Journal of New Music Research
2010, Vol. 39, No. 3, pp. 189–202

The Magnetic Resonator Piano: Electronic Augmentation of an Acoustic Grand Piano

Andrew McPherson

University of Pennsylvania, USA

Abstract

This paper presents the magnetic resonator piano, a hybrid acoustic-electronic instrument extending the traditional grand piano. Sound is produced without loudspeakers using electromagnetic actuators to directly manipulate the piano strings, expanding its vocabulary to include infinite sustain, notes that crescendo from silence, harmonics, and a variety of timbres. A feedback-based approach senses string vibrations from a single pickup on the soundboard, using filters and phase-locked loops to generate signals which reinforce the natural motion of each string. Signal-routing hardware enables coverage of all 88 notes of the piano using a standard digital audio interface. This paper describes the hardware, software, and performing interface of the instrument and presents measurements of amplitude, frequency response, and phase-locked loop performance. Sounds generated by the actuators achieve similar amplitudes to the traditional piano, and while they cannot replicate the rapid attack produced by a piano hammer, they complement the percussive nature of the piano by enabling continuous legato control of amplitude, frequency, and timbre.

1. Introduction

Despite widespread adoption of electronic musical instruments, traditional acoustic instruments remain important to the performance of both old and new works, particularly in the classical concert hall. The advanced state of performance technique on traditional instruments has encouraged the development of electronic interfaces emulating or extending these instruments, particularly those based on the keyboard (Paradiso, 1997). However, many performers find that synthetic instruments still struggle to match the richness and nuance of their acoustic counterparts, which have been refined over centuries to reach their present form.

Recently, interest has been growing in hybrid acoustic-electronic instruments which produce sound by active electronic modification of an existing acoustic mechanism rather than by separate loudspeakers (Besnainou, 1999; Berdahl, Niemeyer, & Smith, 2008b). Electromagnetic actuation is commonly used to modify the resonant properties of an underlying acoustic instrument in real time, retaining the instrument's essential character while introducing new sonic properties impossible in a purely acoustic system. These efforts promise to unite the refinement of the best acoustic instruments with the flexibility of electronic synthesis.

The acoustic piano presents a particularly promising target for electronic augmentation, for both its ubiquity and its surprising versatility. The pianist controls relatively few parameters: a note can be struck with variable velocity, but then cannot be modified until its release. The pedals allow limited opportunities to modulate the sound of the instrument. Despite these limitations, accomplished pianists produce a tremendous range of expression. What could these players do if given the ability to continuously modulate the sound of their instrument?

On the other hand, a hybrid acoustic-electronic piano faces several challenges, most notably cost. Electronic augmentation of the piano, like other steel-stringed instruments, typically uses electromagnets to induce vibrations in the strings through ferromagnetic attraction. As concert grand pianos typically have 243 strings organized into 88 distinct pitches (Raichel, 2000), coverage

Correspondence: Andrew McPherson, Department of Music, University of Pennsylvania, 201 South 34th St., Philadelphia, 19104, USA. E-mail: andrew.p.mcpherson@gmail.com

DOI: 10.1080/09298211003695587

of the entire instrument potentially requires a large number of expensive sensors and actuators as well as considerable computing power. A separate challenge is logistical: a hybrid piano can offer many more dimensions of control than a standard acoustic piano, and these extra dimensions must be integrated into a coherent performance interface that does not impede traditional piano playing.

1.1 Project goals

This paper presents the magnetic resonator piano, a hybrid acoustic-electronic instrument intended as a robust, deployable system for active electronic control of piano strings. The goals of this project are:

(1) continuous control over amplitude, frequency, and timbre of as many as 88 distinct notes;
(2) reasonable hardware and computing costs;
(3) installable in most grand pianos, with setup possible in the rehearsal time before a concert;
(4) intuitive performance interface extending the keyboard model, integrating electronic and hammer-actuated sounds.

1.2 Related work

1.2.1 Electronic control of acoustic instruments

The use of electromagnets to induce vibrations in steel strings is well established. Commercial products for electric guitar such as the EBow (Heet, 1978) sense and reinforce existing string vibrations, creating notes which sustain indefinitely. More recently, Berdahl and Smith (2006), and Berdahl et al. (2008b) have developed tools for comprehensive control of string vibrations using classical feedback topologies, and applied this work to the creation of a hybrid electric guitar (Berdahl, Niemeyer, & Smith, 2008a) which uses collocated optical sensors and electromagnetic actuators to induce novel sustain and damping effects in each string. Boutin and Besnainou have applied similar active control strategies to a xylophone bar (2008a) and a violin bridge (2008b). Stability of feedback control typically requires very low processing latency; Dozio and Mantegazza (2007) provide guidance for the implementation of real-time control systems using general-purpose microprocessors. Lee, Berdahl, Niemeyer, and Smith (2008) present one such system using a standard PC, a specialized data capture card and a real-time operating system, achieving latency as low as 24 μs at a 40 kHz sampling rate. Other implementations use specialized embedded processors (Boutin & Besnainou, 2008a).

Electromagnetic manipulation of piano strings has previously been explored by Berdahl, Backer, and Smith (2005) and Bloland (2007). Their instrument, the Electromagnetically-Prepared Piano (EMPP), uses twelve solenoid electromagnets placed over selected strings. Each electromagnet is driven by an audio amplifier, whose signal originates from one channel of a 12-channel DAC. The EMPP does not use a feedback-based approach; rather, a Max/MSP patch generates periodic waveforms, filtered noise, and recorded samples. The resulting array of sounds blends the natural resonance of the piano with a purity often associated with electronic synthesis.

1.2.2 Expanded keyboard interfaces

Historically, certain keyboard instruments have offered dimensions of performer control beyond those of the acoustic piano. On the clavichord, continued pressure on a key can be used to bend the pitch of a note slightly upward. The Ondioline, an electronic instrument designed in 1941, allows vibrato effects by lateral motion of the entire keyboard and variations in volume by downward pressure (Paradiso, 1997). The desire for expanded control has extended to recent digital controllers. Some MIDI keyboards are equipped with aftertouch, which continuously senses pressure either globally on the keyboard, or individually on each key. Freed and Avizienis (2000) have developed a keyboard which reports the continuous position of each key. Moog and Rhea (1990) created a keyboard which senses the location of a finger touch on the key surface and allows each key to be displaced forwards or backwards. These expanded keyboard interfaces build on well-established performance technique while offering means of controlling more complex musical processes.

1.3 Paper organization

Section 2 presents an overview of the hardware and software design of the magnetic resonator piano, followed by a detailed discussion of each component. Acoustic measurements of amplitude, frequency, and phase tracking are presented in Section 3 accompanied by a discussion of their musical implications. Section 4 discusses efforts to create a comprehensive performance interface to control the new sounds, and summarizes the current state of the project.

2. Design

2.1 Overview

Figure 1 shows a block diagram of the magnetic resonator piano. The design consists of four major components.

(1) **Sensing:** a piezo pickup on the soundboard detects the sum of all string vibrations.
(2) **Signal Processing:** bandpass filters and phase-locked loops isolate and track each pitch, generating waveforms synchronized to the motion of each string.

