A SYSTEMATIC APPROACH TO MUSICAL VIBROTACTILE FEEDBACK

David M. Birnbaum and Marcelo M. Wanderley
Input Devices and Music Interaction Lab
Centre for Interdisciplinary Research in Music Media and Technology
McGill University – Montreal, QC, Canada

ABSTRACT

This paper presents a new approach to the integration of vibrotactile feedback into digital musical instruments. A design strategy for musical vibrotactile systems is developed based on stimulator properties and neurophysiological studies of vibrotactile perception. A software vibration synthesizer driven by perceptual sound features extracted from audio feedback has been created based on these concepts. This framework will help to simplify integration of vibrotactile feedback into instrument designs by defining high-order tactile invariants, avoiding the need to explicitly specify low-level vibration stimulus parameters.

1. MOTIVATION

Touch has been vital to the development of musical skill for millennia, but the recent dominance of digital technology in composition and sound production techniques has separated embodiment from experiences of playing and listening [1]. Electronic music culture has compensated for this missing element in various ways [11]; however, technology is now sufficiently advanced that the haptic channel may be re-engaged. It is well known that acoustic vibrations are utilized for self-monitoring in acoustic performance [28, 8, 37], and that vibrotactile feedback can considerably improve touch perception without adding significant complexity or cost to an interface [27]. This combination of circumstances should make integration of vibrotactile feedback systems into digital musical instruments (DMIs) [41] a top priority for instrument designers. Musical vibrotaction is a high-resolution, high-bandwidth perceptual system that promises nothing less than the reestablishment of embodied experience in electronic musical discourse.

The development of techniques for displaying artificial vibrations in the frequency and amplitude ranges of musical sound is worthwhile because vibrotactile feedback may improve controllability of certain musical processes [25, 33]. However, while simulation of acoustic vibrations is certainly a valid approach, it is not necessary to limit artificial vibrotactile feedback to a simulation model — vibrations with patently un-acoustic properties should also hold potential to complement and reinforce sound feedback. The only defining design requirement of a musical vibrotactile display is the mapping of musical properties to vibrotactile ones.

To explore these issues, the authors have constructed a musical vibrotactile feedback system based on an open-tonehole flute, dubbed the *Tactilicious Flute Display* (*TFD*). [1] Although other instruments (such as a cello) touch the body in more places than a flute [7], the hairless (or *glabrous*) skin of the hand and mouth used in keyless, pre-Böhm flute performance are the areas of the human body most sensitive to vibration [40]. The vibration stimuli produced by an open-tonehole flute thus present clear design criteria, as these flute players are in the unique position of having their highly sensitive fingertips and lips in direct contact with the resonating air column that radiates both sound and vibration.

Representing music with vibration necessitates a cross-modal mapping, giving rise to cognitive and perceptual issues that do not play a role in teletaction or virtual touch. To include a wide range of perceptual variables, the musical output of the TFD has been chosen to consist of breakbeat patterns [10]. The breakbeat idiom exhibits features such as repetition of rhythmic phrases, a sound spectrum saturated with a wide range of frequencies, and a mix of distinct parts or voices. It is important to note that the counterintuitive idea of a flute playing breakbeat music was conceived in two stages — the physical system was designed to maximize vibration sensation, and the vibrotactile feedback was programmed to facilitate perception of separable rhythmic elements.

2. PREVIOUS WORK

One area of tactile interface research that has received considerable attention recently is mobile applications. A crude notification device, the pager motor, is perhaps the most familiar and commonplace use of vibrotactile feedback. Improving on these blunt instruments, systems have recently been developed that are capable of a more nuanced response. A mobile device that provides a background (or “ambient”) channel of information, dubbed a “peripheral awareness display”, was implemented in the *TouchEngine* system. It was shown that the *TouchEngine*’s piezoceramic vibrotactile actuator improved task comple-

[1] The device is a *display* that operates in two modes, sound and vibration. It is not a *controller* because it does not allow user input.

tion time for several mobile interactions [31]. Another vibrotactile display for mobile devices has been proposed where vibration cues are coupled to ringtones. Audio signal is used to drive an actuator directly, with a modest amount of signal processing to boost the tactile range [29].

An early example of vibrotactile feedback based on audio manipulation was a rhythm display for deaf music students. Actuators in the players' chairs were driven by a filtered version of the musical output of their instruments. The students reported a higher amount of "enjoyment" and "appreciation" for music when vibrotactile sensations supplemented their playing experience [12].

Tactile vocoders display vibrotactile feedback to facilitate speech perception [5]. Similarly, structured tactile messages called *tactons*, which use the metaphor of an iconic symbol to represent a concept, have been incorporated into user interfaces [4, 6]. These systems differ from the TFD in that they display symbolic information through sensory substitution and are usually designed for efficiency of communication. In contrast, the TFD supplements the auditory sensory channel with vibrotactile feedback using relatively straightforward manipulations of digital audio signal.

