Determining the Feasibility of Networked Musical Performances over WANs, LANs, and WLANs

Matthew Teeter and Daniel Lindsey
Information & Computer Science

Authors

Matthew Teeter was looking for research topics that combined computers and music when he discovered Professor Dobrian's similar interests. *(continued on page 61)*.

Daniel Lindsey credits his undergraduate research experience with developing his communications skills and work ethic, skills he will use throughout his future years in the workplace. *(continued on page 61)*.

Key Terms

- LAN
- Latency
- MIDI
- Networked Musical Performance
- WAN
- WLAN

Abstract

With the increasing prevalence of broadband Internet connections, people are exploring new applications that rely upon a low-latency communication medium. One such application is networked musical performances, in which physically-separated performers simultaneously play instruments that are connected via the Internet. In this study, we combined empirical data about latency (delays) inherent in the transmission of information via the Internet with psychoacoustic information about the ability of players to synchronize their playing and discern independent musical events. We used this information to decide how feasible it would be to conduct networked musical performances over local-area networks (LANs), wireless local-area networks (WLANs), and even wide-area networks (WANs). The latency data we collected implies that successful networked performances can occur if the network latency is less than the time needed to perceive musical events as simultaneous, and less than the ability of the players to synchronize. These stipulations are usually met with performances between two locations that are less than 400 miles apart (where network latency is below 20 ms). By conducting our tests on commonly-available hardware and software, we have shown that networked performances are accessible to household users and university performers alike.

Faculty Mentor

Increasingly musicians are interested in the potential applications of transmitting musical information wirelessly from one computer instrument to another in real time. The main problem is the delay introduced during transmission, which might cause noticeable timing problems between performers, especially when sending data over the Internet. We researched and tested two related questions: 1) How much timing discrepancy is deemed acceptable by musicians in a normal musical performance situation? 2) Would the timing discrepancies introduced due to wireless data transmission be acceptable to musicians in a real-time networked performance? I&CS students Matthew Teeter and Daniel Lindsey worked with me to review relevant research, design and implement controlled yet realistic experiments to test the questions, and summarize the results.

John Christopher Dobrian
Claire Trevor School of the Arts

Introduction

The concept of networked musical performances has the potential to revolutionize music teaching, rehearsal and performance. For instance, the Yamaha Disklavier piano has been used to enable a piano teacher to give lessons remotely to a student hundreds of miles away (Campbell, 2004). The same technology can allow performers in different locations to play together or one performer to play multiple instruments in different physical locations at the same time—enabling a single musical performance to reach a larger live audience. In a scenario like this, nearly an infinite number of instruments in different locations from the performer could replicate the performance of one pianist. The instrument(s) at the remote location could receive information from the instrument the performer was physically playing, and instantaneously replicate the same musical events locally. In addition, networked musical performances allow musicians to rehearse pieces much more conveniently, since they do not have to travel to the same location to practice together. This benefit saves transportation costs, time, and energy. The feasibility of real-time networked music performance also gives rise to entirely new paradigms for performer interaction, such as group improvisation by performers in remote locations.

All that is needed to create a networked performance is a computer and a MIDI-enabled instrument. MIDI stands for Musical Instrument Digital Interface, and is a standardized way of encoding musical events such as note presses and pedal presses into a concise format. MIDI is commonly used for networked musical performances instead of streaming audio data because MIDI requires much less bandwidth and can allow a performance to be perfectly replicated at another location. Transmitting audio data in real time to distant locations usually does not work as well because of bandwidth limitations.

The ultimate goal of a networked musical performance is a high degree of transparency, which attempts to minimize noticeable problems with a performance. A highly transparent system allows the performers to play together as if they were in the same room. The key to maintaining a transparent system is to keep latency times to a minimum. This is especially important in musical performances because the slightest delay time can propagate back and forth between locations and interfere with accurate musical performance. Previous studies have shown that delays of 200 to 300 ms are the most disturbing to performers, and such delays make it very difficult to play notes in the correct rhythm (Willey, 1990). The best case scenario, however, is to keep delay times as close to 11.5 ms as possible. Previous research at Stanford showed this to be the optimal value for performers attempting to keep an accurate tempo (Chafe et al., 2004).

