



© 1997–2009, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non-commercial use. For other uses, including electronic redistribution, please contact us.

---

November 2003

Features



## Career interview: Audio software engineer

by Rachel Thomas



*Skot McDonald is an audio software engineer, currently building software synthesisers for Expansion, a London-based company specialising in music production tools. He talks to Plus about how he uses mathematics to understand music, and how he managed to combine his passions for music and computing to create a successful career.*

### Creating new possibilities

Most music you hear today uses the kind of technology Skot develops – such as delays and distortions that change the sound of an instrument, synthesisers that recreate old analogue instruments, and autotuners that correct the pitch of a singing voice.



## Career interview: Audio software engineer

To reproduce the sound of a real instrument, Skot creates a mathematical model of its physical properties: the vibrations of all the parts of the instrument – the strings, the sound board and so on – and how the materials that make up the instrument transmit the vibrations from one part to another. The more accurate the model is, the more like the real instrument it will sound. Musicians can then run the model on a laptop instead of lugging the real thing around, "freeing them from the burden of large, heavy equipment and expanding their musical possibilities."

The really interesting thing, says Skot, is that impossible instruments can be modelled too. "You could model a guitar, but instead of having strings that always have the properties of steel or nylon, you could smoothly glide the numbers that model your strings from nylon to steel in real-time, producing some bizarre sound that you couldn't produce with an actual physical instrument."

## From music to mathematics

At school, Skot says, he liked maths "a fair bit" but no more than any other subject, and he didn't have any firm idea of what he wanted to do at university. "Engineering seemed to be the most career-rich option at that time", he says, so in 1992 he enrolled in a double degree of mechanical engineering and computer science at the University of Western Australia, and soon discovered that his choice also offered him the chance to follow his passion for music. "I had a couple of bands back in Perth, none of them particularly successful but all of them fun." As Skot and his band-mates were increasingly writing their own instruments and software, one of these bands, Vellocet, eventually became a software company as well.

In 1995, for his final year project in computer science, Skot investigated automatically transcribing and composing music. "What I wanted to do was analyse goth industrial songs, and transcribe their drums", he says. "Then, using the collections of drum beats that we had pulled out from different songs, we would build a rule base of how Trent Reznor [from *Nine Inch Nails*] and Al Jourgenson [from *Ministry*] liked to put their drums together. And then we would see if we could generate similar-sounding drum beats with suitable dire-sounding chord combinations."

Once Skot discovered that the computer science department would let him create music and call it research, he realised he had found his field. The research also allowed him to play with computers, which Skot says "is another great love of mine, because it lets you build worlds".

## Catching waves

The sounds we hear are combinations of different frequency vibrations, that change over time as the sound changes, say, in pitch. The first step in analysing sound is to split it into the frequencies it consists of. "In biological organisms such as ourselves, this function is done by part of the ear called the cochlea, which is a little snail-shaped thing just behind the eardrum", Skot explains. "It takes an incoming signal and breaks it up into its component frequencies. A lot of sounds that are meaningful for us, for example, speech and music, have a very well-defined frequency structure." A pure frequency has the structure of a simple sine wave. But a tonal sound – such as a note played on a trumpet – is not necessarily a pure frequency. "We hear a single unified sound, but if you look at that in frequency space, you'll actually see a comblike structure."


Tonal sounds produce a *harmonic series*, where the peaks in the comblike structure follow a particular pattern. The first peak occurs at the lowest frequency in the series, called the *fundamental frequency*, and every other peak has a frequency that is an integer multiple of that fundamental frequency.

The nontonal parts of an incoming signal are often described as *noise*. "In the case of the trumpet, the tonal aspect would be all the harmonic peaks, and the noise component might be the little bit of breath at the


## Career interview: Audio software engineer

beginning, and any other random bits of noise", says Skot. One technique is to try to model any incoming sound as a tonal component plus a noise component. This sort of research into analysing audio signals in the frequency domain is useful not only for automatic transcription of music, he says, but also for building hearing aids, and automating speech transcription.

To understand the frequency structure of a sound, Skot uses a mathematical tool called the *Fourier transform*. Fourier, a French mathematician who lived through the French Revolution, realised that any waveform, no matter how complicated, could be written as a sum of sine waves of different frequencies and amplitudes. Finding the Fourier transform of an audio signal means breaking the waveform down into these components, revealing the frequency and amplitude of the underlying sine waves.



Although only a finite part of a signal is recorded (top), the Fourier transform algorithms assume that the signal is infinitely repeating. If the signal is not recorded for an exact number of periods, this would result in a "glitch" where the recorded signal repeats (bottom).



A recorded signal (red) has a windowing function such as a gaussian curve (blue) applied to it, to create a curve (purple) that will repeat smoothly.

But as sound usually changes over time, this analysis can to be done only for short periods of time. "We take small chunks of sound, say 100 milliseconds, and apply a window over the time domain", says Skot. "This window will be something like a Gaussian curve, basically a little lump to say that things at the ends are zero, and the stuff in the middle is one – it's a bit of a fudge." This fudge is necessary as to use Fourier analysis on a wave it must be infinitely long – "it's a small bit of the sound, but we're pretending that it is the whole sound. We just analyse that one little chunk.

"Then, because we want to know how that sound evolves, we move our window slightly down the signal, apply the windowing function again, and see what the next little chunk is doing." For each little slice a spectrum is produced, showing the amplitude and frequency of the component sine waves. As the slice moves with time, these spectra build up a three-dimensional graph of the sound spectrum.

