Interview with Chris Chafe, July 22, 2005

Jeffrey Treviño

Composer/performer Chris Chafe began experimenting with networked musical performance in 1998. In 1999, he received a grant from the National Science Foundation to initiate the SoundWIRE (Sound Waves on the Internet from Real-time Echoes) research group at Stanford's Center for Computer Research in Music and Acoustics (CCRMA). The group develops sonified evaluations of network Quality of Service and experiments in real-time musical performance via networks with high Quality of Service.

The SoundWIRE project has led to several notable collaborative real-time musical performances via high QoS networks. In 2000, the team's real-time networked reverb won the "Most Captivating and Best Tuned" research demo award at the SC2000 supercomputing conference in Dallas, Texas. Chafe played his celletto (an electric cello that he designed and built) in Dallas, sent the audio back to CCRMA's stairwell in Palo Alto, California, and then sent it back again to Dallas for a lush reverb created by a real space miles away. The team expanded their demonstration for SC01 (Denver, Colorado) to include over 320 channels of audio streamed in real-time between Denver and Palo Alto. (All the channels contained plucked string sounds in delay lines caused by the network latency.) 2002 saw the group's first successful multimedia collaboration, with low-latency video by McGill University's Jeremy Cooperstock. For his senior thesis, Stanford undergraduate and SoundWIRE contributor Daniel Walling distributed his dramatic improvisation ensemble between Los Angeles and Palo Alto; the resulting CyberSImps show can be seen online at http://ccrma.stanford.edu/groups/soundwire/cybersimps/. In the spring of 2004, musicians in Palo Alto, California; Missoula, Montana; and Victoria, British Columbia collaborated in real-time for a week to determine the form of an improvisational composition, which was performed at a meeting of CCRMA's industrial affiliates. Acclaimed documentary filmmaker Kris Samuelson joined Chafe and company for a summer 2004 collaboration that paired the improvisations of two duos of musicians, one in Palo Alto and the other in Stockholm, with flowing video of jellyfish and lunar landings. At the Audio Engineering Society's October 2004 convention in San Francisco, Chafe and his colleagues triangulated Mariachi Cardenal de Estanford into three recording studios around the Bay Area. The three studios' sound outputs were mixed back into a mariachi band in a San Francisco concert hall for the conference attendees.

Composer Jeffrey Treviño, Chafe's student while at Stanford, caught up with his former professor in Palo Alto on July 22, 2005. The two discussed Chafe's artistic interests, their relationship to the SoundWIRE project, and future directions for his work in the realm of networked performance.

JT: The last time we talked, your most recent networked performance project involved piping a mariachi band from three different locations around the Bay Area into a performance at the Audio Engineering Society's convention in San Francisco. Was that the most recent major event for SoundWIRE?

CC: Almost the last thing. Roberto Morales and I had a demo when I was in Europe about two months ago. I went to the art institute in Zurich, which is teamed up with the music conservatory, and we wanted to find out if you could play together as a duo between Zurich and here [Palo Alto, California]. Every time you set up for a networked performance somewhere, there's a whole bunch of new problems that you never knew about, you know. This is still kind of the very bleeding edge—this is the hemorrhaging edge, sometimes. The duo was an improvisation with Roberto Morales on flute and electronics and me on celletto. We've been doing a lot of weekly playing together, recording everything we do, so we have this down to where a lot of our reactions and musical thoughts just happen and we're having a good ole time; we're going to keep doing that every week. So it made sense in this case to have Roberto on the California end (since I was traveling in Europe), and we just made a date to try this thing out. There was enough wonderful support on the technical side to get the machines in place and connected up, but then we discovered that, beyond the basics, there was a crummy problem in one direction where packets were being dropped, and—it's interesting, maybe this is a word to the future, you know, for me, note this on a post-it—the thing to really avoid is promising the world to anybody in a show like this before you've actually tried it for real. And I had that misgiving, so I told them, "Don't do any publicity for this demo." You want to say, "Interested and forgiving people are allowed to attend," and it was a good thing I did it, because in this case, we really couldn't spend any time ferreting out the cause of the technical bottleneck. Unfortunately, the audience was in Zurich, and it was the to-Zurich direction that was dropping, whereas back to Roberto was great. At that point, we just yanked it down to one channel of 48

