AN INVESTIGATION INTO NON-LINEAR SONIC SIGNATURES WITH A FOCUS ON DYNAMIC RANGE COMPRESSION AND THE 1176 FET COMPRESSOR

AUSTIN MOORE

A thesis submitted to the University of Huddersfield in partial fulfilment of the requirements for the degree of Doctor of Philosophy

The University of Huddersfield

October 2017

Abstract

Dynamic range compression (DRC) is a common process in music production. Traditionally used to control the dynamic range of signals and reduce the risk of overloading recording devices, over time it has developed into a creative colouration effect rather than a preventative measure.

This thesis investigates sonic signatures, distortion, non-linearity and how audio material is coloured during the music production process. It explores how methodologies used to measure distortion and timbre can be used to define the sonic signature of hardware compressors and other pieces of music production equipment.

A grounded theory and content analysis study was carried out to explore how producers use DRC in their work, how they describe its sound quality, which compressors they frequently use and which audio sources they process with particular types of compressor. The results from this qualitative study reveal that producers use compressors to manipulate the timbre of program material and select specific compressors with particular settings for colouration effects.

Tests were carried out on a number of popular vintage hardware compressors to assess their sonic signature. Firstly, a comparative study was conducted on the Teletronix LA2A, Fairchild 670, Urei 1176 and dbx165A. Secondly a comprehensive in-depth analysis was undertaken of the 1176 to fully catalogue its sonic signature over a range of settings and to compare results from a vintage Urei Blackface 1176 and a modern Universal Audio reissue. Objective analysis was conducted on the compressors using Total Harmonic Distortion (THD), Intermodulation Distortion (IMD) and tone burst measurements. Complex program material was analysed using spectrum analysis, critical listening and audio feature extraction. It was found the compressors all have subtle nuances to their sonic signature as a result of elements in their design colouring the audio with non-linear artefacts. The 1176 was shown to impart significant amounts of distortion when used in its all-buttons mode and with fast attack and release configurations. This style of processing was favoured by producers in the qualitative study.

Table of Contents

List of Tables

List of Figures

Dedications and Acknowledgements

I would like to thank everyone who helped me during the creation of this thesis and in particular my first and second supervisors, Dr Rupert Till and Dr Jonathan Wakefield. I am also very grateful for the support from my family, especially my mother, brothers and sisters and my girlfriend Maria.

This thesis is dedicated to my father, Samuel James Moore who sadly passed away during the final stages of the write up.

Chapter 1: Introduction

1.1 Introduction

This thesis and the accompanying audio files are the results of a study into the non-linear processing effects used in music production. The primary focus of this study is dynamic range compression (DRC) and more specifically the 1176 compressor. It is recommended that the reader makes good use of the audio material that supports the findings presented in chapters 5 and 6. The figures, plots, tables and audio are an important part of this thesis thus it is suggested the reader considers all of these components holistically when reading the report.

1.2 Motivation

In its broadest sense, the motivation behind this study was to add a technical approach to the methodologies used by scholars in the area of music production. This thesis investigates the use of DRC and analyses its sonic effects on a range of audio material. Furthermore, the research sets out to investigate the idea of sonic signatures and develop upon existing work in this area, the most notable being Bernd Gottinger's thesis *Rethinking distortion: Towards a theory of sonic signature* (Gottinger, 2007). The basic premise of sonic signatures is that all audio is shaped by distortion, colouration, and resonances during the tracking or mixing stages of a production, and these artefacts impart a sonic fingerprint on the audio material that was not originally present. In a more technical sense, the idea of sonic signatures can be thought of as the results of non-linear processing, the sometimes subtle and not so subtle effects of small distortions and colouration effects that are part of audio systems.

There has been some research into ways in which distortion in audio equipment can be minimized, but except for of Gottinger's thesis, there has been little research into how this non-linearity can be exploited by the producer or analysis into how it affects program material common to popular music production. Timothy Warner (2009) in his chapter from the book *The Ashgate Research Companion to Popular Musicology* has the following to say on the subject of music production related research and audio processing in particular:

> The changes that all these technological processes bring about are considerable and can directly impact on every aspect of the listener's music experience. Yet there are few psychoacoustic explanations or aesthetic conjectures forthcoming on why the extensive use of technology is so prevalent on modern recordings of popular music. Hence, while it is generally acknowledged by the recording community that judicious use of audio compression and distortion, for example enhance the impact of popular music recording –they have been a fundamental part of modern record

> mastering for several years- the psychoacoustic or musical roles these play in the recording have yet to be thoroughly investigated. (p. 134)

While this thesis does not directly focus on the casual music listener's experience it addresses the aesthetic conjectures of music producers when applying DRC and non-linearity in their productions. Moreover, it is anticipated the work in this thesis will act as a model for other researchers who wish to investigate the area of listener experience. The author presents this work as a first step in the creation of a database of music production sonic signatures. The methodology used in this thesis should be implemented and adapted for further studies that add to this database. Thus, another core motivation is to establish an optimal method and demonstrate how to apply it.

1.3 Other Related Work

In musicology, there is an increasing body of work in the area of music production (Bennett, 2010; Journal on the Art of Record Production, n.d.; Zagorski-Thomas, 2014; Zeiner-Henriksen, 2010). While this work is highly scholarly, it does not have a particular focus on the technical aspects of music production, or at least not to the technical level of this thesis. Other work in the literature such as that by (De Man & Reiss, 2014; Giannoulis, Massberg, & Reiss, 2012, 2013; Marui & Martens, 2005; Schneiderman & Sarisky, 2009) investigates and analyses music production equipment from a technical point of view. While this work provided the author with ideas with regards to technical testing procedures and analysis methodologies, it did not address the idea of sonic signatures or their impact on music production. Perhaps the closest related work to this thesis is the research by Campbell et al. (2014) that looks into distortion artefacts and signal masking as a result of DRC. There are some similarities in the testing methodology used by Campbell et al. but they focus on stereo mixes while this project primarily investigates compression at the channel level. Furthermore, their work is concerned with reducing distortion artefacts while this work examines production techniques that add more distortion. Additionally, this thesis utilises a qualitative approach, in that it incorporates working practices from industry professionals as a basis for some of the testing methodology.

1.4 Developing a Methodology

Zagorski-Thomas (2012) notes in his book *The Art of Record Production* that "The study of record production is not in itself a discipline, a field of study with its own clearly defined methodology" (p. 2). A similar observation is made by Wells in his

PhD thesis when he states that Music Technology (an academic discipline closely related to the study of record production) is a field of study that incorporates elements from a wide range of subjects and is, in essence, multidisciplinary (Wells, 2006). There exists no accepted methodology for the study conducted in this thesis. Therefore, a method had to be developed by the author. Investigating a range of methods from a number of disciplines and implementing or adapting the most appropriate approaches achieved this goal. Consequently, the work in this thesis is an example of a multidisciplinary approach. It utilises methodologies from electronic testing and signal analysis (objective measurements on hardware devices), qualitative research (grounded theory and content analysis), musicology (subjective analysis of music material using critical listening and spectrogram analysis) and timbre studies (extraction and analysis of audio features).

The mixed methodology implemented in this thesis will help academics in the field of music production and musicology. They can build upon the methods used in this study and apply them in their work. Furthermore, this study will appeal to those working within the music production industry and provide them with academic research on areas of their work that is currently available only in anecdotal form, found in online forums, blog posts and magazine articles. To that end elements of this thesis (Chapter 4 in particular) act as a transfer of discourse by helping codify common music production language. As well as being an academic, the author of this thesis has a background in the music production industry and was keen to keep the work focused on real-world music production processes. This philosophy was carried through into all areas of the research, and consequently, this thesis bridges a gap between academic studies and professional industry practice.

1.5 Research Questions

To understand the sonic signatures of DRC, the following core research questions need to be addressed:

1. Why and how do music producers use dynamic range compression and associated non-linear distortion in music production?
2. What is an appropriate methodology for investigating and understanding the sonic signatures of audio processing devices in a music production context?
3. What is the sonic signature of the 1176 FET compressor and how is it utilised within music production processes?

The first question is important as empirical data on the use of DRC is currently limited, particularly with regards to music production and mixing. This area is under-researched, and the results from this study will fill a gap in knowledge. Furthermore, this question is of importance to the objective testing carried out in this project. It gives it context and helps position its findings, so that they are not only of significance to scholars but also professionals working in the industry.

The second research question is highly important. It provides the theoretical background for the methodologies used in this thesis and creates a methodology for scholars in related disciplines to apply in their own work. The author of this thesis does not view this thesis as the end of a study, rather aiming it to be the beginning of many more studies.

Question three is significant as it provides objective empirical testing on one of the most commonly used dynamic range compressors in production history. There has been no other study of this kind thus the work in this chapter fills another gap in knowledge. Chapter 5 presents similar, but comparative, information and results on three other important compressors.

1.6 Chapter Overview

The chapters of the thesis are briefly discussed here to give the reader an indication of the structure and journey being set out by the sequence of chapters.

Chapter one introduces the thesis. Chapter two discusses the methodologies used, details the approach to the qualitative study carried out in Chapter 4, and provides justification for the methodology. Also, it reviews alternative methodologies to present the reader with relevant information they can consider for their scholarly work. The methods used for objective measurements and analysis of the audio material in chapters 5 and 6 are also discussed along with justifications for the approach. This section answers research question 2.

Chapter three addresses the core concept of non-linearity in music production by discussing colouration, sonic signatures, distortion and how these elements are introduced into audio signals and exploited during the production process. It is also an in-depth analysis of DRC design and a critical review and evaluation of the compressor units tested in the thesis. The technical differences between the units are highlighted and help prepare the way for the qualitative and quantitative evaluation of sonic signatures in chapters 4, 5 and 6. A significant function of this chapter is to help identify the key characteristics of compressor design that most affect and shape sonic signatures. Some of the material used in this chapter is based on a paper presented at the 2012 Art of Record Production conference in San Francisco (Moore, 2012). This chapter sets the context for the research.

Chapter four is a mixed methodology content analysis and grounded theory study into the use of DRC in music production. The primary purpose of this study is to investigate how producers use DRC, which sources they use it on, which compressors they use most often, if they have any preferred settings when using DRC and what is the resultant effect on sound quality. Settings gleaned from this study were implemented in the creation of audio material for chapters 5 and 6. The results of this study were also used in the analysis of audio content in an attempt to map the producers' descriptors onto objective audio features. This chapter helps to answer research question 1.

Chapter five is a comparative study of four popular styles of compressor used in music production. In this section, each compressor processes the same audio, and the resultant audio files are analysed and discussed to catalogue the unit's sonic signature. A series of audio extracts accompanies the discussion, which the reader is recommended to listen to while working through the chapter. This chapter highlights the core differences between the 1176 compressor, the primary focus of this project, and the other compressors tested in the section. A conference paper presented at the 2016 AES conference in Paris (Moore, Till, & Wakefield, 2016) adapted some of the material in this chapter. This chapter is a companion chapter for chapter six and helps to provide context on the 1176's sonic signature, particularly how it differs from the three other compressors.

Chapter six is a focused study that specifically analyses the 1176. The aim is to investigate the sonic signature of a range of settings found to be popular amongst producers in Chapter 4. Testing was made on the 1176 using complex audio material and test tones. Two 1176s are examined in this chapter to compare and contrast any differences between them. One is a modern Universal Audio reissue, and the other is a vintage Urei Black Face edition. This chapter helps to answer research question 3.

Chapter seven is the concluding chapter with overall conclusions presented and a summary given of potential areas for further work and investigation.

Chapter 2 : Research Methodologies

2.1 Research Methodologies Introduction

The following chapter considers the mixed methodology used in this thesis. Firstly, there is an exploration of qualitative methodologies, and then the method used in Chapter 4 is presented in detail. Secondly, an explanation of the objective methods employed in chapters 5 and 6 is shown with information on their origins in electronic testing and timbre studies provided. The application of these methods is then displayed to prepare the reader for the analysis of results in chapters 5 and 6.

Various qualitative research methodologies were reviewed to find the most appropriate approach to answer research question one. This question required the use of qualitative research methods as text interviews were to be analysed to explain how and why music producers used compression in their work. Thus, the following sub-chapters present a review of qualitative methods and elucidates the final approach implemented in this thesis.

2.2.1 Descriptors and Metaphors

Interviews with music producers were analyzed to investigate how they used compression and described its sound quality. During this research, it was important to get a solid understanding of what the producers were trying to communicate. Therefore, it was necessary to investigate methods for interpreting communication and define what is meant by communication.

The interests of business and industries influenced early work relating to communication. Research into communication of this kind often treated it as a method to influence the performance of individuals or organizations. Consequently, two significant interests in the field emerged, one relating to personal communication effectiveness, the other proficiency in system-wide communication. Putnam and Cheney (as cited in Putnam, Phillips & Chapman, 2017) define the latter as "the study of messages, information, meaning and symbolic activity that makes up an organization" (p. 127).

The topic of metaphor is often discussed by scholars of communication and is particularly relevant to this study. Lakoff and Johnson (1980) define a metaphor as a way of seeing something as something else. Theorists such as Ortony (1979) state that metaphors are a way to associate abstract concepts with the absolute (cited in Putnam et al. 2017). Thus, the descriptors used by producers in the analyzed texts are examples of metaphors. The producers are using metaphors to describe the objective sound quality of the compressors, and one can consider these descriptors as an example of their interpretative repertoire. They are paradigms of the terminology and phrases available within the culture of music production and are used to describe and understand the world they inhabit. Studying this lexicon can help us comprehend how music producers use compression and what they want to achieve with this process. Furthermore, real meaning can be given to these descriptors when they are mapped onto objective attributes of the audio, one of the core motivations for combining qualitative and objective methodologies in this thesis.

2.2.2 Qualitative Research Methods

As well as inferring the meaning of descriptors it was decided the study would gather quantitative data on how frequently descriptors were used. Additionally, a record was to be kept of how often particular compressors and types of compressors were mentioned in the interviews and used for numerical analysis. This would help to develop theories relating to how professional producers used compression in their work and give the objective testing focus and relevance. Thus, approaches for collecting data of this kind were reviewed to find the most effective methods.

Sommer (2006) notes that *content analysis* "is used to systematically summarize written, spoken or visual communication in a quantitative way." Treadwell (2013) states that content analysis is a quantitative technique because it requires the researcher to count the occurrences of whatever is of interest to the study (p.216). It is an unobtrusive research method, which does not require contact with people. Many content analysis studies involve looking at samples of the media, television, film and the internet. Examples of content analysis can include examining the frequency of occurrence of themes in films and television, keeping a tally of the words used in political speeches to describe a specific policy or counting how frequently a specific artist uses particular colours. Refer to the literature by Busch et al. (1994) and Neuendorf (2002) for more information on content analysis.

Altheide (1987) splits content analysis into two discrete categories, *Quantitative Content Analysis (QCA)* and *Ethnographic content analysis (ECA)* and states the two forms differ "in approach to data collection, analysis and interpretation" (p.66). Altheidie lists a number of differences between QCA and ECA in his paper. ECA is used to catalogue and understand the nuances of communication, proposes the collection of both numeric and narrative data and uses purposive sampling to handpick material that is considered appropriate for the area of research. In contrast, QCA is implemented to validate theories, generates numerical data and uses random or stratified sampling.

Grounded theory is an inductive method, meaning that a theory is developed towards the end of the study and comes as a result of observations made during the process (Goddard & Melville, 2001). Grounded theory gets its name because the theory is grounded in the data, "a grounded theory is one that is inductively derived from the study of the phenomena it represents" (Strauss & Corbin, 1991, p. 23).

The grounded theory methodology is illustrated in Figure 2-1.

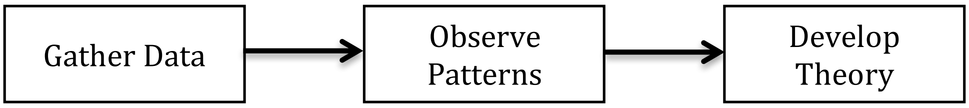

Figure 2-1: Grounded theory process

Grounded theory involves developing a theory that is based on the data gathered during analysis and enables the researcher to create a theory that offers an explanation into the main concern of the population under study. In grounded theory data collection and analysis occur concurrently. Thus, the data collection adapts as a direct result of the analysis. This attribute is another example of why grounded theory is considered inductive. The method is like a feedback loop with the focus of the research shifting depending upon the results of the on-going data analysis. Grounded theorists call this part of the process theoretical sampling. Glaser and Strauss (2009) define it as, "the process of data collection for generating theory whereby the analyst jointly collects, codes and analyses the data and decides what data to collect next and where to find them, in order to develop the theory as it emerges" (p.45). Furthermore, within grounded theory a hypothesis is not required as a starting point, rather the theory comes from the data and new ideas are created as the researcher works through the data. The reader is directed to the work by Charmaz (2006), Glaser & Holton (2007), Glaser & Strauss (2009) and Strauss & Corbin (1991) for more comprehensive detail on grounded theory.

Closely related to grounded theory is *discourse analysis,* but before a discussion of this method, it is worth defining discourse. Schiffrin et al. (2008) state the three main definitions of discourse are "(1) anything beyond the sentence, (2) language use and (3) a broader range of social practice that includes nonlinguistic and nonspecific instances of language” (p.1). The first definition could relate to studies carried out by linguistics into areas such as sentence diagraming or research into how people use sentences and paragraphs together. The second definition refers to language in use and is most relevant to conversation analysis. Description three may describe the work of sociologists in understanding critical issues such as power and how it manifests in language. Jørgensen and Phillips (2002) note that in many debates discourse is considered “the general idea that language is structured according to different patterns that people’s utterances follow when they take part in different domains in social life” and go on to propose a “preliminary definition of a discourse as a particular way of talking about and understanding the world (or an aspect of the world)” (p.1).

Hodges, Cooper and Reeves (2008) note that discourse analysis is the study and analysis of language and identified three variations in how discourse analysis can be implemented. They refer to the different approaches as formal linguistic discourse analysis (FLDA), empirical discourse analysis (EDA), and critical discourse analysis (CDA). FLDA examines texts to uncover rules and may involve coding words and sentences in related categories of meaning and structure. EDA uses conversation and genre analysis to look for significant themes and functions of language in social settings. CDA investigates how issues such as identity, power and social relations are constructed through written and spoken words and can be used to formulate ways in which this critique can bring about social change. A final variation of discourse analysis not identified by Hodges et al. is called *Discourse Tracing* and was introduced by LeGreco and Tracy (2009) to help scholars trace how discourse has changed over time and “to critically analyze the power relations associated with change” (p. 1517). As can be seen, all these variants have similarities and what underpins them is the notion of examining the data to get an understanding of its meaning at a “meta” level.

2.3 Conclusion and Choice of Qualitative Method

The previous sub-chapters reviewed various methods for qualitative research and as it was highlighted there is some overlap between the methods. Therefore, the correct method for a given study will depend upon several factors including the research question, whether there is a hypothesis to test, what form the data takes on and whether the analysis is to be statistical, textual or a mixture of both. The method used for qualitative research in this thesis was a mixed method of grounded theory with some content analysis to count how frequently descriptors, compressors and audio sources were mentioned in the data. Grounded theory is typically selected when there is little dialogue or debate about a subject in the academic literature. This point is exemplified by Suddaby (2006) who states that "When the researcher makes an inquiry when no relevant theory exists, grounded theory gives the researcher a creative approach without confining him or her to an already existing realm of theory." Thus, grounded theory was considered the most appropriate method for this study as there is no theory regarding how professionals use DRC in their work.

2.4 Implementation of the Qualitative Method

The following sub-chapters detail how the mixed method was used to generate the results discussed in Chapter 4. A significant part of the review was to investigate common language amongst engineers when describing the sonic characteristics of DRC. Data for analysis was gleaned from a number of interviews from music production magazines, and this data was analysed to observe if trends emerged in how producers described the sound of DRC. Data was analysed to find portions of the text where producers described the sonic signature of compression. The descriptors they used were extracted to count the frequency of occurrence and to develop categories relating to compression production techniques.

From here, a theory was developed to address a sub-question of research question one, namely what are the popular reasons to apply compression during the production process? Are producers trying to compress program material conservatively, transparently and not noticeably affecting the audio or are they seeking to impart some of the compressor's sonic signature on the material? Furthermore, a study of this kind facilitates an investigation into the descriptors used to describe the effects of compression and allows for mapping of these descriptors onto acoustic properties extracted from the audio in chapters 5 and 6.

2.4.1 Information on the Data Sample

Data in grounded theory can take on a variety of forms but it is typical for it to be word-based. For example, in the social sciences where grounded theory is widely used, data often consists of transcripts of interviews conducted by the researcher. During the collection process, data is analysed and coded and this involves examining the text for salient categories and grouping words (typically adjectives) and short phrases into these categories. Throughout the process, it is important to constantly compare the data and ensure the coding is being applied in a consistent manner. If a new category emerges from retrospective analysis a review of the previous coding should take place and words coded into these new categories accordingly. Categories in Chapter 4 were developed by grouping descriptors into categories relating to compression production techniques e.g. distortion, transient shaping, modulation.

Data collection and coding stops once the researcher has reached a point where it is felt that analysis adds nothing new. This point is called *saturation* by grounded theorists and is the criterion used in data collection to help the researcher judge when they should cease sampling. Glaser and Strauss (2009) note that "saturation means that no additional data are being found whereby the sociologist can develop properties of the category" (p. 61). In the study for this thesis, saturation was reached when no new categories relating to compression techniques emerged from the data.

Once the sampling process has been thoroughly saturated further categorization is carried out to group concepts that appear to relate to inductive categories. This process is called *axial coding*, and here categories are rationalized into related groups. From this point, the categories can be abstracted one stage further to create a core category or categories (typically no more than three), and these are used to underpin and support the narrative of the theory garnered from the research. This process was used to develop the two core categories presented in Chapter 4. These categories give us a better understanding of the main motivation behind the producer's use of DRC. Also, it gave focus to the processing approaches used in testing conducted for chapters 5 and 6.

Data for this study came from a pool of online interviews of recording and mix engineers discussing the production techniques used on well-known popular music tracks. The sample consisted of 100 articles from the *Sound on Sound* magazine series "Classic Tracks", 39 items from another Sound on Sound series "Secrets of the Mix Engineers" and a pool of 140 articles sourced from *The Mix* magazine's "Classic Tracks" series. All items created a total of 279 interviews. The articles spanned 14 years from January 1999 to January 2013. The discussions in these articles explored productions from the 1950s to 2010, and the music genres covered in the interviews were diverse, covering a range of styles including pop, heavy metal, hip-hop, dance music, rock and easy listening. To minimize any potential biases the magazines had for a particular brand of equipment or production style the articles were split almost equally between the two publications. It was decided to focus the search exclusively on these articles as they were grouped into a compendium of links. Thus, they could be easily found if the exercise needed to be repeated or validated at a later date. Other articles of a similar nature were found at the Sound on Sound site. They were rejected for analysis as they were not collated in an easy to find format. Furthermore, a cursory analysis of these articles found they revealed similar trends or focused on less relevant areas of the production process. This final point was true of articles written during the mid to late 1990s where the focus of the magazines was on the amateur bedroom producer, with much discussion on the affordable hardware, synthesizers, and samplers of that era.

The data was initially collected from a wide range of text sources, including popular books on music production and mixing audio. Early data analysis suggested the majority of these texts were not written by active music producers. Therefore, it was decided this material was not entirely appropriate, and the data collection was adjusted accordingly to focus exclusively on interviews with industry professionals. At first, the data was sourced wholly from Sound on Sound articles which provided the author with enough suitable material to construct well-developed categories. However, to judge if collection should be stopped, it was decided to extend the sample to include articles from the Mix online. Extending the search to include these additional texts allowed the author to ensure theoretical saturation had been achieved. Additionally, due to the Mix Magazine being an American publication it gave the search a more global focus. Sound on Sound is a British magazine, and it was thought that some provincial bias might be implicit in its texts. It should be noted that analysing the Mix articles did not create any new codes or categories, but it helped to reinforce those that had already been identified. Including the Mix added substantially to the sample size of the study. Although there are no specific suggestions for correct sample sizes to use in grounded theory studies, Mason (2010) recommends sizes in the range of 30-50. Thus, the sample of 279 used here well exceeds that recommendation.

During early stages of research, participant observation and interviews with producers were considered. However, it was thought the level of engagement by professionals would be small. Therefore, these methods were not pursued. Moreover, this type of research, while of great use, can be time-consuming to organize and at that juncture the author needed the qualitative data to move forward with arranging the measurements session detailed in Chapter 5. It is the author's belief the method used in this thesis generated a larger more diverse pool of data than could be achieved with the other approaches. The author recommends the use of multiple methods for areas of further study to compare and contrast the results.

2.4.2 Recording Descriptors

Adjectives and short phrases were extracted from the literature and coded to develop categories relating to compression production techniques. During this procedure, several rules were imposed on the coding to ensure the study was focused and appropriate. Most importantly a word or phrase was only recorded if it was a descriptor of the sound quality and was used to describe the result of applying compression to a track or drum buss source. Descriptions of parallel compression or mix buss compression were not recorded, as these techniques were not the focus of this study. Furthermore, a word was not coded if it was used as a verb. For example, the word crush was coded if used in a context such as "this gave us a crushing sound" but it was not coded if the producer stated, "I crushed the room mics using a compressor". It is not clear from the second example if the word crush is being used to describe the sound quality or describing the manner in which compression was applied. Multiple repetitions of the same word used within the same thread of discussion were recorded once. Finally, words were only coded if it was clear that the compressor was the piece of equipment having the purported effect on the sound quality. If the statement included other pieces of equipment, a microphone preamplifier, EQ and compressor chain for example, the word was not recorded, as the sound quality could not be attributed solely to the compressor.

2.4.3 Grouping by Compressor, Gain Reduction Type and Source

Searches were carried out in the texts for the words compress, compressor, compressed, opto, FET, VCA, valve, tube, vari-mu. From here relevant areas of the text were examined, and relevant data coded into groups relating to specific instrument sources, particular dynamic range compressors and gain reduction types (FET, Opto, VCA, Valve, other). If the compressor model, type of gain reduction style or source was not discernible the data was put into a general category. Coding into groups allowed the codes to be mapped not only onto compression as a whole but also onto specific types of compressors and audio sources. This process allowed for a lower level of data analysis to be carried out and resulted in a more focused insight into the application of compression.

During data collection, the majority of discussions in the literature used for this study related to drums, bass, and vocals. Thus, to give the study emphasis, coding was focused on vocals, drum busses/rooms/overheads, bass instruments (including bass guitar, double bass, and synth bass) and what can be collectively referred to as membrane-based drums i.e. kick, snare and tom toms. Data was coded into a general category for any discussion of other music sources or where the source was not explicitly named.

This concludes the discussion of qualitative methodologies and the results from the method described in Chapter 2.4 up to this point are discussed thoroughly in Chapter 4. The information garnered from this study was also responsible for the choice of compressors and sources tested in chapters 5 and 6

2.5 Objective Methodologies

The following sub-chapters provide essential background on the objective methodologies used in chapters 5 and 6. They consider how the amount of distortion in an audio system is measured, evaluate the strengths and weaknesses of the methodologies and discuss how successful they are in providing meaningful data, particularly with regards to sonic signatures. They investigate the use of Total Harmonic Distortion (THD), Intermodulation Distortion (IMD) and tone burst measurements. Additionally, they explore how Fast Fourier Transform (FFT) based spectrum analysis and audio feature extraction procedures can be of good use in the study of sonic signatures. These methods were used to build a comprehensive overview of the sonic signature of the compressors tested in this thesis and allowed for objective meaning to be inferred from the descriptors used by the producers in the qualitative study.

2.6 Distortion Measurements

Audio systems can be measured using numerous metrics including amplitude measurements, frequency response, and dynamic range measurements. In addition, a system can be measured for linearity, and this involves measuring the amount of distortion present. As this study investigates the non-linear sonic signature of compression, it was necessary to make distortion measurements on the compressors tested in chapters 5 and 6. This sub-chapter discusses how distortion is evaluated in an audio system.

When measuring distortion, the input is compared to the output and any differences observed at the output that involve additional harmonic content, are considered non-linearity. Other measurements such as the frequency response and phase response can be conducted, but these measurements are not considered measures of non-linearity. Variations of this kind are regarded as a consequence of linear distortion, a process discussed in Chapter 3. The most common way to express distortion in an audio system is THD. In this measure, a single sine wave test tone (typically of 1kHz) is used because it shows if any additional components are added to the output. The THD figure is either articulated per harmonic component (expressed as the number of decibels lower than the test tone) or as a root-sum-square calculation. The calculation is worked out by taking the amplitude of the first *N* harmonics, performing an RMS summation, dividing this by the amplitude of the fundamental and presenting the result as a percentage (Kester, 1993). When measuring light to mild non-linearity, it is only necessary for the lower order harmonics to be included in the THD calculation (Berners, 2009a). The amplitude of the higher order harmonics typically drops off sharply after the first few non-linear components thus rendering them inaudible due to auditory masking. However, more harmonics should be used if the harmonic distortion is extreme and the amplitude of the higher order artifacts is of a considerable level. The measurement will be more accurate and representative of how the non-linearity is perceived by the listener. The psychoacoustic models that affect the perception of non-linearity are outside of the scope of this thesis but the reader is directed to research by Geddes and Lee (2003) and Voishvillo (2007) for more information.

Temme (1992, p.9) states that test tones of varying frequencies should be used for a more comprehensive insight into any frequency related aspects of non-linearity in the audio system under test. Therefore, it is prudent to make some measurements at different frequencies to observe any variation in non-linearity over the frequency spectrum. Outside of time domain analysis, the use of non-sine test tones and complex waveforms is not usually recommended for distortion measurements. These signals may be rich in harmonic content that will mask the harmonic distortion artefacts, making input versus output analysis difficult (Berners, 2009b). Square waves can be used to observe changes in the shape of the wave which can indicate variations in the frequency response. However, as noted by Elliot (2015) a square wave is of no use in testing distortion due to the waveform itself having theoretically 48% THD.

Intermodulation distortion (IMD) is another problematic form of non-linearity. As with harmonic distortion, IMD is measured using sine test tones, but for this test two sine waves need to be used concurrently. There exist two tests that can be conducted for IMD, SMPTE or ITU-R. The administrator of the test should decide which one to use in their test and state clearly the methodology in the results. The SMPTE method uses two test tones of 60Hz and 7kHz with an amplitude ratio of

1:1. The output is then examined by measuring the IMD artefacts (that appear as sidebands) created by the modulation effect of the input tones (Bohn, 2000). High IMD results do not necessarily signify a lot of audible non-linearity and unless the artefacts are of considerable amplitude they will be masked by the test tones. Thus the observation of measurement results can be misleading, and the effects of masking need to be taken into consideration when discussing the findings (National Association of Broadcasters, 2007).

2.7 Observing Distortion

Non-linear distortion has the effect of altering both the time and frequency domain of audio waveforms. Changes in the time domain are observed using a time domain plot, and changes in the frequency domain are observed using a spectrum analyzer, which shows the audio as a function of its frequency. Looking at the audio in the frequency domain enables the frequency components and the magnitude of these components to be analyzed, which can then be used to calculate THD manually. These non-linear components can be used in calculations to create THD percentages or expressed as a decibel down from the test tone frequency. Figure 2-2 shows a square wave analyzed using a spectrum analyzer where the frequency components of this wave can be clearly seen. It is typical for this type of display to be referred to simply as an FFT, named after the Fast Fourier Transform process that is used to transform the time domain into the frequency domain.

Additionally, complex audio signals, particularly those that develop over time, can be observed using a spectrogram which plots frequency on the x-axis; time on the y-axis and represents the intensity of frequencies with colour. Figure 2-3 shows a spectrogram of a middle C piano note where the intensity is colour coded from blue to red with red being the highest amplitude. A thorough discussion of how spectrum analyzers work is beyond the scope of this thesis but a comprehensive overview of the process and the theory underpinning is given by Roads (1996, pp. 1073-1112).

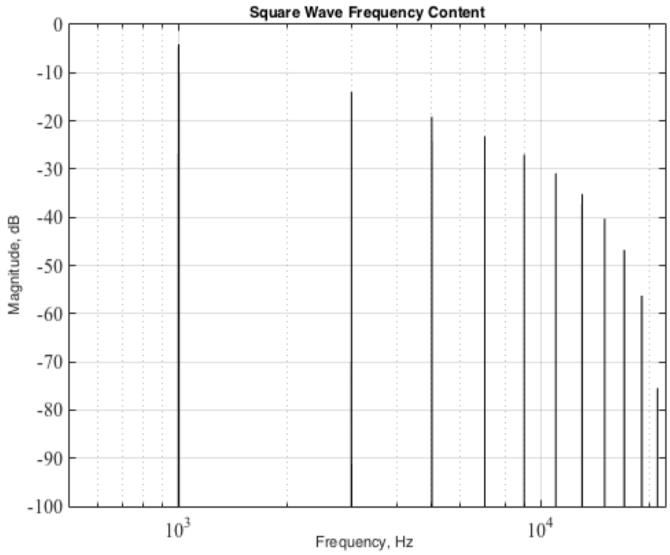

Figure 2-2: FFT of a square wave

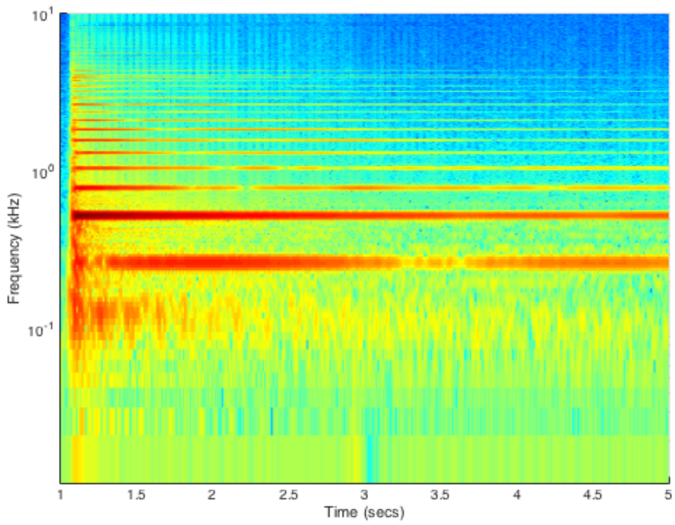

Figure 2-3: Spectrogram of a middle C piano note

2.8 Tone Burst Measurements

Although not a measurement of distortion, tone burst measurements are often conducted on dynamic range compressors to observe their timing behaviour. This measurement is made by creating a burst test tone consisting of low-level and high-level amplitude portions. An example of a burst test tone is shown in Figure 2-4. To make the measurement the compressor is adjusted so only the high-level parts exceed the compressor's threshold. The tone is then run through the compressor's input, and the output recorded to a computer to observe any changes in the tone's shape. Tone bursts with the high-level portion at various durations (fast and slow) can be used to investigate program dependency in the timing characteristics of a compressor. Tone burst measurements are recommended by EBU (1992, p.11) as a method to measure dynamic processors.

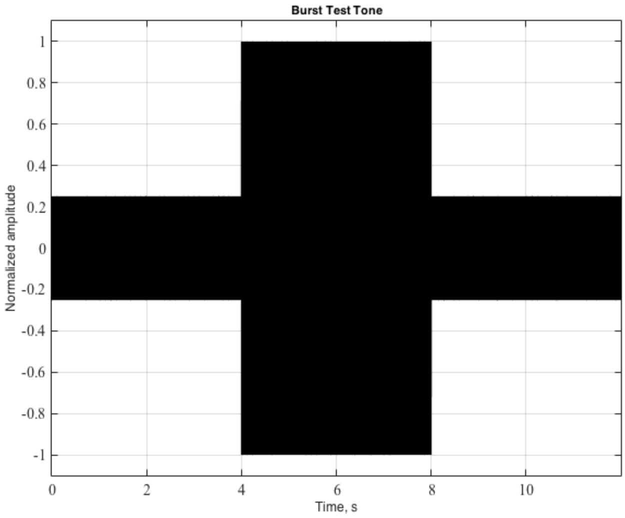

Figure 2-4: Tone burst signal

2.9 Implementation of Objective Measurements

THD measurements were made on compressors in chapters 5 and 6 using a number of different frequencies that are detailed in the relevant chapters. Additionally, IMD was measured on two 1176 compressors using the SMPTE method and the results are presented in Chapter 6. These measurements were made by feeding test tones to the compressors and recording the output into a Digital Audio Workstation (DAW). FFTs of the audio files were created using the spectrum analyser in Izotope RX to observe the distortion components. The peak picker function was used to extract the level of these components, and they were used to manually calculate THD and IMD percentages. The calculations used the root-sum-square method with up to ten of the non-linear harmonic components included in the sum.

In chapters 5 and 6 spectrum analysis plots of complex audio signals (music sources) were created using MATLAB's spectrogram function using a Hann window, an FFT size of 1024 or 2048 (depending on the source) and an overlap of 75%. These spectrograms were used to give a visual representation of the audio and highlight areas of the compressor's sonic signatures.

Tone bursts were employed in chapters 5 and 6. The tone bursts were sent from a DAW into the compressor's input and the output was recorded back to the DAW.

The measurements of non-linearity during release in chapters 5 and 6 were measured by taking an average of the spectral content during the release portion of the tone burst. This was implemented using Izotope RX, which allows the user to select a portion of audio in the time domain, the FFT averages this audio and then displays the spectral components present. The release portion was measured as the point at which the compressor came out of gain reduction up to a steady state level. The averaged spectral components were then used in a THD percentage calculation using the root-sum-square-method.

2.10 Timbre Studies

Analysis methods that have their origins in timbre studies were used for objective measurements to explore the effect of the compressors on complex program material. Thus, it is necessary to define the meaning of timbre and explore methodologies that can be used to categorize it. Timbre is an important aspect of sound quality and it is a key component of many audio related activities, including the identification and perception of sound, speech recognition, orchestration, acoustic instrument design, and synthesis. Timbre is an acoustical parameter that allows the listener to distinguish between two sounds of equal loudness and pitch. For example, we can discern the difference between a violin and a clarinet playing the same note and at the same volume by differences in their timbre. Variations in timbre can occur as a result of non-linearity within pieces of audio equipment, as well as in musical instruments. Case (2007) states:

> Musicians think of timbre as that attribute of sound that separates a piano and a guitar when they play the same note. Engineers must sort out what sonically separates a Steinway from a Bosendorfer among pianos and a Les Paul from a Telecaster among guitars, even as these instruments play the same note. Trained and careful listeners can reliably distinguish very fine properties of sound. The unique distortion traits of each audio device in the studio endow that device with a sort of sound fingerprint related to distortion. (p. 97)

Furthermore both White (2010) and Robjohns (2010) note that harmonic distortion in audio signals results in a perceptual change to the timbre of program material. Thus, timbre research is an important area to consider in an investigation into sonic signatures.

The standard description of timbre that comes from the American National Standards Institute (ANSI) states, “timbre is that attribute of auditory sensation in terms of which a listener can judge two sounds similarly presented and having the same loudness and pitch as being dissimilar" (ANSI, 1973). While the ANSI statement is of some use as a starting point in discussions of timbre it is rather vague. The main area of contention with the ANSI statement is that it focuses more on what timbre is not and not on what it actually is at an objective level. Fitzgerald and Lindsay (2004, p.1) argue this ambiguity arises from the fact that timbre cannot be measured on one scale. Pitch for example can be measured and referenced on the frequency scale and loudness on an amplitude scale. This is not the case with timbre because it constitutes many acoustical properties that act collectively. As a result, it is not feasible to measure timbre on a single scale and a solution is to measure it in a number of dimensions and extract specific features of the audio signal in order to categorize its perceptual sound quality.

2.11 Audio Features

There is much interest in establishing an objective measure of timbre perception. Research into this area involves measuring timbre as a multidimensional property. A popular method is executed by first conducting subjective listening experiments where listeners assign semantic descriptors to the sound (bright, dull, sharp for example) and then carrying out multidimensional scaling analysis of pairs of musical instruments to rate their similarity/dissimilarity in a number of timbre dimensions. These dimensions are called timbre spaces and the distances of sounds within these spaces are computed and analyzed for any relationships. Examples of this type of research can be seen in the literature by (Fitzgerald & Lindsay, 2004) and (Alluri & Toiviainen, 2010). From the research into timbre a number of salient features of sound were discovered (Iverson & Krumhansl, 1993; Lakatos, 2000; McAdams, Beauchamp, & Meneguzzi, 1999) and it has become ever more important for researchers to align subjective descriptors with physical correlates.

The extraction of timbre-related audio features is used outside of timbre research. Computer musicians utilise features to manipulate synthesized sound sources and to automate the process of categorization of music in database collections. Additionally, features have been extracted for research into speech recognition, transient detection algorithms, and for use in legal cases relating to music copyright claims. More relevant to this thesis is the use of feature extraction in the design of software tools to automate tasks in the mixing process (De Man & Reiss, 2014; Fenton, Lee, & Wakefield, 2014) and the investigation of timbral similarity in Electronic Dance Music (Rocha, Bogaards, & Honingh, 2013).

Audio features are extracted using a number of different algorithms and software packages including Sonic Visualizer and VAMP plugins, MATLAB toolboxes, Max/MSP externals and open source tools such as Essentia that was developed for research by the University of Barcelona (Bogdanov et al., 2013).

2.11.1 Audio Features Extraction Process

As can be seen in Figure 2-5, the feature extraction process works by first taking the signal x [n], then a section of it is windowed by w [n] and the spectrum is computed using the FFT. The output of the FFT is the magnitude and phase spectrums and from here the features are extracted.

The extraction process can be based on single frames or multiple frames, meaning some features are computed and extracted with single frames while others need multiple frames for computation. Multiple frames are typically used for features that need to be calculated as they develop over time. Some audio features such as those discussed in chapters 2.12.1 and 2.12.3 can be calculated directly from the time domain.

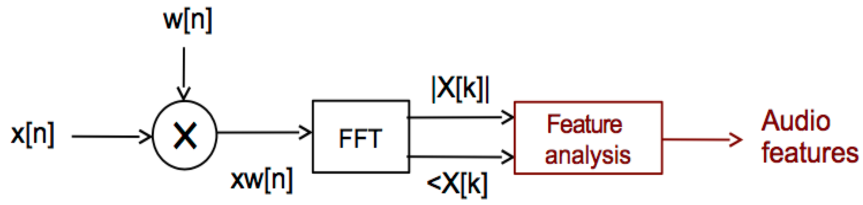

Figure 2-5: Audio feature extraction process

There are many types of features that can be extracted from audio signals and the specific feature or group of features the researcher needs to extract is dependent upon the research aims. One audio feature will not be enough to help describe the timbral characteristics of an audio signal therefore researchers will extract and combine several features to build up a comprehensive picture of the sonic signature.

Audio features are grouped into several categories and they can vary slightly between the tools used for extraction. The following groups utilised in the popular Essentia toolkit provide a good overview of the most common feature groups.

1. Spectral descriptors
2. Time-domain descriptors
3. Tonal descriptors: These descriptors relate to pitch based information and are normally derived from the spectrum of the sound

4. Rhythm descriptors: As the name suggests this group relates to the rhythmic aspect of sound. These descriptors are extracted from the time domain or the spectral domain of the sound.
5. SFX descriptors: These descriptors are typically related to sound characteristics such as the energy or the attack of a sound,
6. High-level descriptors: Descriptors in this category attempt to identify higher-level aspects of the sound (Verfaille, Arfib, Keiler, von dem Knesebeck, & Zölzer, 2011). An example is the danceability feature proposed by Streich and Herrera (2005) which correlates low detrended fluctuation analysis (DCA) values with music that has been subjectively labeled as rhythmic

As a point of comparison, the research group CUIDADO (Content Based Unified, Interfaces and Descriptors for Audio/Music) categorises audio features in the following manner:

1 Temporal shape features such as attack time
2 Temporal features such as zero crossing rate
3 Energy related features such as global energy and noise energy
4 Spectral shape features such as spectral spread
5 Harmonic features such as harmonic/noise ratio and odd to even ratio
6 Perceptual features computed using a model of the human hearing system (Peeters, 2004).

As can be seen, there are differences in categorization between tools but the majority of the features are grouped similarly with only minor differences

2.12 Details on the Features Extracted in This Thesis

In the following sub chapters a brief discussion of the features used for analysis in this thesis will be discussed to justify their use. The following publications should be read if more details on audio features is required (Bogdanov et al., 2013; Olivier Lartillot, Toiviainen, & Eerola, 2008; McKinney & Breebaart, 2003; Mitrović, Zeppelzauer, & Breiteneder, 2010; Peeters, 2004). All audio features extracted in this thesis made use of MIRToolbox for MATLAB (Lartillot, 2011) or using MATLAB scripts.

2.12.1 Energy Related Features

The energy of an audio frame can be computed from the magnitude spectrum and this is achieved by summing over the square of the magnitudes in the time domain. *RMS Energy* is another energy-related measure and like the standard energy measurement it is derived from the magnitude in the time domain. This time it is achieved by calculating the square root of the arithmetic mean of the energy. RMS energy envelopes were extracted for analysis in this thesis as they can reveal temporal changes that are difficult to observe on a simple time domain plot. They were created using a MATLAB script adapted from the work by Verfaille et al. (2011, p. 329)

Low-energy is a feature defined by Tzanetakis and Cook (2002, p.40) as the percentage of analysis windows that have less RMS energy than the average energy across all analysis windows. A small low-energy value suggests a piece of audio is continuous or consistent in its energy distribution. Low-energy was extracted to observe the effect different compressors and settings had on sonic signature of audio material.

2.12.2 Spectrum Related Features

The *spectral centroid* aims to characterize the spectral shape of a sound and gives an indication of where the center mass of the spectrum is located, it is often called the spectral center of gravity and can be a good determinate of around which frequency area most of the sound energy is concentrated. Perceptually it is associated with the sense of brightness and is calculated as the weighted mean of the frequencies present in the signal. The calculation sums over the total spectrum and weights it by every frequency before normalizing it by the total energy. The spectral centroid can be specified as a discrete value, an integrated value or displayed on a time varying plot. The spectral centroid was extracted in this thesis to assess if any colouration had occurred to the audio signal from the use of compression. The main motivation for extracting this feature was to observe any subtle colouration changes to the signal that might be perceived as brighter or warmer. Chapter 4 demonstrated that descriptors of this kind were used to describe the sound quality of DRC.

In addition to spectral centroid, *spectral brightness* was extracted to assess frequency related colouration. The feature is defined as the amount of energy above a given cut-off frequency and the result is expressed as a number between zero and one. The cut-off frequency used for brightness features extracted in chapters 5 and 6 was 1.5kHz which is the default in MIRToolbox.

Spectral flux (sometimes called spectral variation) represents the amount of variation or rate of change in the spectrum of an audio signal over time. In simple terms this can be thought of as a measurement of fluctuations in the spectrum. It is calculated by computing the difference in amplitude of adjacent FFT bins between successive frames and then sums over all the differences. This gives a measure of spectral variation and it can be used to describe the amount of spectral change in a sound. Spectral flux has been used in transient detection algorithms that work by measuring changes in energy at the onset of an audio signal, high changes in energy typically correlate with high spectral flux.

For this thesis the effect of DRC on spectral flux was observed in sub-bands by extracting audio using elliptical filters, this was conducted to investigate the findings of Alluri and Toiviainen (2010) who found a strong correlation with fluctuations in the lower end of the frequency spectrum (below 200Hz) and the subjective quality of fullness. In addition, they found an increase in spectral flux in high frequency bands correlated with a sense of liveliness therefore spectral flux in high frequency bands was extracted for specific analysis in chapters 5 and 6.

Spectral rolloff is used as an indication of how much energy resides in low frequencies. This frequency is the point below which *N*% of the magnitude distribution is concentrated and *N* is typically 85% or 95%. Spectral rolloff was extracted on audio material analyzed in chapters 5 and 6 where it is found that DRC has the effect of increasing the rolloff.

In addition to the features discussed above, other spectral descriptors were used for a *Principle Component Analysis* (PCA) study in Chapter 5. These features were *spectral flatness,* used to characterize how noise-like or tonal-like a sound is, *spectral spread,* a measurement of how spread the spectrum is around its mean value and *spectral kurtosis* which is a measure of how flat the distribution is around the mean. The PCA method is detailed in Chapter 2.13.

As noted by Löfqvist and Mandersson (1987), the *long-term average spectrum* (LTAS) can be used to display the average spectral distribution of a signal over time. It has been used in previous studies such as the classification of classical male singing voices by Johnson and Kempster (2011) and analysing the spectral characteristics of popular music productions by Pestana et al. (2013). Thus, the author proposes its usefulness in classifying the sonic signatures of DRC. The LTAS of audio material processed through compressors was investigated in chapters 5 and 6 and was calculated with 1/16 octave smoothing using a MATLAB function by Hummersone (2016).

2.12.3 Time Domain Related Features

The attack time is a temporal feature that measures how quickly energy increases after the onset of a sound. It can also be defined as the duration it takes for a sound to become perceptually audible and reach maximum intensity. The start point is worked out by estimating a point where the signal has reached 20% of its maximum value (to account for noise) and the end point is estimated to the point where the sound is at 90% of its maximum value (again to account for noise). Eidenberger (2011, p.65) notes that quick attacking sounds are perceived as sharp and loud. Therefore, attack time is an appropriate feature to extract when measuring punch in percussive sounds. Fenton et al. (2014) use a similar approach in their study of punch in drum sounds and extract the attack time feature for analysis of sounds that were subjectively rated as sounding punchy. The attack time feature was extracted on audio in this thesis to compare the transient shaping effect of a number of compressors on snare drums. The results can be seen in Chapter 6.

The *zero-crossing rate* (ZCR) is extracted using the time domain representation of the audio and is measured to give an idea of how noisy the audio signal is. This feature measures the number of times the signal crosses zero. A signal with a high zero crossing rate indicates that it is a noisy sound. ZCR can be used to classify percussive sounds and percussive genres in music classification but it can also be utilised to distinguish patterns in speech, particularly sibilants, plosives and fricatives. ZCR was extracted in this thesis to observe additional noise introduced to the signal as a result of non-linearity. It was found by Campbell et al. (2014) that IMD artefacts created by DRC can sound noise-like.

2.12.4 Psychoacoustic Related Features

The feature *roughness* can be extracted from a signal and is considered to be an important feature related to musical dissonance (Laurier, Lartillot, Eerola, & Toiviainen, 2009). Roughness occurs as a result of envelope modulations in the 20-150Hz range. It has been used in previous studies into music classification (McKinney & Breebaart, 2003) and timbre and rhythm similarity in electronic dance music (EDM) (Rocha et al., 2013). The latter study mapped the descriptor dirtiness to roughness in order to quantify the distorted sound quality of EDM synthesis. Like ZCR roughness was extracted in this study to investigate non-linear distortion artefacts from DRC.

2.13 Principle Component Analysis

Principle component analysis (PCA) was used in Chapter 5 to look for similarities between different compressors. This was accomplished by using data from audio features extracted from compressed complex program material. PCA was implemented because it is a data reduction method that takes a large data set and reduces the dimensionality of the data whilst preserving its most significant variation. To achieve this the PCA algorithm identifies directions in the data in which the variation is at its maximum. To understand the algorithm one can imagine a number of data points in an n-dimensional space. To find where most variation is located on the data points a line is plotted and the amount of variation in data is measured. This process is repeated until the direction of the line with most variation is found and this is called a principle component. The direction of this line is called the eigenvector and the measure of variance in the direction is called the eigenvalue. The eigenvector with the highest eigenvalue is called the first principle component. Carrying out the same process and plotting lines that are orthogonal to the first component locates the next eigenvector. The number of eigenvectors used in analysis is equal to the number of dimensions in the data and when all eigenvalues are summed they account for 100% of variation in the data. However, some eigenvalues will make up for a large proportion of the data while others may account for small proportions. It is common for the first two or three eigenvalues to account for 90-95% of variation in the data and these are typically the eigenvectors used for analysis. The resultant data is then displayed on a 2D or 3D graph with each axis displaying the variation in the data and where it is most spread. The X-axis spans most variation and is known as principle component one and the Y-axis spans the second most variation in the data and is known as principle component two. A full description of the process is beyond the remit of this thesis but a more thorough overview can be seen in the work by Smith (2002).

PCA has been used in several studies that are related to music technology and examples can be seen in the literature by Sandell and Martens (1995), Bech (1999), Loughran et al (2008) and Stowel and Plumbley (2010). The author could not find any papers directly related to this project that used PCA but it is proposed that PCA is beneficial in the study of sonic signatures, particularly if a large number of audio features are extracted and the researcher wants to reduce the number of variables to a lower number. These principle components can then be plotted as samples on a score plot for visual scrutiny and the researcher can look for correlations between the variables and/or the observations. More thorough methods of clustering such as K-Means clustering should be used if the researcher wants to get a better understanding of how the data clusters but PCA is adequate for visualisation purposes and assessing similarities and differences between the plotted samples.

For the PCA study in this thesis, the variation in audio feature data (variables) is plotted in PCA 1 and PCA 2 and any similarities or dissimilarities between the compressors (observations) can be gleaned by looking at where they plot on PCA score charts. This type of analysis is used in the study of genetics to observe how different genes cluster on score charts, which allows the researcher to note the similarity and dissimilarity of genetic clusters. An example of this approach can be seen in the work by Pollen et al. (2014).

2.14 Conclusions on Research Methodologies

Several qualitative methodologies were discussed and evaluated in this chapter and the author believes that all of these methodologies are appropriate for future studies into sonic signatures. It is recommended that methodologies that collect both numeric and narrative data should be used as they allow for objective results to be generated out of any subjective data. Using this approach, a robust and comprehensive set of results will be created that can be presented in a variety of forms such as tables, charts, word clouds, dendrograms and statistically analysed if required. For the work in this thesis, a hybrid method of grounded theory and content analysis was implemented for the study in Chapter 4 to extract data on how producers use DRC and describe its sound quality. These results were then combined with objective testing (particularly audio feature extraction) to get a better understanding of how these processes affect compressed audio at a perceptual level.

The use of the objective measurements THD, IMD and tone bursts were considered in this chapter and justifications for their use given. These measurements are of good use when creating objective data on how a piece of equipment performs but they should be used in conjunction with critical listening and visual FFT analysis to get a better understanding of sonic signatures at an objective and a perceptual level. FFT, IMD and tone bursts were used to assess the non-linearity and timing behaviour of a number of compressors in chapters 5 and 6. FFTs and spectrograms were created of tones and complex program material processed with compressors to illustrate their sonic signature and non-linearity.

Audio feature extraction was discussed and highlighted as a powerful method for audio research. Feature extraction is a useful tool for any researcher hoping to classify the sonic signature of music production devices and serves as a valuable addition to other higher-level analysis options such as spectrograms and time domain plots that do not always correspond entirely with human perception. It should be kept in mind that feature extraction can only give the researcher objective acoustic data. Therefore, it is suggested this method is used in conjunction with critical listening and objective measurements to get a comprehensive picture of the audio's sonic signature.

Finally, the PCA method was detailed, and it was shown how it could be used with an extensive audio feature data set to look for similarities between the devices under test. Audio features were extracted from audio material compressed with a number of compressors in Chapter 5 and PCA was used to look for similarities between the compressors under different states of gain reduction.

Chapter 3 : Non-Linearity, Audio Equipment and Dynamic Range Compression

3.1 Non-Linearity in Music Production

Before an investigation into non-linear sonic signatures can take place, it is first necessary to define non-linearity and detail how it can be introduced to an audio signal. If the transfer function of a piece of equipment is non-linear, then the waveshape at the output will be an altered form of the waveshape at the input (Metzler, p.23). In addition to changing the waveshape, additional harmonic components that were not present in the original signal will be created. This type of non-linearity is referred to as non-linear distortion. Linear distortion, on the other hand, refers to changes in the input signal that while potentially transforming the shape of the wave do not add additional harmonic components (Preis, 1976). Only changes to signal components already present in the signal will take place. Linear distortion can include irregularities or changes to the frequency and phase response. It is important to differentiate between these two forms of distortion at this juncture and state that distortion in this thesis pertains to non-linear distortion unless otherwise indicated.

From a historical perspective, non-linearity was something that equipment manufacturers tried to avoid and the general trend in equipment design was a move towards ever decreasing amounts of distortion and more transparent audio signals (Fenton, Fazenda, & Wakefield, 2011). While it may be fair to say that music producers also wanted this increased transparency from recording equipment there seemed to be a limit to how much distortion they wanted to lose. This point is addressed by Hamm (1973) in a paper that investigates the audible differences between valve and transistor recording equipment. As part of the work Hamm collected a number of quotes from musicians discussing the sound of transistor and valve based equipment and it was found they all had a preference for the valve sound. He stated the reason for their preference was due to the non-linear distortion of valve equipment when compared with the transparency of transistor based devices. Despite Hamm's paper being over 40 years old this opinion has not diminished and more recent literature reveals that audio engineers still have a similar point of view (Anderson, 2006; Kemp, 1999; Liniere, 2013; Vintage King, 2014).

Distortion can be introduced to an audio signal by a piece of equipment, and this affects its sonic signature. In an ideal situation, a theoretically perfect amplifier will not induce distortion; it should be able to accurately reproduce the input signal at the output with no change to the signal (Millett, 2004). The function of an amplifier can be represented as a graph with input vs. output on an X-Y plot, this is called the amplifier's transfer function. The plot in Figure 3-1 shows a perfect transfer function, it is a straight line and the input equals the output. This is considered a linear mapping of the input and output. Notice that the line is at a 45-degree angle. If a hardware designer was capable of producing a perfectly linear amplifier, one that simply scaled the signal up or down, the amplitude would plot as a straight line steeper than a 45-degree angle for an accentuated level or a straight line less than 45-degrees for an attenuated level (Case, 2007, p. 90).

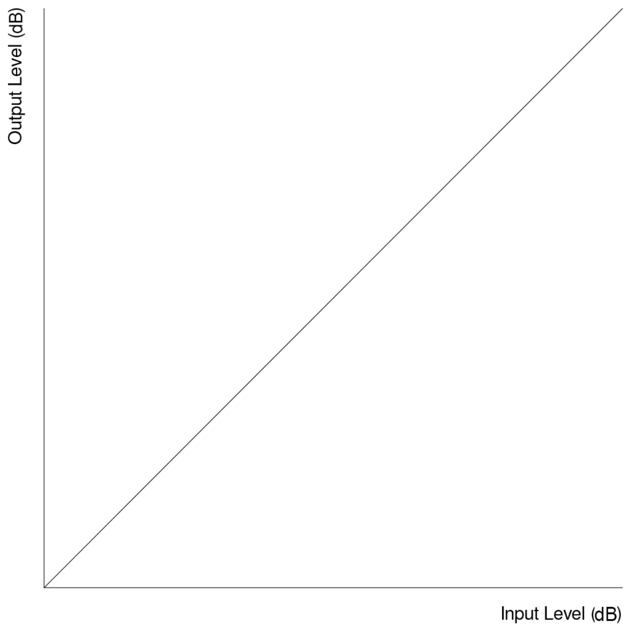

Figure 3-1: Theoretically perfect transfer function

In the real world, no amplifier is perfect and differences in the output compared to the input is viewed as distortion. As noted, non-linearity is introduced to the audio material when a component in the audio path (often an amplifier) distorts the input and adds additional artefacts that were not present in the signal. This distortion can occur when an audio signal is input to an amplifier and boosted to such a degree that it reaches the amplifier’s voltage limit, which is typically governed by its power supply voltage.

In most pieces of audio equipment, the power supply makes up a bipolar rail, which is a fixed positive and negative DC voltage with zero in the middle of the two rails. If a waveform peaks and exceeds these rails, it hits the voltage limit and the result at the output is a distorted waveform. This process is commonly referred to as clipping and gets its name from the shape a clipped sine wave takes on in an oscilloscope display. Complex program material such as music fluctuates considerably in amplitude over time, thus clipping is not static but instead momentary and level dependent, meaning the amount of clipping is dependent upon the varying level of the program material.

Loud audio signals are not the only type to create problems; there are also issues when working with quiet signals. All amplifiers have what is called a *noise floor* that is related to the resistance in the circuit, and if the input signal is too quiet, this noise can become audible at the output. This problem is referred to as a poor *signal-to-noise ratio* and can be an issue when working with quiet microphone levels that require significant amounts of gain. The difference between the noise floor and the maximum output of an amplifier is called its *dynamic range*. In practical terms, all audio equipment has an operating level that is somewhere between the noise floor and the maximum output before clipping, and this is considered the piece of equipment's useable range. A well-designed piece of equipment can have a maximum output level of approximately 20-24dBu and a noise floor as low as -100dBu. Under normal working conditions, this should provide the audio engineer with a significant amount of dynamic range to use. Thus clipped audio signals are typically a result of careless gain staging or implementation of a production technique.

There are differences in how amplifiers or specific components can clip, which in turn play a role in their sonic signature. Clipping behaviour falls into one of two main categories, hard clipping or soft clipping.

3.2 Hard Clipping

Hard clipping can be conceptualized by considering an audio system that upon reaching a set level acts like a brick wall and prevents any further increases in voltage. With a system like this the onset of clipping is not gradual and the signal at the output is considered to be hard clipped. Figure 3-2 is a transfer function plot like that used in Figure 1 but here when the input reaches a certain point it is completely levelled off. How this affects sound quality is easy to imagine if one considers a sine wave driven into an amplifier with a transfer function like that in Figure 2. The result is clipping to the top and bottom of the waveform and the sinusoidal shape transforms into something more like a square wave. Figure 3-3 shows the effect of hard clipping a sine wave. In this figure, the dotted blue line represents the portion of the sine wave that has been truncated by clipping and the solid red line represents the clipped signal.

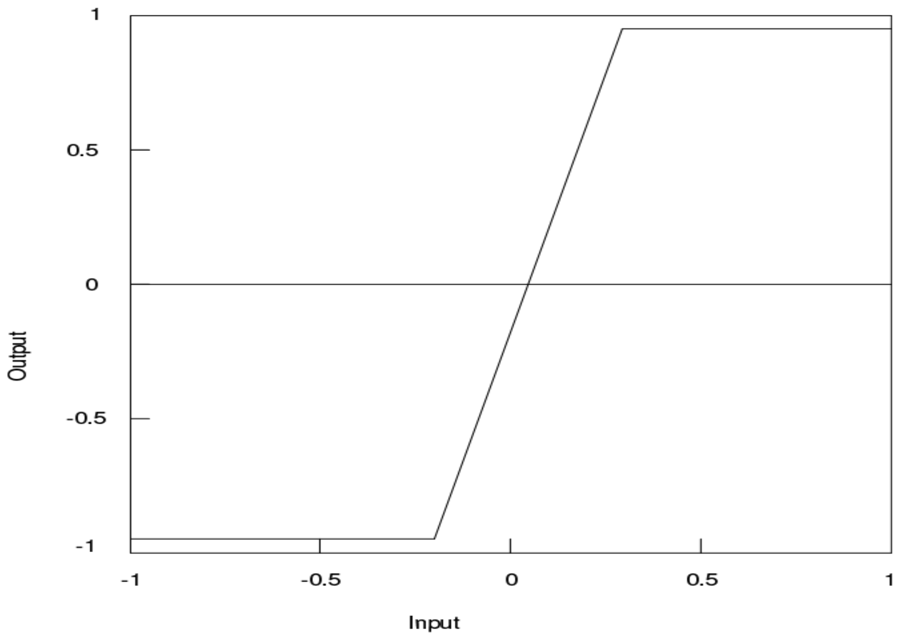

Figure 3-2: Hard clipped transfer function

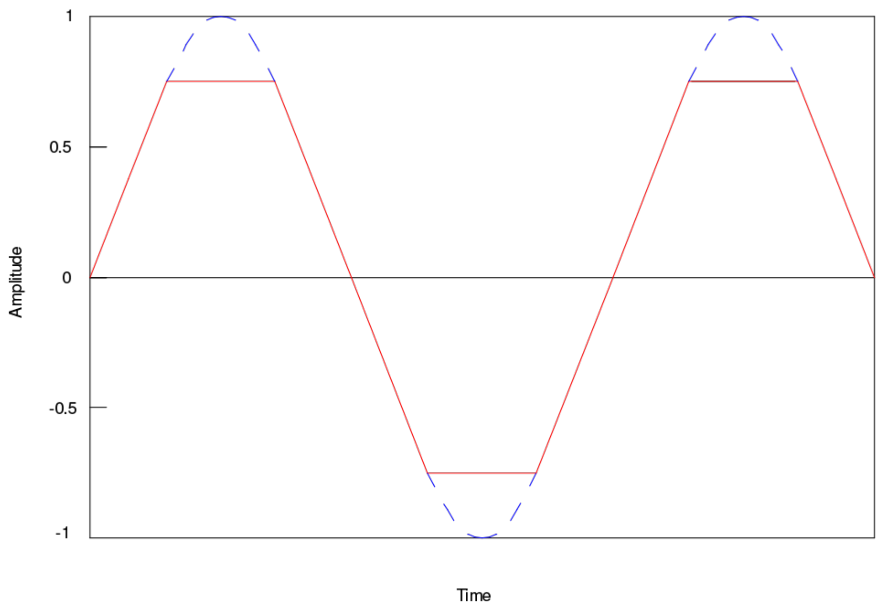

Figure 3-3: Hard clipped sine wave in the time domain

As the waveform begins to clip it not only starts to look like a square wave, it also starts to sound like one. The reason for this perceptual change is as the audio signal clips it generates additional harmonics, which were not present in the original signal. This behaviour can be seen in Figure 3-4 that shows a 1kHz sine driven into a clipped input stage of a dynamic range compressor with the compressor set for no gain reduction. The harmonics generated by this type of clipping are all integers of the fundamental pitch (ignoring the noise floor and low-level mains hum artefacts) and thus harmonically related. This harmonic relationship is why distortion of this kind is referred to as harmonic distortion.

Hard clipping typically occurs in diodes and op amps that have been pushed above a clipping threshold. Negative feedback is a design technique used to reduce distortion in solid-state amplifiers (and also valve amplifiers) but the downside of this approach is that when distortion occurs, the clipping is hard. This is due to clipping growing exponentially as the audio is fed back to the input (Elliot, 2006).

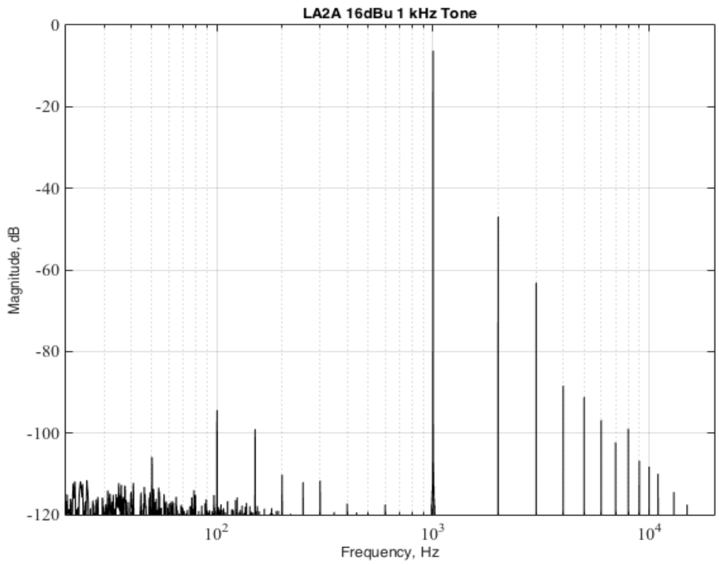

Figure 3-4: Clipped sine in the frequency domain

3.3 Soft Clipping

While hard clipping is an abrupt process with no gradual onset into clipping, there exists another distortion process that has a more measured transition. It is called soft clipping, and can be thought of as a process that squeezes down overloading peaks as opposed to chopping them off (Case, 2007, p. 94). Figure 3-5 shows the effect of soft clipping where the onset of clipping is more gradual than in Figure 3-2. As with hard clipping, the harmonics generated by soft clipping are integers of the fundamental. The introduction of these components is more gradual than hard clipping, however, and the amplitude of the higher order artefacts is typically of a much lower level than those generated by hard clipping.

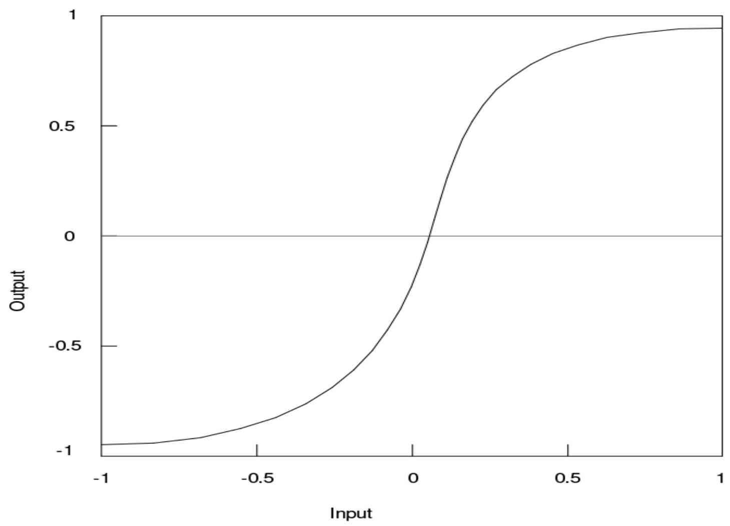

Figure 3-5: Soft clipped transfer function

Analogue components such as valves naturally exhibit a soft clipping character that is often claimed to be an integral part of their sound (Hamm, 1973). However transistor-based equipment can also be designed to clip softly and this is a common design approach in many guitar distortion pedals that aim to emulate the valve amp sound (Elliot, 2006).

Clipping is something to avoid if transparent recordings are the production goal, but music producers can use clipping as a creative effect to colour and sometimes radically transform audio signals. Out of the two forms of clipping discussed so far soft clipping is considered to be more musical sounding than hard clipping, in part due to the lower level of the higher order harmonics. Soft clipping is more suited to creative distortion in many music production styles than hard clipping due to its softer, subtler sound quality (Corey, 2010, p.114). A moderate amount of soft clipping is not necessarily destructive to the perceived quality of an audio signal. Enderby and Baracskai (2012, p.1) note that when audio is clipped to a moderate level, the non-linearity is perceived as a change in timbre as opposed to audible unwanted distortion. Hugh Robjohns from the music production publication Sound on Sound has this to say about clipping with regards to music production:

> As a technical engineer, I'd suggest that clipping is a fault condition that should always be avoided, and that peak control should always be achieved with a fast-acting limiter. However, as a mixing engineer, I know that the harmonic distortion produced through analogue clipping can sometimes be an artistically and musically useful tool in the right context. It produces a very different kind of sound, and at the end of the day, the intended sound is what would determine whether clipping, soft clipping, or limiting is the appropriate form of processing (Robjohns, 2013, para. 5).

The difference between a non-clipped, soft clipped and hard clipped sine can be seen in Figure 3-6. The clipping behaviour of a number of dynamic range compressors was tested for this thesis and the results can be found in chapters 5 and 6.

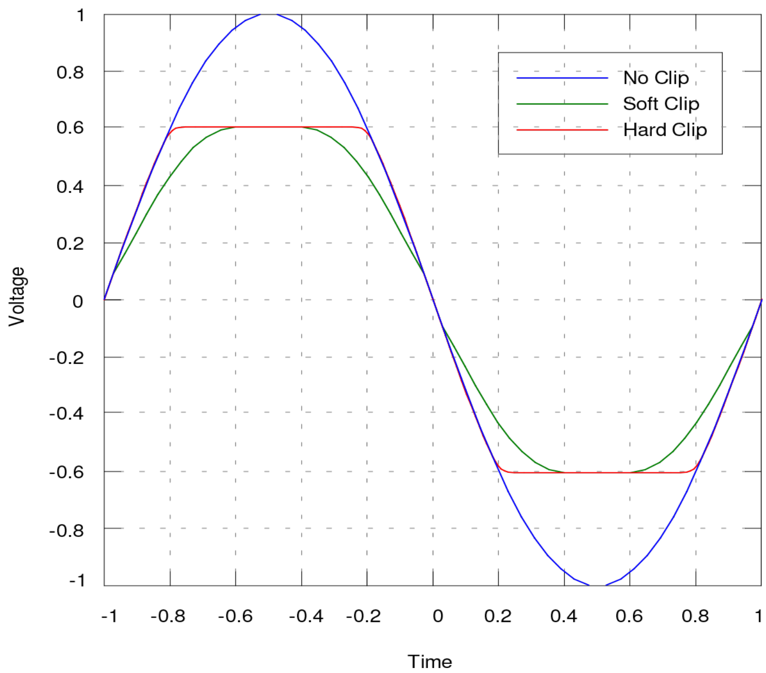

Figure 3-6: Non clipped, hard and soft clipped sine

3.4 Other Ways a Signal Can Be Clipped

Audio signals are usually bipolar and as a result can be clipped in two different ways. Firstly, if both sides of the waveform are clipped identically then the clipping is said to be symmetrical. This is the type of distortion we have seen in Figures 2 and 5. Clipping is not always symmetrical and it is considered to be asymmetrical if the positive and negative sides of the waveform are clipped differently. The character of distortion is connected with the symmetry of the clipped signal. Enderby and Baracskai (2012, p.1) identified that perceptual differences in distortion quality were mainly due to variations in harmonic content. Asymmetrical clipping is considered to be a more musical sounding form of clipping, due to asymmetrically clipped waveforms having a higher proportion of even harmonics than symmetrically clipped waveforms. In addition Brice (2001, p. 410) and Rutt (1984, p.9) state that lower gain valve guitar amps exhibit asymmetrical clipping and argue this is one of the reasons why guitar players prefer valve amplification. The differences between hard, soft and asymmetrical clipping can be seen in Figure 3-7. It should be noted that asymmetrical clipping occurs in both hard and soft clipping despite Figure 3-7 not illustrating this distinction.

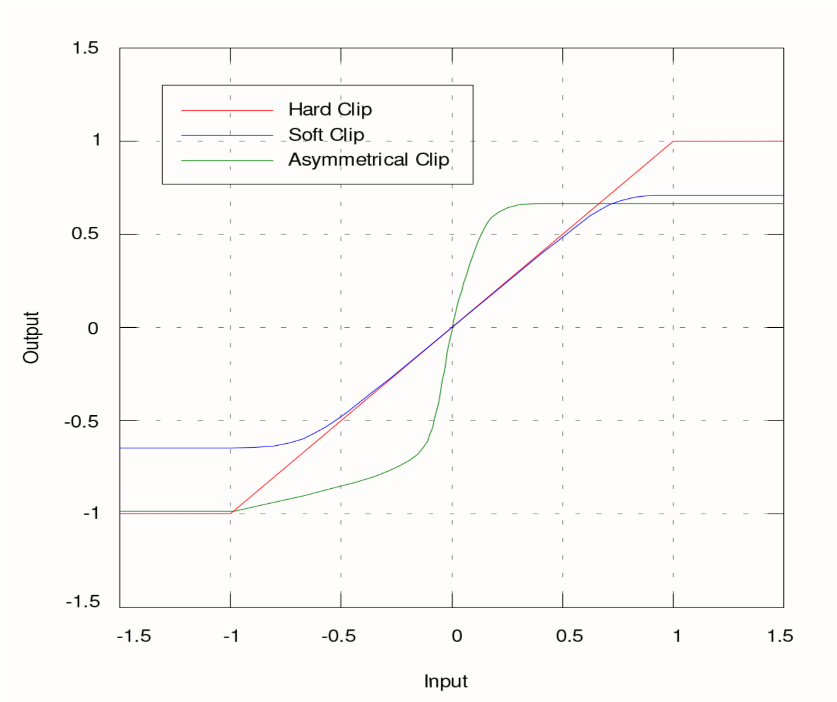

Figure 3-7: Hard, soft and asymmetrical clipping

As previously noted, if harmonic distortion is kept to a relatively low level then its perceptual effects are not always understood as distortion or considered a negative trait, even where transparency is the production goal. It was found in a study by Voisvillo (cited in Czerwinski et al. 2001, p.1014) that small amounts of harmonic distortion were perceived as subtle and pleasant changes to the timbre of audio material. Other forms of distortion can be introduced to audio however, and their effect is sometimes considered to be more objectionable than harmonic distortion. An example of distortion of this kind is called Intermodulation Distortion (IMD*)*, which consists of artefacts generated from the sum and difference of the input frequencies. For example, if two sine waves with the frequencies 100Hz and 250Hz are input to an audio system that exhibits IMD, it will generate additional harmonics at 350Hz (the sum) and 150Hz (the difference) at the output. These harmonic components are no longer integer multiples of the original audio signal and are instead inharmonically related. Additional artifacts can be produced from the sum and difference frequencies, which results in a significant number of inharmonic distortion components that grow quickly as the input is increased (Winer, 2013, p. 44).

IMD is always found in audio systems that have harmonic distortion and as noted by Whitlock (2008, p.285), it can be predicted to be three to four times greater than THD at a given input level. Rutt (1984) investigated intermodulation distortion in his research into valve and transistor guitar amplifiers and discovered that IMD was most common in transistor designs. IMD artefacts increased in both number and magnitude when the amplifier was driven into hard clipping. Furthermore, he demonstrated with listening tests that IMD artefacts were undesirable and had a negative impact on the sonic signature of transistor based guitar amplifiers (Rutt, 1984, p.11). The effects of IMD are shown in Figure 3-8, which illustrates the result of feeding a compressor with a test tone used for IMD measurements.

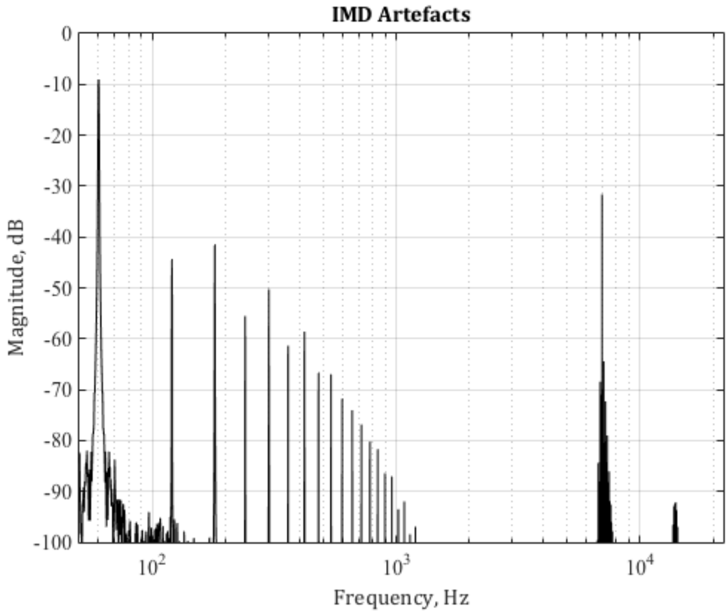

Figure 3-8: Harmonic content generated by IMD

3.5 Other Forms of Distortion

Harmonic and intermodulation distortions are the most commonly found forms of non-linearity in audio systems. However, for completeness it is worth considering other types of distortion that while not as significant can still play a role in a piece of audio equipment's sonic signature.

Crossover distortion is a form of distortion found in push-pull amplifiers and gets its name because it occurs when a signal crosses over from one transistor to another as it passes through zero volts. It occurs as a result of a short delay that is created when the signal passes between the transistors and audibly it produces a short click with high order harmonics. The design of some famous studio equipment was driven by a desire to avoid this form of distortion. Rupert Neve (1998) and (2013) has discussed on several occasions how crossover distortion was of significant concern when designing Neve amp circuits (including the 1073 microphone preamplifier) and that he would go to great lengths to avoid push-pull amplifier designs to prevent the unpleasant clicking artefact.

Transients are particularly susceptible to distortion artefacts. While some audio systems may be capable of handling steady state signals in a relatively linear fashion they may exhibit considerable non-linearity at the transient portion of audio signals; this form of non-linearity is known as transient distortion. Otala (1972, p.234) notes that transient distortion occurs in an amplifier system when the transient response of the power amplifier is slower than the preamplifier. Transient distortion can appear in program material that has many sudden amplitude variations and a steep rise time in the transient. In other words, if a piece of equipment suffers from transient distortion it will be more readily found in music sources that are percussive and of a rhythmic nature

3.6 Audio Equipment and Distortion

This sub-chapter first addresses equipment that was designed specifically for distortion-based tasks before moving onto equipment that was designed with linearity in mind, but has been exploited by engineers to work outside of the linear zone. This discussion is included to give some overall context to the use of distortion in the studio and how it relates to the non-linear DRC techniques discussed in Chapter 4. Additionally, the theory behind DRC and important aspects of compressor design is presented to highlight design traits that are necessary to understand for this study of sonic signatures.

The most obvious form of audio distortion is the guitar distortion pedal. There are many different designs of pedals available, ranging from subtle to extreme. A presentation for the Audio Engineering Society (AES) by Mauser (2010) traces the historical development of guitar distortion pedals and discovers the inspiration for these devices comes from an accident in the recording studio whereby a bass guitar recording for the song "Don't Worry" by Marty Robbins (1961) was distorted by a defective channel on a mixing console. This in turn led to the development of distortion pedals as engineers and musicians tried to recreate the effect.

The nature of clipping exhibited by a guitar pedal is usually dictated by design considerations such as component choice and the sequence of these components on the circuit board. Many overdrive pedals make use of diodes with the basic idea being to exceed the diode's clipping threshold and the user controls the amount of clipping via a gain control. Internally the designer of the pedal can affect the nature and amount of clipping by using multiple diodes, placing resistors between the diodes, filtering the signal and implementing feedback into the design (Keen, 2000). Work related to this thesis by Schneiderman and Sarisky (2009) discusses the effect of modifications to a Boss DS-1 distortion pedal and demonstrates how small alterations to the components used in the pedal can change its sonic signature.

Microphone preamplifiers (or mic pres) are an integral part of any tracking session and it is part of a recording engineer's job to decide which preamplifiers are the most appropriate for a given sound source. The tutorial video *Tracking Rock* (2011) featuring engineer Joe Barresi provides a revealing insight into how mic preamplifiers are used during a typical recording session. Barassi, like many professional producers and engineers, is shown to select a mic preamplifier because of its non-linear behaviour and the colouration effect it has on audio material. As with all audio devices that utilises an amplifier, the sonic signature of a mic preamplifier is significantly affected by the element providing gain change and the design, amplifier class and quality of components all play a role in a mic preamplifier's sonic fingerprint. Mic preamplifiers were originally designed to amplify low level microphone signals as transparently as possible and many popular designs worked hard to minimize distortion while offering high levels of gain. The Neve 1073 for example employs multiple amplifiers that have their gain adjusted over a relatively small range to reduce the possibility of driving one amp stage too hard and into clipping (Squire, 2014). However, many producers still push the 1073 into distortion during tracking.

Mixing consoles (or mixing desks) were designed primarily for mixing and summing audio signals as transparently as possible but this has not prevented engineers from using them to manipulate sonic signatures. It is common for specific mixing consoles to be selected for non-linear behaviour and for mix engineers to exploit this non-linearity by driving the channels and busses into saturation. When certain types of consoles begin to clip, they impart colouration on the program material, and affect the signals dynamic range by gently shaping the envelope. Rock producer Spike Stent is a proponent of driving his console into clipping. In an interview with Tingen (1999) Stent claims that one of the aspects he likes about his SSL console is the manner in which it clips audio signals during tracking and mixing. Steve Orchard, another engineer who favours desk saturation as a non-linear effect, claims the process adds character and energy to his drum sub mixes, particularly when room mic channels are driven hard on his mixing desk (Orchard, 2014).

Valves (or tubes) are still popular in the audio industry due to their non-linear behavior, despite more technologically advanced replacement components being available. Much of the subjectively pleasant distortion in valve guitar amplifiers comes from the way in which a valve clips, soft, asymmetrical and with low-level high order harmonics. Valves are typically made up of two or more electrodes that are enclosed in glass to form a vacuum, hence the name vacuum valve. The number of terminals in a valve can vary but a two-terminal valve is called a diode, a three-terminal a triode, and four and five terminal devices named a tetrode and pentode respectively. Triode valves are the most common in audio systems and inside this device electrons flow from the plate (terminal 1) to the cathode (terminal 2) and the flow of the electrons is controlled by a voltage applied to the grid (terminal 3). Amplification is linear around the idle operating voltage (better known as the bias voltage) but at large signal levels the output will begin to saturate and exhibit soft clipping behavior (Pakarinen & Yeh, 2009, p.87).

