



HAL
open science

Performance cartography, performance cartology: Musical networks for computational musicological analysis.

Jacob Hart

► **To cite this version:**

Jacob Hart. Performance cartography, performance cartology: Musical networks for computational musicological analysis.. Musicology and performing arts. University of Huddersfield, 2021. English. NNT: . tel-04152214

HAL Id: tel-04152214

<https://hal.science/tel-04152214v1>

Submitted on 27 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Title: Performance cartography, performance cartology: Musical networks for computational musicological analysis.

Author: Jacob Hart

University: University of Huddersfield (UK)

Viva

Date: 01/11/2021

Jury members: Simon Emmerson, De Montfort University; Robert Adlington, University of Huddersfield.



Titre : Performance cartography, performance cartology: Musical networks for computational musicological analysis.

Auteur : Jacob Hart

Université de délivrance : University of Huddersfield (Royaume-Uni)

Soutenance

Date : 01/11/2021

Composition du jury : Simon Emmerson, De Montfort University; Robert Adlington, University of Huddersfield.



University of HUDDERSFIELD

University of Huddersfield Repository

Hart, Jacob

Performance cartography, performance cartology: Musical networks for computational musicological analysis

Original Citation

Hart, Jacob (2021) Performance cartography, performance cartology: Musical networks for computational musicological analysis. Doctoral thesis, University of Huddersfield.

This version is available at <http://eprints.hud.ac.uk/id/eprint/35726/>

The University Repository is a digital collection of the research output of the University, available on Open Access. Copyright and Moral Rights for the items on this site are retained by the individual author and/or other copyright owners. Users may access full items free of charge; copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational or not-for-profit purposes without prior permission or charge, provided:

- The authors, title and full bibliographic details is credited in any copy;
- A hyperlink and/or URL is included for the original metadata page; and
- The content is not changed in any way.

For more information, including our policy and submission procedure, please contact the Repository Team at: E.mailbox@hud.ac.uk.

<http://eprints.hud.ac.uk/>

Performance cartography, performance cartology

Musical networks for computational musicological analysis

Jacob Hart

Ph.D. Thesis, Music

University of Huddersfield

School of Arts and Humanities

Department of History, English, Linguistics and Music

September 2021

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 725899).

TABLE OF CONTENTS.

Table of contents	2
List of figures	7
Copyright statement	11
Acknowledgments	12
List of abbreviations	13
Institutions	13
Algorithms	13
Other	13
Multimedia Appendices	14
Abstract	15
Introduction	16
Formal Introduction	16
Area of research and context	16
Research questions and desired outcomes	16
Vocabulary	18
The FluCoMa Project	19
The Musickers	20
Project timeline	24
Thesis Form and Structure	25
Chapter I. Theoretical Context & Methodology	27
1. Theoretical Context	27
a) Analysis	27
b) Creation	37
2. Research Questions	43
a) The question of time	43
b) Musical thought: inscription and materialisation	45
c) Network superposition and interfacing	47
3. Methodology	49
a) Source collection and archiving	49
b) Cartographical Analysis	50
c) Cartological Analysis	55
PART ONE: CARTOGRAPHICAL ANALYSES	
Chapter II. Surface Level Analyses	58
1. Surface Level Analyses of the Performances	58
a) Olivier Pasquet – <i>Herbig-Haro</i>	59
b) Lauren Sarah Hayes – <i>Moon via Spirit</i>	61

c) Rodrigo Constanzo – <i>Kaizo Snare</i>	62
d) John Burton – <i>Line Crossing</i>	64
e) Alice Eldridge – <i>Feedback Cello</i>	65
f) Sam Pluta – <i>Neutral Duo I & II</i>	67
g) Alexander Harker – <i>Drift Shadow</i>	68
h) Hans Tutschku – <i>Sparks</i>	70
i) Richard Devine – <i>Constructors</i>	72
2. Surface Level Analyses of the Instruments and Code.....	73
a) Olivier Pasquet.....	73
b) Lauren Sarah Hayes.....	79
c) Rodrigo Constanzo.....	84
d) John Burton.....	88
e) Alice Eldridge.....	92
f) Sam Pluta.....	93
g) Alexander Harker.....	95
h) Hans Tutschku.....	97
i) Richard Devine.....	100
Chapter III. Network Construction and Grammatisation.....	104
1. Tool Development and Pilot Studies.....	104
a) Requirements and Goals.....	104
b) Tools and Pilot-Studies.....	107
2. Global Level and Local Level Networks.....	115
a) Global level.....	116
b) Local level.....	122
3. Instrumental Level Networks.....	131
a) Static networks.....	131
b) Networks over time.....	138
Chapter IV. Sonorous Potentialities.....	142
1. Tool Development.....	142
a) Requirements and Goals.....	142
b) Tools.....	143
2. The Iterative Approach.....	147
a) Sample libraries.....	147
b) Parametric iterations.....	149
3. The Alternative Performances Approach.....	154
a) Analysis of other recordings.....	154
b) Automated performance – John Burton.....	157
c) Automated performance – Olivier Pasquet.....	159
PART TWO: CARTOLOGICAL ANALYSES	
Chapter V. Navigation of Networks.....	163
1. Static Sound Plots.....	163

a) Principles of Navigation	163
b) Static sound plot analysis.....	167
2. Superimposed sound plots.....	170
a) Rodrigo Constanzo.....	170
b) Lauren Sarah Hayes.....	173
c) John Burton.....	179
d) Olivier Pasquet.....	184
3. Dynamic sound plots.....	190
a) Tool development.....	190
b) Rodrigo Constanzo case study.....	191
Chapter VI. Points of Interface.....	194
1. A typology.....	194
2. Interface analysis.....	198
a) Aesthetic and functional points of interface.....	198
b) Hybrid points of interface.....	203
3. Code design.....	207
a) Digital interfaces.....	208
b) Case studies.....	211
Chapter VII. The Record and Musical Thought.....	215
1. The Record as fragment of sound.....	215
2. The Record as musical event.....	219
3. The Record as aesthetic object.....	222
Conclusion.....	226
Analysis summary.....	226
Analysis recap.....	226
Research questions.....	233
Methodology.....	237
The ontology of networks.....	240
Meaningless axes.....	240
Navigation of networks.....	242
Descriptors and perceptions.....	244
Future Research.....	246
Development of the tools.....	246
Paths for this thesis.....	247
Computational musicology without computers.....	249
Bibliography	251
Research	251
Books.....	251
Chapters	253
Articles.....	253

Conference Proceedings.....	256
Datasets	257
Sources.....	257
Theses	257
Writings.....	258
Blogs.....	259
Cited Works.....	260
Recordings	260
Pieces	260
Performances.....	260
Software.....	261
Appendix	263
1. FluCoMa objects overview.....	263
1.1. Slicing objects.....	263
1.2. Layer-finding objects.....	263
1.3. Object-finding objects.....	263
1.4. Audio description objects.....	264
1.5. Data storage objects.....	264
1.6. Data manipulation objects.....	264
1.7. Classification and regression.....	264
1.8. Audio transformation objects.....	265
2. Performance segmentations.....	266
2.1. Olivier Pasquet – <i>Herbig-Haro</i>	266
2.2. Lauren Sarah Hayes – <i>Moon via Spirit</i>	267
2.3. Rodrigo Constanzo – <i>Kaizo Snare</i>	269
2.4. John Burton – <i>Line Crossing</i>	270
2.5. Alice Eldridge – <i>Feedback Cello</i>	272
2.6. Sam Pluta – <i>Neural Duo I and II</i>	274
2.7. Alexander Harker – <i>Drift Shadow</i>	277
2.8. Hans Tutschku – <i>Sparks</i>	278
2.9. Richard Devine – <i>Constructors</i>	279
3. Software Overview.....	281
3.1. Rodrigo Constanzo.....	281
3.2. John Burton.....	282
3.3. Alexander Harker.....	283
4. Exploration of descriptor, dimensionality reduction and FFT size combinations with the metal resonance corpus.....	285
5. Iterative Approach Sound Plot Clusters: Metal Resonance > Wavefolder.....	289
6. Burton Alternative Performance Patch Modifications.....	291
7. Hayes Pass 1 Full Cluster Analysis.....	292
8. Interview extracts.....	296

8.1. John Burton.....	296
8.2. Olivier Pasquet.....	296
8.3. Lauren Sarah Hayes.....	297
8.4. Rodrigo Constanzo.....	298
8.5. Sam Pluta.....	300
8.6. Alice Eldridge and Chris Kiefer.....	301
8.7. Richard Devine.....	303
8.8. Alexander Harker.....	304
8.9. Hans Tutschku.....	305
9. Multimedia Appendices Index.....	308
9.1. 01_Code.....	308
9.2. 02_Data.....	309
9.3. 03_Audio.....	310
9.4. 04_Video.....	310
9.5. 05_Images.....	311

LIST OF FIGURES.

Figure 1: visualisation of the network superposition analogy.	47
Figure 2: segmentation of <i>Herbig-Haro</i>	60
Figure 3: segmentation of <i>Moon via Spirit</i>	61
Figure 4: segmentation of <i>Kaizo Snare</i>	63
Figure 5: the score generated for <i>Line Crossing</i> on the night of the FluCoMa performance...	64
Figure 6: segmentation of <i>Feedback Cello</i>	66
Figure 7: segmentation of <i>Neural Duo I</i>	67
Figure 8: segmentation of <i>Neural Duo II</i>	68
Figure 9: segmentation of <i>Drift Shadow</i>	69
Figure 10: segmentation of <i>Sparks</i>	70
Figure 11: segmenation of <i>Constructors</i>	72
Figure 12: <i>Herbig-Haro</i> setup.	73
Figure 13: <i>Herbig-Haro</i> software overview.	75
Figure 14: <i>Moon via Spirit</i> stage setup.	80
Figure 15: <i>Moon via Spirit</i> patch overview.	81
Figure 16: <i>Kaizo Snare</i> setup.	84
Figure 17: <i>Kaizo Snare</i> signal paths.	85
Figure 18: <i>Line Crossing</i> setup.	88
Figure 19: <i>Line Crossing</i> patch overview.	89
Figure 20: the subpatcher <i>human_block_param</i>	90
Figure 21: stage setup for <i>Feedback Cello</i>	92
Figure 22: setup for Pluta's <i>Neural Duo</i> pieces.	93
Figure 23: diagram of Pluta's performance software.	94
Figure 24: stage setup diagram for Harker's piece <i>Drift Shadow</i>	95
Figure 25: abstract diagram of Harker's <i>Drift Shadow</i> patch.	96
Figure 26: stage setup diagram given in score for <i>Sparks</i>	98
Figure 27: patch diagram for Hans Tutschku's <i>Sparks</i>	99
Figure 28: Devine's instrumental setup.	101

Figure 29: the <i>Autechre test patch</i> in its original state.....	107
Figure 30: the <i>Autechre test patch</i> after analysis.....	108
Figure 31: trumpet network visualisation	108
Figure 32: trumpet visualisation highlighting flow of energy.....	108
Figure 33: edge-weighted spring-embedded layout of the <i>Autechre test patch</i> network.	109
Figure 34: prefuse force-directed layout of the <i>Autechre test patch</i> network.	109
Figure 35: compound spring embedder layout of the <i>Autechre test patch</i> network.	109
Figure 36: yFiles organic layout of the <i>Autechre test patch</i> network.	109
Figure 37: attribute circle layout for all nodes in the <i>Autechre test patch</i> network.	110
Figure 38: group attribute circles layout for the <i>Autechre test patch</i> network.	110
Figure 39: hierarchical layout of the <i>Autechre test patch</i> network.	111
Figure 40: yFiles tree layout of the <i>Autechre test patch</i> network.	111
Figure 41: yFiles hierarchical layout of the <i>Autechre test patch</i> network.....	111
Figure 42: yFiles radial layout of the <i>Autechre test patch</i> network.	112
Figure 43: yFiles orthogonal layout of the <i>Autechre test patch</i> network.....	112
Figure 44: colours showing the different subpatches of each object in the <i>Autechre test patch</i> network.....	113
Figure 45: node size showing the degree of edges for each node in the <i>Autechre test patch</i> network.	114
Figure 46: filtering nodes by degrees in the <i>Autechre test patch</i> network. From left to right, in yellow: nodes with at least 2 degrees, 3 degrees and four degrees.	114
Figure 47: global-level network displayed with a prefuse force-directed layout.....	118
Figure 48: Bates Mill's Blending Shed performance area floor plan.....	122
Figure 49: <i>Reactive Flows</i> performance network displayed with the prefuse force-directed algorithm.	124
Figure 50: <i>Liminal Spaces</i> performance network displayed with the prefuse force-directed algorithm.	128
Figure 51: raw Max patch network visualisation of Constanzo's <i>Kaizo Snare</i> performance patch displayed with a prefuse force-directed layout.....	133
Figure 52: grammatised version of the <i>Kaizo Snare</i> network visualisation.....	134

Figure 53: the <i>Kaizo Snare</i> network insterted into the local level network.....	135
Figure 54: <i>Kaizo Snare</i> local and instrumental level network displayed with a hierarchical layout algorithm.	136
Figure 55: Metal resonance > wavefolder sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.	152
Figure 56: Raw mic > dirt sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.	153
Figure 57: Raw mic > dirt > wavefolder sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.	153
Figure 58: Hayes full sound plot t-SNE dimensionality reduction.	156
Figure 59: Hayes full sound plot UMAP dimensionality reduction.....	156
Figure 60: <i>Kaizo Snare</i> general segmentation sound plot described with MFCCs and displayed with t-SNE dimensionality reduction with clusters.....	168
Figure 61: <i>Kaizo Snare</i> entire composite visualisation displayed with the UMAP dimensionality reduction.	171
Figure 62: <i>Kaizo Snare</i> corpora composite displayed with the UMAP dimensionality reduction.	171
Figure 63: Pass 1 of Hayes' work displayed with a UMAP dimensionality reduction.....	173
Figure 64: Pass 1 of Hayes' sound plots using a t-SNE dimensionality reduction with clusterings.....	174
Figure 65: pass 2 of Hayes' work displayed with a UMAP dimensionality reduction.....	176
Figure 66: pass 3 of Hayes' work displayed with the UMAP dimensionality reduction.	177
Figure 67: pass 4 of Hayes' work displayed with the UMAP dimensionality reduction.	179
Figure 68: same scores sound plot for Burton's work displayed with the UMAP dimensionality reduction.	180
Figure 69: same human score, generated machine scores for Burton's work displayed with UMAP dimensionality reduction.....	180
Figure 70: both scores generated for Burton's work with a UMAP dimensionality reduction.	182
Figure 71: combined sound plot for the three configurations of Burton's alternative performances displayed with a UMAP dimensionality reduction.....	183

Figure 72: the *params 25* combined sound plot for Pasquet's work using MFCC descriptors, displayed with the UMAP dimensionality reduction algorithm. 185

Figure 73: the *params 25* combined sound plot for Pasquet's work using spectral shape descriptors, displayed with the UMAP dimensionality reduction algorithm. 185

Figure 74: combined *same settings* alternative performances for Pasquet's work using spectral shape descriptors, displayed with the UMAP dimensionality reduction algorithm. 187

Figure 75: *other rhythms* configuration combined with *same settings* for Pasquet's work, described with spectral shape descriptors and displayed with UMAP dimensionality reduction algorithm. 188

Figure 76: random parameters combined with same settings configurations for Pasquet's work, described with spectral shape descriptors and displayed with UMAP dimensionality reduction algorithm. 188

COPYRIGHT STATEMENT.

- i. The author of this thesis (including any appendices and/ or schedules to this thesis) owns any copyright in it (the “Copyright”) and s/he has given The University of Huddersfield the right to use such Copyright for any administrative, promotional, educational and/or teaching.
- ii. Copies of this thesis, either in full or in extracts, may be made only in accordance with the regulations of the University Details of these regulations may be obtained from the Librarian. Details of these regulations may be obtained from the Librarian. This page must form part of any such copies made.
- iii. The ownership of any patents, designs, trademarks and any and all other intellectual property rights except for the Copyright (the “Intellectual Property Rights”) and any reproductions of copyright works, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third Such Intellectual Property Rights and Reproductions cannot and must not be made available for use without permission of the owner(s) of the relevant Intellectual Property Rights and/or Reproductions.

ACKNOWLEDGMENTS.

There are many people I need to thank who have helped me in many ways over the course of this thesis. First, thanks to Pierre Alexandre Tremblay, leader of the FluCoMa project, who made this research possible: I thank him for giving me the wonderful opportunity to work with some amazing people and for his unfaltering support throughout my research. Thanks also to my supervisors Frédéric Dufeu and Owen Green for their precious insights and whose guidance pulled me through the difficult times. I really couldn't have wished for a better team behind me.

I would like to thank all the musicians on the FluCoMa project, my subjects of study: John Burton, Rodrigo Constanzo, Richard Devine, Alice Eldridge, Alex Harker, Lauren Sarah Hayes, Olivier Pasquet, Same Pluta and Hans Tutschku. They are some of the most talented and inspiring people I have ever met: when motivation was low, all I had to do was listen to their music. I thank them for their help and openness to cooperate with a young musicologist.

I would like to thank Michael Clarke, leader of the IRiMaS project, who welcomed me to his sessions and pointed me towards many stimulating ideas that are intertwined within this thesis. I also thank Antoine Bonnet who put me on the path of musicological research and has continued to support me for many years. Thanks also to the staff members at the University of Huddersfield, notably the members of CeReNeM who have created one of the most stimulating environments for a young researcher in music to find themselves in.

I thank my fellow PhD students and friends at Huddersfield: notably James Bradbury who, beyond his precious friendship, played a huge role in making me the programmer I am today; and Laurens Van der Wee, with whom I developed many ideas over numerous English breakfasts. Thanks, and much love to my friends and colleagues Andie Brown, Francesco Cameli, Niamh Dell, and Richard Piatak.

I am lucky to have many people who support me in my life. Thanks to all my friends, notably to Tolga for putting up with my ramblings over all these years and always being interested. Thank you to my family: my parents Georgia and Eddie, who have always supported me, even when I told them I was going to do musicology; my brother Max, and my grandparents Jill and John. Finally, eternal thanks to Anaïs: you are the loveliest person I know, and I am the luckiest person alive to get to spend every day with you.

LIST OF ABBREVIATIONS.

Institutions.

- *AACM*: Association for the Advancement of Creative Musicians.
- *CCL*: Creative Coding Laboratory.
- *CeReNeM*: Centre for Research in New Music.
- *ERC*: European Research Council.
- *FluCoMa*: Fluid Corpus Manipulation.
- *HCMF*: Huddersfield Contemporary Music Festival.
- *IRCAM*: Institut de Recherche et Coordination Acoustique/Musique.
- *IRiMaS*: Interactive Research in Music as Sound.
- *MaMI*: Music and Moving Image.
- *NIGMS*: National Institute of General Medical Sciences.
- *PARIESA*: Practice and Research in Enactive Sonic Art
- *TaCEM*: Technology and Creativity in Electroacoustic Music.

Algorithms.

- *FFT*: Fast Fourier Transform.
- *HPSS*: Harmonic Percussive Source Separation.
- *KNN*: K-Nearest Neighbour.
- *LZW*: Lempel-Ziv-Welch.
- *MDS*: Multidimensional Scaling.
- *MFCC*: Mel-Frequency Cepstral Coefficient.
- *MLP*: Multilayer Perceptron.
- *NMF*: Non-Negative Matrix Factorisation.
- *PCA*: Principal Component Analysis.
- *t-SNE*: t-Distributed Stochastic Neighbour Embedding.
- *UMAP*: Uniform Manifold Approximation and Projection.

Other.

- *ANT*: Actor-Network Theory.
- *CCE*: Creative Coding Environment.
- *CLI*: Command Line Interface.
- *DAW*: Digital Audio Workstation.
- *DOI*: Digital Object Identifier.
- *EDM*: Electronic Dance Music.
- *GUI*: Graphical User Interface
- *HID*: Human Interface Device.
- *HISS*: Huddersfield Immersive Sound System.
- *LDR*: Light Dependent Resistor.
- *LLLL*: Lisp-Like Linked List.
- *MIDI*: Musical Instrument Digital Interface.
- *VST*: Virtual Studio Technology.

MULTIMEDIA APPENDICES.

There are several elements that accompany this thesis in digital form (Hart, 2021), they can be downloaded with the link associated to its DOI¹. Appendix 9 gives the reader an account of the content and hierarchical structure of these elements.

When a digital element is referred to in the text, there shall be a footnote with the indication *D.E.* (Digital Elements), followed by a file path from the root folder *Performance Cartography, Performance Cartology: Multimedia Appendices*. For example: *D.E.: 03_Audio>01_Performance_Audio>01_pasquet.mp3*.

The reader shall be asked to engage with two principal elements for which standalone programs have been created for Mac and Windows (Windows 10+) at the following locations:

- *01_Code>02_Tools>05_Sound_Plot_Navigation_Standalone*.
- *01_Code>02_Tools>06_Segmentation_Explorer_Standalone*.

There are also videos demonstrating how these programs work.

Please note: in some of the tools and demonstration videos, the root folder is referred to as *DIGITAL ELEMENTS*. The name of this folder has since been changed to *Performance Cartography, Performance Cartology: Multimedia Appendices*.

Note that for these elements to function, the entire folder must be downloaded, and its structure retained.

If they are curious, the reader may also consult some of the secondary elements. It may be necessary to obtain the corresponding programs for these secondary elements, notably: Max² (version 8+), Cytoscape³ (version 3.8.0+) (Shannon et al., 2003), SuperCollider⁴ (version 3.9.3+) and a text editor such as Visual Studio Code⁵ (version 1.6+).

There is also a collection of supplementary multimedia appendices on a git repository⁶ of various, eventually unused, pieces of software that were developed during the project – these are experiments, drafts and sketches that are referred to throughout the thesis. These contributed to the development process of the final software that was used, and the reader may consult these if they wish.

¹ <https://doi.org/10.34696/p0y3-7f65>

² <https://cycling74.com/>

³ <https://cytoscape.org/>

⁴ <https://supercollider.github.io/>

⁵ <https://code.visualstudio.com/>

⁶ <https://github.com/jdchart/supplementary-multimedia-appendices>

ABSTRACT

This research seeks to track the creative processes of a group of musicians who were commissioned to make music with a set of digital tools created by the FluCoMa project. These tools offer many solutions for dealing with digital audio, notably large collections of sounds.

The varied and multidisciplinary natures of the case studies mean that traditional forms of analysis would miss essential parts of these practices. Subscribing to contemporary approaches, this research proposes a methodology for analysis articulated around the idea of musical networks. Here, musical practice is conceived of as the configuration of, and existence within, networks of entities where musicking occurs.

The methodology is proposed in two parts: a cartographical, descriptive analysis that seeks to map these networks; and a cartological, interpretative analysis that seeks to inspect the nature of these networks. The methods are illustrated and developed by the case studies, grounding them in real musical practice.

With these methods, this research looks to address three primary questions: how can a network-oriented analytical stance account for temporality in musicking? Can we consider the network as constituting a materialised form of musical thought? And what are the ways in which entities of networks are configured?

Keywords: networks, creative process, electroacoustic music, creative coding, analysis, computational musicology.

Word count: 81114

INTRODUCTION.

Formal Introduction.

I joined Pierre Alexandre Tremblay's Fluid Corpus Manipulation (FluCoMa) project in September 2018. At the time, all I knew was that the two postdoctoral researchers on the project (Owen Green and Gerard Roma) would be creating a set of digital tools for techno-fluent composers, and that a group of nine musicians had been commissioned to use the tools for a performance. At this point, the actual nature of these tools was still a mystery to me (as it was to the composers) – I knew only that my job would be to track the artists' creative processes from a musicological perspective.

Area of research and context.

I consider my primary area of research to be in computational musicology and musical analysis. My task on the FluCoMa project was to track the creative process of two cohorts of musicians as they participated in the development of a set of enabling technologies for dealing with large amounts of digital audio. This context, and the breadth of practices under examination presented clear challenges to the orthodox, primarily text-oriented musicology in which I was trained. To deal with this enlarged scope, I was led to develop new theoretical tools and new methods. These theories and methods form the core outcomes of this research. The analyses of the musicians' work serve as case studies against which I will critically assess the viability, applicability, and possible future directions of these developments.

Research questions and desired outcomes.

The following research questions emerge from a theoretical context that has developed over the course of the research, explained in detail in Chapter 1. I start from an initial theoretical context that embeds itself in reflection around musical creation and musicological analysis. I depart from a contemporary perspective which critiques a traditional approach to analysis that would attempt to fit music into pre-existing models and fetishize the object of the score. This is an approach that does not cater well to contemporary practices, especially those examined through the case studies in this thesis. Indeed, as the research has progressed, I have made sure that it is these musical practices that drive my thinking, allowing me to ground my analysis in reality.

The perspective on musical creation I take in this research follows from and builds on Small's (Small, 1998) depiction of musicking as networks of related and meta-related entities. I build on this in Chapter 1 via related perspectives from Georgina Born, Bruno Latour and Tim Ingold, among others. I argue for an approach to analysis that could be considered by analogy to rambling: to exist in proximity to a musical network and consider it from new perspectives to gain access to other listenings. This has several implications for analysis: an analysis will necessarily become part of this network and have an impact upon it; the network offers itself as a rich site for analysis where the *musical thought* (this term is developed in Chapter 1) of its participants may be inscribed within the configuration of its various physicalities; it is possible to consider musical activity across two modes of engagement with musical objects (or *Records*, this term is developed in Chapter 1): *measurement* and *manipulation*.

From this theoretical context, I draw the following research questions: first, how can this network-oriented analytical stance account better for temporality in musicking? After a historical sweep of this network-oriented approach, I propose that it could be useful to use the network not only as a static account of things that have happened, but as something that allows for things to happen. This reflects my vision of musical practice as being the configuration of and existence within networks, rather than the creation of a single object. How can we approach this in analysis? Is it feasible to try and gain an image of the entire breadth of *potentialities* of a musical network? How can we do this in a tangible way?

Next, can we conceive of musical thought as being inscribed within a musical network? Is it possible to access a materialised form of musical thought through analysis? I propose the idea of examining the relationships between different entities, and how they are measured and manipulated across three dimensions: energetic, temporal, and organisational. With my case study analyses, I wish to examine the validity of the idea that the musical network offers a materialisation of musical thought.

Finally, what are the ways in which musicians configure these networks? I propose that there are critical entities in these networks where energy is translated; these points are interwoven with physicalities where measurement and manipulation occur. Where can these critical entities be found? How can we examine their nature? What is the nature of the translations that occur? Is it possible to draw a categorisation or a typology of them?

Alongside the case studies through which I try to answer these questions, the methodology that has emerged is also a significant output of this research. As with the theory, this methodology is described in detail in Chapter 1, and is developed across the thesis. It is divided into two complementary parts: *cartographical* analyses (visualisation, materialisation of various networks) and *cartological* analyses (reading and finally interpretation of these networks). These analyses are derived from a large archive of various sources: interviews, technical setup videos, digital instruments, and code. A supplementary research outcome is also the collection of various tools that I have created over the research to aid my analytical approach (these are discussed accordingly over the course of the thesis).

Vocabulary.

Before continuing, these are some precisions concerning some of the terminology used across the thesis:

- *Physical* and *material* domains, *physicalities* and *materialities*: here, the material refers to the traditional conception of the word – a tangible, concrete object in the real world. The physical refers to entities in a network which allow for something to happen. For example, a material button and a digital button are both physicalities.
- The *mechanical*, the *mechanics* of a network: this refers to the systematic structure of a network, the way in which entities come together to allow for things to happen.
- *Musickers*, *musicking*: this research deals with a group of artists that have many areas of activity: musical composition, performance, instrument making, coding, teaching, research, and writing. As such, it would be reductive to refer to them simply as *composers*. To address this issue, where possible, I refer to these people as *musickers*, borrowing from Small's notion of *musicking*. In parallel, when possible, I refer to their work not as *pieces*, *compositions*, *improvisations*, or *performances* but as *musicking*.
- *Musical thought*: this term is detailed further in Chapter 1; however, it is necessary to give an initial explanation here. Musical thought can broadly be conceived of as the ways in which humans, and in the case of this research, musickers, will interact with musical objects. It is inherently intangible, but I will argue that it can materialise through the configuration and manipulation of objects.

There are other terms used across this thesis that refer to specific parts of my theoretical context such as: the *Record*, *network superposition*, *points of interface* and *interfacing*. These terms have specific meanings in the context of my research and are explained fully in Chapter 1.

The FluCoMa Project

The FluCoMa project began in 2018 and is a European Research Council (ERC)-funded project at the University of Huddersfield, UK. The tools bring solutions for manipulating large collections of sound (or *corpora*, to use the project's terminology) using a variety of techniques. They're intended to animate future research around the musical questions thrown up by working with corpora. They were developed as a set of agnostic objects in C++, and a framework was created that allowed the team to easily compile them into objects for a variety of Creative Coding Environments (CCEs): Max¹, SuperCollider², Pure Data³ and a Command Line Interface (CLI). The objects were developed as two toolsets that would eventually be combined into one single package. The algorithms that the team coded and brought to these systems already existed but were either unimplemented in these environments or implemented in an unsatisfactory manner. The tools are also divided into real-time and non-real-time objects.

A full overview of the objects can be found in Appendix 1, here, I give a summary. The first package (Tremblay et al., 2019) deals mainly with audio decomposition and description. Decomposition is achieved in three forms: finding slices, layers, and objects. Description is achieved with an implementation of a wide variety of audio descriptors, and with solutions for deriving different types of statistics from them.

The second package (Tremblay et al., 2021) deals with the manipulation of the large datasets that can be derived from the first package. There are objects for storing data, and objects for querying and processing this data with techniques like normalisation, standardisation, and dimensionality reduction. There are also tools for performing regression and classification on the datasets. There are also tools for transformation and concatenation of audio.

¹ <https://cycling74.com/>

² <https://supercollider.github.io/>

³ <https://puredata.info/>

The FluCoMa project operates with two underlying design philosophies (Green et al, 2018): first, it was important for the development to be *practice-driven*. Therefore, a group of musickers was commissioned to use the tools *during* the development process. A forum⁴ was set up that allowed the musickers to give feedback on the tools, not only for bugs and other practical issues, but also on the interface and the way the tools were being used. This was to facilitate a collaborative design of the tools between developers and users.

Second, the team made it clear from the start that they wanted to avoid the idea of *blackboxing* – delivering a set of tools where the user would have very little control over the underlying algorithms and only giving control of a handful of top-level parameters. There is a difficult balance to find between completely open, bare-bones algorithms which would leave the user overwhelmed with choices and that would demand a deep understanding of how the algorithms function; and software that would make all the decisions for you and would assume it knows what you wish to do with it.

It is worth noting also that there is a distinct workflow that is proposed by the project through the various help files and reference documents – that of decomposition, description, manipulation. This workflow is inevitably inscribed within the tools and their presentation and must be considered when analysing their use.

The Musickers.

A first cohort of musickers was commissioned to use the first package of tools, with their work premiered at a concert programmed for November 2019. The second cohort used the first and second packages for a concert programmed for February 2021 (postponed to June 2021 due to the COVID-19 crisis). This biographical information⁵ emerges from three main sources: personal websites, presentations given at the FluCoMa plenaries and information derived from interviews.

John Burton

Burton, a.k.a. Leafcutter John, came to prominence in 2003 with the release of his critically acclaimed album *The Housebound Spirit*⁶. The album is a collage of a wide variety of processing techniques on a broad range of sounds. Computer processing

⁴ <http://discourse.flucoma.org>

⁵ Presented by cohort and in alphabetical order.

⁶ LEAFCUTTER JOHN. *The Housebound Spirit*. Planet-Mu Recordings, 2003.

and specifically Max have always played a big role in his music production: he explained⁷ that he acquired a laptop to do his final project at art school with which he discovered digital audio techniques. Much of his work since⁸ seems to follow a more ambient, tape-music inspired trajectory, or contemporary jazz with his band Polar Bear⁹. He also cites¹⁰ Parmegiani's works from the late seventies as being a big inspiration. In a performance context, in recent years Burton has been known for his bespoke light interface¹¹ – a board of Light Dependent Resistors (LDRs) which runs through Max, mainly triggering sample playback.

Rodrigo Constanzo

Constanzo obtained his PhD (Constanzo, 2016) from the University of Huddersfield in 2016 and is currently the Deputy Head of Popular Music at the Royal Northern College of Music, Manchester. Constanzo trained as a pianist before pivoting into percussion, and now undertakes a wide range of activities from programming, 3D printing and instrument making, to performing and improvising. His thesis was centred around improvisation and proposes a framework for self-analysis of improvised performances. His music can be described as harsh, percussive, and noisy, with a distinct taste for distortion and concatenative synthesis. Each project he works on has a clear identity¹², and he will typically build a bespoke system around each performance.

Lauren Sarah Hayes

Hayes is currently an Assistant Professor of Sound Studies in the School of Arts, Media and Engineering at Arizona State University. She is also the founder and lead researcher of the Practice and Research in Enactive Sonic Art (PARIESA) research group. Hayes trained as a pianist, and now considers herself more broadly as an improviser and sound artist. She has several areas of interest, notably her PhD work on haptic feedback (Hayes, 2014), augmented hybrid instruments and embodied knowledge (Hayes, 2019). She has worked with Max since she was an undergraduate

⁷ Appendix 8.1.1.

⁸ For example: LEAFCUTTER JOHN. *The Forest and the Sea*. Staubgold, 2006; LEAFCUTTER JOHN. *Tunis*. Tsuku Boshi, 2010.

⁹ POLAR BEAR. *Held on the Tips of Fingers*. Babel Label, 2005.

¹⁰ Appendix 8.1.2.

¹¹ LEAFCUTTER JOHN. *Nightless Night*, XOYO, London, UK, 2014.

¹² For example: his 2015 live-sampling and processing software *Cut Glove*; his 2016 performance *Rhythm Wish*.

and has built up a modular performance instrument over the past decade. She has also released several albums¹³ that have gained critical acclaim.

Olivier Pasquet

Pasquet acquired his PhD in musical composition and non-standard architecture from the University of Huddersfield in 2018 (Pasquet, 2018). Pasquet has worked all over the world, and mainly makes installations that combine music and architecture¹⁴. Within his musical practice, he even incorporates the use of architectural tools such as the program Rhinoceros 3D¹⁵. He has also worked in theatre and has been strongly associated as a composer with the Institut de Recherche et Coordination Acoustique/Musique (IRCAM) for many years. His music has a very strong aesthetic profile with a style that is instantly recognisable – he explains¹⁶ that he tends to work at a micro-granular level with sound, with small rhythmic patterns, and layers.

Richard Devine

Devine is a well-known figure in the world of Electronic Dance Music (EDM) and sound design; he is known to have worked on presets for Native Instruments, and to be consulted upon by many creators of digital audio plugins and modular synthesizer modules. He has also collaborated with corporations such as Google and is now a Senior Content Producer for Apple. Devine is also associated with the label Warp Records and has worked with well-known artists like Aphex Twin and Autechre. Devine has much experience working with digital and analogue technologies, but at the time he came to the project, he was working mostly with modular synthesizers, as can be heard in his 2018 album *Sort\Love*¹⁷.

Alice Eldridge

Eldridge obtained her PhD (Eldridge, 2007) in Computer Science and AI from the University of Sussex in 2007 and is now the director of the Sussex Humanities Lab at the same university. She has a long history of using machine learning in her creative practice (Eldridge, 2007). One project that has been at the centre of her practice for the past few years is that of the *Feedback Cello* (Eldridge and Kiefer, 2017): this is a project

¹³ For example: HAYES, Lauren Sarah. *Manipulation*. Pan y rosas discos, 2016 ; HAYES, Lauren Sarah. *Embrace*. Superpang, 2021.

¹⁴ For example: his 2018 installation *Dual Cornographs* at Silos City, Buffalo, NY, USA; his 2014 *hr 8799* installations at the National Taiwan museum of fine arts, Taipei, Taiwan.

¹⁵ <https://www.rhino3d.com/fr/>

¹⁶ Appendix 8.2.1.

¹⁷ DEVINE, Richard. *Sort\Love*. Timesig, 2018.

she works on with Chris Kiefer, also at Sussex, and consists of an augmented cello equipped with electromagnetic pickups allowing the performer to induce feedback to each string independently. She is also known for her work in the field of eco-acoustics (Eldridge et al., 2018).

Alexander Harker

Harker is currently a lecturer at the University of Huddersfield. He gained his PhD in 2011 at the University of York (Harker, 2011). Harker is well-known in the Max community and has participated in the creation of many extensions and externals for it, notably his own set of externals¹⁸; the Huddersfield Immersive Sound System (HISS) tools¹⁹ (Harker and Tremblay, 2012); and more recently his Framelib²⁰ (Harker, 2017) framework. Much of Harker's recent activity has been in coding, however he also has a substantial body of compositional works under his belt. The most recent of these date back to 2010 with *Fluence*²¹ and 2011 with *Fractures*²².

Sam Pluta

Pluta is currently a Professor of Music at the University of Chicago. He received his DMA (Pluta, 2012) in composition and electronic music in 2012 at Columbia University. He is also currently the technical director for and a performer in the *Wet Ink Ensemble*. Pluta is an improviser and performs often with other musicians, performing live-processing on their playing. He has also worked with modular synthesizers. Pluta's typical setup includes a joystick, computer, and iPad which he has been using for some ten years, all running processes in SuperCollider.

Hans Tutschku

Tutschku is currently a Professor of Composition and the director of the electroacoustic studio at Harvard University. He obtained his PhD in composition in 2003 (Tutschku, 2003) from the University of Birmingham. Over the course of his career, he has won several prestigious awards for his work, such as the Hanns Eisler Preis, CIMESP Sao Paulo, Prix Ars Electronica and the culture prize of the city of Weimar. Tutschku trained as a pianist, and the piano often plays an important role in

¹⁸ <http://www.alexanderjharker.co.uk/Software.html>

¹⁹ <https://research.hud.ac.uk/institutes-centres/cerenem/projects/thehisstools/>

²⁰ <https://github.com/AlexHarker/FrameLib>

²¹ HARKER, Alexander. *Fluence*, clarinet and Max, 2010.

²² HARKER, Alexander. *Fractures*, fixed media, 2011.

his works. He often works with ensembles of musicians and live-processing, but also produces works of fixed-media and installations.

Project timeline.

My research timeline has been heavily tied to the FluCoMa project timeline. I have frequently had to reconfigure various steps of my research according not only to the way in which the research itself and the methodology have developed, but also to tie-in with how the project has advanced. There have been some key moments over the project which had to be considered for the progress of my research: notably the access the musickers had to the tools in their various states of development, the series of plenaries which occurred over the three years, and the two concerts which play an important role in my approach. This is a broad overview of how the project developed over the three years I was part of the research team:

- *September 2018*: Plenary #1. The musickers are introduced to the first set of tools.
- *July 2019*: Plenary #2. The first toolbox is presented and made available in a stable state to the musickers.
- *November 2019*: Plenary #3 and Concert #1. The first batch of musickers present their work their work and perform at the first gig as part of Huddersfield Contemporary Music Festival (HCMF) 2019.
- *July 2020*: Plenary #4. The second toolbox is presented and made available in a stable state to the musickers.
- *February 2021*: Plenary #5. The second batch of musickers present their work.
- *July 2021*: Plenary #6 and Concert #2. The second batch of musickers perform at the second gig as part of Dialogues Festival 2021.

I inserted my research into this timeline in a way that somewhat follows the path of this thesis. During the first year, the musickers' engagement with the tools was relatively low, and I spent a lot of time developing the theoretical context which is outlined in Chapter 1. From November 2019, analysis began with the first set of case studies. A lot of this time was spend forging the first iterations of my methodology and developing tools that would help me address the research questions that emerged from the theoretical context. This part of the analysis broadly corresponded to the first part of this thesis, the *cartographical* analyses. The third year was spent developing this methodology and tools with further case studies and performing the *cartological* part of my analyses.

In 2020, the COVID-19 crisis struck, and several things in the project were reconfigured: plenaries 4 and 5 were conducted online, and the second concert that was initially planned for February 2021 was postponed to July. My initial plan was to conduct a first pass of analyses on the first batch of musickers until the second concert, then perform a second pass on the second batch.

This had to be reconfigured for several reasons: considering the postponement of the second concert to late July and my given timeframe, I realised that a full analysis of the second cohort of composers was not going to be feasible. Furthermore, the breadth of my analyses of the first cohort was producing more than enough material to sensibly discuss within the limits of one research project. These first four case studies had sufficed to refine my methodology from its original conception to the point where it is today. Therefore, the main body of my analysis is centred around the first cohort – there are some surface level analyses of the second, but they do not constitute the primary focus of this research.

I performed a series of interviews at various points of the project – notably a first batch at the beginning, and a pass after each of the concerts. This was part of an extensive archiving of several different sources: video presentations, code, sample and rehearsal audio, writings.

Thesis Form and Structure.

As suggested above, this thesis is divided into two large sections which reflect its methodology: *cartographical* analyses, and *cartological* analyses. These terms are explored and explained in greater detail in Chapter 1, however for now we can understand them respectively as *descriptive* analyses and *interpretative* analyses.

The *cartographical* section is comprised of three chapters. The first looks at what I have called *surface level analyses* – comprised of analyses of each of the performances and each of the instrumental setups and the software. The second chapter focuses on network construction – here we shall see how I constructed network visualisations of the networks under analysis and an initial descriptive analysis we can draw from them. Finally, the third chapter is devoted to the idea of *sonorous potentialities*. Again, this will discuss how I made a series of visualisations looking to reveal the sonorous potentialities of each network, and how the methodology developed over the course of each case study.

The *cartological* section is again divided into three chapters. The first looks at the navigation of networks, notably the sonorous networks produced in the previous section, proposing an interpretative analysis of them. The second chapter is focused on the notion of *points of interface*, another term discussed in Chapter 1. Finally, the third chapter is somewhat conclusive in nature, drawing elements from all the previous chapters and interrogating the idea of musical thought and the Record.

All of this is preceded by a chapter that outlines the theoretical context of the project, the research questions that have emerged from it, and outlining in more detail the methodology that has been forged over the course of the research.

Chapter I. THEORETICAL CONTEXT & METHODOLOGY.

1. Theoretical Context.

a) *Analysis.*

The current analytical zeitgeist.

I came to this project with a traditional approach to musicology and analysis. In the context of French academia in which I trained, there seem to be two broad strands along which one must build their analysis: a literary, poetic, even philosophical interpretation of a piece and its aesthetic questions; and the application of a host of methods for forging detailed segmentations and global abstractions of a piece: identification of form, Schenkerian analysis, pitch-class set theory, Neo-Riemannian theory.

Broadly, these analysis systems have two main goals. The first is to be able to talk about a piece in an articulate manner: this is often to the service of the first literary strand of analysis. One task of the musicologist is to convey their analysis of a piece of music in a way that is useful and meaningful in a narrative form. Not only do segmentations and abstractions allow us to refer easily to local events, they also give us the possibility to discuss larger-scale events. Indeed, a music deploys itself through time in a particular manner, and it is useful to have tools that allow us to conceive of and talk about these gestures in an abstract, out-of-time way.

The second goal is to gain access to some inner 'truth', some hidden element in a piece that would reveal its workings. Some analysts are content to take this as a means unto itself, others will go so far as to interpret this in a hermeneutic manner (bringing us again to the literary strand of analysis). The purpose of this research is by no means to discredit these methods – my methodology extracts and uses many of their global principles. However, on a conceptual level it will be useful to consider these characteristically 19th and 20th centenarian methods in this caricatural manner.

At the University of Huddersfield, I was confronted with what I consider to be a more contemporary approach. This environment is very much practice-driven in nature: practice-led research and development are one of the core principles of the FluCoMa project. I was also a regular participant in the weekly sessions of two research groups which share these kinds of ideas and approaches: Centre for Research in New Music (CeReNeM), and the Interactive Research in Music as Sound (IRiMaS) project (Dufeu et al., 2019).

The contemporary approach to musicology and social phenomena in general can be characterised by Born (Born, 2010) who proposes “*moving beyond the terms musicology, ethnomusicology, popular music studies, the sociology and psychology of music [...] to a new, integrated music studies*” (Born, 2010, pp. 205-206). This is achieved through two moves: an expanded conception of the ontology of music, and a departure from traditional musics of study of Western musicology. Contemporary music practices demand of the analyst a broader approach that require skills in fields that depart from traditional musicological approaches. This means that the methodologies for analysis are inherently driven by the music – an approach that starkly contradicts traditional ones, which ultimately seek to frame music within a pre-existing analytical model. This kind of approach can work in the context of a musicology that focuses on a certain type of music written over a very restricted period and geographical location; but when we begin to consider the countless other types of musics that exist and have existed, it quickly crumbles.

Born also evokes the idea of the production of musical knowledge as being *performative* in nature. This is something that echoes Cook’s ideas around performance and analysis (Cook, 1999). Initially, he brings into question the traditional Chomskyan perspective on the dualism of competence as “*abstract knowledge on which any rule-based system depends*” (Cook, 1999, p. 242) and performance as “*the use of that knowledge in any given situation*” (Cook, 1999, p. 242): the idea of performance as being the performance of a Platonic work. This is clearly the perspective held by the traditional musicological approach and inflects its Work-driven approach to analysis. Cook suggests that “*performance should be seen as a source of signification in its own right*” (Cook, 1999, p. 247), giving it a much more important place, and regarding each performance as a separate, aesthetic object for analysis. This steering away from the Work-centric perspective is characteristic of the contemporary approach.

Following from this, when regarding the relationship between analysis and performance, he offers an idea that renders analysis itself as performative. He proposes that what is important in analysis is not “*what it represents but what it does [...] what it leads you to do*” (Cook, 1999, p. 249). To be performative is for something to “*have meaning by virtue of what it does*” (Cook, 1999, p. 255). Analysis should be seen as a “*means of posing articulate questions*” (Cook, 1999, p. 248).

Cook cites David Lewin, who says that analysis is “*not an aid to perception, or to the memory of perception; rather, we are in the very act of perceiving*” (Lewin, 1986, pp. 381-

382). He also cites F.E. Maus who explains that “[analysis] could be regarded as translations. That is, analyses can be seen, not as pale copies of a determinate original, but as ways of exploring musical compositions in an ongoing process in which there is no point in distinguishing between making and finding the qualities of music” (Maus, 1993, p. 70). Here, we have a clear demonstration of how the very ontology of analysis has shifted – we no longer approach a piece of music in search of some truth, to demonstrate the existence of some element that may support our poetic reading of the musical experience; here, analysis is simply conceived of as a means to experience the music differently. As such, our relationship with the music in question evolves, and we may gain access to *new listenings* of a piece.

This type of approach reflects the issues discussed by Abbate (Abbate, 2004). She bases her ideas around the statements of philosopher Vladimir Jankélévitch, who insists that “*real music is music that exists in time, the material acoustic phenomenon*” (Abbate, 2004, p. 505). It is an *irreversible experience* – and therefore is necessarily of the order of the drastic domain, “*connoting physicality, [...] knowledge that flows from drastic actions or experiences and not from verbal reasoning*” (Abbate, 2004, p. 510), and not the gnostic domain, “[drastic’s] *antithesis [...] implies not just knowledge per se but making the opaque transparent, knowledge based on semiosis and disclosed secrets, reserved for the elite and hidden from others*” (Abbate, 2004, p. 510).

Abbate discusses some of Cook’s work (Cook, 2001) and seems to suggest that he misses the mark somewhat. Despite making steps in the right direction in terms of analysis, she brings into question his idea that “*analysing music as performance does not necessarily mean analysing specific performances or recordings at all*” (Cook, 2001, p. 9), and that “*performance [...] has been scripted into the work*” (Abbate, 2004, p. 508). She translates this as the idea that “*musical works take heed of the ‘performance network’ – the channels between composer, performer, material realisation and listener*” (Abbate, 2004, p. 508). For Abbate, this approach is still fatally flawed in as much as it tries to “*domesticate what remains nonetheless wild*” (Abbate, 2004, p. 508).

She remains pessimistic about what a gnostic analysis¹ could bring to the table. She is wary of the way interpreted “*ideas and truths in music are made monumental and given aura by music*” (Abbate, 2004, p. 520). When discussing the dichotomy of music and image in cinema or advertising, she describes music as “*sticky*” (Abbate, 2004, p. 523).

¹ Musical commentary, descriptive analysis, musical hermeneutics – in short, writings or sayings of music.

Words, images, and corporeal gestures stick to it: in memory they become “*part of the music [...] they corrupt the music*” (Abbate, 2004, p. 524). This explains why in her view a gnostic analysis can be undesirable.

I agree with this assessment of musical stickiness – but would instead look to embrace the relationship that is forged between analysis and music as an opportunity for new listenings. There is a fine line between an analysis that would subjugate a music and one that would allow for evolved musical experience from new perspectives. It seems reasonable to suggest that Abbate would agree that no one listening of a music is the same – the drastic nature of music implies the existence of an infinite number of listenings, each unique and inflected by countless parameters. Why, though, would we shy from listenings that may have been reached via written means?

Clearly, this is a complicated question, and one that I do not intend to weigh-in on too heavily. I believe that there is a balance to be achieved between a completely gnostic analysis, one that simply “*aims to expose something imperceptible to the untutored or uninitiated*” (Abbate, 2004, pp. 528-529) ignoring the carnal nature of music; and an analysis that would completely avoid “*tactile monuments in music’s necropolis*” (Abbate, 2004, p. 510) such as recordings and scores and swear by ephemeral experience alone. There is one principle that Abbate proposes with which I fully agree: her suggestion to not take intellectual pleasure from music as a *Work*, but as an *event*.

We must move away from the *Work*-centric perspective of traditional analysis and progress towards Small’s (Small, 1998) concept of *musicking* – considering music not as a *noun*, but as a *verb*. As discussed in the next part of this chapter, I consider musical practice as something that occurs within a network of elements. The listener is emerged within this network, as is the analysis. Recognizing each element as being part of a network and as imparting agency within this network is fundamental. In this perspective, it is important to examine what is *done*, not only the objects that *are*. The performance, the performance as it exists after the fact, the analysis, subsequent performances, the performers, the listeners: these are all objects that affect each other and participate in the *musicking* that occurs.

To conclude this initial highlighting of themes from contemporary musicology, I will turn to an analogy: that of searching for oil versus rambling. The traditional approach can be conceived of as entering an environment and searching for a specific object that we know exists. There is a pre-existing element that is somewhere to be found underneath the surface, and we possess the tools to dig up the ground and find it.

Once we have found the object, our job is done – it is this object that has value, it is this object that we can use to fuel the motors of our poetic readings of a piece. This is an activity that can be useful, and seems objectively, and quantifiably, meaningful – however it is something that can also be detrimental to the environment.

On the other hand, we can enter this environment without the intention of finding anything at all. After all, how could we know if there is anything beneath the surface in the first place? If we were to find nothing, could we resist the temptation to fabricate a false version of the product anyway? No, as ramblers, we navigate this place with the sole intention of existing within it, and what has value is this journey that is taken. We walk around the landmarks, noticing new things each time we do. Sometimes, we may find a new path, and the landmarks may look very different from the different viewpoints we attain. To refer to Kofi Agawu: “*analytical knowledge is not necessarily cumulative. [It can] resist or escape verbal summary. It is a hands-on activity*” (Agawu, 2004, p. 274). Like composition, it is a form of making with infinite variety.

Networks.

This rambling approach to analysis leads us to look for methods that steer away from a Work-centric perspective, and towards a conception of music as event, as verb, as process. When thinking along these lines and beginning to consider how one could attempt to analyse process in a concrete way, it seems useful to not focus on one object, however important this object may seem, but on the relationships between objects. When we take a step back, and following Small’s (Small, 1998) ideas around musicking, we realize that we wish to conceive of the social in the form of networks.

The term *network* is loaded, and as I demonstrate in this section, it has been inflected in several ways over the past fifty years. The concept is at the heart of my methodology, and I develop a particular perspective on the idea that has been shaped by my subject matter.

In the context of musical analysis, a network-oriented approach has traditionally been adopted to approach questions of *genre*, *musical influence* and *listening practices*. For example, Bryan and Wang (Bryan and Wang, 2011) perform network analyses on a pre-existing dataset (whosampled.com) of music which catalogues metadata for songs which sample material from other works. These analyses allow them to apply social network metrics such as Katz centrality to quantify various trends and characteristics, and they even go so far as to draw conclusions on questions concerning influence among artists and genres.

As we shall see, influence is a term that can be assimilated to *agency* and will be at the heart of the network-oriented approach. However, this kind of research tends to work with pre-existing datasets (other examples use big-data datasets derived from platforms such as *Spotify* (Donker, 2019) and *Deezer* (Buffa et al., 2021)) and take a broad perspective around genre and listening practices. I wish to adopt the network-oriented approach at a much more local level and use it as an analytical tool that will go beyond the identification of trends and patterns and help address a music at an aesthetic level. To do this, I will first give a broad overview of the question, before finally proposing my own approach that draws from these theories and seeks to answer some of the issues that I feel they leave open.

I begin with Actor-Network Theory (ANT), a concept that is generally recognized as being developed by French sociologists Bruno Latour (Latour, 1988) and Michel Callon (Callon, 1986) in the 1980s. There aren't many examples of ANT being used in musicology. I draw on Piekut's article (Piekut, 2014), where he gives a good introduction to ANT, and discusses how some of its notions and methods may be useful for musicologists. The context of this article is that of music history: I wish to take these ideas and examine how they might be used in an analytical context².

Piekut starts by summarising the four orders of social mediation offered by Born (Born, 2010): in music, *"we relate to each other as collaborators in the course of a musical performance, in the imagined communities that are animated by these performances, in the identity categories and hierarchies enacted in sonic practices, and in the social modes of its production and distribution"* (Piekut, 2014, p. 191). Essentially, *"music requires collaborators in order to touch the world"* (Piekut, 2014, p. 191). In the perspective of ANT, these collaborators have many forms – not necessarily human. Furthermore, these orders of social mediation, themselves networks, will sit within larger networks, all influencing each other. It becomes clear that an analysis of a social phenomenon conceived of in this way would not only content itself to examine the various elements separately, but also the relationships they have with each other.

Piekut cites a quote from Latour that eloquently summarizes the approach: *"[Society does not exist]. It is a name that has been pasted onto certain sections of certain networks,*

² However, the nature of ANT and the relational stance means that the distinctions between fields like historical musicology and analysis become blurred. We become alert to different ways in which to engage with objects of study.

associations that are so sparse and fragile that they would have escaped attention altogether if everything had not been attributed to them" (Latour, 1988, p. 216). Piekut rightly suggests that the word *society*, here, can easily be replaced with *music*. The perspective is clear, but what are the notions and methods that can be pulled from this in a musicological context? Piekut gives an overview of some of the key concepts with four methodological principles: agency, action, ontology, and performance.

The network is comprised of *actors* – an actor may be human or non-human: ANT makes no judgment of value on actors, there is no pre-existing hierarchy between say, a human and an instrument. This is a heterarchical perspective. Actors are placed on an *"equal ontological footing [...] to avoid absconding, unwittingly, any of them from one's account"* (Piekut, 2014, p. 195). This is contentious, for example he discusses Taruskin's (Taruskin, 2005) point of view on the question who believes that an attribution of agency that would be unmediated by human actors would be a lie. Essentially, agency would require sentience. I place myself in a middleground, believing that each actor will impart agency in some form upon the network, even if this agency can be conceived of in various types and as having various degrees of impact. Agency is *distributed* across the network – the musicking that occurs within a network might not be attributed to the sole agency of one lone actor.

The concept of agency is tightly interwoven with that of action. Piekut tells us that action is always a kind of *translation* (Piekut, 2014, p. 198). Action takes place *"through a chain of translations that disperse, mediate and circulate agency"* (Piekut, 2014, p. 198). Agency, again, is *"not concentrated on a single entity"* (Piekut, 2014, p. 198). In my methodology, it is not the origin of the agency that is important, but the effects on the global network. I am interested in the effect of actions, rather than the causes. Actions will also *produce* actors, a notion I return to in the second part of this chapter.

ANT offers a particular idea about the ontology of an actor within a network regarding its status and the way it imparts agency. This perspective is important, as in ANT a researcher follows how *"different networks of actors constitute or enact different realities; they emphasize specificities and differences rather than universals. [Therefore,] by not deciding ahead of time what we are going to find in the world, we allow entanglements to emerge in all of their messiness"* (Piekut, 2014, p. 199). Essentially, *"being means 'being related' and 'being in the world'"* (Piekut, 2014, p. 200); this has important implications as to how the analyst must proceed.

It suggests that an analysis of performance (in the expanded sense of doing things and as the enactment of musical reality) may be a potential model for analysis in a general sense. There are things to be wary of: for example, Piekut summarises the point raised by Born: “*musicology’s recent interest in performance [...] has led to an over-emphasis on the micro-socialities of musical encounter, at the expense of investigating larger forms of social mediation*” (Piekut, 2014, p. 201). However, I believe that ANT, if taken in a broad way, could provide a good basis for a compromise between an atomic, surgical dissection of micro-socialities and Abbate’s call for a focus on the drastic nature of music, the experience of musical process, of musical enactment.

Piekut concludes by proposing the notion of *musical ecologies*: “*webs of relations, [...] interconnecting and mutually affecting*” (Piekut, 2014, p. 212). They are emergent and hybrid. They have real boundaries, but they are “*variable and open*” (Piekut, 2014, p. 212). When looking to analyse a specific performance, I will draw on this model and propose the idea of the *performance network*, a notion that goes further than the notion proposed by Cook. It is not easy to define the boundaries of this network, and “*distinctions between social, technological or musical domains are difficult to make*” (Piekut, 2014, p. 213), but it is not beyond our grasp. Suchman describes this as “*a practical matter [...] of cutting the network. [...] The relatively arbitrary or principled character of the cut is a matter not of its alignment with some independently existing ontology but of our ability to articulate its basic implications*” (Suchman, 2007, p. 284). In my methodology I approach this question by looking at the processes and functions present at certain parts of the network.

At an abstract level, ANT’s ultimate goal is to “*provide an empirically justified description of historical events*” (Piekut, 2014, p. 193). The key word here is description. This will be one of the limitations that I bring into question with ANT. If we were to limit ourselves to this first, descriptive step, we would miss an opportunity to experience the network under scrutiny in a way described previously with the analogy of rambling. With this map in hand, we could exist in proximity to our object of study and use it to gain access to different perspectives. In this research, I will try to explore the idea of not just empirically describing events with networks, but how we could navigate them in an analytical way.

One major critique that has been made of ANT is its failure to account for time. This is also an important part of the contribution that I wish to propose. However, before

discussing this further, I will go over the ideas of two influential thinkers who have built upon the ANT approach in sociology, and notably the ways they deal with time. First, I will examine some ideas of British anthropologist Tim Ingold who proposes one notable inflection of the network approach: *meshworks*. He demonstrates this with the narrative of an ant (ANT) and a spider (Skilled Practice Involves Developmentally Embodied Responsiveness) discussing social theory (Ingold, 2008).

For him, in ANT, the relation between actors has no material presence, "*materiality of the world [...] is fully comprehended in the entities connected*" (Ingold, 2008, p. 210). Here, he brings into question the idea in ANT that, despite studying the relationships between actors, it is still ultimately the actors which become the object of analytical study. ANT stipulates that action produces actors – the agency of an actor is comprehended *via* the actor that its agency, its action creates. Contrarily, for Ingold, relationships are lines *along* which actors live and conduct perception and action in the world, they have a certain materiality. "*The lines of [the spider's] web do not connect [it] to the fly. [They] set up through their material presence the conditions of entrapment under which such a connection can potentially be established*" (Ingold, 2008, p. 211). In a sense, the connections, the relationships between actors, between entities, are pre-existent. Entities operate within a meshwork that already exists.

He does not imagine a world of entities like ANT which are assembled. For him, "*the web is not an entity*" (Ingold, 2008, p. 211). It is not a "*closed-in, self-contained object*" (Ingold, 2008, p. 211). For Ingold, the world is not "*an assemblage of heterogeneous bits and pieces but a tangle of threads and pathways*" (Ingold, 2008, p. 212), it is a *meshwork*, rather than a network. Therefore, action is not the result of distributed agency, but it "*emerges from the interplay of forces that are conducted along lines of the meshwork*" (Ingold, 2008, p. 212). For things to interact, they must be immersed in "*a kind of force-field set up by the currents of media around them, [without which] they would be dead. [...] The web, [for spider], is the condition for agency, but [...] not itself an agent [...] The essence of action lies not in aforethought [...] but in the close coupling of bodily movement and perception*" (Ingold, 2008, pp. 213-214): action is skilled, and skill is something that is developed. Agency requires skill; therefore, we cannot attribute agency to things that do not grow or develop.

This is not the place to give a detailed critique on Ingold's work. This is a summary of some of his thoughts he develops in other works (Ingold, 2013). These distinctions can seem minor, but they become interesting for us when we consider the question of

time. Where social analysis from the perspective of ANT can seem to forget time and focus on collections of actors in an out-of-time, suspended state, it seems clear that Ingold encourages a focus on *process*, and happenings that pre-exist the arrival of actors which then impart energy across the meshwork. This approach of Ingold's, while useful, is still characteristically a perspective on time that can be characterized as *retrospective* – social phenomena are regarded as happenings that *have happened*. For me, this is still unsatisfactory.

The final modulation I would like to discuss is the thinking of philosopher Peter Sloterdijk (Sloterdijk, 2004), and more precisely Marie-Eve Morin's (Morin, 2009) dissection and interpretation of his theories. Regarding what interests us and the network approach to analysis of the social, we can understand his theory through his metaphor of foam. This is tightly interwoven with his ontological theory of being-in-a-sphere, in *ensouled bubbles* and the trajectory that humankind has undertaken over history towards living in what he calls *global foams* (Morin, 2009, pp. 61-62). The two main points to draw from his metaphor, are that "*foams are loosened structures, multichambered systems whose cells are separated by thin membranes*" (Morin, 2009, p. 67), and that "*foams are processes which tend towards stability and inclusiveness. One recognises a 'young' foam by its smaller, rounder, more mobile, and more autonomous bubbles. With time, each bubble will come to be shaped by the surrounding ones and its interior will stabilise itself. As a consequence of the reciprocal stress exerted by each bubble on the surrounding one, a foam will gain a certain tonicity*" (Morin, 2009, p. 67).

Morin goes on to explain that "*if we apply this metaphor to the social world we can say that 'society' is [...] an aggregate of microspheres. In this aggregate, each bubble is a 'world' [...] each of these worlds is simultaneous and connected to all others, yet at once separated by a transparent and flexible boundary. The result is a system of cofragility and co-isolation*" (Morin, 2009, p. 67). We can take this image as a particular perspective on the network, and as presenting a particular idea of time: everything is interconnected in an unstable structure. I read this as time being equated to decay.

Like with Ingold, time is again considered as unidirectional, and process as inherently interconnected. Everything tends towards a certain decay as a product of all these interconnected bubbles – in these theories of the social there is no space for a vision of time that could examine anything other than what has happened. This is something that I take issue with. Why should social analysis be constrained to operate along this sole vector? I believe that when adopting a network-inspired approach, we have a tool

to regard time in a much more flexible manner – in a manner that could open new avenues for research, and a different understanding of social phenomena.

b) *Creation.*

Ontology of the Record.

I now examine from a broad perspective the nature of the musicking that is occurring in my case studies and demonstrate why the network and rambling approaches I began to announce in the previous section are justified. The traditional, text-centric approach will be difficult to apply to the FluCoMa musickers' practices where there is no strong candidate for a Work-like object upon which to centre an analysis. In this section, I propose a model of *measurement* and *manipulation* of an artifact, the *Record*, which will serve over the course of the thesis as a starting point for lines of questioning and as a tangible object for analysis.

The premise of creation in the context of this research is simple: people engage with things in a *musical* way – this can be considered an offshoot of general artistic practice, where people engage with things in an *aesthetic* way. The musical could be considered as the *sonorous aesthetic*, or attitudes that are loaded with cultural and historical musical baggage. Our tendency in Western philosophy led to our labelling and codifying of what seem to be stable elements of this equation: the three primary musical engagements or attitudes of composition, performance and listening; and fetishised objects such as the score and the musical instrument. These are indeed names that have been pasted onto certain parts of certain networks. Given my ANT-inspired commitment to not entering analysis with strong ontological presuppositions, I need something that can present a flexible object for analysis – what I call the Record.

I propose, then, a perspective of music and creation as a heterarchical network comprised of nodes that interact, where agency is distributed. I have a broad vision of the network, somewhere between ANT, *meshworks* and the musical ecologies that Piekut offers. There are a few elements to unpack. First, I do find it justified to discuss at least two different types of actors in this network: *humans* and *artefacts*. Another is the bounding of a network. This is something that the theories I have presented seem to offer little help with, but is an important question, especially for analysis. I propose to bound a network by considering its process, its function. In each case, I will look to identify the different processes and functions that are occurring, and bound parts of

networks by including actors and relationships which seem to intervene in these processes.

This is something that can be used on various levels: it will allow me to modularise a network (see the next part of this chapter) at a local level; but also, it gives me a broad means to define something like the *performance network* discussed above. The performance network is meaningful in that there is a clear process that is occurring and hence a bounding of it can be justified. By no means is this straightforward, many things could arguably be included within the performance network, but it at least begins to centre focus from the chaotic flux of the entire world to a certain area.

Finally, this perspective allows me to suggest the following hypothesis: musical creation, artistic activity, musicking, can be considered as the creation of, modification of, assemblage of and existence within various networks. There are actors, people that we would traditionally call artists, composers, performers, who actively seek to configure networks of humans and artefacts in a way that adheres to their *musical thought*³. I argue that through these processes of creation, modification, assemblage and existence, musical thought can become inscribed within a network.

I propose that musickers engage with different artefacts in different manners, and that their musical thought tends to be concentrated around certain artefacts, rather than others. When musicking, musickers direct their agency to the measurement and manipulation of a certain type of artefact – the Record. In this context, the Record takes on a larger signification than that of an analogue or digital recording or of a transcription of events. Here, the Record is considered as a conceptual artefact, which contains or symbolises musical knowledge.

A simple example of this is a digital recording: at a material level, it is a collection of 0s and 1s that can be interpreted by a machine and cause a set of speakers to play a sound. However, what is to be understood in the Record, is the *idea* of this digital recording before it is played. It is a symbolic object – in this example with a physical form, but this is not always the case – that contains the *idea* of a certain sound, of a sonorous object. The human actor conceives of this Record in their mind and has an idea of the musical knowledge that it symbolises for them. The human actor knows that they can *do* certain things with this artefact, they will engage with this artefact in a musical way. To demonstrate the breadth of this notion, other examples may be a

³ A brief definition of this term was given in the introduction. It shall be fully explored in Section 2.b. of this chapter.

musical note on a page indicating a certain sound, a certain timbre, a certain length, a certain pitch; a musical instrument, indicating again a certain realm of sonorous potentialities; another musician, who is conceived of in the mind of the musicker as making a certain type of sound in a certain way; a concert venue, which will necessarily offer and limit certain sonorous possibilities.

The Record is a conceptual, *hypothetical* artefact. It is something that exists within the network but does not necessarily possess a physical body. The Record is the *idea* of a sound, it is what one thinks of when one writes a piano note on a staff. The Record is what we *expect* a sound to be, it is what we think of when we see a sound file sitting on a computer. The hypothetical sound, the conceptual fragment of musical knowledge, is brought to life through musicking, through its interaction with actors in a network. It is something that is produced, actions produce actors. It is something that has some form of agency, it is something that can produce other artefacts, and other Records. It is a pivotal element for understanding musicking. Observing the ways in which musickers interact with it will surely give us precious insight on what is happening within a network.

There are two main ways in which musickers engage with the Record: measurement and manipulation. Over the next two sections, I examine these activities in more detail, before finally taking this theoretical context and laying out the research questions that articulate my research.

Measurement.

To discuss the idea of measurement and manipulation of the Record, it is first useful to discuss along which vectors these engagements occur. What are the units that we use to understand musical knowledge? This is something that is closely tied to technology and the tools we have at our disposal to conceive of sound. The first two dimensions, that can be conceived of by the analogy of the spectrogram, are *time* and *energy*. We can instinctively interpret content across these dimensions to discern things like duration, timbre, and pitch.

There is, however, a third, more elusive dimension which adds itself to these two when we start thinking about how we organise sound in a musical way, when deploying musical thought. This dimension is somewhat cultural in nature, and less empirical. It is the dimension from which concepts such as harmony, melody and rhythm emerge. It is what we do with sound, how we put sounds into the context of themselves and others. I call it the *organisational* dimension.

Let us draw a broad historical image of our relationship with the Record, and how it is tied to our *conceptual* access to these three dimensions via technology based on the epochal account provided by Firth (Frith, 1996). This is a notably Western-centric account: in reality technologies and epochs are constantly mutating and porous (if this were not the case, nothing would ever change). This is an abstract tool for understanding the Record and is by no means an exhaustive account of global music history.

We can first consider an *archaic* epoch, where the primary technology is aural/oral transmission – it is effectively difficult to assess this period given the notable absence of historical traces.

Then we can consider the *classical* period, tightly linked to the technology of symbolic musical notation. This period is characterised by a focus on the organisational dimension suggested above. Musical notation, the primary technology that constitutes this epoch, inherently relies on the musical note either representing a discrete sonorous object, or not referring to a specific object at all. With the technology available at the time, people's conception of sound was inextricably linked to the sonorous body. The spectrum of conceived possible timbres to work with was finite. As these timbres are given, musical thought tends to be contained within the realm of the configuration across the organisational dimension of these pre-existing elements. This conception of music, and the technology of notation, professed to the notated score the status of Work – hence we see emerge the despotic figure of the composer. Music *is* the Work, the score, the artistic vision of the composer.

Finally, our *contemporary* period is also constituted by a musical technology: that of sound recording. The role that technology plays here is fundamental: we see that, beyond the technological possibilities it offers, it changes our *perspective* on the reality of what sound and musical knowledge is – even if this reality hasn't essentially changed. Here, the radical change we observe is the shifting from sound as a finite pool of discrete timbres, to sound as an energy that can be controlled, transformed, modulated, fragmented at any point, and can notionally be heard again and again.

It is a discursive shift, rather than an actual one. The artefact that analogue and later digital recording techniques produces, is one that allows the human consciousness to conceive of sound in a completely different way, it opens new avenues and vectors for musical thought. It is no surprise that a large part of contemporary practice involves instrument-making or sound design of some kind. The instrument is no longer this

inextricable discrete object, but rather an element that will perform manipulations on sound in the energetic and temporal dimensions. We can even go so far as to propose the hypothesis that conceptual access to these afore-unattainable parameters may transform the way we think across the organisational dimension⁴.

Now let us regard this through the prism of the first of two major activities of engagement with the Record: measurement. As I have proposed, it is possible to stipulate that conceptual access to the energetic and temporal dimensions of sound was limited to musickers of the classical period, hence a tendency to focus on the organisational dimension. Much musical production across this period attests to this: the primary vectors are pitch, harmony, and rhythm. This extends to a point where the Work, the collectively conceived musicking, lies within the score. This is one of the reasons why traditional musicology focuses on this object, and why performance was once placed below the score in the hierarchy of the Western zeitgeist, as fatally imperfect executions of a piece.

Measurement, or *the conception of a Record in some way*, was somewhat simpler than today (see below). Measurement of timbre is not always a primary vector for the classical composer – the finite pool of orchestral sounds does this work for them. From a compositional perspective, measurement will occur much more strongly in the conception of Records across the organisational dimension.

Conceptual measurement in the contemporary epoch is more complex. I have already mentioned the spectrogram and sound recording which give us literal and conceptual access to conception across the temporal and energetic dimensions. We can also add to this the tens of standardised audio descriptors and the infinitely variable types of statistics we can draw from them. Fast Fourier Transform (FFT) processing lets us construct and measure sound at extremely high resolutions; a sound can be split into atomic grains of mere milliseconds.

The dramatic branching out and explosion of different musical practices in our time attest to the number of avenues that these new perspectives – granted to us by sound recording technologies – offer. Not only has this technology opened two new dimensions across which to perceive and measure sound, but the *resolution* of these

⁴ Indeed, one may argue that conceptual access to different dimensions will influence the way one conceives of sound across other dimensions. For example, the musical thought and the organisational configuration of Harvey's *Mortuos Plango, Vivos Voco* is explicitly derived from energetic measurements of the sound of a bell – it is reasonable to propose that this kind of approach will remain in Harvey's approach to musicking in general, even in the context of pieces that do not directly use the same kinds of techniques.

measurements becomes ever finer. Measurement, our capacity, and tendencies to conceive of sound, can go from simply thinking about the amplitude of a sound or the duration of a motif, to the detailed gathering of thousands of statistics of a mere grain. I hope to have demonstrated that technology has a large impact on this.

Manipulation.

Finally, let us look at the second activity around the Record: manipulation. Again, modern technologies have opened new avenues for contemporary musickers. Whereas before, manipulation of the Record consisted essentially of assembling various Records together, in the contemporary epoch we have the ability to conceptually go into the Record itself and modulate its essence. As mentioned above, this is a reason why much of contemporary musicking occurs in instrument-making and sound design – a role that has historically been conceived of as craft rather than art.

As I outlined above, and as we can see in many of the practices of the FluCoMa musickers, a large part of contemporary practice occurs in the configuration and assemblage of networks, and subsequent existing within these networks during rehearsal and performance. The configuration of these networks is built around the various Records the musicker wishes to manipulate. As I explore fully in Chapter 7, with the Record, I understand not only basic musical units such as a frame of sound, but also broader artefacts such as the hypothetical conception of a collection of sounds or an instrument.

Manipulation, *the conception of what one can do with a Record*, can take various forms: when regarding digital sound, there are a whole host of ways of manipulating the record across the energetic and temporal domains. Temporal manipulations can be seen with processes like time stretching, repetition through echoes and delays, reverberation, splitting audio and playing it back in different orders or backwards, spectral freezing, granular and concatenative synthesis. Energetic manipulations can be seen with processes like distortion, EQing, filtering, and bitcrushing. These categories are porous, but we get an idea of the kinds of tools across these two dimensions that contemporary musickers can conceive of.

It is interesting to also consider manipulations across the organisational dimension, a notion which is much less transparent. Candidates for this would be generative composition systems or works like Jonathan Harvey's *Mortuos Plango Vivos Voco* which draws its musical thought across the organisational dimension directly from

the measurement and manipulation of a single Record. There are also examples of this in more classical works, for example Mozart's *Musikalisches Würfelspiel*.

We can also observe manipulations of Records of other, broader natures. For example, with the perspective offered to musickers through sound recording, Records such as the score and even the Work appear to have become candidates for manipulation. We can think of works such as Brian Eno's *Fullness of Wind*⁵ which takes Johann Pachelbel's Work *Canon in D Major* and operates a large transformation across the temporal dimension.

Before concluding this theoretical context and outlining the research questions I wish to address, it is necessary to make clear that I do not wish to fashion this into a universal, axiomatic model that would allow us to encompass and understand all types of musical practice. It is a theoretical bedrock upon which I build my research, as I believe it will lead me to *pose articulate questions*. It is a perspective on music and creation that allows us to adopt some of the analytical principles outlined in the first section, and more importantly, allows the analysis to be driven by the practice it wishes to observe.

2. Research Questions.

a) *The question of time.*

Having made clear the theoretical context upon which this research builds itself, giving a broad demonstration of my approach to analysis and my conception of musicking, here I present the specific research questions I wish to interrogate. These are a set of questions that try to address some of the more elusive parts of the theories and hypotheses already discussed and are also somewhat methodological in nature.

The first question is around the question of time in the network approach, specifically, how can we examine the nature of a network beyond one temporal occurrence? How can we gain insight on the nature of the possibilities, of the *potentialities* of the network? This stems from an inherent dissatisfaction with the way time is dealt with in the theories discussed above. To recap, ANT's approach has been critiqued as not accounting for time in a satisfactory manner. Other approaches, such as Ingold's

⁵ ENO, Brian. *Discreet Music*. Obscure, 1975.

meshworks and Sloterdijk's *foam* do factor in for this element, but their perspectives keep a unidirectional, retrospective vision of time.

Considering that I wish to approach analysis in the image of rambling, journeying in proximity of and inserting oneself into the musicking network, it will be useful to gain insight not only on what *has happened*, but also what *could have happened*. This would mean that we have a much better understanding of the *nature* of the network, rather than just the nature of one of its occurrences at a given time. This transports the idea of the network as a static account of something that has happened, towards a dynamic conceptualisation of what *could* occur. The network becomes a living site of potentiality, something that sets the conditions for something, or somethings to occur.

This reflects the idea that musicking does not occur within the confines of a lone object, but that indeed, artistic practice is the configuration of networks. Artistic practice is not the delivery of an object, of a Work, but the modulation and configuration of networks, and existence within them so that something may happen. For analysis, this means ceasing to consider networks as timeless, or bound to a single timeline: they are quantum objects with an infinite breadth of existential possibilities. The question is how can we approach this concept in an analytical way? Over the course of the research, I develop several methods that would try and cater to this perspective – some with more success than others.

I propose that time can be viewed as multidimensional, but what does this mean for analysis? What is it that the analyst is actually looking at when attempting to approach a network? First, we need to know the bounds of the network we are looking at. As was discussed above, the way I propose to bound a network is by looking at its function, its process. The focus of analysis becomes this process, but how does one conceive of a process in a tangible way? There are two possibilities: a process can be seen as the changing, creation or destruction of actors within a network – in this case, we consider the process by its result, focusing on the actor; or we can focus on the interaction itself, considering the process through its action, its energy. We only know that a process has occurred if there is a result, yet this result is clearly not the process. Should we consider the process as the hypothetical action, or the trace, the result?

The answer probably lies somewhere in the middle. Ultimately, we are trying to read the shape of a network, its aspect through time, its aspects through times. To do this, we need to understand at once the nature of its structure, and the aspect of the energy it generates. This is reflected in the two parts of the methodology that are discussed

further into the chapter: the *cartographical* and *cartological* analyses. This research asks to what extent it is possible to approach analysis from this perspective by proposing a concrete methodology that is built around this concept of time.

b) *Musical thought: inscription and materialisation.*

Can we conceive of musical thought as being inscribed within the network, and if this is the case, is it possible to access a materialised form of the musical thought through analysis? Musical thought is that which deals with the measurement or manipulation of sound, of Records across one of the three afore-given dimensions: energetic, temporal, and organisational. Musical thought can therefore also be seen as a series of actions, intentions, tasks, and functions, giving a molecular, modular nature to musicking.

To understand this term more precisely we can look to Stiegler's (Stiegler, 2010) idea of *grammatisation*: "*by grammatisation, I mean the process whereby the currents and continuities shaping our lives become discrete elements. The history of human memory is the history of this process. Writing, as the breaking into discrete elements of the flux of speech [...] is an example of a stage in the process of grammatisation*" (Stiegler, 2010, p. 70). This idea is explored in a musical context by Magnusson (Magnusson, 2019). He discusses Tomlinson's (Tomlinson, 2015) work around ideas of development of early instruments, notably a stone flute: "*conscious planning of where the holes should be made in the flute is what Bernard Stiegler calls grammatization, or the process of externalizing our thoughts using discrete systems [...] The location of these holes [...] is abstract, non-semantic, and of no particular function*" (Magnusson, 2019, p. 3).

This is an example of the materialisation of musical thought: it is inscribed within an object in a tangible way. Like the positioning of the holes on an ancient flute give us insight on the maker's musical thought regarding pitch and possibly modes, the physical materiality of a violin, and the way its strings are disposed and tuned, gives us insight on the musical thought of its maker(s). This instrument can be seen as emerging from the development of the tonal system: the difference of fifths between the strings can attest to this. The same idea can be proposed in all the classical orchestra's instrumental components: for example, the keys of a piano being divided into sets of repeating 12 keys, visually built around the *c* and its fifth *g*. It is constructed in a way that allows for execution of musics that follow a certain type of musical thought.

Musical thought is embedded within the tools that are created to allow for this musical thought to manifest itself in musicking. Therefore, an *organological* approach to analysis seems to be a good candidate for accessing an understanding of this musical thought. This is where the network approach, again, becomes useful. What is an instrument if not a network of human and non-human actors which are assembled to execute a specific function, a specific process? A network which is specifically constituted to allow for musicking to occur according to its creators musical thought. This is the vision of the instrument I propose in the context of this research.

