Tidal data sonification as a case-study of the challenges associated with mapping information onto the sound domain

Tanguy Pocquet du Haut-Jussé

Contents

Abstract

Looking at existing data sonification works; at research conducted in sonification, acoustic ecology, acoustics and psychoacoustics; and at my own work[1] exploring the sonification of tidal data, this dissertation is a mixture of theoretical — in an attempt to evaluate the difficulties (both artistic and technical) faced when sonifying data — and practice-based — presenting compositional work reflecting theoretical observations — research.

Introduction

This project is an exploration of the sonification of a very simple dataset as a way to engage with various issues inherent to all sonification, with the aim to create a work of sound art that strikes a balance between the often conflicting analytic and artistic approaches to data sonification.

I will begin with a review of existing sonifications, which all prioritise either of these two approaches, and go through the different shortcomings that some of them face. I will then introduce the dataset, my aims, and the final sonification approach, before reflecting on the various practical, theoretical and contextual issues I have had to face over the creation of this work. These include questions of mapping between data and sound, of legibility of the sound output, of aesthetic value of the resulting work, of semantic coherence between the type of sound used and the object of the sonification, and of ecological associations that those sounds might create and of how to take advantage of them. These reflections will be accompanied by various examples of earlier iterations of the sonification of this dataset, which illustrate some of the problems I had to face and how I came to resolve them.

The first aim of this work is to provide a comprehensive — although not exhaustive — presentation of the areas that I have had to familiarise myself with in order to achieve the practical side of this work, and which are common to many sonification works, regardless of the data they focus on. The second is to present the practical output of this thorough exploration of the various fields that intersect within data sonification. I hope this work can

[1] in the practical component folder, see the ./Tidal_forms.mp4 file for a rendition of the finished sonification piece

display some artistic and some technical value to the listener, as my approach is very much motivated by a need to explore the symbiosis between data and sound, to use data to create art that will in turn give us insights into the information it is derived from.

Review of existing sonification works and approaches

Definitions

The most commonly cited definition of sonification is “non-speech audio to represent information” (Kramer et al. 1999; Hermann, 2008; Walker, Nees in Neuhoff, 2011), but it has also been defined as “the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation” (Kramer et al. 1999); the “data-dependent generation of sound, if the transformation is systematic, objective and reproducible, so that it can be used as scientific method” (Hermann, 2011); or as “the rendering of (typically scientific) data into (typically non-speech) sound designed for human auditory perception.” (de Campo, 2007). The latter is the definition I will use, as I see no reason to necessarily exclude speech from the definition, as it has proved a useful tool in multiple sonifications (Jeon & Walker 2011; Walker, Nance, & Lindsay 2006; Jeon et al. 2009; Nesbitt & Barrass 2002), and I do not believe that usability as scientific method is a necessary condition of sonification: some of the very successful sonifications that will be introduced in this section are interesting artistically, and bring some interesting insights on the data used to the listener, but are not rigorous enough to be useful scientific monitoring tools.

What is clear from all these definitions is that sonification is an inherently multidisciplinary field, involving data science, music and sound art, acoustics, and psychoacoustics, in addition to the fields to which the data sonified might belong. Hence, any sonification work has to balance the priorities (aesthetic, scientific, analytic, etc.) particular to each of these disciplines, as difficult this might be. This is particularly the case when having to prioritise between clarity and aesthetic in the sound output.

Purely artistic approaches

A lot of examples of sonification are purely artistic endeavours, where the clarity of the data being sonified is not as much of a concern as the aesthetic value of the work. They are closer to systematic compositions than to attempts at conveying information through sound. They often feature data from fields that are inspiring artistically and engaging for the general public: the weather (::vtol:: & Alekhina, 2015; Kaeser, 2018); urban activity (Alexander, 2011; Cherouvim, 2017; House, 2012; Reeves, 2014); text (Ruten, 2017; Alhoff, 2022); Covid-19 (Chelidon Frame, 2020); etc.; and result in sound works that are often literal mappings of complex, dense numerical information, making the output entirely unintelligible even with context (Cherouvim, 2017; Alhoff, 2022). Moreover, a lot of these sonifications are accompanied by a visualisation element (Alexander, 2011; Cherouvim, 2017; Alhoff, 2022) that provides the only clues as to the link between data and sound.

I would also include in this category of purely artistic sonifications compositions that few would even qualify as such; ones where the material is somehow derived from non-audio information in a systematic way. This includes works like *Music for viola* (Kyriakides, 2014), in which musical phrases are created from a systematic translation of the letters in the word ‘music’ in various languages into a series of viola harmonics.

Scientific approaches and attempts

Looking at the other end of the spectrum, examples of purely scientific sonification are much rarer. There are nonetheless unmissable examples: heart monitors and Geiger counters, just to name a couple. The main aim in this case being the absolute clarity of the information presented, the sound outputs are extremely bare and simple. Aesthetic concerns are completely absent, except perhaps to ensure the sound is audible and not too irritating.

Another major group of perhaps not strictly scientific, but purely technical sonification examples are earcons. These are defined by Blattner et al. (1989) as “non-verbal audio messages used in the user-computer interface to provide information to the user about some computer object, operation, or interaction”; the sound of paper falling in a bin when deleting a digital file for instance. Again, there are barely any aesthetic concerns at play here.

All of these examples have in common their simplicity; as soon as one looks at sonifying more complex data, the mapping concerns to ensure the result is clear increase exponentially, which means sonification is rarely deemed an efficient enough solution for technical applications that are not the most basic.

Although more complex, and apparently technical sonification works, for instance algorithmic sonification (Blackwell & Young, 2004; Asonitis et al., 2022; Winters et al. 2019), or quantum mechanics sonification (Kontogeorgakopoulos & Burgarth, 2014; Saranti et al., 2010), have been explored, they remain largely restricted to the research world, with almost no example of large scale applications. Moreover, they are generally published by researchers also interested in the sonifications' artistic applications; the research output is often a way to present the scientific aspects of what are pluri-disciplinary arts and science projects. Ironically therefore, complex sonification is more of an artistic field than a scientific one.

Pluri-disciplinary work

Artistic sonifications are not limited to the few examples cited above, where the data is little more than a source of inspiration. There are other examples where the aesthetic concerns did not completely inhibit the clarity of the data sonified, and where the sonified information adds an interesting layer of meaning to the artistic work. These include works that are particularly relevant to my project, sonifying tidal data (Berrick, Cunningham & Griswold, 2022; Myrbeck & Shisko, 2012), water data (Bellona, Park & Reagan, 2016), or data derived from other natural phenomena (Briggs & Corey, 2015; Renard & Le Bescond, 2023) but also projects looking at similar fields as mentioned above, including urban data (Foo, 2015; Myrbeck & Shisko, 2012), or Covid-19 data (Rebelo, 2020).

Most of these examples have in common the simplicity of their input dataset, looking at a single stream of numerical information, behaving somewhat predictably: cyclical tidal data (Berrick, Cunningham & Griswold, 2022; Myrbeck & Shisko, 2012), exponentially increasing cases of Covid-19 (Rebelo, 2020), unidirectional income change by area (Foo, 2015), etc. The exception is Renard & Le Bescond's (2023) work, in which the sound output is made clear not by the simplicity of the data, but by using pre-existing music as the sonification output, and varying parameters within it, making changes predictable, rather than mapping numerical data onto the sound domain directly.

It is also worth noting that most of these examples still have a visual element accompanying the sonification, which does help to contextualise it (Berrick, Cunningham & Griswold, 2022; Bellona, Park & Reagan, 2016; Rebelo, 2020; Foo, 2015; Renard & Le Bescond, 2023), and that some of them would not, in my opinion, be enough to enable a listener to gather precise and accurate information about the input dataset, but then again that is rarely the goal of sonification.

Other approaches suffering limitations

Finally, there are sonification works where the artistic concern was maybe not as prevalent as in the first works cited here, yet where the sonification still suffers limitations. These include localised weather sonifications where it is difficult to take away much from the sound, but where the comparison between two locations is interesting (Bultitude, 2017; McCorkle, 2018; Quintronics, 2014), sonifications in different domains paired with visualisations, where the sonification adds very little to the perceived information (BBC Visual and Data Journalism team, 2020; System sounds, 2018; Holtzman, B. & The Seismic Sound Lab, 2006), and ones, similar to Kyriakides' work mentioned earlier, where the data is used to generate a musical score, but where its complexity means it is less well integrated than for Kyriakides, and brings little to the pieces (Miebach, 2009-2019; Carter, 2013).

Finally, it is also worth noting that there are examples I would consider bad in all regards: *gimmicky* uses of sonification, where the data sonified is not even remotely perceivable in the sound output, and where the aesthetic value of the latter was not particularly considered. These are literal, arbitrary, and often simplistic mappings of data onto the sound domain. A notorious example is NASA's sonifications (NASA, 2019; NASA, 2022), but there are other examples by smaller entities, and individuals (Lorenzo, 2010).

Secondary classifications of sonification work

Beyond this primary classification from artistic to scientific in their aims, there are other useful ways to look at these sonifications, by placing them on secondary spectrums, for instance one looking at the mapping of data onto sound, going from literal to abstract/derivative. Their place on this spectrum depends on decisions on the mapping between data and sound, and impacts the clarity of the data presented, as well as the artistic and aesthetic value of the auditory display. Cherouvim (2017) would be an example of a completely artistic (the goal is not for the data to be fully understood and assimilated), yet completely literal (there is a direct mapping of data onto sound), whereas Renard & Le Bescond's (2023) work would be close to fully scientific (it has no artistic ambition, and the musical dimension is only used to make the data more intelligible), and somewhere between literal and derivative (there is extraneous musical information that is not a literal mapping of the data, although the mapping process remains very clear overall).

