The optimization of a modern day record press

by

Christopher Zaworski

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Masters of Science
in
Physics

Waterloo, Ontario, Canada, 2021

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The WarmTone is a fully automatic, modern-day vinyl record press manufactured by Viryl Technologies. The WarmTone has a wide array of sensors, parameters and settings that can be accessed via it's ADAPT Interface and the front panel. This thesis is a collaboration with Viryl Technologies through the Mitacs Accelerate Fellowship. These advancements in pressing technology enabled a large scale study of the press settings and parameters and how they related to the audio and visual quality of the finished vinyl record. For the purpose of this study a test record was developed and over 400 records were pressed and over 200 recorded. These records were pressed under varying conditions and settings in the press. The surface noise, number of pops and clicks, wow and flutter, stereo bleed and total harmonic distortion are measured from recordings of these records and the audio quality of the record is related back to parameters and settings in the press. Two different PVC compounds and WarmTone presses were used in the testing and overall the difference between the two plastics made the largest difference in the surface noise measured. New clicks were also added to the records as the side b stamper was damaged possibly during the loading of pressing two. When the surface noise is A-Weighted or CCIR/ARM weighted then the difference between press parameters and settings is negligible, press parameters do very little to affect the final audio quality. Unweighted surface noise measurements suggest that cooling in the moulds is the best way to reduce unweighted surface noise and pops and clicks on the records. Coherence in the surface noise is observed when the record audio is separated into "groove segments", segments of audio equal to the period of rotation of the record. This coherence, and other key press parameters affecting surface noise, such as the rotation speed of the extrusion screw, lead to the hypothesis that it is the size of the vinyl pellets and their residual elasticity that is the primary cause of surface noise and pops and clicks on the records.

Acknowledgements

I would like to thank my family for the seemingly infinite and immeasurable amount of support I received from them. My mother Ania, my father Richard, my brothers Eric and Kevin, my aunt, uncle and cousin Eva, Bill and Chelsea, my aunt and uncle Margie and Don, as well as my cousins Ashley, Kailey, Kervin and their children Makai and Bria. I'd like to also thank the non-human members of my family Dante, Hoss, Disney and my favourite organism on Planet Earth, my dog Abby the Labby. My best friends Mat, Lucas, Mike and Chris. My fellow graduate students Tiana, Julia, Bailey, Connor, Michael*2, Tony, Monica, honorary grad student Lindsay and all my other classmates and peers. I hope you find what you're looking for in the stars, particles and forces that make up our Universe the same as I found what I was looking for in the melted plastic and immense pressure of a vinyl record press.

I would like to thank God for giving me the opportunity to study and enjoy His Universe and the beauty designed for my appreciation. I would like to thank Him for placing all of my professors, teachers, family, counsellors and friends in my life at the exact moment I needed them. I would like to thank Him for the uncountable lessons I learned while I studied at the University of Waterloo. I am forever grateful for my life and the opportunities I found within it.

I would like to thank the University of Waterloo for playing such a pivotal role in my life and in the development of who I am as a scientist, musician, physicist and man. To all the professors, teaching assistants and instructors that have taught me over the 8 years I've attended the university– your lessons helped me realize a childhood and lifelong dream to be a physicist. I would like to thank James Hashmi, the CEO of Viryl Technologies for believing in me and trusting me with his vision and project. I hope my research benefits his company and I am able to enjoy the vinyl records that Viryl Technologies produces for many years to come, I truly love the industry and am so thankful for the opportunity and contributions I have made to it. I would like to thank my supervisor John Vanderkooy for being the kindest, warmest and most knowledgeable teacher I have ever had. I believe I will be his last graduate student and I do not hold that title lightly. I wish that I am able to carry on his legacy in my future research, and I hope I am able to instill the values that he instilled in me in my future students. The same values I see when I interact with his other previous students and friends. I am so thankful to consider myself as a part of their alumni, and will do my best for their sake as well as John's. I would like to thank Richard Mann, another member of the Audio Research Group, for first employing me in his research of audio. For teaching me the principles of audio in his course Computer Audio and introducing me to John. I owe him everything. I would like to thank Kevin Krauel, for his contributions to this project and for providing me with friendship and guidance through my graduate journey. I would like to thank my co-supervisor Jim Martin for teaching me how a physicist should behave. He probably isn't aware of this story, but during my first year of my undergraduate studies at the University of Waterloo I became interested in being a researcher and a professor. I decided I would attend the Colloquiums that the Physics Department puts on every week in order to involve myself in the affairs of the professors. Jim was the speaker at the first colloquium I attended, hearing the respect and attention the other physicists in the room gave to Jim and his research is what gave me the confidence to pursue a career in academia and research. The joy with which he taught me Electricity and Magnetism in my undergraduate degree is what gave me the confidence to accept him as my co-supervisor, and I am proud to have studied under him. I would like to thank Brian McNamara, the Chair of the Physics Department for being an invaluable mentor, leader and friend throughout my research. He shares my passion for music and physics, and speaking with him constantly reminded me of why I was at the university in the first place. A shared love and appreciation for the beauty of the Universe. A warm thank you to Chris Yakymchuk of the Earth and Environmental Sciences Department for lending the use of his microscope. I would like to thank Dmitry Pushin, he was the instructor of the very first course I took in my graduate studies, Statistical Mechanics. I was never very good at math and I'm still not, Dr. Pushin would stay late every Friday night with me and a few other students teaching us during his office hours, we would study well into the night. He was the professor who gave me the confidence to know that no matter how I felt about my abilities, I was capable and I would finish my degree and earn my Master's of Physics.

Lastly, myself and my family are forever grateful for the support I received from the University of Waterloo during my time as a student there. I owe the University of Waterloo much that I have in my professional and personal life. This thesis is as much a work of the University as it is my own and I am incredibly proud to have been a student of this wonderful institution.

Dedication

This thesis is dedicated to those who have stood beside me every step of the way. It was not easy, and I had my own unique challenges during this degree. My time at the University of Waterloo is very special to me and I hold everyone who was there for the journey very near and dear to my heart.

I hope that your involuntary dance of life is as beautiful as mine was. I hope that by understanding the physics by which I appreciate music, you too can observe this dance. For all music is beautiful, because it comes from the soul, and all physics is beautiful, because it comes from God.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis is an exploration into the physics of the vinyl pressing process. It's goals are to understand the process by which noise, pops, clicks and other unwanted audio artifacts are introduced into the audio of a vinyl record during pressing, develop a series of quantitative measurements of audio quality of a vinyl record and measure any quantitative difference in audio caused by parameter changes in the press. To achieve this a test record was developed that contained a series of tones, silence and sweeps that allows for the measurement of the records surface noise, wow and flutter, frequency response, stereo bleed, distortion and more. This test record was then pressed under varying conditions and then recorded. The resulting audio is then matched up with the press data at the time the record was pressed and thus any change in audio quality as a result of press parameter changes can be measured quantitatively.

Creating a high quality vinyl reproduction of sound is as much of an art as what gets pressed on the disc. Much of the knowledge and expertise required to create high quality reproductions was lost or forgotten when vinyl records fell out of favour with the public. In modern days, techniques and processes are being relearned and reinvented as sales of records reach highs not seen since the heyday of vinyl [1]. This thesis is part of a new resurgence in interest in vinyl records throughout the world, both commercially as vinyl record sales continue to boom [2], and academically as the industry demands innovation, research and training [3].

This project specifically studies the WarmTone press produced by Viryl Technologies, a company based in Toronto, Canada. The WarmTone is a modern machine, complete with full automation and cloud control, designed from the ground up to press records. The innovations in the WarmTone press and its ADAPT system make this research possible.

A special thank you to the Mitacs Accelerate Fellowship, a research grant that encourages the collaboration between industry and academic partners. The industry partner for this research is Viryl Technologies, the manufacturer of the WarmTone. Viryl Technologies in addition to helping fund the Mitacs Accelerate Fellowship also provided the Technics SL-1200GR turntable and Ortofon Blue cartridge, funded production of the test lacquers, stampers and provided the means and expertise to press the test records at their facilities in Toronto, Canada.

1.1 The WarmTone Press

The WarmTone is a fully-automatic modern record press equipped with Viryl's ADAPT Platform, a customized cloud based interface that allows full access to press's data through a smart device. Two individual WarmTone Presses were used in this thesis at Viryl's Facilities in Toronto, Canada. In addition to the ADAPT data, Viryl Engineers are able to pull data from the various press sensors that are laid throughout these test machines to give insight into the pressing process and conditions. This sensor data can later be matched up with the records beings pressed and recorded as described in Chapter 3.

The WarmTone also features detailed control over the main components and operating parameters. The WarmTone record press is semi-modular, with an extruder, mould, press and trimmer making up the WarmTone. Stampers are cleaned and loaded into the moulds of the WarmTone, the PVC compound starts as vinyl "pellets" that are loaded into the extruder and melted into a vinyl "puck". This puck is then moved into the press and pressed into a record where it is trimmed by the trimmer. The trimmer consists of an automatic arm that lifts records from the press and trims the excess plastic or "flash" and creates a clean edge for the vinyl record. For the most part the trimmer and its associated parameters, settings and sensors are ignored in this thesis as it has little to do with the final audio quality of the record. The extruder, mould and press settings do have an effect on the audio quality measured, as has been noticed by press operators and will be the main focus of research and results in this thesis. Various parameter changes were chosen from the available parameters and settings for the extruder, mould and press and minimum and maximum values were chosen for each parameter.

Overall two pressings were done. The first pressing held parameters and settings constant in the WarmTone and produced over 100 records. This was done to determine how much the audio quality varied inside of the press while the press was held at a steady state. In the second pressing done, the goal was to determine which of these parameters and to what extent these parameters impact the final audio quality. So instead the parameters were varied between maximum and minimum values and at minimum five records were pressed at each of these parameter changes. The results of these tests are given in Chapter 4.

1.2 Research Motivations

With the WarmTone press, press operators have access to various controls such as extruder temperatures, steam heating and cooling cycle times, press pressure and more. These controls are set in order to maximize both the audio and visual quality of the disc. Records are spot checked as the press runs, listened to and inspected visually to ensure that they meet quality standards and that there are no defects. If defects are common or issues with the audio are heard then the WarmTone settings are adjusted and records listened to and inspected until these issues are resolved.

This is a very labour intensive and time consuming process that requires a trained ear. Minute adjustments in the WarmTone's settings are believed to cause changes in the audio such as increasing or decreasing the surface noise or number of clicks. Operators can spend time dialing in these settings in order to reduce noise, warp and other unwanted properties in the records produced. However much skepticism surrounds how much influence the press has over the final audio quality, compared to other steps in the process, as well as whether or not these minor adjustments do significantly influence the audio of the resulting record.

This project's goal is to quantify these parameter changes with changes in the audio. A custom test record was designed and pressed under varying conditions in the press and then recorded. These recordings were analyzed and the settings and sensor measurements in the press were matched with the audio measurements from the recordings. In this way, the audio quality of the recordings measured as a function of parameters and sensors in the WarmTone press.

It's important to note the climate that this research is being conducted in. There has been a resurgence in vinyl record production and sales in the past few years. Crawford reports the 12th consecutive year of vinyl sales growth in 2017 [2], even amidst the pandemic the RIAA reports that in 2020 revenues from vinyl records were larger than from CDs for the first time since 1986. [4]. The RTA School of Media reintroduced vinyl theory and production practices into it's courses [3]. It's clear that the vinyl resurgence is here to stay in some capacity. This research stands at the forefront of improving and understanding the vinyl manufacturing process.

1.3 Manufacturing Process

Before records can be pressed by the WarmTone "stampers" need to be produced. Stampers are negatives of the vinyl record, they have "hills" instead of "grooves", that fit into the tooling of the moulds of the WarmTone press and are what press the grooves into the vinyl puck. In the days of Eargle and Ruda, music would be recorded to a master tape before being cut to a lacquer, in the modern era the vast majority of new vinyl projects are produced digitally using Pro Tools and other DAWs. A digital file is turned into a physical vinyl record through a multi-stage process. Special care has to be taken when creating a file for vinyl playback, as the mastering engineer must be made aware of and correct for the physical limitations of vinyl. These include constraints on the stereo separation, dynamic range and maximum time trade offs, and more including a loss in quality in the inner grooves.[3] The manufacturing process for a vinyl record is multi-staged, with a lot of room for error to be made and for defects in the process to be introduced. At first, the music on the digital file or master tape is cut to the master lacquer by a lathe. Afterward the lacquer is goes through two rounds of electroplating, the first to produce a negative "father" from the lacquer and then a positive "mother" from the father. Stampers are then electroplated from the mother. [5] All this happens before the pressing of vinyl by the WarmTone begins.

There are various stages at which manufacturing defects can be introduced into the pressing process. First a heated stylus cuts the mastered digital file into the lacquer. A lacquer is an aluminum disc with a lacquer coating that the grooves are cut into. Measurements of the audio of test lacquers sent for evaluation are given in Chapter 5. Consultations with our mastering engineer and our experiences with the manufacturing of the lacquers for the test records has corroborated Ruda's claims that the lacquer cutting, including the manufacture of the lacquers themselves, is still very much an art. Much like the pressing itself, the process by which a lacquer is made is still as much of an art as a science.[5]

After a lacquer is cut it is cleaned and sprayed in a chemical bath of silver. It is possible for noise to be added at this stage through residual oils, however in this study it was found that pressed records can indeed be quieter, suggesting that noise isn't simply added at each stage of the process, rather that there is a complicated relationship between the audio cut to the lacquer, the methods by which the stamper are produced and the plastics used to create the vinyl record. All play crucial roles in determining the resulting sound quality. After silvering the lacquer is plated with nickel to create the "father", which is a negative metal mould of the lacquer and can be used as a stamper for pressing records. In this thesis and in most cases a two-step process is used where a "mother", typically a copper disc that can be played on turntables, is made from the "father" and the stampers are then taken from the "mother" as opposed to from the lacquer directly. [1]

During the process by which the lacquer is cleaned and made electrically conductive with a chemical bath of silver, Ruda describes a method by which oils on the surface of the lacquer can produce audible defects during silvering. In this study it was found that pressed records can indeed be quieter than a test lacquer, suggesting that while noise can be added during the silvering, plating and stamper production stages of the production of a record it appears as though the nature of noise on a lacquer is simply different compared to the nature of noise on a vinyl record. Indeed these differences in the lacquer and vinyl spectra are also seen in the spectra produced by Eargle [6]. Of course imperfections at this stage should be avoided. Unfortunately there is no way for us to tell with the current dataset or measurement techniques how much noise is added or removed during the plating process.

Immediately after silvering the nickel plating process begins, plating can produce other audio defects such as echo and ticks. After pre-plating and plating is completed at a thickness of 0.015 to 0.035 inch the negative silver-faced master is "peeled" from the lacquer. [5] Typically the lacquer is then discarded. [1] While reproduction from the silver "father" is possible, the silver negative is very delicate and not built to withstand the manufacturing process and more commonly, as was done with this test pressing, the two step process is used and a second round of plating begins in order to produce a "mother". A stamper averages around 2000 presses for its life. For this thesis four stampers in total were created however only two have been used, 419 records total were pressed. The metal "mother" disc has the potential to cause a molding defect called a fractured groove. Careful polishing removes burrs at the groove edges that are created during the lacquer cutting. After polishing and cleaning stampers are then created from plating the "mother", once the stamper is made from the plating it is prepared for interfacing with the press tooling and moulds. [5]

Inside the press, PVC compound is loaded into a hopper and melted through the WarmTone's extruder in a multi-stage heating process. This vinyl is extruded into a "puck"-sized piece of vinyl in the premoulds before the puck is then loaded into the press alongside the labels and the press compresses the puck into a vinyl record, with the labels embedding themselves into the vinyl, no adhesive required. As the press is closed, steam is first pumped through the moulds heating the vinyl further before cold water is then pumped through the moulds cooling the record. The press opens and the record is picked up by the trimmer arm, moved to the trimmer and the excess flash around the ends of the record are cut off. The record is then stacked and cooling plates are placed on top every few records to ensure the records cool properly. To press one record takes about 30-45 seconds depending on the timing involved, the WarmTone press is fully automatic, only requiring the press operator to supervise and clear any errors or junk records.

1.4 Measurements of Audio Quality

As this is one of the only studies of it's kind into the audio quality of vinyl records manufactured, the standards by which audio quality is to be measured and judged needed to be developed. In order to facilitate these measurements a custom test record was developed, a test lacquer was sent for evaluation, four stampers were produced albeit only two ended up being used and over 400 test records were pressed. Described in Chapter 2 are how each of the tracks on the test records are used to facilitate measurements of surface noise RMS levels, wow and flutter, frequency response, stereo separation and total harmonic distortion. Additionally pops and clicks are detected using an adapted MATLAB version of Audacity's Click Removal tool [7], the details of which are described in Chapter 6. These measurements and the normalization procedure ensure that the measurements are normalized within our data set, and so records can be compared against one another. Additionally, our normalization tone mirrors the NAB Standard of 1 kHz 7 cm/s peak, so in theory the data in this thesis can be compared to other studies that use the same normalization. Eargle uses the same NAB standard for his noise and distortion measurements in his study however a different methodology, and so our results can be compared to his somewhat as is done in Chapter 7.

The NAB Recording and Reproducing Standards Committee produced the Standards for Disc Recording and Reproducing to describe standards for wow and flutter, disc specifications and dimensions and some rumble and high frequency noise standards that use reference tones as the basis for their measurement [8]. In this thesis surface noise measurements are done with reference to a 1 kHz 7 cm/s peak normalization tone, in a normalization procedure described in Chapter 2. RMS levels of surface noise are used and A or Dolby CCIR/ARM Weighting is used to represent the high frequency content and suppress the rumble noise. The Unweighted measurement is considered the measurement of rumble as both A and CCIR/ARM weighted both reduce rumble substantially as discussed in Chapter 5. [9]

Literature review found two previous studies on the AES E-Library done in 1977 by Ruda [5] and 1969 by Eargle [6] which had similar research goals and methods. Whereas those studies looked at the manufacturing process as a whole, including the lacquer cutting and stamper creation, this thesis primarily focuses on the role that the press plays in the manufacturing process. The multi-stage manufacturing process, from lacquer, to father, to mother, to stamper and finally to record has many avenues by which noise and other defects can be introduced. The question asked in this thesis is how much control does the press operator have over the final audio quality, or is the press operator simply at the mercy of the mastering and engineering that produced the stamper. Eargle looked more intently at noise introduced at each stage of the process. He measured the metal mother's spectra additionally. [6]

For the most part however, the measurements in this thesis are specific to this thesis and there are no comparable studies testing the pressing process. This kind of test into the pressing and audio of vinyl records is uncharted territory. The measurements developed for this test record are specific to this test record, however where possible standardized tones and normalization methods were used, as will be described in Chapter 2. The two main areas of focus for audio quality is the surface noise of the record and the number of pops and clicks. Additional measurements of quality are given such as the total harmonic distortion, wow and flutter, and stereo separation alongside visual defects.

Chapter 2

Design of the test record and measurement techniques

A key component of this thesis is to establish a quantitative basis on which the audio quality of test records are to be measured and judged. These measurements of audio quality must accurately reflect how a listener might judge a record. This mirrors the process by which press operators judge the quality of records as they are made. As records are pressed in a commercial run, some sample records are chosen and listened to in a controlled listening environment on professional grade equipment by the press operators. The trained professional then makes a determination to judge whether the press is producing "good" or "bad" sounding records and the press is adjusted accordingly until adequate audio and visual quality is achieved. Such a process can be very time consuming and expensive.

For the purpose of this thesis, a test record was developed in order to allow for a much more quantitative approach to measuring sound quality to be taken. As parameters are being varied in the press real numbers needed to be attached to the audio measurements in order to judge the change in audio due to the parameter change in the press. For this purpose, a custom test record was developed, cut to a lacquer and four stampers were produced in order to create test records. Only two stampers were used for the purpose of this thesis, one for the a-side of the record loaded into the top mould of the WarmTone and one for the b-side of the record loaded into the bottom mould of the WarmTome. The test records contain a series of tones, silence and sweeps that facilitate measurements of the surface noise RMS levels, coherence, wow and flutter, frequency response, stereo bleed, and total harmonic distortion through the development of custom MATLAB scripts. The three main scripts SeperateTracks, RecordProcess and SensorProcess take as inputs wav file recordings of the entire test record and the sensor data from the WarmTone and output a table of audio and sensor measurements on a per track per record basis, so all the audio measurements of each individual track on each test record is recorded in a table, and then matched up with the corresponding sensor data from the WarmTone. This large dataset is further processed into plots and tables and then conclusions can be drawn from them.

This Chapter details the design of the test record, the signals that it contains and how the audio measurements are taken using those signals. It also details the MATLAB scripts that deal with the recordings as well as the press sensor data in order to produce the final data set.

2.1 Tracks on the test record

There are in total 25 tracks on the test record not including the leadin and leadout. The signals found on the record are tones at 100 Hz, 1000 Hz, 3150 Hz and 10 000 Hz, a logarithmic sweep from 16-16000 Hz and deliberate sections of silence. With the exception of the 3150 Hz tone, these signals are repeated throughout the record cut in various ways. First is as a traditional monophonic signal in both stereo channels, where the stylus tracks laterally. Next as a vertical cut, such that the stylus moves vertically as it tracks the signal– as in a Hill and Dale recording. Lastly the tones may be cut to either the left or the right channel. This set of tones and signals is repeated once on the record, with a long stretch of silence called the transition in between. The complete track listing is given in Table 2.1.

The extended periods of "silence" found in the quiet and transition tracks on the record are in fact not silent at all, they contain over to -30 dB of unweighted RMS surface noise. The majority of this surface noise is in the very low, inaudible frequencies peaking at around 12 Hz and is classified as rumble noise. This region of low, inaudible rumble noise also contains the turntable arm resonance, which is right around 20 Hz for the recording setup. It's important to note that this noise is inaudible and to represent this the A or CCIR/ARM weighted measurements will be used. When the surface noise is weighted with A or CCIR/ARM weighting the surface noise RMS is much more reasonable at around -55 dB. In reality these surface noise levels should be interpreted even lower, as our data is normalized to 7 cm/s, whereas a typical maximum signal found on a record is around 40 cm/s, as is the case on the Shure TTR-103, with Alexandrovich reporting unavoidable distortion occurring at 50 cm/s. [10] As such, -15.14 dB should be subtracted from our noise measurements, to normalize to 40 cm/s, the true peak signal found on a record, as opposed to the 7 cm/s peak signal found on our test records. So the true signal to noise ratio of our records are around -70 dB, -15.14 dB better than reported in this thesis due to the normalization. The details of these surface noise measurements and the record's spectrum can be found in Chapter 5. The quiet tracks enable a lengthy measurement of this surface noise which allows the best comparison of RMS noise levels between the different test records. Typically if a press operator wanted to judge a records surface noise they would only have the lead in and lead out grooves and the brief period of silence between tracks to clearly listen to the surface noise without any music playing. There are overall two quiet tracks and one transition track on the record that add up to over 139 seconds of silence on the test record. This enables a lengthy measurement of the surface noise at three different groove radii.

This silence can also be used to measure the record coherence. These coherence measurements are discussed in Chapter 5. This is especially interesting for the transition track. The transition track is somewhat unique on the record, as it is the section of silence that connects the two sets of signals, it contains grooves that are spaced close together as they are in the quiet tracks and grooves that are spaced far apart in order to enable the stylus to travel a fairly far distance on the record surface. A photograph of the test record is included in Figure 2.1. The transition track is clearly visible as the large section of silence that connects the two sets of signals. An interesting conclusion found in this and discussed in Chapter 6, the transition track by far contains the most amount of clicks of all the tracks on the record, by a factor of 10.

The tones included on the record are a 1 kHz normalization tone at 5 cm/s RMS, a tone at 10 kHz, 100 Hz and 3150 Hz. There is a sweep of frequencies from 20 Hz - 16 kHz. The 1 kHz tone along with the sweep is also repeated in the left channel only, the right channel only and as a vertical cut, "Hill and Dale" recording. Hill and Dale recordings modulate the stylus vertically rather than laterally for a signal in both channels, it is a mainly outdated recording technique, it was noted obsolete in 1964 by the NAB. [8] The tones cut to a single channel are meant for measuring stereo separation or stereo bleed. The various other tones are mainly used for distortion measurements. The 3150 Hz is used for Wow and Flutter however in theory any of the tones can be used for wow and flutter.

Lastly the sweep from 20 Hz - 16 kHz enables measurements of the test record's frequency response, a measurement that isn't to be confused with the record's noise spectra. The frequency response discussed in Section 2.4, is a measurement of the records ability to reproduce signals of certain frequencies accurately, whereas the noise spectra are a measurement of the frequencies contained in the record's background surface noise.

Track name	Description	Duration (seconds)	Radius
leadin	The quiet grooves after the needle drop.	–	15 cm
1kHz	1 kHz at 7 cm/s peak velocity 5 cm/s RMS.	60	14.5 cm
10kHz	A 10 kHz tone at -20 dB.	30	14.0 cm
100Hz	A 100 Hz tone at -20 dB.	31	13.8 cm
sweep	20 Hz - 16 kHz.	36	13.5 cm
quiet	A deliberate section of silence.	21	13.3 cm
3150Hz	A 3150 Hz tone.	66	13.0 cm
1kHzL	A 1 kHz tone in the left channel.	20	12.7 cm
sweepL	20 Hz - 16 kHz in the left channel.	37	12.5 cm
1kHzR	A 1 kHz tone only in the left channel.	19	12.3 cm
sweepR	20 Hz - 16 kHz in the right channel.	37	12.1 cm
1kHzV	A 1 kHz tone vertically cut.	19	11.8 cm
sweepV	20 Hz - 16 kHz vertically cut.	37	11.6 cm
transition	A deliberate section of silence.	97	11.3 cm
1kHz2	Identical to track 1kHz.	60	9.2 cm
10kHz2	Identical to track 10kHz.	30	8.8 cm
100Hz2	Identical to track 100Hz.	31	8.6 cm
sweep2	Identical to track sweep.	36	8.4 cm
quiet2	Identical to track quiet.	21	8.2 cm
3150Hz	Identical to track 3150 Hz.	66	8.0 cm
1kHzL2	Identical to track 1 kHzL.	20	7.6 cm
sweepL2	Identical to track sweepL.	37	7.3 cm
1kHzR2	Identical to track 1kHzR.	19	7.0 cm
sweepR2	Identical to track sweepR.	37	6.8 cm
1kHzV2	Identical to track 1kHzV.	19	6.5 cm
sweepV2	Identical to track sweepV.	37	6.3 cm
leadout	The quiet and lockout grooves.	–	5.2 cm

Table 2.1: The tracks found on the test records. L and R designate a signal only found in the left or right channel respectively, V designates a signal that is vertically cut.

2.2 Normalization signal and procedure

One of the most crucial parts in the design of the test record is how the audio levels of the different recordings will be normalized. When recording, the gain on the Focusrite interface is a very key parameter as it sets the signal level and relative "volume" of the recording.

Figure 2.1: An image of the test record, note the two sets of signals and the large space of silence in between them. The transition track is visible as the spaced out grooves in the middle of the record connecting the two sets of signals.

Since records are being compared to one another and assessed, any measurements that are sensitive to the volume level will be influenced by these gain controls. To bypass these controls a method to read the underlying signal level regardless of the gain on our recording system was developed.

To accomplish this a reference signal of known velocity is cut to the disc. The level of this signal is measured during playback, its amplitude in the digital domain directly corresponds to the velocity of the stylus. Since the underlying velocity and the corresponding digital level is known for this signal, the underlying velocity of an unknown signal can be determined from its digital level using this reference signal. The NAB defines a reference tone of 1 kHz tone at 5 cm/s RMS (7 cm/s peak) for a length of 20 seconds. The reference tone used in this thesis is a 1 kHz tone at 5 cm/s RMS velocity for a duration of 60 seconds– however the duration is ultimately unimportant as long as it is sufficient to determine audio levels. All data in this thesis is normalized such that this 5 cm/s RMS velocity (7 cm/s peak) reads an RMS level of 0 dB FS. [8]

This signal was verified through a microscope by measuring the maximum slope along the curve and multiplying by the tangential velocity at that radius on the record. The peak velocity was measured to be 6.8 ± 0.1 cm/s. While the true value doesn't lie within the error bounds, this is considered an acceptable error and the signal is considered verified at 7 cm/s peak. To verify this velocity, an image of the reference tone was taken with a microscope. The radial and the tangential distance that the stylus must travel was determined using the geometry of the grooves. See Figure 2.2, these distances can then be used to determine the approximate distance that the stylus covers as the hypotenuse of this triangle in the following calculation.

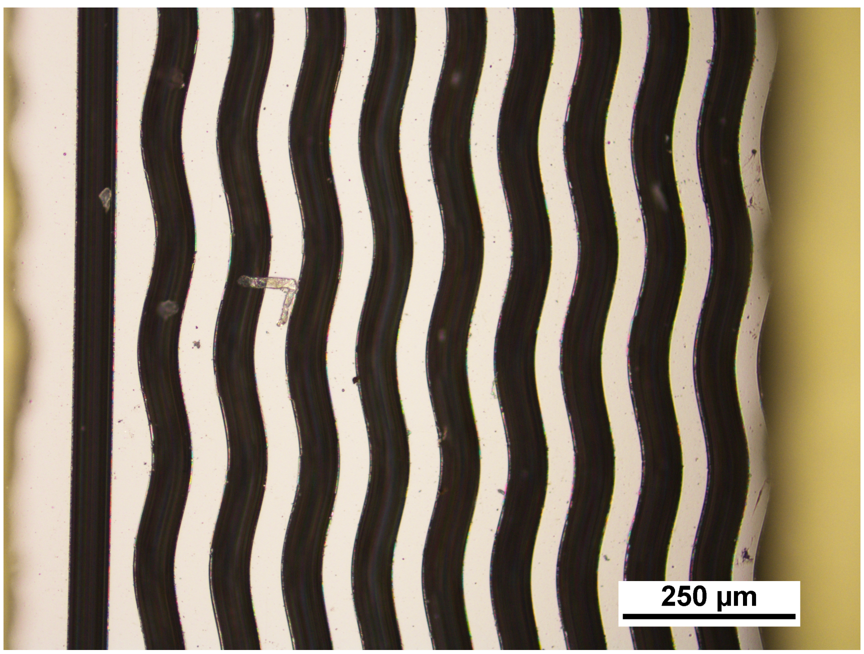

Figure 2.2: The 1 kHz 5 cm/s RMS tone.

Knowing the radius at which these grooves appear on the record and the speed of the rotation of the turntable we are able to determine the velocity of the stylus as it travels this distance. The angular speed of the turntable ω, can be quickly calculated from the rotation speed of the platter in rotations per minute (rpm).

$$\omega = 2\pi * 33\frac{1}{3}RPM$$

$$\omega = 2\pi * 33.\overline{33}\frac{1}{min} * \frac{1}{60}\frac{min}{sec}$$
$$\omega \approx 3.4906\frac{rad}{sec}$$

Since the record rotates at a constant rate it traverses a tangential distance d_t in time Δt. The tangential velocity is simply the angular speed of turntable ω multiplied by the radius of the track R. The distance from the stylus position to the centre of the record.

$$\Delta t = \frac{d_t}{v_t}; v_t = \omega R$$

For a laterally cut signal, the voltage from the cartridge is proportional to the radial velocity v_r which itself is proportional to the radial distance d_r.

$$v_r = \frac{d_r}{\Delta t} = \frac{d_r}{d_t}\omega R$$

Using these calculations and a measurement made of the groove geometry from a microscope image, the velocity of the 1 kHz 5 cm/s RMS signal can be verified.

The Image 2.3 was taken at a radius of 8.975 ± 0.025 cm, the distance d_t was measured at 157.1 ± 0.05 cm and the distance $d_r = 34.1 \pm 0.05$ cm. From these measurements, the velocity of the stylus can be calculated as:

$$v_r = \frac{d_r}{d_t}\omega R = \frac{34.1}{157.1} * 3.4906 * 8.975 = 6.8 \pm 0.075\frac{cm}{sec}$$

So the final calculated peak velocity is 6.8 ± 0.075 $\frac{cm}{sec}$ and the final calculated RMS velocity 4.8 ± 0.05 $\frac{cm}{sec}$. These values are considered acceptable and fall within what we would expect to see for the 1 kHz 5 cm/s reference tone, albeit slightly off.

To verify this measurement process a second test record, the Shure TTR-103 was used. The Shure TTR-103 test record is similar to the test record developed for this research, the TTR-103 is a 45 rpm Trackability Test Record. It is meant to test how well a turntable and cartridge can track a powerful signal produced by the grooves of a record. It's meant for a consumer to test their listening equipment to ensure it meets the trackability standard, to demonstrate that the listener's equipment is able to reproduce the signals found on most if not all commercial vinyl records. For the TTR-103, the peak velocity is 40 cm/sec which is produced by two signals one at 1000 Hz and one at 1500 Hz. [11] This is assumed to be

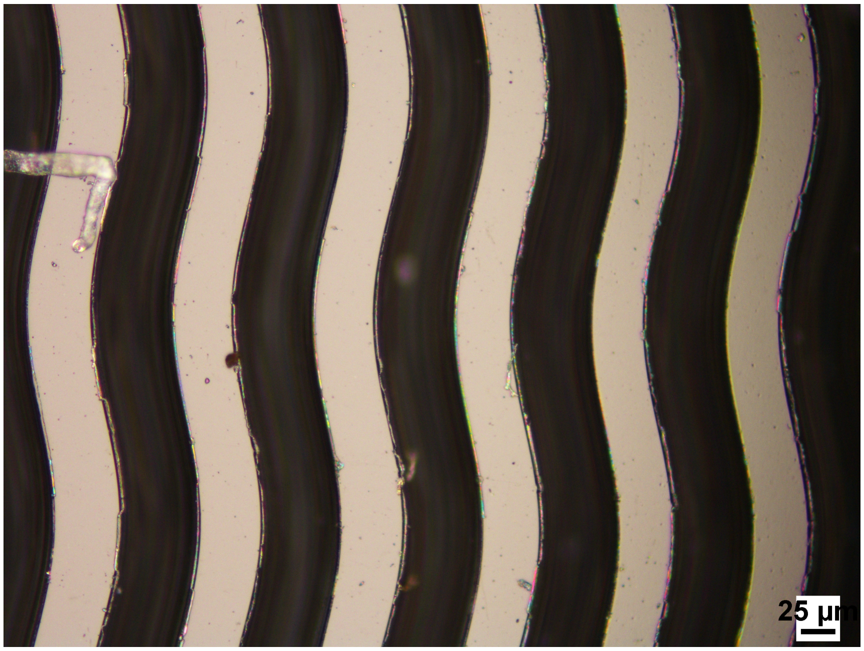

Figure 2.3: The 1 kHz 5 cm/s RMS tone taken at a radius of 8.975 cm.

the maximum stylus velocity for a traditional turntable and indeed 50 cm/s is reported to be too high a stylus velocity to be reproduced without distortion by Alexandrovich. [10]

Repeating the calculation for peak velocity from the measurement and image of the Shure TTR-103 Trackability Test Record we can demonstrate how accurate our calculation of the normalization velocity is.

$$\omega = 2\pi * 45RPM$$
$$\omega = 2\pi * 45\frac{1}{min} * \frac{1}{60}\frac{min}{sec}$$
$$\omega \approx 4.7124\frac{rad}{sec}$$
$$\Delta t = \frac{d_t}{v_t}; v_t = \omega R$$
$$v_r = \frac{d_r}{\Delta t} = \frac{d_r}{d_t}\omega R$$
$$v_r = \frac{d_r}{d_t}\omega * R = \frac{54.0}{59.4} * 4.7124 * 9.00 = 38.6 \pm 0.1\frac{cm}{sec}$$

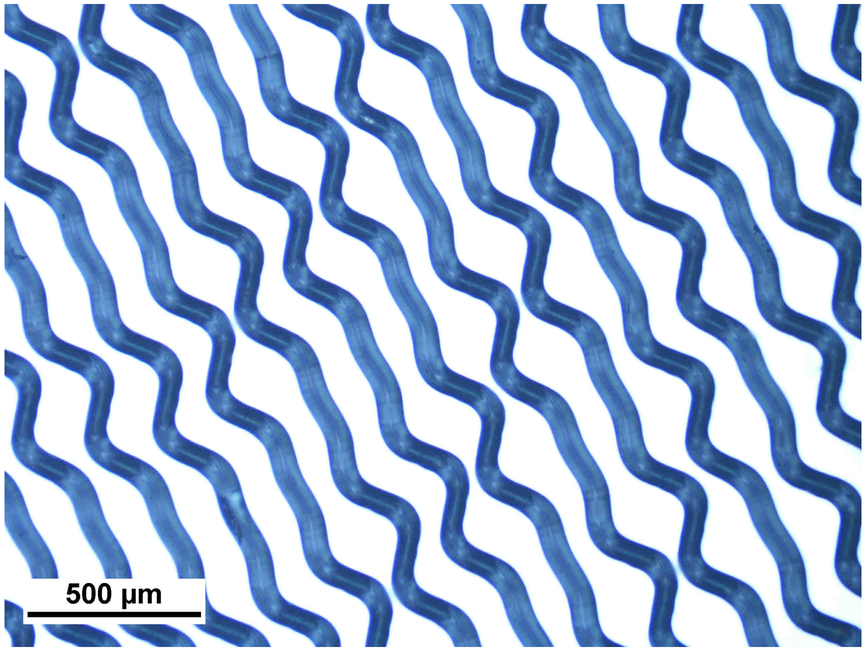

Figure 2.4: An image of the Shure TTR-103 test record taken at a radius of 9 cm.

Once again the calculated velocity from the measurement of the groove geometry is slightly less than the specified velocity, however it isn't off by much. This measurement of the Shure TTR-103 test record provides additional confidence in the measurement of the 1 kHz 5 cm/s RMS 7 cm/s peak normalization signal on the test record. While both measurements are slightly less than the anticipated velocity, and the error bounds do not necessarily capture the true value the result is close and reasonable enough to conclude that indeed the normalization tone on the test records corresponds to the NAB Standard.

Normalizing using this signal of known velocity ensures that RMS levels and other measurements of noise have physical meaning rather than simply being a characteristic level for our measurement system. The 1 kHz normalization tone is tied with the other 1 kHz tracks, those cut vertically or in a single channel, as the loudest track on the record. It's important to note that while the data in this study is normalized to 5 cm/s RMS, the true peak velocity of a stylus tracking a vinyl groove is assumed to be 40 cm/s which is the peak velocity of the Shure TTR-103. As such it is also totally reasonable to normalize the measurements in this research to a peak value of 40 cm/s, however eventually the researchers settled on using the 5 cm/s RMS normalization compared to the 40 cm/s peak normalization.

There were two normalization procedures considered. The first required the use of a Fourier transform to obtain the amplitude of the 1 kHz track, and use that amplitude as the normalization coefficient. For the second method the normalization coefficient is simply equal to the RMS level of the 1 kHz signal. The latter method ended up being much more reliable and accurate than the more complex measurement of the 1 kHz peak in the spectrum, the reasons for which are interesting in and of themselves and are discussed here.

First, a portion of the 1 kHz signal is selected. In order to get an accurate measure of amplitude from the spectrum, the set of samples are windowed with a flattop window. The flattop window used is shown in Figure 2.5. Alongside this window, the signal is divided by a correction factor of 0.2155774 in order to correct for the drop in amplitude caused by windowing. The flat top window was chosen as it preserves the signal amplitude well during a Fourier transform. After windowing the data is Fourier transformed and the amplitude of the 1 kHz peak in the spectrum is measured.

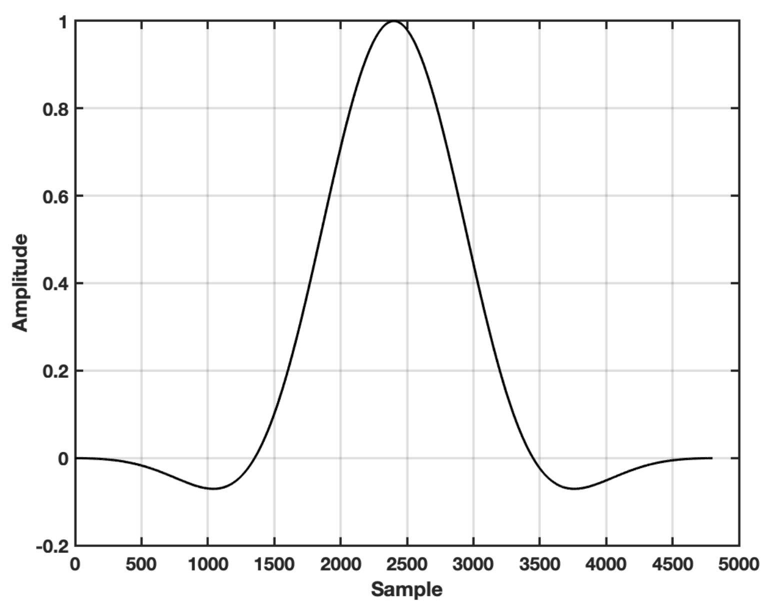

Figure 2.5: The flattop window used on the 1 kHz track during normalization.

The peak of the windowed spectrum corresponds to the peak value of the 1 kHz tone. The normalization procedure is based around the RMS velocity and not the peak velocity, since the incoming signal is a sine wave, the peak of the spectrum is divided by $\sqrt{2}$ in order to get the RMS amplitude and define the normalization coefficient. Each channel is given its own normalization coefficient, as the gain is independently set on either channel. The data is then divided by this normalization coefficient, such that the 5 cm/s RMS signal has an RMS value of 0 dB.