kHz and said, "Okay, we'll drop a certain number of those packets, but at least it'll play." And we played like crazy; it was really fun. The thing that I really cared about musically was going great. We could really get into our thing. Roberto does this fabulous Max-based processing of his flute and my cello and everything goes into the Osterizer, and so on—so sometimes you couldn't tell whether there were dropped packets or not…no, seriously, you could tell, and we weren't playing tight rhythmic-based music, either. We were flying around all over the place, and we played a good half-hour set that people enjoyed—at least they said they did—and we didn't pay them to say that.

JT: So what do you care about musically?

CC: In improvisation, it feels like it's working if you have this causal development of a piece going on, where one thing leads to another thing, you're building up this kind of forward history in the piece, and you're listening like crazy to each other, right? Roberto and I have that. There's a lot of development, there's a lot of common vocabulary, and when you're finally in it, you're playing with all those elements, and you know you are. So we had that going on. It wasn't just sound effects; it was this really strong kind of direction, which I think came through to the fifteen people or so who were there. They really picked up on it.

JT: But you play with him in person a lot.

CC: A lot, yeah.

JT: So, do you think you could get that kind of rapport with somebody whom you've never played with in person, somebody with whom you've only played over a network?

CC: Right, good question. So you only meet them for the first time in a tunnel or something, and then you start to play. It happens all the time when I'm improvising that I find other people who do it, too, out of some sort of weird bodily need or something, and it works, the first time. Well, I saw you and Max Mathews play together once—was that the first time you guys had played together?

JT: At the Cantor Arts Center?

CC: Yeah.

JT: Yeah, that was the first time we'd ever played together—

CC: And it smoked—I mean, it was great. You know, and everybody picked up on it. So it really has a lot to do with just a kind of willingness, and some chops.

JT: But Max and I were physically there.

CC: Yeah. Now the question is, could the same thing happen over a network? The answer is: the technology is successful when it doesn't matter whether you're physically in the same place or you're remote. So that's when we've gotten there. Now, the past history of our demos, and maybe the hype of all this, is that we have a successful technology, probably, because of the kind of high-definition audio that we're doing. And my definition of high-definition in this case is multi-channel, uncompressed, uncorrupted, low-latency audio. And all these things kind of add up. Adding compression adds latency, etc., so you just keep it down to the bare bones: it's here on the computer side, then it comes off the converters, it goes into packets, and it goes on the wire, and there's just nothing else going on. Just do that with lots of channels, and do it in a distance radius in which the delay doesn't impact the type of music you're playing. In Zurich, you're not going to play salsa (at least I'm not going to try), but in Seattle, maybe. So scaling the type of music according to the network's distance radius seems, right now, to be part of that definition of high-definition.

So if two improvisers meet in this tunnel, no video channel or anything like that, would it work? And this seems to be the question everybody's asking right now: how crucial is the visual connection in this world, too? And I can't say one thing or the other. Right now, I think it's desirable, but the musicians, once they're playing the music—it's like the cockpit window on the shuttle: you just want to know if the planet's really out there, you know? Take that window away, and there's a little less of that assurance. Where we've had ensembles that are less used to purely acoustical cueing (like inhaling breath to get a phrase started), where they really need to have a nod, then you have to cover for that. You put in an acoustical nod, otherwise known as an upbeat. These are all questions that we're feeling our way through in this new venue.

