

A Self-Sensing Haptic Actuator for Tactile Interaction with Physical Modelling Synthesis

Matthew Davison
Dyson School of Design
Engineering
Imperial College London,
South Kensington
London SW7 2DB, UK
m.davison23@imperial.ac.uk

Craig J. Webb
Department of Industrial
Engineering
University of Bologna
Viale del Risorgimento 2,
40136 Bologna, Italy
craigjonathan.webb@unibo.it

Michele Ducceschi
Department of Industrial
Engineering
University of Bologna
Viale del Risorgimento 2,
40136 Bologna, Italy
michele.ducceschi@unibo.it

Andrew P. McPherson
Dyson School of Design
Engineering
Imperial College London,
South Kensington
London SW7 2DB, UK
andrew.mcpherson@imperial.ac.uk

ABSTRACT

The use of transducers to excite physical modelling synthesisers with real-world audio signals is a well-established practice within the digital musical instrument design community, yet it is normally presented as a unidirectional process – energy is transferred into the system from human to instrument. In this paper, a novel approach to tactile interaction with physical modelling synthesis is presented, through the use of a self-sensing vibrotactile transducer. This enables simultaneous collocated sensing and haptic actuation with a single moving coil transducer. A current drive amplifier is used for haptic actuation, using signals derived from the physical modelling synthesiser. The varying impedance of the transducer (due to changes in the mechanical damping) enables the sensing of force applied upon the device whilst also acting as a pickup to excite the physical model, all with simultaneous haptic actuation. A digital filter equivalent of the transducer's impedance is used to prevent feedback in the system, allowing simultaneous excitation and haptic actuation without self-oscillation.

Author Keywords

vibrotactile, self-sensing, haptics, tactile, physical modelling, collocation

CCS Concepts

•Human-centered computing → Haptic devices; •Applied computing → Sound and music computing;

1. INTRODUCTION

Tactile interaction with acoustic instruments involves a bidirectional process, whereby the user transfers energy to the instrument – through actions such as hitting, plucking, or bowing – and the instrument simultaneously transfers energy back to the user. This occurs through mechanisms such as mechanical resonances in the instrument's body, the vibration of strings under the fingertip, and resistance to force applied by the user [29]. Such phenomena have been shown to have a noticeable effect on a musician's experience of the instrument [18]. Standard commercial controllers for digital synthesis, on the other hand, lack haptic cues beyond those afforded by their passive structural components, yielding a limited user experience [32, 9] and only enabling a unidirectional transfer of energy.

The addition of haptic feedback to such control interfaces is a possible method of creating more engaging tactile interactions [30]. In this paper, such a system is introduced, using a novel configuration of a single voice coil transducer as a self-sensing actuator. Collocation is achieved through the use of a single transducer, creating deep virtual-physical integration at a single contact point. The voice coil transducer is driven by a current drive amplifier, using a dedicated audio rate simulated force output from a physical modelling synthesiser. The voltage across the transducer is returned to the system as an input signal. Filtering of this signal provides both an audio excitation signal to drive the input of the physical model along with a measurement of the force applied upon the transducer. Both measurements can be achieved whilst simultaneously driving the transducer with vibrotactile haptic signals, enabling a bidirectional haptic coupling between the musician and the physical modelling system.

Unlike more abstract forms of sound synthesis, such as subtractive or modulation synthesis [37, 22], in physical modelling synthesis sound is produced by simulating the behaviour of physical objects such as strings and plates and, hence, can naturally incorporate control gestures such as striking, picking, sliding and bowing [28, 35]. Furthermore, the virtual vibration of the simulated instrument can be converted into control signals with a natural physical inter-



Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

NIME'24, 4–6 September, Utrecht, The Netherlands.

pretation: the net virtual force acting on an infinitesimal of a simulated string can be directly converted into an electrical signal, driving a haptic actuator [8]. This paper seeks to explore how the tactile vibrations of a physical model can be experienced whilst simultaneously providing a rich manner of interacting with the physical model in the tactile domain.

2. RELATED WORK

While predominantly focused upon force feedback, rather than vibrotactile haptic feedback, research into haptic interaction with physical synthesis is a well-established practice. Early experiments can be traced back to the works by the ACROE group using haptic keyboards within the CORDIS-ANIMA mass-spring framework [8, 7]. Within the same framework, a haptic bow design was presented soon after [17]. The vBow [28], a bowed-instrument controller with feedback, was designed to work within the Synthesis Toolkit (STK) [11], and successive controllers included the use of tablets to control digital waveguide models [35]. In spite of such a large body of work, this research did not produce interfaces widely used by the public. This was partly due to the relatively high cost of such early devices, as well as to the proprietary nature of many physical modelling algorithms. This led researchers to explore the possibility of using open-source hardware and software at affordable prices, leading to implementations such as the FireFader [2] and the Plank [39].

