Navigation of Pitch Space on a Digital Musical Instrument with Dynamic Tactile Feedback

Robert Jack
Centre for Digital Music
Queen Mary University of London
London, UK
r.h.jack@qmul.ac.uk

Tony Stockman
Centre for Digital Music
Queen Mary University of London
London, UK
t.stockman@qmul.ac.uk

Andrew McPherson
Centre for Digital Music
Queen Mary University of London
London, UK
a.mcpherson@qmul.ac.uk

ABSTRACT
We present a study investigating the impact of dynamic tactile feedback on performer navigation of a continuous pitch space on a digital musical instrument. Ten musicians performed a series of blind pitch selection and melodic tasks on a self-contained digital musical instrument with audio-frequency tactile feedback that was generated in response to their interaction. Results from the study show that tactile feedback can positively impact a performer's ability to play in tune when the instrument is hidden from sight, however with a temporal impact on performance. Furthermore, several playing techniques were observed that emerged from the performer's engagement with the tactile feedback conditions. We discuss the implications of our findings in the context of tangible interface design and non-visual interface navigation. We also discuss how our implementation suggests guidelines for future instruments and interfaces incorporating dynamic tactile feedback and present a novel tactile feedback technique that uses tactile 'beating'.

Author Keywords
Tangible User Interfaces; Digital Musical Instruments; Haptics; Tuning Assistance; Non-Visual; Beat Frequency; Active Tactile Feedback; Embedded Computing.

ACM Classification Keywords
H.5.5 [Information Interfaces and Presentation] Sound and Music Computing—Methodologies and techniques C.3 [Special-Purpose and Application-Based Systems] Real-time and embedded systems: H.5.2 [Information Interfaces and Presentation] User Interfaces—Input devices and strategies

INTRODUCTION
We present an investigation into how performers navigate the pitch space of a simple musical instrument under different dynamic tactile feedback conditions. In this paper, we use the term *dynamic* tactile feedback to indicate actively-generated vibrations which change in response to user input, this is in contrast with passive or *static* tactile feedback which remains qualitatively similar for any user input. Our aim is to evaluate how dynamic tactile feedback can assist a performer to play in tune within a continuous pitch space, and whether it does so in a way that enhances or detracts from the user experience.

Tactile feedback through an instrument's body is a natural property of most acoustic instruments. Askenfelt and Jansson [1] show that the vibrations of acoustic string instruments can readily be perceived by their performers. Other studies show this feedback to strongly affect evaluations of instrument quality [30], help the performer determine when a note is settled and stable in the instrument [5] and ensure the performer is in tune with neighbouring musicians. Fulford et al. [8] show that double-bass players can check their tuning via beat frequencies between their pitch and the ensemble, the beats transmitted through the instrument, air and floor.

Leman [15, p. 163] describes haptic feedback as "a multimodal prerequisite for musical expressiveness" as it gives the performer a more reliable sense of how gesture translates to sound at the moment of excitation. In addition to sight, it is through active exploration of haptic characteristics that a performer engages with an instrument [27]. This engagement is guided by both static factors (material, weight, arrangement of keys, strings or frets) and dynamic factors (how it responds to energy put in by the performer) [22].

From this we can see that tactile feedback is fundamental to instrumental performance, but *dynamic* feedback in particular is often lacking from digital musical instruments, due to the absence of an inherent mechanical coupling between gesture and sound [5]. This paper contributes to a growing body of work on active feedback in musical instruments by examining how different dynamic feedback conditions affect tuning accuracy and performer experience.

RELATED WORK
Before continuing, a clarification is needed of terms used to describe touch, as there is considerable variety in usage even within the HCI community. *Haptic* perception is the umbrella term for perceptions pertaining to touch. It covers two distinct categories: *kinaesthetic* perception, which is the sensation of the movement and position of one's body parts, and

DOI: http://dx.doi.org/10.1145/2839462.2839503

cutaneous or *tactile* sensation which is related to the perception of stimulation of the cutaneous receptors in the skin [29]. This is both a useful theoretical distinction and a true physiological distinction when considering individual sensory receptors [9]; however, whenever we employ our sense of touch we are always using a combination of these two sensations. Our focus in this paper is upon active tactile feedback, that is sensations felt on the surface of the skin during exploratory movement.

For Gibson [10] the function of the sense of touch depends wholly on active exploration: it is through active exploration that we are able to "isolate invariants" in the flux of incoming sensory information. Gibson demonstrates that when a subject is passively presented with a haptic stimulus, they will describe the object in subjective terms, describing the sensations on the hand for example. In contrast, when a subject is allowed to actively explore an object they will generally report object properties and object identity. Accordingly, active touch generally has higher perceptual performance due to the various human exploratory patterns tailored to each tactile property [14]. In other words: we can quickly identify an object when we are able to actively touch it rather than being touched by it. Tactile feedback itself has been shown to act as a support and source of confirmation when interacting with multi-touch interfaces [35] and both dynamic and static tactile feedback have been shown to aid navigation and orientation on touch screen devices [26].

