Supporting pro-amateur composers using digital audio workstations

Anna Clarke

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

2019

Abstract

This thesis investigates the activity of *pro-amateur* composers in order to identify possible design improvements to a category of composition software called *digital audio workstations*. Pro-amateur composers are composers who are not full-time professional musicians but who have a considerably greater level of expertise than amateurs. In contrast to the collaborative settings that this group is normally studied in this thesis will focus on situations where pro-amateur composers work independently.

Existing research on the use of composition software is reviewed, revealing that the composition process can involve a wide variety of component activities and overarching macro structures, and that other aids are often used in addition to composition software. Studies have also indicated that the design of composition software may constrain the creativity of composers. Four important considerations are identified for studying composers: triangulating multiple data capture methods, avoiding study designs that constrain what activities can be observed, capturing use of any external aids, and studying the use of a variety of composition software (or a prototype design) to mitigate any constraints that are due to the software's design.

Four pro-amateur composers were observed composing in their usual environments using a methodology based on *interaction analysis*. Based on information recorded about the settings, artefacts used, and activities carried out, three major patterns are observed. Firstly, existing tools support different composition activities to varying degrees, with additional support needed for improvisation, reflection, and auditioning incompletely specified material; secondly, composers make coordinated use of multiple representations; and finally, composers make use of strategies that enable selective allocation of time and effort (habituation, limited exploration, and self-constraining).

Previous authors have used many different notions of external cognition when studying the use of composition software. A literature review of such studies identifies techniques that can be applied to improve the representations used in composition software. Seven techniques are described: selective representation, diverse media types, structured representations, incomplete specification, representing alternatives, task lists, and representing history. A detailed review of evidence from the literature and the observational study is used to identify implementation suggestions for each technique. The technique of *task lists* has been studied significantly less in the literature on composition software and appears to be a fruitful avenue for further exploration.

A prototype to-do list website designed for coordinated use with *Ableton Live* is created to further investigate the *task lists* technique by studying how it is used by five pro-amateur composers. Using *thematic analysis* of interviews triangulated with video recordings and logs, four main themes are identified: using to-do lists to plan and focus, changing to-do list items over time, organising to-do lists, and applicability of to-do lists. Seven key patterns of activity that are enabled by task lists are also described: planning activity, journalling activity, interleaving activities, reflection, organising the to-do list, idea capture, and collaboration. Task lists appear to be useful because explicitly representing tasks, processes, and plans helps the composer to consider those subjects; and also because task lists ease many related activities, such as tracking incomplete work, monitoring deviation from a planned creative direction, or identifying and re-using useful strategies. Two important considerations for design of task lists in DAWs are identified: how task lists are integrated with the DAW, and how to increase the visibility of the composer's activity. For both considerations, specific suggestions are made for how these could be achieved.

Acknowledgements

The author is supported by the Horizon Centre for Doctoral Training at the University of Nottingham (RCUK Grant No. EP/G037574/1) and by the RCUK's Horizon Digital Economy Research Institute (RCUK Grant No. EP/G065802/1).

I would like to thank my supervisors, Chris Greenhalgh and Rob Houghton, for seven years of academic and personal support far above and beyond the call of duty. I'm also extremely grateful for the continued support provided by Steve Benford, Sarah Sharples, Emma Juggins, and the other staff at the Horizon Centre for Doctoral Training, and to the assistance I received from the other students in the CDT and the Mixed Reality Lab. A number of my University of Nottingham colleagues freely volunteered their time to discuss my research, including in particular Tim Coughlan, Sean McGrath, and Adrian Hazzard. I am also grateful for the many hours that Jenny Hackett, Nat Titman, Rob Moss, Rob Mitchelmore, and my sister spent together with me discussing ideas, and their invaluable help in figuring out what it was that I was trying to say. This research would not have been possible without the hospitality and generosity of my study participants, who I am indebted to. I would also like to thank Jérémie Garcia and Simon Dixon for kindly providing permission to reproduce in this thesis figures from their previously published work. Nottingham City Autism Service provided seemingly endless practical and emotional support to me, without which this project would never have never stood a chance of being completed. Vital support was also provided by Camilla Löwy, who always made time to listen, and whose enthusiasm during my internship helped enormously.

I'd like to thank my parents for all their support throughout my education and for giving me the start in life that allowed me to come this far. Finally, and most importantly, I must thank Aithne, Richard, and Vin for their incredible patience and love.

Contents

1 Introduction **15**

1.1 Digital audio workstations 15

1.1.1 Historical development 16

1.1.2 Interface design 17

1.1.3 Users 20

1.1.4 Defining the digital audio workstation 21

1.2 Research questions 23

1.3 Methodology 24

1.4 Contributions 26

1.5 Structure 27

2 Studies of composition with computers **29**

2.1 Populations of composers 29

2.2 Methodologies for studying composition 31

2.2.1 Study quality 33

2.2.2 Study settings 34

2.3 Themes 36

2.3.1 Macro structure of the composition process 36

2.3.2 Composition activities 38

2.3.3 External aids in composition 40

2.3.4 Interface issues in composition tools 42

2.4 Conclusion 43

3 Observing the use of existing DAW tools **45**

3.1 Methodology 46

3.2 Data 48

3.2.1 Interviews 48

3.2.2 Artefact-oriented records 49

3.2.3 Person-oriented records . . . 52
3.2.4 Setting-oriented records . . . 54
3.3 Discussion . . . 55
3.3.1 Varying support for composition activities . . . 55
3.3.2 Coordinating and combining multiple representations . . . 58
3.3.3 Selective allocation of time and effort . . . 60
3.4 Conclusions . . . 63

4 Representation design in DAWs **65**
4.1 Introduction . . . 65
4.1.1 Theoretical approaches . . . 67
4.2 Information represented in DAWs . . . 69
4.3 Techniques for designing representations in DAWs . . . 71
4.3.1 Selective representation . . . 71
4.3.2 Diverse media types . . . 75
4.3.3 Structured representations . . . 79
4.3.4 Incomplete specification . . . 83
4.3.5 Representing alternatives . . . 87
4.3.6 Task lists . . . 91
4.3.7 Representing history . . . 93
4.4 Summary . . . 98
4.5 Next steps . . . 101

5 Evaluating a to-do list for a DAW **103**
5.1 Introduction . . . 103
5.2 LiveTodos: interface and evolution . . . 107
5.3 Methodology . . . 109
5.3.1 Experimental protocol . . . 109
5.3.2 Data analysis . . . 112
5.3.3 Recruitment and participants . . . 113
5.4 Data . . . 115
5.4.1 Setting-oriented records . . . 115
5.4.2 Person-oriented records . . . 115
5.4.3 Object-oriented records . . . 117
5.4.4 Task-oriented records . . . 119

5.4.5 Interviews 121
5.5 Themes 122
5.5.1 Using to-do lists to plan and focus 122
5.5.2 Changing to-do list items over time 125
5.5.3 Organising to-do lists 128
5.5.4 Applicability of to-do lists 130
5.6 Activity patterns 132
5.6.1 Planning activity 133
5.6.2 Journalling activity 133
5.6.3 Interleaving activities 134
5.6.4 Reflection 135
5.6.5 Organising the to-do list 137
5.6.6 Idea capture 137
5.6.7 Collaboration 138
5.7 Discussion 139
5.7.1 Why to-do lists are useful in DAWs 139
5.7.2 Integrating to-do lists with the DAW 140
5.7.3 Increasing visibility of activity 142
5.8 Conclusions 143

6 Conclusion **145**
6.1 Summary of preceding chapters 145
6.2 Research questions 147
6.2.1 How do pro-amateur composers use DAWs? 148
6.2.2 How could the representations in DAWs be improved? . . 149
6.2.3 How do pro-amateur composers use to-do lists in concert with a DAW? 151
6.3 Summary of contributions 153
6.4 Critical reflection 154
6.5 Future work 156

7 Bibliography **159**

APPENDICES **179**

A Ableton Live **181**

B LiveTodos technical architecture **185**

C Written materials for studies **187**

List of Tables

2.1 Methods for gathering data on composers writing using music software . . . 31

2.2 Evaluation of study designs against the criteria from Collins (2005) 33

3.1 DAW functionality used by each composer in the study 52

3.2 Data on choosing options from libraries of options 53

4.1 Composition activities from Chapter 2 that are supported by the representation techniques . . . 98

4.2 Allocation strategies from Chapter 3 that are supported by the representation techniques . . . 98

4.3 Summary of the implementation suggestions for each design technique . . . 99

4.4 Composition software that has applied each design technique . . . 100

List of Figures

3.1 Use of artefacts by Composers A, B, and C 50

3.2 Paper artefacts demonstrated by composers A, C and D when asked about their use of paper 50

3.3 Inspirational material in the working environment of Composer D, including a tracklist of a previous album 54

3.4 Multiple screens in use by Composer C (left) and Composer B (right) 55

4.1 Viewing selected parameters of a track in *Ableton Live 9* 73

4.2 Modes of the loops library in *GarageBand 10* that use tags (left) and hierarchy (right) 82

4.3 'Comping' using grouped tracks in *Ableton Live 9* 88

4.4 Drawing multiple possible expressive curves with *Inksplorer*, taken from Garcia et al. (2011) 89

4.5 The 'Undo History' feature in *Sibelius 5* 94

4.6 Part of the *Sonic Zoom* interface, taken from Tubb and Dixon (2014b) 96

5.1 *Live's Locator* and *Info Text* features 105

5.2 Viewing a song (left) and editing a list item (right) in LiveTodos . . 108

5.3 Viewing a song (left) and adding a new list item (right) in the updated LiveTodos version 109

5.4 Equipment used in the study 110

5.5 Timeline of objects in use by each composer 116

5.6 The mode *Live* was used in by each composer 117

5.7 Number of to-do list items created by each composer 119

5.8 Timeline of LiveTodos activity for each composer 120

5.9 Timeline of LiveTodos activity for Composer 5 122

5.10 Detail section of timeline of objects used by Composer 4 123

5.11 Timeline of LiveTodos activity for Composer 4 123

5.12 Updates to to-do list items created by Composer 3 126

5.13 Updates to to-do list items created by Composer 2 126
5.14 Updates to to-do list items created by Composer 4 127
5.15 Detail of part of Composer 3's activity with significant group related events indicated 129
5.16 Number of tracks created in *Live* by each composer 131
5.17 Number of scenes created in *Live* by each composer 131

A.1 Live's *arrangement* mode 182
A.2 Live's *session* mode 182

B.1 Architecture diagram for LiveTodos 186

Lists of Design Suggestions

4.1 Relating to selective representation, based on composition software 74

4.2 Relating to selective representation, based on other software . . . 75

4.3 Relating to diverse media types, based on composition software . . 77

4.4 Relating to diverse media types, based on other software 78

4.5 Relating to structured representations, based on composition software 81

4.6 Relating to structured representations, based on other software . . 83

4.7 Relating to incomplete specification, based on composition software 85

4.8 Relating to incomplete specification, based on other software . . . 87

4.9 Relating to representing alternatives, based on composition software 89

4.10 Relating to representing alternatives, based on other software . . . 91

4.11 Relating to task lists, based on other software 93

4.12 Relating to representing history, based on composition software . 96

4.13 Relating to representing history, based on other software 97

5.1 Relating to integrating to-do lists with the DAW 141

5.2 Relating to increasing visibility of activity using to-do lists 143

Chapter 1

Introduction

This thesis investigates the activity of composers using a type of composition software called *digital audio workstations*.

This chapter will first focus on introducing **digital audio workstations** in more detail, providing an initial overview of the literature, in order to clarify the scope of the thesis. In particular, this thesis will focus specifically on the *design* of the user interfaces of digital audio workstations, considering their use by *pro-amateur* composers, while they are creating music *independently*.

Following sections will list the **research questions** that will be investigated, and summarise the **methodological approach**. Finally, the main **contributions** to the literature of this thesis will be described, along with the **structure** of the remaining chapters.

1.1 Digital audio workstations

Digital audio workstations (DAWs) are software packages designed to recreate and replace many of the technologies found in a traditional recording studio. The DAW provides functionality that previously required devices such as multitrack tape recorders and mixing desks (Bell, Hein, and Ratcliffe 2015), hardware synthesisers and effects, physical hardware racks, modular systems (Barlindhaug 2007), drum machines, and also devices introduced to the studio in the 1980s such as sampling keyboards and MIDI sequencing hardware.

The DAW is a *creativity support tool* (CST): software designed to support creative processes. More information and a more general introduction to CSTs can be found in Resnick et al. (2005) and Hewett et al. (2005), and an up-to-date review of CST research can be found in Frich et al. (2019). Many studies of CSTs have yielded results or proposed approaches that are relevant to DAW design and these will be discussed in later chapters. The research questions and findings of this thesis relate entirely to musicians using DAWs but may be useful to researchers studying related types of CST.

This section will introduce the DAW, describing in overview:

1. the **historical development** of DAW designs
2. the advantages and criticisms of the **design of DAW interfaces**
3. the impact of the technology on the **users of DAWs**
4. the **definition** of a DAW

1.1.1 Historical development

While a range of systems designed to support composition with computers had been created in academia in the 60s and 70s, such as MUSIC-N, Groove, and SSSP (Pennycook 1985), it was not until the 1980s that the direct ancestors of the modern DAW started to appear.

Digital hard disk recorders, initially used as a replacement for tape recording, evolved into software tools for editing an digital audio recording during the 1980s. By the mid-1980s, systems such as Studer MacMix (Yavelow 1989) had introduced a timeline view based on the *multitrack recorder* commonly used in studios, though it had to be used in combination with an expensive hardware audio workstation, limiting the potential audience. By the end of the decade, software such as Digidesign Sound Tools brought limited versions of this functionality to the mass market (Eldridge and Taylor 1990), complete with a software *mixing console* interface with faders. At this stage, systems still required a specialist sound card and were only capable of mixing four tracks of audio at once, but did permit useful features such as non-destructive editing.

MIDI sequencers (both hardware and software-based) were widely adopted by recording studios after their 1983 introduction (Théberge 2004), but early audio editors also provided only text-based interfaces. The *piano roll*, a graphical interface for editing MIDI note data, was included in a commercial MIDI sequencer in the late 1980s and adopted by competitors. After the graphical "Arrange" view adopting the multitrack recorder model was introduced in *Cubase* in 1989, it was rapidly adopted by competitors and became a standard interface (Bell, Hein, and Ratcliffe 2015).

Graphical programming systems such as Max (Puckette 2002) and Patchwork (Assayag et al. 1999), created in the late 1980s, allowed musicians to create their own synthesised sounds and audio effects. Again, these featured a new graphical interface, in this instance allowing users to connect "devices" by dragging "wires", similar to a modular synthesiser hardware.

Starting in 1990, these different approaches started to converge with the creation of integrated DAWs that allowed editing and playback of both audio and MIDI using a graphical interface (Bell, Hein, and Ratcliffe 2015). This brought together the multitrack recorder, mixing desk, and piano roll approaches to the graphical manipulation of music in one common system. This was followed by the widespread adoption beginning in 1996 of a standard interface, *VST* (Tanev and Božinovski 2014), that facilitated the development and use of emulated synthesisers and audio effects for DAWs.

By the late 1990s, home computers had become powerful enough to support the playback of multiple channels of high-quality audio simultaneously (Théberge 2004). Audio could be generated and manipulated by routing it through a *signal flow* network of simulated instruments and effects. The interfaces to control this network varied between DAWs, with some adopting a simulation of the "insert slots" and "sends" found in hardware mixing consoles (Gibson and Polfreman 2011), while others took the approach previously used by graphical programming systems of simulated wired connections (Walther-Hansen 2017).

1.1.2 Interface design

The previous section described how the interfaces of DAWs are heavily based on existing studio hardware, primarily:

1. multitrack tape recorders

2. mixing consoles
3. hardware synthesisers and effects
4. mechanisms for routing audio signals, either through wires or slots and sends associated with each mixing channel.

Even individual hardware components such as sliders and rotary knobs are recreated in DAWs (Duignan et al. 2004).

We can describe this as heavy use of a specific interface design strategy: *interface metaphors*. This strategy involves providing interfaces that have aspects of a familiar entity - in this case, existing studio hardware - but which also have their own behaviours and properties.

> "Interface metaphors have proven to be highly successful, providing users with a familiar orienting device and helping them understand and learn how to use a system. People find it easier to learn and talk about what they are doing at the computer interface in terms familiar to them" (Rogers, Sharp, and Preece 2011)

The DAW interface imitates the older studio hardware that it replaces both in appearance and function. The use of metaphors in this way suggests users would find a DAW and easier to use as a result. However, as Duignan et al. (2004) point out, "an important question that must be answered to validate the dependence on music hardware metaphors is: what proportion of new and potential users have prior experience with music hardware?" A clear example of this is the commonly used metaphor of the traditional analogue mixing console, something which an entire generation of native DAW users will think of as a "museum piece":

> "To the user unfamiliar with the design metaphor of the mixing console presented to them on screen, the console is not a representation of a mixing interface - it *is* the mixing interface" (Bell 2015).

Users of digital audio workstation may not have used a multitrack recorder or oscilloscope, but interfaces based on these technologies still appear in DAWs (Marrington 2016), and are even more unlikely to be familiar with the original piano roll (popular in the early 20th century) when using the piano roll metaphor to edit notes.

If DAWs are not using these metaphors because they will be familiar and intuitive to an audience of users of the original technology, what purpose *do* they serve instead? The metaphors in DAWs are unusual in nature in that they are extremely careful reproductions of the original (in contrast to, for example, the classic "file" metaphor which does not function like real paper files), sometimes incorporating details that actually make the interface more difficult to use. This suggests that composers desire these metaphors because of their "aesthetic characteristics" and have a desire to use "inaccessible iconic technology" (Barlindhaug 2007; Marrington 2016), and that DAWs adopt the aesthetics of older hardware in order to market their product to aspiring music stars.

Of course, other explanations could reasonably be suggested for the considerable similarity between DAW interfaces and older hardware. The common repertoire of metaphors used across different DAWs could make it easier for users to adopt a new DAW, or help composers move between the use of multiple different DAW systems in the course of a project. For composers who wish to use an external physical interface to control some aspects of the DAW, such as a specialised controller with knobs and faders rather than a mouse and keyboard, using a similar interface in the DAW provides a natural mapping from the physical controller to elements of the DAW.

Previous authors have been keen to emphasise that there are significant disadvantages to the common metaphors used in DAWs. The "mixing desk" metaphor has been particularly criticised and a number of alternative interfaces proposed (Gelineck, Büchert, and Andersen 2013; and in great detail in Mycroft 2018, chap. 2). The restrictions of the multitrack recorder metaphor (with a graphical display of tracks as rows) are discussed by Gohlke et al. (2010), and the metaphorical approaches to signal flow can also have disadvantages (Barlindhaug 2007; Gibson and Polfreman 2011; Myllys 2014). In particular, a key concern of authors is that of technological determinism: does the use of a DAW encourage the composers to create music in particular ways? To contribute towards the understanding of this question, this thesis will focus on how composers make use of the interfaces provided in DAWs, and how design changes might influence the activity of composers.

The following section will focus on who these composers are.

1.1.3 Users

Traditionally in a recording studio work was carried by a team of professionals with specialised roles, including songwriters, producers, studio musicians, and sound engineers (Leyshon 2001), with sound engineers being further specialised into recording engineers, sound editors, mixing engineers, and mastering engineers (Pras, Guastavino, and Lavoie 2013). This separation into different roles was so marked that it could extend to confining different personnel to different spaces within the studio (Bell 2014).

However, the advent of the DAW created software packages that support a complete integrated process of music creation: one that enables a single person to carry out the tasks associated with all of the traditional studio roles (Méndez 2015). These specialised roles have disappeared within the studio environment in favour of 'multiskilled' professionals (Pras, Guastavino, and Lavoie 2013), and songwriting and production skills now overlap in many genres of music, a change driven primarily by the ease-of-use and low costs of digital tools (Bennett 2011). As a result of this, as Bell, Hein, and Ratcliffe (2015) points out, "in the digital era, performance, recording, and composition have largely collapsed into a single act."

Hracs (2012) observed that:

> "with professional and even consumer software, recording, editing, mixing, and mastering digitally recorded music has become easy enough for a much larger number of musicians to do on their own… this constitutes a fundamental shift in the working lives of musicians who, under the major label model of music production, allocated the majority of their time to performing creative tasks such as song writing, recording, and performing".

This change has supported the growth and development of a specific class of musicians: one distinct both from novices and professionals. Armed with the digital audio workstation, these composers now are capable of creating polished sounding music without the assistance of specialised professionals at a commercial recording studio. These composers have been referred to variously as prosumers (Cole 2011), recordists (Merrill 2010), hobbyists (Kaloterakis 2013), and pro-amateurs (Méndez 2015), the latter terminology being adopted in this thesis. Pro-amateurs may be engaged in full-time education or work in another area, but are still dedicated and passionate about their composition work.

Pro-amateurs have different practices and experience different challenges to other categories of composers (Hoare et al. 2014). They may create their own working methods or adopt those of friends and other pro-amateur collaborators, in contrast to studio professionals who have "learned their trade by shadowing experienced veterans of the craft" (Bell 2014). As Marrington (2016) points out, "this 'incompetence' where such specialist skills are concerned is… actually advantageous in discovering new creative possibilities". As a result of their novel and distinct characteristics, this group has generated academic interest.

This thesis will focus on pro-amateur composers using studies of their activity to investigate DAW design. Findings that relate to pro-amateur composers may also be relevant to other groups of composers and this will be further discussed in Chapter 6.

Existing studies of pro-amateurs typically focus on collaborative settings in which composers, musicians and/or recording engineers are working together (McGrath, Chamberlain, and Benford 2016; Nabavian 2010; McGarry et al. 2017). By contrast, this thesis will focus on the work of individual composers who work in isolation and are responsible for all aspects of a composition.

In order to study the composers using DAWs, it is important to first define what a DAW *is*. The following section will discuss this question.

1.1.4 Defining the digital audio workstation

Several authors have previously attempted to define the DAW. Duignan, Noble, and Biddle (2005) locate digital audio workstations within a taxonomy of composition tools used for sequencing music events:

1. textual language music tools (such as MUSIC-N, CLM, CSound, and Supercollider)
2. sample and loop triggers (such as MPCs and programmable drum machines)
3. visual programming tools (such as Max, PureData, and OpenMusic)
4. linear sequencers (which includes the subcategory of digital audio workstations)

Zadel and Scavone (2006) echo these distinctions in their description of software for performance control, drawing a distinction between timeline-oriented software (i.e. linear sequencers, including digital audio workstations) and procedural software (textual and visual programming tools). Composition software has also been divided into three broader categories (see Wilkie 2014, 10): music notation software (which may also be referred to as score-writing software); music generation and performance software; and music sequencing, recording, and production software (a category containing digital audio workstations).

It is important to note that while the DAWs of the 1990s might have easily been placed within one category of these models, the modern DAW does not fit as easily. The description by Duignan, Noble, and Biddle (2005) of a linear sequencer in which "all events must be placed at an absolute time location" and "playback proceeds linearly from start to end" do not describe modern DAWs which may also provide non-linear alternatives to the timeline. The distinction drawn by Wilkie (2014) between notation software, performance software, and sequencing software has also been blurred by digital audio workstations that offer support for notation and live performances. Some digital audio workstations have integrated support for textual or visual programming systems.

An alternative definition of the DAW is in terms of the *material* that it manipulates. (Marrington 2016) has taken this approach, stating that "A DAW's essential functionality… is to allow for the manipulation of two main forms of information, MIDI data and digital audio." This definition has the advantage of succinctness, but fails to distinguish the DAW from other tools (for example, music visual programming tools), and ignores the importance of the support other mediums in the DAW (such as video for composers scoring films, and OSC or CV envelopes for composers working with additional hardware and software tools). When considering the DAW in terms of the materials it manipulates even the final output of a DAW, which might originally have been assumed to be a single stereo audio file, has now broadened. DAWs may now be used to create a set of multitrack stems, a live performance, multiple mixes for different surround sound setups, or a library of pre-prepared loops for use by other composers.

It is also important to consider that even describing the DAW as a software package may also be misleading. DAWs are frequently used in combination with hardware devices (Nash 2012, 20), and it is increasingly common for DAW manufacturers to sell dedicated hardware designed for use in combination the software (e.g. Native Instruments' *Maschine* or Ableton's *Push*), with the hardware becoming the primary user interface for the combined system.

Instead, the most useful definition of a DAW may instead be one that considers the capabilities the DAW provides. Bell (2015) describes the DAW as "a categorisation that has evolved over time, with most sharing common capabilities including sequencing, recording, and mixing music, with support for software synthesis". For the purposes of this thesis, we will focus on considering the "archetypal" DAW capabilities: those common capabilities seen across many existing examples, and are likely to continue to be found in DAWs of the future as technology evolves:

1. arranging audio and MIDI data ("clips") within tracks on a graphical timeline
2. virtual instruments and effects whose parameters can be automated on the timeline
3. control of a signal flow network associated with the tracks, instruments, and effects
4. control of volume and stereo panning of tracks
5. playback and recording of audio and MIDI material
6. output of a final audio waveform "mixdown" that is the resulting composition.

While this thesis must adopt a clear definition of what a "DAW" is, in order to limit the scope of what will be studied, it is important to note that this also artificially constrains the corresponding definition of what a "composition created with a DAW" is.

1.2 Research questions

This thesis addresses the following questions:

1. *How do pro-amateur composers use DAWs?*
    - How do these composers work?
    - How is their behaviour similar to and different from composition as described in the literature?
    - What facilities of the DAW do they use?
    - What additional facilities might be useful to them?

2. *How could the representations in DAWs be improved?*
    - What are the problems with current DAW interfaces and representations?
    - What solutions have already been identified in the literature?
    - How do these solutions relate to the work of pro-amateur composers?
    - What are the "low-hanging fruit" for improving DAW designs?

By the conclusion of Chapter 4, using *task lists* in DAWs is found to be both underexplored in the existing literature and also potentially useful for the behaviours seen as important to pro-amateur composers. This leads to a third research question:

3. *How could pro-amateur composers use task lists in concert with a DAW?*
    - How are task lists useful to pro-amateur composers working with a DAW?
    - What activities do pro-amateur composers use task lists to carry out?
    - What further improvements to the representations in DAWs would support these activities?

1.3 Methodology

A study of pro-amateur composers using digital audio workstations at the present moment must navigate a significant disparity in the amount of research: there is a significant body of literature from many years of research of composers using DAWs, but there is very little research specifically on pro-amateurs composing with DAWs and other composition software (see Chapter 2).

As a result, this area requires research that might seem contradictory: both research that is informed by the existing research about *known* important issues for composers, but also *exploratory* research to that uses a grounded approach to determine important issues for pro-amateurs.

In order to resolve this tension, this thesis will take the following approach:

1. initially, a study specifically of the behaviour of *pro-amateurs* to find important unmet requirements for this group;
2. secondly, a detailed literature review of research into DAWs to relate the specific requirements of pro-amateurs found in the study to the existing

work into composition *more generally*;

3. finally, based on both the previous study and the literature review, a previously under-explored but also relevant area is identified ("task lists"), and a study is carried out using a prototype system designed to investigate this area.

The following chapters will describe this approach in more detail.

Some important methodological issues should also be discussed at this point.

Defining creativity. Musical composition is a *creative* activity. Other authors studying composition software (Coughlan 2009, chap. 2; Nash 2012, chap. 3) have provided detailed overviews of the many definitions and theories of creativity that have been proposed, and investigated how the use of software tools impacts the creativity of composers. In this thesis, there will be no attempt made to define "creativity" as a concept or to suggest ways that it might be enhanced or increased. The thesis instead focuses on identifying specific externally observable processes that composers make use of when creating music, how well (or poorly) those processes are supported by the existing software they use, and design changes that might better support those processes.

Unique composition practices. Musical composition is a process that is necessarily idiosyncratic: there is no single way to compose, and composers seek to work in novel ways both for their own enjoyment and to ensure novelty in their creative output. There are also an enormous set of genres each with their own conventions of processes, inputs, and outputs. As a result, this thesis cannot produce a single universal set of design requirements for pro-amateur composers, but instead uses the findings from nine pro-amateur composers to suggest design directions to consider.

Author's experience. Musical composition is an area which requires highly specialised knowledge, and the issues facing pro-amateur composers using DAWs are another specific sub-field within that area. The analysis of composer's activity in this thesis is therefore based on the author's knowledge of and experiences of performing and recording as part of pro-amateur bands using DAWs, along with many years of training in playing classical instruments. The author also has experience of using tracker and notation software in past projects. While the findings from this thesis are based on the research of other composers the conclusions drawn about their behaviour will unavoidably be shaped by the viewpoint of the author.

1.4 Contributions

The main contributions of this thesis are:

1. A *literature review* that identifies the strengths and weaknesses of existing studies in this area. The review identifies what methods have been used to study which groups, what groups have been studied, and what kinds of patterns can be seen. Findings from this review will be of most use to researchers designing studies of composers using software tools.

2. A *fine-grained account* of the activity of four pro-amateur composers in a naturalistic setting, using an *interaction analysis* approach to generate information about the activity from setting, person, object and task orientated viewpoints.

3. A *catalogue of techniques* for designing DAWs in order to support the processes and activities that are used by pro-amateur composers. Specific recommendations are given for how to implement these techniques based on the literature review and study. The techniques that are most ripe for further investigation are also identified. The catalogue of techniques will be most useful to designers of research prototypes or commercial products that are used with a DAW or as part of a DAW.

4. A *fine-grained account* of the activity of five pro-amateur composers using a prototype *task list* system in a lab setting, with *interaction analysis* and *thematic analysis* used as complementary analysis techniques to identify important themes, activity patterns enabled by task lists, and further design recommendations.

Each contribution is described more fully in chapters 2, 3, 4, and 5 respectively.

Findings from this thesis may also be useful for people working in related fields within creativity support tools. In particular, tools that are primarily based around a timeline (such as video editing tools) may benefit from applying the techniques identified in this thesis.

1.5 Structure

Chapter two, *Studies of composition with computers*, reviews the existing literature around composition using computer tools.

The first part describes which *populations* of composers have been studied, what *methodologies* have been used to study them, and what *settings* studies have been carried out in. The second part reviews the *findings* of these studies, discussing the macro structure of the composition process, the component activities of composition, problems in existing tools, and use of external aids such as paper notes. Finally, based on the observations made in the review, some important study design considerations for this research area are summarised.

Chapter three, *Observing the use of existing DAW tools*, details the design and results of a study based on the criteria from the previous chapter. Four pro-amateur composers are studied in their normal working environments using an *interaction analysis* methodology. Consistent with this methodology, *setting-oriented*, *person-oriented*, *object-oriented*, and *task-oriented* data were gathered, along with biographical information gathered through semi-structured interviews.

From this analysis, three patterns are described: varying levels of support for the different activities involved in composition, coordinating and combining multiple representations, and selective allocation of time and effort.

By relating the findings from this chapter to those in the previous chapter, the focus for the remainder of the thesis is then identified: the use of external representations by composers, particularly to support the activities and strategies used by pro-amateur composers.

Chapter four, *Representation design in DAWs*, reviews existing research relating to the use of external representations in composition software to identify techniques for designing DAWs.

A number of techniques are identified which can help support composers: selective representation, diverse media types, structured representations, incomplete specification, representing alternatives, task lists, and representing history.

For each technique, specific suggestions are made of potential ways to implement these techniques that have not yet been explored in digital audio workstations. Based on the lack of existing research and the potential relevance to pro-amateur users, the *task lists* technique is identified as a good candidate for further investigation.

Chapter five, *Evaluating a to-do list for a DAW*, describes a study to gain more insight into the *task lists* technique identified in the previous chapter. The *task lists* technique is implemented in a *prototype* tool that can be used with an existing digital audio workstation, *Ableton Live*.

By analysing video recordings and log data from a composition task and semi-structured interviews with the composers, four common themes are identified, and seven patterns of activity that task lists support are described. Based on these findings, general conclusions are drawn regarding why task lists are useful to composers, and two important considerations for designing task lists in DAWs are highlighted.

Chapter six, *Conclusion*, summarises the findings and contributions of this work. Following this, answers to the research questions are each discussed in turn, and the contributions summarised. The quality of the research in this thesis is evaluated through a critical reflection. Finally, suggestions for future work in this area are proposed.

Chapter 2

Studies of composition with computers

This thesis seeks to first understand the existing working practices of music composers working with digital audio workstations. An initial step along this path was a review of the literature of studies of composers using composition software. The results of this review will inform the design of the studies of pro-amateur composers that will be described in Chapter 3 and Chapter 5.

The first part of this chapter describes the **populations** of composers that have been studied and which **methodologies** have been used to study them.

The second part of this chapter summarises the findings of the surveyed studies, focusing upon findings relating to four key themes found in the literature:

1. what is the **structure** of the composition process?
2. what **activities** are involved in composition?
3. what **problems** exist with composition tools?
4. what **external aids** are used by composers to work around these problems?

Finally, some important study design considerations for this research area are summarised.

2.1 Populations of composers

Composers may work in different musical genres and social contexts and may have differing levels of expertise and training. The studies in this review used participants who could be described as being in one of several populations:

Amateur composers may have limited training and experience or none at all. This group is very infrequently studied.

Children are more frequently studied, usually in the context of the literature on music education, and some of the earliest studies of composition using computer tools are of children (Folkestad, Hargreaves, and Lindström 1998; Seddon and O'Neill 2003; Farbood, Pasztor, and Jennings 2004). Some studies in educational settings, like that of adolescents in Tobias (2013), are of subjects who are old enough to use professional tools and provide results that may be applicable to composition behaviour in general.

Academic composers[1] are students or staff of a university music department and usually have significant formal musical training. They were the most common population used in the studies that were reviewed, and almost all studies published before 2008 were of this group. Most studies of academic composers use a case study approach with only a few composers being included in each study.

Professional composers may write soundtracks for film, music, and games and have significant practical experience. With the notable exception of Collins (2005), studies of professional composers using DAWs were only found after 2008 in this review. Professional composers are often studied working alone, with some exceptions (Auvinen 2019). Some studies are able to recruit larger numbers of composers by speculatively contacting record labels (Gelineck and Serafin 2009) or by using existing personal contacts in the music industry (Duignan 2008; Morey 2017; Auvinen 2019). Roels (2016) has reviewed existing studies of composition by professional composers.

Pro-amateur composers are a group who have significant expertise marking them as clearly distinct from amateurs, but who are not in full time professional musical work, distinguishing them from professional composers. A number of studies of pro-amateur composers have been published since 2015 (Koszolko 2015; Méndez 2015), and in some studies, the terminology "pro-amateur" is explicitly used (McGarry et al. 2017; McGrath, Chamberlain, and Benford 2016). All studies of pro-amateur composers found in this review were of collaborative work. Some studies recruit

[1]As not all studies describe their study participants clearly, any composers described as "contemporary composers", "electroacoustic composers", or users of tools mostly used in academia (e.g. *Max*, *OpenMusic*) were categorised as "academic composers" for the purposes of this review.

composers of a mix of abilities online, resulting in many pro-amateur composers being sampled, and this more indiscriminate approach to recruitment permits studies on a much larger scale of hundreds of composers (Méndez 2015; Etinger and Orehovački 2017) in some cases.

2.2 Methodologies for studying composition

A variety of methods have been used for gathering data on composition using computers: these are summarised in Table 2.1. These have included interviews and discussions; observation and recording of video, audio, screen capture, and screenshots; data logging and analysis of saved data; participatory observation and design; questionnaires; and diary studies. When the composition tool studied is a notation tool, rather than a DAW, expert score analysis is sometimes also used.

Table 2.1: Methods for gathering data on composers writing using music software

Methods	Studies
Audio recording	Amitani and Hori (2002), Bainbridge, Novak, and Cunningham (2010), McGarry et al. (2017), McGrath, Chamberlain, and Benford (2016), Brooker and Sharrock (2016)
Diary study	Bainbridge, Novak, and Cunningham (2010)
Datalogging	Collins (2005), McCulloch (2014), Nash (2012), Jillings and Stables (2017)
Discussions	Coughlan and Johnson (2006), Donin and Theureau (2007), McGrath, Chamberlain, and Benford (2016), Smith (2011), Thiebaut (2010)
Interview	Auvinen (2019), Bainbridge, Novak, and Cunningham (2010), Bell (2014), Bertelsen, Breinbjerg, and Pold (2009), Collins (2005), Coughlan and Johnson (2006), Donin and Theureau (2007), Duignan, Noble, and Biddle (2010), Eaglestone et al. (2001), Folkestad, Hargreaves, and Lindström (1998), Garcia et al. (2011), Gelineck and Serafin (2009), Healey and Thiebaut (2007), Marrington (2010), Morey (2017), Polfreman (1997), Tsandilas, Letondal, and Mackay (2009), Tobias (2013)
Observation	Auvinen (2019), Coughlan and Johnson (2006), Duignan, Noble, and Biddle (2010), Eaglestone et al. (2001), Eaglestone et al. (2007), Folkestad, Hargreaves, and Lindström (1998), Marrington (2010), McGarry et al. (2017), McGrath, Chamberlain, and Benford (2016), Polfreman (1997), Tobias (2013), Farbood, Pasztor, and Jennings (2004)
Online video	Macchiusi (2017), Nash (2012)
Online forums	Macchiusi (2017)
Participatory design	Garcia et al. (2011), Tsandilas, Letondal, and Mackay (2009)
Participatory observation	Koszolko (2015), Brooker and Sharrock (2016)

Methods	Studies
Questionairre	Coughlan and Johnson (2006), Etinger (2016), Gelineck and Serafin (2009), Méndez (2015), Nash (2012), Peterson (2008), Polfreman (1997), Etinger and Orehovački (2017)
Saved data analysis	Bainbridge, Novak, and Cunningham (2010), Collins (2005), Folkestad, Hargreaves, and Lindström (1998), Jennings (2006), Jillings and Stables (2017)
Score analysis	Donin and Theureau (2007), Healey and Thiebaut (2007), McCulloch (2014), Peterson (2008), Smith (2011)
Screen recording	Collins (2005), Donin and Theureau (2007), Eaglestone et al. (2001), Eaglestone et al. (2007), Peterson and Schubert (2007), Peterson (2008), Brooker and Sharrock (2016), Tobias (2013)
Software and customisation analysis	Bertelsen, Breinbjerg, and Pold (2009), Donin and Theureau (2007), Duignan, Noble, and Biddle (2010)
Verbal protocol	Collins (2005), Eaglestone et al. (2007)
Video recording	Amitani and Hori (2002), Auvinen (2019), Bell (2014), Eaglestone et al. (2001), Eaglestone et al. (2007), Peterson (2008), Polfreman (1997), Thiebaut (2010), Brooker and Sharrock (2016), Tobias (2013)

Early studies identified a need for triangulating the results of multiple data collection methodologies (Eaglestone et al. 2001) and almost all studies that were reviewed have taken this approach. This has also been identified as good practice for studies of creativity support tools in general (Hewett et al. 2005). A typical approach used by the majority of reviewed studies is to combine data from discussions or interviews with at least one other method of collecting data. However, some recent studies have now begun to use one method for gathering data without triangulating against another source of data (Etinger 2016; Etinger and Orehovački 2017; Méndez 2015), presumably due to the unprecedented sample sizes of hundreds of composers that this approach permits.

Some authors have given advice on the choice of data collection methods. Bell (2014) and Donin and Theureau (2007) suggest that for more longitudinal studies, it can be useful to enhance interviews and discussions by using a *stimulated recall* approach, in which interviews are carried out while composers examine the physical artefacts they created during the composition process. McCulloch (2014) critique using video/screen capture without also analysing saved data or data logged data, as examining video alone may cause some aspects of composition to be missed, particularly those related to improvisation.

While many studies of composition look at composer's existing behaviour, another common approach is to examine the composer's reactions to a *prototype* composition system (Amitani and Hori 2002; Bainbridge, Novak, and Cunningham 2010; McCulloch 2014) or to a prototype tool that provides an *additional* interface for an existing system (Garcia et al. 2011; Tsandilas, Letondal, and Mackay 2009). More recently, a simple DAW that runs in the web browser and permits data logging of user activity has also been used in Jillings and Stables (2017).

2.2.1 *Study quality*

Collins (2005) makes five specific criticisms of the literature available at the time. For studies which involve observing composers working, these can be used as criteria to evaluate their study design:

1. *longitudinal*: does the study capture larger-scale creative processes that occur over a significant period of time?
2. *microscale*: are finer-grained phenomena described in addition to macro-scale activity?
3. *tools*: are the composers provided with compositional tools of an appropriate quality?
4. *naturalistic*: does the work studied involve realistic tasks carried out in a "real-world" setting?
5. *polyphonic*: are composers asked to write music for more than just a solo instrument?

For the studies of solo composition for which these criteria were applicable, their performance is listed in Table 2.2.

