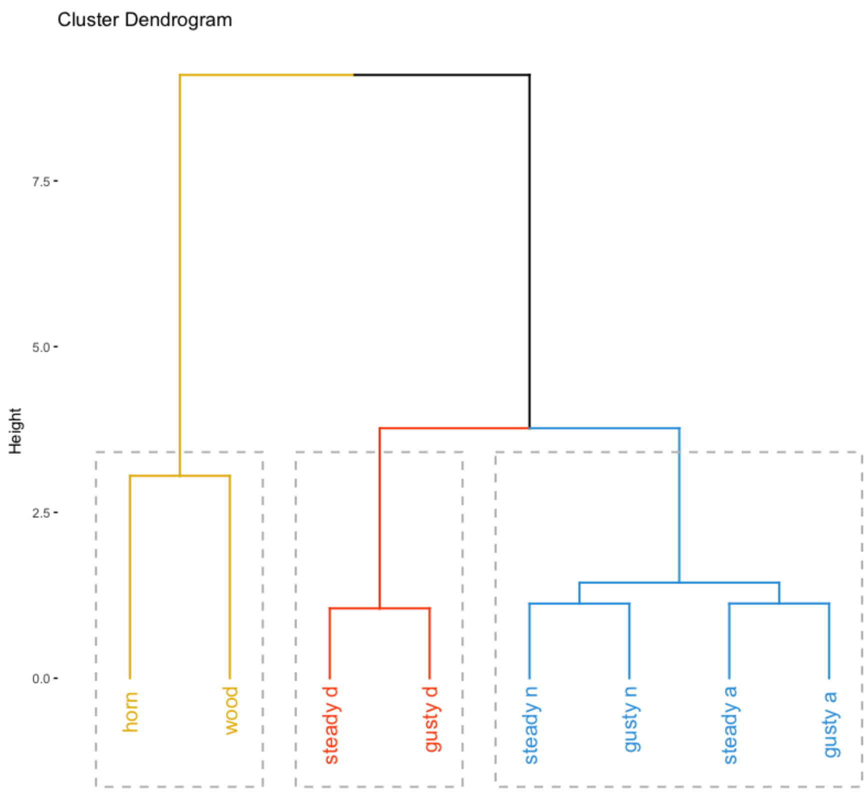

Figure 7.6: Hierarchical clustering of the distance matrix data.

7.2.5.2 Acoustic Factors in Similarity Ratings

In order to discover which acoustical qualities of the sound stimuli could be influencing their perceived similarity to each other, their acoustic features were correlated with their coordinates along the dimensions of the MDS solution (their x and y values in the MDS plots). Due to the time restrictions of this study there are only a limited number of observations available for this analysis (8 points along each dimension, one for each sound),

Acoustic Feature	**Dimension 1 (74.6%)**	**Dimension 2 (12.5%)**	**Dimension 3 (8.2%)**
Brightness	-0.71***	*0.06*	0.65*
Inharmonicity	-0.90***	0.67*	*0.44*
Centroid	-0.86***	0.64*	0.64*
Spread	-0.77**	0.66*	0.68*
Skewness	0.76**	-0.80**	-0.71**
Event Density	*0*	0.62*	*-0.22*

Table 7.4: Table of correlations (Spearman's rho) between the three dimensions of the MDS analysis and the acoustic features of the sound stimuli [$p \leq .001 = ***$, $p \leq .01 = **$, $p \leq .05 = *$]

and so these results should be considered with caution. However, given that the MDS analysis has produced results that confirm the clusters observed in the larger study by Gygi et al. (2007), it is reasonable to produce some observations that could be confirmed with a further in-depth study to evaluate a larger group of sound stimuli. To examine the potential acoustic factors in the results of this listening test, the sound stimuli were acoustically evaluated using the MIR Toolbox in MATLAB (Lartillot & Toiviainen 2007) to produce some numerical measures of their acoustic features. This data was then correlated with the coordinates extracted from each dimension using a Spearman's rho coefficient, as there were a small number of coordinates and they were not normally distributed. The results are shown in Table 7.4.

The correlations show that Dimension 1 seems to be associated highly with spectral measurements of the sounds, including statistical measures of the spectrum, which reflects the results from the more extensive study in Gygi et al. (2007). This suggests that, according to the scree plot presented in Fig. 3, the frequency content or timbral qualities of the sound stimuli were responsible for 74.6% of the variance in the ratings of the stimuli in this study. In particular, the noisiness of the sounds (inharmonicity) and the energy of their spectra (centroid) seem to have been particularly important.

Dimension 2, accounting for 12.5% of the variance in the similarity ratings, shares an association with statistical measures of the spectrum, in particular whether the frequency content of the sounds was clustered towards the low or high frequencies (skewness). It is also singularly associated with the frequency of onsets in the amplitude envelope of the sounds (event density). These associations also reflect the results of the more extensive study by Gygi et al. (2007).

Dimension 3 is most highly associated with the spectral skewness of the stimuli, and is not associated with either noisiness or the amplitude envelope of the sounds. It accounts for a further 8.2% of the variance in the similarity ratings for this study. These associations with spectral features of the sounds reflect the results in the more extensive study by Gygi et al. (2007).

7.2.6 Summary of Findings

This step of the experiment aimed to establish whether there was perceived similarity between the wind-like sounds produced by the acoustic wind machine and the prototype digital wind machine. It is clear from the results that participants clustered all of the wind-like sounds together in a group of continuous sounds, but rated the acoustic wind much more similar to its natural than its digital counterpart. As such, the results did not allow the null hypothesis to be rejected. This suggests that further work is needed to adjust the digital wind model in Max/MSP to bring its resulting sound perceptually closer to that of its acoustic counterpart.

However, the similarity perceived between the acoustic wind machine and the natural wind sounds it imitated confirms the descriptions and claims of the efficacy of this effect method made by historical practitioners that we previously explored in Chapter 3. The MDS analysis procedure has been extended here to encompass environmental sounds along with acoustic and digital effects imitating those sounds. Exploring the potential of this approach further might yield some interesting understandings of the importance of particular acoustic features across types of sound sources (real-world environmental sounds and effects imitating those sounds) when they are rated in terms of similarity.

The limited number of observations from this study required individual differences in participants' mean ratings to be adjusted for a robust analysis, meaning that the final MDS result cannot take any individual differences in the importance participants gave to each factor (acoustic feature) when rating for similarity. A dedicated listening experiment with a fuller corpus of sounds would allow for the use of a more complex MDS algorithm such as INDSCAL, which adds weights for individual differences in responses, or ALSCAL, which requires two separate matrices of similarity data (and therefore two separate listening tests) but will produce a more comprehensive result. The results of this listening step do confirm those found in the more extensive study by Gygi et al. (2007) however, which gives more weight to the correlations between the acoustical features of the sound stimuli and their positions along the dimensions of the MDS solution. A fuller corpus of sounds in a future listening test would allow for further statistical testing to establish which of the acoustic features of the stimuli predicted their ordering along each dimension of the MDS result.

Having reported on the first listening step, the next section outlines the procedure and results of the next part of the experiment, where participants ranked the acoustic and digital wind-like sounds in terms of their perceived rotational speed.

7.3 Perception of the Rotational Speed of a Wind-Like Sound

The second step of the experiment attempted to evaluate both whether participants could perceive different levels of rotational speed (measured in rotations per minute, or RPM) from the resulting sounds of the two wind machines, and also whether those differences in rotational speed could be communicated equally by both the acoustic and digital wind-like sounds. As discussed in Chapter 4, the procedure was devised from the listening experiment outlined in Lemaitre et al. (2009), where the Spinotron ratchet sound model was evaluated through participants estimations of the speed of the ratchet from a sound produced by a specific rotational speed. The study confirmed that the speed of rotation could be perceived from the ratchet sound. This step of the experiment similarly facilitates an examination of whether participants can connect a wind-like sound to a rotational gesture, potentially by imagining a virtual mechanism as the coupling between them.

7.3.1 Stimuli

As outlined in Chapter 4, the acoustic and digital wind-like sounds require manual operation, and cannot be configured with specific real-time parameters or quantisation (as with the Spinotron ratchet model) to produce sound stimuli. Instead, the two wind wind machines were performed at different rotational speeds and their resulting sounds simultaneously recorded for this task. This was facilitated with the recording setup previously described in Section 4.1.5. A microphone captured the sound of the acoustic wind machine via an audio interface into a Pro Tools session, and the prototype digital wind machine sound was also captured through the same interface, routed as line level audio from the separate computer and audio interface running its program in Max/MSP. The acoustic wind machine and prototype digital wind machine were then simultaneously performed by the author to produce a corpus of wind-like sounds produced by specific rotational speeds. During the recording process, the author listened to a split headphone feed, with the acoustic and digital wind-like sounds routed separately to each ear. The program in Max/MSP facilitated monitoring of the real-time rotational speed of the crank handle in RPM during these performances.

The relatively slow and restricted range of RPM afforded by the acoustic wind machine interface produced a limited corpus of sounds at unequal speed intervals. As such, the task was designed to require participants to rank the sounds they heard in order of speed, rather than estimate their speed on a continuous scale. A corpus of eight audio clips of the acoustic and digital wind-like sounds resulting from rotational speeds of 0.5RPM, 1RPM, 2RPM and 5RPM was produced. The sounds were produced by partial rotations of the crank handle in order to ensure a consistency of speed for the full duration of each clip. Following the procedure in Lemaitre et al. (2009), all of the audio clips were of a similar duration, this time a maximum of two seconds in length.

7.3.2 Apparatus

The experiment took place in an acoustically treated room at the Department of Theatre, Film and Television at the University of York. A MacBook Pro running the Python-based Open Sesame experiment platform (Mathôt et al. 2012) presented the audio stimuli and collected ratings data from participants. An RME Fireface 400 audio interface played the 24bit/48KHz audio clips, and participants listened through a closed-back pair of Sennheiser HD280 Pro headphones.

7.3.3 Subjects

All of the participants who had previous participated in the similarity ratings step also participated in this step of the experiment. Data from two of the participants had to be excluded from this listening test due to incomplete data however, and so data from a further three participants was also excluded to ensure that equal numbers of the experimental conditions were used for the final analysis. This brought the total to 46 participants for this listening test. Of these 46, 31 identified as female and 15 identified as male. 36 participants identified themselves as 18-24 years of age, 8 as 25-34 years of age, and 2 as 45-54 years of age. 33 of the participants reported that they had experience playing a musical instrument, with 15 of these designating themselves as beginners, 10 as intermediate and 8 as advanced. All participants reported normal hearing and all were paid for their participation.

7.3.4 Procedure

The acoustic and digital wind-like sounds were presented to participants separately to be ranked in order of their perceived rotational speed. This listening step was based on a repeated measures design, with all participants ranking both the acoustic and the digital wind-like sounds. The order of presentation of the stimuli was randomised, with 23 participants rating the group of acoustic sounds first and 23 rating the group of digital sounds first in order to minimise order effects. The order of the four sounds for each task were randomised in Open Sesame and labelled as *sound 1, sound 2, sound 3*, and *sound 4*. Labels such as *acoustic* or *digital* were not used. Each sound could be triggered with an onscreen button.

Participants were told that each of the four sounds had been made by a handle being turned at different speeds. They were asked to listen to each sound, play them as many times as they liked, and then rank the sounds in terms of the speed of the turning handle on a scale from 1(*slowest*) to 4(*fastest*). Participants were only given one opportunity to rate each kind of sound for perceived rotational speed. Once the participants had finished rating the first group of sounds (e.g. acoustic), they then proceeded to the second group of sounds (e.g. digital) and rated them.

7.3.5 Results and Analysis

The speed ratings were scored with values from 1 to 4 according to their ranking on the scale. The speed ratings for both the acoustic and digital sounds were not normally distributed. An examination of the mean and median speed ratings showed that participants did not consistently rank the sounds in order of the rotational speed that produced them. The mean ratings for the acoustic wind machine sounds showed more variation than those for the digital wind machine sounds, which were grouped closely together (Table 7.5).

Plotting the actual rotational speeds of the corpus of sounds against the speed rankings by participants shows that, despite the different rotational speeds, each sounds rankings were spread over a range, and not focused on one particular rank (Figure 7.7).

RPM	Expected Rank	Sound	Mean	Standard Deviation (SD)	Median
0.5	1	*Acoustic*	1.83	1.0	1
		Digital	2.0	1.19	1
1	2	*Acoustic*	3.17	0.9	3
		Digital	2.76	1.02	3
2	3	*Acoustic*	2.33	1.09	2
		Digital	2.76	1.14	3
5	4	*Acoustic*	2.65	0.99	3
		Digital	2.61	0.95	3

Table 7.5: A comparison of descriptive statistics for the perceived speed of the sounds as ranked by participants.

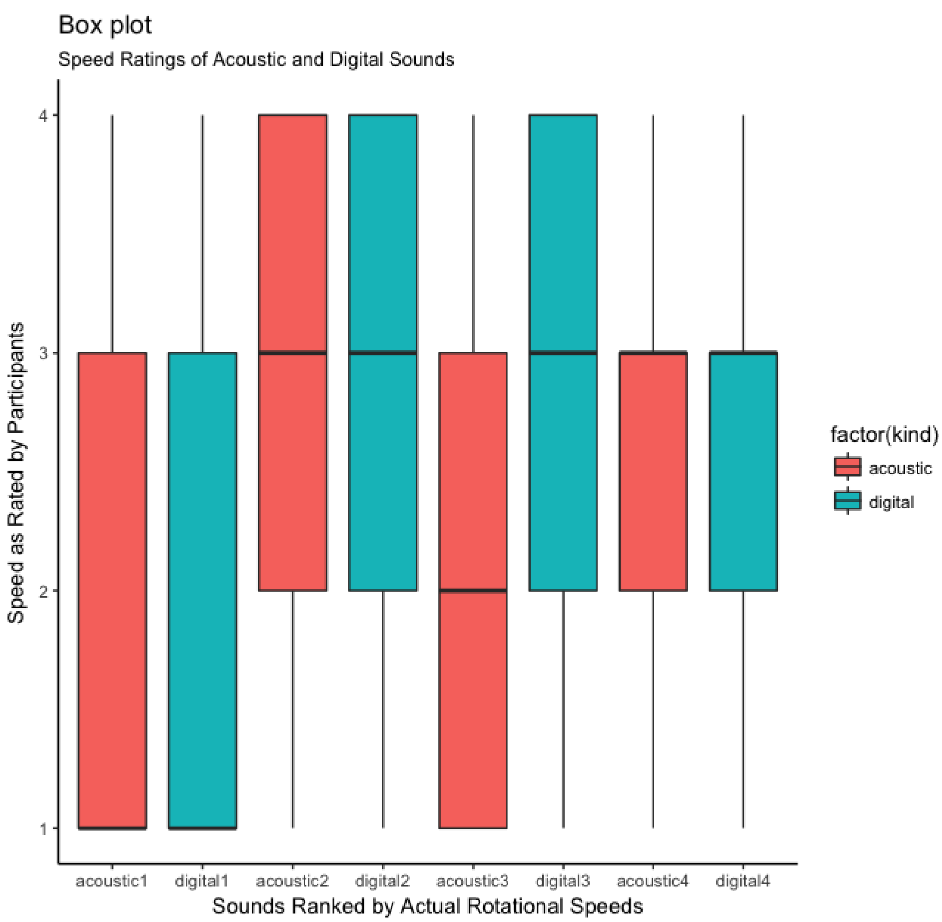

Figure 7.7: Boxplot of the speed rankings of the acoustic and digital wind sounds for each actual speed (0.5RPM, 1RPM, 2RPM, 5RPM).

To compare the rankings given to each sound, a Friedman rank sum test was performed to establish whether they differed significantly from each other. This test showed that

Test: Friedman's ANOVA	**Significance**	**Effect Size**
2(7) = 44.194, Z = -5.2	p <0.01	1.16 (large) statistical power = 0.8
Post-hoc comparisons: Wilcoxon signed rank test	**Significance**	**Effect Size**
Acoustic 1 (0.5RPM) Acoustic 2 (1RPM): Z = -4.13	p <0.01	-0.71 (large) statistical power = 0.8
Acoustic 1 (0.5RPM) Acoustic 4 (5 RPM): Z = -1.95	p = 0.05	-0.46 (medium) statistical power = 0.8
Acoustic 1 (0.5RPM) Digital 2 (1RPM): Z = -3.24	p <0.01	-0.60 (large) statistical power = 0.8
Acoustic 2 (1RPM) Digital 1 (0.5RPM): Z = -3.89	p <0.01	-0.68 (large) statistical power = 0.8

Table 7.6: Results of the statistical testing to compare how the wind-like sounds were ranked in order of speed.

there were statistically significant differences in the way that participants had ranked the various sounds. Post-hoc testing using a Wilcoxon signed-rank test with a Bonferroni correction applied[2] showed that only some of the rankings differed significantly from each other (Table 7.6). No significant difference was observed when comparing the rankings of each of the digital wind-like sounds to each other.

7.3.6 Summary of Findings

This step of the experiment aimed to establish whether participants could perceive different speeds of rotation in a continuous wind-like sound produced through acoustic or digital means, and whether that difference in rotational speed could be communicated equally by both sounds. The results showed that participants did not consistently rank the sounds they heard in the order of the rotational speed of the crank handle that produced them, and so did not allow the first null hypothesis to be rejected (H2a0).

However, statistical testing of the speed rankings did indicate that there were some statistically significant differences in how participants ranked the sounds for speed. These differences were observed when comparing the acoustic wind-like sounds produced with rotational gestures at 0.5RPM and 1RPM, and at 0.5RPM and 5RPM. As the effect size observed when comparing 0.5RPM and 5RPM is medium, it would be interesting to confirm this result in the future with a larger number of participants. Statistically

[2] The Bonferroni correction adjusts the p-value of a test according to the number of comparisons made in order to ensure that multiple comparisons do not produce errors in the results of the analysis.

significant differences were also observed between the rankings for the acoustic wind-like sound and the digital wind-like sound when produced with rotational gestures at 0.5RPM and 1RPM, and 1RPM and 0.5RPM respectively. No statistically significant differences were observed when comparing the rankings between the various digital wind-like sounds. This suggests that the two wind machines are not able to communicate their speed equally, and so the results did not allow the second null hypothesis to be rejected (H2b0).

The range of speed rankings given to each actual RPM value (Figure 7.7) may have been due to the small range of rotational speeds that produced the stimuli for evaluation. Participants may have been presented with relatively small differences between the sounds such that they could not be ranked with confidence. The limited number of sound stimuli may also have been unduly affected by the floor/ceiling effect in experimental data gathering, where scores tend to bunch together at the lower and higher end of a scale. This could be remedied in a future test by increasing the number of sound stimuli available for the experiment. A larger corpus of sound stimuli would also provide more observations for further statistical analysis to see if the actual RPM that produced a wind-like sound might predict how participants rank it for speed. If the wooden acoustic wind machine interface still restricts the available range of speeds to the ones produced for this listening step, a repeated measures experiment design, which allows participants to rank the sounds more than once, could help to validate the speed rankings instead.

The results of this step of the experiment do not allow clear conclusions to be drawn as to whether participants could understand the metaphor of a turning handle creating the sounds that they heard. If this metaphor had been very simple to understand however, the rankings given to the sounds may have been more precisely aligned with the RPM that produced them. The short duration of the audio clips, produced with only partial rotational movements, may have contributed to confusion around how the stated mechanism (the crank handle) produced the sound. In the evaluation of Spinotron, researchers showed that perceptions of what kind of material caused the sound might influence this understanding (Lemaitre et al. 2009, p.982). An extra experimental step to establish what participants might understand as the material causing either the acoustic or digital wind-like sound could help to reveal more about their comprehension of the link between a rotating crank handle and a wind-like sound in a future speed rankings task.

This section has reported on the second listening step of the experiment. The next section details the final step, where participants performed with both the acoustic and prototype digital wind machines and reported on their experiences.

7.4 Continuous Interaction with a Wind-Like Sound

This step of the experiment attempted to compare how participants experienced a continuous sonic interaction with the acoustic and prototype digital wind machines in performance.

7.4.1 Stimuli

Two simple rotational gestures, a slow single rotation and a double rotation at moderate speed, were simultaneously recorded for both the acoustic and digital wind-like sounds to serve as stimuli for this step of the experiment. These were produced using the recording setup introduced in Section 4.1.5. The acoustic wind machine was captured with a microphone and audio interface to facilitate recording in Pro Tools. The prototype digital wind machine sound was also captured through the same interface, routed as line level audio from the separate computer and audio interface running its program in Max/MSP. During the recording process, the author listened to a split headphone feed, with the acoustic and digital wind-like sounds routed individually to each ear. The acoustic wind machine and prototype digital wind machine were then simultaneously performed by the author to produce a corpus of wind-like sounds produced by the specific rotational gestures. Another natural wind sound with an exaggerated and shifting amplitude envelope was chosen from the BBC Sound Effects Library (BBC 1988) for use in the practice step. This stimulus was chosen to encourage participants to perform an exaggerated gesture with each wind machine when they first encountered them, which would give them more of a sense of the sonic range of each wind-like sound when activated by the crank handle in performance.

7.4.2 Apparatus

This step of the experiment took place in the same acoustically treated room at the Department of Theatre, Film and Television at the University of York that hosted the listening tests. A MacBook Pro running the python-based Open Sesame experiment platform (Mathôt et al. 2012) presented the questions and collected data from participants. A second MacBook Pro running the prototype digital wind machine program in Max/MSP with a Mackie Onyx Satellite audio interface fed line level audio to an RME Fireface 400 audio interface connected to an iMac computer running Pro Tools at 24bit/48KHz. An AKGC414 microphone to capture the acoustic wind machine audio was also connected to the RME Fireface 400. Both the Max/MSP patch on the laptop and the Pro Tools session on the iMac were obscured from the participants to ensure they did not receive any additional visual feedback during the session.

The audio stimuli and live audio of participants' performances was delivered to them via Pro Tools through a closed-back pair of Sennheiser HD280 Pro headphones. Participants' performances in response to the sound stimuli were recorded into the same Pro Tools session. The acoustic wind machine was obscured, apart from its crank handle, behind a cardboard screen to ensure that it provided no visual feedback to participants during performance (Figure 7.8). Participants stood facing the crank handle when they performed the sounds. The wind-like sounds were panned to the centre of the stereo field in Pro Tools so that participants heard the sound coming from directly in front of them when holding the crank handle during play.

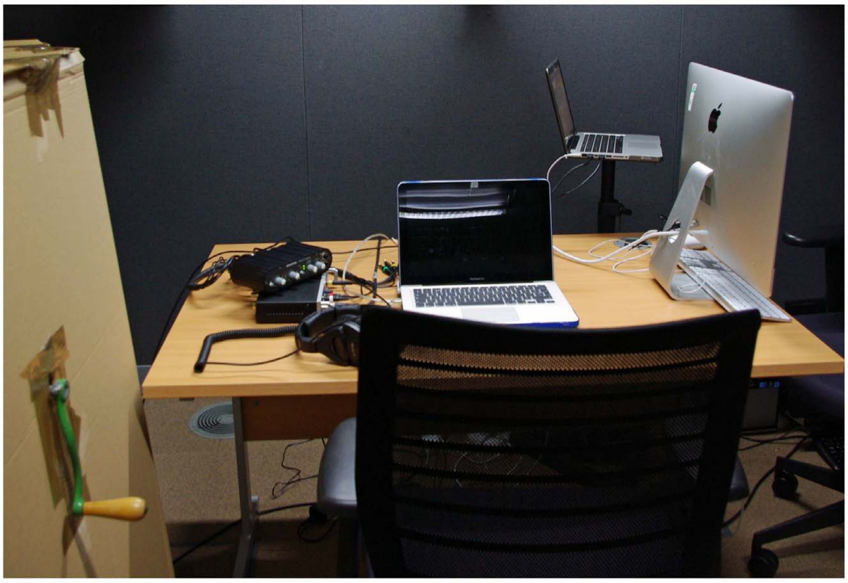

Figure 7.8: Equipment setup for the performance step of the experiment.

As participants would have no visual feedback during their performances, they would be unaware that the acoustic wind machine's handle had a rest position as seen in Figure 7.8, or that the handle could only be turned clockwise to activate the sound. There was therefore the potential that they would turn the crank handle anticlockwise at points even if advised otherwise, and as a result pull the cloth off the acoustic wind machine, which would require it to be readjusted before continuing. To ensure that this would not disrupt the experimental procedure, the loose side of the acoustic wind machine's cloth was tied down to the A-frame with a loose piece of wool. This ensured that the loose side of the cloth still felt less resistive to the cylinders movement during rotation than the tight side, but an exaggerated anticlockwise movement would not be able to completely remove the cloth (Figure 7.9).

Figure 7.9: Tying down the loose side of the acoustic wind machine's cloth.

7.4.3 Participants

All of the participants who had taken part in both listening steps took part in this step of the experiment, but data from 3 of those participants had to be excluded from the final analysis due to technical issues during the experiment. Of the 48 participants included in the final analysis, 32 identified themselves as female and 16 as male. 38 participants designated themselves as 18-24, 8 as 25-34, and 2 as 45-54 years old. 13 participants said they did not have experience of playing a musical instrument, 15 played a musical instrument at beginner level, 12 at intermediate level and 8 at advanced level. All participants reported normal hearing and were paid for their participation.

First System Used In Performance	**Subgroup**	**First Stimuli Presented**
Acoustic Wind Machine	A	acoustic wind-like sounds
	B	digital wind-like sounds
Prototype Digital Wind Machine	A	digital wind-like sounds
	B	acoustic wind-like sounds

Table 7.7: Outline of the different orders of system and stimuli for the performance step.

7.4.4 Procedure

This step of the experiment was based on a repeated measures design, with all participants performing with both the acoustic and prototype digital wind machines in response to all of the stimuli. To avoid order effects, the order of presentation of the acoustic and digital wind machines was randomised. The order of presentation of the sound stimuli was also randomised. This created four groups of twelve participants. Each group had its own order of system performed and stimuli presented (Table 7.7).

For this step of the experiment, participants were presented with the crank handle and advised that they would be able to perform a wind sound by turning it clockwise. They were told that there would be two wind sounds to perform with during this step of the experiment, and that they would get to perform with both of these sounds. The sounds were not identified as *acoustic* or *digital*. Participants were then asked to listen to a wind sound from the group of stimuli played through their headphones, and then try to imitate what they had heard directly afterwards by turning the crank handle. There was a practice step, and then a test step.

During the practice step, participants listened to the natural wind sound, and played one of the wind machines to imitate it. They then answered all of the questions that would be presented in the test step. This was repeated for both the acoustic and prototype digital wind machines. For each changeover between the sounds being controlled, participants were asked to first rotate the handle once only to introduce the sound they would play, and to ensure that they could hear the new sound in their headphones. The experiment then proceeded to the test step, where participants imitated all of the gestures in the corpus of stimuli (whether acoustic or digital) by playing one of the wind machines, and then answered all of the questions about the wind machine they had just performed with. The procedure then repeated with the other wind machine.

The questions presented to participants first asked them to rate how similar their

Test: Kruskal-Wallis	Significance	Effect Size
Acoustic similarity rating by group: $H\ (3) = 6.36$	$p > 0.05$ (not significant)	-0.12 (small) statistical power = 0.8
Digital similarity rating by group: $H\ (3) = 3.04$	$p > 0.05$ (not significant)	0 (none) statistical power = 0.8

Table 7.8: Results of the statistical testing to confirm no order effects on the similarity ratings given to participants' performances with the acoustic and prototype digital wind machines.

performances were to the stimuli they had heard on a scale of 1(*not similar at all*) to 7(*as similar as they can possibly be*). Participants were then asked to rate how far they agreed with the statement "This wind sound is easy to play" on a scale of 1(*strongly disagree*) to 7(*strongly agree*). A list of possible descriptors for the wind sound they had been performing was presented, and participants were asked to describe the wind sound they had played by selecting from these. There was also a space to add a descriptor of their own to this list. Finally, participants were given the opportunity to provide some free description of their experience of playing the sound.

7.4.5 Results and Analysis

7.4.5.1 Perceived Similarity of Performances to Stimuli

Participants' ratings of perceived similarity between the wind stimuli they had listened to and the wind sound they had performed to imitate them were scored according to their place on the scale with values from 1 to 7. A Kruskal-Wallis test was then performed on the similarity ratings given by the participants across each of the groups according to the order of performance system and the order of presentation of stimuli. This test confirmed that there was no statistically significant difference between the ratings given according to the experimental condition, confirming that no order effects had influenced the ratings (Table 7.8).

A summary of the similarity ratings showed that, while there was a range of scores for each of the interactions, the acoustic wind machine performances had a higher mean rating for similarity to the stimuli presented than the prototype digital wind machine performances (Table 7.9).

The similarity ratings were then collated and displayed in bar graphs according to the wind machine being compared to the stimuli. This showed that the acoustic wind machine performances received similarity ratings across the full scale, but participants did perceive

Sound Played	**Mean**	**Standard Deviation (SD)**	**Median**
Acoustic	4.88	1.66	5.5
Digital	2.77	1.51	2.5

Table 7.9: Summary of the similarity ratings for the performed acoustic and the prototype digital wind machine sounds.

their performances with this system to be very similar to the stimuli they had listened to (Figure 7.10).

