A Pre-Characterized Toolkit for Vibrotactile Feedback

Aditya Tirumala Bukkapatnam

Input Devices and Music Interaction Laboratory
Music Technology Area
Schulich School of Music
McGill University
Montreal, Canada

July 2019

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Masters in Arts.

2019/07/27

Abstract

Vibrotactile feedback - a type of haptic feedback, is a popular area of research. Aside from widespread applications in video games and mobile phones, vibrotactile feedback is also implemented in Digital Music Instruments (DMIs) to improve the quality of interaction for performers, or as general purpose interfaces for conducting experimental research. Musical interaction is known to involve high cognitive expertise, requiring responsive and highly accurate tools. When implementing such devices, the outputs of the vibrotactile implementations are characterized and equalized. However, these implementations are often "one-of-a-kind", bulky (>8 kg), expensive (<1000 $), and impractical to some musical applications, whereas, inexpensive and portable alternatives often lack characterized and equalized outputs. The latter also rely on actuators that cannot display the various types of stimuli that are possible in musical interactions. This study describes the design and implementation of a cost-effective and portable vibrotactile toolkit with a characterized output. The goal of this toolkit is to allow users to prototype and study vibrotactile feedback in musical applications. Many perceptual and technological factors influence the design of vibrotactile systems meant for musical applications. By using these factors as guidelines, suitable actuators, and amplifiers are selected for implementing the toolkit. The outputs of these devices are measured to obtain their magnitude-frequency and harmonic distortion characteristics. In this study, harmonic distortion is characterized by Total Harmonic Distortion (THD) for actuators and Total Harmonic Distortion + Noise (THD+N) for amplifiers. All the tested actuators and amplifiers are found to operate in the full vibrotactile bandwidth (40-1000 Hz). The actuators are measured both when placed on a sandbag (unloaded) as well as when attached to an external surface (loaded). The characteristics of the actuator change with loading; Resonant features appear in the frequency response, along with an overall increase in the harmonic distortion. Previously, the resonant characteristics were equalized manually using parametric equalizers. Unlike previous cases, the autoregressive Yule-Walker method is proposed to equalize the resonant characteristics automatically. From the tested actuators and amplifiers, an inexpensive (approximately 200 $) and portable (<40 g) 2-channel toolkit is implemented. The implemented toolkit is characterized, and the results are presented in a datasheet. Lastly, vibrotactile feedback is implemented in the chassis of an unfinished T-Stick DMI using the toolkit, and its frequency response is automatically equalized using an accelerometer measurement.

Résumé

Le retour vibrotactile est un type de retour haptique très étudié dans le milieu de la recherche. En plus d'être utilisé dans les applications pour jeux vidéos et téléphones mobiles, le retour vibrotactile se retrouve aussi dans la conception d'instruments numériques de musique (INM) afin d'améliorer l'interaction entre l'instrumentiste et l'instrument, et également étudié en recherche expérimentale en tant que simple interface. La pratique d'un instrument de musique implique une expertise cognitive élevée, nécessitant des outils justes et précis. Pour réaliser de tels outils avec un retour vibrotactile, les signaux de sortie des systèmes de traitement de signal vibrotactile doivent être caractérisés et égalisés. Cependant, ces implémentations sont souvent à usage unique pour chaque cas d'étude isolé, encombrantes (>8 kg), coûteuses (>1000 $), et non applicables à certains cas d'utilisations musicales, alors que les alternatives peu coûteuses et portables manquent souvent d'égalisation ou n'ont pas été proprement caractérisées au sens du traitement de signal qu'elles emploient. De plus, ces applications utilisent des technologies d'actuation qui manquent de versatilité pour la génération d'une variété de stimuli vibrotactiles nécessaires pour des interactions musicales. Cette recherche vise à concevoir une boîte à outils peu coûteuse et portable pour la génération de retour haptique vibrotactile, dans le but de permettre à ses utilisateurs de prototyper rapidement et étudier différents types de retours vibrotactiles pour diverses applications musicales. De nombreux facteurs perceptuels et technologiques influencent la conception des systèmes vibrotactiles. En considérant ces facteurs, des technologies d'actuation et amplificateurs sont sélectionnés pour concevoir la boîte à outils. Les signaux de sortie de tels outils sont mesurées pour obtenir leurs caractéristiques de magnitude-fréquence et de distorsion harmonique. Dans cette étude, la distorsion harmonique est caractérisée par un taux de distorsion harmonique (THD en anglais) pour les actionneurs et un taux de distorsion harmonique plus bruit (THD+N en anglais) pour les amplificateurs. Tous les actionneurs et amplificateurs testés utilisent la largeur totale de la bande vibrotactile (40-1000 Hz). Les actionneurs sont mesurés dans deux situations : 1) placés sur un sac de sable (non chargé), et 2) fixés à une surface externe (chargée). Les caractéristiques de l'actionneur changent selon la charge, c'est-à-dire que les caractéristiques de résonance apparaissent dans la réponse en fréquence, avec une augmentation globale de la distorsion harmonique. Auparavant, ces caractéristiques de résonance étaient égalisées manuellement à l'aide d'égaliseurs paramétriques. Dans la présente recherche, la méthode autorégressive de Yule-Walker est proposée pour égaliser automatiquement les caractéristiques de résonance. À partir des actionneurs et amplificateurs testés, une boîte à outils à deux canaux, d'environ <200 $ et pesant <40 g, est caractérisée et les résultats obtenus sont présentés dans une fiche technique. Finalement, à l'aide de la boîte à outils, le retour vibrotactile est intégrer dans la structure d'un INM T-Stick, où sa réponse en fréquence est automatiquement égalisées grâce à un accéléromètre.

Acknowledgements

This elaborate journey was nowhere close to being one of valor or heroism - quite the contrary, if anything. Thanks Dad, for the strength to persevere through all the odds. To Mom and my brother Dhruva, for that unrelenting confidence and support. To my partner Gillian, for being an amazing listener and for being the silent clockwork that keeps the universe around me moving incessantly.

To my supervisor Marcelo Wanderley, for the motivation, encouragement and freedom to explore, pursue, fail and therein truly learn. To Philippe Depalle for offering me priceless insight and helping me look in the right places. To Yves Méthot, Romain Dumoulin, Marcello Giordano, and Harish Venkatesan for being ready volunteers in all those times of technical crisis, and to Ian Marci, Christian Frisson, and Catherine Massie-Laberge for all your help with words!

Lastly, my colleagues and dearest friends, both near and far, for making every beer count.

Contents

List of Figures

List of Tables

Abbreviations

AC Alternating Current. 14, 15, 16, 17, 20, 33, 35, 58, 81

DAC Digital to Analog Converter. 25, 64, 81

DC Direct Current. 14, 41, 59, 62

DMI Digital Music Instrument. i, 1, 4, 22, 26, 59, 63, 69, 72

DUT Device Under Test. 20, 22, 31, 32, 33, 34, 35, 39, 40

ENBW Equivalent Noise Bandwidth. 39

ERM Eccentric Rotating Mass. 14, 15, 16, 17, 22, 25

FA-I Fast Adapting-I. 10

FA-II Fast Adapting-II. 10

FFT Fast Fourier Transform. 39, 40

FIR Finite Impulse Response. 64

HCI Human-Computer Interaction. 1, 4

IMD Intermodulation Distortion. 24

LPC Linear Predictive Coding. 55

LRA Linear Resonant Actuator. 14, 16, 21, 25

NP Non-Pascinian. 9, 10

NP-I Non-Pascinian-I. 9, 10

NP-II Non-Pascinian-II. 9, 10

NP-III Non-Pascinian-III. 9, 10

P Pascinian. 9, 10

PC Pascinian-corpuscle. 6, 8, 10, 11, 12, 21, 25

PCB Printed Circuit Board. 59

PSD Power Spectral Density. v, 39, 40, 41, 44, 56, 64, 65, 67, 69

PVC Poly Vinyl Chloride. 63, 64, 72

PWM Pulse Width Modulation. 14, 15, 16

RA Rapidly Adapting. 6, 7, 9, 10, 11, 12

RMS Root Mean Square. 12, 23, 33, 35, 40, 41, 43, 44, 56

SA-I Slowly Adapting-I. 7, 9, 10, 12

SA-II Slowly Adapting-II. 7, 9, 10, 12

THD Total Harmonic Distortion. i, ii, v, vii, 23, 24, 26, 32, 40, 41, 42, 43, 44, 50, 52, 53, 57, 59, 60, 61, 62, 68, 69, 70, 71

THD+N Total Harmonic Distortion + Noise. i, ii, v, vii, 24, 29, 30, 32, 38, 40, 41, 44, 50, 60, 71, 81, 82

VCA Voice Coil Actuator. 14, 16, 17, 18, 19, 20, 21, 25, 27, 28, 36, 42, 43, 50, 52, 53, 59, 60, 61, 62, 63, 64, 67, 68, 70, 71

Chapter 1

Introduction

Haptic feedback is now a popular form of feedback across various disciplines of Human-Computer Interaction (HCI). Vibrotactile feedback is a sub-category of haptic feedback that pertains to using vibrating actuators to stimulate the surface of the skin [1]. From mobile devices to virtual reality and other domains, vibrotactile feedback is widely used to enhance the quality of interaction between humans and technology. Similarly, musical applications have also seen benefits from the inclusion of vibrotactile feedback. An increasing number of Digital Music Instruments (DMIs)[1] are known for including vibrotactile feedback [3], while seeking an improved quality of engagement between the performer and the instrument [4, 5].

However, vibrotactile feedback has perceptual and technological constraints that influence its implementation [6]. Neurophysiological and the psychophysical studies on the sense of touch shed light on the perception of tactile sensations in the human body. Tactile perception is found to be dependent on factors like the pitch of the vibration (frequency), the intensity of vibration, the location of stimulus on the body (body site), the type of skin (hairy or glabrous), the size of the contactor (large or small), and pressing forces on the vibrating surface [7].

In musical applications, there is a need for high output accuracy [8, 9] over a wide bandwidth of frequencies [10], along with the requirement to independently control the amplitude and frequency

[1]Gestural controllers interfaced with sound synthesis/production software, with the help of mapping strategies [2].

of the vibrotactile display [4, 11]. In order to meet these requirements, we begin by selecting actuators that allow us the required degree of control and then characterizing their outputs using accelerometer measurements before determining their suitability to our application [5].

General purpose vibrotactile interfaces play an important role in researching the role of perception for musical applications, where the characterization of the vibrotactile output benefits designers and researchers with design validation and experimental data [8]. However, the existing interfaces are either "one-of-a-kind" implementations [8] that are characterized but are also bulky and expensive [11, 12]. Existing toolkits offer cheaper and more portable alternatives but are neither characterized nor allow independent control of amplitude and frequency [13, 14]. Such factors make the existing implementations impractical [15], particularly in musical applications.

In this study, we explore this requirement for a more portable toolkit that can readily display characterized vibrotactile feedback for musical applications. Following is a chapter-wise summary on the implementation of a toolkit that could address some of these requirements:

Chapter 2 explores the background on the perceptual and technological factors that influence the design of vibrotactile implementations for musical applications. We begin with a brief survey of the neurophysiology and psychophysiology of the sense of touch and some of the perceptual factors that arise. We then survey guidelines on selection of vibrotactile technologies, methods to characterize and equalize the vibrotactile output, and assessment of the fidelity of the system.

Chapter 3 begins implementing the toolkit by selecting amplifiers and actuator technologies that meet the demands of musical applications. Subsequently, output measurements are used to characterize the magnitude-frequency response and fidelity of each of the selected devices, to determine their suitability for the toolkit.

Chapter 4 examines the effect of loading actuators on their charactersitics. Musical interaction involves a high cognitive expertise [16], and tools developed for music require high-accuracy [9]. In this context, we review an automatic method to equalize the frequency characteristics of a vibrotactile implementation to improve its accuracy.

Chapter 5 uses the results from the characterization of the selected devices to implement a toolkit as per the requirements outlined in Chapters 2 & 3. We present the implemented toolkit with a data-sheet that informs an end-user of the capabilities and limitations of the implementation. We then explore using implemented toolkit to display vibrotactile feedback in a potential musical application. The automatic equalization reviewed in Chapter 4 is used here to equalize the frequency response of the implementation.

Chapter 6 presents concluding remarks on the characterization and equalization process, the importance of characterizing distortion in vibrotactile implementations, and brief comments on the observations of characterizing loaded actuators.

Chapter 2

Review of Literature

2.1 The Tactile Channel

Over the years, DMIs have come up as a category of computer-based musical instruments wherein the sound production of the instrument and the controlling interface connect through a set of *mappings* that transform the performer's input gestures into a sound output [17]. Due to this decoupling, an important channel of feedback – the tactile channel [2] – is rendered absent to a performer, which forces reliance on the other modalities of feedback [18].

The haptic channel of feedback is said to be comprised of a *tactile sense* and a *kinesthetic sense* [19]. The tactile sense is associated with perceiving properties of a surface or object, whereas the kinesthetic has more to do with the state of one's body and other properties like position and velocity. Together, the two senses constitute the *haptic channel* in HCI research, and *haptic feedback* pertains to using the haptic channel as a mode of feedback to the user [18]. Haptic feedback is speculated to be among the few modalities that respond fast enough for controlling articulation in music [6, 20]. It is also key to the interaction between the player and the instrument since it renders a quality of interaction that is closer to that of a traditional musical instrument, as opposed to "feeling" like an open-looped computer interaction [3]. Vibrotactile feedback has been explored in attempts to restore this missing channel [5, 18, 21, 22].

2.2 Vibrotactile Feedback

Vibrotactile feedback is an active area of research in musical applications. Marshall [3] recorded 15 DMIs using vibrotactile feedback in the proceedings of NIME[1] between years 2001-2008. Alongside the quest for vibrotactile feedback in DMIs, research also investigates vibrotactile feedback in non-DMI musical applications like educational or learning tools [21, 23] through the use of more general purpose vibrotactile interfaces [24]. Such interfaces help investigate other ways of exploiting the sense of touch for feedback purposes [4, 11]. Characterizing such general purpose interfaces and iteratively improving their design helps produce more reliable and comparable results across devices [8].

A discourse on vibrotactile feedback would remain incomplete without a preliminary discussion on the neurophysiological and psychophysical aspects of the human somatosensory system, followed by the technological considerations that influence the design of vibrotactile systems.

2.2.1 Neurophysiology of Touch

The human sense of touch is essentially a microcosm, where a large number of variables vitally influence something seemingly trivial [4]. The neurophysiological layout of touch helps detect different types of vibrations based on specific characteristics of the stimulus such as the amplitude, the frequency, and the contact area of the stimulus [3]. The following is at best, one distilled discussion on the neurophysiological layout of touch. This insight helps inform oneself of the challenges in the study of vibrotaction, and the design of vibrotactile interfaces.

The understanding of the human sense of touch begins with the study of subcutaneous *mechanoreceptors* among the layers of the skin. Four types of mechanoreceptors contribute to our perception in the tactile channel, which are namely *Merkel disks, Meissner & Pascinian corpuscles* and the *Ruffini endings* [25]. These mechanoreceptors lie innervated with nerve fibers, or afferents[2], which conduct the neuronal impulses of each receptor to the central nervous system [7]. *Mylenated* af-

[1] New Instruments for Musical Expression.

[2] Afferents carry neuronal pulses from the peripherals to the brain; Efferents carry pulses from the brain to the peripherals [26].

ferents in skin bear a high conduction velocity of $35 - 75$ m/s and fiber diameters $6 - 12$ μm. Conversely, *Unmylentated* afferents possess a lower conduction velocity of $0.5-2$ m/s, with smaller diameters of $0.2 - 1.5$ μm, and are more commonly used to conduct stimuli of slow pain and thermal sensations [26]. All touch mechanoreceptors are fused solely with *Mylenated* afferents[3][27]. However, the afferent nerve fibers and their respective mechanoreceptors are each selective to specific information about the stimulus [7].

Across the skin on the human body, glabrous skin (non-hairy skin) is the most sensitive, and much of neurophysiological and psychophysical investigations on tactile perception focus on the glabrous skin of the palm or feet. This focus is because the primary touch interface between humans and their surrounding environment is essentially through glabrous skin [7]. Hairy skin has been understood to have higher thresholds of perception and lower peak sensitivity [6], with physiological differences in the types and distribution of mechanoreceptors [27]. Recent research also shows that stimuli at the fingertip can elicit neuronal activity from pascinian afferent units in the volar arm [28], indicating the presence of more complex phenomena within the tactile channel as a whole. Therefore the following discussions in this dissertation regarding touch and its perception focus on the widely researched effects found in the study of glabrous skin.

The mechanoreceptors and related afferents were studied in awake humans (and earlier in animals) using electrophysiological methods. A probe depressed onto the skin was observed to yield a variety of action potentials among the different nerve fibers, recorded through a microelectrode implanted inside the skin, near the nerve fiber. For a given receptor and its nerve fiber, when the depth of indentation of the probe is held constant, the rate of decrease in the frequency of its action potentials over time is called its *adaptation rate*. Four neural systems are identified based on the types of nerve fibers, the receptors they innervate, their adaptation rate, and the areas of their receptive fields [7]:

1. Pascinian-corpuscle (PC) fibers having large receptive fields, each fiber ending in a single Pascinian corpuscle, and adapting very rapidly [26],

[3] More specifically type $A\beta$.

2. Rapidly Adapting (RA) fibers with small receptive fields, where a single fiber can innervate one or more Meissner corpuscles,

3. Slowly Adapting-I (SA-I) fibers with small receptive fields, where each fiber innervates multiple Merkel disks, and lastly,

4. Slowly Adapting-II (SA-II) fibers with large receptive fields, ending in SA-II end organs[4].

The occurrence of action potentials in the nerve fiber or its stimulated receptors is not single-handedly responsible for tactile perception. There also exist minimum rates of occurrence of spikes in the fibers, for each fiber type, and a minimum number of activated fibers, if a stimulus is to be perceived by its human observer. E.g., One spike per RA fiber and activation of 5-10 RA fibers was found necessary for detection in the RA system. Similarly, several SA-II fibers would need to be activated, with a spike rate of at least five spikes per second in the SA-II system, or 0.8 spikes per second for the SA-I system, to register stimuli respectively [7]. Therefore, the aforementioned neural systems are more closely understood to function as *afferent units*[4, 29], where each system has a threshold of intensity for perception and a frequency selectivity function that is ultimately determined by the frequency selectivity of its end-receptor type [7].

Psychophysical investigations on human perception offer a more holistic overview, unifying the layout of mechanoreceptors and their afferent fibers into independent information-processing channels in the central nervous system. These channels work like filters, responding to very specific aspects of the stimulus before relaying it to the brain [7]. The following section examines the psychophysical aspects of touch based on the channels of perception.

2.2.2 Psychophysics and the Channels of Perception

The introduction of the Microneurography technique in 1970, allowed the direct study of human subjects, with a higher resolution than what was available in the previously electrophysiological experiments. As a result, it became possible to study the psychophysical relations between

[4]Previously thought to be Ruffini endings, which were later found to be missing in glabrous skin [7].

the primary afferent activity and the result perceived by a human[29], shedding light on various perceptual factors relevant to vibrotactile implementations [6].

2.2.2.1 Spatial and Temporal Summation

The effect of ***spatial summation*** was investigated using an apparatus consisting of a hole within a rigid surround, where a circular vibrating actuator could protrude through while a participant's palm rested on the other side of it. Actuators of different surface area ($< 0.02-2.9\ cm^2$) were thus interfaced with the thenar eminence[5] of the palm, using the apparatus. A drop in the threshold of perception was observed, in correlation with an increase in the contact area, i.e., the sensitivity of the tactile system improved for an increase in the stimulated area on the thenar eminence. This effect was only observed for contact area 0.02 cm^2 and higher (tested up to 2.9 cm^2), and frequencies of vibration above 40 Hz [7]. In the range of 40-800 Hz, displacement thresholds are generally observed to decrease by approximately 3 dB, for a doubling of contact area from $0.002 - 5.1\ cm^2$ [30].

Alongside spatial summation, for a 250 Hz stimulus applied to the thenar eminence using a large 3 cm^2 contactor with varying stimulus duration, the threshold of detection decreased for an increase in duration up to 1 second. When the same stimulus was applied through a smaller 0.01 cm^2 contactor, the detection threshold was observed to be independent of stimulus duration. This was referred to as ***temporal summation***. Further evidence for temporal summation is that the phenomenon was found absent for a 30 Hz stimulus on the thenar eminence, which does not stimulate the PC system. The same holds when a 250 Hz stimulus is applied to the dorsal surface of the tongue that is known to lack Pascinian corpuscles [7].

