

Any Sound You Can't Imagine:- Exploring the Soundscape

Has the introduction of electricity into music forced us into a paradigm shift away from the traditional understanding of both music and ourselves? Muddy waters indeed, but ones that are possibly becoming clearer with innovative research around the globe at such institutes as the University of York, where currently a new course has been established to dig around further in the depths of this possible deep misunderstanding.

The Heinemann English dictionary contains the following definition:

synthesis - *noun*

the combination of parts or elements to produce a complex whole. (Greek)

This is the fundamental concept that underlies most, if not all, of sound production by electricity and electronics -the combination of sound elements into a complex tone. Since the start of production of electrical and electronic instruments there has become established two broad groups of people interested in synthesis. The first, their general aim being to produce an electrical or electronic replica of the sound of traditional instruments, The second being those investigating the new and completely novel ways in which humans, sound and electronics can interact. A thought: if synthesis is to be a basis on which we can recreate every aspect of the sound of a traditional instrument, might it not be possible that this is like attempting to produce the Elgin marbles using Lego bricks?

Traditionally the bias has been heavily towards the Lego brick side – in theory, by combining enough different types of brick, you can reproduce anything. The major problem is knowing what the different types of brick are, and in the end, the reproduction will still be made of plastic and won't feel right! Recently, the shift has been towards analysing human interaction with objects and extrapolating that data to musical instruments so that the developed sound synthesis techniques can be explored in far more dynamic ways, and some might say, in ways that are more musical.

In order to examine more closely the issues surrounding the current debate, a little background and history might be useful. The first electrical instruments were developed around the 1900's with some distinct examples emerging which can be used to demonstrate the distinctions between the concepts of design. The [Telharmonium](#), patented in 1897, was to many people a visionary instrument. Although enormous and expensive, it was based on an early form of additive synthesis, and exhibited most of the properties that were the holy grail of later electronic instruments – polyphony, tonal control and almost unbelievably, two touch sensitive keyboards. However, it also marked the start of the almost total and utter domination of the keyboard based synthesiser in later years. One thing to note – it was based entirely on *electrical* equipment and employed no electronics. With the development of the vacuum tube 'valves', the possibilities for development of new and innovative sound production rapidly expanded. The [Theremin](#) is a good example of a device that very early on broke away from the keyboard type of interface, in fact requiring no physical contact at all. This enabled continuous variation of pitch and volume, and whilst this provided incredible flexibility in these areas, the timbre of the instrument was somewhat fixed. With skill this 'restriction' could be circumvented – as

demonstrated by the playing of [Clara Rockmore](#), who lead the way in developing the techniques for playing the theramin in such a way as to elicit what could be described as highly musical sounds, from what was basically a single tuned oscillator. An important point can be made here – the Theremin is a very basic form of device, but so is a string on a violin. Neglecting the effect of the highly developed body of the violin, essentially the principles of playing are the same – you have direct control over volume (speed of bowing) and pitch (position of finger). But without the skill of the player, it becomes simply the equivalent that a lump of rock (the uncarved block for Taoists.. maybe there is more musicality in silence....!) is to the Elgin marbles. The transformation from a simple system for making sounds to a musical instrument could be likened to requiring the skill of the stone mason – a skill honed by years of practice and teaching by masons before. Not entirely unlike a skilled musician. The sound of a violin in the hands of the untrained is hard to forget. Historically, musical instruments are simple but flexible systems, and as a result of the flexibility, take time, effort, training and a little talent just to play, let alone master.

How is it possible to know what makes a good instrument before many years are invested in devolving techniques to play it? A partial answer to this may direct us to the reasons for the keyboard's domination of the electronic musical instrument interfaces. The first half of the answer lies in evolution – if you already have a popular and widely used interface, why modify it? The piano keyboard has a set of distinct advantages over most other interfaces. Peculiarly enough, one of the most important factors is that your target market *already* knows how to play an instrument based on this particular system. It is suggested that the keyboard achieved this over and above much older instruments such as the violin at least in part due to a somewhat bizarre social quirk of the Victorian era in Britain. It was expected that [any aspiring young lady would learn from an early age](#) Demand for [pianos in middle class drawing rooms](#) soared, and in the interests of the market, the development of mass production and the invention of the [upright piano](#), the piano became cheap enough to be afforded by anyone aspiring upper class culture. Hence a massive base existed of people who already had the skills and time invested in learning to play a keyboard based instrument. Why learn a different instrument, when not only could you learn to pick out a tune in a few weeks, but also gain social status as a result?

