SYNTHESISERS

INTERFACE DESIGN & DEVELOPMENT

This article discusses and evaluates the design and implementation of synthesiser control and performance interfaces.

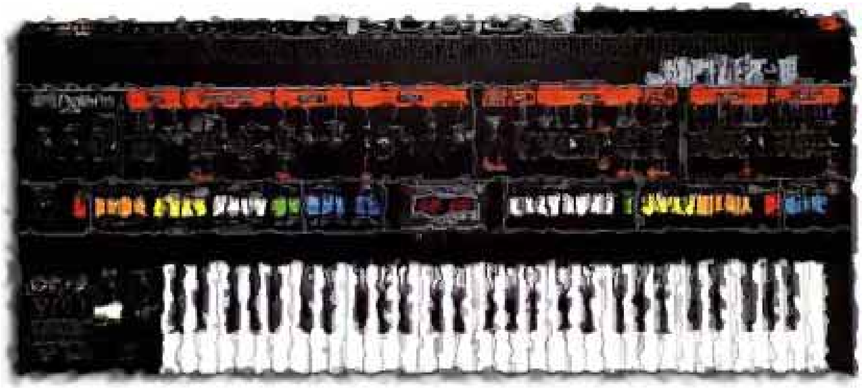

By Paul Hazel 1992

Authors Note

This article dates from 1992. A lot has changed since then. In particular, many of the problems with synthesizers discussed here have been at least partially addressed. In recent years there have been a range of new synths with 'strong' operating systems and a high degree of control (both in terms of the front panel and via MIDI continuous controllers).

However the discussions of interface design itself, and the extensive overview of the history of the synthesizer [that is, 95% of the article] remain accurate and relevant. Enjoy!

Contents

Click on a link to navigate to its page.

Moog Minimoog.

1: Enter

There is a common perception that there is something wrong with synthesizer interfaces. Synthesizers have become difficult and time consuming to use, to the extent that a lot of people have just given up on them. They never attempt to program new sounds or in any way delve into the inner workings of the machine, skimming across the surface of the 'cyberspace' within. If they want new sounds they buy them on RAM cards or CD-ROM's, or maybe just buy a new synthesizer!

Rather naively, I initially took this problem to be simply a function of the control panel: certainly on some synthesizers this may be a significant factor. I soon realised that in order to come to any satisfactory conclusions about the musician-machine interface, the whole philosophy behind contemporary synthesizer design would have to be looked at. This project is consequently divided into three main parts. The first part looks at the interface in its broadest sense, and includes discussions on cognitive psychology, software design principles, and the physical interface itself: the control panel. Within this section there is also a brief look at manuals, and it is considered here that they are in fact part of the interface between man and machine. The second part looks at the design and development of synthesizers and synthesis techniques from the beginning of the century, and the third follows on, taking as its starting point the introduction of all-digital synthesizers and MIDI.

Throughout, discussion will centre around two keynote papers. The first is by Barry Truax (1980), and is called 'The Inverse Relation Between Generality and Strength in Computer Music Programs '. The central thesis of this paper is that a computer system can be classified as existing on a

BACK TO CONTENTS

continuum: at the general end, a system is said to have weak procedures which require a lot of information from the user to generate any results. This type of system has the advantage of being very flexible and of having a high level of applicability. At the strong end, the system tends towards automation, with a consequent lack of choice on the users part. Somewhere in the middle is an area of maximum interaction, ease of use, and productivity for the user. Once aware of this type of dichotomy, it seems to spring up everywhere. In Gardiner & Christie (1987) it is expressed as the Generality v Power law, precisely in the context of interface design. And in Cole (1974) there is a discussion in these terms of the tension between rules and ambiguity in language.

The second key paper is by F. Richard Moore (1988) and is called 'The Dysfunctions of MIDI '. In it he coins the term control intimacy to describe the relationship a skilled performer has with an expressive acoustic instrument such as a violin. This relates to the immediate tactile and auditory feedback the performer would receive from such an instrument: it is this control intimacy that allows expression. That contemporary synthesizers lack such intimacy is understood: as Manning (1985) has put it, "the problems of immediate and effective communication between the composer and his tools...remains perhaps the greatest stumbling block throughout the evolution of the medium". Also in the Moore paper is some discussion of the implications of the MIDI standard on accurately capturing continuous control information generated by a performer: this will be looked at also.

2: The Interface

The modern synthesizer is essentially a highly specialised computer. Typically, one might say there are two main ways in which it would be used: Performance Mode or Edit Mode. In Performance Mode the digital hardware within will be controlled mainly by the keyboard, plus any other performance controllers being used such as the pitch-bend and modulation wheels, sustain pedal, or breath controller. The only functions likely to be accessed from the front panel are voice selection or perhaps master volume adjustments. In Edit Mode however, the front panel controls will become the primary means of interaction with the machine. The present state of the machine will be relayed by some sort of visual display, and the user will make a decision on the basis of that information as to what edits are necessary. Any instructions given to the machine via the controls will be read by the internal processor, which will in turn relay these instructions to the sound-generating hardware: the new status should then be reflected in the visual display and, depending upon the task, in the sound itself. The user will reassess the situation using this new information: if the changes are satisfactory then Edit Mode may be closed, but if not the process will be repeated. The user and the machine form an information loop, and could be said to engage in a dialogue.

If we are going to look at interface design in any detail, from the simple description above we can identify three main areas of concern: the user, the software, and the physical interface of the control panel that lies between them. These are discussed in some detail. There will also be a brief discussion here on manuals, which on certain important occasions will enter into the information loop.

Human Information Processing. The idea of the human being as a processor of information lies, paradoxically, in the

BACK TO CONTENTS

theoretical origins of modern computing. Both Turing and von Neumann compared the computer to the brain, its program to the mind. The claim was that whatever the brain did, it did because it was a logical system: the physics and chemistry, Turing said, were only relevant in so far as they were able to support these discrete states. Other advances in Linguistics, Information Theory, and Artificial Intelligence all helped to legitimise a return to the study of cognition (as opposed to behaviour), and the emergence of Cognitive Psychology as a major current in theory and research.

From the model developed by von Neumann, the human information processing system could be described as being:

- a sensory input system,
- a memory,
- a central processor, and
- a response (verbal or motor, for example).

Modern theory would tend to assume parallel processing rather than a strictly sequential processor, but the basic flow is still considered to be correct. This model will now be applied to the particular context at hand.

1) Sensory Input. The musician at the keyboard is going to be receiving auditory, tactile, and visual information. The editing task at hand will perhaps require some decisions to be made on the basis of auditory stimulus, such as the quality of sound or depth of LFO. Tactile information will constantly be received from the fingers, in particular as they manipulate the controls. But by far the most important information will be received as visual stimuli: the position of the hands in relation to the controls; the position of the controls on the

BACK TO CONTENTS

facia; and, crucially, the status of the machine as relayed by the data on the display.

The essential processes of vision are functions of the brain: the eyes are merely light receptors. Let it be sufficient to say here that the actual receptor organs are the light-sensitive nerve cells embedded in the retina, of which there are two types, rods and cones. In each eye there are approximately 6 million cones and 120 million rods. The cones are responsible for colour perception, but are not very sensitive to light intensity. Exactly the opposite is true for rods. Distribution of the two complementary types of nerve cell across the retina is unequal, with the cones more prevalent in the centre and the rods more prevalent at the periphery. (Explaining why we are able to see light fittings and TV screens flickering 'at the corner of the eye'.) At the very centre of the eye is a small depression called the fovea, particularly densely packed exclusively with cones. Although most cells in the eye do not have direct access to the brain, individual foveal elements do. The result of this is that the eye has a definite area of maximum acuity, located directly on-axis. Only objects focused upon the fovea are perceived clearly, the image become progressively more blurred toward the periphery. In order to construct a wider image in more detail the eyes move four or five times per second: these jumps are called saccades.

We could almost say that we sample the visual environment. How then does perception appear continuous? Well, in the same way that a digital recording system will sample a signal and then hold its value so that it might be measured, so it seems we have a Sensory Information Storage (SIS) system, sometimes called the Iconic Memory, that holds the visual

BACK TO CONTENTS

image and allows it to be processed. It is believed that the duration a detailed image may be retained in the SIS system is related to the photochemical reaction at the retina, with a typical duration of around 250ms and a maximum at around 500ms.

2) Memory. The second level of memory is the Short-Term Memory (STM), or Working Memory. The single most important factor relating to STM that all experimenters seem to agree on is that it has a strictly limited capacity. A rule of thumb measure was demonstrated by Miller (1956) in his widely-reported paper 'The Magical Number Seven Plus or Minus Two'. In a variety of situations he showed that people are only able to recall approximately seven items at any one time. The nature of the information is crucial: it would only be possible to remember seven unrelated 'bits' of information, but if a list or series of numbers could be easily grouped into what he termed 'chunks' then it became possible to recall seven of these chunks, with a corresponding increase in the number of bits recalled.

There are obviously many other factors complicating this. For instance, it has been demonstrated (Gardiner & Christie 1987) that doing mental arithmetic immediately after learning a short list will seriously hamper recall. This is called Interference. Repeating information over and over again in your head aids recall: this is called Rehearsal. Overloading, exceeding the apparent capacity of the STM, causes errors: anxiety further reduces that capacity.

One of the by-products of this limited capacity is that there is a good deal of relief when information relating to a particular task no longer needs to be retained. What happens then is

BACK TO CONTENTS

Roland SH101.

that operations are habitually broken down by the user into a sequence of semi-independent 'unit-tasks', allowing Closure to occur. This allows progress to be monitored more easily (at the same time localising the extent of any errors), and keeps the information load within the capacity of the STM.

The third level of memory commonly identified is Long-Term Memory (LTM). This is generally considered to have an effectively unlimited capacity, although as the brain is itself finite, so presumably is LTM. The problems associated with LTM are not those of what we call remembering, but those of Recall. In other words the information is there but we don't know how to get at it. This highlights the interdependence of encoding and retrieval: the Encoding Specifity Principle postulates that cues and prompts are linked to the stored information at the time of storage, and if these are lost the information will be irrecoverable. Closely allied to this is the use of mnemonics as an aid to memory. Key words, geographic locations, or other meaningful associations can be used to organise otherwise disparate material. This might in turn be linked to other theories relating retrieval to the depth of embedding of the information; quite simply how much processing was done on it.

One other interesting point is that we seem to have a superior memory for pictures (Gardiner & Christie 1987, Lindsay & Norman 1977). Both references attribute this to a double encoding: first as an image, and second as a verbalised abstract derived from the image. Needless to say this begs a plethora of new questions about the form(s) in which our memories are encoded.

