
Sonification, Science and Popular Music: In search of the ‘wow’

MARK BALLORA

The Pennsylvania State University, College of Arts & Architecture, School of Music/School of Theatre, Music Building I, University Park, PA 16802, USA
E-mail: ballora@psu.edu

Sonification is described as an under-utilised dimension of the ‘wow!’ factor in science engagement multi-media. It is suggested that sonification’s potential value, like much of the scientific visualisation content, probably lies less in hard facts and more in how it may serve as a stimulant for curiosity. Sound is described as a multi-dimensional phenomenon, and a number of approaches to creating sonifications are reviewed. Design strategies are described for five types of phenomena that were sonified for works created by cosmologist George Smoot III and percussionist/ethnomusicologist Mickey Hart, most particularly for their film *Rhythms of the Universe* (Hart and Smoot 2013).

1. SEARCHING FOR THE ‘WOW!’

Sonification is currently an under-utilised element of the ‘wow!’ factor of science. Such is not the case for visualisation. Popular science programmes, such as the Public Broadcasting Service’s *NOVA* in the USA, benefit greatly from their eye-popping graphics, which stimulate the imagination with accessible, intuitive explanations of the subject matter. Far from being disregarded as mere eye candy, entertaining presentations are increasingly valued by the scientific community. The ‘wow!’ factor is acknowledged as a critical component for engaging audiences, be they colleagues, young future scientists or potential present-day funding sources.

The benefits of ‘wow!’ are largely intangible. The Pittsburgh Children’s Museum’s exterior façade is the sculpture *Articulated Cloud*, created by Ned Kahn in 2004. It consists of a grid made from hundreds of small white canvas squares, all suspended along the top and loose on the bottom. As the wind blows, the strips flap up and down, creating visible patterns along the surface of the building. I would suggest that the value of this sculpture does not lie in any concrete facts that it might impart. Rather, its benefits lie in how it encourages viewers to see their world differently. Translating the motion of wind into an aesthetic visual display stimulates a healthy, holistic curiosity about a variety of topics, including weather, waves, gravity, sculpture, textile properties and so on.

Most of the research in sonification to date has focused on concrete examples of utility in various contexts.

This focus is quite sensible: if sonifications are to be universally adopted as a means of information display, a high priority should be to show that they can solve well-defined problems.

In many ways, the adoption of sonification seems inevitable. In the information sciences, it is commonly acknowledged that information sources have multiplied faster than our ability to effectively make sense of them. The result is digital landfills, as much of the data from our information age is rarely, if ever, fully studied and understood. As information sources multiply, it seems implausible to expect that the eyes alone can make sense of everything. After all, the eyes and ears play complementary roles in our awareness of everyday life; it seems reasonable to suppose that using multiple senses would be similarly advantageous in order to sort meaningfully through large amounts of data or information.

Yet use of sonification has been slow to gain acceptance. This is probably due to a variety of factors. Historically, printed information has only been transmittable via images, and a standardised vocabulary has been in place for some time. The line graph and bar chart were created in 1786, the pie chart in 1801, and visualisation techniques had been used even earlier as aids to astronomy, mapmaking and navigation (Friendly 2008). We are educated to think in terms of these images, as math and science classes train us from early grade levels to study and create well-established types of graphs. Visualisations are a vital part of everyday life, routinely used for tasks such as navigating the streets with the aid of signs, or checking yesterday’s stock market performance in the morning paper.

Hearing is just too important a sense to disregard when dealing with the task of rendering complex information. It is well documented that the auditory sense is quite sensitive to small changes in periodicity and timing. The auditory system also can more easily follow multiple sources of information than the visual system (Bregman 1990). An even deeper mystery lies in sound’s peculiar power in unlocking sensations well beyond acoustic information, as when a former favourite song, heard years later, evokes emotions

and memories quite strikingly. Listening is a key to emergent experiential properties. Vision does not have this power to the same degree – photographs bring back memories, but typically not with the same emotional embrace. Smell can be similarly evocative, but it is not one of the primary senses for humans, who are audiovisual by nature. (This audiovisual preference is not cultural/societal, but species-based, as indicated by linguistics. While most species rely on taste and smell to a high degree, humans rely mostly on vision and sound: in various languages used worldwide, 75 per cent of words describing sensory impressions describe vision and hearing (Wilson 1998: 165)).

Since the early 2000s, sonification has become a familiar term, even if researchers tend to respond to it with some caution, not quite fully convinced of its effectiveness. I have heard many report that they were put off by preliminary-sounding sonifications, which, like a weak first handshake, left a lasting disinclination to get further acquainted. Papers on sonification are often less than conclusive, summarising the responses of small focus groups to rather simple listening tests. Like the seventeenth-century scientists who felt it was inevitable that some connection would be found between magnetism and electricity, many hope for a sonification ‘killer app’ that makes its use commonplace.

Perhaps sonification’s inevitable acceptance will not come in the form of a sudden quantum leap. Perhaps the effect will be more gradual, as compelling renderings stimulate the thinking of new generations of students, scientists and researchers. Many artists are turning to science for inspiration, including musicians and sound artists. It seems that announcements of concerts and installations informed by scientific data are appearing with ever-greater frequency.

The work described here resulted from an alliance between an astrophysicist and a musician, both of whom are prominent in their fields. The design challenges have been compelling, and my hope is that, even if these sonification examples don’t represent any concrete scientific breakthroughs, they still can provide some food for thought as design models.

To the extent that terminology matters, some might fairly point out that work of this type is more accurately called ‘data music’ than ‘sonification’. While that may be true enough, the distinction is essentially beside the point. Whether the objective in creating a sonification is to better analyse a problem, to make music or to explore auditory perception, success in one area will be likely to spread to the others. People who tap their feet to data music are probably those who will make discoveries using sonifications.

2. SOUND DIMENSIONS

Sonification is a multi-disciplinary pursuit. Effective design draws on musicality, musical acoustics, sound synthesis and human perceptual capacities – particularly in the area of gestalt pattern recognition and auditory streaming (Bregman 1990). Incidentally, I am grateful for having been exposed to micro-tonality, which got me out of the piano-trained habit of thinking of twelve-note octaves as the single basis of pitch. This sound-based training must then be supplemented with sufficient understanding of the knowledge domain that is to be sonified.

