A SYSTEMATIC APPROACH TO MUSICAL VIBROTACTILE FEEDBACK

David 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 instrument designs. A design strategy for musical vibrotactile systems is developed that considers actuator placement, vibration synthesis, and a mapping from audio to vibrotactile feedback parameters utilizing neurophysiological studies. Stimulation takes both physiology and actuator technology into account to maximize sensation. Vibrations are created using perceptual sound features of the audio feedback to drive a vibration synthesizer, producing vibrotactile feedback that is tailored specifically for cutaneous display. This framework simplifies integration of vibrotactile feedback into instrument designs, avoiding the need to explicitly specify vibration stimulus parameters.

1. MOTIVATION

Touch has been vital to the development of musical skill for millennia — yet 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], but technology is now sufficiently advanced that the haptic channel may be re-engaged. It is known that acoustic vibrations are utilized for self-monitoring in acoustic performance [27, 8, 36], and that vibrotactile feedback can greatly improve touch perception during interaction without significantly adding complexity or cost to an interface [26]. This combination of circumstances should make integration of vibrotactile feedback systems into digital musical instruments (DMIs) [39] a high priority for instrument designers. Musical vibrotaction is a high-resolution, high-bandwidth, highly active perceptual system that promises nothing less than the reestablishment of embodied experience in electronic musical discourse.

Development of techniques for producing vibrations in the frequency and amplitude ranges of musical sound is important because vibrotactile feedback increases controllability of certain musical processes [24, 32]. While of acoustic vibrations is certainly valuable, it is not necessary to limit artificial vibrotactile feedback for DMIs to a simulation model — vibrations with patently un-acoustic properties can also complement and reinforce sound feedback. The only defining design requirement of a musical vibrotactile display is the mapping of musical properties to vibrotactile ones.

The aim of this research is to create a musical vibrotactile feedback system based on an open-tonehole flute-like display, called the *Tactilicious Flute Display* (TFD). [1] Whereas 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 pre-Bohm flute performance are the areas of the human body most sensitive to vibration [38]. Therefore a flutist's vibrotactile experiences present clear design criteria — flute players are in the unique position of having their highly sensitive fingerpads and lips in direct contact with their instrument's resonator.

The vibration actuators in the display are located inside the toneholes, and all are driven by the same vibration signal. This investigation does not consider tactual perception of shape, curvature, slip, textured surfaces, direction of motion, multiple body loci, or multiple simultaneous stimuli. Because the points of contact do not change, psychophysical measures of spatial acuity and the spatiotemporal structure of afferent receptive fields do not come into play. However the display is affected by tactile detection thresholds, temporal discharge patterns, pattern recognition, pitch differentiation, magnitude, and adaptation. All of these perceptual variables are relevant to music perception and correlate to vibration stimulus descriptors including frequency, amplitude, duration, waveform, and modulation.

Representing music as 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 many features including 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 notion of a flute playing breakbeat music was chosen strategically to subvert expectations and explore the plasticity of the cross-modal mapping.

[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.

2. PREVIOUS WORK

A vibrotactile display for mobile devices has been proposed where audio signal is used directly, with a small amount of signal processing to boost the tactile range [28]. For complex musical applications, however, music perception is better represented by a generative model that extracts high-order musical invariants and resynthesizes them as tactile stimuli. An older but nonetheless interesting use of the audio manipulation technique was a rhythm transmission system for deaf music students [12]. The vibration actuator was driven by a filtered version of the musical output of the player's instrument, and the players reported a higher amount of "enjoyment" and "appreciation" for music when vibration feedback heightened the sensory experience. A vibration synthesizer for instrumental feedback was programmed by [33], to be discussed in more detail below. Designed for maximum information display, tactile vocoders utilize touch for speech perception [5]. However, tactile vocoders differ from this research in that they communicate symbolic representations with pure substitution. The system described in this paper strikes a balance between relatively straightforward audio-to-vibration manipulations and complex techniques for tactile sensory substitution. Also along these lines, *Skinscape* used a music composition model to develop an approach to tactile composition [16]. Aesthetics, rather than communication, were of primary concern. Still, it is similar to the technique presented here in that low-level synthesis was used to construct vibrotactile metaphors for musical events.