3) Processing and Response. So: the information received by

BACK TO CONTENTS

the sense organs is perceived, interpreted whilst residing in STM, and possibly goes on to be encoded in LTM. These and any other activities occurring within the brain could be termed as Processing. However there is another more specific sense in which the information in STM is 'processed', whereby it is combined and compared with what is already known and used as a basis for decision making. And here, in the context of the synthesizer editing task at hand, we can identify a continuum of skill in problem solving from novice through to expert that depends entirely upon internal processing of this sort. We can differentiate the novice from the expert in two main ways. First, the expert has a large amount of information (of procedures, commands, underlying principles, etc) already in LTM, acquired either through rote learning, direct experience, or both. The novice, by definition, has little or none: each problem has to be solved 'for the first time'. In a worst-possible case, the novice may not even know what the problem is. Second, the acquiring of a skill that involves perceptual-motor movement will almost inevitably result in the user developing a degree of Automaticity. Routine tasks will tend to move outside conscious control, and the process is no longer limited by or resident in STM: we would say the skilled user displays smoothness, control, and economy of effort. This phenomenon has been expressed mathematically as the Power Law of Practice.

Software. The computer program within the synthesizer has two main functions: to operate upon data, and to display status and mode to the user. Data in this case will be related to the sounds themselves, both in the sense of sample data and the tables of individual parameter values that shape that raw data; MIDI commands; Digital Signal Processing (DSP)

BACK TO CONTENTS

algorithms; and various other 'housekeeping' utilities such as data dumping. All of this information has to be accessible to and manipulable by the user, and how this is done will have a major influence on the efficiency of the system. Any problems will be exacerbated by the fact that this large amount of data has to be presented on a relatively small display: it is obvious that the physical constraints of the synthesizer preclude large monitor/ TV screen displays. In practice what this means is that the information is arranged in a strictly hierarchical fashion, with the various levels represented by discrete 'pages' of information. Typically then, editing a sound would require the user to select the relevant sound, enter Edit mode, and then work down through the hierarchy to the necessary level. This may also involve various sub-tasks such as turning off those oscillators not being used, or muting effects. Having carried out the necessary edit, the hierarchy will have to be renegotiated in the opposite direction using the Exit button, at the same time as re-enabling any muted functions.

In terms of the Generality v Strength dichotomy, we would say that such a system is at the very 'general' extremity. It is designed in such a way that all tasks, from the most detailed to the very broad, are executed in the same way: the result of this is a 'weak' operating system that (as Truax points out) is typified by the user having to specify and input large amounts of information in order to complete the task. The user has to identify the problem, locate the function within the hierarchy, and remember the status of various sub-tasks. As described in the previous section, we can see that even such a basic operation as this can severely tax the users working memory. Furthermore, this type of system severely hampers Control Intimacy, a prime prerequisite of which is immediate

BACK TO CONTENTS

feedback to the user: this whole process is extremely time consuming!

The user of the system has a goal, a particular task to achieve: the primary function of the program should therefore be to enable them to achieve this goal as simply and as quickly as possible. The implication of this is that the programmer must start with the user interface as one of the, if not the, core design criteria. So, even given such a laborious general control system as described above, what are the factors governing software design that might go towards making it more efficient and user-friendly?

1) Unity and Form. The program must be well 'thought through', with a logically designed system structure. A common problem with this is that as the software gets updated, new versions of the system are often perceived as having a core program with new modules 'bolted on' around it.

2) Rationality. The program should display the programmers understanding about the way in which the completed system will be used, with the hierarchy of tasks divided into meaningful chunks. This in itself should reduce the number of unit-tasks necessary to achieve a certain goal. Also inherent within this is the idea of program flow, where certain operations are more likely to lead to 'Option X' than are others.

3) Consistency. Following on from the above, this means that the program should not contradict itself or confuse the user: in similar situations the program should behave in similar ways. The effort spent learning how to communicate one

choice or interpret one state will not have to be repeated for the next. Apart from anything else, a consistent program inspires confidence.

4) Communications Conventions. The system as a social animal: we could reasonably describe communication with another human being as a dialogue in natural language. Computers communicate in a synthetic language, and a good measure of the effectiveness of a command syntax would be its approximation to the conventions of human conversation.

There are many problems with this: although human conversation does have strict rules underlying it, it is often the very breaking of those rules that allow for expression. Machines, however, are absolutely literal. Secondly, a human dialogue is likely to include a good deal of non-verbal information derived from physical gestures such as body posture and eye contact. Thirdly, human conversation is what is technically termed as duplex, that is communication is possible both ways simultaneously. Man-machine dialogues are typically half-duplex: communication is two-way, but alternates in a cause and effect cycle.

Finally, the system is likely to use common or everyday terms in a rather specialised way, and the user is faced with the problem of 'translation'. (An existing example of this is the common confusion over the meaning of related terms such as loudness, intensity, and volume.) On another level, the program should also take into account the conventional ways certain data are presented. There is a strong tendency in synthesizer software to value everything from 0-127, regardless of the parameter. This might make life easier for the designer, but it certainly isn't how people work: filter cut-off, for example, should be calibrated in Hertz.

5) Feedback. This should be reliable. In the case of a synthesizer system feedback is always likely to be visual, for the simple reason that auditory feed back is less likely to be effective in a noisy, musical, environment. Accurate feedback has been shown to significantly reduce errors (Card, Moran, & Newell 1983).

6) Timing. This plays a crucial role in the effectiveness of feedback. If events generated by the user do not appear to happen with the human cycle time of around 50ms, the cause-event percept begins to break down. The fact that no feedback occurs to indicate the completion of the action is extremely distracting, destroying the psychological 'flow' of the user (Russ 1988, St Hippolyte 1989).

7) The Window. Earlier in this section we noted the distinction between the data stream operated upon by the program, and the display of status and mode generated by the program to represent that data. A window in this context is simply a space generated by the program to display that representation. It follows that the way in which a programmer designs the software will have a significant effect on the usability of the program, insofar as that data is more accurately, more clearly, or more concisely represented. There is a certain illusory quality to this: the user perceive themselves to be directly manipulating data, whereas there is a good deal of 'hidden' internal processing going on to translate information in and out of different formats purely for presentation purposes. The better the illusion, the more power the user will appear to have.

BACK TO CONTENTS

Visual ergonomics will be discussed in the following section.

The Control Panel. There was a time when it would have been accurate to say that a complex machine would have necessitated a complex control panel. With the introduction of digital technology, however, many of the control systems have been 'moved' onto the software, with the result that much contemporary machinery has an extremely sparse control panel. In a musical context, we have seen a radical change away from basically very simple synthesizers with many controls towards complex synthesizers with few controls. The controlling mechanisms now reside in the 'cyberspace' within the machine, accessed via a small display unit having only a tiny fraction of the available data visible at any one time. Hence that graphic description of synth editing: wallpapering the hallway through the letterbox.

It could be argued that having such a minimal control panel means that its elements have proportionally increased in importance. With so much control, and so few controls, it seems reasonable to assume that these remaining controls will suffer heavy usage. And with so much information now only visible via the display unit, one would like to think that it had been designed with the proper ergonomic principles in mind...

1) Control Layout. Editing a synthesizer is a skilled task, and it is a function of almost all skilled tasks that they depend upon eyesight and manual dexterity. With there no longer being any need for the control panel to reflect or make explicit the inner structure of the synthesizer, there is now absolutely no reason why the ergonomic factors likely to enhance the performance of the user should not take

precedence.

The display should be in the centre of the control panel. Black marks to those few synths (eg the Synclavier) that don't implement this. Ideally it would be recessed and angled so that it faces the user when they're in the normal position seated in front of the keyboard; manufacturing cost is probably the reason why this is so rarely done. The master volume control should be isolated and distinct, and usually is. For the rest of the controls, the general rule is that the most important should be placed centrally. There is a tradeoff here: the closer the controls the smaller movements will have to be, but compacting the controls makes accuracy increasingly important. Without such accuracy, more errors will be made. This is described by Fitts Law, which states the time to move the hand to a target depends only on the relative precision required, the ratio between the distance to the target and its size (Card, Moran, & Newell 1983; Russ 1988). A rule of thumb is that controls should be at least 15mm apart (Grandjean 1980).

2) The Controls. Editing a synthesizer is primarily a cognitive task: the perceptual-motor actions governing the use of the controls are an expression of that cognitive activity. It is important, then, that the physical controls have been selected with this relationship in mind. On a gross level, we could differentiate between the user inputting either discrete or continuous information into the synthesizer, mapped onto corresponding discrete or continuous controllers. In practice this is rarely the case. The all-digital synth has brought with it an information input system based almost exclusively on incremental/ decremental amounts, even though it is undoubtedly best suited to a continuous-controller type

environment. Things are changing for the better, however, and most synthesizers now sport at least one rotary encoder (also called infinity wheels and alpha dials).

Controls movements should comply with our 'stereotyped reactions'. A vertical slider, to take a simple example, should increment its related parameter value as it goes up. But what of a horizontally placed slider? Which end should be 'high' and which 'low'. Although this is more a concern for the industrial designer, it does highlight the types of pitfalls lying in wait for the unwary...

Discrete controllers include a whole range of different types of switches and buttons, too numerous too describe fully. Choice is likely to be largely on the basis of cost, although in certain contexts buttons with status LED's or specially textured surfaces may be required. One type of button that has aroused some controversy is the membrane switch. This was used extensively on the first generation of MIDI-equipped synthesizers, most notoriously the Yamaha DX7, presumably because they were considered to be protection against beer spillage and the like! Whatever, people didn't like them and later models had discarded them. The most common complaint was lack of feedback (Sanders & McCormick 1987). Because key travel was virtually zero users were unsure whether data had been entered or not; the lack of movement also tended to generate an excessive use of force. The use of graphics to 'show' where the key should be pressed, and either a central raised dome or an embossed rim all help to overcome these drawbacks. The most common form of feedback used with electronic equipment of all sorts, the bleep, is inappropriate on a synthesizer, limiting the continued use of this type of switch to more utilitarian functions (cash dispensers, industrial control equipment, almost anything outdoors).

Extensive research has been done on continuous controllers, usually in the context of controlling an on-screen cursor whilst doing text editing tasks (Card, Moran, & Newell 1983; Sanders & McCormick 1987). Of the controllers tested, the only ones of interest to us here are the different types of joystick, the trackerball, and the foot pedal. First, a distinction between two types of joystick: the Isotonic joystick is relatively free moving, and depends upon encoding by displacement. The further it moves, the greater the parameter changes. A second type is the Isometric joystick, which has a very limited movement and depends upon vector forces for encoding: in this case, the harder you push in a certain direction the greater the parameter change. The results of the research were consistent. Continuous controller devices were faster, easier to use (involving less mental effort), and more accurate than discrete devices. Isometric joysticks were more accurate than Isotonic, and trackerballs were the most accurate of all. The results should not surprise us: Isometric joysticks and trackerballs have been the mainstay controllers of the computer games fraternity for a long time. Movement, accuracy, speed, and ease of use are all equally crucial here.