Valve audio systems are not inherently more coloured than solid-state systems and depending upon the design can in fact be more linear. In an IEEE article audio designer John Atwood states:

> Some of the differences in the audio qualities between tubes and transistors have to do with the inherent physical properties of the devices and with the circuit topologies and components used with each type of device. There is no way around it: linear [triode] vacuum tubes have lower overall distortion than bipolar transistors or FETs, and the distortion products are primarily lower-order…the clipping characteristic of tubes is actually not much softer than transistors, but feedback tends to 'square-up' the clipping. Thus, the heavy feedback in most solid-state designs gives them worse overload performance (Barbour, 1998, p. 26)

A well-designed valve audio system is as linear, if not more so, than some solid-state designs and a valve system that introduces noticeable distortion to an audio signal, often has been designed with non-linearity in mind. Nonetheless, valves can impart significant non-linearity to a signal when driven outside of their linear range and this non-linearity can be clearly heard in any overdriven valve guitar amplifier.

Valves are not exclusively used in guitar amplifiers and are also found in studio equipment ranging from equalizers such as the Pultec passive equalizer to compressors such as the Fairchild and the Teletronix LA2A (Barbour, 1998). Well-designed and maintained valve audio processors can be linear when used within their optimum working range but there is still the possibility for colouration if the user drives the input and the valves begin to clip. The soft clipping behavior of valve saturation allows these units to be used for subtle non-linearity and producers often describe this distortion using descriptors such as warm and fat, this will be discussed in Chapter 4. The Fairchild and the LA2A compressors were tested for their sonic signature in Chapter 5 and the author sought them out partly because of their implementation of valve technology.

Valve non-linearity is a case of component level non-linearity and another component that introduces colouration to audio signals is the transformer. Like valves, transformers are found in many equipment designs and they too impart a non-linearity that is deemed to have a positive effect on sonic signatures. A detailed discussion of how a transformer works is beyond the scope of this thesis but broadly speaking a transformer consists of two coils of wire with differing amounts of turns. These coils are called the primary and the secondary. The transformer gets its name because the ratio of primary to secondary turns creates a transformation of the voltage. The voltage applied at the primary and the voltage at the secondary is proportional to the ratio of turns, meaning if there are 500 turns at the primary and 1000 at the secondary then the ratio is 1:2. In practical terms this means that one volt applied at the primary results in two volts at the secondary (Squire, 2010). Additionally, the transformer affects impedance. For example, a 1:10 transformer will have a 1:100 impedance ratio. Note the ratio is the square of the turns (Farmelo & Hampton, 2010). This manipulation of voltage can have an effect on the audio signal that results in colouration. Equipment designer Bruce Rosenblit said the following with regard to transformer colouration, "the slow rise time of the output transformer causes a colouration that I would describe as a smoothing effect (…) the transformer is a nonlinear element that causes alterations of the signal in the time and frequency domain thereby altering the sound" (Barbour, 1998. p. 27).

Transformer colouration is more prevalent at lower frequencies and can be used to increase the low end in productions. Squire (2010) states that the reason for this increase in the low end is due to transformer distortion being inversely proportional to frequency, meaning that distortion increases as frequency gets lower. He notes that transformer distortion is level dependent thus more input level creates more distortion. It has been claimed that the transformers in pieces of equipment such as the 1176 compressor (Fuston, 2012) and some dbx compressors (Sound on Sound, 2009) play a role in the sonic character of these devices. A 1176 and a dbx165A compressor was tested for the work in Chapter 5.

3.7 Compressors

The main focus of this thesis is the use of DRC in music production and its non-linear sonic signatures. In chapters 5 and 6 several compressors process audio material and the audio is analyzed to investigate each compressor's sonic signature. Before that process could take place, it was necessary to examine DRC at a technical level to gain insight into common compressor circuit designs and how they impact sonic signatures. The following sub-chapters explore the information discovered during this research phase. It presents a comprehensive overview of the DRC process, investigates how it is implemented in analogue hardware, provides a detailed discussion on the differences in each compressor's design and notes how these differences affect the sonic signature of the compressors tested in chapters 5 and 6.

The first question to address in an investigation of DRC is why is it needed in the first place? To answer this question, we must consider dynamic range. As noted previously, electronic devices have a finite dynamic range that is measured between the noise floor at the bottom of the scale and the clipping limit at the top. The amount of dynamic range available in an audio system can vary from device to device, and an important role for the audio engineer is to set an appropriate level that maximizes the dynamic range of each piece of equipment and ensures the recording is not compromised with noise or clipping. The process is relatively straightforward if the recorded source is of consistent amplitude, but this is not always the case. During a recording session, the majority of audio material can vary considerably in level, often as a result of the nature of the sound source, musical dynamics and the playing technique of the performer. Thus, audio signals have their own inherent dynamic range that can be thought of as the lowest level to the loudest peak in the signal and this presents an additional challenge for the recording engineer. Not only do they have to maximize the varying amounts of dynamic range offered by the equipment they also have to pack the fluctuating dynamics of audio sources into this equipment. To make this process simpler engineers began to use DRC to automate these tasks.

DRC acts upon the dynamic range of audio sources by restricting the loudest peaks and (often but not always) turning up the signal's overall level thus reducing its dynamic range. The process is achieved by varying the gain of an amplifier, and this is why compressors are referred to as variable gain amplifiers. The design of the compressor can have a profound effect on the device and will be investigated later in this chapter.

A compressor in its most basic sense is a piece of audio equipment that makes loud levels quieter and quiet levels louder, essentially reducing the dynamic range. In simple terms, a compressor is an audio device in which the output signal does not increase as much as the input signal (Dove, 2008, p. 906). The effect of DRC is expressed in a transfer function plot such as that shown in Figure 3-9 where it can be seen the dynamic range of the signal is reduced progressively once the signal reaches a designated level. The level at which the dynamic range starts being compressed is called the threshold, and the amount by which it is compressed is called the ratio. The ratio of a compressor specifies a ratio between input and output and the audio signal is attenuated by this amount.

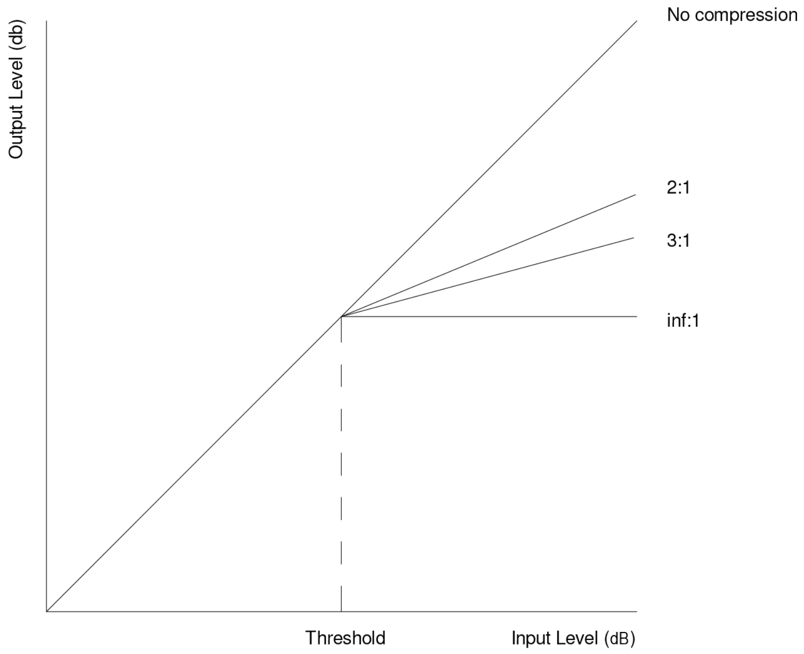

Figure 3-9: DRC transfer function

3.8 Compressor Design

Perhaps the simplest form of dynamic control to conceptualize is the limiter. As explained so far, all audio equipment has a finite amount of headroom and signals exceeding this limit result in clipping and non-linearity. While it was noted earlier that distortion can be used for creative music production effects it is not always desirable, and in many other areas of audio work, it can be a serious cause for concern. For example, over modulated radio signals can cause damage to transmitters and overloaded PA systems can cause hearing damage. In both these cases, the overloaded audio signal needs to be attenuated as quickly as possible with no further increase in level. This form of abrupt dynamic range control is called limiting, and from a design perspective, the simplest way to limit an audio signal is to use a pair of diodes that hard clip the signal (Dove, 2008, p. 898).

While this clipping method will accomplish fast level attenuation, it is only utilised as a basic form of dynamic range control for damage prevention and is rarely if ever used in studio production equipment. To limit the amount of non-linearity created during the attenuation process several other elements can be added to the circuit to control its performance and preserve the integrity of the input signal. The core components implemented in DRC design are shown in Figure 3-10. Here the audio is split as it enters the compressor and is sent to a side chain block consisting of several components including a level detection section. Controls have been added to specify the threshold level, the ratio amount and to control how quickly compression starts and stops being applied, often called the attack and release or the *time constants.* The behavior of the attack and release and the manner in which they work plays a significant role in the sonic signature of a compressor (Sonnox, 2007, p. 25).

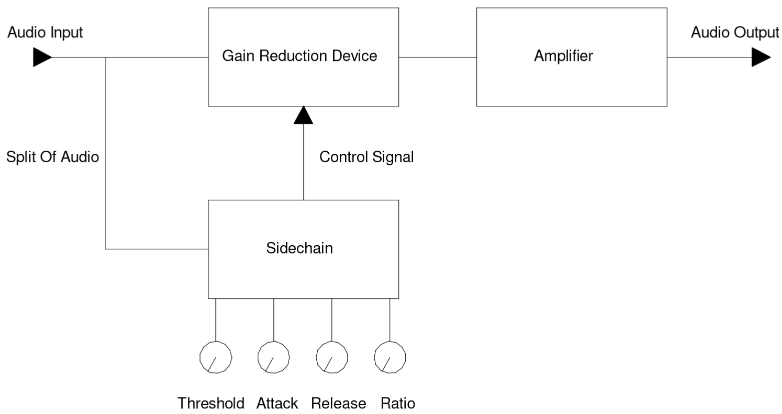

Figure 3-10: Block diagram of a simple DRC

3.8.1 Attack and release

Attack and release controls govern the speed at which the compressor starts and stops attenuating the signal once it exceeds the threshold. Depending upon the design of the compressor this behavior can be executed in a variety of ways but it is typical for attack and release to be implemented by a smoothing detector filter in the form of a variable resistor capacitor network (Giannoulis et al., 2012, p.401). In this design, the charging and discharging of capacitors govern the attack and release timings. In the attack portion, the capacitor is charged through a serial resistor while in the release period the capacitor is discharged (Talbot-Smith, 1999, p. 159). This design can be seen in Figure 3-11.

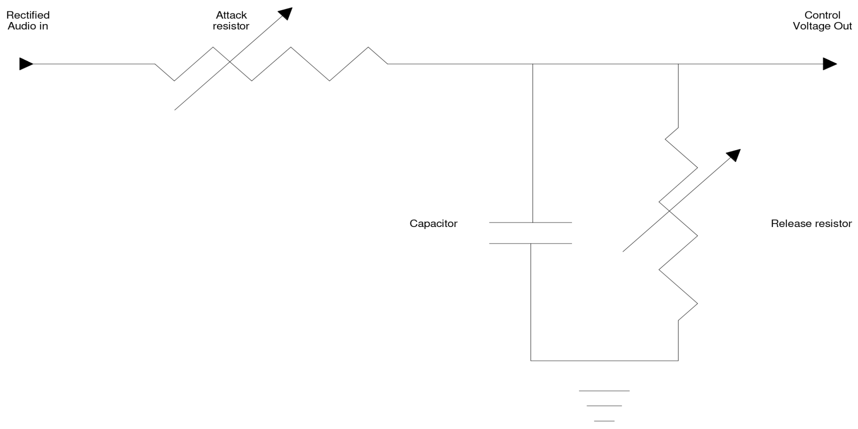

Figure 3-11: Attack and release capacitor resistor network

The attack and release controls are often called time constants, a term which is related to how capacitors charge and discharge. The product of resistance and capacitance dictates the length of time a capacitor takes to charge or discharge to a new voltage, and this is called a time constant. Of importance is the fact that after one time constant the change in voltage is 63% of the total voltage and after five time constants, the voltage change is 99% of its total. What this means for a compressor is if the attack time is 100ms then after one time constant (100ms) the gain change of the compressor will be at 63% of its final level. After 500ms (or five time constants) the gain change will be at 99% (Martin, 2011, p.85).

The timing law of a compressor is another important design consideration, and it can vary from model to model, often due to the designer's taste. Most compressors make use of either an exponential or a linear timing curve, and the former is the most popular.

With an exponential curve, the amount of time required to reach a steady state of compression remains the same regardless of the amount of attenuation. Figure 3-12 shows light, and heavy amounts of gain reduction applied using an exponential timing law. Note the initial rate of change is increased under heavy gain reduction but the time taken to reach a steady state of compression by both amounts of attenuation remains the same.

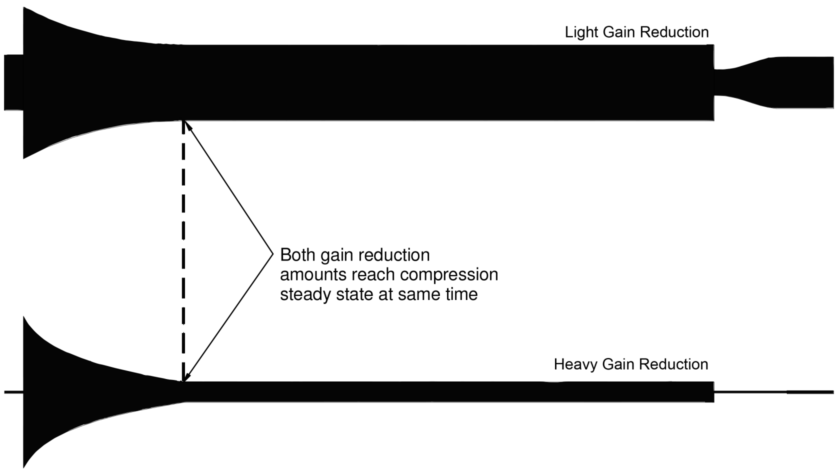

Figure 3-12: Exponential timing law. Light gain reduction on the top and heavy on the bottom

A linear timing law is shown in Figure 3-13 where it can be seen the compression curve changes with the amount of gain reduction. With larger amounts of gain reduction, the compressor takes longer to reach a steady state of compression. Thus, the attack, and release curves are dependent upon the amount of attenuation. The two plots used in figures 12 and 13 are based on diagrams by Sonnox (2007, pp.23-24).

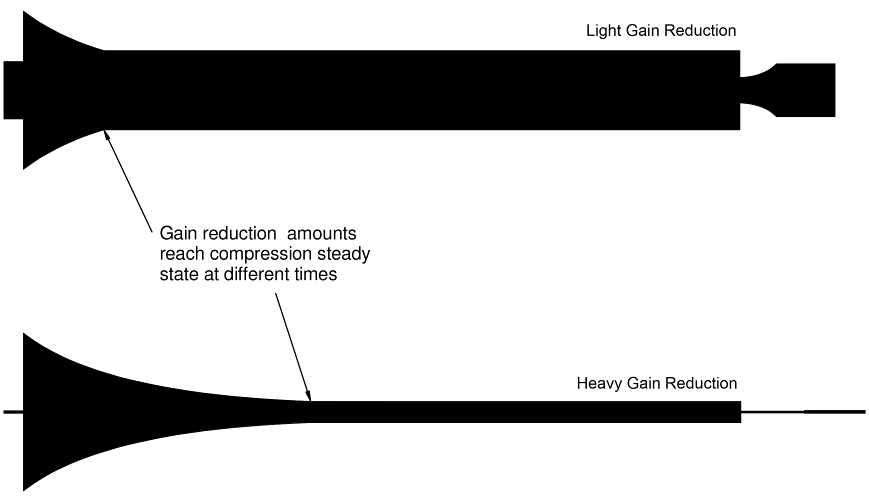

Figure 3-13: Linear timing law. Light gain reduction on the top and heavy on the bottom

The actual range of time constants offered by a compressor can vary considerably. Some units such as the LA2A allow the user no control over time constants while VCA based designs such as the dbx165A allow for a broad range of time constants that can be manually adjusted by the engineer. It is common for modern VCA styled compressors to have attack times ranging from 500 microseconds to 100 milliseconds and release times ranging from 100 milliseconds to 3 seconds and beyond.

Figures quoted by manufacturers in equipment specifications can provide the user with an indication of compression speeds, but they are often not an exact number. Some manufacturers will state speeds as the time it takes to reach a percentage of its final value (as previously noted this is 63% or one time constant). Not all manufacturers adhere to this convention and instead measure speeds in such a way that creates results which look impressive in technical specifications but virtually unachievable under real-world music production scenarios (Jeffs, Bohn, & Holden, 2005).

Compressors can be designed to exhibit a degree of program dependency. Meaning the speed of compression activity is changed by the acoustic properties of the material being compressed. One method is to vary a compressor’s attack and release time depending upon the duration of the signal exceeding the threshold. The motivation behind implementing this behavior is to minimize unwanted artefacts. With a transient, it is desirable to have a fast release to avoid over long periods of gain reduction but with steady state signals, the same release time will lead to audible pumping. Using a program dependent design will vary the release to prevent this issue (Berners & Abel, 2004). Program dependency is implemented in a hardware compressor by having a combination of time constants in the release section of the circuit. This approach is employed in the buss compressor found in SSL mixing consoles (Reiss & McPherson, 2014, p. 149). The effect of program dependency was investigated during testing of the 1176 compressor, and the results can be found in Chapter 6.

3.8.2 Non-linearity From Time Constant Settings

Reiss (2011) notes a number of audible artefacts are caused by the use of inappropriate time constant settings when using a compressor. These artefacts can be considered user errors rather than flaws in compressor design. They are:

1. Dropouts occurring in the audio where the compressor continues to attenuate the signal after it initially exceeds the threshold. This is due to an excessively long release time.
2. Overshoots occurring where the compressor fails to act quickly enough on the initial transient. This is due to a long attack time.
3. Loss of clarity from the use of fast attack times that attenuate the transient portion of the signal

4. "Pumping" (or fast modulation of the signal) as the compressor releases quickly after gain reduction
5. "Breathing" where the noise floor (or other aspects of the recording that aren't specifically the direct source) are raised in level by the time constants, producing audible amplitude movement of the extraneous audio.

In addition to the list above, distortion is introduced to a compressed signal by the use of fast attack and release times. Quick time constants lead the compressor to compress within the period of an audio signal and reshape the output. This process can be explained by considering a low frequency sine wave that has a period length that is longer than the time constant settings. Figure 3-14 shows a 50Hz sine that has been compressed with fast attack and release times, as can be seen, the time constants have changed the shape of the sine wave in a manner that is not unlike soft clipping. This waveshaping becomes more significant as the amount of gain reduction is increased and creates additional harmonic spectra that sound perceptually as audible harmonic distortion.

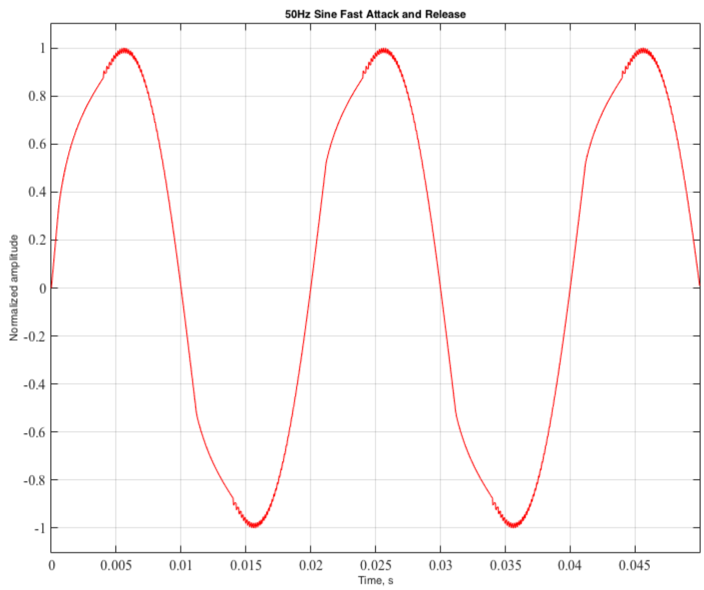

Figure 3-14: 50Hz sine, distorted with fast attack and release time

Both attack and release play a role in this non-linearity, the attack time has the most destructive effect on the signal due to it affecting the crest of each cycle. The release also reshapes the waveform, albeit in a less profound manner than the attack (Dove, 2008, p. 901). This point is not true for all compressors, particularly the 1176 where the release control plays a more significant role in time constant related non-linearity than the attack. This behaviour will be discussed in detail in

Chapter 6 in a test that measures non-linearity as a function of the attack and release.

As well as user errors, there are limitations in design that can generate non-linearity that is associated with release times. These flaws are related to modulations that occur in a waveform by the timing capacitor in the release stage charging and discharging and causing ripple distortion artefacts. This time constant behavior affects audio material by creating amplitude modulation that imparts many spurious harmonics to the output signal. This problem is one of the reasons why longer release times, relative to the attack, is implemented in hardware compressors (Self, 2010, p. 686).

As will be seen in Chapter 4, producers favour the use of fast time constants when using compressors such as the 1176 and the use of these fast time constants is at least in part to add distortion to program material. This non-linearity can be heard in the audio examples included with this thesis and seen in the analysis of audio material in chapters 5 and 6.

3.8.3 Side Chain and Level Detector

It was shown in Figure X that compressors have a block called the sidechain and the purpose of this component is to convert the audio signal into a control voltage to be sent to the gain reduction element for attenuation. The side chain section consists of several stages including: a rectifier; a comparator (this is the threshold value where the compressor compares the side chain signal to the input to decide whether attenuation is required); the time constants (also known as attack and release); and the ratio control that dictates the amount of the control voltage sent for gain reduction. The ratio and comparator section are thought of collectively as the gain computer section of a DRC (Berners, 2005). It is at the sidechain block where the control voltage is sent to the gain reduction component for variations in the audio signal to be attenuated in accordance with the ratio setting and time constant speeds.

Sidechains in a hardware compressor use either peak or RMS sensing to detect variations in the audio signal. RMS sensing is less obtrusive and more transparent than peak sensing and this is due to RMS sensing correlating better with the perception of loudness (Self, 2010, p. 501). Peak sensing derives its level estimate from a short-term peak in the signal, which is then smoothed using a filter and the manner in which this sensing occurs is affected by the time constant speed. For example, a short attack time encourages the capacitor to charge almost immediately while a longer attack time results in the capacitor charging slowly, which smooths out the peaks used in detection and results in an averaged value. True RMS detection (squaring the signal, finding an average of the square and then working out the average of the square) became accessible with the introduction of VCA based compressors (Dove, 2008, p. 908).

The choice of peak or RMS detection plays a role in the sonic signature of a compressor, which in turn dictates its suitability for a given audio source or production technique. Peak detection is used to reduce the peak level of an audio signal without negatively affecting its RMS level. This tactic is applied in mastering to increase the perceived level of audio mixes but has the tradeoff that the crest factor (a measure of RMS to peak) of the program material is reduced. RMS compression, on the other hand, can be used when an engineer wants a compressor to act transparently and reduce the overall dynamic range of the audio. The trade-off with this approach is that short-term peaks can get through uncompressed and create overshoots in the audio. It is uncommon for a hardware compressor to allow users to select between peak and RMS detection thus the detection method used by a compressor is an integral part of its sonic signature

3.8.4 Feed-Forward and Feed-Back Designs

There are two ways in which the sidechain block can derive the control voltage, either from the input signal (feed-forward) or from the output signal (feed-back). These two approaches are shown in Figure 3-15.

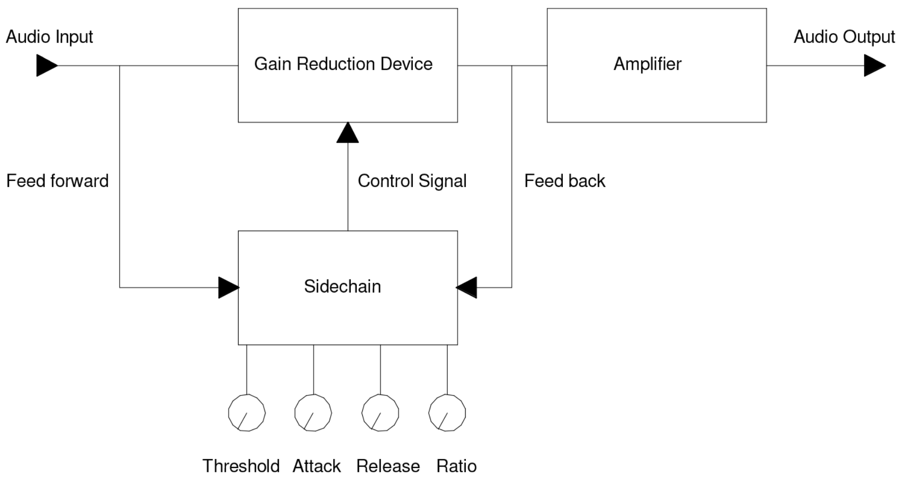

Figure 3-15: Feedforward and feedback compressor designs

Historically these methodologies were implemented because of limitations in the gain reduction component. For example, feed-back was utilised in FET compressors because it helped linearize unmusical aspects of the timing law (Self, 2010, p. 502) and feed-forward was only usable once VCAs had been adopted as the gain reduction element (Dove, 2008, p. 907).

Subjectively there are some audible differences between the two approaches, and feedback designs may have a more transparent sound. The reason for this difference is because the feedback style acts gentler on the audio due to it having a smaller dynamic range which is a result of the audio signal being fed back to the gain reduction stage multiple times. Feedback also has a technical advantage over feedforward because the sidechain can rectify potential artefacts in the gain stage. Feedback design has a technical limitation however that prevents the use of high ratios such as infinity to one (infinite negative feedback is not possible in this design) thus feedback is not used in compressors that necessitate true limiting behavior (Berners, 2006).

3.8.5 The Effect of Knee

The onset of compression can be applied abruptly or gradually. When compression ratio is applied immediately, it is referred to as hard knee compression. When compression ratio is applied gradually and increases proportionally with an increase in attenuation, it is called soft knee compression (Talbot-Smith, 1999, p. 150). Hard knee compression can sound aggressive while soft knee compression can sound subtle and transparent. True hard knee compression is difficult to implement in hardware and is usually only available in feed-forward VCA compressors.

The effect of knee on the sound quality of a compressor is often overlooked, but it is an integral part of a compressor's sonic signature. Soft knee curves in combination with slow acting program dependent timing laws can result in unobtrusive and musical sounding compression behavior while hard knee compressors in conjunction with fast time constants can create aggressive compression styles (White, 2000). The compressors tested in Chapter 5 featured a variety of units that make use of both hard and soft knee compression curves.

3.8.6 Gain Reduction Section

The majority of dynamic range compressors used in music production can be grouped into one of the following design styles, optical compressors, valve compressors, FET compressors and VCA compressors. There are other variations in compressor design, but these four are the most common. There are many differences in sonic signature between these designs, and one of the most popular compressor of each design is discussed in detail in the following sub-chapters.

At this juncture, one difference that can be addressed is the amount of available gain reduction. Table 1 shows figures identified by Ciletti, Hill and Wolf (2008) and should be used as an approximate guide. The amount of available gain reduction will have an effect on sonic signatures and particularly those that require the use of significant amounts of gain reduction for heavy attenuation of program material. The data in Table 1 should be kept in mind when looking at the amplitude statistics and plots in Chapter 5 where there is some variation in how much attenuation the compressors have applied, particularly on the room mic audio source.

Gain Reduction Style	**Amount of Gain Reduction**
Optical	12-25dB
Valve	25-30dB
FET	40-50dB
VCA	100dB

Table 3-1: Amount of gain reduction afforded by style

3.9 Design Analysis of Four Popular Compressors

The following sub chapters investigate the design of the best-known compressors from each of the categories listed above. These compressors were selected for discussion because they were found to be popular in the qualitative study detailed in Chapter 2. The reason for this discussion is to evaluate how compressor design affects a unit's sonic signature and highlight areas of design in which non-linearity can be imparted to audio material. The compressors detailed in these sub chapters were tested to assess their sonic signature and the results are discussed in Chapter 5.

3.9.1 The Teletronix LA2A

The basic design of the LA2A is relatively simple. First the input is passed through an input transformer and into the gain reduction section. Here it is routed through a gain control that drives the signal into some 12AX7 and 12BH7 valves before being sent to an output transformer. The LA2A is referred to as an opto-based compressor because it makes use of an electroluminescent panel that shines a light on a photo resistor to change its resistance. This works as a gain control because more light equals more resistance.

The component used to achieve gain reduction in the LA2A is called the T4 cell, it has unique qualities not achievable with similar photo resistors and is an integral part of the LA2A's sonic signature. One important aspect of the cell is its program dependency with attack, release, and ratio all varying depending upon the level of overshoot and the nature of the program material. For example, short transient material results in a fast release while steady state material results in a longer recovery period. Additionally, the LA2A has a two-stage release with the first stage taking place over the initial 40-80 milliseconds and the remainder lasting up to 2 seconds. This behavior, which is natural in the T4 cell and not engineered by design, allows the compressor to work transparently and unobtrusively and is the type of performance an engineer would want from a compressor if the production goal were to track the amplitude of program material musically. Ciletti et al. (2008) state this program dependent response is one of the reasons why it is difficult to make an optical compressor sound bad.

Aging components can change the sonic signature of analogue equipment over time, and the T4 cell is known to age in a manner that affects attack and release characteristics. Shanks (2003b) notes older cells have faster attack times and modulations (or ripples) during the release period. This modulation effect was observed by the author during testing of an LA2A and can be seen in some amplitude plots in Chapter 5.

The LA2A has the simplest user interface of the compressors discussed in this thesis. It features only an input and output control and a VU meter for monitoring levels. The output control is labeled gain and works as an output trim, and it has no effect on compression activity. The only control the user has to affect the nature of compression is the peak reduction control which when turned clockwise sends a split of the audio signal via an amplifier to the electroluminescent panel. The LA2A's simple interface is as attractive as its sonic signature, and this interface coupled with its program dependency creates a simple, efficient compressor to use during demanding tracking and mixing work.

The LA2A is considered a feedback compressor because the signal driving the sidechain is affected by a gain reduced signal (Universal Audio, 2000, p.11). The sidechain sees the effects of compression on the signal and adjusts the amount of gain reduction accordingly, and this adds to the transparent sound of the compressor.

3.9.2 The Fairchild

The Fairchild is a valve-based compressor and makes use of the *variable-mu* (also known as vari-mu) form of compression. This method works by sending a voltage through some 6386 valves, as the voltage increases beyond a certain point the flow of electrons between the grid and the gate of the valve is restricted thus attenuating the signal. The use of 6383 valves is necessary because other valves such as the 12AX7 will not allow for artefact free gain reduction. Compression happens directly on the audio path rather than being sent to a separate block and the reason for this design is to make the unit work with high current, low impedance and as little noise and hum as possible. However, the Fairchild does create non-linearity, and most of it occurs in the valves and the class-A amplifier due to mismatches in the two sides of the amp. This non-linearity is most significant under heavy use, and distortion artefacts increase when using higher levels of gain reduction.

The Fairchild has two large controls on the front panel. One is an input gain stepped in roughly 1dB increments that allow for a maximum of 18dB of gain. The second is a threshold control, which is continuously variable rather than stepped and turning it clockwise yields more gain reduction. The Fairchild offers more control over compression speed than the LA2A but is limited to 6 options via a stepped control labelled time constants. The positions range from 200 microseconds attack and 300 milliseconds release to 800 microseconds attack and 5 seconds release. There are also two program dependent options at positions five and six. Positions one to three are probably the most used for mixing. Murphy (2013) suggests that these three time constant positions are suited to the coloured and aggressive styles of compression used in rock music production, presumably because of their fast speeds. However, for mastering purposes (a role the Fairchild was originally designed for) the longer release settings and program dependent options may be more appropriate.

The Fairchild does not allow for direct user control over the ratio. Instead, it is changed as the input level increases. The fundamental working principle behind the Fairchild is to set the threshold with the threshold control and then drive it into higher ratios with the input. The factory default setting for this variable ratio is set at +2dBm and signals that exceed this level are compressed with ratios ranging from 1:1 to 20:1 (Fairchild Recording Equipment, 1959, p.4).

Two versions of the Fairchild were made, a mono 660 and a stereo 670 version that is essentially two 660 units stereo linked and controlled globally. The 670 version makes use of approximately 20 valves in various areas of its design including the rectifier, input and output stages (with several valves in parallel) and the side chain. There is no official record of the number of units made but it is estimated by the original designer Rein Narma, that only 40 units of the 660 and 40 units of the 670 were manufactured giving a total of approximately 80 Fairchild compressors (Narna, n.d.). Other manufacturers have developed many clones of the Fairchild (Analoguetube, n.d.; Mercury Recording Equipment, n.d.) and are reputed to have a similar sound quality. The sonic signature of a 670 Fairchild has been analyzed for this thesis and the results are presented in Chapter 5.

3.9.3 The 1176

The 1176 is called a FET compressor because it makes use of a BF245A field effect transistor (FET) in its gain reduction stage. The FET works as a voltage dependent resistor whose resistance is altered by a control voltage applied to its gate. In the 1176 gain reduction is created by shunting the audio signal to ground once it starts to exceed a threshold point that is fixed depending upon the ratio setting. The user can alter this threshold point by changing the ratio controls on the front panel, and it rises commensurate with ratio meaning that higher ratios shift the threshold point higher and vice versa. The knee of compression changes when the ratio is altered, and the curve of the knee gradually turns from soft to hard as the ratio is increased. Internally, the audio signal has to be kept to a small level to stop the FET from distorting during gain reduction activity. Consequently, the 1176 employs a powerful output amplifier to boost the output signal to make up for this initial loss in gain. The 1176 makes use of a 1:1 balanced output transformer that is used to reduce non-linearity and provide impedance matching.

Time constants are implemented by means of resistors and capacitors, which are placed in between some diodes and the gate of the FET. This design configuration alters the speed of gain reduction, which ranges from attack times between 200 and 800 microseconds and release times between 50 milliseconds and 1.1 seconds. In use the range of the attack time is very limited, and as noted by an engineer in the literature it goes from “fast to faster” (Warhead, 2011). The effect of the attack control is investigated in Chapter 6 and the results demonstrate how small the actual range of this control is in practical use. Similarly, the range of release times is limited due to the release pots taper applying most of its change over the first half of the dial.

One particularly important aspect of the 1176’s design is the bias control which is used to ensure the FET is never idling at 0 volts but instead at a voltage just under and slightly into conduction. This bias point results in a smoother transition into gain reduction because the FET is not audibly jumping in and out of attenuation as would be the case if biased exactly at 0 volts. An incorrectly set bias creates significant distortion during gain reduction. Therefore, care needs to be taken to ensure it is correctly calibrated and maintained. Variations in bias settings between 1176s may be one of the reasons for differences in their sonic signature and the degree of non-linearity they produce.

The input control of the 1176 increases the input and also sends more of the control voltage to the FET. As the input is driven, the control voltage is increased and creates more attenuation as it is raised towards the threshold point. The ratio is selected via switches on the front panel that can be depressed in some unorthodox configurations to yield unpredictable compression behavior. Nonetheless, this erratic performance has not stopped producers from using these ratio settings in their work. The most well-known of these settings is called the all-buttons mode and as noted later in Chapter 4 is popular to use on drum busses and room mics. What is not quite so well known is that this effect can be achieved by pressing the outer two ratio switches. This configuration works because the ratio resistors are stacked in series. As well as affecting the resistance in the ratio block, pressing multiple buttons alters the bias of the FET, and this affects the manner in which it applies gain reduction. This change in bias has the effect of increasing non-linearity and as discussed in Chapter 4 is one of the reasons why engineers use this mode for distortion and colouration effects.

Another interesting aspect of the all-buttons mode is the change in meter behavior. In all-buttons the meter shifts from 0VU and into the red portion of the VU meter, meaning all the way to the right. This change in the meter has no effect on the audio and occurs from a second FET in the circuit, which controls the VU meter, having its bias altered, and ballistics set out of calibration. The effect is purely cosmetic and plays no role in the sonic signature, but it may affect how a producer interfaces with the compressor. Engineer Michael Brauer has this to say about the sonic signature of the all-buttons mode and the behavior of the gain reduction meter:

> If you are familiar with the 1176, you basically have two knobs, an in and an out, and you have four buttons. With those four buttons, you can select your compression ratio. What you do is press them all in. Depending on the vintage of the unit, because you can't do this on some of the newer 1176s, hitting these four buttons makes it freak. The compressor needle, or indicator, will slam over to the right. Normally, whenever there is anything going on, the needle does the opposite. This looks really weird, but as long as it slams over this way, you know that it's working.

> This setting gives the sound a certain sense of urgency. It strains it. It's great for a vocal that needs extra urgency. Of course, you are going to be able to control the amount of strain in the voice by the input level. In the beginning, the needle may not move at all, so you have to keep bringing the gain up until the needle starts slamming over to the left. (Brauer, 2001, para. 6)

The 1176 was first introduced in 1967 and since then has been through a series of revisions. The following list has been adapted from the literature by Shanks (2011) and notes each revision and how they differ:

- Revision A was the first 1176 released in 1967
- Revision AB, also released in 1967 was an update that changed resistor values to reduce noise
- Revision B was released in 1968 and had some revisions to the preamplifier
- Revision C (1970) was the first of the blackface editions and introduced the 1176LN (low noise) circuit by Brad Plunkett
- Revision D (1973) featured a redesigned PCB (printed circuit board)
- Revisions E-H (1973 onwards) made a number of changes including a different transformer and output amplifier
- Current UA 1176s are claimed to be a faithful reissue of Revision D/E

There has been some debate amongst engineers over the differences in sonic signature between these revisions, and the author of this thesis made a number of measurements on two different revisions, the original Urei Revision D (aka Blackface edition) and the recent Universal Audio reissue. The results of these tests can be seen in Chapter 6 and heard in the audio material that supports this work. A second Revision D was tested as part of the research, and the results are featured in Chapter 5.

3.9.4 The dbx165A

The dbx165A is the most modern device discussed and tested in this thesis, and this recent design is reflected in its wider range of controls, and the LED that supplements the traditional VU meter. The dbx165A utilises a *voltage controlled amplifier* (VCA) in the circuit for gain reduction, and wholly original units make use of a dbx202 VCA.

A VCA is a special type of amplifier that has been designed specifically to attenuate or accentuate audio levels over a broad dynamic range. VCAs have their gain adjusted by an external DC voltage, and this is an important aspect of their design. Traditionally amplifiers are intended to increase the level of a signal (apply gain) and do not readily lend themselves to the attenuation of signals without potentially becoming unstable and oscillating. A VCA does not suffer from these issues. As well as having little noise, a VCA has a wide dynamic range and as shown in Table 2 yields large amounts of gain reduction, considerably more than the other compressors discussed so far.

The time constants on the dbx165A are expressed in a dB per millisecond scale for the attack and dB per second scale for the release. This specification means an attack time of 1dB per millisecond takes 5 milliseconds to change attenuation by an additional 5dB. While this is a precise way of stipulating attack times it is not intuitive, the maximum attack time figure of 400dB per millisecond is not going to mean a great deal during a vocal tracking session other than it is the fastest attack time offered by the unit.

The dbx165A offers a wide range of attack and release speeds, and this makes it more suitable for creative envelope shaping than the other compressors. The timing law of the time constants is logarithmic, and the VCA design creates precise and consistent compression speed and timing behavior. This performance is unlike the other compressors mentioned so far which have less precise control over their time constants. A short sub-chapter in Chapter 6 compares the range of attack times offered by the dbx165A and the 1176 and demonstrates the difference in envelope shaping possibilities afforded by these compressors. As well as manual time constants, the dbx165A has the option to use auto attack and release.

The most significant differences between the dbx165A and the other compressors mentioned here are it uses a feed-forward design, allows for higher ratios (up to infinity to one) and has an RMS detector that under certain circumstances provide it with a musical and transparent sound. In addition, it features a peak stop setting that can be employed to catch peaks that have been missed by the RMS detector. Therefore, the dbx165A has the option to work rapidly on the detection of program material if need be. The dbx165A allows the user to control the compression knee for settings between soft and hard, although instead of soft knee a gentler curve is called over easy on the dbx165A. Soft knee is implemented in this compressor by using an open-loop diode that creates variation in impedance over a range of voltages, and this means a higher voltage produces a sharper knee characteristic (That Corp, 2009).

Hicks (2009) states that VCA compressors are typically clean and less non-linear than other styles. While this applies mostly to the dbx165A it creates significant non-linearity when used aggressively with fast time constants. This was observed in tests made for this thesis using a dbx165A, and the results can be seen in Chapter 5. There is much online discussion regarding the distortion generated by the dbx165A and one commenter at the popular repforums states that "manual mode on these devices=distortion box" (Schwenkler, 2011). One other possible source of distortion was pointed out by Ciletti et al. (2008) when they claimed distortion could appear in VCAs because of mismatched transistors, adding that even modern designs suffer from spurious distortion with high input signals.

3.10 Conclusions on Non-Linearity and Equipment

This chapter investigated the core concepts concerning non-linearity and its relation to the production process. While clipping in a technical sense is to be avoided, it has not stopped music producers from exploiting it in their productions. They will use non-linearity to change the sonic signature of audio in subtle and not so subtle ways. Producers use a range of pieces of equipment for distortion, and this often includes processors or components that were originally designed for transparency. It was also noted that a dichotomy exists between audio manufacturers and audio producers. One group (designers) work to minimize non-linearity while the other group (producers) work to increase non-linearity. Modern trends in music production (mainly music genres such as rock, metal and some sub-genres of electronic dance music where the use of distortion on many elements is commonplace) have resulted in the lines between the two positions becoming blurred.

Non-linearity can be imparted onto audio material in a variety of forms, and soft clipping, hard clipping, and intermodulation distortion are the most pertinent to this study. The manner in which an audio signal is clipped can have an effect on the perceptual quality of non-linearity. Soft clipping and asymmetrical clipping (usually found together) are perceived as musical sounding while hard clipping and IMD are viewed as more objectionable.

This chapter also discussed several approaches to compressor design that can affect their sonic signature and discussed ways in which non-linearity can be imparted onto audio material when using specific time constant and ratio settings. It was noted the range of time constants, the timing law, feedback or feed-forward designs, the method implemented for sensing (peak or RMS) and the device used for gain reduction all play a significant role in the sonic signature of a dynamic range compressor.

Chapter 4 : How Producers Use Dynamic Range Compression in Music Production

4.1 Introduction to DRC Study

To get a detailed picture of how compression is used in music production, and in particular in rock and pop styles, this chapter explores how music production professionals use the process in their work. This study was conducted using the mixed methodology of grounded theory and content analysis detailed in Chapter 2. As well as gaining insight into the intentions behind the use of compression and settings, common techniques are extracted from this study and provide a basis for the settings and techniques tested in chapters 5 and 6.

This study addresses the following questions which are sub questions of research question one:

1. What are the commonly used compressors in music production?
2. What are the common reasons to apply compression in the music production process?
3. What are the popular sources to compress?
4. Are there any particular types of compressor that are common when compressing a specific source?
5. What descriptors are used to describe the sonic signature a compressor imparts onto a given sound source?
6. Are there any specific settings that are commonly used to achieve these sonic signatures? In particular, are there any patterns that can be observed relating to how time constant settings are used, principally with the 1176 compressor?

4.2 Popularity of Compressor and Gain Reduction Type

The first result worth examining is how frequently each compressor type was mentioned regarding a particular source. The findings can be seen in the stacked bar chart in Figure 4-1. In general, bass and vocal categories, the 1176 compressor is the most frequently mentioned, the only exception to this trend relates to discussions of kick and snare drums. In this category, the dbx 160 and 165 and more broadly other VCA style designs are the most frequently discussed. It is thought the longer attack times offered by these types of compressors (discussed in Chapter 3 and tested in Chapter 6) lend themselves better to transient shaping styles of compression. Furthermore, there is a clear connection with this style of compression, and the word analysis discussed later in this chapter, particularly regarding kick and snare drums. The other VCA category was predominately made up of iterations of the SSL channel compressor. The SSL channel compressor's popularity is in part due to it being found on the channel strip of the popular SSL mixing consoles.

The room mic category is interesting as it shows a preference for the 1176, albeit a relatively marginal preference. From the analysis of the qualitative descriptions of the sound quality, it appears the 1176 is the preferred choice for generating modulating compression (pumping) on room busses while the other compressors are utilised more for general colouration. The fast release times on the 1176 and the non-linear response from the all-buttons mode are most likely responsible for its popularity in this regard.

4.2.1 Grouping by Gain Reduction Type

Figure 4-2 shows a stacked bar chart with the compressors grouped together by gain reduction type and reveals a similar result. FET style gain reduction is the most frequently mentioned in the general, vocals and bass categories. Likewise, the snare and kick category shows that VCA types are the most popular choice when applying DRC to the membrane based drums (this includes kick, snare and toms, so essentially a drum that is not a bus or a cymbal). This popularity is possibly due to the wide range of attack times available on VCA compressors compared with the more limited range offered by the other compressor styles.

Looking at the drum bus/room category, the FET and valve styles of gain reduction are the most popular types with just over 31% and 27% respectively and accounting collectively for almost 60% of the compressors mentioned for this source. Analysis of the qualitative data suggests that modulating effects for the FETs and colouration for the Valve types are the primary reasons for the popularity of these compressors.

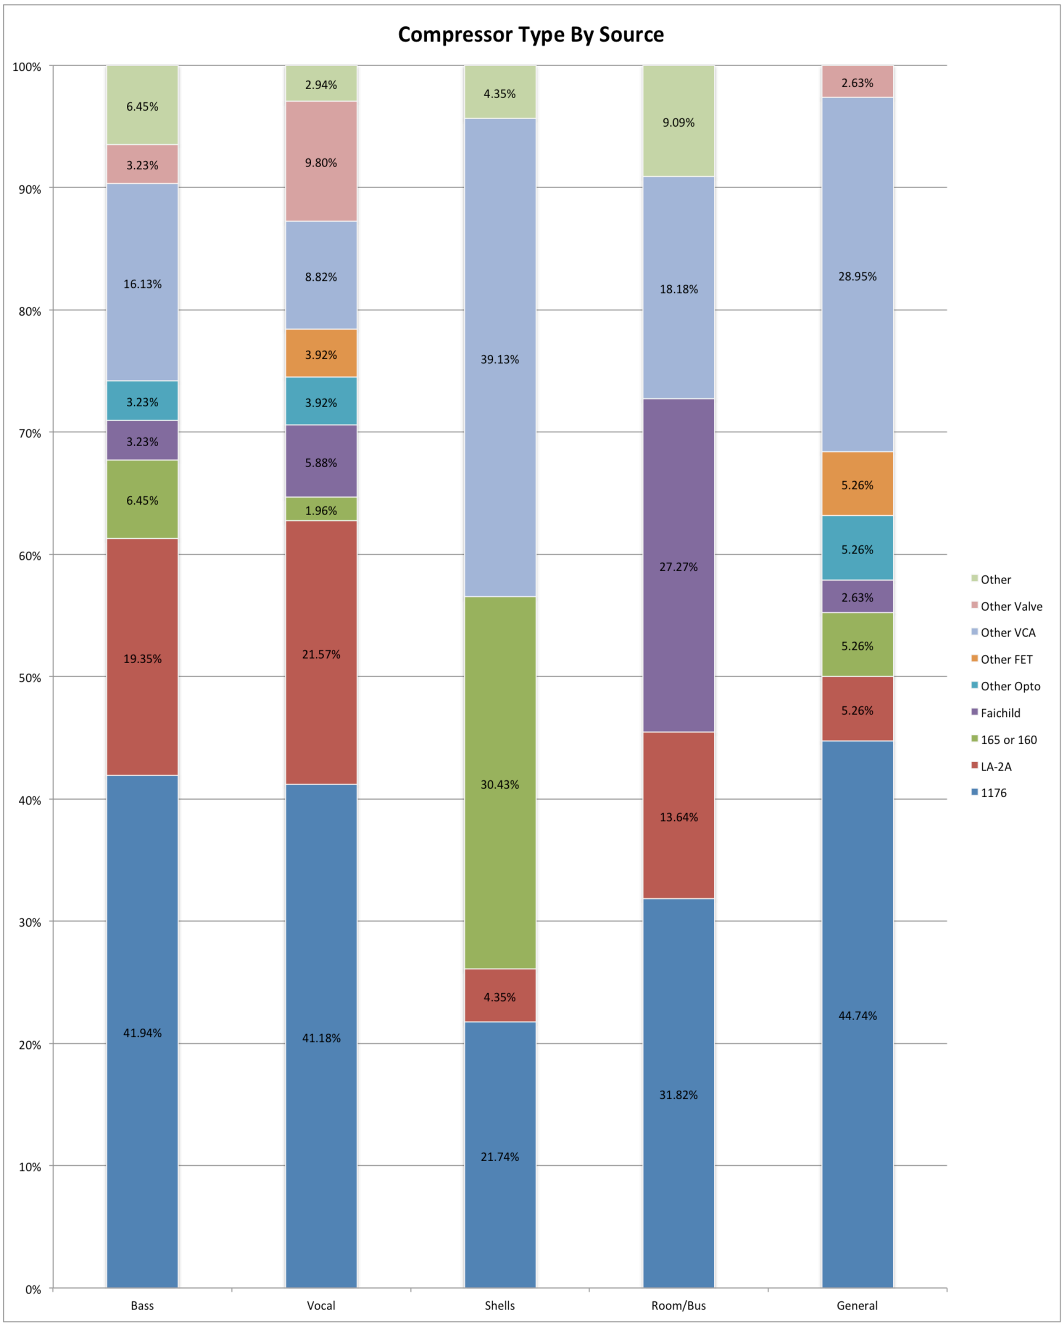

Figure 4-1: Popularity of compressor type

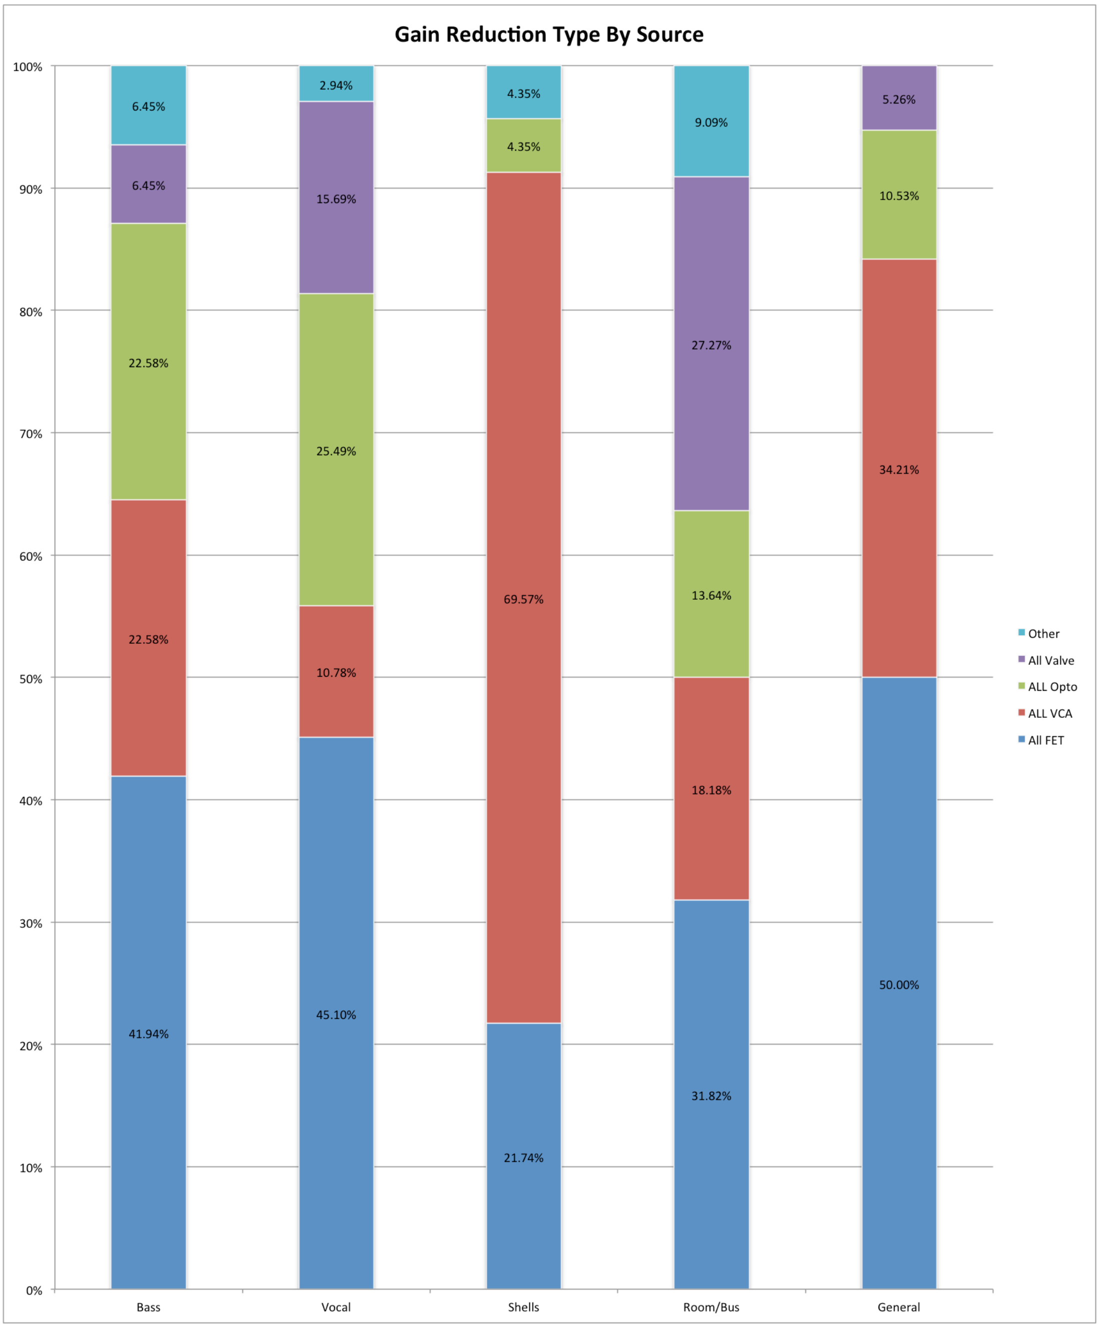

Figure 4-2: Popularity of gain reduction type

4.3 Popularity of Sources

The data was analysed to extract information relating to how frequently a given audio source was mentioned, data having been recorded when producers discussed the perceptual effects of compressing audio sources. Figure 4-3 shows the most common category is for vocals. When calculated as a percentage of the total data, vocal compression makes up 47% of the discussion. The other categories (excluding general) make up 35% of the total and are split relatively equally with bass having a slightly higher value.

Grouping the drum (shells and bus) categories together results in a different picture. With this grouping, the vocals are still the most popular (47%), but are now followed by all the drum sources (21%), the general category (18%) and finally the bass (14%). These sources are frequently discussed by audio engineers because signals of this nature can vary considerably in dynamic range. Nevertheless, as will be demonstrated later in this chapter, the restriction of levels is not the primary focus for engineers when using DRC in mixing.

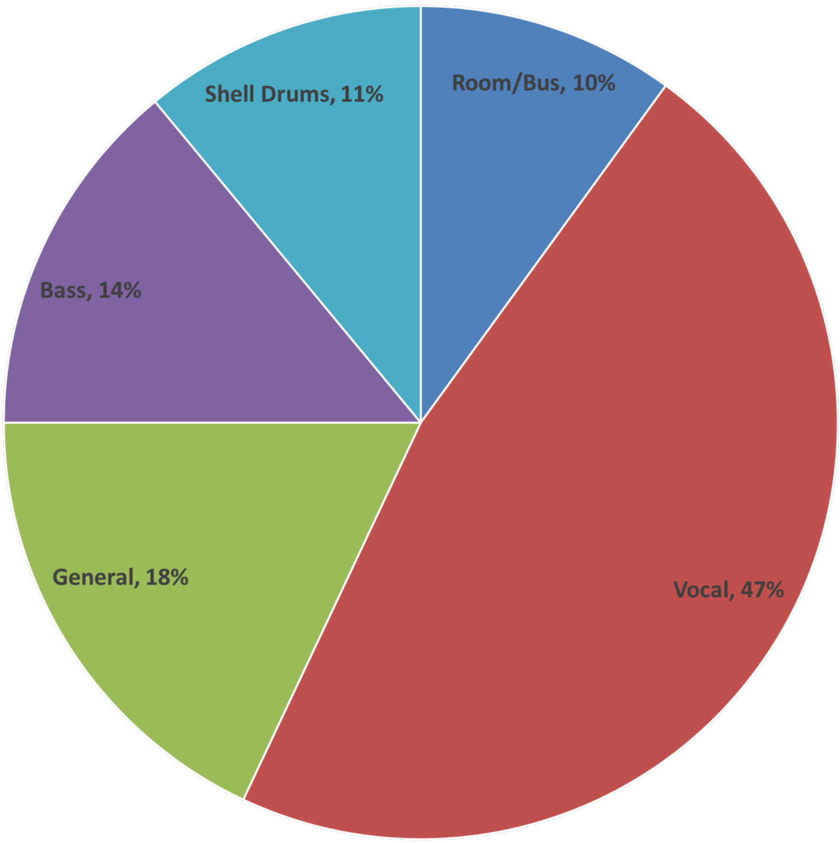

Figure 4-3: Popularity of sources

4.4 Results from the Grounded Theory Analysis

During data collection, the discussion was coded into categories and the results from the coding can be seen in Table 4-1 and the word cloud in Figure 4-4. The table shows the adjectives and phrases grouped together into the main inductive categories (compression production techniques) in the category column, the words extracted from the data are in the descriptors column and the inferred meaning of these descriptors is elucidated in the meaning column. The word cloud represents all the adjectives used to describe the effects of DRC and has the most common words represented in bolder text.

The table shows there is a trend amongst producers to favour the use of DRC for some form of colouration, both frequency based colouration and temporal modulation and transient shaping. The colouration general category is made up of descriptors that did not clearly relate to an obviously quantifiable acoustic property (frequency or temporal for example) but rather suggested a more general timbral change in the program material that could not be ascribed to a particular category. Due to the axial coding and core coding exercises that took place later in this grounded theory study the precise coding of these descriptors at this point was not overly important. However, it was felt prudent to group them into the general colouration category so as not to skew the data at this particular juncture in coding.

The chart in Figure 4-5 shows the initial categories that were created during the coding procedure. It shows the results for all compressors and sources and consists of six main categories with a smaller 1% category that includes all descriptors that were inferred as referring to transparency of some kind, this category is simply named linearity. 50% of the chart is made up of descriptors coded as relating to forms of colouration and distortion. Transient shaping and modulation make up two other categories and collectively they constitute 31% of the result. The dynamic control category includes descriptors relating to more typical dynamic range processing and makes up 18% of the data.

Category	Descriptors	Meaning
Colouration General	Character, colour, open it up, aggressive, edge, attitude, larger, hardness, excitement	General change to the sound, not clearly related to frequency, amplitude or distortion artefacts
Colouration Frequency	Presence, bite, beef, thick, warm	Change to the sound that can be related to frequency
Distortion	Grit, nasty, drive, grindy, rough, tough crunchy, distorted	Pertaining to the sound being distorted as a result of non-linear distortion artefacts
Modulation	Pulsating, sucking, bring out the ambience, pump, sucking, sucked it, pumping	Modulating change to the amplitude of a continuous sound source
Transient Shaping	Reshape, more body, total sustain, punch, snappy	Envelope reshaping of a transient
Dynamic Control	Tighten up, even, control, glue, level it, round, tight	General restricting of levels, not implying any noticeable effect

Table 4-1: Catagories, descriptors and inferred meaning

Figure 4-4: Word cloud of descriptors

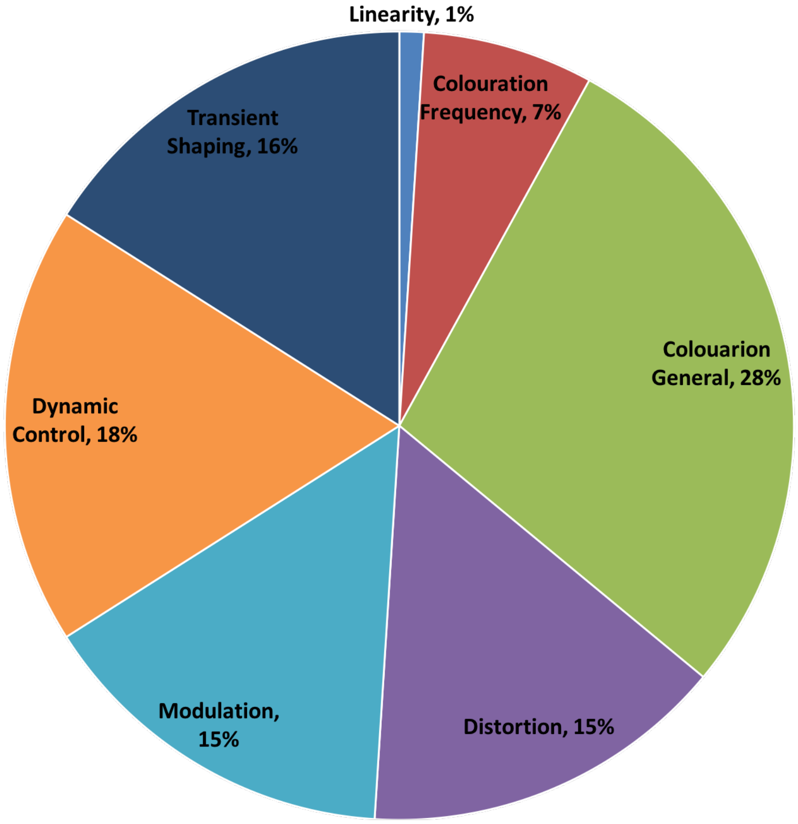

Figure 4-5: Initial categories from coding

4.4.1 Results from Axial Coding

Once the data was coded and these initial categories were created, axial coding was undertaken to explore relationships between categories and to establish any connections that could be used to link them together for further consolidation of the data. Axial coding is a process in the grounded theory methodology where concepts are linked into conceptual families to develop a coding paradigm. The results from this process can be seen in Figure 4-6.

Firstly, colouration general, colouration frequency, and distortion were linked into the category linear and non-linear colouration. This category included descriptors that related explicitly to colouration (e.g. presence and warm), alluded to colouration with descriptors meaning changes to the audio's character (e.g. edge and attitude) or were metaphors for distortion (e.g. grit and crunchy). The codes for modulation and transient shaping were put together into the collective code temporal change as they represented the same production approach but in different contexts. The transient shaping code originally related to descriptors associated with the envelope reshaping of kick, snare and tom sources (words such as punch and splat) while the modulation code related to words describing temporal modulation effects (sucky, sucking, bring out the ambience and so on). The connection between the two is subtle but obvious once explained, they are both reshaping temporal aspects of the audio, in the first instance it is the transient portion of the sound and in the second case it is the steady state portion.

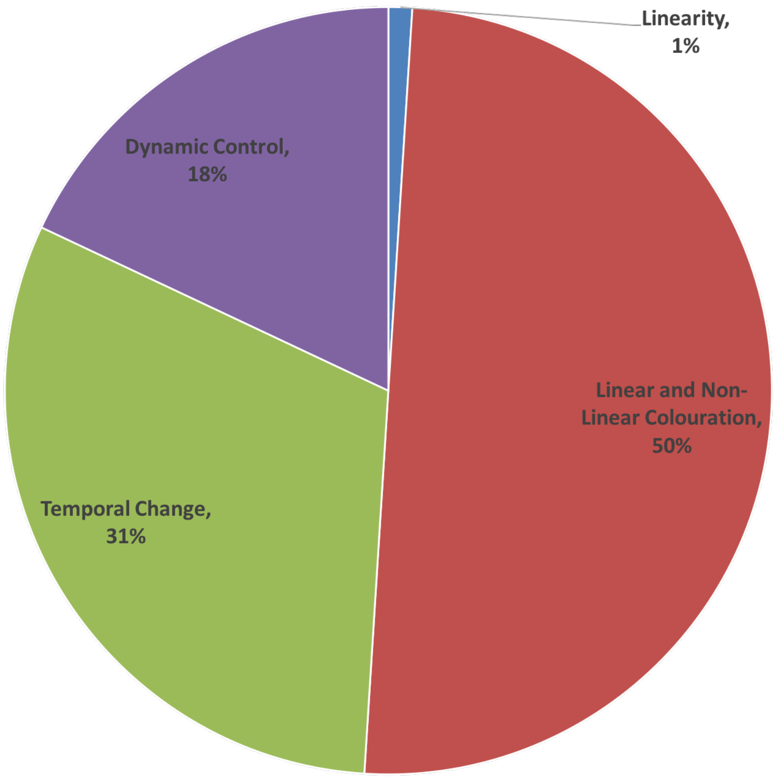

Figure 4-6: Categories from axial coding

4.4.2 Results from Selective Coding and Core Categories

Finally, the interconnected categories were rationalized further by selectively coding them into two main groups that would produce a discursive set of theoretical propositions (Strauss & Corbin, 1991). The two categories can be seen in Figure 4-7, and they have been coded as linear processing and linear and non-linear colouration. The linear processing group includes categories where the desired outcome of compression is transparency or subtle levelling of the signal, while the linear and non-linear colouration category has grouped together all categories that relate to the use of compression where its sonic signature is clearly audible, in other words the use of compression where the process is used as an effect and not as a transparent process. The results here show the dominance of the linear and non-linear colouration category with it constituting 81% of the final result. This finding allows for the proposition of the theory that producers and engineers are using compression in their work to impart non-linearity and colouration on audio sources.

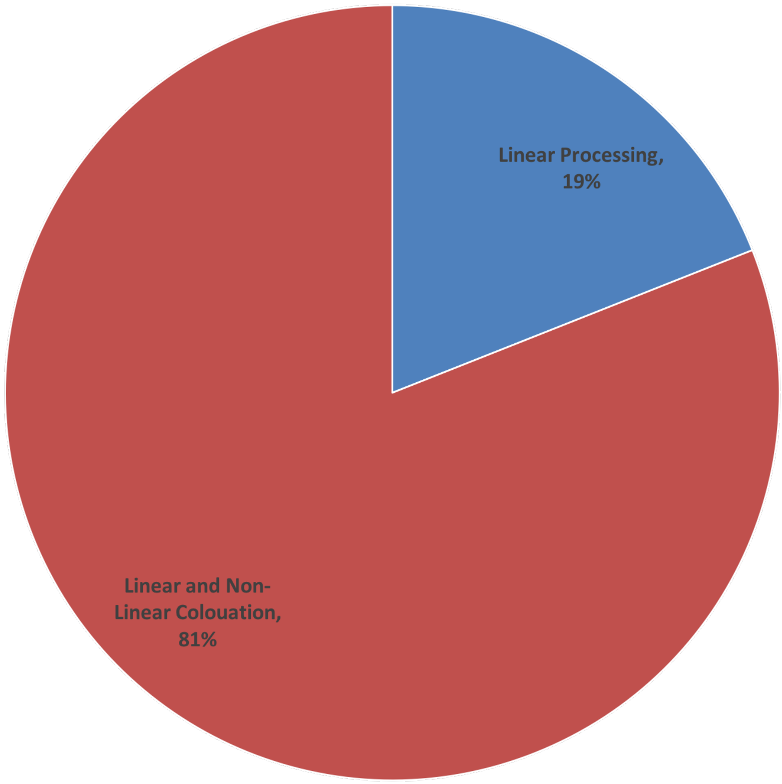

Figure 4-7: The core categories

4.5 Analysis by Source

In the previous sub-chapters, the analysis was carried out on results for all sources. The grouping together of all the sources was conducted because the results for some discrete sources was considered too small to be of use. However, analysing the results by source revealed some interesting nuances and a brief discussion of the most pertinent findings is detailed in this sub-chapter.

The results of the findings can be seen in Table 4-2 that shows there is a general preference among producers to use a compressor to colour the voice rather than for simple dynamic control. The colouration general category includes words such as aggression, edge and attitude thus there is a trend for producers to use DRC to add a sonic quality to the signal they deem to be aggressive in character. It is worth noting that the second most popular reason to apply DRC to the voice is to introduce distortion to the audio.

Category	Vocals	Bass	Rooms	Membranes
Colouration General	52.94%	0.00%	13.04%	11.11%
Colouration Frequency	5.88%	22.22%	0.00%	11.11%
Distortion	23.53%	22.22%	17.39%	0.00%
Modulation	0.00%	11.11%	39.13%	0.00%
Transient Shaping	0.00%	11.11%	13.04%	77.78%
Dynamic Control	17.65%	33.33%	17.39%	0.00%

Table 4-2: Analysis of data by source

The use of compression on bass sources is a little more varied, but there is a bias towards using compressors to colour the signal, typically to add a form of frequency related colouration. This frequency related colouration tends to refer to adding top end to the bass and is evidenced by producers using descriptors such as bite and presence. As well as frequency related colouration, the second most popular category is distortion. Thus, producers appear to be in fact describing the colouration effect of non-linearity. Nonetheless, regardless of how the colouration occurs, the results demonstrate there is a focus on colouration when using compressors on the bass.

Analysis of the texts pertaining to drums reveals a trend for producers to describe the effects of transient shaping. Table 4-2 shows this is their main motivation, with discussion of this kind accounting for 77.7% of the dialogue. Transient shaping techniques are one of the reasons producers favour VCA style compressors on drum sources. The wide range of attack times offered by these compressors allows the user more flexibility when using a compressor as an envelope shaper. On the other hand, when discussing material like compressed drum busses, overheads, and room mics, engineers focus on modulation effects the majority of the time. As will be noted in the next sub-chapter, FET-based compressors (and to an extent the Fairchild) are popular among mix engineers for this role due to their fast release, quick attack, and for example the non-linear artefacts from the 1176 in all-buttons mode.

4.6 Analysis and Overview of Quotes Extracted from the Grounded Theory Research

A number of relevant and informative quotes were extracted during the grounded theory and content analysis study and are presented in this sub-chapter to support the theories developed in the previous sections. The quotes have been analysed by the author to make clear any language or areas of discourse that may be vague or unknown to the reader, particularly if they are not fully acquainted with music production language.

4.6.1 Compression for General Colour and Temporal Changes

The most common theme emerging from analysis was that producers used DRC to impart non-linearity and colouration, particularly during mixing. Furthermore, it is clear that a number of producers are selecting specific compressors for their sound quality and using their colouration as a means to alter the timbre of the audio. Producer Michael Brauer elucidated this point when he stated, "Eighty percent of my compressors are used strictly for tone (…) the compressors are there to give attitude and tone and don't necessarily compress" (2008, para. 31). Brauer is implying the dynamic range control aspect of DRC is not of great importance to him, to what degree is not quite clear but he is certainly putting emphasis on the colouration aspect of DRC rather than dynamic range control.

Engineer Andy Smith is more explicit and notes in his work with Paul Simon he would not "generally use the compression for control of dynamics but more for a little bit of colour. Paul likes the colour of compression" (Smith, 2011, para. 15). This theme of adding colour to recordings with compression is common and there are a number of engineers who stated they ran audio signals through a compressor set to bypass to introduce this colour.

Pop music producer Jaycen Joshua notes, "colour is very important. Sometimes I run a track through a piece of gear, like a Fairchild, without touching the compressor, just to get it to add some colour to the sound" (Joshua, 2010, para. 14). Declan Gaffney makes a similar comment when detailing a recording chain for a U2 guitar part. He states that a coloured guitar tone he tracked for the part consisted of an electric guitar recording chain that "went through the Neve and then an LA2A, though it's not doing anything, it was just there for the sound" (Gaffney, 2009, para. 24).

On a similar theme is the producer Jack Joseph Puig, who also uses compression to alter the character of his recordings but in Puig's case he is using the compressor to change the temporal aspects of the audio. Puig appears to be aiming to shape the rhythmic components of the audio and alter the feel of the music with compression. Puig states the following in this regard:

> Compression is definitely the most musical tool that we have. I don't care for compression as a volume control. Using compression to alter feel and to affect performances has been done for a long time. For instance, it was common in the '80s to take a really fast compressor, like the Dbx 165, set it really aggressively, send a snare drum to it, and then gate that sound as tightly as possible. What you get is a 'kh–kh–kh–kh' sound, just an attack note that's extremely aggressive. You put that under a separate fader, which becomes your attack fader, and you feed that in with the regular snare to get the degree of attack that you want (Puig, 2007, para. 6).

In the second part of this quote Puig is describing shaping the envelope of a snare drum with a slow attack and fast release to pronounce the transient relative to the decay portion. The first two sentences of the quote are particularly interesting because they describe a subtler approach to the application of DRC. Puig explains this nuance more clearly.

> You choose the compressor for what you need. The way a Fairchild wheezes and moves is very different to the attack of a harder compressor, like an SSL. You can then place the front end of the note where you want, and make the track feel different. A Fairchild will give you a more legato effect, with more sustain, and you find the right place for it, in the way that musicians play. So you're actually playing almost like a musician (Puig, 2007, para. 7).

In this quote, Puig is stating the time constants of a compressor alter the articulation of a musical performance and the timing law plays a role in its sonic signature. What is particularly interesting is he notes engineers have to learn to appreciate the subtle differences in character between compressors and that once learned engineers can apply the sound of compression in a more musical manner.

Another reoccurring theme that emerged from the analysis is the use of heavy gain reduction to encourage the colouration process. Tom Lord-Alge notes that for him part of this effect comes from pushing compressors hard and into heavy gain reduction. He states:

> I often compress things really hard and I wish the meters would go round in a circle, because I'd love to see how much compression I'm really using. Each one of my 40 compressors has a different sound, and I generally use them as an effect (Lord-Alge, 2000, para. 18).

Note here how Lord-Alge, like Puig, is stating the compressors he uses each have their own unique sonic signature and he is applying compression as an effects processor and less as a control over dynamic range.

Producers are not only introducing colouration from non-linearity in the circuit they are also shaping audio through the subtle and sometimes not so subtle nuances of the compressor's time constants. These production techniques are examples of the categories highlighted previously in this chapter called non-linear colouration, transient shaping and modulation. This is evidence to suggest producers are using compressors in heavy states of gain reduction to add as much of the compressor's intrinsic sonic signature as possible. Tests made for this thesis noted the full nature of professional compressors' sonic signature is only fully apparent under heavier amounts of gain reduction. This point will be discussed in more detail in chapters 5 and 6.

4.6.2 Compression on Vocals

The quotes relating to vocal compression reveal once again that producers are attempting to add colouration to their vocal recordings. Mix engineer Bob Clearmountain expresses this point succinctly during a discussion of vocal mixing:

> I might compress for the sound, to get a certain kind of effect, but not to level it. I have some old UREI LA3A compressors which have been modified to reduce the noise, and I use a UREI 1176 sometimes, and I have these [*Empirical Labs*] Distressor units that are pretty cool for certain types of things (Clearmountain, 1999, para. 16).

Clearmountain adds a caveat and points out he achieves control over dynamics via desk automation. Furthermore, consideration should be given to the performance quality of the tracks Clearmountain is mixing (it is reasonable to presume they will be of a high standard) and to whether the audio had some compression added during the tracking stage or saturation was implemented from other production equipment. Both of these tactics are common in the professional recording industry during tracking.

Another trend that emerged out of the analysis was the desire to create an aggressive character when using vocal compression, sometimes expressed as the need to bring vocals to the front of the mix and create attitude. Mix Engineer Neil Avron has the following to say about vocal compression in this regard:

> The vocal compression is not so much about level adjusting as about the sound it gives, giving the vocal the right attitude so it sits better in the track. It may have been 4:1 or 6:1 perhaps even 8:1. Typically I'll hit the vocal pretty hard (Avron, 2008, para. 19).

It is worth noting that the tactic of adding an aggressive character on a vocal performance is usually achieved in part by the use of heavy amounts of gain reduction. It seems that this heavy-handed approach to DRC in rock and pop music is common and is part of the standard working procedure when using DRC for colouration and non-linearity

On a similar theme is the comment from producer Joe Chiccarelli who makes a comparable statement on vocal compression and non-linearity. He says:

> When it came to the mix, Jack wanted still more distortion on the vocals and more edge, so I overloaded an LA2A compressor, setting the output to 80. This meant that I was getting the distortion from the last tube stage of the compressor, which creates a really beautiful distortion (Chiccarelli, 2007, para. 23).

Again, we see an engineer deliberately using a compressor outside of its linear area and implementing sound processing that extends beyond the restriction of audio levels. Moreover, this effect is again achieved with considerable amounts of gain reduction in addition to the deliberate overdriving of the input and output stages of the circuit. Under certain circumstances, engineers are extracting as much non-linearity out of their equipment as they can.

4.6.3 Compression on Bass

From the analysis relating to bass instruments, it was shown in Table 3 that 44% of the time engineers discussed the sound quality of compression when it related to colouration of the audio signal. This was equally split between general change in the frequency spectrum and distorted non-linearity. Engineer Peter Mokran said the following about this style of compression:

> I treated it with a Urei 1176 blackface compressor, to get that second and third harmonic, so you could hear it a little bit more in the track. I usually don't hit the compressor too hard on bass. I don't even look at the meter: when it sounds right, I leave it. It's more a matter of tightening the bottom and getting the right amount of tone and grind out of it. The blackface 1176 has a very grindy character, more so than the new silver ones (Mokran, 2009, para. 17).

This quote is an example of an engineer discussing the distortion effects of the 1176 on bass material and noting explicitly how he will use compression in a manner that creates non-linearity. His comment about metering is somewhat vague, but he does express a sentiment similar to many others when he argues for ignoring the meter and focusing on sound. Additionally, he claims to hear audible differences in the sonic signature of the non-linearity. The differences between two 1176 compressors were tested in Chapter 6 and one of the tests makes use of a bass audio source processed with a setting known to introduce distortion to bass program material.

Engineer Dylan Dresdow adds the following about using fast time constants on bass guitar sources: "slam it aggressively with a faster release time than I normally use on a compressor. I make it really aggressive-sounding, giving as much life and character as I can: I'm shoving it over the edge" (Dresdow, 2009, para. 14).

Finally, professional recording engineer and audio mixer Andy Wallace promotes the use of DRC on bass guitars to impart non-linear sonic signatures. Regarding the mixing of bass guitars Wallace states "when I compress the bass, it tends to be for the sound rather than for the levels" (Wallace, 2014, para. 25). While Wallace is not entirely prescriptive about the nature of the sound quality he is seeking to introduce, he is clear that colouration is his primary focus. Like Clearmountain earlier, Wallace adds in this interview that he achieves much of his dynamic control via small adjustments in automation rather than using DRC.

4.6.4 Compression on Drum Sources

During the analysis of drum compression, it became clear that a significant portion of engineers were either discussing the use of DRC to shape transients of drum sources or to compress room mics and drum busses for colouration and modulation effects.

Mix Engineer Glen Ballard says the following relating to the dbx160X, which is a VCA style compressor similar to the dbx165A tested in chapters 5 and 6 of this thesis:

> I think the 160X is the best compressor for drums that I've found. You just give it a little bit on the kick and snare while you're recording and it brings them into focus. I don't think you can find anything better for kick and snare — they're amazing! You can really slam them too! I love the fact that they have a really quick recovery, and you can exaggerate it if you need to (Ballard, 2003, para. 24).

As noted previously the trend amongst engineers when compressing drum sources was shaping transients with a slow attack time to increase the perceived level of the transient relative to the decay or to a lesser extent round off the transient with fast attack times for a softer audio signal. In the final sentence, Ballard alludes to the timing behaviour of the dbx as the reason this compressor is an appropriate choice for this form of audio manipulation.

As well as transient sources, engineers favoured the use of DRC on steady state material to induce modulation effects. Using compressors to deliberately encourage audible pumping is not a new technique and was first used creatively by British engineer Joe Meek in the 1950s. Since then pumping and breathing compression effects have become commonplace, particularly when processing room mics and drum busses. Mix engineer Tom Elmhirst says the following on this style of compression effect:

> I'll have a compressor on buses 3-4, because compressing ambience gives you this lovely sucking sensation. I can pull it back a bit in the verses, so the sound is a bit closer. Using the desk allows me to easily ride the compressor (Elmhirst, 2011, para. 13).

Elmhirst is not only using DRC for a creative effect, he is utilising it to help create interest and development in the mix. Like Elmhirst, other engineers use compressed drum room ambience as a production effect. As noted by Buskin (2011) Hugh Padgham used it in his work with Phil Collins and helped popularise this sound in the 1980s.

Engineer Alan Moulder uses a similar approach to create modulation and movement in drum room mics but supplements the effect with distortion. He says the following comment regarding a drum tracking session for an industrial rock band:

> I was driving them pretty hard on the mic preamps — the PZMs were starting to distort and the 87s were on the edge of distortion. I compressed the 87s with 1176s with all of the buttons in so that they started to really pump, and again we instantly got a fantastic sound. It was really unique (Moulder, 2000, para. 32).

Note how Moulder is using every link in the recording chain to distort the audio signal, mic preamps are driven hard, mics are on the cusp of distorting from loud SPL levels and the 1176 compressor is set to its most non-linear ratio setting with time constants presumably adjusted for maximum dynamic modulation of the audio signal.

Moulder is discussing the production of drums for an extreme music genre, and it might appear this approach is only applicable to music of this kind. However, this is not the case and drums for Electric Light Orchestra's "Don't Bring Me Down" (1979) were captured with overloaded and driven audio equipment. The engineer who tracked the drums states the following:

> The overdubbed kit was in the bathroom, and I just stuck one mic up there and compressed it with a Urei 1176, overloaded. We did that on every album, but on *Discovery* we just recorded the bass drum, snare and toms in there for more control. Otherwise, it would have been too messy (Mack, 2013, para. 32).

The effect of the all-buttons mode on drum material and test tones was tested in Chapter 6 of this thesis where the non-linear effects of this setting can be clearly seen and heard.

4.7 Analysis of Common 1176 Time Constant Settings

As part of the grounded theory research, quotes were extracted from engineers discussing their choice of time constant settings when using the 1176. However, the amount of data generated by this exercise was small and the settings divulged by the engineers non-specific. Typically, they would use words like fast, moderate or slow or synonyms to describe their time constant settings. Despite not providing detailed data this information was useful as it helped build a hypothesis, that engineers were using attack times from the right side of the attack control (positions 1-4) and release times from the left side of the release control (positions 4-7). It is worth noting for the reader that the timing controls on the 1176 work counter clockwise which is the reverse of many other compressors.

It is likely the reason for this approach is to create an attack time long enough to allow as much of the transient portion of the sound through as possible. As discussed in the previous chapter attack times that are set too fast can exhibit a waveshaping characteristic, and it is typically recommended to avoid such behaviour. The results of the grounded theory analysis suggest that under certain circumstances (particular types of bass guitar and vocal production techniques) engineers will use this waveshaping behaviour for effects, but the trend when working with other sources suggests the opposite. However, the attack times offered by the 1176 are all fast and the control has a limited range that (as will be seen in Chapter 6) has little impact on non-linearity or the transient portion of audio material. Therefore, whatever the motivation for producers to set the attack in the 1-4 area it has little effect on the overall sound quality. Nonetheless, a trend was observed, and it was decided that more research was needed to substantiate this observation.

Additionally, the use of short release times can add non-linearity, and two techniques use this approach. Firstly, the use of short release times on continuous audio signals (particularly room mics) is implemented to increase the pumping characteristic identified earlier in this chapter. Secondly, the use of shorter release times is used to make colouration effects more obvious. This applies in particular to vocal sources where the compressor can work within syllables and words to aggressively shape the envelope. This effect is pointed out by engineer and designer Gregory Scott (2010) who argues that if a coloured vocal is the desired sonic signature then a fast release setting with the 1176 is one of the most effective ways to achieve it. The shortest release time available on the 1176 is 50 milliseconds and the longest 1.1 seconds. Experimentation for this thesis revealed that release times beyond position 4 yielded much longer times than settings between 1-4, which created little audible difference. The effect of release on non-linearity is significant, with distortion being most noticeable between positions 1-4 and dropping off sharply as the dial is turned beyond this point. Distortion as a function of release was measured using IMD and THD tones, and the results are discussed in Chapter 6.

4.7.1 Content Analysis Methodology

Focused studies were carried out to explore the 1176's timing settings. The studies consisted of:

- Content analysis on presets designed for Universal Audio Digital (UAD) plugins that emulate the compressors under study. These presets were designed by professional producers and created for the most recent UAD emulations
- Content analysis of video tutorials featuring professional engineers using the 1176
- Analysis of images and videos made during studio tests for this thesis. This material was generated in two professional recording studios and during the sessions experienced engineers were asked to set the 1176 compressor to impart certain sonic signatures on given program material. The author filmed and took images of the procedure for future analysis.

