Title

Determining the Feasibility of Networked Musical Performances over WANs, LANs, and WLANs (Part 1: MIDI)

Authors

Dobrian, Christopher
Teeter, Matthew
Lindsey, Daniel

Publication Date

2011

Determining the Feasibility of Networked Musical Performances over WANs, LANs, and WLANs (Part 1: MIDI)

Matthew Teeter, Daniel Lindsey, Christopher Dobrian
University of California, Irvine
dobrian@uci.edu

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 musical performers to synchronize their playing and discern independent musical events. We used this information to determine the feasibility of conducting networked musical performances over local-area networks (LANs), wireless local-area networks (WLANs), and even wide-area networks (WANs), including performance of music that requires relatively tight synchrony of events. The experimental psychoacoustic and performance data we collected implies that successful rhythmically-synchronized 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 were usually met in performances involving MIDI transmission between two locations that are less than 400 miles apart (where network latency is below about 20 ms). In a future article we will detail the latency characteristics of networked performances involving transmission of audio streams, In this article we restrict our discussion to MIDI streams, which is far less demanding of network bandwidth. 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.

Introduction

Thanks to the increased prevalence of broadband Internet connections, software designers are exploring new applications of the Internet as a low-latency communication medium. One such application of interest in the arts is networked musical performances. In a networked musical performance over the Internet, performers play in two or more physically remote locations. The audience is also usually in two or more locations, possibly including locations where no performers are present, such as via an Internet media stream. This presents many new questions and challenges in performance practice, not the least of which is the technical problem of the delay caused by the transmission of information over the Internet.

Long-distance musical interaction 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, potentially enabling a single musical performance to reach a larger live audience. A large number of instruments, each in a different location, can replicate the performance of one performer. The instrument(s) at the remote location can receive information from the instrument that the performer is physically playing, and instantaneously replicate the same musical events locally. In addition, networked musical performance allows musicians to rehearse pieces without traveling to the same location, saving transportation costs and time. The feasibility of real-time networked music performance has also given rise to 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 has some obvious advantages for networked musical performance compared with streaming raw audio data, in that MIDI requires much less bandwidth and can allow a performance to be perfectly replicated on an instrument at another location. Transmitting audio data in real time to distant locations usually does not work as well for musical situations that require tight rhythmic synchronization. However, as we will discuss in part 2 of this study, low-latency networked audio is increasingly accessible.

A goal of most networked performances is to achieve a high degree of transparency—to minimize noticeable problems that occur due to network latencies. A highly transparent system allows the performers to play together as if they were in the same room. One 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). Previous research at Stanford University has shown that it is desirable to keep delay time as close as possible to 11.5 ms 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 the following scenario. Imagine two performers are trying to play a duet together that involves playing four separate notes per measure. Performer 1 starts off a networked performance by playing the first measure alone. Performer 2 tries to synchronize with Performer 1, and begins playing in the second measure at 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 2 uses the (delayed!) first measure played by Performer 1 as a reference for synchronization. Performer 1 listens to the timing of notes played by Performer 2, and slightly adjusts his/her own tempo in order to stay synchronized with Performer 2. Even if both performers had instantaneous reaction time and played their notes at the exact same time as they heard them played by the other, delays in the network will cause the other person's notes to arrive slightly after the original person played his/her own notes. 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 person is able to distinguish two musically independent events. As long as that is the case, the delay between when the first person plays notes and when the other person's notes arrive will 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 are smaller networks typically found in a home or building. Because they span relatively short distances, delays were expected to be less than 10 ms. 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, WANs span large geographic locations, such as states. As a result, the latency on this network is higher than that of LANs and WLANs. We were most interested in studying latency on WANs because the delays on these networks vary widely depending on distance. We also sought to determine the maximum distance two locations could be separated by 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 1970s, individuals in the League of Automatic Music Composers were investigating the idea 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 perform 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 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 instrumental 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 being 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 often did not require 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 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 simultaneously performed a duo piano concert from two cities, with Akagi playing at UC Irvine and Davis playing at UC San Diego (Dobrian). This networked performance was as tightly synchronized. Although no precise controlled experimental data was recorded, anecdotal evidence during testing in rehearsals showed that 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 UC Berkeley, Stanford University, and Caltech. While musicians may not always think about delays when performing, acoustical delays are present not only in a networked performance, but also in performance settings where all of the musicians are in the same room and close proximity. Players on a stage may also be separated by several meters, in which case they often use visual cues of the conductor to stay 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 the Linux OS has made impressive strides forward in improving the user experience, most musicians desiring to participate in a networked performance would own a computer running Windows or Mac OS. Thus, in order to see how well a networked musical performance could work on the Mac OSX operating system, we conducted our latency tests using Macintosh computers running OSX 10.4. Version 10.4 allows one to create a virtual MIDI device that is connected over the network. Apple claims that OSX's audio platform, Core Audio, was designed with the goal of keeping MIDI latency to a minimum. One of our objectives was to find out if these optimizations in Core Audio would allow a networked musical performance to take place using commonly available Macintosh computers, running Mac OSX 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 whether 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 connection. Wireless capabilities enable many exciting possibilities for computer music concerts, but we will not discuss the significance of these capabilities at this time.