Skinscape used a music composition model to develop an approach to artistic tactile composition [17]. Although primarily concerned with aesthetics, *Skinscape* is similar to the TFD insofar as low-level synthesis is used to create an evolving vibration signal with similarities to music. Another system, called *VR/TX*, was designed to augment feedback in electronic instruments [34]. *VR/TX* utilizes spatiotemporal classification criteria for encoding feedback as *tactile stimulation events*, similar to the vibrotactile events produced by the TFD.

3. RELATED RESEARCH

Just as audio feedback is more effective if it integrates psychoacoustical models, vibrotactile feedback stands to be improved if based on a comprehensive understanding of vibrotactile perception. It is generally accepted that mechanoreceptors in the skin enable tactile sensation, and proprioceptors in joints, muscles, and ligaments give rise to kinesthesia. [2] *Haptic perception* is defined specifically as physical object or event perception external to the body effectuated by a combination of active, exploratory kinesthetic perception and passive cutaneous sensation [23]. Vibrotaction is thus a vital component of haptic perception. Neurophysiological research on vibrotaction has focused on mechanoreceptive nerve fiber response, the cortical entry stage, and subjective judgments of perceptual characteristics such as threshold of detection, magnitude, and frequency [20].

In order to display vibration to a performer that is perceived as meaningful and tightly coupled to sound feedback, we must first examine the capabilities and dimensions of the vibrotactile system. The dimensions included in this model are *pitch*, *loudness*, *brightness*, and *envelope*.

3.1. Pitch

Pitch perception is such a fundamental aspect of musical experience that it seems to naturally command a dominant role in feedback, in both auditory and vibrotactile modes. Vibrotactile "pitch" is a term that highlights sensitivity to the rate of periodic stimuli, as it does in psychoacoustics [35]. Unfortunately, the neural coding mechanisms for signaling information about the frequency of vibrotactile stimuli are not well understood. In audition, pitch is almost exclusively dependent on stimulus frequency; amplitude and waveform have little effect on pitch perception. In contrast, vibrotactile pitch is complicated by several factors, such as the multichannel nature of the cutaneous sense organ [26], amplitude of skin displacement [26], stimulus duration [16], and body locus [18]. A comprehensive theory of vibrotactile pitch would be very useful for interaction design, but attempts at developing such a theory have fallen short of proposing a universal and comprehensive translation scheme from auditory pitch.

Several researchers have proposed that the high interdependency of frequency and amplitude suggests they be considered a single vibrotactile stimulus parameter [4]. Moreover, it has been claimed that the only musical parameters representable by vibration are timing, amplitude, and spectral weighting (relative amount of harmonic content) [7]; frequency is excluded because tactile frequency discrimination has been shown to be poor compared to audition [40]. Yet vibrotactile pitch exhibits an important similarity to auditory pitch — within certain ranges, frequency discrimination fits a critical band model [24]. Furthermore, while frequency perception is known to be contingent on the stimulus properties mentioned above, a feedback system may be designed to incorporate them as independent variables, which could help account for these dependencies. Most significantly, it has been shown that certain frequency ranges give rise to distinct subjective sensations [40]. This seems to imply that while vibration frequency does not directly correlate to vibrotactile pitch, it nevertheless can be exploited for musical feedback as long as it is understood that pitch must be considered dependent on amplitude and that the sensory quality varies dramatically with frequency.

There are two aspects of vibrotactile pitch that are commonly studied: frequency following response, referring to the fidelity of the entrained neural firing pattern to a periodic stimulus, and frequency discrimination, concerned with the just-noticable difference (JND) between pitches and the number of discriminable pitches in the vibrotactile range. The JND increases with frequency [40], leading to

[2] In terms of stimulus characterization, these two modes are not distinct but rather comprise a "kinesthetic-cutaneous continuum", where low frequency, high amplitude stimuli that move parts of the body relative to each other constitute "forces" activating kinesthesia, and higher frequency, lower amplitude stimuli fall under the "vibration" category and activate cutaneous mechanoreceptors [38]. Recent research has further revised this model by showing that cutaneous mechanoreceptors contribute to kinesthesia, for example by responding to internal vibrations and skin stretch [9].