Another common way to assess the aims and priorities of a sonification, theorised by Keller & Stevens (2004), is to look at their place on the analogic-symbolic representation continuum:

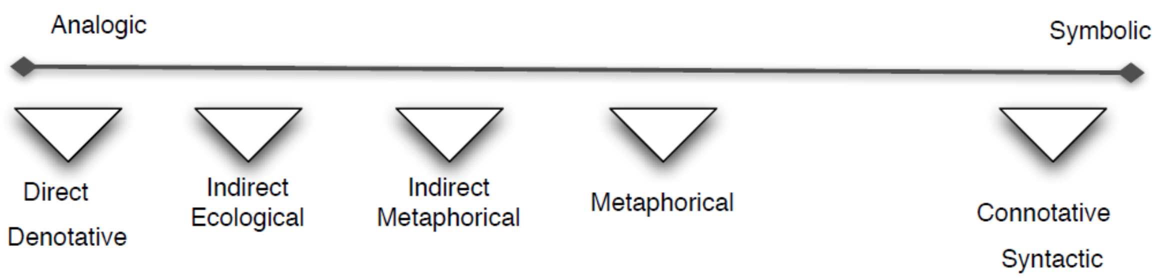

Fig. 1 Analogic-symbolic continuum

Rather than looking at the mapping strategy, this continuum looks at ecological ties between the data and the sounds used to represent them, and can be explained in the following way: "Direct relationships are those in which the sound is ecologically attributable to the referent. Indirect ecological relationships are those in which a sound that is ecologically associated with, but not directly attributable to, the referent is employed (e.g., the sound of branches snapping to represent a tornado). Finally, indirect metaphorical relationships are those in which the sound signal is related to its referent only in some emblematic way (e.g., the sound of a mosquito buzzing to represent a helicopter)" (Nees & Walker in Neuhoff, 2011), whereas a connotative syntactic representation would be devoid of any metaphorical link to the data (a heart monitor, for instance, does not attempt to sound like a heart at all).

Virtually all of the examples we have looked at are on the symbolic end of the continuum (None use field recordings or synthesised sounds ecologically or metaphorically attributable to the object of the sonification), which is a very interesting thing to notice. The value of analogic approaches, but also the difficulties associated with them will be discussed in relation to my own attempts at creating ecologically-motivated sonifications for this project.

These questions of mapping, literalism, acoustic ecology and semantics of sounds used are crucial when designing a sonification, were at the core of the creative process for my own piece, and will be discussed in more details, looking at their theoretical background and application in my work.

Choice of data to sonify and working priorities

This project uses tidal data (Shom, 2023) from 6 locations on the Atlantic coast of France (see Fig. 2), in the form of water-level readings, taken every minute for the first week of 2023, as the object of the sonification (see Fig. 3). I chose this data because of its inherent simplicity. Tide data is essentially sinusoidal (although modulated on longer time scales by another sine function), and therefore cyclical, predictable, and so is its rate of change (the derivative of *sin(x)* being *cos(x)*, a very similar, simple, periodic function).

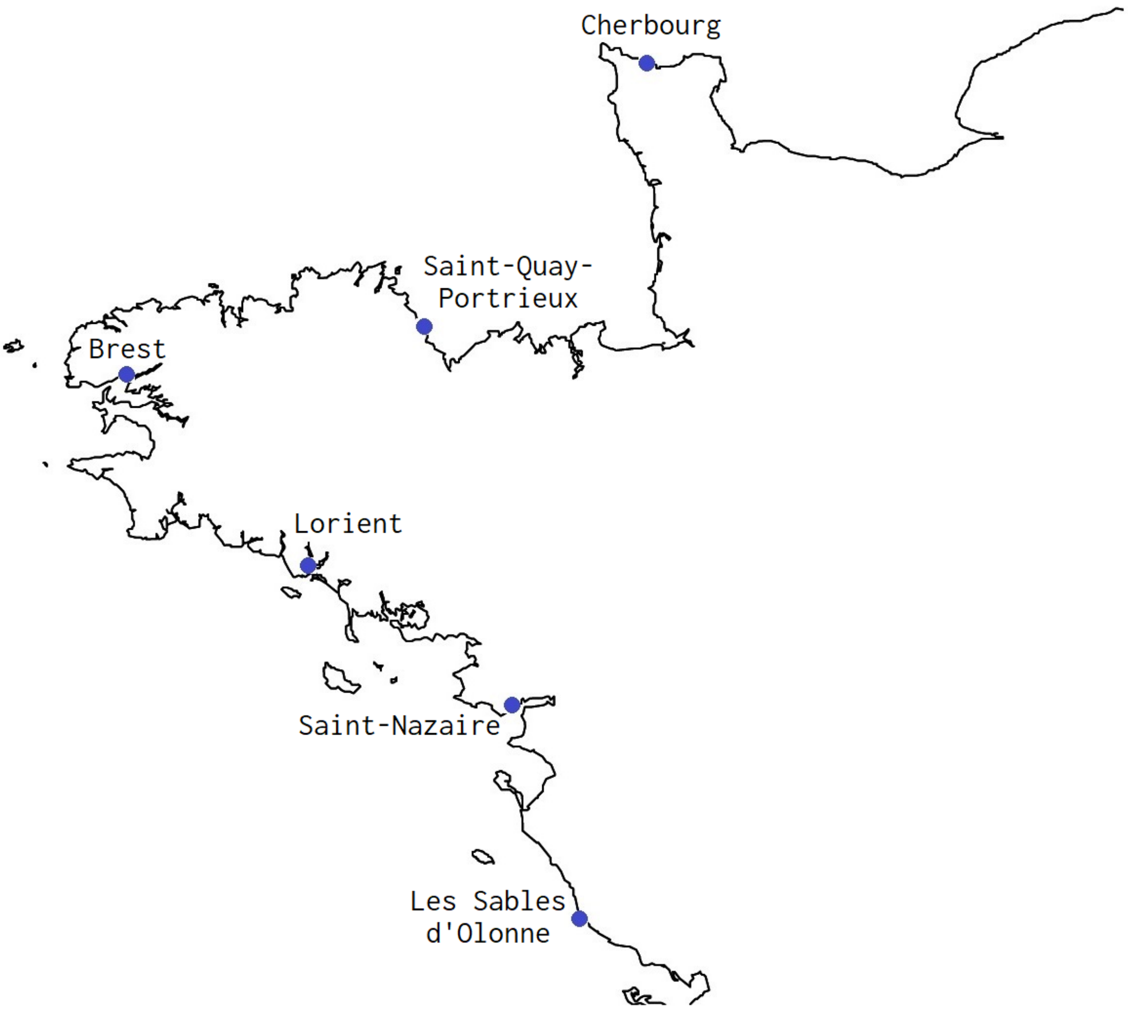

Fig. 2 Map of all six locations whose tide data was sonified

There is therefore very little, if any at all, extraneous information to be derived from the data, which limits possibilities for its sonification, and forces me to focus on making the little information I have as clear as possible. The only thing that adds some very limited complexity to the input data is that it features streams from 6 different location, which although all on the same coast, are out of phase with one another.

My goal with this project is to create an artistic work that I am happy with aesthetically and fits my usual compositional interests (the explorations of slow timbral transformations; non-narrative structures; movement within apparent stasis; noise to foreground or drown out pitch), but is also driven by external information, and through this provides an additional layer of meaning to audiences that other pieces might not. My hope is that with minimal contextual information, a listener would be able to understand how the data relates to the sound they are hearing.

I am also interested in presenting information linked to a familiar phenomenon (tide patterns) whose rate of change makes it difficult to perceive, and render it as a tangible auditory experience, compressing a week into a few minutes. The fact that “When an observer enters a space, some aspects of their environment are obvious, but many phenomena (e.g. barometric pressure, climate change, soil moisture) remain mostly imperceptible because we do not have appropriate biological sensors to detect them, they are too large or small, they change on timescales that are too long or short, or they are beyond our reach.” (Lynch & Paradiso, 2016), and the subsequent exploration of imperceivable and intangible phenomena through their rendering as discernible sound structures is something that is interesting to many people working with sonification — as well as visualisation — and it is definitely the principal reason for my own interest in it.

The position of wanting to create a piece that is a work of art but also has the ambition to present intelligible information is an inherently precarious one, since, as we have seen, it can be difficult to balance the scientific and artistic sides of sonification. I do however believe this is the right approach, firstly because it is challenging, and intellectually more rewarding to try and reach this balance than to choose one side, and secondly because setting ourselves up for some degree of failure is an intrinsic part of sonification; the resulting sound will never be simultaneously as close to our aesthetic aims and as clear informationally as we would wish, mainly because of mapping issues that will be discussed in the next section, whereby “More complex mappings, while more satisfying from a [musical] perspective require careful consideration and tuning in order for the relationship between movement and sound to attain syncretic coherence” (Bencina, Wilde & Langley, 2008). This is a particularly engaging premise from an artistic perspective, and might be the best angle to approach these data structures today: “the rate at which we collect and generate data and information vastly outpaces our ability to process and interpret it. […] We need powerful new paradigms to render these data interpretable. On the path to these new paradigms, these data present themselves as a canvas for the contemporary artist.” (Lynch & Paradiso, 2016).

