Caress: An Enactive Electro-acoustic Percussive Instrument for Caressing Sound

Ali Momeni
Carnegie Mellon University
CFA 300 – 5000 Forbes Ave
Pittsburgh, PA 15213
momeni@cmu.edu

ABSTRACT

This paper documents the development of Caress, an electroacoustic percussive instrument that blends drumming and audio synthesis in a small and portable form factor. Caress is an octophonic miniature drum-set for the fingertips that employs eight acoustically isolated piezo-microphones, coupled with eight independent signal chains that excite a unique resonance model with audio from the piezos. The hardware is designed to be robust and quickly reproducible (parametric design and machine fabrication), while the software aims to be light-weight (low-CPU requirements) and portable (multiple platforms, multiple computing architectures). Above all, the instrument aims for the level of control intimacy and tactile expressivity achieved by traditional acoustic percussive instruments, while leveraging real-time software synthesis and control to expand the sonic palette. This instrument as well as this document are dedicated to the memory of the late David Wessel, pioneering composer, performer, researcher, mentor and all-around Yoda of electroacoustic music.

Author Keywords

Percussion, Resonance, physical modeling, Max, Puredata, Gesture, Rhythm, Finger Drums

ACM Classification

H.1.2 [User/Machine Systems] Human factors H.5.2 [Information Interfaces and Presentation] User Interfaces—Input devices and strategies, Interaction styles, Prototyping H.5.5 [Information Interfaces and Presentation] Sound and Music Computing—Methodologies and techniques, Modeling, Signal analysis, synthesis, and processing, Systems

1. INTRODUCTION

The late musician, researcher, mentor and thinker David Wessel had a saying: he never wanted to *trigger* sounds, rather he wanted to *caress* sounds (D. Wessel, 2009). The desire to caress sounds was rooted in Wessel's approach to control intimacy (Momeni & Wessel, 2003; D. Wessel & Wright, 2002); this position was conceptually flushed out in Wessel's discussion of an *enactive* approach to computer music performance, where he noted significant research linking perception and action in the visual experience, while auditory and musical cognition research has "a lot of catching up to do" (D. Wessel, 2006).

Enactive musical interfaces received significant attention in the previous decade, Essl and O'Modhrain (Essl & O'Modhrain, 2006) discuss a lineage of theorizing on gesture and sound that begins with Cadoz's *instrumental gesture* (Battier & Wanderley, 2000; Cadoz, 1988) and apply it to the analysis of several contemporary gestural software based instruments. This position proposes a layered mechanism of information-transfer between the performer and the instrument (*ergotic*), performer and audience (*semiotic*), and the performance and its context (*epistemic*); Jorda and Armstrong's dissertations from 2005 and 2006 (Armstrong, 2006; Jorda, 2005) make significant strides in helping instrument designers contextualize their work within several decades of *digital lutherie* (Jorda, 2005).

In the realm of electronic percussion instruments, Wessel's comment about controllers that merely "trigger" sounds is especially poignant. Commercially available controllers that trigger sounds are pervasive in the market and on festival stages throughout the world. While some musicians clearly reach a level of virtuosity with these controllers, Wessel's comment points to the limiting inflexibility of sounds that are triggered: First, once the sound is triggered, it is out of the musicians hands; he/she can stop the sound or change some filter parameters but little more. Second, trigger controllers tend to provide very few degrees of freedom: some offer velocity sensitivity for each trigger, but in comparison with the expressivity of an acoustic drum-head, a single 7-bit scalar associated with each hit leaves much to be desired.

Caress (figure 1) pursues an enactive approach to the design of percussive controller by relying on the richness and intimacy of audio signals generated directly through gesture and touch. While sensors tend to only sense what they are made for (e.g. force-sensitive-resistors sense pressure, photoresistors sense light, infrared-proximity sensors sense distance, etc.), microphones capture an enormous range of expressive gestures from tapping with the fingertip to scratching with the nails, to rubbing, and shaking, etc. Caress is designed around the hypothesis that a percussion instrument must leverage this range of expressivity as opposed to limiting it through its choice of sensing technology.