Table 2.2: Evaluation of study designs against the criteria from Collins (2005)

Study	Longitudinal	Microscale	Tools	Naturalistic	Polyphonic
Amitani and Hori (2002)		✓	✓		✓
Bainbridge, Novak, and Cunningham (2010)			✓	✓	✓
Bell (2014)	✓	✓	✓	✓	✓
Collins (2005)	✓	✓	✓	✓	✓
Donin and Theureau (2007)	✓		✓	✓	✓

Study	Longitudinal	Microscale	Tools	Naturalistic	Polyphonic
Duignan, Noble, and Biddle (2010)		✓	✓	✓	✓
Eaglestone et al. (2007)	✓	✓	✓	✓	✓
Jennings (2006)					✓
Jillings and Stables (2017)		✓	✓		✓
Marrington (2010)	✓		✓		✓
McCulloch (2014)		✓	✓		✓
Nash (2012)			✓	✓	✓
Peterson and Schubert (2007)		✓	✓		✓
Peterson (2008)		✓	✓		✓
Polfreman (1997)		✓	✓	✓	✓
Thiebaut (2010)		✓	✓		

Only authors who chose to restrict their design to a case study of one or two composers were able to meet all five criteria, which suggests it may be difficult to ensure all criteria are met in studies with a larger number of participants. It appears that ensuring tasks and settings are naturalistic and that composers are studied over an extended time period are particularly challenging goals to meet. In this review, a similar pattern was also observed in studies of collaborative composition.

Study designs need to consider the benefits and disadvantages of meeting the criteria against a larger sample size. Larger sample sizes may give a broader picture of composition practices, giving results more generalisable to other composers, but this may come at the expense of being unable to observe some types of composer behaviour.

2.2.2 *Study settings*

Studies of composition with computer tools took place in one of two settings until recently:

Field studies take place in the normal environment used by the composer to work. As a result, they are often useful for studies that are highly naturalistic. Methodologies that would be simple to carry out in a field setting are observation, verbal protocol, and capturing a record of electronic and paper representations (for example, by photographs or asking composers to regularly save their work). Examples of field studies include the observations of university music production society jam sessions in Coughlan and Johnson (2006) and of composers working in their home or workplace studios in Eaglestone et al. (2007).

Lab studies have additional advantages and disadvantages while permitting the same methodologies as field settings. An advantage of a lab study is that no disruptive installation of new software or hardware in the composers working environment is necessary; a disadvantage of a lab study is that it requires the composers to travel to the lab, which would reduce the number of potential participants to only those nearby. Some lab studies attempt to increase realism by asking composers to bring their own equipment with them (Garcia et al. 2011). The disadvantages of lab studies appear to be discouraging research in this setting, with a rare example of this approach from the 2010s being McCulloch (2014).

A third approach has also recently emerged. **Online studies** involve gathering data from remote composers via the internet. Jennings (2006) analysed hundreds of children's compositions made in tools *Hyperscore* and *DrumSteps* that had been uploaded to an online gallery. Studies of adult composers were seen in the following decade, such as Nash (2012). Online studies may involve gathering data that already exists on online video sites (Nash 2012) or discussion forums (Macchiusi 2017), asking composers to respond to questionnaires (Méndez 2015; Etinger 2016; Etinger and Orehovački 2017), or data logging of the activity of composers (Nash 2012; Jillings and Stables 2017).

Nash (2012) identifies significant risks that are involved with investigating the use of a new piece of software in an online study: recruiting enough participants to be useful may be difficult, there may be a conflict of interest between producing a tool that is desirable to participants and producing a tool that meets the research objectives, and the data gathered may be difficult to analyse and require custom software to be developed. Both Nash (2012) and McCulloch (2014) were required to create such custom software to analyse their logged data.

2.3 Themes

Four themes commonly appear within the findings of studies in this review: descriptions of the *macro structure* of the composition process (for example, dividing composition up into various stages of activity), descriptions of the *activities* involved in the composition process, details of the *external aids* that are used by composers (such as paper notes), and *interface issues* that have been observed in the composition software used.

Each theme is described in more detail in each of the following sections.

2.3.1 Macro structure of the composition process

In the literature, it is possible to see a division between *staged* and *iterative* descriptions of the macro structure of the composition process over time. Often authors will also describe macro-level processes by drawing a division between *top-down* and *bottom-up* composition, and between *chronological* and *patchwork* composition. Findings relevant to each style of macro structure will now be summarised.

Staged models of composition are those in which a set of "phases" or "stages" are *sequentially* followed.

Some authors describe phases which focus on different levels of hierarchy, with Bainbridge, Novak, and Cunningham (2010) describing "ideation" (fragment generation) and fragment combination phases, and Eaglestone et al. (2001) reporting a "divergent" phase of creating short sound fragments followed by a second phase of sensemaking and recombination.

Other authors divide compositional stages up by the aspects of music that are focused on. Peterson (2008) describes how some composers work features stages where different aspects of the composition (e.g. notes, dynamics) are focused on, as do some of the composers in McCulloch (2014), but both note that this behaviour is not universal.

Nash (2012) divides composition sessions (rather than the entire composition process) into stages based on the amount of time that users spend playing back their composition, using quantitative data to suggest a three-stage model of "preparation", "creative editing", and "evaluation".

A commonly described stage is initial preparatory work carried out before starting work with a DAW (Auvinen 2019; Bell 2014; Donin and Theureau 2007; Tobias 2013). Some composers make use of paper notes, hand-written scores, or notation software for this stage (Healey and Thiebaut 2007; Marrington 2010; Tsandilas, Letondal, and Mackay 2009).

Iterative models of composition are those in which a small number of phases of activity are *repeatedly* iterated through. Gelineck and Serafin (2009) describe exploratory, editing, and pragmatic phases; Garcia et al. (2011) expressing, exploring, and execution; Coughlan and Johnson (2006) ideation and evaluation; and Collins (2005) a cycle between reformulating goals and exploring a solution space. Some studies describe similar iterative patterns of expression and editing on the *micro* scale during improvisation (Bell 2014, 301; McCulloch 2014, 135).

Top-down and bottom-up styles of composition are terms used to describe the macro structure of how composers move between high-level and low-level aspects of the composition. Polfreman (1997) reports some composers first attending to high-level structural aspects, while other composers prioritise low-level aspects such as assembling musical motifs or fragments, and use the terms "top-down" and bottom-up" to describe these approaches. Eaglestone et al. (2008) find a division between similar groups they instead label "refiners" and "synthesisers", linking this to the idea of *cognitive styles* and relating this to the "analytic/global" axis of cognitive styles. Polfreman (1997) notes that some users display both behaviours in parallel, concerning themselves simultaneously with global structure and with specific details. McCulloch (2014) reports quantitative data that indicates that the composers they studied "lean" toward using a "heuristic" approach (low-level first) which might suggest a bias created by the custom composition software they were testing or the study design.

Chronological and patchwork styles of composition are used to describe the way that the temporal arrangement of compositing activity relates to the timeline within the composition itself. *Chronological* or *linear* composition describes composition activity starting at the beginning of the work and proceeding to the end, whereas *patchwork* composition processes create parts of the work out of order. Amitani and Hori (2002) claim that using conventional timeline-based composition packages leads to chronological composition and that alternative designs without a conventional timeline view are needed to change this behaviour. However, other authors specifically observe that composers do not work linearly (Coughlan and Johnson 2006) or that they have seen a mixture of linear and non-linear approaches (Duignan, Noble, and Biddle 2010; McCulloch 2014). McCulloch (2014) notes that for some composers "the development of the form at the global level is non-linear, the ordering of the composer's activity at the keyboard is still fairly sequential". Extensive forward planning can also enable more unusual orderings of the composition process including creating a composition one track at a time (Bell 2014).

It appears that the main conclusion that can be drawn from twenty years of studies, using a wide variety of study designs and methodologies, is that the macro structure of composition *varies* between composers. This could be possibly be explained in several ways. Particular styles of composition activity might be culturally transmitted within communities of composers, or might be deliberately chosen by a composer to match their particular strengths and weaknesses. The structure of composition seen in a particular composer might also reflect the constraints that they work under including deadlines, academic requirements for novelty, or a need for the output to be marketable. A need for novel creative outputs even may lead some composers to deliberately choose novel structures for their creative processes. Finally, the macro structure may be influenced by the composition software being used, an effect which will be discussed further in the later section *Interface issues in composition tools*.

2.3.2 Composition activities

Several authors have attempted to create taxonomies of activities or tasks involved in composition. Amitani and Hori (2002) provide a taxonomy of "unit cognition processes", Polfreman (1997) creates a "task model" in which different activities are catalogued, and Nash (2012) provides a "flow model", while Coughlan and Johnson (2006) merely provide a list of four different main activities. However, all four of these taxonomies contain several component activities not contained in any other taxonomy, suggesting that they all paint an incomplete picture of composition. Many authors deliberately focus on a particular group of activities to avoid this issue.

Some common activities found in this review include:

- *Planning*: making notes about a work on which the main compositional process has not yet started. See also the previous section on macro structure, where a separate phase of research and planning is often suggested.
- *Collecting and capturing material*: collecting of sounds, samples, and musical ideas as an initial planning activity (Polfreman 1997; Coughlan and Johnson 2006). This activity is studied in detail in Bainbridge, Novak, and Cunningham (2010).
- *Auditioning*: reflective listening to the composition in progress. This may allow the composer to evaluate the part of the composition they are working on in the context of the whole or neighbouring sections (Polfreman 1997).
- *Reflection*: reflective activity on composition in progress *without* playing it back. This actually describes a broad category of related activities, including re-reading existing plans and notes (Donin and Theureau 2007), restructuring of a composition (Amitani and Hori 2002), pausing to think before attempting another iteration of improvisation (Bell 2014), or listening to other compositions for comparison or inspiration (Auvinen 2019).
- *Improvisation*: spontaneously creating new musical ideas at high speed (McCulloch 2014; Duignan, Noble, and Biddle 2010). Some authors may use the more colloquial term "jamming" to describe this activity (Koszolko 2015), particularly when studying composition in collaborative settings.
- *Mixing*: controlling the volume, panning, and sometimes effects applied to tracks within a composition. For some groups of composers, this is a very important activity, so studies may describe it as a major stage of composition (Tobias 2013) or focus entirely on it (Dewey and Wakefield 2017; Jillings and Stables 2017).
- *Generating and combining alternatives*: deliberately recording many slightly different takes, and then "constructing performances" from small sections or even individual notes of many separate takes (Tobias 2013; Auvinen 2019).
- *Software management*: activity required to manage the composition software, rather than activity carried out purely to manipulate the composition. This might include navigating around the program itself (as opposed to the composition), configuring settings or accessing help facilities (Nash 2012), or adding identifying labels and colours to tracks within a composition to help understand it (McGarry et al. 2017).
- *Collaboration*: activities that are used to coordinate and manage collaborative

projects, which are out of the scope of this thesis, but are described in detail in other work, e.g. Coughlan and Johnson (2006).

Again, it appears that which activities are seen varies significantly between composers. Some authors have discussed why some activities might be seen in some composers but not in others.

A widely discussed difference is the amount of *auditioning* versus silent *reflection* carried out by different composers, with some suggested explanations being differences in experience level (Peterson 2008), generational differences (McCulloch 2014), and different macro structures of composition (Eaglestone et al. 2008; Nash 2012).

McCulloch (2014) suggests that certain study designs may cause the study to find little evidence of *planning* (due to time and representational constraints) or *improvisation* (due to the choice of data collection and analysis methods, which will often avoid the more time-consuming methods required to study improvisation activity, such as musical transcription).

The composition process appears to involve a large variety of activities, and which are observed can vary from study to study, based on what study design is used and who is studied. Any study of this area will likely not capture evidence of all the compositional activities being carried out, and it is likely that some activities involved in composition have not yet described in the literature.

As a result, an important study design consideration is to try and avoid constraining what kind of activity can be observed. Possible factors to consider are ensuring that the study uses *naturalistic* tasks (to avoid constraining what kind of activities are observed) and ensuring that the study is *longitudinal* (or studies a short time interval within a longer ongoing task) in order to remove artificial time constraints.

2.3.3 External aids in composition

A variety of external aids can be found in use in the reviewed studies. Hardware tools can include *MPC* devices, MIDI controllers, MIDI keyboards, and audio effects (Gelineck and Serafin 2009). Dictaphones and audio recordings can be used as external memory (Coughlan and Johnson 2006), with composers in one study asking specifically for overdubbing facilities to be added to their dictaphones, indicating this can be an important feature (Bainbridge, Novak, and Cunningham 2010). Phones and tablets are often used for capturing rough ideas or recordings of improvisations (McGrath, Chamberlain, and Benford 2016). The computer's file system itself may even form part of the composition process when composers stored information in file names or "notes" metadata (Duignan, Noble, and Biddle 2010).

The most consistent finding across many studies was composers augmenting the computer tools using paper and notebooks, with a considerable variety of different types of use observed. The information recorded on paper includes diagrams and visual representations (Healey and Thiebaut 2007; Polfreman 1997; Coughlan and Johnson 2006), graphs (Tsandilas, Letondal, and Mackay 2009), and mnemonics and codes (Eaglestone et al. 2008). Paper is sometimes used to record information that the composition software cannot model, such as quarter tones and positions of faders on a mixing console (Tsandilas, Letondal, and Mackay 2009), linkages and groupings between different musical elements (Eaglestone et al. 2007), and contextual information (Bainbridge, Novak, and Cunningham 2010). Diagrams drawn upon paper can use visual properties that do not have a consistent semantic interpretation (Healey and Thiebaut 2007). Tsandilas, Letondal, and Mackay (2009) observed printed scores making references to specific files contained in a computer, including poems, drawings, and code. Paper also permits annotation, coloured highlighting, and shading of areas (Healey and Thiebaut 2007), multiple paper types can be used such as graph paper and manuscript paper (Bainbridge, Novak, and Cunningham 2010; Bell 2014; Healey and Thiebaut 2007; Polfreman 1997), and different writing implements can give different advantages[2]. Composers may even print out screenshots of their composition software in order to annotate them (Auvinen 2019).

However, it is important to note that some studies appear to show little or no use of paper, such as the case study of an expert composer in Nash (2012) which appears to describe activity entirely focused on the computer.

It appears that external aids are used to complement the computer by a broad spectrum of composers. These aids are very frequently used to perform functions that are not accommodated by the composition software. Studies in this area

[2]See Garcia et al. (2011), who note that all their composers used pencil instead of pen for this reason.

should take care to record what external aids are used and what they are used for, particularly paper notes, by carefully choosing data capture methods and study designs that allow this.

2.3.4 Interface issues in composition tools

Where authors have evaluated the interfaces of existing composition tools, one criticism is more prevalent than any other: *representational determinism*. Coughlan and Johnson (2006) describe this by stating that the design of composition software constrains composers' creativity by restricting what kind of things can be externalised into the software.

A number of ways in which composition software could potentially restrict creativity by representational determinism have been identified. *Inability to represent material vaguely* could lead to composers being forced to make premature commitments compared to paper representations (Thiebaut 2010). *Inability to represent musical structures* could lead composers to habitually use "bottom-up" strategies when composing (Polfreman 1997). *Enforced use of a linear timeline* which can lead to using "chronological" strategies when composing (Amitani and Hori 2002; Duignan, Noble, and Biddle 2010, 27). *Enforced use of a 'track' metaphor* could lead to composers avoiding instruments, voices, or effects that appear only momentarily (Duignan, Noble, and Biddle 2010, 29). Finally, *rigid temporal grids* can lead composers to produce machine-like rhythms that lack 'groove', with some composers choosing to deliberately use software that lacks this feature (Duignan, Noble, and Biddle 2010, 27; Macchiusi 2017, chap. 5).

Marrington (2010) discusses how the restrictions of particular DAWs and notation software can have both negative but also *positive* effects on students who are developing their composition practice. Similarly, Bertelsen, Breinbjerg, and Pold (2009) studied composers who described how the restrictions imposed by composition software are "a constraint that both composers cited as a challenge and an inspiration". It is also worth noting that while the software used by composers at the time of these studies may have exhibited these restrictions, they may be less evident in more modern composition tools.

It is important to be aware of the ability for composition tools to constrain the creativity composers when designing a study. Some possible ways to mitigate this effect are recruiting composers with *diverse experiences* of composition software, studying the use of several *different tools* for composition, and studying *prototype* tools that remove some or all of the constraints described above.

2.4 Conclusion

This chapter has reviewed studies of the use of composition software.

The composers who have been studied have been recruited from one or more of several populations: professional, pro-amateur, amateur, children, and academics. They have been studied with a wide variety of methodologies, and in both lab and field settings, with recent studies also gathering data via the internet.

Four key themes were seen in the findings of these studies:

1. The macro structure of composition can vary in several ways: staged vs. iterative, top-down vs. bottom-up, and chronological vs. patchwork.
2. There is also a broad range of component activities used in the composition process, which often appear across multiple studies, but where the full range of activities is normally not captured by any individual study.
3. Composers have been observed using a broad range of external aids outside of the software, particularly paper notes.
4. Finally, the design of the tools composers use may constrain their creative work in many different ways.

The studies also suggest important design considerations for studies in this area:

1. **Triangulate** multiple data capture methods. Studies should use a methodology that produces multiple types of data that can be triangulated to produce a better picture of composition.

   The studies described in the following chapters will combine interviews with video and audio recording, observation, and in Chapter 5, data logging.

2. **Avoid constraining study designs** that might prevent certain creative processes or activities from being observed. Studies should be longitudinal, and use either naturalistic tasks, or observe composer's activity on work they are already engaged in.

Chapter 3 will study composers working on tasks of their choice, which may be part of their existing work, including a segment of work in progress.

3. **Capture use of any external aids** apart from the DAW itself such as paper notes.

   Chapter 3 and Chapter 5 will use a combination of video recording, photography, and observation to collect data on the use of any other aids. In Chapter 3, composers are studied in their normal environment so that aids that might not be available to them in a lab setting are still available.

4. **Mitigate constraints imposed by DAWs** that might restrict creativity. Studies should use a variety of DAWs or a prototype designed to avoid some constraints.

   In Chapter 3, the composers all have different experiences of DAWs and use different DAWs of their own choice. In Chapter 5, the composers are observed using a prototype that implements a technique for mitigating some limitations of DAWs.

Pro-amateur composers have been significantly less studied than other groups while being relatively easy to recruit. The studies in chapters 3 and 5 will focus on this group in the hope that new insights will be gained.

Chapter 3

Observing the use of existing DAW tools

In the previous chapter, the literature relating to studies of composers working with composition software was reviewed. The review highlighted some important themes in the existing literature and was used to create guidelines on appropriate study designs for studying composers.

This chapter describes in detail a study of four pro-amateur composers informed by the findings of the review in the previous chapter. From this study, three common patterns found in their working practices are discussed. Based on these findings, recommendations are made to help designers of digital audio workstations to support the needs of pro-amateurs.

Finally, the findings of this chapter and the preceding chapter are related to each other, suggesting that the remainder of the thesis should focus how *external representations* can support the particular activities and strategies used by pro-amateur composers.

This chapter is divided into four parts: a description of the **methodology** and study design, a description of the **data** gathered during the study divided into five different types, a **discussion** of patterns observed in the data, and finally **conclusions** about how these themes can be used in the design of digital audio workstations and to set the direction of the rest of this thesis.

3.1 Methodology

The aim of this study was to understand the existing working practices of pro-amateur composers working alone using commercially available *digital audio workstation* (DAW) software.

The activity of pro-amateur composers has not been frequently studied,[1] particularly when working alone, and it was unknown at the start of the study how their work might be similar to or different from other groups. As a result of this lack of prior studies, the study was designed to identify the important concepts for understanding the work of these composers by adopting an ethnomethodologically informed approach. This approach involved examining in fine detail exactly what these composers do and analysing this data to understand the problems faced by composers and their working practices.

The specific methodology used was based on *interaction analysis* (Jordan and Henderson 1995). Interaction analysis "investigates human activities… the use of artefacts and technologies, identifying routine practices and problems". It uses detailed records of activity, primarily video recordings of "the interaction of people with each other and the environment" (Suchman and Trigg 1995), which are then transcribed and analysed to find recurring themes in the activity.

Four composers were observed, all of whom had created a significant body of music over a period of longer than ten years, with a mixture of genders. None of the composers studied currently wrote music as their main source of income. Two composers (C and D) had taken university music courses involving composition in the past, but had left that environment years before and now released albums on an independent electronic music label.

Composers were recruited either through existing social connections or by "snowball sampling" using social connections of the composers being studied. The composers were not compensated for their time. Ethical approval was received from the University of Nottingham Computer Science Research Ethics Committee (see *Appendix C*) and consent was gained from participants by discussing and signing a study information and consent form with them to indicate their willingness to participate. As most of the composers were studied in their own home, composers

[1] See Chapter 2 for examples of studies of pro-amateur composers.

over the age of 18 were specifically recruited. Personal details of the composers were recorded for the purpose of administering the study but were destroyed after the studies were complete.

The composers were asked to work on a project they would normally be working on at that time, using their normal tools, in the environment they would normally work in. The author deliberately took the role of non-participant observer where possible. The composers were observed working for around an hour each. In most cases, the composer reached a natural stopping point (or completed the project they were working on) after around an hour and did not need to be asked to stop.

Each composer used a different digital audio workstation (respectively Ableton *Live 9*, Apple *GarageBand '09*, Apple *Logic Pro X*, and Steinberg *Cubase 4*) and had differing past experience of DAWs and score writing packages.

Suchman and Trigg (1995) identify four different ways of recording data in interaction analysis:

1. *setting-oriented* records are designed to capture the whole of the physical space in which work occurs. Photographs of the environment the composer worked in were taken. The composers in this study made only minor practical changes to their environment unrelated to the composition work during the sessions (e.g. closing a blind to reduce glare from sunlight, opening a window to cool a room down), so further records of the setting were not necessary.

2. *person-oriented* records are intended to capture the work done by a particular person. A video recording was made of everything that took place on the composer's computer screen during the task. Where physical input controllers were used, the video recording was arranged to include those actions where possible. The video recording also contained an audio recording of everything spoken or played while the composer worked on their chosen task, which was often useful when attempting to transcribe what actions the composer was taking.

3. *object-oriented* records which track particular artefacts. Photographs were taken of any external notes and aids the composer used, and where those notes were added to over time, more photographs were taken. If a composer created notes that were not immediately self-explanatory the composer was

asked to provide some additional explanation.

4. *task-oriented* records are intended to capture how a task is done as it moves between multiple individuals. No special attention was given to making task oriented records as the study was intended to be of composers working alone. However, during the course of the study, it became apparent that this may have been an overly simplistic view: three of the four composers continued to talk at the (deliberately non-participating) observer during their composition process, despite not having been asked to do so, and would also frequently talk to the computer they were using *as if it were a person*.

Before the observation began, a semi-structured interview was carried out, designed to elicit biographical information and information about the context of the session. Audio of this was recorded and later transcribed.

3.2 Data

3.2.1 Interviews

Information about the backgrounds of each composer and what tasks they were attempting was recorded as part of an initial biographical interview.

- *Composer A* had published albums online in a variety of genres and worked both collaboratively and on solo projects. They had strong skills at organ and recorder but often played other instruments. They were observed while they were writing a vocal melody and piano backing to the words from a poem. Before the session, they had already thought about the task, and had some rough ideas for themes and a fragment of a chord sequence.

- *Composer B* had over 25 years of experience as a singer and songwriter. They primarily wrote folk music but had also worked in other genres. They were observed while writing a track for a new album which was constructed from audio fragments from previous works. They described this as using “snippets” to make a “soup” or “musical collage”. Before the session, they had identified three snippets they wanted to use and a title but had done no other preparation.

- *Composer C* had played drums in a variety of musical groups since the mid-1990s. They had collaborated with a singer to compose two albums, had written soundtracks for short films, and ran their own record label. They were continuing work on an existing project, writing a piano part as part of a backing that would later be given to a vocalist.
- *Composer D* was an electronic musician and guitarist who also had experience playing the bass guitar, piano, and performing live with effects and a sampler. They played flamenco guitar as part of a duo in local venues and had released two experimental electronic solo albums on an independent label. They began work on a track for a solo album that was entirely constructed from discarded audio from previous projects.

While in theory, all four were very experienced composers, it was noticeable that some made significantly more errors while working. These composers also made considerably less use of "shortcuts" to speed up repetitive tasks. Expertise in using music software (or expertise in composition in general) is not the same thing as expertise at using a specific DAW software package, and these composers should be considered of a range of levels of expertise at using the software they are using in this study.

3.2.2 Artefact-oriented records

Information about what artefacts the composers used was captured through video, photographs, and notes made while observing the session. Figure 3.1 illustrates this information.

The artefacts used spanned a variety of different mediums:

- *Paper notes* were used by Composer A. As it is a common feature in other studies, all the composers were asked for more information about how they used paper using their composition process, and all of the composers were able to demonstrate paper notes they had used in the past. Photographs of these can be seen in Figure 3.2.

  While the paper notes used by Composer A contained lyrical and melodic information, Composer C had scrapbooks of newspaper clippings which they had assembled *after* the composition process which were thematically

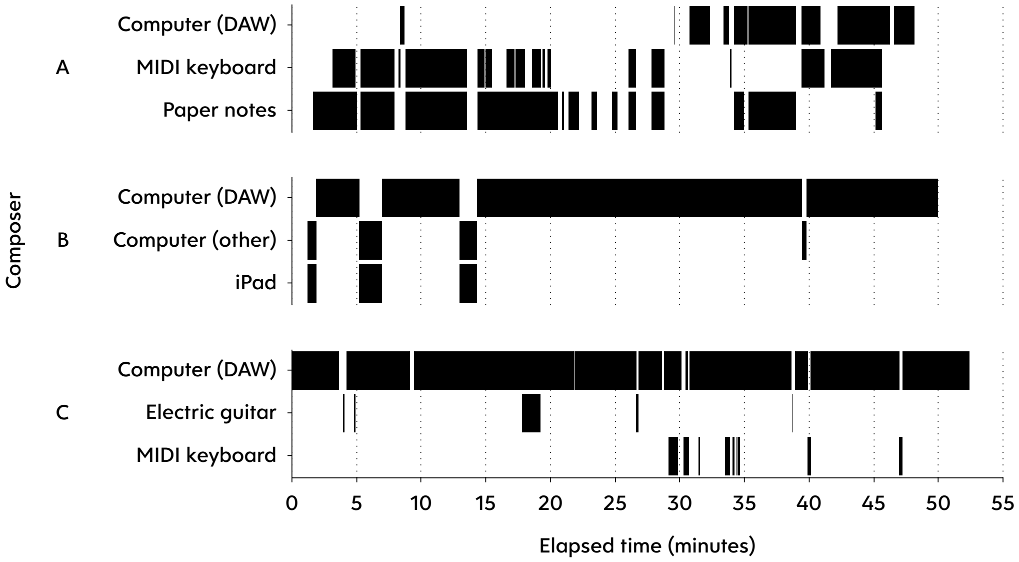

Figure 3.1: Use of artefacts by Composers A, B, and C

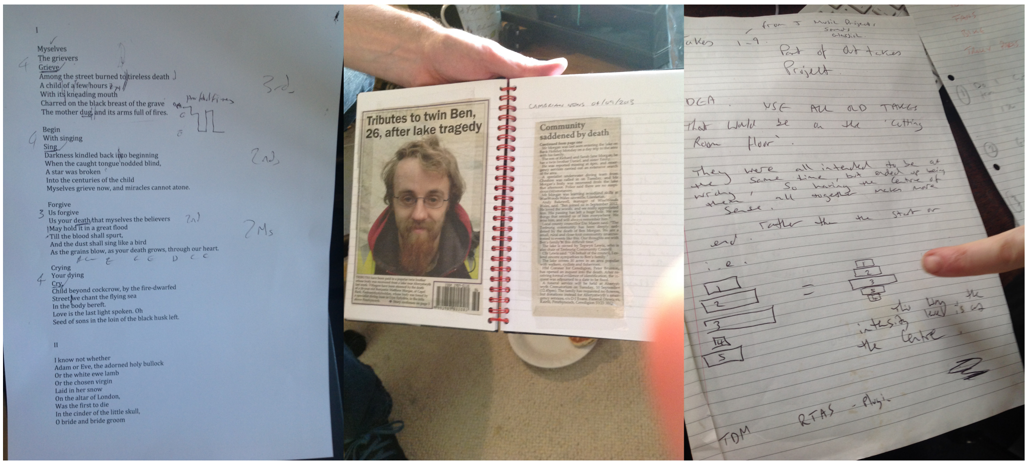

Figure 3.2: Paper artefacts demonstrated by composers A, C and D when asked about their use of paper

related to the work. Composer D had drawn diagrams illustrating the overall structure of a piece, using the same multi-track timeline metaphor used by their DAW.

- *Electronic text notes* were observed being used instead of paper by composers B and C during their sessions, to record information about their future plans. These did not appear to be used in ways that were unique to being held as electronic text - they could have served the same function as handwritten paper notes. Composer C additionally used a built-in ability of *Logic* to store timestamps of interest by adding "markers", and referenced these in their notes.

  In general, the paper and electronic notes were used to represent information about melody, rhythm, note velocity, timestamps of interest, and future tasks. There was no evidence of using these notes to store information about other musical qualities which are normally significant in composition (e.g. dynamics, texture, timbres, tempo, emotional feel, etc.).

- *Musical instruments* were used as a way of aiding the composition work by composers A and C. There were three obvious uses of instruments: improvisation, as a data entry method, and for making recordings.

  When instruments were used for improvisation, composers often did not interact with the DAW at the same time. Improvisation took place over a significant proportion of the time spent by Composer A and a small amount of time spent by Composer C. Both composers did not make recordings while they were improvising, which resulted in them following the improvisation by trying to memorise or recreate what they had just improvised.

  Instruments were sometimes used as a more effective data entry method for pitches, chords, or intervals. Entering notes in this way, rather than using a standard computer keyboard and mouse, leverages the composer's existing familiarity with the instrument to accomplish the task more quickly and easily. Composer C used a particularly complicated method of doing this: they would play a note on the guitar, which was plugged into a guitar tuner pedal which displayed a letter corresponding to the pitch of the note, and then used the mouse to enter that pitch into the DAW.

- *Customisation of the DAW* was observed taking place in limited ways. Customisation tended to take one of two forms: changing the functionality of the mouse pointer, and showing or hiding functionality.

  Logic provided the ability to customise which action was triggered by each mouse button, and Composer C had already set up this customisation before the session started.

  Cubase only provided the ability to customise which action triggered was by the *main* mouse button, which meant that Composer D had to keep changing modes during the session in order to access other functions.

  As the DAWs must provide a large number of features within limited screen space, features can be shown or hidden to provide just that functionality required in order to make the most efficient use of the space. In these observations, the automation controls, media browsers and annotation windows were observed being shown and hidden. This customisation usually took place at the start of the session, and afterwards the composers then almost exclusively used the features that had remained visible on the screen.

3.2.3 Person-oriented records

Information about which features of the DAW the composers used was captured through screen recordings. The features observed can be grouped into "collections" to indicate the kind of work taking place. Table 3.1 indicates which composers made use of features in each collection.

Table 3.1: DAW functionality used by each composer in the study

Collection	A	B	C	D
Timeline controls	✓	✓	✓	✓
Recording audio or MIDI data	✓			
Editing audio clips		✓	✓	✓
Editing MIDI clips	✓		✓	
Editing envelopes			✓	✓
Track controls		✓	✓	✓
Browsing media libraries		✓		✓
Looping a section of the timeline	✓			✓

Collection	A	B	C	D
Exporting an audio mixdown		✓		

In addition to providing features for viewing, playing or editing the music, each DAW provided to the composer libraries of pre-created elements that could be used in the composition process, and composers B and D spent time using these. Depending on the DAW these libraries might include lists of samples, effects, instruments, drum loops, or musical transformations. Some of the composers studied had also created their own libraries.

Table 3.2 indicates how the composers used these libraries. For each choice from a library, the number of possible alternative choices investigated and rejected is listed, along with the total time taken between starting to search until committing to the final choice.

Table 3.2: Data on choosing options from libraries of options

Composer	Library	Choices	Alternatives tried	Time taken (sec)
A	*None used*			
B	Audio loops	Background pad sound	0	25
		Drum beat	5	195
		Glockenspiel loop	0	10[2]
		Bassline	3	31
		Bassline alternate	1	70
C	*None used*			
D	Audio samples	Guitar sample	5	74
		Ambient sample	0	24
	Audio effects	Resample	0	20
		Time stretch	0	12
		Normalize	0	7

[2]This loop was found serendipitously while searching for a drum loop.

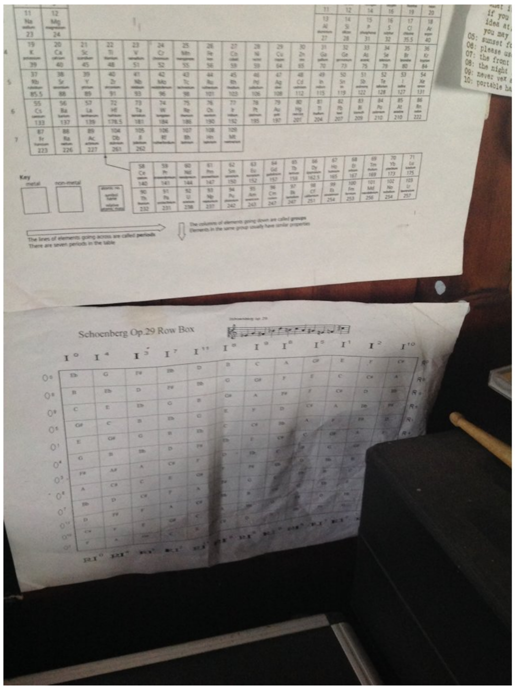

Figure 3.3: Inspirational material in the working environment of Composer D, including a tracklist of a previous album

3.2.4 Setting-oriented records

Information about the settings the composers were observed working in was captured in the interview and using photographs. Composers A, B, and C all composed in their homes, while Composer D used a desk in an artist's studio very near to their home. All four rooms contained a variety of different musical instruments close to hand, a computer, and a stereo system.

Apart from these items, the contents of the rooms were inspirational rather than directly related to the composition in progress. This included artwork from local artists or of admired musicians, tracklists of *previously* completed albums, a periodic table, books about music, and a diagram showing the tone rows used in a Schoenberg serialist composition (see Figure 3.3 for an example). While decorating the room in this way was clearly useful to the composers (perhaps as a way of tracking their musical history), there was no point at which any of the composers were observed to make any use of these things, nor did they describe having used them as part of composition.

While Composer C had a whiteboard above their desk which they said they "always intended to use" as part of their composition, it had actually been used to record information for their record label. Composer D had a portable audio recorder sat

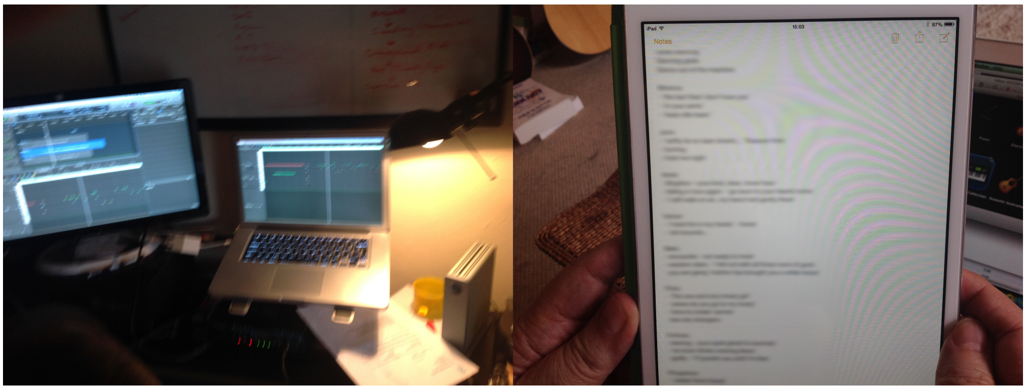

Figure 3.4: Multiple screens in use by Composer C (left) and Composer B (right)

on their desk plugged into the computer but did not make use of it.

Composer C used multiple monitors in their work (see Figure 3.4): one displayed a piano roll view of the part they were writing (the main view used in their work during that session), while the other monitor displayed the main multitrack timeline of *Logic* (which they would intermittently use). Composer B also used two screens - one on their laptop, and the other on their tablet computer; the laptop displayed *GarageBand*, while the tablet was used to display notes they had previously written.

3.3 Discussion

3.3.1 Varying support for composition activities

The previous chapter identified that composition involves many different activities, and evidence of many of these activities was also seen in this study. Both Composer C and Composer A had *improvised* prior to the study session, Composer A and Composer B would quietly *reflect* for periods of time, and Composer B and Composer D *collected and captured* samples from their previous musical activity. Closer examination of how the composers carried out these three activities reveals evidence of a *lack of supporting features* in the DAW to aid these activities and a pattern of composers *choosing to avoid* using the DAW while carrying out these activities.

For three of the composers, a **lack of supporting features** made carrying out these activities more laborious:

- *Composer A* improvised ideas by singing and playing a MIDI keyboard but

Ableton Live provided few features designed to capture a record of improvisation activity. The composer only made use of *Live* at the very end of the improvisation process when they had generated an idea they wished to move forward with. They were required to make several attempts at recording the ideas before they were able to record a performance without any errors. For each attempt at recording the material, they were required to set up the recording again for a new take, interrupting the recording process.

- *Composer B* collected material they wished to use in a song by entering text notes on their iPad to remember which songs they planned to sample. However, as there were no hyperlinks between these notes and the audio files they referred to, the composer was required to search through their *iTunes* music library to find the correct files. Even once they had done this, they had to perform the time-consuming task of manually trimming the audio clips to ensure they started and stopped exactly on the beat.
- *Composer D* collected material at the start of their session. This activity involved searching for audio clips that they had previously recorded but not yet used in any compositions. However, the only method in *Cubase* for finding and adding material to a composition was the "Import Audio" window, a standard Windows file open dialog box to which some playback controls had been added. The composer stored their past compositions in the standard file system hierarchy. They were required to manually search this structure, looking for audio clips they might not have used in previous projects, and auditioning them one at a time to see if they contained appropriate and unused material.

If activities are poorly supported by a DAW, there would seem to be little advantage to using a DAW to do them, and it would be likely that composers would not wish to use the DAW to carry out those activities. In several cases, composers in this study do appear to be **choosing to avoid** using the DAW when carrying out some activities:

- *Composer A* did not use actually use the *Live* interface until over eight minutes into the session (see Figure 3.1). At some periods during the session, they would sit silently *planning* and *reflecting* (visible in the paper notes they were making), not using the DAW at all.
- *Composer B* did significant preparatory work making text notes on an iPad

before they even opened *GarageBand*. They spent the first two minutes of the session consulting these notes and they took three further breaks from using *GarageBand* to consult this list and audition material in *iTunes*. Composer B regularly took brief breaks to *reflect*, saying "I want to make some sense of it… to just consider what I have", but this process involved little interaction with the DAW.

- *Composer C* was observed in a session in which their work was entirely based in the DAW. However, they were working with material that had been generated in a previous *improvisation* session during which the DAW was used purely to create a record of the improvisation.

Composers may deliberately choose to avoid a DAW for many reasons regardless of how well the activity is supported: as a stylistic choice, to deliberately constrain themselves (see *Selective allocation of time and effort* later in this chapter), or to avoid issues related to representational determinism. Nonetheless, the behaviour seen in this study is suggestive of an opportunity for better software support for some common activities composers carry out. This support does not have to be part of the DAW itself, as composers may be better served by a special-purpose tool for the activity, rather than more functionality within the DAW (for example, Apple's *Music Memos* mobile app that complements *GarageBand* by providing idea capture facilities).

Some authors have discussed the opportunity to better support these activities. While *Live* did not have such a feature, Macchiusi (2017) notes that some DAWs now include a "loop recording" or "cycle recording" feature designed for repeated recording of several takes of a section. Coughlan and Johnson (2006) found a similar need for better support capturing and collecting of material (in particular, automated support for aligning audio clips to the beat and removing silences at the beginning and end of clips).

The findings in this study reinforce the argument that study designs should triangulate data from multiple sources. Some activities may involve little interaction with the DAW and many data capture methods (such as automatically logging usage data or screen capture recording) will result in data in which certain activities are not visible. In this study, there were indications of further activities taking place that were not clearly visible in the data that was recorded. As an example of this, Composer A took a break midway through their session to fetch a drink, and on returning immediately wrote down an idea that would be important in the rest of their work. Studies of composition software have noted that "the development of ideas often occur away from intentional periods of work" (Coughlan and Johnson 2009b), and that "much of the creative process is happening away from the computer, e.g. between computer-based composition sessions" (Eaglestone et al. 2007). Some kind of unconscious *incubation* activity - "time away from the problem when conscious work is set to one side" (Collins 2005) - may have been happening away from the DAW. However, with the data capture methods used, it is not possible to determine what the composer was doing or thinking during this period.

3.3.2 Coordinating and combining multiple representations

The composers all attempted to keep track of multiple different ways of describing their composition. These representations often focused on different aspects of the composition or musical process. Their workflow involved coordinating the use of these representations together.

Composer A merged information they stored on paper (about lyrics, syllable structure, and important words) with a separate set of information held in their head (about chords of interest and some melodic fragments). They then translated this into audio recordings, which imposed a fixed rhythm and tempo on the music, and established a fixed single version of the music.

Composer B merged information they held on their iPad (about the mythological inspiration for their work, the emotional feel of their work) with clips from their iTunes library of past work (as audio clips), and translated it into a fixed ordering in time within a *GarageBand* timeline.

Composer C merged information in a recording of a drum improvisation (containing rhythmic information) with harmonic thoughts from past guitar improvisations and improvisation they did during the session. They also produced additional information in the form of text stored within *Logic* (to indicate proposed future work).

Composer D described how in a past composition, they had merged a visual sketch of the proposed structure of a piece of music with a prepared list of audio clips to use from their sample library. They then translated this into a composition within *Cubase*.

