SONIFYING FOR PUBLIC ENGAGEMENT: A CONTEXT-BASED MODEL FOR SONIFYING AIR POLLUTION DATA

Marc St Pierre

Simon Fraser University
mkstpier@sfu.ca

Milena Droumeva

Assistant Professor
Simon Fraser University
mvdroume@sfu.ca

ABSTRACT

In this paper we report on a unique and contextually-sensitive approach to sonification of a subset of climate data: urban air pollution for four Canadian cities. Similarly to other data-driven models for sonification and auditory display, this model details an approach to data parameter mappings, however we specifically consider the context of a public engagement initiative and a reception by an 'everyday' listener, which informs our design. Further, we present an innovative model for FM index-driven sonification that rests on the notion of 'harmonic identities' for each air pollution data parameter sonified, allowing us to sonify more datasets in a perceptually 'economic' way. Finally, we briefly discuss usability and design implications and outline future work.

1. INTRODUCTION

Sonification has, over the last two decades, established itself as a growing modality for conveying information and an increasingly legitimized tool, useful in many different circumstances. It can integrate successfully into workflows for control room monitoring, scientific data exploration, and even physiotherapeutic treatment [1]. Arguably, in any circumstance where there is (ongoing or continual) information that requires perception and/or action, sonification can play a part in its communication, either alone or as a complement to visual displays.

Within the broader collection of ICAD literature, there have been numerous advances in scientific sonification and accompanying issues of auditory stream perception, aesthetics and usability. However, there has been relatively little attention as yet given to circumstances in which scientific data is communicated in the public sphere. That is, when sonifications have been designed to engage a mass audience through the audible representation of scientific data, with the goal of raising awareness and allowing an everyday listener access to challenging and often 'cold' scientific information. As with other contexts, this form comes with its own design ideals, aesthetic affordances, and functional constraints to consider.

Recent sonification projects, which include the discovery of the Higgs Boson [2], Rosetta Comet [3], and gravitational waves [4], have received widespread public appeal. They are among the more salient examples of aesthetically driven mappings that are meant to engage the public in a particular way, with, of course, the openly political motive of trying to foster a public interest in scientific discovery. And also, perhaps, with the underlying political motive of trying to justify funding for costly programs. It is important to acknowledge this, because in circumstances where sonification is used as a 'public relations' tool, there is always a message underneath which serves as the guiding principle for its design: a note made most famously in Alexandra Supper's discussion of the legitimacy of sonification [5].

With sonification's increasing presence in the public domain, one of the most pressing implications is the need to rethink perceptual mappings and training as part of ongoing design developments in sonification to involve an increasingly non-specialized audience that is more used to listening to music than data. The visual equivalent of this would be the rise and popularity of conveying popular scientific/quantitative data through accessible and visually appealing infographics and data visualizations. Here is where sonification, akin to visualization, provides a unique entrypoint into understanding complex and often 'invisible' scientific, but also potentially social, cultural, and political processes. For this reason, it is useful to explore the unique ways in which sonification operates in the context of public relations. This paper will chart a practice-based research approach to a mapping strategy meant to engage the public with issues surrounding air quality in cities across Canada. Sonification, in this case, serves as a tool to communicate a broader message of mitigating emissions to the public, for citizen health, as well as for climate change activism.

2. DATA DOMAIN AND DESIGN PRINCIPLES

Climate and weather data have been a popular resource for sonification research in the past, possessing time-varying and dynamic characteristics, ripe with temporal patterns for the ear to perceive. Previous explorations of climate data have often cited these as reasons for developing new sonification methods [6]. In more recent cases, such as Goudarzi's user-centered approach [7], the stated end-goal of her design methodology is to develop mapping strategies that serve the research needs of climate scientists. The data used is complex and multidimensional, requiring many perceptually diverse mappings sounding in parallel. Scientists are also required to learn the mappings over an extended period of time.

DOI: 10.21785/icad2016.033

If we consider public presentation as the context of reception, the requirements are different. Given time constraints, sonifications often need to be simplified to drive home one or two highly important and impactful points. Fewer mapping strategies are used, and they may be more aesthetically driven. The message behind the sonification needs to be self-evident, and would not require specialized ear training to perceive.

In this project air quality was chosen among the broad spectrum of climate data based on its availability, and its recurring presence in the media's continuous coverage of issues surrounding urban environments. Because of this coverage, a public is primed to understand the connections between air pollution, their health, and the contributions of these emissions to the climate change. A sonification of air quality data, then, serves as an access point for the public to further understand the urgency of our changing environmental condition. It is this underlying outcome that guides our design.