To understand why even small delays are harmful to a networked performance, consider this scenario. Imagine two performers are trying to play a duet that involves playing four (different) notes per measure. Performer 1 starts a networked performance by playing the first measure alone. Performer 2 is trying to synchronize with Performer 1, and begins playing in the same tempo as Performer 1. All networks have some amount of delay due to physical properties of the connection medium and laws of nature; we call this delay network latency. Because of network latency, the notes played by Performer 1 reach Performer 2 after a short time. Likewise, the notes played by Performer 2 reach Performer 1 slightly after they were actually played. Performer 1 listens to the timing of notes played by Performer 2, and slightly adjusts his/her own tempo to stay synchronized with Performer 2. Even if both performers have instantaneous reaction times and play notes at the exact same time, delays in the network will cause the other person's notes to arrive slightly after the original person played his/her own. This results in each performer hearing the notes they played, followed slightly by the notes of the other person. Since this complicates timing and tempo tracking, latency must be kept below the time a human is able to distinguish two musically independent events. In this case, the delay between when the first person plays notes and when the other person's notes arrive would be indistinguishable, giving the impression that both performers were playing together in the same physical location.

We included three different types of networks in our experiments in order to measure and compare the latency differences present in each type of network. We needed this information to determine if networked performances were feasible over LANs, WLANs, and WANs. LANs, or local-area networks, are smaller networks typically found in homes or buildings. Because they span relatively short distances, delays were expected to be less than 10ms. WLANs, or wireless local-area networks, are similar to LANs in many respects except that information is sent via radio waves instead of over a wire. This results in slower transmission time than LANs. Finally, wide-area networks (WANs) span large geographic locations, such as states. As a result, the latency on this network is higher than that of LANs and WLANs. Because of this, we were most interested in studying latency on WANs because the delays on these networks vary widely depending on distance. We sought to determine the maximum distance two locations could be apart while still maintaining a level of latency conducive to networked performance.

Previous Work

The concept of a networked musical performance is not new. As early as the late 1970s, individuals in the League of Automatic Music Composers were investigating the concept of using networked computers to create and perform music (Bischoff and Brown). Members of this group typically brought their computers to the same room and had their programs put on a musical concert. Each person programmed his computer to obtain information from other computers it was linked to, allowing the machines to "improvise" together. The group known as the Hub emerged in the late 1980s, and they were the first to have computers playing music together from different buildings in the same city. The work done by these groups was revolutionary during that time, but their work focused on creating new styles of computer music, instead of allowing synchronized real-time networked performances between human performers.

When the Internet began to be widely adopted in the 1990s, new possibilities enabling long-distance performance emerged. The concept of a networked performance expanded from that of a local event of computers linked by MIDI cables in a room to an inter-city phenomenon in which electronic instruments communicated via the Internet. In contrast to the networked computer performances of the previous decade, which required small amounts of bandwidth and no precise timing, real-time performance consumes more bandwidth and requires higher-speed connections. This is because in a real-time performance, the success of the concert depends upon having information transmitted in a regular, timely fashion. In the Hub's performances, it was not disastrous if a program received information from another computer 200 ms late. This is because the programmers would know about such latency and could compensate ahead of time. On the other hand, delays of this magnitude could ruin a real-time performance in which precise timing is of the utmost importance. Moreover, human performers rely on auditory feedback while playing, and the slightest delays can disrupt one's concentration. As a result, previously-avoidable problems such as propagation delay must be dealt with using more elaborate solutions when the connected devices are miles apart instead of in the same room.

An early demonstration of a real-time networked musical performance occurred in 2001. Jazz pianists Kei Akagi and Anthony Davis performed a dual piano concert from two cities simultaneously, with Akagi playing at UC Irvine and Davis playing at UC San Diego (Dobrian). This networked performance was as synchronized as possible at the time; and, although no precise data was recorded, anecdotal evidence suggested the delay due to Internet latencies was about 10 ms.

Recently, UC Berkeley professors John Lazzaro and John Wawrzynek implemented a system allowing networked performances using the Real-time Transport Control Protocol (RTCP) (Lazzaro and Wawrzynek, 2001). Their software ran on the Linux operating system and was tested between Berkeley, Stanford, and Caltech. While musicians may not always think about it, acoustical delays are present not only in a networked performance, but in real performances as well. Players on a stage may be separated by several meters, so they use the visual cue of the conductor to keep synchronized. Keeping this in mind, Lazzaro and Wawrzynek reasoned that the network delays observed could be combined with information about the speed of sound to determine the "distance" that networked performers would have between them if they were in the same physical location. Therefore, given a latency time in milliseconds, the equivalent distance between the two performers could be calculated. They concluded that such networked performances were feasible because the average observed latency was 14 ms, equivalent to performers being separated by 4.8 meters. Musicians often play together with ease at such distances.