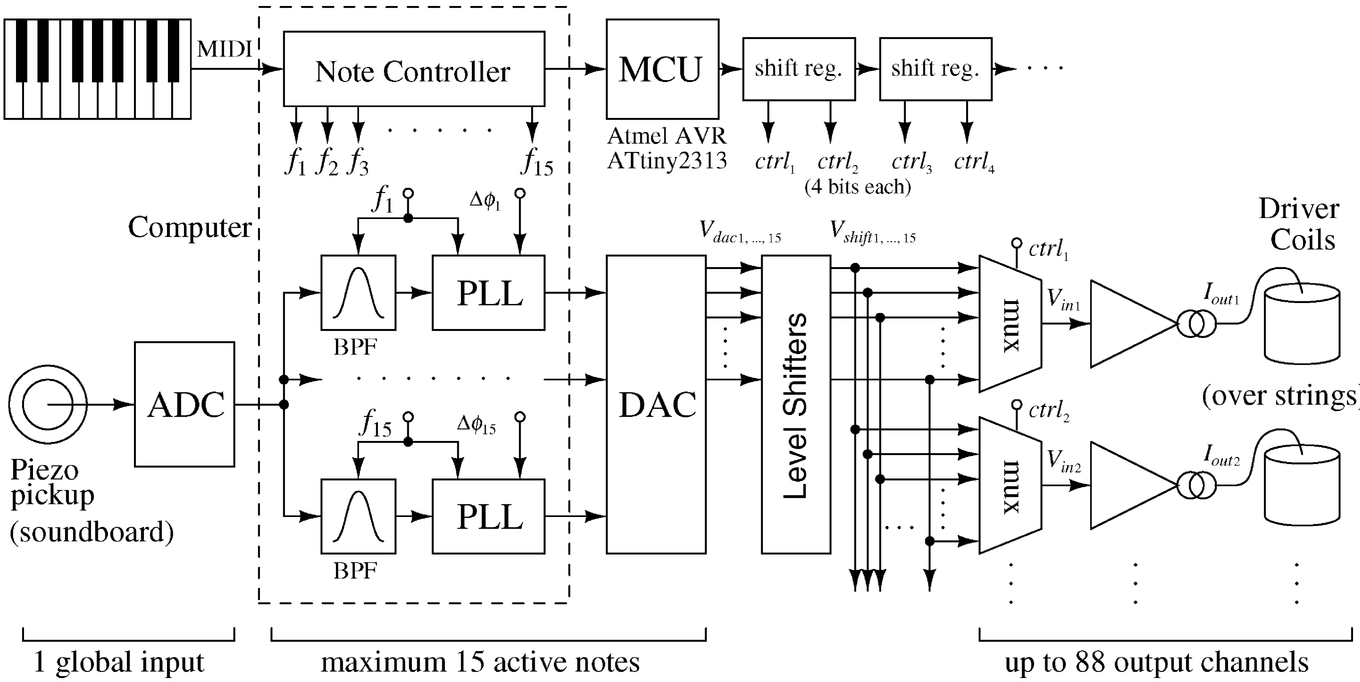

Fig. 1. Block diagram of the magnetic resonator piano.

(3) **Routing:** hardware multiplexers route audio from each DAC channel to the appropriate string. The number of DAC channels determines the maximum polyphony (up to 15 voices), but coverage of all 88 pitches is possible regardless of DAC size.
(4) **Actuation:** electromagnets produce vibrations in each string by ferromagnetic attraction. The mechanical coupling from string to soundboard transfers these vibrations to the pickup, completing a large-scale feedback loop.

This design reflects trade-offs among precision of control, cost, and setup time. True active feedback control in the manner of Berdahl et al. (2008b) or Boutin and Besnainou (2008a) would require a separate sensor for each piano string, but this was judged to be prohibitively expensive and too logistically complex to rapidly install. The phase-locked loop system presented here allows more limited adaptive control using a single global pickup and off-the-shelf audio hardware. At present, the prototype instrument covers 30 strings, but expansion to all 88 strings is readily possible.

In the following sections, actuation will be described first, followed by sensing, signal processing, and routing.

2.2 Electromagnetic actuators

Solenoid electromagnets are used to induce vibrations in the piano strings. An adjustable bracket holds one electromagnet (coil) for each note of the piano, suspending the coils above the strings at a gap ranging from 1/16″ in the upper registers to 3/16″ or more in the bass. The bracket rests on the steel crossbeams of the grand piano, allowing easy installation and removal, with the coils positioned immediately behind the dampers.[1]

Actuation operates on the principle of ferromagnetic attraction, whose application to piano strings is detailed in Berdahl et al. (2005). By modulating the current through a coil at the natural frequency of its associated string (or one of its harmonics), the string will be induced to vibration.

The ideal actuator will produce strong magnetic fields and modulate them as rapidly as possible. In practice, these two goals are opposed to one another. The magnetic field produced by a current I through an ideal solenoid with N turns of wire, length ℓ, cross-sectional area A, and core permeability μ is given by:

$$B = \mu N I / \ell. \tag{1}$$

On the other hand, given a fixed power supply voltage V, we can express the maximum current slew rate as:

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \frac{V}{L} = \frac{V\ell}{\mu N^2 A}. \tag{2}$$

Since the width of a group of piano strings puts a floor on the value of A, the designer must choose between high-current, rapidly-responding coils (small N) or lower-current coils which respond more slowly (large N).

[1] On most pianos, two or three strings in the mid-bass register are physically obscured underneath the wound bass strings; these actuators are typically placed on a separate bracket forward of the dampers.

For this project, several prototypes were explored, including commercial solenoid electromagnets as used in Bloland (2007). Best performance within reasonable power and thermal limits was obtained by a solenoid consisting of 600 turns of 30AWG wire wound on a 3/8″ steel threaded rod which acts as both a high-permeability core and a mounting method. Measurement of 30 coils yielded an average inductance of 18.9 mH and a DC resistance of 9.3 Ω. Average field strength for $V = 5$ V ($I \approx 535$ mA) at a distance of 1/8″ was measured to be 197 gauss. This design also produced the subjectively best musical results across several piano registers, although future work may investigate a different coil design for the copper-wound bass strings, where only a fraction of the total mass responds ferromagnetically.

2.3 Amplifiers

Each coil is driven with its own amplifier, shown in Figure 2. Transconductance amplifiers are used since a coil's field strength is proportional to current rather than voltage; the gain of these amplifiers is given by:

$$\frac{I_{\text{out}}}{V_{\text{in}}} = \frac{1}{R_{\text{sense}}(1 + R_2/R_1)}. \tag{3}$$

2.3.1 Theory of operation

Op-amp U1a drives MOSFET Q1, configured as a source follower. The voltage across R_{sense} is proportional to the output current; U1b amplifies this voltage by a factor of $(1 + R_2/R_1)$. The gain in Equation 3 is inferred by noting that with feedback, the two inputs of U1a must have equal voltage. Resistor R_p is required for stability. Without it, since coil impedance increases with frequency, the voltage gain of the system would approach infinity at high frequencies. R_p creates a first-order high-frequency roll-off in I_{out} with corner frequency:

$$f = \frac{R_s + R_p}{2\pi L}. \tag{4}$$

Equation 2 shows that the current slew rate is limited by coil inductance and supply voltage. It also shows that decreasing the current requires a negative voltage. This is partly provided by the voltage drop $I_{\text{out}}R_s$ and the offset $V_{\text{gs}} \approx 4$ V from gate to source of Q1, but to obtain fast turnoff near $I_{\text{out}} = 0$, a negative op-amp supply is required. This supply only needs to provide sufficient current to power the op-amps, and never powers the load. The unusual asymmetrical power rails reflect a balance of high output power and fast turn-off time, constrained by the maximum differential supply voltage of U1. With the values shown, V_{out} ranges from approximately -13 to $+14$ V.

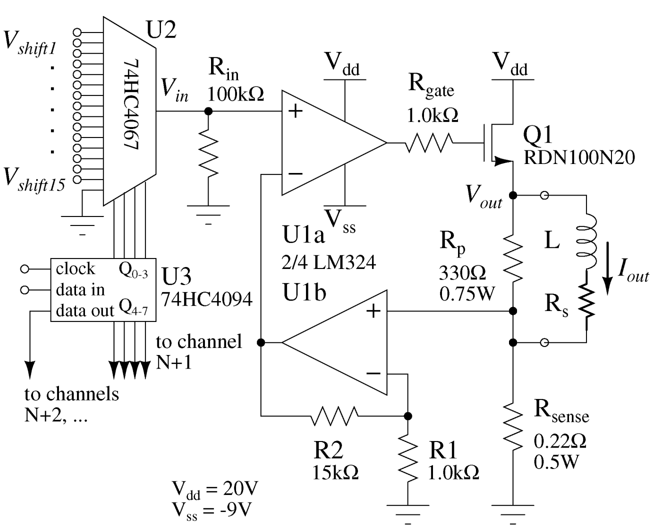

Fig. 2. Amplifier and multiplexer for each piano note.