Sonifications for four Canadian cities, Vancouver, Edmonton, Toronto, and Sarnia during the year of 2014 were analyzed. Data is retrieved from provincial websites for British Columbia, Alberta, and Ontario respectively [8], [9], [10]. Within each dataset, five metrics of air pollution are sonified. These metrics are Ozone (O_3), Particulate Matter ($PM_{2.5}$), Nitrogen Dioxide (NO_2), Carbon Monoxide (CO), and Sulphur Dioxide (SO_2). Each is sampled hourly.

3. A REVIEW OF PREVIOUS METHODS: THE AUDITORY GRAPH

Of the three conventional methods of sonification, audification, model-based, and PMSon (parameter mapping sonification), the latter is most widely used for small to large datasets with multiple data properties or dimensions. deCampo's Design Space Map [11] proves a useful starting point when designing sonifications, however, there appears to be a canonical propensity, in cases where the design space map is applied, to view a one-to-one data-to-pitch mapping as the archetype for an effective sonification, i.e. the auditory graph [12]. It is true that the ear's ability to resolve frequency deviations - the JND - is apt compared to other mapping strategies within certain frequency ranges, however this often results in sonification designs which can be aesthetically challenging to listen to for extended periods of time [13], as well as semantically limited in terms of what pitch change denotes in terms of meaning. Furthermore, if more than one data dimension is sonified, resulting pitch streams must be kept sufficiently separate in frequency so they don't overlap and risk confusion as to which data parameter is changing. These remarks are ultimately situated within one of two typical 'requirements' of sonification listening: 'exact value perception' or a high degree of correct value identification, versus the general perception of significant shifts in a continuous dataset. In the case of raising awareness, rather than scientific identification of data, this challenges perceptual considerations when designing sonifications. Conventional auditory graphs may seem appropriate for air quality data given an exploratory or purely scientific context. However, under the aesthetic and functional requirements for public relations outlined in the previous section, an alternative interpretation of the auditory graph is used, which we outline in the following sections.

4. TIME SCALING AND PARAMETER MAPPING

Before the details of the design are revealed, it is encouraged that you listen to the sonifications for both Vancouver and Sarnia first. They are publicly available at https://soundcloud.com/marcstpierre [14]. We contend that this is meant to demonstrate an implementation which does not require detailed understanding, or prolonged training, in order to holistically perceive which city is more polluted. But it is ultimately in the hands of the listener to support this argument or critique it.

To begin, there are a few stated goals of a sonification such as this one. The first is to holistically map more pollutants in the atmosphere to noise, which takes advantage of its negative connotations. This is an aesthetic decision to help communicate the detrimental effects of emissions to the environment. The second is to be able to differentiate and compare the relative levels of pollutants in each city. The third is to do this with temporal effectiveness. In short, the mapping requires enough aesthetic and functional flexibility to yield five differentiable streams that do not interfere with each other or cause auditory fatigue.

The first and most crucial design consideration is time-scaling. For this sonification, each data value representing an hour of real time, is reduced to 0.2 seconds of sonification time. A 12-hour day is therefore represented in 2.4 seconds, or a year in roughly half an hour. Daily emission patterns in the data are easily perceivable at this scale and rest comfortably within the echoic memory range [11]. With regard to streaming, several fundamental ideas from psychoacoustics and sound synthesis provide a framework for the final design. The first is Bregman's seminal stream segregation grouping cue, which states that "when two concurrent sounds have different fundamental frequencies, the brain can use the fact that the harmonics that comprise each sound will be a whole number multiple of the fundamental." [15]

Considering the design criteria for perceivability, the nature of our dataset, and these perceptual attributes, we chose frequency modulation (FM PMSon) as the most suitable mapping strategy based on its ability to efficiently generate rich harmonic spectra in relation to a fundamental. Frequency modulation is a waveshaping synthesis technique, which uses one waveform to modulate the frequency of another waveform. One wave is called the carrier, the other the modulator. When modulation occurs at frequencies above the audio rate, "sinusoidal sidebands are created at frequencies equal to the carrier frequency plus and minus integer multiples of the modulator frequency". The index of modulation is a ratio that indicates the amount of deviation from the carrier signal, and this value determines the number of sidebands on either side of the carrier, resulting a subjective experience of 'noise' [16].

