

A Low-Cost, Low-Latency Multi-Touch Table with Haptic Feedback for Musical Applications

Matthew Montag, Stefan Sullivan, Scott Dickey, and Colby Leider

Music Engineering Technology Group

University of Miami

Frost School of Music

{matt.montag, stefan.sullivan}@gmail.com, d.dickey@umiami.edu, cleider@miami.edu

ABSTRACT

During the past decade, multi-touch surfaces have emerged as valuable tools for collaboration, display, interaction, and musical expression. Unfortunately, they tend to be costly and often suffer from two drawbacks for music performance: (1) relatively high latency owing to their sensing mechanism, and (2) lack of haptic feedback. We analyze the latency present in several current multi-touch platforms, and we describe a new custom system that reduces latency to an average of 30 ms while providing programmable haptic feedback to the user. The paper concludes with a description of ongoing and future work.

Keywords

multi-touch, haptics, frustrated total internal reflection, music performance, music composition, latency, DIY

1. INTRODUCTION

1.1 Motivation

Multi-touch input devices have the potential to serve as highly expressive musical instruments. The plurality of potential control inputs (e.g., position, relative distance, relative rotation, etc.) and the simultaneity of these inputs provides considerable control of parameters compared with many current electronic or acoustic instruments. As a performance device, this provides great potential for interesting and dynamic audio/musical control. Additionally, the large size of many multi-touch tables accommodates larger gestural movements, helping the performer to physically interact with the music, as well as allowing for collaboration and interaction among multiple users simultaneously.

A multi-touch table surface can also be used as a video projection surface, providing visual feedback to the performer without blocking line-of-sight with the audience. In addition to enhancing the performer's experience, multi-touch tables allow audiences to witness performance gestures, providing increased emotional connection between performer and audience [20]. Recent advances in a number of technologies have also added to the growing do-it-yourself (DIY) multi-touch movement [21], a design philosophy under which the work described here falls.

However, camera-based multi-touch tables suffer from input lag and event quantization imposed by the camera's

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME'11, 30 May–1 June 2011, Oslo, Norway.

Copyright remains with the author(s).

frame-capture rate. The input lag may vary based on idiosyncrasies of the camera and host computer configuration. For example, if JPEG frame compression is performed on the camera in order to transfer a 640×480 image at 30 frames per second (FPS) over a USB 1.0 connection, the frame must be decompressed by the host computer before it is further processed, adding a significant delay to the input path. This motivates an investigation of camera selection and host configuration in the pursuit of minimum latency. Additionally, the table interface provides minimal haptic feedback—only the sensation of touch that occurs when a finger is physically in contact with the table's surface.

1.2 Background

Optical-based touch surfaces use infrared light and an IR-sensitive camera to detect the disturbance caused by a touch event. Two common methods are laser light plane illumination (LLP) [22] [18] and frustrated total internal reflection (FTIR) [7]. The LLP method uses infrared lasers to create a laser plane a few millimeters above a translucent touch surface. When users touch the surface, their fingers are illuminated by the light plane and become visible to the camera below. Because the laser plane is situated above the touch surface, the user is able to create a touch event without actually touching the table. FTIR touch surfaces operate by shining infrared light into the edge of a sheet of acrylic. The light is internally reflected by the acrylic, and it escapes only when diffused at the contact area between the surface and the user's finger.

The current project was informed by lessons learned from our first attempt at a multi-touch table (Figure 1) [18], which was constructed using the laser light plane method. The FTIR and LLP methods are similar in cost and operation, but because the FTIR system registers a touch exactly when the user touches the surface, this method is well-suited for percussive musical interaction under strict latency constraints.

1.3 Multi-Touch Latency, Audio, and Haptics

Audio latency can prove a significant impediment to musical performance [25] [16] [6], and multi-touch surfaces are particularly prone to relatively large latency times, as discussed above. In the realm of music, many authors agree that latencies of up to 30 ms between gestural input and sonic result are allowable in most real-time performance situations [14] [17].

