

Revisiting the Technical Achievements of *Listening Post* Ten Years On

[Wes Modes](#)

University of California Santa Cruz

Abstract

At the turn of the 21st century, a collaboration between a Bell Labs statistician and an experimental sound artist resulted in a mesmerizing visualization of real-time data that resonates still, in equal parts for its conceptual strength, cohesive aesthetics, and technical brilliance. Mark Hansen and Ben Rubin debuted *Listening Post* at the Brooklyn Academy of Music in 2001. The work is a dynamic collage of live conversation drawn from online sources in a room-sized installation featuring 231 small text displays mounted in a grid overlaid with dynamically-generated sound, music, and voice. Much has been written about the beauty of the piece, the rhythms, visuals, and sound, about the conceptual intersections of wired online existence, privacy concerns, and surveillance, and about adding big data to the mix of art and technology that is new media. However, with the exception of some early writings by the artists, little has been written about the still-impressive technical achievement. That the piece is still fresh is proved by a permanent installation of a similar work by Rubin and Hansen in The New York Times Building in New York City just two years ago. Additionally, *Listening Post* is a living piece that Hansen and Rubin have maintained to keep up with the times. In this paper, I will examine the technical details of *Listening Post*, as much as possible in Mark Hansen and Ben Rubin's own words culled from papers on the piece, conference proceedings, workshops, and interviews.

Beyond its far-reaching influence in new media art, and despite the ten years of rapid technological change that have passed since it was debuted, Mark Hansen and Ben Rubin's 2002 installation *Listening Post* remains an outstanding engineering and aesthetic achievement in the realms of data visualization, internet-connected sculpture, and real-time data mining technology.

Keywords: Listening Post, Mark Hansen, Ben Rubin, New Media Art, Data Visualization, Data Sonification

Rubin put it this way: "Dissociating the communication from its conventional on-screen presence, *Listening Post* is a visual and sonic response to the content, magnitude, and immediacy of virtual communication" (Rubin, 2010). Rubin talks of his hopes and intention. "There are an untold number of souls out there just dying to connect, and we want to convey that yearning. I hope people come away from this feeling the scale and immensity of human communication" (Mirapaul, 2001).

Hannah Redler, Head of Arts Projects at The Science Museum in South Kensington wrote, "As a work of art and a piece of technological ingenuity in its own right, *Listening Post* is hard to categorize. An extraordinary investigation into the meaning and malleability of statistics, it combines a Minimal art aesthetic with the elements of chance and randomness common to experimental art from the early 20th century to the present day. But its engagement with media technologies and sophisticated data-analysis techniques differentiates it from traditional visual art. It relies not on the found objects of Modern Art but on found data and extracted thoughts - the very unstill lives of a hundred thousand active minds" (Redler, 2007).

Mark Hansen examined the broader implications of their work early on in a 2000 paper to the IEEE: "The advent of these enormous repositories of digital information presents us with an interesting challenge. How can we represent and interpret such complex, abstract, and socially important data?" Hansen demonstrates that they were not unaware of the political dimension of their work:

As we become more adept at observing, measuring, and recording both our physical and virtual movements, we can expect to see a proliferation of large, complex data sets. These new digital records of human activity often force us to consider difficult societal questions. From the current debate over privacy on the Internet to ethical concerns over the use of genetic information, we recognize that the now simple act of compiling data has serious political implications. To further complicate matters, the scale and structure of these data make them difficult to comprehend in the large, relegating activities like data analysis to a select few (Hansen & Rubin, 2000).

In the last decade, *Listening Post* has been deeply influential in new media in general and data visualization/sonification in particular. Hansen and Rubin's two jointly-written academic papers for the International Conference on Auditory Display (2001, 2002) have been cited in scores of academic and artistic follow-on studies and projects. The piece is hailed as a new media masterpiece. Kenneth Baker, SFGate arts reviewer, wrote "Perhaps no other artwork has evoked human collectivity, and the communicative needs that enrich it, as *Listening Post* does. It is that rarest surprise in contemporary art: a recognizable instant classic. *Listening Post* positions itself so forcefully as a work of the present cultural moment that it can make things on view around it seem sadly peripheral" (Baker, 2013).

Hannah Redler, curator at London's The Science Museum, acknowledged it as a "masterpiece of electronic art; it references issues and themes central to software and interactive art, while subverting notions of interactivity. By anonymously drawing from active public places on the internet for its raw material, using thousands of expressions from thousands of unwitting online contributors, it repositions the point of interaction to the point of source rather than the point of encounter. It is itself as much a voyeur as the gallery audiences to whom it performs its findings" (Redler, 2007).

The piece has evolved over time, slowly becoming more focused, algorithmically simple, and using more sophisticated hardware. Any attempt to describe the precise technical characteristics of *Listening Post* will be like shooting at a moving target, but try we will. Hansen said, "We developed a number of experiments on our way to *Listening Post*. It started as a sound piece with a synthesized piano score together with a computer voice. Over time, and with newer versions of the work, groups of screens were added, first in a small grid and then in a larger one" (Hansen, 2014). In broad strokes, it can be said that *Listening Post* has done a lot with simple tools and improved over time.

