Marrying Data and Sound for Analytical and Creative Composition

Christopher Ritter

Northeastern University, Boston

Abstract

We live in a data driven world. Even our most rudimentary interactions with modern technology are tracked, logged, calculated, predicted, and modeled. I offer a reflection on this fact through two musical projects that integrate data into the fiber of their construction, allowing one to identify patterns and create connections about their use of technology through the use of music. The first “Always On,” looks inwards at myself, with a data sonification of my own personal messages over the past two years. The sonification aims to highlight relationships and provide a snapshot of the way in which I conducted digital discourse with the people in my life. The next piece, “Swipe,” looks outwards, with a composition that aims to accentuate particular interaction schemes and use-cases of technology (in particular, our mobile phones). This piece makes use of aggregated real-time data from the audience to modulate aspects of the music and dictate the sonic avenues taken throughout the piece. The end result paints a picture of how similarly (or differently) we partake in the various affordances provided by these smartphones.

Keywords

Data, composition, sonification, performance, smartphone

Introduction

As I type this out, my phone buzzes and lights up every so often, commanding my attention with each vibration. Connectivity governs our days, and “always on” culture is the norm that pressures us to be available at all hours of the day should anyone send a message and expect a prompt response for whatever reason. So, I pick up my phone, send the response, like a picture, and put it back down.

Behind the scenes, nearly all of our interactions with modern technology generate some piece of data. Whether it is to store a post so others can see it, fuel a business decision, recommend someone relevant content, predict their next action, or identify someone, the data is stored and persisted on a computer system for an indeterminate amount of time. We place a lot of trust in this data and its handling, and generate incalculable amounts of it—constantly.

Data holds an interesting position in the consumption, analysis, and creation of music. Music itself is information, whether it is instructions via musical notes in western notation on how to manipulate the keys on a flute through time, a mathematical function, or a stream of numbers on a hard drive to be reproduced as sound by a digital-to-analog converter. From information, one can extrapolate, interpolate, deduce, and synthesize new data based on what is given. How we characterize and compose music is often derived from the existing body of data, whether it comes from traditional styles, genre tropes, cultural influences, or even things that appeared earlier in the same piece.

By utilizing a broader data space beyond the music itself, a composition can be imbued with a deeper context and understanding. This could be as simple as adding lyrics, borrowing a scale from a different cultural tradition, using an algorithm to generate a melody, or even deriving computational models to generate sound based on a particular dataset. With this project, I intend to reflect on the digital status quo by composing a set of pieces with data at their heart. In particular, the data that comes from the use of our phones and social media.

Background

The compositions presented aim to create experiences to draw attention to various aspects of how we approach technology within our daily lives. I do not intend to create a commentary on whether the ubiquity of interconnectedness and data collection is good or bad, but rather establish a mindfulness of its effects through the music I have written. The interpretation is up to the listener, and I hope that the experiences involved in performing and hearing the pieces can allow people to reflect on their own relationship with technology.

In order to illustrate this, the pieces are presented within two frames of reference. The first piece, “Always On,” is an introspective look at the last two years of my own personal messages. To accomplish this, I am using techniques under the umbrella of “data sonification” in which audio or music is directly driven by data. A formal definition from Hermann (2008) explains that something might be labelled as sonification if it follows four criteria: The sound reflects objective properties of the data, there is a precise definition of how the data causes the sound to change, the sonification is reproducible under the same definitions, and the same system can be used with different data to produce a new result.

For this piece, I have downloaded the data from various sources that have facilitated the majority of my interpersonal digital communication over the years. The amount of data at my disposal is massive and has, in some manner, become a time capsule for who I was at that point in time and how I interacted with the world. Because of this, these sonifications are composed so that

various interpersonal relationships and patterns within the data become clearer to the listener and can provide abstract insight on my state through time. This includes the volume of social media use, the extent of my interpersonal communication with particular individuals, or even my mood at the time.

