Exposing the Scaffolding of Digital Instruments with Hardware-Software Feedback Loops

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

Victor Zappi
Media and Graphics Interdisciplinary Centre
University of British Columbia
Vancouver, BC, Canada
victor.zappi@gmail.com

ABSTRACT

The implementation of digital musical instruments is often opaque to the performer. Even when the relationship between action and sound is readily understandable, the internal hardware or software operations that create that relationship may be inaccessible to scrutiny or modification. This paper presents a new approach to digital instrument design which lets the performer alter and subvert the instrument's internal operation through circuit-bending techniques. The approach uses low-latency feedback loops between software and analog hardware to expose the internal working of the instrument. Compared to the standard control voltage approach used on analog synths, alterations to the feedback loops produce distinctive and less predictable changes in behaviour with original artistic applications. This paper discusses the technical foundations of the approach, its roots in hacking and circuit bending, and case studies of its use in live performance with the D-Box "hackable instrument".

Author Keywords

Feedback, circuit bending, hacking, analog circuits, low latency, embedded hardware

ACM Classification

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.

1. INTRODUCTION

The implementation of an instrument matters. For the designer, the exact selection of materials, circuits or code may be primarily a means to an end, a way of achieving a particular sound. For the performer, though, the subtle idiosyncracies and limitations of how the instrument behaves can become powerful sources of creative inspiration, giving rise to new playing techniques and even new styles of music.

For example, in early guitar amplifiers, distortion was an engineering limitation until guitarists discovered the musical value in pushing the device beyond its limits. The amplifier's internal construction matters just as much as its nominal action of making the guitar louder: tubes and transistors produce roughly equivalent results in linear amplification, but their sound when pushed into distortion is noticeably different, such that tube amplifiers remain dominant despite cost and practical shortcomings.

Similarly, in the 19th century, the high-tension steel strings in the piano were the means to an end of creating loud and sustained tones, but the design also makes possible John Cage's prepared piano, the bowed ensemble pieces of Stephen Scott, and electromagnetic actuation techniques [3, 14]. A hypothetical digital piano which perfectly emulated the sound and feel of an acoustic grand would nonetheless be more restrictive than the original, since these and other interventions would be impossible.

1.1 Implementation of Digital Instruments

The internal implementation of a digital musical instrument (DMI) is often a black box. Even when the nominal relationship between action and sound is obvious, the software and hardware processes that create that relationship may be inaccessible to the performer. This inaccessibility can prevent the kinds of creative misuses and modifications which are common on acoustic instruments [20].

This paper presents a new method of exposing the internal operations of DMIs, allowing performers to modify or subvert the original design for creative purposes. The aim is not to explain or demonstrate the instrument's nominal mapping to the performer or audience, but rather to make the *implementation* of the instrument open to repurposing.

1.2 Appropriation, Modification and Hacking

In an earlier study [21], we showed how even very simple instruments were appropriated by performers for use in unexpected ways. Surprisingly, our results suggested that perfomers saw instruments with *fewer* degrees of freedom as richer, and that tight constraints provided greater incentives to develop creative use cases.

This techniques in this paper widen our study of appropriation to include modifications to the instrument's internal operation, while retaining the essential simplicity of the instrument. Specifically, our techniques aim to:

- Expose the internal state of an instrument's software (not just sensor inputs) to modification *in hardware.*
- Do not limit the allowable modifications to choosing from predefined options created by the designer.
- Make modifications interesting: the instrument should not crash or go silent, and the behaviour when modified should include unusual but repeatable effects.

Our approach is inspired by circuit bending practice (Section 2.4). The technical implementation uses low-latency feedback loops between software and hardware which produce unique characteristic effects when altered. To state the goal more compactly: **Make the instrument simple, but make it break in interesting ways.**

NIME'15, May 31-June 3, 2015, Louisiana State Univ., Baton Rouge, LA.

2. BACKGROUND

The concept of *modularity* has been an important factor in the development of electronic instruments, but it is not the only approach to modifiability. This section examines analog control voltages (CVs) and their influence on DMI mapping before considering limitations and alternatives.

