From Sonification to Sound Art and Music: The interaction of usefulness and aesthetics.

Warren Burt, Faculty of Creative Arts, University of Wollongong

Abstract:

Although scientific data is most often displayed visually, it can often be just as advantageous to sonify data instead. Many problems about sense modalities are thus raised: what we are trained to perceive as useful information, how we can train ourselves to perceive information, how different aspects of information may be emphasized. The author was approached by the team from the Wollongong Room Calorimeter, headed by Associate Professor Arthur Jenkins, and asked if he would like to sonify data produced by the project. This paper follows the course of that work, with many sound and visual examples, discussing the problems involved in sonification, and the many areas where sonification crosses over into the aesthetic realm of sound art and music. One dimensional sonification, a series of pitches on a pulse, gives way to more complex realizations, which draw on ideas from fractal mathematics, such as canons of melodies in which pitch, duration and loudness are all determined by data. The crucial role of timbre will also be discussed: it often influences whether a sound is heard as “musical”, as well as aiding in the useful representation of data. The final outcome of the research is a musical composition, written with interactive software which allows one to explore the music inherent in the interaction of scientific data and musical creativity.

Keywords: sonification, timbre research, microtonality, generative composition, algorithmic composition.

In May 2005, I received an email from Mary Jane Leahy, honorary fellow in the Department of Biomedical Sciences, inviting me to get in contact with Associate Professor Arthur Jenkins and his team at the Wollongong Room Calorimeter project. She said that they had some data that a musician might be interested in using. I replied that I would be interested, and met them and saw their work. The Wollongong Room Calorimeter project consists of a small room in the Department of Biomedical Sciences which can be sealed off so that all activity in the room, specifically intake and output of gasses, can be measured precisely. Motion detectors, built by engineer Harry Battam, are also in the room, and the amount of motion in the room can be monitored and recorded, and correlated with gas usage. The data from the movement detectors was what they thought I might be interested in. The movement in the room is monitored 20 times a second. Output from the detectors can range from 0 to 10 volts, but, in the data I used, ranged between 3.59 volts and 7.05 volts. Participants in the experiment are asked to perform a number of activities, such as sitting still, walking, singing, drinking, typing, and riding an exercise bicycle. Each of these activities produces different data, with different patterning. The data I was given, by postgraduate researcher Femke van Nassau, consisted of a spreadsheet with 18 columns of data, each of which had 1200 elements. That is, the data consisted of readings of the motion of participants doing 18 different activities for one minute each. It was this data set that I was to engage with over the next several weeks, mining it for musical and sonification potential.

Sonification is defined by the International Community on Auditory Display (ICAD) as

“the use of nonspeech audio to convey information. More specifically, sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation. By its very nature, sonification is interdisciplinary, integrating concepts from human perception, acoustics, design, the arts, and engineering. Thus, development of effective auditory representations of data will require interdisciplinary collaborations using the combined knowledge and efforts of psychologists, computer scientists, engineers, physicists, composers, and musicians, along with the expertise of specialists in the application areas being addressed.” (Kramer, et al, 1997)

Sonification has a long history. Some of the scientific achievements in the field, quoted by ICAD include

“the Geiger counter, sonar, the auditory thermometer, and numerous medical and cockpit auditory displays, particularly those designed to present data variations. More recent successes include software that enables blind chemists to examine infrared spectrographic data via auditory presentation (Lunney & Morrison, 1990) and the mapping of data-dependent auditory signals to ongoing processes in dynamic monitoring tasks such as anesthesiology workstations (Fitch & Kramer, 1994) or factory production controls (Gaver, Smith, & O’Shea, 1991).” (Kramer et al, 1997)

There is an equally long and distinguished history, within the field of music, of composers using non-musical data as a source for composition. A few early examples of this rich history would include Brazilian composer Heitor Villa-Lobos’ compositions “New York Skyline” (1939), for piano, and his “Symphony No. 6: On the Profiles of the Mountains of Brazil” (1944) for orchestra, both of which used found object contours to determine thematic material; American composer John Cage’s Atlas Eclipticalis (1961-62) for orchestra; Etudes Australes (1974-75) for piano; and Etudes Boreales (1978) for cello and/or piano, all made by tracing star maps onto score paper; and American composer Charles Dodge’s “The Earth’s Magnetic Field” (1970) for computer, which mapped magnetic field measurements onto the notes of a diatonic musical scale. Also of interest from this early period is the essay “Any Bunch of Notes” (1953) by German emigre composer Stefan Wolpe, which deals with his Bauhaus and Dada influenced ideas about the shaping possibilities of any arbitrarily selected musical material. (Wolpe, 1953) More recent examples would include Czech composer Petr Kotik’s “There is Singularly Nothing” (1971-73) for voices and instruments, which used, among other sources, readings of the brain waves of fruit flies, Australian composer Tristram Cary’s “Contours and Densities at First Hill” (1976) for orchestra, which used topographical data and photographs from the Flinders’ Ranges as source material, and American Tom Hamilton’s “London Fix” (2003) for computer, which used the price of gold in London for the year 2002 as its source set.