While I will be examining technical details, I want in no way to neglect to acknowledge the design

achievement of the artists of *Listening Post*. As Mark Hansen put it recently, "With a few exceptions, data has no natural 'look,' no natural 'visualization,' and choices have to be made about how it should be displayed. Each choice can reveal certain kinds of patterns in the data while hiding others. While these decisions are often made on technical grounds, they are also questions of design" (Hansen, 2013). While data has no natural look, the artists of *Listening Post* make audible the voices crying out into the darkness of the internet for connection.

Figure 2. [Listening Post: "I am"](#) from [Ben Rubin](#) on [Vimeo](#).

Data Sources

Mark Hansen, a statistician at the Bells Laboratories division of Lucent Technologies, met experimental sound artist Ben Rubin, who had collaborated with Steve Reich and Laurie Anderson, at an art and technology event sponsored by Lucent and the Brooklyn Academy of Music. They began working on various projects to sonically represent activity on the Internet. In their experiments, they quickly moved from exploring website navigational data to actual web content. According to Mark Hansen, "We decided that perhaps navigation statistics were less interesting than the substance of their online transactions, the content being exchanged. We also agreed that the act of Web browsing wasn't very 'expressive' in the sense that our only glimpse of the users came from patterns of clicks, lengths of visits, and the circle of pages they requested. These considerations led us to online forums like chat and bulletin boards" (Reas & Fry, 2007).

"We plan on establishing listening posts -- points where people can tap in and listen to the pulse of global information systems," Hansen and Rubin wrote in 2000. The goal, said the artists, "is to make interpretable the thousands of streams of dynamic information being generated on the Web. In so doing, we attempt to characterize a global dialogue, integrating political debates, discussions of current events, and casual exchanges between members of virtual communities" (Hansen & Rubin, 2001). Originally, these points were web-based bulletin boards and Internet chat rooms, but in keeping up with technology and shifts in internet communication patterns, the artists have recently made some evolutionary changes to *Listening Post*.

Data Gathering

Identifying data sources was only a fraction of the challenge in the ever-changing landscape of the Internet. "We first had to create specialized software agents that would both discover new chat rooms

and message boards, as well as harvest the content posted to these sites" (Hansen & Rubin, 2001). Once sources had been discovered and identified, the system had to retrieve the content. "For these situations, we constructed a content agent in Perl, as this language provides us the most convenient platform for managing access details (like cookies). The public chat rooms on sites like chat.yahoo.com can be monitored in this way" (Hansen & Rubin, 2001).

Internet Relay Chat (IRC) is a non-web-based protocol for live interactive text messaging. In the early pre-web days of the internet, it was common way for those had access to net-connected systems to chat across distances in real-time. In 2001, the artists describe the system that mines IRC for content: "For IRC we built a configurable Java client that polls a particular server for active channels. Web sites like www.cnn.com (a popular news portal) and www.financialchat.com (a financial community hosting chat services for day traders) offer several IRC rooms, some of which are tightly moderated" (Hansen & Rubin, 2001).

For more than a decade, *Listening Post* has culled content from IRC chat rooms and online message boards. However, newer technologies eclipse older ones, and to keep the piece vital, Hansen and Rubin have had to reconsider their data sources. Mark Hansen said very recently, "Our early writing on the piece talked about a 'global conversation' and now that seems to be happening more broadly on Twitter and other social media, not IRC" (Hansen, 2014). Later in this article, I'll talk about the future of *Listening Post* and its current adaptations. Looking back on the challenges, Mark Hansen said, "In retrospect, it was pretty easy to create a data stream from these forums, sampling posts from various places around the Web. Doing something with it, representing it in some way, responding to its natural rhythms or cycles, proved to be much harder" (Reas & Fry, 2007).

Data Processing

Given the wide gulf between the expectations of scientific statistical data processing and arts-based data visualization, the amount of care and rigor that Hansen and Rubin put into statistical methods of data gathering and processing for *Listening Post* is a generous surprise. In a 2001 paper to the ICAD, Hansen and Rubin describe the first stage of analysis of their data stream. "In addition to collecting content, each monitoring agent also summarizes the chat stream, identifying basic topics and updating statistics about the characteristics of the discussion: What percentage of visitors are contributing? How often to they contribute and at what length? Is the room 'on topic,' or are many visitors posting comments on very different subjects?"




Figure 3. *The Listening Post*. Photo credit: Graham Peet for the Science Museum

While these metadata were available and gathered in early versions of the piece, they are not used in the current version of *Listening Post*. Mark Hansen said, "We never really incorporated them into a scene except to display a list of disassociated user names... a bit like credits" (Hansen, 2014). Interestingly, this hints at the rich potential of art that takes advantage not only of large public data sets,

but of metadata one can gather *about* those data. This high potential of metadata hasn't been lost on new media artists, corporations, or government security agencies.

Though simple tools were used, the algorithmic complexity of *Listening Post's* data processing has at times been quite sophisticated. For instance, early on, a statistical method called generalized sequence mining, or GSP, was used to choose the topics that *Listening Post* highlights. GSP is a 1996 algorithm developed by two IBM researchers used to discover sequences in a possibly noisy data stream (Srikant & Agrawal, 1996). It can be used to iteratively extend existing sequences to generate new sequences (Pitman & Zanker, 2011). "Topics are derived from the chat stream using a variant of *generalized sequence mining* that incorporates tags for the different parts of speech" (Hansen & Rubin, 2001).