The second frame of reference is from that of the listener. “Swipe” is composed with the smartphone in mind, giving the audience a participatory role in the music being performed so that there is a shared experience beyond just hearing the same piece. Many of the audience interaction schemes for this piece are modeled after how we use our phones regularly, toeing the line between choreographed and open-ended motions of taps, swipes, and tilts. Data flowing from the phones to the main computer where the music is being produced is somewhat representative of what actually happens when utilizing phones, with data flowing to remote data warehouses to be processed.

Concepts and Projects Utilized

Data sonification is a relatively new field of study, but has widespread applications from medicine to music. Our ears are so well tuned to qualities of sound, sonification is a natural way to experience, pattern match, and perceive data. Any number of aspects of the provided data can contribute to the characteristics of the resulting sound.

The four most prominent techniques in sonification are auditory event markers, audification, parameter mapping, and modeling. Auditory event markers are simple sounds that are played to signal an event, like a beep when a button is pressed. Audification is the process of translating a data stream or waveform directly into sound. Parameter mapping is the association of a dimension of the data with an aspect of the sound, such as timbre or pitch. Lastly, model-based sonification is the creation of an acoustic response by interacting with a virtual system generated by the dataset (Hermann, 2011). For “Always On,” the most relevant of those are parameter mapping and event markers. In the past, I have developed several scripts for easily converting structured data into Ableton Live automations and clips. Improved versions of these scripts are the primary composition tool for the first piece. Across the piece, messages are generated as MIDI notes for auditory event markers, and parameter automation is generated to reflect dimensions of the dataset, accomplishing parameter mapping sonification.

Regardless of what methods are used for the sonification itself, the purpose of this sonification is to bring attention to the aspects and relationships of the data that are being sonified while maintaining a high level of accuracy and ease of perception. Artistic decisions are only made in accordance to their function with respect to the data. For instance, I have hand-picked specific sounds for each social media platform such that they are unique and identifiable, as well as chosen effects that I think sound compelling without intruding on other relevant dimensions of the sound that are already occupied. One would not want to parameter map an effect that alters or obscures the fundamental frequency of a note if pitch is used to denote something else within the sonification.

Additionally, as a project for a previous Directed Study, I developed a set of software that I called CrowdMap; It is a plugin for generalized smartphone control of musical performance, written for Max for Live using Max/MSP and JavaScript. This tool allows composers to leverage the many sensors and interaction schemes that the smartphone affords to modulate parameters within an Ableton Live set. The data from these devices can be aggregated in a number of ways, allowing creative applications of the phones as controllers by both the performers as well as the audience. The second piece, “Swipe,” is an interactive fixed-media piece composed for this particular plugin, to be performed on an online live-stream. Individuals can connect and participate from their phones from anywhere in the world, listening to the music being played though the stream. While conceptually

this piece is real-time, in practice there is an unavoidable feedback delay caused by stream transmission and buffering. Despite this, each performance of the piece is different depending on how the users participate throughout its duration.

This concept of smartphone control of musical parameters is not particularly new. However, my implementation is designed specifically for Ableton Live and allows for easily configurable mapping of smartphone sensor data directly to the Live set's parameters, as well as MIDI programmable configuration changes that allow the composer to structure what interactions and sensors are being utilized throughout the duration of the piece. A simple example might be to map the angle at which a phone is held to the pitch of a note. This angle can also be aggregated across all of the connected smartphones, to provide a single value that controls the parameters mapped to that sensor. Ableton was chosen as the target software due to my own experience utilizing Max for Live, as well as the recent addition of Node for Max to the Max environment. Node for Max makes networked musical components much simpler to write and work with by interfacing the Node.js JavaScript runtime engine with Max.

Related Work

In Sonification

There has been a fairly comprehensive development of research into data sonification in recent years. At the forefront is the International Community for Auditory Display (ICAD), with annual conferences that feature some of the latest in data sonification and auditory display. Each conference helps to put a spotlight on various scientific and creative applications of sonification techniques. This past year, at the 2019 conference, an array of papers and presentations from all sorts of backgrounds were introduced. Among them were methods for perceptualizing Artificial Intelligence (AI) assistance in diagnosis of melanoma via data-driven non-speech audio, auditory display and interaction for blind individuals utilizing autonomous vehicles, and a sonification of Twitter data for calming everyday users while keeping them aware of online activity (ICAD, 2019). Even from these three examples one can identify the incredible breadth of application that sonification can provide within medicine, accessibility, and everyday life.