Two remarkable projects that cross the line between scientific sonification and musical composition are the DNA music project of Professor Mary Anne Clark and John Dunn, and the 2004 “Listening to the Mind Listening” project of ICAD. Clark and Dunn’s work, especially their paper “Life Music: The Sonification of Proteins” and their several collaborative CDs of DNA music are paradigmatic for the field. (Clark and Dunn, 1997, 1998; Clark, 2001) Dunn has also developed a number of computer programs which are designed for sonification such as SoftStep, MicroTone, BioEditor, DataBin, and ArtWonk, some of which were used in this project. (Dunn, 1996-2005, van Raaij, 2004) The ICAD project involved composers from around the world using EEG data obtained from a person listening to music to generate further sound and music compositions. Ten of these compositions were selected and performed in the Sydney Opera House during the ICAD 10 conference in 2004. Australian contributors to the concert included Gordon Monro, Tim Barrass, the team of Roger Dean, Greg White and David Worrall, and from the University of Wollongong, the team of Guilliaume Potard and Greg Schiemer. The 10 pieces chosen have an enormous range of style and sound, showing just how wide the possibilities of sonification are. (Barrass et al, 2004)

In each of these compositions or projects, non-musical source material is worked with and the resulting music often has quite unfamiliar aspects to it. Learning to hear this music on its own terms often results in expansion of one's musical tastes and expectations. For example, in Mary Anne Clark's music, the same protein from several different species will be sonified simultaneously. Unison lines happen when the protein sequences are identical, and chords occur where there are differences in the protein between species. Here a musical patterning - the timing of chords and unison lines - is produced by biological information. The result is an unexpected musical timing and a revealing of aspects of a biological structure simultaneously. Similarly, any of the Cage star-map pieces are full of musical surprises - unexpected melodic leaps, sudden textural changes, etc, of the kind not normally found in traditionally produced musics. The joy of all these pieces is both in the finding of familiar musical gestures where none were suspected, and also in the opportunity they give us to discover musical material that we now can learn to hear, and use as a means of expanding our musical tastes and knowledge.

For this project, I wanted to start with as explicit a mapping as I could, in order to hear structures in the data, and to hear comparisons between structures contained in the different data streams. There are, of course, as many ways to sonify data as the imagination can come up with, but for this piece, I was concerned with what kind of mapping might be easiest to hear, and also with what kind of mappings we might be able to learn to hear. This project would be as much an exercise in ear training as in data mapping.

The first question was how to use the data? It consisted of sets of 1200 readings made at 20 readings per second, with readings that varied between 3.59 and 7.05 volts overall. However, only 2 data streams used that full range. The rest seemed to hover somewhere in between, often covering only a very narrow range. If I used the raw data, some data streams would cover a very wide range, and others a very narrow range. On the other hand, if I normalized each data stream independently, so that it covered the complete range the software could handle, finer details in the structure of each data stream might be made audible. Here are some graphics (almost the last ones in this talk) which show these relationships. Notice in the first 2, how "Sit" is very narrow in range, while "Drink" is much wider. Once the data is normalized, as shown in illustrations 3 and 4, each data stream covers the full range, and differences in the internal structure of each data stream are much more easily both seen and heard.

