Queen Mary
University of London

Issue3

AUDIO!

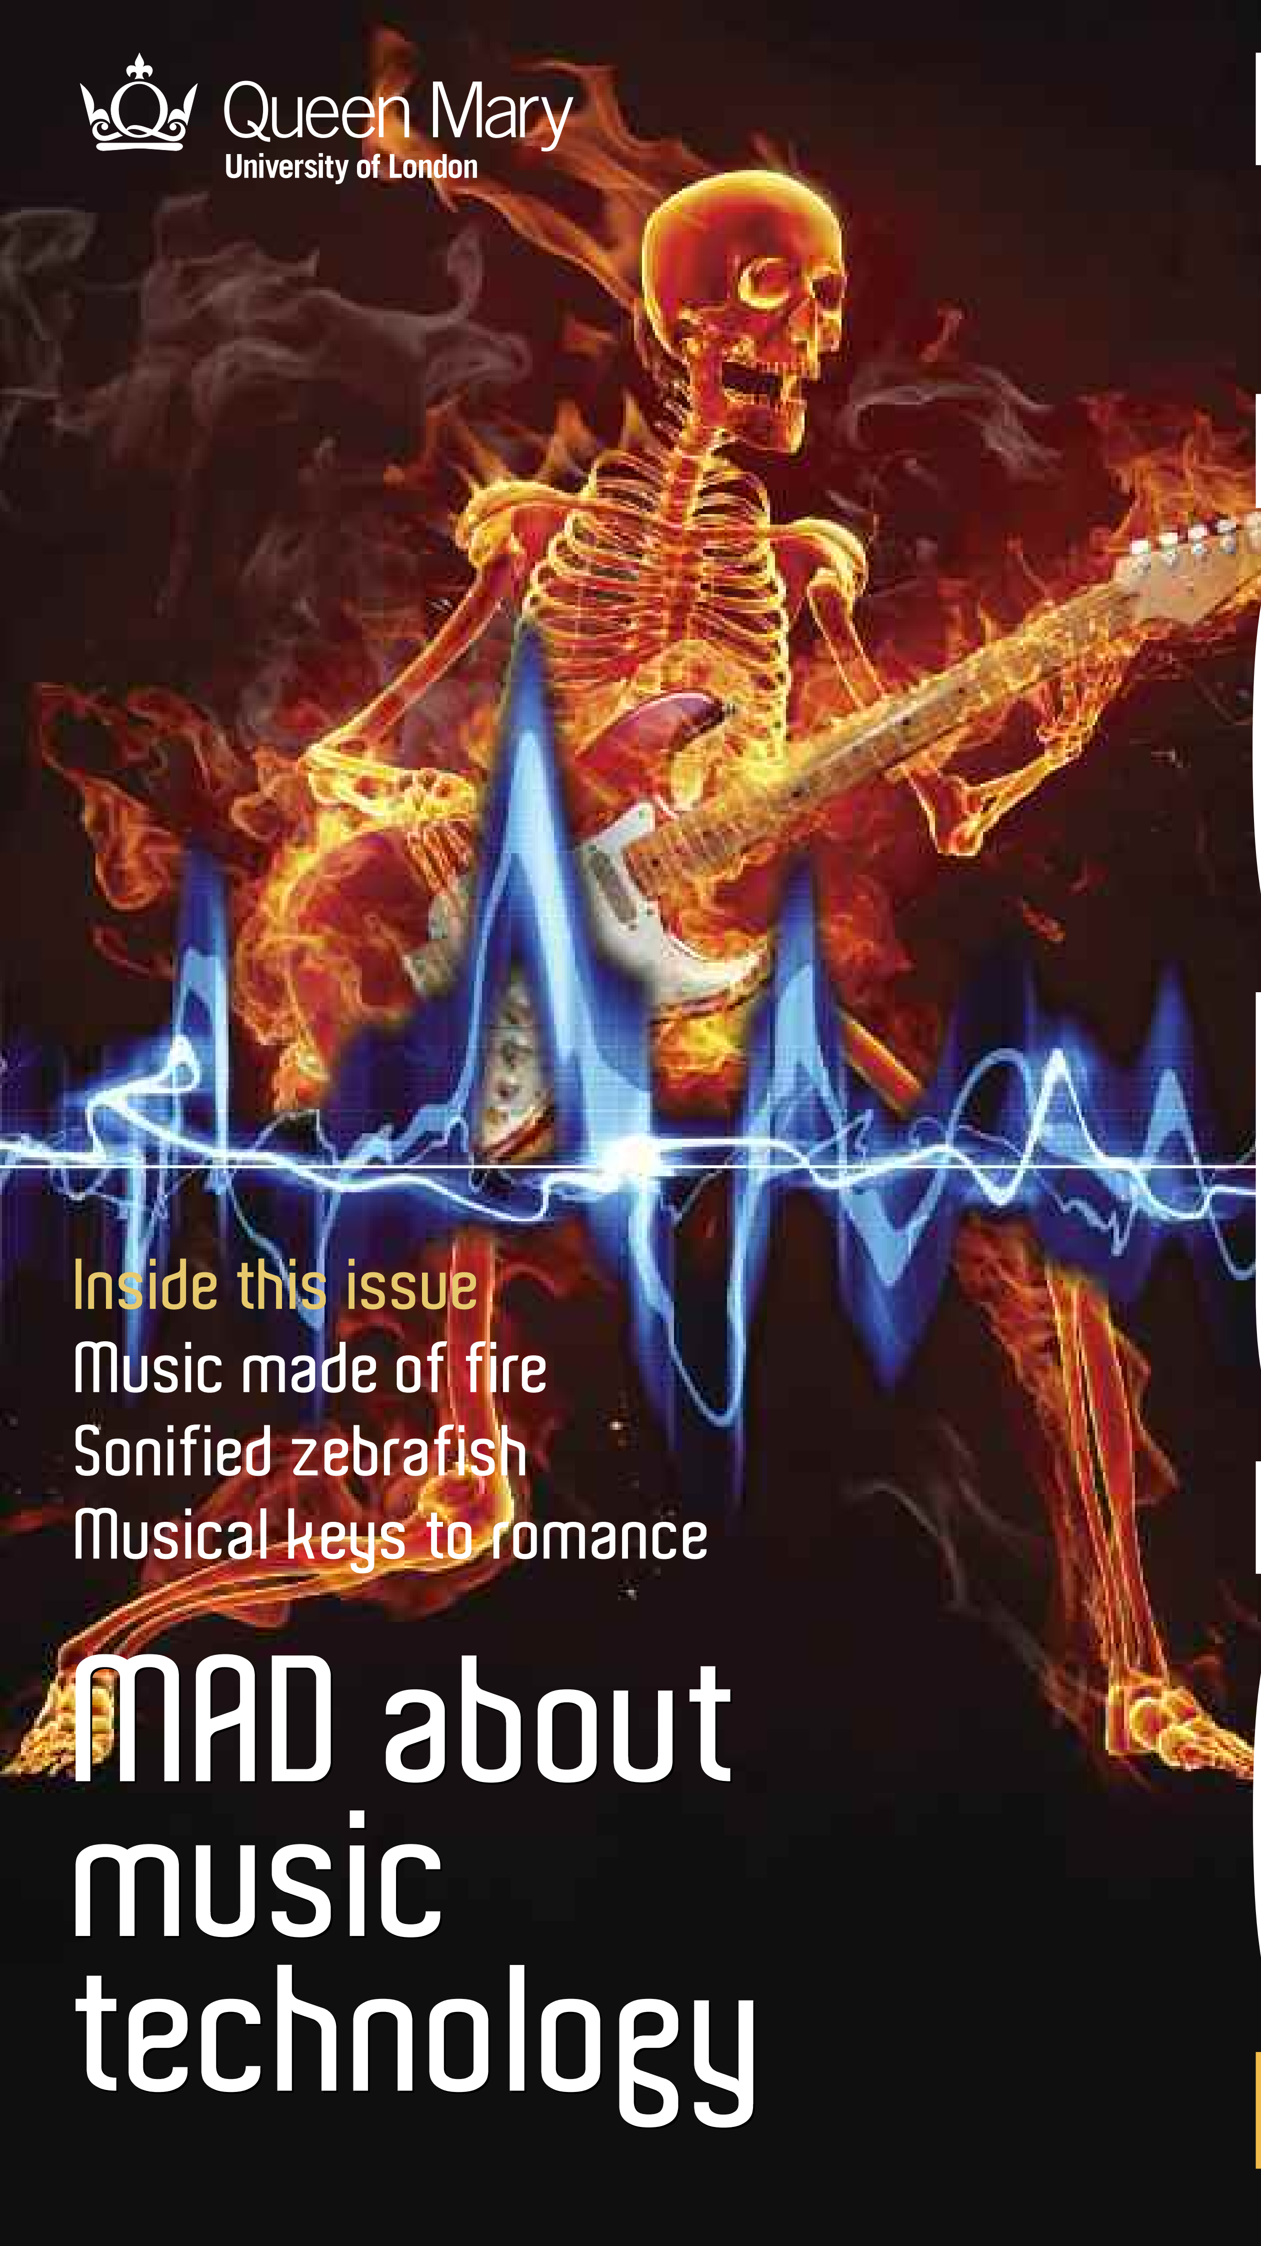

Inside this issue

Music made of fire

Sonified zebrafish

Musical keys to romance

MAD about music technology

Musicology for the masses

Some people listen to music, some want to play music, others – ‘musicologists’ – study it too. They are interested in all sorts of questions. What makes good music and why? What is the history of music? How does it vary across cultures? One branch of musicology is about analysing music, and that is where technology can help...and even spin off whole new industries.

If the audio engineers can come up with programs that make it easy to analyse music, to search in new ways, to visualise it, then professional musicologists, hobby musicians and music lovers alike can benefit. That is what Queen Mary’s ‘Musicology for the masses’ project is exploring.

In this issue of Audio! we explore research that is applying musicology to help the masses: technology that will help bands make better recordings; how to create new formats for storing music; how analysing the acoustics of music venues has led to tools for recreating the sounds of any building; a new way to visualise music collections; and how musicologists are even studying what goes on in schools.

Musicology in school

One place where people talk about their different ideas of music is secondary school music classes so that is where the musicology researchers from Queen Mary are going. They are applying a research technique known as ‘ethnography’: observing classes in depth. The aim is to provide a rich understanding of how people’s common-sense understanding of music relates to the concepts taught in class. They will also look at how the ideas that emerge relate to the information that can be analysed and visualised using digital music technologies. The results will then drive the development of new technologies to support music education.

If you want to make a start looking into music files rather than listening to them then why not download Sonic Visualiser from www.sonicvisualiser.org? It is being extended as part of the Musicology for the Masses project.

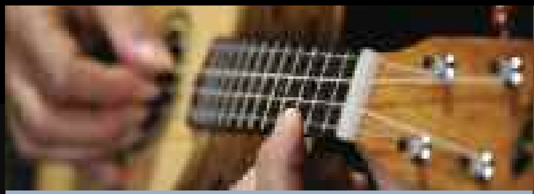

Ukulele Hero

Mastered the entire "Rock Band" series? Achieved global fame touring the world on "Guitar Hero"? If so, perhaps it's time for you to crank the volume and wow those imaginary crowds with your fret-melting chops on ... "Ukulele Hero"!

