

A GUIDE FOR USING THE PAIA 4700/S SYNTHESIZER SYSTEM



There has never been a musical instrument that was conceived overnight and released to the world in an immutable state the next day, each instrument has undergone change and refinement to bring it to its present condition. The same is true of electronic music but the newness of this field coupled with the technology explosion has caused its development to be compressed in time. Whereas the piano has taken centuries to evolve, electronic musical instruments first appeared only a little over four decades ago.

Most of these first instruments were little more than exercises in technology but some were designed to overcome shortcomings in existing instruments. For instance, the piano keyboard is one of the most powerful musical operating systems available but it has one outstanding drawback in that while it provides the musician easy access to all twelve notes of the equally tempered musical scale it prohibits him from using the infinity of musical pitches between those twelve notes. By its very nature it eliminates the possibility of any easy glide from one musical pitch to another (glissando).

One electronic instrument designed to overcome this weakness was the Martinot. The Martinot is similar to modern organs in that a standard keyboard is used to control an electronic oscillator built around a frequency determining capacitance/inductance tank circuit. The inductor is tapped at points that produce frequencies corresponding to the chromatic scale; a rather straightforward, if somewhat simplistic, approach to electronically generating a musical scale. In addition to the keyboard there is a finger ring attached to a slider that controls another oscillator. When properly adjusted this second oscillator produces the pitch corresponding to the keyboard key adjacent to the position of the ring. The combination of keyboard and slider allows the musician to glissando from one note to another or add vibrato with a simple move of the hand without sacrificing the operating ease of the keyboard.

The Ondioline was a contemporary of the Martinot but is significant because it was the first electronic musical instrument to use something other than a sine wave as its basic tone. In the Ondioline a relaxation oscillator controlled by the keyboard produced a sawtooth wave which in turn activated several frequency dividers. The output of the oscillator and frequency dividers were combined using much the same techniques used in some modern organs so that the instrument was capable of generating a great variety of sounds. Observers report that a skilled operator could come close to making the Ondioline talk.

While the Martinot and Ondioline were both designed in France, America's contribution to freeing the musician from the restrictions of the keyboard was probably the most outstanding - not to mention bizarre. A Theremin has no visible means of control at all and is played simply by moving the hand in relation to two metal plates or rods. Inside the instrument are two high frequency oscillators, one shielded from any external influences and the second arranged so that the plate or rod forms part of the frequency determining inductance/capacitance tank circuit. The outputs of these two oscillators are combined in such a way that an audible tone that is the difference between the two frequencies is produced. As the performer's hand is brought closer to the sensing antenna the difference in the two frequencies increases and so does the pitch of the tone. A second circuit allows the performer's other hand to determine the volume of the sound produced. Since there are no frets or keys to

to provide visual or tactile clues to the pitch a Theremin will produce, it is a very difficult instrument to play - but loads of fun.

SYNTHESIZERS

The first equipment that would come close to meeting our current definition of a synthesizer was built by Dr. Harry Olson during the early 1940's.* Produced under the auspices of the RCA Labs, the RCA Mark I and Mark II Synthesizers were something to behold. The Mark I has been disassembled for some time now but the Mark II still exists and is currently being leased to Columbia-Princeton Electronic Music Center, it measures 17 feet long by 7 feet high and is valued at anywhere between \$250,000 and one and a half million dollars depending on who you're talking to.

The average performer might be a little disappointed in the Mark II today because even if there were some way to transport it to a gig, he would find when he got there that he couldn't actually perform a number. The Mark II was simply not capable of real-time operation, each characteristic of the sound the instrument was to produce was laboriously calculated and plotted ahead of time and the result punched into a roll of paper tape. When it came time for the Mark II to do its thing the tape was fed in - like a very large, very expensive player piano - and the results recorded on a multi-track disc, (early 40's, remember, recording tape wasn't so hot in those days.) When all the parts of a number had been recorded on the separate tracks of the disc they were re-recorded on another disc from which a master was made.

You might think that about the only thing that the Mark I and Mark II did that was of any consequence was add the word Synthesizer to our vocabulary but that's not the case at all. They were significant first of all because they were the first to put it all together as far as electronic music production was concerned. All the oscillators, amplifiers and filters needed in one place at one time and best of all some means - no matter how cumbersome - of controlling them all. Secondly, they were the first instruments to utilize white noise sources as part of an electro-musical instrument. White noise will be covered in detail later, for now it should suffice to say that without it sounds like snare drums and cymbals, to mention only two, are impossible.

Don't get the impression that electronic music cannot be produced without a synthesizer, that's not true. Imagine that you are in a laboratory with all sorts of electronic equipment such as oscillators, filters, amplifiers, modulators, tape recorders, etc. You turn on one of the tape recorders and set the oscillator for the pitch you want, twiddle the knobs of the amplifier to shape the loudness contour and play with the filter knobs to adjust timbre. It only takes about six hands and a couple of minutes but when you're through you've got a whole note recorded on the tape. Repeat the process often enough and you've got a whole string of notes. Of course, the tempo is not right and the notes may not be in the right sequence but you can fix that by snipping the tape apart and editing out all the junk before splicing it back together again to produce the desired melody. Now you go back and do the same thing for bass, rhythm and all the other parts. About the only thing you can say for this technique is that it should certainly give you a feeling of accomplishment. Considering the complexity of the process even such monstrosities as dogs barking out the tune of "Away in a Manger" can be forgiven - all that knob twiddling has to do something to a person's mind.

In the early 1960's Dr. R. A. Moog (recognize the name?) began developing and producing a line of electronic music synthesis equipment that revolutionized the field. The feature that made the Moog equipment such a quantum jump in ease of operation sounds almost ridiculously simple, but its implications are so far reaching that it must be stressed; THE KEY PARAMETERS OF THE PROCESSING ELEMENTS ARE A FUNCTION OF THE SUM OF SEVERAL CONTROL VOLTAGES RATHER THAN THE POSITION OF A KNOB.

As an example of the operating ease of voltage control let's see what it does for a relatively simple processing element, an amplifier. As we shall see a little later, one of the things that contributes most to the way an instrument sounds is the manner in which its sound builds up and dies away. When using the classical tape splicing technique these characteristics have to be duplicated manually for every note by turning the volume control of the amplifier. Even though the Mark II allowed for automatic control of the amplifiers, information still had to be punched into its programming tape for each individual note.