2.1 Resonant Feedback Instruments

Closely linked to haptic interaction with physical models is the research area of hybrid resonant instruments [6, 15]. Such systems utilise the acoustic and mechanical properties of physical objects, instruments, and assemblages to create the timbre of the musical output. Actuators, sensors, and a digital processing system are used for both feedback control and the addition of digital audio effects into the feedback loop. The physicality of such systems provides inherent tactile interaction; the musician is able to gain control over the resonant system by varying touch position, touch pressure and grip shape - affecting the resonances and coupling within the system - as well as adding excitation energy into the system. Furthermore, along with direct tactile control of the system, the musician is also able to feel the system's resonances - creating a direct and bidirectional coupling between the musician and the instrument assemblage.

2.2 Audio Excitation of Physical Models

Sensors such as piezoelectric pickups and accelerometers have previously been employed alongside resonant physical modelling synthesisers in DMI design, to enable audio signal excitation of the model [27, 26, 34, 10]. While many designs simply utilise the sensor data as a direct audio input, some also extract further information from the sensor readings, using a single sensor for both audio excitation and parameter control. Mice and McPherson [26] make use of analogue accelerometers, sensing mechanical vibrations to excite a resonant synthesis algorithm. The DC component of the signal (corresponding to the angle of the apparatus) modulates the feedback parameter of the synthesis algorithm. In a similar approach, Schmalfuß et. al. [34] use frequencies below 100Hz to measure change in force on a piezo disc, with frequencies above 100Hz used for excitation of the digital resonator model. The concept of utilising a single sensor signal for both audio excitation and param-

eter control is explored in this paper, whilst also utilising self-sensing techniques to enable simultaneous haptic actuation.

2.3 Self-Sensing Transducers

Self-sensing systems utilise a single transducer for simultaneous sensing and actuation [38]. Such systems have been previously examined in other fields, including active vibration control, noise reduction, and light sensing/emitting [38, 31, 20]. Self-sensing systems provide the benefits of a more compact system and true collocation when compared to systems with separate sensors and actuators. Collocation is particularly important for closed-loop control systems [3]. In the field of digital musical instrument design, Sierra [36] uses a dual voice coil loudspeaker as a self-sensing actuator within a feedback instrument, with one coil actuating the speaker and the second sensing tactile input. Within the field of haptic feedback, studies have explored multiplexing sensor-actuator functionalities with haptic transducers [33, 13, 1]. Such multiplexing prevents truly simultaneous sensing and actuation, though techniques can be utilised to achieve perceptual simultaneity - for example Asai et. al. [1] utilise back-EMF sensing during the low period of the pulse width modulation actuation signal for load sensing with a linear resonant actuator (LRA).

Dementyev et. al. [12] utilise current sensing with linear resonant actuators to achieve a self-sensing configuration with true simultaneous sensing and actuation. The signal level of the measured current is used to detect loading of the actuator due to the change in measured impedance when driving the actuator. A higher loading force upon the actuator reduces the back-EMF, lowering the impedance of the actuator at the resonant frequency.

3. HAPTIC CONTROLLER DESIGN

The haptic controller interface presented in this paper consists of a single voice coil transducer that simultaneously performs three distinct functions achieving the following:

1. Provide vibrotactile feedback to the user's fingertip
2. Sense the force applied to the transducer
3. Act as a pickup for tactile audio excitation signals such as strikes upon the surface of the transducer

The pressure sensing is achieved by exploiting the under-damped nature of the specific transducer - whereby the transducer's impedance at its resonant frequency changes based upon the amount of pressure applied to the surface of the actuator, a technique similar to that of Dementyev et. al. [12]. A current drive amplifier is used for amplification of the haptic actuation signal, with the voltage across the actuator used as an sensor input signal back into the system. A digital filter equivalent of the voice coil's impedance is used to prevent feedback - cancelling the actuation signal from the excitation signal that is returned to the physical model.