The discovery of the spatial[6] and temporal[7] summation phenomena are reason to believe that there exist at least two channels for detecting disturbances on the skin surface. This two-channel system is known as ***the duplex model of mechanoreception***. The two channels could also be

[5] Prominence on the palm, at the base of the thumb.
[6] Spatial Summation: Increased sensitivity to a stimulus, for an increase in the area of stimulus.
[7] Temporal Summation: Decrease in detection threshold, for an increase in duration of stimulus.

stimulated individually, and were outlined for glabrous skin as follows [7]:

1. The Pascinian (P) channel, ending in Pascinian corpuscles, which was capable of spatial and temporal summation, and exhibiting a U-shaped sensitivity curve across frequency, and,

2. The Non-Pascinian (NP) channel, ending in other mechanoreceptors, and whose sensitivity across frequency does not vary with spatial or temporal summation.

2.2.2.2 The Four Channel Theory

In later developments, studies theorized and tested a *Four Channel Theory*, where the sense of touch is shown to function as four channels of perception [25]. The phenomena of enhancement, adaptation, and masking can validate that the sense of touch is stimulatable in four separate channels, i.e., the stimulation in one channel did not affect the other channels [7]. The NP channel of the duplex theory now comprised of three NP systems namely Non-Pascinian-I (NP-I), Non-Pascinian-II (NP-II) and Non-Pascinian-III (NP-III), each processing one unique neural system. The NP-I channel processing the RA neural system and similarly, NP-II with SA-II, and NP-III with SA-I neural system.

In experiments testing the enhancement effect, observers were asked to match the perceived intensities of a comparison stimulus to that of a test stimulus, wherein for some trials, a conditioning stimulus preceded the test stimulus. The purpose of the conditioning stimulus was to assess whether it had any effect on the perceived threshold. It was revealed that observers would set the intensity of the comparison stimulus to a relatively higher level when a conditioning stimulus preceded the test stimulus, revealing an ***enhancement of perceived sensitivity*** due to the presence of a conditioning stimulus. The amount of enhancement[8] was higher when the conditioning and test stimuli were presented nearly simultaneously, showing that the neural effects of a first stimulus could persist up to 400-500 ms after its cessation. Therefore, the enhancement effect is only observed if the second stimulus occurs while the neural effects of the first persist.

[8] Amount of enhancement is defined as the difference between the observed intensities, of the matches made in the presence and absence of the conditioning stimulus [7].

Furthermore, this enhancement effect is only observed when both stimuli excite the same neural system. However, the enhancement phenomenon does not occur as summing of impulses in nerve fibers but rather, higher up in the central nervous system, confirming that channels function in the perceptual domain, and not merely at the physiological level [7].

Experiments with ***masking*** and ***adapting*** stimuli were conducted to determine frequency selection of each of the four channels. Two contactor sizes of 0.008 cm^2 and 2.9 cm^2 were used to produce test stimuli for specific channels. Each test stimulus has a fixed frequency and an intensity that lies well above the threshold of detection for the tested channel. Then a masking stimulus is delivered with varying intensity until the test stimulus is "barely detectable". The intensities where the test stimulus is masked are recorded for different frequencies, establishing the frequency selectivity for the respective channel [7]. The results are tabulated in Table 2.1, illustrating the frequency selectivity of the different channels, as a function of contactor size. Furthermore, masking and adaptation effects occur within channels, and not across channels. Each of the channels has also been understood to play specific functional roles, akin to "feature detectors" i.e., each responding only to specific aspects of a stimulus [7]:

1. The SA-I or NP-III channel, encoding spatial details with strong resolution in detecting "spatially distributed amplitudes of skin displacement",
2. The SA-II or NP-II channel, responsive to skin stretching,
3. The RA or NP-I channel, responding to the velocity of a vibratory stimulus, thereby detecting its waveform rather than the energy of the vibration, and lastly,
4. The PC or P channel, which has a superior capability to detect vibration energy.

The different nomenclature for channels are often used interchangeably. The P and NP convention, are sometimes known by the names of the governing neural system, i.e. the PC, RA, SA-I and SA-II nomenclature. In other cases, the RA and PC channels are referred to as Fast Adapting-I (FA-I) and Fast Adapting-II (FA-II) respectively, on the basis of the adaptation property [1], where the *FA* acronym sometimes expands to *Fast Afferent* [6], as opposed to *Fast Adapting*.

Detecting Channel	Frequency Bandwidth (Hz)	
	Large Contactor	Small Contactor
SA-I	0.4 - 2	0.4 - 2
RA	2 - 40	2 - 100
SA-II	-	>100
PC	>40	Threshold too high for detectable PC thresholds

Table 2.1 Frequency bandwidths and their respective detecting channels, for varying contactor sizes [7].

2.2.2.3 Perception and Vibration Frequency

The frequency bands for each of the channels are not precise and are discrepant across literature [1]. However, the PC channel is widely considered to be responsible for mediating the perception of vibration from 40 Hz and upward [7]. The perception of vibration in humans is also regarded to extend up to 1000 Hz [31], and possibly higher among the hearing-impaired [32]. The sensitivity to a vibrating stimulus is also known to increase monotonically with frequencies up to 250 Hz, and decrease monotonically after that (shown in Figure 2.1) [6, 31], with 12 dB/octave slopes given by the PC channel [7]. These U-shaped contours describe the perceived intensity across different frequencies and have characteristics that are similar to the Fletcher-Munson loudness curves of the auditory modality [6]. However, this response is expected to vary with actuator size [3, 7].

The RA channel perceives the waveform of the stimulus as opposed to the actual vibration and frequencies delivered to it are perceived as "Flutter" [7]. Furthermore, the RA channel's Meissner corpuscles help with 2-point discrimination and location precision [26].

The perception of frequencies is also known to occur in distinct bands. In some cases, 3-5 bands in the range of 2-300 Hz range [18], and 8-10 bands between 70-1000 Hz [34] are informally suggested [6]. With regards to a complex stimulus waveform, a variety of textures ranging from smoothness to roughness were observable when using spectrally rich content, and bursts with fast attacks were discovered to be useful when cueing boundary crossings in "open-air" musical

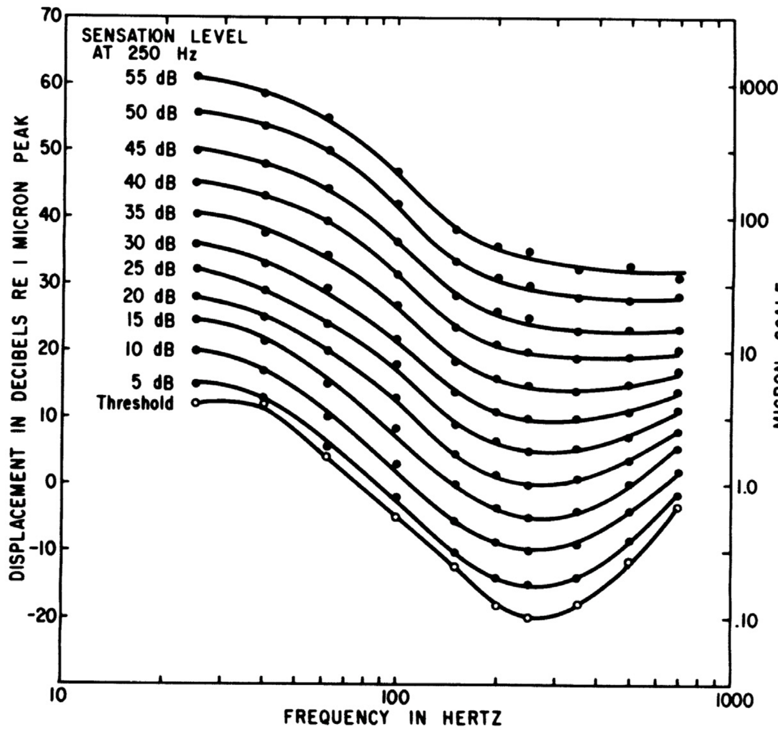

Fig. 2.1 Vibrotactile threshold curves reported by Verrillo [31] showing U-shaped contours that describe the sensitivity perceived glabrous skin to vibrating stimuli at different frequencies. The curves show that sensitivity increases monotonically up to 250 Hz and then decreases monotonically up to 1000 Hz. Reproduced from Verrillo et al. [33] with permission. ©*Psychonomic Society, Inc.* 1969.

instruments [18]. Spectrally rich content is more effectively perceived [31] and the use of amplitude modulation is notably effective in modeling rough features [6].

The following table 2.2, indicates the frequency bands for each of the channels, with their receptive field sizes, and encoded property in the perceptual domain:

2.2.2.4 Perceptual Thresholds

Complex waveforms like square-waves or having harmonically rich spectra are perceived differently than pure sinusoidal stimuli [6, 18], and the perceptual thresholds for such waveforms are found to be consistently lower than those of sine tones at the same frequency [32, 35].

Using a modified version of the TouchBox interface, Fontana et al. [30] found that the finger

Channel	SA-I	SA-II	RA	PC
Receptive Field	Small	Large	Small	Large
Frequency Range (Hz)	<5	15 - 400	3 - 100	10 - 500[a], (>) 1000[b]
Perceived Properties	Pressure, Fine Details	Stretch	Flutter	Vibration

Table 2.2 Channels of vibrotactile perception, their bandwidths, size of their receptive fields, and the perceived property upon stimulation. Adapted from [1, 3, 6, 12, 31, 32].

[a] 500 Hz: [1].
[b] 1000 Hz:[3, 6, 12, 31], 2000 - 4000 Hz: [32].

pressing force could significantly lower the threshold of perception. Participants were required to apply approximately 1.9, 8, and 15 N of force, termed as low, mid and high force conditions, onto the vibrating surface of the TouchBox interface. A load cell placed under the TouchBox interface measured the force, and an LCD screen turned green to signal when the participant reached the target force. The vibrating stimuli presented were either a white noise band-limited between 50-500 Hz or a sine wave at 250Hz. The Root Mean Square (RMS) acceleration of the surface was measured and expressed in dB RMS acceleration using the reference level 10^{-6} m/s^2, indicated in the ISO R 1683 specification [12]. In contrast to earlier research, the observed threshold of perception was found to be much lower for finger pressing conditions. The lowest thresholds were found to be around 68.5 dB RMS acceleration for the 250 Hz sinusoidal stimulus in the high force condition, and around 83.1 dB RMS acceleration for the band-limited white noise stimulus in the low force condition. Overall, both results were much lower than the earlier reported 105 dB RMS acceleration [30, 31].

2.3 Technological Considerations for Vibrotactile Interfaces

Alongside perceptual factors, the type of actuator technology and the signals driving the actuators are essential factors that influence the design of haptic interfaces for musical applications [6]. The interfaces designed for musical applications are ideally unobtrusive, consume low power, offer high accuracy [15, 22], introduce minimum latency [3], and offer decoupled amplitude and frequency controls [5, 11].

2.3.1 Actuator Selection Criteria

Vibration stimuli are typically delivered through miniature vibrating actuators, i.e., transducers that can produce a vibration controlled by an input voltage. However, these actuators have different working principles that make them more suitable for some applications, and less suitable to others. Below is an inexhaustive list of commonly encountered vibrating actuators, and their working principles [3, 8, 12]:

1. *Solenoid*: These actuators produce a longitudinal motion, that corresponds to an Alternating Current (AC) input signal (or sometimes Direct Current (DC) [5, 6].)

2. *Eccentric Rotating Mass (ERM) motors*: These actuators typically consist of an unbalanced mass attached to the shaft of a DC motor, which exhibits vibration when activated with a direct input voltage. When provided with a Pulse Width Modulation (PWM) type of input signal, the amplitude and frequency of the vibration are simultaneously controllable by the duty cycle of the PWM. Due to their reliance on PWMs signals, these actuators are considered to have low power requirements. However, these actuators also take time to reach the target speed due to their inertia, making them unsuitable for displaying rapid or transient information.

3. *Voice Coil Actuator (VCA)*: These actuators typically consist of either a moving coil or moving magnet assembly, similar to loudspeakers, and are controlled by AC signals. Moving magnet VCAs are particularly of interest to vibrotactile applications, due to their ability to produce higher energy in the lower frequency bands.

4. *Linear Resonant Actuator (LRA)*: These actuators are similar in principle to moving mass VCAs, with the principal difference of being tuned to resonate strongly at predefined frequencies (often 250Hz), with the help of a moving mass-spring assembly. These actuators are driven using AC signals and are known to be power efficient, often requiring lower power than ERM motors.

5. *Piezoelectric Actuator*: These actuators are constructed using materials that deform when excited with either a DC or AC signal, but are typically known to require higher orders of voltages, up to a few hundred volts. However, the current requirement is often relatively lower, and there exist driving circuits that can adapt audio signals (AC signals) for use with such actuators. These actuators can often possess a wide bandwidth of operation above 100 Hz, and age minimally due to low-friction operation from the absence of moving parts.

Actuators also typically have unique frequency characteristics and specific driving signal and power requirements. Across the board, solenoids produce vibrations at a single intensity, while the frequency can be varied if being driven by an AC input signal. ERM motors are driven with simpler PWM signals and consume low power (typically 3-5 V) but do not allow elaborate control over amplitude and frequency of the vibration [3]. On the other hand, "loudspeaker-like" actuators require dedicated amplifier and power supply systems to drive them [6], but offer independent control of amplitude and frequency of vibration, better transient response[9] and a superior frequency resolution[10] than solenoids or ERMs [5]. The frequency controllability factor gives rise to two kinds of vibrotactile systems: (1) A variable frequency system and, (2) a fixed frequency system, depending on whether or not the frequency of vibration is controllable [3]. Notably, variable frequency systems are more often driven with complex audio signals (AC signals) rather than PWM signals [3, 6]. Marshall [3] suggests that a variable frequency system also needs to be capable of displaying at least an 18% change in frequency to maximize frequency discriminability, and operate in the vibrotactile bandwidth of 40-1000 Hz with a dynamic range of at least 60 dB to accommodate compensation of the U-shaped frequency response of the skin.

Giordano [12] surveyed various musical features like pitch, rhythm, and timbre, being presented in the tactile channel. A taxonomy was later developed for the different types of vibrotactile feedback encountered in musical interactions, with guidelines informing the selection of actuators for each type, while also suggesting that both amplitude and frequency of the stimulus are sometimes

[9] Transient response tested at 100, 200 and 500 Hz input signals [3].
[10] See Section 2.3.2 for details on frequency resolution.

key parameters when it comes to vibrotactile feedback in musical applications. This taxonomy was also recently updated to form design guidelines for vibrotactile interfaces [6, 24]:

1. *Tactile Notification*: Related to alerting users to events in the surrounding environment or results of interacting with a system, with simple suprathreshold stimuli, or complex temporal envelopes and spatial patterns. This kind of feedback is typically implemented with low power actuators like ERM motors, being driven by PWM signals.

2. *Tactile Translation*: Related to cross-modal mapping of features from sound to touch by leveraging the neurophysiological similarities between the auditory and tactile modalities of feedback (sensory substitution), or simulating the behavior of structures whose vibrational properties are already characterized (tactile stimulation). This type of feedback typically requires “loudspeaker-like” actuators that offer the ability to control the amplitude and frequency of the vibration independently.

3. *Tactile Synthesis*: Related to the creation of “languages solely addressed to the sense of touch” for mediating abstract ideas through the skin, ranging from simple notifications to icons built using frequency, intensity, envelope, spectral content, and the waveform of the vibration, as well as using different locations on the body of the user as building blocks of the developed language. This type of feedback relies on loudspeaker-like actuators so that the various linguistic parameters are deliverable.

As seen in the taxonomy mentioned above, as well as in a later section (2.3.3), there is often a need to be able to vary the amplitude or the frequency of the vibration independently of each other. This is not possible with ERMs, LRAs or solenoid actuators. In such scenarios, we need to rely on “loudspeaker-like” actuators that are AC driven, such as VCAs or piezoelectric actuators.

2.3.2 Characterization of Vibrotactile Systems

The Oxford dictionary defines *characterization* as “a description of the distinctive nature or features of...something” - in this case, the mechanical and electrical characteristics of a vibrating system. Papetti et al. [8] examine the benefits of characterization in interfaces designed for musical research. In vibrotactile applications, characterization can help produce verified data to users, by describing the true nature of the rendered feedback. In turn, this allows refinement and better interpretation of experimental results, particularly those that are psychophysical [8].

An actuator's mechanical characteristics are definable by its frequency response between 40-1000 Hz, its frequency resolution, and the amplitude response over its complete amplitude range[11] [3]. Marshall [3] performed these measurements using a reference accelerometer affixed to a thin wooden board, and the test actuator affixed on the other side of the board, directly opposite the accelerometer. The accelerometer signal was amplified with a gain of 10, using a signal conditioner, which also provides the bias current required for operating the accelerometer. The frequency response of AC controlled actuators was measured by feeding the actuator with control signals of constant amplitude for test frequencies spaced at 1/4th octave between the range of 40-1030 Hz. The amplitude response was measured at a single frequency, and the frequency resolution was measured by incrementing the frequency of the input control signal by 0.1 Hz until a change in output frequency was detected. In the results of Marshall's [3] characterization, VCAs were found to offer the widest bandwidth (40-1000 Hz), while piezoelectric actuators suffered in low and high-frequency regions, below 100 Hz and above 500 Hz respectively. All the tested actuators were found having an adequately low-frequency resolution for frequency discriminability, where VCAs, *tactors*[12] and piezoelectric buzzers offered superior resolutions of 1%, 1%, and 1.2% respectively, compared to solenoids (5.6%) and ERMs (4.3%).

Giordano [12] additionally characterized the frequency content of some amplifiers used for driving such actuators. The *Bryston 2B-LP* reference amplifier did not significantly alter the frequency content of the input signal, introducing a low noise at -120 dB/Hz, and preserving the overall frequency-amplitude distribution. It was found with sufficient fidelity to serve as a reference for characterizing other amplifiers and actuators [12]. In order to test the amplifiers, each

[11] Defined as the range between the minimum amplitude for a perceivable vibration, and the maximum vibration amplitude offered by the actuator.

[12] Specialized actuators designed to produce vibration at a specific frequency, usually 250 Hz [3].

amplifier was input with a 0.2 V logarithmic sinesweep signal, and the output was terminated with a 6Ω load resistance to simulate the impedance of the test VCA[13]. An RME Fireface 400 audio interface operating at a 48 kHz sample rate was used to record the output voltage. The VCAs were measured by feeding each actuator with the same 0.2 V sinesweep used for the amplifiers. A reference accelerometer was placed on the shell of the actuator using an adhesive petrowax such that the sensing axis of the accelerometer and the vibrating axis of the actuator were aligned and the actuator is then mounted onto a plastic surface. The frequency response and the frequency content of the devices were evaluated with the help of a transfer function calculation, computed as the ratio of output to input spectra of the test device.

Papetti et al.'s [8] *Touch-Box* characterizes the effect of finger pressing forces and finger area on the vibrotactile output being delivered (see Section 2.3.3 for details on implementation). This interface was later used by Fontana et al. [30] to measure how pressing down on a vibrating surface influences the vibrotactile sensitivity threshold (see Section 2.2.2.4). The interface has a moving magnet VCA[14] mounted to a plexiglass surface with its vibrating axis set perpendicular to the surface [8]. A 250 Hz sinusoid and a band-limited white noise (50-500 Hz) were used as test signals to test the accuracy of the vibrotactile output. A measurement accelerometer was mounted to the top (vibrating) surface of the interface using double-sided tape, to measure the vibration produced by the interface. An audio interface recorded the signals from the accelerometer, and the RMS acceleration was calculated over an observation interval of 8 seconds and expressed in dB, using a reference of 10^{-6} m/s^2. As mentioned earlier, the finger area in contact with the vibrotactile surface, as well as the influence of force on the vibration are also characterized. The finger's pressing forces were recorded by a CZL635 load cell (capable of measuring 49 N) feeding a 10-bit Arduino UNO microcontroller. The microcontroller transmits the force data to the *PureData* software, which generates vibrotactile responses in the form of audio signals. A Plexiglass with infrared LEDs mounted to its side is used to illuminate the area of the finger in contact with the surface. An infrared camera mounted under the glass captures the area of the finger to determine

[13] The Haptuator Original by Tactile Labs, and a custom-made replica of the same.
[14] *Haptuator mk-II* by *Tactile Labs*.

the area of contact of the finger. The LED light propagates in the plexiglass and is dispersed when a finger comes in contact with the plexiglass surface, allowing the finger area to be characterized [36]. The finger pressing force was simulated using test weights, and the resulting dampening effect was characterized. By characterizing the amount of force applied by a participant and the participant's finger area, the interface compensates for the dampening effect of the pressing action and always records the true RMS acceleration of the surface, as if it were undamped.

