Queen Mary
University of London

Issue 4

AUDIO!

Mad about music technology

Inside this issue

Audio engineering: from the ancients and the animals

Audio! isn't just about music but about audio engineering more generally. That includes anything to do with understanding and manipulating sound: the way we hear it, and the ways we (and even bats) don't.

In this issue we explore some of the science and engineering behind how we might one day make sonic invisibility cloaks and find out how a bunch of female hackers used teapots to mix audio engineering and art. We explore how to correct flaky X Factor singers and even how to make completely virtual singers.

Audio engineering isn't a modern invention. We therefore look at the feats of the engineers of antiquity who created audio effects without modern technology. We also explore animal audio engineering – asking how bats avoid deafening themselves with their own sonar.

Listening to the stars

Listening to stars sing doesn't have to mean listening to the latest talent show. The stars in the night sky sing too!

When you think of astronomy you probably think of the people looking into telescopes, beautiful images of distant nebulae, or maybe the stunning pictures of Saturn's rings taken by probes sent to explore them. Maybe if you are a real enthusiast you will have been up at dawn hoping for a break in the clouds to catch a glimpse of the transit of Venus. Astronomy isn't all about what we can see though. Stars make a lot of noise, and the Kepler spacecraft has recorded the sounds from over 500 sun-like stars so far. Even more surprisingly perhaps, it turns out you can weigh stars based on the sounds they make!

The movement inside a star makes them vibrate like a musical instrument and so make sounds. As Bill Chaplin from the University of Birmingham's School of Physics and Astronomy, who leads the international collaboration doing the research, explains: "If you measure the pitch of the notes produced by an instrument it can tell you how big the instrument is. The bigger the instrument, the lower the pitch and deeper the sound." A big double bass makes a deeper sound than a cello, which in turn is deeper than the much smaller violin. Similarly a big star will make music with a deeper pitch than a small one.

As one of the scientists involved, Graham Verner of Queen Mary, University of London has observed, being able to weigh stars by listening to their songs is exciting because it means we can create a much more accurate picture of the stars in our galaxy.

> "You can weigh stars based on the sounds they make!"

For a BBC news video which includes the sounds of stars, follow the links on the Audio! website.

Naturally splendid acoustics

Stadium concerts are only possible because of modern loudspeaker technology. It means that even people far from the stage can hear the band. Clever as we are, the ancient Greeks solved the problem without the technology. They just used the seats. **Dimitrios Giannoulis**, a researcher at Queen Mary, University of London explains.

The ancient amphitheatre of Epidaurus in Greece has extraordinary acoustics that left researchers baffled for years. The theatre's design somehow made it possible for people at the back to clearly hear the play despite being more than 60 metres from stage. How was it done? Is there something special about the slope of the amphitheatre? Is it something to do with the shape channelling the wind in a way that drives the sound forward? In 2007, researchers from Georgia Institute of Technology finally cracked the problem. Its acoustic excellence is all down to the size, material and row spacing of the seats.

The seat rows play the role of what is known as an 'acoustic filter'. They make different sounds scatter in different directions depending on their frequency. The result is that some sounds, here the lower frequencies, are 'damped': they become much quieter. That means the high frequencies sound louder in relation to the low ones. Filtering out the low frequencies in this way removes background noise such as the sound of the wind, rustling trees and the murmuring of the audience. There is a further boost to the sound reception too, because the higher frequencies also reach the audience by being scattered back from the seats behind. Together these effects improve the clarity of a speaker's voice throughout the theatre, giving it its renowned acoustics.

For the amphitheatre of Epidaurus, the threshold frequency that matters is 530Hz. Lower than that and the sounds are damped. Higher frequency sounds are amplified. That raises a problem though. The lower frequencies of the human voice fall into this lower range too. Why isn't the speaker's voice filtered away? It turns out not to be a problem because our brains cleverly reconstruct the missing sound of our voices from the high frequency information available. This same effect is often used today to create the illusion of bass in sound systems that are not actually capable of producing it.

The acoustics of Epidaurus are unquestionably special and that's why the architects of later ancient Greek and Roman theatres tried to copy the seat row design. Perhaps we should think about doing the same more often too.

“Its acoustic excellence is all down to the seats.”


Image by Fingalo: CCA-SA 2.0 German

Olympic acoustics

The Olympic Stadium in Stratford London Olympics was designed specifically to be loud. Unlike theatres where the aim is to bounce the sound from the stage to the audience, the shape of the Olympic roof was designed to bounce the roar of the crowd back to the athletes.

The Melody Triangle: surprising music

Humans quite like surprises, nice ones at least. We go to lots of effort to make birthday presents a surprise. The best jokes work because their punch lines surprise us and the greatest magic tricks are most magical when something we didn't expect happens at the end. So does the same apply to music? Do we prefer music that is surprising and if so which ways make the nicest surprises? Can we compose interesting music based on its predictability?

Researchers at Queen Mary, University of London have been exploring the links between surprise and music. In particular they have been looking at how one of the most important theories of the information age, Claude Shannon's information theory, can be used to help make music.

Information theory dates back to the 1940s and gives a way to measure information. It was originally invented to explore questions like: how far can we shrink down a message and still restore the original? Being able to shrink data and then get it back is important for doing things like downloading movies quickly. The more the movie can be compressed the quicker it can be downloaded. Information, movies included, is stored in computers as sequences of 1s and 0s. A single piece of information that could be 1 or 0 is called a bit. Information theory boils down to a question of how many bits do you really need to store a piece of information. It turns out that is the same as asking how uncertain a prediction is, and that is linked to how surprising something is.