Early on, *Listening Post* handled loosely-moderated discussion by clustering topics using an algorithm based on Occam's razor. "Because a content stream can in fact support a number of simultaneous discussions (the threads of a bulletin board, say), we employ a soft-clustering technique" (Hansen & Rubin, 2001). The artists used a principle called Minimum Description Length (MDL) to cluster possible topics. MDL says the best model provides the shortest description of a given set of data while still capturing the important features evident in the data (Hansen & Yu, 2001). After topics are chosen and grouped, *Listening Post* pulls content from the stream to represent those topics. "A stochastic framework was developed to sample representative sentences posted to the chat or bulletin board" (Hansen & Rubin, 2001).

The statistical methods used by *Listening Post* render it immune to the notoriously creative spelling and grammar of many internet posts. "Unlike most applications of statistical natural language processing, our content monitors update their summaries each time new material is posted and downweight older contributions. Because our sonification renders these sources in real-time, small mistakes have little effect on the power of the overall display to convey the ideas being discussed" (Hansen & Rubin, 2001). However, filters are used to remove cryptic symbolic content, such as emoji, from the output data stream.

Given the complexity of *Listening Post's* data processing, it is sometimes surprising to find that much of the heavy lifting was done with very basic Unix scripting languages. Hansen and Rubin write, "To make sense of our stream of text data, we relied on Perl for text parsing and feature extraction, some flavor of UNIX shell for process control, and the R programming environment" -- a free software programming language for statistical computing -- "for data analysis, modeling and statistical graphics" (Reas & Fry, 2007).

Early papers on *Listening Post* describe quite complex statistical processing, but the piece has evolved over time. With recent changes in *Listening Post*, statistical analysis is vastly simplified, obviating the need for calculating topics in a chat room or forum. "The scenes do make use of specific kinds of filters as well as text clustering. But [some] processing isn't needed because we are not doing that much per-room summarization," said Hansen. "Our initial experiments with the audio-only version of *Listening Post* and the early grid even had some more advanced processing. As the piece matured, some things simplified. We dropped the technical work that really didn't add to the story we were telling" (Hansen, 2014). The textual streams are gathered, processed, and then used to influence the sonic and visual output of *Listening Post*. "In addition to textual data streams, the statistics engine is also responsible for communicating the various ingredients for the display to our sonification engine, Max/MSP" (Hansen & Rubin, 2001).

Sonification

It is no coincidence that *Listening Post* is a collaboration between a statistician and an experimental sound artist. Sound was not an afterthought to accompany the visual display, but the center around which *Listening Post* was designed. In fact, earlier versions of the project lacked the striking visual

components we associate with the piece today (Hansen & Rubin, 2002). "In the end, the visual component of *Listening Post* acts as a kind of ventriloquist's dummy, which is animated by our sonic design" (Hansen & Rubin, 2002). It is fitting that the first two detailed reports from *Listening Post* by Hansen and Rubin appear in the Proceedings of the International Conference on Auditory Display. *Listening Post*, with its data-sonification roots, offers a multi-layer audio experience, consisting of mechanical noises, sampled sounds, synthesized voices, and a dynamic musical composition.



Figure 4. Early version of *Listening Post* in December, 2001, at the Brooklyn Academy of Music.
Source: Ben Rubin, EAR Studio

Text-to-Speech

The artists realized the need for computer-generated voice early on. "Our starting point is text. Albeit diverse in style and dynamic in character, the text (or transcript) of these data sources carries their meaning. Therefore, any auditory display consisting only of generated tones would not be able to adequately represent the data without a very complex codebook. The design of our sonification then depends heavily on text-to-speech (TTS)" (Hansen & Rubin, 2001). With all the audio output of *Listening Post*, much of it multi-layered and simultaneous, how not to create a confusing mishmash of sound? "Our goal is to create a sonification that is both communicative and listenable. Here we face the additional challenge of incorporating verbal content. With TTS annotations, it becomes more difficult to intelligibly convey more than one layer of information through the audio channel. Our design incorporates spatialization, pitch and timbral differentiation, and rhythm to achieve clarity in the presentation of the hierarchically structured data coming from the statistics engine" (Hansen & Rubin, 2001).

While it might have been sufficient to use the content of Internet forums and chatrooms as a springboard for data-inspired patterns, a source of nearly random data to influence an ever-changing sonic score for the piece, Hansen and Rubin used statistical tools to generate sonification that carefully reflected the character of the discussion of their internet sources. "The incorporation of spoken components in the sound design poses new challenges, both practical and aesthetic. For example, simply voicing every word taking place in a single chat room can produce too much text to be intelligible when played in real-time and can quickly exhaust the listener. Instead, we build a hierarchical representation of the text streams that relies on statistical processing for content organization and summarization prior to display" (Hansen & Rubin, 2001).