It is helpful to think of sound events as being multi-dimensional ‘objects’, with dimensional axes that include pitch, noise content, volume, spectrum, pan position, attack time and distance. The sound qualities to which we have particular sensitivity are best suited for use as *primary cues* – ‘melodies’ of the sonification, wherein small changes are apparent even to untrained listeners. Examples include pitch, tremolo rate, rhythm and attack time. Other parameters are less effective at representing small changes – such as volume, pan position, number of harmonics or envelope shape (following the attack portion). Yet they are helpful both as differentiators of primary cues (panning two melodies to opposite channels, for instance) or as *supporting auditory cues*, which make details more apparent through combinations of acoustic qualities. For example, data points in an increasing pattern may produce a primary cue of ascending pitches, but these pitch changes may be complemented with corresponding changes in harmonic content and/or pan position. Acoustic events that change in pitch usually change in other ways as well; thus these redundancies can add intuitiveness and subtlety to the rendering.

3. SONIFICATION STRATEGIES

Sonification design is usually based in three principal methodologies (Hermann, Hunt and Neuhoff 2011):

- Audification is literal, in that it plays a dataset as an audio file: data points are scaled appropriately and wrapped in an audio file format such as WAVE or AIFF.
- Parameter-based sonification maps data values to synthesised sound parameters, and treats the dataset as something akin to a sequencer file.
- Model-based sonification creates a theoretical acoustic model that corresponds to features of the dataset, and is typically used for interactive exploration of system properties.

In addition to these, I would also list a *diagrammatic* approach, which refers not to sonifying a dataset, but

rather to expressing certain relationships, such as geometry, through sound.

While parameter-based sonification is the approach that is used most commonly, I would suggest that its design methodologies are indebted to the issues raised by model-based sonification. Of particular value are the distinctions raised in (Hermann 2002) between *everyday listening* and *musical listening*, and the various examples showing the auditory system in everyday life to be a finely tuned detector of source information. Through a complex set of frequency changes over time, we infer reliable estimates as to type of source, direction, size, speed, location and so on. Thus, the potential power of sonification lies in its capacity to harness these perceptual strengths that we possess as human listeners.

Model-based sonification also suggests a greater integration of data features within the sound of a single, complex instrument, as opposed to treating data dimensions separately within the sounds of many simple instruments. The latter was the case in much earlier parameter-based work, including my own (Ballora, Pennycook, Ivanov, Glass and Goldberger 2004). The examples that will be described here are certainly much richer than that earlier work due to concepts introduced by model-based sonifications. Use of simple physical models creates a more natural sound, and allows for a greater degree of *control intimacy* (Moore 1990) by assigning the same data stream to multiple characteristics of a synthesised instrument, some primary and some supporting. This allows the possibility for auditory gestalts to form out of emergent activity among the sound dimensions.

In conversations and presentations, I have encountered a number of questions or confusions that have come up repeatedly. Some confuse model-based sonifications with parameter-based sonifications that employ physically modelled synthesis instruments. To the extent (again) that correct descriptors of methodology matter, it is helpful to compare and make distinctions between the two. With both approaches, the data does not directly create the sound signal, but rather represents it at a symbolic level. Both involve the creation of multi-layered instruments that are rich in information, reflecting data dynamics at very subtle levels. But with parameter-based sonification, the sound model is typically based more in electroacoustic synthesis techniques and less in complex mathematics. Model-based sonifications tend to rely on differential equations of motion to describe complex vibrations in a theoretical medium such as a plate or spring mesh. Furthermore, they tend not to be created for use with time-indexed data, but rather toward interactive exploration of system properties, such as the nature of a chemical compound.

Another question that sometimes comes up concerns the integrity of parameter- and model-based sonifications: since audifications are creations of sound made directly from the data, should they not be considered more accurate? After all, the other two approaches create a symbolic ‘instrument’ that is imposed on the data, but not intrinsically derived from it. An answer can be found by recalling how one views locations with an online map. Typically there is a choice between a literal satellite view and a symbolic street view. Neither is considered to have more integrity than the other. Depending on how one needs to navigate the terrain, one of the views, symbolic or literal, will be most useful.

It is important to bear in mind that symbolising is done with all information representation, not just sonification. Many visualisations are quite symbolic and abstract. The photographs of outer space often seen in the media have been coloured, processed and filtered to effectively show various electromagnetic bands that are invisible to the eyes. The l-nu diagram used in helioseismology is a quite ingeniously contrived representation of acoustic resonances within the sun. This image is sometimes accompanied in educational packages (Newhouse 2008) by engaging statements such as ‘The sun is like a piano, and these are its keys’. Such descriptions are more fanciful than factual, and that is precisely the point. They provide orientation to a strange terrain, and are thus quite useful, even if they are not quite literal.

A limitation of audification is that the data must be listened to at an audio rate, typically 44,100 data points per second (although this can be slowed by rendering data with a variable increment rate on a buffer). Since the ear cannot discern fine detail at the millisecond level, a good deal of nuance is lost by this approach. Small variations are averaged perceptually, and the result is usually a kind of coloured noise. To be sure, this approach produces some compelling work, one example being that of Alexander, Gilbert, Landi, Simoni, Zurbuchen and Roberts (2011). But there are also instances where finer resolution is required, just as there are also instances when higher dimensionality is needed. Through the use of sonifications, different data dimensions can be mapped to various acoustic properties of a virtual instrument, creating an organic-sounding blend of characteristics.