This practice of combining multiple different representations has been referred to in multiple places in the literature. Balaban (1996) points out that all representations of music can only represent part of the music:

> "An essential property of the real world problem in music is that it is not conceptualizable in its totality… In music, a complete formal account for the semantical phenomenon is not attainable… The music world described by a system is always partial."

In a later paper, Balaban and Elhadad (1999) describe how composers need to be able to link and merge partial representations of different aspects of the music:

> "These different levels cannot be described independently of one another since they ultimately describe the same ontological object. Instead, they must be viewed as different perspectives on a common, partially known object… At the user-interface level, one must be prepared to integrate and coordinate different media to describe the relations between these different levels."

Bainbridge, Novak, and Cunningham (2010) identify that "ideas can consist of multiple pieces of information, such as raw audio, text, and images combined. Therefore support for directly associating, or grouping different types of media together is needed." Coughlan and Johnson (2006) describe how this practice is significantly muted in users of compositional software: the software users had a "complete representation of the composition that the computer played whereas the composers in the unsupported observations used multiple partial representations of the composition." They reach the same conclusion: "since composers have been found to use both visual and audio representations of their ideas the aim is to link these representations in a suitable manner".

Eaglestone et al. (2007) describe how they observed composers working around the limitations of having one representation in their software by using *multiple* software packages at once:

> "All composers involved within the observations elected to work with

> multiple applications… The observations supported the view that this is not only a phenomenon that composers have learned to live with. It also has an important positive impact on their compositional process and appears to support their creative behaviour."

Eaglestone also observed this process of moving between different representations as having a "catalytic effect" on the composers' creativity.

It appears clear from both the behaviour of the composers in this study and the literature that the ability to coordinate multiple representations is a desirable feature.

3.3.3 Selective allocation of time and effort

Three common behaviours can be observed across the composers studied:

1. *Habituation*: repeatedly doing the same thing or one of a small set of things
2. *Limited exploration*: limiting the time they spend considering possible options
3. *Self-constraining*: imposing additional restrictions upon themselves

All three of these behaviours can be viewed as strategies the composers use in order to be selective about how they allocate their resources of time and energy in their work. These resources are finite and limited, meaning composers need to prioritise how they are spent, but to do this they must inevitably make trade-offs in other areas. The qualities which are sacrificed by the above strategies are respectively *novelty*, *quality*, and *flexibility*.

The "work" of composing in a DAW consists not just of recording MIDI notes or audio tracks, but also of making creative decisions, which can be frequent and numerous. DAWs provide a difficult environment in which to make these decisions: Duignan, Noble, and Biddle (2010) report a composer in their study used the term *option dilemma* to describe the "paralysis caused by the overwhelmingly open design space provided by computer-music systems" that "encouraged endless experimentation and fine-tuning", with Magnusson (2010) also referring to a "practically infinite expressive scope" causing "creative paralysis", and Gelineck and Serafin (2009) noting this issue as one that can "kill creativity". The strategies composers use to control their use of time and effort are applied not just in the activities of recording and editing material, but also in activities relating to decision making.

Habituation involves repeatedly doing the same thing or one of a small set of things to avoid making decisions about unfamiliar options.

Composer D repeatedly used just three of the wide variety of audio effects available in *Cubase*, saying that they were what they "tend to do" or "generally do". Composer B specifically looked for a category of loops they had previously used, immediately rejecting drum beats that were not from that category.

Habituation can be used to limit options to those known to be quick and easy to do, and it guarantees a quick and easy decision making process. Habituation can help to create a "trademark sound" for a composer or an album. Habituation sacrifices *novelty*, but in a way that may help to maintain a consistent level of quality.

Limited exploration means the composer deliberately limiting the time they spend considering possible options. For example, composers can limit their exploration of libraries of effects, loops, and samples, or their exploration of synthesiser and effect control values.

Composers B and D both appeared to make use of limited exploration when browsing libraries, displayed extremely pronounced *satisficing* behaviour (stopping searching alternatives when an acceptable option is discovered). The mean time spent on locating and evaluating a given option was just 19 seconds. In most cases, this would be just enough time to listen to an option once before deciding whether to select or discard it.

Composer C significantly limited their exploration to particular aspects of the music, focusing on altering the pitches and note lengths of the MIDI they were manipulating, while paying no attention to the (unquantised, often out of time) note onset timings.

Composer A represented the specific qualities of interest to them in their paper notes. Composers who limit their exploration to specific musical aspects can encode these limitations in representations they create by selectively representing just those aspects.

The strategy of limited exploration reduces the time and effort costs of making decisions. However, it is less useful as a way of reducing time and effort costs of other activities when compared to habituation, as known cheap options are not repeatedly reused. Limiting exploration in this way deliberately sacrifices *quality*, but it could potentially help to maintain more novelty and flexibility than the other strategies.

Self-constraining is the strategy of a composer deliberately choosing to impose additional restrictions upon themselves. For example, composers may constrain themselves by their choice of creative inputs and outputs, genres, or tools, and also by tying their musical material into conceptual frameworks.

Composer D restricted themselves to creating an entire composition out of two short samples, as part of a larger project that was constrained to use only samples from previous projects.

Composer A constrained their work by choosing to set a poem to music. The poem constrained the overall musical structure of the composition (which needed to match that of the poem). The poem also directed the composer's approach to harmony: the planned mood for the composition was clearly influenced by the themes of the poem, and the composer repeatedly used discordant intervals to achieve this mood.

Composer B used the Roman pantheon as a conceptual framework that constrained the overall form of their album, as well as informing the mood of the specific track worked on in the study. This framework appeared to have appeared organically during composition: "one was called Venus, because I'd just written a song about Venus, and then I did one called Mercury, so I'm doing Roman gods".

Constraints and conceptual frameworks can add additional interest to a composition but can potentially overshadow the music itself: Composer C described their frustrating experience studying at a conservatoire where "the whole thing seemed to be 'how good is your concept' - not - 'what does your music sound like?'… I'm not interested in that, I'm interested in what it sounds like".

When imposing constraints on themselves, composers may find the process of making decisions simpler and less overwhelming. Composers can also use constraints to limit themselves from undertaking activities that are costly in terms of time and effort. Composers using the strategy of self-constraining sacrifice *flexibility*, but a higher quality result could potentially be achieved than by using the other two strategies.

Out of the three strategies the existing literature discusses the strategy of self-constraining in significant detail:

> "a common strategy can be detected, defined here as that of designing constraints… encapsulating a defined space for potential expression." (Magnusson 2010)
>
> "Most subjects set boundaries, rules, dogmas, limiting their options in order to guide or challenge the creative process… musicians want help for making decisions". (Gelineck and Serafin 2009)
>
> "part of the enjoyment… for [composer Andy Carthy] is in having to work with the inherent restrictions, such as the elements in a sample that you wouldn't necessarily choose to include" (Morey 2013)

Healey and Thiebaut (2007) also remind us to avoid thinking that the "ideal situation for creative composition is one in which there are no constraints". It has even been suggested that to avoid confusion, using the more positive word "criteria" might better describe constraints in creative tasks (Candy and Edmonds 1997).

3.4 Conclusions

Three main conclusions were found as a result of this study.

1. Existing composition tools support different composition activities to varying degrees, and some activities could potentially be better supported, including improvisation, reflection, and collecting and capturing material. This suggests composers require either adaptations to the software to support these activities or a second complementary tool that focuses on them.

2. The composers in this study were observed making coordinated use of multiple representations. Evidence of the value of this approach has also been reported by multiple authors in the literature. This suggests that these composers require the ability to create secondary representations which present the composition in different ways.

3. The composers in this study were observed trying to selectively allocate their limited time and energy to prioritise areas of importance to them. This

> suggests that these composers would benefit from software that supported their strategies for focusing their time on the areas of interest to them. Digital audio workstations should allow easy reuse of habitual strategies, linking of music to conceptual frameworks to constrain composition, rapid exploration of libraries, and focusing representations on specific musical qualities.

This chapter has identified some key considerations for supporting pro-amateur composers.

In particular, a common theme throughout both Chapter 2 and Chapter 3 has been the importance of the representations used by composers: in terms of how they may restrict composers through *representational determinism*, the wide variety of representations created on paper by composers seen in the literature, and the need for pro-amateurs to make coordinated use of multiple representations in their work.

As a result, the following chapters will draw particular attention to how representations in composition software can be designed to help support the practices that have been identified so far.

Chapter 4

Representation design in DAWs

The previous chapters have discussed the importance of considering the representations used by composers. This chapter will now focus on the representations used in composition software. In particular, this chapter will identify specific recommendations on how these representations should be designed in DAWs.

A significant quantity of existing research has previously discussed the design of representations in composition software and the theoretical approaches that have been adopted by these previous studies will first be described. After summarising the kinds of information that can be represented in a DAW, a detailed description is then given of a set of representation design techniques relevant to DAWs. For each technique, existing studies and prototypes relevant to the technique from within and outside the field of composition software are described. Based on relevant insights from this prior research, specific suggestions are made regarding how these techniques could be implemented. To aid the reader in navigating this material, a summary section then follows which highlights the key findings.

Finally, future directions for research in this area are discussed, and the *task lists* technique is identified as a good candidate for further exploration. The following chapter will explore *task lists* in more detail.

4.1 Introduction

Chapter 1 identified how the design of DAW interfaces, particularly those which attempt to use interface metaphors that imitate physical studio hardware, can act to constrain the activity of composers.

In Chapter 2, a survey of the existing literature identified how these kinds of criticisms have previously been made using the terminology *representational determinism*. Composers have been observed using a broad range of aids and representations of information outside of the DAW which may allow them to avoid the issue of being constrained by representational determinism.

In Chapter 3, the study of pro-amateur composers found that there was indeed a potential for improvements to DAWs to provide better support for some of the composition activities that the composers were undertaking. Similarly to the strategy described in Chapter 2, these pro-amateur composers made use of additional representations of the composition, coordinating and combining them as required in a way that may help them avoid issues of representational determinism.

This chapter will further discuss this theme of representational determinism, focusing on how composers interact with representations within the DAW, and how these representations affect the way they work. In particular, it is concerned with how design improvements to the representations used in DAWs could better avoid representational determinism by increasing their support for common composition activities.

To begin, the existing research relating to the representations used in composition software will be discussed. The existing research has adopted several different approaches from psychology and human-computer interaction to perform this research. These approaches are closely related and have many similarities, and include *active externalism*, *activity theory*, *distributed cognition*, and *situated cognition*. The next subsection will discuss in more detail how previous authors have applied these approaches.

Following this, the remainder of this chapter will discuss potential design techniques that could be applied to the representations used in DAWs, based on the understanding that improving these representations can affect their creative processes. For each technique, suggestions of approaches from other categories of software that could improve representations within DAWs will be highlighted.

4.1.1 *Theoretical approaches*

Past authors have often used one of several closely related frameworks when analysing the behaviour of composers using DAWs. In general, these frameworks are concerned with how people carrying out cognitive tasks use not just own their mental resources, but also external resources in their surroundings, to achieve their goals. For example, the theory of *embodied cognition* suggests that:

> "We off-load cognitive work onto the environment. Because of limits on our information-processing abilities (e.g., limits on attention and working memory), we exploit the environment to reduce the cognitive workload. We make the environment hold or even manipulate information for us". (Wilson 2002)

While these theories can disagree on important aspects (in particular, whether the environment is used *for* cognition, or whether cognition *is* the combination of internal and external resources), one aspect they have in common is that they describe how we create *external representations* of cognitive processes that aid in activities we are carrying out. For example, we might create a shopping list to reduce the amount we use our memory. In this case, the shopping list is a *cognitive artefact*: an external representation that can be used to enhance our abilities to perform tasks that require cognitive work, and whose use may even change which cognitive tasks we are able to perform (Norman 1993).

In the case of DAW software, we are specifically concerned with people using software tools as cognitive artefacts. One member of this group of related theories of cognition, *external cognition*, is concerned not just with how the physical environment can be used to help us think, but specifically how software tools can be used for this purpose (Scaife and Rogers 1996). From the viewpoint of external cognition, composers use the external representations provided by the DAW software as cognitive artefacts to aid their creative processes. If composition software is being used as part of external cognition, representational determinism restricts not only what can be represented and what creative processes can be performed, but also restricts creative cognitive activities.

Active externalism (Clark and Chalmers 1998) is invoked by Magnusson (2010) as a way of explaining how "live coders" (musicians who, as performance, write software to generate music in front of an audience) use the live coding system as "a scaffold for externalizing musical thinking" which "attempts to ease the live coder's cognitive load".

Activity theory (Nardi 1996) is adopted (and adapted) as an analytical framework in the study of DAW users by Duignan (2008). Duignan created a questionnaire called an "activity interview", inspired by similar checklists proposed by the originators of the activity theory concepts, in order to understand the work and externalisation activity of composers through an activity theory lens. Duignan describes how:

> "one of the most striking features of the computer music production activity is the extent to which activity and mental processes are externalised... In our observations, the iterative process of creative action took place through the medium of the externalised object."

Activity theory is also used as a lens to understand composition using computers by Bertelsen, Breinbjerg, and Pold (2009), who see music software as a *mediating artefact* as described by activity theory. They describe composers as creating chains of artefacts that are used to crystallise or reify actions that are part of their practice. They point out that in composition, the boundaries between "tool" and "object/material" are dissolved: "the traditional HCI triad has become insufficient, because there is no clear-cut domain object and often no well-understood tool".

Distributed cognition (Hollan, Hutchins, and Kirsh 2000) has been used by Thiebaut (2010) to analyse the behaviour of composers. They describe sequences of representations that were used by the composers they studied and observe that "composer's cognition co-evolves with the successive representations". They note that "distributed cognition provided a useful but also in some respects limited framework for an analysis of representations in composition", but note two main conclusions: the importance of *vagueness* in the representation of compositional intention in order to avoid premature commitment, and of *re-interpretation* of representations to generate new ideas. Coughlan and Johnson (2009a) also refer to distributed cognition ideas, noting that composers "generally described their creative processes as a series of transitions where new representations were made based on existing ones, supporting the continuing development of the idea in a *distributed cognitive system*".

Situated cognition has been used in studies of the computer-based compositional activity of teenagers to view that activity as *situated practice* (Folkestad 2011). Folkestad says that adopting this sort of framework leads to a view that "learning is involved in any activity" and that "in music-making, activity, knowledge and learning can be considered as integrated… inseparable facets of a unified whole… to create music also involves learning how to create music". Donin and Theureau (2007) also describe their work using the terminology of situated cognition:

> "the particularity of situated music composition is that many important elements of the composition situation have been constructed in the past by the composer himself. This explains the essential role of memorisation, inscription, and re-reading and their corresponding techniques, which participate in the construction of an ensemble of which the realized work is only one of its most obvious manifestations."

Other related approaches exist, including *joint cognitive systems* (Hollnagel and Woods 2005) and *embodied interaction* (Dourish 2004), but these are normally used to explore the space of physical hardware interfaces rather than conventional music software, as in Bennett (2010).

By surveying the work of the previous authors, either those who have used these frameworks or who have used methodologies that have focused the on the representations used by composers, this chapter will identify a set of practical design techniques that can be used by designers of DAWs.

The previous research described in this section has identified the importance of the use of representations as part of composition activities, and the discussion of these techniques will focus on what kinds of representations composers create, and the kinds of representations that support composers in the activities they commonly perform.

4.2 Information represented in DAWs

Before discussing the kinds of representational techniques that could be used in DAWs, it is important to consider *which* kinds of information are represented by DAWs. While the DAW as a category of software may have initially formed as the combination of editors for waveforms and MIDI data (see *Historical development*), a brief inspection of the interface of any modern DAW would make it immediately apparent that DAWs actually store many types of information, such as:

- *clips* of audio or MIDI data whose start points and duration are specified within the timeline
- *tracks* that group these clips together into conceptual "musical parts"
- the *soundbox*: "features that conventional notation and analysis do not consider but which are, in fact, crucial for the understanding of the musical processes from a compositional and listener's point of view" (Pavese 2016), such as track gain or stereo position
- *routing*: a potentially complex network of signal routing called "sends" and "busses" between tracks, often poorly represented to the user (Koda 2011)
- *instruments* and *effects* that indicate how clips within a track should be processed before sending to signal routing
- *parameters* of these instruments and effects, such as choice of sample set or the amount of filter resonance
- *global parameters* that affect all clips in the composition such as tempo, "swing", "groove", or time signature (and which in some software can be overridden within specific tracks or clips)
- *automation* of parameters and global parameters: often called "envelopes" or "curves", these are usually displayed in DAW software in panels that can be shown or hidden when required

Some kinds of information are not directly associated with a specific composition but instead with the composer's process and may be used across many compositions:

- *saved presets* which may range from a stored set of parameters for a single synthesiser up to an entire network of synthesisers, samplers, and effects used to produce a single "instrument"
- *libraries of samples and loops* that the composer has saved for future use
- *preset arrangements of windows or panels* which the composer uses for specific purposes

As a result, the techniques in this chapter will be considered not just in terms of their use in representations of the audio or MIDI data that are contained in the composition, but also their use in representations of the other elements of compositions, and representations of elements of the DAW itself.

4.3 Techniques for designing representations in DAWs

The types of information described in the previous section can be represented in different ways, changing which compositional processes, activities and focusing techniques they can use. Each of the following sections focuses on a proposed technique for representing these elements that can be used to support particular composition activities.

For each technique, the following sections will describe:

1. a summary of the technique
2. which composition activities are supported by the use of the technique
3. existing use of the technique in research into composition software
4. existing use of the technique in other categories of software outside of music

> As-yet-unexplored avenues in the design of DAWs that warrant further investigation will be highlighted in boxes like this one.

Following this, a summary section will reproduce the most important information on these techniques in a tabular format.

4.3.1 Selective representation

Almost all composition systems cannot display all aspects of a composition at once. A useful technique is to allow the composer to specify which aspects of the work they would like to be represented.

Composers already make extensive use of user interfaces that selectively focus on particular aspects of compositions. Historically, these interfaces have been hardware tools such as mixing decks or four-track recorders, which only provided control over a small set of features of the audio such as panning and volume of tracks. As digital audio workstations have replaced these tools, the interfaces of those physical tools have been reproduced in software. Some authors describe these user interfaces in terms of *user interface metaphors*: sets of mappings from the data comprising the composition onto metaphorical objects represented to the composer (to paraphrase the more sophisticated definitions given in Barr, Biddle, and Noble (2002)). In these mappings, aspects of the composition become properties of the metaphorical objects represented to the user that they manipulate.

However, as Stowell and McLean (2013) note, the aspects of music that comprise the "problem space" of music varies between different genres, or even between tracks. They suggest that as a result, the fixed set of metaphorical mappings that are often provided by composition software are inadequate:

> "such metaphors reflect neither the 'problem space' (the target music domain) nor the breadth of possibilities provided by the computer.... The user may be left feeling as though they are dealing with a nonsensical metaphor, which induces unneeded limitations (such as running out of display space), and embeds now-irrelevant design decisions ... structuring software design around fixed metaphors does not hold cognitive advantage".

As a result, some authors have suggested that composers require the ability to select *their own* sets of music aspects that will be represented to them metaphorically. The pair of composers studied by Bertelsen, Breinbjerg, and Pold (2009) are reported to "typically transgress the metaphors" provided by their software. Stowell and McLean (2013) concluded that "if we assume that everyone has their own systems of metaphor... then we should instead develop interfaces that let people apply their own metaphors". Pearce and Wiggins (2002) claim that creative cognition "is supported by the ability of the composer to simultaneously represent multiple features of the emerging composition and to move flexibly between them during composition."

Some support for selective representation already exists in DAWs, as most give the ability to choose a small subset of automation curves to be displayed for each track (see Figure 4.1).

4.3.1.1 Activities supported

When the user interface focuses on a small set of aspects of the composition, it may ease the process of creative cognition. Tubb and Dixon (2014a) describe how mapping two musical properties to axes on a 2D plane may allow composers to switch from an analytical conscious mode of *explicit* thought into an instinctive *implicit* mode. This would be particularly suitable for supporting *improvisation* activities.

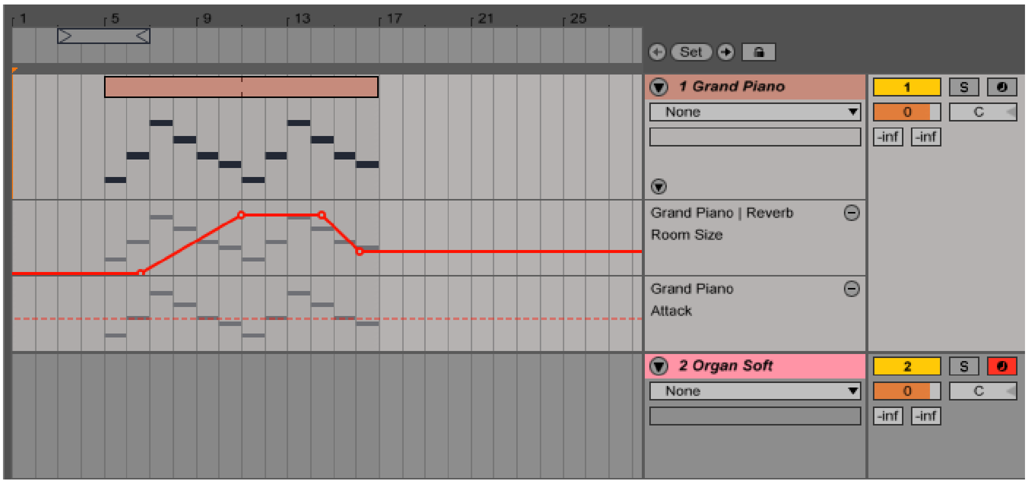

Figure 4.1: Viewing selected parameters of a track in *Ableton Live 9*

Systems that allow composers to view and manipulate a subset of aspects of the composition can assist composers who deliberately engage in *limited exploration* (see *Selective allocation of time and effort*). In the study in Chapter 3, Composer C started the session by using the tools in *Logic* to construct two representations, one focused on pitches and one on note velocities, and then spent the majority of the session focused on these representations. This may be evidence of them deliberately limiting their exploration to these aspects of the music.

Some research studies have created prototypes focused upon the activity of *mixing* by providing user interfaces designed primarily for the editing of stereo position and amplitude. These systems provide fixed mappings from these qualities to a diverse range of visual properties (Dewey and Wakefield 2016, 2017; Gelineck and Uhrenholt 2016). These studies show how representing musical attributes using different visual properties can help solve different questions about the composition: a two-dimensional "stage" metaphor makes it easier to answer questions like "which instruments are panned to the left-hand side?", while a virtual "mixing desk" metaphor makes it easier to find the current stereo position of a specified instrument. Some initial research has started to consider how these representations might be expanded so that composers could edit additional properties of the track, such as effects parameters (Mycroft, Reiss, and Stockman 2016). It is worth noting that the interfaces used in these systems usually do not represent the important dimension of time visually. This has the effect of focusing the composer's attention on a particular moment in time but might also lead to a lack of awareness of the musical context they are working within.

4.3.1.2 *In existing composition software*

Prior research involving composition software other than digital audio workstations involved creating systems that allow users to choose which aspects of a composition should be represented to them. *Scriva* (Buxton et al. 1979) can switch between four different representations that focus on different aspects of the composition being edited. *MusicSketcher* (Thiebaut, Healey, and Kinns 2008) provides a timeline on which the user can draw curves whose height is mapped to one of a small set of attributes of their choice (filter cutoff, pitch, granular synthesis parameters). Some more unusual variants of this technique exist. While *Comp-I* (Miyazaki, Fujishiro, and Hiraga 2004) maps musical attributes to a fixed *single* representation of music in 3D space, the composer may switch between several different 2D projections of the 3D space (e.g. "front view" or "top view"), effectively selecting a subset of these attributes that they wish to be displayed.

The most powerful implementations of this concept provide editable representations that are not displayed on a screen, but which are printed out to *paper*, and edited by drawing with an augmented pen. *MusInk* (Tsandilas, Letondal, and Mackay 2009) allowed users to map X and Y axes of hand-drawn two-dimensional shapes to attributes in *Open Music*. *Inksplorer* (Garcia et al. 2011) provided a similar system but where the axes could be mapped to any *OSC* parameter, in theory allowing control of any music software supporting *OSC*. *PaperComposer* (Garcia et al. 2014a) provided a variety of representations of melodic spaces, such as piano roll, tonnetz grid, and musical staves, that can be mapped to *OSC* parameters.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.1

- Provide sets of *common useful representations*, such as line graphs, points in a 2D plane, musical staves, and tonnetz grids.
- Let composers specify *their own mappings* from musical attributes they are working on to the visual axes provided by these representations.

4.3.1.3 *In other software categories*

Recent research in information visualisation has suggested new interfaces to allow users to easily specify mappings from data attributes to a broad range of visual properties using direct manipulation techniques without requiring any programming (Satyanarayan and Heer 2014). These kinds of interfaces could allow composers' on-screen representations of their work to more closely approximate the diverse representations seen in their use of paper. Mixed-initiative approaches that automatically suggest possible visualisations to users based on the user selecting aspects they wish to be represented (Wongsuphasawat et al. 2016) might further increase the ease with which composers could create new representations.

In the study, Composer C used a multi-monitor setup, and so was able to use one display to view specific aspects of the composition while keeping the full context visible on a secondary monitor. Several visualisation systems (reviewed in Roberts (2007)) investigate the space of how to allow users to coordinate multiple visualisations, and techniques from this area could be re-used within the DAW interface. It is important to note, as Wang Baldonado, Woodruff, and Kuchinsky (2000) warn us, that while providing multiple representations to users can provide cognitive advantages, it can also increase cognitive demands in some situations. Given the fact that the composers studied in Chapter 3 seem to all make coordinated use of multiple different representations, this would seem to suggest that features for *coordinating representations* are important, but that care should be taken to allow representations to be easily hidden when they are no longer useful.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.2

- Allow *coordinated* manipulations of existing DAW representations with these new representations.
- Provide a power-user interface for extending existing representations into additional *custom* representations.

4.3.2 Diverse media types

Most DAWs will allow the representation of common types of information about the composition (such as MIDI notes, effects parameters, and envelopes), but some prototype systems have explored using a broader range of media, including text annotations, photographs, and diagrams, in order to be able to represent additional types of information.

Observing the use that composers make of paper provides many clues about what types of information composers use. The previous chapters have described several examples of composers attempting to represent more unusual types of information. In Chapter 2, it was noted that Tsandilas, Letondal, and Mackay (2009) describe composers using extended musical notations for representing quarter and eighth tones, and Thiebaut, Healey, and Kinns (2008) provides detailed illustrations of custom geometric notations used by a composer to represent audio channel positions in a more sophisticated way than the conventional stereo panning model. In Chapter 3, additional evidence of note making on paper by composers was observed, including the representation by Composer A of "loose" rhythms and generalised intervals from music theory, and the representation by Composer D of graphical plans for a composition (see Figure 3.2).

At the most basic level, some composers report frustration with their DAW software lacking support for text lyrics (Bainbridge, Novak, and Cunningham 2010). McGrath, Chamberlain, and Benford (2016) note that while the representations (either electronic or paper) that are used outside of music tools provide ways of representing the logic and reasoning behind why choices have been made, this information is lost when the composer translates their work into software composition tools.

Nash (2015) notes that while composers require support for "secondary notations", most DAWs only support limited features such as "labelling and colour-coding parts and tracks, and free text… few mechanisms are provided for flexibly annotating the music in any of the sub-notations or views, beyond those forms formally recognised by the program".

4.3.2.1 Supported activities

In Chapter 3, *conceptual frameworks* were described as a mechanism that composers use for *self-constraining*. To represent these conceptual frameworks, composers create paper representations of them; for example, a composer might draw a *magic square* on paper, which they use to determine the phrasing of a piece (Letondal and Mackay 2007). This is a common behaviour for composers using conceptual frameworks: Thiebaut (2010) describes a composer who used series of paper sketches of geometric forms as "a medium between these concepts and the music" representing their conceptual framework, and in the study described in Chapter 3, Composer

D's working environment contained paper representations of conceptual frameworks used in previous compositions including a periodic table and a set of *tone rows*. To represent these (often visual) ideas *within* the DAW instead of on paper would require the DAW to support a broader range of media types such as freehand sketches, diagrams, and lists.

An important activity can be eased using representations that include diverse media types is *collecting and capturing* of material. When inspiration strikes, composers may wish to record ideas using one or more of audio, text, images (Bainbridge, Novak, and Cunningham 2010), chord symbols, tablature notation, or conventional music notation (Bainbridge, Novak, and Cunningham 2012). In Chapter 3, Composer B and Composer D were observed collecting snippets from previous work to incorporate into their new compositions, but were required to leave the DAW in order to do so, as there were no obvious facilities in their tools for collecting such materials.

4.3.2.2 In existing composition software

Research composition systems have attempted to include support for media types beyond those conventionally seen in commercial software. *Frameworks* (Polfreman 2001) supports text, pictures, and diagrams. *Apollo* (Bainbridge, Novak, and Cunningham 2010) permits freehand sketching, allowing users to make annotations. *Sonic Sketchpad* (Coughlan and Johnson 2006) allows compound objects to be created that consisted of both a drawing and recorded audio. *QSketcher* (Abrams et al. 2002) automatically generates and stores additional metadata about the context of composition (for example, the time at which each musical note was entered into the system). *Paper Substrates* (Garcia et al. 2012) takes an alternative approach, using augmented pens to link the custom hand-drawn paper notations created by composers to the more conventional musical notations in the computer, while *Symbolist* (Gottfried and Bresson 2018) allows similar custom graphic notations to be used within *Max* and *OpenMusic* (but not in conventional DAWs).

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.3

- Support *more media types* including text, photographs, drawings, and links to files.

- Record *metadata about the context* musical material was created in, such as date, time of day, and geographical location.

4.3.2.3 In other software categories

As described earlier in this section, *collecting and capturing* is an activity which could be enabled by support for a broader range of media types. This activity is commonly seen across creativity tools, with the review of frameworks for understanding creativity using software tools in Wang and Nickerson (2017) describing "collect potentially relevant information" as one of the four key steps involved in creative process. Coughlan and Johnson (2008) have discussed "collecting" of ideas as a common behaviour that can be seen across many different creative domains.

Research into the creative activities of designers by Keller et al. (2009) and Inie et al. (2018) found that assembling collections of material was a common activity, with designers frequently maintaining both physical and digital collections separately. Keller et al. (2009) suggest that a possible reason for this might be the poor support for visual information in computer tools, and suggests that their designers require the ability to easily create digital versions of their ideas by scanning or photographing. In the context of a DAW, this could include facilities for quickly scanning paper notes into the DAW or importing photographs from a digital photo library.

A common theme in this research is the capturing of ideas when away from the computer. Gross et al. (1998) created *Digital Design Sketchbooks* for mobile capture of ideas by architectural designers and Lee and Klemmer (2005) created the augmented paper *iDeas notebook* for design students to provide a portable mechanism for capturing ideas. In the study in Chapter 3, Composer C made use of musical material they had recorded away from their home studio, while Composer B had noted down ideas for compositions on their iPad, suggesting that mobile capture of ideas might also be relevant to composers.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.4

- Allow composers to *record ideas on the move* away from a desktop computer.
- Integrate the ability to *digitise notes*, e.g. scanning, capturing photos.

4.3.3 Structured representations

The ability to represent structures can be an important tool for composers, allowing the composition to be broken up into smaller sections that can be worked on individually, and enabling the composer to view the composition as a whole.

Musical structures are not just *temporal*: other useful ways of structuring compositions used by composers include grouping tracks by frequency, instrument type, or function (Duignan, Noble, and Biddle 2010; Macchiusi 2017, 140–41). Macchiusi (2017) describes existing features in commercial DAWs that can be used to encode such structures, such as setting colours on tracks and clips and creating groupings of tracks.

Previous authors have discussed how composers create representations of musical structures in their composition practices. Duignan, Noble, and Biddle (2010) describe watching composers who "worked with musical material at varying temporal levels, from short musical riffs… all the way up to large song structures such as the traditional verse, chorus". Tsandilas, Letondal, and Mackay (2009) describe how the paper notes used by composers reference temporal structures and give examples of structural diagrams created by composers.

Eaglestone et al. (2007) stress the importance of the ability to "express free associations between data and tools", particularly *personally meaningful* associations, and described how one composer they studied even created their own software tools to allow them to represent groups of sounds hierarchically.

Representations allowing the composer to represent *structure* can allow the composer to focus on *structural* aspects of a composition. Pearce and Wiggins (2002) argue that creative cognition is "supported by the ability of the composer to represent and process musical information in a hierarchical manner and to attend to the more abstract levels of representation during composition", and that expert composition requires this facility. Similarly, Balaban (1996) argues that "a good representation for structured music pieces must capture temporal hierarchies".

Some of the alternative DAW interface wireframes illustrated in Koda (2011) suggest a possible route for higher-level temporal structures within DAWs, but they were not implemented in a prototype system.

4.3.3.1 Supported activities

Representations which can represent structure be useful for *top-down* composition processes:

> "many computer based systems do not support the top down strategy for composition adequately and in fact, they usually require a bottom up process. The 'top down' composer in this case generally has to develop the high level structure externally to the system, build up the material using the program while manipulating this material to fit the required structure" (Polfreman 1997)

Composers who work in this fashion could make use of representations of structures as part of their *planning* activities. In Chapter 3, Composer A started to work in this kind of top-down way by making use of the pre-existing structure of the text of a poem to provide the structure for a composition that they then composed.

This structure also appeared to help Composer A later on, when they shifted to working in a *patchwork* style. For composers who use *patchwork* composition processes, in which they do not work linearly from the start of the composition to the end, representations of the temporal structure may be useful. Composer B moved forward and back between different sections of their composition during the session that was observed, and when they returned to an earlier section to extend it in length, they found doing so to be difficult. Explicit support for temporal structures, rather than just a timeline, may have helped to make structural manipulation more frictionless.

4.3.3.2 In existing composition software

Temporal structures that go beyond the simple two-level (notes arranged into clips) model used in DAWs have been used by some systems within research into other types of composition software. While more complicated models of the temporal structure of compositions are more common in the field of *generative* music software (Polfreman 2001; Smith 2011), only a few conventional composition systems represent more complex temporal structures, such as *Apollo* (Bainbridge, Novak, and Cunningham 2010) and *SSSP* (Buxton et al. 1978).

It may initially appear that the solution to representing structure is to allow the composer to specify a hierarchy. However, Dannenberg (1993) points out that "music often contains multiple hierarchies… A single hierarchy system is inadequate to represent all these concepts at the same time." Polfreman (1997) echoes this concern, stating that

> "systems should effectively support the definition and manipulation of multiple parallel hierarchical structures, in order to allow composers to manipulate musical information within the system in a manner closely modelling their conceptual view of the music".

Dannenberg (1993) suggests two ways to support multiple hierarchies: named links and tags. However, these remain mostly unexplored in composition systems, with *Apollo* (Bainbridge, Novak, and Cunningham 2010) being notable for being a rare example of using hyperlinks to represent temporal structure.

Some commercial composition software provides support for tagging in their built-in libraries of loops, instruments, and effects. While the instrument browser in *Ableton Live* uses a hierarchical model, for example, *Maschine 2* uses the named tags approach. *GarageBand 10* provides an interface for browsing audio loops that provides a choice between a conventional hierarchy and named tags (see Figure 4.2). However, use of these approaches in other structures within the DAW, including temporal structures, remains unexplored.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.5

- Use a *hyperlinking or tagging mechanism* to implement structures, so that multiple overlapping structures can exist simultaneously, and multiple levels of hierarchy can be created easily.

4.3.3.3 In other software categories

Across many different creative fields, support for creating multiple different levels of structure could potentially aid creators (Shneiderman 2000), but perhaps the most similar field is video editing. Video editing software also involves a timeline with multiple tracks on which material can be edited and played back.

While DAWs have mostly taken an approach of grouping individual notes into "clips",

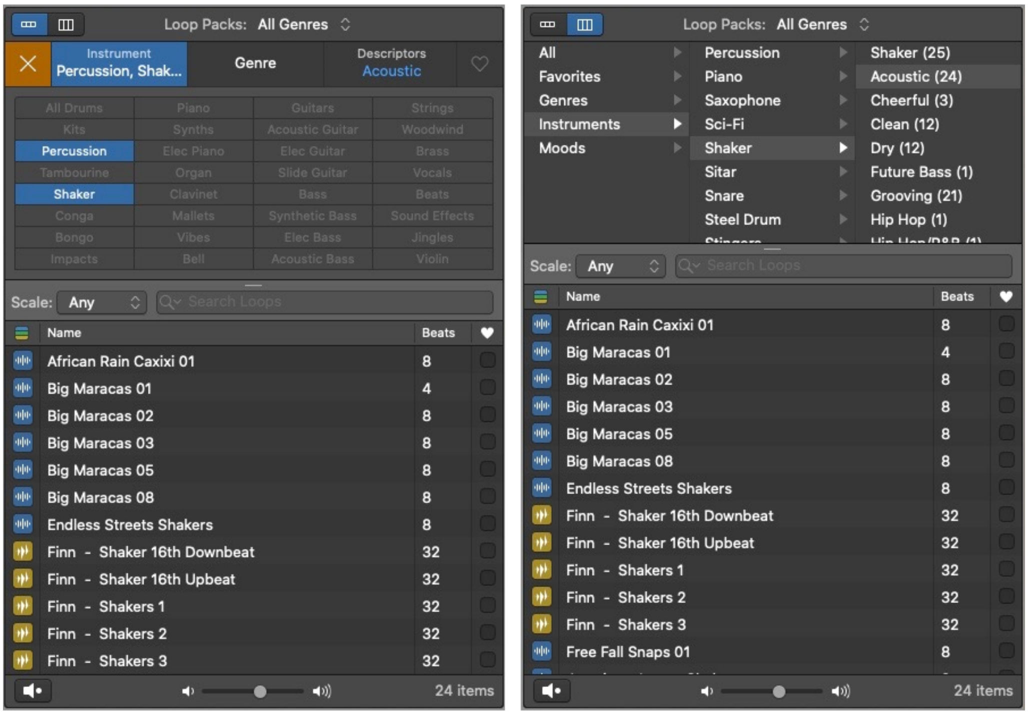

Figure 4.2: Modes of the loops library in *GarageBand 10* that use tags (left) and hierarchy (right)

research involving software has experimented with more levels of hierarchy, and additional view of temporal structure in addition to the timeline. *Silver* (Casares et al. 2002) groups video frames into "shots", which in turn are grouped into "clips". In addition to the traditional timeline, video editors can also edit the temporal structure using the "storyboard" view, which displays clips as an ordered list. This allows easier editing of the high-level structure without needing to view or edit lower-level shots or frames. However, *Silver*'s interface only supports one hierarchy, rather than the multiple parallel hierarchies recommended for music software by Polfreman (1997).

An alternative approach that has been used to represent temporal structures in video is annotations as an alternative to a more formal hierarchical structure. Annotations can be used to provide multiple different perspectives, allowing multiple parallel temporal structures to be created (Costa, Correia, and Guimarães 2002), similarly to the tags and hyperlinks described in the previous section. Interfaces implementing this kind of annotation in video have been used in prototype systems including *TAV* (Müller, Miller, and Fels 2010).

Several research systems have also allowed the timeline to be annotated with transcripts of the video (Berthouzoz, Li, and Agrawala 2012; Casares et al. 2002), and provide a separate view of the transcript that can be interacted with to navigate to locations within the timeline. A similar approach could be used with lyrics in DAWs. In Chapter 3, Composer A used an annotated lyrics sheet to represent the structure of their composition, so this approach could be more natural for composers whose structure is externally imposed in this way.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.6

- Provide *additional temporal representations* in addition to timelines that allow rearrangement of higher-level structures.
- Remove the need to define a formal structure by allowing structural *annotations* for browsing.
- Allow the timeline to be *annotated with lyrics*.

4.3.4 Incomplete specification

Nash (2012) notes that in most composition software, the representations used often assume that the composition will be represented in a form that is completely specified and can be played back by the computer, and that

> "software is thus often exclusively able to support the final stages of creativity: the elaboration, verification or refinement of an idea, and take the 'a-ha' moment of insight as having already happened… Music programs that seek to support exploratory design are often used simply for transcription".

To work around this restriction, composers will often make external representations to describe the composition at a lower level of detail or fidelity to the finished product:

> "notebooks, cameras and voice recorders are used… because initial representations can often be made without requiring the specific qualities of the expected outcome." (Coughlan and Johnson 2009b)

Allowing composers more media types to record musical ideas allows them to capture rough sketches (see *Diverse media types*). However, unlike other material in the DAW, those representations do not support *auditioning* which is a significant drawback, as a key advantage of composition software over paper is that the composer can listen to their composition at any time.

This suggests that representations in a DAW should support *incomplete specification*: that even if they do not yet contain complete instructions on how to play back the composition, the DAW should attempt to do so. While Bainbridge, Novak, and Cunningham (2010) note that in DAWs "musical ideas must be entered with an attention to fine details of performance that is not appropriate to the capturing of nascent musical ideas", and Balaban (1996) suggests that systems should support "partially-specified" and "incomplete" musical elements, information about how incomplete specification might actually be implemented in composition software appears to be mostly absent in the literature.