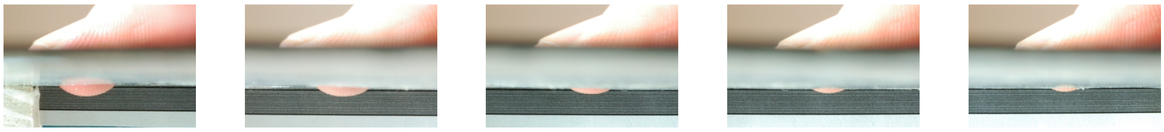

(a) 10mm (b) 9mm (c) 8mm (d) 7mm (e) 6mm

Figure 1. Index fingertip extension, through holes of various diameters (marked).

the design guideline that the lower frequency range should use wider pitch bands than the higher range. Numerous studies have shown that each mechanical afferent type responds strongest within a "primary" frequency range [22]. When subjective magnitude is made equal, the number of discriminable pitches is still dependent on whether pitch is considered as relative or absolute [4]. In [15] it is alleged that up to nine perceivable discrete pitches should be used for symbolic information, and [34] hypothesizes between eight and ten discriminable pitches, but neither of these guidelines seem to be based on formal studies.

3.2. Loudness

Vibrotactile "loudness" refers to the subjective variable that responds most directly to the amplitude of skin displacement. The threshold of perception is the lowest amplitude of periodic displacement that can be detected as a tactile sensation. Within the range of about 20–40Hz, the threshold of vibration perception is independent of frequency. Between about 40–700Hz, however, sensitivity peaks at about 250Hz [21]. The threshold is also responsive to the presence of a non-vibrating element around the contactor, called a *surround*. A surround decreases the threshold in the lower range and increases it in the higher range [14]. This is an important consideration because the body of the TFD is in essence a surround. Stimulus envelope, duration, temporal masking, and skin impedance also affect perceived magnitude [22]. It is suggested in [15] that up to four levels of vibration amplitude are easily discriminated. Loudness must also take into account stimulation boundaries; vibrotactile pulses or events must occur above the threshold of perception, but not be so loud that the sensation is uncomfortable. It is also noted in [40] that with smaller contactor sizes (roughly the same size of the contactors in the TFD), frequency has no discernible effect on threshold or suprathreshold sensation magnitude.

3.3. Brightness

Complex waveforms are not distinguishable by vibrotaction to nearly the same extent they are in audition, but there are waveform properties that can be distinguished, viz., amount of harmonic content, periodicity, and certain ranges of modulation [2]. It has been reported that the spectrum from sine (periodic, no partials) to square (periodic, many partials) to noise (non-periodic) is subjectively sensed as a spectrum from "smoothness" to "roughness" [34]. This and other research suggests that there is a parameter of vibrotactile "brightness" that can be targeted by varying periodicity and harmonic content in the stimulus.

3.4. Envelope

The envelope of a vibration sensation is affected by many factors. Because envelope is time-dependent, adaptation, summation, enhancement, and temporal masking play a significant role in perception; dynamic qualities such as attack, sustain and decay durations should take these into account [39].

3.5. Four channels of mechanoreception

Cutaneous sensitivity differs from hearing in that several more channels mediate sensation at the afferent level. The four-channel model of mechanoreception delimits the neural processing of vibrotaction into four psychophysical channels, identified with the four known mechanoreceptors in glabrous skin [3]. These four neural units are abbreviated as FAI, FAII, SAI, and SAII (see Table 1). FA units adapt quickly to ramp-and-hold stimuli, meaning that they fire while the skin is being moved or deformed, but stop firing when the skin movement stops. SA units adapt slowly, continuing to fire the entire duration the skin is indented. [3] Each of these nerve fiber types activate channels that, when stimulated independently, produce "unitary" sensations. However everyday suprathreshold sensations are a result of the combination of neural activity across the four channels. The framework presented in this paper integrates the multichannel nature of mechanical touch into a vibration synthesis approach, described in Section 4.2.

4. TACTILICIOUS FLUTE DISPLAY

4.1. Actuation

Transducer design affects controllability, and some transducers may be better suited than others to certain musical tasks [42]. The approach to actuator selection and placement presented here is specific to fingertip stimulation for a flute-like interface.

The stimulators used in the TFD are modified voice coils. Miniature stimulators such as these (previously used in [30], to cite one of many examples) are convenient because they are inexpensive, highly efficient and responsive to signal dynamics, and easy to control. Voice coils were, after all, designed for musical applications (e.g.,

[3] Note that because vibration is an alternating signal, the role of the adaptation characteristic is not clear-cut; both FA and SA type afferents continually fire when stimulated with vibration [21].