With voltage control the job of setting the correct time varying amplifier gain can be turned over to an automatic electronic function generator circuit that produces a repeatable, pre-set voltage waveform each time a key is pressed. This voltage is then used to control the amplifier. The musician sets the function generator to reproduce the characteristics of some real or imagined instrument and the electronics will produce that characteristic for each note he plays. If he desires a totally different sound it's simply a matter of re-setting a couple of knobs. Summing the control voltages allows the performer to produce more than one effect from a single processing module. If, in the above example, the operator decides to add a low frequency amplitude modulation (tremolo) to the sound, he needs only to sum a second voltage that is changing at the rate of the desired tremolo into one of the remaining amplifier control inputs. As the control voltage varies up and down so does the gain of the amplifier and the volume of the sound.

THE SOUND OF MUSIC

Anyone can make weird noises on a synthesizer simply by randomly making connections and pushing buttons. It's even fun for the first hour or so, until you begin to think of specific sounds you want to make and can't. If we're going to learn to use a synthesizer rather than just play with it it's important that we understand what sound is and what makes one sound different from another.

If your knowledge makes the following discussion seem trite, read on anyway. We have to start somewhere and if nothing else you can probably find something to disagree with.

Sound travels as waves, waves of pressure in the air. A vibrating string displaces the air around it and the air molecules that the string moves in turn bump into and move other molecules. All the things that these sound waves can bump into and be reflected off of and the effect that this has on the original wave are beyond the scope of our discussion. The only thing relevant to the subject at hand is that if a man is present the pressure of the waves will finally cause a deflection of his eardrum which in turn will vibrate three small bones inside his ear which will in turn cause a disturbance in a fluid medium which in turn excites the auditory nerves which in turn causes the man to say " Hey, listen to that _____. " Whether he fills in the blank with "noise or "music" is personal preference.

The thing that vibrates to produce the sound doesn't have to be a string. It can be a synthetic or organic membrane as in a drum, a vibrating reed as in the wind instruments or the lips of the musician as in the brass instruments. Most important to us, it can also be the cone of a loudspeaker.

When a recording of a musical instrument is made a microphone converts the air pressure waves into exactly analogous electrical voltage waves. If you were to graph the vibrations of the air and the "vibrations" of the voltage side by side they would be identical except that one would be measured in volts and the other in dynes per square centimeter - or something. When these voltage variations are re-played through an amplifier and loudspeaker they are converted from electrical back into sound energy. If all the links in the chain have been faithful in their recording and reproducing functions the pressure waves generated by the loudspeaker will be exactly the same as those originally generated by the musical instrument and the two will be indistinguishable.

Since the thing that an amplifier and loudspeaker works with is not really sound but an electrical analog of sound; and since it is possible to electronically generate any imaginable voltage waveform (difficult in some cases, but possible), it seems only logical that at some point sounds should be generated not by physical musical instruments but by synthesizing their electronic analog and then converting that to sound.

PITCH, DYNAMICS, TIMBRE

There are really only three characteristics that determine what a musical instrument will sound like; pitch, dynamics and timbre. Of the three, pitch probably requires the least explanation.

Pitch and frequency are two words from two different technologies that describe the same thing. When an engineer or technician speaks of 261 Hz. they mean that the thing they are referring to is vibrating 261 times per second. When a musician mentions middle C he is also talking about something that is vibrating 261 times per second. If the musician is dealing with conventional instruments he is probably talking about a string or reed but if he is working with an organ or synthesizer he is likely referring to the same thing that the technologists were talking about, the frequency of the changes of an electrical waveform.

The human ear is more sensitive to changes in pitch than any other musical parameter. The intensity of a sound has to be cut significantly before a listener experiences any decrease in loudness but a skilled musician can tell when a musical semi-tone deviates by as little as 3% of the interval between that note and the next higher tone.

Dynamics is a broad term that refers to the time varying intensity characteristics of the sound; how fast it builds up and how fast it dies away.

The length of time required for a sound to build up to its greatest intensity is called attack time and this one parameter conveys more information about the way an instrument is played than any other. If the attack time is very short the instrument will be in the percussion family where the vibrating member is immediately excited to its maximum amplitude by the deforming action of being plucked or struck with a hammer or mallet. If the attack is relatively slow then the instrument is probably in the reed or bowed string groups where the action of the exciting force - the wind or bow of the performer - takes a short time to fully excite the vibrating element.

If you forget about the talent factor for a moment the primary purpose of the musician in playing most instruments is to serve as an energy source. The performer pumps energy into the system (instrument) and the system dissipates it in some way, usually as either sound or heat. I know what you're probably thinking, Heat? Yes, heat; if you were able to accurately measure the temperature of a drum head you'd find that it gets hotter as you pound on it. The energy that is converted to heat can be thought of as being lost since it does not contribute to the primary object of producing sound.

Very interesting, right? But what has this to do with the sound of a bassoon. Just this, another important characteristic of an instrument is its release time. That is, how fast the sound dies away. Release time is directly related to how much of the energy goes into heat and how much into sound. A vibrating string, for instance, is as close to lossless as you can get and its release time is very long. The stretched membrane of a drumhead on the other hand is very lossy and as a result the release time of drums is very short.

Reed instruments have a short release time because the reeds are relatively lossy and don't continue to vibrate for very long after the musician stops adding energy. Brass instruments have the shortest release time because the performer can force his lips to stop vibrating and the column of air in the instrument is very lossy.

Sustain time is the interval in between attack and release, the steady state response of the instrument. As is obvious, percussion instruments have zero sustain time - as soon as the attack is finished there is no more energy input so it's downhill the rest of the way. Instruments that have some continuous energy input from the performer, in the form of bowing, blowing or even pedaling in the case of some organs, can sustain as long as the energy holds out.

Though attack, sustain and release are the primary phenomena of dynamics there is one other condition that is common. When a percussion instrument is struck very hard the vibrating member will deform beyond the point at which a smooth release is possible, in effect more energy is put into the system than it can handle, with a resulting overload. Under these conditions the system (string, membrane or whatever) will rapidly get rid of the excess energy. With the "overload" dissipated the vibrating element will continue to dissipate the remaining energy in a normal fashion. The result is an initial rapid attack immediately followed by a decay time which is then followed by a normal release. In a natural instrument it would be all but impossible for the decay time to be followed by a sustain interval but with a synthesizer this is simple.