Of particular relevance to "Always On" is a composition presented at the 2019 ICAD conference, titled "We Interact" by Daniel Grayvold (ICAD, 2019, pg. 349). In this piece, Grayvold sonifies a year's worth of text messages such that each measure in the piece corresponds to a consecutive day in 2018, with notes being triggered as text messages came in throughout that day. The length of a particular conversation with a person mapped to the pitch that was played on a piano, which caused easy recognition of patterns for periods throughout the year with flurries of messages and continued conversations that could be heard as upper register note clusters. This piece is a relevant case study for evaluating interaction through a technological medium over a span of time utilizing sonification to present a dense and expansive dataset of over 8000 texts. The listener does not need to be familiar with the data itself to be able to identify points in time with notable conversations (or lack thereof). The creation of the sonification was accomplished with data generated from Pure Data, and then sent as MIDI into the Digital Audio Workstation (DAW) Logic Pro X, where the actual sound is generated. This is similar to "Always On," however, I am utilizing Max/MSP and Ableton Live to accomplish the same thing. My implementation also allows for parameter automation based on the data, whereas Grayvold's piece is restricted to using pitch as the driving parameter (ICAD, 2019).

While "We Interact" is an effective sonification, the dimension being analyzed is somewhat limited and the output is restricted solely to a piano. "Always On" was composed to be more

complex, with several more perceivable dimensions. If the listener is made aware of the aspects being sonified, they would be able to interpret multiple streams of information at once as they listen by picking out the developments over time. While this may make the result more chaotic, I believe it is done in a way that remains musical and engaging to listen to without being overwhelming.

Winters et. al (2018) lay out several strategies for sonifying a social media feed for accessibility users, such as blind individuals. The current status quo for experiencing a social media feed through sound is a straightforward text-to-speech implementation that reads only the text details included in the post. The issue with this, however, is that social media in recent years has trended toward an image-centric paradigm, where images dominate the feeds over text posts. While advancements in AI have allowed for general identification of image content that can be relayed to the listener via text-to-speech, it still lacks much of the impact that can be conveyed through an image. For example, a picture of a woman on the beach may be labelled as “woman on a beach with palm trees” by the image recognizer, however the delivery lacks the emotion or sentiment that a picture of someone enjoying themselves in the sun may impart to an unimpaired viewer.

To offer an alternative, Winters et. al (2018) propose an AI driven sonification strategy for social media feeds. This allows each post to be represented by a soundbite containing a much broader scope of information than a text-to-speech implementation. Their strategy introduces a temporal evolution of each post, with different sections of the resulting sound imparting a different aspect of the image being represented. Using text analytics of the caption, a good or bad sentiment and impact is derived and turned into an auditory cue that is recognized to communicate emotion. For example, if someone is wishing another user a happy birthday, the sentiment will usually evaluate to “happy” and the resulting sound imparts a joyful melody or tone. Next, aspects of the photos are imparted to the user by way of auditory icons and soundscapes. These are best understood as “sounds of the moment,” or sounds that could be heard at the time a particular image was taken. For example, a picture of the woman laughing at the beach would be paired with the sound of a woman laughing, alongside the sound of waves crashing in the background. This is intuitively more impactful and immediate to the user about what exactly is in the picture and what may be understood than the computer simply reading out the words “a woman laughing at the beach.” Any other identifiable objects in the picture are also given their corresponding sounds, like a dog barking or a lawnmower. Once those finish, the environmental sounds of the picture’s setting continue at a lower volume while the post’s text is read to the user, and in some cases a short theme song may even be included here for identifiable celebrities or franchises, like the Game of Thrones theme song playing alongside a post about the show. If no appropriate theme songs could be derived, a theme song to evoke a mood identified in the image is chosen.