With a parameter-based approach, the data functions not as a source of audio samples themselves, but rather as a control signal that can be applied to a sound (or sound stream). This means that the iteration rate can be adjusted, which is similar to zooming in or out on an image. The design challenge is to create a musical sound that is intuitively suitable for the data it is representing, and to modulate the sound with the data values in ways that allow for meaningful discrimination.

```

//Read in data file, save to array of floats
-data=CSVFileReader.read("BHZ-1751-1755.csv");
-bhz=Array.newClear(-data.size)
-data.do({ arg item, i; -bhz.put(i, item.at(1).asFloat) });

//to get magnitudes, square, to make all values positive, then take square
root;
-mags=-bhz**2;
-srmags=sqrt(-mags);

//Copy the array a number of times, rescaling it to derive values suitable for
various auditory parameters
-magnitudeevols=-srmags.normalize(0.2, 0.9);
-rates=-bhz.normalize(0.75, 1.3);
-cutoffs=-bhz.normalize(1000, 4000);
-cutoff2s=-bhzmg-normalize(50, 150);
-freqs=-bhz.normalize(-1.0, 1.0);
-pans=-bhz.normalize(-0.8, 0.8);

//Store an audio file of tongue clicking to a buffer
-tapbuf=Buffer.read(s, "/Users/markballora/Documents/rhythma\ of\ the\
universe\seismology\seismic\ sonifications\tapmono.aif");

//Create a SynthDef
//Sound is a brief click - a single impulse is sent through many bandpass
filters (Ringz and Resonz), and is joined by
//a recording of a tongue click. Control arguments for modulation include
filter upper and lower bandwidth frequency,
SynthDef("seismicquiro", { arg cutoff=4000, cutoff2=50, rate=1, freq=709,
pan=0, dur=0.1;
  var impulse;
  e=Env.perc(dur, 0.01, 1, 1, -4);
  impulse=Impulse.ar(0, 0, 1,
BrownNoise.ar(0.00025)+PinkNoise.ar(0.0005)+GrayNoise.ar(0.00025));
  c=
  {
  (Resonz.ar(impulse, freq*0.0606, dur*0.3, 15)
  +
  Resonz.ar(impulse, freq*0.1146, dur*0.1, 12.5)
  +
  Ringz.ar(impulse, freq*0.1438, dur, -27.dbamp)
  +
  Ringz.ar(impulse, freq*0.148, dur, -20.dbamp)
  +
  Resonz.ar(impulse, freq*0.1799, dur, 2)
  +
  Resonz.ar(impulse, freq*0.2417, dur*0.5, 25)
  +
  Ringz.ar(impulse, freq*0.3, dur, -27.dbamp)
  +
  Ringz.ar(impulse, freq*0.4851, dur, -27.dbamp)
  +
  Resonz.ar(impulse, freq*0.4922, dur*0.7, 15)
  +
  Resonz.ar(impulse, freq*0.7238, dur*0.7, 12)
  +
  Resonz.ar(impulse, freq*0.7348, dur*0.7, 12)
  +
  Ringz.ar(impulse, freq*0.9065, dur*0.5, -18.dbamp)
  +
  Ringz.ar(impulse, freq*0.976, dur*0.5, -18.dbamp)
  +
  Ringz.ar(impulse, freq, dur*0.5, -18.dbamp)
  +
  Resonz.ar(impulse, freq*1.2143, dur*0.5, 10)
  +
  Resonz.ar(impulse, freq*1.4569, dur*0.5, 20)
  +
  Resonz.ar(impulse, freq*1.6995, dur*0.7, 23)
  +
  Ringz.ar(impulse, freq*2.4254, dur*0.5, -18.dbamp)
  +
  Resonz.ar(impulse, freq*2.4287, dur*0.7, 13)
  +
  Resonz.ar(impulse, freq*2.9153, dur*0.3, 13)
  +
  Resonz.ar(impulse, freq*5.5923, dur*0.3, 13)
  +
  Resonz.ar(impulse, freq*6.5657, dur*0.9, 23)
  +
  Resonz.ar(impulse, freq*7.2908, 0.09, 13)
  +
  Resonz.ar(impulse, freq*7.5359, dur*0.9, 23)
  +
  Resonz.ar(impulse, freq*10.2045, dur*0.9, 23)
  +
  Resonz.ar(impulse, freq*12.1486, dur*0.9, 18))
  +
  PlayBuf.ar(1, -tapbuf.bufnum, rate)
  };
  x=Pan2.ar(c, pan);
  x=LFPF.ar(x, cutoff);
  x=HPPF.ar(x, cutoff2);
  Out.ar(0, x*EnvGen.ar(e, doneAction:2)
  ).add;

//Set a tempo
-duration=0.4;

//Create a Task to iterate through the data, creating instances of the SynthDef
for each data point
{
-earthquake=Task({
  1601.do({ arg j; var freq;
    freq=330*(2**(-freqs.at(j)));
    Synth("seismicquiro", [\cutoff, -cutoffs.at(j), \cutoff2,
-cutoff2s.at(j), \rate, -rates.at(j), \freq, freq*2, \pan, -pans.at(j), \dur,
-duration]);
    (-duration*0.25).wait;
  });
});
})
}

//Stop or play the Task
-earthquake.start
-earthquake.stop;

```

Figure 1. SuperCollider code to load dataset, rescale it, create a SynthDef instrument, and sonify the dataset.

A common approach of many early parameter-based sonifications was to map data values to MIDI note numbers in order to create melodies from the data. MIDI's seven bits, however, are much too coarse for many complex datasets. Mapping to MIDI typically involves binning values so that certain ranges fall within a window that corresponds to a MIDI note number. The weakness with this approach is that the boundaries are arbitrary. For example, suppose that values between 0.7 and 0.79 were mapped to MIDI note 60, and values from 0.8 to 0.89 to note 61. If the data values alternate between 0.73 and 0.74, the same pitch is heard repeatedly, but if the values alternate between 0.79 and 0.8, a half-step trill is produced. Equivalent value differences produce quite different sonic results, due only to an arbitrary boundary that does not correspond to any inherent boundary in the data. Since pitch distances far smaller than an equal-tempered half-step are readily perceptible, a much finer increment is desirable for a sonification. One approach would be to add pitch bend values to each MIDI note value for finer frequency resolution. But my preferred strategy is to calculate frequencies, plugging the data values into exponents applied to a power of two, which maps equivalent data spans to equal pitch spans. Examples of this pitch mapping are shown in Figure 1.