The ukulele is a small Hawaiian guitar-like instrument with four strings that originated in the 19th century. Okay, so Ukulele Hero isn’t quite as rock 'n' roll as its guitar-based counterparts. It's also not strictly a game, but rather an interactive virtual instrument application that simulates the strings and frets of a real ukulele. It was created by Duncan Menzies a graduate student on the Media and Arts Technology programme at Queen Mary. Using Duncan’s custom-built hardware controller, the player can strum actual chords as though they were playing a genuine acoustic instrument. By connecting the controller to a computer, the software produces both the audio, and an on-screen visualisation of the frets and strings being played.

Notes are played by pushing 16 buttons that correspond to the positions behind the frets of each string. The button in the top right corner corresponds to the position behind the first fret of the highest string, and so on. You ‘pluck’ strings by moving a slider up and down. It is made from a ‘potentiometer’; which is just a variable resistor that changes the input voltage as you slide the slider.

Ukulele Hero was written using an open source "creative coding" environment called "openFrameworks", which is essentially a set of libraries for the programming language C++. OpenFrameworks is a great toolkit for creating interesting applications quickly, even with little or no programming experience, as it provides an easy way to use powerful libraries of code that other programmers have already written. So if you are into both electronics and music, why not have a go yourself – not at playing Ukulele Hero, but at creating your own digital instrument. For a video demonstration of Ukulele Hero, follow the links from www.audio4fn.org.

Audio! Action

Former Beatle George Harrison was a ukulele fan and had a big collection of ukes.

Romantic songs and the dating game

If you struggle to get dates it might just mean that you play the wrong music! A team of French researchers has shown that playing the right music before asking for a date can increase a person's seductive powers.

Many studies have shown that media can affect our behaviour. Violent video games or music with aggressive lyrics can increase aggressive feelings, thoughts, and behaviours. Do we also act differently when listening to romantic songs? That is what the team from the Bretagne-Sud and Paris-Sud Universities decided to find out. In a previous experiment, they had shown how romantic songs played as background music in a flower shop encouraged men to spend more money. They decided to design a new experiment to test whether songs with romantic lyrics would influence the dating behaviour of young single women.

> "Women gave out their telephone numbers more often if they'd been listening to romantic music."