4.8 Content Analysis of 1176 Presets

A content analysis exercise was conducted on the artist presets available from the UAD 1176 Classic Limiter Plug-In Collection*.* These presets were developed by a number of professional recording engineers and thus were considered to be appropriate for this research.

The results of the analysis for the attack time settings can be seen in Table 4-3. Positions from 1-4 (the left side of the attack pot) are used for the majority of the time on all sources. This is particularly the case for the vocal, membranes and room/bus categories and also stands true when all categories are grouped together in the final overall column. While there is no single setting for each source, bass guitar settings bunch around 4 and 5, vocals 1 and 3, membranes 2 and 3 and room/bus 2 and 4. The general category includes all instrument sources that do not conform to the main types, the majority of the settings relate to acoustic and electric guitars with some all-purpose presets. The raw data in some of these categories is small, for example, only 15 vocal presets were collected, meaning the 6.67% in attack time 5 is made up of only one setting. It is also notable that ratios of 4:1 are used in combination with the slower attack times while the fast attack times are combined with the more extreme limiting ratios.

Attack Time	Bass	Vocal	Membranes	Room/Bus	General	Overall
1	9.09%	46.67%	7.14%	9.09%	25.00%	21.69%
2	9.09%	20.00%	42.86%	45.45%	15.63%	24.10%
3	9.09%	20.00%	35.71%	18.18%	31.25%	25.30%
4	45.45%	6.67%	7.14%	18.18%	9.38%	14.46%
5	27.27%	6.67%	7.14%	0.00%	6.25%	8.43%
6	0.00%	0.00%	0.00%	0.00%	9.38%	3.61%
7	0.00%	0.00%	0.00%	9.09%	3.13%	2.41%
1 to 4	72.73%	93.33%	92.86%	90.91%	81.25%	85.54%
5 to 7	27.27%	6.67%	7.14%	9.09%	18.75%	14.46%

Table 4-3: Popularity of 1176 attack times

Table 4-4 shows the results for the release times, which are more conclusive than the attack times. In the vast majority of instances, the release is set to positions 5 to 7. Again, the sample size for vocal presets means the 6.67% for release time 3 consists of only one setting. This configuration relates specifically to a backing vocal preset. It is not clear to the author why the engineer used this setting, and it is certainly not congruent with any other findings that were collected during research.

Release Time	Bass	Vocal	Membranes	Room/Bus	General	Overall
1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
3	0.00%	6.67%	0.00%	0.00%	0.00%	1.20%
4	0.00%	0.00%	0.00%	0.00%	3.13%	1.20%
5	18.18%	20.00%	7.14%	9.09%	31.25%	20.48%
6	18.18%	20.00%	50.00%	27.27%	25.00%	27.71%
7	63.64%	53.33%	42.86%	63.64%	40.63%	49.40%
1 to 4	0.00%	6.67%	0.00%	0.00%	3.13%	2.41%
5 to 7	100.00%	93.33%	100.00%	100.00%	96.88%	97.59%

Table 4-4: Popularity of 1176 release times

4.9 Analysis of Video Material

A number of video sources were analysed in order to extract information relating to time constant settings. Puremix is a company who have created a significant amount of video content relating to audio production, and two of their productions focus specifically on the 1176. One of the videos is a tutorial on using the 1176 (Pure Mix, 2012) and the second is an interview with a developer of 1176 clones (The Pure Mix Tutorials, 2011). The first presentation features a professional mix engineer (Fab DuPont) compressing a range of sources with the 1176 and discussing his approach. The first vocal setting he demonstrates aims to severely restrict the dynamic range of a vocal with the compressor aggressively working on the program material. To achieve this effect DuPont uses the fastest attack and release settings in conjunction with a 20:1 ratio, essentially utilizing the 1176 as a limiter. What is interesting to note is that DuPont also demonstrates an alternative setting that has an on-screen annotation stating that the release time is set slow. Upon inspection of the video, it can be seen the dial is at position 4, adding evidence to the hypothesis that release settings rarely extend lower than position 4 due to the slow speed of response beyond this point.

An interesting exception to the time constant settings appears in a segment of the video that discusses bass compression. Here DuPont initially sets the attack to position 1 and notes he is doing this to make the compression effect more transparent. He also sets the release to position 1 and his rationale for this release setting is again transparency. He notes that his motivation is to get the compressor to remain in gain reduction and avoid pumping. DuPont's views are similar to the hypothesis suggested earlier, that non-linearity and colouration are more profound when using the fastest time constant settings.

A second source from Puremix was analysed, and featured mix engineer Andrew Scheps (Pure Mix, 2014). In this tutorial, Scheps mixes a modern reggae track and uses the 1176 twice during the mix. Scheps works within the time constant ranges that have been identified thus far. He compresses some transient-heavy percussion with a slow attack and the release at position 5 and compresses a vocal with the attack set at 1 and the release at position 7. Scheps uses the all button mode for the vocal compression and describes the sound quality using the descriptors "spitty" and "aggressive".

The final source from Puremix is an interview with engineer Ben Lindell who demonstrates an 1176 clone he has built. During this presentation, Lindell shows the attack and release behaviour of the 1176. Of particular interest is his approach to using the release control. Apart from in one instance (to demonstrate the range of the control) Lindell never moves it from position 7. In fact, he comments that longer release times are not sonically interesting. The attack time is demonstrated in three positions: 1, 4 and 7. However, the only time Lindell sets it on the fastest position is to specifically showcase the heavy transient attenuation achievable with this position.

One final instructional video that was analysed was a presentation by Chris Lord-Alge for developer Waves audio, in which he demonstrates the process of mixing a rock production (Waves Audio, 2009a, 2009b). During vocal compression Lord-Alge uses the attack on position 3, the release on position 7 and a ratio of 4:1.

He also demonstrates the 1176 on drum sources. For the snare, the compressor attack is set to position 1 and the release to position 7. For both the kick drums used in the mix the compressors are set with the attack at position 3 and the release again at position 7. It is worth noting that despite the attack being at its slowest position for the snare Alge argues that what he likes about this setting is the increase in sustain and the reduction of the transient. This type of approach to rounding off the transient on drums is not commonly expressed in the literature but it is indeed a form of transient shaping that is well suited to the 1176 and makes up part of its sonic signature. This sound quality is discussed in more detail in Chapter 6.

The next section of this presentation focuses on overheads and Lord-Alge uses the 1176 with the release again at position 7 and the attack at 3. From listening to this section, he is using the compressor to increase the sustain of the drums and reduce the dynamic range. He expresses his intent is to use the compressor for gentle bus compression and cohesion. In this case, Lord-Alge is using a 4:1 ratio with a small amount of gain reduction. He then discusses his approach to room mics and demonstrates some compression on a mono room. Here Lord-Alge uses the all-buttons mode in conjunction with the attack set at 3 and the release set yet again at 7. He uses heavy amounts of gain reduction in this setting.

What is perhaps the most interesting piece of information to be extracted is that Lord-Alge uses the same time constant setting for all the sources bar the snare. However, even with the snare drum he still keeps the release fixed at position 7. The use of time constant settings by Lord-Alge gives more support to the hypothesis that producers work within a limited range of attack and release positions when using the 1176.

4.10 Studio Engineer Settings

Videos made by the author during studio sessions for this thesis were analysed retrospectively to investigate how the studio engineers used the 1176. One of these sessions took place in a professional commercial studio in London during late February 2013. During the session, the in-house engineer, who had much experience working with the compressors on a daily basis on professional music productions, was asked to set up the compressors to impart distinct sonic signatures to the signal. While he was setting up the compressors he was filmed adjusting his preferred settings and during the review process these videos and images were examined and his settings were noted. While working with the vocals (the mandate was to achieve an aggressive in your face vocal sound), the engineer sets the release to its fastest position and alters the attack until he is happy with the sound. The attack time is set to position 6 but in combination with a 12:1 ratio, giving support to the idea that faster attack times on vocals are used in combination with the limiting ratios (12:1 and 20:1). What is revealing is that in all instances, aggressive vocals, coloured grinding bass and heavily coloured drum rooms, the engineer (like Lord-Alge) sets the release to the fastest setting. Attack times for the bass and drum rooms are set to positions 5 and 4 respectively.

A second exercise was carried out with an experienced engineer at another professional studio in the UK. The engineer was given the same stipulations as the previous session and was videoed setting the compressors until the required sound characteristics were achieved. His method of adjusting the compressor was the same as the previous engineers, first setting the release to 7 and then adjusting the attack pot around the left-hand side of the attack dial until satisfied with the sound quality. For all sources, the attack is set to position 3 and the release at 7. The 4:1 ratio is used for both the vocal and bass while the all-buttons mode is implemented for the mono room mic. The videos for both exercises can be found in the video folder in the associated memory stick.

4.11 Conclusions on DRC Study

The results found in this chapter reveal a bias towards the use of DRC as a colouration effect. This is evidenced by the popularity of the non-linear and linear colouration category that accounted for 81% of the discourse in the grounded theory study. One possible limitation of this work is that it may be that producers only explicitly describe the sound quality of compression when it deviates from the norm, meaning they are not investing in descriptions of compression when it is simply reducing the dynamic range of a recording transparently. The author of this report feels that this is a valid limitation of the current study, and as an area of further research would like to expand upon this work by conducting a series of interviews with professional producers and mix engineers on their motivation behind the use of compression.

Certain design characteristics limit the number of suitable sound sources for a particular compressor, thus affecting how frequently it is utilised. An engineer seeking to shape the envelope of a snare drum, for example, will typically overlook a compressor with no user controllable time constants in favour of another device with user definable attack and release times. Therefore, the reader is reminded not to simply deduce that sound quality alone is the only determinate over which compressor to use in a popular music production scenario.

From analysing the quotes, it emerged that engineers were very much in favour of working with DRC to impart distortion, colouring audio, reshaping the envelope of transients, modulating the steady state portion of drum audio and selecting different compressors for their sonic signature. Producers reported that they would use compression at the mix stage for these types of techniques and achieve dynamic control by other means such as automation. Much of the audio received by professionals for mixing will have had some DRC applied during tracking, and this must be considered when interpreting the results. Radical fluctuations in level may have been evened out at the tracking stage, which means the mixer can start to explore other reasons to apply DRC.

The results from the various forms of content analysis carried out on time constant settings show that 1176 time constants are typically set with attack times between 1 and 4, and the release times between 5 and 7. The exact positions can vary, but this trend emerged across all the content analysed. The results for the release times in the presets are particularly conclusive with the timings for all sources set between positions 5-7.

The video analysis found that engineers kept the release control at position 7 for a significant portion of the time. There was more variation in the attack time settings, however, due to the limited range offered by the attack control, this does not have a significant effect on the sonic signature. The studio engineers helped to validate the findings from the other forms of content analysis. A limitation of this work was the small number of engineers used in the exercise. It is hoped that in future work this study can be carried out with a larger number of experienced engineers, and such comprehensive results will give additional support to the hypothesis proposed in this chapter.

Finally, it is worth keeping in mind that the results presented in this chapter are culturally specific to the producers in this study. As previously mentioned the interviews featured a diverse range of producers but the majority came from a rock and pop music background. Therefore, it is unclear how the results presented here relate to music producers working in other musical genres. Furthermore, from a social constructionist (Burr, 1995, pp. 1-4) point of view, the way of using compression in these styles of music has been constructed through interaction between music producers. Thus, what they regard as their working procedure is a product of social processes and interactions and may not represent the working procedure of music producers as a whole.

Chapter 5 : Analysis of Four Hardware Compressors

5.1 Introduction to Analysis of Four Popular Compressors

The following chapter discusses audio measurements made on four dynamic range compressor styles found to be popular in Chapter 4. The compressors are the Urei 1176 Revision D, Teletronix LA2A, Fairchild 670 and dbx165A. The measurements were made to get an overview of each compressor's sonic signature and compare and contrast the results. Measurements were made using test tones and complex program material. Test tones are analysed using THD measurements to assess distortion and tone bursts to determine the timing laws. Complex program material is analysed using a mixture of spectrum analysis, critical listening and examination of audio features. Analysing the compressors in this way can help us get an understanding of their sonic signature at an objective level.

The audio was processed through the compressors in a professional recording studio in London (UK) and assistance and guidance was given from the resident in-house engineers who had much experience with the devices under test. All other parts of the signal chain were bypassed for the recording process, and the audio was passed through high-quality Prism ADA8 converters at 44.1kHz 24bit. A higher sample rate was considered, but at this juncture in the research was deemed unnecessary. One possible option for analysis was to create full multi-track mixes using the audio. The issue of considerable processing power when working at sample rates such as 96kHz or 192kHz became a concern, which resulted in the decision to work at 44.1kHz. Recordings made at 44.1kHz 24bit (a standard delivery format) and using high-quality converters such as those employed in this chapter are sufficient for a study of this kind and the sonic signature of the devices under test will be present in the audio.

5.2 Testing for Non-Linearity

To investigate THD, the compressors were fed a 1kHz sine wave of three different amplitudes. This tested for differences in non-linearity as a function of input level. The compressors were sent a signal at 0dBu, +9dBu and +16dBu. An additional tone of +20dBu was sent to the compressors, but the results were ignored as the extreme THD figures led the author to believe clipping had occurred at the digital audio conversion stage. When making measurements to test input-based non-linearity, all the compressors were set to have no gain reduction or set to a compression off mode if the compressor had one. Therefore, only non-linearity that occurred because of the input level was measured. Audio with the compressor engaged in gain reduction was also measured to assess non-linearity when the compressor was compressing the signal and will be discussed later. The results from the non-compression activity tests at all input levels can be seen in Table 5-1 and the FFT plots for the 16dBu tone can be seen in Figures 5-1 and 5-2.

Compressor	THD % 0dBu	THD % 9dBu	THD % 16dBu
1176 Rev D	0.01	0.06	0.25
dbx165A	0.00	0.00	0.00
Fairchild	0.01	0.04	0.08
LA2A	0.20	0.50	0.90

Table 5-1: THD for four compressors at three input levels with a 1kHz tone

Under this test, the dbx165A is very clean with insignificant amounts of non-linearity. The distortion rises by a small amount as the input level is increased but the artefacts are inaudible, and the compressor is transparent. When calculated as a THD percentage, these small amounts of non-linearity result in 0% values across all input levels.

The Fairchild is also clean in this test, with the most prominent harmonics at 2kHz and 3kHz. There are additional low-level harmonics at 50Hz,100Hz and 150Hz that appear to be low-level artefacts from mains hum but they do not rise in level as the input is increased. The harmonics that cluster around the 1kHz test tone are spread apart by 50hz and are most likely sidebands that have been created in the audio as a result of sum and difference components that have been generated as a product of the test tone and mains hum frequencies. These spurious harmonics rise in level in accordance with an increase in test tone level and may play a role in this Fairchild's sonic signature, particularly at more driven input levels. Additional harmonics have been created at integer multiples of the 1kHz test tone but drop off sharply after the 3KHz harmonic. Without the mains hum and the side band frequencies, the Fairchild is notably clean when driven with the +16dBu input. The THD figure with the +16dBu driven input is in the region of 0.07%.

The LA2A has a slight amount of mains hum at integer multiples of 50Hz. The hum artefacts are lower in level than the Fairchild's and they do not modulate with the test tone frequency to introduce side band artefacts. Compared to the Fairchild, the LA2A has significantly more non-linearity at all input levels. At the 0dBu input level, the second and third harmonics are -54dB and -85dB down and rise up to -41dB and -57dB by the loudest input level of +16dBu. This distortion results in an audible amount of THD that is rated at 0.5% and 0.9% for the +9 and +16dBu input levels.

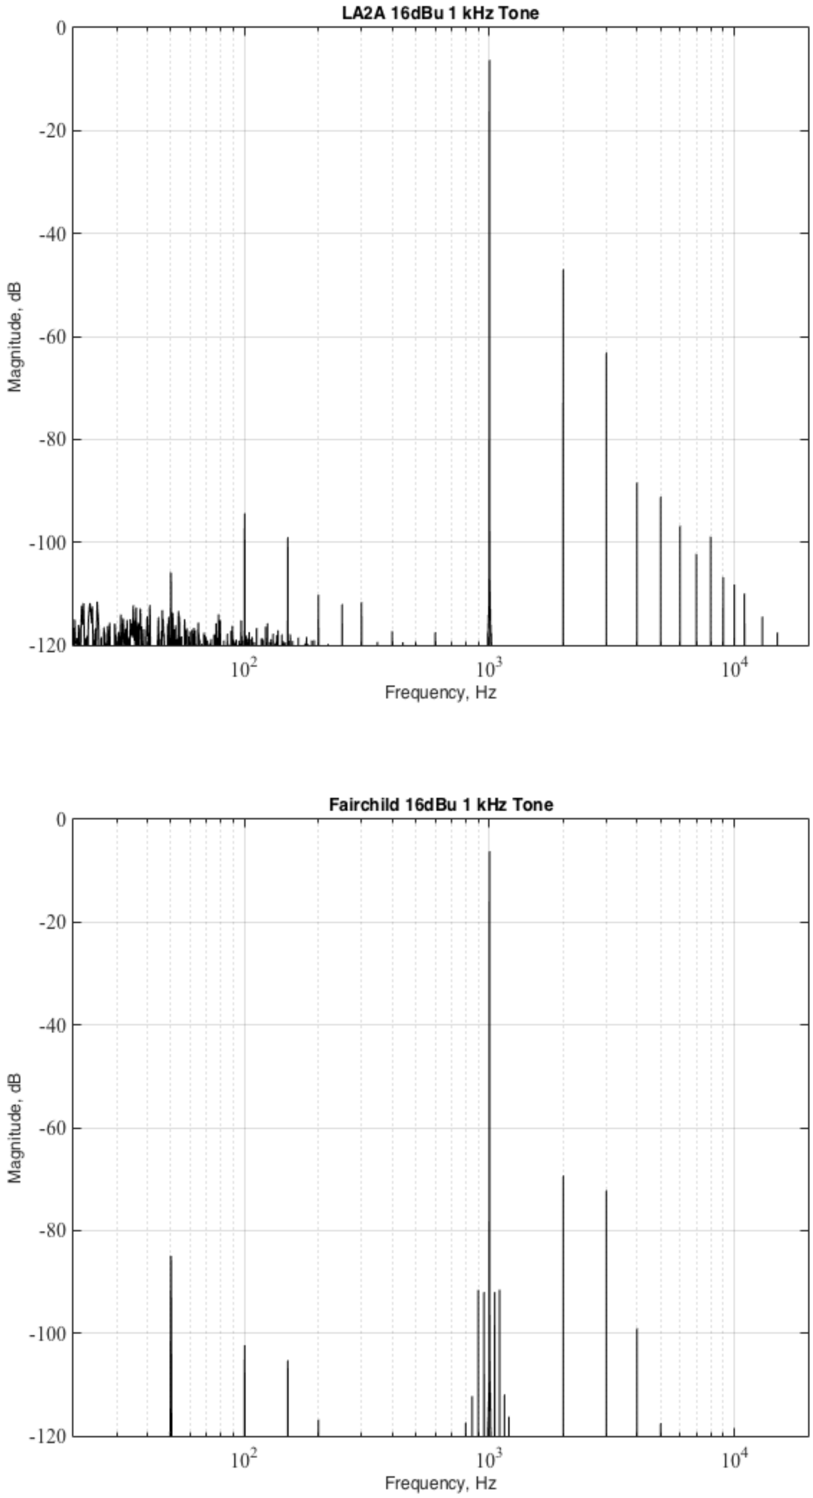

Figure 5-1: THD 1kHz tone with 16dBu input. LA2A top and Fairchild bottom

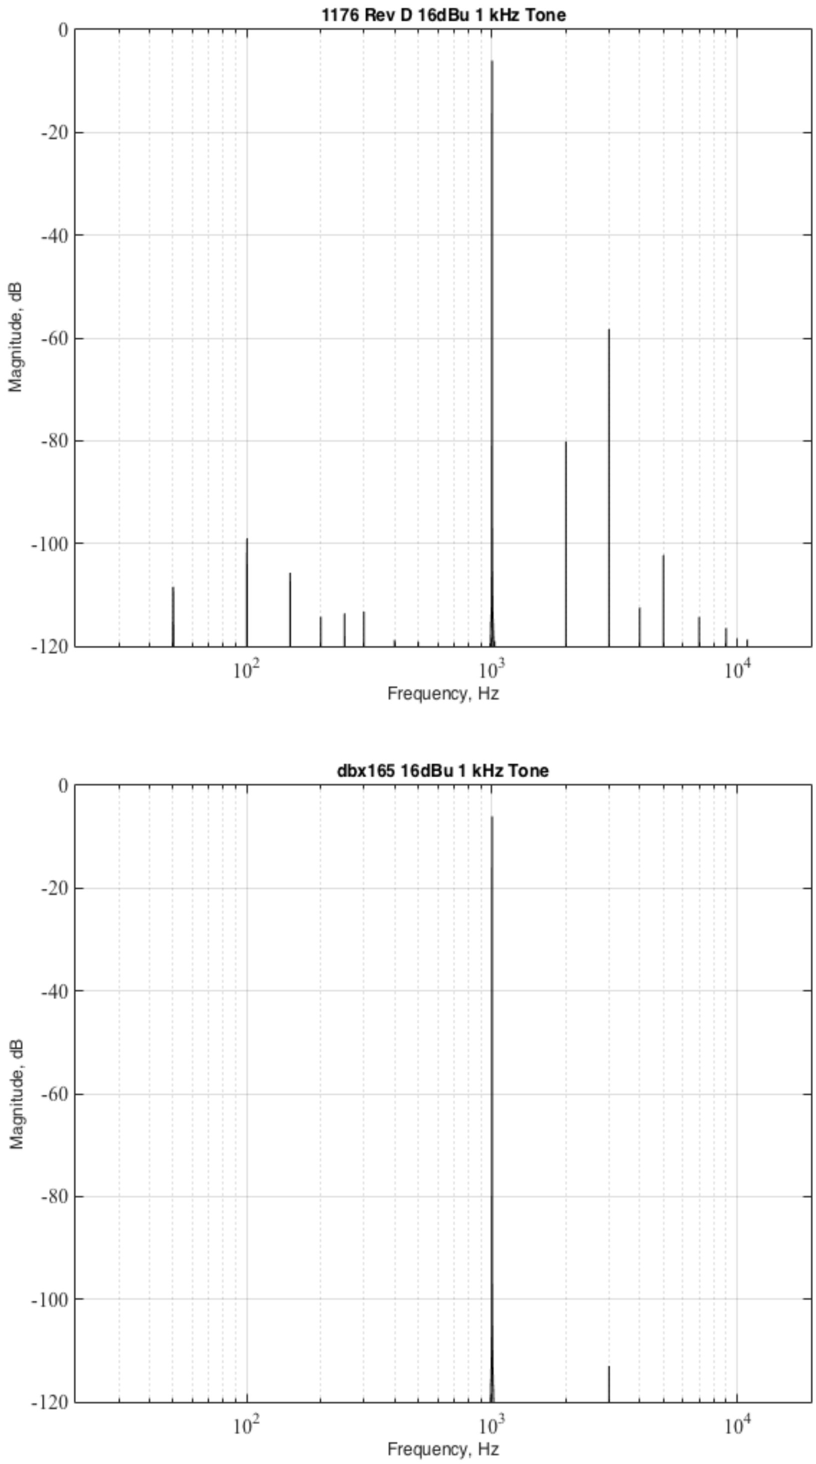

Figure 5-2: THD 1kHz tone with 16dBu input. 1176 top and dbx165A bottom

The 1176 Rev D has a clean THD result until +16dBu where the third harmonic increases in level to -52dB down from the test tone. This harmonic is partly responsible for the THD percentage of 0.25%, which is still a low THD figure for a driven input.

THD percentages during compression and release can be seen in Table 5-2. The plots generated while the compressors were in gain reduction are shown in Figures 5-3 and 5-4 and the plots during release in Figures 5-5 and 5-6. It was hoped that meaningful plots could be generated from the attack stage but the FFT output was too noisy and frequency components were difficult to discern, therefore this behaviour has not been analysed.

Measurements were made at four frequencies, 125Hz, 250Hz, 500Hz, and 1kHz. It was found the most significant differences between frequencies were between 125Hz and 1kHz thus they are the results discussed in this chapter. Visually the only difference between the results is the relative difference of the amplitudes of the non-linear components, therefore, only 1kHz plots are included in Figures 5-3 to 5-6 as they communicate all the necessary information.

The time constants used in this test made use of rock vocal settings that were extracted from the literature and adapted slightly by the engineers to account for any idiosyncrasies of the units being tested. The settings used for the 1176 had the attack at 3, release at 7, the Fairchild set at time constant 1, the dbx165A set for a fast release and moderately fast attack. The LA2A has no time constant settings.

Compressor	**THD % 125Hz In Comp**	**THD % 125Hz In Release**	**THD % 1kHz In Comp**	**THD % 1kHz In Release**
1176 Rev D	0.40	0.27	0.07	0.01
dbx165A	0.27	0.14	0.14	0.12
Fairchild	0.56	0.35	0.07	0.03
LA2A	2.13	1.18	2.64	2.38

Table 5-2: THD for four compressors during compression activity

As can be seen in Table 5-2, the LA2A has the most non-linearity under this test. It has significant amounts of THD for both tones during compression and release. For the 125Hz tone, the Fairchild is the next most non-linear, followed by the 1176 and the dbx165A is the cleanest. The dbx165A is second most distorted with the 1kHz tone, and the 1176 and Fairchild are almost equally as clean. The LA2A is more distorted with the 1kHz tone than the other compressors, which are all cleaner at this frequency. Colouration in this range appears to be one of the reasons why the LA2A was the second most popular vocal compressor in Chapter 4 with descriptors such as crunchy, colour and slightly distorted used to describe this compressor on the voice.

Figures 5-4 and 5-6 reveal interesting behaviour. The dbx165A is exhibiting considerable non-linearity during both compression and release. As well as prominent second and third harmonics, some additional sidebands are clustering around the test tone frequency, spread out in 50Hz increments. Low-level harmonics at the bottom end of the spectrum are evident in this plot and are at 50Hz, 150Hz, 250Hz, and 350Hz. These harmonics are sum and difference frequencies of the non-linearity.

The Fairchild has similar artefacts to the dbx165A. In addition to sidebands clustering around the 1kHz test tone the Fairchild has artefacts that group around the non-linear harmonics. The majority of these artefacts are sum and difference harmonics, of very low level, and outside of an audible range where they would be perceived as distortion. However, they may play a role in the sonic signature of this compressor, fusing with harmonics in complex program material for subtle colouration. The effect of low levels of distortion on audio has been investigated in a range of studies. The results suggest that subtle non-linearity can be perceived as a timbral change and a fusion of the harmonics rather than overt distortion (Gabrielsson & Sjögren, 1972; Gottinger, 2007; Petri-Larmi, Otala, & Lammasniemi, 1978). The sidebands in the Fairchild audio quickly disappear after gain reduction thus they are a creation of the gain reduction element. This behaviour can be seen in the release plot for the Fairchild. The LA2A has sideband harmonics present in the audio (again they are products of the sum and difference of the tone and the hum) but the main colouration effect is due to strong odd order harmonics, with 3kHz and 5kHz being -32dB and -45dB down from the test tone respectively. These harmonics bring THD well within the audible range. The non-linearity is reduced during the release stage, but it is still audible, albeit to a much lesser extent than when in compression.

During compression and release the 1176 produces only harmonic distortion (ignoring the low-level mains hum) and there are no sidebands. When compared to the Fairchild, the 1176 has a similar amount of non-linearity and cleans up with the 1kHz tone. The THD percentages for the 1176 are affected by the third harmonic used in the calculation and differences in results are often due to level variations in this third order component. This is particularly true for the 1kHz tone of the 1176 and Fairchild, which have small THD results for this test frequency.

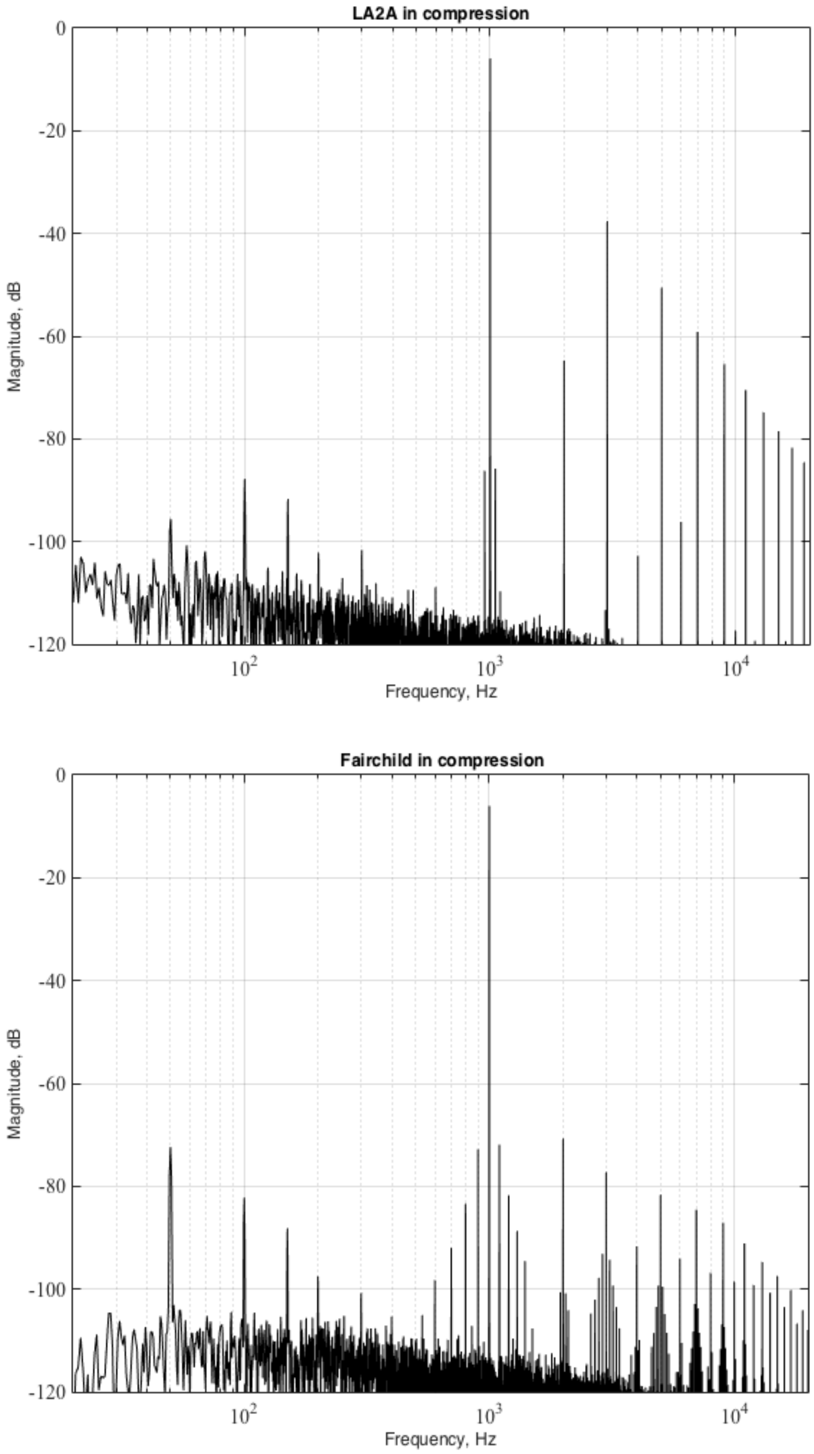

Figure 5-3: THD during compression. LA2A top and Fairchild bottom

It is suggested the reader listens to the static THD test tones in the audio analysis folder of this thesis to hear the sound quality. The levels have been matched between tones for a fairer comparison and the colouration from the non-linearity is clearly audible, particularly in the +9dBu and +16dBu tones.

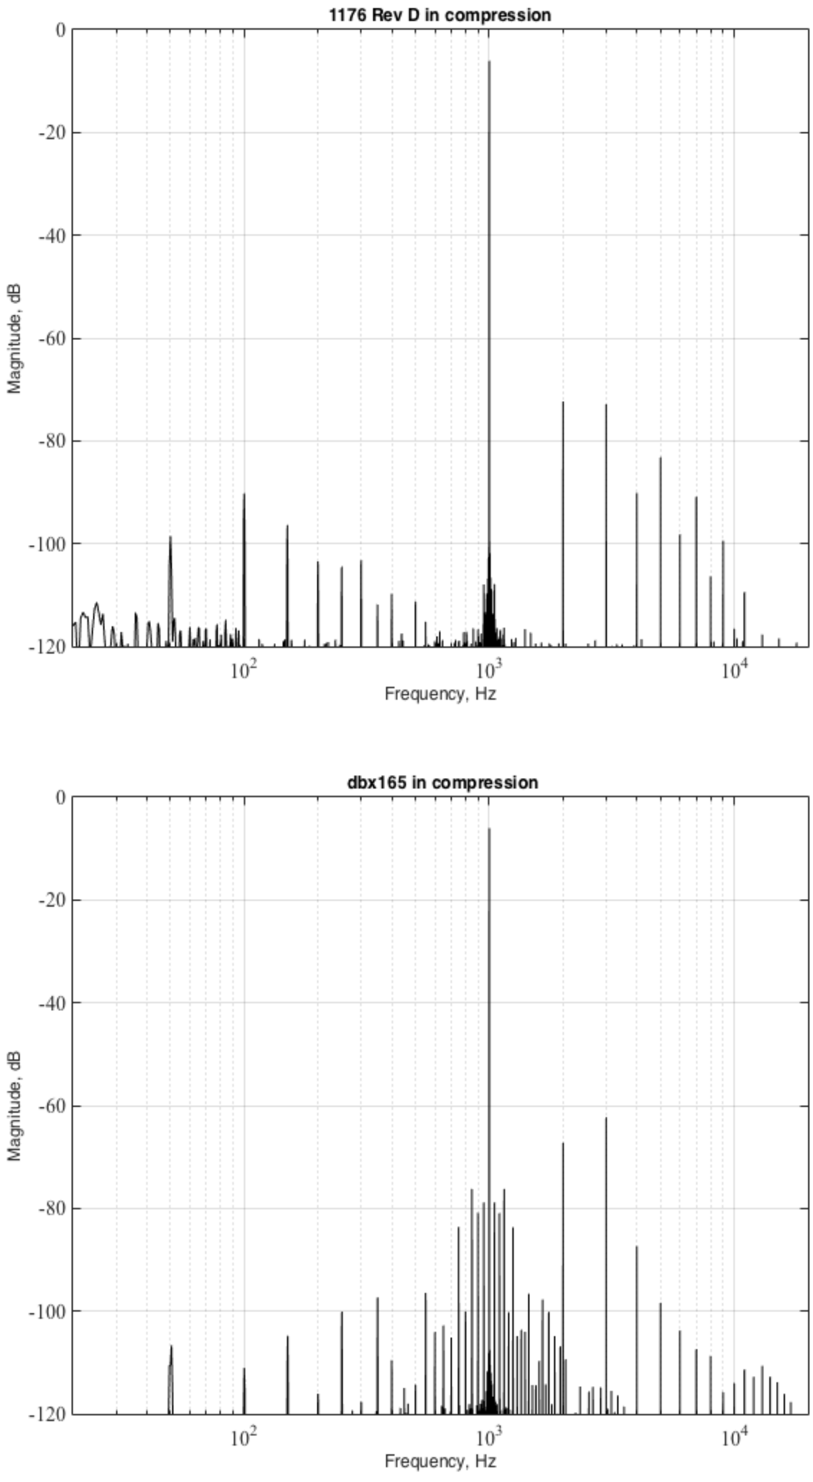

Figure 5-4: THD during compression. 1176 top and dbx165 bottom

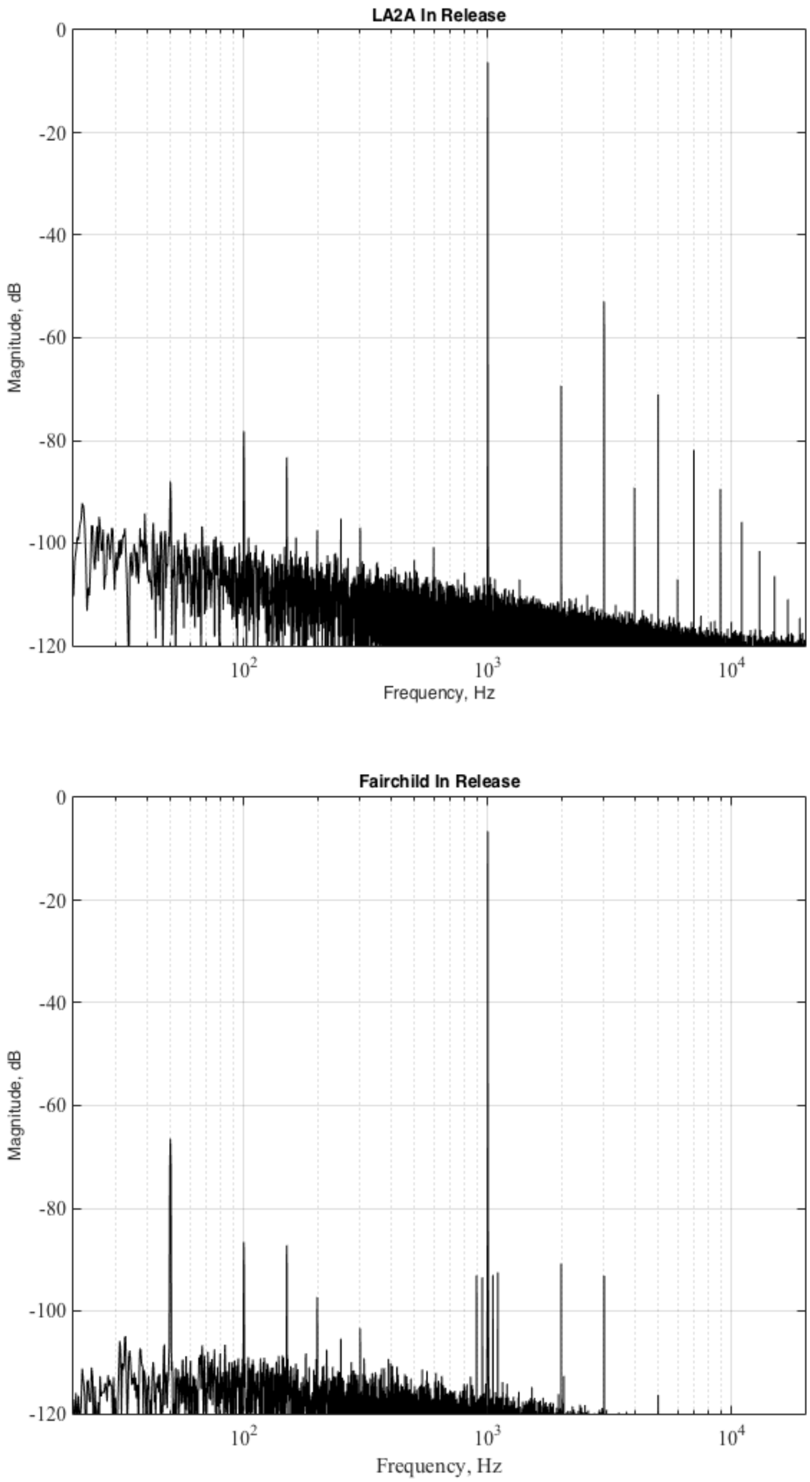

Figure 5-5: THD during release. LA2A top and Fairchild bottom

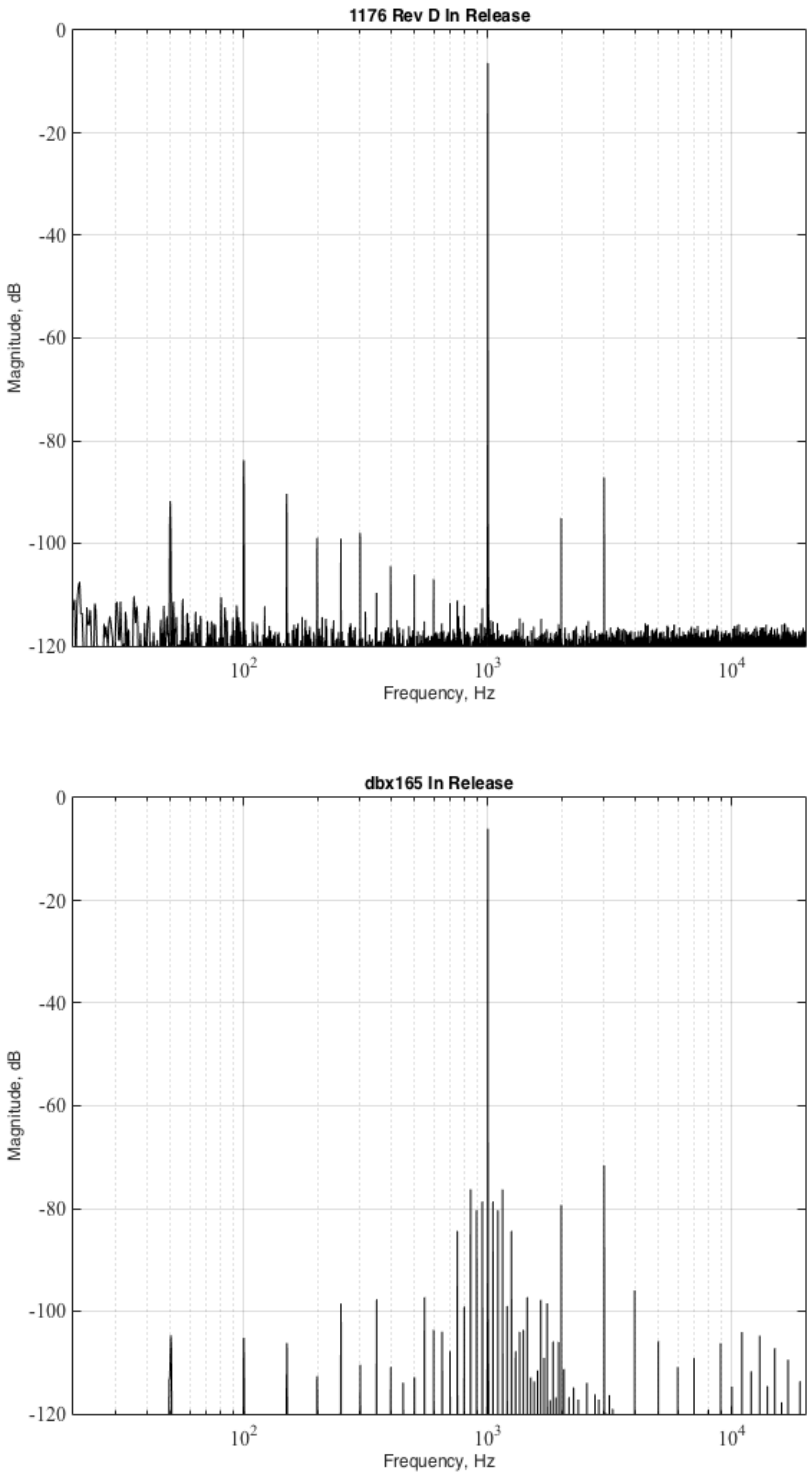

Figure 5-6: Thd during release. 1176 top and dbx165 bottom

5.3 Comparison of Compressor Time Constants

A series of burst measurements using the method described in Chapter 2 were made on the compressors to compare their timing laws. These tests made use of the same compression settings as used previously. The burst signal consisted of a sine wave of three different amplitudes (0dBu, +12dBu, 0dBu), and each of the compressors was set so that the 0dBu signals were just under the threshold point. This meant that the compressors were triggered into gain reduction only during the +12dBu tone. This burst tone sequence can be seen in Figure 5-7 below.

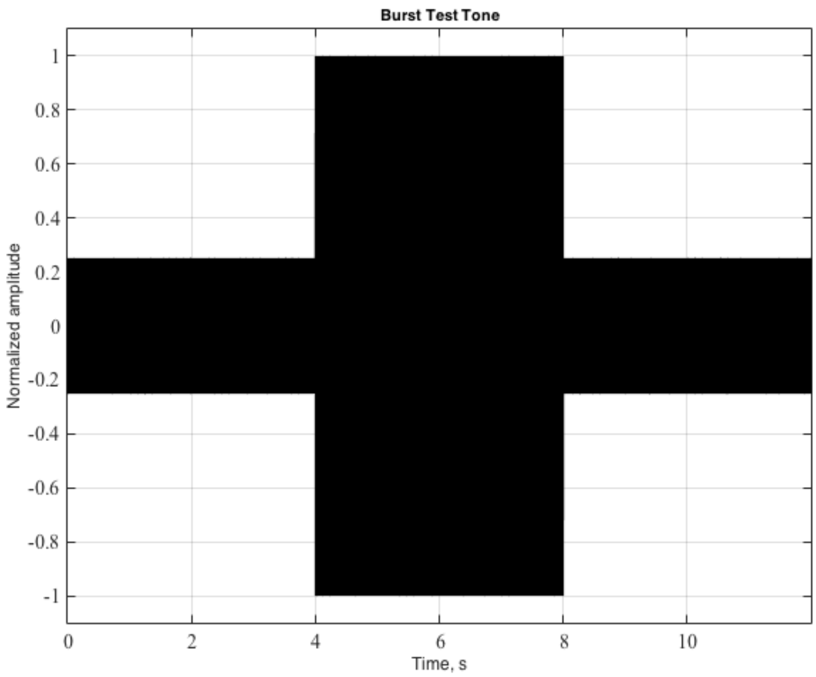

Figure 5-7: Tone burst test tone

The burst was sent to the compressors at four different frequencies to test for any variation in behaviour as a function of frequency. The frequencies used were 1kHz, 500Hz, 250Hz, and 125Hz. Additional measurements were made to check for variations in the time constants as a function of overshoot level. This test used the same burst sequence but with the higher-level amplitude tone set to overshoot the threshold by 3dB and 6dB. The results of this test revealed that there was no significant difference in time constant behaviour as a function of overshoot amount therefore only the 12dB overshoot is discussed in this thesis. The timing curves measured at the output can be seen in the plots in Figures 5-8 and 5-9, each plot features the four different frequencies, plotted with the lowest to highest from left to right.

Both the 1176 and LA2A have a gentler overall attack curve than the dbx165A and the Fairchild. The 1176 has an exponential curve with the attack portion clamping down quickly on the initial overshoot and then applying a more gradual amount of gain reduction over the rest of the steady state signal. The first stage of the attack on the 1176 pulls down the signal over a short period of time (approximately 5ms) but the final resting level for the steady state compression is not reached until approximately 950ms. The release portion of the signal takes approximately 650ms to reach its final level. There is slight variation in overshoot at the different frequencies. The 1kHz tone overshoots by 2dB more for 1ms before the attack time starts to attenuate the signal.

The LA2A has the slowest response of the three units with a slower attack and release. The unit appears to have a multistage attack with 63% of the gain reduction occurring over the first 125ms and all attenuation applied over approximately 1100ms. There are some fluctuations visible during the period of steady state gain reduction, and these small fluctuations may be as a result of variations in the T4 cell activity. This rippling behaviour from aging T4 cells was discussed in Chapter 3. The release portion measured is approximately 1300ms and creates a long, gentle release curve. The LA2A attenuates lower frequencies more than higher frequencies, although the difference is only approximately 0.2dB. The LA2A attenuates the signals more than the 1176 and dbx165A but not as much as the Fairchild. This is despite all the compressors being set for the same amount of gain reduction on their VU meters.

The dbx165A's attack curve is abrupt with this setting, it quickly attenuates the audio over approximately 65ms. The release is aggressive with the entire curve lasting approximately 95ms. The response of the dbx165A is consistent over all the frequencies tested, and there is no observable difference in the attack or release times.

The Fairchild has a consistent behaviour across all frequencies tested, with the attack time falling in the 30ms time span. Like the 1176, there is a slight increase in the overshoot amount with the 1kHz tone and this overshoot lasts for 1ms. Release times are consistent over the frequency range, with the complete release curve measuring approximately 900ms.

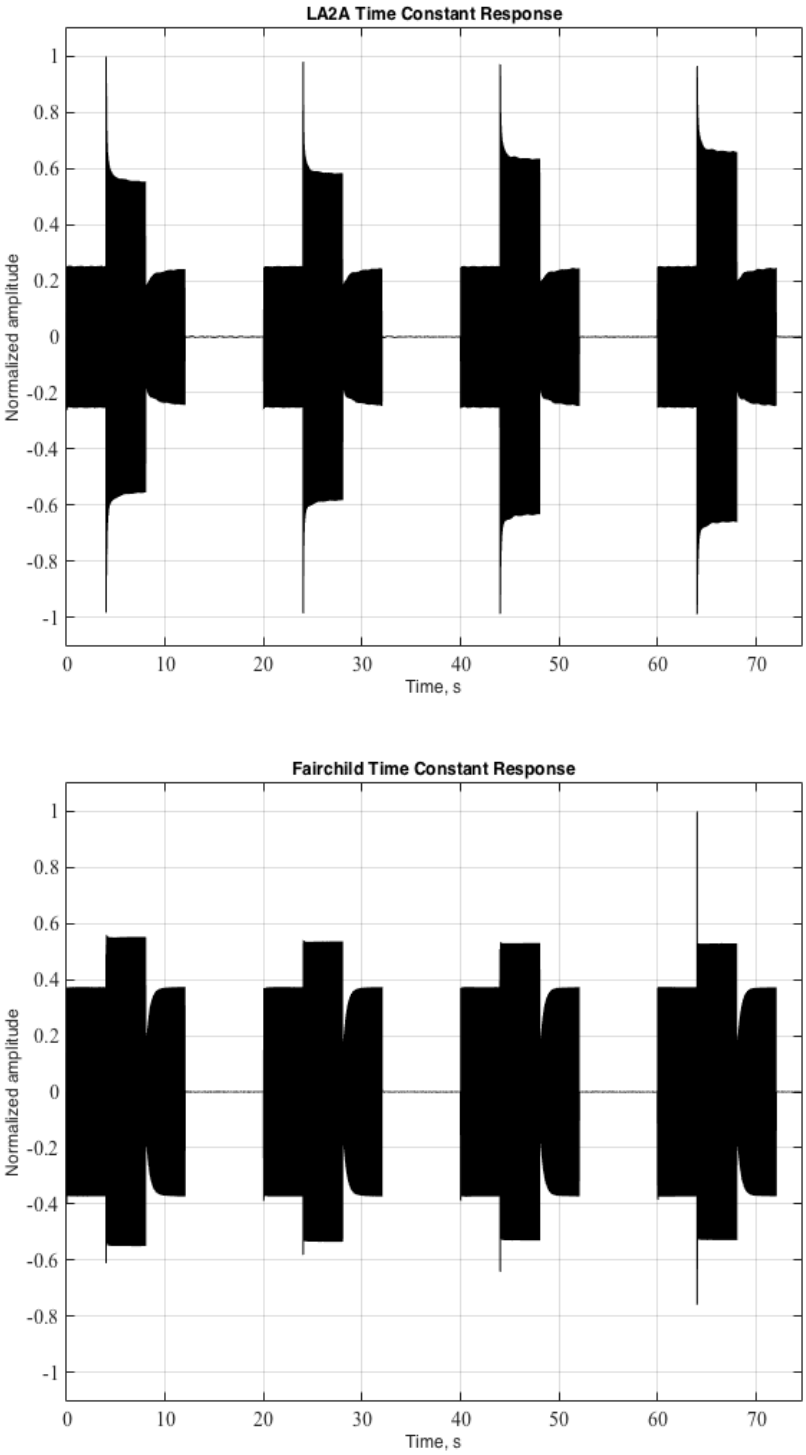

Figure 5-8: Tone burst shaped with compressor's timing curve. LA2A top and Fairchild bottom

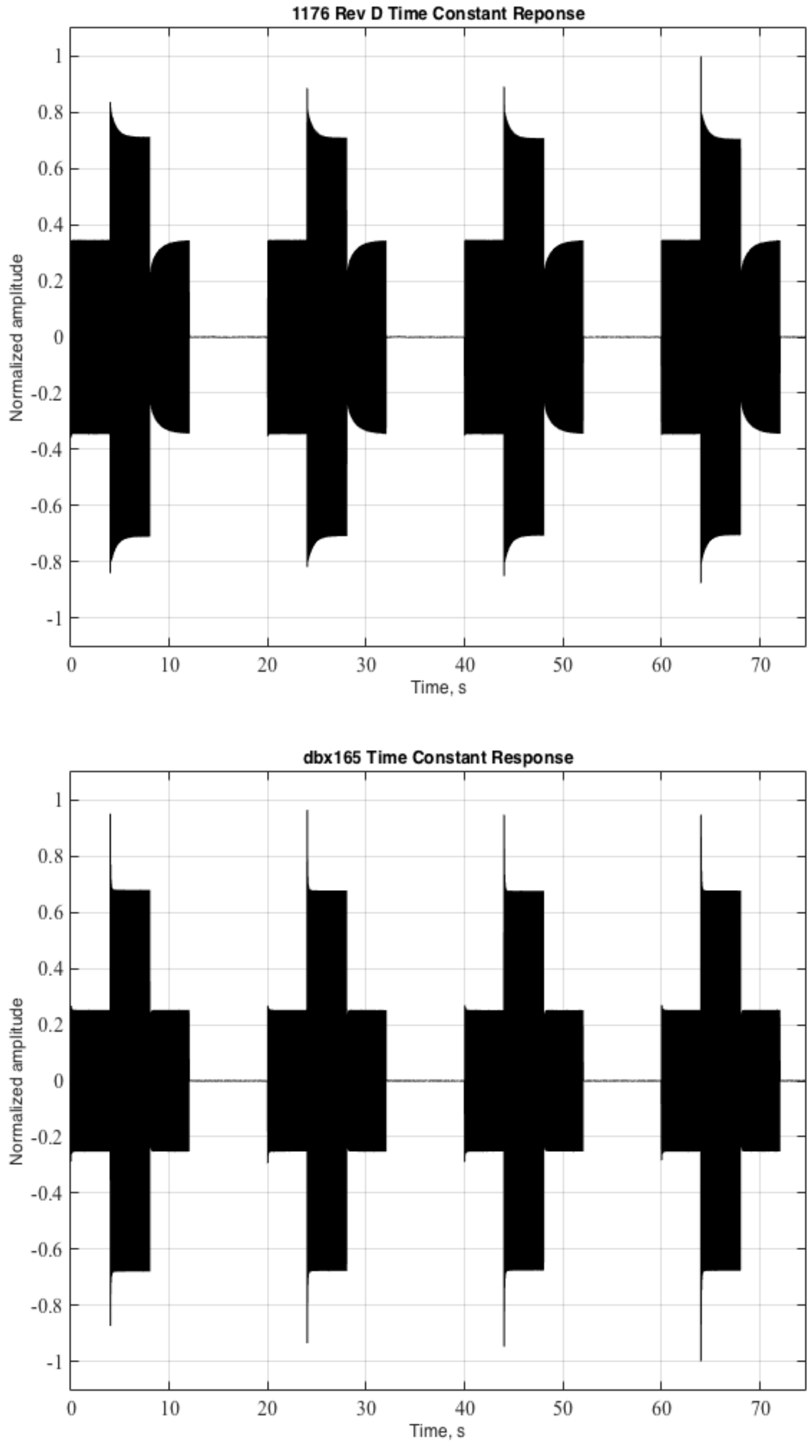

Figure 5-9: Tone burst shaped with compressor's timing curve. 1176 top and dbx165 bottom

The release curve of the Fairchild is similar to the 1176 despite the attack being more aggressive. The Fairchild has also attenuated the signal more than the other compressors in this test. Again, it is worth noting this compressor was adjusted to show the same amount of gain reduction on its VU meter as the other compressors.

There are clear differences in the way in which the compressors apply gain. Although some detail in regards to measured time has been discussed here, the reader is directed to focus on the general curve when looking at the plots. The main idea is to demonstrate the general differences between the units when set to compress a vocal with settings appropriate for that unit. There is minimal variation in the time constant behaviour when using test tone bursts of differing frequencies, and the only observable difference is a slight increase in overshoot for the 1176 and Fairchild compressors when using the 1kHz signal.

5.4 THD and Time Constant Testing: Conclusions

The following conclusions can be made from the tone testing. The THD measurements show there are some similarities in how the compressors impart non-linearity, but there are subtle differences in the harmonic content and amplitude of these artefacts. The dbx165A and Fairchild compressors are cleaner than the 1176 and LA2A when not compressing the signal. During compression, non-linearity in all compressors is significantly more pronounced and the Fairchild and dbx165A have many spurious sum and difference sidebands. The non-linearity for all compressors quickly returns to much lower levels once the compressor has stopped attenuating.

Measurements made at excessive input levels, which is how an engineer might drive the input for heavy colouration, were compromised because of digital clipping, presumably by clipping the input of the audio converters. The omission of hot input levels is a limitation of the study. Furthermore, intermodulation distortion measurements were not conducted as a working IMD methodology had not been developed at this point in research. It would have been interesting to see the results of the Fairchild under IMD tests and to observe the degree of sideband non-linearity. These two points make further investigations worthwhile, and the author hopes to complete this as part of further work.

From the burst measurements, it can be observed that the 1176 and LA2A have similar overall timing curves, but the 1176 is much faster to attenuate the initial overshoot. The dbx165A has consistent timing curves that are applied in a much sharper and aggressive manner than the other compressors. The Fairchild's attack time in this test is faster than the 1176 and the LA2A. However, the Fairchild's release curve is similar to the 1176. The LA2A shows slight modulations in amplitude during steady state compression, and this is most likely a result of modulation in the T4 cell. These small fluctuations in the timing behaviour play a role in the LA2A's sonic signature but it is not possible to quantify to what extent from this small test. The LA2A is somewhat more aggressive in the lower frequencies and attenuates them more than higher frequencies. This performance is most noticeable with the 50Hz test tone. There is a slight increase in overshoot amount when using the 1kHz tone for the 1176 and the Fairchild, but the actual increase in overshoot for the 1176 compressor is in the region of 0.8dB, not a significant difference, but something that may have an effect on its sonic signature. The Fairchild has an overshoot at 1kHz that differs by approximately 3dB, and it is audible when listening to the audio of these amplitude bursts.

The time constants used in this study are one snapshot of the compressors timing behaviour and this is an obvious limitation. The compression speed for all the compressors bar the LA2A can be affected by the attack and release settings. However, the focus of the test was to illustrate differences in timing law with regards to their general timing curve, and to this end the test was successful.

A further limitation of this test is that it did not address the program dependent nature of the compressors. This was not tested in these measurements but program dependency is addressed by an additional bout of testing on the 1176 and is discussed in Chapter 6. An area for further research is to compare the differences in transient and steady state compression behaviour between all four units tested in this chapter.

5.5 Measurements Using Complex Program Material

As discussed in Chapter 2, the use of test tones can only provide the user with a certain amount of useful information on a piece of equipment's sonic signature. With that in mind, a number of audio extracts from a professional production were processed through the compressors with settings extracted from the literature and adapted slightly (to account for variations in the equipment's behaviour) until deemed appropriate by the studio engineers and the author of this report.

The compressors were adjusted with settings for each unit that were applicable for a number of production scenarios. These production scenarios were highlighted in the grounded theory analysis in Chapter 4 as being standard approaches when using compression in mixing. The scenarios were:

- Aggressive and "in your face", heavily compressed rock vocals
- Full bodied and coloured bass part
- Pumping (amplitude modulation of steady state signals) room mics with attenuated transients

The compressors were not set with the same time constants. This was not possible due to the limited stepped controls on the Fairchild and the complete lack of time constant control on the LA2A. Instead, the compressors were set to impart as much of the sonic qualities listed above while working within the restrictions of that particular compressor. In this context, the compressors can be thought of as a black box, the primary concern being the sonic output and not the settings used to achieve it. The settings used are shown in Table 5-3.

Compressor	Vocal	Bass	Drums
Fairchild	Time Constant 1	Time Constant 3	Time Constant 1
LA2A	N/A	N/A	N/A
dbx165A	Moderately Fast Attack, Fast Release, 4:1	Fast Attack, Fast Release, 4:1	Fast Attack, Fast Release, 20:1
1176 Rev D	Attack 5, Release 7, 4:1	Attack 5, Release 7, 4:1	Attack 3, Release 7, All-Buttons

Table 5-3: Settings used to compress the audio

All devices were set to have the same amount of gain reduction with the VU meter on each compressor used as a guide. The 1176's VU reduction meter becomes of little use when in all-buttons mode so this aspect of the adjustments had to be calibrated partly by ears and in consensus with all of the participants in the studio at the time of testing.

5.5.1 Information on the Audio Material

The audio material used is from a professional recording session of the hard rock band Zico Chain. The recordings were made in a professional studio and were tracked by the engineer Joe Barresi, who has worked with artists such as Queens of the Stone Age, Kelly Clarkson, and Soundgarden. Full permission was given by the band to use the recordings in this project. The author of this report was keen to use recordings of a professional standard to keep this thesis focused on professional industry standard working methods.

The recordings had small amounts of compression added during tracking to control the fluctuations of the levels going to the converters. This is common working practice in a professional environment and was discussed in Chapter 4. The audio files used are the lead vocal track, a room mic track, and a bass DI track. The session files for the song included bass amp recordings, but the bass DI was chosen over the amp signals due to them having saturation from the amplifier. Using a recording that already had distortion would make it difficult to focus on non-linearity from the compressors. During the testing stage in the studio, the audio files were processed through the compressors simultaneously and set to achieve light, moderate and heavy amounts of gain reduction. These compression amounts were obtained by setting each compressor's gain reduction meter to fluctuate around -3dB for light, -6dB for medium and -12dB for heavy.

The quotes discussed in Chapter 4 showed that producers use significant amounts of gain reduction when using compression at the mixing stage. Therefore, it was decided the heavy gain reduction material was the most appropriate for the processing scenarios mentioned earlier. These audio files are analysed in detail in this chapter. The analysis of heavily processed material is in adherence with the comments made by hardware designer Paul Wolf, who advises on using significant gain reduction when assessing a compressor for its sound quality (Ciletti et al., 2008). A similar point is made by Campbell et al. (2014) during their research into compressors and signal masking. Initial analysis of the moderate and light files revealed there was some difference between the compressors. This difference was evident under critical listening and inspection of the FFT and time domain plots. Thus, for completeness, a discussion of some of the general differences between the gain reduction settings will be presented in a sub-chapter later in this chapter.

To make analysis easier to observe the full audio files were edited into small portions that featured elements in the music performance that would demonstrate the compressors ability to work over a wide dynamic range and frequency range. For example, a portion of the room mic recording was extracted that featured drums and cymbals as this covered the full frequency range of the drum kit and a section of the vocal was extracted that featured breathing and variation in the dynamics. To make subjective comparisons of the audio fairer all the material was loudness matched to -22 Loudness Units Full Scale (LUFS). Additional information on LUFS is provided by EBU–Recommendation (2011).

5.6 Results from Bass Material

The time domain and spectrogram plots for the bass extracts can be seen in Figure 5-10 (original audio) and Figures 5-11 and 5-12 (compressed audio). Amplitude statistics measured from the audio are presented in Table 5-4. The dynamic range was calculated using the amplitude statistics analysis in Adobe Audition, and this dynamic range calculation measures the difference between the minimum and maximum RMS amplitude (Adobe, 2016). For consistency, it may have been preferable to measure this as the minimum to the maximum LUFS, but the calculation did not offer this option. However, observations by the author suggests that differences in RMS and LUFS measurements made from the material used in the thesis were often small.

Compressor	Dynamic Range (dB)
Fairchild	18.06
LA2A	17.95
dbx165A	14.11
1176 Rev D	14.40
No Comp	21.51

Table 5-4: Amplitude statistics for the bass material

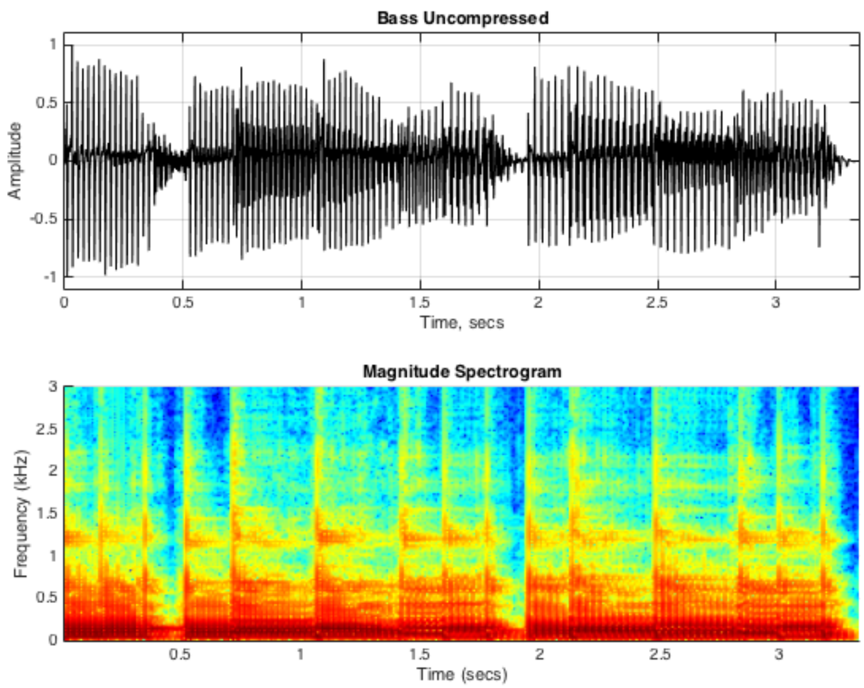

Figure 5-10: Time domain and spectrogram plots of the uncompressed bass material

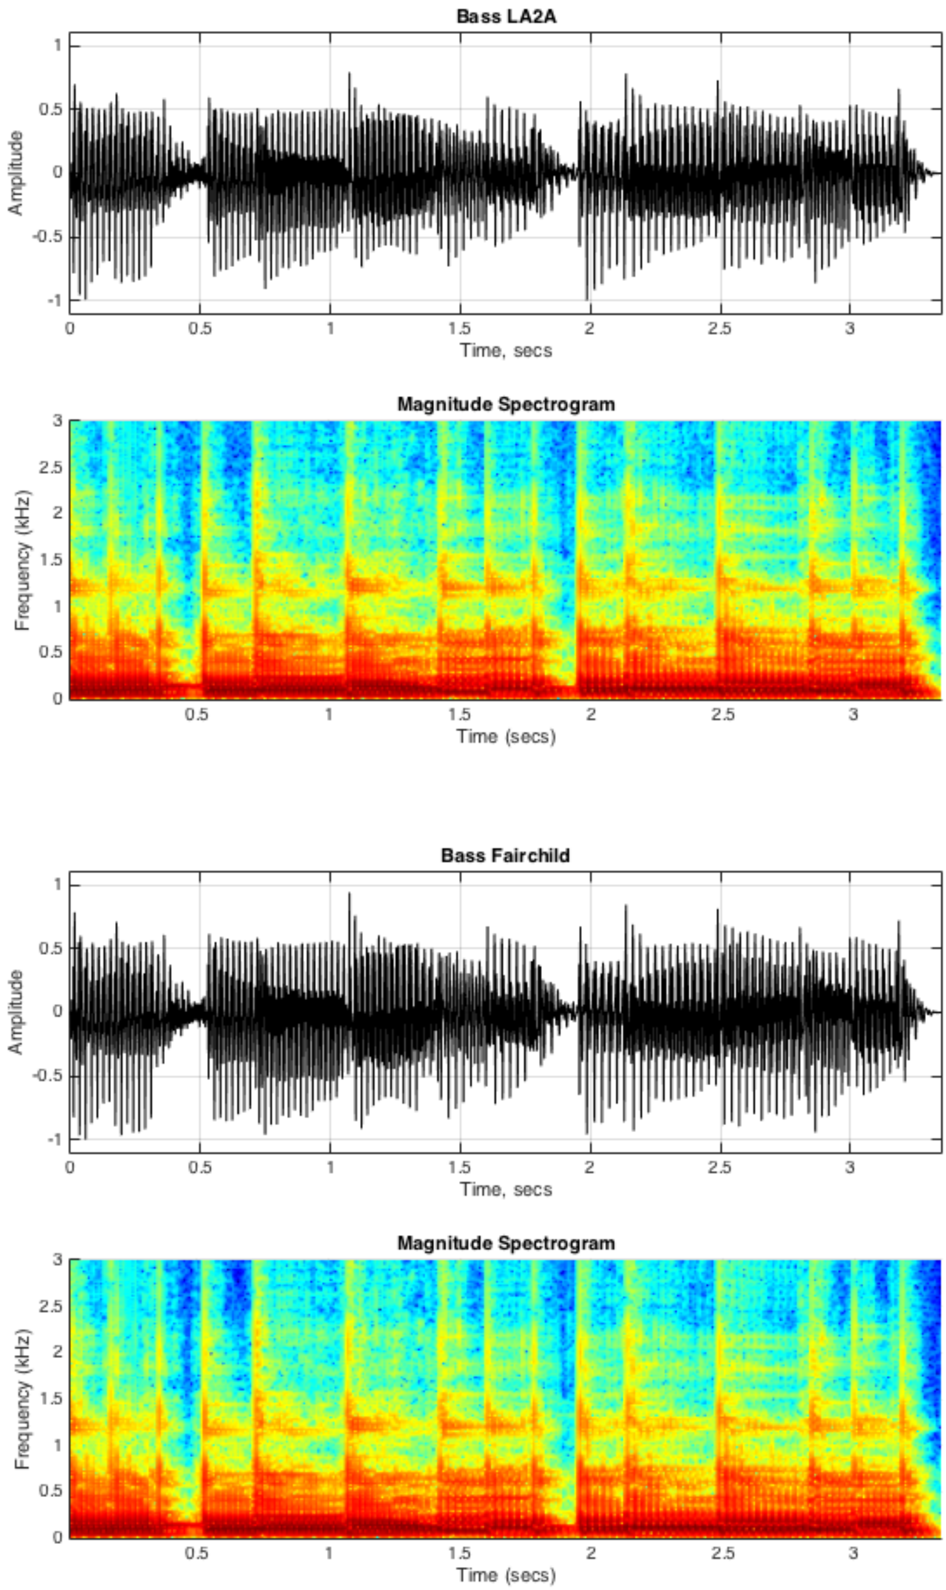

Figure 5-11: Time domain and spectrogram plots of the LA2A (top) and Fairchild (bottom) bass material

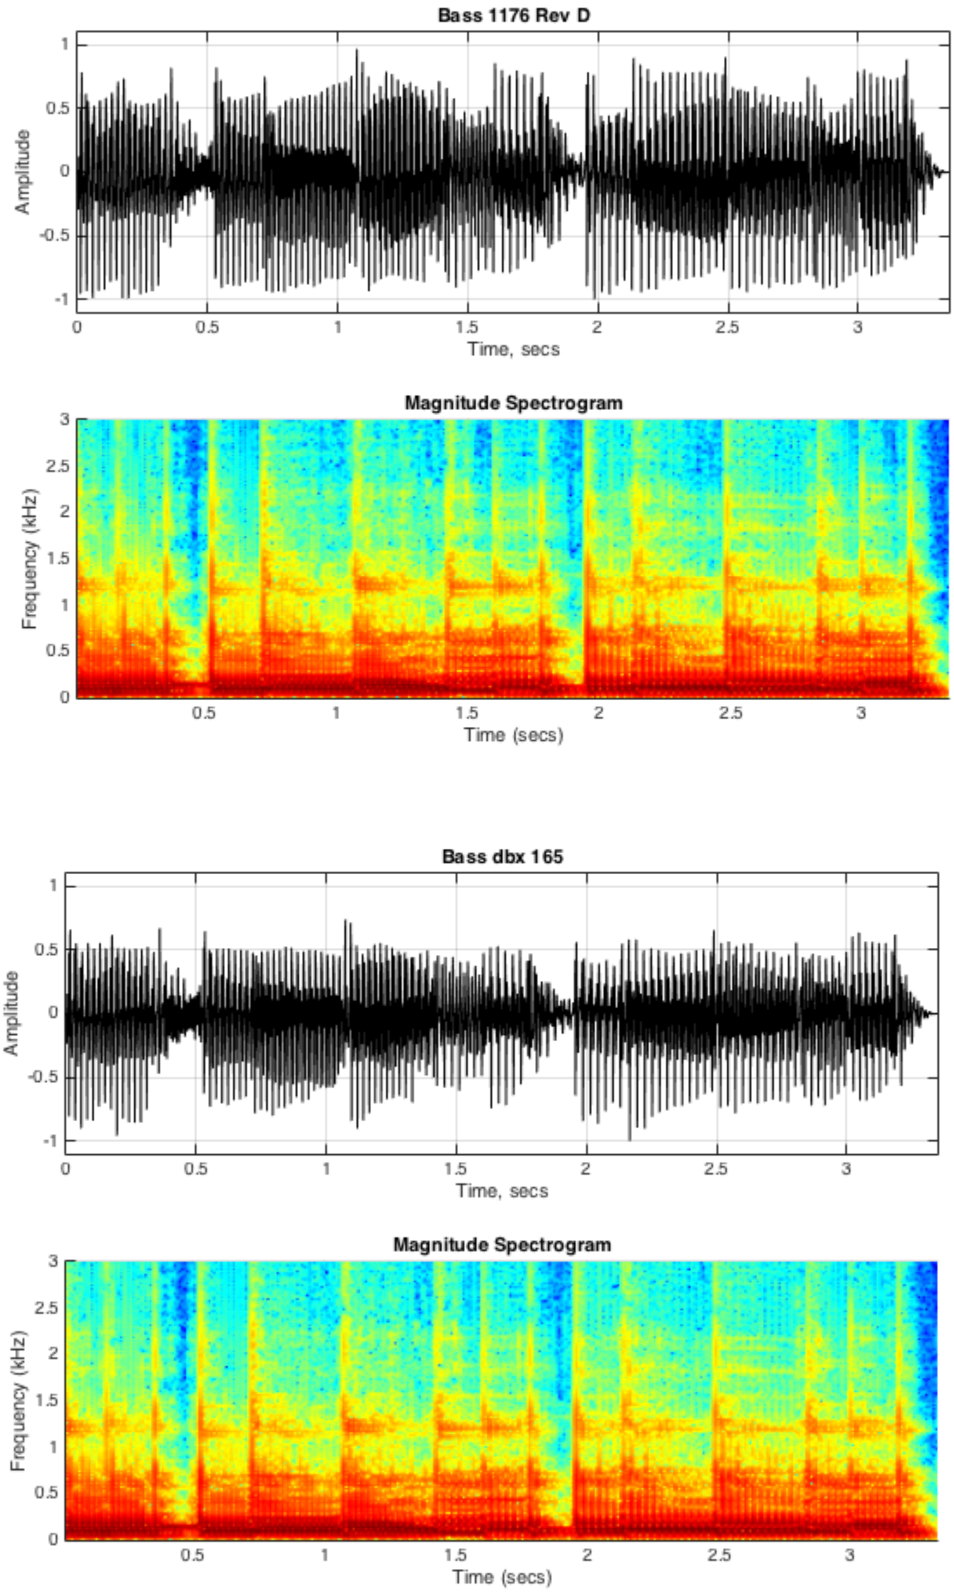

Figure 5-12: Time domain and spectrogram plots of the 1176 (top) and dbx165A (bottom) bass material

All compressors have increased the midrange spectral content between 1kHz and 1.5kHz with the Fairchild, 1176 and dbx165A being the most coloured. When compared to the uncompressed audio, the compressed signals have more energy across all bands. The dbx165A is the most coloured, with distortion that is audible in the audio. Additional non-linearity can be seen in all the spectrogram plots above 2.5Khz with the dbx165A particularly spectrally dense in this area. The time domain plots show an increase in the amplitude of quieter notes for all compressors. The Fairchild and LA2A have some similarities, this is due to the longer release curve of the Fairchild in time constant three being similar to that of the LA2A's timing law. The dbx165A and the 1176 Rev D are the most consistent in terms of amplitude.

The RMS energy of the audio was extracted and plotted over time to investigate differences in the envelope of the compressed material. The energy was plotted using a window size of 4096 and a hop size of 1024. RMS energy plots for all the compressors can be seen in Figure 5-13. There is some similarity in terms of the overall curve between the 1176 and the dbx165 and between the LA2A and the Fairchild. However, each of the compressors has a unique curve in terms of microdynamics, meaning the finer variations in dynamics within the overall envelope. Skovenborg (2014) provides more information on macro and microdynamics.

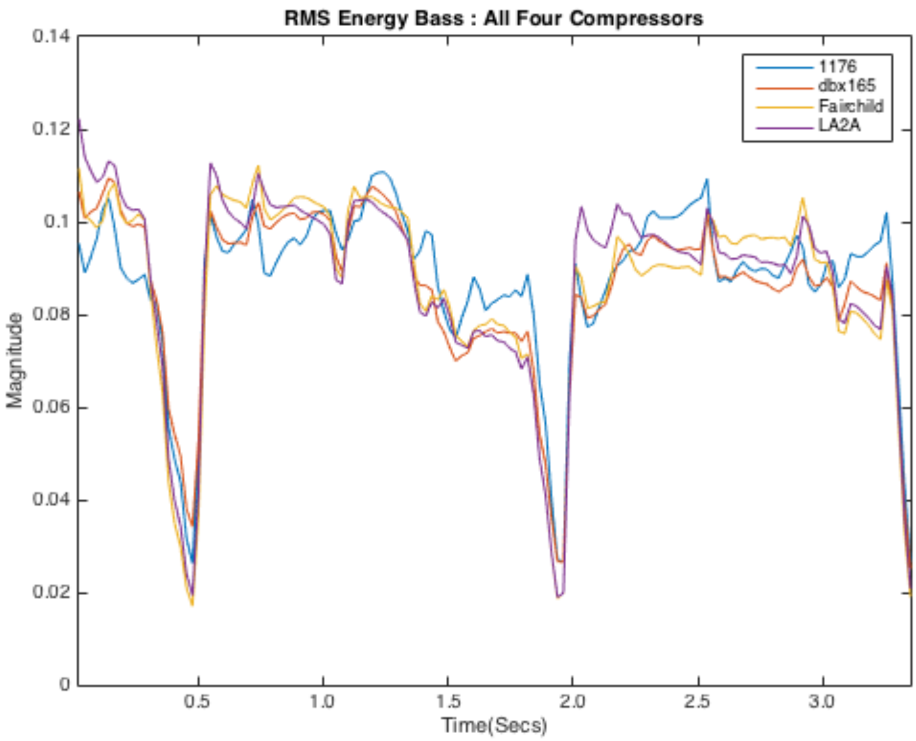

Figure 5-13: RMS energy for bass material. All four compressors

The Long-term average spectrum (LTAS) was extracted from the bass material with 1/16 octave smoothing using the MATLAB function detailed in Chapter 2. The results for the bass audio are shown in Figure 5-14 where the frequency axis of the plot includes only 50Hz-3kHz to focus on the areas most affected by the colouration. The compressed audio groups together closely in this LTAS plot and the dbx165A is shown to have the most energy across the frequency range.

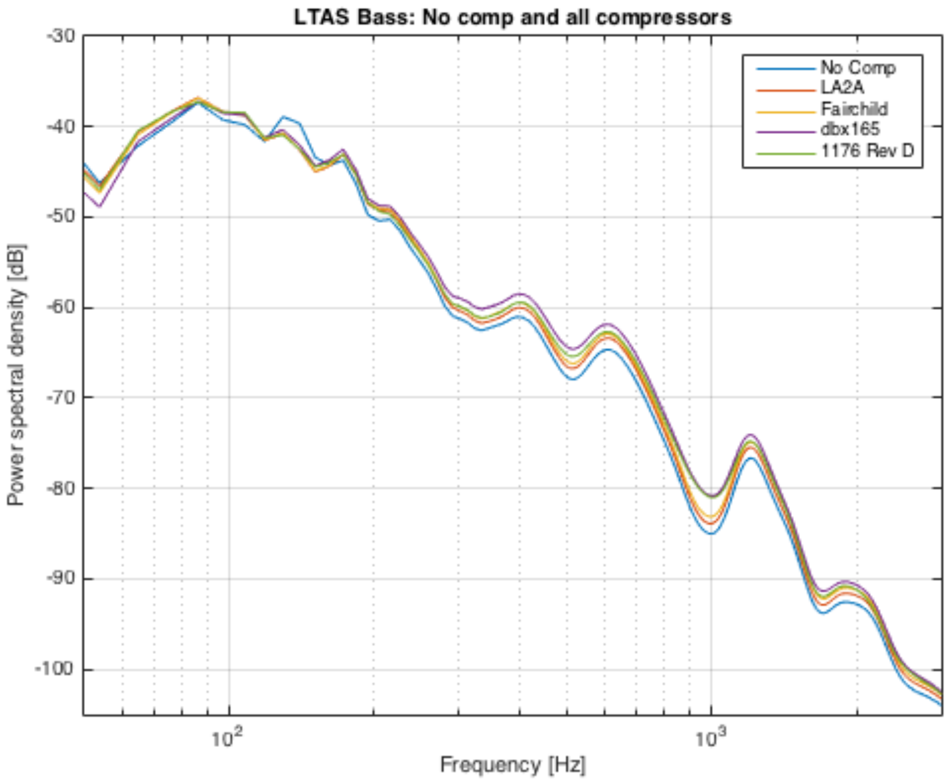

Figure 5-14: LTAS for bass material. All compressors and uncompressed

5.6.1 Audio Features: Bass Material

A number of salient audio features were extracted from the audio to investigate how each compressor affected the timbre of the sound. The results can be seen in Table 5-5.

The audio processed through the compressors has higher values for all the features except low-energy, and the compressors have changed the timbre of the original material by different amounts. All compressors bar the dbx165A (which is slightly higher) have similar spectral centroid and brightness values. This result may be due to the additional non-linear frequency components added during compression fusing with the audio material to create a perceptually brighter output. Zwicky

(2011) argues a side effect of DRC is its tendency to dull audio signals during heavy compression, but under this test, the compressors are having the opposite effect and increasing features that relate to perceptual brightness. Listening to the audio reveals there is a subtle increase to the top end that gives the material more bite and definition.

Compressor	Roll Off	Spectral Centroid	Brightness	Low-Energy
Fairchild	512 Hz	291 Hz	0.017	0.36
LA2A	481 Hz	292 Hz	0.016	0.36
dbx165A	564 Hz	312 Hz	0.018	0.38
1176 Rev D	519 Hz	293 Hz	0.016	0.34
No Comp	420 Hz	256 Hz	0.014	0.41

Table 5-5: Audio features extracted from the bass for all four compressors and the uncompressed material

As shown in Chapter 4, producers used descriptors such as bite and presence to describe the desirable qualities of bass compression. Part of what they are describing seems to be an increase in the spectral centroid and brightness values as a result of non-linear artefacts.

The low-energy feature is less for all of the compressors compared to the value calculated for the original audio material. Lower values of this feature mean that more of the frames of audio taken for calculation are higher than the average energy, resulting in a small low-energy rating, which translates as a more consistent piece of audio regarding the temporal distribution of energy. The 1176 has the lowest low-energy results for this test, which reflects the small dynamic range result for this compressor.

Zero crossing rate was extracted for the audio, but there was no clear trend in the results. This is a disappointing result and it was hoped that the ZCR feature could be used as a reliable measure of noise-like non-linearity when using compression.

Variations in spectral flux in low frequency bands were measured to see how the compressors differed. The audio was split using a filter bank in MIRtoolbox consisting of an octave-scaled second order elliptical filter. A sub-band covering 100-200Hz was examined for spectral flux and the mean and standard deviation values are shown in Table 5-6. This frequency band was chosen because the bottom end in the bass part was present here. All compressors have an increase in spectral flux in this band, and the dbx165A has the largest value. This increase in spectral flux for the dbx165A is because of additional non-linearity in the signal, most probably IMD artefacts. It has been found in a study by Alluri and Toiviainen (2010) that higher values of spectral flux in lower bands correlates with a perceptual sense of fullness, thus these compressors produce an increase in spectral flux and a more consistent dynamic range that will introduce fullness and weight to the bass material. These attributes were noted in the analysis in Chapter 4 as being desirable changes to timbre when using compression.

Compressor	Spec Flux Mean	Spec Flux Std
Fairchild	2.97	2.22
LA2A	2.67	1.89
dbx165A	3.17	2.67
1176 Rev D	3.10	2.5
No Comp	2.33	1.76

Table 5-6: Spectral flux extracted from the 100-200Hz sub band for all compressors and uncompressed bass material

Roughness was extracted to investigate non-linearity and was measured using the same filter bank audio as mentioned previously. The 100-200Hz area is where one would expect to find first order non-linear components on bass recordings, particularly because of waveshaping from fast time constants. The results are presented in Table 5-7.