Audio-Tactile Correspondences

There is a growing body of psycho-physical research which focuses specifically on the perception of musical parameters through the tactile modality, including studies on pitch perception [18], timbre [28], amplitude-modulated vibrotactile stimuli [34], consonance and dissonance [33]. This research generally points to fundamental correspondences of the perceptual systems of touch and audition. Eitan and Rothschild [6, p. 67] write that both touch and audition "are based on receptors that respond to pressure stimuli, transferring them (converted into electrochemical stimuli) through the nerves to the brain for processing; and both process vibrations, analyzing (albeit with different subtlety) amplitude, frequency and waveform, within perceptual ranges and JNDs (just noticeable differences) that are often roughly compatible." Comparatively little research has been conducted on active exploration of systems with dynamic tactile feedback, the focus rather being on passive sensation of vibrating stimuli.

Haptic Feedback in Digital Musical Instruments

Unlike acoustic instruments, digital musical instruments (DMIs) often decouple performance interface and sound source, eliminating a natural channel for haptic cues to reach the performer. This has been identified as a central problem in DMI design [5], resulting in considerable research into how this feedback can be reinstated. Approaches include simulating the feel of existing musical instruments using vibrotactile feedback [4], using haptic force-feedback controllers to perform digital music [2, 24], augmenting existing instruments with vibrotactile feedback [7, 25, 3], and using vibrotactile actuators for the notification of performance parameters and sound properties [13, 11].

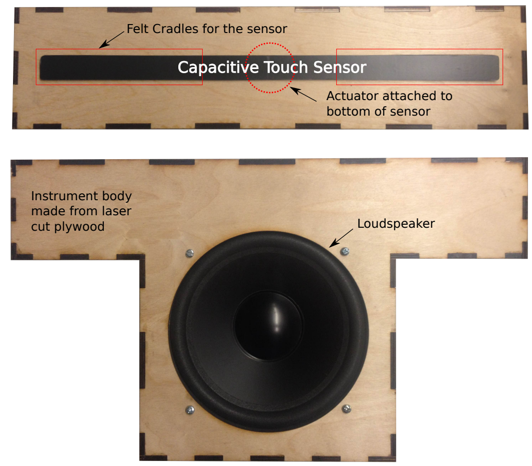

Figure 1. Diagram of the instrument's top surface and front face

Haptic Tuning Guidance

Assisting pitch control through haptic feedback was explored by O'Modhrain [24] who created an instrument similar to the theremin, but with integrated haptic feedback. From studies comparing pitch accuracy under various feedback conditions (including a spring force model) against a case with no haptic feedback at all, she concluded that the existence of force-feedback within a digital musical instrument can marginally improve the musical task of pitch selection. This conclusion has also been reached by Moss and Cunitz [23] in their work on haptic tuning guidance, which used force-feedback to push the musician's finger towards chromatic notes. Berdahl et al. [2] created virtual haptic detents using active force feedback, testing it against conditions similar to O'Modhrain's spring force. They concluded that the haptic detents model improves the pitch selection accuracy in comparison to other feedback conditions or no feedback.

The above studies all rely upon actively guiding a musician's movement through force feedback. In this paper, we present an approach to aid pitch selection based on vibrotactile feedback that is directly related to audio output of the instrument in terms of frequency. A similar paradigm has been explored by Yoo and Choi [32] with their *HapTune* device, where two vibrotactile channels are used to guide tuning on a violin. One of the actuation points is above the elbow and the other below: each are actuated to varying degrees depending on the distance from the desired tuned note. The goal of the study we present differs in desiring to test how much can be achieved through a single channel of vibrotactile feedback that is directly related to the music being made on the instrument.

INSTRUMENT DESIGN

We created a simple digital musical instrument based on the BeagleRT audio and sensor environment [20] for the BeagleBone Black single-board computer. BeagleRT allows for hard real-time, very low latency processing of audio and sensor data. The instrument (Figure 1) is a T-shaped wooden

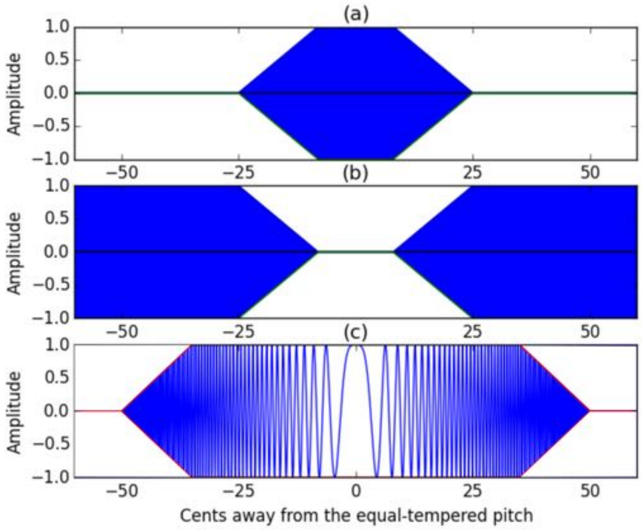

Figure 2. (a) amplitude envelope for the *vibrations when in tune* condition (b) amplitude envelope for the *vibrations when out of tune* condition (c) representation of beat frequency modulation with amplitude envelope in red for the *beat frequency* condition, mathematical description in Equation 2 and 3