3. 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 as referring to a combination of cutaneous and kinesthetic sensations relying on active exploration to perceive distal objects and events [22]. 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 [19].

Because we live in the physical world, our cognitive faculties expect certain input-output correlates. For example, we anticipate that inputting more bodily energy into a musical instrument will bring about a louder sound which contains more sound energy [41]. In order to understand how to display vibration to a performer that is perceived as meaningful and tightly coupled to the music, we must first examine the capabilities of the player's vibrotactile system. The locations and sizes of the stimuli displayed in the TFD are both unchanging, and so have not been included in these considerations. Therefore the dimensions of vibrotactile perception used in this model include *pitch*, *loudness*, *brightness*, and *envelope*.

3.1. Pitch

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 [38]. However vibrotaction is similar to audition in that, within certain ranges, frequency discrimination fits a critical band model [23]. Furthermore, frequency perception is known to be dependent on other stimulus factors such as duration [15], amplitude of skin displacement [25], and body locus [17], which are independent variables in feedback design. Several researchers have proposed that the high interdependency of frequency and amplitude suggests they be considered a single vibrotactile stimulus parameter [4]. Indeed, the periodic functions that musicians are used to experiencing through touch have both variable frequency and amplitude, so we are forced to make do with the large gaps in our understanding of vibrotactile frequency perception. And yet, certain frequency ranges give rise to distinct subjective sensations [38], implying that although vibration frequency may not be fed back preserving all the frequency content of the sound, and does not directly correlate to vibrotactile pitch, it is still a signal property that can be used for communication.

Pitch perception is such a central aspect of musical experience that it naturally tends to play a dominant role in feedback, in both auditory and vibrotactile modes. Vibrotactile "pitch" is a term that highlights the sensitivity to the rate of periodic stimuli, as it does in musical sound [34]. The neural coding mechanisms for signaling information about the frequency of vibrotactile stimuli are not well understood. Whereas the effect of stimulus amplitude and waveform on auditory pitch has a negligible effect compared to frequency, the vibrotactile sort of pitch perception is complicated by a highly dependent yet erratic dependency on amplitude, the multichannel nature of the cutaneous sense organ, and other factors [25]. A comprehensive theory of vibrotactile pitch would be very useful for feedback design, but attempts at developing such a theory have fallen short of proposing a universal and comprehensive translation scheme from auditory pitch (even accepting the inevitable drastic loss in resolution).

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*) and the number of

[2] In terms of stimulus characterization, these two modes are not distinct but rather represent 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 [37]. To further complicate matters, recent research has further revised this model, showing that cutaneous mechanoreceptors contribute to kinesthesia by responding to internal vibrations and skin stretch [9].

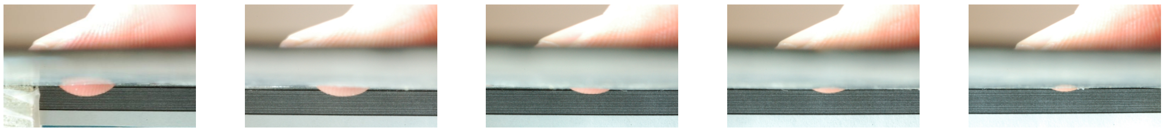

(a) 10mm (b) 9mm (c) 8mm (d) 7mm (e) 6mm

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

discriminable pitches in the vibrotactile range. Numerous studies show the response of each mechanical afferent type responds strongest to a specific frequency range [21]. When subjective magnitude is made equal, the number of discriminable pitches is still dependent on whether pitch is considered as relative or absolute [4]. [14] alleges up to nine perceivable discrete pitches should be used for symbolic information, and [33] hypothesizes between eight and ten discriminable pitches, but neither of these seem to be based on formal studies. According to [38], the *jnd* increases with frequency, lending the design suggestion that the lower frequencies should use wider pitch bands than higher ones. However [38] notes that with smaller contactor sizes (roughly the same size of the contactors in the TFD), frequency had no discernible effect on threshold or suprathreshold sensation magnitude.