2.3.2 Practical considerations

The amplifier of Figure 2 is single-polarity. An applied magnetic field, regardless of polarity, exerts an attractive force on a ferromagnetic object; for this reason, bipolar output would produce an unwanted frequency-doubling effect as both positive and negative half-waves attract the string. Bloland (2007) instead uses permanent magnets to create a fixed offset. Tests showed that with single-polarity amplifiers, incorporating permanent magnets into the coil design had no significant effect on the amplitude of sounds produced.

Since 88 amplifiers are required to cover the whole piano, the design aims for minimum possible cost, requiring only one inexpensive power device, half of a quad op-amp and several passive components. Greater efficiency is possible with switching amplifiers, with a corresponding increase in cost and complexity.

2.4 Piezo sensor

For control purposes, the motion of each string should be measured at the point of its interaction with the actuator. Typically, magnetic or optical sensors are used for this purpose; however, EM interference from the actuators precludes the use of magnetic sensors, and optical sensors are difficult to install in the piano given the close spacing of the strings and limited access to the space beneath them.

Instead, a single piezo pickup affixed to the soundboard detects the sum of all string vibrations. Losses and delays in transmitting vibration from string to soundboard result in attenuation and phase offsets between the motion of each string and the signal detected by the pickup. The motion of the set of strings for piano note k can be described as a sum of sinusoids whose amplitudes

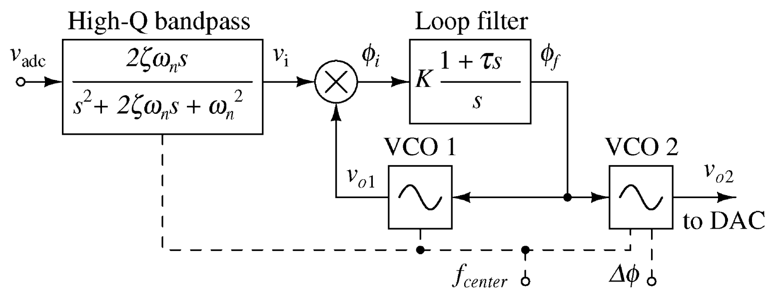

Fig. 3. Input filter and PLL system for one active note.

change in time, but whose frequencies remain approximately constant:

$$\sum_{i=1}^{N} A_{k,i}(t)\cos(\omega_{k,i}t). \tag{5}$$

Ignoring transmission nonlinearities, the pickup signal represents the sum of all strings, scaled in amplitude and shifted in phase:

$$v_{\text{piezo}}(t) = \sum_{k=1}^{88}\sum_{i=1}^{N} g_{k,i}A_{k,i}(t)\cos(\omega_{k,i}t + \phi_{k,i}). \tag{6}$$

The frequencies $\omega_{k,i}$ are approximately known but vary slightly depending on piano tuning.[2] Each phase offset $\phi_{k,i}$ is unknown but time-invariant and is discovered during a calibration process described in Section 2.5.2. $g_{k,i}A_{k,i}(t)$ is irrelevant to the signal processing as long as it is sufficiently above the noise floor.

2.5 Signal processing

The pickup signal $v_{\text{piezo}}(t)$ connects to an ADC. Signal processing is performed digitally at a rate of 44.1 kHz using Csound. The signal processing aims to produce waveforms which remain in phase with the motion of each string at the point of its interaction with the actuator, thereby reinforcing the natural string motion. Two challenges must be addressed: first, each frequency must be isolated from the summed piezo signal, and second, compensation must be made for unknown delays in the system, including digital buffering and mechanical propagation from string to soundboard.

2.5.1 Frequency separation

Each active fundamental frequency is isolated using a second-order bandpass filter with $Q = 50$ (Figure 3). This isolation is far from perfect, as adjacent semitones are attenuated by only 15 dB, but this preliminary filtering provides a cleaner input to the subsequent processing.

[2]For any given string k, the frequencies $\omega_{k,i}$ will be slightly stretched with respect to a standard harmonic relation (Fletcher, Blackham, & Stratton, 1962).

While higher Q filters improve frequency isolation, they do so at the cost of group delay, defined as the derivative of phase response:

$$\tau_g = -\frac{d\phi(\omega)}{d\omega}. \tag{7}$$

For a second-order bandpass with centre frequency ω_n and $\zeta = 1/(2Q)$:

$$\tau_g = \frac{2\zeta\omega_n[(\omega_n^2 - \omega^2) + 2\omega^2]}{(\omega_n^2 - \omega^2)^2 + (2\zeta\omega_n\omega)^2} \tag{8}$$

$$\approx \frac{1}{\zeta\omega_n} = \frac{2Q}{\omega_n} \quad \text{for} \quad \omega \approx \omega_n. \tag{9}$$

For a constant-Q bank of filters, the low frequencies will particularly suffer from large group delay, producing dramatic phase shifts in response to miniscule alterations of input frequency. For $Q = 200$ at C3 ($\omega_n = 2\pi \times 130.8$ Hz), an adjustment of 1 cent (0.08 Hz) will create a phase shift of 13°; an adjustment of 5 cents (0.38 Hz) will create a phase shift of 66°. The calibration procedure adds a fixed phase offset to the output waveform, but this technique has little value for highly frequency-dependent phase offets.

2.5.2 Phase-locked loop tracking

The phase-locked loop (PLL) is the heart of the control system, determining important aspects of the instrument's musical character. The implementation in Figure 3 uses a variation on a classic analogue form described in Best (1993). The loop filter gain K and time constant τ are chosen for a balance between fast settling time and narrow lock range. Since the bandpass filters are imperfect, narrow lock range is important to prevent the PLL from spuriously locking to an adjacent semitone.

VCO1 output v_{o1} is a sinusoid in quadrature with the input v_i. VCO2 is configured identically to VCO1 with the exception of a fixed phase offset $\Delta\phi$, which is chosen experimentally during calibration to offset $\phi_{k,1}$ of Equation 6 plus offsets from digital latency. v_{o2} can assume any periodic waveform, enabling a variety of timbres.

Two control strategies are possible for dealing with multiple harmonic modes within a single piano string. In the first, a single PLL locks to the fundamental frequency of the string, and harmonically-rich waveforms are used for v_{o2}. Owing to frequency-stretching in the piano strings, the upper partials of v_{o2} will differ slightly in frequency from the natural string modes, but empirically, the response is still acceptably strong. The second approach uses separate filters and PLLs for each harmonic of the string, with each v_{o2} configured for sinusoidal output. This is more computationally expensive, but allows interesting multiphonic effects from single strings.

Calibration is an empirical process, though automatic procedures are under investigation. Each note is activated in succession with its damper lifted, and $\Delta\phi$ is adjusted for the loudest tone. A unique value of $\Delta\phi$ is stored for each note of the keyboard, or optionally for each harmonic of each note of the keyboard. Recalibration is only necessary if the pickup or actuators are moved. In performance, the amplitude and waveform of v_{o2} are continuously adjustable independent of the input signal characteristics, allowing a diversity of envelopes and timbres.

2.6 Signal routing

External analogue signal routing allows access to up to 88 amplifiers and actuators using many fewer DAC channels. Each actuator has a dedicated amplifier whose input is driven from a 16-channel analogue multiplexer (Figure 2). The multiplexer selects a DAC channel, with one multiplexer input reserved for muting the amplifier. Thus up to 15 channels can be used, allowing 15-note polyphony.

A microcontroller (MCU) maintains a mapping between DAC channels and amplifiers. The MCU drives a cascade of 8 bit shift registers, each of which controls two multiplexers. Complete keyboard coverage thus requires 88 coils and amplifiers, 88 multiplexers and 44 shift registers, while reducing the more costly DAC channels to 15. A smaller DAC can be used with a reduction in the maximum number of simultaneous notes.

At power-up, all amplifiers are muted. The MCU receives MIDI Control Change messages, where the controller number corresponds to a DAC channel and the value byte indicates the MIDI note number of the piano string. To break a connection and mute the corresponding amplifier, a control change is sent with a value byte of 0. The MIDI Control Reset message mutes all amplifiers. Mappings are generated by the computer, which assigns each new note to the lowest available DAC channel.

2.7 Physical setup

Figure 4 shows a picture of the magnetic resonator piano; Figure 5 shows a closer view of the electromagnetic actuators. The resonator system can be installed in any acoustic grand piano without modification to the piano.