2.1 Dominance of the Control Voltage

Hugh LeCaine's Electronic Sackbut (1945-48) is widely seen as the first voltage-controlled synthesizer [19]. Its keyboard features pitch and volume control from lateral and downward pressure, respectively, while a set of pressure sensors played with the left hand control various aspects of timbre. The modular synths of Moog and Buchla in the 1960's and 70's established control voltages as the dominant paradigm for analog synthesis [17].

The advantage of the CV lies in its conceptual clarity and interconnectability: by describing each musical parameter with a single continuous variable, any signal source can conceivably be attached to any destination with predictable results. Physical patch cords for routing CVs are a staple of analog modular synths, but modular approaches can be found whether or not cords are literally present, for example in modulation matrices or switches for reassigning control parameters. Recent kits such as littleBits and Patchblocks[1] have also featured interconnectable hardware modules.

2.2 Digital Musical Instrument Mapping

Mapping in DMIs refers to the assignment of control (sensor) inputs to instrument outputs (audio, visual, haptic). Many mapping strategies have been proposed, including one-to-one and many-to-many relationships [10], multi-layer approaches [1], and dynamic performer-adjustable mappings [12]. Machine learning is increasingly used for creating mappings [8, 11, 5]. Recent efforts have also focused on making mappings understandable to the audience [2, 16].

Mapping, particularly in its more direct forms, can be seen as a digital extension of the CV: musical parameters are associated with one or more continuous or discrete variables, which are then assigned to various sensors or other signal generators. Like CVs, conceptual clarity and reconfigurability are advantages: the effect of making particular assignments can often be predicted in advance or (using supervised learning) implicitly created by training examples.

2.3 Limitations

A major DMI limitation occurs when the *abstraction* becomes more important than the *implementation*. The piano action is not just a tool for creating decaying musical tones: it is also a combination of steel, wood and felt which can be manipulated in ways unrelated to its original function. Likewise, an analog CV is not only a representation of a musical parameter, it is also an electrical signal with finite impedance and bandwidth, susceptible to noise and interference. The synth designer might seek to minimise these effects, but the performer could choose to accentuate them by floating or shorting cables, using skin conductance to manipulate them, or injecting unrelated electrical signals into the circuit.

No such possibilities are available in CV-inspired digital mappings. Modularity is seen as a way to offer the performer more possibilities, but its design places strict limits on the creative space. Only the patch cords (real or virtual) can be manipulated, and there is no physical substrate whose properties can be exploited for unexpected effects.

Most familiar instruments have undergone many generations of testing and refinement, often incorporating what were originally unexpected usage techniques (e.g. fuzz boxes). The limitations opaque DMIs impose on the performer may thus end up affecting their future development.

2.4 Modularity versus Circuit Bending

Circuit bending [6] rejects the black box model of digital electronic devices, working directly with the circuits and signals inside. Modularity is irrelevant, since any trace on the board that is physically accessible can become part of a modification. Some objects become prized circuit bending targets precisely because of their implementation. The TI *Speak and Spell* (1978) has been used in dozens of hacks to produce new sound effects [9]. A smartphone emulation of the device would hold no interest for circuit bending; arbitrary modifications would lead to a crash or damage to the device rather than creating interesting effects.

It is rare to see a circuit bender working with a latest-generation DMI, other than one they have created. We thus argue for a new approach to DMI design which minimises opaque software processes and exposes as much of the substrate as possible to modification by the performer.

2.5 Feedback in Musical Instruments

Feedback in performance dates back to the 1960's [18], with examples including amp-to-pickup on electric guitar, Tudor's *Microphone* (1973) and other works, Di Scipio's *Audible Ecosystems* [7], Collins' 1974 *Pea Soup*, Nakamura's 1997 *No-Input Mixing Board* [18], and Bowers and Haas's "hybrid resonant assemblages" [4]. Feedback through skin conductance features in Waisvisz's *Cracklebox*.[2]

In comparison to previous work, this paper explores feedback as a control mechanism within a larger DMI rather than as a direct audio output. It also demonstrates the construction of loops between software and analog electronics whose effects could not be created with either modality alone. Finally, it demonstrates a circuit-bending methodology for modifying feedback loops to alter the behaviour of a primarily software-driven instrument.