3.2. Loudness

Like the transference of the word "pitch" from hearing to vibrotaction, "loudness" is used to refer to the perceptual 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 20–40Hz, the threshold of vibration perception is independent of frequency. Between 40–700Hz, however, sensitivity peaks at around 250Hz [20]. The threshold is also responsive to the presence of a non-vibrating element around the contactor, called a *surround*. This is important because the body of the TFD is in essence a surround. Loudness is also dependent on stimulus envelope, duration, temporal masking, and skin impedance [21]. It is suggested by [14] that up to four levels of vibration amplitude are easily discriminated. Vibrotactile pulses or events must occur above the threshold of perception, and sensations must be comfortable.

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, namely 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" [33]. This suggests that there is a sensation of vibrotactile "brightness" that can be targeted by controlling the amount of harmonic and periodic content.

3.4. Envelope

The envelope of a vibration sensation is affected by the dynamic responses of the actuator and the activated receptor system. Because envelope is time-dependent, adaptation and temporal masking play a significant role in perception; sustain and decay durations should take these into account.

3.5. Four channels of mechanoreception

Cutaneous sensitivity differs from hearing because there are several more channels that mediate sensory stimuli at the afferent level. The four-channel model of mechanoreception delimits the neural processing of vibrotaction into four channels associated with a putative receptor structure [3]. The four known mechanoreceptive afferent nerve fibers in glabrous skin are FAI (Fast Afferent I), FAII (Fast Afferent II), SAI (Slow Afferent I), and SAII (Slow Afferent II). These nerves activate channels which, when stimulated independently, produce "unitary" sensations. Suprathreshold sensations are the result of the combination of neural activity across the four channels (see Table 1). This framework integrates the multichannel nature of mechanical touch into a vibration synthesis approach, as described in Section 5.

4. ACTUATION

Transducer design affects controllability, and some transducers seem better suited than others to certain musical tasks [40]. The approach to actuator selection and placement presented in this section is specific to the fingertips in the context of flute performance.

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 [19]. Without such accurate measurements, the actuator system does not account for at-rest static skin pressure, damping, or skin impedance. On the other hand, miniature voice coil stimulators (previously used in [29], to cite one of many examples) are convenient because they are low-cost, high resolution, highly efficient, and easy to control. Voice coils were, after all, designed for musical applications (e.g. loudspeakers). But they are less resistant to interference from external forces exerted by the human body

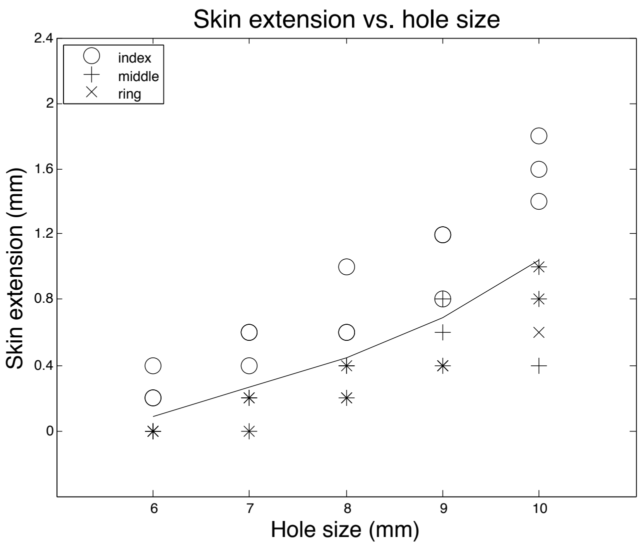

Figure 2. Hole size versus skin extension in three fingers on each of three different people. Because resolution was only 0.2mm, some points overlap. The line represents average skin extension across all fingers.

when compared to actuators with more inertia such as unbalanced motors or cylindrical contactors. Furthermore, a voice coil is not backdrivable; pressing against its surface in order to perceive vibration feedback significantly alters its 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 much improved if placed just close enough to the skin to be felt.

Installing an actuator inside a tonehole raises the question of how far to position the actuator below the surface of the hole; it is essentially a vibration stimulus offset within a surround (Figure 3). Because the deformation qualities of glabrous skin are similar to those of a fluid-filled sack [20], 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, to determine the optimal distance for the actuator to sit from the outer surface of the surround.