Fig. 4. Physical setup of the magnetic resonator piano in a Steinway D concert grand.

Fig. 5. Close view of the electromagnetic actuators.

Aluminium brackets holding the actuators are placed across the steel crossbeams of the piano frame. Detachable ribbon cables connect the actuators to the amplifier and signal routing hardware, which rests on a table near the piano. When moving the system to a new piano, actuator height and spacing must be adjusted to accommodate the geometry of the instrument, a process that takes approximately two hours; however, once configured for a particular piano, reinstallation takes only a few minutes.

System setup also includes placing a Moog Piano Bar sensor on the piano keys and a MIDI keyboard atop the piano. The music rack is elevated on foam spacers to allow space for the MIDI keyboard. Details of the performance interface are discussed in Section 4.

Between installation of actuator and interface hardware and connection of all cables, setup typically requires 20 to 30 min, plus an additional 10 min for calibration and testing. In most situations, this process can be performed in the rehearsal time before a concert.

3. Measurements

The most important metric for any musical instrument is its utility as a creative tool, which is impossible to quantify. In its place, measurements of amplitude, frequency response, and phase tracking will be presented, and the ensuing discussion will highlight their musical relevance. In this section, the term ‘resonator’ denotes electronically-actuated sounds, in contrast to standard (hammer-actuated) piano notes.

3.1 Amplitude

Figure 6 shows amplitude envelopes for the resonators in several registers. Measurements were taken on a Steinway B piano with the lid open full-stick; an omnidirectional measurement microphone was placed just outside the instrument. Four cases were tested: resonator with optimal phase calibration, resonator with deliberately non-optimal phase calibration, resonator with feedback disabled (PLL loop gain $K = 0$), and an ordinary piano note (played roughly *forte*) for approximate comparison. Sinusoidal output waveforms were used for all resonator measurements.

Compared to notes played open-loop (without feedback), optimal phase calibration can produce a gain of up to 10 dB. Non-optimal phase calibration produces results with very low amplitude, 10–20 dB lower than the optimal calibration and typically lower than the open-loop case as well. Surprisingly, in certain cases (C3 and C5) the open-loop configuration produced higher amplitude than the well-calibrated PLL. This may be attributable to calibration error; it can occasionally be difficult to ascertain by ear which value

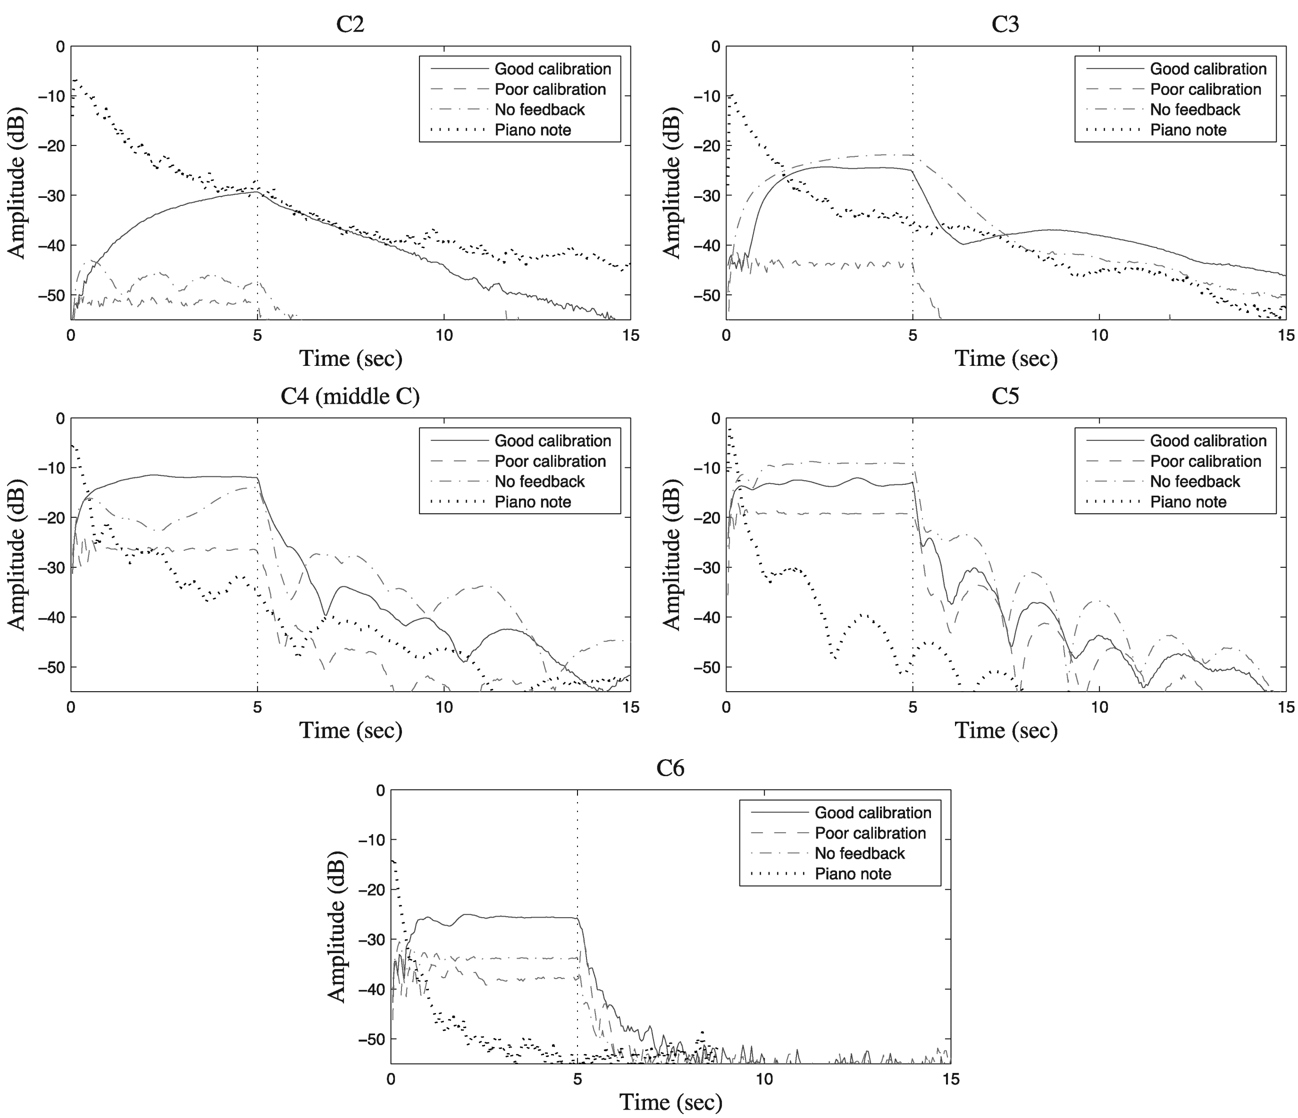

Fig. 6. Amplitude envelopes for three resonator configurations versus a standard piano note, tested in five registers. Resonator tones are played at maximum volume for five seconds; piano note is played approximately *forte*.

of $\Delta\phi$ produces the strongest tone, and a different choice may have resulted in higher amplitude. It is hoped that a future automatic calibration procedure eliminates these errors.

The amplitude of the resonator tones on C2 are both numerically and perceptually lower than those in higher octaves. C2 is a copper-wound string and hence its response to ferromagnetic actuation is weaker.

3.1.1 Comparison with piano

Though the initial impulse of a piano note can be stronger than the loudest resonator tone, the resonators can sustain at a higher level than the piano, which decays rapidly. Nonetheless, the piano is perceptually louder; psychoacoustically, a note that begins loudly is perceived as being loud even if it quickly decays (Weinreich, 1990). The piano gains additional perceptual advantage from its rich spectrum (see Section 3.2).

The rise times of the resonator tones scale inversely with the frequency of the string. It is common for low-pitched acoustic instruments to speak more slowly than those of higher pitch, but the effect here is particularly pronounced for C2, which after 5 s does not reach its full amplitude. These measurements indicate that the resonator cannot replicate the percussive nature of the piano, but can complement it with tones that grow slowly or maintain constant amplitude. Piano and resonator can be actuated simultaneously to produce an envelope with percussive attack and extended sustain.