Text is organized into topics, the words and phrases appearing most frequently, mined from the multiple chat streams. The time spent on each topic reflects the relative attention it is receiving online. "The auditory display cycles through topic clusters, spending relatively more time on subjects being actively discussed by the largest numbers of people" (Hansen & Rubin, 2001). While each topic has a unique pitch to differential it as content is read by the text-to-speech subsystem (Hansen & Rubin, 2001).

The statistical tools run in real-time influence both the visualization and sonification of the data stream. The results are used as the material from which each movement is constructed. For each cluster of topics, the statistics engine sends three streams of information to the sonification engine: A continuously updated list of topics; A selection of sample sentences pulled from the content stream, identified by the statistics engine as typical or representative, in which these topics appear; A value that represents the current level of entropy in the source data (Hansen & Rubin, 2001). To render the data streams comprehensibly, Hansen and Rubin put the spatial characteristics of the speaker layout to full use. For example, in one scene, "the topics are spoken by the TTS system at regular intervals in a pitched monotone, and are panned alternately hard left and hard right in the stereo field, creating a sort of rhythmic 'call and response.' The sample sentences are panned center, and rendered with limited inflection (as opposed to the pitched monotone of the topics)" (Hansen & Rubin, 2001).

Before text is output via TTS, it is filtered to remove emoji and other unpronounceable content. Mark Hansen said, "Our TTS sounds awful on certain patterns of symbols," whether being sources from IRC or newer data sources. *Listening Post* uses a non-commercial TTS system developed by the Bell Laboratories TTS group (Hansen, 2014). The TTS system receives its cues from other components of the system, the content of the spoken text from the display system and the pitch and volume from the sound generation system. The artists worked with a Bell Labs engineer to write a custom C++ wrapper to handle the network communications between the TTS system and the other system components.

Sampled Sounds and Dynamically-Generated Score

Throughout most of the movements, sampled sounds signal significant events in each scene. For instance, in one scene, as a system agent "searches" through posts for a match to a chosen topic, a sampled bell or gong is played when a post is found that matches the topic. Some of the movements of *Listening Post* have a dynamically-generated score, "an accompaniment to the vocal foreground, enhancing the compositional balance and overall musicality of the sound design" (Hansen & Rubin, 2001). The output of the statistics engine sends messages to the sound generation system. Early on, statistic gathered from the data stream had a greater effect on the musical score. In 2001, Hansen and Rubin described the system creating a sophisticated soundscape: "The entropy vector controls an algorithmic piano score. When entropy is minimal and the discussion in the chat room or bulletin board is very focused on one subject, chords are played rhythmically in time with the rhythmic recitation of the topics. As entropy increases and the conversations diverge, a Gaussian distribution is used to expand the number, range and dynamics of notes that fall between the chords. With this audio component, one can easily differentiate a well-moderated content source from a more free- form, public chat without distracting from the TTS annotations." The audio portion of *Listening Post* including sampled sounds and the dynamic musical composition is orchestrated by Max/MSP running on an Apple Mac. The scene management system sent messages to Max via the Open Sound Control (OSC) protocol (Reas & Fry, 2007).

Mechanical Point Sources

Each screen contains a mechanical point source experienced as a click of variable intensity. "Mounted to the back of each display unit is a relay which can be actuated under software control. By varying the duration of the relay actuation pulse, we can control the loudness of a mechanical 'click' from each display," in effect creating a large grid of controllable point-sources for these clicking sounds. "Each

display can then make a click that was loud enough to be heard over the other sounds in the room, drawing a visitor's attention to that spot on the array" (Hansen & Rubin, 2002). Singly, the click immediately draws your attention to the currently active display. In multiples, whether a subtle rustle or a loud cascade, the mechanical click is a critical part of the sonic makeup of the piece. This simple click is surprisingly effective in embodying the output of the visual displays, which might otherwise feel insubstantial.

Other Aural Considerations

Much thought went into the display space for *Listening Post*. Early on in 2002, the artists said, "the aural impact of our display depends on having a controlled acoustic space." From the beginning, Hansen and Rubin carefully mapped out the acoustic needs of the space. Installations include carpeted floors, acoustic treatment on the walls, and a baffled entrance to reduce sound from outside the installation. "These acoustical treatments yield a relatively quiet and non-reverberant room, enhancing intelligibility. 8 of our 10 speakers are mounted behind aluminum panels, and the remaining 4 (as well as the subwoofer) hang high up and out of view" (Hansen & Rubin, 2002). Audio spatial separation is itself a separate subsystem of the piece. "During each scene, the voices and other musical elements move around the room. While Max handles much of this motion, a Yamaha Digital Mixing Engine is also used, which in turn requires a separate program for each of the scenes" (Reas & Fry, 2007).




Figure 5. Early version of *Listening Post* in December, 2001, at the Brooklyn Academy of Music.
Source: Ben Rubin, EAR Studio

Visual Display

The visual centerpiece of *Listening Post* is a large array of 231 small text displays, each measuring 2" by 6", and able to hold four lines of 20 characters. The displays are arranged in 11 rows and 21 columns and the individual displays are set 6" apart vertically and horizontally (Hansen & Rubin, 2002)(Reas & Fry, 2007). The striking visual components of *Listening Post* were not part of the original aurally-focused design. During initial experiments with the sonic piece, "We observed, not surprisingly, that when we added a projected visual display of the four text streams, the audience was much better

able to attend and comprehend the spoken text. When considering scaling the Kitchen experiment from 50 rooms to 5,000, this final observation led us to design a visual element for the *Listening Post* display system presented here" (Hansen & Rubin, 2002). The artists are referring to an early live performance at the Kitchen in NYC using sound, TTS and a single wall-sized projection with several lines of animated text.

