Percussion instruments using realtime convolution: Physical controllers

Roberto Aimi
MIT Media Laboratory
20 Ames st
Cambridge. MA
roberto@media.mit.edu

ABSTRACT

This paper describes several example hybrid acoustic / electronic percussion instruments using realtime convolution to augment and modify the apparent acoustics of damped physical objects. Examples of cymbal, frame drum, practice pad, brush, and bass drum controllers are described.

Keywords

Musical controllers, extended acoustic instruments

1. INTRODUCTION

The purpose of this paper is to present several examples of hybrid acoustic / electronic percussion instruments that use realtime convolution with stored impulse responses to augment the apparent acoustics of damped physical objects.

While the convolution algorithm, prior art, and system architecture will be mentioned very briefly, a more thorough treatment can be found in [2].

1.1 Motivation

In typical digital percussion instruments [13], a set of pads are connected to a tone generator. The intensity of hits is measured by a peak detector and converted into trigger messages that tell the tone generator to play a sound of the specified intensity. Some systems additionally cross-fade between sounds based on measurement of hit position.

Unfortunately, the timbre of the hit is ignored. Whether it was a glancing hit, a soft mallet, or a fingertip, the same intensity results in identical output. Stirring the drum head with brushes won't make clear peaks for the detector to recognize. While players use these digital drum kits to good effect, much of the nuance of an acoustic set is missing.

One way around this is to use the audio signal of a hit directly. Previous applications have used the hit signal to drive waveguides [6] and banks of resonant filters [14]. The technique applied to the instruments presented in this paper is based on realtime convolution with stored impulse responses [4][7].

1.2 System architecture

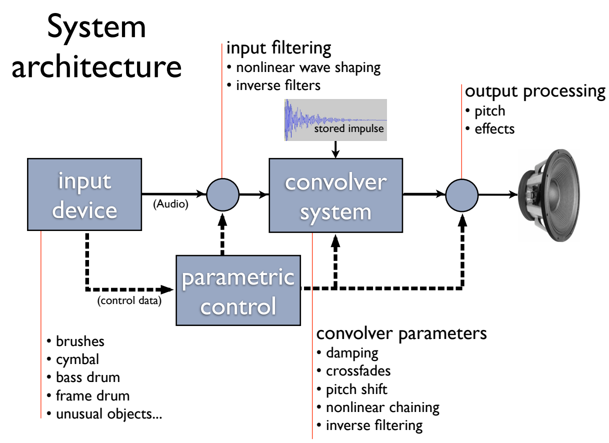

Figure 1: Basic system architecture

Figure 1 shows the basic system architecture for these instruments. Audio from a semi-acoustic damped object ("input device") is convolved with a stored impulse. Additional parametric control over the convolver and other processing can be applied by sensors built into the input device.

This system has been implemented to run in real time both in Max/MSP [3] and Pure Data ("Pd") [12].

2. PHYSICAL CONTROLLERS

Because of the nature of the processing, the physical part of the instrument is at least as important as the algorithmic part. For this section I will refer to the physical part of the instrument as a *controller* for convenience, though its acoustic properties and conception differ from typical controller schema. These controllers are designed to exploit the way the convolver acts as a resonator. By varying the degree of damping, physical resonances can be removed and replaced with any desired resonance.

The controllers described can be represented on a continuum based on the degree to which their own acoustics influences the output. At one extreme, the practice pad controller is highly damped, and although it does impart a "plastic" sound, it is a relatively minor coloration. In the middle, the various brush controllers give a clear impression that the stored impulse is being performed with a brush, taking on the dense time texture of the metal tines. At the other extreme, the cymbal controller provides significant coloration to any sound, enough so that it can sound like a cymbal bolted to a bass drum, or a cymbal attached to a snare when used with those stored impulses.

2.1 Cymbal

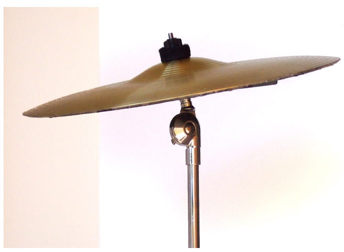

(a) Side view

(b) View of Redel connector, plastic substrate

Figure 2: Cymbal controller

This cymbal controller started out as a budget brass student cymbal, and it is designed to accommodate normal cymbal playing gestures such as hitting the bell or shell and choking the cymbal by grabbing the front of it. Since it is built around a modified real cymbal, it can sit on a standard cymbal stand. A foam grommet limits contact with the cymbal stand, and allows the cymbal to swing normally.