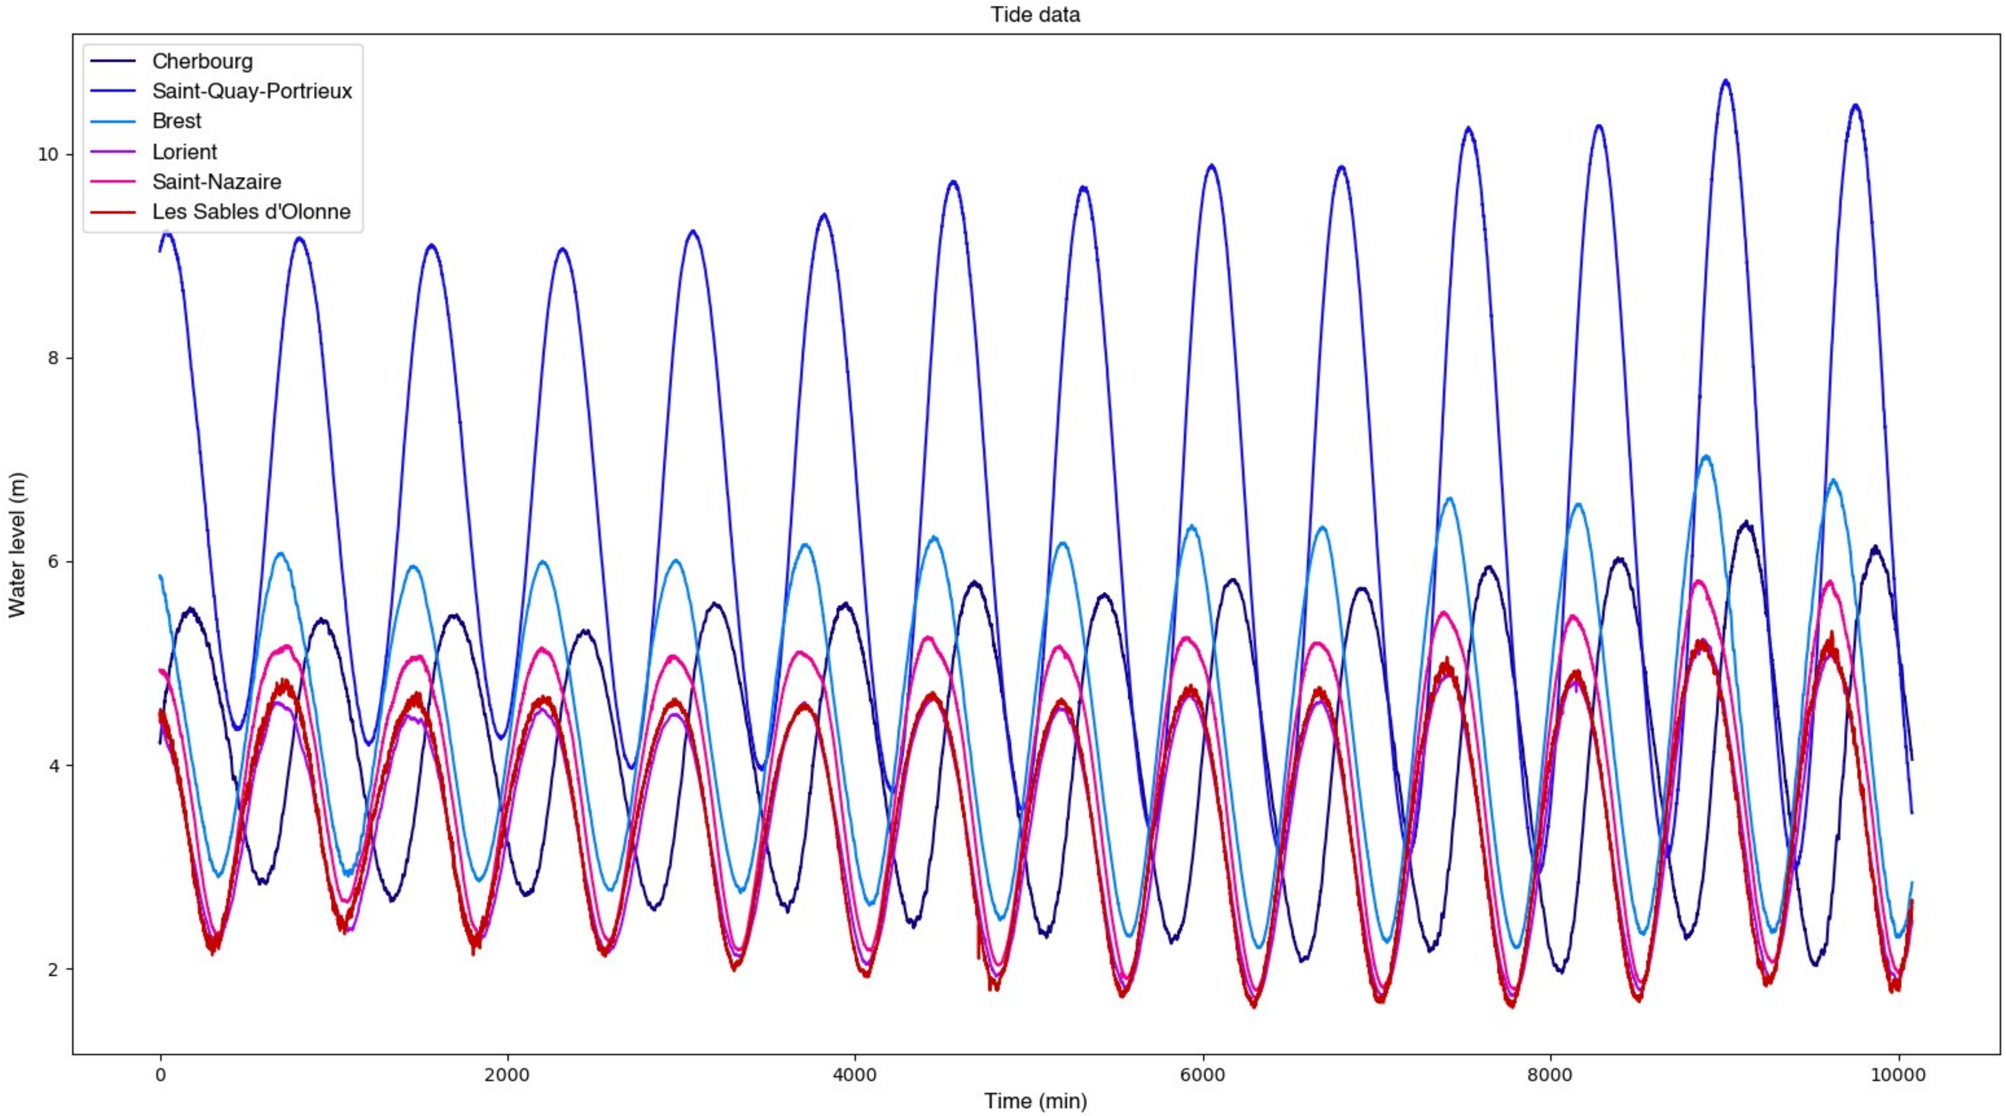

Fig. 3 Tide evolution over a week in the six chosen locations

Description of the final sonification

Over the course of this project I have gone through a series of attempts at sonifying the dataset I just described. I will now present the final version of the sonification, as the rest of this dissertation will be dedicated to a reflection on the various issues (technical, artistic, contextual) I had to face in order to reach this result, alongside descriptions of previous, unsuccessful, iterations of the work.

This sonification makes use of granulation, a process by which sound files are segmented into tiny fragments (grains), that can then be “reassembled in a new time order and microrhythm” (Roads, 2001), enabling the sound to be slowed down or sped up while maintaining its pitch (since the grains themselves are not stretched). It is crucial to note that “granulation is a purely time-domain operation” (Roads, 2001); my focus on this dimension is the result of numerous attempts at exploring other sound dimensions to map this data onto, that were all somewhat unsuccessful. I found this approach particularly effective, in a way that reflected the findings that “auditory scatterplots are quite efficient in conveying sign and magnitude of correlation” (Flowers, Buhman & Turnage, 1997).

Practically then, each data stream (from each location) has a sample associated with it, which is granulated, and the parameters of each granulation engine change with the live data. These samples are six different guitar harmonics, from a guitar tuned in a 7^{th}-limit just intonation system, which, when layered, create a complex, consonant chord. The selection of this specific set of pitches was made with auditory streaming and acoustic masking issues in mind (issues that will be discussed later on), in an attempt to make the resulting auditory landscape both clear (each stream should be discernible) and coherent (the sonification should be perceived as a unified soundscape).

The parameters that are changed with the data are the place of the grains within the sample, their length, and the density of grains. The higher the water-level data, the closer the grains are to the start of the sample (the attack), the shorter they are, and the more are being played. This means a high-tide is represented by a texture of sporadic, independent, discernible grains with sharp-attack and quick decay, whereas the low-tide is represented by a much more continuous, unbroken, drone-like texture, made up of fewer, longer grains, placed further down the sample (in the sustain of the harmonics). I find this approach and this contrast to be extremely effective to sonify a one dimensional change.

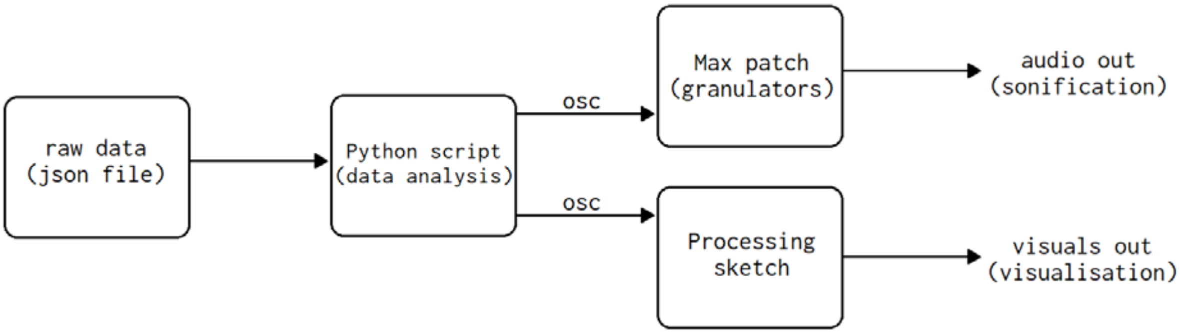

Fig. 4 Path from raw data to sonified output

The data (Shom, 2023) was downloaded as six JSON files[2], which are imported into python [3], and read, with an entry (water level readings taken every minute) being sent via the OSC (open sound control) protocol to a Max patch [4] and Processing sketch every 0.06" (accounting for Python's processing time) (see Fig. 4).

The data, in meters, spanning roughly 0m to 10m, is then scaled within Max to fit the sound dimensions it affects:

- 0-10m are scaled exponentially to 5000-80ms for the starting point of the granulation within the sample (the higher the water, the closer to the attack the grains get, and the more aggressive the texture)
- the data is also scaled exponentially to 20-0s for the size of the window, impacting the length of the grains (the higher the water, the shorter and more aggressive the grain sounds)
- these same 0-10m are scaled linearly to 300-10ms for the interval between two grains (the higher the water, the shorter the time between two grains, the busier the texture)

A similar scaling happens within the Processing sketch [5], and the visualisation, consisting of circles radiating at varying pace (the higher the water, the quicker the generation) from the six

[2] see files in ./data

[3] see ./script.py for the Python script

[4] see ./patch.maxpat for the stereo version of the patch, or ./patch_spat.maxpat for the spatialised version

[5] see ./visualisation/visualisation.pde

locations on a coastal map, varies with the live data. This will be briefly discussed later, but is not the main focus of this dissertation.