The *VibroPiano* is another interface developed by Papetti et al. [8], to assess how the nature of vibrotactile feedback can affect pianists' performance and their perception of quality features. A keyboard component mounted to a plywood board couples higher vibration energy from the actuators to the keys. Two *Clark Synthesis* TST239 VCAs were affixed to the bottom of the wooden board, split between lower and middle ranges of the keyboard. Accelerometer measurements were performed on all the A keys yielding an average magnitude spectrum, where the "keyboard-plywood" system, in combination with the transducer's frequency response, showed spurious resonances with a peak at 300 Hz. These resonances were later equalized (See section 2.4).

The *HSoundplane* is a Soundplane interface[15], augmented with haptic feedback by Papetti et al [8]. Piezo actuators were added to the interface so that it generates localized vibrotactile feedback along its grid-like surface [8]. This device was developed to realize a flexible framework that would help experiment with various audio-tactile mappings and assess their effectiveness in musical practice. Sinesweeps between 20-1000 Hz were used to characterize the final interface's frequency and amplitude content [8]. The final prototype was additionally found producing localized vibration at intensities that are well above the lowest known vibrotactile thresholds [8], recorded between 68-83.1dB RMS acceleration (reference 10^{-6} m/s^2) under the effect of pressing forces [30].

[15] A low-latency (few ms) multi-touch musical interface with pressure sensitive surface and capacitive sensing [8].

2.3.3 Previous Vibrotactile Interfaces and Toolkits

As mentioned earlier, general purpose vibrotactile interfaces are also commonplace in applications for music technology research. We now explore the components that drive the actuators in general purpose vibrotactile interface implementations to help us better illustrate the need for an inexpensive and portable toolkit that readily serves musical applications.

High-fidelity components are used to ensure an accurate output in implementations. In Papetti et al.'s [8] *Touch-Box* interface (see Section 2.3.2 for construction details), the generated audio signals are delivered to the VCAs of the interface via an RME Fireface 800 audio interface and an audio amplifier[16]. Papetti et al.'s [8] *VibroPiano* interface is another such example. The tactile transducers here are driven by the *Yamaha P2700* amplifier, which is fed audio signals via an RME Fireface 800 audio interface. Giordano [12] also used a high-fidelity reference amplifier (*Bryston 2B-LP*) to characterize the frequency content of amplifiers and actuators. An RME Fireface 400 audio interface was used for digital to audio conversion, to produce analog AC signals for feeding the amplifier or actuators. Egloff et al. [11] prepared a similar apparatus, to help present musical intervals as vibrotactile stimuli to participants in an ABX test. The apparatus used a *Dayton Audio* 13 mm VCAs, driven by a similar high-fidelity Bryston amplifier. The amplifier was fed audio signals generated on a Macbook, presented via a Gracedesign audio interface.

Many of the implementations above are seen to rely on high-fidelity audio interfaces and amplifiers. These high-fidelity components are mainly used to ensure a reliably high-fidelity signal to characterize the Devices Under Test (DUTs). However, these devices are expensive, costing in the range of hundreds to thousands of dollars, and also overkill in terms of the power required to drive the vibrotactile actuators. The Yamaha P2700 used by Papetti et al. [8] is a high-fidelity amplifier that outputs 1000 W of power into an 8 Ω impedance, when operating in mono mode, and weighs 24 kg[17]. The *Bryston 2B-LP* amplifier used by Marshall [3], Giordano [12] and Egloff et al. [11] weighs approximately 8 kg[18], and can output up to 100 W of peak power into a 4 Ω

[16] Amplifier model is not mentioned.
[17] `https://www.manualslib.com/manual/339870/Yamaha-P2350.html?page=11`
[18] `http://www.bryston.com/PDF/brochures/2BLP_BROCHURE.pdf`

impedance.

There also exist vibrotactile implementations that are less expensive than the ones mentioned above. The *Techtile Toolkit* is one example of an affordable vibrotactile toolkit that uses a microphone, amplifier, and an LRA signal chain, to recreate textures picked up by the microphone, with the actuator [13]. While the amplifier of this toolkit has a wide bandwidth of operation, it relies on LRA actuators of much narrower bandwidth, ultimately limiting the bandwidth of the system; This can be undesirable in certain musical applications. *Stereohaptics* is a similar portable toolkit developed at *Disney Research* that uses VCAs and a two-channel amplifier. However, this toolkit is not characterized and provides no information on its accuracy or suitability to musical applications. Tools for musical performance require high accuracy [9], which is ascertained through the characterization and equalization of the vibrotactile output when designing vibrotactile interfaces (discussed in Section 2.4) [8].

2.4 Frequency Response Equalization

Compensating the actuator's frequency response plays a vital role in the perception of the vibrotactile feedback. Marshall [3] carried out a study with 10 participants that evaluated a perceived change in frequency on a 5 point Likert scale where the stimuli were presented in three conditions: (1) No frequency compensation, (2) Only actuator response compensated, and (3) Full compensation, where the actuator's and the skin's U-shaped frequency response are both equalized. A significant improvement in frequency discriminability and participant's confidence ratings were observed when some form of frequency compensation was applied. Furthermore, results show a higher percentage of accurate judgments when only the actuator's response is compensated, compared to the case with full compensation. While participants made less accurate judgments in the case with full compensation, they are reported to have shown higher confidence in judgments made for fully compensated trials [3]. In another ABX schemed experiment conducted by Egloff et al. [11], participants were given vibrating stimuli at different frequencies that represented standard musical intervals. The results suggest that participants could have mistaken the variation in perceived intensity for a difference in frequency, indicating the need to equalize the stimuli to neutralize the effects of perceptual effects of the PC channel. This equalization is similar to the full compensation case explored by Marshall [3], which is only possible for actuators that permit control of amplitude, without altering the predetermined frequency of the vibration.

Actuators that permit the independent control of amplitude and frequency have the advantage that their frequency response can be modified or compensated [3]. Ryu et al. [37] attempted to equalize ERM based systems, where the perceived intensity of output vibration was encoded in terms of the required input voltage. This encoding was mainly done to enable users to design with the final perceived intensity as their design parameter, which is decoded into the appropriate input voltage for the actuator. Nevertheless, the fundamental property of the ERM does not change; The controls of amplitude and frequency of vibration ultimately remain coupled.

Giordano [12] and Marshall [3] both developed DMIs where parametric equalizers were used to equalize the output frequency response. Papetti et al.'s [8] *Touch-Box* and *VibroPiano* implementations also implement frequency response equalization using parametric equalizers. In each of the previous cases, the frequency characteristics are analyzed from acceleration measurements of the vibrating interface, which informs a choice of parameters to configure a parametric equalizer. The parametric equalizer is configured to approximately model the inverse of the interface's measured response, which then modifies the signal driving the actuators such that it evens out any resonant features that were observed in the initial response of the actuator [8].

The question is then, how do we perform equalization. While many approaches exist in system identification literature, we have attempted to automate this equalization process using one such method. The approach adopted uses the Yule-Walker equations to equalize the frequency response using a white noise measurement, and the accuracy of the method is expressed using a spectral flatness estimation. The method was proposed in a paper co-authored by Philippe Depalle, Marcelo Wanderley, and myself as first author [38]. This method is revisited in Chapter 5, to equalize a hypothetical DMI implementation using the toolkit developed in this thesis.

2.5 Non-linearity as a Factor

Giordano [12] characterized the frequency content as an indicator of the fidelity of a DUT i.e., either an actuator or an amplifier. Components of the interface were identified introducing harmonic distortion into the signal and understood to be of lower fidelity. In the output of a high-fidelity device, the output is found to be an accurate reproduction of the input and largely undistorted. Giordano [12] compared the non-linear operation of amplifiers and vibrating actuators by analyzing spectrogram plots of the input and output of the DUTs. For a given input sinesweep signal, the output was observed having additional sinusoidal components that varied with the instantaneous frequency of the input. Low power and inexpensive amplifier circuits used for powering vibrotactile actuators were shown to be of significantly lower fidelity and discouraged from use where fidelity is a requirement [12].

This non-linearity is often characterized for amplifiers [39], and actuators [40] in the form of Total Harmonic Distortion (THD). THD is calculated as a ratio of the RMS magnitude of the additionally introduced harmonic frequencies at the output, to the magnitude of a fixed, input fundamental frequency, observed at the same output. While taking into account n harmonic frequencies above the fundamental, the n^{th} order THD relation is given as [41, 42]:

$$THD(\%) = \frac{\sqrt{\sum_{i=2}^{n} a_i^2}}{a_1} * 100$$

where,

a_i = magnitude of the i^{th}frequency component, and

a_1 = magnitude of the fundamental frequency.

On this note, the THD of actuators is typically calculated by considering the first five harmonics above the input fundamental [12, 40]. In a study comparing different vibrotactile systems having different amounts of THD, Maeda and Griffin [41] show that the harmonic distortion in a system can influence the perceived vibrotactile threshold, or more generally, influence the results obtained.

The harmonic distortion in high-fidelity amplifier systems such as the *Bryston 2B-LP* and *Yamaha P2700* is often expressed as Total Harmonic Distortion + Noise (THD+N) by their manufacturers, which is a metric that also accounts for the electrical noise introduced by the system alongside the harmonic distortion [43]. The harmonic distortion introduced by some amplifiers could sometimes result in perceived vibrations due to spectral components introduced by the amplifier being suprathreshold stimuli. In an experiment by Salisbury et al. [35], the authors used two amplifiers of different harmonic distortions to output a 40 Hz pure tone and demonstrated that one of the two systems was perceivable due to a suprathreshold 80 Hz (harmonic) component. Maeda and Griffin [41] and Wyse et al. [32] also reported that vibration perception thresholds could be affected due to additional spectral components introduced by the system in the vibration output.

However, THD or THD+N are not comprehensive measures of non-linearity. More specifically, they quantify only the harmonically related components of the introduced distortion. THD further depends on the magnitude of the input fundamental that can often be outputted at lower levels than the harmonics, possibly leading to falsely high measures of THD, found to be particularly true when operating at lower frequencies [42]. In audio and electroacoustic systems, research assessing non-linearity has extended to both technological and perceptual effects that try to find more accurate ways of measuring the perceptual effects of non-linearity [43]. Intermodulation Distortion (IMD) is another common measurement that characterizes the non-linearity of the system for a two-tone excitation [42], but this characterization is a study in itself, requiring an informed selection of the two pure tones that sufficiently characterize the system. Hence, IMD characterization presently remains outside the scope of this work.

2.6 Discussion

While caring for some of the psychophysical and technological complexities surveyed across this chapter, researchers devise apparatuses and experiments that further investigate perceptual, psychophysical, and technological factors at play, forming an iteratively informed approach to designing for the tactile channel [4, 8, 11, 44]. Stimuli for the tactile notification paradigm can always be implemented by a variable frequency system, with the drawback of needing redundantly higher power requirements than a fixed frequency system. However, for a vibrotactile toolkit, the ability to display a variety of stimuli remains a necessity. Hence, we will focus on implementing a variable frequency system.

While the perception of frequencies may extend as low as 0.4-2 Hz in psychophysical experiments [7], frequencies below 5 Hz are reportedly perceived as pressure (see table 2.2) [1]. It is clear that the PC channel predominantly mediates stimuli from larger contactors (2.9 cm^2 [7]) at frequencies above 40 Hz, perceived as vibration. The PC channel is also the primary mediator of vibrotactile musical stimuli [28]. However, vibration is perceived only for stimuli frequencies in the range of 40-1000 Hz on glabrous skin [31]. Since independent controllability of vibration amplitude and frequency is a critical aspect in our toolkit, we rely on loudspeaker-like actuators for the system, that can operate in the vibrotactile bandwidth, and offer a good transient response and frequency resolution. These requirements can be met by choosing VCAs and piezoelectric actuators.

Power efficient, unobtrusive and high-fidelity interfaces are a growing need in upcoming vibrotactile interfaces [22]. However, current interfaces are either particularly obtrusive due to their reliance on bulky and expensive devices, or are mostly ERM based [10, 24], or LRA based [13]. These bulky systems are not only non-portable but also overkill in terms of power, often driving relatively low-power and miniature actuators. A Digital to Analog Converter (DAC), an amplifier, and an actuator are the primary components required for producing the various types of vibrotactile feedback encountered in musical interactions. However, we have seen that the amplifiers and the DACs are the primary contributors of bulk and cost among implementations. These devices render many of these existing interfaces into "one-of-a-kind" implementations that can only be reproduced with a high budget.

Vibrotactile interfaces are typically characterizable in terms of their amplitude response, frequency response, and fidelity (i.e., frequency content). Low fidelity systems introduce non-linear components into the vibrotactile signal, thereby compromising the integrity of the driving signal and affecting the results obtained with such a system. Therefore the fidelity of a vibrotactile system may further be quantified in terms of THD and is expressed as a function of its driving input voltage.

There is a need for an affordable and portable interface that can produce a reliable[19] output over a wide bandwidth (40-1000 Hz), which relies on loudspeaker-like actuators. In the following chapters, we attempt to implement a toolkit to address this requirement in the following chapters in a step by step approach, first choosing affordable amplifiers that can deliver a low order of harmonic distortion, and a selection of commonly encountered loudspeaker-like actuators. We then characterize these devices for their magnitude-frequency response and harmonic distortion. Based on the results of characterization, we select the devices having an optimal combination of portability, cost, magnitude-frequency and fidelity to implement a vibrotactile toolkit and provide the results of its characterization as technical documentation, for avoiding a "one-of-a-kind" implementation. The toolkit is then used to implement vibrotactile feedback in a DMI chassis, illustrating an automatic equalization approach to equalize the final frequency response of the DMI's vibration characteristics.

[19] With equalized and good fidelity output characteristics.

Chapter 3

Device Selection and Characterization

In this thesis, we are implementing a vibrotactile toolkit for prototyping vibrotactile feedback in musical applications. Given the factors outlined in Chapter 2, the requirements of the toolkit are summarized as follows:

1. *Cost and Portability*: We aim to realize a toolkit that weighs approximately $\leq$1 kg and can be implemented for a few hundred dollars.

2. *Decoupled Amplitude & Frequency Controls*: The toolkit must be able to serve applications where independent control of amplitude and frequency is required, as per the demands of musical applications. To ensure this, we select VCA and piezoelectric actuators that can operate in the vibrotactile bandwidth of 40-1000 Hz.

3. *Characterization*: Characterization offers insight on the accuracy of the system. With musical applications in mind, we aim to characterize our implementation in terms of its magnitude-frequency and harmonic distortion characteristics, in the vibrotactile bandwidth.

4. *Equalization*: Equalization of the vibrotactile output is a key step in musical applications. Our toolkit implementation should explore automated methods to facilitate equalization.

This chapter explores the selection of devices that will help achieve a toolkit within the above requirements. We then explore the characterization of these selected devices, to determine if they are suitable choices for implementing the toolkit. Lastly, we perform vibration and electrical measurements to characterize their magnitude-frequency responses and their respective amounts of harmonic distortion within the vibrotactile bandwidth.

3.1 Device Selection

The implementation of a vibrotactile toolkit that meets the above requirements begins with a selection of devices. We first survey actuators and amplifiers based on the technological requirements mentioned above. This section will outline the selection criteria and present a discussion on the devices selected for the implementation of the toolkit.

3.1.1 Actuator Selection

In order to implement this toolkit system, we begin by selecting actuators that make ideal choices for implementing vibrotactile feedback in musical applications. We select the VCA and piezoelectric actuators as shown in table 3.1, which can serve the vibrotactile bandwidth. The actuators are listed along with their respective costs, loading impedances, maximum rated voltages, operating powers, and the rated bandwidths of operation.

The actuators listed in Table 3.1 can be operated up to 1000 Hz, except for the *TDK PowerHap 7G* piezo actuator, whose range is limited to 500 Hz by its self-heating property. This piezo actuator requires a biasing voltage of around 120 $V_{\text{p-p}}$ for operation, which is provided by additional biasing circuitry such as the *TI DRV2700EVM*[1].

The selected VCAs require approximately 3-6 V_{pk} to operate at their maximum rated voltages, and the biasing circuit would drive the piezo actuators. Using Ohm's law, we can estimate that amplifiers that can output about 1-5 W of power into an 8 Ω impedance or higher will output around 6 V_{pk} unloaded voltage and will have sufficient power to drive the selected VCAs. The

[1] These specifications listed in Table 3.1 are the combined properties of the *TDK PowerHap 7G* actuator and the *TI DRV2700EVM* module, since they need to operate as a unit.

Device	Type	Cost (USD)	Impedance (Ω)	Rated Voltage (V_{pk})	Rated Bandwidth (Hz)
Tactile Labs Haptuator mkII	VCA	200	4.7	6	10 - 7000
Tactile Labs Haptuator Redesign	VCA	80	6	3	50 - 500[a]
TDK PowerHap 7G[b] + TI DRV2700EVM[c]	Piezo	60[b]+150[c]	-	1.8[c], 60[b]	50 - 500[b]

Table 3.1 List of actuators selected for the implementation of the toolkit. The selection focuses on VCA and piezo actuators that can operate within the vibrotactile bandwidth. Their costs, maximum rated voltages, and bandwidths of operation are also listed. Links to actuator datasheets can be accessed in Appendix A.1.

[a] According to the datasheet, driving the actuator outside the range of 50-500Hz doesn't damage it. The transducer can be safely operated beyond 10Hz.
[b] TDK PowerHap 7G Piezo Actuator[1].
[c] TI DRV2700EVM[1].

piezo actuators depend on the *TI DRV2700EVM* module for operation.

3.1.2 Amplifier Selection

As seen in Section 2.3.3, while the amplifiers encountered in some existing vibrotactile implementations for music are bulky and expensive, they possess a high-fidelity output. However, smaller inexpensive amplifiers can sometimes be of lower fidelity and introduce harmonic distortion into the signal [12]. This distortion is characterized by the THD+N (see Section 2.5), which is less than 0.01% for the *Bryston 2B-LP*[2]. It is also common to encounter professional audio amplifiers with THD+N less than 0.001%. However, over the years, the THD+N ratings among commercial audio amplifiers have generally seen a drastic decrease from over 25% to lower than 1% [39]. THD+N less than 1% implies that the introduced harmonic distortion is at least 40 dB below the fundamental at the output.

In this section, we select less expensive (costing <100$) and portable amplifiers (weighing <1 kg) that can be suitable for vibrotactile applications, with THD+N <1% set as an inclusion criterion. Aside from THD+N, the selected amplifiers must also be able to serve the vibrotactile bandwidth and can deliver adequate power to drive the actuators, i.e., being able to output at

[2] `http://www.bryston.com/PDF/Manuals/300004[2BLP].pdf`

least 3-6 V_{pk} of unloaded output voltage.

The high-fidelity amplifiers of the implementations surveyed in Section 2.3.3 are also noted to have more than one channel. In order to retain the flexibility to power more than one channel and be able to compare amplifiers in terms of the cost per channel, we select amplifiers having two or more channels. The selected amplifiers are listed in Table 3.2 along with their cost, weight, THD+N rating, and the maximum rated output power. Since the power rating of an amplifier can vary with the load impedance, we additionally express the rated powers of the amplifiers in terms of their maximum unloaded output voltages, using the rated impedance provided by the manufacturer.

Device	Cost (USD)	Weight (kg)	THD+N (%)	Channels	Output Power per Channel (W)	Max. Unloaded output Voltage (V_{pk})
SureElectronics (WONDOM) AA-AB33182	56	0.417[a]	1	4	64[b]	16
SureElectronics (WONDOM) AA-AB32971	37	0.28	1	2	72[b]	16.97
Tactile Labs HaptuAmp Mini	10	0.0015[c]	0.12	2	2.4[d]	4.38

Table 3.2 List of amplifier boards, selected for implementation in the toolkit. The amplifiers are selected to be less bulky and more affordable than existing reference amplifiers, with THD+N <1%. Links to datasheets of the listed amplifiers are provided in Appendix A.2.

[a] Measured with AccuWeight AW-KS001BB Kitchen Scale; Accurate to 1 g.
[b] 4Ω Load.
[c] Measured with SmartWeigh GEM20 Jewellery Scale; Accurate to 0.001 g.
[d] 8Ω Load.

The selection of amplifiers listed in Table 3.2 was subject to availability at the time of selection. The *HaptuAmp Mini* by *Tactile Labs* was an alternative suggested by the manufacturer to their four-channel alternative, which was unavailable at the time of selection. While this amplifier is much lighter than the other two by *Sure Electronics*, it offers a THD+N rating of 0.12%.

The THD+N ratings provided by the manufacturers for the amplifiers listed in Table 3.2 are typically for an input sine signal at 1 kHz but does not adequately inform us about the fidelity of the device within the bandwidth of vibration perception. In order to better understand the suitability of the selected amplifiers for vibrotactile applications, we measure the THD+N for the selected amplifiers in the vibrotactile bandwidth.

