

Peter Traub

> interview



Sound Study at the Center of The End of Things

In "WoodEar" you embedded microphones and sensors in a living tree, driving an installation which is both physical and online. Furthermore it makes use of the resonance of the trunk and a number of external factors that influence the tree's life (light, external sound, wind, and temperature). Beyond the user's ability to mix those sounds remotely, can we say that the work allows for a tree's proper "telepresence" in our acoustic environment? And are the sounds also in a sense resonating through the networks? I'm not sure what would define the tree's 'proper' telepresence in our acoustic environment, or if there is such a thing. Right now, in the simplest terms, "WoodEar" is a stream of data coming from the tree: temperature, air pressure, light levels (visible and infra-red) from three sensors, and sway. It is then interpreted by the Processing-based application that the user downloads. This first version doesn't have live audio from contact microphones as I had hoped. I built the microphones and tested them, but kept experiencing serious ground-loop issues that I haven't solved yet. There was also the difficulty of streaming live audio into a Processing application, along with having that application generate sound. So the initial version of the application loads MP3s of recordings made underneath the tree, as well as synthesizing sound in realtime using the Beads library for Processing and Java. The synthesized sounds are driven by the live data feed from the tree. The sampled recordings are played randomly



"Solera" is a work about audio memory, which recorded and played in a public space stratifies and intertwines recordings day after day, as with an aging and blending process similar to liquor production (as referred to in the Spanish title). Is there any technical parallel between the process you implemented and how our acoustic memory is structured?

If there is a parallel between the two, I don't think it was really intentional - at least not in the initial installation. I think our acoustic memory functions within different timeframes and depends on the material we're listening to. If you're listening to spoken words, you're probably focusing more on the content and meaning of those words than remembering how they sounded. If you're listening to a concerto, you're remembering the immediacy of sounds, gestures, melodies, etc., but also remembering that within the larger structure of the work. With Solera, I really wanted to make present the memory of the space, and use this repetitive and additive process to show the ebb and flow of human activity within the space over a much larger time span than we're typically aware of. This worked especially well in a building at a university, where classes start and end at the same time every day. The human traffic overlapped at the same times each day - footsteps, voices, cell phones ringing, etc. I'd like to install the piece again one day, but in a different type of space - a stairwell perhaps. I can also imagine a version of the piece where it doesn't record and overlap in 24-hour cycles, but rather in one or two or half-hour cycles (or combinations of those). In its current incarnation, it is not a piece that's meant to be explicitly or closely listened to, but rather one that slowly permeates the acoustic consciousness of the people who work in that space from day to day. In that way, I consider it a type of aural architecture.

A similar physical structure (a room with four speakers in the corners and a microphone in the centre) was made in "Sound Study at the Center of The End of Things" where recordings were made in five-minute cycles (a thirty-second recording which generated a five-minute composition made through a

but are mapped to time-of-day - so recordings made at night are only played at night, and so on. As to the telepresence issue, I think that my network, or my representation of the tree and its environment, is somewhat like looking at very early digital photographs and comparing them to the real thing. The resolution and depth is so poor, that it is only a shallow approximation of the complexity of the real thing. "WoodEar's" representation of the tree is a bit like that. The tree and its environment, even when seemingly still and quiet, are infinitely more complex and rich than a piece like this could ever be. So the telepresence of the tree in our acoustic environment is really more the telepresence of this artist's representation of the tree, and my choices about what to show, how to show it, and how to map the data in a way that makes for a compelling experience.

The "always online" condition, especially in urban areas, is leading us to consider Internet as something at hand and so remoteness through networks seems more and more available. But when it comes to audio remoteness, as in both receiving or sending sounds to remote locations and therefore beyond being transmitted through IP protocols, do you think that networked audio implemented on mass will be able to physically shape objects or processes by converting transmitted sounds into acoustic waves? Does non-networked audio physically shape objects or processes now? Or does it reflect off

of and sound out physical objects and processes? I typically think of sound as a means of revealing objects or processes and affecting our perception of their physicality. In that sense, perhaps audio remotely shapes our perception of physical objects and processes (think of sound effects in a movie for example, and how they can add weight, or texture, or space, to a scene). In that way, at least to me, the meaning of 'always on' audio connecting spaces and people, and how it is altering our perceptions of 'local' and 'remote' is a more significant aesthetic issue than the physical effects of the shared audio.

