

ELECTRONOTES 76

Newsletter of the Musical Engineering Group
203 Snyder Hill Road
Ithaca, N. Y. 14850

VOLUME 9, NUMBER 76

APRIL 1977

GROUP ANNOUNCEMENTS:CONTENTS OF EN#76

Page 2	News and Notes
Page 2	Reader's Questions (Continued, Page 20)
Page 3	The ENS-76 Home-Built Synthesizer System - Part 8, Random Sources -by Bernie Hutchins
Page 3	ENS-76 Random Source - Option 1
Page 5	ENS-76 Random Source - Option 2
Page 10	Further Options
Page 12	Design Example
Page 16	Reader's Equipment
Page 17	Forum - About Our Readers (Results of Survey)
Page 19	Record Review: "Sequencer" by Larry Fast reviewed by Craig Anderton

The ENS-76 presentation in this issue gives the random sources (noise sources) for the system. The first option is conventional for the most part. The final option on the other hand, given as a "Design Example" is an entirely new device, and might have been given as a report rather than an ENS-76 presentation. Reader's interested in Pseudo-Random and partially random processes should study this carefully as there are many results which are related to partially random sequences, and not just to noise sources in the conventional sense.

NEW MEMBERS AND CHANGES (c)

Harold W. Anderson, Jr.	10 Cemetery St., Hopedale, MA 01747
Peter Baltzer	1212 Sir Francis Drake Blvd., San Anselmo, CA 94960
Bill Bauson	2008 So. Goyer Rd., Apt 1, Kokomo, IN 46901
Kenneth L. Caron	1271 Water St., Fitchburg, MA 01420
Robert D. Conkwright	629 W. 33rd St., Baltimore, MD 21211
Michale Convery	18 McGraw Ct., East Greenwich, RI 02818
Michale Coxie	24 Dixie Trail, Raleigh, NC 27607
J. Rodney DeWitt	2210 Wilbur, Kalamazoo, MI 49007
Richard Griffin	4342 Lawrence Ave. East, West Hill, ONT, Canada M1E 2T1
Jeff Grothaus	1265 Iliff Ave., Cincinnati, OH 45205
Ed Hansen Jr	PO Box 2064, Akron, OH 44309 (c)
Jim Kelly	78 Endicott St., Salem, MA 01970
Jim Laird	284 Ridge Rd. E., Rochester, NY 14621
Sergio Menchinelli	Greenhouse Cl, 00040 S. Maria D. Mole (ROMA), Italy

Norm Milliard
 Danny H. Mills
 Stephen M. Moore
 Jeffrey L. Noble
 Richard Olshausen
 John Orpin
 Dino P. Parato
 R. S. Penley
 Håkan Persson
 Phillip Pope
 John M. Prickett
 Fred B. Rayworth
 Richard Robinson
 Mark Sankey
 Peter A. Scheff
 Terry Shultz
 Mark Stockfisch
 R. E. Sullivan
 Ken Swanson
 Emile Tobenfeld
 August Vernieuwe
 David M. Vosh
 Peter Walden
 Dean Wallraff
 Gary Wedge

Star Instruments, PO Box 71, Stafford Springs, CT 06076 (c)
 2100 Howell Branch Rd Apt 48-G, Maitland, FL 32751
 24 Carr Dr., Moraga, CA 94556
 PO Box 326, Rosamond, CA 93560
 16842 Sea Witch Lane, Huntington Beach, CA 92649
 6 Conrad Ave., Toronto, ONT Canada M6G 3G5
 121 Rhode Island Ave., Bay Shore, NY 11706
 604 Quance Ave., Saskatoon, SA Canada S7H 3B4 (c)
 Box 104, 37041 Ramdala, Sweden
 PO Box 5991, Rome, GA 30161 (c)
 3705 Grady St., Ft. Worth, TX 76119
 RFD#1 Box 12, Roanoke, WV 26423
 3065 Brook Dr., Decatur, GA 30033 (c)
 Hyde Park Estates #41, Hyde Park, NY 12538 (c)
 920 W. Oakdale, Chicago, IL 60657 (c)
 13522 E. Penn St., Whittier, CA 90602
 2121 Birch Lane, Park Ridge, IL 60068
 1131 Hill Ave., Fullerton, CA 92633
 20516 Hamlin St., Canoga Park, CA 91306
 314 West 100 St. #85, New York, NY 10025
 Pathoekeweg 40, B-8000 Brugge, Belgium
 8024 Carey Branch Place, Oxon Hill, MD 20022
 19 Tallwood Ave., Eastwood 2122, NSW Australia
 119 Pearl St Apt 2, Cambridge, MA 02139
 134 Milton Ave., Summerside, PEI, Canada

NEWS AND NOTES:

The 1977 International Computer Music Conference will be held at the Univ. of California at San Diego from Oct. 26 - 30, 1977. Papers to be presented should be submitted for evaluation by June 30, 1977. For more information, write to The Second International Computer Music Conference, Center for Music Experiment Q-037, U. of Cal. at San Diego, La Jolla, CA 92093.

A catalog of computer music representing over 100 composers, Computer Music Compositions of the United States 1976, is available at \$5 from Theodore Front, 155 North San Vicente Blvd., Beverly Hills, CA 90211

A new Guitar/Synthesizer interface unit, "The Slavedriver" is being produced by 360 Systems, 2825 Hyans St., Los Angeles, CA 90026, at a cost of \$795. Also available is the "Bass Slavedriver" at the same price. Contact 360 Systems for literature.

The Dept. of Journalism and Communication at the Univ. of Bridgeport and Dondisound Studios, Inc. of Red Hook, NY will be holding the second annual Institute of Audio Studies at Dondisound from June 6 to August 5, 1977. Interested persons should contact David Moulton, Dondisound, Red Hook, NY 12571 or Howard Jacobson, Dept of Journalism, North Hall 214, Univ. of Bridgeport, Bridgeport, CT

Dennis Electronics, 2130 Metcalf St., Honolulu, HI 96822 is changing the front panel designs of their synthesizer modules, and until June 1st will be selling older modules as "prototype" units at reduced price. Write for a list of modules and prices.

READER'S QUESTIONS:

Q: I am using monolithic matched pairs of transistors and the type Q81 temperature compensating resistors in my VCO's, but the oscillators still drift. Why is this, and what can I do about it?

THE ENS-76 HOME-BUILT SYNTHESIZER SYSTEM - PART 8, RANDOM SOURCES:

-by Bernie Hutchins, ELECTRONOTES

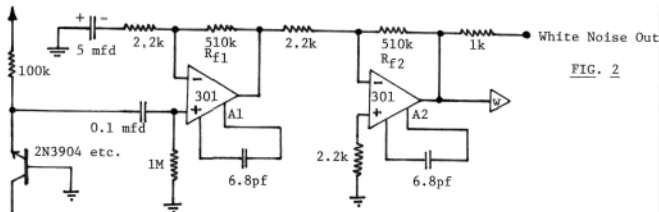
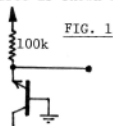
Most synthesizers have one or more "Noise Sources" which provide outputs for "White Noise" and often for "Pink Noise" as well. White noise has a flat spectral density, while pink noise has a spectral density that rolls off at 3db/octave, thus providing equal spectral energy in all bandwidth of the same ratio. Both types have a general character of "static noise" and are very similar to the interstation noise that is heard on an FM receiver. These noise sources can be used to provide basic textural material for sound synthesis, can be used for the synthesis of noise-like sounds (e.g., thunder), or can be converted to control signals with devices like filters and sample-and-hold units.