3.2 Overview of Measurements

In order to evaluate the DUTs for our implementation, we aim to characterize magnitude-frequency characteristics and harmonic distortion characteristics of each DUT across the full vibrotactile bandwidth. The magnitude-frequency characteristics help assess whether a device can operate in this bandwidth, while the harmonic distortion characteristics indicate amount of distortion in the device's output, across the same bandwidth.

In this section, we outline the measurements being performed on the selected devices (i.e., the actuators, amplifiers) to obtain the required characteristics for each device. Figure 3.1 shows the setup used to perform accurate measurements of the test devices. The measurement apparatus includes an *RME Fireface UC* audio interface to generate and feed analog voltage signals to the test devices, a *Bryston 2B-LP* reference amplifier to amplify the signals for driving the test actuators (where required), and a low-mass (0.2 g) *PCB Piezotronics 352C23* ICP accelerometer to measure the acceleration of the test actuator's output vibration. The accelerometer's output voltage signals are conditioned and amplified by a gain factor of 10, using a *PCB Piezotronics 482B11* signal conditioner. Also, a *National Instruments NI USB-4431* data acquisition interface is used to record the voltage signals outputted by the test amplifiers and the accelerometer into MATLAB for analysis. This apparatus is fed by signals designed in the MATLAB software and output to the *RME Fireface UC* audio interface.

Fig. 3.1 Apparatus used for measuring electrical and mechanical characteristics of the selected amplifiers and actuators respectively. Top: From Left to Right, *PCB Piezotronics 482B11* line conditioner for the accelerometer (seen in foreground), *National Instruments NI USB-4431* data acquisition interface, *RME Fireface UC* audio interface; Bottom: *Bryston 2B-LP* reference amplifier. Signals are generated in MATLAB, converted to analog voltage by the audio interface, and amplified by the *Bryston 2B-LP* amplifier are used to drive the actuators. The vibration of the actuators is measured by the accelerometer, conditioned by line conditioner, and acquired in MATLAB by the data acquisition card.

Using the apparatus shown in Figure 3.1, the vibration output of the selected actuators and the voltage output of the amplifiers can be measured to determine their respective magnitude-frequency and harmonic distortion characteristics. Through the remainder of this thesis, the term Device Under Test (DUT) is occasionally used to refer to any of the devices selected for testing i.e., actuators, or amplifiers. When the term DUT is encountered, the reader is requested to interpret the associated discussion as being applicable to each test device within the context of its device type i.e., measurement of a DUT refers to vibration measurements in the case of test actuators or electrical measurements in the case of the test amplifiers. In the context of distortion characteristics, the measurement of the DUT corresponds to THD in the case of the test actuators and THD+N in the case of the test amplifiers.

We now discuss the measurements in more detail. The following subsections will examine the generation of test signals and the measurement of the different DUTs.

3.2.1 Generation of Test Signals

Figure 3.2 illustrates the generation and the flow of signals used for measuring the different DUTs. The test signals used to perform the different measurements are designed in MATLAB and delivered as continuous-time voltage signals by the *RME Fireface UC* USB audio interface. The *RME Fireface UC* audio interface is set to a sampling rate of 48 kHz with a bit-depth of 24 bits and is configured to deliver a line level output, i.e., -10 dBV (approximately 0.32 V_{rms}) at 0 dBFS. When measuring the characteristics of the selected amplifiers, the signals from the audio interface are used to drive the test amplifiers directly. However when testing the selected actuators, the signals from the audio interface were first amplified by the *Bryston 2B-LP* amplifier. The *Bryston 2B-LP* is a reference amplifier chosen to ensure a high-fidelity amplified signal for driving the test actuators.

During the measurements, we test the characteristics of the amplifiers at different output voltage levels, which are also the test input levels for the selected actuators (when output by the *Bryston 2B-LP* amplifier). Since the signals being used are AC signals, we quantify the level of the signal in terms of its RMS level. The RMS level is determined by measuring the signal over 8 seconds, where the signal in MATLAB is scaled to achieve the desired test RMS level. For the measurements in this thesis, the DUTs were tested for increasing voltage levels starting at 0.2 V_{rms} up to the DUT's rated maximum. Some of the selected amplifiers were able to yield higher RMS voltages than required to operate the selected actuators. These amplifiers were tested up to 3.5 V_{rms}.

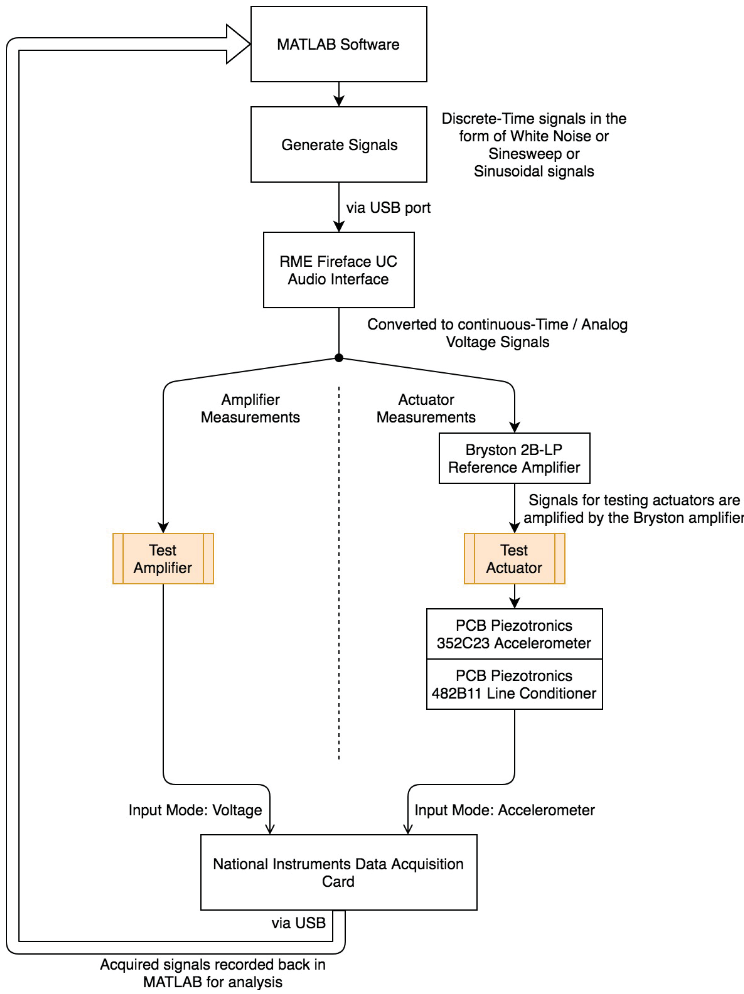

Fig. 3.2 Generation and flow of signals in the measurement apparatus. The signal flow varies slightly for amplifiers and actuators. The dotted line separates the two configurations. The signal flow to the left of the dotted line indicates the setup for testing the selectd amplifiers, and the signal flow for testing actuators is on the right.

As mentioned earlier in Section 3.2, the aim of measuring the different test devices is to characterize their magnitude-frequency and harmonic distortion characteristics. However, different types of test signals are required to measure these characteristics. Measuring the magnitude-frequency characteristics pertains to testing a wide bandwidth of frequencies (40-1000 Hz) while on the other hand, measuring harmonic distortion requires the analysis of the DUT's output for an input of a single frequency. For these reasons, we require two different types of test signals for measuring the different characteristics. Earlier works have analyzed the magnitude-frequency response using sinesweep inputs [12], step-increments of frequency and amplitude [3], as well as white noise[3] [8]. In our case, we choose to use band-limited white noise signals to characterize the magnitude-frequency response. A measurement using white noise also allows us to equalize the DUT's frequency response automatically if it is found necessary [38] (see Chapter 5[4]). On the other hand, we need to use sinusoidal test signals for characterizing harmonic distortion, so that the input to the DUT at any point always has a unique frequency. The two signal types are further discussed in the following subsections.

3.2.1.1 White Noise Test Signals

White noise signals have an infinitely wide bandwidth [45]. For the measurements done in this thesis, the white noise signal for each measurement is 1 second in duration and band-limited to 1000 Hz, with its magnitude set to a constant RMS voltage level. For a given DUT, the measurements performed with the white noise were repeated 5 times, and the process is repeated for increasing test voltage levels (see Section 3.2.1 for RMS voltage levels).

3.2.1.2 Sinusoidal Test Signals

In order to measure the harmonic distortion across the frequency bandwidth, pure sinusoidal test signals were generated at different test frequencies; first at 40 Hz, followed by step-increments of 50 Hz between 50-1000 Hz. Each sinusoidal signal was generated at a constant RMS voltage level for a 3 second duration. For each DUT, the measurements using sinusoidal signals were performed 3 times at each test frequency. After completing measurements for all the test frequencies at one voltage level, the process is repeated for increasing test voltages (see Section 3.2.1 for RMS voltage levels).

[3]Band-limited between 50-500 Hz.

[4]Measurements made with white noise and sinesweeps are also compared in Chapter 5.

3.2.2 Vibration Measurements: Actuators

We measure the vibration of each test actuator using a *PCB Piezotronics 352C23* accelerometer, similar to earlier works [3, 8, 12]. The input of the *National Instruments NI USB-4431* data acquisition card was AC coupled and set to *Accelerometer* mode. The recorded voltage signals were then interpreted in units of acceleration (m/s^2) using the accelerometer's rated sensitivity[5] as specified by the manufacturer.

The test actuators were driven (0.2 V_{rms}, up to the rated maximum) via the *Bryston 2B-LP* reference amplifier. However, the *Haptuator mk-II* was measured up to 2.5 V_{rms} input, which is less than its rated maximum since its moving magnet assembly would get pushed out of its suspension mechanism when driven at higher driving voltages, rendering the actuator inactive.

To ensure accurate measurement of the vibrations, we measure the two test VCAs in two different configurations: 1) Freely suspended, and 2) Placed on a sandbag. The goal was to test the actuator in an unloaded condition, i.e., the actuator is not being used to induce vibrations within an external surface or body. The freely suspended configuration is shown in Figure 3.3 and the sandbag configuration in Figure 3.4. The results from these two configurations are seen to be reasonably identical (as discussed later in Section 3.4), and the sandbag configuration was used for the subsequent measurements in this thesis.

The *TDK PowerHap 7G* piezo actuator is driven directly by the audio interface, via the *TI DRV2700EVM* biasing circuit instead of the *Bryston 2B-LP* amplifier. The biasing circuit delivers 200 $V_{\text{p-p}}$ at its maximum gain setting. In our case, the biasing circuit was configured at the lowest gain setting, delivering an approximately 50 V_{pk} signal to the piezo actuator. The actuator itself has a maximum rated voltage of 60 V_{pk}. Figure 3.5 shows the *TDK PowerHap 7G* piezo actuator when measured on the sandbag.

[5]Sensitivity: 0.5 $mV/m/s^2$.

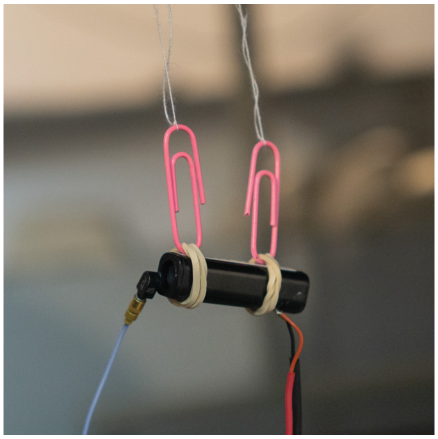

(a) *Haptuator mk-II* VCA, by *Tactile Labs.*

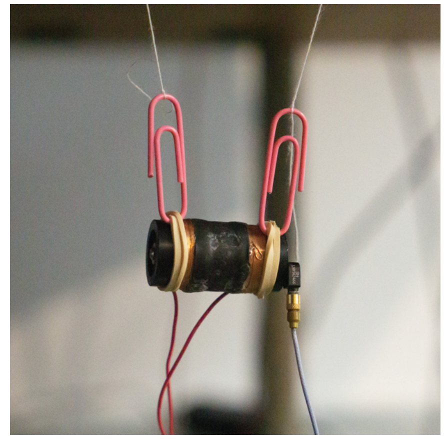

(b) *Haptuator Redesigned* VCA, by *Tactile Labs.*

Fig. 3.3 Vibration measurements of the selected VCAs, suspended freely by their own weight. The apparatus consists of two paper clips that are attached to either end of the test actuator using rubber bands. The paper clips are also connected to each other by a piece of string. The string hangs the actuator from a hook protruding from the edge of a table (not seen in the picture). A lightweight accelerometer is seen glued to one end of the actuator, whose sensing axis aligns with the actuator's axis of vibration.

3.2.3 Electrical Measurements

In order to characterize the selected amplifiers, we record output voltage measurements. For these measurements, the input of the *National Instruments NI USB-4431* data acquisition card was set to *Voltage* mode.

Aside from the amplifiers listed in Section 3.1, the *TI DRV2700EVM* is an electronic circuit that biases the Piezo actuator for operation. The electrical characteristics of this device could not be measured because its output voltage is higher than the rated input voltage of the *National Instruments NI USB-4431* data acqusition card (details in Section 3.2.3.2).

3.2.3.1 Amplifiers

All amplifier measurements were performed by loading each amplifier's output with a 6 Ω test resistor that is capable of dissipating up to 5 W of power. The output voltage across the 6 Ω load resistor was recorded by the *National Instruments NI USB-4431* data acquisition card. Each

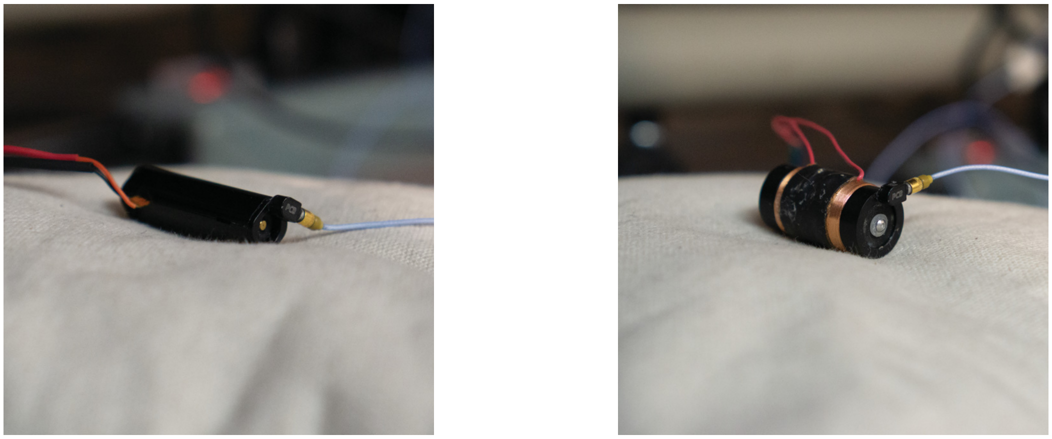

(a) *Haptuator mk-II* VCA, by *Tactile Labs*. (b) *Haptuator Redesigned* VCA, by *Tactile Labs*.

Fig. 3.4 Vibration measurements of the selected VCAs, placed on a sandbag. The actuators are loosely placed on top of a sandbag. A lightweight accelerometer is seen glued to one end of the actuator, whose sensing axis aligns with the actuator's axis of vibration.

amplifier including the *Bryston 2B-LP* was set to specific output levels between 0.2-3.5 V_{rms} except for the *HaptuAmp Mini* by *Tactile Labs*, which was only tested up to its maximum rated power output of 3 V_{rms}.

3.2.3.2 *TI DRV2700EVM* Biasing Circuit

This biasing circuit is configured to drive a piezo actuator with an analog signal at 50 V_{pk}. The measurements in this study rely on the *National Instruments NI USB-4431* data acquisition card, which is not rated to accept voltages greater than 10 V_{pk} at its negative terminal, with respect to the chassis ground [6]. As a result, the characteristics of the piezo driving circuit could not be verified through measurement. However, according to the driving circuit's datasheet (see Appendix A.3), the biasing circuit has a THD+N rating of 1%.

[6] https://booking.cirmmt.org/media/model/466/372485e.pdf

Fig. 3.5 Vibration measurements of the *TDK PowerHap 7G* piezo actuator, placed on a sandbag. The actuator is loosely placed on top of a sandbag. A lightweight accelerometer is seen glued to one end of the actuator, whose sensing axis aligns with the actuator's axis of vibration. This actuator was only tested in the sandbag configuration since it was found to yield identical results to the freely suspended configuration (see Section 3.4.1).

3.3 Analysis of Measurements

In this study, the measurement data obtained from the apparatus for the different DUTs is in the form of analog voltage signals that are digitally sampled and stored on the computer. The spectra of these measurements are analyzed using the Welch Power Spectral Density (PSD) estimate to obtain the magnitude-frequency and harmonic distortion characteristics of the various DUTs.

3.3.1 Welch PSD Estimates

As described by Proakis and Manolakis [45], an Fast Fourier Transform (FFT) helps us analyze the spectrum of a given time-series data by projecting it into a discrete number of frequency bins. Squaring the magnitude of the FFT of a time series signal results in a Power Spectrum, which indicates the distribution of power in the spectral domain but is typically affected by spectral leakage of energy in neighboring frequency bins. This spectral leakage is influenced by the window function chosen to perform the FFT, and can be typically characterized by the Equivalent Noise Bandwidth (ENBW) of the window. The PSD is then defined as a Power Spectrum that is normalized by the ENBW of the window being used to calculate the FFT. The Welch [46] PSD estimate produces a PSD estimate of the signal by splitting it into smaller overlapping chunks, and averaging the PSD obtained from these smaller chunks while using a Hamming window. This approach minimizes the effect of noise and produces smoother frequency characteristics than the traditional PSD estimate [47]. For a given FFT size, smoother PSD estimates can be produced by either using longer measurements or averaging the estimates obtained from multiple measurements [47].

The *pwelch* function in MATLAB produces a Welch PSD estimate using the user-defined FFT size and overlap factor to divide the time-series data into equal length overlapping chunks [7]. Unless explicitly specified, all of the Welch PSD estimates discussed in this thesis for the different DUTs are computed using 8192 FFT points and a 50% overlap factor. The Welch PSD estimate can be converted from units of Power to units of RMS magnitude by taking the square-root[8].

3.3.2 Characterization of Magnitude and Frequency Responses

As mentioned in Section 3.2.1.1, each DUT was measured 5 times with white noise signals of 1 second length at each increasing RMS voltage levels, to measure the magnitude-frequency characterstics. For each DUT, the Welch PSD of the 5 obtained measurements are averaged and the resulting power estimate is converted to units of RMS magnitude. This magnitude spectrum is then expressed in dB RMS acceleration (reference: $10^{-6}\ m/s^2$) for actuators and in dBV (reference: 1 V_{rms}) for the amplifiers.

3.3.3 Characterization of Distortion

As described in Section 3.2.1.2, each DUT is measured 3 times with sinusoidal signals. The harmonic distortion for each given test frequency and magnitude is determined by averaging the Welch PSDs of the 3 measurements.

As mentioned in Section 2.5, the harmonic distortion is calculated in the form of THD for

[7] `https://www.mathworks.com/help/signal/ref/pwelch.html`

[8] For a wide-sense stationary signal (e.g., white noise), the square-root of the power spectrum gives the RMS magnitude spectrum [45].

actuators, and THD+N for the amplifiers. The following subsections explore these calculations separately.

3.3.3.1 THD in Actuators

In this study, the THD is calculated as the ratio of the RMS energy of the first 5 harmonics above the fundamental to the energy of the fundamental, expressed as the percentage of THD for that particular fundamental frequency. The *thd* function in MATLAB performs this calculation from a PSD estimate[9] by ignoring the DC component and additive noise in the PSD.

The *thd* function in MATLAB considers the highest peak in the PSD of the measurement as the frequency of the fundamental, and then calculates the upper harmonics in relation to this frequency. However as shown in Figure 3.10, actuators can sometimes produce harmonic peaks that are higher in energy than the frequency of the input sine tone. This ambiguity causes errors in MATLAB's calculation of THD. To ensure accurate calculation of THD, we modify MATLAB's *thd* function such that the frequency of the fundamental is chosen as given at the input, instead of being automatically inferred from the highest peak of the PSD of the measurement.

The calculation of THD is repeated for each test frequency, for the different tested voltage levels of driving signal to obtain the THD characteristics across magnitude and frequency.