This leads to questions about my analytical approach: should I examine the nature of this network through its constitutive structure, or through its results, through its use? I encountered these questions in the previous section, and again, the answer surely lies between the two: a need to understand not only the static structure and the potentialities it allows, but also how it is used in practice.

Take for example, the trumpet of Miles Davis, and that of a trumpeter in a marching band. Clearly, this object is being used to execute two very different types of musics, two different streams of musical thought. In this case, the trumpet is to be viewed as an actor in a larger instrumental/performance network, and an examination of the way the object is used plays an important role in understanding the musical thought of its creators.

In the context of this research, this question is not so divisive, as we are facing a set of practices where the musickers are all instrument-makers. Every one of them participates in instrument making at some level: from the assemblage of traditional instruments into hybrid instruments, to the fabrication of specific objects that did not exist before. Here we do not need to trouble ourselves so much with the messy untangling of the musical thought of the instrument-maker and that of the performer: they are one and the same.

c) *Network superposition and interfacing.*

In what ways do musickers configure these networks? I deal with this question by looking at the physicalities of a network, and notably the physicalities of points where agency and energy is translated through *points of interface*. This leads us to ask, where can we find these points of interface and how can we examine their nature? What are the physicalities of a point of interface? Would a typology of points of interface be possible? What are the translations that can occur through a point of interface?

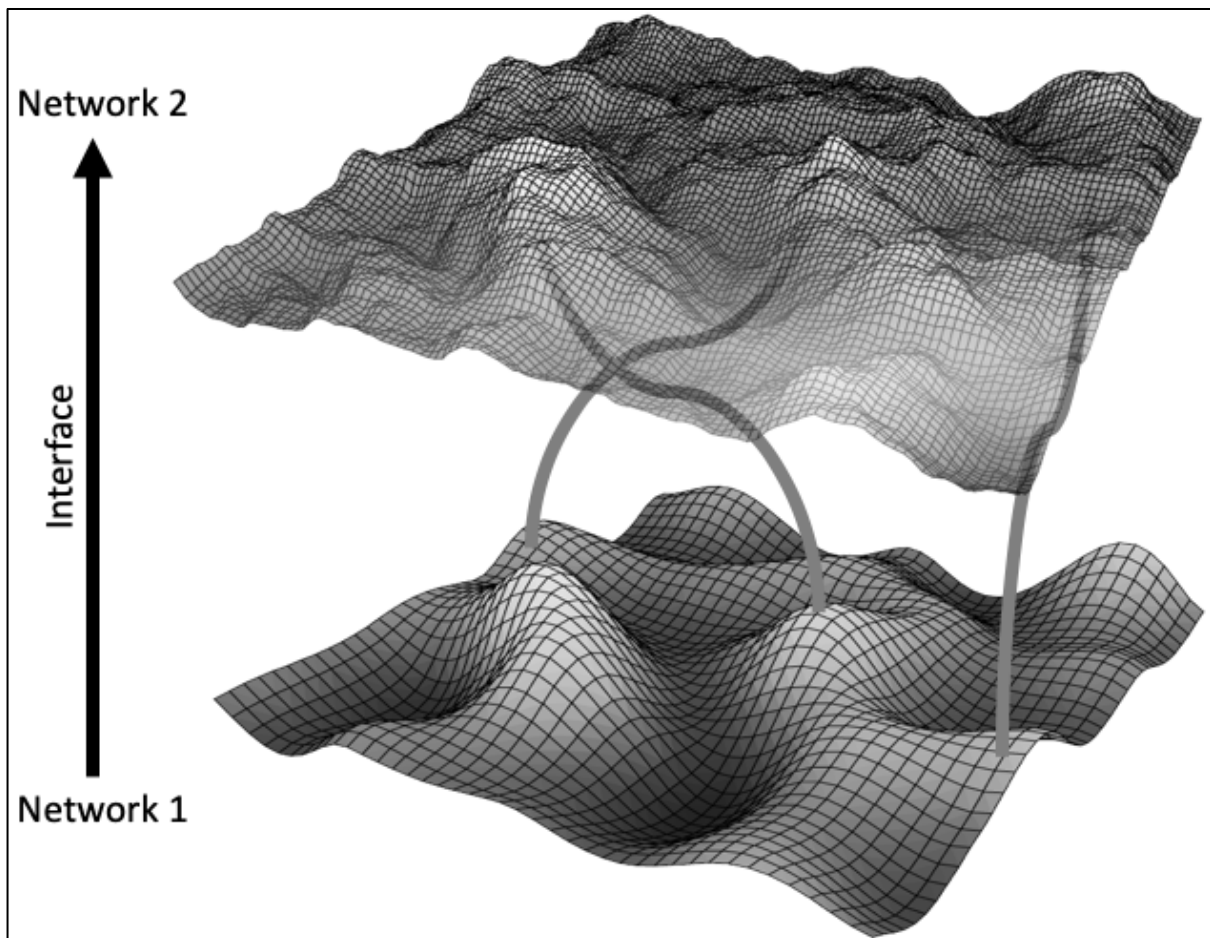


Figure 1: *visualisation of the network superposition analogy.*

To better understand the idea of points of interface, I use an analogy: one can imagine the networks configured as existing within an infinite three-dimensional space. The network can be seen as a plane, a large piece of fabric with a structure that gives it a topographical aspect. When inactive, this sheet lies motionless in space, however, we can pick it up at the corners and impart energy across it, as if spreading a sheet out over a bed. The plane becomes rugged, rough, and constantly moving like the surface of a disturbed body of water. Over time, the relief changes, concentrating around

peaks and troughs before moving to other ones. This is action, this is process. The internal structure of the sheet will drive it to gain certain aspects over time when energy is imparted across it; it is an unstable, unpredictable process. The plane is essentially endless, meaning that for analysis, we can only concentrate on certain parts of it. However, as Stiegler and Magnusson suggest, we can observe patterns and repetition which would suggest a grammatisation of this flux.

How do musickers go about grammatising these infinite planes? What are the tools that they have at their disposal? Following an Ingoldeque approach, we can visualise *meshworks* of matter pre-existing an actor, and propose the following idea: one can structure a network through interfacing of networks between each other. Galloway (Galloway, 2012) talks of points of interface as being “*autonomous zones of activity*” (Galloway, 2012, p. vii). To join with my own perspective, we can see interfacing as points where energy is *translated* from one network to another, from one Record to another. Points where matter is interacted with and translated through measurement and manipulation.

It is where different types of networks meet, we could talk of *network superposition*. To continue my analogy, imagine in this space, two or more sheets being placed upon each other, being stitched together at certain points. The movement of one network will influence the other, the two networks collide and create relief. Energy is translated from one to another, its effects rippling through each. This rippling and constant collision between networks best describes the processes within networks over time and the infinite complex of potential existences these configurations can have. This analogy can extend to demonstrate the complexity of the social fact: multitudes of different networks stacking on top of one another, planes looping round to superimpose themselves upon themselves and networks interweaving.

All of this can be daunting when beginning to imagine an analytical approach that would usefully tackle these complex webs of energy. It would seem useful at this point to move away from this initial theoretical discussion and examine how this approach stands in practice. I wish to embed this research in concrete, tangible analysis through a series of case studies. I add, then, to my previous research questions, the following hypothesis: musical thought may be sought after, as described in the previous section, in the structure and aspect of a musical network over time; and notably at points of interface, where energy is translated from one type of Record to another, where

measurement and manipulation occurs. Let us now begin to forge and test a methodology that would intend to answer some of these questions.

3. Methodology.

Taking this theoretical context and research questions, the question now becomes how can we approach a set of case studies in a methodological manner? The nature of the networks I wish to examine, and the broad stance I adopt regarding analysis suggest that a wide variety of techniques will need to be deployed. I will borrow methods from several fields and modify them to fit my uses. I have also developed several digital tools that have aided my research.

a) Source collection and archiving.

Before discussing the main methodology that structures this thesis, I shall discuss the activities around source collection and archiving that I was led to perform. Tracking the creative process suggests a certain level of archiving, and even if this was not a primary objective of my research, it was a necessary step.

First, I conducted several interviews with each of the musickers at two key points of the project: essentially, the beginning and the end. These interviews were semi-structured, a set of general questions was outlined for the group, and it was also necessary to forge questions specific to each of the musickers' practices. During the interviews, I left a lot of space for the musickers to lead the conversation, allowing them to discuss what they believed was important. This was to gain a better understanding of their own grammatisation of their networks.

I also collected and analysed the various presentations the musickers gave over the course of the project – notably their initial artist presentation and the presentations of the work for the project. This was very useful in gaining a better understanding of the way they conceive of their practices and how they see themselves configuring musical networks. Many of the musickers are also academics: this allowed me to collect, read and comment several different writings they have published concerning their work – some concerning the work they did for the project directly.

I also collected their work for archival and analysis purposes: this was informed by each individual practice, but mainly involved collection of code and any external or third-party libraries required to make it work. Some of the musickers kept an iterative record of the software as it developed; some provided some initial experiments that

didn't necessarily end up in their final projects. Where possible, I also collected audio and video recordings of rehearsals and tests.

On my request, some of the musickers modified their code slightly for the performance in order to record various audio streams or other elements such as the contents of *coll* objects that would later help my analysis. Finally, during the first concert where the musickers were present, I filmed the setup process, allowing me to interview the musickers on the fly as various parts of their setups were put in place.

This left me with a vast amount of data. It has been classified and stored for the FluCoMa project on a cloud-based storage interface. The elements that concern my research directly can be found in the accompanying digital elements, or in the Appendix of this thesis.

This is the formal side of my methodology that doesn't essential constitute analytical activity. I found myself with a wide array of different sources of varying natures: such is the nature of an analysis which looks to examine networks of actors of varying ontologies. Next, I explain the broad structure of my main analytical methodology. This methodology is one of the research outcomes of this project and it has been developed over iterative engagement with the various case studies. This iterative process becomes apparent across the thesis. A reconfiguration of the methods with each case study developed them as research progressed and helped to keep them grounded in real instances of analysis.

b) Cartographical Analysis.

Given my theoretical context, I propose an analytical approach that can be viewed through the metaphor of map-making. I evoked the idea of rambling within the musical network to gain new perspectives on the processes which emerge from it to gain access to new listenings. This suggests that analysis should try and provide an understanding of the structure of the network – this can be understood as trying to construct a map of this environment, this musical ecology. The musical network is not one that is necessarily materially spaced – although this can constitute parts of the structure of the network. The map, here, is seen as a visualisation that reveals the aspect of a network's structure, its centres of activity, distributions of agency, relationships and lead us to pose questions about the ontology of the actors of which it is comprised.

Chen and Floridi describe visualisation today as “a form of ‘computer-aided seeing’” (Chen and Floridi, 2013, p. 3421). They offer a typical *pipeline* that is followed from a source to destination, which passes through several steps: enriching and filtering, visual mapping, rendering, displaying, optical transmission, viewing, perception and cognition. Most of these steps demand human interaction, and we will adopt different perspectives regarding our source with each. Following on from ideas of agency and translation of agency discussed above, note that each of these steps will disfigure the original source in some manner. This means that every step must be approached with careful consideration and must be sensitive to how it relates to the original source.

A first essential step described by Chen and Floridi, especially in a digital context, is the representation of the source as *data*, a concept that is intertwined with *dimensionality*. There are many ways that we can do this: typically, in music we have represented the source with the somewhat low-resolution data structure of timing, pitch and indications of timbre which is later visualised as the score; more recently sound is transformed into data by dividing it into frames, deriving the amplitude and phase components and later visualising this as a spectrogram (see tools such as Queen Mary University of London’s *Sonic Visualiser*⁶ or IRCAM’s *AudioSculpt*⁷). Indeed, the way in which music is translated into data is typically closely intertwined with the way in which this data will be visualised.

Over the course of this thesis, I will examine the relationship between the derivation of multi-dimensional data in music and the dynamic ways in which it may be represented, visualised. This is something that has been examined in a musicological context before: for example, Couprie’s work and the development of his *iAnalyse*⁸ software. This tool is made for musicologists and allows them to construct graphical linear representations of musical scores. Another example is the IRCAM’s *CataRT*⁹ (Schwarz et al., 2006) software, which allows the user to visualise fragments of audio in a two-dimensional space – the axes of which can correspond to any number of scales from file length to audio descriptor data. I shall propose my own techniques for the mapping of sound that draws on these works.

⁶ <https://www.sonicvisualiser.org/>

⁷ <http://anasynth.ircam.fr/home/english/software/audiosculpt>

⁸ http://logiciels.pierrecouprie.fr/?page_id=1794

⁹ <https://ismm.ircam.fr/catart/>

I was initially inspired by Hanninen's (Hanninen, 2012) *A theory of music analysis* – notably by the cover which depicts a topographical map of Kings Canyon National Park in the Sierra Nevada. This naively conjured-up images in my head of an analytical model that would visualise music as a space and give me a framework for drawing maps that would represent the structure of a piece and allow me to observe the musical process as a landscape. This wasn't quite the case in practice, however there are some notions from Hanninen's book that put me on a path for developing my own analytical framework.

Bernstein (Bernstein, 2013) offers an overlook of her ideas, starting with the notion of different *domains* of musical experience: the *sonic* domain, corresponding to relatively large differences in some musical dimension (Bernstein, 2013, p. 2); the *contextual* domain, defined by its associative orientation, allowing to associate segments due to shared features (Bernstein, 2013, pp. 2-3); and the *structural* domain, which can be associated with approaches like Schenkerian analysis or twelve-tone theory, where analysis defines segmentation from its own theoretical framework (Bernstein, 2013, p. 3).

This idea of musical experience is something that will help guide the ways in which I segment, and grammatise not just music, but the networks from which it emerges. I will also discuss the idea of how different human actors can operate different stances and experience and navigate networks in differing ways: for example, there will be a difference between a performer who is entwined within the network, the audience member who has a much more fluid engagement with the network, and the analyst whose position begins to encroach on the omnipotent.

She also offers methods for segmentation: *genosegments* where a segment is defined by a parameter that may be perceived by the listener or not; *phenosegments* which are more "*readily-perceived*" (Bernstein, 2013, p. 3). This is interrogated notably in Chapters 4 and 5 when looking at sound plots and automated versus supervised segmentation of music. Finally, the most promising aspect was that of *associative sets* and *landscapes*: an associative set is a group of segments bound by a certain number of criteria (Bernstein, 2013, p. 3); an associative landscape is the idea of an ecological landscape populated with species and organisms (Bernstein, 2013, pp. 3-4). "*One is invited [...] to admire the interactions of its populations*" (Bernstein, 2013, p. 4) – and, notably, the linear aspect of time is something that is departed from.

Much of my methodology bases itself on these kinds of ideas: in Chapter 5 the idea of associative sets is morphed into that of *clusters* in sound plots, and the sound plot can be conceived of as a type of associative landscape. However, with this methodology I wish to go beyond Hanninen's methods in two ways: first, I wish to draw literal map-like visualisations of the networks under inspection and interrogate how engagement with these visualisations can inform analysis; second, I wish not only to engage with the music solely as its sonorous manifestation at one point, but as a network that allows for sound to emerge.

My methodology, then, is conceived of in two parts that in practice constantly inform each other. The first part is a *cartographical* analysis: here, I wish to construct mappings of networks and sound that are descriptive in nature; the second is a *cartological* analysis: the reading of these maps, navigation and engagement with them, an interpretative analysis that would look to understand to aspect of their structure and the ontologies of its populations.

The cartographical analysis will start in Chapter 2 with a surface level analysis of three key elements: the performance, the material setup on stage, and the software. This has two main goals: the first is to gain a good understanding of the general aspect of the networks and what it is that is under scrutiny; the second is to gain a general understanding of how the networks are initially perceived, allowing us to keep the more detailed, atomic analyses grounded in reality. With these initial surface level analyses, I also propose strands of musical thought which appear to be present, and then examine with the rest of my analysis to what extent these are correct¹⁰.

For the performances, this analysis comes in the form of a traditional abstract segmentation, attempting to identify local level events and an overarching structure. This also allows me to articulate my written analysis. For the instrumental setup, I will draw a broad picture of the various material elements that are present and interacting with each other. For the software, according to the different case studies, it was necessary to pass through an analysis and refactorization of the code at vary degrees of detail. I will present an abstract conceptual image of what is happening inside the code and how it inserts itself among the material elements of the network.

¹⁰ Indeed, we can draw an idea of what an artist was trying to achieve or have an instinctual understanding of their musical thought, however on closer inspection, it can be interesting to discover if this was really the case and interrogate why strands that may eventually seem absent come through as apparent on an initial inspection of the network.

In Chapter 3, I begin construction of the network visualisations and their grammatisation. This passes through three levels of abstraction: the first at a global level similar to the type of analysis given by Piekut (Piekut, 2014), examining the various institutional actors in the wide FluCoMa network; then I draw a local level visualisation of the two performance networks around the two FluCoMa concerts, giving a broad account of the various materialities in the networks; and finally I look at the instrumental level, which examines closer the instrumental setup, and notably conceives of the code in network form.

Through these analyses, I examine the natures of the various types of agency that are circulating, and comment on how each level may interact with each other. It will also be an opportunity to observe how some of the networks developed over time. This chapter begins by explaining how some of the tools for building these visualisations were created, and the use of the program Cytoscape¹¹ (Shannon et al., 2003) in conjunction with Max and an HTML interface. I also propose several ways with which these networks can be engaged through layout algorithms and filtering, borrowing techniques from network theory. I will demonstrate that grammatisation is an inherent analytical technique that must be deployed *during* network construction and attempt to comment on the various grammatisations as analysis progresses.

Finally, in Chapter 4, I present work around the visualisation of sonorous potentialities of a network. I wish to access a picture that would indicate the scope of sonorous potentialities to conceive of the network not in a static manner, but as a system which allows for things to emerge. This comes in the form of the creation of sound plots following a variety of techniques. As with the previous chapter, I begin by explaining the development of the various tools I modified and built to accomplish this analysis.

Then, I will show with four case studies how my methods developed: first with an *iterative approach*, looking to iterate through all the parametric states of an instrument. Learning from this, I then cover three different approaches that take the idea of *alternative performances* as their method of source collection: first by taking audio from rehearsals and combining it with audio from performance, then two different approaches to reconfiguring and recreating performances in an analytical

¹¹ <https://cytoscape.org/>

environment with the possibility of changing points of the network which appear to be critical.

This will give me a number of network visualisations, some of them *spaced*, others not, which will give me a comprehensive descriptive account of the networks presented by my case studies.

c) *Cartological Analysis.*

The next *cartological* analyses will take the network visualisations created in the first part of analysis and read them, draw signification from them, and ultimately draw from them inscribed fragments of musical thought. Chapter 5 takes the sound plots and inspects the possibilities for navigation of this kind of artefact. I draw from three analogous practices – cartography, rambling and video games – to propose several principles for exploration of a 2 or 3D space.

I then take these principles and put them into practice against sound plots of various kinds: first, static sound plots where the dimension of time is notably unidimensional or untimed; then I examine superimposed sound plots where multiple temporalities are visualised in the same space; finally, I examine the idea of dynamic sound plots, where visualisation of the space evolves over time.

In Chapter 6 I examine the idea of points of interface. This involves using the network graphs constructed in Chapter 3 to identify points in the network that seem to be *centred*, or where notable translation of agency from one domain to another occurs. I begin by proposing a typology for categorising and discussing points of interface in an articulate manner: this draws on some of the classifications made in Chapter 3 of various nodes in the network, leaning on the physical composition of nodes, their singular or composite conception, the types of and way they translate data.

I also propose and explore the idea that when agency passes through a point of interface, there is necessarily a phenomenon of *loss of resolution*. Then, I perform an analysis of notable points of interface in the networks, leaning on two articulating questions: the differences between points of interface that are *aesthetic* and *functional* in nature; and the notion of hybrid points of interface, where the network will reconfigure over time.

Finally, in Chapter 7, I conclude my analysis by examining the idea of the Record and musical thought. Drawing on elements from all the previous chapters, I interrogate the model of the Record and the musical thought that can be inscribed around it. I do

this by examining examples of this configuration in the case studies through Records at three levels of abstraction: first, where the Record can be considered as a fragment of audio; second, where the Record can be considered as a larger scale musical event; and finally, I see how far this model can be extended and interrogate the borders between Record in its context of musical thought and broader aesthetic object. I will systematically examine the activities of measurement and manipulation which are engaged with regarding these objects, discuss their physicalities, and propose ways in which musical thought can be conceived of as being inscribed within these parts of the network.

This methodology developed over the course of the research. There are many ways it could be extended further, and it also presents several limitations: these are discussed in the Conclusion. This methodology is not to be taken as a framework that would need to be applied in a step-by-step manner. There are parts of it that can be taken and deployed independently. However, I believe that the configuration I have presented here does offer an interesting set of methods to try and answer the research questions which drive this research.

PART ONE: CARTOGRAPHICAL ANALYSES.

Chapter II. SURFACE LEVEL ANALYSES.

1. Surface Level Analyses of the Performances.

These initial analyses can be viewed as traditional segmentations. Segmentation is a concept that has always been tightly linked to musicological practice: indeed, Lefkowitz and Taavola explain that “*the very word ‘analysis’ means the division of the whole into its constituent parts, segmentation is intrinsic – implicitly or explicitly – to many analytical endeavours*” (Lefkowitz and Taavola, 2000, p.171). In this same article, Lefkowitz and Taavola explain that formal *theories of segmentation* “*arose in response to the highly variegated textures of twentieth-century music*” (Lefkowitz and Taavola, 2000, p.171), citing works by Tenney and Polansky (Tenney and Polansky, 1980), and Hasty (Hasty, 1981). They base some of their ‘algorithmic’ approach to segmentation on these works.

Much of their approach consists of identifying elements in various domains: pitch, timbre, rhythm, and articulation (Lefkowitz and Taavola, 2000, p.176). This is an aspect that is generally shared in approaches to segmentation, for example Hanninen who I have already cited¹ and her *associated sets* and *landscapes*, her *genosegments* and *phenosegments* which are all groupings of similar elements which give grounds for their conceptual segmentation from the whole.

In our contemporary age of musicology, segmentation has become an idea that must be deployed with great care. Indeed, as explored in Chapter 1, we have recognised the problematic concepts behind techniques such as Schenkerian analysis which hold segmentation at their hearts; furthermore, the turn towards the consideration of music as a process is an idea that steers us away from a segmentation approach. As Laske explains, this can be traced to Nattiez who “*thinks that musicology [...] cannot be restricted to the description of immanent structures*” (Laske, 1977, p. 220). Indeed, as explained in Chapter 1, this is an idea that this research bases itself upon; however, this does not mean that segmentation can be entirely evacuated from the equation. The reader shall remark that segmentation is a necessary step that will appear in various forms across the thesis; however, it is never considered as an end unto itself.

¹ See also her 2001 article which offers a *general theory of segmentation for music analysis* (Hanninen, 2001).

In this first chapter, a broad segmentation is performed to try and reveal some immediate strands of musical thought that appear to be present on a first encounter with the networks in question. The segmentations are subjective in nature, and completely open for debate: my goal not being to reveal any particular truth but to try and gain an intelligent abstract view of each performance that will allow me to talk about the pieces in an articulate manner. I will go through each of the performances by cohort of musickers, and in the order they appeared on the night of the gigs.

There being nine case studies, I will be as concise as possible, drawing on a series of diagrams, give a summary of what is seen on stage, and then an inexhaustive list of potential strands of musical thought that may be present. Full descriptions of the segmentations can be found in Appendix 2. The reader may also consult the digital tool to explore the segmentations interactively². There is also a video explaining how to operate this standalone³, and a video showing how I made these interactive tools⁴.

a) *Olivier Pasquet – Herbig-Haro.*

Pasquet's piece is titled *Herbig-Haro*⁵. In contrast with some of the other pieces, Pasquet's piece is much less performative in nature: the room is plunged into darkness, and he is sat behind a laptop on stage with no lights on him whatsoever. During setup, Pasquet placed several spotlights around the performance space which are controlled by his computer, and which flicker intermittently during the performance. The piece feels like an installation: the spatial dimension seems to be an important part of the piece.

- *Space and lighting*: the space within which the piece is being played seems very important. Pasquet took time to place the lighting rig around the space and mapped them to different aspects of the sound. This is reinforced by the

² D.E.: 01_Code>02_Tools>06_Segmentation_Explorer_Standalone

³ D.E.: 04_Video>01_Demos>segmentation_explorer_demo.mp4

⁴ D.E.: 04_Video>01_Demos>traditional_segmentation_tool_demo.mp4

⁵ Run time approximately 12:00.

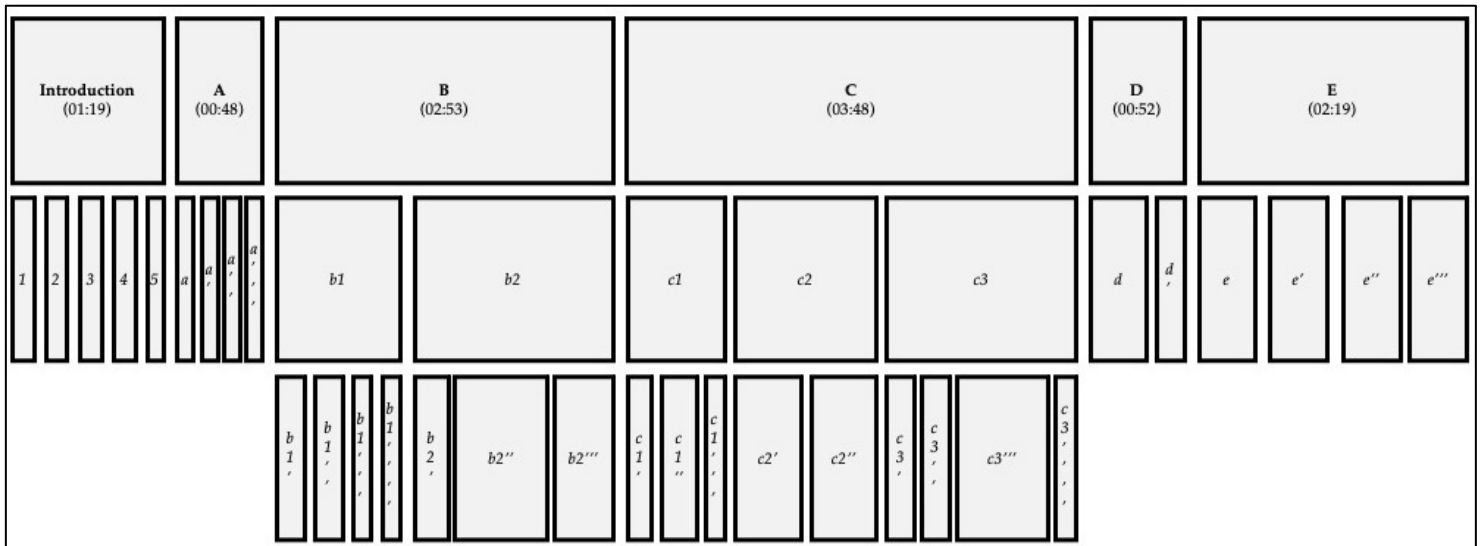


Figure 2: segmentation of Herbig-Haro.

fact that these are the only lights within the dark space, and that his stage is in darkness.

- *Absence of performance*: parallel to this, we note that this isn't a performance in the traditional sense. Pasquet is evidently doing things behind his laptop, however these actions are not to be seen, they are not the focus of attention.
- *Layering*: there are four notable layers to the sounds: clicks, sustained bass tones, pitched clicks, and vocals. These are four vectors which each intervene predominantly at different parts of the piece.
- *Smallest unit into sound*: at several moments we observe the phenomenon of these clicks, considered the smallest possible unit of sound, progressively morphing into other sounds. This is experienced notably towards the end when the pitched clicks morph into longer sustained events, and with the hybridization of the bass and clicks in the middle of the piece.
- *Succession of gestures*: the piece is divided into a succession of gestures, most of which are marked by silences. We can segment the piece into several collections of gestures, where each collection modulates a certain idea, often focusing on one of the four layers.
- *Frequency spectrum extremities*: Pasquet clearly seems to be operating in two frequency realms, each at extreme ends of the spectrum: very low bass tones and very high clicks and scrapes.
- *Dominant frequency*: much of the progression of the piece is driven by oscillations around a dominant frequency of *c sharp*. Departures from this tone are notable and become events in of themselves.

b) *Lauren Sarah Hayes – Moon via Spirit.*

Hayes' piece⁶ is titled *Moon via Spirit*⁷. Hayes is alone on stage surrounded by a whole host of machines: a Moog DFAM⁸, an Elektron Digitakt⁹, a TC-Helicon VoiceLive 2¹⁰, an ES-3 module¹¹, a CME XKey¹² controller and a Korg Nanocroncontrol 2¹³ controller. She holds in her hands a generic game controller and performs mainly with this. This and the various controllers are controlling a Max patch which is sampling and processing her voice. She uses several extended singing techniques. The segmentation for this piece is relatively simple with three clear large sections.

- *Continuous drones*: the pulsar synths are present and continuous throughout much of the piece. They provide a grounding layer from which other events emerge.
- *Synthetic sounds*: there is a heavy amount of characteristically synthetic drones and electronic drums. The piece seems to emerge from the electronic, pop world, taking those sounds as its primary material and manipulating them in unusual ways.

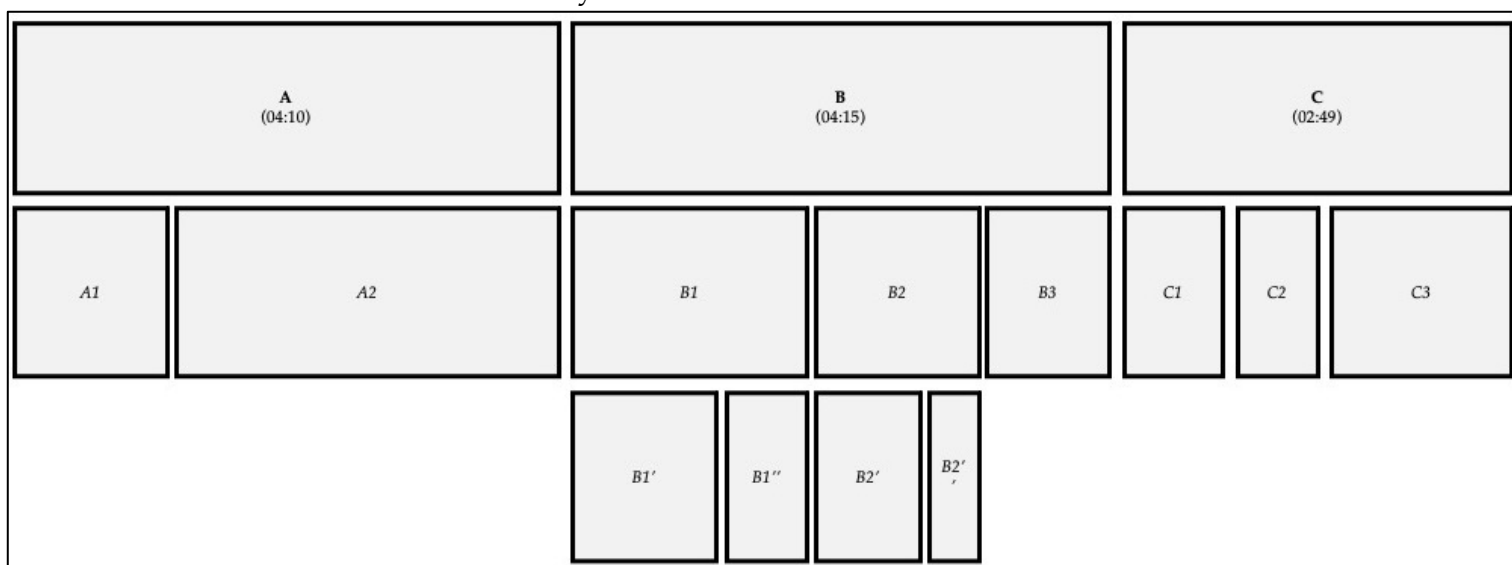


Figure 3: segmentation of *Moon via Spirit*.

⁶ Run time approximately 11:15.

⁷ An anagram of improvisation.

⁸ <https://www.moogmusic.com/products/dfam-drummer-another-mother>

⁹ <https://www.elektron.se/products/digitakt/>

¹⁰ <https://www.tc-helicon.com/product.html?modelCode=P0CMC>

¹¹ <https://www.modularsquare.com/fr/shop/expert-sleepers/es-3-mk4/>

¹² <https://xkeyair.com/xkey-air/>

¹³ <https://www.korg.com/uk/products/computergear/nanokontrol2/>

- *Recording to manipulation*: there is a clear trajectory across the piece that starts from a moment when material is being collected, to it being manipulated. This also appears at a local level, with gestures that take Hayes' singing into the microphone which is then immediately processed.
- *Improvised and extended vocal techniques*: the type of vocals that Hayes supplies is notably improvised and extended in nature. Hayes discussed¹⁴ having recently sought training for these techniques and is looking to extend her vocabulary further.
- *Interconnectedness*: many of the various elements seem to have some effect on the others, for example in Section C where the vocals seem to disfigure the aspect of the synth lines. This also appears in the setup of the various machines and gear.
- *Sampling*: sampling techniques are present at two major levels: live sampling of the vocals which are processed and transformed, and playback of pre-recorded drum samples. In the case of the vocals, sample playback is extremely unstable and segmented in nature, whereas the percussion samples are played in their entirety.
- *Expressive performance*: it is worth noting the expressive nature of the performance that Hayes gives. Her facial expressions and bodily movements are very earnest, especially when compared to others (for example, Constanzo who will remain relatively inexpressive and Burton who has his back to the audience). Some of her movements could even be considered dance-like, and her body seems to move with the musical gestures she is producing.
- *Controller display*: when watching the performance, the position in which Hayes holds the game controller is remarkable. It is turned towards the audience, as if she were explicitly showing the state of her hands to those watching.

c) *Rodrigo Constanzo – Kaizo Snare.*

Constanzo's piece *Kaizo Snare*¹⁵ sees him alone on stage with a kick drum, snare drum, a laptop, and a table of various items for playing the snare drum with – notably a

¹⁴ Appendix 8.3.1.

¹⁵ Run time approximately 11:20.

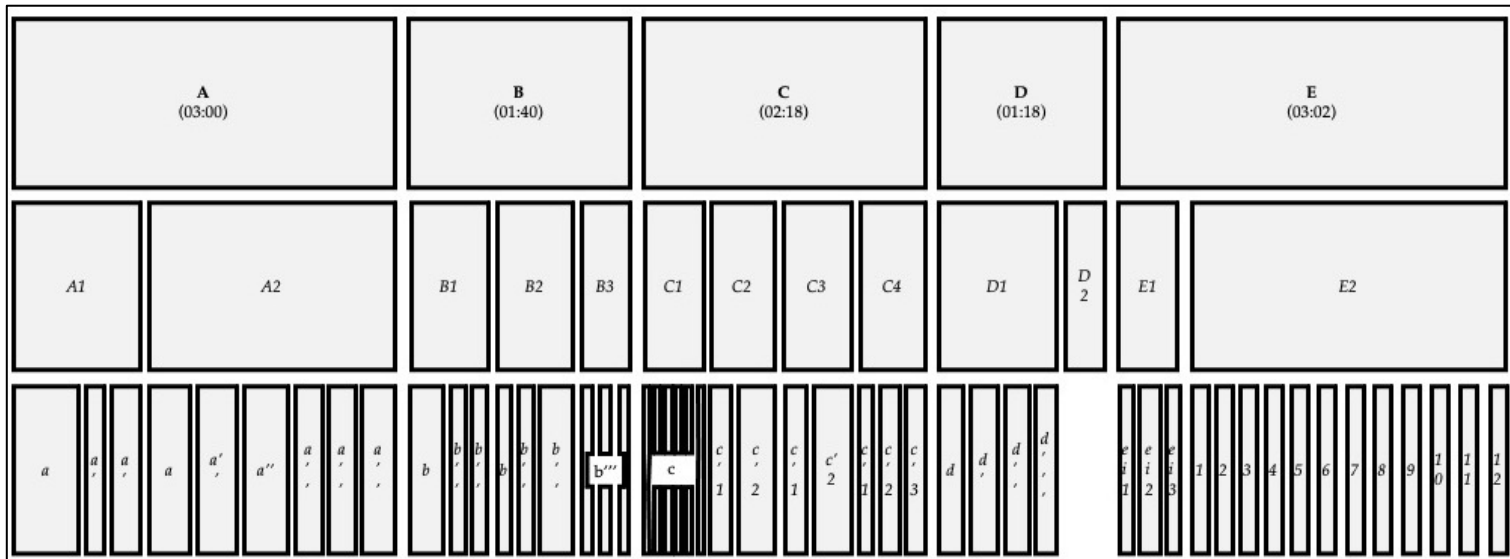


Figure 4: segmentation of Kaizo Snare.

condenser microphone. Attached to the snare drum he also has a small Musical Instrument Digital Interface (MIDI) controller that he uses to control his Max patch. Around and among the audience are a set of six small, hidden robots (solenoids) that tap on crotales of various sizes when triggered either with a plastic or metal rod.

- *Improvisation*: Constanzo talks¹⁶ about the improvised nature of the piece. Constanzo also does not use the entirety of the objects he made available.
- *Turntablism*: this is an import strand that is quite apparent when watching the performance. Here, the snare drum's surface replaces the vinyl. Constanzo, in some cases quite literally, transcribes gestural techniques from turntablism to his setup here.
- *Sampling*: we hear clips of sounds that do not come from the stage. Despite this, it's worth noting that all the samples do seem to be clearly triggered by actions that we see on stage.
- *Found objects*: Constanzo uses a variety of 'non-musical' objects against the snare drum – notably three different combs.
- *Feedback*: there are parts of the piece where feedback cuts into the sound. It is difficult to discern if this is voluntary or not – I discussed¹⁷ this with Constanzo after the performance.

¹⁶ Appendix 8.4.1.

¹⁷ Appendix 8.4.2.

- *Arc structure*: At various macroscopic levels we often discern gestures that are arced in nature.
- *Focus on the performer*: the lighting for the performance is very simple – the room is in complete darkness, save for an unmoving light on Constanzo.
- *Objects within the audience*: the robots are activated twice during the performance, they fuse well into the recorded piece, but on the night, it was quite surprising for the audience to suddenly find themselves surrounded by these sounds.

d) *John Burton – Line Crossing.*

Burton's piece *Line Crossing*¹⁸ is remarkable in that it lends itself very explicitly to various processes of segmentation. Burton is stood before a large screen that is showing a constantly evolving score. The screen and score dominate the stage area, focus for both the artist and the audience is very clearly directed towards it. The score is simple: a play head shows where we are in the piece, when there is a red block, the artist can play something, and their input is recorded. The blue block outlines show when part of this ever-growing recording is played back by the patch in some way.




Figure 5: the score generated for *Line Crossing* on the night of the *FluCoMa* performance.

¹⁸ Run time approximately 09:55.

The score is fundamentally segmented in nature – a series of blocks that follow on from one another (even if many overlap). This nature can be clearly heard in the piece, is visible in Burton’s playing, and necessarily inflects the experience had by the audience. I allow myself to be guided by this segmentation for my own initial segmentation of the piece.

- *Segmentation*: clear-cut segmentation is present at numerous different places within the piece: visually in the score, sonically (between different local events and groups of local events) and in the various instrumental configurations Burton passes through.
- *Arch structure*: I have identified an arch-shaped structure at several different levels: in the instrumental configurations, in the nature of the machine processing and in the amount and frequency of events.
- *Score transparency*: there is a notable act of exposing the score to the audience, and the conscious decision to fix all attention to it – not only because of its imposing position in the room, but because the artist is turned towards it rather than the audience. This suggests that the visual aesthetic of the score is to be considered when considering the experience of the piece.
- *Improvisation and extended techniques*: the nature of the material that Burton produces seems to be improvised¹⁹, and deploys some extended, albeit conventional, contemporary techniques. Also, Burton did not use all the objects that he had made available to himself.
- *Sonic contrast*: the piece seems to broadly contrast two main types of sonic material: long sustained sounds, and clusters of short, chaotic sounds. This is present in the machine blocks and in the material Burton provides with his playing.

e) *Alice Eldridge – Feedback Cello.*

Alice Eldridge was playing her *Feedback Cello*²⁰ setup. She was accompanied by Chris Kiefer who was controlling a modular setup. The piece is composed of long, flowing, sustained gestures which can make a segmentation challenging. However, there are

¹⁹ He gave a demonstration of the piece the previous day during his presentation at the FluCoMa plenary (<https://www.youtube.com/watch?v=r1LsSFEe4yM>), and the material was very different.

²⁰ Run time approximately 13:25.

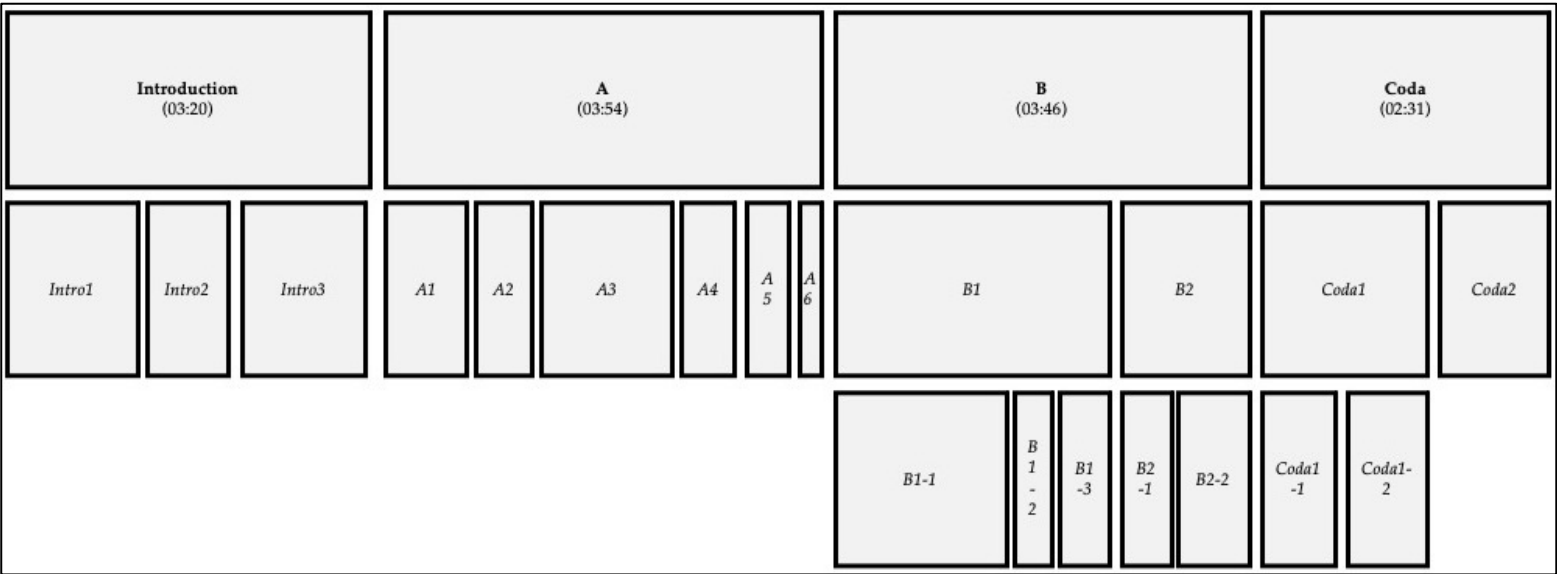


Figure 6: segmentation of Feedback Cello.

elements notably in the ways that the cello is played and the nature of the electronics that can guide a segmentation.

- *Cello versus modular synth*: the audience is presented with two instruments which are known to them, but which conceptually occupy very different cultural and musical worlds. The traditional conception of the cello is also challenged through the nature of the playing, we see some traditional techniques, but also some more extended ones.
- *Face to face*: the musicians are face to face to one another, side-on to the audience. We see them often looking to each other and communicating with facial expressions.
- *Feedback*: as the title would indicate, feedback is obviously an important structural aspect of this piece. There are moments in the performance where we can observe the *Feedback Cello* in action, where it is very clearly visible that gestures Eldridge is performing are generating or killing feedback.
- *Sustained material*: perhaps due to the structural idea of feedback, most of the material that is explored is sustained in nature. The moments where the cello performs content which is more staccato in nature are amplified by this. Often in these moments, they will still be submerged by sustained electronics, and sometimes fight to find space within them.
- *Source/modulation*: there is an idea of play and interaction between the cello and electronics, where sometimes it seems one is generating the other, and finally culminating in a moment where both seem inherently fused together.

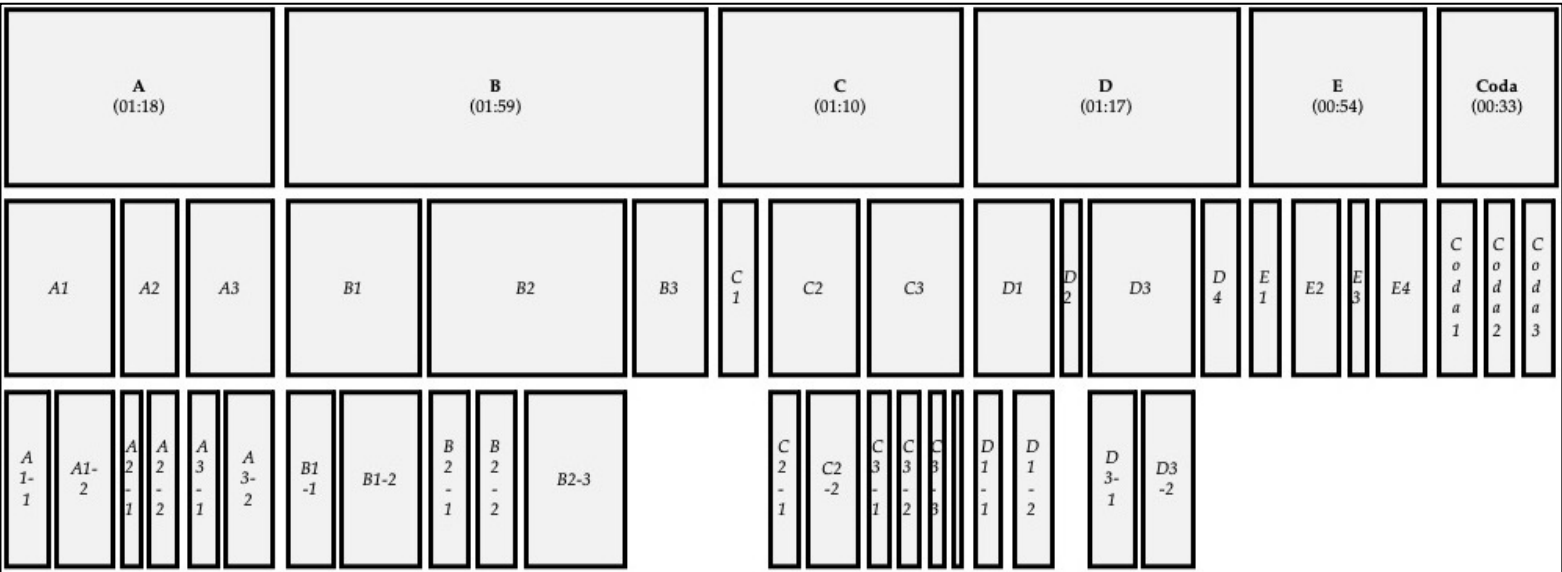


Figure 7: segmentation of Neural Duo I.

- *Bitcrushing*: aside from the feedback, partial-tracking, and reverberation processes which occupy most of the piece, the only other notably process that we can hear is that of some kind of low-resolution, synthetic bitcrushing. This only intervenes from time to time and is starkly amplified by its contrast to the rest of the sonic material.

f) *Sam Pluta – Neutral Duo I & II.*

Pluta played twice during the concert. His performance, a duo with trumpeter Peter Evans, was divided into two parts²¹. They were pre-recorded and played in video form: Pluta is performing with his laptop, iPads and joystick setup; Evans playing his trumpet. Pluta explained²² that each part was played separately: for Part I, Pluta recorded his part, then sent it to Evans to overdub, and vice versa for Part II.

- *Duo*: this duo configuration has the particularity of having been played out of time, and the musicians were not in the same space. Taking the video editing into account, very rarely do we see shots that only show one of the musicians alone – this is reflected in the sound, where both musicians are almost always playing together.

²¹ Run time approximately 07:12 Part 1, 05:30 Part II.

²² Appendix 8.5.1.

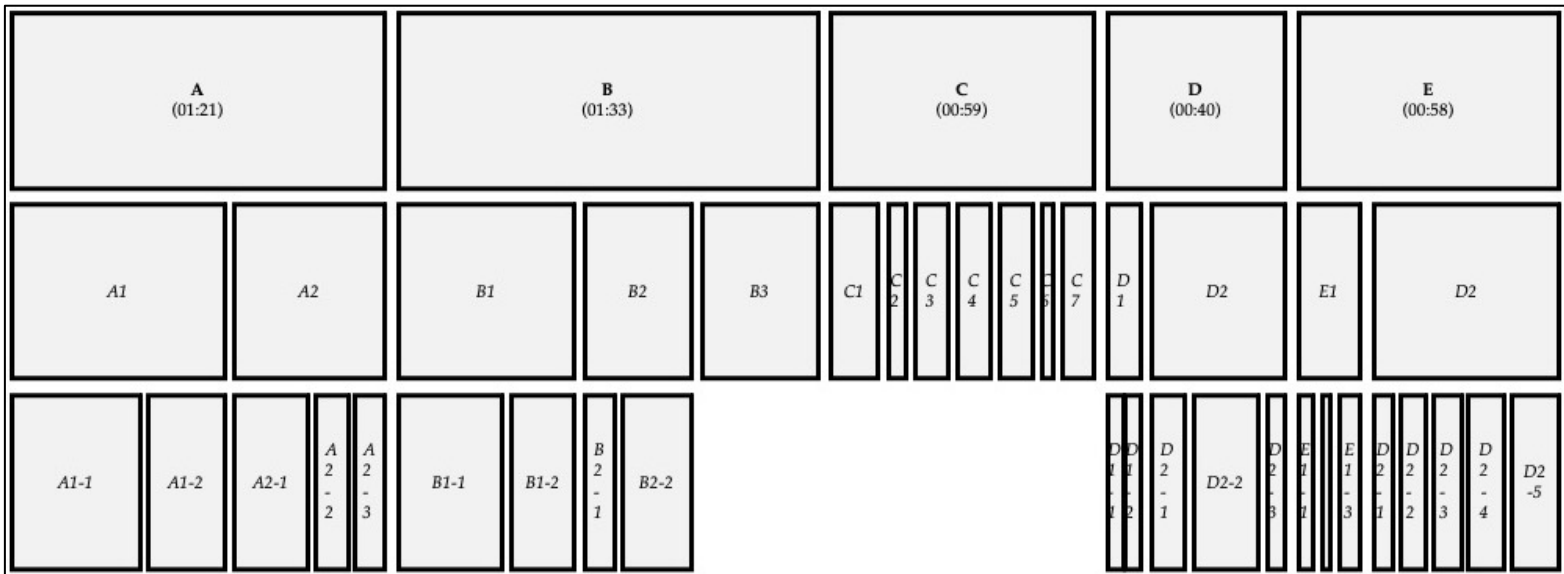


Figure 8 segmentation of Neural Duo II.

- *Counterpoint*: there are clear moments of counterpoint between the electronics and the trumpet: lines answer each other, mimic each other, moments will overlap and intersect.
- *Sonic material*: each performance has a particular set of sonic materials. This is especially notable with the trumpet: in Part I the trumpet is without a mouthpiece, exploring different fast-paced flutters, in Part II there is a mouthpiece, and most of the material is melodic in nature. The electronics are more similar across the two performances, though there is a notable presence of more sweeping, pad-like gestures in Part II.
- *Running fragments*: without close inspection, it is not easy to discern an overarching structure to each of the performances. We get the impression of having caught up with the sound while it was running – this seems to be reflected in the posture of each of the musicians at the beginning and ends of each performance. Events also rapidly flow into each other, with very few moments of suspension or silence.

g) *Alexander Harker – Drift Shadow.*

At a run time of 21 minutes, Harker's piece *Drift Shadow* is the longest piece that was made for the project. The piece is for oboe and live-electronics and was performed on the night by Niamh Dell. Apart from Tutschku's, Harker's piece is the only one where the performer was reading from a score. I have divided the piece into 5 large sections.

- *Multiphonics and partial tracking*: the piece seems to be clearly structured by an exploration of different multiphonic gestures in the oboe. These come in various forms: some sustained, others oscillating, others broken or bending. The electronics seem to work by tracking the partial content of each of these multiphonics.
- *Performer and score*: with Tutschku, this is one of the only pieces on the project with a score, and where the commissioned musician is not performing. Harker was present during the performance and was behind the desk supervising the electronics computer.
- *Modes of interaction*: there seem to be two different modes of interaction that are explored between the oboe and the multiphonics: one where the electronics emerge from the oboe, prolonging the partials found in the oboe, and sometimes even seeming to mimic the gestures that Dell plays; and one where the electronics are sustained, and Dell comes into interaction with them, modulating them by imparting gestures that delicately bend the sustain into new directions.
- *Sustained material*: it is worth noting the most striking feature of the piece. Most of the material in the oboe and the electronics is sustained, resonant

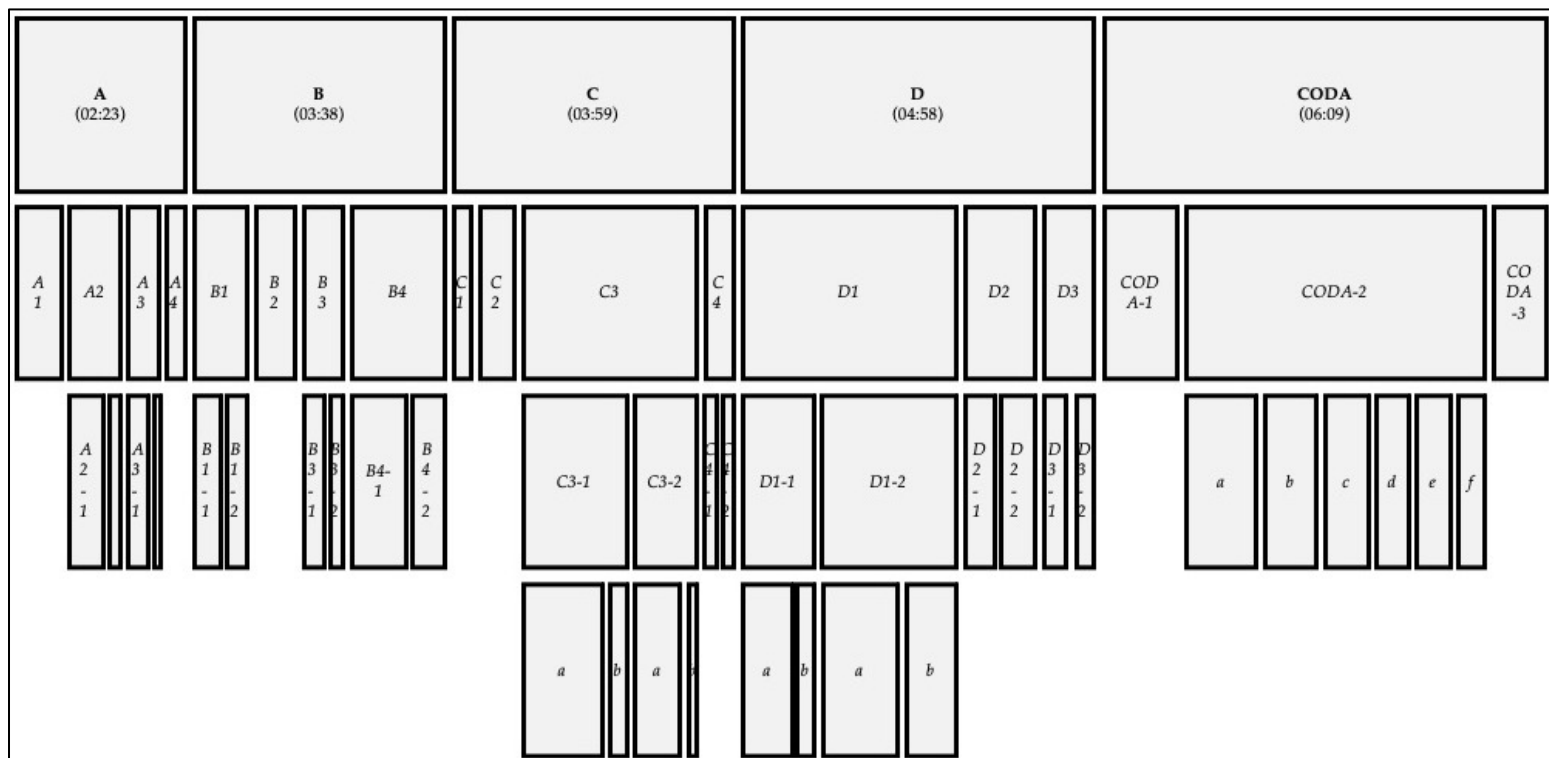


Figure 9 segmentation of Drift Shadow.

and emergent. Most events flow into the next over long periods of time, and anything abrupt is amplified by this.

h) Hans Tutschku – Sparks.

Like Pluta, Tutschku was unable to be present for the performance of his piece *Sparks*²³: if he would have been present, he would have been running the patch and electronics; in his absence, Tremblay performed this task. The piece is for piano and live-electronics, and on the night, it was performed by pianist Mark Knoop. In the preface to his score, Tutschku describes the piece as “*the next step in [his] 30 years investigation of the relationship between piano and live electronics*” (Tutschku, 2021). He explains that in “*an elastic relationship, the pianist is invited to make decisions based on what the electronics are producing at that moment [and that he was] hoping for a two-way interaction between both*” (Tutschku, 2021).

With Harker, Tutschku is the only musician to have produced a score, and despite the announcement in the preface, progression through the piece follows the score quite linearly. Indeed, the elastic relationship stems from the proportional notation of the score, and the presence of blocks where the pianist is invited to improvise in a relatively supervised way. The piece can be conceived of in five large parts.

- *Contrast of material*: in the piano, there is a clear contrast between two different types of events: busy, chaotic events versus sustained notes and

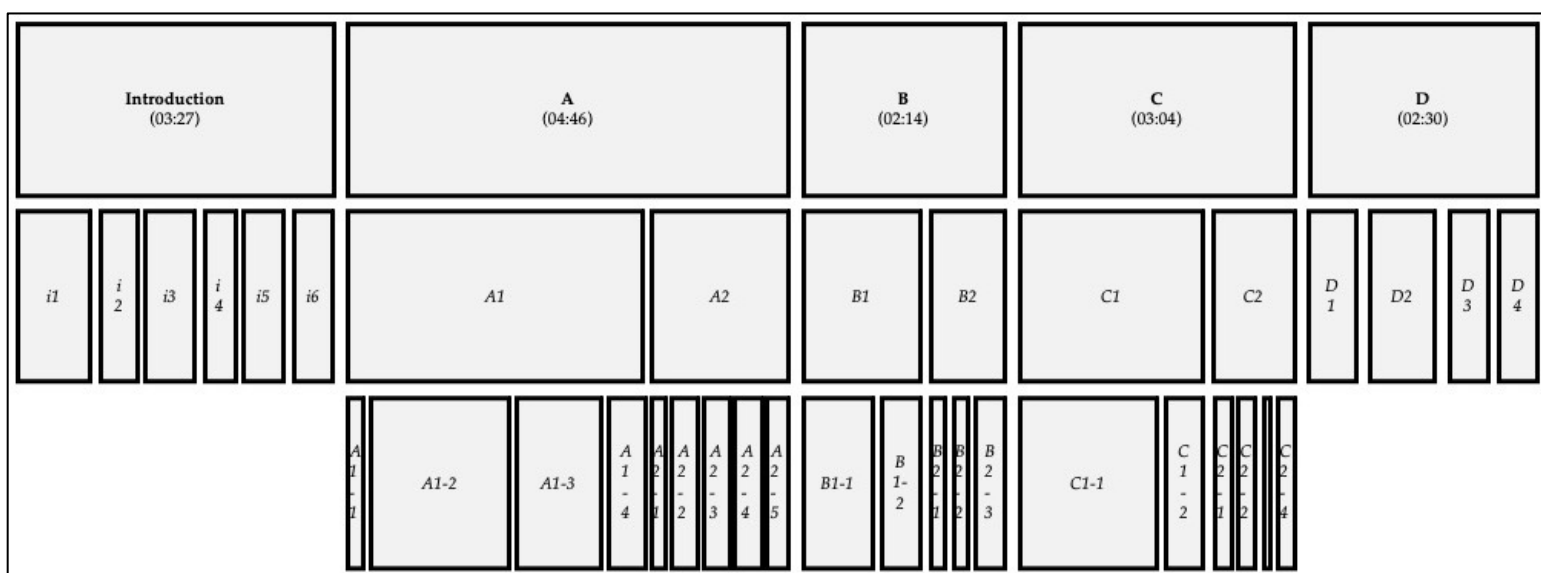


Figure 10: segmentation of Sparks.

²³ Run time approximately 16:00.

chords. The different sections that I have drawn are also structured by these kinds of events.

- *Resonance*: one of the major elements that structures the piece, and the electronics, seems to be this exploration of resonance in different forms. The term resonance is used rather than reverberation as the sustains really do seem to decompose the piano and prolong their resonance rather than being simply projected into a space. The space creates itself from the decomposition and treatment of the piano.
- *Delay*: the other structuring element of the piece and the electronics are the varying forms of delay that are heard: this varies greatly in time and length. The delay also seems to effect different parts of the sound: sometimes it appears to have more effect on the transients than the sinusoidal content and vice versa, participating in the construction of resonance.
- *Performer and score*: it is worth noting that, like Harker, Tutschku leaves the execution of the piece to a performer who is reading a score. Indeed, the two pieces of the project where the musician is not performing, are the two pieces which have a score.

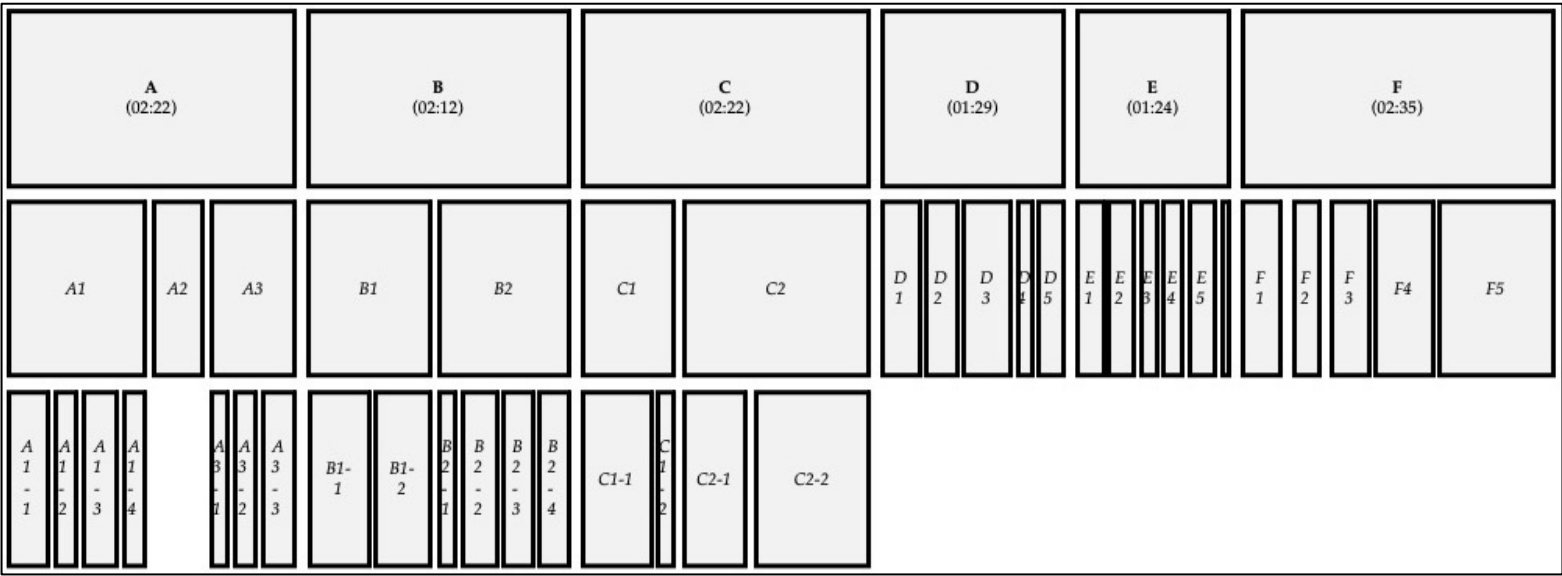


Figure 11: segmentation of Constructors.

i) *Richard Devine – Constructors.*

Devine’s piece *Constructors*²⁴ was a pre-recorded video of him performing with a modular synth setup and a number of bespoke percussion instruments. The piece can be divided into 6 large sections.

- *Modular synth and performance:* the modular synth setup is known by the audience, and comes with a certain cultural baggage. It is also difficult to understand what could be happening faced with the expansive hive of modules and wires. It also appears to be difficult to associate the gestures that Devine makes with a lot of the sounds that are being heard.
- *Reverberation:* Devine makes heavy use of reverberation processing which has different aspects²⁴ across the piece.
- *Suspension:* a lot of the piece seems to be structured by gestures that start with a loud impact, or some kind of chaotic material, which is then suspended and allows the resonance of the triggering sound to be explored.
- *Sonic material:* there are two different types of sonic material which are used: metallic, percussive impacts and chaotic synthesizers, sometimes encroaching on white noise; and low drones and harmonic pads. This second type of material intervenes much less than the first, however the resonances and heavy use of reverberation can also be seen as calling back to it.

²⁴ Run time approximately 12:30.

- *Video editing*: Devine clearly put a lot of effort into the editing of the video. There are several different camera angles, and white flashes between different transitions. The cuts all correspond to different impacts that are heard in the sound, and the changing of angles drives our perception of the piece.

2. Surface Level Analyses of the Instruments and Code.

I continue my surface level analyses with analysis of each of the instrumental setups. Again, I will give a musicker-by-musicker account, following the order of appearances in each of the gigs. I focus on two main elements: the physical setup on stage (or in the studio) of the various objects and instruments; and the structure and functionality of the various pieces of code and software that was written. This broad, surface level account precedes a much more in-depth, network visualisation of these setups (see Chapter 3). The goal of this section is to bring the reader up to speed about what was globally happening on stage. Most of this information was gathered through filming of the setup process, interviews with the musickers and extensive analysis of the softwares²⁵.

a) *Olivier Pasquet.*

There seems to be very little happening on stage in Pasquet's setup. It seems that Pasquet wishes to steer attention away from what is happening on stage

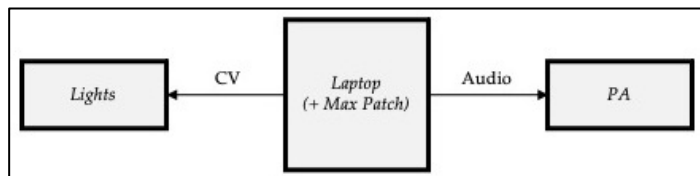


Figure 12: Herbig-Haro setup.

and immerse the audience within a space. It is within this space that things are happening, that the performance is to be experienced. Pasquet is sat behind a desk, barely visible, and is controlling things on his laptop. Sound is coming from the laptop and is fed through twelve channels around the space. Pasquet is also controlling eight spotlights from his Max patch, through a DMX USB Pro²⁶ lighting interface.

In the software, there are two main things to look at: the concert patch, and the generation patches. Pasquet doesn't use any of the FluCoMa tools in the performance

²⁵ The software sources can be found in the digital elements, some of them extensively annotated: *D.E.: 01_Code>01_Sources*.

²⁶ <https://www.enttec.com/product/controls/dmx-usb-interfaces/dmx-usb-pro-interface/>

patch – they were used to create sonic material to be played during the gig (the processed vocals).

Let's look at the performance patch first. There are a number of top-level controls on the parent patcher that he interacts with: they increment through 25 different states of the patch, parameters for rhythm generation, and the parameters of four Modalys²⁷ physical-model, virtual instruments. The patch makes heavy use of Pasquet's JTOL library²⁸. It is a set of objects that is built around the Bach²⁹ (Agostini and Ghisi, 2015) library and is *"dedicated to real-time pattern generation"* (Pasquet, 2012) and *"deals with multi-scaling and multi-dimensions where rhythm is considered to be a skeleton onto [which] everything else is attached"* (Pasquet, 2012).

²⁷ <https://www.ircam.fr/transmission/formations-professionnelles/modalys/>

²⁸ Built in collaboration with J.T. Rinker. <https://www.opasquet.fr/jtol/>

²⁹ <https://www.bachproject.net/>

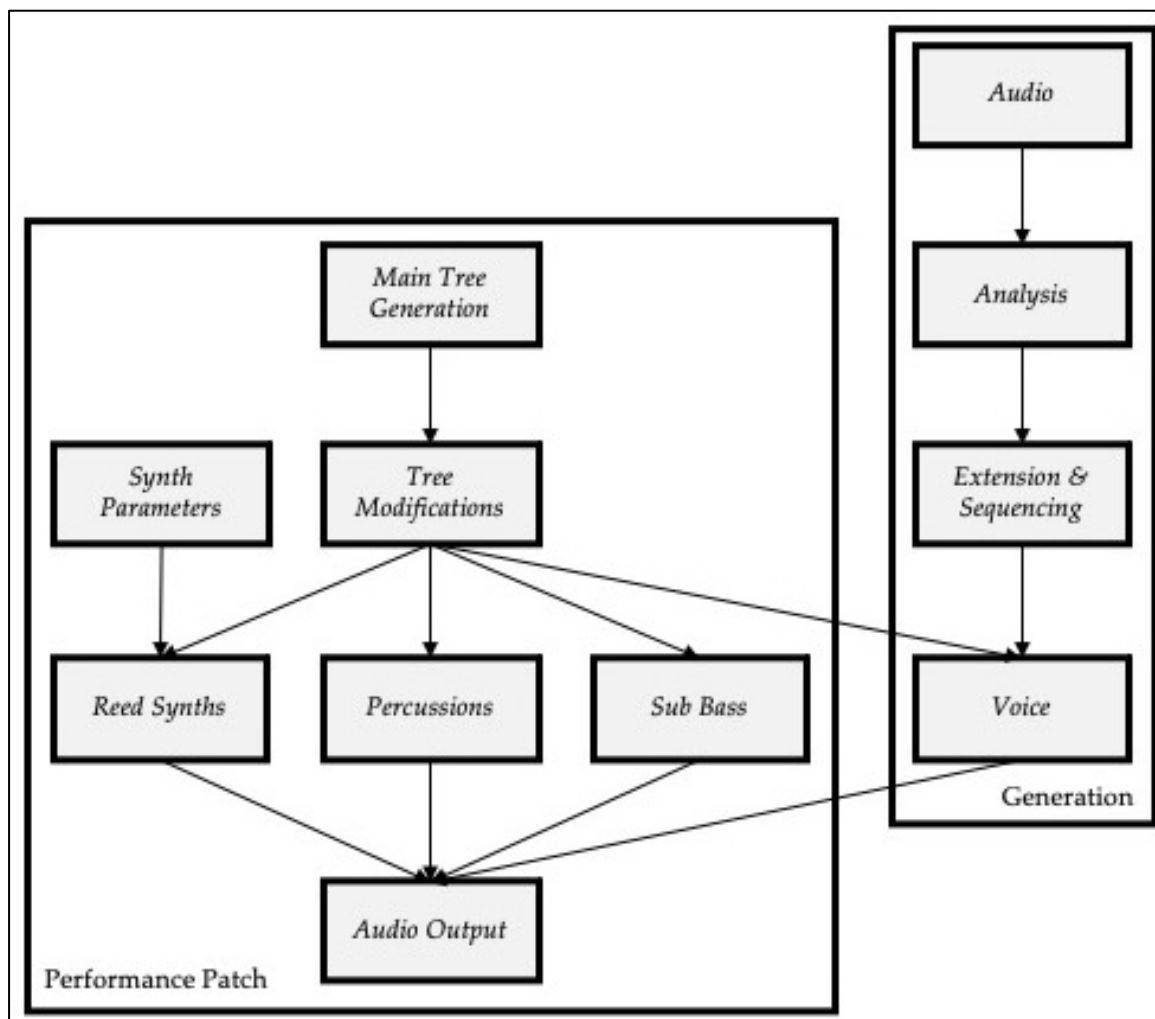


Figure 13: Herbig-Haro software overview.

One can effectively hear in the piece the explicitly rhythmic approach to musicking, with a lot of the material being clicks or very short fragments of sound. The JTOL library essentially makes use of the Lisp-Like Linked Lists (LLLLs) made possible with Bach, a data structure that can essentially be thought of as hierarchical trees. JTOL offers several objects for quick generation of trees, and manipulation, modification, and playthrough of them.

The patch is built around the generation of a main tree, using Pasquet's *jtol.bach.randtree* object – it will always be generated from a 1-2-1-1 base level tree. This tree is taken and then either used as-is or modified (adding silences, taking certain levels of the tree, or adding events) to drive several audio streams. Once playback of this generated tree has finished, a new main tree is generated, and everything starts again. This is what gives the piece the impression of being a series of phrases and gestures.

One of the advantages of the JTOL library is that despite all these manipulations of the main generated tree, everything will remain synchronised, as every level remains proportional to the branch above it. This also means that, despite always being generated from the same 1-2-1-1 rhythm, the scope for variation from that is very large – at the same time, there will also be a certain sense of cohesion, as everything has emerged from the same nucleus.

There are seven audio streams that are having events triggered by the main tree and its inflections: *mcsnare* (essentially a modified *click~*); *synth* (this is divided into four streams, each a different iteration of a same Modalys instrument, a vague reed instrument); *clicks* (another modified *click~*); *kick* (from a 3rd party Virtual Studio Technology (VST) Kick 2³⁰); *beat* (a short enveloped low-frequency sine wave); *sub_octavia* (a low-frequency sub-bass); and *voices* (playback of sonic material generated using the FluCoMa tools).

Pasquet also makes use of the IRCAM's spat³¹ reverberation, with events moving around the reverberation space, sound being output across 12 channels. Control of the spotlights happens in the *lightZ* subpatcher and is tied to events of the main generated tree³². Main control of the patch comes through an iteration through a number of parameter state presets: there is a timer that is restarted every time a new preset is set, and after roughly 28 seconds triggers a bang to flash on screen. Incidentally, 25x28 comes to 700 seconds, or about eleven and a half minutes – which happens to be roughly the duration of the piece.

Finally, we can look at some of the sonic generation patches that Pasquet made to create the voice-like audio that is heard during the performance. There are two main parts to this. First, Pasquet takes some recordings of lone vocals, and resynthesises them in his *hiphop_generator* patch. He uses the *supervp.scrub~* object from the SuperVP³³ library to resynthesise the vocals with varying levels of *sinusoidal*, *noise*, *transient*, *relax* and *error* components, giving the voice a synthetic, choppy nature. He then appears to have taken these recordings and treated them using the *LZW_NMF* patches. On the parent of this patch, Pasquet leaves an interesting comment:

³⁰ <https://www.sonicacademy.com/products/kick-2>

³¹ <https://forum.ircam.fr/projects/detail/spat/>

³² During my patch analysis, I set up some elements to help visualise how things are working, like the state of the lights during playback and visualisations of the various trees.

³³ <http://anasynth.ircam.fr/home/software/supervp>

“Idea: like one sound. Sound as untimed. This gives it a timing.

So physical model control is great.

Vocal synthesis is good because low entropy.

Particularly difficult with voice because of the expectation for meaning.

The voice in electronic music”

This comment seems enigmatic, but we can try and break it down to get an idea of what Pasquet was attempting to achieve. The initial intention seems to be contained within the first line: Pasquet advances the notion of a sound as being untimed – this can be read as: that which makes *one sound*, that which gives a sound coherency as a whole, renders the artefact as timeless. If a sound can be considered *a sound*, then there is no inherent temporal dimension within it. If a sound is contained within itself, then notions such as timing in a rhythmic, musical sense are evacuated. Indeed, I do not believe that Pasquet is advancing the idea that there is no *temporal* dimension in the sense that this research gives the term, but there is no activity across the *organisational* dimension. This is something that seems to be reflected in the nature of the sonic material that Pasquet tends to look to – clicks and short percussive elements. The goal of this patch, then, could be to find a way of offering a timing, offering content across the organisational dimension, to *a sound*, dare I add, to *a Record* that held no inherent deployment across the organisation dimension.

Before speculating further on the signification of this comment, let’s break down what is happening in the patch. Pasquet first randomly takes one of four synthetic speech files (each is between 10-20 seconds long; each gives a couple of hip-hop-related phrases). The file is loaded, then he performs two Non-Negative Matrix Factorisation (NMF) analyses: one looking to find 100 components, then another looking to find 10 that takes the bases from the first analysis that have been reordered according to pitch analysis, essentially looking at the components with the highest pitch. Pasquet then slices these components using the FluCoMa onset detector, and finally uses these fragments to create sequences using his *jtol.bach.lzw_encode* object. This is part of his implementation of Lempel-Ziv-Welch (LZW) compression algorithm (Butterfield et al., 2016).

Essentially, using the FluCoMa NMF and onset tools, Pasquet creates a conception of the audio that can be visualized as a piano roll – a set of pitched, rhythmic, and timed events. Then, using the LZW algorithm, he can encode this data, and using markov

chains, he can “*extend pseudo-repetitive sequences [that were] observed in the initial sequence*” (Pasquet, 2019). The idea is to take pre-existing data to generate new data: Pasquet states that he is interested in “*artificial creativity with autonomous production concepts*” (Pasquet, 2019). The results of his process here are the six sound files found with the performance patch which are played during the piece.

Returning to Pasquet’s initial comment, when he discusses physical model control as being great, and vocal synthesis being good because it has low entropy, we understand that this is essentially tied to the LZW encoding process. When explaining the process, Pasquet discusses the notion of the “*variety of scales*” (Pasquet, 2019) where the smallest units retain their *identity*. The algorithm, being a compression algorithm, will inherently lose data, be inaccurate, and will deteriorate the original’s resolution. He explains that the low amount of entropy in speech, allows for easier identification of *smallest units* (essentially, syllables) that retain identity, and thus proposes that it is a good candidate for this kind of approach. We see, then, that the *sound* that Pasquet mentioned, in the context of speech, is the smallest unit that retains identity: and as he states, *meaning*.

The idea with this patch, seems to be to take this sound, this sound that is mute across the organisational dimension, and give it timing. This can indeed be an interesting lens through which to hear the results. The produced audio is full of fragmented speech and syllables, which, while not conveying literal meaning, are recognizable and verging on the brink of meaning. He seems to run the line between noise and the beginnings of signification. These fragments do seem, however, to be smeared across time. They appear as if stepping on themselves in a very glitchy manner.

Musical thought.

- *Rhythm and trees*: the piece is structured by the constant generation and inflection of a rhythmic tree that always emerges from the following proportionality: 1-2-1-1. This tree structure means that there is a large scope for variation, but that the piece will be held together by a structural, rhythmic coherency.
- *Lighting and space*: we have seen that the lighting is tied to the playthrough of the main generated rhythm tree, indicating perhaps the importance that Pasquet gives to the perception of this element. Indeed, the space within which the audience find themselves is in darkness; it is articulated, it comes to life, *through* this overarching rhythm.

- *Sonic material*: most of the sonic material used is relatively simple, and notably short in nature. The spectrum is filled at its two extremities: clicks and grating high-pitched synths on one end; deep kicks and low sub-basses on the other.
- *Smallest time unit, the sound*: with this is the idea of what Pasquet calls the *sound* – an ‘untimed’ sound, the smallest fragment of sound that can have *meaning*. Pasquet also explores the idea of giving *timing* to this smallest unit, what we would call rhythm, or content across the organisational dimension.
- *Layers and tracks*: the way the patch is designed clearly shows that there are seven different layers of events that can be conceived of as tracks in a mixer.
- *Voice*: the generated voice sequences that intervene from time to time in the piece contrast with the clicky, disparate nature of the rest of the sonic material. They stand out as being much more coherent, long-form fragments of sound, and could be conceived of as occupying a prominent place within the piece. The preparatory work of generation that Pasquet performed also suggests that the treatment of the vocals was a real driving question for the piece’s conception. It’s also worth noting the content of the speech – clichéd hip-hop phrases, and the fact that Pasquet called some of his generative patches *hip-hop generators*. This reflects the place that the vocals would appear to hold in the piece: like in hip-hop, the voice is at the forefront.
- *Audio resolution*: the process that Pasquet performs with the NMF analysis and onset detection is very interesting. He essentially translates the original audio from what it is, to a comparatively low-resolution data structure. He then uses this for manipulation with a compression algorithm. This gives us insight perhaps on his conception of speech as a Record.

b) *Lauren Sarah Hayes.*

Hayes’ stage setup is much more complicated. She is surrounded by an array of machines: one felt that she was caught in the middle of the web of interconnected objects that were constantly sounding around her.