Three recorder players (males aged 26, 30, and 31) were asked 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. [3] Behind the metal surface, a card was placed with horizontal black lines spaced 0.2mm apart (see Figure 1). A high-resolution photograph was taken as the subject pressed his index, middle, and ring fingers down on each hole. Counting the number of lines obscured by the fingertip gives skin extension past the 1mm-thick metal surface within 0.2mm. The importance of considering hole size when placing the actuators is clearly shown in Figure 2. Variability across subjects was significant enough to indicate 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 and not bias the effectiveness of the feedback to the use of certain fingers over certain toneholes, a smaller tonehole size should be used.

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

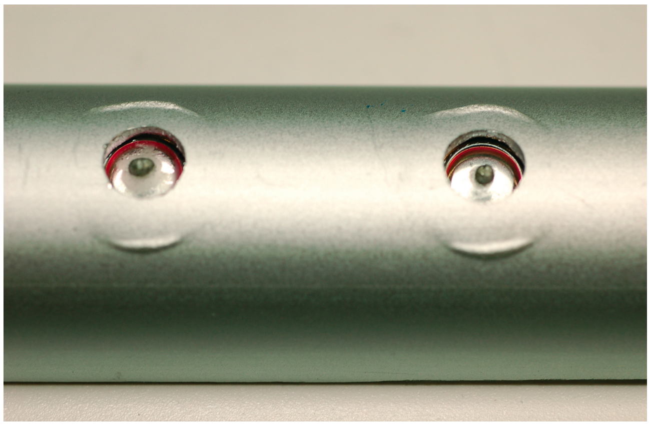

Figure 3. Two toneholes with vibrotactile actuators. The small metal dot in the center of each actuator is a load which lowers the diaphragm's resonant frequency and increases inertia [30].

5. VIBRATION PROGRAMMING

It has been claimed that the rate of receiving information through human skin sensation is about one percent of that of hearing [12]. Although the experiments that prompted that statement did not engage the entire spectrum of sensitivity in each of the multiple spatiotemporal modes of cutaneous sensation, a coarser frequency response and limited sensitivity to harmonic content do suggest that skin has a lower bandwidth than audition. Therefore an accurate model of vibration perception is the key to designing a vibrotactile display that communicates efficiently. This section presents a feedback synthesizer based on the four channel model of vibrotaction (see Section 3.5). It also describes a method for defining vibrotactile stimuli as unitary, sequential events.

To generate stimulation codes — the combined total of which the authors term the "vibration program" — signal parameters must be changed over time. The vibration synth described here uses abstractions of cutaneous perceptual parameters, but it is not the first to do so. Structured tactile messages called *tactons* use the metaphor of an iconic symbol to represent a concept [4], and experiments have been done to evaluate their effectiveness in user interfaces [6]. Another example is the VR/TX system's *tactile stimulation event* (TSE), which utilizes spa-

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
Putative receptor structure:	Pacinian	Meissner	Ruffini	Merkel
Fiber innervation density:[1]	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:[2]	250–300Hz	25–40Hz	150–400Hz	0.4–1.0Hz
Shape of frequency response function:	U-shape	Flat[3]	U-shape	Flat

[1] Human fingerpad, per square centimeter
[2] Defined as best frequencies to lower threshold of perception
[3] Notch at 30Hz

Table 1. Vibrotactile channel characteristics, adapted from [20]

tiotemporal classification criteria for encoding feedback [33].

Deviating from these approaches, the Tactilicious Flute Display maps sound features to weighted combinations of the psychophysical vibration channels that mediate tactile perception (see Section 3). Modeling these channels with a feedback synthesizer allows separate channels to be targeted with specific sound features. Transposing musical signal descriptors into the prime sensitivity ranges of the vibration channels may be an effective technique for encoding music as vibrotactile feedback. Certain ranges of frequency are often subjectively described as having distinct qualities from one another: lower frequency stimuli that are felt as a "flutter" recruit receptors in the cutaneous layer of skin and are easier to localize, whereas higher frequencies felt as a "hum" or "buzz" stimulate Pacinian corpuscules located in the deeper, subcutaneous skin layer.