3.3.3.2 THD+N in the Amplifiers

In the case of the amplifiers, the THD+N is calculated as the ratio of the RMS energy of the first 5 harmonics and the added noise to the energy of the inputted fundamental tone. The *sinad* function in MATLAB is used to calculate THD+N from a Welch PSD estimate. The *sinad* function calculates the ratio of the power of the input fundamental to the power of harmonics and noise; the reciprocal of this *sinad* ratio gives the THD+N [48–50]. The calculated THD+N is then expressed as a percentage of distortion for that particular input fundamental frequency. This methodology is also validated by measuring the THD+N of the *Bryston 2B-LP* reference amplifier[2], shown in

[9] `https://www.mathworks.com/help/signal/ref/thd.html`

Figure 3.6.

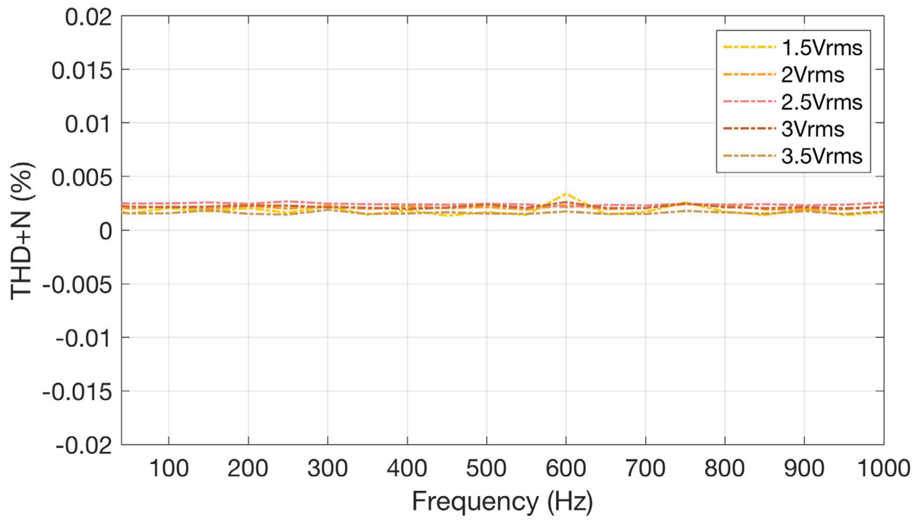

Fig. 3.6 THD+N measurement of the *Bryston 2B-LP* reference amplifier. The measured THD+N of the amplifier is in accordance with the THD+N rating specified by the manufacturer (THD+N<0.01%).

3.4 Results

The above measurement methodology is used to measure and characterize each of the selected devices for the implementation in the toolkit. The results from these measurements are presented in this section.

3.4.1 Actuators

As mentioned in Section 3.2.2, the vibrations of the two selected VCAs were measured in two different configurations. The THD and the magnitude-frequency characteristics of the devices are compared to help determine whether the measurement of the actuator's vibration is repeatable and accurate. Figure 3.7 shows the comparison of vibration measurements of the *Haptuator Redesigned* VCA, and Figure 3.8 shows the comparison of the same for *Haptuator mk-II*. The THD of both actuators appears to be attenuated when the actuator is placed on the sandbag, indicating an attenuation in the energy of the output partials. Both cases show attenuation over different ranges of frequencies when placed on the sandbag, particularly when driven at a low input of 0.2 V_{rms}. Other than these attenuation effects, the measurements from the two configurations show similar characteristics with regards to both magnitude-frequency and THD responses. Our primary objective is to ensure that we use a measurement apparatus that does not influence the results obtained from the test actuator. Since the results from the two configurations are largely similar, we choose the sandbag in subsequent measurements due to its ease and simplicity of configuration, while acknowledging that there will be attentuation in the results obtained.

The magnitude-frequency responses of the actuators were measured from 0.2 V_{rms} up to their maximum rated input voltages, and expressed in dB RMS acceleration, using 10^{-6} m/s^2 as reference. Recently reported perceptual threshold values by Fontana et al. [30] are also indicated with the magnitude-frequency characteristics of the actuators. Under the influence of pressing forces (8±1.5 N), the lowest reported threshold is approximately 72.5 dB RMS acceleration for noise stimuli [30].

The THD plots of the *Haptuator mk-II*, and *Haptuator Redesigned* VCAs are shown in Figures 3.8b and 3.7b respectively. Similarly, the magnitude-frequency characteristics are shown in Figures 3.8d and 3.7d respectively. Figure 3.9 shows the THD and magnitude-frequency characteristics of the *TDK PowerHap 7G* piezo actuator. Among the VCAs tested on the sandbag, THD for both actuators is lowest in the bandwidth 150-1000 Hz and found to be <4.5% for the *Haptuator Redesigned*, and <4% for *Haptuator mk-II*. The THD tapers down to <1% after 300 Hz for *Haptuator Redesigned*, and after 200 Hz for the *Haptuator mk-II*. Both actuators exhibit uniform magnitude-frequency characteristics within the bandwidth 100-1000 Hz. In comparison, the *TDK PowerHap 7G* yields significantly higher THD, and its performance also suffers in the low-frequency band[10].

[10]Poor low-frequency performance was previously observed for piezo actuators by Marshall [3].

The piezo actuator also has THD above 100% for the low-frequency bands because the harmonic peaks generated by the actuator are stronger than the produced fundamental. This effect is seen in the frequency response of the actuator when being driven by a 150 Hz pure tone, as shown in Figure 3.10. A 750 Hz peak is observed having a stronger magnitude than the 150 Hz sinusoid that drives the actuator.

3.4.2 Amplifiers

The THD+N characteristics of selected amplifiers are indicated in Figures 3.11a, 3.12a and 3.13a. Similarly, the magnitude responses of the selected amplifiers are shown in Figures 3.11b, 3.12b and 3.13b respectively. While the *Sure Electronics* amplifiers were measured to a 3.5 V_{rms} output voltage, the *HaptuAmp Mini* by *Tactile Labs* was only measured up to 3 V_{rms}, which lies just below its rated voltage output.

All amplifiers are observed to yield THD+N$<$0.5%, within the vibrotactile bandwidth. Furthermore, all the tested amplifiers yield a relatively flat frequency response up to their maximum tested output RMS voltages, when measured with white noise.

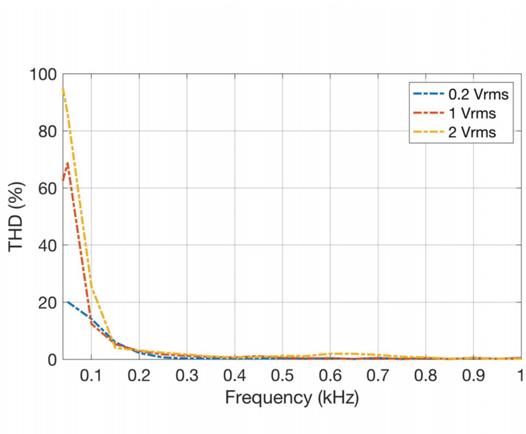

(a) THD of actuator in freely suspended configuration.

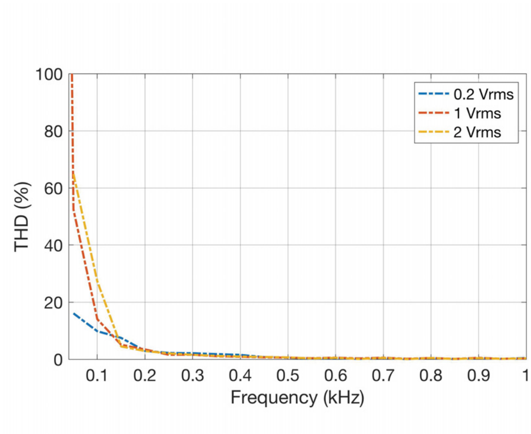

(b) THD of actuator placed on a sandbag.

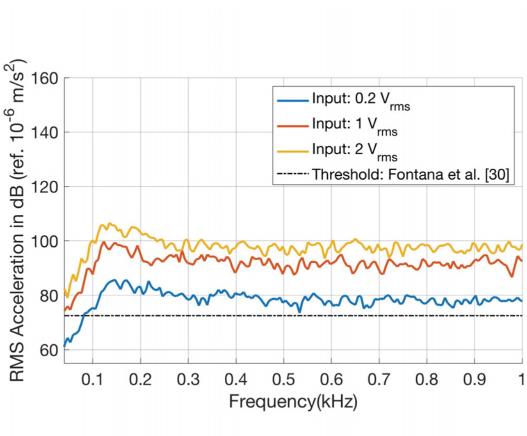

(c) Magnitude-Frequency response of actuator in freely suspended configuration.

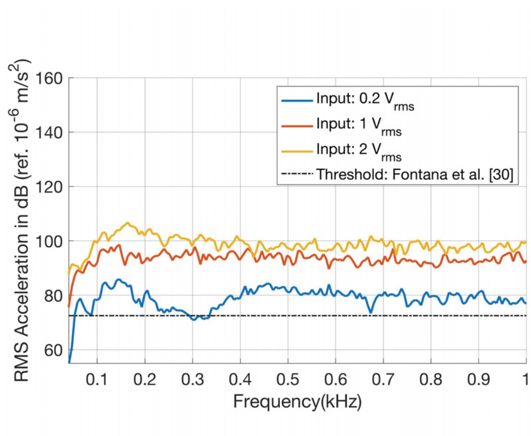

(d) Magnitude-Frequency response of actuator placed on a sandbag.

Fig. 3.7 Comparing THD and Magnitude-Frequency responses of the *Haptuator Redesigned* when measured in two different configurations i.e., in freely suspended configuration, or while placed on a sandbag. THD curves appear attenuated at low frequencies when the actuator is placed on the sandbag. Magnitude-Frequency responses do not indicate any pronounced resonances in both cases, and results appear fairly consistent across the configurations.

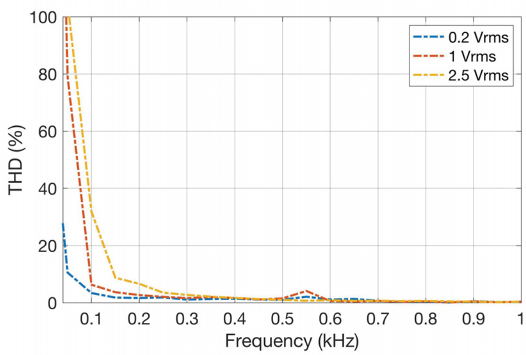

(a) THD of the actuator in a freely suspended configuration.

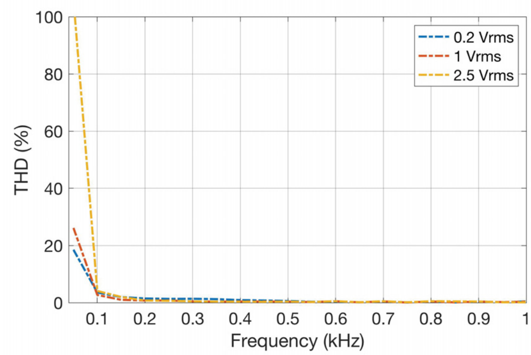

(b) THD of the actuator placed on a sandbag.

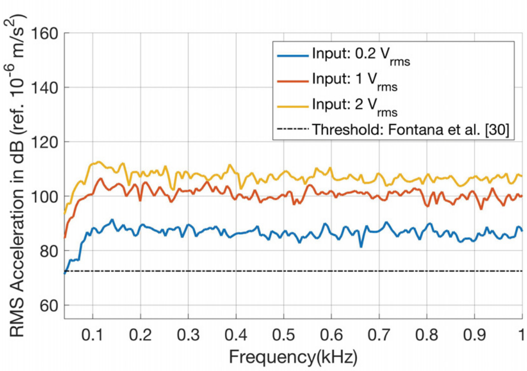

(c) Magnitude-Frequency response of the actuator in freely suspended configuration.

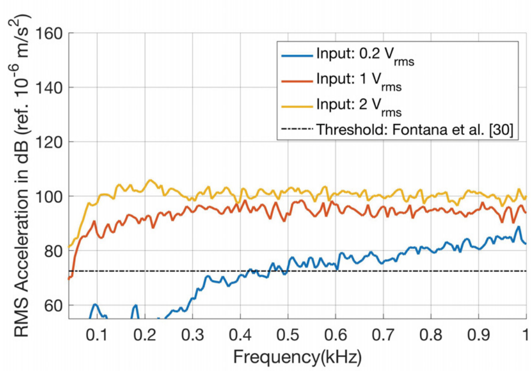

(d) Magnitude-Frequency response of the actuator placed on a sandbag.

Fig. 3.8 Comparing THD and Magnitude-Frequency responses of the *Haptuator mk-II* when measured in two different configurations i.e., in freely suspended configuration, or placed on a sandbag. THD curves appear attenuated at low frequencies, like in the case of the *Haptuator Redesigned* (Figure 3.7). Magnitude-Frequency response of the actuator for a 0.2 V_{rms} input, shows an attenuation at low frequencies. The response does not indicate any resonant characteristics when driven at higher input voltages.

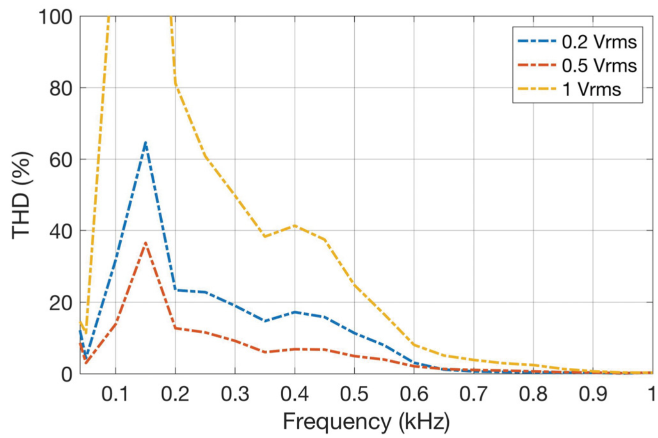

(a) THD of the actuator.

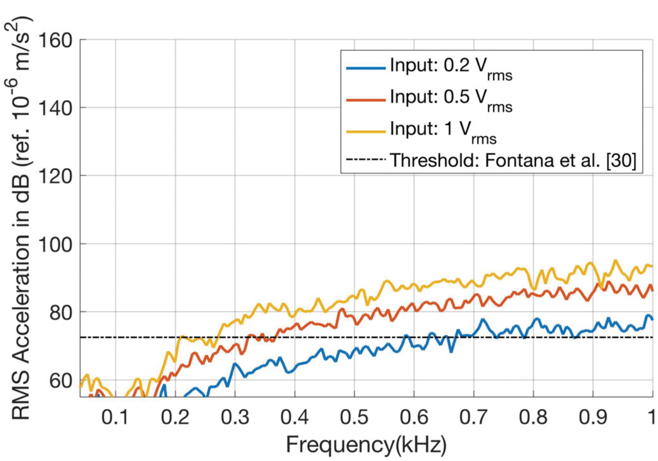

(b) Magnitude-Frequency response of the actuator.

Fig. 3.9 Comparing THD and Magnitude-Frequency response of the *TDK PowerHap 7G* piezoelectric actuator, measured while placed on a sandbag. The THD of the actuator is significantly higher than the other tested actuators, exceeding 100 % at some frequencies, when driven with a 1 V_{rms} input. Magnitude-Frequency responses show a poor low-frequency performance, compared to the tested VCA counterparts.

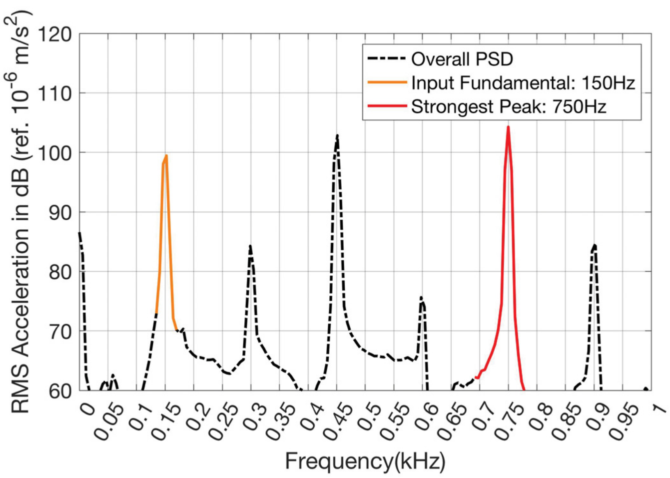

Fig. 3.10 Strong harmonic peaks in the response of the Piezo actuator, observed in the calculation of THD. The response is a mean Welch PSD of 5 measurements. The actuator is driven with 1 V_{rms} sinusoidal signal at 150 Hz. This figure shows a harmonic peak at 750 Hz that is stronger than the input 150 Hz sine tone.

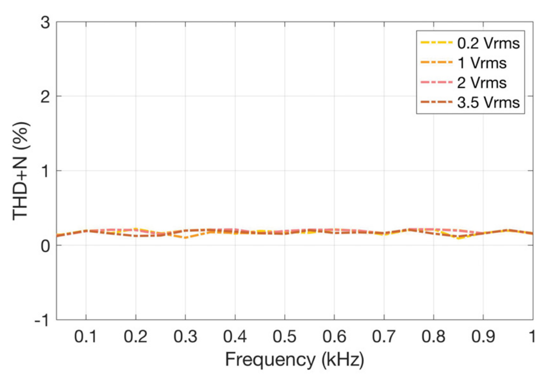

(a) THD+N characteristics.

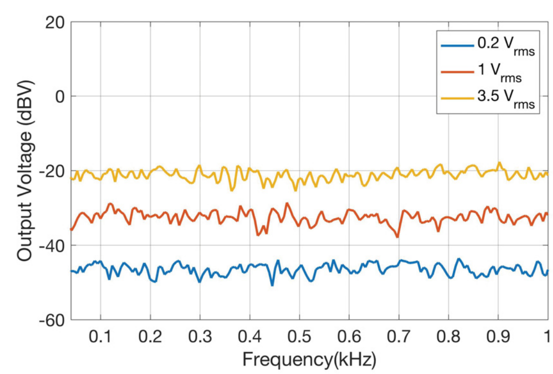

(b) Magnitude-Frequency response.

Fig. 3.11 THD+N and Magnitude-Frequency responses of the 4-channel Sure amplifier AA-AB33182. The THD+N of the amplifier is under 0.5% for output voltages up to 3.5 V_{rms}. The response of the amplifier is linear and flat across the vibrotactile bandwidth, within its operating range.

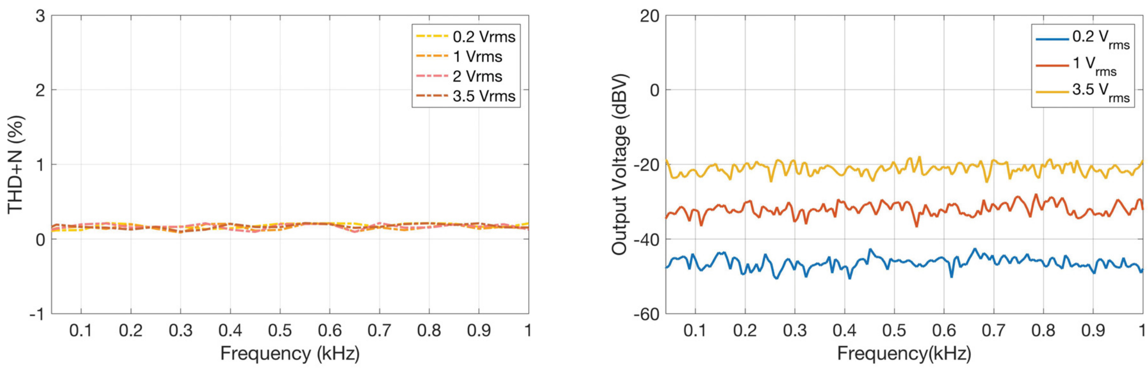

(a) THD+N characteristics. (b) Magnitude-Frequency response.

Fig. 3.12 THD+N and Magnitude-Frequency responses of the 2-channel Sure amplifier AA-AB32971. The THD+N of the amplifier is under 0.5% for output voltages up to 3.5 V_{rms}. The response of the amplifier is linear and flat across the vibrotactile bandwidth, within its operating range.

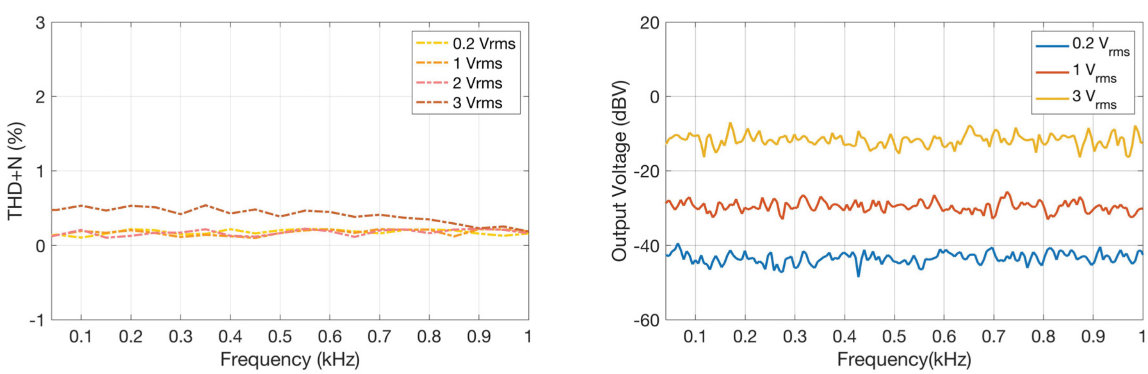

(a) THD+N characteristics. (b) Magnitude-Frequency response.