We can graphically illustrate the conditions discussed by plotting the overall intensity of the sound versus time as shown in figure 1. Since these graphs are drawn to show the peak amplitude of the sound at any given time and therefore "contain" the sound they are often referred to as envelopes.

It is pretty obvious that as important as dynamics is, it doesn't account for all the differences between the sounds of instruments. For instance, the trumpet and french horn are both brass instruments with approximately the same attack, sustain and release characteristics. They even overlap as far as pitch range is concerned but there would be little danger of mistaking the blaring, brassy sound of the trumpet for the mellow, muted tones of the french horn. These differences come about




FIGURE 1. Amplitude envelopes for (A) Percussion (B) reeds (C) attack-decay-release (D) attack-decay-sustain-release

because no musical instrument produces a tone that is composed exclusively of a single frequency. Each note is composed of a number of different frequencies, and the number and amplitude of the various components are what gives each instrument its distinctive timbre.

The concept that a single musical pitch can be made up of more than one frequency can be confusing and needs further attention. The sine wave shown in figure 2 is the basic building block of any imaginable acoustic or electrical wave. It is the only waveform that is composed entirely of a single frequency and, more importantly, any waveform can be built up using nothing but sine waves.




FIGURE 2. Sine wave




FIGURE 3. Fundamental and 2nd. harmonic

To illustrate this look at figure 3. Here we have two sine waves drawn in dotted lines which are labeled "A" and "B". As you can see from the drawing, waveform "B" goes through two cycles in the time that it takes waveform "A" to complete a single cycle. Waveform "B" is therefore twice the frequency of "A" and is said to be the second harmonic of the fundamental frequency "A". If we draw another wave that was three times the frequency of "A" it would be the third harmonic, four times would be the fourth harmonic, five times the fifth and so on.

If at every point in time we sum together the amplitudes of waveforms A and B the result is the waveform shown by the solid line. Note that while the new wave is shaped differently than either A or B it has the same frequency (and consequently pitch) as the fundamental frequency A. If third, fourth, fifth and higher order harmonics were added into this wave the result would continue to change shape but the frequency would remain the same.

It is not necessary that every harmonic of a fundamental frequency be included in a wave and indeed the most musically interesting sounds have certain harmonics deleted. The square wave shown in figure 4 is a good example. It is difficult to




FIGURE 4. Square wave




figure 5.

imagine that the sharp-edged wave illustrated could be built up from smoothly changing sine waves but it can as shown in the progression of diagrams figure 5 (a) through (c). In (a) fundamental frequency is added to its third harmonic producing the waveform shown by the solid line. In (b) the fifth harmonic has been added to the result of (a) to produce the new solid waveform and in (c) the seventh harmonic has been added to all the rest. You can see that the trend as higher order harmonics are added is to steepen the sides of the square and flatten and reduce the ripple in the top. When enough harmonics have been added the result will be a square wave. Notice in particular that not all harmonics are added together for a square wave, only the odd harmonics (3rd, 5th, 7th, etc.) are included.

MAKING WAVES

Now that we have a pretty good idea of why instruments sound the way they do we can look at ways of duplicating these sounds using electronic circuits.

The first method of electronically producing a desired waveform is called additive synthesis and the technique should be obvious from the discussion of harmonic structure. Several oscillators provide a source of sine waves of various harmonically related frequencies and combinations of the outputs are summed together to build up the desired waveform. By changing the amplitudes of the sine waves practically any waveform can be easily produced. One of the problems with this system is keeping all of the oscillators tuned so that they are multiples of one

another. Most electronic organs that use additive synthesis systems get around this problem by using one oscillator for the highest frequency component desired and then producing the other frequencies using a chain of frequency dividers.

The technique used in synthesizers is called subtractive synthesis and can be thought of as just the opposite of additive synthesis. Rather than summing together the frequencies you do want, you start off with a source that is already rich in harmonics and then remove the ones you don't want. This may seem a rather strange way to get from here to there but there is an excellent biological precedent for subtractive synthesis, the most versatile musical instrument of all - the human voice.

There are other reasons for using subtractive synthesis than just pleasing mother nature. If we are going to be consistent in our design of a line of voltage controlled equipment, then everything should be voltage controlled, including the oscillators. Designing a voltage controlled sine wave oscillator is not impossible but it is difficult. Then there is the need to lock all the sine wave oscillators to precise multiples of the fundamental frequency, and a mixing system to combine all the harmonics in the proper amplitude relationships. As you can see, a system of this type would be quite cumbersome and tedious to operate.

From a technological standpoint, it is much easier to electronically generate a complex ramp or square wave than a sine wave. And with the many recent refinements in filter circuits, subtractive synthesis is definitely the better route to take.

Since synthesizers operate with harmonic rich waveforms as their primary signal source there is no need to start out with a sine wave at all. The VCO's supplied with most synthesizers provide a variety of waveforms each of which provides different harmonic structures. Common practice is to use a relaxation oscillator to generate a voltage ramp which is then converted to triangle and pulse waves using simple shaping circuits. In some cases the triangle will also be shaped into a sine wave. These waveforms and their harmonic contents are listed in table 1.

In order to use subtractive synthesis we need some means of getting rid of the harmonics we don't want and to do this we use filters. A filter is quite simply an electronic gadget that eliminates a single frequency or group of frequencies.

wave harmonic	 triangle	 ramp	 square
fundamental	$8/\pi^2$	$2/\pi$	$4/\pi$
2nd	---	$1/\pi$	---
3rd	$8/9\pi^2$	$2/3\pi$	$4/3\pi$
4th	---	$1/2\pi$	---
5th	$8/25\pi^2$	$2/5\pi$	$4/5\pi$
6th	---	$1/3\pi$	---
7th	$8/49\pi^2$	$2/7\pi$	$4/7\pi$
8th	---	$1/4\pi$	---
9th	$8/81\pi^2$	$2/9\pi$	$4/9\pi$

Table 1. Harmonic content of triangle, ramp and square wave
 $\pi=3.142$ $\pi^2=9.872$

Figure 6 shows diagrammatically a representation of the frequency response of a low pass filter. This drawing shows that as the frequency of the signal being fed to the input of the filter increases the amplitude of the filter's output falls off. Note that the filter does not change the frequency of the input signal, only the amplitude. If the input is a complex waveform the filter will of course change the signal's shape as it attenuates the higher frequency components but that is, after all, what we're after.