First they had to set up the experiment. Based on the answers of participants to music questionnaires, they selected a 'neutral' song ("L'heure du thé" from Vincent Delerm) and a romantic song ("Je l'aime à mourir" from Francis Cabrel) to use in the experiment. They also asked a group of women to rate the physical attractiveness of twelve men. For the next stage, they wanted someone not too attractive, but not too repulsive either, so they chose 'Antoine', the man whose score was nearest the average for the twelve.

Over 80 women aged between 18 and 20 took part in the main experiment. First, each woman listened to background music in a waiting room. She was then invited into a separate room to discuss the differences between two food products with a young man (actually Antoine). This was done in the presence of the experimenters but they then left asking the man and woman to wait for a few minutes. This gave Antoine the opportunity to try a straight-forward chat-up line: "My name is Antoine, as you know. I find you very good looking and I was wondering if you'd agree to give me your telephone number. I'll call you and we could have a drink together next week." (It presumably sounds very seductive in French!)

Audio! Action

Mozart once helpfully composed a canon where the second melody was the same as the first but backward and upside down. That meant it could be read by two players from the opposite sides of the music sheet.

Now the intriguing bit. It turned out that when the romantic song had been playing in the waiting room Antoine's chances of getting a telephone number doubled. 52% of the women who heard the Cabrel song agreed to give their number compared to 28% of those who heard the 'neutral' song from Delerm.

These findings fit a general model proposed back in 2006 that aims to explain the effects of media exposure. It suggests that media exposure in general (not only violent or aggressive media) affects people's internal states. 'Pro-social' media (as opposite to antisocial ones) should favour pro-social results.

Why can our behaviours be influenced by music? Certain types of music may induce positive 'affects'. In psychology, an 'affect' is the experience of having a feeling or emotion. Our better receptivity to declarations of love is because romantic songs induce positive affects – they give us a positive experience and that makes us susceptible to positive behaviour. Is it the romantic meaning of the lyrics or the way the instrumental music was composed that matters most? Answering that will take some new studies. In the meantime if you are single, then pay attention to the background music and it may well be love at first sight…

Parts of this article were translated from French from: **www.insoliscience.fr/?Les-chansons-romantiques-le-truc**

Combing for

Your band has played some gigs and now you are going to record a demo. You set up the equipment, scatter microphones around, hit record and off you go. You have a great time and it felt really good. You play back the track and...it sounds horrible! The bass has all but disappeared and what are those other weird sounds? Surely you weren't playing that badly. Probably not, but maybe you need to learn a bit more about microphones and effects called 'microphone bleed' and 'comb filtering'. With any luck though, some time soon Alice Clifford from the Centre for Digital Music at Queen Mary, University of London will have made the problem a thing of the past.

A microphone is really easy to use. It will reproduce any sound in a specific area around it as an electrical signal, but anything from elsewhere isn't picked up. That means to record an acoustical musical instrument (say, a saxophone) you point the microphone directly at it. If it's an electrical instrument (like an electric guitar) you point the microphone at the amplifier. If the sound source falls within the sensitive area of the microphone, the sound of the instrument can be reproduced and so played back through speakers or recorded to hard disk.

There isn't that much that can go wrong with this, but there can be problems. Other instruments could be picked up in the microphone's signal (that's 'microphone bleed'), or there might be some reverberation or random noise on the microphone signal. You solve problems like this using soundproofing and being careful about where you place the microphones. The idea is to make sure the interfering sounds end up in those areas that the microphone can't 'hear'.

Often you want to record more than one instrument at the same time or several microphones need to be placed close together to record one instrument like a drum kit. This leads to a specific problem when more than one microphone picks up the same sound. When the microphone signals are mixed together, copies of that sound can be added together. The trouble is that they may have taken different amounts of time to get to the different microphones, so may not be perfectly synchronised.

This is where 'comb filtering' comes in. Why that name? It has nothing to do with bad hair days, even though that's what it makes your music sound like! It is because the electronic signal involved looks like a comb. In signal processing, a 'filter' is a

sound

device or process that modifies features of a signal. In this case the process of combining the out of time sounds can be thought of as a filter. The sounds interfere with each other in a way that cuts out some groups of frequencies whilst boosting others. Draw a picture of the signal's interference with its spikes and gaps at different frequencies and it looks like the teeth of a comb.

As well as other weird sounding effects, comb filtering can lead to large groups of frequencies being cut out that should have been kept. For example, you may find your bass guitar recording suddenly has no bass frequencies in it. The effects can be great for giving expression to a guitar solo, but it isn't what you want on a classical harp recording.

You can reduce the problem by paying attention to the positions of microphones in relation to each other and to other instruments. Ideally all microphones would be placed exactly the same distance from each instrument but that is impractical. A standard rule used by sound engineers is the '3 to 1' rule: always place a second microphone or instrument at least 3 times the distance from the first instrument to the first microphone. The idea is to make the strength of the sound in the second one low enough that it doesn't affect the first one. Basically, the aim is to reduce microphone bleed simply by placing the instruments and microphones a long way from each other.

So next time you are recording your music band, you should pay close attention to where you place the microphones. Sound engineers painstakingly position microphones for good reasons! If the audio that comes out sounds different to what you expect, try moving microphones as it might improve things. The only trouble is perfect positioning is not always possible. It would be much better if you could just do some clever audio processing after the sounds have been picked up by the microphone.

This is where Alice Clifford's work comes in. She has set herself the problem of how to analyse audio signals and apply processing to remove the comb filtering effect. One way this can be done is by "nudging" audio regions in a digital recorder so visually they line up and so will be played back at the same time. Often, the comb filtering may not be apparent, but as soon as the audio is nudged in line, the kick drum has more of a thump or a guitar has more presence. Of course, nudging regions may not be very accurate. Alice is aiming to come up with a way that solves the problem completely. If successful it will mean that in future sound engineers would not have to move the microphones at all.

So at the moment you have to be extra careful over where you place the microphones when recording your band. With any luck though Alice will solve the problem, and her signal processing software will mean messing about with the position of microphones before recording will be a thing of the past.

MP3: something's gotta change!

The music industry has gone through a massive upheaval as a result of digital music, and more and more interactive music services, from iTunes and spotify to sonarflow, appear all the time. If the revolution is to continue at the same pace though *something's gotta change!*

The original digital music revolution wasn't just on the back of the idea of recording and storing music digitally but because everyone agreed on a common music format - the MP3 format. MP3 is just a particular code for storing music digitally – as strings of 1s and 0s. There are lots of different ways of doing that, and if different record companies were to release music using different formats – a different code – it would be chaos. Everyone would need lots of different players and worse need to know which one was needed to play which tracks. Confusion would reign. Having a common format means that anyone who has an MP3 player can play any music. Better still, it means anyone with a great new idea can create a new service based on it and everyone else will have the capability to use it. The trouble is there is a lot you can't easily do with MP3 so some great ideas just won't get off the ground if we stick with it.