Fig. 3.13 THD+N and Magnitude-Frequency responses of the 2-channel *HaptuAmp Mini* amplifier, by *Tactile Labs*. The THD+N of the amplifier is under 0.5% for output voltages up to 3 V_{rms}. The response of the amplifier is linear and flat across the vibrotactile bandwidth, within its operating range.

3.5 Discussion

In this chapter, we select and characterize amplifiers and actuators for implementing a vibrotactile toolkit. The toolkit needs to have an output whose frequency and amplitude controls are decoupled, and with known output fidelity and magnitude-frequency characteristics across the vibrotactile bandwidth (40-1000 Hz). The toolkit also requires a portable amplifier alternative that does not compromise the fidelity of the output.

In order to accurately measure the actuators, we test them in two configurations: either suspended freely by their own weights or while placed loosely on a sandbag. When the actuators are placed loosely on a sandbag, the THD curves appear attenuated at low frequencies, compared to when they are freely suspended. The damped THD curves suggest that the harmonic peaks produced by the actuator are more damped by the sandbag than the output fundamental. Except for this observation, the two configurations yield similar results. Among the tested actuators, both VCAs show a uniform frequency characteristic within the vibrotactile bandwidth. These actuators possess the lowest THD between 150-1000 Hz, while THD in the low-frequency ranges below 150 Hz is significantly higher. However, as discussed in Section 2.5, THD at lower frequencies can turn out to be falsely high if the actuator's output magnitude is low [42]. The *TDK PowerHap 7G* actuator shows a significantly higher THD and a more reduced low-frequency response compared to the two tested VCAs. For this reason, the piezo actuator is excluded from the toolkit implementation.

Among the chosen amplifiers, the magnitude response and THD+N characteristics of all the amplifiers are well below the 1% benchmark within the vibrotactile bandwidth. However, in terms of portability, the *HaptuAmp Mini* offers the lowest weight and adequately powers both the tested VCAs. The tested amplifiers also show uniform magnitude-frequency characteristics across the vibrotactile bandwidth. While all the selected amplifiers are suitable choices in terms of portability and fidelity, the amplifiers by *Sure Electronics* supply up to 100 W, which is comparable to the *Bryston 2B-LP* reference amplifier in terms of power capabilities. At the time of amplifier selection, fidelity was a primary focus, and as a result, we did not place an upper limit on the amplifier's output range in the hope of maximizing chances of finding a suitable alternative to the *Bryston 2B-LP*. However, with the inclusion of the *HaptuAmp Mini*, we can achieve a comparable performance for a lower cost and better portability than the other two amplifier choices. Henceforth, the primary amplifier choice in our toolkit implementation is the *HaptuAmp Mini* by *Tactile Labs*.

A low-cost (<10$) audio interface was also tested for possible inclusion in the toolkit, but was not undertaken as a thorough investigation. The results of this characterization are presented in Appendix B, where the magnitude-frequency characteristics show a high-pass characteristic, attentuating frequencies below 100 Hz. This result rendered the audio interface unsuitable to our application due to its inability to serve the full vibrotactile bandwidth. However it brings a note of caution when selecting inexpensive audio interfaces as they can sometimes lack adequate low-frequency (<100 Hz) performance to be suitable in musical applications.

In the next chapter, we use these results to implement the vibrotactile toolkit, and the toolkit is then measured again without the *Bryston 2B-LP*reference amplifier.

Chapter 4

Effects of Loading and Equalization

Musical interactions require high physical and cognitive expertise [16], and tools developed for musical applications require high accuracy [8, 9]. In practice, vibrotactile actuators are typically used in a loaded condition i.e., fixed to an external surface like a screen, or the chassis of a musical instrument [1, 4, 21]. In these practical cases, the resulting frequency response is typically observed with resonant features in the frequency characteristics of the output [3].

In Section 2.4, we discuss previous approaches to equalizing these resonant features, for improving the accuracy of the vibrotactile interface. Although equalization is a common step in vibrotactile implementations developed for musical research, it is typically implemented with heuristically configured parametric equalizers. We previously investigated the use of autoregressive parameter estimation to equalize vibrotactile systems [38], where the resonant frequency characteristics of a loaded vibrotactile actuator can be automatically equalized using a measurement of the output acceleration for a white noise input. We include this as a way of achieving an equalized output in implementations made using our toolkit[1].

In this chapter, we first examine how loading an actuator can alter the final frequency characteristics, and review the method for automaticaly equalizing the frequency characteristics of a loaded actuator.

[1] A demo implementation is discussed in Section 5.4.

4.1 Effects of Loading the Actuators

Here we examine the characteristics of the actuator under loaded conditions. The *Haptuator mk-II* actuator is mounted to a stone base, as shown in Figure 4.1. The accelerometer is glued to the shell of the actuator, and the measurements for THD and magnitude-frequency characteristics are repeated. When loaded, driving the actuator with an input of 2.5 V_{rms} or higher repeated the previously discussed issue, where the moving magnet of the VCA assembly would be pushed out of the actuator's housing and render it motionless. Hence the actuator was only driven up to 2 V_{rms} for measurements under this condition.

Figure 4.2b shows the magnitude-frequency characteristics of the VCA when loaded. The magnitude-frequency response has a resonant characteristic, which is to be expected, based on observations by Marshall [3] and Giordano [12] when the actuator is attached to some external surface. Figure 4.2a shows the THD characteristics of the VCA when loaded. When compared to the response of the unloaded actuator shown in Figure 3.8b, the THD shows a significant increase when loaded, only diminishing to under 1% in the bandwidth 700-1000 Hz for lower driving voltages (0.2-0.5 V_{rms}).

Fig. 4.1 *Haptuator mk-II* VCA by *Tactile Labs*, fixed to a stone base to simulate loading. The actuator is glued to the stone base as shown in the figure. A lightweight accelerometer is shown glued to one end of the actuator, whose sensing axis aligns with the actuator's axis of vibration. This configuration is used to investigate how the actuator's output is affected when it is used to induce vibrations in an external object or surface.

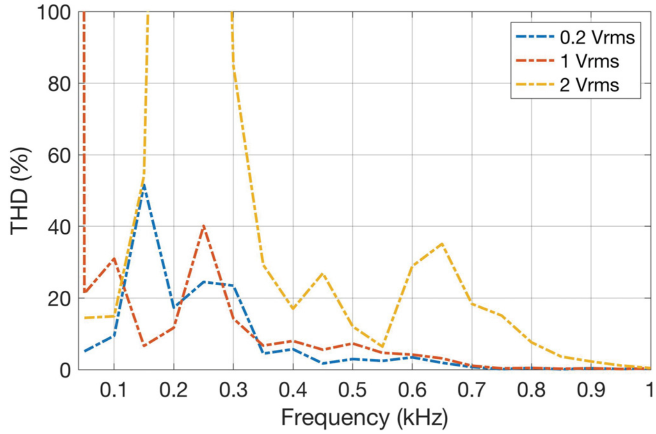

(a) THD of the actuator.

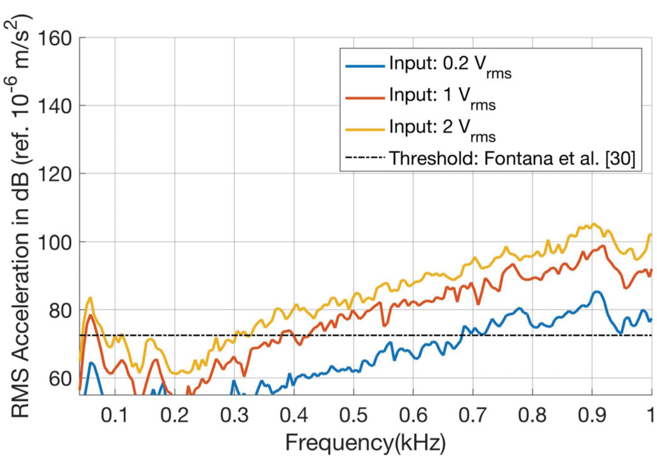

(b) Magnitude-Frequency response of the actuator.

Fig. 4.2 Comparing THD and Magnitude-Frequency response of the *Haptuator mk-II* actuator, measured while mounted to a stone base. The THD of the actuator is significantly higher than the unloaded condition, exceeding 100% at some frequencies. The Magnitude-Frequency response shows attentuation at low frequencies along with expected resonant characteristics.

4.2 The Yule-Walker Method for Automatic Equalization

In this section, we review the Yule-Walker estimation method proposed by us in [38], as an automatic equalization method when using the implemented toolkit. There exist system identification methods that allow identification and inverse filtering of a system's frequency response [51]. Spectral estimation methods allow estimation of a system's frequency response characteristics, as models that can be implemented as all-pole filters. Many methods exist to help identify the frequency response of a system's output, each offering pros and cons in terms of computationally efficiency, sensitivity to noise, or accuracy in identifying the system's parameters [47]. One early example is the compression for speech signals, which were analyzed using a source-filter model. In this process, the Linear Predictive Coding (LPC) method separates the noisy part of a speech signal from the resonant part, by representing the resonant part as an all-pole filter, which could be used as an inverse filter [52]. These identification methods work iteratively, while working to minimize a cost function like in the least-squares approach, or maximizes a likelihood estimator.

Here we explore a parameter estimation method called the *Yule-Walker* method, which is similar to LPC and commonly encountered in system identification applications [47, 53, 54]. The method estimates minimum-phase filters by design [47], which guarantee stability upon inversion [55][2] and therefore particularly useful for the task of equalization. Among various spectral estimation methods, autocorrelation methods attempt to compare the output spectrum of a system, to the spectrum of an input zero-mean white noise, which is generally regarded to be flat and minimum-phase. The Yule-Walker method uses equations that draw a relation between a set of all-pole model parameters $\{a_0, a_1, \cdots, a_p, \sigma^2\}$ and the autocorrelation sequence of the output sequence y_n [47]:

$$R_{yy}(k) = \begin{cases} -\sum_{l=1}^{p} a_l R_{yy}(k-l), & \forall \quad k > 0, \\ -\sum_{l=1}^{p} a_l R_{yy}(-l) + \sigma^2, & \text{for} \quad k = 0. \end{cases} \tag{4.1}$$

[2] `https://ccrma.stanford.edu/~jos/fp/Definition_Minimum_Phase_Filters.html`

R_{yy} in the above equation (4.1), has a Hermitian and Toeplitz structure[3], and the parameters for the model $\{a_0, a_1, \cdots, a_p, \sigma^2\}$ are evaluated by choosing p equations for $k > 0$, in the above relation, where σ^2 is estimated for $k = 0$. These equations are solved in a more computationally efficient method using the *Levinson-Durbin* recursion. The parameters $\{a_0, a_1, \cdots, a_p\}$ are now coefficients of z-transform of the model $\hat{y_n}$, which estimates the output sequence y_n, given by $A(z)$ where [47]:

$$A(z) = 1 + \sum_{k=1}^{p} a_k z^{-k}, \quad \text{where } a_0 = 1. \tag{4.2}$$

Then the inverse of equation 4.2, given by $A^{-1}(z)$, represents the inverse filter for the model that serves to equalize the response of the system [47]. The technique is shown to automatically determine stable inverse filters to equalize an actuator's frequency response with accelerometer measurements of the actuator while being driven with white noise. As opposed to configuring a parametric equalizer, the equalization process reduces to a choice of filter order, which determines the accuracy of equalization. Although a higher order filter improves the accuracy of modeling a given system, the spectral flatness estimate, which is given by the ratio of geometric to the arithmetic mean of the PSD, shows that there is a limit to the pole order beyond which, the amount of equalization begins to saturate. In other words, increasing the pole order beyond this limit only leads to a marginal improvement in the amount of equalization achieved.

4.3 Discussion

Loading the actuators is seen to affect both the magnitude-frequency and harmonic distortion characteristics of the vibrotactile output. When the actuators are tested under loaded conditions, the magnitude-frequency characteristics change from a uniform curve to one containing resonant features, when the actuator is loaded. This result is consistent with observations by Marshall [3], Giordano [12], and Papetti et al. [8].

[3]Elements along any diagonal of the matrix are identical.

In the characterization of the TouchBox interface by Papetti et al. [8], magnitude response attenuation is observed when test loads are placed on the vibrating surface for simulating pressing forces. The acceleration profiles measured on the vibrating surface for test weights of 800 g and 1500 g were found to be approximately 8 and 10 dB lower respectively, compared to that of a 200 g load. This amplitude attenuation is due to a dampening effect that is proportionate to the force exerted on the vibrating surface. In our case for the loaded actuator, the stone pestle (weight = 897 g) to which the test actuator is mounted also exhibits a similar dampening effect, and the order of RMS vibration magnitude observed in our loaded results is comparable with results reported by Papetti et al. [8]. The overall average magnitude response of is attenuated by approximately 20 dB in the loaded condition, when compared to that of the unloaded (sandbag) condition seen in Section 3.2.2.

The THD characeristics are also significantly altered when the actuator is loaded, indicating an increase in the non-linear operation of the actuator. This is consistent with THD measurements among loaded loudspeakers [56]. An important result in the characterization of distortion of the actuators is that the harmonic peaks produced by the actuator can sometimes be stronger than the intended fundamental frequency causing the THD measurement to be $\geq$100%, and the same effect is reproduced when the actuator is loaded. This phenomenon is concerning among vibrotactile apparatuses that are expected to deliver purely sinusoidal vibrotactile stimuli to participants of psychophysical experiments. The THD $\geq$ 100% suggests that the output of the actuator is a spectrally rich waveform, and no longer a sinusoid of a unique frequency. As seen in Section 2.2.2.4, complex (spectrally rich) waveforms have a lower threshold of perception than pure sinusoidal waveforms. Also discussed in the same section; Perceptual thresholds were reported to be consistently lower when participants pressed down on the vibrating surface with a measured force. These results suggest that loading the vibrating surface results in a complex waveform that can in-turn lower the threshold of perception. When the frequency of the test tone and the threshold of perception are both critical variables in an experiment, finger loading can result in inaccurate results within an experiment. The effect of loading makes a case for characterization of THD in the vibrotactile systems, particularly where psychophysical experiments are concerned. We contend that THD of a vibrotactile interface is an accurate indicator of such artifacts.

As seen from Sections 2.4 & 4.1, loading the actuator brings about resonant frequency characteristics, which can be equalized to improve the accuracy of the resulting interface. While depending on AC signal driven actuators, the autoregressive Yule-Walker method can automatically equalize any resonant features that arise in implementations using the toolkit, using white noise measurements. Among the surveyed methods, autoregressive methods rely on a white noise input and choosing a pole order. The pole order can be inferred from the degree of spectral-flatness achieved [38].

The final toolkit implementation along with the equalization method described above are demonstrated in Chapter 5, with the help of an example application.

Chapter 5

The Toolkit and its Application

In this chapter, we first implement our vibrotactile toolkit, measure its characteristics and present a datasheet informing users of its characteristics. Secondly, a practical use of the toolkit is presented by implementing vibrotactile feedback in the body of an unfinished DMI chassis. The output of the implementation is equalized using the *Yule-Walker* method discussed in Chapter 4.

5.1 Implementing the Toolkit

Based on the results of characterization discussed in Chapter 3 we now implement a toolkit that meets the requirements of musical applications (see Chapters 2 and 3).

Among the tested actuators, the *TDK PowerHap 7G* piezo offered an inadequate low-frequency response and a significantly higher THD than the tested VCAs. While this piezo actuator is the most portable among the tested actuators, it depends on an external biasing circuit for operation. Furthermore, each additional channel of piezo actuator would require an additional biasing circuit. When compared to the tested VCA counterparts, the piezo system as a whole does not offer any cost, portability or performance advantage.

All of the tested amplifiers can be suitable choices for implementing vibrotactile applications. The amplifiers by *Sure Electronics* are multi-channel and can also optionally be used to drive loudspeakers, making them ideal for applications where both haptics and audio are desired. For e.g.

The 4-channel amplifier by *Sure Electronics* (AA-AB33182) is particularly useful for implementing 2 channels of haptic feedback alongside driving a stereo loudspeaker system. However, among the tested amplifiers the 2-channel *HaptuAmp Mini* has the lowest cost of 10 USD, and the highest portability, weighing only 15 g (excluding the required 5 V DC power supply). Given the low weight of the *HaptuAmp Mini* amplifier board, we prepared a Printed Circuit Board (PCB) that provides a plug and play interface to the amplifier board. The interface includes a barrel connector for power supply, one 3.5 mm stereo input jack, and a pair of 3.5 mm (switched) output jacks, each with a corresponding header pin terminal. The final amplifier assembly (*HaptuAmp Mini* + PCB with peripherals) weighs approximately 16 g.

Therefore, for the implementation of the toolkit we choose the two tested VCAs and the 2-channel *HaptuAmp Mini*, as shown in Figure 5.1. The *Haptuator mk-II* actuator was not practically usable above 2.5 V_{rms} since the moving magnet assembly would be pushed out of its suspension assembly for higher driving voltages. For this reason, we regard the maximum operating voltage of the *Haptuator mk-II* to be around 2 V_{rms}, which can be delivered by the *HaptuAmp Mini*. The *HaptuAmp Mini* by *Tactile Labs* provides 2 channels of output, with a flat frequency response and a relatively good fidelity (THD+N <1%).

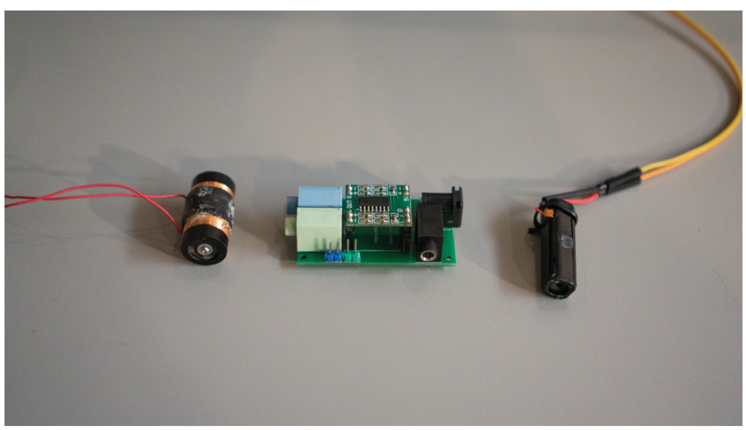

Fig. 5.1 The Implemented Toolkit, consisiting of the two tested VCAs and the *HaptuAmp Mini* amplifier board by *Tactile Labs*. All the shown components collectively weigh approximately 40 g, and can be implemented for around 200 $ (see Section 5.3). An additional 5 V power supply (not shown here) is still required to operate the toolkit.

5.2 Characterization of the Implemented Toolkit

We repeat THD and magnitude-frequency characteristics of the toolkit, i.e., driving the toolkit's actuators using the *HaptuAmp Mini*, to allow us to evaluate the implemented toolkit's performance. The *HaptuAmp Mini* was driven by signals generated in MATLAB and input to the amplifier via the *RME Fireface UC* audio interface, which in turn drive the VCAs placed on a sandbag. The accelerometer is attached to the shell of the VCAs to measure the characteristics of the actuator, as seen previously in 3.

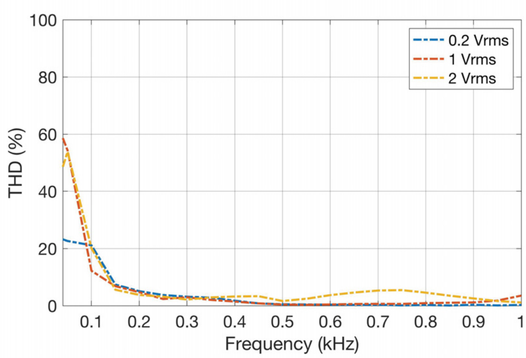

(a) THD characteristics of the *Haptuator Redesigned*. THD <5.6% between 150-1000 Hz.