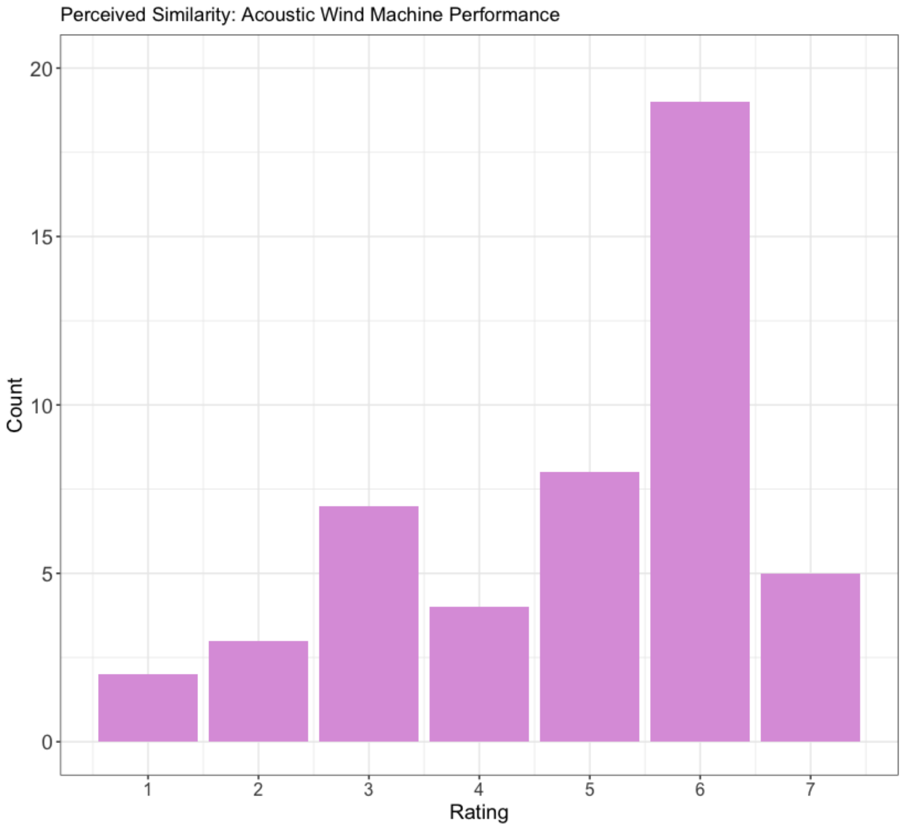

Figure 7.10: Similarity ratings for the acoustic wind machine performances when compared to the sound stimuli they imitated.

By contrast, the similarity ratings for the prototype digital wind machine performances were spread more evenly across the scale, but were weighted towards the low end (Figure 7.11).

A Wilcoxon signed rank test was performed to compare how participants rated each of their wind machine performances in terms of their similarity to all of the stimuli they tried to imitate. This test confirmed that there was a statistically significant difference between the similarity ratings given to the acoustic wind machine performances and

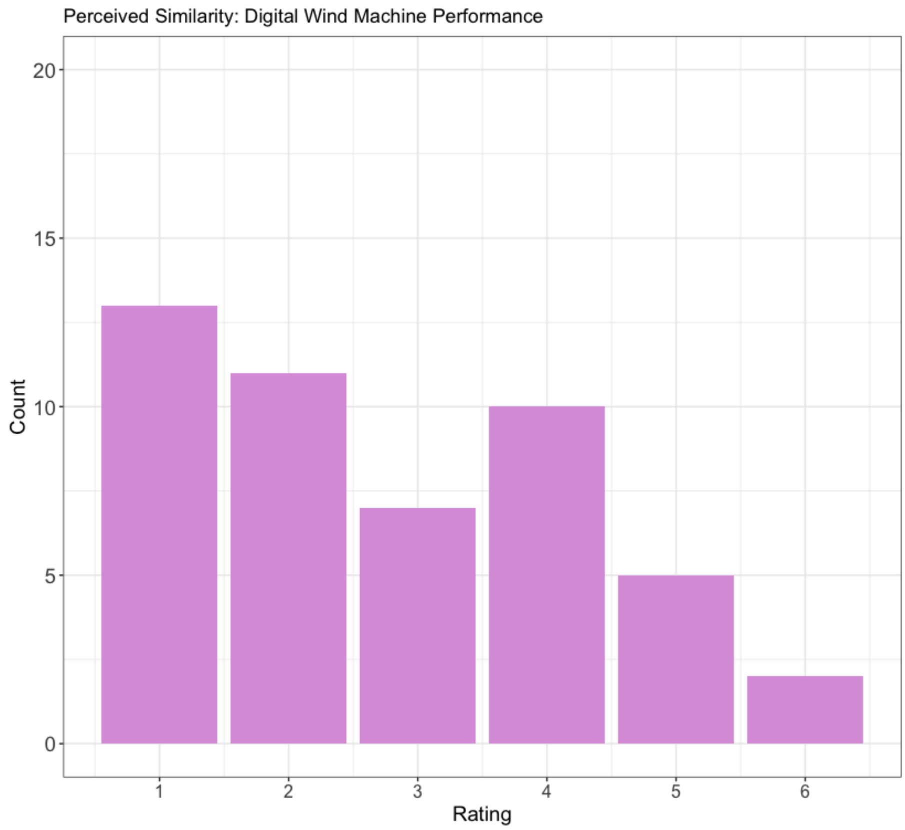

Figure 7.11: Similarity ratings for the prototype digital wind machine performances when compared to the sound stimuli they imitated.

Test: Wilcoxon Signed Rank	Significance	Effect Size
Z = -5.40	p <0.01	-0.78 (large) statistical power = 0.8

Table 7.10: Results of the statistical testing comparing how performances with the two wind machines were rated for similarity to the stimuli.

Test: Kruskal-Wallis	Significance	Effect Size
Acoustic easiness rating by group: H(3) = 5.36	p >0.05 (not significant)	0.03 (none) statistical power = 0.8
Digital easiness rating by group: H (3) = 1.33	p >0.05 (not significant)	0 (none) statistical power = 0.8

Table 7.11: Results of the statistical testing to confirm no order effects on the easiness ratings given to participants performances with the acoustic and prototype digital wind machines.

the performances with its digital counterpart (Table 7.10). Participants therefore rated the similarity of the wind machine performances to the stimuli significantly differently depending on whether they were performing an acoustic or digital wind-like sound.

7.4.5.2 Perceived Easiness of Play

Participants' scores for their responses to the statement "The wind sound is easy to play" were scored were scored according to their place on the scale with values from 1 to 7. A Kruskal-Wallis test was then performed on these easiness ratings given across each of the groups according to the order of performance system and the order of presentation of stimuli. This confirmed that there was no statistically significant difference between the ratings in each group, confirming that no order effects had influenced the results (Table 7.11).

A summary of the easiness ratings showed that the acoustic wind machine had a higher mean rating for ease of play than the prototype digital wind machine (Table 7.12).

Again, these ratings were collated according to the wind machine being rated to create bar graphs of the results. This showed that ratings for the acoustic wind machine were

Sound Played	Mean	Standard Deviation (SD)	Median
Acoustic	4.98	1.19	5
Digital	3.04	1.41	3

Table 7.12: Summary of the ease of play ratings for the performed acoustic and digital wind machine sounds.

Test: Wilcoxon Signed Rank	Significance	Effect Size
Z = -5.62	p <0.01	-0.81 (large) statistical power = 0.8

Table 7.13: Results of the statistical testing to compare how the acoustic and prototype digital wind machines were rated for ease of play.

clustered in the higher scores of the scale (Figure 7.12).

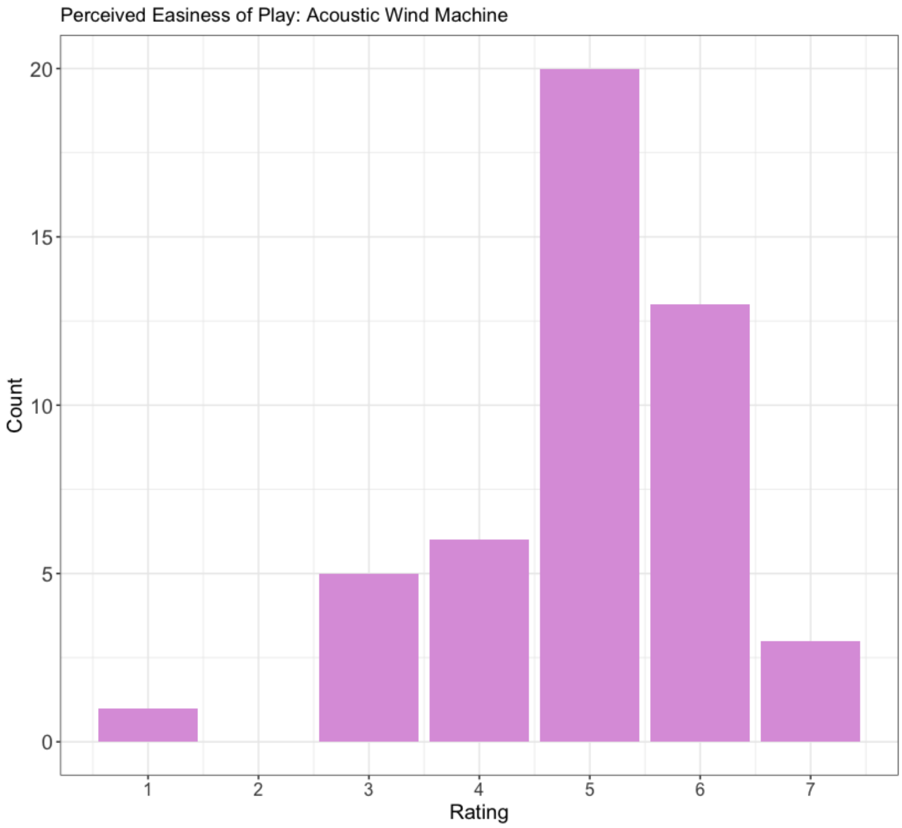

Figure 7.12: Ease of play ratings for the acoustic wind machine.

The bar graph for the prototype digital wind machine ratings showed a clear clustering towards the lower scores on the scale, although the ratings were clustered more evenly across the scale (Figure 7.13).

A Wilcoxon signed rank test was also performed on the easiness ratings to statistically compare the results for each wind machine. The test confirmed a statistically significant difference between the easiness ratings given to the acoustic wind machine and those given to the prototype digital wind machine (Table 7.13). Participants therefore rated the acoustic wind machine as significantly easier to play than the prototype digital wind machine.

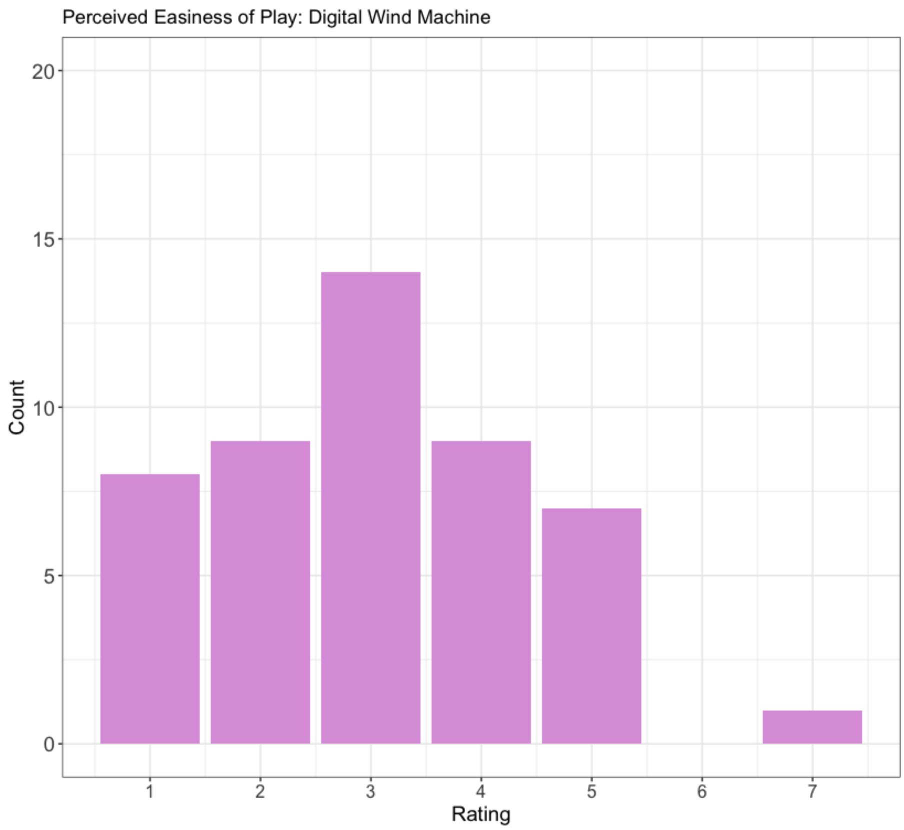

Figure 7.13: Ease of play ratings for the prototype digital wind machine.

7.4.5.3 Descriptions of Sounds

Participants were then invited to describe the acoustic and digital wind-like sounds by choosing as many descriptors as they liked from a list. These descriptors came from a range of categories:

- weather-associated (breeze, gale)
- general descriptions of force (gentle, strong)
- onomatopoeic descriptions of wind (shrieking, howling)
- an action or movement-oriented onomatopoeic historical descriptor (swishing)

Participants' responses were collated to produce a bar graph in R comparing the frequency of the descriptors given to each wind machine (Figure 7.14). Participants chose not to add their own descriptors to the list, but instead chose from the descriptors provided.

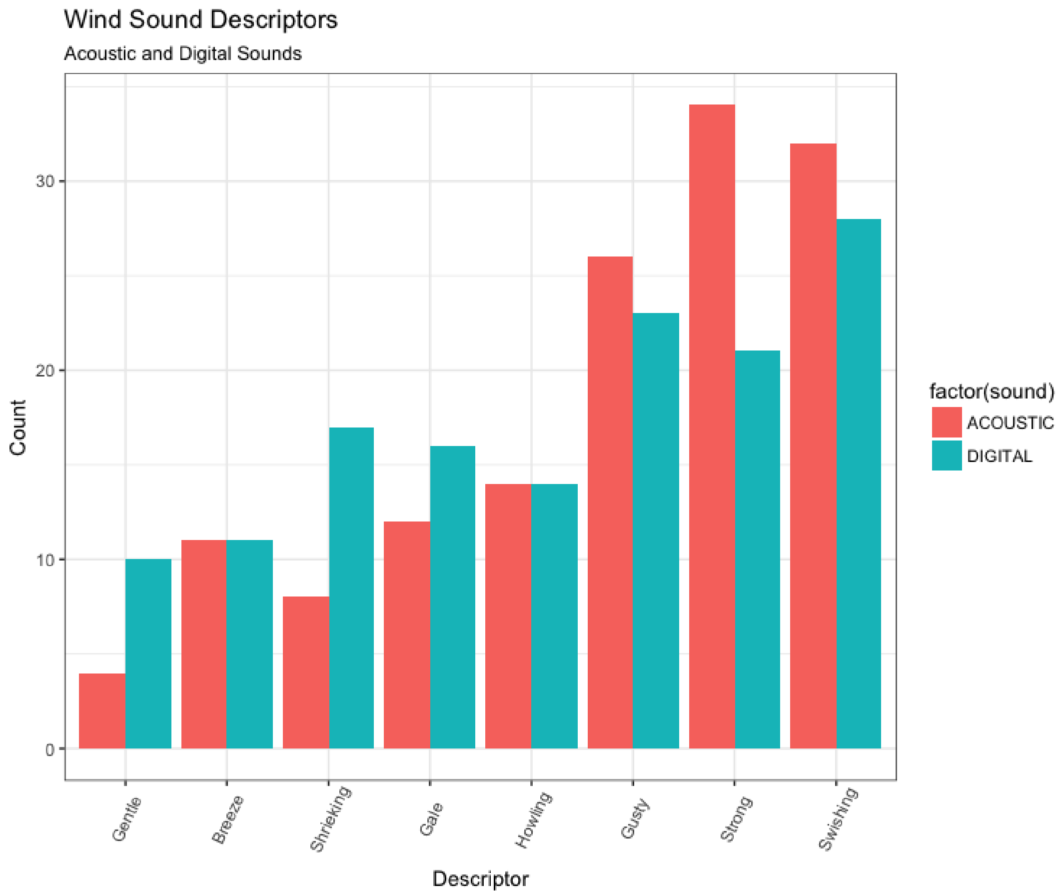

Figure 7.14: Descriptors chosen for the acoustic wind machine and prototype digital wind machine, ordered by frequency.

This showed that the most popular descriptor for both the acoustic and prototype digital wind machine sounds was the action-oriented *swishing*, followed by the force descriptor *strong* and the weather-associated *gusty*. The acoustic wind machine scored more highly across these three descriptors than its digital counterpart. The prototype digital wind machine was described with a fuller spread of adjectives, and was described more often as *shrieking* and *gale* when compared with its acoustic counterpart. Both wind machines were described as *howling* with the same frequency. The digital wind machine was described as *gentle*, *breeze*, *shrieking* and *gale* more frequently than its acoustic counterpart. This may reflect the fact that participants perceived the digital wind-like sound as having a narrower bandwidth of frequencies than the acoustic wind-like sound in performance.

7.4.5.4 Free Description

Participants were also given the opportunity to give some free description of their experience of performing with the wind machines in response to the stimuli. Some participants chose not to offer a comment. The comments that were given were coded to highlight some general trends and a few particularly interesting issues:

- Some participants seemed to distinguish between the sources of the stimuli they were asked to imitate, and reported the matched sounds as easier to imitate (e.g. acoustic wind machine imitating acoustic wind stimuli) than the others.

- Participants readily connected the speed of rotation of the handle with what they described as the speed, motion, rhythm or pace of the resulting wind sound, whether it was acoustic or digital in origin.

- The level of effort and force in the movement of the handle were connected with the quality of the wind sound produced, whether acoustic or digital in origin. For example, increased force in the handle movement was linked to increased force or *strength* in the resulting wind.

- Words like *gentle*, *strong*, *violent* or *gusty* were used to describe the qualities of both the acoustic and digital wind sounds. It is likely that participants acquired this vocabulary from the list of descriptive words they were asked to choose from prior to offering free description. While this may have unduly influenced their descriptions, it is interesting to note that other words from the list (e.g. *howling*) were not used to describe the sonic results of their performances.

- Some participants reported that the handle felt like it became heavier to turn when controlling the acoustic wind sound.

- Despite being informed that they would be playing a wind sound with the crank handle, one participant discussed their digital wind performances as rain sounds.

- One participant highlighted that they perceived a disconnection between the crank handle movement and the resulting digital wind sound.

Wordclouds were produced in R from participants' free descriptions, with very popular but less usefully descriptive words like *sound* or *sounded* removed to allow a closer look at others (Figure 7.15).

Figure 7.15: Wordclouds produced from participants' free descriptions for their acoustic wind machine performances with the acoustic wind machine (L) and prototype digital wind machine (R).

The resulting wordclouds show that while the acoustic wind machine may have *felt like wind* in performance, the crank handle was discussed more frequently in relation to the experience of performing with the prototype digital wind machine. This suggests that while performing with the digital wind sound, the coupling between the handle's movement and the resulting sound and how this might be working became an area of focus for participants, whereas during performance with the acoustic wind machine the wind sound felt more in-hand.

Wind Stimulus	**Rotations**	**Gesture Codes for Performed Responses**	
		Acoustic	**Digital**
Acoustic	1	G1aa	G1da
	2	G2aa	G2da
Digital	1	G1ad	G1dd
	2	G2ad	G2dd

Table 7.14: Gesture codes for the corpus of audio clips produced from participants' performances with the acoustic and prototype digital wind machines.

7.4.5.5 Acoustic Analysis of Performed Sounds

The performance step of the experiment produced a corpus of recordings of participants' performances of the acoustic and prototype digital wind machines in response to both the acoustic and digital wind stimuli. These were exported from Pro Tools as 24bit/48KHz audio clips, and coded for analysis according to the performance gesture (1 or 2 rotations) and the wind machine being performed (acoustic or digital) (Table 7.14).

These audio clips were then analysed in Matlab using the MIR Toolbox (Lartillot & Toiviainen 2007) to produce numerical measures of their acoustic qualities. These measures were first outlined in Chapter 4, but are briefly restated here for the understanding of the reader:

- Spectral Measurements:
  - *Brightness*: the high frequency energy in the spectrum above 1500Hz
  - *Inharmonicity*: how noisy the wind-like sound is relative to its Aeolian tone
  - *Spectral Centroid*: the frequency where the most energy of the sound is located
  - *Spread*: how spread out the spectrum of the sound is across the frequencies
  - *Skewness*: how clustered the spectrum is towards the low and high frequencies
- Amplitude Envelope Measurements:
  - *Event Density*: the average number of onsets in the sounds envelope per second

The resulting numerical values for each feature were then collated together for statistical analysis in R.

To establish whether the kind of stimulus presented to participants might have influenced how they performed the wind sound, gestures performed with the same wind

Test: Wilcoxon Signed Rank	Significance	Effect Size
Acoustic Wind Machine: Event Density for gestures G1aa and G1ad Z = -3.46	$p < 0.01$	-0.49 (medium) statistical power = 0.8
Prototype Digital Wind Machine: Event Density for gestures G1da and G1dd Z = -2.14	$p < 0.05$	0.3 (medium) statistical power = 0.8

Table 7.15: Results of the statistical testing to compare the acoustic analyses of participants' recorded performances.

machine were paired, and the numerical measures of their acoustic qualities statistically compared. For example, two rotations performed with the acoustic wind machine in response to an acoustic stimulus (G2aa in Table 7.14) were compared to two rotations performed with the acoustic wind machine in response to a digital stimulus (G2ad in Table 7.14).

A Wilcoxon signed rank test was then performed to compare each acoustic feature of the gesture (e.g. G2aa's spectral centroid measurements compared to G2ad's spectral centroid measurements). These tests established that no statistically significant difference existed between the gestures paired by stimuli source across all of the spectral measures.

For the measures of event density, no statistically significant difference was found between the gesture pairs for two rotations performed with the acoustic (G2aa and G2ad) and prototype digital (G2dd and G2da) wind machines. However, statistically significant differences were found for measures of event density for the single rotation performed with both the acoustic wind machine and the prototype digital wind machine (Table 7.15). This suggests that the gesture of two rotations performed with either of the wind machines was quite consistent regardless of whether participants had first listened to a stimulus that matched the sound that they were performing. For a single rotation, performances seem to have been more directly influenced by whether the stimulus presented matched the sound of the wind machine being played.

7.4.6 Summary of Findings

This step of the experiment aimed to establish whether there was perceived similarity between the experience of performing with the acoustic wind machine and that of performing with the prototype digital wind machine. The results established that, while the continuous sonic feedback was the only kind of multimodal feedback that changed between these two performance conditions, participants found the acoustic wind machine easier to play and perceived it as sonically similar to the stimuli used to elicit their performances. By contrast, the prototype digital wind machine was rated as significantly less easy to play, and participants found their performances with it to be significantly less similar to the stimuli they were trying to imitate. Statistical testing showed that the ratings for similarity and ease of play were significantly different depending on the kind of wind machine being rated, and so the results did not allow the null hypothesis to be rejected (Hc0).

When asked to choose from a list of descriptors for the acoustic and digital wind sounds, participants preferred the action-based descriptor *swishing* for both sounds, and were more confident in categorising the acoustic wind machine (as *gusty*, *strong* and *swishing*) than the prototype digital wind machine. Some interesting information emerged from participants' free description of their performances, particularly that the crank handle seemed to become an area of specific focus during their performances with the prototype digital wind machine, whereas performance with the acoustic wind machine seemed to place the sound more intuitively in-hand.

Acoustical analysis of the corpus of wind sounds produced from recordings of participants' performances established that there was no statistically significant difference in the acoustical measurements of sounds performed in response to a stimulus that matched the wind machine being played when compared with performances responding to an unmatched stimulus. The exception to this finding was the measurement of event density, or number of onsets in the sounds amplitude envelope per second, which was found to be significantly different for a single rotation performed with the acoustic wind machine between the acoustic and digital wind stimuli. The same pattern was visible for a single rotation with the prototype digital wind machine. This suggests that the type of stimuli had elicited a different kind of performance from the participants. Given the medium effect size observed here, further testing with a larger number of participants would be able to confirm these results.

This section has reported on the final step of the experiment, where participants performed with both of the wind machines and described their experiences. The next section draws together the findings from each step of the experiment and discusses their potential implications.

7.5 Discussion

This chapter has so far reported fully on the procedure and results of each step of the experimental evaluation. This section first reflects on the experimental procedure and apparatus used, and then brings the findings from each experimental step together in order to examine what can be concluded from the results.

7.5.1 Experiment Procedure

The experiment design successfully combined listening and performance tests to evaluate the acoustic and prototype digital wind machines in light of previous research in environmental sound perception, SID and DMI design. Enough participants were recruited for the experiment to facilitate a robust statistical analysis despite some data having to be excluded from each experimental step. An expansion of these experimental steps into dedicated experiments would allow the findings of this research to be explored further. The results produced from this evaluation do suggest that there is scope for further work in this direction.

In the performance step of the experiment, participants seemed to use the vocabulary of descriptors they had been asked to choose from to describe their experience of performing with each wind machine when they wrote their free descriptions. Future experiment designs should ensure that participants are not primed with descriptors like this in order to capture more of their own words when they are reflecting on their experiences. As highlighted in Section 7.3, it would have been pertinent to ask participants what they thought the source of the sounds was as part of the second listening step of the experiment, before they ranked the wind-like sounds for speed. This information on perception of the source of the sounds would also have been interesting to elicit from the performance step of the experiment, and could be included in a future study.

7.5.2 Experiment Apparatus

The experimental apparatus worked well during testing. The OpenSesame experiment platform (Mathôt et al. 2012) proved to be a robust method of presenting stimuli and collecting data, although choices made by the author during the coding process for the program resulted in participants being able to move on to the next section of some parts of the test without first completing all of the data entry in the previous section. This had a particular impact on the second listening step of the experiment, where a few participants did not rank all of the sounds in each section for speed, meaning that their data had to be excluded from the final analysis. This could be remedied for a future experimental procedure.

The acoustic wind machine interface proved to be very robust under repeated performance testing during the experiment, although the longitudinal movement of the steel axle on the A-frame did result in one participant's exclusion from the performance step data collection, as the rotary encoder's gear moved out of range of its counterpart and could not activate the program in Max/MSP. Adjustment and monitoring of the ratchet strap tightening the machine's A-frame together ensured that this was an isolated incident.

The gesture recorder in the Max/MSP program was used during the performance step of the experiment to record data from the rotary encoder for each participant, but this proved too difficult to later sync accurately enough to the recorded performances in Pro Tools to facilitate further analysis. Implementing a shared timecode solution, or using software such as Reaper[3] to record participants' performances along with simultaneous data sent from Max/MSP, could be a future solution to this issue. This would produce a database of gestures for more detailed analysis of participants' performances. It would also allow the prototype digital wind machine to be automatically activated using participants' previous performances, facilitating further calibration and adjustment of the model.

The headphones used for the performance step of the experiment may have had some impact on participants' experiences of the continuous sonic interaction with each wind machine. There will have been some small bleeding of the acoustic wind machine's sound in the testing space into the headphones during the performances, although this will have happened equally across all of the performance conditions. Future work could examine the potential of developing the propagation of the digital wind sound to bring it closer to the sound of the acoustic wind machine being conveyed to participants via a microphone and headphones. This could be achieved via the addition of spatialisation or reverberation to the program in Max/MSP, or the use of a loudspeaker to deliver the digital wind sound via a microphone to participants and replicate some of that same acoustic bleed in the headphones.

[3] Digital audio production software by Cockos.

7.5.3 The Perception of Wind-Like Sounds

The first listening step of the experiment confirmed the wind-like qualities of the sound of the acoustic wind machine, as participants rated it as highly similar to the natural wind gestures it imitated. The scree plot produced as part of the analysis of the similarity ratings suggests that the most important factor in participants' perception of the sounds was their timbre rather than their temporal patterning. This suggests that the acoustic wind was perceived as timbrally similar to the natural wind sounds, confirming its efficacy as a wind *effect* and evidencing the claims of historical theatre practitioners discussed in Chapters 3 and 5.

As timbre was so important to the results of this experimental step, it appears that the *steady* or *gusty* gesture of the wind sounds had less of an influence on the perception of their similarity. The results of this listening test clustered each source of wind together, rather than the winds being clustered by gestural category. The primary importance of the spectral features of the stimuli in their ordering along the dimensions of the multidimensional scaling solution chimes with the previous research by Gygi et al. (2007) on which this listening step was closely based.

Therefore like Gygi et al. (2007), this research does not follow Vanderveer's (1980) previous work, which found that temporal patterning was the dominant feature in perceived similarity of environmental sounds. The results of the analysis did suggest that the event density, a measure of the onsets in a sounds amplitude envelope per second, was an important factor in the ordering of the sounds along the second dimension of the MDS solution. There is more to explore here in terms of clarifying the reasons for these different conclusions. As outlined in Section 7.2, a larger corpus of wind sounds and an expanded listening test to compare their similarity could improve the sensitivity of the statistical analysis and reflect more of the individual differences in the similarity ratings. This should also help to examine whether temporal patterning, or features of the amplitude envelope of a continuous sound like wind, is the second most important feature when it comes to perceived similarity.

By contrast, the digital wind-like sounds were perceived to be somewhat timbrally distinct from their acoustic counterparts, although part of the cluster of continuous sounds. Examining the MDS solution along Dimension 1 (Figure 7.4) shows that the digital wind-like sound is slightly closer to the side of the pitched horn sound, indicating that perhaps the whistling Aeolian tone of the acoustic wind has been overemphasised within the digital model. Potential improvements to future iterations of the program in Max/MSP could be effectively evaluated using the procedure described in Section 7.2 given that this baseline of similarity between the two wind machine sounds has been established.

The disparity in similarity between the acoustic and digital wind-like sounds was also evidenced in the third step of the experiment, where participants performed with both wind machines. Again, participants affirmed the wind-like quality of the acoustic wind machine in performance and rated it as very easy to play, confirming the claims of historical theatre practitioners. The experiment design perhaps missed an opportunity here to elicit participants' own descriptions of their experiences of the acoustic wind by placing the descriptor selection task before the free description task. It is clear however that participants felt that there was a connection between their performance and the acoustic wind sound, although future work should ensure that this is given space to be articulated more fully. A more qualitative procedure incorporating interviewing and filming may help to expand upon what this *feeling* of wind-in-hand is for participants, particularly in relation to other experiences they may have previously had of music or sound performance. This would also enable other aspects that participants highlighted to be explored, including the potential of the sonic qualities of each wind-like sound to in turn influence the feel of the crank handle in performance, something that concurs with previous research on the perception of haptics and movement (DiFranco et al. (1997), Avanzini & Crosato (2006*a*) and Turchet et al. (2013)).

Participants rated their performances with the prototype digital wind machine as less similar to the stimuli they were presented with. The prototype digital wind machine was also rated as not very easy to play, suggesting that participants did not perceive the sonic feedback as directly responding to their performance gestures. This reflects the acoustic analysis of the prototype digital wind machine presented in Chapter 6, where variation in the rotational gesture did not produce response in the model that tracked tightly with the amplitude envelope of the acoustic wind machine.