3. TECHNICAL FOUNDATIONS

The techniques in this paper form the modifiable core of the D-Box "hackable instrument" (Figure 1) [20]. The D-Box is designed to appear simple and limited at first but allow wide-ranging creative exploration through modifications to a network of analog circuits within the case.

3.1 Ultra Low-Latency Audio and Sensors

BeagleRT [15] is a new ultra-low-latency embedded audio platform created specifically for the D-Box. It is based on the BeagleBone Black single-board computer which features a 1GHz ARM Cortex-A8 processor and 512MB of RAM. BeagleRT uses the Xenomai hard real-time Linux kernel extensions to run audio I/O at a higher priority than the kernel itself. This means that other activity on the system does not affect the timing of audio calculations, enabling extremely small buffer sizes which cannot be achieved on any other Linux-based audio platform.

Using a buffer size of 2 audio samples, a latency of under 100μs from input to output can be achieved with no underruns. The D-Box uses a buffer size of 4 audio samples, which is more CPU-efficient, with a latency of 184μs.

A hardware *cape* (expansion board) features 8 channels each of 16-bit ADC and DAC in addition to stereo audio

---

[1] http://littlebits.cc and http://patchblocks.com

[2] That the Cracklebox will only work with 1970's-vintage externally compensated op amps and not with modern chips shows again the importance of the internal implementation.

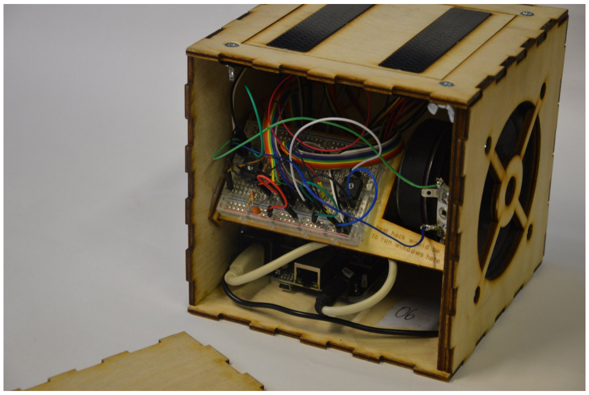

Figure 1: The D-Box, a hackable digital instrument. Circuits on the breadboard form feedback loops with the internal software, and changing these circuits alters the behaviour of the instrument.

input and output. The ADC and DAC are sampled automatically at half the audio sample rate (22.05kHz), resulting in jitter-free alignment of sensor data with respect to the audio signal. 1W amplifiers onboard the cape can drive 8-ohm speakers; a mono speaker is used inside the D-Box.

3.2 D-Box Sensors and Mapping

The D-Box uses two capacitive touch sensors derived from [13]; one sensor contains a pressure sensor underneath. In its initial configuration, touch position on the first sensor controls pitch of a stored sample reconstructed by oscillator bank, and pressure controls volume. Location on the other sensor applies a bandpass filter effect. New samples can be loaded onto the instrument by SD card.

The synthesis is further described in [20]. The mapping is deliberately simple: there are several further parameters not affected by the sensors in any way, including choice of sample, playback speed, wavetable (timbre) of the oscillators, and loop points within the sample. The software is open source [15], but the study in this paper focused only on hardware modifications.

3.3 The Matrix

In the D-Box, the 8 analog inputs and outputs on the BeagleRT cape are connected to a breadboard here called the *matrix* for its network of connections between inputs and outputs. The matrix is the substrate on which the internal operation of the instrument is exposed, and the basis for the feedback loops described in the next section.

In the initial configuration of the D-Box, the matrix is prepopulated with a standard circuit containing several resistors and capacitors (Figure 2). Notably, modification to the matrix is not limited to patch cords or selection amongst discrete options; any electrical components can be used.

4. FEEDBACK AS CONTROL STRATEGY

In most DMIs, signals flow in one direction, from sensor inputs through a mapping layer (usually software, often opaque) to sound outputs. There is an implicit form of feedback in the performer's reactions to the instrument, but less commonly is feedback found internally in the instrument.