Illustration 1: Sit - raw data

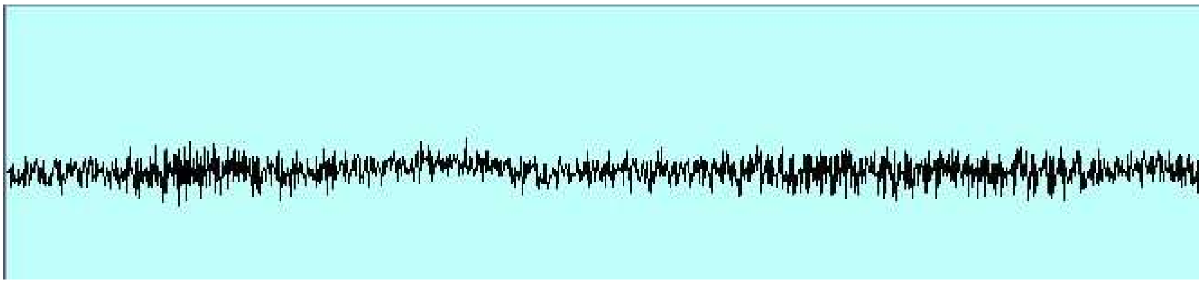

Illustration 2: Drink - raw data

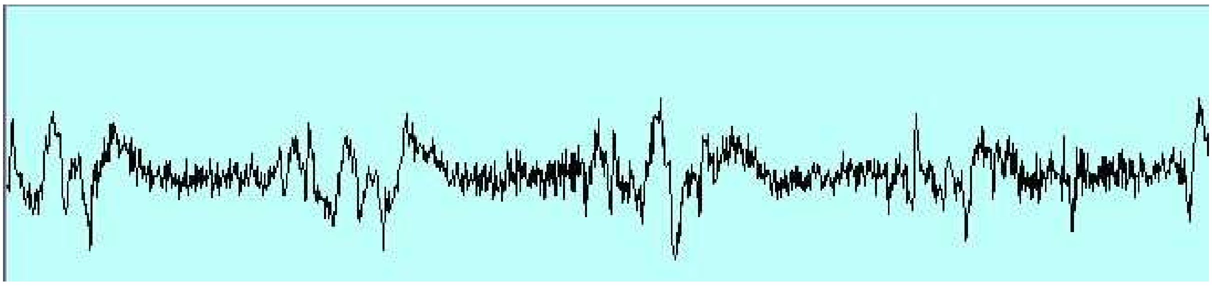

Illustration 3: Sit - normalized data

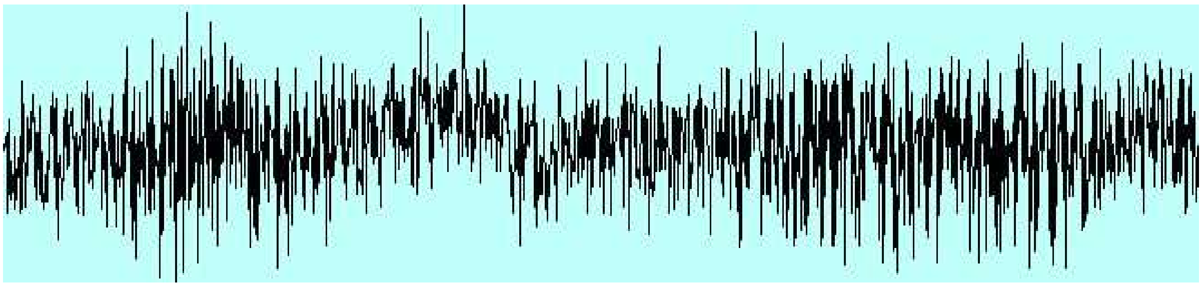

Illustration 4: Drink - normalized data

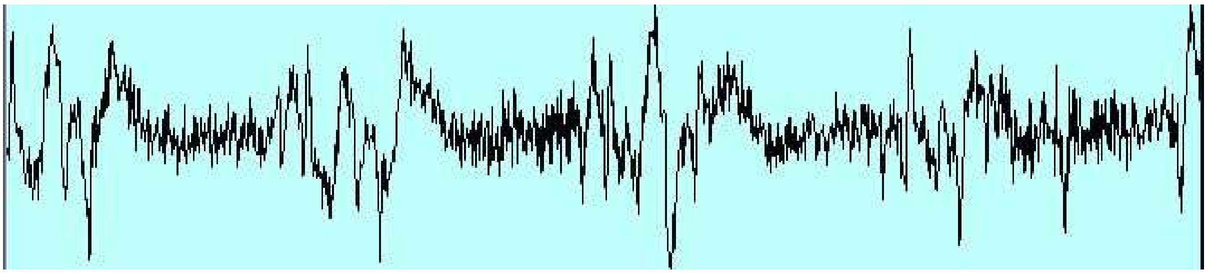

In the end, I decided to sonify both raw and normalized data, since the possibilities for establishing interrelationships between the different kinds of data seemed so rich.

My first impulse was to use the data to control pitch. That is, with the data ranging between 7.05 and 3.59 volts, that gives us, using increments of .01 volts, 346 possible values to apply to pitch. I decided that one could easily apply these values to a microtonal scale, of say, 53 tones per octave, and have a range of just over 6 octaves of pitch to control. Sound Example 1 is the data at real time rate - 20 notes per second - of "Sit" data - which covers a narrow range. Sound Example 2 is again, the data at real time rate - this time with "Drink" data - which covers a much wider range. The differences should be immediately hearable.

Controlling pitch, however, is only one way to sonify data - what if the data were used to control pitch and filtering of a sound derived from sonification of mathematical chaos data? The results would be much more complex. Sound Example 3 is "Sit" used to control this sound. Sound Example 4 is "Drink" used to control the same sound. The results, though sonically fascinating, might not lend themselves to easily hearing relationships within the data, because the sonic material being controlled already has such a complex texture on its own.

Both of the preceding sets of sound examples use the data as signals to control another sound. In the case of Examples 1 and 2, the sounds are electronic waveforms. In the case of Examples 3 and 4, the sound is a sample of an already complex chaotic sonification. What would happen if the raw data itself was used as a waveform. In the case of this data, various noisy waveforms resulted. Sound Example 5 is all the complete data set, all 21,600 values, heard as an audio waveform. Sound Example 6A is "Sit" as a waveform, and Sound Example 6B is "Drink" as a waveform. As can be heard, in this case, all three examples produce quite noisy results. The more varied waveform of "Drink" produces the most irregular sound. This turning data into waveforms is not without interest, however. For example, in Sound Example 7 is a sequence of someone resting, then bicycling at 40 rpm, then resting, then bicycling at 60 rpm. The ratio between 60 and 40 is 3/2. In musical terms, this ratio is heard as a perfect fifth. As one can hear, the cycling sections do indeed have pitch, and the pitches are in the perfect fifth relationship. Here, then, we can hear speed relationships compressed in time, becoming pitch relationships. For this project, however, rather than using the data as a sound waveform, I decided to stick with using the data to control traditional musical parameters, as these seemed to allow the quickest and easiest hearing of the qualities and details of the data.