This is an incredibly useful application of sonification, as it tightly bundles a large amount of information into audio streams that the listener can easily separate and identify simultaneously. Unfortunately, to implement my own sort of AI to interpret my photo data from websites like Facebook and Instagram would be out of the scope for the project. However, the concept of deriving mood and sentiment from a post is one that is incredibly useful for “Always On” in order to capture exactly how I portrayed my mood through messages throughout the two-year timeline chosen for the piece. Instead of writing my own implementation of such an algorithm, I have utilized the Amazon Comprehend machine learning service that can evaluate sentiment for strings of text sent to it. I developed a script to evaluate all of my text-based messages by sending them to Amazon Comprehend to be processed and then store the sentiment data to be used for the sonification.

In Smartphone Control

There has also been a lot of work in the fields of networked audio control, using smartphones as a tool or instrument in performance, and audience participation in music through online means. The Stanford Mobile Phone Orchestra (MoPhO) has been utilizing smartphones in the performance context since 2007. In performances by this group, generally the smartphones are used as the main sound producing instruments as the performers manipulate them using custom software. Similar to CrowdMap, their approach to utilizing the phone as an instrument capitalizes on the numerous sensors and display capabilities that modern phones provide such as touch, gesture, position, microphone, and camera (Wang et. al, 2014). However, unlike CrowdMap, the sounds come from the phones themselves instead of one centralized output.

In the past, the Stanford MoPhO has also utilized phones as a medium for audience participation. Oh and Wang (2011) highlight several different examples and methods that have been utilized in past compositions. Of these, several participation schemes occurred before the pieces were performed, such that the audience was a part of the performance because they were implemented into it ahead of time by submitting recorded video snippets, for example. However, real-time participation was achieved either through leveraging social media integrations or custom interactive web applications (CrowdMap is categorized under the latter). Citing Freeman (2005), they claim that utilizing phones as an interactive tool for performance can satisfy three requirements for successful large-scale audience participation: Each person has a critical role, no one requires prior experience or skill, and the result itself is unified and satisfying.

One particular piece, “Constellation” by Madhavan & Snyder (2016), aims to analyze several different participatory audience roles throughout the runtime of the performance. The piece is divided into four sections with a particular interaction during each one. After the completion of the piece, surveys were administered to the audience to get feedback about the performance and their roles within it. The sections that most resembled the interaction schemes implemented in CrowdMap were cited by the authors as some of the most problematic. However, it seemed that the fault was in the execution rather than the design: The prompts given to the users were overly simplistic. For instance, one prompt was “SHAKE,” which the composers had hoped would be more intentional, like sounding a bell, but instead the audience shook their phones fiercely “like a rattle” (Madhavan & Snyder, 2016). To help against this situation, CrowdMap allows the composer to enter custom prompts for each mapping which can include more specific direction to the users.

Additionally, the issue of latency was raised. Any sort of networked interaction is subject to a degree of delay from input to output, which can be prohibitive if internet speeds are too slow or if the piece requires precise timing of interaction. Since “Swipe” is a live-streamed performance, the issue of latency is even greater than what Madhavan & Snyder (2016) allude to as there is a guaranteed delay of several seconds as the live-stream host processes the streamed audio and video before sending it to the viewers. Following suit with “Constellation,” “Swipe” establishes a sonic environment that does not require precise timing to participate in. Unfortunately, this means that the audience will have even less sense of how they are contributing to the final result, however, each performance will still be unique and subject to the audience’s will regardless of latency.

Matuszewski et. al (2019) have developed a software framework called *Soundworks* which operates in much a similar way to CrowdMap, where actors interact with sound systems by connecting to a central computer that takes and interprets networked controls and directs them to various other nodes in the system or sound-producers. They even utilize the same technologies, networking over WebSockets via Node.js. Their implementation has several additional connection topologies that control how control data is sent throughout their system. According to the taxonomy

described by Matuszewski et. al (2019), CrowdMap is a Situated Networked Music System utilizing a centripetal star graph network topology. This means that all of the involved nodes are connected and communicate with a single central node, instead of with each other. Matuszewski et. al (2019) evaluate a piece that utilizes this topology, “88 Fingers”, as having low perceived freedom, agency, and interactivity compared to some of the other pieces that utilize other strategies.