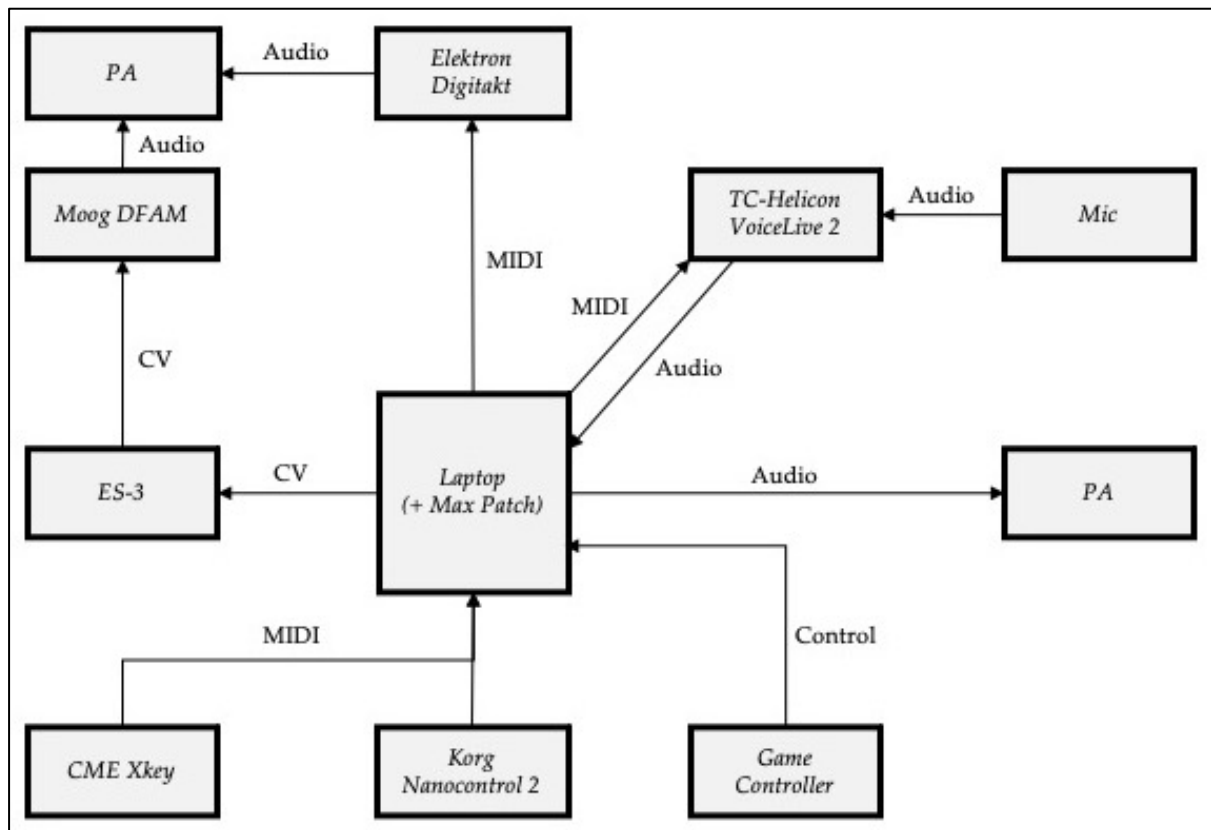


Figure 14: Moon via Spirit stage setup.

Primary control came from the generic GAME³⁴ controller, the Korg Nanocontrol 2, and the CME Xkey. The game controller is used to control the more performative parts of the patch, giving the object a much more sensitive and unstable nature, whereas the Korg and Xkey are more formal controllers for controlling modalities of the patch³⁵.

Hayes also sends control voltage to an ES-3 module to control a Moog DFAM analogue synthesizer which feeds its audio out to the PA. There are also two other sound inputs going to the PA that are controlled by the game controller and Korg through the patch: the Elektron Digitakt drum machine, and the TC-Helicon VoiceLive 2 vocal processor.

Hayes' patch is quite expansive. It was the first patch that I analysed, and I had initially wanted to create an interactive tool for explanation and exploration of the patches³⁶. This can be seen in a short demonstration video³⁷: essentially, I created a

³⁴ This is the name of the UK video gaming shop that Hayes acquired the controller from.

³⁵ I return to this in Chapter 6.

³⁶ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/01_Patch_Tour

³⁷ D.E.: 04_Videos>01_Demos>Patch_Tour_Demo.mp4

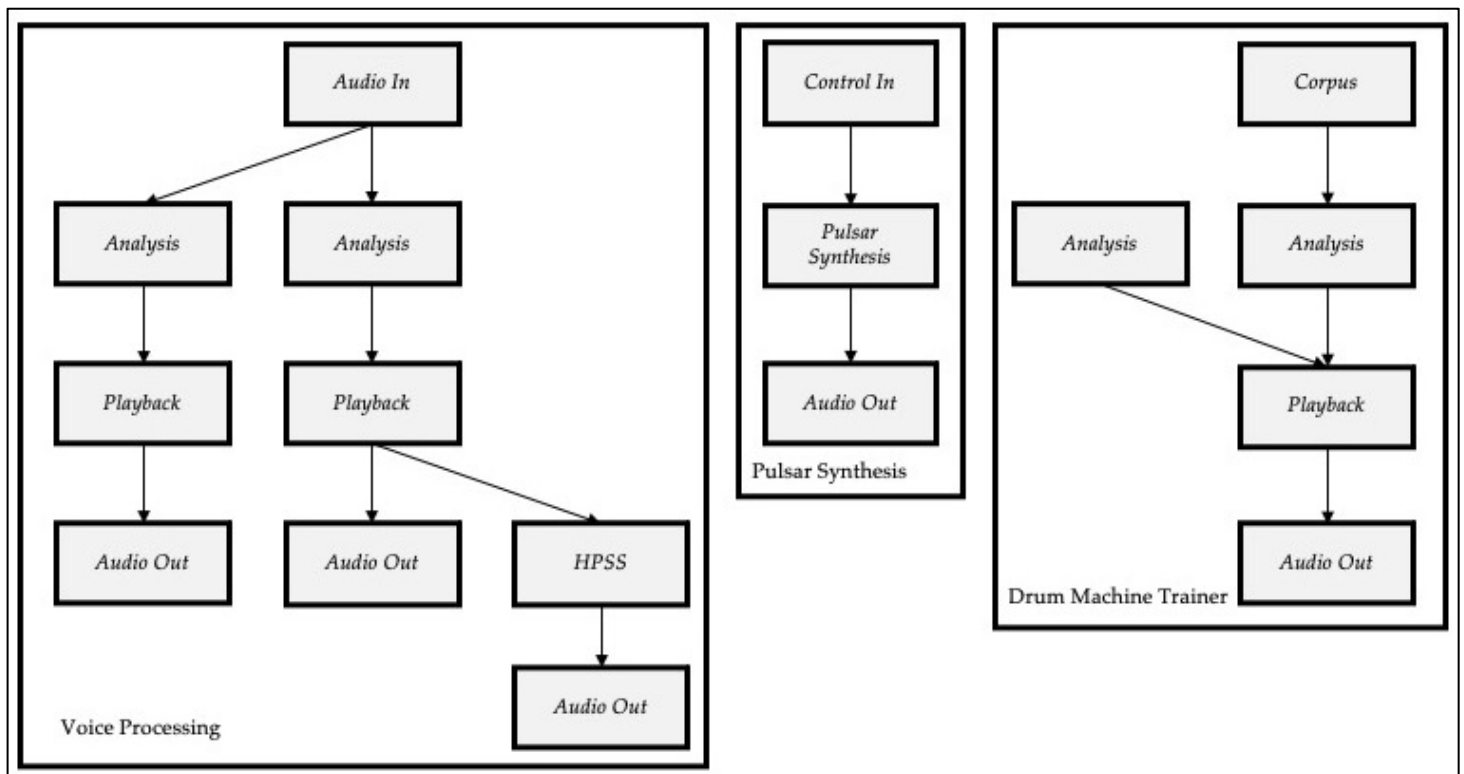


Figure 15: Moon via Spirit patch overview.

version of the patch that incorporated the interview I did with Hayes. As the video played, the things that Hayes opened-up and demonstrated would open in the actual patch³⁸.

We can conceive of the patch in three main parts: voice processing, pulsar synthesis and drum machine training. There is also some interfacing that occurs, allowing the Korg and Xkey to control the two external synthesizers.

Let us look at the voice processing first. This essentially boils down to recording into buffers with various types of event detection; and playback at varying speeds and filtering. The first part can be seen in the *techno2* abstraction. Here, audio is recorded into two buffers when the incoming audio is over a certain amplitude: one buffer is constantly being overwritten, and one will be sliced using the FluCoMa novelty slicer.

Playback of these elements is being controlled by the game controller – notably crossfading between output of the two streams; the rate of playback is controlled by the Xkey. This audio is output, but also branches out to the Harmonic Percussive

³⁸ I believe that this is potentially a very useful tool, however, the process for building this interface was quite a time-consuming one. With more time, I would have optimised the process and would have sought to do this for each of the musician's projects. However, I do not deem it to be essential to this research – it is a promising avenue for future research, notably regarding the pedagogical possibilities it offers.

Source Separation (HPSS) abstraction. Hayes describes³⁹ using this essentially as a filter – she uses the FluCoMa HPSS object to increase the gain of the percussive element. This audio is then also output.

Finally, the voice is also fed into the *knife* subpatcher. Here, audio is again being taken in, and this time sliced using the FluCoMa transient slicer. With the game controller, Hayes can choose which slices are being played, looping them if desired. This audio is filtered through a ring filter, then a high-pass filter, before being output.

Another stream is the pulsar synthesis. In the code Hayes handed-over for the project, she replaced the actual synthesis patch with a demonstration patch, giving a simple form of pulsar synthesis based on *roger.carruthers*⁴⁰ example on the Cycling 74 forum⁴¹. This is controlled by sliders on the Korg.

Finally, we have the *drum machine trainer* section. Here, Hayes uses the FluCoMa NMF implementation to create a corpus of sounds from various sound corpora: prepared piano, bowed saws, buchla sounds and synthetic percussion. Essentially, she is looking to build up a corpus from the thousands of files on her hard drive which are similar to the ones she uses in training by analysis of the training corpus through NMF and looking for similarities in the activations found. These are combined into a buffer, which is then sliced by transients and reordered by pitch using the FluCoMa tools. Playback of this is then controlled by amplitude detection on the pulsar synthesis signal.

There appears to be a difference in the roles of the Korg and the game controller. Indeed, we see that the game controller is used to control recording, selection, and playback of vocal slices. We discussed⁴² the nature of this object and its ties to physicality, and how it can inform her experience as an improviser. I asked her how it compared using a game controller compared to a lot of her other work with haptic feedback – notably the fact that this would be hard to achieve with a controller.

She explained that with the controller, she was looking to find a type of *resistance*, a type of *tension*, but in a different way. Instead of having to physically push back against an object, here she must keep things incredibly still if she does not want to

³⁹ Appendix 8.3.3.

⁴⁰ <https://cycling74.com/author/557c789e1e00771a62041d59>

⁴¹ <https://cycling74.com/forums/page/1>

⁴² Appendix 8.3.2.

create drastic changes in the sound. She discussed Martin Parker's (Parker, 2004) work, who points out the physical ease with which one can move the object of a joystick and explained that she put in place extremely crude scalings that indeed made moving even a millimetre have drastic effects.

This idea of tension with the instrument seems constant across a lot of her work: when looking at this instrument from a broader perspective, one can see the fragility of what is happening, purely through the sheer amount of interconnectedness of the software and hardware. During performance, we can follow her struggle while grappling with this almost autonomous network through the posture she must hold (she mentions⁴³ for example having to hold parts of the controller in her mouth if she wants to change something else) and her facial expressions. This is something that seems to drive her improvisation practice. The Korg, on the other hand, seems to play a much more pragmatic, functionary role, notably in the fact that it is an object that can maintain its state autonomously. The controller is forever bouncing back to its initial configuration, and one must constantly engage with it.

Musical thought.

- *Hardware, materiality and tension:* it is clear that Hayes does not limit herself to the computer and looks to many sources for sound generation. This could also imply that the main focus of her practice does not necessarily revolve around things like sound synthesis, or sound design. She appears to be looking to put in place a space where she can improvise, a place with *tension*. There is some notion of a struggle with the objects she immerses herself within.
- *Event detection:* there are two different types of event detection: onset and transient. These are used in parallel across the patch. One can imagine that they imply for Hayes a fluid conception of what an event to be detected is. We see that these different events are also processed in different ways: the onset detection is crossfaded between itself and passed through an HPSS filter whereas transient detection is played back in extremely short bursts and passed through a ring filter.
- *Modular, audio streams:* a working practice for Hayes appears to be to work in a modular way, conceiving of audio streams that interact with each other.

⁴³ Appendix 8.3.2.

Beyond the hardware, we see also the idea of interconnectedness in the patch with things like the pulsar triggering events in the drum machine trainer.

- *Object ontology*: we see that different types of hardware play different types of roles, which informs us on Hayes' conception of the ontology of these objects. I return to this in Chapter 6.

c) *Rodrigo Constanzo.*

In Constanzo's setup, everything is centralised around a Max patch. The audio that comes into this is from the condenser microphone that is held in his left hand. This microphone is encased within a handle that Constanzo designed himself using 3D printing. The microphone is directed towards the snare drum surface and is applied in various ways: alone on the surface or just above; crushing a small felt pad on the surface of the snare; or against one of the combs. Attached to the snare is a Novation Dicer⁴⁴ MIDI controller which Constanzo uses to control various aspects of the Max patch – the mount for this is again self-designed and 3D printed. A crossfader, again built by Constanzo through a long process of experimentation and 3D printing (Constanzo, 2020), is put through a Teensy⁴⁵ microcontroller, and sends high

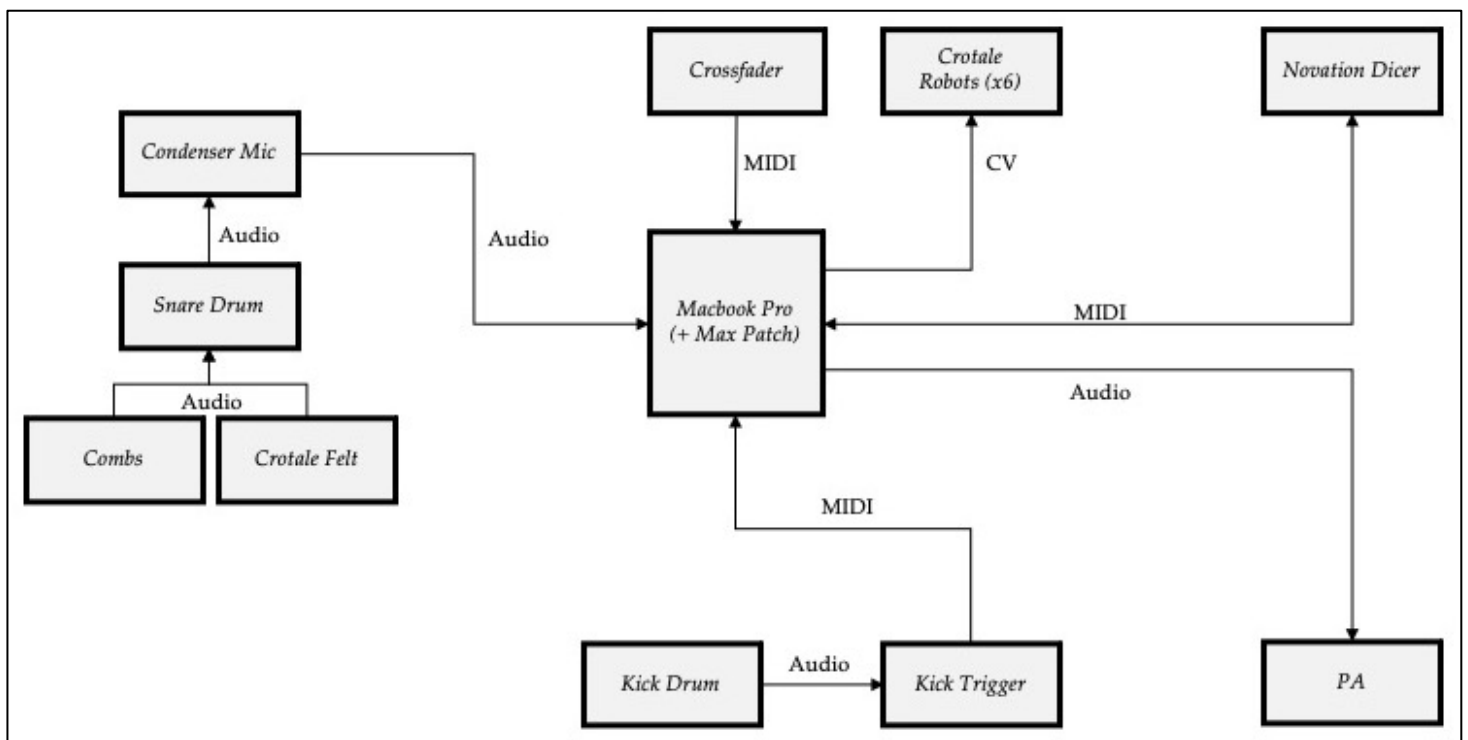


Figure 16: Kaizo Snare setup.

⁴⁴ https://www.musicstore.com/en_OT/EUR/Novation-Dicer-Pair-/art-DJE0003600-000

⁴⁵ <https://www.pjrc.com/store/teensy40.html>

resolution MIDI data to the patch. The six crotales solenoids which are placed around the audience are controlled via control voltage from the Max patch. Finally, there is a kick drum trigger attached to the kick drum which is sent into the Max patch.

Constanzo is a notoriously tidy programmer, and always takes care to make his work very clear and readable. This made analysis of his software a comparatively painless task. To make things easier to follow, I have divided the patch into three main signal paths (SP-A, SP-B and SP-C).

In SP-A, audio from the snare comes in and is split by onset using the FluCoMa onset slicing object. It is then analysed using the FluCoMa loudness and spectral shape descriptors implementation; and pitch, again using the FluCoMa tools. These analyses are used to query two things. First, a sample library of metal resonance sounds that

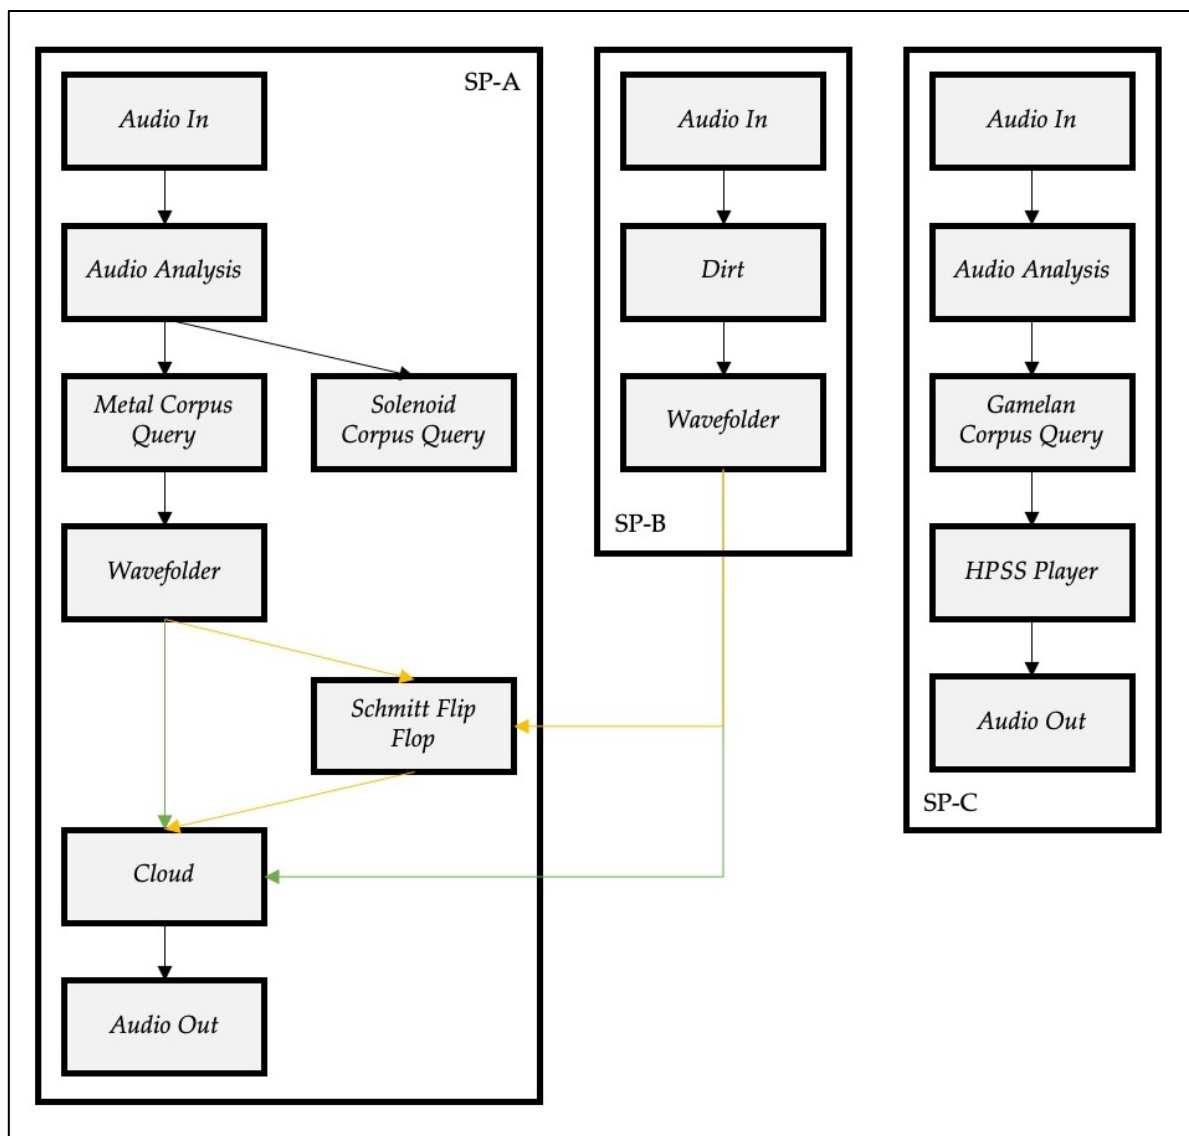


Figure 17: Kaizo Snare signal paths.

have been pre-analysed – Constanzo used Alexander Harker’s *entrymatcher* object to find the closest sample to the given audio. There is also a condition here, if the patch is in *yellow* mode (controlled using the Novation Dicer), samples with ‘bounces’ are allowed to play.

The second corpus they query is a corpus of solenoid recordings. When in the right mode (*Sol1* or *Sol2*), they trigger the action of the solenoid robots which each have a plastic or metal rod (depending on the matched entry). Then, the samples are played, and fed into two parallel *wavefolder* effects (see below). Then, once again according to the mode the patch is in, if the mode is *green* audio gets fed into the *cloud* effect (see below); if the mode is *yellow*, audio gets fed into the Schmitt *flip-flop* subpatcher (see below) before being fed into the *cloud* effect. Finally, audio is output.

In SP-B, the same snare audio comes in, but it is this actual audio that will be processed and used. It is first clipped and filtered, then it is processed by the *dirt* effect module (see below). This audio is then blended between the dry version of itself and a version that has gone through another *wavefolder* effect. It is filtered, then goes in one of two places according to the mode. In *green* mode, it is sent to a subpatcher that adjusts the gain according to the position of the crossfader, and is then fed into the *cloud* effect in SP-A; in *yellow* mode, the audio is fed into the afore-mentioned Schmitt *flip-flop* subpatcher in SP-A.

SP-C deals with audio coming in from the kick drum. It is sliced using the FluCoMa onset detector and is then analysed according to loudness and the spectral shape set of descriptors. It is also ‘auto-scaled’, which scales the description according to what has already been played in the piece. Constanzo then takes this analysis and queries a library of gamelan samples, looking to match loudness. This sample is taken and is transposed according to a random entry in a *coll* called *gong_pitches*. This *coll* is populated in SP-A at analysis when the pitch confidence of a slice is above 0.5. Constanzo explained⁴⁶ that this was to create variation in the gamelan samples, as it is a relatively small library. This is then sent to the *hpssPlayer~ poly~* object, which is built around the *fluid.bufhpss~* object. Essentially, the audio is divided into three streams: the transient component which is played from the start; the harmonic component which is enveloped and overdriven; and the percussive component, which is enveloped in a much more jagged way, and passed through another *wavefolder* effect.

⁴⁶ Appendix 8.4.4.

There are two very important aspects that the instrument seems to be built around: the three effect modules, *wavefolder*, *cloud*, and *dirt*; and the crossfader which controls several things in the patch. The crossfader intervenes in the following places: the speed of its movement controls the *timbre balance*, *panning width* and *activity* parameters of the *cloud* effect; the blend of the dry and *wavefolder* signals in SP-B (from the right to the middle, it fades from silence to the dry signal, from the middle to the left, it fades from dry to wet); and the gain of audio input into the *cloud* effect from SP-B. In his blog post (Constanzo, 2020) about his work, Constanzo discusses the importance and time he accorded to the development of the crossfader: “*I really wanted to optimise that level of control – that level of touch – as it formed the foundation for much of the sonic material in the piece*” (Constanzo, 2020) Indeed, the setup emerged from the idea of a classic turntablism setup, where the vinyl record is replaced by a snare drum surface.

In Appendix 3.1, I give a full rundown about the mechanical workings of the three effects modules – this is important as they will be important elements for analysis in Chapter 4. Here, the reader can understand that the *wavefolder* modules implement exactly that: a fairly standard wavefolding algorithm; the *cloud* modules Constanzo describes as “*a real-time granulator that is triggered by onset detection such that whenever an attack is detected a new bit of audio is written into a circular buffer and bits from that rotating pool of attacks are played back. So, it’s kind of like separate record and play heads for a looping delay that only operate when an attack is detected*” (Constanzo, 2020); and the *dirt* module implements two different distortion algorithms.

Musical thought.

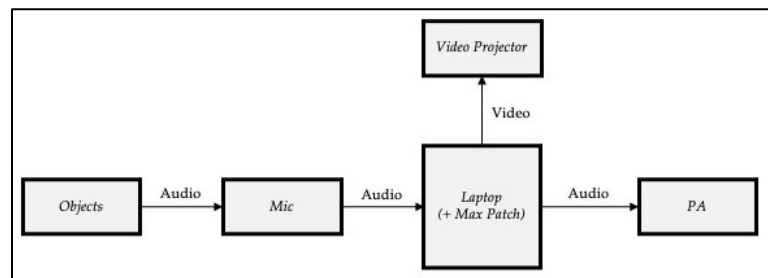
- *Improvisation*: despite the improvised nature of the piece, there seem to be clear, pre-defined moments at which Constanzo wishes to arrive. It seems that Constanzo conceives of a number of moments within which to improvise: these moments are to be conceived of with the two main modes, green and yellow.
- *Drumming, turntablism and hybridisation*: after exploring the patch, we realise that the techniques here sway much more towards turntablism than drumming. In most cases (except for SP-B, although this can be cut off in *yellow* mode), the activity on the snare drum and kick drum are in fact triggers for sample playback. There is, however, a positive connection between the gestures on the drum and the sounds being played through

descriptors and entry-matching. We thus begin to conceive of the idea of hybrid sounds – or at the very least, the drum surface as a *point of interface* between gesture and sound (see Chapter 6).

- *Feedback*: we see the extent to which Constanzo went out of his way to install safeguards against feedback getting out of control. This was a preoccupation of his, and it seems that feedback can be considered an undesirable phenomenon.
- *Harsh noise and distortion*: we see over the patch that there seems to be, if not a study, at the very least a taste for distortion in various forms. There is simple clipping, the *overdrive~* object, the *dirt* effect module and the *wavefolder* module which occurs multiple times with various parameter settings.
- *DIY*: a lot of the objects seen on stage were physically fabricated by Constanzo himself using 3D printing.

d) *John Burton.*

Burton’s stage setup was fairly simple: there was a table of various objects, some of which Burton used during the performance, some not⁴⁷.



During performance, Burton *Figure 18: Line Crossing setup.*

used these objects in various extended ways, and fed the audio into his computer through a microphone. The other notable part of the setup was the large screen that dominated the performance area. This was being projected via a video projector, the signal of which was being emitted from Burton’s Max patch.

Quite a substantial factorisation process had to be done to read Burton’s patch – it was also necessary to make some modifications for an enveloping window to work for some of the processing⁴⁸. Ultimately, the software can be conceived of in two parts: score generation, and audio processing.

⁴⁷ Appendix 3.2.1.

⁴⁸ This didn’t have much consequence to the performance, effectively, this envelope window ended up as being hard-coded as a sine wave.

There are two parts to the score: *human blocks* in red and *machine blocks* as blue outlines of blocks. The data generated for these scores is kept in two corresponding *colls* – making for easy exportation as text files. Each block, each event, has a set of characteristics: type (the shape of the block: ramp up, ramp down or block); onset time, length, and amplitude. On the parent patch, there are two streams of code that run in parallel that each deal with the human and machine score generation – essentially, when one human block is created, a certain number (see below) of machine blocks is created further on, and a new human block is created when the end of a human block is reached.

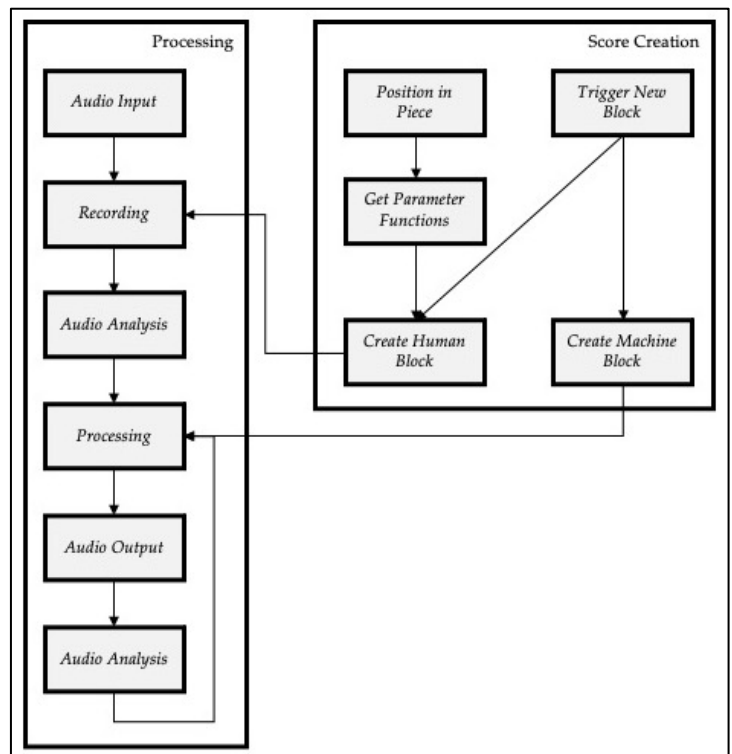


Figure 19: Line Crossing patch overview.

Burton uses weighted random number generation to create blocks as the piece goes on. There is a subpatcher called *human_block_param* where we find a series of *function* objects which control the weightings of various aspects of this generation as the piece progresses. The parameters that Burton controls are human block pause length range (0-400); human block length range (10-50); human block amplitude range (20-100); machine block density (0-16); and clatter (1-7 – this controls the types of processing to be used and is a stepped function, see below).

This is an interesting patch that lets us explicitly see the macroscopic structure of the piece as it was set by Burton himself. This is quite a versatile way of controlling the macroscopic structure: it is very easy to make large-scale changes just by modifying these few top-level parameters. In my initial surface level analysis of the piece, I discussed the idea of an arch-shaped structure to the performance. This indeed seems to be the case for the first *pause length range* and could eventually be said for the others⁴⁹.

⁴⁹ Although the smoothing needed to achieve this arch-shape would have to be relatively high.

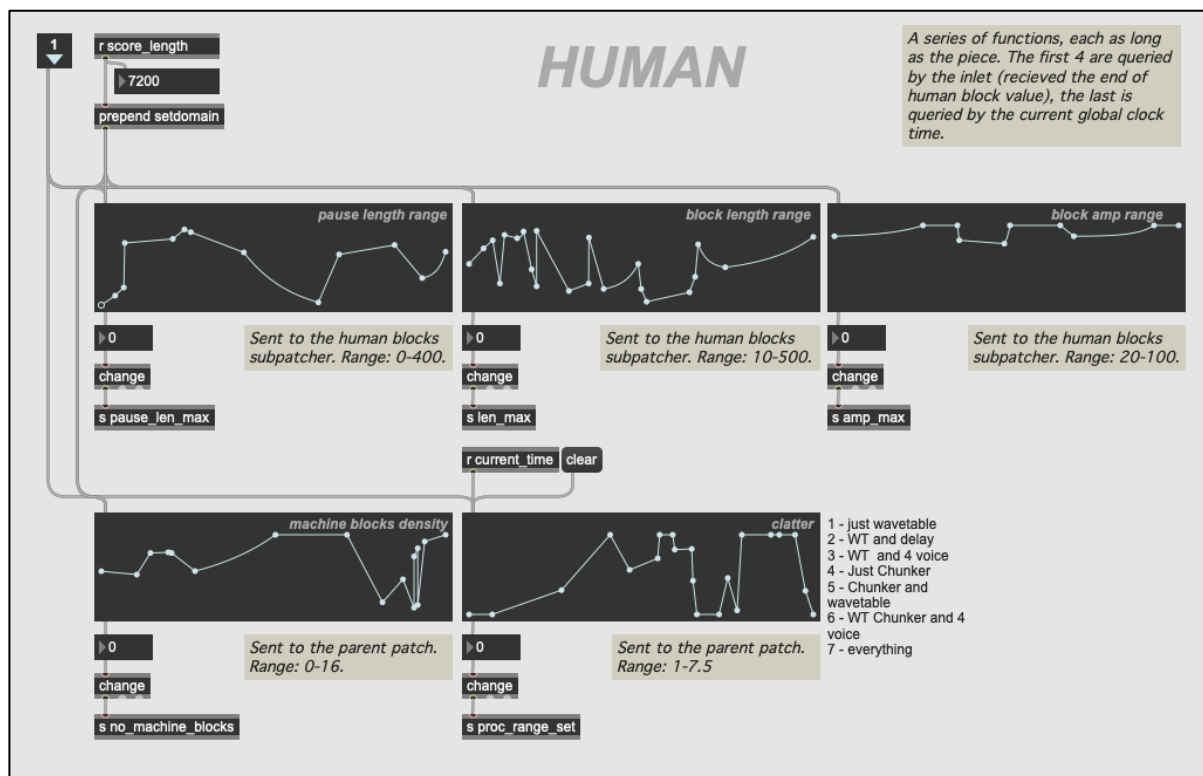


Figure 20: the subpatcher human_block_param.

There is a global clock that drives the patch, ticking 1 pixel every 33ms, giving the global current position in the piece. On each tick, Burton checks a number of things: if input needs to be recorded into the main rolling buffer, if a new set of score blocks needs to be created, and decisions relative to the playback and processing of the input audio buffer. Playback of the main input recording buffer is controlled in the *machine_block_playback* subpatcher. This contains 16 different *voices*, all abstractions called *PlayerBlok2*. Every time there is a new machine block, it is sent to the next available voice.

This abstraction contains 4 different audio processors which play back segments of the recorded audio (see below). The mix between the output of these processors is chosen from the *clatter* function mentioned above where the 7 different positions correspond to the following: 1, wavetable; 2, wavetable and delay; 3, wavetable and 4 voice, 4, chunker; 5, wavetable and chunker; 6, wavetable chunker and 4 voice; 7, everything. As we can see, the wavetable processor has the most chance of being heard.

Each processor is contained within its own *poly~* object (with only one instance of the *poly~* for each voice). See Appendix 3.2 for a description of each processing module.

Finally, we can mention some of the audio analysis which occurs in the patch and use of the FluCoMa tools: first, global audio output is analysed – Burton looks for the mel band content. He looks to see what the loudest mel bands are, and then in some of the audio processing units, he uses this information to control some of the random generation to try and make up for the quieter mel bands. Indeed, balance in sound design seems to be quite important for Burton (see below). Finally, Burton makes use of the *fluid.bufloudness* object to retrieve and order the loudest segments of the main audio buffer and write them to a *coll*. This is then used by the wavetable patch when choosing chunks from the main buffer.

Musical thought.

- *Processing painting*: the various audio processing units are mixed according to 7 predefined combinations which have varying chances of occurring across the piece. The various processors are not considered individually on a macroscopic scale, but rather shades of combinations are considered.
- *Functions & top-level parameters*: Burton has created an interconnected environment where various things control others – however at some point he decided to allow himself control over a small number of key parameters which in turn affect everything else. Considering that these are the only parameters that Burton gives himself access to, we can imagine that they may have either an important aesthetic function or a notable sway on the working of the patch.
- *Controlled chaos*: generation which could seem random at a first glance, is actually quite controlled. Across the piece, the state of a given parameter, while indeed randomly generated, is generated within a constrained range. This allows for variation over iterations, but still some degree of control over what will happen.
- *VSTs*: some VST modules are used within the patch, notably EQs and compressors. These attest to two things: first the desire for a sound quality towards a more commercial level⁵⁰ – there is a clear desire to control and construct the spectrum in a way that will be agreeable to the listener. Second, this shows that Burton is not squeamish about deferring processing to

⁵⁰ This is reinforced by numerous normalisations, filterings and envelopings across the entirety of the patch.

readily available commercial products, which amplifies the significance of the processing he did decide to make himself.

- *Patching style*: it's worth noting that Burton's patch was rather messy. This may of course be mainly due to the deadline-driven nature of the project it was created for, but it can be considered. It seems that Burton is generally a 'messy patcher': some of the patches which seem to have been repurposed from other projects such as the audio processing units are still quite messy; patches he has shown at various plenaries are also messy. This can have significance. There is however some effort to make things readable and make the code visually pleasing to himself (the *score builder* title comment has been multiplied, put in various colours and offset – there is no inherent need for this other than visual taste). This suggests that Burton may be very much at ease with the cluttered layout of the patch.

e) *Alice Eldridge.*

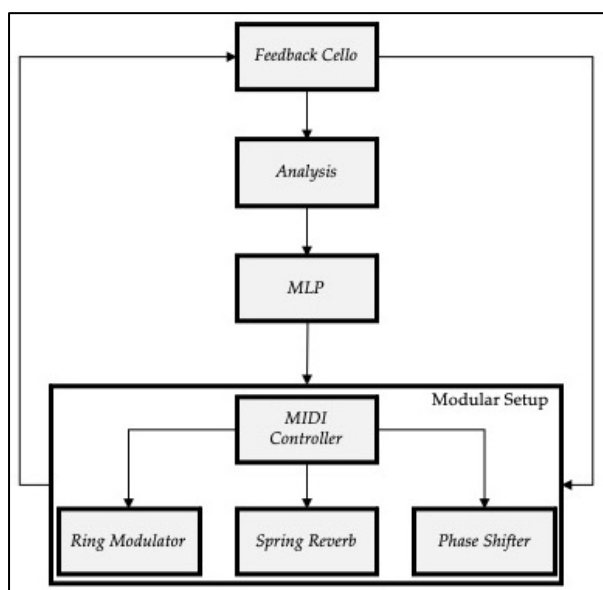


Figure 21: stage setup for Feedback Cello.

Eldridge and Kiefer presented a new configuration of their *Feedback Cello*. The setup is to be conceived of as a *shared instrument*⁵¹. Here the cello is emitting its signal and being sent to the laptop where Kiefer is running a SuperCollider patch. This patch is a controller for his modular setup.

The patch is running the FluCoMa Multilayer Perceptron (MLP) regressor and taking the real-time audio analysis of Mel-Frequency Cepstral Coefficient (MFCC) (Peeters and Deruty, 2010) content from the cello output and using it to control parameters of different modules which are modulating the signal of

the cello. The three main processing modules are a spring reverberation, a ring modulator, and a phase shifter. There is also a set of three MIDI controllers which are controlling the amplitude and low pass filtering of the cello input into these modules. Sound is then fed back towards the cello. Eldridge also had two-foot pedals – although only one was used during the performance. This controlled gain of the cello output.

⁵¹ Appendix 8.6.1.

Musical thought.

- *Shared instrument:* the setup is really meant to be conceived of as a shared instrument – the feedback loop is such that one element cannot be disentangled from the other.
- *Embodied knowledge and live coding:* Kiefer explained the rehearsal and development process⁵². He explained how he gained an intuitive, embodied knowledge about the MLP’s internal structure, and how to get it to do the things he wanted. He explained an iterative process where for several months they had weekly sessions, where Kiefer would live code the MLP and perform experiments.
- *Cello gesture as controller:* Eldridge explained⁵³ that she conceived of her gestures on the cello across two domains: the usual sonorous domain, but also that of control data. She explained that it was necessary to build-up a certain gestural vocabulary around this.
- *Analogue system:* most of what is occurring is in the analogue domain – both musickers confirmed that this was a decision that was taken early on⁵⁴. The computer and the software are used as a controller for the modular setup – no synthesis or processing occurs within it.

f) *Sam Pluta.*

Pluta’s performances were pre-recorded. Each artist recorded each of their parts separately, the first piece Pluta recorded first, and Evans recorded on top (Pluta informed me that Evans found it challenging to play against, explaining why Evans chose to take one broad musical idea and stick with it⁵⁵), and the second piece worked the other way round. Evans was not actually aware that Pluta was not going to be doing any processing on his trumpet⁵⁶.

The stage setup, then, includes the interface that they used to share each other’s

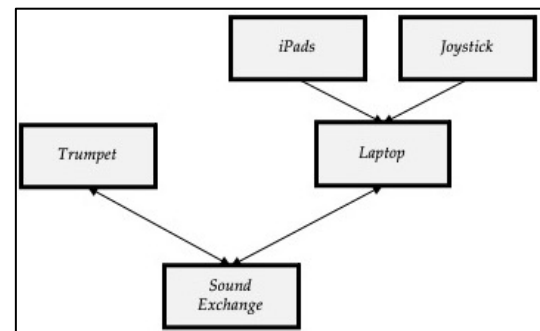


Figure 22: setup for Pluta’s Neural Duo pieces.

⁵² Appendix 8.6.3.

⁵³ Appendix 8.6.2.

⁵⁴ Appendix 8.6.4.

⁵⁵ Appendix 8.5.2.

⁵⁶ Appendix 8.5.1.

recordings. Pluta also has two principal controllers which are controlling his software: 2 iPads running Lemur⁵⁷ patches and a gaming joystick.

Pluta used the software that he has been building up over a large part of his career (Pluta, 2012) and integrated the FluCoMa implementation of the MLP regressor for multidimensional parametric control in a way that he now describes as being an “integral part”⁵⁸ of his software. He uses it to create the mappings between his joystick and parametric spaces of the various synthesizers he uses – many of which, such as the FM7 emulator and cross feedback synths, have a very large number of parameters.

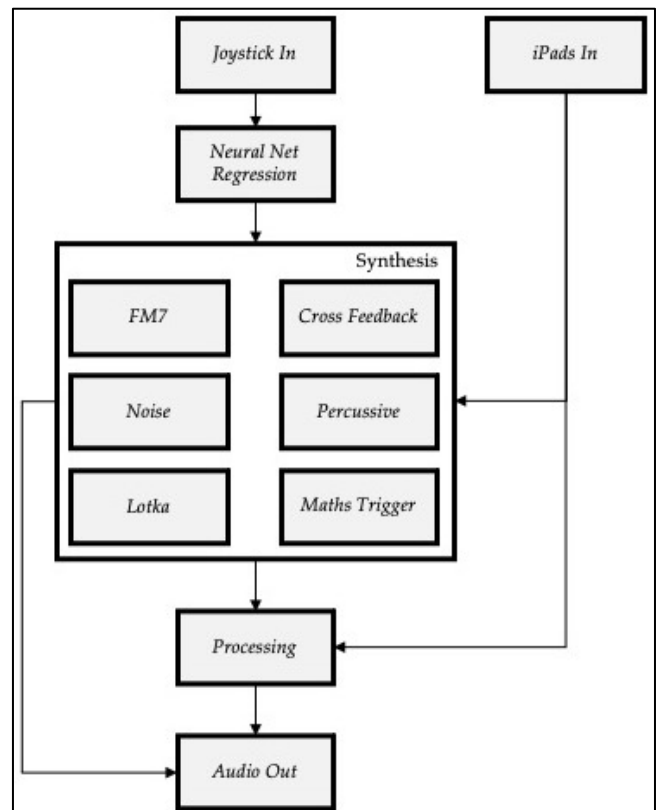


Figure 23: diagram of Pluta's performance software.

Pluta doesn't work with samples: he explained that he found it hard to get them do what he wanted, harder than with synthesis⁵⁹. The sounds from these synthesizers are then either directly routed out or go through the processing part of the software. This uses an interface that he has built up over the years that allows him to build patches of various effects modules and save them all so that he can access them easily during a piece.

Musical thought.

- *Synthesis over samples*: despite trying to work with various interfaces⁶⁰ that would allow him to manipulate samples, Pluta prefers to use sound synthesis.
- *Modular setup*: much of the software's functionality is similar to an analogue modular synthesiser setup, with an interface that allows him to route

⁵⁷ <https://liine.net/en/products/lemur/>

⁵⁸ Appendix 8.5.3.

⁵⁹ Appendix 8.5.4.

⁶⁰ Notably using the FluCoMa tools for dimensionality reduction.

outputs to various places. The notable difference is that his interface allows him to switch instantly to other patch configurations, and he uses this to structure his playing.

- *Source and processing*: Pluta draws a clear difference between the synthesizers as a sound source and the processing which occurs in a second conceptual moment. This may stem from his history of working with other musicians. Indeed, he explained that it was only since the pandemic and having to isolate that he forged a vocabulary for playing alone.
- *Duo*: this is a very particular setup for a duo, in as much as the two musicians did not play together. Pluta commented that he felt that the first performance was convincing regarding the feel of them playing together, but this was less the case for the second.

g) *Alexander Harker.*

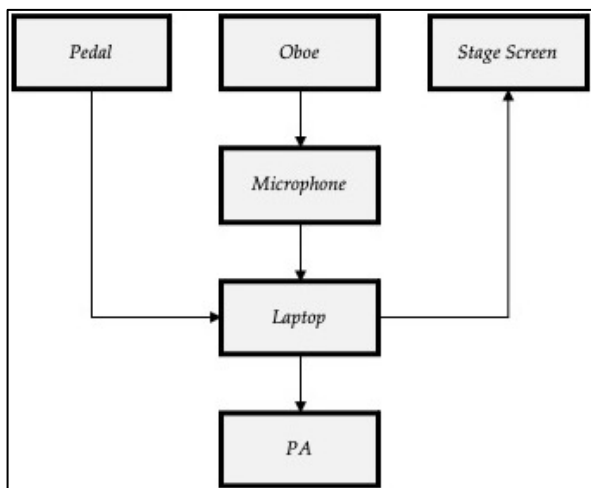


Figure 24: stage setup diagram for Harker's piece *Drift Shadow*.

The stage configuration for Harker's piece is traditional for instrument and live electronics. The oboe is simply playing into a microphone, the sound of which is sent to a laptop. The laptop then sends audio to the PA system, as well as feeding image back to a screen on stage providing Dell with information about position in the piece. There is also a pedal on stage which allows Dell to have some control over the patch.

Harker's software is very extensive, counting over 250 different patches. We can, however, draw an

abstract image about what is happening. Essentially, the patch can be conceived of in two parts: tracking of what the oboe is doing and calculating where that means the player is in the score; and consequent processing of the oboe sound.

Tracking of the oboe consists of determining which multiphonic the performer is currently playing. To do this, Harker pre-recorded all the multiphonics that the oboe would be playing during the piece, and trained a 120 hidden neuron, single layer neural network. Harker used his Framelib library to perform this task frame by frame, essentially getting as fast a response as possible. He applied several filters on the incoming audio to avoid issues that could arise with different microphones (although

Harker did mention⁶¹ that Dell had to use the same type of reed that was used to record the training data).

He also implemented a number of systems to guarantee correct multiphonic detection: first, he modified the FluCoMa Max MLP implementation so that he could assess the energetic values of the output neurons in order to threshold the result and get its ratio compared to the other neurons; he also hard-coded fixed rules dictating that certain detections were impossible compared to the current position in the piece, and generally tried not to repeat multiphonics over the piece; finally, he allowed Dell to have some control over the patch with the foot pedal – she could lock the piece in a section and disallow change as she pleased. This information is then passed on to the score following, where each section of the piece has its own patch that will activate the various DSP modules.

The three spectral freeze, partial freeze, and granular freeze modules are described in Appendix 3.3.

Harker also used some of the FluCoMa software when writing the score for exploring the scope of multiphonic possibilities on the oboe. He departed from using Peter

Veale's (Veale, 1994) multiphonic charts, but found them difficult to work with. He used some of the Uniform Manifold Approximation and Projection (UMAP) (Dorrity et al., 2020) visual mapping techniques offered by FluCoMa in two ways: mapping all of the recorded multiphonics in a space in order to find ones which were sonically similar; and looking at the results of combinations of different multiphonics. It is interesting to see here that Harker is engaging with the oboe

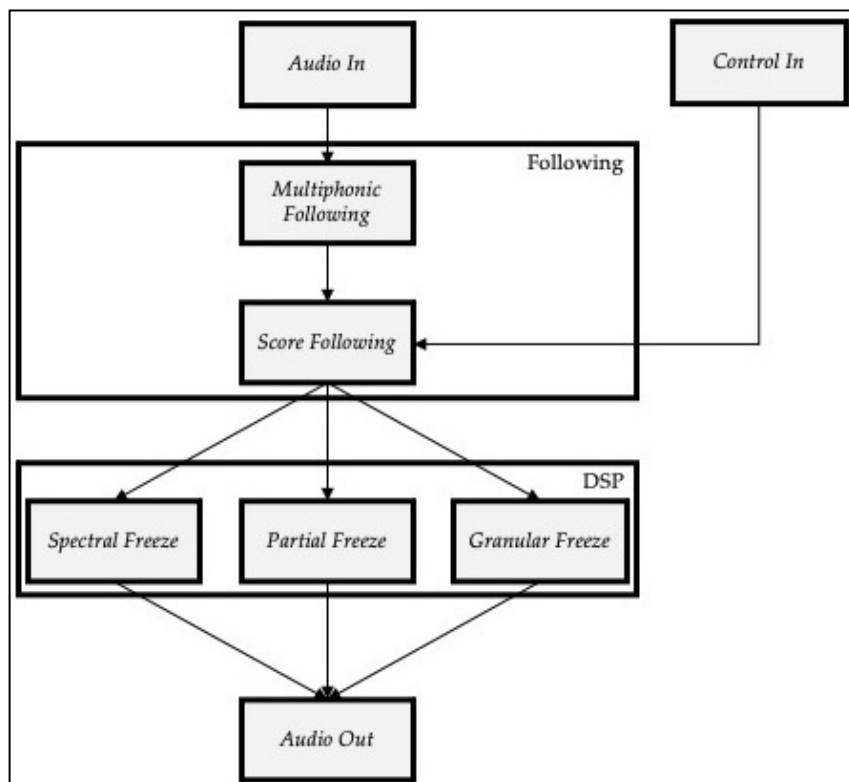


Figure 25: abstract diagram of Harker's Drift Shadow patch.

⁶¹ Appendix 8.8.3.

multiphonic corpus not just as sound, but also interwoven with the fingerings that produce them.

Musical thought.

- *Frame by frame analysis:* Harker is searching to achieve extreme rapidity in processing. This is also seen in his transformation of certain parts of the patch into gen code – optimisation is something that is very important in his practice.
- *Tool modification:* Harker is the only one of the musickers that decided to go into the FluCoMa codebase and recompile the objects in a way that served the uses he wished to make of them.
- *Technical limitations form the composition:* Harker discussed⁶² the fact that repeating multiphonics over the course of the piece would make it more difficult for the tracker to function, and that he composed the piece according to this limitation. This demonstrates a certain idea of how the software conceptually played a role and came first in the aesthetic conception of the piece. He also discussed⁶³ initially wanting to perhaps have more melodic content from the oboe, but this was also disregarded due to the nature of the tools.
- *Freeze:* the major sound processes present are ones that take an analysis of the incoming sound and prolong it in time in various ways.

h) Hans Tutschku.

In his score, Tutschku provides us with a diagram depicting the stage setup for the performance. It is a typical setup for piano and electronics: there are eight speakers surrounding the audience: the front two are routed directly from the two microphones that are used to pick up the sound of the piano. The piano sound is also fed into the laptop which is running Tutschku's Max patch, and passes by a sound card which outputs 8 channels of audio. There is also a sustain pedal which is routed to a MIDI interface which goes into the computer, allowing the pianist to control event changes. On stage there is a monitor for sound feedback and the piano itself which at the venue was a Steinway & Sons Grand Piano.

⁶² Appendix 8.8.2.

⁶³ Appendix 8.8.1.

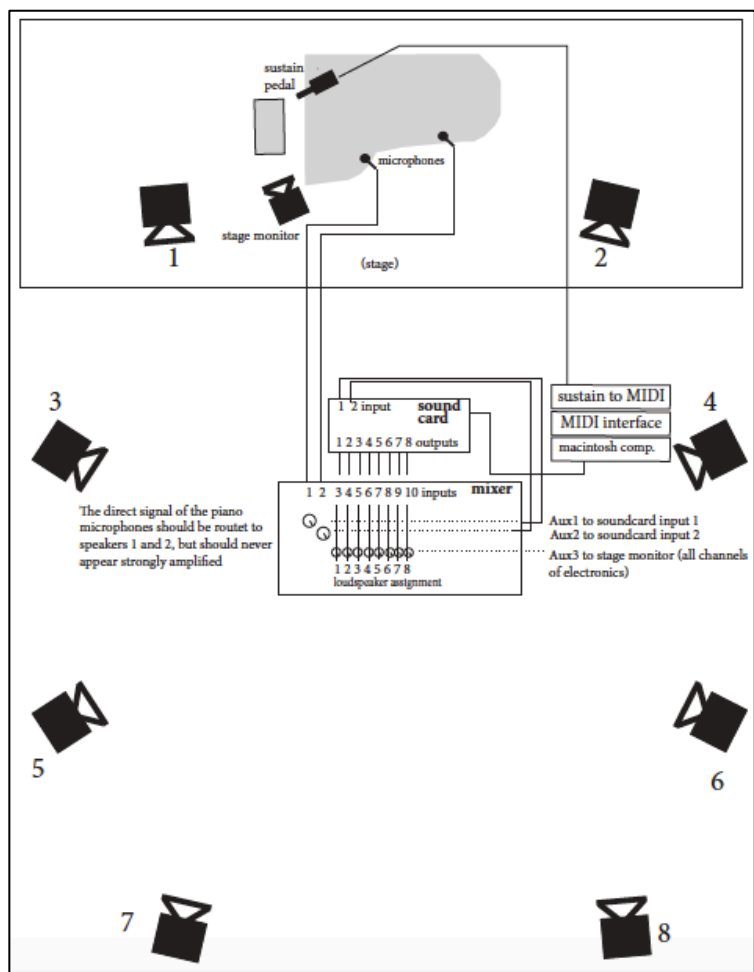


Figure 26: stage setup diagram given in score for Sparks.

In Tutschku's software, audio from the piano is received and followed by attack detection (using simple amplitude detection) and pitch detection (see below). The control data from the sustain pedal is coming in to change events⁶⁴. At each event, beneath the score, several indications are given which describe the setup of the processing modules. At the beginning of the score, Tutschku gives a screenshot which indicates all the different processing modules in the patch. Each event corresponds to a different configuration of them.

There are three different types of effects modules: the first are operating in the temporal and energetic dimensions, and correspond to standard effects such as reverberation, flanger, chorus and spectral freeze. The second are sample

playback modules. These are playing back a number of files from a collection of prepared piano sounds and come in various forms: notably samplers and granular synthesis. The final type of module is those that tend to operate in the organisational dimension: for example, the *slicer* module or the *stubborn delay* module. Essentially, when attack and pitch tracking, Tutschku transcribes this data into symbolic notation using the Bach score writing tools, and then performs transformations – for example inversions, prolongations, or drunk walks between notes – in what could be considered the organisational dimension.

There are moments when drastic changes in configuration can be heard, but Tutschku explained⁶⁵ that he wanted to try and avoid sharp cuts when possible. Therefore, the configuration of modules from event to event can see modules from the previous event staying the same. Each module is independent, and therefore can run on while

⁶⁴ The different events are marked on the score above the stave with an Arabic numeral in a square box.

⁶⁵ Appendix 8.9.2.

others change. In terms of modulation within an event, Tutschku explained⁶⁶ that he defined weighted randomness, giving a minimum and maximum range for parameters to move around in, usually in a glissando (these values can also be set to be the same to fix the value).

Another mode of modulation comes thanks to the NMF pitch-tracker. Indeed, this is where Tutschku makes use of the FluCoMa NMF tool. Essentially, using a corpus of 6000 piano sounds collected from the Pianoteq⁶⁷ VST iterating through four different presets, three different intensities and all the notes on the piano except for the first octave, he obtained an NMF activation model which allowed him to run the

audio from the piano through the FluCoMa NMF match tool, and allowed for surprisingly good pitch detection. The lower the pitch, the less reliable it is, and beyond 4 simultaneous notes it also becomes less reliable, picking up various overtones of the notes and confusing them with what is played. However, this is something he used as a compositional tool.

Indeed, he spent many months developing this tool, and came to know its limits very well. If he knew that the tracking would be uncertain, there were several different things he could do. First, he knew that any confused tones that were tracked from overtones were going to be harmonious in some way with what had been played, and therefore gave him access to useful randomness; if he wanted to insert a certain amount of randomness in the electronics, he could use this limit to his advantage. If the tracking of a certain pitch at a certain moment really was important, he would just hard code the specific detection at that point by using the attack detector to calculate where the pianist was in the piece.

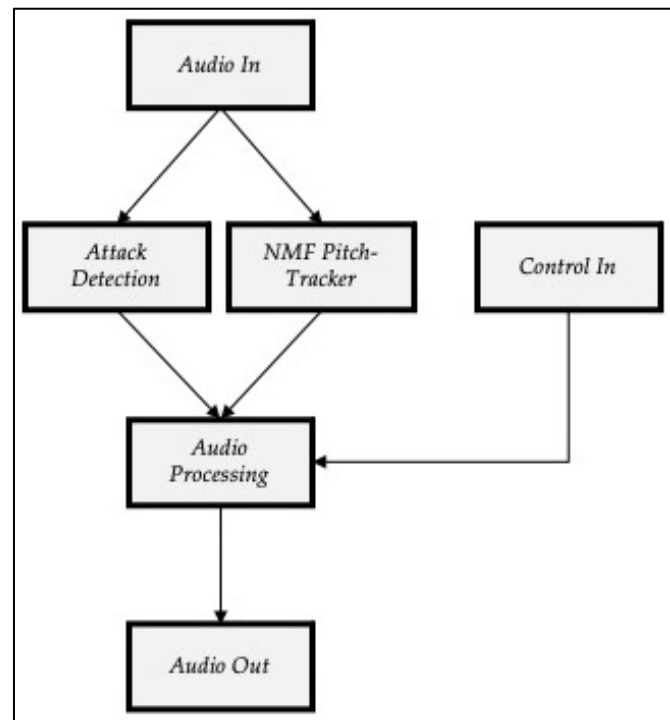


Figure 27: patch diagram for Hans Tutschku's Sparks.

⁶⁶ Appendix 8.9.1.

⁶⁷ <https://www.modartt.com/pianoteq>

Musical thought.

- *Spatialisation*: the sound is spatialised across eight channels. Tutschku explained⁶⁸ that spatialisation wasn't a structuring parameter for his work in the same way as, for example, someone like Stockhausen. Rather than inscribing specific movements and trajectories, he would think of spatialisation in terms of *qualities* of the space: nervous, agitated, slow swinging or lop-sided. This means that specific trajectories are not lost when not sitting in the sweet spot.
- *Modular processing*: Tutschku thinks of processing in a way that is very similar to Pluta. As the piece progresses, he iterates through a series of pre-defined configurations that are taken from a pool of possibilities. This process implied an inherent segmentation of the piece that organises itself by these events, by a series of different sonic fabrics.
- *Randomness*: we see a particularly controlled use of randomness in various forms: random parameter changing set by upper and lower limits and use of the knowledge of a tool and its tendency to yield random or unexpected results in certain situations.
- *Organisational processing*: it is extremely interesting to observe the modules that Tutschku has constructed that seem to operate in what we could consider as being the organisational dimension, making use of symbolic notation. This is an approach to sound processing that is relatively unique.

i) *Richard Devine.*

Devine performed with an extensive modular synthesis system. In the diagram, I have given the main structuring modules that were used to build the patch. There were also several acoustic percussion instruments being triggered by the modular setup.

⁶⁸ Appendix 8.9.3.

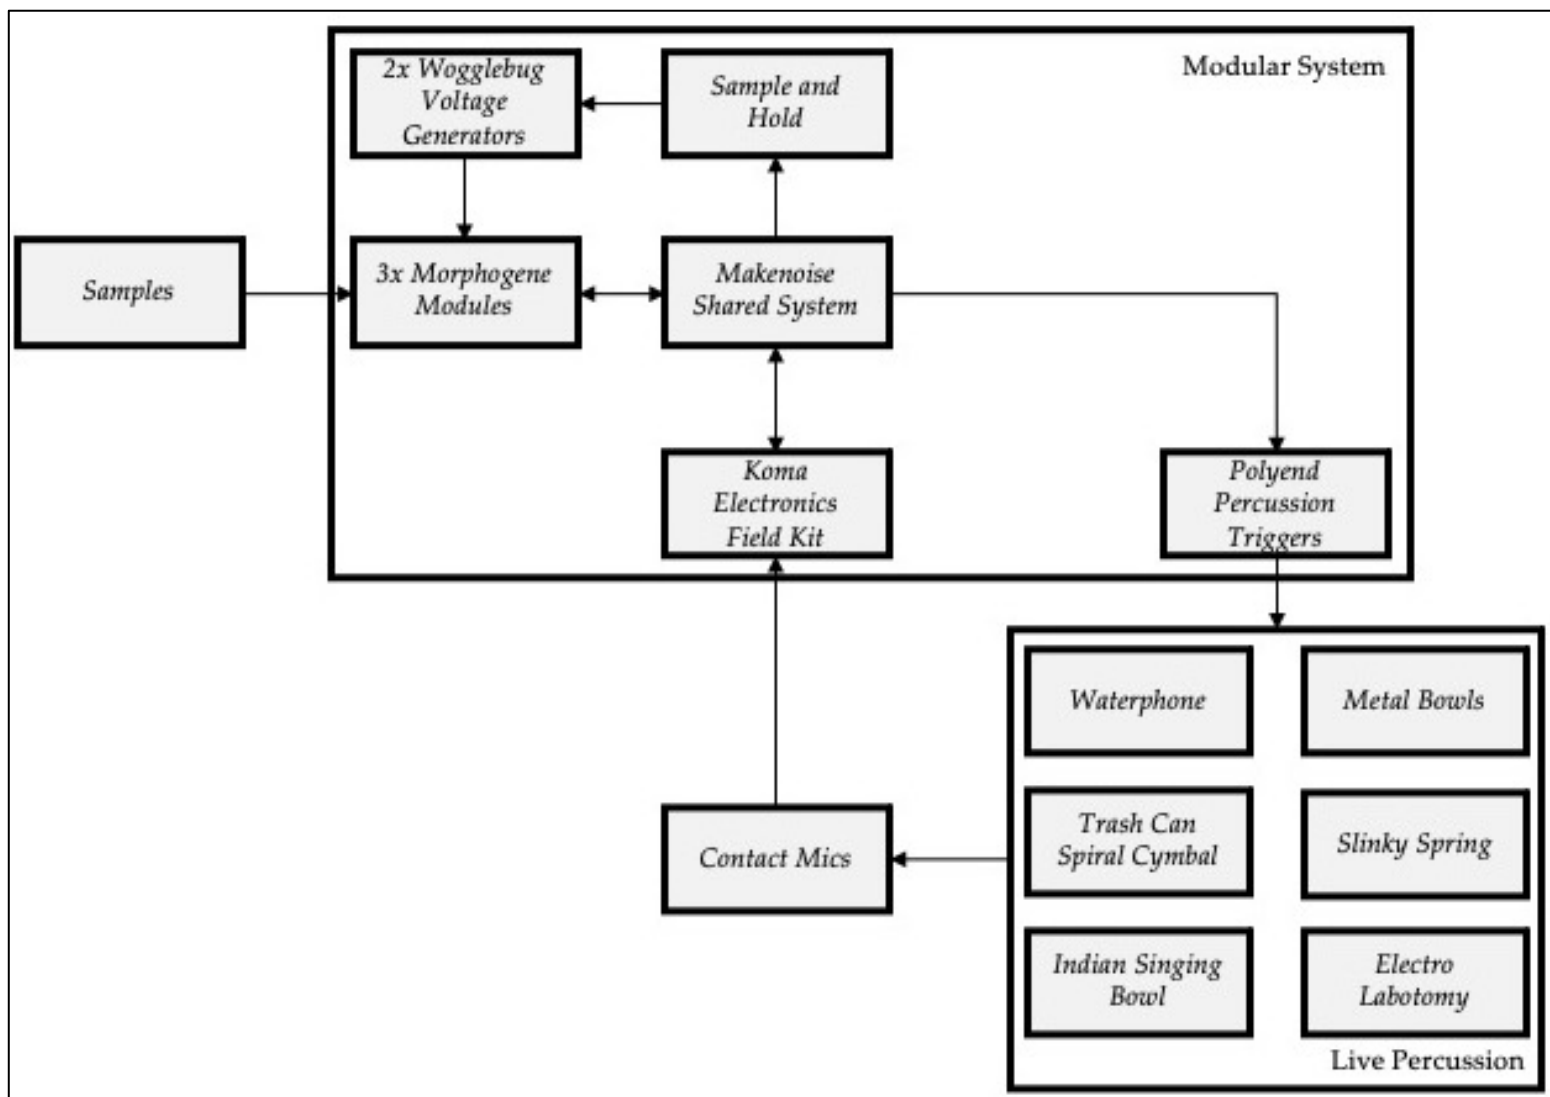


Figure 28: Devine's instrumental setup.

The heart of the system is the Makenoise Shared System⁶⁹ which is where all the processing is happening. Events are triggered by a René 3D cartesian-based sequencer⁷⁰ paired with a linear clock, and notably 2 Wogglebug random voltage generator⁷¹ modules. Devine explained⁷² that it was through interaction with the 2 Wogglebug's that most of the gestural content of the piece was achieved. The *rate* knob allowed him to go from a completely frozen clock up to audio rate, allowing him to create gestures. He discussed how important it was for him to “*break out of the steady*

⁶⁹ <https://www.makenoisemusic.com/synthesizers/black-and-gold-shared-system-plus>
⁷⁰ <https://www.makenoisemusic.com/modules/rene>
⁷¹ <https://www.makenoisemusic.com/modules/richter-wogglebug>
⁷² Appendix 8.7.3.

*pulse*⁷³ that is so common with modular synthesis. The *speed* input of the Wogglebugs is also being controlled by a sample and hold module to give further randomness to the timing of events.

There are two main sound sources in the setup: first, the live percussion. These are being controlled by the Polyend percussion trigger module⁷⁴ and are being picked up by contact microphones and fed back into the system using the Koma Electronics Field Kit⁷⁵ module. The other main sound source is coming from the 3 Morphagene⁷⁶ modules. These are modules upon which the user can load ‘reels’ of up to 3 minutes in length and define slice points. Devine explained that he spent much time constructing the reels (see below), and that he purposefully configured them so that various sounds would be occurring at the same time as others. He said that it was “*very planned out*”⁷⁷, and that having the samples and the main sound source of the patch on these modules allowed him to concentrate on the performance rather than micromanaging various synthesis elements. There was some synthesis going on during the performance; however, this was more to “*fill out the space*”⁷⁸.

The reels for the Morphagene modules were built in the computer – this is where Devine used the FluCoMa tools, notably James Bradbury’s Reaper⁷⁹ implementation ReaCoMa⁸⁰. The sources for the sounds were of three main types: field recordings, synthesized analogue gestures and synthesized and acoustic percussive elements (interestingly, some of these sources were recordings of the same percussion instruments used live during the piece). Devine essentially used the FluCoMa tools to perform slicing tasks, but also to decompose the sound using the HPSS and sine tools.

Musical thought.

- *Samples over synthesis*: Devine prefers to use samples and live sound sources, rather than building up synthesis within the patch, in order to concentrate on the performance⁸¹.

⁷³ Appendix 8.7.3.

⁷⁴ <https://polyend.com/perc-drumming-machine/>

⁷⁵ <https://koma-elektronik.com/?product=field-kit>

⁷⁶ <https://www.makenoisemusic.com/modules/morphagene>

⁷⁷ Appendix 8.7.2.

⁷⁸ Appendix 8.7.2.

⁷⁹ <https://www.reaper.fm/>

⁸⁰ <https://www.reacoma.xyz/>

⁸¹ Appendix 8.7.2.

- *Control of time for gesture*: the main way that Devine controls gestures is by controlling the rate of the various clocks within the patch. He described⁸² this system as allowing him to stop and pause on various sounds, stop the system completely, and move on through the piece as he wanted.
- *Control*: discussing⁸³ the setup of the system, Devine made clear that he worked for 4 to 5 weeks on the patch, desiring to have complete control and understanding about what was happening. Ultimately, he said that it is never exactly the same, and that he only recorded one take when filming, saying that whatever the piece was at that point, is what the piece would be.
- *Gesture over pulse*: Devine also made it clear at several moments that he wished to break away from the steady clock-like pulse that is often found in granular synthesis⁸⁴.
- *Digital Audio Workstation (DAW) over Max*: while building the samples for his performance, Devine opted for the use of Reaper rather than Max⁸⁵.
- *Sampled and live percussion*: the live percussive elements were also part of the recorded sample library. It seems that these elements can be considered as important structural sonic material for the piece. This is reflected in the generally percussive nature of the sound.

⁸² Appendix 8.7.3.

⁸³ Appendix 8.7.1.

⁸⁴ Appendix 8.7.3.

⁸⁵ Which he is nonetheless experienced and proficient in.

Chapter III. NETWORK CONSTRUCTION AND GRAMMATISATION.

1. Tool Development and Pilot Studies.

We now have a well-informed, broad picture of the musicking that occurred during the performances, and during development of the networks. Several possible strands of musical thought have already emerged, which will be useful in guiding our further analyses. Now, I move on to the construction and grammatisation of network visualisations with the goal of accessing a broader picture of all the actors involved: human, non-human, physical and conceptual. Returning to my initial theoretical discussion, I will attempt to construct these networks by interrogating their functions. There will be several different levels of abstraction to be considered. First, however, I explain how this part of the methodology was developed: I will present some of the tools that were used, created, and modified, and some of the pilot studies that drove development.

a) *Requirements and Goals.*

My goal is to represent networks in a digital space. To borrow terminology from network theory, networks are essentially comprised of two elements: *nodes* (or *vertexes*) and *edges*. In the context of this research, a node will represent a component of the musical network – this can be a simple component, or a composite component, made up of other nodes and edges (a sub-network). The notion of sub-networks leads us to modularisation, grammatisation, and making these groupings is an inherently *analytical, interpretative* activity.

It will be useful to be able to attach different classifications to nodes – in my analyses I use several categorisations such as *physical, biological, conceptual, and digital*. Edges represent relationships between nodes and should likewise be able to be categorised. In network theory, there is the possibility to conceive of *directed networks* with *directed edges*. This notion is not uninteresting for our uses, however, in practice it revealed itself as not bringing as much to the analysis as could be expected. Ultimately, the goal is to access a digital representation of a web of actors and to search for musical thought being inscribed within it – adding directed edges from node to node not only makes things needlessly complicated (relationships between nodes are often multi-directional) but can also lead the analyst to make contentious statements about relationships between actors. Indeed, it will be impossible to describe the full extent of relationships between two actors.

It will be necessary for these visualisations to be dynamic at various levels. I will want to be able to consider these networks at various levels of abstraction, various levels of modularisation. Also, I will want to be able to access different kinds of layouts of a same network. Indeed, some of the networks that I am attempting to visualise are abstracted from spatialisation – therefore, the question of how they are to be laid-out and viewed for exploration poses itself. There are a number of algorithms that address this very question and will offer themselves as interesting tools.

The digital form also lends itself to processes like filtering and querying: ways of viewing the information on screen that would reveal elements which are hidden or difficult to find. It will also be necessary to be able to manipulate these network graphs as agnostic data structures in the computer: there are several different ways I wish to approach their exploration, and within a number of different environments. This means that the way they are stored on the computer must be flexible enough to be exploited by a variety of programs. Similarly, a goal of this research will be to find a way of transcribing code, notably Max patches, to be exploited in network form – this technique must marry up with the broader network graphs under construction.

Once these network graphs are constructed there are several ways that can be envisioned for engaging with them in my musicological, organological context. First, in a static way, I will have a collection of visualisations of the setups seen in the previous chapter that is much more detailed, and that can be interacted with dynamically, allowing me to give a much more comprehensive archive of what occurred without being bound by the physical limitations of representation on a page. Second, is the possibility for dynamic clustering and grouping of parts of the network, building up a representation of how parts of the network are perceived by various actors. The idea is that these groupings will reflect how certain actors, notably the musickers who configured them, conceive of the network. This builds upon some of the ideas announced in Chapter 1, where for example, the inclusion of a hand or fingers could potentially be up for debate in its grouping or not with a piano key or a trumpet piston. For this reason, I initially construct the graphs as ‘unclustered’ and interrogate the modularisations progressively. This brings in to question the resolution of a network – how detailed must an initial construction be? This resolution could potentially be infinitely fine and will have to be addressed on a case-by-case basis.

Another way of engaging a network for musicological analysis would be through the idea of *degrees*: this is another term that is borrowed from network theory. This measures the number of edges that are connected to a node. We could expect to find nodes with a high number of degrees as perhaps holding a central position within the network, and perhaps being important points of interface.

Similarly, we could look to inspect the connectivity between nodes – through how many nodes does the path between two nodes pass through? This could suggest higher or lower relationships between actors. For example, imagine the connectivity between a foot and the sound that it eventually triggers. Perhaps if there is a high number of nodes between them, this could be translated into a perceptual disconnection between the body part and what it is doing. It is not argued that this argument could be immediately made, however, it will lead us to *ask* these questions and offer us concrete artefacts of representation to point to and discuss.

Another way could be to examine reoccurring actors across a network or compare networks with similar actors. Perhaps it will be possible to find similar configurations around similar actors, or not. In any case, this approach should be useful for building profiles of different actors and examining their ontology within different contexts. It could also be interesting to try and bring these visualisations back into the context of what they are representing. For example, it could be possible to link a network visualisation to a Max patch and have them interact with each other. We could envision a playback of a performance that would progressively build-up heat maps that would reveal zones of activity within the network.

The methods I propose and the tools I have developed address most of these ideas in some encouraging ways. This research has offered a precious opportunity to try these ideas out on real case studies, and a lot has been learnt. There are other ways of interrogating networks that stem from network theory: eccentricity, stress, betweenness, topological coefficient and radiality. It will be important not to get lost in the analysis of the aspect of a network to a point where I lose sight of reality, as per the previous example of a hypothesis of connectedness between two nodes and perceptual connectedness. It is important to not lose touch of the meaning or function of a network that could potentially become convoluted through either too much resolution or too much abstraction.

b) Tools and Pilot-Studies.

I started by looking at some existing tools and how I could incorporate them into a workflow. The two main network visualisation programs that appeared to me were Gephi¹ and Cytoscape. After some exploration, I judged Cytoscape to be the most robust and stable and decided to start working within its framework.

Cytoscape is a program that was originally built by several research groups; however, its ongoing funding and maintenance is handled by the National Institute of General Medical Sciences (NIGMS). It was originally intended to be used by biologists for genome analysis and other bioinformatics, and broader social network analysis. Aside

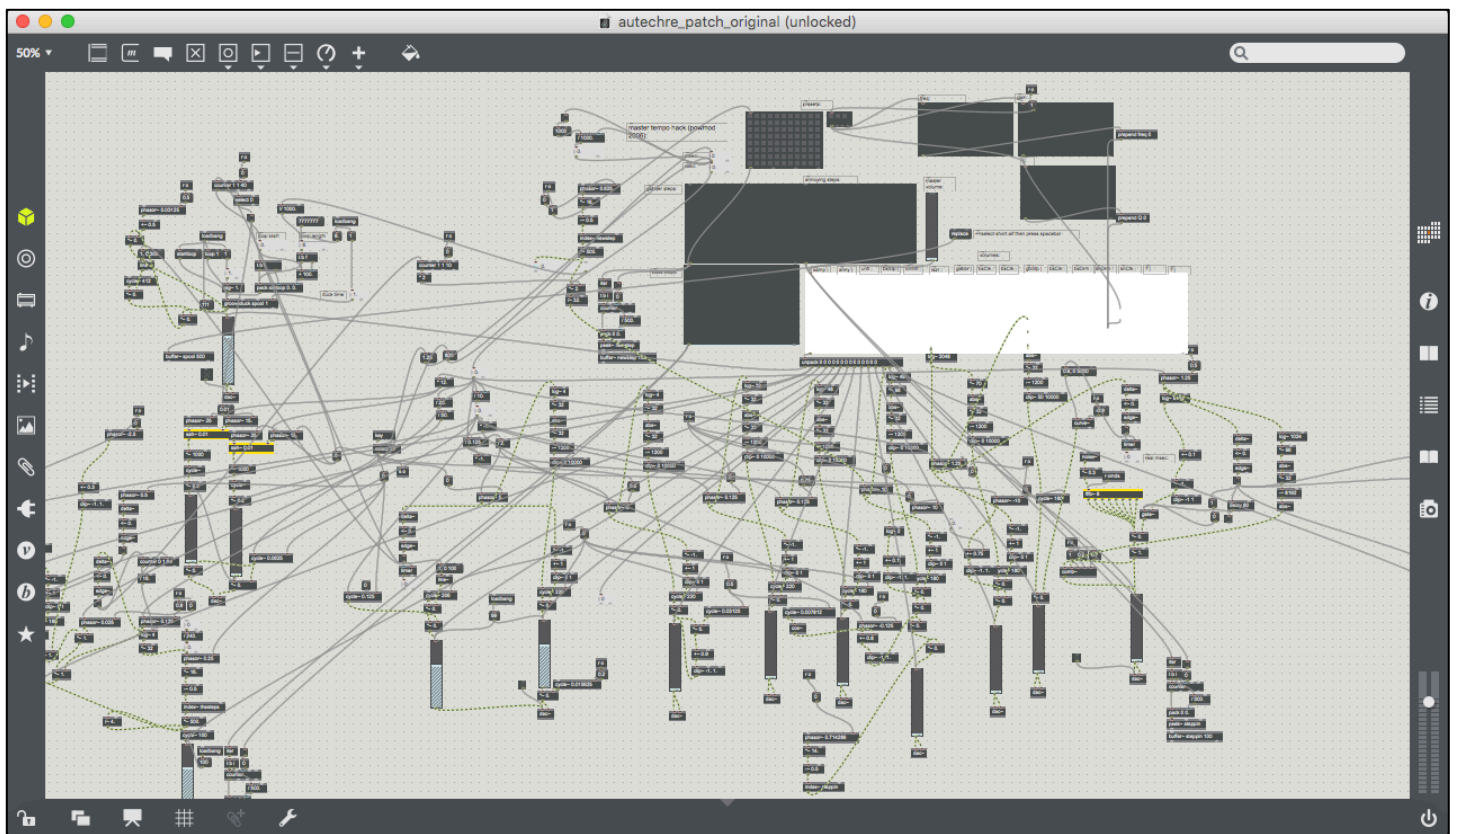


Figure 29: the Autechre test patch in its original state.