The thing about MP3 is that it follows the tradition of CDs and records, and only stores the music in its final mixed form. When music is recorded lots of different tracks are recorded separately. In the simplest situation, for example the vocals and the music might be recorded and stored separately. In reality many more tracks than that are recorded – each separate instrument might be recorded separately for example. The producer then mixes the tracks together to give the final music that you download.

Because all the music has been integrated into mixed tracks before it is distributed means digital music isn't very flexible at the moment. If instead all the original tracks were separately available in the stored music file lots of new applications would be possible. That is what the new 'Interactive Music Application Format' known as 'IM AF' does. It is the brainchild of an organisation known as MPEG - the Moving Pictures and Associated Audio Experts Group – the same group that came up with MP3 in the first place.

IM AF allows you to store and share music as lots of audio tracks and combine it with other information too. That creates lots of exciting possibilities for new ways of experiencing music. Listeners can listen to different music producer mixes of the same music, for example – just store the mix settings in the IM AF file. This kind of multi-track interactive music format allows you to do more than listen passively though. You can also interact with it. For example, why not experiment with your own mixes? Lots of things become simple to do. To create a Karaoke version of a single to sing along to you just pick out the music tracks and drop the vocals. To practice your musical instrument skills isolate the instruments you want to play along with – either the one you are playing or some others to accompany you. You could also emphasize the melody, harmony, or rhythm of a particular piece of music depending upon your personal tastes. If you don't want to do all this yourself, don't worry, as services will spring up to do it at the touch of a button. Ways to legally share your modified versions with your friends also become a possibility if some budding entrepreneur wants to set up a service to support it.

This simple idea of a change in the way music is stored means not only that music listening becomes more flexible, but we could all become creative music producers if we want to. Better still if you want to be a digital music entrepreneur, a whole new space of ideas for services is just around the corner.

Sing in the Royal Albert Hall while staying at home

In the early days of music recording, producers would find the best sounding halls for bands to play in to get the most interesting acoustics on the record. But then the budgets for hiring interesting venues disappeared. Instead vocalists were put in acoustically 'dry' spaces in the studio that had no character. That led to attempts to add character to recordings using ‘artificial reverb’. Early reverb systems used metal plates or springs to emulate the sound of multiple echoes from walls. These didn’t give the fine control needed to copy the sound of a particular room though so ‘digital reverb’ was born.

Digital reverb involves creating those interesting acoustics on a computer and applying them to the recording afterwards. Building a reverb that sounds exactly like a real room is very difficult so it’s taken some clever engineering to design a way to capture a room's sound. The first step is simple. You just play a known sound through a speaker in the room and record it. The sound used is usually what is called a ‘swept sine wave’ – a sound wave that goes from the lowest rumbling frequencies up to the high frequencies only dogs can hear. That kind of sound is easy to identify so we can take it away from what ended up being recorded to obtain a sort of acoustical fingerprint of the room – technically it’s called the ‘room impulse response’. Doing this slowly isn’t too difficult, but some complicated mathematical trickery is used to do it as the recording is happening without a noticeable delay. This process of removing the sweeping sine wave, known as ‘deconvolution’, is a neat trick. It means we can then add (‘convolution’) any dry audio recording to the fingerprint of any room even if the room is on the other side of the world. Our up and coming vocalist’s new album can sound as if it was recorded in any room we like even though she stays in that boringly dry studio. This basic idea has also been used by researchers at the Helsinki University of Technology to create a 'Rehearsal Hall with Virtual Acoustic for Symphony Orchestras': a system for musicians to do real-time performance and practising.

On the other side of the Atlantic, researchers at McGill University in Canada have used the same approach with a different purpose. They’ve created a system for measuring the musicality of musicians as they rehearse. First the room where the musicians will play is acoustically treated so that it doesn't have any big echoes or sound character of its own anymore. Then, lots of loudspeakers are set up around the musicians. Each instrument or instrument group has a microphone and its sound is combined with the acoustics of a different room. That combined sound is fed back to the individual musician through the loudspeakers. The overall effect is that each musician can play listening to the acoustics from their favourite venue. Experiments have shown that this improves musicality in rehearsals, and a happy musician leads to a more exciting performance.

The technology can also be used to pre-practice as if in a venue where you are going to perform, help composers understand what their piece will sound like at its premier and let orchestras choose a different venue for each piece. Have you ever wondered what your voice would sound like in the Royal Albert Hall? Install the technology at home and that is where you could practice your singing from the comfort of your bedroom.

And the beat goes on… forever?

Hip-hop music is usually around 80 beats per minute (bpm). Pop music is faster, around 120 bpm, and dance music can be anything from 120 up to around 200 bpm. But what if you had a beat that kept speeding up forever? How fast could you possibly get without it just becoming a blur?

The composer and computer musician Jean-Claude Risset came up with a way that a rhythm could keep speeding up forever, and yet however much it speeds up it sounds the same. It's actually an audio illusion as even though it sounds like the music is speeding up, you are just listening to the same music on a loop. It's done by combining different versions of the rhythm at different relative speeds (for example, 30 bpm as well as 60 bpm, 120 bpm, 240 bpm...) so that as the beat speeds up, different "layers" can take over as the main rhythm we perceive.

Music researcher Dan Stowell came up with a way to apply Risset's effect to any breakbeat rhythm. Follow the links from **www.audio4fn.org** and listen to the beat - tap along and see how fast you can get...

Music, Fire and Smoke

Music is all about strumming strings, hitting things or huffing and puffing into tubes. Fire and smoke doesn't have much to do with it - well, unless you are at a rock concert where they have decided to set off some pyrotechnic bangs maybe. So it may come as a surprise that people have been making music with fire for at least a century, that smoke may help make better microphones, and that flames give a new way to watch music. Dan Stowell tells us more:

"Music has never been so hot."

Sound Sculpting with fire


Photo CC-by-nc-sa Stephen J Johnson

Ruben's tube: a six-foot tube with flames pouring out, dancing in response to sound. Music has never been so hot.

The normal way to visualise sound is using an EQ meter. It is just a sound meter with bars of lights driven by the sounds. Each bar is linked to ranges of frequencies (essentially groups of notes). The more energy from those frequencies in a sound the more of the lights in the bar that light up.

Ruben's tube does a similar thing. A direct, physical reaction to sound is created by placing a loudspeaker at one end of a horizontal gas-filled tube. Gas comes out of 200 holes in the tube. With no sound it would come out of each hole at the same rate, but sound from the speaker creates standing waves down the tube. They are just waves that stay in the same place. It's a bit like being on a beach where the waves don't rush towards the shore but all just stay put. Now imagine the waves are in a tube. Where the crests are, more water has to compress into the same space than where the troughs are. Having a standing wave in a tube means the pressure is greatest where the crests of the wave are, and least in the troughs. The greater the pressure, the higher the flame is forced out at that point. There's no digital trickery involved, it's driven entirely by the physics of acoustic waves in tubes. The effect is that 200 flames dance around in response to changing sounds made into a microphone. It's an EQ meter made out of a gas fire!