In Hansen and Rubin's early sonic experiments, "we found ourselves constantly referring to a text display I hacked together to monitor the data collection. While we were led to this simple visual device to help make up for deficiencies in the TTS program ('Did someone really type that?'), it soon became an important creative component" (Reas & Fry, 2007). As mentioned above, the artists started with a modest projection with four lines of text at a live performance at the Kitchen in 2000. Soon they created a 10 by 11 suspended grid of flat displays, vacuum fluorescent displays (VFDs) for the 2001 show at the Brooklyn Academy of Music. Finally, they went to an arched grid with 231 VFDs for the Whitney Museum of Art in 2001 (Reas & Fry, 2007). In an earlier version of *Listening Post*, an Intel-based PC running Windows NT and a series of custom drivers was responsible for sending text to the array of VFDs and directing the mechanical sounds made by the clicking relays in the displays. A later source references multiple servers to manage the displays (Reas & Fry, 2007). In the current iteration of *Listening Post*, a single server drives the entire grid of screens. An ethernet-to-serial device exposes the VFDs to the display server. The screens are grouped in seven groups of three vertical columns each, consisting of 33 screens, each screen independently addressable via attached dip switches that assign an address from 0-32 in the circuit (Hansen, 2014).

Presentation Structure

Part of the genius of *Listening Post* is its theatricality, epitomized by its presentation structure, tellingly referred to by the artists as scenes. "*Listening Post* cycles through a series of seven movements (or scenes) each with a different arrangement of visual, aural, and musical elements and each with its own data-processing logic" (Reas & Fry, 2007). Ben Rubin said, "It was very conscious. We tried to structure each of these movements in the piece, to some degree theatrically, or think about them like a song so there is structure. They have some beginning, some build up, an ending to some degree. Each one is pursuing its own individual compositional logic" (Rubin & Hansen, 2013).

For example, the scene "I Am/I Like" is based on Mark Hansen's observation as he performed statistical analysis of all the chat that they were getting. Ben Rubin said Hansen asked the question, "What are the most common first few words of any given chat utterance?" As Rubin put it, "It turned out to be 'I am' was number one, day after day, hour after hour. Based on that, we have an algorithm that just scoops up all these phrases that begin with 'I am.' It organizes 80 of them in order of shortest to longest then strings them together and puts them up in that scene" (Rubin & Hansen, 2013). Rubin offered another example: "In another part of the piece in which there is a kind of chanting, words go around the room. The logic is that it is looking at the least frequently occurring words in the last two hours. It just runs that list, some hundred least-frequent words or words that only show up once" (Rubin & Hansen, 2013).

The visuals follow their own scene-based compositional logic as well. In the "I Am" scene, for instance, Rubin describes the movement of the scene across the array of screens. It "starts out in the center and then fills up kind of randomly. At the end the whole thing kind of irises down to the central screen. Others of them sweep across in one direction or another. Another one kind of comes down" (Rubin & Hansen, 2013). Mark Hansen in a workshop described another scene that we will look at in some detail: "It starts slowly with text that scrolls on each screen. Tick tick tick tick tick. It locks and then you hear a tone and then a voice narrates it. That one's trying to build up clusters of text that are all about the same topic. So it's doing a sort of dynamic cluster thing" (Rubin & Hansen, 2013).

Hansen is describing the most complex of the seven scenes that comprise *Listening Post*, one

designed to highlight content. In their 2002 ICAD paper, Hansen and Rubin offer a detailed description of the viewer experience of this scene, revealing the thoughtful attention and algorithmic integrity with which they created *Listening Post*:

This scene is best characterized in terms of an agent that manipulates text and effects changes in the audio/visual display. An agent associates itself with one of the locations in the array. At the beginning of the scene, this choice is random. When the agent chooses a location, text will scroll rapidly on the chosen display.

The agent is rapidly looking through the content stream of posts to find one that matches the text displayed in the screens around it. During this phase, the agent is refreshing the screen with new posts 10 times per second with a corresponding low-volume relay click.

With the 10 Hz refresh rate, the series of clicks becomes a whirl or a fluttering sound. In the world of this scene, that fluttering sound is tightly linked to this searching process. If a match is found within 20 seconds the scrolling/fluttering stops, a loud click is heard from this display unit, and a single message is held on the display.

As the agent finds a match, it triggers a pitched sample tone of a bell or gong and a monotone voice or matched pitch reads the text on the display. Then the agent jumps to an empty neighboring screen if one is available and starts the search pattern all over again, looking for a match among its new neighbors. If no match is made within a number of seconds, the agent jumps at random to another part of the array and begins again.