By contrast, the D-Box contains several internal hardware-software feedback loops. A signal generated by software is sent to a DAC channel, whereupon it is acted on by a network of analog components. The output of the network is read by an ADC channel and processed again by the software. Either positive or negative feedback can be used, and the signal processing can be either linear or nonlinear. Positive feedback tends to create oscillators; negative feedback can help pull a control signal toward a desired value.

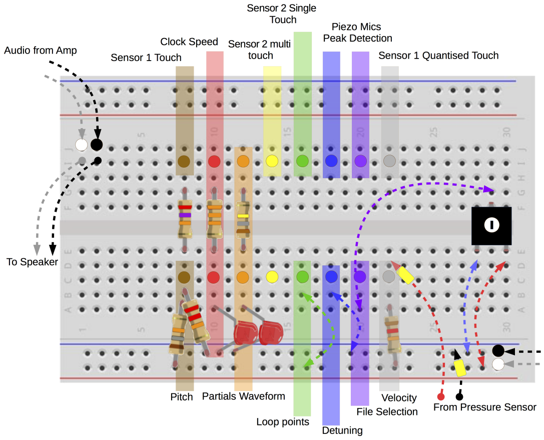

Figure 2: The matrix, a breadboard of hackable analog circuits within the D-Box. DAC outputs at top, ADC inputs at bottom.

Rather than using the loop as a direct audio signal generator, features are extracted from the signal that emerges which then become part of the instrument's mapping. Alterations to the loop produce effects which behave differently to CV control: changes are less linear, more chaotic, include time-variant effects, and often exhibit an interesting boundary between stability and chaos.

The following sections demonstrate several examples of how feedback is used in the D-Box.

4.1 Adaptive Pitch Control

Pitch on the D-Box is controlled by finger location on a touch sensor. Though this mapping could be performed entirely in software, this offers no opportunity to alter its response. Instead, touch position generates an analog voltage from the DAC. This voltage runs through a resistor network which shifts and scales it, before it is read by the ADC (Figure 3).

Changing resistor values changes the pitch range, and portamento effects can be obtained by including a capacitor in the circuit. Alternatively, the ADC input can be left floating which makes the pitch highly sensitive to nearby stray signals, or cross-wired to other circuits (Section 4.4).

A separate DAC output uses negative feedback to pull the pitch toward the nearest semitone. At sample n, the ADC input is compared to the nearest value which would

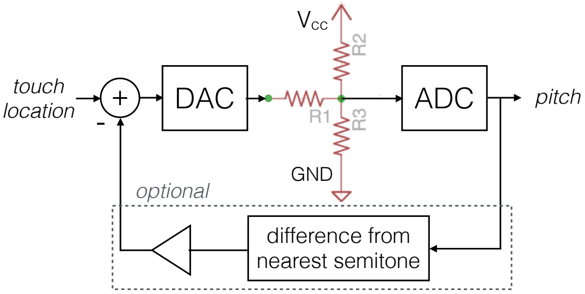

Figure 3: Software-hardware signal flow for controlling pitch. Circuit in the middle can be changed.

produce an equal-tempered semitone, and the difference is calculated. At sample $n+1$, a fraction of the difference is subtracted from the DAC output, such that the pitch gradually converges to equal temperament.

4.2 Phase-Shift Oscillator: Waveform Control

The phase-shift oscillator is a classic analog design consisting of three first-order lowpass filters in a closed loop. At the frequency where the total phase shift reaches 180 degrees, the system will self-oscillate.

The D-Box implements a hybrid analog-digital phase shift oscillator (Figure 4), using the fact that a constant delay produces a linear phase shift with respect to frequency. The DAC output passes through a first-order RC network and is read by the ADC. In software, the ADC signal is high-pass filtered to remove DC offset and sent back to the DAC. The RC phase shift plus the buffering delay produce 180 degrees of shift at some frequency, at which the system will oscillate.

The oscillator does not feed directly into the D-Box audio output. Instead, the waveform of the oscillation is extracted, stretched to occupy 1024 samples, and used as the waveform for each oscillator of an oscillator bank synthesizer. Because of nonlinearities and limited voltage range, the waveform is never perfectly sinusoidal, with smaller capacitors providing a square-like oscillation (Figure 5). Changing the circuits therefore changes the timbre of the instrument. If the RC network is disconnected entirely, the oscillation stops. To prevent the instrument from going silent, a sine wavetable is used when no oscillation is detected.