7.5.4 Virtual Rotating Mechanisms

The second step of the experiment showed that participants could not perceive clear differences between different rotational speeds of the crank handle when presented with the acoustic and digital wind-like sounds that they produced. While the results may have been influenced by some of the factors previously discussed in Section 7.3 (including the duration of the sound stimuli), it is clear that had the stimuli clearly evoked the metaphor of a turning handle, participants would have been able to differentiate between the speeds more easily. It is unclear whether *rotation* can be perceptually understood as directly producing *wind*, or perhaps more accurately in this case given the use of partial rotations - *scraping*.

As this procedure was developed from the experimental evaluation of the Spinotron sound model (Lemaitre et al. 2009), it is useful to reconsider the stimuli used for that study in light of the findings here. The Spinotron sound was based on a ratchet model, which rather than being a continuous scraping sound like that evaluated in this study, consists of many impacts being triggered one after the other. The ratchet model also produced changes in loudness and modulation depending on the wheel position (2009, p.981). The impacts or clicks produced by Spinotron's virtual rotating ratchet wheel may have been much simpler for participants to determine speed of movement from, as this is a stepped sound rather than a continuous scrape that only evokes speed from a change in pitch. The timing between each click provides a reference for participants as to the speed of movement, but a continuous scrape may require another reference scrape to establish a perceptual notion of speed. In a future experimental evaluation, this could be achieved with a practice step before the test step, to help participants establish a sense of a scale of speed in much the same way as they practiced rating for similarity before the first step of this experiment.

In addition, given that the Spinotron's ratchet model was a sequence of impacts, an understanding of the metaphor of *turning* may not have been required to accurately rate its sounds for speed. Instead, the speed of a sequence of impacts could also be perceived or understood as the result of a linear movement, produced by sliding a finger along the teeth of a comb at a particular speed for example. Some of the free descriptions of the cause of Spinotron's sound elicited before proceeding to the speed ratings did not focus on the metaphor of a turning wheel, for example. Some participants described hollow objects, or other metaphors such as a bouncing ball (2009, p.981).

There is more to explore here in order to establish whether a sound can really evoke a virtual mechanism, or movement in a particular direction (*turning* compared with *up/down* for example) rather than just an event of a particular magnitude (*speed*). It may have been the reference to a turning handle that was a factor in the lack of clarity of participants' responses to the speed-ranking task. It is proposed that an earlier iteration of the prototype digital wind machine discussed in Chapter 6, that produced clicking at high rotational speeds, would be an ideal stimulus to investigate this possibility further. A further dedicated listening test, with some space for participants to describe what they perceive to be the source of the sound stimuli, might yield more conclusive results on the ability of a *scraping* or *wind* sound to evoke the rotational mechanism that produced it.

7.5.5 Action Coupled to a Wind-Like Sound

The performance step of the experiment produced a corpus of recordings of participants' performances of two very simple gestures - a single rotation and two steady rotations - with both wind machines. As reported in Section 7.4, an acoustical analysis of these sounds showed a statistically significant difference between the frequency of onsets in the amplitude envelopes of the single rotations. This suggests that participants played the wind machines quite differently depending on the kind of wind stimulus (acoustic or digital) presented to them to elicit their performance. However, this difference was not evident in the gestures of two steady rotations. It is possible that participants understood the wind stimuli of two steady rotations much more easily, but given the lower ratings for similarity and easiness participants gave to the prototype digital wind machine, it is unlikely that the digital stimuli were so simple to imitate. The difference between the single and double rotational gestures must therefore be due to another factor.

It is proposed that the continuity of gestural response evidenced across the categories of stimuli to elicit the performances of two steady rotations has more to do with the flywheel action of the acoustic wind machine interface than the responses of participants. As explored in Chapter 5, the acoustic wind machine's ability to store energy in its cylinder and axle mechanism due to rotational inertia has a profound effect on the perceptual experience of performance. With a single rotation, the cylinder does not have time to accumulate rotational energy and push forward from the movement of the performer's hand. However, with a gesture of two rotations, the moving cylinder must be imposing more of its dynamics, and hence some regularity, on the performer's rotation of the crank handle. This could be confirmed with a further performance experiment examining a broader range of gestures. In particular, a robust method of recording data from the rotary encoder would help to illustrate the influence of the cylinder's rotational inertia on the performer's movement in the continuous sonic interaction.

This section has discussed the results of the experimental evaluation and developed some conclusions from the findings presented in this chapter. This chapter will now be summarised.

7.6 Summary

This chapter has reported on the procedure and results of an experiment conducted to evaluate the wind-like sounds and sonic interactivity of the acoustic and prototype digital wind machines with participants. The results of this experiment show that further work is required to bring the sonic response of the program in Max/MSP closer to that of the acoustic wind machine it models. The acoustic wind machine itself was rated as highly similar to the recordings of natural wind that it imitated. It was also rated as significantly easier to perform with when compared with its digital counterpart. The speed of a gesture of rotation with the crank handle was difficult for participants to clearly perceive from the resulting wind-like sounds produced by both the acoustic and prototype digital wind machines.

The next chapter restates the research questions first presented in Chapter 1 of this thesis and discusses the research methodology and findings in light of the theoretical background presented in Chapter 2. Suggestions for future work are also presented.

Chapter 8

General Discussion and Findings

This thesis has presented a research enquiry into the sonic interactivity of late nineteenth and early twentieth century theatre sound effects, and how their enactive properties might help to bridge the gap between simple hand actions and the vast potential of digital soundmaking. This chapter now draws together the work presented in the previous chapters, and considers the research findings in light of the research questions first outlined in Chapter 1.

Section 8.1 re-states the research questions set out at the beginning of this thesis. Section 8.2 summarises the methodology that was used to answer these research questions. Section 8.3 discusses the work undertaken as part of this research and considers the main themes and issues arising in light of theoretical background introduced in Chapter 2. Section 8.4 summarises the findings of this research in relation to the research questions. Section 8.5 looks ahead to potential future work arising from this research. Section 8.6 summarises this chapter and concludes this thesis. Section 8.6 concludes this chapter, and this thesis.

8.1 Research Questions

The main research questions investigated in this thesis were:

- How can historical theatre sound effects be fully investigated as interactive sounding objects?
- How can the enactive properties of historical theatre sound effects be captured in the design of a digital sonic interaction?

At the heart of these research questions was the proposal that late nineteenth and early twentieth century theatre sound effects practitioners had already solved some of the design problems inherent in present-day Sonic Interaction Design (SID) and Digital Musical Instrument (DMI) design practice, which enabled them to create continuously interactive sounding objects that allowed the performer to enactively build skill in the performance of an everyday sound through a process of exploration and practice. It was argued that if the sonic interactivity of these historical sound effect designs could be somehow defined and captured, these qualities could then be used to create a similarly immediate and intuitive encounter with digital sound.

As part of the theoretical background of this thesis (Chapter 2), the phenomenology of tool use was extended to the perceptual experience of soundmaking with objects, introducing the concept of *sound-in-hand* to describe how an everyday sound could feel embodied by a performer during a simple continuous sonic interaction. It was proposed that by engaging with historical theatre sound effects, digital sound could itself become the focus of bodily intent and fully in-hand during an interaction. It was also proposed that engaging with this historical context of soundmaking in late nineteenth and early twentieth century theatres could offer an opportunity to examine and reflect on the design conventions inherent in present-day SID and DMI design, as well as reveal more about how these digital sound disciplines could draw on the embodied expertise of present-day Foley artists or other performative soundmaking practitioners in the future.

The following sub-questions were developed from the main research questions in order to direct the stages of work reported in this thesis:

- What theories, methods and frameworks are available for studying historical theatre sound effects as interactive sounding objects?
- Which historical theatre sound effect designs offer the most interesting action-sound configuration to the operator?
- How might a historical theatre sound effect be made digital?
- How can a historical theatre sound effect be best evaluated?

This section has re-stated the research questions and central premise of the enquiry presented in this thesis. The next section revisits the methodology pursued to answer the research questions.

8.2 Research Methodology

This research implemented a transdisciplinary methodology that converged the fields of Sonic Interaction Design (SID) and Digital Musical Instrument (DMI) design, and integrated approaches from theatre history, the study of human perception, remaking as enquiry, research through design, and both acoustical and experimental evaluation. Chapter 1 argued that this integration of distinct disciplines was crucial in order to facilitate an exchange of designerly knowledge between late nineteenth and early twentieth century theatre sound effects design practice, a hitherto unexplored context of performative soundmaking, and present-day strategies from SID and DMI design for interacting with digital sound.

A literature review was carried out to establish the theoretical background of this thesis (Chapter 2). The term *interactive sounding object*, a designed object for acoustic or digital soundmaking through human manipulation, was introduced to define the point of convergence between previous work in SID and DMI design and bind it to the historical context of soundmaking in late nineteenth and early twentieth century theatre. Previous research examining human perceptual experience of continuous sonic interactions was explored, highlighting work to examine how certain correspondences between actions and sounds might be more meaningful to performers, and examining how sound itself might suggest an action to a listener through a "sonic affordance" (Altavilla et al. 2013). The complexity of transitioning from acoustic to digital soundmaking was discussed, and strategies for designing digital interfaces and sounds that captured some of the material resistances of real-world sonic interactions were considered. Theatre sound effects were then introduced, and defined as interactive sounding objects capable of complex soundmaking through their use of acoustic mechanisms, and contextualised in terms of a simple everyday sonic interaction with a piece of crumpled paper and a digital interface for performing a digital sound. It was argued that theatre sound effect designs might be particularly useful in examining the perceived sonic affordances of particular sounds, developing an understanding of how mechanisms might add richness to simple performance interfaces, exploring how complex sounds can placed *in-hand* for a performer, and designing evaluations of continuous sonic interactions.

Historical research was then undertaken to draw many sources together and establish a fuller picture of the practice of theatre sound effects design and performance in the late nineteenth and early twentieth century from the perspective of a present-day sonic interaction design practice (Chapter 3). Theatre sound effects design and performance practice was established as an early site of performative soundmaking with objects, and linked to that of later media that grew in popularity in the twentieth century, such as silent cinema sound effects performance or present-day Foley. Some potential reasons for the hitherto neglected status of theatre sound effects as a focus of research were also considered. Two existing frameworks were then used to connect many fragments of information on this hitherto ill-documented soundmaking practice together in light of the concerns of this research. First, Cross' "five aspects of designerly ways of knowing" (2006, p.29) focused a general discussion of the design practice for creating sound effects in theatres, and highlighted how practitioners translated often abstract requirements for specific qualities of sounds into actions, mechanisms and materials. This revealed a rich and interesting design practice based on an embodied knowledge of carpentry and mechanisms, the influence of materials and actions on resulting sounds, and the importance of context and the role of the audience's perceptual experience in the success of an effect. Secondly, the most commonly used sound effect designs of the late nineteenth and early twentieth century were collected together and considered in more detail. The Sound Design Toolkit (SDT) taxonomy of low-level sound events and basic textures (Baldan et al. 2017) was used to classify these theatre sound effect designs in terms of their action and sound correspondences, opening up a new design framework to realise enactive recreations of these historical effects using the SDT digital synthesis algorithms. A specific design, a theatre wind machine, was then discussed in detail as a particularly interesting sound effect to examine more closely and attempt to capture digitally.

The work undertaken to construct the acoustic wind machine was presented in Chapter 5. Several historical sources were brought together to inform the initial design to help expand on the limited instructions provided by historical practitioners. This approach built a connection to prior scholarship in making and the digital humanities. Extending Elliott et al.'s (2012) remaking and prototyping of stage illusions to the study of the wind machine allowed not only its historical design instructions to be critically examined, but also an exploration of the perceptual experience of its performance. Despite the apparent simplicity of the mechanisms and materials used to create theatre sound effects, they were clearly the result of a long and iterative process of imagining, prototyping, testing, calibrating and re-testing. Refining the wind machine design and making it work revealed more of the tacit knowledge of historical practitioners through the details omitted from their instructions and the techniques implied by their descriptions. These included an understanding of carpentry and fixings, as well as how to effectively facilitate only the desired movement in the final construction through mechanical design. As explored in Chapter 5, a small amount of extra weight added to the cloth, a slight inaccuracy in how the cylinder and axle mechanism moved, or too much distance between the crank handle and pivot point could have a profound impact on the experience of performance. If these issues were not resolved, the rotational gesture might require too much effort on the part of the performer, thereby limiting the sonic possibilities of the wind machine. Once constructed, the author explored the acoustic wind machine as an interactive sounding object. The acoustic wind machine was found to offer a compelling, rich and complex continuous sonic interaction, and produced a convincing wind-like sound in performance rather than just a complex scraping sound. The dynamics of the wind machine's cylinder-and-axle mechanism were investigated further, and found to have flywheel properties, allowing rotational energy to be stored and changes in rotational speed during performance to be resisted. The acoustic wind machine was performed to imitate field recordings of natural real-world wind sounds (BBC 1988), which revealed more about the importance of an irregular rotational gesture in producing a convincing wind effect. An entity-action model was then produced to describe the soundmaking components of the acoustic wind machine.

The insight gained from remaking and performing with the acoustic wind machine was then applied to the "wicked" problem (Rittel & Webber 1973) of capturing its sonic interactivity digitally (Chapter 6). Both the hardware and software components of the prototype digital wind machine developed through several iterations. The movement of the acoustic wind machine's cylinder was captured as data using a configuration of a continuous rotary encoder and laser cut gearing. This data was used to drive a prototype digital model programmed in Max/MSP using the SDT physical model of friction, and a mechanical mapping was implemented to allow the acoustic wind machine interface to simultaneously control the digital model through a rotary encoder coupled to its mechanism. The digital wind-like sound was produced by modelling two main soundmaking components of the acoustic wind machine - its slatted cylinder and its rough cloth. The two wind machines were then simultaneously performed and acoustically evaluated to compare their sounds and sonic response in performance. As the prototype digital wind machine was producing promising results when tracking with its acoustic counterpart at slow and regular speeds, it was argued that it would be appropriate to proceed to an experimental evaluation with participants. A parallel piece of work was undertaken to create a digital crank interface in order to activate and modulate the digital wind-like sound without the acoustic wind machine. The qualities of this digital crank were explored in performance, and its multisensory feedback was compared to that of the acoustic wind machine interface. This process revealed some other qualties of the acoustic wind machine's crank handle. In particular, it was observed that the digital crank shrank the perceivably large rotation afforded by the acoustic wind machine, and in doing this increased the possible speed of rotation. This suggested that the *large* and *slow* rotational gesture of the acoustic wind machine may be a critical part of how it offers articulation and expression to a performer in performance.

Chapter 7 then reported on an experiment designed to evaluate the acoustic and prototype digital wind machine with participants. The first listening step of the experiment, informed by previous research into environmental sound perception, asked participants to rate the acoustic and digital wind-like sounds for similarity in the context of some other everyday sounds. The stimuli used for this step of the experiment were based around *wind gestures*, where the acoustic and prototype digital wind machines were simultaneously performed to imitate recordings of real-world wind. Participants rated the acoustic wind machine sounds as most similar to the recordings of natural wind. However, the analysis of participants' similarity ratings of these sounds suggested that their spectral features were the most important factor in their ratings. The wind sounds were grouped together not by their *gesture*, but by the source of the sound (e.g. acoustic vs. digital). The second listening step, informed by prior research in SID, asked participants to rank the acoustic and digital wind-like sounds in terms of the perceived speed of the rotation that produced them. The test produced some interesting results, with participants unable to consistently rank the wind-like sounds in order of speed. This suggests that *rotation* could not be perceptually understood as causing *scraping*,[1] and that the idea of a virtual mechanism perceptually connecting the action and sound was potentially more complex than the test had managed to explore. A final performance step then guided participants to perform simple gestures with the acoustic and digital wind machines, and give some descriptions of their experiences of each continuous sonic interaction. The data produced from each step of the experiment was examined and subjected to statistical

[1]As outlined in Chapter 7, *scraping* is perhaps a more accurate term for the stimuli used for this step of the experiment, produced as they were by partial rotations of the acoustic wind machine interface.

analysis. Participants' performances were also recorded and acoustically evaluated to produce further data for analysis. The two systems were rated significantly differently. The acoustic wind machine was more highly rated for similarity to the stimuli participants were given to imitate, and as much easier to play than its digital counterpart. An examination of the free description offered by participants suggested that they did perceive their movement of the acoustic wind machine's crank handle as resulting in the performance of *wind*. So while participants did not directly perceive a specific *action* or *gesture* in the wind-like sounds during the two listening steps of the experiment, they somehow understood how their action linked to that sound in performance. An acoustic analysis and statistical comparison of participants' recorded performances showed that although the source of stimulus (acoustic or digital wind) had a significant effect on the gesture performed to imitate one rotation, this was not evident for the gesture of two rotations. This suggested that the gesture of two rotations activated the flywheel more fully, allowing it to guide participants' performances.

This section has summarised the methodology pursued as part of this enquiry in order to answer the research questions. The next section discusses how the research was pursued and considers the findings and issues arising in light of the theoretical background presented in Chapter 2.

8.3 General Discussion

Engaging with theatre sound effects practice from the late nineteenth and early twentieth century was introduced at the beginning of this thesis as a potential way to find solutions to the design problem of creating embodied, intuitive, and satisfying encounters with performable digital sounds. Now, following the completion of the stages of work described in Section 8.2, it is argued that the gap between simple hand actions and the potential of creative digital soundmaking has not yet been bridged. However, this encounter with a very different context of sonic interaction design has shed light on the design and evaluation conventions of present-day SID and DMI design. It has afforded a glimpse of some of the complexities of defining particular qualities of real-world sonic interactivity and of capturing them digitally. It has also framed creative soundmaking (whether acoustic or digital) as a *designerly* practice - one that will benefit from prior research in the broader field of design or from approaches such as research through design. It has also uncovered many of the *wicked* problems inherent in this practice. Indeed, the process that aimed to capture the acoustic wind machine as a digital interactive sounding object also captured many of its inherent design problems digitally, which in turn drew out interesting findings from the process of evaluation. We have so far revisited the research questions of this enquiry and the methodology used to guide the work undertaken to answer them. This section reflects on the research approach more broadly and considers the findings in light of the theoretical background presented in Chapter 2, as well as the potential contributions theatre sound effects might make to the design and evaluation of interactive sounding objects, as first presented in Section 2.4.

8.3.1 Wind-Like Sonic Affordances

Can the sound of real-world natural wind directly suggest an interface design or an action to a performer? The historical research presented in Chapter 3 did not reveal that practitioners designed their acoustic and mechanical effects in response to a perceived affordance of the real-world sound they were trying to imitate. It is not clear that the sound of real-world rain directly suggested a rotary mechanism and crank handle to practitioners creating a rain machine, for example. The "ways of knowing" ((Cross 2006, p.29)) glimpsed across many sources of information on the practice suggest that a more iterative and explorative method was at work in the creation of these acoustic and mechanical interfaces. With further historical research, perhaps focused on an earlier era of sound effects design, the specific origins of the designs described within this thesis could be more firmly established.

While it has not been possible to conclude whether the sound of real-world wind directly influenced the design of the acoustic wind machine, the working example created as part of this research facilitated an exploration of whether a wind-like sound could directly suggest a performance action to a listener through a "sonic affordance" (Altavilla et al. 2013). Both the author's exploration of the acoustic wind machine discussed in Chapter 5 and the results of the experiment with participants reported in Chapter 7 suggest that the rotation action afforded by the crank handle and its resulting wind-like sound do "correspond" (Hug 2008, p.15) during performance, in that the link between a *turning* action and the resulting wind-like sound are perceptually meaningful. If a wind-like sound can be the subject of ergoaudition (Chion 2015, p.671), and the rotation of a crank handle can meaningfully couple to the performance of a continuous wind-like sound, then that wind-like sound might contain a perceivable *affordance* that suggests the same performative action. The sound of real-world wind itself, or indeed a digital wind-like sound, might contain the same affordance.

The first listening step of the experiment reported in Chapter 7 attempted to explore this potential. However, when participants were asked to rate wind-like sounds for similarity, the *gesture* of the natural, acoustic and digital winds did not prove to be the most important factor in how the sounds were rated for similarity. Instead, the spectral features of the sounds were the most important, confirming more extensive research into environmental sound perception (Gygi et al. 2007). This implies that participants did not perceive the performance gesture of the wind sounds as their most immediately important feature, casting doubt on whether the sound of wind could directly suggest a sonic affordance of rotation to a listener. Similarly, the second listening step of the experiment explored a particular magnitude of the performance gesture - speed - to examine if different examples of wind-like sounds could be perceived as *faster* or *slower* than others. This built on previous research which evaluated a digital synthesis model of a ratchet with participants as part of the development of the Spinotron (Lemaitre et al. 2009). The results showed that participants could not consistently rate examples of partial rotations of the acoustic and digital wind machines by the speed that produced them. While some potential issues with the design of this step of the experiment have already been discussed in Chapter 7, this suggests again that the potential of a continuous wind-like sound to offer a perceived *sonic affordance* to a listener is in doubt.

This raises some interesting issues. Perhaps only certain sounds can suggest a performance action to a listener, or perhaps *continuous* sound is a special case when it comes to *action-listening*. Perhaps listening to sounds produced through action is qualitatively different depending on whether that sound is self-produced during performance or heard as recorded audio. Perhaps certain weather-associated "everyday sounds" (Gaver 1993*b*) evoke a particular response in a listener that by way of recognition are not immediately connected to human action. There is much still to explore here. It is proposed that this research has revealed that meaningful action-sound correspondences and sonic affordances are much more complex than has been suggested by previous research. In particular, these framings of human experience of action-activated sounds do not take take the qualities of an interactive sounding object - the facilitator of a meaningful performance gesture - into account. These qualities may be a critical part of what creates a meaningful encounter with a continuous sound. We will now move on to examine what this research has revealed about the potential meaning created by the interactive sounding object itself.

8.3.2 Enactive Sound Machines

Which specific qualities of the acoustic wind machine design make it *enactive*, allowing the performer to acquire new bodily knowledge about its sonic potential through rehearsal? As previously explored in Chapter 2 this research explicitly engaged with the potential of Franinović's "enactive sound design" approach and applied it to the design of interactive sounding objects. This is restated here for the reader:

- Sound affects the user's sensorimotor activity, guiding them to learn by building on previously accumulated tacit knowledge.
- The object only produces sound in response to the user's movement, thereby engaging their "willed action."
- Audition is not separated from other senses (e.g. touch) in order to enhance multisensory experience.
- Sound is a direct and continuous response to a user's movement. (2013, p.21)

Chapter 3 proposed that the acoustic wind machine was an interesting design to explore in light of this definition, particularly given its continuous and complex sound produced by a simple rotation action. As part of the framework for the enactive recreation of historical sound effects using the SDT suite of physical modelling synthesis algorithms (Baldan et al. 2017) presented in the same chapter, the wind machine design was classified with a simple rotation action to correspond with the sound produced by the [*sdt.friction~*] and [*sdt.scraping~*] objects in Max/MSP.

However, the remaking work presented in Chapter 5 of this thesis revealed that this gesture of rotation was actually much more perceptually complex in performance than historical descriptions of the effect suggested. The study of the working acoustic wind machine discovered that the freely spinning wooden handle of the crank continuously reorients the hand, wrist and arm during rotation, giving the sensation of somehow stirring in a figure-8 motion as the gesture progresses. Due to the influence of the tight side of the cloth against the cylinder and the need to move against its frictional force during the first half of the rotation, the gesture is in fact comprised of two parts, with the first half requiring more effort than the second half. In addition, the cylinder and axle design is in fact a flywheel mechanism, which accumulates movement energy through a high rotational inertia and resists changes in speed imparted by the performer through the crank handle. These features give responsive vibrotactile feedback, and a subtly shifting resistance to the movement of the crank handle, making the simple and constrained single degree of freedom (DOF) movement into an unexpectedly rich and compelling performance gesture. The complexity of the feeling of the crank handle when performing with the acoustic wind machine draws the experience of the apparently simple gesture perceptually closer to that of "sound tracing" (Godøy et al. (2006*a*), Caramiaux et al. (2011) and Caramiaux et al. (2014)) when continuously shaping and modulating the wind-like sound. This embodied experience had not been captured in historical texts.

Given the richness of the acoustic wind machine as a performance interface, it was used to drive the digital model in Max/MSP in order to preserve all of the same tactile and kinaesthetic feedback across encounters with both the acoustic and digital wind-like sounds. The digital model described in Chapter 6 was focused on recreating the main soundmaking components of the acoustic wind machine, which were defined as its slatted cylinder and moving cloth. Despite this approach, the results of the performance step of the experiment reported in Chapter 7 showed that participants rated the acoustic wind machine as significantly easier to play than its digital counterpart, suggesting that although the main soundmaking components of the acoustic wind machine had been prototyped digitally and the acoustic wind machine's tactile and kinaesthetic qualities were in place, there was something that this digital interaction had failed to capture.

This is particularly interesting as the prototype digital wind machine meets Franinović's criteria for an "enactive sound design" (2013, p.21) - its sound affects sensorimotor activity, continuous sound is the direct result of a user's movement, and performing with it is a multisensory experience. It is proposed that the critical factor here may be the lack of a "continuum of energy" (Cadoz 2009, p.218) from performance gesture to resulting sound. In the case of the prototype digital wind machine, although the real-world flywheel mechanism was part of the interaction it afforded, this flywheel was not also digitally modelled and connected to the real-time activation and modulation of the soundmaking components in Max/MSP. It is perhaps the absence of the flywheel regulating variations in the speed of the crank handle that can be seen in the acoustical analysis of the prototype digital wind machine to a rotational gesture of varying speed

(Chapter 6). It is proposed that if this mechanism is more clearly modelled digitally, the sonic response of the digital wind-like sound will be brought closer to that of its acoustic counterpart in performance, and the continuity of energy exchange across the acoustic and digital interactive sounding objects will be improved. As such, the design framework for enactive recreations of historical theatre sound effects presented in Chapter 3 has not adequately captured how the rotation action corresponds with the resulting wind-like sound. This could be expanded with a further layer of information to begin to account for the gestural complexity of the acoustic wind machine (Figure 8.1).

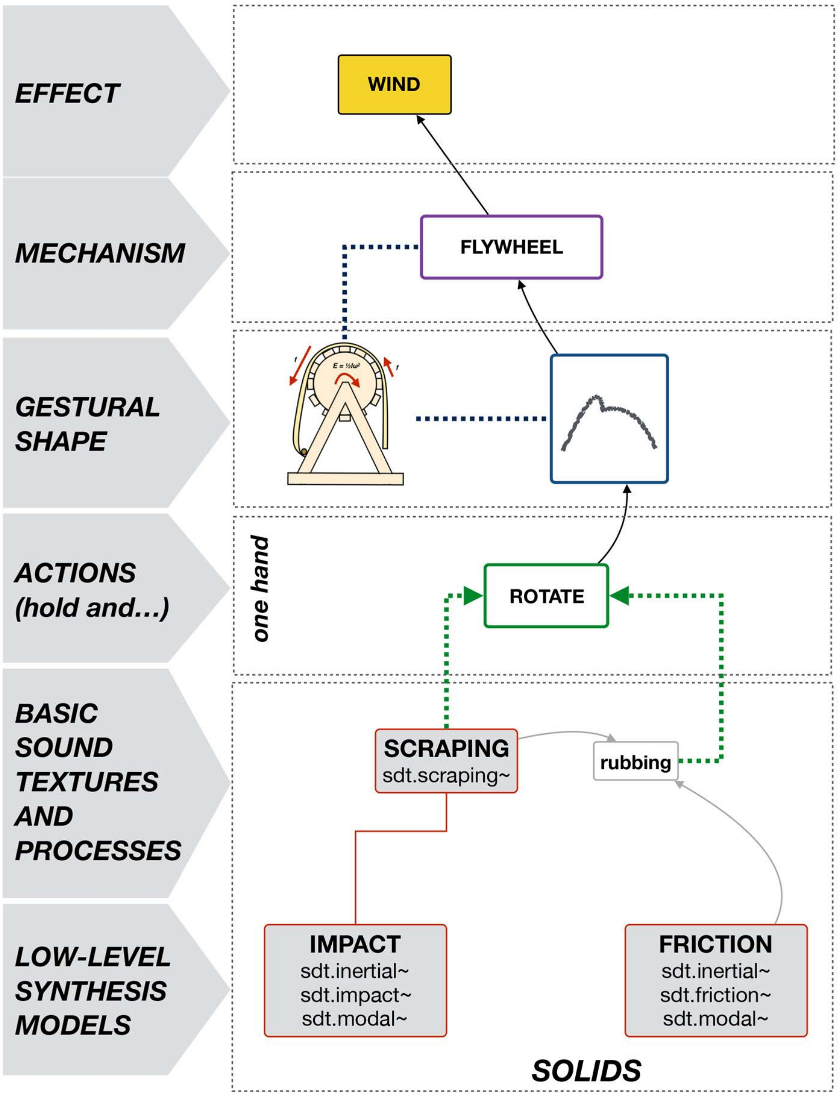

Figure 8.1: The SDT framework from Chapter 3 showing an extra layer of information for the influence of the cloth and flywheel on the gesture of rotation.

The inclusion of this information in the framework begins to account for the specific role of the material and mechanical properties of the interactive sounding object in linking action to sound in a meaningful way. It is proposed that this research has evidenced the specific importance of the mechanism of the acoustic wind machine in its ability to afford an *enactive* experience to a performer. The wind machine design gives the performer the experience of directly playing the wind-like sound, but in fact they are working in collaboration with or in opposition to the motion of the flywheel and cloth. As the mechanical *feel* of the acoustic wind machine is such a critical factor in connecting the crank handle to the wind-like sound in a perceptually meaningful way, this property should be explored further as a potential way to make action-sound correspondences meaningful.

Indeed, the need for a *continuity of energy* suggested here may be a requirement for a successful "enactive sound design" (Franinović 2013, p.21). The lack of *felt* energy exchange in a wind-like sound heard as recorded audio may be the reason why a sonic affordance of *turning* is not immediately perceivable to a listener. The complex *feel* of the machine may also be an important factor in its success as a wind *effect* rather than it being perceived as the direct result of of materials (cloth and wood) interacting. This research has clearly evidenced the complexity of theatre sound effects as interactive sounding objects, and the importance of their material qualities and mechanisms to their potential as enactive sound machines. In particular, the multisensory richness of a simple turning handle as a compelling performance interface suggests that the historical designs presented in Chapter 3 may offer more interesting ways to study enactive sound performance and create sensory complexity from apparently simple configurations of materials and mechanisms. Next, we will examine the way evaluation has been deployed as part of this research and discuss its implications.