This normalization procedure should result in an accurate measurement of the cartridge's peak velocity and thus its RMS velocity, however the procedure reliably produces a normalization coeffecient that is too small and overcorrects the amplitude. It was discovered that this is due to the normalization signal's wow and flutter.

Wow and flutter is a modulation of pitch that most commonly occurs on a vinyl record due to the centre hole being offset from the geometrical centre of the record. This in turn causes the speed of the record underneath the cartridge to vary from slow to fast as the radius of the cartridge moves inwards and outwards. This speeding up and slowing down of the record causes an oscillation in pitch called wow. Flutter has a similar effect on pitch, however is typically a feature of magnetic tape playback.

This variation in pitch can be seen as a widening of the peak in the spectrum. As the signal is no longer a pure tone of 1 kHz, but instead varies from say 0.990 - 1.010 kHz. The smeering due to wow causes problems when measuring the amplitude of the sinusoidal signal. Instead of all the total power of the signal being captured in the 1 kHz bin, as it would be if the normalization tone's frequency were a constant 1 kHz, it is spread out across multiple bins around 1 kHz, resulting in a spectral peak that is lower than the true power in the signal. This lower peak results in a larger normalization factor and an over-correction of the amplitude of the signal.

In order to measure the amplitude of the normalization tone accurately using its spectrum, the power in this region around 1 kHz needs to be summed together somehow. This is a very difficult process to reliably automate as the Wow, and thus the width of the frequency peak and number of bins used in the sum, varies for each record, so the much simpler method of normalizing to the RMS level of the normalization signal is used instead.

2.3 Wow and flutter measurement

If the centre mounting hole is not precisely in the centre of the record, then the listener will hear a rising and lowering of pitch when the record is played. This change in pitch is called wow and flutter. Traditionally, flutter is an effect found in tape playback. In this thesis, we will be focusing on the wow measurement that is most often due to this centre hole offset. Mullen defines wow as below 10 Hz and flutter as above 10 Hz [12]. Wow and flutter is quite audible and can be a large factor in a listener's determination of audio quality in records and tapes and so various standards exist surrounding its measurement. There are various ways to define and measure wow. This section describes the methodology and definition of wow used in this thesis.

The first step in the process of measuring wow is to obtain a rough idea of the frequency of the signal. This test frequency is determined decimated down from 96000 Hz to 9600 Hz, a factor of 10. After decimation, the Hilbert transform of the signal is taken, using the MATLAB function "hilbert". The phase of the transformed signal is obtained by using the MATLAB functions "angle" and "unwrap". From this array of the phase of the Hilbert transformed signal, we take the derivative of the phase to obtain an array that contains the instantaneous frequency of the signal.

After obtaining this array it is filtered by a 2nd order Butterworth lowpass filter generated by the "butter" command and the frequency array is filtered in both forward and reverse time using the "filtfilt" command. Finally, DC is removed from this array by subtracting the average value. Now that we have obtained the instantaneous frequency of the signal, what is left to measure the wow is to take the RMS response of this instantaneous frequency.

This RMS response can be both weighted and unweighted. The weighting is given in Figure 2.6. It is used to give a rough estimate of the audibility of the wow, however for accurate measurements of the centre hole offset, the unweighted wow should be used. Both the weighted and unweighted RMS measurements are included. The peak variation in frequency is used as the measurement for wow.

2.4 Log sweep

A sweep from 20 Hz to 16 kHz is included on the record intended to measure the frequency response of the records. It is hard to quantify a measurement of frequency response especially on the large scale for test records, it is moreso a characteristic that needs to be examined on a per record basis. The spectrum of the digital file used as the basis for the sweep is provided in Figure 2.7. This sweep is made up of a sinusoidal tone for which the logarithm of the instantaneous frequency varies linearly with time, and can be shown to result in a spectral roll off of 3 dB per octave. In the digital domain this sweep results in a very clean frequency response, all frequencies are reproduced faithfully in the digital domain. Once cut to a record however, the sweep is once again at the mercy of the vinyl playback system, an analog system. This analog system imparts its own characteristics onto the spectrum of the sweep, changing the frequency response. An example spectrum of the sweep is given in Figure 2.8.

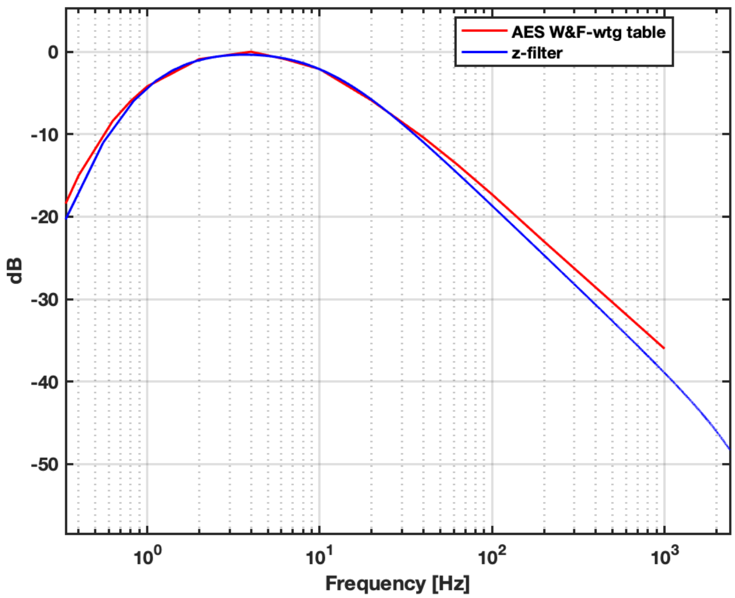

Figure 2.6: The Weighting Filter used in the Wow and Flutter Measurements.

The cartridge is a physical device, itself with internal resonances and variations in the way it responds to the certain frequencies, the playback system itself has other "frequency responses" that impart their own characteristics on the spectrum. The RIAA Equalization Curve discussed in Section 3.3 is one specific instance of a frequency response being directly imparted on the system, however frequency responses can be imparted either directly, as is the case of the RIAA equalization curve, or indirectly, as is the case of the physical resonance of the cartridge.

As each aspect of the system imparts in own characteristics on the frequency response of the system as a whole, the sweep tracks of the record allow measurement of this overall frequency response, from groove to digital file. The sweep cut to the record is meant to produce a 3 dB per octave slope to the spectrum. So any deviations from flat would be heard as a definite change in the music's character, The audibility of these changes is meant to be captured by the A-weighting filter. However when trying to understand the vinyl system as a physical process, it's best to look at the unweighted spectrum.

It's important to note the distinction between the noise spectrum and the frequency response of the record. The noise spectrum shows the frequencies present in a groove that contains no signal. This is the spectrum of the background surface noise of the vinyl. The

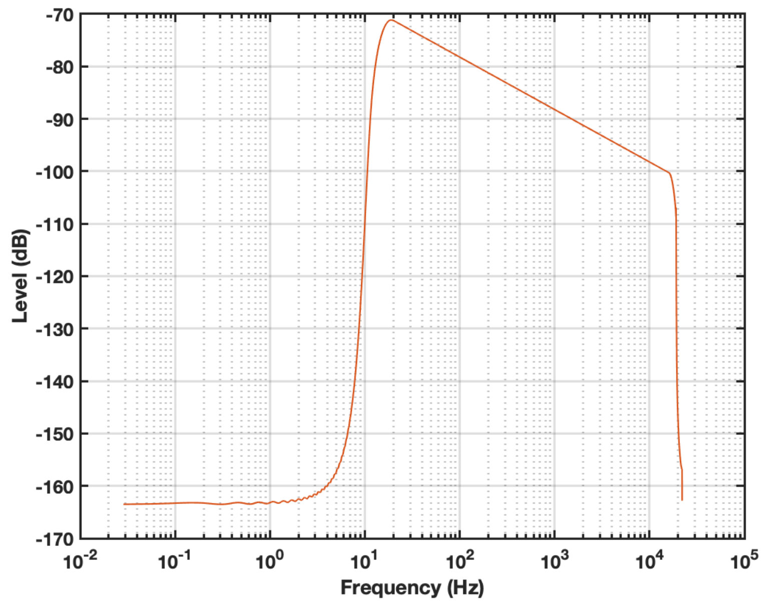

Figure 2.7: The Spectrum of the Digital Sweep cut to the records as a digital signal. This signal has a gentle roll off of 3 dB per octave.

frequency response of the sweep tracks represent vinyl's ability to reproduce uniformly any frequency in the audible range.

2.5 Stereo Separation

Stereo separation is another important metric of audio quality. In a digital file, such as the WAV files used in this research, the recording can be thought of as truly stereo. The file contains distinct audio information in each channel, while there are elements that will be similar between the two channels, those are an artistic or purposeful choice. The two channels are truly separate, and any audio common to the two channels is meant to be there. In audio systems, there is usually an effect called stereo bleed.

Stereo bleed can occur in a variety of audio equipment. In stereo amplifiers for instance, it can typically occur because the electrical components and circuits that amplify each channel are contained in one single unit. The electrical currents and voltages carrying the audio information in one channel can interact with the currents containing audio

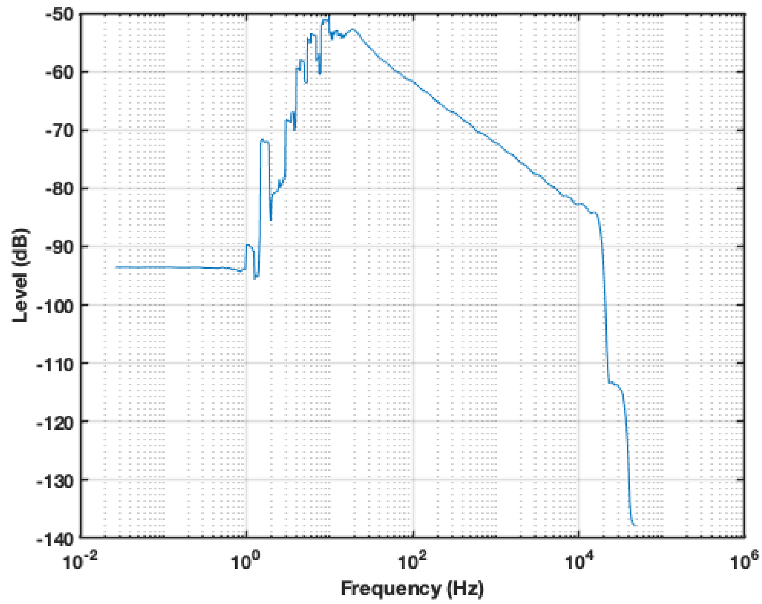

Figure 2.8: The Spectrum of the Analog Sweep recorded from a vinyl record. Note the noise and distortion that occurs from the vinyl playback system that is not found in the digital file.

information in the other channel. This interaction bleeds information from one channel in with the other and now there is audio information common to both channels that was not originally there.

In the vinyl system, signal can bleed across channels in a variety of ways. The audio is encoded in stereo in the construction of the groove. Grooves on vinyl records are V-shaped, each wall that forms the V represents a channel. The inner groove wall contains the left-channel audio information and the outer groove wall contains the right channel. The stylus sits in the middle of this groove, and is modulated both back and forth and up and down as the groove moves. A signal common to both channels oscillates the stylus laterally, while an opposing signal in the channels oscillates it vertically. This physical motion moves the stylus needle, which is attached to coils that sit in a magnetic field. The coils in the cartridge translate the motion of the stylus to a voltage, which travels through to the amplifier. There is an initial bleed between the stereo channels, based on the physics of the stylus in the groove and the motion inside of the cartridge.

Stereo bleed is measured quite simply in the test records, certain signals were cut in either only the Left or Right Channel, see Figure 2.9 as an example of a 1 kHz tone cut to only a single channel. The signal level in the opposite channel is then measured during playback and the ratio is taken as the stereo bleed measurement. There is a limitation with this method, and in hindsight a narrow band measurement would have been better. As the signal bleeding into the other channel is around the same level as the unweighted surface noise, the noise in the channel threw off our wide band measurements.

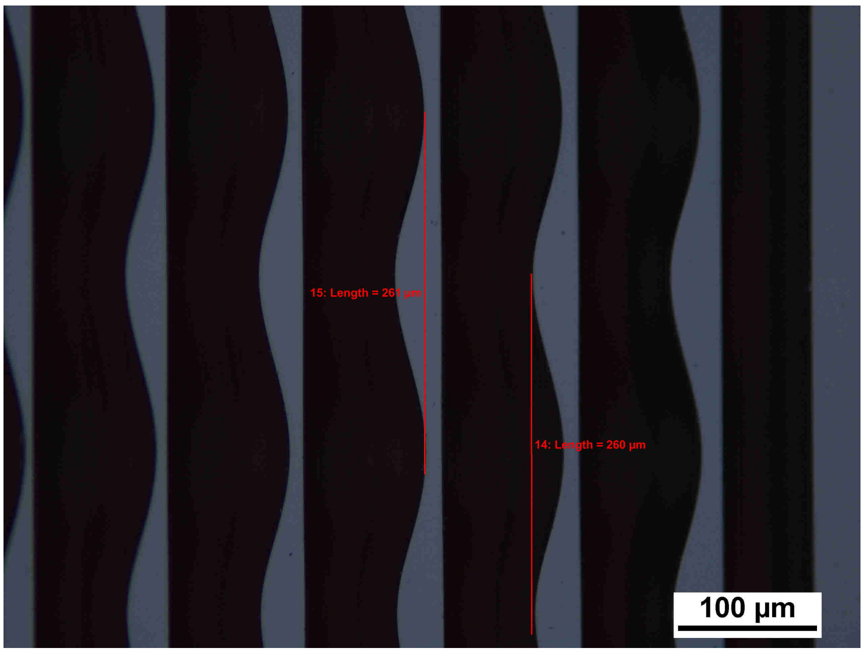

Figure 2.9: The 1kHz sine wave cut only to the right channel.

2.6 Total harmonic distortion

The last measurement is that of distortion. Distortion is a process by which upper harmonics are introduced into the audio signal. There are various sources of distortion during the vinyl playback system, the two most common cause of distortion is tracing distortion. This is the result of a round or spherical stylus tip being used in a groove that was cut with a triangular cutting head. As a result, the contact points of the elliptical stylus are not the same as the contact points of the cutting head, and timing errors are induced in the reproduction of the audio signal and non-linear distortion can be heard in the resulting playback. [8]

There are other causes of distortion too in the Vinyl System such as groove deformation from stylus playback however that isn't a concern for this thesis as the majority of test records were only played and recorded a single time. The total harmonic distortion measurement is handled by the MATLAB function "thd". This function takes a sinusoidal signal, it's important to note that all signals on the record are sinusoidal, and "thd" returns the total-harmonic distortion in dBc. It describes the total harmonic distortion of the first five harmonics. The function does so using a periodogram of the same length as the input signal.

2.7 Separate tracks script

Throughout this thesis various MATLAB scripts were developed in order to process the hundreds of recordings of test records. All of this code is included in Appendix A. These scripts are split into various stages of the analysis. The scripts combine data from the recordings and the press sensors and settings into a large table that is able to be filtered by track and side. This allows easy comparison and plotting of any of the sensor measurements against any of the audio measurements.

This section details the scripts used and the algorithms developed to process both the data coming from the WarmTone press, the audio from the recorded records and the final sorting of all the data into a master table. The inputs to the record script are the sensor readings from the WarmTone, the settings pulled from Viryl's ADAPT system and the audio files itself. There are several intermediate tables made, namely the AudioTable and the SensorTable.

The separate tracks script takes a complete 16 minute recording of a test record and outputs a MATLAB *map container* that allows easy access to the tracks within the long recording via *key*, *value* pairs in MATLAB. It works using labels in the filename that mark the timestamp that the stylus needle enters the lockout groove. This is a much more reliable place to line up the records when compared to the beginning of the record as the lockout groove is always in the same position as compared to the needle drop, which has a variety of groove segments and an infinite amount of places it can land on the record. The SeperateTracks.m file is include in Appendix A.1.

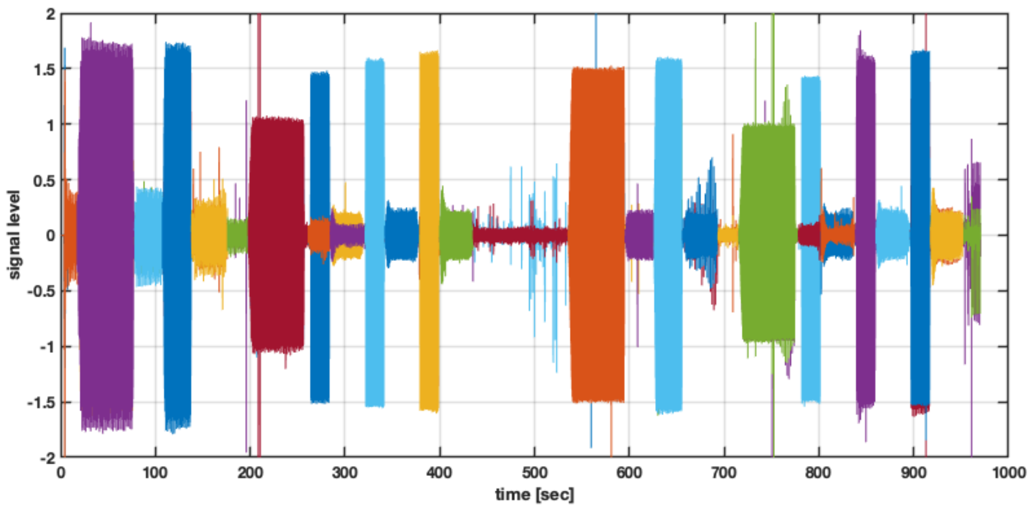

Figure 2.10: An example of a recorded test record with each track plotted as a different colour.

2.8 Record processing script

The record processing script takes the tracks *map container* that is the output of the SeperateTracks.m file, and outputs one row for every track on the record, 27 in total including the leadin, leadout and transition in the AudioTable. The AudioTable lists every audio measurement for each track and is quite a large table. An example of AudioTable data is given below in Table 2.2.

pressing	record	side	track	lagdiff	normalizatio	normalizatio	RMS_L	RMS_R	A_L	A_R	CCIR_L	CCIR_R
midpoint	-01a1558.88	a	100Hz	-78528	0.44501485	0.27249242	-0.3672525	0.21728925	-19.511477	-18.927891	-26.043136	-25.461038
midpoint	-01a1558.88	a	100Hz2	-78528	0.44501485	0.27249242	-0.2618131	0.31169745	-19.405103	-18.831727	-25.851874	-25.277364
midpoint	-01a1558.88	a	10kHz	-78528	0.44501485	0.27249242	-19.328994	-18.225766	-21.608992	-20.598443	-17.451354	-16.442228
midpoint	-01a1558.88	a	10kHz2	-78528	0.44501485	0.27249242	-20.563175	-19.546341	-22.524033	-21.541443	-18.536053	-17.551213
midpoint	-01a1558.88	a	1kHz	-78528	0.44501485	0.27249242	-0.3610613	-0.290967	-0.3546106	-0.2912656	-6.2523676	-6.1995242
midpoint	-01a1558.88	a	1kHz2	-78528	0.44501485	0.27249242	-0.5866886	-0.1118463	-0.5724404	-0.0988981	-6.4401013	-5.9738882
midpoint	-01a1558.88	a	1kHzL	-78528	0.44501485	0.27249242	-0.8748379	-27.009445	-0.8615363	-29.054156	-7.2860039	-35.267271
midpoint	-01a1558.88	a	1kHzL2	-78528	0.44501485	0.27249242	-0.9614598	-27.808053	-0.9464499	-28.49034	-7.4809643	-34.777239
midpoint	-01a1558.88	a	1kHzR	-78528	0.44501485	0.27249242	-31.029645	0.2582717	-35.092648	0.27175209	-40.776333	-5.6945295
midpoint	-01a1558.88	a	1kHzR2	-78528	0.44501485	0.27249242	-28.642703	0.00855828	-29.56349	0.02309406	-35.398278	-6.0532854

Table 2.2: An example of the AudioTable. The output of the RecordProcess.m script.

Record Process outputs a large amount of columns as it records many measurements for the two stereo channels and a large amount of rows as it needs to record each of these measurements for each individual track per record. Crucially all the RMS levels, numbers of clicks, distortion, wow and stereo bleed measurements are recorded for each track. This gives a ton of data per record and is a full assessment of what the test data "sounds like" quality wise to the MATLAB script. Batch Process is responsible for processing a large amount of recordings at once with the record process script and automating the process of generating the AudioTable from the output of Record Process. Its code and the code for Record Process are included in Appendices A.3 and A.4.

2.9 Sensor processing script

The WarmTone press has various sensors within it and outputs its sensor readings in csv files that measure all sorts of parameters in the press. The raw data isn't too useful, it needs to be parsed and sorted first before the readings in the press can be matched the audio of a recorded test record.

The press samples its sensors at a rate of 25 times per minute. This is independent of what the press is doing. There are no measurements done on a per record basis, as the measurements are instantaneous to that timestamp. Meaningful measurements data and points from these sensors need to be made in order to tie what is going on in the press to any one individual record. This is done by the SensorProcess.m script and is included in Appendix A.2. The SensorProcess.m script outputs the SensorTable. The SensorTable is listed by pressing number and does track the records in the press retroactively, each record is detected as being pressed based on the position sensor of the press ram. The opening and closing action of the press is detected by the script and a record is considered pressed based on the time the press spends closed and a buffer of samples once it's opened.

The sensor data is then looked at for each record pressed and the maximum and minimum value of each sensor measurement for each record is taken. This is to distill the large amount of sensor data per record into a much smaller set of readings that can be useful when looking at the larger dataset. Pressing a record is a dynamic, time-based process that involves cycles of heating and cooling. These cycles are sampled by the SensorProcess script and the maximum and minimum measurements of each sensor during each cycle is taken.

An important piece is missing when only looking at the minimum and maximum values, especially for the mould temperatures. Importantly, operators do not have control over their steam temperature through the WarmTone, instead that is controlled by the boiler set up at the facility. Operators have control of the heating profile however, that is the length of

RecordTimeStamp	PressPosition_Inches	PressForce_Ton	MouldSteamIn_PSI	MouldSteamIn_F	MouldSteamOutTop_F	MouldSteamOutBottom_F
2019-12-06 10:26	0.059	1	0	342.6	183.7	185.9
2019-12-06 10:26	0.059	1	0	341.3	220.1	231.4
2019-12-06 10:26	0.059	1	0	342.6	245.3	249.4
2019-12-06 10:26	0.059	1	0	342.6	232.3	232.2
2019-12-06 10:26	0.059	1	0	344.1	187	191.8
2019-12-06 10:26	0.059	1	0	345.4	153.7	158.2
2019-12-06 10:26	0.06	1	0	346.2	131.5	135.9
2019-12-06 10:26	0.059	1	0	346.8	116.8	119.8
2019-12-06 10:26	0.059	1	0	347.4	106.2	107.8
2019-12-06 10:26	0.058	1	0	347.6	101.7	106.7
2019-12-06 10:26	0.059	1	0	347.6	100.6	105.8
2019-12-06 10:26	0.059	1	0	347.6	99.9	104.9
2019-12-06 10:26	0.059	1	0	346.2	100.9	104
2019-12-06 10:26	2.984	0	0	342.4	101.3	102.4

Table 2.3: An example of raw SensorValue data from the WarmTone press. The input of the SensorProcess.m script.

PressingNumber	maxPressForce_Ton	maxMouldSteamIn_F	minMouldSteamIn_F	maxMouldSteamOutTop_F	minMouldSteamOutTop_F
1	1138	347.6	338.3	245.3	99.9
2	1138	351.8	338.1	230.7	76.3
3	1138	353.1	338.6	230	69.6
4	1138	350	344	245.8	96.1
5	1138	346.8	339.3	240.6	101.3
6	1138	349.4	337	237	102.4
7	1138	350.9	346.1	248.4	102.6
8	1138	348.1	342.5	245.7	104.5

Table 2.4: An example of the SensorTable, the sorted sensor data from the WarmTone. The output of the SensorProcess.m script.

time that steam is pushed through the moulds, held there and how long water is pumped through the moulds. These correspond to the heat, dwell and cool times. The overlap time is similar to the dwell, it is discussed in more depth in Chapter 4. It is this heating profile that is of most interest to the operators, and the maximum and minimum value's cannot accurately capture every detail that the entire profile would provide. However the maximum and minimum values do still give an indication of what the conditions were like when the press was closed and the record was made as the temperatures measured inside the press do vary.

Chapter 3

Recording and experimental setup

For the purpose of analysis, the records must be recorded digitally. The recording setup was chosen to accurately and faithfully recreate the audio on the test records and to have as low a noise floor as possible. The Scarlett 2i2 was set to operate at a sample rate of 96000 Hz and a bit depth of 24, ensuring these digital recordings have a nyquist frequency of 48 kHz and a theoretical digital noise floor and dynamic range of -144 dB. This is well and above the required sampling rate and bit-depth that a digital system needs in order to measure a vinyl system, the RIAA preamp isn't specified above 20 kHz and weighted surface noise levels are around -55 dB on a vinyl record.

The test records were sleeved and brought to the University of Waterloo's Audio Lab in order to be recorded. The recording setup consisted of an Orotofon Blue cartridge, a Technics Sl-1200GR turntable, a modified Technics SU-9070 preamp and a Focusrite Scarlett 2i2 Gen 2 interface connected to a Windows PC. The PC was a laptop with an Intel (R) Core (TM) Duo CPU T450 @ 2.00 GHz, only 2 gb of RAM and is running Windows 7, using the open-source program Audacity to record audio. The cartridge, turntable, recording interface, amplifier– everything in the signal chain will impart a certain characteristic upon the signal. It's important that these effects and characteristics are known or otherwise accounted for during recording, in order to get a true sense of the quality of the underlying record.

This Chapter details the equipment used to record the test records, the required gain structure, the turntable's specifications and real-world performance and the RIAA preamp used. Additionally sample dropouts were experienced in the recordings that threw off the coherence measurements in Chapter 5 however they have little impact on the overall results and measurements of audio quality as those results averaged large sections of data. This

Chapter discusses those sample dropouts and their impact on the data.

3.1 Focusrite interface, noise floor and sampling rate

The entry point of audio into the digital system is the Focusrite Scarlet 2i2 Gen 2 audio interface. The tape record output from the Technics SU-9070 runs into the Focusrite sound card's inputs which is connected via USB to a Windows laptop running Audacity. Crucially the Focusrite interface contains the two gain pots that are the only volume controls on the entire recording system. These set the overall recording level into the interface, which is somewhat arbitrary.

Indeed these gain controls are somewhat of a nuisance on the recording setup, as the RMS levels of noise, distortion and other aspects of records are being compared; the system requires normalization in order for recordings made at different gain settings to be compared against one another. The normalization procedure is described in Chapter 2. Using normalization these gain pots can be set at a comfortable level to capture the required dynamic range and forgotten about. For the purpose of this testing the recording level was set so that the 1 kHZ normalization tone read around -6 dB in Audacity during recording. This gave adequate headroom for pops and clicks that might induce amplitudes above 7 cm/s, and still ensures that there is enough dynamic range to capture the surface noise accurately.

The Focusrite was set to a sampling frequency of 96 kHz and a bit-depth of 24 for the recordings made in this thesis. These were chosen as the Focusrite was capable of making the recordings at this quality. In hindsight however, this high sampling rate and slow computer may be the cause of the sample dropouts that are described in Section 3.4.

3.2 Technics SL-1200GR turntable, specifications and performance

Several tests were performed in order to measure the accuracy and consistency of rotation of the SL-1200GR platter. This is important as any speeding up, slowing down or cogging in the electric motor will modulate the pitch of the recordings. This can knock the records out of alignment, and cause issues with syncing up the audio of the different records. This speeding up and slowing down will also be heard by the listener as wow.

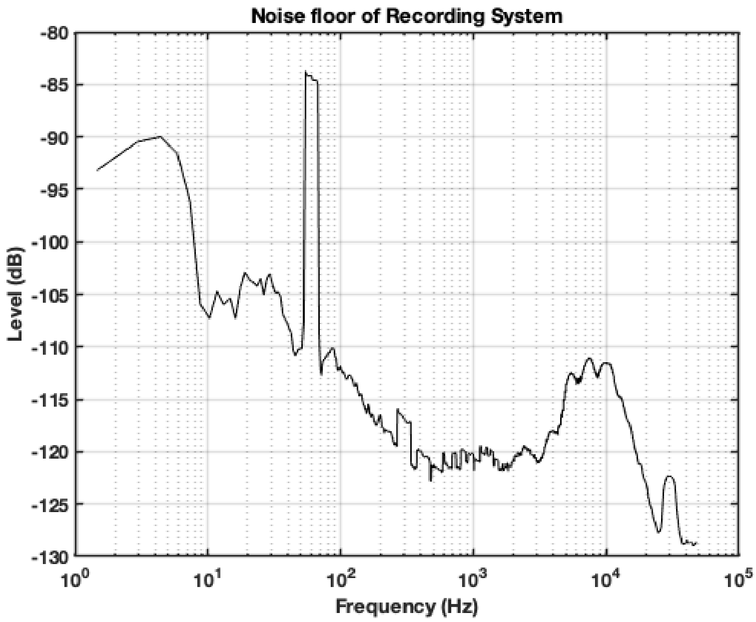

Figure 3.1: The noise floor of the recording setup from the Ortofon Blue Cartridge, Technics SL1200GR, Technics SU-9070 into the Focusrite interface recorded at 96 kHz 24-bit.

An important measure of a record's quality is its wow. Wow is a modulation of pitch that can occur a variety of ways. The most notable way it is generated in this test is via the centre hole. If the centre hole is offset from the geometrical centre of the record, then the speed of the groove beneath the cartridge modulates from slow to fast in a predictable almost sinusoidal manner with a period that matches the period of the turntable. When the record is playing a pure tone, this wow can be measured as a modulation of the frequency of the tone. The details of this measurement process can be found in Section 2.3.

Here, we will study the wow caused by the rotation of the turntable itself. The turntable is powered by a crystal-controlled electric motor, that rotates a 1.7 kg platter. The turntable has a specified wow and flutter of 0.01% measured using the a signal from the frequency generator from the turntable, 0.025% WRMS (JIS C5521) and ±0.035 % peak (IEC 98A Weighted). An attempt was made to examine these standards and compare these values with our own measurements– however the standards were inaccessible.

The wow and flutter of the turntable assembly can be assessed using a reference recording. The main cause of wow is the centre hole offset, so using a record with the centre hole as close to the centre of the record is required if the wow of the turntable is to be measured. For this purpose, the wow and flutter was measured using a lacquer. Lacquers are typically made to a much higher standard than a record. They are what the audio is cut into and ultimately cloned into every record. The centre hole and the lathe that cuts the lacquer are made to much tighter tolerances and specifications than a record. Due to this, it is assumed that the centre hole of the lacquer has no or little offset, and a measurement of the wow of the lacquer is directly a measurement of the wow of the turntable.

The frequency variation of the 3150 Hz track is given as a function of time in Figures 3.2 and 3.3. Figure 3.2 is the raw frequency variations, while Figure 3.3 has a wow and flutter filter applied. The applied filter is specified in IEC386. This filter was recreated digitally. It's frequency response is given in Figure 2.6.

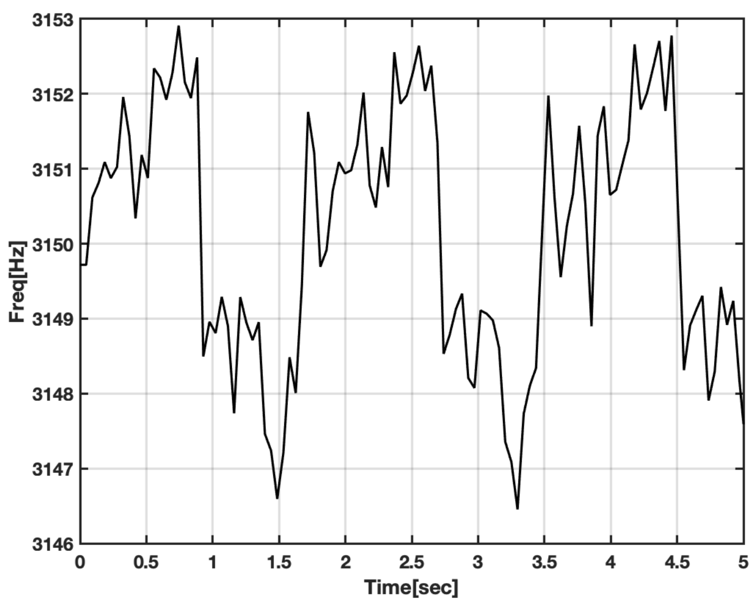

Figure 3.2: The frequency variation of the 3150 Hz signal measured on a lacquer, showing the wow and flutter of the Technics SL-1200GR turntable.

The frequency variations as a function of time measured on the lacquer's unfiltered 3150 Hz track is given in Figure 3.2. What is seen here is the variation in speed of the turntable platter, caused by cogging in the electric motor. Notice the periodic nature of the variations. Ultimately in the 3150 Hz unfiltered signal, there is a peak wow of approximately 3 Hz corresponding to 0.10% wow. In the filtered 3150 Hz signal, the peak

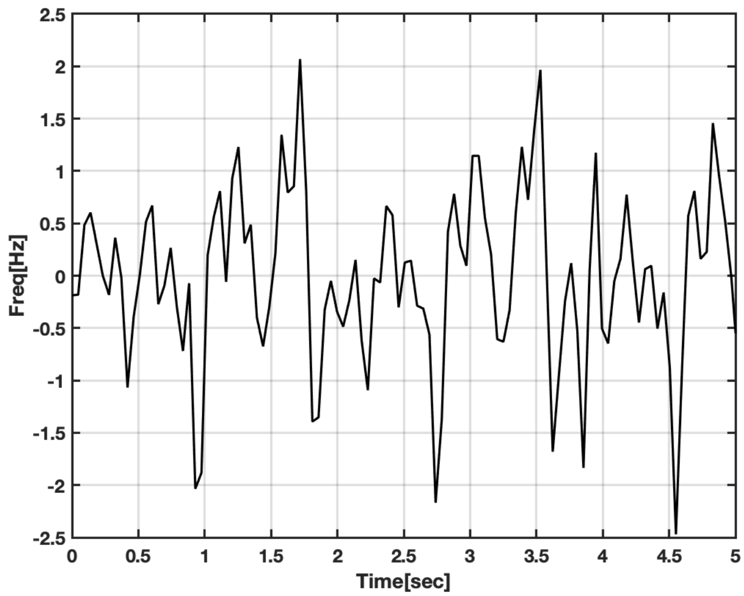

Figure 3.3: The frequency variation of the 3150 Hz filtered signal measured on a lacquer, showing the wow and flutter of the Technics SL-1200GR turntable.

wow is approximately 2 Hz, corresponding to 0.06% wow. These values are higher than the specified wow of the turntable, however significantly lower than the average wow of the test records–which is typically around 0.3% peak unfiltered. The large discrepancy between the specified wow and the measured wow is most likely due to a small centre hole offset of the lacquer.

Regardless of the influence of the lacquer, the cogging of the Technics motor can be seen in Figures 3.2 and 3.3 as the much smaller, higher frequency oscillations above the low frequency fundamental. It's important to note that the oscillations are much lower than the effect of the centre hole offset, even when measured on the lacquer. However they are still there, and influence the measured wow. Depending on how the centre hole is offset, they can contribute constructively or destructively to the measured wow of a record, and as such those wow and flutter measurements need to be understood in the context of approximately 0.1% peak error.

3.3 Technics RIAA preamp

To properly play a record requires an RIAA preamp. The RIAA standard is a set of two equalization curves, one that is applied just before the signal is cut to lacquer and another that is applied as the record is played back. These two curves are designed such that they cancel each other out during playback and only the original audio is heard.

There are several reasons the curves are necessary. The purpose of RIAA equalization is to control the displacement, velocity and acceleration of the stylus to ensure playability and sufficient playing time. The groove-stylus system is physical, the audio heard translates directly from the motion of the stylus. The grooves are shaped by the frequencies present in them and the stylus reacts to these frequencies. Low frequencies in music are typically much higher in amplitude than high frequencies. This combination of low and loud signal results in a large groove– large in the sense that the range of motion required for it's oscillation takes up a considerable amount of space on the record. Additionally, the stylus would have to travel along this high amplitude path, something that is difficult for it to track and can cause tracking errors, distortion and in extreme cases skipping of grooves.

Issues with both the groove shape and the motion of the stylus could arise with the presence of these low frequencies. As such, the RIAA curves were developed in order to de-emphasise the low frequencies and emphasise the high frequencies during cutting and to do the reverse during playback. This sidesteps having to cut low frequencies to the disc, saving space and helping cut down on distortion.

The RIAA playback curve is defined by the following time constants $T_1 = 3180\mu s$, $T_2 = 318\mu s$ and $T_5 = 75\mu s$. [13] It's important to note that RIAA curves are only defined between 20 Hz-20 kHz. For the recordings in this study, the Technics SU-9070 preamp was used as the RIAA preamp, using its direct out. The frequency response of the Technics SU-9070 preamp is given in Figure 3.4.

An important note, our Technics was modified to go down to 5 Hz, to include more rumble. This does not influence A or CCIR weighted noise levels, since the low frequencies are heavily down weighted.

3.4 Sample Dropouts

Lining up the audio of the records is a crucial step in the analysis of a record. A test record is recorded the whole way through (approximately 16 minutes) and the timestamp of the tick that occurs when the needle enters the lockout groove at the end of the record is

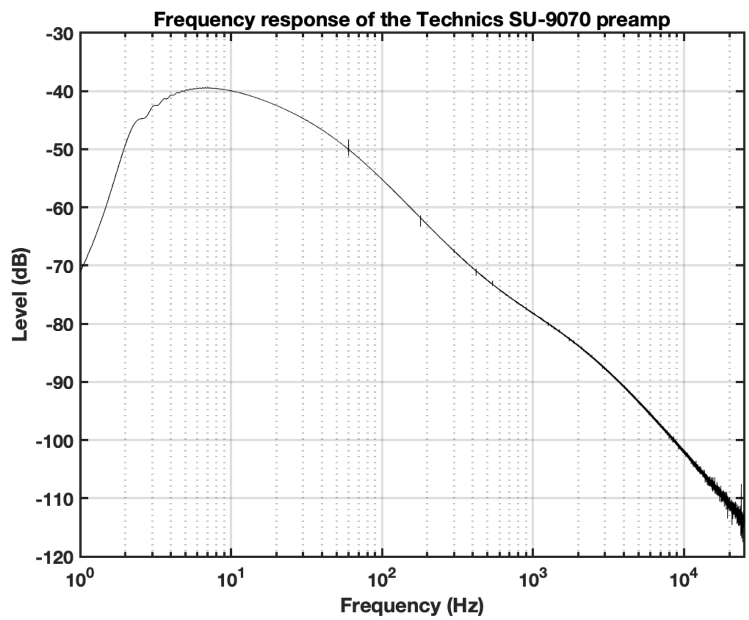

Figure 3.4: Frequency response of the Technics SU-9070 phono inputs.

recorded in the filename. Since every record comes from the same lacquer, this tick should occur at the same spot in each record. So by correcting the time differences between these ticks, we should expect that the audio of two recordings that have been corrected should line up perfectly over top of one another. This is not the case however, it was consistently observed that lining up records in this manner did not guarantee that the audio would be lined up at an earlier point on the record.

There are a few plausible causes for the difference in these track lengths to be observed. One is simply a mechanical error, perhaps a skipped groove. If this were the case then then there would be a large artifact in the audio caused by the violent motion of the stylus as it skips grooves. Furthermore the time differences would be a multiple of 1.8 seconds– which is not the case. The conclusion that was eventually drawn is that these timing errors are due to sample dropouts in the recording. At some point in the recording process, data is being lost. As will be shown later however, sample dropouts could not entirely explain all the timing errors observed.

It is difficult to determine precisely where these dropouts occur and how much data is lost, however the amount lost is significant– several seconds worth. An example of an observed sample dropout is given below. Figure 3.5 is a very extreme example of data loss.

It shows a difference in track length between two recordings to be about 2 seconds. It's difficult to say how much audio is lost in a single dropout, all that can be concluded is that in this problematic example over the course of about 20 seconds, 2 seconds total were lost due to sample dropouts.

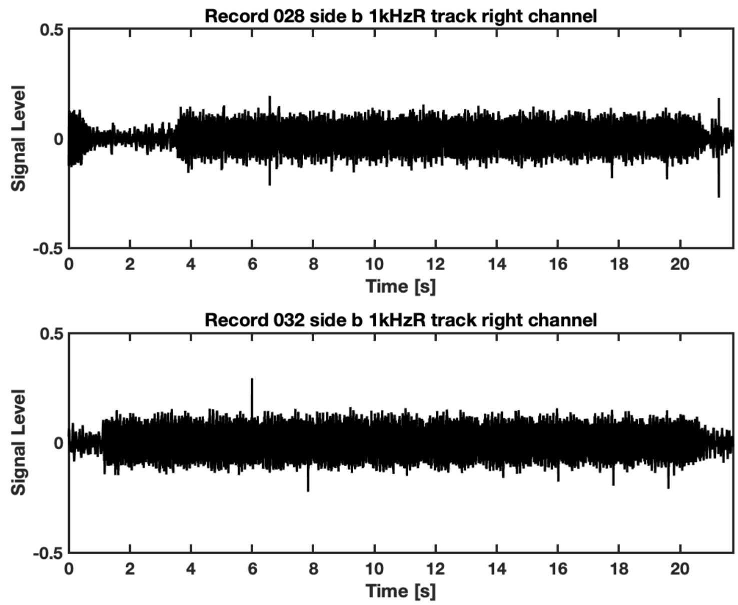

Figure 3.5: A plot of the 1kHzR track for two recordings, one with a sample dropout and one without.