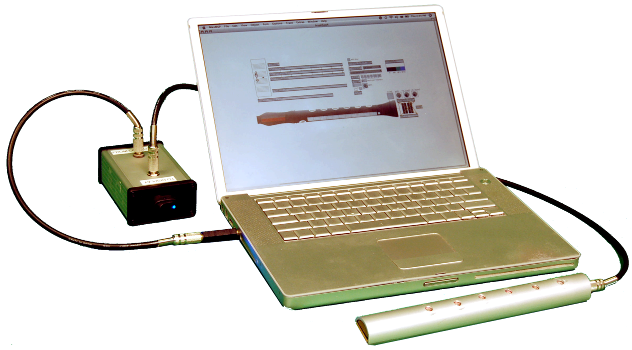

Figure 2. The TFD connected to the audio out of a laptop through a current amplifier.

Figure 3. Closeup of a tonehole with vibrotactile actuator. The metal ball in the center of the diaphragm is a mass load that lowers the diaphragm's resonant frequency and increases inertia [31].

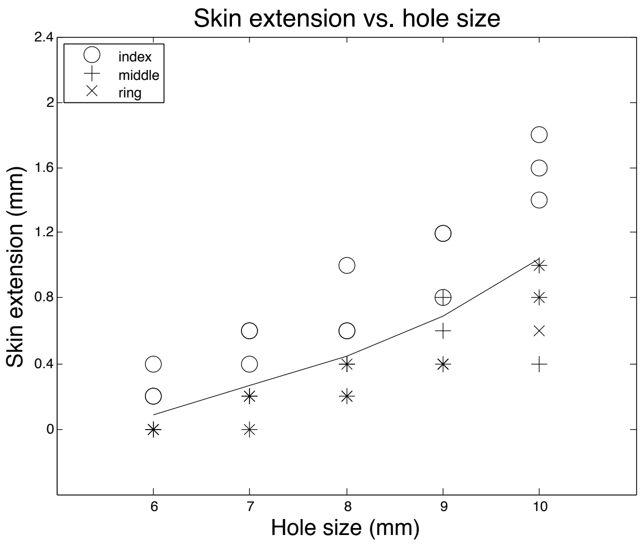

Figure 4. Hole size versus skin extension in three fingers on each of three different people. Because resolution was only 0.2mm, some points overlap. The line represents average skin extension of all fingers.

loudspeakers). However, they are less resistant to interference from external forces exerted by the human body when compared to actuators with more inertia such as unbalanced motors or cylindrical contactors. These voice coils are not backdrivable; contact with fingertips significantly alters their output. Therefore sensation is maximized if the actuator is placed at an optimal distance from the skin's surface so that the skin is maximally stimulated by the actuator and the actuator is minimally dampened by the skin. The tactile response of a voice coil will be dramatically improved if placed just close enough to the skin to be felt. [4]

Bearing this in mind, the question is raised of how far to position the actuator below the surface of the hole (see Figure 3). Because the deformation quality of glabrous skin is similar to that of a fluid-filled sack [21], pressing on a tonehole causes the skin to extend down past the surrounding surface a distance that is determined by the pressure applied and the size of the hole. A small experiment was conducted to relate tonehole size to skin extension, in order to determine the optimal distance between the contactor and the outer surface of the surround for this particular application.

Three recorder players (males aged 26, 30, and 31) were directed to press their index finger down on a rigid 1mm-thick metal surface with five drilled holes, measuring 6, 7, 8, 9, and 10mm in diameter. [5] Behind the metal surface, a card with horizontal black lines spaced 0.2mm apart served as a reference (see Figure 1). A high resolution photograph was then taken as the subject pressed his index, middle, and ring fingers down on each hole. The number of lines obscured by the fingertip in the photograph were counted to determine the depth that the skin extended down below the surface. The importance of considering hole size when placing the actuators is clearly shown in Figure 4. Variability across subjects was significant enough to suggest that actuator placement may be further improved by interface personalization. However there was less variation amongst subjects and their individual fingers when the hole size was smaller, suggesting that if an interface is to be used by multiple players without biasing the feedback for the use of certain fingers over certain toneholes, a smaller tonehole size should be used.

4.2. Vibration programming

To generate stimulation codes — the combined total of which the authors term the "vibration program" — vibration signal parameters must be changed over time. The goal of modeling psychophysical vibration channels with

[4] The actuator system does not account for at-rest static skin pressure, damping, or skin impedance. A stimulator with an integrated non-zero-force indicator would be necessary to place the actuator against the surface of the skin with the least amount of static pressure, and a vibrometer to sense the stimulator's position would allow for tuning of absolute skin displacement [20].

[5] Applied pressure was not measured; instead the subjects were directed to press with the amount of force they would typically use to cover a tonehole. The effect of dynamic pressure would be an interesting topic for future study, but it was excluded from these tests.

