Queen Mary
University of London

Issue2

AUDIO!

Inside this issue

Video game audio

Keeping conversations secret

Digital songwriting

MAD about music technology

Merely making music?

Music – that's an arty subject isn't it? Why would budding musicians be interested in electronic engineering, physics and computer science? Well classical music still rules for many but electronic music is where all the action is.

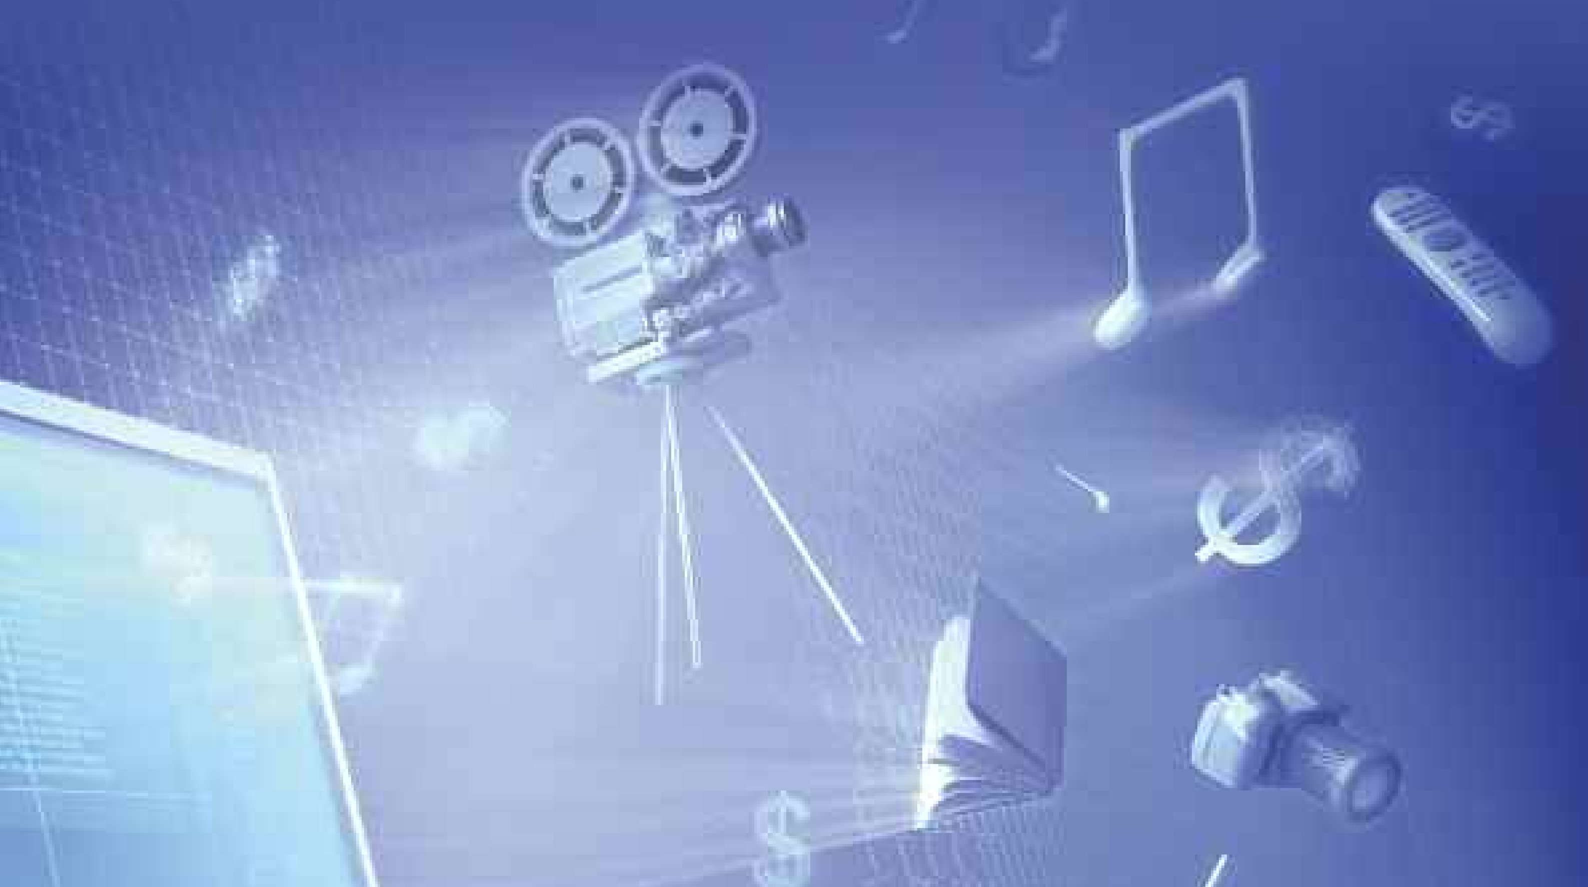

Audio engineers are interested in doing more than just playing music. They want to create new sounds and even new instruments to make them. It is audio engineers who are responsible for writing systems that, for example, allow out-of-tune pop prima donnas to actually sound great, and allow inventive musicians to create new sounds. Mere musicians might strive to compose new music but audio engineers aim higher – they might aim to write programs that can compose for them!

Audio is about more than just music though. Once sound has gone electronic it opens up new ways for creating video game soundtracks, creating bubbles of sound in which your secret conversations are safe, and quickly spotting ill animals (and maybe ill people too). Even the team trying to beat the land speed record need to know about the physics of sound if their driver is to survive the experience.

Audio engineering is an exciting place to be right now. Read on to find out more.

The classical composer browser

Just as the original World Wide Web links documents with other documents, the Semantic Web links data with other data. It is the next step in the evolution of the web towards the creation of an 'intelligent web' where the web starts to understand the material it is made up of.