box containing a 20cm speaker for audio output, a 40cm-long touch surface for pitch selection, a vibration transducer mechanically coupled to the touch surface for haptic feedback, and a BeagleBone Black for data processing. The instrument is designed to be played sitting on the lap of the performer, with the speaker resting between their legs.

Audio is sampled at 44.1kHz, and touch data is collected at 200Hz. Audio synthesis and vibrotactile control algorithms were built in Pure Data (`https://puredata.info/`) which was converted to C++ code using the Heavy compiler (`https://enzienaudio.com/`).

Sensors and Audio Mapping

The form of the instrument is inspired by the D-Box [21], but with a reduced control space: the performer only has controls for frequency and amplitude of a monophonic sound (a sawtooth wave bandpassed at 330Hz, $Q = 0.5$). The instrument has two 20×2.5cm capacitive touch sensors derived from [19] on the top surface that each sense touch position and contact area (which roughly corresponds to finger pressure). The two sensors are joined into a single continuous strip 40cm long.

A simple one-to-one mapping relates touch location to frequency and contact area to amplitude. Frequency is mapped logarithmically, with a range of two octaves plus one semitone centred an octave below middle C (C3). The centre point can be felt as a small ridge where the two touch sensors join. The distance between semitones is approximately 1.6cm.

Tactile Feedback Conditions

Tactile feedback was produced using a HiWave HIAX25C10-8/HS actuator that was affixed to the bottom of the touch sensor strip, with the magnet able to vibrate freely inside the box. The sensor strip itself was also able to vibrate freely, as it was suspended on a felt cradle attached to the top panel of the box.

The actuator is driven with an audio-frequency signal, similar to a speaker. The frequency of tactile actuation ranged from approximately 65Hz to 260Hz, following the frequency of the main audio output. The range of frequencies were chosen to specifically target the Pacinian receptors whose range of sensitivity is between approximately 50Hz and 350Hz [12]. Approximate equalisation was applied via a series of four bandpass filters to the tactile feedback signal to ensure equal intensity across the frequency range.

We considered four tactile feedback conditions, generating a signal for the actuator based on touch location. The purpose of the feedback was to communicate to the performer when they were on or near a diatonic pitch of the C major scale. The conditions were:

1. No feedback

This was the control condition. No signal was sent to the actuator, but vibrations from the speaker can still be felt through the case and sensor, albeit at a substantially lower intensity than the other conditions.

2. Vibrations when in tune

In this condition, the vibration actuator engages when the pitch is near an (equal-tempered) diatonic note. As seen in Figure 2(a), the amplitude of the vibrations is 0 when the note is more than 25 cents out of tune, ramping up to maximum when the performer is within 8 cents of the correct pitch.

The frequency of the actuator signal is that of the correct pitch (i.e. the closest equal-tempered note), which will be close but generally not identical to the frequency played through the speaker. The waveform was a sine wave. Compared to other work with fixed-frequency actuation signals [32], we chose to match the frequency to the nearest note to reinforce the relationship between tactile and audio output. This also led to occasional gentle beat frequencies due to the difference in frequency of actuator and speaker output. The beating effect inspired the deliberately exaggerated beats of Condition 4.

3. Vibrations when out of tune

As shown in Figure 2(b), this feedback condition applies a similar technique to Condition 2, but reversed: the actuator is at maximum amplitude when the performer plays a note more than 25 cents from the nearest diatonic pitch, reaching zero amplitude when the pitch is within 8 cents of the target. The frequency of the actuator signal is identical to the frequency played through the speaker and the waveform a sine wave.