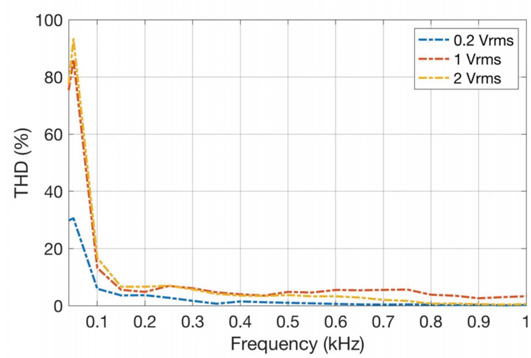

(b) THD characteristics of the *Haptuator mk-II*. THD <7% between 150-1000 Hz.

Fig. 5.2 THD of the two actuators included in the toolkit, when driven with the *HaptuAmp Mini* amplifier module. An increase in THD is observed when compared with results of the same actuators driven by the *Bryston 2B-LP* reference amplifier (see Figures 3.7b & 3.8b respsectively).

When driven with the *HaptuAmp Mini* amplifier module, the measured THD of the VCAs sees an overall increase, compared to the results obtained by driving the same actuators with the *Bryston 2B-LP* reference amplifier. When driven by the *HaptuAmp Mini*, the THD of the *Haptuator mk-II* is <7% between 150-1000 Hz, dropping to <1% from 800 Hz and above. The results for *Haptuator Redesigned* show that THD is <5.6% between 150-1000 Hz and only drops to 1.1% around 1000 Hz.

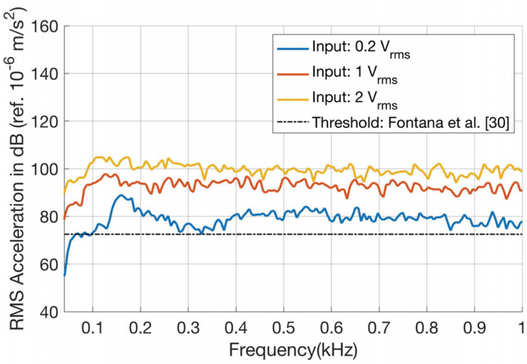

(a) Magnitude-Frequency characteristics of the *Haptuator Redesigned*.

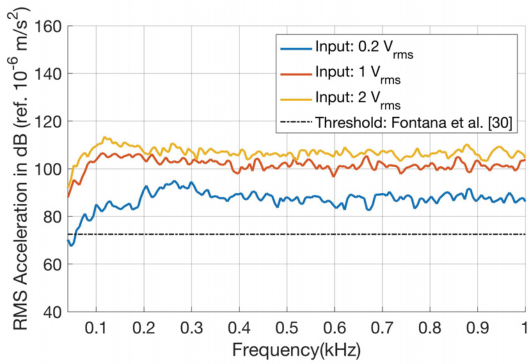

(b) Magnitude-Frequency characteristics of the *Haptuator mk-II*.

Fig. 5.3 Magnitude-Frequency of the two actuators included in the toolkit, when driven with the *HaptuAmp Mini* amplifier module. The results are identical to the results obtained when driven using the *Bryston 2B-LP* reference amplifier (Figures 3.7d & 3.8d).

5.3 The Final Toolkit

The implemented toolkit includes an inexpensive (10 $) and portable (approximately 16 g) amplifier, and two VCA alternatives: The *Haptuator Redesigned* and *Haptuator mk-II* VCAs by *Tactile Labs*. The toolkit also requires an external 5 V DC power supply. While the amplifier supports a lower supply voltage, the amount of gain delivered by the amplifier may be altered. These details along with other parameters inferred from the results of Chapter 3 are tabulated below in the form of a datasheet.

5.3.1 Datasheet

We list a datasheet for each of the two variants of the toolkit, in order to inform an end user of its characteristics.

Implemented Toolkit: *HaptuAmp Mini* with Haptuator VCAs		
Quantity	Typical Value	Unit
Amplifier Input Impedance	10	kΩ
Amplifier Output Voltage	3	V_{rms}
Output THD	$<6^{a,c}$, $<7^{b,c}$	%
Total Usable Bandwidth	40 - 1000	Hz
THD Bandwidth	150 - 1000	Hz

Table 5.1 Datasheet of the Implemented Toolkit.

[a] *Haptuator Redesigned*.
[b] *Haptuator mk-II*.
[c] See THD Bandwdith Values.

5.4 Application of the Toolkit

In this section, we demonstrate a hypothetical application of the toolkit. The toolkit is used to implement vibrotactile feedback in the chassis of a T-Stick DMI by Malloch and Wanderley [57], using the *Haptuator Redesigned* VCA. This implementation demonstrates equalizing the resonant frequency characteristics that may arise in an end-application of the toolkit. The system is equalized by using accelerometer measurements of the T-Stick's chassis when the mounted VCA is driven with a white noise input. The resonant characteristics of the frequency response are then equalized using the autoregressive Yule-Walker equalization method discussed in Chapter 4.

5.4.1 Equalization with the Implemented Toolkit

The Haptuator Redesigned VCA is mounted inside a long Poly Vinyl Chloride (PVC) tube, simulating the chassis of an unfinished T-Stick DMI [57]. The PVC tube with actuator mounted inside is shown in Figure 5.4. Also shown is the lightweight accelerometer mounted on the side of the tube, such that the axis of accelerometer's sensing is aligned with the axis of the actuator's vibration.

We use the Yule-Walker equalization method to determine the inverse filter for this configura-

(a) *Haptuator Redesigned* fixed to the inside of an unfinished T-Stick chassis.

(b) Accelerometer attached to the side of the chassis, in the axis of vibration of the actuator.

Fig. 5.4 Demonstrating the application of the implemented toolkit, by using it to produce vibrotactile feedback in an unfinished T-Stick chassis (i.e. a PVC tube). The *Haptuator Redesigned* VCA is shown stuck inside the T-Stick's chassis, with the axis of vibration parallel to the axis of the tube. The lightweight measurement accelerometer is also shown stuck to the edge of the pipe where the sensing axis of the accelerometer and the axis of the tube are in parallel.

tion and equalize the frequency characteristics.

5.4.2 Equalization Methodology and Results

In Section 3.2, we described the measurement methodology used to obtain the magnitude-frequency characteristics of each actuator. Similar to the methodology adopted in [38], the apparatus shown in Figure 5.4 was measured 5 times for a white noise input. To design the inverse filters using the Yule-Walker method, we use the measurements made for the VCA attached to the PVC tube with white noise signals at 1 V_{rms}. The Welch PSDs estimates are obtained for the 5 measurements and then averaged. The averaged PSD is expressed as a magnitude in dB RMS acceleration, and multiplied with a random phase[1], as per the methodology adopted in [38]. This process yields a real signal when an inverse Fourier transform operation is applied, without altering the PSD. The autocovariance matrix of this resulting time-series signal is used to obtain an autoregressive Yule-Walker parameter estimate. The result is an all-pole filter that models

[1] The phase of a stationary random process is also random.

the actuator's frequency characteristics whose inverse is an Finite Impulse Response (FIR) filter that equalizes the actuator's frequency characteristics. The actuator is equalized by digitally filtering the driving signal using the obtained inverse filter, where the generated driving signal is downsampled to 3 kHz, inverse filtered, and then upsampled back to 48 kHz before being output by the *RME Fireface UC* DAC. This method for equalization requires us first to select a pole order that achieves an optimum degree of equalization. In order to determine the proper pole order, we use a spectral flatness estimate, obtained as the ratio of the geometric to arithmetic means of the PSD of the measurement [38, 58]. We estimate multiple Yule-Walker inverse filters for increasing pole order, and each inverse filter is used to equalize the T-Stick apparatus. The proper pole order for equalizing the assembly is determined heuristically as the point of saturation of the spectral flatness estimate.

This equalization method relies on the system being driven with a zero-mean white noise input. However, in comparison with the sinesweep methodology adopted by Giordano [12], the resulting magnitude-frequency response when using white noise includes the variance of the white noise. Here we additionally procure sinesweep measurements on the equalized T-Stick to more accurately evaluate the quality of equalization being achieved, and compare the equalization result with Giordano's [12] results.

5.4.3 Model Order Selection

The response of the T-Stick is estimated with Yule-Walker inverse filters of increasing pole orders, starting from 4 poles to 40 poles, ascending in multiples of 4. Each inverse filter was used to filter the 1 V_{rms} noise signal driving the actuator. In order to accomplish this in each case, the white noise signal was first downsampled to 3 kHz, filtered with the inverse filter, and then upsampled back to 48 kHz[2]. The average Welch PSDs was obtained for the 5 white noise measurements of the actuator, for each inverse filter model. The averaged Welch PSD and then used to calculate the spectral flatness estimate. This process repeats for each of the ten inverse filters of each actuator,

[2]The original sample rate.

to assess the minimum pole order required to inverse the response. The effects of pole order on the spectral flatness estimate of the actuator are in Figure 5.5.

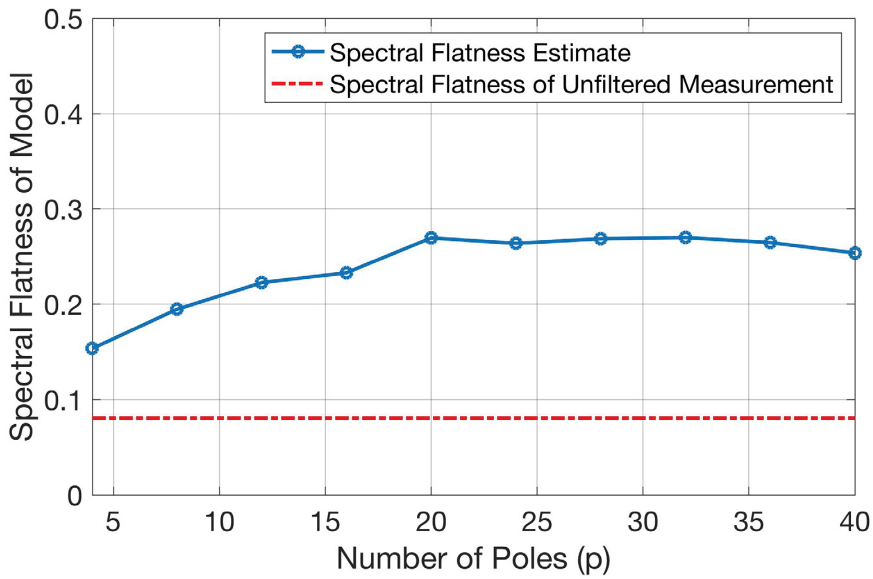

Fig. 5.5 Selecting the proper pole order to equalize the T-Stick's vibrotactile response, based on the spectral flatness estimate obtained for various inverse filters of increasing pole order between 4 and 40 poles. The spectral flatness saturates for a 20-pole inverse filter, which is selected to equalize the T-Stick's frequency response.

From the results in Figure 5.5, we infer the point where the spectral flatness estimate saturates as being 20 poles. We select this saturation point as the minimum pole order required to equalize the actuator.

5.4.4 Characteristics of the Equalized T-Stick

Having selected an inverse filter for equalizing each actuator, we now observe the effect of such equalization on the magnitude-frequency characteristics of the vibrotactile output. The measurements for magnitude-frequency characteristics were repeated as described in Section 3.2.2. The only difference in performing these measurements is the equalization process, which is implemented by downsampling the driving signal to 3 kHz, inverse filtering, and upsampling back to 48 kHz before driving the actuator with it.

Figure 5.6 depicts the magnitude-frequency characteristics of the T-Stick assembly before and after equalization.

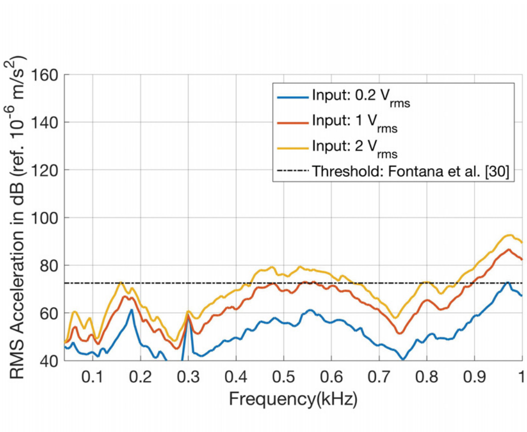

(a) Before Equalization.

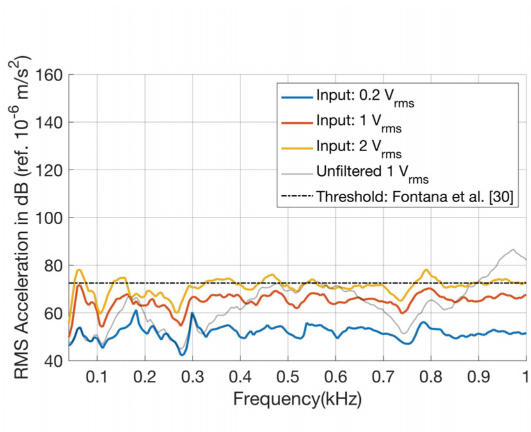

(b) After Equalization with 20-pole inverse filter.

Fig. 5.6 Magnitude-Frequency characterstics of the T-Stick before and after equalization. The indicated thresholds are as reported by Fontana et al. [30]: Lower threshold at approximately 72.5 dB RMS Acceleration, observed for noise stimuli bandlimited between 50-500 Hz at a pressing force of 8±1.5 N.

5.4.5 Sinesweep Measurements and Results

The T-Stick apparatus was driven with a 1 V_{rms} logarithmic sinesweep signal, varying from 1-1500 Hz over 10 seconds. The 20-pole inverse filter obtained from the Yule-Walker estimation in Section 5.4.2 is used to filter the sinesweep before driving the VCA to obtain its equalized sinesweep response. Figure 5.7 shows the results obtained from sinesweep measurements.

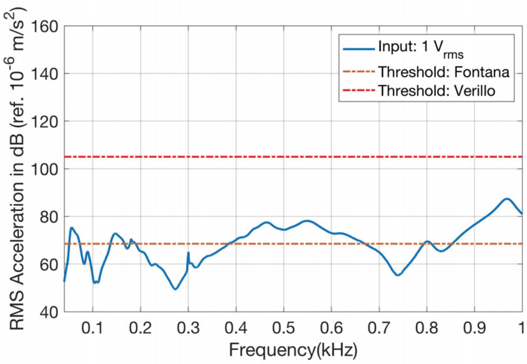

(a) Before Equalization.

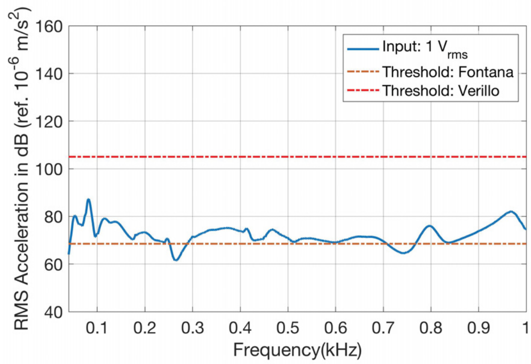

(b) After equalization with the 20-pole inverse filter.

Fig. 5.7 Magnitude-Frequency characteristics of the T-Stick before and after equalization, obtained from sinesweep measurements. The sinesweep is of a 10-second duration, varying from 1-1500 Hz. The figure shows the response between 40-1000 Hz. Obtained from Welch PSD estimates analyzed with 32768 FFT bins. Indicated thresholds are as reported by Fontana et al. [30]: Lower threshold at 68.5 dB RMS Acceleration, observed for sinusoidal stimuli of 250 Hz for a pressing force of approximately 15 N; The upper threshold line at 105 dB RMS Acceleration was reported by Verrillo [31] for sinusoidal stimuli at 250 Hz (no reported pressing force). There is a visible improvement in the flatness of the frequency response curve, from approximately $\pm$35 dB RMS acceleration before equalization, to $\pm$20 dB RMS acceleration after equalization, including equalization of the resonant features seen before equalization.

5.5 Discussion

While all of the tested amplifiers are suitable choices based on fidelity, we select the *HaptuAmp Mini* by *Tactile Labs* for our toolkit due to better portability and lower cost. Among the tested actuators, both the tested VCAs are observed to have similar characteristics and are both included in the toolkit despite the *Haptuator Redesigned* being the less expensive option among the two. The *TDK PowerHap 7G* piezo actuator is excluded from the toolkit for its relatively lower fidelity and poor low-frequency response. Together, these devices constitute our toolkit.

The final toolkit including the actuators is portable (weight: <40 g), low-cost (90-210 USD per channel) and has the lowest characterized THD between 150-1000 Hz. The toolkit also serves a wide bandwidth (40-1000 Hz) and is suitable for implementing vibrotactile feedback in most types of musical applications.

The implemented toolkit is then tested and characterized by driving the actuators with the *HaptuAmp Mini* and comparing the results to those obtained when the *Bryston 2B-LP* reference amplifier drove the actuators. A minor increase in THD of the actuators is observed, while the magnitude-frequency characteristics remain mostly similar. The THD of the final toolkit is found to be $<7\%$. However, the THD specifications of the toolkit no longer hold when the actuator is loaded. The toolkit also currently relies on being driven by a high-fidelity audio interface, without which the reported magnitude-frequency and THD characteristics are invalid. An audio interface that can at least serve the vibrotactile bandwidth is required for the datasheet provided in Table 5.1 to be valid.

A practical application using the toolkit is then demonstrated by implementing and then equalizing vibrotactile feedback in the chassis of a T-Stick DMI. In this case, the pole order required to equalize the T-Stick's response was inferred from the spectral flatness of the actuator's response, after equalization with different inverse filters of increasing pole order. A pole order of 20 was identified as the point of saturation in the spectral flatness estimate.

When using white noise, the Welch PSD curves are seen to be affected by the variance of the white noise, which can obscure the exact frequency response. However, the noise artifacts are not observed when the measurement is performed using a sinesweep, as seen in Giordano's results [12]. Here we repeat measurements of the equalized T-stick chassis using sinesweep signals for comparison. The results of sinesweep measurements after equalization indicate that the equalized result is flat within ± 20 dB RMS acceleration for a 1 V_{rms} sinesweep input, which is approximately a 43% increase in the flatness of the overall frequency characteristics. Sharp resonances that are not equalized by the inverse filter are more easily observed in the sinesweep measurement, particularly at the low and high-frequency bands below 100 Hz and above 800 Hz. However, there is still an overall visible improvement in the flatness, and reduced variance in the overall frequency response.

Chapter 6

Conclusions

In review, the goal of this study is to implement a vibrotactile toolkit that offers the benefits of a lower cost, better portability, and multiple characterized output channels. The toolkit allows users to prototype vibrotactile feedback in musical applications easily.

The perceptual and technological factors influencing the design of interfaces are surveyed. For supporting musical applications, the toolkit needs to allow decoupled controls for the output's amplitude and frequency and be able to operate across the full vibrotactile bandwidth (40-1000 Hz). The harmonic distortion of the toolkit should additionally be characterized, as an indicator of its fidelity. A selection of actuators and amplifiers are tested and evaluated for implementing such a toolkit.

To moving magnet VCAs and a piezoelectric actuator were preferred for their versatility of control (i.e., decoupled controls) and wide bandwidth (40-1000 Hz). A zero-mean white noise input is used to measure the magnitude-frequency characteristics of the actuators. Under unloaded conditions, the selected VCAs show uniform frequency characteristics across the vibrotactile bandwidth and linear amplitude characteristics. The THD of the actuators is characterized across the vibrotactile bandwidth to determine the usable range of each actuator. The VCAs were found to possess the lowest order distortion in the bandwidth 100-1000 Hz. The piezo actuator is excluded from the toolkit implementation due to its inadequate low-frequency response and significantly higher THD characteristics, with no particular cost-benefit.

The loading of the actuator was seen to influence the characteristics significantly, introducing resonant characteristics into the vibration response and causing a significant increase in THD characteristics. Additionally, this loading effect also observes cases where THD is occasionally $\geq$100%. In these cases, the harmonic peaks produced by the actuator are much higher than the fundamental, which can be a cause for concern in applications requiring a highly accurate vibrotactile output. In psychophysical experiments where the frequency of vibration and the threshold of perception are critical experimental variables, the THD response of the interface plays a vital role in obtaining accurate results. The THD response can inform an experimenter of the usable bandwidth of the system where the frequency artifacts introduced by the system are at least consistent across the bandwidth, if not minimal. The Yule-Walker method is proposed with the toolkit to simplify the equalization process for an end-user. The spectral flatness estimate helps assess the quality of equalization in the end application.