Changing circuit elements can produce distinctive transient behaviour (Figure 5d), including unstable oscillation patterns or ramps in oscillation amplitude. Suddenly removing the capacitor from the network can trigger one of several high-frequency oscillations (Figure 5b-c). The same action, executed multiple times, can produce different oscillation patterns (hence different timbres), but interestingly, each pattern appears to be stable once it has been initiated.

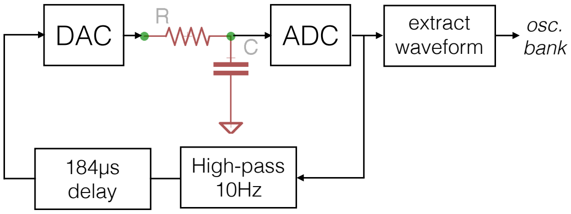

Figure 4: Software-hardware phase-shift oscillator controlling oscillator bank wavetable.

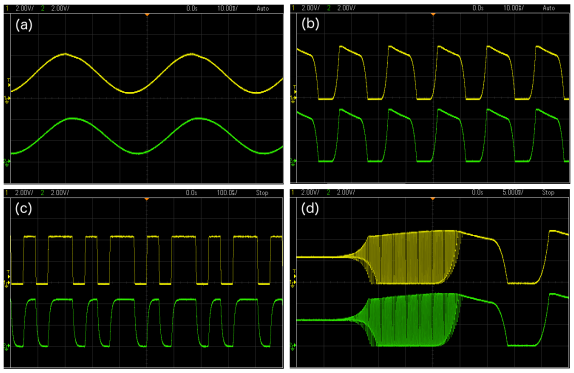

Figure 5: Phase shift oscillator waveforms. Yellow plot (top) = DAC, green (bottom) = ADC. (a) quasi-sinusoidal oscillation (R=3.3k, C=10μF); (b,c) two results when capacitor removed (R=180k); (d) transients on circuit change.

4.3 Relaxation Oscillator: Playback Speed

Relaxation oscillator circuits have been used since the earliest days of electronics. In a typical implementation, a comparator with hysteresis (Schmitt trigger) is used with feedback through an RC network. The output of the comparator can have two states (high and low), neither stable: a high output causes the capacitor to charge, eventually flipping the output to the low state, after which the process repeats in reverse. Relaxation oscillators produce non-sinusoidal outputs, in this case a square or triangular wave depending on which point in the circuit is probed.

The D-Box creates a hybrid feedback loop where the comparator with hysteresis is implemented in software, but the RC network in hardware (Figure 6). The oscillator frequency determines the rate of playback of the stored sample, within a range of 0.1x to 10x normal speed. With default component values, the network oscillates at 160Hz, corresponding to 1x speed. Changing either the resistor or capacitor changes the oscillation frequency (Figure 7). Alternatively, other time-varying signals could be fed into the ADC input from which the playback speed would be derived.

4.4 Cross-Wiring Feedback Circuits

The parameters these feedback loops control– pitch, waveform, playback speed– could easily be controlled entirely in software or using analog CVs. The unique value in the feedback loop approach comes from the unorthodox ways in which these circuits can be manipulated and subverted.

All the voltages on the breadboard “matrix”, including the speaker output, take the range 0-5V. Voltages from any part of the matrix can thus be attached to any other part. For example, the speaker output can be fed back to any of the feedback networks, creating a chaotic result with different sound qualities depending on which network is chosen (Figure 8, also Section 5.4). One major motivation for using feedback is in part that the designer does *not* specify the complete range of possible effects. Many side-effects of changing the feedback networks will be unstable or unpredictable, yet repeatable and far from random.

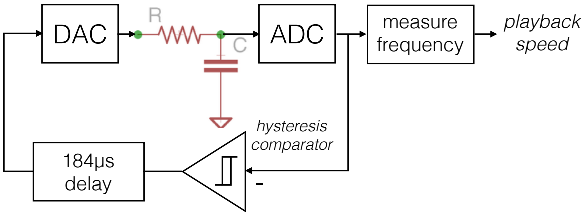

Figure 6: Software-hardware relaxation oscillator controlling sample playback speed.