It is worth noting that the minimum and maximum points for each dataset are different, as some locations will have much higher readings on average than others (as is clearly shown in Fig. 3). Hence the above scales will vary slightly between locations. Given that my goal is to bring out patterns in the tide cycles and phasing effects between locations, rather than provide a comparative analysis of tide levels across a coast, I believe it is much more useful to use relative mappings for each location.

Time was scaled down by a factor of 1000; a reading corresponding to 1’ in the original data is sent out every 0.06”; this compresses the week-long dataset to a little over 10’. The extracts for the failed sonification attempts that will be described in the text are shorter, scaled down by a factor of roughly 3000 (each extract is approximately 3 minutes long), so that one can get an idea of their structure more quickly.

On mapping

The first concern in sonification is mapping, the link between the data and the sound, the “designed connection between signals and parameters that creates the interdependencies between otherwise independent parts, making them into a whole” (West, Wanderley & Caramiaux, 2020). This is not a concern solely in sonification, and has been widely studied with regards to electronic musical instruments, thinking about links between control parameters and output signal (Bencina, Wilde & Langley, 2008; Hunt, Wanderley & Paradis, 2002; de las Pozas, 2020).

The way a mapping is designed is often seen as inherently arbitrary: “gesture-sound relationships which comprise a mapping strategy are arbitrarily assigned by the designer based on their aesthetic values or practical needs” (de las Pozas, 2020), which allows for artistic licence from the person designing the mapping, and leaves room for very different approaches, which may be appealing to some users and not others. Yet in sonification, “there seems to be some agreement among listeners about what sound attributes are good (or poor) at representing particular data dimensions” (Nees & Walker in Neuhoff, 2011), for instance, pitch is generally good for representing temperature, whereas tempo is not as effective (Walker, 2002).

Knowing what sound-data relations might be more intuitive to listeners helps reaching the ultimate goal of mapping: to “maintain a balance in the relationship between [data] and resultant sound that is easy to perceive for the audience” (Bencina, Wilde & Langley, 2008), and achieve the highest level of clarity possible between input and output. This was particularly studied in relation to the oldest tools used to transform human input into sound output: musical instruments. A violin’s string, for instance, is “both part of the control mechanism and the sound generator” (Hunt, Wanderley & Paradis, 2002), but in the digital world, “in stark contrast with acoustic instruments where the dependencies between parameters are unchangeable, cross-coupling between variables can easily be created or destroyed” (ibid), and our job is to design connections between the source and the sound-product that are as clear as possible.

The way this balance is achieved is generally by trial and error, and mapping-design is an inherently experimental process: “the designer supposes that a certain change to the mapping might improve it, then they test their hypothesis by making the change and playing with the modified [parameters]” (West, et al. 2020). West et al. (2020) have found that in their reflection after an experiment asking them to come up with efficient sound mappings, “participants especially used the terms trial and error (6/9), experimentation (4/9), discovery (3/9), and/or exploration (3/9)”.

Attempt #1: pitch sonification

Given that sound is an experience of frequency changing over time, and can be defined solely by those two dimensions, and given that the latter will always be a component of any sonification, as one cannot experience sound outside of it — “Sound exists in time, over space. Visual objects exist over time, in space” (Gaver, 1989) — then perhaps the most obvious other sound dimension one would think of mapping information upon is pitch. In the case of this work, the data I am looking to sonify is also two-dimensional, and also includes a temporal axis. Therefore, it made sense to begin with an attempt to map the other dimension (water level) onto pitch.

I approached this by taking the maximum and minimum water levels in the datasets of each location, assigned those to a high and low frequency, framing a comfortable area of the hearing range (50Hz – 2000Hz), and mapped the water data of each location exponentially onto the frequency of a different sine-wave oscillator, so that linear changes in water level correspond to exponential changes in frequency (which we experience as linear) [extr 1a] [6]. Given that each location is represented by a sine tone, it is very difficult to perceive individual streams, and the result is quite unclear. Panning the locations left and right depending on their relative position only makes this slightly better [extr. 1b] [7]. It is also difficult to keep track of longer-form data changes (is the peak you hear higher or lower than the last one you heard; is it always the same location that reaches a peak or trough first; are there recurring phasing patterns between locations; etc.).

As I mentioned for the final version of the sonification, each mapping is relative to the minimum and maximum values for each individual location. I believe this pitch sonification provides a good illustration of why I made this decision: I am providing a version of this first mapping where all locations' data is scaled in the same way, and where because one location happens to have much higher readings, it is the only one going over 400Hz, while the rest form an unintelligible, bass-heavy backdrop [extr. 1c] [8]. This difference in average water-levels could be due to geographical and oceanographical features of the location where the reading was taken, and therefore not as interesting to me as the relative highest and lowest points that each location will get to, how fast or abruptly they will reach those extremes, how much earlier or later than the surrounding locations, etc. (in this case, as is made clear in Fig.3, the second location, Saint-Quay-Portrieux, has much higher readings, because it is located in the Mont-Saint-Michel Bay, where the Atlantic and the English Channel join, and where the average difference between high and low tides is one of the highest in Europe (Marchand, 2020)).

Specific practical concerns when designing a sonification

There are multiple factors that come into play when designing a mapping. As we mentioned, some sounds are generally seen as better to represent certain types of information than others. Some are also seen as poor mediums to convey any information: "some sound dimensions (e.g. loudness) are not very effective in auditory displays for practical design reasons"

---

[6] see ./extracts/1a_oscillator.wav

(all sound examples for past sonification attempts quoted in the text are in the 'extracts' folder)

[7] see ./extracts/1b_oscillator_stereo.wav

[8] See ./extracts/1c_oscillator_stereo_same_scale.wav

(Neuhoff, Kramer, & Wayand, 2002). Choosing which audio components (pitch, volume, timbre, grain, etc.) to use is therefore an essential concern.

There are various possible approaches to sonification, the most basic, *Parameter mapping* sonification, “represents changes in some data dimension with changes in an acoustic dimension to produce a sonification” (Nees & Walker in Neuhoff, 2011); the first mapping I described above is an example of this approach. *Model-based* sonification is another type of sonification, where a model is designed for the listener to interact with, and where “the sonification is the reaction of the data-driven model to the actions of the user” (Hermann, 2002). With this work, I have chosen an approach that is at the intersection of these two. I have not simply mapped data dimensions onto sound dimensions, but created granular synthesisers and associated multiple of their parameters to one data dimension. In one sense this engine fits the model definition, however it lacks the interactivity generally required to fit the model-based sonification definition.

Beyond the choice of which sound dimension(s) to use and of which data dimensions to associate with them, another concern when looking at mapping numerical information onto the sound domain is the practical parameters under which this translation will occur. Of these parameters, polarity and scaling are the most important. As we have seen with different sound dimensions, Walker (2002) has also found that certain polarities were generally perceived as clearer by listeners; hence they might find that pitch should increase in order to represent increasing temperature (a positive mapping polarity), while at the same time feel that pitch should decrease in order to represent increasing size (a negative polarity). Likewise, Walker (2002) has also estimated scaling factors for several mappings. Knowing how to scale the incoming data to fit whichever sound dimensions was chosen to represent it is a crucial part of the sonification work. This comes with added challenges when what humans perceive as linear changes in pitch and loudness are in fact exponential changes, and this needs to be taken into account when mapping data onto those dimensions.

In the case of this sonification, I decided to associate higher water-levels with a busier, more aggressive, more granulated texture, and lower water-levels with sparser, drone-like, sustained and continuous sounds. I believe this greater closeness of the sound when the tide is high is an immediately intuitive way to represent the data aurally. Again, the scaling of the water-level data was made for each location individually, to better understand the local maximums and minimums in each location, and bring out phasing patterns between the six data streams. The scaling down of time by a factor of 1000 makes these patterns clear enough to be perceived, while still slow enough to appreciate slight phasing differences from one high-tide to the next.

Attempt #2: substractive synthesis sonification

Given the difficulties encountered with the pitch mapping described earlier, my second instinct was to keep the pitch fixed, choose a different one for each location to differentiate them more easily, and vary another aspect of the sound, for instance volume [extr. 2] [9]. As we have seen, Neuhoff, Kramer, & Wayand (2002) consider loudness to be a poor way to convey information, and indeed this seems to be proven by my own exploration of it. I was looking at the use of subtractive-synthesis, cutting out six partials out of a noise signal. The partials' pitches are derived from the same justly-intoned harmonic relationships as in the final version of the work. I have outlined my reasons — and will describe them in more detail — for choosing those particular pitches, but the fact that they blend so well together means that limiting the sonification to a loudness change of individual partials results in a perceived colour change, more than a real harmonic transformation of the soundscape. The partial corresponding to each location gets louder as the water level data for it rises, but this is quite difficult to pick up.

Ecological and artistic context

Having looked at concerns to do with the intelligibility of the sonification, I would now like to focus on the thoughts and decisions I have had and made over the creation of this work that relate to the ecological context of the data used and the field of acoustic ecology, as well as to more artistic, aesthetic and semantic concerns.

Although we have addressed questions relating to what data-to-sound relationships made sense from a mapping perspective — looking at sound dimensions, polarity or scaling —, the question of which sound would strike the audience as the best representation of the data sonified remains a much more difficult one to solve. It is inherently difficult to find a universally (temporally, geographically and culturally) understood representation, when "All acoustic

---

[9] see ./extracts/2_volume_reson.wav

symbolism, even that associated with archetypes, is slowly but steadily undergoing modification" (Schaffer, 1977). However, it is interesting to look at the source of the data itself, before we try to find the best way to symbolise it, and I therefore believe that an ecological look on the sonification is particularly fitting. Moreover, I do believe there are many interesting crossovers between soundscape composition (listening to natural sounds as composed sound) and data sonification (listening to the sound rendition of data as composed sound).