Some of that linked data on the Semantic Web is about music, so scientists at the Centre for Digital Music at Queen Mary have been working out the best ways to describe and interact with such music-related data. One interface allows users to visually explore the connections of influence between classical composers. Because the work of Fredrick Chopin had an influence on the work of Claude Debussy, we have a connection between these two composers. Such connections can be used to create a "composer space" where similar composers are found nearer to each other in the visual picture of the composer space. An interested listener can then follow the connections between composers that are near to each other to learn about new and interesting music.

Browsing the musical web is about to get a whole lot cleverer!

Video game audio

Video game audio has come a long way since Pac-Man and Super Mario Brothers, yet the sounds created by those games have become iconic. Video games are now a multi-billion pound industry and as video game graphics are getting more realistic, the audio for video games is also advancing. Players are expecting sounds that make them believe they are in the middle of the action.

Not like a day at the movies

Movies have a set plot that the composer and sound effects engineers work hard to enhance. They want you to be interested in the story and feel empathy with the characters you see on the screen. Whilst video games also have plotlines and characters, the creators don't want you to be merely empathetic, they want you to experience the game as though you were part of the story.

Movies are something you just watch, but video games are interactive. That means that the sequence of sounds the player hears in the game depends on what the player does and can't be controlled by the game developers. You can watch your favourite DVD over and over and it will be the same every time, but game creators want you to have a new experience each time you play their game. They don't want you to be bored, after all!

Start your engines

Unlike movies, games constantly need new sound effects that respond to what the player does. One example of a sound heard constantly throughout a video game is the sound of the car engine in a racing game. The player needs to be able to hear the engine revs and gear changes. You can also play racing games from the viewpoint of sitting inside the car or alternatively be racing whilst outside and above the car. The audio needs to reflect these changes in the environment and the engine.

A recording engineer might record the sounds of a car engine at different RPMs from idle to full throttle, accelerating and decelerating. The engine might also be recorded from different angles, perhaps in front, at the sides, from the back of the car and from the driver's seat. The audio programmer for a game then instructs the gaming console how and when to play the recorded samples according to what the player is doing in the game.

Alternatively, a common trick to reduce the number of sound recordings, or samples, needed is to modify the ones you have during gameplay, so that the same recording can be used to generate a range of different sounds. The samples can be changed with filters to modify the frequency content and then can be mixed together in varying combinations to create an array of sounds.

An alternative to storing a sound recording is to use a computer to synthesise sounds. Instead of recording a car engine, a computer uses maths to model what it sounds like at different speeds. This saves on computer memory, which is important, as gaming consoles have limited space to store sound recordings. Using this computer simulation trick also allows a lot more flexibility, as the sounds aren't limited to what was recorded. They have the disadvantage though of requiring a faster processor, to do all the extra computation. That trade-off often happens in computing – to use less memory, you need more speed and vice versa.

Games of the future

So the next time you play your favourite video game, try turning the sound off and listening to the difference it makes. The gaming experience is much more than just the graphics. If programmed well, the sound effects and music make the adventure more immersive and believable. So believable in fact that if done well you'll stop listening and just be sucked into the game.

Getting that digital voice

Anyone who's listened to chart, pop, rap, R&B or dance music in the past few years will know about the influx of songs where the voice sounds digital and slides up and down between the notes. Some big stars including Kanye West, T-Pain and Rihanna are using this on their records. These are all great singers, but with this effect people really don't have to be. They just use a plugin on their computer called Auto-Tune. It was originally used by recording engineers to correct any tuning problems in the singer's performance.

Oops! Sorry we let the secret out! Yes, your favourite star may not actually be able to always sing in tune! The engineers fix it for them, tweaking the pitches graphically to what the notes should have been so that we listeners don't realise the vocal was a bit off-par.

Auto-Tune also has an automatic mode, and somewhere along the line someone realised that the digital side effects of using the auto mode were not actually unpleasant. In fact they created a good effect. Those familiar with musical techniques will realise that all it does is pushes the note up or down to be perfectly in tune with a certain musical scale.

Auto-Tune is an expensive piece of software and boasts that it is used by major recording studios worldwide, but there's actually a free alternative for those of us dreaming of being the next rap star or producer. It's called GSnap and is available from www.gvst.co.uk. According to the manual it can be used to subtly correct the pitch of a vocal, or, with more extreme settings, to create a robot-voice effect. That sounds like a free version of Auto-Tune to us!

GSnap has many more features. It allows you to set how strong the tuning effect is, what musical scale you would like to use and how fast the notes are moved up/down to the correct pitch. GSnap can be used in lots of your favourite Windows-based recording software, including Audacity (which is another free piece of music making software – see www.audio4fn.org for a download link). So give it a try on your new track, it may just add the effect that you need.

Audio ! Action

The first devices for recording sound could only produce a few copies – and the copying process destroyed the original. To get the most out of each session, a singer was often surrounded by as many recording machines as possible, and had to record the same song over and over to produce many batches of originals.

Synthetic speech

Computer-generated voices are encountered more and more frequently in everyday life, not only in automated call centres, but also in satellite navigation systems and home appliances.

Although synthetic speech is getting better, it's still not as easy to understand as human speech, and many people don't like synthetic speech at all. Maria Klara Wolters of Edinburgh University decided to find out why. In particular she wanted to discover what makes synthetic speech difficult for older people to understand, so that the next generation of talking computers will be able to speak more clearly.

She asked a range of people to try out a state-of-the-art speech synthesis system, tested their hearing and asked their thoughts about the voices. She found that older people have more difficulty understanding computer-generated voices, even if they were assessed as having healthy hearing. She also discovered that messages about times and people were well understood, but young and old alike struggled with complicated words, such as the names of medications, when pronounced by a computer.