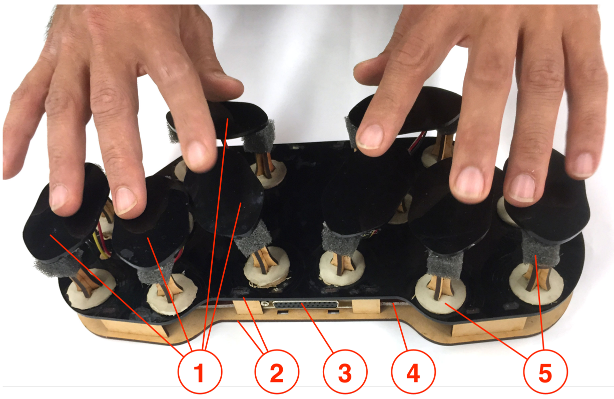

Figure 1. Caress (prototype 4): 1) four finger pads under each hand, each with a piezo-mic underneath, 2) parallel base plates, 3) DB25 multi-channel audio connector, 4) analog mic-preamp circuit board, 5) vibration isolation mechanisms

Finally, by relying primarily tactile interaction rooted in acoustic percussion techniques, Caress seeks intimate musical control, sonic versatility and high-ceiling for virtuosity (D. Wessel & Wright, 2002), what Golan Levin calls Instantly Knowable, Infinitely Masterable (Levin, 2000).

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

2. BACKGROUND

Caress continues along a trajectory of electroacoustic percussion research that has received significant attention in the past decades. Related research may be classified into several categories discussed below.

2.1 Controllers For Percussive Sounds

The commercial side of the percussion controller market is dominated with drum triggers in a variety of sizes and shapes. Wildly popular controllers like the Ableton Push, the Novation LaunchPad, Native Instrument Machine Studio, Arturia SparkLE, Akai MPC Element, Korg padKONTROL, Alesis Performance Pad Pro, Roland Octapad and the Vestax PAD-One have been improved over the past two decades to provide more features (primarily additional continuous controllers like knobs, sliders and x-y controllers), smaller and lighter form factors, and lower costs (many options under $50). The dominant paradigm for these controllers is to produce MIDI notes with pitches and velocities that reflect some aspect of the triggering gesture. A small number of manufacturers have forayed into alternative percussion controllers; notable among them are the Roland HandSonic and Korg Wavedrum that add continuous control to the standard pitch-velocity input, using proximity sensors that report the position of the hands above the controller. Within computer music research, the monumental Radio Baton (Boulanger & Mathews, 1997) has inspired many noteworthy projects; Schloss et al.'s work (Jones & Schloss, 2007; Schloss, Kapur, Tzanetakis, & Tindale, 2005; Schloss & Driessen, 2001) is particularly relevant as it also pertains to the notion of gesture-as-audio (see 2.5), as is Wessel's latest iteration of his controller, the SLAB (D. Wessel, 2009).

2.2 Sensing For Percussive Instruments

This category covers a wide range of related work within the various branches of computer-human-interaction applied to real-time control of percussive sounds. Several invaluable literature reviews for sensing applied to percussion instruments have been published in the past few years (Collicutt, Casciato, & Wanderley, 2009; Bouënard & Wanderley, 2009; Tindale, Kapur, Tzanetakis, Driessen, & Schloss, 2005). Sensing paradigms may be broadly separated into the categories of direct sensing (Tzanetakis, Kapur, & Tindale, 2006; Roberto Aimi, 2007; Robeto Aimi & Young, 2004; Kapur, Essl, Davidson, & Cook, 2003) where sensors on the controller itself convey information about the performers gestures; or indirect sensing where external sensing mechanisms like a depth camera (Schloss et al., 2005; Trail, Dean, Tavares, & Odowichuk, 2012) or a multi-touch screen (Ren, Mehra, Coposky, & Lin, 2012) are employed to acquire gesture information. Since indirect sensing for percussion instruments tends to detract from the enactive quality of the interaction, by removing materiality and haptic feedback. Caress employs direct sensing at audio-sampling rates.