¹ <https://gephi.org/>

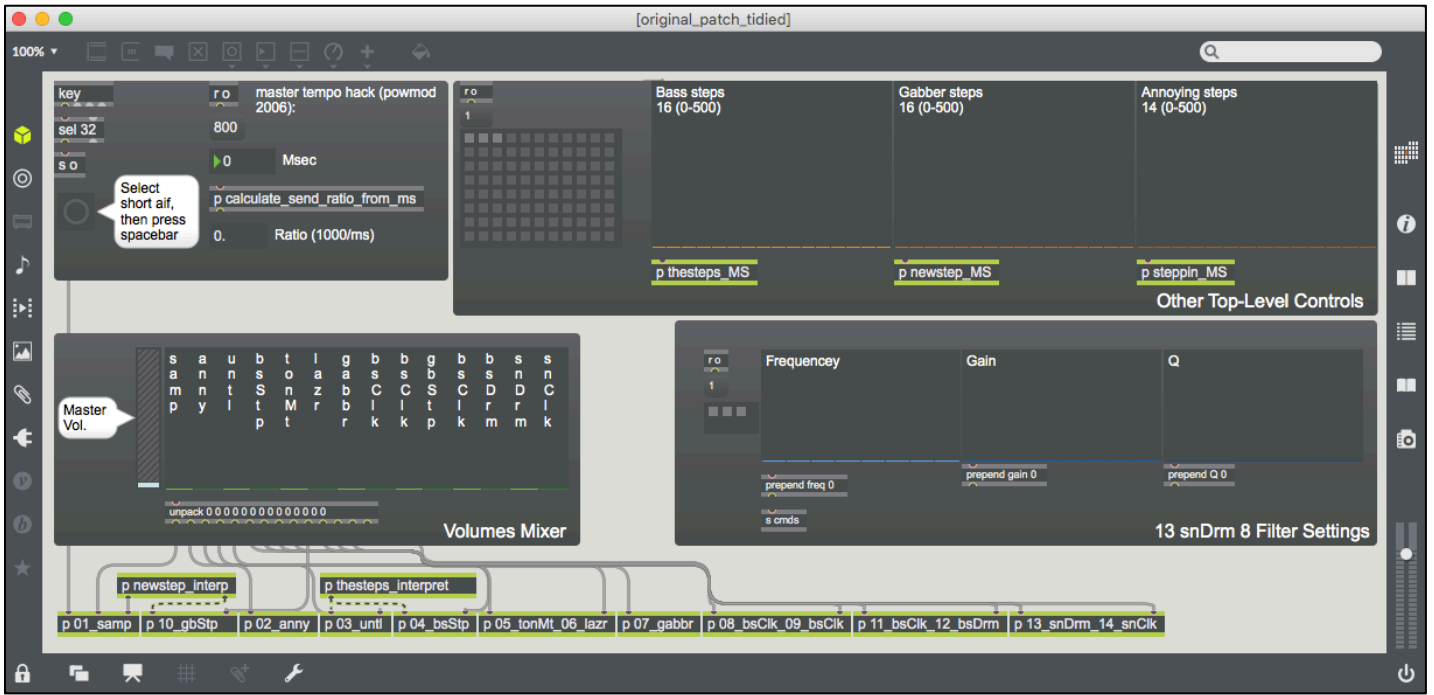


Figure 30: the Autechre test patch after analysis.

from the standalone desktop program, there is also an extensive library of *apps* from various sources that extend the program's functionality, supporting Java, and with the possibility of scripting in Python, Ruby and JavaScript. There is also a browser-based library version of the software which allows for network visualisation in the browser. I started by creating some test networks by hand – notably a network representing a trumpet. Despite the network being quite small, it was still a time-consuming process.

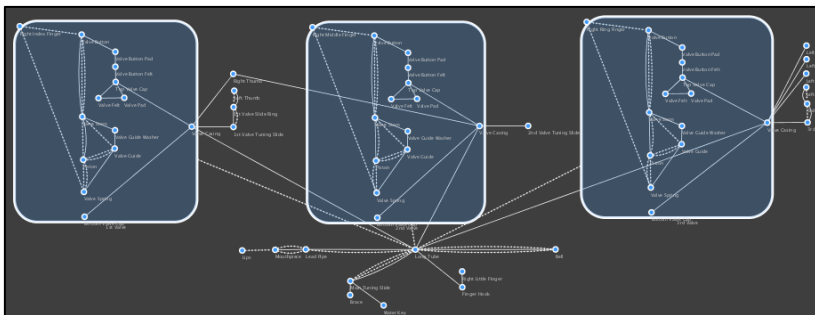


Figure 31: trumpet network visualisation.

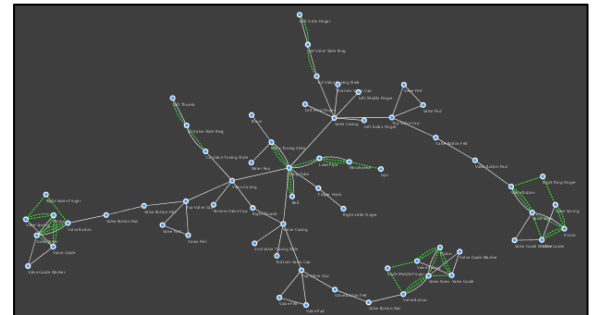


Figure 32: trumpet visualisation highlighting flow of energy.




Figure 33: edge-weighted spring-embedded layout of the Autechre test patch network.




Figure 34: prefuse force-directed layout of the Autechre test patch network.




Figure 35: compound spring embedder layout of the Autechre test patch network.




Figure 36: yFiles organic layout of the Autechre test patch network.

These initial tests didn't bode well considering the breadth of some of the networks I would be facing. Some positive points were the fact that the networks could be exported to several different formats, notably JSON which is a human-readable format, and a data structure that was present in many different computer programs. The functionalities in Cytoscape for clustering and creating sub-networks worked well. Also, even with the small trumpet network, some of the layout algorithms seemed very encouraging.




Figure 37: attribute circle layout for all nodes in the Autechre test patch network.

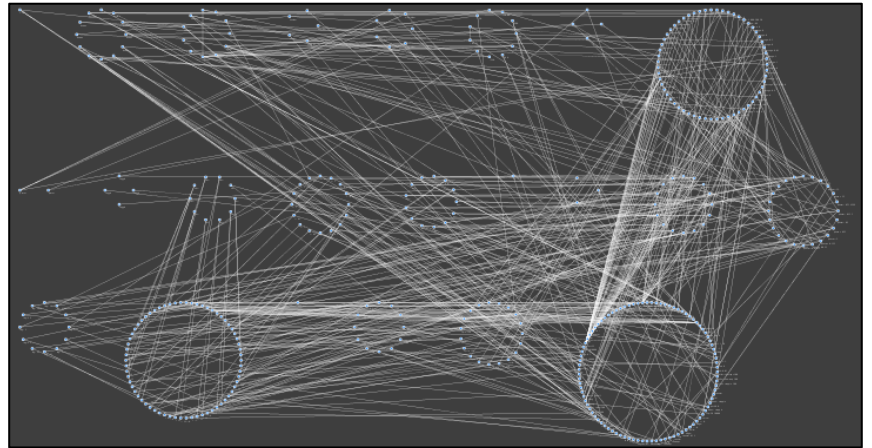


Figure 38: group attribute circles layout for the Autechre test patch network.

Soon, I wanted to examine how I would be able to transpose something like a complicated Max patch into a form that could be used by Cytoscape. To test this, I carried out a pilot study on a patch² from the Cycling 74 forums – known there as the *Autechre patch*³. It was a good candidate for testing out Max patch analysis.

The patch is messy. Like with the software made by the FluCoMa musickers, my first job was to do an initial surface level analysis of the patch, and a general tidy-up to get a good idea about what is going on. This is a necessary step, as the functionality of a network is something that we need to know, at least broadly, *before* we wish to represent it. The goal of a network visualisation is not to directly inform us of the functionality of a network – it is a tool that we can use to engage with the full extent of the actors and be used to generate other questions. It is the *exercise* of creating a network that will lead us to try and understand the various grammatisations, and the important materialities and points of interface inscribed within it. It allows us to gain insight on musical thought that may be inscribed within it. I do not mean to say that a collection of nodes and edges *is* a concrete representation of musical thought. This process will lead us to a conceptually materialised understanding of this *saying* of

² D.E.: 01_Code>01_Sources>10_Other

³ <https://cycling74.com/forums/autechre-patch/>




Figure 39: hierarchical layout of the Autechre test patch network.




Figure 40: yFiles tree layout of the Autechre test patch network.

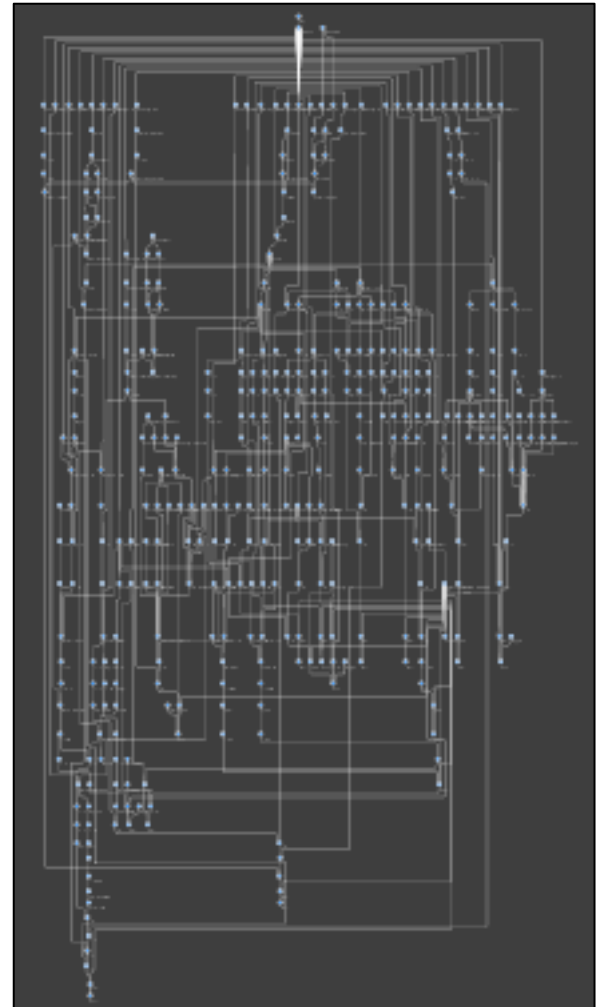


Figure 41: yFiles hierarchical layout of the Autechre test patch network.

music. It will help us discuss the materialities within which this musical thought is intertwined.

Hence, I performed my initial surface level analysis of the *Autechre patch*, and then sought to find a way of bringing it into Cytoscape. When opening a *.maxpat* file in a text editor, we discover that everything is stored in JSON format. I started working on a tool⁴ that would take this Max-specific format, and parse it into a more general, network format, which had two large collections: nodes and their properties; and edges and their properties. I made this in a way that would allow for the file to be opened in Cytoscape. The process has become streamlined: a demonstration of this process can be found in this video⁵.

There were various issues to overcome with the process of parsing a Max patch in this way: issues with hierarchy, purifying the network of things like subpatchers,

⁴ D.E.: 01_Code>02_Tools>01_Network_Construction

⁵ D.E.: 04_Video>01_Demos>Max_to_Cytoscape_Demo.mp4




Figure 42: *yFiles* radial layout of the Autechre test patch network.

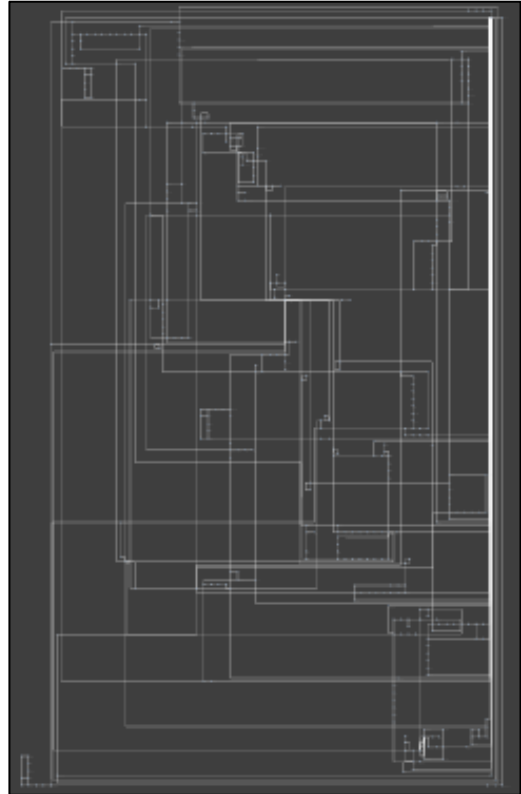


Figure 43: *yFiles* orthogonal layout of the Autechre test patch network.

abstractions, inlets and outlets, *poly~* objects and *pfft~* objects; and *sends* and *receives* in order to directly link up the nodes⁶.

With this first, complex case in Cytoscape I investigated some of its features. First, I wanted to examine the question of classification of nodes and edges mentioned above. For this, I automated the process in the parsing script: elements are automatically classified by things that could be useful (abstraction or subpatcher, parent patches, GUI elements and later FluCoMa objects, audio patchcords and Max patchcords).

Next, I explored some of the layout algorithms that are included with the software: some of these can be seen in the images in this section. There are three different types of layout algorithm: *force-directed* algorithms (Fruchterman and Reingold, 1991), which treat the network as a physical object, where the edges can be thought of as springs. In Cytoscape, we see several implementations of these: edge-weighted spring-embedded layout “nodes are treated like physical objects that repel each other [...] The connections between nodes are treated like metal springs attached to the pair of nodes. These

⁶ The system needs to be handled with care; however, it is at a point where this can be achieved without too much hardship. Sometimes it is necessary to rectify errors in Cytoscape which may have occurred during the parsing process.




Figure 44: colours showing the different subpatches of each object in the Autechre test patch network.

springs repel or attract their end points according to a force function. [It] sets the positions of the nodes in a way that minimises the sum of forces in the network" (Shannon et al., 2003); prefuse force-directed layout by Jeff Heer; compound spring embedder; or the *organic yFiles* (see below) layout. These algorithms tend to yield networks which branch out from central nodes that have high degrees of edges and could potentially be useful for inspecting the idea of centrality and important points of interface in a network.

Next there are *attribute-driven* layouts that organise the space according to the nature of the nodes. In Cytoscape they tend to be implemented as attribute circles: there are two types, one where all nodes are grouped into one circle, with nodes that are of the same type being adjacent to one another; and one where nodes of the same type are grouped into different circles. These could be useful for determining the quantities and proportions of different types of nodes, which could inform us about the nature of the musical network.

Finally, there are *hierarchical* layouts – these tend to inform on the internal structure of the network and could be good candidates for examining the paths between different nodes. There is a sub-library of layout algorithms called the *yFiles* library that produce some visually pleasing layouts and seem to be exceptionally effective regarding hierarchic layouts. This can yield layouts such as in Figure 44 which recall some of the *synoptic visualisations* proposed by Dufeu for presenting a patch's structure (Dufeu, 2013).

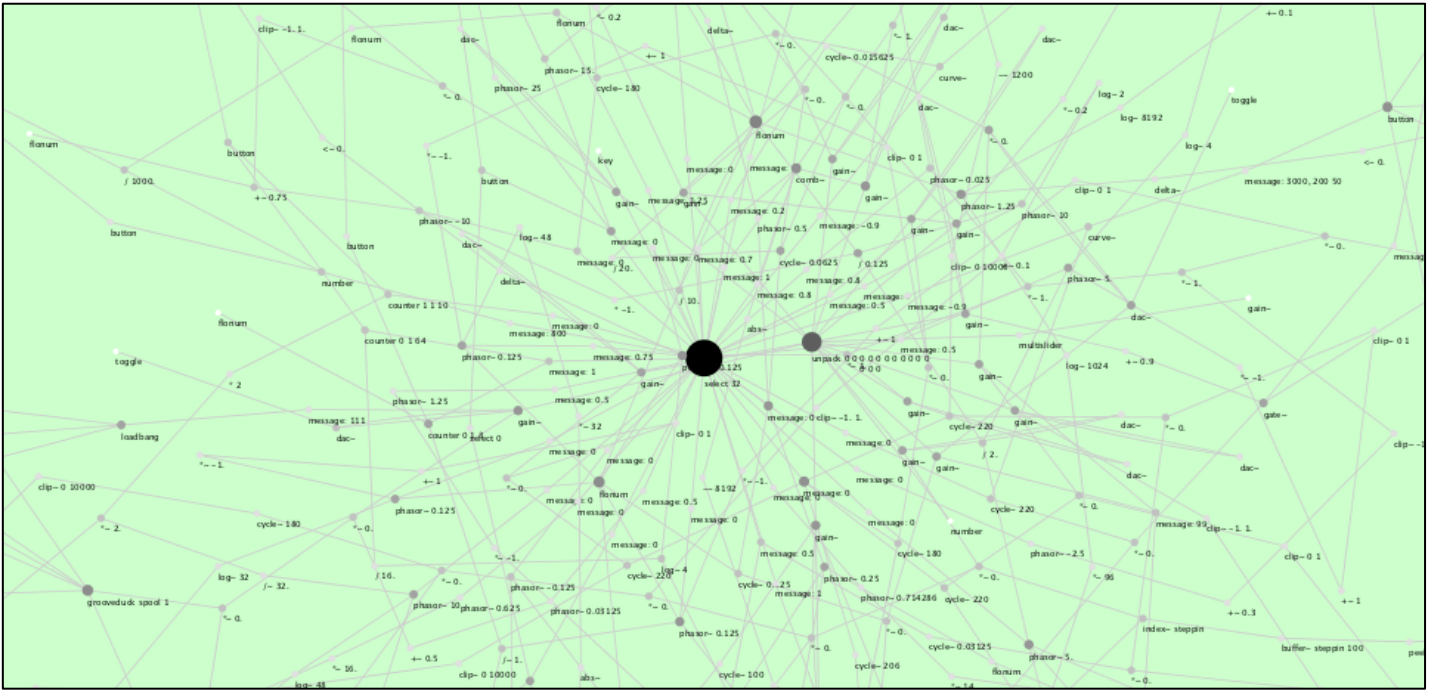


Figure 45: node size showing the degree of edges for each node in the Autechre test patch network.

I then experimented with various configurations of the *network style* tab, seeing which configurations could best highlight various things one would wish to highlight in a network. I also experimented with some of the filtering and querying options, as can be seen in some of the figures. I tired showing the audio flow within the network by highlighting all nodes with the `~` character in their name, and all edges with the type




Figure 46: filtering nodes by degrees in the Autechre test patch network. From left to right, in yellow: nodes with at least 2 degrees, 3 degrees and four degrees.

*continuous*⁷. I also tried filters such as the *degree filter* which revealed, for example, that the busiest nodes in a force-directed layout tend to be concentrated towards the centre but are much more evenly distributed across a hierarchic layout.

I also tried some of the 3rd party tools, such as Pesca3.0 which allows visualisation of the shortest path from one node to another, or the average path length from a node to all other nodes, indicating its centrality⁸.

At this point, I was satisfied that this software would be good for the type of analysis I wished to perform. There are two final points that I do wish to address. The first is the possibility of bringing these networks back into Max and having them *dynamically* linked in some way. I started building quite a substantial Max patch⁹ built mainly in JavaScript designed to open these graphs, and modify them directly in Max¹⁰. Ultimately, the size of the networks that I was intending to bring in was too much to be handled by my JavaScript interface.

Finally, there is the idea of visualisation of these networks in a form that would reveal the entire scope of sonorous potentialities of the network. This is the subject of Chapter 4 and was subject to a whole strand of research and development.

2. Global Level and Local Level Networks.

I shall now present the network graphs created for my analysis. As discussed in Chapter 1, I look at various levels of abstraction: global level, local level and in Section 3 the instrumental level. The actors in these networks have been categorised into four main types: *digital*, actors which exist within the digital domain (computer programs, digital audio signals, sound files, Max objects); *physical*, actors that exist in the physical world (musical instruments, stage gear, places); *biological*, actors that are comprised of biological matter (humans, body parts, animals); and *conceptual*, actors that don't have any physical existence and exist as ideas (philosophical concepts, groups of people). These categories can be porous, and a single actor may belong to more than one. This

⁷ The classification I gave to signal patchcords, opposed to *periodic* for Max message cords.

⁸ The most central node in the *Autechre* patch was a *select 32* object with an average path length to all nodes of 4.11 nodes. This object triggers when the space bar is pressed: it is the key that stops and starts processing.

⁹ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/05_maxNetwork

¹⁰ I had even started to implement some of the layout algorithms found in Cytoscape, the goal being to emulate many of the functionalities of this program in Max.

categorisation is neither exhaustive nor rigorously scientific: it helps me conceptualise and organise my networks.

As discussed in Section 1, the categorisation of edges, and notably *directed* and *non-directed* edges can be contentious. Sometimes, the distinction is easy to make: for example, the edge between two Max objects can be conceived of as directed as there is a clear logical flow between the two¹¹. In broader contexts, this question becomes more complicated: take the edge between the biological actor of a musicker and the conceptual actor of their piece. It can be argued that the direction of this edge is inherently unstable – each constantly imparting agency upon the other. Indeed, situations where the direction of an edge is unclear can be candidates for interesting discussion. I address these questions as they arise, but it is taken that in these visualisations, edges are not represented as directed. They suggest some connection of agency, the nature of which in most cases is evident by reading the nature of the source and target nodes¹².

These networks can be found in the digital elements¹³ as Cytoscape files and web interfaces. There is also a video¹⁴ demonstrating how to operate them.

a) *Global level.*

As explained in Chapter 1, the bounding of a network will be suggested by the process, or processes, it appears to make possible. Its function gives it meaning. For my first, global-level network¹⁵, I interrogate the general functionality of the FluCoMa project. There are three main strands, three main processes that I argue justify the delimitation of this network: the production of digital tools for fluid corpus manipulation; the creation of musical networks by a group of musickers that would use these tools; and the imparting of energy within these networks, concentrated around two moments, two performances, two concerts.

This is the level of abstraction that will produce a network similar to the example of the New York Jazz Composer's Guild, founded in 1964, imparting agency upon Chicago's Association for the Advancement of Creative Musicians (AACM) – founded

¹¹ Although it could be argued that the target object imparts agency upon the source by its very existence being the condition for the source's existence.

¹² Even if, consequentially, these are bad terms to use.

¹³ D.E.: 02_Data>01_Cytoscape_Networks

¹⁴ D.E.: 04_Videos>01_Demos>Cytoscape_web_demo.mp4

¹⁵ D.E.: 02_Data>01_Cytoscape_Networks>02_Global

in 1965, after the Jazz Composer's guild was disbanded – and thus musical practice in Chicago, drawn by Piekut (Piekut, 2014). Indeed, with this network, I observe translations of agency across various types of actors at a characteristically institutional level.

An image of this network can be found below, using the prefuse force-directed layout¹⁶. The network was built by hand, and the very process of its construction informed me greatly on some points of the methodology. I experienced first-hand the extent to which an atomic method of this kind can lead one down never-ending paths of nodes that spread further and further from the main body of the network.

Take, for example, the biological node *Sam Pluta*. In this network, it has eight degrees: *FluCoMa Project Commissioned Composers*, *Creative Coding Laboratory (CCL) Symposium* at which he has presented, *Electric Spring* at which he has performed, *Ted Moore* for whom he was the research supervisor, *More is More* which is a musical label he has released music on, the *University of Chicago* where he is based, *Wet Ink Ensemble* for whom he has composed and finally *Neural Duo I* and *Neural Duo II* which are the two pieces he created for the FluCoMa Dialogues gig. I hesitated for a long time with a large number of other nodes that could have been included in the network: *New York Philharmonic*, *International Contemporary Ensemble*, *Grossman Ensemble*, *Yarn/Wire* which are all ensembles he has written for; *Columbia University*, *University of Birmingham*, *University of Texas*, which are places he studied; *George Lewis*, *Brad Garton*, *Tristan Murail*, *Fabien Lévy*, all people who have taught him; *Evan Parker*, *Ikue Mori*, *Ingrid Laubrock*, *Anne La Berge*; all people with whom he has collaborated in music. The list could conceivably go on endlessly.

Why, then, are certain nodes included and others not? For example, it can at first seem incidental that Pluta composed and performed with the *Wet Ink Ensemble* in the context of this network – however, the links do become apparent when looking at all of *Wet Ink Ensemble's* neighbours: *Pierre Alexandre Tremblay* who has also composed music for them and is the leader of the FluCoMa project; *HCMF* where they have played and which hosted the first FluCoMa concert; *More is More* where *Wet Ink Ensemble* have

¹⁶ It was necessary to make a few modifications to the default parameters of the algorithm to make the nodes visible, I notably changed the number of iterations to 1000 and the default spring length to 150.

was of course the second half of Pluta's duo for his FluCoMa piece.

Essentially, the nodes that are to be included in a network are understood as the analysis progresses, and there is a time-consuming process of adding one node, then another and examining the links between the two, then adding another and examining the links between the three and so on. Through my research and practice of this technique, I have come to consider two basic rules regarding the inclusion of nodes within a network: first, if the node has at least two neighbours, then there is a chance that it will be significant within the network; second, the node must conceivably impart an important amount of agency upon its neighbour. Of course, this process is subject to human error – there are undoubtedly nodes that are missing from this network. However, I believe that I have managed to draw a satisfactory visualisation of the web of actors that comprise this network that is delimited by the functionality discussed above.

The question now becomes, how can we talk about this artefact in a musicological analysis? I must find strategies to guide the reader's exploration. I wish to interrogate the form, the aspect of the network, so I will focus on structural points which appear to be constitutive of its morphology. First, we can start by looking at the nodes that directly address the bounding functionality that was first expressed: indeed, construction began around nodes which seemed to be most closely tied to this. The temptation could be to call them *central*, but that would go against the heterarchical nature of the approach. Let us call them, then, *centred*.

The first bounding function that was given was around the production of digital tools for fluid corpus manipulation: for this, we can observe the *FluCoMa* project node. When the network is displayed with the prefuse force-directed algorithm, this node indeed takes a *centred* spatial position, and it has thirteen neighbours. First, there are nodes for which it can be considered the source, for which agency can be conceived of as emerging from the node: the two *FluCoMa Toolsets*, the *FluCoMa project Commissioned Composers*, the *FluCoMa plenaries*, the *Reactive Flows* and *Liminal Spaces* concerts. Then, there are nodes where the *FluCoMa Project* appears to be the target: the *FluCoMa Grant Agreement* and *Pierre Alexandre Tremblay*. Finally, there are the nodes where agency seems to flow in both directions: *the University of Huddersfield*, *Owen Green*, *Gerard Roma*, *Alex Harker* and myself, *Jacob Hart*.

We can examine the nature of these nodes: 5 biological, 7 conceptual and 2 digital. As we could expect from the initial process, the two toolsets are present, along with the

principal materialities that drive the project: the core team members and the two institutions of the University (the geographical place where it is based) and the ERC's funding agreement (the actor that funds it). Neither the pieces, nor the individual commissioned composers are immediate neighbours. What could this mean in terms of agency that the *FluCoMa Project* deploys? Perhaps that the research ideas of the project are mediated, pass through the interface of the tools they create, and perhaps that commissioned composers can be considered as a group, rather than individuals. Do not forget that this is a biased perspective that stems from my visualisation of the network, yet it does lead us to pose interesting questions: for example, is it possible that the core team are more inclined to be influenced by the project than the commissioned composers? The proximity of the various nodes in question in this visualisation could suggest this; however, this is to be taken as an artefact that will generate these kinds of questions to alter our perspective, not demonstrably answer them.

When looking at the actual *Toolset* nodes, as the network has been constructed, the *Commissioned Composers* node is actually the same distance away as the core team members, however, instead of passing through the interface of the *FluCoMa Project*, they pass through each of their individual pieces. This suggests that the agency that the composers deploy upon the toolsets is through the medium of their work. We can imagine that the type of translation that occurs through these two different types of nodes is different – indeed, there is surely a big difference between inflecting the development of a toolset through a direct, overarching project philosophy and through an artistic practice.

The second process given was for the creation of musical networks by a group of musickers that would use the tools. The nodes that are probably the best candidates to examine are those of the pieces themselves. There isn't anything too surprising here: the nearest neighbours are the musicker in question, and performers that were involved in execution of the performance, the gig, and the toolboxes used to create them. Similarly, when we observe the nodes that would best be centred for the process of imparting energy within these networks, concentrated around two moments (the gigs themselves), it is also predictable in terms of nearest neighbours: the pieces and the festivals that played host to them. These three groups of nodes are all also very close together, with either one or two nodes separating them, and all occupy relatively *centred* positions within the prefuse force-directed layout.

Another way of approaching a network of this nature could be to examine nodes with exceptionally high degrees of edges that have become centred. The most centred node is that of the *University of Huddersfield* with 38 neighbours: 29 biological, 8 conceptual and 1 physical. The biological, human nodes are mainly staff members and research graduates; the conceptual nodes are the various research groups and festivals that occur at the university; and the physical one is *Huddersfield* itself.

These research groups can all be related back to the FluCoMa project in some way, and all their various research interests and approaches can be found intertwined with those of the FluCoMa project. There are some relatively self-evident ones, like Huddersfield's CCL and its symposium, but others that would seem further away. For example, Michael Clarke's IRiMaS and Technology and Creativity in Electroacoustic Music (TaCEM) projects, mainly musicological in nature, can be linked to the project through multiple paths – notably Frédéric Dufeu who is a core team member of them, who is my supervisor, myself being a member of the FluCoMa project. There is also CeReNeM, a research group with which many of the FluCoMa musickers are linked and whose approaches can be thought of as surely trickling, filtering through into the project in some form.

Another interesting approach can be to speculate on the agency that seemingly *decentred* nodes could have on centred ones. For example, on the outskirts of the network we find the node *Ursula von der Leyen*, who is the president of the *European Commission* node¹⁷. The shortest path to the *FluCoMa Project* passes through the *ERC*, then through the *FluCoMa Grant Agreement*. Indeed, it seems strange to discuss the agency that Ursula von der Leyen could have on a musical performance in Huddersfield, or even something like the European Union which occupies a position of equal distance; but it is true that these actors play a role in *allowing* for the subsequent actors to exist. Without the EU, there is no European Commission, there is no ERC, there is no grant agreement, there is no FluCoMa project, there is no Reactive Flows concert, there is no toolset, there is no piece. The type of agency they deploy may indeed only be financial or *permissive* in nature, yet the agency is still there. It is for this reason that I also chose to include the various institutions at which all the actors are based – there are no less than 18 universities in this network and a whole host of research projects receiving funding from a broad spectrum of funding

¹⁷ Von der Leyen occupied this position since 2019, at the time of the FluCoMa grant request, Jean-Claude Juncker occupied position.

bodies – all of them can be linked to everything else in some way, imparting some kind of agency.

Before moving on, we can touch on the modularisation, the grammatisation of the network into a form that is more conceivable. Indeed, with these different levels of abstraction, it could be said that I am already making some jarring grammatisations. At this global level, I have already performed a modularisation in containing the local level network *within* the node of each gig. Nodes such as the ERC, the University of Huddersfield, HCMF, and indeed any node within this network, are essentially modularisations. It must be understood that these modularisations have been made with thought and in the goal of constructing a useful tool for musicological analysis. It is not useful for analysis, at least in the current context when wishing to examine the afore-given processes, to draw a picture that would understand, for example, every human actor's body part, or every institution's member.

b) *Local level.*

We can now zoom in to a much more local level. The two networks¹⁸ I examine here can be conceived of as the sub-networks contained within the *Bates Mill's Blending Shed, Edinburgh College of Art, Reactive Flows, Liminal Spaces* performance nodes, and the various piece nodes from the network in the previous section. Once again, I start by stating the functionality, the processes that bound these networks: here, we are looking at a very focused process of musicking occurring at a given time within a given space.

As can be imagined, these networks are comprised mainly of physical typed nodes, compared to the previous network that mostly contained conceptual nodes. Here we begin to see the real materialities that have been configured by the various

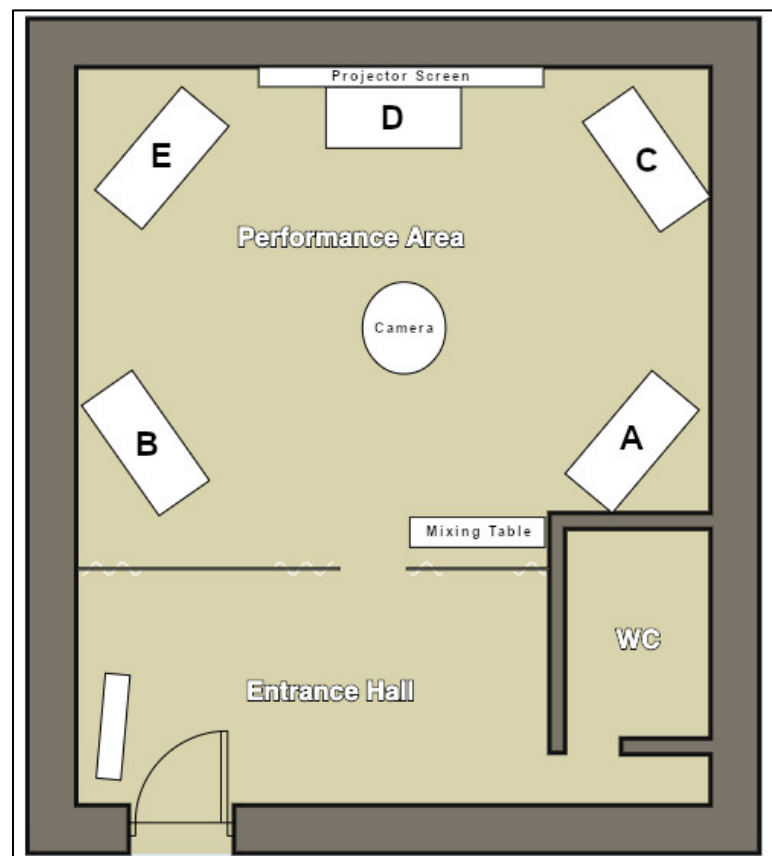


Figure 48: Bates Mill's Blending Shed performance area floor plan.

¹⁸ D.E.: 02_Data>01_Cytoscape_Networks>03_Local

musickers and start to see how their physical bodies interact with them. I will present my analysis of these two networks¹⁹ in a similar way that the last, focusing on what appear to be structural elements of the networks that could inform us, or drive our thinking towards interesting places. The networks are just about readable in the images, however the reader is invited to explore the networks on their own machine for a more agreeable exploration experience.

I start with the first *Reactive Flows* network. There are two very obviously *centred* nodes that come as no surprise: the *Performance Area* and *BMBS Mixing Table* nodes – all the performances were held in the same place, and all ran through the same mixing desk. All the immediate neighbours of the mixing table are the various speakers around the space that it is feeding audio to, and the various inputs from audio interfaces and microphones. Two nodes to be noted are those of the *Sound technician* and *Pierre Alexandre Tremblay*: before sound is output to the speakers, it passes through the mixing desk which in turn is mediated by these two people – this allows us to visualise the important role and the amount of agency that these two actors have on the musicking that occurs.

¹⁹ One for the *Reactive Flows* network, one for the *Liminal Spaces* network.

The *Performance Area* node's immediate neighbours are notably the five different stages and the six *Solenoid Base* nodes for Constanzo's performance. Indeed, the audience, finding themselves in the performance area (or not, see below), have a very short path before arriving at the solenoids; and the solenoids seem to constitute part of the environmental structure of the performance area, the same way that the stages




Figure 49: Reactive Flows performance network displayed with the prefuse force-directed algorithm.

and other objects in the room build up the space. The other objects that Constanzo controls tend to be more centred around him, rather than the performance area.

Let us address the bounding process by looking at some of the musickers and their places in the network, continuing first with Constanzo: we see that he imparts agency across the network via his four immediate neighbours: his two hands and feet. Not including Constanzo, these nodes have the following immediate neighbours: left hand, *condenser mic casing*, *thick comb*, *small comb* (x2), *dicer buttons*; right hand, *crossfader knob*, *thick comb*, *small comb* (x2); left foot, *expression pedal footboard*; right foot, *kick pedal footboard*.

These nodes could be good candidates for important points of interface within the instrument. Regarding the translations that occur, there seem to be two types of nodes: objects which imply *gestural* engagement, and objects which imply *sonorous* content. I discuss this further in Chapter 6, notably if we can consider these artefacts which demand gestural engagement as inherently musical or containing musical knowledge of some kind. For now, let us consider the target to which they appear to translate – we see that the left hand has a much wider reach, and that the hands are much busier than the feet. This reflects the importance and time that Constanzo attributed to the development of the condenser microphone and crossfader knob points of interface. Ultimately, their reach ends up in the *Macbook Pro* node, suggesting that the result of translation could be found there (see Chapter 4). Constanzo's *Macbook Pro* node is notably busy with 19 different edges and can be seen as *centred* for Constanzo's part of the network. Again, this leads us to consider that the *Macbook* could be a critical point of interface in his network.

How could we see a physical and conceptual node of the audience coming into Constanzo's part of the network? One of the reasons for which an *Audience* node has not been included, is that its place in the network would seem difficult to locate: does it coincide with the place where it is found (the *Performance Area*), or can it be viewed as some kind of super-node, that is somehow detached from this network with a view and connections to various points at once? This is an idea that would differentiate it from a node such as *Rodrigo Constanzo*, who seems to be physically and immediately intertwined with the materialities of the network – he becomes *part of the materialities*. He is less able to navigate some of the more conceptual, ephemeral nodes at play.

We could also assimilate our position of analysis to this: there is a certain degree of *omnipotence*, perhaps even more so than that of the audience, where our place in the

network can be seen as particularly fluid and intangible. We are in a position where we can focus our attention on any node at any point, something that seems difficult for the performer during performance.

For now, let's consider this from a pragmatic angle: let us examine Constanzo's network while coinciding the audience with the performance area. If we look at the lengths of the various paths to critical points, we discover that there are 4 nodes that separate the performance area and each of the hands and feet, 3 nodes to the *Macbook Pro* and 3 nodes to Constanzo himself. Does this mean that the audience feels closer to Constanzo himself and the laptop than the hands and feet during performance? Perhaps, but this could also very easily be argued against. It seems that both the *omnipotent* nature of the audience node and its materialities should be considered.

Looking at Hayes' part of the network, we see that her node comes into contact much quicker than Constanzo's with nodes that immediately emit sonorous content. Does this reveal something about the nature of the network, or is the methodology to be brought into question? There are elements of both. Indeed, Constanzo's network is much more atomic in nature, each composite part of the instrument seems to be cited, whereas with Hayes, there are nodes that could arguably be divided into many more composite nodes – for example, why have I chosen to represent a composite node called *Gamepad*, rather than giving a network with nodes like *joystick*, *button*, and *plastic casing*? This is certainly an approach that could yield interesting answers, especially when we recall the delicate relationship that Hayes holds with the object and the notion that minute movements will evoke big changes in the sound. However, we also need to consider the grammatisation that the musickers operate with their networks.

Constanzo clearly gives much more importance to the atomic composition of the objects he is manipulating²⁰, whereas Hayes is more on the side of taking objects as given²¹. It seems fair to suppose that the way she conceives of these objects is a lot less molecular than Constanzo. This, of course, is entirely debatable, but constitutes the type of decision one must make when using methods like this.

Parallel to this, we can also remark the comparative proximity of Hayes' sonorous nodes to the speakers and to the hypothetical *audience*. Perhaps there is something to

²⁰ Illustrated by the fact that he built much of the elements in his network himself.

²¹ Illustrated by the presence of the multiple commercial hardware synthesizers.

be said of the nature of Hayes' network that is visually more familiar to an audience than one such as Constanzo's, where the hybridised instrument retains a part of mystery. Hayes presents a setup that – while remaining complicated and not immediately revealing the sound that is going to be heard – does still live up to certain expectations that we can imagine the audience having²². There is also a very immediate connection between her singing, front of stage, into the microphone and immediately hearing the sound of her voice, compared to Constanzo performing gestures that are hidden, and hearing noises that are completely unexpected in the context of a snare and kickdrum.

When we look to the *Olivier Pasquet* node, we instantly remark its isolation on the outskirts of the network. Indeed, as we have seen, there are very little physical elements to his setup – the only notable contribution he brings to the environmental structure of the space are the *Spotlight* nodes which are mounted on some of the speaker stands – and these are all detached from him as a person in the space.

Again, we really see the nature of his performance reflected, and his history as a musicker revealing itself: the audience is immersed within something much closer to an installation, and things are happening around them without the presence of a performer. In an offhand comment, Pasquet admitted to me that he takes a certain pleasure from wondering around some of his installations, unseen and unrecognised, and listening to some of the comments that people make to each other when experiencing his work. This may even offer him a certain level of omnipotence like that discussed before of the audience – perhaps he prefers to disentangle himself from the materialities of the network, and be able to connect with decentred, or more conceptual nodes.

²² The network of various electronic objects on the floor and flight-case, the computer in plain view, the multitude of wires and a game controller come with a cultural baggage that suggests some kind of electronic music. This can be contrasted with Constanzo's setup which largely suggests acoustic musicking.

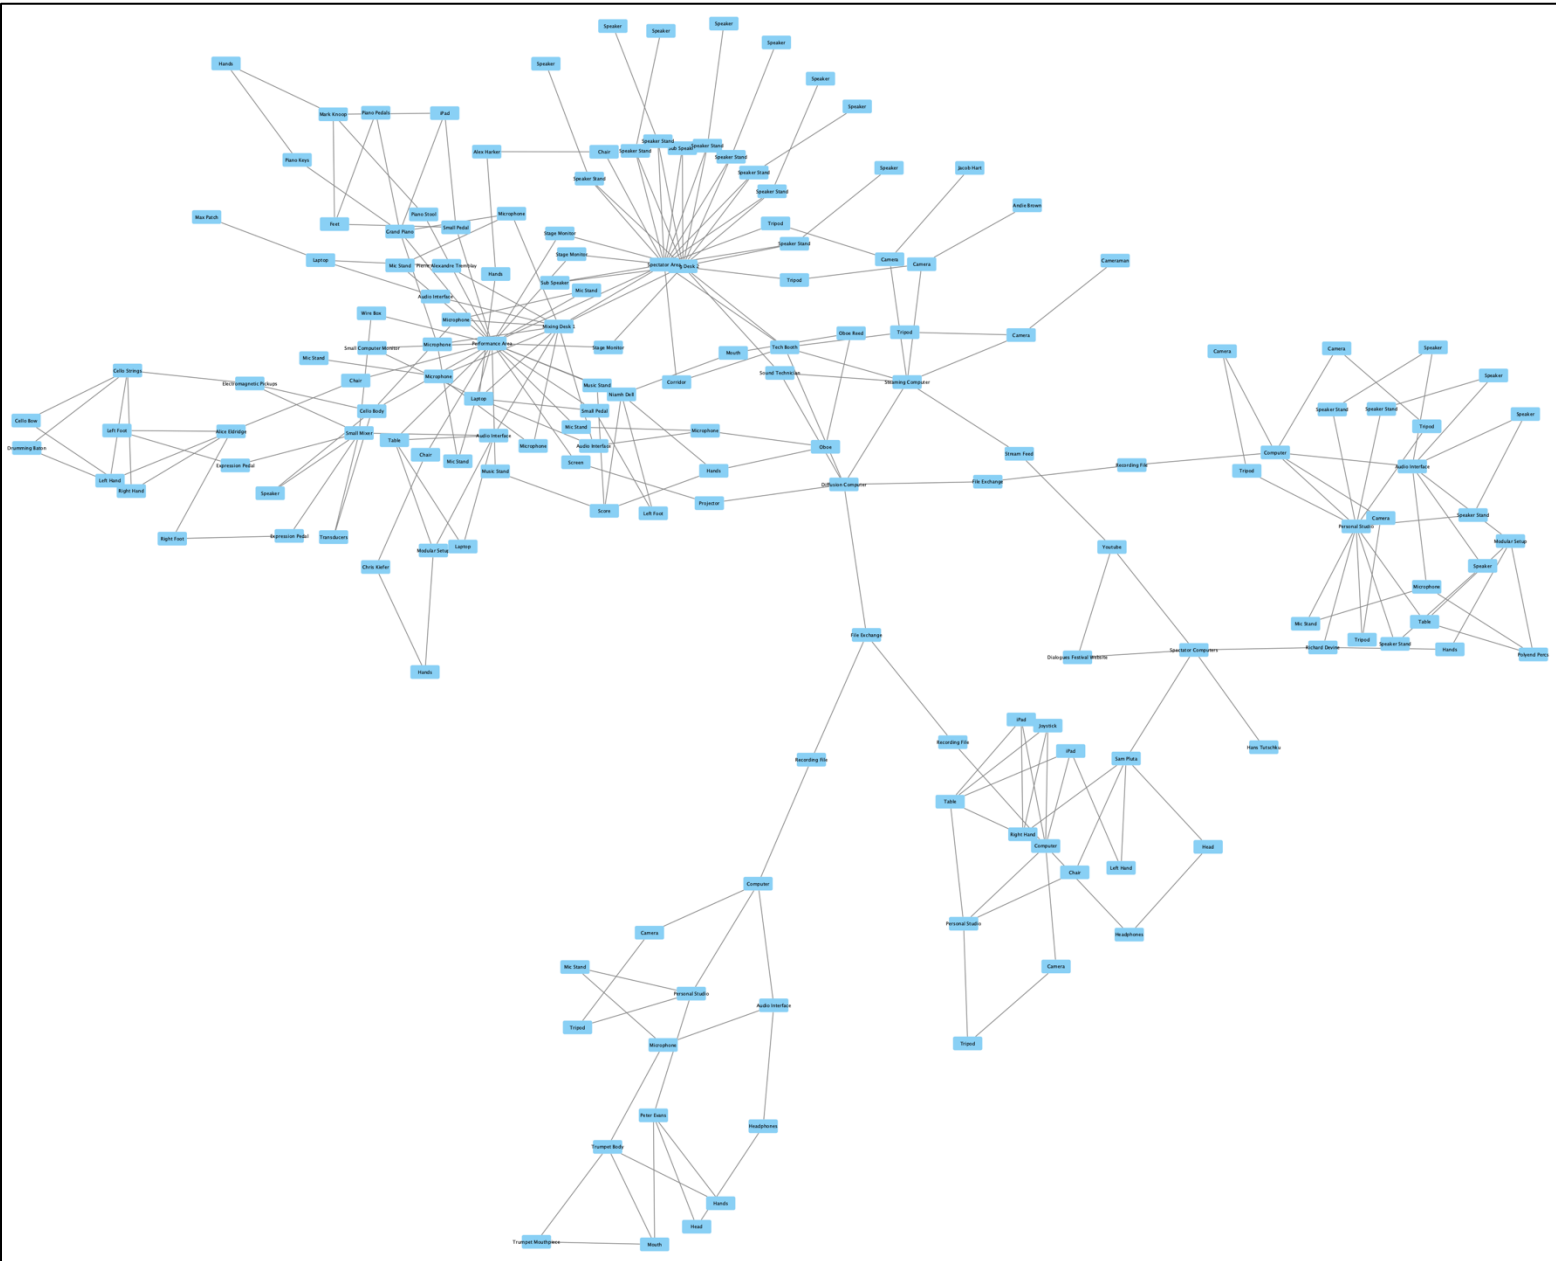


Figure 50: Liminal Spaces performance network displayed with the prefuse force-directed algorithm.

Let's now look at the second, *Liminal Spaces* network. It is comparatively simpler because the five concerts share the same stage area. We are immediately struck by the four large clusters of nodes. This is because, due to the COVID-19 health restrictions imposed at the time, three of the performances were diffused via video. They were pre-recorded, and we see that the three²³ are isolated from the main body of the network, passing through *File Transfer* digital nodes which link the musician's computers and the *Diffusion Computer* node. This is an interesting situation to

²³ Or rather two, the networks centred around the *Sam Pluta* and *Peter Evans* nodes essentially constituting a same network.

consider, and it must also not be forgotten that time is largely expanded here – Pluta and Devine’s pieces were recorded long before the time of their diffusion.

Looking at Pluta’s piece, we see that there is another *File Transfer* node that links the nodes centred around him and around *Peter Evans*, the trumpeter with whom he performed the duo. These are also the only musickers that have *Headphones* in their parts of the network – the conditions of performance were very different. As seen in Chapter 2, Pluta discussed the notion of performing in this displaced manner. Pluta remarked²⁴ that when watching the video back – especially for the first piece – he heard no inherent difference to an improvisation that could have been performed by the two musicians in the same room, at the same time. This has led him to reflect on his performance practice in general.

In any case, the types and extent of agency of actors such as the audience, the *Performance Area*, and the *Cameras* must differ greatly to Pluta and Evans’ usual concert experience. The same can be said when observing the part of the network that displays the materialities of the streaming process. The musickers are much further away from their usual audience, and the performance passes through a whole host of nodes before arriving at the *Spectator Computers* node: *Mixing Desk 1*, *Mixing Desk 2*, *Streaming Computer*, *Stream Feed*, *Youtube* and finally *Dialogues Festival Website*.

All these nodes impart agency and some kind of translation on the agency that is passing through them. To further illustrate this idea, the *Hans Tutschku* node is completely on its own, isolated in its own corner of the network – indeed, his piece was performed live, but *Pierre Alexandre Tremblay* was running the patch, and apart from an anonymous watching of the stream feed, Tutschku was essentially absent from the network.

The main centred nodes of the network are *Mixing Desk 1*, *Mixing Desk 2*, *Spectator Area* and *Performance Area*. We can observe the neighbouring nodes of the two mixing desks: *Mixing Desk 2* which was up above the spectator area in the tech booth, is directly connected to all the *Speaker* nodes; whereas *Mixing Desk 1*, down at the back of the spectator area, is linked to the various audio interfaces and microphones. There is a splitting of two very different tasks – input and output.

Different profiles of people are found at the two desks: up in the tech booth there is the sound technician, whereas down in the spectator area, with the inputs, there are

²⁴ Appendix 8.5.5.

musickers such as Pierre Alexandre Tremblay and Alex Harker. This suggests that *Mixing Desk 1* would impart more *artistic, aesthetic* agency within the network, whereas *Mixing Desk 2* would be more permissive and functional in nature. We remark the spatial position of these actors – the musickers are downstairs, immersed in the *Spectator Area* where all the speakers are placed.

Here, there are several configurations where two musickers are involved with the performance – the only musicker who performs alone is Devine. We have already seen Pluta and Evans' relationship which passes through many points of interface, what about the others? Harker and Dell are 4 nodes away from each other, Eldridge and Kiefer are also 4, and Tutschku and Knoop are 9 (although in this context, it would also be reasonable to consider Knoop and Tremblay, who are just 5 nodes from each other).

Musickers who are either performing together, or where one is performing for the other, seem to always be around 4 or 5 nodes away from each other. Again, this could be down to the analyst's grammatisations, but I believe there is some truth to this. One would image that perhaps Eldridge and Kiefer could be expected to be closer together, given their proximity on stage and their idea of a shared instrument; however, when one considers the nature of the network configuration, this doesn't necessarily ring true. Consider the complexity of the network, and the sheer amount of materialities on stage. Just to look at them, we clearly see that they are deeply entwined within their own sub-networks, Eldridge within the Feedback Cello, Kiefer within the complicated modular setup. It can be argued that again, this immediate physicality and proximity – almost symbiosis – with the physical objects, demand of them a focus of attention towards their bodies, therefore creating a certain space between the two.

With these networks, we begin to understand the materialities of the configurations that the FluCoMa musickers have configured, and the varying ways with which they exist within them. Some networks are explicitly configured in ways that demand immediate physicality, which created space and disconnection from other parts of the network, whereas others demand less physical engagement, and allow the musicker to adopt a posture that becomes like the audience, according to them a certain omnipotence and ability to engage with nodes further away in the network of a more conceptual nature. Next, I go further still, taking one more step in my levels of abstraction and examine some of these networks at an instrumental level.

3. Instrumental Level Networks.

Here, I move away from Piekut's (Piekut, 2014) network as a historical account of the broad-scale agency that institutions and people can have on musical production, and really focus on the materialities and composition of the instruments themselves. These networks are permissive in other ways – where an institution such as the ERC is permissive in as much as its financial backing allows for the happening of a music, here, something as seemingly inconsequential as the physical composition of a slider can have effects on the musicking which occurs.

As discussed in the Introduction, I have not chosen to give a detailed account of each of the networks at this level – most of these networks comprise several thousands of nodes, and each configuration could be the subject of a whole thesis. What's more, restrictions in the timeline meant that a detailed examination of the networks at this level was only possible for the first cohort of musickers. Instead, I will take some examples that I believe to be significantly interesting to discuss, and that push the methodology into interesting places. Again, networks of this size are not something that can be fully engaged with in image form: the reader is invited to explore the networks that accompany this text at their leisure²⁵.

Here, the methodology consists of transcribing the code written by the musickers in Max into network form and using some of the tools offered by Cytoscape for analysis; also, I insert them into some of the broader networks I have already presented. I divide this section into two sub-sections that present two uses for this kind of approach: examining networks as static and examining a network's progression over time.

a) Static networks.

Let us take a first example with Constanzo's work. There are three network visualisations to observe here: the first is a raw network visualisation of the Max patch itself; the second is a further grammatised version of this network, where the various modularisations that Constanzo made (which symptomatically reveal something about Constanzo's conception of this network) have been applied by grouping nodes into sub-networks; and finally, the second network that has been inserted into the previous local level network, revealing some of the points of interface between the physical and digital domains.

²⁵ D.E.: 02_Data>01_Cytoscape_Networks>04_Instrumental

The bounding processes for this network are constant across the three: taking input from physical objects and using that to trigger and process sounds – essentially looking at inputs and outputs. In the first network, we immediately observe three main regions, three networks that are isolated from one another. At the bottom, the two smaller networks correspond to the batch processing and analysis of folders of sound located in the pre-processing subpatcher; the bigger network corresponds to the rest of the patch which deals with the performance itself.

In the performance patch part, we can observe some centred nodes which reveal some important points of interface: first, the *s dicerButtons* node, which corresponds to messages being received from the Dicer MIDI controller and sending them around the patch. This has 18 immediate neighbours, 9 of which can be discounted as corresponding to routing from the Dicer to the send (the *prepend* nodes). It could be tempting to try and assess the coverage and reach of these nodes by looking at the iterations of their immediate neighbours, and their immediate neighbours and so on. Indeed, when taking the Dicer receives, in only ten iterations a large majority of the




Figure 51: raw Max patch network visualisation of Constanzo's Kaizo Snare performance patch displayed with a prefuse force-directed layout.

nodes in the network can be reached. However, I do not believe this to be an insightful

exercise, as the internal structure of a network of this kind means that this process could be repeated with almost any node (with the exception maybe of nodes placed on the very outskirts of the network). Through my research, I have discovered that the most insightful ways of approaching a network like this are by looking at immediate neighbours and paths between nodes – the further one gets from a node, the less meaningful the data becomes, as the chance of hitting another centred node becomes higher.

That being said, we can draw meaning from the number of degrees a node has, and the scope of directions as neighbours begin to branch out. Many nodes will follow a path leading in one direction for several iterations before hitting a centred node. Centred nodes, however, tend to branch out in multiple directions, informing us

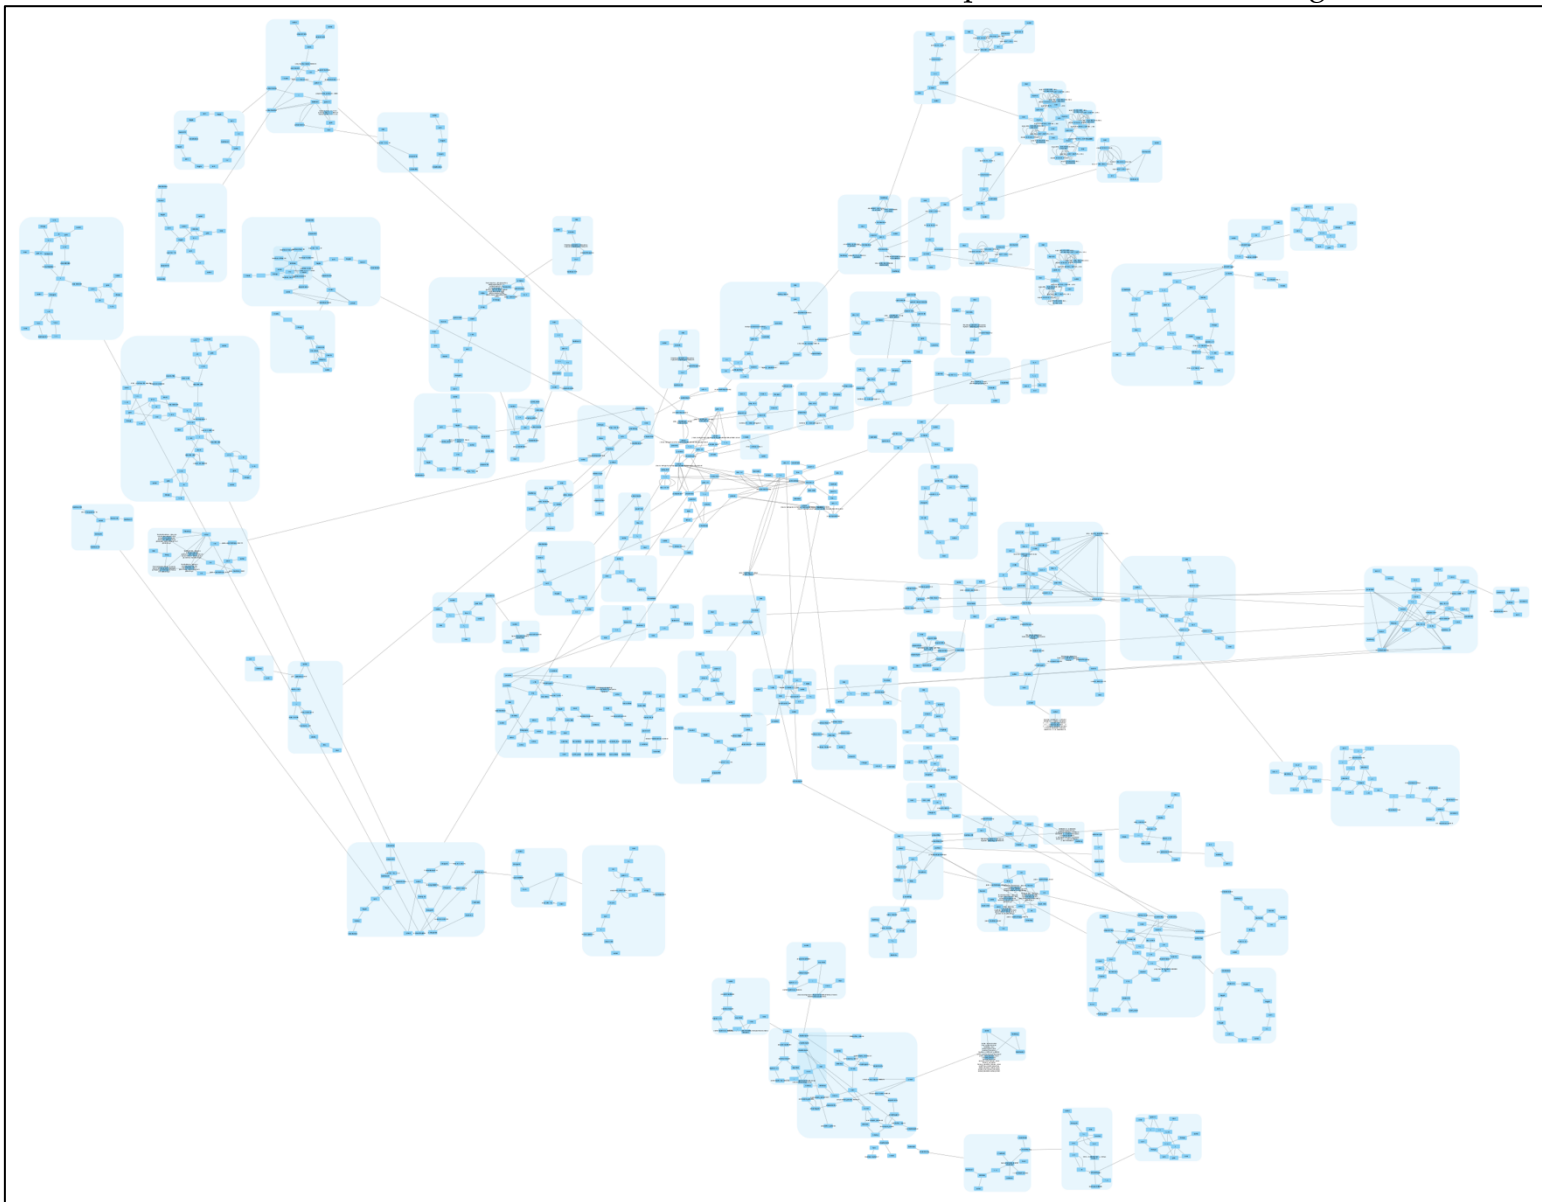


Figure 52: grammatised version of the Kaiso Snare network visualisation.

perhaps on a wider distribution of agency across the network. Indeed, in this case when considering the functionality of this node, we see that it is certainly not unreasonable to hypothesise that it will have a great impact across the network: the Dicer is used to change functionality and modes in a variety of places in the patch.

As can be imagined, other notable centred nodes, are the points of interface between the digital part of the network and the physical part: between the *Audio Interface* and

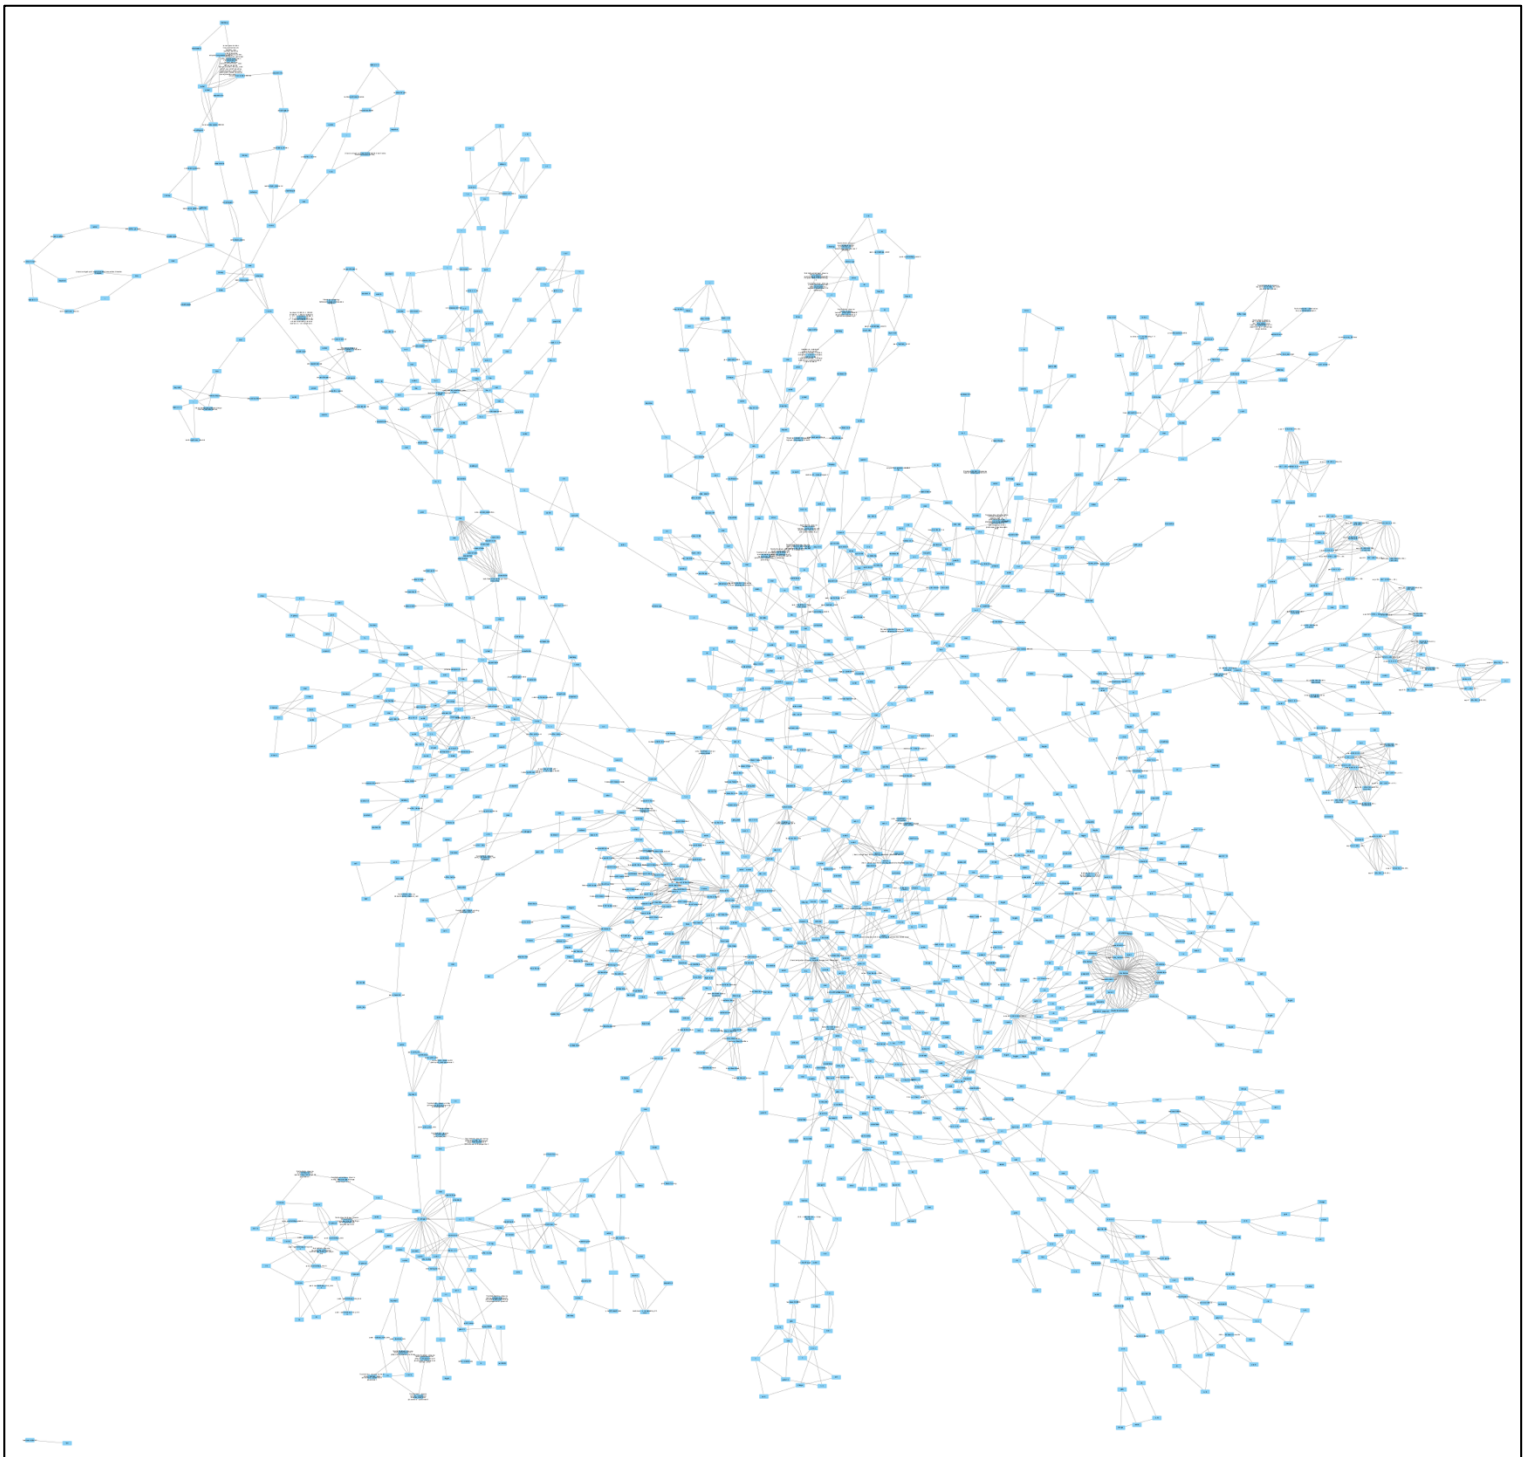


Figure 53: the Kaizo Snare network insterted into the local level network.




Figure 54: Kaizo Snare local and instrumental level network displayed with a hierarchical layout algorithm.

the various *adc~* and *dac~* nodes stemming from the *Condenser mic cable* and *DDrum Shot Trigger*; between the *Solenoid Holders* and the *noteout "dadamachines automat"* node; and between the *Teensy Microcontroller* and the *teensyMicrocontroller* abstractions. It can be interesting to examine the paths between these points of interface into the digital domain, and the paths they can take before becoming sound.

If we look at the *hierarchical* layout, the connectivity between Constanzo's hands and feet and parts of the patch becomes apparent. In the annotated image, we see the position of the two hand nodes in light green and the two feet nodes in dark green. The various *dac~* nodes, where sound emerges from the digital domain, are in purple, and the object that sends control data out to the solenoids is in red. Despite being very busy and being connected to more nodes, hands have a less direct connection to actual

sound output. Essentially, they may control more things, but their agency is perhaps diluted within the patch.

Another way we can engage with this type of network is by looking at specific actors – for example, some of the FluCoMa objects. This can also quickly inform us on some of the tastes the musickers had for the tools, and perhaps inform on which parts of the FluCoMa approach they feel is compatible with their own. When we run a search for the FluCoMa tools in Constanzo’s patch, there are 29 instances that appear: 1 *fluid.ampslice~*, 1 *fluid.bufcompose~*, one *fluid.bufhpss~*, 4 *fluid.bufloudness~*, 1 *fluid.buфонsetslice~*, 3 *fluid.bufpitch~*, 4 *fluid.spectralshape~*, 12 *fluid.bufstats~*, 1 *fluid.buftransients~*, and 1 *fluid.pitch~*. It is remarkable that in these 29 instances, there is only one use of a real-time object with *fluid.pitch~* – an object which notably produces a descriptive piece of data on the incoming audio. All of the other objects operate on buffers, and most of them (*fluid.bufloudness~*, *fluid.bufpitch~*, *fluid.spectralshape~* and *fluid.bufstats~*) are objects that, again, produce and manipulate descriptive data.

This is interesting when we consider Constanzo’s practice which seems very immediate and real-time in nature. As we saw in our initial analysis, Constanzo often defects to using rolling buffers – indeed, he discusses²⁶ the need for very low latency, and it can be assumed that the real-time objects did not offer a latency that was low enough for his immediate needs. These offline configurations can, then, perhaps be considered as workarounds for what he was really aiming for.

Compare this to somebody like Hayes: 1 *fluid.pitch~*, 2 *fluid.ampslice~*, 4 *fluid.bufcompose~*, 2 *fluid.bufnmf~* 2, 1 *fluid.bufnoveltyslice~*, 3 *fluid.bufpitch~*, 2 *fluid.bufstats~*, 2 *fluid.buftransientslice~* and 1 *fluid.hpss~*. Again, a surprising amount of the objects operate offline, although we do see more real-time objects. Here, we observe a much wider use of the slicing algorithms, as well as some of the descriptive objects, revealing some of the priorities in Hayes’ practice.

Another example: Burton only makes use of the *fluid.bufloudness~* and 2 *fluidmelbands~* objects. Indeed, we saw in our initial analysis that he used the tools to assess the global output and fill in the quieter parts of the spectrum through his processing – he appears to have no need for slicing tools whatsoever. This kind of statistical data that these

²⁶ Appendix 8.4.3.

networks allow us to quickly access does indeed allow us to articulate some questions around the artistic project of the musickers.

b) *Networks over time.*

Finally, I examine how these kinds of network visualisations can help us understand the development of networks over time. Some of the musickers gave me previous versions of their code as it progressed over the development process. Let us start by looking at Burton's work. If the reader consults the video presenting the progression of the patch²⁷, they will see that Burton gave me 37 different iterations.

By looking through this video, we notice three main periods: from *ScorePage* to *ScorePage019* the network grows progressively, starting at 87 nodes and finally gravitating around 1300 nodes – the aspect of the network is quite unstable at first, but then from around *ScorePage013* onwards, it seems to stabilise more (the files skip here, there is no *ScorePage012*). The network does generally retain a structure of having one large group of nodes above a smaller one. From *ScorePage007* onwards we also remark arc shapes that start branching out. Then, at *ScorePage020* there is a sudden shift: the size of the network increases greatly with 3350 nodes and 6014 edges, with the network spreading out from the middle in spokes. This continues through to *ScorePage023*. There is a final period where the network is reduced in size once more and starts resembling those at the end of the first period. The network will continue to grow along the same aspect, with the only notable change occurring at *ScorePage031*, where a distinctive ladder-like structure appears on the edge of the network and continues until the end.

We see that this method allows us to observe big structural changes in the code very quickly, but what do these elements correspond to? The patches from the beginning up to *ScorePage019* are progressively building up the score generation part of his final patch. It is interesting to note that this is where Burton began, rather than the audio processing. The score was visibly very important, and perhaps not only drove what was happening in performance, but also development of the sound processing in the patch. Some of the branching arc structures that were noticed correspond to sending the score dimensions to various parts of the patch and sending the current time to other parts. Much more of the patch seems to be affected and structured by calculations regarding the spatial content of the score, rather than the temporal

²⁷ D.E.: 04_Video>02_Network_Visualisation>ScorePage_progression.mov

dimension of the current place in the piece. Also, by *ScorePage019* we are already beginning to see the code that deals with analysis of output audio with the loudness of mel bands.

The major change from *ScorePage020* that was noticed is due to passing from 4 machine block subpatchers to 16. It is safe to say that this major structural change could have gone unnoticed had I not had access to these tools – my detailed analysis of the code for the final iteration took a very long time, and it would have been impossible to repeat the process for all 37 iterations in the time I had available to me. This was hidden in a subpatcher on the parent patcher, then in another subpatcher from there – indeed, the importance it had regarding the structure of the overall network is certainly not apparent at first glance. Thanks to this sweeping network overview, this change was clearly visible. Likewise, the change of structure from *Score024* onwards is due to the placement of this processing in the newly introduced *PlayerBlok* abstractions. We observe this initial playback of audio transform first by multiplying, then by articulating itself as an abstraction where processing becomes more complicated. Burton effectively goes from score to the notion of simple play back of sound, to more complicated and precise processing of this sound.

Finally, the ladder-like structure that was identified at the top of *ScorePage031* and onwards, corresponds to some number visualisation that is placed in the mel bands analysis part of the patch. This visualisation, however, is not actually connected to the analysis – we can suppose that at this point in the development process, he needed some feedback from the algorithm but didn't need it for the piece.

To finish, let's look at some of Pasquet's work²⁸. There is a very clear progression throughout the development process of his performance patch. The network starts out as one body, then there is one major mutation in *HH_5*: a big circular clump of nodes branches out of the main network and will stay isolated until the final iteration. From there, the rest of the network progressively starts to divide itself into two groups which are similar in size. In one of these groups, there is a node that seems to be an interface for a particularly high amount of activity, notably from *HH_7* onwards. Let us investigate to see what all of this means.

From the very beginning, we find two elements which are important in the final iteration: the main rhythm tree generation and the Modalys reed instrument. I suspect

²⁸ D.E.: 04_Video>02_Network_Visualisation>HH_progression.mov

that the Modalys instrument was imported from another one of Pasquet's projects. One interesting element to note is that the main tree rhythm has much less scope than one would expect within the network. This is not to mean that this node does not impart a great deal of agency within the network, indeed it does; however, it is filtered through and translated in many ways, meaning that it doesn't necessarily manifest itself as particularly important when just looking at the visualisation.

The large circular clump that separates itself from the rest of the network from *HH_5* onwards that joins the rest of the network through only one *qmetro* node, corresponds to the control interface of number boxes that is routed to control the Modalys reed instruments. This switched from two dials controlling the radius and reed weight for each instrument to four float boxes controlling radius, weight, area, and aperture. The two groups of the network finally organise themselves with the actual processing of each of the four Modalys instruments (it happens that the centred node identified above is actually two closely knit nodes routing parameters to each of the instruments) in one part of the network, and the rest of the patch in the other.

Working from sources of this iterative nature seems to be a potent source for my network approach. Over this section, we have seen that this method allows us to access and quantify information about the networks quickly. It also allows us to visualise aspects from various domains at the same time (institutional and social, physical, digital).

I have also been able to pose informed and articulate questions about the structure of a network: how certain nodes interact with each other and the types of agency which are circulating. However, it has also become apparent that it is always necessary to refer to the initial surface level analyses to avoid losing a grounding in reality when studying this mapping.

Indeed, this is the type of artefact that we are dealing with: a map. There are two important things to bear in mind when engaging with this kind of object: first, no map, as detailed as one may make it, is fully representative of reality, and while revealing elements which may have been hidden when faced with reality, can also be read in a way that would lead one astray from it; second, maps are made by human beings – their resolution, and what is chosen to be represented is governed by the cartographer. I return to this in Chapter 5.

I return to some of these visualisations in following chapters, however, I will now move on to another type of network. Here we have been giving accounts of the

structure and composition of these networks in the digital, physical, and social domains. Now I wish to return to some of the ideas discussed in Chapter 1 and try to consider them in a temporally multidimensional manner.

Chapter IV. SONOROUS POTENTIALITIES.

1. Tool Development.

After having constructed a collection of network representations that intend to represent some of the *physicalities* of the networks in question, I shall now present construction of network representations that seek to represent the *sonorous potentialities* of these networks. This begins to approach one of my principal research questions – when considering a network not as static, but as multidimensional in time, as something that allows for things to happen, how can we access and analyse the entire scope of sonorous potentialities of a musical network? Immediately, this strikes us an inherently impossible task – a musical instrument could be used in an infinite number of ways. Therefore, the first principle by which I abide, is that analysis should be guided by the use of a network – through configuration and existence within it, the musician has inscribed musical thought into a network and we can use this to our advantage when slimming the range of potentialities.

This does not mean that we can content ourselves to study one performance, one iteration of this network in motion – but instances of its iteration must not be ignored as they attest to inscribed musical thought. The goal is to access the potentialities that a network *could* make according to the broad framework of the musical thought inscribed within it, not to blindly explore any and all potentialities that exist.

There are countless ways that one could go about trying to accomplish this task. The methodology refined itself as it encountered different case studies. In this chapter, I present the construction of *sound plots* for several of the networks and explain how the methodology has developed. Then, in Chapter 5, I explain how I engage with these artefacts that I have created in my *cartological* analysis. First, I discuss how I developed a set of tools to help me accomplish these tasks.

a) *Requirements and Goals.*

When seeking to access the entire scope of sonorous potentialities of an instrument, there are three elements that are quite certain: I will be dealing with large amounts of audio; a solution must be found for exploring these large amounts of audio in a way that is useful for analysis; and it will be necessary to manipulate the musical networks in some automated way. This last point means that I will have to find a way of interacting with the network myself. For the case of a musician like Olivier Pasquet, this doesn't seem particularly difficult as his whole instrument is contained within the

digital domain. Essentially, if I possess all his code, it should be possible to recreate sounds that emerged from his network. There are other networks, however, where there are a whole host of external, physical elements that I may not necessarily have access to for analysis: Burton, Constanzo, Devine, Eldridge, Harker, Hayes, Pluta and Tutschku all integrate acoustic or electronic instruments into their networks. I will therefore need to find strategies for recreating a situation that is faithful to how these networks were deployed during performance in an analytical setting. Similarly, it will be necessary to find ways of keeping aspects of the network's use when drawing the sonorous potentialities from the network.

The first two points are also important for the methodology: indeed, if my analysis consists of ending up with hours and hours of audio without any intelligent way of going through it other than listening to it second by second, then I won't have gotten very far. Indeed, I believe this to be a primary question for contemporary *computational* musicology which can precisely levy the computer's ability to archive large amounts of audio: how can these huge amounts of data be deployed in analysis in a way that has meaning, in a way that does not overwhelm the human analyst?