3.1 An Electrical Model of a Voice Coil

To accomplish all three functions simultaneously, an equivalent electrical model of a moving coil transducer must be considered, as shown in Figure 1. The components, and their respective values, within this circuit will affect the impedance of the transducer [5]. Due to electrical impedance being a complex value across frequency, both the magnitude and phase elements must be considered.

The measured magnitude across frequency of the voltage across the selected transducer used for this system can be seen in Figure 2 which, due to the use of a current drive amp, is directly proportional to the transducer’s impedance. Notably, the graph displays a resonant peak at 350Hz and a rise in level at higher frequencies. The rise at high frequencies is due to the series inductance of the transducer’s coil. The resonant peak is caused by the parallel RLC circuit formed by the compliance, moving mass, and mechanical damping components. At the resonant frequency, the reactances of the capacitor and inductor cancel, leaving only the mechanical damping resistive component. Changing the values of the compliance and moving mass alter the frequency of the resonance whilst changing the mechanical damping changes the Q-factor and peak impedance at the resonance. It is this particular characteristic that enables the force applied to the transducer to be measured: an increase in the force applied upon the transducer will increase the mechanical damping. This leads to a reduced peak impedance.




Figure 1: An electrical equivalent of a moving coil transducer, modelling both the electrical and mechanical characteristics of the transducer. Adapted from Borwick [5].




Figure 2: Frequency response of the voice coil transducer’s transfer function when driven with a current source measuring the voltage across the transducer in both a damped (red) and undamped (blue) state. Note that the low frequency roll-off seen on both plots is due to the AC-coupling of the measurement interface.

Through driving the transducer with a current drive amplifier, and measuring the voltage across it as an input, any changes in the transducer’s impedance can be detected as a change in the transfer function from input to output, due to Ohm’s law.

3.2 Transducer Selection

Overviews of vibrotactile actuator types and considerations can be found in [6, 23] however it is worth briefly considering the design criteria for an effective self-sensing actuator in a back-emf sensing circuit configuration that is capable of direct representation of audio signals. Actuator requirements for direct representation of audio signals in the haptic

domain are listed by Marshall and Wanderley [23]: the ability to independently control the frequency and amplitude of the haptic signal over a wide range along with a fast transient response. This set of requirements differs significantly from the requirements when choosing a suitable actuator for more rudimentary tasks, such as a notification method in portable devices. In these instances, the priorities are often cost, efficiency and size.

Unlike standard haptic feedback implementations, to satisfy the requirements of a self-sensing transducer capable of pressure sensing, a suitable actuator for this design is required to have an under-damped resonance. This ensures that there is a measurable difference in the actuator’s impedance at the resonant frequency when freely vibrating compared to when pressure is applied by a fingertip, due to the back-emf. Previous studies have utilised a linear resonant actuators [13, 12, 1]. Such devices are efficient to drive and exhibit a fast transient response however they have a limited frequency response. Given the above criteria, a voice coil surface transducer was selected as the most suitable actuator.

After preliminary tests, a Tectonic TEAX25C10-8/SP¹ moving coil surface transducer with a resonant frequency of 350Hz, as seen in Figure 4, was selected, due to its prominent change in impedance due to mechanical damping – shown in Figure 2. This particular type of transducer is designed to output audio when mounted onto a surface, creating a mechanical coupling that further amplifies the audio output. This mechanical coupling increases the mechanical damping experienced by the transducer, leading to a reduced resonant peak and a more even frequency response. While this is a desirable feature when used as a conventional surface transducer, an under-damped transducer is required for the purposes of the system presented in this paper. To achieve this, the transducer is used ‘upside down’ to its usual mounting orientation (as shown in Figure 4), with the haptic interaction occurring directly upon the voice coil enclosure.

3.3 Hardware and Electronics

An overview of the hardware involved in the haptic interaction setup can be seen in Figure 3. The physical modelling synthesiser is run on a laptop, due to it being implemented as an audio plugin and its computational complexity requiring significant processing power. The self-sensing haptic system is based around a Bela Mini low latency embedded computing platform [25] along with a Bela Mini Multichannel Expander to enable three audio inputs and outputs to be used. This system provides a low latency, deterministic platform for signal processing, which is a requirement of the feedback cancellation process. The first audio input on the Bela Mini is used to receive the driving signal from the physical modelling synthesis running on a laptop via a USB audio interface. Similarly, the first output of the Bela Mini system is used to send the processed excitation signal to the physical modelling synthesiser, via an input on the USB audio interface. Output 2 of the Bela Mini sends the processed driving signal to the current drive amplifier – a transconductance amplifier design modified from McPherson [24]. Input two receives the voltage across the transducer.