Truax (2008) gives a list of principles of soundscape composition: "(a) listener recognisability of the source material is maintained, even if it subsequently undergoes transformation; (b) the listener's knowledge of the environmental and psychological context of the soundscape material is invoked and encouraged to complete the network of meanings ascribed to the music; (c) the composer's knowledge of the environmental and psychological context of the soundscape material is allowed to influence the shape of the composition at every level, and ultimately the composition is inseparable from some or all of those aspects of reality; and ideally, (d) the work enhances our understanding of the world, and its influence carries over into everyday perceptual habits". I believe that all of these principles, without exception, can be applied to sonification, and become an interesting prism to experience sonified works through, if one replaces source with data. The prerequisite that "the original sounds must stay recognisable and the listener's contextual and symbolic associations should be invoked for a piece to be a soundscape composition" (Truax, 2008) is very close to my goal of making the data intelligible, and appealing to the listener's contextual knowledge of the data source.

Therefore, I used this conception of sound as a premise when working on this sonification, and hoped to create a work that would invoke a landscape related to the data-source. I have made attempts to sonify the data that used field recordings of seaside environments, precisely to facilitate the emergence of this perceived landscape, and to take advantage of the fact that "The sounds of the environment have referential meanings. For the soundscape researcher they are not merely abstract acoustical events, but must be investigated as acoustic signs, signals and symbols" (Schaffer, 1977).

Attempts #3;4;5: ecological approaches

My first attempt to feature a more ecologically-minded approach, in an effort to link the semantic associations of the sounds used to the object of the sonification, hoping this would provide added information and clarity for the listener, once again looked at a very simple volume mapping. This time however, I tried to change the volume of semantically-significant field recordings, one for each location, picked and recorded specifically because of cultural and geographical features of each place. I looked at what industry were prevalent in each area: industrial docks, sailing, thalassotherapy tourism, etc.; what were the defining geographical and urban planning features of each place: beach coast line, concrete industrial ports, smaller or older ports; more raw natural landscapes; etc. and selected recordings accordingly [extr. 3a] [10]. Again, this was not particularly clear, as it was difficult to know which field recording was which in the mix.

I then added this and the pitched material from earlier together [extr. 3b] [11], and crossfaded between the two: the higher the water-level, the louder the pitch material and the softer the field recordings, and vice versa. Again however, this was not particularly effective; it was difficult to identify one field recording as being associated to a particular data-stream, and both recordings and pitched-materials were constantly present making it difficult to clearly perceive changes in the data.

In an attempt to better blend the two parts of this approach (pitch and noise), I made use of subtractive synthesis on the recordings themselves; instead of cutting out pitches from a noise signal, and superimposing the resulting pitches onto field recordings, I directly filtered pitches out of the field recordings. The higher the water-level, the narrower the Q-value (the bandwidth of frequency cut out from the recordings), and therefore the clearer the pitch [extr. 4] [12].

I then made a similar attempt, using a slightly different process: instead of filtering pitches out of the recordings, I used convolution (multiplying the frequency spectra of two audio sources, the input signal and the impulse response), using the recordings as the input, and recordings of guitar harmonics with the same pitches used before as the impulse response. I then changed the dry/wet ratio, to once again make the overall signal more pitched as the water-level value rose [extr. 5] [13].

Both of these last methods were interesting sonically, but disappointing to keep track of the data. Once again, because the fully-pitched sound was mapped to the highest value in each dataset, which is only reached once, the resulting soundscape is a rather unintelligible blend

---

[10] see ./extracts/3a_volume_samples_stereo.wav

[11] see ./extracts/3b_pitch_noise_crossfade_stereo.wav

[12] see ./extracts/4_Reson_sub_synthesis.wav

[13] see ./extracts/5_convolver.wav

of constant noise and pitch, where changes of balance between both are hardly discernible. Even using an exponential mapping did not resolve the issue. I also doubled this convolution approach with the subtractive synthesis on a noise signal mentioned earlier, with the hope that interpreting the same streams in two ways would allow me to reiterate and clarify the sound-output [extr. 6] [14]. I do not believe this made it clear enough either, however.

Limits of a direct-ecological approach and benefits of ecological concepts in a connotative approach

The direct-ecological approach, using recorded material tied to the object of the sonification, greatly limits the sound-processing possibilities of the project, as I have no choice but to use these very specific sounds, even when they might not be the greatest mediums to convey information (recordings of the sea are for instance quite close acoustically to noise signals, which greatly limits their potential as information-conveyers). With a connotative approach, on the other hand, one can select the sounds that will enable them to create the clearest sound output possible, even if ecological implications are not as immediately perceivable. I have however really considered how I could induce strong metaphorical links between the sound output of a symbolic sonification and its object.

My approaches using sounds that were quite literally ecologically-linked to the source were somewhat unsuccessful, but the final version of the sonification was still built with multiple ideas from acoustic ecology in mind. Acoustic ecologists are by no means the only ones to have used concepts such as environment and landscape when listening to music or sound-art. This ambition to conceive of 'environment as music', of "listening to the environment (either where one is situated or via a recording) as if it were music – that is, with the same level of attentiveness one might ascribe to musical listening" (Truax, 2015) goes hand in hand with the opposite initiative, to listen to "the landscape of a sound-image as the imagined source of the perceived sounds" (Wishart, 1996). With the latter, every sound has the potential of creating an imagined or perceived environment: "The landscape of the sounds heard at an orchestral concert is 'musicians-playing-instruments'. The landscape of the same concert heard over loudspeakers is also 'musicians playing instruments'" (ibid).

---

[14] see ./extracts/6_reson_and_convolver.wav

In an eternal need to balance clarity and aesthetic, I have settled on a more symbolic, granulation-based approach, using acoustic instrumental sounds, rather than on a direct ecological one, making use of field-recordings or natural sounds. I am not, however, evoking the instrumental landscape we mentioned, in fact I am to some extent annihilating it — by using instrumental sounds that are obviously not played by an instrumentalist; “recognisable sounds or their relationships point outside themselves to ideas and relationships which do not reside essentially in the aural landscape.” (Wishart, 1996). Interpreting a sound as belonging to a certain environment goes beyond a simple connection between the sound and an imagined source: “The sound-object is to be analysed for its intrinsic acoustic properties and not in relation to the instrument or physical cause which brought it into being” (ibid). Hence, my goal is more to create a metaphorical landscape than a literal one, to evoke ideas linked to philosophical perceptions of the sea more than to use recordings to force recollections of physical experiences of it.

With this priority in mind, I have been using notions very present in acoustic ecology. One question that is particularly prominent in that field is that of perspective, of figure versus ground: “According to the gestalt psychologists, who introduced the distinction, figure is the focus of interest and ground is the setting or context.” (Schaffer, 1977). In sound, “the figure corresponds to the signal or the soundmark, the ground to the ambient sounds around it” (ibid). This is a particularly thought-provoking way to think about my work, given the hazy distinction between figure and ground in it.

This duality is paralleled in a less metaphorical, and more musical sense by the contrast between gesture and texture, where “Gesture is the name we can give to the unique event, the solo, the specific, the noticeable; texture is then the generalised aggregate, the mottled effect, the imprecise anarchy of conflicting actions.” (ibid). Again this is not so immediately and easily noticeable in my work, but that is not to say that it is not present at all. I believe that part of the strength of the granulation approach is that it brings figure out of ground, gesture out of texture. The fact that a simple change in volume and timbre can enable a change from texture to gesture, that perspective has an appropriate aural analogue in dynamics, means that this piece constantly brings parts of the background into the foreground, as grains become shorter and denser, and dynamics bulge out to make individual elements come out of the whole. This to me is a great way to illustrate a rising tide, where the changes are barely perceivable on an immediate time-scale, but undeniable on a longer one; these parallels and metaphorical associations were a major reason for me to look at this time-domain timbral transformation.

This appeal of granulation, its ability to “clearly juxtapose the micro and macro levels, as the richness of the latter lies in stark contrast to the insignificance of the former” (Truax, 1990) has been explored extensively in acoustic ecology by Truax (1990; 1994; 2008; 2015). Beyond these bulges in volume that enable a continuous movement between background and foreground, granulation enables a ‘zooming in’ effect: “the overlay of up to 20 simultaneous versions of such sound per stereo pair of tracks, each with its own variations, produces a "magnification" of the original sound” (Truax, 1994). This was again particularly appealing to me, in terms of the semantic links this zooming in has to the tide rising, gradually approaching and getting in focus, and therefore contributed to the creation of the metaphorical landscape I wanted to convey. This abstract and somewhat loose connection is made clearer by the fact that many have compared sound grains to individual droplets, accumulating to form a stream. This thought that “the enormously rich and powerful textures [granulation] produces result from its being based on the most "trivial" grains of sound suggested a metaphoric relation to the river whose power is based on the accumulation of countless "powerless" droplets of water” (Truax, 1990), might reinforce the idea of a perceived landscape for the listener, and strengthen the semantic and ecological relation between the data and the sonified output, even when the approach is extremely symbolic and connotative.

Superimposition of data-streams, auditory streaming, and masking

Having looked at concerns to do with the intelligibility of the sonification with regards to clarity of mapping, and at the ecological and artistic contexts and motivations of my work, the final practical issue left to discuss in this project concerns how I dealt with multiple simultaneous data-streams. I am using data from six locations, but from a single time frame. Hence, I have to manage the superimposition of these streams, attempting to make each clear, while also retaining their cohesiveness in the final sound-art work. This adds a layer of complexity with regards to the concerns mentioned above, as “multiple simultaneous streams of data present a unique mapping challenge” (Schlei, 2010).

It has been found that “differences in timbre and spatial location are parameters that sonification designers can often use simply and effectively” (Nees & Walker in Neuhoff, 2011), and with this piece in particular, spatialisation of each stream, in a way to approximate the six locations’ relative position, was particularly fitting. Yet, “in some cases (for some tasks), it is important to be able to perceptually separate or segregate the different [data-streams], whereas in other cases it is preferable for the two streams of data to fuse into a perceptual whole" (Walker & Nees in Neuhoff, 2011). The six locations I have chosen, although definitely independent, and relatively distant from each other, are still on one continuous coastline, all facing the Atlantic ocean. Hence, I thought that being able to focus in on one location as a discrete data-stream by spatialising the sound, while also allowing for a cohesive whole to emerge when listening all six sound-sources was an elegant solution. The piece was therefore conceived to be experienced in an ambisonic space, and more specifically the RNCM's Studio 8 [15].