Psychophysically defined channel:	**P**	**NPI**	**NPII**	**NPIII**
Full name:	Pacinian	Non-Pacinian I	Non-Pacinian II	Non-Pacinian III
Physiological type:	FAII	FAI	SAII	SAI
Fiber innervation density (fingertip, per cm^2):	21	140	49	70
Subjective sensation:	"vibration"	"flutter"	(unknown)	"pressure"
Frequency range:	40–500Hz	2–40Hz	100–500Hz	0.4–3.0Hz
Prime sensitivity range:[1]	250–300Hz	25–40Hz	150–400Hz	0.4–1.0Hz
Shape of frequency response function:	U-shape	Flat[2]	U-shape	Flat

[1] Defined as best frequencies to lower threshold of perception
[2] Notch at 30Hz

Table 1. Vibrotactile channel characteristics, adapted from [21].

a feedback synthesizer is to investigate how separate channels are excited by the extracted perceptual sound features. Transposing musical signal descriptors into the prime sensitivity ranges of these channels may be an effective technique for encoding music as vibration. This section outlines some of the characteristics of each of the channels and investigates their role in mediating TFD vibrotactile feedback.

4.2.1. P

Higher frequencies (40–500Hz) are felt as a "hum" or "buzz" and excite the P channel, which is thought to be the system most directly responsible for vibrotactile perception. There is evidence that this channel integrates stimulus energy over time [2], and its peak sensitivity occurs at about 250Hz. It has a U-shaped equal sensation magnitude contour, similar to the Fletcher-Munsen equal loudness contours of psychoacoustics, except that it does not flatten out as intensity increases [40]. It may thus be beneficial for feedback targeting this range to be filtered to account for this curve. Because the magnitude of the sensory response is directly dependent on amplitude of skin displacement, an accurate filter model would need to include a translation from sound intensity to a spatial skin displacement measure such as microns. The filter used for the TFD does not do this and instead is based on subjective reports of equal sensation magnitude.

4.2.2. NPI

With the highest innervation density in the human fingertip, it follows that the NPI channel is highly responsive to feedback targeting this location. There are approximately 140 FAI neural units per cm^2 in glabrous skin, making the NPI twice as "sensitive" as the next most innervated channel, the NPIII. If innervation directly affects perceived magnitude, the intensity of stimuli in the NPI range (2–40Hz) should be de-emphasized to account for this heightened sensitivity. A flattening function is not vital because the response of the NPI is naturally flat, excluding "notch" at 30Hz.[6] The TFD may engage the NPI with variations in vibrotactile brightness as this channel has been found to be particularly well suited for encoding stimulus waveform [2].

[6] It may be reasonable to include a peak filter at 30Hz to remove this nonlinearity. The notch was not accounted for in this DSP framework.

4.2.3. NPII

The frequency following response of the NPII (100–500Hz) lies within that of the P channel, but it is particularly sensitive to lateral skin stretch. Its high vibrotactile threshold characteristic makes its role in vibrotactile coding difficult to discern [35]. However the four-channel model implies that vibrotactile feedback can excite this channel with suprathreshold stimulation, allowing the NPII's unitary subjective sensory quality to serve as a viable mediator of musical feedback. The mechanical stimulation of the NPII necessarily activates the P channel well above its threshold, which raises some interesting questions about how information might be displayed to the NPII. However the actuators used in this implementation are neither accurate nor powerful enough to directly engage the NPII, and so while the NPII probably does mediate TFD feedback, it is left for future research to explore how and to what extent it does so.

4.2.4. NPIII

The NPIII is chiefly responsive to pressure or very low frequency periodic skin displacement. It would perhaps be possible for a custom actuator to display information to the NPIII through the use of step functions, or multiple levels of sustained pressure, and simultaneously display periodic stimuli. However the actuators used in the TFD cannot output a high-amplitude, sustained offset stimulus above the NPIII threshold; a transducer with the combined ability to display low frequency offsets and high frequency vibrations would be an excellent tool for tactile interfacing.

4.3. Software implementation

An abstract model of musical vibrotactile perception has been programmed in Max/MSP, outlined in Figure 5. An analysis layer extracts musical information from the audio signal, which is then mapped to vibrotactile perceptual parameters, consisting of *pitch*, *loudness*, *brightness*, and *envelope trigger*.

External objects presented in [19] are used to extract audio features. The *noisiness∼* object, which outputs a measure of spectral flatness, is mapped to vibrotactile brightness. The brightness of the vibration signal is determined by an equal power crossfade between a sine wave and a square wave. More tonality in the sound feedback is thus represented with a richer harmonic spectrum in the vibrotactile domain.