FIGURE 6. Low-pass filter

Figure 7 shows the frequency response of a high pass filter. In this case the amplitude of the output falls off as the input frequency decreases.




FIGURE 7. High-pass filter

Notice that in both of these filters the response curve is flat either up to or beyond some definite frequency. This is the frequency at which the filter begins to "take hold" and is designated the cutoff frequency or f_c . One other important parameter associated with low pass and high pass filters is the roll off rate, ordinarily measured in units of db/octave. This sounds complicated but it's really not. A decibel (db) is a measure of electrical level and when you're talking about voltage, a change of 6 db, corresponds to halving (if - 6 db.) or doubling if (+ 6 db.) the original reference level. Octaves are of course frequencies that are double some reference frequency; therefore a low pass filter that "rolls off" at 6 db./octave simply means that every time the frequency is doubled the output of the filter falls by 1/2.

Figure 8 shows the frequency response of a band pass filter. As the diagram implies, a band pass filter attenuates all frequencies above and below a certain frequency while allowing the frequency of interest (or frequencies close to it) to pass without being effected. The frequency that is allowed to pass without attenuation is quite logically called the center frequency and is also designated f_c . There are parameters that can be used to specify how well the filter does its job of rejecting frequencies outside of its pass band but none of them are very easy to understand and for our purposes we will confine ourselves to speaking of the "Q" (quality) of the filter. The higher the "Q" the greater the frequencies outside the pass band will be attenuated.




FIGURE 8. Band-pass filter

As the frequency response curve of figure 9 shows, you can think of a notch filter as being the opposite of a band pass filter. Instead of allowing frequencies around the center frequency through, the notch filter blocks these and allows all others to pass.




FIGURE 9. Notch filter

CONTROLLERS

It is about time that we looked at a problem that has plagued instrument makers since the first caveman beat on a hollow log - how to control the instrument in such a way that you realize its full potential. With most conventional instruments the control system is obvious. You control some of the elements of the dynamics by how hard you blow, pick, or strike the instrument and you control the pitch by the positions of your hands and/or lips. Timbre is in most cases a quality of the instrument and is therefore beyond the control of the performer.

This is not the case with a synthesizer; you have at least the theoretical capability of controlling and varying every characteristic of the sound. Some characteristics you can pre-set by the position of a knob and some you can turn over to automatic function generating equipment. Some parameters are varied with a manual controller such as a keyboard and some, unfortunately, you wind up forgetting about because there are no more controllers available.

Before examining some of the types of controllers that are available, make sure that you have firmly implanted in your mind that a controller for a synthesizer does only one thing; it provides a voltage proportional to some parameter that is physically changed by the performer. While in most cases the voltage produced by the controller will subsequently be used to set the pitch of a VCO, this will not always be the case. Depending on the sound being produced, the controller may also be used to set the center frequency of a band-pass filter, roll off rate of a low pass filter or any number of other things.

KEYBOARDS

When used to control a piano, a keyboard is one of the greatest inventions of all time. When used with a synthesizer it is at best a compromise.

Musicians are used to keyboards being connected to polytonic instruments; that is, instruments that are capable of playing as many notes at one time as the number of keys being pressed down. With most basic synthesizers, this is not the case. Basic keyboards are by design a monotonic controller capable of producing one control voltage output at a time. This means the synthesizer is to be played like a saxophone or trumpet - one note at a time. Some intermediate line synthesizers use a clever switching arrangement to produce two notes at a time, but this is still far from a true polytonic system. Development of micro-processors (a computer in an Integrated Circuit) will be the technological step which will finally allow development of fully polytonic keyboards - playing more than one note at a time, and having individual voicing for each additional note being played.

Since the electronic organ has become commonplace, performers have gotten used to the idea that they can't control the dynamics of their instrument by varying the striking force on the keys, this is also true of most synthesizers. Other than triggering signals that are generated when any key is pressed, the only control voltage that most keyboards produce is proportional to the location of the key being activated. One manufacturer has a keyboard that is an exception to this rule; in addition to the standard control voltage it also generates two voltages proportional to the velocity of the key as it is pressed down and the final pressure on the key as it is held down. This is a significant improvement since it allows the performer to directly influence three musical parameters by pressing a single key.

Strangely enough, the original objection to a keyboard that was mentioned in the first part of this booklet (unavailability of pitches between semi-tones) is not a great problem on a synthesizer. Most keyboards provide a "pitch" knob that allows some variation in tuning of the instrument and many provide for an automatic, variable rate glissando.

In spite of its drawbacks, the standard keyboard has one big thing in its favor - familiarity. It is similar to a thing that the musician already knows how to use and re-training time is therefore reduced.

LINEAR CONTROLLERS

These are electrically and mechanically the simplest of all controllers. Most consist of a long strip of electrically resistive material with a voltage applied to each end. The potential difference between the two ends distributes evenly along the length of the strip so that the voltage between any point and electrical ground is proportional to the position of that point on the strip. When the performer presses on the controller, a conducting metal band makes contact with the resistance element and picks off the voltage present at the point of contact.

Linear controllers are generally not intended as substitutes for keyboards for a number of reasons. First, it is technically difficult to automatically produce a trigger pulse whenever the controller is pressed. This function has to be performed manually with a separate switch that must be closed for each note or run that is to be played. Secondly, using a linear controller for pitch is like playing a fretless instrument such as a violin, it requires considerable experience to know what pitch is going to be produced at a given location.

These devices come into their own when used in conjunction with a keyboard. In this application they can provide an auxiliary control for some parameter other than pitch, like manually sweeping a filter or controlling the amount of noise mixed into a sound. The control voltages produced by this unit can also be summed into one of the VCO control inputs to produce a manually controlled glissando or vibrato or can be used with a VCA to give manually controlled tremolo.

FOOT PEDALS

Foot pedals allow you to control additional musical parameters with your feet. They are similar to the expression pedals on electronic organs except that instead of controlling only the volume they can be used to control filters, oscillators or amplifiers.