3.2 Frequency response

3.2.1 Harmonics

A piano string can be modelled as a parallel set of high-Q bandpass filters, one for each harmonic (Berdahl et al., 2005). The string will respond strongly to signals at these

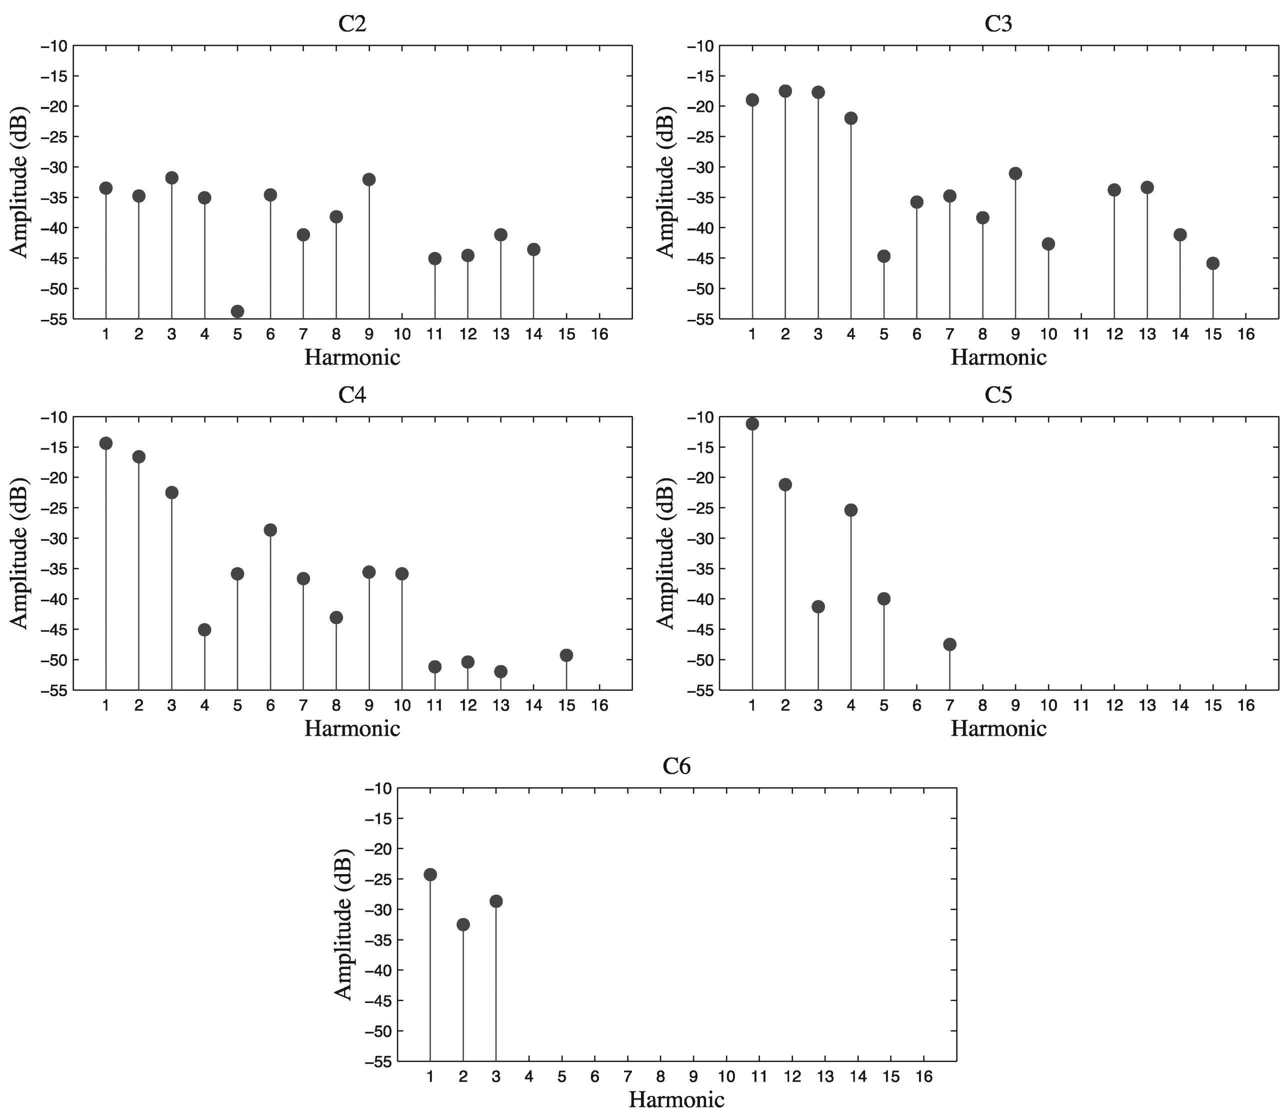

Fig. 7. Maximum strength of harmonics 1–16 for resonator tones in five registers.

frequencies and negligibly to everything in between. Frequency response can thus be represented by the strength of individual harmonics. Figure 7 shows the maximum amplitude of the first 16 harmonics of several strings. For each string, higher harmonics generally have lower amplitudes, reflecting reduced actuator performance owing to coil inductance. Qualitatively, eight or more harmonics of the lower strings have sufficient amplitude to project individually in a concert space, while four or fewer are usable in the upper octaves.

Each data set of Figure 7 has a harmonic with greatly reduced amplitude compared to its neighbours (5th partial for C2 and C3, 4th partial for C4, 3rd partial for C5). A harmonic is masked when the actuator is placed at one of its nodal points along the string. The actuators were placed immediately behind the dampers, which was approximately 1/5 of the string length in the bass but less than 1/3 of the length in the treble. Placement closer to the midpoint is possible, subject to constraints of physical accessibility.

3.2.2 Spectral measurements

Figure 8 compares the spectra of resonator and piano tones on middle C (C4). Using a sinusoidal output, the resonator is capable of more intense energy at the fundamental frequency than the piano, but the actuators are mostly ineffective above the 8th harmonic. By contrast, the piano produces considerable energy extending beyond the 16th partial. This spectral richness underlines an important timbral distinction between conventional piano and resonator and demonstrates why the piano is perceptually louder despite similar numerical amplitudes.

The waveforms in Figure 8 are simple, consisting of one or two summed sinusoids, yet the output spectrum is quite complex. There are several places for distortion in the system. Most significantly, when operating at maximum amplitude, clipping and slew rate limitations in the amplifiers significantly change the shape of the output waveform. Additionally, the ferromagnetic force on the piano string is heavily dependent on the distance from string to actuator, which varies over the course of an oscillation. Finally, transmission of vibrations in the bridge and soundboard may introduce distortion. The claim is made here that this distortion is desirable. Spectral purity has little musical relevance in this context; rather, like distortion in guitar amplifiers, the nonlinearities in the system introduce a complex yet intuitive relationship between amplitude and timbre.