Other studies discuss the perception of haptic-audio asynchrony (e.g., [1] [24]) and audio-visual-haptic asynchrony (e.g., [8]). One study notes that haptic-audio asynchrony can be detected with latencies of as little as 2 ms [24] with respect to event time-order, while others [15] [1] note just-noticeable differences (JNDs) of 18–42 ms to haptic-audio events. These numbers give us a reasonable target. The

audio latency of our first-generation system was over 100 ms, giving us much room for improvement in building a new table.

2. OUR MULTI-TOUCH TABLE

In this project, we attempted to leverage an FTIR-based optical tracking surface with a high-frame-rate camera. The system output that drives the audio display is used to simultaneously drive the haptic display, resulting in no latency between audio and haptic feedback systems.

2.1 System Overview

Incorporating lessons learned from our previous multi-touch table and recent literature, we constructed a second-generation table. The table surface is 90×67 cm and was constructed using off-the-shelf materials for around US\$ 300, plus the cost of a short-throw projector that drives the display (Table 1). We edge-light a 3/8" acrylic touch surface with a strip of 850 nm infrared LEDs. The surface is overlaid with a sheet of silicone-treated vellum, which serves as a projection surface and compliant layer. The compliant layer helps touches appear brighter to the camera underneath. Our new table incorporates a PlayStation Eye camera running at 100 FPS. Community Core Vision (ccv.nuigroup.com) is used to process the video input and generate TUIO (tangible user interface object) messages [13]. Max/MSP or Processing receives the TUIO messages, and this software layer performs event logic for turning these messages into audio output. Other software configurations are the subject of ongoing research. The host computer is connected to a multichannel audio interface that drives a 10.2-channel audio display comprised of ten Genelec 8020A active loudspeaker monitors and two Genelec 7050B active subwoofers.

2.2 Adding Haptic Feedback

Haptic feedback can be used to provide an additional dimension of feedback for the performer on a multi-touch table. The inclusion of haptic feedback has been shown to increase performance accuracy significantly [19], and many recent musical instruments engage haptics as a central feature of their interface [4] [10] [5] [3]. Haptic interaction in multi-touch tables has previously been described in the context of physical “pucks” that a user places on the table [23] [12]. More recent haptic multi-touch displays use electrical



Figure 1: Our first-generation multi-touch table, using the Laser-Light Plane (LLP) method.

Table 1: Construction Costs

Component	Cost (US\$)
Polished Acrylic	\$100
LED Strip (2 m)	\$96
Compliant Surface	\$20
Wooden Frame	\$25
Short-Throw DLP Projector	\$400
Playstation 3 Eye Camera	\$27
Power Supply	\$16
Amplifier	\$25
Actuator	\$10
Total	\$719

fields [2], pneumatic pressure [9], and magnetic fields [11] to provide tactile feedback. Our implementation is unique in that we couple a single large tactile transducer capable of producing up to 20 foot-pounds (89 N) of force onto the surface of the table itself, which can achieve a variety of tactile effects ranging from subtle to startling.

We found that the most efficient method of vibrating the multi-touch surface was to couple a tactile transducer (the Aura AST-1B-4) directly to the surface. We had tested haptic feedback with a DC motor driven by a microcontroller that was connected to the host computer via USB, but this configuration exhibited noticeable latency. On the other hand, the tactile transducer which is driven by an audio signal provides near-zero latency and guarantees that haptic events are synchronous with audio output.

The question of when and how to provide haptic feedback to the user is a function of the software application, and may vary from a simple touch response to a method of indicating perceived edges or proximity when the user drags a finger across a virtual surface boundary. For instance, our first approach was to generate a short transient “bump” on the surface every time a touch event occurred.



Figure 2: Second-generation multi-touch table, incorporating FTIR sensing, haptic feedback, and a simplified two-channel audio display used during development. Note the transducer coupled directly to the surface of the table at the top right.

