The following appeared in edited form in Muzik ETC in Graham Collins' Synthesizer Basics column.

In our last instalment, I promised something a little more esoteric. This month we delve into a concept that has been getting a fair amount of press lately - granular synthesis. Grain? As in.. wheat? Not exactly, no. A little history is in order at this point so lets dig in...

As early as the late 19th century, it was suggested that sound might be able to be broken down into fundamental building blocks much tinier than musical notes or intervals. We are talking here about sound at (for lack of a better term) an atomic level. Matter is composed of atoms, but what exactly is sound made of? Shortly after WWII British physicist Dennis Gabor posited the notion that any sound might be constructed using the correct number of sonic 'grains' - very short sound snippets that are nearly inaudible on their own. Thus, suggesting a kind of aural quantum mechanics. Early experiments with tape recorders splicing minute bits of tape together were both time-consuming and labour-intensive, producing rather disappointing results for the man-hours required. It wasn't really until the advent of computers and digital synthesis that the idea of assembling 'grains' of sound became practical. The fascinating Iannis Xenakis ultimately was the man who became known for applying and advancing the concept of granular synthesis. Xenakis was a colourful character - a Romanian-born Greek nationalist engineer who fought in the resistance against the Nazis and ultimately was exiled and emigrated to Paris to study musical composition with the likes of Darius Milhaud and Olivier Messiaen. Oh yes, he also founded the *School of Mathematical and Automated Music.* Did I mention he was also a modernist architect and worked with Le Corbusier? Sheesh. It was Xenakis who advanced the theory and championed the cause of granular synthesis as a practicable source of sound-creation.

So what the heck is a 'grain' of sound? How do I make one? Are they sold by the pound? Er, no. The theory is really surprisingly simple. If you take a basic waveform, and give it a very short duration with an amplitude envelope approximately rectangular in shape and a slight fade in/out, you would have formulated a basic sonic grain. Opinions differ on how long the envelope should be, but somewhere in the 10-20 millisecond range seems to be fairly consensual. If you go much shorter than 5 msec or so, sounds are perceived merely as 'clicks'. If you go much longer than say 100 msec or so they become very short 'notes'. The idea essentially is to keep them short enough just to hover on the horizon of human perception. This is all well and good, but what can you do with this little grain? On its own, not much. What we need to do is to formulate these basic elements of sound into a larger array of some kind. The term often used is a granular 'cloud'. Think of it this way. It's sort of like the way you can see a cloud in the sky, but not the individual drops of water it is composed of. Sonic 'grains' are the metaphysical trees we are *supposed* to miss for the forest. Each grain has its own pitch, tone, and place in time and space in a much larger sonic 'cloud' that we can create. Yes I said time and space. One of the most fascinating aspects of granular synthesis is the relative ease with which creation of sounds involving stereo and surround imaging can be accomplished. Another way to think about this is as a form of animation. Individually, movie frames show no movement when viewed one at a time. They must be viewed in succession very rapidly for our brain to make any real aesthetic sense of the big picture.

There are at least two current applications for this type of technology. Firstly, as a synthesizer - building sound from the ground up. There do exist a number of standalone applications as well as VST plugins kicking around the Internet that use this technology to one degree or another. There are I believe also a couple of recent hardware devices playing with this technology, and their numbers will only increase over the next little while. Depending on the application, you may have a variety of parameters to play with. Remember, *granular synthesis* is a *concept* or *class* of synthesis in the same way *analog subtractive synthesis* is merely a concept and not an actual instrument. Every instance of the technology will not necessarily have the same layout as any other. At least, no more than say a Moog modular synth must be similar to a Korg Polysix. There are a number of programs that approach this 'technology of the musically very small' and can elicit very interesting results.

One other form of this technology that has become very popular in the past couple of years is a form of granular *re*-synthesis. Digital recording and sampling software that performs real-time pitch shifting without changing loop tempo, or conversely changes sample duration without changing pitch, use

this very technology. By breaking a sample down into very small component 'grains' of sound, it becomes possible to give the perception of changing pitch by lowering the pitch of the grains but not altering their placement. Conversely, by spreading out the grains a bit further from each other than originally formed without altering their individual pitches, we can seem to 'slow down' a sample. It's pretty groundbreaking stuff and allows us to think of sound, as being somewhat more elastic than we previously thought was possible. Pitch and time are indeed directly related, but there is plenty of room in our little heads for perceptual trickery that make it seem not so.

Graham Collins is an Ottawa composer/synthesist currently producing an instructional DVD on synthesis. His website is www.pongthrob.com, and he can be reached at graham@pongthrob.com