The *brightness∼* object calculates spectral centroid, a metric that has been shown to act as a determinant of drum part separation in percussion listening [13]. The spectral centroid is scaled with the lowest frequency mapped to a 40Hz vibrotactile pitch and the highest frequency mapped to 400Hz, so that the different drums in the breakbeat are represented by relative vibrotactile pitch aimed at the P channel. It will be recalled that the JND of vibrotactile pitch is larger in the lower ranges; thus a logarithmic frequency scale is applied to assure lower pitches include more frequencies than high ones. Of the drum loops tested, the typical distance between the lowest vibrotactile pitch and the highest is about 100–300Hz.

The output of the *loudness∼* object is mapped directly to the amplitude of the vibration waveform. The narrower dynamic range of vibrotaction was compensated for in the post-processing stage with a peak compressor.

The onset detection external *bonk∼* [32] is used to drive a simple envelope generator with an adjustable decay to create the sensation of discrete vibrotactile pulses with the above characteristics.

After a vibrotactile event is synthesized with the above characteristics, a bandpass filter is applied to remove frequencies out of the vibrotactile range. The signal then passes through a second filter acting as a frequency flattening function to compensate for the nonlinear response of the upper ranges of vibrotaction (P channel). Dynamic range is then reduced using the *omx.peaklim∼* object so that quieter vibrotactile events are not lost. Extra-vibrotactile frequencies are then filtered out again.

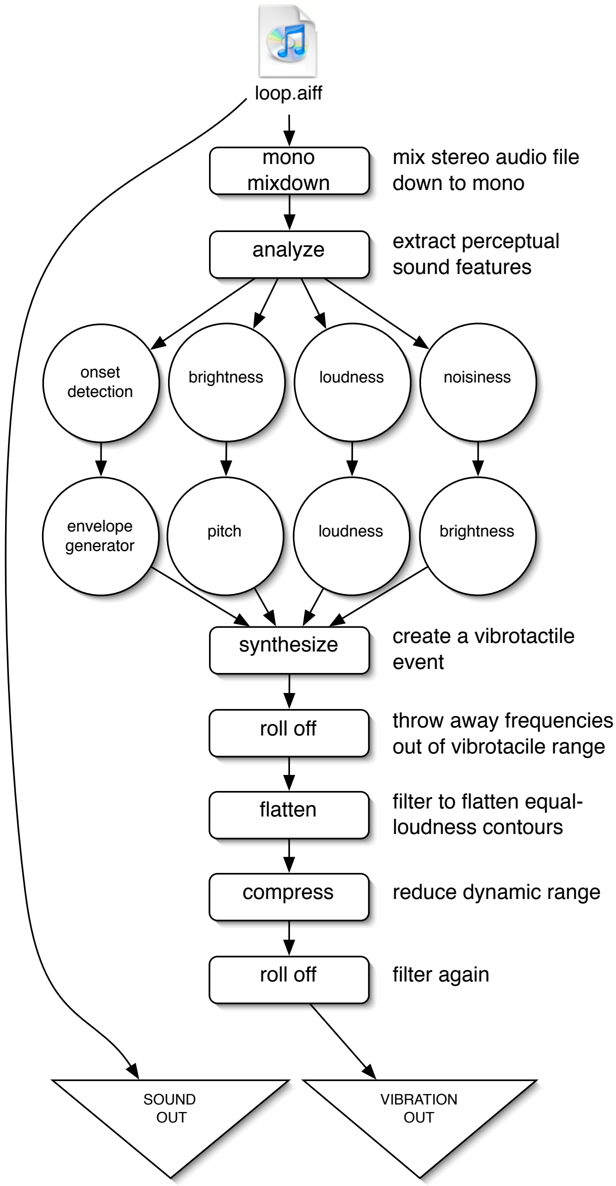

Figure 5. Flow chart illustrating audio feature extraction and vibration mapping.

5. DISCUSSION: INHERENT OR AUGMENTED?

One way feedback can be characterized is by whether it is interpreted as task-intrinsic (*inherent feedback*), or as incorporating external information (*augmented feedback*) [36]. In a sense, acoustic instruments provide vibration feedback that is tightly coupled to the musical output "for free", i.e. the same resonant system excited by the performer determines both the sound and the vibration properties of the instrument. If an accurate simulation of acoustic vibrations is desired, vibrotactile stimuli outside the acoustic vibrational range constitute noise, and so should be minimized; an understanding of what is perceived as the "inherent vibrational properties of resonating objects" must play a role in the vibration program.

With DMIs, however, the issue becomes complicated because the useful capabilities of vibrotaction extend beyond acoustic musical experience. Describing vibrotactile feedback as inherent may be taken to imply that the parameters of stimulation are within the range of acoustic vibrations, or that the vibration signal mimics the sound "accurately" according to a musician's preexisting cognitive model of musical vibrotactile feedback. Augmented feedback, on the other hand, may lie outside of the musical range and depend on other modes of human information processing, for anything from the abstraction of harmonic content to score-level cues. Ultimately, the usefulness of augmented feedback for musical applications will depend on the musician's bandwidth for feedback perception during the given task, and whether it is significantly wide to accommodate multiple modes of information processing.