There are two common approaches to the generation of random signals. First, you can use true random noise generated in an electrical device as a result of some physical process that is random. A back-biased transistor junction is a good source of such noise if it is properly amplified. The second approach is the "Pseudo Random" (PR) approach which employs digital feedback shift registers to generate a long series of digital states in a sequence that "seems" random to non-specialized electrical instruments and to the subjective human ear.

In our approach, we shall offer several random units. A simple white noise source based on the back biased transistor will be presented for simplicity and economy. We shall also offer some PR devices with additional features. Finally, we want to look at devices for generating sounds that have both random (or PR) aspects as well as periodic aspects to them. These latter devices are new and unique devices which we have not discussed before.

ENS-76 RANDOM SOURCE - OPTION 1

The simplest form of the back-biased transistor junction noise source is shown in Fig. 1. Note that it consists of only a single transistor and a current limiting resistor. The noise voltage generated at the junction is not large enough for synthesizer use (although it has a very wide bandwidth), so for an actual source, we need quite a bit of amplification so that we get something like a five volt level. It is difficult with these sources to set a level since the signal is random, but the user should have little trouble getting a satisfactory signal. The circuit for the White Noise Source is shown in Fig. 2



The circuit of Fig. 2 is virtually unchanged from previous designs. A1 and A2 form a two stage amplifier. The gain of these amplifiers is adjusted so that the noise level at the output is sufficient. This depends on the individual transistor

selected for T1, so you may have to do some adjusting of the feedback resistors R_{f1} and R_{f2} , upward for more gain, and downward for less gain. If desired, trim pots can be used to make this easier. We have selected type 301 amplifiers here (or substitute type 748) because these are easily compensated for higher gain giving a larger bandwidth. Type 741, 307, should not be used here. Type 556 is acceptable, but the 301 or 748 will probably have a "crisper" sound.

It is typical to do some processing of the white noise output to make other types of random signals available. These are driven off amplifier A2 (from the little triangle with the "w" inside). Below we show three processors. Fig. 3 shows a pink noise filter and a low-frequency random signal processing filter. There are from a design submitted by Terry Mikulic and first presented in EN#34. A graph of the frequency response of the pink noise filter can be found in EN#64 (14), and the report on noise generators given in that issue may be useful reading at this point. Note that the random signal is driven off the pink noise output. It can also be driven directly off the white noise, although some gain adjustments may be needed in this case. Be a little careful adjusting the level of the L.F. Random Signal as it has a relatively small bandwidth (below 10 Hz) and should be observed for 10 seconds or more before you reach any conclusions about the signal levels you might expect. If you observe for too short a period, you may set the gain too high. The third processing unit is the 7 Hz bandpass filter shown in Fig. 4. This is a high-Q bandpass filter which is excited by the white noise. Its output depends on the exact setting of the Q control, but tends toward a 7 Hz sinusoidal as Q is increased. We put this in to provide a somewhat random vibrato signal. With the Q turned to the max position, the circuit oscillates providing a steady state vibrato signal. The design of the filter is discussed in EN#73 (14).

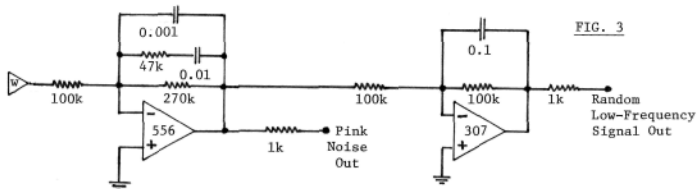
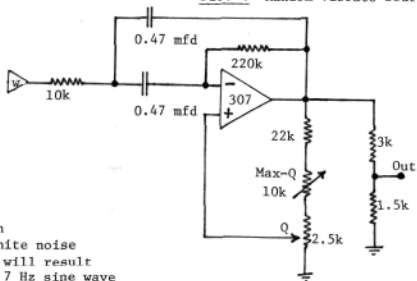


FIG. 4 Random Vibrato Source

The bandpass circuit is adjusted by first grounding the input. Then the Q control is set all the way up, and the Max-Q trimmer is then set so that the circuit oscillates with only a small amount of clipping. It is possible to adjust the Max-Q control so that a reasonably pure sine wave output is achieved, but the level will not be stable. Thus, a small amount of clipping should be used so that oscillation is steady. Now connect up the white noise source. Decreasing the Q control will result in a signal that is essentially a 7 Hz sine wave with randomly varying amplitude, and as the Q is



further decreased, the output becomes more and more random in frequency as well as in amplitude. Also, the amplitude will drop as the Q is decreased, but this can usually be compensated for in normal use by changing the input gain of the circuit being modulated by the vibrato signal.

ENS-76 RANDOM SOURCE - OPTION 2

In Option 2, we will be replacing the white noise circuit of Fig. 2 with a Pseudo-Random (PR) shift register. This is more complex and more expensive than the simple back-biased transistor. So why do it? There are several reasons for doing this (or for considering doing this - as the case may be). The most important reason is that the resulting outputs will be repeatable, although seemingly random. If something of particular musical use does occur in the output, we will at least have the potential of recovering this and doing it again. This is done by resetting the generator or letting it go through a full cycle. The reader unfamiliar with PR generators should consult a text on feedback shift registers, and/or read the discussion of noise generators in EN#64. A second reason for considering a PR generator is that it is much easier to set up and maintain (from unit to unit) signal levels in the generators since the noise output does not depend on the individual diodes (back-biased transistors) used. A final reason for choosing the PR generator is that a lot of additional units can be driven from the PR generator. We will be considering these units in final options. They will include units for different statistical distributions of amplitudes, recycle units for random selection of timbre and for randomly changing timbres, and units for periodic changes of random timbre.

The basic white noise generator of Option 2 is shown in Fig. 5. The circuit uses a 24 stage shift register with Modulo-2 feedback according to the setup required for a maximal length sequence. [See Chapter 5h of MEH]. The total length of such a sequence is $2^n - 1$ where n is the number of stages - hence this 24 stage shift register produces a PR sequence of length 16,777,215. At the 36 kHz rate of clocking shown, the sequence will cycle after a period of nearly 8 minutes. This is more than long enough for most purposes. We selected this length because it was exactly the length you get using three type 74164 chips, and two chips (shift register length 16) is not long enough. The start up circuitry shown is needed to prevent the shift register from starting up in the "All Zero" state, from which it can never emerge without special resetting. The output of the Exclusive-OR gates, or any other stage of the shift register, is the PR Binary Sequence that we will be using. The output sequence could in fact be used directly as a white noise source. However, while the spectrum is right for white noise (i.e., flat over the audio range), the statistical distribution of amplitude levels consists of only two possible values, unlike the back-biased junction source that is approximately Gaussian (see EN#62). It is well known however that low-pass filtering of the PR Binary Sequence will result in an amplitude distribution that has continuous levels, and approaches Gaussian. The 1k and 0.01 mfd R-C filter gives a cutoff of about 16 kHz, which is enough in this case, although not really enough to assure the best approach to a Gaussian distribution. If the best Gaussian distribution is desired, the 0.01 mfd capacitor in the TTL Clock should be changed to 0.002 mfd, which speeds the sequence, so that filtering is effectively lower. Here we really have in mind that we want to go over to an external clock input anyway (a VCO) so our major concern is setting the filtering at a 16 kHz low-pass cutoff.