4.1. Creating "Harmonic Identities" Using Stream Segregation

The five metrics are sonified in parallel using a simple mapping strategy with a positive polarity (larger data values equal to larger acoustic values). 4 of the pollutants, CO, O_3, SO_2, and NO_2, were scaled and mapped to the modulation index of an FM synth in SuperCollider. Each pollutant was given its own fundamental frequency and carrier-to-modulator ratio in different regions of the auditory spectrum. This mapping possesses multiple affordances: the first, given different c/m ratios, each pollutant occupies a fixed fundamental frequency in the auditory spectrum which remains unchanged throughout the sonification. This means that once the pollutant positions are known, it becomes very easy to identify which one is changing at any given time. Furthermore, because of the different ratios, each pollutant also possesses a distinct array of sidebands that are harmonic multiples of the modulator, creating unique timbral structures, or *harmonic identities*, for each pollutant. Importantly, this is what allows for the superimposition of streams on top of each other without perceptual and cognitive occlusion. As the modulation index goes up for each of the pollutants to the point of overlap between the sidebands, the streams remain differentiable, based on the consistent harmonic relationship to an unchanging and unique fundamental. In the same way that you are perceptually able to 'parse' out the sounds of individual instruments playing together in an orchestra, pollutants which are sonified using FM can be segregated (to a degree) based on their harmonic identities. As pollutant levels go up, the sidebands increase, becoming, fuller, brighter, and ultimately noisy. The effects of this mapping are evidenced by comparing a relatively less polluted city like Vancouver to an industrial town like Sarnia, where one sounds much more distorted and aesthetically 'harsh' than the other.

4.2. Particulate Matter ($PM_{2.5}$)

The fifth pollutant, $PM_{2.5}$, is measured differently than the other chemicals and therefore receives a different mapping. Particulate matter is commonly cited as the most dangerous air pollutant among those measured [17]. Instead of a single chemical, $PM_{2.5}$ is composed of multiple substances, some of which are quite toxic, that penetrate deep into the lungs causing cancer and other related diseases. Because of this $PM_{2.5}$ is mapped to a granular synth whose click rate increases as the amount of particles increase. This is meant to evoke the sonic archetype of a Geiger counter, where the increasing click rate signifies increased urgency / proximity, and in this case danger to listeners who experience it.

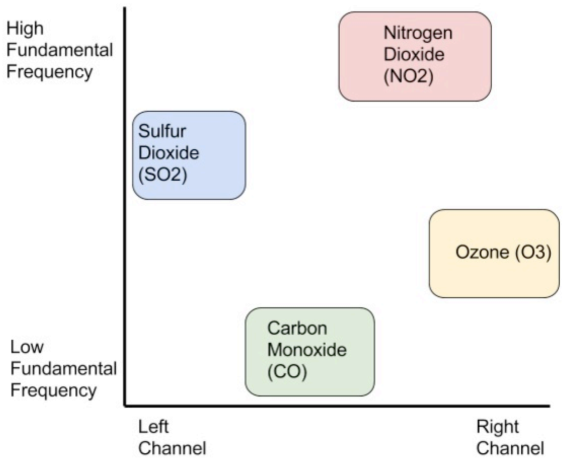

Figure 1: An approximate representation of the model's frequency domain and stereo-space mappings ($PM_{2.5}$ not represented here)

```
(
SynthDef.new(\granular, {
    arg amp, out, clickRate, pan;
    var sig;
    sig = WhiteNoise.ar(1);
    sig = GrainIn.ar(2,Impulse.kr(clickRate), 0.003, sig, pan);
    sig = sig * amp;
    Out.ar(out, sig);
}).add;

SynthDef.new(\fm, {
    arg freq, carPartial, modPartial, index, amp, pan;
    var mod, car;
    mod = SinOsc.ar(freq * modPartial, 0, freq * index);
    car = SinOsc.ar((freq * carPartial) + mod, 0, amp);
    Out.ar(0, Pan2.ar(car, pan));
}).add;

Pdef(
    \sonification,
    Ppar([
        Pmono(\fm,
            \freq, 50,
            \carPartial, 1,
            \modPartial, 2,
            \index, Pseq(~coindex, 1),
            \amp, Pseq(~coamp, 1),
            \dur, 0.2,
            \pan, 0,
        ),
        Pmono(\granular,
            \clickRate, Pseq(~pm25rate, 1),
            \amp, Pseq(~pm25amp, 1),
            \dur, 0.2,
            \pan, 0,
        ),
    ]);
).play;
)
```

Figure 2: Example synth definition and parameter mapping for CO and $PM_{2.5}$ implemented in SuperCollider.

4.3. The Importance of Redundancy in Communicating a Message

Stream segregation between the four pollutants (CO, O_3, NO_2, and SO_2) is further reinforced by the redundant encodings of the same data in different mapping strategies. Data is encoded onto amplitude, so that pollutants become louder and more salient as they increase, as well as the stereo field, so that each pollutant occupies a fixed position in space, making it accessible to a wide array of commercially available (non-specialized) speaker arrangements and listening conditions. Previous research supports the idea that redundant integral mapping strategies improve performance in auditory graph comprehension [18].