It happens that many of the tools produced by the FluCoMa project aim to offer solutions to this very problem; therefore, it was a great opportunity to engage with the tools not only as a subject for my analysis, but as tools for accomplishing it. In parallel to this, it can be expected that when examining the sonorous potentialities of a network, there will be a lot of redundant data: there may be a lot of repetition and things that sound the same. It will be useful to have a way of quickly being able to discard things that aren't of a great deal of interest and be able to discover the various extremities of a network's sonorous potentialities.

Other methodological questions emerged as each case study was encountered and I will go over them as they arise. First, I present some of the tools that I have developed for my analysis.

b) Tools.

At the second FluCoMa plenary, Gerard Roma presented a 2D sound plot where a collection of sounds had been distributed into the space according to their timbral data (Roma et al., 2019). This struck me as proposing a physical manifestation of Hanninen's (Hanninen, 2012) *associative landscapes*. I had already seen things like this, notably the IRCAM's CataRT software. In CataRT, the user can distribute sound

across the space using two user-defined axes of descriptor data, or other metadata like length.

What enticed me with Roma's approach, was the fact that the axes were not organised according to two sets of parameters but were essentially meaningless. He had used *dimensionality reduction* to arrive at a distribution in space where the axes held no signification; but position in space, and proximity with other nodes, was meaningful. This approach was particularly encouraging, as it bypassed the element that I found most problematic for analysis using something like CataRT: the reductive definition of axes. Here, it was possible to draw data from a whole host of different descriptions and draw a distribution from them that goes far beyond just 2 dimensions.

I began to explore the idea of implementing this into my own analysis. The FluCoMa project essentially proposes the following process for dealing with large amounts of audio: slicing (in time and energy), description, dimensionality reduction, exploration. I decided that it would be useful, not just for myself, but for the project and for other musicologists, to build a wrapper around the FluCoMa objects with an intuitive user interface¹.

This has gone through many iterations. For the first version, I recycled some of the JavaScript UI library I had built-up for creating my Cytoscape emulation (see Chapter 3) in the goal of making a single object that would give control over all the FluCoMa objects with an intuitive UI. This was called *FluidControl*².

I next built on this in a project called *amat*³ (Automated Musicological Analysis Tool) which grouped together a few things I had been working on: a more in-depth version of the *FluidControl* tool with presets for the objects and a dynamic FFT visualisation, a slicing visualisation tool, and the beginnings of a segment visualisation tool. Up to here, I was focusing mainly on the slicing aspect of the methodology.

Next, I began to implement descriptors when starting to look at Constanzo's analysis: in a project called *instrumentAnalysis*⁴, I built a patch that could take slices and generate descriptor data using not only the FluCoMa tools, but also Alex Harker's

¹ D.E.: 01_Code>02_Tools>02_Fluid_Musicology

² Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/02_fluidControl

³ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Unfinished/03_amat

⁴ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/03_instrumentAnalysis

implementation of them. This patch also gave a great deal of control over the settings of the various objects: the user could combine different sets of descriptors and statistics and the patch dynamically displayed explications of the algorithms at play.

My approach then became a lot more modular: incorporating fundamental principles of programming, such as modularity and interchangeability of data. Over the previous iterations, I had built my tools such that the results gathered were difficult to take again and apply in other contexts. Also, the user had to perform each process from beginning to end and would have to start over again if they wanted to change a parameter. To work efficiently, it would be necessary to have a system where all the data procured was agnostic to the environment it was created in and easily usable in multiple environments (I eventually settled on JSON files); and where each step was separate, so that the user could start from any point in the process without having to do everything again.

It would also be useful to implement these tools not only in Max, but in at least one written programming language. I began building a project called *computationalMusicology*⁵ where each tool was a different module implemented in Max and Python, and data could be freely exchanged between them. I used James Bradbury's wrapper for the CLI version of the FluCoMa tools⁶ which I used heavily in the Python parts of the code⁷.

We begin to see the real skeleton of the methodology being fleshed out, with separate tools for slicing, description, dimensionality reduction, file conversion and visualisation of dimensionality reductions (I implemented a 2D sound plot⁸ patch using JavaScript in Max, and Van der Wee added a 3D plotter that he was developing for his research at the time).

Finally, I arrived at the form with which I performed a majority of my analyses. *FluidMusicology* has the same essential skeleton as the *computationalMusicology* project. I made a wrapper for all the FluCoMa tools where the attributes are dynamically displayed and changeable, and built everything up into a large, tabbed patch where data can be saved at each step. The dimensionality reduction part of the tools finished

⁵ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Unfinished/04_computational_musiocology

⁶ <https://github.com/jamesb93/python-flucoma>

⁷ I also credit Laurens Van der Wee with whom I initially started this project.

⁸ D.E.: 04_Video>01_Demos>max_plot_demo.mp4

in Python: this is because the FluCoMa tools only implemented Multidimensional Scaling (MDS) and Principal Component Analysis (PCA) dimensionality reduction algorithms at the time, and I found that they were yielding unsatisfactory results. There are many other algorithms that exist, many are implemented in the scikit-learn⁹ machine learning Python library. Also, the visualisation part of the patch would only ultimately be used for Constanzo's analysis, as it was necessary to do this in Max (see below). The others were done in a JavaScript, HTML environment I created for my thesis¹⁰ for which the reader may consult a video demonstration¹¹.

To finish, I shall go over the large steps of the methodology and how these tools make them possible. The reader may also consult a video demonstrating this¹². The user can first import any number of files into a format that is used by the rest of the software – the corpus is concatenating into one big file upon-which operations are performed. Then, the user can slice this data, using either one of the FluCoMa automatic slicing algorithms, slicing by file¹³ or a recursive slicing that aims to use one of the FluCoMa algorithms and find a final number of slices between a given minimum and maximum range by iteratively modifying a given parameter and evaluating the result¹⁴.

This slice data can then be described using the FluCoMa-implemented audio descriptors. These descriptions are represented for each slice with several statistics: mean, standard deviation, skewness, kurtosis, low, mid, and high. This data can then be taken and used for dimensionality reduction – essentially taking many dimensions for each slice and crunching them into 2 or 3 dimensions so that distribution in space will yield groupings of similar data. As mentioned above, this does have a Max implementation, but I used my Python wrapper for the scikit-learn implementation which gave me a greater number of reductions. This can then be viewed and navigated in a 2D space, either in the Max patch, or in the HTML environment mentioned above. Ultimately, this yields a sound plot where the various slices are placed in a space that can be explored and grammatised. It allows us to see the extent of the timbral content of a collection of sounds.

⁹ <https://scikit-learn.org/stable/>

¹⁰ Supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/06_PhD_Browser/PhD_Browser

¹¹ D.E.: 04_Video>01_Demos>phd_browser_demo.mp4

¹² D.E.: 04_Video>01_Demos>fluid_musicology_demo.mp4

¹³ A useful feature when the files were already divided into meaningful slices when collecting the initial data.

¹⁴ This is currently broken.

These tools offer robust solutions for the first two goals outlined in the above requirements and goals, but the notion of automated collection of data from the networks and creating of actual material for analysis remains an open one. This will depend greatly on the network in question and is outlined in what follows. There have been four main approaches to this question, led by the case studies which have driven the development of the methodology. I believe this research has led me to a broad framework, or a set of principles informed by practice, that could be applied to a wide number of networks.

2. The Iterative Approach.

The first iteration of analysis is demonstrated with Constanzo's work. Constanzo's instrument is based around three main effects modules: *wavefolder*, *cloud* and *dirt*. In addition to this, Constanzo's patch is essentially *modulative*, in as much as it takes sound in, processes it, then spits it out, rather than synthesizing sound. During the performance, Constanzo manipulated three principal sound sources: what I call the *raw mic* corpus – sounds picked up by the condenser microphone on the snare's surface; the *metal resonance*¹⁵ commercial sample library; and the small *gamelan* sample library.

The core of this first, iterative approach, will be to try and gain access to the scope of sonorous potentialities of the network by literally iterating through all the parametric states and all the sounds that can go through them. This is a heavy-handed approach but is certainly one that deserves a thorough investigation. Before I can start iteration through the various parametric configurations of the patch, I first need to gain a good overview of the possible input audio.

a) *Sample libraries.*

The *metal resonance* library is the biggest collection of sounds with 3094 different elements. Each sound is a metal object impacting against various surfaces. Many of the sounds are similar, but there is certainly a wide range of timbres across the corpus. It was necessary to try and build a much smaller collection of sounds that fairly represented the collection at large to keep the datasets as small as possible. I turned to my corpus manipulation tools to try and manage this problem. Perhaps due to the

¹⁵ <https://hissandaroar.com/v3/soundlibrary/sd013-metal-resonance/>

number of samples that are similar in nature, arriving at a distribution that portrayed the range of timbres was more difficult as first anticipated.

This was an opportunity to test out different combinations of audio descriptors, dimensionality reductions and FFT settings and see what seemed to work well. The reader may consult the accompanying elements¹⁶ to explore the sound plots themselves. There is also a video¹⁷ demonstrating how to operate this tool. In Appendix 4, the reader will find a detailed account of testing of various combinations for the metal resonance corpus. Despite many of the reduction and description techniques yielding unhelpful results, many of them do allow us to access the scope of timbres, pitches, and loudness within the sample library. I spent much time going through various combinations of settings to find ones that would be useful for analysis. As research progressed, I fell progressively into a set of combinations that always seemed to work well. However, it was certainly useful to start with this thorough investigation of the breadth of possibilities.

I repeated the exercise for the *gamelan* sound corpus. I built the *raw mic* corpus myself – Constanzo recorded his performance through multiple streams: raw microphone input, microphone input post-processing, main metal samples, gong samples, cloud output left, cloud output right, raw crossfader control data, scaled crossfader control data.

I was able to use this source to build-up a corpus of Constanzo's input audio. It does not constitute every single sound that could be made with the condenser microphone, combs and felt on the snare surface – however, this isn't necessarily a limitation. One issue I wished to address in this process was that of retaining an aspect of the network's use, and as long as the corpus demonstrates a wide enough spectrum of sounds, this is enough to work with.

I imported these recording streams into Reaper. I made use of James Bradbury's ReaCoMa tools and made a few modifications to the Lua scripts for my use-case¹⁸ and experimented with various slicing techniques. First, I tried using some of the slicing algorithms on the control data, then sliced the microphone input according to that. I

¹⁶ D.E.: 01_Code>02_Tools>05_Sound_Plot_Navigator_Standalone

¹⁷ D.E.: 04_Video>01_Demos>sound_plot_navigator_demo.mp4

¹⁸ The modifications I made were, when a media item is sliced by one of the FluCoMa slicing algorithms, instead of directly slicing the item, markers are added at the slice points to the Reaper project. In turn, I also added a script that allows the user to slice an item by marker. With these minor modifications, it was possible gather slice points from one track, and slice another from them.

did find some correlations, however, it ultimately yielded unsatisfactory results, seeming to cut events in the middle of a gesture.

Eventually, I used the same algorithm that Constanzo used in his code directly on the microphone input: the onset slicer. This gave me a corpus of 493 files that gave a broad range of different gestures and timbres. To get a better understanding of the corpus I performed the same kind of operations discussed above. Regarding the combinations of descriptors, dimensionality reduction algorithms and FFT sizes, I found many similar results: t-Distributed Stochastic Neighbour Embedding (t-SNE) (van der Maaten and Hinton, 2008) appears to work well in terms of meaningful clustering of sounds and timbral similarities, as well as macroscopic structure. I also tried the configuration of descriptors that Constanzo used in the network: loudness, centroid, and flatness. This configuration offered some useful results, but I didn't find them quite as effective as the MFCC analysis.

b) Parametric iterations.

With a set of sound plots that give a good overview of the three main corpora, I started iterating through the various parametric configurations through which they are processed. Before going over this process, it is first necessary to discuss some limitations to this approach: computing power, analogue resolution, and infinite variables.

First, computing power: many of the afore-mentioned slicing, description and reduction processes can demand a lot of processing power from the computer and take a long time to perform¹⁹. Furthermore, there is only a certain amount of audio data that can be loaded into memory at a given time. I had limited computing resources at my disposal and had to make decisions and compromises to try and yield the best results considering the hardware I had at hand²⁰.

Next, analogue resolution: in the following analysis, I iterate over the different possible parametric states of various effects modules in the network. The number of states is a Cartesian product of all the various parameters, and thus grows exponentially in quantity. There are also certain parameters that are not simply binary *on* or *off* but linear: therefore I must chose a resolution of granularity for the jumps

¹⁹ For example, the MFCC description of the 3000 *metal resonance* samples on my 2012 Macbook Pro can take several hours.

²⁰ This is also an issue to be considered when thinking about further uses by a wider audience of this methodology beyond this research project.

between different states. For example, a slider between 0 and 1 could have a resolution of 0.25, giving it 5 possible states (0, 0.25, 0.5, 0.75, 1). Depending on the resolution chosen, the number of different states one finds themselves dealing with can be gargantuan²¹. This means that I will have to make intelligent choices to access information that is useful.

Finally, infinite variables: as much as we can try to account for everything, it will always be impossible to account for every single variable in a network. This method will deliver a rough approximation of the ensemble of sonorous potentialities – however, it will be sufficient to begin to understand the extremities and broad structure of the networks.

Remaining grounded in the reality of the network, I followed four signal paths: *Metal resonance* > *Wavefolder* > *Cloud* (SP-A); *Raw mic* > *Clip and Filter* > *Dirt* > *Cloud* (the left strand of SP-B); *Raw mic* > *Clip and Filter* > *Dirt* > *Wavefolder* > *Cloud* (the other strand from SP-B); and *Gamelan* > *Wavefolder* (from SP-C). Most of these paths contain multiple effects modules. There are over 3000 metal resonance sounds, and over 400 raw microphone sounds. The *cloud* effect has 2 toggles, 1 7-state toggle and 5 linear parameters: with a resolution of 0.25 this gives us 87500 different states, with a resolution of 0.5 there are 8604. The *wavefolder* has 6 linear parameters: 0.25 resolution yields 15625 different states; 0.5 yields 729. Finally, the *dirt* module has 1 toggle and 4 linear parameters: 0.25 yields 1250 states, and 0.5 162.

The number of states is multiplied by the number of samples that are fed through them, as an effect may not react in the same way on one sound than another. Therefore, I began by drawing visualisations of the large corpora, searching for gross clusterings of sound: I needed to discover a small number of sounds that could account for most of the corpus.

For the metal resonance sample library, my initial investigation of the sound plots informed me that there were grossly two different categories of sound: long, clear, and resonant sounds; and bounced, dry, muffled sounds. I have chosen two sounds from the MFCC 128 FFT t-SNE visualisation that appeared to best represent this: *METAL RESONANCE HITS Suspended ROD BAR HEAVY 03 – 2727.wav* and *METAL RESONANCE DROPS Hard concrete GIRDER 03 – 74.wav*.

²¹ When preparing the afore-mentioned MaMI presentation, I found that a standard *Xbox 360* controller, with a resolution of 0.25 for the linear parameters, has more than 512 million different states.

From my initial investigation of the gamelan library, I found that there wasn't a large variation in timbre between samples – the main difference is in pitch. Therefore, I chose a sound from the middle of the pitch t-SNE visualisation: *Kempur-A119_HPSS.wav*.

Finally, the sounds representing the extremities of the raw microphone corpus are: crunches (*raw_mic-471.wav*), comb pops (*raw_mic-305.wav*) and microphone gestures (*raw_mic-016.wav*). I will also have to perform a similar exercise of clustering and representation at each stage of the signal paths.

I built a patch²² around each of Constanzo's effects modules called *FXIter*. The user gives the number of variables and their types (toggle, linear, or stepped linear) and the desired resolution for the linear parameters to a piece of JavaScript which then calculates the Cartesian product of these parameters and prepares a list of all states. Then, this list is iterated through: on each iteration, the state of the module is set, a piece of audio is fed through the module, this is recorded and saved, then we pass to the next state.

Once the process is finished, all the files can be fed into the corpus manipulation tools discussed above. Here, I give a full rendition of the first signal path discussed above, SP-A.

Having already chosen the metal resonance samples, this involves another intermediary analysis of a sound plot after the first step, where the samples from the metal resonance library have been fed through the *wavefolder* module in the fashion discussed above. Representative sounds are then taken from this sound plot and fed through the *cloud* module. I give a more interpretative analysis of the results in Chapter 5, here I am explaining how the sound plots were constituted. The intermediary steps for the other signal paths can be found in the accompanying elements.

Let us assess the results of the metal resonance samples passed through the *wavefolder* module. Learning from earlier experiments with various combinations, the MFCC audio descriptor, 1096 FFT size and t-SNE dimensionality reduction is what has been used – this combination appears to offer a good insight on the timbral content, which

²² D.E.: 01_Code>02_Tools>04_Constanzo_Iterator




Figure 55: Metal resonance > wavefolder sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.

also assess pitch and activity over time, and t-SNE gives us distributions which are placed across the space and that form clear clusterings of similar-sounding sounds. An analysis of the eleven identified clusters can be found in Appendix 5.

As we see, there are a few ways we could categorise the many sounds that have emerged from this process. Because of the reasons and compromises I must make discussed above, I have to make a categorisation that is very rough. For the next step

of the analysis, I used the *BAR_463.wav* file from Cluster 10 where we find several elements: the hyper-distortion of the bar sample that starts as white noise, from which the strong partials of the sample emerge, and a percussive oscillation, followed by a synthetic bass tone; and *GIRDER_522.wav* from Cluster 4: a girder drop sample which is distorted in a middleground distortion, followed by the higher-pitched synthetic sound.

Without giving a full analysis of sound plots, for the other signal paths, the following intermediate samples were retained: first there is the *raw mic > dirt* step which is at the beginning of two signal paths where I kept *mic_crunch_78.wav* from Cluster 6, a microphone crunch that has been distorted to clicks; *comb_pop_129.wav* from Cluster 4, a comb pop that has been slightly distorted, but not in a particularly significant way; and *mic_gesture_142.wav* from Cluster 8, here there is a microphone gesture that seems to have been filtered or EQed rather than distorted.

Then, for the following *dirt > wavefolder* section of the third signal path, the following files were retained for the final iteration: *mic_gesture_142_726.wav* from Cluster 7, finding the distortion and the long synthetic bass notes, as well as a higher-pitched

bass note; *mic_crunch_78_416.wav* from Cluster 1 where we find a quick and clicky burst; and finally, *comb_pop_129_155.wav* from Cluster 4 which gives a distorted comb pop, followed by some clicks.

For the final step of the first three signal paths concerning the *cloud* module, to reduce the amount of computing power required, I decided to only use the first windowing option the module offers (the one used during the performance). This greatly reduces the number of possible states, passing from 6804 to 972.

After all this, I end up with a set of four sound plots, one for each signal path. These can be navigated using either my Max 2D sound plot, or my HTML interface. In my Max implementation, I also linked it back up to the effects modules to check that the process was working properly²³: the user can switch between playback of the recorded samples or live processing of the original sample through the effect module set to the parameters that were used to make the sample.

In the Chapter 5, these sound plots are used in a variety of ways: they can be explored as static plots, grouped together, grouped together with a plot of the original recording to examine where the performance went compared to where it could have gone, and even in a dynamic plot that shows where in the plot Constanzo is during

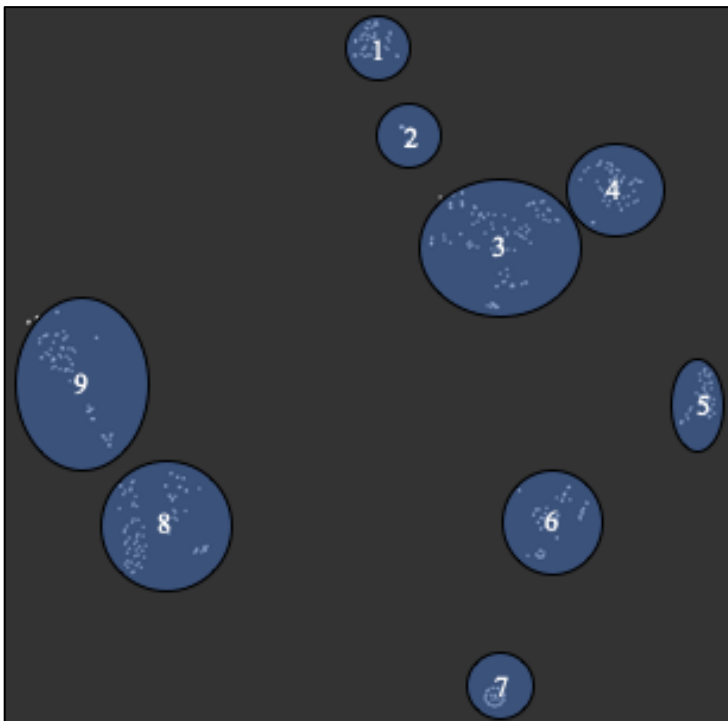


Figure 56: Raw mic > dirt sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.




Figure 57: Raw mic > dirt > wavefolder sound plot from MFCC audio descriptors and t-SNE dimensionality reduction.

²³ D.E.: 04_Video>01_Demos>sound_plot_fx_module_comparison.mp4

the performance. Cartological analysis of these elements is presented in Chapter 5, where I also discuss some of the limitations of this method; next, I go over three more iterations of this cartographical analysis.

3. The Alternative Performances Approach.

Next, I developed an approach that is more informed by the performance, that tries to inscribe elements from the organisation dimension – and questions of musical thought which are present in the network – into the sound plots of the sonorous potentialities. To achieve this, I decided to see if it were possible to gain access to alternative performances. This also addresses some of the issues with the previous, iterative method which demanded some quite scathing compromises.

This method gives a much more comprehensive and informed way of visualising the scope of sonorous potentialities. The method was achieved through tackling some of the limitations presented by the case studies: indeed, the three networks I look at cannot be reduced to a set of three effects modules and four signal paths. Their configuration is comparatively more complicated, rendering an iterative approach difficult. One could argue that iteration through the whole patch could be achieved, however, the more varied nature of the sound sources and the quantity of parameters across the whole patch would leave us with an unworkable number of different states.

This part of the methodology is illustrated across three case studies: first, an analysis of Hayes' work and an approach that gathers its source material from recordings of previous iterations of performance and rehearsals; then, I take Burton's work and the concept of recreating performances programmatically in the computer; finally I look at Pasquet's work, following the same process of computer-driven alternative performances while simultaneously trying to address certain issues around segmentation that arise in the approach around Burton's work.

a) Analysis of other recordings.

The first step of this alternative performance approach is to look at recordings of other performances and takes from rehearsals. This is an approach to case studies where automated or even human-driven recreation of performances is not feasible. This was the case with Hayes' work: the network that she constructed for the project included many external hardware instruments. Without all this hardware at hand, a programmatic recreation of her performance was not possible. One could argue that I could have approached her work in a fashion similar to Constanzo, taking

representative samples as inputs. I chose not to do this for two reasons: first, the variety of the different inputs was far too broad; second, the interconnected nature of her network meant that elements from the patch modified elements of the hardware in a feedback loop. What's more, it is useful for the development of the methodology to try other approaches.

Hayes provided me with a large collection of rehearsal recordings – over two and a half hours of content. She also provided me with some examples of the kinds of samples that some of her drum machines would trigger. I also had the recording of her performance. This meant that I had three corpora to work with. Working with these sources also pushed some of my tools to their limits – I first tried to deal with networks of over 31000 nodes. In some cases, my tools failed, and I was forced to find ways to work around these limitations.

This also informed me for further research regarding the way that I should try and build-up future tools: too late into the analysis process I realised that I was being limited, not necessarily by the number of nodes, but by the length of the audio files. I adopted a method from the FluCoMa help files that flattens audio down into one long file – this makes things easier for lots of the processing implementation; however, when trying to load files of over two and a half hours long into my various navigation tools, I encountered problems.

My system was already built around this structure: modification of this fundamental system would have meant rebuilding every step of the process. For future research, and for those wanting to implement tools like these themselves, I would advise constructing a system that can dynamically load many small files, rather than working with one large one. This will necessarily induce problems such as latency in navigation when clicking on nodes and hearing a sound, however, this kind of latency is less problematic in an analytical context than it is for performance or composition.

In the previous approach, I had not yet encountered the question of segmentation – I was manipulating sound files that were all considered as their own event. Here, I was faced with long recordings of audio, and segmentation by hand would have taken far too much time. I explored some of the possibilities for automated segmentation offered by the FluCoMa tools. After experimentation, I found that there is no one slicer that will give a perfect segmentation for all the audio content across a performance. I will try to address this issue further in the third iteration of this approach, however




Figure 58: Hayes full sound plot t-SNE dimensionality reduction.




Figure 59: Hayes full sound plot UMAP dimensionality reduction.

for this and the next iteration, I drew a compromise with the novelty slicer evaluating MFCC content, drawing its slicing when it judges that it has found a new event.

The MFCC descriptor is broad enough to give an informed idea of various aspects at once: timbre, but also gesture and pitch. I couple this with an MFCC description of the slices: this is a logical way to describe these fragments that were drawn from the same process. In terms of dimensionality reduction, I found that t-SNE, UMAP and PCA work well for the ways in which I wish to engage with these plots in Chapter 5.

I did my analyses in chunks²⁴. I performed four passes over the rehearsal audio, and separate passes over the sample and performance corpora. I then made some scripts that allowed me to join these analyses together and make them work with concatenated audio.

When I tried to use this combined picture in my navigation tools, my computer, once again, failed. This was frustrating, as the results looked extremely interesting. Indeed, I coloured the nodes so that audio from the rehearsals appeared in red, samples in green, and the performance in blue. The sample and performance nodes are quite evenly distributed amongst the rehearsal nodes – this is exactly the kind of sound plot

²⁴ When I tried to process all three corpora at the same time, my computer crashed.

I was aiming for. As examined in Chapter 5, one of the main methodological ideas is to observe what occurred at the same time as what could potentially have occurred. Unfortunately, these plots of over 31000 nodes were closed-off to me, and I had to find other solutions.

I realised again that my only option was to approach my sources in chunks. I built-up four combinations where the sample and performance data were added to each of the four rehearsal passes. This worked, and what's more, the visualisations and the clustering seemed to have worked well. Some of the segmentation left something to be desired – some of the events were very short – however, there was enough material for me to be able to draw a meaningful analysis.

I now had several sound plots that showed what had occurred in performance, some of the source material for sound generation, and elements from other performances. In each of the visualisations, the performance material was distributed across the space, suggesting that there were many more things that the instrument *could* have done. In Chapter 5, I explain my approach for engaging with and analysing this data.

b) *Automated performance – John Burton.*

Next, I attack the idea of alternative performances where the analyst has control over certain parameters. I begin with the case study of Burton's work which lends itself well to this approach. In his patch, a score is generated in two parts: one for indicating and recording human input; one for triggering playback of this input through four main processing units. Many things in this network are generated with weighted randomness that was defined using a set of *function* objects.

For this analysis I wished to see how the performance could have happened differently at several levels: if the same human and machine scores were created; if the same human score but a different machine score was created; and if a different human score and a different machine score were created. At each level, I move further from what occurred during the performance²⁵. I intend to see to what extent things really change, and if there are things that remain the same.

The primary question here is how to set up an alternative performance. This follows methods suggested by Baudoin's *faktura* approach (Baudoin, 2009): following steps of *transcription*, *concordance*, and *exploitation* to recreate a piece in analytical conditions.

²⁵ There are even differences to be remarked at the first level, as the choice of which parts of the recorded input and the types of machine processing are also randomly generated.

Here, the main issue is that the input sound was generated by recording of acoustic instruments. Burton exported the recording buffer of the input after his performance and give it to me as a source for analysis. I split this into the various files that would correspond to each block, using the human score data which he also provided me with. I wrote a Python script²⁶ which took this data and split the audio, giving me a useable version of the input corpus from the night. Next, I set up a Max patch that would play the right segment at the right time as the piece played out.

Next, I needed to make sure that Burton's original patch was working correctly on my machine. Modifications were that needed for each of the levels. These are described fully in Appendix 6, and there is also a video demonstrating them²⁷.

With my system in place, I started through my iterations. Contrarily to the work I did around Constanzo's piece, I had to perform these in real time. This was to be able to capture the video of the score as it was being created.

Considering that there was video involved, I wanted to see if there could be a way of examining these generated sources where I could inspect the videos simultaneously. I built a patch²⁸ in Max that could play four videos simultaneously, all of them synchronised to be playing the same part of the piece. Clicking on one video mutes the audio from the others: for each stream of playback, the user can choose between the many audio streams I collected when recording the performances (full audio, only input audio, full processed machine audio and a stream for each of the 16 processing units). I ultimately didn't use this tool much for analysis, but I do consider it to be a useful contribution. A video demonstration²⁹ can be viewed.

I took these results and followed the same process of corpus manipulation as with Hayes' work. In my visualisations, I coloured each performance differently which reveals that slices from a same performance tend to cluster together; but again, things do tend to be well distributed. I also constructed some visualisations that put slices from different levels in the same visualisation to examine the differences between them. My cartological analysis of these artefacts is taken-up in the next chapter.

²⁶ D.E.: 01_Code>02_Tools>03_Burton_Score_Transcribe

²⁷ D.E.: 04_Video>01_Demos>burton_alternative_performances_setup_demo.mp4

²⁸ D.E.: 01_Code>02_Tools>07_Burton_Video_Player

²⁹ D.E.: 04_Video>01_Demos>burton_video_player_demo.mp4

c) *Automated performance – Olivier Pasquet.*

Finally, I present the last iteration of this method with my case study of Pasquet's work. There are two notable differences between this, and the work performed around Burton: first, here everything is generated by the software, meaning that I can access a completely new and comparatively complete generation of the piece each time; second, here I examine the question of segmentation that has been a limitation over the previous iterations of the method.

As useful and general as the MFCC novelty segmentation can be, it has two major problems: sometimes what can be perceived as a musical event may be cut-off mid-way; and many of the slices are extremely short in duration – some only tens of milliseconds – rendering them difficult to deal with in a musicological manner. We can attempt to remedy this by lowering the threshold of the novelty slicer, however, then we find ourselves with slices which are too long.

In addition to this, these slices may steer us away from the segmentation that the musickers themselves make of their pieces. Indeed, though the flow of audio, it is difficult to discern what collections of gestures the performers consider comprising a slice, especially in the absence of a score. With this analysis of Pasquet's work, I wish to address this question by looking to the network and seeing if it is possible to find this kind of segmentation – this musical thought – inscribed within it, and to apply it to my analysis. It happens that this kind of segmentation is transparent within Pasquet's software.

It was first necessary to go through the same process of making sure that Pasquet's software ran correctly on my machine. The only issue was running a 3rd party VST, the Kick 2 kick drum virtual instrument. This required a hacky workaround to make sure that I was using the same settings as Pasquet: I had to open the original version of the patch in a text editor, find the *vst~* object, and copy this into my version of the patch, modifying the *text* field of the object to add the path to the VST. In his original patch, Pasquet had not included the path name in this way, therefore, when I loaded the VST on my machine, it would reinitialise the saved state of the VST. With this workaround, the settings that Pasquet used during performance are retained.

With this in place, the other point to address was that of multi-channel audio. Pasquet's patch outputs audio across twelve channels. I debated whether to carry out my analyses in stereo or multi-channel. I decided on a compromise: I recorded both

the multi-channel and a stereo fold-around output and performed the descriptor analysis on the multi-channel audio but played back the stereo output when navigating the networks. The fold-around solution gives the user the choice to re-route the channels in either a 6-6 mapping or alternating each channel from left to right. I used the 6-6 configuration.

For the iterations of performances, I examined three different configurations: the first uses all the same parametric settings that Pasquet used during performance; the second uses the same parameters with different rhythms to generate the main tree³⁰ – I experimented with 1-1-1-1 and 2-1 which are the two examples that Pasquet gives in his help file for *jtol.bach.randtree*; finally, a configuration that randomly changes the main parameters during the piece: all of the parameters controlling the four Modalys instruments, and parameters for tree generation³¹.

To incorporate Pasquet's musical thought around gestures into the slicing of my sources, I used the beginning and end of playback of the main generated tree, which drives playback of everything else. It can be argued that this represents a gesture, or a coherent musical event for Pasquet: at the beginning, everything is generated, the gesture is played, and then a new gesture is generated, with a whole new set of generated parameters. I wish to examine how it will affect my analysis to take these slicings for my visualisations, rather than a novelty slicer or any other automated slicer. My hope is that considering this strand of musical thought which is inscribed within the network will grant us a better-informed visualisation.

Playback of a piece iterates through the 25 preset states that Pasquet defined. Pasquet implemented a timer on each state change that sets off a bang roughly 25 seconds after a state has been changed. I implemented a system that plays gestures until this bang is flashing – when the bang is flashing, the patch changes to the next state on the next gesture end. I performed 5 passes of the *same settings* configuration, 2 passes each of the *1-1-1-1* and *1-2* configurations, and 2 passes of the *random parameters* configuration. The reader may consult a video³² demonstrating these processes.

³⁰ Pasquet used the 1-2-1-1 proportion in his original patch.

³¹ I created a patch that randomly chooses a value between a given minimum and maximum, waits to change a random amount of time between 3 and 20 seconds, and changes the parameter to the target over a random proportion of this time. I wanted to create an emulation of how a human might change a parameter.

³² D.E.: 04_Video>01_Demos>pasquet_alternative_performances_setup_demo.mp4

I generated a large number of different sound plots – all using the slices generated in source collection and describing with MFCCs and the collection of spectral shape descriptors. I used the spectral shape descriptors as I was dealing with significantly longer slices and believed that it would be important to grasp a picture of gesture and rhythm over time: this is something that the spectral shape descriptors could potentially transcribe better than MFCCs which, while still possibly informing us on this front, are inherently better suited for describing general timbre.

The different sound plot configurations are as follows: all of the *same settings* configurations together; each of the 25 different sections for all configurations together; 2 global comparison plots that group together the first and then second passes of each configuration; the two rhythmic configurations together grouped with the first two same settings configurations; and the two random parameter configurations together, again grouped with the first two same settings configurations. Each plot was described once by MFCC content, once by spectral shape descriptors, and each time this was reduced using t-SNE and UMAP for a total of 120 sound plots³³.

These analyses have given me a vast number of new artefacts to analyse. In Chapter 5, I discuss how I approached analysis of these plots in a methodological way. One of the things that the computer can offer the analyst as a tool is a way to deal with large amounts of source data – however, the analyst must also be wary not to go too far when given a tool that also allows them to create vast amounts of source material.

Indeed, it must not be forgotten that the aim is to be able to deal with large bodies of data, not to create a situation where I am even further submerged. Until now, my approach has typically been to produce as much as I can, with the idea that nothing will be missed. My job now, as an analyst, is to identify what data will inform my analysis, what data will not, and decide how to approach this data in a methodical way that will inform my analysis as a whole.

³³ I created a small Python script to speed up some of the process of constructing the HTML page, given the large number of the to be created: see supplementary multimedia appendices: https://github.com/jdchart/supplementary-multimedia-appendices/tree/main/Finished/06_PhD_Browser/PhD_Browser/tools

PART TWO: CARTOLOGICAL ANALYSES

Chapter V. NAVIGATION OF NETWORKS.

1. Static Sound Plots.

Now that my networks and sound plots have been constructed and grammatised to a certain level, I can begin to approach a cartological analysis: an analysis which wishes to draw meaning and signification from these plots, an analysis which wishes to understand the structure of these networks. There are two types of networks: the first set constructed in Chapter 3 are *unspaced*, the positions of the nodes and edges have physical spatiality, but their representations do not; the second set, the sound plots from Chapter 4, are *spaced* – the nodes do not have a physical spatiality, however their representation does.

I begin by examining the question of navigation of *spaced* networks – what are the methods we can deploy when interrogating, exploring, and navigating a place that has meaningful structure across a number of dimensions? The question of time is also inherently interwoven into my network visualizations – some are *static*, representing a unidimensional representation of time, or elements that are out of time; others are *superimposed*, representing multiple temporal dimensions at once. This is another dimension that the analyst needs to consider when approaching the navigation of a network. I start by discussing some methodological principles for network navigation, then illustrate them in action through various examples from my case studies.

a) *Principles of Navigation*

My approach to analysis is assimilated to cartography, suggesting that we can look to the real world of cartographic practice to borrow principles for analysis. In this section, I have drawn five principles from the observation of three practices: cartography, rambling, and video games. In addition to this, and to return to some of the ideas explored in Chapter 1, I shall examine these principles through the prism of some of Ingold's ideas around the notion of *wayfaring* (Ingold, 2011).

For the purposes of this section, we can understand wayfaring as a certain mode of existence and of conceiving of the world. Ingold explains that the paths we trace through wayfaring, through movement from place to place, are what create the *meshwork* discussed in Chapter 1. Existence is understood through this movement – rendering existence a process, a sensory and engaged act – rather than the momentary

pauses that occur in *places*. This he opposes to *transport*, where existence is – for Ingold, regrettably – understood through the prism of these places, these sources and destinations. He proposes that it is unfortunate to configure one’s perspective of the world around these places, which are but mere pauses in our wayfaring, which are far removed from what existence is (Ingold, 2011, p.145-155). As I explain in Chapter 1, I am not entirely content to adopt this model as it remains an essentially *retrospective* perspective on the social: however, it is useful to take this polarity and interrogate the questions as to why and how we would look to map our world.

The first principle I wish to discuss is that of *abstraction*. In cartography, it is understood that the representation the cartographer is looking to make is an abstract version of the real, physical space it visualises. A map can be detailed, but it will only ever *indicate* the reality it represents. This is something that has been discussed during network construction – when building the visualisation, it is necessary to make choices regarding the *resolution* of the representation, meaning that atomic details will be absent or understood as composite elements of what is present. This infers the notion of abstraction: when faced with a representation, it is understood that this is an abstract object that aims to access *essential* elements of what it represents.

This principle of abstraction, and notably what elements are deemed to constitute the essence of an environment, will vary according to the person who is drawing the map. In video games, developers create worlds in computers where only what is essential to meaning and signification is produced: in the 3D universe of an open-world game, we navigate through an abstract representation of a world. We do not represent the soil beneath the earth as this is not something that is essential to draw meaning from the visualisation – it is understood to be a composite of the distorted plane that creates the ground.

A Rambler may not be attentive to every atomic detail they pass – they may conceive of things at various levels of abstraction: they may see a tree, not a collection of bark, leaves, branches, and moss; a path, not a flux of pebbles; a hill, not soil, water, rocks, and grass. Inherently, in our perception of the world, we construct abstract *measures*, grammatisations of the world. It is important to keep this idea in mind.

This principle, this essential aspect of the map object means that we must interrogate ourselves as map makers and users about the kind of knowledge that is being represented, the kind of knowledge that is being drawn from the environment. This is reflected first in the second principle: the principle of *clustering*, of grouping, of

grammatisation. This could be assimilated to Hanninen's notion of *associative sets* (Hanninen, 2012). I have discussed this across the thesis, and it prolongs the idea of abstraction. To make sense of an environment, it is necessary to identify which elements compose disparate objects. For the cartographer, some of these clusterings are made for them: for example, the artificial boundaries of countries that are an important part of a map's structure. Other things are more delicate: for example, when does a forest end? Is a slightly isolated tree part of that forest or another object in itself? When does a city end and where does the countryside begin? The regions that the cartographer must draw are things that inherently overlap and undulate over time.

The idea of groupings and clustering also goes beyond spatial position and interrogates the ontology of the elements themselves. Indeed, these clusterings are not elements that pre-exist the elements themselves but *emerge* through their wayfaring. In video games, we often encounter maps with filters that allow the user to engage with elements that are grouped together by their ontology – entities that are the same do not necessarily group themselves together in space, however, their grouping by ontology may give their spatial positions meaning. Indeed, the very idea that elements stemming from a same process are *not* together in space leads us to interrogate *why* they have adopted a certain spatial configuration – is there meaning to be drawn not only from proximity, but also from distance?

Another type of knowledge can be understood in a third principle of *landmarks*. When one goes out rambling, they will often think of a landmark – “I'm going for a walk around Castle Hill”, “I'm going to walk around the South Downs”, “I'm going to climb up that mountain”. We conceive of environments and places around notable objects, usually things that will dominate the surrounding area, that stand out through their large scale.

They condition our moving through these environments: once a rambler reaches the summit of a landmark, they will stop and rest; if they reach a viewpoint that offers a view over a landmark, they will stop and look; paths are created around landmarks, leading to and from them. Cartographers will clearly mark landmarks on their maps. Game developers will construct worlds around them: for example, the linear levels of a game like *The Last of Us*¹ have no maps, but make use of the notion of *signposting*, where the player's destination is revealed at critical points in the distance and the path

¹ *The Last of Us*. Naughty Dog, 2013.

is indicated, driving the player towards it. It can be imagined that these landmarks would reveal themselves quite transparently and have a conditioning effect on the network around them.

The idea of landmarks goes somewhat against the ideas that Ingold proposes; they can be conceived of as sources and destinations, as places. These pauses in our wayfaring are not something that can be completely evacuated from our consideration: indeed, without these points and knots along the paths of the meshwork, there would be no relationship between entities. If not taken as an end unto themselves, they can help us a great deal to understand the *motivations* of our wayfaring. That being said, it remains imperatively important to examine the wayfaring of the elements in a map, hence the next principle.

The fourth principle is that of *movement*: movement as wayfaring rather than transport. An environment demands that we move through it in certain ways – in video games, actors are configured in a way that pushes the user to move through the world following a certain path. In the networks I wish to examine, this notion of movement is an interesting one. There are certain movements that are inscribed within a network – for example, in a sound plot, one can suppose movement through the space according to the order in which fragments were heard during a performance. There is a structure in the movement through time across the network. This movement can be subjected to analysis.

It also begs the question of other paths that may be traced through the space. Elements of an environment's structure will be the condition for certain types of movement: roads tend to be constructed following the easiest route – why carve through a mountain when it is possible to circumnavigate it? Similarly, in musical networks, their materialities will condition some of the movements that are possible throughout them. I examine this idea in Chapter 6.

Finally, there is the fifth general principle of *ontology*. This could be assimilated to Hanninen's (Hanninen, 2012) idea of *genosegments* and *phenosegments*. An environment is filled with different kinds of elements; it is important to identify them and understand their differences. A map is usually accompanied by a legend that gives a typology of the different kind of elements that can be found in the environment. In a video game, different kinds of elements that fill the world serve different functions. These elements manifest themselves differently at different times – a rambler may know an environment well and be used to the structure of the place.

However, in winter, the trees will take on a different aspect, or the fauna will behave differently. Each of these elements can potentially have an infinite number of ways it could manifest itself across multidimensional hypothetical timelines.

These are general principles, drawn from fields that deal with space, and representation and exploration of space. In this chapter, the first part of my cartological analysis will take some of these principles in a somewhat literal manner; as analysis continues over the following two chapters, they will become more conceptual.

b) Static sound plot analysis.

Let us take an example of a *static* sound plot with Constanzo's work. This first sound plot² is a simple visualization of the general segmentation of the performance audio discussed in the previous chapter, without any of the iterative material. Its temporal dimension is simple: a unidirectional account of what happened at that time. The main body of our analysis will concern superimposed networks: this analysis essentially serves as an introduction into the approach.

The segmentation was made semi-automatically, using the modified ReaCoMa tools to perform an initial novelty MFCC slicing which was then corrected by hand. The description was done with MFCCs and the dimensionality reduction with the t-SNE algorithm.

² D.E.: 01_Code>02_Tools>05_Sound_Plot_Navigation_Standalone>Constanzo: Performance Segmentation>mfcc_TSNE.json

Broadly, I have identified two general vectors across the space: from top to bottom there is a clear movement from extreme saturation and distortion to clear, dry sounds. There is also a vector going from the right with clearer, more sinusoidal sounds towards nosier, more chaotic sounds on the left. At a more detailed level, I have identified 11 different clusters:




Figure 60: Kaizo Snare general segmentation sound plot described with MFCCs and displayed with t-SNE dimensionality reduction with clusters.

- 1: harsh, scraping against the snare surface.
- 2: progressively softer rubbing against the snare surface.
- 3: very soft sounds, slices of resonance, and light crushing of felt pad.
- 4: dry impact sounds, lots of resonance and silence.
- 5: various resonances.
- 6: resonances once more, with synthetic-sounding sounds of the type found with the previously seen wavefolder iterations.
- 7: general scratches of varying velocities.
- 8: more scratches of varying velocity, with a notable presence of lower frequencies.
- 9: long, gestural, noisy scrapes.
- 10: more general impacts which are dry. Here, there is a clear presence of low frequencies.
- 11: more general scratching noises which get more and more distorted towards the top.

From a methodological point of view, the descriptions of nodes are tightly interwoven with the way the sound is produced. There are several ways of approaching sound analysis – one is the spectromorphological (Smalley, 1997), reduced listening (Schaeffer, 2017) approach which would describe sounds clinically regarding a grid of possible spectral aspects. Another is to examine the source of the sound, an approach that would incorporate the notion of gesture which can sometimes be forgotten in the former approaches. I take both approaches into account – I examine the scope of sonorous potentialities, therefore a spectromorphological vocabulary is useful in this context; however, I am also examining the network as a configuration that allows for sounds to emerge, meaning an examination of the source is important.

With this clustering I examine the plot through the principles of clustering and ontology. What of the other principles? Contrarily to some other plots I examine, there are no apparent landmarks in this network. This reveals notions around the abstract structure of the network – the morphologies of sounds merge quite uniformly into each other. There is no group of sounds that is isolated from the rest which would indicate that parts of the network produce starkly different noises to others.

In Chapter 2, I proposed the idea that Constanzo examined distortion in various forms, and that it could potentially be a structural aspect of the piece. The structure of

this plot could suggest that, indeed, Constanzo attempts to cover a broad spectrum from completely dry to heavily distorted sounds. The nodes are distributed evenly across the space, with no one concentration of sounds in any particular spot; the borders of the various clusters I have made are relatively porous. It is perhaps strenuous to suggest that Constanzo has examined distortion evenly across all its forms, but this plot certainly demonstrates that it was incorporated to a large scope of degrees across the piece.

The idea of movement is examined in Section 3 of this chapter where I produce a visualization that allows me to examine where in the space we are as the performance progresses. With this initial static analysis, we do learn things about the nature of the timbral and gestural space that Constanzo navigated – however, I wish to go further and see if this method can allow us to understand the whole scope of sonorous potentialities of a network.

2. Superimposed sound plots.

a) *Rodrigo Constanzo.*

I continue with further analysis of the sound plots produced through analysis of Constanzo's network and the iterative approach. I have already discussed some of the intermediate steps of the signal paths, now I can examine some of the final results.

I met a limit when attempting to combine results from my iterative approach with the general segmentation of the performance. In Figure 61, we find a composite network including this segmentation, the results from the three signal paths and the three original corpora: the general segmentation is in white, the raw microphone corpus in red, the gamelan corpus in pink. The metal resonance corpus is in orange, the *gamelan>wavefolder* results are in green, the *metal resonance>wavefolder>cloud* results are in purple, the *raw mic>dirt>cloud* results are in light blue and the *raw mic>dirt>wavefolder>cloud* results are in dark blue.

Here, the dimensionality reduction seems to be working too well – the results of the various iteration processes are distributed evenly across the plot; however, the general segmentation of the piece occupies its own cluster with very little intersection with the other corpora. It can be supposed that differences in the spectral aspect of sounds recorded with the room microphones compared to the pristine sound captured in the computer are too great. For future research, a possible solution could be through filtering of the audio to try and recreate the same spectral aspect, but a better solution




Figure 61: Kaizo Snare entire composite visualisation displayed with the UMAP dimensionality reduction.

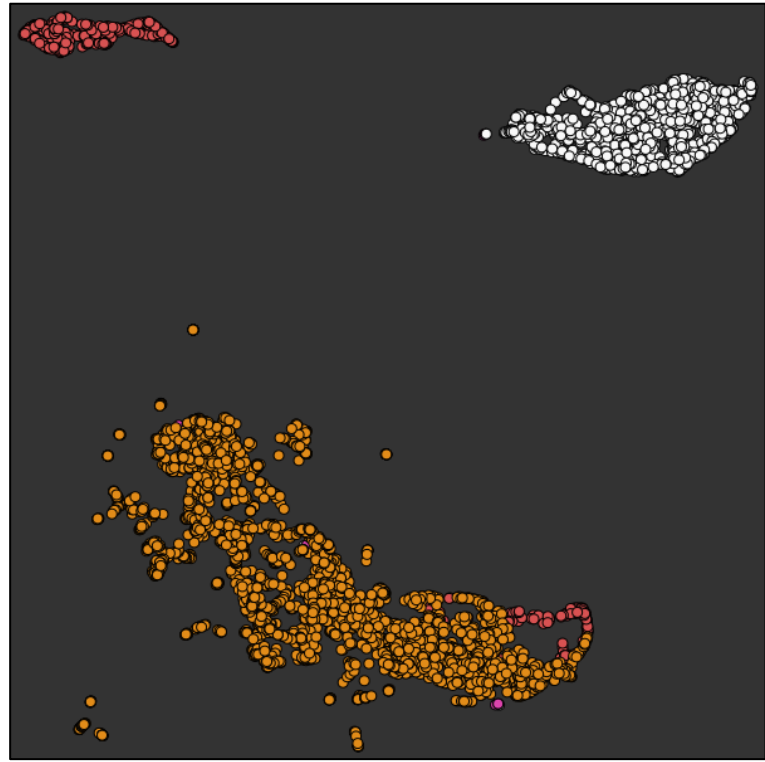


Figure 62: Kaizo Snare corpora composite displayed with the UMAP dimensionality reduction.

would appear to be to proceed along the alternative performances approach where the analyst has full control over every aspect. Despite this, we can still draw some interesting ideas from these results.

Figure 62 displays an image showing just the three corpora and the general piece segmentation. There is very little intersection between the metal resonance (orange) and raw microphone (red) corpora. Both corpora remain grouped together in tight clusters, except for a small strand of the raw microphone corpus being found among the metal resonance. The small gamelan corpus is engulfed within the metal resonance corpus – this is to be expected as the corpus is comprised of metal objects being impacted.

One idea that we can draw from this, is that there is a great difference between the sample libraries in the computer and the sounds that are coming in from the snare surface (with a few exceptions). There are, essentially, two bodies of sound profiles. In the entire composite network displaying the results of my iterative analyses, the results of the signal paths are bridging the gap between the two clusters – especially the *raw mic>dirt>wavefolder>cloud* signal path in dark blue. Indeed, most of the results can be found between the two groups. This is intersected with the *raw mic>dirt>cloud* signal path in light blue and *metal resonance>wavefolder>cloud* signal path in purple,

which also stretch far beyond these groups. Finally, the green sounds, corresponding to the *gamelan>wavefolder* results tend to skirt around the two groups, and around the edges of the plot.

This can inform us about the nature of these signal paths, and the nature of the *Kaizo Snare* network. For example, the two blue groups of sound emerge *from* the raw microphone corpus and tend *towards* the metal resonance corpus. This could indicate that the *dirt>wavefolder>cloud* effect configuration is intended to manipulate the incoming sounds into sounding like the metal resonance corpus, and underlines the importance of the *wavefolder* effect which, being absent from the light blue signal path, causes a distribution of sounds which is much more broadly distributed across the space, being less concentrated towards the orange sounds.

Similarly, we can argue for a vector moving from the metal resonance corpus towards the red nodes, corresponding to the signal path that treats the metal resonance corpus with a *wavefolder>cloud* configuration. Here there is an absence of the *dirt* effect, leading us to believe that this is not necessary when looking to transform sound from the metal resonance corpus towards sounds that resemble the raw microphone. These vectors do correspond to the input source of origin in their signal paths, but how much can we allow ourselves to draw from this representation?

Indeed, it is an intriguing artefact, however, the results are unverifiable due to my navigation system not being robust enough to handle what is well over ten hours of audio. Furthermore, there are some issues to address in the collection process of my iterative approach. Upon closer inspection, looking at results from the two blue corpora³, we realize that there is very little difference in timbre from cluster to cluster. The clusters seem to be defined according to the source audio file, the only real differences in timbre being some quite inconsequential differences in distortion and filtering. This is probably due to the nature of the iterative approach: the filtering effect of having to choose representative slices of sound at each step has got to a point where the differences in timbre are minute.

Indeed, with each pass of this iterative approach, the scope of the picture that can be conceived becomes more restricted. This is the fundamental flaw of this approach, and one that is circumnavigated with alternative performances. This does not mean that

³ D.E.: 01_Code>02_Tools>05_Sound_Plot_Navigation_Standalone>Constanzo: raw mic > dirt > wavefolder > cloud; 01_Code>02_Tools>05_Sound_Plot_Navigation_Standalone>Constanzo: raw mic > dirt > cloud

the results I have obtained are completely meaningless – there is content from the previous layers of iteration that remain and structure the plot. However, it does strike us that a big part of the picture may be missing.

The starting point of the alternative performances approach was to address this problem, while simultaneously gaining a better-informed plot that is also structured by the musical thought that was present in the network. It offers systemic solutions for the technical limitations of having to deal with large amounts of audio – it is an approach where meaning can be drawn from much smaller sets of data. Let us now examine the artefacts derived from this approach.

b) Lauren Sarah Hayes.

I created a series of sound plots that display a combination of MFCC novelty slicings of a series of rehearsal recordings that Hayes gave to me, several examples of the types of samples she used, and an MFCC segmentation of her performance. The main goal here was to visualize what Hayes played against the rehearsal recordings, and as we see across the four sound plots, the recording is quite evenly distributed across the network. I begin with an analysis of the first pass.

Both the t-SNE and UMAP dimensionality reductions offer interesting artefacts, where distribution seems to

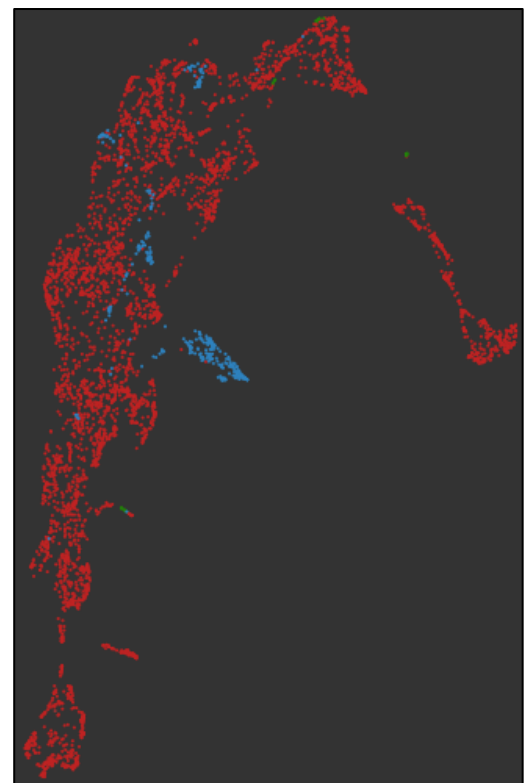


Figure 63: Pass 1 of Hayes' work displayed with a UMAP dimensionality reduction.




Figure 64: Pass 1 of Hayes' sound plots using a t-SNE dimensionality reduction with clusterings.

have formed similar groupings. I worked with the t-SNE document as the clusters are not quite so tightly distributed, making it easier to work with.

The full cluster analysis can be found in Appendix 7, and the network can also be navigated by the user⁴. A lot of the data gathered can seem redundant. Indeed, we do get a clear picture of some of the timbres that are used in the piece and those that are not – for example, the highly synthetic, distorted, and compressed sounds from rehearsal do not appear to be present in the piece; likewise, the looped vocals below chaotic low-pitched percussion is not found within rehearsal. However, lots of the segments are very short, and some of the gestural nature of Hayes' network is lost. We see that the automatic slicing used to make this plot – slicing according to a mathematical model – will struggle to make slicings of a more conceptual nature like slicing by gesture (I return to this problem in the following section).

Another conclusion we can draw from this plot, is that in the performance, we do not hear any dry, acoustic sounding segments, like the ones present in Cluster 1 from rehearsals. There is a definite abstract structure to be found within the plot: there appears to be one main, central continent around which gravitate islands of sounds and from which sprout small peninsulas.

From left to right, this continent broadly moves from resonant sounds towards short bursts of distorted synthetic noise. A large portion of the central continent is populated mostly by rehearsal content, with segments from the performance being found more towards the left, resonant side. The short synthetic parts of the performance are found towards the right, but isolated from the main continent. This reveals that a whole region of distorted, noisier content was not explored during the performance, but that the instrument has the capacity to produce it. Finally, I found no slices in the rehearsal or performance content that give exact renditions of the acoustic samples – indeed, they are completely transformed or reproduced synthetically.

⁴ D.E.: 01_Code>02_Tools>05_Sound_Plot_Navigation_Standalone>Hayes: Pass 1

I won't give such a detailed account of the other sound plots, but I will present some of the main points to take away. The second pass, this time analysed with a UMAP dimensionality reduction, is much more spaced-out. The segments from the piece can be found towards the bottom, and skirt around the border of a cluster of rehearsal segments that is comprised of synthetic sounds that get brighter and more harmonic towards the bottom right. These longer, higher-pitched, and partial-rich synth notes contrast with what was heard during the performance, which remained much murkier in nature.

The rest of the rehearsal segments are found much further away from the performance. The cluster found immediately above the performance segments is mainly comprised of dryer, acoustic sound bursts with a notable presence of dry vocals. The length of these segments, however, is very short, leaving us in doubt about their gestural content.

This is not the case for the arch-shaped cluster towards the top, with long gestures of glitchy vocals that are different in nature to what was heard in performance. Here, the vocals seem to be on their own, chopped-up and rearranged, and many of them have been sped-up so that the pitch is greatly increased. The only part where this kind of thing was heard in the performance was in a much more rhythmic way, with pauses between the vocal bursts. In this cluster, we also find some samples present, the gestures of which seem to be imitated with the processing. The other, outer clusters in this plot are redundant, their length being extremely short.

The third pass was much more successful in terms of slices and presents many clusters of

sound that are much more gestural in nature. Here, we really begin to see the scope

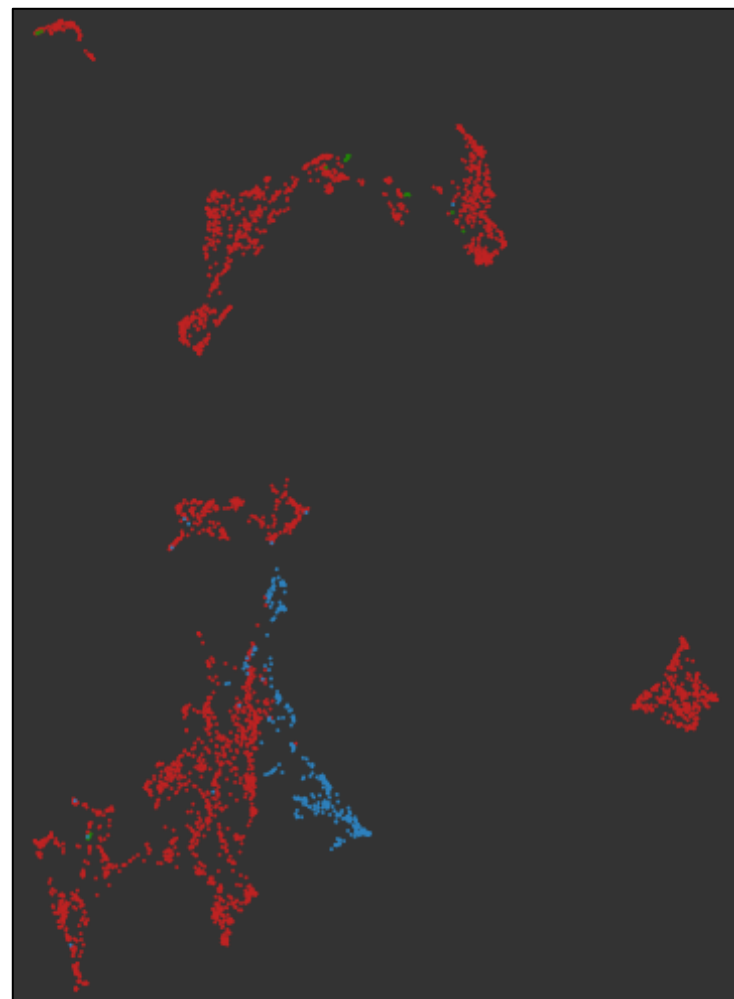


Figure 65: pass 2 of Hayes' work displayed with a UMAP dimensionality reduction.

of the instrument and understand the extent of its versatility. I have highlighted some clusters which are particularly interesting and diverse:

- 1: a collection of synthetic kick sounds with loud bass frequencies. It intersects the performance segments at the point that can be conceived of as Clusters 5a-5d from the first pass (short bursts of vocals or synth that are introduced with a heavy click or percussive noise). In the cluster we find kick sounds of many different sorts with differing gain envelopes, pitches, and levels of distortion.
- 2: a large cluster comprised of vocals that have been heavily processed. Notably down the middle spine of the cluster, they are looped very rapidly, and there is a clear mid-range, stable *d sharp* pitch that rings out across the cluster.
- 3: very sharp percussive sounds of many sorts. The bottom left isolated cluster is comprised of very high-pitched clicks and samples of wood blocks – some of these clicks have a long, reverberating resonance after them. The

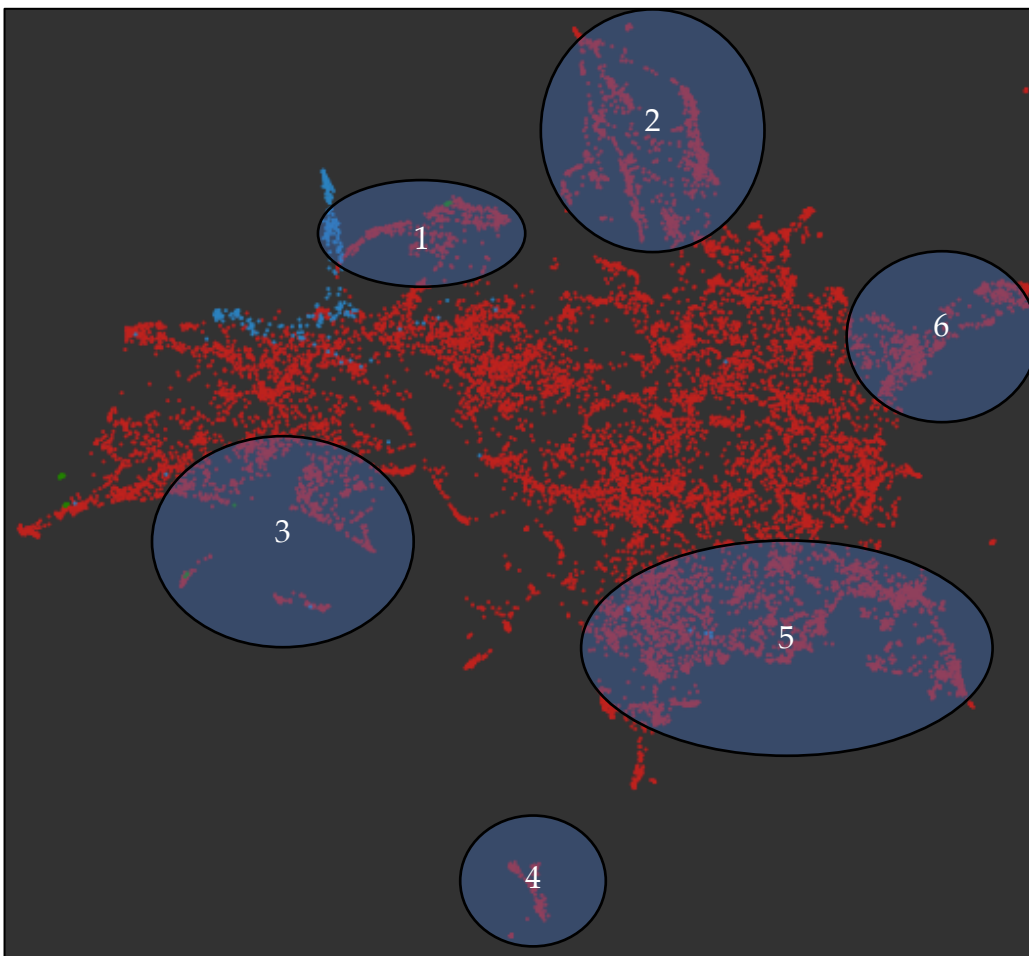


Figure 66: pass 3 of Hayes' work displayed with the UMAP dimensionality reduction.

other isolated cluster is sharp hits of what sounds like a classic synthetic cowbell. The part of the cluster attached to the main body of the plot has various sharp clicks that get progressively duller higher into the plot.

- 4: an isolated set of synthetic scratch sounds that play high pitched *a sharp* to *b* notes.
- 5: lots of vocal segments and synthetic segments where clear syllables can be heard. From the bottom up there are lots of *ts* syllables, moving towards the top right we hear more *ah* syllables, towards the top we have more *ee* and *eh* syllables. There are three slices from the performance here which fall in the *ts* section.
- 6: many gestures of what sound like synthetic percussive samples that do not have much processing on them, but that are looped rapidly.

Finally, the fourth pass. Here, the performance segments are surrounded by a large cluster of rehearsal segments. From left to right across this cluster, there is a real trajectory of pitch going from low to high, indicating that the performance tended to stay around the mid-range of the network's potentialities. From top to bottom, there is also a clear trajectory from clearer sounds down to chaotic, extremely noisy ones, finishing at the tip with what sounds like various synthetic percussions being looped at an extreme rate. This can indicate that Hayes also stayed on the clearer side of the networks' potentialities.

This is where other questions of agency can intervene: why not push the instrument into these extreme limits of pitch and of spectral aspect? Did Hayes choose to remain relatively tame? Did the context of the prestigious HCMF, or the large audience offered by the BBC recording have any kind of sway in this? It's entirely possible, and the decision may have been a conscious or subconscious one. However, it's also worth considering that during rehearsals, Hayes was possibly trying to test out the limits of her network, with the specific goal of finding where she wanted to go in the performance, and where she did not.

As I have discussed, the interconnected nature of her network renders it a difficult instrument to steer and control. We can imagine that despite her statement about




Figure 67: pass 4 of Hayes' work displayed with the UMAP dimensionality reduction.

wanting to create a certain fragility in the system her goal was to be able to impart a certain level of agency within it.

This kind of approach inflects our perception of the piece and some of the musical thought that is inscribed within the network. There are still some limitations that have been encountered: notably the amount of very similar and sometimes redundant data. It is also still possible that the differences between room recording of the performance and the computer may be

biasing the sound plots, even though it didn't seem quite so apparent in this case. Next, I examine a case where everything was reproduced in the computer.

c) *John Burton.*

In Chapter 4 I presented three different configurations for passes of alternative performances: both scores staying the same, the same human score but with a newly generated machine score, and both scores newly generated. These were used to produce a series of sound plots, here I will look at: the five passes of the first configuration combined with the input content; the five passes of the second configuration combined with the input content; the five passes of the third configuration; and a combined plot of the first two passes of each configuration.

Looking at the first plot representing the five passes of the generated performances with the same score, unsurprisingly, each of the five colours (one colour per pass, the grey representing the input material) is represented evenly across the space. Apart from the bottom right-hand corner where there is a string of slices from pass three

protruding out of high-pitched drone sounds, it seems that there is no place of the plot where one pass is present more than others.

This allows us to suggest the conclusion that the aspect of the score is deterministic not only in the configuration of sounds across the organisational dimension as the piece progresses, but also for the aspect of timbres across the performance. It would seem, then, that the random distribution of voices between the four processing units

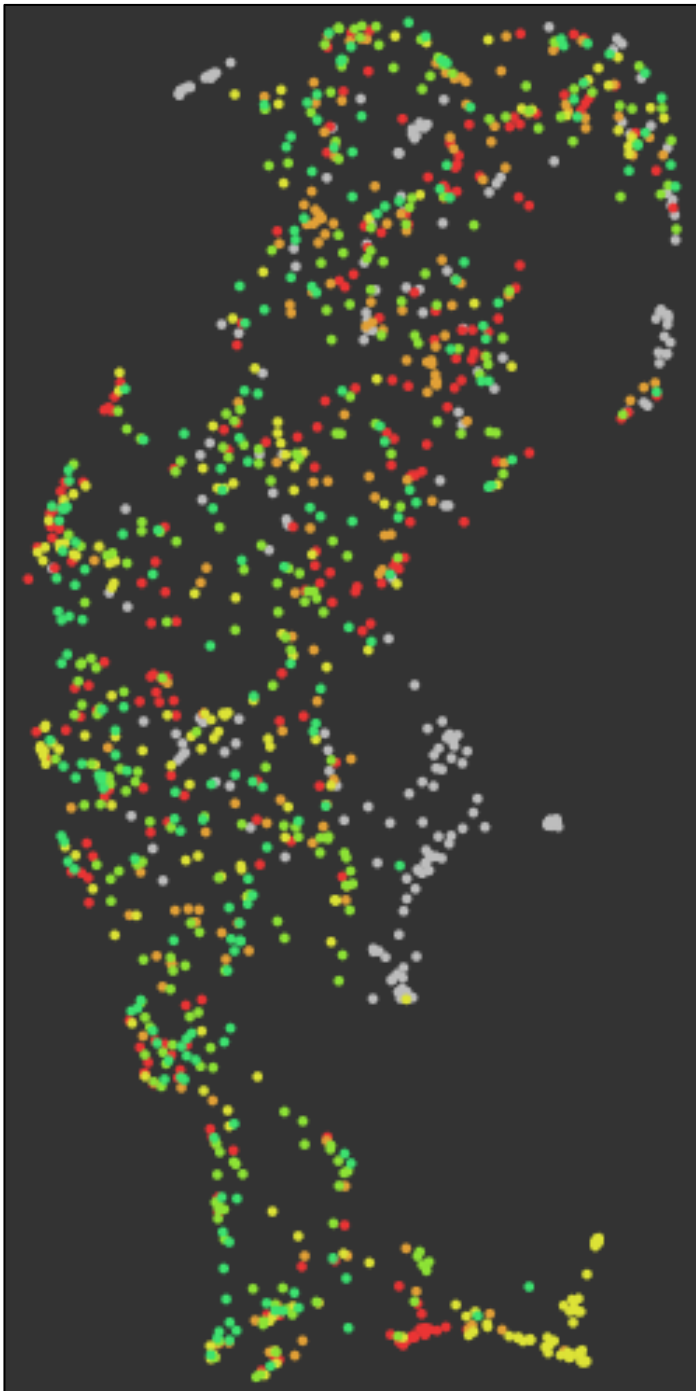


Figure 68: same scores sound plot for Burton's work displayed with the UMAP dimensionality reduction.

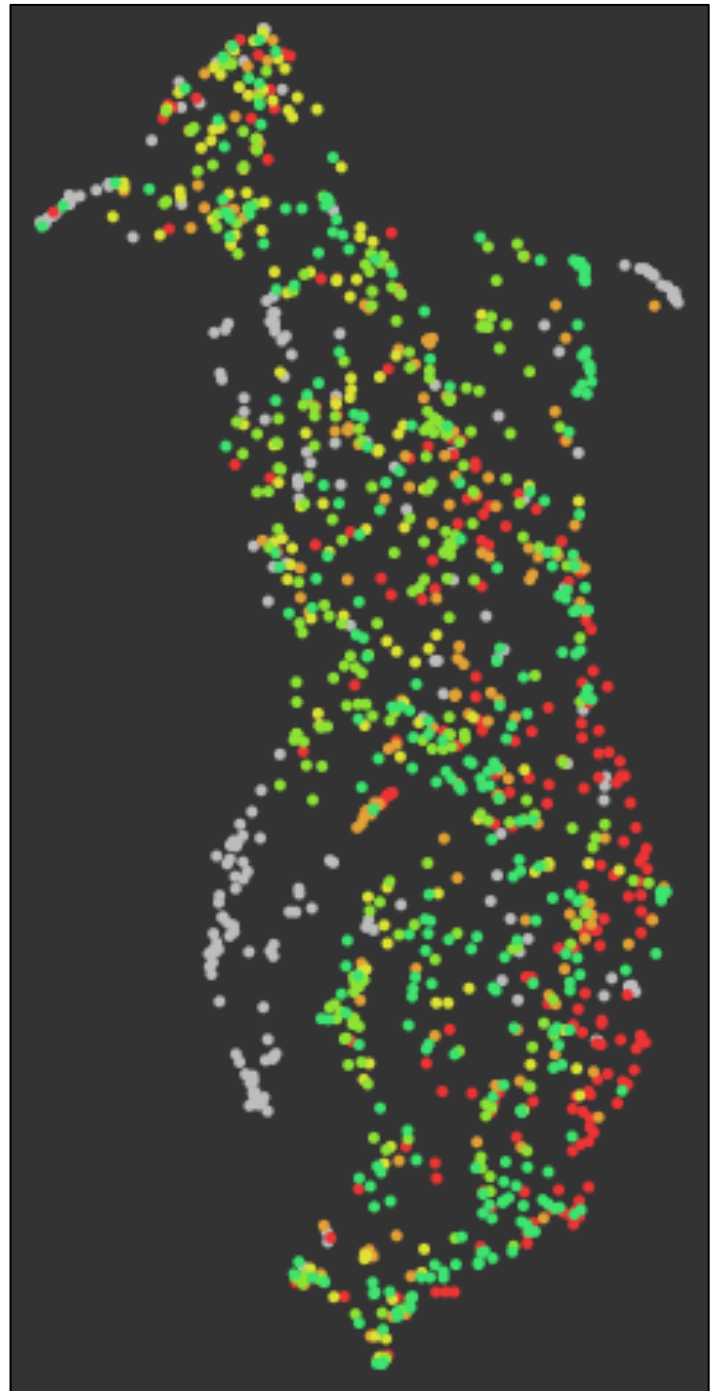


Figure 69: same human score, generated machine scores for Burton's work displayed with UMAP dimensionality reduction.

in the patch doesn't necessarily have much effect on the overall sound. Perhaps greater differences can be found when different scores are generated.

Before moving on to examine this idea, a quick overview of this sound plot and the trajectories that are present within it: from left to right, events go from short bursts to much longer gestures and prolonged tones. From top to bottom, we move from noisy, chaotic sounds to pure, sinusoidal sounds. This is indeed the impression that I got from my initial listening: I discussed the idea of two clear, starkly different sound worlds.

What this plot can add to this notion is the very smooth transition between the two, with sound distributed evenly across the two axes. What's more, there is a slight bulge at the top end, suggesting that there is much more chaotic and noisy content. This, however, is another aspect that one must be wary of when using this technique – indeed, there are more segments in the chaotic section, simply because the nature of these sounds causes the novelty slicing algorithm to make many more segmentations. The number of nodes does not necessarily correspond to the amount of time heard in performance.

Let us now examine if letting the machine score generate itself with the same human score has much effect. First, the visualisation has been flipped across the vertical plane, with segments going from long gestures and sustains on the left to shorter bursts on the right, and the top to bottom trajectory seems to have remained the same.

There are a few differences: where in the first configuration, all the colours were quite evenly distributed, here, there is one that does not span across the whole plot. Pass 1, in red, is much more concentrated towards the bottom right with fewer events at the very top and the left-hand side of the plot. Given the trajectories we have discussed, this could suggest that this first pass was much noisier and chaotic than the others. Does this have any correlation with the data from the scores generated?

It is difficult to discern many differences purely from the images, however I did export each of the scores in text form when iterating⁵. The first pass does have fewer generated machine events than most of the other scores: 509 – most of the other scores are around 520. We could draw a correlation between the number of events and the plot, as the top part, where the red pass is mainly absent, corresponds to the area of longer gestures of chaotic noise. The only exception to this is Pass 5, in dark green,

⁵ D.E.: 02_Data>04_Burton_Scores

which only had 488 events. These are very evenly distributed across the plot, ultimately suggesting that machine block quantity may not correspond to the final aspect of the performance.

All things said, there isn't much difference to be noted between the two plots: when the human score remains the same, with the same inputs, the aspect of generated audio remains very similar, even when a new machine score is generated. Let us examine, then, the final configuration where both scores were generated.

There are clearly zones of this sound plot which are much more monochrome, and the events are spaced out much further apart. The main structuring trajectories that can be identified are: from top to bottom we clearly move from very loud gestures down to much quieter ones, and from left to right purer sounds to much noisier ones.

Ultimately, it appears that these sound plots inform us less about the scope of sonorous potentialities of this network – which appears to be (and certainly intentionally) limited – and more about the balance of agency that has been setup within it. We understand quite clearly that the machine score does not have a great deal of agency, the various configurations of processing are not what is structuring the network. What really appears to give the network its shape is the human score, and the types of inputs that are being fed into it.

This informs us on strands of musical thought that Burton has inscribed – it explains perhaps why there are only four real sound processing units, and why so much time was spent on development of score generation as discovered in Chapter 3.

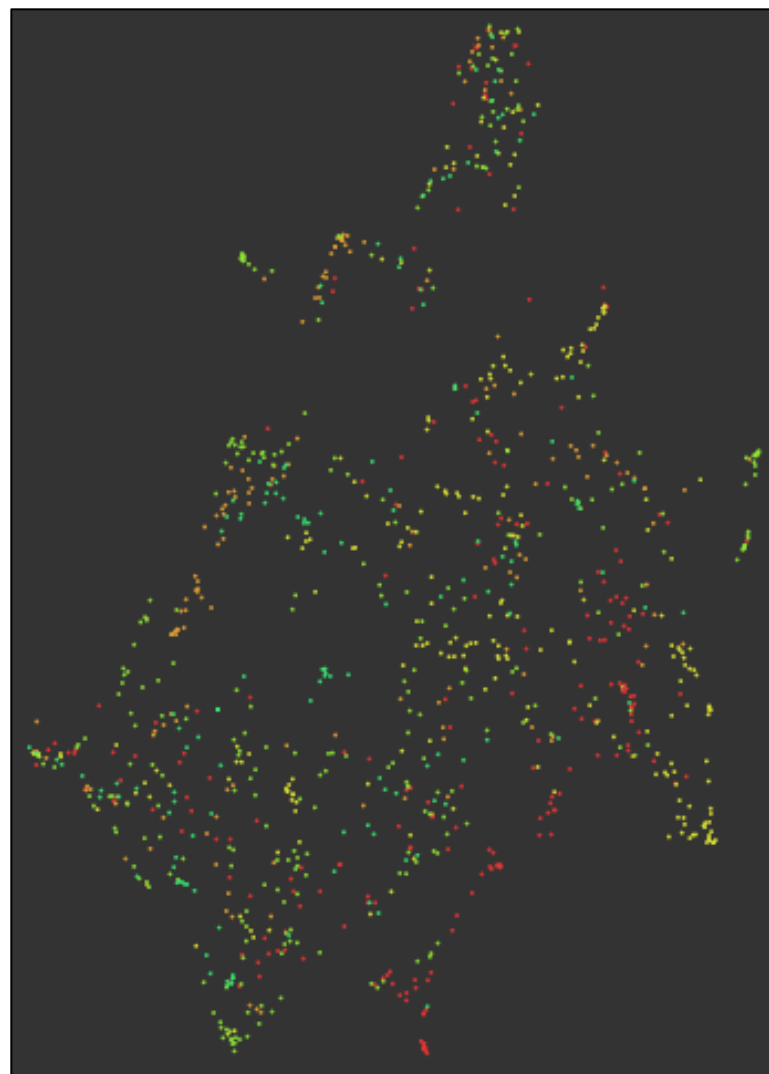


Figure 70: both scores generated for Burton's work with a UMAP dimensionality reduction.

Looking at the sound plot that visualises the three configurations together, we remark that elements from the three configurations are present in most of the different clusters we could make. This suggests that Burton had an idea about how the piece will sound, and whatever is generated in this seemingly random system, will ultimately shape into something that resembles this.

The *all generated* content in green spans much more broadly across the plot compared to the blue sounds which represent both non-generated scores – of course, there is a level of difference between the two. However, the difference is certainly not as great as one might expect it to be when taking an initial look at the performance or indeed the patch.

I will take this opportunity to discuss the novelty slicing on MFCCs paired with the MFCC description. This technique often yields slices of audio that are very short, which can render the writing of an analysis difficult – as we have seen, we can either give a meticulous cluster-by-cluster account or talk about broad trajectories of sound




Figure 71: combined sound plot for the three configurations of Burton's alternative performances displayed with a UMAP dimensionality reduction.

across the space. The inherent problem with giving a cluster-by-cluster account, is that the written analysis becomes overly verbose. Ultimately, if the analyst is not wary in their writing, we can come round in a full circle and find ourselves in the situation that these techniques were striving to avoid in the first place: having too much information to deal with.