In your opinion, how might networked 3D printing techniques that combine remoteness and the creation of precise objects affect sound art practices? Do you think that 3D printing a "record" at a distance would make sense at a time when we are seeing increased digitalization processes in music?

I'm not sure. I've honestly never thought about remote 3D printing in relation to sound art, in part because my work typically involves taking the physical - such as a tree, a hall, a public space, etc. - and transmuting certain unseen qualities of it into something acoustically perceptible and compelling. I can imagine some physical interpretation/manifestation of an event that is then mapped to the formation of a 3D object (in much the same way that we now map data into visual or sonic representations), though I'm not sure that 3D printing is an area I'll be working in any time soon.



Solera

spectral manipulation). How did people react to this work? And was it intended for visitors to become used to the cycle after a while?

It was interesting to hear how people reacted. The work was done as part of a group exhibit of mostly visual art called "The End of Things", which was set up in a vacant building that was scheduled to be demolished (but it actually still stands today - things didn't quite end!). There was this very small room with a big window that overlooked the gallery floor (it was formerly a furniture showroom), and I wanted to play with elements other than just sound - trying to create a unique and memorable space. I ended up bathing the room in red light with a single shaft of white light beaming down from the centre, perhaps making the space like the heart of the building. I also wanted the generative composition to be loud - so it gradually builds until it is very loud (not quite painful, but close), and then cuts off. Some people just hung out in the room and had conversations, even while the piece was going. I think most of the visitors hadn't really been exposed to sound installations before - they were mostly college students or people from the visual arts community - so their reactions were pretty varied. The piece ran for a week, but I wasn't there most of the time so I only got to see and hear a few people experience it.

Panta Rhei (developed with Lanier Sammons) is a more participatory work in which an audio/visual combination of LEDs, photo-resistors and sounds create an emergent

system where human interaction is not "triggering" a system reaction, but rather "disturbing" the system in its own functioning. Was the work aimed to challenge standard conventions about interaction? And did the spectators have opportunity to understand their role?

I think the work was primarily intended to create a system that showed emergent properties through implementing some simple rules, loosely based on cellular automata. At the time, Lanier and I were still in the Ph.D. Composition and Computer Technologies program at the University of Virginia, and we were participating in a joint Music and Architecture class with a number of architecture undergraduates. That people could interrupt the system and affect its behavior was almost secondary, though it did become an integral part of the piece. I think the piece did challenge interaction conventions as the interaction was not very one-to-one: placing one's hand between the lights and the photosensors interrupted the 'flow', but not necessarily in a predictable way. There were a lot of analog variables in Panta Rhei - for example, slight variations in the angles of the lights or their intensity on their corresponding analog photo-resistors, or variations in the resistance of each photo resistor to a given light level, and so forth. The system already had a lot of analog variation in it, and so introducing a hand or other object to interrupt the light would create even further unpredictable variation in the light and sound behavior.


You have also investigated social networks in your artworks. In ItSpace you created MySpace profiles made out of household objects with recordings of the sound they produce. As fellow abstract musicians, they were all friends (allowing visitors to discover them one by one) and visitors were encouraged to add new profiles with their own pictures and sounds of objects and grow this sub-network. Were many new pages created? Did this produce any new outcomes?

The funny (and perhaps slightly disappointing) thing, is that ItSpace has had more success as a physical installation - or at least it has

reached a wider audience. The original piece that was done within MySpace back in 2007 (a little before Facebook completely overtook it as the social network of choice) did actually receive other contributions, and at least one composer created several 'friends' for the nine core pieces. At present there are somewhere around 40 ItSpace objects still hanging around on MySpace. It did receive some good coverage - and particularly a story about it that was broadcast on National Public Radio. I think the barrier to participation in the piece was a little high. Contributors needed to be fairly technically experienced, i.e., they needed to know how to create a new MySpace page, record live sounds, edit and mix those sounds into a piece, and upload that sound to the page. After contributions died down, I wanted to make a version of ItSpace that could bring the pieces out to performances, galleries, etc. I created a physical installation that embodied each object in the form of a photo board with a push button embedded in it. Pressing the push button generates a piece made from sounds of that object. The pieces are generated algorithmically, so they sound quite different from the original ItSpace pieces. This version of the piece was shown at SIGGRAPH 2011 in Vancouver, BC and then at Pixilerations 2011 in Providence, Rhode Island. It ends up being fairly different from the social network version (the social network aspect of the piece is essentially absent), but it also allows people to physically connect to the piece through an almost comically limited interface. This version is more about taking objects from my house and re-mapping their relationships through multi-media.