However, it is my belief that this is more related to the process of the piece rather than the audience interaction scheme itself. In “88 Fingers”, members of the audience get to choose one key on the piano, and when they tap their screen, the note is played. This is essentially a playground for collective improvisation, however, the improvisation is restricted to how well you coordinate with the audience around you. If no one presses their key, there is no performance. If everyone presses their keys as they please, the sense of unity of the final result is lost in the clunking of nearly all the keys of the piano being played at once. In contrast, “Swipe” allows for voluntary participation as there is still a driving performance that will continue if people do not wholly participate. Also, the degree of freedom of expression is greater due to the breadth of available interactions compared to the simplistic action of pressing a single key. Unfortunately, the perceived individual contribution to the result using CrowdMap would likely still be low due to the fact each individual’s data is averaged. For future versions of the software, a “solo” mode could be added in which each connected person gets sole control over the mapping they are assigned.

Methodology

The two compositions presented here attempt to manifest the cycles that many friendships and relationships endure—getting to know someone, talking, drifting away, reaching back out, moving on and so forth. While they lean heavily on this concept, they aim to be compelling and engaging to listen to without context. The process for creating these pieces emerge directly from their context, however, and are technologically dense.

“Always On,” a Data Sonification

“Always On” is a data sonification of two years of digital messages, from 2018 to 2020. This date range was chosen primarily since it was as far back as the data backups on my phone would go for my text messages, but I also underwent noteworthy personal development through that time and moved between various friend groups. That was something I wanted to explore within the sonification, so the following goals were established: It would identify the size of my circle of friends, highlight the conversations with the people closest to me, and underline the state of my mood through that timeframe.

The initial step was to plan the entire sonification process before even beginning it. I identified the platforms which I used most frequently to contact my friends. These would be Facebook, Instagram, SMS/MMS, and GroupMe. Snapchat would have been a valid option if I was able to retrieve more than just the last month of messages. Then, starting with the assumption that each message would simply be denoted by auditory event markers, I laid out what parameter mappings would be applied to these markers to convey the data (Table 1). In order to maintain a level of perceptibility, I kept the number of dimensions being analyzed small so that it would be easier to identify specific patterns in the data more quickly, particularly since the density of the note events will be rather high throughout the piece.

Parameter Mappings

Parameter	Representation
Pitch	Closeness: Middle C represents my own messages. Individuals are ranked by the total number of messages shared across all platforms, and the rank determines the number of semitones away from middle C that individual's specific pitch is. The closer to middle C, the more messages I have sent to and received from that individual. Ideally each person receives one pitch, although I have messaged more people than the 127 MIDI specification's number of pitches available, so all individuals in the outermost ranks outside the range are given either 0 or 127.
Velocity	Time since the individual's last message on a particular platform, logarithmically decreasing on a scale of one week. Therefore, messages from people I had not heard from in over a week will have a velocity of 1, and messages sent within minutes of each other will have a velocity of 127. The logarithmic scale corresponds to a fast dropoff of velocity for smaller time frames.
Positive Sentiment	The amount of signal that should be sent to a bus to a reverberation effect. Therefore, the higher the positive sentiment, the more prominent the reverb effect.
Negative Sentiment	The amount of signal that should be sent to a bus to a distortion effect. Therefore, the higher the negative sentiment, the more prominent the distortion effect.
Number of likes	For platforms that support liking messages (that is, GroupMe and Facebook Messenger), the number of likes a message has corresponds to a lengthening of the release time of the amplitude envelope of the sound.

Table 1

After evaluating the desired mappings, I downloaded all of the relevant data from each platform. Because of the recent General Data Protection Regulations (GDPR) from the European Union, gathering this data is as simple as requesting a download of my data from each website's settings or support sections. Some of these datasets are rich, with a lot of metadata about each message that could be used to drive various parameters. However, the only relevant information from these messages for this sonification are the message body, the sender/receiver, and the timestamp.