Gambling surprises

A key idea of information theory is 'entropy'. It's a way of actually putting a number to how certain a prediction is. If a gambler predicts the toss of a coin, he is choosing one of only two equally likely possibilities: the entropy is low. If he were to predict the top card of a shuffled pack of cards he is less certain of the prediction, now choosing one of 52 possibilities: the entropy is higher. Predicting the next character in a message can also be measured in terms of entropy, as can predicting what comes next in a piece of music.

That's where the Queen Mary project comes in. They are exploring what information theory can tell us about how people listen to music. They are studying several different kinds of predictability, and so potential surprise, in musical patterns: how we perceive them, and how they shape our experience listening to music. If the music is very easy to predict after a short period of listening, are we more likely to find it boring? If the notes that come next are more surprising, will we find it more interesting?

How entropy changes through a piece of music is one kind of predictability they are looking at. Another is 'redundancy' – is the same information available throughout the music. The more redundancy, the more that long sequences of notes repeat. A third kind of predictability they consider is about how our uncertainty about the music remaining changes as we hear more music.

Making music

They have created a way for non-musicians to make music based on these ideas. Their 'Melody Triangle' randomly generates new pieces of music given values for the different kinds of predictability. Different points on the triangle correspond to the different values of predictability. For example, one corner of the triangle is maximum repetition – a single note repeating over and over. Another corner represents noise – where the next note is completely random and cannot be predicted – there is maximum entropy. The final corner represents periodic music where longer series of notes are repeated – maximum redundancy. Points in the middle of the triangle represent trade-offs between those extremes.

> "If you want to try the Melody Triangle app follow the links from ee4fn.org"

You can compose melodies by moving tokens around the Melody Triangle. By adding different tokens you can combine piano, drums and bass, as well as different rhythms to make novel, interesting musical textures. There is an Android mobile phone app version you can download and try yourself, and the team also created a Kinect version where different people became the tokens themselves, making interweaving melodies and changing them as they move around a triangle on the floor, and where arm signals change the tempo or instrument.

The team have gathered some initial feedback from people using the Kinect version of the Melody Triangle. It suggests it could be useful as a performance tool or composition aid perhaps as an idea generator. They are now using it to explore our attitudes to surprising music. It shouldn't be a surprise that a key theory about information that is more than a half a century old is driving new music research. That is the way research works.

Audio ! Action

The frequency of sound waves, that is, the number of times they repeat per second, is measured in Hertz (Hz). These are named not after the car rental company, but after the German scientist Heinrich Hertz, who tragically died aged only 37. He did however have a crater on the far side of the Moon named after him. So that's all right then.

Hacking teapots

Tea drinking – not the most obvious thing to inspire technological innovation. Despite that, teapots were the focus for 20 female interaction designers, sound artists, bio-artists, architects, performance artists and technologists who met at the Victoria & Albert Museum in London.

The women involved were challenged to mend, modify or enhance an ordinary teapot, turning it into something new. Teapots created for the event included a Suicidal Teapot that hits itself with nails and a group of teapots that moved in synchrony as a flock. There was also a bio-hacked tea cosy made of fabric covered with mould and bacteria grown from tea, biscuits and cake, interacting with the antibiotic effects of green tea.

G.Hack, a women-only hacking club founded by electronic engineering and computer science research students at Queen Mary, University of London, were also there. They wanted to do more than just hack a teapot. They wanted a social adventure too. So armed with audio recorders, headphones, video cameras, iPhones and digital cameras, they set off on a mission to record the soundscapes at London's favourite teahouses, enjoying the centuries-old tea-drinking culture as they went.

The result was a tea table covered with a map of London incorporating a recording from a teahouse in each of the 33 London boroughs. They marked the teapot with an infrared LED allowing it to be tracked by a computer vision system. As the teapot is moved around the map, the vision system works out which borough of London it has been placed on and plays the related teahouse soundscape.

The resulting exhibit went on display at the Victoria & Albert Museum. It was the process of creating it that will be most memorable for the G.Hack team, though. It was both physically and technically challenging to create, but in the end they delivered an intriguing interactive experience. It was popular with the museum visitors, inspiring many to share stories about their own favourite teahouses.

Tea: not just calming, it can be creative too!

The Red Special: creating your own unique sound

What does it take to get started on the road to being a major rock star? Some skill in singing or playing an instrument and a unique sound, of course, but an interest in electronics might help too. That's what got Brian May of Queen up and running. His guitar, 'the Red Special', was unique. It defined his style and he has used it when recording and touring throughout his career. He didn't buy it though. He made it himself with help from his dad. You could make your own electronic instrument too if you wanted...

Brian May decided to make an electric guitar partly because he couldn't afford to buy one and partly for the fun of it. He wanted one with a unique sound that was better than anything that existed at the time.

What do you need to make a serious musical instrument? Well, for May's the parts were scavenged from junk and things around the house. The wooden parts were from a mantelpiece being thrown out, so his guitar comes complete with wormholes. He also used bits of shelf edging, a knife for the tremolo, some old motorbike springs, a knitting needle and some bits from a bike! For the electronics he adapted bought circuitry, such as using a commercial pickup that he modified to improve it.

A pickup is a gadget that converts the vibration of the strings into an electrical signal that can then be amplified. It's essentially just an electromagnet made of a coil of wire round a permanent magnet. When the string (a metal wire) vibrates nearby it affects the magnetic field of the magnet and that changing field creates a current in the coil. For the Red Special, May rewound the coil and covered it in glue as a way of reducing unwanted noise generated by the pickup, leaving just the good sounds.

> "Have a go at making our Arduinitar: an electric guitar with a difference"

As it's partially hollow, the Red Special is a semi-acoustic guitar that gave him a new sound. It was more tunable with more sounds than anything else out there at the time and also allowed him to control the feedback effects guitarists like Jimi Hendrix were having to force from their guitars.