The circuit in Fig. 5 is a good white noise source, and can be used to drive the processing units of Fig. 2 and Fig. 3, replacing the noise source of Fig. 2. We are however using it mainly as an example, because if you are going to go to the trouble of building this more complex circuit, you will certainly want to add on some of the extra features we are about to describe.

The first feature we will consider is the external clock input. There are numerous circuits for interfacing an external clock to TTL. Here we choose the simple 555 timer shown in Fig. 6. The timer in this case is configured as a simple Schmitt

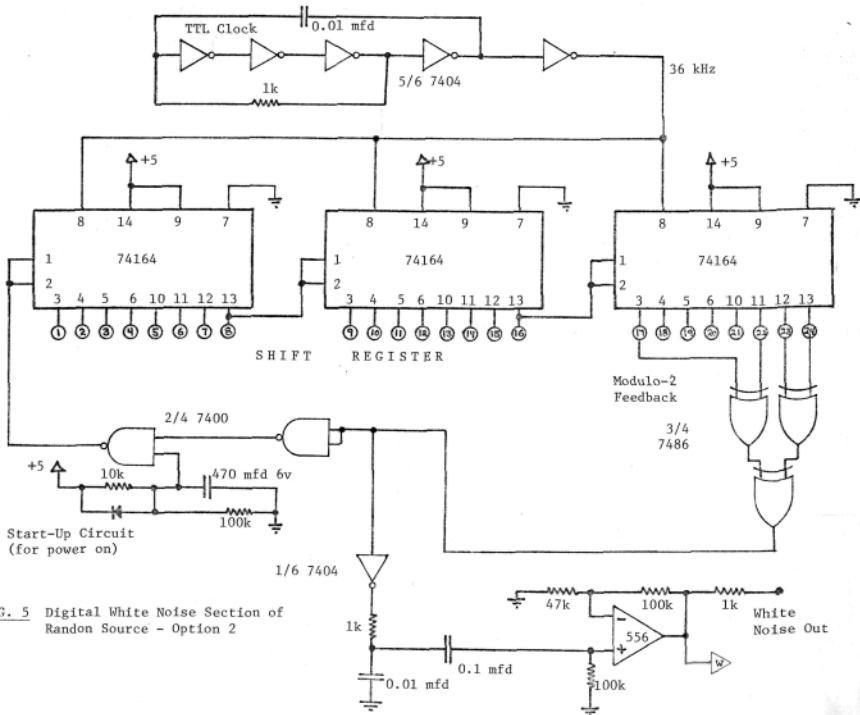


FIG. 5 Digital White Noise Section of
Randon Source - Option 2

trigger (See Application Note #31). This circuit requires that the input voltage pass through the region from 1.66 volts to 3.33 volts, which is the case for all our standard synthesizer waveforms. The Schmitt trigger action of the 555 means that we can use any waveform that is convenient from the external source - it need not be rectangular.

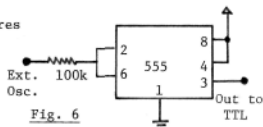


Fig. 6

The use of the external input with a quality musical VCO will become important when we consider the use of PR sequences as timbral elements. For the moment however, we note that with a wide range VCO attached, we can slow the sequence down into the sub-audio range, and the output can be used as a sequencer. Of course, if we want we can do the same general thing with a sample-and-hold on the full audio bandwidth output, but several interesting things happen when we try to use the lower clocking rates. First note that if we still wanted to use the sample-and-hold, we would have to use a much lower cutoff frequency on the low-pass filter, which is not really practical. Instead, it is practical to use the discrete time properties of the shift register itself rather than some external sample-and-hold. Thus we will consider groups of stages of the shift register to be digital words, and will convert these words to voltages using some sort of D/A converter. There are several ways to do this.

First, suppose we attach a standard 8-bit D/A converter to the first 8 stages of the shift register. We are interested in the different voltage levels that come out with each clock cycle. In particular, we inquire about the statistical distribution of levels and the expected changes in voltage between consecutive clock cycles. To do this properly, we have to know quite a bit about PR Sequences, but it will suit our purposes here to just know that on the average we can expect all possible combinations of 8 bits to appear in our 8-bit window with approximately equal probability. This means that all possible 8-bit numbers are equally likely, and the amplitude distribution is uniform. This is indicated in Fig. 7a. Thus, a binary weighted (standard) D/A attached to the shift register results in a uniform distribution. Next, we ask what would happen if we used equal weighting of all eight bits? The result is about the same as we would get if we shook up a box of eight pennies and counted all the heads equally. We expect that very often we get 4 or 5 heads, but cases where we get 8 heads, or only 1 head, are rare. The distribution is binomial, not uniform, and is a good approximation to the Gaussian distribution. Thus uniform weighting results in a Gaussian-like distribution (Fig. 7b). Finally we want to look at the variation of the output during consecutive clock cycles. One clock pulse shifts the bits of the shift register one place to the right. Thus in the case of uniform weighting, only the first and last bits will have any effect on the sum during any one clock pulse, and the maximum change is one bit. Thus, large changes are inhibited, and this means low-pass filtering of the sequence. If you are familiar with discrete time filters, you see that we have formed a low-pass transversal filter. Thus, the D/A connection implies some bandwidth reduction that must be allowed for, but this is of little importance for control sequences. Note also that with the binary weighted D/A there is also some low-pass filtering, although the exact function is different. Note further that if the taps for the D/A are not consecutive shift register positions, the degree of low-pass filtering is reduced, and larger changes are allowed (the maximum changes is allowed if there is at least one space between all taps). Note finally that if you have not understood a thing about this paragraph, you can find more information in EN#62, or you can read on for another page or so, and then come back, at which time your understanding may be increased by the examples we will consider.

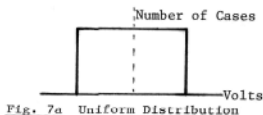


Fig. 7a Uniform Distribution

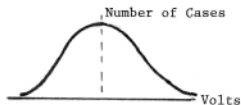


Fig. 7b Gaussian Distribution

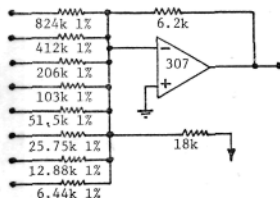


Fig. 8a Standard Binary Weighting
Uniform Distribution

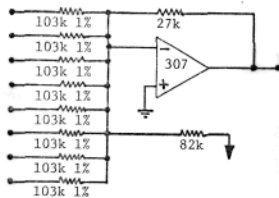


Fig. 8b Uniform Weighting
Gaussian Distribution

Two circuits for D/A converters for the digital noise source are shown in Fig. 8a and Fig. 8b. These circuits should be interfaced with the 74164 using some sort of CMOS chips (inverters are fine - the inversion makes little difference with noise sources). The CMOS gives better defined digital levels. I used 74C00 chips for my interface, and the 103k resistors used just happened to be some I had on hand. You can scale the whole network for whatever you have on hand.