The vibrotactile feedback scheme presented here, which uses high-level audio feature extraction to drive subsequent low-level signal synthesis, tends toward the inherent pole. Because the synth is continuously driven by musical signal, allows no way to define vibrotactile events independently of psychoacoustic events, and does not incorporate score-level or environmental awareness, it is a model of an inherent feedback system.

6. CONCLUSION

The human vibrotactile system involves a complex interplay between a vast number of perceptual variables, making it difficult to unravel the mechanisms involved in musical vibrotaction. This paper integrates literature from digital musical instrument design and physiology to develop a framework for musical vibrotactile feedback design. Vibrotactile digital instruments promise to be significantly more *like* their acoustic predecessors. It is not necessary, however, to limit the approach to "acoustic vibration simulation" in order to model musical vibrotactile perception in a useful way. Instead, psychophysics and stimulator design must be considered as co-dependent systems. For complex musical applications, vibrotactile perception is best represented by a generative model that extracts high-order musical invariants and resynthesizes them as tactile stimuli tailored for cutaneous display. Vibration can thus be synthesized organically to communicate relevant performance feedback.

7. ACKNOWLEDGMENTS

Thanks to Darryl Cameron, Ioana Dalca, Denis Lebel, Joseph Malloch, Mark Marshall, Alexandre Savard, Andrey Ricardo da Silva, and Steve Sinclair for their helpful contributions and suggestions.

8. REFERENCES

[1] Curtis Bahn, Tomie Hahn, and Dan Trueman. Physicality and feedback: A focus on the body in the performance of electronic music. In *Proceedings of the Multimedia Technology and Applications Conference*, 2001.

[2] Sliman J. Bensmaïa and Mark Hollins. Complex tactile waveform discrimination. *Journal of the Acoustical Society of America*, 108(3):1236–1245, September 2000.

[3] Stanley J. Bolanowski, Jr., George A. Gescheider, Ronald T. Verrillo, and Christin M. Checkosky. Four channels mediate the mechanical aspects of touch. *Journal of the Acoustical Society of America*, 84(5):1680–1694, November 1988.

[4] Stephen A. Brewster and Lorna M. Brown. Non-visual information display using tactons. In *Proceedings of CHI*, pages 787–788, 2004.

[5] P. L. Brooks and Barrie J. Frost. Evaluation of a tactile vocoder for word recognition. *Journal of the Acoustical Society of America*, 74(1):34–39, July 1983.

[6] Lorna M. Brown, Stephen A. Brewster, and Helen C. Purchase. A first investigation into the effectiveness of tactons. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 167–176, 2005.

[7] Chris Chafe. Tactile audio feedback. In *Proceedings of the International Computer Music Conference*, pages 76–79, 1993.

[8] Chris Chafe and Sile O'Modhrain. Musical muscle memory and the haptic display of performance nuance. In *Proceedings of the International Computer Music Conference*, 1996.

[9] David F. Collins, Kathryn M. Refshauge, Gabrielle Todd, and Simon C. Gandevia. Cutaneous receptors contribute to kinesthesia at the index finger, elbow, and knee. *Journal of Neurophysiology*, 94:1699–1706, May 2005.

[10] Nick Collins. Algorithmic composition methods for breakbeat science. In *Proceedings of Music Without Walls*, 2001.

[11] Mine Dogantan-Dack. The body behind the music: precedents and prospects. *Psychology of Music*, 34(4):449–464, 2006.

[12] Masashi Ezawa. Rhythm perception equipment for skin vibratory stimulation. *IEEE Engineering in Medicine and Biology Magazine*, 1988.

[13] Derry Fitzgerald, Bob Lawlor, and Eugene Coyle. Prior subspace analysis for drum transcription. In *Proceedings of the Audio Engineering Society 114th Convention*, 2003.

[14] George A. Gescheider, A. J. Capraro, Robert D. Frisina, Russell D. Hamer, and Ronald T. Verrillo. The effects of a surround on vibrotactile thresholds. *Sensory Processes*, 2(2):99–115, June 1978.

[15] John Gill. Guidelines for the design of accessible information and communication technology systems. Royal Institute of the Blind, 2004.

[16] Alan K. Goble and Mark Hollins. Vibrotactile adaptation enhances frequency discrimination. *Journal of the Acoustical Society of America*, 96(2):771–780, August 1994.

[17] Eric Gunther. Skinscape: A tool for composition in the tactile modality. Master's thesis, Massachusetts Institute of Technology, 2001.