4. Beat frequency

This condition uses the difference between the target note and the played note to create haptic ‘beating’ [31], taking inspiration from the accounts of double-bassists mentioned above [8]. The beats are generated by interference patterns between two closely-tuned oscillators, according to the following trigonometric identity:

$$\sin(u) + \sin(v) = 2\sin\left(\frac{u+v}{2}\right)\cos\left(\frac{u-v}{2}\right) \qquad (1)$$

Because the audio frequency differences involved are small, especially in the instrument's lower octave, the natural beat frequencies will be quite slow. To exaggerate the effect, we applied a warping to the reference frequency as follows:

Let f_{spk} be the frequency of the speaker output, f_{tuned} be the frequency of the closest diatonic pitch and f_{beat} be the desired beat frequency at 50 cents away from the tuned pitch. Then f_{ref} is the reference frequency used in the calculation of the warped beat frequency as follows:

$$f_{ref} = 2 f_{beat} \mid f_{spk} - f_{tuned} \mid - \mid f_{spk} - f_{tuned} \mid \qquad (2)$$

$$\text{warped beat frequency} = (f_{spk} \pm f_{ref}) - f_{tuned} + 1 \qquad (3)$$

We chose to create a beat frequency of 60Hz when the performer was 50 cents away from the tuned note, reducing to beating of 1Hz when the performer played perfectly in tune. The amplitude envelope applied to the actuator signal is shown in Figure 2(c). Though on acoustic instruments, the beating disappears entirely when two notes are precisely in tune, to implement this behaviour would create a subtle problem: depending on the relative phase of the two signals, the in-tune condition could be either a maximum or a minimum in amplitude, which could be confusing. Instead we chose to limit the minimum beat frequency to 1Hz so that performers would be able to feel a slow pulsing when in tune.

Referring back to Equation 1, we see the beating is a form of amplitude modulation. Here the frequency of the modulator is the beat frequency, the primary frequency that the performer perceives, which ranges from 1Hz to 60Hz. The frequency of the carrier is also variable (the mean of the reference and speaker frequencies) and provides additional haptic information. Since the beat frequency alone does not distinguish whether a note is sharp or flat, we hoped this additional carrier frequency information would help the performer distinguish between these cases.

STUDY

Experiment Design

Method

Our experimental method was derived from Berdahl et al.'s study on pitch selection with force-feedback haptic assistance [2]. To assess the impact of tactile feedback on pitch accuracy, we asked participants to perform a series of pitch selection tasks and play melodies using each tactile feedback condition. We recruited 10 participants (3 female) from our host university. All ten identified themselves as musicians; eight participants had 10+ years of instrumental experience, and the other two participants had 3+ years each. The study lasted around one hour and fifteen minutes per participant.

Setup

Our experiment setup consisted of the musical instrument and a wooden panel which hid the instrument from sight while the participant completed the task as can be seen in Figure 3. The

Figure 3. Participant performing on the instrument. The instrument is hidden from sight by a wooden panel

Figure 4. The second simple melody that participants had to perform. Transposed up one octave for ease of reading.

instrument produced sound through the speaker, but the participants wore noise-cancelling headphones during the experiment, through which they could hear the same audio signal played by the speaker. This was to avoid any residual sounds from the tactile actuator influencing their performance.

The study procedure was as follows:

1. Each participant was introduced to the instrument without tactile feedback. The procedure of the experiment was explained, and they were given guidelines about how to play the instrument: only one hand was to be used on the sensor although they were free to use any finger and change fingers as they saw fit. Participants were also encouraged to rest their other hand on the body of the instrument to enhance the perception of tactile feedback.

2. For each feedback condition, the participant was first given a note-finding exercise. Six single tones were played in isolation, and the participant had to match the pitch on the instrument. This task served two functions: to improve the participant's familiarity with the instrument under a given feedback condition, and to provide us with a metric of their pitch-finding skill.

3. The participant was then asked to play scales and melodies. The scale was a 2-octave C major scale, first ascending and then descending. They could rehearse this until comfortable and it was then recorded. For melodies, they would first listen to the excerpt and be provided with a score (e.g. Figure 4). Participants were advised to pick an appropriate tempo that allowed them to maintain accurate tuning with clear stable notes as well as a steady beat. After practicing until they felt comfortable, they recorded the excerpt three times and chose their preferred take after doing so.

4. Every participant performed with all four feedback conditions, presented in a counterbalanced random order between participants to minimise learning effects. The scales and melodies were however presented in the same order.

5. Upon completing the set of pitch selection tasks for each feedback condition, the participant completed a questionnaire about that condition, which included questions about their perceived tuning accuracy, their comprehension of the feedback and the mental effort it required. At the end of the study, the participants filled out an exit survey asking their favourite feedback condition amongst other questions.

Data collection

Alongside the questionnaire results we also collected performance data from the instrument: the speaker signal, the actuator signal, touch position on the sensor, contact area on the sensor, the computed closest equal tempered note and the computed beat frequency for the beat frequency feedback condition. The two audio streams were sampled at 44.1kHz while the sensor streams were normalised to a range of $[-1, 1]$ and up-sampled to a rate of 22.05kHz.

RESULTS

Mean absolute pitch error

Based on the position reading of the touch sensor we compared the pitch of the actual note played by the performer against the target equal-tempered pitch, which was measured in semitones (logarithmic with respect to frequency). Performances were manually annotated, first segmenting into notes, then identifying the region within each note when the performer settled on a stable pitch. A mean pitch value was then calculated for each stable section by averaging the pitch across the stable region. The absolute difference between this value and the target pitch values for each melody was calculated. This gave a measure of mean absolute pitch error for every note performed under each tactile feedback condition.