Over the four-minute duration of this scene, more and more agents appear. By the time a dozen or more agents are working simultaneously on the array, the irregular patten of clicks, tones and voices takes on an arrhythmic musical pattern. The layers of pitched voices take on the quality of a chant or litany as they blend with each other and with the reverberating tones. By the end of the scene, as many as forty voices can be heard and about 2/3 of the 110 displays show messages.

Conducting this process, the scenes are directed by an Intel-based PC running Linux, "acting as a kind of coordination engine that orchestrates communication between the various audio and display components" (Hansen & Rubin, 2002). Again, relatively simple tools, protocols, and procedural rules are put to elegant use to create emergent behavior. "Each new scene consists of a Perl program orchestrating the visual elements and controlling the overall scene structure and a Max patch/MDE program pair creating the scene-specific audio. At this point, we treat the VFD grid and the TTS engine as fixed-output devices whose programming does not change with scene; they respond to a predetermined set of commands" (Reas & Fry, 2007).



Figure 6. *Listening Post* at the MIT List Center for Visual Arts, Cambridge, Massachusetts, in 2004.
Source: Ben Rubin, EAR Studio

Systems and Software

While later versions may have updated some of the software and hardware, substantive versions of *Listening Post* including the 2002 Brooklyn Academy of Music with 110 screens used humble computing hardware and software. In the 2002 ICAD paper, Hansen and Rubin detail the systems and software that made up *Listening Post* at that time, a network of 3 PCs and a Mac G4. One PC running Windows NT controlled the array of text screens including the drivers that controlled the point click sound of the relays. Another PC running Windows 2000 controlled the TTS engine that produced as many as forty voices in the room at a time. The Mac ran Max/MSP and was responsible for all the other sounds in the room. A final PC running Linux acted as the coordination engine responsible for communications between the various audio and display components.

The subsystems of *Listening Post* function independently and as peers to one another, passing messages back and forth to coordinate timing, content, actions, and dynamics. Hansen and Rubin use an open source protocol tailored for multimedia devices and sound synthesizers developed at UC Berkeley called Open Sound Control (OSC). "We wrote a general purpose OSC client in Perl so that Max can communicate with the other pieces of the system. We created a sequence of OSC devices that specify scene type and parameter values" (Hansen & Rubin, 2002). As much as possible, the artists designed the peer subsystems of *Listening Post* to be decoupled from one another. Mark Hansen said, "You can think of *Listening Post* as a kind of instrument, with systems that render the output of our simple language processing algorithms in sound and on the screens" (Hansen, 2014).

For instance in the content-highlighting scene described in detail above, "when an agent identifies a sample to display, several events are triggered simultaneously: 1) the relay on the display makes a loud clicking noise; 2) Max generates a pitched tone, and 3) the TTS engine reads the content displayed on the screen in a monotone voice pitched to match that of the introductory tone" (Hansen & Rubin, 2002). In their second ICAD paper, Hansen and Rubin go beyond viewer experience to describe in technical detail the network and protocol negotiations as *Listening Post* composes and executes the scene. To start the scene, the controller on the Linux PC sends an OSC message to a port on the NT computer corresponding to this scene. The message specifies how long the scene

should run for. When the scene starts, the program on the NT computer sends Max an OSC message indicating that the scene has begun. It also starts a single agent scrolling on the display, and gradually introduces more as the scene progresses. When Max receives notification that this scene has begun, it sends the TTS engine an OSC message specifying the pitch and volume that the next voice should speak at.

The system sends a message over the network using the Open Sound Control protocol, of the form "/lp/content/pitchvol p v" where p and v are integers. In terms of the OSC message, it is a symbolic address structured to represent the project (lp for Listening Post), the scene (content), the parameters to be changed (p and v), and the new values. When one of the agents finds a match in the data stream, it sends the message to the display along with the specification that a loud click be issued. It also sends signals to Max and the TTS engine. The latter message consists of the text the TTS engine is to speak (at a pitch and volume previously specified by Max). The notice to Max is of the form "/lp/content/pulse".

Periodically, the audio system will receive messages that report the activity on the displays, the number of displays that currently hold content. The audio system (running Max) uses this information to set the volume of the sound in the room. These messages are of the form "/lp/content/activity a", where a is an integer specifying the number of active displays. When Max receives notice that a match happens, it plays a sample with the pitch sent previously to the TTS engine. It also sends an OSC message to the TTS engine, giving it the volume and pitch of the next voice.

This system of interconnected but independent subsystems sharing a simple network protocol is sophisticated yet simple and elegant. The temptation to update and revise the technology that *Listening Post* relies on has clearly been there for the artists. In an interview in 2007, Mark Hansen mused about how he would do things today. For example, "since 2000, the language Python has emerged as a strong competitor to Perl in this kind of application. If our development were taking place today, we would have to think seriously about programming in Python instead of Perl. The lesson here is that programming tools, and information technologies in general, are constantly in flux. If you choose software as a medium, your practice has to keep up with these changes" (Reas & Fry, 2007). This is a temptation that the artists have not failed to succumb to, as *Listening Post* appears to have been continuously improved upon, in terms of systems, software, performance, and aesthetics, since its first appearance in 2001. The latest changes with *Listening Post's* visit to Montpellier, France change the texture of the global conversation from which it samples.