And it is a venue; I think it is, anyway. I define it that way, because it's really not like playing in a tunnel. A tunnel has very describable acoustics. I've been hiking through the underpass of a freeway this summer. I think it's a hundred feet long, and it's just a tube. Normally a river goes through it, but during the summer, hikers go through it, and it's really narrow. And if you're in that tunnel, there's a very peculiar acoustic to a conduit like that, and that's part of the sound of what you're doing. If you're in a room together, you can't avoid the sound of the room: it's a physically consistent ambience that has the players and their reflections all in the right place. How would you simulate that electronically? You'd have to build

a tunnel out of a computer music reverb that includes the regional distance of the delay as part of the acoustic. So we've been playing some of those games, and that leads to what's been going on since the mariachi demonstration for the AES convention last fall. We've been looking at distributed reverb, which allows you to form a tunnel surrogate in a computer music reverb algorithm. The transit time from one end to the other is actually incorporated in the algorithm, so it's a distributed signal-processing algorithm. It's a reverberator that has components on both ends and uses the network as part of the delay structure of the reverb. And, if you do it with multiple channels of audio, you can do it in a way that keeps the reverberation reflection angles, player positions, all of that stuff, consistent with the physics. So it would be like going into a tunnel. We aren't really there yet, and I think that may be another step towards making this venue become even more of a natural performance space. Finally, though, if I walk into a tunnel with a new player, and we just start playing, I think our performance is enhanced by the fact that our interactions in the same space are physically meaningful, even if the space is synthetic. I should probably do the experiment with a willing stranger under that freeway sometime, just to see, but I probably won't. It's more likely that we'll do it between here and, who knows, Los Angeles or some place. So we have to develop that. I think contributions from various quarters will be necessary to get the physically consistent ambience part of the technology going right.

JT: You said that you could have a shorter delay time to Seattle than you did to Zurich. In your physical model, then, would you basically change the distance of your tunnel based on your latency?

CC: Yes. If we wanted the hundred-foot tunnel in all cases, we could artificially lengthen the delay to Seattle so that it matches the bare-bones delay that you get to Zurich. That'd be one way to do the same performance to both directions. Unfortunately, we can't go the other way (get the Zurich dimension to be as tight rhythmically as the Seattle one) until we figure out how to beat the speed of light. It's quite a differential: just over ten milliseconds to Seattle, and just under a hundred to Zurich. That's what we're dealing with on these round planets, darn it.

JT: And you don't think that that's going to get any faster?

CC: It will, slightly. The basic speed law is at work here, but what's been nice is that—for reasons other than music, obviously—people have been gnawing away at the transit times on the Internet, so that these router delays are shrinking substantially. I think we had something on the order of twenty routers in the Zurich experiment. Each router's delay time is under a millisecond now, and that's really cool. You'll probably still have twenty routers in a lot of these cases, but as the router delay time decreases, those twenty milliseconds of latency will go away. That part gets good, and we have optical-based router switching and all these things coming around the corner—again, not because of us, but we can use it musically. And those twenty milliseconds are going to be significant for the extremes, for both the low latencies in Seattle and the larger times in Zurich. If Zurich comes down from a hundred milliseconds to eighty milliseconds, it may not ever get you into this range where you're really cooking on the rhythmic thing, but when some of the more local delay times change from fifteen milliseconds down to zero milliseconds, that puts it into the extremely close range. That's less than the five feet between us talking; that's five milliseconds. And that was what happened with Mariachi Cardenal de Stanford at the AES convention. That was the first time we had actually heard a distributed ensemble for which the radius, in terms of the acoustical delay between the ends of the ensemble, was much smaller than the room that we were listening in. It was like a little egg inside this bigger natural room, the concert hall that the audience was in. That was kind of inside out for me, because most of the time the delays are bigger than those of the listening space of the audience. So we're getting there.

JT: Going back to Max, I was talking with composer Justin Yang earlier about how we admired the musicality of people like George Lewis and Max Matthews, who build a system or an instrument and then stop development to take time and learn how to play it. If you as an artist were to stop at certain points throughout the entire development from 1998 on, as things have changed, how did "what you would do" change with the technology as it developed?