Methods

When developing our experiments, there were three factors that we needed to test: (1) the average delay times over computer networks; (2) the average precision with which two performers could synchronize their playing; (3) 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 environment 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 were either a visual metronome or an audible metronome, or both. 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 the musicians cued one another with a nod of the head (which is common practice in chamber music). When using the head nod, pianists were specifically instructed not to follow a 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 (bpm). We tested a variety of pianists ranging from casual players to those majoring in Piano Performance at UC Irvine. Each of the test subjects had played the piano for at least 5 years and many had taken formal lessons during that time. To keep the performance material simple, we tested the pianists' abilities using only notes from 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 pianist synchronization ability to determine conditions which 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 perfectly identical to a conventional performance, which is the ultimate goal of this system. This is due to the fact that humans will not notice any latency in the system, since the delay time required to transmit the musical events from one location to the other is less than the time humans can even perceive such events. In order to test how well the human ear can distinguish separate musical events, we conducted a test where listeners closed their eyes and listened to two piano tones which began a few milliseconds apart. If the listener subjects believed the sounds to be distinct in their starting times, they would raise their hand. In contrast, if subjects believed the tones sounded simultaneously, then they did not raise their hand. One experimenter controlled the program which generated the tones, and observed the subjects' 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 was played.

Next, in order to observe the amount of delay inherent in networks across the United States, we used the standard network ping command to record the roundtrip times between UC Irvine and a variety of locations. We developed a Visual Basic program which utilized 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 9AM to 9PM. To allow for weekly variances, we ran the program every day for a full month. We tested network latencies by pinging the following areas: the same building at UC Irvine, across the UC Irvine 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 which 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 utilizing 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 Yamaha 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 Powerbook computers running Mac OS X 10.4.5. First, we tested latency when the computers were connected to the LAN with an Ethernet Cat 5 cable, and then we tested again when the computers were connected through the wireless LAN using the Apple Airport wireless Ethernet card. In this configuration, we pinged the other computer repeatedly, tried sending a three-byte MIDI message once every second using the MXJ net.udp.send/recv object (in Max/MSP), and also tested latency by sending a three-byte MIDI message once a second using the operating system's built in MIDI networking technology. The MIDI tests were conducted using 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 over six pianists, we recorded the ability of two pianists to synchronize their playing when using various cues in Table 1:

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

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.

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 (Tanaka 2000) and Winckel (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 network musical performance. We should, however, keep in mind that this perception test was an artificially controlled situation, where the subjects were concentrating on listening for two notes, instead of a typical musical setting where many more 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 UC Irvine 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:

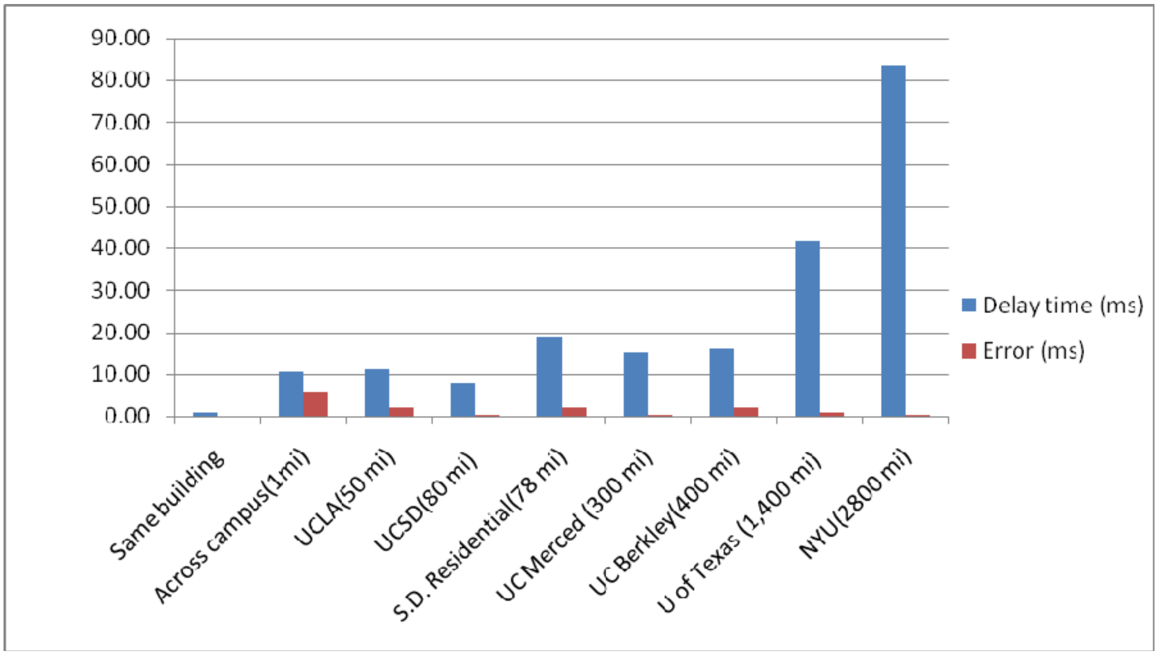

Figure 1. Average roundtrip time to various locations from UCI, determined using the ping command.

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, however, were occasionally slightly longer on weekday mornings (9:00 AM), but since most performances would likely occur later in the day, this does not appear to be a significant concern in most cases.

We see that the average network delay times for locations in California are less than 20 ms, which indicates that a networked musical performance is certainly possible when performing with someone who is 400 miles away. On the other hand, the average delay to out-of-state destinations, including 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 will 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, possibly to the extent of resulting in lack of synchronization.

Test Method	Wired Delay using LAN (ms)	Wireless Delay using WLAN (ms)
Ping	0.4	3 with occasionally spikes of 24
Max/MSP UDP objects	7	11
OSX MIDI Networking	4	7 with occasional spikes of 28

Table 2. Network latency determined using various methods on a LAN and WLAN.

It is also important to notice that the OSX 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 (Internet Control Message Protocol) Echo Request and Reply messages, which are small and require little processing. Typical network applications, however, 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 Mac OSX 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 also 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 OSX MIDI networking latency.

Finally, we come to the question of wired vs. wireless. Not surprisingly, we found that wireless communication is slightly slower than a wired medium, but also found that moving to a wireless medium only increased latency by a few milliseconds (see Table 2). Occasionally, however, there would be spikes in latency when using the wireless medium, likely caused by collisions of packets. Because of this, to ensure the most reliable connection a wired network should be used.

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 widespread and escalating adoption of broadband Internet connections an increasing number of amateur and professional musicians will utilize networked musical performances to take advantage of the many benefits such technology brings.

Our work for this article 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. As higher-speed networks continue to evolve, and more efficient video and audio codecs are developed, future musicians may utilize audio and video streams for even more immersive, networked musical performances.

Acknowledgements

We would like to thank everyone involved with this project, including our correspondent Ben Israel from Yamaha Corporation of America, Andrew Dahlin for assistance with wireless testing, reviewers of the UCI Undergraduate Research Opportunity Program (UROP) and the pianists who generously volunteered their time to participate in our study.

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 2 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. The results are shown in Tables A1 ~ A3.

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 scale upward with their right hand. One pianist would start at middle C, while the other would start at the C two octaves down. 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 3 times per pianist group, using a metronome speed 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 3 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 one, except the pianists relied upon both a visual cue and an audible cue. Thus, the metronome was playing at the same time as the pianists were observing the laptop conductor. This experiment was repeated 3 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 one, except the pianists were instructed not to follow a tempo. One pianist would use the “chamber music head nod” and press one note, and the other pianist would watch his or her partner and try to play their own note (2 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. Thus, we measured 57 starts in total per pianist group.