2.3 Synthesis Techniques For Percussion Instruments

Digital audio synthesis research for percussive sounds is particularly rich in applications of physical modeling (Eckel, Iovino, & Causse, 2014; Polfreman, 2002; Bilbao, 2010; Bilbao, 2011; Avanzini & Marogna, 2010; van den Doel & Pai, 2003; Bruyns, 2006; Peltola, Erkut, Cook, & Välimäki, 2007; Cook & Scavone, 1999). Bruyns's work on model synthesis (Bruyns, 2006) provides an excellent review of relevant existing work. Chuchacz et al. provide a framework for designing controllers dedicated to physically modeled percussion instruments (Chuchacz, O'Modhrain, & Woods, 2007). Other approaches include real-time convolution (Roberto Aimi, 2007), concatinative synthesis (Schwarz, Beller, Verbrugghe, & Britton, 2006; Schwarz et al., 2006; Schwarz, 2012) and analysis-resynthesis (Macon, McCree, Lai, & Viswanathan, 1998; Jehan, Freed, & Dudas, 1999). Mogee[1], a recent UK-based startup, is an excellent convergence of sophisticated analysis-resynthesis, industrial design and mobile performance; their product "turns everyday objects into musical instruments" using physical modeling and piezo-microphone excitation.

2.4 Gesture As An Audio Signal

In their 2003 publication Nevile et al. (Nevile, Driessen, & Schloss, 2003) recognize that "if what is desired is to actually *play* the machine in the sense of a traditional instrument, the interfaces currently available to control the sound synthesis are unsatisfactory". Caress builds on this claim by leveraging the infinite richness and variability sounds generated directly by the hands—as compared with gesture captation mediated by a controller. A handful of research projects have stepped away from *triggering sounds* and towards *caressing sounds* by focusing on continuous and high-sample-rate control of sounds beyond the moment of initiation. This work reaches for modes of expressivity beyond what is offered by MIDI: trigger a particular sound with a desired "pitch" and "velocity", along with a tad of continuous sound "bending" beyond the initial trigger moment. Nevile et al. (Nevile, Driessen, & Schloss, 2003) propose a technique for capturing gesture information from a controller as modulations of audio signals, thus allowing a computer's analog-to-digital converters to serve as very high-sampling-rate and low-latency gesture capitation interface. Avizienis et al. (Aviziensis & Wessel, 2000) introduced a hardware platform for achieving very low latency and low jitter gesture data by multiplexing the audio, gesture data, and MIDI data into a single stream of UDP packets. Tindale's work on extended percussive gesture (Tindale, 2007) applies audio analysis towards timbre recognition of percussive sounds; the timbre information is then used to modify synthesis parameters of a physical model. While this approach shares a number of features with Caress (e.g. gesture as audio, physical modeling synthesis), it differs in an important way: where as Tindale extracts control information from the audio signal through audio analysis and timbre recognition, Caress performs no audio analysis to the incoming signal, but rather directly excites a resonance model with it.

Highly relevant research exists in the HCI literature: Deyle et al. (Deyle, Palinko, & Poole, 2007) describe an impressive bioacoustics gesture interface that employs two piezoelectric sensors placed on the wrist or ankles and reports information about the users gestures based on audio analysis of sounds generated by the body movement and transmitted to their device with bone conduction. In their paper describing "acoustic barcodes"—an ingenious technique for low-cost gesture recognition—Harrison et al. (Harrison, Xiao, & Hudson, 2012) employ "structured patterns of physical notches that, when swiped with e.g., a fingernail, produce a complex sound that can be resolved to a binary ID".

3. IMPLEMENTATION

This section describes the hardware and software design for Caress (figure 1), followed by more details about intermediary prototypes and techniques.

[1] http://mogees.co.uk/

3.1 Hardware Design

Caress' hardware is comprised of: a physical chassis, analog circuitry, audio interface and computing platform.

3.1.1 Chassis: Base-plate and Finger-pads

The base plates are made of 1/8" hardboard (actual thickness of 0.115"), an affordable composite wood material well suited for rapid prototyping with a laser cutter, and separated from one another with threaded rod. The top plate holds several playable finger-pads which are acoustically isolated from one another; the earlier prototypes achieve this isolation with a combination of materials (e.g. foam) and structural features (see section 3.5.3 and figure 5). The space between the two base-plates houses the analog circuitry and the DB-25 connector for transmitting audio to the audio interface.

3.1.2 Analog Circuitry

Connecting piezo-microphones directly to an audio interface or mixer results in a notable bias in the frequency spectrum; musicians recognize this as effect as the typical "tinny" sound of contact microphones. This effect is caused by the way extremely high impedance of contact microphones (upwards of 10MOhms) interacts with mixers and audio interfaces. A custom circuit board was developed based on a design by Alex Rice and modified by Zach Poff[2] that employs three J-FET amplifiers to match the transducer's impedance to that expected by standard audio interfaces (Figure 2). An improved design is under development that adds tone control and much higher analog gain using a high-quality audio amplifier component. Designs for the printed circuit board in figure 2[3] as well as the latter high-gain version[4] are available to the public on the community based low-cost PCB manufacturing site OSHPARK.com. Note that this design accommodates 2- and 3-wire piezos; if the piezo is 2-wire, the "GND" pin for the piezo input must remain disconnected.