Anything you can say about linear controllers applied to foot pedals, they're intended to be used in conjunction with a keyboard.

JOY STICKS

These are the wackiest controllers imaginable and as you would expect are similar to the joy sticks used in airplanes.

The biggest thing going for this type of controller is that it offers the possibility of directly controlling four musical parameters simultaneously. One parameter could be controlled by moving the stick forward and backward, another by moving the stick from side to side, a third control voltage could be generated proportional to vertical motions (along the long axis of the stick) and a fourth proportional to

the rotation of the handle. If you like you could even put a switch on top to control such vital functions as self-destruct.

A joy stick seems like a valid concept but anyone that could use one properly probably wouldn't be able to communicate with earth people.

ENVELOPE GENERATORS

Envelope generators are automatic controllers that electronically generate a time varying voltage as pre-set by the positions of knobs or sliders. An envelope generator ordinarily responds to a trigger pulse by generating an electrical waveform that rises to some pre-set value in a pre-set time, sustains that level as long as the trigger pulse is present (or for a pre-set time in some cases) and then falling back to zero in a pre-set time. Some envelope generators are capable of producing the attack, decay, sustain, release type functions discussed earlier.

The output of the envelope generator can be used in the same ways any other control voltage source should, but these items find their most common application in controlling dynamics and time varying timbral qualities of a sound.

A low frequency oscillator can also serve as a control voltage source to provide cyclicly varying voltages for vibrato, tremolo or filter sweeping.

SEQUENCERS

Sequencers fall under the category of control devices much like keyboards. The main point of a sequencer is that it is programmed ahead of time to provide a specific pattern of control voltages. On command, the series of voltages appears at the output, and this pattern can be repeated many times with high consistency. Standard analog sequencers are programmed via potentiometers to provide the desired voltage for each stage. After programming, the module is turned on either manually or from remote trigger signals, and the internal clock steps through the sequence to "read" the programmed voltages. Digital technology has aided in the development of digital sequencers which differ in one main area. Rather than manually programming the voltage for each stage, an input is provided to accept control voltages from a keyboard or other control device. As the keyboard is played, the various control voltage steps are translated into a digital word and stored in a memory system. On command, the sequencer reads its output information from the memory rather than from a string of potentiometers. These devices are still costly, but as computer technology increases the cost will continue to drop.

Although sequencers are usually thought of as being used to control an oscillator (repeating bass lines, high speed arpeggios etc.) many interesting effects can be obtained by using them as controllers for filters, VCA's or any other modules requiring a control voltage.

DIGITAL COMPUTERS

As short as two years ago, thoughts of digital computers being interfaced with an electronic music synthesizer were reserved for well funded college music departments or government subsidized artistic research programs. Since then, computers have grown smaller, more inexpensive, and basically more accessible. Electronic hobbyists are delving into these machines as they did with Ham radio ten years ago. Realizing a computerized music system is

practical today. Initially, we will see computers being used to control the analog sound generating and processing modules of the synthesizer. Then, as prices decrease, fully computerized synthesizers will be seen which will generate any waveform, changing waveform, amplitude, and other parameters as you play - depending on how the computer is programmed. Eventually, we may even see computer music systems which can "listen" to a sound and reproduce the same sound under operator command. Or how about a machine that monitors the operators brain waves and vital signs, and composes music specifically to fit his moods/thoughts. The list is endless.

THE EQUALLY TEMPERED SCALE

This is as good a place as any to bring up the subject of the equally tempered or chromatic scale. As anyone who is reading this knows there are 12 semi-tones in each octave of the chromatic scale, 7 naturals labeled A through G and 5 accidentals that are designated as either sharps or flats of the naturals. With two exceptions the sharp of one note is identical to the flat of the next highest note, the exception being that there are no accidentals between B and C, or E and F so that B# is the same as C and F^b is the same as E.

For each octave increase in the musical scale the frequency of the note doubles so that since middle C corresponds to 261.6 cycles per second the next C above middle C is 523.2 cycles per second.

Somewhere back in antiquity (around the time of J. S. Bach) some genius decided that since there are 12 semi-tones to the octave and each octave doubles the frequency, each note should be related to the note directly below it in the scale by a factor of the twelfth root of two. Just in case you're not used to working out the twelfth root of numbers in your head, this translates to 1.059 times the frequency of the note directly below it. The significance of this is that as pitch increases, the difference between adjacent notes in the scale also increases.


All this may seem like academic trivia until you realize one point. All voltage controlled oscillator designs produce a device whose output frequency is directly proportional to the control voltage, identical control voltage changes produce identical frequency changes.

An example will most readily demonstrate the significance of these facts. Suppose that we have a keyboard that produces a control voltage of .625 volts when its lowest C is pressed. The voltage corresponding to the next C up is quite logically 1.25 volts but don't fall into the trap of thinking that the voltage corresponding to the third C is 1.25 plus .625 or 1.875 volts because it's not, it should be twice the voltage required for the second C or $2 \times 1.25 = 2.5$ volts.

Many synthesizer designers use an electronic conversion device to get around this difficulty. This device converts a linear controller output voltage (1v. for the first C, 2v. for the second, 3v. for the third C, etc.) to the octavely related voltage required by the VCO.

Unfortunately, the exponential converter circuits (as these devices are known) are not only expensive but also quite often they tend to drift so that even for a fixed input voltage the output voltage (and of course pitch of the VCO) wanders from one value to another.


A simple means of getting around this is to have the keyboard generate octavely related voltages in the first place. This eliminates the need for a separate exponential converter on each oscillator and filter. And when you're talking about 3 or 4 VCOs and a couple of filters, this is a big savings in cost and drift problems.



4710 BALANCED MODULATOR

Conventional synthesizer components provide the user a multitude a sources of sound waveforms that are rich in harmonic-related frequencies, but only the Balanced Modulator provides the basis for generating sounds with component frequencies spaced according to non-harmonic intervals. While natural sounds with these properties are less than numerous (chimes, gongs, certain man-made whistles, insects and others), the sound of the suppressed carrier waveform can be one of the most distinctively "electronic".