More surprisingly, she found that the ability of her volunteers to remember speech correctly didn't depend so much on their memory, but on their ability to hear particular frequencies (between 1 and 3 kHz). These frequencies are in the lower part of the middle range of frequencies that the ear can hear. They contain a large amount of information about the identity of speech sounds. Another result of the experiments was that the processing of sounds by the brain, so called 'central auditory processing' appeared to play a more important role for understanding natural speech, while peripheral auditory processing (processing of sounds in the ear) appeared to be more important for synthetic speech.

As a result of the experiments, Maria drew up a list of design guidelines for the next generation of talking computers: make pauses around important words, slow down, and change to simpler forms of expressions (e.g. "the blue pill" is much easier to understand and remember than a complicated medical name). Such simple changes to the robot voices could make an immense difference to the lives of many older people. They will also make services that use computer-generated voices easier for everyone to use. This kind of inclusive design benefits everybody, as it allows people from all walks of life to use the same technology. Maybe Maria's rules would work for people you know too. Try them out next time grandpa asks you to repeat what you just said!

Audio ! Action

The older you get, the less you can hear high frequencies. That's why some people use a "mosquito tone" as a text message alert so their teachers can't hear!

Saving our

Pigs can't fly, and right now they are responsible for quite a lot of humans failing to fly too! Why? To try and stop the spread of swine flu. Recent research to help keep pigs healthy may soon catch sick people who do try to fly despite having flu. They could well be hauled out of departure lounges by the flu police working on tip-offs from computers.

What does flu have to do with sound, though? Well, people (and pigs) with flu cough, so spot the coughers and you can spot the flu. The trouble is that healthy people also clear their throats in a cough-like way and cough after choking on a sandwich that goes down the wrong way. If you can find a way to tell the different kinds of cough apart though, maybe a computer system could be designed to automatically recognise flu carriers.

Back in 1998 a team from Nippon Medical School in Tokyo showed that different kinds of cough sounds might be automatically recognised. They recorded a series of ill people coughing and compared them with a series of healthy participants who coughed voluntarily. They found that by analysing the coughing sounds they could tell the difference between different coughs. The telltale signature was in high frequency sounds in the middle part of the cough.

Sounding perky

Move on 10 years to 2008 and a joint team from the Faculty of Veterinary Medicine in Milan and the Department of Biosystems from Katholieke Universiteit Leuven in Belgium showed that a similar thing could be done with the coughs of pigs. It is vital that ill pigs are identified quickly before the illness spreads through the herd. Coughs are one of the main way farmers can do this. So could a computer system be programmed to tell the difference between a cough caused by a lung infection and one from a healthy animal? The team recorded pigs with pneumonia out in the field as well as healthy ones made to cough in a lab.

Rather than looking at just one feature of the sound, they looked at lots of different properties and found a series of clear differences between sick and healthy coughs. Healthy coughs last a few tenths of a second longer than infectious ones on average, for example. Similarly, healthy coughs come further apart than sick ones. The frequency of

> Put all the different properties together and you have the signature of an ill pig's cough.

the sound differs too: infectious coughs peak at a much lower frequency to healthy ones. Put all the properties where there is a clear difference together and you have the signature of an ill pig's cough.

Once the scientists have done their classification experiments, it's time for the audio engineers to step in. Their job is to turn this into a practical system, work that is still in progress. The aim is to create a classification program that constantly monitors the sounds from a pig house and can automatically hear when a pig is unhealthy, allowing it to be quickly treated.

Not to be sneezed at

Farmers may worry about their pigs, but the rest of us are probably more worried about the human form of swine flu. Right now it is invisibly spreading across the world as people move around. One holiday flight and a flu bug can travel from one side of the world to the other inside a day.

Some airports are trialling ways to stop ill travelers boarding flights. One method that has been partially successful is to use infra-red scanners to spot when a passenger leaving a plane has a fever. Engineers at BioRICS, a spin-off company from Katholieke Universiteit Leuven, where the pig research was done, think the pig approach could work in airports too. It wouldn't be looking for coughing pigs of course but coughing humans (a coughing pig wandering through Departures wouldn't be hard to spot, after all). If they can make it work their system will use microphones scattered around airport lounges that sound an alarm if they detect someone coughing in the wrong way.

So the coughing pigs may well save our bacon yet, and as long as the audio engineers get it right you won't be hauled off by the flu police just for clearing your throat.

Audio! Action

Barn Owls don't just have a pretty face. Their feathery facial discs act like reflectors and allow the Owl to precisely detect their prey in the dark from the sound signals the Owls send out.

Sounds funny in space

Space explorer Richard Garriott, recently returned from the International Space Station, is a second generation spacefarer. His dad, Owen Garriott, was a NASA astronaut who served aboard Skylab. Dad Owen had played a prank when he was in space by taking a recording of his wife, Richard's mum, which he played back over the radio link. That made it sound, to the people in ground control, that she had popped in to deliver a home cooked meal for the astronauts.

Richard decided to play the same trick but using newer technology. His computer contained a series of pre-recorded snippets of his mum talking, and as the conversation with ground control developed he was able to select the appropriate reply to send back to earth. The result was that they believed Richard's mum had again decided that a home cooked meal was needed in space. Mmm, tasty.

Sounds like magic

Three matchboxes lie on a table. You pick up the middle one and shake it.
You can clearly hear that it contains some matches rattling round inside. You pick up and shake the other two, and it's clear they are empty. You then move the boxes round, and ask your spectator to follow the filled box. At the end of the movements they are to say which box has the matches. They select a box, you shake it, and they are wrong. You shake a different box and everyone can hear that's where the matches are. You do it all again, but however closely they watch they can never find the box with the matches. Magic or cunning audio misdirection? Find out how on the Audio! website.