Fire Music: the Pyrophone

Photo CC-by-sa Thomas Goller

"The idea of making music with flames might seem pretty steampunk, but it's been around a while. The German composer Wendelin Weissheimer was playing a Pyrophone in the late 19th century."

In Ruben's tube fire is used to make sound visible. Flames aren't just affected by sound, they can be used to make sound too. A pyrophone is a flame organ that uses small explosions or combustion to generate musical notes.

A "normal" organ uses air blown across large cylindrical pipes to start the air columns vibrating and so create sounds, just as blowing across the top of a glass bottle creates a note. For some people that's not exciting enough though. A pyrophone usually involves connecting a supply of gas to each pipe. A burst of flame in the bottom of a pipe is used to set the air vibrating and create each note.

To see and hear a pyrophone in action, follow the links from www.audio4fn.org

A Microphone made of smoke

What good would a microphone made of smoke be? Potentially incredibly sensitive it turns out. Microphones turn sound – waves of movement of the air – into electrical signals. Most microphones have a moving part inside them, a diaphragm whose movement is used to detect the moving air and so pick up the audio signal. Rather than use a diaphragm, father-and-son team David and Daniel Schwartz have invented a microphone that uses a single plume of smoke instead.

Why? The physical mass of the diaphragm in a normal microphone limits how responsive it can be. Soft sounds may not be strong enough to push the diaphragm so would not be detected. David Schwartz noticed that smoke rising from an oil lamp moved in response to his voice, and wondered if that movement could be detected using a laser, to create an ultra-lightweight way of sensing sound. Son Daniel wondered if he could actually make one. A few months later and the Schwartzes had made a prototype of their new microphone.

The device shines a laser pointer through a thin plume of smoke rising through a tube. Any sound disturbs the smoke and makes it vibrate, and that changes the amount of laser light that can pass through. A sensitive light detector on the other side detects the light and turns it into an electrical signal. Out of the smoke you have a microphone!

To see a video of the smoke microphone, follow the link from the Audio! website: www.audio4fn.org

Zooming in on music

Digital music collections, both personal and professional, are constantly growing. It's getting increasingly difficult to find the track or sample you want amongst the tens of thousands you own, never mind the millions of titles in online music stores. And how on earth do you browse music for particular sounds? Are we doomed to constant frustration, poring through lists and lists of tracks failing to find songs we know are in there somewhere? Perhaps not.

Spot that tune

At the moment, the main way of searching collections is based on textual information added to the tracks that describe them. This text is known as 'metadata'. It means you can search for keywords in that textual information and get Google-like lists of results back – you are just doing text searches. The trouble is the searches can only be as good as the metadata that has been added.

A more interesting approach is to actually search the sounds based on analysing the audio data itself – search the music not the text that has been written about it. Over the last decade, algorithms have been developed that can automatically detect the rhythm, tempo, genre or mood of a song, for example. This is known as 'semantic audio feature extraction'. Once you can detect these features of music you can use them to work out how similar two pieces of music are. That can then be the basis of a search system where you search for tracks similar to a sample you have already, browsing through the results. This approach can also be used in music recommendation systems to suggest new tracks similar to one you already like.

Photo courtesy of Spectralmind OG

Exploring bubbles

Having broken away from searching text comments, why not get rid of those lists of results too? That is what the company Spectralmind's 'sonarflow' does. Spectralmind is a spinout company of Vienna University of Technology turning the university's research into useful products by combining work on the analysis of music with innovative ways of visualising it.

In particular what sonarflow does is provide a way to browse tracks of music visually based on how similar they sound. When you start to browse you are given a visual overview of the available content shown as "bubbles" that roughly cluster the tracks into groups reflecting musical genres. Similar music is automatically clustered into groups but those groups are clustered into groups themselves, and so on, so once you have chosen a genre you can zoom into it to reveal ever more bubbles that group the tracks with finer and finer distinctions until eventually you get to the individual tracks. Any available textual information (i.e., metadata) about the music is also clustered with the bubbles to enrich the browsing experience and provide navigation hints.

At any time immediate playback of the objects is possible at any level. This means that you can choose between playing individual tracks or entire clusters of music with a similar sound or mood. Sonarflow can also be used to recommend new music because it is easy to explore the neighbourhood of well-known songs to discover new titles that have a similar sound.

All the problems may not have been solved as searching through millions of versions of anything isn't easy, never mind when you can't see the thing you are looking for, but by combining fun ways to view music with clever ways of recognizing the audio features, searching through those endlessly frustrating lists may soon be a thing of the past.

There are currently sonarflow apps for the web, iPhone and iPad - you can try a version of sonarflow for free from the Apple App Store. See also **www.sonarflow.com**

What's all this 'metadata'?

The Greek prefix meta is used to mean 'about'. So, metadata is 'data about data': who produced it, when, why, and so on. Metadata is commonly used in digital audio file formats like mp3 to describe the tunes. A song's name and the artist or band that composed it, are examples of metadata typically embedded in music files to identify the tracks. When you browse through alphabetical lists in online music collections you are reading metadata.

Sonifying zebrafish biology

Biologists often analyse data about the cell biology of living animals to understand their development. A large part of this involves looking for patterns in the data to use to refine their understanding of what is going on. The trouble is that patterns can be hard to spot when hidden in the vast amounts of data that is typically collected. Humans are very good at spotting patterns in sound though – after all that is all music is. So why not turn the data into sound to find these biological patterns?

In hospitals, the heartbeats of critically ill patients are monitored by turning the data from heart monitors into sounds. Under the sea, in (perhaps yellow) submarines, "golden ear" mariners use their listening talent to help with navigation and detect potential danger for fish and the submarine. They do this by listening to the soundscapes produced by sonar built up from echoes from the objects round about. This way of using sounds to represent other kinds of data is called 'sonification'. Perhaps similar ideas can help to find patterns in biological data? An interdisciplinary team of researchers from Queen Mary including biologist Rachel Ashworth, Audio experts Mathieu Barthet and Katy Noland and computer scientist William Marsh have been trying the idea out on zebrafish.

Why zebrafish? Well, they are used lots for the study of the development of vertebrates (animals with backbones). In fact it is what is called a 'model organism': a creature that lots of people do research on as a way of building a really detailed understanding of its biology. The hope is that what you learn about zebrafish helps you understand the biology of other vertebrates too. Zebrafish make a good model organism because they mature very quickly. Their embryos are also transparent. That is really useful when doing experiments because it means you can directly see what is going on inside their bodies using special kinds of microscopes.

The particular aspect of zebrafish biology the Queen Mary team has been investigating is the way calcium signals are used by the body. Changes in the concentration of calcium ions are important as they are used inside a cell to regulate its behaviour. These changes can be tracked in zebrafish by injecting fluorescent dyes into cells. Because the zebrafish embryos are transparent whatever has been fluorescently labelled can then be observed.