8.3.3 Evaluating the Task of Soundmaking

Have the evaluation strategies explored within this research revealed how to best evaluate theatre sound effects as interactive sounding objects? Section 2.4 reflected on the fact that while SID conducted evaluations informed by procedures from HCI research, there was still little agreement among DMI researchers on the best way to evaluate their designs (Barbosa et al. 2015). This research drew together procedures from previous studies in environmental sound perception (Gygi et al. 2007), SID (Lemaitre et al. 2009) and DMI design (Poepel 2005) in order to examine the acoustic and prototype digital wind machine in an experiment with participants. As previously discussed in Chapter 7, the experiment conducted here has established a baseline of metrics such as *similarity* and *ease of play* between the acoustic wind digital wind-like sounds that can be reexamined following further development of the digital model in Max/MSP. Adjustments to the order of tasks within the experiment design itself may also improve upon the data captured. Ensuring that participants are not given descriptors for the wind-like sounds before soliciting their free descriptions, or placing the second listening step of the experiment after the performance step to build on participants' familiarity with the wind-like sounds may produce further useful findings.

It is argued that this research has evidenced that theatre sound effects are a particularly interesting subject for evaluation. Both sound and performance action, both material interaction and *effect*, the wind machine design is firmly between the clear categories that often characterise experimental evaluation of sound and perception or interface interaction. As discussed previously within Section 8.3.1, the listening steps of the experiment produced some interesting results that suggest the wind-like sounds created through action require the *feel* of that action to be clearly understood as such. While the acoustic wind machine presented an opportunity to control only the sonic feedback element of the interaction, building on previous research in SID (Giordano et al. 2012), the results of the performance step of the experiment suggest that the digital wind-like sound may have nevertheless influenced participants' tactile and kinaesthetic experience of the acoustic wind machine interface itself. This potential connection with previous research on the influence of sonic feedback on haptic or movement perception (DiFranco et al. (1997), Avanzini & Crosato (2006*b*) and Turchet et al. (2013)) has not been explored as part of this research, and should be examined further. As will be explored further in Section 8.3.4, the work presented in this thesis has afforded a deeper examination of potential *continuousness* in embodied interactions with sound.

It is proposed that the "operationalization" (Poepel 2005) strategy pursued for the performance step of the experiment has proven to be an effective way to evaluate the enactive potential of a theatre sound effect design, and this should be developed and expanded upon, particularly as the prototype digital wind machine is developed and refined further. The use of sound stimuli to elicit performances from participants with varied levels of musical experience produced useful results, in particular highlighting the importance of the acoustic wind machine's flywheel mechanism in guiding and creating meaning through those performances. This produced insight into the embodied experience of historical theatre practitioners as they performed wind-like sounds. This finding was obtained from statistical analysis of the acoustical features of the corpus of participants' performances, and underlines the importance of more formal evaluation methods in discovering potential factors that participants themselves find challenging to define or describe. A dedicated performance experiment could allow for a more focused examination of participants' experiences, and help to define the experience of a sound *in-hand* more clearly. Although participants rated the acoustic wind machine as significantly easier to play, the specific reasons for this have not yet been fully captured. It is proposed that subjecting the acoustic wind machine itself to a dedicated study will be an effective method of exploring this further, as the comparative focus of the experiment required both wind machines to be examined in arguably less depth.

Finally, this research has underlined the importance of a qualitative and designerly exploration of theatre sound effects as a way to examine the experience they afford more closely, and thereby refine a strategy for a more formal evaluation. The explorations of the acoustic wind machine and its prototype digital counterpart reported in Chapters 5 and 6 were an important part of informing how best to examine them further in an experiment with participants. It is proposed that this work could be refined further to inform the design of more qualitative examinations of participants' experiences of sound performance, thereby drawing out more insights into the perceptual experience afforded by theatre sound effects. We will now explore the theme of *continuousness* that has emerged throughout the work this thesis.

8.3.4 Capturing Continuousness

What has this research discovered about designing interactions with continuous sounds? Chapter 2 established a frame for the research presented within this thesis, and focused the exploration on the creation of a new *continuous* sonic interaction - one where action reliably, directly and constantly activates and modulates a digital sound. In this encounter with a historical context of creative soundmaking, some new insights and challenges have emerged that have the potential to inform the ways in which designers can facilitate an embodied encounter with the digital. In particular, the work undertaken to capture the acoustic wind machine's enactive qualities digitally revealed unexpected points of transition, or edges, in between some of the stages of *sounds-in-hand* first presented in Chapter 4. These further points of transition have captured more of the *in-betweenness* of the incremental shifts from a continuous *scraping* interaction with crumpled paper through to an open-handed gesture activating a digital sound with sensors. It is argued that examining these shifts more closely will offer more insight into the implications of particular interactive sounding object designs for a performer's embodied experience, and perhaps a broader understanding of how meaning might be created during a sonic interaction. Design conventions around digital sound may emerge more clearly to be examined and critically interrogated.

The acoustic wind machine augmented with a rotary encoder and laser cut gearing, as the bridge between acoustic and digital sound performance, troubled the previous strategies for DMI design introduced in Chapter 2. While some aspects of the design required new connections to be created, other already-existing connections proved challenging to pull apart and isolate. It is argued that this is the inevitable and desirable result of sound created through action, mechanism and material encountering sound created through audio recording, programming, and the creation of connections between discrete elements.

As its cloth was a critical part of the tactile and kinaesthetic experience it afforded, the acoustic wind machine could not produce a digital wind-like sound without also simultaneously producing an acoustic wind-like sound. This required participants in the experiment reported in Chapter 7 to experience both wind-like sounds in headphones, albeit with a continuity of bleed from the acoustic wind-like sound during both interactions. However, this also changed the acoustic wind-like sound from something listened to in a room to something listened to via a single microphone. Similarly, the simultaneous recording process for the acoustic and digital wind analysis and stimuli described in Chapter 4 required careful consideration. Sometimes, the line between acoustic and digital was drawn through the centre of the author herself, with the acoustic sound in one ear of a pair of headphones and the digital sound in the other ear. This particular *in betweenness* may have had a direct influence on the quality of the stimuli produced for the experiment - did the author become a better performer with the acoustic wind machine rather than its prototype digital counterpart, and therefore produce stimuli more weighted towards the progress of that sound? There is more to explore here.

As a wind *effect*, the acoustic wind machine also formed a bridge between perceptual experiences of various categories of sounds. The listening steps of the experiment reported in Chapter 7 presented some of these incremental shifts for participants to examine. The first step attempted to examine perceived similarity between wind *effects* and real-world wind sounds, but by including acoustic and digital sounds within the corpus suggested the potential of a perceptual evaluation space that transitions from natural real-world sounds through to increasingly abstract digital sounds. The second step evaluated the potential perceptual transition between continuous *scrape* sounds of different speeds, and the results highlighted that even within the category of continuously activated and modulated sounds there is a transition point - from sequences of impacts such as Spinotron's ratchet model (Lemaitre et al. 2009) to the continuous scrape produced by the acoustic wind machine. It is proposed that the efficacy of the wind machine in producing the *effect* of natural wind lies in the fragility of human perception when attempting to determine the material sources of sounds (Lemaitre & Heller 2012). However, when the acoustic wind machine transitions to producing a single continuous scrape, the perception of its sound is very different. The programming work in Max/MSP also reflected this transition between a material *scrape* and the complex sound of wind the sound of the prototype digital wind machine was improved by reducing the audible stepping first produced by the configuration of friction models (Chapter 6).

It is argued that the approach taken by this research to extend the hand through very simple real-world sonic interactions to increasingly complex stages of acoustic and then digital interactive sounding objects has been a valuable way to allow perceptual experiences of the acoustic to speak clearly to those of the digital. This has reaffirmed the importance of very simple and intuitive explorations, such as the continuous sonic interaction with crumpled paper described in Chapter 2, in examining and defining human experience of fully embodied sounds. This opens the way for practitioners from beyond the audio programming, interface design and evaluation skillsets of SID and DMI design to explore soundmaking and perceptual experience in a designerly way. Some potential avenues for this future work will be discussed in Section 8.5. How human perception calibrates between some of the transitions mentioned here, for example between an effect and a real-world sound, or between a microphone feed and a digital sound, may be revealed through future work informed by research into multisensory integration, which examines how various sensory inputs are reconciled together (Spence & Squire (2003) and Talsma et al. (2010)). This may allow a deeper examination of the perceptual experience of an *effect*, or indeed how the complex sensation afforded by the acoustic wind machine's crank handle binds to that effect during an interaction. By explicitly moving from the hand *towards* the digital, the tensions between the continuous nature of real-world interaction and the discrete steps and conventions of digital systems have been evidenced within this thesis. The discussion will now shift to reconsider the potential of digital *sounds-in-hand*.

8.3.5 Digital Sounds-In-Hand

Can digital sounds be placed in-hand? Through its exploration of historical theatre sound effects practice, this research has attempted to define and describe the experience of acoustic effects performance and capture this digitally. As has already been discussed within this chapter, this work has produced new and interesting design problems to explore, but the experience of a digital *sound-in-hand* remains elusive. This was evidenced by the results of the experiment with participants reported in Chapter 7, which suggest that the digital wind-like sound was not experienced as in-hand during their performances. However, the insights generated by the work undertaken to model the acoustic wind machine digitally point to the development of a specific design approach to guide the creation of new interactive sounding objects to place digital sound *in-hand* for a performer. These threads are connected here to establish this approach, building explicitly from Franinović's "enactive sound design" (2013, p.21):

- The interactive sounding object should offer a tactile and kinaesthetic *handle* or material to the performer to hold and perform a gesture.
- The interactive sounding object should be created specifically to produce digital sound through human action.
- Digital sound should always be the direct and continuous result of the performer's gesture.
- There should be a continuity between the tactile and kinaesthetic *feel* of the interactive sounding object and the energy perceivable in the digital sound.
- During sound performance, the interactive sounding object should perceivably *disappear* and extend the performer's intent directly towards the sound itself.
- The continuous sonic interaction emerges and develops in a bidirectional way. Digital sound should *fold back* appropriately upon the interactive sounding object and enhance or cohere with its physical qualities rather than conflict with them.

- The interactive sound object should be clearly *enactive*, and teach the perfomer to perform the sound through a process of exploration and rehearsal.

These definitions establish an outline for future work to refine and explore the potential of placing digital sounds in-hand for a performer. With further research, it may be possible to expand on these guidelines - working to create an in-hand experience for a performer using open handed gestures and worn sensors, for example. The next stages of the development of the *sounds-in-hand* framework first presented in Chapter 4 to explore and develop this design approach will be considered further in Section 8.5.

This section has drawn together the various stages of work presented in this thesis and discussed them in relation to the theoretical background presented in Chapter 2. The next section summarises the main findings of this research.

8.4 Summary of Main Findings

The main findings of this research in relation to the research questions will now be summarised.

The first research question was:

- How can historical theatre sound effects be fully investigated as interactive sounding objects?

This research question was answered by taking a transdisciplinary approach to the investigation of available historical sources on the practice that created theatre sound effects in the late nineteenth and early twentieth century. Rather than considering the history of theatre sound effects as a chronology, two frameworks were applied to bring together a clearer picture of practitioners' work as a design practice, and contextualise some specific sound effect designs as potentially meaningful correspondences between actions and sounds that could be realised as digital enactive recreations. The author's own embodied knowledge of soundmaking was also incorporated through a design-led enquiry to remake a working theatre wind machine from historical design instructions. By focusing on this single working example, a fuller investigation of its properties as an interactive sounding object was facilitated, allowing a digital model of it to be programmed in Max/MSP. The sonic interactivity of this acoustic sound effect, and its digital counterpart, were also fruitfully evaluated in an experiment with participants.

This research has established the importance of designing, making and experiential exploration as a way to fully investigate theatre sound effects as interactive sounding objects. The work presented within this thesis has evidenced that historical texts have not fully documented all of the interactive qualities of historical effects that will be of interest to designers. The importance of this design-led approach in producing insights through a simple process of listening and performing opens the way for designers with expertise beyond soundmaking, or creative soundmaking pracitioners like Foley artists, to undertake further explorations of theatre sound effects and thereby produce new knowledge. The *sounds-in-hand* framework produced as part of this enquiry clearly connects very simple real-world continuous sonic interactions to increasingly complex encounters with digital sounds. Theatre sound effects have been established as interactive sounding objects that afford satisfying and simple encounters with self-produced everyday sounds, and as such a way to engage participants and practitioners from within and beyond the fields of SID and DMI design in the creation of new knowledge about human perception of sound and action.

The second research question was:

- How can the enactive properties of historical theatre sound effects be captured in the design of a digital sonic interaction?

This research question was answered through a detailed study of the sonic interactivity of the working acoustic wind machine, which informed a programming approach to modelling its soundmaking elements in Max/MSP using the SDT physical model of friction. The acoustic wind machine was itself used as an interface to perform the digital model, facilitating an acoustical comparison of simultaneously recorded acoustic and digital performances. This helped to develop the digital model, and the configuration also allowed just one element of the multimodal feedback in performance - the wind-like sound - to be examined more closely in an experiment with participants. The results of the design and evaluation work have established the acoustic wind machine as an interactive sounding object that facilitates enactive learning, and a more detailed picture of the components of its soundmaking has emerged.

The work presented here has highlighted the complexity of capturing the *continuous* nature of an acoustic theatre sound effect digitally. Programming within Max/MSP has been deployed as an important method of enquiry, and the practical synthesis approach used here will be of interest to SID and DMI designers working to explore the potential of interactions with physical modelling synthesis algorithms. This research has evidenced the importance of a continuity of energy between human action, interactive sounding object and digital sound when attempting to capture real-world enactive properties digitally. It has also revealed the inherent perceptual complexity of apparently simple constrained interfaces, such as the crank handle at the centre of this thesis. This research will also be of interest to designers and practitioners in the wider field of embodied interaction working to facilitate interaction with sound in other contexts such as VR and gaming in its comparative study of acoustic and digital encounters with performable sounds.

This section has summarised the main findings of this research in relation to the research questions and examined which fields might benefit from the insights generated. The next section looks ahead to potential future work arising from this thesis.

8.5 Future Work

This section discusses the potential future work arising from this research in light of the discussion and findings presented in this chapter.

8.5.1 Soundmaking as Design Practice

This research has framed performative soundmaking as a design practice in its own right, while also rooting prior research in continuous sonic interactions in very simple real-world encounters with objects that produce sound. This has opened the way for potential collaborations between design practitioners without previous knowledge of sound as a material and those working in the more audio-focused disciplines of SID and DMI design. Similarly, the framework of designerly ways of knowing used within this research could help to draw the individual approaches of present-day practitioners, such as Foley artists, together to form a more unified picture of how sonic interactivity is negotiated and implemented in their work. This could help facilitate a closer exchange between SID and the implicit knowledge of soundmaking practices like Foley, perhaps even as far as collaboration in the creation of new interactive sounding objects for digital sound performance. Historical information on theatre sound effects design and performance could be used as a way to guide and direct this work.

Many of the sources gathered and examined in this encounter with a historical site of soundmaking have been manuals specifically written by sound effects practitioners to communicate, and teach, their skills. This knowledge could be applied to the development of a new pedagogical approach to performative soundmaking across acoustic and digital contexts. Some potential areas to explore include:

- How prototyping with materials could be used to design new sounds.
- How an embodied soundmaking practice could be taught and cutivated.
- How abstract instructions might be translated into new designs for new interactive sounding objects.
- How a sound might be designed through an action or mechanism.
- How performances of digital sound effects might be devised and cued.

Such a bodily and interactive sounding object-centred practice could eventually help to open up digital sound design as a medium for exploration by practitioners outside the more usual programming and interface design skillsets of sound and music computing. For example, this research has highlighted the potential of a crank handle as an interesting and rich control interface for the performance of digital sound. This could be an interesting way to expand on the data mapping used within this research and bring some digital synthesis methods more usually used in music technology settings to practitioners more familiar with acoustic soundmaking, like Foley artists. One particularly interesting area of digital synthesis to explore is Frequency Modulation (FM) synthesis. A single FM sound can incorporate several oscillators together, with control available over the amplitude envelope of each. These may lend themselves well to a similar approach of creating many data streams from one rotary encoder.

8.5.2 Developing the Prototype Digital Wind Machine

The results of the experiment reported in Chapter 7 confirm that the prototype digital wind machine is not yet perceptually similar enough to its acoustic counterpart to be used as a substitute in a future experiment. Some potential improvements to the program in Max/MSP have already been proposed in Chapter 6 and work should be undertaken to explore these further. More complex implementations of the digital waveguide approach used in this research to create the cloth model, such as waveguide networks (Serafin et al. 2002) or meshes (Murphy et al. 2007) could prove useful in capturing how the irregular undulations of the cloth produces those resonances specifically responsible for the acoustic wind machine's characteristic whistling. The Max/MSP program produced as part of this research has required careful management of DSP due to a full twelve friction models being implemented within it. The efficiency of the prototype may be improved further by moving on from Max/MSP to another programming language such as C++. This could help to expand on the programming approach developed in this research through the eventual creation of dedicated audio software for sound effects creation with physical models.

Alternatively, the use of a an embedded system rather than a computer might also improve the real-time efficiency of the digital model, and could also allow for graphical programming using the Pure Data environment (Moro et al. 2016). An embedded system could also develop the prototype digital wind machine into a distinct interactive sounding object in its own right, and free it from the acoustic wind machine on which it is based. This would involve the fabrication of an enclosure to house the embedded system hardware and facilitate control with a digital crank handle. It could also be designed to directly output sound through an embedded speaker system (Berdahl 2014), bringing the experience of using the digital model closer to that of the acoustic wind machine. Removing the digital wind machine from the constraints of its current acoustic interface would allow many different qualities of wind-like sound to be produced, and a fuller exploration of the sonic potential of a physical model of friction when activated with a crank handle.

The way the acoustic wind machine's motion is captured could be improved further with a higher resolution encoder and metal gearing. This would help to make the connection between the mechanism and the program in Max/MSP even more robust. The very simple digital crank controller developed as part of this research should be augmented further with the addition of haptic feedback to add some of the resistance of the acoustic wind machine's handle. This could be achieved through the use of a motorised rotary encoder rather than the passive component used as part of the first prototype presented in Chapter 6. With the right choice of motorised rotary encoder, it could be possible to prototype different kinds and levels of haptic feedback and explore their perceptual potential. This kind of prototyping work has already been used in product design for creating haptic dials for electrical appliances (Kim et al. 2008).

Work could also be undertaken to create a simple flywheel mechanism to add some of the sensation of movement of the acoustic wind machine to the digital crank handle. This could proceed much like the approach used to create the Gyrotyre (Sinyor & Wanderley 2005) DMI introduced in Chapter 2, where a spinning bicycle wheel provided vibrotactile feedback for the performer through the wheel's rotational inertia. The fully digital crank controller is no longer bound by the innate limitations of the acoustic wind machine's crank-and-axle mechanism, and so it could be possible to design a new crank handle-based interface with further movement potential. For example, being able to push the handle forward and back longitudinally while rotating it might give an interesting additional dimension of control. While this research has reused crank handles manufactured for other purposes, such as table winding and meat grinding, there is the potential of manufacturing a new design specifically for the performance of sound. This could take new requirements into account, such as changes to the shape of the handle itself, or accommodate some of the augmentations previously suggested here. Recent developments in 3D printing technology, which allow components with moving parts to be printed (Calì et al. (2012) and Krassenstein (2015)) could facilitate this.

More broadly, the approach of programming of a *prototype* to generate further design problems as part of an enquiry could be developed as part of future research within SID and DMI design. Building on the use of the SDT physical model of friction, work could also be undertaken to create a tangible computing environment for prototyping and investigating sound effects. The potential of physical modelling synthesis as a material for digital sound design through performance could be developed further. Particular sound effect design could also be investigated as enactive recreations, while also focusing on modern fabrication methods and the mechanical qualities of performance interfaces. This would contribute to developing scholarship that takes a craft approach to the design and creation of DMIs (Armitage & McPherson 2018).

8.5.3 Exploring Sounds-In-Hand

As discussed in Section 8.3, the *sounds-in-hand* framework produced as part of this research focused on one particular transitional stage between the performance of an acoustic wind-like sound and a digital wind-like sound, both using the wooden acoustic wind machine as an interface. This approach has revealed many new transitional shifts in the stages from fully acoustic to fully digital soundmaking, opening up several other potentially interesting areas of enquiry into human perceptual experience of sonic interactions that have not been fully explored within this thesis. These transitions can be added to expand the original framework and produce a template for further design and evaluation of continuous sonic interactions with a wind-like sound (Figure 8.2). These additional transitions will now be considered in more detail.

- *A: From scraping to wind-like:* The historical enquiry presented in Chapter 3 showed that theatre practitioners were aware of a connection between simple scraping sounds produced with hand actions and materials, and the more complex wind-like sound of the acoustic wind machine, produced through a combination of several wooden slats rubbing a cloth. However, the results of the second step of the listening experiment reported in Chapter 7 showed that the speed of the rotational gesture was not consistently perceivable in the resulting acoustic wind-like sound. However, the results of the performance step of the experiment showed that participants found the acoustic wind machine easy to play. As previously explored within this chapter, this suggests that the perceptual connection between action and sound may not need to be so simple and direct to create meaning for a performer during an interaction. The perceptual transition between a simple scraping sound and an acoustic wind-like sound should be investigated further to establish if an effect produced with materials could link the imitation of a sound-producing action to sound tracing describing a complex sound. An initial step in this direction arising from this research could be a listening experiment to evaluate and compare the perceived speed of a group of scraping and a group of acoustic wind-like sounds, as could a performance experiment to elicit scraping sounds and wind-like sounds of different speeds from participants.

- *B: From acoustic mechanism to digital crank interface:* Chapter 5 explored the perceptually rich experience of performance with the acoustic wind machine's crank handle. How the sensation of movement and resistance it offers might be captured in the design of a fully digital crank interface has already been considered in this chapter. Evaluation could be an important part of the design of this new interface. The experience of performing a digital wind-like sound with the acoustic wind machine, and then with the digital crank interface, could be compared to examine how much richness could be produced through haptic feedback, for example. The two interfaces could also be compared to examine how the perceivably large and slow performance gesture of the acoustic wind machine could be captured digitally.

- *C: From freely moving to static handle:* A further dimension of the experience of

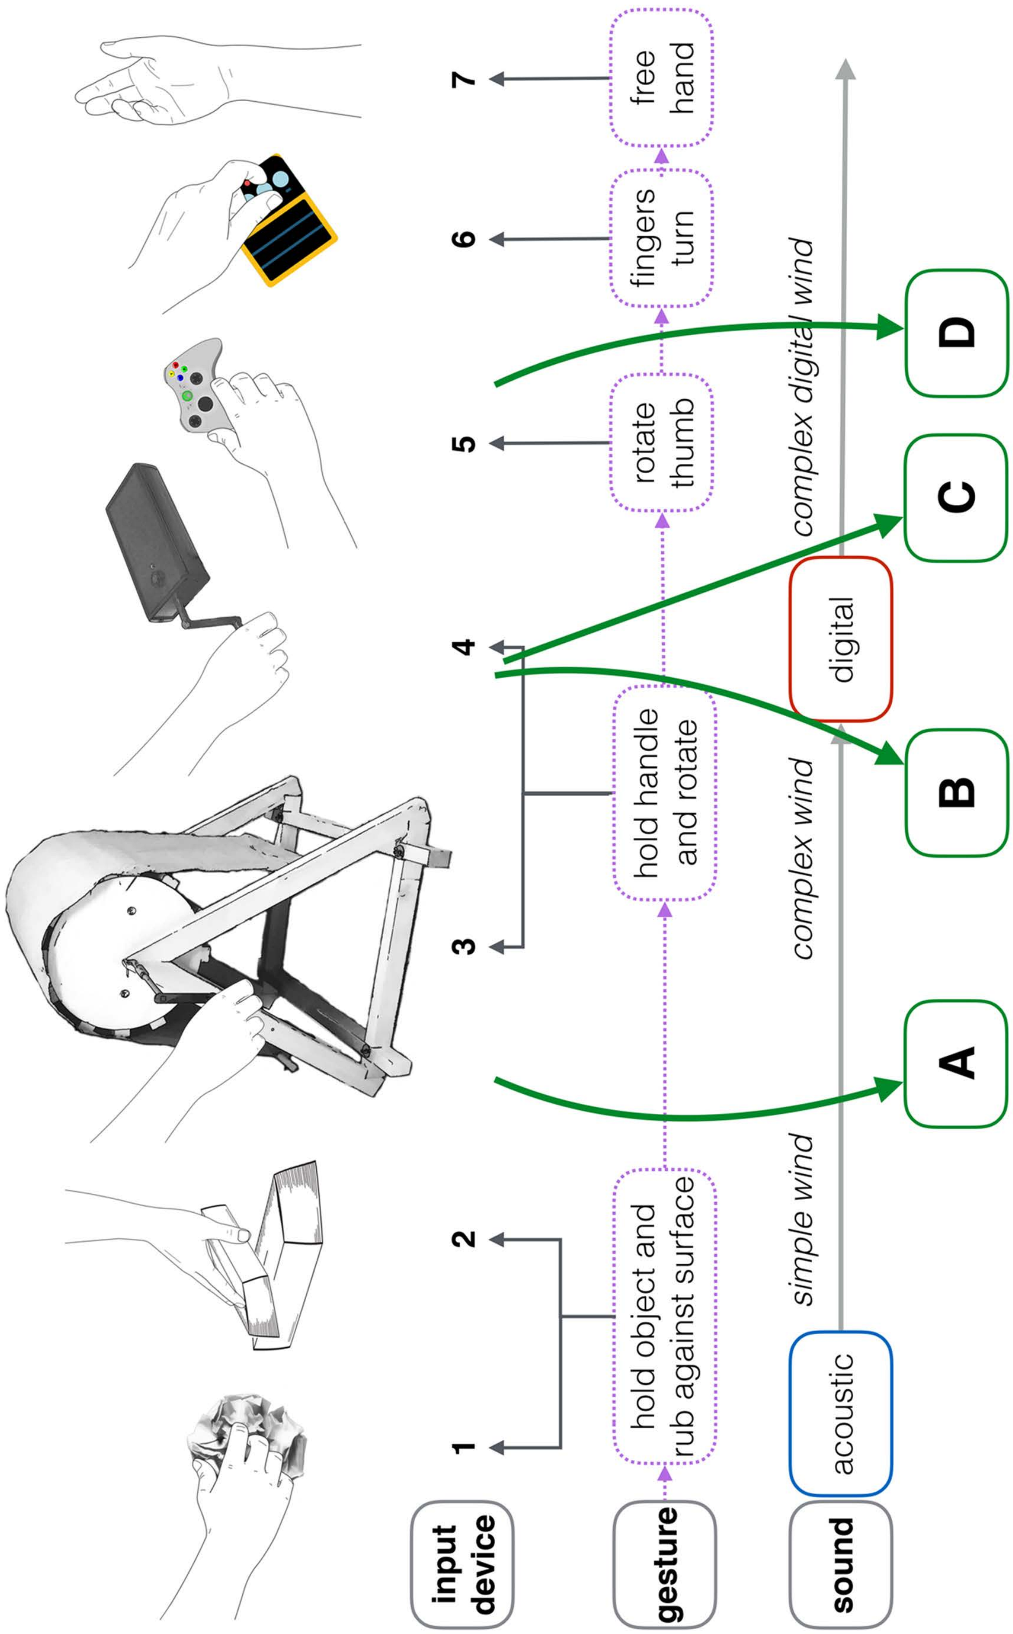

Figure 8.2: Expanding on the *sounds-in-hand* framework introduced in Chapter 4 to highlight some additional transitions for evaluation.

performance with the crank handle was revealed in the course of this research. The freely moving wooden handle of the crank on the acoustic wind machine continuously reoriented the hand during performance, allowing it to move more fluidly (Chapter 5). By contrast, the smooth metal crank used to create the first iteration of the digital crank interface did not allow for this kind of movement (Chapter 6). A comparison between the crank with and without a moving handle in an experiment would allow the perceptual importance of this additional kinaesthetic feedback to be investigated more fully. This could provide useful information to inform the design of a new 3D printed crank interface for sound performance.

- *D: From hand to fingers:* The work to build the acoustic wind machine revealed the importance of simple mechanical considerations in performance interface design (Chapter 5), adding a further dimension to the previous research into the design of interfaces for digital sound and music performance discussed in Chapter 2. For example, the distance of the crank handle from its pivot point had a profound effect on the level of effort required to perform a rotational gesture with the acoustic wind machine. This finding could be usefully applied to the augmentation of interfaces for sound performances that follow established music technology conventions. A useful first step in this direction would be to compare the perceptual experience of sound performance with the smooth crank handle interface to that of sound performance with a small crank held between the fingertips to activate a rotary encoder.

This approach, which establishes useful points of comparison between soundmaking contexts, could be fruitfully applied to the performance of other kinds of acoustic or digital sounds. In particular, this framework allows findings from investigations into acoustic modalities of soundmaking and interface design to be applied to related digital modalities. This could have great potential in future evaluations of other theatre sound effects, and also help to draw insights gained from simple intuitive encounters with acoustic materials towards new and enhanced ways to encounter acoustic and digital soundmaking.

This section has described some potential avenues for future work arising from this research. This chapter, and this thesis, will now be concluded.

8.6 Conclusion

This chapter has revisited the main questions and methodology of this research, and discussed the findings in light of the theoretical background underpinning this work. This research has clearly established late nineteenth and early twentieth century theatre sound effects design and performance as a practice centred on the design of continuous sonic interactions, and as such an interesting and rich source of knowledge for present-day research in SID and DMI design. It has also reframed performative soundmaking as accessible, intuitive and simple in its multisensory complexity, creating a space for real-world interactions to be explored by a broader range of practitioners without prior knowledge of audio recording and programming.

The work presented in this thesis has evidenced how a transdisciplinary approach can pose interesting new design problems for the development of digital interactive sounding objects. While the enactive qualities of the acoustic theatre wind machine, and the compelling nature of its sonic illusion, have been conclusively proven here, not all of its features have been fully captured digitally. The design and evaluation of the prototype digital wind machine have been instrumental in bringing the importance of the acoustic wind machine's mechanism to light however, and this work has opened up new dimensions to explore in the design of interactive sounding objects for performative acoustic and digital soundmaking.

Appendix A

Wind Machine Descriptions

This appendix contains the full text descriptions of the wind machine designs discussed in Chapter 5 of this thesis.