However, abstraction is also problematic in the context of my research, as the idea is not to find any particular truth within the music, but to exist within the network and allow the scope of our listenings to grow from it. This, coupled with the constrained length of the slices, means that the best way to really draw information from these plots appears to be to navigate them oneself.

Indeed, many of the small slices can seem similar, but there are always minute differences as one travels across the network that escape written analysis. I believe these notions to be both limitations and exciting new avenues for analysis: I discuss this further in the Conclusion. For now, I depart from this automated slicing technique, and examine a final case study where the slicing was supervised.

d) Olivier Pasquet.

In Chapter 4 I created three configurations: all of the same settings; the same settings with different main tree generation rhythms (1-1-1-1 and 2-1); and one where most of the parameters on the parent patcher were being changed randomly in a human-like manner. I made many sound plots⁶, each time using MFCC and spectral shape descriptors, using t-SNE and UMAP dimensionality reductions. The sound plot configurations were: 1 for each of the results of the three configurations, two that take passes from each configuration and combine them together, and one for each of the 25 parametric modes which combine results from all configurations. The slicing was done using the beginning and end of each rhythmic tree, which was deemed to be the inscribed beginning and end of gestures.

⁶ D.E.: 01_Code>02_Tools_05_Sound_Plot_Navigation_Standalone>Pasquet





Figure 72: the params 25 combined sound plot for Pasquet's work using MFCC descriptors, displayed with the UMAP dimensionality reduction algorithm.

Figure 73: the params 25 combined sound plot for Pasquet's work using spectral shape descriptors, displayed with the UMAP dimensionality reduction algorithm.

Supervised slicing means that the choice of audio descriptor is open: the idea behind the use of the spectral shape descriptors was that they could potentially inform me about the shape and aspect of a longer gesture, rather than just accounting for overall timbre⁷. Let us take an example with a small data set to examine how this has played out: the sound plot for the 25th parametric preset for each of the three configurations. Both sets of audio descriptors have yielded very similar distributions: two clusters that are separated by a relatively large space. The MFCC distribution, however, does find one of its clusters expanded into two, even three, sub-clusters.

On the MFCC distribution, the first main cluster is on the left of the space, in a long line. There is a top cluster, a bottom cluster, and two isolated sounds between the two. In the top cluster, the bass is either very quiet, or only intervenes at medium gain for part of the gesture. The clicks are either slowly rhythmic or sparse – we note also only slices from the same settings passes (in blue) and the rhythmic iterations (1-1-1-1 in red and 2-1 in orange). The bottom cluster is similar, comprised of sounds from the

⁷ Which could conceivably be very similar across two different gestures.

same configurations, with the only notable difference being in the bass which is slightly louder than the upper cluster.

In the cluster on the right of the space, the notable difference is that the bass is heavily distorted, and much louder. There are variations in the clicks – they become less present towards the bottom, and their general aspect is more stable and less rhythmic. We also find passes from the random settings configuration here in green.

Compare this to the spectral shape distribution, which has two clusters that are much more tightly distributed. When taking an initial look through the slices, we understand that the distribution has a different structure entirely as the notably distorted bass slices that structured the MFCC distribution can be found in both clusters. Therefore, this is a parameter that must be evacuated from our ontological considerations of the sound plot. Therefore, we need to start lending attention to things other than timbre.

Looking at the bottom left cluster, the bass tends to play sustained notes throughout the whole gesture with its varying timbral profiles, either one long sustained note for the whole slice, or a long one that gives in to silence⁸. In terms of clicks, we have long, constant steady rhythms with very little rhythmic variation. Now, looking at the top left cluster, for the bass, there is an internal trajectory that sees us going from distortion up to clear sub bass at the top, and the gestures seem to be the same, sustained notes. However, there is a difference in the clicks: here we have much more rhythmic and unstable content, with accented clicks suggesting clear grooves. Slices from all the different configurations are found in both clusters.

Parameters for listening are not the same when looking at the two sound plots. The border between timbre and gesture is a bulging one, however I do believe that for this kind of supervised slicing, the spectral shape descriptors offer a more revealing picture regarding the gestures and the rhythmic content. Essentially, this is a set of descriptors which is much more inclined to offer us a path of reflection around the organisational dimension, and it is content across this dimension that I try and comprehend going on.

⁸ We can essentially consider this as being the same gesture.

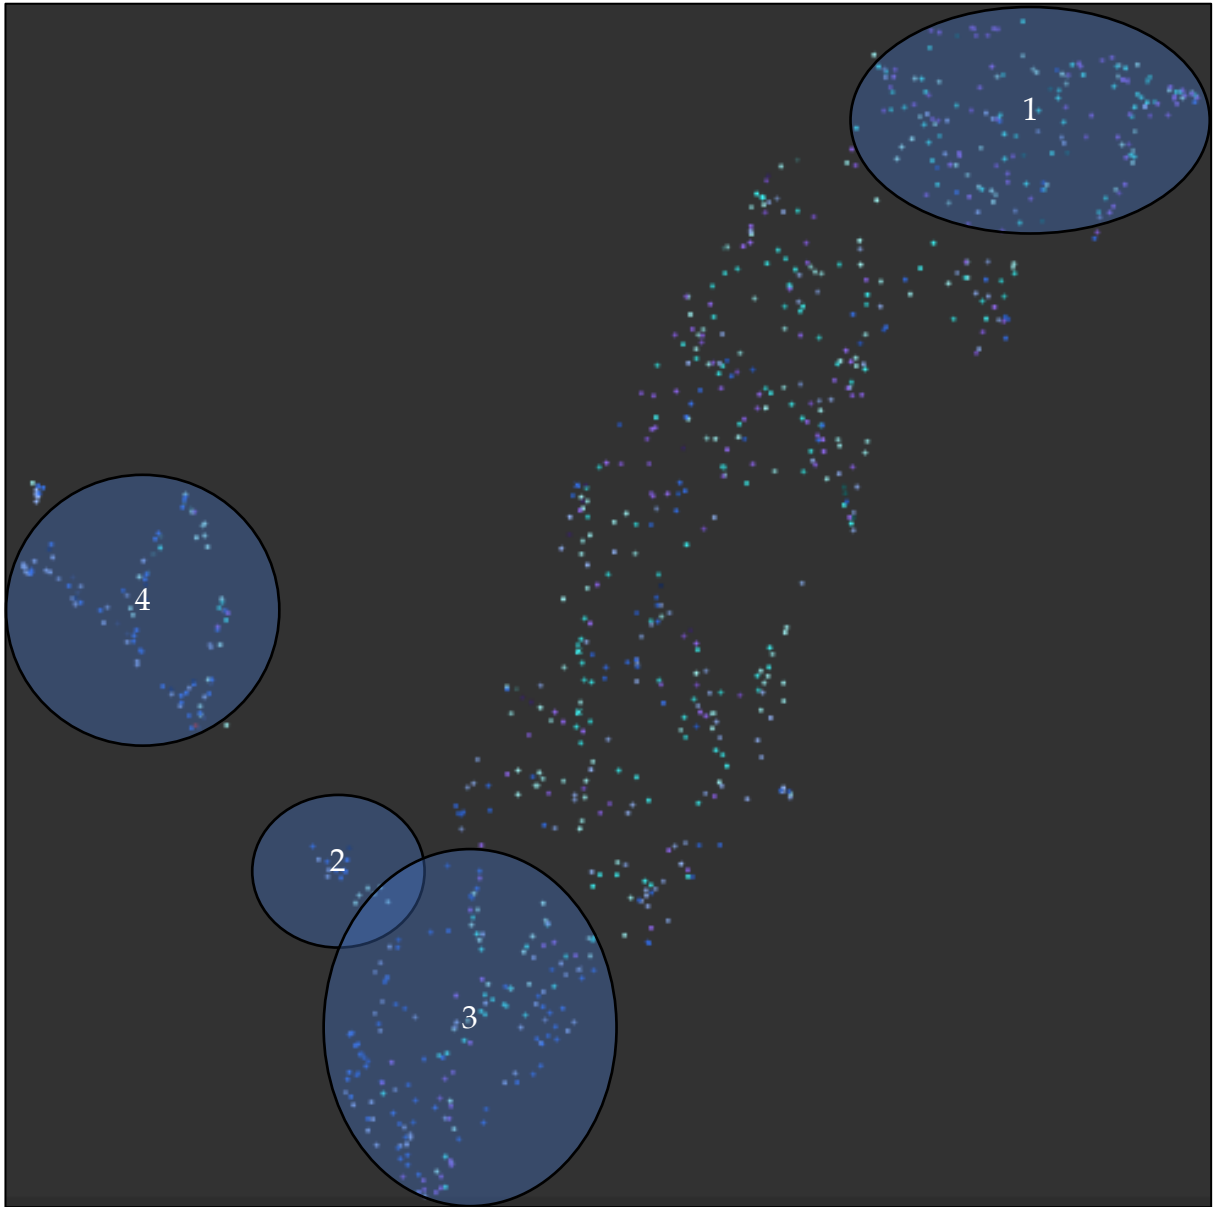


Figure 74: combined same settings alternative performances for Pasquet's work using spectral shape descriptors, displayed with the UMAP dimensionality reduction algorithm.

Looking at the sound plot⁹ for the configuration with all the same settings, there are two clear clusters that can be identified, the larger one divided into several sub-clusters. When moving down the large, right-most cluster, when examining the clicks and other high-pitched content, there is a clear trajectory from very chaotic rhythms at the top down to much less busy gestures with fewer elements and rhythms which are easier to follow. Cluster 1 shows the area where the rhythms are extremely chaotic,

⁹ D.E.: 01_Code>02_Tools_05_Sound_Plot_Navigation_Standalone>Pasquet: Same Settings

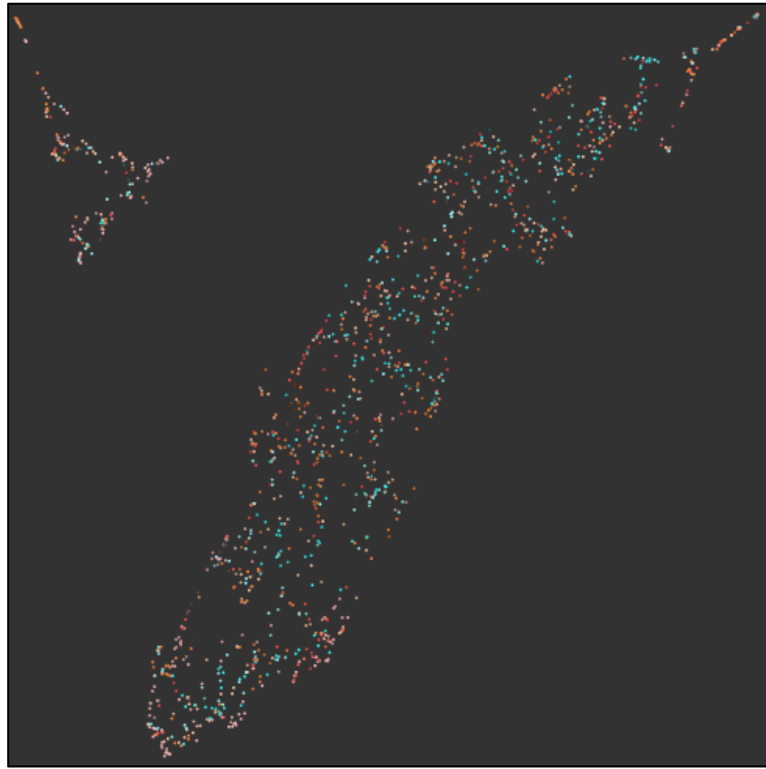


Figure 75: other rhythms configuration combined with same settings for Pasquet's work, described with spectral shape descriptors and displayed with UMAP dimensionality reduction algorithm.

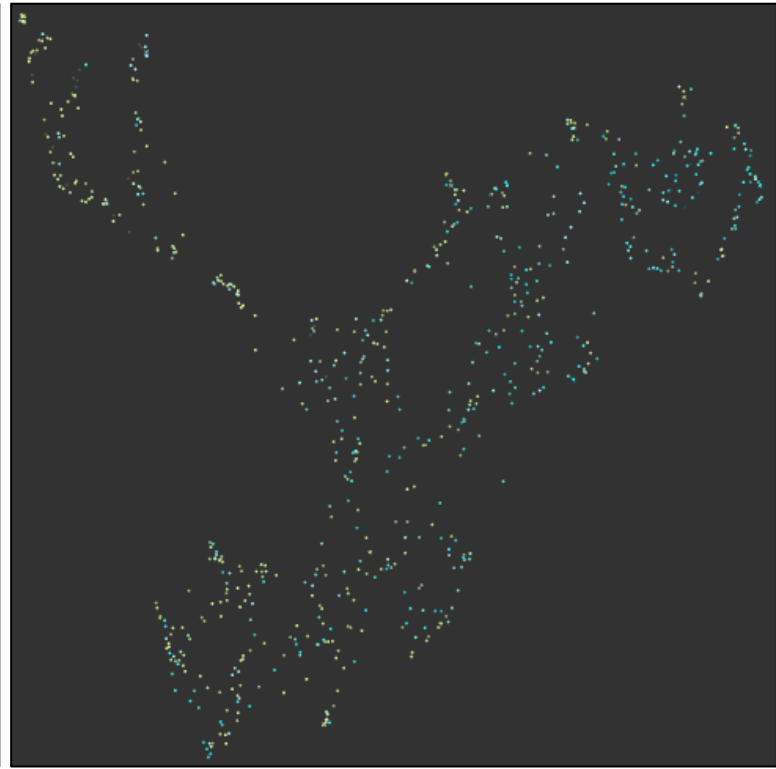


Figure 76: random parameters combined with same settings configurations for Pasquet's work, described with spectral shape descriptors and displayed with UMAP dimensionality reduction algorithm.

and it is slightly isolated from the rest of the cluster. The progression down to simpler rhythms is progressive; however, there is a slightly displaced shifting point.

In these plots, contrarily to the previous case studies, the quantity of slices can give a good indication of time in the performance as most of the slices are of an equal length.

Cluster 2 branches out towards the other main Cluster 4 on the left where there is a gesture of constant, steady rhythm that is found not only in the clicks, but also in the processing of the bass. From Cluster 3 downwards the scraping, high-pitched synth sound (corresponding to the Modalys virtual instrument) is much less present, and the rhythms are primarily driven by the clicks. Moving up this cluster, the bass gets louder and more distorted, and more sustained.

Cluster 4 is characterised at the bottom by a notable relative absence of clicks and percussive elements, with the very top seeing a return of highly chaotic rhythms. The bass in this cluster seems to remain in the sustained gestures. In terms of the parametric presets, this cluster is comprised primarily of slices from the 20s and up.

There are other concentrations of preset configurations across the plot that can indicate movement throughout the performance: for example, the top right tip of the large cluster is comprised of many slices from 10-14 presets; Cluster 2 of 23-25; or the

bottom left tip of the cluster of slices below 5. Many of the micro clusters are comprised of slices from the same pass, suggesting a continuity from gesture to gesture.

One hypothesis we could make is that the main rhythm generation will have a big difference on the types of gestures that are generated. When looking at the sound plot that displays the slices from this configuration, we remark that it is very similar in aspect to the first plot where the rhythm was the one used during the performance. Furthermore, the slices from the 1-1-1-1 and 2-1 configuration (in red and orange) are distributed evenly with the slices from the performance configuration (1-2-1-1 in blue).

Perhaps, then, the random reconfiguration of parameters across the performance will have more of an impact on the scope of the network. Looking at the plot, there is indeed a different aspect, even if it could be argued that it could present itself as a morphed version of the previous two. Interestingly, there are much more clearer groupings of gestures that lend themselves to segmentation.

At the bottom part of the plot, the random parameters (in green) and the same settings slices (in blue) seem to intersect evenly, but we then remark two clusters: one at the top left, which is mainly comprised of random parameter sounds; one at the top right which is mainly comprised of same settings sounds. The top left cluster could notably inform us about further scope of the sonorous potentialities of the instrument, which here is interwoven with the gestural potentialities. In this cluster, the bass and the Modalys virtual instruments, seem to perform a type of gesture that we haven't encountered yet – one where their playback is accelerated and slowed, creating a rolling effect. This also influences the perceived timbre of the Modalys instrument. The bass also tends to sway towards the more distorted side of the spectrum. Slices from the same settings configuration that are present are ones on the rhythmic side, with clear rhythms being accented, without being too chaotic.

The top-right cluster, comprised primarily of slices from the same settings configuration, is very interesting as this is where we find most of the slices of extreme chaotic activity in the clicks and percussive elements. One could imagine that, when randomly setting parameters, gestures would tend towards this chaotic kind of activity – however, it is indeed the supervised same settings configuration that ventures here. We can suggest, then, that this is a specific type of gesture that Pasquet was looking to create.

When referring to Pasquet's code, I identified a departure from one basic rhythmic tree to articulating it in different ways to reach more complicated things. It seems that

in Pasquet's practice complication is something that is built-up through a series of processes, of modulations of a basic rhythm. When first listening to the piece, this certainly isn't a notion that would necessarily clearly manifest itself.

This kind of approach can inform us about musical thought, not only in the timbral content of the sounds, but also in the configuration of the network that has created it, and indeed the gestural potentialities that are present. It must be said that this is a very specific case where the timbral aspect is tightly interwoven with the gestural content: this is due to the extreme poles of the spectrum that are explored. To generalise an approach that would allow a spectral shape description to give us access to gesture is perhaps a step too far. However, judicious choice of different slicing, descriptor and dimensionality reduction algorithms can inform us about the configuration of a network and the musical thought inscribed within it. I now finish this chapter by taking these visualisations a step further and examining the notion of dynamic sound plots.

3. Dynamic sound plots.

a) *Tool development.*

This was a later addition to the methodology that holds promise for future research. I developed a tool using a test corpus that I used frequently when building other tools for my research, a piece by French electronic musician Mr Oizo¹⁰ that has two useful characteristics: it is relatively short, and there is a high degree of timbral scope.

The idea of dynamic networks could take two major forms: the first would be to take the networks developed in Chapter 3 and have them linked up in a dynamic way to the networks they represent. The other is to see in real-time the paths that are taken through a sound plot while viewing a performance. In both cases, the abstract idea is to gain a visualisation of the network that shows how it is navigated over time that reveals various zones of activity. I have not addressed the first form of this – further development of the methodology would be required to bridge programs such as Max and Cytoscape. I did not put a priority on this approach as the network visualisations from Chapter 3 have been sufficient to gain a good understanding of the various points of interface.

¹⁰ Mr Oizo. *Lamb's Anger*. Ed Banger Records, 2008.

I have, however, begun to work on the second approach. I created a script¹¹ in Python that works in the following way: first, the user loads the data files for the distribution they wish to visualise, and the slice data that indicates when each slice is heard during the piece. Then, the script builds a video that displays a yellow dot moving around the space, and displays the trace of this dot as it moves in order to better comprehend the movement that is occurring.

A video¹² showing the first example I made using the Mr Oizo data demonstrates the extent to which this is an extremely compelling artefact. The music is repetitive, with clearly defined timbral zones; there are clear patterns of movement to be identified. There are parts of the plot which are navigated only at certain times. When we hear a modulated gesture in the music, this corresponds very clearly to modulations in the movement on screen.

This was very encouraging. The major limitation was the time it took to construct the video file. I rely on some external libraries for constructing images which aren't necessarily designed for this kind of application. The video is very high resolution, especially in time – to gain as smooth a movement as possible, the images are built-up at 60 frames per second. With smaller networks and shorter audio files this isn't a significant problem, however, with Constanzo's work, the process took about 24 hours to render.

There are many ways that one could approach this problem in the future: for example, with a robust network visualisation system in Max, this kind of process could be screen captured in real-time. This is an avenue for future research, and I have been able to do some of the groundwork as part of this project. The initial results are promising. I believe it to be something well worth exploring.

b) Rodrigo Constanzo case study.

In the video¹³ that was produced for Constanzo's work, we see video footage from the performance on the left, and the network navigation video on the right. The plot used to make the video was the same as the general performance segmentation t-SNE plot analysed in Section 1, and we can draw the same 11 clusters for our analysis of the

¹¹ D.E.: 01_Code>02_Tools>02_Fluid_Musicology>movie_maker.py

¹² D.E.: 04_Video>01_Demos>dynamic_sound_plot_demo.mp4

¹³ D.E.: 04_Video>02_Network_Visualisation>kaizo_snare_dynamic_plot.mp4

movement through the space. There are two different kinds of movement gestures observed throughout the performance: large scale and local scale.

Large scale gestures:

- *Macroscopic vectors:* when taking an abstract view over the whole performance, there is a clear general shift in concentration of activity from the bottom towards the top of the space.
- *Venturing from a home position:* this can be observed across the whole of the performance. We see that Constanzo seems to make various gestures across the space, always returning to the same zone: Cluster 4. He departs from this zone in varying gestures, jumping linearly to a spot (often around Cluster 10), and following the line back down, sometimes going in a circle, for example from 8 to 9 to 2. This kind of gesture is not surprising: all the slices of silence and resonance seem to be concentrated in Cluster 4, and the beginning of the piece is characterised by clear phrases that end in silence and resonance. Moving through the piece, this anchor disperses into Clusters 3 and 5, then flips at the final section to anchor in Clusters 10 and 11.
- *Progressive venturing out:* between 3:12 and around 3:50, we can note a similar phenomenon. Here the anchor is around Cluster 3, and Constanzo performs gestures that grow in length and in space covered. With each gesture, we venture further from the anchor point and spend more time away, until we reach the top of the space and jump into a chaotic passage around 3:50.
- *Avoiding clusters:* for example, between 6:45 to 7:00. Constanzo seems to traverse a contour around the middle Cluster 7, without venturing into it.

Local scale gestures:

- *Static:* for example, from 5:07 to 5:13, Constanzo remains in a very constrained space (here, the top of Cluster 3).
- *Loop:* from 7:54, Constanzo starts at a position in Cluster 8, then loops around the space in a circle to return to the same position.
- *There and back:* from 5:28, we see the simple gesture of starting at one point, moving far away, then returning immediately after – this can be considered as the *venturing from a home position* macroscopic gesture on a smaller scale.

- *Movement across clusters*: from 4:03 we see an example of this. With each event, Constanzo moves between two clusters (2 and 7), while travelling in a specific direction.
- *Cluster swapping*: at 6:06 Constanzo switches between two clusters with each event (Clusters 8 and 7).
- *Growing steps*: from 6:28 to around 6:40, we see a series of disparate events with a considerable amount of time between them. The distance between each event gets further and further as the gesture progresses.
- *Shape translation*: at 7:48 for example – Constanzo moves through a part of the space in a square shape, this shape is then translated directly towards the top.

With this tentative attempt to build-up a vocabulary of gestures from observation of the artefact, we realise that this kind of analysis would benefit from borrowing notions from choreographic studies. In musicology, despite our experience with dealing with time and simultaneous vectors of gestures, we seem to be mal equipped to deal with movement across a 2 or 3D space. For future research, it would be useful to work with colleagues from the dance world to see if they could lend a terminology and methods that would develop this part of the methodology.

This method certainly offers some answers when considering the principles of movement and landmarks in the sound plot. Here, the theoretical conception of musical creation as the configuration of networks, and the role of the analyst as cartographer takes on a very literal form.

After these extensive analyses of sound plots in the goal of better understanding the scope of sonorous potentialities of the networks, I next start bridging the gaps between these notions and the physicalities of the networks.

Chapter VI. POINTS OF INTERFACE.

1. A typology.

I now return to some of the physicalities of these networks and examine the notion of points of interface. Before an analysis of my case studies, it is necessary to spend some time equipping myself with a vocabulary to talk about points of interface in an informed manner. This typology has emerged from the analysis of case studies to come, and a broader reflection on contemporary instrument making and issues around our relationship with technology in general¹.

Broadly, a point of interface is understood as an area of a network where agency is translated in a certain way. When adopting this perspective, it can be argued that every node in a network is a point of network – this is true; however, some points of interface will be more interesting for us to observe from a musicological perspective than others. Here, I am interested in those that deal with the Record and allow engagement in the two primary activities I discussed previously: measurement and manipulation.

There are certain preconceptions around the idea of interface that are necessary to first address. Serres (Serres, 2007) gives an interesting analogy: *“systems work because they don’t work. Nonfunctioning remains essential for functioning. This can be formalised. Given two stations and a channel. They exchange messages. If the relation succeeds, if it is perfect, optimum, and immediate; it disappears as a relation. If it is there, if it exists, that means it has failed”* (Serres, 2007, p. 79).

This is an interesting point that goes against what we may initially think of when thinking of interface: for example, in digital interface design, there is a prevalent idea around functionality, around immediate translation of agency within the system. This is a design principle that is understandable in a commercial, functional market for objects such as computers: companies will sell the idea that interfacing, engaging with their system is seamless.

However, when considering the idea of interface in an aesthetic context, this idea becomes problematic; especially when considering the types of practices studied over the course of this research. In my case studies, many of the musickers concentrate

¹ However, a point of interface is not something that is exclusively reserved for the analogue to digital world and vice-versa.

effort into the development and gestural interaction with interface objects. They tend to work *with* the materiality of the object, rather than seeking to remove it.

We see the point of interface becoming a central part of many practices – interface not as a seamless portal for imparting agency within a system, but as a rough, malformed tool for translating agency into another form. It is a material object towards which one concentrates thought and bends into transcribing some trace of that thought into a network. Interface is failure – failure is an inherent part of its essence. Some of my case studies take this notion quite literally, for example Hayes who seeks to grapple with the delicate and fragile balance around the object of the game controller.

It is important to understand this perspective. It is important also to understand the point of interface as a physicality. Here, the term physicality is used to differentiate it from the notions of the Record and of musical thought. The Record is a hypothetical object, an idea about sound or music with which we engage across two attitudes: measurement and manipulation; the *physical act* of this is translated through points of interface. Musical thought is also conceptual, it is inscribed and materialised within a network through points of interface and their configuration.

A point of interface is *something* that the human actor will have to *grapple* with. With a point of interface, there will always be a process of translation, of encryption, of interpretation between actors. There is no clean, perfect passing on of agency – if this were the case, there would effectively be no point of interface.

To gain a better understanding of the scope of this term, let us draw a typology of the different kinds of points that one may encounter. Again, this is a typology that has been dressed through engagement with the case studies of this research, however a much broader application is entirely possible. The classification I propose also draws from the classifications of various nodes made in Chapter 3.

First, let us consider *material* points of interface. These are subject to many levels of classification: first, they may be composite or singular. A composite material point will be comprised of more than one singular point: some examples are a MIDI keyboard controller, a modular synthesizer module, a games controller, an iPad, a computer, or a cello.

Despite being comprised of many points, the ontology of a point depends on its conception at a given time by the person engaging with it – we have already discussed our tendency to paste certain names onto certain parts of networks, our tendency for

grammatisation. A singular point of interface will again, mechanically, be composite – for example a crossfader is comprised of a knob, a casing, and a potentiometer. We grammatise this as a singular point of interface through the functionality of this part of the network, of the process it is conceived of as performing.

Material points of interface translate agency from one domain to another – in the case of Human Interface Devices (HIDs), agency is transcribed from the physical, gestural domain to the digital domain. In the case of an acoustic instrument like a cello, agency is again transcribed from the physical, gestural domain, to the physical, sonorous domain. This is what is happening mechanically, but as we shall see, it can be considered that, through their conception, agency can be seen as being transcribed much further than the immediate domain its mechanics dictate – for example, the slider on a MIDI keyboard is mechanically translating energy from the gestural domain to the digital domain, however it could be conceived of by the user as translating agency straight into the *sonorous* domain.

These types of points can also be categorised according to the nature of the information that they translate. There are two primary types of information that can be conceived of being translated through a point of interface: *analogue* and *digital*, or to avoid confusion and to retain terminology from Chapter 4, *linear* and *stepped* (or in many cases, *binary*). There is a great difference between a point, such as a slider, a dial, an analogue stick on a game controller, that is conceived of as offering an infinite scope of possibilities between two limits; and a button or a switch that can only ever be in one of a few predefined states.

A point of interface is something that is queried, it adopts different states at different times, and the potential nature and number of those states is important to consider. These states change over time: changing of state can be *interpolated*, for example, a linear slider which will travel through intermediary states over a certain period before reaching its target; or *immediate*, for example a linear slider that would pass directly to its target without any intermediate steps.

Next, we can discuss *digital* points of interface. These will always be composite in some manner, as they are dependent on a mediating object such as a computer or tablet. They can emulate physical objects, such as a material linear slider or dial, or a material binary toggle or button.

There are two modes of interaction with them: direct interaction through the manipulation of a material point of interface – for example, moving a slider with a

mouse, or a MIDI keyboard slider that is mapped to a digital slider; or indirect interaction, where the interaction is mediated by other points of interface – for example, a button that will trigger or a slider that will be set programmatically or procedurally. This kind of indirect mode of interaction is also something that can occur in the material domain: for example, the strings of Eldridge’s Feedback Cello, a linear material point, are interacted with in an indirect manner through the transducers that cause the strings to vibrate.

Another important kind of digital point are those that provide feedback. Again, these will always be composite, dependant on something like a screen, or a light interface. However, once again, it is how the point is grammatised in the mind of the user which is important to consider – the musicker does not consider themselves to be interfacing with a collection of differently coloured pixels on a screen, but with a slider. Here, the point will translate agency from the digital domain to the conceptual – broad, often encrypted data will be translated in some way, and the encryption of the information is interpreted by the human actor. We touch again on this fundamental act of translation, of the *loss of resolution* that occurs when agency passes through the filter of a point of interface.

Finally, there are the *biological* points that we are as human beings. When agency is fed out of the computer towards the conceptual domain to be engaged with by the human actor, this agency will then pass through the human, once again be translated, and in turn fed out to other domains. The human actor is but one of many within the network, and the principles of points of interface and translation of agency apply to them in the same way.

The human actor is by no means the only type of actor from which agency can appear to emerge: for example, think of a generative composition system that will make decisions of its own accord. We could argue that the only reason this system is making decisions is because it was designed by a human actor, and that fundamentally it is their agency which is manifesting itself. This is true; indeed, my entire research bases itself on the notion that musical thought is inscribed within a network. However, this does not mean that this is any different for the human actor.

To lean on broader social and psychological schools such as social cognitive theory: what are we but the products of the environmental conditions that have led us up to this point in our lives? The human actor is a point of interface where agency from an infinite number of sources is combined, interpreted, and translated. The nature, the

essence, the mechanical aspect of a point of interface is formed through the amalgamation of agencies that have come together to produce it.

From a scientific point of view, I am hesitant to draw categorical lines. Hence, despite the affirmations I have just discussed, I do believe in a special, inexplicable, intangible aspect in the kind of agency that humans deploy. It is the reason why I choose to examine the musical thought inscribed in the network *by human beings*; it is what makes artistic practice so fascinating to participate in, to create, to study. We are complicated points of interface: it may be entirely possible to one day access a comprehensive explication of the intricate web of agencies that disfigure themselves, fuse, shatter, fight, connect and combine to produce the nature of ourselves and the agency we deploy – however, we are not yet at that point.

To summarise, for my analysis I consider three broad types of point of interface: material, digital and biological. These points can be composite or singular. A point is always in a particular state, and the nature of this state can be somewhere on a linear plane, or one of a stepped pool of possibilities. This state can be queried and can change either through interpolation or immediacy. The modes of interaction that points can have between each other can be direct or indirect. Finally, to pass through a point of interface is to move agency from one domain to another, rendering it a conceptually physical phenomenon. It is to encrypt, to translate, to interpret – the agency that is passed on will always be disfigured in some way, there is always a loss of resolution. Next, I examine some of my case studies through this lens.

2. Interface analysis.

a) Aesthetic and functional points of interface.

First, I present some examples that demonstrate the idea of aesthetic versus functional points of interface. This was first evoked in Chapter 2 when looking at the initial surface level analysis of Hayes' network. I discussed the potential difference between the way she uses the Korg controller which appeared to have a more pragmatic, functionary role, and the game controller which is something towards which a much more performative attitude seems to be directed.

They are both material composite points of interface. The Korg is divided into eight tracks, each with their own linear slider, linear dial and three binary buttons (*S*, *M* and *R*). There are 11 other binary buttons which are labelled as having various roles, but this is a generic controller that can be reconfigured in an environment such as Max.

The game controller has two linear joysticks which can be conceived of as controlling two linear scales in an xy configuration, and 12 binary buttons. Likewise using the *hid* object, this object is reconfigurable in Max.

I discovered that the Korg was used to control many elements across the patch: gain levels, synthesizer parameters, crossfading between sources. It is difficult to discern the scope of the node on the network graph as it is found in its own *bpacther*, but on the patch we observe that there are no less than 39 different routings from the object. Compare this to the game controller which only has 17 routings. Despite this, during the performance I observed that Hayes tends to use the game controller much more than the Korg – she spends most of the performance with the controller in her hands, momentarily adjusting things with the Korg, with one final section at the end where she abandons the game controller for the Korg. The game controller was principally used to control the speed of playback of recorded audio, notably voice, and trigger playback of samples.

A big difference in these two composite objects can be found in the domains from which they transfer agency. The game controller is clearly translating agency from the gestural domain, whereas the Korg appears to be translating agency from the conceptual domain. This can be tied to the materialities of the objects – the Korg lies on a table, immobile, and retains its state without supervision. The game controller must be held, and its state must also be held to stop it from reverting to an initial state.

The resolution of data from the two objects is not the same: both can be conceived of physically as having linear sliders – however mechanically, they are stepped. The Korg's sliders give a very low, typical MIDI stepped resolution of 0-127, whereas the game controller joysticks have a much finer resolution at 0-255. Observe that this object with which Hayes engages gesturally is the one that offers the highest resolution and can be conceived of as better translating the subtle movements she performs.

The same difference can be found in Constanzo's network – here we can give the examples of the composite material Dicer controller which is primarily functional, and the singular material crossfader which is primarily aesthetic. Constanzo developed this crossfader himself and went to great lengths to draw as much resolution out of the object as possible: he explained that he *“really wanted to optimize that level of control – that level of touch – as it formed the foundation for much of the sonic material of the piece”* (Constanzo, 2020).

Once again, we can conceive of the object as being linear; however, as is the case with anything that translates towards the digital domain, it will always remain mechanically stepped. Looking at the network graphs and the patch, the Dicer controller has more scope than the crossfader, but the crossfader has much more resolution (the Dicer has a stepped binary state type, the crossfader, passing through a Teensy and paired with an additional Adafruit ADS1115 ADC², had a final resolution of 0-65536). The Dicer is used to set various coloured modes within the code, and the crossfader is used to control not only the gain of inputs into effects modules and blending between signal paths, but also its speed is used to control parameters of the *cloud* effect.

Consider the mode of state changes in these points of interface. The functional points tend to be immediate – the Korg and Dicer are mainly comprised of buttons that have immediate effects; whereas the aesthetic points tend to have interpolated modes of change (although there are the sliders and dials of the Korg which are also interpolated).

This notion about the materialities of the object and how it effects changes of state was pivotal for Constanzo over the development and configuration of his network. He explains how the crossfader “*formed the material language for everything, as in the local to medium scale musical events and gestures that were generated as well as a very distant, tangentially-removed stylistic, idiomatic aspect – the nature of a fader, how it moves, how it behaves, how specifically applying this kind of technique to it commands a type of language*”³. He went on to give the example of a 2D xy pad and suggested that he could *hear* it in the music. The same phenomenon occurs with the crossfader – if the fader is at one extremity, it *physically has to* move to the other. There is a linearity to the movement and in turn to the sounds. This linearity is something that Constanzo explores and is a structuring element of his musical thought.

With Constanzo we can also examine the domains of translation for these points: we see that the Dicer tends to transcribe agency from the conceptual domain, and the crossfader from the gestural domain. We realise that the line between an aesthetic and functional point of interface is blurred – indeed, why would a conceptual idea about

² <https://www.adafruit.com/product/1085>

³ Appendix 8.4.5.

how the general modality of the network should be at a given point not be classed as aesthetic?

The categories are indeed porous. The Korg and the Dicer are used to transcribe aesthetic ideas, however, they are more a means of triggering musical thought that has already been inscribed within the network. The agency that they transcribe is immediately functional, they serve as a means for triggering something aesthetic inscribed beforehand.

During the moment of musicking, the mode of engagement is shifted when looking at the game controller and crossfader. The networks have been configured in a way that accords them somewhat of a creative agency *during* the performance. Like the others, they have been configured in a way that inscribes musical thought; however, the material interaction with these objects is also a subject of, if not musical, aesthetic thought. The agency that they impart within the network is *immediately* aesthetic.

Burton's network is interesting, as the main point of interface into the digital domain is the singular material microphone. However, it can be considered that the points of interface are the various instruments on the table, and the agency is translated in an indirect manner into the patch. The microphone is once again a stepped point of interface, albeit with a very high resolution.

There is an interesting point to discuss when regarding the nature of the instruments as points of interface that brings us back to the discussion around grammatisation. Mechanically, many of them can be considered as linear: a recorder's hole can be covered to an infinite number of degrees, a peg can be placed on a guitar string in an infinite number of positions. The ways with which one can engage with these objects is theoretically infinite. This contrasts a point of interface that translates towards the digital domain, which, mechanically, will always have to translate into a stepped form of data.

There is a mismatch when taking this into account and the material nature of the objects: the necessarily stepped translation of the slider is attached to a material object that suggests linear interaction, whereas the linear data translated by the musical instrument such as the recorder or the guitar, is heavily grammatised in a way that suggests stepped, highly segmented thought. When the chain of agency is followed in both directions, we observe that the performer and the perceiver (both musickers who can be same or different people) conceive of the stepped slider as being linear; and the linear material instrument is widely conceived of as being stepped. Even when, for

example with Burton's performance, an instrument such as a guitar is played with extended techniques, the pool of sonorous potentialities is conceived of to a large extent as stepped.

This is probably due to the amount of cultural baggage that is inscribed within these instruments. It is interesting to observe how the agency translated from these points are used further within the network. In Burton, Constanzo, and Hayes' software, recorded audio is always subject to slicing, to onset detection of some form – there is an inherent notion of segmentation around the flux of audio. On the other hand, as we have seen in Constanzo and Hayes' networks, things like digital sliders and joysticks are used for things like blending and controlling conceptually linear parameters. We begin to understand how the material aspect of a point of interface – and how it is understood at the crossroads of the musician's mind and all the actors that have contributed to its constitution – has a profound effect on the nature of the network's configuration.

The musical instruments in Burton's network certainly seem to serve as aesthetic points of interface, clearly translating agency from the gestural domain to the sonorous. They are one of two major points of interface, the other being the visual score which, as we saw in Chapter 3, was the principal element from which the rest of the piece tended to emerge. I discuss below the way in which these points interact with each other and the idea of hybrid points of interface.

In Pasquet's network there are no points with which he appears to engage with in a performative manner – there is no apparent gestural content to be translated. Does this mean that all his points are functional in nature? When looking at his network, the only point with which he imparts agency within the network is the composite, material point that is his computer – comprised of the singular binary keys, and the linear (but transcribing towards digital) material mouse trackpad.

Out of these four case studies, Pasquet is the only one who interacts somewhat directly with a GUI in his software. His GUI is comprised of number boxes and dials. The dials can only change state through interpolation; however, he can control their state indirectly with number boxes which can change state either immediately or through interpolation.

From this perspective, these points would seem to share many of the characteristics we have seen for functional points: translating agency from the conceptual domain and with potential for immediate changing of state. Most of the GUI is devoted to

controlling parameters for the four Modalys instruments, and remarkably, there are four different ways of controlling them: directly with the dials, indirectly with the number boxes, indirectly with other number boxes which control all the other number boxes at once, and somewhat directly with the changing of the 25 different preset modes during the performance.

As discussed in Chapter 3, these GUI elements were present from a relatively early stage during the development process, therefore it is possible to argue that they may not have all been used during performance. In any case, there was a moment where Pasquet at least experimented with different ways of controlling these parameters. This could suggest an aesthetic nature in the development of these points. This reveals that a classification regarding the aesthetic or functional nature of a point can be informed by the way in which it was developed – this steers us back towards the development process of Constanzo’s crossfader, or the extended techniques Burton deploys in regard to his instruments.

b) Hybrid points of interface.

Next, I examine how points of interface can come into interaction throughout a performance, how they can morph over time into hybrid points, and what this could reveal about musical thought. A basic hybridization can be discussed around the human actor’s biological body and its relationship with the various material objects. There are certainly different levels to this: I have already discussed the idea of space and proximity. Regarding my case studies, there is certainly a difference in the proximity between an example such as Constanzo and the hybridized relationship between his hand and the condenser microphone or the crossfader (even if I would argue that there is a higher degree of space between him and the crossfader than with the microphone) and an example such as Hayes’ body and its intermittent relationship with the Korg sliders.

Again, this is inherently conditioned by the material nature of the point of interface: the Korg and the crossfader can hold their states without material intervention; the condenser microphone constantly requires supervision. What are the other modes of hybridization that can be found in these networks that go beyond the relationship of the musician and the material object, and that supersede the notion of composite points of interface?

Possible instances of this can be found in Constanzo's network with interaction between the combs, the felt, the snare surface, and the condenser microphone. These are all singular material points, and all are engaged with in a gestural, direct manner with the exception of the felt which is engaged with indirectly through the condenser microphone.

Leaving the snare surface to one side for a moment, there are two possible hybrid configurations here: the first is that of the microphone and the felt. When Constanzo crushes the microphone against the felt and moves across the snare surface with it, we conceive of a hybridization of the two elements. There does seem to be, however, an unbalance in terms of agency between these two composite objects: the felt is understood as something which modulates the signal of the microphone. The sound of the microphone is still the main element that structures this point of interface – here the Record that is the microphone's signal is engaged with through the attitude of manipulation via the felt. The felt also contains musical thought; however, the Record that it constitutes *depends* on that of the microphone. This leads us to conclude that there is no real hybridization in this instance.

The case of the combs and the microphone is different. Most of the time, Constanzo holds the comb in one hand, the microphone on the other, and they are both engaging with the snare surface in a seemingly separate manner. Other times, Constanzo brings the comb and the microphone together, and they begin to act as one hybrid point. The main difference between this configuration and that of the microphone and the felt, is that of balance of agency. Here, there appears to be no immediate Record that is being modulated, rather, they are each equally co-dependent. It is impossible to unravel the two signals: it is not a case of the microphone signal being modulated by another element; here, the content emerging from the comb and the microphone are inherently intertwined. This becomes a Record unto itself, and it is this hybrid object that Constanzo engages with and seeks to measure and manipulate.

Finally, we can bring the snare surface back into the picture. There are several interactions to speak of: the snare surface and the microphone (the microphone sometimes modulated by the felt); the snare surface and the combs; and the snare surface and the hybrid comb/microphone object. Are these configurations hybrid points where the agency of each object becomes interwoven, or is there an unbalance in the agency which is deployed? Are these composite objects with which Constanzo

interacts with, or are they two objects which Constanzo is causing to interact between themselves?

I would argue that it is impossible to unravel the agency that emerges from the snare surface and the various configurations of the microphone, suggesting a hybrid point of interface; but when the comb is acting upon the snare's surface, we are dealing with a modulation, a manipulation of this Record.

The gestural vocabulary that Constanzo deploys does not suggest that we are invited to hear the snare as an object unto itself: what is being deployed is the signal that emerges from the hybridization of microphone and snare surface. This allows us to analyse this hybrid object as a single, composite material point of interface: like other material instruments it offers a linear plane of sonorous potentialities, with modes of interaction that can pass from state to state both through immediacy and interpolation.

How far can this hybridization extend? For example, how does the crossfader enter this equation? The crossfader is primarily used for the setting of parameters, blending between signals, and setting of gain. I previously defined it as a singular composite, aesthetic point of interface. From Constanzo's point of view, it is logical to consider that the crossfader and the afore-mentioned point centred around the snare surface are separate: clearly, the strands of musical thought that are conceived of in these two objects as Records are different in nature. It could be said that, further down the chain of agency, the crossfader is essentially manipulating the signal from the snare surface hybrid. This is not false; however, it also operates on other sound sources in the network, and the quantity of aesthetic thought that is inscribed within it – notably gestural in nature – is sufficient to create an unbalance of agency.

That being said, this may not be the case if we regard the network from other perspectives. First, again from Constanzo's point of view, there is an inherent relationship and synchronisation in the gestures that he performs around these two objects. With this network, Constanzo set out to emulate a turntablism setup: in this kind of network there are serious grounds for conceiving of a hybrid point between the vinyl and the crossfader. From our detailed and omnipotent point of view of analysis, we can conclude that in Constanzo's network this is not the case.

However, there is also a difference to be drawn about how the network is grammatised by Constanzo during performance, and during development. It is entirely possible that Constanzo adopted the perspective we have described during the configuration of the piece, but that during performance, he may shift to a

conception where the crossfader and snare surface become hybridised. This kind of conceptual shift can also occur during the performance: Constanzo's interactions with these objects are not constant – sometimes he adopts a more traditional turntablism gestural vocabulary, sometimes not.

Another perspective that is interesting to consider is that of the audience. They are unaware of the mechanical workings of the intricate network of code and material instruments. Faced with Constanzo's setup, they could be forgiven for expecting to hear a drum solo and being surprised to hear the types of sounds to emerge from the performance. They may recognise some of the turntablism techniques being deployed, but this may escape notice. When the audience sees a solo performer on stage, playing what seems to be a single instrument, it is reasonable to consider that they will consider the performer as engaging with a single, hybrid point of interface.

Not all hybrid points of interface necessarily suppose proximity in space. Let us take the example of Burton's network, and examine the relationship between the graphic score, the instruments, the microphone, and the performer. As we have seen, hybridisation can be identified by looking at the balance of agency between actors that appear to be working together.

If we start by examining the relationship between instruments and microphone, these are two points of interface of different natures. We previously identified the instruments as being aesthetic points, and we can understand the microphone as being a singular material, functional point. When regarding a combination of points of different natures, it seems difficult to think of hybridisation. Indeed, in the examples above, all Constanzo's objects were notably aesthetic in nature. Here, the microphone serves as a means of translating the agency to the digital domain from the instruments that have been translated from Burton's gestural domain to the sonorous domain.

The diffusion of the graphic score is an interesting phenomenon. We see a visual manifestation of a Record that is also simultaneously being heard. From the audience and Burton's point of view, this visual representation and the sounds of the piece are interwoven into a hybrid point. Each depends on the other, each is driving the perception of the other.

More interesting still is to consider the relationship between the hybrid score/sound feedback and the hybrid Burton/instruments points of interface. Consider the agency that the various elements of the score translate towards Burton – the length, shape, size, and position regarding other events displayed on the score deploys agency which

is tightly interwoven with the agency that Burton sends for translation within the network via the microphone. There is a notable notion of feedback in this chain of agency which starts neither with the score nor Burton.

Finally, let us take another example of potentially hybridised points of interface where spatial proximity is not an immediate factor. Again, from a perceptual perspective, we can consider the interactions between the sounds emerging from Pasquet's software and the lights that have been set up around the space. If we examine the lights as a point of interface, what is the agency that is being translated? The flashing of the lights is mapped to the playback of the main generated tree. This is effectively translating agency from the patch into the spatial, material domain; it also translates agency into the conceptual domain of Pasquet and the audience (in Chapter 3 I proposed the idea that Pasquet adopts an attitude that is close to that of an audience member).

Much of the sonorous material that is output by the network is chaotic, especially in the high-frequency range. This accenting of the main structuring rhythm from which everything is derived and to which everything is synchronised can serve as a guide to the listener. Can we talk of hybridisation between the lights and the sound? Mechanically, the lights *depend* very much on the sound, and they are both products of the system from which they emerge. However, I argue that the extent of the aesthetic thought being imparted by the lights upon the network is such that it becomes not only something that depends on the sound, but something which is inherently interwoven with it. The agency that these two objects impart is hybridised as they constantly inform one another.

A point of interface, then, can expand to various levels of abstraction. The various grammatisations that are constructed of them can differ according to the actor that is interpreting them (this actor, once again, constituting themselves a point of interface). To conclude this analysis, I now look at a very particular set of points that play structural roles in these practices via code design. I inspect not only the idea of GUIs, but the structure of the code itself.

3. Code design.

The computer is a composite, material object and any points of interface such as a GUI are engaged with in an indirect manner, either through the computer's hardware, an external controller, or algorithmically. Most points of interface in a piece of code will not actually be wrapped by a GUI element, and their translation is effectuated within

the stream of execution of the code. Before looking at some concrete examples from the case studies, I propose a discussion around the design of interfaces in the digital domain that calls back to some of the ideas discussed at the beginning of this chapter.

a) *Digital interfaces.*

Fuller (Fuller et al., 2008) proposes a framework for interfaces in computing: essentially, there are two sides to a computer program: the underlying *system* that performs a series of tasks; and the *interface* that the human actor will interact with to engage with this system. The underlying program is in fact composed of a series of interfaces: these terms are merely a result of grammatisation that can be readjusted at any point. Fuller describes the parts of a network that the human agent interacts with to engage with the underlying system as “*user-interfaces*” (Fuller et al., 2008, p.150) – the grammatised, composite points of interface which perform a task. He states that they fall “*under the entry on language*” (Fuller et al., 2008, p.150).

All programmers have dealt with interface in some form: from the design of a simple GUI to a basic syntax they use to perform tasks. The general conception of a good interface intersects with notions of clarity, readability, understanding and ease-of-use. It is generally considered that a well-designed interface will reveal the workings and possibilities of the underlying system without the user having to get their hands dirty and immerse themselves within its mechanics. An interface is a *way into* this system at a high conceptual level. Indeed, a ‘well-designed’ interface is *useful*, especially regarding tools that we use in our day-to-day lives. Fuller describes this as “*the paradigm of ‘user-friendliness’*” (Fuller et al., 2008, p.151).

These are interfaces of the functional nature discussed above. I do not deny the usefulness of these kinds of interface – but they do possess certain characteristics to be discussed: they are notably standardised, utilitarian, and fixed.

I illustrate this with some examples. If one compares a blank Word document between the first version of the program from the early 90s and the current version, we remark that in nearly 30 years, the interface has changed very little. The paradigm of a top toolbar above a central workspace is something that has been present from very early on and can be found in most digital interfaces.

We could also take an HTML page which explicitly segments the data on screen into hierarchical elements such as *headers*, *footers*, and *body*: within this programming framework, it is difficult to break free from this kind of layout. From a commercial

point of view, moving away from this kind of template is considered a risk. This is tightly linked with the notion of software as a product.

Common themes over the 20th century, with thinkers like Adorno (Adorno, 2002) and Benjamin (Benjamin, 2008), have been around standardisation, and how capitalist ideologies can dissuade artistic innovation. Software does not escape these fatalities – this is clear when looking at Word; however, it is also present in programs that go beyond text formatting. In our musical context, programs like DAWs all share very common features. Consider the Sibelius⁴ music notation software whose developers went so far as to purchase the license from Microsoft that allowed them to incorporate the *ribbon* toolbar in their interface.

Word, Sibelius and DAWs are generally considered to be *tools*. They are not intended to give us access to an authentic, temporal, inner mode of existence – they are spatial and utilitarian. They are things which aid us in our measurement and segmentation of the world, a means to an end. However, could we consider an interweaving of these two aspects? Could a working environment explore notions of temporal ontology? Could our relationship with the tools that we use interrogate these kinds of notions?

Before examining this idea closer in the digital realm, we can take a material musical example which has a long history. We could consider the score as an interface, as a means of engaging with an underlying musical system, as a way of triggering processes from a musician. How has this object been treated over history? It appears to have emerged from a utilitarian functionality: the need to bear a trace of musical thought for practical purposes of transmission. Quickly, however, the Work began to coincide with the score. In the contemporary epoch, still, there are composers who lend particular attention to the score: Feldman, for example, was particularly attentive to the graphic layout of his scores believing that the music to emerge from it would draw from this (Cline, 2016). Are there analogues to be drawn between this aesthetic treatment of the score and the design of digital interface?

The notion of standardisation of interface design necessarily influences how the software is used. The more interface design becomes standardised, the less scope there is for divergent use of the underlying system. We become lazy in our use of technologies, to a point where the conceived shape of the underlying system is morphed by its interface and the idea of pushing the system becomes impossible. In

⁴ <https://www.avid.com/sibelius>

the contemporary epoch, the general perspective seems to be that the interface is static, and that artistic creation is to be performed *within* it, *despite* its boundaries. However, the practices I have seen in this research and many others seem to wish to break away from the fatalistic conditions of contemporary interface design.

A first way of approaching this is through rebellion and subversion, through a refusal to accept the interface that one is given and taking things into one's own hands to extend or modify it. Consider, for example, circuit-bending: taking apart various devices and modifying the internal mechanics to extend its sonorous potentialities. This is a modification of the underlying system, but it also usually passes through modification of interface, adding new dials and sliders, reconfiguring the mappings of existing points of interface. This practice is often conveyed with a particular visual aesthetic – warping children's toys into grotesque creations, leaving parts of the machine bare: a clear statement of wanting to depart from the clear-cut, neatly packaged boxes within which things are sold.

This follows on to an aesthetic taste for the surreal, the absurd. An interesting musical example is the AudioNerdz VST Delay Lama⁵. This is a VST developed by a group of students from the Netherlands: a vocal synthesis synthesizer with a stereo delay which allows the user to control pitch and voxel sound. The interface is remarkable: it is difficult to recognise some of the sliders as sliders, none of the controls have any labels, and most of the space is dominated by an amusing, poorly 3D-modeled Buddhist monk who moves his mouth as the user plays notes.

The team of programmers were all experienced, and surely had experience dealing with concepts like interface design. When one sees the interface, it is difficult not to see it as an intentional joke, a statement against the slick, functional interface designs that we are used to seeing in VSTs.

Another form that could be considered is through *resistance* – the interface as resisting the agency of the user. This is a notion that Hogg explores in his work: he explains that “*resistance of materials to energy, things to action, objects to movement, animate bodies to external forces, is [...] of primary importance in the development of human consciousness*” (Hogg, 2013, p. 257).

This is starkly opposed to the paradigm of clarity and ease of use. There is something to be said for an interface that would be unstable, that would be difficult to control,

⁵ https://www.kvraudio.com/product/delay_lama_by_audionerdz

that would render the nature of the instrument opaque, where control is not an immediate structuring parameter. This is something that can come to address the notion of disembodiment in contemporary interface design discussed at the beginning of this chapter: when there is no resistance when interfacing with the underlying system, the physicality and materiality of the process is lost. A resistant interface “*is a powerful mechanism through which cognition is inextricably related to embodiment*” (Hogg, 2013, p. 258), it suggests deployment of knowledge that escapes functional, utilitarian thought and knowledge – it can demand a much more performative, embodied, and temporal engagement.

b) *Case studies.*

Are there places where this idea of resistance can be found in my case studies? I could arguably place Constanzo and Pasquet’s work on one side of a spectrum, and Hayes and Burton’s work on the other. I have discussed the notable tidiness of Constanzo’s coding, and how he is clearly wary of his code being usable by other people. He packaged his three main structuring effects modules into Max4Live⁶ objects, and all his code is clearly commented, colour-coded and immaculately modularised.

What can we draw from this? Constanzo tends to work his code into the functional, utilitarian model: we also recognise the standardised nature of his interfaces, falling in to the Max4Live, VST plugin template with clearly labelled dials. Some of the dials – for example the *timbre* dial in the *cloud effect*⁷ – evoke some of the seemingly problematic parameter labels encountered in commercial VSTs. It is common to see arbitrary names pasted onto parameters or collections of parameters: things like *brightness, tone, shape, or intensity*.

This kind of phenomenon can be conceived of as being parallel to notions of blackboxing, and detrimental in as much as they bar the user’s access not only to separate parameters that may be controlled by one top level parameter; but also, because they hinder the user’s understanding of how the instrument works. Indeed, one could wonder why it would be easier to call a parameter *timbre* rather than *HPSS source separation blend* – why hide this information from the user?

Putting questions of elegance and conciseness aside, perhaps this phenomenon could be viewed through the lens of some of the ideas discussed above. Perhaps this could

⁶ <https://www.ableton.com/en/live/max-for-live/>

⁷ Mechanically corresponding to a blending of various HPSS and transient separations.

be a move towards a more aesthetic conception of the underlying system, one that goes beyond cold lists of mechanical parameters. Perhaps a move towards more poetic concepts in parameter naming begins to give us access to a perception of this VST and the algorithm as an aesthetic object rather than a tool? My stance is as follows: this is certainly an avenue that could be explored, maybe with more effect when pushed to much more extreme and dynamic measures; however, from the examples that we see in commercial VSTs, and things like Constanzo's Max4Live objects, it is more a case of ease of use, and wanting to engage users that don't necessarily possess the technical knowledge to understand how certain parameters will affect the sound.

On the surface, Pasquet's code seems very well structured and modularised – this is the case in presentation mode, however as soon as the user begins to open some of the subpatchers, they understand that the lower levels of the patch are much messier. There are also elements, such as the four sets of dials for controlling the various Modalys instruments that were evident candidates for modularisation, that have been copied and pasted across the main patch.

In this case, I do not believe that it is useful to read too much into the messiness of the code for two reasons: first, Pasquet built the patch in a very short period of time; and second, this project is clearly an amalgamation of fragments of other projects that have been collaged together. We also know for a fact that, when working within his own time scale, Pasquet is a very meticulous coder: evidence for this can be found in his JTOL library, with extensive and clearly constructed help files. The very structure of this set of objects attests to a working knowledge of modularity. He also promotes and shares it on his website as a tool that is intended for use by others.

The modular conception of his code is also read through the layout of his parent patcher in presentation mode – there are clear streams that emerge from the main tree structure that all occupy separate parts of the screen. I would argue that, like Constanzo, Pasquet can be found in the field of transparency and ease of use in terms of interface design.

On the other side of this conceptual spectrum, we find Burton and Hayes' work. When looking through the patches, we are immediately struck by the seemingly chaotic nature of their coding. The essential question when regarding something like this, is to discern to what extent the musician still has control and a comprehensive understanding about what is happening. As seen in Chapter 3, the development of the patch is an iterative process: over a certain period, different elements will be

constructed, added on to what exists already and modify the system as it stands at that moment. The developer will have a unique position in being conscious of a system and a spatial coding of the patch from which they draw signification – they will place certain elements in certain places within the space and visualise a grammatisation that may be invisible for someone who did not assist in development.

It is possible, however, that through this iterative process, certain aspects may escape the comprehensive understanding of a developer, and the system becomes something that supersedes them. There can be a progressive shift from building-up a system of elements where one is fully conscious of what each part is doing; to a system that functions, to which we have accorded top-level parameters, but has gained agency through its escape of control.

This could be hypothesized in Burton's network. Development began with score generation, then shifted towards sound processing once this was completed. I also found that Burton had set-up a set of top-level *function* objects to control generation of this score which left many parameters and local level processes of their own accord.

In such a complex system, it can be supposed that towards the end of development and during the performance, Burton was no longer conscious of the local level decisions being made by his algorithm. He grammatised the system into a top-level interface and purposefully allowed himself to conceive of the network in this abstract way. Subsequent additions to the network are thus entering into a system that has its own agency, and amendments to previous code become far less well implicated and optimised than if they were incorporated from the beginning. One adds on to the system, *grafting* new nodes into the network.

Even if the musicker may be able to find their way around the patch and the interface like no other person may thanks to the afore-mentioned embodied knowledge about the network that emerges through iterative development, it is possible that parts of the system may become impenetrable to them. Indeed, Hayes discussed the idea of resistance in her network – she was discussing the mapping of the game controller at the time. However, we can extend this notion to the rest of her network. It is a complicated web of interconnecting parts, a sonorous body that will produce sound without her intervention; small changes in her gestures are amplified through a chain of connections of agency into sonorous results that can be unpredictable. This is something that she actively looks to create. When looking at the complicated state of

the Max patch, we can conceivably extend this idea to the code itself: an impenetrable object, full of agency, which the musicker must resist and fight to find a place in.

Chapter VII. THE RECORD AND MUSICAL THOUGHT.

My methodology has offered me a detailed view through various lenses on the aspects and the workings of the networks configured by my case studies. In this final chapter, I take a broader perspective and take stock of what I have learned by focusing on my third primary research question – that of musical thought and the Record. Through three conceptual levels of abstraction, I identify Records around which the musickers have inscribed musical thought and examine the activities of measurement and manipulation that are deployed around them.

I draw on many elements from the previous chapters, simultaneously interrogating my methodology and its effectiveness regarding this question. The three levels of abstraction I examine are: the Record as a hypothetical fragment of sound; as a larger scale musical event; and as a broad musical or aesthetic object.

I will begin by reviewing the notion of the Record outlined in Chapter 1. The Record is a node in the network that is understood as a hypothetical idea of sound or music. It can be interwoven with a physical element, but it is not to be considered as the physical manifestation of this element: for example, a sound file on a computer is considered as the physical manifestation of a Record, but not the Record itself. It is the expectation of a sound; it is what we conceive of when thinking of it.

It can be engaged with in two principal ways: measurement, from the hypothetical perception of the Record to a detailed statistical representation with audio descriptors; or manipulation, from a simple playback of the sound to a complicated modulation of its spectrum across various dimensions (energetic, temporal, and organisational).

I present examples of musical thought being deployed through the measurement and manipulation of Records across these three dimensions. These processes are performed physically through points of interface. I go through my examples systematically, interrogating each time: the identity of the Record, the points of interface at play, the processes of measurement and manipulation across the three dimensions, and the abstract conclusions around the musical thought that are inscribed in that part of the network.

1. The Record as fragment of sound.

I start with an interesting instance of a Record being engaged with in Burton's network: the entire output of sound. Here, the hypothetical sound is extremely small,

being considered almost frame by frame. The points of interface concerned are the sound inputs and outputs: Burton uses the *adoutput~* object to feed the outgoing audio with a delay of one signal vector into another point of interface, the FluCoMa *melbands* object¹.

Measurement is performed in the energetic dimension, looking at the mel band content of each slice. We can argue that there is also content from the temporal and organisational dimensions, as the temporal and organisation structure of the piece will be found accumulatively over time in the outgoing audio. This measurement is then used for manipulation – the loudest bands are taken and used by the audio processing units to make decisions about parametric settings. Burton looks to measure outgoing audio, see what parts of the spectrum are quieter, then build-up those parts with the processing. The manipulations he performs with this Record can be read across the energetic and organisational dimensions.

What can we say about the musical thought that is inscribed here? As suggested in Chapter 2, it appears that Burton wishes to try and find a balance in the audio spectrum. I also discovered that he uses external VSTs to keep the EQ in a balanced state, and I found various instances of normalisation of audio. There is a clear desire to avoid harshness in the audio spectrum.

In the organisational dimension, we understand that one of the structuring ideas is that if a sound has lots of energy concentrated in a certain part of the spectrum, this will soon be counterbalanced further into the piece. We could even go so far as to suggest the origin of this tendency towards balance. I evoked the idea of sound painting and building-up different layers of sound. This could be traced back to Burton's studies at art school, or to his career as an independent musician (rather than operating in the academic world). In a more commercially driven context, notions of balance and normalisation of sound are important.

Next, I consider an example from Constanzo's network. The Record I define here is sound from the snare drum surface. I consider this not sliced by the gestural events which occur around it – this would be a Record that would correspond to the second level of abstraction below – but as an *untimed*² slice of this signal. The primary point

¹ Configured to a latency of 1024 frames.

² In Pasquet's sense of the word.

of interface here, discussed in Chapter 6, is the condenser microphone and the various hybrid configurations that occur around it.

We can tangibly point towards measurement in the temporal domain with use of onset slicing, and in the energetic domain with the loudness, spectral shape, and pitch descriptors. Constanzo is constantly looking to obtain a very complete timbral picture of this sound via this statistical data for each slice which he defines as beginning with an onset. The manipulations are numerous: first in SP-B, where through a series of signal modulations the sound is transformed and output. In Chapter 5 it was suggested that these transformations could be performed with the goal of making the Record sound like the sounds from the metal resonance library. Interestingly, another manipulation around this Record triggers playback of sounds from this library that are matched based on the afore-mentioned measurements. Indeed, this Record is tightly interwoven with another.

At certain points in the piece³, this content will also trigger playback of the solenoid robots. There are also the material manipulations that occur on the snare surface itself with the combs and the felt – Constanzo measures, conceives of the sound in a conceptual manner and then manipulates it. This Record has multiple implications over what is heard during the piece – most of what we hear is the result of a measurement, of a conception of this sound.

This leads us to question several things: first, the notion of how Constanzo conceives of this Record, the hypothetical sound of the snare drum. We know that Constanzo is primarily a percussionist, and that he has worked with this instrument many times, in many ways. Why does he feel the need to transform the sound like this? What role does the metal resonance library play in the articulation of this sound for him? It would be possible for him to abandon the snare drum completely and build a setup of various metal objects and start from that point.

There are several possible explications: first, there are the practical implications – it would be arduous to build and transport this kind of setup, and through his experience with this instrument Constanzo can deploy a virtuosic skill. Indeed, we have noted how important gestural vocabulary is for Constanzo – perhaps he draws a dissociation between gesture and the sound they produce. Over development, the

³ Controlled by the Dicer point of interface.

gestural possibilities of the network and the sounds they produce belong to their own streams of development.

One final point is that most of these manipulations operate in the energetic domain, and that control over the organisational dimension is left with Constanzo and his immediate gestures. This attests to a certain conceptual separation of these dimensions in Constanzo's approach.

With Hayes, we can look at the NMF matching drum machine. The Record here will essentially be a sound file in one of the numerous sample libraries she uses. She measures this with the FluCoMa NMF tool, looking for similarities in activations in other Records. This process of *matching* is one that occurs frequently across the various case studies and can perhaps be explained through the agency of the FluCoMa project and the kind of workflow that it suggests.

Here, the point of interface is exploded across several elements: the buffers within which the data is contained, the analysis of the pulsar synthesis signal which will trigger playback. It is the NMF activations that are primarily used for the matching process – this is a measurement that is operated primarily in the temporal dimension. This suggests that it is the gestural envelope of a sound that interests Hayes when matching. Evidence for this was discovered in Chapter 5 when looking at the sound plots: we found numerous examples of segments from the performance and rehearsals that emulated the gestural profile of examples from the sample libraries without having the same energetic content. Parts of the performance and rehearsal audio seemed to be *imitating* the gestures of samples.

This could reveal a great deal about Hayes' musical thought and general conception of sound, and the sounds of sample libraries. Constanzo used sample libraries to provide audio content for his piece; this is true in other parts of Hayes' network, but we also see the sample being considered as a shape, as something to be imitated by the network.

Finally, an example from Pasquet's work. The Record here is that of the hypothetical sound of the Modalys instruments. Playback of these sounds is being triggered by a transformed version of the main generated tree (I return to this as a Record below); what interests me here is the sound of the instrument in its *untimed* form.

There is no statistical measurement of this record as we have seen in previous examples – measurement passes through the same point of interface which translates

its manipulation. The dials and number boxes display several perceptual parameters concerning the sound: the *radius*, *weight*, *area*, and *aperture* parameters of the instrument. These are parameters that perform manipulations across the temporal and energetic dimensions, but through the feedback of their state, they also allow Pasquet to draw a conceptual measurement of how this Record will sound. This is interesting, as contrarily to other examples, the measurement is formulated through high-level, perceptual parameters like the ones discussed in the final section of Chapter 6.

Perhaps Pasquet tends to consider sound in a more abstract, conceptual way. In the performance part of his network, there are no examples of using descriptors or slicing or any of these algorithmic, large-scale statistic and data collection techniques. He puts his trust in his knowledge of the instrument and the kinds of sound that it will make according to a set of four parameters.

This also attests to the spatial nature of his work discussed when noting his history of installations. Here, the Modalys instrument is indeed a physical emulation of an acoustic instrument; the parameters that he is changing are aspects of this emulated physical object. This suggests a tendency to regard sound much more in terms of acoustics and space, rather than a Record that would be purely sonorous.

2. The Record as musical event.

I now take a step back of abstraction and examine examples of the Record as musical event. In the previous section the Records tended to be measured across the temporal and energetic dimensions, here we expect to see things deploying themselves across the organisational dimension: untimed sounds begin to be deployed into a context as musical events.

I begin by examining a Record with Burton: a red block in the score. The first point of interface is the score and its diffusion on the screen discussed in the previous chapter. Measurement occurs across several dimensions: locally, it is conceived of across the temporal dimension, having a certain length, amplitude, and shape. Simultaneously, these measurements are understood in the context of the rest of the piece: the length draws its signification from the lengths of other blocks around it – this is constantly reconfigured as the score constructs itself over the course of the performance.

Manipulation of this Record can be read through the filling of these blocks with sound, bringing a second point of interface into the picture: the hybrid performer/instrument. The hypothetical form of this musical event is being taken and

manipulated through play. This process in turn creates another Record which is then measured and manipulated by the audio processing units.

It is interesting to ask if these manipulations are having a large impact upon the organisation dimension across the piece. As the input buffer gets filled with audio, the performer will necessarily play differently according to content that has already been performed. This was seen through the structure of the instrumental configurations examined in Chapter 2. Furthermore, as audio builds-up, this will affect choices made in the audio processing to fill out the spectrum. However, the generation of these Records, these red blocks, and indeed the machine blocks, is not tied to their manipulation in any way. As we saw, they are tied to a series of pre-defined *function* objects, and not anything derived from audio analysis.

How can we articulate the musical thought inscribed here? Like the previous discussion around Constanzo's work, we can conceive of a conceptual separation between agency being deployed across different dimensions. Here, the broad organisational structure of the piece is decided ahead of time, and the local records which are measured and manipulated during the piece have a relatively small degree of agency on it. Ideas do articulate themselves across the organisational dimension at a more local level with the red blocks, but the balance of agency is stemming *from* the broad organisational structure, *towards* the temporal and energetic manipulations of these Records.

Burton has created a structure, a shell that can be filled out in any number of ways. I discovered in Chapter 5 that the sonorous scope of the network is limited, and it is the human blocks which appear to have the most agency in the structure of the piece. Here, we understand that it is the broad organisational structure that has a high level of agency over the red blocks, suggesting a linear, fixed state of the network during performance.

In Constanzo's network I take the example of the various effects modules that have been examined across this thesis. The Record is the musical event that occurs when sound is passed through, transformed, and output of one of the modules. The points of interface are the composite modules themselves. In terms of measurement, we see a situation like that of Pasquet's Modalys instruments: the sound is conceived of through the visual feedback of the state of parameters.

I consider this instance at this level of abstraction because this Record is taking a stream of audio that has already been configured in some way across the

organisational dimension, and manipulates its energetic content, rather than being at the source of an audio stream as in the case of the Modalys instruments. These modules are also operating on a rolling buffer of recorded audio, hence interweaving content from the organisational and temporal dimensions through slice selection which extends the audio from a fragment of sound to a musical event. There are also manipulations that reconfigure the incoming audio across the temporal dimension, such as with the *cloud* module which triggers multiple bursts of audio.

Manipulation is deployed through the various processes, each module considered as a composite point of interface that translates agency from the sonorous domain to the sonorous domain. The settings of the modules pass either through preset configurations, or through manipulation of another point of interface – the crossfader (I return to this below).

Observe how these disparate Records are configured within the software: compared to some of the other case studies, Constanzo's network is clearly grammatised, and each centred Record is conceptually and visually contained within its own space. There is a clear flow of Records being measured and manipulated, new Records being created, passed on, and measured and manipulated once more. This refers back to one of the initial streams of musical thought suggested in Chapter 2: the idea of streams of audio – separate, linear paths that are followed and that come together at certain points to constitute the final outgoing sound. This is characteristic of Constanzo's practice which is notably constructive in nature: unlike a network such as Hayes' where sound seems to precede her intervention, here sound is constructed by Constanzo and is fused together into a final composite object.

Looking to Hayes' network, I evoke a similar example which is the external pieces of hardware. Notably the sonorous ones: the Moog, Digitakt and TC-Helicon. Despite the interconnected nature of the network, these elements are conceived of as residing in their own corner: they all have their internal working logic, they are connected to the rest of the network through thin strands of edges. The conceptual space between them and the centred part of the network, the Max patch, suggest that the audio they propose and the agency they impart within the network can be seen as musical events, not just fragments of sound. There are examples where the measurement of these Records is dealt with in a more atomic way through transient detection, however, I argue that during the performance, Hayes will consider these composite material points of interface as translating full musical events.

During performance, Hayes lends little attention to these objects: she is fully engaged with the game controller, the Korg and the computer screen. We know that events in these objects are being triggered by the patch, and incoming audio is dealt with in the patch essentially via the game controller point of interface. This suggests that Hayes leaves these elements running in a somewhat autonomous manner. Gestures will eventually be translated into imparting agency upon them, but in a very indirect manner, filtering through numerous points of interface.

This prolongs the idea I have formulated of Hayes intervening in a sonorous network. The isolated nature of these Records and the indirect nature of the measurements and manipulations that are engaged with them suggest a structuring stand of musical thought that calls to notions of assemblage, interconnectedness and parts of the network which will make decisions in a manner disconnected from the performer.

In Pasquet's work, the Record I examine is the main generated rhythm. This is a Record from which much of the organisational structure of the piece is derived. Pasquet retains a notion of how this rhythm will broadly manifest itself through the basic rhythm that is used to generate the new tree. Measurement of this record physically manifests itself through LLLs and can be visualised as proportional tree structures. Manipulation is affected through weighted random modulation of these lists.

This informs us about several notions concerning musical thought: Pasquet conceives of the broad organisational structure of the piece primarily across the organisational dimension; the music can be conceived of in a way that would be like a Schenkerian approach to harmony in regard to rhythm with various levels of abstraction always stemming from a basic structure. Each strand of rhythm is synchronised together. This attests to a great deal of control and a high degree of agency from the performer/instrument maker. Contrarily to a network such as Constanzo's, here there is no mechanical scope for organisational divergence from the content that has been inscribed by Pasquet through configuration of the network.

3. The Record as aesthetic object.

Finally, I propose a third level of abstraction that considers Records in terms that may extend beyond what is traditionally considered musical. Here I wish to examine how far the model can be taken – with some of the examples, this section explores the idea

of taking elements which may not appear inherently musical and grant them a certain musicality through the nature of their measurement and manipulation.

I start with an example in Burton's work which can still abstractly be considered musical: the Record as the broad organisational structure of the piece itself. It is something that is clearly transcribed and materialised within the network. Through the various points of interface, notably the *function* objects and the diffusion of the score, this is an object that Burton engages with in several ways.

Measurement of this broad organisation structure is visualised through the aforementioned points of interface, but also through the iterative process of rehearsals and software development. All the decisions that Burton makes when playing the instruments during performance can be considered as manipulations of this Record.

In terms of musical thought, we return to the idea of this Record's high degree of agency. Having examined the agency it deploys over all aspects of the network, can it be considered as inherently musical? It conveys a structure that is destined to be interpreted in a musical way, however, this structure is inherently silent and non-sonorous. Development started with the score generation system; this in an abstract, out-of-sound way. It was the visual aspect of the score that was initially important, which was then filled – as it is in performance – with sonorous interpretations of its elements in a second moment. Ultimately, I argue that this is a musical Record, but one that is constituted of an inherently aesthetic – rather than musical – structure.

Another Record we can consider to be seen as constitutively non-musical is the crossfader in Constanzo's network. I have identified it as a point of interface, and examined how it intervenes in the manipulation of other Records, but can we consider it as a Record in its own right? The development of this object was interwoven with constant feedback of musical thought – through an arduous testing process, Constanzo refined the object so that it may translate his gestures at as high a resolution as possible. It is manipulated in several ways: direct mappings, measurement of its speed, mappings to several parameters at once. But when we look at the activity that occurs around it, is the manipulation musical, or gestural? Is it possible to conceive of the musical and the gestural as being inherently interwoven?

I argue that in this case, they are. Constanzo's musical practice is an expanded world of object fabrication, programming, gestural content, sound generation, sound modulation and performing. The crossfader can be conceived as a musical Record because the gestural data it translates is inherently interwoven with musical thought

in the context of this network. The musical thought that is inscribed in this part of the network could be articulated in a conception of gestural vocabulary as being inherently musical – that sound and movement are synonymous. It is not surprising to see this from someone who was trained as a percussionist.

Going further still from the inherently musical, I would like to discuss the idea of Hayes' body and face as a musical Record. This is a Record that is primarily engaged with, although not exclusively, by the audience. I have remarked how the facial expressions and the choreographic evolution of Hayes' posture seem to be driving perception of the piece. Let us consider this in my theoretical framework: first, the point of interface would be interwoven with Hayes' biological body. The audience has a hypothetical idea about the performer as something from which sound emerges⁴, which is interwoven in musical thought. Measurement emerges from a reading and an interpretation of Hayes' expressions and gestures. Manipulation would be emerging from Hayes, imparting agency upon the audience to guide their experience through the network.

It is the tight relationship with sound that gives me grounds to consider this as a Record. Once again, there is nothing that could be called inherently musical in the composition of this Record; however, all the engagements that are suggested are intertwined with musical thought. Hayes clearly accords a strong place to the materiality of her body on stage within the music. She offers her body as an object for interpretation and uses it to drive perception.

This is the case for all performers. However, in my other case studies, the bodies of the performers are consumed within the network – Constanzo becomes part of a composite object with his drum set, Burton keeps his back to the audience and pushes focus towards the score, and Pasquet is sat behind a desk in darkness, throwing focus to the physical space around the audience. I take the example of Hayes' network as the treatment of the body as Record is particularly striking and clear, notably its immediate and exposed relationship to the sound. However, the ways in which the bodies of the other case studies are configured are no less interesting – they all drive the perception of the piece in their own way.

Finally, I examine the idea of space with Pasquet's work. Is it possible to conceive of the performance space as a Record? This is something that preoccupies Pasquet, as he

⁴ Hayes is the only one of my case studies who uses her voice during performance.

explained in an early interview⁵. Points of interface here are exploded across many: as we have seen there are the lights and the various speakers around the space; there are also constitutive nodes of the performance area: the space itself, beanbags, the disposition of the various stages, the position of each human actor within the space and regarding each other. If we take the space as a Record engaged with by Pasquet, there are concrete examples of manipulations: the use of multi-channel output, the use of the spat5 reverberation object moving in a virtual space, mapping to the lights across the space.

If there is musical thought to be found here, it is intertwined with agency from a large number of actors. For example, Pasquet had no say in the disposition of the stages or beanbags. It can be said, however, that Pasquet conceives of a hypothetical space when configuring the network through the various multichannel operations within his patch. The question of spatialisation in musical thought is a vast one, and one that certainly goes beyond the scope of this research. In this specific case, we can observe that there is somewhat of a dissociation between generation and organisations of sound, and its diffusion through the space. Sound is first conceived, then this is taken and translated by the various spatial points of interface. Spatialisation is evidently an important aspect of Pasquet's practice, however it seems unreasonable to consider the space as a Record which is inherently musical.

⁵ Appendix 8.2.2.

CONCLUSION

Analysis summary.

To conclude, I will go over the various steps of the analysis and highlight the key ideas that emerged. After this, I return to my three initial research questions and discuss the extent to which they may have been answered by my analyses. Finally, I take a broad look at the methodology that has emerged and suggest its advantages and limitations.

Analysis recap.

I started by highlighting some of the major themes in contemporary musicological analysis, subscribing to a turn away from Work-centric approaches, looking to interrogate the drastic nature of music (Abbate, 2004) and understand music as a process (Small, 1998). I looked to various perspectives where networks are an essential part of analysis, notably with ANT (Piekut, 2014) and Ingold's *meshworks* (Ingold, 2008).

Drawing from these approaches, I presented a theoretical framework within which I could examine my case studies. I proposed the perspective that musical practice is the configuration of, and existence within networks. To focus my analysis on something tangible, I proposed looking at a certain type of actor: the *Record*. This is a flexible object for analysis, a conception of sound in a broad sense with which musickers engage through *measurement* and *manipulation*. I presented my conception of sound being measured and manipulated across three dimensions: temporal, energetic, and organisational – an elusive and non-empirical dimension of sounds being deployed in a musical context, in the context of themselves and other sounds.

I presented the three research questions which would drive the research. First, I wished to inflect an aspect which I found to be untreated to its fullest in the network approaches I had discussed: that of time. I wished to consider the network not as a static account of what had happened, but as a structure which allows for things to happen. In my musical context, I suggested that it would be useful to access an image of the entire scope of *sonorous potentialities* a network could create.