In our study, we sought to determine if such musical performances were feasible using computer hardware and software that is more readily available to end users. Although Linux has made impressive strides in improving the user experience, most musicians desiring to participate in a networked performance would own a Windows or Mac computer. Thus, to see how well a networked musical performance could work on the Mac OS X operating system, we conducted our latency tests using Macs running OS X 10.4. This version allows one to create a virtual MIDI device that is connected over the network. Apple claims that OS X's audio platform, CoreAudio, was designed with the goal of keeping MIDI latency to a minimum. Would these optimizations in CoreAudio allow a networked musical performance to take place using commonly-available Macintosh computers, running Mac OS X 10.4?

Our study further sought to gather and thoroughly analyze empirical data concerning the cognitive, physical, and technical latencies involved throughout the entire process of a networked performance. We also investigated if extremely long-distance communication was feasible (ranging from hundreds to thousands of miles apart) and, in addition, studied how feasible a networked performance would be over a WLAN. Wireless capabilities enable many exciting possibilities for computer music concerts, but we will not discuss the significance of these capabilities at this time.

Methods

There were three factors that we needed to test: the average delay times over computer networks, the average precision with which two performers could synchronize their playing, and the average ability of listeners to discern separate musical events. Our reasoning was that if the average delays over a network were less than the time needed for pianists to synchronize, and those network delays were also less than the time needed for performers to perceive independent musical events, then network delays should not impede networked musical performances.

To test how well pianists could synchronize their playing in the best-case scenario (i.e. in the same room), we developed a program in the Max/MSP programming language that would allow us to store the time discrepancy in milliseconds between two performers pressing the same note several octaves apart on a keyboard. The measurement program also allowed us to control various cues from which the performers would set their tempo. The cues we used included a visual metronome, audible metronome, or both at once. The audible cue was similar to a metronome, whereas the visual cue was similar to a conductor. The pianists observed a laptop that displayed a red circle in one of four locations to indicate the beat. Finally, to determine how well pianists could begin playing together, we kept track of their ability to press a note simultaneously after they cued each other with a head nod. When using the head nod, pianists were specifically instructed not to follow a set tempo, so that their ability to start a performance could be observed. The pianists attempted to synchronize their playing using the cues at 80, 100 and 120 beats per minute. We tested a variety of pianists ranging from casual players to those majoring in Piano Performance. Each of the test subjects had played the piano for at least five years and many had taken formal lessons during that time. To keep the performance material simple, we tested the pianists' abilities using only a C major scale. For each test, the pianists played the C scale up an octave, then back down, and repeated this three more times. Thus, for each run, the time discrepancy between a total of 57 notes was recorded. For a more detailed description of this experiment, please see Appendix A.

The concept of propagation delay was not considered in this experiment because the performers were in the same room. Because of this, each performer heard the notes played by the other performer immediately. Thus, the growing note-transmission delay phenomenon described previously did not arise in this situation. We did not test the ability of performers to synchronize in the presence of growing delay because it was already known that such delays make network performances impractical (Willey). Instead, we wanted to use information on pianists' synchronization ability to determine conditions that permit successful networked performances (in terms of distance, medium, delay in milliseconds, etc).

A networked performance requires a continuous stream of musical information to be delivered to all participants. But how far apart can these musical events be without noticeably affecting the performance? If the network delay time is less than the time needed to perceive separate musical events, then the performance would appear to be identical to a conventional performance, which is the ultimate goal of this system. To test how well the human ear can distinguish separate musical events, we conducted a test in which listeners closed their eyes and listened to two piano tones that began a few milliseconds apart. Listener subjects raised their hand if they believed the sounds to be distinct in their starting times. One experimenter controlled the program that generated the tones and observed the listeners' responses. We shifted the delay between the tones from 10 ms to 30 ms. This listening test used sampled piano sounds to ensure a fast attack time. A sound with a slower attack time might have skewed the results because it would be more difficult to tell exactly when a sound occurred.

Next, to observe the amount of delay inherent in networks across the United States, we used the ping command to record the roundtrip times between UCI and a variety of locations. We developed a Visual Basic program that used batch commands to organize and record ping results. These results were then used to determine the average delay times between various locations. To allow for variances in daily Internet traffic, we ran the program five times per day, evenly spaced out from 9 A.M. to 9 P.M. To allow for weekly variances, we ran the program every day for one month. We tested network latencies by pinging the following areas: the same building, across the UCI campus, UCLA, UCSD, a residential area in San Diego, UC Merced, UC Berkeley, University of Texas, and New York University. We included the residential area to provide insight into what kind of delays would be involved when communicating with a location off of the high-speed Internet2 network that links the universities. Typical home users would not have access to such a high speed network, and we wanted to observe the extent to which latency increased when using a slower, residential network.