3.3 Phase tracking

3.3.1 Instantaneous frequency

Figure 9 plots instantaneous frequency for the same test cases as Figure 6. These results indirectly illustrate

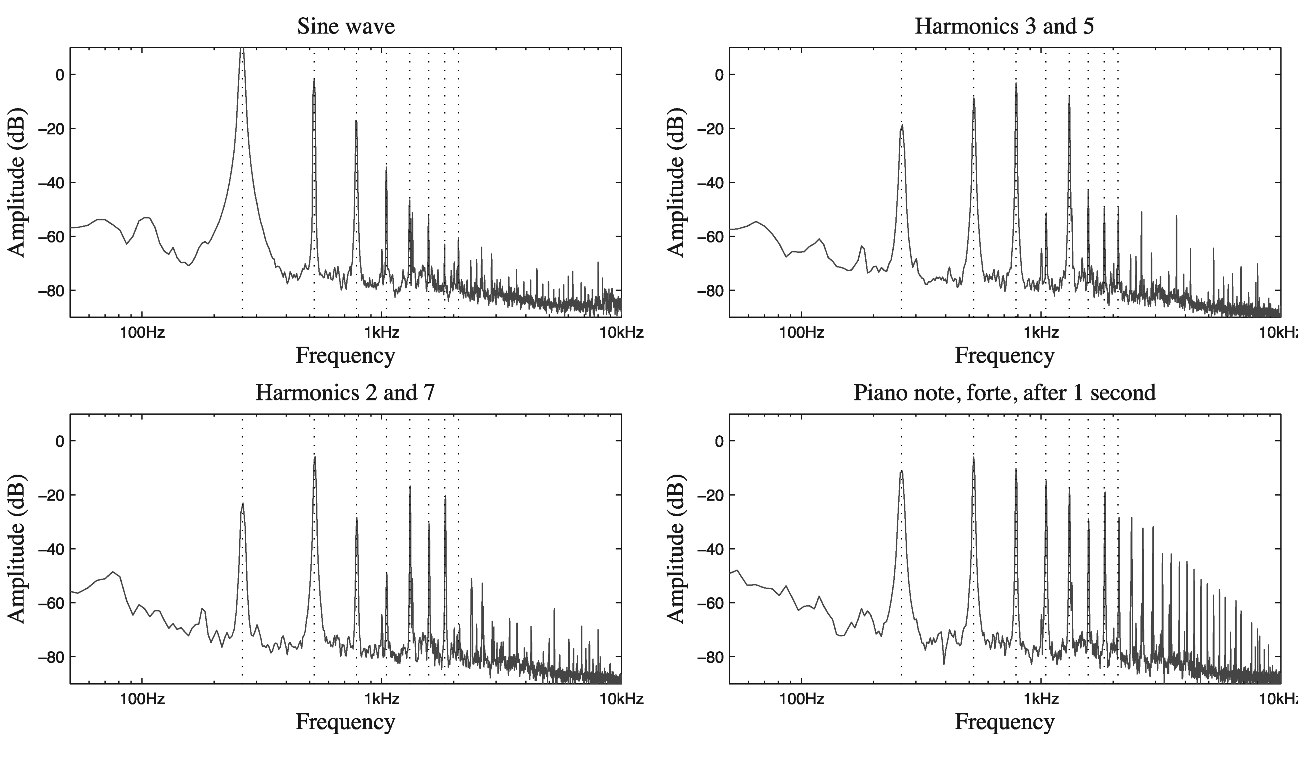

Fig. 8. Spectra for three resonator waveforms compared with a standard piano note, middle C (C4). 8192-point FFT, Hanning window. Vertical lines indicate the first 8 harmonics of the string.

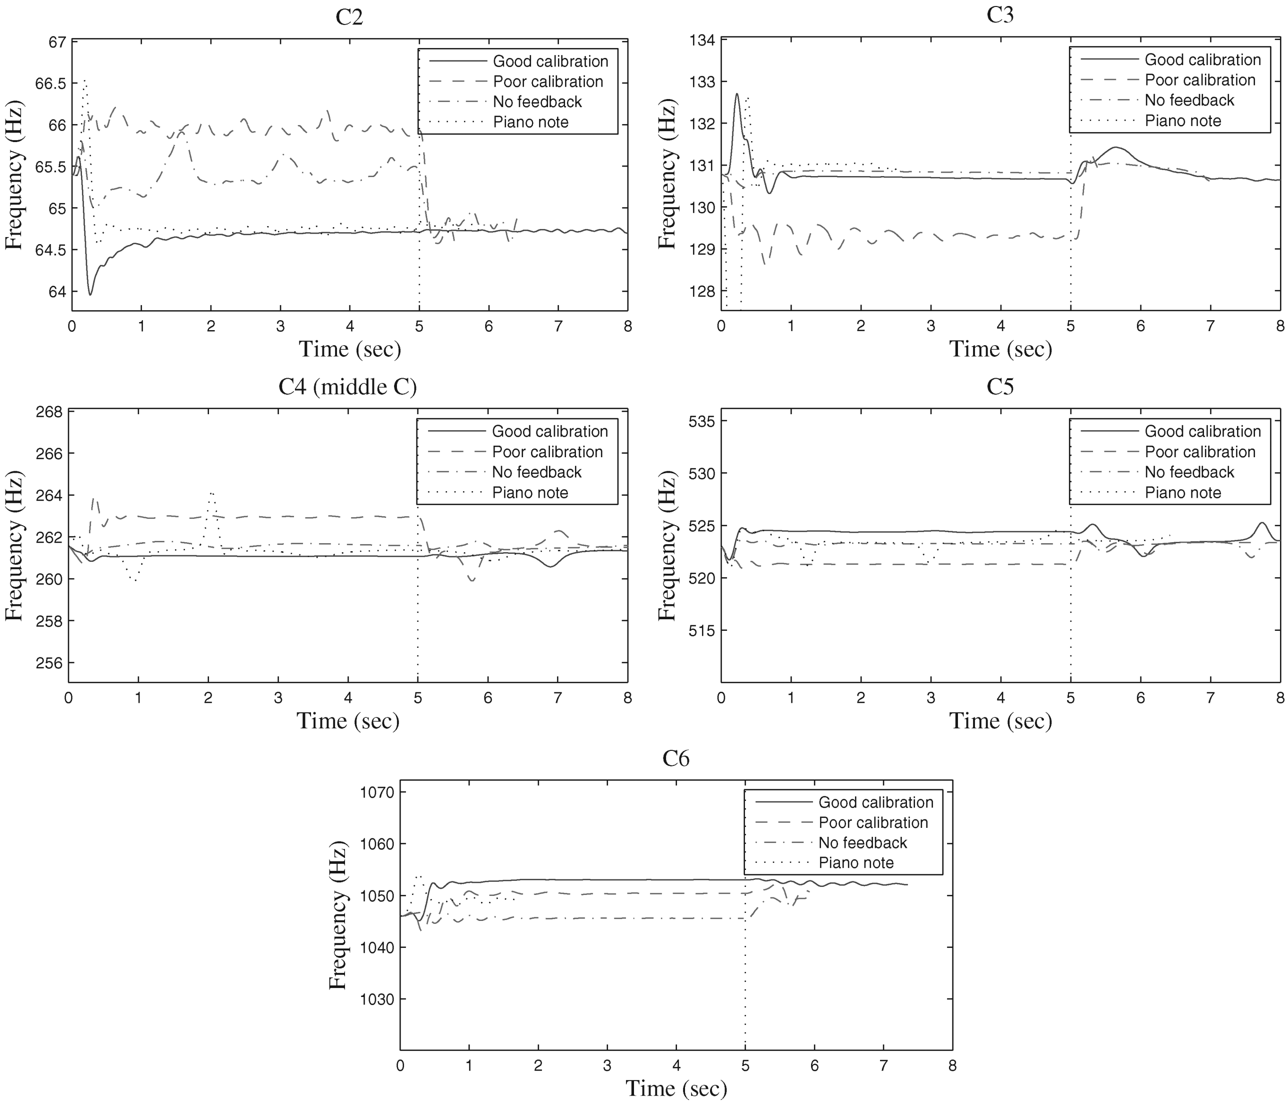

Fig. 9. Instantaneous frequency measurements for three resonator configurations and a standard piano note. Resonator tones last 5 s. Only the first partial of each note is measured.

phase-tracking behaviour: instantaneous frequency is the derivative of phase, and in a well-configured system, it would be expected that the frequency would remain similar to that of the piano note. Indeed, the well-calibrated PLL maintains consistently better frequency-matching properties compared to the other two resonator configurations.

Without feedback, the actuator output frequency will always conform to an equal-tempered scale tuned to A440. Because the piano string may be tuned differently, interactions between actuator and string can produce beating, as seen in the case of C2 (reflected also in the amplitudes of Figure 6). Beating is also evident in the poorly-calibrated PLL, which produces a warbling effect that is musically compelling though difficult to predict or control. In certain cases, the natural piano decay also evinces beating, owing to variations in tuning among multiple strings for a note.