We used accuracy on the single-note pitch matching test as a screen for reliability of the rest of the participant's data. Figure 5 shows that there was at times large differences between participants' performance. We excluded participants who achieved less than 80% accuracy in the pitch-matching task. One participant was excluded (Participant 7): as can be seen in Figure 5 their performance across most of the conditions was markedly worse than other participants.

Overall, averaging all participants and melodies, we found that in relation to the 'audio only' condition, the differences in absolute mean pitch error are marginally significant ($p < 0.08$) in the case of the 'vibrations when in tune' condition and significant ($p < 0.05$) in the case of the 'beat frequency' condition (Figure 6(a)). For 81% of the melody-participant pairs the mean pitch error was smaller for the 'vibrations when in tune' than the 'audio only' and 'vibrations when out of tune' condition. For 76% of the melody-participant pairs the 'beat frequency' condition yielded a smaller mean pitch error than the 'vibrations when out of tune' and 'audio only' conditions.

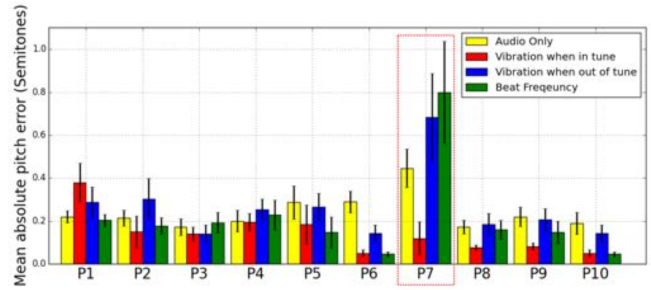

Figure 5. Mean absolute pitch error for each feedback condition for each participant for melody 1. Error bars represent standard error.

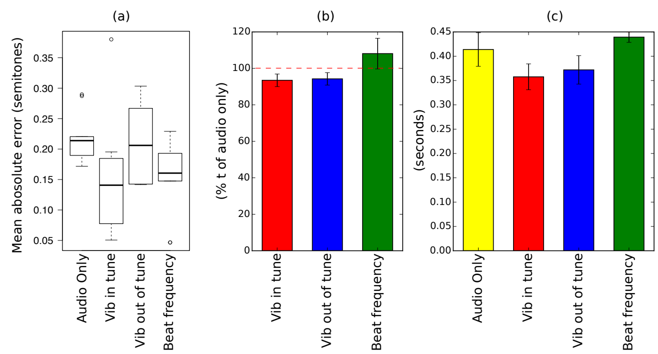

Figure 6. (a) Mean absolute pitch error for each feedback condition across 9 out of the 10 participants for both melodies. (b) Percentage of mean tempo in relation to the audio only condition. (c) The mean searching gesture duration (time before reaching a stable note). Error bars represent standard error.

Timing impact

The average difference in timing for all participants and melodies was calculated by measuring the duration of each melodic passage minus the first and last note. Figure 6(b) shows that the 'vibrations when in tune' and 'vibrations when out of tune' were on average performed faster by participants than the 'audio only' condition (83% and 73% of the melody-participant pairs respectively). The 'beat frequency' condition however was generally slower than the 'audio only' condition (73% of melody-participant pairs).

To examine temporal performance in more detail we then measured the duration of stable notes and the duration of the gestures used to reach them, the latter of which can be seen in Figure 6(c). In the 'beat frequency' condition there is a smaller mean stable note duration than in any of the other conditions even though the time taken to perform each of the melodies is generally longer: the *searching* gestures before a stable note were the longest for this condition.

Learning effects

The effect of learning was examined by comparing the pitch accuracy from the first and last feedback conditions each participant encountered. For 5 of 9 participants, the feedback condition that they performed with *first* was their most accurate, suggesting that learning effects due to the order of conditions are minimal in this study. The counterbalanced random

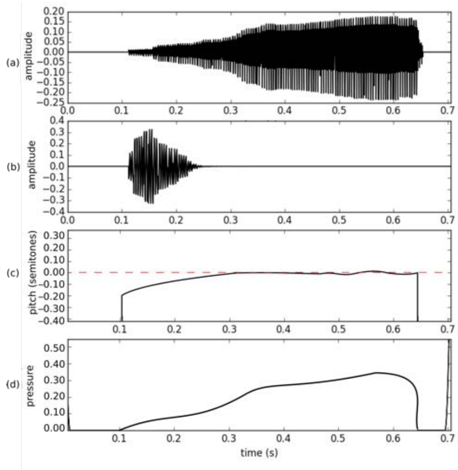

Figure 7. Note finding playing technique for 'vibration when out of tune' feedback condition. Shows the relationship of finger position and pressure to tactile and audio feedback. (a) audio channel, (b) actuation channel, (c) pitch from the sensor, (d) pressure (touch size from the sensor).

order of conditions and the opportunity for practice on each melody also help reduce bias from learning effects.