[18] Justin A. Harris, Irina M. Harris, and Mathew E. Diamond. The topography of tactile learning in humans. *Journal of Neuroscience*, 21(3):1056–1061, February 2001.

[19] Tristan Jehan and Bernd Schoner. An audio-driven perceptually meaningful timbre synthesizer. In *Proceedings of the International Computer Music Conference*, pages 381–388, 2001.

[20] Edward F. Kelly, Debra F. McLaughlin, W. J. Ross Dunseath, Stephen Folger, Franky Jones, Jr., and H. Kenneth Hudnell. Frequency-domain measurement of vibrotactile driving responses in first-order afferent populations. *Experimental Brain Research*, 109:500–506, 1996.

[21] Lawrence Kruger, editor. *Pain and Touch*. Handbook of Perception and Cognition. Academic Press, 1996.

[22] Ki-Uk Kyung, Minseung Ahn, Dong-Soo Kwon, and Mandayam A. Srinivasan. Perceptual and biomechanical frequency response of human skin: Implication for design of tactile displays. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2005.

[23] Jack M. Loomis and Susan J. Lederman. *Handbook of human perception and performance*, volume 2, chapter 31. John Wiley and Sons, 1986.

[24] James C. Makous, Robert M. Friedman, and Charles J. Vierck, Jr. A critical band filter in touch. *Journal of Neuroscience*, 15(4):2808–2818, April 1995.

[25] Mark T. Marshall and Marcelo M. Wanderley. Vibrotactile feedback in digital musical instruments. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression*, pages 226–229, 2006.

[26] John W. Morley and Mark J. Rowe. Perceived pitch of vibrotactile stimuli: Effects of vibration amplitude, and implications for vibration frequency coding. *Journal of Physiology*, 431:403–416, 1990.

[27] Allison M. Okamura, Jack T. Dennerlein, and Robert D. Howe. Vibration feedback models for virtual environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1998.

[28] M. Sile O'Modhrain. *Playing by feel: Incorporating haptic feedback into computer-based musical instruments*. PhD thesis, Stanford University, 2000.

[29] Conor O'Sullivan and Angela Chang. Dimensional design; explorations of the auditory and haptic correlate for the mobile device. In *Proceedings of the 11th meeting of the International Conference on Auditory Display*, 2005.

[30] Patrick E. Patterson and Judd A. Katz. Design and evaluation of a sensory feedback system that provides grasping pressure in a myoelectric hand. *Journal of Rehabilitation Research and Development*, 29(1):1–8, 1992.

[31] Ivan Poupyrev, Shigeaki Maruyama, and Jun Rekimoto. Ambient touch: Designing tactile interfaces for handheld devices. In *ACM UIST*, 2002.

[32] Miller S. Puckette, Ted Apel, and David Zicarelli. Real-time audio analysis tools for Pd and MSP. In *Proceedings of the International Computer Music Conference*, pages 109–112, 1998.

[33] Miller S. Puckette and Zack Settel. Non-obvious roles for electronics in performance enhancement. In *Proceedings of the International Computer Music Conference*, pages 134–137, 1993.

[34] Joseph Rovan and Vincent Hayward. Typology of tactile sounds and their synthesis in gesture-driven computer music performance. In *Trends in Gestural Control of Music*. L'Institute de Recherche et Coordination Acoustique/Musique, 2000.

[35] Elizabeth A. Roy and Mark Hollins. A ratio code for vibrotactile pitch. *Somatosensory & Motor Research*, 15(2):134–145, 1998.

[36] Richard A. Schmidt. *Motor Control and Learning: A Behavioral Emphasis*. Human Kinetics, 4th edition, 2005.

[37] Johan Sundberg. Chest wall vibrations in singers. *Journal of Speech and Hearing Research*, 26(3):329–340, September 1983.

[38] Hong Z. Tan, Nathaniel I. Durlach, William M. Rabinowitz, and Charlotte M. Reed. Information transmission with a multi-finger tactual display. In *Proceedings of the International Sensory Aid Conference*, 1996.

[39] Jan B. F. van Erp. Guidelines for the use of vibrotactile displays in human computer interaction. In *Proceedings of EuroHaptics 2002*, 2002.

[40] Ronald T. Verrillo. Vibration sensation in humans. *Music Perception*, 9(3):281–302, 1992.

[41] Marcelo M. Wanderley and Philippe Depalle. Gestural control of sound synthesis. In *Proceedings of the IEEE*, volume 92, pages 632–644, April 2004.

[42] Marcelo M. Wanderley, Jean-Philippe Viollet, Fabrice Isart, and Xavier Redet. On the choice of transducer techonlogies for specific musical functions. In *Proceedings of the International Computer Music Conference*, 2000.