Of the other continuous controllers, the foot pedal did significantly better than any discrete device, although it is less accurate and slower than the hand devices described above. Partly this is because of the larger mass of the foot, but perhaps more importantly the foot pedal is not likely to be visible in normal use, and the reliance of skilled tasks on sight has already been noted.

BACK TO CONTENTS

Korg Prophecy.

3) Text. When reading, we actually carry out two cognitive tasks in conjunction: a primarily visual one, used for searching and perceiving, and a primarily cognitive one that absorbs and comprehends the text. We tend to pass over individual letters and words in favour of larger meaningful units: we recognise grammatical situations not by detail but by phrase. New or unfamiliar words or constructs slow the reader down. We read quickly when there is a high degree of predictability, and therefore redundancy, in the text. On the basis of this information, we can make some decisions about the display.

First, the font used should be sans serif for clarity, and the information should be presented in upper and lower case letters. The ascenders and descenders lend words a characteristic contour, making them easier to recognise. Text should be of a size that can clearly be read from the usual operating distance. Areas of text with blank spaces amongst it have been found to be much easier to read (van Nes 1991). Space also helps keep search times to a minimum: paragraphs or headings can be highlighted by outlining, bold letters, capitals, italics, or a colour that differs from the background. The spatial grouping of text should reflect meaning.

These types of measures, in the context of a word processing task, have been shown to yield measurable improvements (Card, Moran, & Newell 1983; van Nes 1991).

4) Graphics. Images may present certain information more immediately: envelope and wave shapes are the obvious examples in the present context. As noted earlier, we seem to

be able to remember images very effectively, and there is also some evidence to suggest that our brains are able to process images very quickly (Gardiner & Christie 1987). Is this anything to do with the 'direct line' from the fovea to the processing centres?

A study to determine the effectiveness of symbols on machines (Harvard et al 1991) reported that symbols enhanced glance and distance legibility, but became more meaningful to users when used in conjunction with a text label. Symbols were better for speed, words better for accuracy. The major drawback seemed to be in finding symbols that had 'response consistency': this is exacerbated by the multinational nature of modern consumption, for it was found that cultural differences significantly affected interpretation.

5) Colour. Aesthetics aside, this is most useful when seeking to add contrast to, or differentiate areas of, the control panel and display. One particularly interesting and potentially useful aspect of colour for the designer is that its connotations tend to remain stable from one object to another: the obvious example is red, which is typically associated (in the industrial sense) with warnings of danger, fire, or just plain STOP! There are two main problems associated with the use of colour in this way. First, the user must be aware of the code and the way in which it works. Second is the problem of colour-blindness amongst users. This would obviously render any control panel coding invisible, but it is most crucial in the area of safety. Although not directly applicable to synthesizer interfaces, it is interesting to note how the colour codes of the three-pin plug have been designed with colour-blindness in mind: one dark

wire, one light, and one patterned.

On the display colour can be used to accent key words or phrases, helping the user to search for particular text items, and it is possible to group related items of text together with coloured areas. However, it is easy to over-accentuate, in which case the eye tends to be continually attracted to the coloured areas, which hinders reading: its indiscriminate use leads to a fragmentation of meaning. Finally, it may also be used with text to colour code meaning, over and above that actually within the text.

On the control panel, as well as being used for text-labelling of controls, colour can either be used to separate or associate different areas; for coding individual control functions; for accentuating small controls; and for grouping sets of controls.

6) Aesthetics. The control panel is a blend between aesthetics and technology. Although we would expect it to function correctly, there is also a strong element of pleasure to be obtained from a machine that pleases the eye. There is a tension here between designing from function and designing for purely aesthetic reasons. The designer must also take into account pressures from the marketplace, typically more ephemeral, fashionable, elements.

Most synthesizers are black. Those that aren't differ usually because the manufacturer has a 'house' style for their products. A good example is Akai: literally everything they produce is a light grey-blue colour with the Akai logo picked out in red. There are several reasons why black has come to predominate, none of them entirely convincing in their own right. There is the perception that black products are professional, expensive, powerful, even masculine, with the corollary that colour is frivolous, feminine, weak, cheap. Colour is also linked with toys and children. Whether or not in practice a black instrument is an advantage to the professional is open to debate. Certainly on stage it will tend not to reflect light, it won't show dirty marks so easily. Finally, referring again to the idea of products as multinational and cross-cultural, could black be seen as being merely neutral?

Simplicity and complexity are uneasy bedfellows. We have seen the synthesizer develop from being simple and knobular, to complex and minimal. Whilst the 'Jodrell Bank Approach' is obviously unsuitable for a modern digital synthesizer, there is the feeling that the 'Less is More' approach of recent years is equally misguided. There has certainly been a perception that hi-tech equipment doesn't needs knobs, and the more powerful is that equipment the less knobs it'll need (especially if it's black!) Here more than anywhere else we can see the effects of marketing on design, the triumph of aesthetics over function, almost completely divorced from the needs of the user.

The Manual. Whoever it was who claimed that computers would bring about the disappearance of paper from our lives was, at the very least, premature. Buy a computer system or a piece of software and you're likely to be confronted with at least one, and possibly several, dense and weighty tomes. These will have been translated from the language they were written in, probably Japanese, German, or Californian, and any attempt to read them will only confuse, dismay, and disorient you.

This common perception of 'user' manuals is probably a little out of date; manufacturers are much more aware of the problems involved. The manual should not be designed around the technical features offered by the system. In this case, the manual is merely an extension of the system. Rather it should be written from the point of view of the user, and must take into consideration the way in which the user will be likely to use that system: it must be as task oriented as the user. Central to this is the idea of an overview of the system, both in terms of its structure and in terms of its general context. In other words a sampling system manual should begin with some description of what sampling actually is, how it goes about doing it, and then finally how the system is structured internally. A completely separate section should give an index of step-by-step instructions for doing individual tasks.

There are two underlying ideas here, both related. First, that the user will initially be lacking in skill and will not use many of the more esoteric functions for some time. Consequently these should be 'out of sight' until needed: the users first priority is to become conversant only with the broad workings of the system. Secondly, as Wright (1988) has reported, the user principally learns about the system by interacting with it. Thus the manual should be arranged in modules, with the most basic immediately allowing and encouraging use.

Other prerequisites for a good manual might include: an extensive Glossary; comprehensive and accurate cross-indexing; examples related to skill level; clear diagrams; colour; and a professional and well designed layout.

BACK TO CONTENTS

Summary. Communication between man and machine takes the form of an information loop, a dialogue. The more free-flowing is that information, the more expressive the dialogue, then the better an interface could be said to be. This involves the concept of transparency, where the interface does not interfere with information flow and the user has the illusion of direct manipulation of data.

The primary limiting factor on the users performance is that of STM, or 'working memory'. It is crucial that the system works in a consistent and coherent way, and that it allows tasks to be readily and logically broken down, reducing demands on STM allowing closure. Furthermore, most of the users time will be spent doing very few tasks very often. It is important for the system designer to accurately specify these high-frequency tasks and implement them in the most efficient way possible.

On the control panel, it has been shown that continuous controller devices are much faster, more accurate, and cause the user least stress when inputting information. Whilst not suitable for all applications, it is proposed here that they are absolutely vital in a musical context, especially in terms of control intimacy during performance. Closely related to this, we have identified a typical modern synthesizer as having a very general, weak, operating system, requiring the user to input large amounts of data often. Given that music is an art that exists 'in time', these time consuming processes are rarely appropriate.

3: Design and Development.

Having now established a framework within which to discuss certain elements of

BACK TO CONTENTS

synthesizer design, most notably those that concern interaction with the user, this next section will look at the development of the ideas and technology lying behind synthesis and synthesizers. A selective history covers the period roughly from the turn of the century up to the end of World War II. The period from then until 1983 is divided into separate sections covering developments in Europe and America. An example of a particular production synthesizer is included, chosen to be representative of a particular stage in the development process.

1896-1945. "The first twenty-five years of the life of the archetypal modern artist, Pablo Picasso - who was born in 1881 - witnessed the foundation of twentieth century technology for war and peace alike: the recoil operated machine gun (1882), the first synthetic fibre (1883), the Parsons steam turbine (1884), coated photographic paper (1885), the Tesla electric motor, the Kodak box camera and the Dunlop pneumatic tyre (1888), cordite (1889), the Diesel engine (1892), the Ford car (1893), the cinematograph and the gramophone disc (1894). In 1895, Roentgen discovered X-rays, Marconi invented radio telegraphy, the Lumiere brothers developed the movie camera, the Russian Konstantin Tsiolkovsky first enunciated the principle of rocket drive, and Freud published his fundamental studies on hysteria. And so it went: the discovery of radium, the magnetic recording of sound, the first voice radio transmissions, the Wright brothers first powered flight (1903), and the annus mirabilis of theoretical physics, 1905, in which Albert Einstein formulated the Special Theory of Relativity, the photon theory of light, and ushered in the nuclear age with the climactic formula of his law of mass-energy equivalence, E = mc2. One did not need to be a scientist to sense the magnitude of such

BACK TO CONTENTS

Thaddeus Cahill inventor of the telharmonium.

The telharmonium.

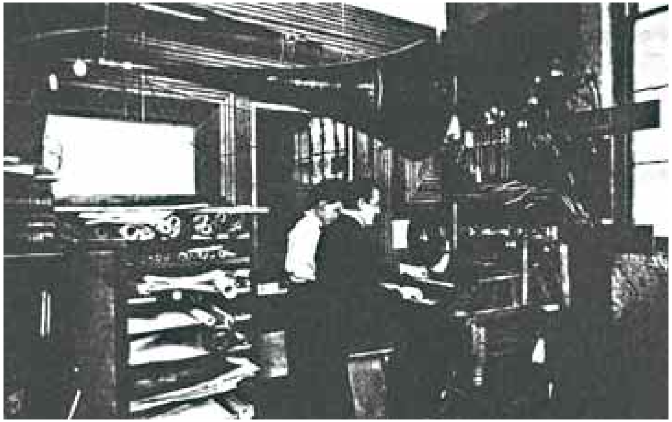

changes. They amounted to the greatest alteration of man's view of the universe since Isaac Newton". - Robert Hughes (1981)

In 1896 Thaddeus Cahill patented an electrically based sound generation system. It used the principle of additive tone synthesis, individual tones being built up from fundamentals and overtones generated by huge dynamos. Unbelievably huge: the instrument weighed 200 tons and was 60 feet in length. It had a conventional piano-type keyboard and was even polyphonic. First publicly demonstrated in 1906, this remarkable machine became known as the Dynamophone or Telharmonium. Cahill's vision was to sell production models of the machine to all the large cities in America, and to have concerts of 'Telharmony' broadcast into homes, hotels, theatres, and restaurants via the telephone networks. Needless to say cost, and the fact that it actually interfered with the normal workings of the network, meant that this grandiose scheme never came to fruition.