Playing Technique Observations

Aside from the instruction to play with one hand, the participants were not instructed to use a particular playing technique. We observed a wide variety of playing techniques that varied both among participants and feedback conditions. Techniques included the following: sliding from note to note on a single finger; various fingering positions on the sensor using different fingers with sliding used for small corrections or jumps of a large interval; detached playing with one finger where the performer lifts the finger between each note. These examples show how each feedback condition engendered its own playing techniques that participants discovered through experimentation with the instrument. This finding is similar to that of Marshall and Wanderley [16, 17] during an equally open-ended study which evaluated sensor choice for parameter modulation in digital musical instruments.

Note finding

An additional performance technique that was observed under certain tactile feedback conditions was a method for searching out the correct note before playing (Figure 7). In this case the performer would lightly touch the sensor and adjust their position in response to the tactile feedback before applying more pressure and increasing the amplitude of the audio output to its full level. It is worth noting that the finger movement on the sensor in Figure 7 happens over a distance of ~3mm and lasts ~200ms. We did not mention the amplitude mapping to the participants, again this technique appeared as an intuitive response to the feedback condition.

Participant tactile feedback preferences

A summary of the results from the surveys is presented in Table 1. In general, the 'vibrations when in tune' condition scored best on most metrics, though all tactile feedback

	Feedback Condition			
Survey Question	*Audio*	*In Tune*	*Out of Tune*	*Beating*
How successfully did you play in tune? (1: Very badly – 10: Perfectly)	4.2	**5.82**	4.78	5.1
How hard was it to play in tune? (1: Very easy – 10: Very Hard)	6.2	**5.2**	6.8	6.0
Were you able to maintain your desired tempo? (1: Not at all – 10: Completely)	**7.3**	6.1	5.0	4.8
How mentally demanding the tactile feedback? (1: Not at all – 10: Very)	N/A	**5.9**	6.44	6.7
How much did the tactile feedback assist tuning? (1: Not at all – 10: Very)	N/A	**6.9**	5.5	6.5
Which was your preferred condition? (Number of participants)	0	4	3	3

Table 1. Summary of the responses from the survey conducted at the end of each feedback condition. The low and high limits of the 10 point metric are listed under the questions in the left hand column.

conditions appear to reduce participants' reported ability to maintain their desired tempo. For 7 of 9 participants, the preferred feedback condition also yielded their most accurate performance. There was no clear favourite feedback condition across all participants. Participants were asked to explain the reason for their preference. For participants who preferred 'vibrations when in tune', they stated that they liked the affirmative nature of the feedback, that they were provided with a clear signal of when they were playing in tune and could play detached notes and know immediately whether they were playing the correct note.

Participants who preferred 'vibrations when out of tune' stated that they preferred being 'buzzed' when out of tune as this reflected the way they would normally think about pitch selection on their instrument: when they are playing in an ensemble and are in tune they don't think about their tuning; it is rather when they are out of tune with the ensemble that they become aware of their tuning and know they must correct it.

Reasons for selecting the 'beat frequency' tactile feedback as the preferred method had to do with the variety and amount of information it provided. Participants who preferred it stated that this condition had the most potential for long-lasting engagement, as it allowed micro adjustments to tuning to be performed and helped maintain a focus on the tactile feedback. However, it was acknowledged that this condition was the most difficult to play melodies with, and that it would take a longer time than was available in the study to take full advantage of it. This was confirmed by the results of the individual condition questionnaires where ‘beat frequency’ was consistently rated as the most mentally demanding.

Discussion

The study examined how dynamic tactile feedback altered the navigation of pitch space, focusing on two main questions: first, how does dynamic tactile feedback impact pitch selection accuracy when compared to an audio-only condition; and second, what impact does dynamic tactile feedback have on the performer’s actions and experience? Considering the first question, we found that the ‘vibration when in tune’ and ‘beat frequency’ conditions both provide improvements to tuning accuracy although with different temporal costs.

The ‘vibrations when out of tune’ condition, although structurally similar to ‘vibrations when in tune’, does not improve tuning performance across our participants, suggesting that removing tactile stimulation for task confirmation rather than adding could impact feedback effectiveness. Although this condition did not improve tuning accuracy in comparison to the ‘audio only’ condition, participant preference for a particular *polarity* of tactile feedback highlights an important consideration when designing interfaces with tactile feedback: should the user be informed of successful or unsuccessful task completion? From this small sample we cannot conclude which is generally preferable, but it seems that such a guideline would need to be informed by both subjective user preference and objective measurement of performance.