While all the selected amplifiers are afforded within a few 100 USD and serve as cheaper and more portable alternatives to the existing "one-of-a-kind" implementations, there are a few factors to consider in terms of their performance. Among the selected amplifiers, the two amplifiers by *Sure Electronics*were selected to be in the league of the *Bryston 2B-LP* reference amplifier in terms the capacity of 100 W per channel, in the hope of finding better fidelity at lower powers of operation. Subsequently, the *HaptuAmp Mini* by *Tactile Labs* offered a performance that paralleled the two *Sure Electronics*amplifiers in terms of fidelity, but for a much lower cost and higher portability. This performance was unexpected based on results from earlier tests of low-cost portable amplifiers in this output power range [12], and most of the amplifiers surveyed at the time of selection within this weight and price range reported THD+N <10%, beyond the <1% benchmark. However, the *HaptuAmp Mini* remains in the selection due to its low cost and better portability. The *HaptuAmp Mini* is a 2-channel amplifier costing 10 USD, with adequate power to serve the selected VCAs. From the results of characterization, the THD+N was found to be <0.5% across the entire vibrotactile bandwidth for increasing levels of input, indicating that all the noise and distortion components introduced by the amplifier lie at least 46 dB below the amplified input signal. Lastly, all the tested amplifiers had a flat output frequency response (i.e., no resonant characteristics) in the vibrotactile bandwidth.

A low-cost audio interface was also characterized to explore the possibility of making the toolkit more compatible with different platforms. However, the audio interface was found to have an inadequate low-frequency response, unable to reproduce an output over the lower part of vibrotactile bandwidth. This limitation is essential to consider when attempting to implement low-cost vibrotactile interfaces. In light of this result, the characteristics of the implemented toolkit currently rely on a high-fidelity audio interface, particularly for use in psychophysical experiments.

The final toolkit is used to implement vibrotactile feedback in a PVC tube that emulates the chassis of the T-Stick DMI. A 20-pole inverse filter, produced from the Yule-Walker estimation is used to automatically equalize the resonances observed in the frequency response of the T-Stick. The analysis of spectral flatness for varying pole orders indicates that the 20-pole filter offers an optimal quality of equalization.

The equalization of the T-Stick is also compared with measurements performed with sinesweep input signals. When compared to the measurements made using the white noise input, sharp resonances that were not equalized by the inverse filter are observed more easily in the sinesweep measurement, particularly at the low frequency and high-frequency bands below 100 Hz and above 800 Hz. These sharp resonances may also be present due to non-linear operation, which is not typically visible in white-noise measurements and remains uncompensated when using the Yule-Walker equalization method [45]. However, an overall visible improvement in the flatness is still evident in the form of reduced variance in the overall frequency response.

6.1 Future Work

While the current toolkit readily implements a wide variety of vibrotactile feedback for musical applications, a few improvements are well deserved.

Pressing forces currently remain uncharacterized in the toolkit, offering little insight into the change in characteristics when a finger interacts with the vibrating surface. The use of load cells like in the example of the Touch-Box by Papetti et al. [8] would be a significant addition to the toolkit.

The present equalization method offers a quick approach to equalizing response of the implementation but currently relies on a reliable accelerometer measurement. Including a reliable and inexpensive accelerometer with the toolkit like in the case of Stereohaptics by Israr et al. [14], would help equalize the toolkit's end-applications more readily.

Bibliography

[1] S. Choi and K. J. Kuchenbecker, "Vibrotactile Display: Perception, Technology, and Applications," in *Proceedings of the IEEE*, vol. 101, 2013, pp. 2093–2104.

[2] M. M. Wanderley and P. Depalle, "Gestural Control of Sound Synthesis," *Proceedings of the IEEE*, vol. 92, no. 4, pp. 632–644, 2004.

[3] M. T. Marshall, "Physical Interface Design for Digital Musical Instruments," PhD thesis, McGill University, 2008, p. 307.

[4] D. M. Birnbaum, "Musical Vibrotactile Feedback," PhD thesis, McGill University, 2007, p. 83.

[5] M. T. Marshall and M. M. Wanderley, "Vibrotactile Feedback in Digital Musical Instruments," in *Proceedings of NIME 2006*, 2006, pp. 226–229.

[6] M. Giordano and M. M. Wanderley, "Perceptual and Technological Issues in the Design of Vibrotactile-augmented Interfaces for Music Technology and Media," in *Haptic and Audio Interaction Design. 8th International Workshop, HAID 2013, Daejeon Korea.* I. Oakley and S. Brewster, Eds., vol. LNCS 7989, Berlin: Springer, 2013, pp. 89–98.

[7] G. A. Gescheider, J. H. Wright, and R. T. Verrillo, *Information-Processing Channels in the Tactile Sensory System*, 1st Editio. New York: Psychology Press, 2009, pp. 1–.

[8] S. Papetti, M. Fröhlich, F. Fontana, S. Schiesser, and F. Avanzini, "Implementation and Characterization of Vibrotactile Interfaces," in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 13, pp. 257–282.

[9] C. B. Medeiros and M. M. Wanderley, "A Comprehensive Review of Sensors and Instrumentation Methods in Devices for Musical Expression," *Sensors*, vol. 14, no. 8, pp. 13 556–13 591, 2014.

[10] I. Hattwick, I. Franco, and M. M. Wanderley, "The Vibropixels: A Scalable Wireless Tactile Display System," in *Human Interface and the Management of Information: Information, Knowledge and Interaction Design. HIMI 2017*, S. Yamamoto, Ed., vol. LNCS, vol, Vancouver, Canada: Springer, Cham, 2017, pp. 517–528.

[11] D. C. Egloff, M. M. Wanderley, and I. Frissen, "Haptic Display of Melodic Intervals for Musical Applications," in *2018 IEEE Haptics Symposium (HAPTICS)*, San Francisco, CA: IEEE, 2018, pp. 284–289.

[12] M. Giordano, “Vibrotactile Feedback and Stimulation in Music Performance,” PhD thesis, McGill University, 2016.

[13] K. Minamizawa, Y. Kakehi, M. Nakatani, S. Mihara, and S. Tachi, “TECHTILE Toolkit: A Prototyping Tool for Design and Education of Haptic Media,” in *Proceedings of the 2012 Virtual Reality International Conference (VRIC ’12)*, Laval, France: ACM Press, 2012.

[14] A. Israr, S. Zhao, K. Mcintosh, Z. Schwemler, A. Fritz, J. Mars, J. Bedford, C. Frisson, I. Huerta, and M. Kosek, “Stereohaptics: A Haptic Interaction Toolkit for Tangible Virtual Experiences,” in *ACM SIGGRAPH 2016 Studio, Jul 2016*, Anaheim, United States, 2016.

[15] K. E. MacLean, “Foundations of Transparency in Tactile Information Design,” *IEEE Transactions on Haptics*, vol. 1, no. 2, 2008.

[16] S. O’Modhrain and R. B. Gillespie, “Once More, with Feeling: Revisiting the Role of Touch in Performer-Instrument Interaction,” in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 2, pp. 11–28.

[17] E. R. Miranda and M. M. Wanderley, *New Digital Musical Instruments: Control And Interaction Beyond the Keyboard*, Vol. 21. Middleton, Wisconsin: AR Editions, Inc., 2006.

[18] J. Rovan and V. Hayward, “Typology of Tactile Sounds and their Synthesis in Gesture-Driven Computer Music Performance,” in *Trends in Gestural Control of Music*, M. M. Wanderley and M. Battier, Eds., IRCAM, 2000, pp. 297–320.

[19] A. Bongers, “Tactual Display of Sound Properties in Electronic Musical Instruments,” *Displays*, vol. 18, no. 3, pp. 129–133, 1998.

[20] M. Puckette and Z. Settel, “Non-Obvious Roles for Electronics in Performance Enhancement,” in *Proceedings of the International Computer Music Conference*, 1993, p. 134.

[21] M. Giordano and M. M. Wanderley, “A Learning Interface for Novice Guitar Players Using Vibrotactile Stimulation,” in *Proceedings of the Sound and Music Computing Conference, SMC*, Padova, Italy, 2011, pp. 368–374.

[22] S. Papetti and C. Saitis, “Musical Haptics: Introduction,” in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 1, pp. 1–10.

[23] C. Seim, T. Estes, and T. Starner, “Towards Passive Haptic Learning of Piano Songs,” in *IEEE World Haptics Conference, WHC 2015*, 2015, pp. 445–450.

[24] M. Giordano, J. Sullivan, and M. M. Wanderley, “Design of Vibrotactile Feedback and Stimulation for Music Performance,” in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 10, pp. 193–214.

[25] S. J. Bolanowski, G. A. Gescheider, R. T. Verrillo, and C. M. Checkosky, “Four Channels Mediate the Mechanical Aspects of Touch,” *The Journal of the Acoustical Society of America*, vol. 84, no. 5, pp. 1680–1694, 1988.

[26] L. S. Costanzo, “Neurophysiology,” in *Physiology*, 6th Editio, Philadelphia: Elsevier, Inc., 2018, ch. 3, pp. 69–116.

[27] K. S. Chen and E. L. Feldman, "Peripheral Neuropathies," in *Youmans and Winn Neurological Surgery, Volume 3*, 7th Editio, Philadelphia: Elsevier, Inc., 2017, ch. 246, pp. 1985–1995.

[28] V. Hayward, "A Brief Overview of the Human Somatosensory System," in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 3, pp. 29–48.

[29] Å. B. Vallbo and R. S. Johansson, "Properties of cutaneous mechanoreceptors in the human hand related to touch sensation.," *Human Neurobiology*, vol. 3, no. 1, pp. 3–14, 1984.

[30] F. Fontana, S. Papetti, H. Järveläinen, F. Avanzini, and B. L. Giordano, "Perception of Vibrotactile Cues in Musical Performance," in *Musical Haptics*, S. Papetti and C. Saitis, Eds., Cham, Switzerland: Springer, 2018, ch. 4, pp. 49–72.

[31] R. T. Verrillo, "Vibration Sensation in Humans," *Music Perception*, vol. 9, no. 3, pp. 281–302, 1992.

[32] L. Wyse, S. Nanayakkara, P. Seekings, S. H. Ong, and E. A. Taylor, "Palm-area sensitivity to vibrotactile stimuli above 1 kHz," in *NIME'12 Proceedings of the International Conference on New Interfaces for Musical Expression*, Michigan, 2012.

[33] R. T. Verrillo, A. J. Fraioli, and R. L. Smith, "Sensation magnitude of vibrotactile stimuli," *Perception & Psychophysics*, vol. 6, no. 6, pp. 366–372, 1969.

[34] R. T. Verrillo and G. A. Gescheider, *Perception via the Sense of Touch*, I. Summers, Ed. London: Whurr Publishers Ltd., 1992, pp. 1–36.

[35] C. Salisbury, R. B. Gillespie, H. Z. Tan, F. Barbagli, and J. K. Salisbury, "What You Can't Feel Won't Hurt You: Evaluating Haptic Hardware Using a Haptic Contrast Sensitivity Function," *IEEE Transactions on Haptics*, vol. 4, no. 2, pp. 134–146, 2011.

[36] M. Yamaoka, A. Yamamoto, and T. Higuchi, "Basic analysis of stickiness sensation for tactile displays," in *Haptics: Perception, Devices and Scenarios. EuroHaptics 2008. Lecture Notes in Computer Science, vol. 5024*, M. Ferre, Ed., Berlin, Heidelberg: Springer, 2008, pp. 427–436.

[37] J. Ryu, J. Jung, S. Kim, and S. Choi, "Perceptually Transparent Vibration Rendering Using a Vibration Motor for Haptic Interaction," in *RO-MAN2007 - The 16th IEEE International Symposium on Robot and Human Interactive Communication*, Jeju, Korea: IEEE, 2007, pp. 310–315.

[38] A. T. Bukkapatnam, P. Depalle, and M. M. Wanderley, "Autoregressive Parameter Estimation for Equalizing Vibrotactile Systems," in *International Workshop on Haptic and Audio Interaction Design (HAID)*, Lille, France.: HAL, 2019.

[39] B. Duncan, "Specifications and Testing," in *High Performance Audio Power Amplifiers*, Elsevier, 1996, ch. 7, pp. 287–288.

[40] E. Rustighi, W. Kaal, S. Herold, and A. Kubbara, "Experimental Characterisation of a Flat Dielectric Elastomer Loudspeaker," *Actuators*, vol. 7, no. 28, pp. 1–14, 2018.

[41] S. Maeda and M. J. Griffin, "A Comparison of Vibrotactile Thresholds on the Finger Obtained with Different Equipment," *Ergonomics*, vol. 37, no. 8, pp. 1391–1406, 1994.

[42] W. Klippel, "Tutorial: Loudspeaker Nonlinearities - Causes, Parameters, Symptoms," *AES: Journal of the Audio Engineering Society*, vol. 54, no. 10, pp. 907–939, 2006.

[43] A. Voishvillo, "Assessment of Nonlinearity in Transducers and Sound Systems – From THD to Perceptual Models," in *121st AES Convention on Loudspeakers*, 2006.

[44] C. Branje, M. Maksimowski, M. Karam, D. I. Fels, and F. Russoy, "Vibrotactile Display of Music on the Human Back," in *3rd International Conference on Advances in Computer-Human Interactions, ACHI 2010*, 2010, pp. 154–159.

[45] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms & Applications*, 4th ed. Upper Saddle River, N.J.: Pearson Prentice Hall, 2007.

[46] P. Welch, "On the Variance of Time and Frequency Averages over Modified Periodograms," in *ICASSP 1977. IEEE International Conference on Acoustics, Speech, and Signal Processing*, Hartford, CT, USA: IEEE, 1977, pp. 58–62.

[47] S. M. Kay and S. L. Marple, "Spectrum Analysis-A Modern Perspective," *Proceedings of the IEEE*, vol. 69, no. 11, pp. 1380–1419, 1981.

[48] W. Kester, "MT-003 Tutorial: Understand SINAD, ENOB, SNR, THD, THD + N, and SFDR so You Don't Get Lost in the Noise Floor," Analog Devices, Inc., Tech. Rep., 2009. [Online]. Available: `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.476.1903&rep=rep1&type=pdf` (visited on 01/11/2019).

[49] Audio Precision, *More about THD+N and THD*, 2013. [Online]. Available: `https://www.ap.com/technical-library/more-about-thdn-and-thd/` (visited on 01/11/2019).

[50] MathWorks®, *Distortion Measurements*, n.d. [Online]. Available: `https://www.mathworks.com/help/signal/ug/distortion-measurements.html` (visited on 10/25/2018).

[51] L. Ljung, *System Identification Theory for User*. Englewood Cliffs, New Jersey: PTR Prentice-Hall, 1987.

[52] J. Makhoul, "Linear Prediction: A Tutorial Review," *Proceedings of the IEEE*, vol. 63, no. 4, pp. 561–580, 1975.

[53] G. U. Yule, "On a Method of Investigating Periodicities in Disturbed Series, with Special Reference to Wolfer's Sunspot Numbers," *Philosophical Transactions of the Royal Society of London. Series A*, vol. 226, no. 642, pp. 267–298, 1927.

[54] G. Walker, "On Periodicity in Series of Related Terms," *Proceedings of the Royal Society. Series A*, vol. 131, no. 818, pp. 518–532, 1931.

[55] J. O. Smith, *Introduction to Digital Filters with Audio Applications*, 2007th ed. online book, 2007. [Online]. Available: `https://ccrma.stanford.edu/~jos/fp/Definition_Minimum_Phase_Filters.html` (visited on 10/13/2018).

[56] S. Temme, "Why and How to Measure Distortion in Electroacoustic Transducers," in *AES 11th International Conference: Test & Measurement (May 1992)*, Portland, Oregon, 1992.

[57] J. Malloch and M. M. Wanderley, "The T-Stick: From Musical Interface to Musical Instrument," in *2007 Conference on New Interfaces for Musical Expression (NIME07)*, 2007, pp. 66–70.

[58] A. H. Gray and J. D. Markel, "A Spectral-Flatness Measure," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-22, no. 3, pp. 207–217, 1974.

Appendix A

Links to Device Datasheets

A.1 Actuators

Device	Link to Datasheet
TactileLabs Haptuator mkII (TL-002-09-C)	http://tactilelabs.com/wp-content/uploads/2012/07/TL002-09-A_v1.01.pdf
TactileLabs Haptuator Re-design (TL002-14R)	http://tactilelabs.com/wp-content/uploads/2012/07/HaptuatorRedesignSpec_v1.0.pdf
TDK PowerHap 7G	https://product.tdk.com/info/en/documents/data_sheet/20/10/ds/PowerHap_7G.pdf

Table A.1 List of actuators selected for testing and links to their respective datasheets.

A.2 Amplifiers

Device	Link to Datasheet
SureElectronics (WONDOM) AA-AB33182	`https://www.parts-express.com/pedocs/manuals/320-335m.pdf`
SureElectronics (WONDOM) AA-AB32971	`https://www.parts-express.com/pedocs/specs/320-3340--2%20X%20100Watt%20T-AMP%20AA-AB32971%20MidPowerStereoSeries%20Manual%20(3).pdf`
TactileLabs HaptuAmp Mini	`http://tactilelabs.com/wp-content/uploads/2016/12/MiniHaptuampDatasheet_1_1.pdf`

Table A.2 List of amplifiers selected for testing and links to their respective datasheets.

A.3 Biasing Circuitry

Device	Link to Datasheet
TI DRV2700EVM	`http://www.ti.com/lit/ds/symlink/drv2700.pdf`
DRV2700 circuit	`http://www.ti.com/lit/ds/slos861b/slos861b.pdf`

Table A.3 Link to Datasheet of DRV2700 Piezo Biasing circuit board.

Appendix B

Characteristics of a Low-Cost Audio Interface

Audio interfaces are a common component in the implementation of vibrotactile interfaces (see Section 2.3.3). They play the vital role of a DAC, translating digitally generated signals into continuous-time AC signals. We originally envisioned that the characteristics of the implemented toolkit could be better guaranteed by including an inexpensive audio interface with it. A low-cost (approximately 10$) and portable (approximately 0.4 g) USB audio interface is evaluated while exploring the possibility of including it in the toolkit. The *Trond AC-2* audio interface did not have available data-sheets for making a more informed selection, and was chosen purely for reasons of portability and cost.

Since including an audio interface was not among the primary goals of this study, the results of characterizing the selected audio interface are presented here in this appendix. This appendix presents the characteristics obtained from electrical measurements of the *Trond AC-2* audio interface.

B.1 Measurement Overview

When performing measurements of the audio interface, the output of the interface was fed directly to the input of the *National Instruments NI USB-4431* set to *Voltage* mode. The interface was made to output the white noise and sinusoidal test signals of 0 dBFS and -10 dBFS, to obtain its magnitude-frequency and THD+N characteristics. The analysis of the measurements is done with the same approach as the one discussed in Chapter 3, for characterizing the selected amplifiers.

B.2 Results of Characterization

The THD+N and magnitude-frequency characteristics of the *Trond AC-2* USB audio interface are measured at its maximum output. The voltage output by the audio interface was directly acquired by the *National Instruments NI USB-4431* data acquisition card, at a sample rate of 48 kHz. The THD+N characteristics are shown in Figure B.1a to be consistently less than 1% from 150-1000 Hz. The magnitude response (Figure B.1b) shows a high-pass characteristic that filters out frequencies below approximately 100 Hz. The response after that is observed to be flat, i.e., without visible resonant features.

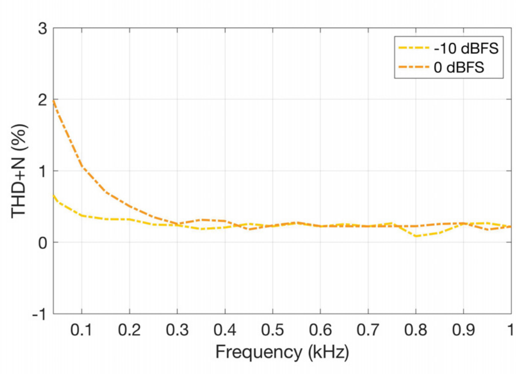

(a) THD+N characteristics.

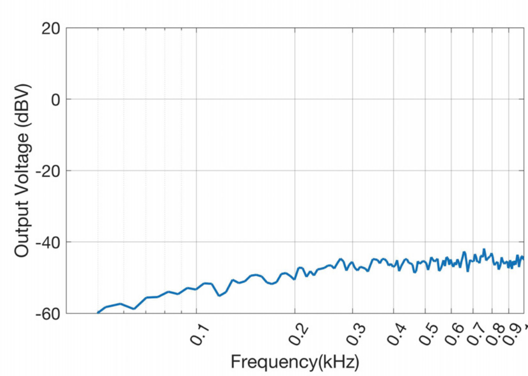

(b) Magnitude-Frequency response.

Fig. B.1 THD+N and Magnitude-Frequency responses of the *Trond AC-2* audio interface. The THD+N is under 1% only within the bandwidth 150-1000 Hz. The Magnitude-Frequency characteristics show a poor low-frequency response below 100 Hz.