Compressor	Roughness Mean	Roughness Std
Fairchild	1.91	2.46
LA2A	1.78	2.09
dbx165A	2.93	3.72
1176 Rev D	2.18	2.52
No Comp	1.52	2.01

Table 5-7: Roughness extracted from the 100-200Hz sub band for all compressors and uncompressed bass material

The dbx165A has the highest roughness figure, and this is consistent with the audible non-linearity that can be heard in the audio. The 1176 has the second largest roughness value followed by the Fairchild and the LA2A. The LA2A rates lowest for this measurement because its slower attack time is not distorting the audio material as much as the faster attack time constants found in the other compressors. Non-linearity was noted in Chapter 4 as a motivation for using compression on bass material, and this test demonstrates that compressors with fast time constants are the most effective at achieving this type of colouration effect.

5.7 Results from Vocal Material

A vocal track was processed through the compressors and the time domain and spectrogram plots are shown in Figure 5-15 (uncompressed) and Figures 5-16 and 5-17 (compressed audio). The amplitude statistics were extracted from the audio material and are presented in Table 5-8. The time domain plots for the compressed audio reveals the average level of the vocal has been raised and the dynamics levelled out. The breaths in between words have been increased in level and the envelope of the breaths at approximately 4 seconds have been reshaped by both the Fairchild and the dbx165A to create more consistency. The level of the final word starting at approximately 6 seconds is more consistent and this is the case for all the compressors tested.

Looking at the spectral content reveals there has been an overall increase in intensity during plosives and breaths. There is also a slight amount of attenuation in the top end, and this is because of the compressor's attenuating the signal after short bursts of plosives and sibilants. Perceptually this attenuation in the high end does not noticeably affect the spectral characteristics of the compressed material thus changes visible in the spectrograms appear to be more significant than they are.

Sonically all of the compressors sound similar, with the exception being the dbx165A that sounds significantly more distorted than the other devices. It is fuller in the midrange, and this is because of the additional non-linearity that has been added in this area. This accentuation in the midrange can be observed in the spectrogram in the 1.5-2.5kHz band from approximately 2.5 seconds to the end.

Despite the LA2A sharing similar amplitude statistics with the Fairchild and the 1176, it is not as aggressive in its application of compression. However, the LA2A has a fuller character to its sound and more body than the other compressors. The time domain plots for the 1176 and Fairchild are similar thus their timing curve is playing a role in the aggressive sonic signature. The Fairchild was used with its fastest time constant in this test thus this 200-microsecond attack means it is working at a similar speed to the 1176. The dbx165A has audible non-linearity in the audio, suggesting the compressor is distorting the voice. The differences between all the compressors are small and subtle, but they provide the original vocal with varying amounts of extra clarity and an aggressive character.

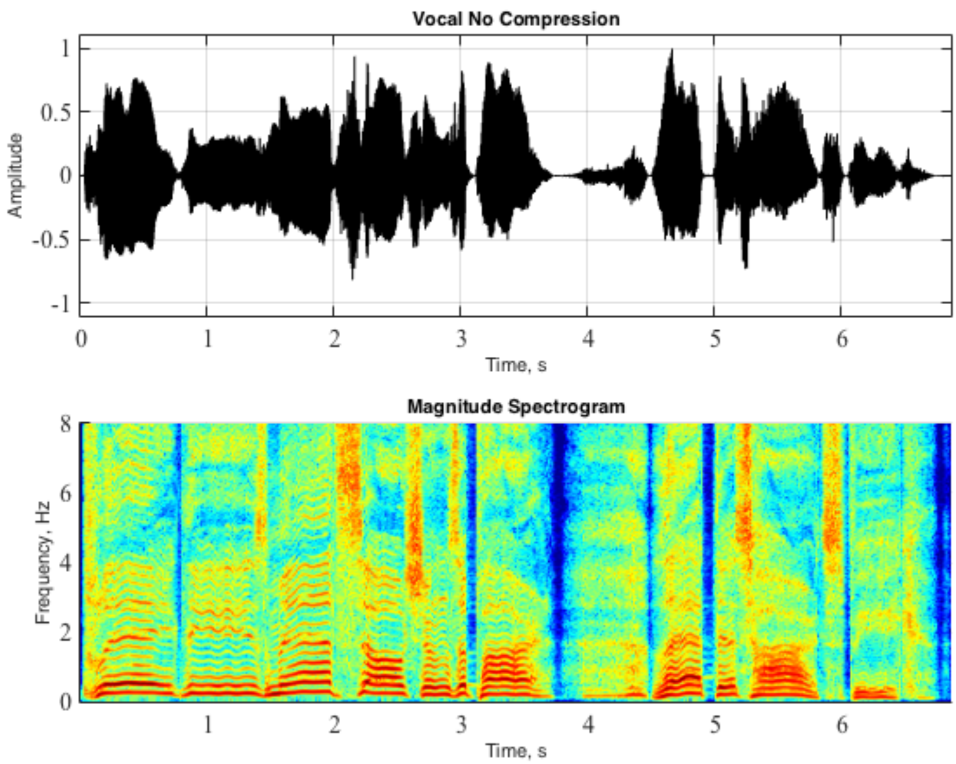

Figure 5-15: Time domain and spectrogram plots of the uncompressed vocal audio material

Compressor	Dynamic Range (dB)
Fairchild	29.75
LA2A	29.84
dbx165A	19.99
1176 Rev D	30.79
No Comp	34.48

Table 5-8: Amplitude statistics for the vocal material

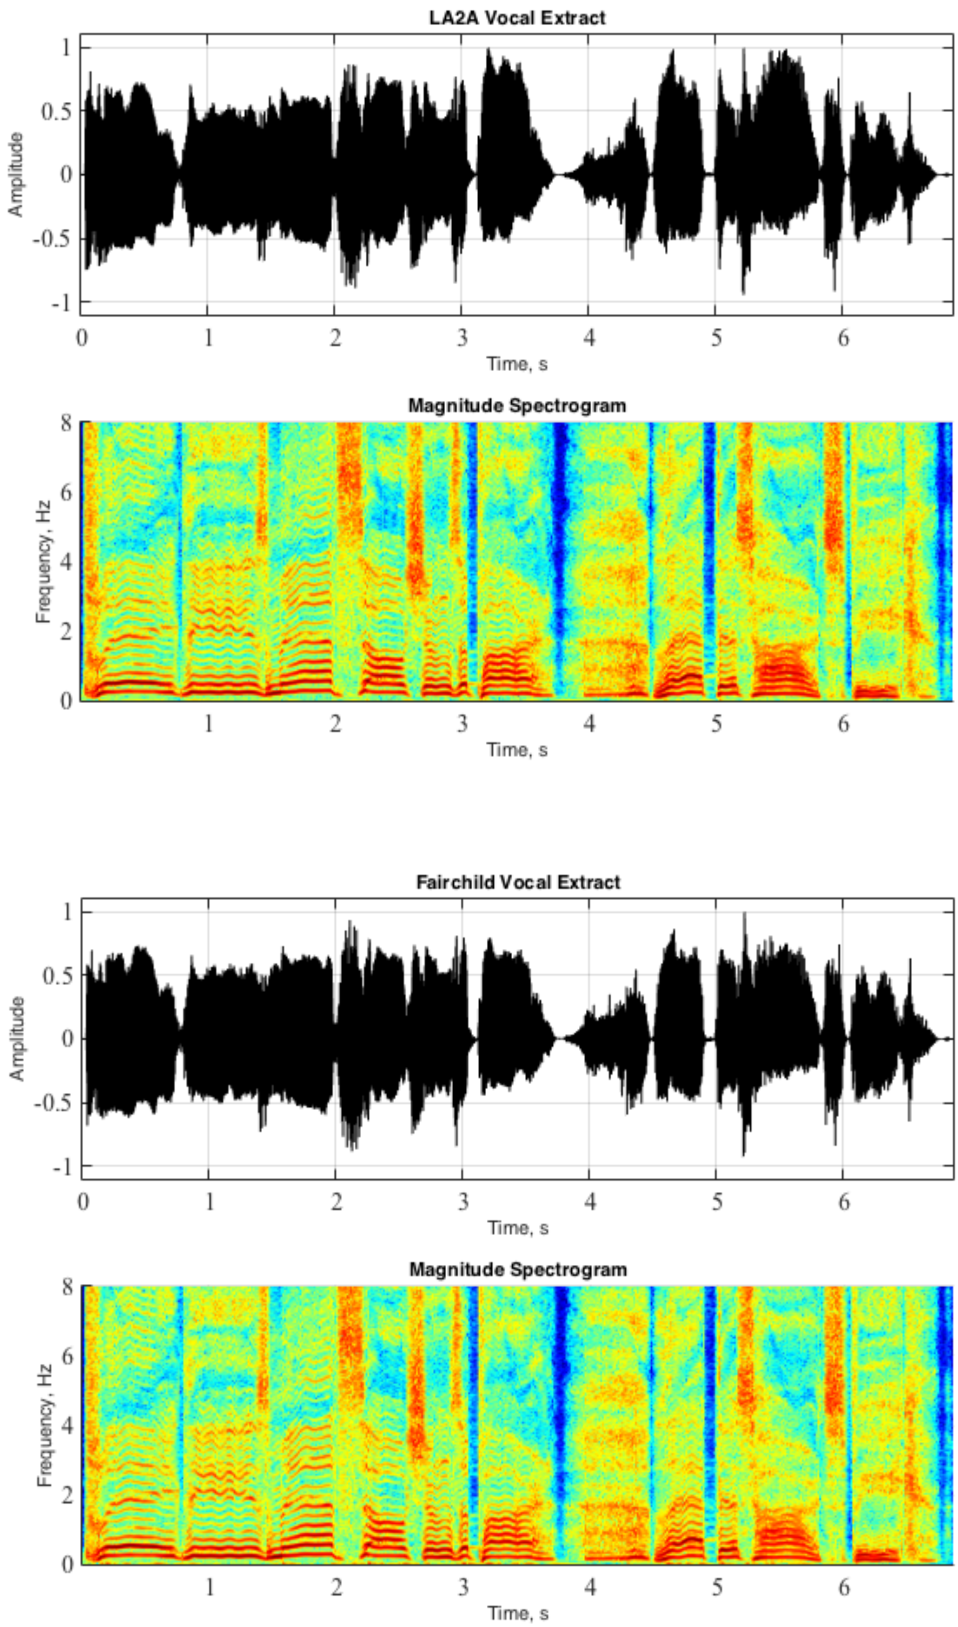

Figure 5-16: Time domain and spectrogram plots of the LA2A (top) and Fairchild (bottom) vocal material

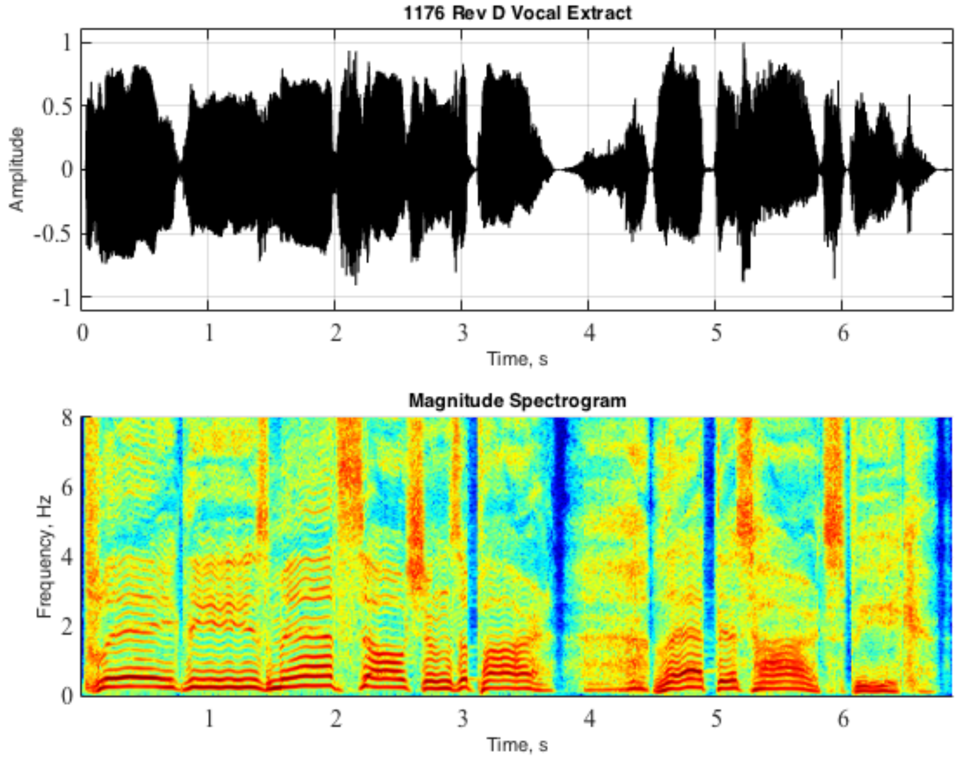

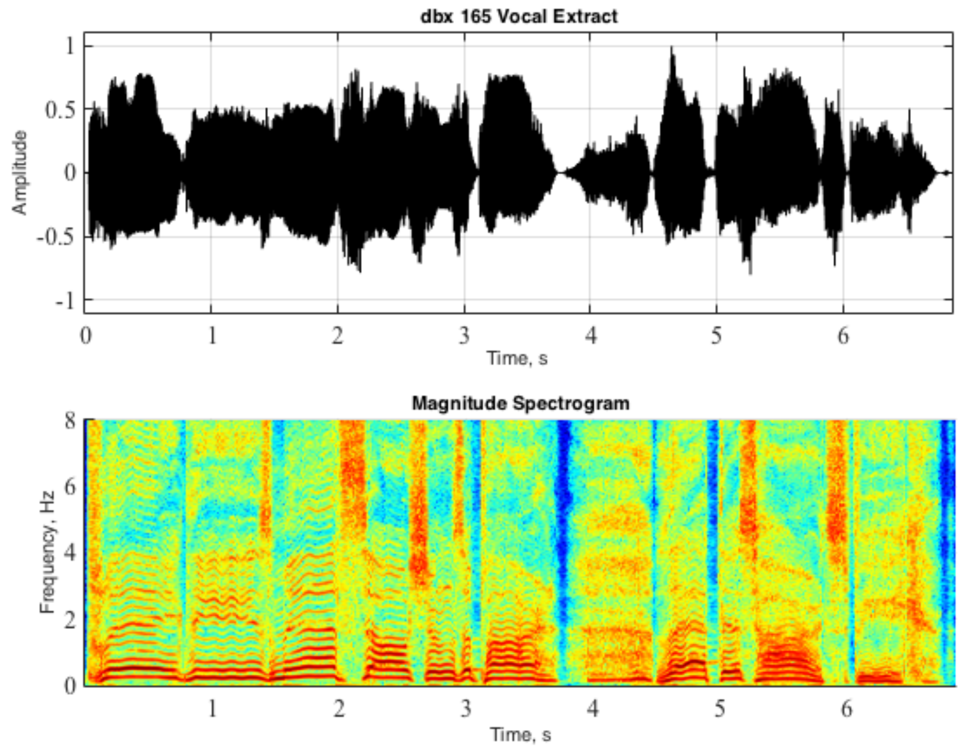

Figure 5-17: Time domain and spectrogram plots of the 1176 (top) and dbx165A (bottom) vocal material

The RMS energy envelope was calculated for the compressors and is shown in Figure 5-18. It can be observed the compressors have many similarities and only small differences in the microdynamics and around the start of phrases. The Fairchild and the 1176 are most similar in this plot and have a tendency to group closely together.

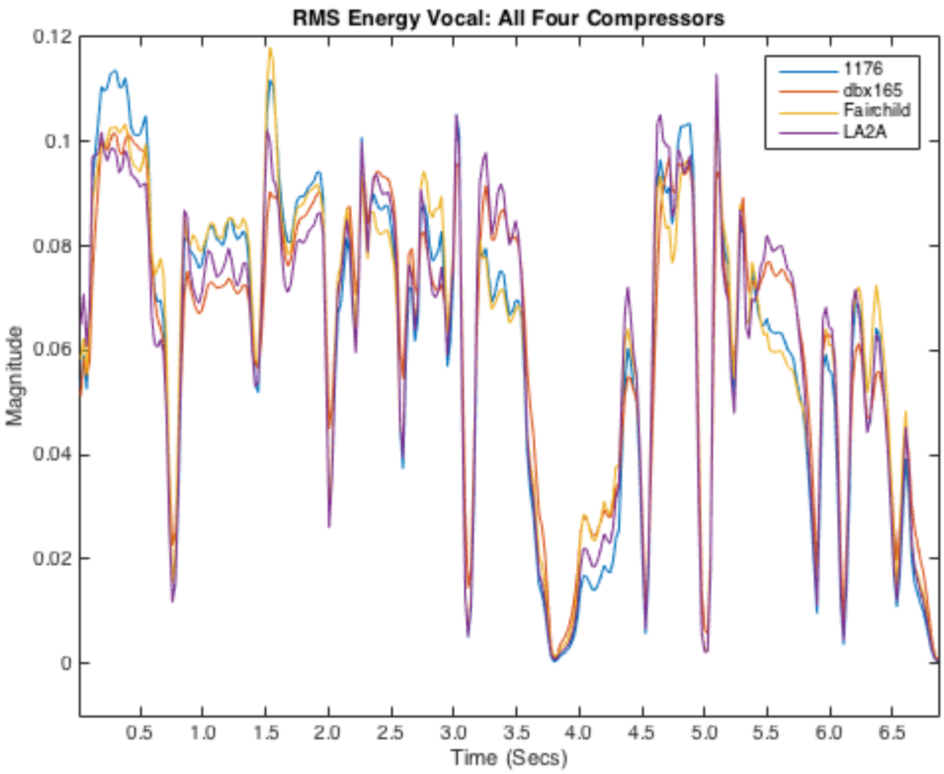

Figure 5-18: RMS energy for vocal material. All four compressors

The LTAS was calculated from the vocal material and can be seen in Figure 5-19. All compressors have a similar increase in the bottom, and top ends of the frequency spectrum and the LTAS curve has the compressors grouped closely suggesting there are only small frequency related differences between the units. The dbx165A has the brightest top end at about 1kHz in the plot, and this is because of the distortion the compressor is adding to the audio material

5.7.1 Audio Features: Vocal Material

Audio features were extracted from the signals and are shown in Table 5-9. The first results to note is all the compressors have similar low-energy results thus the temporal distribution of energy is comparable in all compressed material.

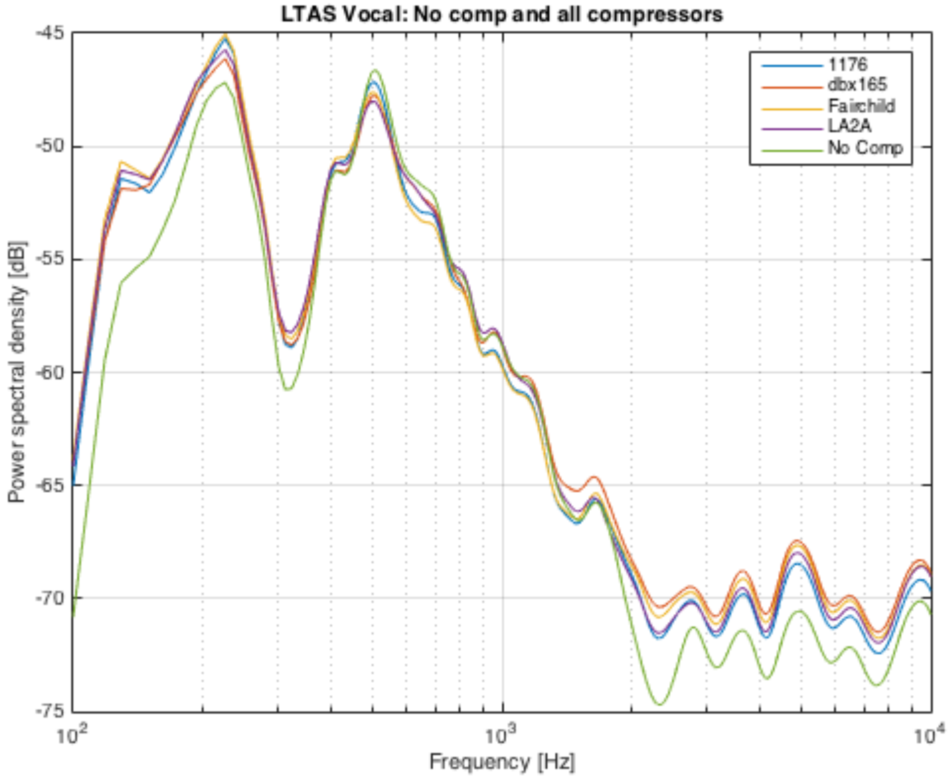

Figure 5-19: LTAS for vocal material. All compressors and uncompressed

Compressor	Roll Off	Spectral Centroid	Brightness	Low-Energy
Fairchild	10467 Hz	5075 Hz	0.61	0.41
LA2A	10540 Hz	5063 Hz	0.59	0.40
dbx165A	10500 Hz	5109 Hz	0.61	0.41
1176 Rev D	10361 Hz	4917 Hz	0.59	0.40
No Comp	10203 Hz	4655 Hz	0.55	0.53

Table 5-9: Audio features extracted from the vocal for all four compressors and the uncompressed material

Features relating to perceptual brightness are increased for all compressors but the variation between each of the units is small, and there is no clear pattern or trend. The Fairchild and dbx165A have marginally higher results for spectral centroid and brightness. The 1176, on the other hand, has a smaller spectral centroid value than the other compressors but it is not clear why.

Audio in the 200-400Hz frequency range was extracted using a filter bank, and spectral flux was calculated from this material. This area was chosen because it is the low end of the vocal frequency spectrum. The results can be seen in Table 5-10, which show a slight increase for the compressed audio. The dbx165A and the LA2A have the highest amount of spectral flux, and the 1176 is only marginally higher than the uncompressed audio. Again, the difference between compressors is small and apart from the distorted dbx165A, there is not a great deal of difference between the units when measured for spectral flux in this frequency range.

Compressor	Spec Flux Mean	Spec Flux Std
Fairchild	7.03	8.16
LA2A	7.37	9.21
dbx165A	7.40	8.22
1176 Rev D	6.89	8.51
No Comp	6.83	9.64

Table 5-10: Spectral flux for all compressors and uncompressed vocal material extracted from the 200-400Hz sub band

5.8 Results from Room Mic Material

Figures 5-20 (uncompressed) and 5-21 and 50-22 (compressed audio) depicts the results of room mic audio that was processed through the compressors, and a significant amount of variation is visible between the units with this source.

The dbx165A is aggressive in its application of compression. This effect is visible in the time domain plot and the amplitude statistics in Table 5-11 where it is shown to have severely restricted the dynamic range of the recording. There is much audible distortion in the audio and this combined with an increase of lower level components results in a spectrogram with significantly increased energy across the frequency range.

The Fairchild's amplitude plot resembles the 1176 Rev D, the two compressors sound similar, but the 1176 modulates the amplitude (pumps) more than the Fairchild. It is worth noting that the Fairchild was set to its fastest time constant that yields a fast 200-microsecond attack bringing it within the range of the 1176's attack.

The dbx165A has less of the audible pumping and is a more consistent block of audio, this is not to say that the dbx165A is not capable of achieving a similar sound quality to the 1176 but under these circumstances, it does not exhibit 1176 style modulation. Modulating effects like those offered by the 1176 are possible somewhere within the dbx165A's wide attack and release range but they are not as easy to implement as when using the 1176.

The LA2A has some of the pumping effects, but it is more reserved because of its program dependent release discussed in Chapter 3. Instead, the LA2A allows the transients of the audio material through which results in a punchier sound that is suitable for some production scenarios but not heavily compressed modulating styles.

Compressor	Dynamic Range (dB)
Uncompressed	20.55
LA2A	14.93
Fairchild	13.13
dbx165A	11.10
1176 Rev D	10.30

Table 5-11: Amplitude statistics for the room mic material

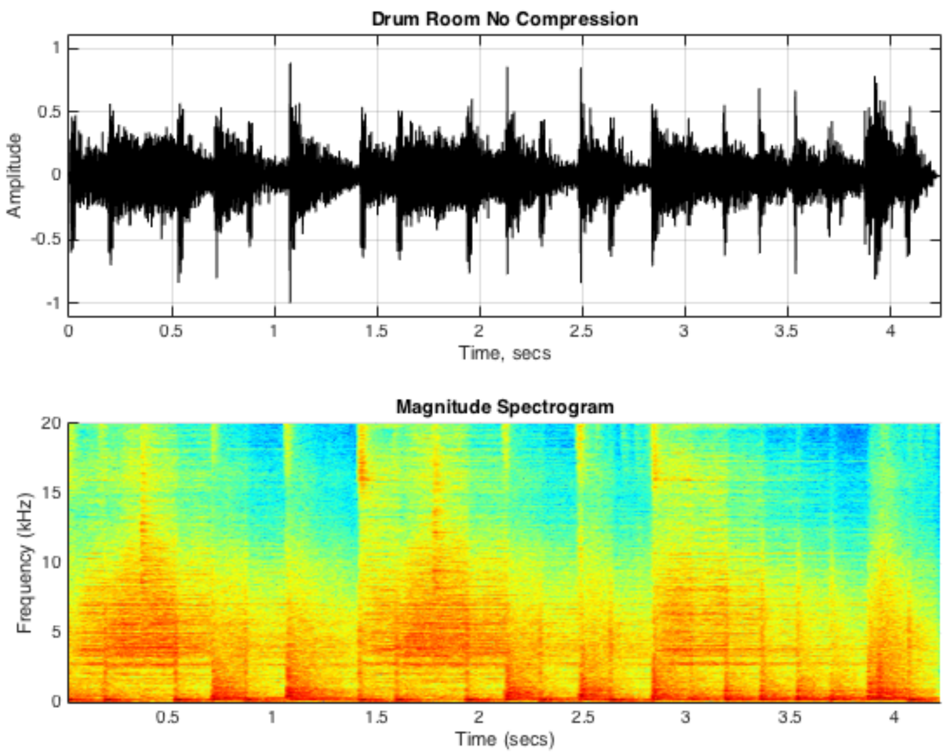

Figure 5-20: Time domain and spectrogram plots of the uncompressed room mic audio material

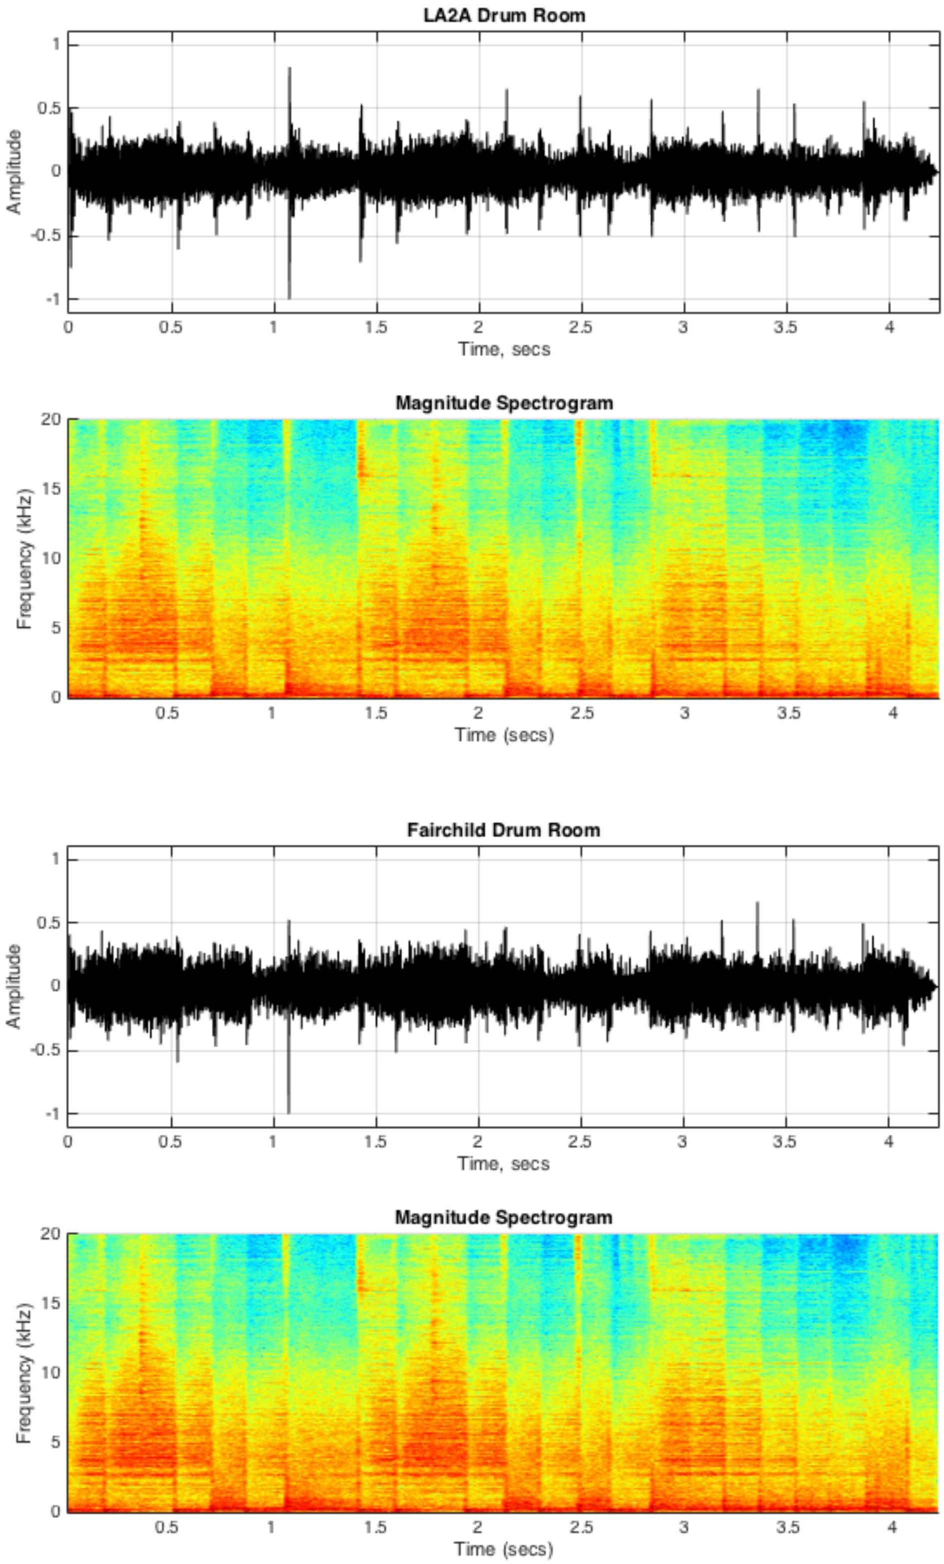

Figure 5-21: Time domain and spectrogram plots of the LA2A (top) and Fairchild (bottom) room mic material

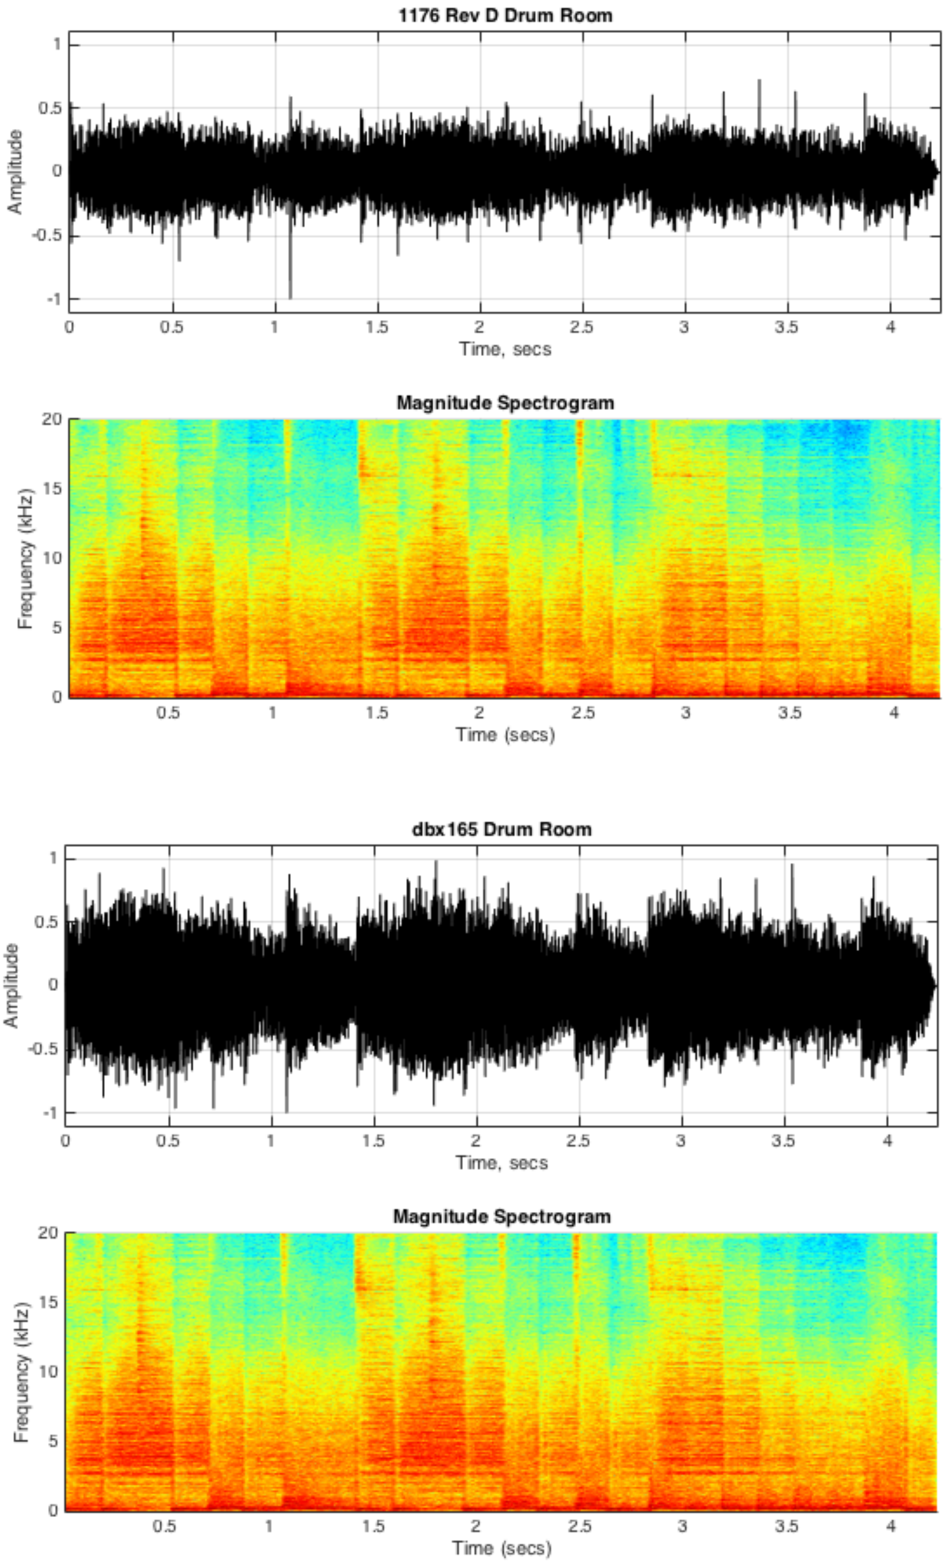

Figure 5-22: Time domain and spectrogram plots of the 1176 (top) and dbx165A (bottom) room mic material

Figure 5-23 shows RMS energy for all the compressors. The LA2A has the most variation in this plot followed by the Fairchild, the 1176 and finally the dbx165A. This plot demonstrates how the LA2A's slower attack time allows for punchier room mic compression while the other compressors act to attenuate the transient portion of the audio aggressively. It can also be observed that the LA2A does not raise low-level material between drum hits as much as the other units. The increase in the low-level detail is because of the faster release times of the other compressors.

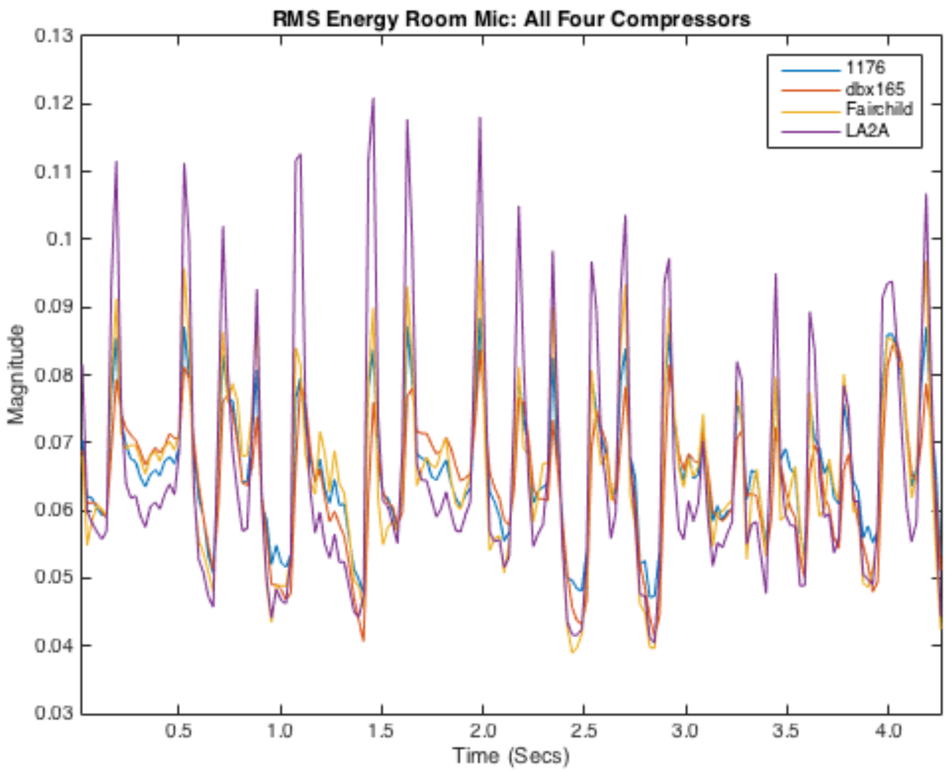

Figure 5-23: RMS energy for room mic material. All four compressors

Changes to the spectrum of the audio are shown in the LTAS plot in Figure 5-24 where it can be seen the compressors add energy to the audio particularly from 1kHz upwards. The 1176, dbx165A and Fairchild group together tightly, but the dbx165A has more energy due to the high levels of distortion it is producing with this setting.

Amplitude statistics were calculated for a section in the audio that featured a long crescendo drum roll and are presented in Table 5-12. They are shown to demonstrate how heavy some of the compressors have worked on the audio material.

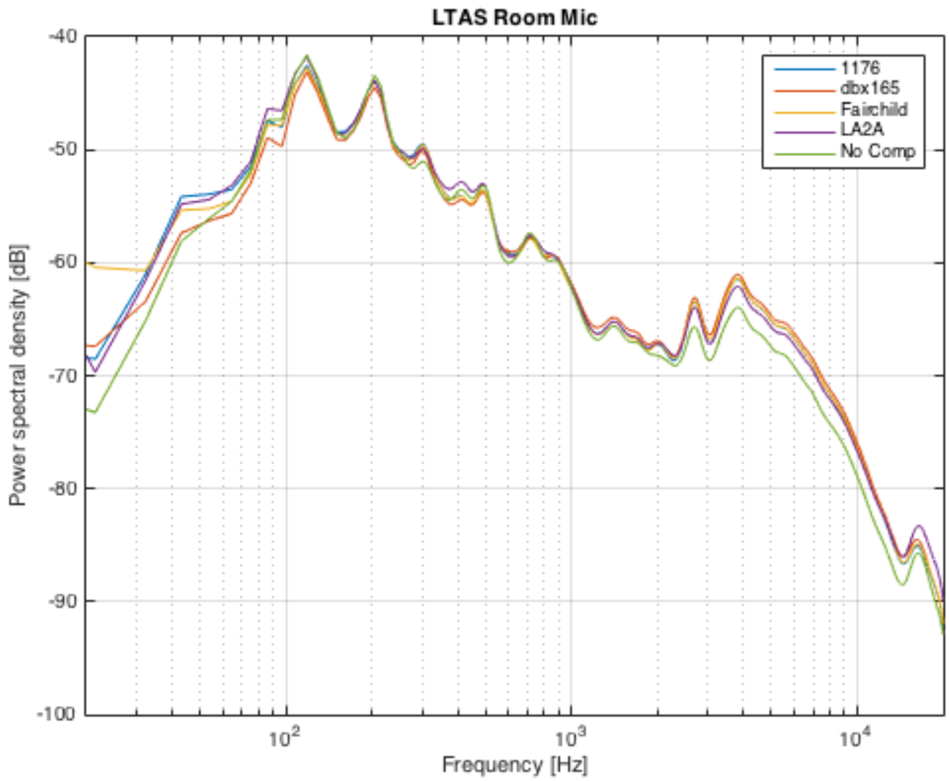

Figure 5-24: LTAS for room mic material. All compressors and uncompressed

The dbx165A has a very small dynamic range during this section and has reduced the uncompressed audio by 8.97dB. The 1176 has also reduced the dynamic range by a considerable amount, but it is approximately 1dB less than the dbx165A. The Fairchild is not working on the audio as aggressively as the other two compressors, and this is because of its slower time constants allowing for more variation in dynamics. The LA2A in comparison is the gentlest and this partly due to its slower attack time allowing transients from the drum hits through uncompressed. The results here are congruent with the information shown in Table 3-1 that presented the maximum amount of gain reduction offered by each compression style.

Compressor	Dynamic Range (dB)
Uncompressed	12.95
LA2A	8.56
Fairchild	6.99
dbx165A	3.98
1176 Rev D	4.99

Table 5-12: Amplitude statistics for the drum room extract

5.8.1 Audio Features: Room Mic Material

Table 5-13 presents the results from audio features that were extracted the room mic.

Compressor	Roll Off	Spectral Centroid	Brightness	Roughness	Low-Energy
Fairchild	8613 Hz	4598 Hz	0.68	403.9	0.47
LA2A	8936 Hz	4662 Hz	0.66	310	0.66
dbx165A	8700 Hz	4693 Hz	0.70	441.2	0.47
1176 Rev D	8574 Hz	4579 Hz	0.68	406.9	0.51
No Comp	8345 Hz	4296 Hz	0.62	212.9	0.64

Table 5-13: Audio features extracted from the room mic for all four compressors and the uncompressed material

In the table, the LA2A has the highest result for the low-energy feature. This data suggests it is the least consistent regarding energy distribution and this is due to the slow attack allowing uncompressed drum hits. The Fairchild and the dbx165A have similar ratings, with the 1176 creating slightly higher values that reflect the small microdynamic fluctuations and pumping that were described earlier.

ZCR was extracted from the material, and despite non-linearity being audible in the audio, it did not appear to correlate with the ZCR data. It was found the uncompressed audio rated higher for ZCR than the compressed audio. However, roughness correlated much better with the audible distortion thus it is a more appropriate feature to extract when measuring non-linearity of this kind. The roughness data correlates with perceptual levels of distortion, and for this feature the compressors were ranked, dbx165A, 1176, Fairchild, LA2A and uncompressed in highest to lowest order.

Spectral flux was extracted from the 100-200Hz band, and the results are significantly lower for all the compressors bar the LA2A. The compressors with the most moderate amounts of spectral flux are the compressors that offer the fastest time constants and the most aggressive pumping style. Spectral flux results can be seen in Table 5-14 where the LA2A has the highest value. This result indicates the LA2A is the most useful for room mic and drum buss compression styles that necessitate the preservation of transient hits.

As noted in Chapter 2 increased levels of spectral flux in high frequency bands is found to correlate with a sense of liveliness in sound sources. To that end, spectral flux was measured in the drum room audio from a band extracted between 1.6kHz-3.2kHz to measure if DRC was affecting this feature. It was anticipated this measure would have higher values for compressors with the most aggressive time constant behaviour and audible pumping. Table 5-15 shows the compressors processing the audio in this style and with fast time constants (all compressors except the LA2A) have the highest values, but the difference between them is subtle.

Compressor	Spec Flux Mean	Spec Flux Std
Fairchild	9.23	8.33
LA2A	10.24	11.70
dbx165A	8.90	6.80
1176 Rev D	9.08	7.85
No Comp	9.46	11.68

Table 5-14: Spectral flux for all compressors and uncompressed room mic material extracted from the 100-200Hz sub band

Compressor	Spec Flux Mean	Spec Flux Std
Fairchild	10.10	6.50
LA2A	9.41	5.72
dbx165A	10.64	6.05
1176 Rev D	10.35	5.87
No Comp	7.26	4.27

Table 5-15: Spectral flux for all compressors and uncompressed room mic material extracted from the 1.6-3.2kHz sub band

5.9 Differences in Gain Reduction

The previous sub-chapters focused on differences between the compressors when they were set for heavy amounts of gain reduction. For completeness, it was decided to carry out feature extraction on the medium and light gain reduction audio to observe any differences in timbre characteristics. Tables with audio features for all compressors over the three gain reduction amounts can be seen in Appendix 3 where there is some variation in the data. The manner in which this variation occurs is currently unclear, and it is thought that a much larger series of measurements is needed for more thorough data analysis. However, some observations can be gleaned from the current data set and are discussed in the following subchapters. Differences in spectral centroid and sub-band spectral flux data are addressed before a discussion of principle component analysis (PCA) to look for any similarities between the compressors over the three gain reduction amounts

5.9.1 Differences in Bass Material

Figure 5-25 depicts a general trend for the spectral centroid to increase with lower amounts of gain reduction, and this is most noticeable in the LA2A. There is a dip in spectral centroid for all compressors bar the LA2A between the high and medium amounts of gain reduction. Looking at the table of features in Appendix 3 reveals this slight dip in feature data between high and medium settings occurs in a number of the features. It is not clear what is responsible for this behaviour but it is a peculiar aspect of the compressors' sonic signature.

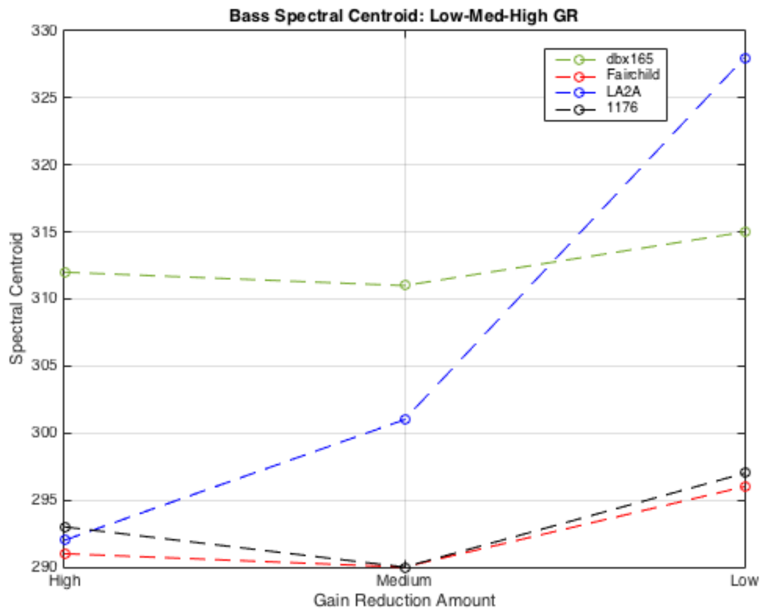

Figure 5-25: Spectral centroid for bass material and all compressors over three gain reduction amounts

Figure 5-26 illustrates the differences in spectral flux in the 100-200Hz band between the three amounts of gain reduction. The 1176, has a sharp drop between high to medium gain reduction, while the other compressors show a rise. The dbx165A, LA2A, and Fairchild have a slight increase from high to medium and then drop for light gain reduction. A decrease or increase in any of these features (in particular with the small amounts observed here) is not necessarily a positive or an adverse change. A further study containing a critical listening test should investigate the perceptual effect these small changes have on the listener.

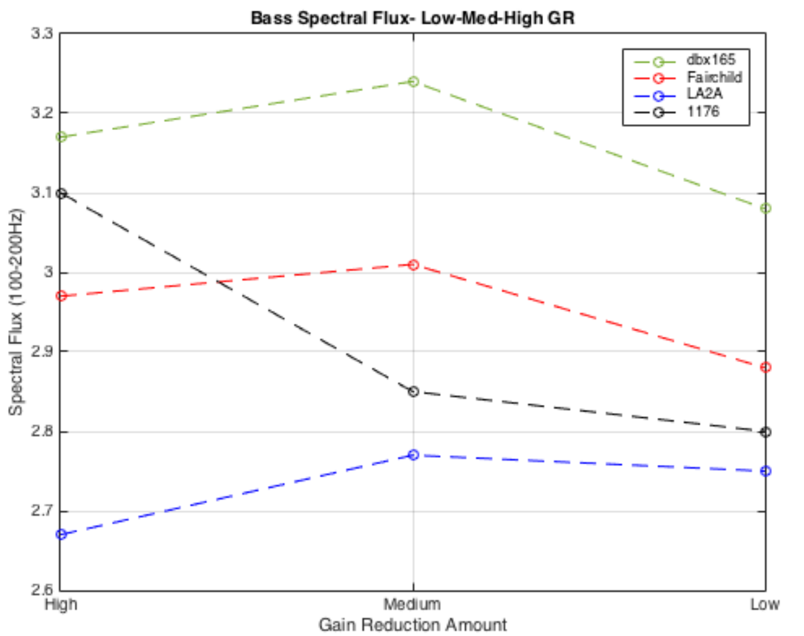

Figure 5-26: Spectral flux in the 100-200Hz band for bass material and all compressors over three gain reduction amounts

5.9.2 Differences in Room Mic Material

The plot for spectral centroid in Figure 5-27 indicates a reduction in value for the 1176 and dbx165A as the amount of gain reduction decreases. The LA2A and the Fairchild have an increase between high and medium settings but drop for low gain reduction. Out of all the compressors, the dbx165A has the most amount of variation, and this is due in part to the high order non-linearity it introduces with heavier amounts of compression diminishing as the gain reduction is reduced.

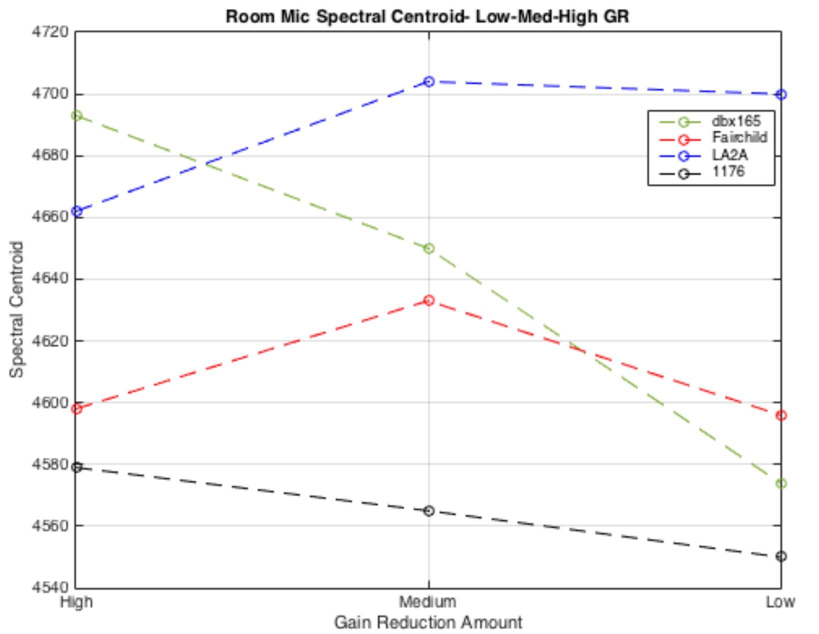

Figure 5-27: Spectral centroid for room mic material and all compressors over three gain reduction amounts

Spectral flux between 100-200Hz is presented in Figure 5-28. As shown in the figure there is a slight upward trend for the 1176 and dbx165A. On the other hand, the LA2A has a slight decrease between high and medium amounts of gain reduction and the Fairchild has an increase from low to medium gain reduction, before decreasing with higher values.

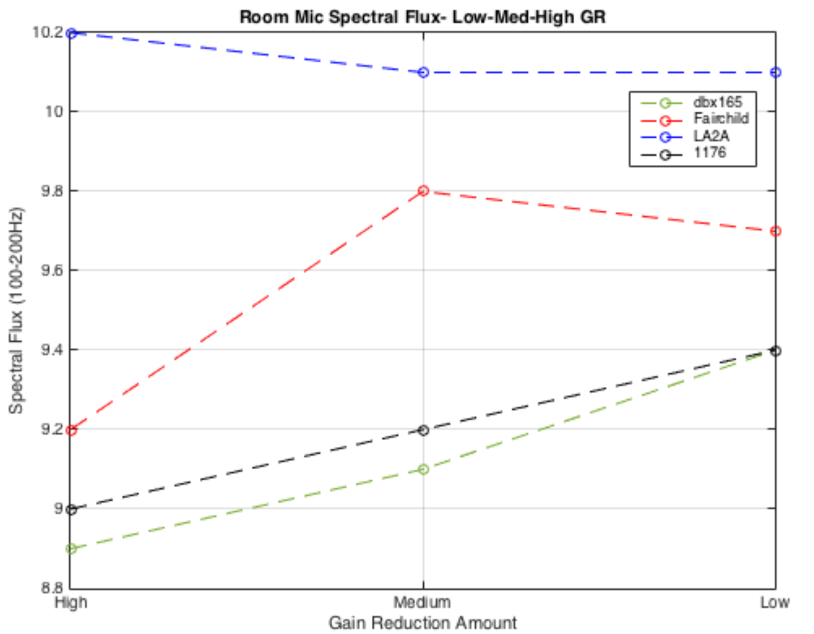

Figure 5-28: Spectral flux in the 100-200Hz band for room mic material and all compressors over three gain reduction amounts

Figure 5-29 plots Spectral flux in the 1.6-3.2kHz band and depicts a decrease for the 1176 and the dbx165A, a rise and a drop for the Fairchild and a more consistent result for the LA2A.

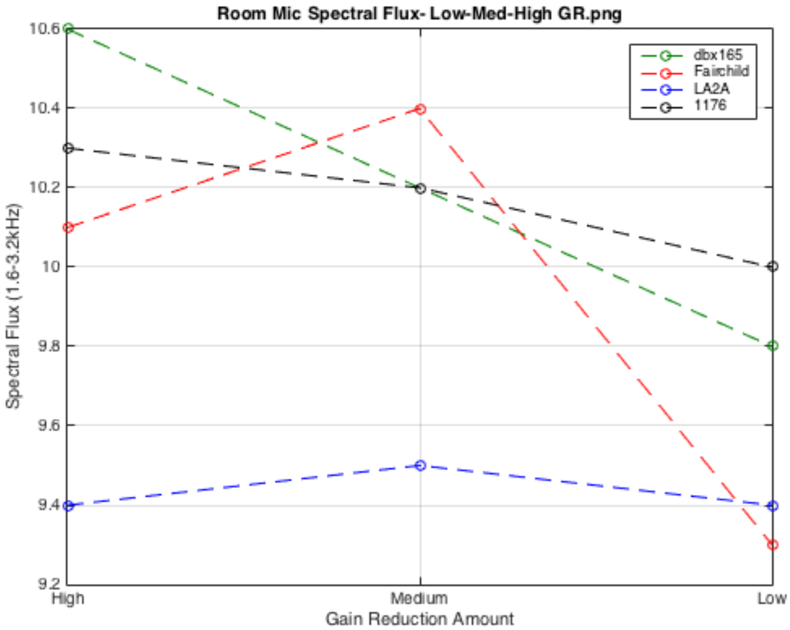

Figure 5-29: Spectral flux in the 1.6-3.2kHz band for room mic material and all compressors over three gain reduction amounts

5.9.3 Difference in Vocal Material

As illustrated in Figure 5-30 there is a general decrease to the spectral centroid of vocal material as the amount of gain reduction is decreased. The LA2A and the 1176 have a small rise between the medium and high gain reduction settings but then drop between medium and low settings. The increase in spectral centroid is slight and presumably will have little effect on the sonic signature.

Figure 5-31 shows Spectral flux in the 200-400Hz band, and a similar trend can be observed, all the compressors have lower values for this feature as the amount of gain reduction is reduced. As with spectral centroid, the 1176 has a slight increase in value between heavy and medium settings before decreasing as the gain reduction is lowered.

5.10 Conclusions on Different Amounts of Gain Reduction

In conclusion, it has been demonstrated that given the same input material there are slight changes to the audio features when varying the amount of gain reduction. The changes discussed are relatively small, and it is unclear how much an effect they will have on the perceptual qualities of each unit's sound quality. Nonetheless, the changes are present in the audio material and are an intrinsic aspect of each compressors sonic signature.

.

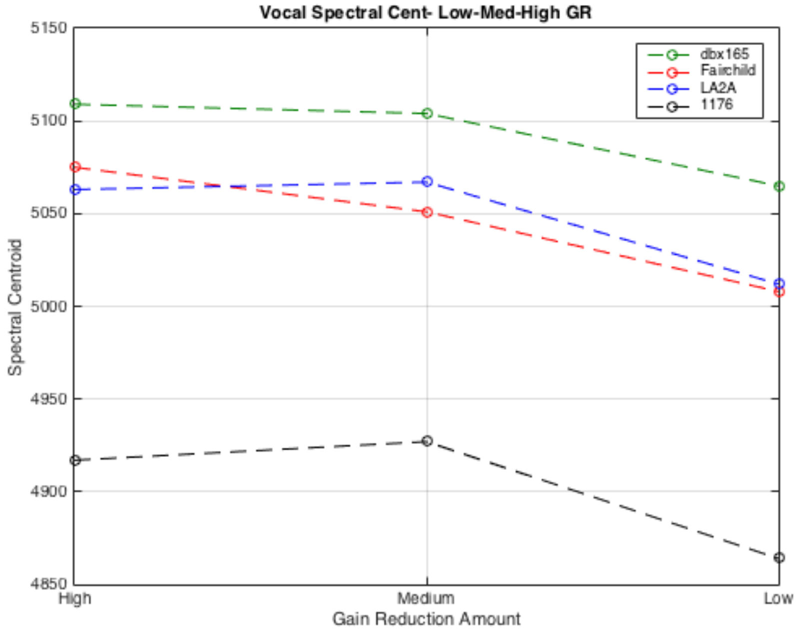

Figure 5-30: Spectral centroid for vocal material and all compressors over three gain reduction amounts

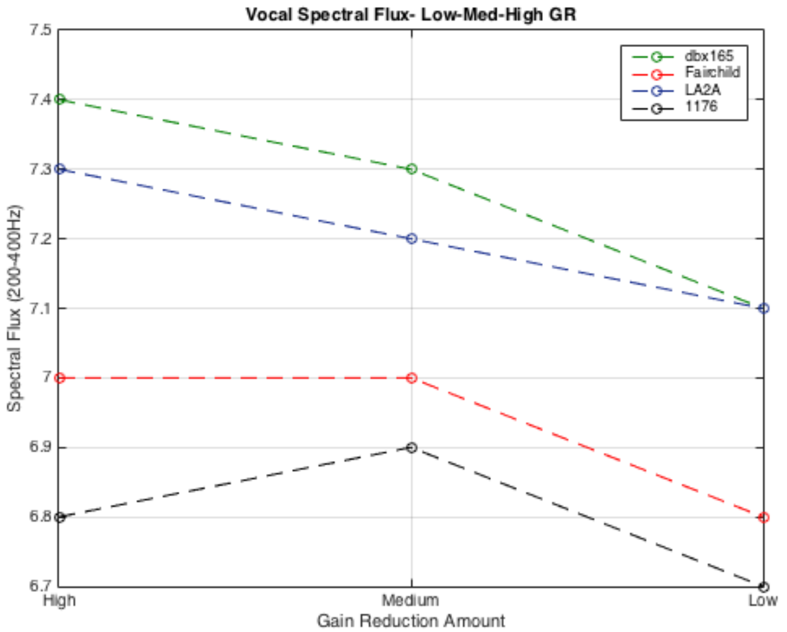

Figure 5-31: Spectral flux in the 200-400Hz band for room mic material and all compressors over three gain reduction amounts

5.11 Principle Component Analysis of Audio Features

One final piece of research was carried out on the audio to assess how the compressors differed. This exercise was conducted using principle component analysis (PCA), a process detailed in Chapter 2. The statistical package R (R Core Team, 2016) and the prcomp function from the stats library was used for PCA. Plots were created using the ggbiplot function from the devtools library (Wickham & Chang, 2016). As the audio features were not measured on the same scale or using the same unit of measurement the data was scaled using the prcomp function. Normalizing data before carrying out PCA is important because the process is based around finding data with the most variation. If data is not normalized, the data with the largest values and most variance will take up the majority of the first principle components creating an unrepresentative result.

The data used for PCA analysis can be seen in a table in Appendix 3 and includes all features for all compressors and uncompressed material over the three gain reduction amounts. The scree plots (a visual representation of the percentage of variance in each principle component) and the loadings (a term used for the weightings of the variables) for each of the principle components is included in Appendix 1. The current study focuses on PCA score, plots but PCA score plots with loadings can be seen in Appendix 2 if the reader wishes to look at this information.

5.11.1 Bass Material: PCA Results

The score plot for the bass material can be seen in Figure 5-32 that shows the uncompressed material located a distance from all the compressed audio, although it is similar to LA2A medium and heavy gain reduction and 1176 light on PC2. The LA2A light setting plots away from all the other program material suggesting that with light settings the LA2A takes on some unique characteristics. The same observation can be made for all the compressors, with light gain reduction versions positioned lower on PC2. The dbx165A medium and heavy settings cluster tightly while the light gain reduction setting is located further away on PC2. This appears to be as a result of lower roughness values that have a positive loading on PC2. 1176 heavy setting clusters with both Fairchild heavy and medium settings, suggesting a similarity between these compressors when they are set with specific amounts of gain reduction. The settings that cluster towards the negative end of PC2 (particularly the LA2A light setting) are located here because of larger spectral spread results.

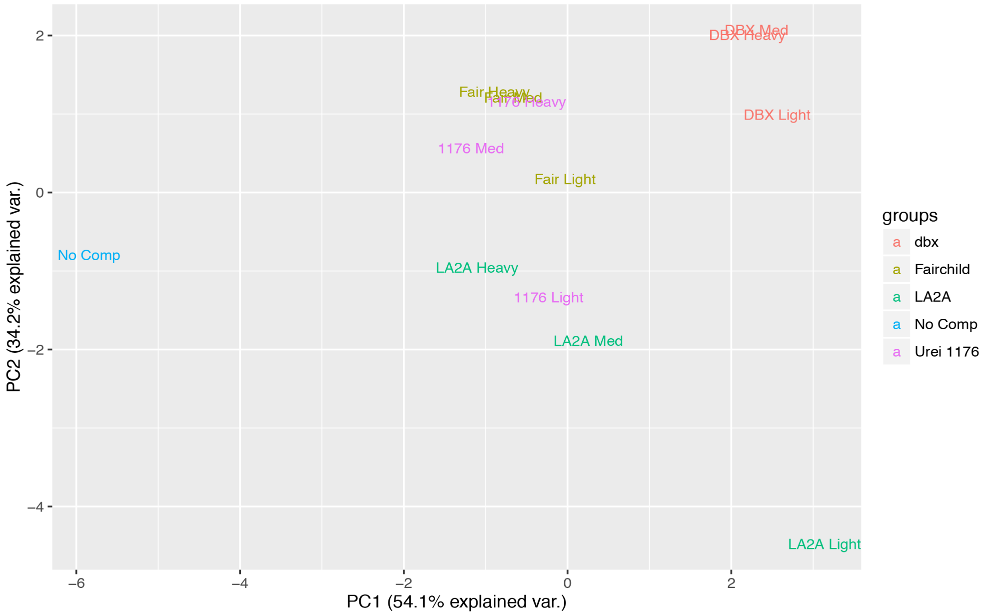

Figure 5-32: PCA score plot for bass material

5.11.2 Vocal Material: PCA Results

The score plot for the vocals in Figure 5-33 shows the compressed vocal material clusters a considerable distance from the uncompressed audio. The compressors fall into discernible groups, and the light versions of the compressed material for each compressor are positioned further away than the heavy and medium version on both principle components. This result suggests that with lower amounts of gain reduction all compressors take on a sonic signature that is more different than the medium and heavy settings. The loadings in Appendix 2 show that program material, which is located towards the negative end of PC1, is positioned there because of lower 100-200Hz spectral flux values. Material located towards the positive end of PC2 is located there because of higher low-energy values. This data is also the reason why the uncompressed audio is located positively at the extreme of the PC1 axis.

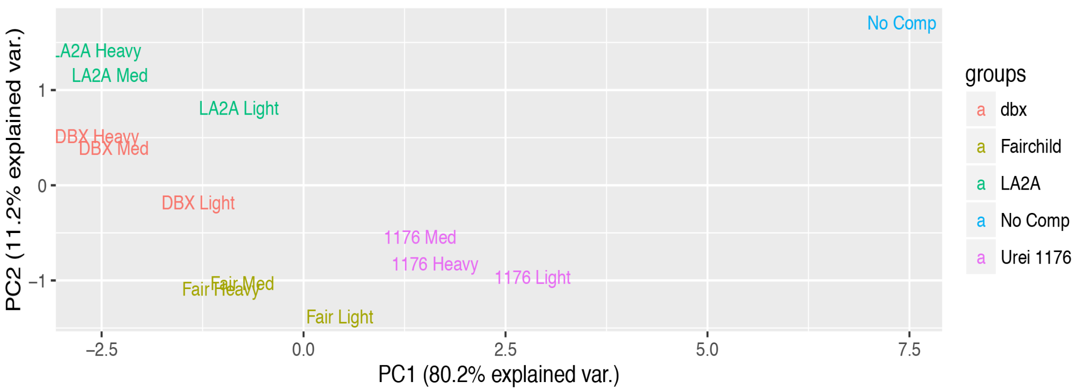

Figure 5-33: PCA score plot for vocal material

5.11.3 Room Mic Material: PCA Results

The score plot for the room mic material in Figure 5-34 shows differences in the compressors with some variation in where they plot as the gain reduction is altered. All the compressors vary mainly on PC2. The loadings in Appendix 2 show that lower values for roughness and 1.6-3.2kHz spectral flux are responsible for the lower gain reduction material being positioned towards the negative end of PC2. The Fairchild heavy setting is plotted to the positive side of PC1 and this is due to a higher spectral kurtosis value for this setting

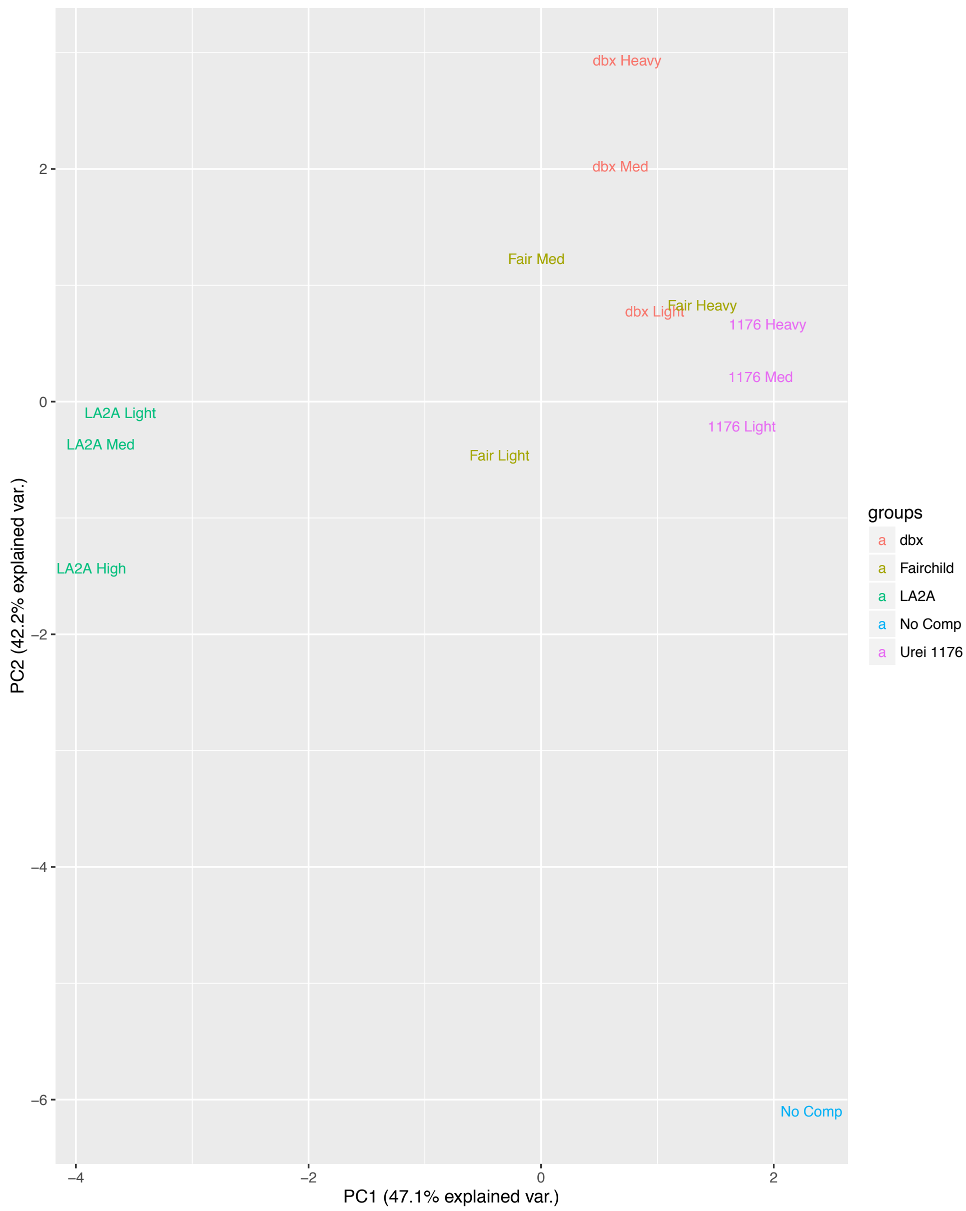

Figure 5-34: PCA score plot for room mic material

5.12 Conclusions from the Chapter

This study set out to determine the sonic signature of four classic compressors and compare and contrast the results. Many noteworthy findings were discovered and are presented in this summary.

For bass audio material, all compressed signals have an increase in the perceived brightness of program material. This is often because of the non-linearity introduced by the compressor during gain reduction and because of fast time constants. This sound quality is particularly evident in the audio processed through the dbx165A and to a lesser extent the 1176. The compressors impart unique temporal changes to the overall envelope of the audio material and to the microdynamics of individual notes.

All compressors increase the spectral flux of the audio in lower bands, and this correlates with perceptual fullness. The dbx165A has a higher result for this measurement, and it is thought that this is because of its non-linearity. All compressed signals have a lower low-energy rating, which suggests that the temporal distribution of the energy is more evenly distributed under DRC. The 1176 and the dbx165A have lower ratings than the Fairchild and LA2A. This effect creates a greater sense of perceived fullness. Analysis of the light and moderately processed files reveals small differences between the compressors and the heavy and medium settings tend to group together tightly when using PCA.

A limitation of the bass study is that it only used one style of bass part played in a relatively simple yet consistent rock style. Processing a wider range of bass styles through the compressors may expose additional qualities of the sonic signatures.

Critical listening of the vocal audio reveals that when working under moderate conditions, all of the compressors share similarities. When pushed into extreme gain reduction the following aspects of their sonic signature becomes audible. The dbx165A works aggressively on the signal to restrict its dynamic range more than the other compressors. It adds non-linearity to the signal that sounds like a pronounced midrange. The LA2A has the least aggressive compression signature, and despite having a restricted dynamic range, it sounds less compressed and "in your face" than the other compressors. The Fairchild and the 1176 have a comparable overall sound quality and similarly shape the vocal envelope. All the compressors bar the LA2A have the effect of aggressively raising the level of breaths in the vocal performance. The LA2A creates this effect also, but the effect is not as pronounced as with the other compressors. This manipulation of the breaths in a performance is a desirable characteristic amongst engineers who wish to add an aggressive character to their vocal productions (Case, 2007; Gearslutz, 2014).

Extraction of audio features reveals little difference between the compressors, but they all have the effect of increasing the perceived brightness and spectral centroid of the original audio file and increasing levels of spectral flux between 200-400Hz. This correlates with fullness in audio material and is most significant in the LA2A. Results from PCA shows that the heavy and medium settings for each compressor group together closer than the light settings. This result demonstrates the compressors have a different character when lightly compressing.

The vocal study was limited to only one style of vocalist, and while it is thought these results will transfer to other vocalists in the rock and pop genres, it would be interesting to assess the effects of the compressors on other vocal styles.

In summarizing the tests made using drum audio the following points can be made. Under heavy room mic compression, the 1176 exhibits some unique modulation of the envelope that manifests as pumping and subtle variations in microdynamics. The Fairchild is similar in sonic signature to the 1176 but its time constants cannot achieve the same pumping effect as the 1176 when used in all- buttons mode. The dbx165A achieves more extreme gain reduction than the other compressors and radically flattens the amplitude and introduces audible non-linearity. The LA2A does not introduce any of the non-linearity or modulating characteristics of the other compressors. Its slower attack time has the effect of allowing drum hits through uncompressed resulting in transient heavy room mic compression. This type of processing was not considered as desirable as pumping modulation in Chapter 4.

All compressors have the effect of increasing the spectral centroid and brightness measures. The dbx165A has the most significant change to the audio in this regard, and this is due to the high levels of non-linearity it adds to the audio. PCA reveals that the LA2A clusters further away from the other compressors on the score chart. This result highlights it has a different sonic signature than the other compressors in this test.

The room mic study was limited to a specific type of room mic processing, and it would be interesting to investigate how the compressors respond under different circumstances, particularly when processing a mono room mic to increase the perception of punch and articulation of transients.

To conclude, the results presented in this chapter add substantially to our understanding of the sonic signatures of the compressors tested and add to the body of literature available to researchers in music technology related subject areas. Future work should use a wider range of sources and contexts to develop a fuller picture of each compressor's sonic signature and build upon the methodology established here.

Chapter 6 : An In-Depth Investigation into the 1176

6.1 The 1176 Analysis: Introduction

Although some testing on the 1176 was carried out in the previous chapter a more focused examination of its sonic signature is now presented. This section provides a full characterisation of the 1176's sonic signature and is a case study example of a methodology for sonic signature analysis. The study includes the use of IMD test tones and a method to test the 1176's program dependency.

The analysis is conducted on audio processed with a range of time constant settings to get an overview of the characteristics of the 1176 and how its sonic signature changes with different time constant settings. The specific settings used for each source will be detailed in subsequent sub-chapters, but they were gleaned from the analysis carried out in Chapter 4. The complex program material used for testing is the same as the material used in the previous chapter.

The audio was compressed using two separate 1176s simultaneously with the same time constant settings and they were adjusted to have the equivalent amount of gain reduction showing on the VU meters. The reason for this approach was to compare differences between two 1176s when set with the same settings and analyse the variation in sonic signature between the units. The differences between the compressors is not intended to be the main focus of the testing but is included to get an idea of how much variation there is between the two models when set in a like-for-like configuration. The 1176s tested are a modern Universal Audio 1176 reissue and an original Urei 1176 Revision D Blackface. Robjohns (2001) describes the Blackface as the best sounding 1176 revision in use. During this chapter, the Universal Audio reissue 1176 is called the Reissue, and the Urei 1176 is called the Blackface.

6.2 Test Tone Methodology

The following sub-chapters discuss the results of tests made using test tones. The testing was conducted in accordance with the standard distortion measurements detailed in Chapter 2. The measurements discussed in this chapter include hotter input levels to assess THD, IMD measurements, measurements made at different time constant settings and program dependency measurements. Program dependency was tested on the 1176's using a method detailed by Berners and Abel (2004) of Universal Audio that will be discussed in due course.

The measurements were made using Prism Sound ADA-8XR converters, and the audio material was recorded at 44.1kHz 24bit resolution. The input level of all test tones and complex program material was monitored using a Phonic AA3 audio analyser to ensure consistency of levels between tests. The audio portions analysed in the following chapters have been extracted and are included as discrete files in the data drive that supports this thesis.

6.3 Testing Program Dependency: Two 1176s

The program dependent nature of the 1176 was explored using a sine burst consisting of a 10ms transient followed by a 500ms portion of lower level audio. This was then followed by 2000ms of higher-level audio and finished with 2000ms of lower-level amplitude. The higher-level portions were 8dB louder than the lower level and the compressor was set to apply gain reduction only during louder levels. This test was designed to investigate the difference in time constant behaviour between transient and steady state material. The test signal can be seen in Figure 6-1 and is based on the aforementioned test by Berners and Abel (2004).

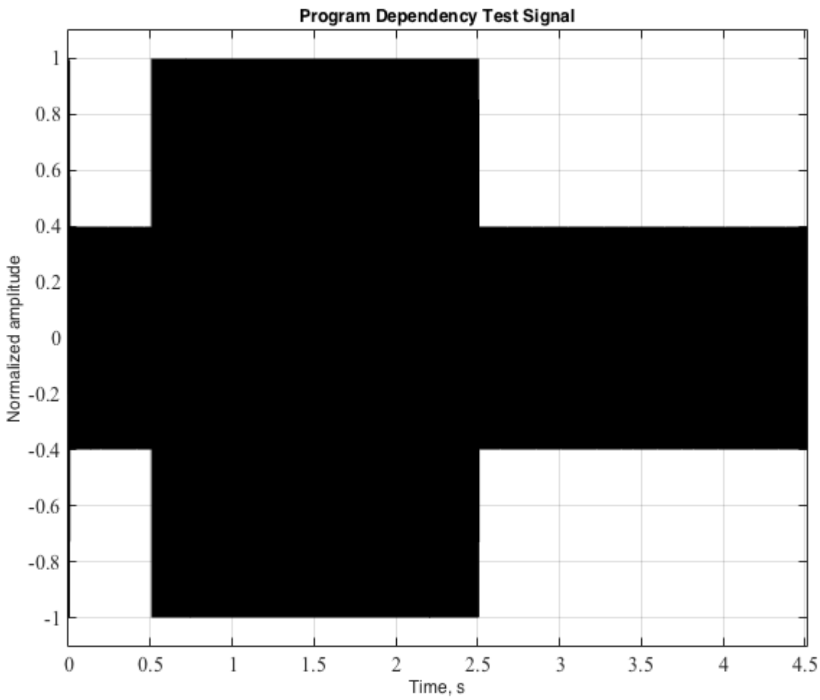

Figure 6-1: Program dependency test tone

The test tone was sent to the compressor at 1kHz and 50Hz. The 1kHz tone was used as it is a standard in test measurements. The 50Hz tone was selected to observe how the compressors behaved with low frequency material in a range similar to that of a bass guitar and detect non-linearity added to low frequency signals when using fast time constants. The compressors were adjusted to have 10dB of gain reduction.

The results for the 1kHz tone are illustrated in Figure 6-2 on the next page with the Reissue on the top and the Blackface on the bottom. As can be seen, there is a difference in how the time constants respond to the transient material, the release is much quicker after the transient when compared to the steady state portion where the release time is noticeably longer. This type of behaviour is desirable in a compressor to avoid unwanted amplitude modulation in complex program material and is an integral component of the 1176’s sonic signature.

When comparing the two compressors, there is a difference in the attack and release curves. The Blackface has a gradual attack after the overshoot at 0.5 seconds. The Reissue has a longer and gentler release curve and is slower to release over the period from 2.5 to 4.5 seconds while the Blackface is back to a steady level by 3.5 seconds. Settings used for this test had attack set at 3 and release at 7 with a 4:1 ratio, therefore, it is not possible to adjust the Reissue to match the release speed of the Blackface (7 is the fastest release time offered by the 1176). This timing behaviour may translate into a more pronounced form of compression under normal studio-based working conditions, and the Blackface may be perceived as having a more aggressive sonic signature.

It can also be seen there is some difference between the amount of gain reduction and level at the output of the compressors. The Blackface is attenuating the audio more, despite both compressors being set to the same ratio and reading the same amount of gain reduction on the meters. Additionally, the Reissue is applying more output gain to the signal after compression, again despite both compressors being adjusted to have the same output on their meters.

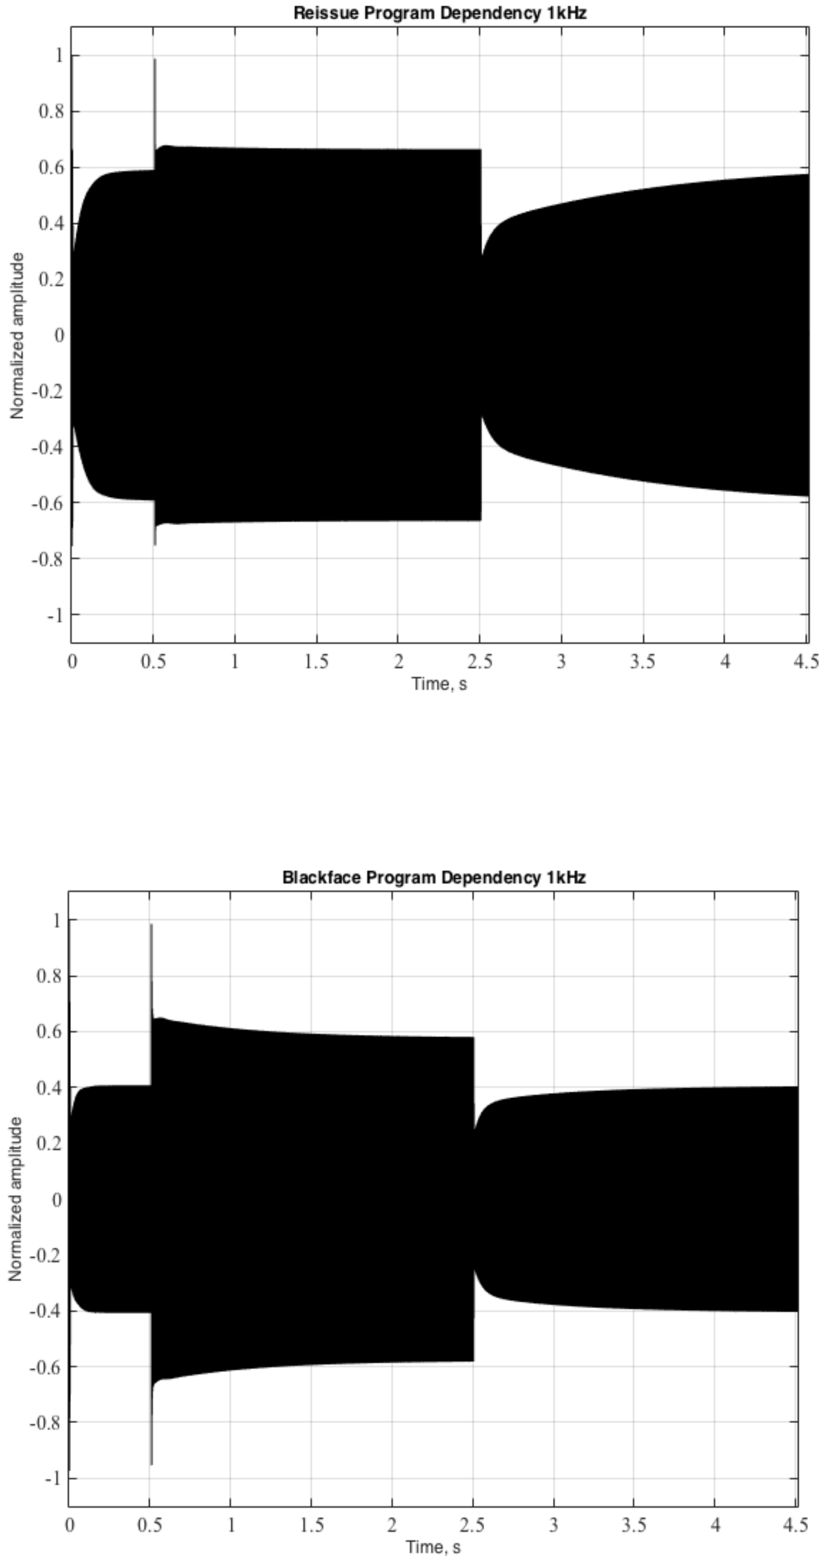

Figure 6-2: Results from program dependency 1kHz tone. Reissue on top and Blackface on bottom

To investigate the difference between the compressors further, the plots in Figure 6-3 show the compressors set with the attack and release at their fastest position and the ratio at 20:1. This configuration was one of the settings used on bass material in a test described later. It was used to look for non-linearity when using fast attack times. The results are consistent with the previous time constant setting. The Reissue has a more aggressive attack, but the release curve of the Blackface is noticeably quicker. There is some room to adjust the attack speed of the Reissue to match that of the Blackface somewhat but its release speed cannot be matched by the Reissue (it is at 7, the fastest position).