In Figure 6, open-loop tones often reach their steady-state amplitude more rapidly than the well-calibrated PLL. Figure 9 shows why: at the beginning of the notes with phase tracking, the frequency is unsteady for a fraction of a second while the PLL settles. This fluctuation has aesthetic value, enlivening the onset of each note in a way that makes it subjectively sound more acoustic than electronic. The above discussion indicates that all three PLL configurations (well-calibrated, poorly-calibrated, and open-loop) have distinct musical properties which can find use in various contexts.

3.3.2 Interference rejection

Section 2.5.1 noted that the bandpass filters produce imperfect attenuation of neighbouring semitones. Thus, the input to the PLL may contain significant energy not only at the desired frequency, but at the frequencies a semitone above and below it. Figure 10 shows simulations of PLL behaviour given an input of the form:

$$s(t) = \cos(\omega_n t) + B\cos\left(\sqrt[12]{2}\,\omega_n t\right). \quad (10)$$

The second term of this equation represents an interfering signal one semitone higher. In these tests, the PLL has a centre frequency 1 Hz below ω_n to simulate real-world conditions where the string frequency is not precisely known in advance. For $B = 0$, the PLL readily acquires and maintains lock. As B grows, the tracking signal exhibits beats of growing amplitude until, for $B > 1$, the PLL spuriously locks to the interfering signal. Thus, with input filters having $Q = 50$, an interfering signal one semitone away which is 15 dB stronger than the desired signal will cause spurious lock.

Spurious lock, from the performer's perspective, results in a note that fails to play, since the actuator is operating at the wrong frequency for the string. The PLL can return to proper operation if the interfering signal reduces sufficiently in amplitude, which creates the impression of an unpredictable delayed onset. These behaviours, of course, are unacceptable.

No single set of PLL parameters will maintain lock under strongly unfavourable interference conditions, so a dynamic loop gain strategy is used to prevent spurious lock. Additional bandpass filters are operated one semitone above and below the target frequency, and as the output of these filters approaches the level of the main filter, the loop gain is gradually reduced to zero. The effect of this strategy is to disable phase tracking whenever a strong interfering signal is present and re-enable it when the interference subsides. The transition between these two states is gradual and generally not audible.

3.4 Limits of active control

The use of a single centralized pickup and standard audio hardware with several milliseconds of latency substantially reduces system cost and complexity, but it places limits on the precision of active string control. Unlike other hybrid instruments (Berdahl et al., 2008a), active damping of string vibrations is not possible. In general, it is easier to add energy to a system than to remove it, and delays and interference preclude the generation of a signal whose phase relationship to the string motion is consistent enough to actively cancel it.

However, the mechanics of the piano make active damping challenging under any circumstances. Piano notes with three strings can feasibly have only one pickup and actuator, making the control problem more difficult. Furthermore, each string can vibrate horizontally or vertically, with the interaction of these modes producing the piano's characteristic sharp attack and long decay (Weinreich, 1990). Though vibrations on the vertical axis couple more strongly to the soundboard than those on the horizontal axis, both

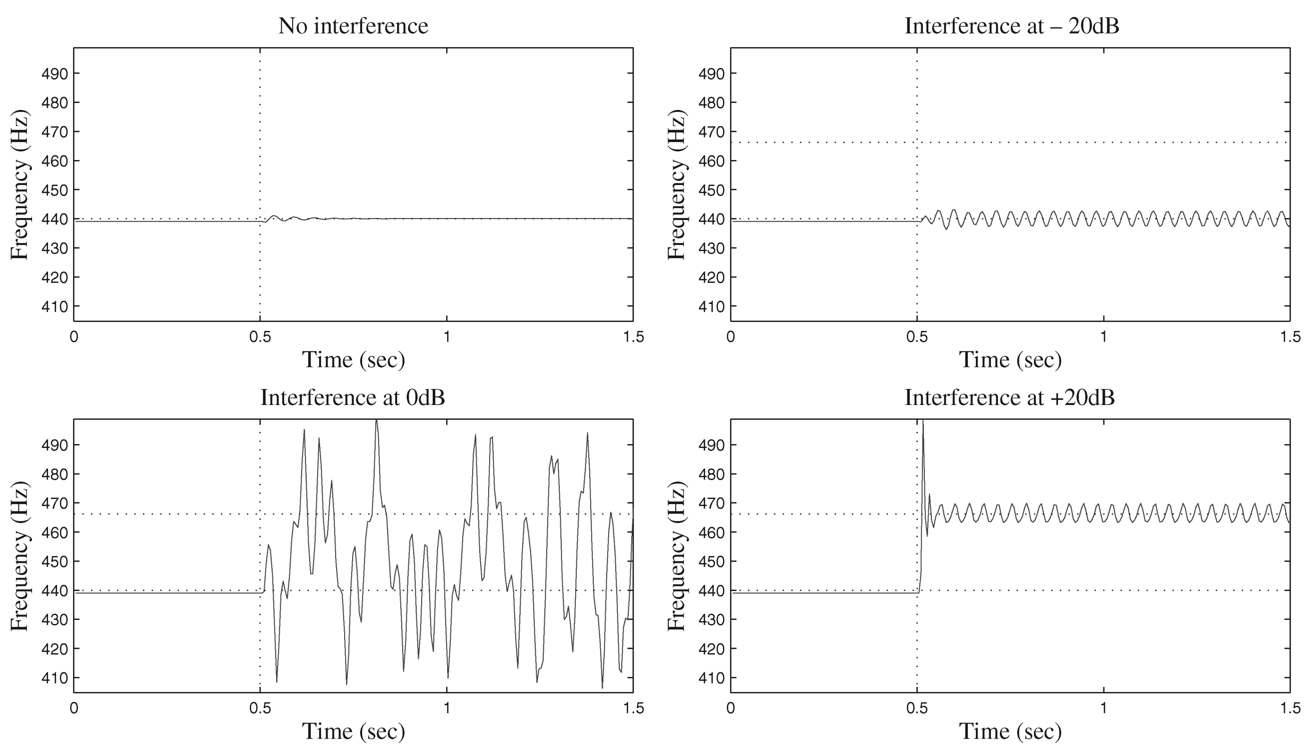

Fig. 10. Simulation of PLL behaviour at A4 in the presence of interference a semitone away (B♭4).

contribute to the radiated sound. The orientation of the actuators only allows control on the vertical axis, making perfect cancellation of all radiated sound impossible.

4. Performance interface

The identity of a musical instrument lies as much in its performance interface as in its sound production mechanism. This section summarizes current efforts to create an integrated performance interface extending traditional piano technique.

4.1 Goals

The ideal performance interface should have the following three characteristics:

(1) continuous control of amplitude, frequency, and timbre for each note (within the physical constraints of the piano string) using natural physical gestures;
(2) seamless integration with the piano keyboard, unifying piano and resonator sounds even within individual notes;
(3) compatibility with existing piano performance technique, in no way impeding traditional piano playing.

4.2 Keyboard-based approach

Two keyboards are used to control the extended vocabulary of the magnetic resonator piano. Except on special MIDI-enabled pianos such as the Yamaha Disklavier or Bösendorfer SE (Moog & Rhea, 1990; Coenen and Schäfer, 1992), the primary piano keyboard is the only method of triggering hammer-actuated piano sounds. The goal of seamless acoustic-electronic integration thus requires that the resonators also be controlled from the primary piano keyboard. This is accomplished by a Moog Piano Bar (Piano Bar, 2005), which enables the piano keyboard to act as a MIDI controller. Each key press triggers a resonator event for the corresponding string.

Some sounds, such as notes which grow from silence, do not use the hammer action at all. These sounds can be triggered from a secondary MIDI keyboard placed above the piano keyboard in the manner of organ manuals. A pianist can easily play with one hand on each keyboard. When playing on the secondary keyboard, it is necessary to use the other hand or the pedal to lift the dampers on the strings to be played. This function could be automated on a MIDI-enabled piano.

In the current system, harmonics, timbres, and dynamic envelopes are selected in advance like programs on a MIDI synthesizer or stops on an organ. Limited real-time modulation of amplitude or timbre is possible using a MIDI control wheel on the secondary keyboard. Although this configuration does not provide true continuous control in every dimension, it allows performance of a wide range of sounds using standard keyboard technique.