The third input and output are utilised as an audio loop-back to account for the transfer function of the AC-coupled I/O during the feedback cancellation process. While the amplifier design is DC-coupled, the analogue to digital (ADC) and digital to analogue (DAC) audio converters on the Bela

¹<https://www.tectonicaudiolabs.com/product/teax25c10-8-sp/>




Figure 3: A block diagram of the hardware and processing setup used for simultaneous tactile excitation, force sensing, and haptic feedback to/from the physical modelling algorithm



Figure 4: A Tectonic TEAX25C10-8/SP voice coil actuator

Mini board are AC-coupled. At low frequencies, the DC blocking capacitors will affect the magnitude and phase of the transfer function of the system. By using the loopback signal as a reference signal, the same filtering from the DC-blocking capacitors will be applied to both the actuator signal and the reference signal.

4. TRANSDUCER SIGNAL PROCESSING

To achieve all three uses of the transducer simultaneously, signal processing is applied using the Bela Mini system to both the current output to the transducer and the voltage input back from the transducer, as shown in Figure 3. Features of the transducer’s impedance explained in Section 3.1 are utilised to both measure damping at the resonant frequency and prevent feedback between haptic actuation and tactile excitation.

4.1 Haptic Actuation and Audio Excitation

As the same transducer is also be used as a vibrotactile haptic actuator simultaneously, with the driving signal being directly derived from the physical modelling synthesis process, processing must be applied to prevent feedback.

The issue of feedback and self-oscillation is mitigated utilising the self-sensing qualities of the transducer – a known transfer function can be applied to the driving signal to determine the expected measured voltage signal. As the system is driven by a current drive amplifier and the voltage is measured, the transfer function will match the impedance of the transducer. An equivalent digital filter that matches the transducer’s impedance can be created and applied to the reference loopback signal, along with a wideband matching gain adjustment to account for the amplifier gain. The filtered signal is subtracted from the measured voltage across the transducer, cancelling the driving signal from the measured voltage. The remaining signal will only contain any excitation due to the user tapping or hitting the transducer.

The significant features of the transfer function noted in Section 3.1 of the resonant peak and high frequency rise must be considered in the filter design. The resonant peak in the magnitude response is due to the parallel RLC circuit, as described in Section 3.1, which creates a second order resonant filter. In the phase response, the capacitive reactance of the RLC filter causes the current to lag behind the voltage below the resonant frequency, causing a positive phase angle. The phase angle at the resonant frequency is zero, due to the capacitive and inductive reactances cancelling. Above the resonant frequency, the current leads the voltage causing a negative phase angle. Both the magnitude and phase response of this circuit can be modelled using a biquad filter topology, configured as a peak filter. The gain and Q-factor parameters of the filter are adjusted to match the resonance of the transducer.

The rise at high frequencies due to the coil’s series inductance must also be considered as part of this modelling filter, however Figure 2 shows the high frequency rise to be insignificant below 1kHz. This is of significance as during preliminary tests it was noted that most of the energy in the excitation audio signal, when using the transducer as a pickup for excitation of the physical model with taps and hits, was below 500Hz. When analysed in the frequency domain, a peak in level at 150Hz was noted, with a 20dB reduction by 500Hz and 40dB by 1kHz. This is due to the actuator being constructed to be optimised as an audio actuator rather than sensor; the moving coil’s high mass causes it to be ineffective at picking up high frequencies.

Furthermore, the upper limit of frequencies within the

range of human vibrotactile perception is approximately 1kHz [21]. Given this, the actuation signal does not need frequency content above 1kHz, nor does the transducer pick up high frequencies, thus the decision was taken to apply 2nd order low pass filters at 1kHz to both the actuation signal to drive the transducer and the excitation signal picked up by the transducer. This avoids the need to accurately model the series inductance as an additional digital filter. Additionally, the low pass filter on the actuation output signal minimises the level of audible sound from the transducer.

One notable caveat is that the resonance filter is modelled upon the transducer in an undamped state (i.e. no external force is being applied). When a force is applied, the increased mechanical damping causes a reduction in the magnitude of the resonant peak. This will cause the modelled filter to no longer match the actual transfer function of the system. To mitigate this occurrence a notch filter at f_0 is applied to remove any signals at the resonant frequency from the excitation signal.