In total, each of the four pollutant chemicals possesses four dimensional attributes. Two of them: spatial position and fundamental frequency, are meant to facilitate fast and easy identification of the pollutant. The other two, loudness and number of sidebands, afford the perception of change. What is unique and promising in this design – and particular dataset used – is that a redundant one-to-many parameter mapping actually becomes a perceptual strength instead of weakness, owing to FM's maintenance of consistency of *harmonic identities,* promoting coherence of the overall listening experience.

4.4. A Brief Note on Spatialization for Public Sonification

A practical consideration, chronically under-addressed by sonification work is that the context in which sonifications are presented, at conferences, in auditoriums, classrooms etc. offer widely differing conditions for both audio quality and spatialization. Designing for public presentation means designing for a variety of conditions in mind. In many cases, good quality speakers are not readily available and they may not be spaced at a wide enough distance to encompass the entire audience within an immersive 'sweet spot'. Under these constraints, sonifications that rely solely on spatial mappings struggle to produce meaning for the audience, unless they are already encoded redundantly to other parameters. That is to say, redundant mappings are important not only because their integration with other auditory dimensions emphasize perceptibility, but also because in contexts where one mapping fails, the data can remain comprehensible. Designers in visual modes of representation, for instance, routinely account for this: it is important to choose colours that, when printed in black and white, still offer enough contrast to delineate the image. Auditory representations in practical contexts can and should operate under the same principle.

5. PRELIMINARY FINDINGS

While we have not yet had the opportunity to formally 'user test' this sonification with an actual public of everyday listeners, we have informally presented and tested it in several working groups, including two sets of audiences of 20+ each, and a sonification working group of 5 participants. Here we report first impressions from these presentations to a portion of a 'general public'. We discuss the design in terms of perceivability, but also in terms of affective engagement with the issues at hand and its utility in generating fruitful dialogue about climate issues and city infrastructures as a result of accessing complex data in this way.

Ultimately, a successful sonification design solution reveals auditory gestalts, perceptible artifacts, or patterns in the data that hopefully inspire further research directions in the respective data domain. Based on the listening experiences during the 3 working group presentations of this sonification, the dataset anecdotally generated these results:

- Most everyone was able to readily identify the most polluted city (Sarnia) and the least polluted city (Vancouver)
- Most everyone was able to comfortably identify if not interpret the four harmonic identities of the chemical pollutants, and their spatial position in the recording
- The 'Geiger counter' mapping of particle matter pollution seemed intuitive
- With some re-listening, most people were able to identify a rhythmic pattern in the datasets in terms of ozone ebbs and flows related to solar activity

What is more interesting to us in terms of auditory displays as a form of public engagement is that listening collectively (rather than individually) proved a rich way of raising and discussing a number of questions related to specific patterns of air pollution in these geographic areas. Sharing the sonification experience – and having it as a reference to come back to – allowed us to communicate important additional information to the publics we interacted with, who otherwise would not have volunteered to learn more about air pollution or its health and climate repercussions. Further, while parameter mapping might seem a 'technical' detail, folks raised a number of interesting perspectives on the use of 'logical' and counter-logical approaches to representing degrees and type of pollution. One person suggested that while noisy sidebands convey the idea intuitively, a more impactful and artistically driven approach could be using silence somehow – to signify the loss of healthy air and environments. Comments such as these opened an interesting discussion of not only semantic mappings of data, but potential connotations of mapping choices as culturally specific, and as part of a larger ecology of accessible information visualization and public knowledge translation.

One specific instance of having a fruitful and informative discussion based directly on the shared listening experience comes from listening to air pollution in Edmonton. This dataset exhibits a unique temporal emissions pattern that is not present in any of the other cities. Short bursts of emissions from different sources can be heard at extremely regular and predictable intervals throughout longer sections of the sonification for Edmonton. Our collective discussion specifically included brainstorming about possible causes of this pattern implying a spark of interest for further research into the infrastructural and industrial, as well as seasonal and environmental character of the city. Rush hour, although initially thought to be the cause, is not a possible explanation for the bursts; other cities would yield the same pattern if this was the case. What is curious is that, even though the short bursts sound at regular intervals, the source of the burst changes between the pollutants. At times there are prominent bursts in SO_2 and CO, but they will then switch to O_3 and NO_2 and back again. One current collectively-generated hypothesis is that the bursts are a result of a shared industrial practice across multiple sources of these pollutants, e.g. an active industrial complex or factory perhaps interacting with the weather. The point is that while we are not familiar enough with the data domain to make assumptions about causes, the listening experience of the sonified data provided not only enough recognition of relevant shifts and patterns, but also generates unique questions for further scientific inquiry and more importantly – public engagement.