It's important to note that no two recordings were ever perfectly lined up, at least when measured at a point far away from the marks used to line up records. These errors can range from a few samples to a few seconds. The dropouts of a few seconds are not uncommon, and nearly every recording has some kind of timing error.

The extent of these timing errors, and a general location of where they happen on the record can be tested. This is done using the various sweep tracks on the record. Sweeps, being non-periodic, allow any offset in their position in time to be measured. Measuring the peak of the cross correlation of the sweep tracks between different recordings will show how much one signal lags behind the other. The location of the peak will give the total time difference between the two tracks.

These measurements of track lengths via the cross correlation of the sweeps were made on two different sets of recordings. The first set of 5 recordings were made all using the same record, recorded multiple times over. The second set of 5 recordings were made using different records. In either case, one of the recordings was chosen to be the reference, this was the record that was the longest with the fewest timing errors. The lag that corresponds to the peak of the cross correlations for each sweep track on the record is given in Tables 3.4 and 3.4.

For recordings of the same record, in Table 3.4 it's safe to say timing errors only occur in the first recording . Somewhere between the tracks sweepL and sweepR there is a sample dropout of around 1000 samples. The small sample differences of less than 100 samples are of no concern. Such differences are so small when compared to the 96000 Hz sample rate of the recordings, they could easily be due to wow and flutter or other natural variations in timing during playback.

Interesting to note is the extent of the timing errors in the second set of recordings. These recordings are from five different records of the first pressing. These are very large variations and do not entirely fit with what would be expected if sample dropouts were the only contributing factor to these timing errors. It would be expected that any timing errors would follow a similar pattern to Table 3.4, with relatively few dropouts that occur at specific points on the record. However as shown in Table 3.4, the dropouts have different characteristics. As the tracks progress further and further from the point where records are lined up, they are knocked further and further from alignment. Always in the negative direction, the records are always getting shorter compared to the reference.

Recording	sweep	sweepL	sweepR	sweepV	sweep2	sweepR2	sweepL2	sweepV2
1	-1050	-1025	-54	-45	-18	-7	-6	-3
2	-114	-88	-80	-72	-43	-33	-30	-28
3	-117	-90	-81	-72	-45	-35	-32	-30
4	-54	-27	-18	-10	17	28	30	33

Table 3.1: The lag difference calculation from the peak of the cross correlations of the sweep tracks for the records that will be shown in Figures 5.18 and 5.17.

It is clear that a sample dropout is the cause of the timing differences observed in Table 3.4, where the recordings compared were made from the same record. However the timing differences in 3.4 seem to be of an entirely different origin. There are several interesting characteristics to note about the data in this table. As mentioned previously, the records are lined up at the lockout groove at the end of the record. As such it is expected that the last sweep track, the one closest to this lockout groove, is lined up the closest. Indeed this is the case, however the lining up is nowhere near as accurate as in the other example.

Recording	sweep	sweepL	sweepR	sweepV	sweep2	sweepR2	sweepL2	sweepV2
1	-72924	-64466	-56284	-48094	-31490	-16254	-15962	-7899
2	-66196	-57783	-49930	-49778	-25177	-16463	-8586	-207
3	-73854	-65305	-57230	-49224	-24563	-8374	-8106	-218
4	-58488	-42176	-41867	-33782	-16982	-8462	-8172	-7882
5	-68456	-67727	-54784	-54484	-27793	-13802	-13496	-13218

Table 3.2: The lag difference calculation from the peak of the cross correlations of the sweep tracks for the records that will be shown in Figures 5.18 and 5.17, from the first set of recordings.

Recording	sweep	sweepL	sweepR	sweepV	sweep2	sweepR2	sweepL2	sweepV2
1	46	63	80	65	66	-7	-24	-54
2	37	4	8	20	-17	14	-4	-85
3	-16	32	14	-32	-40	-49	-50	-132
4	35	56	63	76	50	80	78	-28
5	40	7	6	19	-11	11	30	-85

Table 3.3: The lag difference calculation from the peak of the cross correlations of the sweep tracks for the records that will be shown in Figures 5.18 and 5.17, from the second set of recordings.

The sample dropouts have been directly observed. One of the dropouts that occurs in Figure 3.5 is plotted below in Figure 3.6. It is difficult to say how much data is lost in a single dropout, or how often they occur. However as shown in Figure 3.5, over a 20 second period the amount of data lost can add up to over 2 seconds.

The dropouts sound like small ticks in the audio, quite hard to detect for the listener, certainly less noisy than a traditional click. Certainly however, they would be noticed by the listener as a surprising amount of data is lost in these dropouts. These dropouts are considered low impact on the main conclusions drawn from this thesis. In all measurements, only the middle two-thirds of the track was used, thus ensuring that no part of the previous track or quiet portions make it into the measurements of RMS levels. Where these sample dropouts become significant is in the measurements of the coherence, as detailed later on in Chapter 5. The most likely cause of these sample dropouts is the Windows laptop used for recording, it is quite old and may be the source of these recording dropouts. It has an Intel (R) Core (TM) Duo CPU, a T450 @ 2.00 GHz, only 2 gb of RAM and is running Windows

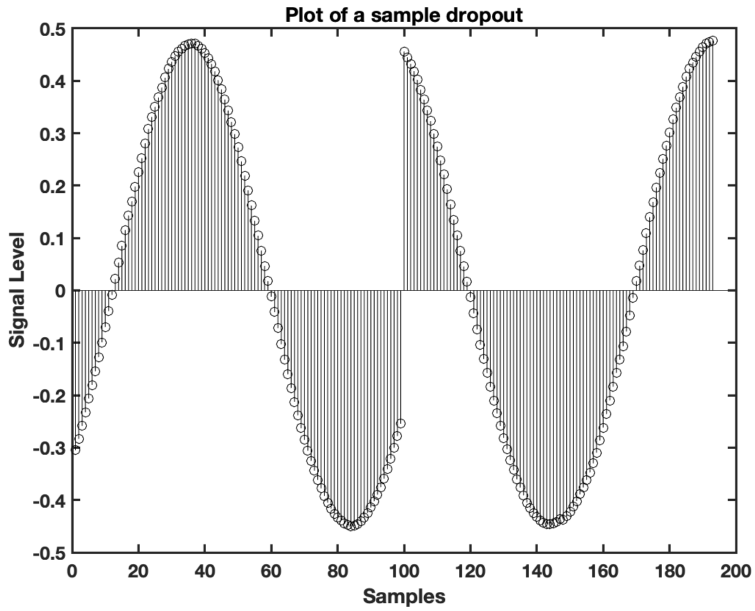

Figure 3.6: A stem plot of a sample dropout. It is impossible to say how many samples were lost in this dropout.

7– which are fairly outdated components and may not have been up to the recording task.

Chapter 4

Comparison between pressings and results

The ultimate goal of this thesis is gain an understanding of how the vinyl press affects the audio of the record. There are various controls and parameters available to operators while records are being pressed. Qualitatively it is known that these controls influence the audio of the final pressed records, increasing or decreasing perceived noise levels, visual defects and other effects.

In the first pressing these controls were kept constant, in order to measure how much variability there was in the audio quality of records pressed under steady conditions. In the second pressing, these parameters were instead varied as much as possible to capture the extremes of the press and see how these changes in parameters affected the audio of the records. In the following sections, the differences between these two pressings are compared and conclusions are drawn regarding the press's effect on the audio quality of a record.

The following will be looked at as examples of the record's audio quality, the RMS noise measurements, numbers of pops and clicks, the wow, the total harmonic distortion measured and the stereo bleed. During sleeving and packaging the records were inspected visually and any common visual defects were recorded as well as an estimate of how warped or cupped the record is. This is all done in order to determine which parameters and setting inside the WarmTone press affect the final audio quality and how, to help operators determine the optimal settings for maximizing the sound of their records.

4.1 Press Parameters and Settings

Two sets of pressings were done at Viryl's facilities in Toronto. Pressing one was done over the course of two days on December 19th and 20th, 2018 and pressing two was completed on December 8th 2019. The two pressings were done for different purposes.

The first set of records pressed were to be a quality control. Before test records of varying parameters could be pressed, the press had to be measured at a steady state. The group of records pressed under these constant parameters serve as a control group and will be used to determine how consistent the records from the WarmTone are. A total of 119 records were pressed at a steady state with 97 a-sides and 93 b-sides recorded. This large sample of records was meant to provide a baseline measurement of the distribution of record quality.

The second set of records pressed are meant to test the WarmTone press parameters, to see what influence changing the press parameters has on the final audio quality of the records. The most important parameters were chosen and records were pressed at minimum and maximum values of each parameter, resetting to the midpoint after each parameter change. A total of 300 records were pressed however only 98 a sides and 93 b-sides were recorded once again. As a lot of the records pressed were at the midpoint setting, and so weren't a part of any parameter change.

One thing to keep in mind surrounding the manufacture of the test records is that two different presses and PVC plastic compounds were used between the first and second pressing. The same plastic was used for the whole second pressing while the parameter changes were happening of course, so it's only a difference between the first and second set of records. This is fortunate as the differences between two PVC compounds can be measured. Unfortunately two different presses were also used between pressings one and two. It would have been much preferred to keep the presses the same, however as parameters were not changed during the first pressing this isn't too big of an issue.

4.1.1 First Pressing

There are three main components of settings that we are concerned with in the WarmTone press. They are the extruder, mould and press. The final component to the WarmTone is the automatic trimmer of course, however that part of the machine has no bearing on the final audio of the record, as it only cuts off the excess compound and creates a smooth edge for the record, so it's settings and parameters are ignored for this thesis. The settings chosen for the first pressing are typical for the vinyl compound used and were, with the

Barrel Zone 1	Barrel Zone 2	Barrel Zone 3	Die Zone	Premould
270 °F	275 °F	280 °F	275 °F	150 °F

Table 4.1: The heat settings chosen for the first pressing within the Extruder in Fahrenheit.

Pucksize	Operating Speed
1.5	45 rpm

Table 4.2: The values for the Extruder settings for the first pressing of test records. Pucksize is a setting inside the WarmTone that determines how much vinyl is extruded into the vinyl "puck". The operating speed refers to the rotation speed of the extruder screw.

exception of the mould settings at the beginning of the pressing when records were sticking to the top mould, kept constant throughout the first pressing.

The first component of the WarmTone is the extruder. Vinyl "pellets" are loaded into a hopper and melted and extruded through a three-level heating system into a hockey "puck"-sized chunk of vinyl that is then loaded into the press to be pressed into a record. The extruder temperatures are typically not a setting that is adjusted on the fly, usually a good temperature is found for each vinyl compound and those values are kept to. Indeed for this pressing the extruder temperatures were held constant. Table 4.1 lists the temperatures that the extruder was set at for the Black PVC compound used for this pressing.

The Barrel Zone 1-3 settings have to do with the main heating elements in the extruder, the die zone is the point where the extruded vinyl leaves the extruder and begins to form the puck, the premould is the heating of the premould that forms the vinyl puck and the operating speed is the speed of the extrusion screw inside of the barrel zone.

After the vinyl "puck" is extruded into the machine it is loaded into the press to be pressed by the stampers into a record. The press closes to a maximum pressure with two stages to it's speed, the first stage is the close to speed and is the speed the press closes at before the vinyl puck makes contact with both stampers and the second stage is the initial speed and is the initial speed at which the press closes while the puck is under pressure. The press closes to a maximum press force and held at that pressure while steam and cold water are pumped through the moulds. Table 4.3 details the settings chosen for the press for this pressing.

Settings were chosen for the first pressing based on the plastic used. The initial planned press settings for this first run of records needed to be changed during the pressing as quite

Press Force	Close to Speed	Initial Speed
1700	50	80

Table 4.3: The values for the press settings for the first pressing of test records.

Heat/Steam	Dwell	Cool	Press Open Delay	Overlap
12.5	0	11.5	6	0

Table 4.4: The values for the mould settings for the first pressing of test records.

a few records were sticking to the top mould and were kept as junk records instead. The issue with records sticking to the top moulds required the mould settings to be adjusted during the start of the pressing process. Records were found to be sticking to the top mould and were interrupting the automatic operation of the press. In order to eliminate this issue the heating and cooling times of the moulds needed to be adjusted.

The mould settings refer to times during the heating and cooling cycle in the moulds. There are five different "times" that can be set in the press. The heat, dwell, cool, press open delay and overlap times all affect the heating profile of the moulds while the press is closed. Heat refers to the amount of time steam is pumped through the moulds, dwell represents the amount of time the steam is held in the moulds, cool is the amount of time that cold water is pumped through the moulds, press open delay represents the amount of time that the press waits after cold water has been pushed through before opening. Overlap represents an overlap of time between the delay and the cool settings.

Initially, when pressing records the test records were sticking to the top mould and paused the pressing. The heating and cooling times inside the moulds were adjusted until records stopped sticking to the top mould. Figure 4.1 demonstrates how the heat and cool times were adjusted in order to reach the ideal settings. Ultimately landing on 12.5 seconds for heating and 11.5 second for cooling. After that point the press parameters were held constant. Table 4.4 lists the rest of the heating and cooling cycle times for the first pressing.

4.1.2 Second Pressing

The goal of the second pressing was to vary the parameters in the WarmTone press as much as possible and compare the relative audio measurements of the records to see what

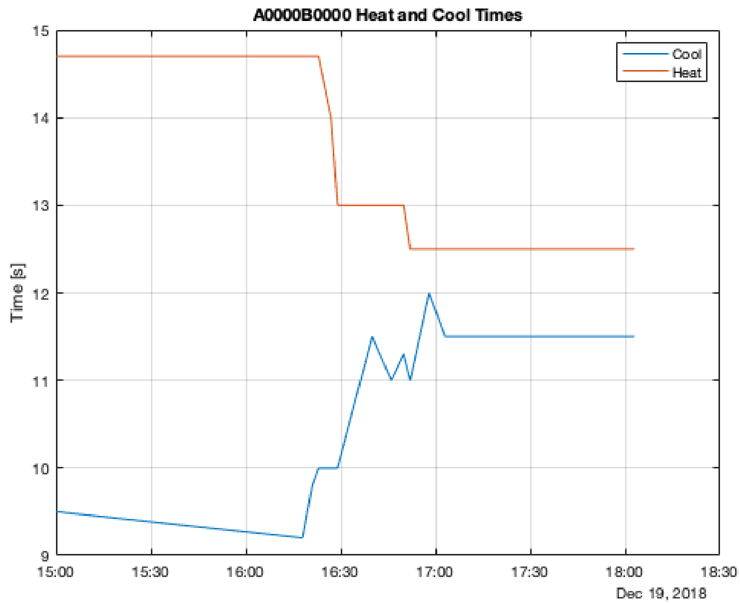

Figure 4.1: The varying heating and cooling times attempting to stop the record from sticking to the top mould during the first pressing.

effect these parameter changes had on the audio quality. For each of the 14 parameters a maximum, minimum and middle value were chosen to represent the typical range of values that the WarmTone would operate at. To keep the number of records pressed at a minimum, and to isolate each parameter change, each parameter was set to either its minimum or maximum value and at least five records pressed. All settings were then reset to their midpoints, and at least five more records would be pressed. This was done in order to ensure the press had time to "reset" after each change. This process was repeated until all the minimum and maximum values were pressed.

Special care was taken in order to ensure that the press had enough time to settle back to its midpoint before another parameter change took place. When appropriate, ample time was given to allow the press to either heat up or cool down. In all cases, every parameter was reset back to its midpoint and at least five more records pressed before another parameter change took place. The order that the parameter changes took place were also significant. When needed the press was given time to adjust to the new settings. Some parameter changes such as puck size or steam time are instantaneous, some such as the barrel zone temperatures took a few hours for the desired temperature to be reached.

In all cases adequate time was given for each parameter change.

The complete list of parameter changes and the order that the changes were pressed in is given in Table 4.5. The press parameters are separated into three categories– extruder, press and mould. The extruder is responsible for melting the vinyl and there are various temperature controls for the different stages in the extruder. The press parameters determine the closing action of the press. The mould parameters deal with the steam and cold water that are fed through the moulds.

Description	Barrel Zone 1 (F)	Barrel Zone 2 (F)	Barrel Zone 3 (F)	Die Zone (F)	Premould (F)	Puck Size	Operating Speed (rpm)	Press Force	Close to speed	Initial speed	Heat/Steam (sec)	Dwell (sec)	Cool (sec)	Press open delay (sec)	Overlap (sec)
min pucksize	260	267.5	275	267.5	157.5	2	48	1850	75	100	10.5	2.5	12	2.5	3
max pucksize	260	267.5	275	267.5	157.5	7	48	1850	75	100	10.5	2.5	12	2.5	3
min cool	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	8	2.5	3
min press force	260	267.5	275	267.5	157.5	3.5	48	1300	75	100	10.5	2.5	12	2.5	3
min steam	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	8	2.5	12	2.5	3
max press force	260	267.5	275	267.5	157.5	3.5	48	2400	75	100	10.5	2.5	12	2.5	3
max steam	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	13.5	2.5	12	2.5	3
min operating speed	260	267.5	275	267.5	157.5	3.5	40	1850	75	100	10.5	2.5	12	2.5	3
max operating speed	260	267.5	275	267.5	157.5	3.5	50	1850	75	100	10.5	2.5	12	2.5	3
min press open delay	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	0	3
max press open delay	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	5	3
min dwell	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	0	12	2.5	3
max dwell	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	5	12	2.5	3
min close to speed	260	267.5	275	267.5	157.5	3.5	48	1850	50	100	10.5	2.5	12	2.5	3
max cool	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	16	2.5	3
max close to speed	260	267.5	275	267.5	157.5	3.5	48	1850	100	100	10.5	2.5	12	2.5	3
min initial speed	260	267.5	275	267.5	157.5	3.5	48	1850	75	50	10.5	2.5	12	2.5	3
min barrel zones	245	250	255	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
Max overlap	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	7
Min overlap	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	0
min die zone	260	267.5	275	235	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
mid point	260	267.5	275	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
min premould	260	267.5	275	267.5	145	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
max premould	260	267.5	275	267.5	170	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
max die zone	260	267.5	275	300	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	3
max barrel zones	280	290	300	267.5	157.5	3.5	48	1850	75	100	10.5	2.5	12	2.5	3

Table 4.5: A complete list of the planned parameter changes during the pressing in the order that the changes were made. The green colour represents the mid point settings. The press was returned to these settings and 5 records were pressed at these settings every time a parameter change took place. The yellow in the table highlights the parameter changes from this midpoint.

4.1.3 Extruder parameters

The important extruder parameters are barrel zones 1-3 temperatures, the die zone temperature, the premould temperature, puck size and the operating speed. The extruder temperatures take time to reach those temperatures. In every case the press was given enough time to ensure that the set temperature was achieved before every record was pressed. The three barrel zones were grouped and changed together. Testing the heating profile of the extruder while interesting would be way too time consuming. So instead a gentle heating ramp was chosen and the maximum and minimum temperatures were chosen for each barrel zone to accommodate that ramp. Table 4.6 gives the values for the maximum and minimum barrel zone and the ramp chosen.

The remaining extruder parameters are given in Table 4.7. They represent typical

Pressing	Barrel Zone 1 (F)	Barrel Zone 2 (F)	Barrel Zone 3 (F)
min barrel zones	245	250	255
max barrel zones	280	290	300
midpoint	260	267.5	275

Table 4.6: The range of temperatures chosen for the barrel zone tests.

maximum and minimum temperature ranges for a record pressed with this type of PVC compound in the WarmTone.

Pressing	Die Zone (F)	Premould (F)	Puck Size	Operating Speed(rpm)
min pucksize	267.5	157.5	2	45
max pucksize	267.5	157.5	7	45
min operating speed	267.5	157.5	3.5	40
max operating speed	267.5	157.5	3.5	50
min premould	267.5	145	3.5	45
max premould	267.5	170	3.5	45
min die zone	235	157.5	3.5	45
max diezone	300	157.5	3.5	45
midpoint	267.5	157.5	3.5	45

Table 4.7: The range of parameters chosen for the extruder settings tests.

4.1.4 Press parameters

The important press parameters are the press force, close to speed and initial speed. Their minima and maxima for the purpose of testing is given in Table 4.8 Note the initial speed was already at the maximum setting for the midpoint so no maximum initial speed setting was tested only a minimum. The press has a two stage closing process, the initial speed and the close to speed. The close to speed is the speed the press closes at before the puck makes contact with the stamper and the initial speed is the starting speed of the press ram as it compresses the vinyl puck into a record.

4.1.5 Mould parameters

The important mould parameters are given in Table 4.9. This pressing didn't have the issues of the first with records sticking to the top moulds, which is surprising considering

Pressing	Press Force (Ton)	Close to speed	Initial Speed
min press force	1300	75	100
max press force	2400	75	100
min close to speed	1850	50	100
max close to speed	1850	100	100
min initial speed	1850	75	50
midpoint	1850	75	100

Table 4.8: The range of parameters chosen for the press settings tests.

how varied the mould settings were. A standard heating profile was chosen for the moulds, from there the minimum and maximum cool, steam, dwell, press open delay and overlap times were chosen.

Pressing	Heat/Steam (sec)	Dwell (sec)	Cool (sec)	Delay (sec)	Overlap (sec)
min cool	10.5	2.5	8	2.5	3
max cool	10.5	2.5	16	2.5	3
min steam	8	2.5	12	2.5	3
max steam	13.5	2.5	12	2.5	3
min dwell	10.5	0	12	2.5	3
max dwell	10.5	5	12	2.5	3
max overlap	10.5	2.5	12	2.5	7
min overlap	10.5	2.5	12	2.5	0
min press open delay	10.5	2.5	12	0	3
min press open delay	10.5	2.5	12	5	3
midpoint	10.5	2.5	12	2.5	3

Table 4.9: The range of parameters chosen for the press settings tests.

4.2 Audio, Sensor and other Measurements and Plots

This thesis deals with a very large number of data points, every record has two sides, 49 tracks (two sets of signals and the transition) and associated with each track are 12 sensor readings with an associated minimum and maximum value, 14 press parameters and 29 audio measurements. Distilling all this data into a useable format and visualizing it properly is a difficult task to represent all that is going on in the data.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	-30.96	-29.14	-54.19	-52.51	-55.71	-53.43
transition	-31.95	-30.54	-52.86	-52.69	-55.41	-52.54
quiet2	-25.71	-23.82	-55.45	-53.71	-55.11	-52.46

Table 4.10: The average Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the first pressing side a.

In this Chapter the data is represented primarily by pressing, the records have been grouped based on similar parameters. This is to give the best idea as to which parameters might be affecting the final surface noise. Another way to view the data is to look at the sensor data inside the press. As discussed in Chapter 2, the maximum and minimum of each sensor reading is associated with each record as it pressed by the WarmTone for the duration of that records cycle in the machine. These maximum and minimum sensor readings can then be plotted against the audio measurements of the same record and any relationship between the conditions in the press at the time the record is pressed and the resulting audio can be spotted.

Unfortunately no such relationship has been seen. There is no relationship between any of the sensor readings and the audio measurements as can be determined visually. Some plots are included in this section to demonstrate this. The remaining plots that couldn't make it into this section are included in Appendix B.

4.3 Noise Measurements

There are three places where the surface noise on the test record can be measured, the two quiet tracks and the transition track. The quiet tracks are simply 30 seconds of silence on the record, that appear in two locations on this disk. Once in the first set of signals on the outer tracks and once on the second set of signals on the inner tracks. The transition track is an extended period of silence that contains both closely spaced grooves and farther spaced grooves. It is the middle track that connects the two sets of inner and outer signals on the record.

The average Unweighted, A weighted and CCIR RMS noise measurements and the standard deviations for both pressings is given in Tables 4.10- 4.15. It's important to note how vastly different the different weighting measurements are. CCIR weighting appears to be the most "stable" and consistent measurement of record noise, the Unweighted and A

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	-30.15	-28.99	-52.96	-51.01	-55.47	-52.86
transition	-32.63	-31.19	-53.89	-51.77	-55.59	-52.44
quiet2	-33.12	-31.82	-53.92	-47.89	-54.76	-51.90

Table 4.11: The average Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the first pressing side b.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	-32.39	-31.14	-46.48	-45.90	-52.96	-51.74
transition	-34.55	-33.15	-53.21	-53.55	-55.58	-53.59
quiet2	-35.88	-35.13	-55.61	-54.19	-55.75	-53.02

Table 4.12: The average Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the second pressing side a.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	-31.17	-30.00	-46.27	-45.69	-52.86	-51.64
transition	-31.37	-30.13	-53.72	-52.93	-55.72	-53.60
quiet2	-33.45	-32.69	-54.54	-46.57	-55.45	-52.97

Table 4.13: The average Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the second pressing side b.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	1.80	1.51	0.58	0.50	1.11	0.70
transition	1.60	1.47	0.84	0.46	0.76	0.62
quiet2	1.73	1.77	0.77	1.03	1.17	1.03

Table 4.14: The standard deviation of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the first pressing side a.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	2.22	2.42	3.30	2.94	1.06	0.92
transition	1.76	1.82	2.01	1.56	0.82	0.86
quiet2	2.39	2.31	3.80	3.95	2.70	1.10

Table 4.15: The standard deviation of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the first pressing side b.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	2.21	2.05	0.37	0.40	0.58	0.67
transition	2.04	1.89	1.30	0.90	1.04	0.97
quiet2	1.52	1.57	0.85	1.30	1.06	1.27

Table 4.16: The standard deviation of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the second pressing side a.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	2.68	2.65	1.61	1.65	1.37	1.42
quiet2	1.31	1.32	2.81	2.83	1.30	1.37
transition	1.90	1.75	2.51	1.81	1.10	1.03

Table 4.17: The standard deviation of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet tracks and the transition track for the second pressing side b.

weighted results both show strange outliers. For the first pressing quiet2 track side a the unweighted noise is very substantial, however this increase isn't seen in the A weighting and CCIR weighting measurements suggesting it is rumble noise or arm resonance that is the cause of this increased noise.

Comparing pressing one to pressing two it is clear that pressing two has more noise when looking at the weighted measurements, however the unweighted measurements show that pressing two has less noise than pressing one. This leads to the conclusions that pressing two might have less rumble noise or arm resonance, however more actual surface noise. Looking at the A-weighted results, there are many oddities that exist. There is a strange result in the right channel A weighted measurements, the quiet2 track on side b shows almost a +8 dB increase in noise compared to the left channel. Additionally the quiet track for pressing two shows a similar increase of about +8 dB in A weighted noise as compared to the quiet track in pressing one. These large jumps in noise are real data, not outliers and perhaps shows the limitations of using A weighting on record noise. CCIR weighting by contrast shows very little variation amongst the different tracks and pressings for it's measurements of noise. Perhaps this "stability" is an indication that the CCIR-weighted measurements are the most reliable and representative of vinyl noise and least influenced by the rumble.

Tables 4.18 and 4.19 show the difference in level between the two pressings. What's interesting is if you average the differences noise measurements across the two sides and all three tracks the average difference between the two pressings +0.11 dB in favour of pressings one. This is a negligible difference in average level and shows that indeed the surface noise of the two pressings are comparable. Indeed any difference between the two pressings appears primarily in the changes in levels depending on the weightings. This suggests a complex relationship between the surface noise and the PVC compound being used. The spectrum appears to shift with the change in plastic, however the overall surface noise remains largely the same.

Overall side b is much worse quality wise compared to side a. The noise is higher in nearly all cases and much more significant, the standard deviation of the measurements are much higher too. Which points to not only more noise on side b, but indeed more inconsistent noise as well. Side a and side b are of course made by independent stampers which both ultimately come from the same master lacquer so should have comparable noise levels. Side a represents the stamper loaded into the top mould and side b represents the stamper loaded into the bottom mould. It's hard to say whether it is this configuration of the stampers that is causing this difference in noise, or if it's a difference between the two stampers. However side b of the records manufactured for both pressings, shows more noise and more inconsistent noise than side a.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	1.43	2.00	-7.71	-6.61	-2.75	-1.69
transition	2.60	2.61	0.35	2.89	-1.82	1.05
quiet2	10.17	11.31	0.16	0.48	0.64	0.56

Table 4.18: The difference in noise measurements of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet and the transition tracks between the two pressings for side a. A positive value means pressing two is louder, a negative value means pressing one is louder.

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	1.02	1.01	-6.69	-5.32	-2.61	-1.22
transition	-1.26	-1.06	-0.17	1.16	0.13	1.16
quiet2	0.33	0.87	0.62	-1.32	0.69	1.07

Table 4.19: The difference in noise measurements of the Unweighted, A and CCIR weighted RMS noise levels of the two quiet and the transition tracks between the two pressings for side b. A positive value means pressing one is louder, a negative value means pressing two is louder.

There is a very interesting result that's occurring in pressing two, a sharp increase in noise in the first quiet track's A weighting is observed, similar the increase in noise in the right channel quiet2 track on side b that was observed in the first pressing. Note the same increase in the A-weighted right channel quiet2 track is observed here. It is a real and very strange result. The CCIR weighted measurements appear to be the most constant over the record. There is very little change in the CCIR weighted RMS noise level across either of the tracks. Indeed it appears as though the CCIR weighted result is the most stable, perhaps an indication that it is the most accurate measurement of the surface noise and least affected by rumble noise and the turntable arm resonance.

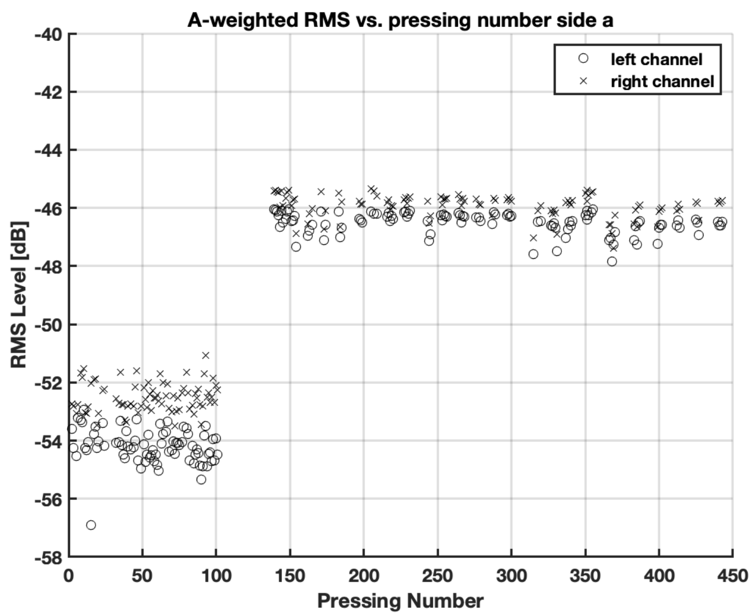

Figure 4.2: The A-weighted RMS level of the quiet track for side a of all records recorded.

Figures 4.2 and 4.3 show the A-weighted RMS level of the quiet track for sides a and b respectively, arranged in the order they were pressed. It's important to note that in displaying the data this way, each clump of records represents a change in a press parameter. The first 137 records (here only up to record 102 was recorded) represent the first pressing run where parameters were kept constant after the first few pressed records. The second group of records represent the second pressing run where the parameters were changed to observe differences in the audio quality. For the first quiet track it's quite obvious there is a large difference between the two pressings in noise. This appears in all three measurements however it is most pronounced in the A-weighted measurements. The

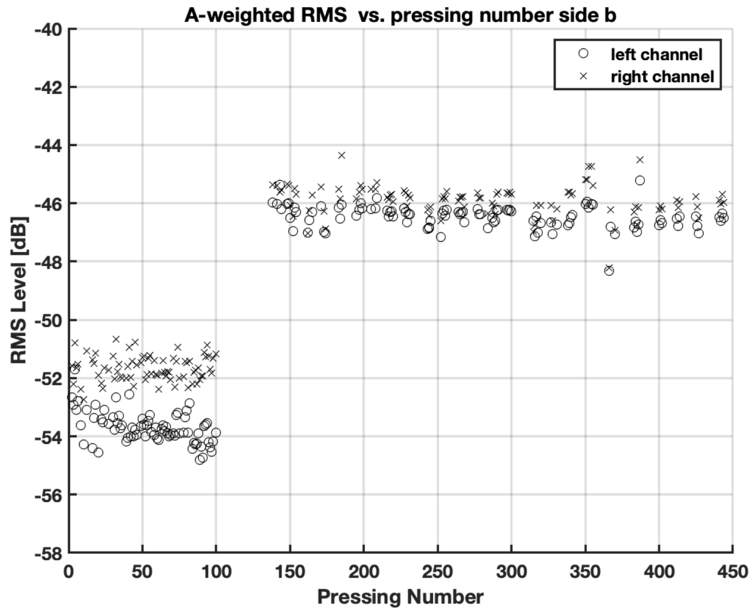

Figure 4.3: The A-weighted RMS level of the quiet track for side b of all records recorded.

Unweighted measurements show a very large variance between records, while the weighted measurements both calm these values and give much more consistent measurements.

Overall it appears that there is much more noise in pressing two compared to pressing once, especially looking at the A-weighted results. However as mentioned previously, averaging all the measurements of noise we see that there isn't a large difference in the average RMS's measured when multiple tracks and weighting measurements are considered and averaged. So the difference between the plastics appears in the spectra in a complicated manner. Overall it's important to note that the differences when comparing pressing one to pressing two are much more pronounced as compared to the differences between any of the parameter changes.

Now we can begin to look at the parameter changes for press two. Figures 4.7 - 4.9 show the Unweighted, A and CCIR weighted RMS noise in the quiet track for each pressing. There are a few interesting results. Side a and side b once again show drastically different characteristics. Side a shows much more variance in the levels of surface noise measured although for the most part the majority of press runs are within the same range of values around -42 dB, there are a few exceptions. Minimum cooling seems to produce much more noise, on par with the side b measurements. Interestingly maximum overlap and minimum press force surpass the original noise measurements of the first pressing with minimum barrel zones close to that baseline as well.

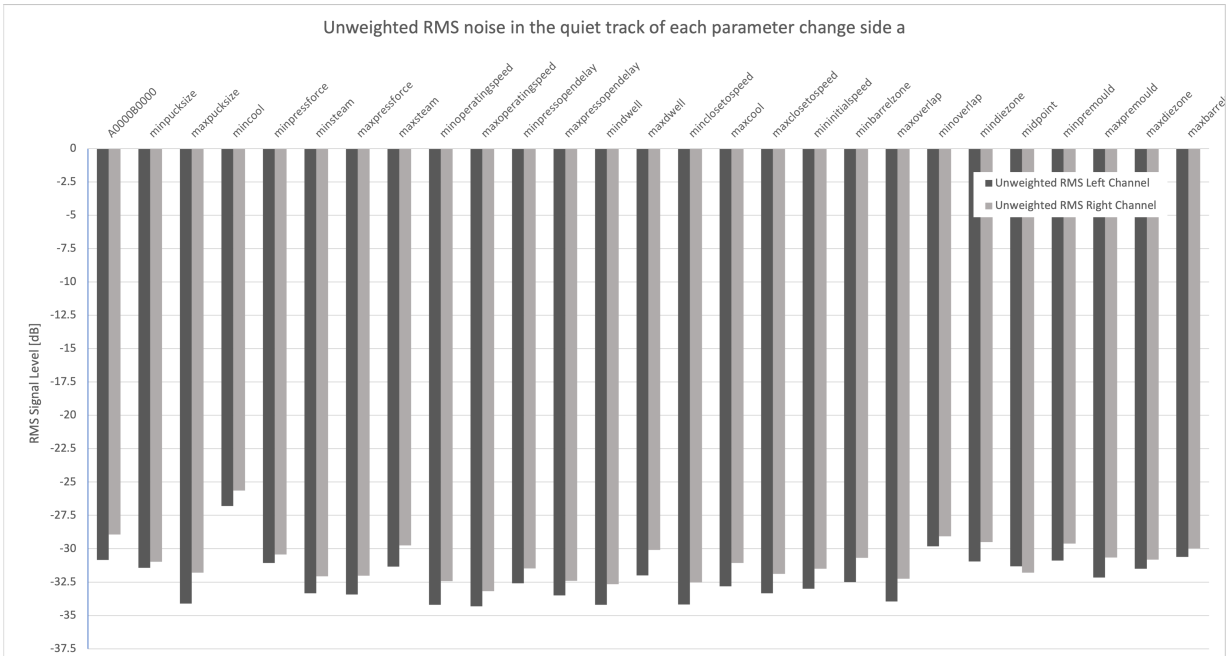

Figure 4.4: The average Unweighted RMS level of the quiet track for side a of all the different parameter changes recorded.

Side b shows that the noise measurements were on average much higher and much more static. Nearly all the pressings have the same noise measurement, a rather high -26 dB. This appears to be a real result and shows drastically different behaviour compared to side a. Oddly, both minimum and maximum barrel zones show lower noise than the average and so does max steam. Maximum overlap and minimum press force, which were the lowest noise pressings for side a, appear to be an average pressing on side b.

Overall the main conclusion that must be drawn is that it is the two different plastics used that caused the largest difference in surface noise but not much else in the other measurements of audio quality. The difference between the surface noise measurements of the two pressings is much more pronounced than any difference between any individual pressing especially when looking at the weighted results. To look at the difference between the parameter changes and surface noise, the unweighted results must be used as the weighted results show almost no difference between the parameter changes. These unweighted results should be taken with a grain of salt, as they mostly deal with rumble

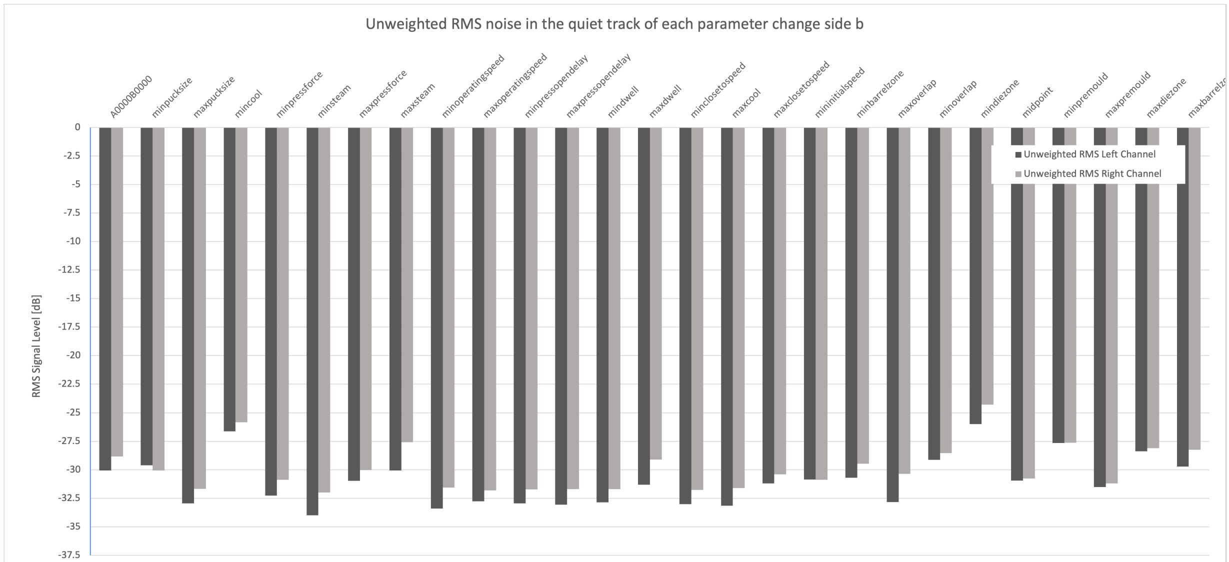

Figure 4.5: The average Unweighted RMS level of the quiet track for side b of all the different parameter changes recorded.

and arm resonance noise which are largely inaudible and not a focus in a press operators assessment of quality, however they are still an indication of the noise on the record.

Looking at the unweighted surface noise measurements there are a few conclusions that can be drawn based on our measurements to inform the manufacturer about the effect the press parameters have on the surface noise of the record. For the extruder parameters the maximum barrel zone shows increased unweighted noise compared to the midpoint while the minimum barrel zone shows decreased unweighted noise as compared to a midpoint. Suggesting a colder temperature is best for the extruder with this kind of PVC compound. Maximum die zone and minimum die zone show noise very close to the midpoint, however the maximum is lower noise and minimum is higher noise in this case. Contrary to what the barrel zones show. Overall the extruder temperatures have little effect on the surface noise, with the exception of the operating speed and pucksize. Both the maximum and the minimum operating speed show greatly decreased unweighted noise compared to the midpoint. Maximum puck size shows less noise than minimum puck size. The operating speed is an interesting parameter to have such a large effect on the surface noise of the record, this might be an indication that it is once again the structure of the vinyl pellets that cause the surface noise of the record as the operating speed no doubt determines how mixed or unmixed the pellets of melted vinyl are in the extruder. Perhaps evidence that

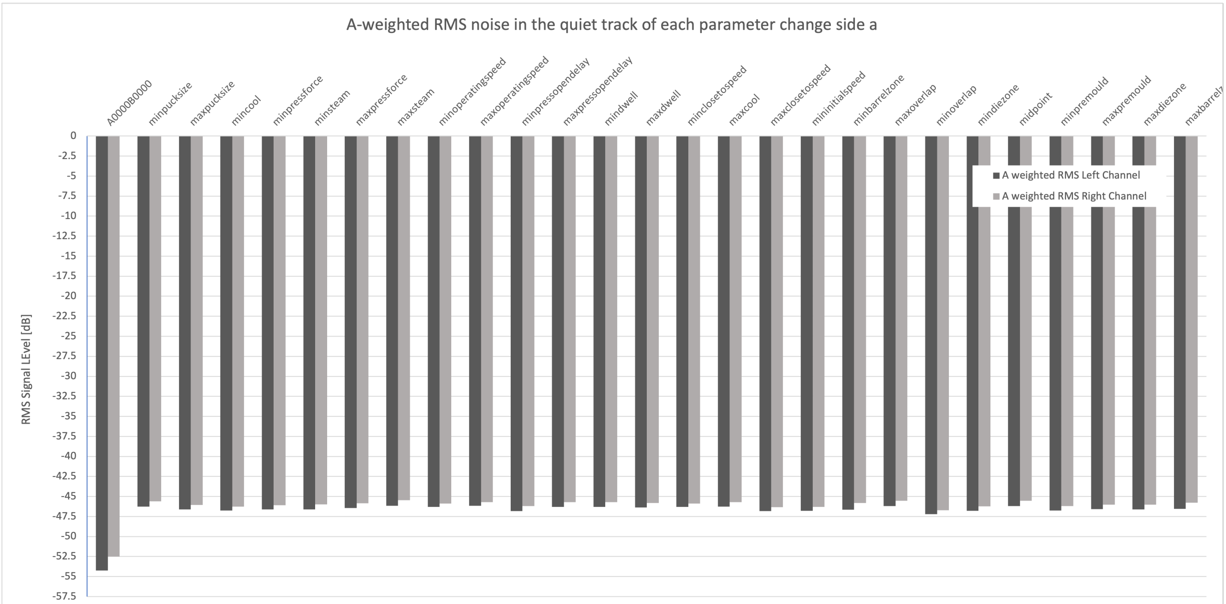

Figure 4.6: The average A-weighted RMS level of the quiet track for side a of all the different parameter changes recorded.