The message content itself can provide an interesting and expansive amount of information that can be utilized. In order to illustrate my mood over time and highlight how I outwardly presented to particular people that I spoke to, I conducted sentiment analysis on my messages utilizing the natural language processing service Amazon Comprehend. This service allows me to input my thousands of data points and receive estimates for how positive, negative, neutral, and mixed each one is by assigning scores ranging from 0 to 1. For example, if I texted a friend "I love you" it would rate highly in the "positive" metric, whereas a message like "you make me very stressed" would rate highly as "negative." Depending on how this is represented as sound, this can be very indicative of both my mental state as well as the sort of conversation I was having without giving away the details of what was actually being said. Unfortunately, the service is not perfect and may not pick up on contextual or linguistic cues that help determine sentiment. To help protect against that, the value

modulating the parameters are actually a running average of the sentiment for the last 25 messages, so that if I was really negative during a particular time it would become apparent through an increase in the average negativity.

In order to assign pitches to the MIDI notes, I had to first calculate each person's "closeness". This involved summing up the number of messages I had sent and received with each person throughout those two years across all platforms. Then, each person was ranked by the number of messages, and assigned a pitch such that the higher ranks are closer to middle C, with middle C representing my own messages that I sent.

Next, each message and parameter automation was inputted into the Ableton Live set according to those rules. To accomplish this, I used Node.js scripts within Max for Live to parse the data and interact with Ableton Live directly. Every dataset is structured differently, so each one required its own script to parse it. Lastly, I tweaked the sounds and effects of everything to my liking without treading on any of the dimensions being sonified. Additionally, the tracks corresponding to a particular platform were panned left or right arbitrarily so that it was easier to identify certain platform-specific sounds. One can isolate their listening to the left or right channel to hear a less dense collection of sounds or hone in on a particular platform without missing anything important in the other channel. This also has the added effect of having the positive sentiment reverberation effect bleed into the center channel and become more prominent when listening in stereo.

"Swipe," an Interactive Performance

Online dating is undoubtedly integrated into the current-day dating and hookup cultures of society. Even if the conversation began offline, the digital space remains to be the facilitator in many aspects of the development of a relationship. The first piece written for CrowdMap is a fixed media piece titled "Swipe." The audience will partake in replicating the subconsciously rehearsed motions of digitally getting to know someone.

Due to the global circumstances of the COVID-19 crisis, this piece in its original form as a performance in a concert space was not possible. However, it was adapted to be playable via online live-stream. The medium of live-streamed performance has a number of considerations, particularly if the performance is interactive. Audience members can join and leave at any time. They can also coordinate various interactions from the live-stream's chat feature. The video component of the livestream can be leveraged in a number of ways relevant to the performance. Lastly, and most importantly, is the issue of latency. As discussed previously, "Swipe" deals with latency by manifesting as a textured and arhythmic piece that does not require precisely timed interactions.

"Swipe" utilizes CrowdMap to drive the direction of the performance. Throughout the piece, audience members will receive prompts to perform various actions with their phones. Primarily, it allows users to experience a "choose your own adventure" selection of phases of the piece. Mimicking the infamous swipe gesture found in various dating apps such as Tinder and Bumble, at various crossroads users will receive a prompt to swipe in a particular direction on their screen based on a prompt, and whichever direction receives the majority is the sonic pathway that the piece goes to (Fig. 1). Left pathways are more abrasive and industrial, while right pathways are more harmonious and shimmering. During the first half of both phases "2a" and "2b," the audience is prompted to enter a message or a word. During the latter half, they are shown these words and are prompted to swipe left (negative) or right (positive) depending on how they feel towards that word. Offensive or profane words are filtered out before this step.

CrowdMap is also utilized for modulating the textures heard in each phase by mapping sensors in the phones to parameters in the effects that are applied to the actively playing tracks. This allows

the audience to have a direct effect on the rich and varied nature of the resultant sounds, while reinforcing the various smartphone interaction schemes that they are used to performing such as gestures, taps, and motion.