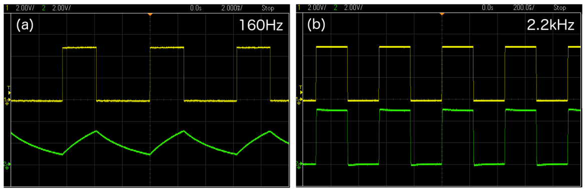

Figure 7: Relaxation oscillator waveforms. Yellow (top) = DAC, green (bottom) = ADC. (a) standard setting (R=33k, C=0.1μF); (b) no capacitor.

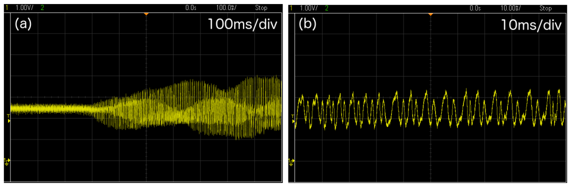

Figure 8: Chaotic result of pitch input when cross-wired to the speaker.

5. USER STUDY

We conducted an extended user study with the D-Box to see how performers would make use of its modifiable circuits. The reaction of any given musician towards the D-Box and its hackable design is not easily predictable, depending on diverse factors such as musical background, electronics expertise or rehearsal context. We gathered 14 volunteer performers, a heterogeneous group ranging from experienced circuit benders to classical instrumentalists with no knowledge of electronic components. Each performer received an identical D-Box in its original circuit configuration, along with an identical bag of components to use for hacking. We met each performer separately and asked them not to talk to one another about the D-Box during the study.

Participants were asked to prepare two solo performances over a period going from 20 to 62 days, according to their availability. (Due to personal circumstances, only 10 of 14 performers completed the study.) The first performance was a private recording, while the second one gathered all the participants in front an audience (80+ people), who voted for the best act. Both the performances were followed by an interview session. As incentive to participate, performers could keep their D-Boxes at the end of the study.

Participants were not required to modify the D-Box, but 8 out of 10 applied substantial hacks on the matrix and on the body of the instrument. In the next subsections, we analyse some of the most interesting hacks featured during the two performances, specifically focusing on the usage of feedback loops. Aside from showing different technical apporaches, these use cases suggest how the hackable design of the instrument could affect musicianship, also in relation to performers' skills and backgrounds.

5.1 Adaptive Pitch Hack

Participant ID 5 (P5), a guitar and piano player, decided to build his performances around a main melody and sing along with it. Despite his limited familiarity with electronics, he intuitively used an adjustable trimmer to modify the pitch circuit and adapt the instrument's pitch range to his needs. This allowed to find what he called a *"sweet spot"*, i.e. a value granting the pitch extension needed for the composed melody and a comfortable note distribution over the sensor to play it. Eventually, he added marks to the instrument's top plate which highlight the exact position of notes over the sensor, effectively acting as visual guitar frets.

After recording the first performance, P5 literally hacked his own modifications. While probing the matrix, he inadvertently connected the circuit responsible for sample selection to the pitch trimmer, by means of two buttons and a Light Dependent Resistor (LDR). The resulting circuit generates a bidirectional feedback between pitch and sample selection: when only one button is pressed, the sample remains unchanged and the LDR can be used to manipulate the pitch. When both buttons are pressed, the second sample is activated and its attack is enhanced by an abrupt pitch shift, acting as a fast envelope. These features, discovered by serendipity, inspired the addition of rhythmic and non-melodic sections to the piece.

P5 modified some constraints of his D-Box to play it as a familiar instrument and preserve his approach to music making. However, P5 declared that the unexpected behaviour of his hacks pushed him to experiment with new interaction techniques and musical genres he was not accustomed to, enlarging his creative horizons.