5.1. P

There is evidence that the P channel integrates stimulus energy over time [2], and is probably stimulated most by the TFD's amplitude of vibration output. Its peak sensitivity occurs at about 250Hz. It has a U-shaped frequency following response similar to the equal-loudness contours in audition except that it does not flatten as intensity increases.

5.2. NPI

With the highest innervation density in the human fingerpad, it follows that the NPI channel is highly responsive to feedback targeting this location. There are 140 FAI nerve endings per cm^2, making the NPI twice as "sensitive" as the next most innervated channel, the NPIII. If innervation directly affects perceived magnitude, a fingerpad stimulator displaying vibration frequencies within the NPI range (between 2–40Hz) should be tailored to account for this heightened sensitivity; if all vibration channels are to be engaged equally loudly, average vibration intensity should be de-emphasized in this frequency range. A flattening function is not vital because the response of the NPI is naturally flat, excluding "notch" at 30Hz. [4] The TFD presents vibrotactile brightness to the NPI because this channel has been found to be particularly well suited for encoding stimulus waveform [2].

5.3. NPII

The frequency following range of the NPII lies within that of the P channel (100–500Hz), but it is particularly sensitive to lateral skin stretch. Its high vibrotactile threshold characteristic makes its role in vibrotactile coding difficult to discern [34]. However the four-channel model implies that a vibration program can engage this receptor structure with suprathreshold stimulation, so that the NPII's unitary subjective sensory quality could theoretically serve as a viable mediator of musical feedback. Because mechanical stimulation of the NPII necessarily activates the P channel well above its threshold, crosstalk is inevitible. This raises some interesting questions about what kind of musical information could be displayed to the NPII; however the actuators used in this implementation are neither accurate nor powerful enough to utilize the NPII, and so while the NPII probably does mediate TFD information, it is left for future research as to how and to what extent it does so.

5.4. NPIII

The NPIII is chiefly responsive to pressure or very low frequency periodic skin displacement. It is imaginable that a custom actuator could be used to display information to the NPIII through the use of "step functions" or multiple levels of sustained pressure, at the same time as displaying periodic stimuli. Such an actuator could also make use of relatively common signal manipulations to remain readily controllable. However the actuators used in the TFD cannot produce a high-amplitude, sustained offset function above the NPIII threshold; a transducer that combines this ability with vibration capabilities would be an excellent tool for vibrotactile feedback design.

[4] For a more accurate model it may be reasonable to include a peak filter at 30Hz to remove this nonlinearity. Short of empirical evidence, however, the notch was not accounted for in the synthesis technique presented here.