Although conceptually very advanced, the Telharmonium was already old technology: in 1907 Lee de Forest invented the vacuum tube. It primarily provided a compact means of generating continuous radio waves and of amplifying and detecting radio signals, but by extension also solved the problem of producing, amplifying, and processing all sorts of signals. 1907 also saw the publication of Busoni's influential 'Sketch of a New Aesthetic of Music ', which whilst it does not specifically refer to the production of music by electronic or mechanical means, exhorts modern composers to take the next step into "abstract sound, to unhampered technique, to unlimited tonal material". In 1910 the Futurist Balilla Pratella published 'The Technical Manifesto of Futurist Music ', a

The theremin.

clarion call to the composer as city dweller, openly embracing the machine age and all its implications. Their work, he said, should reflect "the musical soul of crowds, of great industrial plants, of trains, of transatlantic liners, of armoured warships, of automobiles, of aeroplanes". Whilst the Futurist romanticisation of war has rightly always been criticised, they nonetheless caught the spirit of the age with their frenetic and breathless art. In 1912 another Futurist Luigi Russolo published 'The Art of Noises ', a somewhat more considered and technically informed text than Pratella's, following it up in 1914 with what was possibly the first successful performance of absolute 'new music' at the Teatro dal Verne in Milan.

Partly because of the influx of Europeans, and partly because it had escaped the mass destruction of the First World War, the focus of development shifted to America. In 1924, Russian physicist cum instrument designer cum virtuoso violinist Leon Theremin demonstrated his new invention, variously known as the Aetherphone, Thereminvox, or more usually, simply the Theremin. A direct result of vacuum tube technology, the instrument remains unique in that it is played without being touched! It has two antennae that propagate low-power, high-frequency electromagnetic fields. Each field may be altered by the performer moving their hands within it. These alterations are then amplified and used to control the pitch and volume of sounds generated using a beat-frequency or heterodyning oscillator. (The difference between two supersonic frequencies creates the audio). The pitch antenna is a straight rod on the right side of the console, whilst the volume antenna curves like a shepherds crook and projects horizontally on the left. From this brief description it is obvious that the instrument needs performers of extreme skill and a very good ear for pitch, as there are no physical guides like frets or keys. It is probably this factor that limited the widespread use of the Theremin, along with the lack of original material for it. Having said that, the instrument is now available again with digital control and a decent MIDI spec from Bob Moog's company, Big Briar!

The twenties and thirties saw a number of other, largely unsuccessful, instruments being built: the Ondes Martinot, Dynaphone, Trautonium, Warbo Format Organ, Spharophon, and Givelet. Typically these were conventional keyboard type machines with a limited tonal repertory based around additive synthesis principles via sine wave generation, although some (such as the Warbo Formant Organ) did allow for reasonably complex filtering. The Givelet was unusual in that it combined electronic sound production with control by pre-punched tape. Oscar Vierling's Electrochord and the Miessner piano used strings to produce sounds, with movable capacitor pickups allowing tonal variation. In 1927 Les Paul built his first solid-bodied electric guitar.

Significant progress was being made in other quarters. Electrical recording, gramophones, and radio had all developed hand in hand. In 1935 the German company AEG produced the Magnetophon, the first modern tape recorder. Although still of relatively poor quality, it used plastic tape coated with ferrous particles as its recording medium. This was a vast improvement over the steel tape used previously: it could be cut and therefore edited, it was much lighter, and much safer. (Breakages of steel tape were notoriously hazardous.) In America Dudley Homer at Bell Labs developed the Voder, and then the Vocoder in 1936. These were offshoots from the telecommunications industry, Bell Labs being the research arm of the American Telephone and Telegraph company (AT&T). What the Vocoder did was to analyse speech sounds using an array of bandpass filters, and then generate a series of control voltages from envelope followers. The idea was that it would be these control voltages that would be transmitted, rather than the speech itself, and these would then be decoded at the receiving end. Because these control signals had a much lower bandwidth than speech, it was hoped the system would greatly increase effective channel capacity. Although of only limited use in a purely musical sense, the importance of this work lies in its relation to acoustics and psychoacoustics, information theory and sampling theory, all just around the corner…

Europe. World War II had in itself been a spur to technological innovation: much progress had been made with radio and radar, von Neumann and Turing were laying the foundations of modern computing, and of course the atomic age had been born. In Europe, rebuilding was the order of the day.

Pierre Schaeffer was an electronics engineer who had risen through the ranks at Radiodiffusion Television Francaise (RTF) in Paris. As early as 1942 he had persuaded the corporation to support research into musical acoustics. Inspired somewhat by the Futurists, Schaeffer developed the technique of recording naturally produced sound events, and in 1948 embarked upon a series of compositions using these sound events as source material. Hence the music of the 'Paris School' came to be known as Musique Concrete, and the recorded sounds came to be known as Objets Sonores. On a theoretical level, concrete can be taken to represent the opposite of abstract. To make a parallel with painting, we could say that a typical Mondrian or Kandinsky was totally abstract. A Cezanne landscape or Matisse interior, on the other hand, could be likened to objets sonores: whilst they do not directly represent the real world, they are nonetheless deliberately derived from it, but transformed by the painter into something unique, personal, and only 'of itself'.

Recording equipment consisted initially of direct-to-disc cutting lathes. Schaeffer experimented with removing attack portions of sounds by manually manipulating a volume control between the microphone and the recorder, simply not recording them. He played discs backwards and at differing speeds, re-recording the results onto another disc. With the arrival in 1951 of tape machines and a brand new studio, new techniques were developed. The Morphophone was an early tape echo machine, with a row of twelve playback heads instead of the usual one. Two other machines, called Phonogenes, were designed to play back pre-recorded tapes at different speeds; one had a continuously variable pitch range, the other was controlled by a conventional keyboard. Some experiments were also carried out with sound diffusion, using a sound projection aid called the potentiometre d'espace. This would be used to manually control the movement of one channel of audio on a five-track tape. The other four tracks were each sent to one of four loudspeakers. It is interesting to note here that rather than providing 'surround sound' in the arrangement we know as quadraphonic, one of the speakers was placed on the ceiling, allowing the illusion of vertical as well as horizontal movement to be created.

In 1948 Dr. Werner Meyer-Eppler, then director of the Phonetics Department at Bonn University, was visited by Homer Dudley and given a demonstration of the Vocoder. Suitably impressed, he used the machine in the creation of a tape illustrating a lecture on electronic sound production. In attendance was Robert Beyer of North-West German Radio.The pair struck up a relationship, with Meyer-Eppler the theoretician and Beyer the technician. They were joined by another influential figure in Germany at that time, Herbert Eimert, a radio producer for West German Radio (WDR) in Cologne, music critic, and composer. It was he who was instrumental in WDR broadcasting 'The Sound World of Electronic Music' in October 1951. The programme featured a discussion and tapes of sounds 'constructed' by overdubbing the simple tones generated by a Melochord (designed by American Harald Bode). On the same day WDR agreed to establish an electronic music studio for them. In 1953 Karlheinz Stockhausen joined the studio.

The whole approach to the creation of sound and of composing was radically different than the Paris School: in fact, initially, they were diametrically opposed. Whereas Schaeffer was taking complex sounds and transforming them, the Cologne studio affected a 'Year Zero' approach. Complex sounds were laboriously built up by overdubbing simple tones, initially using only the most basic equipment: tape machines; a single sine oscillator; a white noise generator; filters; and later, reverberation. Stockhausen describes one such process during the production of Gesang der Junglinge, quoted in Kurtz (1992):

"I invented completely different processes in which the three

BACK TO CONTENTS

of us - myself and two musical and technical collaborators - each used a different piece of equipment. One of us had a pulse generator, the second a feedback filter whose width could be continuously changed and the third a volume control (potentiometer). I drew graphic representations of the processual forms. In one such form, lasting twenty seconds, for example, the first of us would alter the pulse speed, say from three to fourteen pulses per second, following a zigzag curve; the second would change the pitch curve of the feedback filter, in accordance with another graphic pattern; and the third - using yet another graphic - would change the dynamic curve....So we sat down to realise one of these processual forms, one of us would count 3, 2, 1, 0, then off we went. The stopwatch was running, and at the end of twenty seconds each of us had to be finished."

This exact method of sound synthesis was also applied to composition. Eimert based his compositions on measure and number; Meyer-Eppler proposed statistical compositional techniques derived from information theory, and in his classes students were encouraged to create texts using cards, lotteries, roulette, or telephone directory numbers! Underlying it all were the serial techniques of Schoenberg and, especially influential at this time, Webern.

Although the Paris and Cologne schools started out from opposed positions, as time went on the hardline stances were softened until a 1967 piece like Stockhausens Hymnen is using all available techniques regardless of their origin. From being a primarily studio bound medium many experiments were carried out with mixed live and electronic performances, sound diffusion scores, and aleatory performance scores. As the 1950's progressed many more

BACK TO CONTENTS

studios were started: in 1955 the Studio di Fonologia Audizioni Italiane in Milan; in 1956 Japanese Radio (NHK) in Tokyo; 1957 saw new studios in Warsaw, Munich (Siemens), and Eindhoven (Philips). A studio in Stockholm and the BBC's Radiophonic Workshop in London followed in 1958.

A final word on French composer Edgard Varese. A unique and uncompromising personality, his work is at once typically modern and yet completely distinct from any movement or school. His work in the twenties used conventional orchestral resources (often scoring unorthodox playing techniques in order to coax new sounds from traditional instruments, and then adding sirens and a whole scrapyard of drums and percussion), culminating in Ionisation of 1931: percussion as pure sound. For the next twenty years he composed virtually nothing, desperately trying to find the money to build his own studio: he even approached one of the Hollywood film companies. By the 1950's everyone had caught up with him and he started composing again; Deserts of 1951-54 for orchestra and prepared tapes; and finally Poeme Electronique, an all-tape piece commissioned by Le Corbusier for the Philips Pavilion at the Brussels World Fair of 1958. One of the great men of modern music.

America. The musical climate in America was very different from that in Europe. The war years had seen another influx of intellectuals and artists into the country with Stravinsky, Schoenberg, Bartok, Hindemith, Milhaud, Krenek, Martinu, and Varese being the most prominent names amongst the composers. Perhaps the most profound effect of this was that it left a vacuum behind in Europe, which, as we have seen, sucked into it the most adventurous and forward-looking of the new composers and theoreticians, now unhindered by the

weight of tradition. This European tradition now took root in the Universities of America, and it meant that for a long time the established music departments had relatively little to do with the emergent new music: that was left to the scientists.