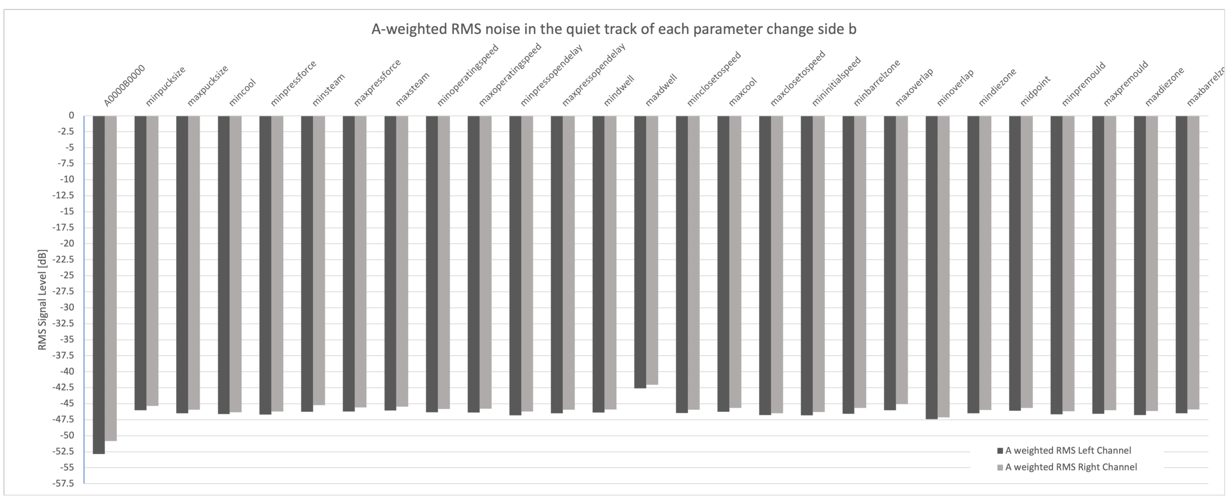

Figure 4.7: The average A-weighted RMS level of the quiet track for side b of all the different parameter changes recorded.

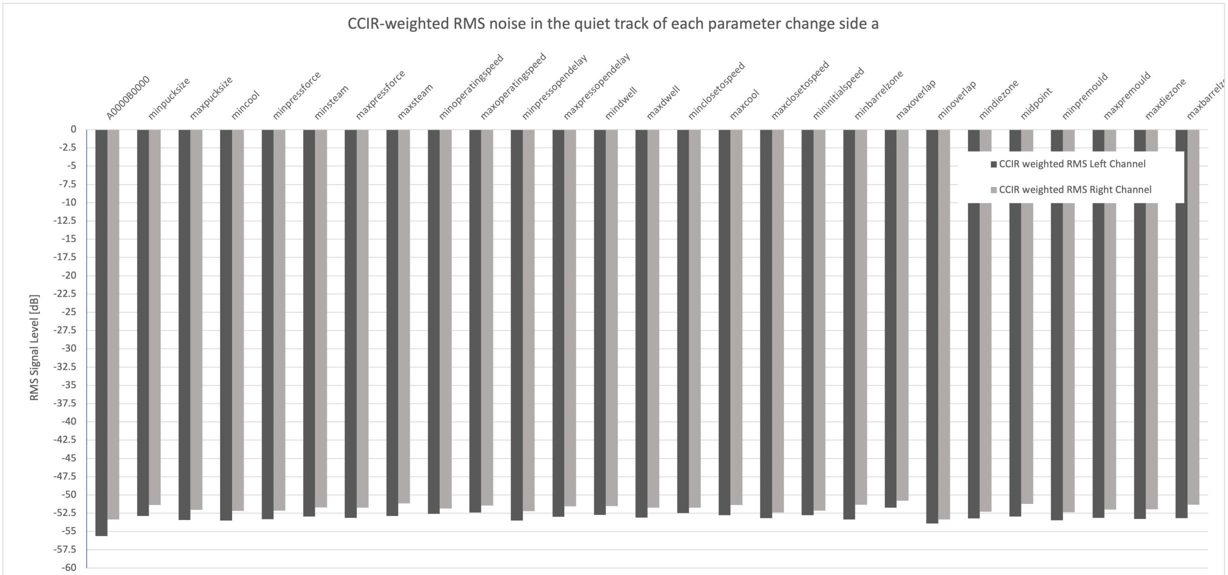

Figure 4.8: The average CCIR-weighted RMS level of the quiet track for side a of all the different parameter changes recorded.

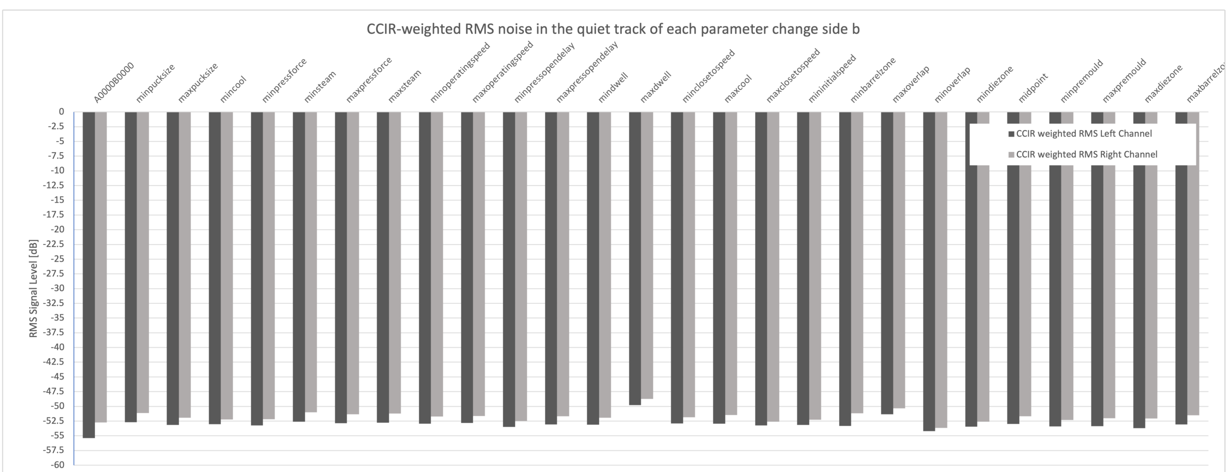

Figure 4.9: The average CCIR-weighted RMS level of the quiet track for side b of all the different parameter changes recorded.

it is the size and shape of the pellets loaded into the extruder and their residual elasticity that is the primary cause of surface noise on vinyl records.

For the mould settings there is a clear result. Minimum steam shows the least unweighted noise and maximum steam is fairly high on the noise measurements too suggesting a lower steam time reduces noise. Minimum dwell similarly shows low noise, remember dwell is the amount of time steam is held in the moulds without being replaced by new steam. So it is related to the steam time. This suggests that too much heating in the moulds causes increased surface noise and should be avoided. Maximum cool shows less noise than minimum cool, in fact minimum cool is one of the noisiest pressings. This suggests that it is the relaxation and cooling of the vinyl that is possibly the cause of and the key to reducing surface noise, once again evidence of the residual elasticity of the vinyl compound informing the amount of surface noise. Cooling in the moulds appears to be the best way to reduce unweighted surface noise.

The press parameters also appear to impact the unweighted surface noise, particularly the press force. The minimum press force shows the largest amount of noise while the maximum press force shows low noise. Minimum and maximum close to speed are fairly similar to one another, however the minimum shows less noise. The minimum initial speed shows less noise compared to the midpoint, which was already set at the maximum initial speed. This suggests a high press force and a low press speed will achieve the best results for reducing unweighted surface noise.

Of course all of this is looking at the unweighted surface noise. The weighted results show no way to reduce or change the surface noise on the record. This is perhaps the more significant result, as it more accurately reflects real world listening conditions. In practice there is little to no perceived change in either the A or CCIR weighted surface noise levels when comparing the different parameter changes, the difference in plastics dominates. Additionally, moving inward on the record also increases the surface noise, with the inner quiet track quiet2 seeing on average an approximately 2 dB decrease in surface noise simply by having the stylus moving inward towards the centre of the record, once again demonstrating the change in vinyl quality as the stylus plays the inner grooves.

4.4 Pops and Clicks

Tables 4.20, 4.21, 4.22, 4.23 show the average number of clicks found per track for the first and second pressing a sides and b sides respectively. For the most part the records were found to be click free, for the first pressing there was an average of 49 clicks pressed into the left channel and 27 clicks in the right channel, for the second pressing this number jumps to 68 clicks in the left channel and 29 clicks in the right channel. Over the approximate 16 minute run time of the record, this equates to an average of 3.7 clicks per minute, in the left channel and 1.8 clicks per minute in the right channel. The increase in clicks between the two pressings is believed to be contamination or damage caused to the stamper during the second pressing as groups of common clicks were found in the second pressing that did not appear in the first due to this damage or contamination.

It's interesting to note that the transition track contains the majority of clicks found on the record, almost ten times as many. It's difficult to say what causes this large increase in the number of clicks, perhaps it is a limitation in the click detector. The transition track is the longest track on the record, so it is expected that it would feature the most clicks. However the number here appear disproportional and no explanation can be offered as to why. As it's typical to find only one or two clicks in the other tracks with their 30 second run time, the transition track can have upward of 60 clicks in it's 97 second run time.

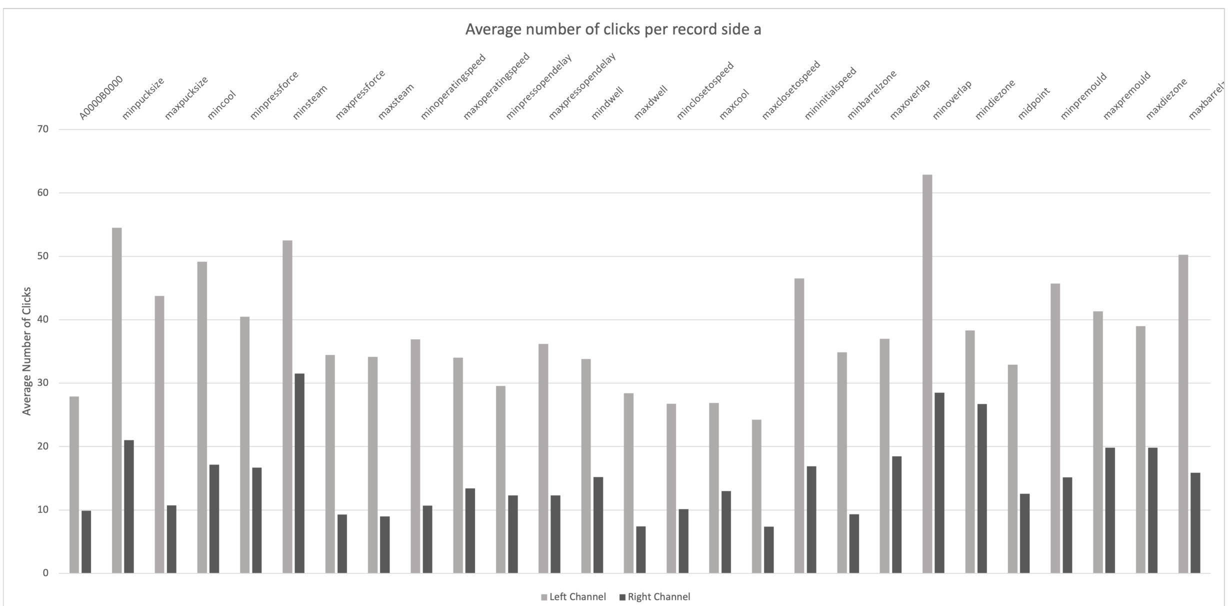

Figure 4.10: The average number of clicks per track for side a of all the different parameter changes recorded.

The number of clicks do seem to vary with the press parameters in a logical way, at least for side a. Maximum cooling shows the lowest amount of clicks, alongside the first pressing and the midpoint while minimum cooling shows one of the highest amount of clicks. Minimum overlap shows the most amount of clicks while maximum overlap shows few clicks. Dwell and steam appear to have no influence on the number of clicks and are both middle values, as they appear relatively close to one another. Max press open delay

Track name	left channel	right channel
1kHz	0.0	0.2
10kHz	0.1	0.1
100Hz	0.0	0.1
quiet	1.0	1.2
sweep	0.2	0.3
3150Hz	0.1	0.2
1kHzL	0.2	1.0
sweepL	0.5	1.0
1kHzR	0.6	0.0
sweepR	1.7	0.3
1kHzV	0.2	0.1
sweepV	0.3	0.1
transition	46.5	12.3
1kHz2	0.2	0.1
10kHz2	0.1	1.0
100Hz2	0.0	0.0
quiet2	1.4	0.5
sweep2	1.0	0.3
3150Hz2	0.2	0.1
1kHzL2	0.8	0.9
sweepL2	0.3	1.1
1kHzR2	1.0	0.0
sweepR2	1.4	0.1
1kHzV2	0.0	0.0
sweepV2	0.8	0.2
Total number of clicks	58.7	21.3

Table 4.20: The average number of clicks per track for pressing one side a.

Track name	left channel	right channel
1kHz	0.5	0.5
10kHz	0.4	0.3
100Hz	0.0	0.0
quiet	1.3	1.4
sweep	0.6	0.5
3150Hz	0.2	0.4
1kHzL	0.3	1.3
sweepL	0.3	1.9
1kHzR	0.8	0.0
sweepR	2.7	0.3
1kHzV	0.1	0.1
sweepV	0.6	0.4
transition	15.9	7.6
1kHz2	0.2	0.1
10kHz2	0.3	0.4
100Hz2	0.0	0.0
quiet2	8.2	5.0
sweep2	0.4	1.2
3150Hz2	0.2	0.3
1kHzL2	0.3	2.7
sweepL2	0.3	3.5
1kHzR2	0.7	0.1
sweepR2	2.5	0.2
1kHzV2	0.1	0.0
sweepV2	0.4	0.5
Total number of clicks	37.4	28.6

Table 4.21: The average number of clicks per track for pressing one side b.

Track name	left channel	right channel
1kHz	0.1	0.0
10kHz	0.3	0.1
100Hz	0.7	0.1
quiet	2.6	0.8
sweep	0.5	0.2
3150Hz	0.2	0.3
1kHzL	0.3	0.6
sweepL	0.5	1.5
1kHzR	0.5	0.0
sweepR	3.8	0.4
1kHzV	0.3	0.1
sweepV	1.3	0.2
transition	47.2	11.0
1kHz2	0.2	0.1
10kHz2	0.4	0.9
100Hz2	0.1	0.6
quiet2	4.0	2.4
sweep2	1.8	0.6
3150Hz2	0.5	0.5
1kHzL2	1.4	1.2
sweepL2	1.8	5.8
1kHzR2	1.2	0.2
sweepR2	3.6	0.6
1kHzV2	0.1	0.0
sweepV2	1.1	0.2
Total number of clicks	68.5	29.9

Table 4.22: The average number of clicks per track for pressing two side a.

Track name	left channel	right channel
1kHz	0.6	0.3
10kHz	0.5	0.4
100Hz	0.1	0.0
sweep	1.3	0.7
quiet	2.3	2.8
3150Hz	0.7	0.6
1kHzL	1.6	3.7
sweepL	2.6	1.9
1kHzR	1.3	0.1
sweepR	4.1	0.5
1kHzV	0.1	0.0
sweepV	1.2	0.4
transition	16.2	6.4
1kHz2	0.4	0.1
10kHz2	1.1	0.2
100Hz2	0.1	0.6
sweep2	2.0	0.3
quiet2	7.9	6.1
3150Hz2	1.5	0.2
1kHzL2	1.3	1.0
sweepL2	2.2	3.1
1kHzR2	2.2	0.2
sweepR2	8.8	1.0
1kHzV2	1.1	0.0
sweepV2	1.2	0.9
Total number of clicks	62.2	31.4

Table 4.23: The average number of clicks per track for pressing two side b.

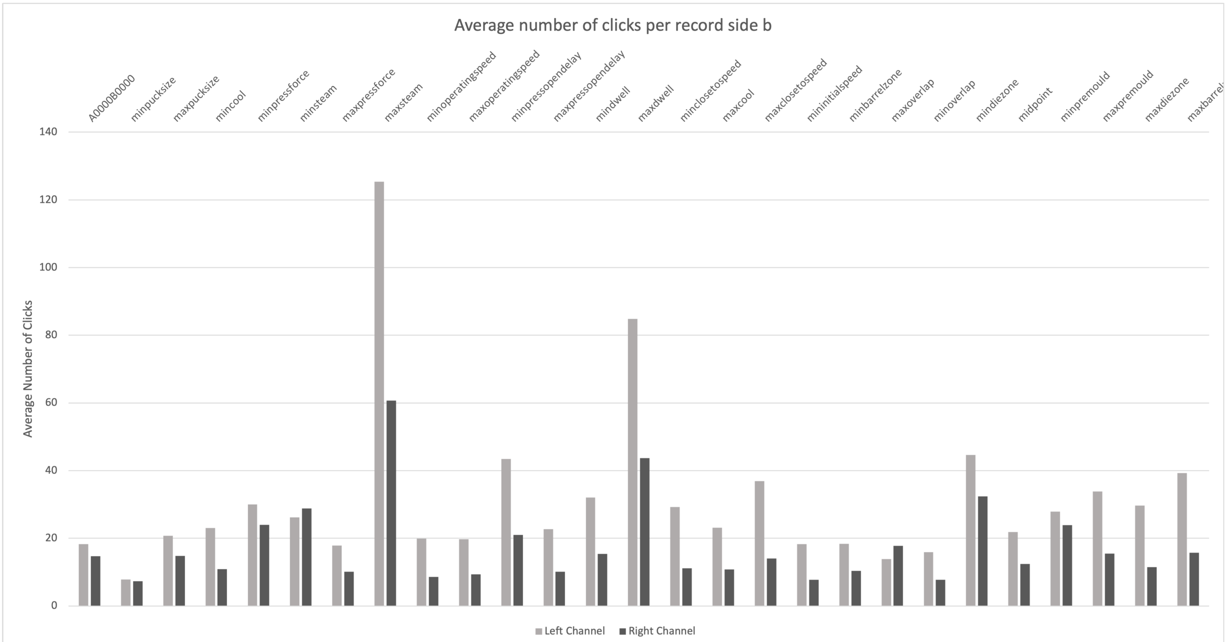

Figure 4.11: The average number of clicks per track for side b of all the different parameter changes recorded.

shows less clicks than the minimum press open delay. This follows the same pattern as we observe in the unweighted surface noise more cooling in the moulds means less pops and clicks on the record. At least for side a, side b shows very little difference between any of the pressings for the number of clicks. The operating speed appears as a middle value here with little effect on the number of clicks.

Maximum press force shows fewer clicks than minimum press force, max close to speed shows fewer clicks than minimum close to speed, suggesting that there is an inverse relationship between press force and speed and the surface noise and number of clicks. As press force increases, you can reduce surface noise however the number of clicks are greatly increased. The maximum barrel zones compared to the minimum barrel zones shows one of the largest increases in number of clicks, similar to the effect that maximum and minimum cooling. The extruder temperatures here affecting the number of clicks is interesting, as the die zone and premould both appear to do very little to change the number of clicks on the record. Hotter barrel zones in this case increase the number of clicks, while a cooler barrel zone reduces the number of clicks.

This is only for side a however, side b shows a different story and in many instances the exact opposite story. Whereas minimum overlap shows a large amount of clicks for side a, it shows some of the lowest amounts of clicks for side b. Additionally whereas minimum steam shows some of the highest clicks for side a, maximum steam shows the highest amount of clicks for side b. This mismatch between parameters and the cause of clicks means that these results should be taken with caution, as results are not mirrored between sides, in fact the second side shows much less variation in the number of clicks and much less clicks overall and often contradicting patterns.

One pattern than can be noticed is time under pressure having an effect on reducing the number of clicks at least for side a. Particularly if this time under pressure doesn't involve steam or dwell. Maximum dwell and steam both produced a lot of clicks on side b however on side a it produces among the lowest. Minimum overlap and steam produced the highest on side a. This general trend of time under pressure reducing number of clicks is seen as minimum cool, steam and overlap all produce a large amount of clicks for side a and so do minimum steam and press open delay produce a lot of clicks for side b. Once again this isn't a hard and fast rule, as the number of clicks on each side do not seem to correlate with one another. For instance, minimum overlap shows among the lowest clicks for side b. Moreso this is simply a general trend noticed in the data and a broad recommendation to operators.

4.5 Wow

The wow doesn't seem to change much with either pressing or parameter change. An expected result, however once again side b appears to have something weird happen it appears as though the stamper has shifted. You can see this in the increase in wow between Appendix Figures B.7 and B.8. It's important to note that in fact two different presses were used for the testing, so perhaps the difference in wow observed here is the result of the difference in the mould alignments between the two presses. That seems the most likely scenario. As expected the inner wow track 3150Hz2 has much more wow than the outer wow track 3150Hz.

Wow can indeed vary on a per record basis. Warpage is known to cause a wow-like modulation of frequency [14] that would be noticed by our scripts. Additionally the size of the centre hole does indeed vary. Some records can be produced with a relatively wide centre hole and some with a relatively tight centre hole. Knowing this, it isn't hard to imagine that some records can be produced with offset centre holes causing a difference in the wow measured. It's impossible to distinguish between these two effects however warpage was independently measured when the records were inspected visually as is reported in Section 4.9

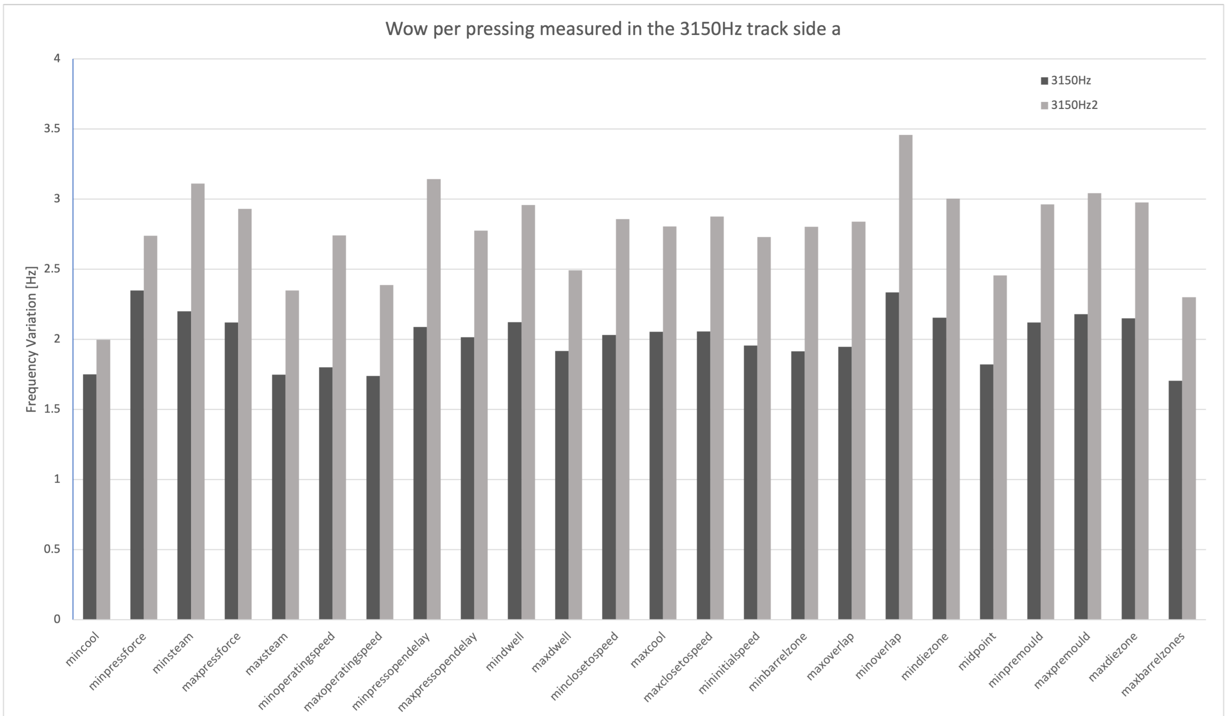

Figure 4.12: The average wow in the first 3150Hz track of each pressing for side a.

Warpage can cause some wow to be introduced however it is not the main cause as that would be the centre hole offset. [15] It's important to note that the size of the centre hole does change record to record and so it is easy to imagine that the centre hole offset too might shift in this way too. However the wow measurements can be lightly interpreted as measuring some form of warpage in the audio. Looking at side a all minimum mould settings appear to cause high wow, minimum overlap, press open delay and steam appear to cause high wow. Looking at side b the opposite is true as maximum dwell, steam and overlap appear to cause high wow while minimum cool, overlap, and press open delay produce a low wow.

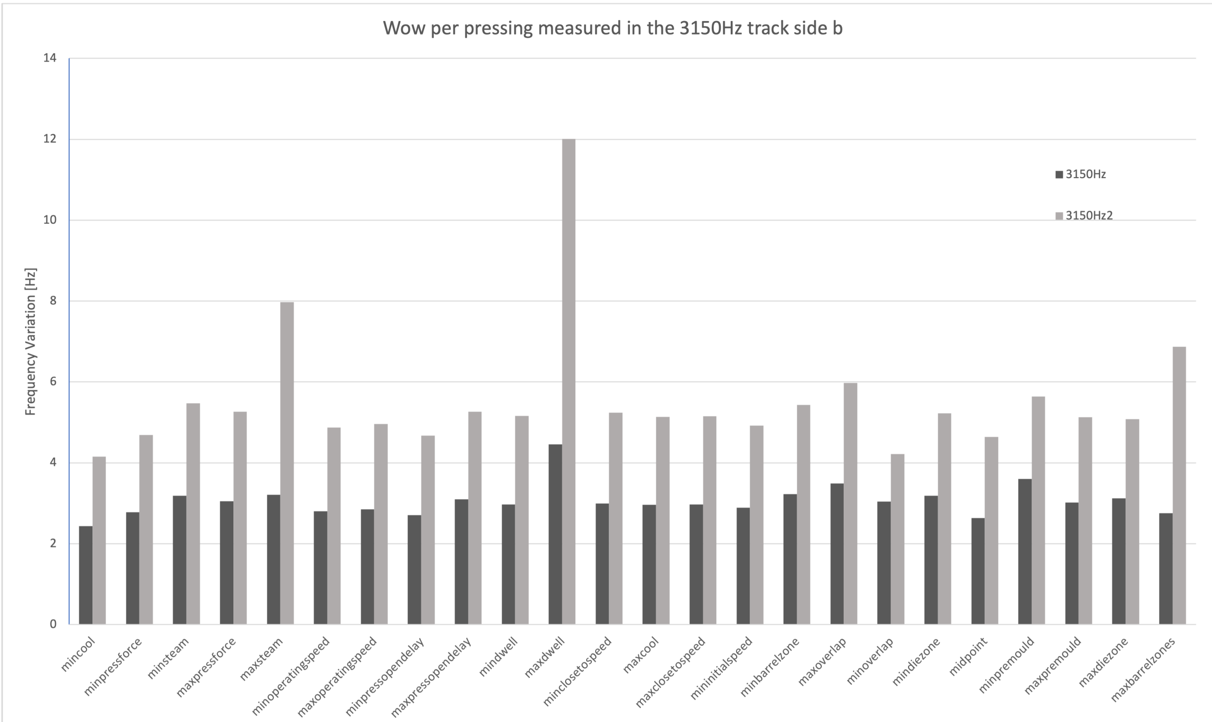

Figure 4.13: The average wow in the first 3150Hz track of each pressing for side b.

4.6 Distortion

There are many places to measure distortion and many tones are included on the test record however for this section the 1 kHz tone will be used for demonstration purposes as it's unclear which settings affect distortion on the record. The distortion measurement is too inconsistent to draw any accurate conclusions from. Figure 4.15 demonstrates the average distortion in the 1 kHz and 1 kHz2 tracks for each pressing. It's clear to see that diameter losses, the stylus moving closer to the centre of the record is the dominant effect in distortion, it's not uncommon to see a +10 dB increase in distortion on the same record simply moving inward along the grooves.

A puzzling observation however, the total harmonic distortion of the inner grooves does not seem to match the distortion of the outer grooves, the two appear unrelated. A record with a low amount of distortion in the outer 1 kHz track can have a high amount of distortion in the inner 1 kHz track. This is independent of the increase in distortion that would be expected as the stylus moves towards the centre of the record. In fact that is often the case that there is a large discrepancy in distortion between tracks.

Distortion is not a simple parameter that can be easily related back to any setting or process in the press. There is too large of a variance between channel, side and track measurements of distortion for the distortion measurement to be of any use. No single measurement of distortion can be reliably associated with a test record in order for a meaningful comparison to be made between test records.

The other possibility of course is that the process by which distortion is embedded on the grooves is complicated and random. Random enough that there is too large of a variance for any of the press parameters to make any difference in how distortion comes about in a vinyl record. Still the diameter losses, the increase in distortion as the stylus moves inward is by far the most dominant effect. There don't seem to be any press parameters that affect distortion to a noticeable degree.

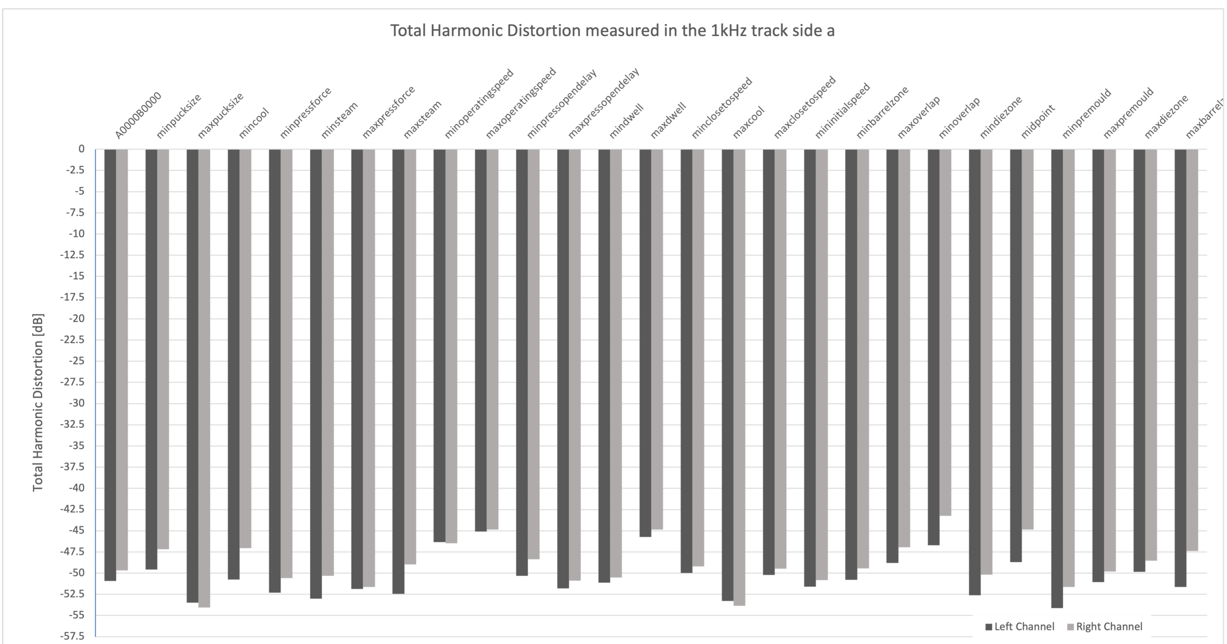

Figure 4.14: A measurement of distortion in the 1kHz track for side a.

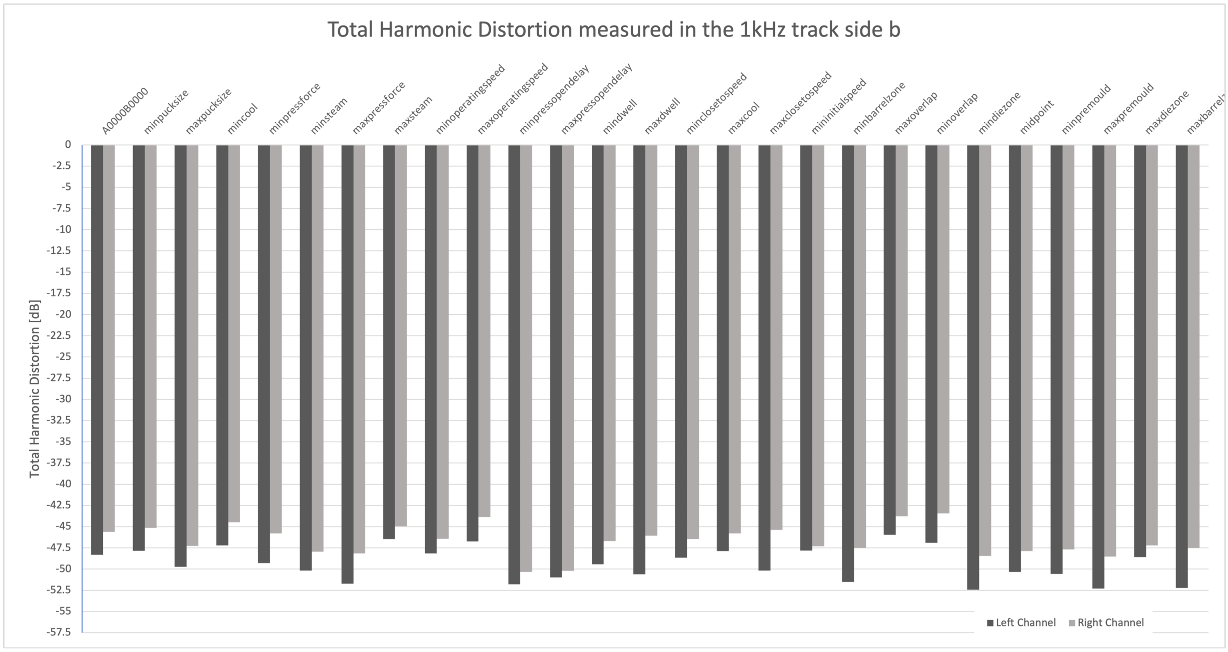

Figure 4.15: A measurement of distortion in the 1kHz track for side b.

4.7 Stereo Bleed

Very surprisingly the stereo bleed shows a very large variation between pressings and records. The stereo bleed measurement varies quite substantially from each pressing parameter change which is a very interesting result. It would be assumed that the stereo bleed, like distortion, is largely a factor of the playback equipment. Here however it's very obvious that the record plays a large role in determining the amount of stereo bleed. An unfortunate oversight, this large variance in the stereo bleed measurement is most likely due to the surface noise levels being comparable to the signal level bleeding into the other channel. To counteract this a narrow band measurement of the stereo bleed would have been more accurate, unfortunately that wasn't done in this instance.

Indeed Figures 4.18 and 4.19 lists the stereo bleed measured in the 1kHzL track for each record and shows a large variance in the stereo bleed per record. This variance is most likely due to the low signal level in the channel and the surface noise interfering with the measurement. Note that the reported stereo bleed values are close to the unweighted noise measurements, suggesting that the noise in this channel is greater than the signal bleeding from the other channel and interfering with the measurement. Stereo bleed intuitively

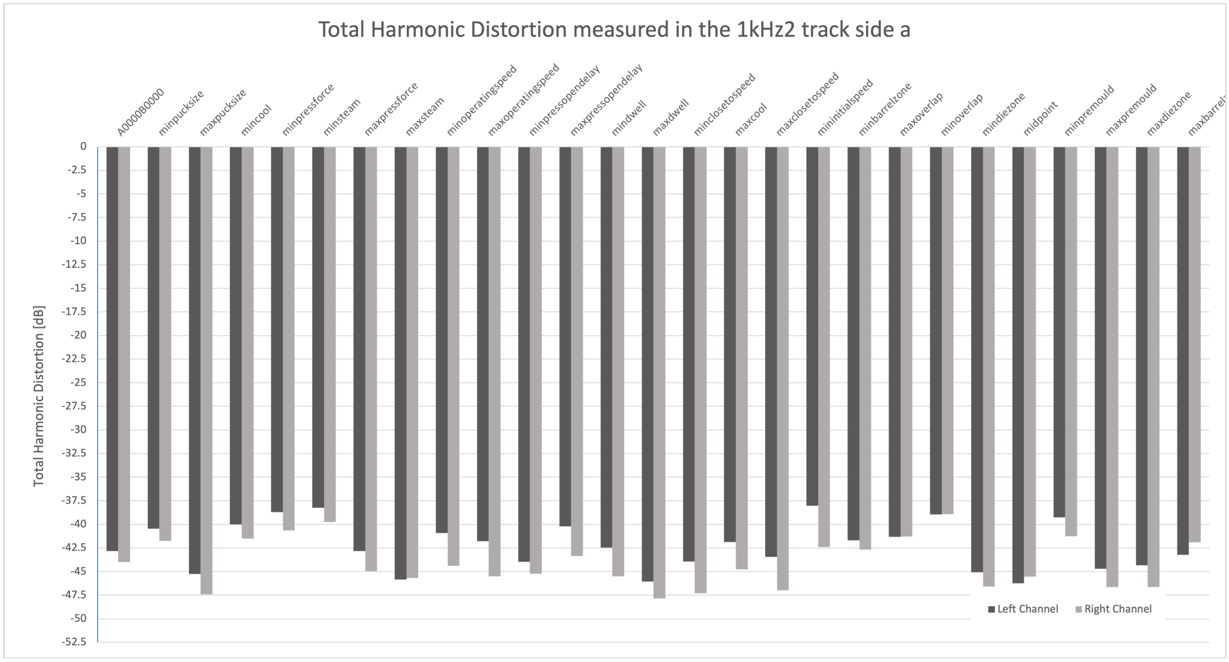

Figure 4.16: A measurement of distortion in the 1kHz2 track for side a.

should be a measurement of quality that would be similar for all records, as all records are formed from the same stamper they should have the same rough groove shape, and since all records were played and recorded on the same playback system then the stereo bleed would be thought to only be determined by the playback system. Indeed, each individual record appears to have its own stereo bleed measurement.

Figures 4.18 - 4.21 show the stereo bleed for two tracks, 1kHzL and 1kHzR for all press parameter changes side a and side b. They aren't sorted nicely like the other Figures as for the most part, stereo bleed appears to be a random distribution of values. There aren't any patterns to look for in the parameter changes either as they too seem largely random. Even when sorted to look like the other figures, the stereo bleed seems to be a wildly varying measurement with no cause.

Again similar to distortion it looks like the process by which stereo bleed is "embedded" into the record's grooves by the press is entirely a random process. There is a large amount of variance and the different pressings while they do change the stereo bleed do not do so in any discernible pattern.

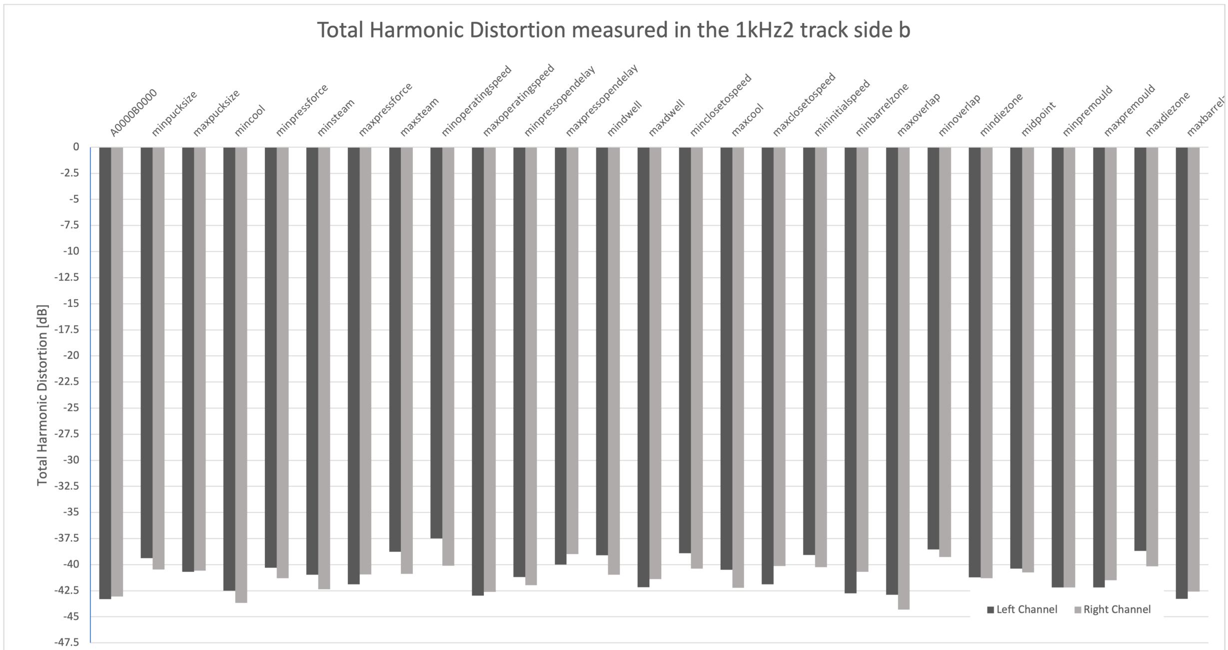

Figure 4.17: A measurement of distortion in the 1kHz2 track for side b.