You can now buy a replica Red Special if you have the money, but then that sound is no longer unique, of course. You'd need a more novel instrument to do that. May had help from his father to make his. If you want to have a go at making your own novel electronic instrument and need help getting started, why not have a go at making our Arduinitar first: an electric guitar with a difference. After that you could perhaps experiment and build something of your own design that has a completely new sound. Rock stardom may then be in sight. You will just then need to put in the hours and hours of practice to become a virtuoso musician.

To find out more about the Arduinitar and how to make it, go to www.audio4fn.org.

Pots fixing problematic acoustics

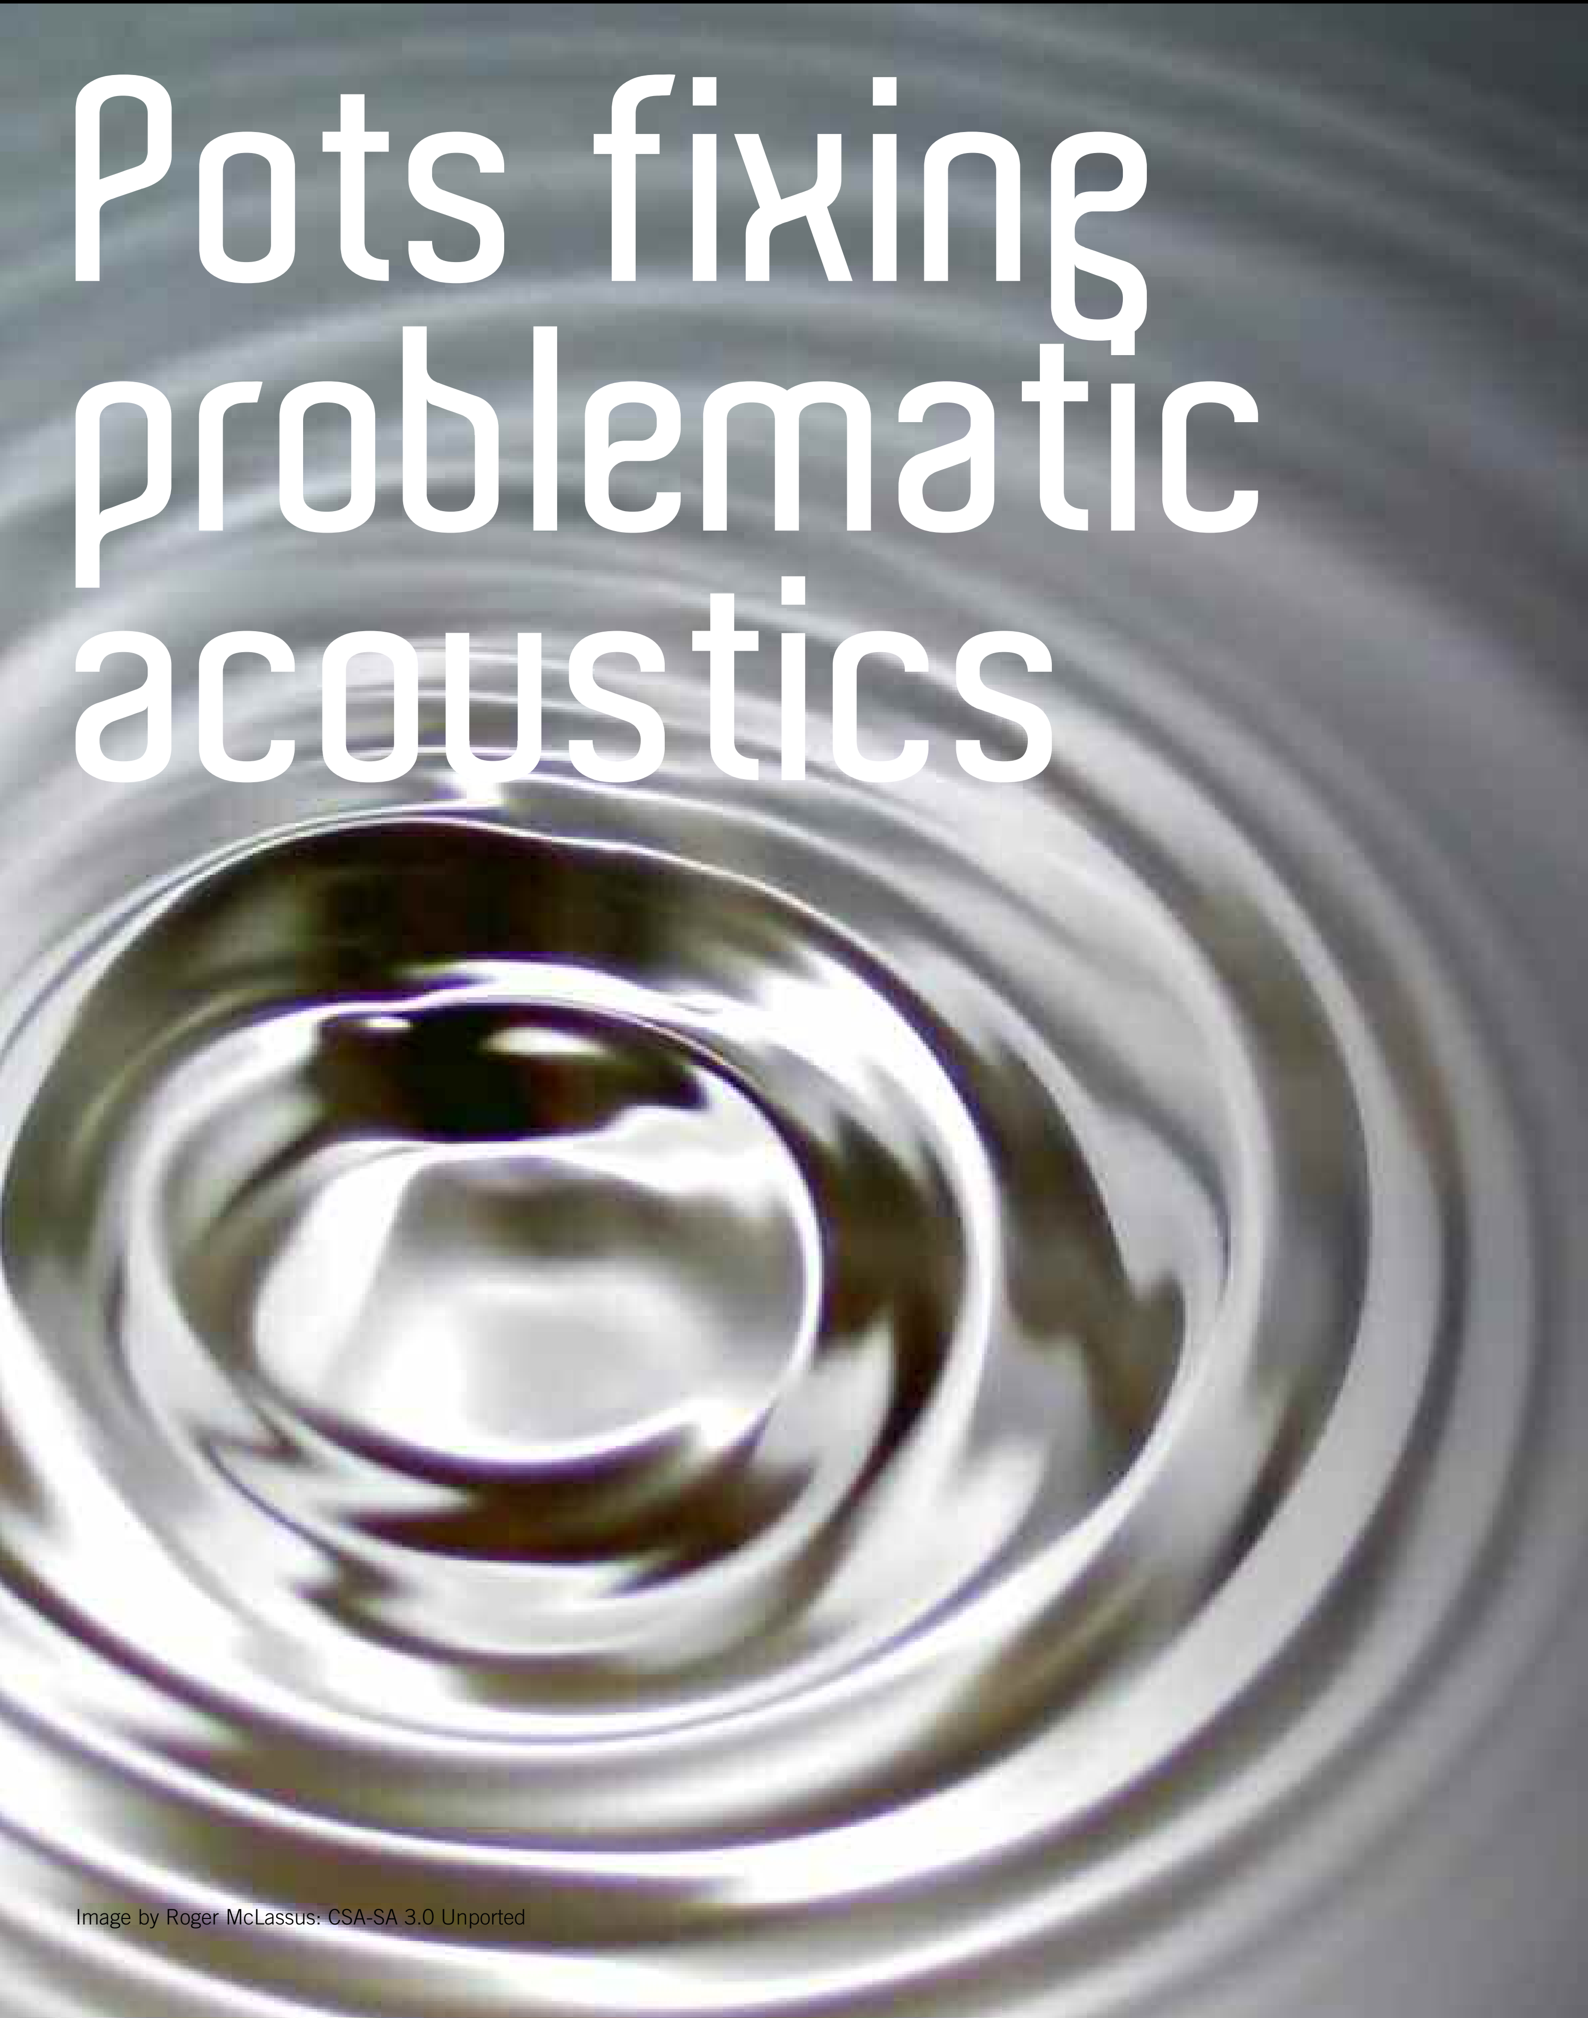

Image by Roger McLassus: CSA-SA 3.0 Unported

Pots are buried in the walls of medieval churches and monasteries across Europe: in the UK, Sweden, Denmark and Serbia. Why? Are they just a weird form of decoration? Actually, they are there to fix problematic acoustics. **Dimitrios Giannoulis** of Queen Mary, University of London explains.

The problem

First of all, what do we mean by 'problematic' acoustics? When sound waves move around a room they reflect off the walls in a way that creates strange sound effects when they meet their reflections.

It happens because of what are called 'standing waves'. Imagine dropping a pebble into a bath. The ripples create patterns in the water where they interfere with those that have bounced off the sides. As the two ripples pass in opposite directions, if the movement pushing the molecule up from one ripple exactly cancels out the movement pushing down from the other and keeps doing so, then at that point the molecules remain still. On either side the two ripples reinforce each other rather than cancelling out, giving the peaks and troughs of the combined wave. The result is the ripples appear to stop moving forward: a standing wave.