In hearing sonifications, however, musical considerations immediately present themselves. For example, three of the most important aspects of musical sound are timbre, tuning and mode. These are heard by people without their even knowing what they are, and they are often what makes a piece immediately memorable. As anecdotal proof of this, consider any of those contests where people are called on to identify an old rock song. The real experts successfully identify the tune often after only 2 or 3 seconds. I maintain that an essential part of what they are remembering is the timbre of the tune, even before they remember the melody. Those wishing to see a more scientific approach to this question should refer to (Scheirer, Watson and Vercoe, 2001). The choice of timbre, tuning, and mode will be critical in the musical affect / effect of the sonifications. Sound Example 8 is a short melody played with a harp timbre, Sound Example 9 is the same melody played with a bad trumpet timbre, and Sound Example 10 is the same melody playing a set of rock drums. All three sound very different.

Further, tuning and mode have a great effect on how we hear. The next two Sound Examples are of a short harp melody played first, in a familiar sounding major scale in a familiar sounding tuning (Sound Example 11 - Scale 1 - Diatonic 31), and then in a similar major scale, but in a different tuning (Sound Example 12 - Scale 2 - Diatonic 27). Can you hear the difference?

If you couldn't hear the difference between those two scales, you'll certainly be able to hear the difference between the next two examples, which play the exact same little harp melody in two different modes and tunings (Sound Example 13 - Scale 3 - Chromatic 24, and Sound Example 14 - Scale 4 - Enharmonic 25). As you can hear, the same melody played in different scales and modes can have quite different emotional affects / effects.

So far, we've only used the data streams to control pitch - all rhythm has been on a pulse, and every note is the same loudness. Here are four sound examples which show aspects of this. All four of these examples use the same data stream - "Sit", and map it onto different sets of pitches. This first two play the data at real-time speed - 1200 notes per minute, while the second two play it at half speed - 600 notes per second.

Sound Example 15, which we've already heard, is the raw data version of "Sit" at 1200 notes a minute, mapped to a 53 note to the octave scale.

Sound Example 16 is the same data, but now mapped to a C-major scale. The 346 possible pitch levels are here reduced to the 28 possible levels of a diatonic scale covering 4 octaves. This is then played by a synthesizer with long decays, so that chords are allowed to form by the overlapping decays of the one note at a time sequence. I mapped the data to a diatonic scale in homage to Charles Dodge's "Earth's Magnetic Field", cited earlier. I immediately liked the sound of this. The familiarity of the diatonic harmonies seemed more "meaningful" (whatever *that* means) than the chromatic mapping of the previous example. I should note that I am inherently distrustful of what we perceive as immediately graspable meanings, if only because they many hide from us the delights of what we do not yet know.

Sound Example 17 is the normalized version of "Sit", played at 600 notes per minute. The normalization maps the full range of the data to the integer values 0-127. This is then mapped to several octaves of a 23 tone per octave scale.

Finally, Sound Example 18 is the same normalized data as the previous example, but mapped to a C-Major scale, again. At this point, I was beginning to wonder if the prettiness of the major scale and the sustaining timbres would get in the way of hearing data relationships, or if they would, in fact, make them more comprehensible. As stated before, I am not a great believer in the exclusive use of familiarity as a learning tool.

So far, we've only controlled pitch with our data. We could use it to control a number of other aspects of sound as well. In chaos science, a single stream of data is often mapped onto an x-y axis by using two slightly differently delayed versions of the data. That is, every Nth element of a sequence is mapped onto the x axis, while every N+1th element is mapped onto the y axis. The resulting "delay maps" often reveal interesting aspects of the structure of the data. (Peak and Frame, 1994) Similarly, in music, we could use delayed versions of the same data stream to control different aspects of sound, for example, pitch, duration and loudness of each sound event. Such a 3 dimensional sonic fractal has a name - in music, it's called a "melody." The next four sound examples create "3D" melodies using delayed data streams. Each pitch is determined by scaling element N, each loudness by scaling element N+1, and each duration by scaling element N+2 of the data stream.

In Sound Example 19, a raw data stream has pitch mapped to a 51 note per octave scale. We're using scales with so many notes because we have to fit 346 possible values into the range of hearing. This wouldn't be possible with only 12 notes per octave, unless we filtered the data, as we did for the diatonic scale examples heard earlier. Dynamics are made by dividing 346 levels by 2.73, giving us 128 possible levels, the full dynamic range of the MIDI computer musical control standard. Durations are made by dividing the 346 values by 6, and adding 3, giving us 58 different durations from very short to somewhat longer. Most of the values cluster around the middle of the range in this example. If we had used the same value for pitch, duration and dynamics, we would have always had short, soft and low notes, and long, loud, and high notes. By using the delayed data streams, we avoid this simple unvarying relationship, although in this data, usually values don't change widely from moment to moment, so some amount of correlation still occurs, as in occasional flurries of low, soft, fast notes produced when the data stream dips to lower values for an extended period of time.