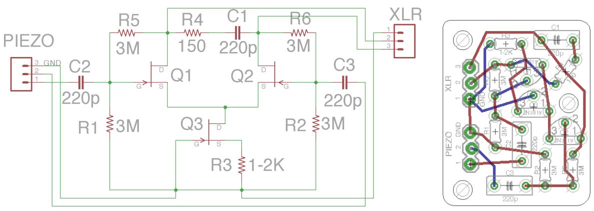

Figure 2. Analog audio preamplifier schematic (left) and board (right); Q1-3 are 2N3819 are JFET amplifiers

3.1.3 Multi-channel Audio Interface

A Metric Halo ULN8 firewire audio interface was used to digitize the signal from the multi-channel preamp circuitry. While the eight-channels of ultra-low-noise preamps and a DB-25 connector for mic-inputs make this interface very well suited to this application, a more modest audio interface would also be sufficient for creating this instrument.

3.1.4 Computing Platforms

Max offers musicians a world of possibilities within synthesis, control, composition and performance. Max for Live was introduced in order to combine these possibilities with the optimized workflow of a powerful digital-audio-workstation designed for production and live play. The software behind Caress was developed in Max and implemented as a Max for Live plugin to get the best of both worlds. All initial tests were performed using custom software written in Max using the *resonators~* object. Future development will include deploying the synthesis on embedded platforms like the Raspberry Pi, Android and iOS devices using *libpd* (Brinkmann, Kirn, & Lawler, 2011). The portability and cost benefits of these latter platforms will come with the trade-off of limitations on the number of available audio channels as these smaller and less robust devices typically only handle stereo input and output.

3.2 SOFTWARE DESIGN

Caress's audio synthesis scheme relies on resonance synthesis (Jehan et al., 1999), a highly efficient way to synthesize percussive sounds with sharp attacks and exponentially decaying amplitude envelopes. This approach offers several notable benefits:

Analysis-resynthesis: Synthesis of sounds based on analysis of existing sounds is a powerful rubric. Environments like Diphone Studio's ResAn[5] perform complex multi-step analysis of an audio sample and produce a text file containing a large number of frequency-amplitude-decay-rate triplets (sometimes upwards of several hundred resonances). In practical terms, this dramatically expands the palette of possible sounds by allowing composers and musicians to create their own resonance models by analyzing samples they themselves record, or those from cherished sample libraries. The *resonators~*[6] external for Max provides a highly optimized means for doing resonance model synthesis in real-time.

Real-time control: Resonance model synthesis allows for very rich manipulation of the sound during synthesis. Since the synthesis is based on known frequencies/amplitude/decay-rates, simple arithmetic operations allow for complex changes to the resultant sound, including pitch shifting, spectral shaping, changes in spectral density, as well as ways of "squelching" the sound that are more analogous to the physical world (think of hitting a cymbal and dampening it with your hand). The restransform[7] external for Max provides a means for controlling many of these transformations in real-time.

Actuation with audio: Resonance models can be excited with another audio signal; Caress leverages this feature of resonance model synthesis to bring very sensitive and intimate gestural control of sound synthesis. Since the palette of possible sounds one can create with ones fingers and a contract microphone is extremely broad, the resultant synthesis-space is extremely rich and malleable.

Computational efficiency: Whereas a time-domain sample requires significant storage space, a resonance models require only a tiny text file with just a few lines of text. These files are human readable and hand-editable. A modern laptop computer can easily synthesize dozens of resonance models, each with hundreds of resonance frequencies.