We have to consider here just what we expect to do with the random source to see what is required of the D/A's. In general, of course, we can not make 8 bit D/A's with 1% resistors (since 8 bits gives 256 levels). However, if the output of the source is to be used as audible noise, it is difficult to argue that errors caused by the 1% resistors will matter, since the effect of the errors is mainly additive white noise. Similarly, we would not insist that the op-amps in Fig. 2 be of the low noise type. What happens if we use the random source as a sequencer, driving a VCO for example. Ideally, the levels are equally spaced, and if we want 12-tone equally tempered scales, this gives about 120 levels over the audible 10 octave range. This is slightly less than 7 bits, and probably something like 6 bits or fewer is more like what we need for most musical purposes. We might therefore suppose that for sequencer operation, 6 bits could be used (see later design example), and that either the 1% resistors are close enough, or we could select, trim, or buy resistors of greater precision if necessary. If desired, the D/A converter used with the ENS-76 digital keyboard can be used in place of the one in Fig. 8a. This circuit is shown in EN#68 (16).

Next we want to consider how the taps on the shift register should be arranged, and the differences between the different outputs of the two different D/A's. We will look at this with regard to the effect these things have on the nature of the sequences that result, and we are generally thinking of the sequence as being converted to a sequence of pitches by a VCO, since this is a standard application. We can arrive at four cases for combinations of the two D/A's and close or wide spacing of the shift register taps. These four cases are diagrammed in Fig. 9. You can see that there are considerable differences between the four cases. These differences can be described in terms of the amplitude distribution and the degree of low-pass filtering, but we find it useful to combine these into one term which we will call "Expectation" which is a somewhat subjective feeling about what will come next in the sequence based on what has come before. It is related to our ability to find trends even in the absence of clear patterns. Narrow distributions like the Gaussian add Expectation because we do not expect extreme values. Low-pass filtering adds Expectation because we do not expect large changes. It is of some interest that the standard method of sampling white noise results in a moderate case of expectation, and is thus not as subjectively random as some users assume.

FIG. 9

BINARY WEIGHTING (Fig. 8a) UNIFORM DIST.

UNIFORM WEIGHTING (Fig. 8b) GAUSSIAN DIST.

CONSECUTIVE
TAPS ON THE
SHIFT
REGISTER

1. DISTRIBUTION: Uniform
2. LOW-PASS FILTERING: Some
3. EXPECTATION: Consecutive samples tend to cluster together, except when one of the more significant bits changes, at which time large jumps are possible.



MODERATE EXPECTATION

1. DISTRIBUTION: Gaussian
2. LOW-PASS FILTERING: Substantial
3. EXPECTATION: Large jumps are ruled out due to transversal low-pass filtering, and values far from zero are rare due to Gaussian distribution.
4. NOTES: Similar to Sample-and-Hold with white noise in and slew-limited S & H.



HIGHEST EXPECTATION

WIDELY SPACED
TAPS ON THE
SHIFT
REGISTER

1. DISTRIBUTION: Uniform
2. LOW-PASS FILTERING: Little or None
3. EXPECTATION: None



LOWEST (NO) EXPECTATION

1. DISTRIBUTION: Gaussian
2. LOW-PASS FILTERING: Little or None
3. EXPECTATION: Large or small jumps are possible, but extreme values far from zero are rare due to Gaussian distribution.
4. NOTES: Very similar to process of using a standard Sample-and-Hold on a standard white noise source.



MODERATE EXPECTATION

There are no exact rules on how the taps on the line should be placed, even when we know which of the cases in Fig. 9 we are trying to get. As a start, the taps listed below can be tried. Probably you would want to have several different sets of taps available, so a 8-pole, 4-position switch, or the equivalent digital logic, could be used.

Close Spacing: 1, 2, 3, 4, 5, 6, 7, 8

Spaced Taps - 1: 1, 2, 3, 4, 6, 9, 13, 18

Spaced Taps - 2: 1, 6, 10, 13, 15, 16, 17, 18

Wide Spacing: 1, 3, 6, 10, 12, 15, 19, 21

Two different sets of "Spaced Taps" are given. These are combinations of consecutive taps and widely spaced taps. These give an in between case of Expectation. Two sets are given because it makes a difference with the binary weighted D/A.

FURTHER OPTIONS

We want to discuss here some recycle devices that can be added on to the digital shift register for some new devices. First however, we should recall that when a PR sequence is of relatively short length, or even when it gets up to a length of a thousand or so, it still may be heard as a pitch when cycled rapidly. The exact reason for this is not completely understood, but probably the ear is not remembering the whole pattern and recognizing the repetition. Rather, the ear is probably recognizing the periodic cycling of some of the major "landmarks" in such a sequence. In any event, it is of interest that for short sequences we hear a sound that is not a great deal different than what we hear from a sawtooth or a pulse waveform. At the other extreme, we hear white noise for very long sequences. Clearly we are curious about what happens for medium length sequences. The area between random sounds and perfectly ordered sounds is little explored. There are many fascinating sounds to be heard for sequences from shift registers of between 10 and 20 stages. There are many more sounds available for non-maximal sequences. One can spend hours playing with a fully programmable sequence generator (see Mid-Month Letter 75A).

However, one thing can be concluded from a study of medium length sequences: We do not really get the in between sounds we were thinking we might get. Rather they sound like mixtures of random and ordered sounds - that is, noise and periodic clicks pops and dings. Interesting, possibly useful, but what else can we do?

Since short sequences result in interesting timbre generators (See R. Burhans, "Pseudo-Noise Timbre Generators," J. Aud Eng Soc. Vol. 20, No. 3, April 1972, pg. 3 and R. Burhans, "Harmonic Structure of PN Sequences," JAES Vol. 20) we can consider that short "captured" sub-sequences from longer sequences can also be used. This is easily implemented by adding an extra length of shift register to the output of the PR generator, and adding provisions so the contents of this added length can be loaded and then recirculated. The basic scheme should be evident from Fig. 10.

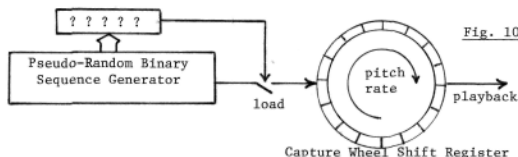
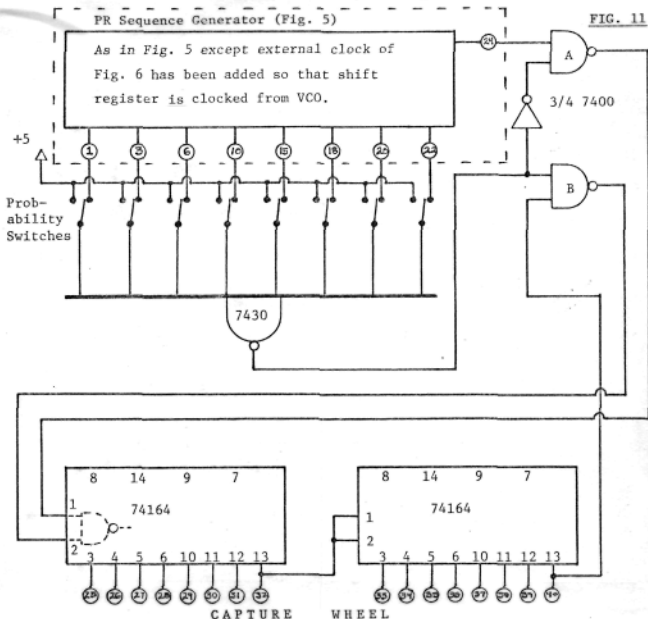


Fig. 10 shows a 16 stage shift register arranged as a "capture wheel" which can be loaded from the PR sequence generator. If we close the load switch, white noise is loaded on the wheel and read off at the other side. If we then open the load switch, the 16 states are fixed and one pitch cycle occurs on each rotation of the wheel. Thus we can easily obtain timbres somewhat at random, subject of course to the limitations of rectangular waveforms in general. We can also if we choose, periodically change one state of the shift register, or part of it by opening and closing the switch at regular intervals. However, by far the most interesting process is the one implied by the box with the question marks in it in Fig. 10. This permits the states on the wheel to be changed or not according to certain random (PR) conditions on the main PR sequence generator.