Figure 7. Installation of Listening Post at the San Jose Museum of Art in 2006

The Future of Listening Post

Inevitably, new tech replaces old. This is especially apparent in aging new media works that have lost the cutting-edge and appear prematurely dated. In 2002, Hansen and Rubin mined Internet forums and IRC chatrooms. More than a decade later, while IRC and forums still exist, their use is increasingly limited. Has the decline of *Listening Post's* original sources diminished the vitality of the piece? Ben Rubin addressed this very question in a workshop in 2008. "That death of the piece will probably occur before the hardware and software give out. Mark likes to say the Internet itself may not exist five years from now in any recognizable form. It may have transitioned so totally to other modes of communication that we can certainly see the end for this kind of live chat. It's not clear. Maybe live public chat will find other technologies and continue to migrate so that we can continue to adapt *Listening Post* to monitor it" (Rubin & Hansen, 2013).

Other forms of "microblogging" have come to overshadow these older technologies. One wonders whether Hansen and Rubin have been tempted to update *Listening Post* to encompass or even wholly replace their original forums and chat rooms with contemporary microblogging sources such as Facebook and Twitter, and wide-reaching dynamic forums such as Reddit. In 2007, when the London exhibit of *Listening Post* first opened, Hannah Redler, curator at The Science Museum made clear, "*Listening Post* has a finite life span. The messaging phenomena that it feeds upon were enabled by the evolution of networks and mass access to continual bandwidth over HTML bulletin boards and internet relay chat (IRC). Changes to the text-based nature of these environments - the proliferation of video, graphics and animation - are in turn bound to radically change the content sources that *Listening Post* relies on, perhaps even rendering it silent one day" (Redler, 2007).

Ben Rubin addressed his concerns about *Listening Post* data sources in 2008: "We put the piece up, and then it goes back in the box for a year, and then we put it up again. And every time we put it up, I think Mark and I both get this feeling, like 'Oh god, is it all going to be different now? The world is changed. Now we have blogs and Facebook. Who's left chatting?' But somehow it hasn't transformed fundamentally." (Rubin & Hansen, 2013). Ben Rubin said, "When we started this piece in 1999, and even when it first premiered in 2001 and 2002, there were no social networking sites. There were no blogs. There was no Twitter. There was no Flickr. There was no Facebook. None of those things existed. So if we were to make the piece again today, it would be a very different piece" (Rubin & Hansen, 2013). As an aside, one surprise is that the oldest technologies that *Listening Post* relies upon still endures. In 2008, Ben Rubin said, "Amazingly, [Internet Relay Chat] which is a technology that predates the [web], that sort of very low-tech live chat protocol, is still around and proven to be more robust than we ever imagined. But one day, surely it won't be there anymore except for a few fanatics."

One wonders what happens to *Listening Post* when it no longer has anything to listen to. "At that point, we would switch the piece over to run in a sort of playback mode, and all the years of data it's been collecting will simply play back. The piece will become an archive of the moment of chat" (Rubin & Hansen, 2013). It seemed clear from Ben Rubin's comments in the late 2000s that he and Mark Hansen had no intention of updating *Listening Post's* sources. With the demise of IRC chat and online forums, the possibility loomed that *Listening Post* would be switched into "playback mode." At that point, the meaning of *Listening Post* would shift from a glimpse of the incomprehensible real-life, real-time communication of the entire Internet, to a snapshot of a time when Internet Relay Chat was king.

But in the decade since *Listening Post* debuted, Hansen and Rubin have not been idle. For a dozen years, they didn't let aging computing hardware cripple the system and they were not about to let shifting internet trends render *Listening Post* obsolete. In 2013, when it toured from its place in the permanent collection of the San Jose Museum of Art to Montpellier, France, Hansen and Rubin made a major change in how the piece views the on-going discussion on the internet. In an interview in late 2013, Mark Hansen told me that *Listening Post* can now run on Twitter feeds. Mark Hansen said, "When we first made listening post the data collection required some effort. These days with Twitter

and social media people often refer to a global conversation... All this seems less magical now" (Hansen, 2014).

Though for many, the global conversation and *Listening Post's* distillation of it still hold plenty of magic. Hannah Redler again, "For now, and as long as the sources it depends upon are available to its constant trawling, *Listening Post* remains an astonishing, awe-inspiring and strangely humbling 'instrument of mass, if random, surveillance and a chapel to the human need for contact' (Smith, 2003). Hansen and Rubin's creation can at times seem like a modern-day oracle, a snapshot of the text-based internet as we know it today or a monument to the ways we find to connect with each other online" (Redler, 2007).

Conclusions

Far from dispelling the magic, a plunge into the technical details of *Listening Post* reveal yet more layers of extraordinary craftsmanship and careful aesthetic consideration. We find that *Listening Post* is addressing complex ideas while combining sophisticated theory, simple systems, and elegant execution. Because *Listening Post*, still seems so current, the temptation is there to revise and update, to use newer, more powerful computers, newer programming languages, and shift to more current communications sources. This is a temptation that is obviously shared by the authors who've continued to bring *Listening Post* forward in time.