One exception is Buxton et al. (1978), who suggest that requirement that a music system should be "capable of coping with incompletely specified data" and describe an approach for providing this functionality. Their approach separates the process of specifying a composition into different activities (such as defining a palette of timbres, entering notes, or assigning notes to specific instruments). They propose that by using a "system of defaults", composition systems could support auditioning when only some of these activities had been carried out, regardless of the order in which those activities are performed.

4.3.4.1 Supported activities

Support for incomplete specification would be useful for supporting the activity of *auditioning* at earlier stages in the creation of compositions. Peterson (2008) argues that allowing auditioning of incompletely specified material may have disadvantages, suggesting that interfaces allowing playback of incomplete music with only one aspect entered (such as notes without expression markings) might encourage an undue focus on that aspect of the music. However, based on the findings in Chapter 3 (see *Selective allocation of time and effort*), we can also look at this as a positive effect: auditioning of incompletely specified material can assist composers in deliberately *limiting exploration* and focusing on those aspects that are of interest to the composer.

4.3.4.2 In existing composition software

Based on their experience of working with composers who use paper scores, Tsandilas, Letondal, and Mackay (2009) created *MusInk*, a composition system in which visual elements within the representation of the composition are not required to have specific musical meanings, and visual objections can be created that do not describe any particular pitches, effect parameters, instruments, or other information required for playback.

Their research illustrates an important point: that to support the ability to audition an incomplete composition, the software should not just allow the auditioning of symbolic elements which are incompletely specified, but also tolerate the presence of symbolic elements so incompletely specified they *cannot* yet be auditioned at all.

Describing this approach as "semi-structured delayed interpretation", the approach used by Tsandilas, Letondal, and Mackay (2009) depends on two features: the ability to add symbolic elements to the representation of a composition without specifying a musical meaning, and the ability to add a defined meaning to those elements at a later time. Their implementation of the approach uses digital pens with which composers draw onto paper representations, allowing composers to sketch envelopes, groups, ranges, or markers onto a printed score, but the approach itself need not be specific to that medium and could be applied in DAWs in future.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.7

- *Do not require musical interpretation* to be specified when envelopes, groups, ranges and markers are created.
- Allow *mappings at a later time* from these objects to a musical property.

4.3.4.3 In other software categories

In the more general domain of creativity support tools, many authors have identified the ability to "sketch" incompletely specified representations as a useful ability. As Mangano et al. (2014) observe, designers "draw what they need, and no more. Few sketches are created with extensive detail; rather, designers create sketches with the detail and notation necessary to help them reason." Gharib (2014) reports that designers avoid CAD systems because they require "a high level of accuracy" while "sketching needs an easy, fast, and intuitive way to express ideas that arising quickly in mind… the designer uses sketching to record his ideas quickly", and Mangano et al. (2014) agree that "low detail enables sketches to be created quickly and modified easily, providing rapid feedback". An advantage of sketching is the ability to easily try many ideas by giving more flexibility for experimentation, avoiding *premature commitment* to particular ideas. Mangano et al. (2014) note that adopting a formal notation too quickly results in "less exploratory and broad search for solutions" and Walther, Robertson, and Radcliffe (2007) claim that too much detail results in "premature fixation" on certain solutions.

As discussed earlier in this section, Buxton et al. (1978) suggest that composition software should support incomplete specification by using a "system of defaults" for missing data. However, research in other fields of software has suggested that providing default settings can have unwelcome effects. The presence of a default option in a user interface can significantly affect behaviour through the "default effect" (Herrmann et al. 2011; Shah and Kesan 2006), steering users subconsciously towards the default option in ways that the user is often unaware of. Nicoll and Keogh (2019) describe how when small teams and hobbyists are designing games using the engine Unity, their need to prioritise their use of resources and time can lead to overuse of default values, leading to games that feel derivative and unprofessional. As a result, it appears using default options for incomplete data does present a danger of the choice of default values influencing the composer's work.

The current approach commonly used to mitigate this issue within composition software is a library of presets, which can remove the need for spending time and energy on fully specifying a choice by allowing the composer to select an acceptable pre-existing option. However, this may still result in derivative or unprofessional sounding results (Paterson 2011), and some creativity support researchers have suggested that a large variety of presets can be overwhelming for less experienced users (Benedetti et al. 2014). Experience from the field of game development suggests that even when presets are intended to be used only as temporary placeholders that are later replaced they may still cause issues such as creators becoming overly attached to the placeholder content, accidentally retaining the placeholder, and difficulty in replacing the placeholder with an item of better quality. These issues can be partially mitigated by clearly indicating placeholder content in the interface (Zagal and Altizer 2015).

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.8

- Clearly *indicate temporary work* that composers intend to complete later.

4.3.5 *Representing alternatives*

As Coughlan (2009) note, “few painters would ruin a canvas they valued just to explore an alternative idea, but with computers the cost of repairing a mistake has commonly been reduced to a keyboard shortcut.” Composition software offers the possibility of representing *multiple* possible versions of sections of a composition simultaneously.

Coughlan and Johnson (2006) suggest that composition software should support the representation of multiple possibilities for comparison, evaluation, and generation of further ideas. Duignan (2008) reports observing composers creating workarounds that allow them to represent multiple possibilities in a DAW. These composers would make a copy of a track’s current state and store it on a temporary muted track so they could safely experiment with other possibilities. Other composers found being required to store multiple “raw takes” as separate tracks to be confusing and resorted to storing alternate versions in an entirely separate copy of their project. McGarry et al. (2017) similarly report observing frequent duplication of files in the recording sessions they observed; again, this workaround is used to reduce risk and to allow safe experimentation. Macchiusi (2017) points out that some DAW software (*Pro Tools*, *Logic*, and *Cubase*) provides some limited support for this technique by allowing users to record multiple takes of a performance, an ability that is useful for common recording techniques of “comping” and “punching” (Phillips 2010, 257–64). Even in DAWs that do not directly support this behaviour, composers can use workarounds such as grouping a set of tracks each containing a different take (see Figure 4.3).

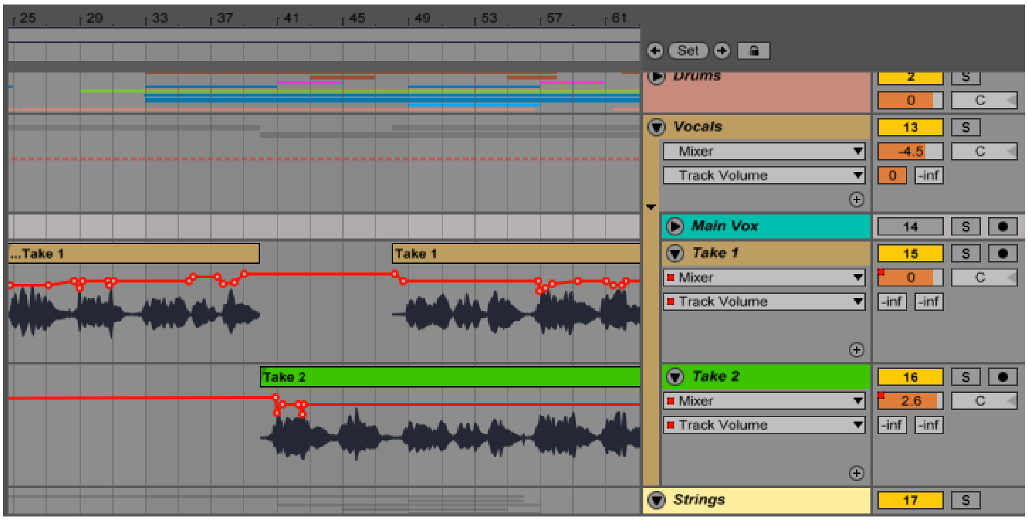

Figure 4.3: 'Comping' using grouped tracks in *Ableton Live 9*

4.3.5.1 Supported activities

It has been suggested that representing alternatives may be important for supporting *iterative* musical processes. Hailpern et al. (2007) suggest a requirement for creativity tools to "keep multiple design ideas visible simultaneously" because it allows a "rapid cycle of divergent/convergent thinking", suggesting that fast iterative creative processes may be aided by the ability to represent alternatives. In terms of the activities described in Chapter 2, the ability to represent alternatives helps in the activity of *combining alternatives* used in the convergent phase of iterative processes.

Continuing to represent alternatives after they have been discarded may be useful for composition processes that involve recycling material. Donin and Theureau (2007) describe how the composer in their study "recycles" material written for a previous movement, which was not judged appropriate for that context, by including it in later movements. Keeping a library of material which was considered "good but not the right option" may allow composers to quickly solve later creative problems by recycling previously abandoned work. This suggests that representing alternatives continues to be useful, even after a decision about which alternative to use has been made. This kind of recycling behaviour does appear to be relevant to pro-amateur composers: in Chapter 3, both Composer B and Composer D were observed making use of fragments of previous work in the compositions they were creating.

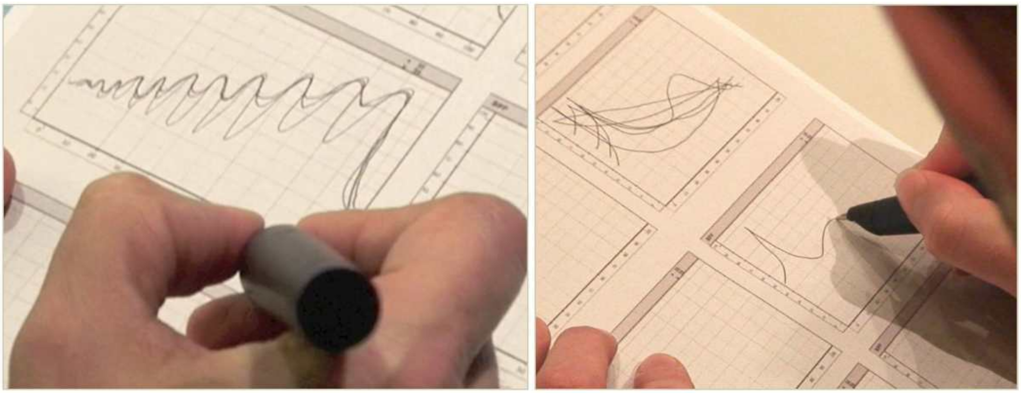

Figure 4.4: Drawing multiple possible expressive curves with *Inksplorer*, taken from Garcia et al. (2011)

4.3.5.2 *In existing composition software*

Garcia et al. (2011) describe how composers using their augmented paper system *Inksplorer* used the interactive paper to explore ideas by drawing multiple possibilities (see Figure 4.4).

> MM drew several long curves on top of each other to evaluate different alternatives in the afore-mentioned composition, each providing incremental corrections… He used layers of curves to guide each refinement, explaining "It's a kind of guide that lets you correct it next time." (Garcia et al. 2011)

This approach was further developed in a subsequent system by the same researchers, *Polyphony* (Garcia et al. 2014b).

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.9

- Allow composers to create multiple *alternative versions* of elements within the composition.
- Where values are graphically displayed and edited, *display multiple versions* simultaneously.

4.3.5.3 *In other software categories*

The need to represent multiple alternatives is a recognised issue in creativity tools. Terry and Mynatt (2002) state that the prevalent interaction model that "requires a document to be in one, and only one, state at any particular time… is typically a poor match to the non-linear, experimental processes characteristic of creative endeavours", and suggest creativity tools should support the ability to represent multiple possibilities at once.

The existing research into representations of multiple alternatives focuses almost exclusively on graphical media such as bitmap editing and 3D modelling. These mediums have the advantage that alternative options can be easily previewed side by side and the effects of changes can be made instantly visible. Musical compositions, on the other hand, can only be played back one alternative at a time, and the composer must spend time auditioning each alternative. Mediums such as music composition, which are much less visually and more temporally focused, are rarely discussed in the existing work on exploring alternatives. A notable exception is Hartmann et al. (2008), which briefly discusses the exploration of alternatives in parameters affecting the behaviour of interactive user interface elements.

Some applications of alternatives to specific problems seen in existing papers could potentially translate well to the DAW. Allowing composers to create alternative orderings of structural sections could be implemented using the approaches that have been used for this task in document writing (Elkhaldi and Woodbury 2015) and storyboarding (Tharatipyakul et al. 2016) software. Signal flow between instruments and effects in a DAW is comparable to a dataflow graph, so the approaches to exploring alternatives in dataflow graphs that have been used in generative visual art (Zaman et al. 2015) and game development (MacCormick and Zaman 2019) might be helpful.

While the ability to represent multiple options *within* a composition is important, it is important to note that this is not the only way that composers to achieve their aims. Discussing the need to experiment with alternative options in creative software, Terry et al. (2004) identified four methods for doing this: throwaway sketching and prototyping of alternatives in different documents, history and version control for keeping copies of past considered versions, using ad-hoc workarounds within a document (for example, using multiple layers and turning them on and off), and "what if" tools that represent multiple versions of a document. For early stages in creative processes, throwaway prototypes may be sufficient: Mangano et al. (2014) describes how designers use multiple sketches representing alternatives that they can then juxtapose and synthesise. Within the context of a DAW, the ability to link together multiple sketch compositions as one "project" could assist composers in this behaviour.

As yet unexplored design possibilities for this technique include:

Design suggestions 4.10

- Provide the ability to create multiple *alternative orderings of temporal structures*.
- Allow composers to *create projects* that group throwaway experiments together with a composition.
- Consider exploring *alternatives within signal flow chains* of effects and instruments.

4.3.6 *Task lists*

Lists of tasks are a type of representation used to record planned work on a composition and monitor the progress of the composition activity over time. Composers might call these representations by various names such as checklists, to-do lists, or just 'reminders'.

Abrams et al. (2002) report that "composers often cover their desk or walls with… todo notes" and Nash (2015) describes the use of secondary notations such as freeform notes for a variety of purposes including to-do lists. In Chapter 3, Composer C was seen to use the plain text notes feature in *Logic* to maintain lists of instrumental recordings that needed to be re-recorded.

4.3.6.1 *Supported activities*

The most obvious use of task lists is as a memory aid, and in the literature on external representations, to-do lists are often an archetypal example used to explain reducing memory load by cognitive offloading. In terms of the composition activities described in Chapter 2, task lists are a representation that could support composers in *planning* within the DAW, without requiring them to link plans to specific positions in the timeline.

In the context of composition, they could also help composers to retain focus by supporting review of the composition and prioritisation of particular aspects of it, supporting the strategy of *limited exploration* of other aspects.

4.3.6.2 In existing composition software

Even though the use of task lists by composers has been documented, and the apparent utility of task lists across a wide variety of domains and groups of users outside of composers, there does not appear to be any existing discussion of how task lists might be integrated with DAWs or other composition software.

4.3.6.3 In other software categories

Outside of the field of creativity support tools, the use of task list representations is described using the terminology *personal task management* (PTM) or *electronic personal task management* (e-PTM). Many potentially relevant design suggestions can be found in this literature including sorting, filtering, highlighting, and grouping by properties of tasks (Bellotti et al. 2004, 740; Haraty and McGrenere 2016), automatically suggesting sub-tasks (Gil et al. 2012; Kokkalis et al. 2013), a single text field for both search and input of tasks to avoid accidental duplications (Conley and Carpenter 2007), capturing history of edits to the task list (Bellotti et al. 2004, 740), support for entry away from the desk including use on mobile devices (Bellotti et al. 2004, 740), and automatic detection of when tasks have been completed (Gil et al. 2012).

A potential pitfall of postponing work is that composers need to re-familiarise themselves with the task and its context when they resume work. Salvucci (2010) view "task resumption as a process of reconstruction. In reconstruction, the user visually re-encodes the task environment to reconstruct the task context immediately prior to interruption." This suggests that providing the facility for users to restore the representations they were using when a problem was discovered and added to the list may be helpful. As Rule, Tabard, and Hollan (2017) note:

> "images of past work may help users not only remember suspended activities, but also cue reconstruction of the complex network of thoughts they had while performing them, easing resumption of those activities… Helping users make sense of past activities is an essential step in supporting the reconstruction of past mental contexts and ultimately

resuming suspended activities."

Finally, it is worth noting that records of problems and mistakes can be useful for many different and sometimes not immediately obvious reasons, as Kim, Bagla, and Bernstein (2015) suggest. Lists of mistakes can be used as inspirational material for future work (Eaglestone et al. 2007) or categorised to spot common patterns of errors.

As yet unexplored design possibilities for this technique include:

Design suggestions 4.11

- Provide a *list or database* facility for storing tasks.
- Let composers easily *restore representations* they were using when a task was created.
- Consider *solutions from PTM literature* to find useful interface patterns.

4.3.7 Representing history

In addition to representing plans for future activity, it is also helpful to be able to represent past composition activities and states of the composition.

As has been previously described in Chapter 2, composers often make use of paper in varied ways as part of their composition processes. Past authors have documented, using photographs of the paper notes used by composers, the way the permanent nature of these physical notes creates a representation of the historical development of a composition over time (Healey and Thiebaut 2007; Garcia et al. 2011; Rutz 2014, ch. 4).

With the proliferation of digital studio tools, composers have come to rely on the ability to use "undo" features as a similar mechanism by which they can revisit historical states of the composition, a significant change from previous eras of recording technology in which non-destructive editing was not available (Macchiusi 2017, 68; Phillips 2010, 266).

However, undo/redo features are limited in the ways that they can be used:

> "Widely used multi-track editing systems, such as Protools or Ardour, only provide support for the standard desktop application metaphor

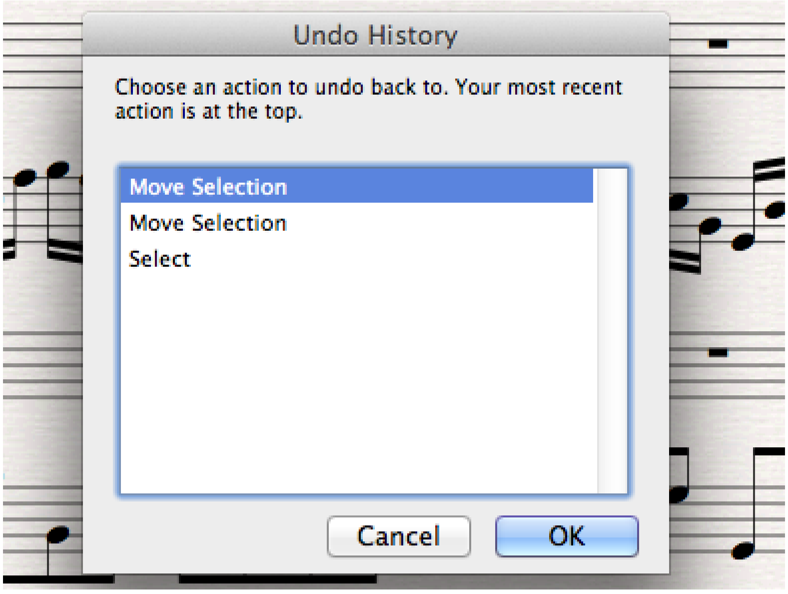

Figure 4.5: The 'Undo History' feature in *Sibelius 5*

> of an undo-history… The purpose of the undo-history, as its name suggests, is merely to be able to correct recent mistakes by making the last editing steps undone." (Rutz 2014, 44–45)

Designed as a specialised tool intended purely to remedy errors, these undo history facilities only display a list of past *actions* that have been carried out, not past *states* of the composition (for example, the interface in Figure 4.5). An undo history appears limited when compared to the detailed visual record of a composition provided by paper notes. Even if composers do attempt to return to previous states within the undo history, the state of the entire composition is altered, making it difficult to identify actions taken on individual elements of the composition.

4.3.7.1 Supported activities

Representations of historical activity would be particularly useful for supporting the *habituation* strategy observed in the study in Chapter 3 (see *Selective allocation of time and effort*). In that chapter, Composer D used a habituation approach, re-using the same techniques repeatedly. However, due to the constraints of the software they were using, they had to repeatedly locate and reselect the same options in the *Cubase* menu system, rather than being able to easily able to access them. Duignan (2008) observed that composers were so keen to "fine-tune or recreate the same process at a later point" that they would document the history of a process. Koda (2011) notes that "a simple listed history of processes applied to individual audio regions would be immensely useful in documenting the concoction of processes that might create similar audio effects." Another potentially useful effect of displaying habitual behaviours is to allow composers to avoid those behaviours to ensure novelty in their work (Nakamura et al. 2018). Outside of the specific context of composition, Nancel and Cockburn (2014) summarise some of the existing research into interaction histories designed to enable reuse.

Access to historical interaction data could also be a powerful tool for helping composers in the activity of *reflection*. *iMPULS|IVE* (Nash 2012) and *Omaggio* (McCulloch 2014) are tools designed to visualises the history of interactions with specific composition software packages. While these visualisations were intended solely for use by researchers, they suggest useful properties that can be derived from composition session data such as "distraction events" and "uncertainty", and these might also aid composers in their own reflective activities.

4.3.7.2 In existing composition software

Rutz (2014) introduces *Mellite*, a music programming environment that captures a full historical record of past states, and art installations which partly created using the tool. However, while they describe in detail how the software was implemented, and the novel creative avenues it has been used to explore, there is little discussion of how the historical information it records might be represented to the composer.

Sonic Sketchpad (Coughlan and Johnson 2006) features a slider control which is used to access all the previous states of the composition, but has an unusual design quite different to a DAW, with no timeline or tracks.

QSketcher (Abrams et al. 2002) provides an undo history feature that allows a composer to view and restore previous configurations of windows and tool palettes the composer has used.

Sonic Zoom (Tubb and Dixon 2014b) is an interface for controlling a synthesiser that visually represents the path the composer has taken while exploring the space of possible timbres. The *Sonic Zoom* system represents historical parameter values as a white trail in a 2D space (see Figure 4.6).

In the *Sonic Zoom* interface, blue circles used to represent saved presets. Wexelblat and Maes (1999) note that history can contain both *active* and passive elements,

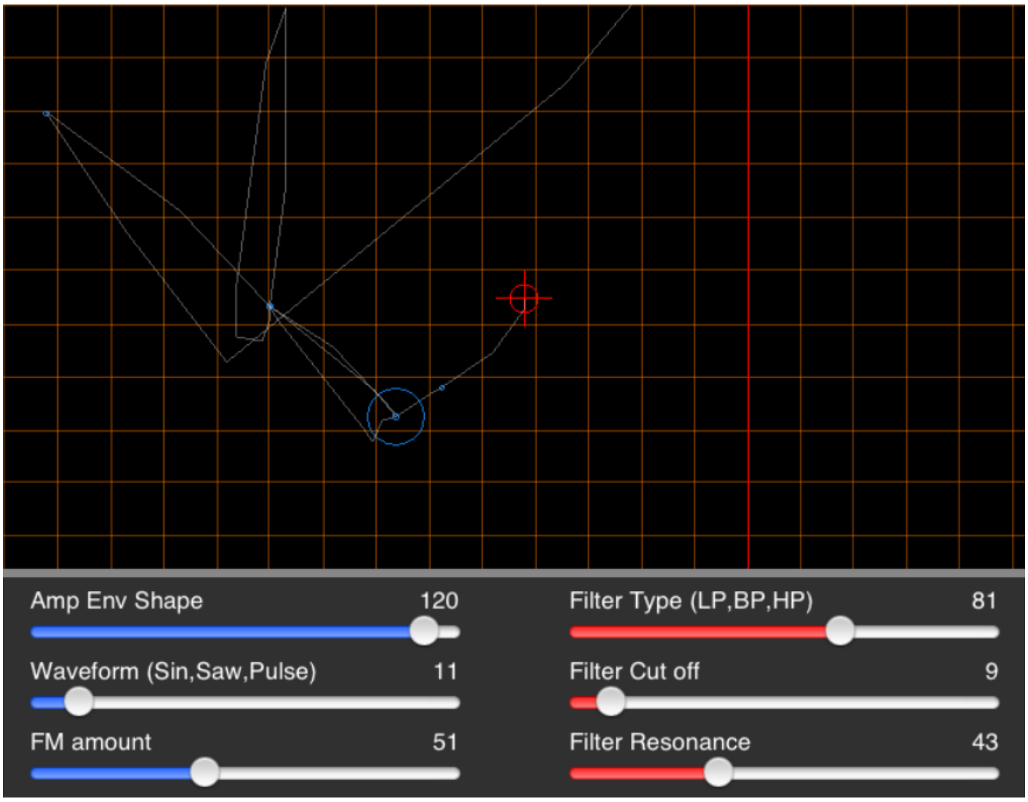

Figure 4.6: Part of the *Sonic Zoom* interface, taken from Tubb and Dixon (2014b)

contrasting the automatically recorded passive history used by Undo and Redo features with the actively created bookmarks and favourites that are deliberately selected by users. In *Sonic Zoom*, displaying and highlighting presets in the same display as historical values integrates both active and passive histories into a single representation.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 4.12

- Allow the user to *revisit past states* of the composition.
- Display *historical values in visualisations* of parameters in instruments and effects.
- *Highlight saved custom presets* in these visualisations.

4.3.7.3 In other software categories

The ability to trace the history of the creative process has identified as an important aspect of creativity support, with guidelines for creativity support tool design such as Resnick et al. (2005) recommending designers to consider it.

A similar approach to the previously described *Sonic Zoom* interface, in which previously explored settings are visually displayed, has also appeared in other fields. Feng et al. (2017) have discussed how the exploration of information visualisations could be enhanced by displaying what areas have already been encountered. They review existing approaches to visualising the history of a user's exploration, focusing on systems that directly encode this history as a *visual variable* of a visualisation, and report that the prototype systems they created to explore this area appear to "nudge" users to perform more exploration.

While *Sonic Zoom* only shows the history of explored values within a visualisation of continuous parameters, Feng et al. (2017) note that many other interfaces for exploration could be enhanced in a similar fashion, including some standard user interface widgets found in DAWs. Controls for setting parameters such as knobs, sliders, and drop-down boxes could be augmented with visualisations of history using the approach of *Scented Widgets* (Willett, Heer, and Agrawala 2007). Views of the composition that allow panning and zooming can display also use a similar technique to display *visit wear* in scrollbars (Skopik and Gutwin 2005). Interfaces for browsing and choosing items from lists can benefit also from a similar approach (Gutwin and Cockburn 2006). Another area that should be considered is the re-use of items from within hierarchical structures such as loop, instrument, or effect libraries. Existing interfaces designed to assist users in revisiting parts of the file system hierarchy may be useful guides here. The approach of highlighting previously opened file system folders used by Fitchett, Cockburn, and Gutwin (2013) might be useful for composers like Composer B in Chapter 3, who was observed spending considerable time trying to relocate a previously used category folder within the *GarageBand* loops library.

A final consideration is whether the record of a user's historical activity should be carried over between different compositions. Doing so would support users like Composer B (from Chapter 3), who reused options between multiple songs within an album to obtain a consistent feel.

As yet unexplored design possibilities for this technique include:

Design suggestions 4.13

- Use *augmented user interface elements* that display the history of previously viewed and used options.
- Show *historical choices from other compositions* in addition to the current composition.

4.4 Summary

Chapter 2 described component *activities* that are seen in composition. In this chapter, a number of techniques for designing representations in digital audio workstations have been described, and many of these techniques help support the composition activities that were previously described. This is summarised in Table 4.1.

Table 4.1: Composition activities from Chapter 2 that are supported by the representation techniques

Activity	Techniques supporting the activity
Planning	Structured representations, Task lists
Collecting and capturing	Diverse media types
Auditioning	Incomplete representation
Reflection	Representing history
Improvisation	Selective representation
Mixing	Selective representation
Combining alternatives	Representing alternatives

Chapter 3 described three *strategies* used by composers to selectively allocate their time and effort, and these techniques are also supported by several of the techniques, as is summarised in Table 4.2.

Table 4.2: Allocation strategies from Chapter 3 that are supported by the representation techniques

Strategies	Techniques supporting the strategy
Limited exploration	Selective representation, Incomplete representation, Task lists
Conceptual frameworks	Diverse media types
Habituation and reuse	Representing alternatives, Representing history

Based on the study in Chapter 3 and sources from the literature, this section has made specific suggestions on how each technique should best be implemented, and these suggestions are summarised in Table 4.3.

Table 4.3: Summary of the implementation suggestions for each design technique

Technique	Based on composition software	Based on other software categories
Selective representation	Provide sets of *common useful representations*, such as line graphs, points in a 2D plane, musical staves, and tonnetz grids. Let composers specify *their own mappings* from musical attributes they are working on to the visual axes provided by these representations.	Allow *coordinated* manipulations of existing DAW representations with these new representations. Provide a power-user interface for extending existing representations into additional *custom* representations.
Diverse media types	Support *more media types* including text, photographs, drawings, and links to files. Record *metadata about the context* musical material was created in, such as date, time of day, and geographical location.	Allow composers to *record ideas on the move* away from a desktop computer. Integrate the ability to *digitise notes*, e.g. scanning, capturing photos.
Structured representations	Use a *hyperlinking or tagging mechanism* to implement structures, so that multiple overlapping structures can exist simultaneously, and multiple levels of hierarchy can be created easily.	Provide *additional temporal representations* in addition to timelines that allow rearrangement of higher-level structures. Remove the need to define a formal structure by allowing structural *annotations* for browsing. Allow the timeline to be *annotated with lyrics*.
Incomplete specification	*Do not require musical interpretation* to be specified when envelopes, groups, ranges and markers are created. Allow *mappings at a later time* from these objects to a musical property.	Clearly *indicate temporary work* that composers intend to complete later.
Representing alternatives	Allow composers to create multiple *alternative versions* of elements within the composition. Where values are graphically displayed and edited, *display multiple versions* simultaneously.	Provide the ability to create multiple *alternative orderings of temporal structures*. Allow composers to *create projects* that group throwaway experiments together with a composition. Consider exploring *alternatives within signal flow chains* of effects and instruments.

Technique	Based on composition software	Based on other software categories
Task lists		Provide a *list or database* facility for storing tasks. Let composers easily *restore representations* they were using when a task was created. Consider *solutions from PTM literature* to find useful interface patterns.
Representing history	Allow the user to *revisit past states* of the composition. Display *historical values in visualisations* of parameters in instruments and effects. *Highlight saved custom presets* in these visualisations.	Use *augmented user interface elements* that display the history of previously viewed and used options. Show *historical choices from other compositions* in addition to the current composition.

This chapter has identified existing composition software which have explored applying one or more of the techniques. These are summarised in Table 4.4.

Table 4.4: Composition software that has applied each design technique

Technique	Relevant systems
Selective representation	*Comp-I* (Miyazaki, Fujishiro, and Hiraga 2004), *Inksplorer* (Garcia et al. 2011), *MusicSketcher* (Thiebaut, Healey, and Kinns 2008), *MusInk* (Tsandilas, Letondal, and Mackay 2009), *PaperComposer* (Garcia et al. 2014a), *Scriva* (Buxton et al. 1979)
Diverse media types	*Apollo* (Bainbridge, Novak, and Cunningham 2010), *Frameworks* (Polfreman 2001), *Paper Substrates* (Garcia et al. 2012), *QSketcher* (Abrams et al. 2002), *Sonic Sketchpad* (Coughlan and Johnson 2006), *Symbolist* (Gottfried and Bresson 2018)
Structured representations	*Apollo* (Bainbridge, Novak, and Cunningham 2010), *SSSP* (Buxton et al. 1978)
Incomplete specification	*MusInk* (Tsandilas, Letondal, and Mackay 2009)
Representing alternatives	*Inksplorer* (Garcia et al. 2011), *Polyphony* (Garcia et al. 2014b)
Task lists	
Representing history	*iMPULS|IVE* (Nash 2012), *Mellite* (Rutz 2014), *Omaggio* (McCulloch 2014), *QSketcher* (Abrams et al. 2002), *Sonic Sketchpad* (Coughlan and Johnson 2006), *Sonic Zoom* (Tubb and Dixon 2014b)

4.5 *Next steps*

For the purposes of this thesis, it is important to now determine a specific representation technique or technique to investigate further. When considering this issue, various factors are relevant to the decision, including:

1. how much could the technique support the specific requirements of pro-amateur composers?
2. how much existing research has already been carried out relating to the technique?
3. how technically feasible is a further investigation of the technique?

In Chapter 3, the particular needs of pro-amateur composers were identified. Pro-amateur composers have a need for better support for some composition activities, and all of the techniques appear to be useful as a means of better supporting at least one type of activity that composers perform. Pro-amateurs also make use of multiple secondary representations in addition to the composition itself, and the techniques of *structured representations* and *task lists* would allow further explanation of this theme. Strategies for selective allocation of time and effort found to be important to pro-amateur composers, but while most techniques could be used as part of those strategies, *structured representations* might be less useful.

While the first three techniques have been prototyped and studied to at least some extent, there has been little investigation of the remaining five techniques, with the *task lists* technique implemented in no prototypes found in the review. Representations using the *task lists* technique are discussed in very limited ways both in the reviewed sources and in the creativity support tools literature despite the apparent usefulness of the technique, suggesting that the *task lists* technique is under-explored.

Different techniques present different challenges in terms of the technical complexity required for further investigation. After building some initial prototypes, it was determined that building a DAW that could implement one of the techniques would be significantly more difficult than building an additional interface that could be used to interact with an existing DAW. This meant that techniques that could require significant modifications to the data model used by DAWs, such as *representing alternatives* or *incomplete specification* techniques, would present more challenges.

Additionally, these prototypes demonstrated that the APIs which DAWs provided for extensions did not provide control of all aspects of the DAW, and when used to control the DAW could involve significant lag, which could make it difficult to adequately implement techniques involving interactive editing of data in the DAW, such as *selective representation* or *structured representations*.

Based on this analysis, the *task lists* technique appears to be a fruitful area for further investigation: the technique is immediately relevant to pro-amateur composers because it can be used as part of strategies for selective time and effort, it allows further investigation of the co-ordinated use of multiple representations, it has been little explored within the existing literature on composition software, and it is technically feasible to explore within the constraints of a Ph.D study.

Task lists can be used by composers to represent many different types of information, such as plans for material on which work has not yet begun, potential improvements to existing material, past composition activities, and incomplete areas of a composition. While many other techniques described in this chapter can represent *some* of these things, the other techniques cannot be flexibly used to do *all* of them in one representation, giving task lists unique advantages that make it a particularly interesting area to explore.

The following chapter will investigate the technique in more detail.

Chapter 5

Evaluating a to-do list for a DAW

In the previous chapter, design techniques for supporting composers' external cognition were proposed, and *task lists* identified as a technique that might be fruitful for further exploration. This chapter further investigates the *task lists* technique by studying pro-amateur composers using a prototype that integrates with the *Ableton Live* DAW. Based on this study, the advantages of using the *task lists* technique are discussed, and more recommendations are identified for implementing the techniques from the previous chapter.

The chapter is divided into nine main parts: an **introduction** to the study; a description of the user interface of the prototype system **LiveTodos**; a description of the **methodology** including details of the experimental protocol, participants, recruitment, and data analysis approach; a description of the **data** gathered during the study sessions; an account of the four main **themes** generated from the data analysis; seven **patterns of activity** observed; a **discussion** of how these themes and patterns relate to findings from previous chapters; and a summary of the **conclusions** of this chapter.

5.1 Introduction

The aim of this study was to gain more information about how DAWs should be designed to support composers' working processes and practices through the use of the *task lists* technique.

In particular, it aims to answer these questions:

1. How are to-do lists useful to pro-amateur composers working with a DAW?
2. What activities do pro-amateur composers use to-do lists to carry out?
3. How could the representations in DAWs be changed to better support these activities?

The methodology used to answer these questions involved creating a prototype system to provoke composers. In order to see how composers would react to the prototype, they were asked to perform a composition task using the prototype as part of a lab-based study at the University of Nottingham. A detailed record of the use of the prototype was gathered using video cameras and automatic logging, and after the study, the composers' thoughts on the technology were captured using a *semi-structured interview*.

This approach was inspired by *technology probes*:

> "Technology Probes are low-fi technology applications designed to collect information around use, explore usability issues, and ultimately provide inspiration for a new design space. The developers of the Technology Probe concept caution that this is not a form of iterative design for advancing prototypes but instead introduces a novel technology to track how users respond to and engage with it over time" (Boehner et al. 2007)

Introducing a prototype implementing the "task lists" technique to composers helps elicit their thoughts and suggestions about these kinds of representations. It also allows gathering data about how they do (and do not) use these kinds of representations as part of their work.

To create the prototype, it was decided to create a system to extend the DAW *Live* created by *Ableton AG*[1]. *Live* was chosen as a system to extend because it is widely used by pro-amateur users, has a more modern design compared to competitors like *Logic* or *ProTools*, can be extended using a public API without licensing requirements, and has limited existing support for the *task lists* technique.

The only existing options available to composers making notes of tasks to return to later are to make their own notes outside of *Live*, to create a *Locator* (a text string

[1] For readers who are unfamiliar with *Live* or its functionality, a brief summary can be found in Appendix A, which may be useful for understanding the terminology used in this chapter.

Figure 5.1: *Live's Locator* and *Info Text* features

displayed at a particular point in the timeline of the Arrangement mode), or to use the *Info Text* feature to attach a text note to a track, scene, locator, or clip (which has the significant disadvantage of only being visible when selecting that object). Figure 5.1 shows these features in use.

To provide more useful functionality for these purposes, the prototype was created consisting of two parts: the *LiveTodos* website, and a supporting plugin for *Live*.

There were several requirements for this prototype:

1. **Simple functionality**: In order to gain insight into what features composers actually felt were important - and avoid drawing their attention towards particular functionality - the prototype should implement the minimum set of features possible. However, in order to stimulate useful feedback from composers, it should still provide at least some functionality not available with pen and paper notes.

   To meet this requirement, LiveTodos provides a minimal set of features. LiveTodos has the same basic abilities as a paper to-do list - the ability to make and edit a to-do list relating to a particular song, add additional text notes to a list item, and mark each list item done or not done. It extends the abilities of paper to-do lists in one way only: the ability to associate each list item to a specific location and track within a composition in *Live*. A number of alternative extensions were considered before selecting this choice, including extensions directly inspired by the findings in the *Task lists* section of Chapter 4. After a series of technical feasibility prototypes, this extension was chosen, because it appeared to be the most useful and practical to implement while not conflicting with the other requirements.

2. **Provides a secondary representation**: Chapter 3 identified that providing secondary representations to composers was an important feature requiring more investigation, so the prototype should provide an additional interface in addition to the one found in Live. The second finding in Chapter 3 was that the composers studied carried out significant activities away from the DAW so it should be possible to use this secondary representation separately to Live.

   To meet this requirement, LiveTodos provides a web page which acts as a secondary representation in addition to Live. The plugin allows coordination of shared objects between the two representations - for example, if a track is renamed in *Live*, to-do list items in LiveTodos referring to that track will be automatically updated with the new name. In addition to being used in conjunction with *Live*, the LiveTodos website can be used on mobile devices or other computers away from *Live*, as information about the most recent state of the composition can be stored on the LiveTodos web server. This also allows the use of LiveTodos when *Live* is not currently running.

3. **Suitable for a longitudinal study**: Following the technology probe approach, and based on the findings in Chapter 2, the prototype needs to be able to support a study of prolonged use by the composers. The prototype should support use in a field setting, on a composer's own computers, and as they work on their own compositions, while still collecting useful data.

   To meet this requirement, and be suitable for longitudinal use, the prototype also needs to be easy to install onto a composer's own devices and to not interfere with their work. As the LiveTodos user interface is a web page, no installation is required and system requirements are minimal. The plugin can be installed onto any computer running *Live* by copying a folder into *Live's* library and adjusting one setting in *Live's* preferences. To reduce the risk of negative effects on the composer's activity or corruption of the composer's work, the plugin works silently, and does not modify the compositions created by the composer[2]. The LiveTodos website automatically logs each action taken on the web page, and the plugin also records information about the state of *Live* which it regularly sends to the LiveTodos server to be logged, allowing the prototype to record data in a field setting.

5.2 LiveTodos: interface and evolution

Users sign up to LiveTodos through the website and are given instructions to download and install the plugin. Once this is complete, the main LiveTodos website becomes available.

The LiveTodos website displays to-do lists to composers and provides some generic to-do list features (see Figure 5.2). Each to-do list item has a "Done" checkbox, a title, and a set of associated tags. At the top of the to-do list, a text field can be used to enter new list items. When this field is selected, a second field is revealed in which the composer can type in tags to add to the list item. Selecting a list item allows the composer to view and edit metadata (including an 'additional notes' field) and to delete the list item.

The plugin supports the additional functionality that LiveTodos provides that would not be available in a general-purpose to-do list tool. The plugin sends information to the LiveTodos website about the current Live Set, such as the name of each Live Set that is opened, and what tracks or scenes it contains. This permits more advanced co-ordination features:

1. each Live Set opened in Live is automatically associated with a new to-do list on the LiveTodos website
2. to-do list items are automatically annotated with the currently selected track, scene, or loop region in Live

[2] Due to restrictions in the *Live* MIDI remote scripts API, the LiveTodos plugin must add a small amount of data into the key-value store that is contained within each Live Set. This does not affect the composition itself and is invisible to the end user.