2.1.1 Assembly

The cymbal controller is assembled in layers, from top to bottom, as shown in figure 3, the layers are:

- A real brass cymbal
- PVDF element (MSI FLDT1-052K [8]) bonded to the cymbal underside, away from the playing area
- A thin foam layer to damp the cymbal and transfer choke force
- Force sensing resistor (interlink #406 [5]) to detect choke force at edge of the playing surface and control damping in the convolver
- Molded plastic cymbal substrate (pintech XT practice cymbal [11]) to support the assembly and further damp vibration. Edges are sealed with silicone caulk.

The FSR is connected directly to the computer audio interface. Audio signals are sent through the FSR and change in the signal levels is measured to determine the sensor's resistance. An advantage of this approach is that no additional hardware is needed, but it does take up another channel of input and output(figure 4).

The signals used have been in the 150 – 500 Hz range to minimize capacitive coupling while maintaining sufficient time resolution for controlling the damping parameter.

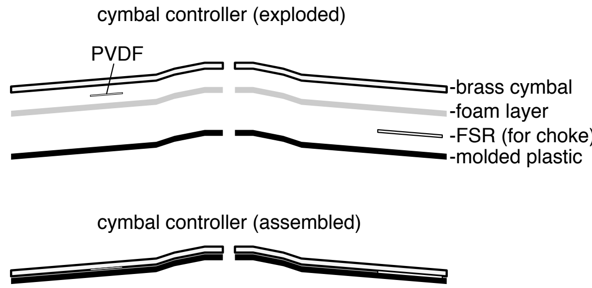

Figure 3: Cymbal controller assembly

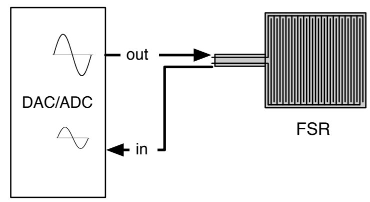

Figure 4: Using the audio interface to measure the resistance of the FSR

2.1.2 Function

Since there is significant spectral contribution from the cymbal, hits on the bell, rim, or edge sound substantially different from each other. I originally expected to need multiple contact microphones to get enough variation from hits in different locations, but it turns out that one microphone is sufficient because of the range of sounds achievable by hitting different parts of the cymbal. When convolving with a cymbal sound, the effect is that the lost resonance of the cymbal (due to damping) is restored, and it is quite surprising when the processing is turned off to hear that the real cymbal only sounds like a dull clank.

One drawback to allowing the controller to provide more of the spectrum is that while it heightens the realism of cymbal sounds, it will always impart a cymbal-like quality, even to non-cymbal sounds. For example, when convolved with a concert bass drum sound, the output sounds as if a cymbal was physically joined to the drum head.

2.1.2.1 Extensions.

To allow for cymbal crashes, two convolvers can be chained, approximating some of the nonlinearity of the real cymbal (described in detail in [2]). In addition to the FSR circuit, the surface of the cymbal was also electrically connected to the audio interface to pick up the 60Hz hum from when the player touched the surface. The hum was filtered and the envelope was used to control damping. Even though it provided essentially only one bit of data, having the cymbal be sensitive to touch over its entire surface proved to be more important than having a range of damping in one location. A potentiometer knob was added to the top of the cymbal (figure 5) to control pitch. The knob's resistance was measured by Pd using the same method as for the FSR. This allows the player to dial in a particular cymbal sound from the cymbal itself.

Figure 5: Cymbal pitch controlled with knob

2.2 Brushes

Two kinds of brush controllers were developed for use with this system, one wireless and one tethered. Drum brushes in either configuration were fitted with a PVDF contact microphone to pick up the sound in the metal tines. Any surface can be played with the brushes, and the resulting output sounds as if the sampled instrument is being played with brushes, but has the texture of the surface being played. By stirring the brush on a surface, a sustained broad band noise can be produced that results in quite different timbres than were observed with the pads or cymbal controller. Different combinations of surface textures, brush movements and stored impulse are possible.