4. SUPERCOLLIDER

Given the wide variance in dataset formats from different fields, attempts to create general-purpose sonification software have not been widely successful (de Campo 2009). Many 'sonificationists' seem to prefer using general-purpose software sound synthesis programmes and tailoring them to fit their needs. My software of choice is SuperCollider. Besides its computational efficiency and flexibility in synthesis, its Task scheduler is well suited for controlling iteration through the dataset. While SuperCollider's Pattern classes also offer various means of iteration that are applicable to sonification, Tasks (to my experience, at least) allow more precise control of iteration and positioning within a data sequence, allowing the dataset to be 'scrubbed' with a slider. (While none of the work described here uses graphical user interfaces, Tasks, rather than Patterns, remain my habitual approach.)

When datasets are to be audified, it is easy enough to load the dataset into a Buffer object and read through it. When datasets are to be sonified, the sounds themselves are created as SynthDefs – collections of unit generators, akin to the contents of an .orc file in Csound. I often derive these by recording the sound of a particular acoustic instrument (usually percussion), viewing the recording in a spectral analyser and resynthesising something similar in a SynthDef. The Klank

unit generator, which allows an arbitrary number of filters, all with their own centre frequencies, amplitudes and ring times, is invaluable. (Anecdotally speaking, many of the more interesting instruments created with SuperCollider seem to be based on a Klank in some way.) A SynthDef should be designed with a good number of control inputs to allow numerous mappings of the data to various aspects of the sound. Even when the changes may be negligible – such as very minute changes in pan or amplitude – it costs little computationally to implement them, and it can never hurt to add layers of dimensionality.

In summary, the methodology employed in these sonifications is to:

- create SynthDefs,
- read the datasets into SuperCollider as arrays,
- rescale them appropriately,
- iterate through the arrays with a Task, with each iteration either updating properties of a sounding Synth or creating a short instance (a ‘note’) of a SynthDef.

5. ROCK ET SCIENCE

Work I have described previously explored informative sonifications of computer network activity (Ballora and Hall 2010; Ballora, Giacobe and Hall 2011; Ballora, Cole, Kreusi, Greene, Monahan and Hall 2012). The work described here was created for the twenty-one-minute film *Rhythms of the Universe*, a co-creation of former Grateful Dead percussionist and ethnomusicologist Mickey Hart and cosmologist George Smoot III of Lawrence Berkeley Labs. It is a poetic and scientific speculation on humankind’s innate desire to understand the cosmos, featuring narration by Smoot and Hart, along with a shifting collage of music, videos, visualisations and sonifications. The focus is different when designing sonifications as an engagement element rather than as a research element. While they need to be accurate and informative, their success depends primarily on their musicality.

The first version of *Rhythms of the Universe* was shown in January 2010 at the Congreso Cosmologia conference held in Playa del Carmen, Mexico, and featured audifications created by Keith R. Jackson of Lawrence Berkeley Labs. The final version of the film premiered on 29 September 2013, at the Air and Space Museum at the Smithsonian Institute in Washington, DC.

The work is a natural extension of Mickey Hart’s earlier work in ethnomusicology and ‘world’ music. It is also well aligned with the mission of the Berkeley Center for Cosmological Physics (BCCP), an outreach initiative created by George Smoot after he became a co-recipient of the Nobel prize in 2006. While a variety of producers create science engagement media,

the BCCP is unique in its emphasis on using sonification as well as visualisation in its productions.

What follows are specific descriptions of five of the sonifications created for *Rhythms of the Universe* and related projects. The balance of informativeness versus musicality varies. Some are more informative; some are more metaphoric. I chose these five as examples of various types of design approaches. Other examples can be heard online (Ballora 2011; Ballora and Smoot 2013).

Example audio files may be downloaded from <http://www.personal.psu.edu/meb26/Ballora-Organised-Sound-sonifications.zip>.

6. SCHUMANN RESONANCES

The air space between the earth’s surface and the ionosphere, some 80 km above the ground, functions like a waveguide, a giant spherical ‘flute’ for electromagnetic waves that reach us from the sun and outer space. This electromagnetic energy immerses us constantly. There is particular interest in the seven Schumann resonances, which appear at 7.8 Hz, 14 Hz, 20 Hz, 26 Hz, 33 Hz, 39 Hz and 45 Hz, although the frequencies vary plus or minus 5 per cent due to changes in the day–night cycle. The activity of these resonances aids the detection of distant lightning storms, information that is critically important to air traffic systems. Their activity also correlates with changes in the atmospheric temperature, helping scientists gain insights into climate change (Schlegel and Füllekrug 2002).

A real-time Schumann resonance sonification project is described in (First 2003). An additive synthesis approach was taken to create complex tones based on the shifting amplitudes of certain frequency ranges, which were transposed up to audible frequencies. For purposes of this film, it was sufficient to create an additive synthesis sonification derived from average spectral measurements. The seven principal resonances are played as partials of a fundamental, with the fundamental transposed up three or four octaves, and the partials kept at equivalent ratios to it (that is, 1.79, 2.56, 3.33, 4.23, 5.0, 5.76). Each of the seven resonances is synthesised as a sine oscillator and filtered noise combination, with the frequency and amplitude of both the oscillator and the filter derived from one of the Schumann frequencies, and with the bandwidths of the filtered noise matching the resonances’ bandwidths (although this is the least perceptible part of the sound).