Sound Example 20 is the same raw data stream mapped to a diatonic scale, a simple set of regular rhythmic durations, and only six loudness levels. The raw data version of "Sit" just stays around middle values, and because we're only mapping duration and dynamic to six values, the result is a moderately soft, pretty rhythmically stable melody. With the normalized versions of the data set, you'll get much wider variation in values.

Sound Example 21 is the normalized version of the "Sit" data, mapped to a 23 tone per octave scale, the full possible dynamic range, and durations from 6 - 31 pulses - very short notes to notes about 5 times as long as the shortest values. The slower tempo and the normalization allows one to hear internal relationships in each data set a bit more clearly. Again, you should be able to hear the correlation of low notes with softer and shorter values, but because of the normalization, there will be more large leaps between subsequent items of data, so the correlation is less pronounced.

Finally, in this set, Sound Example 22 is again, "Sit" normalized, but now mapped to a diatonic scale with 8 loudness levels and 8 durations. To my ear, we're beginning to approach one possible idea of "traditionally musical" melody here.
In chaos theory, the use of delayed data streams are called "delay maps." In music, the use of delayed melodies is called a "canon." "Row, row, row your boat" and "Frere Jacques" are two examples of canons everyone knows from childhood. Since we used delayed data streams - in musical terms, a canonic relationship, to generate the melodies, we could also delay the resulting melodies against themselves, and get sound examples with the musical form of the canon. One problem with our data is that it doesn't have repeating patterns in it - so it's hard to hear things that are interesting when new information is always happening - a canon allows one to hear material again, often at a different pitch level. This kind of repetition allows one to hear striking aspects of the material again, while allowing the data to progress unhindered. Sometimes the use of contrapuntal techniques like the canon can obscure things. Here, I hope I'm using it as an aid to comprehension. The next four Sound Examples are simply the previous four examples with the melody delayed against itself.

Sound Example 23 is the raw "Sit" data in 53 tones per octave as a 2 voice canon.
Sound Example 24 is the raw "Sit" data mapped to a C-major scale as a 2 voice canon.
Sound Example 25 is the normalized "Sit" data, in 23 tones per octave, as a 2 voice canon.
Sound Example 26 is the normalized "Sit" data, mapped to a C-Major scale, as a 2 voice

canon. Notice that the range of the melody is wider here than in Sound Example 24, and that the melody is more rhythmically irregular.

In the previous sound examples, the use of diatonic C-Major pitch set was just to give a sense of familiarity. But I found that I liked the sounds of the smaller pitch set better than the chromatic sound of the larger pitch sets. So I decided to use mapping to diatonic scales for the piece I would make with the data. However I didn't want to use the normal C Major scale and 12 tone tuning. I wanted to keep the familiarity of diatonic, but also explore other kinds of tunings. It turns out that there are a number of different scales that have the structure of the diatonic scale (2 kinds of intervals arranged A A B A A A B), and which exist in other equal temperaments than 12 tone tuning. In fact, it turns out that there are exactly 18 of these kind of tunings (technically called "Moment of Symmetry 5 + 2" scales) contained in the equal temperaments between 21 and 31 notes per octave, and they range from the familiarity of Sound Example 11, to the unfamiliarity of scales such as Sound Example 14. Appendix A has a complete listing of these scales.

To explore this set of 18 data sets and 18 possible tunings, I made an interface with John Dunn's data sonification software ArtWonk. It allows me a number of controls for exploring these musical potential of this material. Data Sets and Scales are chosen with the controls shown in Illustration 5. Horizontal control sliders allow selection of the following:
-Pitch mode - what note does the chosen scale begin on?
-Rhythm mode - the pitch levels are used to determine durations - which level of pitch are we using to start our set of durations with?
-Base octave - what is the lowest note used?
-Duration multiplier - to get slower tempi, if desired.
-Tempo in BPM - set the overall tempo. A tempo of 300 with a duration of 6 = 1200 notes per minute - the speed of the original data.
-Canon Delay in Ticks - length of delay between the first and second voice (the second voice is always an octave higher than the first voice).
-Timbre Select Voice 1 - choose between 18 possible timbres, produced by Martin Fay's software synthesizer Vaz Modular, made specially for this piece.
-Timbre Select Voice 2 - choose between 18 possible timbres, produced by Martin Fay's software synthesizer Vaz Modular, made specially for this piece.
-Counting by... sets how the data is stepped through. Counting by 1s gives every element in order. Counting by 3s gives every third element, in order, etcetera. Larger values of counting give shorter pieces and cruder scans through the data set.

Illustration 5 - the performance interface made with John Dunn's ArtWonk sonification software.