The inputs and controls of the PAIA 4710 Balanced Modulator have been arranged to provide for this module's convenient use as a Voltage Controlled Amplifier as well as a sine wave frequency doubler.




SPECIFICATIONS

Power Requirements:	+18v. @ 10 ma. +9v. @ 1.5 ma. -9v. @ 7 ma.
Input Impedance:	Carrier - 5K nominal Modulation - 50K nominal Control - 680K
Output Impedance:	3K nominal
Frequency Response:	15 Hz. to 30 kHz.
Carrier/Mod. Rejection:	50 db. typical

4711 MIXER

If there is one module that the developing synthesist never seems to have enough of it is mixers. Increasing expertise at the instrument produces increasing sophistication and subtlety in the sounds produced. In most cases the more subtle and "natural" the sound the more individual component sounds are required to produce it.

The PAIA 4711 Mixer allows for the mixing of up to 4 channels of program material with individual level controls for each channel. Pan pots allow the 4 input channels to be split between the two output channels in any proportion from 0 to 100%. Master gain controls on the output channels allow for raising and lowering the total output levels while maintaining constant percentage of program material from the 4 inputs.




SPECIFICATIONS

Power Requirements:	+9v. @ 4 ma. -9v. @ 4 ma.
Input Impedance:	10K ohms
Output Impedance:	less than 100 ohms
Freq. Response: (3 db, 10K load)	5 Hz. to 20 kHz.
Max. Gain:	6 db.
Max. Input Level:	10v. p-p
Cross Talk:	-60 db.

4712 REVERB MODULE

A reverberation module is a natural addition to any synthesizer system since it provides an electro-mechanical equivalent to the delays and echos that define the acoustics of a large auditorium. The 4712 Reverb Module optimizes control flexibility by allowing for the mixing of two separate input signals as well as post-mixing of the input signals with the processed reverb signal. Front panel switch allows for instantaneous, pop-free termination of effects without affecting normal signal flow through the module.




SPECIFICATIONS

Power Requirements:	+9v. @ 3.5 ma. -9v. @ 3.5 ma.
Inputs:	2; mixing
Input Impedance:	100K min.
Output Impedance:	less than 1K
Reverberation Time:	3 sec. nom.

4720 VOLTAGE CONTROLLED OSCILLATOR

The 4720 Voltage Controlled Oscillator features simultaneous ramp, triangle, sine and pulse/square wave outputs; manually set or voltage controlled pulse width and 16 Hz. to 16 kHz. frequency response corresponding to 5 volt peak control voltage. Front panel initial pitch control allows multiple oscillators to be offset and track one another chromatically over the entire operating range. A self-zeroing front end and active, on-board voltage regulation combine to make the 4720 extremely linear and stable.




SPECIFICATIONS

Power Requirements:	9 volts @ 20 ma.
Frequency Range:	16 Hz. to 16 kHz.
Output Level:	.5v. peak to peak
Output Waveforms:	ramp, triangle, sine, square/variable pulse
Control Voltage/Frequency Response:	Linear
Frequency Control Input Impedance:	270K
Pulse Width Modulation Input Impedance:	68K
Pulse Width Range:	0 to 100% duty cycle

4730 MULTI-MODAL FILTER

A good filter is the cornerstone of a music synthesis system. The PAIA 4730 has features that are ordinarily reserved for devices costing 5 to 20 times more. Including simultaneous low, high, and band-pass outputs and extreme linearity allowing the filter to track the keyboard for numerous interesting effects. Yet, like all PAIA modules, the calibration procedure for this device has been refined so that the only equipment required is a Volt-Ohm meter.




SPECIFICATIONS

Power Requirements:	+9v. @ 15 ma. -9v. @ 25 ma.
Outputs:	Low Pass, Band-Pass, High Pass Simultaneously available.
C. V./Freq. Response:	Linear
Ranges:	
Freq.:	16Hz. to 16 kHz.
Q:	.5 to 150
Input Impedance:	200K minimum
Control Input Impedance:	270K

4740 ENVELOPE GENERATOR

The 4740 ADSR Envelope Generator provides the synthesist with a triggerable source of programmable time varying voltage that is useful in controlling filters, oscillators, amplifiers or other voltage controllable elements. Front panel controls allow independent, non-interactive adjustment of the output voltage's rise time (attack), initial decay, sustain level and final release time. Electrical and manual triggers are provided for and front panel light emitting diodes indicate status of module.



SPECIFICATIONS


Power requirements:
Trigger input impedance:
Output impedance:
Peak Output Voltage:
Attack, Decay, Release
Time Range:
Triggering voltage:
Sustain Level Range:
Front Panel Indicators

+18v. @ 17 ma.
66,000 ohms nominal
less than 1K ohms
5 volts nominal

2 ms. to 1.5 sec. nom.
3.5 volts nom.
.5 to 5 volts
attack and decay/ sustain
Light Emitting Diodes


2720-5 CONTROL OSCILLATOR/ NOISE SOURCE

The 2720-5 Control Oscillator/ Noise Source is a dual function module providing a source of slowly varying sinusoidal control voltage for tremolo, vibrato, filter sweeping, etc. and a white noise source useful in a variety of special effects including wind and surf sounds.



SPECIFICATIONS

Module Power Requirements:	18v. @ 1.75 ma.
Oscillator:	
Output:	Sinusoidal
Output Impedance:	Nominal 5K or less
Output Amplitude:	One fixed 5v. p-p output One variable 0 to 5v. p-p
Output D. C. Bias:	50% of peak to peak value
Frequency Range:	1 to 25 Hz. in three overlapping ranges
Noise Source:	
Output:	Broadband noise nominal 100 Hz. to 10 kHz.
Output Impedance:	Nominal 1K
Output Amplitude:	.5v. p-p
Output D. C. Bias:	None




4761 ROAD CASE

The PAIA 4761 wing cabinet is a rugged module case, capable of withstanding the abuse encountered when equipment is taken on the road. When used in a studio setup, the 4761 offers the versatility of rearranging the system in a "building block" fashion to suit the individual synthesist's needs. The 4761 front panel provides ground connections between cabinets and increased patching flexibility.