A.1 Jean-Pierre Moynet (1874)

From Moynet's "L'Envers du Théâtre" (2015):

> There are several kinds of machine in use to imitate the sound of wind, but this is the most usual: a solid framework supports the axle of a cylinder resting upon two trunnions. The cylinder is made up of individual sections each of which in cross-section looks like the tooth of a cog wheel, and form projections on the surface of the cylinder. Generally, there are between fifteen and twenty of these ridges fastened to the turning cylinder. Strong silk fabric on top of the frame runs over the cylinder and small bolts, which can be tightened as required, allow it to be tensioned. When the handle of the cylinder is turned, the friction of the silk over the ridges produces a continuous sound exactly like that of wind whistling in chimneys or corridors (Baugh & Wilmore 2015, p.169).

A.2 Van Dyke Browne (1913)

From Browne's "Secrets of Scene Painting and Stage Effects" (1913):

> Stage wind is produced by means of a large drum or wheel made of slats of wood, with about three inches of space between the slats. This drum is

fastened to two uprights and has a large handle attached to the centre, so that it can be turned round and round very easily. Above the drum is fixed a stout rod, from which hangs a large piece of moiré silk. This silk hangs over the drum, and when the drum is turned the sound made by the slats of wood touching the silk produces the sound of wind. The illustrations show how the appliance can be made (Browne 1913, p.70).

A.3 Garrett Leverton (1936)

From Leverton's "The Production of Later Nineteenth Century Drama: A Basis for Teaching" (1936):

> The various quotations would seem to indicate that the method of producing the wind sound was almost standard. Phonograph records can be purchased today to give wind effects, but the majority of stages still use the drum and silk. The construction of such a piece is very simple and the effects are astounding. A little practice enables the operator to produce winds of various degrees of intensity and pitch. The principle involved is that of the friction of wood on cloth. A cylindrical frame is made and a piece of two-inch by one-inch batten fastened to the edge of this cylinder. A piece of cloth is made fast over the cylinder. Different kinds of cloth will vary the sound of the wind produced. A crank is provided at one end for turning the cylinder in its frame (Leverton 1936, p.50).

A.4 Frank Napier (1936)

From Napier's "Noises Off," first published in 1936 (1962):

> It consists of a wooden drum, composed of two circles of wood connected by strips, the outer corners of which have been chamfered. The drum is mounted to rotate on a wooden stand, and a strip of sail-canvas is stretched over it. A handle is fitted to the end of the drum's axle. When the drum is turned, the chamfered strips rub the canvas and make it shriek. In the construction of this machine, particular attention should be paid to the joints of the stand and the axle-bearings. Constant and violent use puts great strain on these,

so that they work loose in time and add unwanted bumps and squeaks to the sound of wind. It is well to use bolts at the joints, which can be tightened as "play" develops (Napier 1962, p.51).

Appendix B

Prototype Digital Wind Machine Programming

A guide to the coding work in the Arduino IDE and Max/MSP that created the final prototype digital wind machine. This appendix supplements the discussion and work presented in Chapter 6.

B.1 Arduino Program

This simple program reads data from the rotary encoder connected to pins A2 and A3, and sends it via the computer's serial bus to be received by Max/MSP (Figure B.1). As detailed in Chapter 6, the rotary encoder is turned by the laser cut gearing coupled to the acoustic wind machine's A-frame and cylinder.

B.2 Max/MSP program

This is a full guide to the main components of the program in Max/MSP. These have been discussed in detail in Chapter 6, but are presented here to give the reader further insight into how the processes and connections have been implemented in this programming environment.

Section B.2.1 presents the main GUI of the program and describes how the signal from the digital slats is processed. Section B.2.2 discusses how the rotary encoder data is connected to the program. Section B.2.3 outlines how the inertia model is implemented to shape the perceptual *shape* of each digital slat scrape. Section B.2.4 describes the layers

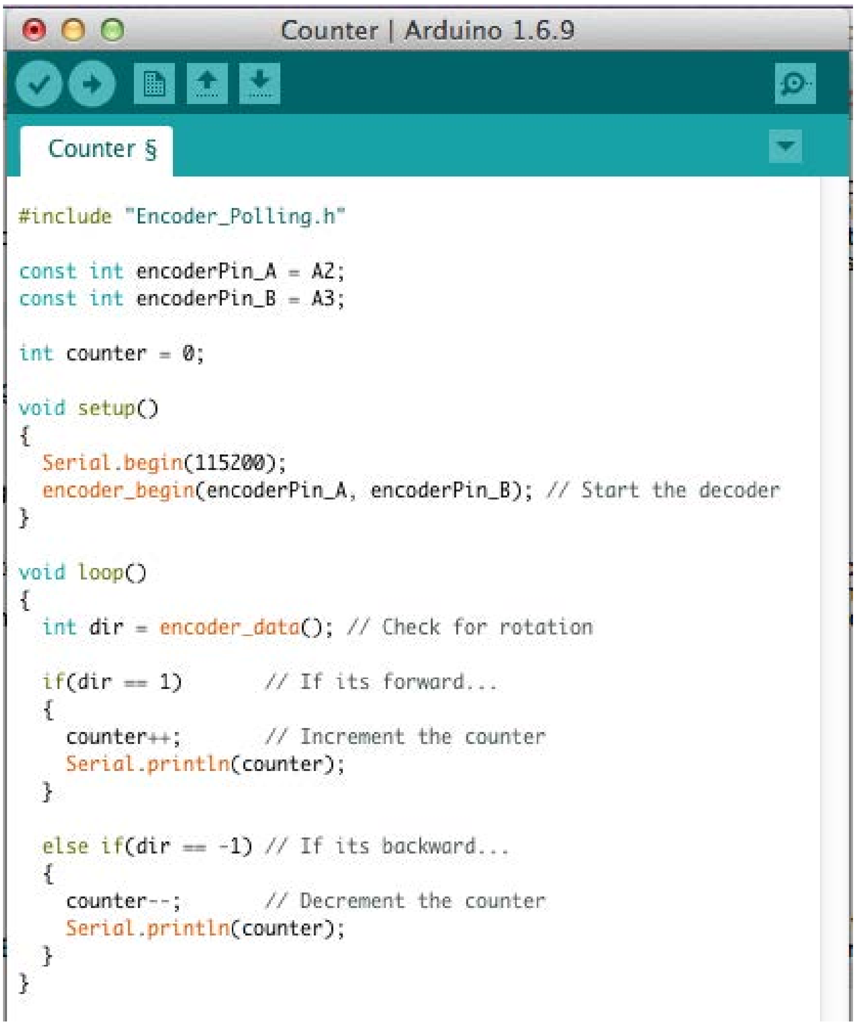

Figure B.1: Rotary encoder program in the Arduino IDE.

that make up the digital slat model, creating and shaping the real-time digital slat sound with the SDT friction model.

B.2.1 Main GUI and Master Processing

This is the program GUI in running or presentation mode (Figure B.2). The main functions of the program that need to be configured for performance, or that give information about important functions, are displayed here. These are made visible in *subpatches*, the inner workings of which will be described here.

The top left of the patch displays the rotary encoder data being recorded from the main slat, which tracks with the acoustic wind machine's crank handle. To the right of this display are the readouts for rotation count, degrees, degrees per second and RPM, which are calculated directly from the rotary encoder's data stream. The dial display shows the real-time position of the crank handle.

Each [*poly~*] object contains one digital slat model, and each sends its audio to one of the twelve faders at the bottom of the GUI, allowing for quick troubleshooting. The top right of the patch allows the serial port for the Arduino (the USB connection) to be selected, and clicking *begin* starts data from the rotary encoder flowing into the patch. Sound only begins when the crank handle is moved, however.

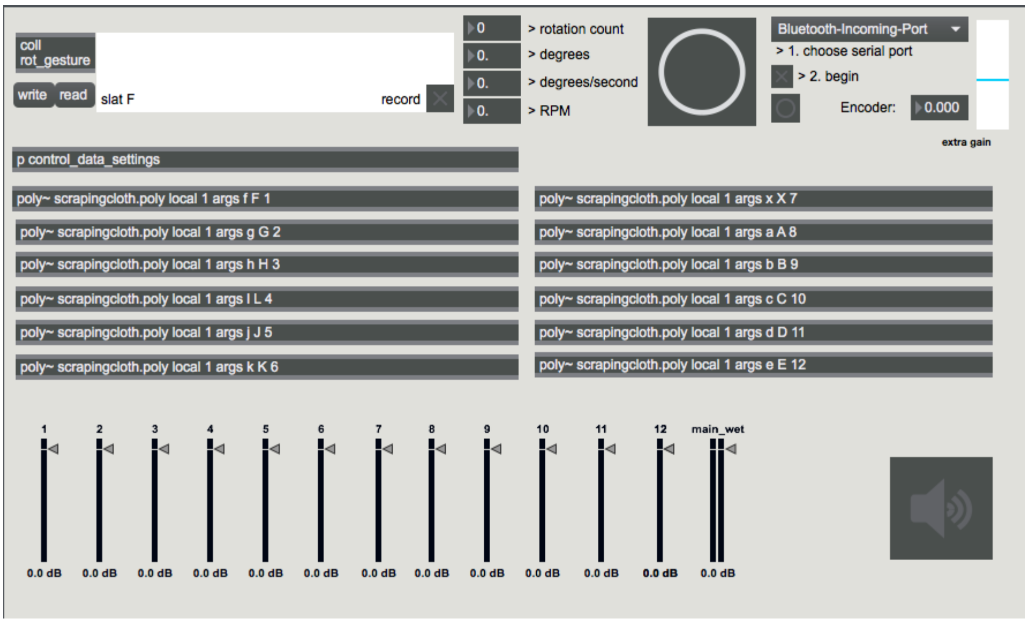

Figure B.2: Graphical User Interface (GUI) in Max/MSP.

Opening the program in patching mode (Figure B.3) shows the master processing on the summed output of the slat models. First, the friction sounds produced by the slat models passes through [*ftk Shriek2Dirs~*], the main model of the acoustic wind machine's cloth. Then [*mdePowerPanning~*] pans the sound of the slat models to the centre of the stereo field. As discussed in Chapter 6, this was to give participants in the perceptula experiment reported in Chapter 7 the impression that the digital wind-like sound was coming from the direction of the crank handle. Then, the [*vst~*] object loads a simple compression plugin to control the dynamics of the digital wind-like sound so it does not become uncomfortably loud for participants.

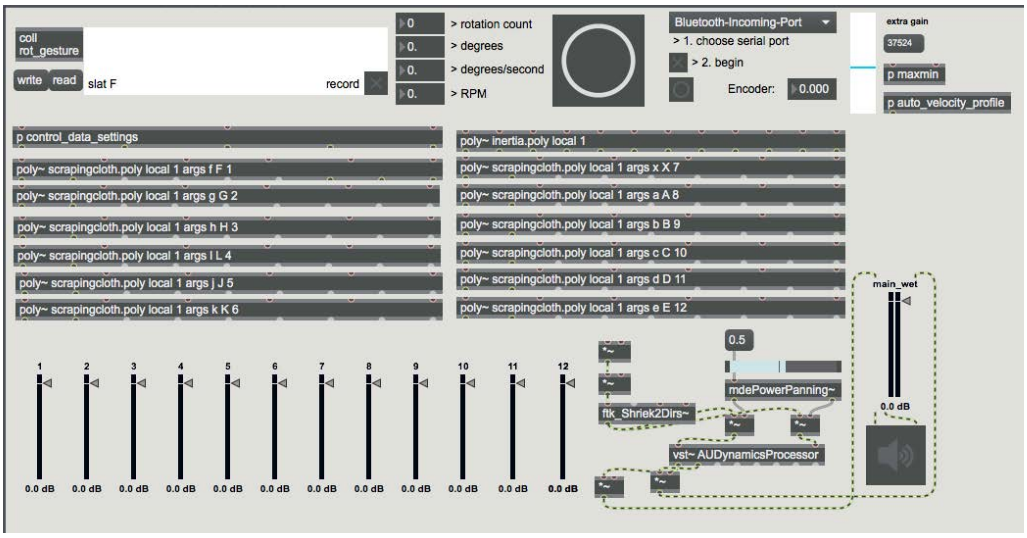

Figure B.3: Master Processing in Max/MSP.

Inside [*ftk Shriek2Dirs~*], the components of the global cloth model can be seen (Figure B.4). To the left of the patch are the components that model the tight and loose sides of the acoustic wind machine's cloth, as discussed in Chapter 6. To the right, data from the main digital slat is smoothed and configured to modulate the [*tapout~*] delay objects according to the crank handle's position in rotation, which recreates some of the acoustic wind machine's characteristic whistling in performance.

B.2.2 Rotary Encoder Data Processing

This patch reads data from the rotary encoder sent over the serial bus by the Arduino program previously described in this section (Figure B.5). Its *choose serial port* and *begin* controls are exposed to the main GUI to allow the rotary encoder data to be activated. The main object in this patch is the [*serial*] object, which reads the Arduino data. The objects below it parse the packets of data into a steady stream of numbers, which is all that is needed for the program.

This is the internal view of the top left of the main GUI in Max/MSP (Figure B.6). The rotational data sent out from the previous patch is send here, and each full rotation of the encoder is counted. The data is then mapped to a continuous rotation of 360.0°, and smoothed by the [*dot.unwrap*] abstraction to remove any sudden jumps in values, which will be heard in the sound of the friction models. From this smoothed rotation, the visual dial display is calculated. Velocity, acceleration and RPM are calculated from this data to be displayed in the main GUI. To the bottom right of this patch, the [*coll*] object records the main rotational data stream (tracking with the crank handle) and the velocity and acceleration of the encoder.

This is the control data settings subpatch from the main GUI (Figure B.7). The velocity data from the rotary encoder and the degree position of the main slat in rotation are brought in here and scaled before being sent out to control the main amplitude of the summed slat models. The velocity data modulates the amplitude to ensure lower amplitudes at slower speeds, and the degree data is tracked to ensure that when the handle stops, the sound also stops. This is achieved by tracking changes in degree values and then triggering a [*line*] object, which ramps the main amplitude up or down as needed. The [*initialise*] subpatch sets up a chain of processes to activate downsampling and assign values to the digital slat models within each [*poly∼*]. The [*settings*] subpatch contains the initial numerical values for the parameters each SDT object in the program, as detailed in Chapter 6.

B.2.3 Logarithmic Filter (Signal Inertia Model)

The main [*poly∼ inertia*] subpatch shown in the main GUI runs twelve individual signal inertia models to logarithmically filter the data stream to each slat. As the signal inertia model runs at signal rate, encapsulating these within a [*poly∼*] object allows them to be downsampled just like the slat models, improving the efficiency of the overall program in terms of its demand on the computer's CPU.

The logarithmic filter itself was developed from code by jvkr on the Cycling '74 forum (Figure B.8). Data from each digital slat's individual rotational path (parsed from the main 360.0° data stream) passes through and is converted to a signal by the [*line∼*] object. The [*tapin∼*], [*tapout∼*] and [*slide∼*] objects slightly delay and logarithmically smooth the signal, changing the gestural shape that activates each digital slat. The [*snapshot∼*] object

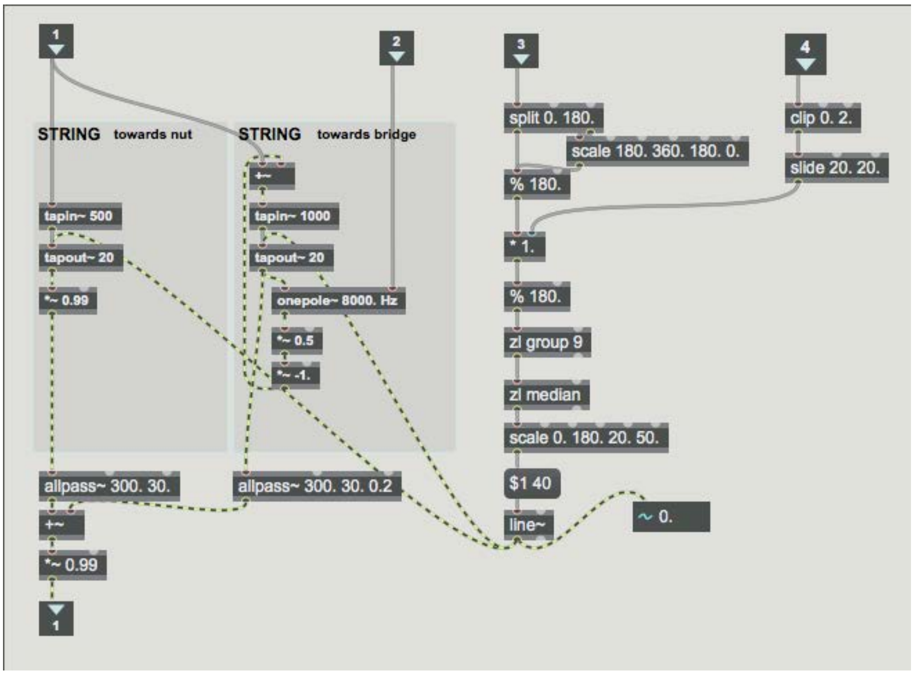

Figure B.4: Global Cloth Model in Max/MSP.

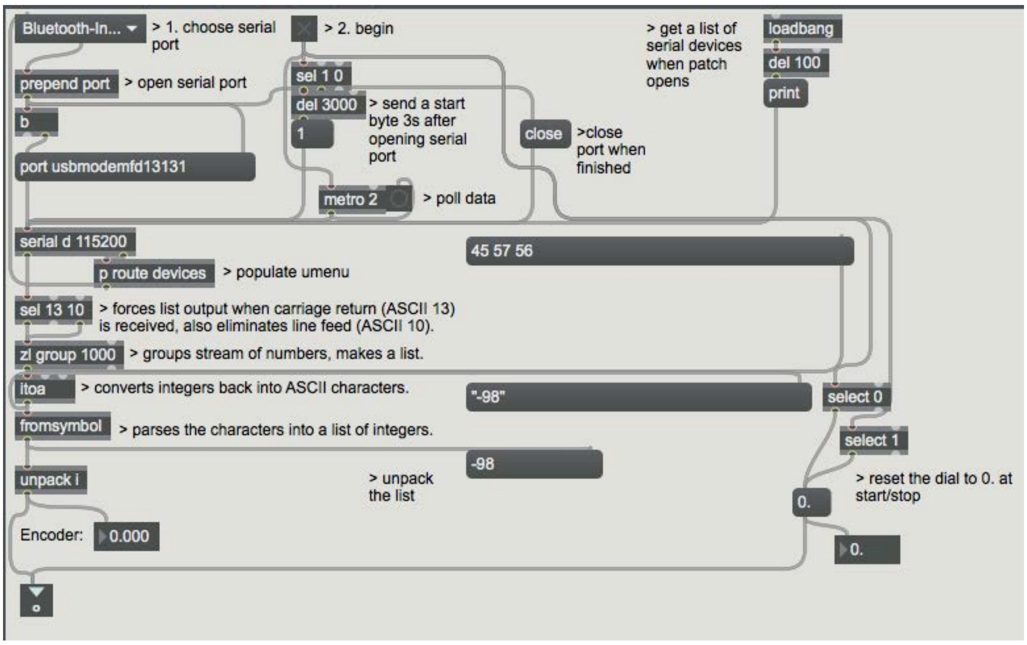

Figure B.5: Rotary Encoder Data to Max/MSP.

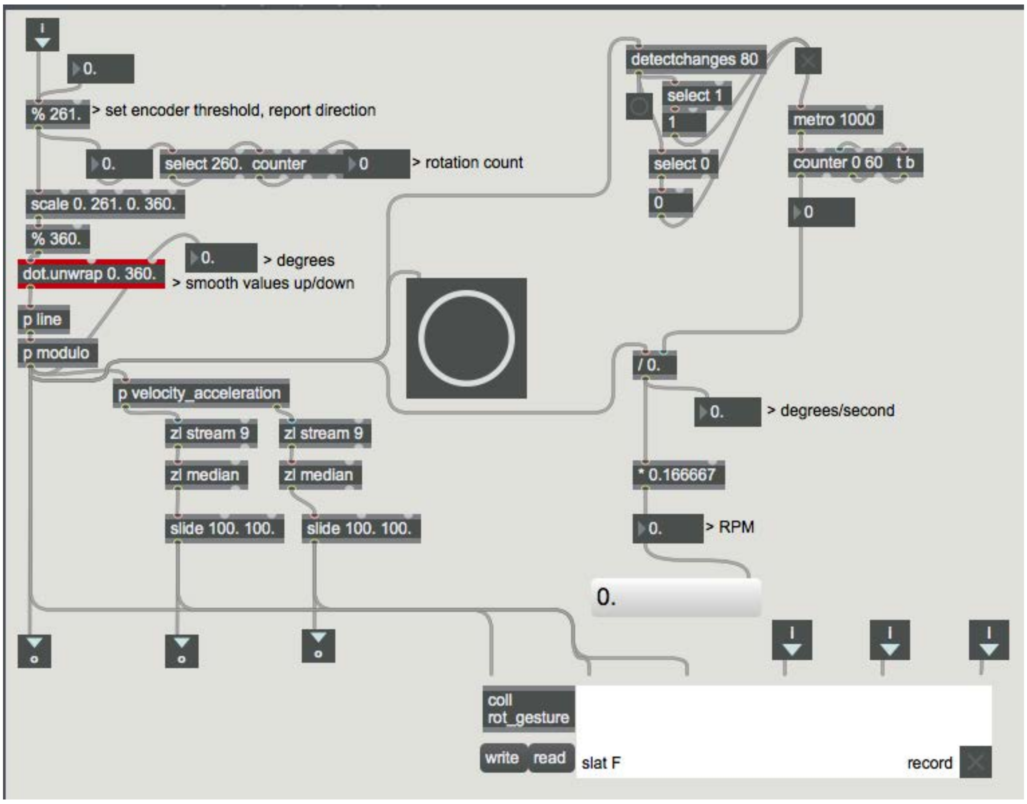

Figure B.6: Rotational Smoothing and Data Recording in Max/MSP.

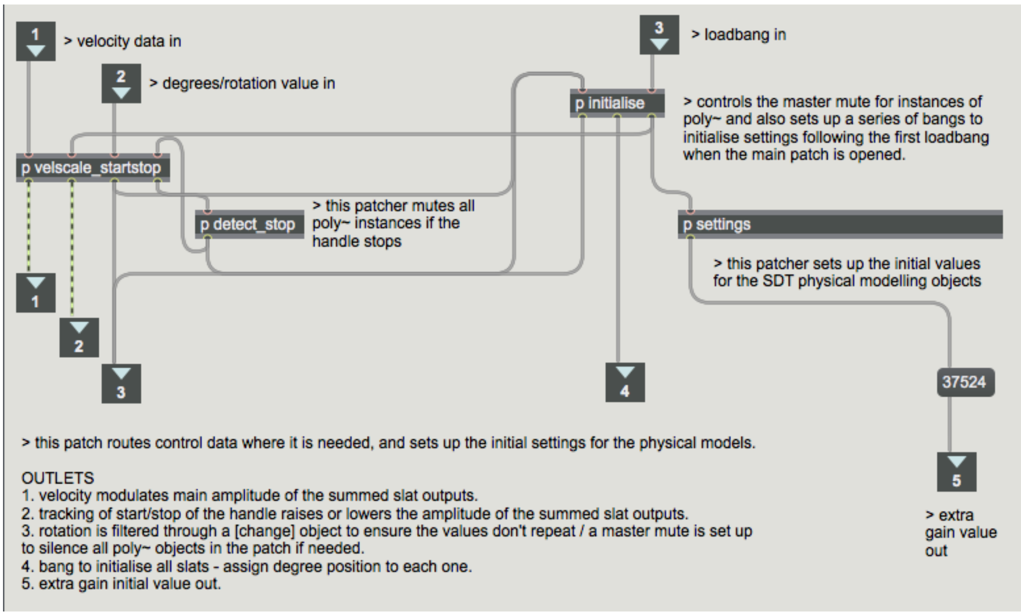

Figure B.7: Control Data and Settings in Max/MSP.

then converts the signal back to data, and passes it out to the velocity parameter of the digital slat it is connected to.

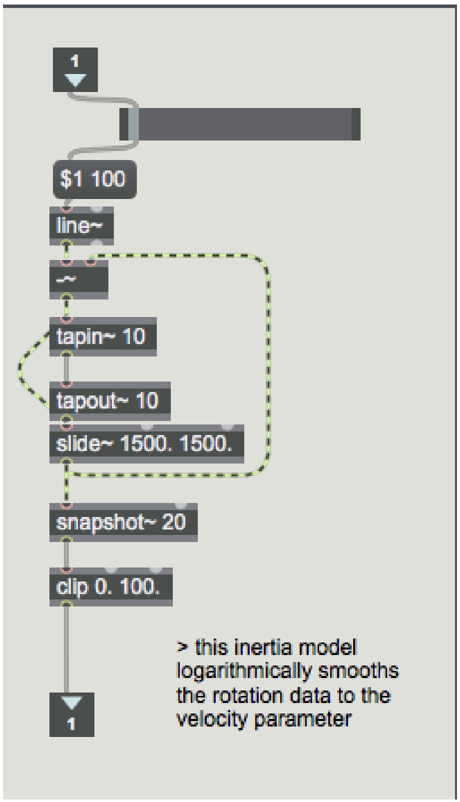

Figure B.8: The Logarithmic Filter/Signal Inertia Model in Max/MSP.

B.2.4 Digital Slat Model

Figure B.9 shows the inside of each [*poly∼ scrapingcloth.poly*] implemented within the program. Each of these objects creates the initial position of the digital slat within the 360.0° circle, passes real-time parameters to the SDT friction objects, and activates their sound when the slat should be in contact with the cloth. Down the left side of this patch, the rotational data is shaped to offset the amplitude of the digital slat during the second half of the rotation when the cloth is looser and exerting less frictional force on the cylinder. The overall velocity of the rotary encoder also scales this resulting value. Data then passes out of this object to be shaped by the inertia model described in Section B.2.3.

The [*selector*] subpatch, to the top right of [*poly∼ scrapingcloth.poly*], passes an appropriate degree value to the slat model to allocate it one of the positions of the slats on the acoustic wind machine (Figure B.10). These are selected depending on a number given to the [*poly~*] when it is set up. The values correspond to the degree values on the acoustic wind machine's slats, as described in Chapters 5 and 6.

The [*slat trajectory*] subpatch, to the top left of [*poly~ scrapingcloth.poly*], tracks the position of the digital slat as its degree value changes throughout the rotation, activating and muting the sound of the digital slat as its counterpart passes in and out of range of the cloth on the acoustic wind machine. This ensures that, like on the acoustic wind machine, only seven digital slats are *in contact* with the cloth and producing sound at any one time.

The probe width and force parameters to the SDT objects are modulated in real time, according to the degree position of each digital slat and acceleration of the main slat respectively. This is achieved by combining a simple [*cos fuction*] (Figure B.12) with an initial value (Figure B.15). Some phase shifting is also added to the [*cos fuction*] with a [*random*] object.

The SDT objects to create the friction for each digital slat are implemented within the [*ftk SlatScrapeRT~*] subpatch. This takes initial values upon startup, and then allows activation and modulation data to pass through.

Finally, some dispersion is added to each digital slat with an instance of [*ftk Dispersion~*] (Figure B.14). This is set up to allow some very slight modulation according to the individual degree position of the digital slat and its real-time velocity parameter after filtering by the inertia model.

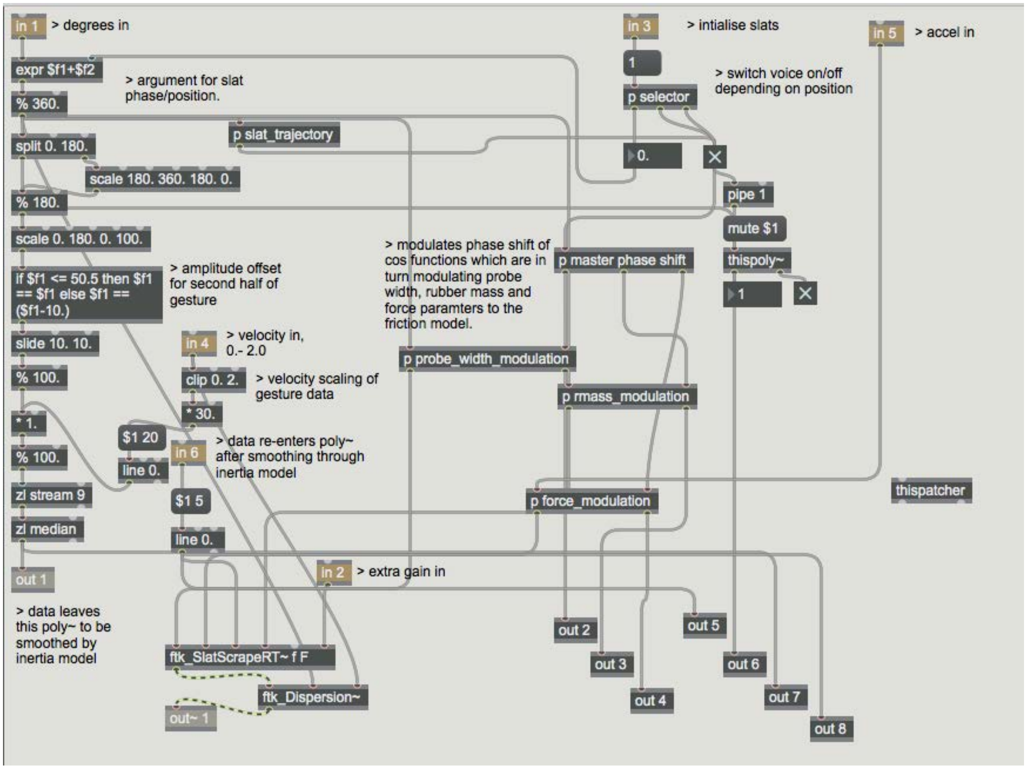

Figure B.9: The Digital Slat Model in Max/MSP.