The Queen Mary team has developed software that detects calcium changes by automatically spotting the peaks of activity over time. They relied on a technique that is used in music signal processing to detect the start of notes in musical sequences. Finding the peaks in a zebrafish calcium signal and the notes from the Beatles' Day Tripper riff may seem to be light years apart, but from a signal processing point of view, the problems are similar. Both involve detecting sudden bursts of energy in the signals. Once the positions of the calcium peaks have been found they can then be monitored by sonifying the data.

What the team have found using this approach is that the calcium activity in the muscle cells of zebrafish varies a lot between early developmental stages of the embryo and the late ones. You can have a go at hearing the difference yourself – follow the links at **www.audio4fn.org** where you can hear the sonified versions of the data.

Audio Action!

Recordings of piano music played backwards sound like they are played on an organ, as the tones start quietly and grow in volume.

Moving forward by modelling the past...

Digital audio has brought the ease of manipulating sounds to everyone with a computer, but it seems musicians want things from the past not the future: the sound of the first 'Fender Stratocaster' guitar, the 'pre-EB Stingray' bass guitar, the 'Fairchild 670 compressor' and the 'Pultec EQP-1A' for example all make aficionados drool.

There is a huge market for recreating the sound of these classics and the most affordable way is using software rather than rebuilding the instruments. Two promising ways are described here. One way, 'physical modelling', is to build a mathematical model in software based on the hardware counterpart. This can de done in varying amounts of detail and sophistication ranging from modelling a section of a hardware device down to modelling the components used. When building the physical models the researchers measure how the original alters the signal of various audio sources. After this they create a mathematical formula to use in the software that best matches the performance of the actual hardware. It is important that they also do listening tests to fine tune the model as it is the 'sound' of the original that the user wants and not just a matching graph!

Compared to the original, creating an instrument in software with physical modelling has some distinct advantages: you are not limited to the number of pieces of kit you own, just click to add another plugin. Furthermore, software is usually a lot cheaper than hardware and needs none of the latter's hands on maintenance (well ok, so there might be an update occasionally to install!). There is one drawback to the modelled version in software though. Because the developers are always trying to match the characteristics of the original there always comes a point where the software designers say "that sounds close enough", but it's never 100% exact.

> "The most affordable way is to do it using software."

There is a second method that has been used by the company Focusrite called 'dynamic convolution' and a slightly different version by the company Acustica-Audio. It is based on the idea of taking an acoustical fingerprint of a room – extracting the acoustics of a room so that virtual versions of it can be created (see page 7). Dynamic convolution takes a fingerprint not of the room but of the gear in question for the different combinations of knob position and input level. As there are lots of combinations this leads to a huge amount of information being stored. Often only a few positions are recorded to save time and file size, but the more combinations stored the more accurate the result. Settings between those stored can also be reproduced by combining those from the knob positions on either side. There are at least 44100 samples for one second of CD quality sound and each distinct sample can use a different one of those virtually created settings. This method means that things like distortion and other characteristics of audio hardware can be reproduced better. Computers can only process so much information at once though so the accuracy of the fingerprints sampled, which in theory are perfect, has to be cut down to a point where it is deemed that people can't hear the difference. This is called 'truncation'. So like the modelling method there is also a point where the software is as close as possible, but it will never be 100% perfect.

Which method is best? There's no clear answer. Dynamic convolution is the most likely to deliver the closest sound to the hardware as measurements from the exact unit are used by the software, but it can be very CPU-intensive to use – it takes lots of computing power. That's why Focusrite sell it in a piece of hardware rather than as

software to run on your own computer. On the other hand, if done well, using physical modelling can come extremely close to the real deal and very often doesn't use much computing power at all so lots of virtual hardware can be used at once. Either way people always seem impressed and excited at a software version of a hardware classic, even though they have probably never heard (nor could afford!) the original which means they won't know if it's not 100% perfect anyway!

A soundly engineered waterfall

You don't need mikes, mixing desks and digital music to do audio engineering. In the 17th century they did it just with rainwater. At least that's what a 17th century French audio engineer called Grillet, did for the 1st Duke of Devonshire.

Grillet was hired by the Duke to create a water feature for his garden – not the kind of water feature you buy in garden centres nowadays but a gigantic architectural feature. After all, the garden in question wasn't out the back of a suburban semi, but the hillside garden behind Chatsworth: the grand Derbyshire stately home of the Devonshires. Grillet applied his previous experience working for the Sun King, Louis XIV of France to create a waterfall with an Audio twist.

The resulting 'Cascade' is an artificial stream, dropping a full 200 feet, from the hillside above down the garden behind Chatsworth House. It consists of groups of steps, each designed so that as you walk up the cascade the sound of the water subtly changes. Of course this wasn't done with some amazing advanced technology. After all it was the late 1600s so all that was available were stone-masons, his engineering know-how and a good dollop of creativity. He created the effect by structuring each group of steps slightly differently. He varied the number of steps in the group, their heights, distances apart, and also added features like ledges or rounded lips. The result is that the water makes a slightly different sound as it flows over each group. That means as you walk up the path besides the cascade, you find you are walking through a gentle, natural but designed soundscape. The cascade is much greener than most of our modern audio-engineering too. It produces its effect without the need for a power guzzling energy source as it is gravity-fed by rainwater from a man-made lake at the top of the hill.

Today's audio artists make use of modern technology to create all kinds of wonderful effects, but it is worth remembering that sound art can be created just by working with nature...even if you then wonder what you could do by combining digital effects with a cascade.

Eye wobbling hauntings

Sound doesn't just vibrate your eardrums, it can also potentially vibrate your eyeballs, causing blurry vision. Some scientists think that this type of inaudable, eye wobbling, low frequency bass sound, called infrasound, might be responsible for spooky hauntings. The stories of strange figures in a supposedly haunted laboratory in Warwick were, after investigation, discoved to be due to a newly installed extractor fan producing low frequency infrasound waves.