The improved accuracy in combination with the negative timing impact and performer survey results for the ‘beat frequency’ condition suggest that this tactile signal is either too nuanced or unfolds too quickly to be useful in a timing critical situation like the performance of a musical instrument, at least with the limited learning times that we had in this study. Nevertheless, the fact that participants found the condition engaging as well as generally comprehensible means that this technique could possibly be fruitfully employed in a tactile interface in which fine accuracy is required but there are not the same time-pressures as a musical performance.

The reported reduced ability to maintain tempo when using vibrotactile feedback could be explained by the sensorimotor latency required to process haptic feedback and then act on it. It could be posited that expert musicians act in a *feed-forward* mode when performing, where planning and execution of musical passages happens at too high a rate to process the note-by-note responses from their instruments. Novice musicians on the other hand could be described as taking a *feed-back* approach where the sound and feel of each note are attended to. In the case of this study there was perhaps simply not enough time for the participants to advance to the feed-forward mode, to a point where the vibrotactile feedback was supporting their expectations rather than determining their note-by-note performance on the instrument.

For the second question, we saw several emergent performance techniques that were directly influenced by the tactile feedback. The technique of lightly searching for a pitch before committing with greater pressure suggests that tactile feedback could be engaged immediately on contact with the instrument, with the sound only starting once the performer applies further pressure. This would allow a performer to confidently find the right pitch before producing a sound.

The high variability in results amongst performers precludes a statistically definitive answer to which condition is best; however, the alignment of accuracy and survey data yields support for the simpler cases, at least for situations where the user only has a short time to work with an interface.

CONCLUSIONS

We presented a study examining the impact of dynamic tactile feedback on the navigation of pitch space on a self-contained digital musical instrument. Our results show an improvement in pitch selection accuracy with certain types of vibrotactile feedback for our participants: in general accuracy was improved when participants performed with the ‘vibrations when in tune’ condition in comparison to the ‘audio only’ condition. An improvement in accuracy was also achieved in the ‘beat frequency’ condition however with a negative impact on timing and user response. Testing more subjects along with a more complete statistical analysis will help us solidify our conclusions which are limited by our sample size.

We also observed a set of emergent gestures linked to the type of feedback, suggesting that haptic information has a strong influence on how a performer conceptualises a new instrument. Our pitch accuracy results are necessarily limited to the particular instrument and participants in this study, but they reinforce findings from previous studies, especially [2]. Notably, where previous studies used force feedback to push the performer to the right pitch, our interface requires an active correction by the performer. It is interesting that both methods are successful in improving accuracy. Further investigations into the generalisability of these findings may have significant implications for the understanding of multimodal and cross-modal interaction design, and for the design of new musical instruments.

ACKNOWLEDGMENTS

Thanks to the members of the Augmented Instruments Laboratory and to Christian Heinrichs, Adib Mehrabi and Giulio Moro for their expert assistance. This work was supported by the EPSRC Centre for Doctoral Training in Media and Arts Technology, Queen Mary University of London.

REFERENCES

1. Anders Askenfelt and EV Jansson. 1992. On vibration sensation and finger touch in stringed instrument playing. *Music Perception* 9, 3 (1992), 311–349.

2. Edgar Berdahl, Günter Niemeyer, and Julius O Smith. 2009. Using Haptics to Assist Performers in Making Gestures to a Musical Instrument. In *Proc. International Conference on New Interfaces for Musical Expression*. Pittsburgh, PA, United States.

3. DM Birnbaum and MM Wanderley. 2007. A systematic approach to musical vibrotactile feedback. *Proc. International Computer Music Conference* (2007).