4.2 Force Sensing

Force sensing utilises the change in mechanical damping of the transducer when force is applied. This change in mechanical damping changes the impedance of the transducer at the resonant frequency, as shown in Figure 2. To measure this change, a notch filter at f_0 is applied to the haptic actuation signal. During prototyping, this was found to reduce the variability of the force reading from the transducer, as it reduces the variation in level at frequencies around f_0 . After this notch filter, an oscillator signal at f_0 is added into the actuation signal. This provides a constant measurable level at the resonant frequency. The oscillator level is adjusted to be at the threshold of perception thereby maximising the available signal to noise ratio, whilst avoiding the addition of a perceivable constant haptic vibration.

The amount of force applied to the transducer is then measured by taking the difference signal (with feedback cancellation applied) and applying the Goertzel algorithm to the signal at frequency f_0 [19]. This provides a computationally efficient method of measuring the signal level at a particular frequency; equivalent to a single bin of a discrete Fourier transform. The Goertzel algorithm is also applied to the haptic actuation driving signal and the difference is taken between the two results, to account for any variation in level within that particular frequency bin in the driving signal. This final result is then scaled from 0 to 127 to fit the range of a MIDI control change value, with the upper and lower bound determined empirically by observing the measured values with the actuator in a damped and undamped state. MIDI is used in this particular configuration to enable mapping to game parameters.

5. VIRTUAL INSTRUMENT DESIGN

The audio excitation signal returned from the haptic interface is used as an input signal to the physical modelling system. The basic building block of the physical model considered here is a vibrating plate, which is most naturally set into vibration by a strike or a tap among the common physical modelling systems. The virtual instrument, however, does not comprise one plate alone; rather, it is a system of two connected plates. The connections, which will be briefly discussed below, have a nonlinear character and, hence, can turn the bandlimited input signal into a wideband, noise-like output signal. This helps to mitigate the issue that the transducer used in this system is ineffective at sensing

high frequency signals. Figure 5 depicts the simulated system comprising two plates and the nonlinear connections: these can either be springs or rattles presenting a gap and colliding intermittently.




Figure 5: System of two plates connected via springs (left) or rattles (right)

The connections are here energy-storing devices, for which the resulting force may be given as the gradient of a potential as:

$$f(\eta) = -\frac{d\phi}{d\eta}, \quad (1)$$

where η is the elongation of the connection. In the rattle configuration, intermittent contact is permitted such that:

$$\phi(\eta) = \frac{K}{\gamma+1} [|\eta| - \beta]_+^{\gamma+1} \geq 0. \quad (2)$$

Here, $K \geq 0$ is a stiffness constant, $\gamma \geq 1$ is a nonlinear exponent, and $\beta \geq 0$ is a gap. Figure 6 represents the potential and the resulting force. Note that no force is exerted by the spring when $|\eta| < \beta$, resulting in a rattling-type force [4]. More details regarding the model and the time-




Figure 6: Example of nonlinear potential (2) and corresponding force (1). Here, $K = 10^3$, $\gamma = 1.2$, $\beta = 0.1$. The shaded area, whose width is given by 2β , represents a dead zone (no force exerted), yielding intermittent contact between the plates.

stepping routine used to solve it can be found in [14]. Examples of typical resulting waveforms are given in Figure 7. As the nonlinearity increases, common nonlinear phenomena ensue, such as increased wave propagation speed, modal couplings, and energy cascades yielding the crashing and shimmering sound typical of gongs and cymbals.

5.1 Implementation

The system is implemented as an audio instrument plug-in, i.e. a VST3i, but with additional busses to allow a signal from the connections to be extracted as a 3/4 stereo output and sidechain audio to be inserted via the Digital Audio Workstation (DAW). This setup is shown in Figure 8. These additional busses are defined in the processor engine of the plug-in. Whilst providing additional outputs is straightforward, sidechaining input busses is more complicated due to the way different DAWs handle this setup for instruments




Figure 7: Time domain and frequency domain signals. The input signal is an example of the normalised recorded voltage across the voice coil transducer when tapped with a fingertip. The output signals are the normalised velocities recorded from the virtual plate surface, using a small and a large value of the nonlinear parameter K in (2). Note that, as the nonlinearity grows, the bandwidth of the resulting output signal expands due to the mechanical nonlinearity in the physical model.