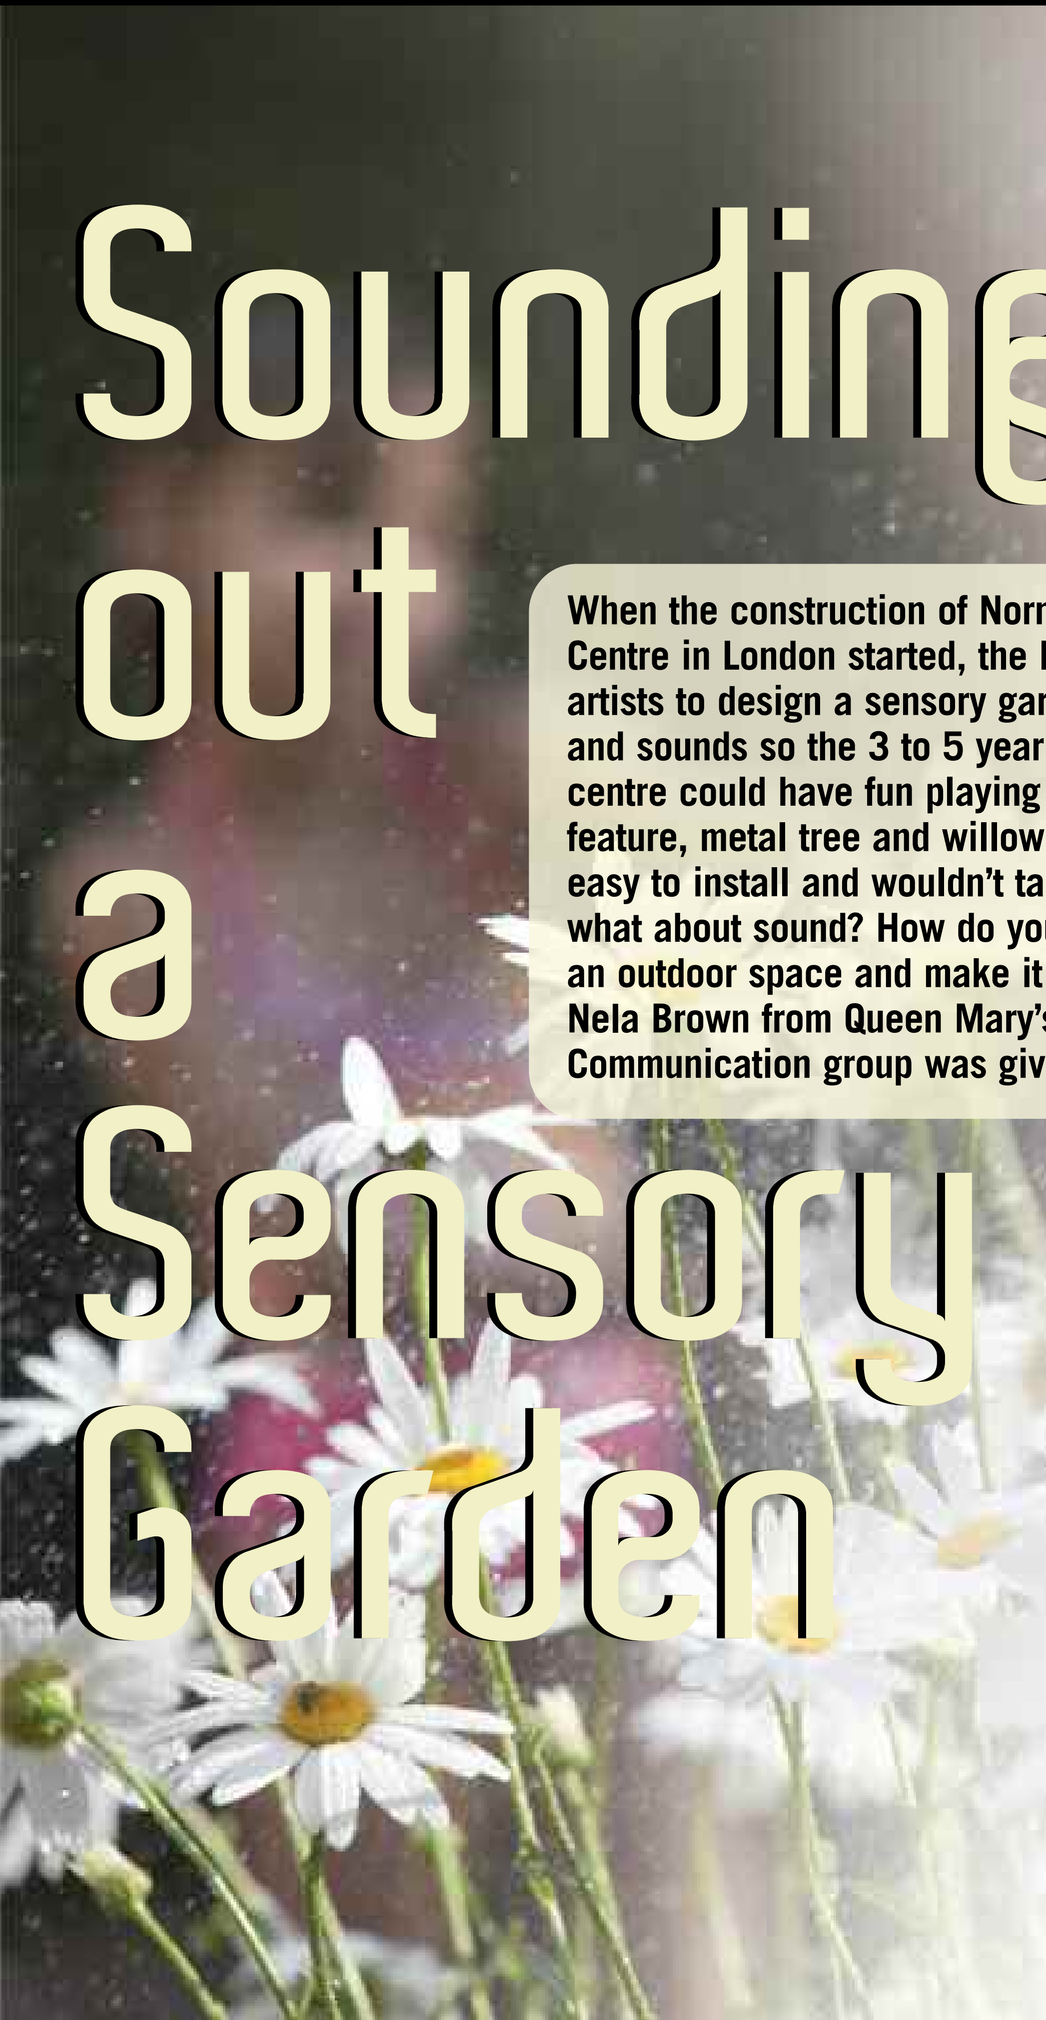

Sounding out a Sensory Garden

When the construction of Norman Jackson Children's Centre in London started, the local council commissioned artists to design a sensory garden full of wonderful sights and sounds so the 3 to 5 year old children using the centre could have fun playing there. Sand pit, water feature, metal tree and willow pods all seemed pretty easy to install and wouldn't take much looking after, but what about sound? How do you bring interesting sound to an outdoor space and make it fun for young children? Nela Brown from Queen Mary's Interactive Media and Communication group was given the job.

After thinking about the problem for a while she came up with an idea for an interactive sound installation. She wanted to entertain any children visiting the centre, but she especially wanted it to benefit children with poor language skills. She wanted it to be informal but have educational and social value even though it was outside.

You name it, they press it!

Somewhere around the age of 18 months, children become fascinated with pressing buttons. Toys, TV remotes, light switches, phones, you name it they want to press it. Given the chance to press all the buttons at the same time in quick succession, that is exactly what young children will do. They will also get bored pretty quickly and move on to something else if their toy just makes lots of noise with little variety or interest.