The second half of the answer, as alluded to above, lies in the ease of learning such an instrument. A simple (though not mechanically) system such as the keyboard allows a first time user to pick out a recognisable tune within a very short time of meeting the instrument – it has instant appeal. Whether the user then improves or not is an entirely different matter. Research at the University of York has provided some evidence that human learning is not quite as expected with regards to interfaces for musical instruments. Evidence has been shown that when presented with a simple, direct mapped (the parameters of the sound are associated directly with physical controls) control interface for audio, a user can use it almost immediately to attempt to replicate sounds presented to them, but there is little or no improvement shown in the short term, and only after extensive practice do the results improve. Using a complex mapped system, where the individual properties of the sound are not directly controlled but each affects the other in some non-linear way, then even though the results are at first worse than the direct mapped system, the user learns much faster and improves to a level far beyond that of the direct mapped system. Particularly effective systems are produced when some concept of “energy” e.g. Faster movement, tighter grip, controls primarily an “energy” type attribute, for example volume. To a lesser extent, it is also shown that a cross coupling of volume with tonal change is also “preferred” by users, it being critical in the description of the device as an “instrument”, rather than just a “sound control device”.

Why then is the keyboard so prevalent? Perhaps it is that it is a reasonable compromise between these two factors – the historical and social influence and the ease of learning. At first glance the piano is a simple, direct mapped system. You press a key and a given, specified note is produced with a volume related the the velocity (or ferocity) of the impulse given to it. This exercise is repeated for every key on the keyboard, and the keyboard includes all notes in a western even-tempered scale. When a closer look is taken, it can be argued that the piano in its current form would not have gained its popularity if that was its only properties. The energy input to the key not only affects the volume, it also affects timbre and to a minor degree, pitch. The coupling is not direct by any means. The development of the electronic

piano gives an insight into the problems of discovering an apparently simple system: the [first attempts](#) were not velocity sensitive and as a result were quite inexpressive and unpopular. Velocity sensitivity rapidly became standard, but still lacked “feedback” - the feel of how loud a note would sound. Aftertouch was developed as a means of further enhancing the basic keyboard, but still the preferred keyboard feel for non-organ (organs never had any of the above in the first place) based keyboards are hammer action keyboards – directly emulating the striking action of the hammers in a piano. They are described as having greater feel and feedback than other systems, and users find them more natural and expressive to play. These factors have conspired with market forces to ensure that just about every developed synthesis technique developed has been implemented on a piano keyboard basis, regardless of suitability. As will be seen later, only very recently have the possibilities of developing instruments that may better suit the synthesis technique been considered. In some cases this can be extreme as trying to carve the Elgin marbles using a plastic spoon. The synthesis technique has massive potential but we lack the tools to exploit it properly.

Having discussed the historical aspect to the interface, how about the innards of these instruments? Many different methods of sound synthesis have been developed and implemented in various forms in the last one hundred years. Both mechanical, electrical and electronic systems have been produced – a pipe organ could easily be considered one of the earliest forms of sound synthesis, and can be an entirely mechanical device. Electronics have allowed some of the greatest exploration of sound synthesis techniques, and now take a large part of research into sound production and reproduction. The earliest methods of electronic synthesis concerned the adding of very simple waveforms to form more complex tones, a technique used in the Telharmonium using sinusoidal waveforms. This is called additive synthesis, and is the most basic form. Subtractive synthesis is the inverse of additive synthesis – a complex tone containing many harmonics is produced, often by means of a distorted oscillator. The result is then passed through a filter to suppress or enhance the harmonics of the signal to vary its tone. This was used to great effect in combination with additive synthesis in the original [Moog Synthesiser](#). These simple forms of synthesis remained the staple of synthesisers in musical instruments until the advent of Frequency Modulation (FM) Synthesis. This is essentially the control of the frequency of one oscillator with the output of another. The use of coupled oscillators in this way could produce radically different and new sounds – the downside being that the results produced can be quite unpredictable. A small change in the frequency of one oscillator can produce quite radically different outputs. Implemented in analogue form using Voltage Controlled Oscillators and in digital form popularised by the [Yamaha DX7](#), it has the advantage that it is a relatively inexpensive method of synthesis, both in terms of processing power (for digital systems) and complexity (for analogue systems). At this point a separation can be made of digital systems that emulate concepts that are parallels of their analogue counterparts and those systems that use techniques or achieve results only possible in the digital domain. FM, additive and subtractive synthesis all have their roots in analogue electronics and as such their digital counterparts can only really improve on them in terms of performance and complexity. Synthesis techniques such as sampling, direct drawing of sound waveforms and physical modeling can only be practically achieved in the digital domain, with few exceptions. Granular synthesis is a good example of a synthesis method only practically achievable in the digital domain. Granular synthesis is based on the use of tiny 'grains' - sections of sampled or generated sound, which are then overlapped, looped, expanded and contracted to produce an overall effect. It can be seen that whilst this technique may never produce the perfect “piano” sound, it allows development and investigation of soundscapes at their lowest level, and can provide completely new sounds with relatively simple means. It does raise a question though:- how could you “play” granular synthesis system in the same way as one would “play” an FM Synthesiser? Investigations into the human control of systems like these are only recently producing usable systems that are “playable” as opposed to programmable, allowing much more intuitive interaction with the sound.