In terms of user experience, this provides only marginally more information than the mechanical feedback of touching a motionless table, but informal tests suggest this improves users' perception of multimodal simultaneity. We describe ongoing efforts to incorporate haptic feedback into a musically expressive device later in this paper.

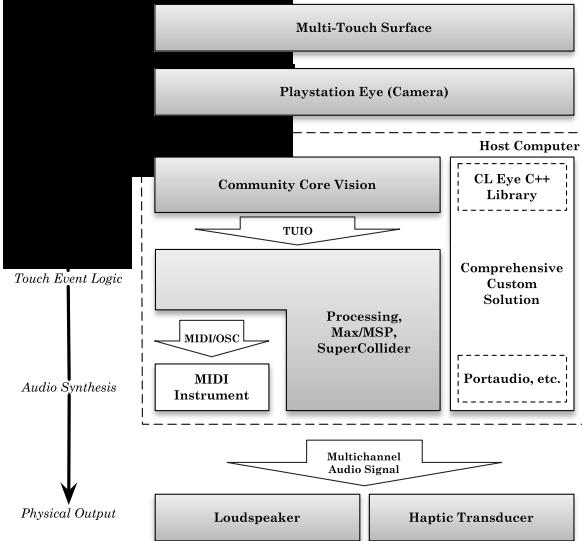


Figure 3: System-level overview of our multi-touch table. Boxes in gray are currently implemented, and boxes in white are potential alternatives.

3. LATENCY TESTS

To assess the audio and video latency present in this new multi-touch table, a variety of tests were performed on the table and other existing systems for comparison.

3.1 Video Latency Tests

Video tests were performed prior to audio latency tests to determine the delay contributed by different cameras and display devices. The test was carried out in the following manner. The camera and display device under question were connected to the host computer. A frame-counter window and camera monitor window were shown side-by-side on the display device. The camera was aimed at the display so that both the frame counter and the video-overlay window were visible, creating a video loop. It was then possible to photograph the display, recording the counter and a time-delayed version of the same counter in one image, allowing a simple subtraction of these two counters to determine the video latency. We tested the Unibrain Fire-i firewire camera (maximum 30 FPS) and the Sony PlayStation Eye USB 2.0 camera (maximum 100 FPS) with a Toshiba TDP ET-10 DLP projector (60 Hz refresh rate), a Dell E173FP LCD display (60 Hz), and an E-machines CRT monitor (100 Hz). The results were averaged over a minimum of 25 measurements. This testing setup is shown in Figure 4.

3.2 Video Latency Results

The results of the video latency tests showed that the CRT was the fastest display device, and the PlayStation Eye was the fastest capture device. The results can be seen in Figure 5. The Unibrain Fire-i at 30 FPS and PlayStation Eye at 100 FPS resulted in a video loop latency of 70 ms and



Figure 4: Experimental setup for testing latency times for our second-generation multi-touch table and other devices.

10 ms, respectively. It is clear that the Unibrain Fire-i suffers more input lag than can be explained by the frame rate alone, which accounts for only 23 ms of the difference. Although the display device itself does not impact latency of the audio output, it is important to note the large disparity in latency among the tested displays. The DLP projector exhibited a video input lag around 80 ms. This delay is imparted by a particular digital image-processing circuit in this Toshiba projector and is not inherent to all projectors or DLP technology.

3.3 Audio Latency Tests

We measured the delay between input touch event and audio output of our low-latency multi-touch table configuration and benchmarked our system against several other multi-touch and MIDI devices. The following devices were tested: Korg Triton keyboard, USB MIDI keyboard, Apple iPad, Apple iPod Touch, HTC Hero Android-based smartphone, and our new multi-touch table. Multi-touch and MIDI tests were conducted with an Intel Macbook 2.2 MHz Core 2 Duo running Windows XP, connected to an Echo AudioFire 12 firewire audio interface configured with a 256 sample buffer at 44.1 KHz. Apple iPad tests were performed with the apps Drum Kit Pro and I Can Drum. The Apple iPod was tested with Drum Kit Pro; the Android phone was tested with DrumKit.