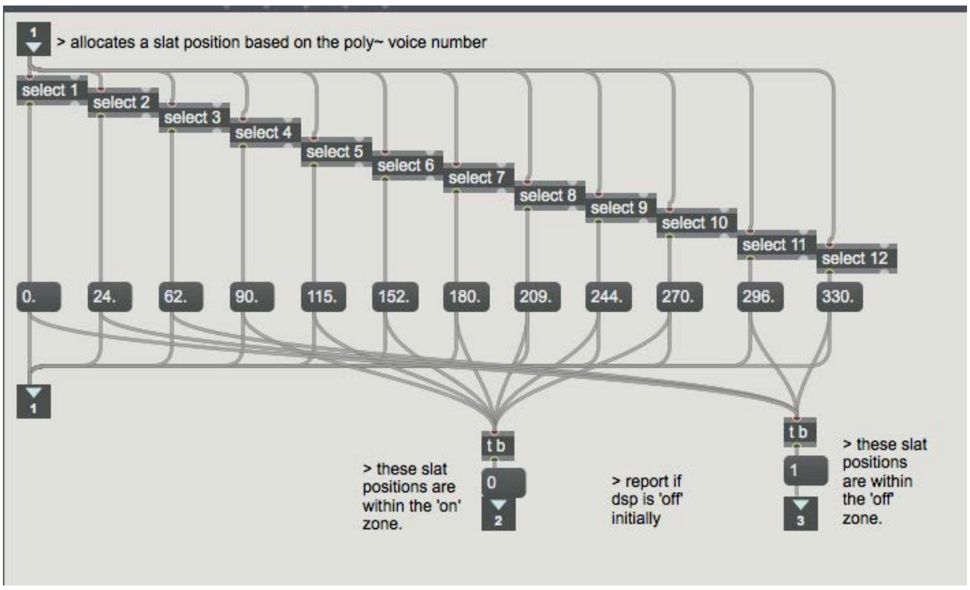

Figure B.10: Allocating a Digital Slat Position in Max/MSP.

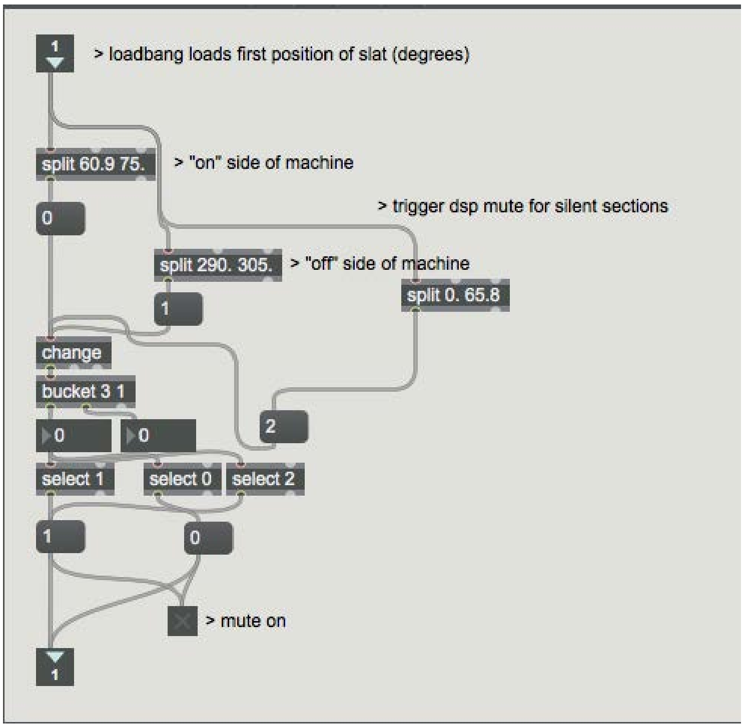

Figure B.11: Activating and Muting the Digital Slat in Max/MSP.

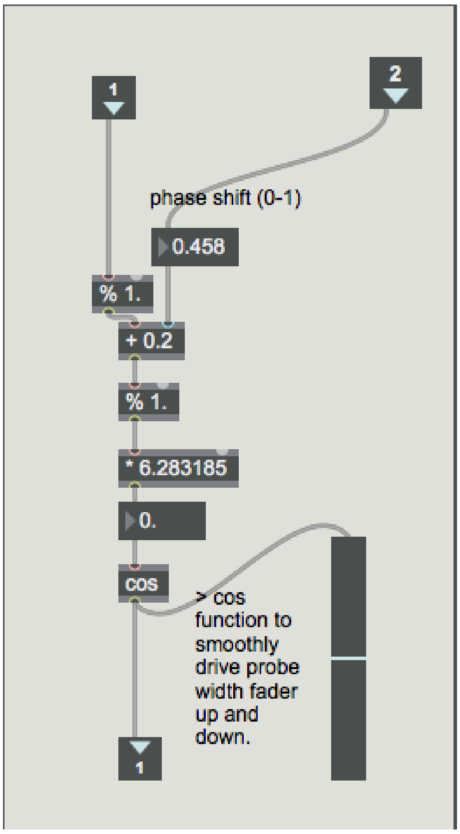

Figure B.12: Cos Function for Modulation in Max/MSP.

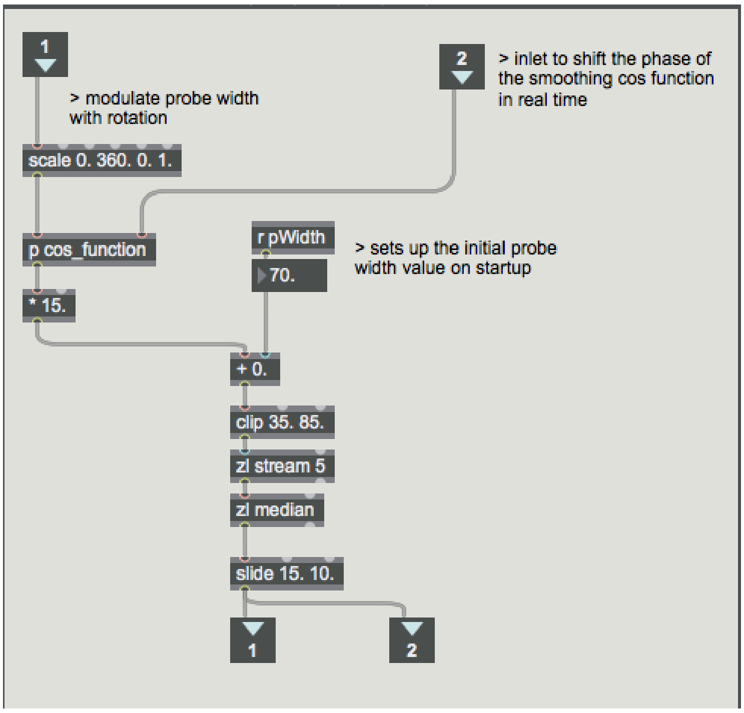

Figure B.13: Real-Time Modulation of Probe Width Parameter in Max/MSP.

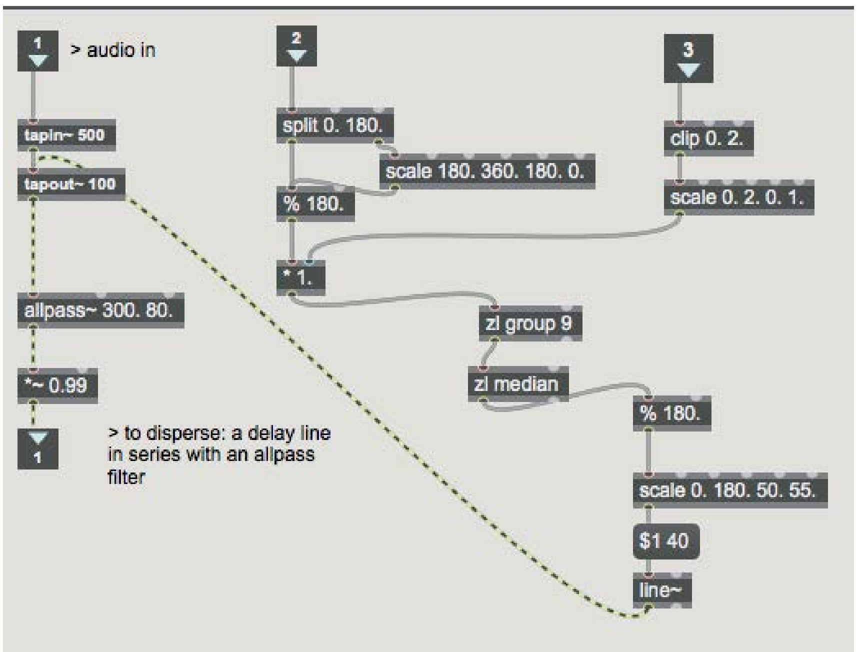

Figure B.14: Dispersion for Each Digital Slat in Max/MSP.

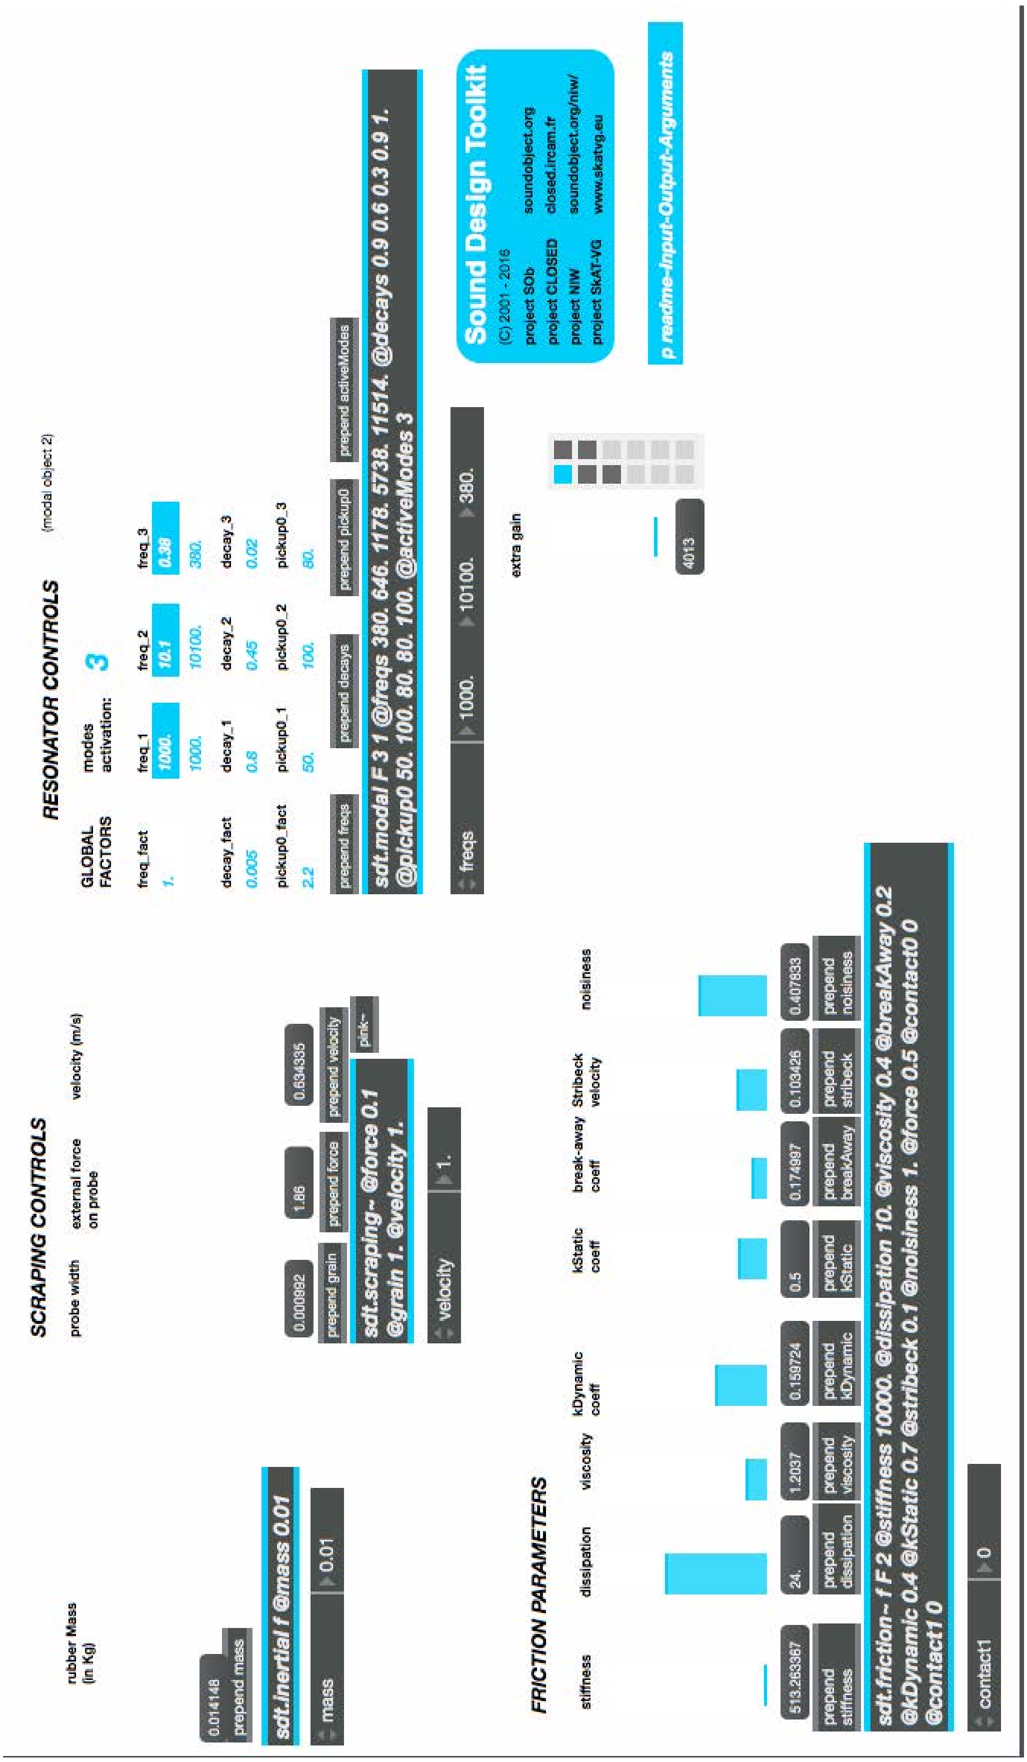

Figure B.15: The SDT Friction Model in Max/MSP.

Appendix C

Acoustical Analysis

Chapters 5 and 6 of this thesis present acoustical analysis and comparisons between real-world natural wind and the acoustic and prototype digital wind machines. This appendix expands on that discussion to include the full results of the acoustical comparison between the acoustic wind machine and a natural wind sound, and also each iteration of the analyses that informed the calibration of the final prototype digital wind model in Max/MSP. The audio recordings used for these analyses are included on the USB media drive that accompanies this thesis.

C.1 Comparing the Acoustic Wind Machine to Natural Wind

The acoustic wind machine was performed to imitate a natural wind sound with an exaggerated and frenetic gesture, and the results recorded, before being analysed with the MIR toolbox in Matlab. The process of performance, and a comparison of the amplitude envelopes of the two resulting sounds, were reported in Chapter 5. The results of the acoustical analysis are fully reported here.

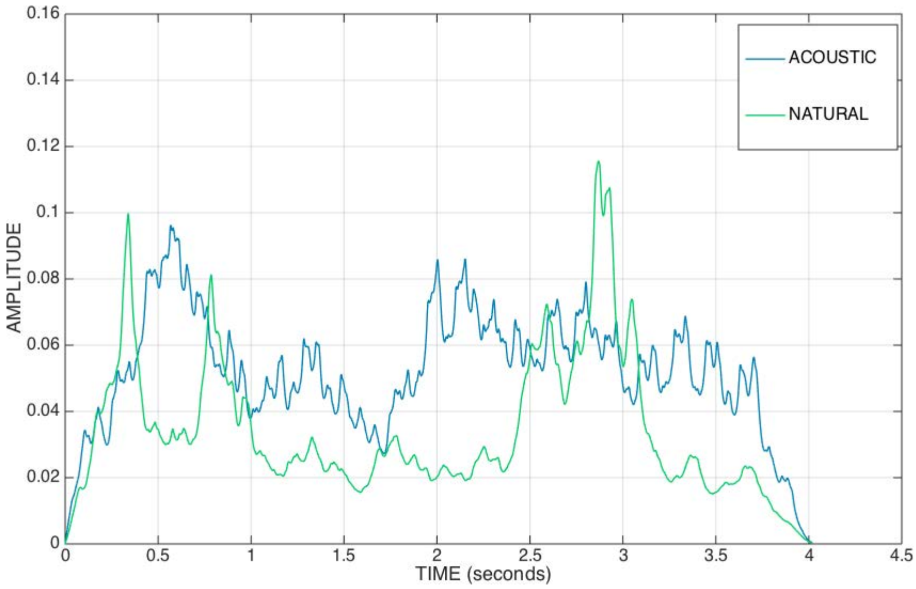

Figure C.1: Amplitude envelopes for the same gesture with the acoustic wind machine (blue) and the natural wind (green).

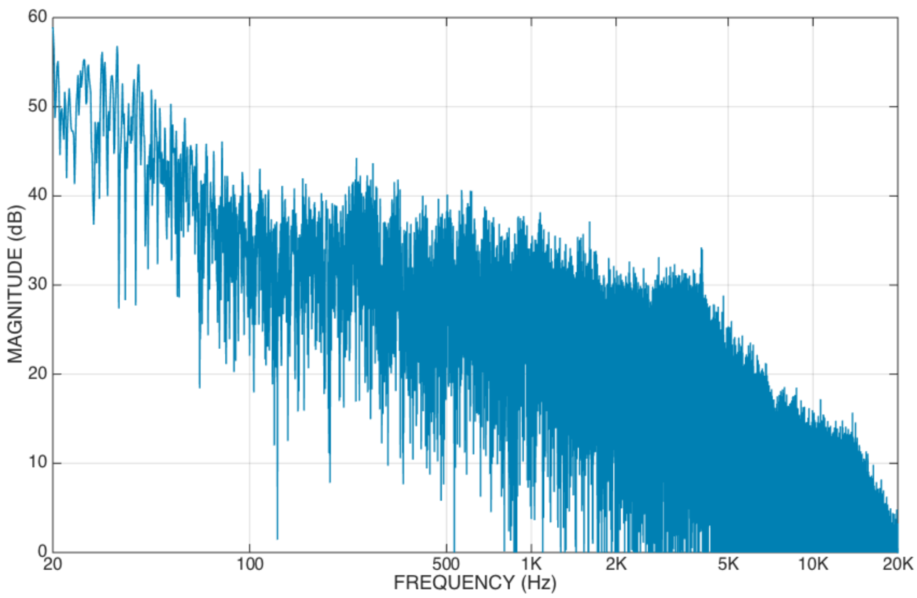

Figure C.2: Spectrum of the acoustic wind machine.

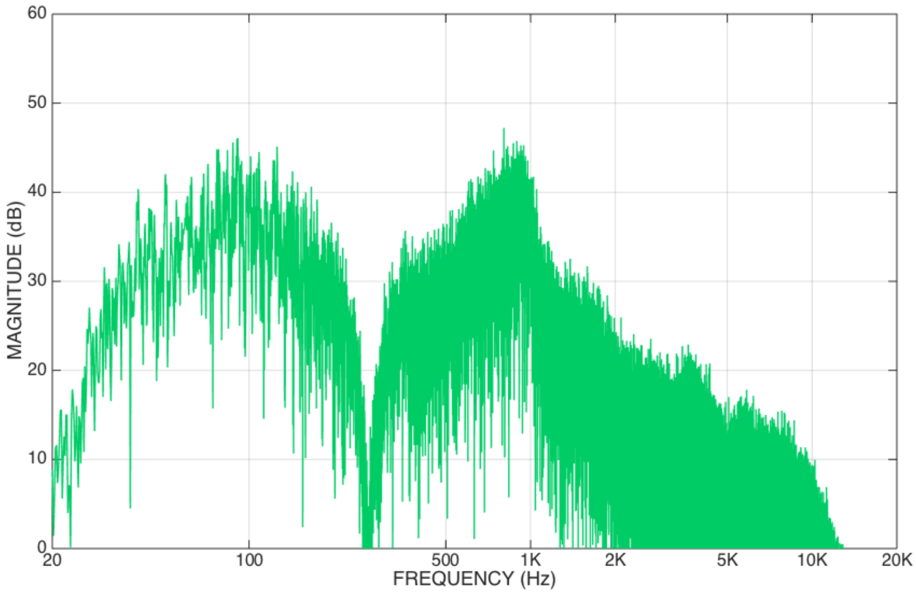

Figure C.3: Spectrum of the natural wind.

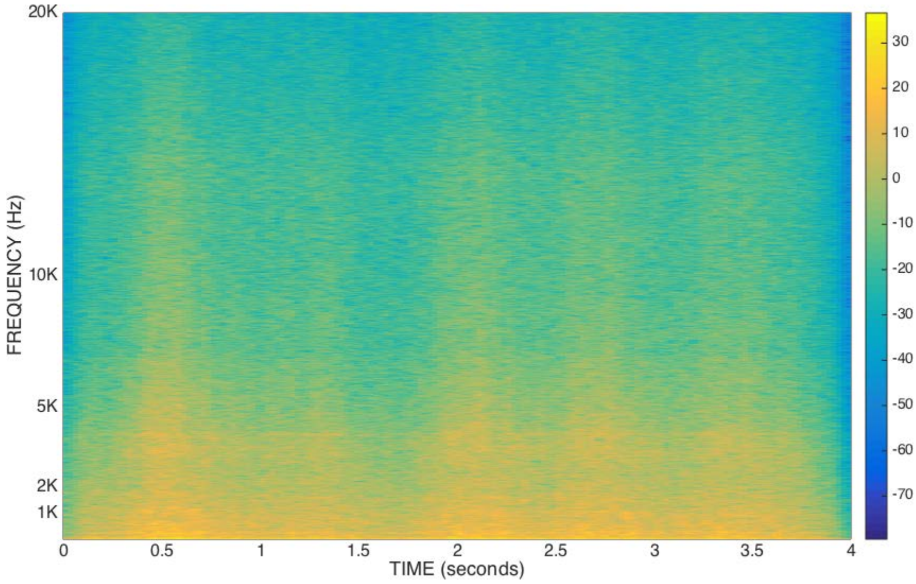

Figure C.4: Spectrogram of the acoustic wind machine.

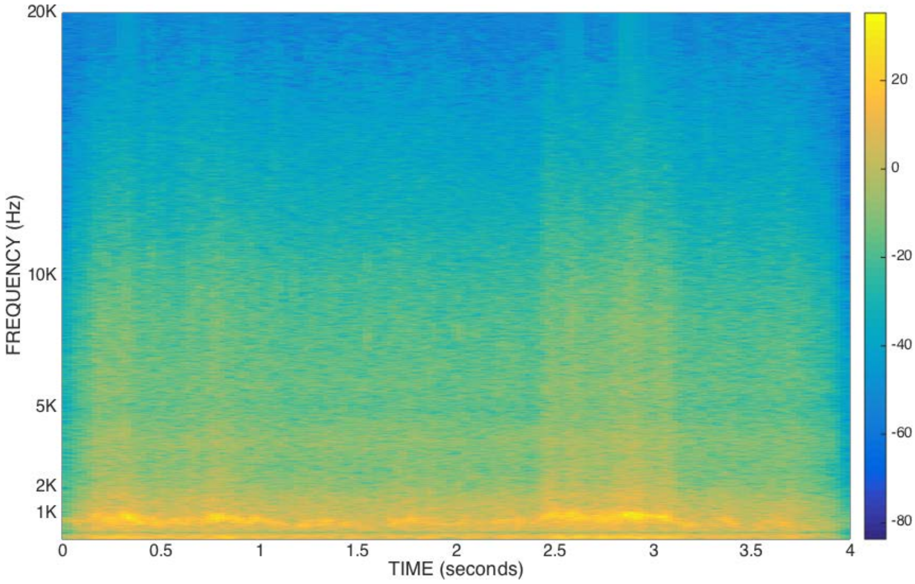

Figure C.5: Spectrogram of the natural wind.

C.2 Development of the Digital Wind Model

Chapter 6 describes the final iteration of the Max/MSP program created to model the sound of the acoustic wind machine in detail. This work was informed by previous acoustical analysis and comparison work published in Keenan & Pauletto (2016) and Keenan (2016). This section presents those acoustical analyses, as well as the parameters to the SDT model that produced each iteration of the digital wind-like sound.

C.2.1 Iteration 1

This version of the prototype digital wind machine was activated and modulated in performance with the IMU sensor coupled to the acoustic wind machine's cylinder. These results were previously published in Keenan & Pauletto (2016).

Parameters to [sdt.scraping∼]	
Surface profile (a signal)	noise∼
Probe width or grain (density of micro-impacts)	0.004
Velocity (m/s)	*real-time data: orientation angle*
Force	*real-time data: accelerometer to torque equation*
Parameters to [sdt.inertial∼]	
Mass of inertial object (kg)	0.01
Fragment size (to simulate crumpling)	1
Parameters to [sdt.friction∼]	
External rubbing force	*signal from sdt.scraping∼*
Bristle stiffness	1000
Bristle dissipation	1
Viscosity (speed of timbre evolution and pitch)	0.3
Amount of sliding noise (perceived surface roughness)	0
Dynamic friction coefficient (high values reduce sound bandwidth)	0.1
Static friction coefficient (smoothness of sound attacks)	0.5
Breakaway coefficient (transient of elasto-plastic state)	0.1
Stribeck velocity (smoothness of sound attacks)	0.1
Parameters to [sdt.modal∼]	
Active modes	3
Frequency of each mode (Hz)	400, 720, 1200
Decay of each mode (s)	0.04, 0.02, 0.03

Table C.1: SDT parameters for the first iteration of the digital slat model.

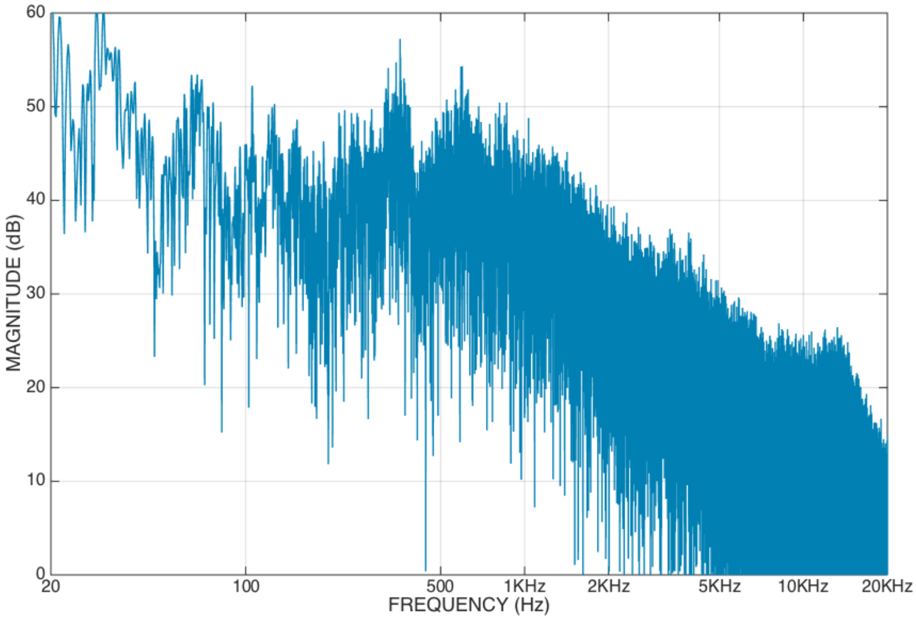

Figure C.6: Spectrogram of the acoustic wind machine.

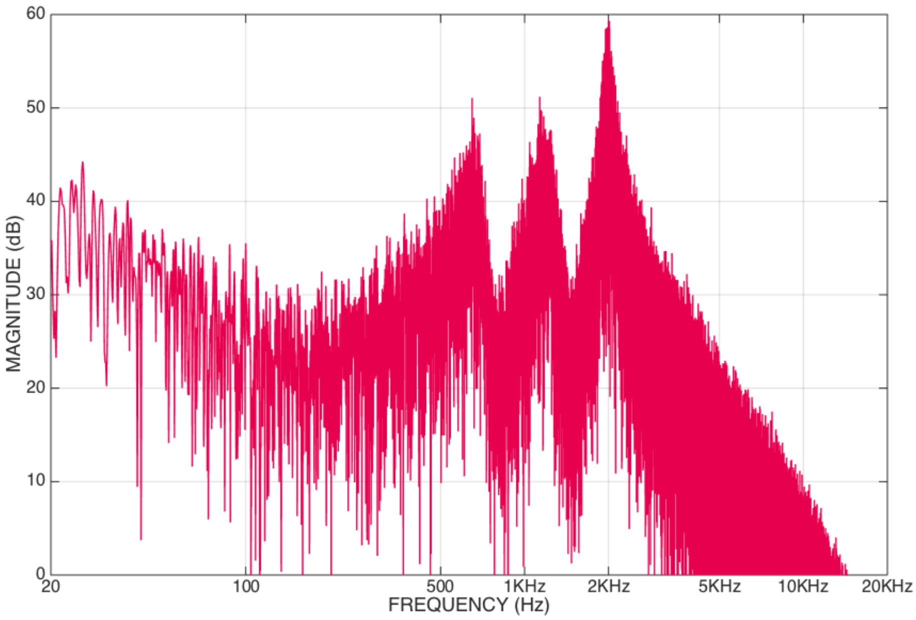

Figure C.7: Spectrogram of the first iteration of the prototype digital wind machine.

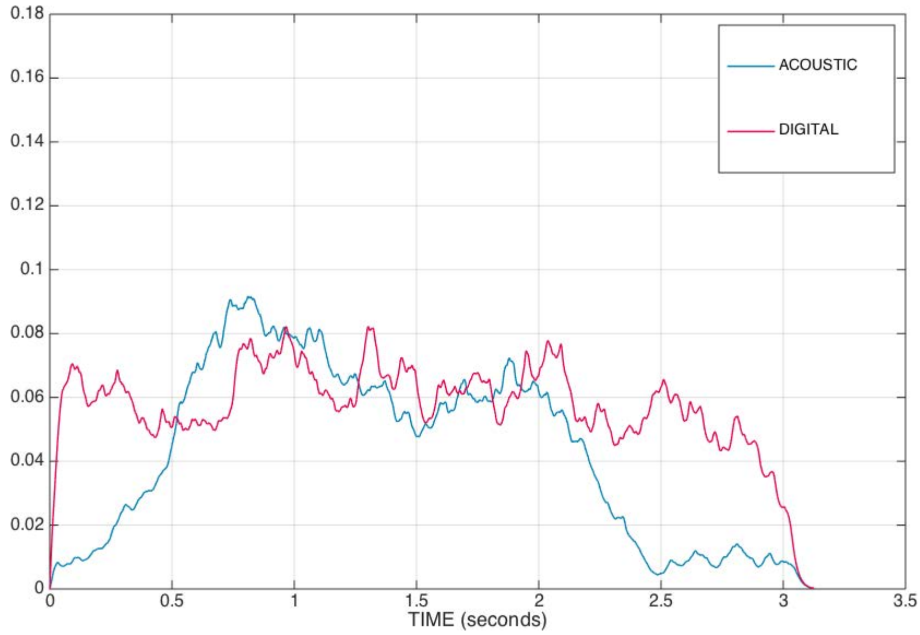

Figure C.8: Comparison of the amplitude envelopes for 1 rotation.