I then presented the idea that *musical thought* – the intangible way a musicker conceives of music – becomes inscribed within the network. It materialises itself through the physicalities upon which it depends, through the engagements of measurement and manipulation that the musickers allow to be performed. Finally, I

suggested that these physicalities could be examined through the lens of *points of interface*: physical actors in the network which translate agency from one domain to another.

Analysis started with a set of surface level analyses of the performances, the material setup, and the code for each case study. This allowed for two things: to discuss the networks in an articulate manner with precise segmentations of the performances and a comprehensive understanding of the mechanical aspect of the instruments; and to identify strands of musical thought which would guide the following analyses.

The nature of these strands differed for each case study: I discussed the sonic nature of each performance, the ways the musickers carried themselves on stage, the nature of the objects around them and the ways they were configured, and the apparent structural aesthetic ideas for each network.

Next, I presented construction and grammatisation – “*the process whereby the currents and continuities shaping our lives become discrete elements*” (Stiegler, 2010, p. 70) – of network graphs at various levels of abstraction. I explained that grammatisation is inherent in any network, and even at a very detailed level of abstraction, it is something that is constantly engaged with by the analyst when constructing network visualisations.

I began by looking at the singular global level network. At this broad level of abstraction, I was able to interrogate the different institutional agencies that were in circulation. For example, the ERC deployed *permissive* agency through its financial backing; HCMF deployed agency that was permissive, but also aesthetic through the cultural baggage it deploys and the nature of its audience; the numerous universities and research projects deploy agency through the natures of their research interests and approaches to musicking.

In keeping with the heterarchical nature of the approach, I proposed the notion of *centred* nodes. These were actors with high degrees of edges and much reach around the network. One example of this was the *FluCoMa Project* and its immediate neighbours. I examined the nature of these neighbours, which led me to examine the idea of agency being filtered through nodes (examined again when looking at points of interface). I discussed the different ways in which the FluCoMa project’s agency was constituted and translated: *towards* the musickers through the toolsets, *from* the musickers through their pieces.

Analysis continued with local (centred around the two concert spaces) and instrumental (centred around the material setups and the software) level networks. These revealed the material structure of the networks and indications about critical points of interface that would be returned to in Chapter 6.

At the local level, I was notably able to examine the positions that various actors hold within the network. This could span from an actor such as Constanzo, who was demonstrated as being intertwined within the immediate materialities of the network; to Pasquet who occupies a much more detached, fluid position like that of an audience member; and that of the analyst who can fluidly focus their attention on any actor at any time. I discussed the consequences this has on the type of musical practice that can emerge – notably the gestural nature of the performance and the articulation of space.

At the instrumental level, I first took some examples of static networks, and used statistical evaluations to inform me about the musickers' approaches. For example, the quantity and types of various FluCoMa tools could indicate what parts of the FluCoMa workflow were important to the musickers and informed me about their practice. This led to surprising results: for example, musickers like Constanzo and Hayes who appear to operate in real-time use much more of the non-real-time objects.

I then looked at some examples of the development of networks over time, examining how the structure evolves, and identifying key points in the timeline. For example, the order in which Burton proceeded over development: starting with the score generation, then adding processing as simple playback of events, then finally adding more bespoke processing towards the end. With Pasquet's case, I was able to observe the development of the various ways of controlling the Modalys instruments which was discussed again in Chapter 6.

In Chapter 4 I approach the idea of sonorous potentialities. The core problem addressed in this chapter was that of source collection: how can one proceed to gather a collection of sounds that is representative of the entire scope of sonorous potentialities of an instrument? Across four case studies I demonstrated first an *iterative technique*, then three approaches to *alternative performances* techniques.

The iterative technique, performed with Constanzo's work, seeks to explore the entire parametric space of a network, iterating through all possible states and sound sources. I identified the different signal paths in the network and took representative sound sources from the various sample libraries and live recorded material. At each step of

the signal path, it was necessary to cull results from the previous step to a small number of representative sounds.

This was achieved by building 2D sound plots using audio descriptors and dimensionality reduction. After much experimentation of different combinations, I found that the MFCC audio descriptors and t-SNE reduction algorithm gave me plots which allowed me to quickly find clusters of sounds with similar timbral profiles. This allowed me to gain a good understanding of the various sources and the three key effects modules in the network; however, the further I went down the signal path, the less reliable the analysis becomes.

I addressed these issues with the idea of alternative performances. This approach not only reduces the amount of audio to analyse, but also retains a sense of the network's structure by examining sounds across organisational as well as the temporal and energetic dimensions. The first instance of this approach was performed with Hayes' work, taking rehearsal recordings, and combining them with the performance audio. This simple configuration of the alternative performance approach led me to investigate automatic slicing algorithms to constitute the sound plots of these corpora that would be investigated in Chapter 5. Again, I found a novelty slicing on MFCC content to be the most effective.

Then, I took the approach of reconfiguring entire performances in analytical conditions of Burton and Pasquet's work. I set up several configurations which progressively moved further away from the configuration in performance. For Burton's work, I began with alternative performances using the same human and machine scores; then the same human score but a newly generated machine score; and both scores newly generated. I again used MFCC novelty slicing to constitute the sound plots of these performances which would be used in Chapter 5.

With these automatic slicings, I gained a good picture of the timbral space, however it appeared that content in the organisational dimension and the gestural nature of the music became somewhat lost. Therefore, in the final iteration with Pasquet's work, the musical thought around gestural slicing that was inscribed into the network was accounted for as part the source collection process: I identified what appeared to be gestural slicing with the beginning and end of playback of the main generated tree and used this to break up the performance. I did three different configurations of performances: first using the same settings; then using different proportional rhythms

for main tree generation (1-1-1-1 and 2-1); and a configuration where all the top-level parameters are randomly modified over the course of the performance.

This left me with a large set of sound plots, all allowing for different kinds of approaches and revealing different kinds of elements in Chapter 5. This is where my cartological, interpretative analysis began. To approach these artefacts in a structured manner, I borrowed notions from the fields of cartography, rambling and video games to define some principles of navigation: *abstraction*, *clustering*, *landmarks*, *movement*, and *ontology*. I then apply these principles to three different types of sound plots: *static*, *superimposed*, and *dynamic*.

In a static sound plot, time is conceived of as unidirectional or out-of-time. In such a plot, it is possible to construct clusterings of sounds and to draw conclusions about their ontologies. I was able to give an account of the spectromorphological profile of the slices found in Constanzo's performance and noted for example the even distribution of sounds from dry acoustic slices to heavy distorted processing. This allowed me to inflect one of the first identified strands of musical thought: an exploration of distortion in various forms. As useful as this static network was, other navigational principals such as movement and landmarks were absent from the analysis.

Superimposed sound plots conceive of multiple temporalities at once and allow me to examine a wider range of these navigational principles. I was able to propose ideas around musical thought: for example, in Constanzo's network the idea of movement and transformation of sounds from one timbral profile to another through different configurations of effects modules; or in Hayes' network the idea of sounds imitating other sounds from sample libraries.

I was also able to examine the scope of sonorous potentialities and see to what extent different configurations of a network could affect the constitution of these spaces. For example, in Burton's network, it was discovered that the different machine scores had very little agency in the timbral profile of the piece compared to the human scores. I was also able to gain a better understanding of how the musickers' composed the networks: for example, it was seen in Pasquet's network that the random modification of parameters steered the network away from some of the more chaotic gestural motifs, suggesting that these were precisely engineered by Pasquet.

Finally, I presented the idea of a dynamic sound plot – allowing an examination of the navigation of a network in time. This allowed me to perform a choreographic analysis

of movement over Constanzo's performance, with the identification of large and local scale gestures in the space. This presented itself as a promising avenue for future research.

In Chapter 6, I returned to the physicalities of the network and explored in more detail some of the points of interface that were drawn from earlier analyses. I proposed a typology for talking about points of interface in an articulate manner that draws on the earlier classification of network nodes. The points of interface can be considered by their ontology at various levels: material, digital or biological; they can be conceived of as singular or composite (comprised of many actors); they can be discussed in terms of the way they translated agency – linear (on a bounded infinite plane) or stepped (a pool of finite possibilities) and through interpolation or immediacy; modes of interaction can be direct (agency is immediately translated through one point) or indirect (agency is translated through a chain of points). A point of interface was described as translating agency from one domain to another – this translation is understood as encryption and interpretation of data, and the content will always be disfigured in some way, having suffered a loss of resolution.

With this typology, I presented an analysis of some of the points of interface that had previously been identified: first through the lens of aesthetic versus functional points. Across several examples, I discovered correlation between the constitution of the points, and the ways they are used: aesthetic points tend to require supervision to hold their states and primarily translate agency from the gestural domain; functional points can retain their state and translate agency from the conceptual domain. Aesthetic points also tend to receive much focus in the development of networks, and their physical constitution can structure the broad musical thought in the piece: for example, Constanzo's crossfader which "*commands a type of language*"¹, or Hayes' game controller (see below).

Then I examined the idea of hybrid points of interface, allowing me to observe how a network would reconfigure itself at different moments. I gave a detailed account of the example of Constanzo's condenser microphone and the snare surface and the various configurations that appear in the performance: it became apparent that an important factor for identifying a hybrid point of interface was the possibility or not to disentangle agency from the various objects. I went on to look at other examples:

¹ Appendix 8.4.5.

Constanzo's crossfader, Burton's score, performer, and instruments, Pasquet's sound output and lighting setup. I also discussed how these points are perceived by various actors: the analyst, the musicker and the audience.

Finally, I opened a discussion around code design, and proposed the importance of the notion of *resistance* in interface. This followed from ideas from Serres (Serres, 2007) where if a relation between two actors is "*optimum and immediate; it disappears as a relation*" (Serres, 2007, p. 79). To ignore resistance in interface is to ignore the essence of interface. It is also a means of addressing materiality in a digital context. I examined how the musickers treated this question in an aesthetic way: for example, Hayes' configuration of the game controller and a slight shift in its held state that can have drastic consequences on the sound; or in the constitution of the code itself, with the example of Burton's code building up iteratively and reaching a point where it shifts into a system against which the musicker must struggle to impart agency.

In Chapter 7, I drew on elements from all these analyses to address the question of musical thought and the Record. I approached each example with a set of elements to identify: the nature of the Record in question, the points of interface that allowed for engagement, the nature of the various activities of measurement and manipulation across the energetic, temporal, and organisational dimensions, and finally the way in which the musical thought that is conceived of as being inscribed in that part of the network could be articulated.

I first considered the Record as a fragment of sound. With Burton, I explained how each slice of outgoing audio is measured in the energetic and temporal dimensions and used to configure the piece across the organisational dimension. This reflected the ideas of spectral balance which seemed important in his practice. With Constanzo, I took the example of the snare drum and its various transformations, allowing me to discuss the reasons behind them and suggest a conceptual separation between his gestural vocabulary and sound they may produce.

I also took the example of Hayes' NMF drum machine, which demonstrated a use of samples different to Constanzo: rather than providing audio content, they are taken as gestural shapes to be imitated. Finally, I took the example of the hypothetical sound of a Modalys instrument in Pasquet's network and remarked the interesting configuration of measurement and manipulation through the spatial parameters of the instrument's physical body.

Next, I considered the Record as a musical event. With Bruton, I took the example of a red human block. This allowed me to discuss the idea that Burton created an abstract shell to be filled, noting that the local sonorous events that filled these blocks didn't influence the broad organisational structure of the piece which was decided ahead of time. For Constanzo, I took the example of an effect module. The musical thought that was interesting here was the clear grammatisation present within the network, and a conception of sounds as being created, modulated, and passed on, rather than a continuous flow.

For Hayes, I looked to the external synthesizers which prolonged previously discussed ideas of assemblage, and Hayes intervening in a network which has its own agency, which will sound of its own accord. Finally, I took the example of the main generated tree in Pasquet's network. This actor, and the way it is treated informed me about Pasquet's conception of his piece: like a Schenkerian reduction – stemming from a basic rhythm at a background level.

Finally, I examined how far the model of the Record could be pushed, considering it as a broad aesthetic object. With Buton, I considered the broad organisational structure of the piece, materialised notably through a series of *function* objects. This allowed me again to question the balance of agency between this silent structure and the sounds that would come to occupy it. With Constanzo, I considered the crossfader and discussed if it were possible to consider this agnostic artefact as inherently *musical*, as a Record. In the context of this network, I argued that it was, considering the nature of its development and how its physical constitution influences the musical thought inscribed in the piece.

Then, I examined the example of Hayes' body and face as a musical Record. I examined these as points of interface, and how they drive the perception of the piece. Finally, I took the example of the performance space as Record in Pasquet's network. The conceptual distance between generation of sound, and its consequent distribution in space led me to conclude that it was difficult to argue for this as being an inherently *musical* Record, however it was clearly an important part of Pasquet's network.

Research questions.

I now return to my initial research questions and examine them through the lens of these analyses. The first was to address the notion of time in the network approach, which I suggest is yet unexplored to its fullest potential. Notably, can we work

towards the ambitious goal of analysis that draws on the entire scope of potentialities of a network – in my musical context, the entire scope of *sonorous* potentialities?

I stress that this idea is not to be taken literally: the potentially infinite number and nature of variables and parameters that compose a network in the real world would render a literal interpretation of this approach physically impossible. It is necessary to work with these limitations and make informed decisions about what to include or not in a source collection. One major factor that helps with this is to conceive of the network as being shaped by its use – it is something that is configured and engaged with by specific human actors; to ignore their conception and use of the network would be to miss an important constitutive part of its structure. The question then becomes, what is the entire scope of sonorous potentialities of this network as it is configured and deployed *by the musician*?

Over the course of analysis, it became apparent that in some cases, it was more interesting to approach this question in terms of the generative process. With the first Constanzo and Hayes examples, the sound plots that were produced were somewhat sterile in nature – I was able to perform analysis of the sound plots that enabled me to understand the timbral nature of the different sounds that could be produced by the network, create groupings, and begin to examine the shapes of some processes; however, it seemed that this is as far as the method can take us.

The way the alternative performances around Burton and Pasquet's work were configured allowed me to perform this kind of analysis and more. Approaching this question by configuring alternative performances not only informed me about the scope of sonorous potentialities; it also gave me a much deeper understanding about how the network worked, and the scope of the sonorous potentialities in the context of the conditioning system.

This provided an interesting perspective for analysis, however the methodology of its deployment must be nuanced: for analysis it is useful to be able to compare different points of reference, to intelligently define corpora that, when distributed together into a 2D space, will inform us about the aspect of the network.

Next, I proposed the idea that musical thought is inscribed within the network, and that approaching analysis from this perspective would provide tangible elements that would allow me to discuss musical thought which can often escape writing. One goal of this research was to adopt this perspective, and the Record model and examine their validity.

In practice, the idea of musical thought is something that was mentioned throughout each analysis. It was a useful tool for discussing the results of each process in a coherent way across the project. However, to articulate a real discussion around this notion, there were a certain number of prerequisite analyses that were necessary to perform. For example, it is not possible to properly discuss the Record without a comprehensive understanding of the points of interface it involves; it is also necessary to have a comprehensive understanding about how the mechanics of a network work.

In Chapter 7, I addressed the notion of the Record explicitly, but it did intervene in a fundamental way across the research. For example, would the construction of sound plots have occurred in the same manner without the knowledge that the analysis was building towards an interrogation around the Record? Do the network visualisations produced in Chapter 3 make sense without this underlying perspective in mind? The idea of actors which engage with each other in various attitudes and the translation of agency are fundamental principles of these ideas. Even if the model of the Record was only explored explicitly at a late stage in the analysis, it is effectively what has structured the analysis as a whole. Also, despite the structure of this thesis, in practice analysis did not proceed in this linear fashion: each section developed alongside the other, with analyses constantly informing each other.

This still leaves the question open as to the justification of adopting this perspective. As I hope to have demonstrated, the methodology and the theoretical context of my research was drawn directly from contact with the case studies. Many fragments of the various practices that I have encountered seemed to suggest to me this kind of model. The notion strikes me as being explicit in the digital domain: in a typical Max patch, the world is populated with objects. All these objects are conceived of as pertaining musical knowledge in some way – seeing them on the screen, we have an idea in our minds about the kind of musical agency they deploy. The digital domain is populated with data, and functions which measure or manipulate this data. Can this idea be transcribed into the material and conceptual domains?

I believe that I have demonstrated this to be the case, at the very least with the networks I have examined. This perspective has allowed me to ask interesting questions and have the desire to develop and deploy methods that have created interesting new artefacts that allowed me and the reader to experience the networks in a way that suggests new listenings. My goal is not to generalise the Record model – there is certainly not enough evidence here to suggest that all music must be

regarded in this fashion. However, it is a tool that is open enough to be applied to different case studies – and perhaps in other contexts – that will certainly lead the analyst to pose articulate questions.

Finally, I wished to interrogate the way musickers configure networks. I suggested that a fruitful way of approaching this question would be to look at critical points of interface. Across the case studies, a whole host of *ways of configuring networks* has emerged – there are as many ways as there are musickers. Are there any general principles that could be drawn?

One recurring element observed were the differences of engagements with different types of point of interface. I proposed a differentiation between aesthetic and functional points of interface and remarked the correlations between the types of engagement and the domains from which agency is transferred and the physicalities of the object (see Section 1.1). There also tend to be correlations in the resolution of the data they can translate: aesthetic points of interface have high, sometimes linear, resolutions; functional ones lower. Another correlation can be found in the mode of state changes: aesthetic points are often interpolated; functional points are often immediate.

Another general principal that can be observed, this time across many of the sound plots, is the nature of the landscapes in terms of groupings of nodes. Especially when looking at superimposed sound plots, I often remarked that there were rarely more than two or three different large clusters of sounds – this was true for both timbrally and gesturally-structured plots. Often, there is a main, large cluster with an evolution of the aspects of its sounds across various axes, and smaller ones which are isolated or branch off from the main cluster in some way.

If we interrogate why this may be the case, it must first be noted that the pieces asked of the musickers were relatively short. However, it is interesting to consider the idea that these networks are often primarily constituted as focusing on the many ways of articulating a certain timbre or a certain type of gesture that all share inherent characteristics; alongside this, there will be smaller groups of sounds that are also explored to a certain degree that will contrast in nature to the main cluster. This may not be immediately discernible on a first listening of the piece – especially when the various manipulations concern ideas in the gestural domain across the organisational dimension as was the case with Pasquet. Further case studies would be required to verify the tenacity of this idea, but it is an interesting notion. It is something that would

intuitively seem to make sense in the context of musical practices such as these where instrument making plays a large role.

Finally, we can discuss the idea of materiality. Here, this term can be seen as opposing the digital and has been articulated in several different ways across the analyses. It is a question that is often thrown-up when discussing the practices of musickers where a lot occurs inside the computer. It seems that the contemporary paradigm sees a need for these types of musickers to bring back a certain materiality into their practice. It would seem a necessary part of musicking, that a performance without materiality is somehow alienating, or lacks a certain instrumentality (Hardjowirogo, 2017) that is conceptually interwoven with music.

I can observe this question in my case studies. There are musickers where this does seem to be a priority: Hayes, for example, has explored this notion across her whole career and brought another perspective to this project. She was searching for a resistance in her network (as we saw, in the points of interface, in the general aspect of her network and even in the design of her digital interface).

What about the practice of someone such as Pasquet? Certainly, it could be said that the idea of instrumentality seems to be absent from his practice when watching him perform, but what of materiality? I would argue that Pasquet's network is the one which takes the idea of materiality to its heart – it is a network which explicitly reconfigures the space around him and the audience. It could be argued that to see him sat behind a computer – with no perceivable link between his gestures and the sound – that his performance may have been alienating. However, this was far from what I personally experienced during his performance. To be in such an immersive, engulfing space led me to experience every aspect in an aesthetic way. During none of the other performances did I feel quite as submerged by the musical expression Pasquet was conveying, and so grounded in a material, tangible music.

This remains an open question; however, from the case studies that I have treated over this research, I can say that they all treated materiality in different ways.

Methodology.

I finish this formal conclusion by looking at the methodology that has emerged from this research and examine how it could be used in the future as well as its limitations. We can start by drawing a broad abstract picture of the various steps.

The goal of this methodology is to find ways of tangibly visualising and manipulating social, musical networks. This suggests two main activities: network visualisation and navigation; and source collection. Network visualisation and navigation suggest a need for dynamic ways of exploring data structures. There have been two main instances of this across this thesis that are informed by the types of networks: visualisations that show the physicalities of networks, and visualisations that show the sonorous content of networks.

For showing the physicalities of networks, I used a pre-existing software: Cytoscape. This gave me a robust system that allowed for network construction within the program, and for importation of networks of other formats. The import system was open enough for me to be able to construct tools for translation of Max patches. I could also export networks constructed in this program to an HTML interface. For showing the sonorous content of networks, I used 2D sound plots that were populated by slices of sound. This was initially done in Max but was also translated to a robust HTML environment.

Once these visualisations are constructed (see source collection below), I explored many ways of navigating these networks. The physicality networks could be distributed according to several different layout algorithms, each revealing aspects about the network's structure. It was often useful to examine nodes which appeared to be centred, and paths between critical nodes. It is also possible to run searches, apply filters and display the networks in different ways according to which aspects the analyst wishes to demonstrate.

One limitation is that these tools are all contained within Cytoscape. It is possible to display networks within Cytoscape's JavaScript system, however the standalone program is required for dynamic and fluid analysis. With these visualisations, it is also possible to gain a picture of the network very quickly and they allow for processes such as observation of the general shape over development iterations of a network. This is useful – however it was necessary to export each network as an image and composite the video myself. It would be useful to have this kind of process under hand in the working environment.

I proposed a set of principles for navigation of sound plots. These principles offered a methodological way of approaching an abstract space. Solutions need to be found given the nature of these plots – for example, I discuss the idea of *meaningless axes* in the next section of this Conclusion. The sound plots offer ways of quickly navigating

large collections of sounds and can propose groupings that would be difficult to make by simply listening through each of the files. However, there were some technical limitations to this – the primary one being the amount of audio that can be loaded into memory: a more atomic approach would be better than the flattened-corpus approach I used.

There are also choices to be made regarding the navigational possibilities given to the user: in my implementation, the user had to click on the sound to trigger playback. I made other prototypes where the user could click anywhere in the space and the nearest sound would be played. This point and click system is something that could be examined, and have alternative solutions offered to. This way of navigation offers a high level of control to the user which may appear useful for analysis; however, how else could one approach the navigation of this space? For example, if the user was constricted to physically moving a point in the space that would be controlled by arrow buttons or a game controller, maybe this could allow new ideas to emerge regarding the aspect of the network. This also brings into question the 2D form of the sound plots – there are evidently many more ways this kind of idea could be implemented.

Source collection was also a large part of the methodology and had much more sway on the results of the analysis than initially anticipated. First, with the physicalities visualisations, it was necessary to transcribe real world networks into a very abstract form. As we saw, the analyst is constantly in a process of grammatisation, and must always make decisions about abstraction and degrees of resolution. The tools for transcribing Max patches were useful, but the system could be more streamlined regarding Cytoscape². It was possible to retain information regarding patch hierarchy with abstractions and subpatchers, however the user is required to create groups of nodes manually in Cytoscape if they wish to view the network at a higher level of abstraction. A large part of source collection for the methodology must also be performed in the field, collecting meticulous accounts of hardware setups and performance spaces. It is also necessary to gain a knowledge of the various institutional actors that intervene in the network.

² Once the user has run the patch through the transcription script, it is necessary to import it into Cytoscape via two different node and edge lists, which then need to be configured to display the information in a form that is readable and calls back to the patch.

Source collection regarding the sonorous content of networks depends greatly on the nature of the network under analysis; however, I did reach some conclusions to help future research. Ultimately, it seemed that the iterative method through the parametric states of the instrument was problematic. This is mainly due to the vast amounts of data that it entails, but also because it evacuates a large part of the network's nature that is revealed through its use. This method could be applied to less complex networks and would require a more robust navigation system. It is the method that offers the closest idea of the entire scope of sonorous potentialities, but as we saw, it can be misguided to aim for this.

A balance between the infinite plane of potentialities and the network deployed through its use offers a picture that not only explores the sonorous potentialities of the network, but also retains elements of the networks' organisational structure within it. This was the goal with the alternative performances approach. The main issue here was around slicing. Both automated slicing and supervised slicing offered different perspectives, both useful in their own ways. Ultimately, I believe it is more useful for analysis – and provides a more informed picture of the structure of a network – to proceed with a supervised slicing. This means that during source collection, it is necessary to identify inscribed musical thought around grammatisation and understand how disparate musical events are perceived by the musicker.

This research has provided a good methodological starting point for the use of a network approach in musical analysis. Further work is required, but I certainly believe that I have opened-up interesting avenues for research, and covered the ground work for a set of tools that is informed by real analytical practice.

The ontology of networks.

Meaningless axes.

Before concluding this thesis on a more open-ended discussion, I will address some of the issues that have arisen throughout the course of research that would also serve as interesting subjects for further exploration. The methodology strives to create new artefacts that are inserted into the network. All analysis can be seen as doing this in some way; however, here is a very particular situation in as much as analysis is no longer directly focused on the object that is the music itself. Analysis can be conceived of as focusing on the *musicking*. When the music-as-noun as a subject for analysis has

been evacuated, it is necessary to identify what objects we can focus our attentions on. The perspective of music as process, as verb, expands greatly the potential candidates. I wanted to achieve an analysis that would still focus of the Work of the musicker, disentangling the Work from its traditional signification. In Chapter 1, I suggested that the Work is the artistic *practice*, and manifests itself through the configuration of networks and existence within these networks.

However, a network and an existence within a network, are intangible objects. The network is considered not as an account of one social occurrence, but as a system that allows for things to happen according to its aspect and internal structure. This is a candidate for analysis; however, before analysis can commence, it is necessary to create artefacts that will then become the subject of analysis. An important job for the analyst is to ensure that the artefacts that are created are objective and representative of the network in question.

As discussed in Chapter 5, these representations will only ever be abstract *maps* of a network. However, the methodology I have proposed allowed for the creation of informed and representative visualisations. Some aspects of these visualisations demand closer inspection. The first aspect that can appear problematic, is the absence of meaningful axes. This is the case for both networks representing physicalities and for sound plots.

Networks of physicalities could be distributed according to several different types of layout algorithms. These algorithms reveal things regarding the structure of the network. In the real world, these nodes are spaced. Even the digital nodes from Max patches can be placed in a spatial dimension. When using the layout algorithms, this spatial content is lost. The question is, is it something that needs to be considered for analysis? Indeed, it may be important to know if one critical point of interface is spatially close to another or not.

This aspect was somewhat absent from my research, and further research would be required to examine the importance of this notion. That being said, nodes that interact directly with each other do tend to share edges and will tend to be close to one another in most layout algorithms. However, as soon as more than two nodes are concerned, distances can grow substantially. Therefore, I focus on things like the degree of edges rather than spatial distance. One solution could be to use dimensionality reduction algorithms, like for the sound plots, where real spatial measurements could be dimensions. I am dubious about the effectiveness of this approach: many numbered

dimensions would have to be defined, and there would be non-linearities between different types of nodes³.

Then there are the sound plots. With previous methods and tools, such as CataRT, we have been accustomed to assigning specific scales to the x and y axes. It can appear bizarre, then, to focus on plots which do not have inherently meaningful axes such as the ones procured from dimensionality reduction. This idea becomes even more strenuous when one considers that different algorithms will produce different distributions.

This is something of which I was conscious during research, and which I came to terms with. My approach is to create artefacts that enter the network in question that are representative in *some way* of its structure. The goal of analysing the artefacts is for me to gain new perspectives, and not to find specific things. A network may be represented any number of ways. My criterion is this: if a representation allows me to pose articulate questions, then it is an object that is useful for analysis.

I did try and keep a uniformity across the different examples, principally using t-SNE and UMAP reductions, and MFCC descriptions and slicing. I do not go so far as to argue that this allows for comparison between plots. However, it does mean that they all yielded representations that allowed me to deploy a same and coherent methodology. Ultimately, the sound plot is to be regarded as an artefact unto itself that allows us to conceive of what it represents from a new perspective.

Navigation of networks

I also proposed a set of principles for the exploration and navigation of these artefacts. I implemented a simple point and click system for exploring the sound plots. This is a skill that most people have developed through years of working with computers where it is the principal paradigm, along with keyboard typing, for navigating a digital space. The material object of the computer screen and its flat, rectangular plane suggest that fluid control over a point across two axes is a useful way of navigating space. Useful is the key word here: this can be considered as a *functional* point of interface. Indeed, when looking at the materialities, the trackpad is the same shape as the space that is being navigated and moving a point seems intuitive.

³ For example, comparing the spatial data of a material object and a digital object would be illogical.

Why is this the case in modern computers? The technology exists for us to create, for example, 3D interfaces; but intuitively, this strikes us as being less functional. For tasks such as writing a text document, doing accounts in Excel, submitting forms on the internet, the seamless, ease-of-use paradigm seems logical. However, as we have seen, there are now cases where we wish not to engage with the computer in a functional way, but an aesthetic way. I discussed how ideas of resistance in interface design appeared in the case studies; could this translate to an analytical context?

When I decided to use 2D sound plots, the decision was intuitive – it didn't occur to me to do it any other way. It was only after having engaged with my case study networks that I realised that this perspective may be limiting in some way. Before taking this approach apart, let us explore how a 2D plot with a point and click navigation system seems intuitively useful. First, human beings can only deal with a certain amount of information at once. This is the very premise of why we would turn to dimensionality reductions. When we are looking to draw classifications, I would argue that it is difficult to deal with more than three dimensions in an intelligent way – this is certainly the case for me. There is a third dimension in the sound plots – the colours of nodes reveal something of the ontology or the genesis of the slice. It could have been possible to add another dimension in the dimensionality reduction and map this to the scaled size of the node, and even another if the sound plot were distributed in a 3D space. I have experimented with this, and my colleague Laurens Van der Wee was also working on this as part of his PhD research for the IRiMaS project; however, I felt that I could not draw intelligent decisions regarding classification from it. There were just too many possible vectors to deal with.

This is only a problem when the goal of network navigation is to draw concrete structures such as a classification. This is a functional engagement with the computer. In this case, we are looking to the computer to aid us in drawing a thinking of music which is transcribable in written form. One of the initial stances on analysis announced in Chapter 1 was the idea of existing in proximity of networks and drawing new perspectives from that. It is entirely possible to envisage an aesthetic representation and navigation system for network representation that escapes reifiable thought or any kind of written form.

Over the course of the project, I have constructed prototypes for representations and navigations of networks that would escape any kind of engagement that could be written down or explained scientifically. This renders such an approach difficult to

present in a musicological context; it does not, however, render the analysis any less interesting. It would require being open to analysis as being truly performative and operating in a way much closer to musical experience. We would be abandoning the privileged, omnipotent position of the analyst and purposefully inserting resistance into the interface to gain another kind of knowledge.

The goal of this research was not to explore this idea to this extent: the concept would be grounds for further research. This is something that I am interested in developing in the future, but for now it is something that at least needs to be pointed out, as it informs and recontextualises the shape of this thesis itself.

Descriptors and perceptions

Finally, I will address some ideas around audio descriptors. The way that descriptors were used in this research was somewhat hands-off. I spent time familiarising myself with how to use them and gaining a better embodied understanding of what they could represent. However, during the research, they were used in quite a broad, abstract way where descriptors such as MFCCs were generally considered to describe timbre and spectral shape gestural content.

They are of course much more nuanced – however, this wasn't so much an issue for me as I used them to deliver artefacts without seeking to represent any specific thing. Providing that the artefact they delivered allowed me to gain new perspectives on the piece, their work was done. The same goes for dimensionality reductions – I have very little idea about how the underlying algorithms work, however, through experience, I do have an idea about the broad kind of artefacts that they can deliver.

Descriptor information is articulated with a number of different statistics. I took all the types of statistics that could be drawn from the FluCoMa tools over the entire slice to use as data. I intuitively supposed that the best approach to gain a coherent picture was to draw a large number of statistics from each description. This may not necessarily be the best approach: a different set of descriptors and statistics for each case study would certainly yield results that are better tailored for each type of audio.

It is also possible to conceive of weighting various statistics or descriptors to gain a picture which is informed more by one description than we would deem more important than another. This, however, requires a pre-requisite analysis of the general nature of the audio which went against my analytical model drawing on the notion from ANT to not evacuate elements from one's objective account, and treat everything

equally. Perhaps a second pass of sound plots, drawing from the first, would be a solution

When using descriptors, especially when paired with something like dimensionality reduction, a part of the analysis is being taken out of the analyst's hands. We put our trust in an automated process and draw from sets of numbers which require interpretation through ourselves, points of interface, to be made sense of. These numerical descriptions can sometimes seem unhelpful, or unable to convey certain ideas about music and sound that are important parts of perception. This is because audio descriptors are tools which measure sound across the energetic and temporal dimensions, but which are unable to access information in the organisational dimension.

To demonstrate this, consider a sample of a snare drum. This sample is played once at the beginning of a piece, and once at the end. The energetic and temporal content of the sound are the same, however its content across the organisational dimension is different. The sound is placed into the context of the piece, at a different time, having been preceded by different things, having been put into the context of its own repetition. This is something that descriptors are unable to account for.

I wanted to try and address this. I was inspired by a presentation that Robert Adlington gave at an IRiMaS session where he presented some work in progress around an analysis of Saariaho's *Du Cristal*: he presented a graph giving along the x axis the progression of the piece, and along the y axis a normalised scale for several different measures. The graph displayed a series of lines, each representing a notion such as *brightness*, *weight*, *tensions*, *dissolution*, or *thinness*, measured as the piece progressed. These measures are all subjective values that refer to subjective human perception of a music as it progresses. These measures inherently account for the organisational dimension: they draw from the context of what has preceded or proceeds them.

I built a prototype in Max that would allow the listener to listen to the piece and set any number of 'perceptions' with a slider as the piece progressed – there is a video demonstration of this⁴. These perceptions were recorded at audio rate, which meant that it was possible to treat them in the same manner as any kind of descriptor and insert them into a workflow such as the one presented with this methodology. We can

⁴ D.E.: 04_Videos>01_Demos>perceptions_demo.mp4

conceivably use the same slicing algorithms, the same statistics, and the same dimensionality reduction algorithms to create visualisations from these perceptions. Furthermore, these perceptions could be used simultaneously with traditional audio descriptors to retain the advantages of these precise measurements of the energetic and temporal domains and add content from the organisational domain to them.

This is a promising avenue for future research, and yet another that I wish to explore.

Future Research

I conclude by looking to the future and discussing how the research could be developed. I begin by discussing the methodology and the tools that were created. Then I look at the theoretical context and broad analytical model and how it could be taken further and into other contexts. Finally, I open the discussion by trying to propose an abstract view of the computational approach that has been taken in this research, examining what is happening conceptually at a fundamental level and looking to see if there is a general mindset that would not be fatally tied to the use of computers.

Development of the tools.

To develop the methodology further, the first task would be to reconfigure the tools. The tools that I developed to aid me in my research work well for my specific use cases, however work would need to be done if they were to be used by other people. Much of the processes happen on a CLI, much of the code is uncommented or difficult to read. It is important to bring these kinds of tools to a much larger audience of musicologists that may not be familiar with using Max or written languages like Python. This would mean packaging everything into a form that would be accessible with a clear GUI. It would also be necessary to render some of the processes more user-friendly – for example, in the current state, the user must save the data for each step of the process themselves. This means that if they wish to use the batch processing functionalities, they must keep a coherent naming convention of files with specific suffixes. This is the kind of task that could be made automatic.

My intention is to create a *Computational Musicology Toolkit* built in Max, Python and HTML/JavaScript, and geared towards integration with the IRiMaS software. Data could be passed around each environment, allowing the user to operate wherever they feel comfortable. A full HTML/JavaScript implementation would also allow for the tools to be supported online which is another aspect I find important.

Taking the Max paradigm as an example, I would build the project in a way that allows the user to take each step separately and incorporate steps into other projects, but also with a set of user-friendly standalone patches that can be used out of the box. The toolkit would notably include: network visualisation and navigation tools; the slicing, description, dimensionality reduction tools; the traditional segmentation tool used in Chapter 2; the patch tour tool presented in Chapter 2; the perceptions tool discussed above; tools for gestural recording and making the configuration of alternative performances easier; a project that would allow better understandings of things like descriptors and FFT; Max patch to network visualisations that could be linked to Max patches; dynamic network visualisation linked to video; bespoke network representation and navigation tools; and more.

This research has allowed me to cover the groundwork on how the internal structure of this environment would work and allowed me to build the prototypes for many of these ideas. There would also be scope for bringing the tools to fields further than musicology – I discussed how I have used some of these tools for gestural analysis in video games. Lots of the slicing tools could be useful for cinematographic and dance studies. Finally, the dynamic network visualisation tools could be taken-up in the field of social studies – I intend to propose a tool that, while certainly not being as complete as Cytoscape, would notably offer ways of reconfiguring a network over time, a feature which is notably absent from the program.

Paths for this thesis.

With a development of these tools, it would be much easier to propose the methods I have employed to a wider set of researchers. However, there are still methods that can be drawn from this thesis in the state it is today. I will discuss further paths for this thesis and its methodology, first by inspecting how the analysis could be taken further; then look into other fields of musicology and the social sciences.

With more time, I would have liked to have been able to perform an analysis of the same kind of detail around the work performed by the second cohort of musickers. There is no doubt that with each case study, the methodology would have developed further. Notably, instances of musickers working with other performers or of pre-recording their pieces. Some of the networks were quite different in nature to those I have presented: for example, Harker and Tutschku's pieces were more traditional, with an acoustic instrument being augmented with live electronics, reading from a score. It would have been interesting to try and configure alternative performances of

these pieces. I have archived all the sources, which means that these elements will be available for further analysis.

It would also be interesting to return to earlier case studies, such as Constanzo's work, and perform a second pass of analysis informed by the methodology in the state it is now. I would also have liked to explore further the idea of dynamic sound plots and networks, and the idea of aesthetic, resistant interfaces for analysis.

There are also sources that have been exploited less than others: for my research I decided to focus on the instrumental setup, the code that was written and the audio of performances and other elements. However, sources such as the forum that accompanied the project where musickers shared experiments, reported issues and helped each other would make a very interesting source not only for a musicological analysis, but also in a sociological or anthropological context.

My approach was also primarily focused on the networks in the state that they were found at the moment of their execution in performance. I did touch on the networks as they developed in Chapter 3, however a much more in-depth analysis of these elements could also be interesting. The same goes for the various experiments and tests that the musickers did. With such a vast amount of source material, I had to make choices to make this thesis coherent and adequately detailed.

It would be interesting to take the results of my research and put them to the musickers in question. This would not be with the idea of seeking validation, but to examine the effects that this research would have by inserting itself back into the networks it took as its subjects. There is no doubt that, like every node in the network, this would impart a degree of agency: it would be interesting to see how the musickers, as points of interface, would translate it. We can imagine a second wave of performances, where the musickers could return to the networks – it would be fascinating to see what would change.

Looking further afield, I believe that interrogation around the concept of time in the network approach is a useful contribution. This is a notion that could be explored in much more detail from a theoretical point of view, however I believe that I have demonstrated that when it is taken as a fundamental principle for analysis, it provides a framework that leads to interesting questions. The Record model, and notably the idea of engagements of measurement and manipulation is also an interesting perspective that could be taken further than this musical context. When exploring this concept, we saw that it is these processes that give the Record its musical quality – we

could envisage a generalisation of the idea by replacing the Record with the Artefact and proposing that it is fundamentally structured by these activities of measurement and manipulation.

It would also be interesting to see to what extent the idea of resistant, aesthetic interfaces for analysis would survive in a context other than the social sciences. For example, if this idea was put to a physicist or a biologist, would they be resistant to the idea? It is difficult to imagine a case where an aesthetic point of interface could coexist in a functional mindset, but maybe there are cases in the hard sciences where an aesthetic approach could be useful.

The idea wouldn't be to generalise this across all scientific practice, but maybe there would be space for allowing embodied knowledge to participate in hard scientific research. Indeed, the goal of science is essentially to understand what is not yet understood – maybe an aesthetic perspective on a subject in the context of the hard sciences could allow the human beings performing the analysis to unlock new perspectives on their subject.

Computational musicology without computers.

Finally, I will open a discussion around the nature of computational musicology and frame it around the physicalities with which it appears to be intertwined. Is a computational approach possible without computers? I breach this question for two reasons: first to try and gain access to the fundamental principles that drive the computational approach, to understand the nature of this perspective at an abstract level; and also to frame this research in the context of a world where heavy use of computers is something that can be put into question, both on an ecological level; and in regard to the analyst's well-being with a progressive shift of consciousness out of the physical and into the digital domain.

What does the computational approach allow us? Computers first allow us to deal with extremely large amounts of data. It is conceivable that the quantities of data could be measured and manipulated by hand, but the tasks performed on a dataset in the digital domain in a matter of seconds could take months or years to calculate by hand. Similarly, computers allow us to conceive of sound at a very high resolution at a typical sample rate of 44100 frames per second. Again, it is possible to conceive of music atomically, but we can get nowhere near this kind of resolution without computers.

So, computers allow us to engage with orders of data which are otherwise unattainable or impossible to engage with. They essentially give us access to different types of information: information that is extremely high in resolution, and information that is vast. These collections are measurements of Records – they are another way of conceiving of the linear flux of music in an abstract way. When a Record is conceived of or translated by a point of interface, there is always a loss of resolution. This is the case with computers; however, the loss of resolution appears to be lower than when the Record is measured by a human point of interface. The computer therefore gives us access to expanded capacities of measurement and manipulation.

In practice, what do we do with these capacities? All musicological analysis is geared towards passing back through the human point of interface. It is a form of measurement that seeks to impart agency regarding the Record towards the reader that will allow them to measure the Record in different ways. There will still be loss of resolution compared to immediate musical experience, but perhaps less so, or perhaps a loss of resolution in different forms. This manifests itself through interfaces such as waveforms, spectrograms and artefacts like network visualisations and sound plots. These are all points of interface that emerge from a high-resolution measurement and manipulation of a Record, through which we engage with the original Record in an indirect manner. Computational musicology can be seen as the production of Records which allow us to conceive of the Record in different ways.

This process is common to all musicological analysis – even the idea of creating artefacts from processes that work at higher resolutions is something that is found in a traditional context. Working from a score allows us to slow the music down and access a more atomic perspective than that of in-the-moment perception. Nevertheless, there are aspects in musical experience and perception which are unable to be transcribed in a written, analytical form. The loss of resolution that occurs through analysis of any form is not only across the energetic, temporal and indeed organisational dimensions, but also in the affective dimension of music which arguably escapes any kind of written form.

I have proposed a way in which this notion could be addressed in the computational approach with aesthetic, resistant forms of interface. The same concept could be envisaged in traditional, analogue forms of analysis. This puts into question much less the nature of the analytical artefact than the nature of our engagement with it.

To summarise, what really seems to differentiate the computational approach are the orders of magnitude of data which are being treated. We have seen this kind of approach growing in analogue musicology – a desire to look beyond the score and consider other objects for analysis such as the performance, organology, sketches and a composer's notes. It seems that the forms of measurement and manipulation offered by computers are inherently intertwined with its physicalities; however, we can try to expand this idea of an enlarged set of sources, and indeed the *creation* of sources as an inherent part of analysis.

The analyst can take a much more active part in creation and play a role in the configuration of networks to examine their nature. This is a paradigm that does not require the presence of computers, it is an approach that can be applied to any kind of analysis and that also serves the turns that have been observed in contemporary musicology. It is an approach that can be extremely fruitful: it brings us into proximity with the network and allows us to gain an embodied understanding of its various elements and the relationships between them. It is an approach that greatly expands the possibilities of analytical activity and proposes a promising avenue for the future.

BIBLIOGRAPHY

Research

Books

- ADORNO, Theodor, LEPPERT, Richard, and GILLESPIE, Susan. *Essays on Music*. Berkeley, California, USA: University of California Press, 2002.
- BENJAMIN, Walter. *The Work of Art in the Age of Mechanical Reproduction*. Translated by James Amery Underwood. London, UK: Penguin, 2008.
- CLINE, David. *The Graph Music of Morton Feldman*. Cambridge, UK: Cambridge University Press: 2016.
- FRITH, Simon. *Performing Rites: On the Value of Popular Music*. Oxford, UK: Oxford University Press, 1996.
- FULLER, Matthew, MALINA, Roger and CUBITT, Sean. *Software Studies: A Lexicon*. Cambridge, Massachusetts, USA: MIT Press, 2008.
- GALLOWAY, Alexander. *The Interface Effect*. Cambridge, UK; Polity, 2012.
- HANNINEN, Dora. *A Theory of Music Analysis: On Segmentation and Associative Organization*. Rochester, New York, USA: University of Rochester Press, 2012.
- INGOLD, Tim. *Being Alive: Essays on Movement, Knowledge and Description*. UK: London: Taylor & Francis Group, 2011.
- INGOLD, Tim. *Making: Anthropology, Archaeology, Art and Architecture*. Abingdon, UK: Routledge, 2013.
- LATOUR, Bruno. *The Pasteurization of France*. Cambridge, Massachusetts, USA: Harvard University Press, 1988.
- MAGNUSSON, Thor. *Sonic Writing: Technologies of Material, Symbolic, and Signal Inscriptions*. New York, New York, USA: Bloomsbury Publishing USA, 2019.
- SCHAEFFER, Pierre. *Treatise on Musical Objects: An Essay across Disciplines*. Translated by Christine North and John Dack. Oakland, California, USA: University of California Press, 2017.
- SERRES, Michel and SCHEHR, Lawrence. *The Parasite*. Minneapolis, Minnesota, USA: University of Minnesota Press, 2007.
- SLOTERDIJK, Peter. *Spheres*. Los Angeles, California, USA: Semiotext(e), 2004.

SMALL, Christopher. *Musicking: The Meanings of Performing and Listening*. Hanover, New Hampshire, USA: University Press of New England, 1998.

SUCHMAN, Lucille Alice. *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge, New York, USA: Cambridge University Press, 2007.

TARUSKIN, Richard. *The Oxford History of Western Music*. Oxford, New York, USA: Oxford University Press, 2005.

TOMLINSON, Gary. *A Million Years of Music: The Emergence of Human Modernity*. New York, New York, USA: Zone Books, 2015.

Chapters

COOK, Nicholas. 'Analysing Performance and Performing Analysis', in *Rethinking Music*. Oxford, UK: Oxford University Press, 1999.

HARDJOWIROGO, Sarah. 'Instrumentality. On the Construction of Instrumental Identity'. In *Musical Instruments in the 21st Century: Identities, Configurations, Practices*. Singapore: Springer Singapore, 2017: 9-24. https://doi.org/10.1007/978-981-10-2951-6_2.

INGOLD, Tim. 'When ANT Meets SPIDER: Social Theory for Arthropods'. In *Material Agency*. Boston, Massachusetts, USA: Springer US, 2008: 209-15. https://doi.org/10.1007/978-0-387-74711-8_11.

STIEGLER, Bernard. 'Memory'. In *Critical Terms for Media Studies*. Chicago, Illinois, USA: The University of Chicago Press, 2010.

Articles

ABBATE, Carolyn. 'Music - Drastic or Gnostic?' *Critical Inquiry*, vol. 30, no. 3 (2004): 505–36. <https://doi.org/10.1086/421160>.

AGAWU, Kofi. 'How We Got out of Analysis, and How to Get Back in Again'. *Music Analysis*, vol. 23, no. 2/3 (2004): 267–86.

- BERNSTEIN, Zachary. 'Review of Dora Hanninen "A Theory of Music Analysis: On Segmentation and Associative Organization"'. *Music Theory Online*, vol. 19, no. 4, (2013).
http://www.academia.edu/5487212/Review_of_Dora_Hanninen_A_Theory_of_Music_Analysis_On_Segmentation_and_Associative_Organization_.
- BORN, Georgina. 'For a Relational Musicology: Music and Interdisciplinarity, Beyond the Practice Turn: The 2007 Dent Medal Address'. *Journal of the Royal Musical Association*, vol. 135, no. 2 (2010): 205–43.
- CALLON, Michel. 'Éléments pour une sociologie de la traduction: la domestication des coquilles Saint-Jacques et des marins-pêcheurs dans la baie de Saint-Brieuc'. *L'Année sociologique (1940/1948-)*, vol. 36 (1986): 169–208.
- CHEN, Min, and FLORIDI, Luciano. 'An Analysis of Information Visualisation'. *Synthese*, vol. 16 (2013): 3421–3438.
https://www.researchgate.net/publication/257666827_An_Analysis_of_Information_Visualisation.
- COOK, Nicholas. 'Between Process and Product: Music and/as Performance'. *Music Theory Online*, vol. 7, no. 2 (2001).
<https://mtosmt.org/issues/mto.01.7.2/mto.01.7.2.cook.html>.
- DONKER, Silvia. 'Networking Data. A Network Analysis of Spotify's Socio- Technical Related Artist Network'. *International Journal of Music Business Research*, vol. 8 (2019): 67. <https://research.rug.nl/en/publications/networking-data-a-network-analysis-of-spotifys-socio-technical-re>.
- DORRITY, Michael, SAUNDERS, Lauren, QUEITSCH, Christine, FIELDS, Stanley, and TRAPNELL, Cole. 'Dimensionality Reduction by UMAP to Visualize Physical and Genetic Interactions'. *Nature Communications*, vol. 11, no. 1 (2020).
<https://doi.org/10.1038/s41467-020-15351-4>.
- DUFEU, Frédéric. 'Comment développer des outils généraux pour l'étude des instruments de musique numériques?' *Revue Francophone d'Informatique et Musique*, vol. 3 (2013). <https://revues.mshparisnord.fr:443/rfim/index.php?id=255>.
- FRUCHTERMAN, Thomas and REINGOLD, Edward. 'Graph Drawing by Force-Directed Placement'. *Software: Practice and Experience* vol. 21, no. 11 (1991): 1129–64.
<https://doi.org/10.1002/spe.4380211102>.

- HANNINEN, Dora. 'Orientations, Criteria, Segments: A General Theory of Segmentation for Music Analysis'. *Journal of Music Theory*, vol. 45, no. 2 (2001): 345–433. <https://doi.org/10.2307/3653443>.
- HASTY, Christopher. 'Segmentation and Process in Post-Tonal Music'. *Music Theory Spectrum*, vol. 3 (1981): 54–73. <https://doi.org/10.2307/746134>.
- HOGG, Bennett. 'When Violins Were Trees: Resistance, Memory, and Performance in the Preparatory Experiments for Landscape Quartet, a Contemporary Environmental Sound Art Project'. *Contemporary Music Review*, vol. 32, no. 2–03 (2013): 249–73. <https://doi.org/10.1080/07494467.2013.775811>.
- LASKE, Otto. 'Toward a Musicology for the Twentieth Century'. *Perspectives of New Music*, vol. 15, no. 2 (1977): 220–25. <https://doi.org/10.2307/832820>.
- LEFKOWITZ, David, and TAAVOLA, Kristin. 'Segmentation in Music: Generalizing a Piece-Sensitive Approach'. *Journal of Music Theory*, vol. 44, no. 1 (2000): 171–229. <https://doi.org/10.2307/3090673>.
- LEWIN, David. 'Music Theory, Phenomenology, and Modes of Perception'. *Music Perception: An Interdisciplinary Journal*, vol. 3, no. 4 (1986): 327–92. <https://doi.org/10.2307/40285344>.
- VAN DER MAATEN, Laurens, and HINTON, Geoffrey. 'Visualizing Data Using T-SNE'. *Journal of Machine Learning Research*, vol. 9 (2008): 2579–2605.
- MAUS, Fred Everett. 'Response to Lawrence Rosenwald'. *The Journal of Musicology*, vol. 11, no. 1 (1993): 66–72. <https://doi.org/10.2307/764153>.
- MORIN, Marie-Eve. 'Cohabiting in the Globalised World: Peter Sloterdijk's Global Foams and Bruno Latour's Cosmopolitics'. *Environment and Planning D: Society and Space*, vol. 27, no. 1 (2009): 58–72. <https://doi.org/10.1068/d4908>.
- PEETERS, Geoffroy and DERUTY, Emmanuel. 'Sound Indexing Using Morphological Description'. *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3 (2010): 675–87. <https://doi.org/10.1109/TASL.2009.2038809>.
- PIEKUT, Benjamin. 'Actor-Networks in Music History: Clarifications and Critiques'. *Twentieth-Century Music*, vol. 11, no. 2 (2014): 191–215. <https://doi.org/10.1017/S147857221400005X>.

- SHANNON, Paul, MARKIEL, Andrew, OZIER, Owen, BALIGA, Nitin, WANG, Jonathan, RAMAGE, Daniel, AMIN, Nada, SCHWIKOWSKI, Benno, and IDEKER, Trey. 'Cytoscape: A Software Environment for Integrated Models of Biomolecular Interaction Networks'. *Genome Research*, vol. 13, no. 11 (2003): 2498–2504. <https://doi.org/10.1101/gr.1239303>.
- SMALLEY, Denis. 'Spectromorphology: Explaining Sound-Shapes'. *Organised Sound*, vol. 2, no. 2 (1997): 107–26. <https://doi.org/10.1017/S1355771897009059>.
- TENNEY, James, and POLANSKY, Larry. 'Temporal Gestalt Perception in Music'. *Journal of Music Theory*, vol. 24, no. 2 (1980): 205–41. <https://doi.org/10.2307/843503>.

Conference Proceedings

- BAUDOIN, Olivier. 'La Faktura, "Outil conceptuel d'analyse" – Illustration avec *Stria*, de John Chowning'. *Journées d'Informatique Musicale*. Grenoble, France: 2009. <https://hal.archives-ouvertes.fr/hal-03132802>.
- BUFFA, Michel, CABRIO, Elena, FELL, Michael, GANDON, Fabien, GIBOIN, Alain, HENNEQUIN, Romain, MICHEL, Franck, et al. 'The WASABI Dataset: Cultural, Lyrics and Audio Analysis Metadata About 2 Million Popular Commercially Released Songs'. *Extended Semantic Web Conference*. Heraklion, Greece: 2021. <https://hal.archives-ouvertes.fr/hal-03282619>.
- BRYAN, Nicholas, WANG, Ge. 'Musical Influence Network Analysis and Rank of Sample-Based Music.'. *12th International Society for Music Information Retrieval Conference*. Miami, Florida, USA: 2011. https://www.researchgate.net/publication/220723768_Musical_Influence_Network_Analysis_and_Rank_of_Sample-Based_Music.
- COUPRIE, Pierre. 'IAAnalyse 5 : Le Développement d'un Logiciel Pour La Musicologie Numérique'. *Journées d'informatique Musicale*. Bayonne, France: 2019. <https://hal.archives-ouvertes.fr/hal-02133538>.
- DUFEU, Frédéric, TAKAHASHI, Keitaro, ROEBEL, Axel, and CLARKE, Michael. 'The IRiMaS Software: Integrating Interactive Listening and Play into Musicological Research'. *International Computer Music Conference Proceedings*. New York, New York, USA: 2019. <https://hal.archives-ouvertes.fr/hal-02465938>.

- GREEN, Owen, TREMBLAY, Pierre Alexandre and ROMA, Gerard. 'Interdisciplinary Research as Musical Experimentation', *Proceedings of the Electroacoustic Music Studies Network Conference*. Florence, Italy: 2018.
- ROMA, Gerard, GREEN, Owen and TREMBLAY, Pierre Alexandre. 'Adaptive Mapping of Sound Collections for Data-Driven Musical Interfaces: The International Conference on New Interfaces for Musical Expression'. *Proceedings of the International Conference on New Interfaces for Musical Expression*. Porto Alegre, Brazil: 2019: 313–18.
- SCHWARZ, Diemo, BELLER, Grégory, VERBRUGGHE, Bruno, and BRITTON, Sam. 'Real-Time Corpus-Based Concatenative Synthesis with CataRT'. *9th International Conference on Digital Audio Effects (DAFx)*. Montreal, Canada: 2006: 279–82. <https://hal.archives-ouvertes.fr/hal-01161358>.
- TREMBLAY, Pierre Alexandre, GREEN, Owen, and ROMA, Gerard. 'Digging It: Programmic Data Mining as Musicking', *International Computer Music Conference Proceedings*. Santiago, Chile: 2021: 295-300.
- TREMBLAY, Pierre Alexandre, GREEN, Owen, ROMA, Gerard, and HARKER, Alex. 'From Collections to Corpora: Exploring Sounds through Fluid Decomposition', *International Computer Music Conference*. New York, New York, USA: 2019: 223-228.

Datasets

- HART, Jacob. 'Performance Cartography, Performance Cartology: Multimedia Appendices'. University of Huddersfield Research Portal, 2021. <https://doi.org/10.34696/p0y3-7f65>.

Sources

Theses

- CONSTANZO, Rodrigo. 'Composition, Performance, Improvisation, and Making Things, Sitting in a Tree: Me-Me-Me-Me-Me-Me-Me.' University of Huddersfield, 2016. <http://www.rodrigoconstanzo.com/thesis/>.

- ELDRIDGE, Alice. 'Collaborating with the Behaving Machine: Simple Adaptive Dynamical Systems for Generative and Interactive Music'. University of Sussex, 2007.
- HARKER, Alexander. 'Portfolio of Compositions'. University of York, 2011.
- HAYES, Lauren. 'Audio-Haptic Relationships as Compositional and Performance Strategies'. University of Edinburgh, 2014. <http://hdl.handle.net/1842/9481>.
- PASQUET, Olivier. 'Automatic versus Automatic, Materialized Fiction as a Confrontational Compositional Process'. University of Huddersfield, 2018.
- PARKER, Martin. 'Spectral Tourist: Joystick Operated Sound Production Software for Macintosh Computers' University of Edinburgh, 2004. <https://era.ed.ac.uk/handle/1842/33458>.
- PLUTA, Samuel Francis. 'Laptop Improvisation in a Multi-Dimensional Space'. Columbia University, 2012. <https://doi.org/10.7916/D8S188KG>.
- TUTSCHKU, Hans. 'Portfolio of Compositions'. University of Birmingham, 2003.

Writings

- AGOSTINI, Andrea, and GHISI, Daniele. 'A Max Library for Musical Notation and Computer-Aided Composition'. *Computer Music Journal*, vol. 39, no. 2 (2015): 11–27. https://doi.org/10.1162/COMJ_a_00296.
- BUTTERFIELD, Andrew, NGONDI, Gerard and KERR, Anne. 'Lempel–Ziv–Welch Compaction'. In *A Dictionary of Computer Science*. Oxford, UK: Oxford University Press, 2016. <http://www.oxfordreference.com/view/10.1093/acref/9780199688975.001.0001/acref-9780199688975-e-2881>.
- ELDRIDGE, Alice, GUYOT, Patrice, MOSCOSO, Paola, JOHNSTON, Alison, EYRE-WALKER, Ying, and PECK, Mika. 'Sounding out Ecoacoustic Metrics: Avian Species Richness Is Predicted by Acoustic Indices in Temperate but Not Tropical Habitats'. *Ecological Indicators*, vol. 95 (2018): 939–52. <https://doi.org/10.1016/j.ecolind.2018.06.012>.

- ELDRIDGE, Alice, and KIEFER, Chris. 'The Self-Resonating Feedback Cello: Interfacing Gestural and Generative Processes in Improvised Performance'. *Proceedings of New Interfaces for Music Expression*. Copenhagen, Denmark: 2017. <http://homes.create.aau.dk/dano/nime17/papers/0005/index.html>.
- HARKER, Alexander. 'FrameLib: Audio DSP Using Frames of Arbitrary Length and Timing'. *43rd International Computer Music Conference and the 6th International Electronic Music Week*. Shanghai, China: 2017: 271–78.
- HARKER, Alexander and TREMBLAY, Pierre Alexandre. 'The HISSTools Impulse Response Toolbox: Convolution for the Masses'. *Proceedings of the International Computer Music Conference*. Ljubljana, Slovenia: 2012: 148–55. https://github.com/HISSTools/HISSTools_Impulse_Response_Toolbox.
- HAYES, Lauren. 'Beyond Skill Acquisition: Improvisation, Interdisciplinarity, and Enactive Music Cognition'. *Contemporary Music Review* (2019): 1–17. <https://doi.org/10.1080/07494467.2019.1684059>.
- VEALE, Peter. *The Techniques of Oboe Playing: A Compendium with Additional Remarks on the Oboe d'amore and the Cor Anglais*. Kassel, New York, USA: Bärenreiter, 1994.

Blogs

- CONSTANZO, Rodrigo. 'Kaizo Snare', blog post (2020). <https://rodrigoconstanzo.com/2020/08/kaizo-snare/>.
- PASQUET, Olivier. 'Jtol', blog post (2012). <https://www.opasquet.fr/jtol/>.
- PASQUET, Olivier. 'Jtol_lzw', blog post (2019). https://www.opasquet.fr/jtol_lzw/.

CITED WORKS

Recordings

- DEVINE, Richard. *Sort\Love*. Timesig, 2018.
- ENO, Brian. *Discreet Music*. Obscure, 1975.
- HAYES, Lauren Sarah. *Manipulation*. Pan y rosas discos, 2016.
- . *Embrace*. Superpang, 2021.
- LEAFCUTTER JOHN. *The Housebound Spirit*. Planet-Mu Recordings, 2003.
- . *The Forest and the Sea*. Staubgold, 2006.
- . *Tunis*. Tsuku Boshi, 2010.
- POLAR BEAR. *Held on the Tips of Fingers*. Babel Label, 2005.

Pieces

- HARKER, Alexander. *Fluence*, clarinet and electronics, 2010.
- . *Fractures*, fixed media, 2011.
- . *Drift Shadow*, score and electronics. 2021.
- HARVEY, Jonathan. *Mortuos Plango Vivos Voco*, fixed media. 1980.
- MOZART, Wolfgang Amadeus. *Musikalisches Würfelspiel*, KV 516f, score. 1787.
- PACHELBEL, Johann. *Canon in D Major*, PXC 37, T.337, PC 358, score. 1680.
- SAARIAHO, Kaija. *Du cristal*, ensemble. 1989-1990.
- TUTSCHKU, Hans. *Sparks*, score and electronics, 2021.

Performances

- CONSTANZO, Rodrigo. *Rhythm Wish*, performed by CONSTANZO, Rodrigo. Borealis Festival, Bergen, Norway, 2016.
- . *Kaizo Snare*, performed by CONSTANZO, Rodrigo. HCMF, Huddersfield, UK, 2019.

DEVINE, Richard. *Constructors*, performed by DEVINE, Richard. Dialogues Festival, Edinburgh, UK, 2021.

ELDRIDGE, Alice and KIEFER, Chris. *Feedback Cello*, performed by ELDRIDGE, Alice and KIEFER, Chris. Dialogues Festival, Edinburgh, UK, 2021.

HARKER, Alexander. *Drift Shadow*, performed by DELL, Niamh. Dialogues Festival, Edinburgh, UK, 2021.

HAYES, Lauren Sarah. *Moon via Spirit*, performed by HAYES, Lauren Sarah. HCMF, Huddersfield, UK, 2019.

LEAFCUTTER JOHN. *Nightless Night*, performed by LEAFCUTTER JOHN. XOYO, London, UK, 2014.

———. *Line Crossing*, performed by LEAFCUTTER JOHN. HCMF, Huddersfield, UK, 2019.

PASQUET, Olivier. *hr 8799*, performed by PASQUET, Olivier. National Taiwan Museum of Fine Arts, Taipei, Taiwan, 2014.

———. *Dual Cornographs*, performed by PASQUET, Olivier. Silos City, Buffalo, NY, USA, 2018.

———. *Herbig-Haro*, performed by PASQUET, Olivier. HCMF, Huddersfield, UK, 2019.

PLUTA, Sam and EVANS, Peter. *Neural Duo I and II*, performed by PLUTA, Sam and EVANS, Peter. Dialogues Festival, Edinburgh, UK, 2021.

TUTSCHKU, Hans. *Sparks*, performed by KNOOP, Mark. Dialogues Festival, Edinburgh, UK, 2021.

Software

AGOSTINI, Andrea, and GHISI, Daniele. *Bach*, (2015). <https://www.bachproject.net/>.

BRADBURY, James. *python-flucoma*, (2020). <https://github.com/jamesb93/python-flucoma>.

———. James. *ReaCoMa*, (2020). <https://www.reacoma.xyz/>.

CONSTANZO, Rodrigo. *Cut Glove*, (2015). <http://rodrigoconstanzo.com/cut-glove/>.

COURNAPEAU, David. *Scikit-learn*, (2010). <https://scikit-learn.org/stable/>.

CYCLING '74. *Max 8*, (2018). <https://cycling74.com/products/max/>.

DUFEU, Frédéric, TAKAHASHI, Keitaro, ROEBEL, Axel, CLARKE, Michael. *TiAALS*, (2021). <https://research.hud.ac.uk/institutes-centres/irimas/tiaals/>.

FRANKEL, Justin. *Reaper*, (2006). <https://www.reaper.fm/>.

HARKER, Alexander. *Alex Harker Externals*, (2011). <https://www.alexanderjharker.co.uk/Software.html>.

———. *Framelib*, (2017). <https://github.com/AlexHarker/FrameLib>.

HARKER, Alexander and TREMBLAY, Pierre Alexandre. *HISSTools Impulse Response Toolbox*, (2012). https://github.com/HISSTools/HISSTools_Impulse_Response_Toolbox.

IRCAM. *Madalys 3.6*, (2021). <https://forum.ircam.fr/projects/detail/modalys/>.

———. *Spat 5*, (2021). <https://forum.ircam.fr/projects/detail/spat/>.

———. *SuperVP*, (2021). <https://forum.ircam.fr/projects/detail/superop-for-max/>.

LIINE. *Lemur 5*, (2021). <https://liine.net/en/products/lemur/>.

MCCARTNEY, James. *SuperCollider 3*, (2020). <https://supercollider.github.io/>.

MCNEEL, Robert et al. *Rhinoceros 3D*, (2012). <https://www.rhino3d.com/>.

MICROSOFT. *Visual Studio Code*, (2015). <https://code.visualstudio.com/>.

MODARTT. *Pianoteq*, (2006). <https://www.modartt.com/pianoteq>.

PASQUET, Olivier. *Jtol*, (2013). <https://www.opasquet.fr/jtol/>.

PUCKETTE, Miller. *Pure Data*, (1996). <http://puredata.info/>.

SCHWARZ, Diemo. *CatRT*, (2006). <http://imtr.ircam.fr/imtr/CataRT>.

SHANNON, Paul, MARKIEL, Andrew, OZIER, Owen, BALIGA, Nitin, WANG, Jonathan, RAMAGE, Daniel, AMIN, Nada, SCHWIKOWSKI, Benno, and IDEKER, Trey. *Cytoscape*, (2003). <https://cytoscape.org/>.

SONIC ACADEMY. *Kick 2*, (2016). <https://www.sonicacademy.com/products/kick-2>.

TREMBLAY, Pierre Alexandre, GREEN, Owen, ROMA, Gerard, and HARKER, Alex. *Fluid Corpus Manipulation Toolbox*, (2020). <https://www.flucoma.org/download/>.

UNIVERSITE DE TECHNOLOGIE DE COMPIEGNE. *Gephi*, (2008). <https://gephi.org/>.

APPENDIX

1. FluCoMa objects overview.

Non-real-time objects have the *Buf* prefix.

1.1. Slicing objects.

- *AmpGate* and *BufAmpGate*: amplitude-based slicing using an envelope threshold with lookahead and lookback.
- *AmpSlice* and *BufAmpSlice*: amplitude-based slicing using detrending.
- *OnsetSlice* and *BufOnsetSlice*: onset detection in the spectral domain.
- *NoveltySlice* and *BufNoveltySlice*: slicing based on novelty detection according to various user-chosen descriptors: the magnitude of the spectrum, MFCC content, pitch description and confidence and true peak and loudness descriptors.
- *TransientSlice* and *BufTransientSlice*: slicing based on transient modelling.

1.2. Layer-finding objects.

- *Hpss* and *BufHpss*: decomposing the sound into harmonic and percussive layers.
- *Sines* and *BufSines*: decomposing the sound into a sinusoidal layer and everything else.
- *Transients* and *BufTransients*: decomposing the sound into a transient layer and everything else.

1.3. Object-finding objects.

- *BufNmf*: a machine-learning algorithm called NMF which essentially returns a user-defined number of components in the data it is fed. This object runs the algorithm in non-real-time and produces a set of bases and activations: bases are the spectral templates of the various components and activations are the gain envelopes for each base.
- *NmfMatch*: process incoming real-time audio against a pre-processed set of NMF activations and bases.
- *NmfFilter*: similar to *NmfMatch*. This returns a resynthesis of the incoming audio based on pre-processed NMF activations and bases.

1.4. *Audio description objects.*

- *Loudness* and *BufLoudness*: amplitude and true peak audio descriptors.
- *Melbands* and *BufMelbands*: amplitude across a number of equally spread perceptual bands.
- *Mfcc* and *BufMfcc*: MFCC spectral content.
- *Pitch* and *BufPitch*: three different pitch descriptors and their confidence.
- *SpectralShape* and *BufSpectralShape*: a set of seven spectral shape descriptors: spectral centroid, spectral spread, normalized skewness, normalized kurtosis, rolloff, flatness and crest.
- *BufStats*: allows the user to gather a number of statistics on this data across the whole buffer, as well as any number of derivatives: mean, standard deviation, skewness, kurtosis, and user-defined minimum, median and maximum percentiles.

1.5. *Data storage objects.*

- *Dataset*: server-side storage of the data.
- *Labelset*: contains labels for each datapoint in a dataset.

1.6. *Data manipulation objects.*

- *DatasetQuery*: for performing operations like selecting datapoints above or below a certain threshold, adding and removing columns.
- *Normalize*: normalize entries of a dataset.
- *Standardize*: standardize a dataset, rescaling the datapoints using the mean and standard deviation in each dimension.
- *KdTree*: essentially performing nearest-neighbour searches within the dataset.
- *Pca*: principal components analysis dimensionality reduction of a dataset.
- *Mds*: multidimensional scaling dimensionality reduction.
- *Umap*: uniform manifold approximation dimensionality reduction.

1.7. *Classification and regression.*

- *KnnClassifier*: classification using KdTree nearest-neighbour.
- *KnnRegressor*: regression using KdTree nearest-neighbour.
- *MlpClassifier*: classification using multi-layer perceptron neural network.
- *MlpRegressor*: regression using multi-layer perceptron neural network.

1.8. *Audio transformation objects.*

- *BufCompose*: used to combine the contents of buffers.
- *AudioTransport* and *BufAudioTransport*: interpolation between the spectra of two sounds using Optimal Transport.
- *BufNmfCross*: produce a hybridisation of two buffers using NMF.
- *NmfMorph*: cross-synthesis using NMF and Optimal Transport.

2. Performance segmentations.

2.1. Olivier Pasquet – Herbig-Haro.

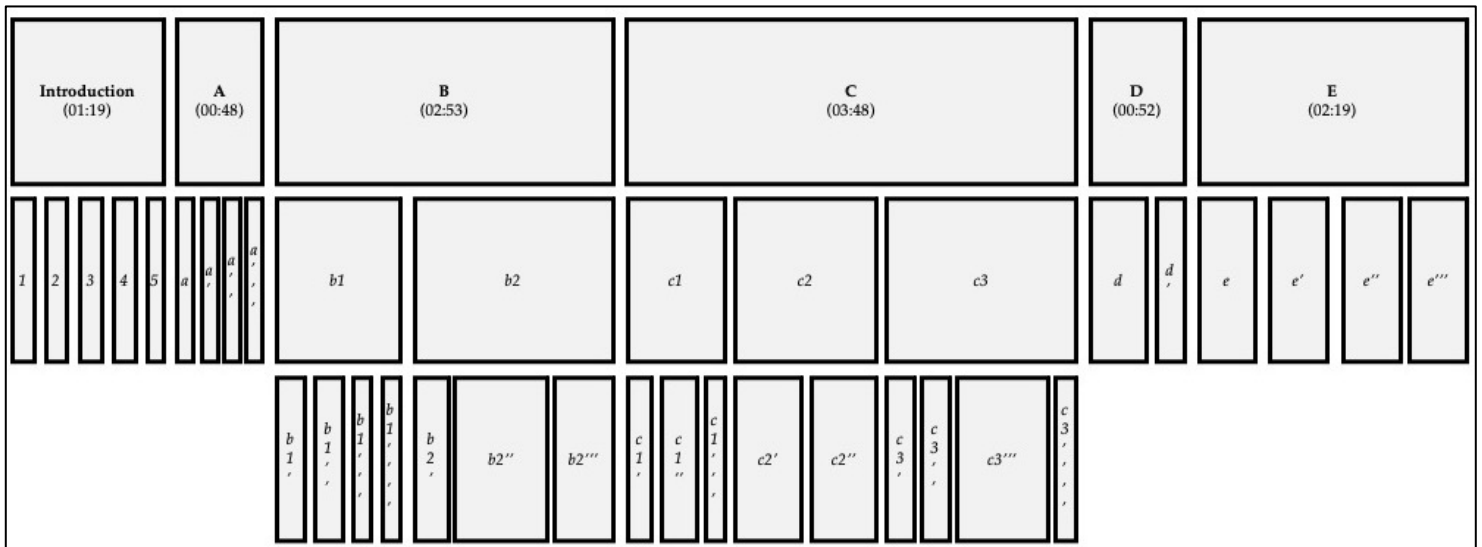


Figure 1: segmentation of Herbig-Haro. (Figure from Chapter 2).

The whole piece can be seen as a series of gestures. The first minute and 19 seconds are dedicated to somewhat of an introduction. From the darkness a series of clicks emerge that grow in intensity. Within the clicks, we remark several distinct layers: light clicks, heavy clicks, bursts of clicks almost resembling white noise, and what we shall refer to as pitched clicks which are introduced in Section A.

From then on, the sections present us with collections of gestures, usually marked by silence. Each section tends to focus on a specific layer. In Section A, Pasquet introduces the pitched clicks and a low sustained bass tone in a series of gestures. Here, the bass remains relatively stable around *c sharp*. With this section, the lights begin to flicker, and we notice that they seem to be being triggered by the pitched clicks. In Section B, the bass becomes more mobile, oscillating between different tones around the *c sharp*. Through this section, the bass and the clicks seem to hybridize and become one body – the modulations in the bass being felt within the clicks. There is also another layer in this section of some heavily processed vocals which come and punctuate the whole.

In Section C, the nature of the gestures changes somewhat. Before, the clicks were very chaotic, whereas here they are much more rhythmic. The gestures are being driven by the bass line, whose vocabulary does not only deploy itself across the pitched dimension, but across others like gain. Towards the end of the section, the dominant tone of *c sharp* seems to distance itself, with several gestures ending in suspension and

the tones being unclear and murky. We then pass through a short section of transition, until arriving at the final Section E where focus is now upon the pitched clicks. Here, they become much more chaotic, and progressively morph into longer events. This resides upon a sustained bass that now remains fixed on the *c sharp*. This eventually fades out, shortly followed by the clicks which in turn eventually dissipate.

2.2. Lauren Sarah Hayes – Moon via Spirit.

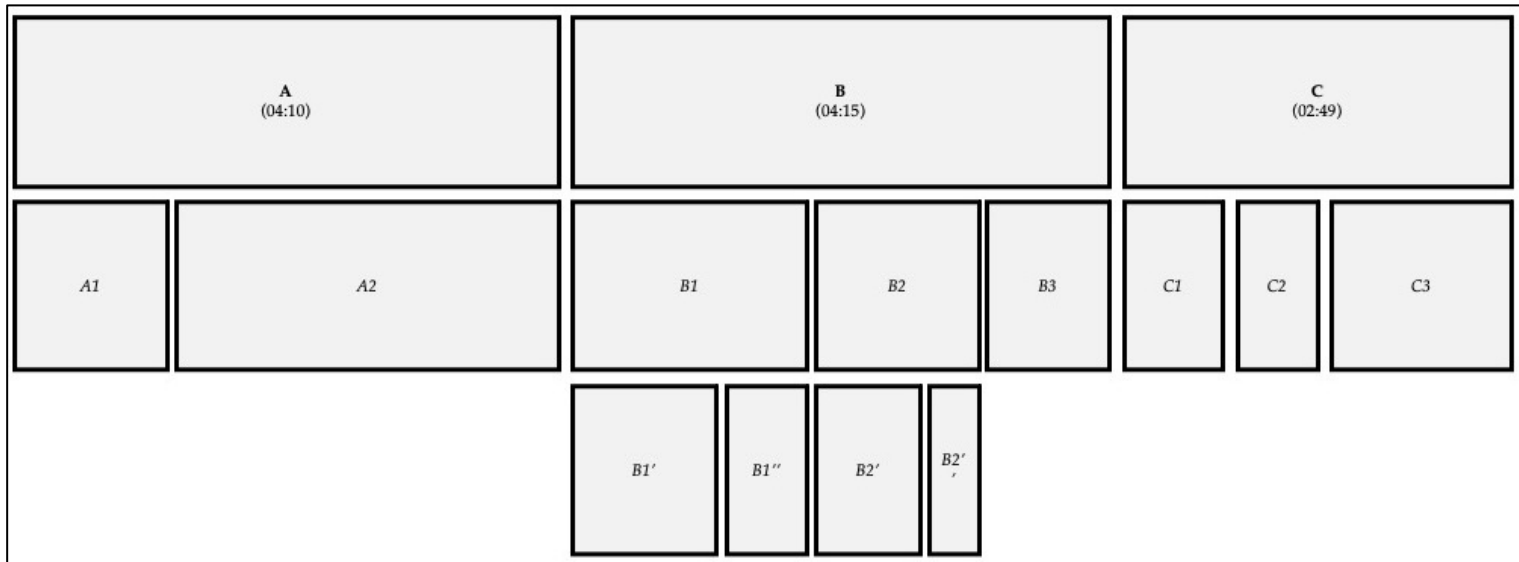


Figure 2: segmentation of Moon via Spirit. (Figure from Chapter 2).

Section 1 is seen as an introduction. It's characterized by sustained, low frequency drone synths – probably pulsar synthesis – which is later accompanied by vocals. This section also seems to serve as the recording of the vocal corpus that will be played back further into the piece. In A1, there are two layers: a low-pitched synthetic pad voice with a relatively pure frequency spectrum that is constantly playing and relatively stable around an *e flat*; and a higher-pitched, noisier synth that intervenes intermittently at first, and becomes more and more present. The frequencies of the two synths seem to be linked, however it's the timbral aspect of this second synth that seems to articulate its interventions: the nature of the distortion seems to be modulated by what sounds like an LFO that is constantly changing. I shall refer to these synths as the *pulsar synth*.

In A2, Hayes starts singing and the pulsar synth which was previously in two distinct layers is now perceived as one. Indeed, the vocals seem to replace the role of the previous noisier layer. The vocals are composed of two parts: Hayes' voice in real-time which seems to be distorted, and the immediate playback of grains at varying

speeds, evoking a spectral freeze effect. As the section progresses, the vocal interventions become longer in duration, and their playback more present.

Section B is the main body of the piece. Hayes plays with the playback of the previously recorded voice in various manners, plus other drum samples interspersed with pulsar synthesis. We can conceive of two large, symmetrical sections that commence with chaotic material (B1' and B2') from which repeated gestures emerge (B1'' and B2''); and one final section which seems to mimic gestures that have come before it and act as the resonance of these two preceding sections. Here, the pulsar synth becomes much more distorted and lower-pitched to a point where it is difficult to discern pitch. The vocals are now what sound like a sequenced play back of grains that were recorded beforehand, either live during Section A or before the performance. This glitchy vocal layer in the mid to high range seems to be constantly modulated and treated with various processes. The new sound source – the samples of synthetic percussions – are played sporadically. The speed and processes of these samples also seems to be being controlled dynamically. The gesture of chaotic attack and resonance is seen at various levels: across the whole of section B, but also more locally in sections B1 and B2.

Finally, in Section C the performance draws to its conclusion. Musical ideas from the beginning of the piece are taken up again, before leading into a final resonant section that fades out to silence. In C1, we begin with a passage that recalls the beginning of the piece, the synth layer is now on its own, but in a much higher frequency (a constant glissando of a semi-tone between *b flat* and *b*). Unlike Section A, there is only one synth layer that is perceptible, and the variation in timbre is more subtle and less abrupt. There are a few sporadic interventions from previous elements, a burst of voice here and there, a quick fade-in and -out of a modulating, dirtier synth line. In C2, again mirroring the beginning of the piece, this synth is then accompanied by a real-time, heavily processed vocal. The vocals seem to modulate and bend the synth line, and when the vocals are not present, the synth seems detuned and warped as if tarnished by the relatively aggressive cyber-sprechgesang vocal line. Finally, in Section C3, Hayes abandons the game controller, and controls the instrument using what appears to be only the Korg MIDI controller. We hear playback of grains from Section C2 which Hayes is constantly modulating with the Korg. The playback has an especially poppy, choppy nature, and gets more and more percussive towards the end of the piece. The bursts are less and less dynamic towards the end.