## Maths comes in handy



Despite using a lot of mathematics now, Skot, like many others, didn't really appreciate its importance until he had a reason to use it. "Engineering puts you through two years of maths hell, because they get you up for eight o'clock lectures every morning", he says. "In second year we actually did Fourier transforms, but I had no idea what was going on. It wasn't really until I was interested enough to create effects in frequency space, that I got into it and started to understand it." Needing the tools that mathematics provides gave Skot a reason to study the subject, but appreciating its usefulness doesn't necessarily make it his favourite subject. "You bash it out over many coffees, and my favourite bit as a computer scientist is you lock that away in a piece of code that knows what to do, and you never go near it again! I have my magic black box, and I never want to learn that again!"

When he finished his degree, Skot went on to a PhD in a field called *source separation*. "You have a band, which consists of a number of musical sources. When they are all mixed together on a CD track, it would be nice to separate them back out again so that you can just listen to the drums, or the bass guitar. If you have done this separation you can then transcribe the parts individually."

For his research, Skot was specifically interested in identifying the rhythmic aspect of the music, "finding the positions of drums in an audio track – when they started playing, working out what sort of drum it is. Often songs won't just be drums in isolation, they will have some other tonal instruments playing over the top, and for my purposes any other instrument is noise. You have to remove the background tonal instruments and just be left with the percussive track. To do that we need to model the tonal instruments, and once we have a pretty good idea of what the contribution of the tonal instruments to the mix is, we can remove that to just leave the drums behind."

## New York, New York



Image [Freeimages](#)

In 1999, while he was still doing his PhD, Skot was hired by New York-based company *tomandandy*. "I put the drums on hold for a while. I was getting interested in modelling tonal sounds at that time, and having someone come along and pay me, especially as my scholarship was running out, was good!"

For the first six months, Skot and his colleagues analysed the singing of opera and jazz singers, who would come in to the *tomandandy* offices to perform. "First of all we would produce a really accurate frequency trace of what they had sung. Then we split that up into notes – 'this bit is definitely an A here, this bit is a trill!'" The work required a lot of musical understanding, as well as mathematical understanding. "You don't always want every little note that someone sings to be fully realised as a note on the score. Sometimes they are elaborating – there's a lot of little rules of thumb."

Later on, Skot's work at *tomandandy* moved on to analysing music at the level of phrases – trying to identify the overall structure (like alternating sections ABAB) in a piece of music. "It wasn't enough to just identify notes. For example, we had to recognise chord patterns, recognise the overall feel of a particular section of music. It's a really fuzzy thing." Recurrent sections in a piece of music aren't always exactly the same, and have to be identified by similarity modelling. "There is a whole bunch of statistics about the piece of music at a certain time. You know that leading away from that moment, in forwards and backwards time, there are particular characteristics – a certain pattern of drums, a certain pitch structure. It might also have other, more airy-fairy characteristics, like it seems grainy, or is particularly busy or unbusy." These less obvious qualities have to be defined mathematically. "What does a busy piece of music sound like? Does it mean that the noise envelope is moving around a lot? Does it mean that there are lots of different fundamental frequencies?"

In 2002, Skot left *tomandandy*, and came to work in London for *Expansion*, a company that creates synthesisers and music effects software. "*tomandandy*'s stuff was good because it was very academic, and close to what I was doing for my PhD", he says. "But at heart I would like to get back to being a practising

musician again." He felt that working Expansion would give him an opportunity to be in close contact with other musicians, and "write synthesisers for my own use, as much as for the people who buy them".

## Making music of your own

Skot has achieved what most of us dream of, combining his passion with a successful job. But it took some lateral thinking to follow his path. "There's not many undergraduate courses in music technology and signal processing being offered", says Skot. "It really is the kind of thing that you have to do at the fourth year or postgrad level, which is a shame. But this will probably change. I am coming in as an old duffer, at the bone-shaking age of 29.

"The use of computers in music seems to be everywhere, but it is still a relatively new thing at your average uni. And being the slow momentum-carrying beasts that they are, they're not going to suddenly come up with something as nice as a music DSP [digital signal processing] course. There are well established computer music research centres, such as CCRMA at Stanford, IRCAM in the Pompidor in Paris, and many newer, smaller labs springing up catering for post-grad work, like the Digital Music Lab at Queen Mary college at the University of London."

Another possibility might be a course in cognitive science. "You do neurobiology, a bit of psychology, and you tend to do a lot of computer modelling of these processes, which is good from a music analysis point of view." But often the way into this area is to get the background in the mathematics, and link up with a member of staff who is interested in music. "Electronic engineering departments and computer science, all the mathsy courses, seem to have the odd professor or two who are secretly music heads."

It's good to know that even academics can have passions!

## Links

- [Natural frequencies and music](#), Issue 4 of *Plus*
- [Designing loudspeakers](#), Issue 4 of *Plus*
- [Self-similar syncopations: Fibonacci, L-systems, limericks and ragtime](#), Issue 10 of *Plus*
- [Online course in digital signal processing from www.bores.com](#)

---

## About this article

[Rachel Thomas](#) is news editor of *Plus*. For this article, she interviewed [Skot McDonald](#) at the offices of [Expansion Audio UK](#) in London.

---



*Plus* is part of the family of activities in the Millennium Mathematics Project, which also includes the [NRICH](#) and [MOTIVATE](#) sites.