The sound of secrets

Shhh, it's a secret! Who knows who could be listening? There are many occasions when you might want to keep something a secret, from national security to business confidentiality. The problem is that it's easy to start off chatting in an open plan office or canteen about normal stuff and then drift into talking about confidential stuff, which we would rather remain between you and I. That's where, in the future, we might want to raise our 'conversational shields'.

Who is listening?

Researchers at the MIT Media Lab in the US have applied for a patent on a new way to keep secrets in the office. Called a conversational shield, the system works by having your office filled with hidden loudspeakers and sensors, perhaps built into the ducts that carry the power for computers round the office. If you are worried that what you're saying is better kept a secret you can simply call up your sound bubble to conceal your chat.

Building a sound bubble

Software locates all the other people in the office who are within earshot and then plays a masking sound through the loudspeakers closest to them. This masking sound can be a combination of general office background sounds mixed with white noise, a kind of special hiss that's good for covering sounds up but is rather annoying on its own. The researchers found that this bubble of masking sound greatly reduces the ability of any potential eavesdroppers to make out what is being said.

Moving on

Better still, your sound bubble can move with you. As you wander through the office the smart software can ensure that the area of masking sound follows, switching on loudspeakers to frustrate any new eavesdroppers who have moved into range. Your bubble is exactly where it's needed for maximum effect. You could say this software gives a whole new meaning to the idea of following a conversation!