4. Nicolas Castagné, Claude Cadoz, Jean-Loup Florens, and Annie Luciani. 2010. Haptics in computer music: a paradigm shift. *arXiv preprint arXiv:1005.3182* (2010).
5. Chris Chafe. 1993. Tactile audio feedback. *Proc. International Computer Music Conference* (1993), 76–79.
6. Zohar Eitan and Inbar Rothschild. 2010. How music touches: Musical parameters and listeners' audiotactile metaphorical mappings. *Psychology of Music* (2010).
7. Federico Fontana, J Hanna, Federico Avanzini, Francesco Zanini, Valerio Zanini, and Musica Cesare. 2014. Perception of Interactive Vibrotactile Cues on the Acoustic Grand and Upright Piano. *Proc. International Computer Music Conference - Sound and Music Computing 2014* September (2014), 948–953.
8. Robert Fulford, Jane Ginsborg, and Juliet Goldbart. 2011. Learning not to listen: the experiences of musicians with hearing impairments. *Music Education Research* 13, 4 (Dec. 2011), 447–464.
9. Alberto Gallace and Charles Spence. 2014. *In touch with the future: The sense of touch from cognitive neuroscience to virtual reality*. Oxford University Press.
10. James J Gibson. 1962. Observations on active touch. *Psychological review* 69, 6 (1962), 477.
11. Marcello Giordano and MM Wanderley. 2013. Perceptual and Technological Issues in the Design of Vibrotactile-Augmented Interfaces for Music Technology and Media. *Haptic and Audio Interaction Design* (2013).
12. Michael J Griffin. 2012. *Handbook of human vibration*. Academic press.
13. Lauren Hayes. 2011. Vibrotactile Feedback-Assisted Performance. In *Proc. International Conference on New Interfaces for Musical Expression*. Oslo, Norway, 72–75.
14. Susan J Lederman and Roberta L Klatzky. 1987. Hand movements: A window into haptic object recognition. *Cognitive psychology* 19, 3 (1987), 342–368.
15. Marc Leman. 2008. *Embodied music cognition and mediation technology*. MIT Press.
16. Mark T Marshall, Max Hartshorn, Marcelo M Wanderley, and Daniel J Levitin. 2009. Sensor choice for parameter modulations in digital musical instruments: Empirical evidence from pitch modulation. *Journal of New Music Research* 38, 3 (2009), 241–253.
17. Mark T Marshall and Marcelo M Wanderley. 2006. Evaluation of sensors as input devices for computer music interfaces. In *Computer Music Modeling and Retrieval*. Springer, 130–139.
18. Saúl Maté-Cid. 2013. Vibrotactile Perception of Musical Pitch. PhD thesis, Unviersity of Liverpool. (2013).
19. Andrew McPherson. 2012. TouchKeys: Capacitive multi-touch sensing on a physical keyboard. In *Proc. International Conference on New Interfaces for Musical Expression*.
20. Andrew McPherson and Victor Zappi. 2015a. An Environment for Submillisecond-Latency Audio and Sensor Processing on BeagleBone Black. In *Audio Engineering Society Convention 138*.
21. Andrew P McPherson and Victor Zappi. 2015b. Exposing the Scaffolding of Digital Instruments with Hardware-Software Feedback Loops. *Proc. International Conference on New Interfaces for Musical Expression* (2015).
22. Eduardo Reck Miranda and Marcelo M Wanderley. 2006. *New digital musical instruments: control and interaction beyond the keyboard*. Vol. 21. AR Editions, Inc.
23. William Moss and Bryan Cunitz. 2005. Haptic theremin: Developing a haptic musical controller using the sensable phantom omni. *Proc. International Computer Music Conference* (2005), 275–277.
24. Maura Sile O'Modhrain. 2001. Playing by feel: incorporating haptic feedback into computer-based musical instruments. PhD thesis, Stanford University. (2001).
25. Stefano Papetti, Sébastien Schiesser, and Martin Fröhlich. 2015. Multi-point vibrotactile feedback for an expressive musical interface. In *Proc. International Conference on New Interfaces for Musical Expression*.
26. Martin Pielot, Anastasia Kazakova, Tobias Hesselmann, Wilko Heuten, and Susanne Boll. 2012. PocketMenu: non-visual menus for touch screen devices. In *Proc. MobileHCI*. ACM, 327–330.
27. Pedro Rebelo. 2006. Haptic sensation and instrumental transgression. *Contemporary Music Review* 25, 1-2 (2006), 27–35.
28. Frank a Russo, Paolo Ammirante, and Deborah I Fels. 2012. Vibrotactile discrimination of musical timbre. *Journal of experimental psychology. Human perception and performance* 38, 4 (Aug. 2012), 822–6.
29. Salvador Soto-Faraco and Gustavo Deco. 2009. Multisensory contributions to the perception of vibrotactile events. *Behavioural Brain Research* 196, 2 (2009), 145–154.
30. Indiana Wollman, Claudia Fritz, Jacques Poitevineau, and Stephen McAdams. 2014. Investigating the Role of Auditory and Tactile Modalities in Violin Quality Evaluation. *PLoS ONE* 9, 12 (2014), e112552.
31. Shiyan Yang, Kathryn Tippey, and Thomas K Ferris. 2014. Exploring the Emergent Perception of Haptic Beats from Paired Vibrotactile Presentation. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 1716–1720.

32. Yongjae Yoo and Seungmoon Choi. 2014. An Initial Study on Pitch Correction Guidance for String Instruments Using Haptic Feedback. In *Haptics: Neuroscience, Devices, Modeling, and Applications*. Springer, 241–247.

33. Yongjae Yoo, Inwook Hwang, and Seungmoon Choi. 2014. Consonance of vibrotactile chords. *IEEE Transactions on Haptics* 7, 1 (March 2014), 3–13.

34. Gareth W Young, David Murphy, and Jeffrey Weeter. 2015. Auditory Discrimination of Pure and Complex Waveforms Combined With Vibrotactile Feedback. In *Proc. International Conference on New Interfaces for Musical Expression*.

35. Simone Zimmermann, Sonja Rümelin, and Andreas Butz. 2014. I Feel it in my Fingers: Haptic Guidance on Touch Surfaces. *Proc. International Conference on Tangible, Embedded and Embodied Interaction* (2014), 9–12.