Figure 8: Virtual instrument overview with 2 output busses, and sidechain input from the Bela used to drive the top plate of the physical model. A control rate MIDI signal is used for additional parameter control.

which exist on software rather than audio tracks. The audio rate signal passed to the 3/4 output bus is the velocity of the connection elongation, $d\eta/dt$, where η is as per (1). This allows the user to feel the movement of the spring or rattle as each of the plates is forcing it. The transducer output is fed back into the model as the forcing signal is applied to the top plate. Thus, a feedback loop is created via the transducer, controlled by the user interaction.

The use of an audio plugin format enables mapping of the force level MIDI output from the Bela Mini to a physical modelling parameter. During preliminary tests of the system, mapping to either the plates damping controls afforded the most natural interaction experience – likely to to an equivalent physical system exhibiting similar behaviour to damping when a continuous force is applied to it – though further exploration of such mappings is required within a fully-featured instrument design. A video demonstration of the system is available at: <https://youtu.be/xsxedAsg9B4>.

6. DESIGN REFLECTIONS

Our motivation in this work is enable richer integration of physical and virtual elements in DMI design, where the coupling between domains can be continuous and bidirectional without resorting to simplified symbolic controls such as MIDI triggers, nor artificially separating different aspects of interaction (e.g. excitation, damping, haptic sensing) into different physical locations or modalities. The designs presented here reflect a process of exploration, iteration and fine tuning, during which we balanced technical goals with the experience of interacting with the combined system of transducer and physical model.

Audio-rate signals, both for excitation and haptic feedback, are one crucial element promoting deep physical-virtual integration. During development, we found that the continuous nature of such signals encourages a more constant and sustained form of interaction with the physical model when compared to interaction using controllers with discretised input events (such as MIDI keyboards). In our explorations, the system was particularly suited for techniques that involved exciting the synthesiser with a strike, then controlling the damping parameter using the force sensing during the decay of the plates. This playing technique is demonstrated in the accompanying demonstration video, where the tactile manipulation of the damping parameter can be observed after striking the transducer. In an echo of enactive design principles [16], this technique benefits from intuitions from the physical world: striking a physical plate then applying force onto the surface to dampen its resonance is a natural interaction technique with such objects. We found that the additional tangibility through the vibrotactile feedback created a more conducive environment for intimate control over the other synthesis parameters – manipulating the controls whilst interacting with the sound rather than tuning in a sound first in a “set and forget” manner.

A notable caveat of the design is the reduction of the force sensing from an audio rate sensing signal to a control rate signal over MIDI. While this is not an intrinsic requirement for controlling the physical modelling synthesiser, this is necessary due to the signal detection implementation employed; an audio rate signal may afford additional nuances

in the control possibilities. The significant difference between audio and control rates is the bandwidth afforded. The physical model paired with the haptic controller in this system is able to blur the distinction between the two – using variable non-linearities to recreate a full bandwidth signal from control signals and a limited bandwidth input, with a range of different timbres available by adjusting the physical parameters used to define the plates.

Though the tactile system presented is able to accomplish three distinct interaction tasks simultaneously, the true collocation of sensing and actuation afforded by a self-sensing approach does impose further design limitations, including a limited audio sensing bandwidth. Additionally, the necessity of parameterising the force measurement into a control signal constrains the richness afforded by the force sensing of the interface and the limited range over which the actuator is able to detect changes in force makes precise control of parameters challenging. Such constraints impose certain design considerations upon any prospective instrument designs, revealed through the system presented in this paper – considerations such as using signal processing techniques that include non-linearities to mitigate the transducer’s limited frequency response and constraining the force control’s range across a parameter to a limited range. In this sense, the function of haptic actuation and the reproduction of a convincing tactile experience collocated with the sensing mechanism was prioritised as a feature of the self-sensing system, with fewer design constraints placed upon the haptic actuation.

7. CONCLUSIONS

A self-sensing vibrotactile haptic transducer configuration has been presented, capable of vibrotactile actuation, force sensing, and acting as a sensor for audio excitation signals simultaneously. The haptic system is designed to integrate tightly into the physical modelling synthesis environment, where haptic actuator and excitation of the physical model both occur at audio rate and the force applied to the transducer is mapped to a control parameter on the physical model. The work presented in this paper aims to contribute a deeply integrated approach to physical modelling synthesis between the physical and digital domains, enabling a close tactile coupling through bidirectional channels between instrument and musician, providing localised haptic feedback and a compact footprint – only requiring a single voice coil transducer at each contact point. The use of physical modelling synthesis affords intuitive tactile interaction due to the system having a direct physical equivalent. The presented system is not a complete musical instrument design but rather a technological enabler for further exploration of tightly coupled haptic interaction, demonstrating the possibilities of a self-sensing transducers within the discipline of digital musical instrument design.