Just as the decision to attempt to create a sonification that is both an engaging artwork and an interesting display of scientific information, the decision to simultaneously present individual data-streams and a coherent whole is a risky one. This is due to issues of individuality of sound sources that have been a major concern of mine over the course of this project. First introduced by Bregman, the concept of auditory streaming is a psychological effect that causes listeners to interpret audio information differently when the rate at which this information is presented to them becomes too high: "when listeners were presented with an endlessly repeating loop of tape on which were recorded a sequence of six different tones, three high ones and three low ones [and] when there was only one-tenth of a second between the onsets of consecutive tones, the listeners did not actually hear the tones in the correct order, 142536. Instead, they heard two streams of tones, one containing a repeating cycle of the three low pitched tones, 1-2-3- (where dashes indicate silences) and the other containing the three high ones (-4-5-6)." (Bregman, 1990). In order to avoid this concern, the sonification of each location's data only contains one pitch, which should prevent any major auditory streaming effect, as these six pitches are constantly present, with changes only occurring on the temporal scale, as changes in density of notes, attack/sustain times, etc.

This is however not the only psychoacoustic effect that I had to be mindful of. The pitches I am using are derived from justly-intoned intervals (see Fig. 5 for the frequencies used and their breakdown as products of simple fractions from a reference tone), and this can add to the uniformity of the overall sound: "Partials are more likely to be assembled into a single auditory image when they conform to a harmonic series. If two complex tones are related by simple integer frequency ratios, their partials are likely to align and so fuse as a single auditory image" (Huron, 2016). Again, this is part of my reason for using these tones, in order to maintain an overall coherence between data-streams, and enable the creation of a rich, consonant chord, that is the principal feature of the resulting artwork. This also motivated my use of harmonic

---

[15] see ./Spatialisation_demo.mp4 for a quick video demonstration of the project running in Studio 8

sounds: “Harmonic sounds tend to produce partials that are (1) easier for the auditory system to assemble into auditory images, (2) produce a subjectively clearer or cleaner sound, (3) sound more tonelike, and (4) produce clearer, less ambiguous pitches” (ibid). I explore this contrast between harmonic and inharmonic by putting the latter in the foreground when the density of sounds is lower (representing low-tide), and the sustain of the guitar harmonics is prominent, and focusing in on the former when the density of sounds is higher (representing high-tide), and repetitive attacks create a more ‘typical’ granular sound.

$$E_3^{+29} = 440 \cdot \frac{3}{4} \cdot \frac{64}{63} = \frac{7040}{21} \approx 335.238 \text{ Hz}$$

$$G_3^{-4} = 440 \cdot \frac{8}{9} = \frac{3520}{9} \approx 391.111 \text{ Hz}$$

$$B_3^{+31} = \frac{7040}{21} \cdot \frac{3}{2} = \frac{3520}{7} \approx 502.857 \text{ Hz}$$

$$D_4^{-2} = 440 \cdot \frac{4}{3} = \frac{1760}{3} \approx 586.667 \text{ Hz}$$

$$F\#_4^{+33} = \frac{7040}{21} \cdot \frac{9}{4} = \frac{5280}{7} \approx 754.286 \text{ Hz}$$

$$C\#_5^{+33} = \frac{5280}{7} \cdot \frac{3}{2} = \frac{7920}{7} \approx 1131.429 \text{ Hz}$$

Fig. 5 breakdown of the 6 pitches used (as guitar harmonics) within the granulation for each location

This use of similar sounds between streams, and more generally the superimposition of sound sources, comes with the risk of losing some of the acoustic information however: “In noisy environments, otherwise easily heard sounds can become difficult or impossible to detect. This interference is called masking. Masking is an auditory phenomenon, not an acoustic phenomenon; in the real world, sounds rarely obscure each other” (Huron, 2016). This effect is a common by-product of any complex musical or sound work: “When a full orchestra is playing, nearly all of the sounds produced by the various instruments (thousands of partials) are fully masked. What we hear is a small subset of partially masked sounds that manage to escape from the cacophony of acoustic activity.” (ibid)

Sounds are not all masked by every other one, however, and there are ways to limit masking. Generally, masking happens between tones that are within a *critical band* (Fletcher, 1953) (Fig. 6): “When two pure tones lie within a critical band of each other, they produce a measurable amount of mutual masking. When separated by more than a critical band, pure tones do not interfere with each other” (Huron, 2016). These bands are linked to the way the ear physically interprets pitch, and each critical band corresponds to roughly 1 millimetre separation along the basilar membrane.

Fig. 6 Approximate size of critical bands represented using musical notation

Hence in this piece, I have used pitches that are roughly a critical bandwidth apart, to limit the masking effect (see Fig. 7). Of course, this will not completely avoid it, as the sounds I use are not pure tones, and there will always be masking between partials and because of noise. Moreover, these pitches are just about a bandwidth apart, not more, as a wider chord lost the cohesiveness I am simultaneously trying to achieve. In fact, composers often space tones in a chord in a way that is coherent with critical band distances, and an even spreading of tension along the basilar membrane. Hence, this choice of tones allows a compromise to avoid some masking, but also retain a unified sound.

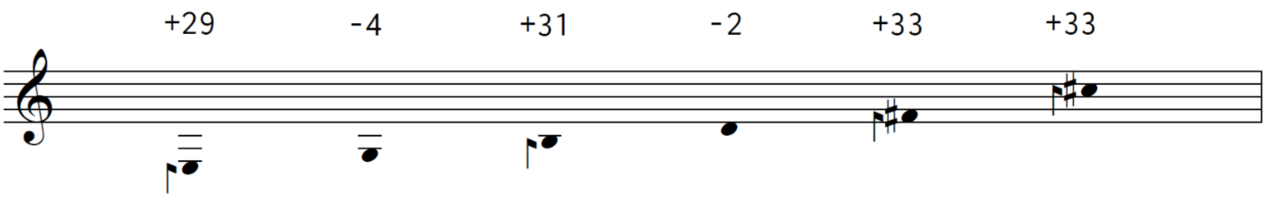

Fig. 7 the 6 pitches used within the sonification notated using musical notation

Overall, however, I am happy to create a soundscape that will inevitably lead to some masking, as I am also spatialising sources, and that is by far the best way to allow for individuality of sounds, as one can literally isolate a signal by getting closer to the source.

Visualisation as a reiterative tool

The practical output of this work includes a visualisation of the data, as well as its sonification. As we have seen with existing examples, visualisation and sonification are often coupled, in order to reiterate the information and add clarity to the conveyed information (Flowers & Hauer 1995; Flowers, Buhman & Turnage, 1997). Although this is not the focus of this dissertation, I thought it would be interesting to take it as an opportunity to create a visualisation to go along with the sound output, and further enhance the clarity of the data.

The visualisation features the six locations from which the data is derived, on a map of the portion of the west coast of France on which they are all located. Circles radiate from each location (each having a different colour, the same as the lines in Fig. 3), growing and eventually fading away, and the rate at which these circles appear and expand is dependent on the live data; the higher the tide, the quicker the circle generation and expansion.

This reiteration of the sound mapping in the same polarity (the higher the water the busier the sonification and the busier the visualisation) helps clarify how what one is hearing relates to the data, but it also helps to mentally separate the various audio streams when one is listening to the stereo version of the sonification, rather than the spatialised one.

A possible criticism of using visualisation in a sonification project — which I mentioned with regards to some examples cited in the first section — would be that it hints at the latter’s incapacity to convey the desired information on its own. However in this case, I believe the sonification is successful on its own, particularly when spatialised, and that the visualisation was an interesting research avenue to explore, to compare the effects of the sound with and without it. It is also worth noting that most of the sonifications I introduced at the beginning of this dissertation were accompanied by a visual element.

Conclusions and discussion

With this project, I have explored multiple different avenues to map the dataset I was working with onto the sound domain. Beginning with simple pitch and volume mappings, which were not so effective to perceive changes in individual data streams, to the ecologically-motivated use of field recordings, and their filtering in different ways, which led to hardly noticeable changes in the overall soundscape. It made me witness the limitations of each method, but allowed me to identify effective mapping parameters, deciding on a polarity that would associate higher water-levels with a busier texture, and using relative scaling for each stream.

I settled on a granular approach, due to the clarity of density changes in this time-domain process to represent one-dimensional changes, but also as a compromise between symbolic and analogic sonification approaches. Going back to the analogic-symbolic representation continuum (Fig.1), my approach might be placed on the connotative side, yet when listening to the sound object resulting from the sonification, my hope is that the perceived source of it has strong metaphorical links to the data it is representing. Moreover, the parallels between the nature of foreground and background in the auditory landscape and in the tides themselves (there is no discernible point at which foreground becomes background, yet when listening to each extreme, the difference is obvious, much like one can tell whether the tide is low or high, but the point at which this change happens is much more difficult to identify) will hopefully make this connection even clearer.

This decision was motivated by concepts linked to acoustic ecology applied to sonification, looking first at a direct-ecological approach, using field recordings, but finding that this came with practical limitations that impacted the clarity of the resulting soundscape. I therefore instead focused on the creation of an imagined landscape (Wishart, 1996) for the sound-image resulting from my sonification that would have strong metaphorical implications, induced by the processes rather than ecological recordings. While the sounds used are quite symbolic, and disconnected from the object of the sonification, I have used multiple processes and strategies to ensure an analogic link between the tidal forms and patterns and the perceived soundscape.

This choice of processes also enabled me to create a piece that is entirely focused on the data and derived from it, that hopefully brings insights into the relations between the different data streams, yet that fulfils the aesthetic ambitions I had for it, and fits with my musical vocabulary (very limited harmonic content; focus on subtle timbral transformations; static, repetitive or cyclical structures; etc.). Once again going back to the opening section of this dissertation, I hope this piece is successful both on the artistic and scientific (or analytic) sides, bringing concrete knowledge of the phenomenon sonified, framed in an interesting sound art work. My ambition was to reach the best possible compromise between aesthetic and clarity, while very aware that it could never be entirely successful on both fronts, and to make use of the 'canvas' that Lynch and Paradiso (2016) talked about. I believe I was successful in that regard.