.

Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo(bpm):	80	Tempo (bpm):	80	Tempo:	80
Std. Dev (ms):	20.556	Std. Dev (ms):	23.176	Std. Dev (ms):	21.864
Mean(ms):	25.643	Mean(ms):	28.276	Mean (ms):	23.828
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	100	Tempo (bpm):	100	Tempo:	100
Std. Dev (ms):	20.833	Std. Dev (ms):	32.686	Std. Dev (ms):	16.463
Mean(ms):	20.69	Mean (ms):	33.759	Mean(ms):	20.055
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	120	Tempo (bpm):	120	Tempo:	120
Std. Dev (ms):	11.465	Std. Dev (ms):	20.107	Std. Dev (ms):	12.444
Mean(ms):	17.448	Mean (ms):	22.034	Mean (ms):	14.996
Cues:	head, Subject 1 leading	Cues:	head, Subject 2 leading		
Tempo:	free	Tempo:	free		
Std. Dev (ms):	22.822	Std. Dev (ms):	22.614	Overall	
Mean(ms):	28.545	Mean (ms):	33.179	Average (ms):	24.40481818

Table A1. Subject 1 and 2

Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo(bpm):	80	Tempo(bpm):	80	Tempo:	80
Std. Dev (ms):	21.968	Std. Dev (ms):	31.403	Std. Dev (ms):	25.187
Mean(ms):	26.474	Mean(ms):	37.729	Mean(ms):	33.328
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	100	Tempo(bpm):	100	Tempo:	100
Std. Dev (ms):	12.737	Std. Dev (ms):	19.496	Std. Dev (ms):	14.168
Mean(ms):	17.638	Mean(ms):	21.259	Mean(ms):	18.069
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	120	Tempo(bpm):	120	Tempo:	120
Std. Dev (ms):	13.293	Std. Dev (ms):	25.024	Std. Dev (ms):	11.757
Mean(ms):	17.086	Mean(ms):	21.69	Mean(ms):	14.362
Cues:	head, Subject 3 leading				
Tempo:	free				
Std. Dev (ms):	30.896			Overall	
Mean(ms):	39.724			Average(ms):	26.05775

Table A2. Subject 3 and 4

Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo(bpm):	80	Tempo(bpm):	80	Tempo:	80
Std. Dev (ms):	25.589	Std. Dev (ms):	49.053	Std. Dev (ms):	34.428
Mean(ms):	38.386	Mean(ms):	55.966	Mean(ms):	56.345
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	100	Tempo(bpm):	100	Tempo:	100
Std. Dev (ms):	33.231	Std. Dev (ms):	35.742	Std. Dev (ms):	26.305
Mean(ms):	33.07	Mean(ms):	45.931	Mean(ms):	36.241
Cues:	sonic metronome	Cues:	visual metronome	Cues:	sonic and visual
Tempo:	120	Tempo(bpm):	120	Tempo:	120
Std. Dev (ms):	17.987	Std. Dev (ms):	25.372	Std. Dev (ms):	27.954
Mean(ms):	24.632	Mean(ms):	44.155	Mean(ms):	41.724
Cues:	head, Subject 5 leading				
Tempo:	free				
Std. Dev (ms):	27.189			Overall	
Mean(ms):	41.089			Average(ms):	42.31725

Table A3. Subject 5and 6

References

Bischoff, J., and C. Brown 2006. *Indigenous to the Net - Early Network Music Bands in the San Francisco Area.* 10 June 2006, http://crossfade.walkerart.org/brownbischoff

Campbell, C. 2004. *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, C., M. Gurevich, G. Leslie, and S. Tyan 2006. *"Effect of Time Delay on Ensemble Accuracy."* P*roceedings of the International Symposium on Musical Acoustics*, March 31st to April 3rd 2004, Nara, Japan.

Dobrian, C. 2006. *The Gassmann Electronic Music Series Program*. 10 June 2006, http://music.arts.uci.edu/dobrian/gemseries01-02.htm

Lazzaro, J., and J. Wawrzynek. 2001. *"A Case for Network Musical Performance" The 11th International Workshop on Network and Operating Systems Support for Digital Audio and Video*. June 25-26, 2001, Port Jefferson, New York.

Tanaka, A. 2000. *Speed of Sound*. Machine Times, NAI/V2_Organisation, Rotterdam, 2000.

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

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

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