The partials of this bell-like hum are panned evenly across the stereo field to ‘open’ the sound up. The 5 per cent variation in each frequency is simulated by slow, linear sample and hold low-frequency oscillators. These are applied to each oscillator-noise combination to create a random 5 per cent drift in both the pitch

Table 1. Creation of a scale of pitches based on relative distances of planets from the sun

	semi-major axis (km)	inverse axis	inverse/ fundamental	transposed ("normalized")	
Mercury	57910000	1.72682E-08	101.992402	1.593631281	six octave transposition (multiplication by 0.015625)
Venus	108210000	9.24129E-09	54.58257093	1.705705341	five octave transposition (multiplication by 0.03125)
Earth	149600000	6.68449E-09	39.48114973	1.233785929	five octave transposition (multiplication by 0.03125)
Mars	227920000	4.38775E-09	25.91426816	1.61964176	four octave transposition (multiplication by 0.0625)
Jupiter	778570000	1.28441E-09	7.586190066	1.896547517	two octave transposition (multiplication by 0.25)
Saturn	1433530000	6.97579E-10	4.120164908	1.030041227	two octave transposition (multiplication by 0.25)
Uranus	2872460000	3.48134E-10	2.05620966	1.02810483	one octave transposition (multiplication by 0.5)
Neptune	4495060000	2.22466E-10	1.313971337	1.313971337	no transposition
Pluto	5906380000	1.69308E-10	1	1	no transposition

The scale, in ascending order:

1	1.028	1.03	1.23	1.313	1.596	1.619	1.706	1.8965
Pluto	Uranus	Saturn	Earth	Neptune	Mercury	Mars	Venus	Jupiter

and the pan positions of the partials, giving the sound a subtly shifting quality (Sound example 1).

7. THE PLANETS OF THE SOLAR SYSTEM

It seemed appropriate for *Rhythms of the Universe* to reference Pythagoras and the Music of the Spheres. The origin of our major scale, Pythagorean intonation, was created as a diagrammatic sonification. Its ratios, related by perfect fifths, were attributed by the Pythagoreans to the relative distances between the earth and the bodies that were thought to orbit it: Sun, Moon, Mercury, Venus, Mars, Jupiter and Saturn. They would demonstrate these ratios on an instrument called the monochord, on which proportional lengths of a plucked string were compared along with their relative pitches, giving Western civilisation its early lessons in musical intervals.

The website SolarBeat (<http://www.whitevinyl.com/solarbeat>) illustrates the solar system with an animation showing accelerated orbits. Each planet has a tone that is played when it passes through the 3 o'clock position, thus illustrating the relative orbit times. I decided to try to expand on this with a diagrammatic sonification: by calculating pitches that were based on the planet's relative distances from the sun, and also including the elliptical shape of the orbits, along with their inclinations relative to earth's orbit.

NASA's JPL website provides measurements of the planets' semi-major axes around the sun, as well as their eccentricities and inclinations. Bearing in mind the Pythagorean monochord, the semi-major axes of each planet can be considered to correspond to string length, with the inverse of each axis corresponding to a relative frequency value. Intervals can be derived from the ratios between the different orbits. By transposing these ratios by successive octaves (multiplications of

either 0.5 or 2), they can be expressed as pitch classes, falling within the span of an octave – 1.0 to 2.0 (Sound example 2). The process is shown in Table 1.

The relative distances of the planets are mapped to their pitches, and the relative orbit times are mapped to their repetition rates. I used a Karplus-Strong pluck generator to play each frequency repeatedly, at rates based on relative orbit times. As each planet's orbit passes a polar coordinate of 0°, a pluck is heard at that planet's particular pitch. The earth's period is set at 1 second, and the other planets pluck at rates that match their relative orbit times: Mercury's pluck is heard 4.15 times for every single earth pluck, while Pluto's pluck is heard every 247 seconds.

The relative distances are also applied to more subtle sound parameters. SuperCollider's Pluck ugen has a coefficient argument that falls between 0 and 1.0. Lower values produce a harsh, hard plucked sound, while higher values are softer and muted. The coefficient for each planet's pluck is a mapping of its relative distance from the sun, so that the inner planets sound more 'present' while the outer planets sound more 'distant'.

Another sound layer along with the plucked string sound is a spinning whistle created from filtered noise, based at the same pitch as the pluck, which represents the orbital path. Using the axis length and eccentricity values from NASA, I plotted distances for a sequence of 360° for each orbit. The whistling sound event occurs for each position. The whistle 'notes' overlap, forming a continuous swirling sound, with pan positions that cycle through the stereo field once every 360°.

The rotation rate and pitch vary slightly according to changing distance, which simulates Doppler effects. This was done by taking the distance from the sun (one of the ellipse's foci) at each of the 360° positions and dividing this distance by the distance

value for the last degree position, then multiplying the inverse of this ratio by the planet's base frequency. Since the planets move more quickly when they are closer to the sun, the frequency of the swirling noise varies according to position, going up slightly when the planet is near the sun, and falling slightly when it is more distant. The iteration rate is also mapped from these distances, as the traversal through the 360° is slightly faster when the radius from a focus is smaller, and slightly slower when the radius increases. Most planets' orbits are close to being circular, so this is a subtle effect, but it adds a small amount of variety to the sound.

The inclination of each planet's orbit relative to the plane of the earth's orbit is reflected in a band-limited impulse generator. This unit generator is similar to Csound's *Buzz*, which plays all harmonics at equal amplitude. The pitch is a transposition by two octaves down of the planet's basic pitch. The number of harmonics is a mapping of the inclination angle, and varies at a rate that matches the planet's orbit time, so that its sound gets slightly brighter and more muffled once per orbit (Sound example 3).

8. EARTHQUAKES

Audifications of seismic phenomena are described in the earliest papers on auditory display (Hayward 1994). It is an approach that makes good intuitive sense, given that seismic activity consists of acoustic waves in the surface of the earth. Listening to them gave seismologists the ability to hear certain patterns that were difficult to detect visually. However, the limitation of this approach is that seismic waves are low in frequency, and typically sampled at rates well below those of digital music recording. As discussed earlier, audifications must compress the timescale, playing the audification much more quickly than the actual event occurred. While the overall shape of the wave may be apparent, nuances are lost in the high-speed music playback. Symbolic sonifications offer the ability to render the earth's motion in the same amount of time it took place, or slower.