A TTL realization of Fig. 10 is shown in Fig. 11, where the PR Sequence generator is the one shown in Fig. 5. In normal operation, assume that the output of the 7430 8-In NAND gate is high so that NAND gate B is passing the sequence from stage 40 back through the input of the first of the 74164's in Fig. 11, thus recycling the sequence into stage 25 again. If the output of the 7430 actually goes low, then the recycle process is ended and new bits from the PR Sequence are input to stage 25. When will the output of the 7430 go low? When all eight inputs are high. This depends on the settings of the eight probability switches. Assume that two of the switches connect



the inputs to the +5 supply, while the other six switches are connected to the stages of the shift register. Now, any given stage of the PR Sequence generator may be high or low with equal probability. Thus, the probability that all six inputs are high is $1/2^6$. If this happens (on the average, once in $2^6 = 64$ clock pulses), then the output of the 7430 gate goes low, and the contents of stage 24 are fed onto the capture wheel (to stage 25). The chance that this new state is different from the state currently on the capture wheel at that time is $1/2$, so the overall chance of a change of the particular bit on the capture wheel is $(1/2^6) \cdot (1/2) = 1/2^7 = 1/128$. Thus we could expect a change in one bit on the capture wheel for one clock cycle out of every 128, which for the 16 stage capture wheel is a change of one bit for every eight pitch cycles, on the average. Of course, this is only one example, and the probabilities can be changed by opening or closing more of the probability switches, and by adding more such switches if desired. Also, note that for the 16 stage capture wheel, the change of any one bit is a change of one part in 16, which is generally audible as a change of spectrum, but not always a major change. We can if we wish use fewer stages in the recycle wheel, or can add more shift registers. We suggest that as a starter, the 16 stage wheel works well, and a switch for changing from 8 to 16 stages might be useful.

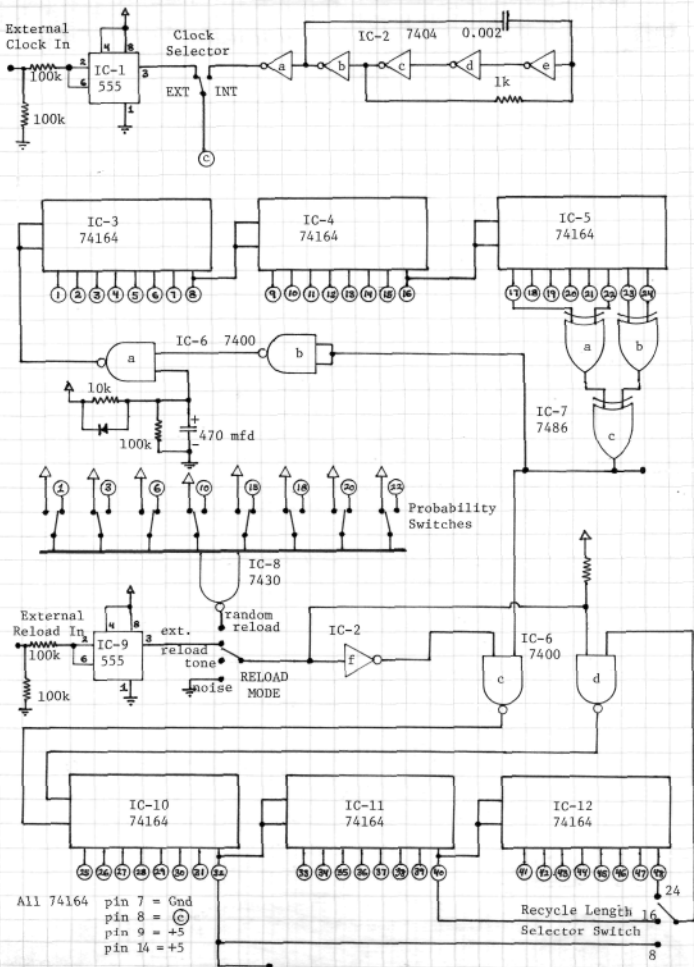
This process works quite well in general. It is felt that the sounds produced by this method do give sounds that are more like those in between random and perfectly periodic rather than mixtures of noise and periodic. What are the sounds like? Well of course it depends on the settings, since for certain probability settings the changes of timbre occur only once a second or so, while for others changes of timbre are quite rapid, and for the extreme of all switches connected to +5, we get white noise out. But in general, the sounds have a very broad spectrum, like a sharp pulse or a sawtooth waveform. Ralph Burhans (second reference mentioned above) showed that the PR sequences have a spectrum that includes all harmonics except those that are integer multiples of the sequence length, so the spectrum includes nearly all harmonics for practical purposes. Thus, the changes that occur when only one bit in 16 is changed may in general be expected to be subtle, although some are dramatic. Of course, the sequence can be further processed with filtering, and the VCO may be modulated as well for additional enrichment of the final spectrum.

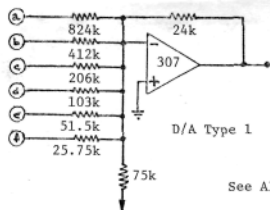
DESIGN EXAMPLE

Since there are so many possibilities for a random source based on the principles given above, it is necessary to limit some of the features so that the resulting unit is not too complicated to be of use. The exact way this is limited is up to the individual designer, so we do not want to pin things down too much at this point. For this reason, instead of calling the following design Option 3, we prefer to call it just a design example.

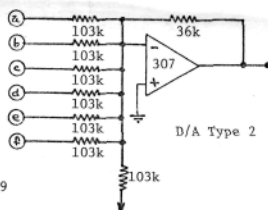
The digital part of the design example is shown in Fig. 12 while the analog part is shown in Fig. 13. The circuit features either an internal clock (IC-2) or an interface for an external clock (IC-1). IC-3, IC-4, IC-5, and IC-7 form a 24 stage PR shift register with IC-6a,b forming the start up circuitry. IC-10, IC-11, and IC-12 form the recirculating or "capture wheel" part of the design. The length of the wheel is adjustable by the "Recycle Length Selector Switch" as 8, 16, or 24 stages. Note that the capture wheel is reloaded whenever a low state exists at the input of IC-2f, since this condition blocks the recycle through IC-6d, and feeds in PR bits from IC-7c through IC-6c. There are four positions on the "Reload Mode" switch. In the lowest position, the input of IC-2f is always low, and the PR sequence is passed onto the capture wheel at all times, and no periodic tone is produced. The second position (Tone) leaves the input of IC-2f always high, and whatever happens to be on the capture wheel at the time it goes high will be recirculated, thus forming a tone. Thus, the user can obtain somewhat random timbres by first loading noise (lower position) and then capturing and holding it (Tone position). There are also two automatic