The same test signal but using a 50Hz tone was sent through the compressors with the fastest time constants and ratio at 20:1. The results depicted in Figure 6-4 reveals the Reissue's timing curve is not tracking the changes as effectively as the Blackface. The portion of the test tone making up the first 500ms has had its envelope radically reshaped resulting in much smaller variation in level between the material in the first 2.5 seconds. Looking closer at the plot shows the fast transient of the test tone has overshot the compressor and the quiet portion after the transient has been aggressively ramped up in level. This result from the Reissue is unexpected, and by comparison, the Blackface is behaving more predictably, the shape of the test tone at the Blackface's output is much like one would expect. This result highlights how there can be some unexpected variation in compression behaviour between two models of the same compressor, which in turn will have an impact on the unit's sonic signature.

Listening to the audio reveals the first 500ms of the Reissue's output is significantly more distorted than the Blackface, which is clean by comparison. This difference in non-linearity can be observed in the plots in Figure 6-5 that shows the FFT of the first 500ms of the test tone through the compressors. Note the considerable amount of distortion in the Reissue's plot compared to the cleaner result from the Blackface on the bottom.

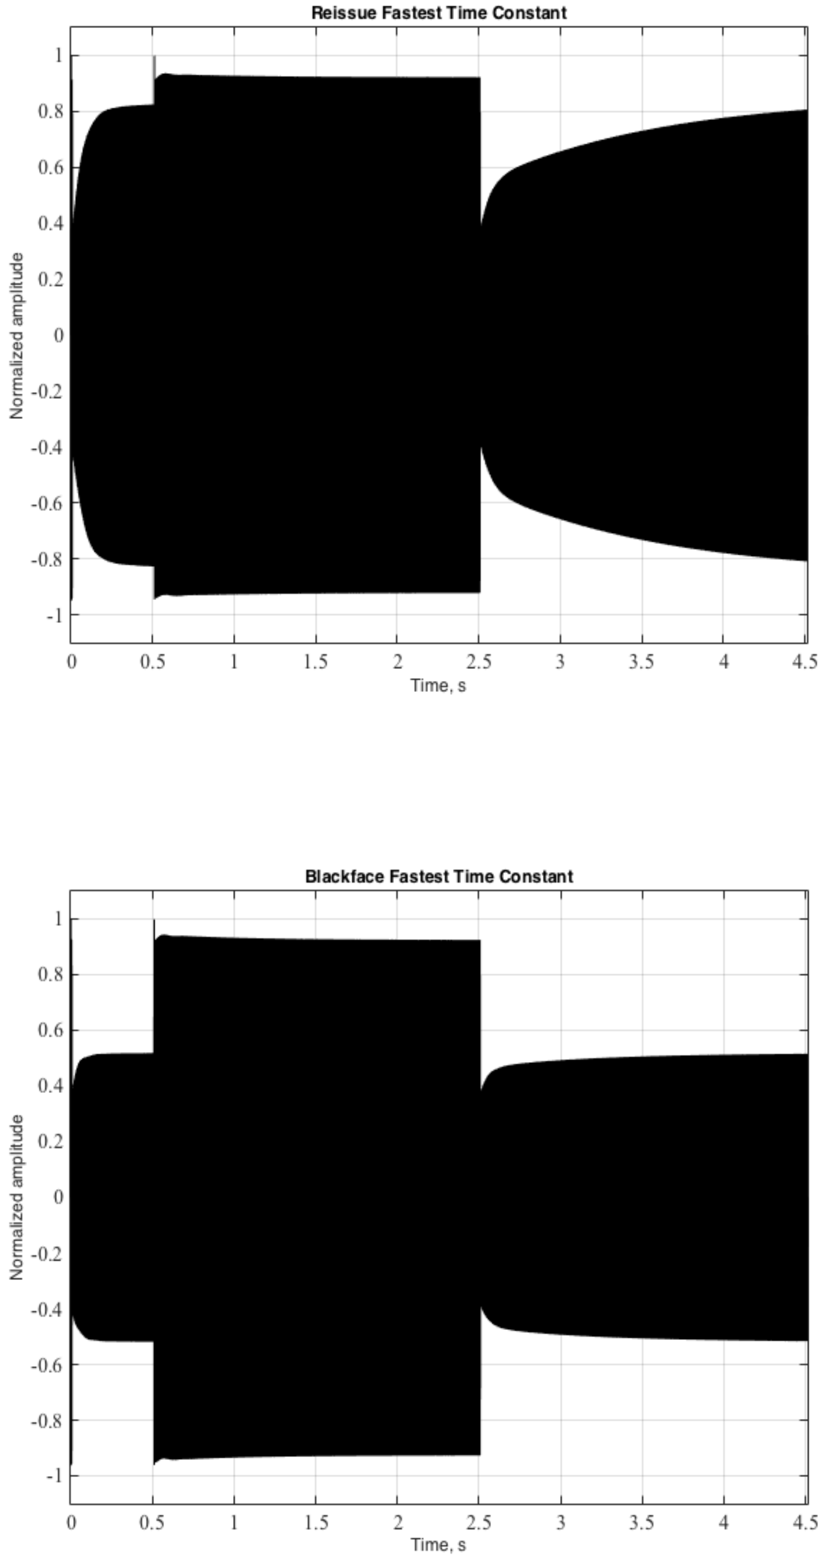

Figure 6-3: Results from program dependency 1kHz tone using the fastest attack and release speeds. Reissue on top and Blackface on bottom

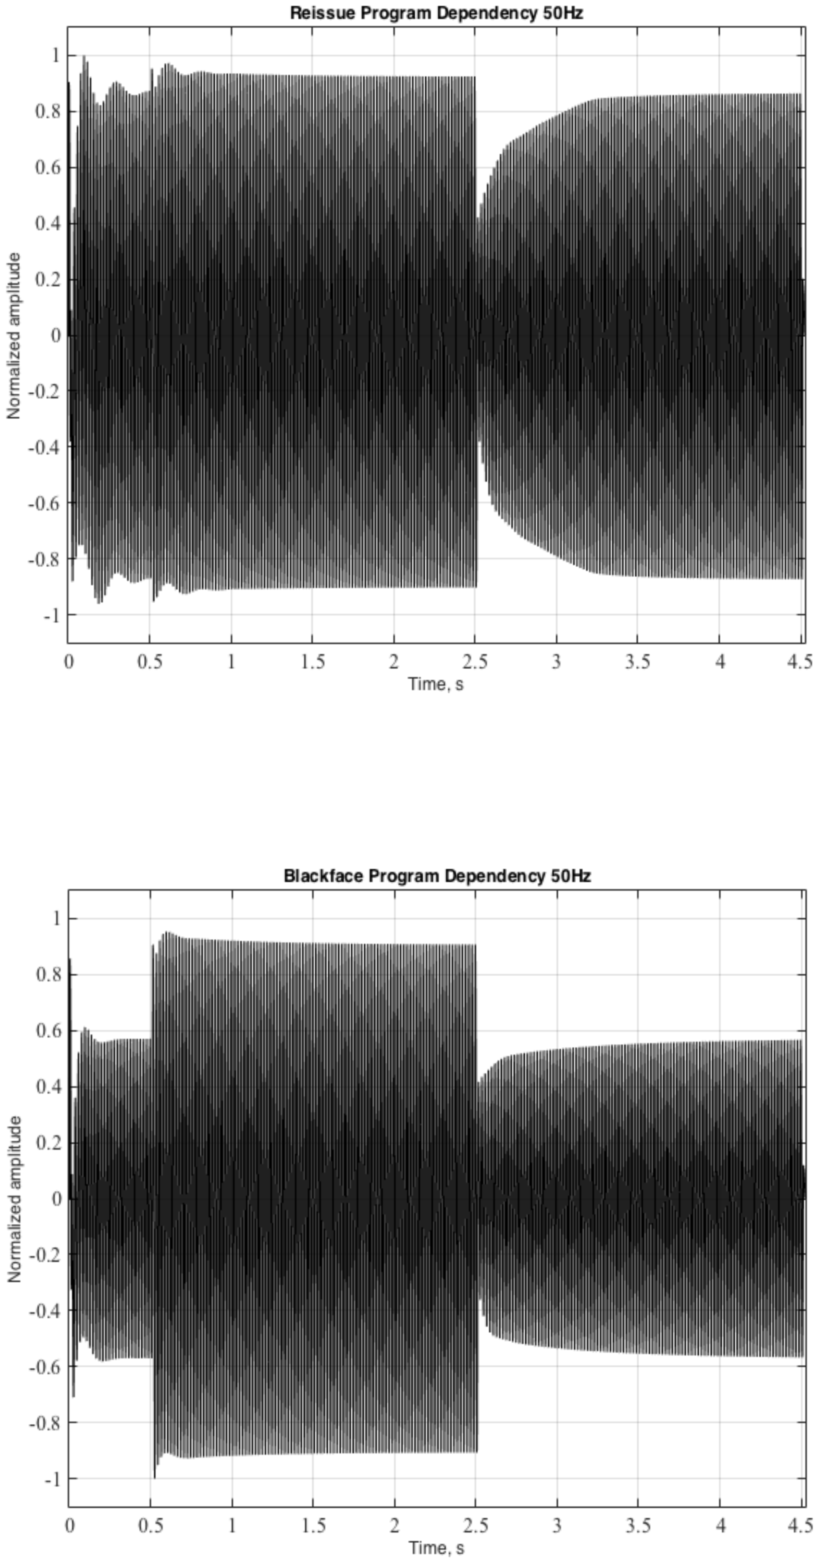

Figure 6-4: Results from program dependency 50Hz tone. Reissue on top and Blackface on bottom

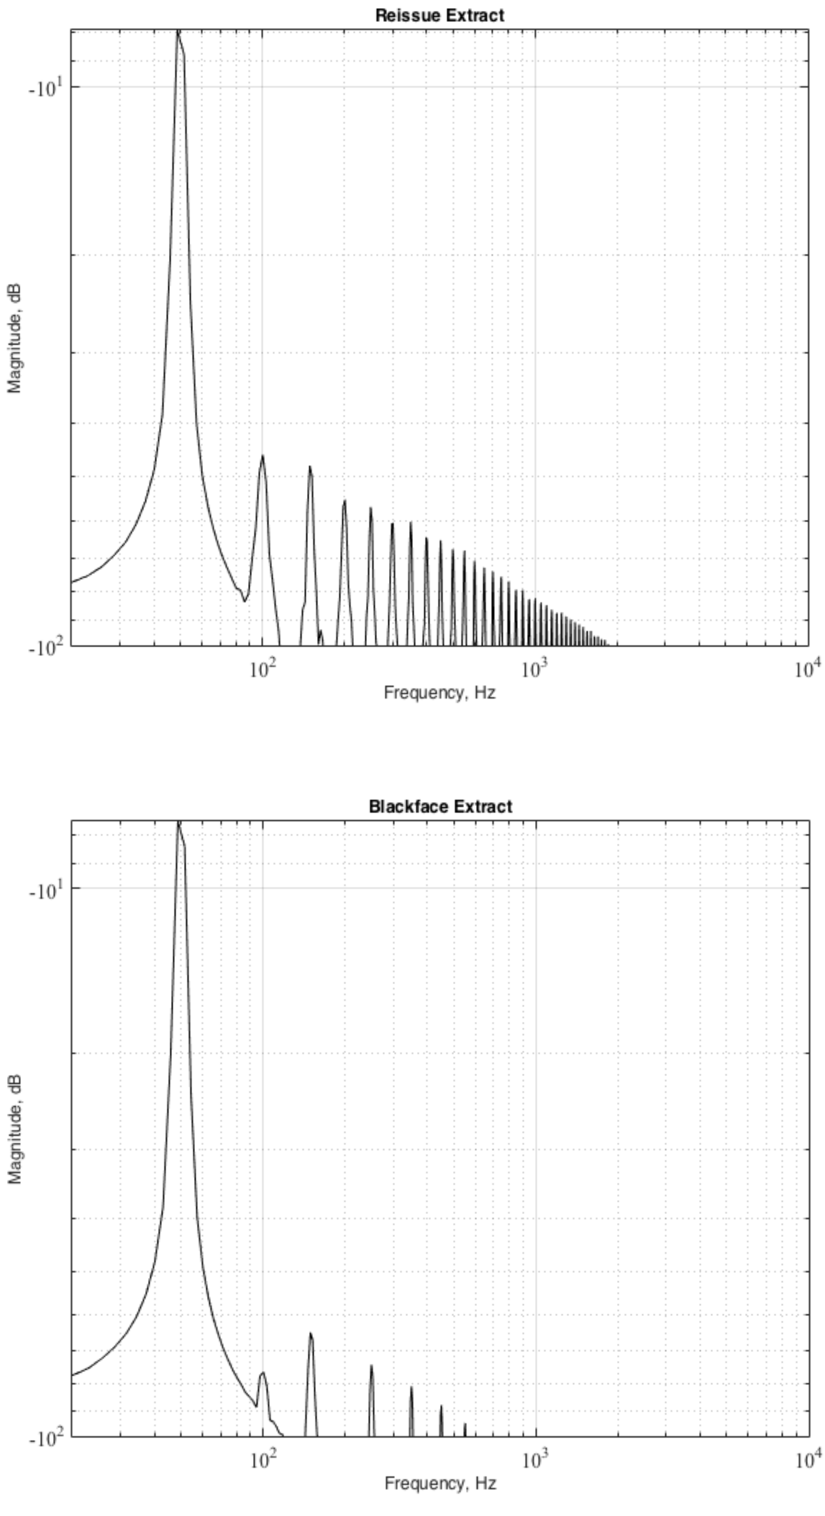

Figure 6-5: Spectrum of the audio during the first 500ms of the 50Hz test tone. Reissue on top and Blackface on bottom

To further explore the result the same test was retrospectively conducted on a second 1176 Universal Audio reissue. The result can be seen in Figure 6-6 where the first 500ms has been controlled more effectively than the other Reissue, although the response after the transient is still not as quick as the Blackface. This test indicates there is some variation in how different models and instances of the 1176 respond under similar working conditions and it is probable these subtle differences will play a role in each compressor's sonic fingerprint.

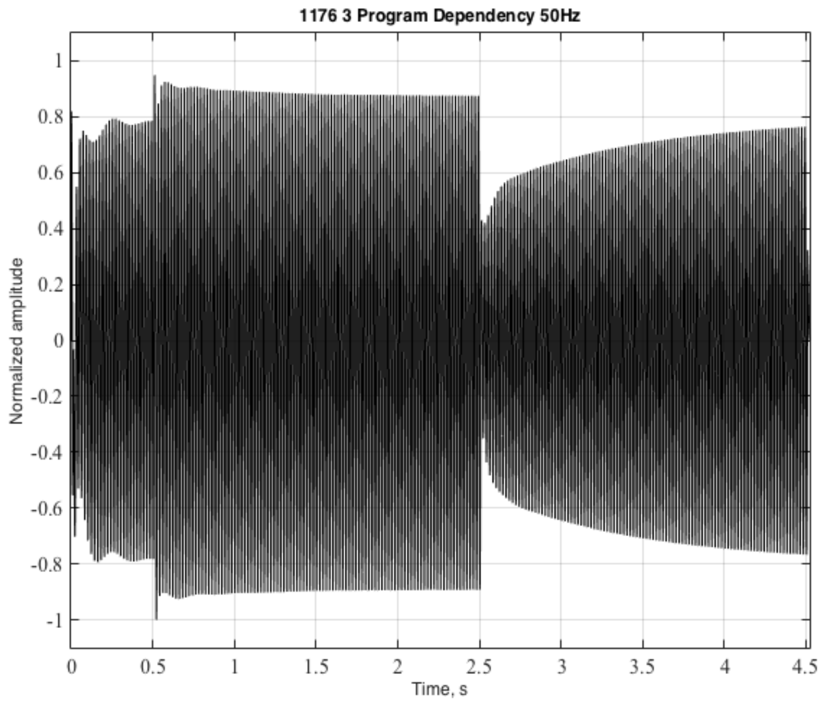

Figure 6-6: Results from program dependency 50Hz tone using a second modern 1176 reissue

6.4 THD at Input: Two 1176s

As with the analysis in Chapter 5, sine waves of varying amplitudes were sent through the 1176s with the gain reduction circuit turned to the off position. The reason for conducting this test again was firstly to compare the difference in non-linearity between the two units and secondly to drive the input with hotter signals than the previous tests.

The test tone was sent to the compressors at +16dBu with the input and output dials set to positions 30 and 18 respectively, this is the typical starting position for gain staging with the 1176 and results in approximately unity gain. From there, the input on both devices was increased to the position with a dot exactly between 30 and 24 on the input dial and then up to 24. The change of the input at these two positions yielded an increase in gain of approximately 5dB per position, thus at position 30 and 24 the input signal was +21dBu and position 24 the signal was +26dBu. These signals are very hot but are the type of levels used by producers to drive input stages for colouration and distortion effects. Table 6-1 shows the THD percentages for the three input levels where input level 1 refers to +16dBu, input level 2 refers to +24dBu and input level 3 refers to +26dBu. The input and output controls of the 1176 are shown in Figure 6-7 to make clear how the levels were set for this test.

Figure 6-7: Image of the 1176 input and output controls

The tones were sent to the compressors at 1kHz and 50Hz to test for differences in non-linearity as a function of frequency. Measurements at a larger number of frequencies were considered but not conducted, as it was found during provisional testing that there was minimal variation in non-linearity over the frequency range. The 50Hz tone was used to encourage transformer-based distortion, but none of this non-linearity is present in the results.

The results from the 1kHz tests can be seen in Table 6-1 and the plots in Figures 6-8 and 6-9 present the FFT of the 1kHz tones. Figures 6-10 and 6-11 illustrates the FFT of the 50Hz tones. The table data shows the Blackface has greater non-linearity than the Reissue for inputs one and two. There is a considerable increase in THD at input level two for the Blackface while the Reissue yields consistent THD levels. The increase in THD at input level three for both compressors is due to hard clipping.

The THD result for input three shows the Reissue generates heavier hard clipping than the Blackface, it is clean up to a point but then quickly transitions into clipping. This result highlights that under certain circumstances the Blackface tested in this study is more flexible for music production techniques that require the use of a driven input stage. The reason for this additional flexibility is due to its input allowing more control over a wider range of non-linearity.

Compressor	Input Level	THD %
Reissue	1	0.15
Blackface	1	0.53
Reissue	2	0.17
Blackface	2	1.85
Reissue	3	7.59
Blackface	3	5.53

Table 6-1: THD for both 1176s at three input levels using a 1kHz tone

The plots depicted in Figure 6-8 show the relative levels of harmonics in the Reissue's 1kHz tone varies between inputs one and two. A reduction in amplitude of the second harmonic is partly responsible for the similar THD percentages shown in Table 6-1 between the Reissue's input levels one and two.

The nature of the hard clipping created at input level three by the Reissue is harsher and brighter than that generated by the Blackface. Note there are some small differences between the 1kHz and 50Hz plots and the most noticeable difference is the level of the second harmonic in the Reissue's plots. It is also shown that even harmonics are higher in amplitude than odd harmonics for all inputs bar input three for the Reissue and this is visible at both frequencies.

The reader is directed to the audio folder for these files where the differences in non-linearity between inputs one and two is clearly audible for the Blackface but virtually inaudible for the Reissue.

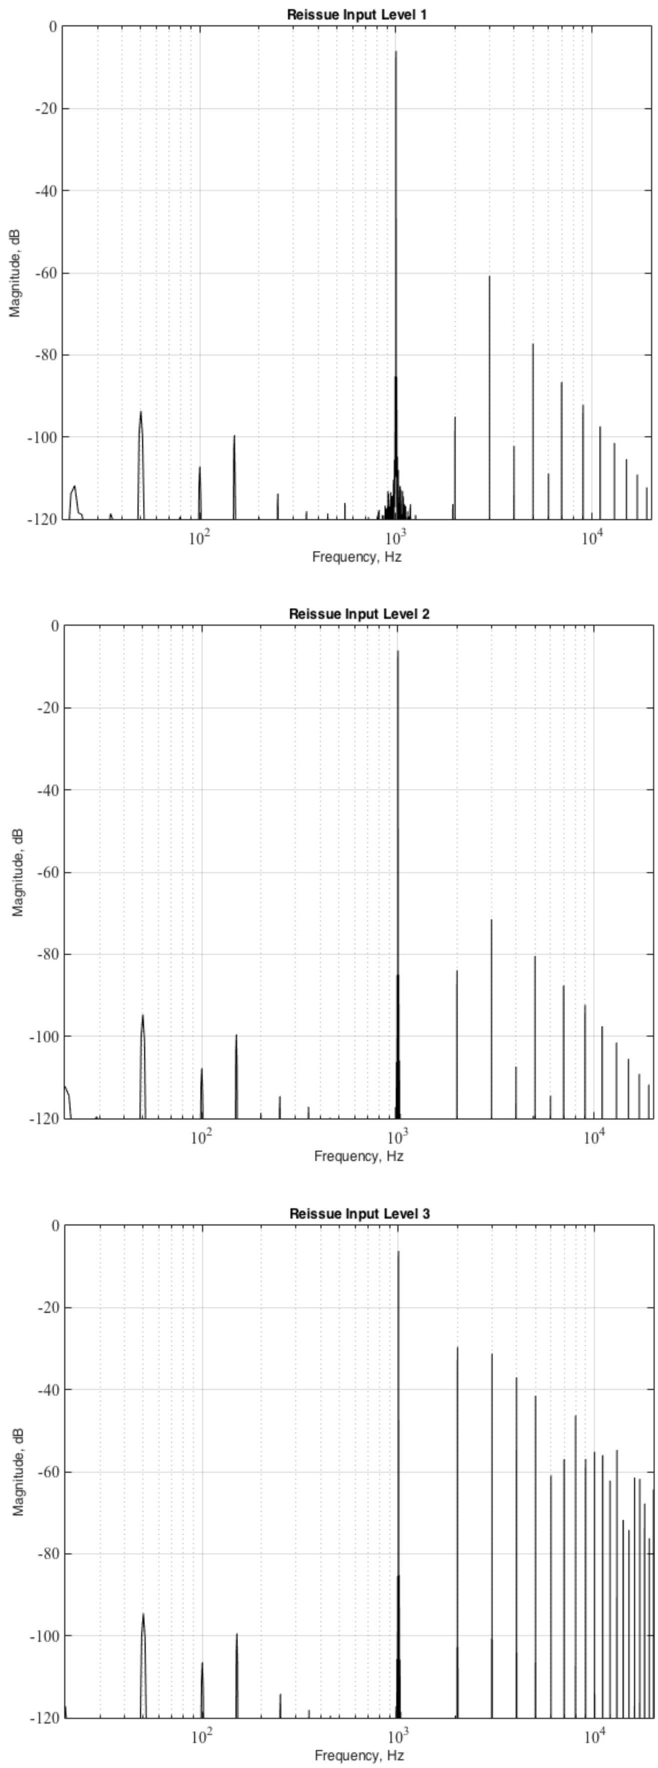

Figure 6-8: THD for Reissue at three input levels and 1kHz tone. Top is low, middle is medium and bottom is hot

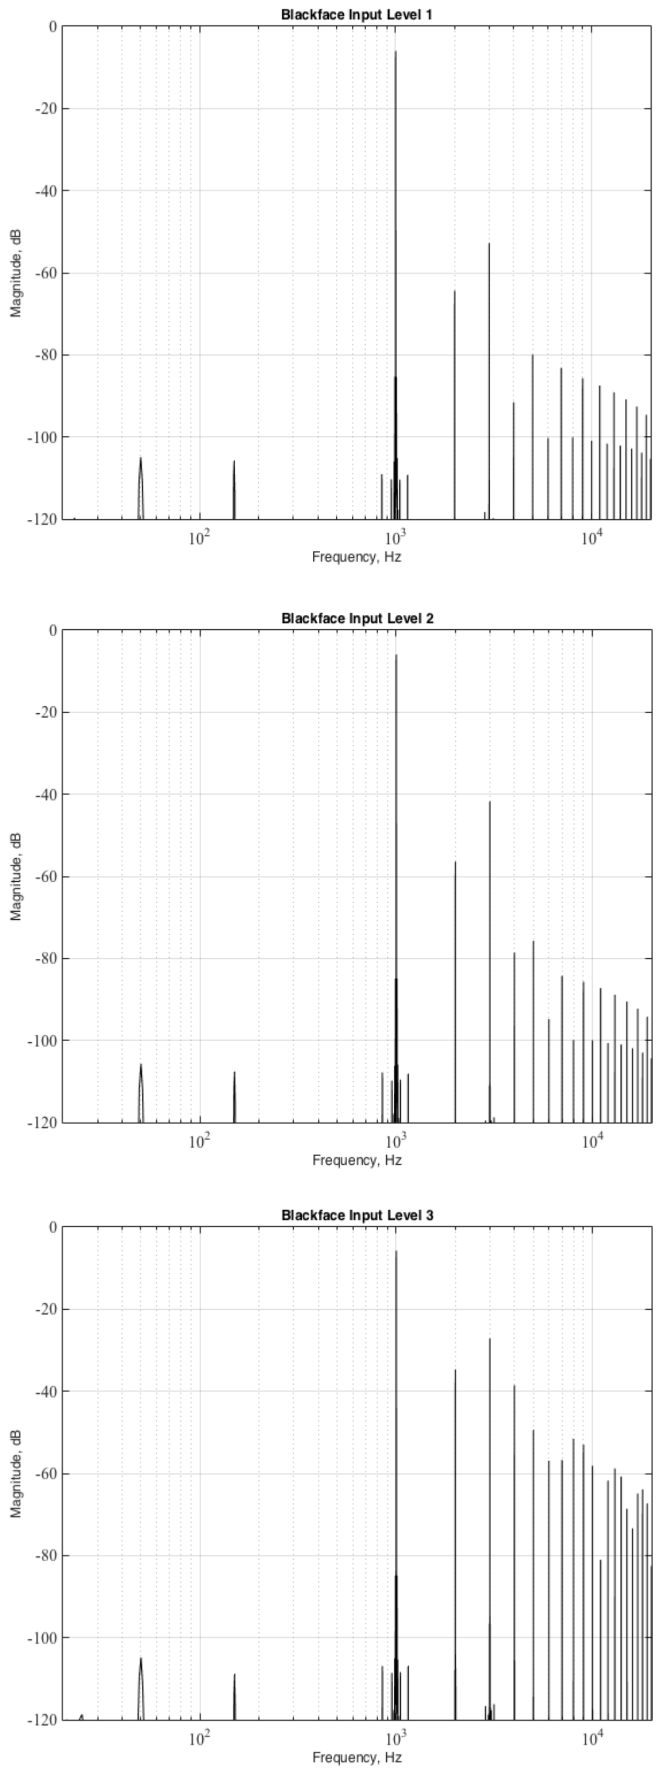

Figure 6-9: THD for Blackface at three input levels and 1kHz tone. Top is low, middle is medium and bottom is hot

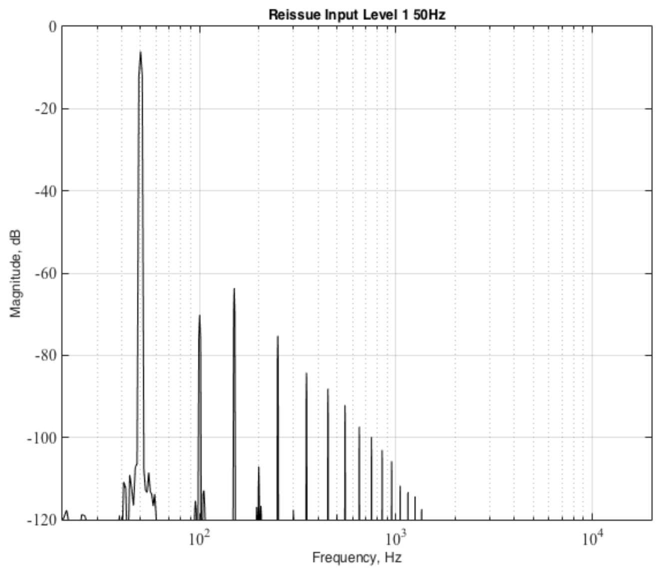

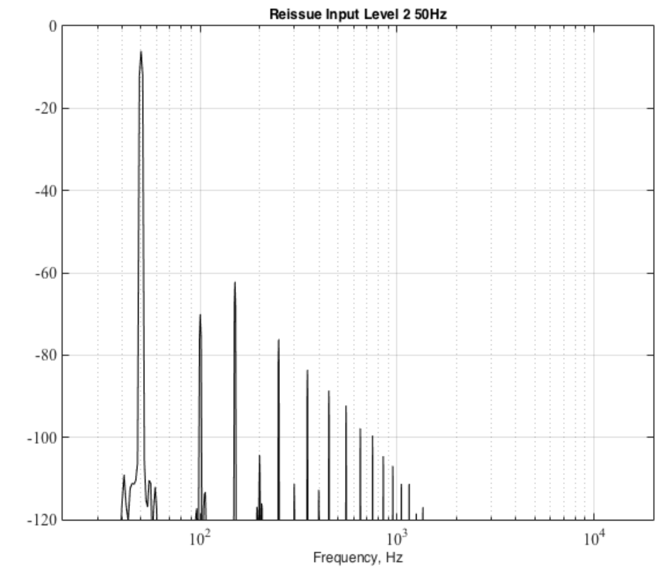

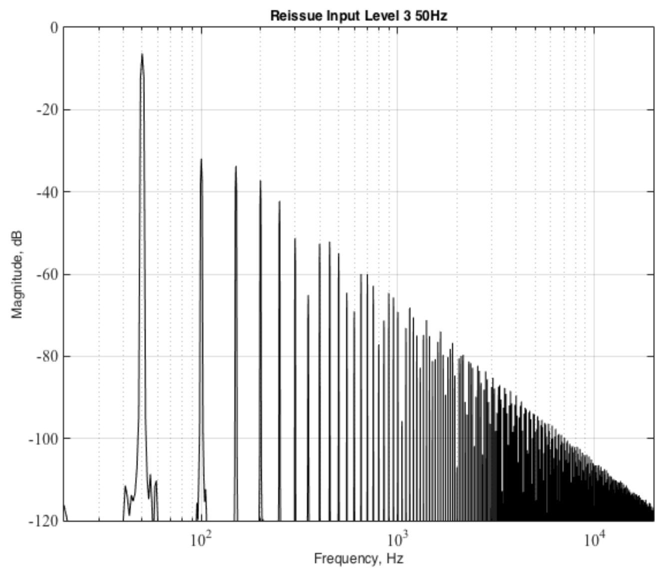

Figure 6-10: THD for Reissue at three input levels and 50Hz tone. Top is low, middle is medium and bottom is hot

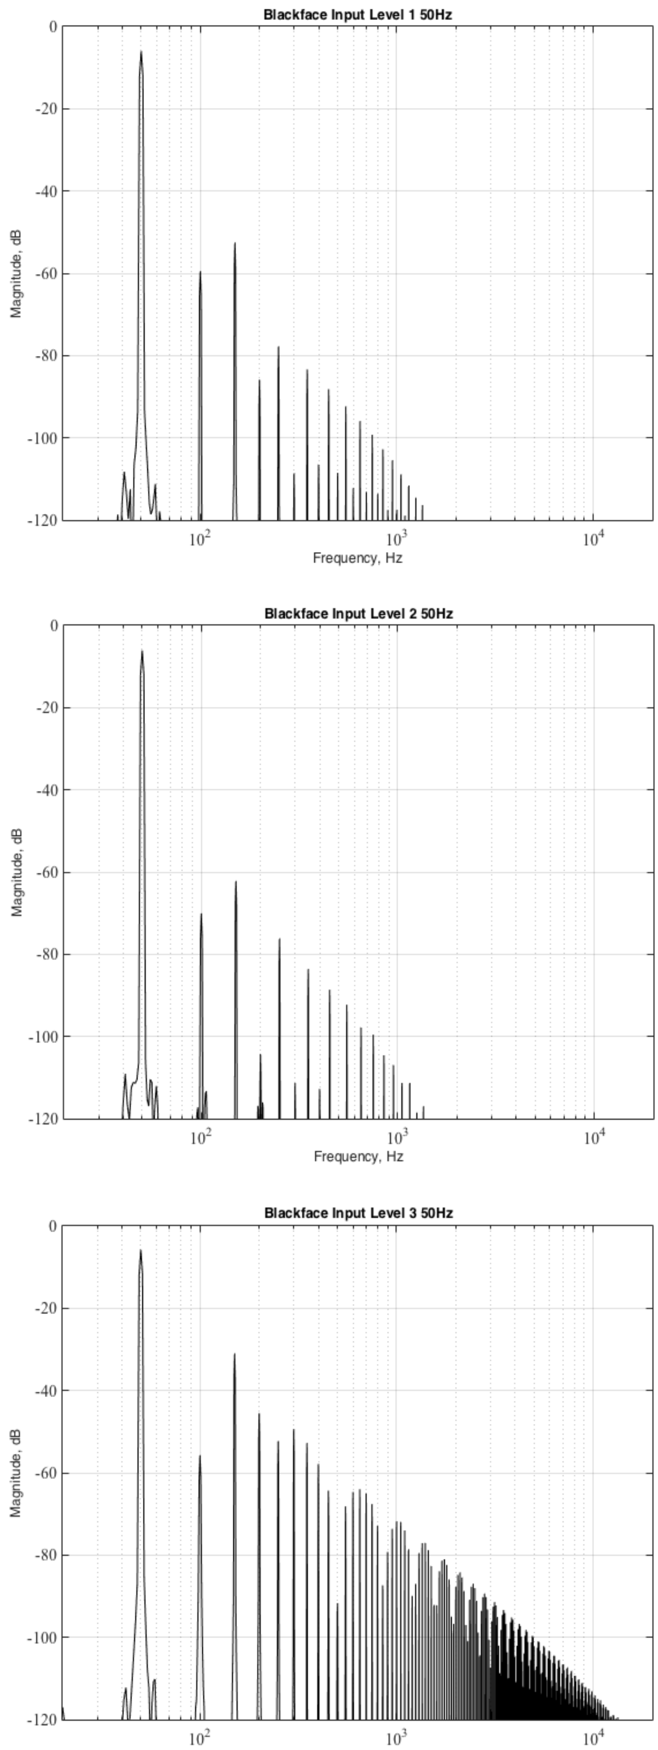

Figure 6-11: THD for Blackface at three input levels and 50Hz tone. Top is low, middle is medium and bottom is hot

6.5 THD During Compression Activity: Two 1176s

The effect of non-linearity during gain reduction and release activity was measured on the two compressors to compare the results and investigate the difference between 1kHz and 50Hz. The 50Hz tone was again selected to reveal transformer based distortion and also distortion from fast time constants. The tones were sent to the compressors at a +4dBU level and then the input control was adjusted to achieve -7dB of gain reduction on each compressor's VU meter. The results are presented in Table 6-2 for the 1kHz tone and Table 6-3 for the 50Hz tone. FFT plots are depicted in Figures 6-12 to 6-15 where the top plots show the compressors during gain reduction and the bottom plots show the compressors in release. Only plots generated from the fastest attack and release settings and a 20:1 ratio have been included here as variations in compression related non-linearity (meaning how it changed between compression activity and release) remained consistent for all other settings.

Comparing Figure 6-12 to 6-13 shows the Reissue produces higher-level harmonics than the Blackface. The first, second and third harmonics are -56dB, -50dB and -63dB down from the test tone for the Reissue, while the Blackface generates the same material at -58dB, -58dB and -72dB down. This distortion equates to THD percentages of 0.37% and 0.16% for the Reissue and the Blackface respectively. Perceptually this difference sounds like a slight increase in brightness, a small change but nonetheless audible under listening.

The plots during release reveal both the compressors are cleaner there than under gain reduction. There are significantly less high order harmonics during release and the distortion now only consists of the first two harmonic components, which are at a lower level than during gain reduction. Comparing the compressors shows the Reissue has more non-linearity than the Blackface. The THD calculation shows the Reissue has 0.19% while the Blackface has 0.05%.

The THD percentages for the 50Hz tone in Table 6-3 and the THD plots in Figures 6-14 and 6-15 depict more non-linearity for both compressors during gain reduction when compared to the equivalent 1kHz plots. The relative level of harmonics varies between the two units, particularly the components between 200-400Hz. The first harmonic in the Reissue is approximately 7dB louder than the equivalent harmonic in the Blackface. Sonically the non-linearity on both devices is very noticeable with the distortion creating a pleasant thickening of the sound, not only a thickening of timbre but an audible addition of distortion. The Reissue sounds more distorted than the Blackface, by a small amount but noticeable nonetheless. The THD figures show the Reissue has 3.81% and the Blackface has 2.49% thus commensurate with what is being heard.

Both compressors are cleaner during release, and the difference between them is more audible than with the 1kHz tone. The first, second and third harmonics are approximately 30, 20 and 50dB louder for the Reissue compared with the Blackface and the drop off in the amplitude of harmonics is more extreme for the Blackface with only artefacts up to 250Hz in the audible range. The high amplitude artefacts produced by the Reissue extend up to around 700Hz, and this creates a more distorted sound quality. The THD values during release are 1.14% for the Reissue and 0.12% for the Blackface.

The results from this test reveal the Reissue is more distorted than the Blackface during compression activity, this is the reverse of the result obtained from THD at input testing.

Compressor	Comp/Release	THD %
Reissue	In Comp	0.37
Blackface	In Comp	0.16
Reissue	In Release	0.19
Blackface	In Release	0.05

Table 6-2: THD during compression activity 1kHz tone for both 1176s

Compressor	Comp/Release	THD %
Reissue	In Comp	3.81
Blackface	In Comp	2.49
Reissue	In Release	1.14
Blackface	In Release	0.12

Table 6-3:THD during compression activity 50Hz tone for both 1176s

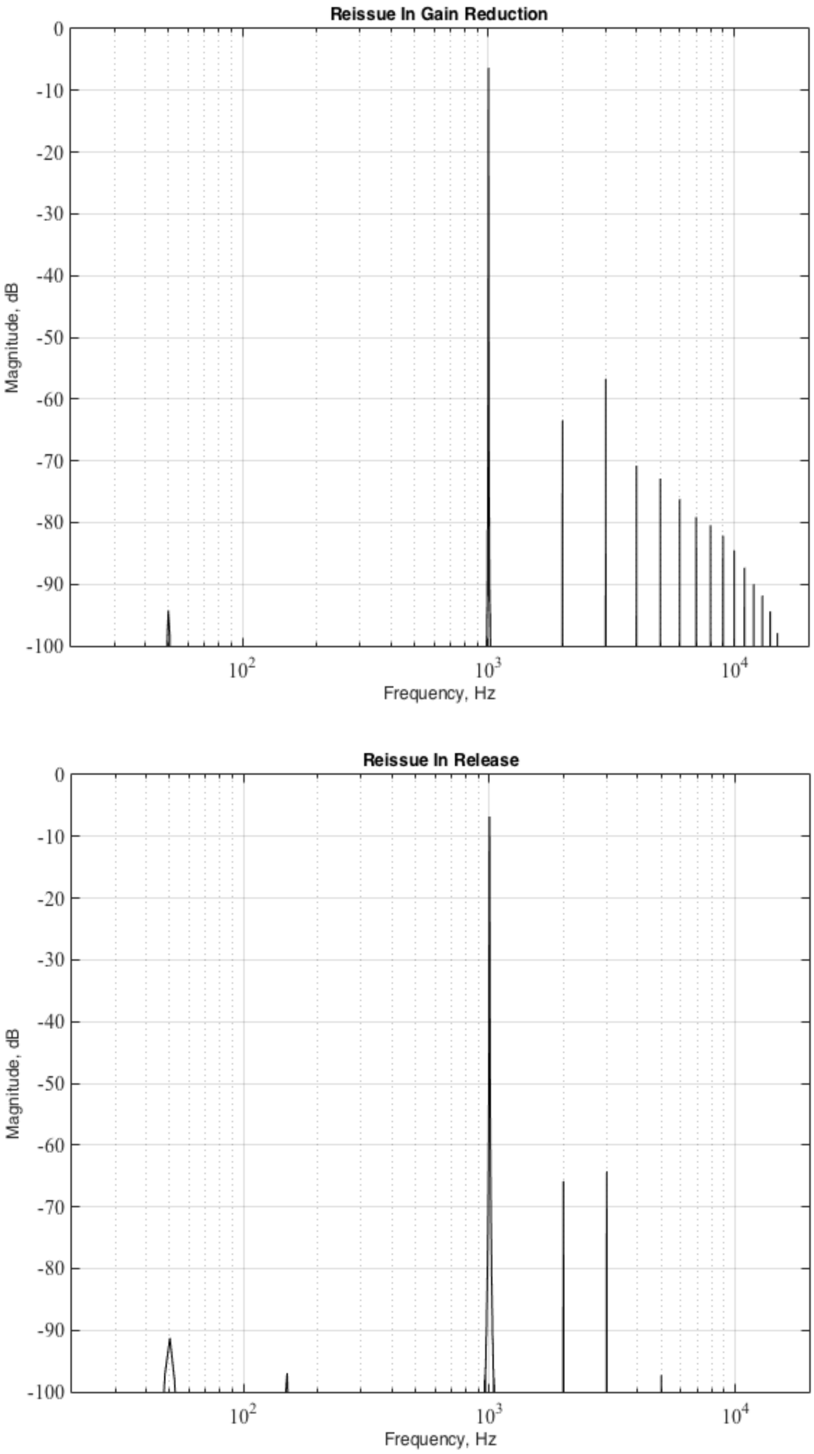

Figure 6-12: THD during compression top and release bottom for Reissue with 1kHz tone

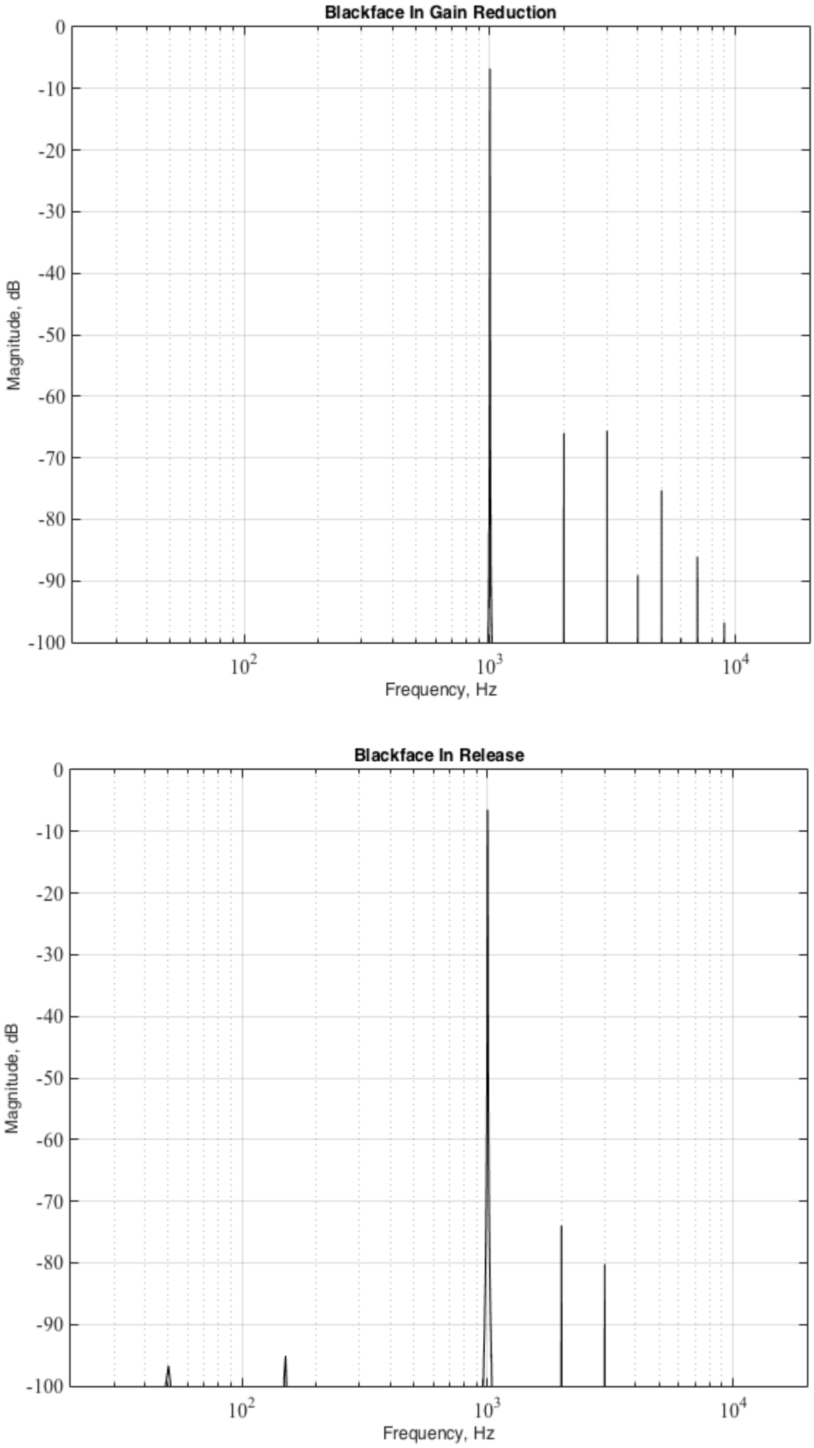

Figure 6-13: THD during compression top and release bottom for Blackface with 1kHz tone

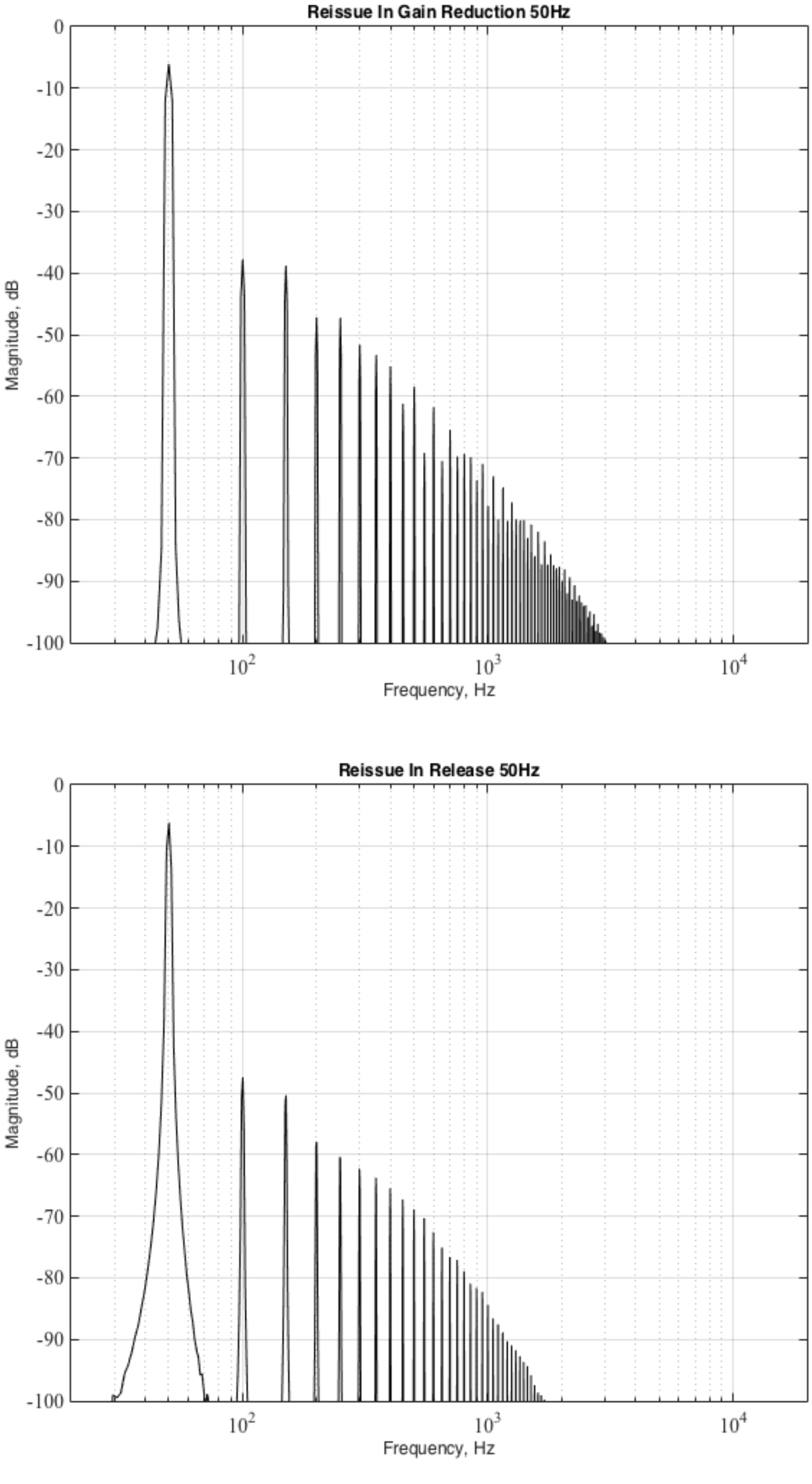

Figure 6-14: THD during compression top and release bottom for Reissue with 50Hz tone

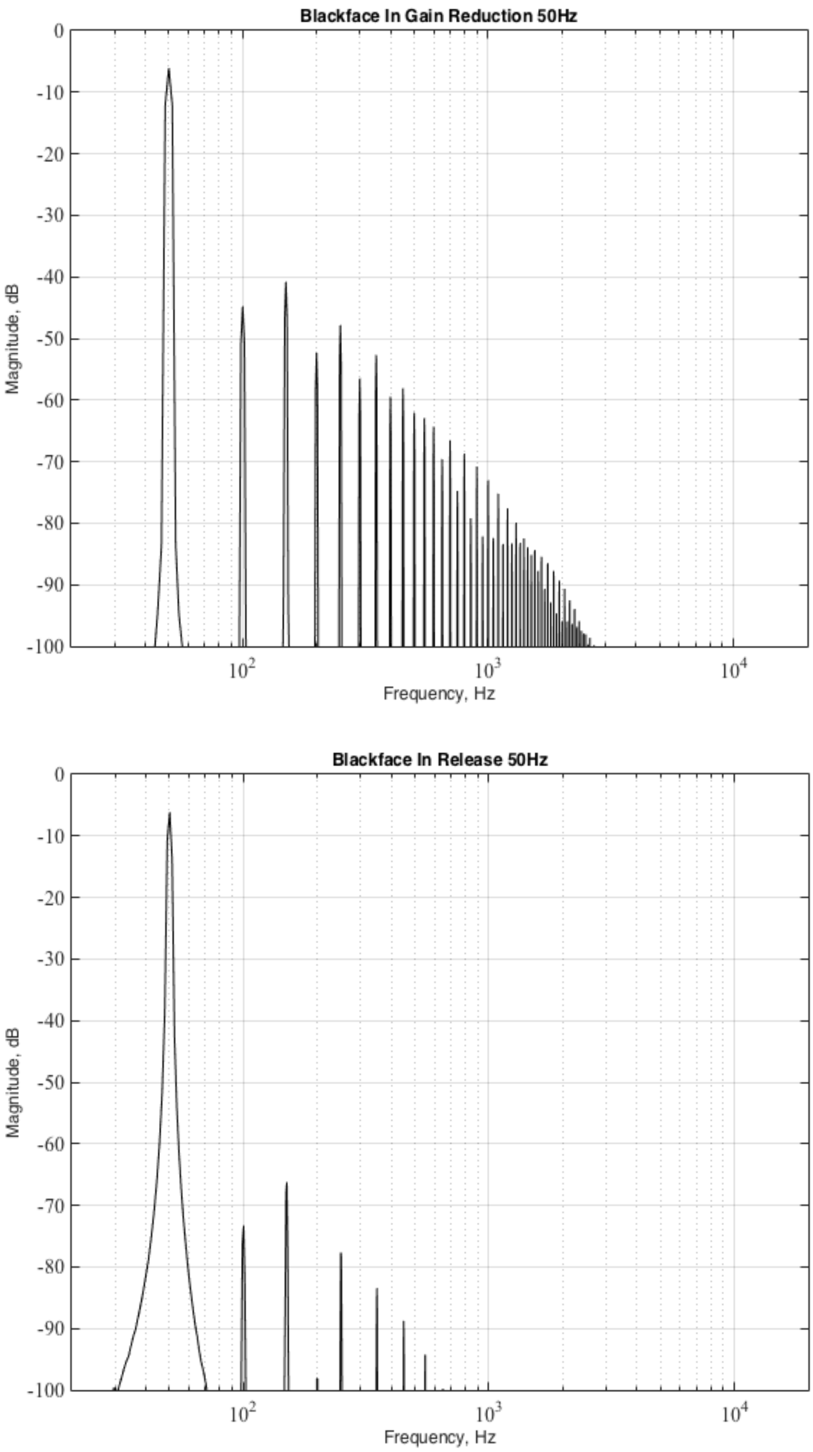

Figure 6-15: THD during compression top and release bottom for Blackface with 50Hz tone

6.6 Testing Intermodulation Distortion (IMD): Two 1176s

One final sine based measurement was made using the SMPTE intermodulation tone consisting of two sine waves at 60Hz and 7kHz with a 4:1 amplitude ratio. This analysis is used to reveal non-linearity generated from intermodulation distortion and gives a more revealing picture of the type of non-linearity that will be present when using complex program material. The measurement was made at a number of typical time constant settings to better understand IMD with more realistic compression activity. The time constant parameters were gleaned from the research in Chapter 4. The results are presented in Table 6-4 where A=attack and R=release positions. The amount of gain reduction was set to read -7dB on the VU meter to achieve a moderately high amount of compression for all the measurements made.

Setting	IMD % The Reissue	IMD % The Blackface
A1R7 4:1	2.89	2.52
A7R7 All	31.73	18.53
A3R5 4:1	0.71	0.52
A3R7 4:1	2.15	2.2
A7R7 4:1	2.33	2.1
A7R7 20:1	6.97	2.88

Table 6-4: IMD for both compressors over a range of compression settings

Table 6-4 indicates there are some subtle differences between the two compressors and the Reissue is more distorted than the Blackface, although the Blackface is marginally more distorted in the fourth setting. It is shown by comparing Table 6-2 and 6-4 that IMD is higher than THD for the 1kHz tone. Whitlock (2008) states that IMD can be predicted to be 3 or 4 times greater than THD but under these tests, it is 18.8 times greater. This figure was calculated by dividing 6.97% (taken from the Reissue's A7R7 20:1 result) by 0.37%, which is from the 1kHz THD result in Table 6-2 that used the same time constant and ratio setting. However, compared with the THD figure at 50Hz in Table 22 for the same configuration, IMD is only 2 times greater. This variation in THD to IMD ratio may be because of complex interactions between different components in the gain reduction circuit of the 1176.

Two aspects of behaviour were noted during generation of the IMD figures. Firstly, release plays an important role in non-linearity, more so than the attack. Secondly, non-linearity increases with higher ratios and most significantly in the all-buttons mode. To test the observations a series of measurements were made on another Universal Audio Reissue 1176 the author had access to during research. The author of the report would have liked to have made the measurements on the Reissue and the Blackface tested in the rest of this chapter, but that option was not viable.

6.6.1 The Effect of Ratio on Non-Linearity

To test how ratio (and more broadly gain reduction) affected non-linearity a number of measurements were made at all ratio settings. Six different amounts of gain reduction were measured at positions on the VU gain reduction meter relating to -1, -3, -5, -7, -10. The sixth position is referred to as pinned which means the input was adjusted until the gain reduction meter was at its most extreme position. The test tone was sent to the 1176 at 0dBu. In all-buttons mode, the VU works out of calibration so although the measurements were taken with the meter at the same positions as other ratios it does not necessarily mean the amount of gain reduction is comparable. For completeness, the tests were made using the SMPTE IMD test tone and also 1kHz and 50Hz tones. The results for all tones are presented in a number of figures that plot IMD and THD as a function of gain reduction. All ratios are included in the plots but each relevant figure features two plots with the top plot excluding all-buttons, this is to make differences between the other ratios easier to visualize.

It can be seen clearly in Figure 6-16 that IMD grows with higher ratios and larger amounts of gain reduction. There are some slight differences between ratios with regards to how this non-linearity develops but the general trend is an upward slope with a small plateau between -10 and pinned positions. The all-buttons mode has a rise to the -5 mark and then dips off. When using all-buttons IMD is extremely high, demonstrating why this mode is popular with engineers who want to colour audio material with distortion. This distortion technique was mentioned by the producers in Chapter 4.

The same test was conducted in two separate tests using two sine test tones at 1kHz and 50Hz to compare and contrast the results. Again, there are two plots from the same test with the top plot excluding the all-buttons mode. The results from the 1Khz tone are depicted in Figure 6-17 where a similar trend is visible and shows the THD rising as gain reduction is increased and ratio changed. The trajectory this time is straighter, and there is a sharp incline between -10 and pinned instead of the plateau seen in the IMD results

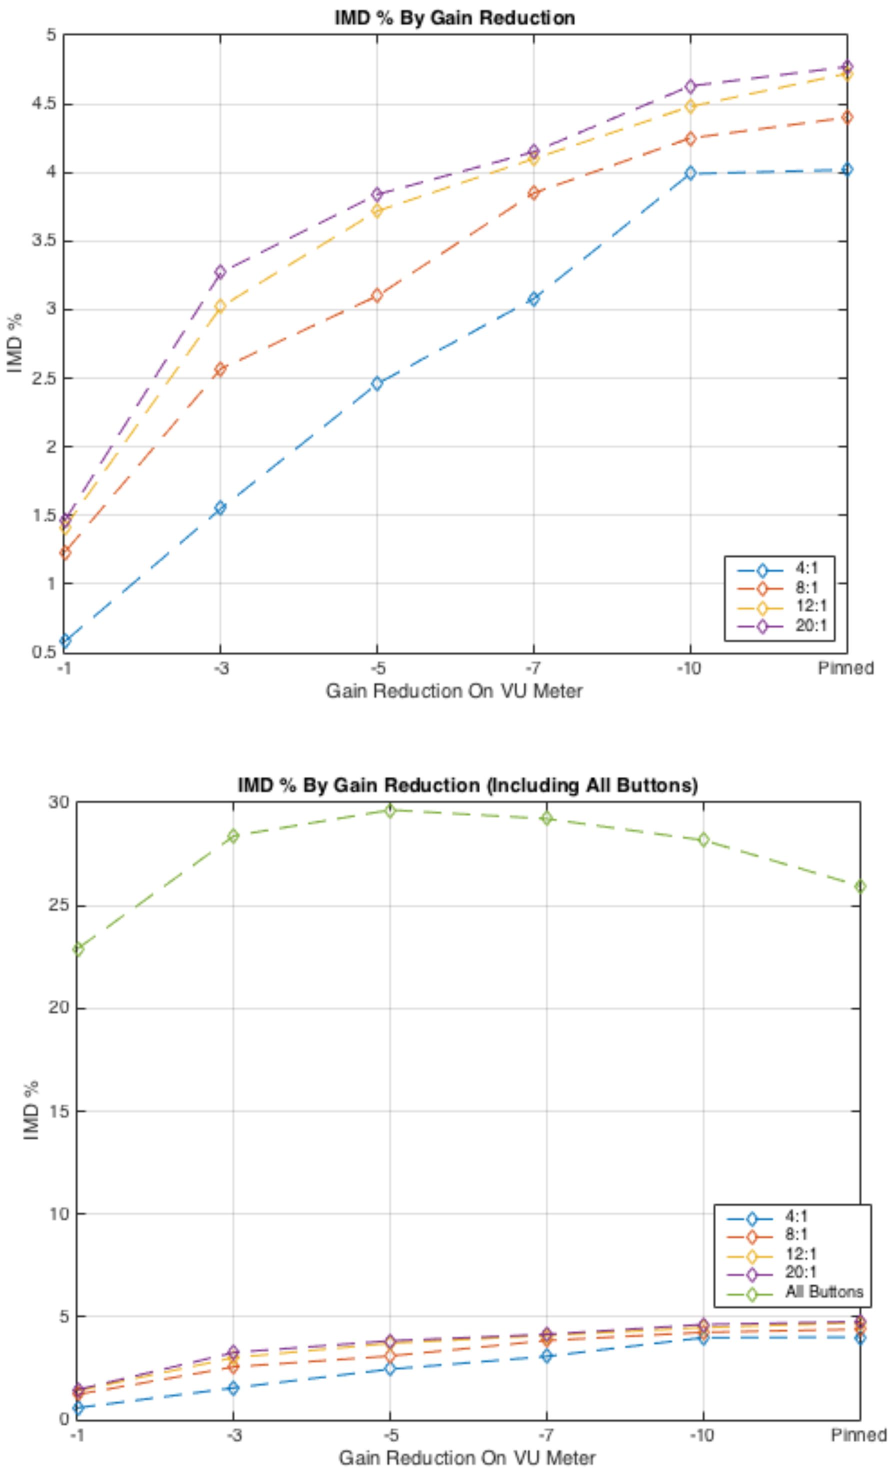

Figure 6-16: IMD as a function of gain reduction. The top plot excludes all buttons mode

As with the IMD tests the THD results indicate the all-buttons mode is significantly more non-linear than the other ratios and has the same drop off at the -5 measurement.

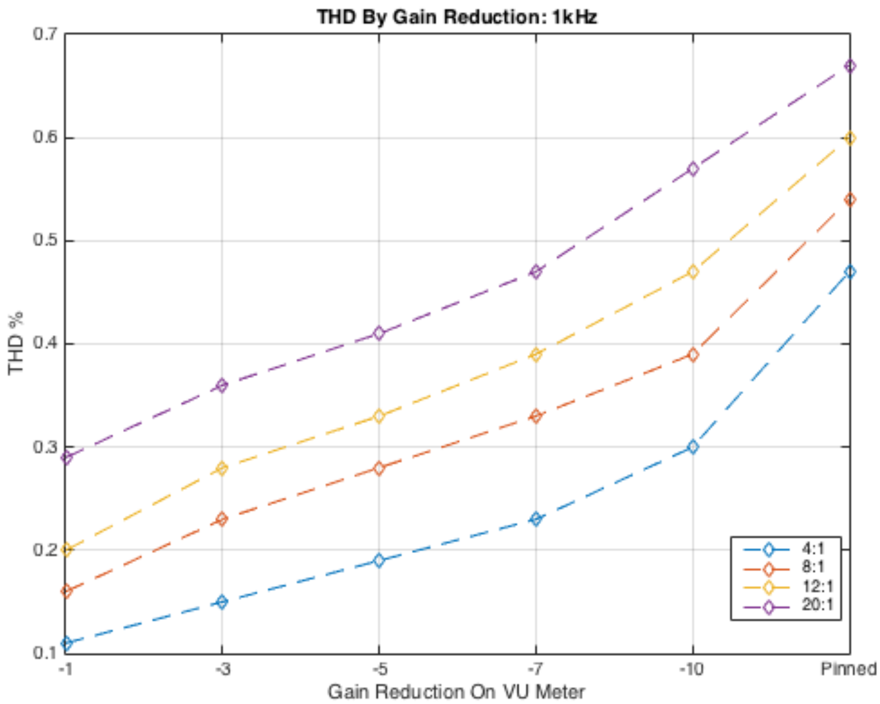

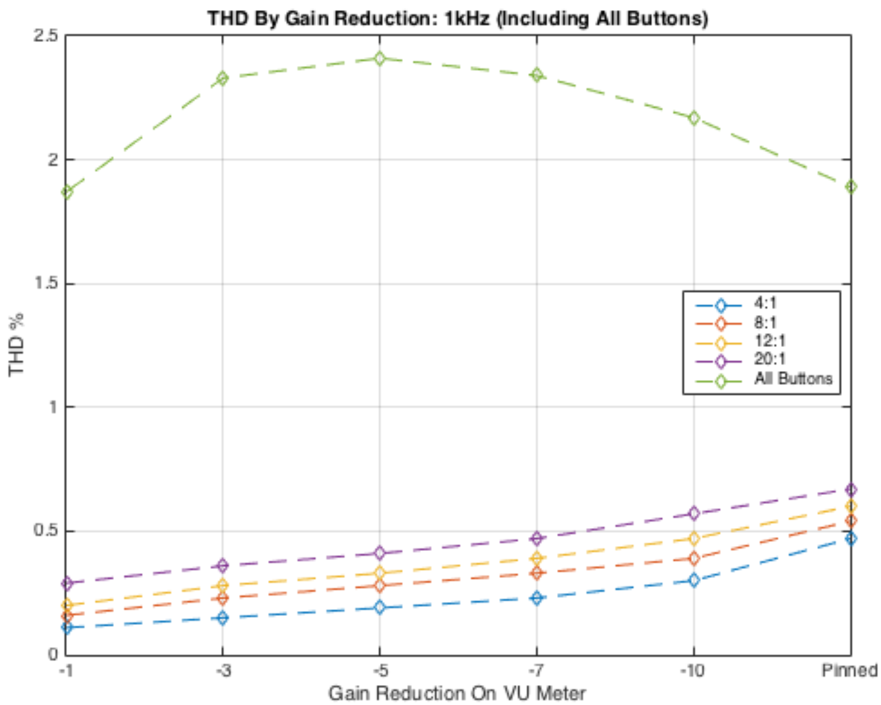

Figure 6-17: THD as a function of gain reduction for a 1kHz tone. The top plot excludes all buttons mode

Figure 6-18 illustrates the results from the 50Hz tone, and a familiar upward trend can be observed. However, the curve this time is a little different with some leveling off at higher ratios between -10 and pinned measurements. The dip between -3 to -5 for the 12:1 ratio is because of a drop in the second harmonic

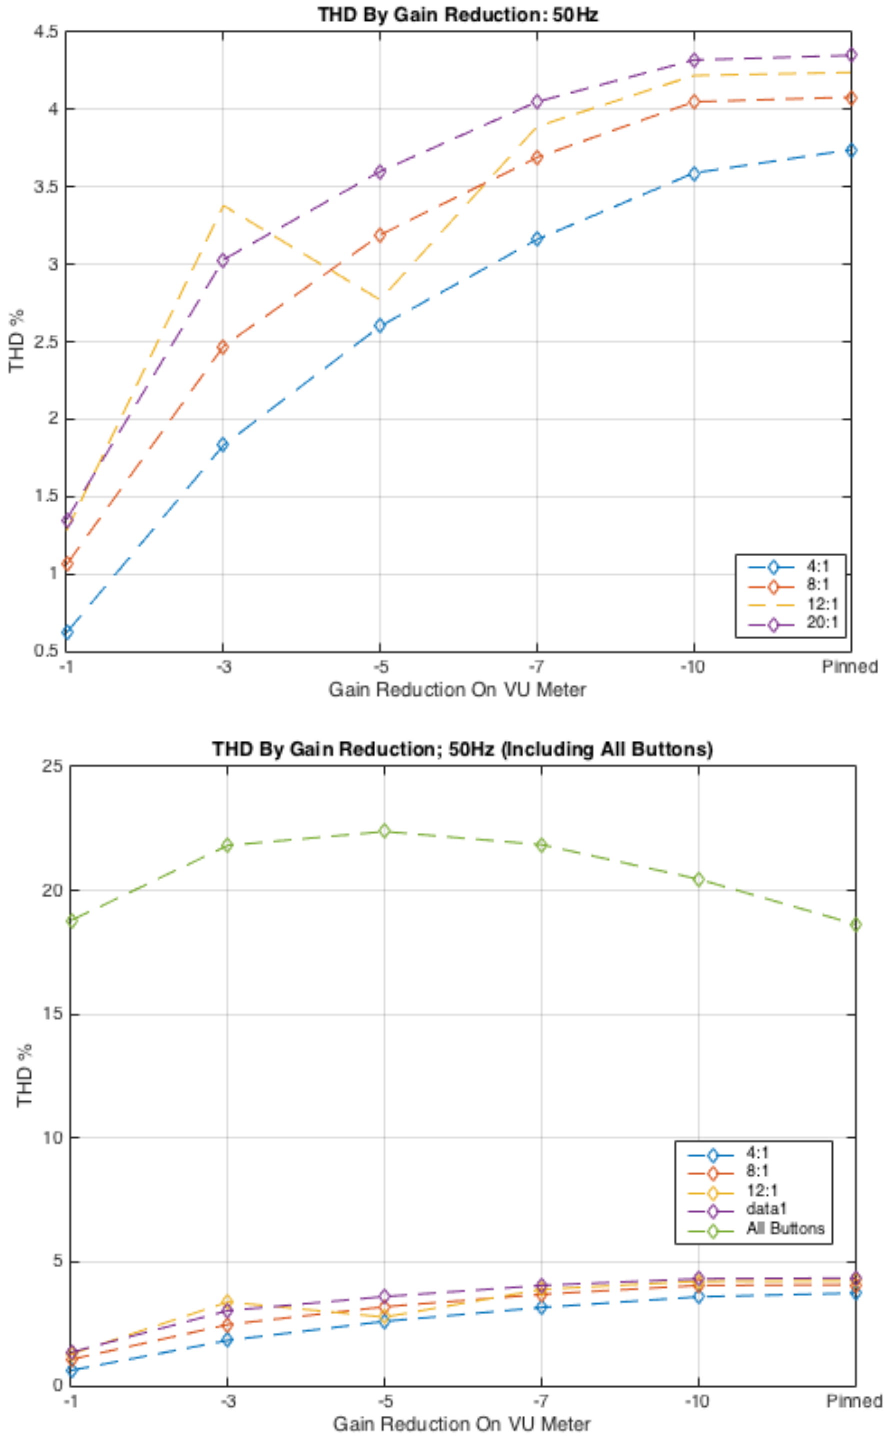

Figure 6-18: THD as a function of gain reduction for a 50Hz tone. The top plot excludes all buttons mode

Comparing Figure 6-17 to Figure 6-18 shows the non-linearity for the 50Hz tone is much larger than 1kHz, and this is due to waveshaping from the fast time constant setting. Again, the all-buttons mode has significantly more non-linearity than the other ratio settings and the plot shows the familiar curve that drops off at the -5 measurement.

The THD and IMD figures for this 1176 are different than those of the Reissue and the Blackface measured in the studio. This result shows there is variation in the non-linearity between all the units tested and presumably this is a common trend with the 1176 compressor. It was shown in Chapter 4 that producers had a preference for certain models of the 1176 and this variation in non-linearity may be one of the reasons why they subjectively prefer one model over the other.

6.6.2 The Effects of Time Constants on Non-Linearity

In the previous sub-chapter, it was noted that higher amounts of ratio and gain reduction had the effect of increasing non-linearity. While making measurements for use in the previous tests, the release control had one of the most significant effects on the amount of non-linearity. More specifically it was noted that the non-linear effects of the time constants are at their maximum when the attack and release are at the fastest positions. Non-linearity is steady until the release control is set to the sixth position where the harmonics drop sharply in level. The effect of the attack on non-linearity is not nearly as significant as the release. Harmonics roll off as the attack is set to slower positions but these artefacts are higher order harmonics, and their effect on audio quality is not significant due to the masking effects of lower order higher amplitude harmonics. Odd order harmonics drop in level by a much greater amount when slowing the release control compared with slowing attack. For example, the third and fifth harmonics with the fastest attack and slowest release are almost 20dB and 17dB less respectively when compared to the slowest attack and fastest release configuration.

A series of measurements were made on the Universal Audio Reissue 1176 at different attack and release settings to investigate these observations thoroughly. The results can be seen in Figures 6-19 that plot IMD as a function of attack and release. The top plot in the figure shows IMD in all-buttons, and the bottom plot shows IMD using a 4:1 ratio.

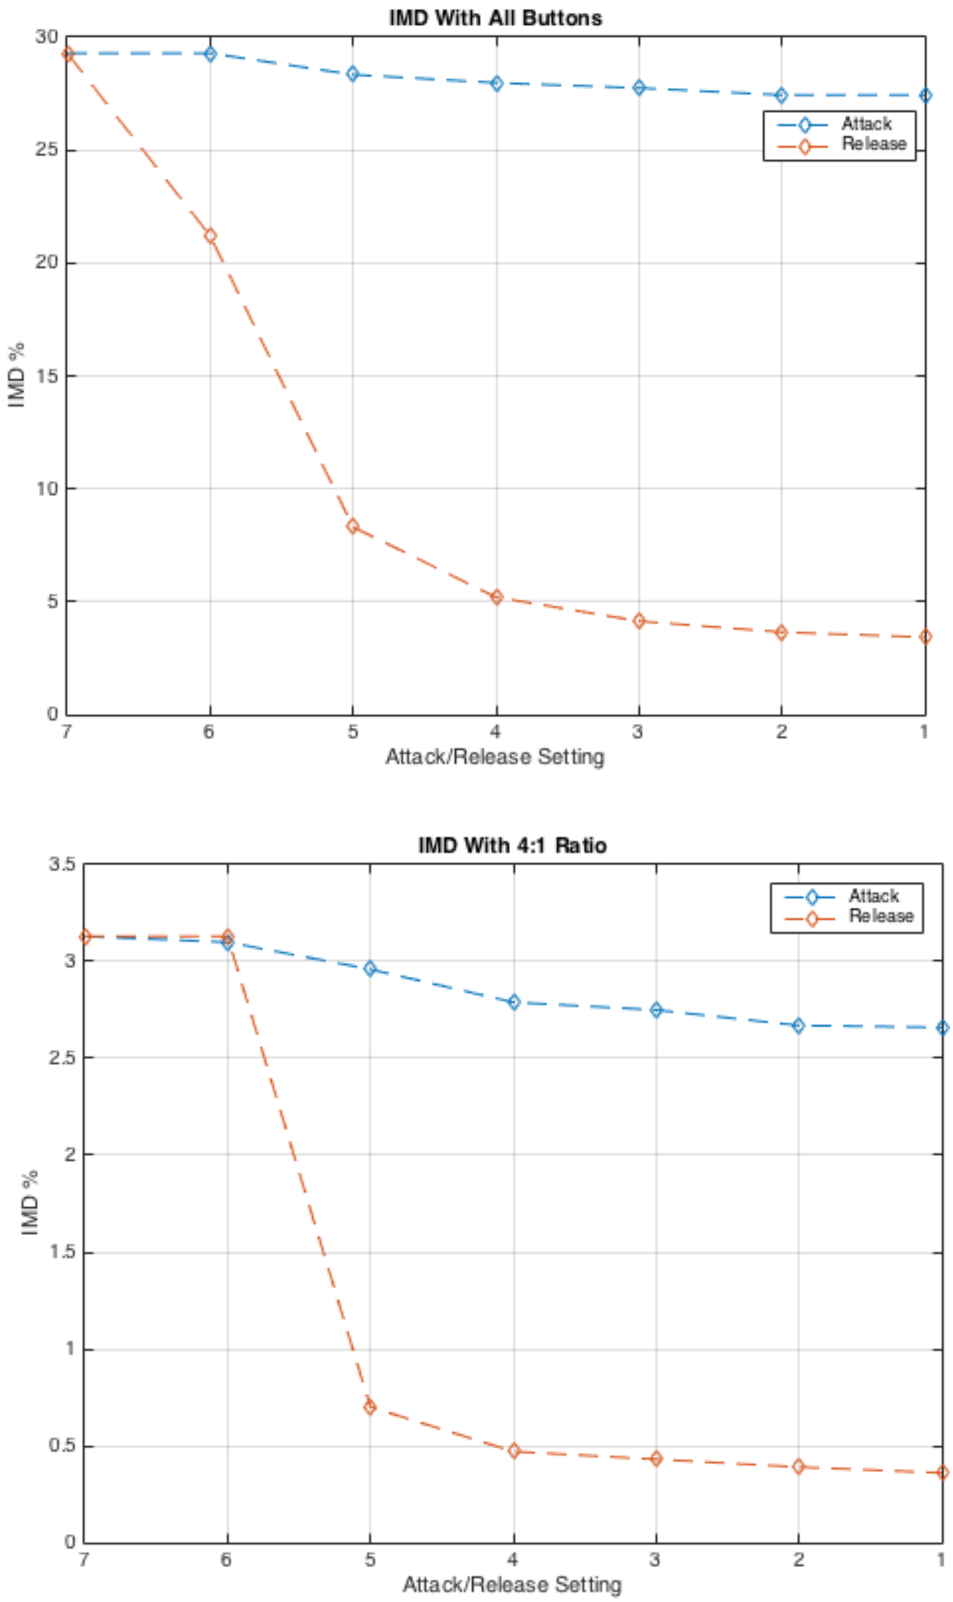

Figure 6-19: IMD as a function of attack and release. Top all buttons and bottom 4:1 ratio

It can be seen in Figure 6-19 that reducing the release speed plays a more significant role in the reduction of non-linearity than decreasing attack. The effect of reducing the attack over the full range of its control has little effect in reducing non-linearity. Given the popularity of the fastest release setting in the study in Chapter 4, engineers appear to be deliberately using this non-linearity in their productions. They are perhaps not consciously aware of what is happening to the signal and are instead preferring the colouration created with the fast release configuration at a perceptual level. This behaviour may be responsible for the sound quality that producers were describing with descriptors such as "bite" and "crunchy" in Chapter 4.

THD was measured as a function of attack and release at 1kHz, and 50Hz and the results are illustrated in Figures 6-20 for the 1kHz tone and Figure 6-21 for 50Hz. As before, the 50Hz tone was selected to observe time constant induced waveshaping at low frequencies. The results correlate with the IMD measurements in that the most significant drop off in distortion occurs at release position 6. There are slight variations in how the distortion diminishes before this position but this sharp decline at position 6 remains a noticeable trend in all the measurements.

6.7 Conclusions on Tone Based Testing

In conclusion, there is variation in the time constant response between the compressors. The second reissue model tested is a new unit that has not been used in a studio, only for research of this thesis. The difference in response between the reissues (that one might expect to behave the same) is possibly due to one being used in a busy tracking and mixing studio. This variation demonstrates how the sonic fingerprint of a piece of equipment can change with age, maintenance, and use. Surveying a much larger range of 1176s to test this variation is recommended for further study.

The THD measurements show that when 1176s are driven hard (meaning with signals above +20dBu) they produce significant non-linearity, which abruptly turns into hard clipping with signals at +25dBu and above. The input of the 1176 is clean and only becomes non-linear when deliberately driven with hot levels. The response to signals at the two different frequencies is consistent, but there are some subtle differences in the non-linearity. There are differences in how the two 1176s respond to the input levels. The Blackface imparts more audible non-linearity with lower input levels while the transition into hard clipping with the Reissue is not as gradual. These differences will no doubt play a role in the subtleties of each unit's sonic signature and will affect the manner in which music producers utilise the devices in their productions. One producer in Chapter 4 described Blackface editions of the 1176 as having a "grindy" sound quality, and this non-linear behaviour may be what he is describing.

When measured for non-linearity during compression, the compressors are more distorted when in gain reduction compared with release. Fast time constants introduce a significant amount of non-linearity to the 50Hz tone, and this is due to the effect of fast time constants discussed in Chapter 3. Non-linearity is also present in the 1kHz tone, but it is not nearly as abundant as in the 50Hz tone. Additionally, this test revealed differences between the two units, particularly when measurements were made during the different stages of compression. The Reissue is the most coloured compressor in this test and this will have an effect on its sonic signature, particularly on low frequency content when using fast time constant settings.

Using higher ratios has the effect of adding more distortion. The effect of the all-buttons mode is considerable, and the amount of non-linearity generated when in this mode is large. This high level of distortion may be one of the reasons why producers use this mode to distort audio material, particularly when seeking to impart a coloured and aggressive sonic signature.

The release control also plays a significant role in creating non-linearity. Setting the release at position seven, which was found to be common in Chapter 4's study, adds the most non-linearity of all the time constant settings measured in the tests. In comparison, the attack has a smaller effect on distortion, the amount of non-linearity differed only slightly over the complete range of the attack control.

Finally, there is some variation in non-linearity between all the 1176s measured for IMD and THD in this thesis. This difference is shown in terms of the amount of distortion they generate and also the manner in which this non-linearity varies between input levels and compression activity. These variations affect the sonic signature of the compressors and explains why users find specific units to have unique characteristics that lend themselves to specific production scenarios.