In all test cases, we placed a microphone near the control surface and recorded a stereo audio file, with the touch/ key-press noise on the left channel and the corresponding audio output on right channel. We then used a sound file editor to measure the time elapsed between the trigger event and the sound output, averaging at least 25 trials. In the case of the keyboards, the impact of the key on the keybed was chosen to mark the trigger event. This process is illustrated in Figure 6.

3.4 Audio Latency Results

The results of the audio latency tests show that the Korg Triton had latency between 0 ms and -1 ms. Note that this is a valid result given the testing method, although it indicates that the electrical contact for the Triton's keys occurs somewhere before the fully depressed position. The Triton has a known latency of less than 2 ms, and our result is included as a reference point for the test method. The Android and Apple devices were shown to have average laten-

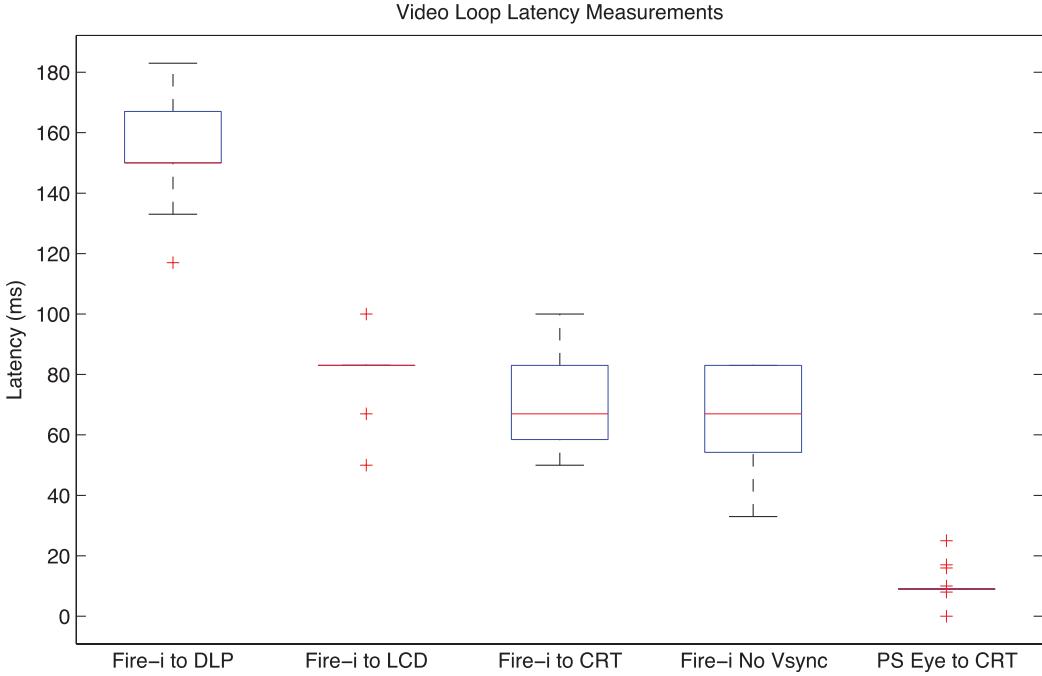


Figure 5: Results of video latency tests.

cies above 50 ms, depending on the application. The multi-touch table with Fire-i camera exhibited a higher latency than the Apple iPad and Apple iPod Touch. The measured latency of the multi-touch table using the PlayStation Eye was much lower and surprisingly comparable to the USB MIDI keyboard. Our average overall latency for the table was 30 ms. These results are summarized in Figure 7.

4. ONGOING AND FUTURE WORK

Our team is currently working on several improvements and related projects. These include custom music applications that take advantage of haptic multi-touch interaction, improvements to the table itself, assessment and improvement of the table's haptic feedback system, and other long-term projects. Each of these is described below.