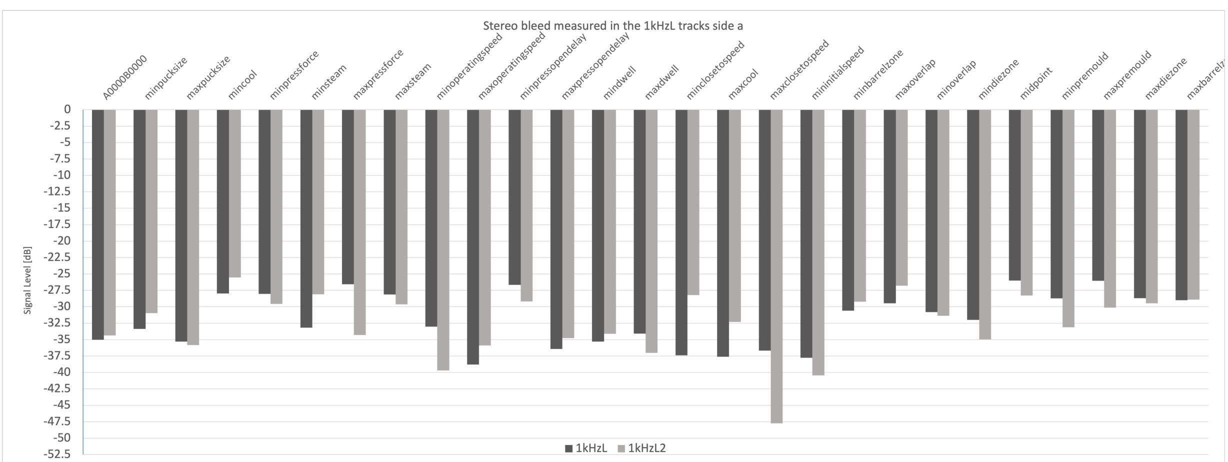

Figure 4.18: A measurement of stereo bleed in the 1kHzL and 1kHzL2 tracks for side a.

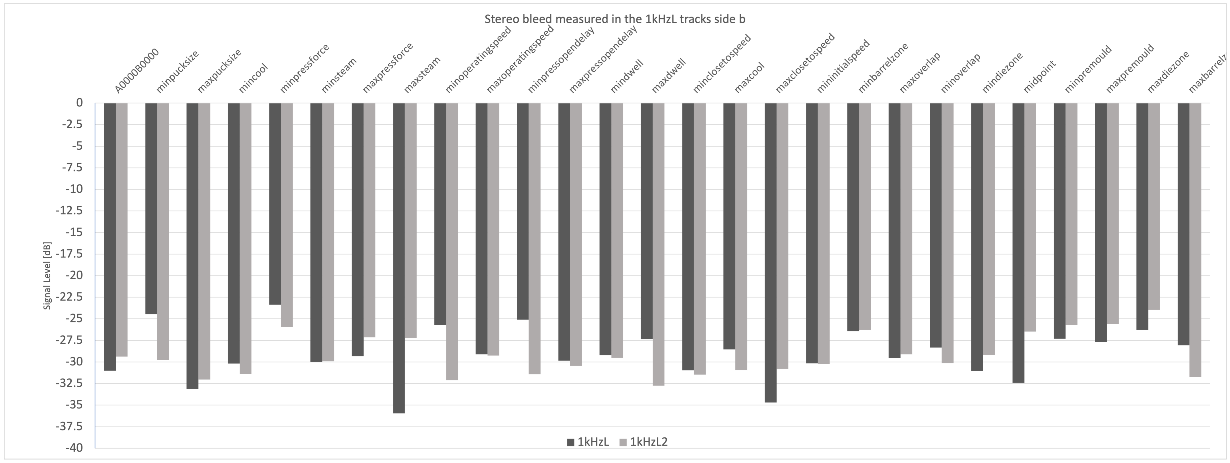

Figure 4.19: A measurement of stereo bleed in the 1kHzL and 1kHzL2 tracks for side b.

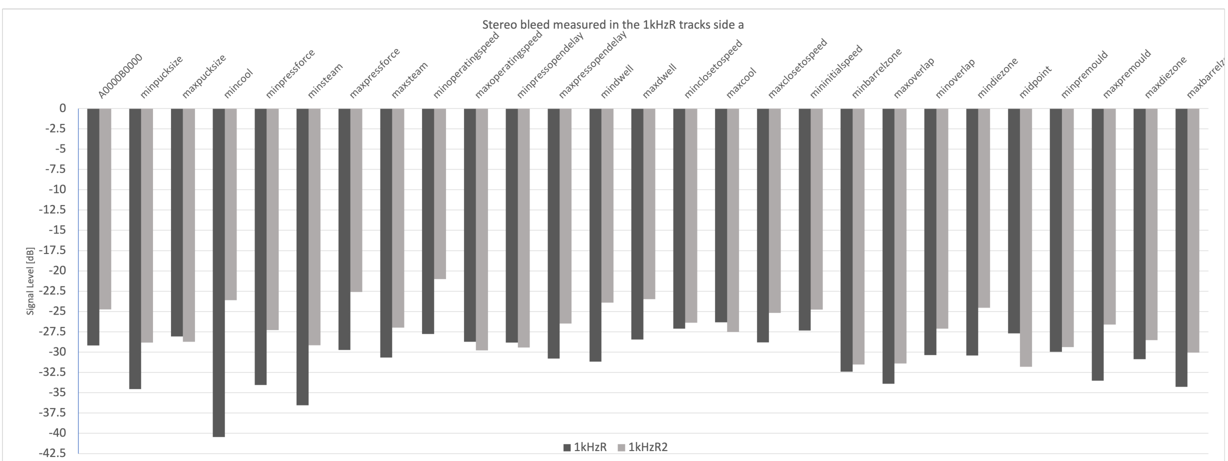

Figure 4.20: A measurement of stereo bleed in the 1kHzR and 1kHzR2 tracks for side a.

4.8 Stitching, Nonfill and Other Defects

As important as the audio quality is the visual quality of the records. There are many visual defects that cause records to not leave the pressing room. While the second set of records were being sleeved and packaged after pressing they went through a manual visual inspection. This visual inspection was looking for a few specific and common visual issues and defects on the records. These defects, including warping and cupping, would cause a record to fail quality control and be considered as a junk record. For the purpose of this

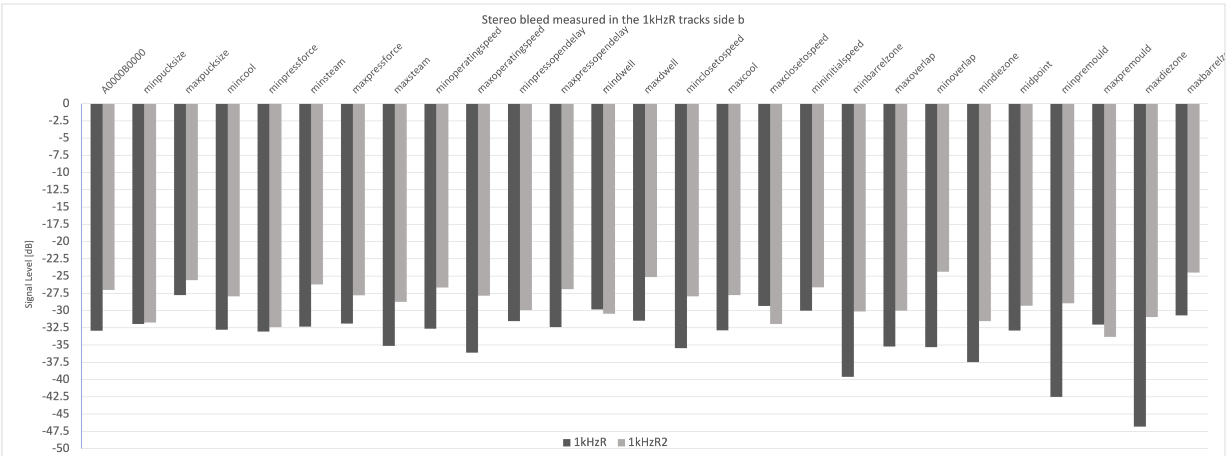

Figure 4.21: A measurement of stereo bleed in the 1kHzR and 1kHzR2 tracks for side b.

project however and with the limited amount of records pressed that was not the case and oftentimes these records had to be kept and recorded instead.

This section details the visual and other defects that were recorded during the manufacture of the test records. While not the focus of this thesis, the visual aspect of record quality is taken into consideration during the manufacture of records. These visual defects such as stitching or nonfill, can also show up in the audio as well as increased noise or other defects such as a pop or click.

There are three main kinds of visual defects that were being tracked during the visual inspection. The first is staining or clouds, a general discolouration in the surface of the vinyl. With black PVC compounds there can be some significant staining that occurs and ruins the pristine black look of the records. Similarly stitching can be seen as a pattern of dots or lines along the black vinyl and isn't to be confused with nonfill, which is a more general term for vinyl not forming proper grooves on the record. Nonfill is typically a discoloured area of vinyl. These visual defects usually carry along with them higher noise and other audio issues as the stylus passes by. Lastly a label issue may have the label cutting into the grooves, not being centred, being doubled up or some other issue. An example of nonfill is shown in Figure 4.22.

Looking at the press parameters that seem to cause visual defects it's interesting to note that running hot extruder temperatures appears to cause staining and clouds across the board. Maximum barrel zones, maximum die zone and maximum premould all show high amounts of staining. What's odd is so too does the minimum die zone pressing, while minimum premould and maximum premould both show less visual defects than

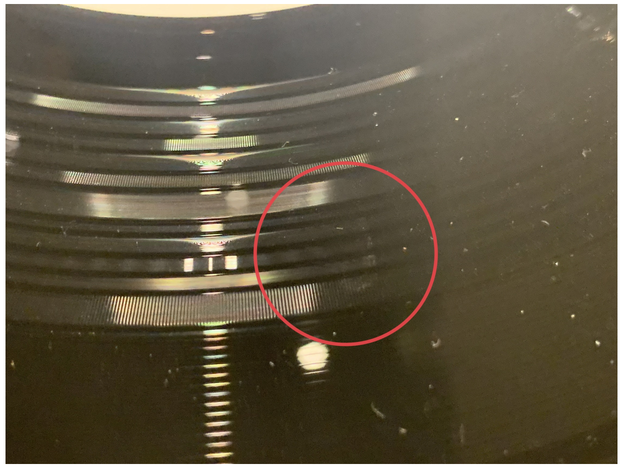

Figure 4.22: An example of nonfill, which is seen as the grey discolouration highlighted by the red circle.

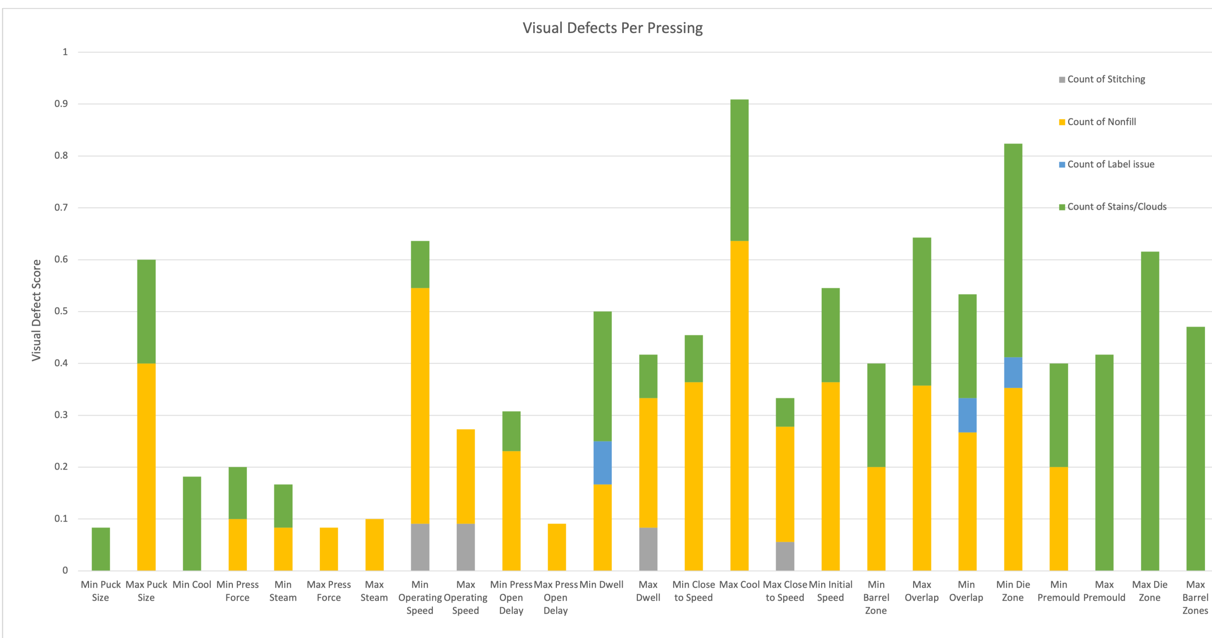

Figure 4.23: A count of the number of visual defects per pressing normalized by number of records in that pressing.

their maximum counterparts, suggesting that too hot an extruder is a cause of stains that operators should be wary of.

Nonfill seemed to be a common defect appearing in some capacity in all parameters. It likely to be caused by mould and press parameters it seems, maximum cool shows the most amount of nonfill followed by both minimum and maximum overlap contrasted with minimum cool showing very few visual defects. It's interesting to note that while maximum cooling shows some of the most visual defects, it shows the among the lowest noise, warpage and numbers of clicks. Perhaps this is a trade off that press operators should be wary of.

Stitching appears to be a rare visual defect, found in only a handful of pressings. From this study no cause of stitching can be determined other than to say it also appears in minimum and maximum operating speed. Two relatively low noise measurements.

Ultimately there aren't many hard and fast conclusions to draw from the distribution to visual defects other than to say that they are not necessarily correlated with worse or better sounding records. This isn't to say that the visual defect itself doesn't have some effect on the audio. It's possible that the defect simply wasn't located in the position where our audio measurements were made or our measurement techniques aren't sensitive to the measuring the sound of such defects. It's also important to note that no pressing was free from visual defects however maximum steam, press open delay and press force alongside minimum puck size all had the fewest defects.

4.9 Warpage and Cupping

Many vinyl listeners know the pains of warped records. Warpage can be caused by improper storage or during the manufacturing process. To prevent warpage or cupping the records are stacked by the trimmer and cooling plates are placed in between each record. In both pressings a cooling plate was placed every 5 records. These cooling plates serve several purposes however they help to manufacture flat records. Despite this however, there is the still the possibility of turning out warped or cupped records by the WarmTone.

Records were given a warpage or cupping score from 0-5 when they were visually inspected and sleeved before transport to the University. This is a largely subjective score however strict care was taken when determining these scores and keeping the measurement consistent across all records. So the results reported here are a reasonable approximation of the amount of warpage and cupping observed when inspected visually. The sum of these warpage and cupping scores are given in Figure 4.24.

Residual elasticity in the compound is known to cause warpage as reported by Ruda.[5] This is interesting when looking at the press parameters that cause high warpage. The maximum operating speed, minimum puck size and minimum initial speed all cause high warpage. What is interesting to note is that maximum operating speed and minimum puck size are extruder parameters, in fact no mould parameters show a high degree of warpage. Operating speed has to do with the speed of the extrusion screw inside of the extruder, a high extrusion screw speed quite possibly supports the hypothesis that it is this residual elasticity in the PVC compound as being quite a large factor in the warpage measured here. A high speed would twist and mix the material more than a low speed, creating a greater difference in the "pressed" versus the "relaxed" shape of the record and cause more warpage as the twisting of the material in the extruder introduces twists in the vinyl puck and eventually the record itself that causes warpage. Minimum puck size causing warpage is a somewhat expected or intuitive result, as there is less material for the vinyl to hold to it's shape after pressing.

Minimum premould, press open delay and die zone all have high cupping. Alongside this so does maximum premould and barrel zones. The minimum premould and minimum die zone having high cupping is once again interesting as those are extruder temperatures, the stages of heating just before the record goes into the press. They are all also minimum temperature parameters, enough to say that too cool of vinyl going into the WarmTone is a possible cause of cupping. It's important to note too that maximum barrel zones also has a high amount of cupping, alongside maximum premould, perhaps an indication that premould and die zone temperatures need to be adjusted alongside the barrel zone temperature to ensure even heating stages.

Once again evidence of this residual elasticity in the PVC compound is a cause of this warp. It's very odd that none of the mould parameters, show up in the visual inspection as warped. Neither maximum or minimum steam, dwell or cool have interesting features to note about how they produce warped or cupped record. Indeed the mould parameters do appear to produce warped and cupped records, just not to the same degree as the extruder parameters seem to.

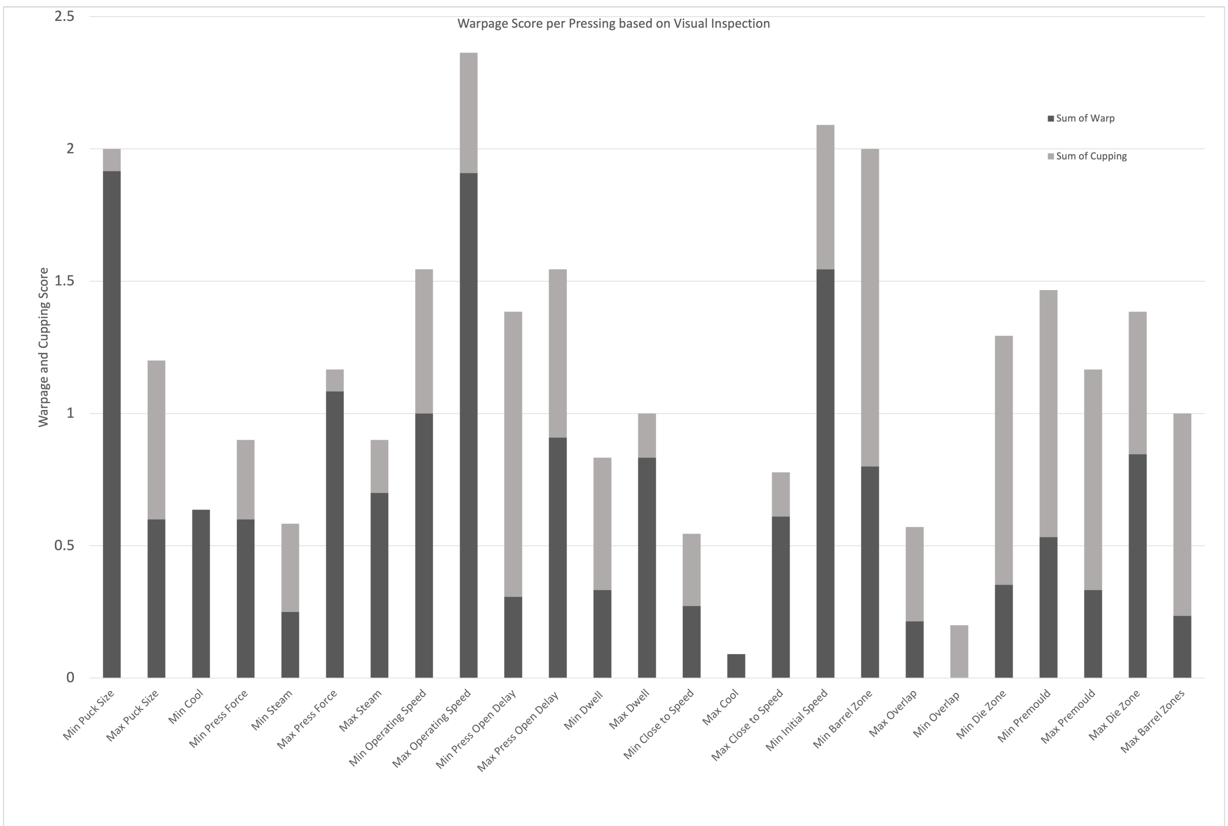

Figure 4.24: The sum of the warpage score per pressing weighted by number of records in that pressing.

Chapter 5

Surface Noise

Surface noise is one of the main contributors to the perceived quality of a record. It is a constant background signal produced as the stylus slides along the groove. Generally unnoticed by the listener while music is playing, it can be heard in the lead in and lead out grooves as well as the quiet portions in between tracks. Measurements of the surface noise gives the signal to noise ratio of a record, an important measure of the recording medium's quality.

The following Chapter explores the nature of this surface noise. Its spectrum, as well as its non-random and periodic nature will be explored. To accomplish this, the spectrum as well as the correlation of vinyl surface noise will be studied. Additionally, there were various timing errors that occurred in the recordings of the records, this will be explored in depth in Section 5.5, as the impact of these errors is most prominently seen in studying the correlation of the surface noise.

An important piece to note about the following section all measurements, unless otherwise stated, have been normalized to a sine wave at 1 kHz with a 7 cm/s peak lateral velocity. Such a signal is described by the NAB standard [8] and has been used in previous studies by John Eargle as a normalization constant when measuring vinyl noise [6]. This is significant as 7 cm/s isn't the maximum signal that is typically cut to a vinyl record, that maximum value is closer to 40 cm/s. As is found on the Shure TTR-103 test record [11] and confirmed by Alexandrovich finding distortion at 50 cm/s [10]. Due to this normalization to 7 cm/s and not 40 cm/s, the reported signal-to-noise measurements in this thesis are around 15.14 dB lower than the true signal-to-noise floor of a typical vinyl record.

5.1 Surface Features

As the stylus passes through a groove, it is modulated back and forth by the groove that represents the audio and smaller variations present on the groove wall's surface. The modulations caused by these small variations make up the majority of surface noise heard. Surface noise is also believed to be proportional to friction and stylus pressure, as testing has shown in Section 5.6, however no clear relationship was shown beyond rumble noise and arm resonance. For the most part, surface noise on a vinyl record is embedded in the grooves. Other factors such as dust and debris also contribute, however their contributions will be shown to be small once we start to look at the coherence of the noise in Section 5.3. The coherence illustrates that the noise on a vinyl record is nearly identical on each playback, and indeed it is also somewhat periodic with the rotation of the record, suggesting that there is a way in which the noise of one segment of the vinyl record is related to noise along another groove segment on the record.

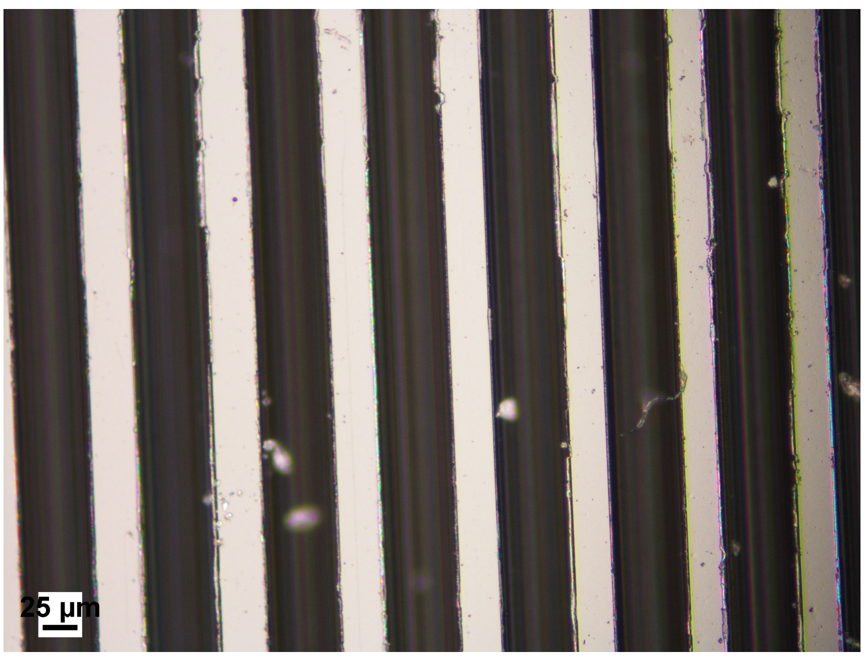

Figure 5.1: Quiet grooves on a record, note the presence of debris as well as the uneven surface of the groove walls—both contribute to the noise heard.

Figures 5.1 and 5.2 illustrate the surface features of both unmodulated and modulated grooves. Note that despite being cleaned with the brush, dust clumps on the scale of about

10 μm are still present in these grooves. These dust clumps are damaging to the record in other ways than simply surface noise. It is static charge that holds the dust to the surface of the record, generally the particles consist of silica, cellulose fibres, carbonaceous material such as soot and other kinds of dirt and contamination. These particles when exposed to the friction of the diamond tip stylus can be ground into the surface of the record and cause defects that create pops and clicks. [16] These pops and clicks will be examined in more depth in Chapter 6.

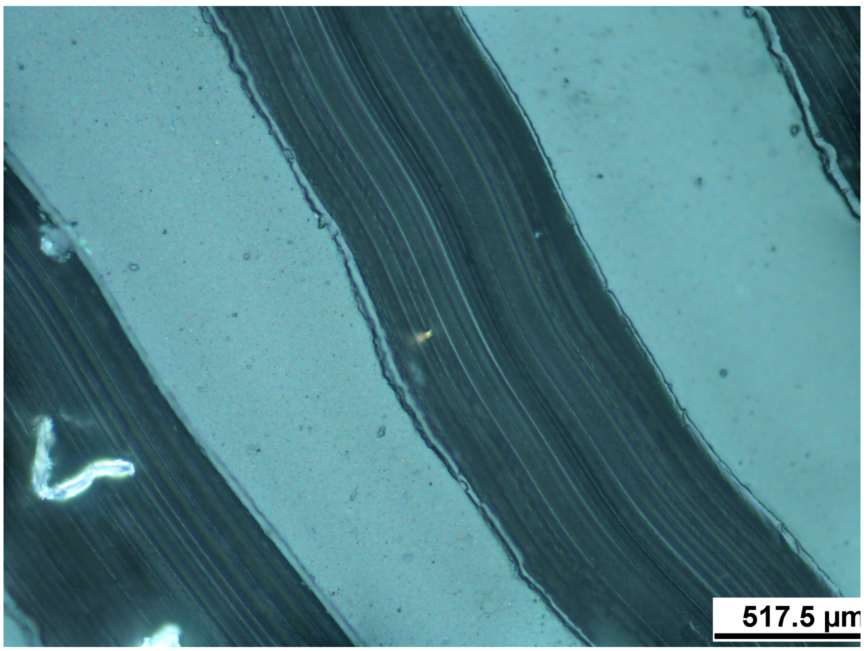

Figure 5.2: A combination focus image of vinyl groove with the microscope. Again, note the uneven surface of the groove wall as well as the debris present in the grooves.

Figure 5.1 is a section of silence taken from the record, these grooves are meant to represent silence on a record, however due to the analog and physical nature of the system these grooves are not entirely "silent" and so are referred to as "unmodulated" or "quiet" grooves instead. These unmodulated grooves illustrate how even quiet portions of vinyl can produce a surprising amount of signal and variation in the stylus path. The groove walls are not perfectly formed, these variations show up in the audio and cause the surface noise that is heard during "quiet" portions of the record.

There are minor variations along both the silent and 1 kHz grooves. These modulations jostle the stylus as it moves through the grooves. Any contributions to the vinyl noise due to this phenomenon are embedded on the grooves themselves and so should be present and identical every time the record is played, assuming there is no groove deformation that occurs. It's important to note that much of the literature and research surrounding groove deformation has been focused on the effect that deformation has on the distortion of the record see [17] and [18]. Simple speculation would lead one to believe that groove deformation and friction leads to a smoother vinyl surface as a record is played more and more– contact from the diamond stylus would wear grooves smooth so less surface noise however higher distortion. This wasn't tested however.

The diagram in Figure 5.3 illustrates the geometry of a stylus sitting in a groove. A groove in a record is V-shaped, according to the NAB standard the groove shape requires an angle of $90° \pm 5°$ [8], the outer groove wall contains the right channel information and inner groove wall contains the left channel information. This allows the motion of the stylus to be stereophonic, in Figure 5.2, the normalization signal is laterally cut as a traditional monophonic signal. Also included on the test record is a 1 kHz tone cut as a "Hill and Dale" recording, vertically. A stereo signal is produced by combining both lateral and vertical stylus motions.[19]

5.2 Spectrum

The spectrum of the noise represents the distribution of power into frequency bins of the time based signal. It takes the audio signal recorded and represents it as components of frequency, to get a sense of what frequencies are present in the signal. The spectrum of the quiet tracks can be taken as the spectrum of the surface noise of the vinyl. Figure 5.4 is an example of a record noise spectrum.

A fundamental characteristic of the surface noise is the presence of the turntable arm resonance and a large amount of rumble noise. The tonearm resonance is at around 20 Hz and was determined based on a dip in the coherence in Section 5.3. The peak at around 12 Hz in the spectrum appears to be an actual peak in the rumble characteristics of the record. Tonearm resonance is determined by the effective mass of the tone arm and the dynamic compliance of the phonographic cartridge and is unique for each turntable-cartridge configuration. [15] The majority of the noise is below 20 Hz, below the audible range as it is intended. This rumble is mostly inaudible. The other prominent feature of the record noise is a hump at around 100 Hz or so.

To represent the audibility of the spectrum A-Weighting and CCIR/ARM-Weighting are used. A-weighting is a very common weighting filter used to help portray the relative loudness of a signal as it is perceived by the human ear. Human hearing has a

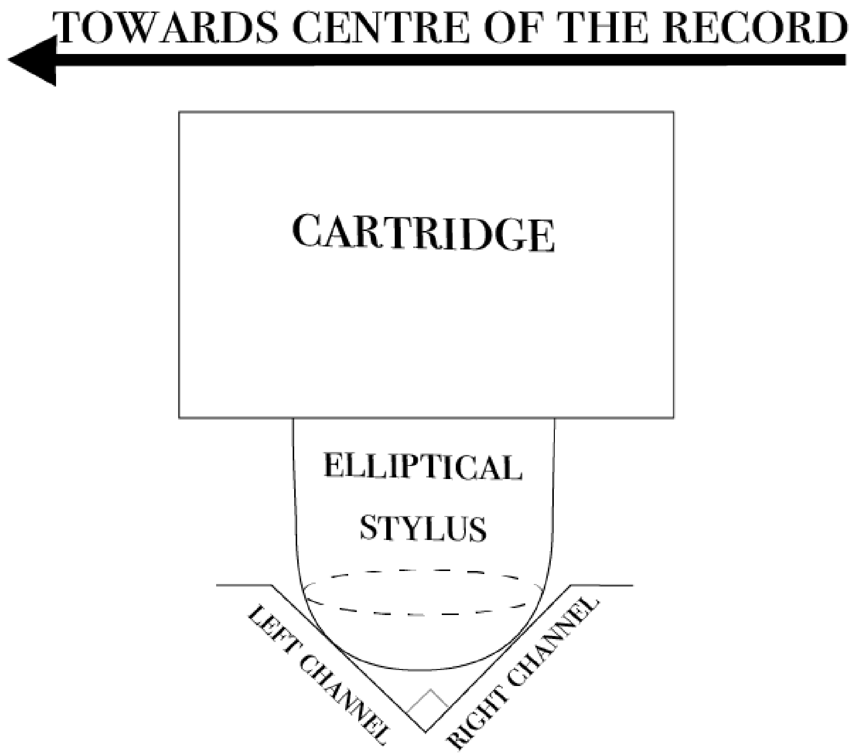

Figure 5.3: A diagram of the stylus sitting in the groove. The wall on the side of the stylus facing the centre of the record contains the left channel information.

specific frequency response, certain frequencies are perceived to be louder than others. A-weighting is very commonly used in the measurement of sound pressure levels to represent this frequency response of the human ear. In this instance the A-weighting is used to represent the audibility of the vinyl surface noise. Alongside A-weighting we also present the the CCIR/ARM-weighting curve. CCIR/ARM noise weighting was developed in 1978 by Dolby Laboratories to better represent the audibility of amplifier, tape and FM noise and so is presented alongside the A-weighting measurements in this thesis as an alternative measurement. [9]

The two weighting curves are given in Figure 5.5, note that both weighting curves achieve the desired effect and drastically reduce the arm resonance in the surface noise spectrum. As both A and CCIR/ARM weightings have hard cutoffs below 20 Hz as such infrasonic frequencies are inaudible. So either of the weighted measurements are typically better representative of the audible surface noise levels.

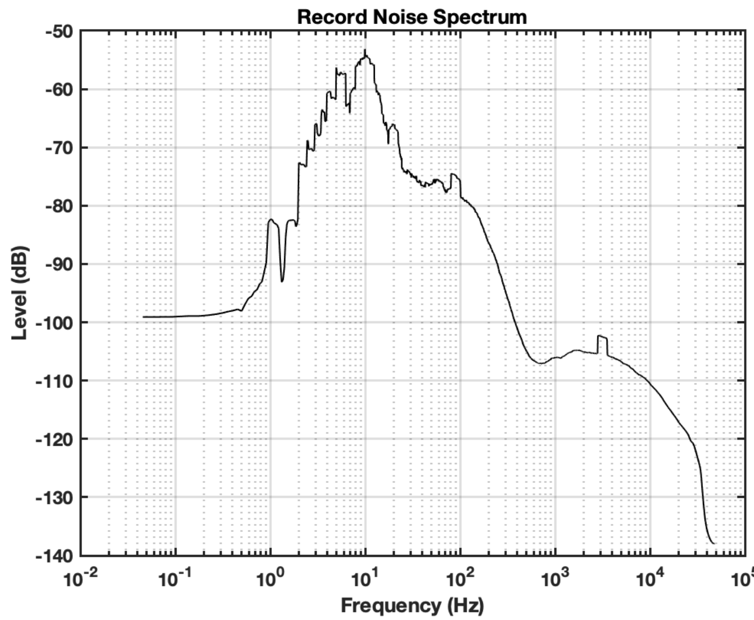

Figure 5.4: An example of a record's noise spectrum. This particular example is from the quiet2 track of record 240 side a in the left channel.

Applying Either A or CCIR/ARM weighting reveals some prominent features of the surface noise. Figure 5.6 reveals two visible humps at 100 Hz and 16 kHz that is visible in the unweighted spectrum also. These humps are significant as they are a distinct feature of the vinyl surface noise, and do not occur in the measurements of the lacquer spectrum.

5.3 Groove Coherence

Coherence is a measurement of how similar the noise in two signals is. We observe a large amount of coherence in various ways on the vinyl records. When looking at large sections of quiet tracks "groove segments" can be divided out in sections of 1.8 seconds, a length of time consistent with the period of the rotation of the record at 33 1/3 RPM. When this is done, the noise follows a periodic pattern consistent with the rotation of the record. This can be easily viewed by plotting each of these groove segments over top of one another as in Figure 5.7.

Dividing up the audio this way allows for each 1.8 second chunk of audio to represent

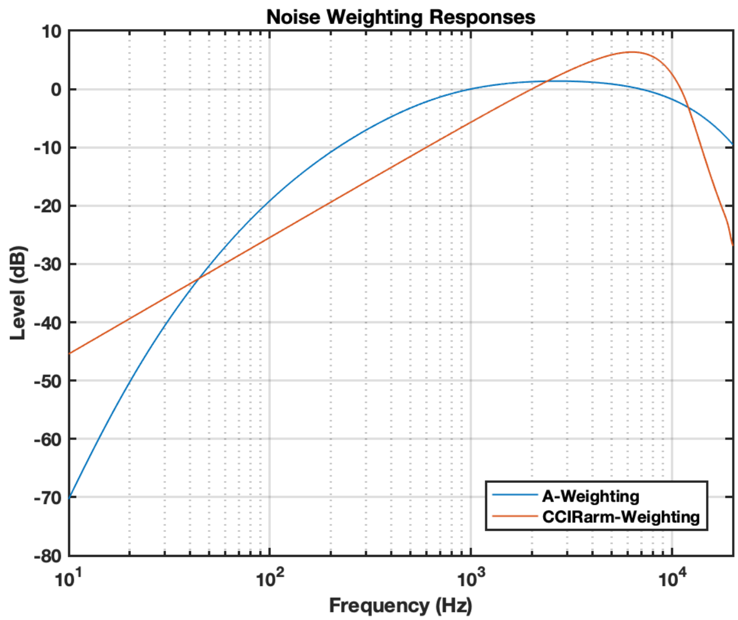

Figure 5.5: The frequency response of the A and CCIR/ARM Weighting filters used in this thesis.

one full rotation of the record. This allows study on consecutive groove segments to probe how the noise in these groove segments are related to one another. In particular different combinations of groove tracks can reveal different features of the playback of the stylus and the vinyl surface beneath it. Taking the 10 groove segments in Figure 5.7 and choosing the first groove segment to be the reference, we can measure the coherence in the following groove segments as compared to this first reference groove segment to see how the coherence falls off as a function of "groove radius". Where groove radius represents the distance from that groove segment to the centre of the record. This is done in Figure 5.9, which shows a clear and steady drop in coherence for each subsequent groove segment when compared to the reference.

Figure 5.10, instead looks at the coherence in each adjacent groove segment. As opposed to choosing a reference groove and measuring the drop in coherence, the coherence between each groove segment compared to the next adjacent groove segment is plotted. When done in this way it is very clear that the relationship between each groove segment is the same, the coherence is approximately 90% between 20-80 Hz for every adjacent groove. This causes the drop in coherence observed in Figure 5.9 as each subsequent groove segment

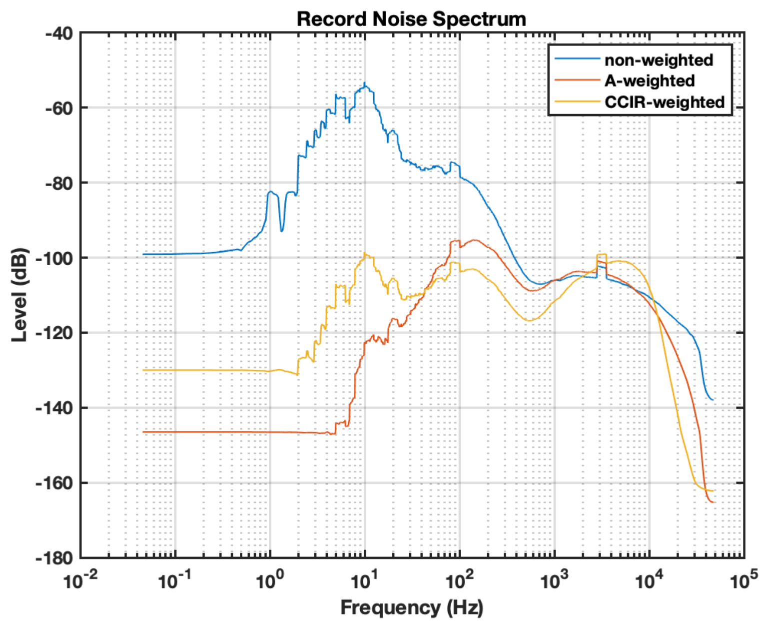

Figure 5.6: The record noise spectrum unweighted and with both A and CCIR/ARM Weightings.

has it's coherence decreased by a constant amount.

Remember back to the geometry of the stylus sitting inside of a groove in Figure 5.3, the outer groove wall represents the right channel and the inner groove wall the left. There are a few interesting stereo measurements of noise coherence. The first is to look at the coherence between the two stereo channels for the same groove segment. Doing so as in Figure 5.11 reveals that a large amount of vinyl noise is common to both channels, roughly 80% between 20-80 Hz. It's also interesting to note a distinct peak in the frequencies that occur at 200 Hz. These humps in coherence in the region around 40-500 Hz match the bump in the record spectrum which peaks right around 100 Hz examined in Section 5.2. It's interesting to note that there is an additional hump at the very upper frequencies, around 30 kHz. This is possibly a resonance within the stylus needle itself as it sits in the groove between the two channels.