An unusually strong earthquake centred in Virginia took place on 23 August 2011, causing considerable damage in the Washington, DC area. (As a result, the Washington Monument remains closed at this writing, over two years later.) Data on this quake (and other quakes) is available at the Incorporated Research Institutions for Seismology (iris.edu). The programme manager provided datasets measured in Standing Stone, PA; Albuquerque, NM; and College Outpost, AK. Seismologists take great interest in comparative earthquake measurements from different locations. At close proximity to a quake's epicentre, many frequencies occur simultaneously as a transient event. But as the waves propagate through, across and

over the earth, components travel at different speeds and arrive at different times at different places, producing longer and more complex vibrations than those felt near the source.

When I audified these three datasets, they had a rapping sound, which was hardly surprising considering the percussive nature of a seismic event. Thus, this approach was useful for creating the sound of the earth as a drum, with slightly different rapping characteristics produced by data taken at different locations (Sound examples 4–6).

I also sonified the datasets by creating a *SynthDef* that was based on the spectrum of a guiro – a hollow, fish-shaped wooden instrument with a corrugated surface, played by rubbing a wooden stick along it. The dataset, consisting of measurements in the up and down motion of the earth, was used as a control signal to change the guiro-like instrument's pitch, filter cutoff and stereo pan position. The flexible iteration approach of this kind of sonification allowed the motion to be played in actual time, and produced a much more evocative and 'shakier' sound than the audifications (Sound example 7).

The SuperCollider code showing how to read the data file, prepare it for use, define an instrument and play the dataset is shown in Figure 1.

9. THE GOLDEN GATE BRIDGE

The Exploratorium, San Francisco's science museum, has an ongoing programme about the Golden Gate Bridge, with exhibits such as a 'The Bridge as a Thermometer', wherein its small changes in height due to temperature are measured periodically with a laser and GPS. Another exhibit shows plots of the side-to-side motion of the bridge.

The Mickey Hart Band performed at a commemoration of the 75th anniversary of the Golden Gate Bridge on 27 May 2012. For the free outdoor performance, the Exploratorium (where Hart is a board member) created a 23-foot flexible model of the bridge with a segmented 'roadway' that could be shaken and twisted. The sculpture was designed and fabricated by Dave Fleming, with electrical design and interface by David Torgersen. It was outfitted with a series of flex resistors, piezo transducers, force sensing resistors and accelerometers, and was controlled by an Arduino Mega. It transmitted 18 MIDI event types, some of them are MIDI note numbers and others control change streams.

To create the sound, we used datasets provided by the California Strong Motion Authority, which has sensors in various locations on the bridge that monitor its motion. We were given displacement data, describing 40 seconds of motion during an earthquake that took place in August of 1999, for forty-six bridge locations. We were also given datasets

from twenty-five locations for 2.5 minutes of normal activity that took place in July of 2011. Each sensor's measurements were along a single axis – up–down, north–south or east–west. Some locations had multiple sensors to measure different directions of motion; other locations had only one sensor.

Audifications of the datasets produced a variety of textures. The earthquake datasets were percussive; the ambient datasets were more akin to slowly rotating a container of sand. As discussed earlier, audifying them reduced their dimensionality, since each file only represents motion along one dimension. It would be possible to combine dimensions by playing audifications on separate stereo channels, for instance, but doing this would tend to confuse any perceptual clarity that a one-dimensional audification might have had.

When the datasets were sonified, the motion in different directions could be combined into different characteristics of the same sound. I synthesised an instrument based on a Thunder Tube, a percussion instrument made by the Remo Corporation that consists of a small hollow tube, open on one end, and with a drum skin affixed to the other, and a spring coil attached to the drum skin. When the instrument is shaken, or the coil is rubbed or plucked, low, creepy warbling sounds are produced. When studying the sound, I found it to be similar in evolution to the sound of a gong, in that there were discernable 'rumble' and 'sizzle' stages. I mapped the up–down data values to the pitch, tremolo rate and tremolo amplitude. North–south values were mapped to the volume of the sizzle and the sizzle's centre frequency. The east–west data was mapped to the filter's bandwidth, the reverb level and the pan position of the sound (Sound examples 8 and 9).

My sonifications and audifications were incorporated into sound designs created by Jonah Sharp. During the performance, MIDI from the bridge's pads and touch slider were sent to Ableton Live, where various sounds were triggered by the different MIDI messages. A MIDI splitter also sent the MIDI to my computer, which was running Max/MSP and SuperCollider. The control change values from the accelerometers were read and rescaled in Max/MSP, then wrapped in OSC messages that were sent to SuperCollider, where they were applied as a control signal to a SynthDef that was based on the Thunder Tube-like definition that had been used for the sonifications. A moaning, windy sound was produced that made Hart's manipulations of the bridge audible (Hart 2012).

The model (which Hart named 'Bridget') will be housed at the Exploratorium. For the permanent exhibit, the electronics will be simplified so that the public can interact with the model and hear its motion. A single accelerometer on the span will send

MIDI pitch bend values for up–down and left–right motion. For sound, David Torgersen requested two monophonic six-minute audio files that could be controlled by pitch bend, making the motion along the two axes audible. To make these sounds, I used Logic's multi-track timeline to stitch together a montage made up of the various audification files. I used pitch shifting and EQ to differentiate between the two sounds. While the result is a similar sound to filtered noise, it has the integrity of being noise that was derived from actual motion of the real bridge. When the public interacts with Bridget, the motions will send pitch bend information to a Max/MSP patch, where the two audification-montage files will be loaded into `sfplay~` objects. The pitch bend values will be rescaled and sent into the `sfplay~` speed inlets, modulating the pitch of the two audio files.

10. GALAXIES

The spectra of various galaxies were sonified for *Rhythms of the Universe*. The datasets are publicly available online at the NASA/IPAC Extragalactic Database (NED).

A common approach to rendering two-dimensional datasets such as these is to map the y-axis intensity levels to pitch, and the x-axis values to time. The result is a melody based on the shape of the data curve. This is the approach taken with software such as xSonify (Candey, Schertenleib and Diaz Merced 2005), as well as in the musical work of Marty Quinn at the Design Rhythmics Sonification Research Lab (<http://drsrl.com>).

I took a similar approach, although my goal was to use more precise values than were obtainable with MIDI-based data. I also pursued a notion of galaxies sounding like distant wind chimes.