Sound waves are like water waves except that the air molecules vibrate from side to side rather than up and down as water molecules do. The same effects therefore happen when sound waves meet and standing waves can form. This is bad for two reasons. Standing waves take more time to die away after the sound source has been silenced than other sounds. Worse, the sound's volume varies around the room depending on whether it is a point where the waves cancel out (no sound) or where they enhance each other (loud). That's 'problematic' acoustics!

These acoustic problems ultimately come about because of what is known as 'resonance'. That is where a sound repeatedly bounces back and forth across a space at a particular frequency. Frequencies that are directly tied to the room's dimensions cause most problems. Called the 'resonant frequencies' they involve a whole number of wave troughs and crests fitting in the space between the walls. That is what leads to standing waves as the original and reflected wave coincide exactly. The lowest resonant frequency of a wave is also called the 'fundamental frequency'. It's the one where a single wave (a single trough and crest) fits in the space.

There are three different types of resonances developed in a room from sounds bouncing off the walls, called axial, tangential and oblique modes. Axial modes result from a sound bouncing back and forth between two facing walls. Tangential ones happen when the waves reflect around all four walls. Oblique modes are the most complicated and result from sound bouncing off the roof and floor too. Of all these, it turns out the worst are the axial modes. To improve the acoustics of a room you need to absorb the sounds at these resonant frequencies. But how?

The solution

OK, now we know the problem, but how do we deal with it? A solution is the 'Helmholtz resonator', named after a device created by Hermann von Helmholtz in the 1850s as part of his studies to identify the 'tones' of sounds. A Helmholtz resonator is just the phenomenon of air resonating in a cavity. It is the way you get a tone from blowing across the mouth of an empty bottle. The frequency of the tone is the resonant frequency of the bottle. If you change the volume of the air cavity or the length or diameter of the neck of the bottle you change its resonant frequency and so the tone.

A Helmholtz resonator actually absorbs sound at its resonant frequency and at a small range of nearby frequencies. This happens because when a sound strikes the resonator's opening, the air mass in the neck starts to vibrate strongly at that resonant frequency and tries to leave. That makes the pressure of the air in the cavity lower than the outside. As a result it draws the air back into the cavity. This process repeats but energy is lost each time, which causes the wave of this particular resonant frequency to dissipate. That means that specific sound is absorbed by the resonator. Helmholtz resonators also reradiate the sound that is not absorbed in all directions from the opening. That means any energy that wasn't absorbed is spread around the room and that improves the room's acoustics too.

So back to those pots in the walls of medieval churches. What are they for? Well they would have acted as Helmholtz resonators so they presumably were designed to remove low-frequency sounds and so correct the acoustic of the vaults and domes. Ashes have been found in some of the pots. That would have increased the range of sound frequencies absorbed as well as helped spread the unabsorbed sound. St Andrew's Church in Lyddington, Rutland, built in the 14th century, has some of the finest examples of this kind of acoustic jars in the UK. Helmholtz resonators obviously predate Helmholtz, actually going back to the ancient Greeks and Romans. The pots in churches are thought to be based on the ideas of Roman architect Vitruvius. He discussed the use of resonant jars in the design of amphitheatres to improve the clarity of the speakers' voices.

Designers of acoustic spaces like concert halls now use a variety of techniques to fix acoustic problems including Helmholtz resonators, resonant panels and tube traps. They're all efficient ways for absorbing low-frequency sounds. Helmholtz resonators though have the particular advantage of being able to treat localised 'problematic' frequencies.

Those church designers were apparently rather sophisticated acoustic engineers. They had to be, of course. It would have been a little unfortunate to build a church so everyone could hear the word of God, only to have those words resonate with the walls rather than with the congregation.

What bats can't hear

We humans are getting pretty good at creating technologies that extend our sense of hearing. We are still playing catch-up with bat technology though. Bats are able to pluck a moth out of the air at high speed in complete darkness. Everyone knows bats 'see' by sending out high-pitched sounds and listening for echoes returning from their prey – echolocation – but is it as straightforward as that? In particular, why do bats go to great lengths to make sure they can't hear themselves speak?

Bats somehow manage to turn the stream of echo information they hear into a mental image of the world around that is detailed enough to take action. They do it brilliantly, even working out not just where their prey is but what it is too. There is a more mundane problem bat evolution had to overcome first, before dealing with such sophisticated issues, though. How does a bat stop itself being deafened by its own shrill squeaks while listening out for those much quieter returning echoes?

A hardware solution

Many bats deal with the deafening problem by physically closing their ears as they make the shrieks. That raises a new problem of how they hear the echoes when they aren't actually listening. They get round that by sending lots of very short pulses of sound, listening for the echoes in the gaps.

Doing it in software

Another way has evolved in a different family of bats that is a little more sophisticated. Rather than separating their shrieks and the echoes in time they separate them by pitch. These bats send out a continuous sound and rather than closing their ears, listen constantly. Their ears are instead tuned very finely to particular frequencies – just the ones of the echoes not the frequency of the original shrieks. Just as we aren't deafened by bat calls as we can't hear them, the bats aren't as they can't hear their own calls either. The original sound is not at a frequency their ears can hear. They can hear the different echoes, though, as they have different frequencies.

Bunching up

It takes some serious audio engineering to get this kind of echolocation to work. These bats have to deal with something called the Doppler shift. It is the reason why sirens sound higher pitched as they come towards you than when going away. If an object is moving quickly towards you the sound source catches up a little with the crest of the previous sound wave before sending out the next one. This means the sound waves are bunched closer together than they would be if the sound source wasn't moving. If the wave crests are closer together, then their frequency – how often wave crests reach a person listening – is increased. A listener will hear a different note to that produced. Similarly when an object is moving away from the listener the last wave crest is that bit further away when the next one is emitted. The frequency of the sound is lower as the wave crests reaching the listener are more spread out. A similar thing applies if the listener is moving too.