These humps and the coherence found in these groove segments can be shown to be a feature of the vinyl surface noise and not a feature of any other aspect of the playback system as this coherence pattern cannot be replicated on recordings made using a lacquer. The coherence, or rather lack of coherence found on the lacquer and the relevant Figures

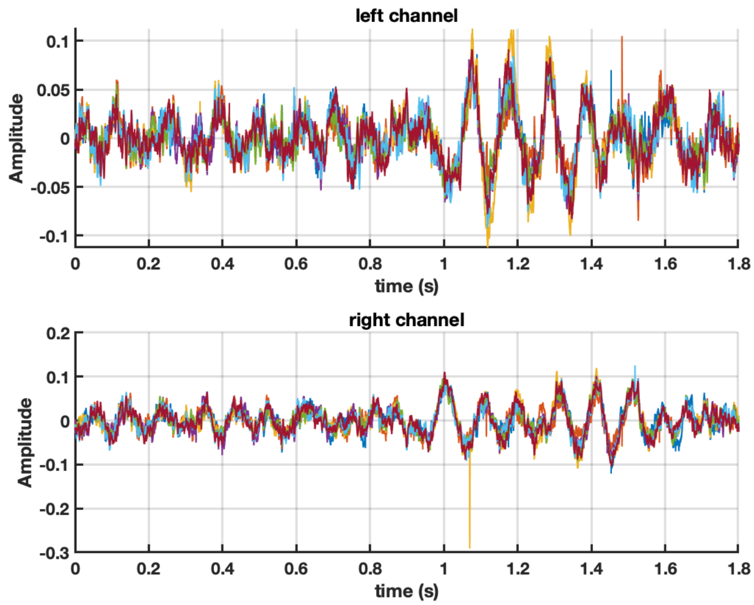

Figure 5.7: Multiple adjacent groove segments plotted on top of one another. There are 10 total. Note how similar the audio waveforms are.

can be found in Chapter 5. This points to noise being added at some point during the vinyl process between the lacquer being cut and the record being pressed. As it stands, these humps seem to be a characteristic of only vinyl surface noise and so should be an area of focus when trying to reduce the noise on the vinyl record.

As discussed in Chapter 4, the differences in plastic caused the largest difference in audio quality. The surface noise measurements for the quiet track of each pressing is given in Tables 5.1 and 5.2 for side a and side b of both pressings respectively in either the left or the right channel. Note that the left channel is consistently around 1 dB lower in noise than the right channel. This appears to be a feature of the record itself and not due to any anti-skate correction. The noise is consistently lower in the left channel as compared to the right channel.

It's interesting to see the behaviour of the A-weighting and the CCIR/ARM-weighting curves in action in these measurements of surface noise. As noted, there are two distinct humps in the spectra of vinyl records, a band around 80 Hz and another around 2000 Hz. It's interesting to note that the CCIR/ARM weighting filter accentuates the second hump around 2000 Hz much more than the first, whereas the A weighting seems to fit

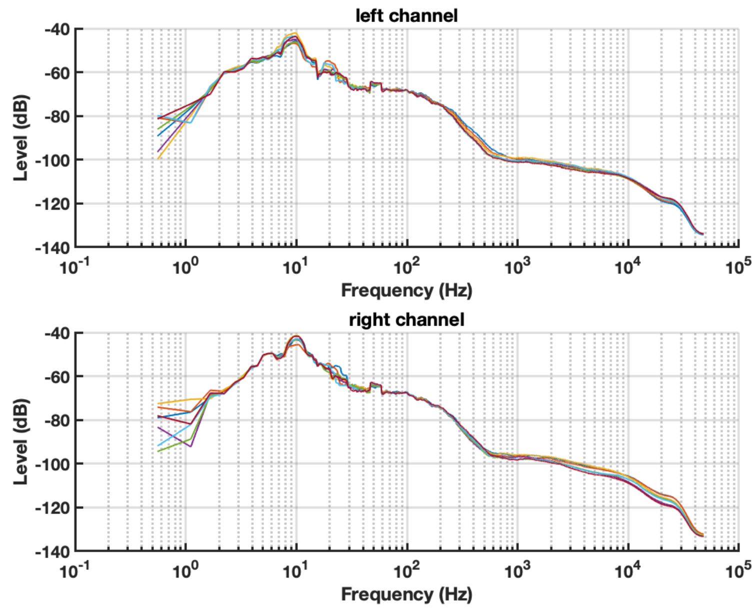

Figure 5.8: Multiple adjacent groove segment's spectra plotted on top of one another. There are 10 total.

perfectly around the first hump in a way. This difference causes a drastic difference in A weighting surface noise, in the second pressing there is much more surface noise in the A weighting measurement than the CCIR/ARM weighting measurement. This demonstrates the difference between Plastic A and Plastic B compounds as the B compound used in the second pressing has much more mid-frequency noise as compared to the Plastic A's high frequency noise.

	Non-Weighted	A Weighting	CCIR/ARM Weighting
First Pressing	-30.9 dB, -29.2 dB	-53.8 dB, -52.2 dB	-55.5 dB, -53.4 dB
Second Pressing	-32.5 dB, -31.3 dB	-46.5 dB, -45.9 dB	-52.9 dB, -51.7 dB

Table 5.1: The surface noise levels for the quiet tracks for side a.

Of course this difference can only be spotted when taking the average of all measurements for all the recorded test records. When examining individual spectra it's very hard to notice a difference between either press run. In Figure 5.12, the spectra of both Plastic A and Plastic B records are shown alongside the A and CCIR/ARM weighted spectra. It's very hard to spot a difference in the spectra visually with only a single sample however it

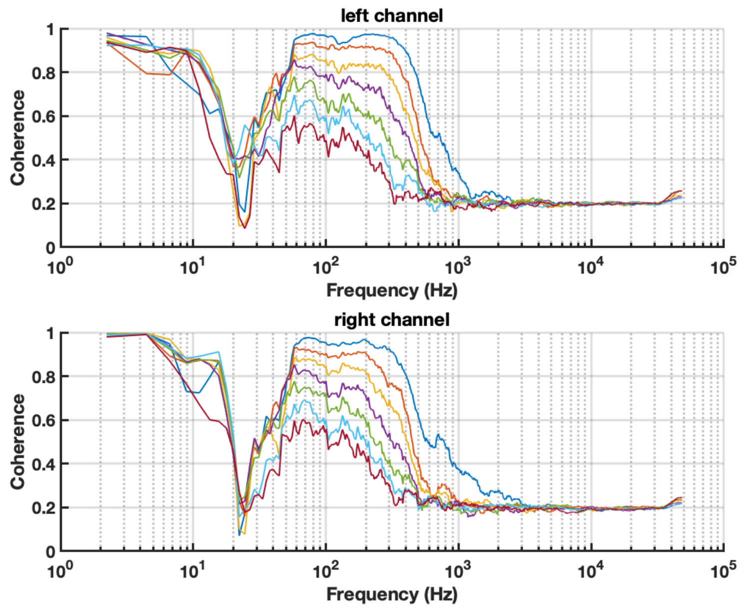

Figure 5.9: The coherence of nine subsequent groove segments as compared to a reference groove segment.

	Non-Weighted	A Weighting	CCIR/ARM Weighting
First Pressing	-30.1 dB, -29.0 dB	-52.9 dB, -51.1 dB	-55.4 dB, -52.9 dB
Second Pressing	-31.3 dB, -30.2 dB	-46.3 dB, -45.8 dB	-52.9 dB, -51.7 dB

Table 5.2: The surface noise levels for the quiet tracks for side b.

is clear that the Plastic B has less low and mid frequency noise than the Plastic A.

The differences in plastics goes further, we can examine a difference in the coherence patterns of the two vinyl compounds surface noise. Figure 5.13 shows the difference in coherence between adjacent groove segments in both pressings. The two coherence patterns while similar, show differences that indicate that the plastic plays some role in determining the surface noise and any coherence observed. This difference between the two plastics is apparent in Figure 5.14, a plot of the coherence between the left and right channels of the same groove segments for the two kinds of plastic. The Plastic A and Plastic B share a similar shape to their stereo coherence patterns, however Plastic A's shape is much more drastic.

One very important track on the test records is the transition track. This is an area

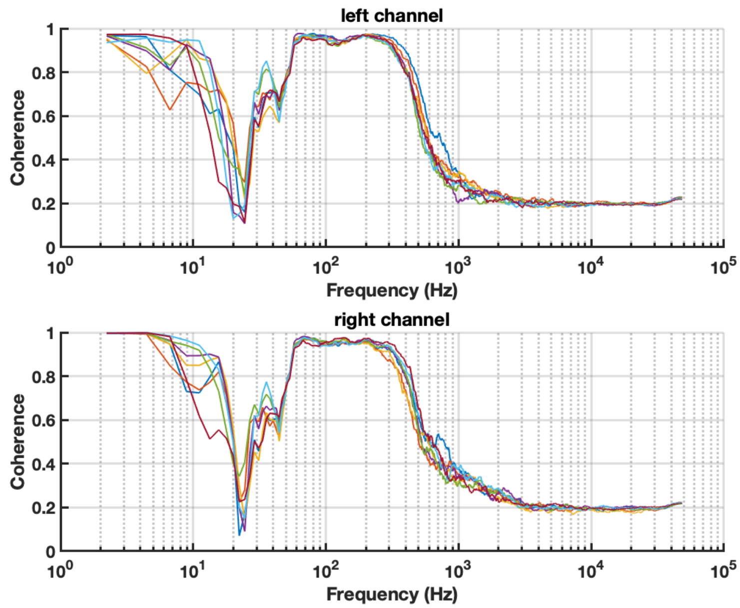

Figure 5.10: The coherence of each groove segment as compared to the previous groove segment.

of extended silence that connects the two sets of signals on the record. What's important about this track, and what separates it from the other quiet tracks on the record is that it contains both closely spaced groove segments and groove segments that are spaced much farther apart. In the traditional quiet track, the grooves are packed tightly together as they would be if they contained signal. The transition track is unique as it covers a large distance on the record as it transitions between the two sets of signals on the test record. This transition contains a section of quiet grooves packed together as they are in the quiet track then a section of quiet tracks that are spaced much farther apart as the needle travels across the record to the second set of signals.

Crucially the coherence of the vinyl noise that is observed between adjacent groove segments isn't present on the farther spaced grooves. Figure 5.15 demonstrates the drop in coherence that is observed across the farther spaced groove segments. The first two segments plotted in Figure 5.15 were closely spaced, the other groove segments plotted were farther spaced and show almost no coherence pattern. This coherence and it's drop-off when grooves are spaced farther apart suggests that vinyl surface noise is locally related. There appears to be a "region of coherence" in the surface noise, a certain surface area on

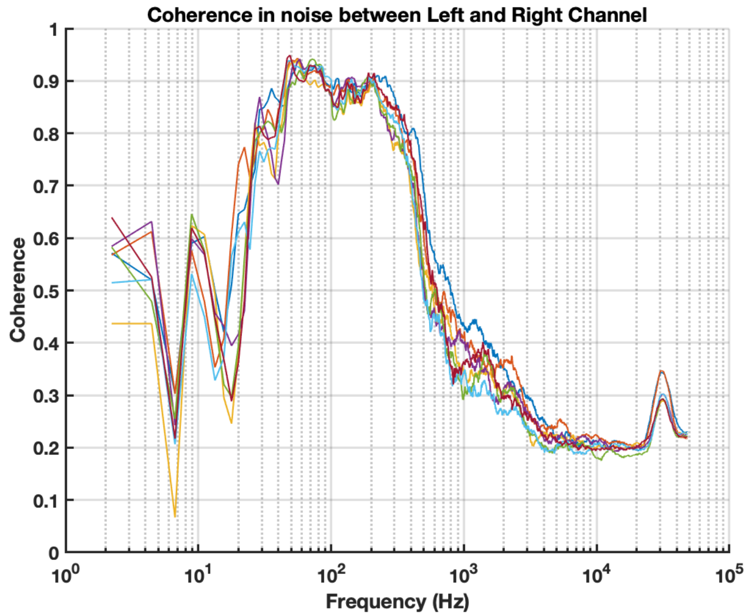

Figure 5.11: The coherence between the left and right channel for each groove segment.

the vinyl record where the noise shows this correlation and coherence.

As it stands it is unknown what is causing the coherence observed in the vinyl surface noise, however the existence of a coherence pattern in vinyl noise points to a possibility of the noise being targeted and reduced in the future. The surface noise seems to be locally related on vinyl records in a way that isn't observed in the lacquer. Considering measurements in this study already found that vinyl records can have noise floors lower than lacquers, perhaps this coherence points to an avenue of exploration where the process that creates a rough groove edge can be studied and improved upon. During testing there were a few hypotheses as to what could be causing this coherence, one belief was that the PVC is loaded into the press as pellets and melted by the extruder. Perhaps uneven heating in the extrusion process left the pellets somewhat intact in the extruded vinyl as local temperature variations, which causes the radial coherence pattern observed. Perhaps the coherence pattern is formed as the vinyl cools and relaxes into its final shape after it has been pressed.

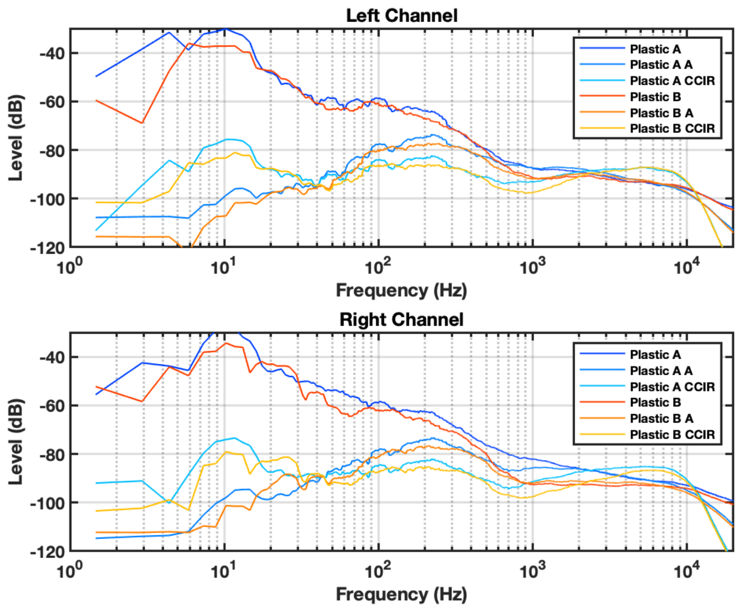

Figure 5.12: The unweighted, A-weighted and CCIR/ARM-weighted spectra for both press runs.

5.4 Vinyl pellet hypothesis

There is a hypothesis being put forward about the nature of the surface noise on vinyl records in this thesis. As reported in Chapter 4, the difference in PVC compounds causes by far the largest change in audio quality. There are a few characteristics in the coherence of the audio and the press parameters that affect the final audio quality to suggest that the size and shape of the vinyl pellets loaded into the extruder are important in determining record sound quality.

Suppose V_p represents the volume of the pellet and the record has a thickness of T, then as the record is pressed each pellet spreads out an area given by V_p/T. If the pattern of pellets is simplified to a square array then the distance between pellet centres is $\sqrt{V_p/T}$. If the record is playing and the stylus is a distance R from the centre, ω is the angular rotation rate of the platter and the frequency at which pellet centres pass the stylus will be given by the following equation.

$$f = \omega R / \sqrt{V_p/T}$$

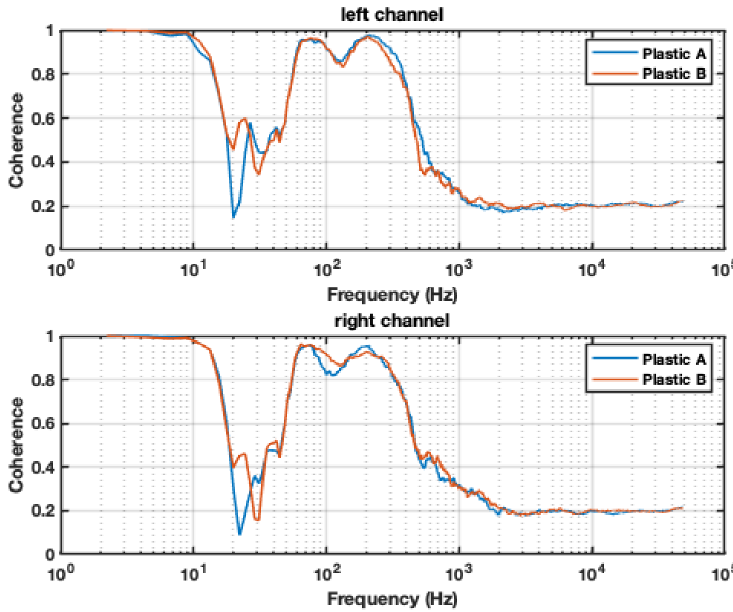

Figure 5.13: The coherence patterns of adjacent groove segments for pressing one done with Plastic A and pressing two done with Plastic B.

Plugging in real world values, a typical pellet has a volume of 0.001 cm^3, a record has a thickness of 2 mm and so the frequency f would have a value of 178.9 Hz for the first quiet track at 23 cm, 163.3 Hz for the transition track starting around 21.0 cm and 141.6 Hz for the quiet2 track at 18.2 cm. These values are much too high to explain the low frequency rumble noise however they are just around the right range of values to describe the coherence in the surface noise and the bump in the record frequency as compared to the lacquer. The bump in the surface noise and range of maximum coherence is between 60-130 Hz, around the same scale size as a pellet.

Figure 5.16 is a microscope image of the extruded vinyl "snake", a long piece of vinyl that leaves the extruder and is typically formed into the vinyl puck in the premould. This vinyl "snake" was taken from the LiteTone, the semi-automatic version of the WarmTone. It's provided here to give an idea of how rough the surface of the vinyl really is before it's been pressed into a record and to give a sense of the structure of the vinyl as it leaves the extruder. It's been reported that residual elasticity in the vinyl materials can be accounted for warpage in the records [5], perhaps on the micro scale the vinyl is returning to the kind of bumpy shape seen in Figure 5.16 and this relaxation on the micro scale is causing the

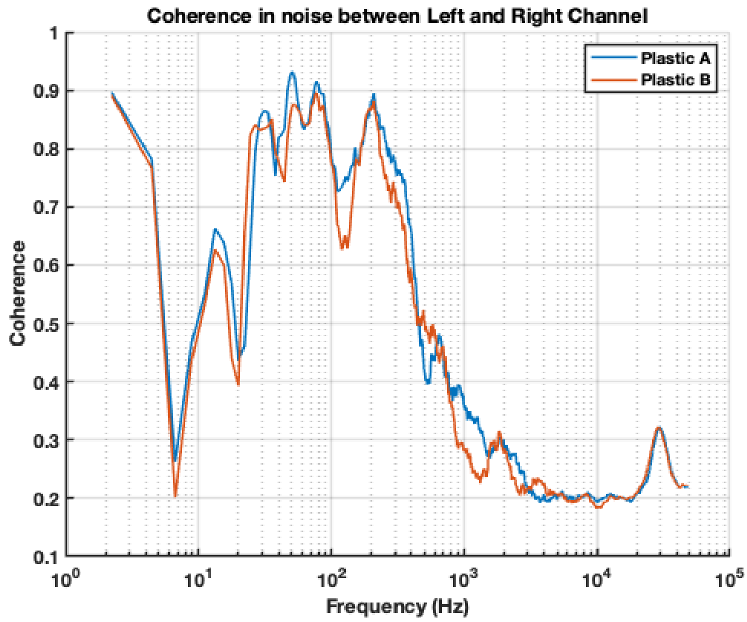

Figure 5.14: The coherence patterns of the left and right channels of the same groove segments for pressing one done with Plastic A and for pressing two done with Plastic B.

surface noise observed in the 60-130 Hz range and the resulting coherence.

5.5 Random noise and timing errors

Looking at different recordings of the same record will measure how similar the noise is on subsequent playbacks. This will give insight into how much noise is embedded on the record itself– rather than due to features of the playback. One would hope that the noise is identical each time a record is played– however with vinyl that is not the case. Recording a record multiple times and then taking the coherence of those recordings compared to a reference recording helps answer the question of how much noise on the record is random, and how much is imprinted on the record. To test this five recordings were made of the same record. Figure 5.17 shows the spectra and Figure 5.18 shows the coherence of these multiple recordings.

As shown the noise is not identical each time a record is played, rather it has slight variations, particularly in the upper frequencies. The coherence appears to be unity from

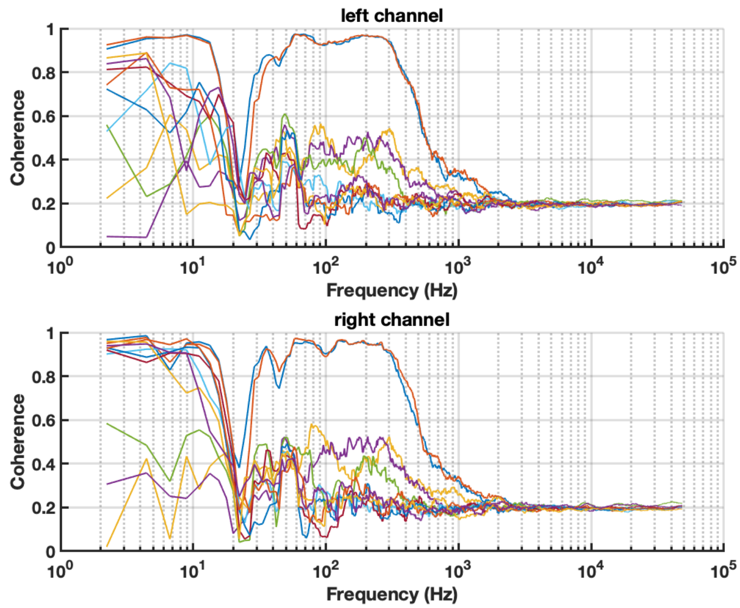

Figure 5.15: The coherence of the closely spaced and farther space groove segments as they move along the transition track.

20-100 Hz, where it then drops to 95% between 100-1000 Hz where it then sees a slight rolling off. Interesting to note is the dip in coherence at approximately 12 Hz and 18 Hz. The 12 Hz dip in frequency corresponds to the peak seen in the noise spectrum in Figure 5.17.

Recording 3 appears to be an outlier. The measurement of coherence is very sensitive to how the recordings are lined up. Any slight variation or offset can cause a drop in coherence, especially in the high frequencies. One discovery made during this thesis was the fact that there were a plethora of timing errors in the recordings of the records. As discussed in Section 3.4, great pains were taken to reduce these timing errors however nothing could be done to eliminate them completely. The effect of these timing errors and sample dropouts can be seen in Recording 3, as this has thrown off the coherence measurement of this recording.

What this test shows is that the majority of vinyl noise is identical each time the record is played, with only about 5% of the noise above 100 Hz to be different on subsequent playbacks of the record. For the most part, surface noise is related to the surface features and is embedded onto the record. The random noise is most likely to come from the

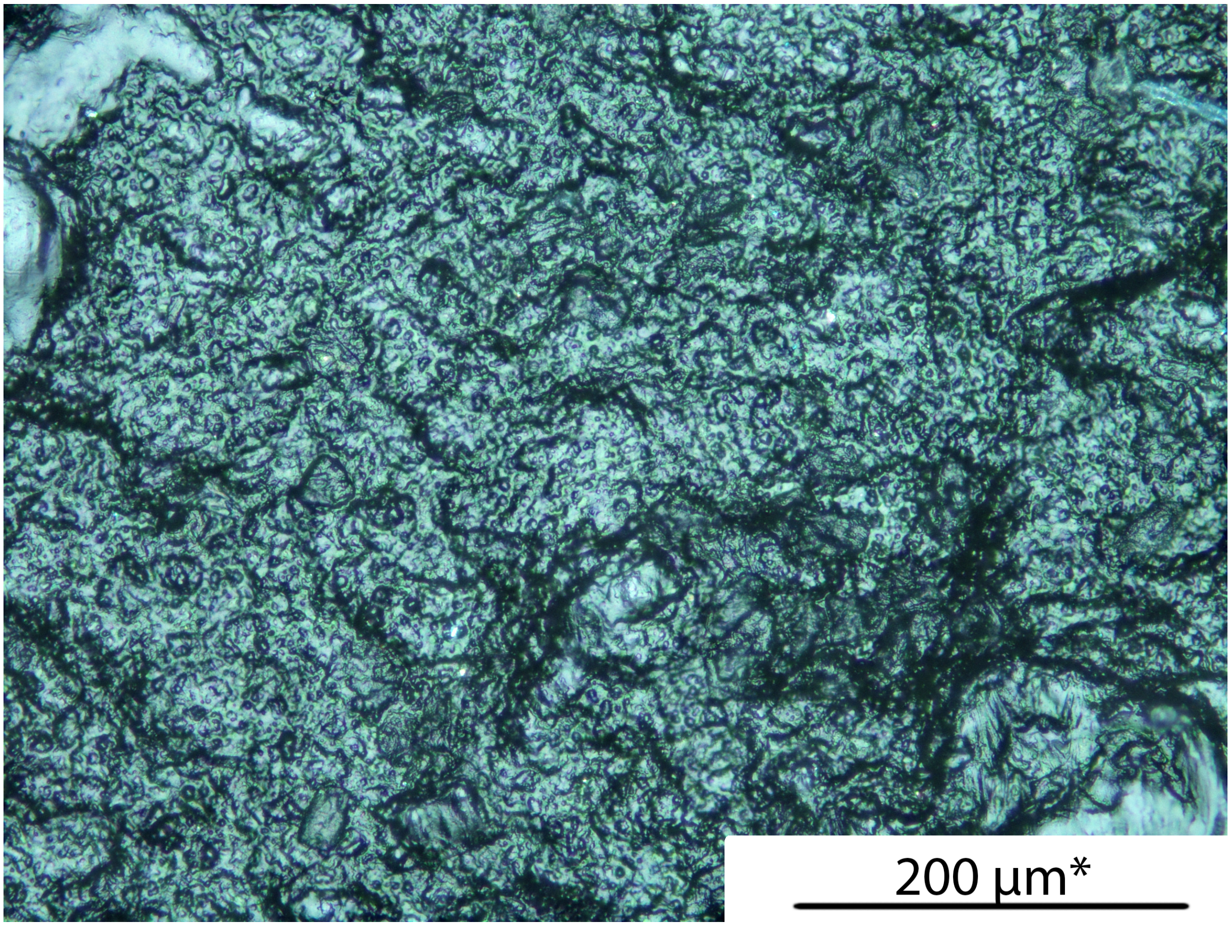

Figure 5.16: An image of the surface of extruded vinyl taken with a microscope. Perhaps the ridges in this image are the pellets taking shape within the extruded vinyl. *suggested scale only

electronic components and cables used in the recording setup as well as the random friction between the stylus and the groove.

5.6 Friction and stylus forces

The act of running a stylus through a groove at speed involves incredible friction, enough to cause plastic deformation in the groove and wearing on the diamond-tipped stylus. This

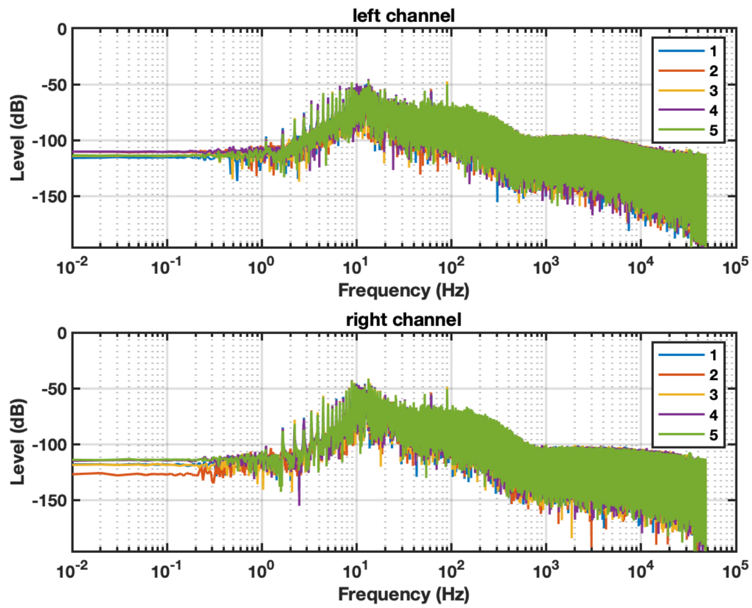

Figure 5.17: The spectrum of five recordings of the same record.

friction was found to correlate with levels of surface noise heard on playback. [20] On a turntable there are various adjustments that can be made to both the stylus tracking force and what is called the anti-skate force. Both of these forces can be set by the user, and on the Technics. Tracking force is typically measured in grams and in the case of the Technics Sl-1200 the anti-skate force is simply a setting that you match based on the tracking force that was set. In this thesis, the tracking force was set to 1.8 g and the anti-skate at the corresponding 1.8 setting– which is the recommended tracking force for the Ortofon Blue cartridge.

Early measurements have indicated that the surface noise is proportional to the force of friction between the stylus and the groove. On the Technics Sl-1200 and most turntables, this force of friction can be set by the needle tracking force and anti-skate forces—both of which are able to set by the user. The needle tracking force is typically set by a counterweight, on the opposite end of the tone arm. Typical tracking forces range from 1-4g. Many audiophiles and vinyl connoisseurs try to use as little a tracking force as possible, so as to reduce stylus and record wear.

The question of stylus force has minimal bearing on the measurements made here, as we are concerned with the record itself not the playback equipment. However it is

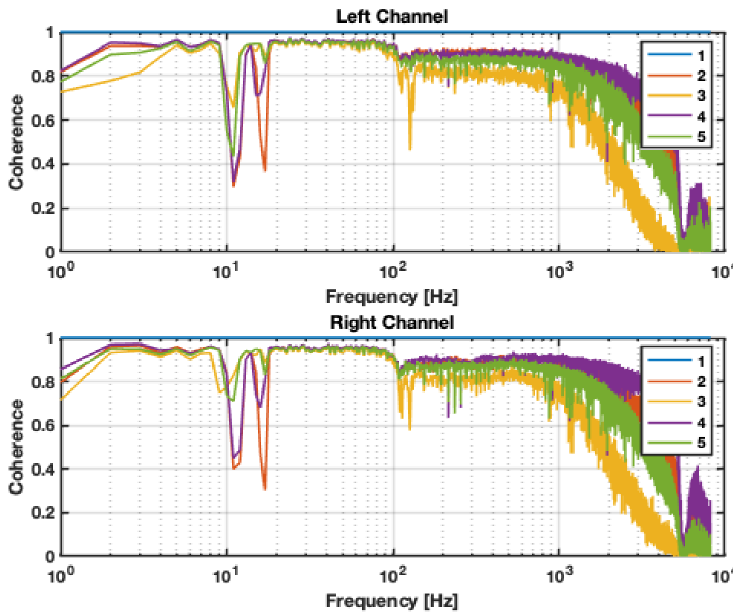

Figure 5.18: Coherence of multiple recordings, note that recording 1 is the reference recording and its coherence is unity.

interesting to study regardless. To test what a change in tracking force does to the noise of a record, the quiet track of one of the records was recorded at different tracking forces. These measurements are normalized to the 7 cm/s peak 1 kHz tone recorded at a stylus force of 1.8 grams. Figures 5.19 and 5.20 show the RMS levels measured at various stylus forces. To obtain these measurements, the stylus force was increased in increments of half a gram, the anti-skate force was set to the appropriate setting for each stylus force and the quiet track of the record was recorded three times so that the mean of those measurements could be taken.

In Figure 5.19 a clear relationship is shown between the unweighted RMS noise and the stylus tracking force. Noise can be seen increasing in both channels, however the more drastic increase is in the left channel. A weighting the results by contrast shows almost no relationship between stylus force and RMS noise, a very strange result. Increasing the stylus tracking force seemed to reliably increase the rumble and arm resonance heard, however it has seemingly no effect on the A-weighted noise. The A-weighting is the more important of the two measurements as it more prominently represents the two "humps" characteristic in vinyl surface noise.

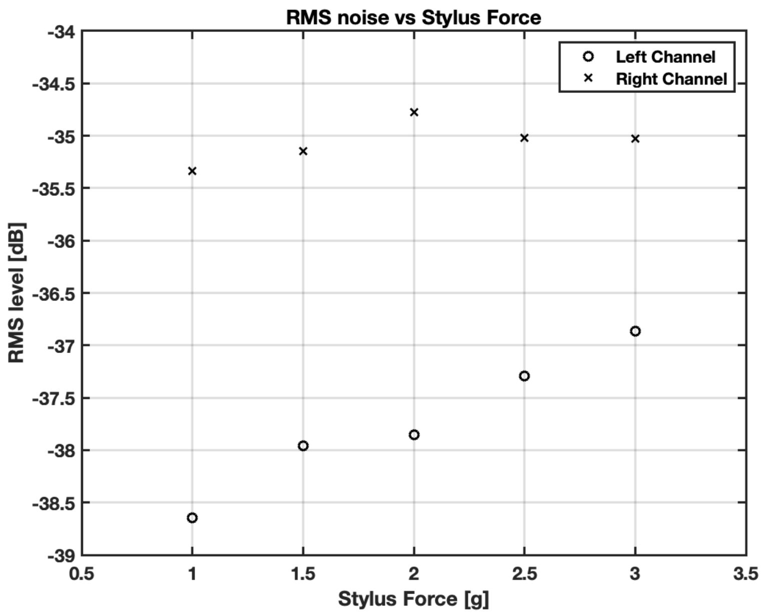

Figure 5.19: RMS level of a quiet track plotted against stylus force.

For the most part the issue of unpredictable or increased noise due to the stylus force is bypassed as all of the other recordings collected for this research had the stylus force set to 1.8 g. It's important to note that varying the stylus force achieved a variation in the noise on the order of around 2 dB, which is an appreciable amount when compared to variations in noise seen between different records.

5.7 Test Lacquer Measurements

In the process of designing the test records, two reference lacquers were cut. These reference lacquers are not the same ones used to create the test records– that one was cut at a later date and destroyed during the stamper making process. They are lacquers cut by a mastering engineer and sent to the customer to be evaluated for quality. Test lacquers are meant to be played back a handful of times before they are unusable. One test lacquer was able to be recorded in this thesis that contained the test signals. It's measurements of surface noise are given in Table 5.3.

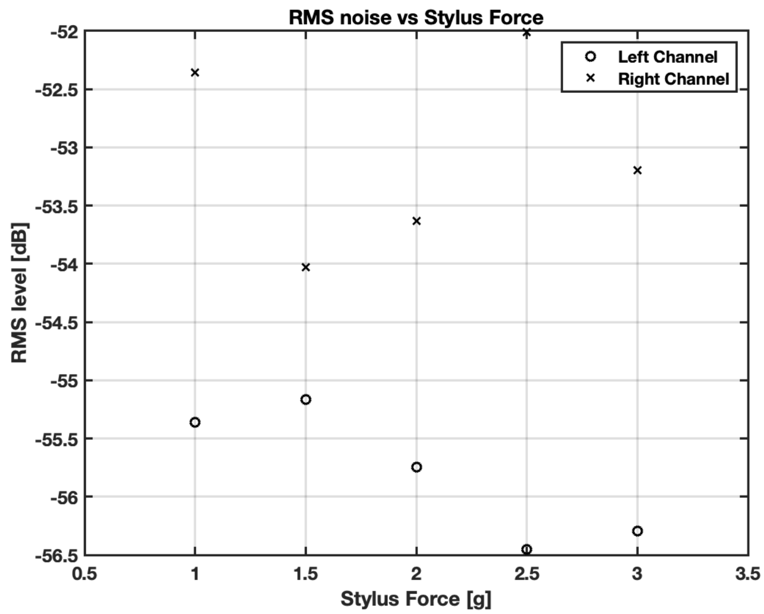

Figure 5.20: A weighted RMS level of a quiet track plotted against stylus force

track	RMS_L	RMS_R	A_L	A_R	$CCIR_L$	$CCIR_R$
quiet	-20.09	-19.37	-35.27	-24.45	-47.71	-33.53
quiet2	-21.94	-22.26	-59.92	-48.54	-71.39	-57.63

Table 5.3: The average Unweighted, A and CCIR/ARM weighted RMS noise levels of the two quiet tracks for the test lacquer measured.

5.8 Lacquer Spectrum

The lacquer's spectrum is given in Figure 5.21. It is a much flatter spectrum with no distinct humps as in the record's noise. This is an interesting feature that is distinct to the vinyl surface noise and hints at noise being added during the pressing process. These two humps in the vinyl surface noise are the two areas of interest, as they indicate surface noise that could be reduced or eliminated based on settings in the press.

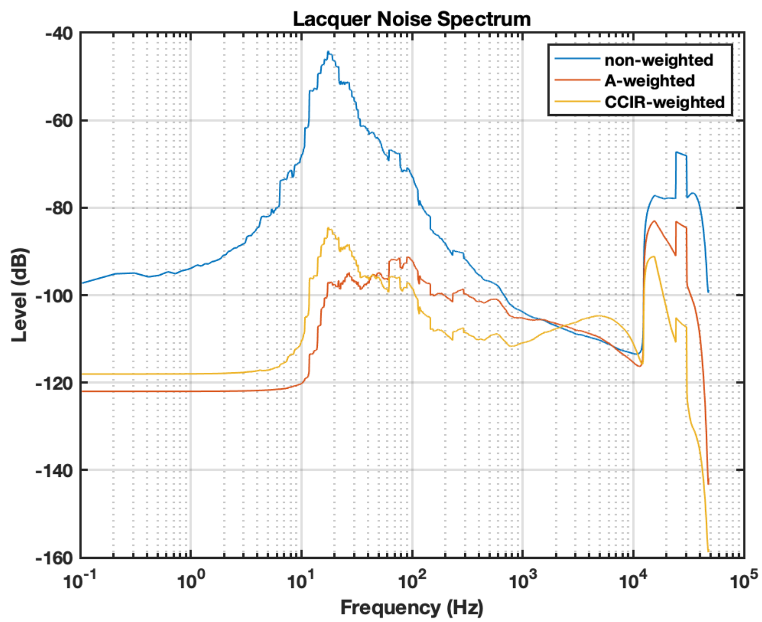

Figure 5.21: The noise spectrum of the Lacquer unweighted and with both A and CCIR/ARM Weightings.

5.9 Lacquer Coherence

It's very interesting to note that there is no coherence observed in the spectrum of the test lacquer as compared to the spectrum of a vinyl record. Applying the same method of separating out each groove segment in lacquer noise we observe no discernible coherence pattern in the quiet grooves of a lacquer. The spectra are similar, as shown in Figure 5.22. Note the lack of a hump in the 12-80 Hz range. This lack of a hump indicates that there is a lack of energy in the region where we'd expect to see coherence.

Indeed the lack of a hump in the lacquer's spectrum predicts that there should be no coherence pattern and this is confirmed in Figure 5.23. Whereas the record sees a nice region of around 80-90% coherence between 12-80 Hz as discussed in Chapter 5, the lacquer sees no such coherence. The coherence does seem to asymptotically approach 0.2 however this is due to the smoothing algorithm and the positive-definite nature of the coherence. The coherence is random at these frequencies, and that random noise is being smoothed to a non-zero value. This lack of coherence demonstrates that indeed the coherence observed in the surface noise of vinyl is a feature of the vinyl record and not a feature of phonographic

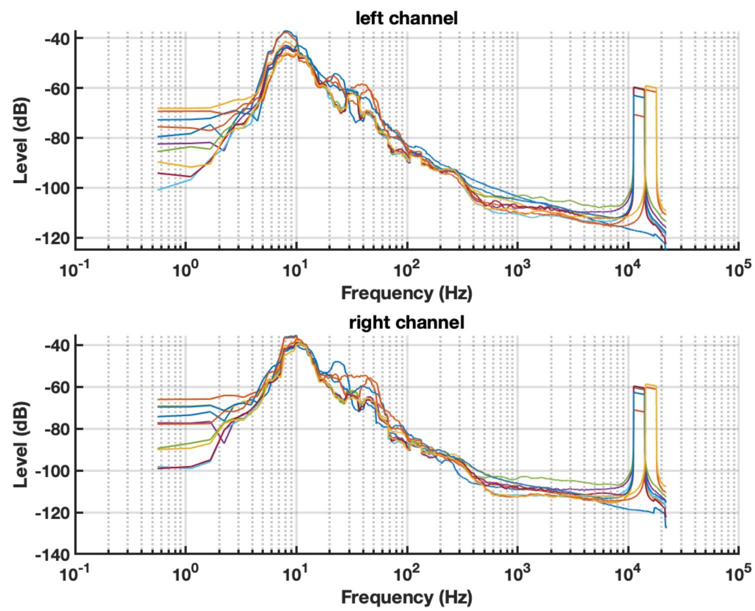

Figure 5.22: The spectra of each groove segments on the lacquer plotted on top of one another.

playback or the playback system used for testing.

Similarly there is no coherence between the left and right channels of the same groove segment as shown in Figure 5.24. Indeed no pattern of coherence can be found anywhere in lacquer noise. The lack of coherence in the lacquer noise suggests that noise levels on vinyl records can be reduced if the cause of the coherent noise can be found in the vinyl records and eliminated in the manufacturing process. It's important to note also, the lacquer is not necessarily quieter than vinyl. Indeed the first pressing of records has an average A weighted RMS noise level of -53 dB, which is much quieter than the RMS noise levels of -35 dB and -48 dB measured on the lacquer, the second pressing had an average A weighted RMS noise level of -46 dB which is again comparable to the lacquer.

Overall, the pressed records seemed to be quieter than the test lacquer, suggesting that it could be possible to achieve even lower surface noise levels if the cause of this coherence is understood and controlled.

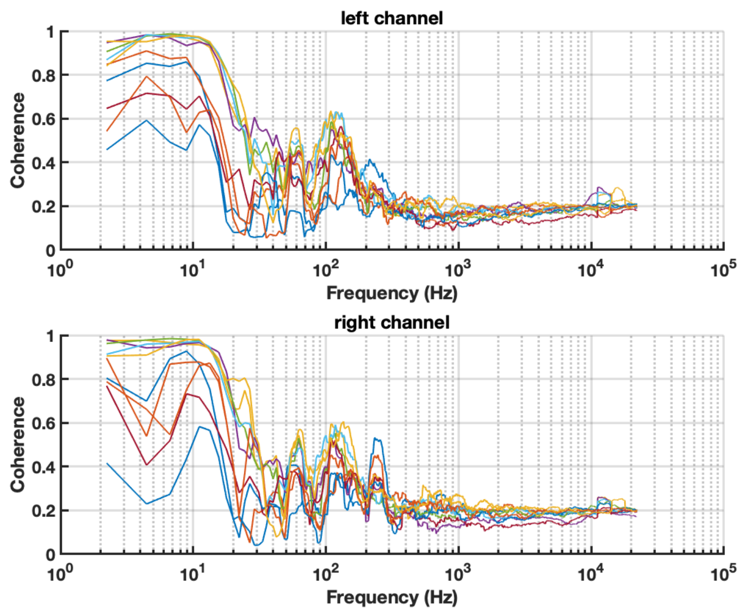

Figure 5.23: The coherence of each groove segment with the adjacent groove segment on the lacquer.