Nela had to use her experience and understanding of the way children play and learn to work out a suitable 'user interface' for the installation. That is she had to design how the children would interact with it and be able to experience the effects. The user interface had to look interesting enough to get the attention of the children playing in the garden in the first place. It also obviously had to be easy to use. Nela watched children playing as part of her preparation to design the installation both to get ideas and get a feel for how they learn and play.

Sit on it

She decided to use a panel with buttons that triggered sounds built into a seat. One important way to make any gadget easier to use is for it to give 'real-time feedback'. That is, it should do something like play sound or change colour as soon as you press any button, so you know immediately that the button press did do something. To achieve this and make them even more interesting her buttons would both change colour and play sound when they were pressed. She also decided the panel would need to be programmed so children wouldn't do what they usually do: press all of the buttons at once, get bored and walk away.
Nela recorded traditional stories, poems and nursery rhymes with parents and children from the local area, and composed music to fit around the stories. She also researched different online sound libraries to find interesting sound effects and soundscapes. Of the three buttons, one played the soundscapes, another played the sound effects and the last played a mixture of stories, poems and nursery rhymes. Nela hoped the variety would make it all more interesting for the children so keep their attention longer and by including stories and nursery rhymes she would be helping with language skills.


Photo: Nela Brown

Can we build it?

Coming up with the ideas was only part of the problem. It then had to be built. It had to be weatherproof, vandal-proof and allow easy access to any parts that might need replacing. As the installation had to avoid disturbing people in the rest of the garden, furniture designer Joe Mellows made two enclosed seats out of cedar wood cladding each big enough for two children, which could house the installation and keep the sound where only the children playing with it would hear it. A speaker was built into the ceiling and two control panels made of aluminium were built into the side. The bottom panel had a special sensor, which could 'sense' when a child was sitting in (or standing on) the seat. It was an ultrasonic range finder – a bit like bat-senses using echoes from high frequency sounds humans can't hear to work out where objects are. The sensor had to be covered with stainless steel mesh, so the children couldn't poke their fingers through it and injure themselves or break the sensor. The top panel had three buttons that changed colour and played sound files when pressed.

Interaction designer Gabriel Scapusio did the wiring and the programming. Data from the sensors and buttons was sent via a cable, along with speaker cables, through a pipe underground to a computer and amplifier housed in the Children's Centre. The computer controlling the music and colour changes was programmed using a special interactive visual programming environment for music, audio, and media called Max/MSP that has been in use for years by a wide range of people: performers, composers, artists, scientists, teachers, and students.

The next job was to make sure it really did work as planned. The volume from the speakers was tested and adjusted according to the approximate head position of young children so it was audible enough for comfortable listening without interfering with the children playing in the rest of the garden. Finally it was crunch time. Would the children actually like it and play with it?

The sensory garden is making a difference – the children are having lots of fun playing in it and within a few days of the opening one boy with poor language skills was not just seen playing with the installation but listening to lots of stories he wouldn't otherwise have heard. Nela's installation has lots of potential to help children like this by provoking and then rewarding their curiosity with something interesting that also has a useful purpose. It is a great example of how, by combining creative and technical skills, projects like these can really make a difference to a child's life.

The panels in each seat were connected to an open-source electronics prototyping platform, Arduino. It's intended for artists, designers, hobbyists, and anyone interested in creating interactive objects or environments, so is based on flexible, easy-to-use hardware and software.

Audio! Action

It's often claimed that there exists a particular sound frequency that makes humans lose control of their bowels. Appropriately enough it's called the brown note. However after repeated experiments there is no scientific evidence to support this effect.
The brown note is pants!

How to make Google beatbox?!

Can Google Translate make music? It turns out it can – it can beatbox!

Beatboxing is a kind of vocal percussion used in hip hop music. It mainly involves creating drumbeats, rhythm, and musical sounds using your mouth, lips, tongue and voice. So how on earth can Google Translate do that? Well a cunning blogger has worked it out and it's easy and fun to do. Once on the Google Translate page you first set it to translate from German into German. Next type the following into the translate box:

pv zk pv pv zk pv zk kz zk pv pv pv zk pv zk zk pzk pzk pvzkpkzvpvzk kkkkkk bsch;

When you click on the "Listen" button. Google Translate will beatbox.

Convincing?! Now, you can try to make your own funky beats and have the computer perform them for you…

So how do programs like Google Translate that turn text into speech do it? The technology that makes this possible is called 'speech synthesis': the artificial production of human speech. To synthesise speech from text, words are first mapped to the way they are pronounced using special pronunciation ('phonetic') dictionaries – one for each language you want to speak. The 'Carnegie Mellon University Pronouncing Dictionary' is a dictionary for North American English, for example. It contains over 125 000 words and their phonetic versions. Speech is about more than the sounds of the words though. Rhythm, stress, and intonation matter too. To get these right, the way the words are grouped into phrases and sentences has to be taken into account as the way a word is spoken depends on those around it.

There are several ways to generate synthesised speech given its pronunciation, information about rhythm and so on. One is simply to glue together pieces of pre-recorded speech that have been recorded when spoken by a person. Another way uses what are called 'physics-based speech synthesisers'. They model the way sounds are created in the first place. We create different sounds by varying the shape of our vocal tract, and altering the position of our tongue and lips, for example. We can also change the frequency of vibration produced by the vocal cords that again changes the sound we make. To make a physics-based speech synthesiser, we first create a mathematical model that simulates the way the vocal tract and vocal cords work together. The inputs of the model can then be used to control the different shapes and vibration frequencies that lead to different sounds. We essentially have a virtual world for making sounds. It's not a very big virtual world admittedly – no bigger than a person's mouth and throat! That's big enough to generate the sounds that match the words we want the computer to say, though.

These physics-based speech models give a new way a computer could beatbox. Rather than start from letters and translate them into sounds that correspond to beatboxing effects, a computer could do what the creative beatboxers do and experiment with the positions of its virtual mouth and vocal cords to find new beatboxing sounds.

Beatboxers have long understood that they could take advantage of the complexity of their vocal organs to produce a wide range of sounds mimicking those of musical instruments. Perhaps in the future Artificial Intelligences with a creative bent could be connected to physics-based speech synthesisers and left to invent their own beatboxing sounds.

> "If you are mad about music technology, then visit www.audio4fn.org for more on the fun side of audio engineering."

This issue of Audio! is edited by Paul Curzon, Mathieu Barthet and Simon Dixon with contributions from Dan Stowell, Martin Morrell, Nela Brown, Peter McOwan, Panos Kudumakis, Emmanouil Benetos, Duncan Menzies, Alice Clifford, Thomas Lidy and François Grandemange. It has been supported by the EPSRC Musicology for the Masses project. Audio! is part of the cs4fn family of magazines on the fun side of science, maths and engineering. **cs4fn** is funded by EPSRC with support from Google and ARM.

EPSRC Engineering and Physical Sciences Research Council

Queen Mary University of London

www.audio4fn.org