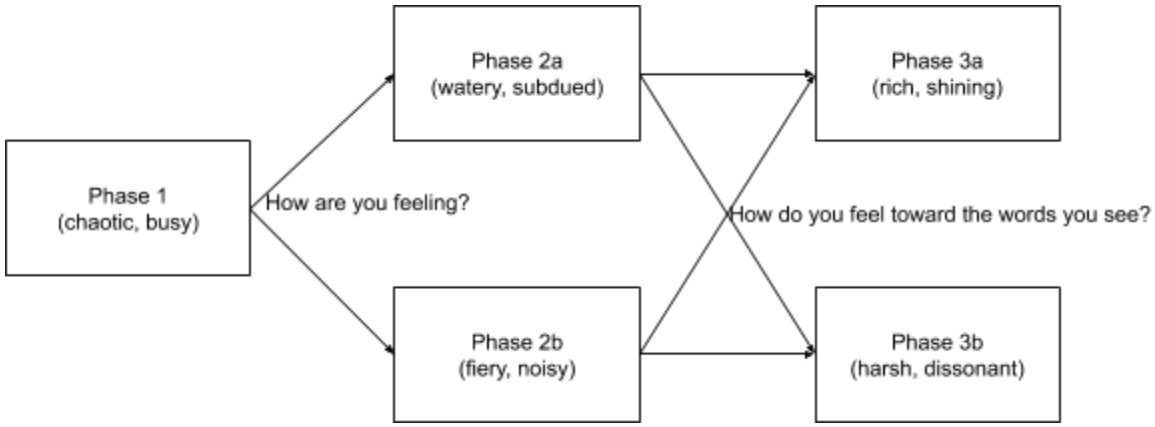

Fig. 1

Discussion

"Always On"

As a sonficiation, "Always On" turned out to be effective in portraying the information it was designed to. Listeners can pick up on various patterns of use of my social media and social life as it plays. One interesting aspect is the waxing and waning of friendships. A trained ear can listen for specific pitches across the piece and hear periods where I more actively spoke with that individual. The most identifiable example of this is the incessant minor second chord that dominates the majority of the piece. When the conversation with this person finally ends, it is marked by a pause and subsequent sudden collapse of other notes as I begin to message other people that I hadn't spoken to in some time.

There are some other interesting patterns that can be identified with a close listen. One is that I have a rather small inner circle of friends due to the fact the majority of the notes occur clustered near middle C, as I generally only message a few of my closest friends on a given day. It is also notable that my MMS messages are particularly reverberant, meaning they are typically positive in sentiment. I attribute this to the fact that my main use of MMS messages is for talking in my family group chat, which is reassuringly positive. Alternatively, Facebook messages are generally harsh and negative. Unfortunately, this may also be the fault of my choice of instrumentation to represent Facebook messages, which is spectrally rich and may be more subject to the distortion effect than less complex sounds.

Listening to the full result allows one to hear every message simultaneously, but it is sometimes more useful to isolate one platform at a time. When evaluating this piece, I found it much easier to identify patterns in particular streams of data by soloing one or two tracks at a time, which considerably lowered the density of the flurry of notes and made it more straightforward to parse. It also let me more directly compare results for particular sets of data. One last fault with the piece is that mid and higher frequency notes are much more perceptible and piercing than the low ones. Lower frequencies require higher sound pressure levels to be perceived as the same volume as mid

frequencies (Fletcher & Munson, 1933). This fact was not directly accounted for within the piece, so when the occasional distant person reaches out to me, the corresponding low frequency note may not be perceptible against the constant clusters of mid frequency notes.

“Swipe”

“Swipe” has not been fully debuted at this time, so its effectiveness in performance cannot be evaluated at this time. However, the process of composing the piece using Ableton Live and CrowdMap allowed me to reflect on what the plugin and workflow succeeds and fails with.

For one, the process of using CrowdMap for parameter mapping was easy and straightforward, and the fact that mappings are saved with the live set meant that I didn’t have to remap anything. On the other hand, once a mapping was established it was occasionally difficult to find exactly what parameter in the set it was mapped to, as the only thing visible from the plugin is the name of the parameter. This can be fixed by utilizing a naming convention for various parameters to identify them, or modify CrowdMap to display the path of a parameter to make it easier to find.