Lastly, we wanted to determine how the connection medium would affect network latency. Wireless Internet access is becoming more and more commonplace, especially with new musical instruments like the Disklavier Mark IV, the first piano with built-in wireless communication capabilities (Yamaha Corporation). To see if wireless communication would impede a networked performance, we ran tests on two Macintosh G4 Powerbooks running Mac OS X 10.4.5. First, we tested latency when the computers were connected to the LAN with an Ethernet cable, and then we tested again when the computers were connected through the wireless LAN using the Airport wireless Ethernet card. In this configuration, we pinged the other computer repeatedly, tried sending a three-byte MIDI message once per second using the MXJ net.udp.send/recv object (in Max/MSP), and also tested latency by sending a three-byte MIDI message once per second using the operating system's built-in MIDI networking technology. The MIDI testing used Max/MSP version 4.5.5. For each test, we took three minutes of data and averaged the results.

Results

After several days of experimentation with more than six pianists, we recorded the ability of two pianists to synchronize their playing when using various cues (Table 1).

Table 1

Average discrepancy between two pianists attempting to synchronize using various cues. These results are the averages of all pianist groups who participated in the study.

Cue Type	Avg. Discrepancy (ms)	Standard Dev. (ms)
Sonic	24.56	6.97
Visual	34.81	12.05
Sonic and Visual simultaneously	29.22	13.90
Head nod	36.45	5.62
Overall	30.06	11.89

On average, the pianists could play a note together within approximately .03 seconds (30 ms) of each other. If the network delays were less than 30 ms, the quality of a networked performance would not improve, since the pianists could be the limiting factor in that case. Thus, network delay times greater than 30 ms do indeed pose a problem for networked musical performances.

We combined this information with the results of our musical perception test. The subjects we tested were able to distinguish musical events that were approximately 20 ms apart, but failed to do so if the musical events were less than 20 ms apart. These results agree with earlier studies done by Tanaka (2000) and Winckel (1967), which produced similar findings. Therefore, it is reasonable to assume that if the average network delay is less than 20 ms, we can expect to have a high-quality networked musical performance. However, this perception test was an artificially controlled situation, in which the subjects were concentrating on listening for two notes, instead of a typical musical setting where many notes are heard in rapid succession. We can assume that humans can notice a difference of 20 ms in a controlled environment, but that slightly longer delays would be tolerable in a more complex musical context.

After a month of testing latency in the networks to various universities from UCI with the ping command, we averaged the roundtrip times. The roundtrip time is the time between when a packet is sent from the local computer and when the remote computer's response arrives at the local computer. These times for each location are shown in Figure 1.

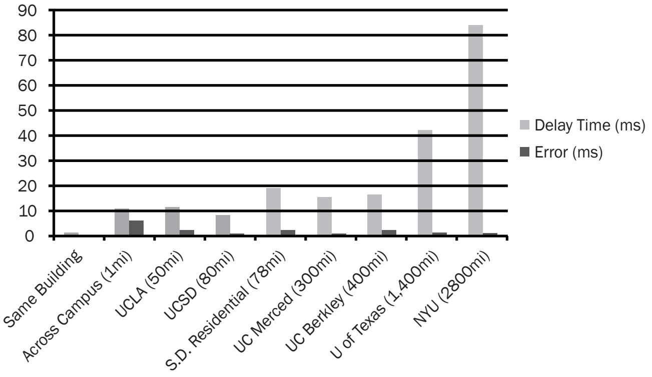

Figure 1

Average roundtrip time to various locations from UCI, determined using the ping command

Because of very little variance in the data between times of day, only the overall average delay times are shown here. In addition, there was no noticeable difference between weekends and weekdays. The delays were occasionally slightly longer on weekday mornings (9 A.M.), but since most performances would likely occur later in the day, this is not a major concern.