2.2.1 *Wireless brush*

Figure 6: Wireless brush

The wireless brush used the circuit board and part of the enclosure of a handheld VHF wireless microphone (Nady DKW-1H [10]) to transmit its audio signal. Up to four wireless brushes can be used simultaneously on four different VHF channels. As with the wired brush, a piece of PVDF (Digikey MSP1006-ND) was threaded through the tines to pick up the brush sound. Kapton tape was used to protect the piezo element from abrasion from the brushes.

2.2.2 *Wired brush*

The wired brush controller started with a rubber-handled drum brush, and added a 3" bend sensor [1] to detect when the brush was pressed against a surface. The bend sensor was placed in line with the tines, while a PVDF tab was threaded through the tines. The rubber covering was split to make room for the wiring.

Figure 7: Wired brush

A Redel-compatible connector was added to the end of the brush to allow quick connections to a multi-conductor cable. This connector was common to several of the controllers built, allowing easy interchangeability.

2.3 Pad

Figure 8: Percussion pad controller

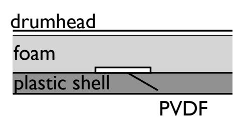

Figure 9: Percussion pad cross section

This is a simple controller derived from a drum practice pad. Since one of the goals of a practice pad is to be quiet, it was already well damped. A piece of PVDF foil [9] was applied under a layer of foam located beneath the drumhead and was connected directly to the audio interface (figure 9).

The pad proved a surprisingly versatile controller, working well with most impulses. Due to the head material and the high degree of damping, treble had to be boosted to maintain reasonable sound. Unfortunately this also made it more susceptible to noise.

The practice pad had a less-variable sound due to its thick plastic head and highly damped design. The foam itself made some noise when compressed, creating unrealistic artifact for louder hits. Players had to work to produce a meaningful range of impulses. Sanding the head helped the sound somewhat, as did maximizing the tension of the head. Hitting the metal tension ring around the perimeter of the head gave more of a metallic clank, which was quite different from any of the sounds achievable by hitting the drum head.

2.4 Frame drum

Based on the preliminary results of early versions of the cymbal controller, I wanted to apply the same technique of

Figure 10: Frame drum controller

using more of the acoustic response of the physical object to the construction of a drum controller. Starting with a wooden frame drum, I added contact microphones, damping material, and pressure sensors (figure 11). This drum was much less damped than the practice pad, ensuring that more of the spectrum of the drum was carried through the processing.

Drums struck in different locations can excite different modal structures. For example, striking location helps create the differences between Djembe bass, tone, and slap sounds. Unfortunately, the convolution system is limited to the one set of modes that are in the sampled sound. One way around this problem is to run multiple convolutions at once, and to have contact microphones at multiple locations on the drum head, or one could also track the location of the hit and control a cross fade between convolvers. In this case, multiple contact microphones were used to be able to process hits on the center and edges of the drum differently.

2.4.1 Assembly

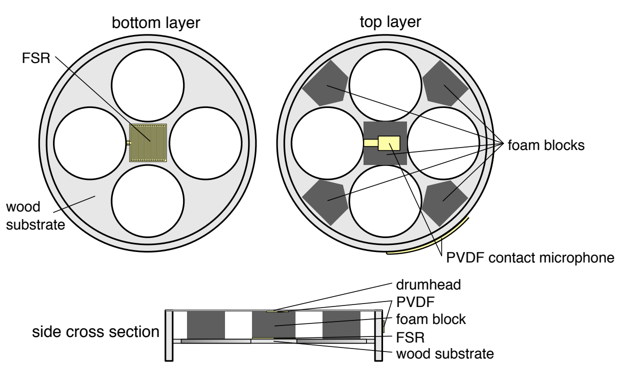

Figure 11: Frame drum controller assembly

One PVDF element (MSI FLDT1-052K [8]) was mounted to the underside of the center of the drumhead, and another was mounted to the frame.