3.3 Design Goals and Challenges

3.3.1 Portability

The rise of laptop musicianship has been accompanied by a complimentary growth in gestural controller offerings that approximately match the size and shape of the laptop itself. This trend is notable in off-the-shelf controllers that appear throughout studios, festival stages and electronic-music classroom alike; examples include affordable USB powered

---

2 http://www.zachpoff.com/diy-resources/alex-rice-piezo-preamplifier/

3 https://oshpark.com/shared_projects/vi96HYmw

4 https://oshpark.com/shared_projects/n0J0Yx2E

5 http://anasynth.ircam.fr/home/english/software/diphone-studio

6 http://cnmat.berkeley.edu/downloads

7 http://cnmat.berkeley.edu/downloads

and driverless controllers (see section 2.1), controllers made by the leading software companies (e.g. Ableton Push or Native Instruments' Traktor Kontrol). The design space for these controllers is largely focused on knobs, sliders and buttons. In designing Caress*,* the author recognizes the affordances of a small and flat form factor, like that of a laptop, in terms of portability and performance set-up. The most recent iteration of *Caress* is therefore designed to be about 9" wide, 5" long and 1" inches thick.

3.3.2 Vibration Isolation

Vibration Isolation has been by far the most difficult challenge in designing Caress. This instrument tries to fit eight contact microphones with high-gain in a very small form-factor, while minimizing cross-bleeding of signals. While the author envisions some possibilities in abetting this challenge in software, the effort thus far has been focused on the physical. Vibration isolation is achieved in three ways: material, structural and kinematic. In terms of material, the various iterations of Caress utilize a range of sound absorbing materials to dampen and isolate signals. Specifically, the author experimented with a variety of extruded foams that dampen the signal from each finger pad from one another. The structural design provides additional vibration isolation through networks of staggered and concentric cuts in the base-plate that surround each finger pad (figure 5, right) that dampen transmission of signal through the solid medium.

3.4 Iterative Prototyping

The challenges to designing this instrument have required a great deal of trial and error. The approach to creating the instrument has therefore been one that couples iterative digital fabrication, with systematic variation of materials, fasteners and features. To these ends, the author employed Grasshopper[8], a parametric design plugin for the popular Rhinoceros[9] computer-aided-design environment in order to make systematic changes between iterations.

3.4.1 Prototype 1: Quick Proof of Concept

This iteration was constructed by hand and with experimentation in mind (figure 3).

The finger pads were made of different materials to see whether one offered advantages over the other. This prototype had finger-pads, hence five independent channels of audio. The vibration isolation (see section 4.3) was implemented only with pieces of foam beneath each finger pad.

Notable lessons included:

- Wide range of expressivity in dynamic range, timbre, and very tight temporal control
- Variations in materials for the finger pad did not offer significant gains
- Cabling management was a major issue to address; movements of the audio cables connected to each contact microphone produced considerable unwanted noise
- Vibration isolation with foam was helpful but insufficient

[8] http://www.grasshopper3d.com/
[9] https://www.rhino3d.com/

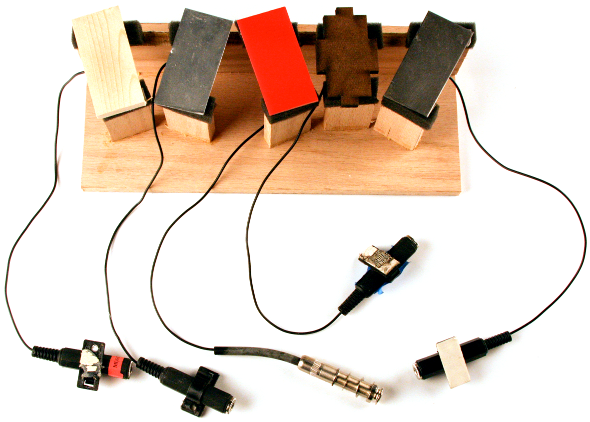

Figure 3. Prototype 1: Five finger pads, five independent piezo contact microphones; finger pads made of different material (hard wood, plastic, composite wood, steel).

3.4.2 Prototype 2: first laser cut model

This prototype was the first fabricated from a parametric design in Rhino/Grasshopper (figure 4).

Figure 4. Prototype 2: Laser-cut from Masonite and acrylic plexi; Ten finger pads all made of acrylic plexi and arranged to match finger positions. Audio cables exit form the front of the instrument independently.