Speeding problems

This matters to the bats as they are moving when chasing prey. When they listen for their echoes they are listening for sounds from a moving source as well as being a moving listener which means the echoes can be separated by frequency from the calls. That of course is what allows the system to work in the first place. However, the frequency of the echoes will change as the bats change their speed in flight. That could mean their finely tuned ears won't always hear the echoes. To overcome this they carefully change the sounds they send out to match their speed so the echoes, after the Doppler shift is taken into account, are in exactly the right frequency range to be heard.

Bats are thus not only brilliant at hearing but also brilliant at not hearing too. Next time you hear a tape of yourself and decide you don't like the sound of your own voice, think about the bats. They have the perfect solution.

> "Bats go to great lengths to make sure they can't hear themselves speak"

Can you hear that diagram?

What does a diagram sound like? What does the shape of a sound feel like? Researchers at Queen Mary, University of London are finding out.

At first sight listening to diagrams and feeling sounds might sound like nonsense, but for people who are visually impaired it is a practical issue. Even if you can't see words, you can still listen to them, after all. Audiobooks were originally intended for partially-sighted people, before we all realised how useful they were. Screen readers similarly read out the words on a computer screen making the web and other programs accessible. Blind people can also use touch to read. That is essentially all Braille is: replacing letters with raised patterns you can feel.

The written world is full of more than just words though. There are tables and diagrams, pictures and charts. How does a paritally-sighted person deal with them? Is there a way to allow them to work with others creating or manipulating diagrams even when each person is using a different sense?

That's what the Queen Mary researchers, working with the Royal National Institute for the Blind and the British Computer Association of the Blind have been exploring. Their solution is a diagram editor with a difference. It allows people to edit 'node-and-link' diagrams, like the London Underground map, for example, where the stations are the nodes and the links show the lines between them. The diagram editor converts the graphical part of a diagram, such as shapes and positions, into sounds you can listen to and textured surfaces you can feel. It allows people to work together exploring and editing a variety of diagrams including flowcharts, circuit diagrams, tube maps, mind maps, organisation charts and software engineering diagrams. Each person, whether fully sighted or not, 'views' the diagram in the way that works for them.

The tool combines speech and non-speech sounds to display a diagram. For example, when the label of a node is spoken, it is accompanied by a bubble bursting sound if it's a circle, and a wooden sound if it's a square. The labels of highlighted nodes are spoken with a higher pitched voice to show that they are highlighted. Different types of links are also displayed using different sounds to match their line style. For example, the sound of a straight line is smoother than that of a dashed line. The idea for arrows came from listening to one being drawn on a chalk board. They are displayed using a short and a long sound where the short sound represents the arrow head, and the long sound represents its tail. Changing the order they are presented changes the direction of the arrow: either pointing towards or away from the node.

For the touch part, the team use a PHANTOM Omni haptic device, which is a robotic arm attached to a stylus that can be programmed to simulate feeling 3D shapes, textures and forces. For example, in the diagram editor nodes have a magnetic effect: if you move the stylus close to one the stylus gets pulled towards it. You can grab a node and move it to another location, and when you do, a spring like effect is applied to simulate dragging. If you let it go, the node springs back to its original location. Sound and touch are also integrated to reinforce each other. As you drag a node, you hear a chain-like sound (like dragging a metal ball chained to a prisoner!). When you drop it in a new location, you hear the sound of a dart hitting a dart board.

The Queen Mary research team have tried out the editor in a variety of schools and work environments where visually impaired and sighted people use diagrams as part of their everyday activities and it seems to work well. It's free to download so why not try it yourself. You might see diagrams in a whole new light.

> "It allows blind people to edit 'node-and-link' diagrams like the London Underground map"

Simon says: no to Auto-Tune

The papers were alive with the sound of Auto-Tune when it was suggested that the talent show X Factor had used a digital sound processing technique called autotune to improve the sounds of the singers on the show.

The story raged for a few days. Was using computer technology to improve singing right? The argument was that it's all over the charts already. Singers from Kanye West to One Direction use it, so what's the problem? But the public was having none of it. Finally Simon announced there would be no more autotune in X Factor... but the question still remains: what is it and how does it work? So let's explore the most controversial tool in the music producer's arsenal.

Setting the tone

Auto-Tune is the name of a commercial software plug-in developed by Antares Audio Technology in 1997. In theory it looks for off-key notes in either vocal or instrumental tracks, and then slides these bum notes to where they ought to be. The computer science behind this is based on something with a far more interesting name: a phase vocoder. The fundamental idea of this technology is that any sound can be broken down into a sum of pure tones, much like any colour can be created by adding a mixture of red, green and blue. Sound is a pattern of air compressions, and pure tones are sounds that have a regular rhythmic repeating pattern. These air waves are defined by their repeating frequency (literally how often a new wave comes along), their amplitude (which is how strong the waves are) and a phase (when they start and finish relative to other waves).

Break it all down

You can take any sound signal, whether it's a voice or an instrument, and sample it in a computer. The microphone detects the changes in air pressure and turns them into a series of numbers that change over time. Then the cunning processing begins. You can take that bit of complex sound and break it down into its pure tones, so you know their frequency, amplitude and relative phases. So now you've got all the pure tones that make up one complex sound – let's call them A, B, and C. You could build the sound back up again, if you wanted, by adding them together: just add A + B + C and you're back to where you were.