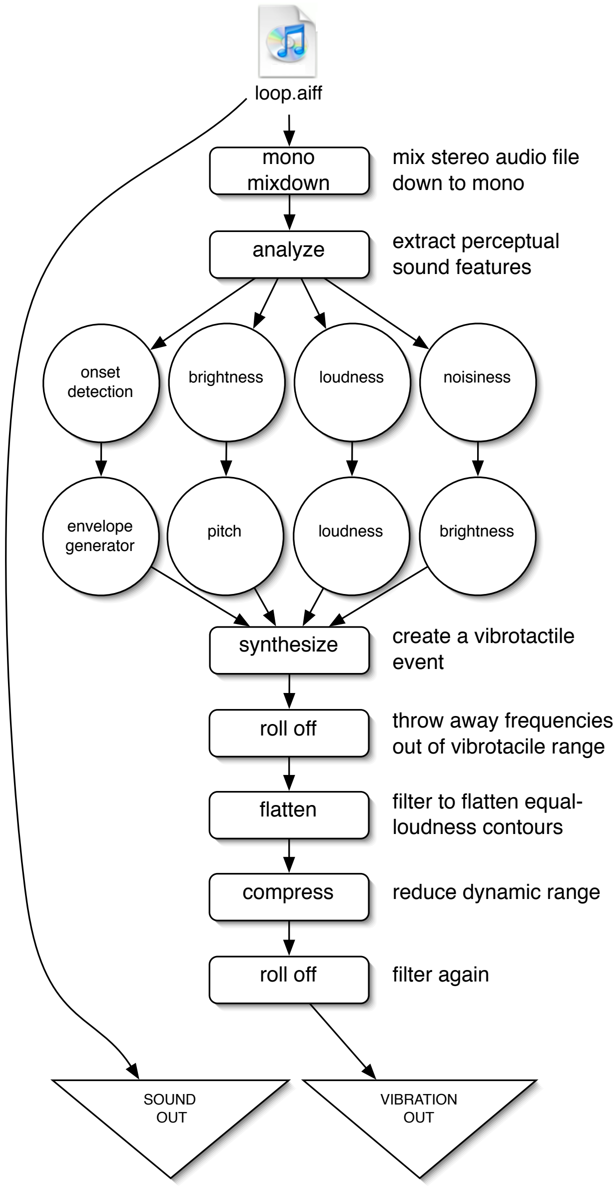

Figure 4. Flow chart illustrating audio feature extraction and vibration mapping.

6. SOFTWARE IMPLEMENTATION

A model of vibration stimulation has been programmed in Max/MSP, outlined in Figure 4. An audio analysis layer first extracts musical information from the audio signal, which is then mapped to vibrotactile perceptual parameters. Parameters of control include *loudness*, *pitch*, *brightness*, and *envelope trigger*.

External objects presented in [18] were used to extract audio features. The *noisiness*∼ object, which outputs a perceptual measure of amount of tonality, is used to control vibrotactile timbre. The vibration signal consists of a mix between sine and square wave generators, where more tonality is represented with a richer harmonic spectrum in the vibration. The *brightness*∼ object, which outputs a psychoacoustic metric that is used for drum part separation in percussion listening [13], was mapped to vibrotactile pitch so that the kick and snare drums would be represented by separate pitches. The pitch space uses a logarithmic frequency scale so that lower frequencies exhibit more variation than high ones. The output of *loudness*∼ was mapped directly to the amplitude of the vibration waveform (the reduced dynamic range of vibrotaction was compensated for in the post-processing stage). The onset detection external *bonk*∼ [31] was used to drive a simple envelope generator with an adjustable decay to create the sensation of discrete vibrotactile pulses with the above characteristics.

A filter removing frequencies out of the vibrotactile range is applied to remove unnecessary spectral content. The signal is then run 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.

7. 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*) [35]. 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 range constitute noise (whether resulting from the physical interface or the vibration signal), and so must 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 musicians 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 several 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 the musical signal, offers no way to excite vibrotactile events independently, and does not incorporate score-level or environmental awareness, it is a model of an inherent feedback system.

8. CONCLUSION

The nature of the human vibrotactile system is one of complex interplay between a vast number of perceptual variables, making it exceedingly 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, however, necessary to limit the approach to "acoustic vibration simulation" to model musical vibrotactile perception in a useful way. Instead, psychophysics and stimulator design must be considered as co-dependent systems. Vibration can then be synthesized organically, driven by high-order musical perceptual parameters, to communicate relevant musical information.

9. 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.

10. 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] S. J. Bolanowski, Jr., G. A. Gescheider, R. T. Verrillo, and C. 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] PL Brooks and BJ 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] D. F. Collins, K. M. Refshauge, G. Todd, and S. 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] John Gill. Guidelines for the design of accessible information and commuication technology systems. Royal Institute of the Blind, 2004.

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

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

[17] 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.

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

[19] 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. *Exp Brain Res*, 109:500–506, 1996.

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

[21] 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.

[22] Jack M. Loomis and Susan J. Lederman. *Handbook of human perception and performance*, volume 2, pages 1–41. John Wiley and Sons, 1986.

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

[24] Mark T. Marshall and Marcelo M. Wanderely. Vibrotactile feedback in digital musical instruments. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression (NIME06)*, pages 226–229, 2006.

[25] J. W. Morley and M. J. Rowe. Perceived pitch of vibrotactile stimuli: Effects of vibration amplitude, and implications for vibration frequency coding, 1990.

[26] 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.

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

[28] 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.

[29] P. E. Patterson and J. 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.

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

[31] Miller S. Puckette, T. 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.

[32] 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.

[33] 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.

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

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

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

[37] 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.

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

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

[40] 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 2000 International Computer Music Conference*, 2000.

[41] Todd Winkler. Making motion musical: Gesture mapping strategies for interactive computer music. In *Proceedings of the 1995 International Computer Music Conference*, pages 261–264, 1995.