Because America did not have a state-sponsored radio network it meant that support for electronic music studios was hard to find. The only place that had enjoyed consistently supported research into the applications of electronics in music was Bell Labs, the research arm of AT&T. The technological influence of the scientists at Bell Labs cannot be underestimated. By 1945, Harry Nyquist had already outlined sampling theory. In 1947 John Bardeen, Walter Brattain and William Shockley had invented the solid-state transistor, a contemporary 'inventing of the wheel' without which the modern communications revolution would not have happened. To put the cap on it, in 1949 Claude Shannon published his work on information theory.

Unlike Europe, very little work was done in America using tape. Louis and Bebe Barron had a studio in New York from 1948, where people such as John Cage and Morton Feldman became involved in a short-lived project called 'Music for Magnetic Tape'. Varese re-worked the tape sections for Deserts here, but the studio is probably most famous for the film soundtrack to Forbidden Planet (1956). From 1951 through to 1959 Vladimir Ussachevsky and Otto Leuning worked on various tape pieces without ever actually ever being able to secure enough funds to build a studio. However persistence won the day, and the Rockefeller foundation eventually provided $175, 000 for the foundation of the Columbia-Princeton Electronic Music Centre. The system was to be based around the RCA synthesizer.

BACK TO CONTENTS

Started in the late 1940's, the RCA synthesizer was designed by two electronic engineers Harry F. Olson and Herbert Belar. Apparently inspired by Shannon's work, the machine was designed to generate compositions based upon statistical probability and to play them back (monophonically). The basis for the compositional system lay in the analysis of the statistical characteristics of Stephen Foster's folk songs. Based on vacuum tube technology, the machine generated sound via sets of tuned oscillators, filters, LFO's, and resonators: two control channels were available. Everything was controlled from a punched paper tape, which was manufactured via a typewriter-style keyboard. Each tape had 36 columns of information (= 36 rows of holes), 18 for each control channel. Once running, electrical contacts were made between a brush and a drum through the punched holes. The machine had direct outputs to loudspeakers, and the results could be recorded onto a direct-to-disc cutting lathe. Version 2 (1959) had expanded voice facilities, now controlled by two synchronised paper tapes, and had a four-track tape machine instead of the cutting lathe.
As a synthesizer, the machine was very limited. As a compositional tool it was deeply flawed by the superficial level of the analysis of the songs, especially the rhythmic aspects. Inputting material was presumably laborious, definitely non-interactive: and like Cahill's Telharmonium in its own time, it was already a technological dinosaur. What is interesting about the machine is that it is the first 'music workstation', bundling sound manipulation, note (event) sequencing, and master recording, all into one centrally controlled unit. Great idea, shame about the music!

Meanwhile, back in Bell Labs, violin-playing electrical engineer Max Mathews was beginning what was to become the mainstream of the American new music. In 1957, using an IBM 704 valve computer, Mathews developed a program called MUSIC I: it generated an equilateral triangle waveform which was converted into audio by an Epsco 12-bit vacuum tube digital-to-analogue converter (DAC). The user could specify pitch, amplitude, and duration for each note. From this incredibly primitive beginning (it must have seemed amazing at the time...) Mathews quickly developed MUSIC II and then MUSIC III in 1960. MUSIC III was notable in that it was written for the first transistorised computer, the IBM 7094. It was also the first to introduce the concept of the unit generator. These were basic 'building blocks' corresponding to the functions now commonly associated with analogue synthesizers, such as oscillators, adders, noise generators, and attack generators: thus the user could build up their own orchestra, as Mathews termed it. The main problem associated with computer music was the vast amount of data that had to be input by the user to generate an event. In order to specify a particular sound at a sampling rate of 30 kHz, for example, 30,000 numbers would have to be supplied every second to determine the pressure fluctuations alone. Multiply this by the number of other parameters that are needed for filters, LFO's, pitch changes, and the like, multiply it again for the number of voices used, and then multiply it again by the number of seconds that elapse whilst the piece plays...and all this had to be input from a QWERTY keyboard. Should there be any mistakes, or should the results not be quite what was intended, it would have to edited and re-computed. We are in deepest Truax territory: these are the ultimate in 'weak', general systems, where literally every last detail of a sound has to be determined and encoded. Unit generators were one way of alleviating the problem by having 'off the shelf' modules at the users disposal, making the system 'stronger' at the expense of some flexibility.

A whole range of other systems were developed on mainframe computers at Universities throughout America, almost all based on Mathews original model. In general the music that came out of these studios could be said to be primarily concerned with texture and timbre, having relatively few 'note-events'. Extensive research was carried out on the analysis and resynthesis of instrumental timbres, and later developments allowed the digital sampling of natural sounds using an analogue-to-digital converter (ADC), and at Stanford John Chowning developed a synthesis technique called Frequency Modulation (FM). As computer processing power increased so the systems became more sophisticated, allowing input from piano-type keyboards, graphic displays of information, cross-system portability, and easier programming.

A good example of a mature computer music piece is Mike McNabb's Dreamsong, realised at the Centre for Computer Research in Music and Acoustics (CCRMA) at Stanford. The piece took two years to complete, and employs techniques such as FM, sung vocal processing and resynthesis, and additive synthesis, with crowd and speech sounds processed by flanging, comb-filtering, Doppler shifting, and panning. The program used is MUS10, a descendant of Mathews' MUSIC IV, computed on a DEC KL-10 mainframe. Unit generators are still in evidence. The voice sounds were recorded digitally and analysed: in the finished piece the original recording and the resynthesised voice are both used. Stunning sonic transformation are made between bells and voices, voices and other purely synthetic sounds. There is a limited amount

BACK TO CONTENTS

of harmonic material, but on the whole the piece flows organically, "from the real world to the dream realm of the imagination, with all that implies with regard to transitions, recurring elements, and the unexpected" (McNabb 1981).

In 1964, Robert Moog presented a paper entitled 'Voltage-Controlled Electronic Music Modules' at the annual convention of the Audio Engineering Society (AES). This became the blueprint for a whole new generation of relatively affordable, portable, synthesizers. The miniaturisation of the electronic components that had been made possible by transistors allowed Moog to develop the concept of a modular synthesis system. These modules could then be 'patched' together in user-determined combinations, the common link between them being control voltages (cv's). Very quickly other companies (Buchla, EMS, ARP) picked up on the idea, and for the first time electronic synthesis reached the general public. The machines were generally of two types: large systems which tended to have the keyboard and electronics separate, and modules which had to be physically patched together with leads; and much smaller machines with a hard-wired control flow and a 21/2 or 3 octave keyboard built in. An example of the former would be the Moog 3C; an example of the latter the MiniMoog. They were all monophonic, with non touch-sensitive keyboards; they tended to be unreliable; most had problems with oscillators drifting out of tune. From a purist point of view, another problem could be that the mass production of these machines defeated the purpose of having a synthesizer. There were now studios and individuals the length and breadth of the country with identical machines: with mass production came the idea of manufacturers dictating a single design policy. In practice, the vast majority of commercial machines were subtractive synthesizers, using filters to remove harmonic material from relatively complex wave shapes such as sawtooth, square, and triangle.

Development of this type of 'analogue' synthesis, as it became commonly known, continued throughout the 1970's. The market came to be increasingly dominated by Japanese manufactured instruments. Synthesizers became polyphonic, and just as they were becoming unwieldy they developed digital control. This was primarily dependent upon multiplexing, yet another telecommunications spin-off that allowed a single processor to do lots of jobs (as opposed to lots of processors all doing one job each, which was far too costly). Digital control in its turn brought with it voice memories, stable oscillators, increased polyphony, and, eventually, digital communication: MIDI.

This is not to say that there was no common ground between the essentially academic/ scientific world of computer music and the increasingly 'pop'-oriented market of analogue synthesis. In 1970 our old friend Max Mathews developed GROOVE, a hybrid system consisting of a 'minicomputer' connected to and controlling an analogue synthesizer. What made it so interesting was that it was an attempt to solve the problems of performer interaction with a computer system. Compositions could be made in the normal way. What the system then encouraged the user to do was to play back the composition, at the same time recording performance gestures enacted via a joystick and rotary controllers. These performance gestures could then be edited if necessary, a process which was aided by a graphic display. The emphasis throughout was that of interaction.

Another interesting system was MUSYS III, developed in

BACK TO CONTENTS

London by Peter Zinovieff using the profits made from his EMS company. The system used two PDP 8 computers controlling a bank of 252(!) oscillators, 64 band-pass filters, 12 tuneable filters, a white-noise generator, a percussion generator, 9 digitally controlled amps, envelope shapers, plus a whole array of manually controlled equipment like ring modulators, filters, and reverb units. Information could be input via a normal keyboard, QWERTY keyboard , or a special console with a 'spinwheel' that could be used to manipulate the sequencer register position. This incredible system was dismantled in 1979 through lack of funds, with only a handful of works having been completed on it.

The Sequential Circuits Pro-1

Example 1: Sequential Circuits Pro-1. Sequential Circuits were a well-respected and innovative American synthesizer manufacturer, releasing the Pro-1 in 1981. It is a small monophonic subtractive synthesizer, comparable with the MiniMoog or Roland SH-101 in terms of voice architecture and features. Measuring 65 x 40 x 12 cms, the frame of the machine is made from a single piece of folded steel, which has been covered with some sort of black vinyl. A 3-octave non touch-sensitive keyboard sits in the front, with pitch-bend and modulation wheels to the left: all 'feel' pretty awful. The black control panel is made of fibre board and is simply dropped into the frame and held in place with four screws: the electronics are mounted on the back of the panel, with its array of knobs on the front. The knobs are all absolutely identical to one other, and laid out 'by module'. For example, the filter controls are grouped together and 'roped off' with a white band on the panel. Each individual control within the group is labelled. The groups themselves are laid out to represent a control flow from left to right across the panel, with the exceptions of the modulation and master control (volume and tuning) groups, which act as bookends. The machine is finished off by screw-on wooden end-pieces, serving no function other than aesthetics. All in all, this is a seriously tacky piece of equipment!

The low build quality is compensated for in other ways: it sounds excellent. The Pro-1 has two beautifully warm sounding audio oscillators, an excellent resonant filter, a single LFO with extensive routing capabilities, glide, and a primitive but nonetheless useful sequencer/ arpeggiator. It suffers from the usual problem of tuning drift: the owner tells me has has to regularly take off the front panel and manually adjust the oscillators. The Pro-1 also has a reputation for unreliability. But the beauty of this type of machine is that any and all edits can be made immediately, and words like interactive, intuitive, and 'enjoyable' spring to mind when using it. In psychological terms, the machine puts a very small load on working memory. Single functions are represented by single controllers; whilst playing, movement can be created in sounds by control manipulation; immediate audio feedback allows you to get the sound just so, in a way that normally isn't possible with digital machines. In this respect, control intimacy is very high. However, the lack of a touch sensitive keyboard and the appalling controller wheels on this particular synthesizer count against it. In terms of the generality v strength dichotomy we would say this machine is very strong. It does a limited number of things, but allows the user to do them quickly and easily.