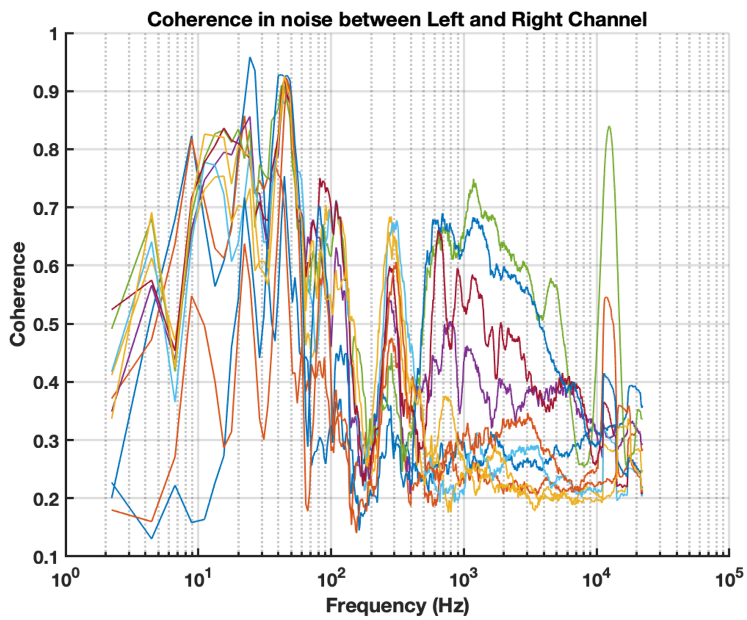

Figure 5.24: The coherence between the left and right channel noise of each groove segment on the lacquer.

Chapter 6

Clicks and Pops

Perhaps the most notable part of vinyl noise are the clicks and pops that jostle the stylus and produce sharp peaks in the audio waveform that interrupt the music. Most often caused by scratches, dust or debris in certain extreme cases a click can be so severe it causes the needle to lose tracking.

The term "click" here refers to a short spike in the waveform of audio resulting in an audible tick or pop. The study of these clicks is integral to the study of vinyl records as they are the culprit of many listener's disdain for vinyl. An abrupt interruption to the music is the last thing record manufacturers want. One of the primary measurements of quality will be the amount of clicks found on each record. The following chapter studies these clicks and pops in depth, discussing their nature, characteristics and possible causes. The algorithm for detecting clicks, as well as common clicks (clicks that are found across multiple pressings of the same record), is explored as well.

Clicks and pops are perhaps the focus of the manufacturer when improving the audio quality. As shown in Chapter 4, the number of clicks and pops do vary on a per record basis– implying that the press plays some role in the production of these defects.

6.1 Click Detection

Many click detectors work based on a few very simple principles. The click detector employed in this study is taken from Audacity's Click Removal effect, and rewritten in MATLAB. Audacity is an open source digital audio workstation, that is free to use. The MATLAB adaptation is included in Appendix A.7.

The Audacity click removal tool has several parameters two of which are set by the user in Audacity, and others which can be manipulated within in the code. Within Audacity the user may select the desired threshold and the click width. Both adjust the sensitivity of the detector and can impact the number of clicks detected. The Audacity Manual cautions that the detector can detect false positives and negatives if the settings are not tuned properly. Thus these parameters must be properly set in order to detect the right number of clicks. In order to do this several tests were run on the click detector to find the best settings.

Figure 6.1 demonstrates the number of clicks detected per the threshold and click width setting on the click detector. The click width setting was not found to do too much however the threshold setting does determine the number of clicks detected to a substantial degree. A threshold of 200 and a click width of 20 were chosen to be used for this test. 200 seems to be a reasonable value, the threshold can afford to be so low as the click detector is mainly designed to work with music. Since the test record contains mostly pure tones, there are very few transients in the audio that could be mistaken for a click by the detector.

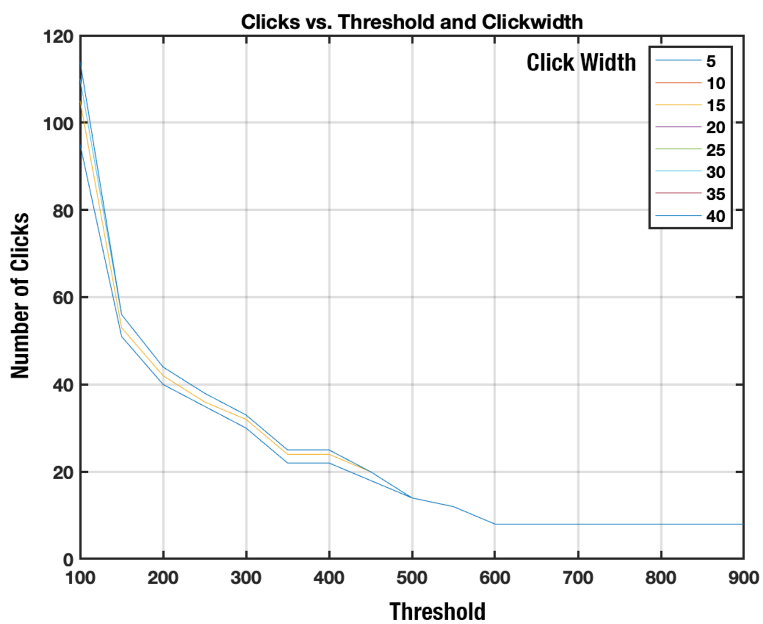

Figure 6.1: The number of clicks detected at varying levels of threshold and click width.

The click detector is limited, however the settings were kept constant. Thus the number of clicks detected in this study can be interpreted as the total number of clicks that meet the criteria laid out by the click detector, namely its threshold and click width. The Audacity manual warns that the click remover does not work on clicks longer than 10 ms. As such, the detector used in this study shares that limitation and will neither detect or correct those clicks [7].

6.2 Click Correction

The Audacity tool on which the click detection is based is primarily used to remove clicks. As such, built into the detection algorithm is the ability to correct these clicks. The correction is very simple, a linear interpolation is done between the two samples where the click begins and ends. This correction is fine for the purposes of this thesis, as the main reason for removing these clicks is to ensure their power is not measured as part of the surface noise measurements. To that end, a simple linear correction achieves the goal nicely. Figure 6.2 is a typical click detected by the detector, Figure 6.3 shows how the click detector removed the click and corrected the audio where the click was.

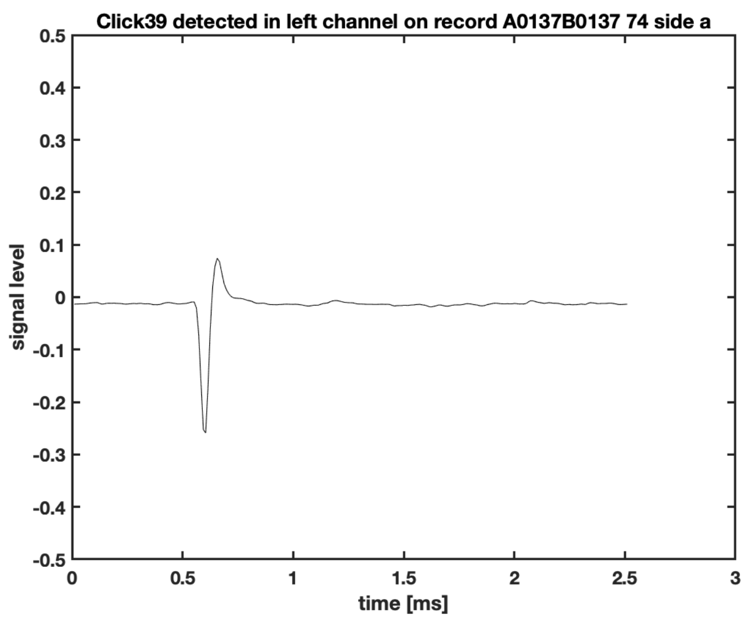

Figure 6.2: A typical click detected by the detector.

Note that for the purpose of testing, the clicks were removed from the test record audio at the detection stage. This was to ensure that there would be no contributions due to

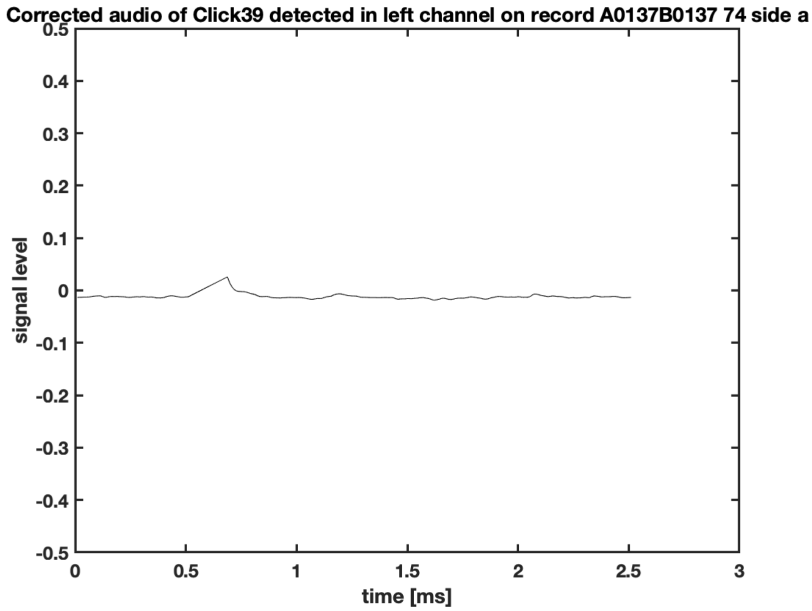

Figure 6.3: The corrected segment of audio where a click was.

click to the surface noise measurements. The corrected audio also showed better coherence without the clicks.

6.3 Common Clicks

One might ask if it's possible that by studying the clicks of the records we might be able to probe different stages in the record making process, to discover the source of these defects. By determining which clicks are common between different test records the origins of these clicks, whether they be the lacquer, press or the stamper can be discovered. It is impossible to study the lacquer that produced the test records directly– as it is destroyed during stamper creation. However, following a series of logical steps one can determine the origin of clicks whether it be the lacquer, stamper or press. While test lacquers were sent by the mastering engineer and evaluated, these were not the same lacquers used to create the stampers.

The click detector returns both a click-corrected audio file and an array with the locations of all the clicks that were detected and removed. Since the recordings are time matched, if the two waveforms were laid on top of one another then any clicks that are common to both recordings would lie on top of each other. Thus their locations in the audio would be precisely the same. In practice this doesn't happen, as no matter what was tried, records could not be time-matched exactly. See Chapter 2. A relaxation parameter was set to account for this mismatch. Thus a common click algorithm was developed that allows for clicks to fall within a range set by this relaxation parameter. This means that the common relaxation parameter must be tuned to ensure it captures all clicks that are common within sample error and does not include clicks which are not common but occur in a similar place in the audio. Below is pseudo code for implementing the common clicks algorithm.

```
reference_clicks = array of click locations from reference file
clicks = array of click locations from file under test
for each click in reference_clicks
    common_click = reference_clicks(i) - relaxation \\
                        < clicks < reference_clicks(i) + relaxation
    if common_click is not empty, count a common click
end
```

As with the click detector, the data here is smoothed by the relaxation parameter. In both cases the data is analyzed in bins of 2048 samples. The ideal relaxation was independently observed to be around 1000 samples, which correspond to bins of 2000. Two samples of audio are given below in Figure 6.4 and Figure 6.5, the clicks are highlighted in green. Note that there are a number of clicks common between the two records, clicks that occur in the same location.

Of course the number of common clicks detected varies according to the relaxation parameter. Figure 6.6 demonstrates how adjusting the threshold value can determine the number of common clicks detected between Figures 6.4 and 6.5. In this example a relaxation parameter of 750 captures all the clicks common between the two test records however a relaxation parameter of 2000 samples was kept in the actual testing.

With the ability to count common clicks, a few measurements come to mind. If two clicks are common between the a-sides of two records, then that click must be imprinted on the stamper used to press the record. If two clicks are common between the a-side and the b-side of two records (or the same record) then that click must be imprinted on the lacquer used to create the stampers. Clicks that are not common between any records, whether they are the a-side or the b-side, must be imprinted on the record during pressing.

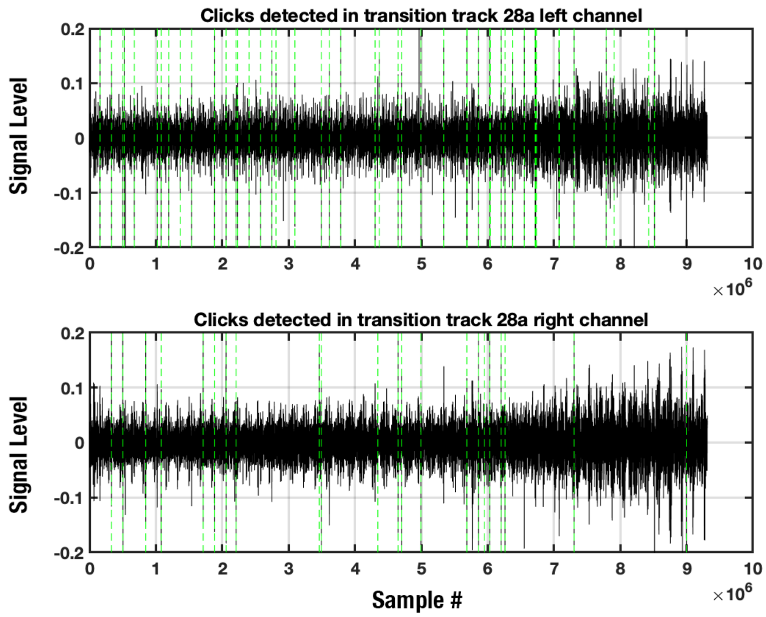

Figure 6.4: The audio from the transition track of record 28 side a with the clicks highlighted in green. Note that there are clicks that appear in both record 28 and 70 in the same place.

Using this logic, the number of clicks can be accounted for during each stage of the record making process.

The number of clicks that are on the lacquer is simply the number of common clicks found between an a-side and a b-side. The number of clicks unique to the stamper is equal to the number of clicks common between the same side subtract the number of clicks on the lacquer. When this is done there is only one area on the record that consistently shows common clicks. That is the transition track in the middle of the record, the long period of quiet that connects the two sets of signals. The transition track is unique as it contains quiet grooves that are spaced both close together and far apart, as the stylus has quite a far distance on the record to travel during the transition track.

The lack of common clicks found on the records points to stamper defects or clicks found on the master lacquer as not being an issue in the pressing of the record. Indeed most clicks were found to be random, which points to dust or defects during the pressing to be the main contributer to pops and clicks. Nearly all pops and clicks are unique to each record, which points to the pressing process as the main culprit for the production

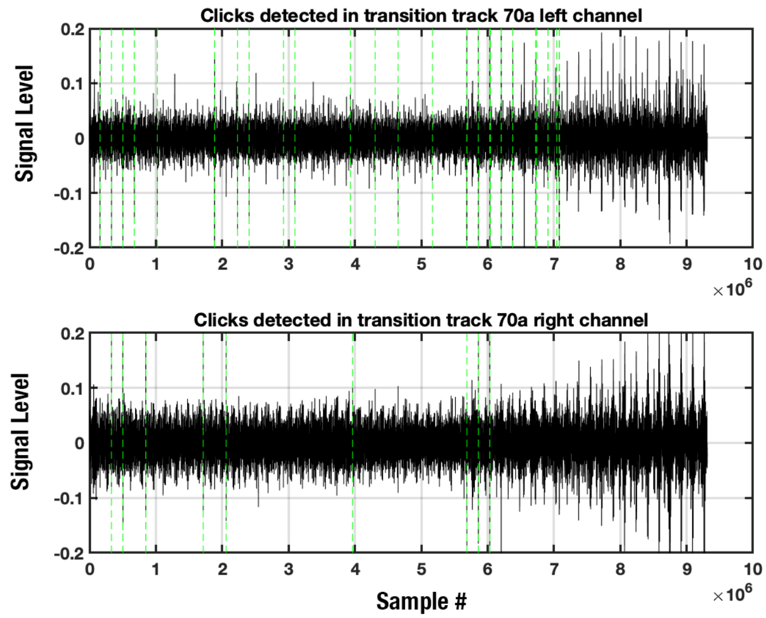

Figure 6.5: The audio from the transition track of record 70 side a with the clicks highlighted in green. Note that there are clicks that appear in both record 28 and 70 in the same place.

of these clicks. Dust is known to produce clicks and defects and indeed despite the best efforts at cleaning the records dust was still observed on the surface of the records.

There are a few groups of clicks on the record however that were found to be common. The largest occurs in the transition track of pressing two. Examples of these clicks have been shown in Figures 6.4 and 6.5. When measuring common clicks two reference records were used, a reference record from the first pressing and a reference taken from the second pressing. This was to test whether clicks or pops could be added to stamper during the loading process and indeed that is what we believe has happened. On side a in the transition track of all the records pressed in pressing two, there is an average of 32 common clicks found in the left channel and 4 common clicks found in the right channel. Also in pressing two there is a series of 5 clicks that appear in the sweepR2 track in the left channel. It is believed that the appearance of these common clicks in only the second pressing points to damage done to the stamper either in storage or during the loading of the stamper into the press.

It is unknown what could have caused this damage and it's very interesting to note

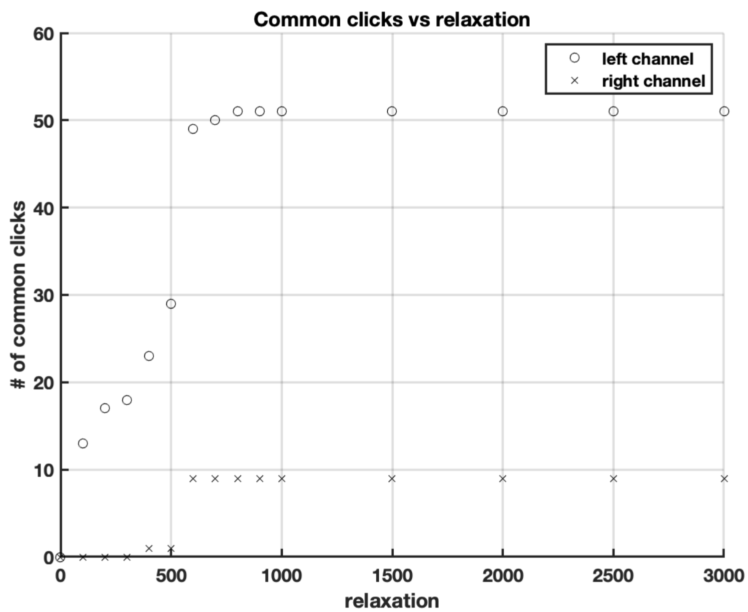

Figure 6.6: The common clicks found between transition track of record A0000B0000 70 side a and A0000B0000 28a side at varying levels of the relaxation parameter.

that it is primarily the left channel that contains the newly found clicks. I find this is a very strange result, as I believe that any contamination on the backplate that holds the stamper would damage the stamper in such a way that any clicks or pops would appear in both channels at the same time as a “stereo” click. Perhaps the appearance of these clicks in only a single stereo channel indicates that it was the groove side of the stamper that was contaminated.

6.4 Difference between the left and right channels

Similar to the surface noise, the two stereo channels display distinct differences when it comes to clicks. The left channel shows substantially more clicks than the right channel. As per the NAB standard, the outer groove wall contains the right channel information and the inner groove wall contains the left channel information [8]. As such, substantially more clicks are found on the inner groove wall when compared to the outer wall.

When records are pressed, the vinyl puck starts in the centre of the record and the vinyl flows outward as the press closes. The outer groove wall, which contains the right channel, has the vinyl flowing into it and pressing up against the wall of the groove to form that groove. The other half of the groove is backfilled by the vinyl, there is no outward flow against the inner groove wall, that contains the left channel. The difference in these two physical processes leaves their impression on the audio. It is believed that more clicks should appear in the groove that is backfilled rather than the one that is facing the direction of vinyl flow, as the vinyl flow outward should smooth the surface of the right channel groove wall and prevent any defects or clicks from forming. It's interesting to note that as reported in Chapter 5, there is less surface noise in the left channel despite the left channel having more clicks due to this backfill. This might suggest that it is the relaxation of the vinyl as it cools that determines the surface noise level, as despite the right channel facing the flowing vinyl it has more noise and thus a rougher edge as compared to the left channel that is backfilled. So this noise in the right channel might be added after the record has left the press and the vinyl is left to cool. Then the rough edges that form the surface noise develop.

Chapter 7

Conclusions

As discussed in the introduction, this thesis is an exploration into the pressing process of vinyl records. It's research goals were to uncover the process by which noise, distortion, pops and clicks and other unwanted audio artifacts and qualities are imprinted onto vinyl records during the pressing stage of production. Specifically, which parameters and settings within the WarmTone press affect the final audio quality and how press operators should be setting their WarmTone presses to create the ideal sounding record.

In conclusion, there is very little a press operator can do to affect the final audio quality of their records. There is no clear difference in audio quality measured between the different parameter changes. This is both a blessing and a curse. It doesn't seem that there is much a press operator can do to change the audio quality of the records being produced by their press in a reliable way. However when time, material and energy costs are taken into consideration, operators can instead choose to focus on minimizing the cost of these parameter changes rather than maximizing any gains in audio quality they hope to achieve with their machine. Instead visual inspection of the records should be the focus of any quality control, as the audio measurements were too unreliable to draw any meaningful relationships from and are very time consuming, resource intensive and specialized when contrasted with simple visual inspection.

In fact all measurements point to the fact that beyond choosing a type of plastic, there is little that press operators can do to impact the final audio of records coming off their press. Interestingly however there does seem to be many indications that perhaps the extrusion process and the plastic chosen does play a large role in determining surface noise and perhaps further innovation in the processing of PVC compounds and the extrusion process in pressing records will lead to improvements in overall noise levels. Coherence measurements seem to suggest that vinyl pellets and their residual memory might be the cause of coherent surface noise on the record and an area of improvement and innovation for future vinyl records.

7.1 Properties and possible nature of vinyl surface noise

One discovery of this thesis is the correlation of the surface noise when the record's surface noise are separated by groove segments, this is discussed in Chapter 5. This correlation corresponds to a hump in the spectrum of surface noise at and around 60-100 Hz and quite possibly represents an area of interest in the surface noise of vinyl records. The hypothesis put forth in this thesis is that this hump is due to the vinyl "pellets" retaining some memory of their shape through the extrusion and pressing process and so returning to that pocketed and bumpy shape as the record cools and producing this surface noise. This hump doesn't appear in the lacquer's spectrum, a result which is mirrored in Eargle's study [6]. The hump also roughly corresponds to the overall pellet size before the PVC is loaded into the extruder.

Overall the records were found to have relatively low weighted surface noise, comparable to the test record measured, suggesting it is not a simple relationship between lacquer, stamper and record. Noise isn't simply added during the vinyl pressing stage, it appears as though the noise is fundamentally of a different nature. Whether this bodes well or not for improving the sound of records or further research into the topic is yet to be determined.

7.2 Optimal Press Parameters and Recommendations

One of the primary goals of this thesis and research is to develop a handbook of sorts on pressing records, to relate each process and setting in the WarmTone press to a noticeable effect or result in the audio of the test records, unfortunately none of the settings or parameters in the WarmTone had any real effect on the resulting audio. When measuring weighted surface noise, there was almost no difference between the different parameter changes, unweighted surface noise is a little bit of a different story but those differences may simply be due to random processes. For the other parameter changes within the press random processes seem to dominate. For the distortion measurements and stereo bleed surprisingly there is a difference measured on each record, the distortion and the stereo bleed are varying on a per record basis. Unfortunately these changes do not appear to correspond with any definitive parameter or setting changes in the WarmTone.

That being said an attempt will be made here to provide optimal press setting recommendations and possible causes of defects and issues with records. It appears as though there is very little the press operator can do about the amount of surface noise on the record. It's possible that the rumble noise may be changing but there don't seem to be any press parameters that point to the operator having control over the amount of surface noise. Minimum overlap, steam and puck size have the most amount of clicks while maximum close to speed, dwell, and minimum close to speed have the least. Ignoring the close to speed results as it's unclear how those might influence the number of clicks then we're left with maximum dwell, cool, steam and minimum press open delay as giving a very small amount of clicks, showing that perhaps it is time under pressure that is the best remedy for clicks on a record. All these are maximum mould times and so would increase the amount of time the press is spent closed. This perhaps fits the hypothesis as the number of clicks found in the left channel nearly doubles the number of clicks in the right. Remember as it is the left channel that is backfilled when a record is pressed, more time under pressure allows more time for vinyl to fill the spaces that cause clicks on the record in the left channel.

Stereo bleed, distortion and wow all appear to vary from record to record but there does not seem to be a pattern to their changes, showing that the aspects and qualities of records does vary during a pressing is important, as not every record pressed is exactly the same. So press operators need to be careful when choosing which and how many records to spot check, as they need to ensure they are getting an accurate glimpse of the pressing as a whole and not just that one record, which may be an outlier. Looking at the visual defects it's clear that running too hot of an extruder is a possible cause of staining and clouding, maximum cooling shows a large amount of non fill and stitching seems to have occured in only a handful of pressings of unrelated parameters.

Maximum operating speed alongside minimum initial speed and puck size show by far the most amount of warpage. This aligns with expectations as operating speed would twist the vinyl material into the extruder. Perhaps it retains this twisted shape into a warp as the record cools. A small puck size means very little vinyl material to form a record and a minimum initial speed means a slow press to compress the shape, all signs of a weak base shape for the vinyl and thus the production of a warped record. None of the parameters seems to reduce warpage significantly, or there are no patterns to the ones that do. However these parameters do seem to cause warpage. Cupping is slightly different as both too cool and too hot of an extruder cause cupping. It could also be an uneven heating pattern in the extruder as the components might prefer to be changed in temperature together, with a hot premould matching hot barrel zones for instance.

Overall time under pressure to reduce clicks, reduce cooling or extruder temperatures to reduce the number of visual defects, increase puck size or decrease operating speed to reduce warpage and experimentation with different plastics shapes and sizes inside the extruder are the best recommendations from this thesis. Records do vary wildly from one another in a single pressing, however in the average overall change between two pressings is minuscule. Plastics make a difference in terms of character of sound, however not in overall noise levels in the two plastics measured, at least when measured at multiple locations on the disc and weighted appropriately. Overall there is very little a press operator can do to affect the quality of records in a press from an audio perspective. However, visual inspections are easy, quick and do not require any equipment. As such it is recommended that press operators perform visual inspections of their records and adjust press parameters to remove visual defects, rather than try to improve the audio quality from the press.

7.3 Comparison with previous research

As discussed in the introduction the other studies into pressing quality were done in 1977 by Ruda [5] and 1969 by Eargle [6], in particular Eargle use the same NAB normalization standard as our test discs, so our data can be compared to his in some respects. Eargle measures groove diameter as opposed to our groove radius, so multiplying our groove radius measurements by two and converting to inches we find our quiet, transition and quiet2 tracks are at groove diameters of 10.5", 8.9" and 6.4". Compared to Eargle's measurement points of 11.5", 8.5" and 5.5" our tracks are at similar diameters on the test records, about an inch off from Eargle's disc. Our measurement techniques are different from the NAB standard, which describes a flat band of noise from 500 to 15000 Hz. [8] Eargle had high frequency noise measurements of -66.5 dB, this was done using a different methodology and measurement technique from this thesis, compliant with the NAB standard. The shape of Eargle's spectra also match the spectra found in this thesis, both with the lacquer and test records measured. The same hump in frequencies is seen in Eargle's plots, as is the lack thereof in the measurements of the test lacquer. However in this thesis a large hump in the upper frequencies was noticed in the lacquer spectrum that isn't in Eargle's measurements or the measurements of the test record. It's unsure what's the cause of this high frequency noise in the spectrum, but it is a real measured effect in our data.

7.4 Next steps

Moving forward, any practical measurements of the quality of records should primarily be visual rather than auditory. The audio measurements do seem to vary on a per record basis, however there appear to be no indications as to which if any of the press parameters can influence this final audio quality. The large variability in audio measurements does point to a real difference from one test record to another, noting that indeed press operators are measuring real differences when they are judging sound quality.

One limitation in this research is the way in which press parameters were changed. Time, cost and labour constraints meant that the press parameters were tested in a limited, albeit still comprehensive way. By varying between maximum and minimum values it was thought that large differences in audio would be seen and would be noticeable, however this wasn't the case. Perhaps the press parameters are more intimately related and need to be changed in groups of changes to have any meaningful impact. This is particularly true of the mould parameters, which describe the heating profile of the moulds.

Most signs in this thesis point to innovation in the extrusion process as a possible direction for future research into improving vinyl records. By far the difference in plastics measured between the two pressings causes the greatest difference in surface noise levels. The plastics didn't do much to change the distortion or the stereo bleed. Those seem to largely be a random process. It was discovered that the nature of surface noise appears different between lacquers and records, indeed some records pressed had lower surface noise levels than the test lacquer measured. Future research should focus on the shape and size of vinyl pellets and examining the correlation between records.

References

[1] F. Rumsey, "The vinyl frontier," *AES: Journal of the Audio Engineering Society*, vol. 62, no. 7-8, pp. 559–562, 2014.

[2] E. Crawford, "Nielsen US End of Year Music Report 2017," pp. 1–33, 2017.

[3] M. J. Murphy, "Back to the Future: Reintroducing Analog and Vinyl Audio Production to Digital Natives," *Audio Engineering Society*, vol. 8, no. 4, pp. 318–322, 2021.

[4] J. P. Friedlander, "Year-end 2020 RIAA Revenue Statistics," pp. 4–6, 2020.

[5] J. Ruda, "Record Manufacturing: Making the Sound for Everyone," *Journal of the Audio Engineering Society*, vol. 25, no. 10/11, pp. 702–711, 1977.

[6] J. Eargle, "Performance Characteristics of the Commercial Stereo Disc," *Journal of the Audio Engineering Society*, vol. 17, no. 4, pp. 416–422, 1969.

[7] "Audacity manual," 2020. https://manual.audacityteam.org/index.html, (accessed: 11.11.2020).

[8] N. Recording and R. S. Committee, "NAB Recording and Reproducing Standards for Disc Recording and Reproducing.," 1964.

[9] R. Dolby, D. Robinson, and K. Gundry, "Ccir/arm: A practical noise measurement method," vol. 1353, no. 1353, 1978.

[10] G. Alexandrovich, "A stereo groove problem," *AES: Journal of the Audio Engineering Society*, vol. 691, 1939.

[11] SHURE, "DATA SHEET Model TTR-103 PHONO CARTRIDGE (45 rpm) Trackability Test Record," 1979.

[12] J. T. Mull and B. C. Enterprises, "Measurement of Flutter and Wow in Magnetic-Tape Instrumentation Recorders," *AES: Journal of the Audio Engineering Society*, vol. 3, no. 3, pp. 151–154, 1955.

[13] S. P. Lipshitz, "ON RIAA EQUALIZATION NETWORKS.," *J Audio Eng Soc*, vol. 27, no. 6, pp. 458–481, 1979.

[14] R. Kilmanas, "Tonearm geometry and frequency-modulation distortion," *AES: Journal of the Audio Engineering Society*, vol. 30, pp. 574–579, 1982.

[15] L. Happ and F. Karlov, "Record Warps and System Playback Performance.," *AES: Journal of the Audio Engineering Society*, vol. 24, no. 8, pp. 630–638, 1976.

[16] G. P. Humfeld, "Control of static electricity on a phonograph record*," 1962.

[17] D. A. B. Garside and G. R., "Groove Deformation and Distortion in Records," *AES: Journal of the Audio Engineering Society*, vol. 26, no. 7/8, pp. 498–510, 1978.

[18] J. V. White, "An Experimental Study of Groove Deformation in Phonograph Records," *AES: Journal of the Audio Engineering Society*, vol. 18, no. 5, pp. 497–506, 1970.

[19] M. John T., "Monogroove stereophonic disk recording*," 8 1954.

[20] F. V. Huntt, "On Stylus Wear and Surface Noise in Phonograph Playback Systems," *AES: Journal of the Audio Engineering Society*, 1955.

[21] "Warmtone processing and troubleshooting cheat sheet," 2018. Viryl Training Handout.

[22] MATLAB, "Matlab documentation," 2020b.

APPENDICES

Appendix A

Matlab Code

Plenty of MATLAB code was developed for the analysis required for this thesis. This code is included in the appendix.

A.1 Separate Tracks Script

The SeparateTracks.m scripts reads an audio file with the following naming convention NNNsXXXX.XXX.wav. Where the "NNN"'s represent the record number being recorded, "s" represents the side that was recorded and XXXX.XXX is the timestamp that the final pip occurs at used to calculate the lag difference between this record and the reference in order to separate the tracks with the proper timestamps.

```

function [output, info_array] = SeperateTracks(file)
            %~~~~~~~~~~~~~~~~~ LOAD REFERENCE ~~~~~~~~~~~~~~~~~%
                % try
                addpath('/Users/cz/Code/vinyl-research/matlab_code/audio_functions')
                disp('SEPERATE TRACKS CALLED')

            %~~~~~~~~~~~~~~~~~~ Reference info ~~~~~~~~~~~~~~~~%

                timestamps_ref = [0, 60, 90, 122, 158, 180, 246, 266, 304, 324, 362, 382, 417.5];
                % this is how many seconds each signal is according to Chris Muths track listing
```

```
                lengths = [60, 30, 31, 36, 21, 66, 20, 37, 19, 37, 19,
   37, 19]; %starts with 1kHz
                signal_names = {'leadin',     % 1
                                '1kHz',       % 2
                                '10kHz',      % 3
                                '100Hz',      % 4
                                'sweep',      % 5
                                'quiet',      % 6
                                '3150Hz',     % 7
                                '1kHzL',      % 8
                                'sweepL',     % 9
                                '1kHzR',      % 10
                                'sweepR',     % 11
                                '1kHzV',      % 12
                                'sweepV',     % 13
                                'transition',% 14
                                '1kHz2',      % 15
                                '10kHz2',     % 16
                                '100Hz2',     % 17
                                'sweep2',     % 18
                                'quiet2',     % 19
                                '3150Hz2',    % 20
                                '1kHzL2',     % 21
                                'sweepL2',    % 22
                                '1kHzR2',     % 23
                                'sweepR2',    % 24
                                '1kHzV2',     % 25
                                'sweepV2',    % 26
                                'leadout'     % 27
                };
                % timestamps =       [[0, 61],     % 1. 1 kHz
                %                     [61,91],     % 2. 10 kHz
                %                     [91,121],    % 3. 100 Hz
                %                     [121,159],   % 4. sweep
                %                     [159,180],   % 5. quiet
                %                     [180,245],   % 6. 3150 Hz
                %                     [245,267],   % 7. 1 kHz left
                %                     [267, 302], % 8. sweep left
                %                     [302, 325], % 9. 1 kHz right
                %                     [325, 361], % 10. sweep right
                %                     [361, 383], % 11. 1 kHz vertical
                %                     [383, 418], % 12. sweep vertical
                %                     [418, 515], % 13. transition
                %                     [515, 578], % 14. 1 kHz
                %                     [578, 608], % 15. 10 kHz
```

```
                %                      [608, 639], % 16. 100 Hz
                %                      [639, 676], % 17. sweep
                %                      [676, 698], % 18. quiet
                %                      [698, 760], % 19. 3150 Hz
                %                      [760, 785], % 20. 1 kHz left
                %                      [785, 820], % 21. sweep left
                %                      [820, 842], % 22. 1 kHz right
                %                      [842, 878], % 23. sweep right
                %                      [878, 900], % 24. 1 kHz vertical
                %                      [900, 938]];% 25. sweep vertical
                %                      % [938, 950]];
                %                      %% dont forget lead in and leadout

            timestamps =            [[14.624, 75.229],     % 1. 1 kHz
                                    [75.229, 104.512],     % 2. 10 kHz
                                    [104.512, 135.275],    % 3. 100 Hz
                                    [135.275, 173.224],   % 4. sweep
                                    [173.224, 195.111],   % 5. quiet
                                    [195.111, 258.293],   % 6. 3150 Hz
                                    [258.293, 281.625],   % 7. 1 kHz left
                                    [281.625, 318.352], % 8. sweep left
                                    [318.352, 339.302], % 9. 1 kHz right
                                    [339.302, 376.067], % 10. sweep right
                                    [376.067, 396.919], % 11. 1 kHz vertical
                                    [396.919, 432.435], % 12. sweep vertical
                                    [432.435, 532.324], % 13. transition
                                    [532.324, 592.655], % 14. 1 kHz
                                    [592.655, 623.211], % 15. 10 kHz
                                    [623.211, 653.123], % 16. 100 Hz
                                    [653.123, 690.572], % 17. sweep
                                    [690.572, 712.537], % 18. quiet
                                    [712.537, 775.368], % 19. 3150 Hz
                                    [775.368, 799.286], % 20. 1 kHz left
                                    [799.286, 835.582], % 21. sweep left
                                    [835.582, 857.235], % 22. 1 kHz right
                                    [857.235, 893.610], % 23. sweep right
                                    [893.610, 914.697], % 24. 1 kHz vertical
                                    [914.697, 950.037]];% 25. sweep vertical
                                    % [938, 950]];
                                    %% dont forget lead in and leadout

```

```


        [data, fs] = audioread(file);

        %~~~~~~~~~~ MANUAL LINEUP ~~~~~~~~~~~%
        timestring = file(end-11:end-4); % get the

        timepip = str2num(timestring(1:2))*60 + str2num(timestring(3:end));

        timestringref = '1558.066';
        timepipref = str2num(timestringref(1:2))*60 + str2num(timestringref(3:end));



        timediff = timepipref - timepip;

        lagdiff = floor(timediff*96000);

        %~~~~~ Correlation correction ~~~~~%

        ref = audioread('/Volumes/AUDIOBANK/audio_files/A0000B0000/031418_A0000B0000r028a1558.066/leadout.wav');
        timepipref2 = 7.934;

        lockout = timepipref2;
        refLockout = ref(timepipref2*fs-0.25*fs:timepipref2*fs+0.25,:);
        dataLockout = data(timepip*fs-0.25*fs:timepip*fs+0.25*fs,:);

        [acor_L,lags_L2] = xcorr(refLockout(:,1),dataLockout(:,1));
        [M_L,I_L] = max(abs(acor_L));
        lagdiff_L2 = lags_L2(I_L);
        lagdiff2 = lagdiff_L2;
        lagdiff2 = 0;
        timediff = timediff + lagdiff2/fs;
        %~~~~~ Correlation correction ends ~~~~~%

        timedata = (0:length(data)-1);% + timediff;
        timestamps = timestamps - timediff;
        if timestamps(1,1) < 0;
            timestamps(1,1) = 1;
```

```
                end

                disp(strcat('timediff...',num2str(timediff)))
                %~~~~~~~~~~ MANUAL LINEUP  ENDS ~~~~~~~~~~~%




            %~~~~~~~~~~~~~~~~~~ NORMALIZATION RMS ~~~~~~~~~~~~~~~~~~~~%


                t = 1;
                sigtime = timedata(floor(timestamps(t,1)*fs):floor(timestamps(t,2)*fs));
                sig = data(floor(timestamps(t,1)*fs): floor(timestamps(t,2)*fs),:);
                N = 3*fs;
                seg = sig(0.33*length(sig):0.33*length(sig) + N - 1,:);



                disp(strcat('RMS before norm... ', num2str((rms(seg)))))
                disp(strcat('dB... ', num2str(20*log10(rms(seg)))))

                normalization = rms_response(seg);

                normalization_L = normalization(1);
                normalization_R = normalization(2);


                data(:,1)=data(:,1)./normalization_L;
                data(:,2)=data(:,2)./normalization_R;

                disp(strcat('normalization_L...', num2str(normalization_L)))
                disp(strcat('normalization_R...', num2str(normalization_R)))


                sig = data(floor(timestamps(t,1)*fs): floor(timestamps(t,2)*fs),:);
                N = 3*fs;
                seg = sig(0.33*length(sig):0.33*length(sig) + N - 1,:);

                disp(strcat('RMS after norm... ', num2str((rms(seg)))))
```

```
            disp(strcat('dB... ', num2str(20*log10(rms(seg)))))

        %~~~~~~~~~~~~~~~~~~ NORMALIZATION RMS ENDS ~~~~~~~~~~~~~~~~~~~~%



            signals = cell(length(signal_names),1);
            signal_times = cell(length(signal_names),1);

            for t = (1:length(signal_names))
                track_name = signal_names{t};
                disp(strcat('track  ...',track_name))


                csig = [];
                CLICKS_R = [];
                CLICKS_L = [];
                RMS_L = [];
                RMS_R = [];
                THD_L = [];
                THD_R = [];

                if t == 1
                    sig = data(1 : floor(timestamps(1,1)*fs),:);
                    sigtime = timedata(1 : floor(timestamps(1,1)*fs));
                    % refT = ref(1 : floor(timestamps(1,1)*fs),:);
                elseif t == length(signal_names)
                    sig = data(floor(timestamps(end,2)*fs) : length(data),:);
                    sigtime = timedata(floor(timestamps(end,2)*fs) : length(data));

                else
                    sig = data(floor(timestamps(t-1,1)*fs) : floor(timestamps(t-1,2)*fs),:);
                    sigtime = timedata(floor(timestamps(t-1,1)*fs) : floor(timestamps(t-1,2)*fs));

                end

                signals{t} = sig;
                signal_times{t} = sigtime; % not currently assigned to output
```

```


            end

            disp('ASSIGNING OUTPUT')
            output = containers.Map(signal_names, signals);
            info_array = [lagdiff, normalization_L, normalization_R];
            disp('EXITING SEPERATE TRACKS')


        end
```