Notable improvements included:

- All finger pads were made of the same material (1/8" acrylic), and the base-plates made of ¼" masonite
- A resting pad for the palms of the hand was added as the floating had position from prototype 1 proved tiring
- Better vibration isolation using material choices (foam) as well as structural design features (see 4.3)

Shortcomings included:

- Cumbersome audio wiring
- Palm rest position ineffective as during performance one has the tendency to move hands about freely
- Difficult to pack and carry about
- Ten finger-pads is inconvenient for most 8-channel interfaces

3.4.3 Prototype 3: Improved vibration isolation and audio wiring

Notable improvements in this prototype (figure 5) were:

- Eight finger pads; smaller form-factor
- Connectors beneath each finger-pad allow for quick replacement of piezo-microphones
- Single DB25 connector for all eight audio signals

Figure 5. Prototype 3: Top: Eight finger pads, DB-25 connector carrying all eight balanced signals exists from front of instrument; larger finger pads for easier interaction; Bottom Left: audio cables from each microphone with JST connectors to enter the space between the base plats where they converge into the DB-25 connector; Bottom Right: better isolation vibration achieved with a series of concentric arcs around the base of each finger-pad

3.4.4 Prototype 4: Improved vibration isolation and analog audio circuitry

The most notable change in prototype four (figure 1) was the addition of dedicated analog audio preamplifing circuitry for each piezo-microphone (see 3.1.2). This feature improved the audio quality by removing the spectral bias imposed by the high impedance piezo input, reduced distortion from extremely high voltage peaks, and allowed for increased gain levels in the audio interface to amplify the signal further. This prototype also integrated a rubber-base for each finger pad that provided additional vibration isolation.

4. CONCLUSIONS

With regard to hardware design, the development and initial evaluation of Caress point to the intersection of percussion, digital synthesis, custom digital fabrication and gesture-as-audio as an exciting area of research through design. Most commercial percussion controllers fail to capture the nuances and subtleties of human touch that make so many percussion instruments so expressive. As computing hardware becomes increasingly powerful, and digital fabrication more approachable, computationally intensive systems like Caress that utilize multichannel audio, real-time synthesis, treatment of gesture as audio-rate data and delicate custom fabrication become more feasible.

With regard to software, the use of analysis-resynthesis in a tactile instrument that accepts a very broad range of interactions has the potential to offer instrument builders a more open approach to instrument-specific sound design. Analysis-resynthesis based on models of resonance offers a wide and expressive sounds space to the performer in two ways: First, by offering recorded sounds as potential source material dynamic percussive performance; second, by opening the interaction to any kind of touch (as opposed to a particular gesture suited to a particular sensor).

Although Caress has gone through several iterations of prototypes, a great deal of work remains to be done. The author considers two areas as priorities in future development:

1) Mobile: A smaller and more portable version of Caress is in the works as a hardware accessory for a mobile phone, and the phones standard sound inputs and outputs from

2) Gesture Extraction: A number of recent papers point to tremendous possibilities in extracting more information about the performers gestures directly from the produced audio signal, or through coupling alternative techniques with this instrument's existing hardware; most notably using swept-frequency acoustic sensing (see (Ono, Shizuki, & Tanaka, 2013); (Honigman, Hochenbaum, & Kapur, 2014)) or various forms of capacitate sensing (see (Große-Puppendahl, Berghoefer, Braun, Wimmer, & Kuijper, 2013); (Savage, Zhang, & Hartmann, 2012); (Hudson & Mankoff, 2006).

5. ACKNOWLEDGMENTS

A range creative exploration and invention over several decades at CNMAT, most notably by Adrian Freed and the late David Wessel, inspires this instrument. Drew Anderson and Luke Loeffler assisted in the fabrication of Caress's 2nd and 4th prototypes.

6. REFERENCES

[1] Aimi, Roberto. (2007). Percussion Instruments Using Realtime Convolution: Physical Controllers. *Nime*, 154–159.

[2] Aimi, Robeto, & Young, D. (2004). A New Beatbug: Revisions, Simplifications, and New Directions [electronic resource].

[3] Armstrong, N. (2006, April 15). *An Enactive Approach to Digital Musical Instrument Design*. Princeton, NJ.

[4] Avanzini, F., & Marogna, R. (2010). A Modular Physically Based Approach to the Sound Synthesis of Membrane Percussion Instruments. *Audio, Speech, and Language Processing, IEEE Transactions on*, *18*(4), 891–902. http://doi.org/10.1109/TASL.2009.2036903

[5] Aviziensis, R., & Wessel, A. F. T. S. D. (2000). Scalable Connectivity Processor for Computer Music Performance Systems [electronic resource].