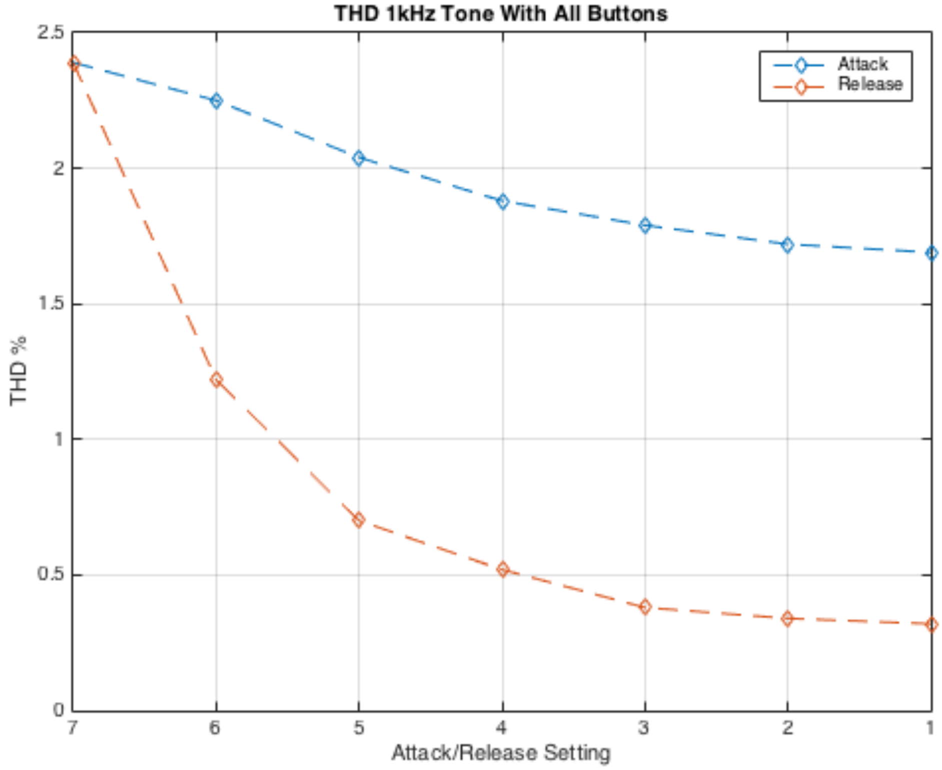

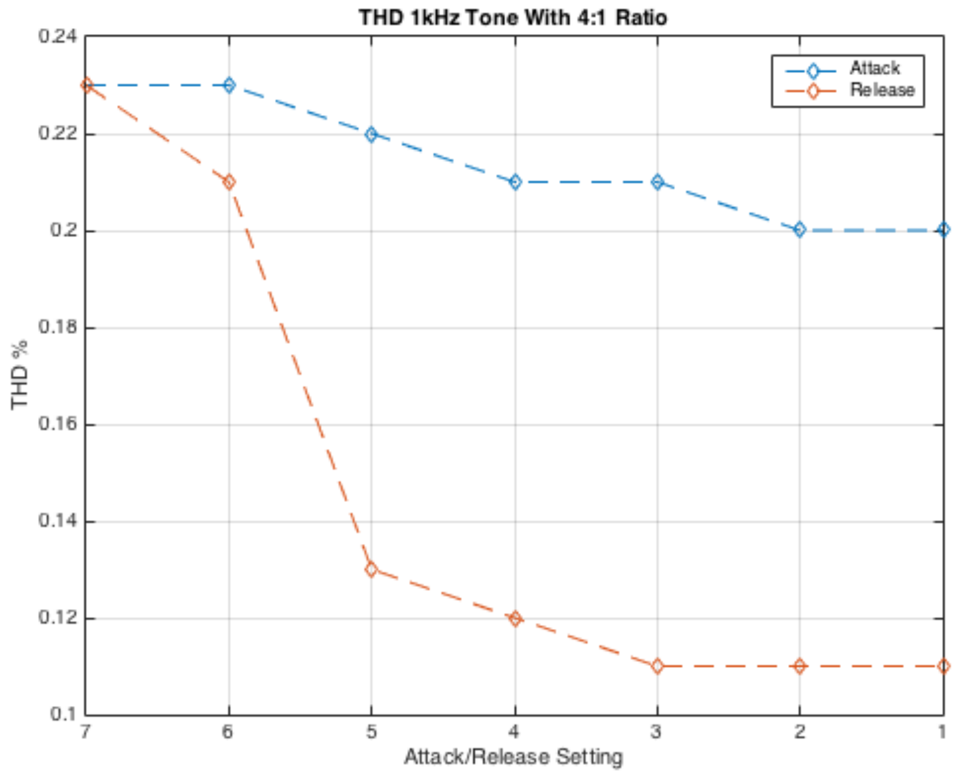

Figure 6-20: THD as a function of attack and release for 1kHz tone. Top all buttons and bottom 4:1 ratio

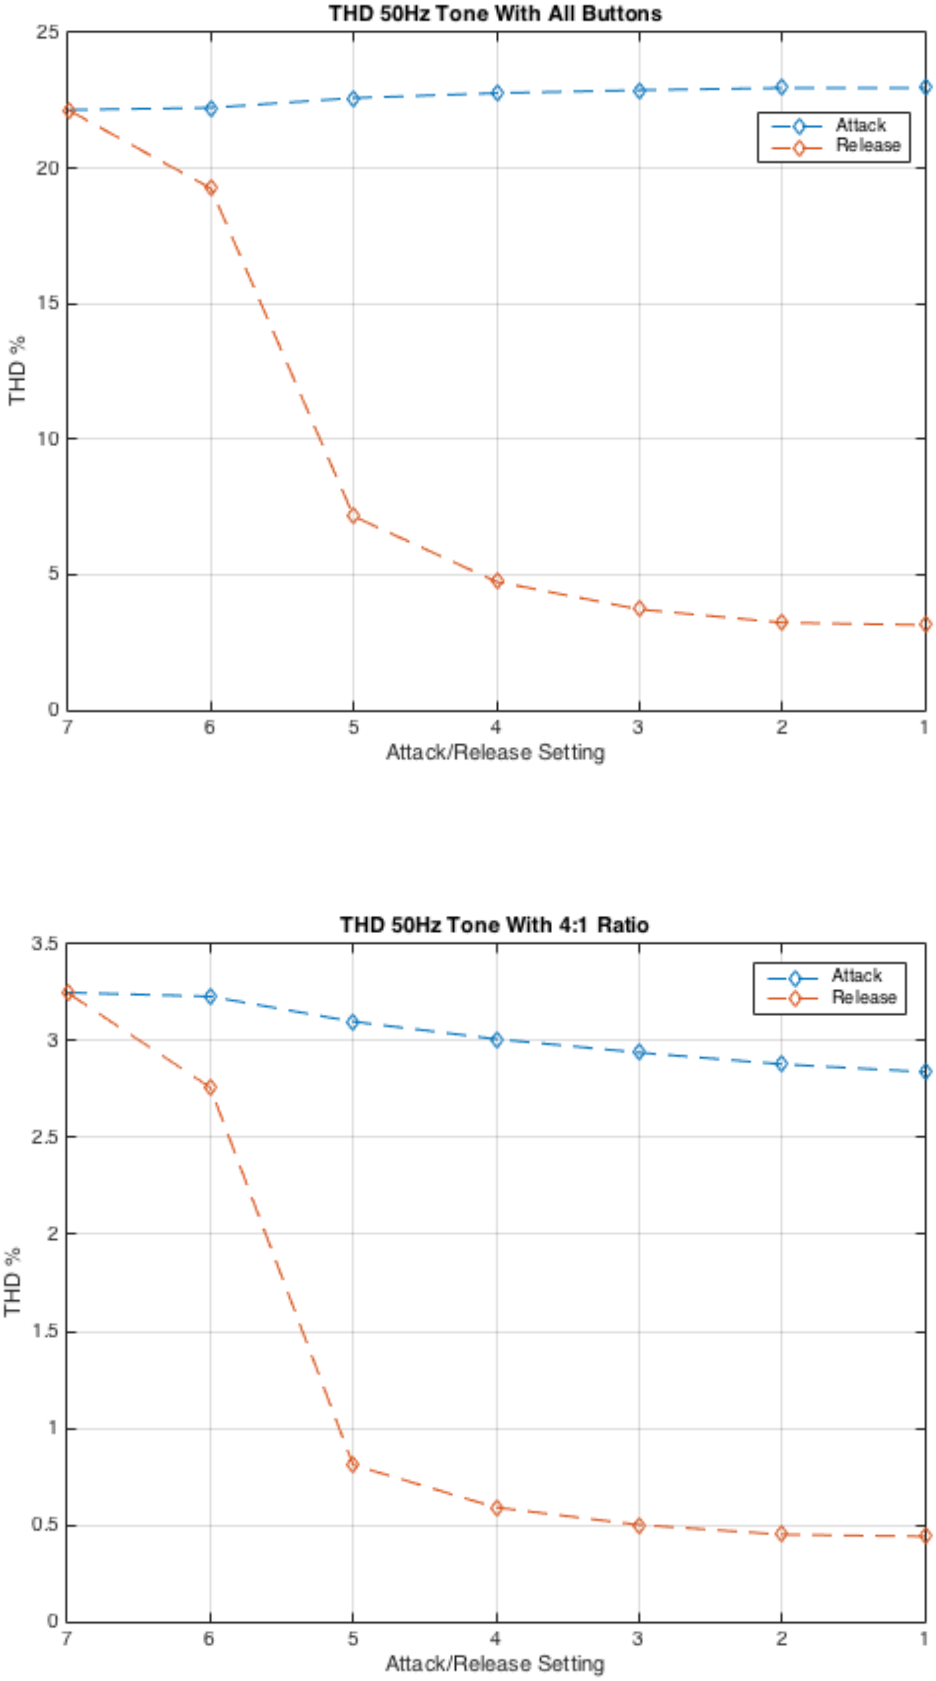

Figure 6-21: 83: THD as a function of attack and release for 50Hz tone. Top all buttons and bottom 4:1 ratio

6.8 Complex Program Material Analysis Introduction

As a continuation of the testing in the previous subchapters, complex program material was processed through the two 1176s and compressed with a number of popular time constant settings discovered from analysis in Chapter 4. The program material used in Chapter 5 was used again to allow for consistency between the tests.

In addition to analysis of popular sources, it was decided to run a test on a source shown to be less popular in Chapter 4's study. In this case, snare drums were used to give insight into why the 1176 is not commonly used on drum sources for transient shaping, particularly when the goal is to accentuate it. See Chapter 4 for more detail on this point. With that in mind, a snare track was processed through the 1176 and then through two popular VCA based compressors to compare the range of transient shaping options available on the devices. The VCA compressors used in this experiment were the on-board channel compressor on the SSL G Series console and the dbx165A, two much used VCA compressors for transient shaping techniques. Additionally, the snare material was processed through a third 1176, the limited edition 1176AE. This 1176 has a special attack mode called "slo" (intentional spelling) that allows for attack times in the region of 15ms.

Audio processed through the two 1176 compressors was set to have in the region of -6dB to -12dB of gain reduction showing on their VU meters. Heavy gain reduction was used as it was shown previously that large amounts of gain reduction is needed to make sonic signatures more discernible.

The analysis in the forthcoming sub-chapters takes on a similar format to Chapter 5. Firstly, the audio is analysed by detailing some of the characteristics that are audible through listening and the discussion is supported with spectrogram and time domain plots. Secondly, a number of audio features are analysed for information considered relevant to the study of sonic signatures. To conclude the chapter a test exploring the 1176's attack time is presented before the results of the tests made using snare material is discussed.

6.9 The 1176 On Vocals

The vocal track was processed using three time constant settings (TC) based on the review in Chapter 4. These parameters represented two common approaches to compressing vocals. TC 1 and 2 are typical vocal compression settings while TC 3 is a more aggressive limiting setting.

The time constants can be seen in Table 6-5. TC 1 is a setting found to be favoured by engineers in Chapter 4, and TC 2 is an often-quoted setting called "Dr. Pepper". The only thing that differentiates the two settings is the release speed. Therefore, it was thought these two settings would make for a good comparison, as they highlight the effect of release on vocal based material. Each time constant setting will be discussed separately in their sub-chapter. The time domain and spectral plots for the uncompressed audio can be seen in Figure 6-22 and the compressed audio plots for TC 1 is depicted in Figure 6-23. The portion of audio extracted for analysis is the same as the audio analysed in Chapter 5. Amplitude statistics for all the time constants are presented in Table 6-6.

Time Constant	Attack	Release	Ratio
1	3	7	4:1
2	3	4	4:1
3	7	7	20:1

Table 6-5: Compression settings used on the vocal material

Compressor	Dynamic Range (dB)
Reissue TC1	37.83
Reissue TC2	41.91
Reissue TC3	37.95
Blackface TC1	39.80
Blackface TC2	41.75
Blackface TC3	40.34
No Comp	44.84

Table 6-6: Amplitude statistics generated from the vocal material. The table includes all settings for both compressors

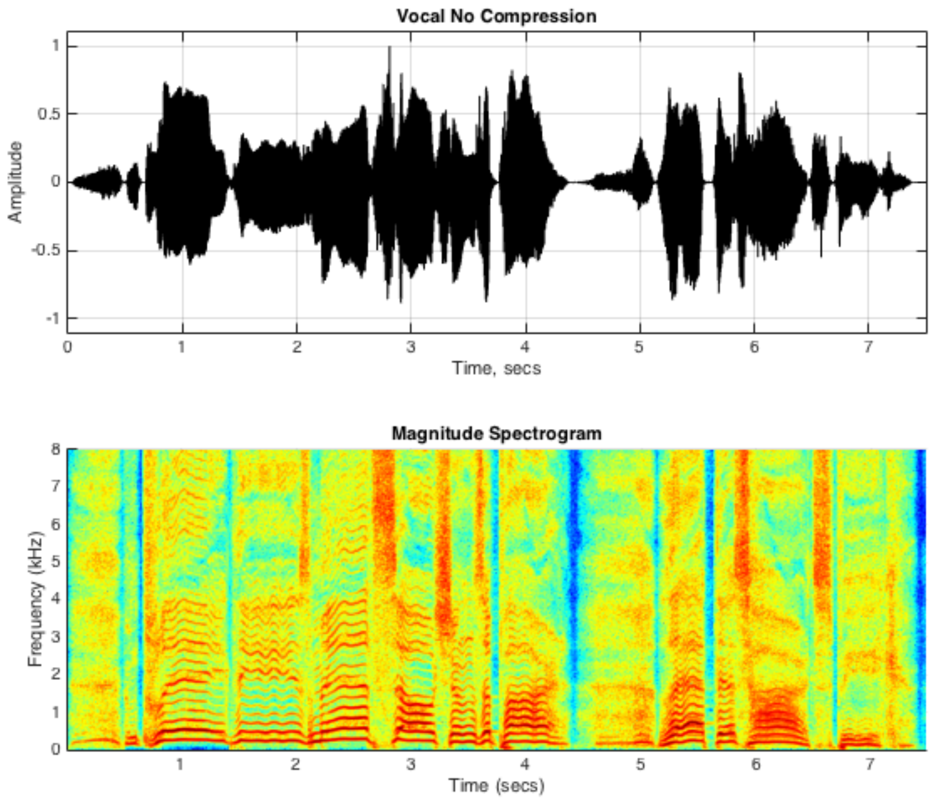

Figure 6-22: Time domain and spectrogram plots of the uncompressed vocal audio material

6.9.1 Vocal Time Constant 1

Comparing Figure 6-22 with 6-23 and listening to the audio files illustrates how aggressively TC1 has compressed the vocal. The majority of low-level breath sounds are now significantly louder than before. Raising the breathiness in the vocal is at least in part responsible for the "in your face" description of DRC that producers used in Chapter 4. The vocal is also more consistent and strident than the uncompressed version, this can be heard in the word "these" and is visible in the time domain plot at approximately 1.5 seconds. Regarding the spectrograms, there is more broadband energy in the breathing and the high frequency band, between 7-8kHz, has more energy. An increase in energy in this area results in a brighter overall timbre, and this can be heard in the audio files where the compressed versions have more bite and presence than the uncompressed material. This brightness seems to relate to producers using descriptors such as "edge" in Chapter 4. The consistency may be another reason why producers like the sound of the 1176 on vocals, it places the audio "in your face" and gives it "attitude." These were descriptors used by the producers in Chapter 4.

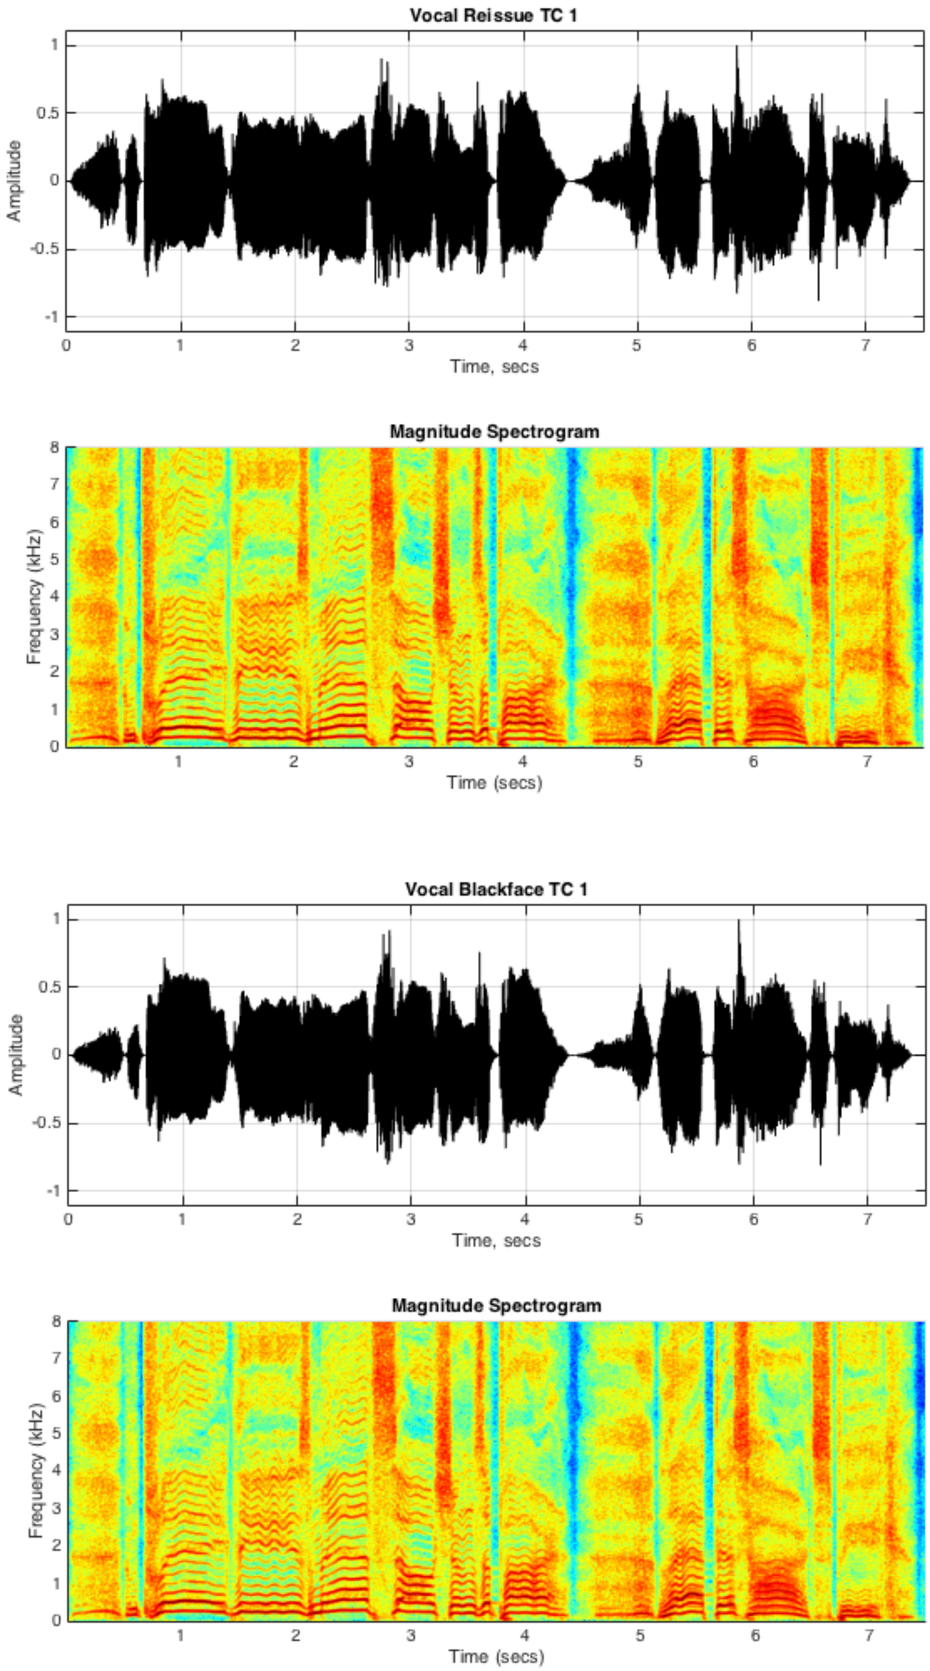

Figure 6-23: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) vocal audio TC1

Inspecting the overall macrodynamics of the part in the time domain plot reveals the vocal has been tracked musically by the compressors, meaning there is still good dynamic shape left in the vocal. It is compressed and obviously processed but not dynamically squashed. When comparing the two units, the Reissue is working more aggressively, and the compression effect is more pronounced, both audibly and visually. The VU meters were set for the same amount of gain reduction, so the difference may be as a result of the meters not having the same ballistics. The amplitude statistics generated from the audio in Table 25 seems to substantiate this point. Close inspection of the amplitude reveals that both compressors are similarly tracking the phrases.

The RMS envelope of the signals is shown in Figure 86 below where the compressors are indeed similar and any differences typically appear as slight fluctuations in microdynamics.

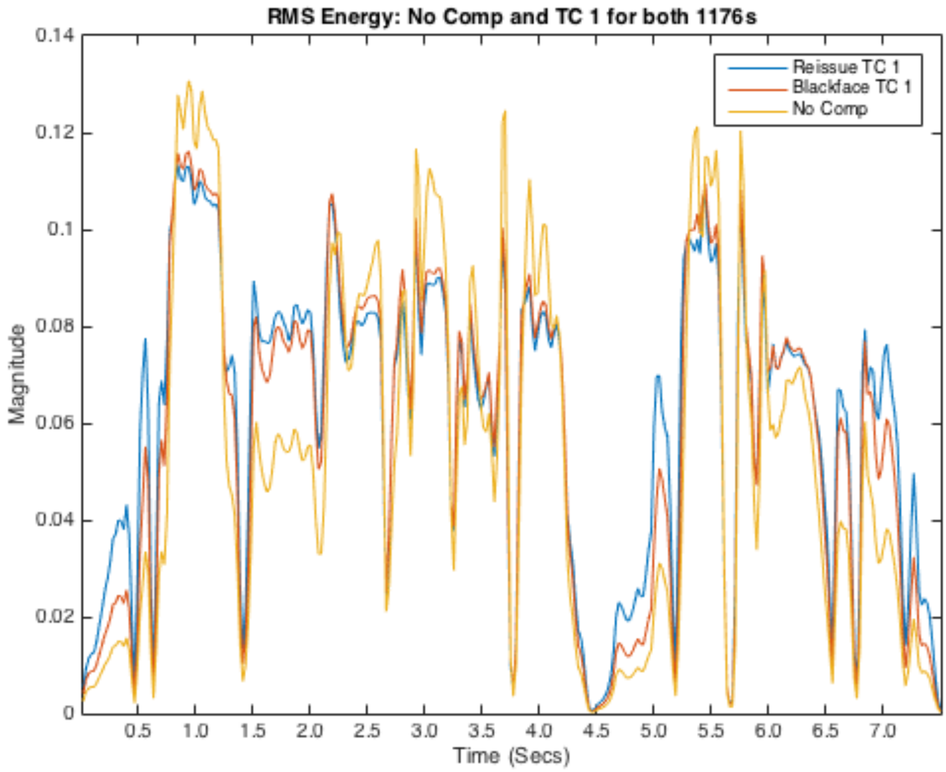

The LTAS for the compressors is shown in Figure 87, and there are some subtle differences in the bottom and top end of the two compressors with the Reissue being marginally brighter and having more energy in the low end.

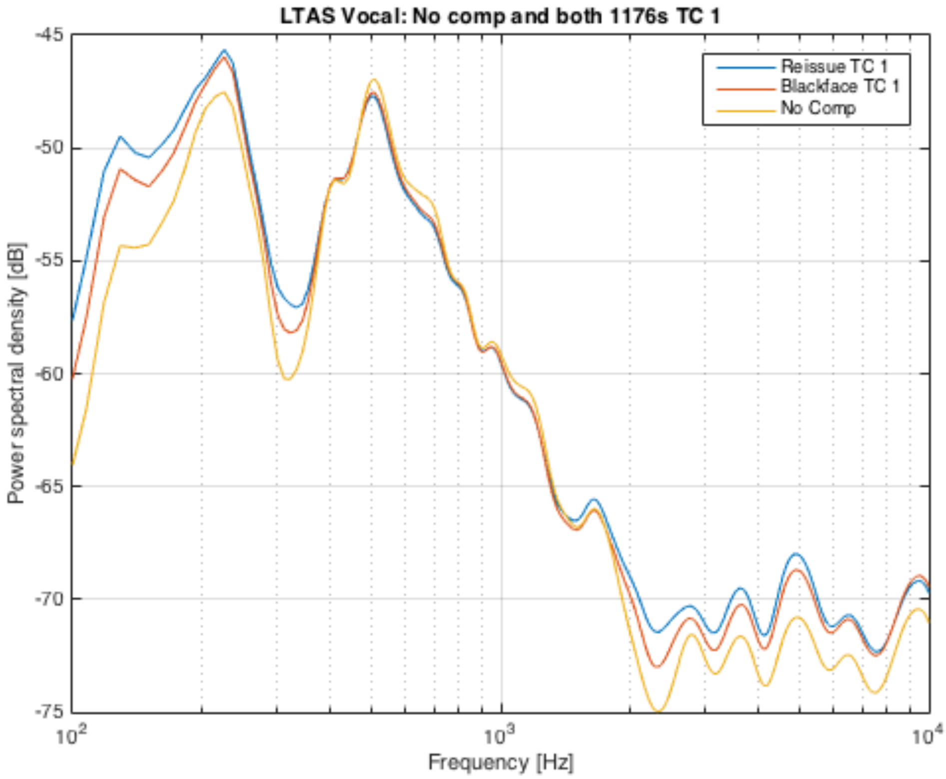

Figure 6-25: LTAS for both 1176s and uncompressed TC1 vocal

6.8.2 Vocal Time Constant 2

Time constant 2 differs from TC1 by making use of a slower release time. Hicks (2012) refers to this setting as "Dr.Pepper" and states that it is popular for vocal compression. The time domain and spectrogram plots for this time constant are illustrated in Figures 90 and 91. Visually the setting has a similar response to TC 1 and comparison of the time domain plots for TC1 and TC2 reveal similarities, but TC2 has some subtle differences, primarily around the end of the audio from five seconds onwards where the phrase has a wider dynamic range. However, the difference is small and most of the overshoots in level are approximately 0.2dB greater than TC 2.

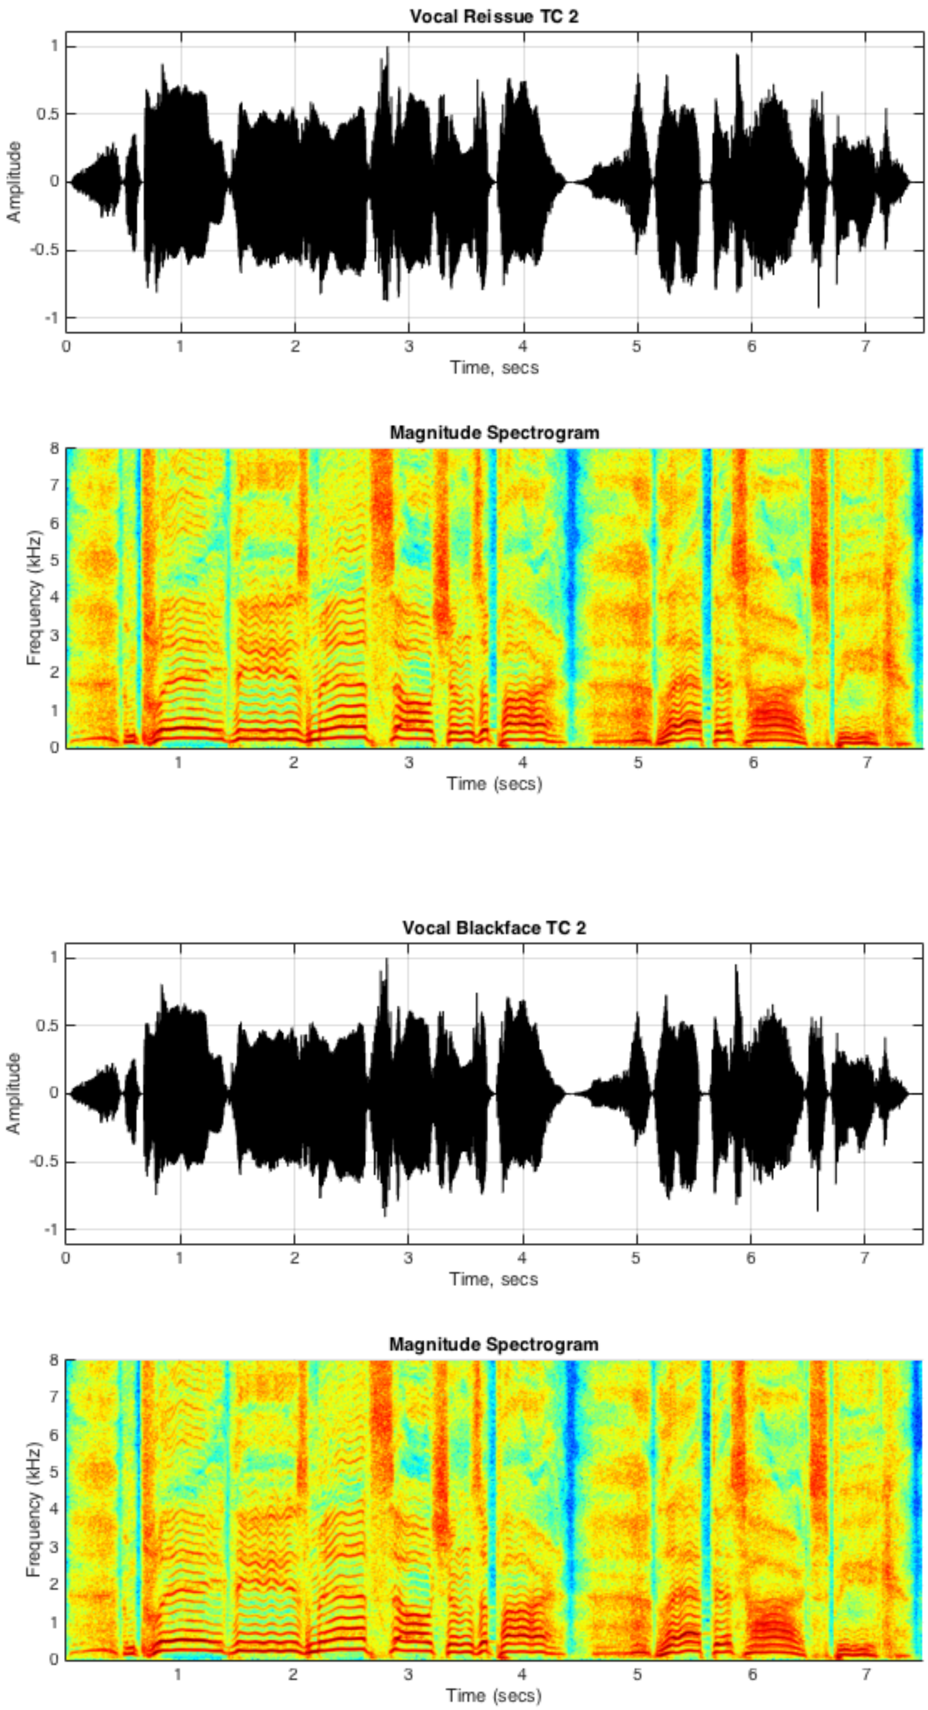

Figure 6-26: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) vocal audio TC2

An inspection of Figure 6-31 that shows the RMS envelope plots for all the TC settings for the Reissue substantiates this point. It suggests that the differences are mainly due to TC2's wider dynamic range. Sonically TC 1 sounds fuller and has more body than TC2 thus its reduction in dynamic range and faster release has produced a thicker texture than the other setting. The amplitude statistics for TC 2 show a wider dynamic range than TC 1, and this correlates with what is being heard in the audio.

Comparing the audio processed through the two compressors reveals a little amount of difference. They both similarly track the envelope, and this is evident in the RMS plot in Figure 6-27 where the main difference is in some overshoots on breath sounds.

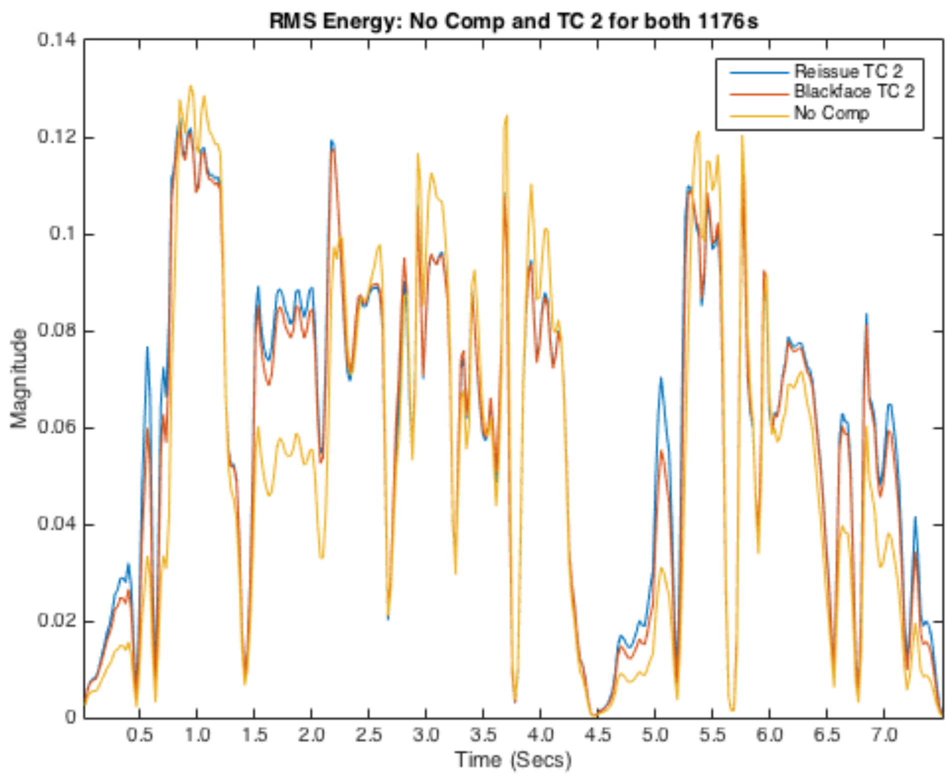

Figure 6-27: RMS energy both 1176s and uncompressed TC2 vocal

Spectrograms in Figures 6-23 and 6-26 illustrate the similarity between TC1 and TC 2. They show that small variations in energy around the breaths at the start of the phrase at 4.5 seconds are the main differences between the two settings. The LTAS for TC2 is displayed in Figure 6-28 and highlights there are only subtle differences in the top and bottom end between the compressors. The Reissue is shown to have a small amount of extra energy in the high and low end when compared with the Blackface.

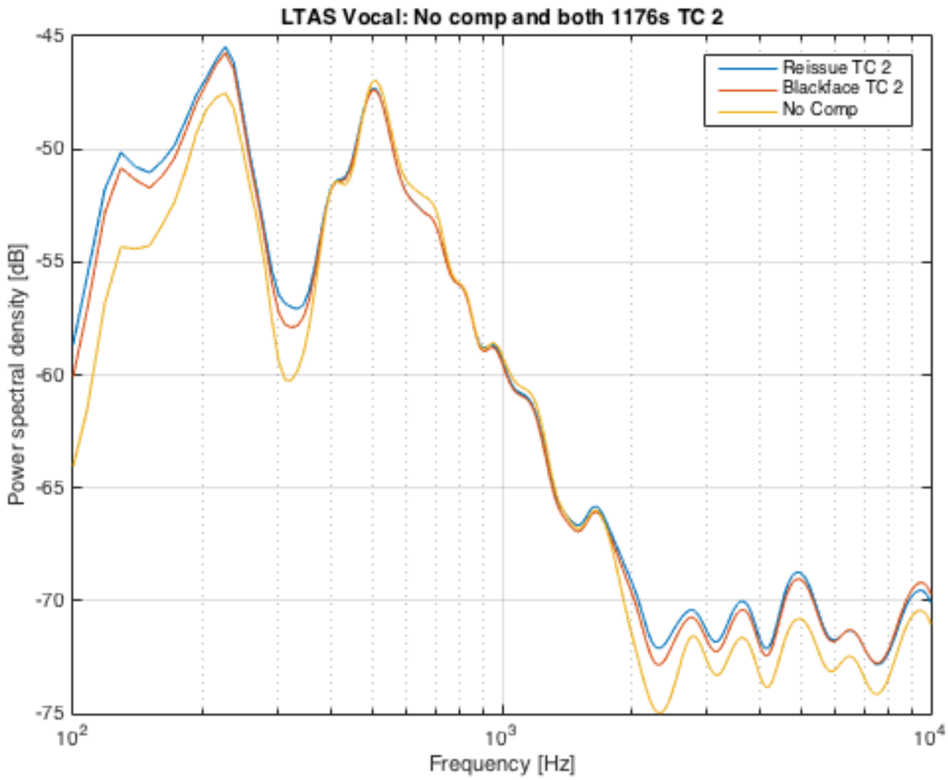

Figure 6-28: LTAS for both 1176s and uncompressed TC2 vocal

6.8.3 Vocal Time Constant 3

The third time constant setting was used to test the effect the 1176 had on audio when working as an aggressive limiter. In this test, time constants were set to their fastest positions and the ratio at 20:1.

Figure 92 shows the compressors are operating harder in this setting and listening to the audio reveals the compression effect is more audible. The overall macrodynamics of the phrase are more consistent with the fast attack and release times allowing the compressor to quickly grab and release the audio for aggressive limiting. Regarding the microdynamics, there is more consistency, the sustained note between 1.5-2.5 seconds has less fluctuation in level.

Focusing on the frequency domain in the spectrograms illustrates there is some difference compared to the original. Perceptually the differences are more audible in the lower end where it sounds like additional fullness has been added to the voice. This thickness has been present in all time constant settings and the results of this test suggest it is an integral part of the 1176's sonic signature.

Comparing the two compressors reveals differences are mainly due to variations in gain reduction. The amplitude statistics generated for the audio depict the dynamic range is approximately 3dB less for the Reissue. Therefore, it appears this compressor is attenuating the signal more aggressively than the Blackface despite both units showing the same amount of gain reduction on their VU meters. This result may be due to calibration issues, possibly with the FET in the meter section.

Additionally, the attack time may be playing a role in the smaller dynamic range results for the Reissue. As noted previously, it has a more aggressive attack characteristic than the Blackface thus the attack time may be allowing the Reissue to attenuate the audio more than the Blackface. However, the results for TC 2, which is the only setting not to use the release at 7, shows very little difference between the two compressors, therefore, it is likely the attack characteristic is compounded with the use of a fast release.

The RMS Energy plots in Figure 93 show there is more consistency in the compressed audio regarding both macro and microdynamics. When comparing the two compressors, it can be observed that there are some slight differences between the plots but the differences are subtle, and the Blackface has a more varied dynamic range than the Reissue. This result correlates with the dynamic range statistics in Table 25.

To make comparison between settings easier the RMS energy plots for all three time constants of the Reissue have been collated and are presented in Figure 94. Here the reader can see the similarities and note how differences in the settings reveal themselves as slight variations in microdynamics and additional dynamic range.

The LTAS for TC3 is featured in Figure 95 and shows the Reissue has more bottom end than the Blackface and that both compressors have more energy in the higher end than the uncompressed audio. The LTAS for all time constants of the Blackface is depicted in Figure 96 where it can be seen how closely the frequency content for all of the settings is grouped together, illustrating how consistent the 1176 is when compressing vocals. This consistent behaviour over a range of settings is presumably one of the reasons why the 1176 was found to be a popular vocal compressor in Chapter 4.

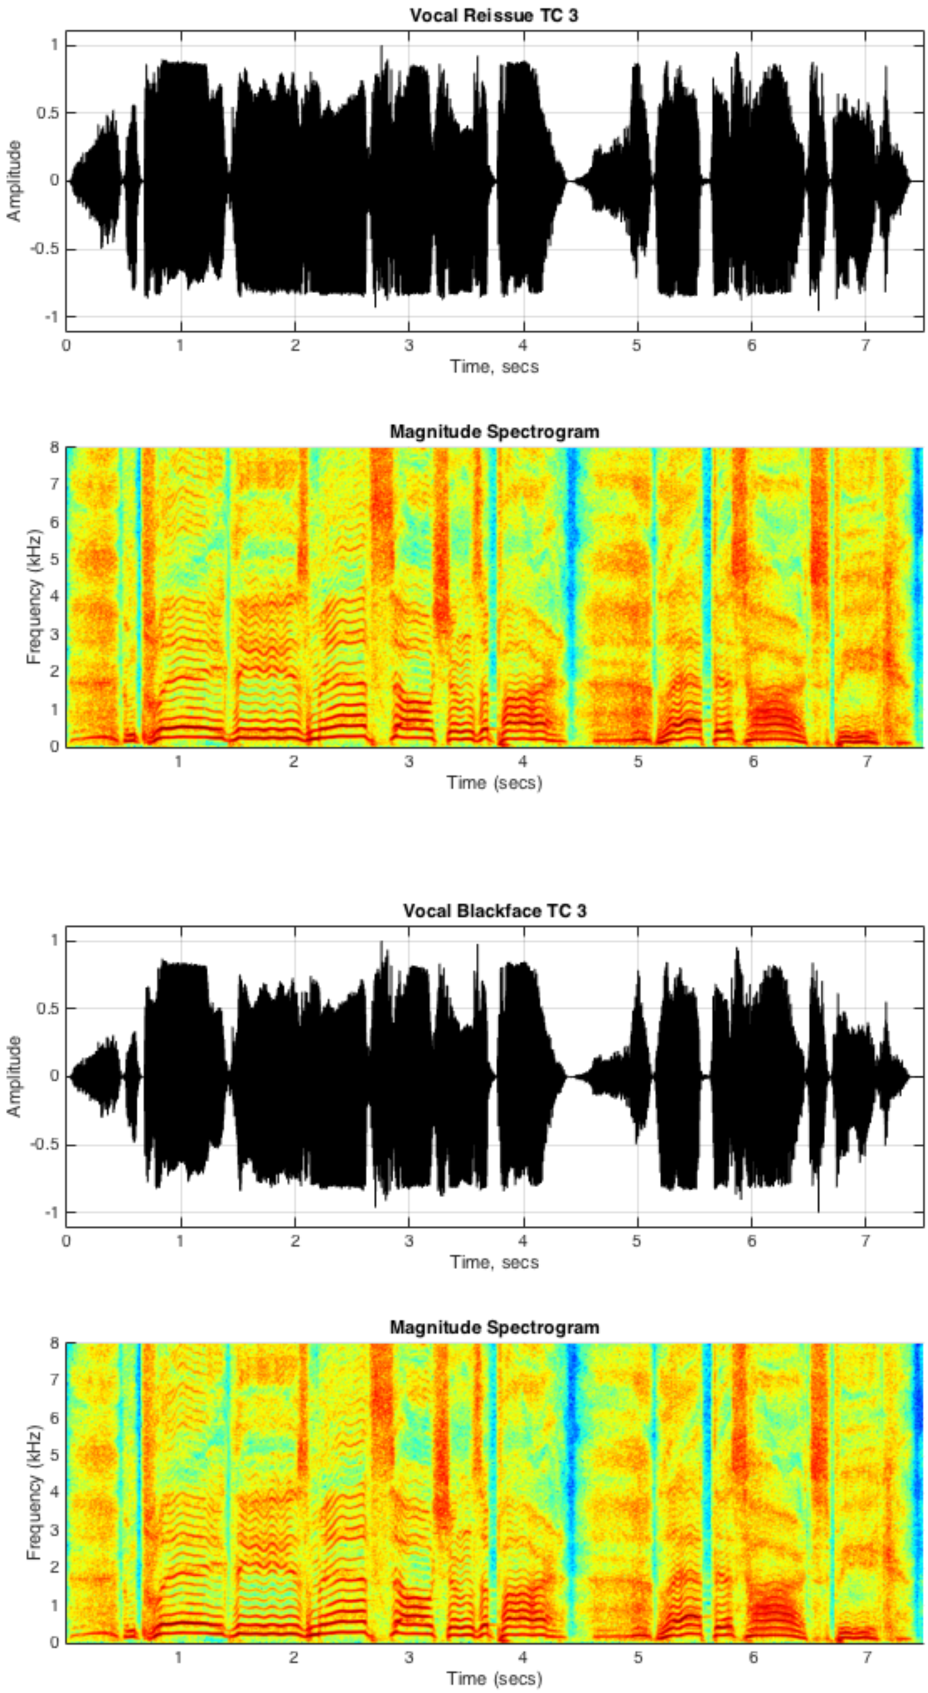

Figure 6-29: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) vocal audio TC3

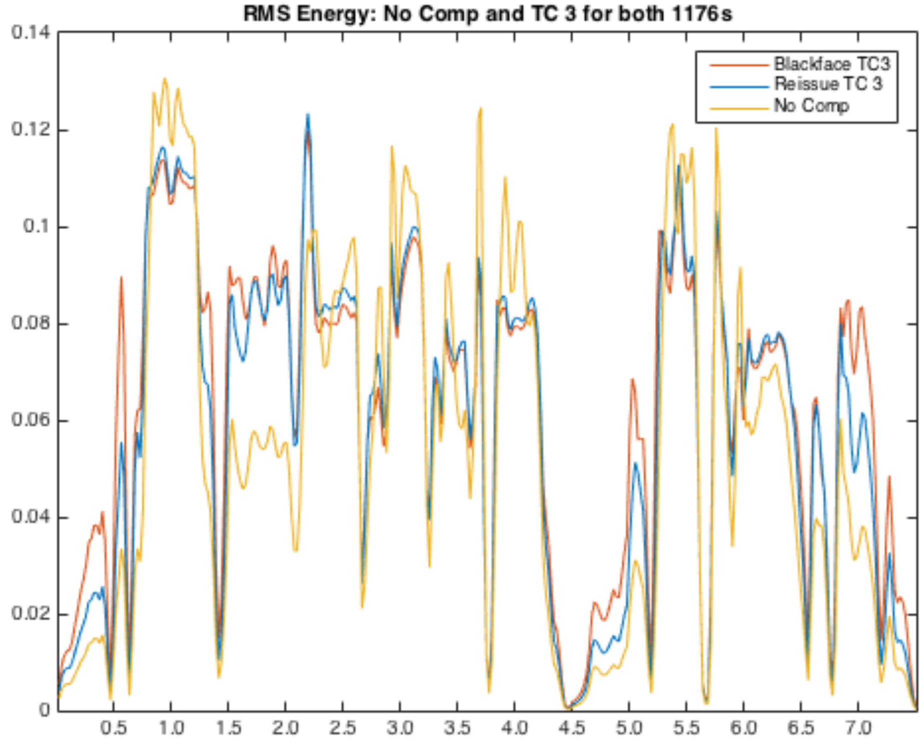

Figure 6-30: RMS energy both 1176s and uncompressed TC3 vocal

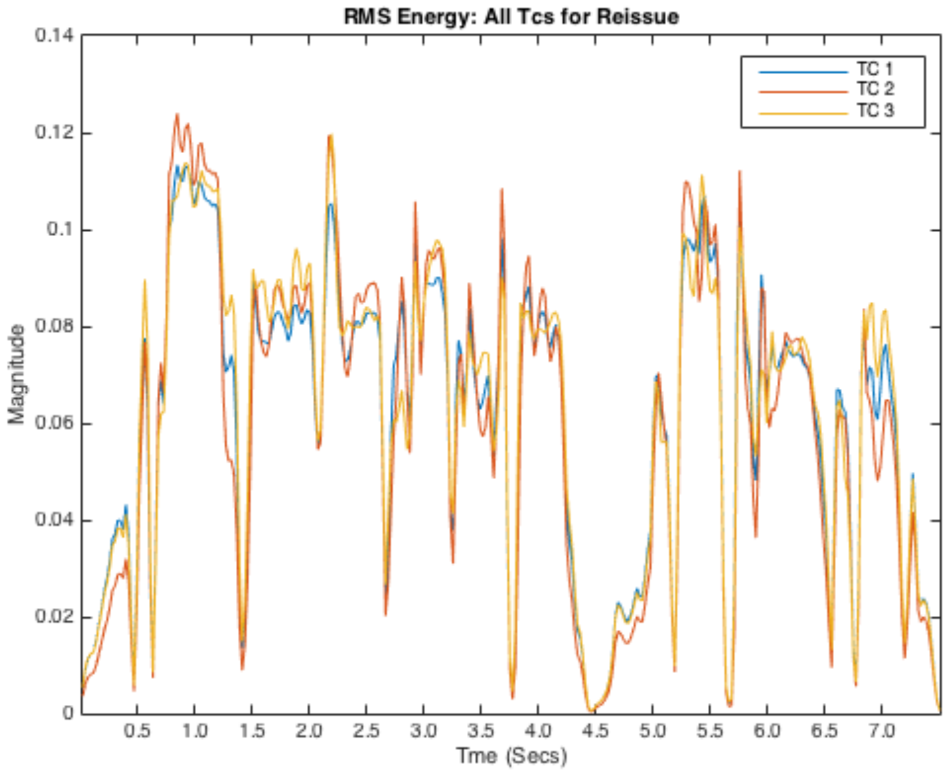

Figure 6-31:RMS energy for Reissue at all time constants

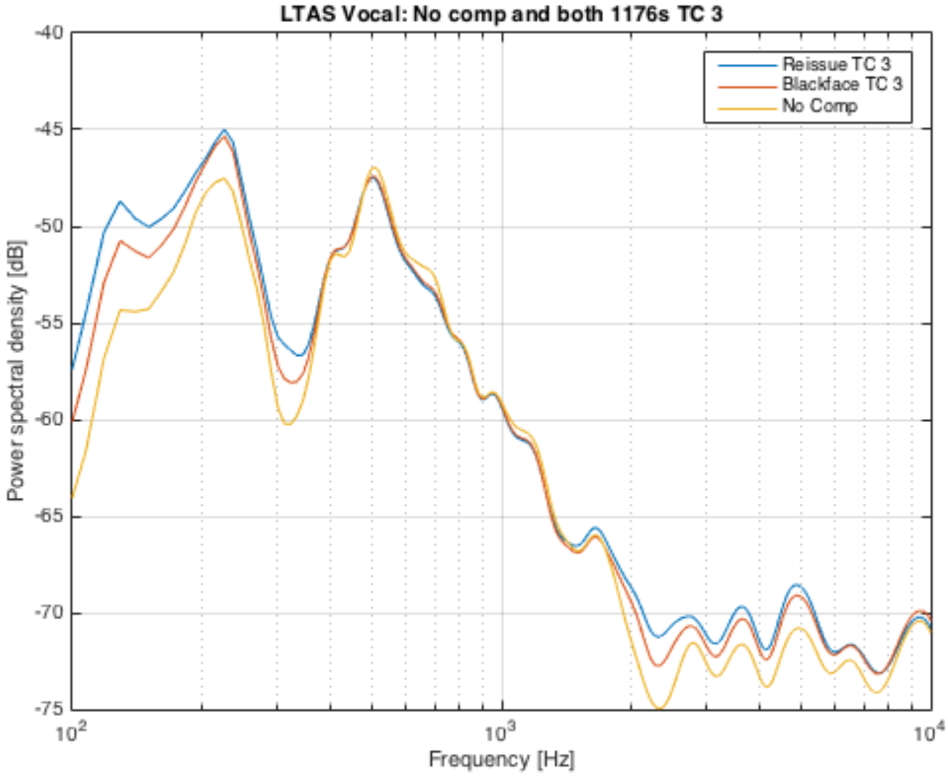

Figure 6-32: LTAS for both 1176s and uncompressed TC3 vocal

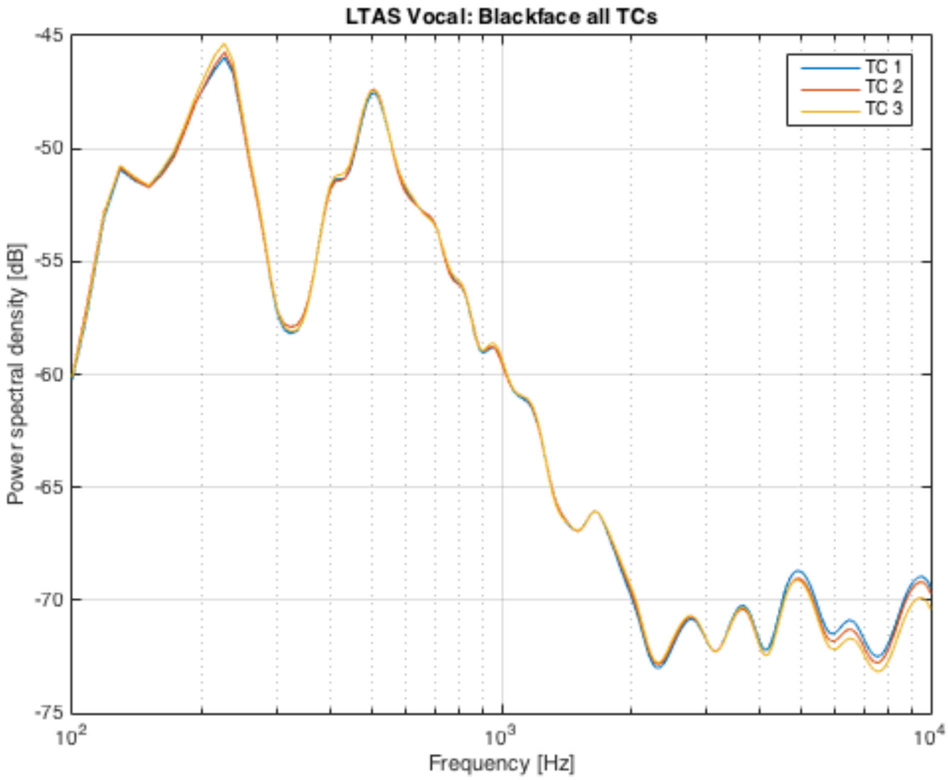

Figure 6-33: LTAS for Blackface at all time constants

6.9.4 Audio Features: 1176 on Vocals

To conclude the vocal testing a number of audio features were extracted from the audio to compare changes to the timbre and the results are displayed in Table 6-7.

Compressor	Roll Off	Spectral Centroid	Brightness	Low-Energy
No Comp	10197Hz	4655Hz	0.55	0.53
Reissue TC1	10332Hz	4894Hz	0.59	0.38
Reissue TC 2	10219Hz.	4746Hz	0.57	0.42
Reissue TC 3	10149Hz.	4631Hz	0.56	0.41
Blackface TC1	10537Hz	5078Hz	0.59	0.41
Blackface TC 2	10452Hz	4967Hz	0.58	0.44
Blackface TC 3	10366Hz	4827Hz	0.57	0.41

Table 6-7: Audio features extracted from the vocal material for all time constants, both 1176s and uncompressed material

All settings have the effect of increasing the perceived brightness of the audio with the most significant difference been in the first time constant setting. TC 3 has the lowest brightness, spectral centroid and roll off value for both 1176s. Thus it can be argued the fast attack time and ratio combination used in this setting has resulted in a duller sonic signature. This result is unexpected as it was predicted this setting would increase features relating to perceptual brightness due to the fast compression speed and high ratio adding non-linearity. The highest value for spectral centroid, brightness and roll off can all be found in the TC 1 setting. These results correlate with the LTAS in Figure 6-33 that shows TC1 has the most high end of all the settings tested here and TC 3 the least. There is a small reduction in spectral centroid for the Reissue TC 3 compared with the uncompressed material yet the brightness feature is marginally higher than the original. The reasons for this are not clear. Nonetheless, the increase in the spectral centroid and brightness results appears to be creating some of the perceptual colouration that producers referenced in Chapter 4 as a motivation to compress vocals.

The lowest low-energy result for both the compressors is not TC3 as one might expect but TC1, although the difference between TC1 and TC3 for the Blackface is small. Chapter 4 showed TC1 is a popular vocal compression setting. Therefore, the increase in brightness and fullness may be partly responsible for its popularity.

Spectral flux in the 100-200Hz band was extracted from the vocal material to investigate if it correlated with the textural thickness heard in the audio. The results are presented in Table 6-8 and show all time constant settings have increased spectral flux in this band, but there is no clear trend with regards to which time constant has the most significant effect. The only conclusion that can be made from the data is that fast release times result in higher values. One can state this assertion because the lowest result for all compressors is the only setting to make use of a release time slower than position seven, the fastest release offered by the 1176.

Compressor	Spec Flux Mean	Spec Flux Std
No Comp	8.41	11.18
Reissue TC1	10.57	13.86
Reissue TC 2	10.00	14.00
Reissue TC 3	11.60	14.10
Blackface TC1	11.60	14.10
Blackface TC 2	10.18	13.67
Blackface TC 3	10.59	13.76

Table 6-8: Spectral flux for all TCs, both 1176s and uncompressed for vocal material extracted from the 200-400Hz sub band

6.10 The 1176 on Bass

A bass part was processed with three time constant settings found to be popular in use amongst engineers in Chapter 4. These parameters are included in Table 6-9 that shows the first two settings are the same as TC 1 and 3 employed in the vocal tests. The final TC makes use of the longest attack time offered by the 1176. This time constant was selected to observe the effect a long attack time had on the transient portions of the bass part. Case (2007) notes that long attack times can help in adding punch to bass. However, as will be seen in Chapter 6.13 the range of the attack time is very small therefore it was predicted that its effect would be minimal. Nevertheless, this setting was selected to observe how it affected the timbre of bass audio material. The time domain and spectrograms for the uncompressed audio material are illustrated in Figure 6-34 and amplitude statistics for all the settings are presented in Table 6-10.

Time Constant	Attack	Release	Ratio
1	3	7	4:1
2	7	7	20:1
3	1	7	4:1

Table 6-9: Compression settings used on the bass material

Compressor	Dynamic Range (dB)
No Comp	22.24
Reissue TC1	16.01
Reissue TC2	17.63
Reissue TC3	16.13
Blackface TC1	17.51
Blackface TC2	18.51
Blackface TC3	17.82

Table 6-10: Amplitude statistics generated from the bass material. The table includes all settings for both compressors

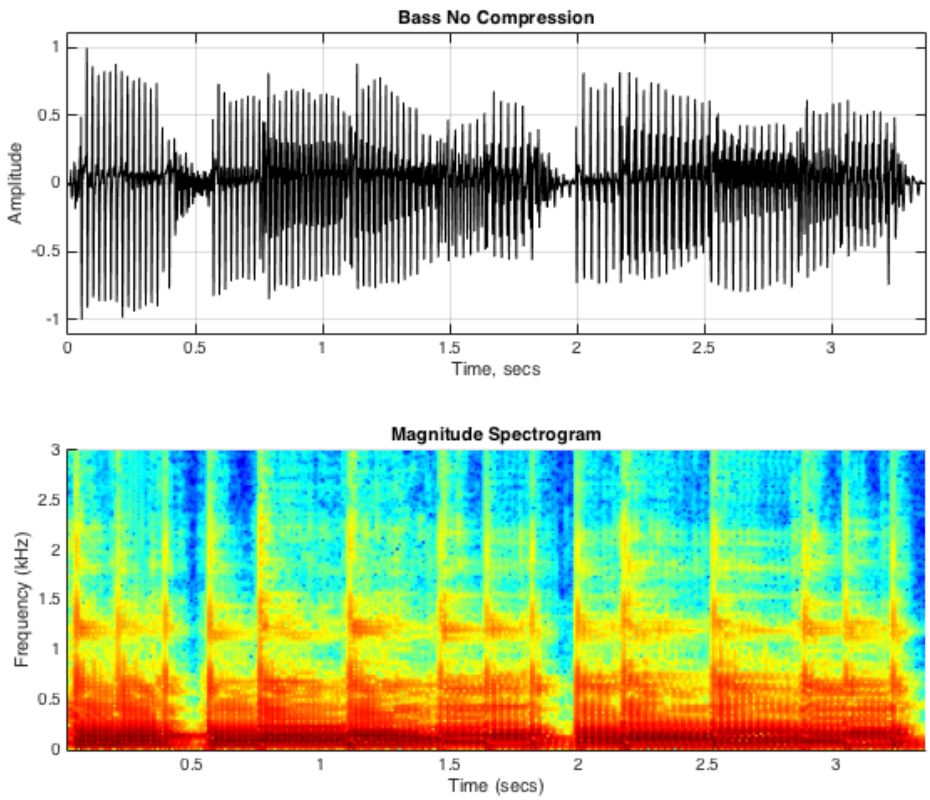

Figure 6-34: Time domain and spectrogram plots of the uncompressed bass audio material

6.10.1 Bass Time Constant 1

Figure 6-35 shows the time domain and spectrogram plots for this setting and the results are quite subtle. Comparing the uncompressed plots in Figure 6-34 with the compressed material reveals there are only small differences around the transients of notes. The trend with this setting is broad control of the dynamics, meaning the 1176 tracks the overall envelope of the audio to preserve its macrodynamic shape while raising some of the lower level material. The accentuation of low-level audio is due to the fast release time, and perceptually it gives the sound a consistent and aggressive character. This characteristic was noted as a positive sound quality by producers in Chapter 4. The amplitude statistics generated from this portion of audio are presented in Table 6-10 and substantiate the observations made from looking at the time domain plots. The dynamic range of the audio has been reduced by approximately 5dB.

The disparity between the two units is small with little audible differences. This similarity can be seen in the plots. The strange response exhibited by the Reissue in the program dependency tests (using this setting and a 50Hz test tone) has not noticeably affected the bass material. Thus, this behaviour may be sporadic and appear more easily with amplitude variations such as those in the test signal.

The RMS energy plots for the two compressors are shown in Figure 6-36 where again there is little difference between the units. There are some slight variations in the microdynamics of the envelope but these are not particularly significant. The difference between the uncompressed and the compressed envelopes is more visible with the uncompressed envelope showing heavier transients in the first two notes and more variation in dynamics across the full phrase. The attack of the 1176 is still fast despite being set at one of its slower speeds. Listening to the audio extract of the compressed material reveals the 1176 adds textural thickness, additional body and a sense of consistency to the bass part. These attributes are closer to colouration than a change in the envelope. Table 4 in Chapter 4 showed that colouration and control over the dynamic range were two important reasons to compress a bass instrument.

Figure 6-37 shows the LTAS for this time constant setting. There is little difference between the uncompressed and the compressed material bar a small amount of additional energy in the lower midrange and less high frequency in the compressed audio.

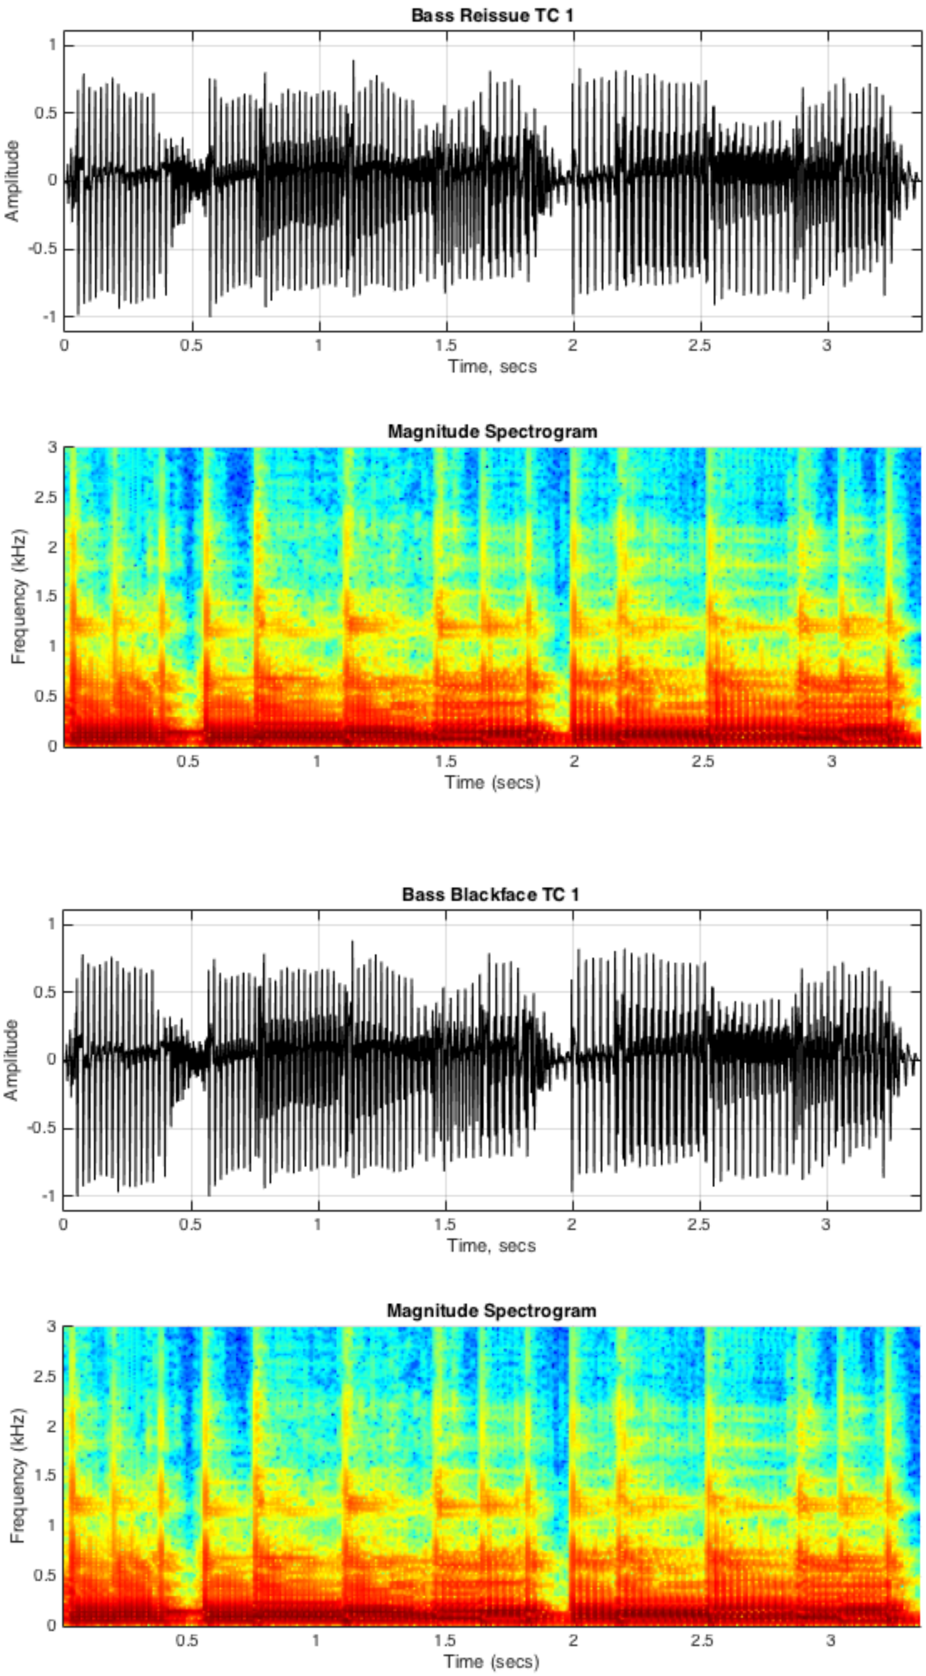

Figure 6-35: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) bass audio TC1

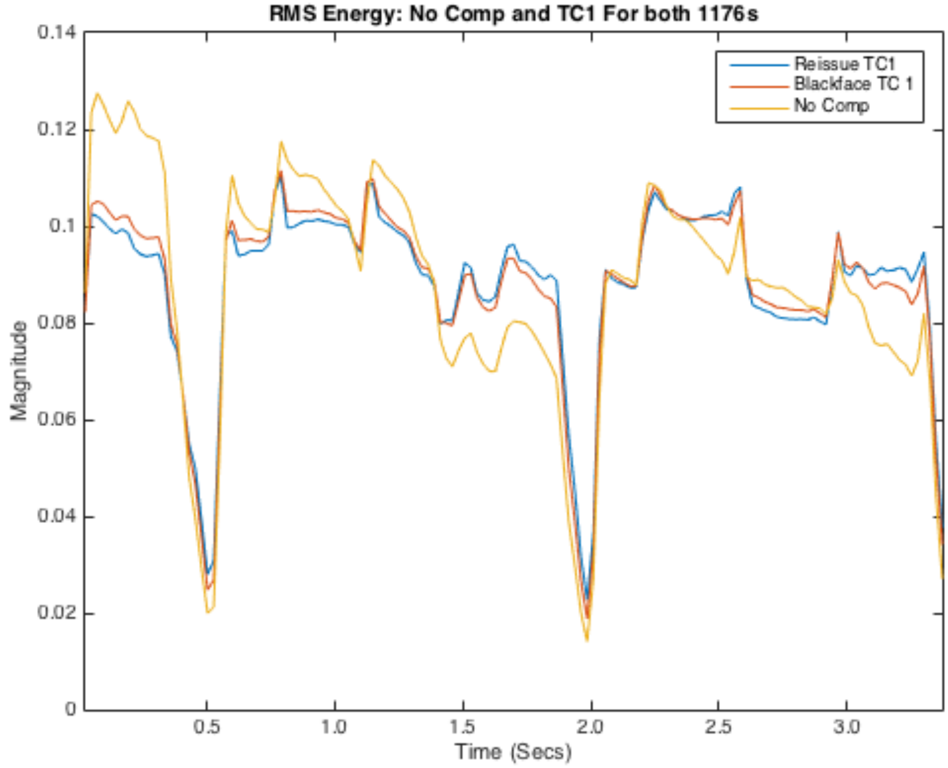

Figure 6-36: RMS energy both 1176s and uncompressed TC1 bass

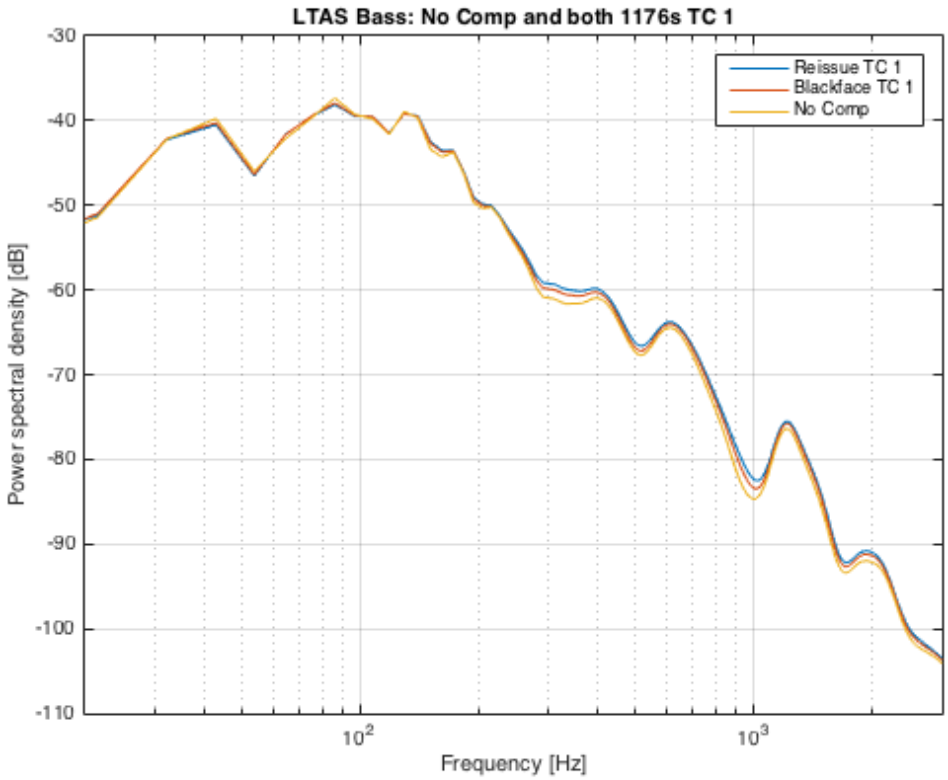

Figure 6-37: LTAS for both 1176s and uncompressed TC1 bass

6.10.2 Bass Time Constant 2

The second setting made use of the fastest available attack and release settings with a ratio of 20:1. This configuration was included to test for non-linearity when using fast time constants on bass sources. In Chapter 3 it was noted that settings such as this could result in low frequency distortion and the audio of this test reveals that distortion is audible. However, the audio still has its transients mostly left intact, the fast time constants are waveshaping and creating distortion but they are allowing much of the original audio envelope to remain unaffected by the compressor. This result is not something that would usually happen when a signal is distorted. It is more common for the transient portion to be softened. Table 4-2 in Chapter 4 showed that distorting bass instruments was an important consideration for producers, and this behaviour by the 1176 makes it a suitable compressor for this technique.

Figure 6-38 depicts the results of this setting and the additional harmonic content created by this configuration can be seen in the spectrogram plots, where the spectral content in the 2-3kHz range is more intense. As with the previous setting (and also in the vocal material), there is a sense of fullness to the texture of the audio. This suggests that fullness is an integral part of the 1176's sonic signature.

Listening to the audio material highlights that the compressors sound similar. There is no noticeable difference in sonic signature between the units and the clipping artefacts sound alike. An inspection of the time domain and spectral plots in Figure 6-38 reveals there is no discernible differences between the two compressors and substantiates the perceptual observation. The dynamic quality of this setting sounds similar to TC1 and to highlight the similarity the RMS energy for TC 2 is plotted with the Reissue's TC 1 in Figure 6-39 where it can be noted the envelopes are alike and include similar variations in microdynamics. RMS energy for TC 2 is plotted for the two compressors and the uncompressed audio in Figure 6-40. The wider dynamic range of the uncompressed material, especially during the first note, is clearly visible.

The LTAS for this setting is shown in Figure 6-41 where it can be observed there is a marginal change in high frequency content when compared with the uncompressed audio. The Reissue is slightly brighter than the Blackface in the LTAS around the 1kHz area.

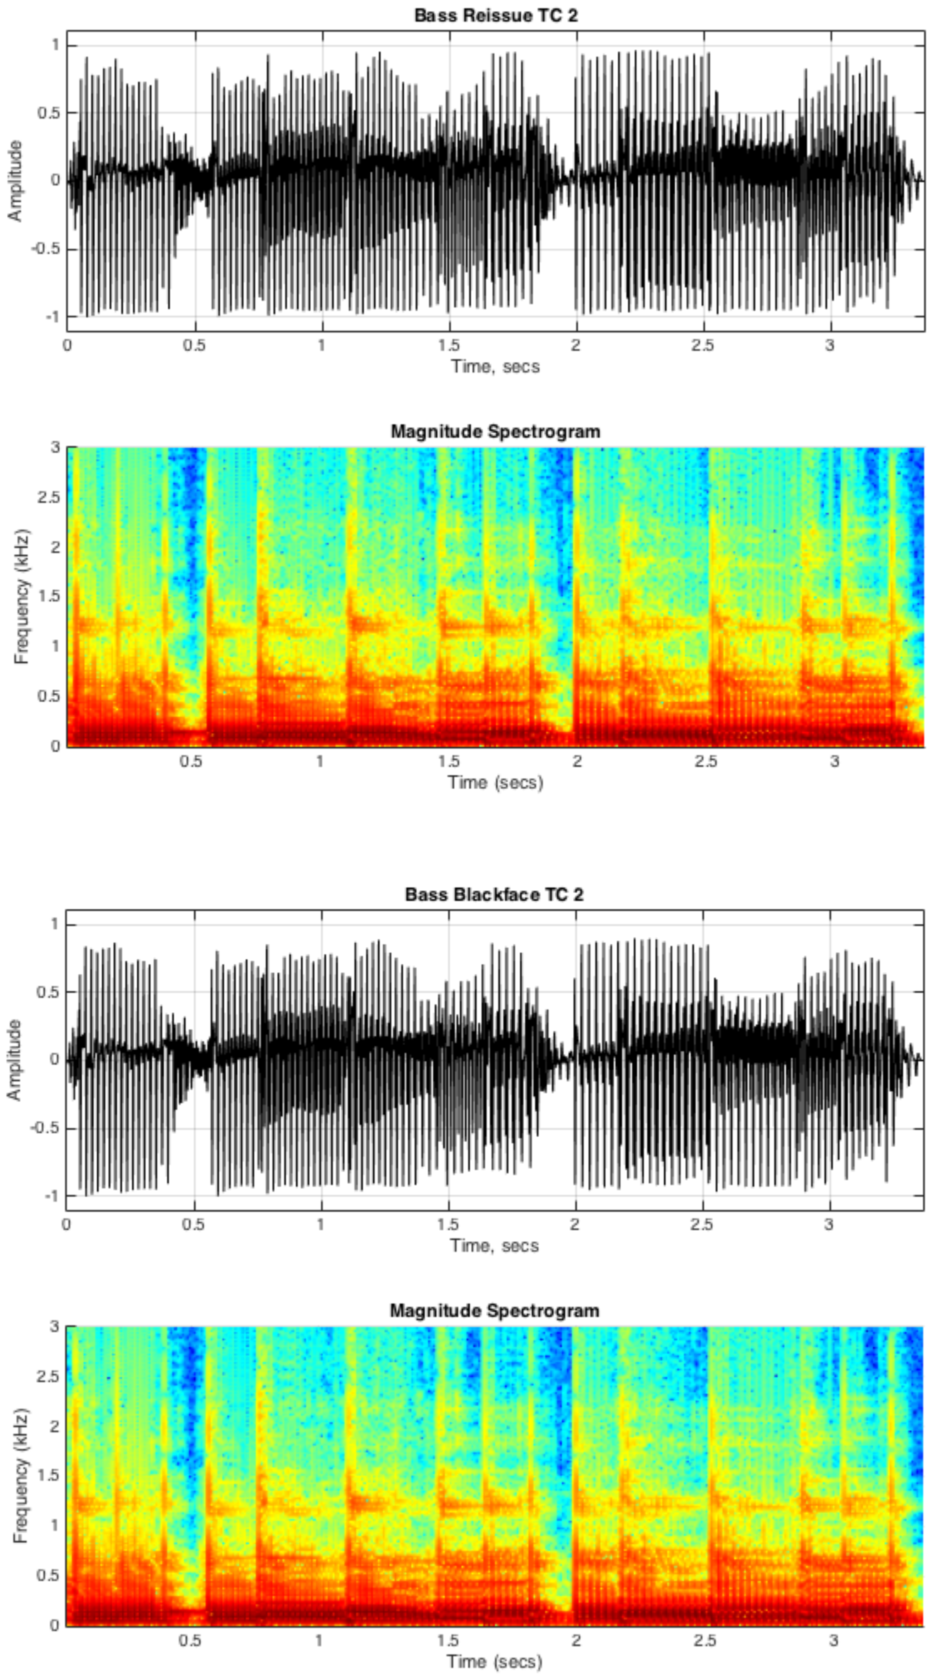

Figure 6-38: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) bass audio TC2

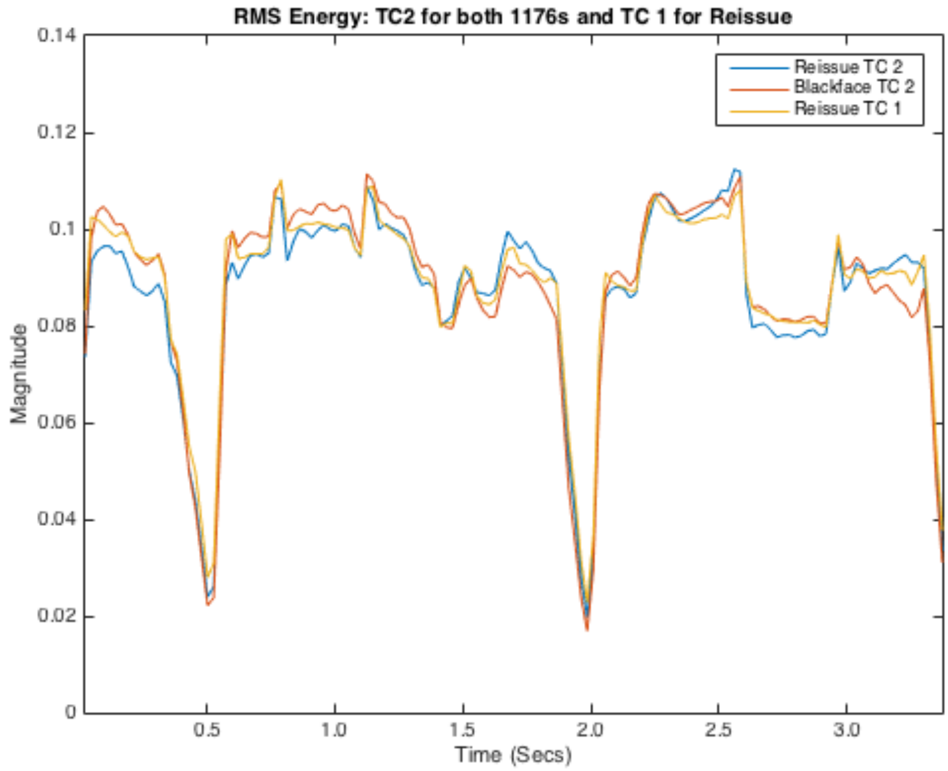

Figure 6-39: RMS energy both 1176s TC2 and TC1 for Reissue bass

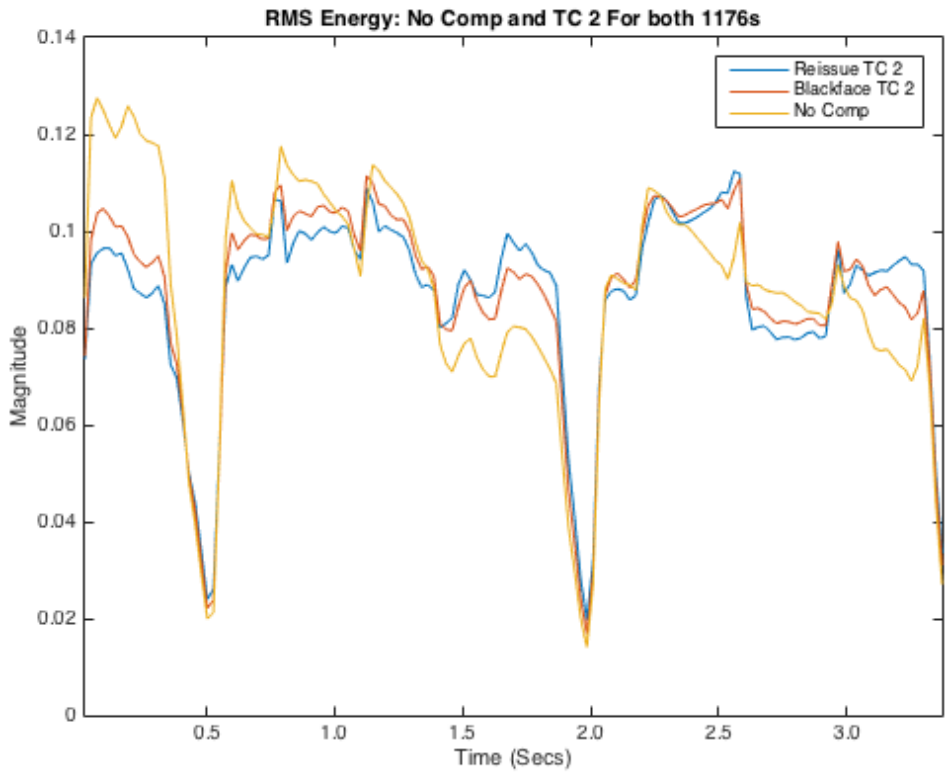

Figure 6-40: RMS energy both 1176s and uncompressed TC2 bass

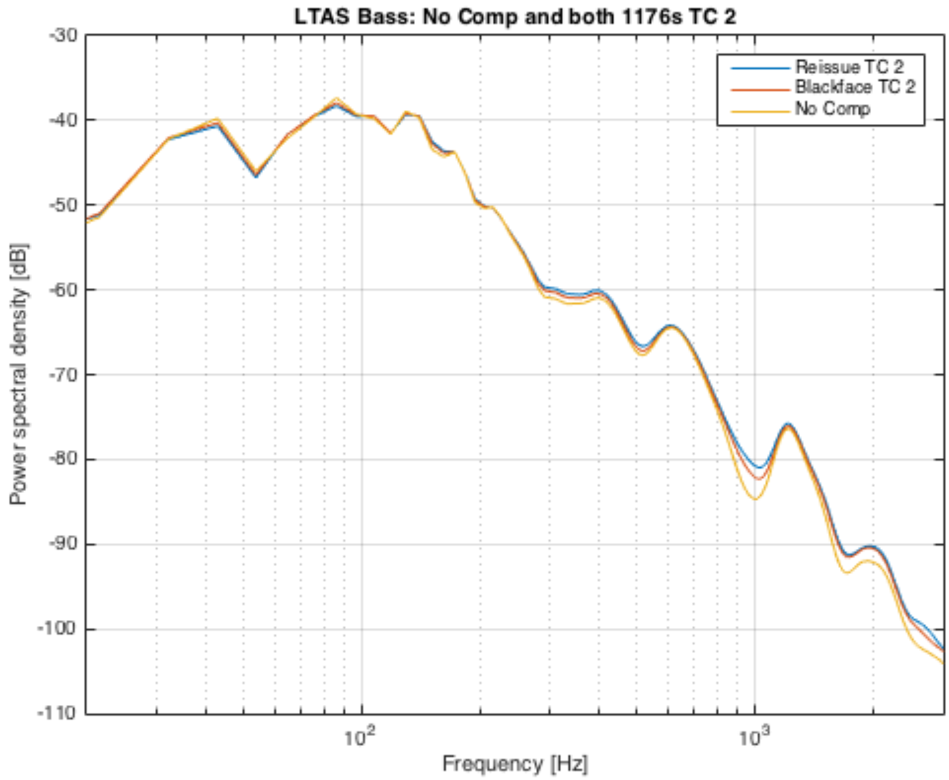

Figure 6-41: LTAS for both 1176s and uncompressed TC2 bass

6.10.3 Bass Time Constant 3

The final time constant tested on the bass made use of the longest attack in combination with the fastest release and a 4:1 ratio. This setting was selected because attack times between positions 1-4 were found to be the most common for 1176 bass compression in Chapter 4. The only thing that differs between TC1 and TC3 is the attack setting. Therefore, this setting highlights how the attack affects the timbre of audio with predominantly low-end frequency content.

The time domain plots in Figure 6-42 show that the results are similar to those presented in Figure 6-35, there is little noticeable differences between TC1 and TC3 in the time domain. The only slight difference is in the front end of the transient where it is more pronounced in material compressed with TC 3. This is to be expected, as the longer attack time will allow for more of the audio to pass through the compressor unaffected by gain reduction. However, the difference is small, with the first 6ms of the transient having approximately 0.3dB more amplitude in TC3 when compared with TC1.