CC: I'm hoping to reach that stage where, for my less improvisational music, I start to actually structure stuff that lives only in this disconnected, remote world. That's part of the musical form, and it becomes one of the things that I'm designing with musically. So the technology needs to sit still enough for me to reach a point at which I can play with those designs. It's exactly put the way you said it. But I haven't had that opportunity yet. I haven't written specifically for this medium, let's say, whereas I guess I've got projects going on for other media that are sitting still, and I'm having that kind of enjoyment. I haven't reached the point at which you cease the technical introspections, the "make it work" part, and really get into the musical materials.

JT: When you say that you eventually want to make the remote and disconnected tangible in these projects, it sounds like a potential outgrowth of your *Ping* project and other collaborations with UC Berkeley's Greg Niemeyer. Can you talk about the connections between your networked improvisations, projects like *Ping*, and the idea of making tangible something that's normally not?

CC: There are four projects with Greg Niemeyer in which we're making tangible some sort of flux that's inherent in really commonplace stuff, but not apparent. In *Ping*, it was the behavior of network traffic. Everybody's got these wires running around them, and there are packets flowing all the time, but we're not really aware of all the funny rhythms and intricacies of traffic jams on the Internet. That was a way of making that tangible. Also like that was the *Oxygen Flute*, which monitored carbon dioxide levels in a plant growth chamber. You walked in, and you became sensible of your gas coming in and out of your mouth and exchanging with the leaves and bamboo inside the chamber. It makes tangible a very necessary exchange going on in our world: we breathe because plants breathe, and we wanted to make that kind of tangible as well, to bring it to the surface.

If you look at those projects, they're in the sonification world. They take data sets and dress them up musically so that, using your musical listening, you can pull out patterns from these data sets. This is interesting in and of itself, because you can use it to better appreciate the dynamics of some sort of system. But for me, in those pieces, it's much more about the music that comes out of them, because they're not all that different from the equations that I play with and jam with, which, in their first principles, really resemble the chaotic systems that are going on in an Internet traffic simulation. So the artistic perspective I have on sonification music is, again, really different from this kind of perspective that I've got on the telecommunications stuff right now. These are worlds that will probably couple together at some point—who knows, at the moment? I'd love to see that. It would be really fun to know what that means.

JT: So right now, it is really more of a telecommunications project.

CC: Pretty much. You can sort of look over the hill in your imagination and say what this might be, in terms of new musical avenues and forms and fun music to make. At the first go-round, benchmarking it against reality is an important thing to do, too. Say I'm going to split an ensemble into two rooms, have these folks either in different parts of the country or different parts of a building, and find out what happens to ensemble playing. What we're learning about are some underpinnings of the psychophysics of this weird beast, the ensemble. We often study players in isolation, but ensembles are really different beasts. They have these coupled behaviors that I don't know much about myself, and I don't think these have been teased out terribly well elsewhere. So as soon as we stick a wire in the middle and cause that separation, we've exposed some of the dynamics of those ensembles. But better to understand it a little bit before I go too far in tweaking this behavior to my own nefarious needs. That's going to happen, too, but it's a little bit like violin acoustics: a lot of time is spent trying to make a software violin from algorithms that sounds exactly like the real thing. Well, that's going to be hard to do, and we're not going to get there any time soon. But the research itself is very informative. The closer you get, you pull out answers, which then become modules for manipulation. You can create weird violins with tuba sprouts on them. We do that, obviously. It's that dual nature of research and creation. I go into the research to learn more about the goods that we're going to play some games with later, and I think this distance stuff is really still in that first stage, you know; we don't know enough to start playing. The technology is not done, by any stretch of the imagination. Dropping packets one way on the Zurich thing: that's broken, that's just absolutely not ready, and we have to find out how to cover for that. On the other hand, Daniel Walling's CyberSImps show, done a couple years ago, is a perfect example of a form that came out of separating the ensemble and crafting improv sketches that took advantage of the fact that they'd been split apart between Los Angeles and here. That was a major tour de force on the technical side, for him to get that going, but he closed the technology and then started working on the show. And that was great. That's really what you want to see happen more in the future, too. We will.