8. ACKNOWLEDGMENTS

Matthew Davison and Andrew McPherson received funding from the UKRI Frontier Research (Consolidator) grant (EP/X023478/1) “RUDIMENTS: Reflective Understanding of Digital Instruments as Musical Entanglements.”

Michele Ducceschi and Craig Webb received funding from the European Research Council (ERC) under the Horizon2020 framework, grant number 950084-StG-NEMUS.

9. ETHICAL STANDARDS

This work is an empirical study of a technology system and involved no research participants.

10. REFERENCES

- [1] Y. Asai, K. Hirata, and T. Ota. Amplitude Control Method of Linear Resonant Actuator by Load Estimation From the Back-EMF. *IEEE Transactions on Magnetics*, 49(5):2253–2256, May 2013.
- [2] E. Berdahl and A. Kontogeorgakopoulos. The FireFader: Simple, Open-Source, and Reconfigurable Haptic Force Feedback for Musicians. *Computer Music Journal*, 37(1):23–34, Mar. 2013.
- [3] E. Berdahl, J. O. Smith, and G. Niemeyer. Feedback control of acoustic musical instruments: Collocated control using physical analogs. *The Journal of the Acoustical Society of America*, 131(1):963–973, Jan. 2012.
- [4] S. Bilbao, M. Ducceschi, and C. Webb. Large-scale real-time modular physical modeling sound synthesis. In *Proceedings of the International Conference on Digital Audio Effects (DAFx-19)*, Birmingham, UK, September 2019.
- [5] J. Borwick. *Loudspeaker and Headphone Handbook*, 3rd Edition. Routledge, 3rd edition edition, 2012.
- [6] J. Bowers and A. Haas. Hybrid Resonant Assemblages: Rethinking Instruments, Touch and Performance in New Interfaces for Musical Expression. In *New Interfaces for Musical Expression*, London, UK, June 2014.
- [7] C. Cadoz, L. Lisowski, and J.-L. Florens. A modular feedback keyboard design. *Computer Music Journal*, 14(2):47–51, 1990.
- [8] C. Cadoz, A. Luciani, and J. Florens. Responsive input devices and sound synthesis by simulation of instrumental mechanisms: The Cordis system. *Computer Music Journal*, 8(3):60 – 73, 1984.
- [9] C. Chafe. Tactile audio feedback. In *Proceedings of the International Computer Music Conference*, Tokyo, Japan, 1993.
- [10] P. J. Christensen, D. Overholt, and S. Serafin. The Daïs: A Haptically Enabled NIME for Controlling Physical Modeling Sound Synthesis Algorithms. In *NIME’20*, Birmingham, UK, July 2020.
- [11] P. R. Cook and G. P. Scavone. The synthesis toolkit (stk). In *International Conference on Mathematics and Computing*, Beijing, China, 1999.
- [12] A. Dementyev, P. Getreuer, D. Kanevsky, M. Slaney, and R. F. Lyon. VHP: Vibrotactile Haptics Platform for On-body Applications. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pages 598–612, Virtual Event USA, Oct. 2021. ACM.
- [13] A. Dementyev, A. Olwal, and R. F. Lyon. Haptics with Input: Back-EMF in Linear Resonant Actuators to Enable Touch, Pressure and Environmental Awareness. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pages 420–429, Virtual Event USA, Oct. 2020. ACM.
- [14] M. Ducceschi, S. Bilbao, and C. Webb. Real-time modal synthesis of nonlinearly interconnected networks. In *Proc Digital Audio Effects (DAFx-23)*, pages 53–60, Copenhagen, Denmark, 2023.
- [15] A. Eldridge, C. Kiefer, D. Overholt, and H. Ulfarsson. Self-resonating Vibrotactile Feedback Instruments ||: Making, Playing, Conceptualising :|. In *NIME 2021*,

Shanghai, China, June 2021. PubPub.