A.2 Sensor Process Script

The SensorProcess.m script takes the XXXXXX_SensorValues.csv file from the WarmTone press and outputs the SensorTable, a table that lists each record pressed along with the maximum and minimum value of each sensor value during the time that record was pressed.

```
% SensorProcess.m
% Created by Chris Zaworski
% Input is the ########_SensorValues.csv
% Output is the SensorTable.csv

%#####_SensorValues
% tracks the data from the sensors in the press via timestamps
selectedColumns ={'id',
'RecordTimeStamp',
'PressPosition_Inches',
'PressForce_Ton',
'MouldSteamIn_PSI',
'MouldSteamIn_F',
'MouldSteamOutTop_F',
'MouldSteamOutBottom_F',
'ExtruderFeedthroatTemp_F',
'ExtruderBarrelZone1Temp_F',
'ExtruderBarrelZone2Temp_F',
'ExtruderBarrelZone3Temp_F',
'ExtruderDieZoneTemp_F',
'ExtruderPremouldTemp_F',
'ExtruderMeltTemp_F'};

opts = detectImportOptions('########_SensorValues.csv');
getvaropts(opts,selectedColumns);
```

```
opts = setvartype(opts,selectedColumns,'string');
opts.SelectedVariableNames = selectedColumns;
SensorValues = readtable('########_SensorValues.csv',opts)

height(SensorValues)
SensorValues.PressPosition_Inches = str2double(SensorValues.
    PressPosition_Inches);
SensorValues.PressForce_Ton = str2double(SensorValues.PressForce_Ton);
SensorValues.MouldSteamIn_F = str2double(SensorValues.MouldSteamIn_F);
SensorValues.MouldSteamOutTop_F = str2double(SensorValues.
    MouldSteamOutTop_F);
SensorValues.MouldSteamOutBottom_F = str2double(SensorValues.
    MouldSteamOutBottom_F);
SensorValues.ExtruderFeedthroatTemp_F = str2double(SensorValues.
    ExtruderFeedthroatTemp_F);
SensorValues.ExtruderBarrelZone1Temp_F = str2double(SensorValues.
    ExtruderBarrelZone1Temp_F);
SensorValues.ExtruderBarrelZone2Temp_F = str2double(SensorValues.
    ExtruderBarrelZone2Temp_F);
SensorValues.ExtruderBarrelZone3Temp_F = str2double(SensorValues.
    ExtruderBarrelZone3Temp_F);
SensorValues.ExtruderDieZoneTemp_F = str2double(SensorValues.
    ExtruderDieZoneTemp_F);
SensorValues.ExtruderFeedthroatTemp_F = str2double(SensorValues.
    ExtruderFeedthroatTemp_F);
SensorValues.ExtruderPremouldTemp_F = str2double(SensorValues.
    ExtruderPremouldTemp_F);
SensorValues.ExtruderMeltTemp_F = str2double(SensorValues.
    ExtruderMeltTemp_F);


i = 1;
press_close = 1;
press_open = 1;
lower_bound = 1;
upper_bound = 1;
record_number = 0;
recordNumbers = [];
recordArray = [];
while i < height(SensorValues)
    if SensorValues.PressPosition_Inches(i) < 1
        i = i + 1;
        continue
    else
        press_close = i;
```

```
        press_open = i;
        while SensorValues.PressPosition_Inches(press_close) > 1
            press_close = press_close + 1;
        end
        lower_bound = press_open - 1;
        while SensorValues.PressPosition_Inches(lower_bound) < 1
            lower_bound = lower_bound - 1;
            if lower_bound < 1
                lower_bound = 1;
                break
            end
        end
        lower_bound = lower_bound + 1;
        while SensorValues.PressPosition_Inches(press_open) > 1
            press_open = press_open + 1;
        end
        upper_bound = press_open;
        while SensorValues.PressPosition_Inches(upper_bound) < 1
            upper_bound = upper_bound + 1;
            if upper_bound > height(SensorValues)
                recordArray = [recordArray; [record_number,lower_bound,
    press_close, press_open, height(SensorValues)-1]];
                break
            end
        end
    end
    i = upper_bound;
    record_number = record_number + 1;
    recordArray = [recordArray; [record_number, lower_bound, press_close,
    press_open, upper_bound]];
end

recordTable  = array2table(zeros(height(SensorValues),1), 'VariableNames'
    , {'recordNumber'});
recordArray(length(recordArray),5) = recordArray(length(recordArray),5) -
    1;

for i = (1:length(recordArray))
    for j = (recordArray(i,2):recordArray(i,4))
        recordTable(j,'recordNumber') = num2cell(recordArray(i,1));
    end
end

recordTable = [recordTable SensorValues];
writetable(recordTable,'#######_SensorTable.csv'));
```

A.3 Record Process Script

The RecordProcess.m file crucially performs all the audio measurements on the record. It loops through all the tracks in the "tracks" *struct* and outputs all the measurements in the output array, which are the rows that get appended to the AudioTable in the BatchProcess.m script.

```

function output = recordProcess(file)
            output = [];
            fs = 96000;
            [tracks, info_array] = SeperateTracks(file);

            info_array
            lagdiff = info_array(1);
            normalization_L = info_array(2);
            normalization_R = info_array(3);

            signal_names = tracks.keys;
            signals = tracks.values;

            display('LOADING REFERENCE')
            if ismac() == true
                if file(length(file)-4) == 'a'
                    reference = ('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028a.wav');
                    %% Reference 02072019_A0000B000r27a.wav
                    offset = 15;
                elseif file(length(file)-4) == 'b'
                    disp('PC')
                    reference = ('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028b.wav');
                    %% Reference 02072019_A0000B000r27b.wav
                    offset = 13.1;
                else
                    disp('NO SIDE FOUND, USING SIDE A REFERENCE')
                    reference = ('d:/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028a.wav');
                    %% Reference 02072019_A0000B000r27a.wav
                    offset = 15;
                end
            end
```

```
        if ispc() == true
            disp('IS PC')
            if file(length(file)-4) == 'a'
                reference = ('d:/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028a.wav');
                %% Reference 02072019_A0000B000r27a.wav
                offset = 15;
            elseif file(length(file)-4) == 'b'
                disp('PC')
                reference = ('d:/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028b.wav');
                %% Reference 02072019_A0000B000r27b.wav
                offset = 13.1;
            else
                disp('NO SIDE FOUND, USING SIDE A REFERENCE')
                reference = ('d:/OneDrive - University of Waterloo/School/Vinyl_Project/audio_bin/A0000B0000/031419_A0000B0000r028a.wav');
                %% Reference 02072019_A0000B000r27a.wav
                offset = 15;
            end
        end


        for t = 1:length(signal_names)
            track_name = signal_names{t};
            sig = signals{t};
            disp(strcat('track  ...',track_name))


            csig = [];
            CLICKS_R = [];
            CLICKS_L = [];
            RMS_L = [];
            RMS_R = [];
            THD_L = [];
            THD_R = [];

            % if t == 1
            [csig(:,1), CLICKS_L] = ClickDetect(sig(:,1));
            [csig(:,2), CLICKS_R] = ClickDetect(sig(:,2));

```

```
                REFSa1_L = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/A0000B0000r028a1558.066clicks_L.mat');

                REFSa1_R = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/A0000B0000r028a1558.066clicks_R.mat');

                refsa1_L = REFSa1_L.clicks_L(signal_names{t});
                refsa1_R = REFSa1_R.clicks_R(signal_names{t});

                REFSb1_L = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/A0000B0000r028a1558.066clicks_L.mat');
                REFSb1_R = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/A0000B0000r028a1558.066clicks_R.mat');

                refsb1_L = REFSb1_L.clicks_L(signal_names{t});
                refsb1_R = REFSb1_R.clicks_R(signal_names{t});


                REFSa2_L = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/131a1552.480clicks_L.mat');
                REFSa2_R = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/131a1552.480clicks_R.mat');

                refsa2_L = REFSa2_L.clicks_L(signal_names{t});
                refsa2_R = REFSa2_R.clicks_R(signal_names{t});


                REFSb2_L = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/075b1559.017clicks_L.mat');
                REFSb2_R = load('/Users/cz/OneDrive - University of Waterloo/School/Vinyl_Project/data/clicks/075b1559.017clicks_R.mat');

                refsb2_L = REFSb2_L.clicks_L(signal_names{t});


                refsb2_R = REFSb2_R.clicks_R(signal_names{t});


                % need to do the reference here by track
                commonclicksa1_L = CommonClicks(CLICKS_L, refsa1_L);
                commonclicksa1_R = CommonClicks(CLICKS_R, refsa1_R);
```

```
                commonclicksb1_L = CommonClicks(CLICKS_L, refsb1_L);
                commonclicksb1_R = CommonClicks(CLICKS_R, refsb1_R);

                commonclicksa2_L = CommonClicks(CLICKS_L, refsa2_L);
                commonclicksa2_R = CommonClicks(CLICKS_R, refsa2_R);
                commonclicksb2_L = CommonClicks(CLICKS_L, refsb2_L);
                commonclicksb2_R = CommonClicks(CLICKS_R, refsb2_R);



                clicks_L = length(CLICKS_L);
                clicks_R = length(CLICKS_R);

                rmssig_L = rms(sig(:,1));
                rmssig_R = rms(sig(:,2));
                rmscsig_L = rms(csig(:,1));
                rmscsig_R = rms(csig(:,2));

                RMSclicks_L = 20.0*log10(sqrt(rmssig_L^2-rmscsig_L^2));
                RMSclicks_R = 20.0*log10(sqrt(rmssig_R^2-rmscsig_R^2));


                RMS_L = 20.0*log10(rms(csig(:,1)));
                RMS_R = 20.0*log10(rms(csig(:,2)));

                Aw = audio_Aweighting(csig);
                CCIRw = audio_CCIRweighting(csig);

                A_L = 20.0*log10(rms_response(Aw(:,1)));
                A_R = 20.0*log10(rms_response(Aw(:,2)));

                CCIR_L = 20.0*log10(avg_response(CCIRw(:,1)));
                CCIR_R = 20.0*log10(avg_response(CCIRw(:,2)));


                THD_L = thd(csig(:,1),fs);
                THD_R = thd(csig(:,2),fs);


                if ismember(signal_names(t), {'1kHzL', 'sweepL', '1kHzL2', 'sweepL2'})
                    stereo_bleed = StereoBleed(csig,1);
                elseif ismember(signal_names(t), {'1kHzR', 'sweepR', '1kHzR2', 'sweepR2'})
                    stereo_bleed = StereoBleed(csig,2);
```

```
                else
                    stereo_bleed = 0;
                end

                if ismember(signal_names(t), {'3150Hz', '3150Hz2', '1kHz', '1kHz2', '1kHzL', '1kHzL2','1kHzR', '1kHzR2', '1kHzV', '1kHzV2'})


                    [test_freq_L, wfreqspecamplitude_L, freqrms_L, WFrms_L] = WowFlutter(csig(:,1));
                    [test_freq_R, wfreqspecamplitude_R, freqrms_R, WFrms_R] = WowFlutter(csig(:,2));


                else
                    test_freq_L = 0;
                    test_freq_R = 0;
                    wfreqspecamplitude_L = 0;
                    wfreqspecamplitude_R = 0;
                    freqrms_L = 0;
                    freqrms_R = 0;
                    WFrms_L = 0;
                    WFrms_R = 0;

                end

                track = signal_names(t);

                output = [output; track, lagdiff, normalization_L, normalization_R, RMS_L, RMS_R, A_L, A_R, CCIR_L, CCIR_R, clicks_L, clicks_R, commonclicksa1_L, commonclicksa1_R, commonclicksb1_L, commonclicksb1_R, commonclicksa2_L, commonclicksa2_R, commonclicksb2_L, commonclicksb2_R, RMSclicks_L, RMSclicks_R, THD_L, THD_R, test_freq_L, wfreqspecamplitude_L, freqrms_L, WFrms_L, test_freq_R, wfreqspecamplitude_R, freqrms_R, WFrms_R, stereo_bleed];


            end
    end
```

A.4 Batch Process Script

BatchProcess.m looks at a folder of recorded test records that follow the proper naming convention and runs RecordProcess.m on each record in the folder appending the output of RecordProcess.m to the AudioTable which it eventually saves as the PRESSING-AudioTable.csv file.

```
folder = ('/audio_files/PRESSING/')
RecordNumbers = readtable('PRESSING_RecordNumbers.csv')

pressingID = 'PRESSING';

disp(['loading folder...:', folder])
files = dir(fullfile(folder,'*.wav'))


AudioTableHeaders = {'pressing','record', 'side','track', 'lagdiff', 'normalization_L', 'normalization_R','RMS_L', 'RMS_R', 'A_L', 'A_R', 'CCIR_L', 'CCIR_R','clicks_L', 'clicks_R', 'commonclicksa1_L', 'commonclicksa1_R','commonclicksb1_L', 'commonclicksb1_R', 'commonclicksa2_L', 'commonclicksa2_R','commonclicksb2_L', 'commonclicksb2_R',  'RMSclicks_L', 'RMSclicks_R', 'THD_L', 'THD_R', 'test_freq_L', 'wfreqspecamplitude_L', 'freqrms_L', 'WFrms_L', 'test_freq_R', 'wfreqspecamplitude_R', 'freqrms_R', 'WFrms_R','stereo_bleed'};


% check if there is already a csv file to append to
try
    disp('trying...')
    strcat(folder,strcat(pressingID,'-AudioTable.csv'))
    AudioTable = readtable(strcat(folder,strcat(pressingID,'-AudioTable.csv')))
catch
    disp('csv file not found, creating one...')
    AudioTable  = cell2table(cell(0,length(AudioTableHeaders)), 'VariableNames', AudioTableHeaders);
end
for i = (1:length(files)) %%loop through records

    filename = files(i).name;
    files(i);
    disp(['opening file...:', filename])
```

```
    if ismember(filename, AudioTable.record)
        disp('record already processed...')
        continue
    end

    file = strcat(files(i).folder,'/',filename);
    date_recorded = 0;
    pressing = 0;
    top_stamper = 0;
    top_hits = 0;
    bottom_stamper = 0;
    bottom_hits = 0;
    record = 0;
    side  = 0;
    track  = 0;

    recordid = str2num(filename(1:end - 13));
    disp([strcat('...recordid:', recordid)])


    pressing = RecordNumbers(RecordNumbers.RecordID == recordid, :);
    pressing = pressing.pressing{1};

    record = filename;
    side = filename(end-12);

    disp([strcat('...pressing:', pressing)])
    disp([strcat('...record:', record)])
    disp([strcat('...side:', side)])

    infoCell = {pressing, record, side};
    AudioOutput = RecordProcess(file);

    infoCell
    len = size(AudioOutput);
    len = len(1);
    for j = (1:len)
        AudioTable = [AudioTable; cell2table([infoCell, AudioOutput(j,:)
    ], 'VariableNames', AudioTableHeaders)];
    end
    disp('SAVING CSV')
    writetable(AudioTable, strcat(folder, pressingID,'-AudioTable.csv'));
end
```

A.5 Pressing Analysis Script

The final stage of the data processing, PressingAnalysis.m takes the AudioTable, SensorTable, RecordNumbers and SettingsTables and combines them into one master Pressing.csv file that lists all the available data on a per record basis. RecordNumbers.csv and SettingsTable.csv are two hand made csv files that track the settings in the press during the pressing process by matching up the pressing number with the planned parameters in the press so that the press settings can be recorded alongside the sensor and audio data.

```

AudioTable = readtable('PRESSING-AudioTable.csv');
SensorTable = readtable('PRESSING_SensorTable.csv');
RecordTable = readtable('PRESSING_RecordNumbers.csv');

SettingsTable = readtable('PRESSING_SettingsTable.csv')

pressruns = unique(AudioTable.pressing);
tracks = unique(AudioTable.track);

% clean up AudioTable
AudioTable.RecordID = erase(AudioTable.record,'.wav');

for i = 1:height(AudioTable)
    AudioTable.RecordID{i} = str2num(AudioTable.RecordID{i}(1:3));
end

AudioTable.RecordID = cell2table(AudioTable.RecordID);
AudioTable.RecordID = AudioTable.RecordID.Var1;



%~~~~~~~~~~~~~~~ MERGING TABLES STARTS ~~~~~~~~~~~~~~%

    Tbl = outerjoin(RecordTable, SensorTable);
    Tbl.Properties.VariableNames([1]) = {'PressingNumber'};
    writetable(Tbl,'Tbl1.csv')

    head(Tbl)
    head(AudioTable)

    Tbl(isnan(Tbl.RecordID), :) = [];
    % join in AudioTable
    Tbl = outerjoin(Tbl, AudioTable, 'Keys', {'RecordID', 'RecordID'});
    Tbl.Properties.VariableNames([1]) = {'PressingNumber'};
```

```
    writetable(Tbl,'Tbl2.csv')

    Tbl.Properties.VariableNames([3]) = {'pressing'};

    % join in SettingsTable
    Tbl = outerjoin(Tbl, SettingsTable);
    writetable(Tbl,'Tbl4.csv')

    Tbl.Properties.VariableNames([1]) = {'PressingNumber'};
    Tbl.Properties.VariableNames([33]) = {'track'};
    Tbl.Properties.VariableNames([3]) = {'pressing'};

    Tbl.Properties.VariableNames([2]) = {'RecordID'};

    Tbl.PressingNumber_SensorTable = [];
    Tbl.pressing_AudioTable = [];
    Tbl.pressing_SettingsTable = [];
    Tbl.RecordID_AudioTable = [];
    writetable(Tbl,'PRESSING.csv')
```

A.6 Wow and Flutter

The WowFlutter.m file takes an audio file that is a pure tone recorded from one of the test records and returns the measured frequency variation in the tone, the wow.

```

function [test_freq, wfreqspecamplitude, freqrms, WFrms] = WowFlutter(
   data);


    fs = 96000;
    tone = data(length(data)/3:(2/3)*length(data),:);
    lr=1; %1=left, 2=right
    tone=tone(:,lr);%now mono
    Nt=length(tone);
    t=(0:Nt-1)/fs;
    f=(0:floor(Nt/2))'*fs/Nt;
    TONE=fft(tone);
    [~,k]=max(abs(TONE(1:floor(Nt/2+1))));
    test_freq=(k-1)*fs/Nt;
    disp(['rough test freq [Hz]: ' num2str(test_freq)])

    decfactor=10;
```

```
tone=decimate(tone,decfactor);
fs=fs/decfactor;


toneh=hilbert(tone);
phase=unwrap(angle(toneh));
freq=(1/(2*pi))*diff(phase)*fs;%length Nt-1
freq=[freq;test_freq];% revert to length Nt
Nt=length(freq);
t=(0:Nt-1)/fs;
f=(0:floor(Nt/2))*fs/Nt;


[b,a]=butter(2,200*2/fs); %%% Change this number and try it out
freq=filtfilt(b,a,freq);%this now has some end effects

disp(['fs: ' num2str(fs) '  N_decimated: ' num2str(Nt) '  duration: ' num2str(Nt/fs)])

%% -----------------------AES WF-wtg table ----------------------------------
fr=[0.1 0.2 0.315 0.4 0.63 0.8 1.0 2.0 4.0 6.3 10.0 20.0 40.0 63 100 200 1000];
dBWFtable=[-57 -30.6 -19.7 -15 -8.4 -6.0 -4.2 -0.9 0 -0.9 -2.1 -5.9 -10.4 -13.7 -17.3 -23.0 -36.0];
%------------------------------------------------------------
%coefficients for use at fs=9600Hz, decimate has set that up
f1 = 11.0;%HF rolloff
f2 = 0.50;%LF rollup
f3 = 0.60;%LF rollup
f4 = 0.90;%LF rollup
WF4 = 1.10;%sets dB gain
X=[f1 f2 f3 f4 WF4];
%-------analog-digital W&F-weighting filter from filter convolution ---------
NUM = X(5)*[(2*pi)^3*X(2)*X(3)*X(4) 0 0 0];% s^3 character
DEN = conv(conv(conv([1 2*pi*X(2)],[1 2*pi*X(3)]) ,[1 2*pi*X(4)]), [1 2*pi*X(1)]);
% Bilinear transformation of analog design to get the digital filter.
[b,a] = bilinear(NUM,DEN,fs);

[H,w]=freqz(b,a,floor(Nt/2+1));
%----------------remove DC from freq---------------------
ns=round(Nt/10);
nf=round(9*Nt/10);
```

```
    freq=freq-sum(freq(ns:nf))/(nf-ns+1);% remove DC from WF sampled
   deviation array
    %-----------------get fundamental wow or flutter amplitude---------
    % 'freq': demodulated frequencies, sampled at fs, no DC, Nt points
    FREQ=(1/0.2155)*(2/Nt)*fft(freq.*flattopwin(Nt));%extra factor for
   flattop
    wfreqspecamplitude=max(abs(FREQ(10:round(Nt/2)-10)))%remove DC &
   Nyquist
    fw=(0:floor(Nt/2))*fs/Nt;%fft frequencies
    %--------------apply W&F weighting---------------
    WFfreq=filter(b,a,freq);%fs=9600Hz
    %----------------characterize W&F result-----------------
    freqrms=rms_response(freq(ns:nf));

    disp(['rms unweighted freq variation [Hz]: ' num2str(freqrms)])
    disp(['rms unweighted W&F: ' num2str(freqrms/test_freq)])
    disp(' ')
    %---------------now do WF weighted---------
    WFrms=rms_response(WFfreq(ns:nf));
    disp(['rms weighted freq variation [Hz]: ' num2str(WFrms)])
    disp(['rms weighted W&F: ' num2str(WFrms/test_freq)])
    disp('-------------------finished--------------------')
end
```

A.7 Click Detect

The following is a MATLAB adaption of Audacity's Click Removal Tool adapted to count the number of clicks and remove them.

```
function [csig, clicks] = ClickDetect(sig)
    threshold = 200;
    clickwidth = 20;
    csig = sig;

    sep = 2048;
    s2 = sep/2;

    b2 = sig.^2;
    ms_seq = b2;

    for ii = (1:floor(log2(sep)))
        i = 2^ii;
        for j =(1:length(sig)-sep)
```

```

            ms_seq(j) = ms_seq(j) + ms_seq(j+i);

        end
    end

    ms_seq = ms_seq./sep;



    clicks = [];
    left = 0;
    wrc = clickwidth/4;
    while wrc >= 1
        wrc = wrc/2;
        ww = clickwidth/wrc;
        for i = (1:length(sig)-2*sep);
            msw = 0;
            for j = (1:ww)
                msw = msw + b2(i + s2 + j);
            end
            msw = msw/ww;
            if  msw >= threshold*ms_seq(i)/10
                clickdetected = 0;
                if left == 0
                    left = i + s2;
                end
            else
                if(left ~= 0 && floor(i-left+s2) <= ww*2)
                    lv = sig(left);
                    rv = sig(i+ww+s2);
                    for j = (left:i+ww+s2)
                        if clickdetected == 0;
                            clicks = [clicks, j];
                            clickdetected = 1;
                        end
                        csig(j) = (rv*(j-left) + lv*(i+ww+s2-j))/(i+ww+s2
   -left);
                        b2(j) = csig(j).^2;
                    end
                    left = 0;
                elseif left ~= 0
                    left = 0;
                end
            end
```

```
        end
    end

end
```

A.8 Audio Functions

There are a few miscellaneous audio functions used in the code that require defining, namely the A and CCIR/ARM weighting filters and a spectrum shortcut function. They are defined in the audio_functions.m file and imported into the code when needed.

```
audio_functions.m

function data_A = audio_Aweighting(data)
        fs = 96000;
        f1 = 20.598997;
        f2 = 107.65265;
        f3 = 737.86223;
        f4 = 12194.217;
        f4 = 14100;
        A1000 = 1.9997;
        NUM = [ (2*pi*f4)^2*(10^(A1000/20)) 0 0 0 0 ];
        DEN = conv([1 +4*pi*f4 (2*pi*f4)^2],[1 +4*pi*f1 (2*pi*f1)^2]);
        DEN = conv(conv(DEN,[1 2*pi*f3]),[1 2*pi*f2]);
        % % Bilinear transformation of analog design to get the digital
    filter.
        [b,a] = bilinear(NUM,DEN,fs);
        data_A = filter(b, a, data);
    end


function data_CCIR = audio_CCIRweighting(data)
       fs=96000;
       frdc=[0 31.5 63 100 200 400 800 1000 2000 3150 4000 5000 6300 7100
    8000 9000 10000 12500 14000 16000 20000 25000 30000 fs/2];
       CCIR=[-inf -35.5 -29.5 -25.4 -19.4 -13.4 -7.5 -5.6 0.0 3.4 4.9 6.1
    6.6 6.4 5.8 4.5 2.5 -5.6 -10.9 -17.3 -27.8 -35 -50 -inf];
       Wn=2*frdc/fs;
       CCIRmag=10.^(CCIR/20);
       [b,a]=yulewalk(12,Wn,CCIRmag);
       [d,c]=butter(1,2*750/fs,'high');
       fb=conv(b,d);ea=conv(a,c);
       data_CCIR=filter(fb,ea,data);
```

```
    end


function [amp_coh, freq_coh] = audio_mscohere(data1, data2, fs)
    nfft=2^14;
    [amp_coh, freq_coh] = mscohere(data1,data2,hanning(nfft),nfft/2,nfft,fs);
end


function [data_fft, freq_fft] = audio_spectrum(data, fs, start_sam, n_sam);
    data_fft = (fft(data(start_sam:start_sam+n_sam-1, :))/n_sam);
    data_fft = data_fft(1:floor(n_sam/2)+1,:);
    data_fft(2:end-1,:) = 2*data_fft(2:end-1,:);
    freq_fft = fs*(0:(n_sam/2))/n_sam;
end


function [smoothed_tf]=pwroctsmooth(freq_tf,octave_width)

if octave_width>=2.0
   octave_width=0.33;
   disp(['oct_width set to ' num2str(octave_width)])
end
an=2^(octave_width/2);

N=length(freq_tf); % should work for even or odd
smoothed_tf=zeros(size(freq_tf));% making this complex causes long execution!
old_lo=0;old_hi=0;
for I=1:floor(N/2)+1 % fast smoothing algorithm
   lo_bin=round((I-1)/an)+1;hi_bin=round((I-1)*an)+1;
   if hi_bin==lo_bin
      pwr_sum=abs(freq_tf(I))^2;% no change
   else
      for J=old_lo:(lo_bin-1)
        pwr_sum = pwr_sum - abs(freq_tf(J))^2;
      end
      for K=(old_hi+1):hi_bin
        pwr_sum = pwr_sum + abs(freq_tf(K))^2;
      end
   end
    smoothed_tf(I) = sqrt(pwr_sum/(hi_bin-lo_bin+1));
    old_lo=lo_bin;old_hi=hi_bin;
```

```
end
smoothed_tf(floor(N/2)+2:N)=conj(smoothed_tf(ceil(N/2):-1:2));% ensure
    conjugate even
if N/2==floor(N/2) % even N
  smoothed_tf(N/2+1)=abs(smoothed_tf(N/2+1));% Nyquist should be real
end

function[smoothed_tf]=pwroctsmooth_singlesided(freq_tf,octave_width)
% function[smoothed_tf]=pwroctsmooth(freq_tf,octave_width)
%
% fractional-octave pwr preserving smoothing. Max width<2.0 octaves
% freq_tf is complex single- sided spectrum, length N, end @ Nyquist.
% smoothed_tf is returned, real single-sided, size N.
% We assume that the double-sided array is even, so Nyquist is N.
%
% John Vanderkooy Oct, 2021
if octave_width>=2.0
   octave_width=0.33;
   disp(['oct_width set to ' num2str(octave_width)])
end
an=2^(octave_width/2);
N=length(freq_tf); % should work for even or odd N
% Nyquist will be at Matlab index N, DC is index 1.
smoothed_tf=zeros(size(freq_tf));% if complex causes long execution?
old_lo=0;old_hi=0;
for I=1:N % fast smoothing algorithm
   lo_bin=round((I-1)/an)+1;
   hi_bin=round((I-1)*an)+1;
   if hi_bin==lo_bin
      pwr_sum=abs(freq_tf(I))^2;% no change
   else
      for J=old_lo:(lo_bin-1)
        pwr_sum = pwr_sum - abs(freq_tf(J))^2;%forget these
      end
      for K=(old_hi+1):hi_bin
        index=K+(K>N)*2*(N-K);%makes it 2N-K if K>N
        pwr_sum = pwr_sum + abs(freq_tf(index)^2);%add these
      end
   end
    smoothed_tf(I) = sqrt(pwr_sum/(hi_bin-lo_bin+1));
    old_lo=lo_bin;old_hi=hi_bin;
end
```

Appendix B

Additional Plots and Tables

There are a lot of plots and data needed for the analysis and conclusions in this thesis. Multiple tracks, sides, sensor measurements and more mean that a lot of conclusions need to be drawn from a large pool of data. In Chapter 4 the results from the measurements of the two test pressings are given. Additional conclusions are drawn in that Chapter that require many more plots to show. This Appendix contains those plots.

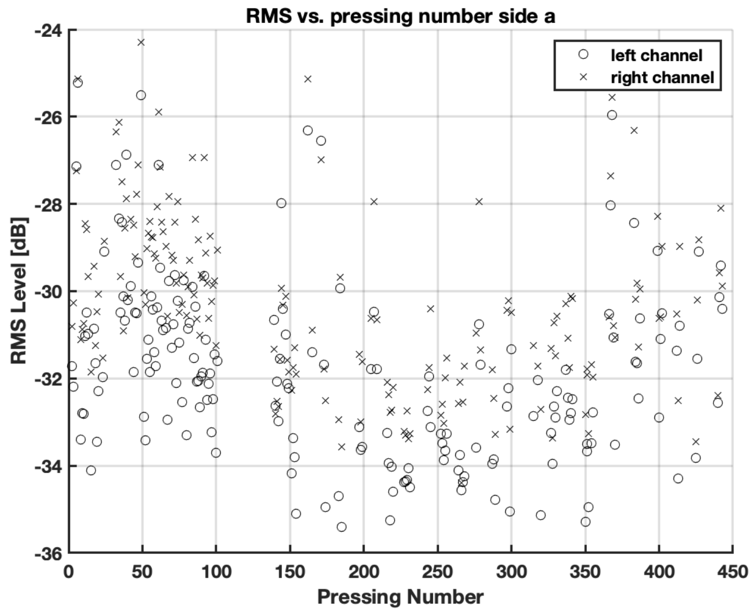

Figure B.1: The Unweighted RMS level of the quiet track for side a of all records recorded.

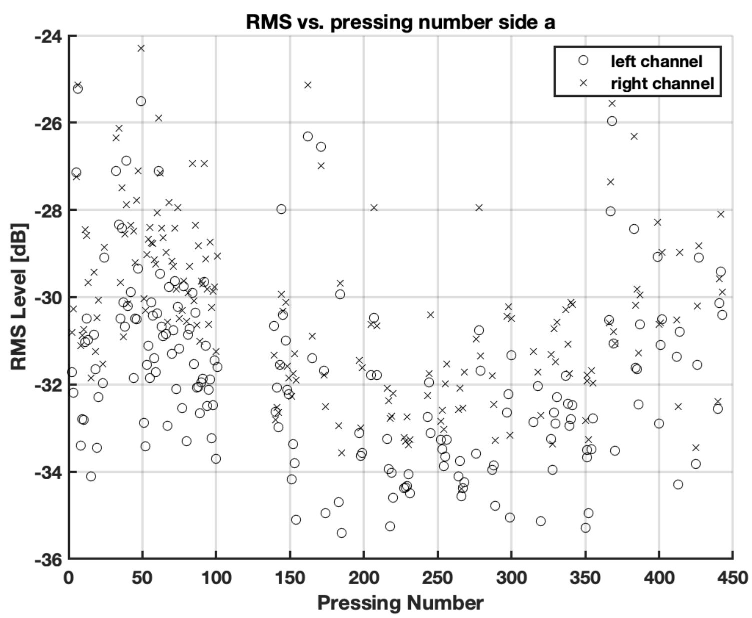

Figure B.2: The Unweighted RMS level of the quiet track for side b of all records recorded.

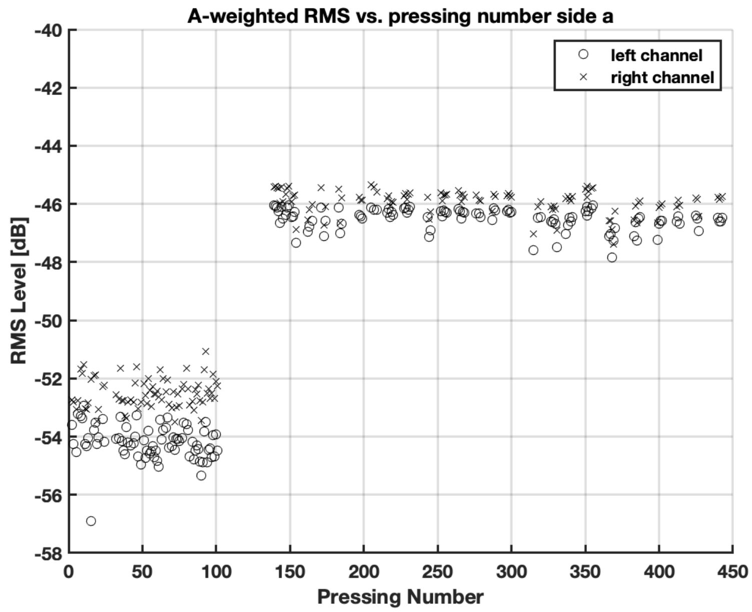

Figure B.3: The A-weighted RMS level of the quiet track for side a of all records recorded.

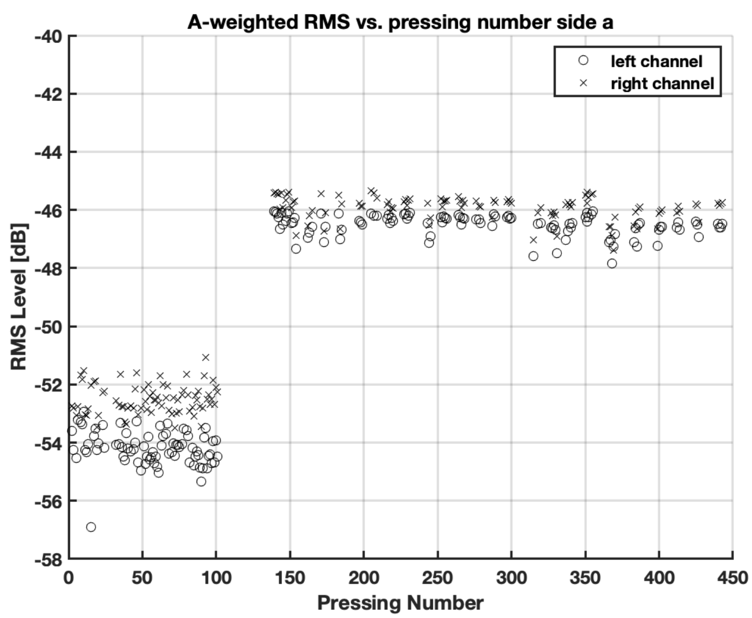

Figure B.4: The A-weighted RMS level of the quiet track for side b of all records recorded.

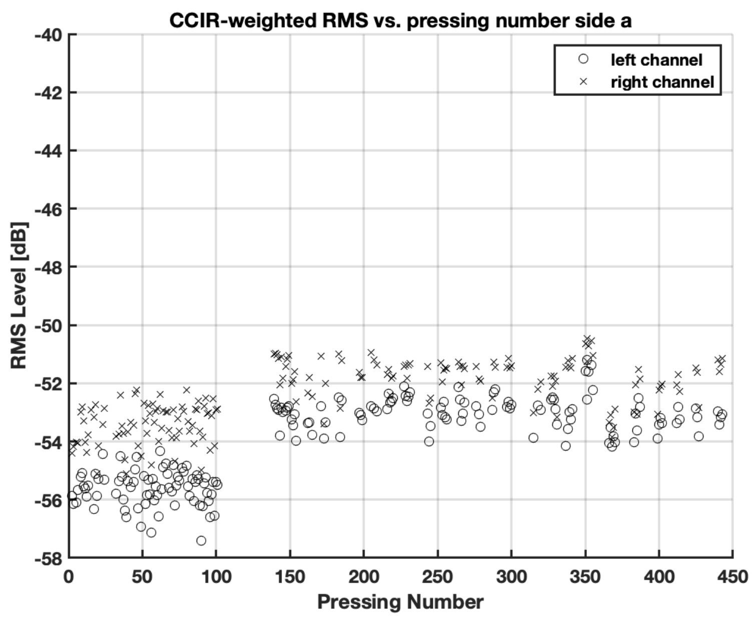

Figure B.5: The CCIR-weighted RMS level of the quiet track for side a of all records recorded.

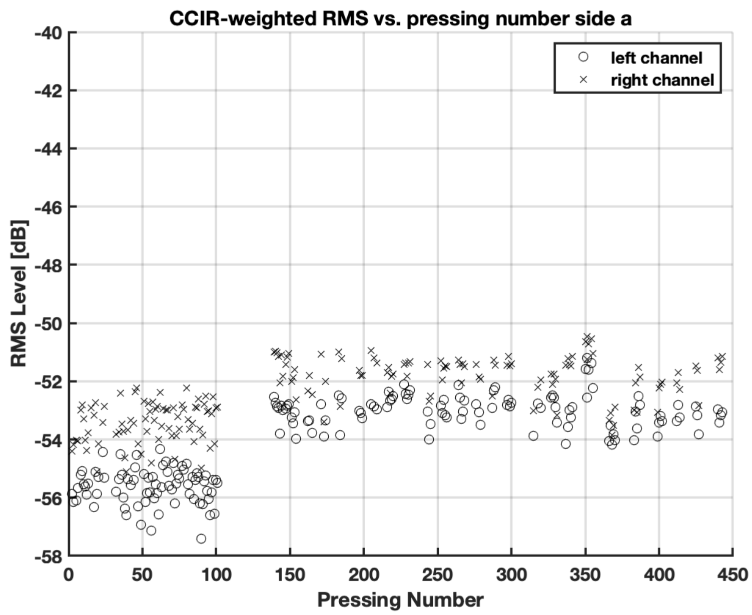

Figure B.6: The CCIR-weighted RMS level of the quiet track for side b of all records recorded.

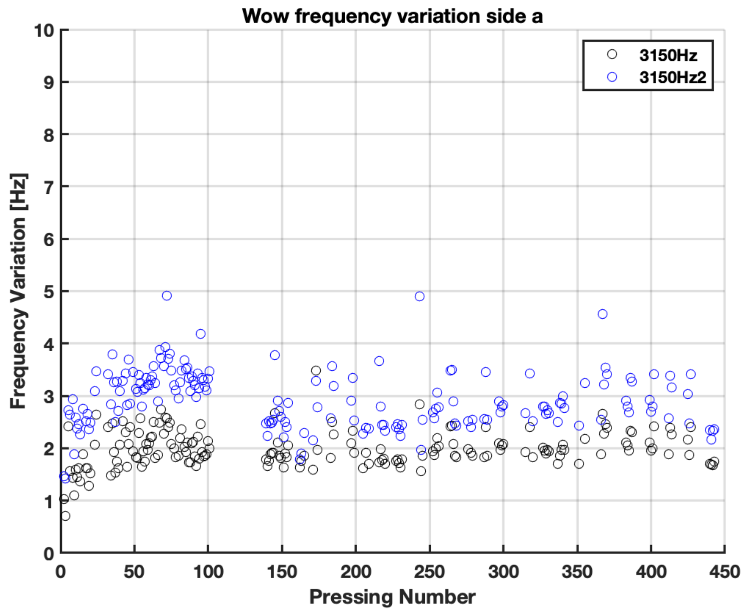

Figure B.7: The wow of all the records listed in pressing order for side a.

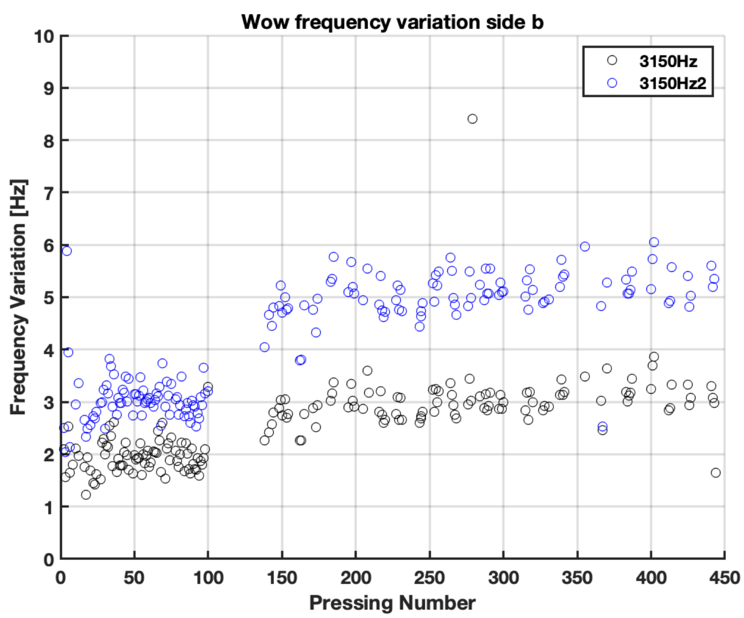

Figure B.8: The wow of all the records listed in pressing order for side b.