JT: So what's going to happen next with all this?

CC: I don't know exactly what's up, although it seems to involve a couple points in Europe. The folks in Zurich would like to do something else, if we get it figured out. There's a really neat possibility that the group in Belfast, Ireland at Queens University, will start to do some stuff with a new group starting there. This group is fun to describe, on two fronts, because it has a little bit of its technical motivation, but it is really more than that. On one front, there have been collaborating haptic instruments coming out of that group. They control synthesis with extremely simple stuff, like stirring your finger in a pan of little pebbles, and the music that you get from the system fits that motion. It's absolutely simple: a microphone pickup on the stones—not tracking every stone, or

anything like that. And the sound is really good. Now, the question is: what happens if you have one tray of stones on one side, one on the other, and you cross their synthesis and send it across the network? You have networked moving pebble music. That has haptics, sound synthesis, and some ensemble questions to it. I think those experiments are just ripe, ready to go. We don't need fancy stuff; we can use kind of a bare-bones signal transport to get that going. On the second front, both of our groups are starting to work with large-array multi-channel surround-sound-type things. At Belfast, they have a concert hall with a grid floor, a couple hundred seats, and they can surround them with an array of loudspeakers, including under the floor, and you can joystick sounds around a full-sphere projection. With the pebbles folks, and some of their new ideas, they're going to be joysticking in haptic ways. You're going to have a very tactile sense of this ambience as well as an instrument that you can feel; it touches back to you, it responds. So it's a full-picture thing. We just installed a room that has a hole in the floor like that, too, so the idea is to connect these kind of full-picture things over the network, doing music that can tolerate the hundred-millisecond delay one way, and getting the haptics involved. These are all pretty happening, I think.

JT: Since it was a telecommunications problem and not a musical one, why did you start this whole interest?

CC: It was a lot of fun for me right at the outset, actually, probably because I was a Ham Radio operator when I was kid.

JT: It all makes sense.