Different spectral regions of different files varied considerably in terms of how much variety they had. To gain some modularity, I broke each spectrum file into smaller sections that could be mixed and matched in a sonification collage.

In harmonic sound spectra, harmonic numbers that are powers of two represent octaves of the fundamental (harmonics 1, 2, 4, 8, 16, ...). I decided to subdivide these galactic spectra into 'octaves' by creating files based on index numbers that were powers of two. Of course, these weren't really octaves, since visually there's no equivalent to an octave. But conceptually this seemed like an integrated way to proceed. Data divisions based on lower octaves contained far fewer members than those at higher octaves. ('Octave' 3 had data indices 16–31, 'octave' 4 had indices 32–63, 'octave' 5 had indices 64–127, and so on. The datasets contained 8–11 'octaves'.)

Each octave was given a fundamental pitch, and all intensities of that octave's dataset were pitched as

partials of this fundamental. As the data incremented through ‘octaves’, the corresponding fundamental ascended by a Pythagorean half-step (ratio of 256:243). In an audio spectrum, each of these fundamentals would have been an octave apart. But this pitch compression was necessary because the sound would quickly have become shrill had each data ‘octave’ corresponded to a pitch octave. Thus, listenable pitch space was preserved by shortening the intervals between octaves. This meant that the pitch space of the partials overlapped somewhat between renderings of different octaves. But the difference in fundamental maintained a slightly different tonal environment for each ‘octave’.

Based on an example posted by Alan Russell to the SuperCollider mailing list in October 2010, I created a bell-like sound from harmonically related sine waves that were slightly detuned and had percussive envelopes. Iterating through each data file at a constant rate produced a sound that quickly became monotonous and irritating (I call this ‘the woodpecker effect’), and did not sound anything like a wind chime. So for further variety, I varied both the time increments between chimes and their volume levels, depending on the difference between the current and the last intensity values. Greater differences caused longer delays between chimes, as well as louder chime strikes. (Smaller datasets, which represent the lower frequency ranges, were iterated at a slower speed than the larger datasets, which represented higher frequencies.)

Thus, the contour of the spectra was the basis of both the melody and the rhythm of each sonification. Large spectral changes were quite perceptible. In addition to the main bell sound, a soft wash of filtered noise gave a subtle background indicator of the frequency region, with different filter frequencies and noise types representing ultraviolet, optical, infrared and far-infrared regions.

For subsequent sonifications of different galaxies’ spectra, I created different chime-like sounds as a basic instrument, keeping the same ‘wind chime’ approach of subdividing them into ‘octaves’ and to mapping the intensity values to pitch, volume and timings (Sound examples 10–14).

11. FINAL THOUGHTS AND ASSESSMENTS

An exposition of exploratory work such as this necessarily concludes with assessments. Here, I tread cautiously, recognising that any work in development must be assessed and improved. But I am sometimes troubled by quantified assessment procedures that are commonly applied in the education and research industries. The most valuable information is typically not expressible by numbers. Thus, the results are often highly reductive, and the process is analogous

to the refinement of wheat into white bread. As summarised by Stanley Fish in the *New York Times*, ‘We’re probably measuring the wrong things and the right things are not amenable to measurement’ (Fish 2013).

I often recall hearing Morton Subotnick describe the creation of his early compositions in the 1960s. He lived in New York City at the time, and every so often he would go to the street, introduce himself to passers-by, and invite them to his loft to hear portions of what he was working on. After people would listen, he would thank them and show them out without any further conversation. There was no need for a discussion or a questionnaire; he got the assessment he needed simply by sitting with listeners.

As is the case with the *Articulated Cloud* sculpture, music-related work defies quantitative assessment. Assessments are likely to be anecdotal. Conductor Benjamin Zander put it aptly in his 2008 TED presentation (Zander 2008) when he defined a performance’s success according to ‘how many shining eyes I have around me’.

What can be said with quantitative certainty is that Mickey Hart and George Smoot are exposing sonification to new and numerous audiences. It has become an important theme in Hart’s music. Sonifications created for *Rhythms of the Universe* are also part of his sound library, from which he draws during live performances and for recordings, such as the Mickey Hart Band’s albums *Mysterium Tremendum* (2012) and *Superorganism* (2013). While they are not always at the foreground of his complex soundscapes, he finds sonification important enough to describe from the stage in his concerts, in interviews and more fully in his website (www.mickeyhart.net/mysterium). Thus, he has made sonification a topic of conversation among the Grateful Dead audience, which is a far-reaching cultural phenomenon that has been acknowledged by many artists, sociologists, musicologists, financiers and politicians.

The premier of *Rhythms of the Universe* generated a good deal of interest. (The event may be seen online at the URL included in the reference.) Far from being formal and perfunctory, the preliminary remarks and question-and-answer period after the film were lively and informative. It was gratifying to hear sonification discussed at the Smithsonian.

The sonifications intended for this piece have expanded beyond this particular project and overlapped with others. Hart is also actively working with scientists at University of California San Francisco and the Gladstone Institute on neurological signals, rhythms and the ‘sounds’ of the brain. George Smoot is involved in creating a planetarium show on dark matter that will feature visualisations and sonifications of astronomical phenomena, as well as of particle collisions in the Large Hadron Collider at

the European Organisation for Nuclear Research (CERN) in Geneva. The goal is for the show to be made available at little-to-no cost to planetariums worldwide. Since nearly 110 million people attend planetarium shows annually, this is an extremely promising avenue for science education and outreach.

Mickey Hart states categorically that his work is not a science experiment, but music (that is hopefully enjoyable) based on universal rhythms. By the same token, I am not a scientist, but, rather, a translator. Although the sonifications are designed to have the capacity to provide new scientific insights, we have not yet put effort into studying these datasets critically with sound. While we cannot state with objective certainty that this work will shape the next generation of scientists, the number of shining eyes we have seen in audiences makes us optimistic that we are headed down a healthy path.