Antisound

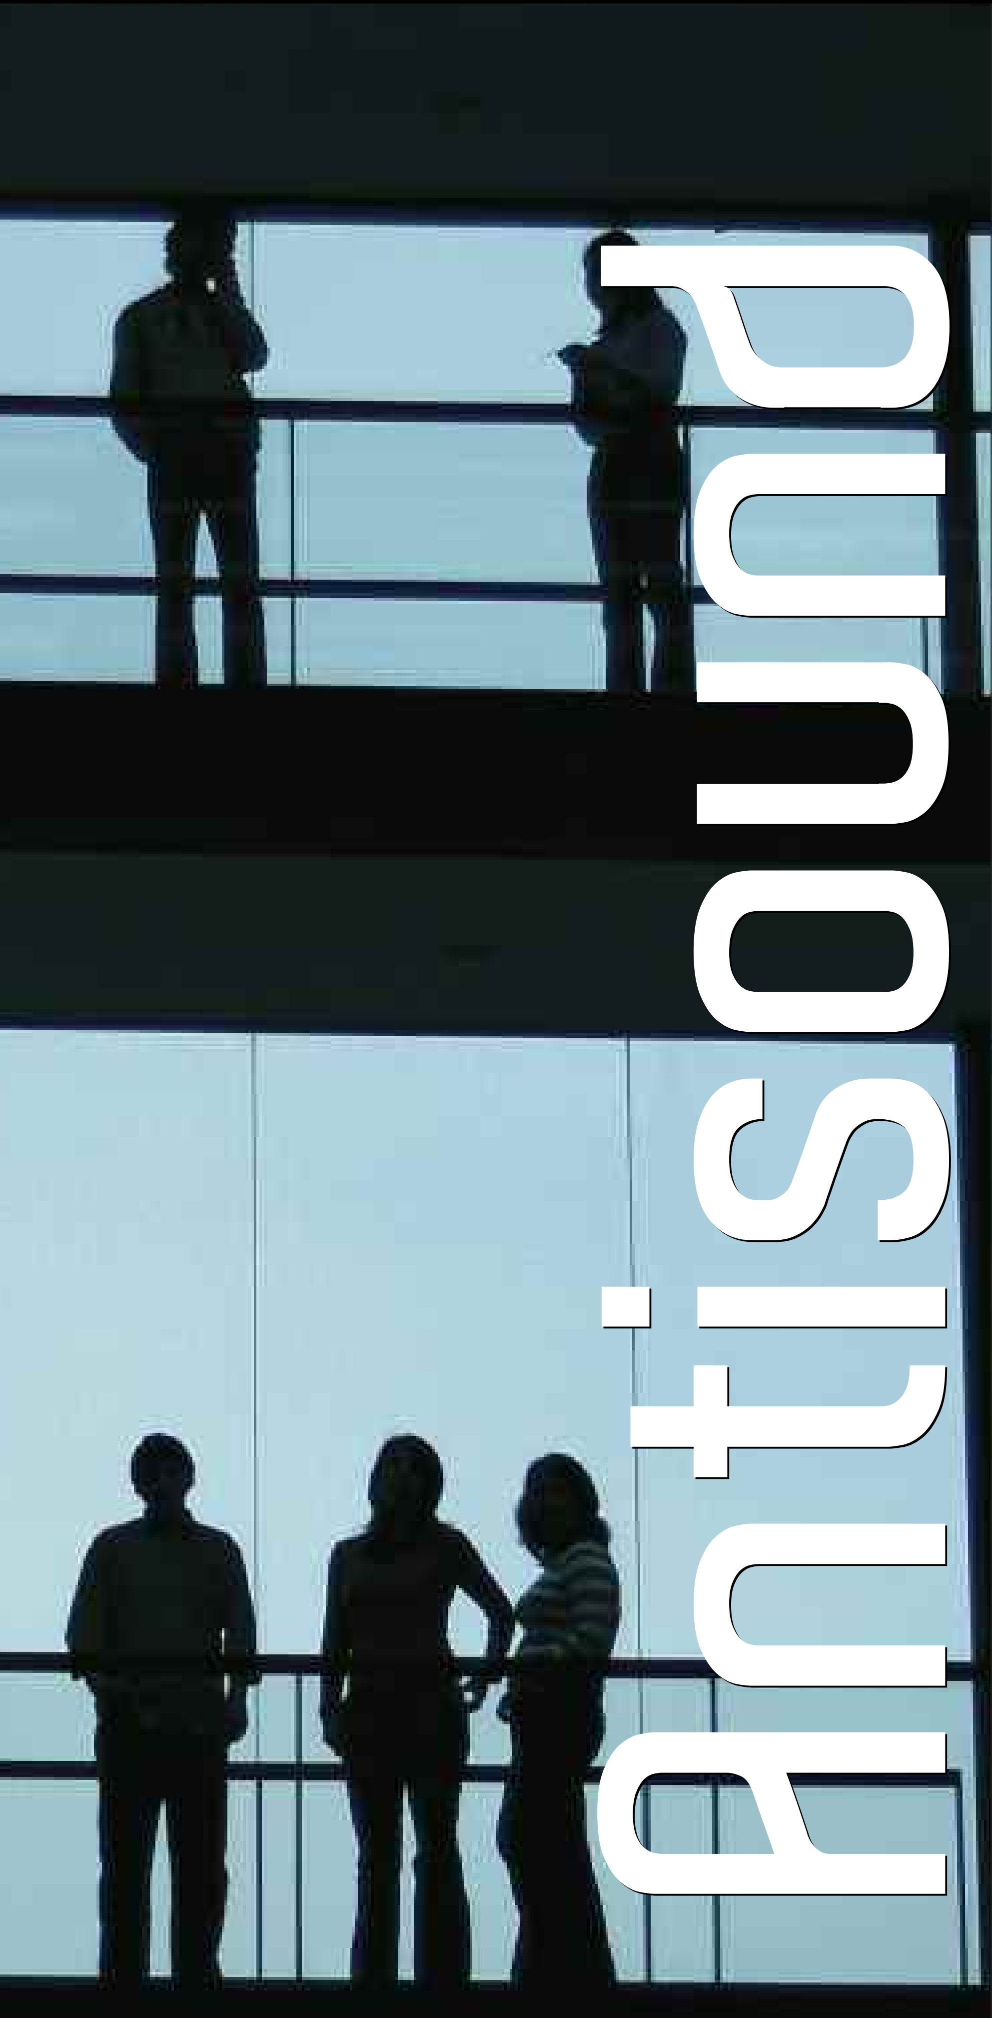

The conversation shield is one example of using audio technologies to help keep things secret, but there are other technologies that can be used in improving the world around us. This technology is called antisound. The basic idea appeared in the 1957 Arthur C Clarke science fiction story Silence Please, where it was called a 'Fenton silencer'. It works by canceling out soundwaves. Sound is a pressure wave in the air – when we speak, for example, the vibrations from our mouths either increase or decrease the air pressure. These puffs then travel through the atmosphere to the ears of others, who are able to detect the changes in pressure. The brain decodes them back into speech. It's the same with any sound; it's just a pressure wave in the air. With antisound we want to make this noise vanish!

Nothing to hear here

To create antisound is simple. You record the sound that's being produced then quickly play it back, but sneakily you invert it: the loud bits get soft and the soft bits get loud. The sound and the antisound pressure waves mix, and where the original sound wave had a high-pressure area, there is a corresponding antisound low-pressure area. When combined, the two produce an area of normal pressure. Normal air pressure is, well, it's just air. Nothing to hear here. Move along please. So by generating an antisound wave that is the exact opposite of the original, the sound can vanish.

Antisound is good if you want to keep secrets but also good if you have a very noisy factory, say, or the throaty roar of your sports car is stopping you hearing the quiet music track you are listening too. Want to quiet the engine down a bit? Just play its antisound at the same time and hey presto – silence.

Sounds like a good idea

Sound can be horrible, pleasant or even secret. Today we are able to use digital technologies to capture sound waves and manipulate them in real time. This gives us new ways to change the soundscape around us and create new ways to listen (or not, as the case may be). Through understanding how to manipulate one of our basic senses, hearing, researchers can change our world. Remember, you heard it here first.

How *fast* is than sound

The speed of sound is pretty fast. Not the sort of thing any one person could measure in their bedroom...or is it easier than you might think? And why are a team chasing a land speed record interested in your measurements?

Speed is just a measure of how far something travels in a given time: how many metres per second, for example. That means you can calculate how fast something is going just by dividing a distance travelled by the time taken to do it.

Cannon fire

One way to measure the speed of sound would be to stand a long way from a helper who fires a cannon. You start a stopwatch when you see the cannon flash. Ignoring the time the light takes to get to you, that is the time when the sound sets off. You then stop timing when you hear the bang. Measure the distance in metres you are away from the cannon and divide that distance by the time in seconds you recorded. Easy! You have the speed of sound in metres per second (m/s).

Well sort of. It's a good start and it is the way scientists did measure the speed of sound for a long time, but it isn't going to be very accurate is it? Your reaction time will mess things up a bit for one thing and a stopwatch isn't very accurate. Also cannons aren't so easy to come across these days, and you would be in big trouble if the cannonball hit something!

Fastest car on earth

You must need some pretty clever equipment to measure it more accurately mustn't you? Well perhaps not. We were talking to Ian Galloway, who is part of the Bloodhound SSC team (see cs4fn issue 9, available at www.cs4fn.org) and he pointed out there is an easy way for anyone to get a really accurate measurement of the speed of sound. The team are building a hi-tech car called Bloodhound to break the land speed record by the year 2011. Actually, they intend to completely smash the old record of 763 miles per hour and drive the car at over 1000 mph.

The Bloodhound team are also building a national database of the speed of sound at different temperatures, pressures and humidity levels, all of which change the speed. They would also like you to contribute by taking measurements for them using Ian's method.

s faster
d?