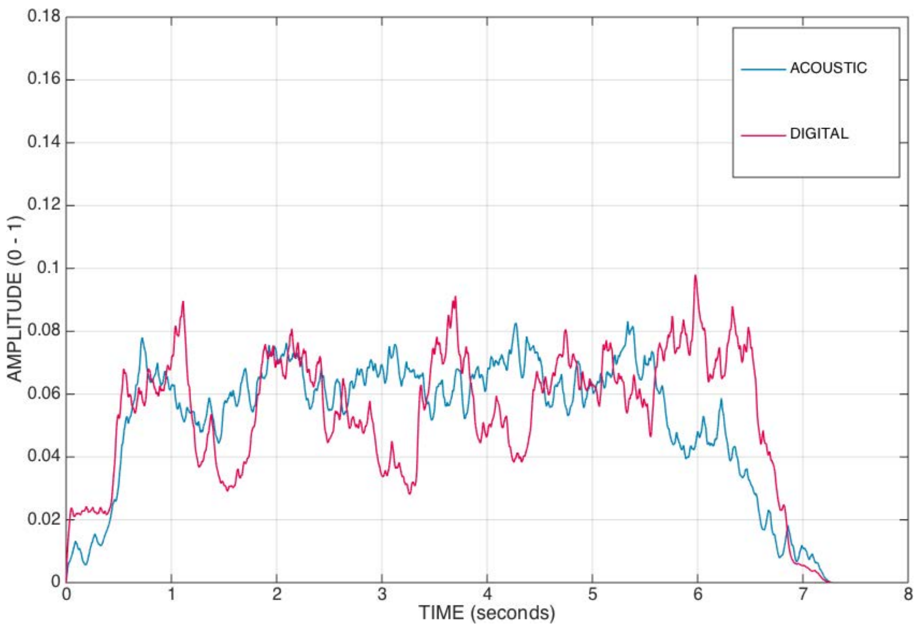

Figure C.9: Comparison of the amplitude envelopes for 5 rotations.

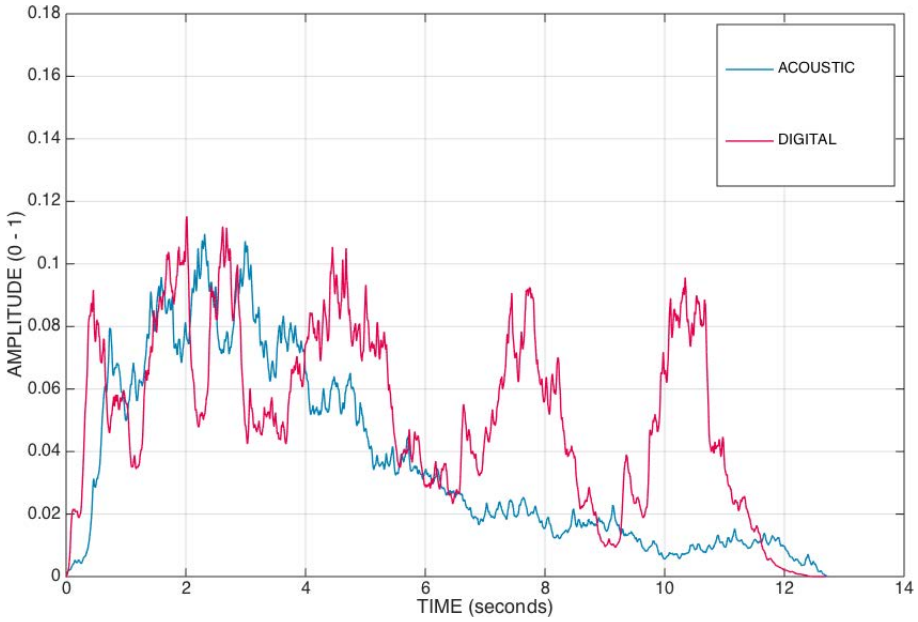

Figure C.10: Comparison of the amplitude envelopes for 10 rotations.

C.2.2 Iteration 2

This version of the prototype digital wind machine was activated and modulated in performance with the rotary encoder and 3D printed gearing coupled to the acoustic wind machine's cylinder and A-frame. These results were previously published in Keenan (2016)

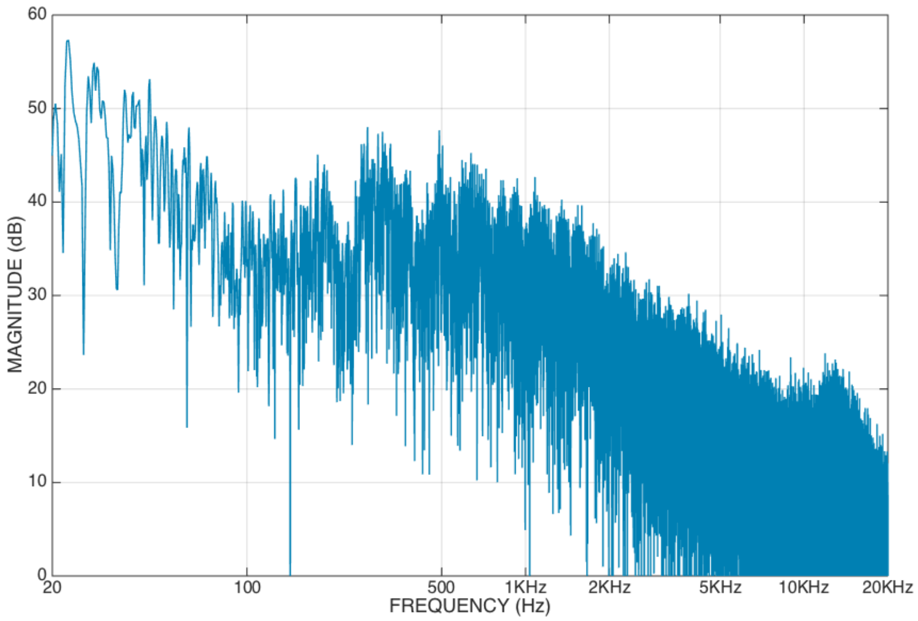

Figure C.11: Spectrogram of the acoustic wind machine.

Parameters to [sdt.scraping~]	
Surface profile (a signal)	noise~
Probe width or grain (density of micro-impacts)	*0.08, modulated by encoder data*
Velocity (m/s)	*real-time encoder data*
Force	*0.54, modulated by encoder data*
Parameters to [sdt.inertial~]	
Mass of inertial object (kg)	0.01
Fragment size (to simulate crumpling)	1
Parameters to [sdt.friction~]	
External rubbing force	*signal from sdt.scraping~*
Bristle stiffness	500.
Bristle dissipation	40.
Viscosity (speed of timbre evolution and pitch)	1.2
Amount of sliding noise (perceived surface roughness)	0.6
Dynamic friction coefficient (high values reduce sound bandwidth)	0.15
Static friction coefficient (smoothness of sound attacks)	0.5
Breakaway coefficient (transient of elasto-plastic state)	0.17
Stribeck velocity (smoothness of sound attacks)	0.103
Parameters to [sdt.modal~]	
Active modes	3
Frequency factor	1
Frequency of each mode (Hz)	380, 836, 1710
Decay factor	0.005
Decay of each mode (s)	0.8, 0.45, 0.09

Table C.2: SDT parameters for the second iteration of the digital slat model.

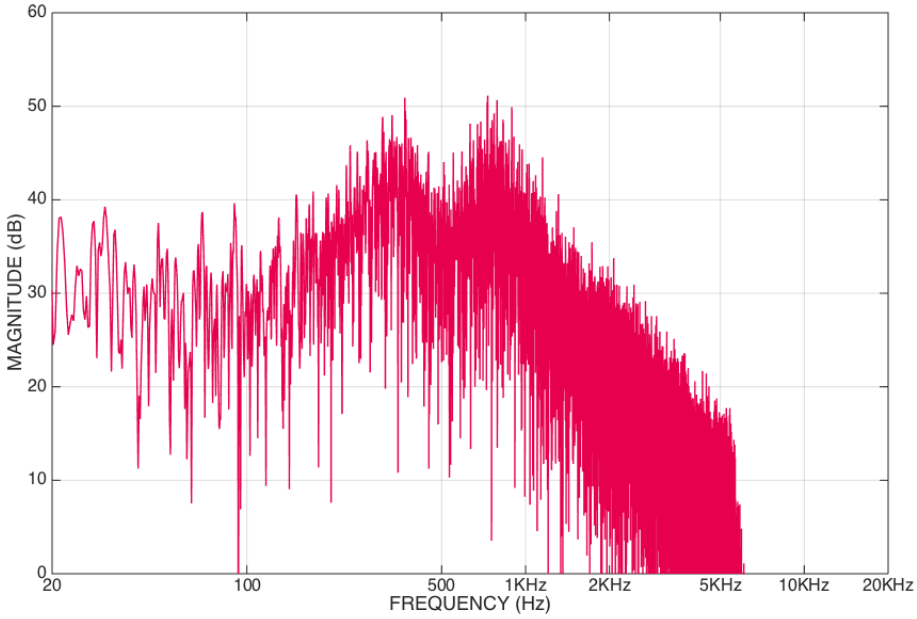

Figure C.12: Spectrogram of the second iteration of the prototype digital wind machine.

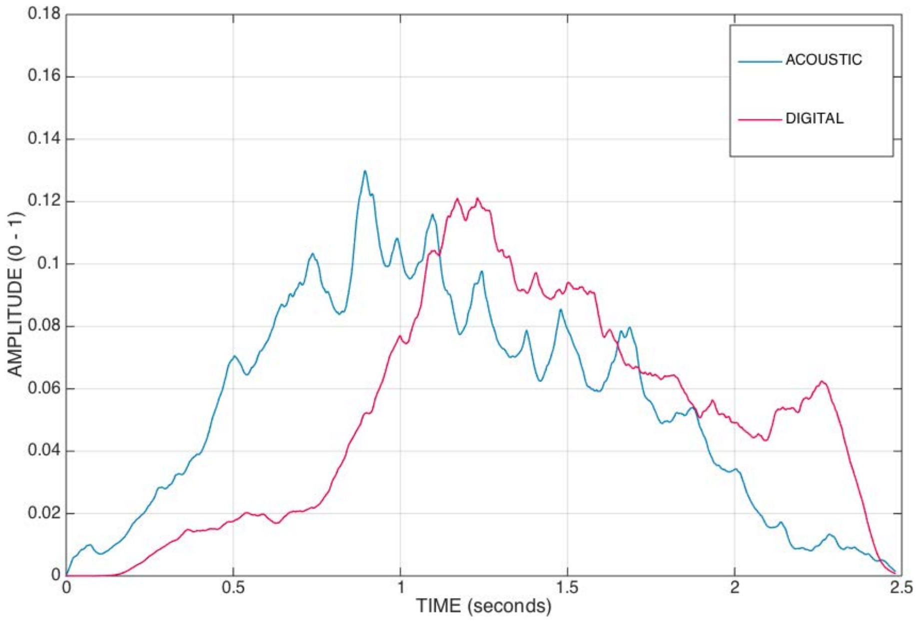

Figure C.13: Comparison of the amplitude envelopes for 1 rotation.

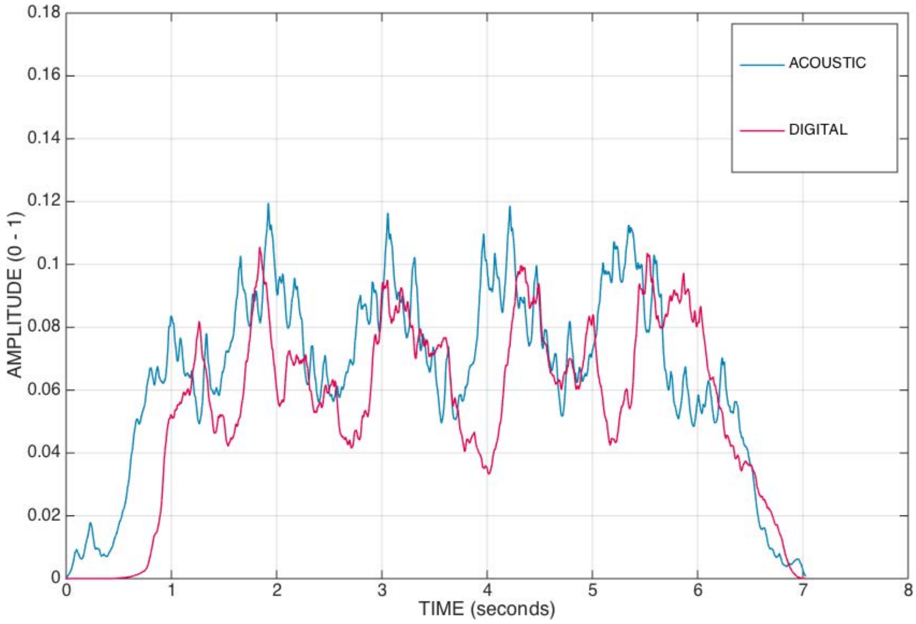

Figure C.14: Comparison of the amplitude envelopes for 5 rotations.

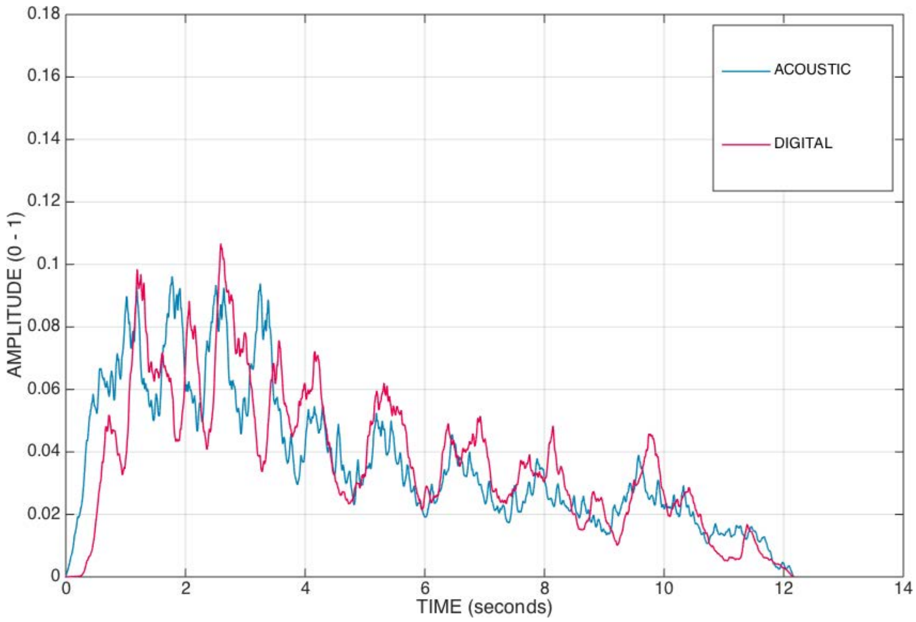

Figure C.15: Comparison of the amplitude envelopes for 10 rotations.

Appendix D

Experiment Documentation

This appendix gathers together the documentation from the experiment reported in Chapter 7 of this thesis. Section D.1 outlines the full protocol that was developed to guide the experimental evaluation. Section D.2 presents the text of the project information sheet and consent form that participants completed before taking part in the experiment. Section D.3 presents the questions that participants answered during the experiment.

D.1 Experiment Protocol

This protocol was developed to guide the experimental evaluation of the two wind machines.

D.1.1 Background

Goal: Confirm the assumptions that historical sound effects designers have made about the acoustic wind machine in operation (that it produces a varying sound with changes in rotational speed, etc.); confirm that the digital synthesis engine wind sound is *similar enough* to the acoustic wind machine sound for it to be used as a substitute for the acoustic version in a future experiment.

Main Hypotheses:

- Hypothesis 1 (H1): There is perceived similarity between the acoustic wind machine sound and the prototype digital wind machine sound.
- Hypothesis 2a (H2a): Participants can perceive different rotational speeds in a

continuous wind-like sound produced through acoustic or digital means.

- Hypothesis 2b (H2b): The perception of rotational speed is equally accurate when the continuous wind-like sound is produced through acoustic or digital means
- Hypothesis 3 (H3): There is perceived similarity between the experience of performing with the acoustic wind machine and that of performing with the prototype digital wind machine.

Null Hypotheses:

- Null Hypothesis 1 (H10): There is no perceived similarity between the acoustic wind machine sound and the prototype digital wind machine sound.
- Null Hypothesis 2a (H2a0): Participants cannot perceive different rotational speeds in a continuous wind-like sound produced through acoustic or digital means.
- Null Hypothesis 2b (H2b0): The perception of rotational speed is not equally accurate when the continuous wind-like sound is produced through acoustic or digital means
- Null Hypothesis 3 (H30): There is no perceived similarity between the experience of performing with the acoustic wind machine and that of performing with the prototype digital wind machine.

D.1.2 Design

Independent/Dependent Variables:

1. *Treatment 1:*
   (a) Independent variable is the pair of audio files being played; dependent variable is the perceived similarity between them.
   (b) Independent variable is the audio file being played; dependent variable is the perceived rotational speed of the stimulus.
2. *Treatment 2:* Independent variable is the audio stimulus used to elicit a performance; dependent variables are the perceived similarity of the wind machine being performed to the stimulus, and the perceived easiness of the wind machine to play.

Variables to Be Controlled:

- *All hypotheses:* The difference in the experience of listening to the acoustic wind-like sound and that of the digital wind-like sound heard through loudspeakers will be eliminated through the use of headphones to listen to both sounds during performance.
- *Hypothesis 3 (H3):* Visual feedback will be eliminated as a factor in the performance of the acoustic and digital wind sounds through the use of an enclosing screen.

Population to Be Studied: University of York students and staff.

Participant Selection: Volunteers will be recruited with an email asking them to sign up for a timeslot on a specific date via a Doodle Poll.

Objects of Study:

1. A corpus of audio files of recorded natural wind (from the BBC sound effects library), acoustic wind machine and digital wind machine performances, and an impact sound and a vocalisation/harmonic sound (from sound effects library).
2. A corpus of audio files of acoustic wind machine and digital wind machine performances at particular speeds (RPM values).
3. Acoustic wind machine and digital wind machine in operation, a corpus of audio files of natural, acoustic and digital wind to imitate in performance.

Treatment 1: The listening experiment will be run on the Open Sesame platform. Participants will listen to all sounds through headphones.

1. Selected pairs of audio files of natural, acoustic and digital wind, along with a very different pair of sounds to contextualise these (an impact, and a vocalisation or harmonic sound (Gygi et al. 2007)), will be played to the participants via a computer. The order of the audio file pairs will be randomised for each participant to minimise order effects (Bonebright et al. 2005). Participants will be asked to rate the pairs in terms of their similarity on a 7-point scale (Gygi et al. 2007). Each audio file in the corpus will be paired with each other audio file to facilitate analysis.
2. Participants will listen to audio files of acoustic wind machine and digital wind machine performances at different speeds. They will be told that the sounds have

been created by turning a handle, and will be asked to rank the audio files by the speed of each sound on a 1 - 4 scale.

Treatment 2: The acoustic wind machine will be positioned behind an enclosing screen, with only its handle protruding through so participants cannot see the cylinder and cloth in motion during operation. It will be mic-ed with an AKG microphone fed through to a pair of headphones. The acoustic wind machine's movement will be coupled to that of a rotary encoder connected to an Arduino, which will drive a digital wind synthesis engine in Max/MSP. The digital wind audio will also be fed through to the same pair of headphones. Participants will perform the acoustic and digital wind sounds, which will be recorded into Pro Tools, along with the data from the rotary encoder movement, which will be recorded by Max/MSP.

1. Participants will be invited to turn the wind machine's crank handle and hear the resulting sound (acoustic or digital).

2. Participants will be played a short clip of natural (for the practice step), acoustic or digital wind (for the main step).

3. Participants will be asked to play the wind machine handle to imitate the clip they have just heard. The order in which they hear the sounds, and the order in which they play the acoustic or digital sound, will be randomised. There will be a training step, and then a test step, for each performed sound. Audio and data from all performances will be recorded.

4. Participants will be asked to describe their experience of performance of each wind sound by completing a questionnaire.

   (a) Do they think they managed to successfully imitate the wind sounds that they heard?

   (b) Did they find the wind sound easy to play?

   (c) Participants will be asked to describe the kind of wind they have performed by selecting descriptors from a list.

   (d) Participants will be asked to freely describe their experience of the performance.

D.1.3 Data

Preparation of Audio Files:

1. Some recordings of natural wind sounds will be selected. The acoustic wind machine and prototype digital wind machine will be performed to imitate the natural wind sound; these will be recorded. An impact sound and a vocalisation/harmonic sound will be selected from a sound library to contrast with the wind sounds. All sound files will be edited and catalogued appropriately to serve as stimuli. Their duration, acoustic qualities (spectral centroid, temporal properties) and classification for this experiment will be established.

2. Some recordings of performances with the acoustic wind machine and digital wind machine will be created, and the RPM of each performance will be established using Max/MSP. Four clips of each of the sounds at varying RPMs will be selected for the experiment.

Preparation of Acoustic and Prototype Digital Wind Machines: A screen will be constructed to eliminate any visual feedback in the performance of the acoustic and digital wind sounds. An audio interface will be configured to accept signal from an AKG microphone for the acoustic wind and digital audio from a MacBook running the digital synthesis engine in Max/MSP.

Collection: Participants will be asked to complete a consent form before taking part in the experiment.

A questionnaire will be written for participants, and provided to them as an Open Sesame program. Questionnaires will be anonymised, but age range and gender data will be gathered from participants. Participants will be asked to report if they consider themselves to have normal hearing.

Audio recordings will be made of the acoustic and digital sounds in performance, along with data recorded from the rotary encoder.

Storage: Following this experiment, the questionnaire data from Open Sesame will be exported as a .csv spreadsheet and stored in Google Drive. This will be imported into R for statistical analysis. The recordings of participants' performances will be trimmed and classified by gesture to faciliate batch analysis in Matlab with the MIR Toolbox.

They will also be stored in Google Drive.

Analysis:

1. Standard descriptive statistics.

2. Multidimensional scaling (MDS) to examine the perceptual space between the natural, acoustic and digital wind-like sounds, in addition to the impact and vocalisation or harmonic sound. The acoustic analysis of the sounds will also be examined to see if it correlates with the ordering of the sounds along the dimensions of the MDS analysis.

3. Friedman rank-sum or ANOVA test to discover if there are significant differences between how the wind-like sounds at the various RPMs were ranked for speed. Wilcoxon signed-rank test for post-hoc testing of these results.

4. Kruskal-Wallis test to confirm that the order of presentation of stimuli in the performance test did not influence how the wind machines were ranked in terms of similarity or ease of play. Friedman rank-sum or ANOVA test to discover if there are significant differences between how the wind machines were ranked for similarity to the stimuli they imitate, and for ease of play. Wilcoxon signed-rank test for post-hoc testing of these results.

5. Wordclouds to summarise participants' free descriptions.

6. Wilcoxon signed-rank test to discover any significant differences between the acoustical descriptors of participants' performances depending on the stimulus that elicited them.

Statistical Power and Participants: As this is the first perceptual study of a theatre wind machine design, a minimum of 50 participants will be sought to ensure a statistical power of 0.8 when detecting a large effect size ($r = 0.5$) at $\alpha \leq 0.05$.

Study Limitations: The performance task with the acoustic and digital wind sounds will not strictly measure accuracy over time (e.g. using timecode) at this stage, but the audio recordings will be compared to establish whether performance with one sound was very different to the other.

Reporting: This experiment will be fully reported in my PhD thesis, and the results will also inform a published journal article.

D.2 Project Information Sheet and Consent Form

This is the text of the information sheet and consent form that participants completed before taking part in the experiment.

PhD Working Title: Enactive Sound Machines: Theatrical Strategies for Sonic Interaction Design

Project Description

This research project examines the design and performance of late nineteenth and early twentieth century theatre sound effects to reveal new strategies for linking sound and action together in performance in an intuitive way.

The aim of the listening test is to examine the perception of sounds, and also the perceived similarity between specific pairs of sounds.

The aim of the performance test is to examine how performance feels with different sounds. Audio recordings will be made of the performances.

The data collected throughout these tests will be analysed and the results reported as part of the PhD project thesis in question. The data might also be used for conference presentations, publications and for the general dissemination of the research project.

Duration: The tests will take approximately 45 minutes to complete. There will be two sections, one for the listening test and one for the performance test, with a short break between them.

If you have any questions or need more information on the project, you can contact the researcher Fiona Keenan at fiona.keenan@york.ac.uk.

Consent Form

1	2	3	4	5	6	7
Not similar at all						As similar as they can possibly be

Table D.1: The similarity scale for the first listening step.

I confirm that I have read this information and I am aware of the description of the research project and what the data will be used for. (**YES/NO**)

I confirm that I have been informed that by completing the listening test and performance test I am providing my consent to participate in the study (**YES/NO**)

I confirm that I give my consent to be recorded anonymously. (**YES/NO**)

Signed by:

D.3 Experiment Questions

The questions that participants answered during the experiment for data collection are detailed here.

1. **Preliminary Questions**

   (a) What is your gender? (M/F/Other/Prefer not to say)

   (b) What is your age range? 18-24/25-34/35-44/45-54/55-64/65+/Prefer not to say

   (c) Do you consider yourself to have normal hearing? Y/N

   (d) Do you play a musical instrument? Y/N If so, how would you rate the level you play at? Beginner/Intermediate/Expert

2. **Perceived Similarity:** Listen to this pair of sounds. How similar are they to each other? Rate them on the scale below (Table D.1).

1	2	3	4
Slowest			Fastest

Table D.2: The speed ranking scale for the second listening step.

1	2	3	4	5	6	7
Strongly disagree						Strongly agree

Table D.3: The ease of play ratings scale.

3. **Perceived Speed:** These sounds are all made by turning a handle. Listen to them all as much as you like, and then rank them in order of their speed on the scales below (Table D.2).

4. **Performances:** Turning this handle will create a sound. You will be played a sound through the headphones. When it finishes, try to imitate the sound you have just heard by turning the handle. There will be a chance for you to practice first.

   Now that you have finished playing, please answer these questions:

   (a) This wind sound is easy to play. Rate how far you agree with this statement on the scale below (Table D.3).

   (b) Did you manage to play a wind sound similar to the wind you listened to first? Rate the similarity on the scale below (Table D.4).

   (c) Describe the wind sound you played (tick all that apply):

   - Gentle
   - Gusty
   - Breeze
   - Gale
   - Swishing
   - Shrieking
   - Howling

1	2	3	4	5	6	7
Not similar at all						As similar as they can possibly be

Table D.4: The similarity ratings scale for the performance step.

- Fast
- Slow
- Continuous
- Other (please specify)

(d) Please describe your experience of performing this wind sound.

Thank you for your participation!

Appendix E

Attributions for Vector Graphics

This appendix credits the designers of the free vector graphics used within this thesis. All images were used under a Creative Commons *free with attribution required* licence. The other components of the images produced for this thesis not listed here were the sole work of the author herself.

- Xbox controller by Open Clip Art.org
- Hands by Vecteezy.com
- Crumpled paper sourced from Freepik
- Audio production devices designed by Zarubin-Leonid / Freepik
- Laptop by Open Clip Art.org
- Hand drill by Gustavo Rezende at Open Clip Art.org
- Gearing sourced from Freepik.
- Arduino from Ui-Ex.com

Appendix F

Guide to USB Media

This appendix lists the folders included on the accompanying USB drive.

1. A copy of this thesis as a .pdf file.
2. Audio files of the acoustic wind machine and a recording of natural wind used for the acoustic analysis in Chapter 5 and Appendix C.1
3. Audio files of the same performance gestures with the acoustic wind machine and prototype digital wind machine, used for the acoustic analyses in Chapter 6 and Appendix appendix:third:section2:
   (a) 1 rotation, 5 rotations and 10 rotations performed with the acoustic wind machine and the first iteration of its digital counterpart. (Appendix C.2.1)
   (b) 1 rotation, 5 rotations and 10 rotations performed with the acoustic wind machine and the second iteration of its digital counterpart. (Appendix C.2.2)
   (c) 1 rotation, 5 rotations and 10 rotations performed with the acoustic wind machine and the third and final iteration of its digital counterpart (Chapter 6).
4. Stimuli used for the experiment reported in Chapter 7:
   (a) Audio files for the first listening step, and a guide to their pairs.
   (b) Audio files for the second listening step, and a guide to their RPMs
   (c) Audio files for the performance step, and a guide to their gestures
5. Data gathered from participants during the experiment reported in Chapter 7:

(a) Participants' similarity ratings.

(b) Participants' speed rankings.

(c) Results of the performance step.

(d) Folder of recordings of participants' performances.

(e) Results of the acoustic analysis of participants' performances.

References

Adams, S. & Bigand, E. (1993), *Thinking in sound: The cognitive psychology of human audition*, Oxford: Clarendon Press.

Altavilla, A., Caramiaux, B. & Tanaka, A. (2013), Towards gestural sonic affordances, *in* 'New Interfaces for Musical Expression (NIME)', Daejeon, Korea., pp. 61 – 64.

Altman, R. (2007), *Silent film sound*, Columbia University Press.

Ament, V. T. (2009), *The foley grail: The art of performing sound for film, games, and animation*, CRC Press.

Armitage, J. & McPherson, A. (2018), Crafting digital musical instruments: An exploratory workshop study, *in* 'New Interfaces for Musical Expression (NIME)'.

Avanzini, F. & Crosato, P. (2006*a*), Haptic-auditory rendering and perception of contact stiffness, *in* 'International Workshop on Haptic and Audio Interaction Design', Springer, pp. 24–35.