The average network delay times for locations within California were less than 20 ms. This means that a networked musical performance is certainly a possibility when performing with someone who is 400 miles away. On the other hand, the average delay to the out-of-state destinations, the University of Texas and NYU, was about 40 ms and 80 ms respectively. While it may be possible to conduct a networked performance with higher delay times like these, such a performance would lack the seamlessness and fluidity that performers and audiences expect of a conventional performance. It is also interesting to note that delay times of the residential area in San Diego were about twice as high as those of UCSD. This implies that a networked musical performance may still be possible over a typical Internet connection if the performance locations are close enough (100 miles or less), but will likely encounter a higher degree of delay, resulting in lack of synchronization.

Finally, we come to the question of wired vs. wireless. We found that wireless communication is slightly slower than a wired medium, but that moving to a wireless medium only increased latency by a few milliseconds (Table 2). Occasionally, however, there would be spikes in latency when using the wireless medium, likely caused by collisions of packets. Because of this, we recommend using a wired connection for networked performances when at all possible to ensure the most reliable connection.

Table 2

Network latency determined using various methods on a LAN and WLAN

Test Method	Wired delay using LAN (ms)	Wireless delay using WLAN (ms)
Ping	0.4	3 (occasionally spiked to 24)
Max/MSP UDP objects	7	11
OS X MIDI Networking	4	7 (occasionally spiked to 28)

It is also important to notice that the OS X MIDI networking and Max/MSP MIDI objects had higher delays than the ping times. This is due to a greater amount of overhead involved in the transmission protocol. For example, the ping command uses ICMP Echo Request and Reply messages, which are small and require little processing, whereas typical network applications use TCP (Transport Control Protocol) or UDP (User Datagram Protocol), which require additional time to package and process because of built-in mechanisms for error correction, flow control, and congestion control. The Max/MSP UDP send and receive objects tended to have slightly more latency than OS X's MIDI networking capability, which suggests that the Max/MSP UDP objects (mxj net.udp.send and mxj net.udp.recv) may send information less frequently than the Max/MSP MIDI objects (notein, noteout, midiin, midiout, etc.). This would explain why the latency times were reduced when using virtual MIDI devices, even though the Max/MSP environment was used for testing both the UDP objects and the OS X MIDI networking latency.

Conclusion

This study sought to produce empirical data about cognitive, physical and technical delays involved in a networked musical performance to determine if such performances were feasible across various types of networks. Latency hampers a smooth networked performance, and is caused by delays in the instrument itself, processing time in the computer, ability of the players to synchronize, and delays in the network. We found that listeners are only able to cognitively discern independent musical events when the events are at least 20 ms apart. Performers were only able to play within approximately 30 ms of each other in the same room, although this reached as little as 14 ms depending on skill level and tempo. There were only a few milliseconds of computer latency involved in processing the incoming messages. Finally, the network delays ranged from less than 10 ms on a LAN or WLAN to greater than 30 ms using the WAN. Collectively, the data implies that successful networked performances can occur if the network latency is less than the time needed to perceive musical events as simultaneous, and less than the ability of the players to synchronize. These stipulations are usually met with performances between two locations less than 400 miles apart (network latency < 20 ms).

By conducting our tests on commonly-available hardware and software, we have shown that networked performances are accessible to both household users and university performers alike. We hope that with the escalating adoption of broadband Internet connections, an increasing number of amateur and professional musicians will use networked musical performances to take advantage of the many benefits such technology brings.

Our work has primarily focused on MIDI streams that require low amounts of bandwidth. Streaming of audio and video signals continues to be a more challenging problem because these streams consume much more bandwidth than MIDI does. As higher-speed networks continue to evolve, and more efficient video and audio codecs are developed, future musicians may use audio and video streams for even more immersive networked musical performances.

Appendix A: Detailed Experiment Protocols

We conducted the following tests on pianist synchronization ability to determine how well pianists can play together in the best circumstances (i.e. in the same room). The program we made would watch for the first note in the sequence to be played by either person. It would record the time when this happened, and wait for the same note to be played two octaves apart by the second pianist. The time difference would be recorded. This process would repeat for every note in the sequence. At the end of the test run, the delay times would be averaged.

Part 1: Audible Cue Only

Two pianists sat side by side on a piano bench in front of a keyboard. They were instructed to play the C major scale upward with their right hand. One pianist would start at middle C, while the other would start at the C two octaves lower. For each test, the pianists played the C scale up an octave, then back down, and repeated this three more times. Thus, each person played a total of 57 notes each test. Pianists were instructed to play one note per beat. They tried to play each note together, as closely as possible. There were also two measures of lead-in for each test, so the pianists started playing at the beginning of the third measure. For this test, pianists listened to an audible cue program, which played a high C on beat 1 and a C an octave lower on beats 2, 3, and 4. For each experiment, the program that displayed the collected data for that run was hidden, so that the pianists would not concentrate on judging their performance while playing. This experiment was repeated three times per pianist group, using metronome speeds of 80, 100, and 120 beats per minute (bpm).