4.1 Software Applications

Our primary ongoing project is to create software applications that take advantage of the multi-touch table as a musical performance instrument. We emphasize applications that use multi-touch input in such a way that could not be easily duplicated by point-and-click interaction, for example, the simultaneous manipulation of several control points along a virtual vocal tract, or the simultaneous control of several harmonic overtones in an additive-synthesis instrument. We also wish to take advantage of our table's low-latency characteristics and we are creating responsive virtual instruments for live performance.

4.2 Table Improvements

As seen in the video latency results above, there is considerable room for improvement in the high latency of the tested DLP projector. The low-latency audio signal should be accompanied by low-latency visual feedback. We are actively

looking for low-cost short throw projectors with exceptional latency characteristics.

We are also currently working by trial and error to improve the performance of the compliant surface used in our FTIR table. Silicone applied to the bottom of the vellum occasionally sticks to the acrylic surface in the area of a touch, preventing the recognition of subsequent touches in that area. This deficiency impacts musical performance, but could be resolved by incorporating the laser light plane method in future experiments.

4.3 Assessment and Improvement of the Haptic Subsystem

The table described here delivers a transient vibration on its surface each time a user touches it, and, as mentioned, this anecdotally improves perception of event simultaneity. A more meaningful implementation might be to provide varying degrees of haptic feedback as a musician drags a finger across the table, to indicate certain particular tasks/events. For instance, the table might vibrate every time a certain sound event begins and/or terminates, or every time a finger moves from one parameter to another, or indicating certain thresholds of parameters/events. We are also investigating haptic feedback to impart additional information about a sound as it is being “scrubbed.” For example, search time when dragging across a graphical waveform representation of a long-duration sound file to find a particular section may be minimized by using haptic displays of simulated surface texture and intensity to represent a particular feature vector, e.g., local novelty. Finally, we are also investigating the creation of physical “buttons” on our multi-touch surfaces by using the haptic subsystem to generate Chladni patterns on the surface of the acrylic.

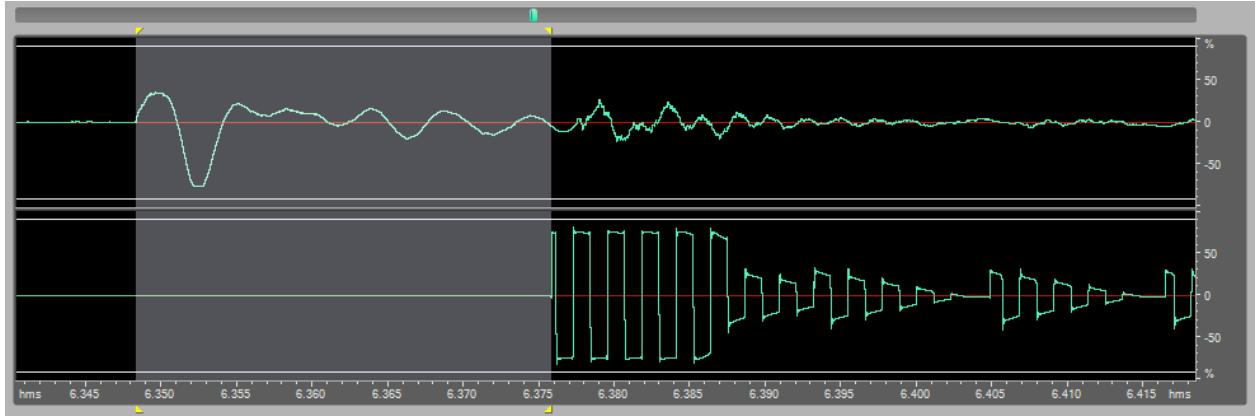


Figure 6: Measuring a 27-ms duration between a finger tap on the table surface, recorded in the left channel, and the synthesized sound output, recorded in the right channel.