Why are a team working to break the land speed record interested in sound stuff? Well, when you travel as fast as Bloodhound is going to, you push through the sound barrier. That just means that the car will be travelling faster than the speed of sound. Sound is made by changes of pressure moving through the air in waves, and as you start to go faster than those waves are travelling you generate shock waves.

It is important that the team understand all the science behind what will happen to the car and the air around it as it goes supersonic. For example, shock waves can also do weird things to the car, like make the back end lift...not something that Andy Green, the driver, is very keen to happen. That could turn the car into something closer to a dril, tunnelling supersonically into the ground! Not good for Andy.

So how can you measure the speed of sound? All you need is a standard computer, some free audio software, a metre ruler and some old walkman earphones! Go to the Audio! website to find out more.

That's science

Coming up with creative new ways to answer apparently difficult questions, then backing them up with careful measurements is one of the things scientists have been doing through history. Welcome to the club!

Audio! Action

High intensity sound waves cause some liquids to give out light. This is called sonoluminescence, where a bubble forms and pulses with light as the sound wave passes. It may one day be a new source of power.

Silent drum

We know the music world can be pretty crazy sometimes, but would you think it's a world where drums are silent and shadows make sounds? It is now.

Inventors are continuously coming up with new technology that enhances the musical experience of performers and composers. 2009 was no different, as shown by the Guthman Musical Instrument Competition at Georgia Tech. It showcases new musical instruments whether played by human, computer or robot. The competition scoring is based not just on musicality but also on design and engineering. Nearly 30 inventors from seven countries battled for $15,000 in prizes with Jaime Oliver (a Peruvian computer scientist, not the TV chef), the ultimate winner for his Silent Drum.

Jaime is a PhD student at the University of California, San Diego, where he is researching the use of gestures to control electronic sound. He developed the Silent Drum with Matthew Jenkins. The original idea was that it would be hit by a mallet. As they experimented with it, though, they realised that you could get much more subtle control of the sounds if they used hand gestures instead.

Physically the Silent Drum is just a drum shell with an elastic spandex head, lit from the inside. It's electronic, so it doesn't need physical motion to create a sound. Instead it uses shapes and shadows to compute and control the sound it makes.

When you press the drum surface, either with a mallet or your hand, you can create all sorts of different shapes. They are captured by a computer vision system. It analyses the images from a video camera that is part of the drum and translates different shapes into different sounds. That means the drum really is silent. It is the computer that generates the sounds based on what it 'sees'. The sound depends on things like the position of the player's hand, the shapes it makes, changes in direction of the hand and so on. Each such 'parameter' can be used to control the sound in different ways.

Essentially what the silent drum does is turn shadow puppetry into a musical art form instead of a visual one. Shadows never need to skulk around silently again!

The 8-bit larynx

There were lots of imaginative entries to the Guthman Competition. For example, there was a tongue music system, an instrument that makes music from the way a player solves a Sudoku, and a whimsical star-and-circle-shaped contraption with exposed wires that uttered slow, mysterious sounds. Dan Stowell, whose work was profiled in the first issue of Audio, was also there and made the finals demonstrating his system for controlling a synthesiser with his beatboxing voice.

Dan's system picks up his vocal articulations with a microphone and converts them into control signals for his computer simulation of a classic 8-bit sound synthesiser. That means he can control its quirky sounds with his own language of vocal acrobatics. See Dan performing with his '8-bit larynx' by following the link on **www.audio4fn.org**, the Audio website.

Image credit: Rain Rabbit, Flickr

Sounds like a good illusion

We all know optical illusions are meant to fool our senses. Straight lines look curved, big circles look smaller. Illusions of this type happen when our brains make the wrong assumptions.

Each of our senses has a really tough job to do, and most of the time we get it right, but occasionally researchers can create particular patterns that cause our brains to make mistakes, to try and measure something that's not quite as we expect. Hearing is a human sense too and our brain also makes assumptions about what it expects to hear. There are lots of fascinating audio illusions that we can create which fool our hearing. To try them out for yourself, just follow the link from the Audio! website at **www.audio4fn.org**

Anything that's rock n' roll

In 1991 Johnny Depp starred in the video for a song called Into the Great Wide Open, by Tom Petty. He played Eddie Rebel, the main character of the song, who had a guitar but not a clue. Then again, you don't need much more than a guitar and some chords to play most rock songs.

"She had a guitar and she taught him some chords, the sky was the limit..."

People all over the world want to know the chords to their favourite songs, so they can easily play them on their guitar or with their band. Unfortunately, songwriters like Tom Petty don't usually publish the chords they base their songs on. What happens instead is that people try to figure out which chords he played by listening to the song. Professional transcribers sell their chords in sheet music books, while others just put up a page on the Internet for everyone to use. For rare songs, finding either of these can be a problem, so hobby musicians often end up having to transcribe the chords themselves. Unless ... they can get a computer to do it for them!

This is what Matthias Mauch and Simon Dixon from the Centre for Digital Music at Queen Mary have been doing. It's a three-step process. First you have the computer calculate the pattern of what goes on at every beat of a song. Since that gives you a very messy view of the music, the second step orders the beats into bars, eliminates short 'phantom' chord changes in the middle of bars, assigns the key to the song and so forth. To do that all simultaneously Matthias uses a computational model developed by statisticians called a 'dynamic Bayesian network' and feeds it what we know about chords, bars and keys in general. That means that when the model runs, it comes up with results that make musical sense. The third step is the easiest: write the results of the process in a way that musicians can use it, known as a lead sheet.

So, one day soon you may no longer need to search the web looking for chords. Your chord-savvy computer may just listen to the music, hand you the lead sheet, and off you go. Of course that doesn't guarantee you will sound as good as Tom Petty, so don't blame the computer if your audience ends up telling you: "Don't come around here no more".