[6] Battier, M., & Wanderley, M. (2000). Trends in gestural control of music. Ircam Publ (CDROM).

[7] Bilbao, S. (2010). Percussion Synthesis Based on Models of Nonlinear Shell Vibration. *IEEE Transactions on Audio, Speech, and Language Processing*, *18*(4), 872–880. http://doi.org/10.1109/TASL.2009.2029710

[8] Bilbao, S. (2011). Time domain simulation and sound synthesis for the snare drum. *Journal of the Acoustical Society of America*, *131*(1), 914–925. http://doi.org/10.1121/1.3651240

[9] Bouënard, A., & Wanderley, M. M. G. S. (2009). Advantages and Limitations of Simulating Percussion Gestures for Sound Synthesis [electronic resource].

[10] Boulanger, R., & Mathews, M. (1997). The 1997 Mathews Radio-Baton and Improvisation Modes [electronic resource].

[11] Brinkmann, P., Kirn, P., & Lawler, R. (2011). Embedding pure data with libpd. ... *Of the Pure Data ...*.

[12] Bruyns, C. (2006). Modal Synthesis for Arbitrarily Shaped Objects. *Computer Music Journal*, *30*(3), 22–37. http://doi.org/10.1162/comj.2006.30.3.22

[13] Cadoz, C. (1988). Instrumental Gestures and Musical Composition (pp. 1–14). Presented at the ICMC.

[14] Chuchacz, K., O'Modhrain, S., & Woods, R. (2007). Physical Models and Musical Controllers - Designing a Novel Electronic Percussion Instrument. *Nime*, 37–40.

[15] Collicutt, M., Casciato, C., & Wanderley, M. M. (2009). From Real to Virtual: A Comparison of Input Devices for Percussion Tasks. *Nime*, 1–6.

[16] Cook, P. R., & Scavone, G. (1999). The synthesis toolkit (stk). *... Of the International Computer Music ...*.

[17] Deyle, T., Palinko, S., & Poole, E. S. (2007). Hambone: A bio-acoustic gesture interface. *... Computers*. http://doi.org/10.1109/ISWC.2007.4373768

[18] Eckel, G., Iovino, F., & Causse, R. (2014, December 24). Sound synthesis by physical modelling with Modalys. Retrieved December 24, 2014, from http://iem.at/~eckel/publications/eckel95b.html

[19] Große-Puppendahl, T. A., Berghoefer, Y., Braun, A., Wimmer, R., & Kuijper, A. (2013). OpenCapSense: A rapid prototyping toolkit for pervasive interaction using capacitive sensing. *PerCom*, 152–159. http://doi.org/10.1109/PerCom.2013.6526726

[20] Harrison, C., Xiao, R., & Hudson, S. E. (2012). Acoustic barcodes: passive, durable and inexpensive notched identification tags. *Uist*, 563–568. http://doi.org/10.1145/2380116.2380187

[21] Honigman, C., Hochenbaum, J., & Kapur, A. (2014). Techniques in Swept Frequency Capacitive Sensing: An Open Source Approach.

[22] Hudson, S. E., & Mankoff, J. (2006). Rapid construction of functioning physical interfaces from cardboard, thumbtacks, tin foil and masking tape. *the 19th annual ACM symposium* (pp. 289–298). New York, New York, USA: ACM. http://doi.org/10.1145/1166253.1166299

[23] Jehan, T., Freed, A., & Dudas, R. (1999). Musical applications of new filter extensions to max/msp. *International Computer Music ...*.

[24] Jones, R., & Schloss, W. A. (2007). Controlling a Physical Model with a 2D Force Matrix. *Nime*, 27–30.

[25] Jorda, S. (2005). *Digital Lutherie: Crafting musical computers for new musics performance and improvisation*. PhD diss.

[26] Kapur, A., Essl, G., Davidson, P., & Cook, P. R. (2003). The Electronic Tabla Controller. *Journal of New Music Research*, *32*(4), 351–359. http://doi.org/10.1076/jnmr.32.4.351.18854

[27] Levin, G. (2000). *Painterly Interfaces for Audiovisual Performance*.