Like much work in the arts, the work described here is a leap of faith. We are not doing it because of a pre-defined, quantitative agenda; we are doing it because, at a visceral level, we feel compelled to. Besides the musical uses, we have confidence that it will yield insights beyond what we can predict at this point. The long-term, underlying goal is to pursue an integrated approach to learning, wherein emergent properties reveal themselves – sometimes in surprising ways – through the engagement of multiple senses.

Acknowledgements

Much of the pleasure in creating these works comes in delving into areas where I have no prior knowledge. With each new topic, I need to beg for help from someone in the field who has expertise. Making acquaintances with such a wide variety of scientists and researchers is perhaps the most rewarding part of this work. Keith R. Jackson of Lawrence Berkeley Labs took the initial steps for the first version of the film, then invited me to participate when they took it further. Sadly, cancer took him from us, at far too young an age, just days before the Smithsonian premiere. The work is dedicated to him. Other help came from John Taber, E&O Program Manager at Incorporated Research Institutions for Seismology (<http://iris.edu>), Phil Slack, Tony Shakal from the California Strong Motion Authority and Derek Fox from the Penn State Astronomy Department. The sonification work was made possible by the support of Penn State's College of Arts and Architecture and NC2IF Center.

Supplementary Materials

To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/S1355771813000381>

REFERENCES

- Alexander, R.L., Gilbert, J., Landi, E., Simoni, M., Zurbuchen, T.H. and Roberts, D. 2011. Audification as a Diagnostic Tool for Exploratory Heliospheric Data Analysis. *Proceedings of the 17th International Conference of Auditory Display (ICAD 2011)*.
- Ballora, M. 2011. Opening Your Ears to Data. Presentation given at TEDxPSU. <http://www.youtube.com/watch?v=aQJfQXGbWQ4>.
- Ballora, M., Cole, R.J., Kruesi, H., Greene, H., Monahan, G. and Hall, D.L. 2012. Use of sonification in the detection of anomalous events. *Proceedings of SPIE 8758 Next Generation Analyst*. doi:10.1117/12.2016019.
- Ballora, M., Giacobe, N.A. and Hall, D.L. 2011. Songs of Cyberspace: An Update on Sonifications of Network Traffic to Support Situational Awareness. *Proceedings of SPIE 8064 Multisensor, Multisource Information fusion: Architectures, Algorithms, and Applications*. doi:10.1117/12883443.
- Ballora, M. and Hall, D.L. 2010. Do You See What I Hear? Experiments in Multi-Channel Sound and 3D Visualization for Network Monitoring. *Proceedings of SPIE 7709, Cyber Security, Situation Management, and Impact Assessment II; and Visual Analytics for Homeland Defense and Security II*, 77090J; doi:10.1117/12.850319.
- Ballora, M., Pennycook, B., Ivanov, P.C., Glass, L. and Goldberger, A.L. 2004. Heart Rate Sonification: A New Approach to Medical Diagnosis. *Leonardo* 37(1): 41–6.
- Ballora, M. and Smoot, G.S. 2013. Sound: The Music of the Universe. *The Huffington Post*. http://www.huffingtonpost.com/mark-ballora/sound-the-music-universe_b_2745188.html.
- Bregman, A.S. 1990. *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.
- de Campo, A. 2009. *An Interdisciplinary Approach to an Environment for Sonification of Scientific Data*. PhD thesis, Institute of Electronic Music and Acoustics – IEM, University for Music and Dramatic Arts Graz.
- Candey, R.M., Schertenleib, A.M. and Diaz Merced, W.L. 2005. *Sonification Prototype for Space Physics*. Abstract #ED43B-0850. Presentation at American Geophysical Union, Fall Meeting, San Francisco, CA.
- First, D. 2003. The Music of the Sphere: An Investigation into Asymptotic Harmonics, Brainwave Entrainment and the Earth as a Giant Bell. *Leonardo Music Journal* 13: 31–37.
- Fish, S. 2013. The Two Cultures of Educational Reform. *The New York Times*, 26 August. <http://opinionator.blogs.nytimes.com/2013/08/26/the-two-cultures-of-educational-reform/?ref=opinion&r=0>.
- Friendly, M. 2008. A Brief History of Data Visualization. In C. Chen, W.K. Härdle and A. Unwin (eds.) *Handbook of Data Visualization*. Heidelberg: Springer.
- Hart, M. 2012. Mickey Hart Performs at the Golden Gate Bridge 75th Anniversary Celebration (Video). Exploratorium TV. <http://www.exploratorium.edu/tv/index.php?project=7&program=1358>.
- Hart, M. and Smoot, G.S. 2013. *Rhythms of the Universe*. A 360° Production in association with the Berkeley

- Center for Cosmological Physics at the Lawrence Berkeley National Laboratory. Funding provided by the Richard Lounsbery Foundation. <http://www.ustream.tv/recorded/39395730>.
- Hayward, C. 1994. Listening to the Earth Sing. In G. Kramer (ed.) *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Proceedings volume XVIII, Santa Fe Institute Studies in the Sciences of Complexity. Reading, MA: Addison-Wesley.
- Hermann, T. 2002. *Sonification for Exploratory Data Analysis*. PhD dissertation, Bielefeld University.
- Hermann, T., Hunt, A. and Neuhoff, J.G. (eds.) 2011. *The Sonification Handbook*. Berlin: Logos Publishing House.
- Moore, F.R. 1990. *Elements of Computer Music*. Englewood Cliffs, NJ: Prentice-Hall.
- Newhouse, M. 2008. Solar Music: Using Helioseismology to See Inside Our Sun. http://www.noao.edu/education/ighelio/solar_music.html.
- Schlegel, K. and Füllekrug, M. 2002. 50 Years of Schumann Resonance. *Physik in unserer Zeit* **33**(6): 256–61 Translation: Catarina Geoghan, 2007, at www.hese-project.org.
- Wilson, E.O. 1998. *Consilience: The Unity of Knowledge*. New York: Vintage.
- Zander, B. 2008. The transformative Power of Classical Music. TED2008. http://www.ted.com/talks/benjamin_zander_on_music_and_passion.html.