Audio ! Action

Frequency, the number of times a signal changes in a given time, is measured in Hertz. It is named after German physicist Heinrich Rudolf Hertz, who tragically died at the age of 37. His son went on to invent the basics of medical ultrasound and his nephew won a Nobel Prize.

In the drummer's seat

Picture yourself as the drummer of your favorite band, holding the drum sticks and ready to kick off the gig.

Imagine that the pace at which you play the drums controls various loops of music and syncs them with you. It's now up to you to keep the appropriate tempo. The computer will follow your every move. This is the research topic of Andrew Robertson, who devised the B-Keeper software to do just that. By listening to the kick and snare drum, the software makes the music respond to you as a player, so that when you speed up, the music speeds up with you. The trouble with electronic music of the past is that the computer starts to take over. B-Keeper puts you back in charge when playing with loops and samples. Recently the system was demonstrated at the Cheltenham Science Festival, allowing anyone to experience what it is like to have a band playing with them.

Composing music with home-made software

Music is often referred to as the most scientific of the arts, because of the vast knowledge a composer must possess in order to create original music. Sinan Bokesoy, a Turkish composer studying at the University of Paris VIII, France, raises the scientific dimension of composing music to a new level, though. He believes that innovation is key to music composition, and recently submitted his PhD on his approach to sound generation, 'macro sound'.

He started to develop his own programs back in 2001 as a way of creating unique sounds that he could not obtain using available music tools. He was used to creating sounds with commercial synthesizers and software, but he felt that these sounds were not original enough. As a result he created his own called Stochos. It is a kind of synthesizer, but instead of creating sounds from scratch, it uses sound samples to generate complex and evolving sound textures. The name Stochos refers to the stochastic (ie random) music process that was used by Iannis Xenakis in his music in the 1950s. The program recombines and organises the repetition of a sound sample, such as the sound of a clarinet or a soprano voice, to create rich sound textures, sometimes with unexpected results. He refers to these unexpected results as 'emergence', a feature that he has been specifically pursuing when developing Stochos. He claims that mixing these sounds with traditional instruments enables him to create a 'psychoacoustic melody' by contrasting the timbres of the sound. This work has not only led to a PhD thesis but also an album, "Tales of Future" combining traditional music with that of Stochos. You can see his video "Ghost of the Street" by following the link on **www.audio4fn.org**

Just about everyone listens to music, and most people can sing a bit (even if they would never do so outside the privacy of their own shower). But how many people can write songs? Composing music takes a special type of creativity and years of musical training. It's not something that just anyone can do. Or is it?

Digital song

A group of amateur musicians have developed software to enable even novices to write their own songs. It's called MySong and it was developed by Ian Simon, a graduate student from the University of Washington, while working at Microsoft Research with Dan Morris and Sumit Basu. While they don't expect it to produce a hit song any time in the near future, they see its value in enabling untrained users to share in the enjoyment of the creative process of songwriting without struggling up the steep learning curve of music theory and mastering an instrument. After all, everyone has a voice.

How does a musician usually write a song? There are no hard and fast rules, but typically they first come up with a tune in their head, and then grab a guitar or piano and try out which chords sound best with the melody. If they are worried about forgetting the new tune, they might first rush to their PC, start up their favourite recording software, and commit the idea to hard disk before it's too late.

In any case, many amateurs find it difficult to discover the "right" chords for their new tune, so this is where MySong steps in. The user has to do little more than sing their melody, and MySong does the rest.

Is it really that easy?

This is how you use MySong: first you choose a tempo, using the tempo slider to set the digital metronome. Then you hit the record button, sing your tune, and hit the stop button. MySong then comes up with a sequence of chords to match your melody. Sounds simple?

But how does it work?

Before it can choose the chords, MySong has to work out what notes you sang. This sounds easy, but sung notes are not pure tones: a lot of singers, at all ends of the ability spectrum, sing with vibrato, which can confuse a computer when it is trying to recognise your notes. And then there are those notes that didn't quite come out right (see page 4 for how to fix them), and the fact that the notes might not align with the notes on a piano, but fall somewhere in between. MySong assumes that you get all the notes wrong by roughly he same amount, so it shifts everything by that amount and then locks it on to the nearest musical note.

MySong has a database of 300 songs from which it learns how to associate chords with any given melody. It does this by extracting two types of information: what chords sound good after each other (this is called the 'transition matrix') and what kinds of melody notes are usually associated with each chord. Obviously you don't want the same chord every time you sing a particular note, so the chords are chosen using random numbers and probabilities. Also, you can influence the way chords are chosen using the 'Happy Factor' Slider: this tells the software how much to favour major (happy) or minor (sad) chords. Another control is the "Jazz Factor" slider, which determines how often surprising or unconventional chords appear in the song.

Finally MySong's chords are output to PG Music's Band-in-a-Box software, which creates a complete accompaniment from the chord sequence. In less than 5 minutes of work, you've got a complete song.

writing

Does it sound good?

For a given melody there are many possible chord sequences, and there are no "right" or "wrong" answers. The only way to assess the system is to ask people to listen to the melody and chords together, so MySong's authors put it to the test, asking experienced musicians to rate its output on a scale of 1 to 10.

For comparison, they also rated the chord sequences generated by Band-in-a-Box and some that were generated manually (by musicians). To remove any bias, the tests were "blind" (the listeners didn't know whose chords they were listening to). As you might expect, the musicians got the highest ratings, but MySong ran a very close second place. In many cases MySong's output was preferred over the chords written by humans, although it must be said that the musicians only had 5 minutes to write the chords to each song. Why not listen and judge for yourself by following the link at **www.audio4fn.org**

MySong is never going to be an automatic songwriter of number one hits. But it might achieve the more modest goal of enabling ordinary people, who could never write a song at all, to experience the fun of composing their own music.