4.3 Beyond the conventional keyboard

Electromagnetic actuation permits several dimensions of continuous string control, but conventional keyboard interfaces report note onsets and releases as discrete events, providing only a single velocity measurement for each note. On the other hand, this instrument aims to be an extension rather than a replacement of the piano, so a performance model based in keyboard technique is desirable. The magnetic resonator piano is thus a natural application for the extended keyboard interfaces discussed in Section 1.2.2.

Ongoing work explores adapting these extended keyboard paradigms to the piano keyboard itself, providing the performer with additional dimensions of control between note onset and release. Continuous control over actuator behaviour will allow the performer to shape the expression of notes and phrases in a way impossible on the acoustic piano, enabling greater differentiation among voices and new forms of legato connection by matching amplitude between the end of one note and the onset of the next.

Additional control dimensions alone do not create an intuitive interface. Depending on the context, key motion could be used to control the amplitude, frequency, or timbre of the resulting notes. A truly context-aware interface might also control actuator behaviour in response to groups of notes to shape the sound of entire phrases. Thus, developing mappings between continuous keyboard input and parameters of actuation will be an important area of future work.

5. Musical discussion

The acoustic measurements of Section 3 alone do not entirely describe the aesthetics of the instrument, much less its potential as a creative tool. This section presents musical observations which, though inherently subjective, suggest potential uses of the magnetic resonator piano.

5.1 Aesthetic observations

Electronically-actuated tones have an ethereal, mellow character owing to their slow attack and spectral purity; at the same time, their radiation within a concert space closely resembles that of the acoustic piano, blending naturally with hammer-actuated notes. Sympathetic vibration contributes heavily to the perception of resonator tones: when a single note is sounded electromagnetically without other dampers lifted, the result sounds flat and synthetic in comparison to the depth and richness of the same tone sounded with pedal. Therefore, the resonator system is at its best in musical contexts that admit the use of pedal and freely-sustaining notes.

When electromagnetic actuation is used sparingly in a primarily hammer-actuated passage, the piano can seem to grow in size and intensity without fundamentally changing in character. Used more heavily, the piano begins to take on the sustaining characteristics of an organ or woodwind instrument, pushing its percussive nature to the background even when hammer strikes are used to articulate the beginning of each note.

5.2 Compositional experiences

The author has used the magnetic resonator piano in two compositions: *Secrets of Antikythera* and *d'Amore*, both composed 2009. *Secrets of Antikythera* is scored for a single pianist playing both traditional piano and resonators using the dual-keyboard interface described in Section 4.2. After a first movement consisting entirely of electronically-actuated sounds, the subsequent movements tightly integrate traditional piano technique and electromagnetic actuation. Compositional techniques include:

- doubling piano and resonator to create notes with extended sustain. This is used particularly in the bass register to create long pedal tones;
- doubling piano with harmonically-rich resonator waveforms to modify the timbre of the instrument and distinguish among voices;
- generating a glissando through a string's harmonic series shortly after it has been struck with the hammer, creating notes that begin normally but evaporate into the upper octaves;
- following a resonator-induced crescendo from silence with a staccato hammer strike, simulating a time-reversed piano note.

By contrast, *d'Amore* is scored for viola with automatically-controlled magnetic resonator piano. The resonator effects simulate the viola d'amore, a Baroque-era instrument with a collection of strings under the bridge whose only purpose is to vibrate sympathetically. In *d'Amore*, the piano pedal is fixed down, and selected notes from a pickup placed on the viola activate the corresponding piano string, creating resonant harmonic clouds behind the viola part. No traditional piano sounds are used in *d'Amore*, and no pianist is required.

In both pieces, the resonator blends naturally with the acoustic sounds of each instrument, suggesting that integration into ensemble settings would be straightforward. The delicacy of many resonator effects further suggests that without amplification, the instrument would be more successful in chamber ensembles than orchestral settings.

5.3 Conclusion

This paper has presented the magnetic resonator piano, a hybrid acoustic-electronic instrument augmenting the capabilities of a traditional grand piano. The project aims to unlock new degrees of expressive freedom beyond those offered by the fundamentally percussive nature of the piano. An electromagnetic actuation system has been presented which allows real-time control of the motion of each piano string. This system has been designed for versatility, low cost, and ease of installation. Sounds are controlled from the piano keyboard or a separate electronic keyboard, allowing control of a wide variety of sounds using standard keyboard technique. It is hoped that this instrument will be a valuable creative tool for both performers and composers.

Acknowledgements

The author gratefully acknowledges Youngmoo Kim and Daniel Lee for helpful discussions during the design and construction of this project. Funding for this project was generously provided by an interdisciplinary fund sponsored by the Graduate and Professional Student Assembly and the Office of the Provost at the University of Pennsylvania.

References

Berdahl, E., Backer, S., & Smith, J.O. (2005). If I had a hammer: Design and theory of an electromagnetically-prepared piano. In *Proceedings of the International Computer Music Conference*, Barcelona, Spain.

Berdahl, E., Niemeyer, G., & Smith, J.O. (2008a). Active control of a vibrating string. In *Second Joint Conference of the Acoustical Society of America (ASA) and the European Acoustics Association (EAA)* (also known as *Acoustics '08, #3120*), Paris, France.

Berdahl, E., Niemeyer, G., & Smith, J.O. (2008b). *Feedback control of acoustic musical instruments* (Technical report STAN-M-120). Stanford University: CCRMA.

Berdahl, E., & Smith, J.O. (2006). Some physical audio effects. In *Proceedings of the 9th International Conference on Digital Audio Effects*, Montreal, Quebec.

Besnainou, C. (1999). Transforming the voice of musical instruments by active control of the sound radiation. In *Proceedings of the International Symposium on Active Noise and Vibration Control*, Fort Lauderdale, Florida.

Best, R. (1993). *Phase-locked Loops: Theory, Design, and Applications*. New York: McGraw-Hill.

Bloland, P. (2007). The electromagnetically-prepared piano and its compositional implications. In *Proceedings of the International Computer Music Conference*, Copenhagen, Denmark.

Boutin, H., & Besnainou, C. (2008a). Physical parameters of an oscillator changed by active control: Application to a xylophone bar. In *Proceedings of the 11th International Conference on Digital Audio Effects*, Espoo, Finland.

Boutin, H., & Besnainou, C. (2008b). Physical parameters of the violin bridge changed by active control. In *Proceedings of the Acoustics '08 Conference*, Paris, France.

Coenen, A., & Schäfer, S. (1992) Computer-controlled player-pianos (review). *Computer Music Journal*, *16*(4), 104–111.

Dozio, L., & Mantegazza, P. (2007). General-purpose processors for active vibro-acoustic control: Discussion and experiences. *Control Engineering Practice*, *15*(2), 163–176.

Fletcher, H., Blackham, E.D., & Stratton, R. (1962). Quality of piano tones. *Journal of the Acoustical Society of America*, *34*(6), 749–761.

Freed, A., & Avizienis, R. (2000). A new music keyboard featuring continuous key-position sensing and high-speed communication options. In *Proceedings of the International Computer Music Conference*, Berlin, Germany.

Heet, G. (1978) *String instrument vibration initiator and sustainer*. US Patent 4,075,921.

Lee, N., Berdahl, E., Niemeyer, G., & Smith, J.O. (2008). Practical implementation of low-latency DSP for feedback control of sound. In *Proceedings of the Acoustics '08 Conference*, Paris, France.

Moog, R.A., & Rhea, T.L. (1990) Evolution of the keyboard interface: The Bösendorfer 290 SE recording piano and the Moog multiply-touch-sensitive keyboards. *Computer Music Journal*, *14*(2), 52–60.

Paradiso, J.A. (1997). Electronic music: New ways to play. *IEEE Spectrum*, *34*(12), 18–30.

Piano Bar. (2005). Products of interest. *Computer Music Journal*, *29*(1), 104–113.

Raichel, D. (2000). *The Science and Applications of Acoustics*. New York: Springer.

Weinreich, G. (1990). The coupled motion of piano strings. In A. Askenfelt (Ed.), *Five Lectures on the Acoustics of the Piano* (pp. 73–81). Stockholm: Royal Swedish Academy of Music.