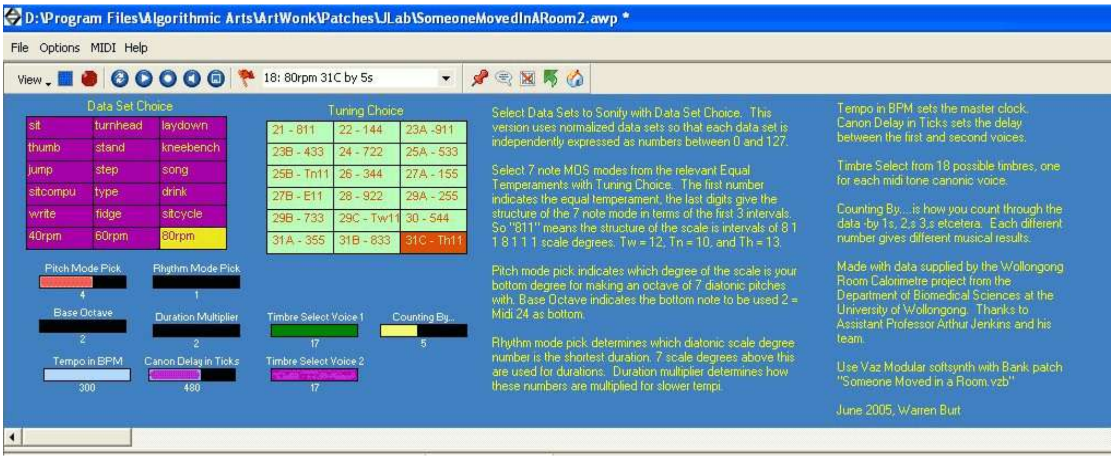

D:\Program Files\Algorithmic Arts\ArtWonk\Patches\JLab\SomeoneMovedInARoom2.awp *

18: 80rpm 31C by 5s

Data Set Choice

sit	turnhead	laydown
thumb	stand	kneebench
jump	step	song
sitcompu	type	drink
write	fidge	sitcycle
40rpm	60rpm	80rpm

Tuning Choice

21 - 811	22 - 144	23A -911
23B - 433	24 - 722	25A - 533
25B - Tn11	26 - 344	27A - 155
27B - E11	28 - 922	29A - 255
29B - 733	29C - Tw11	30 - 544
31A - 355	31B - 833	31C - Th11

Pitch Mode Pick 4
Rhythm Mode Pick 1
Base Octave 2
Duration Multiplier 2
Timbre Select Voice 1 17
Counting By... 5
Tempo in BPM 300
Canon Delay in Ticks 480
Timbre Select Voice 2 17

Select Data Sets to Sonify with Data Set Choice. This version uses normalized data sets so that each data set is independently expressed as numbers between 0 and 127.

Select 7 note MOS modes from the relevant Equal Temperaments with Tuning Choice. The first number indicates the equal temperament, the last digits give the structure of the 7 note mode in terms of the first 3 intervals. So "811" means the structure of the scale is intervals of 8 1 1 8 1 1 1 scale degrees. Tw = 12, Tn = 10, and Th = 13.

Pitch mode pick indicates which degree of the scale is your bottom degree for making an octave of 7 diatonic pitches with. Base Octave indicates the bottom note to be used 2 = Midi 24 as bottom.

Rhythm mode pick determines which diatonic scale degree number is the shortest duration. 7 scale degrees above this are used for durations. Duration multiplier determines how these numbers are multiplied for slower tempi.

Tempo in BPM sets the master clock. Canon Delay in Ticks sets the delay between the first and second voices.

Timbre Select from 18 possible timbres, one for each midi tone canonic voice.

Counting By....is how you count through the data -by 1s, 2,s 3,s etcetera. Each different number gives different musical results.

Made with data supplied by the Wollongong Room Calorimetre project from the Department of Biomedical Sciences at the University of Wollongong. Thanks to Assistant Professor Arthur Jenkins and his team.

Use Vaz Modular softsynth with Bank patch "Someone Moved in a Room.vzb"

June 2005, Warren Burt

With this set up, which gives me 18 data sets, 18 scales, 18 timbres (of which you can have two at once), 10 ways of scanning through the data, a variety of tempi, modes and octave choices, it's obvious that a very large number of different pieces can be made. To work within this infinity, I decided that sets of pieces should be generated, which would allow me to hear the relationships between the different data sets, using the following rules:

1) It's important to use the same algorithm - the same set of rules for all the pieces in a set - in order to really hear differences between the data sets - these then have the potential to become musically meaningful differences.

2) A set of pieces should use the same rate of stepping through the data set. In the end, I made two suites, one with each of the 18 normalized data sets stepped through by 3s, and one with each of the 18 raw data sets stepped through by 8s. This gave a normalized data suite of moderate length (47 minutes) and a raw data suite of shorter length (22 minutes). As well, a version of "Drink" normalized was made stepping through the data by 1s, which took just over 9 minutes. This was done because I felt that the "Drink" data set had the most interesting musical material contained within it.

Sound Examples 27 - 30 give the first few seconds of the normalized "Drink" data stepped through by 1s, 2s, 3s, and 5s. It is immediately obvious that different melodies are generated with each different stepping through the data.

Sound Example 27 - Stepping through Normalized "Drink" by 1s.
Sound Example 28 - Stepping through Normalized "Drink" by 2s.
Sound Example 29 - Stepping through Normalized "Drink" by 3s.
Sound Example 30 - Stepping through Normalized "Drink" by 5s.

When I started out this project, I was fairly committed to maintaining the "integrity of the data." With all this work with different tunings, timbres, and rates of scanning through the data set, I, at least, can still hear that each data set produces music with distinctly different qualities. However, at this point, I do have to ask myself what the concept of "integrity of the data" means in this context, anyway.