I spent a lot of time working on making the sound output as clear as possible, avoiding well-known psychoacoustic effects that could have affected that clarity. I chose pitches that complemented this blend between background and foreground, by being just at critical bandwidths of each other, hence making them discernible as individual streams, but also forming a cohesive and unified landscape. I avoided issues of auditory streaming by keeping these pitches constant.

I spatialised the resulting sonification, yet again as a way to preserve the individuality of each data stream, and allow the listener to focus in on one, but also to listen to the whole room, and experience the whole landscape.

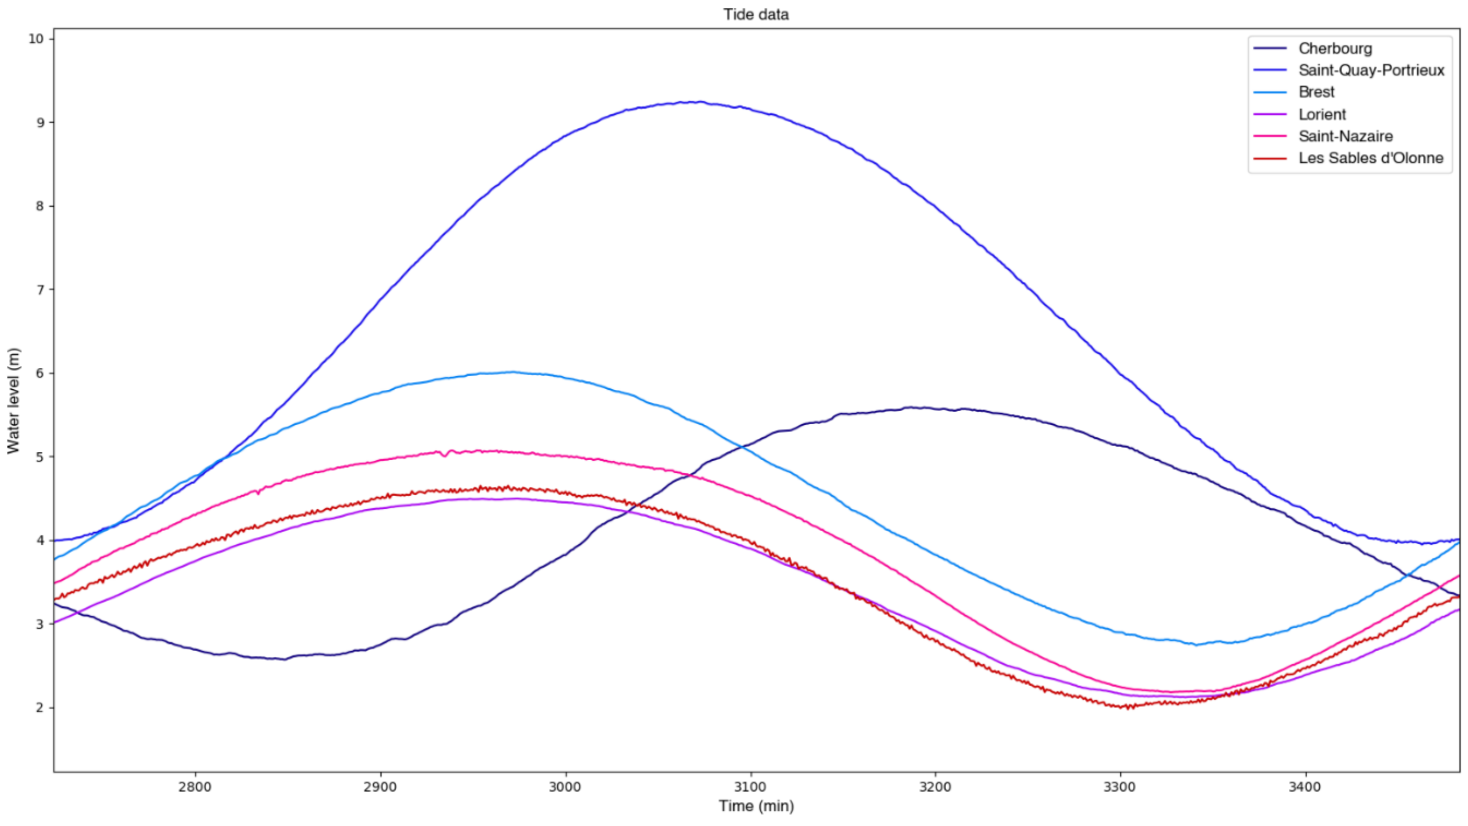

Fig. 8 Single period of the data

I believe the resulting sonification makes the phasing patterns between locations quite clear, that the asynchronicity between the first two locations (Cherbourg and Saint-Quay-Portrieux) and the remaining four (Brest, Lorient, Saint-Nazaire, Les Sables d'Olonne), as well as the asynchronicity between Cherbourg and Saint-Quay-Portrieux themselves, come across well, as the sound in the room is panned back and forth between the speakers corresponding to each location (see Fig. 8 for a single period of the data and an illustration of these asynchronicities). The periodic nature of the data is clear too, and so is the upward trend over the whole dataset.

As I mentioned when originally describing the data, tidal patterns can almost be described as sine-waves on shorter time spans, but over longer periods of time they are modulated by longer sine waves. This causes neap tides (more moderate tides, when the sun and moon are at right angles to each other relative to the Earth) and spring tides (stronger tides, when the sun and moon are aligned with each other relative to the Earth), each happening twice a month. Hence the upward trend over the dataset indicates that the week started in a neap tide and ended in a spring tide. This project was a first attempt at the sonification of tidal data, and therefore this shorter span of the data made sense, but I would be interested to sonify longer datasets, and explore these more complex modulation patterns within them. Artistically, this would also enable me to create longer-form installation works, which I am interested in exploring in the rest of my compositional work.

A further project could also involve data from more locations, exploring phasing patterns in more detail, and uncovering macro-structures along a vast coastal area, or micro-structures looking at many locations in close proximity. This would be particularly fitting if coupled with a longer time-scale. I would also be interested to delve once more into questions of ecological implications of the sounds used, but the project would have to be a longer one for me to explore these issues in a way to overcome the shortcomings mentioned here; I would want to record materials in each location, and curate them to not simply have generic field recordings loosely related to each city, but tailored shorter-span sounds made on site.

Finally, I briefly considered looking at meteorological data to drive sound ideas parallel to the main tidal sonification, but this would have hindered the focused, almost minimalist approach I had set out. A future project could couple this tide data with other, related data streams (geographical, meteorological, etc.), in order to build a much larger database, and really engage with "the rate at which we collect and generate data and information [that] vastly outpaces our ability to process and interpret it" (Lynch & Paradiso, 2016), creating a similarly vast sonification work, too vast for this project. Overall, I believe this was a successful exploration of both the theoretical and practical sides of sonification, and that the technical, aesthetic and contextual considerations I encountered in this work will act as a door to many future sonification projects, on various scales and topics.

Word count: 9920

Bibliography

::vtol:: & Alekhina, A. (2015) r x2. [Online]
Available at:
https://vtol.cc/filter/works/r-x2

Alexander, C. (2011) mta.me [Online]
Available at:
http://mta.me/

Alhoff, J. (2022) Sonification of Twitter data with hashtag #photography. [Online]
Available at:
https://soundcloud.com/joachim-allhoff/sonification-of-twitter-data-with-hashtag-photography

Asonitis, T., Allmendinger, R., Benatan, M., Climent, R. (2022) SonOpt: Sonifying bi-objective population-based optimization algorithms.

Barrass, S. & Kramer, G. (1999) Using sonification. Multimedia Systems, 7(1):23–31.

BBC Visual and Data Journalism team (2020) Coronavirus: How can we imagine the scale of Covid's death toll? BBC News. BBC. [Online]
Available at:
https://www.bbc.co.uk/news/resources/idt-7464500a-6368-4029-aa41-ab94e0ee09fb

Bellona, J.; Park, J.; Reagan, J. (2016) Aqua•litative. [Online]
Available at:
https://aqualitative.org/

Bencina, R.; Wilde, D.; Langley, S. (2008) Gesture=Sound experiments: process and mappings. Proceedings of the International Conference on New Interfaces for Musical Expression, Genoa, Italy, 197–202.

Berrick, S.; Cunningham, R.; Griswold, E. (2022) Sounding Tides.
Available at:
https://www.clockedout.org/soundingtides

Blackwell, T., & Young, M. (2004). Self-organised music. Organised Sound, 9(02).

Blattner, M.; Sumikawa, D.; Greenberg, R. Earcons and icons: Their structure and common design principles. Human Computer Interaction, 4(1):11–44, 1989.

Bregman, A. S. (1990). Auditory Scene Analysis: The Perceptual Organization of Sound. Cambridge, MA: MIT Press.

Briggs, J. & Corey, M. (2015) The Oklahoma shakes. [Online]
Available at:
https://revealnews.org/article/listen-to-the-music-of-seismic-activity-in-oklahoma/

Bultitude, R. (2017) The conditional orchestra. [Online]
Available at:
https://theconditionalorchestra.com/

de Campo, A. (2007). Toward a data sonification design space map. Proceedings of the International Conference on Auditory Display (ICAD2007) (pp. 342–347), Montreal, Canada.

Carter, E. (2013) Barographic. [Online]
Available at:
https://edcarter.net/portfolio/barographic-sage-gateshead/

Chelidon Frame (2020) Covid-19 data sonification. [Online]
Available at:
https://covid19-data-sonification.space/

Cherouvim, G. (2017) Geolocation.
Available at :
https://vimeo.com/218106546

Edworthy, J. (1998). Does sound help us to work better with machines? A commentary on Rautenberg's paper 'About the importance of auditory alarms during the operation of a plant simulator'. Interacting with Computers, 10, 401–409.

Fletcher, H. (1953). Speech and hearing in communication. New York: Van Nostrand.

Flowers, J. H., & Hauer, T. A. (1995). Musical versus visual graphs: Cross-modal equivalence in perception of time series data. Human Factors, 37(3), 553–569.

Flowers, J. H., Buhman, D. C., & Turnage, K. D. (1997). Cross-modal equivalence of visual and auditory scatterplots for exploring bivariate data samples. Human Factors, 39(3), 341–351.

Foo, B. (2015) Two trains. [Online]
Available at:
https://vimeo.com/118358642

Gaver, W. (1989) The SonicFinder: An Interface that Uses Auditory Icons, Human-Computer Interaction, 4, pp. 67-94.