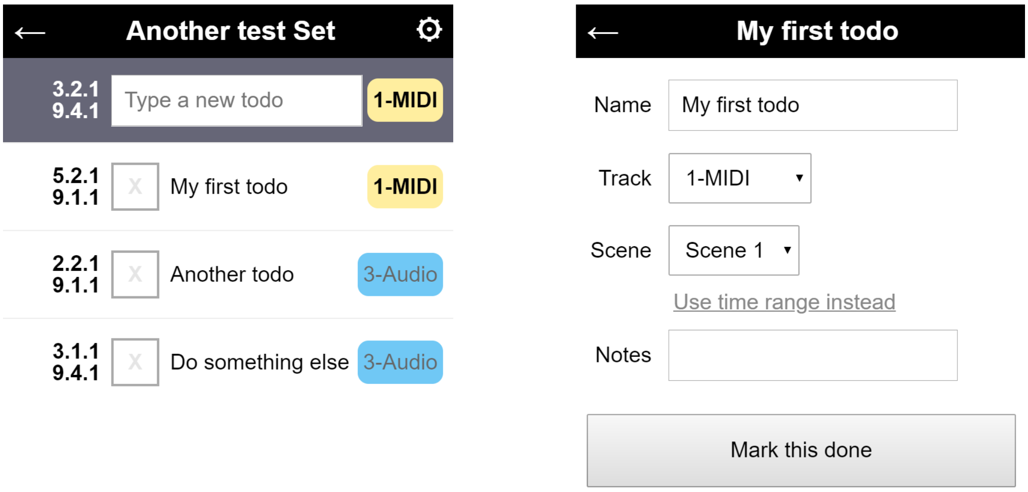

Figure 5.2: Viewing a song (left) and editing a list item (right) in LiveTodos

3. when the name, colour, or position of a track or scene is changed in Live, the LiveTodos website automatically updates to reflect the change # when the currently selected track or scene changes in Live, the LiveTodos website automatically highlights to-do list items related to that track or scene

Once the first composer was observed using LiveTodos, it was apparent both from the feedback received and data recorded that composers needed to be able to incrementally add more detail to their list items over time. An updated version of LiveTodos was created to meet this need which was used by the remaining composers.

The new version (see Figure 5.3) provided a feature where to-do list items could be tagged with tracks, scenes, time ranges, or custom text. In this new version, LiveTodos *suggests* tags to the user based on their current selections in *Live*, rather than adding metadata automatically, allowing details to be gradually added.

As a result, the combined plugin and website now provided the following co-ordination features:

1. each Live Set opened in Live is automatically associated with a new to-do list on the LiveTodos website
2. tags are automatically suggested based on the currently selected track, scene, or loop region in Live
3. tags for tracks and scenes are visually highlighted in LiveTodos when selected in Live

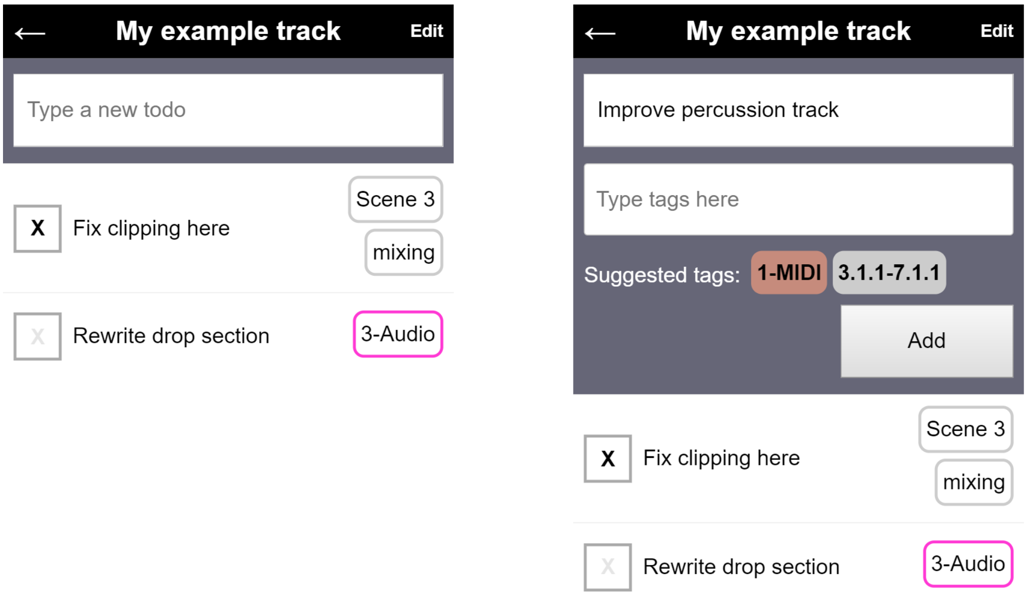

Figure 5.3: Viewing a song (left) and adding a new list item (right) in the updated LiveTodos version

4. when the name, colour, or position of a track or scene is changed in Live, the corresponding tag in LiveTodos automatically updates to reflect the change

The remaining four composers all used versions of this interface. Minor tweaks were made over time as user interface issues were spotted. The composers assigned tags to nearly half of all list items they created. The most commonly created type of tag was tags representing tracks in *Live*. Some composers made use of the ability to tag a list item with more than one track or scene at a time. Composers rarely created their own custom tags.

5.3 *Methodology*

5.3.1 *Experimental protocol*

The composers were all provided with a laptop running Live, a 49-key MIDI controller keyboard, and a pair of over-the-ear headphones. The composers were also provided with a tablet computer on which to use the LiveTodos website.

All sessions of the study were held in a small rectangular room with a circular table, around which chairs were placed (see Figure 5.4). On the table was placed the equipment listed above, along with printed instructions for the study, and an audio recorder used to capture a record of the study. On the wall in front of the

Figure 5.4: Equipment used in the study

composer, there was a whiteboard.

In the previous study described in Chapter 3, a significant problem was the quality of the video data captured, which did not always adequately record the activity of composers. In this study, a small "action camera" was used in order to avoid this issue. The camera was capable of recording 4K resolution video with a wide field of view and when placed on a tripod to the left of the composers' seat allowed almost the entire area of the table to be recorded in detail.

At the start of each study, the laptop was placed in the centre of the table, with the external keyboard in front of the laptop, and the mouse slightly to one side. A chair was positioned in front of this, and upon entering the room the composer was offered that chair to sit in. The composers in the study placed the iPad directly in front of them at the beginning of the study. At 90 degrees around the table, to the right, a second chair was placed, with the MIDI keyboard in front of it.

The study took around two hours in total to complete. Initially, the study was intended to also incorporate a longitudinal element in which the composers made use of the technology probe in their own practice while LiveTodos would gather information through automatic logging. However, feedback from the composers was that the probe was not yet ready for use on their own projects, and this element was abandoned. A description of the study follows, and a detailed time schedule can also be found in Appendix C.

After each composer had read an information sheet and privacy notice, and read and signed a consent form, they participated in a set of introductory tasks that taught them how to set up and use LiveTodos.

After each composer had read an information sheet and privacy notice, and read and signed a consent form, they participated in a set of introductory tasks that taught them how to set up and use LiveTodos. The composers then read the instructions for the main task. For this task, the composers were given three pictures and instructed to produce a short 2-3 minute piece of music to accompany one of them. Composers were instructed to use LiveTodos as part of the task. Around 60 to 70 minutes total was allotted for the composers to work, but some elected to stop earlier as they felt they had completed the task.

Based on the observation in Chapter 3 that for one composer a break in activity appeared to have a significant effect, the study was designed to simulate a natural disruption to the task. The composers were deliberately interrupted once in the middle of their work, and were encouraged to move about and leave the room for a short period, as well as being asked for some biographical questions.

After the composition work was complete, a semi-structured interview asking the composers for thoughts about their use of to-do lists and LiveTodos took place. Composers were interviewed about their use of to-do lists, how LiveTodos affected their work, and for positive and negative thoughts and suggestions about LiveTodos.

A problem experienced in the study in Chapter 3 was that biographical questions had taken up significant time. Biographical questions in this study were designed to obtain the most relevant information quickly, based on the most useful information obtained in the previous study:

1. How long have you been active in making music and performing music?
2. Have you ever worked writing, performing, or recording music as a full-time job?
3. What kind of music do you normally make?

The closing interview involved six questions about LiveTodos and to-do lists:

1. Do you normally use to-do lists when you are writing music?
2. Did you use any other written or electronic notes when you were using LiveTodos?
3. Did you find that using LiveTodos changed how you worked?
4. Can you describe any occasions where you found using LiveTodos was helpful?
5. Can you describe any occasions where you found using LiveTodos was less helpful?
6. Do you have any other suggestions for how LiveTodos could be improved?

If it appeared that it might be fruitful, composers were encouraged to continue speaking about each question after they initially stopped speaking. In some cases, the composers were asked further clarifying questions to provide more detail about their answer.

At the end of the study, some composers were offered the opportunity to continue using LiveTodos in their own work, and instructions on how to do so. None of the composers decided to do so. The reluctance of the composers to make further use of LiveTodos will be discussed in more detail in Chapter 6 (see *Critical reflection*).

5.3.2 Data analysis

In contrast to the study in Chapter 3, which primarily found patterns that emerged from an analysis of video recordings, the patterns found in this study primarily emerged from an analysis of interview data.

This analysis took two parts:

1. a *thematic* analysis to identify the key themes that emerged in the responses of composers
2. grouping the previously coded phrases into *patterns of composer activity*

The initial analysis phase used the six-stage approach to thematic analysis described by Braun and Clarke (2006), and their terminology will be used here.

The entirety of the interviews (including where relevant, the biographical information questions) were used as the data set for analysis. An *inductive* approach was used, based on a more "grounded" approach of describing patterns found in the data, rather than using a pre-existing theoretical framework, with the intention of producing a rich description of the data set (as opposed to concentrating on particular aspects). *Semantic* themes (describing what the composers said, rather than theorising about the underlying patterns behind them) were generated.

For each pattern the interview data was triangulated with multiple additional sources of data gathered through automatic logging of user activity on the LiveTodos website, logging of the state of Live carried out automatically by the LiveTodos plugin, and video recording of the composer at work and of the laptop screen. LiveTodos logged the creation of each to-do list and to-do list item, changes made to lists and list items, and any tags and notes that were attached to list items. Using the plugin, data about the state of Live was automatically logged, including Live's current mode ("Arranger" or "Session"), the number of tracks and scenes, playback state (playing or stopped), the current loop region, and the selected track and scene. Timelines of this data were created using Vega-Lite (Satyanarayan et al. 2017) to aid this work, and are used to illustrate the patterns in the following sections.

5.3.3 Recruitment and participants

Composers were recruited either through existing social connections or by advertising in local groups for musicians. The composers were each compensated for their time with £20 of shopping vouchers. Composers were required to be over the age of 18.

Ethical approval was received from the University of Nottingham Computer Science Research Ethics Committee. Details of the study were provided to participants through an information sheet and privacy notice, and consent was gained from participants by their signing of a consent form. Copies of these documents can be found in Appendix C. Anonymised data was published on the University of Nottingham Research Data Management Repository.

Composers were chosen to ensure variety in age, gender, and the genres they worked in. Five composers were observed in total:

- *Composer 1* had 35 years of experience writing and performing music, and had previously worked full time in this field, but had not been involved in full-time projects for a number of years. Most of their work typically involves a computer-based DAW. They work in a variety of styles and depending on the project may use either primarily electronic or traditional instruments. They did not have prior experience of using Live.

- *Composer 2* had 15 years of experience writing and performing music, but had never done so as a full-time job. They reported that they worked in a range of styles, including guitar-based indie, electronic, and recording and remixing live jazz. They normally worked using *Live* and *Max* along with software synths and a full-sized MIDI controller keyboard.

- *Composer 3* had 10 years of relevant experience including releasing records on commercial record labels and creating commercial sample packs. During this period they had sometimes worked full-time, sometimes working on projects when they received a commission rather than as a full-time job, and sometimes working on music part-time along with jobs unrelated to music. They had not worked full-time on music for over a year. They used a variety of genre names to describe the range of their musical work and musical influences, including house, disco, Balearic, dub-techno, and broken techno. They had been using *Live* for over 10 years.

- *Composer 4* described themselves as being involved in making and performing music for "ten to fifteen years", but never as a full-time job. As part of a band they had been commissioned to create work by art organisations but currently were releasing solo work. They described this work as being of a "techno" nature. *Ableton Live* had been their "main hub for a long long time", but their process had evolved to now use iOS apps for generating material, which they would then bring back to *Live* for mixing.

- *Composer 5* had previously been a professional touring pop-rock performer, but had shifted career paths five or six years previously, and no longer performed or recorded professionally. They described their current creative output as being "atmospheric" "soundscapes" consisting of repeating loops created with iPad apps and ROLI Blocks which they had not yet released publicly. They had only a small amount of experience using *Live* but had previously used a variety of tools for recording music including *Cubase Sequel* and *ProTools*.

5.4 Data

5.4.1 Setting-oriented records

The video camera recording of the composition sessions was analysed to identify what use the composers made of the setting.

Composers 3 and 5 rearranged the objects on the table in order to be able to simultaneously make use of the MIDI keyboard, by placing the MIDI keyboard in front of the laptop, which necessitated placing the iPad off to one side in both cases. Composer 4 made a small adjustment early on as they preferred using the built-in keyboard on the laptop, so placed the external computer keyboard off to one side to provide more room to place the iPad. The remaining two composers made no significant adjustments to the positions of objects.

No composers made use of the pens, blank paper, or whiteboard during the composition period. Composer 5 requested a microphone to record vocals as part of their composition, but as it was not possible to quickly locate a suitable microphone, they used the internal microphone in the laptop to complete their work.

Composer 4 did bring their own MIDI controller with them to the study, but decided not to plug it in or use it, and instead used the controller that was provided.

All apart from Composer 3 made use of the opportunity to take a break and leave the room for a period.

5.4.2 Person-oriented records

Using the video recording, a record of what objects the composers were using at each point in the study was generated.

Figure 5.5 visualises this data graphically. In Figure 5.5 and the following figures, grey shading in the background of figures indicates the period during which the composer was interrupted from their work. Due to a technical problem with the camera, footage of Composer 1 was not captured during part of their study session, so the data is partially missing.

Most composers spent most of their time using LiveTodos on the laptop, but repeatedly switched from using *Live* to actively using LiveTodos for a brief period

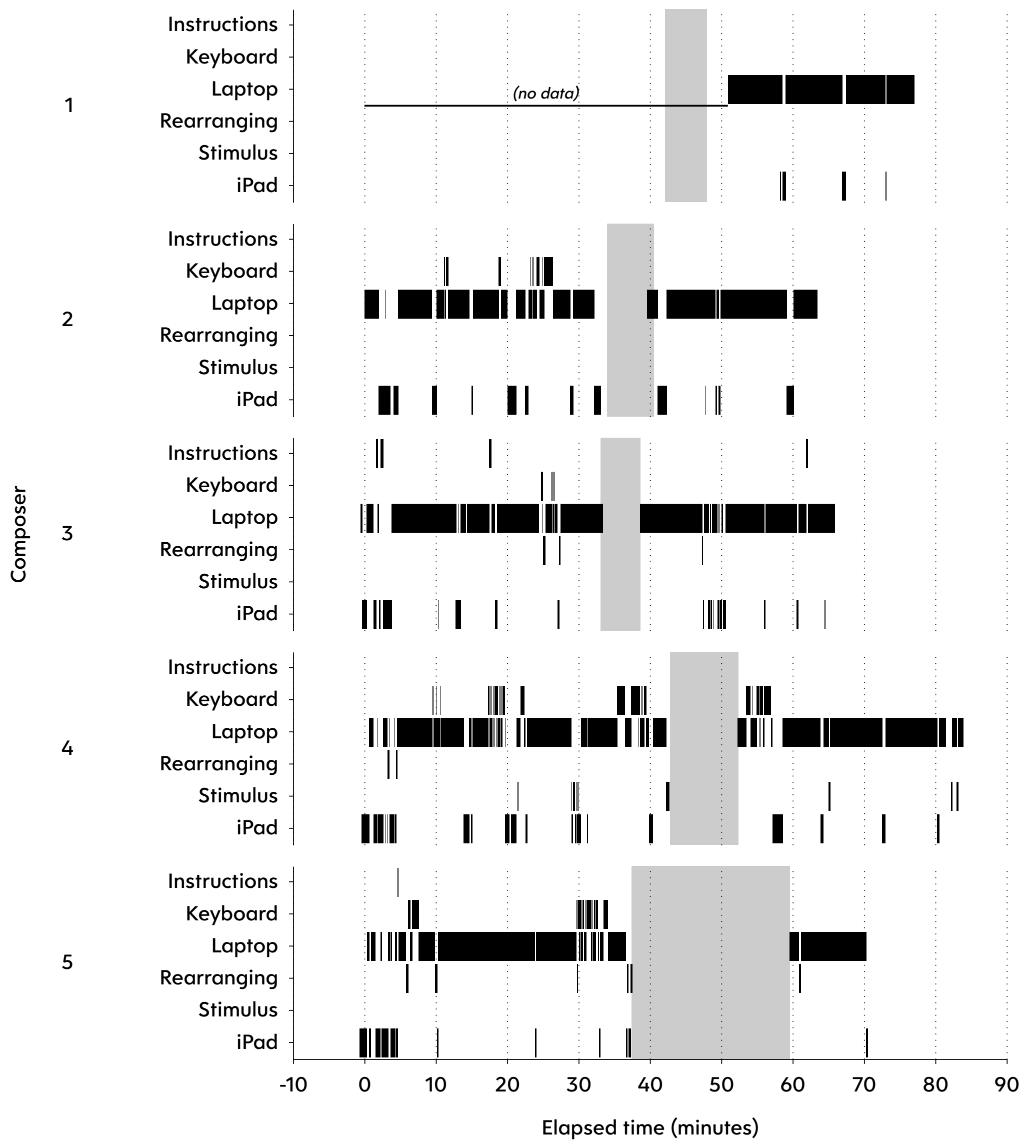

Figure 5.5: Timeline of objects in use by each composer

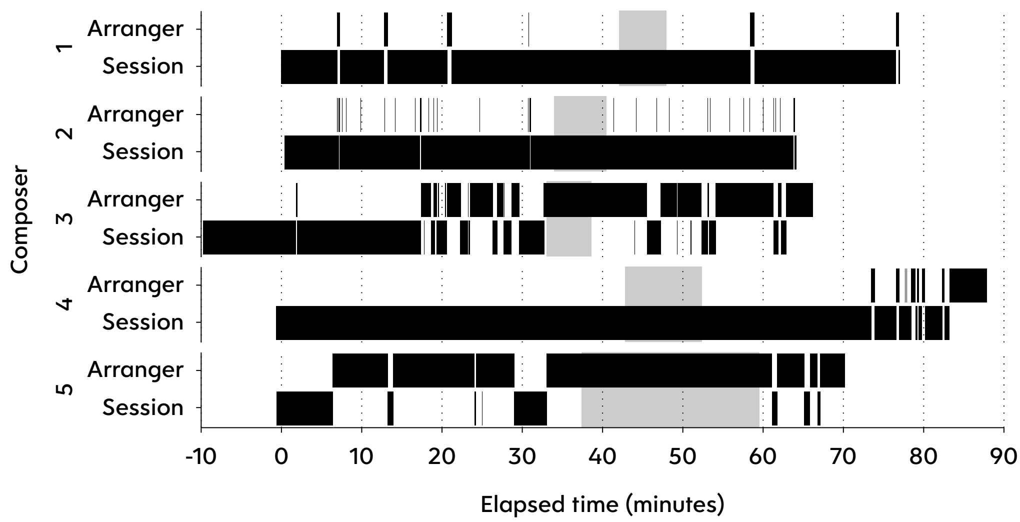

Figure 5.6: The mode *Live* was used in by each composer

every 5-10 minutes. Occasionally, composers would switch to using LiveTodos in longer more focused bursts. These patterns will be discussed further in the section *Using to-do lists to plan and focus*.

Unexpectedly, even though participants were instructed to compose music to accompany a picture that acted as a stimulus, only Composer 4 looked at the stimulus after starting to compose.

5.4.3 *Object-oriented records*

Due to the fact that composers exclusively did their work through computerised representations in *Live* and LiveTodos, it was possible to create detailed logs of their activity in both tools.

The plugin in Live allowed the logging of *Live's* current mode (Arranger or Session), the number of tracks and scenes, playback state (playing or stopped), the current loop region, and the selected track and scene.

Upon each occasion that one of these properties changed, a log entry was automatically created containing the current values of all the properties, the current time and date, the current position in time of the playback cursor, and the current time signature.

This allowed the production of very detailed graphics showing how *Live* was being used.

Surprisingly, this data revealed a variety of unexpected strategies for switching between *Live's* Session and Arranger modes, which are shown in Figure 5.6. Rather than initially "jamming" in Session mode, and then switching to Arranger mode when a musical structure was being created, the composers would often switch repeatedly back and forward between modes. One reason for doing this was that the Session view provided useful information unavailable in the Arranger view, such as the large level meters, and the optionally revealable peak dB level for each channel. This suggests that some of the composers were *already* comfortable with rapidly switching between different representations with different advantages.

The LiveTodos website logged every creation of a new list item or song, changes made to list items or songs, and any tags and notes that were attached to list items.

Each log entry contained information about the current time and date, and an ID number identifying the composer who was currently using the system.

Using this data, it is possible to create a detailed timeline of how the composers made use of LiveTodos, which is shown in Figure 5.8. In that figure, horizontal lines represent the interval of time in which composers interacted with each list item, and vertical marks represent each time a list item was created or edited. The period in which composers were interrupted and were not composing is shown with a grey background. The logged data makes it clear that all composers made a reasonable degree of use of LiveTodos (see Figure 5.7).

In some cases, it was possible to use this data to infer which to-do list item was being worked on at a given time, though this was only possible in limited cases where to-do list items unambiguously referred to work on a specific track within Live.

A number of manual corrections to the logged data were made after reconstructing the activity of composers from the video recording. For Composer 3, no information was logged at all for the second half of the session due to a bug, and this data was manually created from the video record. Some data was misleading due to a UI problem, where after entering the name and tags for a list item, composers would then forget to press the button that would submit the data and create the list item. As a result, the logged data would show the list item as being created at a later point in time when the composer had returned to the tool again and remembered to press

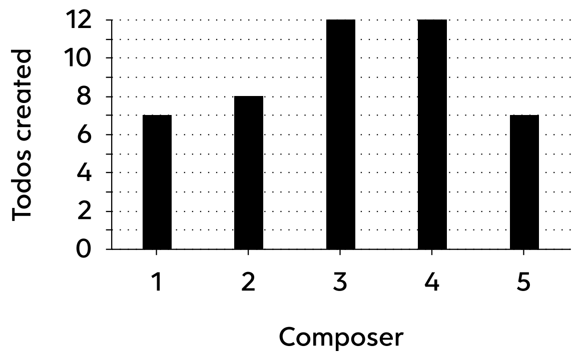

Figure 5.7: Number of to-do list items created by each composer

the submit button. To better represent the way users were actually using LiveTodos, the data was altered to read that the list item was created when the data for that list item was originally entered.

5.4.4 Task-oriented records

In common with the study described in this chapter, the study described in Chapter 3 also involved an observer who was present in the room during the study. One observation in Chapter 3 is that this kind of study involves a risk of the observer becoming co-opted as an unintended participant in the composition activity. This can have unintended effects like the composer narrating their activity to the observer.

The problems described in Chapter 3 are likely to have been exacerbated in that study by the biographical interview that took place before composition started. To avoid similar issues occurring in this study, some precautions were taken: interviews were not scheduled before composing started, and both the observer and the composer wore headphones during the composition work. This strategy appeared to be mostly effective as significantly reduced levels of interaction were observed between the observer and composer compared to the previous study.

However, there were several occasions where it became necessary for the observer to intervene significantly to keep the study running smoothly. LiveTodos bugs occasionally resulted in the observer having to tell the composer what they should do in order to be able to continue working. Problems involving *Live* being misconfigured stopped the MIDI keyboard from working in one case (Composer 4), which required intervention to diagnose. Inexperience with *Live* meant that some composers needed to occasionally ask questions about how to do particular actions. A request for a microphone by Composer 5 also resulted in the study being paused

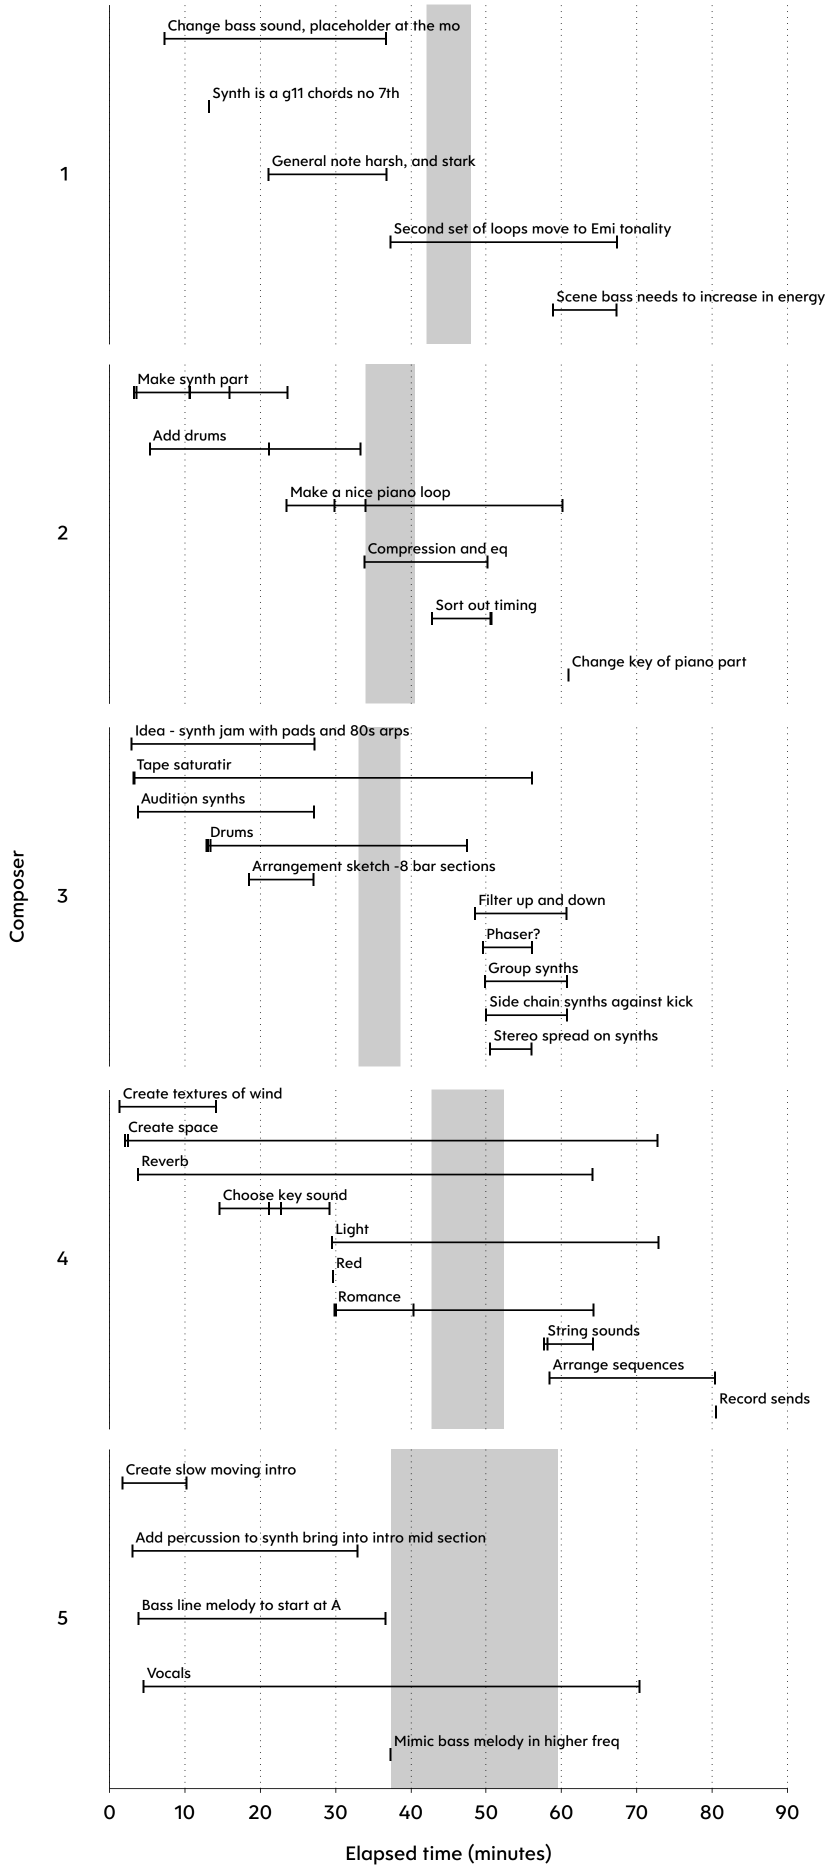

Figure 5.8: Timeline of LiveTodos activity for each composer

while this was discussed and a failed attempt to locate a microphone nearby was made.

Composers were given verbal warnings of how long they had left to complete the task, both at the start of the second composition period (25 minutes remaining) and at the point where 10 minutes remained.

5.4.5 Interviews

Familiarity with the interview material was gained through being present at the interview, during the transcription process, and through repeated close re-reading. 24 initial codes representing patterns found in the interview data were generated, which were then organised into a smaller number of themes.

- *Organisation*: "Ordering notes", "Grouping and connecting notes", "Spatial organisation"
- *Different needs in different situations*: "Differing usefulness at different process stages", "Differing usefulness on different types of project", "Collaboration"
- *Anchoring the creative process*: "Anchoring activity to an initial idea", "Evolving notes", "Signature sounds"
- *Managing time use*: "Managing time use", "Prioritising notes"
- *Reflection*: "Reflection on activity", "Iteration"
- *Augmenting memory*: "Remembering things", "Remembering context"
- *Problems and suggestions*: "UI issues", "Suggested tags", "Connecting to physical notes", "Automatic update of done state"

These themes were validated by triangulating them with the interaction analysis data, as described in the section *Data analysis*. Following this, initial drafts of possible sections corresponding to each theme were produced. These drafts were reflected on, and themes were merged, split up, and renamed to create a new set of themes with improved distinctness and relevance.

The new set of four themes were:

1. *Using to-do lists to plan and focus*
2. *Changing to-do list items over time*
3. *Organising to-do lists*
4. *Applicability of to-do lists*

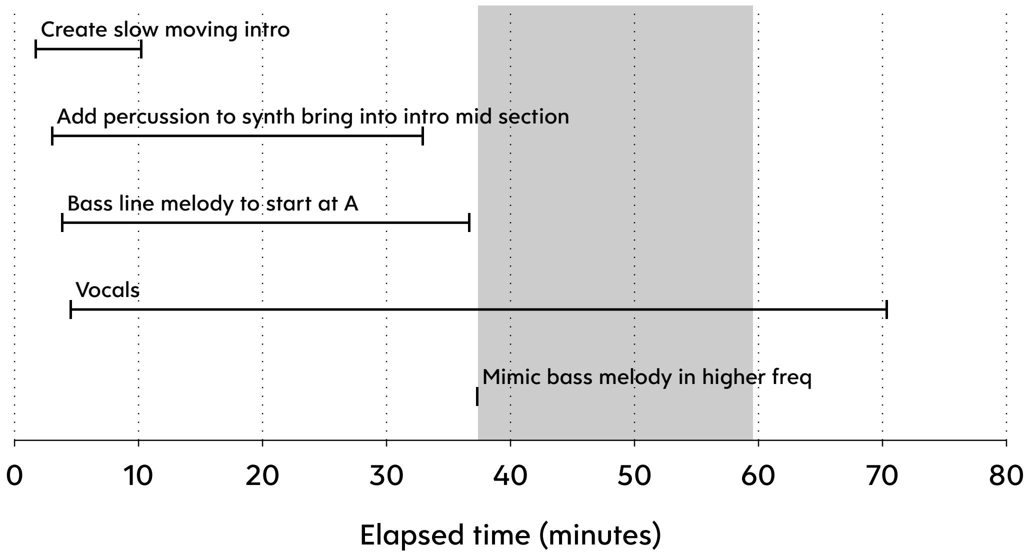

Figure 5.9: Timeline of LiveTodos activity for Composer 5

The following sections provide a more detailed examination of each of these themes.

5.5 Themes

5.5.1 Using to-do lists to plan and focus

In the interviews, composers repeatedly spoke about how LiveTodos was a helpful tool for thinking about what they wanted to make and how to make it.

> "It's so weird to think about this process and like, think about what I actually do, before making the thing… It made me more conscious of what I needed to do to make a thing… I'd never have worked in a way where I consciously think about what I'm doing" (Composer 5)

Composer 5 immediately created notes to indicate what their creative intent was at the start of the session (see Figure 5.9) and then proceeded to carry out their planned work based on their initial plan.

Composer 4 did not immediately note down their creative intent at the start, but the record of what tools they were interacting with (see Figure 5.10) suggests they used an iterative approach with more than one planning phase. Two distinct periods at 29-30 and 57-58 minutes occur, where the use of *Live* stops entirely for a sustained period of time and the composer focuses entirely on using LiveTodos.

Cross-referencing these time periods with the record of edits made to tasks (see

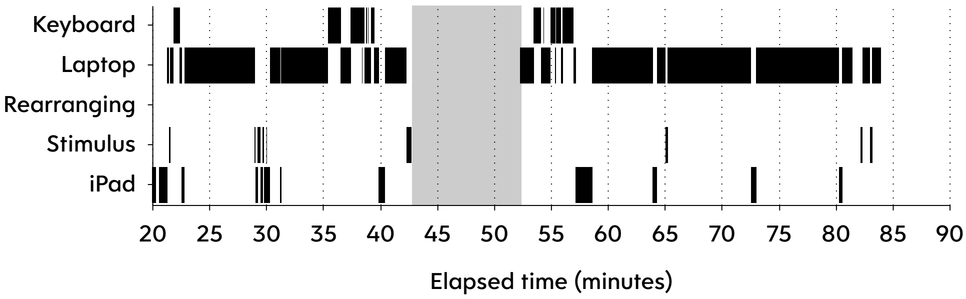

Figure 5.10: Detail section of timeline of objects used by Composer 4

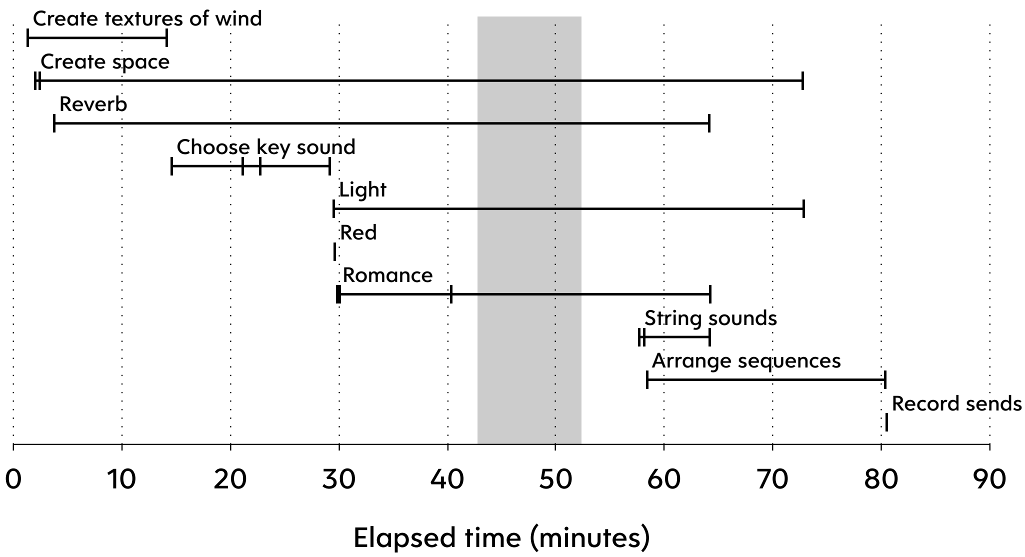

Figure 5.11: Timeline of LiveTodos activity for Composer 4

Figure 5.11) helps to reveal more information about the activity of the composer. In the first period, the composer refers to the stimulus image to decide what aspects of it they plan to focus on representing musically, creating three list items (“Red”, “Light” and “Romance”). In the second period, the video recording shows the composer pausing to think and entering new list items (“Arrange sequences” and “String sounds”).

The composers appeared to be recording their creative intentions for the composition into LiveTodos, though how specific this record was varied between concrete plans of action and recording aspects of interest. Composer 4 described how they found they “make notes about colour, and light, and things to actually try and represent, rather than… ‘chord here’, or um, ‘sort out the arrangements’ ”, while Composer 2 recorded specific actions they intended to carry out.

The comments made by the composers in their interviews made clear they had two main reasons for recording their creative intentions using LiveTodos.

The composers spoke about how it helped them *stick to their original idea*:

> "I'll have an idea and then forget about it, and then go off in a different direction, I think it's quite useful to be able to, um, you know, get ideas out and have them actually in front of you." (Composer 3)

The composers also talked about how it helped them *manage their use of time* by avoiding being "lost" or "stuck in a rut":

> "you can get kind of… swept away in sort of countless possibilities… rather than kind of keeping your mind on the task" (Composer 2)

> "made me more efficient, because I wasn't then just getting lost in the piece", "it wasn't like… noodling around forever finding a decent sound that I needed", "I think it's really useful in giving pace" (Composer 5)

It is important to note that the nature of the study may have made these issues prominent in the minds of the composers. The task involved keeping their work related to a specific picture (for which sticking to their original idea would have been useful) and in a restrictive time period (for which managing their time would be important).

However, in the study described in Chapter 3, which imposed neither of these constraints on the composers, a similar pattern of composers trying to focus their activity and selectively minimising exploration also emerged. This suggests that for at least some pro-amateur composers, these are important concerns. The section *Applicability of to-do lists* will describe how the composers studied in this chapter talked about time constraints on their work.

The comments the composers made indicate that they did not normally record their creative intention even though they clearly found it valuable to do so. This is despite the fact that *Live* already had some facilities that would have allowed composers to keep track of their creative intention. As described earlier in this chapter (see *Introduction*), it was possible to indicate work to be done by adding Locators and Info Text to the timeline. It was also possible to create silent clips whose names gave further information about musical intentions.

In many cases, the composers appear to have actually carried out *more* work by creating to-do list items. Composer 3 created the list item "Group synths" despite the fact that it would have required less effort to just create the group, and instead of creating list items like “Reverb”, Composer 4 could have just added the effect straight to the track in question just as easily. Many of the composers created list items when they could have simply created a new track in Live: “Drums”, “String sounds”, “Vocals”, “Add drums”, “Make synth part”, and “Make a piano loop” are all list items that refer to creating and populating a new track.

As the composers indicated a significant change in the amount they recorded their creative intention compared to their normal practices, LiveTodos must have encouraged composers to do so in ways that *Live* did not.

5.5.2 Changing to-do list items over time

Composer 1 explained how the way LiveTodos automatically linked to-do list items to tracks and scenes in *Live* didn’t fit their way of working:

> “throughout, I was kind of thinking, I don’t need things - at this stage, I don’t need things aligned to particular sections, what I need is a general notepad… maybe an option to be able to take general notes which later on you could assign to particular tracks or locations in a track might be useful”

This prompted the redesign of the tool described in the earlier section *LiveTodos: interface and evolution*.

Despite the design improvements, the study session with Composer 2 featured more discussion about this topic. They explained how the notes they created to represent their idea needed to evolve as their creative idea did:

> “I think that as things evolve… things come up… things needed to change, really, around a little bit” (Composer 2)

This echoed similar comments that Composer 1 had expressed about how the meaning of a list item could change over time:

> “things aren’t necessarily Done, they’ve just been changed… for example… I did [this to-do list item] slightly, but I wasn’t actually happy with what I did there at all. But it was some, movement towards what was in my head. So I marked it off as done, but it’s not Done… in a

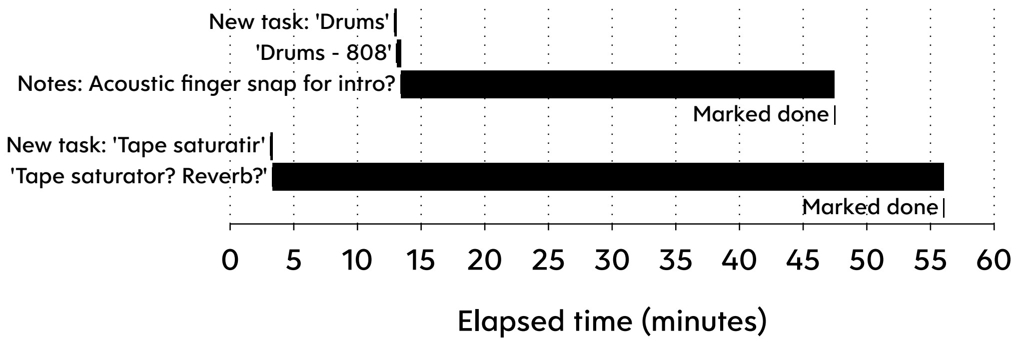

Figure 5.12: Updates to to-do list items created by Composer 3

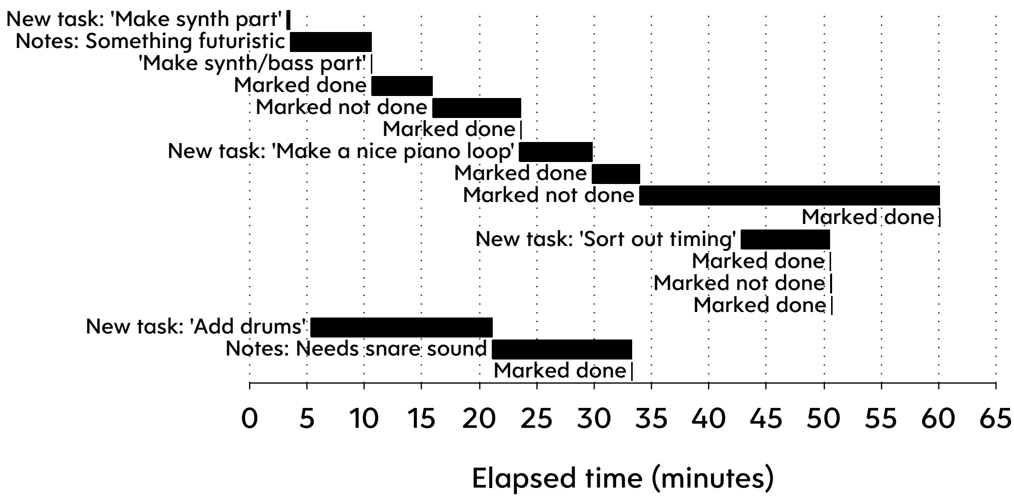

Figure 5.13: Updates to to-do list items created by Composer 2

> sense of looking at my statement I'd done that, but not in terms of what ultimately I wanted to work to" (Composer 1)

The logged data allows studying how composers altered their list items over the course of sessions and illustrates several different kinds of activity.