A force sensing resistor ("FSR", interlink #406 [5]) was mounted to a wooden substrate at the center of the drum and covered by a foam block to provide control of damping. The compressibility of the foam block allowed for a greater displacement of the drum head over the active range of the sensor, and also served to protect the sensor by spreading any forces over its whole area. The order of assembly at the center, from top to bottom is:

- Drumhead
- PVDF
- Foam block
- FSR
- Wooden substrate

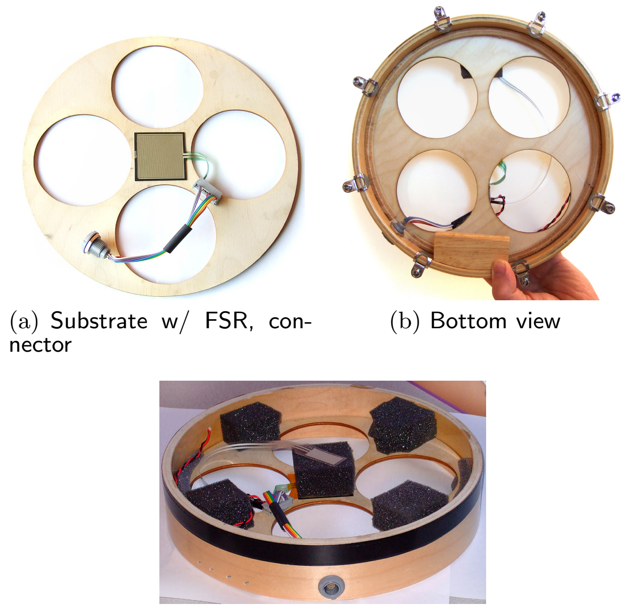

(a) Substrate w/ FSR, connector

(b) Bottom view

(c) Drum without head

Figure 12: Frame drum controller

2.4.2 Function

Since the FSR is mounted at the center of the drum, it responds to pressing anywhere on the drumhead (although much more strongly at the center). This gives good subtle control of damping by pushing at the edges, while still allowing sudden and immediate damping by pushing at the center.

Pushing on the drum head also raises the pitch of the drum slightly. Originally I had intended to have a small pitch bend controlled by a second pressure sensor, but for many drum sounds, there is enough of a pitch effect due to the changes in tension in the real drum head, even though the stored impulse is not shifted.

Separate processing of the rim signals from the center works particularly well for djembe sounds. Since there is an increase in low frequency output of the center PVDF sensor when it is hit directly, I found that I could combine djembe bass and tone sounds into one sample, and obtain more of one or the other sound entirely based on where and how the drum was hit, while using the edge sensor just for djembe slap sounds.

2.5 Bass drum with speaker

For this controller, I was interested in having the sound emit from the object, to provide a stronger illusion that the

Figure 13: Bass drum controller

player was interacting with a physical object rather than a computer. I converted a bass drum shell into what is essentially a speaker cabinet in which the speaker is located behind the drum head. This provided both a sonic and tactile feedback to the player.

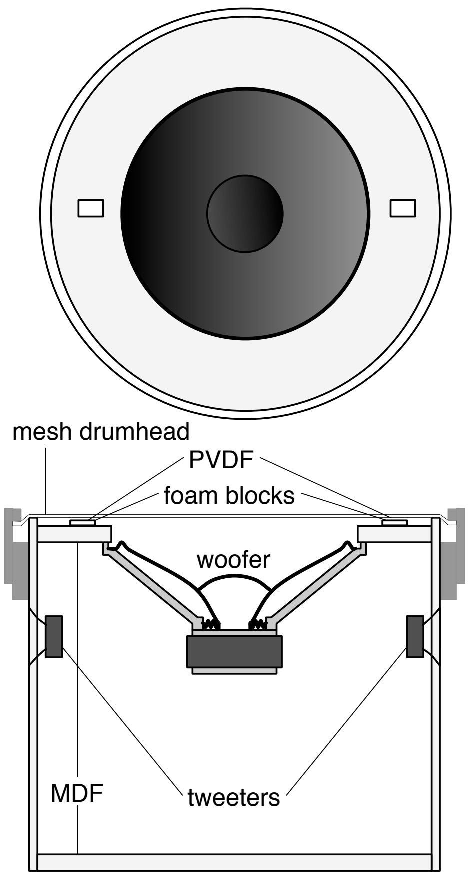

Figure 14: Bass drum controller

2.5.1 *Assembly*

The assembly is shown in figure 14. A circle of medium density fiberboard (MDF) was used to seal one end of the drum, and an MDF ring supports and centers a 15-inch bass speaker at the other end. Internal MDF bracing was also added. A nylon mesh drumhead was stretched over the end with the speaker, allowing sound to pass through the mesh. Vibration in the mesh is picked up by two PVDF elements (LDT0-028K, [8]) supported by foam glued to the MDF ring.