SPECIFICATIONS

Case Dimensions:	15-1/8" X 7" X 9-1/4"
Module Space:	13 single width modules



4782 KEYBOARD / CASE

The 4782 is a self contained keyboard controller housed in a heavy duty roadworthy case. The three octaves of key selectable control voltage is supplemented by variable tuning, sample and hold, and variable glide circuitry. The self contained power supply eliminates the need for multiple power jumpers running to other cabinets. Glide circuitry allows sweeping tones such as trombone slides or steel guitar.

SPECIFICATIONS



Dimensions:	31" X 13" X 4"
Power Requirement:	110 volts AC
Internal Power Supply Output:	±9 volts @ 25 ma, +18 volts @ 5 ma.
Keyboard Style:	3 octaves of standard AGO configuration from C to C
Pitch control Range:	1 octave minimum
Slowest Glide Rate:	2.5 seconds across entire keyboard
Control voltage output:	.3 volts to 5 volts minimum
Trigger Voltage Output:	5 volt pulse on key depression 5 volt step during key depression
Tuning:	Each key tunable
Hold Time:	20 seconds minimum after key release
Output Impedance:	100 ohms nominal

4770 WATT BLOCK

The 4770 is the basis of the power supply system for the 4700/S. Three of these modules provide the necessary DC voltages for proper operation of the modules. As the synthesizer system is expanded, the power supply modules are easily relocated, expanded or rewired as needed to suit the new system. A power switch is provided on the front panel as well as an LED indicator lamp to signify operation. Also mounted on the front panel are two sets of three jack multiples - one for control voltages, one for audio signals.

SPECIFICATIONS

Power Requirements: 117 v. 50/60 Hz.
Power output: +9v. @ 100 ma.
-9v. @ 100 ma.
or
+18v. @ 100 ma.



4780 SEQUENCER

In basic terms, a sequencer is simply another control voltage and trigger source. Unlike other controllers, however; it allows the user to pre-program an extended sequence of control voltage variations that can be called up on command and allows the variations to be executed at speeds beyond human capabilities. In addition to their obvious applications in generating fast arpeggios and automatic bass lines, sequencers can also be used as programmable function generators, counters or clock sources.

Considerable time and effort has gone into making the 4780 as versatile as possible without losing the user in a hopeless tangle of operational details. Many of the internal housekeeping functions are automatic yet externally available inputs and controls permit by-passing automatic functions in favor of manual control.

SPECIFICATIONS

Power RequirementsL +18v. @ 50 ma.
±9v. @ 1 ma.
Clock Rate: 2/sec. to 32/sec.
Control voltage range: 0 - 5 v. nominal
Additional Features: Built in glide, Light Emitting Diode indicates clock rate, LED stage status indicators, Manual and electrically activated controls, 2 to 12 event re-circulating count capabilities.

THE 4700/S AS A SYSTEM

The 4700/S is an assortment of the previously discussed modules which has capabilities competitive with commercial systems selling for several times the price. In fact, there are very few, if any, synthesizers in the "under \$2000" bracket that include a sequencer as a standard function. Also, many low cost filters provide only low pass output, or a choice of low pass or band pass. Rarely do you see three simultaneous filter outputs available. The module complement of the 4700/S is such that the sequencer can control a full synthesizer by itself, and you still have a full array of modules for use with the keyboard.

Another concept of much debate is the modular patchable system versus the pre-patched switch operated system. There is no argument that the normalized synthesizers fill a very useful role - specifically the musician who wants to make ultra-fast changes on stage while he is performing, or someone who wants to add some new sounds to their repertoire without the need to fully understand what they are doing. But for those of us who want to really learn about the nature of sound, the patchable system can't be beat. Being able to utilize any module in any configuration draws every drop of performance out of the circuitry. After mastering the operation of the modules, most any sound that you can imagine can be duplicated on a patchable system. Further, the modular concept implies an obsolescence proof system, for as your ideas grow you can expand your system with a few more modules rather than selling the old system and replacing it with a bigger better model. Most people find that after using a patch system, it doesn't take that much longer to insert a few cords than it does to flip several switches. Thus, more and more patchable systems are being seen on the road with musicians.

The following section contains several examples of patches for the 4700/S system. Starting with these sounds will help explain the uses and capabilities of the modules. After initially setting up the patch, experiment with altered settings on the controls and listen to the changes that occur. This type of experimentation is very helpful in becoming more familiar with the system and how it works.

The symbology used in the following diagrams to represent the various modules is fully explained in the PAIA manual "A Schematic Symbology For Synthesizer Patching Arrangements" which is provided with this manual. We have chosen to represent synthesizer patches in this manner because many people physically arrange their modular systems in different configurations. The symbology system is universal and can be adapted to any synthesizer.

Best of luck in your experiments with the 4700/S; welcome to a whole new world of sound.

EFFECTS USING WHITE NOISE AS A SIGNAL SOURCE

The randomness and wide frequency content of white noise makes an ideal basis for many sounds both musical and special effect. The most obvious is simulating the wind or the sound of the surf. The noise output of the 2720-5 Control Oscillator Noise Source is patched to the 4730 filter input. The band-pass output is the output of the patch. As a control source for the filter you can use a randomly tuned sequencer with a lot of glide, an envelope generator with long attack and release times, the keyboard with long glide time, or the control oscillator output. Each of these controllers gives a different effect, but the common drawback is that they aren't completely random like Mother Nature's wind and surf. To increase randomness, feed more than one of these controllers into the three control voltage inputs of the filter. The random sequencer pattern can be occurring continuously while the keyboard output is user controlled to make the wind sweep higher or lower depending on the location of the key pressed. Additionally, an envelope generator can be triggered by the keyboard to cause an instantaneous gust of wind settling back to the level determined by which key was pressed.

Increasing the Q of the filter causes a narrower band of frequencies to be passed through the filter. With Q at maximum and the keyboard as the only control source, you can generate a melodic wind and actually play melodies.

Using an envelope generator to control a filter with low Q setting and range switch at LOW, some very effective thunder and explosion effects can be added. Also, try the low pass output for these type of effects. The envelope generators will most approximate an explosion with a minimum attack and long release.

Using the noise source through a VCA which is controlled by an envelope generator will produce hi-hat cymbal effects. Set the envelope attack at minimum and the release at about 20%. Patch the keyboard pulse trigger to the input of the envelope generator. Every time you hit a key the cymbal sound will be produced.