The real X Factor

Let's stick with the bits for a moment though. Each of your pure tones, A, B and C, has a particular frequency and phase. The lowest frequency in the sample, called A here, is known as the fundamental frequency and this is what gives the sound its musical pitch. But suppose that segment of music you've processed is a bum note: it should have been a higher or lower pitch. The computer knows what the sound should have been and can change and slide the components A, B and C around to make sure the mix gives the sound you wanted in the first place. So for example, you could slide A back in time with respect to B and C, shifting the relationship between the frequencies. Or, you could replace A with the new correct fundamental frequency – let's call it X. Then the computer system adds X, B and C to produce the pitch-corrected sound in almost real time.

The numbers game

Whether it's music, images or video, once information enters the digital domain it is just a pile of numbers. As such, it can be manipulated, improved or faked. What we decide to do with this way to change reality is up to us. It can make the world better, be wonderfully creative or even let some people cheat. Like any technology, it's the way we use it that's important.

Soundly Invisible

Making things invisible is the stuff of science fiction. Our fascination is with the idea of not being able to see things that are before our eyes, but invisibility in a different sense may be useful too. Harry Potter's invisibility cloak allowed him to disappear at will, but the sounds he made could easily have given him away when he was wandering those quiet night-time corridors. He needed to be 'acoustically invisible' too: he needed the sounds he made to disappear. In fact, researchers have been working on the technology to do that for almost a century and it hasn't been totally science fiction for a long time. Since World War II, technology has existed to make submarines all but disappear from sonar. A new material may even make it possible to make an invisibility cloak that makes things truly acoustically invisible.

Deathly silence

Back in World War II, U-boats were central to the German superiority in the Atlantic, but sonar was giving them away. Remember all those naval war films where everyone in the submarine falls silent as we listen to the pings of the sonar? The German Kriegsmarine badly wanted technology to make their U-boats invisible to those deadly pings.

There are two kinds of sonar: active and passive. In passive sonar, the submarine killer just uses underwater microphones to listen for any noise coming from the submarine to work out where to drop the depth charges – that's why they shut off the engines and everyone falls silent in the movies. Any sound could give them away. Active sonar removes even this possibility of escape by sending out sound waves – the sonar 'pings' – and listening for their echoes. That way, even a totally silent sub can be found. When the pings hit it they bounce back to the listeners, revealing the sub's location in the same way that bats find their prey.

Bubble, bubble

The technology the Kriegsmarine developed involved covering the submarines with special tiles. These so-called 'anechoic' tiles are covered with lots and lots of small bubble-like holes. They absorb the sonar pings, turning the sound energy into heat rather than reflecting it. That means the ping doesn't return to the submarine killer, giving the game away. Different sound frequencies at different depths are best distorted by different-sized holes. That means the tiles have to have a specially engineered mixture of holes if the submarine is to stay hidden. They can also be designed to absorb specific sounds of the submarine like the engine noise, or maybe even that spanner dropped by the clumsy sailor at the critical moment in the films! World War II ended before the technology was perfected, but the Soviets went on to make the idea work during the Cold War. Now everyone's military submarines are covered with anechoic tiles.

This kind of technology works well for submarines because there is only deep ocean and so silence around them. Up on the surface, silence itself could give the game away. Imagine a person walking through a park covered in an acoustic cloak. The park is full of birdsong and the whispering of the leaves. As the person passes by, the cloak absorbs the sounds hitting it and so a person they pass would hear the birds and trees beyond momentarily fall silent. It is the sonic equivalent of a person being invisible to the eye but still leaving a shadow – just a shadow of no sound rather than no light. A ghostly silence would pass by.

Bending sound

True acoustic invisibility involves the sounds around a cloaked object not just being absorbed but apparently passing through the object as though it wasn't there. How could that be done? Well, essentially the sound waves would need to be bent around the object. The cloak would need to take the waves in and spit them out the other side in exactly the direction, not to mention frequency and volume, they had been originally traveling.

José Sánchez-Dehesa and Daniel Torrent at the Polytechnic University of Valencia have invented a material that may one day be the

> "A ghostly silence would pass by"

basis of cloaks able to do just that, and it isn't actually all that hi-tech. It involves sheets of perspex covered with a hexagonal pattern of holes. Whilst the material itself isn't that hi-tech, it was actually inspired by some rather hi-tech material physics: that of graphene. Graphene is a form of carbon where the carbon atoms are laid out in sheets of hexagons. Electrons move through graphene in an unusual way, being drawn into the hexagons as they flow over the surface. That inspired the Spanish researchers to see if sound would behave in the same way as electrons... and it turns out it does. Rather than just distorting sound, turning it into heat, as the rows of holes in the early anechoic tiles do, the special hexagonal pattern means the sound is drawn down into them unchanged. That's what a sonic invisibility cloak would need to do, pull the sound into holes around it then transport it down specially positioned tubes that release it unchanged at the other end to continue on its way. Whilst a cloaking device hasn't yet been engineered out of a material like this, in principle at least it could, so acoustic cloaking may one day be a reality.

In future, miscreants wandering school corridors after curfew may not have to worry about leaving a ghostly silence as they invisibly pass by – if of course we can make them invisible to the eye too... but even that isn't as science fiction as you probably thought.

Audio! Action

The pipe organ laid claim to being the most complex device ever made by humans until the telephone exchange was invented.

The truly virtual singer

Hatsune Miku is a well-known singer in Japan... but she is a cartoon. **Dan Stowell** tells us how her voice was made by researchers in Barcelona.