You never know, as with any research it may be a stepping stone for something greater: that great compositional intelligence that is yet to come. Perhaps you might be one of the audio engineers who will help create it!

The colour of music

Music is all about emotion. From a slushy love ballad to the head-banging, no nonsense roller coaster of thrash metal, the songs we listen to can make or change our mood.

We have favourite tracks to play when we need cheering up, even maybe some music that keeps melancholic company with us when we are lost in thought. Emotions are what it's all about, and with portable music players built into many of today's mobile phones we can listen to music whenever and wherever we want. So how do you ensure that the music you are listing to is right for the mood you are in? Of course you have playlists of your favourite songs, which you can flick through to find the right song for the right mood, but is there a better way? This is the question Queen Mary computer science student Daniel Fiegenschuh asked, and his answer was a smart and colourful one. Go to **www.audio4fn.org** to find out how he created a system that used artificial intelligence to find music that fits a person's current mood.

Moon Droppings

When Apollo 11 touched down and Neil Armstrong first set foot on the Moon he officially uttered the historical words,

"That's one small step for [a] man, one giant leap for mankind."

So why is the 'a' in brackets? Well it seems he had planned to say for 'a man' but in the excitement he wasn't sure if he had said it. Whatever happened on the day, it still produced a memorable historical quote. Computer scientists have analysed the audio from the momentous day and claim that the 'a' was said but is missing because of problems with the quality of the audio feed. After all the Moon is far away and it was the 1960's. Dropped 'a' or not, the Moon landings were a pinnacle of human scientific achievement and involved some very clever NASA electronic engineers and computer scientists too.

Dream jobs: all you need is ears

A career as a pop musician may be the ultimate dream for many, but Marco Migliari has found his dream job on the other side of the studio, as a sound engineer at Real World Studios in Box, Wiltshire. Although not as high-profile as their pop-star counterparts, the engineers who record, mix, and produce music are just as essential to the final product that you hear.

Tucked away in idyllic countryside surroundings near Bath, Real World Studios was established by rock musician Peter Gabriel about 20 years ago. Since then it has become one of the world's most famous high-tech recording studios, particularly known for its vibrant and eclectic world music. Whether it is working with established names such as Van Morrison, Robert Plant or Peter Gabriel, or new names to the Western world such as Nusrat Fateh Ali Khan, Yungchen Lhamo and Youssou N'Dour, the engineer's task is to make them sound good.

Marco's career interests began in his hometown in Italy, where he worked at a local radio station, learning to use equipment such as their mixing console and reel-to-reel tape player. After studying electronic engineering in Italy, he moved to the UK to study music technology. Part of his programme of study included an internship at Real World Studios, where, after the day's work was done, he was allowed to "play" in any studio that was not in use. Before long he was part of the Real World team, and hasn't looked back since.

The job of a sound engineer involves a lot of technology and technical decisions, such as choosing sampling rates, sound encoding schemes and storage formats, and setting gains, EQ, dynamic ranges, and so on, in order to transform the electronic signals into beautiful music. But the real key to success lies not on the technological side, but the artistic: deciding how to position performers and microphones in order to capture their voices and instruments best. Experience doesn't provide all the answers, as Marco is often faced with exotic instruments that he's never seen before. He describes how you have to use your ears to find the place where the instrument sounds best, and then "Put a mike where your ears are". Sounds simple, but we're sure those ears have learnt a lot over the years.

For more on Marco and his job follow the link on our website: **www.audio4fn.org**

Become a real guitar hero

Guitar Hero and Rock Band are popular games whatever the console: from Wii to PS3, DS to Xbox 360. With only a few buttons to press it makes the Aerosmith solos appear easy to play, but many of us who actually play guitar can only dream of playing as fast as Joe Satriani, as smooth as Carlos Santana or to finger tap like Eddie Van Halen. Is there an easy way to become the next real guitar hero?

If you've ever had music lessons, you have probably heard your teacher saying "Practise it slowly until you get it right, and then speed up!"

Unfortunately that's a bold statement and only part of the reason why you should play slowly. Not everyone's teacher tells them about muscle memory and economical playing. When you play a passage over and over you are not just memorising a sequence of notes, but your fingers 'remember' how you played it and they will use these memories next time you go to play the same passage. So if you're one of the people who only plays half-heartedly when playing slowly (be honest, most of us do) and without the proper technique then you'll never play it properly at full speed. Muscle memory can be used to your advantage by practicing a riff or passage you find difficult slowly and precisely so that when you play it faster you play each note the way it's meant to be.

We mentioned using the proper technique when programming a riff into our muscle memory, but make sure that you take economical playing into account too. It simply means doing enough to get the job done, but nothing more! When anyone starts playing guitar it's difficult to press down hard enough for the notes to sound, but be careful as you progress, as you can press down too hard too. If you use too much force, you won't be able to move your fretting fingers fast enough to play the solos you're dreaming about. Try these tips out for yourself and see how playing lighter and smarter can make your solos sound more ferocious.

This doesn't just apply to guitar playing. Economical muscle movement can be applied to lots of different activities, both musical and non-musical, like typing at a computer. If you're a one-finger-jabber who presses the keys like a sledgehammer, you'll never learn to type quickly. Be economic and use more fingers, and only use enough force that's needed to press the key. You'll soon see yourself speed typing.

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 Simon Dixon and Paul Curzon with contributions from Peter McOwan, Jean-Baptiste Thiebaut, Jonathan Black, Rebecca Stewart, Martin Morrell, Kurt Jacobson and Matthias Mauch. **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.

EPSRC Engineering and Physical Sciences Research Council

Queen Mary University of London

www.audio4fn.org