FIG. 12 DIGITAL PART OF DESIGN EXAMPLE



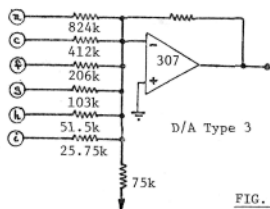


D/A Type 1

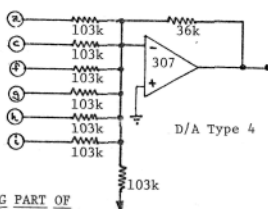


D/A Type 2

See Also Fig. 9

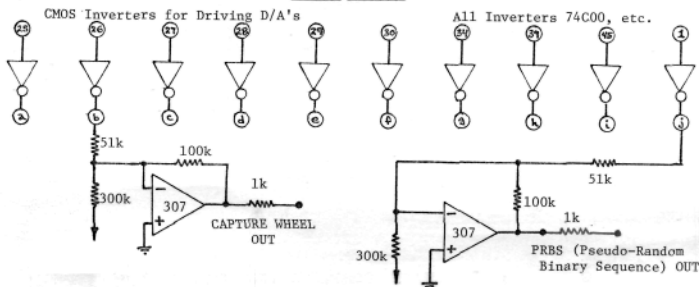


D/A Type 3



D/A Type 4

FIG. 13 ANALOG PART OF
DESIGN EXAMPLE



reload controls. The external reload will be activated and start reloading noise on the capture wheel whenever the input goes above 3.33 volts. An external clock with variable pulse width is particularly useful here, since that will control both the rate at which reload is initiated, and the number of samples that are reloaded. Also available is the PR reload feature through IC-8 and the probability setting switches.

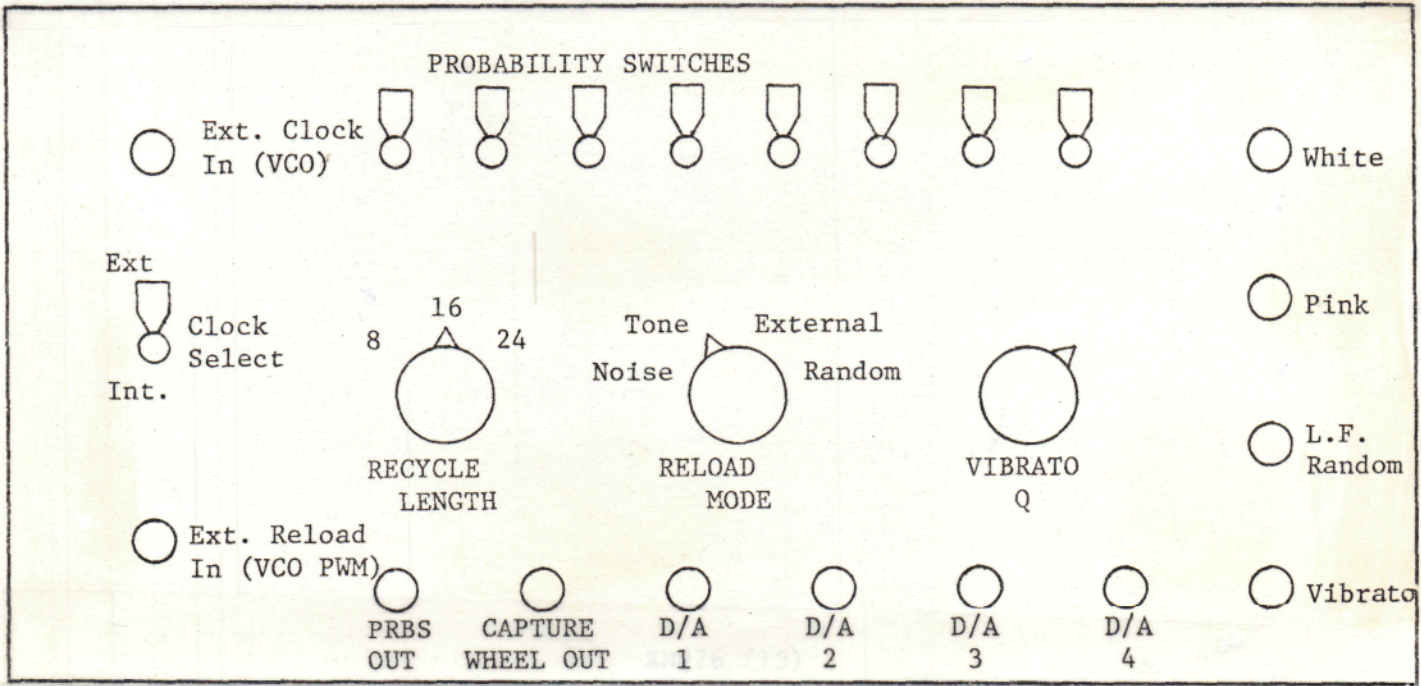
Before going any further, we want to try to head off any possible further confusion concerning the two shift registers. The first (IC-3, IC-4, and IC-5) is connected so as to always produce noise. It runs continually, and cycles after more than 16 million clock pulses. The reload function of the circuit in no way changes what is going on in this shift register. In particular, it does not change the length of the PR Sequence. We could obtain pitched sounds by shortening the PR Sequence and clocking it rapidly, but this is not what is done here. The second shift register (IC-10, IC-11, and IC-12)

while constructed in the same manner as the first, is really quite different in function. Its cycle length is never more than 24 clock pulses. If it is not in the recycle mode, it is simply a 24 stage extension for readout purposes only as far as the first shift register is concerned. It does not lengthen the PR Sequence because it follows the last Exclusive-OR gate, and there is no feedback from it to the first shift register. Once again, consulting Fig. 10 may be of help here.

The analog portion of the design example is shown in Fig. 13. Here we have made provisions to bring out different sequences, and have also added four simple D/A converters. Two of the D/A converters are uniformly weighted and two are binary weighted. We use two of each because we want to use both close and wide spacing of taps on each one, and it is really simpler to use separate D/A's rather than to work out a fancy switching arrangement. Also, we have four simultaneous outputs available. The D/A converters were cut back to 6 bits in this case to keep things simple. We have in mind that these will be used for random voltage sequences, but you can also take noise out of these if you wish. Note that the D/A's have been connected to the capture wheel shift register rather than to the PR Sequence shift register. This has been done so that captured sequences can also be read out through the D/A. By putting the reload mode switch in the noise position, we are effectively using the PR Sequence, only it is slightly delayed. Note that for wide spacing, some of the taps extend beyond IC-10 into IC-11 and IC-12. There is no real problem with this even when the recycle length switch is set to eight. In this case, IC-11 and IC-12 have the same contents as IC-10 but the spacing of the taps is such that no duplicated output is used for an input to a D/A. However, the fact that the sequence in IC-10 is duplicated in IC-11 and IC-12 has the effect of reducing the spacing of taps. This may be useful however since it has the effect of changing the spacing of the taps without actually changing them.