A lot of the effort in developing synthesis techniques has been directed at trying to replicate the sound of traditional instruments, and one of the latest developments originally directed at this area is [physical modeling](#) (synthesis by rule). The fundamental principle behind physical modeling as a synthesis technique is to produce a model of a physical system inside a computer system, and then to interact with it. The model knows nothing about the sound it should produce, it simply obeys the laws of the environment programmed around it. For example, a simple physical modeling system might contain a model for a string, mounted between two fixed unmovable points. This could be considered a basic analogue for a string on a stringed instrument, for example a violin string. Indeed the output from even a basic model such as this, when stimulated in a way modeling plucking, sounds not unlike a plucked violin string. The model can be refined, limited only by the data available on the physical environment and the computing power available. It is possible to produce models that are extremely accurate for many instruments such as flutes,

saxophones, cellos, drums etc, the detail of them only limited by processing power. This could be likened to recreating the Elgin Marbles by creating a model of a molecule from marble, then building the statues molecule by molecule. The accuracy of the reproduction is only limited by the accuracy of the model and the number of molecules you can produce.

But the point of physical modeling is not to produce another violin – you might as well have a real one as you would need just as much skill to play the model as to play the original. Using these techniques, it is possible to experiment with and create instruments that would be considered completely impossible to create or play in the real world – how about a forty meter long trumpet? Or more bizzarely, a four dimensional gong? How would one interact with such an instrument?

If a woodwind or brass type blown instrument is being modeled, how would you play it with a keyboard? Very few alternatives have been used historically for playing with synthesis techniques, leading to a speculation as to what fantastic instrument are possible using alternative input systems. There exists on the market very few accomplished alternatives to a keyboard. The first, drum pads like the Roland PAD8 allow a fairly basic interface with one basic parameter, the power of the striking of the pads, to be categorised. This maps well to simple drum based models, where after the strike is made, no damping takes place, but physical modeling allows for much more – using a simple pressure sensitive pad, it is possible to provide inputs directly to the model. For example, by pressing on the pad in one area whilst striking, the areas on the model are depressed and struck as well, creating a far more dynamic and expressive instrument. The [Korg WaveDrum](#) was the first and the only attempt to produce a device like this. . The alternative is a breath controller such as the [Yamaha WX7](#) . These controllers aim to give as much flexibility to the user whilst retaining the familiarity of woodwind instrument. The results of coupling a device like this to a real time physical modeling system such as the [Yamaha VL/1](#) can be astounding.

Both these input systems rely on an already established and familiar interface, as do the keyboard based systems. They allow us to apply skills that we may have heavily invested time in learning to extend our capabilities but at the same time, restricting us to these. The Yamaha VL/1 and the Korg WaveDrum died a death in the marketplace due to their high prices. However – they were cheaper than a professional quality traditional instrument such as a violin, and were also professional *instruments* not just an electronic drum or another synthesising keyboard. It takes time to learn to control and play a new instrument, but the market is fickle and guided towards instant results, hence new and possibly fantastic instruments become stillborn because no-one can teach someone how to play an instrument no-one has seen before and to learn from scratch takes a lot more effort.

Hence the search for the perfect piano sound goes on. The keyboard is dead. Long live the keyboard - Pass me the Lego.