5.2 Skin Conductance and Audio Feedback

For the first performance, P3 prepared his D-Box to allow for chaotic control of sound texture through his own skin conductance. To do so, he detached the wires belonging to pitch and sample selection circuits from the cape's ADCs. While triggering the samples with one hand, he rhythmically touched, bent and twisted these floating wires with the other one. The sonic outcome of these interactions remarkably varied according to touch area, pressure and duration, and to the specific part of the wire where the manipulation occurred (e.g., the metallic tip, the centre of the plastic body, the end closer to the ADC). The performer made use of these bizarre techniques to control transitions between unpredictable sonic states of the system.

While rehearsing for the second performance, P3 completely discarded matrix and touch sensors from interaction. Instead, he focused on the two piezo microphones the D-Box is equipped with [20]. He detached the microphones from the inner plates of the instrument, moving them freely to play with audio feedback generated by varying the distance between the microphones and speaker.

Differently from P5, P3 found the instrument unique in its genre and did not try to capitalise his performances on known familiar schemes. Instead, he enjoyed a deep free exploration of the instrument, that led to the discovery of hidden affordances made available by the hardware and software interconnections of the D-Box. This is a strong example of instrument appropriation, far from the paths originally intended by the designers [21].

5.3 Cross-Wired Oscillator

P8's piece capitalises on different usages of the relaxation oscillator. He started focusing on playback speed control, by means of changing on the fly the capacitor in the oscillator circuit. He defined a set of discrete speed steps by preselecting capacitors of different sizes from the given bag of components. Then, to facilitate the swapping operation, he added an LED to light up the interior of the instrument.

In the second part of the piece, P8 rewired the circuit to make the pitch oscillate according to the playback speed. To introduce a further degree of control, he connected to the oscillator the signal coming from one of the two piezo microphones, attached on a side plate. Whenever the performer taps on the side plate, the abrupt voltage transient detected by the microphone freezes the oscillator circuit, slowing down the playback speed by 10 times. This state persists until the oscillator's capacitor is removed and repositioned, or changed with a different one.

The final configuration of P8's instrument included several other modifications, carefully chosen to support a personal playing style. For instance, he used pressure to quickly select sample and designed an original mapping using position on the second touch sensor to control loop points. Overall, the complexity of his hack was a real surprise, considering that the performer had no knowledge of electronics at all. The interview revealed that this fine tuning process helped P8 master the instrument in only few weeks of practice, and compose and perform a very original piece. His virtuosity was recognised by the audience as well, who selected P8 as best performer.

5.4 Additional Feedback Loops

P10 is an experienced hacker and circuit bender.He attached 4 potentiometers on top of a cardboard plate, configured as variable resistors (i.e. using just 2 terminals). One end of each potentiometer is fixed to a specific part of the matrix: the relaxation oscillator, the phase-shift oscillator, the sample selection circuit and the amplified audio output. The other end of each potentiometer can be connected anywhere on the matrix by means of crocodile clips.

P10's performance revolved around the indeterminacy triggered by wiring unrelated parts of the matrix on the fly and modulating the resulting feedback by means of the potentiometers. This generated additional oscillating circuits and unexpected software/hardware loops triggering time-varying and highly unstable electrical/sonic configurations.

Unlike the hacks previously introduced, this instrument is characterised by a remarkable degree of self-agency. The performer has almost no instantaneous control over the system state. The consequences of interaction are often non-linear and, even when interaction itself is suspended, sound output constantly evolves in search of a stable configuration.

P10 modified his D-Box using classic circuit bending techniques and tools, and during performances maintained a sonic register typical of this kind of practice. We gladly acknowledged that, for the first time, the performer managed to use his hacking skills to probe and modify a DMI, rather than cheap/old instruments and electronic toys.

6. DISCUSSION

The D-Box, though it presents a very simple interface to the performer, is a complex instrument internally, implementing oscillator bank resynthesis of sampled sounds with the potential for many degrees of freedom: pitch, volume, filtering, oscillator waveform, detuning, sample selection, playback speed, and loop points. Our previous study [21] suggests that breaking all these dimensions out to separate controls risks undermining the clarity of the instrument and reducing the motivation for discovering personal use cases.

On the other hand, having all extra parameters locked inside a complex piece of software limits the performer's ability to discover new sounds and techniques. Even if editing the code is possible, it lacks the immediacy and the exploratory value of pulling wires and swapping components. Several participants in our user study, including those without electronics expertise, rewired the D-Box live in performance, showing not only the immediacy but also the trust they placed in this modification process.