[28] Macon, M. W., McCree, A., Lai, W.-M., & Viswanathan, V. (1998). Efficient analysis/synthesis of percussion musical instrument sounds using an all-pole model. *1998 IEEE International Conference on Acoustics, Speech, and Signal Processing*, *6*, 3589–3592 vol.6. http://doi.org/10.1109/ICASSP.1998.679653

[29] Momeni, A., & Wessel, D. (2003). Characterizing and Controlling Musical Material Intuitively with Geometric Models. *Nime*, 54–62.

[30] Nevile, B., Driessen, P., & Schloss, W. A. (2003). A new control paradigm: software-based gesture analysis for music. *Audio, Transactions of the IRE Professional Group on*, *1*, 360–361. http://doi.org/10.1109/PACRIM.2003.1235791

[31] Ono, M., Shizuki, B., & Tanaka, J. (2013). Touch & activate (pp. 31–40). Presented at the the 26th annual ACM symposium, New York, New York, USA: ACM Press. http://doi.org/10.1145/2501988.2501989

[32] Peltola, L., Erkut, C., Cook, P. R., & Välimäki, V. (2007). Synthesis of Hand Clapping Sounds. *IEEE Transactions on Audio, Speech, and Language Processing*, *15*(3), 1021–1029. http://doi.org/10.1109/TASL.2006.885924

[33] Polfreman, R. (2002). Modalys-ER for OpenMusic (MfOM): virtual instruments and virtual musicians. http://doi.org/10.1017/S1355771802003126

[34] Ren, Z., Mehra, R., Coposky, J., & Lin, M. C. (2012). Tabletop Ensemble: touch-enabled virtual percussion instruments. *I3d*, 7–14. http://doi.org/10.1145/2159616.2159618

[35] Savage, V., Zhang, X., & Hartmann, B. (2012). Midas: fabricating custom capacitive touch sensors to prototype interactive objects. Presented at the UIST '12: Proceedings of the 25th annual ACM symposium on User interface software and technology, ACM Request Permissions. http://doi.org/10.1145/2380116.2380189

[36] Schloss, W. A., & Driessen, P. (2001). New algorithms and technology for analyzing gestural data (Vol. 1, pp. 39–42). Presented at the Communications, Computers and signal Processing, 2001. PACRIM. 2001 IEEE Pacific Rim Conference on, IEEE. http://doi.org/10.1109/PACRIM.2001.953517

[37] Schloss, W. A., Kapur, A., Tzanetakis, G., & Tindale, A. R. (2005). Indirect Acquisition of Percussion Gestures Using Timbre Recognition. *Cim*, 1–8.

[38] Schwarz, D. (2012). The Sound Space as Musical Instrument: Playing Corpus-Based Concatenative Synthesis (pp. 1–6). Presented at the NIME 2012.

[39] Schwarz, D., Beller, G., Verbrugghe, B., & Britton, S. (2006). Real-Time Corpus-Based Concatenative Synthesis with Catart

[40] (pp. 1–7). Presented at the Int. Conference on Digital Audio Effects DAFx-, Montreal, Canada.

[41] Tindale, A. R. (2007). A Hybrid Method for Extended Percussive Gesture. *Nime*, 392–393.

[42] Tindale, A. R., Kapur, A., Tzanetakis, G., Driessen, P. F., & Schloss, W. A. (2005). A Comparison of Sensor Strategies for Capturing Percussive Gestures. *Nime*, 200–203.

[43] Trail, S., Dean, M., Tavares, T., & Odowichuk, G. (2012). Non-invasive sensing and gesture control for pitched percussion hyper-instruments using the Kinect. *Nime 2012*.

[44] Tzanetakis, G., Kapur, A., & Tindale, A. R. (2006). Learning Indirect Acquisition of Instrumental Gestures using Direct Sensors. *Mmsp*, 37–40. http://doi.org/10.1109/MMSP.2006.285264

[45] van den Doel, K., & Pai, D. K. (2003, January 17). Modal Synthesis for Vibrating Objects. Retrieved December 24, 2014, from

[46] Wessel, D. (2006). An Enactive Approach to Computer Music Performance (pp. 1–6). Presented at the Le Feedback dans la Creation Musical, Lyon, France.

[47] Wessel, D. (2009). *SLABS Touch Pads*. YouTube. Retrieved from https://www.youtube.com/watch?v=q_mtCZqN0Ms

[48] Wessel, D., & Wright, M. (2002). Problems and prospects for intimate musical control of computers. *Computer Music Journal*, *26*(3), 11–22.