CC: Who's out there, CQ, CQ—is anybody listening? I think I told you this, but it was literally instigated by a woman, Elizabeth Cohen, who was working with the Audio Engineering Society as president during that time. Betsy had been part of a group who had been just commissioned to look at how Internet2 might serve the audio community. Lo and behold, after looking into that for a bit, they happened to note that the AES and Internet2 fall meetings were both in San Francisco at the same time and, coincidentally, something like a block apart. So Betsy said, "You can't miss this opportunity. We want to cross-connect engineers from both sides and talk about the problems." She called me up and said, "By the way, CCRMA should be there and do a demo." I said, "Oh great," you know, and <ponder> and, "What do I have for this?" There's this basic tenet that I was taught years ago: don't do a demo that makes music sound worse. So I bagged it. I said, "No, there's nothing here. I don't have anything to show." But it got the wheels turning, and this was at the right point in other work that I was doing; I was trying to figure out, "Okay, you have this odd idea of sending MIDI data from one place to another." Of course, lots of people had been working on this, but I hadn't really spent much time in the shower thinking about it before. I was doing a lot of music with feedback algorithms, particularly in MIDI, and I realized, "Okay, you could just get a couple Disklaviers, and you could have a feedback loop, and then they'd both blow a fuse; it'd be really fun." I immediately translated that into an audio picture, which is more of what Betsy was talking about. What if you had a feedback loop, but it was audio feedback? It would incur this network delay, and, with regular deliver and high signal quality, you'd have a delay line. That immediately grabbed me as a weird way to make a plucked string. You can use this delay line in a simple physical model, and if you can use it in a simple physical model, you can use it in anything. It's a delay line. I was also kind of going around and proselytizing at that point in time that delay is everything, and not just because of my administrator side of life. The idea is: all wave motion that we're used to, except for direct sound (which is almost completely missing in a lot of the things that we do), everything that makes a pitch, everything that has an echo, anything that has rhythmic systems—anything—is all based on time delay. I was trying to hammer this into some of my teaching. I also began thinking of the Internet as kind of a weird acoustical medium that has the possibility of reflections. All of a sudden, it became a full-fledged medium, just like air, water, or earth: you bang on it, and it reverberates. I took that interest to a networking group at NSF that I'd just learned about and said, "Hey, we can use these funny reverberating impulses to listen to Quality of Service on the Internet. A slightly changing delay time is going to create a pitch change, or a dropped packet is going to create some kind of crusty string sound." And the proposal floated, which surprised the heck out of me, and it actually turned out to be really out on a limb for that networking group. But it turned out really good, because we had a lot of students join this project. They did all this fabulous work to set up the streaming, which was really hard to do in 2000. It took a lot of special code and inventiveness to get low latency streaming, and we got it going. It was great. Next, we started using our system to split ensembles. We've always had this dual nature in the project. One side of it is experimenting in this odd acoustical medium called the Internet, and the other is fun with ensembles. It hasn't really changed. That's the telecom answer. The telecom approach to me is, "Eyes open, what are the qualities of this weird, acoustical medium?" It's certainly different from air; there's no doubt about it. As far as I can tell, it's the only medium that has a varying speed of sound, although air may be changing a little bit over certain time scales. The Internet

is jittery. You don't want it to be jittery. When we're trying to do these shows, we try to factor that out, but its nature is that it's jittery. The other weird thing about it is that it's asymmetrical. And that's more like a violin top plate, actually, because the speed of sound along the grain and across the grain is different; but end-to-end, bi-directionally, I don't know if there are any media that are asymmetrical like that besides the Internet. So it's this kind of funny beast that we're just playing games with right now. That's the short answer, told long.

Letter from the Editor

Over the past twenty years, Array has been a reflection of the interests and issues surrounding the International Computer Music Association. Periodically, the editors of Array have focused on the status of women in computer music. It has been seven years since the publication of Bonnie Miksch's letter and the responses to it from women working in the field of computer music. Continuing in this tradition, I have asked Gregory Taylor to write an open letter to the community, and I invite responses to his letter. Some people were concerned by my choice—they thought I should have invited a woman to write a statement about the female gender. I strongly believe that the lack of equality is not just a women's issue; it affects all members of the community. Gregory Taylor is an advocate for women in the field, programming many works by women on RTQE, a radio program of electronic, classical, ethnic, improvised and experimental music that has aired on Sunday evenings in Madison, Wisconsin since 1987. He has studied feminist theory and has a unique perspective on the computer music community because of the diversity of his background.

Recently, Harvard University President Lawrence Summers issued an apology for comments he made at an academic conference on women and science suggesting that "innate differences" between the sexes may account for fewer numbers of women in elite math and science academic positions. This created a firestorm in the media, and many articles were written containing possible explanations as to why the percentage of women earning doctorates in science and engineering is considerably higher than the percentage of women professors.

Computer music straddles two worlds: science and art. The number of women in academic positions in art and music is much higher than in science and engineering, but there is still a bias toward men in the arts. Of the 861 works that Christie's, Sotheby's and Phillips de Pury & Company offered over three days starting May 10 2004, a mere 13 percent were by female artists. Sixty-one pieces were assigned an estimated price of $1 million or more; of those, only 6 were by women. Of course, the fields of art and music are vastly different, and it is difficult to put a value on art. I mention this case merely to show a concrete example of difference in gender and the arts.

Computer music exists at the intersection of the two male-dominated fields of science and art, resulting in a subgroup that inherits