Compared to CV-based modular synths, instruments built on feedback loops will be inherently idiosyncratic. Feedback loops are not transferrable between one part of the instrument and the other, as different loops are implemented on different principles. There is also no guarantee that a small change an input produces a correspondingly small change in output: in certain regions, chaotic effects can emerge [18]. Modifications, however, are repeatable, with configurations that are reliably stable and others that are reliably chaotic. These results align with our goals to design a distinctive form of instrument control whose space of possibilities is not pre-scripted by the designer, nor by the musician. As suggested by the experience of the performers, this scenario is likely to benefit music variety, creativity and virtuosity.

7. ACKNOWLEDGMENTS

This research was supported by EPSRC grants EP/K032046/1 and EP/K009059/1. Thanks also to London Music Hackspace for their support and collaboration.

8. REFERENCES

[1] D. Arfib, J. Couturier, L. Kessous, and V. Verfaille. Strategies of mapping between gesture data and synthesis model parameters using perceptual spaces. *Organised Sound*, 7(2):127–144, 2002.

[2] F. Berthaut, M. T. Marshall, S. Subramanian, and M. Hachet. Rouages: Revealing the mechanisms of digital musical instruments to the audience. In *Proc. NIME*, 2013.

[3] P. Bloland. The electromagnetically-prepared piano and its compositional implications. In *Proc. ICMC*, 2007.

[4] J. Bowers and A. Haas. Hybrid resonant assemblages: Rethinking instruments, touch and performance in new interfaces for musical expression. In *Proc. NIME*, 2014.

[5] B. Caramiaux, N. Montecchio, A. Tanaka, and F. Bevilacqua. Adaptive gesture recognition with variation estimation for interactive systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 4(4):18, 2014.

[6] N. Collins. A solder's tale: Putting the "lead" back in "lead users". *Pervasive Computing*, 7, 2008.

[7] A. Di Scipio. A constructivist gesture of deconstruction. sound as a cognitive medium. *Contemporary Music Review*, 33(1):87–102, 2014.

[8] R. Fiebrink. *Real-time human interaction with supervised learning algorithms for music composition and performance.* PhD thesis, Princeton Univ., 2011.

[9] R. Ghazala. *Circuit Bending: Build Your Own Alien Instruments.* Wiley, 2005.

[10] A. Hunt and M. Wanderley. Mapping performer parameters to synthesis engines. *Organised Sound*, 7(2), 2002.

[11] C. Kiefer. Musical instrument mapping design with echo state networks. In *Proc. NIME*, 2014.

[12] J. Malloch, S. Sinclair, and M. M. Wanderley. Libmapper:(a library for connecting things). In *Proc. CHI Extended Abstracts*, 2013.

[13] A. McPherson. TouchKeys: capacitive multi-touch sensing on a physical keyboard. In *Proc. NIME*, 2012.

[14] A. McPherson and Y. Kim. The problem of the second performer: building a community around an augmented instrument. *Computer Music Journal*, 34(4), 2012.

[15] A. McPherson and V. Zappi. An environment for submillisecond-latency audio and sensor processing on BeagleBone Black. In *Proc. AES 138th Conv.*, 2014.

[16] O. Perrotin and C. d'Alessandro. Visualizing gestures in the control of a digital musical instrument. In *Proc. NIME*, 2014.

[17] T. J. Pinch and F. Trocco. *Analog days: The invention and impact of the Moog synthesizer.* Harvard University Press, 2002.

[18] D. Sanfilippo and A. Valle. Feedback systems: An analytical framework. *Computer Music Journal*, 37(2):12–27, 2013.

[19] G. Young. *The Sackbut Blues: Hugh Le Caine, Pioneer in Electronic Music.* National Museum of Science and Technology, 1989.

[20] V. Zappi and A. McPherson. Design and use of a hackable digital instrument. In *Proc. Live Interfaces*, 2014.

[21] V. Zappi and A. McPherson. Dimensionality and appropriation in digital musical instrument design. In *Proc. NIME*, 2014.