2.3. Rodrigo Constanzo – Kaizo Snare.

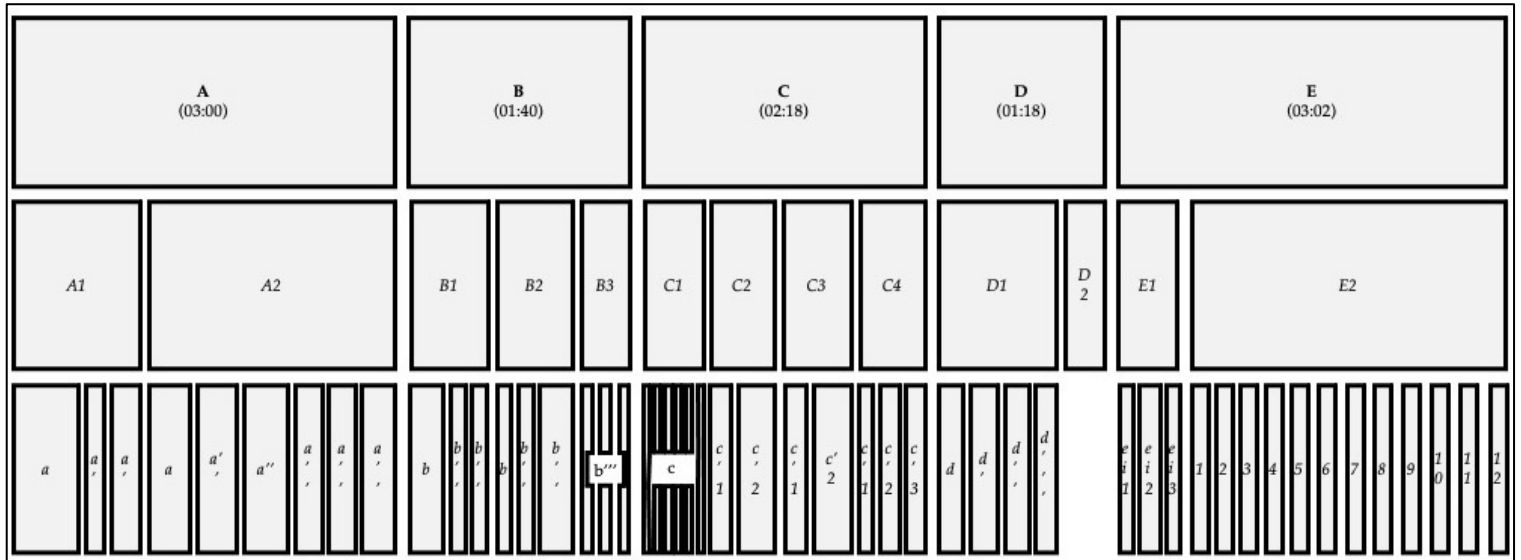


Figure 3: segmentation of Kaizo Snare. (Figure from Chapter 2).

There are four main streams that we can concentrate on: snare drum activity, sample-playback activity, kick drum activity and crotale solenoid activity. The piece can be divided into five large parts.

In Section A we have an introduction in two long gestures that start in similar fashions. Constanzo introduces us to the kind of gestural and musical vocabulary we can expect for the rest of the piece. In Section B, Constanzo again performs two long gestures, this time introducing a new tool: the comb. He ends this section with three long drags, before launching into the central Section C. This central section is more energetic, with shorter phrases and a variety of different playing techniques and types of sounds – we hear the kick drum being used more here. In Section D, we find somewhat of a suspension – indeed, two phrases of activity followed by long moments of suspension. Then, in the image of what precedes Section C, we drag into the final section. Section E feels like the culmination of the piece. We have a short introduction that leads into a suite of 12 gestures that are built upon a pitched low bass melody (procured from a mix of kick drum and low gong samples) which also marks a second solenoid moment.




Figure 4: the score generated for Line Crossing on the night of the FluCoMa performance. (Figure from Chapter 2).

2.4. John Burton – Line Crossing.

The score to Burton's piece is simple: a play head shows where we are in the piece, when there is a red block, the artist can play something, and their input is recorded. The blue block outlines show when part of this ever-growing recording is played back by the patch in some way. Helpfully, Burton implemented a system to save the contents of the *coll* objects that contain the score information. We can use this to reference each block, each has a precise onset time and length, and amplitude (from 0-100). Each block also has a shape type: 0, ramp up; 1, ramp down; 2, block. There are in fact two scores – the human score and the *machine* score.

Bearing in mind that each line on the score is the same length, we can remark just by looking at the score that there is an arch-shaped structure in the amount and frequency of events. Second, it seems that the visual construction of the score influences greatly Burton's activity during the performance – he changes instrument configuration at almost every line break. In this instrumental stream, we also notice a certain arch-shaped structure, starting and ending with cymbals, a central part with the guitar and cymbals, and transitions between the two with the recorders. From the score image, we can see that the frequency of the machine blocks follows this arch structure, and the quantity gets greater across the piece. It is difficult to discern the individual

streams as they often emerge simultaneously, however we can give a broad overview line-by-line to get an idea about what's happening:

- *Line 1:* here the machine blocks start with long, sustained notes, two very high-pitched ones, then a low followed by mid-range ones.
- *Line 2:* more sustained notes in the mid to low range, and the last machine block of this line is a higher-pitched sustained note similar to the beginning of the piece. The machine is giving lots of space for the human to play.
- *Line 3:* the line begins with a sustained bass note, over which we hear a distorted delay. Then we have another distorted delay, growing in intensity, accompanied by sustained notes of varying pitch. The line ends with a very ethereal sustain.
- *Line 4:* here, the machine blocks grow much more in intensity and give less space to the musician who seemingly must fight to pierce through the sound. The line is characterised by 5 or 6 pulsations of sound, emergent sustained notes that are accompanied by growing distortions or delayed playback of attacks.
- *Line 5:* here, we seem to linger in the resonance of the previous line, the sounds are more continuous and less emergent. The end of the line is quieter – at this point Burton starts playing the guitar, seeming to pass into a new space.
- *Line 6:* here, things are much more chaotic and there are many simultaneous events. We seem to discern a lot more identifiable events, rather than long sustained playback of grains. The line is characterised by what sounds like a slightly distorted, clunky delay playback that jumps around the stereo space.
- *Line 7:* the chaotic nature of Line 6 is continued, until the end of the line which is marked by a suspension over a sustained note which recalls the beginning of the piece.
- *Line 8:* for the first half of this line, we return to the sustained notes of the beginning of the piece. Then there is a marked impact that fades into a more ethereal sustain. From this emerges a rise of activity that leads into the next line.
- *Line 9:* the beginning of the line is marked by a large, bassy impact and from its resonance emerges a hive of activity recalling the middle of the piece. This is suspended in a higher-pitched note, then the same thing happens again.

We note that the musician now has a lot more room to play again – the chaotic machine blocks emerge from the musician’s sustained recorder notes. The nature of the ensemble remains very segmented, with clear different events.

- *Line 10*: the line starts with no machine blocks, then sees a return of sustained, ethereal grains in the higher-pitched register. There are a few accents of chaotic activity which fade out quite quickly. The piece ends on a machine block sustain which cuts out violently.

2.5. Alice Eldridge – Feedback Cello.



Figure 5: segmentation of Feedback Cello. (Figure from Chapter 2).

I have divided the piece into four large sections. The piece begins with an introduction, introducing us to the various elements: in Intro1 Eldridge starts with long, sustained, unstable notes which hover between harmonics – she is barely touching the strings with her left hand. In Intro2, electronic feedback begins, and the cello gives more grating, lower gestures. This sustained feedback continues into Intro3 where the cello plays more stable notes and gives fragments of more traditionally melodic gestures.

The feedback fades out at the start of Section A and will be articulated differently across the section. In A1 and A2 it intervenes momentarily, pulsating, and then begins to sustain and grow in intensity until Section A5 where it seems to peak before fading away in Section A6. Across Section A, the cello explores different strikes and plucks, moving away from the sustained and bowed content of the introduction. In A1

Eldridge plucks the strings, in A2 she taps the strings with her bow, in A3 she taps the strings again, this time with a percussive ball stick, in A4 she begins tapping the body and the bridge of the cello, in A5 this continues with more intensity giving a more percussive aspect to the attack, and finally in A6 she taps the strings below the bridge. There are also notable moments where we hear emergences of a low resolution, bitcrushed treatment of the signal.

Section B is comprised of two similar, somewhat symmetrical gestures: these are comprised of high amounts of feedback being generated from plucks or bowed cello, which are then left to ring out, and begin to appear to be modulated by the cello. Similar to what shall be seen in Harker's piece, there seems to be a play here between interchanging the source/modulation dynamic: sometimes the cello generates the electronics and feedback, sometimes the cello is modulating the feedback. Finally, the Coda is comprised of two main parts: Coda1-1 and Coda1-2. We reach a point of culmination, first with a section of high feedback and bitcrushed sounds now being merged with the cello; and then a somewhat chaotic section of more traditional melodic gestures from the cello, and bitcrushed material where the electronics and cello seeming to fight to match each other's partial content. The final Section Coda2 sees the cello remain silent as the electronics perform one final swell before ringing out into silence.

2.6. *Sam Pluta* – Neural Duo I and II.

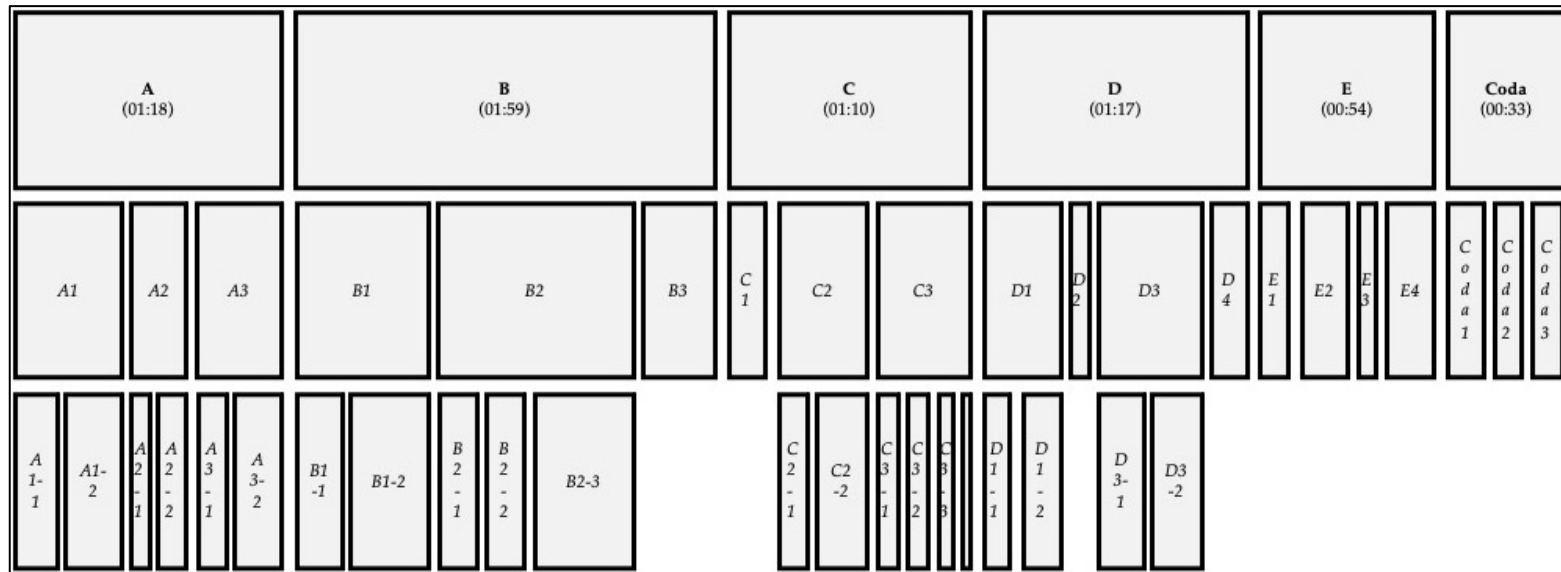


Figure 6: segmentation of Neural Duo I. (Figure from Chapter 2).

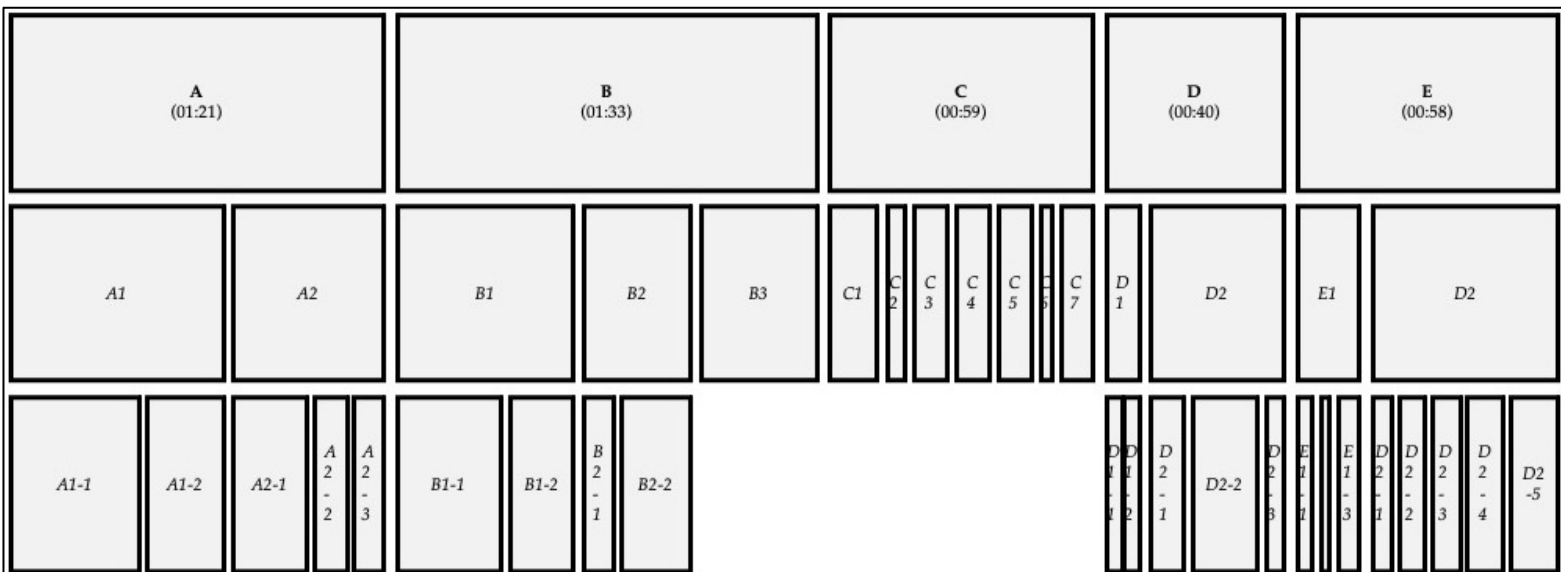


Figure 7: segmentation of Neural Duo II. (Figure from Chapter 2).

Part I can be divided into 6 large parts. Section A introduces us to the kind of material we shall be hearing. Evans will play his trumpet without a mouthpiece, and begins this section with short, staccato blown flutters. The electronics gradually come in, with a low-pitched layer at first, then adding high-pitched noise that somewhat mimics the rhythms of the trumpet. At the start of A2 both the electronics and the trumpet speed up, before going into longer gestures. In A3, the trumpet returns to the

staccato flutters, but the electronics remain sustained, finishing in a suspension with sustained trumpet gestures and long, mobile electronics in a high register.

Section B returns to staccato gestures in both instruments which are continued in the electronics while the trumpet moves to a slower, steady repeated rhythm. In B2, this is repeated in each subsection, the trumpet rhythm getting faster each time and the electronics getting more chaotic and noisier. Finally, in B3, the trumpet progressively slows down while the electronics continue their intensity.

C1 begins with the trumpet returning to material similar to the beginning, and the electronics sustain a steady looping of a phrase which gets pitch shifted. Then this gives way to C2, where the electronics give slow, sustained, and mobile notes; the trumpet a steady slow pulse. This pulse gets louder, and the electronics get progressively more chaotic, shifting violently between extreme high and low registers. C3 is characterised by the trumpet which get steadily more intense – it starts with a whistled blowing, then strained blowing and finally chaotic gestures. The electronics have an inverse gesture, starting chaotic and falling into long downward sweeps and sustains.

Section D1 is characterised by an extremely fast, steady, continuous trumpet rhythm, the electronics move from sustain to chaotic. D2 gives a long steady electronic sweep, and the trumpet breaks into a chaotic gesture, before falling back into the steady rhythm in D3. As D3 progresses, the trumpet gets more and more unstable, and the electronics remain in a high register sustain. In D4, the trumpet has fallen into slower gestures, and the electronics still sustain, and a gradual introduction of downward gliding gestures.

Across Section E, the electronics become chaotic, getting progressively more phrased. The trumpet moves from a steady, slow rhythm, to strained sustained notes, to short attacks, to disparate phrases that complement the ones heard in the electronics.

Finally, the Coda finishes with a long suspension: the electronics give a low pulse, with a few high-pitched glitches here and there. The trumpet moves from a quiet sustained blowing, to pulsing blows, to soft windy blowing.

Part II is shorter in length, and changes in character – notably due to the fact that Evans is now playing with a mouthpiece, changing the nature of the trumpet material significantly. I have divided it into 5 large sections. A1 is characterised by windy, pink noise-like electronics getting progressively louder, and the trumpet starting with a soft

melodic line, before playing more sustained notes interspersed with a few low attacks. In A2, the electronics cut, and then give low bass tones with long release times while the trumpet blows softly. The electronics then sweep to a higher register, then cut. The section ends with sustained trumpet.

B1 starts like A1: building electronics which are now brighter over a soft trumpet melody, before moving to sustained trumpet and a sustained clicky synth. This synth continues into B2, becoming glitchier and brighter, the trumpet continues soft melodic gestures. B3 gives a long trumpet melody characterised by its cupped modulation; sweeping electronics finally filter out into silence.

The glitchy nature of the electronics returns for Section C and is modulated in different ways across the section: sometimes more sustained, sometimes more mobile. The trumpet is very mobile across the section and structures the segmentation: first it gives a high, soft sustain, then 3 hard high notes, then a soft melody, then a series of notes from extreme low to extreme high, then melodic upward gestures, then a low windy blow and finally a series of hard, long sustained high notes. At the end of the section everything cuts.

D1 sees each musician play in turn: first the trumpet with soft, mid-register sustained notes, then the electronics with a growing impulse. Into D2, this impulse then continues, and the trumpet plays melodic gestures on top. These gestures continue and get busier, the impulses get held into a sustain, before finally the section ends with the electronic sustain on its own.

E1 begins with a repeated trumpet gesture of long then short notes accompanied by chaotic clicks in the electronics. The clicks continue and get dryer and glitchier, the trumpet goes up and down a scale. Finally, E1 ends with muffled sustained trumpet, and the clicks slowing down and becoming more intermittent. The final subsection E2 begins with the clicks still playing alone, which are then accompanied by a short, soft trumpet melody, then low sustained trumpet, then a higher trumpet melody and finally a long, almost multiphonic trumpet sustain.

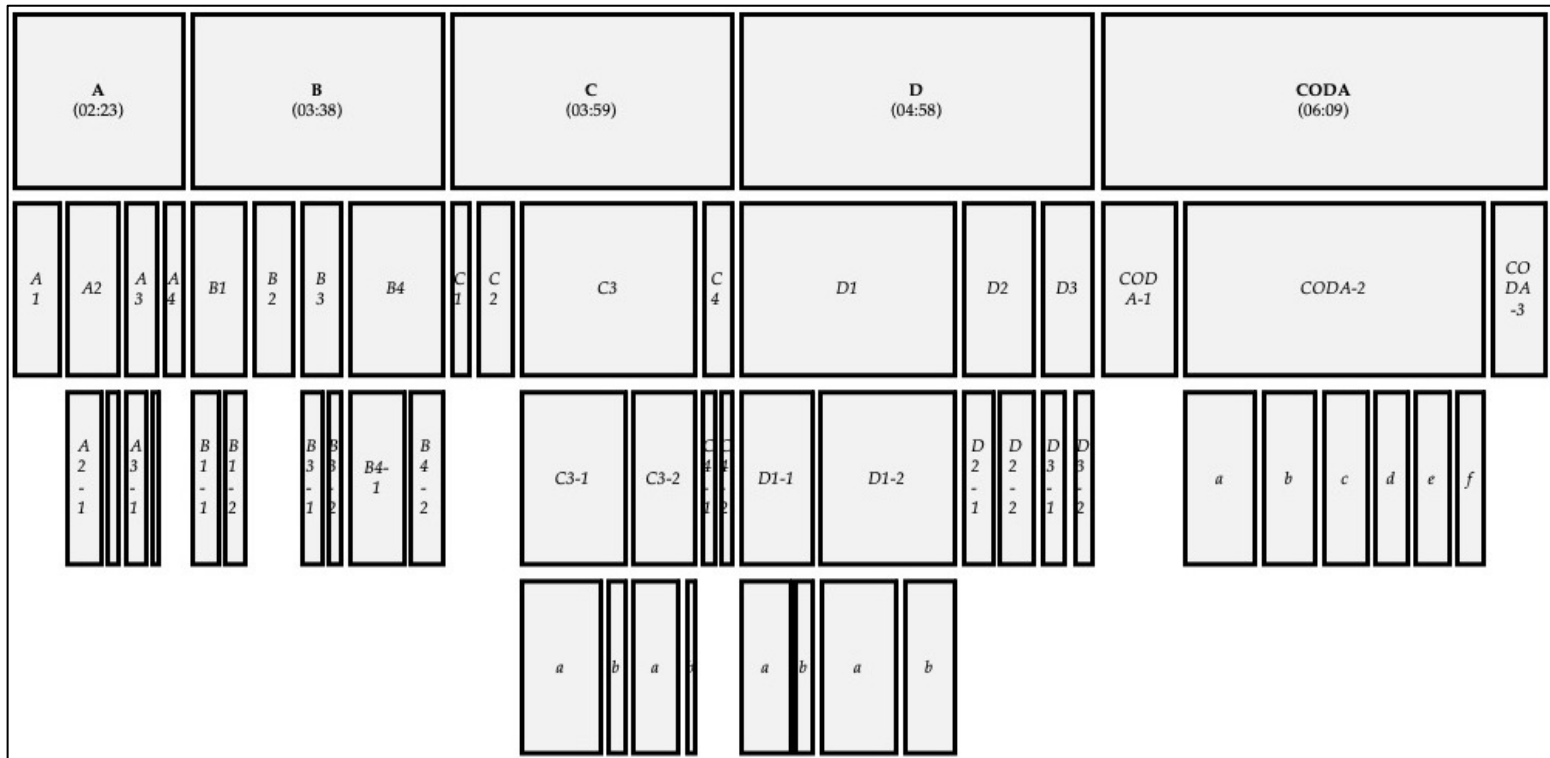


Figure 8: segmentation of Drift Shadow. (Figure from Chapter 2).

2.7. Alexander Harker – Drift Shadow.

Section A can be seen as somewhat of an introduction. The piece opens by presenting us with the kind of material that will be found during the performance: sustained oboe multiphonics, and electronics which appear to prolong the gestures that are made by the oboe. In A1 Dell gives a series of sustained, high-pitched multiphonics that are prolonged by the electronics. This is then followed by two gestures that are similar in nature: A2-1 and A3-1 have less stable oboe flutters with the electronics perduring in the high register, and A2-2 and A3-2 give lower, scraping oboe gestures. Finally, in A4 the electronics fade out with a long, somewhat broken sustained high-pitched oboe gesture in somewhat of a suspension.

Section B starts with the oboe playing in a lower register with more bends and broken material – this is accompanied by more dynamic electronics with the presence of more notable wind-like textures. As the section progresses, we move towards two gestures, B3 and B4 where the electronics and oboe are very interwoven, the electronics prolong the oboe greatly, and each subsection fades to somewhat of a suspension.

Section C starts with the oboe once again becoming more mobile, with few electronics, and then a moment of polyphony between the oboe and electronic partials. This is followed by two similar gestures in C3, long, rich sustains this time in the electronics

which here seemed to be prolonged by the oboe, then a hard oboe gesture which kills the electronics.

Section D starts with D1 which is comprised of gestures where the electronics appear to mimic the oboe in some way. D2 sees a gesture which is similar to that found in C3 and stops on somewhat of a pause before going into D3, where here the long sustains of the electronics appear to be modulated by the oboe. This announces material to come in the code: Indeed, the main part of the Coda, Coda 2, is playing with an ostinato in the electronics. Dell appears to add herself to the sound and modulate it slightly by imparting energy into it. Each of the six gestures finishes with a suspension in the electronics, and each iteration grows further and further from the ostinato with longer suspensions. The piece finishes with a short section on somewhat chaotic, pulsating gestures in both electronics and oboe, before cutting abruptly.

2.8. Hans Tutschku – Sparks.

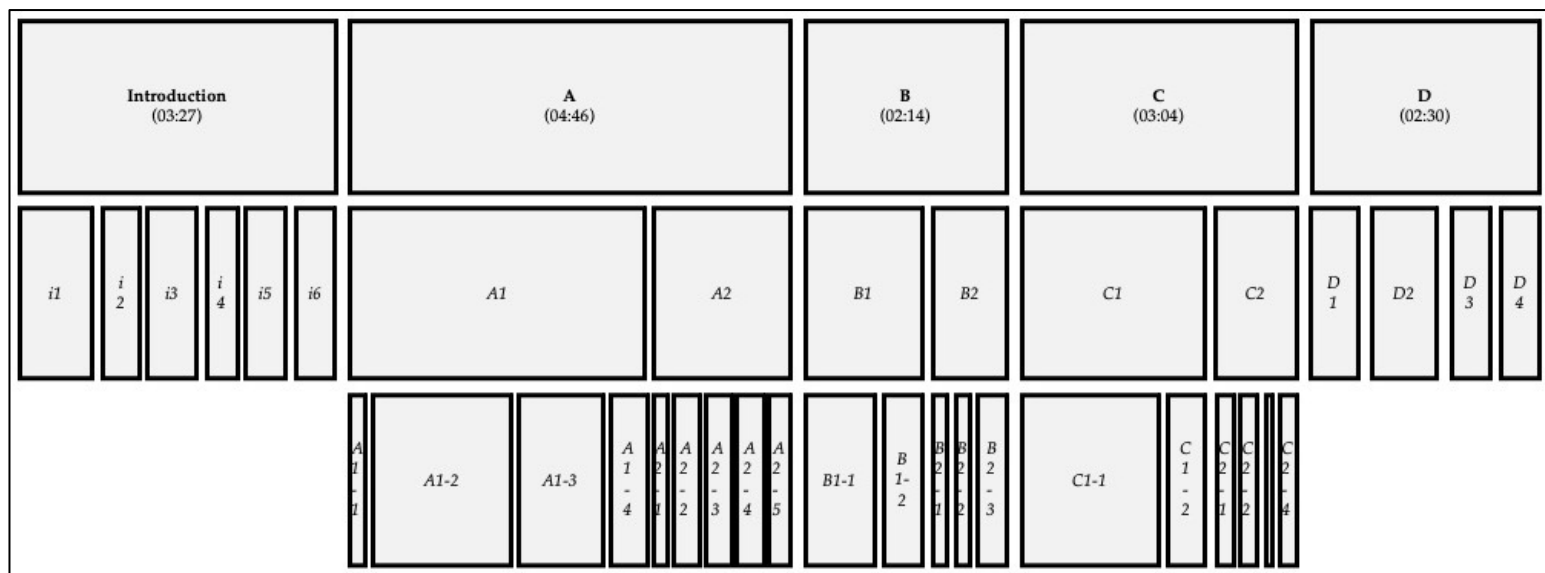


Figure 9: segmentation of Sparks. (Figure from Chapter 2).

The Introduction presents the various types of material that shall be heard throughout the performance: in the piano, there shall be contrasts between sections of somewhat busy, chaotic events, and longer, sustained notes and chords; in the electronics, the two main focuses are on the delay which can vary in time and length and vary in its percussive nature, and the prolonging of resonances of the piano. We are introduced to all of these elements in some form during this first part.

Part 1 is divided into two parts which are somewhat symmetrical: each part starts with a fairly chaotic section, and then suspends over several series of repeated notes, in A1

these repetitions have relatively few electronics, in A2 the electronics seem to draw out the resonance more.

The central Section B follows the model seen in A, this time with a chaotic starting section which is longer in length, and characterised by a long descending bass line, before being suspended in the resonance and repetition of a high *c sharp* in B2. In C, this structure seems to have flipped, resting in a resonant section C1 for a time, (but with a more mobile piano that moves around and plays trills), before a short, busier staccato section in C2 which then suspends itself towards the end.

Section D sees the piece slowly winding down, starting with steady sustained notes which have long, sometimes scraping, resonances, and progressively filtering out towards the end.

2.9. Richard Devine – Constructors.

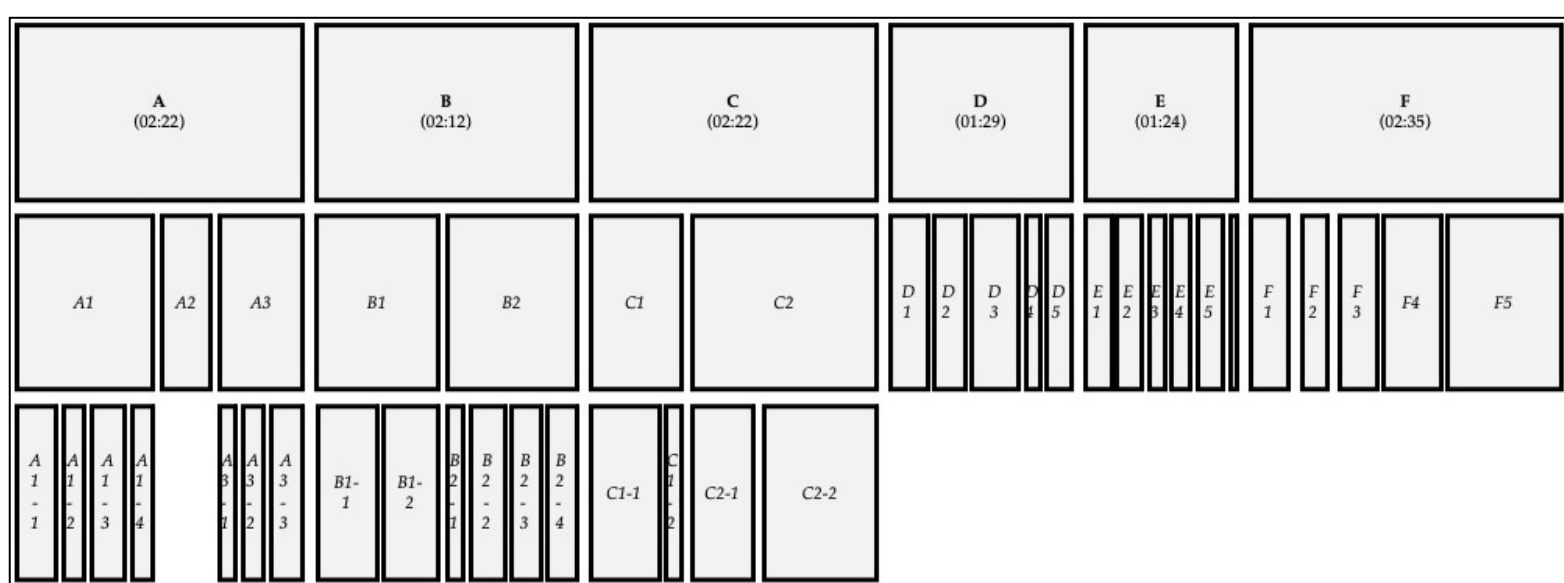


Figure 10: segmentation of Constructors. (Figure from Chapter 2).

Section A is comprised of two large subsections A1 and A3. A1 starts with a series of impacts – some of which appear to be pre-recorded, others with the live percussive objects – processed with heavy reverberation. These continue and are accompanied by a low drone, then the drones get higher in register, and the impacts become longer, downward gliding gestures. A1 ends with a suspension with a synth fading out followed by a string-like swell. A2 is similar to the gestures of A1-3 with downward sliding gestures, and then moves into A3 which is a series of suspensions, each triggered by a high-pitched gesture, or an impact.

At the end of Section A we hear a harmonic pad which continues into Section B. B1 starts with the pad accompanied by various chaotic noises, then the pad dissipates and the noises continue in a much dryer reverberation space. Then we have B2 which is comprised of a series of similar gestures: starting with noisy impacts which suspend, then give way to dryer noises. Section C begins with C1, which transforms the triggering impacts of B2 into a chaotic section, then resides in a long resonance – the gesture of B2 is expanded. C1-2 prolongs the previous resonance, inserting a synth with a reverberation with a short delay time on top. C2 repeats this gesture, with the prolonged resonance being expanded even further.

Section D begins with a suspension, and a series of very clear, bright pads with high reverberation. This progressively gets more chaotic, and then moves towards a suspension in D4, before falling into the final section D5 which is characterised by a looping pattern of scraping noises with chaotic noise on top. Section E also sees several suspensions: the first gives way to fluid, flowing synth lines, the second to a more erratic, white noise rhythm, then a low impact which gives way to silence in E3. The next half of the section begins with synths giving upward glides, moving from a resonant to a dryer reverberation, then gives two more suspensions, the first with a noisy synth that progressively gets slower, the second with an impact in the high register.

The final Section F begins with chaotic, noisy material that will be treated in a number of ways. First, in F2 it is taken up again, but with much more space between gestures, then in F3 it is taken and parts of it are frozen, and grains seem to be looped. F4 and F5 repeat this, but with much more reverberation. F5 finally gives way to a series of impacts, which are accompanied by a low drone, before the end of the piece.

3. Software Overview.

3.1. Rodrigo Constanzo.

3.1.1. Wavefolder module.

The *wavefolder* effect module that appears so often in the patch is relatively simple. Incoming audio is first equalised according to the *bass* and *treble* parameters and is then fed into a *gen~* patch which implements a fairly standard wavefolding algorithm (the signal is multiplied by a pre-gain value, set to -48 to 48 dB, it is then passed through a sigmoid function, and finally multiplied by a post-gain value set by the *gain* parameter. The signal is then multiplied again by the *gain* parameter, we add the *offset* parameter and finally multiply it by the *fold* parameter. Finally, the signal is passed through a *sin* function and divided into two paths where everything below 0 is multiplied to be above 0 and vice versa – the depth of that is controlled by the *crease* parameter). Finally, the audio is sent through a standard *overdrive~* object (the factor of which is again controlled by the *gain* parameter) and output. The parameters of these effects remain fairly static and depend on the mode: when green, *gain* 40, *fold* 25 and *offset* 70; when yellow, *gain* 35, *fold* 20 and *offset* 50. There is a part in SP-C, where the *wavefolder's* *treble* is controlled by audio analysis and the others are fixed to *crease* 0, *gain* 69.8, *bass* 37.8, *fold* 85 and *offset* 50.

3.1.2. Cloud module.

With the *cloud* effect, audio is input, and whenever an onset is detected, 180ms of that audio is collected from there. This fragment is analysed, Constanzo looks for the mean of the first derivative of the loudness descriptor and the time centroid in order to get an impression of how long the fragment is. Then, the fragment is trimmed down, and prepared for playback. It is split using the *fluid.buflpss~* and *fluid.bufltransients~* objects. Playback can be in one of two modes: *chunk* where one trigger is activated; and *grain* where a *cloud* of triggers is activated (the *activity* parameter determines the number of triggers). On each trigger, three things are chosen at random: amplitude (variation in this is controlled also by the *activity* parameter); pitch (transposition variation is controlled by the *pitch* parameter, and a *reverse* toggle can activate reversed playback); and position and duration, again based on the *activity* parameter. Then, the sample is played, outputting the three streams created above (harmonic, percussive and residual from HPSS, and transient-percussive). The *width* parameter pans the playback, and finally, the *timbre* parameter blends the three streams together in the

following way: harmonic, 0% = 150, 50% = 0, 100% = 0; percussive, 0% = 0, 50% = 100, 75% = 100, 100% = 0; transient, 0% = 0, 50% = 100, 75% = 100, 100% = 100.

3.1.3. *Dirt module.*

Finally, there is the *dirt* effect module which implements two different distortion techniques. When in *soft* mode, the distortion is a *gen~* cubic non-linear distortion taken from Julius O. Smith's *faust* code. When in *hard* mode, the distortion is a variable-hardness clipping function by Laurent de Soras, also in *gen~*.

3.2. *John Burton.*

3.2.1. *Index of objects.*

Used:

- A standard crash cymbal.
- A smaller bowl-shaped cymbal.
- A standard violin bow.
- A drumming brush.
- Various drumming batons.
- A standard tenor recorder.
- A smaller recorder.
- A wooden ornamental shaker.

Not used:

- A set of coloured party balloons (he did use one of these the previous day when demonstrating his patch at the FluCoMa plenary).
- Another crash cymbal with a deeper bell.
- Some finger cymbals.
- A set of metal balls attached to each other on a string.
- A fibreglass board covered in copper upon which Burton drew a coil and was destined to be used in a percussive way, dragging things across the ridges.

3.2.2. *PPwaveTable.*

The patcher receives the play command, upon which we generate a certain number of values for the unit: x table frequency, low pass filter cutoff frequency and sin frequency. The table makes use of *Max's* native *2d.wave* object. The wavetable is chosen

at random from one of the points defined by the audio analysis and onward 50 samples. The audio is passed through a low pass filter, and gain-enveloped. There is also a sine wave which is output at the frequency of the x position read speed.

3.2.3. *PPdelay4eva.*

The patcher receives a play command. We once again choose a random point in the input buffer from the analysis data. This is sent into a stereo delay line which contains a variable bandpass filter and a feedback line (the level of which is controlled by the *FB_curve* buffer). Delay time is randomly generated and modulated by noise.

3.2.4. *PP4voice.*

The patcher receives a play command, upon which we generate a number of parameters. We choose a random chunk of the input buffer for playback and choose from a set of 4 presets of parameters. The processing is essentially playback of the segment in four voices and varying speeds, each being enveloped and filtered. Global audio is also filtered before output.

3.2.5. *PPchunker.*

When the play command is received, a *metro* of varying speed (generated every 20ms and varies between 25 and 250ms) bangs processing within the patch. There are four streams that effectively perform the same processing: choosing a random chunk from those made available from the audio analysis, and these chosen points are played back through *groove* objects which are filtered before output. The playback speed of 3 of the *groove* objects is modulated by the audio output of the *wavetable* processing unit (the last one is fixed at 1).

3.3. *Alexander Harker.*

3.3.1. *Processing Modules.*

There are three main sound processing modules in Harker's software: a spectral freeze, a partial freeze, and a granular freeze. The spectral freeze module is based on the stochastic phase vocoder that Harker created for Framelib and for the HISS tools. For each spectral bin, 4 numbers are calculated: the mean amplitude, the average standard amplitude deviation, the mean phase deviation and the average standard deviation of this mean's phase deviation. Resynthesis is then performed using these and random values, and changes in resynthesis are crossfaded, allowing for gestural shapes rather than clean cuts. The partial freeze works by reconstituting the

recognised multiphonics with 50 sine waves – the volume and pitch of which have been pre-analysed on the recorded training material. Harker can perform gestures within these parameters using his *gesture generator* patch and crossfade between either the reconstitution as a sine wave or as filtered noise, allowing for what he described as “*noise bursts*”¹. Finally, Harker described the granular freeze as fairly simple², with the notable presence of a random filter for each grain.

¹ Appendix 8.8.4.

² Appendix 8.8.5.

4. Exploration of descriptor, dimensionality reduction and FFT size combinations with the metal resonance corpus.

The first collections shown uses the mel bands and MFCC descriptors with quite a standard FFT size of 1024. When looking at the PCA and MDS reductions of both datasets (the dimensionality reductions available with the FluCoMa tools at the time), we see that they yield big, tight lumps of slices with various strands tapering from them. With the PCA reductions, there do seem to be some groupings in these extremities, but amid the large central hive of slices, adjacent samples can vary greatly in nature. The MFCC dataset is slightly less dense, but, overall, these large groupings don't appear to be very useful to us. The MDS reductions yield similar results in shape and structure, however the proximity of samples does appear to be more meaningful. Unsatisfied with these results, it was at this point that I decided to look to scikit-learn to see what other dimensionality reduction algorithms would yield.

First, Isomap yields results which are similar, however the strands do seem to have more structure (for example, the various arms protruding from the central mass are denser), and the central hives are not quite as consequent. Next, KMeans clustering yields interestingly shaped networks, however the relationships here between slices seem to be more sporadic. Next, Locally Linear Embedding: this is interesting as it produces a very tightly knit network with long protruding arms which do have very meaningful spatial relationships. Next is the self-organising map which was promising at first as it is designed to deliver a space that is entirely filled by the data – in practice I discovered that slices seem to have very poor timbral relationships regarding proximity. Finally, there is the t-SNE algorithm which yields some very interesting results: the data seems to be well distributed across the space, and we discover very evident small clusterings of sounds. This algorithm seems to work very well indeed – when consulting the file names of all the various slices, we indeed see that sounds of identical sources have been grouped together. This is encouraging, except perhaps for one aspect: the number of clusters is still relatively high, a problem that shall be discussed again below.

Before trying other combinations of descriptors, I tried some different FFT sizes for the audio descriptors to see what difference that would make. I decided to look at two extremes: a low 128 frames, and a high 4096 frames. For the mel bands PCA combination, FFT size didn't seem to make much of a difference, however for the MFCCs an interesting phenomenon occurs: with an FFT of 128 the big grouping splits

into two distinct lines, the members of which unfortunately have no real distinctive features, and at 4096, the data concentrates into one very dense line, the left of which seems relatively bright, the right extremity being relatively dull. For the MDS reduction of mel bands, the 128 FFT size yields a central mass that disintegrates, and the 4096 another mass with slices tapering off to the right – in both cases the spatial relationships aren't very pronounced. For the MFCCs, a similar phenomenon to the PCA occurs – at 128 FFT, there are two, perhaps three, linear groupings – the timbres of which do share some vague similarities; at 4096, we have a large central hive that goes from long and resonant on the left to muffled on the right, with a few outlying slices whose decentralised positions don't appear to be notably justified.

At 128 frames, the Isomap distribution for mel bands concretises into a three-spoked layout – longer, more resonant sounds seem to be situated along the arms, muffled sounds towards the middle; at 4096 frames, we have the same phenomenon but with only two arms. Regarding MFCCs, the 128 FFT size yields a ring-shaped layout with a definite continuity in the timbres, whereas the 4096 size yields a much more evenly distributed spacing. For Locally Linear Embedding reductions, changing the FFT size doesn't seem to yield very different results, nor does it for the self-organising map whose results appear to be just as fruitless. For the Kmeans reductions we get similar results – the higher FFT size, the tighter the shapes seem to become. Finally, for both mel bands and MFCCs, the t-SNE FFT changes yield very similar results, with the exception of MFCCs at 128 FFT: here we find ourselves before the same small groupings of sounds, but this time grouped into three clear regions. It would seem that t-SNE is showing itself to be a potentially very useful algorithm for the uses we wish to make of it.

I also experimented with some other descriptors. First, the collection of spectral shape descriptors. The first PCA reduction yields a fairly evenly distributed space, but with spatial relationships being more or less meaningless. The MDS reduction yields a slightly tighter-knit distribution where local proximity does seem to have meaning, but the macroscopic structure can be sporadic. The Isomap distribution yields a relatively evenly-spaced mapping, with sounds with bounces seemingly placed towards the right – incidentally, this is a distinction that Constanzo made when sorting through samples. It is not surprising to see samples distributed this way, due to the presence of descriptors such as *spectral crest*. The Locally Linear Embedding yields interesting results that are also distributed according to bounces with two long

strands that meet at one extremity. We find bounces articulating the structure again with the KMeans distribution, the closer toward the top-right of the shape, the bouncier the sound. The self-organising map has once-again failed, yielding a very evenly spaced but unhelpful mapping. Finally, the t-SNE reduction yields another useful mapping, clearly articulated by bounciness and by timbre.

Next, I tried the loudness descriptor to see what kind of results that would yield on this kind of audio. If we listen to the sound plot in terms of perceptual loudness or timbre, the PCA reduction seems to yield some quite disappointing results. The MDS reduction has the same evenly distributed result, however samples do seem to become somewhat perceivably louder towards the top of the plot: here, loudness seems to coincide with bounciness as this can be perceived as a more violent impact. The Isomap reduction is, again, fairly evenly distributed, with loudness seeming to increase towards the bottom of the plot. The Locally Linear Embedding reduction again gives two lines that meet at one extremity, the larger line seemingly comprised of louder sounds, the smaller one quieter, longer sounds. The KMeans distribution also seems to hold meaningful relationships, however everything is contained within quite a close cluster. The self-organising map yields yet again unhelpful results. Once again, the t-SNE reduction seems to yield very meaningful results with the louder sounds towards the top of the plot, but once again with lots of small clusterings.

Finally, I did some experiments with the pitch descriptor, on the one hand including pitch confidence in the reduction, on the other without. We would perhaps expect to see noisier samples grouped together when considering confidence, however this does not seem to have been the case. With confidence, the PCA reduction yields a space of two extremities with some slices stretched between the two, and without confidence an evenly distributed space. For computing pitch, the network without confidence seems to work best – we note that from the top (lower pitches) to the bottom (higher pitches), the span of the pitches isn't very large. The MDS reductions yield very similar results – here the plot with confidence does seem to regroup a number of the noisier sounds in the left-hand clustering. The Isomap reduction with confidence seems to have two axes: the x axis being noisy to clear, the y axis high pitch to low pitch; without confidence the plot does a fairly good job of going from low pitches at the top to higher pitches at the bottom. The Locally Linear Embedding technique yields some quite bizarre results – with confidence we find two vertical lines where most samples have a similar pitch, on the left they are noisier and on the right

clearer; without confidence data seems to be distributed along one diagonal line in a sporadic fashion. The KMeans and self-organising map reductions yield unhelpful results again. Finally, the t-SNE reductions yield again lots of small clusters, however the meaningfulness of them this time seems unclear.

5. Iterative Approach Sound Plot Clusters: Metal Resonance > Wavefolder.

- 1: here we find long, very distorted, high-pitched sounds (the original source was the bar file), that end with a quick, low-pitched bass tone.
- 2: here we have the same kind of hyper-distortion, with a mix of sources. When the source is the bar file, there is no longer the final bass tone, when it is the girder there is.
- 3: this cluster is composed of the girder drop, getting progressively more distorted as we move from the top to the bottom.
- 4: this is composed of hyper-distorted girder drops, and, similarly to Cluster 2, there is a synthetic note as the sound cuts off, this time much higher pitched, seeming to follow on from the sample.
- 5: again, this cluster is composed of distorted girder drops. As we move down this long sweep, we find more and more distortion, and rediscover the final bass note at the end which gets more and more intense.
- 6: this cluster is composed of bar samples. At the top, the samples are dry,

and get progressively distorted as we move towards the bottom. At the bottom, we find the sample distorted to a point where the beginning of the sound is almost white noise from which the strong partials of the sound slowly emerge.

- 7: the top part of the cluster is comprised of relatively undistorted girder drops, the bottom of heavily distorted bars (but not as distorted as the bottom of Cluster 6).

- 8: this is comprised of bar samples. We find an effect similar to Cluster 6, however we find a new phenomenon: what seems




Figure 11: Metal resonance > wavefolder sound plot from MFCC audio descriptors and t-SNE dimensionality reduction. (Figure from Chapter 4).

like a more percussive oscillation that emerges from the sound as the samples progress.

- 9: this cluster is comprised of heavily distorted girder drops. Here, the distortion seems to very much bring out the higher tones. This sound is followed by the low-pitched, synthetic bass note.
- 10: this cluster is comprised of distorted bar samples. Here, we find more of these first very noisy samples from which emerge the strong partials and more percussive oscillations.
- 11: this cluster is very similar to Cluster 10.

6. Burton Alternative Performance Patch Modifications.

For the first level, I had to disconnect all the events regarding score generation, as I would be loading in the scores generated on the night. I had to be wary that the setup routine for the patch still ran everything else it needed to. I also had to set up a system for loading the score, and one that would send out triggers and block count during the piece so that the input audio would play at the right time.

For the second level, I had to reconnect all the parts of the score generation for the machine – this meant modifying the previous system so that whenever a human block was reached, it also triggered creation of machine blocks like in the original patch.

For the last level, the original score creation system was restored; however, I needed to make some modifications to the playback of audio. Now that the human score would be modified, logically, it could be supposed that Burton wouldn't have necessarily played the same thing whenever there was a human block. There are various ways that one could approach this question: I decided to implement a simple system where the length of the human block would be read, and a random slice would be chosen that was of a similar length and sped up or slowed down to fit accordingly. Note that, with this method, we lose an aspect to the human input of intentionality in the type of instrument played, and perhaps the nature of the playing from block to block. With more time, we could imagine a system that would be aware of its own playing and try to construct some kind of logic across events. I believe, however, that the system I produced was good enough for this iteration of analysis.




Figure 12: Pass 1 of Hayes' sound plots using a t-SNE dimensionality reduction with clusterings. (Figure from Chapter 5).

7. Hayes Pass 1 Full Cluster Analysis.

I have identified 16 clusters: as the reader shall notice, the first 6 are quite apparent when glancing at the plot, the others, belonging to a larger mass, required closer inspection to be identified.

- 1: here we find only content from the rehearsals. It is comprised of very light pops, resembling a filtered woodblock followed by silence. The closer we get towards the bottom, the lighter the sound is, until finally falling into silence at the very bottom.
- 2: this second cluster, also comprised of only content from the rehearsals, seems to prolong the content from the first cluster. The higher we get in the plot, the more synthetic the sounds become. From the middle of this cluster up, the sounds get repeated into longer phrases, giving the aspect of a highly filtered synth note.
- 3: in this third cluster, we find some elements from the samples and from the performance interspersed between content from the rehearsals. The sounds generally retain aspects from the previous clusters, being even more synthetic still, and the repeating of the initial percussive elements into a prolonged note becomes more frequent. The notable segment of the performance is one of the final moments where Hayes let a light synth with an upward glissando ring, triggered by a slightly percussive noise. The samples nearer this element are bowed saws, also with glissandi, and the others are prepared piano impacts, resembling the percussive impact of the rehearsal content.
- 4: this cluster is comprised solely of rehearsal content. It is comprised of very synthetic sounds, that are very loud compared to what we have heard up to this point. On the left-hand side of the cluster, the sounds are short and percussive like in the other clusters towards the left, but as we move progressively towards the right, the sounds get longer until reaching longer, glissandi gestures.
- 5: this cluster, up at the top of the plot, is where most of the segments of the performance are found. This cluster can be further segmented still, as seen in the image:
 - *a*: very short bursts of low-pitched, synthetic sound.
 - *b*: very short bursts of high-pitched, synthetic sound.
 - *c*: very short bursts of cleaner, low-pitched synthetic sound.
 - *d*: very short bursts of cleaner, high-pitched synthetic sound.
 - *e*: bursts of noisier, grainier synthetic sound on top of low-pitched vocals.

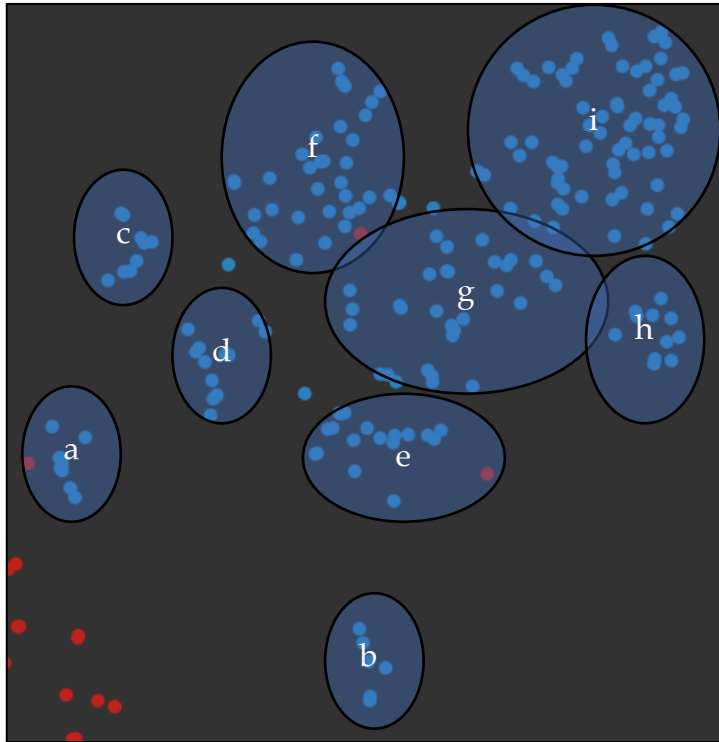


Figure 13: a zoom into Cluster 5 of pass 1 of Hayes' work.

- *f*: more gestural content. Mostly triggered by a percussive sound above either high-pitched vocals or a synth glissando.
- *g*: similar to the content of Cluster *f*, but the vocals and synth sounds are lower pitched.
- *h*: mid-range vocal and synth segments with distinctive *hissing* syllables.
- *i*: longer gestures, with the glitchy vocals being looped in the mid-range, and low-pitched percussion on top.
- 6: only content from the rehearsals. Short bursts of very low-pitched, muddy synthetic sound.

- 7: short bursts of synthetic sounds –

the sound almost begins to sound like it was produced through a very low-resolution FFT process.

- 8: this cluster is centred around a group of segments from the performance. All the content sounds like reverberating resonance – longer resonant spaces where we can sometimes hear remnants of synthetic lines.
- 9: centred around another set of segments from the performance. This time, they appear to be fragments of the final synth glissandi, starting with clicks. The surrounding rehearsal content is similarly longer synthetic lines containing clicks at the beginning.
- 10: this cluster is comprised only of rehearsal content. They are primarily short, synthetic bursts that prolong the content from Cluster 7, getting longer and brighter. It also leads into the clicky content of Cluster 9.
- 11: again, this cluster is centred around segments from the performance. Like Cluster 8, these appear to be resonances, only here, they are much more charged in the low end of the spectrum.
- 12: again, this cluster is comprised only of rehearsal content. From left to right, this cluster eases from the resonant, bassy content of Cluster 11, into the left-most short synthetic bursts of Cluster 5.

- 13: here, we hear vocal processing in the rehearsal content for the first time. It comes in short bursts in the mid-range.
- 14: here, the vocal content from Custer 13 seems to progressively fuse from top to bottom with synthetic content, the vocals becoming less and less present towards the bottom, until we can only really discern high-pitched synthetic noise.
- 15: this cluster, comprised only of rehearsal content (save for one performance segment) presents bursts of heavily pitched synthetic sound. Moving from left to right, the pitch of these sounds gets higher and purer. As we move across the cluster, there is also a level of distortion which grows from left to right.
- 16: finally, this cluster presents rehearsal content of short bursts of synthetic content which seems very compressed and slightly distorted. The pitch tends to be high.

8. Interview extracts.

8.1. John Burton.

8.1.1. 24th April 2018 (*Interview with Owen Green*).

OG: How do you see yourself as a musician? How do you make music?

JB: [...] Then I went to art school. Again, my parents were supportive. I ended up making work about being at art school, which was lovely, self-referential stuff. I ended up making a little museum of what it's like to be at art school. That was the final show, and it had an album as a part of it. I was really into Bonnie 'Prince' Billie and those kind of Americana things. Then I bought a very simple PC to do my dissertation in 1997. I went to one of those shops on the high-street where you could buy a PC. That's when I discovered sampling – that you could take sounds in and then process them.

8.1.2. 24th April 2018 (*Interview with Owen Green*).

OG: Do you still use offline programs these days?

JB: [...] I had a bit of a change in my aesthetic. On the [first] four or five albums, I had really been on this personal journey into transforming sounds. Like Parmegiani's 70s music which I really liked. I could hear all these stretched things and beautiful layers of stuff, but what I never thought about was the quality of that sound – how tape makes it saturated in a certain way. There aren't ridiculous amounts of high frequencies going on, there's no spectral bit where you'd get these peaks coming out. It never occurred to me that that was quite an important part of that sound.

8.2. Olivier Pasquet.

8.2.1. 21st September (*Artist presentation at FluCoMa plenary*).

OP: If you look closer [at the piece Alpha Centauri], you can see that you have a lot of patterns, like millions of dots of things. If you look even closer, you can see layers overlapped. [...] If you look closer [at the piece Japan] you can see that you have another kind of layering with very minimalistic things, with dots, just dots. And you have some kind of rhythms.

8.2.2. 8th November 2018 (*Interview with Jacob Hart*).

OP: *Si j'envoie de grains dans une salle, il va y avoir une réverbération tellement grande qu'on va rien entendre. Donc ça fait une espèce de granulaire de base ou un brouilla ambient. Et comme j'ai fait beaucoup de pièces de théâtre, j'ai l'habitude de travailler avec des contraintes. Et cet environnement extérieur à mon processus compositionnel – il fait partie maintenant de*

mon processus compositionnel. J'ai plein de fantasmes, et après je les met en confrontation avec une réalité. [Depuis 3 ans] je fais pas mal de pièces site-specific [...] et c'est hyper intéressant. [Mes pièces] ne se sont jamais tenues dans le format CD stéréo, je n'arrive pas à entrer dans ce moule, même si ça me passionne et j'ai très envie de faire en sorte que ça sonne super bien au casque. Mais mon truc en réalité – peut être que c'est une manière de me cacher – c'est de faire en sorte que mes pièces existent dans un espace ou dans des conditions particulières.

[...] Ça donne un aspect social de la musique hyper-intéressant et très différent du paradigme de concert, de l'écoute CD. [...] [Je voudrais faire quelque chose] qui amène les gens non seulement dans la temporalité de la pièce mais aussi dans l'espace lui-même.

Translation by the analyst: OP: If I send grains into a room, there can be such a big reverberation that we can hear nothing. It becomes a basic granular synthesis or an ambient blur. As I've worked a lot in theatre, I'm used to working with constraints. This environment which is outside of my compositional process, has now become part of my compositional process. I have lots of fantasies and formalisms, and they come into confrontation with reality. [Over the past 3 years] I've done a lot of site-specific pieces [...]. [My pieces] hardly fit into the stereo CD format, I can't fit myself into this mould, even if it's great and I'd love to make things sound great in headphones. But my thing in reality – maybe it's a way of hiding myself – is to make it so that my pieces exist within a space, a social environment; in particular and singular genuine conditions.

[...] It gives a really interesting social aspect to musique which is very different from the concert paradigm and listening to a CD.

8.3. *Lauren Sarah Hayes.*

8.3.1. *22nd November 2019 (Interview with Jacob Hart).*

JH: I was also curious to know about your vocal vocabulary. [...]

LSH: When I came to academia and started doing this stuff, I just stopped using vocals at all – it was just too strange to do that. But then I started to incorporate it just with vocal utterances and things like that. I've done various workshops over the years and had various instructors. As part of the International Showcase in 2017, at HCMF, I got a bursary to use. So, I worked with a woman called Micaela Tobin, her performance name is White Boy Scream. She's a Filipino American who teaches at Cal Arts. She's a trained opera singer, [...]but she's also a noise musician. So, she does incredible operatic noise, vocal stuff. I've been having some lessons with her, she led two

workshops a few weeks ago where she's given me a lot more ideas and vocabulary and extended techniques and how to do them safely. But I've not really incorporated them yet, because I was focusing on the code. I want to work much more with that in future and yeah, I mean... She was just like, "take off so much processing of your voice, let's just hear the voice more and do some looping, do you have a looper?" No, I don't even have a looper on my patch! So, we were working with a looper just looping vocals. And that's something I'll work towards more, but at the moment, it's just really using it as a sound source. And because so much of my patch is attack-based, I have to do a lot of stuff that's like *ts ch*, like those things to get something to trigger.

8.3.2. 22nd November 2019 (*Interview with Jacob Hart*).

JH: You've done a lot of work with haptic feedback in these instruments to get that kind of kick back and that resistance from them. With the game controller [...] I assume it's more difficult to get?

LSH: That's why it's the opposite. So, the game controller, Martin Parker has written about this [...] He talks about using a joystick, being so easy to move and so easy to reach the extremes. [...] Some of the things I want to achieve require holding it incredibly still. If I move a millimeter, it's going to drastically change things. I find that resistance by the opposite. Instead of having to push really hard, it's actually having to hold it. When you have those things that move really easily, that becomes difficult to do. You can hear, you can see, when it's playing back sounds really, really fast. I have crude scaling – it's like 0-255 on those joysticks. I'm not smoothing it because I love having those kind of digital steps that you can hear. So, if I just move it one, the pitch is going to completely change. [...] So, it's holding that is difficult. And then if I want to do something else, then sometimes I'll have to use my mouth to hold it, or something like that. So yeah, that's how I create those kind of tensions.

8.3.3. 22nd November 2019 (*Interview with Jacob Hart*).

LSH: Then, just using HPSS as a filter for percussive elements.

8.4. *Rodrigo Constanzo.*

8.4.1. 22nd November 2019 (*Interview with Jacob Hart*).

JH: Could you talk about what the piece was for you? [...]

RC: [...] I had x amount of behaviours, x amount of sections, or types of processing material, and to a certain extent, musical material. But, with this one there were,

because of some of the little robot stuff as well I had... And because I didn't want to saturate that, I knew I wanted to do that twice, and the purpose of that. And I knew one of those would happen at the end, I didn't know if it would be the very end, but it would be very near the end if it wasn't. So, I kind of front-loaded that as an aesthetic thing. And I guess to a certain extent because I didn't want that stuff to be something that was just kind of peppered throughout the performance. Whereas some of the other things I was doing, I did allow myself free rein to go in out of some of these processes. The crotales I did not want that to be the case, I didn't want to tip the hand that it was going to happen before it happened. So, it happens once and everyone's like "oh, what?" And then, when it happens again, [it becomes] a much more impactful moment. And then beyond that, I guess, medium scale structure wise, I left really open, knowing that there would be some ebbs and flows, building towards a large section at the end. [...] Yeah, I think structurally, that's about where I was at: I had a couple signposts, and I knew that there was going to be some sort of trajectory that happened within that.

8.4.2. 22nd November 2019 (*Interview with Jacob Hart*).

JH: During the gig there was some feedback. What were your strategies for that?

RC: Yeah. I mean, there's an aspect of where I was playing with it. There's some moments when I was focusing more on just purely that kind of language. I didn't want feedback there. [...] At some point, I kind of felt a little bit coming up, so I would kind of ride that. But then there's other aspects where I just played with it. [...] If I had an issue with feedback I could have I could have killed it elsewhere. [...] At the very end of it [...] I remember this cool feeling, and it's something I normally didn't do while playing but it worked really well: I had a long bass drum going along a sample decay that was fading away. And I think I might have had the wavfolder on and I was sort of getting feedback from the microphone module doing a vibrato. [...] While I was doing it, I felt like a fancy man just sitting here [doing a very overtly musical] kind of movement. But yeah, so some of this was actually playing with feedback. [...] I think there was more feedback than I would have liked, but not enough to be problematic.

8.4.3. 20th November 2019 (*FluCoMa Plenary presentation*).

RC: [Concerning latency] 512 samples was about as far as I was willing to go. I'd like to narrow that down.

8.4.4. 22nd November 2019 (*Interview with Jacob Hart*).

RC: I don't have a pitch value coming in from the bass drum, but I do have a pitch value associated with that gong. So, if there were three pitches of gong, then what I do is I take a random entry from the gong pitches. [...] And then I transpose it by this much. So basically, even though I have three notes of samples from my sample library, I'm correcting those to be whatever notes from here.

8.4.5. 22nd November 2019 (*Interview with Jacob Hart*).

RC: [The crossfader] formed the material language for everything, as in the local to medium scale musical events and gestures that were generated as well as a very distant, tangentially-removed stylistic, idiomatic aspect – the nature of a fader, how it moves, how it behaves, how specifically applying this kind of technique to it commands a type of language.

8.5. *Sam Pluta*.

8.5.1. 2nd August 2021 (*Interview with Jacob Hart*).

SP: Because of the circumstances of us being in different countries, across an ocean and not really being able to play live – one person playing and another person blowing over the top of that – it forced me to only play the instrument that was commissioned. I made that aesthetic choice. I'm not doing any processing on Peter; I'm only playing the synths. I'm playing them through my instrument – so I'm doing processing on them, and that's about it. [...] The loud, fast one, I play first, and Peter played second, and then the other was the opposite – Peter threw down the track and then I played over the top. Now he didn't know what I was doing, he thought we were doing just a normal set; I didn't tell him that I was only going to play the synths. So, the second one is really weird in a way because of that.

8.5.2. 2nd August 2021 (*Interview with Jacob Hart*).

SP: I wanted there to be two pieces. For the first one I went for the all-out, seven-minute noise fest. Peter particularly said that it was challenging to play with because it was so on all the time – there wasn't much room for him. So, instead of going all over the place, he just stuck with one idea, and I think that worked really well. Then, the second one, he plays all these different, disjunct objects – some of which are very tonal. So, I kind of tried to play an accompanist role [...] trying to make a contrapuntal music that isn't necessarily connected, then seeing what that result was.

8.5.3. 2nd August 2021 (Interview with Jacob Hart).

SP: For me, the piece is really the software, and the piece is this instrument that is now part of my setup. That, to me, is what was commissioned – not necessarily the performance that we saw. It's something that's now part of my setup and isn't going to go away. It's really integral to what I do at this point.

8.5.4. 2nd August 2021 (Interview with Jacob Hart).

SP: I don't use samples. I have a hard time getting them to do what I want them to do. Whereas the synths do what I want them to do. I had a version that played a bank of 3000 samples where I used the PCA to get them into a 2D space and played the 2D space with the joystick. But I was never happy with that because to me it sounds like samples.

8.5.5. 2nd August 2021 (Interview with Jacob Hart).

SP: [Part II] I think definitely feels like we're in separate rooms, but I like that. The first one I feel is almost as if we're in the same room, and it's kind of freaky in a way where it's like: "oh [no], maybe the thing that we're doing in a live setting isn't that impressive anyway, because this is tricking me, and we're not even live!"

8.6. Alice Eldridge and Chris Kiefer.

8.6.1. 10th September 2021 (Interview with Jacob Hart).

JH: There were moments where the cello seemed to be a source that were prolonged in the electronics, other moments where the feedback was rolling, and the cello came and modulated the signal. Is there a duality to be thought of, or is it something that is fused into one system?

AE: Yes, that was totally intentional. That was how we designed the system [...] we ended up thinking that a *shared instrument* was the best way of describing it. Because the cello signal go through the modular, and Chris can [mess] with them; but the cello is also effectively controlling the modular – it's controlling the neural network which is controlling the modular. Then Chris is also controlling the neural network which is controlling the modular.

CK: To think of it as a duality, you'd have to think of it as two separate systems. But it's not – I conceive of it as a whole with some complicated parts in it.

8.6.2. 10th September 2021 (Interview with Jacob Hart).

AE: What the MLP gave us was this extra layer where it wasn't just Chris fiddling with knobs and finding a good space: we've got another layer up where the cello is controlling. [...] That's why there's bits, like after the opening section, where I do rhythmic, banging stuff, and there's a big pizz twang – that was a bit, throw caution to the wind, expecting that that twang is going to poke the MLP and push the system into a new space. You start making a new musical language on the cello as controller. [...] Your musical gestures are control signals. [...] It has a funny kind of dual role.

8.6.3. 10th September 2021 (Interview with Jacob Hart).

CK: The role of the MLP in this was just to add interesting complexity dynamics into this feedback loop and make it another agent in the way the thing works, that would do interesting, serendipitous [things]. Sometimes it did, sometimes it didn't. We had a few golden moments where [...] the system would start modulating against itself – you would trigger it, then it would go off down this wonderful path of stuff and sounds amazing. We'd seen that in training, and that was where we were trying to get to.

So, when we first trained it, it sounded really good, and then we had some setbacks where it just wasn't working. [...] This was to do with gain structure and making the system really predictable [...] but also taking a systematic way into training. So, we explored lots of different methodologies for trying to get to this network which is a wonderfully serendipitous, self-modulating system. You can't really define that: it's not like normal AI training where you can train it to recognize cats and dogs. So, creating the dataset for training the network was a really interesting part of this composition. So, we started doing these systematic methods where we'd say, [...] "Ok, this time we're going to collect lots of little ten-second bursts of sound and log an average out and connect them with different spaces; or, let's do just two different pulls and long sounds from the cello [...]". All these different strategies.

That was the kind of choreography of training this network, that was a kind of improvisation: how do we systematically improvise to get something that might be good? Or is more likely to be good than not good. I think in the end, I got quite an intuitive feel for it, and the model we made – I said to Alice, just play, and then I collected some stuff, and just got a feel for when and where to go, and we got one that sounded quite interesting dynamically. We were like, Ok, that's it, that's the one for the gig, we're not [messing] with it anymore!

8.6.4. 10th September 2021 (Interview with Jacob Hart).

AE: We hadn't explicitly planned to use an MLP before this. We took this invitation as an opportunity, and we actually started with the instruments. [Chris] said he wanted to do modular [...].

CK: I think you might have said that you wanted to do analogue processing.

AE: Oh yes.

CK: And I said I'd just built this modular in lockdown – let's do it!

AE: So that was a key thing – there was a commitment to analogue that was in this shared system. The only code is the controller.

8.7. *Richard Devine.*

8.7.1. 4th August 2021 (Interview with Jacob Hart).

RD: I spent probably four or five weeks on this patch preparing it, getting it just exactly the way I wanted it. To capture the right performance. It took many performances to get it right, and each time you do it it's not exactly going to be the same. So, when we did it on camera, I knew exactly the certain things I want to do. [...]

JH: Did you do many takes for the final performance?

RD: [...] I knew that we were going to be doing the recording and filming, and I was just like, whatever it is, is what it'll be. That's kind of the way you have to think about it with the modular system.

8.7.2. 4th August 2021 (Interview with Jacob Hart).

RD: [With the Morphagene] I could kind of create the evolution of how the piece was going to happen. So, all the sounds were moving across this timeline, and I can manipulate and play each one, and interject and perform each piece as it moves from left to right, interjecting each of these pieces of audio. So, for each reel, there were groups of sound which were all meant to be played with each other from the other Morphagene. I had matching sounds that would work at specific parts. It was very planned-out. If you listen to the piece, there are very deliberate parts [...] because I knew exactly what sound would be happening next. It gave me a great way to plan each gesture, each movement, each transition point.

It's something that's really hard to do when you're doing it with all synthesis. There was synthesis actually happening as well and other things that were going along to

fill the space, but because I have the ability to use the Morphogenes – they’re such powerful microsamplers – you can really organize things in a way where you can really perform things. I don’t have to worry as much about what’s coming next in the performance, I can concentrate more of the performance aspect.

8.7.3. *4th August 2021 (Interview with Jacob Hart).*

RD: With the Wogglebug version 2 it allows you to send a combination of a steady clock out, and a random burst generator out [...]. There’s a rate knob on the bottom lower left-hand side of the module which I love for live performances because I can drive things up to rates that are so fast that it’s at audio rate. Or, I could turn it counter-clockwise all the way down and stop the whole sequence completely. I have the ability to have a wide range of different tempos, or I can speed up or completely stop the whole piece. This allows me to create pauses and lets parts of the composition hang and stop and I can create tension. Instead of it being just this steady thing that’s happening all the time. I’ve spent many years trying to figure out how to develop ways of interrupting the clock. [...] I’m interested in breaking out of the steady pulse which so many modular people are accustomed to.

8.8. *Alexander Harker.*

8.8.1. *30th July 2021 (Interview with Jacob Hart).*

AH: The very starting idea of the piece was to explore the oboe from a timbral multiphonic point of view as an instrument that can create these sound complexes, or something other than notes. Early on in the process I didn’t know whether that would be the whole focus of the piece: there was the idea that there still could be melodic-type material with more concentration on line.

8.8.2. *30th July 2021 (Interview with Jacob Hart).*

AH: One of the important things – that was almost an accident – that developed as the score was developing: I started to think about not repeating multiphonics across the score lots of times. Because if that happens, the way that it figures out where [Dell] is in the piece gets more complicated. So, there are a few multiphonics which appear more than once, and they did create more problems. But basically, it knows for most multiphonics that it can say “well, we’re here, if I’m hearing that I must be here”. On top of that I then tell the computer what sections can go to what sections. So, if it’s like, “oh, you’ve jumped from the beginning of the piece to the end”, then the computer

knows that no, that's not possible. So, it'll just ignore any multiphonics that aren't within the rules.

8.8.3. 30th July 2021 (*Interview with Jacob Hart*).

AH: On the whole [the tracker] is pretty accurate. There were two issues that arose: one, embouchure and tuning become really important for the materials that I was dealing with and the fineness of the analysis. In the end [Dell] played on the reed that she recorded on which was always the most accurate. And she had to learn a little bit for certain multiphonics which direction to push the intonation to make sure the computer recognized it.

8.8.4. 30th July 2021 (*Interview with Jacob Hart*).

AH: [Concerning the partial tracker], there's two versions for the playback of the sine wave stuff, one of them has a filtered noise as well as a sine wave. You can crossfade between the two. So, you can have these gestures that come out that have these noise bursts emerging, these waves of noise. It's an individual partial turning from a very narrow thing into a noise thing.

8.8.5. 30th July 2021 (*Interview with Jacob Hart*).

AH: The granular stuff is fairly simple: a granular patch built in Framelib. The only addition, which is one that I've enjoyed for a long time, is that the voice has a filter. So, each grain has a filter, and that's randomized. But it's fairly simple. The parameters are just randomized between two points, there's no ability to control some of those things with gestures.

8.9. *Hans Tutschku*.

8.9.1. 27th August 2021 (*Interview with Jacob Hart*).

HT: Within this event, the network of the processes remains the same, but local parameters might change. Meaning, that while a module is still doing some processing, the parameters of the fine-tuning of that module might not be fixed but might move around. This moving around can be controlled globally by two different processes: one is what I call for myself, random guided spaces. If you think about one parameter, say the pitch of a transposition: I give that pitch a range, it has an upper and a lower limit; and then I give it a random choice between these two limits. If I bring these two limits to the same value – you don't have randomness anymore. So, for me, that's the way of saying that there is no randomness, the value is what it is.

Then I can open that up into both directions, only moving the lower boundary or the upper boundary etc. Those are boundaries that can shift over time. I can start out an event knowing exactly what the values are, and then over time – let's say a minute into that event – those boundaries might open-up. So, I know where I am more or less, but there is freedom of choice for the machine.

[...] The other one is by making those choices dependent on what comes into the microphone. [...] Looking at when an attack comes in, and that moment can then trigger something. [...] And the other process that I apply a lot in this is pitch tracking.

8.9.2. 27th August 2021 (*Interview with Jacob Hart*).

HT: We need to figure out how [...] to have a good combination of processes which uses the CPU in a manageable way where changing from one preset to another remains musical. I'm saying this because, if you just abruptly switch from one thing to the next, then we might not get into an interesting musical flow. And another thing that I've been working on over the years: I don't want that every time the pianist presses the sustain pedal to go from one event to the next – that each of those changes is perceived as a new musical section. So, there are very consciously built-in moments, certain events which just change one or two parameters. They make a slight and almost unperceivable slow migration to another musical fabric. And then there are moments where it's really a huge contrast.

[...] So, each module is completely autonomous. They're organized into different types of treatments. [...] They all have the ability to be soloed or muted. That, for me during the experimentation phase, [I can build-up musical fabrics], and then go in and solo one, mute one, almost like effects of a DAW.

8.9.3. 27th August 2021 (*Interview with Jacob Hart*).

HT: [Concerning spatialization] we learn that there is always a huge difference between what we hear when we compose the piece and what we hear when we are in the real hall. I'm using ambisonics here because that helps to broaden the sweet spot a bit. But me personally, I don't find compositional use for figures in the space or trajectories. Many electroacoustic composers – Jonathan Harvey being one of them – have composed trajectories, Stockhausen is also one. And they perceive them while they sit in the sweet spot. They have musical meaning, and they contribute to our understanding of the structure of the piece. Because it is something that is very fragile

in our perception – it breaks immediately when we are not sitting in the sweet spot – I don't rely on that. I think about space much more like a quality.

When I compose my spaces, I give them adjectives: this is a nervous space, this is an agitated space, this is a slow swinging space, this is a lop-sided space etc. I give them some kind of movement quality. Then, just by sitting in a hall and changing my positions during a rehearsal I can confirm that this particular spatial quality is still present even if I'm sitting closer to the speakers or further away. So, that's the way I think of space.

9. Multimedia Appendices Index

The following is an index of the digital elements found in the multimedia appendices.

9.1. *01_Code*

9.1.1. *01_Sources*

01_Burton

The annotated performance patch for Burton's piece.

02_Constanzo

The raw performance patch for Constanzo's piece.

03_Devine

The raw experimental patches made by Devine.

04_Eldridge

The raw performance code for Eldridge's piece.

05_Harker

The dependencies and the raw performance patch for Harker's piece.

06_Hayes

The annotated performance patch for Hayes' piece.

07_Pasquet

The annotated performance patch for Pasquet's piece and the raw generation patches.

08_Pluta

The dependencies and raw performance code for Pluta's piece.

09_Tutschku

The raw performance patch for Tutschku's piece and the raw generation patches.

10_Other

The raw and annotated Autechre patches used for testing.

9.1.2. *02_Tools*

01_Network_Construction

The Python tools used for transcribing Max patches to network format.

02_Fluid_Musicology

The Max and Python tools used for construction of 2D sound plots.

03_Burton_Score_transcribe

The Python tool used for transcribing Burton's scores into CSV files and audio slices.

04_Constanzo_Iterator

The Max tools used for iterating through the various effects modules from Constanzo's patches.

05_Sound_Plot_Navigation_Standalone

A standalone tool for Mac, Windows and Max for navigating the 2D sound plots used for research.

06_Segmentation_Explorer_Standalone

A standalone tool for Mac, Windows and Max for exploring the segmentations of the performances.

07_Burton_Video_Player

A Max tool for viewing the videos of Burton's alternative performances.

9.1.3. 03_Alternative_Performances_Setups

01_Burton

The Max setups for creating the alternative performances of Burton's piece.

02_Pasquet

The Max setups for creating the alternative performances of Pasquet's piece.

9.2. 02_Data

9.2.1. 01_Cytoscape_Networks

01_Global

The Cytoscape file and HTML interface for the global level network.

02_Local

The Cytoscape files and HTML interfaces for the two local level networks.

03_Instrumental

The Cytoscape files and HTML interfaces for the instrumental level networks of Constanzo, Hayes, Burton and Pasquet's pieces.

9.2.2. 02_Sound_Plots

The JSON files for constructing the 2D sound plots of Constanzo, Hayes, Burton and Pasquet's analyses.

9.2.3. 03_Performance_Segmentations

The JSON files for constructing the interactive performance segmentations of each of the pieces.

9.2.4. 04_Burton_Scores

The exported text format scores for each of burton's alternative performances.

9.3. 03_Audio

9.3.1. 01_Performance_Audio

Recordings of each of the performances.

9.3.2. 02_Sound_Plot_Audio

Audio for all of the 2D sound plots for Constanzo, Hayes, Burton and Pasquet's analyses.

9.3.3. 03_Pasquet_Alternative_Performances

Audio for each of the alternative performance configurations for Pasquet's piece.

9.4. 04_Video

9.4.1. 01_Demos

A collection of video demonstrations of various tools.

9.4.2. 02_Network_Visualisation

A collection of videos of various aspects of network visualization, notably networks over time and a dynamic network demonstration.

9.4.3. 03_Performances

A collection of videos of each of the concert performances.

9.4.4. 04_Burton_Alternative_Performances

Videos of each of the Burton alternative performance configurations.

9.5. 05_Images

9.5.1. 01_Burton_Alternative_Performances

A collection of images of the final score from each of Burton's alternative performances.