The completed computer music composition, "Someone Moved In A Room" consists of the two suites mentioned above, plus the 9 minute version of "Drink" made by stepping through every value. The completed work is on a CD with 37 tracks with a total duration of 80 minutes. It is available from me by request. In it, the listener can hear which data sets produce traditionally meaningful musical gestures, and which produce gestures that we might come to know as musical, given repeated future listenings.

I would like to thank the members of the Wollongong Room Calorimeter project for inviting me to do this project, and for their help along the way. Thanks to Associate Professor Arthur Jenkins, and Dr. Guy Plasqui, Mary Jane Leahy, Femke van Nassau and Harry Battam for their help. Thanks also to Brian McLaren for pointing out the research of Scheirer, Watson and Vercoe.

Appendix A: Scales used in “Someone Moved In a Room” for computer, Warren Burt, 2005

This list shows all the eighteen 7 Note Moment of Symmetry scales that have either 5L+2S or 2L+5S structures found in 21 through 31 tone equal temperament. Only the form with the smaller generator has been shown (each scale has a partner made with it's inversional interval counterpart, but the result is just a mode of the shown scales). The scale is made by piling up intervals of the generator’s size, and then collapsing everything within an octave, just as the diatonic 12 tone tuning scale can be made by piling up perfect 4ths or perfect 5ths.

All these generators hover around (mostly above) a Perfect 4th, from about 496 to 580 cents. This suggests that lots of diatonic MOS scales exist which are not related to equal temperaments, but which can be found with generators about this size.

The scales fall into 4 qualities - I’ve called these Enharmonic, Chromatic and Major and Minor Diatonic after the ancient Greek genera which they sound similar to. The Enharmonic scales have a pattern of two very small intervals followed by a much larger one, the Chromatic scales has 2 intervals, hovering around a minor 2nd, followed by a minor 3rd type interval within them, and the Diatonic scales have scales of alternating quasi Major 2nds and minor 2nds. All these kinds of scales have very distinctive sounds, but the Enharmonic scales sound more like each other, than they sound anything at all like the Diatonic or Chromatic scales, and vice versa.

21 tone equal temperament
Generator - 10 steps:
Pitch levels: 0 8 9 10 18 19 20 21 = 2L+5S intervals
Interval structure in scale steps: 8 1 1 8 1 1 1 - Enharmonic quality.

22 tone equal temperament
Generator - 9 steps:
Pitch levels: 0 1 5 9 10 14 18 22 = 5L+2S intervals
Interval structure in scale steps: 1 4 4 1 4 4 4 - Major Diatonic quality

23 tone equal temperament A
Generator - 10 steps:
Pitch levels: 0 4 7 10 14 17 20 23 = 2L+5S intervals
Interval structure in scale steps: 4 3 3 4 3 3 3 - Enharmonic quality

23 tone equal temperament B
Generator - 11 steps:
Pitch levels: 0 9 10 11 20 21 22 23 = 2L+5S intervals
Interval structure in scale steps: 9 1 1 9 1 1 1 - Minor Diatonic quality

24 tone equal temperament
Generator - 11 steps:
Pitch levels: 0 7 9 11 18 20 22 24 = 2L+5S intervals
Interval structure in scale steps: 7 2 2 7 2 2 2 - Chromatic quality

25 tone equal temperament A
Generator - 11steps:
Pitch levels: 0 5 8 11 16 19 22 25 = 2L+5S intervals
Interval structure in scale steps: 5 3 3 5 3 3 3 - Minor Diatonic quality

25 tone equal temperament B
Generator - 12 steps:
Pitch levels: 0 10 11 12 22 23 24 25 = 2L+5S intervals
Interval structure in scale steps: 10 1 1 10 1 1 1 - Enharmonic quality

26 tone equal temperament
Generator - 11 steps:
Pitch levels: 0 3 7 11 14 18 22 26 = 5L+2S intervals
Interval structure in scale steps: 3 4 4 3 4 4 4 - Major Diatonic quality

27 tone equal temperament A
Generator 11 steps:
Pitch levels: 0 1 6 11 12 17 22 27 = 5L+2S intervals
Interval structure in scale steps: 1 5 5 1 5 5 5 - Major Diatonic quality

27 tone equal temperament B
Generator - 13 steps:
Pitch levels: 0 11 12 13 24 25 26 27 = 2L+5S intervals
Interval structure in scale steps: 11 1 1 11 1 1 1 - Enharmonic quality

28 tone equal temperament
Generator - 13 steps:
Pitch levels: 0 9 11 13 22 24 26 28 = 2L+5S intervals
Interval structure in scale steps: 9 2 2 9 2 2 2 - Chromatic quality

29 tone equal temperament A
Generator - 12 steps:
Pitch levels: 0 2 7 12 14 19 24 29 = 5L+2S intervals
Interval structure in scale steps: 2 5 5 2 5 5 5 - Major Diatonic quality

29 tone equal temperament B
Generator - 13 steps:
Pitch levels: 0 7 10 13 20 23 26 29 = 2L+5S intervals
Interval structure in scale steps: 7 3 3 7 3 3 3 - Chromatic quality