In 2000, Hansen and Rubin wrote about the political implications of gathering personal data on the internet. This has become ever more politicized since September 11th 2001, the PATRIOT Acts I and II, domestic spying, blanket wiretaps, FISA court abuses, recent NSA revelations, and so many government and corporate privacy abuses the media seem to have grown weary of dragging out George Orwell. *Listening Post* exists at the intersection of art and privacy, surveillance, and the use of personal data. In *Listening Post*, these issues lie just below the surface but are not explicitly examined. As *Listening Post* makes clear, people are constantly reaching out to others over the internet. This is more true now than it was when *Listening Post* was launched. Since then, Facebook and Twitter have set off a seismic shift in people's willingness to share their lives with strangers and prompted Facebook CEO Mark Zuckerberg to claim that privacy was no longer the social norm (Johnson, 2010).

Perhaps it is reassuring that with all that is said on the internet, that someone, somewhere is listening. Maybe it is a universal desire to be seen and heard for who we are. But by whom, and for what purpose? And is listening the same as connecting or engaging? Mark Hansen and Ben Rubin's piece is a listening post, a sentinel set up merely to observe, to categorize, to sort, to sift, and to report. *Listening Post* dutifully echos what it hears on the internet, but is resolutely neutral about issues of privacy, ethical use of our data and metadata, and the implications of gathering it. Despite its advanced age for a high-tech new media work, more than ten years on *Listening Post* is still going strong. Versions of *Listening Post* are in permanent collections in San Jose and London, and are in constant demand for exhibitions around the world. *Listening Post* is still a touchstone of possibilities in new media, and certainly still an impressive technical and aesthetic achievement.

References

Baker, K. (2007, August 4). '*Listening Post*' brings the Internet into view - SFGate. Retrieved from <http://www.sfgate.com/entertainment/article/Listening-Post-brings-the-Internet-into-view-2549303.php>

Hansen, M., & Rubin, B. (2000). The audiences would be the artists and their life would be the arts. *MultiMedia, IEEE*, 7(2), 6,9. doi:10.1109/93.848417

Hansen, M., & Rubin, B. (2001). Babble Online: Applying Statistics And Design To Sonify The Internet. *Proceedings of the 2001 International Conference on Auditory Display*.

Hansen, M., & Rubin, B. (2002). Listening post: Giving voice to online communication. *Proceedings of the 2002 International Conference on Auditory Display*.

Hansen, M., & Rubin, B. (2003). *Mark Hansen and Ben Rubin : listening post, December 17, 2002-March 9, 2003 from Whitney Museum of American Art*. Retrieved from <https://archive.org/stream/markhansenbenrub1624hans>

Hansen, M., & Yu, B. (1998). Model Selection and the Principle of Minimum Description Length. *Journal of The American Statistical Association*, 96(454), 746-774.

Hansen, M. (2013, June 20). Data-Driven Aesthetics. *New York Times* [New York]. Retrieved from <http://bits.blogs.nytimes.com/2013/06/19/data-driven-aesthetics>

Hansen, M. (2014, January 20). *An Interview With Mark Hansen* ← *modes.io* . Retrieved from <http://modes.io/mark-hansen-interview>

Johnson, B. (2010, January 10). *Privacy no longer a social norm, says Facebook founder | Technology | theguardian.com*. Retrieved from <http://www.theguardian.com/technology/2010/jan/11/facebook-privacy>

Mirapaul, M. (2001, December 10). Making an opera from Cyberspace's Tower of Babel. *New York Times* [New York]. Retrieved from <http://www.nytimes.com/2001/12/10/arts/music/10ARTS.html>

Pitman, A., & Klagenfurt, A. K. (2011). An Empirical Study of Extracting Multidimensional Sequential Rules for Personalization and Recommendation in Online Commerce. *Wirtschaftsinformatik*.

Reas, C., & Fry, B. (2007). Listening Post: Interview with Mark Hansen. In *Processing: A programming handbook for visual designers and artists* (pp. 515-517). Cambridge, Mass: MIT Press.

Redler, H. (2007). *Listening Post Curatorial Statement - Science Museum Arts Project*. Retrieved from http://www.sciencemuseum.org.uk/smap/collection_index/mark_hansen_ben_rubin_listening_post.aspx

Rubin, B., & Hansen, M. (2013). Takeawayfestival 08 - Mark Hansen and Ben Rubin, part 1 [Video file]. Retrieved from <http://www.youtube.com/watch?v=eKHoulXleEE>

Rubin, B. (2010). *Listening Post - EAR Studios*. Retrieved from <http://earstudio.com/2010/09/29/listening-post/>

Smith, R. (2003, February 21). Mark Hansen and Ben Rubin "Listening Post". *New York Times* [New York]. Retrieved from <http://www.nytimes.com/2003/02/21/arts/design/21GALL.html>

Srikant, R., & Agrawal, R. (1996). Mining Sequential Patterns: Generalization and Performance Improvements. doi:10.1007/BFb0014140

Wes Modes is a Santa Cruz artist and MFA student in the Digital Art and New Media program at the University of California Santa Cruz. In various lives, he is a sculptor, writer, performer, software engineer, and community organizer. Portfolio at modes.io

[About](#) | [Issues](#)

© NMEDIAC & individual NMEDIAC authors, editors, and programmers.