Correcting data entry mistakes: Composer 3 only made updates to tasks immediately after creating them (see Figure 5.12). These updates were to immediately add details that were originally omitted or to fix typos they had made.

A number of Composer 3's list items ("Tape saturator? Reverb?", "Phaser?", and the notes on "Drum - 808": "Acoustic finger snap for intro?") contain question marks. These could potentially represent uncertainty about the task the list item referred to, or could be an instruction to explore one or more possible alternatives. Interestingly, the composer did *not* update these list items later on as they made decisions.

Indicating remaining work: Composer 2 updated to-do list items to indicate that a

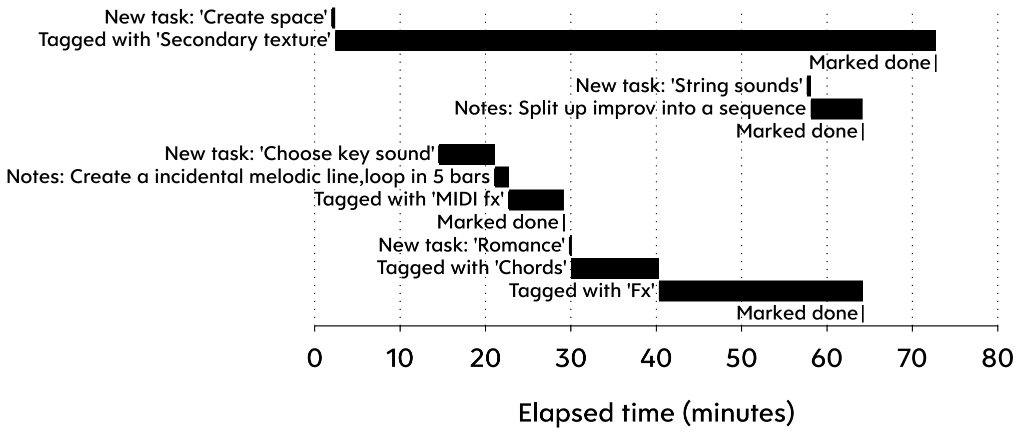

Figure 5.14: Updates to to-do list items created by Composer 4

task was only partially complete (see Figure 5.13). A good example of Composer 2's use is "Add drums"; after starting work on the Drums task they did not finish it completely, so needed to add the note "Needs snare sound" to indicate the remaining work left to do, which they returned to do at a later time.

Associating several to-do list items: Composer 4 updated tasks to indicate what action they were about to carry out next (see Figure 5.14). "Choose key sound" is a good example of this: after performing the original task in the title of the list item, they add the notes "Create an incidental melodic line", a *different* but related task. They then immediately perform that task, and add a "MIDI fx" tag - again indicating a different task - and then immediately add MIDI fx to the track.

Composer 4 never used updates to tasks as reminders for work to return to in future, and always updated to indicate work they were just starting. As they clearly were not using the to-do list items to record intended future actions they presumably felt that list items were useful as a record of previous actions. Two composers did talk about how a diary of their past to-do list items would be useful. Composer 1 wondered about the possibility of "some kind of version history", and started to hypothesise about an interface with "kind of snapshots as things develop and whether or not, you know, they could, (pauses), you could fold back via this or via that, and kind of see… how things have developed and what you've done". Composer 3 also wanted to be able to revisit their to-do lists at a later date to reuse ordered to-do lists between different compositions in order to achieve a personal "signature sound".

A final less interesting cause of updates by Composer 2 was problems with the user interface: the version of LiveTodos that Composer 2 used did not immediately update the Done state of to-do list items when it was changed, which lead to tasks being incorrectly marked Done and Not Done due to the lack of immediate feedback.

It appears that composers might wish to change their to-do list items over time for many different reasons: changing plans, indicating remaining work, associating several list items, keeping an accurate record of the activity they carried out, and correcting mistakes.

5.5.3 Organising to-do lists

All but one composer, when asked for suggested improvements, spoke about the need for more organisational features. Ordering, grouping, linking, and spatial positioning were all suggested.

In most cases, it remains unclear from the interview data exactly what the composers wanted to use organising to-do lists to achieve. The exception is Composer 5, who suggested ordering or grouping by priority as something that was of interest to them.

Composer 5 also described how being able to organise their notes might have advantages over paper notes. With handwritten notes they had created some time ago, they found the notes "had no context, and no idea like what that linked to, or what song that was for, or whether I had lyrics for that or not, and it was really like (stutters) I decided to like put all of it in the bin". For Composer 5, a notes tool with organisational features might allow them to preserve this context.

Composer 3 made use of *Live's* ability to create groups of tracks, probably the most visually obvious organisational feature provided by *Live*, to create two groups of tracks. However, creating these groups does not appear to have been primarily motivated by a need to organise the tracks. Groups of tracks in *Live* also serve a second purpose: signal routing.

Most DAWs, including *Live*, provide an audio bus system that allows the creation of *return tracks* that imitate a similar feature in hardware mixing desks. Return tracks combine signals from a number of input tracks that *send* signals to them. Effects can then be applied to the return track, acting on the combined signal from all the inputs, which is then sent to the master output. A *group* in *Live* provides very similar

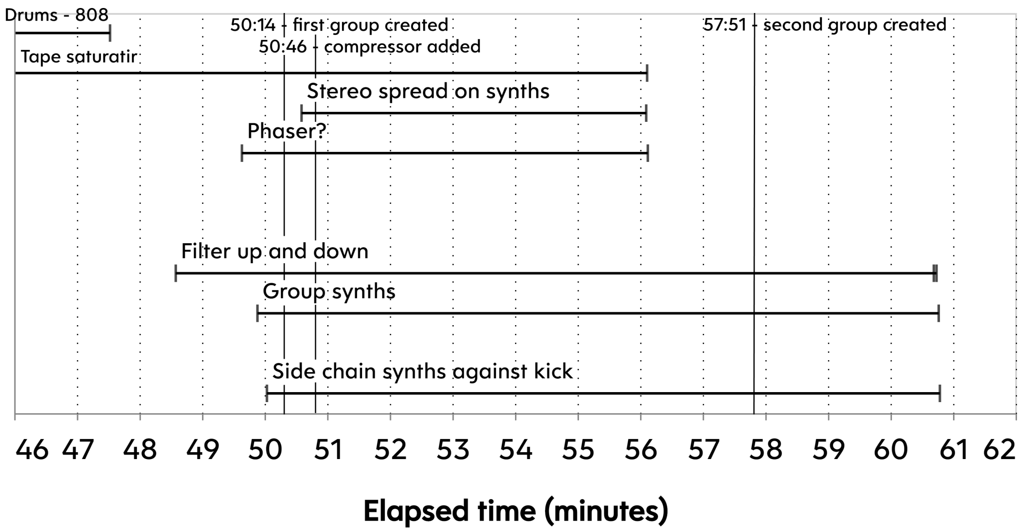

Figure 5.15: Detail of part of Composer 3's activity with significant group related events indicated

functionality to a return, but is represented to the user in a different way. A group track provides the routing functionality of a return, but appears to visually contain the "child" tracks that send signals to it, rather than imitating hardware mixers in a skeuomorphic way.

Composer 3 made multiple to-do list items that referred to track groups in *Live*, all created around 50 minutes into composition (see Figure 5.15). Around the same time, they grouped four tracks in *Live* ("Fog Pad", "Bass", "Arp", and "Melody"), and later also created a second grouping consisting of an 808 drum machine and a triangle loop (the two remaining tracks in their Live Set that had not been assigned to a group).

It can be assumed that "Group synths" refers to the group of four synth instruments which is created shortly after the to-do list item. As soon as Composer 5 finishes creating the list items they almost immediately add a Compressor effect to the group in order to make use of the signal routing described earlier. Later, when the second group is created, a second compressor is again applied to that grouping. Finally, the task referred to in the "Side chain synths against kick" list item was completed by adding an effect on the synths group that controls the output volume based on the volume of the 808 drum machine track.

In this situation, the groups clearly are not just logical groupings, but are also significantly used as signal routing paths. Perhaps other signal routing interfaces in DAWs (e.g. return/send and graph-based) might also be thought about by users as groupings. Coordinating organisational groupings in another software tool with those used in DAWs may be more complex than it initially appears.

5.5.4 *Applicability of to-do lists*

A recurring theme was composers describing how they had different needs in different situations. Comments were made about how LiveTodos might be more or less useful depending on the *stage of composition*, whether the work was being carried out *solo or collaboratively*, the *complexity of the project* being worked on, the *locations* that work took place, and whether there were *time constraints* on the work.

Several composers remarked how this kind of tool might make more sense at specific **stages of composition**. Composer 1 pointed out that tagging with Scenes or timestamps made less sense at the start of composition, when a structure did not yet exist, but did think to-do lists were still useful at that stage. Composer 5 also described not having a structure in early stages, and that they "just let it happen, um, and it's only when I come to record it, that I think, that I have to create this song structure." Composer 2 mentioned that their existing use of paper notes was exclusively in the final stages of composition to make checklists of things to complete, suggesting that they had more of a need at that point. Composer 3 thought they would use it later when refining a track "to drill down and sort of, you know, address, why the track works or doesn't work", and that it would "really come into its own" at that point, while at earlier stages their method was to "sort of throw a lot of things together, so, when I'm in that kind of mindset, in the kind of open mode, um, having to… try and keep track of that doesn't really flow". Composer 3 also spoke about how earlier in composing a piece, a general note pad was more useful, but that later on, the ability to coordinate LiveTodos with Ableton to link to specific tracks and locations in time would become helpful. Composer 3 barely used the tagging facility, tagging a to-do list item only once.

The composers also talked about the difference between how useful LiveTodos might be depending on whether they were working on a **solo or collaborative** project. Composers 3 and 4 both spoke about how they thought LiveTodos would be especially useful in collaborative projects.

Composer 4 also stated that LiveTodos might be more useful on larger projects, suggesting the **complexity of the project** might also be a relevant factor.

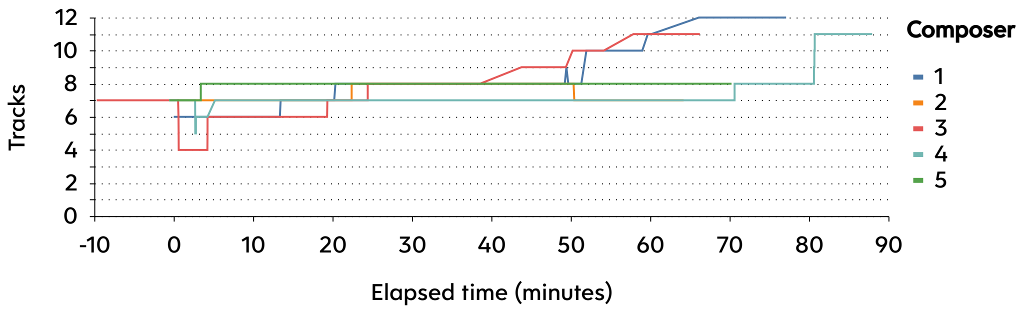

Figure 5.16: Number of tracks created in *Live* by each composer

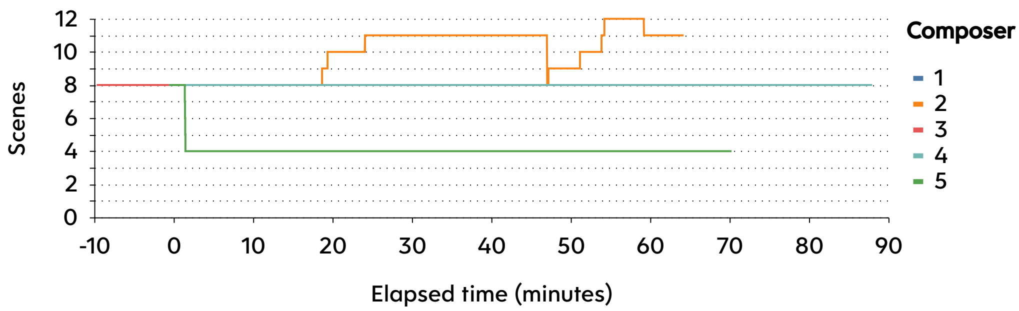

Figure 5.17: Number of scenes created in *Live* by each composer

> "this is quite minimal… there's only what, um, seven channels here… in, perhaps, more expansive Live Sets, where there's say, over twenty channels, it, it would certainly be more useful in that context." (Composer 4)

All of the composers kept the compositions they were creating relatively simple, presumably in response to the limited time they were given to work in, and the Live Sets they created all contained a small number of tracks.

This is shown in Figure 5.16 and Figure 5.17. Note that two 'return' tracks and the 'master' track are created by *Live* by default, so the number of 'instrument' tracks would normally be three less than the number shown in the figure. Composer 1 left two tracks entirely empty. By default *Live* creates new Live Sets containing eight scenes and the majority of the composers used only a small number of these scenes and did not delete them or create additional scenes.

Another important factor to composers was **location**, with working away from a desk was clearly part of the practice of the majority of the composers. Composers 4 described how they used a lot of software on mobile devices because they needed to be able to work in cramped circumstances while travelling, whereas for Composer 5 their use of mobile devices was because their practice initially was "very much an acoustic process" working with instruments with recording software like a DAW only being involved later on. Composer 2 also worked in a different location initially, recording samples of live jazz they would then bring back to Live.

Finally, **time constraints** were a common concern for the composers and may affect how useful LiveTodos is to them. Describing the composition they did, the composers mentioned paid work on contracts and commissions for arts organisations, companies, and record labels, and also working both on their own and as part of groups purely for their own pleasure. Some types of work would involve deadlines and time pressures, either required by a client, or as part of working on a project where material had to be provided to collaborators. Some composers also spoke about the time pressures created by working an unrelated full-time job in addition to their composing (Composer 3 and 4). Composer 4 mentioned that in their normal compositional practice they would normally not create to-do lists or checklists unless there was a deadline they were working towards. They saw LiveTodos as useful for commissioned projects with deadlines and suggested they would be unlikely to use it in more personal projects.

Composition by pro-amateur composers appears to be carried out in a variety of situations in which different features may be more or less applicable. Designers of DAWs need to be aware that features of to-do lists they implement may be particularly required for certain situations but also be less helpful in others.

5.6 Activity patterns

After finishing the initial thematic analysis the phrases that had previously been coded were then re-analysed by grouping them together if they related to similar activities.

Seven common patterns of activity that composers undertook using LiveTodos were identified based on the interview data:

1. planning activity
2. journalling activity
3. interleaving activities
4. reflection
5. organising the to-do list

6. idea capture
7. collaboration

The following sections describe each pattern in turn.

5.6.1 *Planning activity*

Some composers used LiveTodos to plan their work ahead of time, such as Composer 5, who created an initial plan at the start of the session. As previously mentioned (see *Using to-do lists to plan and focus*), writing a to-do list can help make composers more aware of their own activity, and in particular, writing down their plans as a to-do list seems to make composers more aware of how much time they spend investigating possibilities.

> "it helps check myself, so it's like, 'oh, I'm spending too long on this thing that isn't working, what else am I doing' (Composer 4)

However, a clear drawback of creating a detailed plan of work is that it requires a significant upfront investment of time and effort. As the composition evolves over time and new possibilities are discovered, the composer's intention may change causing the original plan to become redundant, making it even more difficult to justify the initial costs. One strategy that was used by some composers to mitigate this issue was *multiple* planning phases throughout the composition process: for example, Composer 3 created a second batch of list items at 47-51 minutes to add further detail to their initial plan.

5.6.2 *Journalling activity*

While many to-do list items are reminders of plans for action at some point in the future, composers also created list items that described something they were *already* doing, or something they were about to immediately begin to do.

Composer 4 created a list item "Choose key sound" which illustrates this well. They entered "Choose key sound", then immediately choose a keyboard sound. Following this, they then added the note "Create an incidental melodic line", and then recorded a melody line. They started to look for a MIDI effect, they then added the tag "MIDI fx", then added a MIDI effect to the track. On each occasion, LiveTodos was not used to create a *reminder* of something to do in future: it was used to capture a *record* of the work being carried out at that time.

A clear advantage of this kind of behaviour is that the record created is an accurate and exact one, as opposed to a plan of intended work, which might not accurately represent what was eventually carried out. Composer 3 spoke about how useful an accurate record could be:

> "because I think something that I, have found to be quite important is having a… a sort of consistent workflow, and maybe being able to have a list of to-dos in a certain order, would help you to achieve that, that kind of consistency? That's how, I think um, that does help you to get a consistent, sound, almost like a signature sound that you have, I think? And perhaps, having as I say, the kind of um, list of to-dos that you can reuse for the next track, um, would be quite useful towards that, maybe." (Composer 3)

This desire has been also noted by Duignan (2008) who reported that composers were so keen to "fine-tune or recreate the same process at a later point" that they would document the history of a process. While it might be possible to use automatically collecting logs of activity for this purpose (for example, Koda (2011) suggest that "a simple listed history of processes applied to individual audio regions would be immensely useful in documenting the concoction of processes that might create similar audio effects"), to-do list items describe activity at a significantly higher level of abstraction that may be more useful than automatically logged data.

5.6.3 *Interleaving activities*

Composer 2 would often start work on a different to-do list item before they had finished their current task, and then alternate between different activities, using LiveTodos as a record of unfinished work to return to.

> "I'm sort of… um… wandering off the to-do list a bit, because I sort of made my synth sound and then I started making drums and then I realised I've got to actually sort out the timing of my synth sounds… so I need to finish… doing that before I do the drums…" (Composer 2)

In the above situation, Composer 2 marked the list item for the synth part "not done". This allowed them to suspend this work while ensuring it would be done later.

Composer 2 interleaved composition activities to a considerable extent, pausing and resuming work on the drums part on three different occasions to work on other parts of the composition.

> "so I'm working on the drums, and I thought 'oh, this is going to need some compression, at some point but I didn't want to do it now', to get kind of my loops, and timings kind of a little bit more organised so there's that instance where I kind of, you know, once I did that I was like 'ooh, yeah, now I'll do that', so yeah." (Composer 2)

Frequently switching between incomplete tasks in this way requires the composer to repeatedly reconstruct the context of the work they left off, something that other composers also mentioned using LiveTodos as part of:

> "you kind of go off on some rambling random path of things and often maybe it's a case of following how something develops and you kind of think 'oh, that's interesting, I'll' - you know - and you go off there, so, sometimes it's quite useful to be able to think 'actually where was my head at when I first started thinking about this?' " (Composer 1)

To help composers reconstruct the context they were working in, it may be useful to automatically capture information about the state of the DAW when a to-do list item is created. This can be useful for presenting to composers when they return to working on a list item for context reconstruction, but also as Composer 5 points out, it can be useful for finding other relevant list items related to that context:

> "the fact that you could like pick up the tags and channels and stuff, and that it was just there automatically, that was really really nice, because… I can then search be like 'OK, what ideas did I have around a melody', or 'what ideas did I have around this' " (Composer 5)

5.6.4 Reflection

Some composers remarked that LiveTodos was useful because it encouraged them to reflect more about their activity (see *Using to-do lists to plan and focus*). For one composer, to-do lists were part of a very active process of iterative refinement and reflection:

> "[the musician Object will] listen to a track when it's at the refinement phase, he'll listen to a track and make a list, do the things, make a list, and repeat the process hundreds of time… later on when it came to refining a track, or (inaudible)ing a track, I would use it [LiveTodos] a lot more, when it was sort of time to drill down and sort of, you know, address, why the track works or doesn't work" (Composer 3)

Another composer noted how it would also be useful in more passive reflective activity as well:

> "there's something about, this being a moment of reflection, and it's quite good that it happens away from the working window… something that would be quite nice, again I don't know what that would look like, but is a kind of an offline, I could imagine, sitting on the sofa in front of the telly, watching something and thinking" (Composer 1)

As Composer 1 notes, LiveTodos could be useful 'on the sofa': not all reflective activity needs to involve the DAW or a conventional desktop computer (see *Applicability of to-do lists*). Other studies have noted reflective activities that composers carry out which do not require the DAW, such as re-reading existing plans and notes (Donin and Theureau 2007) or listening to other compositions for comparison or inspiration (Auvinen 2019). A to-do list that is separate from the DAW could provide more support for these kinds of activities. When asked about what tools they used, some composers described work away from a conventional computer and DAW, such as Composers 4 and 5:

> 'I use a lot of iOS stuff, because I'm always travelling, so it wasn't practical to take a laptop with me on the train very often, I don't know if you know what trains are like nowadays, makes sense to first have a tablet in the cramped spaces' (Composer 4)

With composers subject to physical constraints such as working on public transport, composers are likely to value to-do lists due to the fact they can more easily be used on a mobile device, and because they can be used for reflective activities that do not require the full DAW interface.

5.6.5 *Organising the to-do list*

As mentioned in the earlier section *Organising to-do lists*, the composers actively requested more facilities to organise the to-do list. LiveTodos did not provide obvious mechanisms to allow composers to organise to-do lists. Composer 2 was observed during the study to be spontaneously attempting to drag to-do list items to reorder them before discovering that this was not possible.

Some studies (Duignan, Noble, and Biddle 2010; McCulloch 2014) have indicated that many composers using DAWs work linearly along the timeline from the start of the composition to the end, which suggests that ordering to-do lists by time is important, but also found that other composers work in a "patchwork" style (out of order) which suggests that different orderings might also be helpful.

LiveTodos did provide one method that composers could use to group their to-do lists, though it was not immediately obvious, and had not been pointed out to them. It was possible to tag multiple to-do list items with the same custom text string to indicate they were all part of a group. Composer 4 appears to have been attempting to do this, tagging several to-do list items with very similar custom tags ("Fx", "FX", and "MIDI fx"). However, as the to-do list could not be sorted, grouped, or filtered by tag, their tagging would not have provided many advantages to them, and they did not mention the feature at all in the closing interview.

5.6.6 *Idea capture*

Composers expected a considerable amount of flexibility, and as well as the purposes immediately suggested by the "to-do list" label, also described how LiveTodos was useful as a general notebook:

> "it was quite useful to sort of, just have an idea, write it down, have an idea, write it down, then edit through them afterwards… it's useful to be able to just sort of like sketch out to sort of jot down ideas, and use it almost as a scratch pad" (Composer 3)

> "often I find, because of the way I write music is so freeform, sometimes I'll have an idea that I think is amazing, and then I'll forget it, just completely gone! (laughs) and so having a tool that would allow me to actually write those down whether it be just jotting down a melody or

> keep it in one place" (Composer 5)

List items were used to capture materials ranging from rough notes on qualities the music should include ("Red", "Light", "Romance", "General note harsh, and stark") through to specific chords, notes, or key signatures to use later.

As well as using LiveTodos as a scrapbook for idea capture, composers used the to-do list in a more conventional way to describe planned future actions ("Audition synths", "Change key of piano part"). It might appear that two distinct tools are needed: one a scrapbook for storing material, and another for noting down future activities. However, this appears to be a false distinction, as list items can change function over time between these two categories, evolving from captured ideas into specific actions. For example, the idea "Romance" created by Composer 4 would later be annotated with the additional notes "Chords" and "FX" that represent actions to achieve the desired quality.

Some list items entered by composers were so vague ("String sounds", "Tape saturator? Reverb?") that it is difficult to distinguish whether they might be considered a captured idea or a plan for future action. As composers in this study appeared to not to draw a clear distinction between idea capture and planning activity when they made use of LiveTodos, it appears important to consider idea capture a part of to-do list functionality within the context of composition.

5.6.7 Collaboration

Composers 3 and 4 both spoke about how they thought LiveTodos would be especially useful when collaborating with others:

> "there's been so many times where I'd been working on a track with somebody in a studio or gone to visit them maybe, um, and… we've sort of like gone our separate ways with a half-finished track, to then have that [LiveTodos] as a kind of like shared idea kind of space, would be very very useful, we could literally just say 'this part, this is what needs to happen' and have it, you know, closely tied to it" (Composer 3)

> "in a collaboration, then it's important that, that we're on the same page, it doesn't wander off, there's nothing more frustrating than waiting for somebody else to finish what they're doing before you move on to the

> next stage, so, uh, to be professional, it would be helpful to, try and keep more focus, have any aids like this to keep more focused" (Composer 4)

As composers were observed working independently in this study, there is limited evidence of to-do lists being used for the co-ordination of musical collaboration from which to draw conclusions, and it will not be discussed further here. For further investigation of collaborative practices in composition and the artefacts that are used to support them, interested readers may wish to consult Nabavian (2010).

5.7 Discussion

The study described in this chapter has some obvious limitations: it is a small scale study of only a small number of composers (despite composers being a highly diverse group with idiosyncratic needs) and there is no consideration of collaborative activity (despite the obvious importance of collaboration to some or perhaps even most composers). As a result of these limitations, the data provides an unavoidably incomplete catalogue of the activities and needs of composers relating to PTM, but which suggests that to-do lists in DAWs appear to be useful to many composers despite their differences.

This section will now use the data to contribute a proposed explanation of *why* PTM features in DAWs might be useful, and based on this model, suggest specific features that might be important to consider when designing PTM features in DAWs.

5.7.1 Why to-do lists are useful in DAWs

The previous chapter, *Representation design in DAWs*, discussed composition tools as external representations as part of the cognitive processes of composers. Similarly, a software to-do list also acts as a cognitive artefact and an external representation.

This section applies this approach to LiveTodos to help understand the findings from the study.

Considering LiveTodos as an external representation leads to the question of what LiveTodos is *a representation of*. Compared to representations in DAWs like *Live*, a to-do list is a representation that tends to emphasise process rather than product, past and future tasks rather than current state, and abstract notations rather than auditionable music. There are two clear benefits to composers from using representations with these emphases:

Firstly, they *prime composers to think in terms of activity*, suggesting reflecting on tasks, processes, and plans. The composers described how they paid much more attention to recording their creative intention when they used LiveTodos than they would normally. Having a representation of their tasks and plans may increase how much attention composers pay to how their time is spent, how their work has diverged from their initial intention (see *Planning activity*), and their awareness of their own working practices (see *Reflection*).

Secondly, they *afford particular patterns of use*, making operations relating to tasks, processes, and plans easier to carry out. Having a representation of tasks and plans may make it easier to prioritise tasks (see *Organising to-do lists*), track partially completed work (see *Interleaving activities*), and repeat patterns of actions across compositions (see *Journalling activity*).

Conversely, the lack of to-do list functionality might also affect composers by encouraging them to not reflect on activity, and discourage some activities where no convenient way can be found to undertake them.

However, it is important to note that given the great degree of variation seen between the composers that were studied, the applicability of to-do lists will inevitably vary between composers and even between projects. Even within the five composers in this study, composers had a variety of experience of professional work and formal training, had different types of employment and commissions, did or did not compose for live performance, and had differing skill levels at particular DAW features, MIDI input controllers, and conventional instrumental performance. With such variety, to-do lists obviously may not be relevant to certain composers, certain stages in the processes of a composer, and certain types of composition, and DAWs should not be designed so that composers are required to use to-do lists.

5.7.2 Integrating to-do lists with the DAW

From initial conception, LiveTodos was designed to integrate with Live to provide an advantage over using paper notes. Integration is important as the representation of the composition in Live contains many parts that can be renamed or rearranged, and without such integration, composers would likely be required to repeatedly update to-do list items. Composers did not respond negatively to these features, and some composers made positive remarks about the integration. Within the context of Live, this is perhaps not surprising, as composers using Live already make use of switching between two different coordinated views of the composition that afford different patterns of use (the "Session" and "Arrangement" views).

Composers described their desire to make use of a to-do list tool for many different purposes, and carefully considering how integrating the to-do list with the DAW could support each of these activities potentially has significant benefits to composers. When *organising activities*, a useful ordering is by instrument or by group of instruments (see *Organising to-do lists*), which requires awareness of the track hierarchy within the DAW, as well as other methods of grouping in the DAW such as track colours. Another useful ordering is by timeline order (see *Organising the to-do list*), which would also require integration with the DAW, as DAWs frequently allow manipulation and reordering of the timeline. *Reflection* often takes place away from the DAW, and the ability to easily play back a rendered copy of the latest state of the composition would be particularly useful in this situation (and may provide the additional advantage of helping focus a composer's attention by removing the ability to be distracted by editing the composition). *Planning* and *idea capture* tasks could benefit greatly from the ability to refer to elements within libraries in the DAW (such as particular effects, instruments, or clips) that the composer intends to use, along with the ability to use the facilities in the DAW for recording rough takes. Finally, *interleaving activities* would benefit from automatically recording the context of current work in the DAW to aid the composer in reconstructing context when work is resumed (see *Interleaving activities*).

However, it is important to remember that the processes used by composers may involve moving between multiple DAWs, and to properly support such composers it may be necessary to support integration with (and migration between) multiple DAWs.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 5.1

- Allow to-do list items to be *annotated* with links to objects in the composition or in DAW libraries
- Allow links from list items to *groups* created in the DAW, not just atomic objects
- Automatically *update* information in the to-do list when properties of these objects change
- Provide *playback of a rendered composition* for use away from the DAW
- Record *contextual information* about the circumstances list items are created or modified in

5.7.3 Increasing visibility of activity

Composition is an open-ended problem in which there is no one "correct" solution and composers are required to make use of decision-making techniques to avoid "analysis paralysis". Strategies for managing time use appeared to be very relevant to the composers in this study, as they described work involving deadlines and time pressures imposed by clients and schedule commitments to collaborators on musical projects, along with time pressures created by working an unrelated full-time job in addition to their composing (see *Recruitment and participants*).

General-purpose to-do list tools often try to support time allocation by providing features for prioritisation, but these features are often of limited utility in real use (Kamsin, Blandford, and Cox 2012), and composers in this study appeared to have a greater interest in increasing awareness of their own activities than in explicitly prioritising tasks (see *Organising to-do lists*). This suggests that composers may benefit from to-do list features that increase the visibility of composition activity, such as recording current activity, overviews, and filtering of lists.

Allowing composers to record their current activity may be useful as composers appear to be interested in recording not just plans for future work, but also their current and past activity. To-do list items created through journalling activity could be used by the composer to help navigate and understand their interaction history by providing a higher-level description in the composer's own language (a similar approach for navigating interaction histories has been suggested by Akiyama and Oore (2014)). Composers could also use these to-do list items to more easily reuse successful techniques from previous compositions (see *Journalling activity*).

While the to-do list itself could become long and unwieldy, the ability to view an overview of the entire list (which Bellotti et al. (2004) would describe as a *task vista*) may help to keep perspective on what is left to be done and what has been accomplished. This increased awareness of remaining time could allow composers to deliberately limit the time they spend considering possible options they are not focused on, and instead prioritise musical aspects they care about most. For similar reasons, composers also may benefit from the ability to filter a long to-do list: not just as a search facility, but also as a way to represent the musical aspects that the composer's activity is currently focusing upon, such as a particular instrument, type of effect, or something uniquely specific to that composer's process.

In summary, this section suggests these possibilities for DAW design:

Design suggestions 5.2

- Allow composers to easily record their *current* activity
- Provide an *overview* of the whole to-do list
- Provide the ability to *filter and search* the to-do list
- Provide facilities to *reuse* past to-do list activity

5.8 Conclusions

This chapter has described a study of how pro-amateur composers respond to a prototype system providing a representation that uses the *task lists* technique, aiming to answer three main questions (see *Introduction*):

How are to-do lists useful to pro-amateur composers working with a DAW?

To-do lists appear to be useful to the composers studied, with all composers able to use the system as part of their working process and discuss situations in which it might be applicable to their work. While the composers stated that to-do lists might only be applicable to particular types and contexts of composition they carried out, all felt that they would be useful in at least some situations. Task lists representations should be designed with the assumption that they will be used in many different contexts, including away from the DAW itself, and that different features will be useful to the users in each context.

This study suggests that to-do lists are useful because they **prime composers to think in terms of activity**, and because they **afford particular patterns of use**. In interview, the composers spoke about how to-do lists were useful in activities relating to planning and focusing, the importance of being able to change to-do list items over time, and the usefulness of organising the to-do list.

What activities do pro-amateur composers use to-do lists to carry out?

Activities observed in this study that were being supported by to-do lists were planning, journalling, interleaving of work, reflection, idea capture, and collaboration, as well as organising the to-do list itself. To best support these activities, it is important to provide functionality *integrating to-do lists with the DAW*, and to provide features for *increasing visibility of activity*.

How could the representations in DAWs be changed to better support these activities?

Nine specific design suggestions have been made, based on these observations:

1. Allow to-do list items to be *annotated* with links to objects in the composition or in DAW libraries
2. Allow links from list items to *groups* created in the DAW, not just atomic objects
3. Automatically *update* information in the to-do list when properties of these objects change
4. Provide *playback of a rendered composition* for use away from the DAW
5. Record *contextual information* about the circumstances list items are created or modified in
6. Allow composers to easily record their *current* activity
7. Provide an *overview* of the whole to-do list
8. Provide the ability to *filter and search* the to-do list
9. Provide facilities to *reuse* past to-do list activity

Chapter 6

Conclusion

The previous chapters have described a survey of the existing research into composition software use, a study of composers at work, a literature review used to produce a set of techniques for designing representations of music in composition software, and a detailed study of one such technique and recommendations for designing software that implements it.

This chapter will summarise the **preceding chapters**, discuss the **findings** and **contributions** of this thesis, **critically evaluate** the research activities undertaken and suggest lessons learned for future research, and describe **future work** remaining in this area beyond that described in this thesis.

6.1 Summary of preceding chapters

Chapter 1 discussed the historical context and development of DAWs, their user interfaces, and the types of people who have used them. The scope of the thesis emerged from this discussion and narrowed from digital audio workstations in general to their *design*, for their use by *pro-amateur* composers, in particular those working *independently*. The three main research questions of the thesis ("How do pro-amateur composers use DAWs?", "How could the representations in DAWs be improved?", and "How could pro-amateur composers use task lists in concert with a DAW?") were stated. Finally, the overall methodological approach, which combines literature review with studies of working composers, was described.

Chapter 2 reviewed the literature relating to composition software. The review illustrated the complexity of software-based composition through the diverse types of composers and methods of composing that have been previously observed. Specific recommendations for studies of composition software were proposed, which would be used in subsequent chapters to inform the design of two studies of composers: the triangulation of multiple data capture methods, avoiding study designs that might artificially constrain composers, capturing the use of external aids outside of the DAW, and mitigating constraints the DAWs impose on composers. The review also identified several common themes in existing research. The issue of the design of representations emerged as an important theme in the review, with composers often seen creating their own representations outside of the composition software, and with those representations found in the DAW software having the potential to constrain the work of composers through representational determinism.

Chapter 3 described a study of pro-amateur composers who were observed using DAWs in a naturalistic setting while carrying out their choice of composition task. Composers made use of several strategies to selectively allocate their time and effort in order to focus on musical aspects that were important or interesting to them. In agreement with Chapter 2, composers in this study were again observed to make use of representations other than those in the DAW. Composers also *co-ordinated* their use of these representations together with their use of the DAW. An explanation given for this behaviour was that the composers avoided the use of the DAW for some composition activities, potentially because those activities appeared to be poorly supported by the DAW, suggesting that combining the benefits of different representations together allowed them to avoid being constrained.

Chapter 4 first introduced the various approaches that have been used to discuss how composers work with the representations around them in their environment as part of their creative processes. The chapter described a set of representation design techniques that could be used to better support many composition activities and strategies:

1. *selective representation*: allowing composers to select which musical aspects are to be represented
2. *diverse media types*: allowing composers to use more types of media in their representations
3. *structured representations*: allowing composers to define structures of objects

within representations

4. *incomplete specification*: allowing composers to represent elements of the composition without giving full details on how they might be played back
5. *representing alternatives*: allowing composers to represent multiple possible alternative options
6. *task lists*: allowing composers to keep a list of planned future activities
7. *representing history*: representing the previous activity of the composer and states of the composition

For each technique, the chapter provided specific suggestions for how they could be implemented in a design, based on a review of the approaches used in previous research into composition software and creativity tools, and these were summarised in tables at the end of the chapter. The technique of *task lists* was identified as relatively unexplored and a promising area for further research.

Chapter 5 focused on the specific technique *task lists* in more detail, using a to-do list website that integrated with the DAW *Ableton Live* to investigate how composers would respond to the technique. A study of five composers concluded that the technique is useful because a to-do list primes composers to think in terms of activity and affords particular patterns of use. A mechanism by which this may function is the to-do list representation increasing the visibility of activity to the composer, and suggestions were made for how designers could best achieve this. In order to achieve these benefits, good integration between a to-do list and the DAW is an important consideration, and further design suggestions were given for how to integrate to-do lists and DAWs.

This chapter will now further discuss the findings of these chapters and their relevance to the wider literature.

6.2 Research questions

This thesis asked three questions:

1. how do pro-amateur composers use DAWs?
2. how could the representations in DAWs be improved?
3. how do pro-amateur composers use to-do lists in concert with a DAW?

The following sections will discuss these questions.

6.2.1 How do pro-amateur composers use DAWs?

Chapter 2 describes how the process of composition using a DAW can vary between composers, with many different macro structures having been observed in the composition process, described in this thesis as staged, iterative, top-down, bottom-up, chronological, and patchwork. Across Chapters 2 and 3, this thesis suggests that pro-amateur composition also involves a variety of different component activities, including planning, collecting and capturing, auditioning, reflection, improvisation, mixing, combining alternatives, software management, and collaboration. Some of these activities, such as planning, capturing ideas, and reflecting, were also observed in the activity of composers in Chapter 5 during their use of a DAW with an integrated to-do list.

Co-ordinated use of multiple representations was observed being used as part of the processes of composers using DAWs in Chapter 3, who created secondary representations of musical information which used different media and which presented information in different ways. In Chapter 5, composers were asked to use the LiveTodos to-do list website as a second representation alongside the DAW, and all commented that they would find this useful at some points during their composition process. The findings in these chapters that co-ordinating representations is important to composers using software match the findings in Bainbridge, Novak, and Cunningham (2010) of a need for "directly associating, or grouping different types of media together". In Chapter 4, interfaces to co-ordinate multiple representations together were discussed (see *Selective representation*), and in Chapter 5 the importance of integration between representations was discussed in detail. Most of the research composition systems described in the literature reviews in Chapter 2 and Chapter 4 have replaced the DAW interface entirely, rather than augmenting it with an additional representation, and creating representations that co-ordinate and integrate well with the existing representations in a DAW appears less explored. Several composers in Chapter 5 discussed their regular use of mobile devices in music making, and future studies could potentially further explore the possibilities of additional representations presented on mobile devices.

Selectively allocating time and effort was observed as a common behaviour for the composers in Chapter 3, with three main strategies used by composers towards this end: habituation, limited exploration, and self-constraining. Composers in the study in Chapter 5 appeared to have similar priorities, speaking about the positive effects of representing their use of time and effort so that it is visible to them (see *Using to-do lists to plan and focus*). Of the strategies for allocating time and effort that were identified in Chapter 3, self-constraining is most discussed in the creativity tools literature (Candy and Edmonds 1997), with some discussion of reuse of material and techniques (Akiyama and Oore 2014; Visser 2006), and very little discussion of how creative practitioners might choose to deliberately limit their exploration of a creative space. When discussion of this behaviour of limiting exploration has occurred, previous authors have often viewed it as an undesirable trait, aiming to "decrease the level of exploratory satisficing" (Adams, Gonzalez, and Latulipe 2014) and to solve the problem that "people are so often averse to exploration" (Ngoon, Walker, and Klemmer 2019). This thesis has discussed satisficing in creative exploration as a *positive* strategy instead; one that allows composers to focus their time use to their benefit. A useful balance between these viewpoints is put forward by Gonzalez (2015): "Users should only be satisficing because of external constraints and preferences, not because the parameter space exploration is difficult, slow, or tedious". The design of representations in creative tools can negatively affect users by encouraging them to satisfice on things they do not want to; but it should also assist them to satisfice where they wish to, in order to focus time and effort on those areas they wish to more fully explore.

6.2.2 How could the representations in DAWs be improved?

The review in Chapter 2 and field study in Chapter 3 identified that existing representations of compositions used in composition software often do not fully support the diversity of processes, activities, and focusing techniques used. In Chapter 4, a set of seven techniques are proposed for expanding these representations: *selective representation*, *diverse media types*, *structured representations*, *incomplete specification*, *representing alternatives*, *task lists*, and *representing history*.