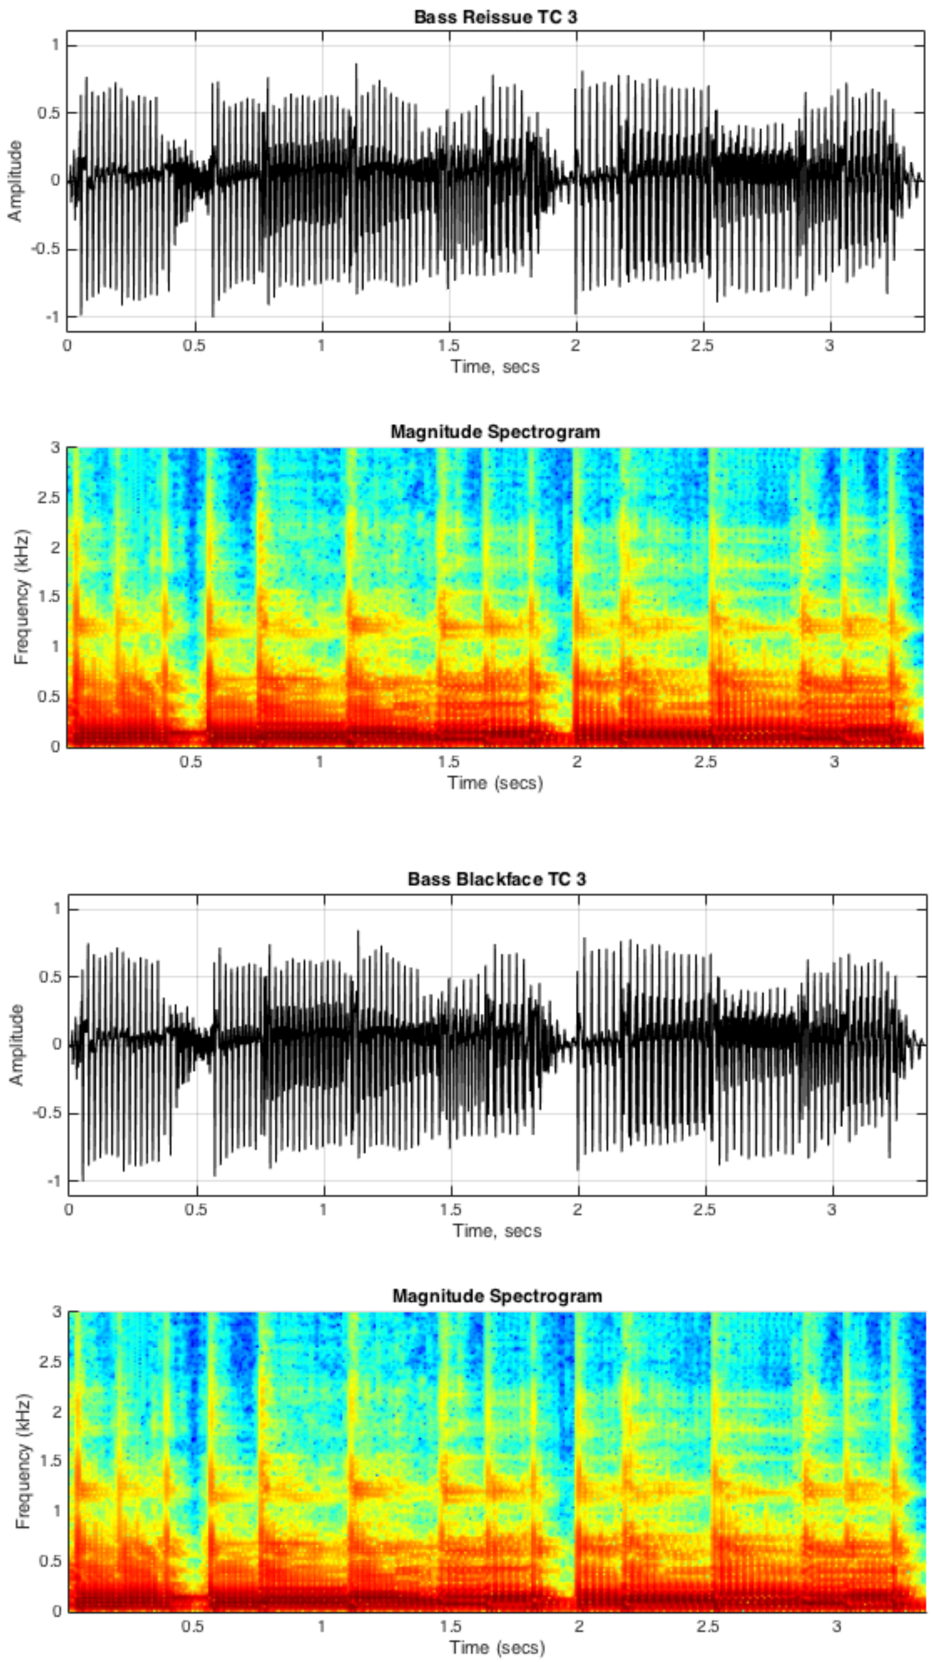

Figure 6-42: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) bass audio TC3

The spectral plot reveals a similar picture. The two settings have comparable frequency content with only marginally more energy between 1-1.5kHz for TC 3 to differentiate the two settings. A comparison of the two 1176 units reveals close similarity. The only difference being at the transient of notes where the Reissue has a little more amplitude. This was measured as approximately a 0.2dB difference in the first 6ms of the first transient thus the difference is largely insignificant. The effect of the attack time is minor in this test and is commensurate with the results of a more thorough attack time test discussed in Chapter 6.13.

RMS energy has been plotted for TC 3 in Figure 6-43 and shows a small difference between the compressors with the Blackface having a wider dynamic range than the Reissue. RMS energy for TC 1 (the Reissue) and TC 3 (both compressors) is depicted in Figure 6-44 and reveals the envelopes for TC 1 and 3 are very similar, particularly when comparing the results for the Reissue.

The LTAS for TC 3 is shown in Figure 6-45 where little difference is visible. A LTAS focused between 500-3000Hz has been plotted for all time constants and the uncompressed audio and is shown in Figure 6-46. This result shows the setting making use of a 20:1 ratio (TC2) has the most energy from 1kHz. This is due to the high ratio and fast time constants creating non-linearity.

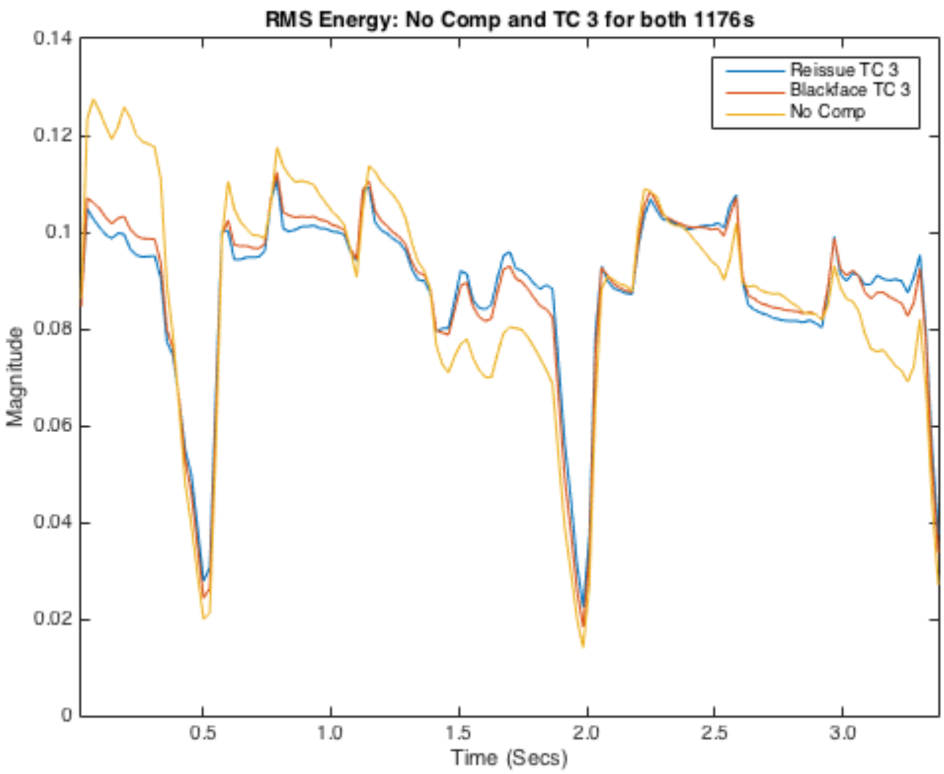

Figure 6-43: RMS energy both 1176s and uncompressed TC3 bass

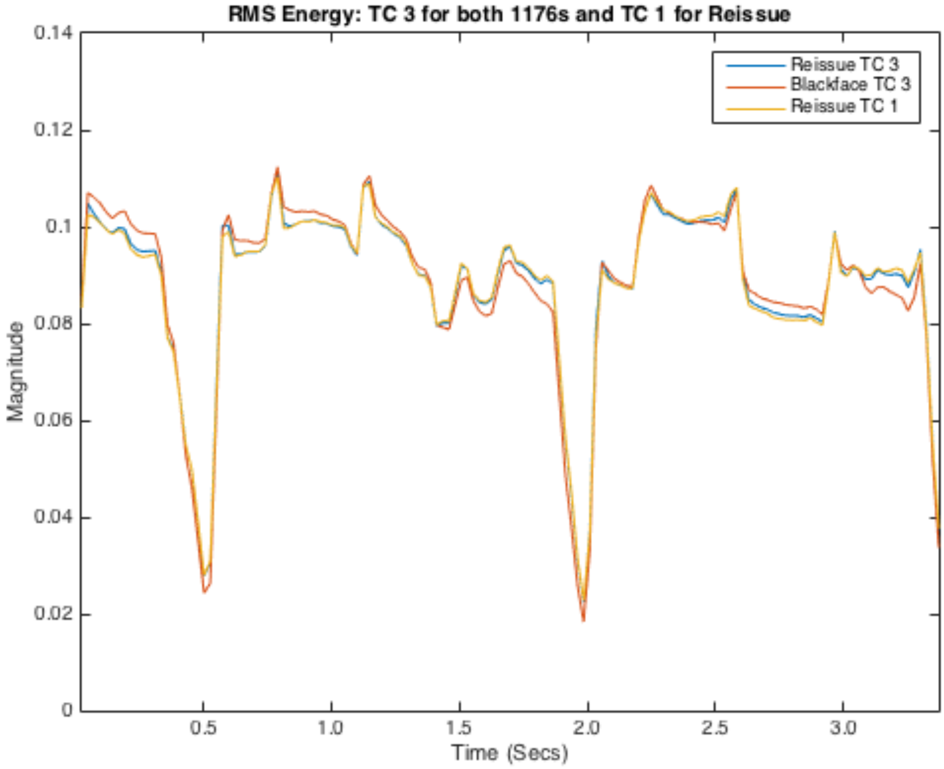

Figure 6-44: RMS energy both 1176s TC 3 and Reissue TC1 bass

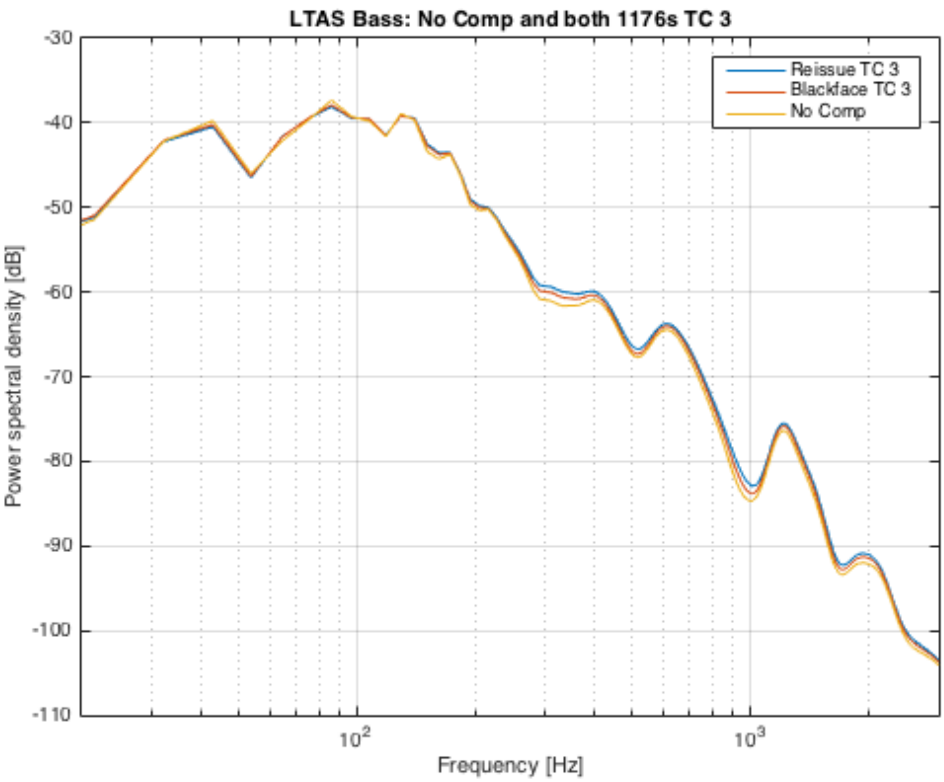

Figure 6-45: LTAS for both 1176s and uncompressed TC3 bass

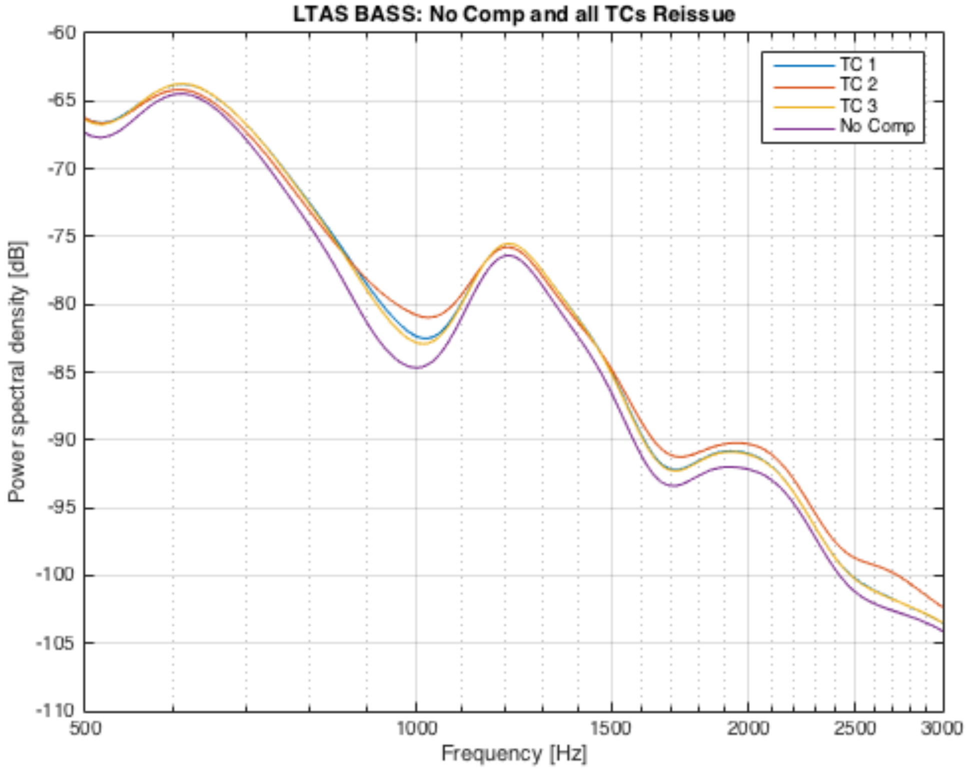

Figure 6-46: LTAS between 500-3000Hz both 1176s and uncompressed all TCs bass

6.10.4 Audio Features: 1176 On Bass

Features were extracted from the bass audio, and the data is presented in Table 6-11. It depicts an increase in values for all the features and a lower low-energy result. *Inharmonicity* was measured from a band between 1.6-3.2kHz extracted using a filter bank, and it was found that this measure correlated well with non-linearity heard in the audio in TC2. The feature inharmonicity is part of MIRToolbox and estimates the number of partials that are not multiples of the fundamental (Lartillot, 2015). The result for this feature for both compressors is similar. Thus, perceptually their non-linearity in this regard is comparable, and this is indeed the case when listening to the audio material.

The feature roughness was extracted from the audio, but it did not correlate with the audible non-linearity in a meaningful way. There was a general increase of roughness for all the compressed material, but TC2 did not have the highest value, demonstrating that roughness is not a suitable measure for the type of non-linearity produced by the 1176 when set for deliberate fast time constant based waveshaping.

Compressor	Roll Off	Spectral Centroid	Brightness	Low-Energy	Inharmonicity 1.6-3.2kHz
No Comp	420Hz	257Hz	0.014	0.41	0.463
Reissue TC1	441Hz	277Hz	0.015	0.34	0.467
Reissue TC 2	448Hz	290Hz	0.019	0.33	0.492
Reissue TC 3	440Hz	277Hz	0.015	0.32	0.464
Blackface TC1	427Hz	276Hz	0.015	0.31	0.464
Blackface TC 2	431Hz	282Hz	0.018	0.36	0.492
Blackface TC 3	425Hz	275Hz	0.015	0.30	0.464

Table 6-11: Audio features extracted from the bass material for all time constants, both 1176s and uncompressed material

Spectral flux in the 100-200Hz band was extracted from the audio, and the results can be seen in Table 6-12 that shows an increase in this feature, albeit small. This indicates that audio processed through the 1176 has an added thickness to its sonic signature.

Compressor	Spec Flux Mean	Spec Flux Std
No Comp	24.17	19.31
Reissue TC1	25.19	20.86
Reissue TC 2	25.27	19.89
Reissue TC 3	25.24	21.12
Blackface TC1	24.73	20.10
Blackface TC 2	24.54	19.22
Blackface TC 3	24.78	20.28

Table 6-12: Spectral flux for all TCs, both 1176s and uncompressed for bass material extracted from the 100-200Hz sub band

To conclude this sub-chapter, the data for different time constant settings shows there are some subtle variations but they are slight and certainly not as significant as the difference between the uncompressed and compressed material, demonstrating again the 1176 behaves consistently over a range of settings.

6.11 The 1176 On Room Mics

Room mic audio was tested with two popular settings extracted from the content analysis in Chapter 4, and the parameters are noted in Table 6-13. It can be seen the only difference between the settings is the ratio, both the attack and release remain fixed. The 4:1 ratio was selected as it represents a gentler form of room mic compression while the all-buttons mode was chosen as it was the most discussed room mic setting in Chapter 4. Furthermore, by using these two settings a direct comparison of the effect all-buttons mode has on audio, particularly the manner in which it adds distortion and modulation effects, can be made. The analysis in Chapter 4 revealed that producers were using the 1176 for its pumping effects and distorted sound quality when compressing drum room mics. Table 4-2 in Chapter 4 showed that distortion and modulation compression techniques accounted for 56% of the reasons why producers would compress room mics.

The original uncompressed room mic is presented in Figure 6-47, and amplitude statistics for all settings are noted in Table 6-14.

Time Constant	Attack	Release	Ratio
1	3	7	All-Buttons
2	3	7	4:1

Table 6-13: Compression settings used on the room mic material

Compressor	Dynamic Range (dB)
Reissue TC1	5.87
Reissue TC2	14.46
Blackface TC1	6.54
Blackface TC2	14.79
No Comp	19.57

Table 6-14: Amplitude statistics generated from the room mic material. The table includes all settings for both compressors

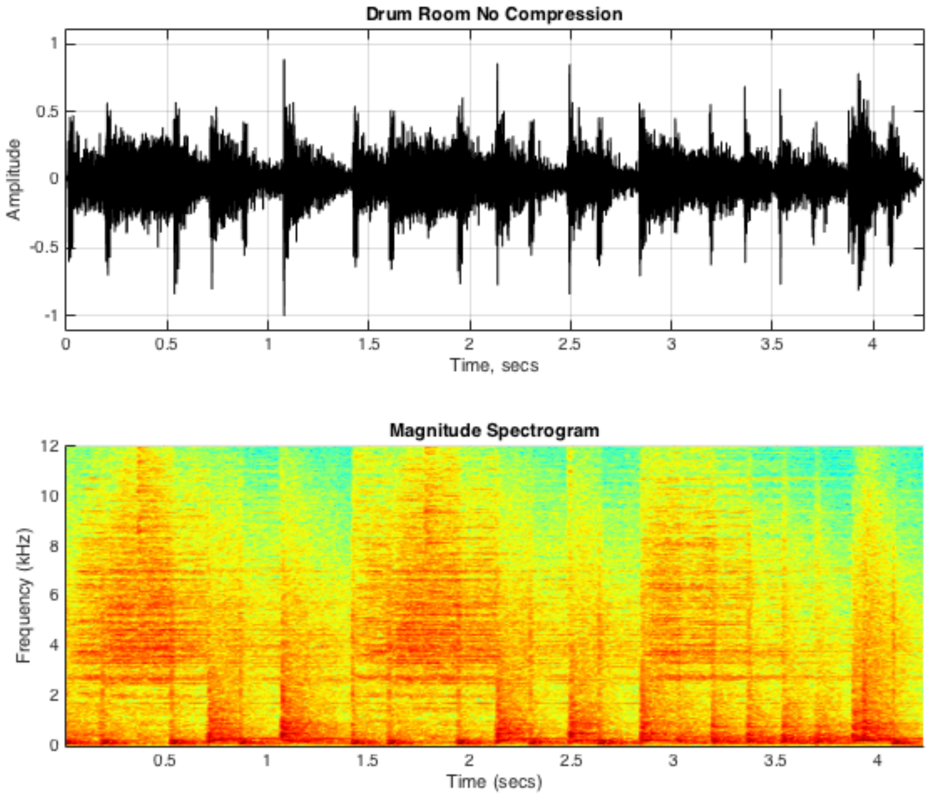

Figure 6-47: Time domain and spectrogram plots of the uncompressed room mic audio material

6.11.1 Room Mic Time Constant 1

Time domain and spectrogram plots of the compressed audio in Figures 6-48 and 6-49 show the effect of the all-buttons mode is very radical and has had a profound effect on the temporal and frequency-related properties of the audio. The all-buttons mode has heavily coloured the signal and has severely restricted the dynamic range. Close inspection of the uncompressed and compressed time domain plots shows the signal has been aggressively limited and only a few overshoots are left in the signal occurring at the start of the audio and at a snare hit during approximately 1100 milliseconds. The amplitude statistics in Table 6-14 depict the compressed signals to have a dynamic range approximately 13-14dB less than the uncompressed audio.

Comparing the time domain plots for the Blackface and the Reissue shows little difference. The Reissue is working a marginally harder on the signal, and this can be seen in the amplitude statistics where it is compressing the audio by approximately 0.5dB more than the Blackface. Due to the unstable performance of the VU meter in all-buttons mode, the compressors had to be set by ear for this test. Thus it is thought these small differences in dynamic range are at least in part due to slight variations in the amount of gain reduction applied by the two compressors, which were not discernible on the meters. The spectral plots for the two compressors look similar, but there is a general increase in spectral energy in the Reissue's spectrogram that is most noticeable up to just below 100Hz.

Listening to the uncompressed and compressed audio exposes a significant change in the sonic signature. The compressed audio has an aggressive character with the cymbals modulating and pumping to create a sense of movement and pulse that was not previously audible in the audio. The decay of snare drum hits is significantly extended, and the room ambience is raised in level to such an extent that it sounds almost as loud as the transient and steady state portions of the drum hits. This modulation effect was shown to be of importance to producers in Chapter 4's study.

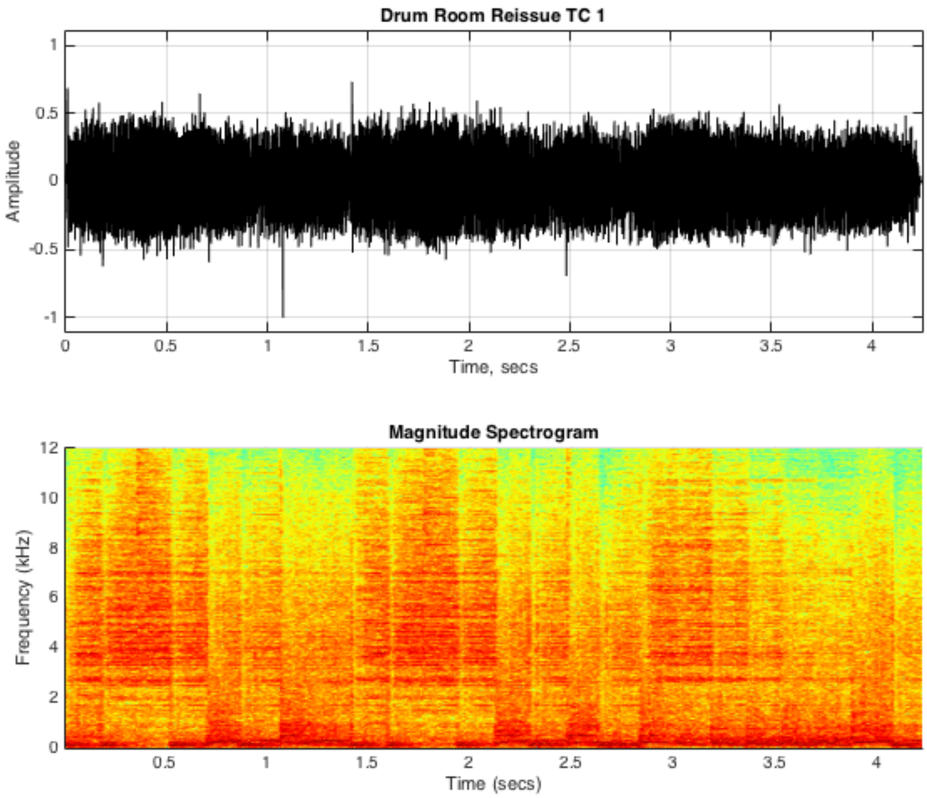

Figure 6-48: Time domain and spectrogram plots of the Reissue TC1 room mic

Distortion is clearly audible in the audio and can be heard throughout the extract. It is particularly distinct on snare hits, especially during the short tom-tom crescendo at the end of the extract. Sub-bands were extracted from the audio using a filterbank and critically listened to, to explore the audibility of this distortion. Distortion was noticeable in all bands but especially audible between 400-800Hz. During listening the Reissue had slightly more distortion than the Blackface. The non-linearity from both units has a gritty sound quality and may be responsible for the colouration effect that producers described as crunchy in Chapter 4.

RMS energy for the uncompressed and compressed audio is shown in Figure 6-50 where the aggressive nature of this setting is visible. The uncompressed material has significant variations in amplitude, and the transients of the audio are defined. The compressed audio, on the other hand, is shown to have its dynamic range aggressively restricted and the transients heavily attenuated to leave little variation in amplitude between them and the rest of the audio.

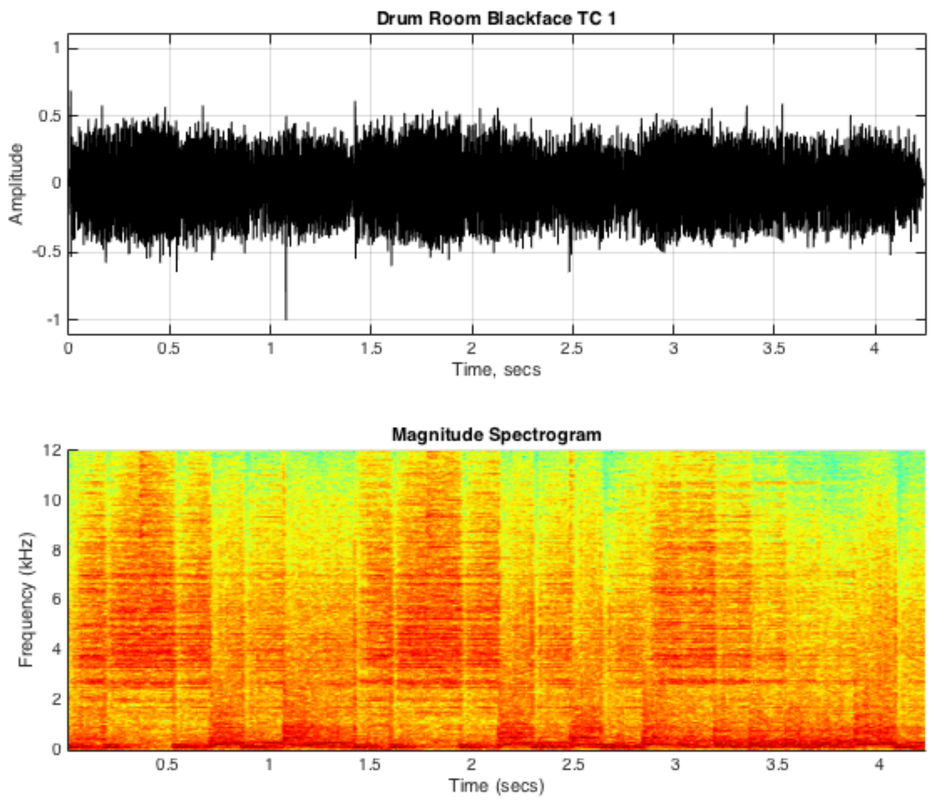

Figure 6-49: Time domain and spectrogram plots of the Blackface TC1 room mic

The LTAS for TC 1 is featured in Figure 6-51 and shows the compressed material has more energy above 2kHz and below 50Hz but some attenuation between 50-600Hz. The slight increase in the lower end of the Reissue that was seen in the spectrograms can also be seen in the LTAS plot extending up to just below 100Hz.

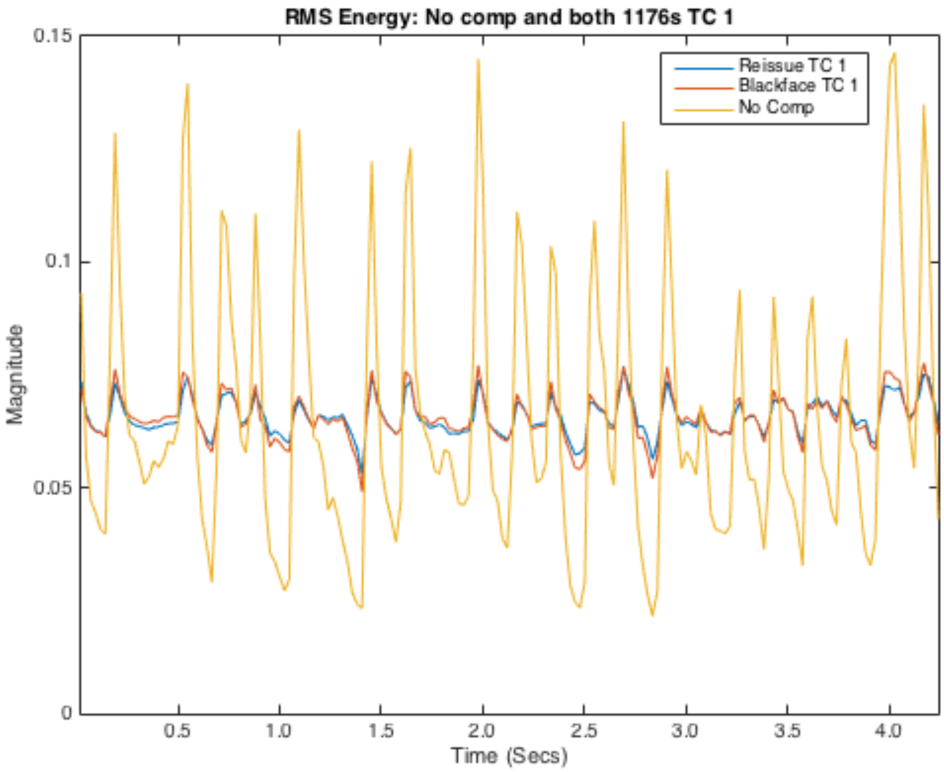

Figure 6-50: RMS energy both 1176s and uncompressed TC1 room mic

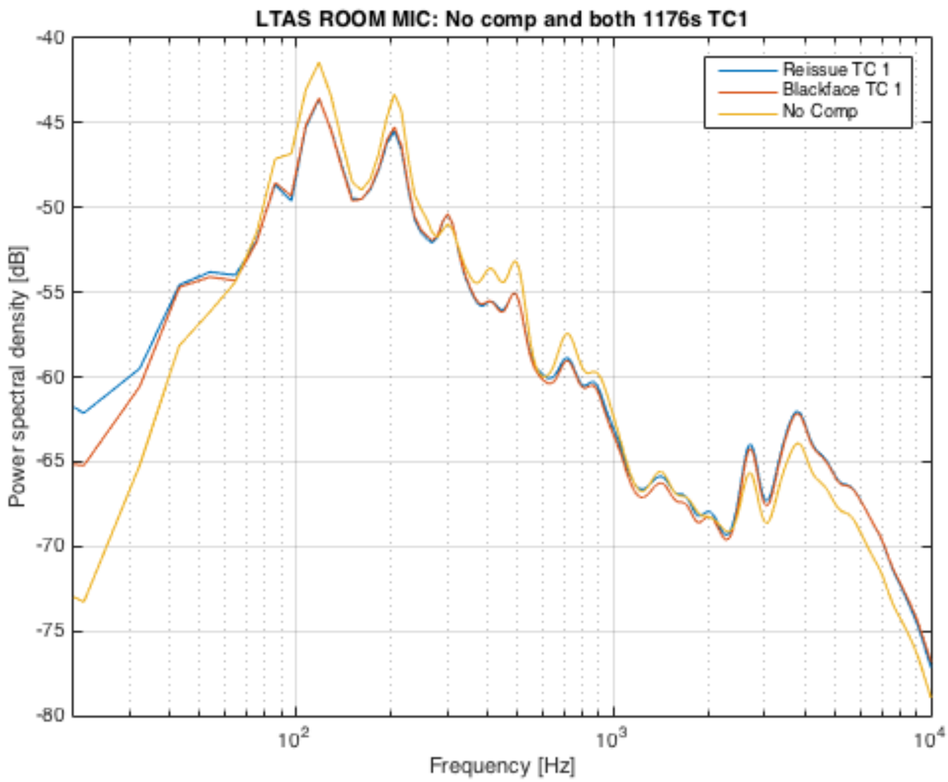

Figure 6-51: LTAS for both 1176s and uncompressed TC1 room mic

6.11.2 Room Mic Time Constant 2

The second time constant setting is shown in Figure 6-52 where it is apparent it does not colour the audio as much as TC 1. When compared with the uncompressed signal there is still a slight increase in energy in some areas of the frequency range, and this is due in large part to an increase in the decay portion of the drum hits. Listening to the audio reveals the uncompressed material has more transient attack on the kick and the snare and while the compressed audio preserves much of this articulation it has the effect of softening the transients due to the fast attack clamping down on the drum hits. The release time is helping raise the level of the subtler drum hits and room ambience but to a lesser extent than with the all-buttons mode. Aside from changes to the transients and increase in drum decay and room ambience levels, this setting adds the thickness that was audible in the other compressed material, although it is not as noticeable as when using all-buttons mode.

This setting does not have the crunchy timbre that was audible with the previous time constant. Audio was extracted using a filter bank between 400-800Hz for critical listening, and unlike TC1, this setting does not have as much of an increase in drum hit decay and room ambience or its distorted sound quality. When comparing the two units perceptually, the author struggled to hear any significant difference between the two with this setting. The plots and features extracted from the audio help verify this assertion. The difference in bottom end between the Reissue and the Blackface that was audible in all-buttons mode is not discernible with this setting.

The RMS energy plot for this time constant can be seen in Figure 6-53 and shows both units have tracked the original envelope accurately and preserved much of its original shape. The most noticeable differences are an increase in level to low-level detail between hits, attenuation of transients and small amounts of microdynamic shaping.

The LTAS for this setting is featured in Figure 6-54 where it can be seen the two compressors group together closely. The only differences between the compressed and uncompressed material are small changes to the bottom and top end. A LTAS plot in Figure 6-55 includes both time constant settings for the Reissue and is included to illustrate the differences between the two settings. The LTAS shows that TC1 has additional low end below 80Hz and above 1kHz.

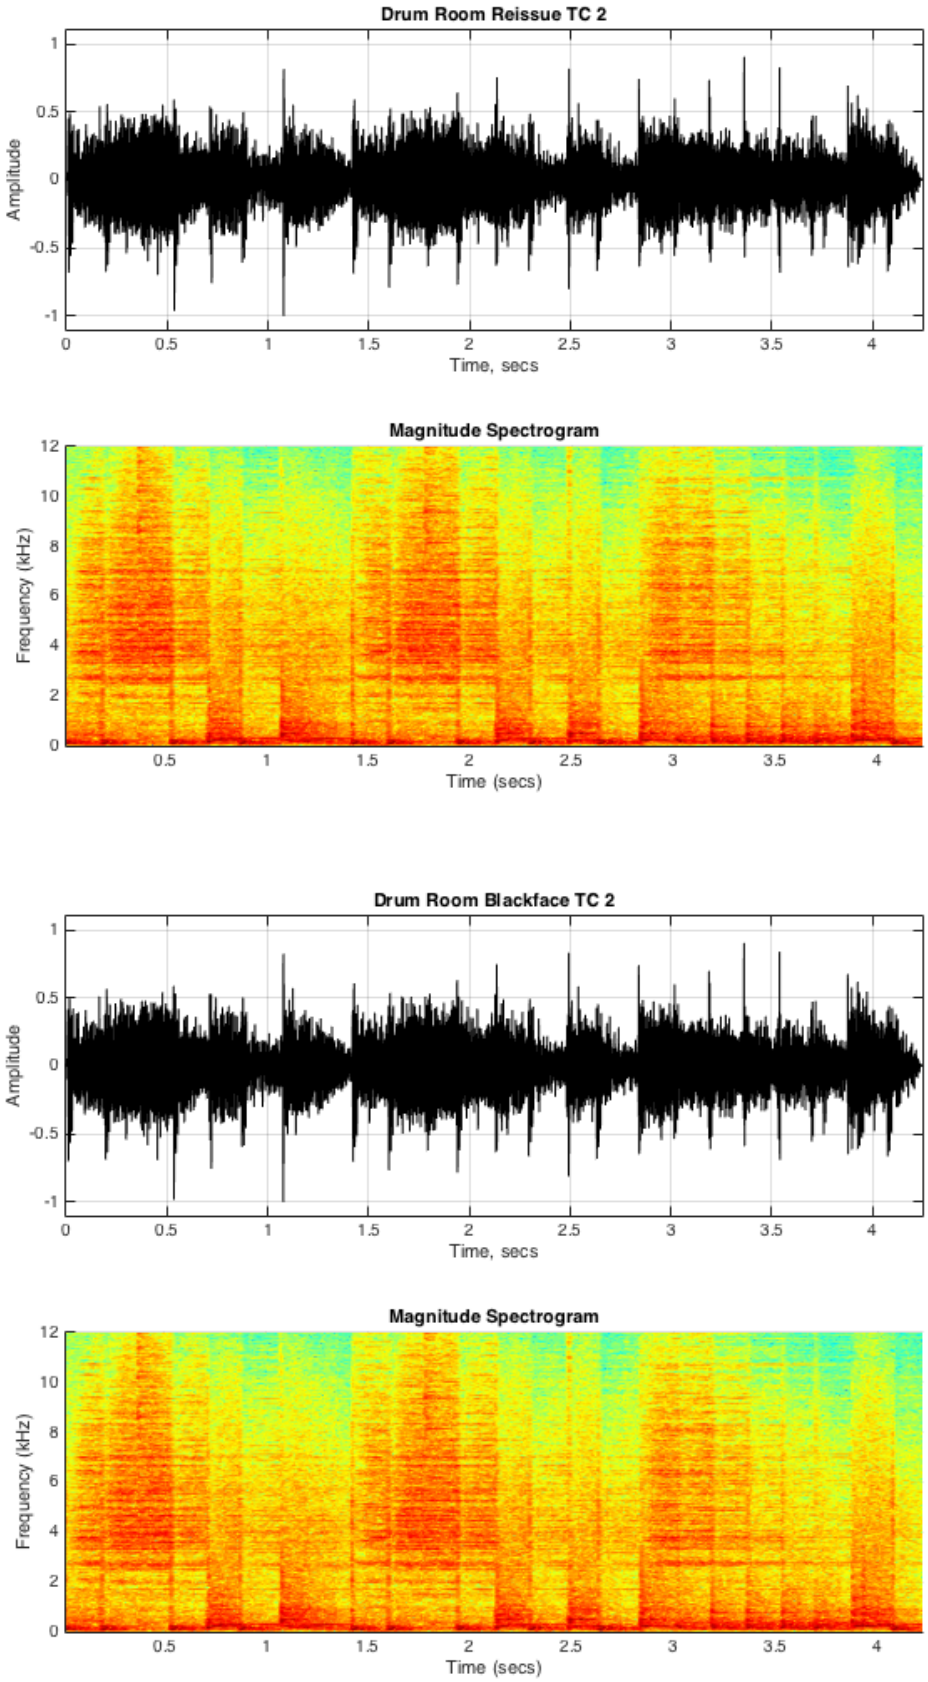

Figure 6-52: Time domain and spectrogram plots of the Reissue (top) and Blackface (bottom) room mic audio TC2

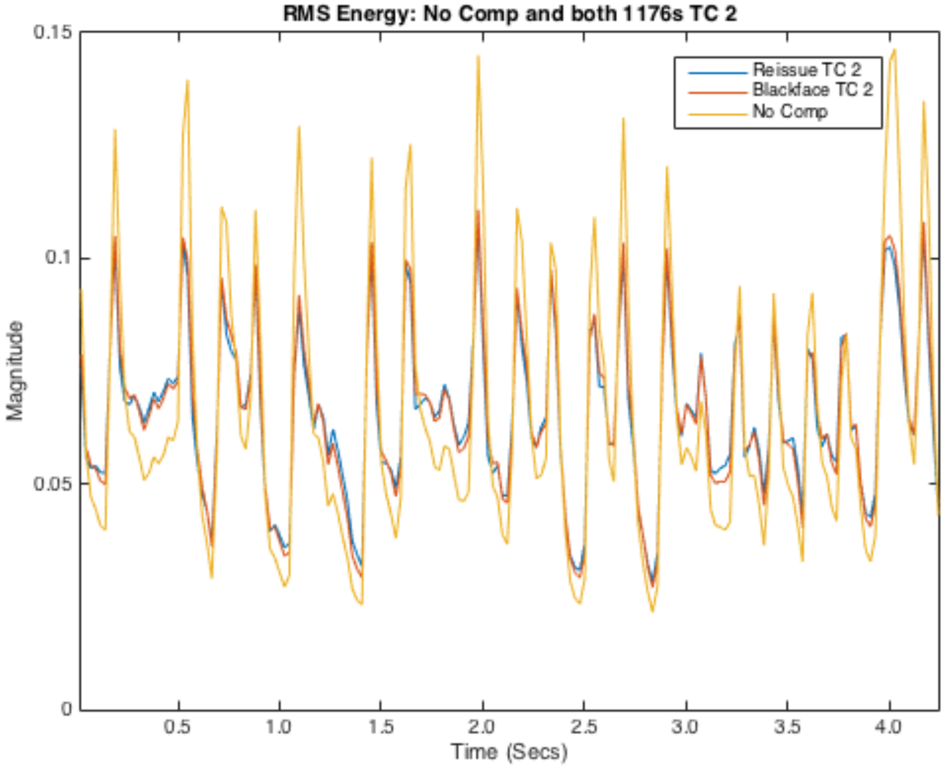

Figure 6-53: RMS energy both 1176s and uncompressed TC2 room mic

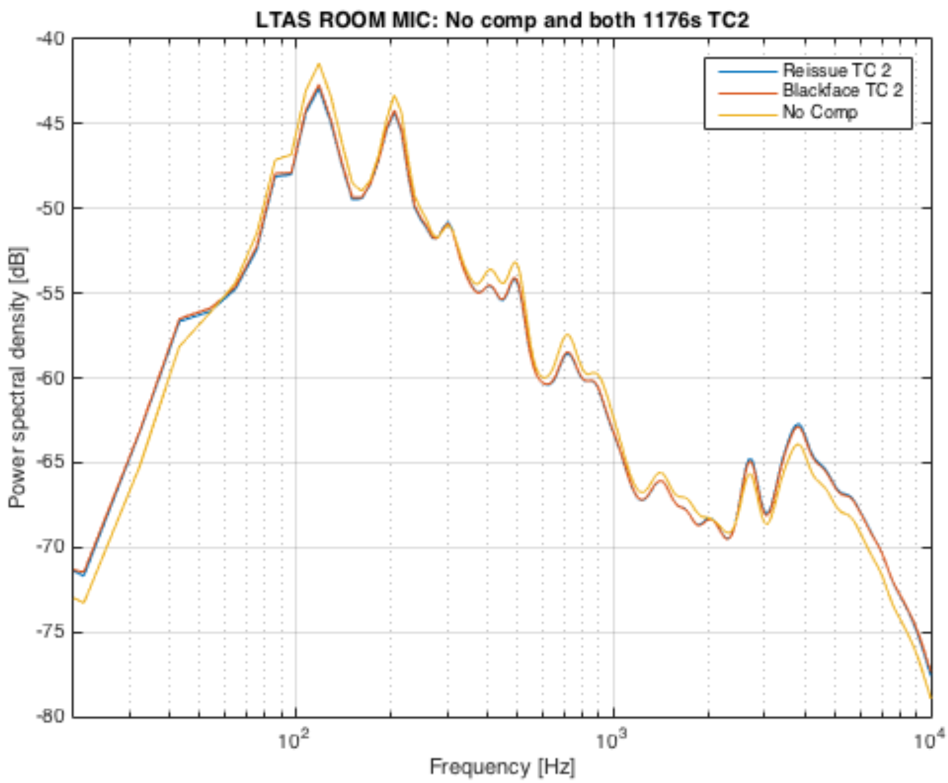

Figure 6-54: LTAS for both 1176s and uncompressed TC2 room mic

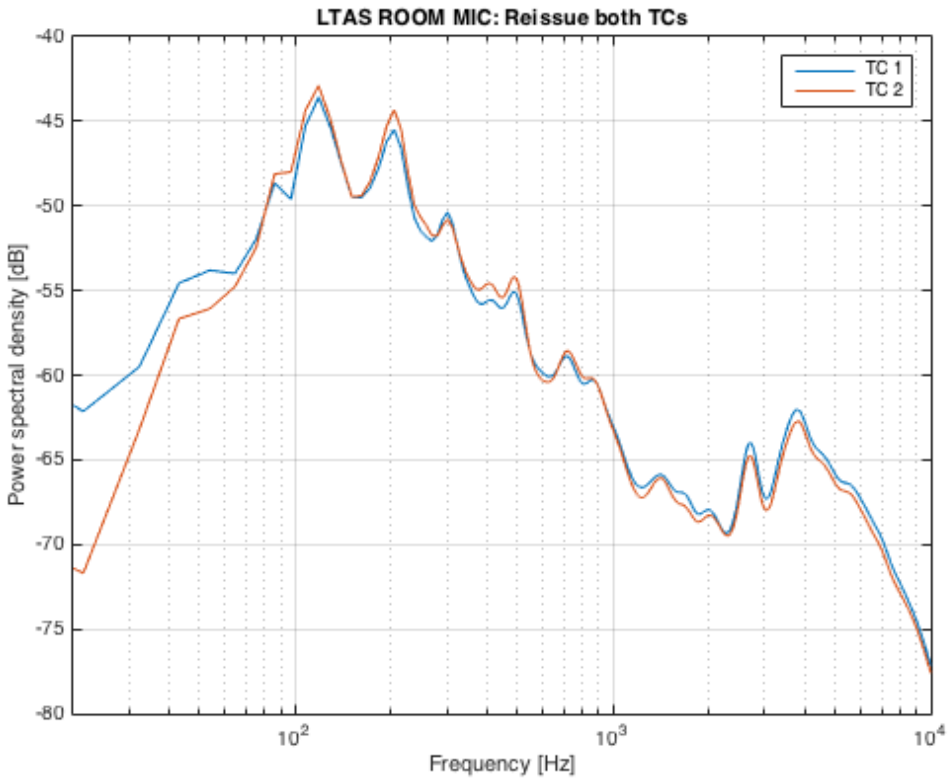

Figure 6-55: LTAS for both time constants Reissue room mic

6.11.3 Audio Features: 1176 Room Mic

Audio features were extracted from the room mic material and the results are presented in Table 6-15 where it can be seen there is an increase in all values of the compressed features bar low-energy, which as one would expect is lower.

Compressor	Roll Off	Spectral Centroid	Brightness	Roughness	Low-Energy
No Comp	8428Hz	4332Hz	0.626	32.80	0.63
Reissue TC1	8659Hz	4668Hz	0.694	34.79	0.56
Reissue TC 2	8697Hz	4591Hz	0.673	32.71	0.52
Blackface TC 1	8980Hz	4821Hz	0.698	34.50	0. 52
Blackface TC 2	8975Hz	4699Hz	0.673	32.62	0.52

Table 6-15: Audio features extracted from the room mic material for both time constants, both 1176s and uncompressed material

TC 1 (the all-buttons mode) has resulted in marginally higher values for the majority of the features, but the differences are small. The 1176 imparts some timbral change, but the variation in timbre between settings is not significant. There is a general brightening of the audio as evidenced by the increase in spectral centroid, roll off and brightness figures. The low-energy feature for the Blackface has the same result for both settings, but the reason for this is unclear.

To further investigate the audible non-linearity in the compressed material the feature roughness was extracted and is included in Table 6-15. All-buttons mode has the highest result for this feature thus it appears to correlate well with all-buttons non-linearity. The Reissue has higher results for this feature, which makes sense when considering the THD and IMD results that were discussed earlier in this chapter.

Spectral flux in the 50-100Hz band was measured, and the results are shown in Table 6-16 that illustrates a slight increase in this feature when in all-buttons mode (TC 1) but little change (a small decrease in fact) in TC 2.

Compressor	Spec Flux Mean	Spec Flux Std
No Comp	32.80	18.92
Reissue TC1	34.79	19.05
Reissue TC 2	32.71	20.31
Blackface TC1	34.50	19.49
Blackface TC 2	32.62	20.07

Table 6-16: Spectral flux for all TCs, both 1176s and uncompressed for room mic material extracted from the 50-100Hz sub band

6.12 Conclusions on Music Testing

Part of the 1176's popularity when compressing vocals is due to the thick texture it imparts to program material and also the transparent, unobtrusive manner in which it tracks and largely preserves the original envelope. This thick texture correlates with an increase in spectral flux in lower frequency bands of the audio. When set for hard limiting the 1176 still has its identifiable sonic signature and features the desirable attributes of the less aggressive compression settings. This colouration is part of the reason why this compressor was a popular choice of professionals in the study in Chapter 4. They had both a preference for FET style compressors and a proclivity for colouration when compressing vocals in the mix. Another reason for the 1176's popularity as a vocal compressor is its consistent behaviour over a range of settings. Arguably this is why the 1176 is described as a fool-proof compressor in a popular music production forum used by professionals (Gearslutz, 2013). The difference between the two units when using vocals is small, and it is thought any differences are mainly due to variation in the ballistics of the VU meters and the attack time.

When compressing bass, part of the 1176's popularity is due to the thick texture it imparts to program material. This increase in texture correlates with an increase in spectral flux in lower frequency bands, and the 1176 is consistent across all the time constants tested in this regard. The compressor tracks the dynamics of bass material in an almost transparent manner and largely preserves the original envelope shape of the audio. The effect of a longer attack is negligible both objectively in the data and also subjectively under listening. When set with a high ratio and its fastest attack and release settings, the 1176 introduces a significant amount of non-linearity to the audio. However, yet again it preserves the general envelope shape of the program material. This effect is perhaps why engineers discussed employing this effect on bass tracks in Chapter 4. The 1176 creates distortion that does not round off the transients as would typically occur when distorting low frequency sources.

When compressing drum room mics and using the all-buttons mode, the 1176 increases energy in the midrange related to the decay of the drums and also the room ambience. This appears to be at least in part responsible for the gritty, crunchy colour that was discussed by producers in Chapter 4. The dynamic range of drum audio compressed using the all-buttons mode is severely reduced, and any transients in the program material are radically attenuated. Listening to the compressed audio suggests that the all-buttons mode in combination with a quick release introduces significant non-linearity onto program material. This is in accordance with the type of non-linearity that was observed under the IMD tests and translates objectively as higher values of the roughness feature. When not using the all-buttons mode, the 1176 is largely transparent. Even when compressing room mics with considerable gain reduction it tracks the envelope in a manner that preserves much of the original shape while raising the low-level detail and decay between drum hits. The difference between the two 1176 units tested is small. There is a marginally higher result for the spectral centroid and brightness features extracted from material compressed with the Blackface, and this may translate as a brighter sonic signature. The Reissue had a little more bottom end when using all-buttons mode, but this was not present with the 4:1 ratio. When compared with the uncompressed audio there is little change to the sub band spectral flux for TC2. There is a slightly higher value for the all-buttons mode (TC1) and this may be another reason why this setting is commonly used.

6.13 Attack Time and Snare Testing Introduction

So far, the focus of testing has been on program material that was found to be popular when using the 1176. From the study in Chapter 4, it was established that while producers occasionally use the 1176 to compress drum spot mics, it is not their first choice, particularly when using the compressor to enhance the envelope of kick and snare drums. Cursory testing suggested that one reason for this may be due to the limited attack times available on the 1176. Thus, a series of measurements was made to observe the actual range of the 1176's attack control. Analysis consisting of a test tone (that will be described later) was used to measure the range of the attack and the results are discussed in this sub-chapter.

In addition to test tones, a portion of complex program material consisting of a spot mic snare recording was processed through a standard 1176, an AE edition and two VCA compressors, namely the SSL G Series channel compressor and the dbx165A compressor. The motivation behind this test was to investigate the range of the 1176's attack on complex program material and to explore the transient shaping options offered by VCA compressors. VCA compressors were selected for comparison as they were found in Chapter 4 to be the most popular type of compressor when implementing this style of production technique.

The AE edition was added to the test to get a fuller picture of the 1176 units available to the music producer. The AE is a limited-edition model designed with a special 10-millisecond attack mode. According to Universal Audio (2008), this attack mode opens the 1176 up to other forms of compression, presumably transient shaping styles that allow the user to accentuate the transient portion of snare and kick drums.

6.14 The 1176's Attack Time

The test signal used for this examination was a 100ms burst tone with a 20ms transient at the front end. The tone was at 50Hz and can be seen in Figure 6-56.

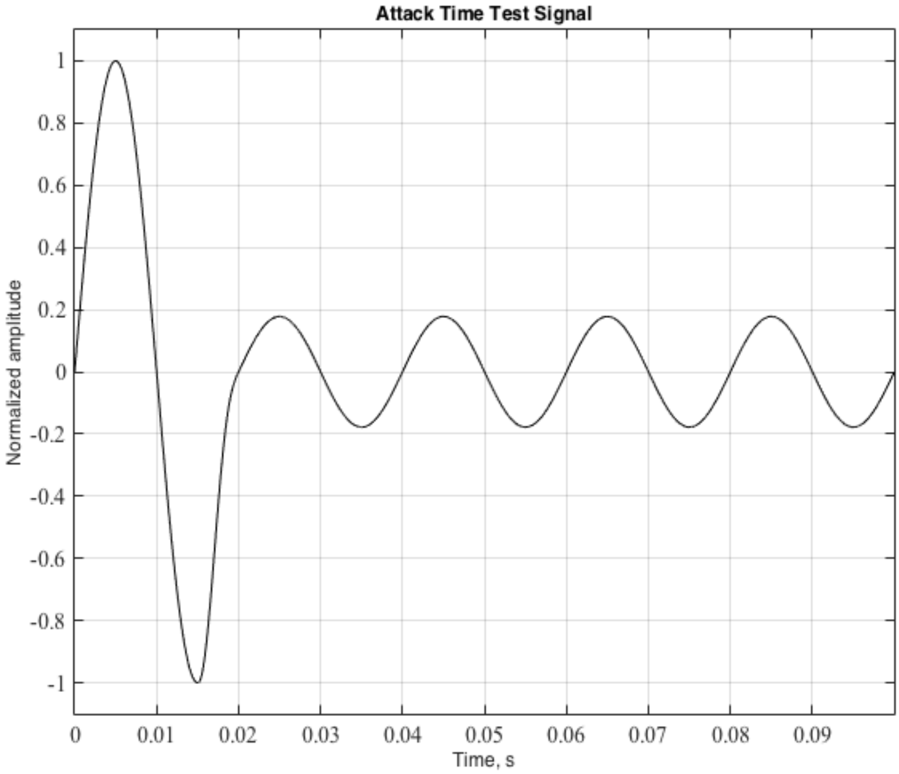

Figure 6-56: Test tone signal used to investigate the range of the 1176's attack control

This tone was compressed with the 1176 release set to its fastest position, the ratio set to 4:1 and input set for approximately 6dB of gain reduction. A series of measurements were made at all the numerical positions as the attack control was turned counter clockwise. The compressed signals were recorded into a DAW for analysis, and the results from all the attack positions are shown in Figure 6-57. Table 6-17 shows the results of measuring the linear attack time of the compressed material using the mirattacktime function from MIRToolbox.

The first piece of information to glean from this table is the small range present. The differential between position 1 to position 7 is five milliseconds. There is no change between positions 4 and 3 and one millisecond of difference between positions 2 and 1. This data confirms the author's initial suspicion that the attack time offered a limited range of control.

Attack Position	Linear-Attack Time (Seconds)
7	0.025
6	0.024
5	0.023
4	0.022
3	0.022
2	0.021
1	0.020

Table 6-17: Linear attack time feature extracted from the audio of each attack time measurement

The top plot in Figure 6-57 shows the difference in attack time between positions 7-4 and the bottom plot shows the difference between positions 4-1. The results have been split into two plots to make it easier for the reader to see the differences. Looking at Figure 6-57 shows the fastest position (attack 7) has produced some wave shaping in the first 40ms. The alteration to the shape of the signal from 20ms onwards remains consistent over the full range of attack times and the only difference between the times is a small variation in the level of the portion after the transient.

The most significant difference between all positions occurs during approximately the first 5ms of the signal as the cycle swings towards positive. This can be thought of as micro-level transient shaping with the front of the short transient being reshaped to have a more pronounced and aggressive front end around the onset. The rate of change is more significant between positions 7 to 5 than compared with positions 4 to 1. In Chapter 4 attack times between 1 and 4 were found to be commonly used by producers and this demonstrates they are restricting themselves to work within a limited range that has a small effect on audio material. It is also noticeable that a combination of the fastest attack and release raises the level of the portion after the transient more than the other positions. The small differences in timbre are just about audible between positions 7 to 5 and practically inaudible between positions 4 to 1. The perceptual effects of slowing the attack time on the test tone between positions 7 to 5 sound like an increase in brightness and some additional thickness.

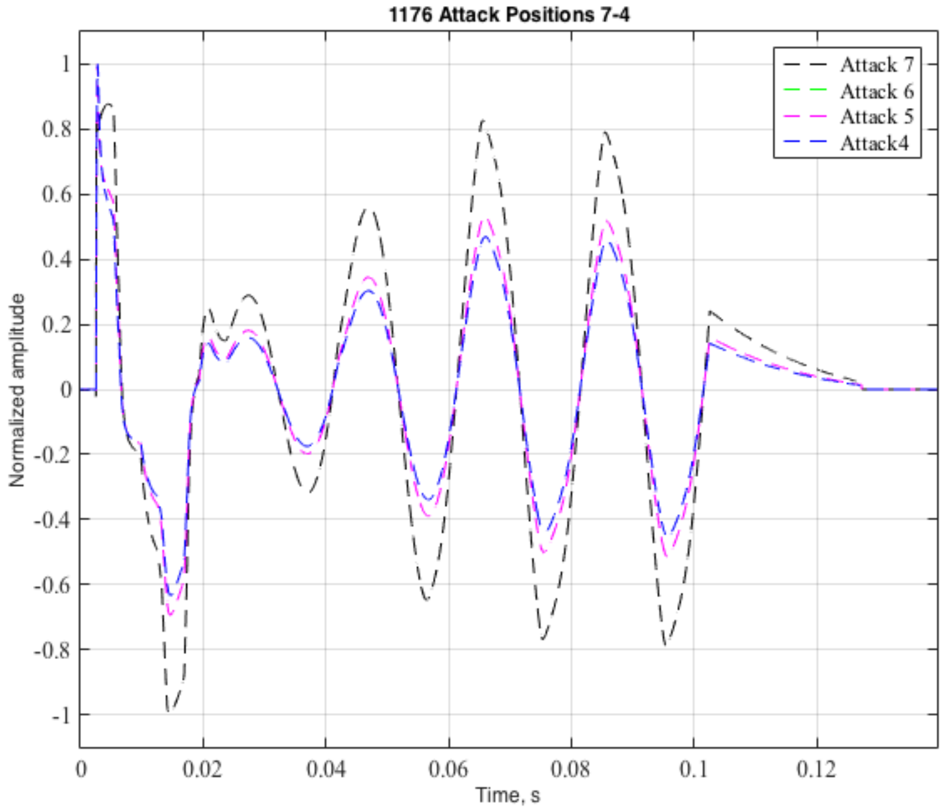

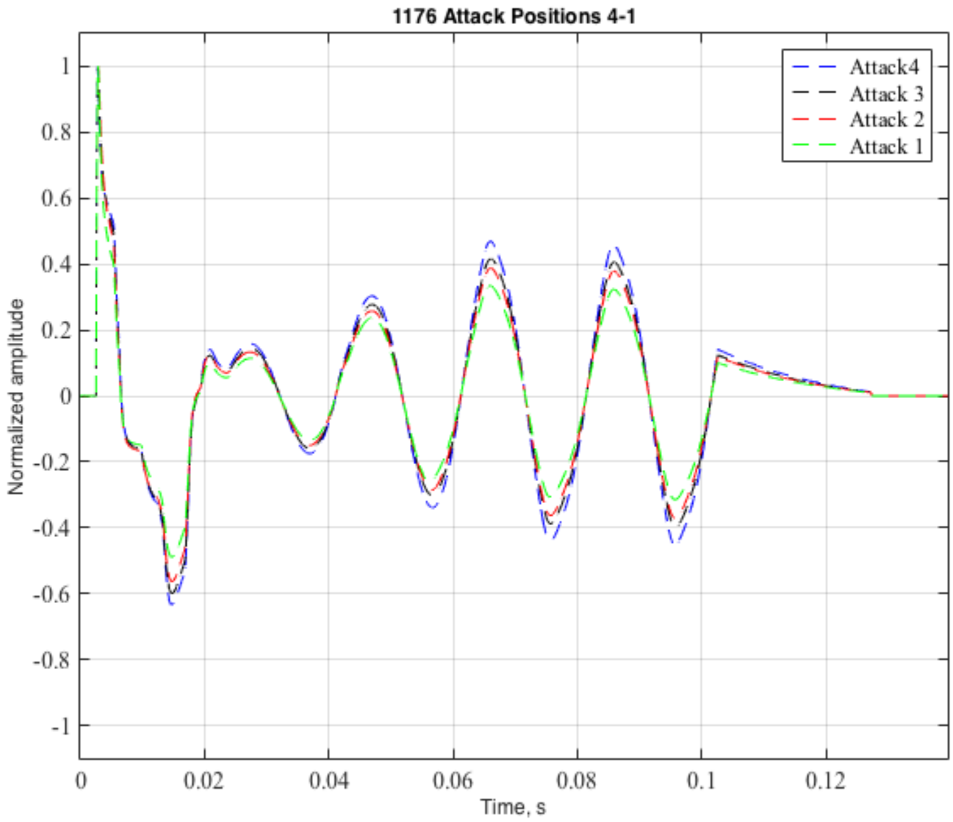

Figure 6-57: The effect of the attack time on the test signal. The top plot shows positions 4-7 and the bottom plot positions 1-4

6.15 Comparison of the 1176 and VCA Compressors

Snare drum material was compressed through two 1176s and two VCA based compressors. The tests were carried out on the 1176 making use of the fastest and slowest attack times to investigate settings at opposite ends of the attack range. Table 4-4 in Chapter 4 showed that only release positions between five and seven were used on drum sources. Thus, measurements were made at release positions 7,6 and 5 in conjunction with the two attack times.

For VCA compression, the snare was processed through the channel compressor on an SSL G series mixing console making use of both the fastest and slowest attack modes in conjunction with its fastest release of 100ms. Snare material was processed with a dbx165A making use of its fastest and slowest attack times. The dbx165A had its release set to the fastest position, which is quoted as 400 dB per millisecond.

The uncompressed snare can be seen in Figure 6-58. The SSL and dbx165A slow attack results are shown in Figure 6-59 and the standard and AE 1176 slow attack results in Figure 6-60. The fast attack settings for the compressors are depicted in Figures 6-61 and 6-62 and the compressors are presented in the same order as the slow attack figures.

Listening to the audio highlights the dbx165A has the most pronounced effect. This can be clearly heard in the audio extracts in the data folder. The transient has been reshaped to such a degree that it now has a large spike at the front end of the audio. The first 20 milliseconds of the transient is allowed through the compressor before it rapidly grabs onto the decay portion. The SSL does a similar job when set to its slowest attack position and allows a little more of the transient portion through, but it does not react as quickly as the dbx165A after this overshoot, and consequently the transient is not as radically accentuated.

The standard 1176 at its slowest attack and release at 7 is not capable of reshaping the envelope to the same extent as the two VCA compressors. The 1176 softens the transient of the snare resulting in a less articulate timbre. Additionally, the 1176 has increased the level of the decay relative to the transient thus perceptually reducing the sense of attack. This production technique can be used for transient softening effects, but it does not lend itself to the creation of accentuated transients. Slowing the attack to position 6 helps address this problem, and the snare drum begins to sound more articulate but not like the VCA compressors.

Setting the release to position 5 reshapes the snare envelope, so it sounds somewhat punchier but at the expense of a slow release speed that audibly reshapes the decay portion with a ramp up in amplitude. Table 5 showed a release of 6 was the least typical of the settings used by producers on drum sources and the audible artefacts are presumably why. Figure 6-63 shows the standard 1176 set with its slowest attack speed and release at positions 6 and 5.

The "slo" attack mode of the AE edition allows for more flexible transient shaping. The AE attack resembles that of the dbx165A and SSL at their slowest positions. However, the program dependent behaviour of the 1176 has resulted in the release stage working aggressively on the signal and has raised the spill to such a degree that it is more noticeable than in the dbx165A or SSL files. This type of behaviour is useful for evening out low level detail or to impart pumping but it is not desirable for envelope shaping compression styles. Adjusting the release speed to a slower position helps minimize this issue, but positions beyond 6 introduce the pumping discussed previously. To illustrate this point, the RMS envelope of the uncompressed audio, SSL and 1176 versions are shown in Figure 6-64 where it can be seen the spill between drum hits has been significantly raised in level by the 1176 while the SSL has not affected it in the same manner.

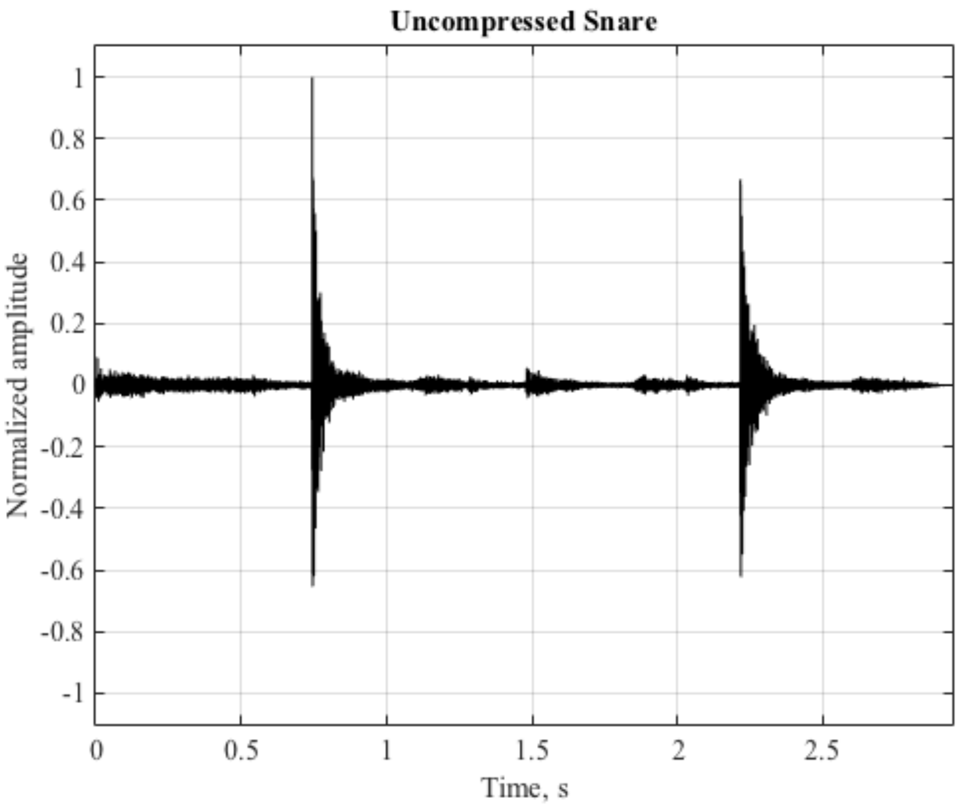

Figure 6-58: Uncompressed snare hits

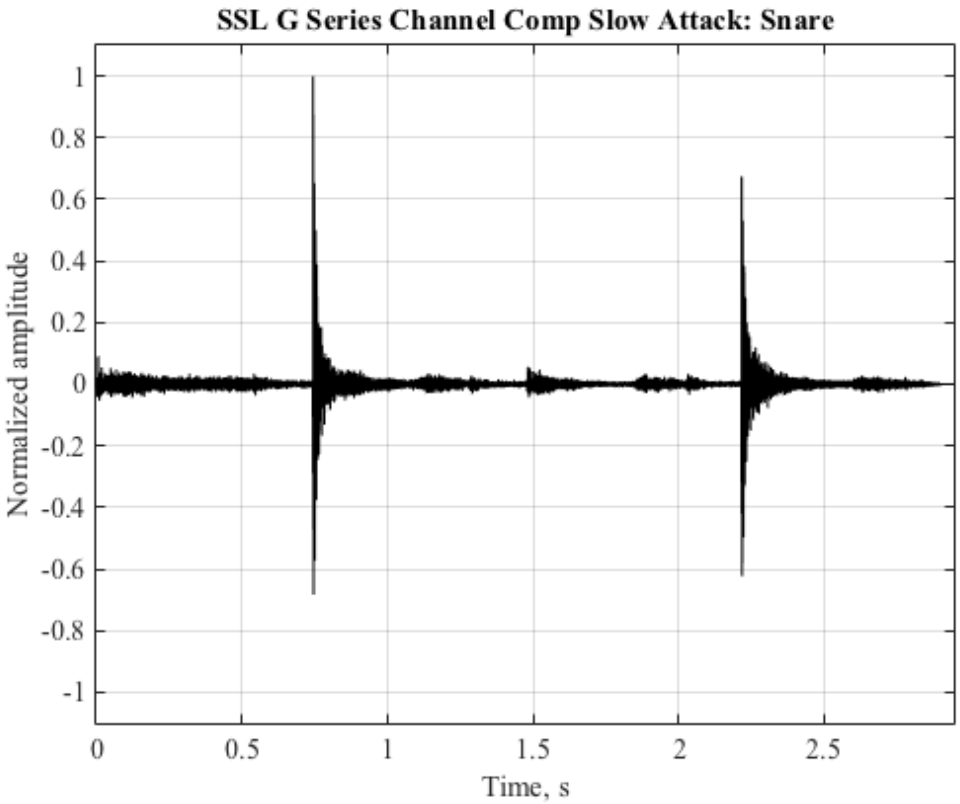

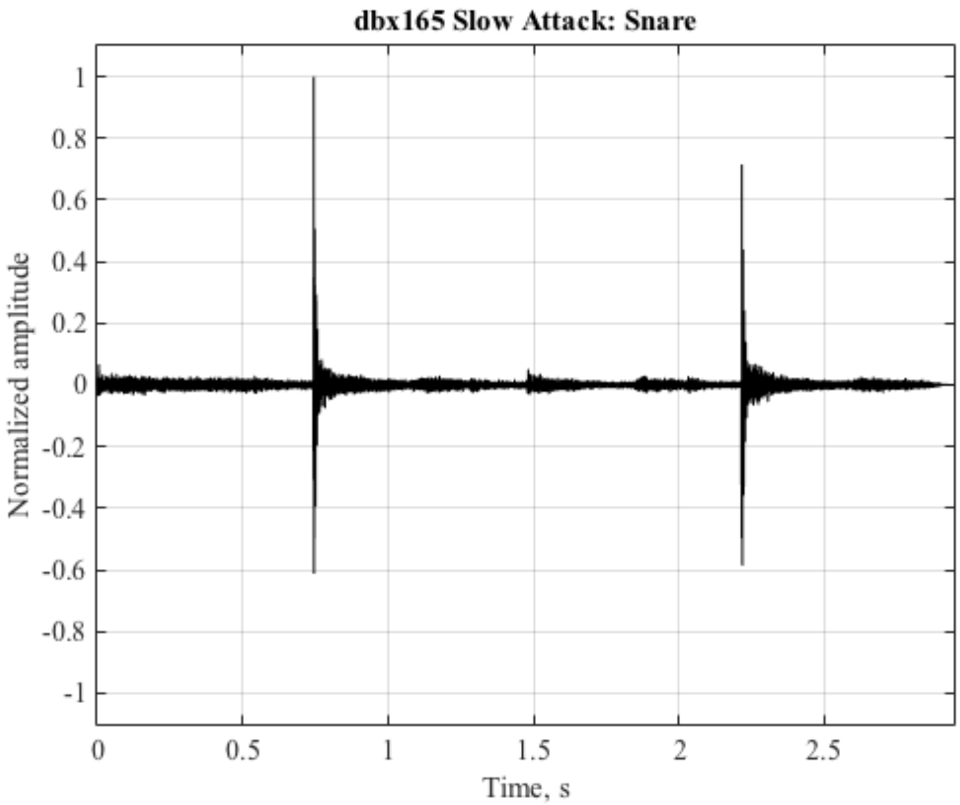

Figure 6-59: Compressed snare hits using slow attack times. SSL channel compressor on top and dbx165A on bottom

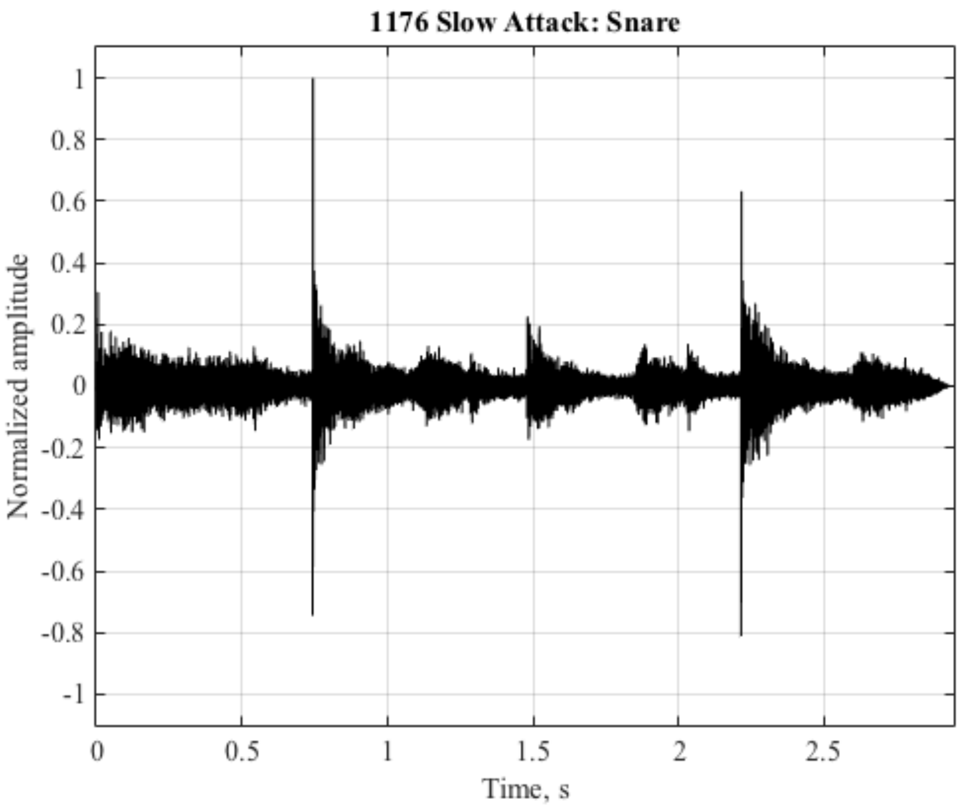

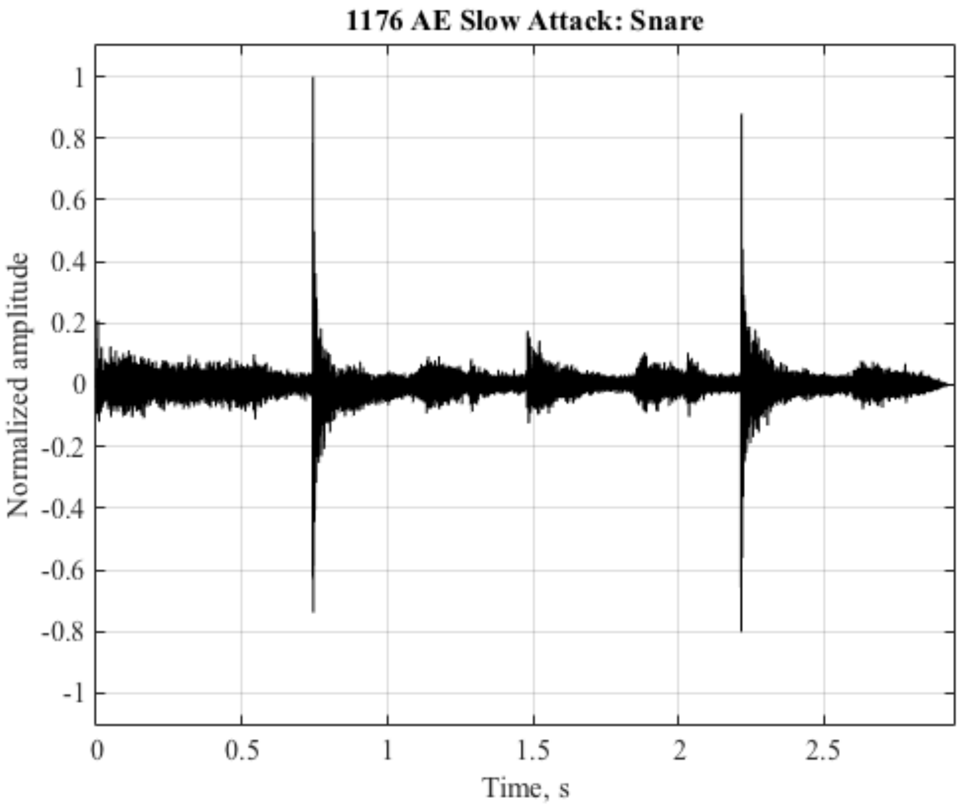

Figure 6-60: Compressed snare hits using slow attack times. Standard 1176 on top and 1176 AE on bottom

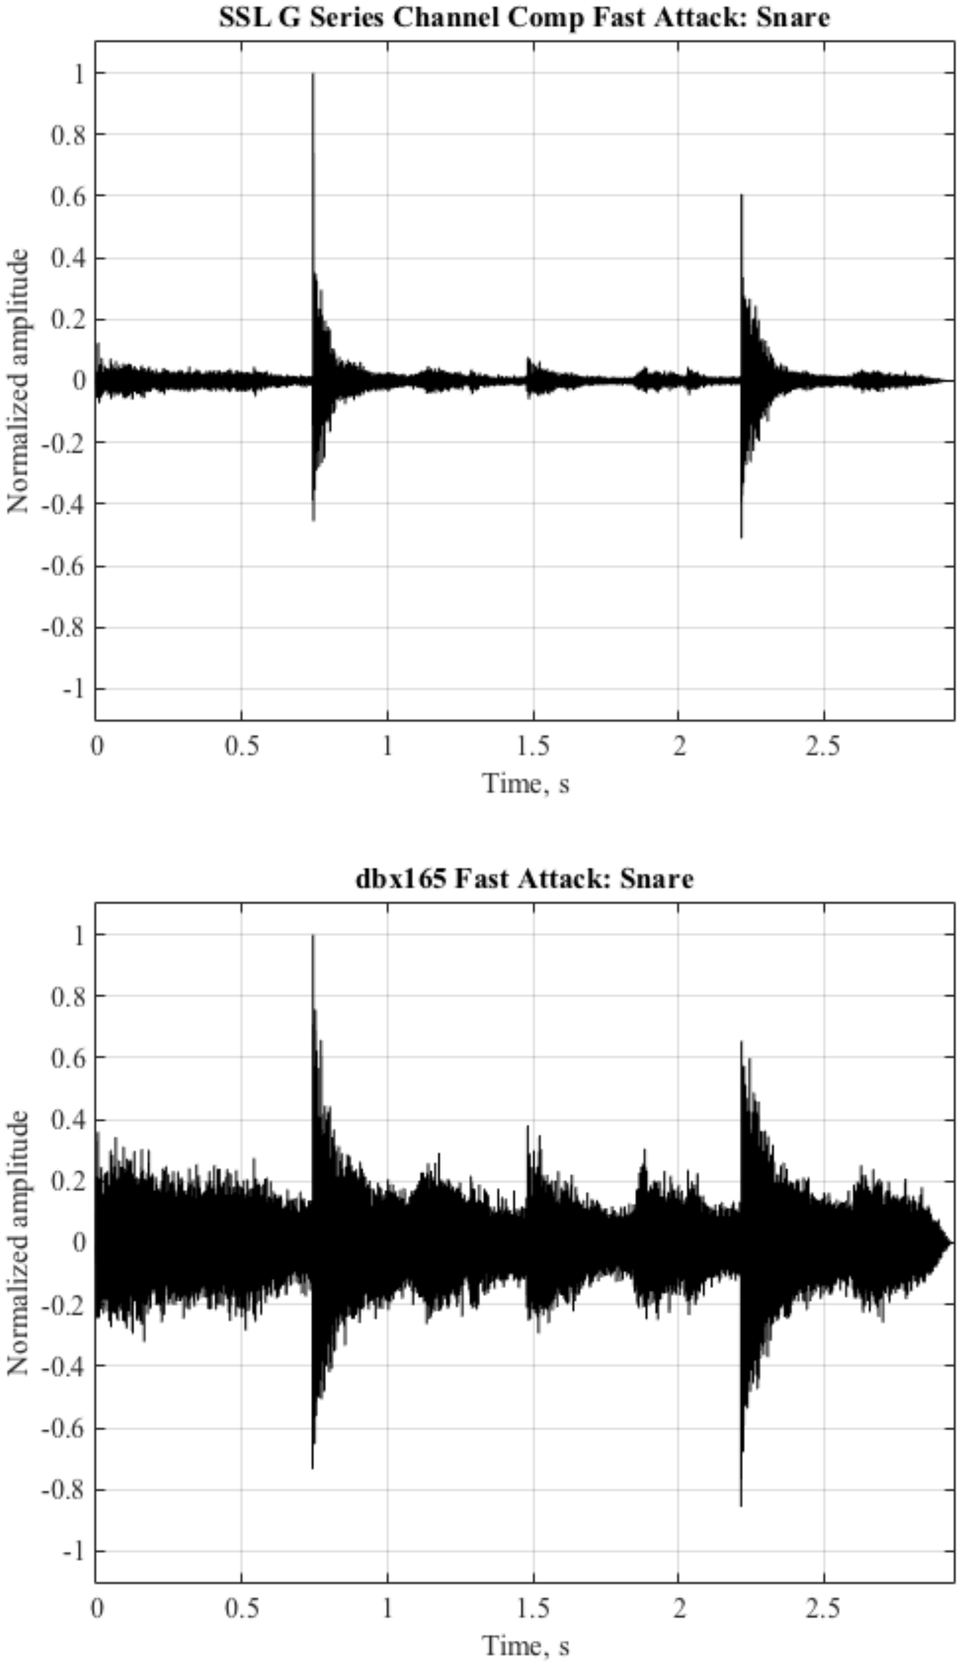

Figure 6-61: Compressed snare hits using fast attack times. SSL channel compressor on top and dbx165A on bottom