Hermann, T. (2002). Sonification for exploratory data analysis. PhD thesis, Faculty of Technology, Bielefeld University.

Hermann, T. (2008). Taxonomy and definitions for sonification and auditory display. Proceedings of the 14th International Conference on Auditory Display, Paris, France.

Hermann, T. (2011). *Sonification – A Definition.* [Online]
Available at :
http://sonification.de/son/definition

Holtzman, B. & The Seismic Sound Lab (2006) SeismoDome. [Online]
Available at:
https://www.vice.com/en/article/bmy3e4/listen-to-seismic-earthquake-waves

House, B. (2012) Quotidian record [Online]
Available at:
https://brianhouse.net/works/quotidian_record/

Hunt, A.; Wanderley, M.M.; Paradis, M. (2002) The importance of parameter-mapping in electronic instrument design. Proceedings of the International Conference on New Interfaces for Musical Expression, Dublin, Ireland, 88–93.

Huron, D. (2016) Voice leading: the science behind a musical art. Cambridge, Massachusetts: MIT Press.

Jeon, M., Davison, B., Nees, M. A., Wilson, J., & Walker, B. N. (2009). Enhanced auditory menu cues improve dual task performance and are preferred with in-vehicle technologies. Proceedings of the First International Conference on Automotive User Interfaces and Interactive Vehicular Applications (Automotive UI 2009), Essen, Germany.

Jeon, M., Walker, B. N. (2011). Spindex (speech index) improves auditory menu acceptance and navigation performance. ACM Transactions on Accessible Computing, 3, Article 10.

Kaeser, A. (2018) Weather thingy. [Online]
Available at:
https://ecal.ch/en/feed/projects/5527/weather-thingy/

Keller, P., & Stevens, C. (2004). Meaning from environmental sounds: Types of signal-referent relations and their effect on recognizing auditory icons. *Journal of Experimental Psychology: Applied*, 10(1), 3–12.

Kontogeorgakopoulos, A. & Burgarth, D. (2014) Sonification of controlled quantum dynamics. Proceedings of ICMC|SCM|2014.

Kramer, G., Walker, B. N., Bonebright, T., Cook, P., Flowers, J., Miner, N., et al. (1999). The Sonification Report: Status of the Field and Research Agenda. Report prepared for the National Science Foundation by members of the International Community for Auditory Display. Santa Fe, NM: International Community for Auditory Display (ICAD).

Kyriakides, Y. (2013) Music for viola.

Lorenzo, L. (2010) 3000 hours of the astronomical tide in Venice, Italy arranged for marimba. Stony Brook University, Music Department. [Online]
Available at:
https://www.math.stonybrook.edu/~tony/tides/venice/venice3000.html

Lynch, E.; Paradiso, J. (2016) SensorChimes: musical mapping for sensor networks. Proceedings of the International Conference on New Interfaces for Musical Expression, Brisbane, Australia, 137–142.

Marchand, T. (2020) Le Mont Saint-Michel et sa baie, au rythme des marées. Planet Terre, Ecole Normale Supérieure de Lyon. [Online]
Available at:
https://planet-terre.ens-lyon.fr/ressource/MSM-maree-mascaret.xml#:~:text=Le%20Mont%2DSaint%2DMichel%2C,%C3%A0%20plus%20de%2014%20m%C3%A8tres

McCorkle, M. (2018) Weather Machine. [Online]
Available at:
https://weathermachine.info/

Miebach, N. (2009-2019) The Weather Score Project.
Available at:
https://www.nathaliemiebach.com/work/changing-weather

Myrbeck, S. & Shisko, E. (2012) Good fences make good neighbours. [Online]
Available at:
https://shanemyrbeck.com/portfolio/good-fences-make-good-neighborhoods/

NASA (2019) Sonification of a Hubble deep space image. [Online]
Available at:
https://www.youtube.com/watch?v=H-Ci_YwfH04&t=32s&ab_channel=NASA

NASA (2022) 5,000 Exoplanets: Listen to the Sounds of Discovery (NASA Data Sonification). [Online]
Available at:
https://exoplanets.nasa.gov/resources/2322/video-5000-exoplanets-listen-to-the-sounds-of-discovery-nasa-data-sonification/

Nees, A., Walker, B. (2011) 'Theory of sonification' in Neuhoff, J. (ed.) The sonification handbook. Berlin: Logos, pp. 9–40.

Nesbitt, K. V., and S. Barrass. 2002. Evaluation of a multimodal sonification and visualisation of depth of market stock data. In Proceedings of the 2002 International Conference on Auditory Display, Kyoto, Japan, July 2–5.

Neuhoff, J. G., Kramer, G., & Wayand, J. (2002). Pitch and loudness interact in auditory displays: Can the data get lost in the map? Journal of Experimental Psychology: Applied, 8(1), 17–25.

de las Pozas, V. (2020) *Semi-automated mappings for object-manipulating gestural control of electronic music*. Proceedings of the International Conference on New Interfaces for Musical Expression, Birmingham, UK, 631–634.

Quintronics (2014) weather for the blind. [Online]
Available at:
http://weatherfortheblind.org/

Rebelo, P. (2020) COVID-19 Data Sonification #65. [Online]
Available at:
https://vimeo.com/401314240?embedded=true&source=vimeo_logo&owner=958720

Renard, B. & Le Bescond, C. (2023) Finding hidden patterns in the data using principal components analysis. [Online]
Available at:
https://globxblog.inrae.fr/pca/

Reeves, G. (2014) AERO [Online]
Available at:
https://gregoryreeves.bandcamp.com/album/aero

Roads, C. (2001) Microsound. Cambridge, Massachusetts: The MIT Press.

Ruten, D. (2017) Sonic word clouds. [Online]
Available at:
https://danielruten.wordpress.com/2017/04/15/sonic-word-clouds-an-experiment-with-data-sonification-part-i-introduction/

Saranti, A., Eckel, G., & Pirrò, D. (2010). Quantum Harmonic Oscillator Sonification. Lecture Notes in Computer Science, 184–201.

Schaffer, R.M. (1977) The soundscape: our sonic environment and the tuning of the world. Rochester, Vermont: Destiny books.

Schlei, K. (2010) Relationship-based instrument mapping of multi-point data streams using a trackpad interface. Proceedings of the International Conference on New Interfaces for Musical Expression, Sydney, Australia, 136–139.

Shom (2023) Données marégraphiques (REFMAR)
Available at:
https://data.shom.fr/

Smith S., Levkowitz H., Pickett R.M., Torpey M. (1994) A System for Psychometric Testing of Auditory Representations of Scientific Data. In: Kramer G, Smith S (eds) Proceedings of the Second International Conference on Auditory Display ICAD'94, Santa Fe, New Mexico.

Smith, D. R. & Walker, B. N. (2005). Effects of auditory context cues and training on performance of a point estimation sonification task. Applied Cognitive Psychology, 19(8), 1065–1087.

Truax, B. (1990) Composing with real-time granular sound. Perspectives of New Music, 28(2), 120-134.

Truax, B. (1994) Discovering inner complexity: time shifting and transposition with a real-time granulation technique. Computer Music Journal, 18(2), 38–48.

Truax, B. (2008) Soundscape composition as global music: electroacoustic music as soundscape. Organised Sound, 13(2), 103-109.

Truax, B. (2015) Paradigm shifts and electroacoustic music: some personal reflections. Organised Sound, 20(1), 105-110.

System sounds (2018) SYSTEM Sounds. [Online]
Available at:
https://www.system-sounds.com/

Walker, B. N. (2002). Magnitude estimation of conceptual data dimensions for use in sonification. *Journal of Experimental Psychology: Applied*, 8, 211–221.

Walker, B. N., Nance, A., & Lindsay, J. (2006). Spearcons: Speech-based earcons improve navigation performance in auditory menus. Proceedings of the International Conference on Auditory Display (ICAD06) (pp. 63–68), London, UK.

West, T.J.; Wanderley, M.; Caramiaux, B. (2020) *Making mappings: examining the design process*. Proceedings of the International Conference on New Interfaces for Musical Expression (NIME-20), Birmingham, UK, 291–296.

Winters, R. M., Kalra, A., Walker, B. N. (2019) Hearing artificial intelligence: Sonification guidelines & results from a case-study in melanoma diagnosis. Proceedings of the 25th International Conference on Auditory Display (ICAD 2019). pp 23-27.

Wishart, T. (1996) *On Sonic Art.* Amsterdam: Overseas Publishers Association.

Appendix: note on the practical component of this dissertation

I am attaching the practical output of this project, described throughout this dissertation, alongside it. This comes in the form of a stereo video rendition of the sonification and visualisation piece (*Tidal forms*) and of the code (Python script, Max patch, and Processing sketch) written to create them.

I am also including the raw data that I sonified needed to run the code (in JSON format); the extracts of previous sonification attempts I referred to throughout this text; a README file explaining how to run the code (in Markdown format); and a short video demonstrating the spatialisation of the work in the RNCM's Studio 8.

The files included are as follows:

> [data]	

> Brest_2023.json
    > Cherbourg_2023.json
    > Lorient_2023.json
    > SablesdOlonne_2023.json
    > Saint-Nazaire_2023.json
    > Saint-Quay-Portrieux_2023.json
	Raw data

> [extracts]	

> 1a_oscillator.wav
    > 1b_oscillator_stereo.wav
    > 1c_oscillator_stereo_same_scale.wav
    > 2_volume_reson.wav
    > 3a_volume_samples_stereo.wav
    > 3b_pitch_noise_crossfade_stereo.wav
    > 4_Reson_sub_synthesis.wav
    > 5_convolver.wav
    > 6_reson_and_convolver.wav
	Extracts from previous sonification attempts

> [visualization]	
> France.svg > visualisation.pde	Visualisation sketch and required companion file
> patch.maxpat	Stereo Max patch (for general use)
> patch_spat.maxpat	Spatialised Max patch (for Studio 8)
> README.md	File detailing how to run the code
> script.py	Python script
> Spatialisation_demo.mp4	Short extract demonstrating the spatialisation of the work in Studio 8
> Tidal_forms.mp4	Sonification and visualisation rendition (main practical output)