For each technique, these findings have previously been discussed in the context of the existing literature of research into composition software and creativity tools, as can be found in the subsections *In existing composition software* and *In other software categories* for each technique listed in Chapter 4.

The activity of and comments made by the composers in the study in Chapter 5 provide additional evidence for the usefulness of many of these techniques. Composers in Chapter 5 talked about their desire to reflect on history and reuse parts of to-do lists in future projects (see *Journalling activity*), were observed creating to-do list items that may denote multiple possible alternative possibilities (see *Changing to-do list items over time*), and talked about their need to group and link to-do list items to organise their work (see *Organising to-do lists*). The data (see *Object-oriented records* in Chapter 5) indicates composers comfortably switching repeatedly between different views in *Live* which represent different sets of information and allow composers to perform different types of actions, suggesting that these composers also find *selective representation* useful.

Further requirements for the *task lists* technique were identified in Chapter 5, and these can be found in the sections *Integrating to-do lists with the DAW* and *Increasing visibility of activity*.

Chapter 5 recommends careful consideration of the integration between composition software and representations of activity. As has been previously explained in *LiveTodos: interface and evolution*, LiveTodos made use of only a small amount of data from the DAW to provide an integrated to-do list representation. Read-only access to the object model of the DAW, and the ability to subscribe to (or poll for) changes in properties of those objects, was almost all that was required to create LiveTodos. This suggests that an alternative approach to directly supporting *task lists* within the DAW is to provide support within a DAW for integrating with an *external* to-do list tool.

In order to provide this functionality, the DAW would be required to expose a list of objects that could be referred to in to-do list items (ideally through a standard mechanism such as a URI), each populated with common properties that assist users in identifying those objects (such as a human-readable name, colours, or thumbnail images), and properties that can be used for sorting or grouping (such as by category, or a numeric, interval or time value). None of these requirements are specific to a DAW, suggesting that this may be a useful technique to use in other types of creative software.

6.2.3 *How do pro-amateur composers use to-do lists in concert with a DAW?*

Chapter 4 discussed some initial ideas regarding how to-do lists could be used by composers, including planning future work, reviewing existing plans to decide on areas to focus on, and supporting the postponing and resumption of work. The study in Chapter 5 provided more detailed information on how composers use to-do lists, with composers in that study observed to use to-do lists with a DAW in order to keep a record of work that had been carried out in the past, for interleaving different activities together, for reflecting on their work, and as a way of organising their plans. Chapter 5 also discusses the use to-do lists along with the DAW as a way of keeping track of ideas that did not fit well into the rigidly defined and temporally-focused representation of the DAW, such as for capture of thoughts and rough ideas. In general, composers appear to use to-do lists along with the DAW because the to-do list allows them to more easily carry out these kinds of activities. They also find to-do lists useful because their use influences them to consider the activities involved in their work.

As in previous chapters, this could again be thought about in terms of representational determinism: the presence of a particular representation can affect what composers can think about or do easily, and this may direct composers towards particular behaviours. While Chapter 5 discussed the potential positive influence of using to-do lists to represent activity, the preceding chapters of this thesis repeatedly discussed the potential negative effects of representational determinism, and these are also important to consider. A previous study of the use of representations of activity (including to-do lists) has suggested that such representations can influence people both to *feel* busy and *enact* busyness: "even when engaged in fun or leisure activities, our participants seemed task-oriented" (Leshed and Sengers 2011). A potential pitfall of representing activity to composers is that it could cause them to become more focused on getting tasks done, and less on experimentation, or on enjoying the act of composition.

Comparing existing studies of personal task management tools with the study in Chapter 5, two clear contrasts are visible. Studies of these tools normally recruit participants to study task management in the context of employment (Bellotti et al. 2004; Gonzalez, Galicia, and Favela 2008) or academic study (Haraty et al.

2012; Kamsin 2014), whereas this thesis has deliberately avoided studying task management as part of someone's main employment. Chapter 5 identified common patterns in the use of to-do lists by composers, and while many of these can also be seen in the lists of activity patterns in task management identified by Gonzalez, Galicia, and Favela (2008), Kamsin (2014), or Haraty et al. (2012), those authors place significantly more emphasis on prioritisation and scheduling of tasks than Chapter 5 does.

A possible explanation for these difference in emphasis may be found in the nature of tasks that composers perform. Many of the tasks that composers undertake have a different character: their duration is often *flexible* depending on how composers choose to weigh the importance of a task, and how the composer chooses to allocate their time and effort. As a result, instead of a conventional prioritisation approach in which less important tasks are postponed or left unfinished, composers will instead simply spend less time on those tasks.

As has previously been described in Chapter 3, composers use techniques such as habitual re-use to do this. Chapter 5 has also suggested that good visibility of their activity can help composers selectively allocate time and effort. Interestingly, these kinds of behaviours are *not* exclusively limited to a creative or pro-amateur context: Yli-Kauhaluoma (2009) describes how chemistry research and development professionals similarly use techniques such as re-use and representing (or "documenting", in their terminology) activity to manage their use of time and to increase the amount of time they have available to them. This suggests that while there is less emphasis on these strategies seen in other contexts, they are still important to consider in the design of task management tools. An alternative explanation for the different emphasis in activity patterns seen in Chapter 5 can be found in the theses of Kamsin (2014) and Wu (2005), who found in their studies that significantly different task and time management behaviours were found when comparing experienced senior academics and less experienced and more junior staff, so differences in behaviours between pro-amateur composers and full-time employees might reasonably be expected.

While Haraty, McGrenere, and Tang (2016) recommend that "Non-PTM tools should offer basic support for PTM", research in this area is limited. Most studies of task management tools have been in collaborative settings, in which an obvious choice is to integrate task management support into the communication system used for collaboration, which has usually been an email application. This thesis has focused on individual work: this has instead lead to integration with a *domain-specific* tool, in this case, a DAW.

6.3 Summary of contributions

This thesis has made four main contributions.

1. *Literature review*: Chapter 2 contributes a review of the literature in this area, and identifies which populations of composers have been studied, what methodologies have been used to study them, and what settings studies have been carried out in. The review also summarises the findings of existing studies on what the structure of the composition process is, what component activities are involved, problems observed with existing DAW tools, and how composers use other aids apart from the DAW. Finally, guidelines are suggested for the design of studies of composition software.
2. *Fine-grained description of composition*: Chapter 3 contributes a detailed account of four composers working for a short period. Three main patterns are identified in their activity: the lack of support provided by the DAW for many of the activities that were part of the creative process of these composers, the coordinated use of multiple representations by the composers, and the tactics used by composers to selectively spend their time and effort on areas of interest to them.
3. *Catalogue of techniques*: Chapter 4 contributes a catalogue of techniques that may be used in DAWs to support particular process structures, activities, and focusing techniques that are part of the work of the pro-amateur composers studied. For each technique, suggestions about how designers of DAWs can support the technique are given based on a review of relevant literature.
4. *Fine-grained description of reactions to a specific technique*: Chapter 5 contributes a detailed description of the use of a to-do list website integrated with a commercial DAW by five pro-amateur composers, along with their responses to subsequent interview questions, and further explores the *task list* technique described in Chapter 4. Four main themes were produced from the thematic analysis, and seven common patterns of activity identified, with

specific suggestions provided for how to design to-do lists for composers.

6.4 Critical reflection

This thesis has brought together studies of composer activity with multiple in-depth reviews of the literature on composition software and creativity tools. In this process, unexplored areas such as personal task management in creative activities have been identified and explored, and findings from these areas could be relevant to researchers outside of composition software.

The studies described in this thesis involved a case study approach of nine composers. Considering the nine composers studied, a great amount of variety was seen between the composers, including but not limited to their prior experience of full-time composition work and academic musical training, their expertise with particular features within DAWs, how they release music to the public (if at all), how much of their composition is done as part of part-time or commission-based employment, whether their compositions are intended to be used in live performance, and how much their process involves skills at playing other instruments or singing. This thesis has investigated composers from many different points within this space to help ensure that the findings have general relevance to other pro-amateur composers, but with the degree of diversity seen in pro-amateur composers there is considerable room for further work in this area, and some potential starting points are described in the following section *Future work*.

In Chapter 2, the studies found in the literature were assessed to determine the quality of their study design, using a set of desirable attributes derived from Collins (2005). The studies in this thesis featured composers using professional-quality tools, and observed composers working across a variety of musical styles and with a diverse set of musical backgrounds, producing data that has illustrated the diversity of approaches used by pro-amateurs. However, the studies in Chapter 3 and 5 did not include a longitudinal element, limiting which behaviours could be observed to those that were evident within a relatively brief period. This suggests there may be room for further investigations with this population, and while attempts were made to encourage the composers in Chapter 5 to participate in further research by using LiveTodos at home, this was unsuccessful. Interestingly, similar problems have also been seen in other studies within related areas of research. Bernstein et al.

(2008) describe a similar study to that described in Chapter 5, in which a prototype tool for managing personal information was created, and a longitudinal approach was desired. The participants in this study displayed similar behaviour and did not incorporate the use of the tool into their routine once they left the lab environment. As a result, that study failed to gather data on longitudinal use. Reflecting on the study, they describe four important factors that study designers should consider in order to avoid a similar lack of acceptance by users: research scope, the design and prototyping process, the study methodology, and the population chosen. Considering these factors, future studies of composition tools that wish to study use over an extended time period might benefit from deliberately studying a specific class of composer (for example, in a specific genre), using a *participatory design* process, or assuming that initial participants using early software versions will not continue to use software outside a lab context.

A number of challenges encountered in the studies of composers related to data capture. Video recording of the composer proved challenging in the study described in Chapter 3, as composers would rapidly switch between using difficult instruments, devices, and paper notes, requiring a time-consuming reconstruction process using audio recordings to generate data about events that had happened off-camera. The study in Chapter 5 significantly improved upon the study described in Chapter 3 by making use of an *action camera* to capture a high-resolution video recording with a wide field of view. This provided significant benefits during the data analysis, but further challenges with data reconstruction occurred when attempting to analyse the data from the study in Chapter 5: synchronising multiple streams of captured data. McCulloch (2014) also reported this issue when studying the use of composition software. Some steps were taken to avoid these problems, including deliberately clapping before activity started so that audio and video recordings could be easily be synchronised, but more up-front planning could have made this synchronisation much easier.

This thesis has contained significant literature review content, including much of both Chapter 2 and Chapter 4, and the extent to which the findings from user studies are related to the literature is a strength of this thesis. The literature that was surveyed included not only research on composition software but also significant review of research into creativity tools. In addition to more general findings about composition software, these reviews have identified opportunities for further research, and provided specific suggestions for designers of digital audio workstations and future researchers of composition software.

6.5 Future work

In the previous section, an important set of considerations proposed in Bernstein et al. (2008) for longitudinal studies of personal information tools was discussed. If future studies of pro-amateur composers use similar methodologies, involving building prototype systems for longitudinal tests in the wild, these studies will need to pay particular attention to those issues. However, it is not sufficient for other studies merely to run a longitudinal study; it also important to consider that *many* compositions may be in progress at any time (Bainbridge, Novak, and Cunningham 2010, 151). This thesis has not investigated how compositional processes interact with each other beyond mentioning re-use between projects. Future work in this area could potentially reveal more important activities that are involved in composition by seeking to understand how composers manage multiple compositions at once.

A question remains as to whether the composers observed in Chapter 3 were actually working on their own. The composers in this study were in isolation for the duration of the study. However, the composers were all either using source material that was either previously produced in collaboration with others, or using art created by others as an input to their musical processes. A possible area for future investigation that was not covered by this study is how composers that primarily work alone use recycling processes to add a collaborative element to their work - is collaborating with your past self similar to collaborative music making, or does it have important differences?

Chapter 3 suggests that composers make use of a number of strategies to help make creative decisions. This thesis does not investigate the literature around *decision-making* and this may be a fruitful avenue for explaining the strategies that composers use and may potentially allow the identification of further strategies.

Rather than describing a clear population of DAW users, the label pro-amateur seems instead to describe a space containing many different types of composers, potentially with quite different needs. While the nine composers studied in this thesis provide a variety of different points within this space, more insights might be gained from studying composers with different experiences, genres, processes, and types of output.

A further possibility for improving the representations available in DAWs was suggested by the literature: *transformation* of representations. In transformation processes, representations of music are used to generate dramatically different musical output by altering the semantic interpretation of the representation. In the small set of composers studied in Chapter 3, and the restricted task and environment in Chapter 5, there was insufficient evidence of this activity to confidently recommend that it was a useful technique. However, evidence from the literature studying academic composers (Thiebaut, Healey, and Kinns 2008; Healey and Thiebaut 2007) suggests this might be an important part of composition. Future studies of pro-amateur composers - particularly those taking a more longitudinal approach - should pay particular attention to whether this practice is part of pro-amateur activity.

The technique of *task lists* was identified as particularly unexplored, and then further investigated in Chapter 5, but relatively little research has been carried out into using *incomplete specification* and *representing alternatives* in a DAW. For both techniques, Chapter 4 has referenced authors discussing how useful these techniques could be across many different creative domains, suggesting that these would be productive avenues for further exploration.

Finally, while this thesis spent considerable time discussing the design of DAWs, one route for researching the design of DAWs was not used: the developers of the DAWs were not contacted or asked for information about the reasons behind their design decisions. This avenue of investigation could provide useful information about DAW design, but in addition to not being used in this thesis, no studies found in the literature review have yet used this approach.

Chapter 7

Bibliography

Abrams, Steven, Ralph Bellofatto, Robert Fuhrer, Daniel Oppenheim, James Wright, Richard Boulanger, Neil Leonard, David Mash, Michael Rendish, and Joe Smith. 2002. 'QSketcher: An Environment for Composing Music for Film'. In *Proceedings of the 4th Conference on Creativity & Cognition*, 157–64. C&C '02. Loughborough, UK: ACM. https://doi.org/10.1145/581710.581734.

Adams, Alexander Travis, Berto Gonzalez, and Celine Latulipe. 2014. 'SonicExplorer: Fluid Exploration of Audio Parameters'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 237–46. CHI '14. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2556288.2557206.

Akiyama, Yasushi, and Sageev Oore. 2014. 'Designing Tools to Support Multimedia Authoring by Incorporating Problem-Solving Strategies'. *Journal of Multimedia Theory and Applications* 1: 11–19. https://doi.org/10.11159/jmta.2014.002.

Amitani, Shigeki, and Koichi Hori. 2002. 'Supporting Musical Composition by Externalizing the Composer's Mental Space'. In *Proceedings of the 4th Conference on Creativity & Cognition*, 165–72. C&C '02. Loughborough, UK: ACM. https://doi.org/10.1145/581710.581735.

Assayag, Gérard, Camilo Rueda, Mikael Laurson, Carlos Agon, and Olivier Delerue. 1999. 'Computer-Assisted Composition at Ircam: From Patchwork to Openmusic'. *Computer Music Journal* 23 (3): 59–72. https://doi.org/10.1162/014892699559896.

Auvinen, Tuomas. 2019. 'The Music Producer as Creative Agent: Studio Production, Technology and Cultural Space in the Work of Three Finnish Producers'. PhD thesis, University of Turku.

Bainbridge, David, Brook J. Novak, and Sally Jo Cunningham. 2010. 'A User-Centered Design of a Personal Digital Library for Music Exploration'. In *Proceedings of the 10th Annual Joint Conference on Digital Libraries*, 149–58. https://doi.org/10.1145/1816123.1816145.

Bainbridge, David, Brook J Novak, and Sally Jo Cunningham. 2012. 'A Spatial Hypertext-Based, Personal Digital Library for Capturing and Organizing Musical Moments'. *International Journal on Digital Libraries* 12 (2-3): 89–103. https://doi.org/10.1007/s00799-012-0090-3.

Balaban, Mira. 1996. 'The Music Structures Approach to Knowledge Representation for Music Processing'. *Computer Music Journal* 20 (2): 96–111.

Balaban, Mira, and Michael Elhadad. 1999. 'On the Need for Visual Formalisms in Music Processing'. *Leonardo* 32 (2): 127–34.

Barlindhaug, Gaute. 2007. 'Analog Sound in the Age of Digital Tools: The Story of the Failure of Digital Technology'. In *A Document (Re)turn: Contributions from a Research Field in Transition*, edited by Roswitha Skare, Nils Windfeld Lund, and Andreas Varheim, 73–93. Peter Lang.

Barr, Pippin, Robert Biddle, and James Noble. 2002. 'A Taxonomy of User-Interface Metaphors'. In *Proceedings of the SIGCHI-NZ Symposium on Computer-Human Interaction*, 25–30. CHINZ '02. Hamilton, New Zealand: ACM. https://doi.org/10.1145/2181216.2181221.

Bell, Adam Patrick. 2014. 'Trial-by-Fire: A Case Study of the Musician–Engineer Hybrid Role in the Home Studio'. *Journal of Music, Technology & Education* 7 (3): 295–312.

———. 2015. 'Can We Afford These Affordances? GarageBand and the Double-Edged Sword of the Digital Audio Workstation'. *Action, Criticism, and Theory for Music Education* 14 (1): 44–65.

Bell, Adam Patrick, Ethan Hein, and Jarrod Ratcliffe. 2015. 'Beyond Skeuomorphism: The Evolution of Music Production Software User Interface Metaphors'. *Journal on* *the Art of Record Production* 9.

Bellotti, Victoria, Brinda Dalal, Nathaniel Good, Peter Flynn, Daniel G. Bobrow, and Nicolas Ducheneaut. 2004. 'What a To-do: Studies of Task Management Towards the Design of a Personal Task List Manager'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 735–42. CHI '04. Vienna, Austria: ACM. https://doi.org/10.1145/985692.985785.

Benedetti, Luca, Holger Winnemöller, Massimiliano Corsini, and Roberto Scopigno. 2014. 'Painting with Bob: Assisted Creativity for Novices'. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, 419–28. UIST '14. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2642918.2647415.

Bennett, Joe. 2011. 'Collaborative Songwriting – the Ontology of Negotiated Creativity in Popular Music Studio Practice'. *Journal on the Art of Record Production* 5.

Bennett, Peter. 2010. 'The Representation and Control of Time in Tangible User Interfaces'. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*, 307–8. TEI '10. Cambridge, Massachusetts, USA: ACM. https://doi.org/10.1145/1709886.1709956.

Bernstein, Michael, Max Van Kleek, Deepali Khushraj, Rajeev Nayak, Curtis Liu, MC Schraefel, and David R Karger. 2008. 'Wicked Problems and Gnarly Results: Reflecting on Design and Evaluation Methods for Idiosyncratic Personal Information Management Tasks'. MIT-CSAIL-TR-2008-007. Massachusetts Institute of Technology.

Bertelsen, Olav W., Morten Breinbjerg, and Søren Pold. 2009. 'Emerging Materiality: Reflections on Creative Use of Software in Electronic Music Composition'. *Leonardo* 42 (3): 197–202.

Berthouzoz, Floraine, Wilmot Li, and Maneesh Agrawala. 2012. 'Tools for Placing Cuts and Transitions in Interview Video'. *ACM Transactions on Graphics* 31 (4). https://doi.org/10.1145/2185520.2185563.

Boehner, Kirsten, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. 'How HCI Interprets the Probes'. In *Proceedings of the SIGCHI Conference on Human*

Factors in Computing Systems, 1077–86. CHI '07. San Jose, California, USA: ACM. https://doi.org/10.1145/1240624.1240789.

Braun, Virginia, and Victoria Clarke. 2006. 'Using Thematic Analysis in Psychology'. *Qualitative Research in Psychology* 3 (2): 77–101.

Brooker, Phillip, and Wes Sharrock. 2016. 'Collaborative Music-Making with Digital Audio Workstations: The 'Nth Member' as a Heuristic Device for Understanding the Role of Technologies in Audio Composition'. *Symbolic Interaction* 39 (3): 463–83. https://doi.org/10.1002/symb.238.

Buxton, William, William Reeves, Ronald Baecker, and Leslie Mezei. 1978. 'The Use of Hierarchy and Instance in a Data Structure for Computer Music'. *Computer Music Journal* 2 (4): 10–20.

Buxton, William, Richard Sniderman, William Reeves, Sanand Patel, and Ronald Baecker. 1979. 'The Evolution of the SSSP Score Editing Tools'. *Computer Music Journal* 3 (4): 14–60.

Candy, Linda, and Ernest A. Edmonds. 1997. 'Supporting the Creative User: A Criteria-Based Approach to Interaction Design'. *Design Studies* 18 (2): 185–94.

Casares, Juan, A. Chris Long, Brad A. Myers, Rishi Bhatnagar, Scott M. Stevens, Laura Dabbish, Dan Yocum, and Albert Corbett. 2002. 'Simplifying Video Editing Using Metadata'. In *Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, 157–66. DIS '02. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/778712.778737.

Clark, Andy, and David Chalmers. 1998. 'The Extended Mind'. *Analysis* 58 (1): 7–19.

Cole, Steven James. 2011. 'The Prosumer and the Project Studio: The Battle for Distinction in the Field of Music Recording'. *Sociology* 45 (3): 447–63. https://doi.org/10.1177/0038038511399627.

Collins, David. 2005. 'A Synthesis Process Model of Creative Thinking in Music Composition'. *Psychology of Music* 33 (2): 193–216. https://doi.org/10.1177/0305735605050651.

Conley, Kenneth, and James Carpenter. 2007. 'Towel: Towards an Intelligent to-Do List.' In *AAAI Spring Symposium: Interaction Challenges for Intelligent Assistants*, 26–32.

Costa, Miguel, Nuno Correia, and Nuno Guimarães. 2002. 'Annotations as Multiple Perspectives of Video Content'. In *Proceedings of the Tenth ACM International Conference on Multimedia*, 283–86. MULTIMEDIA '02. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/641007.641065.

Coughlan, Tim. 2009. 'Understanding Creative Interaction: A Conceptual Framework for Use in the Design of Interactive Systems for Creative Activities'. PhD thesis, University of Bath.

Coughlan, Tim, and Peter Johnson. 2006. 'Interaction in Creative Tasks'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 531–40. CHI '06. Montréal, Québec, Canada: ACM. https://doi.org/10.1145/1124772.1124854.

———. 2008. 'Idea Management in Creative Lives'. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, 3081–6. CHI EA '08. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1358628.1358811.

———. 2009a. 'Designing Personal Information Management Systems for Creative Practitioners'. *Journal of Digital Information* 10 (5). https://journals.tdl.org/jodi/index.php/jodi/article/view/440.

———. 2009b. 'Understanding Productive, Structural and Longitudinal Interactions in the Design of Tools for Creative Activities'. In *Proceedings of the 7th ACM Conference on Creativity & Cognition*, 155–64. C&C '09. Berkeley, California, USA: ACM. https://doi.org/10.1145/1640233.1640258.

Dannenberg, Roger B. 1993. 'Music Representation Issues, Techniques, and Systems'. *Computer Music Journal* 17 (3): 20–30.

Dewey, Christopher, and Jonathan Wakefield. 2016. 'Novel Designs for the Audio Mixing Interface Based on Data Visualization First Principles'. In *Proceedings of the 140th AES International Convention*. Audio Engineering Society.

Dewey, Christopher, and Jonathan P Wakefield. 2017. 'Formal Usability Evaluation of Audio Track Widget Graphical Representation for Two-Dimensional Stage Audio Mixing Interface'. In *Proceedings of the 142nd AES International Convention*. Audio Engineering Society.

Donin, Nicolas, and Jacques Theureau. 2007. 'Theoretical and Methodological Issues Related to Long Term Creative Cognition: The Case of Musical Composition'. *Cognition, Technology & Work* 9 (4): 233–51. https://doi.org/10.1007/s10111-007-0082-z.

Dourish, Paul. 2004. *Where the Action Is: The Foundations of Embodied Interaction*. The MIT Press.

Duignan, Matthew. 2008. 'Computer Mediated Music Production: A Study of Abstraction and Activity'. PhD thesis, Victoria University of Wellington.

Duignan, Matthew, James Noble, Pippin Barr, and Robert Biddle. 2004. 'Metaphors for Electronic Music Production in Reason and Live'. In *Computer Human Interaction*, edited by Masood Masoodian, Steve Jones, and Bill Rogers, 111–20. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-27795-8_12.

Duignan, Matthew, James Noble, and Robert Biddle. 2005. 'A Taxonomy of Sequencer User-Interfaces'. In *Proceedings of the 2005 International Computer Music Conference*, 725–28. International Computer Music Association.

———. 2010. 'Abstraction and Activity in Computer-Mediated Music Production'. *Computer Music Journal* 34 (4): 22–33. https://doi.org/10.1162/COMJ/_a/_00023.

Eaglestone, Barry, Nigel Ford, Guy J. Brown, and Adrian Moore. 2007. 'Information Systems and Creativity: An Empirical Study'. *Journal of Documentation* 63 (4): 443–64.

Eaglestone, Barry, Nigel Ford, Peter Holdridge, Jenny Carter, and Catherine Upton. 2008. 'Cognitive Styles and Computer-Based Creativity Support Systems: Two Linked Studies of Electro-Acoustic Music Composers'. In *Computer Music Modeling and Retrieval. Sense of Sounds*, 74–97. Springer. https://doi.org/10.1007/978-3-540-85035-9_5.

Eaglestone, Barry, Nigel Ford, Ralf Nuhn, Adrian Moore, and Guy Brown. 2001. 'Composition Systems Requirements for Creativity: What Research Methodology?' In *Proceedings of MOSART Workshop on Current Research Directions in Computer Music*, 7–16. Audiovisual Institute Pompeu Fabra University.

Eldridge, William, and Benedict Arnold Taylor. 1990. 'Digidesign Sound Tools for Apple Macintosh Computers'. *Computer Music Journal* 14 (4): 77–81.

Elkhaldi, Maher, and Robert Woodbury. 2015. 'Interactive Design Exploration with Alt.text'. *International Journal of Architectural Computing* 13 (2): 103–22. https://doi.org/10.1260/1478-0771.13.2.103.

Etinger, Darko. 2016. 'Tools of the Trade: Digital Audio Workstation Usage Antecedents'. *Informatologia* 49 (1-2): 61–73.

Etinger, Darko, and Tihomir Orehovački. 2017. 'The Usage of UTAUT Model for Digital Audio Workstation User Experience Evaluation'. In *Advances in Usability and User Experience: Proceedings of the AHFE 2017 International Conference on Usability and User Experience*, 559–68. Springer. https://doi.org/10.1007/978-3-319-60492-3_53.

Farbood, Morwaread M, Egon Pasztor, and Kevin Jennings. 2004. 'Hyperscore: A Graphical Sketchpad for Novice Composers'. *IEEE Computer Graphics and Applications* 24 (1): 50–54. https://doi.org/10.1109/MCG.2004.1255809.

Feng, Mi, Cheng Deng, Evan M. Peck, and Lane Harrison. 2017. 'HindSight: Encouraging Exploration Through Direct Encoding of Personal Interaction History'. *IEEE Transactions on Visualization and Computer Graphics* 23 (1): 351–60. https://doi.org/10.1109/TVCG.2016.2599058.

Fitchett, Stephen, Andy Cockburn, and Carl Gutwin. 2013. 'Improving Navigation-Based File Retrieval'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2329–38. CHI '13. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2470654.2481323.

Folkestad, Göran. 2011. 'Digital Tools and Discourse in Music: The Ecology of Composition'. In *Musical Imaginations: Multidisciplinary Perspectives on Creativity, Performance and Perception*, edited by David Hargreaves, Dorothy Miell, and Raymond MacDonald, 193–205. Oxford University Press.

Folkestad, Göran, David J Hargreaves, and Berner Lindström. 1998. 'Compositional Strategies in Computer-Based Music-Making'. *British Journal of Music Education* 15 (1): 83–97.

Frich, Jonas, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. 'Mapping the Landscape of Creativity Support Tools in HCI'. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 389:1–389:18. CHI '19. Glasgow, Scotland UK: ACM. https:

//doi.org/10.1145/3290605.3300619.

Garcia, Jérémie, Theophanis Tsandilas, Carlos Agon, and Wendy Mackay. 2012. 'Interactive Paper Substrates to Support Musical Creation'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1825–8. CHI '12. Austin, Texas, USA: ACM. https://doi.org/10.1145/2207676.2208316.

Garcia, Jérémie, Theophanis Tsandilas, Carlos Agon, and Wendy E. Mackay. 2011. 'InkSplorer: Exploring Musical Ideas on Paper and Computer'. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 361–66. NIME 2011. Oslo, Norway.

———. 2014a. 'PaperComposer: Creating Interactive Paper Interfaces for Music Composition'. In *Proceedings of the 26th Conference on l'Interaction Homme-Machine*, 1–8. IHM '14. Villeneuve d'Ascq, France: ACM. https://doi.org/10.1145/2670444.2670450.

———. 2014b. 'Structured Observation with Polyphony: A Multifaceted Tool for Studying Music Composition'. In *Proceedings of the 2014 Conference on Designing Interactive Systems*, 199–208. DIS '14. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2598510.2598512.

Gelineck, Steven, Morten Büchert, and Jesper Andersen. 2013. 'Towards a More Flexible and Creative Music Mixing Interface'. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, 733–38. CHI EA '13. Paris, France: ACM. https://doi.org/10.1145/2468356.2468487.

Gelineck, Steven, and Stefania Serafin. 2009. 'From Idea to Realization - Understanding the Compositional Processes of Electronic Musicians'. In *Proceedings of the Audio Mostly Conference*. Glasgow.

Gelineck, Steven, and Anders Kirk Uhrenholt. 2016. 'Exploring Visualisation of Channel Activity, Levels and EQ for User Interfaces Implementing the Stage Metaphor for Music Mixing'. In *Proceedings of the 2nd AES Workshop on Intelligent Music Production*. Audio Engineering Society.

Gharib, Islam. 2014. 'Design of Sketch-Based Interface to Enhance Creativity in Conceptual Design'. *International Design Journal* 4 (2): 261–76.

Gibson, Darrell, and Richard Polfreman. 2011. 'An Architecture for Creating Hosting Plug-Ins for Use in Digital Audio Workstations'. In *Proceedings of the 2011 International Computer Music Conference*, 507–10. International Computer Music Association.

Gil, Yolanda, Varun Ratnakar, Timothy Chklovski, Paul Groth, and Denny Vrandecic. 2012. 'Capturing Common Knowledge About Tasks: Intelligent Assistance for to-Do Lists'. *ACM Transactions on Interactive Intelligent Systems* 2 (3): 15:1–15:35. https://doi.org/10.1145/2362394.2362397.

Gohlke, Kristian, Michael Hlatky, Sebastian Heise, David Black, and Jörn Loviscach. 2010. 'Track Displays in DAW Software: Beyond Waveform Views'. In *Proceedings of the 128th AES International Convention*. Audio Engineering Society.

Gonzalez, Alberto. 2015. 'Explorability, Satisficing, and Satisfaction in Parameter Spaces'. PhD thesis, University of North Carolina at Charlotte.

Gonzalez, Victor M., Leonardo Galicia, and Jesus Favela. 2008. 'Understanding and Supporting Personal Activity Management by It Service Workers'. In *Proceedings of the 2nd ACM Symposium on Computer Human Interaction for Management of Information Technology*, 2:1–2:10. CHiMiT '08. New York, NY, USA: ACM. https://doi.org/10.1145/1477973.1477976.

Gottfried, Rama, and Jean Bresson. 2018. 'Symbolist: An Open Authoring Environment for User-Defined Symbolic Notation'. In *Proceedings of the International Conference on Technologies for Music Notation and Representation*, edited by Sandeep Bhagwati and Jean Bresson, 111–18. Montreal, Canada: Concordia University.

Gross, Mark D., Ellen Yi-Luen Do, Raymond J. McCall, Wayne V. Citrin, Paul Hamill, Adrienne Warmack, and Kyle S. Kuczun. 1998. 'Collaboration and Coordination in Architectural Design: Approaches to Computer Mediated Team Work'. *Automation in Construction* 7 (6): 465–73. https://doi.org/10.1016/S0926-5805(98)00055-7.

Gutwin, Carl, and Andy Cockburn. 2006. 'Improving List Revisitation with Listmaps'. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, 396–403. AVI '06. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1133265.1133347.

Hailpern, Joshua, Erik Hinterbichler, Caryn Leppert, Damon Cook, and Brian P.

Bailey. 2007. 'TEAM STORM: Demonstrating an Interaction Model for Working with Multiple Ideas During Creative Group Work'. In *Proceedings of the 6th ACM SIGCHI Conference on Creativity & Cognition*, 193–202. C&C '07. Washington, DC, USA: ACM. https://doi.org/10.1145/1254960.1254987.

Haraty, Mona, and Joanna McGrenere. 2016. 'Designing for Advanced Personalization in Personal Task Management'. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 239–50. DIS '16. Brisbane, QLD, Australia: ACM. https://doi.org/10.1145/2901790.2901805.

Haraty, Mona, Joanna McGrenere, and Charlotte Tang. 2016. 'How Personal Task Management Differs Across Individuals'. *International Journal of Human-Computer Studies* 88: 13–37. https://doi.org/10.1016/j.ijhcs.2015.11.006.

Haraty, Mona, Diane Tam, Shathel Haddad, Joanna McGrenere, and Charlotte Tang. 2012. 'Individual Differences in Personal Task Management: A Field Study in an Academic Setting'. In *Proceedings of Graphics Interface 2012*, 35–44. GI '12. Toronto, Ont., Canada, Canada: Canadian Information Processing Society.

Hartmann, Björn, Loren Yu, Abel Allison, Yeonsoo Yang, and Scott R. Klemmer. 2008. 'Design as Exploration: Creating Interface Alternatives Through Parallel Authoring and Runtime Tuning'. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology*, 91–100. UIST '08. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1449715.1449732.

Healey, Patrick GT, and Jean-Baptiste Thiebaut. 2007. 'Sketching Musical Compositions'. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 1079–84. The University of California, Merced. https://escholarship.org/uc/item/3kc0t66r.

Herrmann, Andreas, Daniel G. Goldstein, Rupert Stadler, Jan R. Landwehr, Mark Heitmann, Reto Hofstetter, and Frank Huber. 2011. 'The Effect of Default Options on Choice — Evidence from Online Product Configurators'. *Journal of Retailing and Consumer Services* 18 (6): 483–91. https://doi.org/10.1016/j.jretconser.2011.06.005.

Hewett, Tom, Mary Czerwinski, Michael Terry, Jay Nunamaker, Linda Candy, Bill Kules, and Elisabeth Sylvan. 2005. 'Creativity Support Tool Evaluation Methods and Metrics'. In *NSF Workshop Report on Creativity Support Tools*, 10–24.

Hoare, Michaela, Steve Benford, Rachel Jones, and Natasa Milic-Frayling. 2014.

'Coming in from the Margins: Amateur Musicians in the Online Age'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1295–1304. CHI '14. Toronto, Ontario, Canada: ACM. https://doi.org/10.1145/2556288.2557298.

Hollan, James, Edwin Hutchins, and David Kirsh. 2000. 'Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research'. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7 (2): 174–96. https://doi.org/10.1145/353485.353487.

Hollnagel, Erik, and David D. Woods. 2005. *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. CRC Press.

Hracs, Brian J. 2012. 'A Creative Industry in Transition: The Rise of Digitally Driven Independent Music Production'. *Growth and Change* 43 (3): 442–61. https://doi.org/10.1111/j.1468-2257.2012.00593.x.

Inie, Nanna, Allison Endo, Steven Dow, and Peter Dalsgaard. 2018. 'The Problem Solver and the Artisan Designer: Strategies for Utilizing Design Idea Archives'. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*, 397–406. NordiCHI '18. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3240167.3240215.

Jennings, Kevin. 2006. 'Computer Graphical Interfaces, Reflection and Music Composition–A Holistic Study.' PhD thesis, University of Dublin.

Jillings, Nicholas, and Ryan Stables. 2017. 'Investigating Music Production Using a Semantically Powered Digital Audio Workstation in the Browser'. In *2017 AES International Conference on Semantic Audio*. Audio Engineering Society.

Jordan, Brigitte, and Austin Henderson. 1995. 'Interaction Analysis: Foundations and Practice'. *The Journal of the Learning Sciences* 4 (1): 39–103.

Kaloterakis, Stefanos. 2013. 'Creativity and Home Studios: An in-Depth Study of Recording Artists in Greece'. *Journal on the Art of Record Production* 8.

Kamsin, Amirrudin. 2014. 'Improving Tool Support for Personal Task Management (PTM)'. PhD thesis, University College London.

Kamsin, Amirrudin, Ann Blandford, and Anna L. Cox. 2012. 'Personal Task Management: My Tools Fall Apart When I'm Very Busy!' In *CHI '12 Extended Abstracts on*

Human Factors in Computing Systems, 1369–74. CHI EA '12. New York, NY, USA: ACM. https://doi.org/10.1145/2212776.2212457.

Keller, Ianus, Froukje Sleeswijk Visser, Remko van der Lugt, and Pieter Jan Stappers. 2009. 'Collecting with Cabinet: Or How Designers Organise Visual Material, Researched Through an Experiential Prototype'. *Design Studies* 30 (1): 69–86. https://doi.org/10.1016/j.destud.2008.06.001.

Kim, Joy, Avi Bagla, and Michael S. Bernstein. 2015. 'Designing Creativity Support Tools for Failure'. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity & Cognition*, 157–60. C&C '15. Glasgow, United Kingdom: ACM. https://doi.org/10.1145/2757226.2764542.

Koda, Alan Hugh. 2011. 'Cutting the Cables: Developing DAWs Beyond Analog Methods'. PhD thesis, Savannah College of Art and Design, Georgia.

Kokkalis, Nicolas, Thomas Köhn, Johannes Huebner, Moontae Lee, Florian Schulze, and Scott R. Klemmer. 2013. 'TaskGenies: Automatically Providing Action Plans Helps People Complete Tasks'. *ACM Transactions on Computer-Human Interaction (TOCHI)* 20 (5): 27:1–27:25. https://doi.org/10.1145/2513560.

Koszolko, Martin K. 2015. 'Crowdsourcing, Jamming and Remixing: A Qualitative Study of Contemporary Music Production Practices in the Cloud'. *Journal on the Art of Record Production* 10.

Lee, Brian, and Scott R Klemmer. 2005. 'Notebooks That Share and Walls That Remember: Electronic Capture of Design Education Artifacts'. In *Adjunct Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology*. UIST '05. Seattle, WA: Association for Computing Machinery.

Leshed, Gilly, and Phoebe Sengers. 2011. '"I Lie to Myself That I Have Freedom in My Own Schedule": Productivity Tools and Experiences of Busyness'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 905–14. CHI '11. New York, NY, USA: ACM. https://doi.org/10.1145/1978942.1979077.

Letondal, Catherine, and Wendy E. Mackay. 2007. 'The Paperoles Project: An Analysis of Paper Use by Music Composers'. In *Proceedings of the 2nd International Workshop on Collaborating over Paper and Digital Documents*. CoPADD'07.

Leyshon, Andrew. 2001. 'Time–Space (and Digital) Compression: Software Formats, Musical Networks, and the Reorganisation of the Music Industry'. *Environment and Planning A: Economy and Space* 33 (1): 49–77. https://doi.org/10.1068/a3360.

Macchiusi, Ian Andrew. 2017. '"Knowing Is Seeing": The Digital Audio Workstation and the Visualization of Sound'. PhD thesis, York University Toronto.

MacCormick, Daniel, and Loutfouz Zaman. 2019. 'SuBViS: The Use of Subjunctive Visual Programming Environments for Exploring Alternatives in Game Development'. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. FDG '19. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3337722.3337740.

Magnusson, Thor. 2010. 'Designing Constraints: Composing and Performing with Digital Musical Systems'. *Computer Music Journal* 34 (4): 62–73. https://doi.org/10.1162/COMJ/_a/_00026.

Mangano, Nicolas, Thomas D. LaToza, Marian Petre, and André van der Hoek. 2014. 'Supporting Informal Design with Interactive Whiteboards'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 331–40. CHI '14. Toronto, Ontario, Canada: ACM. https://doi.org/10.1145/2556288.2557411.

Marrington, Mark. 2010. 'Experiencing Musical Composition in the DAW: The Software Interface as Mediator of the Musical Idea'. In *Proceedings of the 6th Art of Record Production Conference*. Association for the Study of the Art of Record Production.

———. 2016. 'Paradigms of Music Software Interface Design and Musical Creativity'. In *KES Innovation in Music II*, edited by R Hepworth-Sawyer, J Hodgson, R Toulson, and J L Paterson, 73–93. Future Technology Press, UK.

McCulloch, Peter. 2014. 'Thema: A Software Framework for the Quantitative Study of Compositional Process'. PhD thesis, New York University.

McGarry, Glenn, Peter Tolmie, Steve Benford, Chris Greenhalgh, and Alan Chamberlain. 2017. '"They're All Going Out to Something Weird": Workflow, Legacy and Metadata in the Music Production Process'. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 995–1008. CSCW '17. Portland, Oregon, USA: ACM. https://doi.org/10.1145/2998181.2998325.

McGrath, Sean, Alan Chamberlain, and Steve Benford. 2016. 'Making Music Together: An Exploration of Amateur and Pro-Am Grime Music Production'. In *Proceedings of the Audio Mostly Conference 2016*, 186–93. AM '16. Norrköping, Sweden: ACM. https://doi.org/10.1145/2986416.2986432.