Things will perhaps be clearer if we look at a possible panel diagram for the random source design example. This is shown in Fig. 14. With clock select in the internal position, and the Reload Mode switch in noise, all outputs are some form or audio noise, except for the Low Frequency Random and Vibrato outputs, which are similar to the outputs in Option 1. Let's now assume that an external musical quality VCO is input to the external clock input, and the external clock is selected. This would cause the bandwidth of the audio noise signals to rise and fall with VCO frequency. Now we listen to the Capture Wheel Out or any one of the D/A outputs, and switch the Reload Mode switch to tone. The output is a tone of high harmonic content. If we switch the Reload Mode

FIG. 14 POSSIBLE PANEL LAYOUT FOR DESIGN EXAMPLE



to noise and then back to tone again, we get a slightly different tone, and so on. Next we connect a second VCO to the External Reload input. We use a pulse-width modulated signal in here. We switch the Reload Mode switch to External, and adjust the second VCO to about $\frac{1}{2}$ Hz and the pulse duty cycle to about 1/2. We get out repeating bursts of noise 1 second long, and each of these noise bursts is followed by a tone of slightly different timbre. If the pulse width is changed so that it is nearly all high, we get mostly noise. If it is changed so that it is nearly all low, we get tones which change timbre every few seconds. This is because the external load signal is changing one bit on the capture wheel thus changing the timbre. If we switch the Reload Mode switch to Random, we get white noise out if all the probability switches are up (connected to +5). As we start setting the probability switches down one at a time, bits of tone start to appear and linger longer and longer. With all switches down, we have a tone of a constant pitch but one which is changing timber somewhat at random.

Next we can consider using the Random Source as a sequencer. For this, we can do all the same sort of things, but will be using the external clock at a very slow rate, probably 1 - 2 Hz, and will attach a VCO to one of the D/A outputs. With the Reload Mode switch in Noise, we get an apparently random tone sequence. Trying different D/A outputs will result in different degrees of Expectation. Now, here are a couple of points to consider if you want to use the Random Source as a sequencer. First, it might be well to replace the Recycle Length three position switch with a multi-position switch with more terminals so that sequences of lengths other than 8, 16, and 24 can be selected. Secondly, you may want to add another position to the Reload Mode switch. This position would be fed by a counter working off the system clock. This counter could reload a bit every 13th clock pulse for example with a recycle length of 10, resulting in a sequential evolution at regular intervals. This would supplement the Random position where the sequence would evolve at random times. In either case, the exact notes of the evolving sequence are apparently random.

This is about all we want to say about this sort of device at this time. They are very simple to build, are not overly difficult to learn to use, give some very useful new capabilities, and work well with modular systems. Note in particular that the device is useful as both a sound generator and a sequencer, adding to its utility.

* * * * *

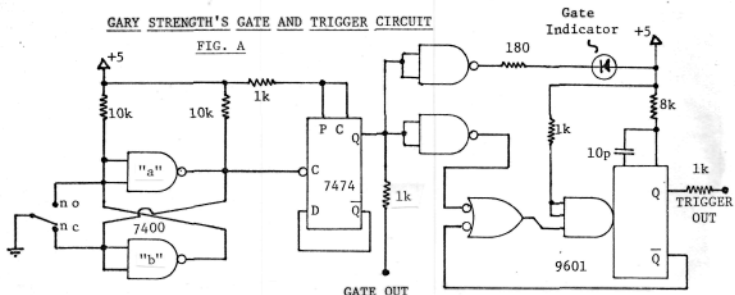
READER'S EQUIPMENT:

The circuit shown in Fig. A was submitted by Gary Strength and is a push-on, push-off type of Gate and Trigger unit. We can suppose that the switch shown is a pushbutton. When the switch is pushed the first time, the R-S Flip-flop formed from gates "a" and "b" debounces the switch and feeds the single transition thus obtained to "c" of the 7474, thus flipping the 7474 flip-flop. This turns the gate on, and the front edge of the gate triggers the monostable producing a 25 millisecond trigger. The next pressing of the pushbutton flips the 7474 to the state where the Gate is low, but this transition does not produce the trigger. Thus the first push of the pushbutton results in a gate and a trigger, but the second push results in only the disappearance of the gate, no additional trigger. The circuit would seem to have several applications. Of course, it would be ideal for testing envelope generators which require gate and trigger timing signals, particularly if a second "trigger only" switch were added so that ADSR generators could be tested. Also, the push-on push-off feature might make it useful as a footswitch control.

The circuit of Fig. B was submitted by Ellis Cooper and is a simple (possibly the simplest) circuit we have seen that produces an envelope of the general form of the ADSR. It may be useful in some simple synthesizer designs, and the general form

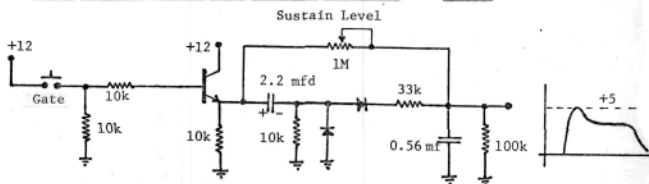
might be used in a fixed mode to provide envelopes of constant parameter such as you might want for certain starting transients. With this sort of arrangement, the envelope generators would be permanently wired to the controller, and all the user would have available would be the output - a constant envelope, take it or leave it. Not a bad way to use that extra panel hole, or with the sustain control, two extra holes. There are probably many times we could use such an extra envelope.

GARY STRENGTH'S GATE AND TRIGGER CIRCUIT



ELLIS COOPER'S SIMPLE ADSR ENVELOPE GENERATOR

FIG. B

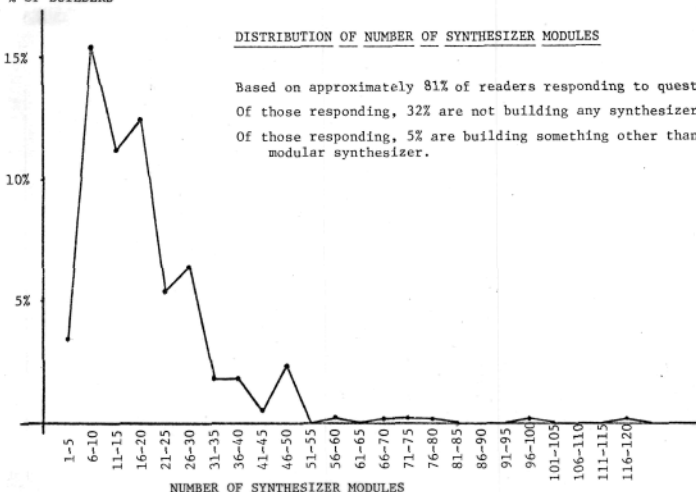


FORUM - ABOUT OUR READERS:

We took a little time to compile the results of the questionnaires that were included with the renewal slips at the end of the last volume. We found out a few things of interest, and since it is always nice to find out how others feel about things you have given your opinion on, we are giving the results below.

The first question concerned the size of equipment being built by our readers. We asked if readers were building a synthesizer, and if so, about how many modules there were in the system. The results are shown in the graph on pg. 18, which includes only those readers who answered the question (81%). Of those responding, about 2/3 were building synthesizers. Note that many builders are building in the range of 10 to 20 modules, while there are still a large number of builders in the 20 to 50 range, and a few very large systems. I think this is about what we would have expected, but there are probably more builders in the 30 to 50 range than we might have expected.

% OF BUILDERS



The second question concerned whether or not more basic type material should be included in the newsletter. Evidently, we could have formulated this question better, because I'm not sure if everyone had in mind what I had when I asked it, and a "write-in" vote turned out to be very popular. The results are:

% Responding to this question: 72%
% Asking for more basic material: 52% of those responding
% Asking for less basic material: 32% of those responding
% Saying OK as is (Write-in): 16% of those responding

The final questions concerned the record reviews of Craig Anderton. We asked first if the reader had read the reviews, and secondly if he thought the reviews should be continued in the future. The results are given below:

% Responding to question of reading reviews: 81%
% Who said they read the reviews: 94% of those responding
% Who said they had not read the reviews: 6% of those responding
% Responding to question of continuing reviews: 73%
% Who favored continuation of the reviews: 82% of those responding
% Not favoring continuation of reviews: 18% of those responding

An earlier discussion of the question of reviews is found in EN#73 (3).