Figure 1 - WIND/SURF




Figure 2 - CYMBAL

To add more expression to the sound keep your hand on the release control of the envelope generator and vary the setting for different beats. This gives the effect of the hi-hat cymbal opening and closing.

A keyboard controlled snare drum can be conceived which allows the user to play a melody on the drum. The keyboard control voltage output sets the pitch of the VCO to duplicate the "strike tone" normally provided by the tuned drum head. The triangle or sine VCO output is mixed with white noise at the mixer and processed through a VCA. The envelope generator for the VCA should have minimum attack and decay, maximum sustain and medium release. Triggering the envelope generator with the pulse keyboard trigger insures that the drum sound will have no sustain. As an option, the noise source can be filtered prior to being mixed with the strike tone. Experiment with assorted combinations of filter output and control sources such as keyboard, envelope generator, or fixed bias (as outlined in the string patch section).




Figure 3 - MELODIC SNARE DRUM




Figure 4 - FLUTE

COMMENTS:

- 4782: Pitch - Set to upper octave.
Glide - Off.
- 4720: Initial Freq. - Around 30 to 50%.
- 4740: Attack - 20%.
Decay - Minimum.
Sustain - Maximum.
Release - 20%.

This is a basic patch using the soft sine wave of a VCO to duplicate the sound of a flute. The medium attack and release on the envelope duplicates the effect of the build-up and collapse of the air column in the flute. The envelope generator, being triggered by the step trigger from the keyboard, will generate a sustained signal as long as a key is depressed. The balanced modulator is being used as a VCA for this patch.




Figure 5 - FUNK BASS

COMMENTS:

- Keyboard: Set to low octave.
- 4720's: Both VCO's use ramp output.
VCO's tuned in unison at lowest setting.
- 4711: Mix two inputs equally.
- 4740: Attack - Minimum.
Decay - 20%.
Sustain - 60%.
Release - 30%
- 4730: Lowpass Output.
"Q" Control - 50% to 100%.
Initial Freq. - Below 50% if range switch is "high".
Maximum if range switch is "low".

This patch gives the funky "wow" bass sound made popular by Stevie Wonder, the Bee Gees and many others. The 4740 variable output can be changed manually while you are playing for added expression and variation.




Figure 6 - CHIMES (AND OTHER METALLIC SOUNDS)

COMMENTS:

- Keyboard: High octave, no glide.
- 4740: Attack - Minimum.
Decay - Minimum.
Sustain - Maximum.
Release - 50 to 100%.
- 4720's: Tune triangle VCO to an augmented fourth above the pulse VCO.
Pulse VCO pulse width - 30%.

This is a general purpose patch for bells. The settings shown will provide a sound of tubular chimes or church bells. Experiment with various permutations of this patch. How about a chime with sustain? To do this, use the keyboard step output rather than the pulse. Also play with the 4740 to get chimes with slow attack, etc. The biggest variables in this patch are the VCO's.




Figure 7 - STRINGS

COMMENTS:

- Keyboard: Pitch control - high range.
No glide.
- VCO's: Tuned in unison, pulse output.
Initial Freq. - Approximately 30%.
Initial Pulse Width - 10 to 20%.
- LFO (Low Frequency Oscillator):
Approximately 10 Hz.
Very small output amplitude - 10%.
- Mixer: All signals mixed equally.
- Filter: Low pass output; Range- maximum; Q Control - 50%.
- ADSR: Attack - 20%.
Decay - 30%.
Sustain - 60%.
Release - 50%.
- Bias Supply (See Text) - Approximately 3 volts.

The sound of violins is a pleasing addition to any type of music. Strings are basically an easy sound to synthesize, but as you can see this patch uses some modules which aren't available in the 4700/S --- or are they? This patch was initially developed on a joint 4700/2720 PAIA system, but with some special patching, all functions can be obtained on the 4700/S. Hopefully, this will show the importance of experimenting with the modules to find applications which aren't immediately apparent from reading the front panel or instruction manuals. Performance above and beyond the call of duty, you might say.

The basic string sound comes from the narrow pulses generated by the VCO. Narrow pulses have a very widely distributed harmonic series and give a hollow, nasal effect. Three VCO's are used to get the effect of a whole string section playing at once. All oscillators are tracking in unison, and are mixed in equal proportions. The low pass filter takes the sharp edge off the sound to duplicate the natural harmonic roll-off of a violin around 8 kHz. The settings given for the envelope generator are intermediate settings. Longer attacks can be used for "background" type strings, or shorter attacks for pizzicato strings. The release setting can be varied to simulate the reverberation of the concert hall - longer release times simulate more reverberant rooms, shorter times for smaller or more absorbant rooms. Addition of the 4712 reverb unit (included in the 4700/S) or external echo units, phasers, or flangers at the output of this patch will increase the apparent size of the group of strings and add more depth to the string sound.

Now we get to the special patching effects. In a real string section, each violinist will generate his own different vibrato. This generates a complex random system of beat frequencies within the sound. The 4700/S has one control oscillator which can be patched directly to one of the 4720 VCOs. To generate the other vibrato signal we will use the 4780 sequencer. Only the first two stages of the sequencer will be used, so patch the second stage output to the load input. Switch the clock to conditional run, and press the load button. The first two lights should be blinking alternately. Set the clock rate about 60%, and turn the tuning controls for the first two stages completely to minimum. Patch the sequencer output to a control input on another VCO. Advance the tuning control on the first stage until you begin to hear an up/down trilling effect. Fine tune this setting to get the amount of vibrato depth you will want. Now, advance the glide control in the output section until the vibrato smooths out to an approximate sine wave. Fine tune the vibrato rate (clock rate), vibrato depth (stage one tuning) and glide rate. If you are more ambitious, more stages of the sequencer could be used to provide a pseudo-random vibrato signal.

The fixed bias supply can be achieved by patching the full level output of the unused envelope generator to its own trigger input. Set all controls to minimum except sustain at maximum, and press the manual trigger button. The 4740 should latch in the sustain mode, and the proper LED should remain lighted. The variable output now becomes a source of 0 to 5 volt regulated bias voltage. This voltage is patched to the control input of the 4730 filter and adjusted for the "stringiest" sound. Some may prefer to have the Q control higher than suggested, but as always - experiment will all controls to adjust the sound to your own taste.