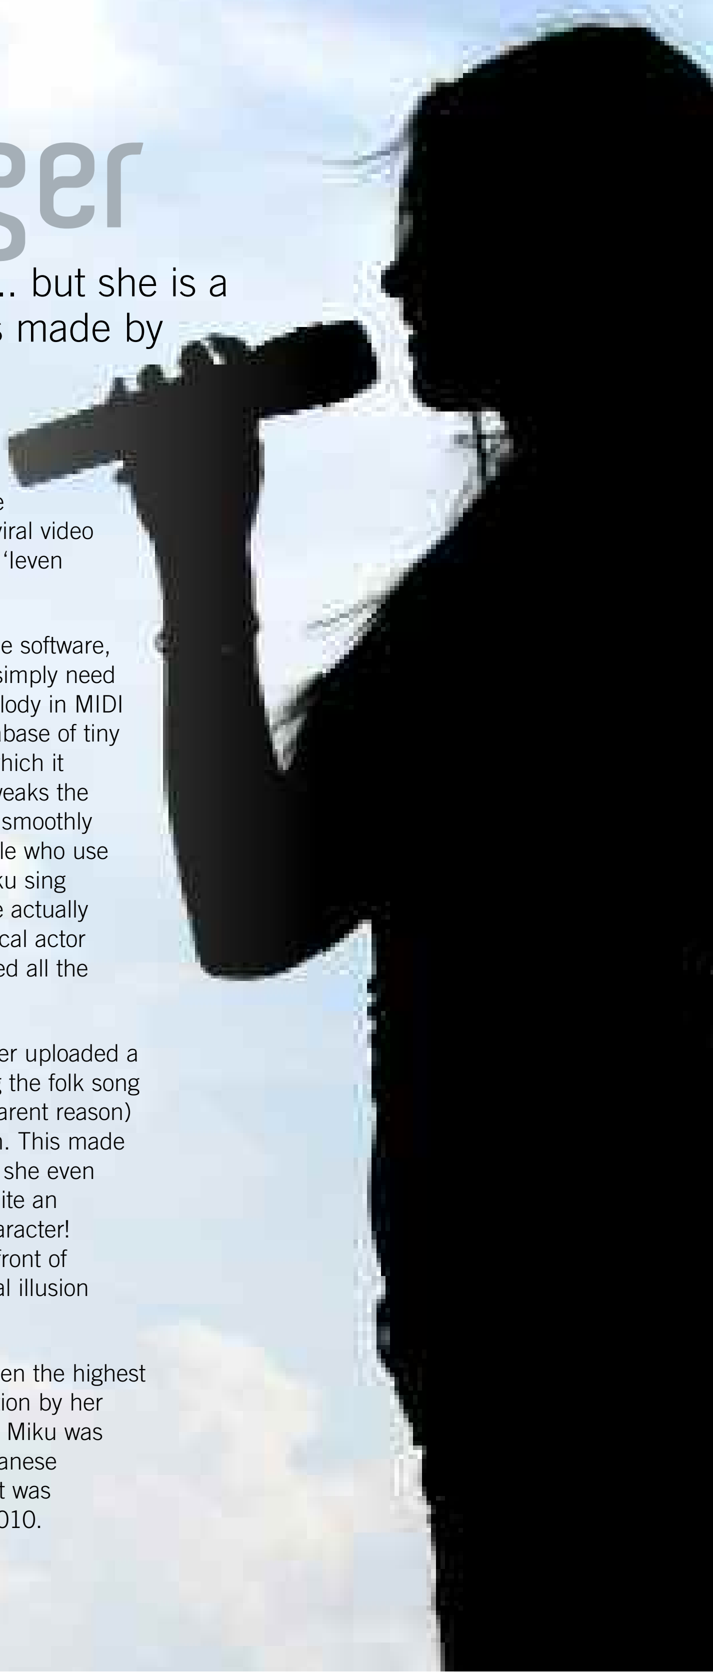

In the UK we have one cartoon band – Gorillaz – but the Japanese are more used to pop stars who happen to be animations. Hatsune Miku is different: her singing voice is generated by a computer program sold by Yamaha. You can make her sing anything you like, which might explain why her biggest hit is a Swedish folk song in which she waves a leek!

A singing voice, just like a speaking voice, has a very complex sound. You've probably heard automatic speech on the phone, or on buses and trains, and you might have noticed it sometimes sounds unnatural. Synthetic singing is even more difficult to create, as it needs to have the right melody as well as sound like a voice. But researchers in Barcelona, at the University Pompeu Fabra, applied similar techniques as used for speech synthesis, and developed computer software that can sing any words you like, with any melody.

This ‘singing voice synthesis' was taken up by the music technology company Yamaha. Cleverly, they sold it not just as a piece of software, but as a set of vocal ‘characters' you can buy, each with their own visual appearance. Yamaha sell a variety of these characters, but Hatsune Miku is by far the most popular as a result of a viral video hit with the Swedish folk song ‘Ieven Polkka' in 2006.

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

To record a vocal track with the software, which is called Vocaloid, you simply need a file with the lyrics, and a melody in MIDI format. The system has a database of tiny fragments of recorded voice which it connects together, and also tweaks the sound so that the pitch varies smoothly in the right musical way. People who use Vocaloid to make Hatsune Miku sing might not realise it, but they're actually manipulating the voice of a vocal actor called Saki Fujita, who recorded all the voice samples.

When a Japanese Vocaloid user uploaded a video of Hatsune Miku singing the folk song and waving a leek (for no apparent reason) it was a big hit online in Japan. This made Hatsune Miku so popular that she even performed live in concert – quite an achievement for a fictional character! She appeared live onstage in front of hundreds thanks to a theatrical illusion called Pepper's Ghost.

The live appearance wasn't even the highest point of her fame. After a petition by her fans, an engraving of Hatsune Miku was even included on board a Japanese spacecraft called Akatsuki that was launched into space in May 2010.

This issue of **Audio!** is edited by Paul Curzon and Jonathan Black with contributions from Dan Stowell, Nela Brown and Dimitrios Giannoulis. It has been supported by the EPSRC via projects including the Digital Music Platform grant and Information Dynamics of Music grants. 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.