The manual, by Stanley Jungleib, is very well written. There are voice and modulation flow charts, good pictures, and a circuit diagram. Factory presets in the shape of control panel layouts are given at the back, and blanks are provided for the users own sounds. There is even a bibliography. There are things missing that I would like to have seen. First, an overview of the way the instrument works, including some discussion of the idea of subtractive synthesis (a phrase which doesn't occur in the manual). Second, a glossary explaining the meaning of all the technical terms, and for obvious reasons: the considerable problems encountered nowadays in this area must have been worse ten years ago when the technology was new.

Summary. The twentieth century has seen an unprecedented acceleration in the rate of technological change. This has transformed society and, as a result, changed the way artists work. For the musician, usable and affordable technology in the form of tape machines and primitive synthesizer systems did not arrive until the 1950's. Most of the core research was completed in America, especially at Bell Labs.

The mainstream of European new music became centred around state-owned radio stations, and initially focussed on composition using magnetic tape. The two main studios in Paris and Cologne had very different theoretical outlooks to start with, but soon coalesced within a much broader concept of Electro-Acoustic Music.

In America the emphasis was on computers in music, both as an analysis/ resynthesis tool and as a composition medium. Development in the musical domain therefore became very closely linked to, and dependent upon, new discoveries in computing and the telecommunications industry in general. Thus as computers became smaller and faster, and as memory grew, so the music itself became more complex and technologically demanding. The parallel development of modular synthesizer systems allowed mass production, bringing with it a certain homogenisation of techniques.

Digital electronics had by the late 1970's reached the point where it became cheap enough and small enough to start appearing in commercially available synthesizers as a control mechanism. The next stage was the totally digital instrument.

4: MIDI and Beyond

In this section we look at the events leading up to the introduction of the first official MIDI spec, plus some discussion of the technical issues involved. A brief history of Japan's economic recovery since the World War II is included, hopefully shedding light on some of the forces that have gone towards shaping the contemporary synthesizer. Where appropriate, examples illustrate this continued development. Finally, some suggestions are made concerning possible future directions.

MIDI. Sometime at the beginning of 1981, the idea was in the air that synthesizers could be designed that would have control facilities enabling them to 'talk' to each other, in the same way that computers did. Someone at Roland spoke to someone at Oberheim who eventually spoke to Dave Smith, then president of Sequential Circuits. In October of 1981 he made a proposal at the AES convention for a Universal Synthesizer Interface (USI). It specified a serial data format at 19.2kBaud (19, 200 bits per second), with connections via standard quarter-inch jack plugs. This was followed up by a meeting at the January 1982 National Association of Music Merchants (NAMM) show of almost all the current synthesizer

BACK TO CONTENTS

manufacturers, American and Japanese, expressly to discuss the USI. Opto-isolation was added, to prevent ground loops, and the serial rate was upped to 31.25kBaud to try and eliminate timing delays inherent in a serial system. Then, strangely, the American companies seemed to lose interest in the idea, and Smith then carried on development primarily with the Japanese companies, especially Roland. They suggested an extension of the USI standard that included the separation of data and status bytes, and for a while it became known as the Universal Musical Instrument Interface (UMII). Smith eventually came up with Musical Instrument Digital Interface (MIDI), partly on the basis that some legal complication would be likely using UMII. MIDI was publicly announced in Bob Moog's column in October 1982's issue of Keyboard magazine. At the January 1983 NAMM show, a Roland JP-6 was successfully hooked up to a Sequential Circuits Prophet 600: all very symbolic.

The first official MIDI spec was released in August 1983. Briefly, MIDI is an open ended system of serial transmission. Information is divided into two classes, data and status, and is always sent in byte-sized chunks. The initial specification was primarily to transmit note-on and note-off information: in other words it was designed to be an event-oriented system. Since its inception it has continued to be expanded, with features such as MIDI Time Code (MTC) for synchronisation purposes, MIDI song files, sample data-dump standards, and more recently the definition of real-time controller numbers for timbral modifications.

The main problem levelled at MIDI has been serial transmission induced delays. At the most basic level, notes struck simultaneously on a keyboard will become serialised

BACK TO CONTENTS

Yamaha DX7.

within the data stream, with the effect that temporal smearing occurs. This has been discussed at some length in Moore (1988), and in particular he demonstrates that for accurately capturing the information generated by a skilled human performer the MIDI bandwidth is too low. There is a certain implication in this piece that if synthesizers were ever to approach the level of control intimacy enjoyed by an expressive acoustic instrument, then it would far exceed MIDI's channel capacity. In practice, this level of detailed timbral control is rarely used: a chicken and egg situation?

In practice, MIDI delays remain below the level of perception. A typical system with many MIDI channels being run from a sequencer has a resolution sufficient for most purposes. With a MIDI command normally requiring 32 bits, and with MIDI running at 31.25kBaud, it is capable of resolving successive events to approximately 1ms. Lennard (1992) has shown that a far more significant timing error is being generated by the synthesizers themselves: the delay between them receiving a note-on command and actually generating a sound. Depending on the machine, delays of between 2 and 14ms were recorded when playing a single sound; in multi-timbral mode, however, these rose to between 14 and 40ms!

Example 2: Yamaha DX7. Although released only two years after the Pro-1, the DX7 seems to come from another technological era altogether. The first mass produced all-digital synthesizer, it remains to this day the best selling synth ever. It is a very unusual dark green colour, and comes in a pressed steel casing which hinges at the back. The top is held down with four large screws, clearly visible, a strangely anachronistic detail. It measures 102 x 33 x 10cms, and weighs in at a massive 14.2kg. There is a 2-line 16 character

BACK TO CONTENTS

Liquid Crystal Display (LCD) and a 2-digit Light Emitting Diode (LED). Apart from the data entry and master volume sliders, all other front panel functions are accessed by membrane switches. It is equipped with a spartan MIDI spec that only allows transmission on Channel 1, but it has the full complement of In, Out, and Thru ports. The 5 octave keyboard is excellent, and is both velocity and after-touch sensitive. It is 16-note polyphonic, which must have seemed miraculous in 1983.

The operating system of the DX7 is divided into four modes, each accessed by their own dedicated function button: Play, Edit, Function, and Store. Once in a mode, the 32 membrane switches to the right hand side of the display call up the individual parameters. For instance, in Edit mode there are 6 switches to turn the operators on and off, a switch each for LFO wave, speed, delay, PMD (= frequency deviation), AMD (= depth modulation), and sync; and so on. Each of the mode buttons is a different colour: Play is green, Edit is blue, Function is brown, and Store is red. The 32 switches on the right each have three sets of labels (the fourth, Store, doesn't use any of these switches), with each of the labels related to a mode by colour. So in Edit mode a switch does what the blue label will says it will do, and in Function mode it does what the brown labels says, etc.

The interface is actually very good! Because there are lots of switches, the operating system is 'broad but shallow'. You don't have to keep digging down within a software hierarchy to adjust a parameter: almost everything is one switch away. Almost all the information you're likely to need is displayed on the front panel in a well-organised format. You don't ever have to remember the individual functions of the switches.

Data entry can all be carried out with the slider, and also it's possible to assign the mod wheel for this purpose. I personally even like the membrane switches, great for running your finger down. Overall the machine manages to strike a reasonable balance between generality and strength, greatly helped by its monotimbrality which limits it nicely.

It isn't perfect: the LCD is not back-lit, limiting the viewing angle considerably. Once programmed, it is not really possible to interact with sounds. The potentially very interesting (in terms of control intimacy) breath controller is limited to boring old LFO depth. The colour coding on the front panel is well implemented only providing you recognise it as being a code: the person who gave me access to this machine didn't. Even so, why has this machine got the reputation for being hard to program? The answer lies in the synthesis system itself.

The DX7 produces its sounds using FM, a technique developed by John Chowning at the Stanford Artificial Intelligence Laboratory. At its most simple, FM consists of the modulating together of two sine waves: its beauty is this economy of means. On the DX7 these sine waves are generated using 8-bit wave tables. One wave is called the carrier, the other the modulator. A carrier, then, is modulated (reasonably enough) by a modulator, and the amount of frequency deviation generated in the carrier is dependent upon the amplitude of the modulator. From this, another measure called the Modulation Index is calculated, from the peak deviation in the carrier being divided by the frequency of the modulator. For values above 0, harmonic spectra are generated equally either side of the carrier frequency. As the modulation index increases, so does the bandwidth of the generated spectra:

energy is 'stolen' from the carrier and distributed amongst the harmonics. The problem is that increasing the modulation index linearly does not bring about a similar linear increase in the amplitudes of the harmonics. These are determined by Bessel Functions, and they progress in a decidedly non-intuitive way.

The only way to really get to grips with FM programming is with experience, and plenty of it. It is interesting to see that Yamaha have to a certain extent given up on FM, and most of their current synthesizers are hybrid instruments with large sample memories and a much watered-down FM capability.

The original DX7 manual is a brief 30-page affair. There is a bare-bones description of FM that could do with some flesh on it. There is no explanation of technical terms and no mention of the control panel colour coding. A brusque dash through "Let's actually create a sound" and that's it...

Far better is 'The Complete DX7 ', by Howard Massey (1986). Massey teaches at the Public Access Synthesizer Studio in New York, and this book was developed from his courses there. It shows. Chapter 1: Basic Audio Theory; Chapter 2: Front Panel Operation; Chapter 3: The Operator; this is more like it! There is a quick reference guide, good pictures, and each chapter includes a number of hands-on exercises ranging from 'Brightening sound of E. Piano 1' through to 'Animating a sound by using a carrier in sub-audio fixed frequency'. The only blot in his copybook is that there is no glossary.

Japan since 1945. Since the introduction of MIDI, the Japanese manufacturers have dominated the 'hi-tech' music

BACK TO CONTENTS

market. So much so, that it is now impossible not to view most contemporary synthesizers as 'Japanese objects', which is certainly strange in view of their Western development history. Of course, looking more widely we can see Japanese manufactured objects all around us: TV's, VCR's, radios, tape decks, microwaves, CD players, cars, plant-hire equipment, generators,...you name it. It is interesting to note that absolutely all of this is dependent upon micro-electronics, either directly or indirectly: those objects that are not electronic in nature owe their build and quality to Computer Aided Design (CAD) techniques and Computer Aided Manufacturing (CAM). How? Why?