Avanzini, F. & Crosato, P. (2006*b*), 'Integrating physically based sound models in a multimodal rendering architecture', *Computer Animation and Virtual Worlds* **17**(34), 411–419.

Avanzini, F., Serafin, S. & Rocchesso, D. (2005), 'Interactive simulation of rigid body interaction with friction-induced sound generation', *Speech and Audio Processing, IEEE Transactions on* **13**(5), 1073–1081.

Baldan, S., Delle Monache, S. & Rocchesso, D. (2017), 'The sound design toolkit', *SoftwareX* **6**, 255–260.

Barbosa, J., Malloch, J., Wanderley, M. M. & Huot, S. (2015), What does "evaluation" mean for the nime community?, *in* 'NIME 2015-15th International Conference on New Interfaces for Musical Expression', Louisiana State University, pp. 156–161.

Bartholomew, D. J., Steele, F., Galbraith, J. & Moustaki, I. (2008), *Analysis of multivariate social science data*, Chapman and Hall/CRC.

Baugh, C. & Wilmore, D. (2015), *Backstage in the theatre: scenes and machines by Jean-Pierre Moynet*, Theatreshire Books, Yeadon, West Yorkshire, UK.

Bax, P. (1936), *Stage management*, London : Lovat Dickson, London.

BBC (1986), 'Comedy, fantasy and humour', CD.

BBC (1988), 'Weather 1', CD.

Bennett, P., Ward, N., O'Modhrain, S. & Rebelo, P. (2007), Damper: a platform for effortful interface development, *in* 'Proceedings of the 7th international conference on New interfaces for musical expression', pp. 273–276.

Berdahl, E. (2014), How to make embedded acoustic instruments, *in* 'New Interfaces for Musical Expression (NIME)', pp. 140–143.

Beyer, R. T. & Raichel, D. R. (1999), *Sounds of our times, two hundred years of acoustics*, ASA.

Bonebright, T. L., Miner, N. E., Goldsmith, T. E. & Caudell, T. P. (2005), 'Data collection and analysis techniques for evaluating the perceptual qualities of auditory stimuli', *ACM Transactions on Applied Perception (TAP)* **2**(4), 505–516.

Born, G. (1995), *Rationalizing culture: IRCAM, Boulez, and the institutionalization of the musical avant-garde*, University of California Press.

Bottomore, S. (1999), 'An international survey of sound effects in early cinema', *Film History* pp. 485–498.

Bottomore, S. (2001), The story of Percy Peashaker: Debates about sound effects in the early cinema, *in* R. Abel & R. Altman, eds, 'The Sounds of Early Cinema', Indiana University Press, pp. 129 – 142.

Brown, R. (2010), *Sound: A reader in theatre practice*, Palgrave Macmillan.

Browne, V. D. (1913), *Secrets of scene painting and stage effects*, Routledge.

Cadoz, C. (1988), Instrumental gesture and musical composition, *in* 'ICMC 1988-International Computer Music Conference', pp. 1–12.

Cadoz, C. (2009), 'Supra-instrumental interactions and gestures', *Journal of New Music Research* **38**(3), 215–230 0929–8215.

Cadoz, C. & Wanderley, M. M. (2000), 'Gesture-music', *Trends in gestural control of music* **12**, 101.

Cage, J. (1937), The future of music: Credo, *in* 'Audio Culture: Readings in Modern Music', Continuum.

Calì, J., Calian, D. A., Amati, C., Kleinberger, R., Steed, A., Kautz, J. & Weyrich, T. (2012), '3d-printing of non-assembly, articulated models', *ACM Transactions on Graphics (TOG)* **31**(6), 130.

Caramiaux, B., Bevilacqua, F., Bianco, T., Schnell, N., Houix, O. & Susini, P. (2014), 'The role of sound source perception in gestural sound description', *ACM Transactions on Applied Perception (TAP)* **11**(1), 1544–3558.

Caramiaux, B., Susini, P., Bianco, T., Bevilacqua, F., Houix, O., Schnell, N. & Misdariis, N. (2011), Gestural embodiment of environmental sounds: an experimental study, *in* 'New Interfaces for Musical Expression (NIME)', Citeseer.

Chadabe, J. (1997), *Electric sound: The past and promise of electronic music*, Pearson.

Chion, M. (2015), Epilogue. audition and ergo-audition: Then and now, *in* D. Daniels & S. Naumann, eds, 'See this Sound: Audiovisuology: a Reader', Buchhandlung Walther König, pp. 670–684.

Choi, I. (2000), 'Gestural primitives and the context for computational processing in an interactive performance system', *Trends in Gestural Control of Music* pp. 139–172.

Collins, N. (2015), Semiconducting: Making music after the transistor, *in* G. Borio, ed., 'Musical Listening in the Age of Technological Reproduction', Ashgate Publishing, Ltd.

Collison, D. (2008), *The sound of theatre*, London: Professional Lighting and Sound Association.

Collison, D. (2010), 'What were we doing in 1960?', *Theatre Design and Technology* **46**(3), 12 – 17.

Collison, D. & Hall, P. (1976), *Stage sound*, Studio Vista.

Cook, D. (1876), *A book of the play: Studies and illustrations of histrionic story, life, and character*, Tredition.

Cook, P. R. (2002), *Real sound synthesis for interactive applications*, AK Peters Wellesley.

Counsell, C. (2013), *Signs of performance: an introduction to twentieth-century theatre*, Routledge.

Crevoisier, A. & Polotti, P. (2005), Tangible acoustic interfaces and their applications for the design of new musical instruments, *in* 'New Interfaces for Musical Expression (NIME)', National University of Singapore, pp. 97–100.

Cross, N. (2006), *Designerly ways of knowing*, Springer.

Culver, M. K. (1981), 'A history of theatre sound effects devices to 1927', PhD Thesis.

Curtin, A. (2011), Noises on: Sights and sites of sound in Apollinaire's The Breasts of Tiresias, *in* D. P. Roesner & L. Kendrick, eds, 'Theatre Noise: The Sound of Performance', Cambridge Scholars Publishing, pp. 125 – 146.

Curtin, A. (2014), *Avant-garde theatre sound: Staging sonic modernity*, Palgrave Macmillan.

Davies, H. (1994), 'The sound world, instruments and music of luigi russolo', Online.
URL: *creativegames.org.uk/modules/Art$_T$echnology/theory/authors/hugh$_d$avies.htm*

DiFranco, D. E., Beauregard, G. L. & Srinivasan, M. A. (1997), Effect of auditory cues on the haptic perception of stiffness in virtual environments, *in* 'American Society of Mechanical Engineers, Dynamic Systems and Control Division (Publication) DSC', Vol. 61, pp. 17–22.

Dobrian, C. & Koppelman, D. (2006), *The'E'in NIME: musical expression with new computer interfaces*, IRCAM—Centre Pompidou.

Dodge, W. P. (1912), Staging a sandstorm, *in* 'Theatre Magazine', Vol. 15.

Dourish, P. (2001), *Where the action is: The foundations of embodied interaction*, MIT Press.

Dupont, P., Hayward, V., Armstrong, B. & Altpeter, F. (2002), 'Single state elastoplastic friction models', *Automatic Control, IEEE Transactions on* **47**(5), 787–792.

Durrant, A. C., Vines, J., Wallace, J. & Yee, J. S. (2017), Research through design: Twenty-first century makers and materialities, *in* 'Design Issues', MIT Press.

Eastaugh, N., Walsh, V., Chaplin, T. & Siddall, R. (2008), *Pigment compendium*, Routledge.

Edge, H. (1991), 'Vol. 2 sound effects', CD.

Electric, S. (1947), 'A completely new glossary of technical theatrical terms'.

Elliott, D., MacDougall, R. & Turkel, W. J. (2012), 'New old things: Fabrication, physical computing, and experiment in historical practice', *Canadian Journal of Communication* **37**(1).

Essl, G. & O'Modhrain, S. (2006), 'An enactive approach to the design of new tangible musical instruments', *Organised sound* **11**(03), 285–296.

Farnell, A. (2010), *Designing sound*, MIT Press Cambridge.

Feld, S. (2015), 'Acoustemology', *Keywords in Sound* pp. 12–21.

Field, A. & Hole, G. (2002), *How to design and report experiments*, Sage.

Field, A., Miles, J. & Field, Z. (2012), *Discovering statistics using R*, Sage publications.

Fitzgerald, P. H. (1881), *The World Behind the Scenes*, Chatto and Windus.

Fletcher, N. H., Tarnopolsky, A. Z., Lai, J. et al. (2002), Australian aboriginal musical instruments–the bullroarer, *in* 'Acoustic Innovation in Acoustics and Vibration Annual Conference of the Australian Acoustical Society', pp. 186–189.

Fontana, F. & Rocchesso, D. (2003), *The sounding object*, Mondo Estremo, Italy.

Franinović, K. (2009), 'Amplified movements: An enactive approach to sound in interaction design', *New Realities: Being Syncretic* pp. 114–117.

Franinović, K. (2013), 'Amplifying actions - towards enactive sound design', PhD Thesis.

Franinović, K. & Serafin, S. (2013), *Sonic interaction design*, MIT Press.

Frayling, C. (1993), 'Research in art and design', *Royal College of Art London* .

Gassner, J. & Barber, P. W. (1953), *Producing the play*, Dryden Press.

Gaver, W. W. (1993*a*), 'How do we hear in the world? explorations in ecological acoustics', *Ecological psychology* **5**(4), 285–313.

Gaver, W. W. (1993*b*), 'What in the world do we hear?: An ecological approach to auditory event perception', *Ecological psychology* **5**(1), 1–29.

Gelineck, S. & Serafin, S. (2009), A quantitative evaluation of the differences between knobs and sliders, *in* 'New Interfaces for Musical Expression (NIME)', pp. 13–18.

Gibson, J. J. (1977), The theory of affordances, *in* R. E. Shaw & J. Bransford, eds, 'Perceiving, Acting and Knowing: Towards an Ecological Psychology', Lawrence Erlbaum Associates.

Gibson, J. J. (1979), *The ecological approach to visual perception*, Psychology Press.

Giordano, B. L., Susini, P. & Bresin, R. (2013), Perceptual evaluation of sound-producing objects, *in* K. Franinović & S. Serafin, eds, 'Sonic interaction design', MIT Press, pp. 151–197.

Giordano, B. L., Visell, Y., Yao, H.-Y., Hayward, V., Cooperstock, J. R. & McAdams, S. (2012), 'Identification of walked-upon materials in auditory, kinesthetic, haptic, and audio-haptic conditions', *The Journal of the Acoustical Society of America* **131**(5), 4002–4012.

Godøy, R. I., Haga, E. & Jensenius, A. R. (2006*a*), 'Exploring music-related gestures by sound-tracing: A preliminary study sound-tracing: A preliminary study'.

Godøy, R. I., Haga, E. & Jensenius, A. R. (2006*b*), Playing "air instruments": Mimicry of sound-producing gestures by novices and experts, *in* 'Gesture in Human-Computer Interaction and Simulation', Springer, pp. 256–267.

Green, M. (1958), *Stage noises and effects*, Herbert Jenkins: London, London.

Grein, J. T. (1928), 'Realism in sound behind the scenes: Railway noises "off"', *Illustrated London News* pp. 516–517.

Gurevich, M. & Treviño, J. (2007), Expression and its discontents: toward an ecology of musical creation, *in* 'Proceedings of the 7th international conference on New interfaces for musical expression', ACM, pp. 106–111.

Gygi, B., Kidd, G. R. & Watson, C. S. (2007), 'Similarity and categorization of environmental sounds', *Perception and Psychophysics* **69**(6), 839–855.

Hayes, L. (2011), *Vibrotactile feedback-assisted performance*, Citeseer.

Heidegger, M. (1962), *Being and time*, New York: Harper Row.

Heinrichs, C. (2018), 'Human expressivity in the control and integration of computationally generated audio', PhD Thesis.

Heinrichs, C., McPherson, A. & Farnell, A. (2014), 'Human performance of computational sound models for immersive environments', *The New Soundtrack* **4**(2), 139–155.

Hopkins, A. A. (1920), 'Noise makers of the stage', *Scientific American* **122**.

Houix, O., Lemaitre, G., Misdariis, N., Susini, P. & Urdapilleta, I. (2012), 'A lexical analysis of environmental sound categories', *Journal of Experimental Psychology: Applied* **18**(1), 52 categories

Hug, D. (2008), Towards a hermeneutics and typology of sound for interactive commodities, *in* 'Proc. CHI Workshop on Sonic Interaction Design', pp. 11–16.

Hug, D. & Kemper, M. (2014), 'From Foley to function: A pedagogical approach to sound design for novel interactions', *Journal of Sonic Studies* **6**(1).

Hunt, A., Wanderley, M. M. & Paradis, M. (2003), 'The importance of parameter mapping in electronic instrument design', *Journal of New Music Research* **32**(4), 429–440.

Ihde, D. (2016), First encounters with Husserl's phenomenology, *in* D. Ihde, ed., 'Husserl's Missing Technologies', Oxford University Press.

Ingold, T. (2006), 'Walking the plank: meditations on a process of skill', *Defining technological literacy: Towards an epistemological framework* pp. 65–80.

Ingold, T. (2013), *Making: Anthropology, archaeology, art and architecture*, Routledge.

Inventor Tells the Secret of Amazing Stage Effects (1913), *New York Times* **November**(9th).

Jackson, A. S. & Wilson, M. G. (1976), *French theatrical production in the nineteenth century: "L'Envers du Théatre" by M.J. Moynet, 1873*, Vol. 10, Max Reinhardt Foundation with the Center for Modern Theatre Research, State University of New York, USA.

Jay, M. (2011), 'In the realm of the senses: An introduction', *The American historical review* **116**(2), 307–315.

Jensenius, A. R. (2007), 'Action-sound: Developing methods and tools to study music-related body movement', PhD Thesis.

Jensenius, A. R. (2012), 'Disciplinarities: intra, cross, multi, inter, trans'.

Jensenius, A. R. & Lyons, M. J. (2017), *A NIME Reader: Fifteen years of new interfaces for musical expression*, Springer.

Jensenius, A. R. & Wanderley, M. M. (2010), Musical gestures: Concepts and methods in research, *in* 'Musical Gestures', Routledge, pp. 24–47.

Juslin, P. N. (2003), 'Five facets of musical expression: A psychologist's perspective on music performance', *Psychology of Music* **31**(3), 273–302.

Kahn, D. (1999), 'Noise, water, meat: A history of sound in the arts', *Massachusetts Institute of Technology* .

Karjalainen, M., Välimäki, V. & Tolonen, T. (1998), 'Plucked-string models: From the Karplus-Strong algorithm to digital waveguides and beyond', *Computer Music Journal* **22**(3), 17–32.

Kassambara, A. & Mundt, F. (2016), 'Factoextra: extract and visualize the results of multivariate data analyses', *R package version* **1**(3).

Keenan, F. (2016), 'A theatre wind machine as interactive sounding object'.

Keenan, F. & Pauletto, S. (2016), An acoustic wind machine and its digital counterpart: Initial audio analysis and comparison, *in* 'Interactive Audio Systems Symposium (IASS)', University of York, York, UK.
URL: *http://www.york.ac.uk/sadie-project/IASS2016.html*

Keenan, F. & Pauletto, S. (2017*a*), Design and evaluation of a digital theatre wind machine, *in* 'Proceedings of The 17th International Conference on New Interfaces for Musical Expression (NIME 17)', Copenhagen, Denmark.

Keenan, F. & Pauletto, S. (2017*b*), A mechanical mapping model for real-time control of a complex physical modelling synthesis engine with a simple gesture, *in* 'Digital Audio Effects (DAFx)'.

Kendrick, L. (2017), Theatre aurality: Beginnings, *in* 'Theatre Aurality', Springer, pp. 27–50.

Keysers, C., Kohler, E., Umiltà, M. A., Nanetti, L., Fogassi, L. & Gallese, V. (2003), 'Audiovisual mirror neurons and action recognition', *Experimental brain research* **153**(4), 628–636 recognition

Kiefer, C., Collins, N. & Fitzpatrick, G. (2008), HCI methodology for evaluating musical controllers: A case study, *in* 'New Interfaces for Musical Expression (NIME)', pp. 87–90.

Kim, L., Han, M., Shin, S. K. & Park, S. H. (2008), A haptic dial system for multimodal prototyping, *in* '18th international conference on artificial reality and telexistence (ICAT 2008)

Kranich, F. (1929), *Bühnentechnik der gegenwart*, München Berlin, München Berlin.

Krassenstein, E. (2015), 'Fully functional platform jack is amazingly 3d printable as one piece - no supports required'.
URL: *https://3dprint.com/84958/3d-printed-platform-jack/*

Krows, A. E. (1916), *Play production in America*, H. Holt and Company.

Krows, A. E. (1928), *Equipment for stage production: A manual of scene building*, D. Appleton Company.

Lahav, A., Saltzman, E. & Schlaug, G. (2007), 'Action representation of sound: audiomotor recognition network while listening to newly acquired actions', *The journal of neuroscience* **27**(2), 308–314.

Lartillot, O. & Toiviainen, P. (2007), A matlab toolbox for musical feature extraction from audio, *in* 'International Conference on Digital Audio Effects (DAFX)', Bordeaux, France, pp. 237–244.

Laumann, E. M. (1897), *La machinerie au théâtre depuis les Grecs jusqu'à nos jours*, Maison Didot.

Lefèvre, W. (2004), *Picturing machines 1400-1700*, MIT Press.

Lemaitre, G. & Heller, L. M. (2012), 'Auditory perception of material is fragile while action is strikingly robust', *The Journal of the Acoustical Society of America* **131**(2), 1337–1348.

Lemaitre, G., Houix, O., Visell, Y., Franinović, K., Misdariis, N. & Susini, P. (2009), 'Toward the design and evaluation of continuous sound in tangible interfaces: The spinotron', *International Journal of Human-Computer Studies* **67**(11), 976–993.

Leverton, G. H. (1936), *The production of later nineteenth century American drama: A basis for teaching*, AMS Press.

Logan, O. S. (1871), *The mimic world*, New-World: Philadelphia.

Logan, O. S. (1874), 'The secret regions of the stage', *Harper's New Monthly Magazine* **48**.

Lopez, M. (2015), 'Using multiple computer models to study the acoustics of a sixteenth-century performance space', *Applied Acoustics* **94**, 14–19.

Malloch, J., Sinclair, S. & Wanderley, M. M. (2007), A network-based framework for collaborative development and performance of digital musical instruments, *in* 'International Symposium on Computer Music Modeling and Retrieval', Springer, pp. 401–425.

Manning, P. (2013), *Electronic and computer music*, Oxford University Press.

Marshall, M. T. & Wanderley, M. M. (2011), Examining the effects of embedded vibrotactile feedback on the feel of a digital musical instrument, *in* 'New Interfaces for Musical Expression (NIME)', pp. 399–404.

Mathôt, S., Schreij, D. & Theeuwes, J. (2012), 'Opensesame: An open-source, graphical experiment builder for the social sciences', *Behavior research methods* **44**(2), 314–324.

McPherson, A. P., Jack, R. H. & Moro, G. (2016), Action-sound latency: Are our tools fast enough?, *in* 'New Interfaces for Musical Expression (NIME)'.

Medeiros, R., Calegario, F., Cabral, G. & Ramalho, G. (2014), Challenges in designing new interfaces for musical expression, *in* 'International Conference of Design, User Experience, and Usability', Springer, pp. 643–652.

Menzies, D. (2002), 'Composing instrument control dynamics', *Organised Sound* **7**(3), 255–266.

Merleau-Ponty, M. (2002), *Phenomenology of perception*, London: Routledge.

Miranda, E. R. & Wanderley, M. M. (2006), *New digital musical instruments: control and interaction beyond the keyboard*, Vol. 21, AR Editions, Inc.

Moro, G., Bin, A., Jack, R. H., Heinrichs, C. & McPherson, A. P. (2016), Making high-performance embedded instruments with bela and pure data instruments with bela and pure data, *in* 'International Conference on Live Interfaces (ICLI)'.

Mott, R. L. (1990), *Sound effects: Radio, TV, and film*, Focal Press London.

Moynet, J. (1874), *L'envers du théâtre: machines et décorations*, Hachette.

Müller-Sievers, H. (2012), *The cylinder: Kinematics of the nineteenth century*, Univ of California Press.

Murphy, D., Kelloniemi, A., Mullen, J. & Shelley, S. (2007), 'Acoustic modeling using the digital waveguide mesh', *IEEE Signal Processing Magazine* **24**(2), 55–66.

Murphy, D. T., Shelley, S. B., Foteinou, A., Brereton, J. S. & Daffern, H. (2017), 'Acoustic heritage and audio creativity: the creative application of sound in the representation, understanding and experience of past environments', *Internet Archaeology* .

Nagler, A. M. (1952), *A source book in theatrical history*, Courier Dover Publications.

Napier, F. (1962), *Noises off: A handbook of sound effects*, London: JG Miller Ltd.

Noë, A. (2004), *Action in perception*, MIT Press.

Noise Makers for the Futurist Concert of Noises (1914), *Sketch* **86**(JC 17).

Norman, D. A. (1999), 'Affordance, conventions, and design', *Interactions* **6**(3), 38–43.

Norman, D. A. (2010), 'Natural user interfaces are not natural', *Interactions* **17**(3), 6–10.

Norman, D. A. (2013), *The design of everyday things: Revised and expanded edition*, Basic Books.

Ovadija, M. (2013), *Dramaturgy of sound in the avant-garde and postdramatic theatre*, McGill-Queen's Press.

Pauletto, S. (2013), 'Film and theatre-based approaches for sonic interaction design', *Digital Creativity* **25**(1), 15–26 Design

Pauletto, S. (2017), The voice delivers the threats, foley delivers the punch: embodied knowledge in foley artistry, *in* M. Mera, R. Sadoff & B. Winters, eds, 'The Routledge Companion to Screen Music and Sound', Routledge.

Peterson, A. (1934), 'Stage effects and noises off', *Theatre and Stage* **1 and 2**.

Poepel, C. (2005), *On interface expressivity: a player-based study*, National University of Singapore.

Polanyi, M. (1967), *The tacit dimension*, University of Chicago Press.

Pollard, H. A. (1929), 'Show boat'.

Rath, M. & Rocchesso, D. (2005), 'Informative sonic feedback for continuous human–machine interaction—controlling a sound model of a rolling ball', *IEEE Multimedia Special on Interactive Sonification* **12**(2), 60–69.

Rittel, H. W. J. & Webber, M. M. (1973), 'Dilemmas in a general theory of planning', *Policy sciences* **4**(2), 155–169.

Rocchesso, D. (2014), 'Sounding objects in europe', *The New Soundtrack* **4**(2), 157–164 Europe

Rocchesso, D., Bresin, R. & Fernstrom, M. (2003), 'Sounding objects', *MultiMedia, IEEE* **10**(2), 42–52

Roesner, D. P. & Kendrick, L. (2011), *Theatre noise: The sound of performance*, Cambridge University Scholars.

Ronan, D. (2010), 'The physical modelling of a sitar'.

Rose, A. (1928), *Stage effects: How to make and work them*, Mackays Ltd, Chatham.

Rovan, J. & Hayward, V. (2000), 'Typology of tactile sounds and their synthesis in gesture-driven computer music performance', *Trends in gestural control of music* pp. 297–320.

Russolo, L. (1913), The art of noises, *in* 'Audio Culture: Readings in Modern Music', 2009 edn, Continuum, chapter 2, pp. 10–14.

Ryan, J. (1991), 'Some remarks on musical instrument design at steim', *Contemporary Music Review* **6**(1), 3–17 0749–4467.

Sayers, J. (2015), 'Prototyping the past', *Visible Language* **49**(3).

Sayers, J., Elliott, D., Kraus, K., Nowviskie, B. & Turkel, W. J. (2015), 'Between bits and atoms: Physical computing and desktop fabrication in the humanities', *A New Companion to Digital Humanities* pp. 1–21.

Selfridge, R., Reiss, J. D., Avital, E. J. & Tang, X. (2016), *Physically derived synthesis model of an Aeolian tone*, Audio Engineering Society.

Serafin, S. & De Götzen, A. (2009), 'An enactive approach to the preservation of musical instruments reconstructing russolo's intonarumori', *Journal of New Music Research* **38**(3), 231–239.

Serafin, S., de Götzen, A., Böttcher, N. & Gelineck, S. (2006), *Synthesis and control of everyday sounds reconstructing Russolo's Intonarumori*, IRCAM—Centre Pompidou.

Serafin, S., Erkut, C., Kojs, J., Nilsson, N. C. & Nordahl, R. (2016), 'Virtual reality musical instruments: State of the art, design principles, and future directions', *Computer Music Journal* **40**(3), 22–40.

Serafin, S. & Nordahl, R. (2005), Russolo's intonarumori: acoustical description and real-time simulation using physical models, *in* 'Australian Computer Music Conference (ACMC)'.

Serafin, S., Wilkerson, C. & Smith, J. O. (2002), Modeling bowl resonators using circular waveguide networks, *in* 'Proceedings of the 5th International Conference on Digital Audio Effects (DAFx-02)', pp. 117–121.

Shirley, F. A. (1963), *Shakespeare's use of off-stage sounds*, University of Nebraska Press.

Sinyor, E. & Wanderley, M. M. (2005), Gyrotyre: A dynamic hand-held computer-music controller based on a spinning wheel, *in* 'New Interfaces for Musical Expression (NIME)', National University of Singapore, pp. 42–45.

Smith, B. R. (1999), *The acoustic world of early modern England: attending to the O-factor*, University of Chicago Press.

Smith, J. O. (2010), *Physical audio signal processing: For virtual musical instruments and audio effects*, W3K Publishing.
URL: *https://ccrma.stanford.edu/ jos/pasp/*

Somerfield, J. (1934), *Behind the scenes*, Thomas Nelson and Sons Ltd.

Southern, R. (1944), 'The stage groove and the thunder run', *Architectural Review* **95**, 135–6.

Southern, R. (1952), A bristol theatre royal inventory, *in* 'Studies in English Theatre History', Society for Theatre Research, London, pp. 98 – 113.

Sparkfun (2015), 'Xbee shield hookup guide'.
URL: *https://learn.sparkfun.com/tutorials/xbee-shield-hookup-guide*

Spence, C. & Squire, S. (2003), 'Multisensory integration: maintaining the perception of synchrony', *Current Biology* **13**(13), R519–R521.

Staley, D. (2018), On the 'maker turn' in the humanities, *in* J. Sayers, ed., 'Making Things and Drawing Boundaries: Experiments in the Digital Humanities', U of Minnesota Press, pp. 32 –41.

Stappers, P. J. (2013), 'Prototypes as central vein for knowledge development', *Prototype: Design and craft in the 21st century* pp. 85–97.

Stappers, P. J. & Giaccardi, E. (2017), 'Research through design', *The encyclopedia of human-computer interaction* pp. 1–94.

Stappers, P. J., Sleeswijk Visser, F. & Keller, A. (2014), *The role of prototypes and frameworks for structuring explorations by research through design*, Routledge, pp. 163–174.

Stember, M. (1991), 'Advancing the social sciences through the interdisciplinary enterprise', *The Social Science Journal* **28**(1), 1–14.

Sterne, J. (2003), *The audible past: Cultural origins of sound reproduction*, Duke University Press, Durham and London.

Stowell, D., Robertson, A., Bryan-Kinns, N. & Plumbley, M. D. (2009), 'Evaluation of live human-computer music-making: Quantitative and qualitative approaches', *International journal of human-computer studies* **67**(11), 960–975.

Svanæs, D. (2013), 'Interaction design for and with the lived body: Some implications of merleau-ponty's phenomenology', *ACM Transactions on Computer-Human Interaction (TOCHI)* **20**(1), 8.

Talsma, D., Senkowski, D., Soto-Faraco, S. & Woldorff, M. G. (2010), 'The multifaceted interplay between attention and multisensory integration', *Trends in cognitive sciences* **14**(9), 400–410.

Tanaka, A. (2000), 'Musical performance practice on sensor-based instruments', *Trends in Gestural Control of Music* **13**(389-405), 284.

The Elements on the Stage: Thunder, Lightning, Wind, Rain and Fire (1913), *Scientific American* **108**(17), 2.

Thompson, E. A. (2004), *The soundscape of modernity: architectural acoustics and the culture of listening in America, 1900-1933*, MIT Press.

Townsend, K. (2013), 'Adafruit 10-dof imu breakout: User's guide'.
URL: *https://learn.adafruit.com/adafruit-10-dof-imu-breakout-lsm303-l3gd20-bmp180*

Turchet, L., Serafin, S. & Cesari, P. (2013), 'Walking pace affected by interactive sounds simulating stepping on different terrains', *ACM Transactions on Applied Perception (TAP)* **10**(4), 23.

Turkel, W. J. & Elliott, D. (2014), Making and playing with models: Using rapid prototyping to explore the history and technology of stage magic, *in* K. Kee, ed., 'Pastplay: Teaching and Learning History With Technology', University of Michigan Press.

Turnbull, R. B. (1951), *Radio and television sound effects: With drawings by the author*, Rinehart.

Vanderveer, N. J. (1980), 'Ecological acoustics: Human perception of environmental sounds', PhD Thesis.

Varela, F. J., Rosch, E. & Thompson, E. (1992), *The embodied mind: Cognitive science and human experience*, MIT Press.

Vincent, H. (1904), 'Stage sounds', *Strand* **28**, 417 – 422.

Ward, N. & Torre, G. (2014), Constraining movement as a basis for DMI design and performance, *in* 'New Interfaces for Musical Expression (NIME)'.

Woodhouse, J. (1992), 'Physical modeling of bowed strings', *Computer Music Journal* **16**(4), 43–56.

Wright, B. (2014), 'Footsteps with character: the art and craft of Foley', *Screen* **55**(2), 204–220.

Young, D. & Serafin, S. (2003), Playability evaluation of a virtual bowed string instrument, *in* 'Proceedings of the 2003 conference on New Interfaces for Musical Expression (NIME)', National University of Singapore, pp. 104–108.

Young, G. W. & Murphy, D. (2015), Hci models for digital musical instruments: Methodologies for rigorous testing of digital musical instruments, *in* 'International Symposium on Computer Music Multidisciplinary Research (CMMR)'.

Zimmerman, J. & Forlizzi, J. (2014), Research through design in HCI, *in* 'Ways of Knowing in HCI', Springer, pp. 167–189.

Zimmerman, J., Forlizzi, J. & Evenson, S. (2007), Research through design as a method for interaction design research in HCI, *in* 'Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, pp. 493–502.