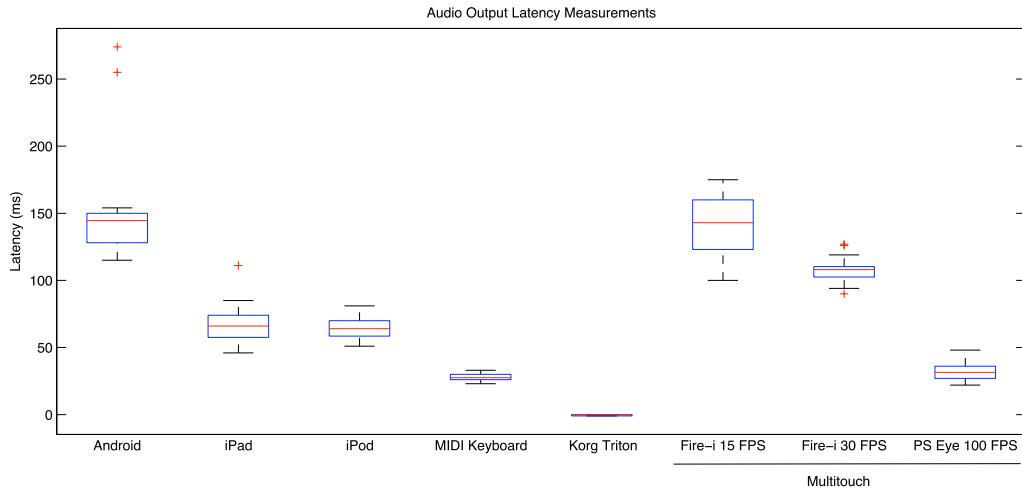


Figure 7: Results of audio latency tests.

4.4 Future Applications

This multi-touch table with haptic feedback has the potential to be deployed in educational environments, in addition to its current use as a software-synthesis performance controller. We are beginning to explore ways in which haptic-multi-touch interfaces can lead to engaging, fun, and collaborative music-making for children, in both schools and museums. We are also beginning work on an interactive soundscape-exploration system in which geographical maps, satellite images, and multi-channel soundscape recordings can be quickly navigated, explored, compared.

5. CONCLUSIONS

Our goal was to produce an economical multi-touch table with haptic feedback that exhibited low enough latency to be useful as a musical instrument. Our system succeeded in reducing average latency to 30 ms. There still seems to be some debate as to the exact JND for multimodal feedback, but subjective reports indicate that the inclusion of haptic feedback can improve the perception of simultaneity in new musical interfaces. This paper demonstrates proof-of-concept of the multi-touch table as a low-cost computer-music instrument with haptic feedback. We are creating a suite of software instruments for the table that we hope

will leverage the multi-touch control paradigm with new tools for musical performance and composition, as the integration of multi-touch technology with haptic feedback provides many opportunities for creative exploration.

6. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant No. IIS 0757552 and by grants from the University of Miami. Thank you also to Pat O’Keefe and Mark Freeman for their work on our first multi-touch table.

7. REFERENCES

- [1] B. Adelstein, D. Begault, M. Anderson, and E. Wenzel. Sensitivity to haptic-audio asynchrony. In *Proceedings of the 5th international conference on Multimodal interfaces*, pages 73–76. ACM, 2003.
- [2] O. Bau, I. Poupyrev, A. Israr, and C. Harrison. Teslatouch: electrovibration for touch surfaces. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 283–292. ACM, 2010.
- [3] E. Berdahl, H. Steiner, and C. Oldham. Practical hardware and algorithms for creating haptic musical

- instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME-2008), Genova, Italy*, 2008.
- [4] D. DiFilippo and D. Pai. The AHI: An audio and haptic interface for contact interactions. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 149–158. ACM, 2000.
 - [5] M. Eid, M. Orozco, and A. El Saddik. A guided tour in haptic audio visual environments and applications. *International Journal of Advanced Media and Communication*, 1(3):265–297, 2007.
 - [6] C. Gunn, M. Hutchins, and M. Adcock. Combating latency in haptic collaborative virtual environments. *Presence: Teleoperators & Virtual Environments*, 14(3):313–328, 2005.
 - [7] J. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 115–118. ACM, 2005.
 - [8] V. Harrar and L. Harris. The effect of exposure to asynchronous audio, visual, and tactile stimulus combinations on the perception of simultaneity. *Experimental Brain Research*, 186(4):517–524, 2008.
 - [9] C. Harrison and S. Hudson. Providing dynamically changeable physical buttons on a visual display. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 299–308. ACM, 2009.
 - [10] G. Huang, D. Metaxas, and M. Govindaraj. Feel the fabric: an audio-haptic interface. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 52–61. Eurographics Association, 2003.
 - [11] Y. Jansen. Mudpad: Fluid haptics for multitouch surfaces. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems, CHI EA '10*, pages 4351–4356, New York, NY, USA, 2010. ACM.
 - [12] S. Jordà, G. Geiger, M. Alonso, and M. Kaltenbrunner. The reactable: exploring the synergy between live music performance and tabletop tangible interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction, TEI '07*, pages 139–146, New York, NY, USA, 2007. ACM.
 - [13] M. Kaltenbrunner, T. Bovermann, R. Bencina, and E. Costanza. TUOI: A protocol for table-top tangible user interfaces. In *Proc. of the The 6th International Workshop on Gesture in Human-Computer Interaction and Simulation*. Citeseer, 2005.
 - [14] N. Lago and F. Kon. The quest for low latency. In *Proceedings of the International Computer Music Conference*, pages 33–36. Citeseer, 2004.
 - [15] D. Levitin, K. MacLean, M. Mathews, L. Chu, and E. Jensen. The perception of cross-modal simultaneity. *International Journal of Computing Anticipatory Systems*, 2000.
 - [16] T. Maki-Patola and P. Hamalainen. Effect of latency on playing accuracy of two gesture controlled continuous sound instruments without tactile feedback. In *Proc. Conf. on Digital Audio Effects, Naples, Italy*, 2004.
 - [17] T. Maki-Patola and P. Hamalainen. Latency tolerance for gesture controlled continuous sound instrument without tactile feedback. In *Proc. International Computer Music Conference (ICMC)*, pages 1–5. Citeseer, 2004.
 - [18] A. Mattek, M. Freeman, and E. Humphrey. Revisiting Cagean Composition Methodology with a Modern Computational Implementation. In *Proc. NIME 2010*, 2010.
 - [19] S. O’Modhrain and C. Chafe. Incorporating haptic feedback into interfaces for music applications. In *Proceedings of the International Symposium on Robotics with Applications, World Automation Conference*, 2000.
 - [20] K. Scherer and M. Zentner. Emotional effects of music: Production rules. *Music and emotion: Theory and research*, pages 361–392, 2001.
 - [21] J. Schöning, J. Hook, T. Bartindale, D. Schmidt, P. Oliver, F. Echtler, N. Motamedi, P. Brandl, and U. Zadow. Building Interactive Multi-touch Surfaces. *Tabletops-Horizontal Interactive Displays*, pages 27–49, 2010.
 - [22] A. Teiche, A. Rai, C. Yanc, C. Moore, D. Solms, G. Cetin, J. Riggio, N. Ramseyer, P. D’Intino, L. Muller, et al. Multi-touch technologies. *NUI Group*, 2009.
 - [23] L. Terrenghi, D. Kirk, H. Richter, S. Krämer, O. Hilliges, and A. Butz. Physical handles at the interactive surface: exploring tangibility and its benefits. In *Proceedings of the working conference on Advanced visual interfaces, AVI ’08*, pages 138–145, New York, NY, USA, 2008. ACM.
 - [24] K. Walker, W. Martens, and S. Kim. Perception of Simultaneity and Detection of Asynchrony Between Audio and Structural Vibration in Multimodal Music Reproduction. In *Proceedings of the 120th Convention of the Audio Engineering Society, Paris, France*, 2006.
 - [25] E. Wenzel. Analysis of the role of update rate and system latency in interactive virtual acoustic environments. *AES Preprint*, 1997.