On the 15th of August, 1945, Japan surrendered to America and brought to an end World War II. Japan's cities and industries had been all but destroyed by bombing. There was virtually no food because of their heavy reliance on imports, no skilled labour, and no raw materials. Under General MacArthur, the Americans assumed power, and almost immediately began to pump aid into the country, partly because of the fear of Chinese Communism: the Americans wanted a Far Eastern base. $2 billion went into industry. They provided new technology, machinery, and manufacturing processes, and in 1950 a statistician named W. Edwards Deming was sent to Tokyo to supply management and quality control training. Deming and the Japanese were made for each other, and they instinctively understood the things he was teaching. Virtually single-handedly Deming introduced the 'quality revolution' into Japanese manufacture, and with this went a whole philosophy about how companies can be run by worker involvement, cooperation, iterative product improvement, flexibility, and an innovative approach to problem solving. The key phrase is integrated manufacture.

This fitted neatly with the traditional Japanese work ethic, derived from Buddhist and Confucian ideals of loyalty, respect, and pride in ones work.

Another crucial factor in the rebuilding program was the Ministry of International Trade and Industry (MITI), which had overall control of allocating government monies. In 1950 MITI drew up a shopping list of preferred Western technologies, and over the next 15 years they steadily bought up licences and patents, with communications as one of its priorities. The emphasis was on miniaturisation and quality, miniaturisation being particularly important because of export costs. By the early 1960's companies such as Sony were making inroads into the American market with things like transistor radios and portable TV's. By 1962 Japan's economy had overtaken that of Britain, and by 1967 West Germany's (Evans 1991). The oil crisis of 1973, however, sent shudders through it. Because they rely almost exclusively on the import of raw materials and fuels, the soaring cost of oil threatened to halt the recovery. The shipbuilding industry was destroyed, but those out of work were immediately re-trained in new industries. The government made a conscious decision to get out of heavy industry altogether, and manufacturing was henceforth to be concentrated almost exclusively on small, high value, high technology, low material cost products for the export markets.

During the 60's and 70's the Americans led the field in computing and semiconductor technology, but most of the work was carried out in the service of either the military or the space programme. MITI has since instigated the creation of a fully Japanese semiconductor industry: in 1976 they applied considerable pressure to the top five electronics

BACK TO CONTENTS

companies (Fujitsu, NEC, Hitachi, Toshiba, and Mitsubishi Electric) to cooperate in developing high quality and powerful semiconductors for household and ordinary consumer use, and stressed the importance of becoming involved in the 'knowledge intensive' industries. As a direct result of this, the sort of equipment now available on the Japanese domestic market is mind boggling: hand-held photocopiers; the "Animan" robotic pet; a palm-top computer with a touch screen and fax facilities; a TV with a 250cm razor sharp colour LCD screen; and now available in England is a hand-held Geographical Positioning Satellite (GPS) receiver, that will tell you where you are on Earth down to the nearest ten yards!

Korg Wavestation.

The point is this: there is no excuse for the type of non-interactive, inert, unmusical synthesizer designs we are used to. They have the technology.

Example 3: Korg Wavestation. Korg have two R&D centres, one in Tokyo and one in San Jose on America's West Coast. The San Jose centre is peopled with engineers who used to be with Ensoniq and Sequential Circuits. (It is ironic that Sequential Circuits went bust not long after the inception of MIDI.) The Korg Wavestation was partly developed at the San Jose centre under the direction of none other than Dave Smith, and the instrument owes a debt to Sequential Circuits 'VS' synthesizers.

The Wavestation is designed to be purely a synthesizer, in that it has none of the trappings of the ubiquitous 'workstations': no sequencer, no disk drive, no drum patches. What it has instead are two unusual synthesis techniques, called Vector Synthesis (hence 'VS') and Wave Sequencing.

Vector synthesis relies on the dynamic control of oscillator amplitude, carried out by a four-pole isotonic joystick mounted on the top left of the front panel. At its most simple, one each of the four oscillators (in reality samples) are assigned to a pole: moving the joystick modulates the respective amplitudes. It is also possible to program in mix envelopes; these allow the user to adjust the mix percentages at predetermined points using a graphic display. Whilst simple to actually do, the principle is somewhat obscure, and is not helped in this case by an unhelpful description in the otherwise very good manual. As is so often the case, the user is told how to do something, but not actually why they might want to do it, or what it will achieve. Another criticism of the vector synthesis on this machine is that it is not controllable over MIDI, as it is on the Yamaha SY22 for instance. This is a shame, effectively limiting its use to on-stage performance: recordable and editable performances in a studio setting would be equally desirable.

Wave sequencing allows the sort of sonic transformations discussed with Mike McNabbs Dreamsong. The Wavestation lets the user chain together up to 255 separate waveforms and set cross-fade times between them, also allowing pitch-shifting, detuning, duration, loop points etc., with each Patch allowing 4 wave sequences to be layered together: and it can be synced to MIDI. Although sample based, playing higher pitches does not result in them 'shortening', as powerful DSP algorithms manipulate the sample data in real-time to keep their durations even. Incredible stuff. Occasionally small glitches can be heard, but on the whole the 49mHz 20-bit(!) processor manages to keep up.

Another excellent feature of the Wavestation is that it has two completely programmable multi-effects units inside it which can be used either in series or parallel, and it is 16-part multi-timbral. It has an exhaustive MIDI spec and a very good keyboard, which is velocity but not after-touch sensitive. It responds to key pressure, which like the mod wheel, is assignable to almost anything (such as filter cut-off, LFO, effects etc.)

Given such complexity, the Wavestation is a potential user-interface disaster area. Fortunately the designers have done a very good job. The front panel is pretty minimal, and apart from the joystick the other obvious feature is a rotary encoder, which with the ten key pad and the Inc/Dec buttons, allows data entry to be carried out in three different ways. Apart from master volume and the cursor buttons, almost everything else is carried out in software, and access to that is via a large pale-blue back-lit LCD with a set of associated 'soft' keys underneath it. The LCD is 64 x 240 pixels, and offers 8 lines of 40 characters each. The software is superb: it is consistent, well developed, and perfectly logical. There are lots of little shortcuts, parameter macros, graphic displays, and other situation-specific software tools: the Jump/ Mark facility, for instance, allows you to save six user settings. Thanks to the LCD, the pages are large enough to get meaningful chunks of information on. If there is a problem with the interface then it's the cursor buttons. Mounted to the left of the display, they are four separate arrow shaped buttons pointing up, down, left, and right. In operation, the user has to look away from the display every time a cursor movement has to be made to find the right button, which is very annoying and distracting. It is interesting to note that the module version of the Wavestation, released some time later, has replaced these

BACK TO CONTENTS

buttons with a single four-way rocker switch (very similar to the controller on a Nintendo Game Boy!)

OK, so I like it. It's a beautifully designed and built object. In terms of applied technology, the Wavestation is as far ahead of the DX7 as the DX7 is ahead of the Pro-1. It is far too easy to take for granted the almost miraculous level at which modern electronics operates. But it has to be said that as a musical instrument the Wavestation leaves a lot to be desired. In terms of the Truax dichotomy, it is still far too 'general'. The vector synthesis and wave sequencing capabilities are not particularly compatible; one is suited to live performance, one to studio work. In which case, why not have them in separate machines? In the case of the vector synthesis, removing multi-timbrality and wave sequencing would leave a vast computational potential for interactive performance controllers, extra large sample memory, whizzo effects units, or whatever. In other words, design a performance synthesizer. Removing vector synthesis, keeping wave sequencing and multi-timbrality, would at least allow for the inclusion of a decent filter with resonance (which it doesn't have at present). Maybe have some of the filters controls synced to MIDI also, like LFO speed or envelope parameters. Again, the details are unimportant, but design a specialised studio-based system that reflects current practice in that area.

It is as if every new design has to be all things to all men all of the time, rather than designing machines that fulfil a specific and unique function. Also much current instrument design is incredibly conservative, based on ideas that date from the early days of MIDI. Things have changed drastically since then: with computer sequencing packages having become so sophisticated, and computing power relatively affordable, how many people really want a sequencer-equipped synthesizer? And how many people want a synthesizer with which one can edit every last minute detail? There are far too many general, weak, systems being bought and sold: there is a yawning gap for a musician-oriented, quick-edit, powerful, interactive, and control intimate synthesizer.

...and Beyond. Here's one I prepared earlier: this imaginary synthesizer is a monotimbral performance synthesizer. It has 4 Meg of ROM-based samples as its sound sources, with each Patch made up of 4 such sources. 256 on-board Patches are stored in volatile RAM, with the provision for two card slots adding a further 128. The 7 octave keyboard is velocity, after-touch, and pressure sensitive. Pressure and mod wheel will be fully assignable controllers. The usual pitch-bend wheel. Sustain and volume pedals would be supplied as standard. It would be black.

The front panel is centred around a large 12 line 40 character colour LCD. To the right of this is a rotary encoder, and to the left a trackerball. A master volume slider is off to the far left. Beneath the display are five 'semi-soft' keys. A 'compare' button is beyond the top left of the display, an 'exit' button to the bottom right. A larger (red?) 'save' key is below and to the right of the rotary encoder. To the left and right below the trackerball are a set (say 12) of large and clearly labelled buttons (akin to drum-machine pads), each with an associated LED. Only one of these may be active at any one time.

BACK TO CONTENTS

Powering up, the machine would immediately be in Play mode. Sounds would be selected via the rotary encoder. Because there would be no 'bank' system, simultaneously pressing the 'exit' button would mean that the rotary encoder stepped through in 10's, speeding up access. The function of the trackerball would be determined by the large pads: these might include filter, attack time, decay, vector synthesis, effects parameters, etc. Because the trackerball is a four-pole controller, it would enable both filter cut-off and resonance to be manipulated simultaneously. Similarly, something like flanging modulation speed and depth could both be controlled: another possibility would be to have a much more complex filtering system than normal. The use of a trackerball also circumvents one of the usual problems associated with digital control and continuous controllers: when a parameter is called up, the position of the controller does not necessarily match the parameter value. Because the trackerball has no absolute position, this will never occur. It should be noted that this system of 'hit the pad and move the trackerball' is the only means of editing timbre available. Finally, any changes made can be compared and saved, if required, by simply pressing 'compare' or 'save'.

The 'semi-soft' keys beneath the display are labelled Edit, MIDI, Global, FX1, and FX2. Whichever sound is currently active will be affected by these buttons. When entering one of these modes, the LCD will change colour: the rotary encoder will be used for data entry, the trackerball for cursor movement.

1) Edit. This allows the four sample/ oscillators to be chosen for a Patch, plus various other parameters such as LFO shape and assign, oscillator detune etc. When in edit, the same

BACK TO CONTENTS

button now shows the legend 'envelope'. Pressing it enters a new page where the overall Patch envelope is drawn using the trackerball. It would not be possible, therefore, to enter numbers numerically, nor would it be possible to assign a different envelope to each oscillator.