6. CONCLUSIONS

By charting a practice-based design methodology for sonification, we demonstrate an array of aesthetic, functional, and practical design choices, which coalesce to produce sonic information. In the case of public engagement, our design work begins at a high level with a message and aesthetic we wish to convey, then works down to lower level decisions that reference synthesis methods and psychoacoustics. An enquiry into the use of sonification in the context of public relations reveals certain generalizable principles: at the most basic level, it illustrates how semiotic decisions, such as pollution-to-noise mappings, can be an integral and effective part of collectively interacting with information. Making a comparison to visualization practices for climate data, there are arrays of similar cultural/semiotic decisions that are not typically discussed as part of public presentation: e.g. the colour red is often chosen to demonstrate the most detrimental environmental effects of pollution. What we are getting at is that these decisions exist already in the contexts of public knowledge translation, and that they must be acknowledged as relying on already established archetypes and 'perceptual mappings'. Listening to sonifications, as a novel form of public engagement opens the door to having these sorts of conversations collectively and bringing attention to the culturally-specific and semiotically-driven mappings of data-to-modality. The alternative is that these decisions are ignored at the cost of truly understanding what constitutes an effective information design strategy, visual, aural, or otherwise. Our assertion is that opening up a sonification design conversation in this way can uniquely and meaningfully inform the diversity of design choices for sonification, taking also in account the context of reception and variety of listening outcomes at hand. In that sonification sheds light on the ways we choose to communicate scientific data for all potential listeners, including the broader public. Living in an ocularcentric world, it is oftentimes easy to forget that all the graphs, Venn diagrams, box plots, infographics, and visualizations we use are full of calculated design choices meant to engage a viewer in particular ways – from scaling, to color theory, to graph and chart shapes. Sonification is no different. By exploring sensory modes alternative to visual designs, we can begin to rediscover the latent practices that govern how we communicate knowledge. Sonification opens the doors to critique them, and offers solutions to improve them.

7. REFERENCES

[1] T. Hermann, A. Hunt, and J. Neuhoff, *The Sonification Handbook*. Logos Publishing House, Berlin, 2011.

[2] http://geant3.archive.geant.net/Media_Centre/Pages/Higgs-like-Boson-Sonification.aspx

[3] http://blogs.esa.int/rosetta/2014/11/11/the-singing-comet/

[4] http://createdigitalmusic.com/2016/02/a-major-breakthrough-in-physics-is-heard-not-seen/

[5] A. Supper. “Sublime frequencies: The construction of sublime listening experiences in the sonification of scientific data,” *Social Studies of Science*, February 2014, 44(1), pp. 34-58.

[6] J. H. Schuett, R. J. Winton, and B. N. Walker. “Comprehension of Sonified Weather Data Across Multiple Auditory Streams,” in *Proceedings of International Conference on Auditory Display*, 2014.

[7] V. Goudarzi, K. Vogt, and R. Höldrich. “Observations on an Interdisciplinary Design Process Using a Sonification Framework” in *Proceedings of International Conference on Auditory Display*, 2015.

[8] http://www.bcairquality.ca/assessment/air-monitoring-data.html

[9] http://airdata.aemera.org/

[10] http://www.airqualityontario.com/history/index.php

[11] A. deCampo, “Toward a Data Sonification Design Space Map,” in *Proceedings of the 13th International Conference on Auditory Display*, Montreal, 2007.

[12] J. H. Flowers, “Thirteen Years of Reflection on Auditory Graphing: Promises, Pitfalls, and Potential New Directions,” in *Proceedings of the International Conference on Auditory Display*, Montreal, 2007.

[13] https://soundcloud.com/syssonproject/toa-in-and-out-radiation

[14] https://soundcloud.com/marcstpierre

[15] S. Carlile, *The Sonification Handbook*. Logos Publishing House, Berlin, 2011, ch. 3 – Psychoacoustics, pp. 41-61.

[16] P. R. Cook, *The Sonification Handbook*. Logos Publishing House, Berlin, 2011, ch. 9 – Sound Synthesis for Auditory Display, pp. 197-235.

[17] http://www3.epa.gov/pmdesignations/faq.htm#0

[18] S. C. Peres and D. M. Lane, “Auditory Graphs: the Effects of Redundant Dimensions and Divided Attention”, in *Proceedings of the International Conference on Auditory Display*, Limerick, 2005