- [16] G. Essl and S. O’modhrain. An enactive approach to the design of new tangible musical instruments. *Organised sound*, 11(3):285–296, 2006.
- [17] J.-L. Florens. Expressive bowing on a virtual string instrument. In A. Camurri and G. Volpe, editors, *Gesture-Based Communication in Human-Computer Interaction*, pages 487–496, Berlin, Heidelberg, 2004.
- [18] F. Fontana, S. Papetti, H. Järveläinen, and F. Avanzini. Detection of keyboard vibrations and effects on perceived piano quality. *The Journal of the Acoustical Society of America*, 142(5):2953–2967, Nov. 2017.
- [19] G. Goertzel. An Algorithm for the Evaluation of Finite Trigonometric Series. *The American Mathematical Monthly*, 65(1):34, Jan. 1958.
- [20] S. E. Hudson. Using light emitting diode arrays as touch-sensitive input and output devices. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*, pages 287–290, Santa Fe NM USA, Oct. 2004. ACM.
- [21] L. A. Jones. *Haptics*. The MIT Press Essential Knowledge Series. The MIT Press, Cambridge, Massachusetts, 2018.
- [22] M. K. Klingbeil. *Spectral Analysis, Editing, and Resynthesis: Methods and Applications*. PhD thesis, Columbia University, 2009.
- [23] M. T. Marshall and M. M. Wanderley. Vibrotactile Feedback in Digital Musical Instruments. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression (NIME06)*, Paris, France, 2006.
- [24] A. McPherson. The Magnetic Resonator Piano: Electronic Augmentation of an Acoustic Grand Piano. *Journal of New Music Research*, 39(3):189–202, Sept. 2010.
- [25] A. P. McPherson and V. Zappi. An Environment for Submillisecond-Latency Audio and Sensor Processing on BeagleBone Black. In *Audio Engineering Society Convention 138*, Warsaw, Poland, May 2015.
- [26] L. Mice and A. P. McPherson. Super Size Me: Interface Size, Identity and Embodiment in Digital Musical Instrument Design. In *CHI Conference on Human Factors in Computing Systems*, pages 1–15, New Orleans LA USA, Apr. 2022. ACM.
- [27] R. Michon and J. O. I. Smith. A Hybrid Guitar Physical Model Controller: The BladeAxe. In *Proceedings of the 2014 International Computer Music Conference*, Athens, Greece, 2014.
- [28] C. Nichols. The vBow: a virtual violin bow controller for mapping gesture to synthesis with haptic feedback. *Organised Sound*, 7(2):215–220, 2002.
- [29] S. O’Modhrain and R. B. Gillespie. Once More, with Feeling: Revisiting the Role of Touch in Performer-Instrument Interaction. In *Musical Haptics*, Springer Series on Touch and Haptic Systems, pages 11–28. Springer International Publishing, Cham, 2018.
- [30] M. G. Onofrei, F. Fontana, and S. Serafin. Perceptual Relevance of Haptic Feedback during Virtual Plucking, Bowing and Rubbing of Physically-Based Musical Resonators. *Arts*, 12(4):144, July 2023.
- [31] K. Oshima, S. Fujii, and T. Mimura. Noise Reduction by a Loud Speaker Based on Self-Sensing Actuation. *IFAC Proceedings Volumes*, 33(25):47–50, Sept. 2000.
- [32] S. Papetti and C. Saitis. *Musical haptics*. Springer Nature, 2018.
- [33] G. Savioz and Y. Perriard. Towards self-sensed drives in linear haptic systems. In *2009 International Conference on Electrical Machines and Systems*, pages 1–5, Tokyo, Nov. 2009. IEEE.
- [34] P. Schmalfuß, M. Neupert, and B. Kessler. Efficient Snare Drum Model for Acoustic Interfaces with Piezoelectric Sensors. In *Proceedings of the 23rd International Conference on Digital Audio Effects (DAFx-20)*, Vienna, Austria, Sept. 2020.
- [35] S. Serafin, M. Burtner, C. Nichols, and S. O’Modhrain. Expressive controllers for bowed physical models. In *Proceedings of the COST G-6 Conference on Digital Audio Effects (DAFX-01)*, Limerick, Ireland, 2001.
- [36] J. D. Sierra. SpeakerDrum. In *NIME’20*, Birmingham, UK, June 2020. Zenodo.
- [37] J. O. Smith. Viewpoints on the history of digital synthesis. In *International Conference on Mathematics and Computing*, Montreal, Canada, 1991.
- [38] M. Verma, V. Lafarga, and C. Collette. Perfect collocation using self-sensing electromagnetic actuator: Application to vibration control of flexible structures. *Sensors and Actuators A: Physical*, 313:112210, Oct. 2020.
- [39] B. Verplank, M. Gurevich, and M. Mathews. The Plank: Designing a Simple Haptic Controller. In *New Interfaces for Musical Expression*, pages 59–70, Dublin, Ireland, May 2002.