2) MIDI. Self explanatory: basic MIDI receive and transmit, programme change on/ off etc. As a performance synthesizer, it should have a patch chain facility.

3) Global. Pitch bend range, memory protect, velocity response curve, etc.

4) and 5) FX1, FX2. Selecting effect type and parameters. Fully programmable.
Press 'exit' to exit any mode. It can be seen that the only mode that has more than one page is Edit, and this is simply to allow for envelope drawing. This in itself is much better, allowing for more organic, curved, and variant envelopes to be generated.

Hopefully the system strikes a meaningful balance on the generality v strength continuum: the operating system is shallow and easy to access; multi-functionality is kept to a bare minimum; interaction with any parameter can be immediately called up and data can be entered in a real-time/ continuous fashion. Control intimacy would be expected to be high: processing power would be concentrated on achieving smoothness through high resolution. It is constructively limited by its monotimbrality and lack of detailed editing.

A studio-based system derived from this would add multi-timbrality, and remove the keyboard. The module would have

BACK TO CONTENTS

no buttons or display: control would be via a dedicated plug-in laptop unit. Real-time parameter control would still be possible using the trackerball (and recordable over MIDI). The operating system would necessarily be more complex, alleviated by the provision of a much larger (touch?) screen.

Summary. We have seen how MIDI developed, primarily at the instigation of Sequential Circuits supremo Dave Smith and the Japanese synthesizer manufacturers. Some problems with MIDI, most notably timing errors resulting from the serial transmission format, have been discussed. With the acceptance of MIDI as an international standard and the complete digitisation of synthesizers, some explanation has been given as to why the market quickly came to be dominated by the Japanese.

There has been no let up in the rate of technological change. In the few years since the inception of MIDI synthesizer design has continued to change, primarily as a result of sheer processing power. However, it has been suggested here that the underlying assumptions behind current synthesizer design are in need of a radical rethink. MIDI has itself brought about new ways of working, and these changes need to be taken into account at the design stage.

5: Exit

It is almost a truism to talk in terms of the Global Village, multinationals, the world at our back door. However, the fact that it has become a truism should not blind us to the profound effect it has on our everyday lives. For the contemporary musician, it should be of some concern that an absolutely identical instrument to his own is being used on the other side of the planet, and at all points inbetween. It should further concern this musician that the vast majority of the instruments turned out by other manufacturers also sound the same as his. With mass production has come a homogenisation of design, manufacturing, and production techniques, a fast-flowing mainstream cutting a deep swathe through what should be, what promised to be, a riot of diversity.

99.9% of synthesizers now use samples as the basis for their sound generation in a quasi subtractive synthesis environment. There is nothing inherently wrong with this: it is a logical and cost effective strategy. However, what is lacking is a means by which this raw data can be personalised by the user: current machinery simply does not allow this! Partly, this is a matter of a simplified and standardised internal synthesizer architecture. The history of synthesis is littered with discarded and undeveloped ideas, and there is absolutely no reason why some of these can't be recycled using digital technology and re-employed. It is also a matter of personalisation through interaction, developing an individual style, sound, or means of expression to the same level that even a mediocre saxophonist or violinist aspires to.

This level of skill and control, automaticity, can only be achieved through practice, and that takes time. Everything is working against this. Firstly, products, synthesizers, are marketed (and therefore perceived) as having a strictly limited lifestyle: another one will be along in a minute. Musicians are actively discouraged from building up a long-term relationship with a piece of equipment, because if they do that they won't buy next years model. Second, the fact that a working musician is likely to have N pieces of equipment means they are unlikely to develop a special relationship with any one. Simply the time taken learning how to use all this

BACK TO CONTENTS

gear will prevent it: bearing in mind here that mastering an instrument is only a means to an end, not the end itself. All this has to be done and then make some music.

Of course this is not just the manufacturers cynically manipulating the hapless musician. With the seemingly ever-increasing rate of technological change, there is always the sense that the goalposts are being moved: just as we got used to tape, here comes hard disk recording, oops, here comes optical disk recording etc. Whilst this has its disadvantages, it obviously has brought many good things to compensate. With the Wavestation, for example, the wave sequencing technique requires immense processing power, and it's all yours in an immaculate high-quality keyboard form for a modest £1000. A few years ago the internal effects units alone would have cost you more than that.

An efficient and user-friendly interface is also a function of processor power: fast chips allow resources to be diverted for graphics, colour, and clarity. Hopefully the next stage will be the introduction on production synthesizers of extensive real-time control devices and an improved, possibly semi-intelligent, operating system. Certainly the aforementioned definition of MIDI real-time controllers for timbral modification hints that things might be moving in a more interesting direction. However, it might also be that the market will divide up into those who buy sounds on RAM card or CD, and those that edit via computer, with on-board editing becoming largely a thing of the past. Control intimacy may not even get on the agenda. Whatever, there is no doubt that the interface has become An Issue, and I for one will be interested to see how things develop.

Acknowledgements

Thanks to: Dave Marshall at Roland UK; Graham Down and Paul Bundock at Korg UK; Andy Hyde for access to the Pro-1; Mark Greaves for access to the DX7.

References

Akita, M. 'Design and Ergonomics'. In Ergonomics, Vol 34, No 6, pp 815-824, (1991).
Bode, H. 'History of Electronic Sound Modification'. In Journal of the Audio Engineering Society, Vol 32, No 10, pp 730-738, (October 1984).
Card, Moran, & Newell. The Psychology of Human-Computer Interaction . Lawrence Erlbaum Associates (1983).
Chowning, J. 'The Synthesis of Complex Audio Spectra by Means of Frequency Modulation'. In Journal of the Audio Engineering Society, Vol 21, No 7, pp 526-534, (September 1973).
Cole, H. Sounds and Signs. Oxford University Press (1974).
Colbeck, J. 'Analogue Synths in the Digital Age'. In Sound On Sound (August 1991).
Darter, T. (Ed). Synthesizers and Computers. GPI Publications (1985).
Evans, S. Contemporary Japanese Design. Quarto (1991).
Flurscheim, C.H. (Ed). Industrial Design in Engineering. The Design Council (1983).
Gardiner & Christie (Eds). Applying Cognitive Psychology to User-Interface Design. John Wiley & Sons (1987).
Grandjean, E. Fitting the task to the Man. Taylor and Francis (1980).
Griffiths, P. A Guide to Electronic Music. Thames and Hudson (1979).
Harvard, O'Boyle, Eastman, Andre, & Motoyama. 'The relative effectiveness of symbols and words to convey photocopier

functions'. In Applied Ergonomics, Vol 22, No 4, pp 218-224, (1991).
Harvey, J. 'Mortuos Plango, Vivos Voco: A Realisation at IRCAM'. In Computer Music Journal, Vol 5, No 4, pp 22-24, (Winter 1981).
Hughes, R. The Shock of the New. BBC (1981).
Johnson-Laird, P.N. The Computer and the Mind. Fontana (1988).
Kaplan, S.J. 'Developing a Commercial Digital Sound Synthesizer'. In Computer Music Journal, Vol 5, No 3, pp 62-74, (Fall 1981).
Kidd, A.L. Man-Machine Dialogue Design. British Telecom Research Laboratories (1982).
Kurtz, M. Stockhausen: A Biography. Faber and Faber (1992).
Lennard, V. 'Delaying Tactics'. In Music Technology (August 1992).
Lehrman, P. 'The Myths of MIDI'. In Sound On Sound (October 1991).
Lindsay & Norman. Human Information Processing. Academic Press (1977).
Lorenz, C. The Design Dimension. Basil Blackwell (1990).
Lucie-Smith, E. A History of Industrial Design. Phaidon (1983).
Machlis, J. Introduction to Contemporary Music. Dent (1982).
Manning, P. Electronic and Computer Music. Oxford University Press (1985).
Massey, H. The Complete DX7. Amsco Press (1986).
Mathews, M. The Technology of Computer Music. MIT Press (1977).
McNabb, M. 'Dreamsong: The Composition'. In Computer Music Journal, Vol 5, No 4, pp 36-53, (Winter 1981).
Moore, F.R. 'The Dysfunctions of MIDI'. In Computer Music Journal, Vol 12, No 1, pp 19-28, (Spring 1988).

BACK TO CONTENTS

Moore, G.D. 'The Future of Keyboard Technology'. In Sound On Sound (April 1990).
Morita, A. Made In Japan. Fontana (1987).
Norman, M. 'Developments in computing and the user interface - emerging issues in end-user interface design'. In Designing end-user interfaces, eds. Heaton & Sinclair, pp 85-95. Pergamon Infotech (1988).
Ouellette, F. Edgard Varese. Calder & Boyars (1968).
Pierce & Noll. Signals. Scientific American Library (1990).
Pressing, J. 'Cybernetic Issues in Interactive Performance Systems'. In Computer Music Journal, Vol 14, No 1, pp 12-25, (Spring 1990).
Rhea, T. 'Clara Rockmore: The Art of the Theremin'. In Computer Music Journal. No date information.
Roads, C. 'Interview with Max Mathews'. In Computer Music Journal, Vol 4, No 4, pp 15-22, (Winter 1980).
Rubine & McAvinney. 'Programmable Finger-tracking Instrument Controllers. In Computer Music Journal, Vol 14, No 1, pp 26-40, (Spring 1990).
Russ, M. 'The User Interface'. In Sound On Sound (August 1988). 'MIDI: Past, Present, and Future'. In Sound On Sound (June 1992).
St. Hippolyte, M. 'Designing the user interface'. In Accessing C, ed. Strawberry Software, pp 19-34. Von Nostrand Reinhold (1989).
Sanders & McCormick. Human Factors in Engineering and Design. McGraw-Hill (1987).
Shackel, B. 'Introduction to the design of end-user interfaces'. In Designing end-user interfaces, eds. Heaton & Sinclair, pp 97-109. Pergamon Infotech (1988).
Thimbleby, H. 'Can Humans Think?' In Ergonomics, Vol 34, No 10, pp 1269-1287, (1991).
Truax, B. 'The Inverse relation Between Generality and

BACK TO CONTENTS

Strength in Computer Music Programs'. In Interface, Vol 9, pp 49-57, (1980).
van Nes, F.L. 'Visual Ergonomics of Displays'. In The Man-Machine Interface, ed. Roufs. MacMillan (1991).
Wishart, T. 'The Composition of Vox-5'. In Computer Music Journal, Vol 12, No 4, pp 21-27, (Spring 1989).
Wright, J. 'Communicating with the user'. In Designing end-user interfaces, eds. Heaton & Sinclair, pp 121-129. Pergamon Infotech (1988).