Symbolizing the constraints we have in transmitting audio over networks, ground loops: for solo percussion and internet sent percussion sounds to three different servers in U.S. (San Diego, CA, Hanover, NH, Charlottesville, VA), obtaining distortions through compression and feedback in the loops. Do you think that compression, which is still the price we pay for transmitting sounds over networks, has affected our listening abilities?

Yes, absolutely. Not just listening but seeing as



well. We've become so accustomed to seeing compressed images, streaming video, etc. that we are hardly bothered by the artifacts anymore. When it comes to music, I think many people have come to ignore artifacts that result from heavy compression, or the flatness that heavily compressed music seems to exhibit. On the other hand, I wonder to what degree our brains make up for the shortcomings of compression, as they are pretty good at 'filling in the blanks' (which is the basis behind psychoacoustic compression anyway - allowing our brains to make up the information that the audio lacks). It's interesting that in an age of digital precision, so many people still value analog recordings on vinyl. Aside from the embodiment of the sound wave on an engraved disk, there are artifacts such as crackles, warps, scratches, etc. that create such sense of nostalgia for so many listeners that we've brought these artifacts back into our music through samples and even digital effects. So I wonder if - at some point in the future - people who grew up in the late 90s to early 00s will be nostalgic for the limited spectra of heavily compressed MP3s. Will they use compression-simulator effects to make their percussion or vocals sound early-21st-century? (Musicians who make low-bit music do this already - playing on nostalgia for computer sounds from the 80s and early 90s).

After bits & pieces (one of your first works which searched and collaged various strategic audio files found on the net), you collaborated

on two other projects in the early 2000s (sibling revelry and NetSong) which aimed to retrieve (or search for by keywords) sounds on the net and play them in various ways. Do you think that now, after a decade, the immense quantity of sounds and music available on the networks still makes sense as material for an artwork? And if you thought about re-doing this project today, what kind of sounds would you use for collages?

That's a good question and I'm not really sure. As we have such an abundance of print material, does it make sense as a material for painters or sculptors or other visual artists who make physical collages? I think that they would answer yes, so I wonder how it is different for sound on the Internet. When I made bits & pieces, audio was still in its infancy on the web compared to today. We also had a large choice of search engines, such that in its premiere, 'bits' would randomly query AltaVista, Yahoo, Lycos, Excite, etc. (do any of those names make you nostalgic?). The results that came back - especially if considered over a large span of time - were probably a decent weighted representation of the types of audio files on the web. Today, I imagine 'bits' would find mostly MP3s of pop music, or excerpts of MP3s. With the abundance that's out there now, it would probably be more interesting to focus on a particular type of sound or music or recording. 'bits' was intentionally very vague with its searches, using terms like 'music', 'sound', 'aiff', 'wav', and so on. It would be interesting to do a collaging piece now just based on material on a site like Freesound, or Soundcloud, or other sound-rich sites that are driven by communities and whose content would reflect the interests and concerns of that community.

We're still missing an effective way to search the Internet for and with sounds (as, instead, it's possible with pictures), while on the other side, voice has been successfully integrated in the interface. Do you think this is a paradox or just a specific industrial policy?

I think the technology is probably there, but perhaps no one has really found the 'killer app' for it - the closest thing I can think of is a

service like Shazam, and I'm guessing that 90+ percent of their traffic is people trying to identify a pop song they're listening to. For it to become something as mainstream as Google image search, it needs to be able to make someone somewhere rich. This is just the nature of the world we live in right now. A company isn't going to put money into researching and refining a tool like that unless it will make them a lot of money. They have, to some degree, already found the aspect of sound searching that most people find valuable: "I like this song, show me other songs that are similar" or "What are my friends who have similar taste to me listening to?" This is cheap, doesn't involve processing individual sound files, and builds from the collective perception of the masses. I think a sound-searching tool like you describe is probably most valuable to a limited set of users - those who work with sound for music, video, games, etc. In that sense, it would be very cool to be able to upload a sound to a site like Freesound and have it show me a list of other sounds that sound similar (through some spectral analysis and comparison I'd assume). Such a tool could be extremely valuable for composers or producers working with large libraries of sounds (especially if poorly categorized). Do you know of one out there?

Peter Traub is a composer and artist currently living in Charlottesville, VA (U.S.A.). He completed his Ph.D. in 2010 in the Composition and Computer Technologies program at the University of Virginia. He has composed numerous works of electronic music and several internet-based and physical sound installations.