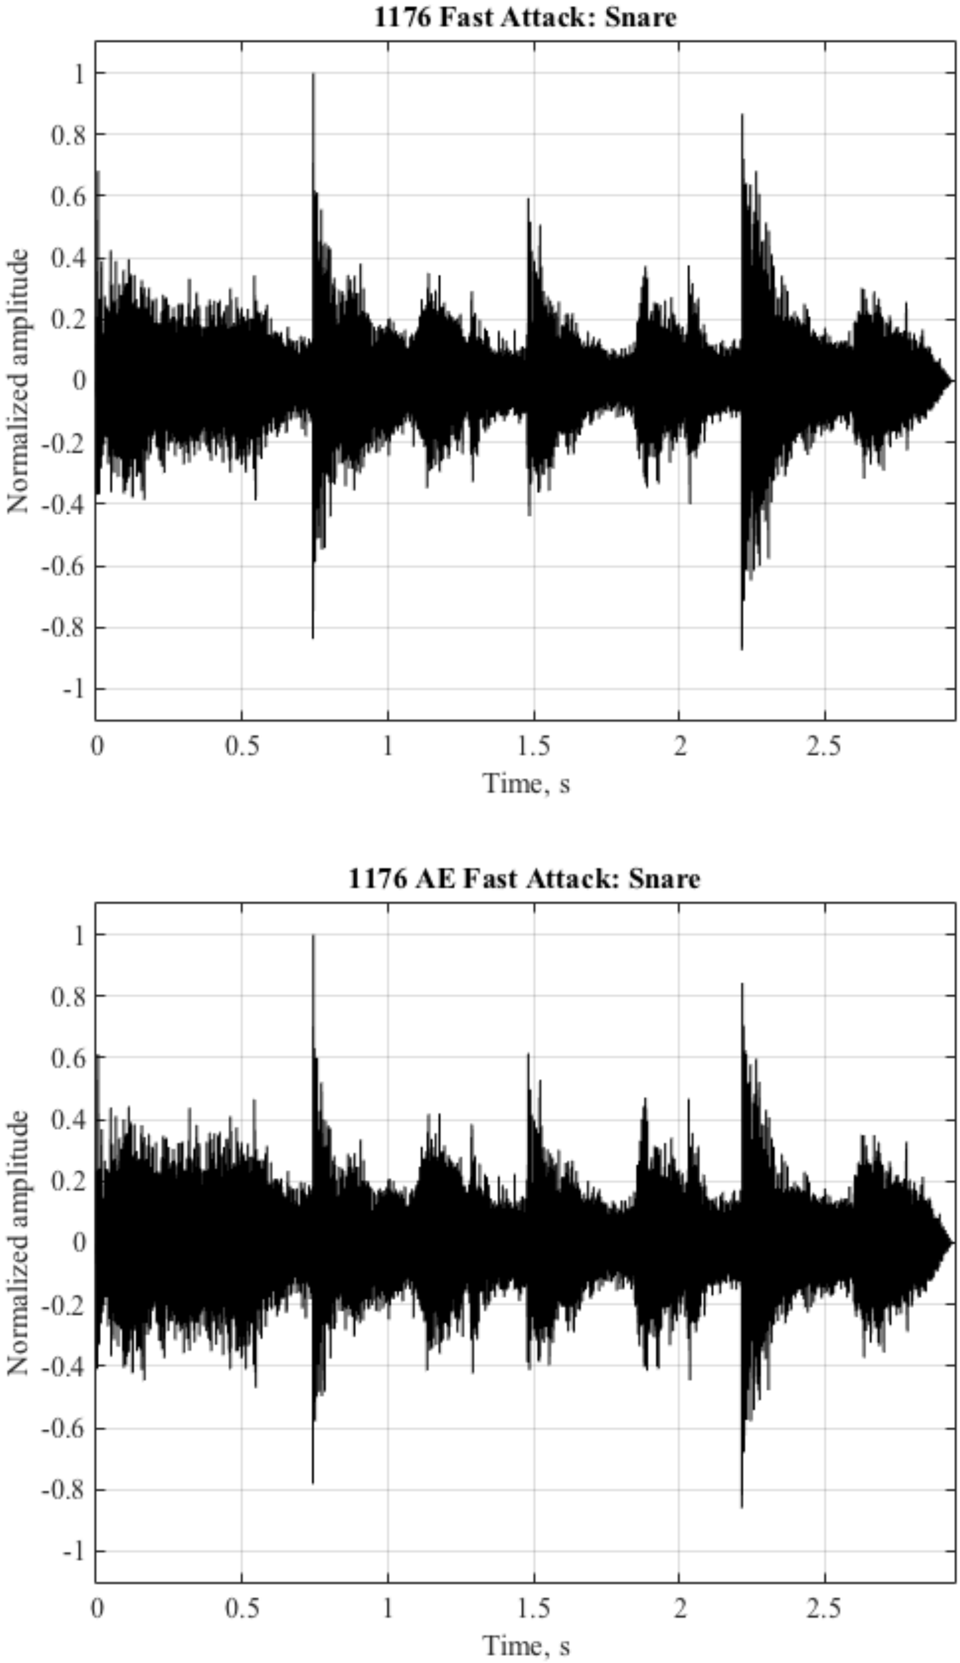

Figure 6-62: Compressed snare hits using fast attack times. Standard 1176 on top and 1176 AE on bottom

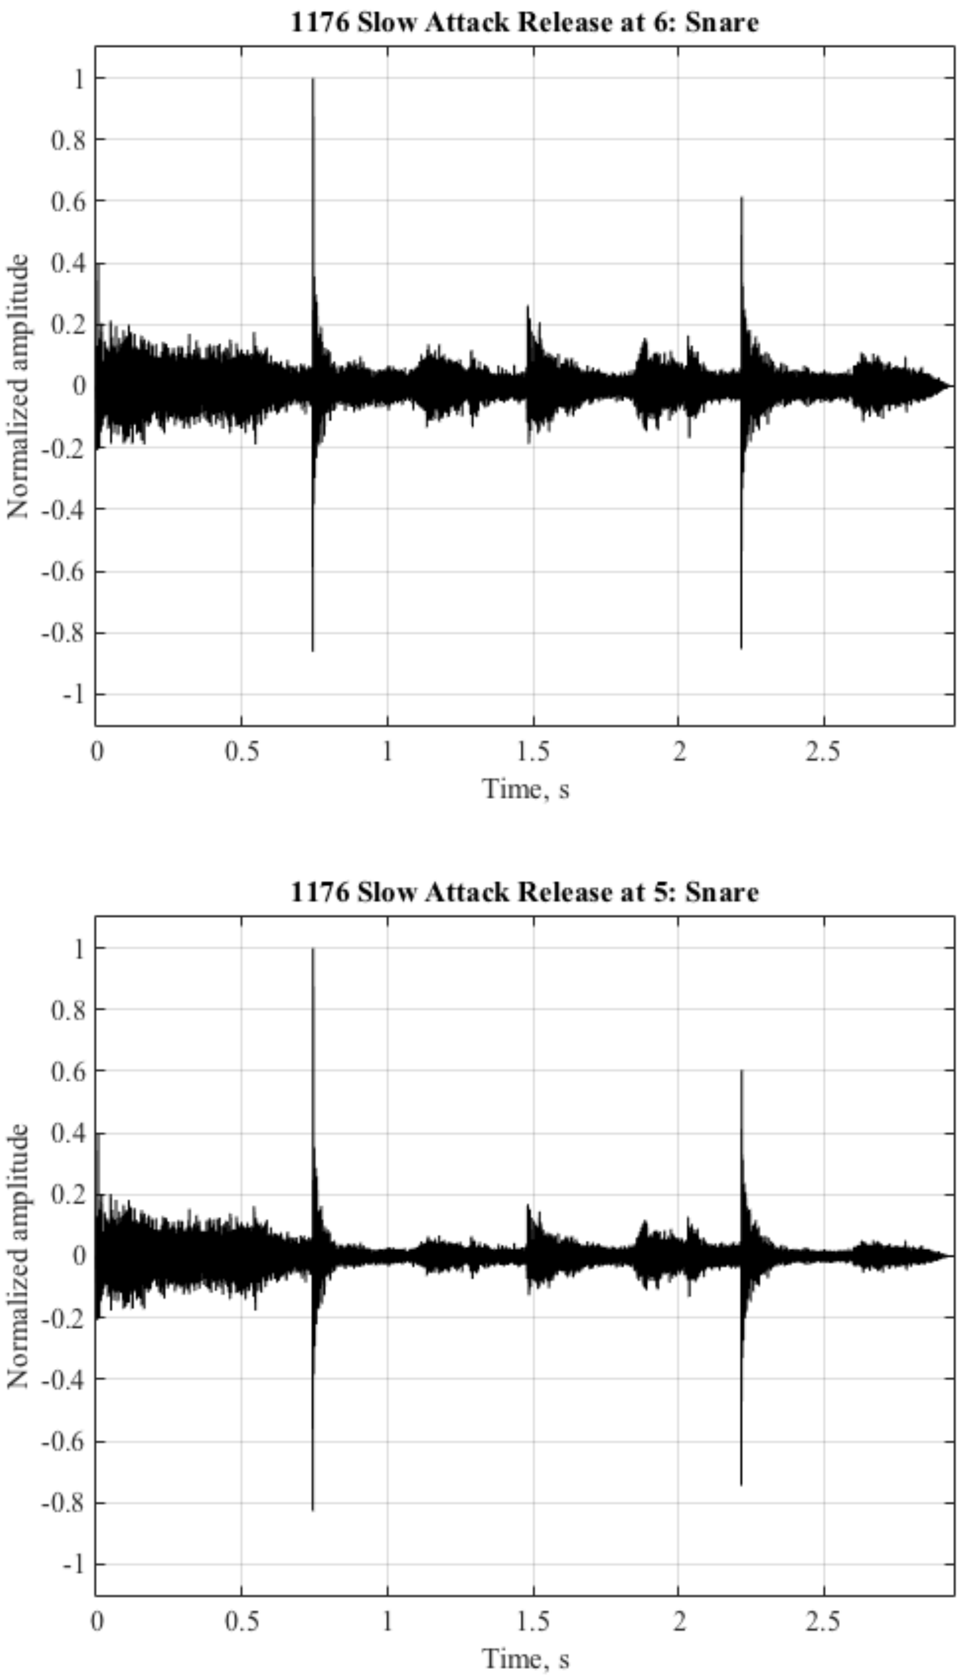

Figure 6-63: Compressed snare hits using slow attack and positions 6 (top) and 5 (bottom) release on a standard 1176

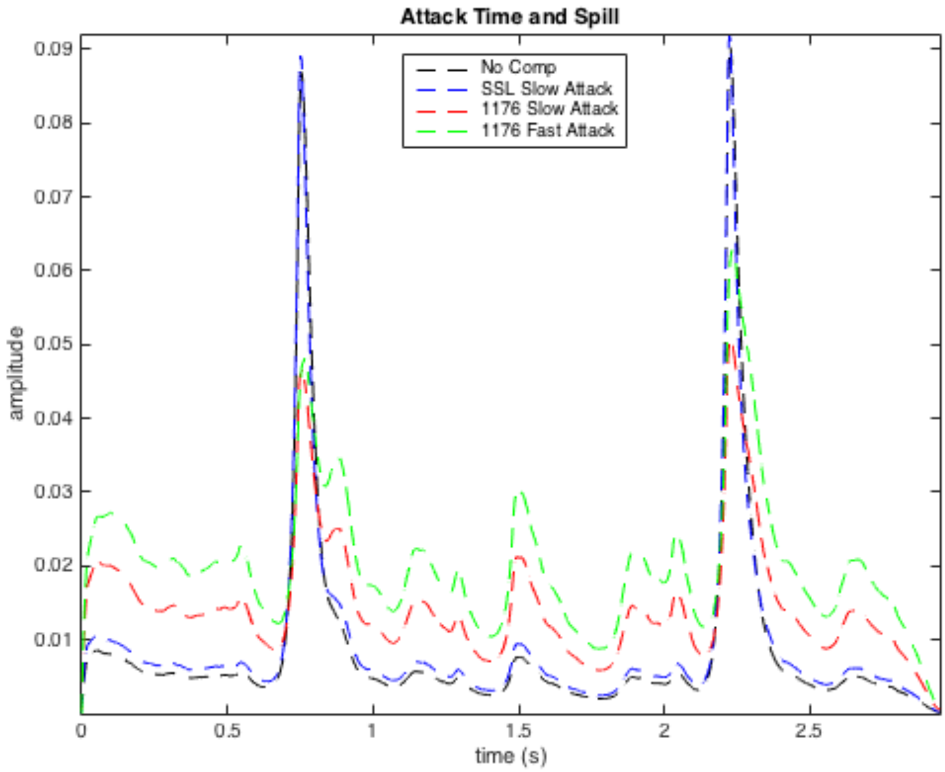

Figure 6-64: RMS energy of uncompressed and compressed drum hits for the SSL and 1176 compressors

6.16 Conclusion on Attack Time and Snare Testing

The tests made on the 1176's attack time and snare material demonstrate that the 1176 has a very fast attack time and the attack control is limited in range. Perceptually most of the change in attack occurs between positions 4 to 7. The variation over the rest of the range is limited with little audible difference. It can be stated again that this small range is perhaps one of the reasons why the 1176 is so popular. There is such a small margin for error regarding attack time choice that it is difficult to set it incorrectly. Given the trend shown in Chapter 4 for engineers to limit themselves to a small number of release times (often leaving it fixed on 7) it can be concluded that the main challenge when using the 1176 is deciding upon the correct amount of gain reduction. Perceptually the change in attack time from the 1176 sounds like a slight increase in brightness and clarity with little effect to the envelope the 1176 has already applied to the audio. This supports the claims made by engineers who state the attack can be considered a clarity control (Gearslutz, 2013).

When using the 1176 on drum sources, it has limited use in reshaping the transient to accentuate it. The 1176 can be used to an extent for this role (particularly with the attack at 1 and release at 6) but there is a limit to its useable range. This finding supports the results in Chapter 4 that revealed the 1176 is not as popular as VCA compressors for this style of compression.

The 1176 can be used to suppress the transient portion of sources such as drums. However, care needs to be taken to ensure extraneous noise, spill and room tone are not raised to such an extent they are as loud as the main source. A slower release can help, but there is a limited range available before pumping is introduced to the signal. An audio engineer can employ aggressive gating to counter this problem, but this can have an impact on the sonic signature of audio material that is not always desirable. A VCA based compressor such as the dbx165A, with a broad range of time constants, will allow for much more flexibility over transient shaping, both pronouncing and suppressing transients. This is made clear in Figures 122 and 124 that show the wide range of the dbx165A's attack control. This result demonstrates why engineers favoured VCA styles of compressor for spot mic drum sources in Chapter 4.

Chapter 7 : Overall Conclusions

7.1 Final Conclusions

This thesis set out to investigate the use of non-linearity in music production and in particular the use of DRC and the 1176 compressor. The study found that a significant body of music producers actively seek out non-linear distortion (and also linear distortion) in their music production equipment and the main motivation behind this is to impart unique sonic signatures on their audio. This approach was in contrast to audio equipment designers who had historically tried to minimise the amount of distortion in their circuits.

Music producers select dynamic range compressors based upon elements of their design and specifically how their sonic signature colours program material. A further study analysed the design of four popular compressors used in music production and discovered that certain aspects of their design, often down to the components used in the circuit, played a role in the equipment's sound quality.

To build a methodology for collecting information on sonic signatures, research was undertaken into audio system testing and timbre studies to develop an appropriate mixed methodology. This approach was successful in extracting salient properties of the equipment's sound quality that related to non-linearity and timbral change. This method was applied in a series of tests on compressors that were identified as being popular in Chapter 4, and their sonic signatures were discussed in chapters 5 and 6 and presented in a collection of short audio examples in the supporting data drive.

There are indeed differences between the compressors, but they are often small and not easily discernible until the amount of gain reduction is considerable, or the piece of equipment is used in a unique setting such as the 1176's all-buttons mode. It was found that design limitations play a role in a compressor's sonic signature and this, in turn, affects how appropriate the unit is for a specific compression related task. The 1176 was tested most thoroughly, and it was found in Chapter 6 that while its effect on dynamics is often transparent, it has the effect of changing the sonic signature of audio material by imparting subtle changes to its timbre. This compressor can also severely distort signals when used in its all-buttons mode and radically softens drum transients because of its fast acting time constants.

The grounded theory and content analysis study in Chapter 4 revealed that producers and engineers in this study are interested in timbral change rather than simple control of dynamics when using compression, particularly during the mix stage of a production. The timbral change they are seeking varies depending upon the program material, and some of these changes are more popular than others. In particular, it was found that when using vocal and bass guitar program material non-linear effects and colouration are important considerations, with steady state drum sources amplitude modulation (more commonly known as pumping) is the primary concern and for spot mic drum sources, envelope shaping, particularly pronouncing transients, is of paramount importance.

7.2 Answers to the core research questions

A number of questions were highlighted at the start of this thesis. The first question addressed how music producers apply DRC in their productions and what is the motivation behind its use? This question was answered in Chapter 4 where it was discovered that producers have some set approaches when applying compression. From the data analysed in the study, it was found that producers were more concerned with imparting various forms of colouration and non-linearity to program material than with simple control over dynamic range. This was particularly true during mixing due to producers using small amounts of compression when tracking to control variations in dynamics.

The second question asked how can audio be analysed to get meaningful information that relates to sonic signatures? This question was addressed in Chapter 2 and then applied in the objective testing of the compressors discussed in chapters 5 and 6. It was shown that a mixed methodology of audio testing (THD, IMD and tone bursts), feature extraction, time domain analysis, spectral analysis and critical listening is an effective methodology to adopt when investigating sonic signatures.

The third question asked what is the sonic signature of the 1176 compressor? This issue was addressed in Chapter 5 and more significantly in Chapter 6 where it was noted that much of the 1176's sonic signature comes from its time constants, colouration (perceived as textural thickness and bite), the limited range of the attack time and the non-linearity generated by all-buttons mode. Additionally, given the different compressors available in a professional recording studio this project asked what the differences are in their sonic signature and how do these differences manifest in audio material? Much of the non-linearity and colouration in a compressor comes from the device used for gain reduction and also the speed of the time constants. Differences in sonic signatures were addressed in objective testing revealing that under normal working conditions professional compressors are all quite similar, whereas more profound differences occurred during heavy gain reduction and under specific forms of application such as envelope shaping, deliberate distortion and amplitude modulation (pumping) of steady state program material.

The findings from this study help enhance our current understanding of how music producers use DRC in the music production process and adds substantially to our understanding of what is happening to audio signals at an objective and subjective level when they are processed with a range of compressors and specifically the 1176. Furthermore, it adds to the body of literature on sonic signatures and has made noteworthy contributions to developing a methodology that can be used by scholars in other related fields of study.

7.3 Areas for further research

In the future, it is hoped that sonic signatures can be catalogued to train engineers and students in their sound quality and act as a database for scholars and researchers. This catalogue of sonic signatures should cover a range of music production equipment and should be used to create a database of sonic signatures, made up of hard data such as plots, THD figures and audio features extracted from the equipment. Once a large enough sample of data has been collected, analysis should be conducted to group the devices into dimensions much like the PCA carried out in this thesis and the multidimensional scaling (MDS) used by scholars in timbre studies (Elliott, Hamilton, & Theunissen, 2013; Le Bagousse, Paquier, & Colomes, 2014).

The author of this current study believes that using a systematic approach to the collection of sonic signatures is paramount and it is hoped that working with the methodology developed in this thesis and building upon it will provide a systematic method of cataloguing the sonic signatures of music production equipment. This type of cataloguing has been an important part of the music production process in the industry, but the collection of data has rarely been systematic. Information on sonic signatures has typically been passed down as word of mouth from the experienced engineer to inexperienced engineer, with much of the information being anecdotal or gleaned retrospectively from recording sessions that were not fully documented. Thus, the information is unreliable, dependent upon opinion and the memory of the participant disseminating the data. By using a systematic approach, the unreliability of anecdotes can be removed and replaced with rigor and quantifiable data.

The importance of cataloging sonic signatures and applying this knowledge to the production process has been suggested in the past. Producer, engineer, and author Michael Stavrou (2003) attempted to create a method in his book *Mixing with your Mind* whereby recording engineers could perceptually rate microphones on a scale ranging from soft to hard and then use this information during tracking sessions to colour audio sources. While Stavrous' idea is novel and relates to timbre and sonic signatures (although he does not describe the process explicitly as such), it leaves many unanswered questions. Most importantly, what are the objective properties of sounds recorded with these microphones that make them subjectively soft or hard? Thus, the methodology used in this thesis can be applied to help get a better understanding of subjective descriptors of this nature and generate robust data on the acoustic properties impacting sonic signatures.

With that in mind, a future area of research should be to conduct listening experiments to investigate the perceptual effects of sonic signatures on groups of listeners and map the acoustic parameters to a number of perceptual qualities. Groups of listeners can be used from several distinct populations, such as experienced engineers, scholars in music production, and the general population with no experience in the production process. This will provide a more comprehensive picture of how end listeners perceive sonic signatures. There exist many methodologies for tests that will work for this experimental stage, but the sample size needs to be quite considerable, particularly for the inexperienced listeners. Any reader who wishes to move forward with experiments of this kind is directed to the work of Bech and Zacharov (2006), Alluri and Toiviainen (2010), De Man and Reiss (2014), Fenton, Lee and Wakefield (2014) and Moore and Wakefield, (2017) for ideas of appropriate listening test methodologies.

Pendersen and Zacharov (2015) are working towards the development of a sound wheel that presents various descriptors of sound quality in a layered format. The inner layers of the ring represent the main core groups, the middle rings include connected categories, and the outer ring features the discrete attributes. This design can be applied to the study of sonic signatures and the work in chapters 5 and 6 of this thesis can be developed to incorporate aspects of Penderson and Zacharov's approach, in particular, the use of a listening panel to establish the most appropriate descriptors and information on the lexicon. The author of this thesis is currently working on similar areas as part of ongoing research.

One final benefit of the systematic collection and creation of a sonic signatures database is that it facilitates the preservation of the unique fingerprints (or sonic DNA samples one could say) of famous pieces of music production equipment. Furthermore, it is argued that a study should be undertaken to create a central reference point for the storage of sound samples and measurement data extracted from these seminal pieces of equipment. The benefits of such a research project may impact a broad range of areas ranging from scholars of music to hardware and software developers to music historians and journalists.

Finally, settings used when compressing audio material with the 1176 are included in Table 7-1 for any producer or scholar hoping to recreate the sonic signatures discussed in this thesis. It is recommended the user adjusts the amount of gain reduction to vary the nature of sonic signatures and adjust the attack time to manipulate the amount of bite in the audio. This will differ for some sound sources, for example, bass requires a fast attack for additional brightness while the vocal requires a slower attack time. Users are free to experiment with the release time at other settings, but it is demonstrated in this thesis that the release set at 7 is the most common position across a range of sources thus it is an integral part of the 1176's sonic signature.

Source	Time Constants	Ratio	Sonic Signature
Vocal General	Attack 3 Release 7	4:1	Thick and full
Vocal Distorted	Attack 7 Release 7	All-Buttons	Distorted and aggressive
Bass General	Attack 3 Release 7	4:1	Thick and full
Bass Distorted	Attack 7 Release 7	20:1	Thick, full, biting and distorted
Drum Room Gentle	Attack 3 Release 7	4:1	Cohesive and consistent
Drum Room Distorted	Attack 3 Release 7	All-Buttons	Distorted, pumping, aggressive

Table 7-1: 1176 settings derived from research in this thesis

Appendices

Appendix 1

The scree plots created during PCA that feature the loadings of each principle component are included here.

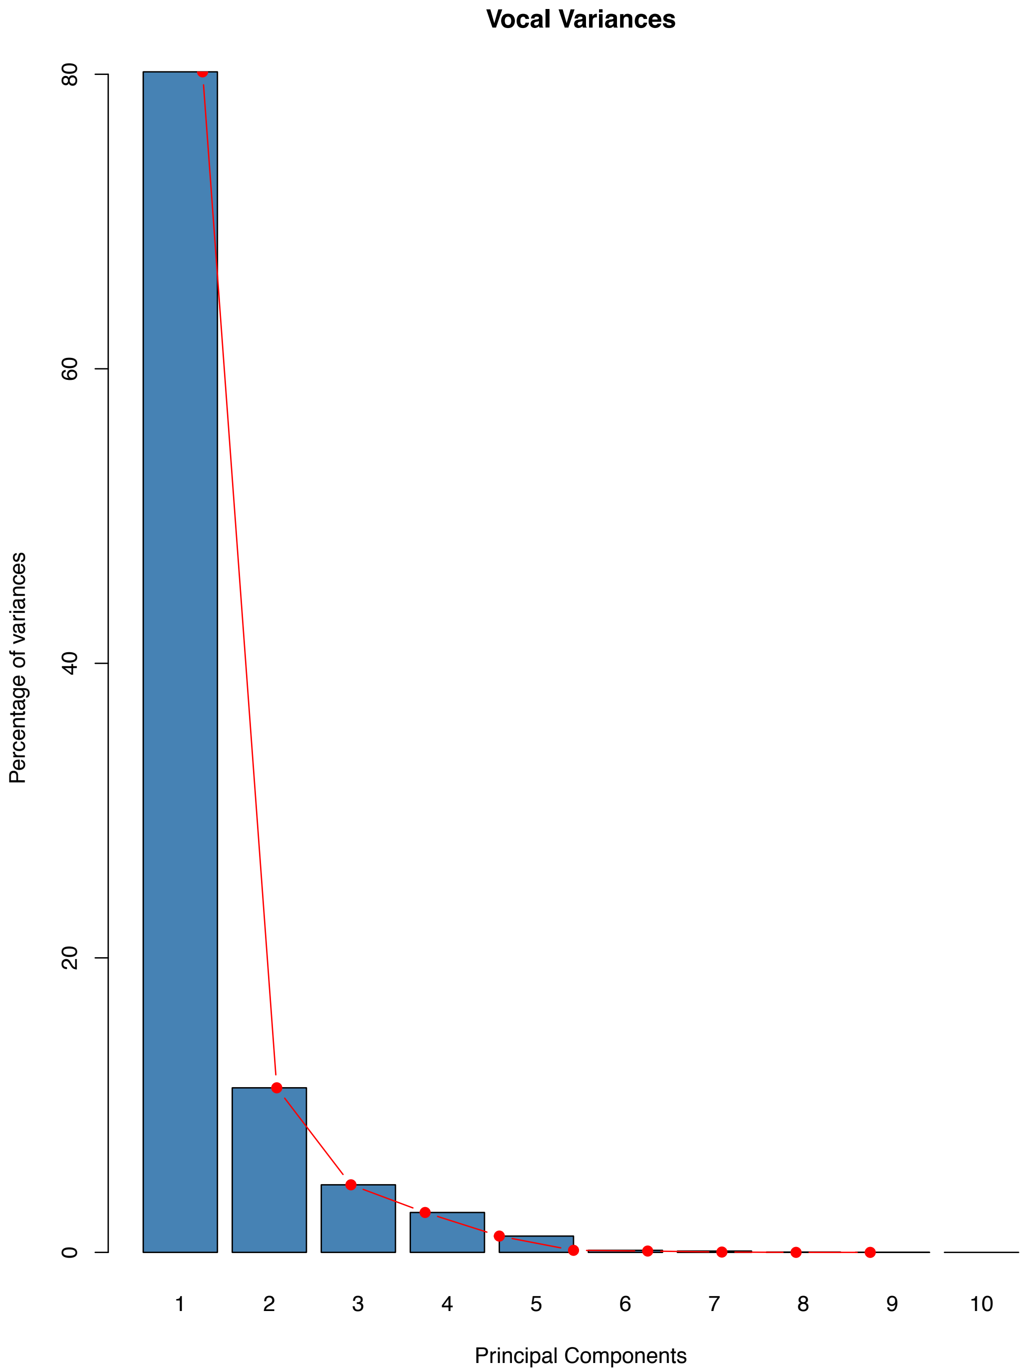

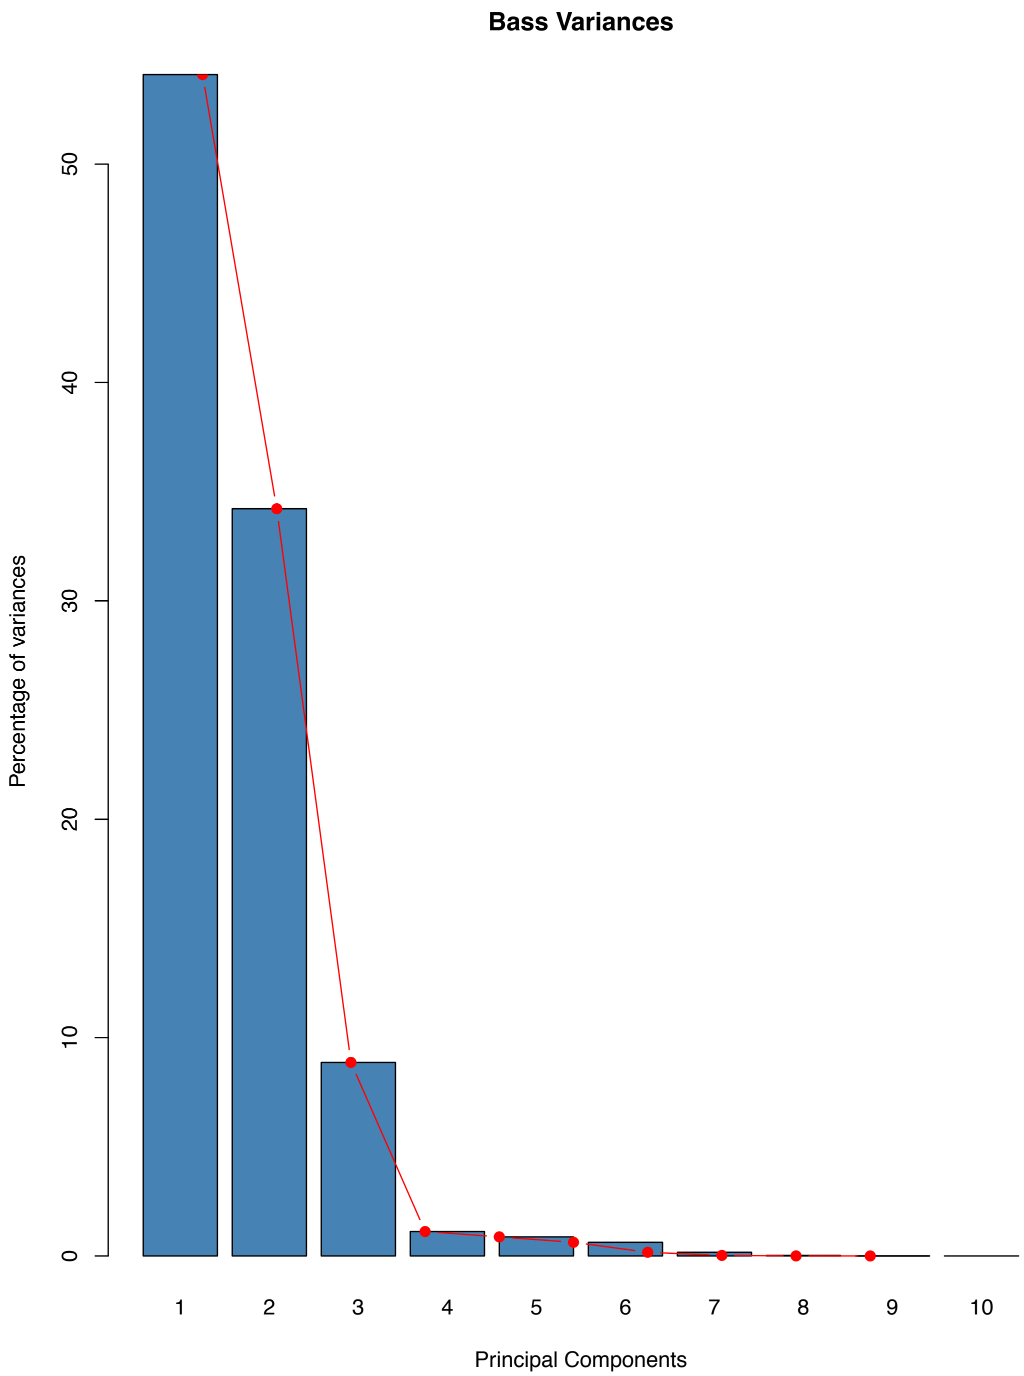
Bass Variances
Percentage of variances
0
10
20
30
40
50
1
2
3
4
5
6
7
8
9
10
Principal Components

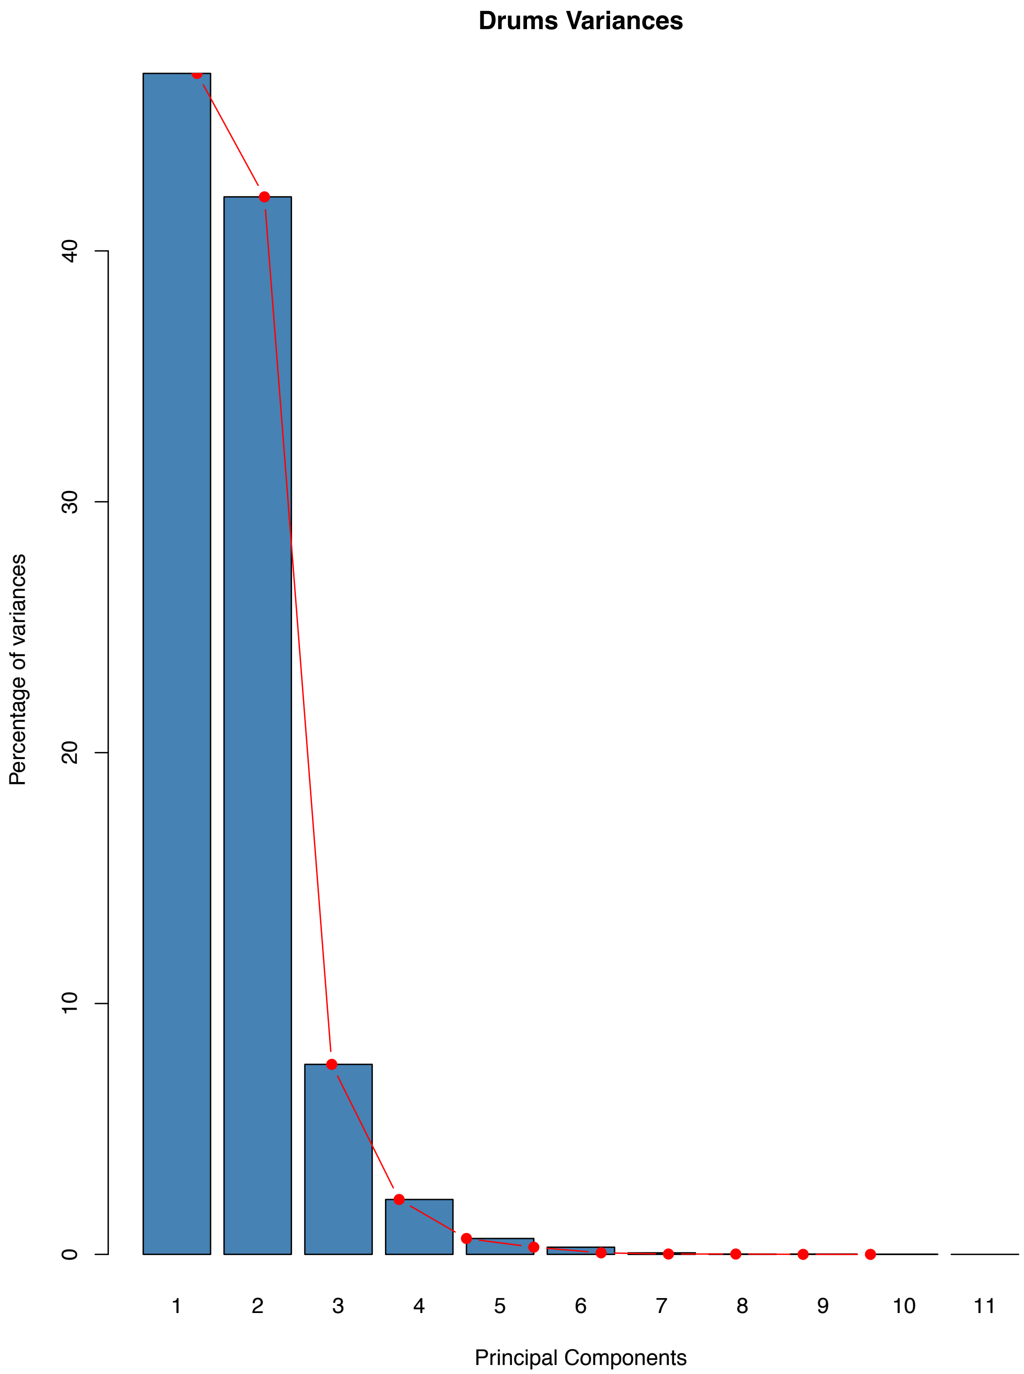
Drums Variances
Percentage of variances
0
10
20
30
40
1
2
3
4
5
6
7
8
9
10
11
Principal Components

Appendix 2

Score plots with loadings for the room mic audio is included here.

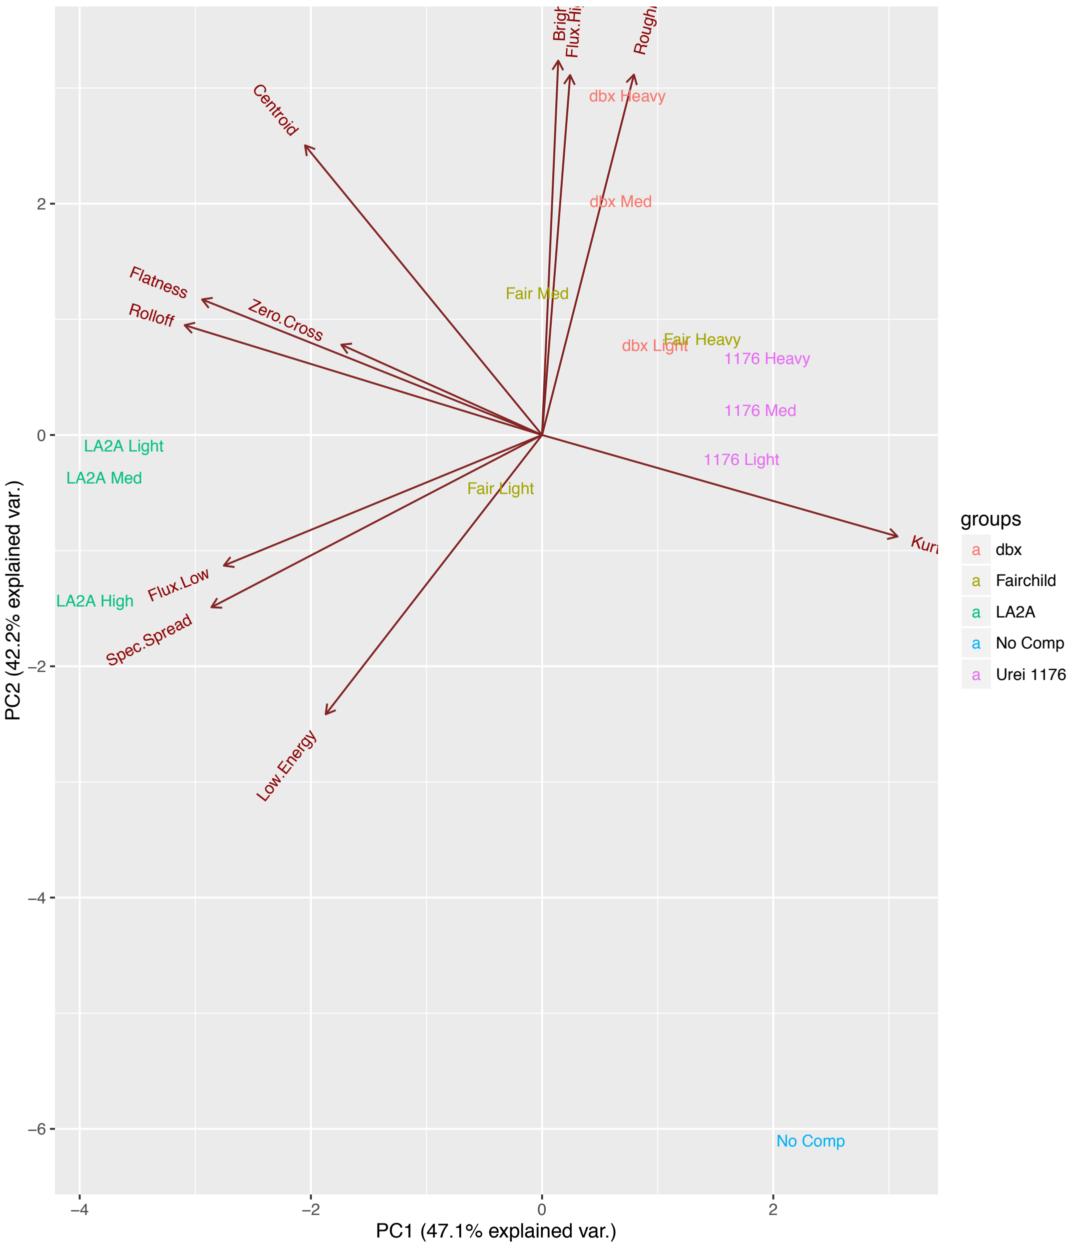

Score plots with loadings for the bass audio is included here.

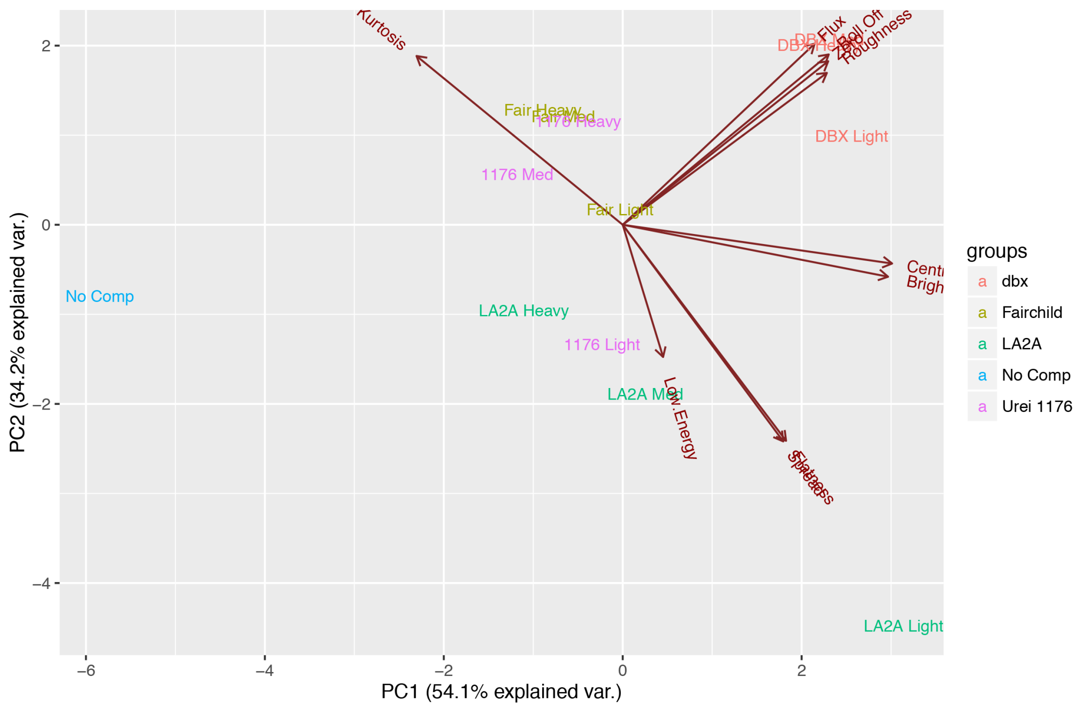

Score plots with loadings for the vocal audio is included here.

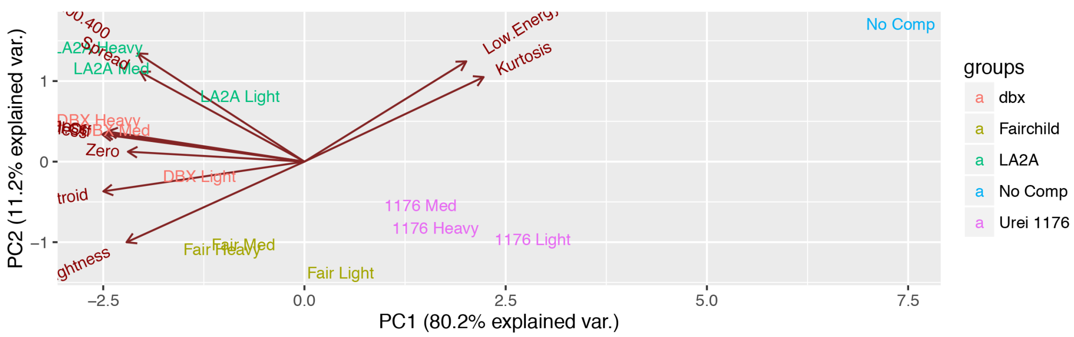

Loadings of each feature in the first three principle components are included here.

Bass Features	PC1	PC2	PC3
Centroid	0.42449553	-0.07667446	0.039870189
Brightness	0.41808031	-0.10284485	-0.034470753
Flux	0.30240528	0.35822626	0.137642924
Spread	0.25296986	-0.42869582	0.151889673
Low.Energy	0.06365512	-0.26165664	-0.912465724
Roughness	0.32169819	0.30053109	-0.246265525
Zero	0.32378175	0.32360547	-0.108038641
Roll.Off	0.32464287	0.33703387	-0.006351299
Flatness	0.25737793	-0.42752532	0.110213022
Kurtosis	-0.32489534	0.33401264	-0.195207799
Statistics	**PC1**	**PC2**	**PC3**
Standard deviation	2.326	1.8498	0.94151
Proportion of Variance	0.541	0.3422	0.08864
Cumulative Proportion	0.541	0.8832	0.97183

Drum Features	**PC1**	**PC2**	**PC3**
Centroid	-0.2768277	0.3573378	0.07466824
Brightness	0.0188533	0.4618109	0.0337568
Flux.Low	-0.3716511	-0.1611288	0.20907302
Flux.High	0.0327537	0.4440337	0.27682615
Roughness	0.1072552	0.4445643	0.13101412
Rolloff	-0.4174603	0.1354774	0.08287712
Low.Energy	-0.2527092	-0.3445656	0.24865477
Zero.Cross	-0.2345225	0.111364	-0.88239667
Spec.Spread	-0.3861807	-0.2125189	0.07974539
Flatness	-0.3971209	0.1672172	-0.02774714
Kurtosis	0.4151307	-0.1251685	-0.03640399
Statistics	**PC1**	**PC2**	**PC3**
Standard deviation	2.2756	2.1534	0.91276
Proportion of Variance	0.4708	0.4215	0.07574
Cumulative Proportion	0.4708	0.8923	0.96803

Vocal Features	PC1	PC2	PC3
Centroid	-0.346953	-0.13727433	-0.07056845
Brightness	-0.3069131	-0.37224288	-0.36190281
Spread	-0.2834318	0.41447443	0.58401393
Zero	-0.3046101	0.04587703	-0.43536477
Low.Energy	0.2791542	0.46286315	-0.4094918
Roughness	-0.3390274	0.13845295	-0.19806811
Spec.Flux.200.400	-0.2879204	0.49970788	-0.2541701
Roll.Off	-0.3454524	0.13189973	0.2002503
Flatness	-0.3488354	0.12581103	0.04626027
Kurtosis	0.3090612	0.3903227	-0.140266
Statistics	**PC1**	**PC2**	**PC3**
Standard deviation	2.8313	1.0571	0.67744
Proportion of Variance	0.8016	0.1118	0.04589
Cumulative Proportion	0.8016	0.9134	0.95928

Appendix 3

The raw data from the features extracted for Chapter 5 and used for PCA are included here.

Features for all the drum settings are in the table below.

Comp	Centroid	Brightness	Flux Low	Flux High	Roughness	Rolloff	Low-Energy	Zero Cross	Spec Spread	Flatness	Kurtosis
dbx Heavy	4693.528	0.701	8.906	10.644	441.2	8700.759	0.452	2228.6	4411	0.357	4.79
dbx Med	4650.347	0.694	9.121	10.298	422.2	8670.478	0.476	2231.3	4414	0.353	4.79
dbx Light	4574.633	0.681	9.436	9.877	394.98	8599.486	0.464	2229.4	4426	0.348	4.82
1176 Heavy	4579.598	0.683	9.089	10.359	406.97	8574.757	0.517	2177.3	4431	0.348	4.88
1176 Med	4565.945	0.68	9.289	10.215	400.94	8572.065	0.529	2178.3	4436	0.343	4.88
1176 Light	4550.9	0.677	9.476	10.045	387.63	8569.541	0.535	2177.7	4436	0.341	4.86
Fair Heavy	4598.188	0.683	9.235	10.176	403.98	8613.281	0.488	2199.3	4453	0.347	4.85
Fair Med	4633.311	0.683	9.854	10.474	425.21	8726.835	0.488	2206.2	4481	0.349	4.74
Fair Light	4596.044	0.674	9.714	9.381	351.01	8742.312	0.535	2204.1	4512	0.347	4.7
LA2A High	4662.397	0.665	10.245	9.415	310.04	8936.447	0.672	2218.4	4721	0.369	4.62
LA2A Med	4704.327	0.673	10.123	9.525	336.22	8984.056	0.619	2224.3	4691	0.37	4.6
LA2A Light	4700.186	0.672	10.114	9.484	344.83	8998.691	0.577	2220.5	4671	0.37	4.57
No Comp	4296.827	0.626	9.463	7.268	212.94	8345.798	0.625	2199.7	4553	0.329	5.01

Features for all the bass settings are in the table below.

Comp	Centroid	Brightness	Flux	Spread	Low-Energy	Roughness	Zero	Roll Off	Flatness	Kurtosis
No Comp	256.51	0.0141	2.33	538.88	0.413	1.52	105.41	420	0.00269	5.79
DBX Heavy	312.39	0.0184	3.17	649.64	0.38	2.93	135.83	564	0.0046	4.19
DBX Med	311.87	0.0183	3.24	659.12	0.383	2.99	136.02	562	0.00468	4.20
DBX Light	315.04	0.0185	3.08	711.7	0.406	2.89	136.19	559	0.00552	3.89
Fair Heavy	291.99	0.017	2.97	599.97	0.36	1.91	127.61	512	0.00369	4.69
Fair Med	290.25	0.0167	3.01	610.29	0.39	2.34	129.29	493	0.00379	4.73
Fair Light	296.11	0.0172	2.88	676.73	0.398	2.38	128.26	493	0.00491	4.21
LA2A Heavy	292.68	0.0168	2.67	723.53	0.368	1.78	120.68	481	0.00555	3.99
LA2A Med	301.02	0.0175	2.77	808.79	0.39	2.04	121.74	470	0.00707	3.36
LA2A Light	328.08	0.0199	2.75	1010.87	0.421	2.11	119.4	484	0.01108	2.29
1176 Heavy	293.64	0.0167	3.1	652.42	0.348	2.18	121.24	519	0.00446	4.45
1176 Med	290.02	0.0162	2.85	677.59	0.333	2.21	121.66	486	0.00478	4.36
1176 Light	297.42	0.0172	2.8	768.45	0.39	1.91	121.77	480	0.00613	3.68

Features for all the vocal settings are in the table below.

Comp	Centroid	Brightness	Spread	Zero	Low-Energy	Roughness	Spec Flux 200-400	Roll Off	Flatness	Kurtosis
DBX Heavy	5109	0.61781	4785	2003.1	0.41759	223.799	7.4024	10500.04	0.375	2.81
DBX Med	5104	0.61741	4784	2003.3	0.41892	222.1681	7.3357	10493.65	0.374	2.814
DBX Light	5065	0.61212	4768	2003.6	0.42155	199.8776	7.148	10460.84	0.368	2.813
Fair Heavy	5075	0.61297	4781	1962	0.38866	187.5992	7.0345	10467.91	0.365	2.789
Fair Med	5051	0.61115	4770	1960	0.38533	198.673	7.0398	10444.02	0.363	2.805
Fair Light	5008	0.60423	4754	1973	0.41267	169.917	6.8454	10416.09	0.353	2.793
LA2A Heavy	5063	0.59898	4853	1997	0.40103	216.9979	7.3706	10540.75	0.377	2.819
LA2A Med	5067	0.59943	4854	1993	0.39925	206.8682	7.2942	10544.28	0.378	2.817
LA2A Light	5012	0.59257	4824	1985	0.4325	184.7813	7.1399	10490.37	0.367	2.805
1176 Heavy	4917	0.59496	4742	1956	0.40192	167.6394	6.8944	10361.42	0.352	2.86
1176 Med	4927	0.59588	4746	1967	0.41524	158.8697	6.9608	10371.26	0.353	2.854
1176 Light	4864	0.58762	4714	1979	0.43716	142.5697	6.7845	10312.72	0.341	2.849
No Comp	4655	0.55626	4704	1909	0.53204	117.1486	6.8308	10203.87	0.322	2.99

Bibliography

Adobe. (2016). Analyze phase, frequency, & amplitude with Audition. Retrieved July 19, 2016, from https://helpx.adobe.com/audition/using/analyzing-phase-frequency-amplitude.html

Alluri, V., & Toiviainen, P. (2010). Exploring perceptual and acoustical correlates of polyphonic timbre. *Music Perception: An Interdisciplinary Journal*, *27*(3), 223–242.

Altheide, D. L. (1987). Reflections: Ethnographic content analysis. *Qualitative Sociology*, *10*(1), 65–77.

Analoguetube. (n.d.). Analoguetube Fairchild 670 and 660 compressor recreations | NEWS. Retrieved July 16, 2016, from http://www.analoguetube.com/#!blog/c22eg

Anderson, J. (2006, September 6). An Evening With Chris Lord-Alge | Emusician. Retrieved August 22, 2014, from http://www.emusician.com/gear/1332/an-evening-with-chris-lord-alge/36961

ANSI. (1973). S3. 20. *New York, NY: American National Standards Institute*.

Avron, N. (2008). Secrets Of The Mix Engineers: Neal Avron. Retrieved January 10, 2016, from https://www.soundonsound.com/sos/jan08/articles/itavron.htm

Ballard, G. (2003). Glen Ballard. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/mar03/articles/glennballard.asp

Barbour, E. (1998). The cool sound of tubes [vacuum tube musical applications]. *Spectrum, IEEE*, *35*(8), 24–35.

Barresi, J. (2011). *Joe Barresi Tracking Rock* [Digital Download]. Pasadena.

Bech, S. (1999). Methods for Subjective Evaluation of Spatial Characteristics of Sound. In *Audio Engineering Society Conference: 16th International*

Conference: Spatial Sound Reproduction. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=8010

Bech, S., & Zacharov, N. (2006). *Perceptual audio evaluation: theory, method and application*. Chichester, England ; Hoboken, NJ: John Wiley & Sons.

Bennett, S. (2010). *Examining the emergence and subsequent proliferation of anti production amongst the popular music producing elite*. University of Surrey.

Berners, D. (2005). UA WebZine "Ask the Doctors!" July 05 | Peak vs. RMS Detection. Retrieved October 26, 2015, from http://www.uaudio.com/webzine/2005/july/text/content2.html

Berners, D. (2006). UA WebZine "Analog Obsession" April 06 | Compression Technology and Topology. Retrieved October 28, 2015, from http://www.uaudio.com/webzine/2006/april/text/content4.html

Berners, D. (2009a). Total Harmonic Distortion - Blog - Universal Audio. Retrieved July 16, 2014, from http://www.uaudio.com/blog/total-harmonic-distortion/

Berners, D. (2009b). Total Harmonic Distortion - Blog - Universal Audio. Retrieved July 16, 2014, from http://www.uaudio.com/blog/total-harmonic-distortion/

Berners, D., & Abel, J. (2004). UA WebZine "Ask the Doctors!" February 04 | Drs. David P. Berners and Jonathan S. Abel Answer Your Signal Processing Questions. Retrieved September 24, 2015, from http://www.uaudio.com/webzine/2004/february/text/content2.html

Bogdanov, D., Wack, N., Gómez, E., Gulati, S., Herrera, P., Mayor, O., … Serra, X. (2013). Essentia: an open-source library for sound and music analysis. In *Proceedings of the 21st ACM international conference on Multimedia* (pp. 855–858). ACM.

Bohn, D. (2000). Audio Specifications. Retrieved October 23, 2015, from http://www.rane.com/note145.html

Brauer, M. (2001). Michael H. Brauer -- MixOnline: Pushing the Envelope. Retrieved November 22, 2015, from http://www.mbrauer.com/articles/compression.asp

Brauer, M. (2008). Secrets Of The Mix Engineers: Michael Brauer. Retrieved January 10, 2016, from https://www.soundonsound.com/sos/nov08/articles/itbrauer.htm

Brice, R. (2001). *Music engineering: the electronics of playing and recording* (2nd ed). Oxford ; Burlington, MA: Newnes.

Burr, V. (1995). *An Introduction to Social Constructionism*. Routledge. Retrieved from https://books.google.co.uk/books?id=7du5229Pc4EC

Busch, C., De Maret, P. S., Flynn, T., Kellum, R., Le, S., Meyers, B., & Palmquist, M. (1994). *Content Analysis. Writing@ CSU. Colorado State University*.

Buskin, R. (2011). John Cougar: "Jack & Diane." Retrieved November 1, 2015, from http://www.soundonsound.com/sos/sep11/articles/classic-tracks-0911.htm

Campbell, W., Paterson, J., & Toulson, R. (2014). A Quantitative Evaluation of Signal Masking in Summed and Compressed Audio.

Case, A. U. (2007). *Sound FX: unlocking the creative potential of recording studio effects*. Amsterdam ; Boston: Focal Press.

Charmaz, K. (2006). *Constructing grounded theory*. London ; Thousand Oaks, Calif: Sage Publications.

Chiccarelli, J. (2007). Secrets Of The Mix Engineers: Joe Chiccarelli. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/oct07/articles/insidetrack_1007.htm

Ciletti, E., Hill, D., & Wolf, P. (2008). GAIN CONTROL DEVICES, SIDE CHAINS, AUDIO AMPLIFIERS. Retrieved September 26, 2014, from http://www.tangible-technology.com/dynamics/comp_lim_ec_dh_pw2.html

Clearmountain, B. (1999). BOB CLEARMOUNTAIN: Master Mixer. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/jun99/articles/bobclear.htm

Corey, J. (2010). *Audio production and critical listening technical ear training*. Burlington, MA: Focal Press. Retrieved from http://public.eblib.com/EBLPublic/PublicView.do?ptiID=610530

Czerwinski, E., Voishvillo, A., Alexandrov, S., & Terekhov, A. (2001). Multitone Testing of Sound System Components'Some Results and Conclusions, Part 1: History and Theory. *J. Audio Eng. Soc*, *49*(11), 1011–1048.

De Man, B., & Reiss, J. D. (2014). Adaptive Control of Amplitude Distortion Effects. In *Audio Engineering Society Conference: 53rd International Conference: Semantic Audio*. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=17118

Dove, S. (2008). *Handbook for sound engineers*. (G. Ballou, Ed.) (4th ed). Amsterdam ; Boston: Focal Press.

Dresdow, D. (2009). Secrets Of The Mix Engineers: Dylan "3D" Dresdow. Retrieved January 10, 2016, from https://www.soundonsound.com/sos/jul09/articles/it_0709.htm

EBU. (1992). Tech. 3219-6- Audio measurement techniques for television tape–recorders with FM audio channels. Retrieved from https://tech.ebu.ch/docs/tech/tech3219.pdf

EBU–Recommendation, R. (2011). Loudness normalisation and permitted maximum level of audio signals.

Eidenberger, H. (2011). *Fundamental media understanding: the common methods of audio retrieval, biosignal processing, content-based image retrieval, face recognition, genome analysis, music genre classification, speech recognition, technical stock analysis, text retrieval and video surveillance* (2. ed). Norderstedt: Books on Demand.

Electric Light Orchestra. (1979). *Don't Bring Me Down*. Jet.

Elliot, R. (2006). Soft Clipping. Retrieved January 4, 2016, from http://sound.westhost.com/articles/soft-clip.htm

Elliot, R. (2015). Squarewave Testing. Retrieved from http://sound.whsites.net/articles/squarewave.htm

Elliott, T. M., Hamilton, L. S., & Theunissen, F. E. (2013). Acoustic structure of the five perceptual dimensions of timbre in orchestral instrument tones. *The Journal of the Acoustical Society of America*, *133*(1), 389–404.

Elmhirst, T. (2011). Tom Elmhirst: Recording Adele "Rolling In The Deep." Retrieved January 10, 2016, from https://www.soundonsound.com/sos/sep11/articles/it-0911.htm

Enderby, S., & Baracskai, Z. (2012). HARMONIC INSTABILITY OF DIGITAL SOFT CLIPPING ALGORITHMS. Presented at the Int. Conference on Digital Audio Effects, York.

Fairchild Recording Equipment. (1959). Fairchild 670 Manual. Retrieved from http://thehistoryofrecording.com/Manuals/Fairchild/Fairchild_670_stereo_limiting_amplifier_instructions_Schematics.pdf

Farmelo, A., & Hampton, S. (2010). Using Transformers: to Transform Audio | Tape Op - the Creative Music Recording Magazine. Retrieved October 6, 2014, from http://www.tapeop.com/tutorials/76/transformers/

Fenton, S., Fazenda, B., & Wakefield, J. (2011). Objective Measurement of Music Quality Using Inter-Band Relationship Analysis. Presented at the Audio Engineering Society Convention 130, Audio Engineering Society. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=15861

Fenton, S., Lee, H., & Wakefield, J. (2014). Elicitation and Objective Grading of Punch within Produced Music. In *Audio Engineering Society Convention 136*. Audio Engineering Society.

Fitzgerald, R. A., & Lindsay, A. T. (2004). Tying semantic labels to computational descriptors of similar timbres. *Proceedings of Sound and Music Computing, 2004*.

Fuston, L. (2012). Universal Audio's Classic 1176 Compressor — A History - Blog - Universal Audio. Retrieved October 11, 2014, from http://www.uaudio.com/blog/analog-obsession-1176-history/

Gabrielsson, A., & Sjögren, H. (1972). Detection of Amplitude Distortion in Flute and Clarinet Spectra. *The Journal of the Acoustical Society of America*, *52*(2A), 471–483. https://doi.org/http://dx.doi.org/10.1121/1.1913124

Gaffney, D. (2009). Secrets Of The Mix Engineers: Declan Gaffney. Retrieved January 10, 2016, from https://www.soundonsound.com/sos/jun09/articles/itu2.htm

Gearslutz. (2013). 1176 attack - Gearslutz.com. Retrieved November 11, 2015, from https://www.gearslutz.com/board/so-much-gear-so-little-time/817016-1176-attack.html

Gearslutz. (2014). Help me define some subjective terms relating to compression - Gearslutz.com. Retrieved November 7, 2015, from https://www.gearslutz.com/board/so-much-gear-so-little-time/937491-help-me-define-some-subjective-terms-relating-compression.html

Geddes, E. R., & Lee, L. W. (2003). Auditory Perception of Nonlinear Distortion - Theory. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=12464

Giannoulis, D., Massberg, M., & Reiss, J. D. (2012). Digital dynamic range compressor design—a tutorial and analysis. *Journal of the Audio Engineering Society*, *60*(6), 399–408.

Giannoulis, D., Massberg, M., & Reiss, J. D. (2013). Parameter automation in a dynamic range compressor. *Journal of the Audio Engineering Society*, *61*(10), 716–726.

Glaser, B. G., & Holton, J. (2007). Remodeling grounded theory. *Historical Social Research/Historische Sozialforschung. Supplement*, 47–68.

Glaser, B. G., & Strauss, A. L. (2009). *The discovery of grounded theory: strategies for qualitative research* (4. paperback printing). New Brunswick: Aldine.

Goddard, W., & Melville, S. (2001). *Research methodology: an introduction*. Lansdowne: Juta.

Gottinger, B. (2007). *Rethinking Distortion: Towards a Theory of'sonic Signatures'.* (Doctor of Philosophy). New York University.

Hamm, R. O. (1973). Tubes Versus Transistors-Is There an Audible Difference. *J. Audio Eng. Soc*, *21*(4), 267–273.

Herrera, P., & Streich, S. (2005). Detrended Fluctuation Analysis of Music Signals: Danceability Estimation and Further Semantic Characterization. In *Audio*

Engineering Society Convention 118. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=13125

Hicks, M. (2009). Audio Compression Basics - Blog - Universal Audio. Retrieved September 12, 2013, from http://www.uaudio.com/blog/audio-compression-basics/

Hicks, M. (2012). Tips & Tricks — 1176 Classic Limiter Collection - Blog - Universal Audio. Retrieved July 22, 2016, from http://www.uaudio.com/blog/1176-collection-tips/

Hodges, B. D., Kuper, A., & Reeves, S. (2008). Discourse analysis. *BMJ*, *337*.

Hummersone, C. (2016). Long term average spectrum - File Exchange - MATLAB Central. Retrieved July 18, 2016, from https://www.mathworks.com/matlabcentral/fileexchange/55212-long-term-average-spectrum

Iverson, P., & Krumhansl, C. L. (1993). Isolating the dynamic attributes of musical timbrea). *The Journal of the Acoustical Society of America*, *94*(5), 2595–2603.

Jeffs, R., Bohn, D., & Holden, S. (2005). Dynamics Processors -- Technology & Applications. Retrieved October 26, 2015, from http://www.rane.com/note155.html

Johnson, A. M., & Kempster, G. B. (2011). Classification of the classical male singing voice using long-term average spectrum. *Journal of Voice*, *25*(5), 538–543.

Jørgensen, M. W., & Phillips, L. J. (2002). *Discourse Analysis as Theory and Method*. SAGE Publications. Retrieved from https://books.google.co.uk/books?id=zw7njKgg_pwC

Joshua, J. (2010). Secrets Of The Mix Engineers: Jaycen Joshua. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/aug10/articles/it-0810.htm

Journal on the Art of Record Production. (n.d.). Journal on the Art of Record Production » ISSN: 1754-9892. Retrieved November 8, 2015, from http://arpjournal.com/

Keen, R. G. (2000). Guitar Effects FAQ. Retrieved March 9, 2015, from http://www.geofex.com/effxfaq/fxfaq.htm

Kemp, M. J. (1999). Analysis and Simulation of Non-Linear Audio Processes Using Finite Impulse Responses Derived at Multiple Impulse Amplitudes. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=8261

Kester, W. (1993). Analog Devices: Analog Dialogue: Ask The Applications Engineer - 13. Retrieved September 3, 2014, from http://www.analog.com/library/analogdialogue/Anniversary/13.html

Lakatos, S. (2000). A common perceptual space for harmonic and percussive timbres. *Perception & Psychophysics*, *62*(7), 1426–1439.

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

Lartillot, O. (2011). MIRtoolbox - File Exchange - MATLAB Central. Retrieved July 18, 2016, from http://uk.mathworks.com/matlabcentral/fileexchange/24583-mirtoolbox

Lartillot, O. (2015). *MIRToolbx 1.6 Manual*. Denmark: Aalborg University. Retrieved from https://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials/mirtoolbox/MIRtoolbox1.6guide

Lartillot, O., Toiviainen, P., & Eerola, T. (2008). A matlab toolbox for music information retrieval. In *Data analysis, machine learning and applications* (pp. 261–268). Springer.

Laurier, C., Lartillot, O., Eerola, T., & Toiviainen, P. (2009). Exploring relationships between audio features and emotion in music.

Le Bagousse, S., Paquier, M., & Colomes, C. (2014). Categorization of Sound Attributes for Audio Quality Assessment—A Lexical Study. *J. Audio Eng. Soc*, *62*(11), 736–747.

LeGreco, M., & Tracy, S. J. (2009). Discourse Tracing as Qualitative Practice. *Qualitative Inquiry*, *15*(9), 1516–1543.

Liniere, D. (2013). Thinking Outside The Box: Part 3 Stripping The Channel. *Audio Technology*, (6).

Löfqvist, A., & Mandersson, B. (1987). Long-time average spectrum of speech and voice analysis. *Folia Phoniatrica et Logopaedica*, *39*(5), 221–229.

Lord-Alge, T. (2000). TOM LORD-ALGE: From Manson To Hanson. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/apr00/articles/tomlord.htm

Loughran, R., Walker, J., O'Neill, M., & O'Farrell, M. (2008). Musical instrument identification using principal component analysis and multi-layered perceptrons. In *Audio, Language and Image Processing, 2008. ICALIP 2008. International Conference on* (pp. 643–648). IEEE.

Mack, R. (2013). Classic Tracks: Electric Light Orchestra "Don"t Bring Me Down'. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/sep13/articles/classic-tracks-0913.htm

Martin, G. (2011). Introduction to Sound Recording. Retrieved October 26, 2015, from http://www.tonmeister.ca/main/textbook/

Marui, A., & Martens, W. L. (2005). Timbre of nonlinear distortion effects: Perceptual attributes beyond sharpness?

Mason, M. (2010). Sample size and saturation in PhD studies using qualitative interviews. In *Forum Qualitative Sozialforschung/Forum: Qualitative Social Research* (Vol. 11).

Mauser, G. (2010). *Evolution of analog guitar FX*. Presented at the AES PNW. Retrieved from http://www.google.co.uk/url?sa=t&rct=j&q=evolution%20of%20guitar%20fx%20aes&source=web&cd=1&cad=rja&uact=8&ved=0CCEQFjAA&url=http%3A%2F%2Fwww.aes.org%2Fsections%2Fpnw%2Fppt%2Fgmauser%2Fevolution_of_analog_guitar_fx.ppsx&ei=5QowVNWsKIzXarW5gPAH&usg=AFQjCNHVgHv2FHo74hBAhTBwjEz94e2XJw

McAdams, S., Beauchamp, J. W., & Meneguzzi, S. (1999). Discrimination of musical instrument sounds resynthesized with simplified spectrotemporal parameters. *The Journal of the Acoustical Society of America*, *105*(2), 882–897.

McKinney, M. F., & Breebaart, J. (2003). Features for audio and music classification. In *ISMIR* (Vol. 3, pp. 151–158).

Mercury Recording Equipment. (n.d.). Mercury Recording Equipment Company: Mercury 66 Studio Limiting Amplifier. Retrieved July 16, 2016, from http://www.mercuryrecordingequipment.com/products/mercury-66/

Millett, P. (2004). *The Sound of Distortion*. Presented at the European Triode Festival, Germany. Retrieved from www.pmillett.com

Mitrović, D., Zeppelzauer, M., & Breiteneder, C. (2010). Features for content-based audio retrieval. *Advances in Computers*, *78*, 71–150.

Mokran, P. (2009). Secrets Of The Mix Engineers: Peter Mokran. Retrieved January 10, 2016, from https://www.soundonsound.com/sos/aug09/articles/it_0809.htm

Moore, A. (2012). All Buttons In: An investigation into the use of the 1176 FET compressor in popular music production. *Journal on the Art of Record Production*, (6).

Moore, A., Till, R., & Wakefield, J. (2016). An Investigation into the Sonic Signature of Three Classic Dynamic Range Compressors. In *Audio Engineering Society Convention 140*. Audio Engineering Society.

Moore, A., & Wakefield, J. (2017). An Investigation into the Relationship between the Subjective Descriptor Aggressive and the Universal Audio of the 1176 FET Compressor. In *Audio Engineering Society Convention 142*. Audio Engineering Society.

Moulder, A. (2000). ALAN MOULDER: Recording Nine Inch Nails & Smashing Pumpkins. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/may00/articles/alan.htm

Murphy, R. (2013). Analoguetube AT-101 - (Fairchild 670 recreation) review by Ronan Chris Murphy - YouTube. Retrieved October 28, 2015, from https://www.youtube.com/watch?v=dfC-NhHqp9Q

Narna, R. (n.d.). http://www.knifaudio.com/pdf/InterviewFairchild.pdf.

National Association of Broadcasters. (2007). *Engineering handbook* (10th ed). Amsterdam ; Boston: National Association of Broadcasters.

Neuendorf, K. A. (2002). *The content analysis guidebook*. Thousand Oaks, Calif: Sage Publications.

Neve, R. (1998). Audio Technology Magazine Issue 2. Retrieved from https://www.google.co.uk/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&uact=8&ved=0ahUKEwjzq-_RwpLKAhXHtRQKHdU4BlkQFggcMAA&url=https%3A%2F%2Fcommunity.klipsch.com%2Findex.php%3Fapp%3Dcore%26module%3Dattach%26section%3Dattach%26attach_id%3D85447&usg=AFQjCNEFuNy5sJQtBuMxYfgZzHSQtWyRsw&bvm=bv.110151844,d.d24

Neve, R. (2013, February 15). Rupert Neve - Pensado's Place #102 - YouTube. Retrieved August 23, 2014, from https://www.youtube.com/watch?v=k6OwY0XbqV4

Orchard, S. (2014). Watch Steve Orchard's interesting drum recording technique | Record Production Producer Interviews Episodes | Music Videos | Blip. Retrieved October 10, 2014, from http://blip.tv/recordproduction/steve-orchard-s-interesting-drum-recording-technique-5851230

Otala, M. (1972). Circuit design modifications for minimizing transient intermodulation distortion in audio amplifiers. *Journal of the Audio Engineering Society*, *20*(5), 396–399.

Pakarinen, J., & Yeh, D. T. (2009). A Review of Digital Techniques for Modeling Vacuum-Tube Guitar Amplifiers. *Computer Music Journal*, 85–100.

Pedersen, T. H., & Zacharov, N. (2015). The Development of a Sound Wheel for Reproduced Sound. In *Audio Engineering Society Convention 138*. Audio Engineering Society.

Peeters, G. (2004). {A large set of audio features for sound description (similarity and classification) in the CUIDADO project}.

Pestana, P. D., Ma, Z., Reiss, J. D., Barbosa, A., & Black, D. A. (2013). Spectral characteristics of popular commercial recordings 1950-2010. In *Audio Engineering Society Convention 135*. Audio Engineering Society.

Petri-Larmi, M., Otala, M., & Lammasniemi, J. (1978). Threshold of Audibility of Transient Intermodulation Distortion. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=2962

Pollen, A. A., Nowakowski, T. J., Shuga, J., Wang, X., Leyrat, A. A., Lui, J. H., … Chen, P. (2014). Low-coverage single-cell mRNA sequencing reveals cellular heterogeneity and activated signaling pathways in developing cerebral cortex. *Nature Biotechnology*, *32*(10), 1053–1058.

Preis, D. (1976). Linear distortion. *Journal of the Audio Engineering Society*, *24*(5), 346–367.

Puig, J. J. (2007). Secrets Of The Mix Engineers: JJ Puig. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/nov07/articles/insidetrack_1107.htm

Pure Mix. (2012). UA 1176 Compressor Settings Tutorial | puremix.net. Retrieved November 1, 2015, from http://www.puremix.net/video/how-to-use-an-1176-compressor.html

Pure Mix. (2014). Andrew Scheps Mixing Fly Rasta In The Box | puremix.net. Retrieved November 1, 2015, from http://www.puremix.net/video/andrew-scheps-mixing-fly-rasta-in-the-box.html

Putnam, L. L., Phillips, N., & Chapman, P. (2017). Managing Organizations: Current Issues. In pages 125-158. London: SAGE Publications Ltd. Retrieved from http://sk.sagepub.com/books/managing-organizations

R Core Team. (2016). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Reiss, J. (2011). *Design of a Dynamic Range Compressor*. Presented at the 131st AES Conference, New York. Retrieved from http://c4dm.eecs.qmul.ac.uk/audioengineering/compressors/documents/Reiss-Tutorialondynamicrangecompression.pdf

Reiss, J. D., & McPherson, A. P. (2014). *Audio effects: theory, implementation, and application*. Boca Raton: CRC Press.

Roads, C. (1996). *The computer music tutorial*. Cambridge, Mass: MIT Press.

Robins, M. (1961). *Don't Worry*. Columbia.

Robjohns, H. (2001). UNIVERSAL APPEAL. Retrieved November 8, 2015, from http://www.soundonsound.com/sos/jun01/articles/universal1176.htm

Robjohns, H. (2010). Analogue Warmth. Retrieved August 22, 2014, from http://www.soundonsound.com/sos/feb10/articles/analoguewarmth.htm

Robjohns, H. (2013). Q Is there a difference between clipping and limiting? Retrieved October 11, 2014, from http://www.soundonsound.com/sos/aug13/articles/qanda-0813-2.htm

Rocha, B., Bogaards, N., & Honingh, A. (2013). Segmentation and Timbre Similarity in Electronic Dance Music. *Proceedings of the SMC*.

Rutt, T. E. (1984). Vacuum Tube Triode Nonlinearity as Part of the Electric Guitar Sound. In *Audio Engineering Society Convention 76*. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=11614

Sandell, G. J., & Martens, W. L. (1995). Perceptual Evaluation of Principal-Component-Based Synthesis of Musical Timbres. *J. Audio Eng. Soc*, *43*(12), 1013–1028.

Schiffrin, D., Tannen, D., & Hamilton, H. E. (2008). *The Handbook of Discourse Analysis*. Wiley. Retrieved from https://books.google.co.uk/books?id=6RfarwereacC

Schneiderman, M., & Sarisky, M. (2009). Analysis of a Modified Boss DS-1 Distortion Pedal. In *Audio Engineering Society Convention 126*. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=14987

Schwenkler, B. (2011). DBX 165a help. Retrieved from http://repforums.prosoundweb.com/index.php?topic=1776.0

Scott, G. (2010). release time on vocal compressor - Gearslutz.com. Retrieved January 9, 2016, from https://www.gearslutz.com/board/so-much-gear-so-little-time/523071-release-time-vocal-compressor.html

Self, D. (2010). *Small signal audio design*. Oxford ; Burlington, MA: Focal Press.

Shanks, W. (2003a). Compression Obsession: The Amazing Release Character of the LA-2A. Retrieved October 28, 2015, from https://www.uaudio.com/webzine/2003/june/text/content4.html

Shanks, W. (2003b). UA WebZine "Compression Obsession" July 03 | Long Live the T4 Cell! Retrieved October 28, 2015, from https://www.uaudio.com/webzine/2003/july/text/content4.html

Shanks, W. (2011). 1176 and LA-2A Hardware Revision History - Blog - Universal Audio. Retrieved from http://www.uaudio.com/blog/1176-la2a-hardware-revision-history/

Skovenborg, E. (2014). Measures of Microdynamics. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=17464

Smith, A. (2011). Phil Ramone & Andy Smith: So Beautiful So What. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/oct11/articles/it-1011.htm

Smith, L. I. (2002). A tutorial on principal components analysis. *Cornell University, USA*, *51*, 52.

Sommer, B. (2006). Content analysis: Overview. Retrieved June 9, 2015, from http://psychology.ucdavis.edu/faculty_sites/sommerb/sommerdemo/content/

Sonnox. (2007). Dynamics Plug-In Manual. Retrieved from http://dload.sonnoxplugins.com/pub/plugins/manuals/SonnoxDynMan.pdf

Sound on Sound. (2009). Classic Compressors. Retrieved October 11, 2014, from http://www.soundonsound.com/sos/sep09/articles/classiccompressors.htm

Squire, R. (2010). On The Bench-Transformers, (55). Retrieved from http://www.proharmonic.com/articles/AT55%20ON%20THE%20BENCH.pdf

Squire, R. (2014). ON THE BENCH: MICROPHONE PREAMPS | AudioTechnology Magazine. Retrieved October 6, 2014, from http://www.audiotechnology.com.au/wp/index.php/on-the-bench-microphone-preamps/

Stavrou, M. P. (2003). *Mixing with your mind*. Mosman, NSW: Flux Research Pty Ltd.

Stowell, D., & Plumbley, M. D. (2010). Cross-associating unlabelled timbre distributions to create expressive musical mappings. In *WAPA* (pp. 28–35).

Strauss, A. L., & Corbin, J. M. (1991). *Basics of qualitative research: grounded theory procedures and techniques* (3. printing). Newbury Park, Calif.: Sage.

Suddaby, R. (2006). From the editors: What grounded theory is not. *Academy of Management Journal*, *49*(4), 633–642.

Talbot-Smith, M. (1999). *Audio engineer's reference book*. Oxford [England]; Boston: Focal Press.

Temme, S. (1992). Audio distortion measurements. *Application Note, Bruel & Kjar*.

That Corp. (2009). That Corp Design Note 00A Basic Compressor/Limiter Design. Retrieved from www.thatcorp.com/datashts/dn00A.pdf

The Pure Mix Tutorials. (2011). Building 1176 Compressor Clone: Interview - YouTube. Retrieved November 1, 2015, from https://www.youtube.com/watch?v=H1Q-xkX8oNI

Tingen, P. (1999). SPIKE STENT: The Work Of A Top-flight Mixer. Retrieved October 10, 2014, from http://www.soundonsound.com/sos/jan99/articles/spike366.htm

Toulson, R., Campbell, W., & Paterson, J. (2014). Evaluating harmonic and intermodulation distortion of mixed signals processed with dynamic range compression.

Treadwell, D. (2013). *Introducing Communication Research: Paths of Inquiry*. SAGE Publications. Retrieved from https://books.google.co.uk/books?id=Mb0gAQAAQBAJ

Tzanetakis, G., & Cook, P. (2002). *Manipulation, analysis and retrieval systems for audio signals*. Princeton University Princeton, NJ.

Universal Audio. (2000). LA2A Compresor Manual. Retrieved from https://media.uaudio.com/assetlibrary/l/a/la-2a_manual.pdf

Universal Audio. (2008). 1176 AE Edition Manual. Retrieved July 22, 2016, from http://media.uaudio.com/assetlibrary/1/1/1176ae-manual.pdf

Verfaille, V., Arfib, D., Keiler, F., von dem Knesebeck, A., & Zölzer, U. (2011). Adaptive Digital Audio Effects. In *DAFX: Digital Audio Effects* (pp. 321–391). John Wiley & Sons, Ltd. Retrieved from http://dx.doi.org/10.1002/9781119991298.ch9

Vintage King. (2014, January 30). Deconstructed: Chris Lord-Alge and Bob Clearmountain - YouTube. Retrieved August 23, 2014, from https://www.youtube.com/watch?v=tOB5TBaUKWc

Voishvillo, A. (2007). Measurements and Perception of Nonlinear Distortion—Comparing Numbers and Sound Quality. In *Audio Engineering Society Convention 123*. Retrieved from http://www.aes.org/e-lib/browse.cfm?elib=14232

Wallace, A. (2014). Inside Track: Linkin Park's The Hunting Party. Retrieved January 10, 2016, from http://www.soundonsound.com/sos/jul14/articles/it-07-14.htm

Warhead. (2011). Gearslutz Pro Audio Community - View Single Post - 1176 - release and attack times, can't figure it out? Retrieved July 16, 2016, from https://www.gearslutz.com/board/6948609-post20.html

Warner, T. (2009). Approaches to Analysing Recordings of Popular Music. In D. B. Scott (Ed.), *The Ashgate research companion to popular musicology* (pp. 131–147). Farnham, Surrey ; Burlington, VT: Ashgate.

Waves Audio. (2009a). CLA live at Mix LA Part 1/2 - YouTube. Retrieved November 1, 2015, from https://www.youtube.com/watch?v=7heuq2lV3h4

Waves Audio. (2009b). CLA live at Mix LA Part 2/2 - YouTube. Retrieved November 1, 2015, from https://www.youtube.com/watch?v=7ULCEMu9pHo

Wells, J. J. (2006). *Real-time spectral modelling of audio for creative sound transformation*. University of York.

White, P. (2000). Advanced Compression Techniques, Part 1. Retrieved October 28, 2015, from https://www.soundonsound.com/sos/dec00/articles/adcompression.htm

White, P. (2010). Distortion In The Studio. Retrieved September 28, 2014, from http://www.soundonsound.com/sos/apr10/articles/distortion.htm

Whitlock, B. (2008). *Handbook for sound engineers*. (G. Ballou, Ed.) (4th ed). Amsterdam ; Boston: Focal Press.

Wickham, H., & Chang, W. (2016). *devtools: Tools to Make Developing R Packages Easier*. Retrieved from https://CRAN.R-project.org/package=devtools

Winer, E. (2013). *The audio expert: everything you need to know about audio*. Waltham, MA: Focal Press.

Zagorski-Thomas, S. (2014). *The musicology of record production*. Cambridge, UK: Cambridge University Press.

Zeiner-Henriksen, H. T. (2010). *The" PoumTschak" pattern: correspondances between rhythm, sound, and movement in electronic dance music*. Department of Musicology, Faculty of Humanities, The University of Oslo.

Zwicky, C. (2011). Softube - EQ Before Compressor or Compressor Before EQ? Retrieved July 18, 2016, from http://www.softube.com/index.php?id=eq_before_compressor