29 tone equal temperament C
Generator - 14 steps:
Pitch levels: 0 12 13 14 26 27 28 29 = 2L+5S intervals
Interval structure in scale steps: 12 1 1 12 1 1 1 - Enharmonic quality

30 tone equal temperament
Generator - 13 steps:
Pitch levels: 0 5 9 13 18 22 26 30 = 2L+5S intervals
Interval structure in scale steps: 5 4 4 5 4 4 4 - Minor Diatonic quality

31 tone equal temperament A
Generator - 13 steps:
Pitch levels: 0 3 8 13 16 21 26 31 = 5L+2S intervals
Interval structure in scale steps: 3 5 5 3 5 5 5 - Major Diatonic quality

31 tone equal temperament B
Generator - 14 steps:
Pitch levels: 0 8 11 14 22 25 28 31 = 2L+5S intervals
Interval structure in scale steps: 8 3 3 8 3 3 3 - Chromatic quality

31 tone equal temperament C
Generator - 15 steps:
Pitch levels: 0 13 14 15 28 29 30 31 = 2L+5S intervals
Interval structure in scale steps: 13 1 1 13 1 1 1 - Enharmonic quality

References:

Barrass, Stephen, et al (2004), Listening to the Mind Listening at ICAD 2004, Sydney
http://www.icad.org/websiteV2.0/Conferences/ICAD2004/concert.htm

Cage, John (1961-62) Atlas Eclipticalis, for Orchestra,
a)Orchestra of SEM Ensemble, Petr Kotik, conductor. Asphodel #72000, 2000
b)Orchestra of SEM Ensemble, Petr Kotik, conductor. Wergo 6216, 1995

Cage, John (1974-75) Etudes Australes, for piano, Stefan Schliermacher, piano, DABRINGHAUS & GRIMM MDG 613 0795-2 (3 CD's), 2001

Cage, John, (1978) Etudes Boreales,for cello or piano or cello and piano
a) Mode Records, LP 1 / 2 1985
b) Cello version alone: Etcetera CD 2016, 1992

Cary, Tristram (1976) Contours and Densities at First Hill, for orchestra,
http://www.amcoz.com.au/comp/c/tcary.htm

Clark, Mary Anne (2001) Ancient Airs: A Musical Reading of Protein Structure
http://www.whozoo.org/mac/Music/SantaFe/AncientAirs.html

Dodge, Charles (1970) The Earth's Magnetic Field, for computer
a) Nonesuch LP H71250, 1971 - both diatonic and chromatic movements
b) New World CD 80521-2, 1998 - diatonic movement only

Dunn, John and Clark, Mary Anne (1997), The Sonification of Proteins
http://mitpress2.mit.edu/e-journals/Leonardo/isast/articles/lifemusic.html

Dunn, John and Clark, Mary Anne (1998) Life Music - 3 CDs - http://algoart.com/music.htm

Dunn, John (1996-2005) ArtWonk, SoftStep, DataBin, MicroTone, BioSon, algorithmic composition software. http://algoart.com

Fay, Martin (1998-2005) Vaz Modular synthesis software.
http://www.software-technology.com

Hamilton, Tom (2003) London Fix, Muse Eek CD MSK 118, 2003

Kotik, Petr (1971-73) There is Singularly Nothing, 21 solos, version for solo flute on "First Record" LP Cramps Records 5205-114, Italy, 1977

Kramer, G. et al (1997): Sonification Report: Status of the Field and Research Agenda
http://www.icad.org/websiteV2.0/References/nsf.html

Fitch, T., & Kramer, G. (1994). Sonifying the body electric: Superiority of an auditory over a visual display in a complex, multi-variate system. In G. Kramer (ed.), Auditory display: Sonification, audification and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD) 1992, 307-326. Reading, MA: Addison-Wesley.

Gaver, W. W., Smith, R. B., & O'Shea, T. (1991). Effective sounds in complex systems: the ARKola simulation. In Proceedings of CHI '91. New York: ACM.

Lunney, D., & Morrison, R. (1990). High technology laboratory aids for visually handicapped chemistry students. Journal of Chemistry Education 58, 228.

Peak, David and Frame, Michael (1994), Chaos Under Control, New York, Freeman and Co. pp. 204 - 216

Scheier, E., Watson, R. & Vercoe, B. (2001) On the Perceived Complexity of Short Musical Segments, http://web.media.mit.edu/%7Eeds/icmpc00.pdf

Van Raaij, Hans (2004) Listening to the Mind Listening - A Sonification / Composition
http://www.icad.org/websiteV2.0/Conferences/ICAD2004/concert/HansVanRaaij.pdf

Villa-Lobos, Heitor (1939) New York Skyline, Etcetera KTC 1123, 1991

Villa-Lobos, Heitor (1944) 6th Symphony: On the Profiles of the Mountains of Brazil, CPO Records 999 517-2, Germany 2002

Wolpe Stefan (1953) "Any bunch of notes", ed. A. Clarkson, Perspectives of New Music 21/1-2, 295-310 (1982-83).