Part 2: Visual Cue Only

This experiment followed the same format as the previous one, except the pianists relied upon a visual cue instead of an audible cue. A laptop was placed in front of the pianists, which ran a program that imitated a conductor. A large red dot appeared in one of four locations (bottom, left, right, top), indicating the beat. This experiment was repeated three times per pianist group, at 80, 100, and 120 bpm.

Part 3: Audible and Visual Cue

This experiment followed the same format as the previous ones, except the pianists relied upon both a visual cue and an audible cue. The metronome was playing at the same time as the pianists were observing the laptop conductor. This experiment was repeated three times per pianist group, at 80, 100, and 120 bpm.

Part 4: Ability to Start in Unison

This experiment followed the same format as the previous ones, except the pianists were instructed not to follow a set tempo. One pianist would use a head nod and press one note, and the other pianist would watch his or her partner and try to play their own note (two octaves apart) at the same time. After a short pause, this process would be repeated, entirely out of tempo, because we were trying to measure how well pianists could start a performance together. We measured 57 starts in total per pianist group.

Matthew Teeter Biography

(continued from page 55) He promptly started working on this project, which he says has helped him refine his interests and discover a topic about which he is deeply passionate. One of the highlights of Matthew's research experience has been the relationship he has developed with Professor Dobrian throughout the project. Matthew hopes to work in the music software industry, creating software that promotes musical education and music in general. His advice to potential researchers is, "By all means, go for it!"

Daniel Lindsey Biography

(continued from page 55) For his project, working alongside his peer, Matthew Teeter, Daniel designed new experiments aimed toward discovering knowledge, rather than following established techniques. He enjoyed working with MaxMSP to create the software they used in their studies, and found the weekly meetings with his fellow researcher and their mentor to be consistently valuable. Daniel's advice to new researchers is to make sure to find a project you really enjoy, something that will be "more of a hobby and less of a chore."

Acknowledgements

We would like to thank everyone involved with this project, including our correspondent Ben Israel from Yamaha Corporation of America, Andrew Dahlin for the use of his laptop, and the talented pianists who generously volunteered their time to participate in our study. Finally, this study could not have been completed without the kind help and advice of our faculty mentor, Professor Christopher Dobrian.

Works Cited

Bischoff, John, and Chris Brown. "Indigenous to the Net - Early Network Music Bands in the San Francisco Area." 10 June 2006 <http://crossfade.walkerart.org/brownbischoff/>.

Campbell, Colin. "Remote Piano Lessons, in Real Time". The New York Times, Thursday, March 11, 2004. 16 March 2006 <http://www.acadiau.ca/musicpath/main.htm>.

Chafe, Chris, Michael Gurevich, Grace Leslie, and Sean Tyan. "Effect of Time Delay on Ensemble Accuracy." Proceedings of the International Symposium on Musical Acoustics, March 31st to April 3rd 2004 (ISMA2004), Nara, Japan. 6 June 2006 <http://ccrma.stanford.edu/~grace/chafeISMA.pdf>.

Dobrian, Christopher. The Gassmann Electronic Music Series program. 10 June 2006 <http://music.arts.uci.edu/dobrian/gemseries01-02.htm>.

Lazzaro, John, and John Wawrzynek. "A Case for Network Musical Performance". The 11th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 2001) June 25-26, 2001, Port Jefferson, New York. 6 Jan 2006 <http://www.cs.berkeley.edu/~lazzaro/sa/pubs/pdf/nossdav01.pdf>.

Tanaka, Atau. "Speed of Sound". Machine Times, NAI/V2_Organisation, Rotterdam, 2000. 25 May 2006 <http://www.csl.sony.fr/downloads/papers/2000/AtauSpeedOfSound.pdf>.

Willey, Robert. "The Relationship Between Tempo and Delay and its Effect on Musical Performance." The Journal of the Acoustical Society of America 88.1 (1990).

Winckel, F. Music, Sound and Sensation. New York: Dover, 1967.

Yamaha Corporation of America. "Disklavier Technology Press Release - November 18, 2004." 18 May 2006 <http://www.prnewswire.com/mnr/yamaha/20680/docs/Mark_IV_Technology_Release-FINAL.doc>.