Working in Scene View of Ableton Live and organizing sections of the piece as clips rather than along a defined timeline was not as easy as I had thought it would be. Occasionally it was difficult to keep track of the playhead and start playing a certain part of the piece directly. This became more straightforward as I developed a separate Max for Live plugin to automate and control the playback and flow of the piece rather than manually triggering the scenes. This also allowed me to use MIDI clips for each scene to control CrowdMap, which was an effective way of maintaining the current stage of interaction as the playback continued.

Conclusion

Both “Always On” and “Swipe” reflected aspects of our infatuation with technology and connectedness, but rather than critiquing it they embrace its function. These two compositions have been an engaging and unique exploration into integrating data with compositional techniques. Future compositions can easily build upon and improve the techniques utilized here. These pieces each accomplished the perceptual and interactive goals that were established for them.

For future sonification efforts, utilizing different mediums and techniques of sonification would be an interesting way of further exploring the processes. For one, a more interactive listening space where the listener can isolate or “zoom in” on particular sounds would be helpful for further identifying noteworthy patterns. Also the creation of model based sonification techniques would be an interesting way of evaluating a person’s online footprint. Theoretically, differences between particular people’s models of online interaction would become immediately clear as one interacts with an individual's sonification model compared to another’s.

Additionally, this work helped to identify new features and workflows that would be beneficial to implement in future versions of CrowdMap. This includes better navigation of established mappings, more varied interaction schemes, and modes that allow individual contributions to the overall sound to be more apparent. Additionally, it would be useful to see how both “Swipe” and CrowdMap are conducted and utilized within a live, in-person performance space compared to the live-stream.

References

Fletcher, H. & Munson, W.A. (1933). Loudness, its definition, measurement and calculation. *Journal of the Acoustical Society of America* 5, 82–108.

Freeman, J. (2005). Large Audience Participation, Technology, and Orchestral Performance. http://distributedmusic.gatech.edu/jason/publications/pdf_files_of_publications/large_audience_participatio.pdf

Hermann, T., Hunt, A., & Neuhoff, J. G. (2011). *The Sonification Handbook*. Berlin: Logos-Verl. https://sonification.de/handbook/download/TheSonificationHandbook-HermannHuntNeuhoff-2011.pdf

Hermann, T. (2008). Taxonomy and Definitions for Sonification and Auditory Display. *Proceedings of the 14th International Conference on Auditory Display*. https://smartech.gatech.edu/bitstream/handle/1853/49960/Hermann2008.pdf

Madhavan, N, & Snyder, J. (2016). Constellation: A Musical Exploration of Phone-Based Audience Interaction Roles. *Proceedings of the 2nd Web Audio Conference (WAC-2016)*. https://www.researchgate.net/publication/311793642_Constellation_A_Musical_Exploration_of_Phone-Based_Audience_Interaction_Roles

Matuszewski, B., Schnell, N., & Bevilacqua, F. (2019). Interaction Topologies in Mobile-Based Situated Networked Music Systems. *Wireless Communications and Mobile Computing*, 1–9. http://downloads.hindawi.com/journals/wcmc/2019/9142490.pdf

Oh, J. & Wang, G. (2011). Audience-participation Techniques Based On Social Mobile Computing. *Proceedings of the International Computer Music Conference 2011*, 665–672. https://ccrma.stanford.edu/~jieun5/publications/2011-icmc-audience.pdf

Proceedings of the 25th International Conference on Auditory Display (ICAD 2019). (2019). *International Community on Auditory Display*. https://icad2019.icad.org/proceedings/

Wang, G., Essl, G., & Penttinen, H. (2014). The Mobile Phone Orchestra. *The Oxford Handbook of Mobile Music Studies*, 2, 453–469. https://www.researchgate.net/publication/259327976_The_Mobile_Phone_Orchestra

Winters, R. M., Joshi, N., Cutrell, E., & Morris, M. R. (2018). Strategies for Auditory Display of Social Media. *Ergonomics in Design: The Quarterly of Human Factors Applications*, 27(1), 11–15. https://www.microsoft.com/en-us/research/uploads/prod/2019/01/WintersED2019.pdf