Méndez, Francisco José Cuadrado. 2015. 'The Use of Sequencer Tools During the Composition Process: A Field Study'. *Journal of Music, Technology & Education* 8 (1): 55–70.

Merrill, Bryce. 2010. 'Music to Remember Me by: Technologies of Memory in Home Recording'. *Symbolic Interaction* 33 (3): 456–74. https://doi.org/10.1525/si.2010.33.3.456.

Miyazaki, R., I. Fujishiro, and R. Hiraga. 2004. 'Comp-I: A System for Visual Exploration and Editing of MIDI Datasets'. In *Proceedings of the 2004 International Computer Music Conference*. International Computer Music Association.

Morey, Justin. 2013. 'Copyright Management and Its Effect on the Sampling Practice of UK Dance Music Producers'. *Journal of the International Association for the Study of Popular Music* 3 (1): 48–62.

———. 2017. 'A Study of Sampling Practice in UK Dance Music, 1987-2012'. PhD thesis, Leeds Beckett University.

Müller, Stefanie, Gregor Miller, and Sidney Fels. 2010. 'Using Temporal Video Annotation as a Navigational Aid for Video Browsing'. In *Adjunct Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, 445–46. UIST '10. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1866218.1866263.

Mycroft, Josh. 2018. 'The Design of Audio Mixing Software Displays to Support Critical Listening'. PhD thesis, Queen Mary, University of London.

Mycroft, Josh, Joshua Reiss, and Tony Stockman. 2016. 'Visually Representing and Interpreting Multivariate Data for Audio Mixing'. In *Proceedings of the Sound and Music Computing Conference*, 332–37.

Myllys, Petri. 2014. 'User Interface Paradigms in Digital Audio Workstations: Examining and Modernising Established Models'. Master's thesis, University of the Arts

Helsinki.

Nabavian, Shahin. 2010. 'Distributed Cognition in Joint Music Composition: Exploring the Role of Language and Artefacts in Multi-Session Creative Collaborative Work'. PhD thesis, Queen Mary, University of London.

Nakamura, Hiromi, Tomoyasu Nakano, Satoru Fukayama, and Masataka Goto. 2018. 'ChordScanner: Browsing Chord Progressions Based on Musical Typicality and Intra-Composer Consistency'. In *Proceedings of the 2018 International Computer Music Conference*, 250–55. International Computer Music Association.

Nancel, Mathieu, and Andy Cockburn. 2014. 'Causality: A Conceptual Model of Interaction History'. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 1777–86. CHI '14. Toronto, Ontario, Canada: ACM. https://doi.org/10.1145/2556288.2556990.

Nardi, Bonnie A. 1996. *Context and Consciousness: Activity Theory and Human-Computer Interaction*. The MIT Press.

Nash, Chris. 2012. 'Supporting Virtuosity and Flow in Computer Music'. PhD thesis, University of Cambridge.

———. 2015. 'The Cognitive Dimensions of Music Notations'. In *Proceedings of the First International Conference on Technologies for Music Notation and Representation*, edited by Marc Battier, Jean Bresson, Pierre Couprie, Cécile Davy-Rigaux, Dominique Fober, Yann Geslin, Hugues Genevois, François Picard, and Alice Tacaille, 190–202. Paris, France.

Ngoon, Tricia J., Caren M. Walker, and Scott Klemmer. 2019. 'The Dark Side of Satisficing: Setting the Temperature of Creative Thinking'. In *Proceedings of the 2019 ACM SIGCHI Conference on Creativity and Cognition*, 591–96. C&C '19. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3325480.3326581.

Nicoll, Benjamin, and Brendan Keogh. 2019. 'Grain: Default Settings, Design Principles, and the Aura of Videogame Production'. In *The Unity Game Engine and the Circuits of Cultural Software*, 63–80. Springer.

Norman, Donald A. 1993. *Things That Make Us Smart: Defending Human Attributes* *in the Age of the Machine*. Basic Books.

Paterson, Justin. 2011. 'The Preset Is Dead; Long Live the Preset'. In *Proceedings of the 130th AES International Convention*. Audio Engineering Society.

Pavese, Errico. 2016. 'Discussing Three Models for the Visual Representation of a Recorded Song'. *El Oído Pensante* 4 (1): 178–90.

Pearce, Marcus, and Geraint A. Wiggins. 2002. 'Aspects of a Cognitive Theory of Creativity in Musical Composition'. In *Proceedings of the European Conference on Artificial Intelligence 2002 Workshop on Creative Systems*, 17–24.

Pennycook, Bruce W. 1985. 'Computer-Music Interfaces: A Survey'. *ACM Computing Surveys* 17 (2): 267–89. https://doi.org/10.1145/4468.4470.

Peterson, John. 2008. 'Computer Notation-Based Music Composition and the Delayed Introduction of Musical Expression Markings'. In *Proceedings of the 6th International Conference on Education and Information Systems, Technologies and Applications*. EISTA 2008. International Institute of Informatics and Systemics.

Peterson, John, and Emery Schubert. 2007. 'Music Notation Software: Some Observations on Its Effects on Composer Creativity'. *Proceedings of the Inaugural International Conference on Music Communication Science*, 127–30.

Phillips, William J. 2010. 'Making Tracks: Digital Recording Technology and the Democratization of Cultural Production'. PhD thesis, New York University.

Polfreman, Richard. 1997. 'User-Interface Design for Software Based Sound Synthesis Systems.' PhD thesis, University of Hertfordshire.

———. 2001. 'Supporting Creative Composition: The Frameworks Approach'. In *Les Actes Des 8eme Journées d'Informatique Musicale*, 99–112.

Pras, Amandine, Catherine Guastavino, and Maryse Lavoie. 2013. 'The Impact of Technological Advances on Recording Studio Practices'. *Journal of the American Society for Information Science and Technology* 64 (3): 612–26. https://doi.org/10.1002/asi.22840.

Puckette, Miller. 2002. 'Max at Seventeen'. *Computer Music Journal* 26 (4): 31–43. https://doi.org/10.1162/014892602320991356.

Resnick, Mitchel, Brad Myers, Kumiyo Nakakoji, Ben Shneiderman, Randy Pausch, Ted Selker, and Mike Eisenberg. 2005. 'Design Principles for Tools to Support Creative Thinking'. In *NSF Workshop Report on Creativity Support Tools*, 25–36.

Roberts, Jonathan C. 2007. 'State of the Art: Coordinated & Multiple Views in Exploratory Visualization'. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization*, 61–71. CMV 2007. IEEE. https://doi.org/10.1109/CMV.2007.20.

Roels, Hans. 2016. 'Comparing the Main Compositional Activities in a Study of Eight Composers'. *Musicæ Scientiæ* 20 (3): 413–35. https://doi.org/10.1177/1029864915624737.

Rogers, Yvonne, Helen Sharp, and Jenny Preece. 2011. *Interaction Design: Beyond Human-Computer Interaction*. John Wiley & Sons.

Rule, Adam, Aurélien Tabard, and Jim Hollan. 2017. 'Using Visual Histories to Reconstruct the Mental Context of Suspended Activities'. *Human–Computer Interaction* 32 (5-6): 511–58. https://doi.org/10.1080/07370024.2017.1300063.

Rutz, Hanns Holger. 2014. 'Tracing the Compositional Process. Sound Art That Rewrites Its Own Past: Formation, Praxis and a Computer Framework'. PhD thesis, University of Plymouth.

Salvucci, Dario D. 2010. 'On Reconstruction of Task Context After Interruption'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 89–92. CHI '10. Atlanta, Georgia, USA: ACM. https://doi.org/10.1145/1753326.1753341.

Satyanarayan, Arvind, and Jeffrey Heer. 2014. 'Lyra: An Interactive Visualization Design Environment'. *Computer Graphics Forum* 33 (3): 351–60. https://doi.org/10.1111/cgf.12391.

Satyanarayan, Arvind, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. 'Vega-Lite: A Grammar of Interactive Graphics'. *IEEE Transactions on Visualization and Computer Graphics* 23 (1): 341–50. https://doi.org/10.1109/TVCG.2016.2599030.

Scaife, Mike, and Yvonne Rogers. 1996. 'External Cognition: How Do Graphical Representations Work?' *International Journal of Human-Computer Studies* 45 (2):

185–213. https://doi.org/10.1006/ijhc.1996.0048.

Seddon, Frederick A, and Susan A O'Neill. 2003. 'Creative Thinking Processes in Adolescent Computer-Based Composition: An Analysis of Strategies Adopted and the Influence of Instrumental Music Training'. *Music Education Research* 5 (2): 125–37. https://doi.org/10.1080/1461380032000085513.

Shah, Rajiv C., and Jay P. Kesan. 2006. 'Policy Through Software Defaults'. In *Proceedings of the 2006 International Conference on Digital Government Research*, 265–72. Dg.o '06. San Diego, California, USA: Digital Government Society of North America. https://doi.org/10.1145/1146598.1146670.

Shneiderman, B. 2000. 'Supporting Creativity with Powerful Composition Tools for Artifacts and Performances'. In *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*. https://doi.org/10.1109/HICSS.2000.926896.

Skopik, Amy, and Carl Gutwin. 2005. 'Improving Revisitation in Fisheye Views with Visit Wear'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 771–80. CHI '05. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1054972.1055079.

Smith, Jeffrey A. 2011. 'A Cognition-Analogous Approach to Early-Stage Creative Ideation Support in Music Composition Software'. PhD thesis, University of Saskatchewan.

Stowell, Dan, and Alex McLean. 2013. 'Live Music-Making: A Rich Open Task Requires a Rich Open Interface'. In *Music and Human-Computer Interaction*, edited by Simon Holland, Katie Wilkie, Paul Mulholland, and Allan Seago, 139–52. London: Springer. https://doi.org/10.1007/978-1-4471-2990-5_8.

Suchman, Lucy, and Randall Trigg. 1995. 'Understanding Practice: Video as a Medium for Reflection and Design (Excerpt)'. In *Human-Computer Interaction: Toward the Year 2000*, edited by Ronald M. Baecker, Jonathan Grudin, William A. S. Buxton, and Saul Greenberg, 233–40. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Tanev, George, and Adrijan Božinovski. 2014. 'Virtual Studio Technology Inside Music Production'. In *ICT Innovations 2013*, edited by Vladimir Trajkovik and Misev Anastas, 231–41. Heidelberg: Springer International Publishing. https://doi.org/10.

1007/978-3-319-01466-1_22.

Terry, Michael, and Elizabeth D. Mynatt. 2002. 'Recognizing Creative Needs in User Interface Design'. In *Proceedings of the 4th Conference on Creativity & Cognition*, 38–44. C&C '02. Loughborough, UK: ACM. https://doi.org/10.1145/581710.581718.

Terry, Michael, Elizabeth D. Mynatt, Kumiyo Nakakoji, and Yasuhiro Yamamoto. 2004. 'Variation in Element and Action: Supporting Simultaneous Development of Alternative Solutions'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 711–18. CHI '04. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/985692.985782.

Tharatipyakul, Atima, Hyowon Lee, Shengdong Zhao, and Richard C. Davis. 2016. 'Supporting the Comparison of Alternative Stories'. In *Proceedings of the 28th Australian Conference on Computer-Human Interaction*, 266–70. OzCHI '16. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3010915.3010963.

Théberge, Paul. 2004. 'The Network Studio: Historical and Technological Paths to a New Ideal in Music Making'. *Social Studies of Science* 34 (5): 759–81. https://doi.org/10.1177/0306312704047173.

Thiebaut, Jean-Baptiste. 2010. 'Sketching Music: Representation and Composition'. PhD thesis, Queen Mary, University of London.

Thiebaut, Jean-Baptiste, Patrick GT Healey, and Nick Bryan Kinns. 2008. 'Drawing Electroacoustic Music'. In *Proceedings of the 2008 International Computer Music Conference*. International Computer Music Association.

Tobias, Evan S. 2013. 'Composing, Songwriting, and Producing: Informing Popular Music Pedagogy'. *Research Studies in Music Education* 35 (2): 213–37. https://doi.org/10.1177/1321103X13487466.

Tsandilas, Theophanis, Catherine Letondal, and Wendy E. Mackay. 2009. 'Musink: Composing Music Through Augmented Drawing'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 819–28. CHI '09. Boston, MA, USA: ACM. https://doi.org/10.1145/1518701.1518827.

Tubb, Robert, and Simon Dixon. 2014a. 'A Four Strategy Model of Creative Parameter Space Interaction'. In *Proceedings of the Fifth International Conference on Computational Creativity*, edited by Simon Colton, Dan Ventura, Nada Lavrač, and Michael Cook, 16–22. Ljubljana, Slovenia.

———. 2014b. 'The Divergent Interface: Supporting Creative Exploration of Parameter Spaces'. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 227–32. NIME 2014. Goldsmiths, University of London.

Visser, Willemien. 2006. 'Designing as Construction of Representations: A Dynamic Viewpoint in Cognitive Design Research'. *Human–Computer Interaction* 21 (1): 103–52.

Walther, Joachim, Brett Robertson, and David Radcliffe. 2007. 'Avoiding the Potential Negative Influence of CAD Tools on the Formation of Students' Creativity'. In *Proceedings of the 18th Conference of the Australasian Association for Engineering Education*. University of Melbourne.

Walther-Hansen, Mads. 2017. 'New and Old User Interface Metaphors in Music Production'. *Journal on the Art of Record Production* 11.

Wang, Kai, and Jeffrey V. Nickerson. 2017. 'A Literature Review on Individual Creativity Support Systems'. *Computers in Human Behavior* 74: 139–51. https://doi.org/10.1016/j.chb.2017.04.035.

Wang Baldonado, Michelle Q., Allison Woodruff, and Allan Kuchinsky. 2000. 'Guidelines for Using Multiple Views in Information Visualization'. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, 110–19. AVI '00. Palermo, Italy: ACM. https://doi.org/10.1145/345513.345271.

Wexelblat, Alan, and Pattie Maes. 1999. 'Footprints: History-Rich Tools for Information Foraging'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 270–77. CHI '99. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/302979.303060.

Wilkie, Katie. 2014. 'Conceptual Metaphor, Human-Computer Interaction and Music: Applying Conceptual Metaphor to the Design and Analysis of Music Interactions'. PhD thesis, The Open University.

Willett, Wesley, Jeffrey Heer, and Maneesh Agrawala. 2007. 'Scented Widgets:

Improving Navigation Cues with Embedded Visualizations'. *IEEE Transactions on Visualization and Computer Graphics* 13 (6): 1129–36. https://doi.org/10.1109/TVCG.2007.70589.

Wilson, Margaret. 2002. 'Six Views of Embodied Cognition'. *Psychonomic Bulletin & Review* 9 (4): 625–36. https://doi.org/10.3758/BF03196322.

Wongsuphasawat, Kanit, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2016. 'Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations'. *IEEE Transactions on Visualization and Computer Graphics* 22 (1): 649–58. https://doi.org/10.1109/TVCG.2015.2467191.

Wu, Dezhi. 2005. 'Supporting Individual Time Management Through the Capture and Display of Temporal Structures'. PhD thesis, New Jersey Institute of Technology.

Yavelow, Christopher. 1989. 'Musicians and Microprocessors: MIDI and the State of the Art'. In *The Music Machine: Selected Readings from Computer Music Journal*, edited by Curtis Roads, 199–235. The MIT Press.

Yli-Kauhaluoma, Sari. 2009. 'Time at R&D Work: Types and Strategies of Time in the Collaborative Development of a Chemical Technology'. *Time & Society* 18 (1): 130–53. https://doi.org/10.1177/0961463X08099948.

Zadel, Mark, and Gary Scavone. 2006. 'Laptop Performance: Techniques, Tools, and a New Interface Design'. In *Proceedings of the International Computer Music Conference*, 643–48.

Zagal, José P., and Roger Altizer. 2015. 'Placeholder Content in Game Development: Benefits and Challenges'. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*, 745–50. CHI Play '15. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2793107.2810319.

Zaman, Loutfouz, Wolfgang Stuerzlinger, Christian Neugebauer, Rob Woodbury, Maher Elkhaldi, Naghmi Shireen, and Michael Terry. 2015. 'GEM-Ni: A System for Creating and Managing Alternatives in Generative Design'. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 1201–10. CHI '15. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2702123.2702398.

Appendix A

Ableton Live

This appendix introduces the basic functionality provided by *Ableton Live* and the associated terminology used to describe it.

Compositions in *Live* are known as *Live Sets*. A Live Set consists of several *tracks*, which are each attached to a *chain* of instruments and effects, and contain either audio or MIDI *clips*.

Live operates in two main modes, *session* mode and *arrangement* mode.

In **arrangement mode**, *Live* is set up like a conventional timeline-based DAW (see Figure A.1). Tracks are represented by rows, with time progressing left-to-right. A particular time interval can be played repeatedly by indicating it using the *Arrangement Loop* tool above the track display.

In **session mode**, *Live* is set up for live performance or jamming (see Figure A.2). Each track is displayed in the main panel as a column, with each row representing a *scene* (a section of the composition, such as an intro, verse, or chorus). The composer can trigger each clip individually, or cause an entire scene to play, by clicking on the corresponding play buttons.

The order of the rows in Session mode does not always represent the order in which scenes are played back. By default, scenes are played only when the composer triggers them manually. The composer may optionally set up rules to describe when scenes are played, for how many repetitions each scene is played, and even include rules with elements of random chance.

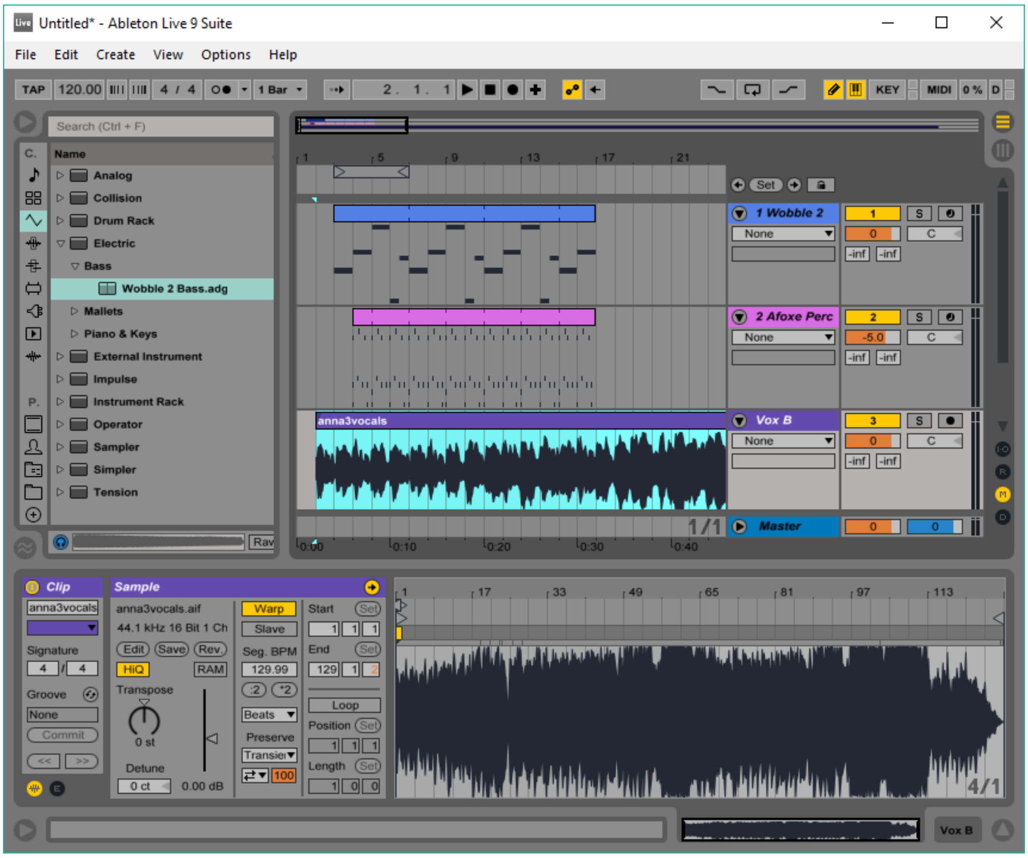

Figure A.1: Live's *arrangement* mode

Figure A.2: Live's *session* mode

In both modes, additional panels provide some shared functionality:

- a panel at the left of the screen contains a *browser* for *Live's* library of reusable objects and presets
- a panel at the bottom of the screen can be switched between either displaying the currently selected clip or the currently selected track's chain

Appendix B

LiveTodos technical architecture

The architecture of the LiveTodos system is illustrated in Figure B.1.

Ableton Live provides a Python interpreter that runs scripts which can access an API to control *Live* and read information about *Live's* state. These scripts are called *MIDI remote scripts*, as they were intended to allow manufacturers of hardware MIDI controllers to create scripts that would enable more complex control over *Live* from their devices, though they can be used for other purposes.

A **custom MIDI remote script** running in *Live's* Python interpreter reads in information about *Live's* current state and attaches event handlers that are triggered when changes occur to objects of interest. The script regularly sends this information to the LiveTodos web server.

The **web server** is implemented as a Node.JS web application using the Express framework. Users are authenticated using Google's OAuth2 support (omitted from the diagram for clarity). Information about users, their Live Sets, and their to-do lists are stored in a SQLite database. The server provides a REST web service API for accessing and updating this information.

The **web page** is a web application implemented in JavaScript using jQuery. It provides a UI to display and edit the information that is provided through the REST API.

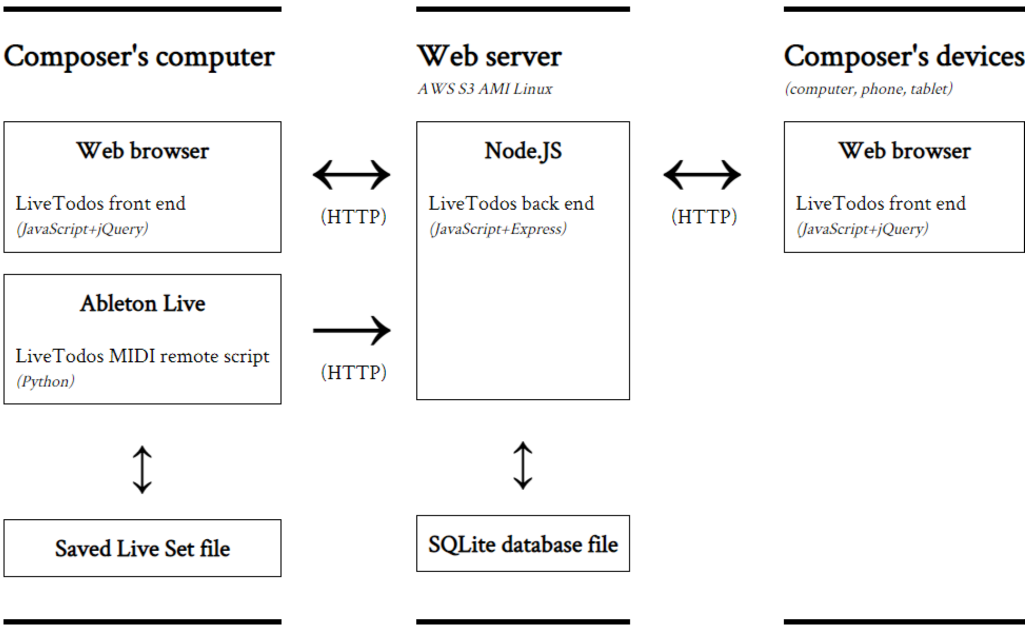

Figure B.1: Architecture diagram for LiveTodos

Appendix C

Written materials for studies

The following pages contain additional written material related to the studies studies described in chapters 3 and 5.

These materials were used during the ethical approval process, recruitment process, and the studies themselves.

Information and Consent Form

This study is to investigate the working processes of people composing music using music software. It is designed to discover how you think about composing, what situation you compose in, and what methods you use to do it.

How the study will work

You will be visited a small number of times for one hour sessions in the location you normally compose in, and observed engaging in your normal composition work. Data from these sessions will be recorded using:

- screen capture video of the computer you are using
- photographs of the activities you do, the tools you use, and any sketches, lists, diagrams, or drawings you make
- audio recording of what you say and what you are listening to during the study, which may include your musical work in progress
- written notes made by a researcher

As part of this study you may also be asked questions about your musical techniques and tools, the requirements for your composition, and the context you are composing in. Your answers to these questions may be recorded by the researcher.

If you agree to it later in this document, you may also choose to be recorded with a video camera or provide data files from your work.

If you are no longer able to participate in some or all of the study, please email me at psxjc7@nottingham.ac.uk to rearrange or cancel.

Your data and withdrawing

Data collected will be held in a secure and safe manner in accordance with the Data Protection Act 1998.

In addition to the data listed above, *personal data* will be collected for the purposes of administering the study: your name, a contact telephone number and/or email address, and any addresses we have agreed to meet at. You are free to withdraw at any time and have your personal data erased from our records. To withdraw from the study, email the researcher at psxjc7@nottingham.ac.uk.

Any data relating to you will be deleted if you withdraw before the end of the study.

Consent

I understand that the data gathered will be used in the study and associated research. I understand that excerpts from the data gathered in the study and conclusions based on that data may be published in an anonymised form (for example, in a scientific journal, conference or thesis). The data may also be shared or discussed with other researchers as part of the research. Electronic copies of these publications may also be made available to the public online.

Please write your initials next to these statements to indicate if you agree:

_______ I confirm that I am over 18.

_______ I understand that I am free to withdraw from the study at any time, without giving a reason, and will not be penalised for it.

_______ I have asked any questions I have about the study and been given any answers I need.

_______ I understand the purpose of this study and what it will involve.

_______ I voluntarily agree to take part in the study.

If you are happy to do so, it would be helpful to get permission from you to do the following, but they are not necessary for the study. Initial these statements if you agree to them.

_______ Additionally, I also agree to allow the researcher to video my working practices.

_______ Additionally, I also agree to provide the researcher with copies of files containing my composition or related electronic notes

Signed _____________________________

Date _____________________________

Name

Address

Contact Tel No

Email

Studying the music composition process of musicians using music software

This study is to investigate the working processes of people composing music using music software. It is an exploratory study to direct the future direction of my research. I intend to do this by going to the location people normally conduct their composition work in (which may be their home, workplace, or a higher education institution) and observing their normal composition activities with their usual tools, over a small number of one hour sessions.

Participants will be recruited from groups that adult (18+) musicians might participate in, which may include (for example) courses at higher education institutions, university societies, and user groups for music software.

I will contact whoever at these groups is responsible for distributing information to members and ask them whether it would be appropriate to advertise to the group. If permission is given, I will ask for advice on what medium is normally used to contact members (mailing list, notice board, etc.), and then send out a short summary of the purpose of the study and what it will involve, along with instructions on how to contact me.

To ensure that participants are able to meaningfully consent, they will be verbally asked questions to ensure that they do understand what will happen during the study and the procedures for withdrawing from it.

Participants will not be compensated for their time as they will be observed carrying out their normal activity.

I am looking to specifically gain information about:

- what activities are involved in their compositional process
- when they make use of secondary artefacts (e.g. paper notes or diagrams) when working with computer music making tools, and what they use them for
- how they describe the task they are engaged in, and how they plan future actions and consider different alternative actions

I will be gathering some or all of:

- personally identifying data for the purpose of running the study, which will be destroyed after the study is complete: personal name, a contact telephone number and/or email address, and the address of a place we have arranged to meet for the study (which may be their home address)
- screen capture video of any computers they are using
- video and photographs of the musicians and their composition method, including their use of physical controllers
- photographs of any sketches, lists, diagrams, or drawings they make on paper or on other electronic devices
- copies of any data files they produce during the session, if permission is given, including those representing the music itself they are composing
- an audio recording of the session. This will include parts of the music they were listening to out loud during the session itself, and also any comments they make verbally or answers to questions I ask them for about the context of their musical

work (e.g. requirements of the project they are making music for, or what their tools are)

- a written record of my observations during the course of the session

Information sheet

Date 2018-08-28 Version v.1.0.0

This is a study of users of Ableton Live, in which we are asking you to use a new "todo list" website that is integrated with Ableton Live.

From this study, we hope to learn:

- about strengths and weaknesses of, and possible improvements to, of the todo list website
- how the todo list website fits into the process of composition, including what kinds of todo list entries are created, when they are created, and why
- when and why other forms of notes are used instead
- how composers co-ordinate use of a second tool together with their use of Ableton Live

You have been asked to participate because you are an experienced user of Ableton Live.

There are three parts to this research, and you may participate in one or more of them:

1. You may be asked to use the software for 1-2 hours at the university. This task will involve composing music using materials provided to you, while you are filmed.
2. You may use the software in other situations, e.g. at home. Data on how you use the software will be automatically logged while you do this.
3. You may be asked to give feedback on your use of the software when you have finished using it (this should take less than half an hour, and will be recorded.)

In exchange for your time, you will be compensated £20 in Amazon vouchers.

You will need a Google account to participate in this study.

Please be aware that the todo list website is **experimental** and while we will try to avoid this, there is a small possibility that todo list items you create on the website could be accidentally lost. If you experience any problems using the software, please email psxjc7@nottingham.ac.uk

Personal data collected will be held in a secure and safe manner in accordance with the Data Protection Act 1998. However no computer system is perfectly secure and it is always possible that a third party might gain unauthorised access to the collected data. Participants are free to withdraw at any time and their personal data will be excluded from further use. If you need or wish to, you may withdraw from the study at any time by emailing psxjc7@nottingham.ac.uk

The information that is recorded in this study will be anonymized before it is stored by either removing personally identifying information, or by replacing it with an ID number. This anonymous data, or information produced from analyzing it, may be published in a PhD thesis, journal article, or on the Nottingham Research Data Management Repository website.

This study is being carried out by Anna Clarke, PhD student in the School of Computer Science at the University of Nottingham. This study has passed ethical review and been approved by the University of Nottingham Computer Science Research Ethics Committee. It is funded by RCUK through the grants "Horizon Centre for Doctoral Training at the University of Nottingham" (RCUK Grant No. EP/G037574/1) and "Horizon Digital Economy Research Institute" (RCUK Grant No. EP/G065802/1).

If you have any further questions about this research, please contact Anna Clarke at psxjc7@nottingham.ac.uk

PRIVACY NOTICE

The University of Nottingham is committed to protecting your personal data and informing you of your rights in relation to that data. The University will process your personal data in accordance with the General Data Protection Regulation (GDPR) and the Data Protection Act 2018 and this privacy notice is issued in accordance with GDPR Articles 13 and 14.

The University of Nottingham, University Park, Nottingham, NG7 2RD is registered as a Data Controller under the Data Protection Act 1998 (registration No. Z5654762, https://ico.org.uk/ESDWebPages/Entry/Z5654762).

The University has appointed a Data Protection Officer (DPO). The DPO's postal address is:

Data Protection Officer,
Legal Services
A5, Trent Building,
University of Nottingham,
University Park,
Nottingham
NG7 2RD

The DPO can be emailed at dpo@nottingham.ac.uk

Why we collect your personal data. We collect personal data under the terms of the University's Royal Charter in our capacity as a teaching and research body to advance education and learning. Specific purposes for data collection on this occasion are to gain information about the working practices of composers creating music and to evaluate a prototype composition tool.

The legal basis for processing your personal data under GDPR. Under the General Data Protection Regulation, the University must establish a legal basis for processing your personal data and communicate this to you. The legal basis for processing your personal data on this occasion is Article 6(1e) processing is necessary for the performance of a task carried out in the public interest.

Where the University receives your personal data from. During this study personal data may be gathered in the form of:

- Your name, email address, or telephone number, for the purpose of administering the study
- Audio and video recordings of your voice and image

- Automatically recorded information through your use of the software, including IP addresses that you access it through, and information about your identity provided by Google that is used to authenticate you

How long we keep your data. The University may store your data for up to 25 years and for a period of no less than 7 years after the research project finishes. The researchers who gathered or processed the data may also store the data indefinitely and reuse it in future research.

Who we share your data with. Your data may be shared with researchers from other collaborating institutions and organisations who are involved in the research. Extracts of your data may be disclosed in published works that are posted online for use by the scientific community. Your data may also be stored indefinitely by members of the researcher team and/or be stored on external data repositories (e.g., the UK Data Archive) and be further processed for archiving purposes in the public interest, or for historical, scientific or statistical purposes.

How we keep your data safe. We keep your data securely and put measures in place to safeguard it. These safeguards include anonymization of data and storing the data on secure servers.

Your rights as a data subject. GDPR provides you, as a data subject, with a number of rights in relation to your personal data. Subject to some exemptions, you have the right to:

- withdraw your consent at any time where that is the legal basis of our processing, and in such circumstances you are not obliged to provide personal data for our research.
- object to automated decision-making, to contest the decision, and to obtain human intervention from the controller.
- access (i.e., receive a copy of) your personal data that we are processing together with information about the purposes of processing, the categories of personal data concerned, recipients/categories of recipient, retention periods, safeguards for any overseas transfers, and information about your rights.
- have inaccuracies in the personal data that we hold about you rectified and, depending on the purposes for which your data is processed, to have personal incomplete data completed
- be forgotten, i.e., to have your personal data erased where it is no longer needed, you withdraw consent and there is no other legal basis for processing your personal data, or you object to the processing and there is no overriding legitimate ground for that processing.
- in certain circumstances, request that the processing of your personal data be restricted, e.g., pending verification where you are contesting its accuracy or you have objected to the processing.
- obtain a copy of your personal data which you have provided to the University in a structured, commonly used electronic form (portability), and to object to certain processing activities such as processing based on the University's or someone else's legitimate interests, processing in the public interest or for direct marketing purposes. In the case of objections based on the latter, the University is obliged to cease processing.

- complain to the Information Commissioner's Office about the way we process your personal data.

If you require advice on exercising any of the above rights, please contact the University's data protection team: data-protection@nottingham.ac.uk

CONSENT FORM

Date: 2018-08-28

Project: LiveTodos

School of Computer Science Ethics Reference: [Insert ref number]

Funded by: This research is supported by the Horizon Centre for Doctoral Training at the University of Nottingham (RCUK Grant No. EP/G037574/1) and by the RCUK's Horizon Digital Economy Research Institute (RCUK Grant No. EP/G065802/1).

Please tick the appropriate boxes	Yes	No
1. Taking part in the study		
a) I have read and understood the project information sheet dated 2018-08-28, or it has been read to me. I have been able to ask questions about the study and my questions have been answered satisfactorily.	☐	☐
b) I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	☐	☐
c) I understand that in taking part in the study there is a potential risk of information entered into the todo list website becoming unavailable.	☐	☐
d) I understand that taking part in the study requires me to provide data and that this may involve audio and video recording of me composing music, audio recordings of interviews, automatic collection of data about my use of the todo list website, and the taking of photographs of materials produced during composition.	☐	☐
2. Use of my data in the study		
a) I understand that data which can identify me will not be shared beyond the project team.	☐	☐
b) I agree that the data provided by me may be used for the following purposes:		
– Presentation and discussion of the project and its results in research activities (e.g., in supervision sessions, project meetings, conferences).	☐	☐
– Publications and reports describing the project and its results.	☐	☐
– Dissemination of the project and its results, including publication of data on web pages and databases.	☐	☐
c) I give permission for my words to be quoted for the purposes described above.	☐	☐
d) I give permission for my visual image contained in photos or video gathered during the research to be used for the purposes described above.	☐	☐

Please tick the appropriate boxes	Yes	No
3. Reuse of my data		
a) I give permission for the data that I provide to be reused for the sole purposes of future research and learning.	☐	☐
b) I understand and agree that this may involve depositing my data in a data repository, which may be accessed by other researchers	☐	☐
4. Security of my data		
a) I understand that safeguards will be put in place to protect my identity and my data during the research, and if my data is kept for future use.	☐	☐
b) I confirm that a written copy of these safeguards has been given to me in the University's privacy notice, and that they have been described to me and are acceptable to me.	☐	☐
c) I understand that no computer system is completely secure and that there is a risk that a third party could obtain a copy of my data.	☐	☐
5. Copyright		
a) I give permission for data gathered during this project to be used, copied, excerpted, annotated, displayed and distributed for the purposes to which I have consented.	☐	☐

6. Signatures (sign as appropriate)

____________________	____________________	________
Name of participant (IN CAPITALS)	Signature	Date

If applicable:

For participants unable to sign their name, mark the box instead of signing ☐

I have witnessed the accurate reading of the consent form with the participant and the individual has had the opportunity to ask questions. I confirm that the individual has given consent freely.

____________________	____________________	________
Name of witness (IN CAPITALS)	Signature	Date

I have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

____________________	____________________	________
Name of researcher (IN CAPITALS)	Signature	Date

7. Researcher's contact details

Name: Anna Clarke

Email: psxjc7@nottingham.ac.uk

Provide the participant with a copy of the completed form either by email or hard copy as they prefer.

Instructions

Getting started

LiveTodos is a todo list website designed to work with Ableton Live.

1. Open LiveTodos by going to **livetodos.co.uk.**
2. Sign into LiveTodos.
3. Follow the instructions to install the plugin.
4. Now open **livetodos.co.uk** on the phone or tablet.
5. Sign in to LiveTodos.

Let's pretend you're working on a new song.

1. Open the "example-set.als" file on the desktop. This is the Live Set you've been working on.
2. In **LiveTodos**, you should see a new "Untitled Set" automatically appear.
3. In LiveTodos, click on the link to choose to work on this new Set.
4. Give your new song a name by clicking on the Edit link in the top right hand corner, then typing a name in the box.
5. Click the back button (⟵) in the top left to return to the Set you're working on.

Pretend that so far you've been sketching out ideas for a bit and you've managed to get a basic groove going for a verse and a chorus. You want to move right on to writing some lyrics while you're feeling inspiration, but before you forget, you want to make a note for yourself later on - to remember to make the hand percussion you're using on the chorus sound a bit more varied.

1. In **LiveTodos**, click on the box at the top of the screen and type in "make percussion more interesting".
2. You'll see that LiveTodos suggests some tags you can attach to your current todo. These are the current selected track, and the current selected scene (in Session mode) or the current loop region (in Arrangement mode).
3. In **Live**, create a new track for hand percussion, and make sure it is selected.
4. Now in **LiveTodos**, click on the todo you were creating again. You should now see that your new track has been suggested. Click on the suggested tag to add it to the tags list.
5. Click in the tags box and start typing "Chorus" - LiveTodos should suggest this tag to you as it's a name of a Scene you created. Click on the suggestion to add it.
6. Click in the tags box and type "tweaks" and press Return. This creates a new tag for you.
7. Now press Return again to create the todo.
8. Practice this by making more todos for different scenes and time ranges.

You can also add some extra information, if you like.

1. In LiveTodos, click on a todo.
2. Type in the "Notes" field to add more information if you like.
3. When you're finished, click the back button (⟵) in the top left to return to the Set you're working on.

When you've finished recording vocals, you go back and improve the chorus percussion by adding some hand claps. Time to mark the todo as done.

1. In **Live**, go to Session mode, and select the chorus Scene.
2. In **LiveTodos**, you'll see that the "Chorus" tag has automatically been highlighted so you can find todos with that tag more easily.
3. In LiveTodos, click or tap in the box next to the todo name to mark it complete.

You should know everything you need to know to use LiveTodos now. Let's get rid of the stuff we've been practicing with.

1. In **Live**, close the Live Set we're working on.
2. In **LiveTodos**, click Edit in the top right hand corner.
3. Click the "Remove this Live Set" button, and confirm your decision by pressing OK.

Today's task

Your friend is curating an exhibition at a gallery in Nottingham. In this exhibition, photos and paintings are being exhibited, with each item accompanied by a specially written piece of music created by a musician from the local area.

They've asked you to contribute a short piece of music, around two minutes long, to be played alongside an exhibit. Three of the exhibits don't have music yet (you can find these in a folder on the Desktop of the laptop) - pick which one you think inspires you the most and create something to go with it.

You've got a busy afternoon, but have found about an hour and a half to make something. Try and stick to **only** using the built-in instruments and effects that are already in Live's library to save time. You can record in new material using the MIDI keyboard if you like.

Please use LiveTodos to keep track of your plans and ideas as you work. You're expecting to get interrupted later on - someone from the gallery will be visiting to interview you for a quick biography to appear next to the exhibit. (This will actually happen - expect to need to take a break in your work!)

Study plan

-0:30 - 0:00: Setup room and equipment

0:00 - 0:10: Paperwork

Read and sign:

- consent forms
- privacy notice
- information sheet

0:10 - 0:30: Participant follows “Getting Started” tasks

0:30 - 1:10: Participant follows “Today’s Task” task

Start video, screen, and audio recording here

1:10 - 1:20: Scheduled interruption and brief biographical interview

- How long have you been active in making music and performing music?
- Have you ever worked writing, performing, or recording music as a full time job?
- What kind of music do you normally make?
- Invitation to take a comfort break and/or get a drink/snack

1:20 - 1:45: Second composition period

1:45 - 1:55: Closing interview questions

- Do you normally use todo lists when you are writing music?
- Did you use any other written or electronic notes when you were using LiveTodos?
- Did you find that using LiveTodos changed how you worked?
- Can you describe any occasions where you found using LiveTodos was helpful?
- Can you describe any occasions where you found using LiveTodos was less helpful?
- Do you have any other suggestions for how LiveTodos could be improved?

1:55 - 2:00 Final activities

- Stop recorders
- Compensation paperwork
- Providing instructions for at home use
- Reminder of ability to leave study at any time.

2:00 - 2:30: Packing up room and equipment