Side-mounted piezo horn tweeters (Pyle PSN1167) were added to improve the system's high frequency response.

Audio output from the computer was routed to the speaker in the drum, semantically re-coupling the resonator to the playing surface, though thanks to the mesh they stayed essentially acoustically uncoupled.

2.5.2 *Function*

The bass drum controller, because of its appearance, loud output, and low bass extension, was well suited for the obvious role of large drum sounds, along with thunder, prepared piano soundboard, as well as for large gongs and cymbals. Due to the resonance of the mesh, some equalization was necessary to control feedback, making it an ideal candidate for using deconvolution to pre-filter a typical hit from the the stored impulses. One surprising outcome is that it is actually well-suited for snare drum sounds, provided that the head is given a high enough tension to provide proper stick bounce.

3. SUMMARY

The fundamental trade-off for this system is that for the output to sound exactly like the stored sample, one would like the input to be a perfect impulse function with no timbral contribution from the physical controller. But for there to be sufficient variation in the timbre, the acoustic contribution of the controller has to be significant. In a system like this, the placement and design of the secondary controls such as pressure, bend, and touch sensors not only have to be consistent with the use of the instrument, but have to allow the controller to still function as an acoustic object.

These controllers differ greatly in how their own acoustics influence the final sound. For the bass drum and pad, I saw that influence as a potential liability. The range of timbres was small, and the typical timbre had strong resonances requiring work through equalization and filtering to mitigate its impact. For the frame drum and cymbal, it was possible for the player to extract a much broader variation of timbre, giving an extra element of realism and variation to the final output.

Since the acoustic qualities of these controllers are so critical to their function, these examples represent a tiny slice of what can be realized through extended development. In the same way that existing percussion instruments have constantly been extended and refined, the physical controllers and their associated processing can also benefit from time and iteration.

4. ACKNOWLEDGMENTS

Many thanks to the percussionists who helped test and reshape these designs:

- Rakalam Bob Moses
- Jamey Haddad
- Curt Newton
- Dave Flaherty

5. REFERENCES

[1] Abrams/Gentile Entertainment, Inc. New York, NY. `http://www.ageinc.com`.

[2] R. Aimi. *Hybrid Percussion: Extending Physical Instruments Using Sampled Acoustics*. PhD thesis, Massachusetts Institute of Technology, 2006.

[3] Cycling '74. *Max/MSP*. `http://www.cycling74.com/`.

[4] W. Gardner. Efficient convolution without input-output delay. *Journal of the Audio Engineering Society*, 43(3):127, 1995.

[5] Interlink Electronics, Inc. Camarillo, CA, `http://www.interlinksensors.com`.

[6] Korg corporation. *Wavedrum user's guide*. Tokyo, Japan, 1994.

[7] D. S. McGrath. Efficient convolution method and apparatus. United States Patent 6,574,649, June 22 2001.

[8] Measurement Specialties, Inc. Hampton, VA. `http://www.meas-spec.com`.

[9] Measurement Specialties, Inc. *Piezo Film Sensors Technical Manual*, 1999.

[10] Nady systems, inc. Emeryville, CA, `http://www.nady.com`.

[11] Pintech USA, inc. Greenville, SC `http://www.pintechworld.com`.

[12] M. Puckette. "pure data.". In *Proceedings, International Computer Music Conference*, pages 269–272, San Francisco, 1996. International Computer Music Association.

[13] Roland corporation. *percussion sound module TD-20 owner's manual*, 2004.

[14] K. Van Den Doel and D. Pai. *Audio anecdotes*, chapter Modal synthesis for vibrating objects. A.K. Peter, Natick, MA, 2003.