RECORD REVIEW:

-by Craig Anderton

"Sequencer", by Larry Fast. Passport Records PPSD-98014

You'd have to be quite a cynic not to enjoy this album, if for no other reason than Larry Fast seems to be having a wonderful time and that feeling is quite infectious. My first reaction was mixed, but I never trust first reactions anyway...and subsequent listenings have decidedly won me over.

"Sequencer" is Larry's 2nd solo album, and it shows more maturity, musical style, and overall polish than his first album, "Synergy". The first improvement worth noting is the order and balance of the pieces; like good theater, good music requires proper pacing. The 1st side is predominantly up-tempo, bright, and melodic; it even closes with a re-working of the hit "Classical Gas". Turning the record over to find more of the same would have been a mistake--a mistake Larry avoids by putting mostly evocative, lower key, and somewhat less "commercial" material on side 2.

Another plus is the greater familiarity with, and range of, synthesis equipment. The difference is subtle, but sounds have much finer shadings and complexity than the sounds on "Synergy". He has also added a greater number of non-tonal sounds to provide mood-shaping effects, such as background noises and oscillator sweeps that neither begin nor end on a specific note. I'm glad to see someone else exploring this new vocabulary.

One persistent problem not successfully overcome is the subject of rhythm. The extensive use of sequencer type equipment gives a definite sense of tempo, but the actual sounds used for purely percussive purposes are limited to synthesized drum sounds. Some of these are quite good; especially the floor tom like sound on "Icarus" and some of the cymbal effects. But as sounds to me like nothing electronic can sound like an acoustic drum; and even if we could synthesize a drum perfectly, it's doubtful whether any of our equipment would have sufficient dynamic range to capture it. What's needed is a pursuit of what types of sounds lend themselves to percussive synthesis. White noise is a start, but there are also ring-modulated oscillators, plucking muted strings of acoustical instruments, processed voice, and so on. Although "Sequencer" has a beat, it doesn't have percussion to stay on top of that beat and thus turn the beat into rhythm. This has been a problem with many electronic albums.

Another suggestion: a few timbres that aren't produced electronically can give an extra dimension to all-electronic albums; the repetitive, no-surprises aspect of electronic sound can get to be a bit insistent after a full album's worth of listening. I know one of the points of this album is what you can do with electronic keyboards; but some processed acoustic or electronic signal sources would give just that much more variety and texture of sound.

About the individual pieces: "S-Scape" is a gem. Its transitions are smooth and delicately handled, as opposed to some of the splice-like transitions on the first album's compositions. The last part of the piece makes a strong case for automated mixdown, as it builds and builds...only to fade out. "Chateau" has a lack of dynamics, something all the more curious considering the extremely varied dynamics of the rest of the album. There is a superb super-synthesized harpsichord sound; I'm glad that Larry didn't just try to get a good harpsichord effect, but amplified and magnified it. Nonetheless, to me this is the weakest cut on the album, especially after repeated listenings. "Cybersports" is an orchestra of synthesizers, yes; but with sounds that don't necessarily try to duplicate conventional instrumentation...which I like. "Classical Gas" deserves mention because it is, to me, an excellent interpretation of a song I don't care that much for in the first place. The original version by Mason Williams used classical guitar as the focal point, which is no match for a battery of synthesizers in terms of drama or recordability. Although I generally prefer to hear original compositions, this version really does justify the re-working Larry has given it.

On the second side, the Dvorak excerpt is somewhat paradoxical; the electronics give an automated feel (which demonstrates why multiple people in real time can give more expressiveness than multiple tracks in tape time), but the depth of feeling that Larry superimposes on the piece is considerable. It's clear Larry must really love this piece, and that transcends the automation. "Icarus" (by Ralph Towner) is another delightful piece, ideally suited to electronic realization. A lot of synthesized orchestra albums I've heard show an amazing lack of sensitivity, favoring bombast over all else. These last two cuts, and "Sequencer" as a whole, show you don't have to club someone in the ears to make your point. "(Sequence) 14" is the most adventurous piece on the album. It commands your attention from the start, which is just as good because the fadeout on Icarus isn't particularly good. It works its way through non-tonal sections, melodic sections, a massive buildup, and so on, employing a variety of textures rather than melodies to carry the composition. Larry combines these various electronic concepts to create visceral, as well as intellectual, reactions.

As much as I enjoy this album, the type of music Larry is pursuing is still very new. Perhaps ten years from now this album will look primitive in technique, and maybe even in content. But for right now, it demonstrates taste, style, and musical maturity applied to '70s technology. Hopefully other composers will get away from the "gee-whiz" school of synthesis and see what these things can really do.

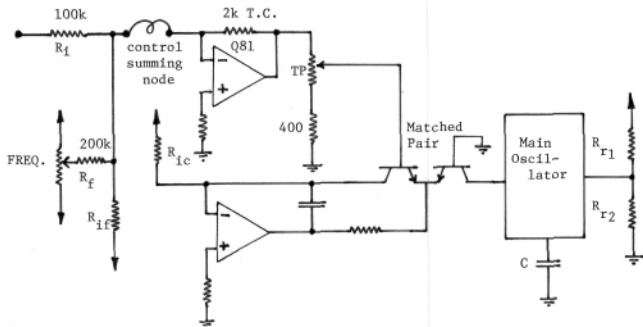
* * * * *

READER'S QUESTIONS (Continued from page 2)

A: You have to keep in mind that the matched monolithic pairs and Q81 resistors are intended to temperature compensate the exponential current stage, and not the entire oscillator. Without the matched pair, drift would be so bad that your oscillator would probably be drifting continually. Without the Q81 resistor, you would have a 1 part in 300 drift term for each °C which will still cause some problems in professional type applications of VCO's, and in some less demanding cases as well. With both these corrections made, your VCO will still drift at least some (nothing is perfect) and may drift very badly. What do you do about it? You find the cause and cure it - easier said than done, but not impossible.

Of all possible drift mechanisms, you will likely be concerned only with drifts related to temperature. This is the type of drift you have cured with the matched monolithic pair and the Q81 resistor. Any you have cured it for all practical purposes. If there is still a substantial drift, you won't find its cause by experimenting more with the matched pair and Q81. Assume that this section is rock stable and you will probably be in a better position to attack the problem. What's left? Well, first off, any device that inputs a current to the control voltage summing node is suspect. This includes all external control devices which may be connected, and unfortunately, any sort of course or fine frequency controls which may seem to be internal to the VCO but which are actually external to the main oscillator. Drifts of power supplies may change reference currents to exponential current sources, reference voltages for waveform levels, and reference currents (initial currents) into the control voltage summing node.

The first step is naturally the isolation of the problem. This is a matter of either identifying the heat sensitive component directly, or determining the pattern of the drift. First, you should isolate the control summing node as much as possible. This means that all external control voltages are removed (including Course and Fine frequency controls). This means that the control sum is zero, and hence the voltage applied to the base of the converting transistor (see figure on page 21) is also zero. In this case, the Q81 is not doing any correction, or you might say that there is no correction to do in this case because the V_{BE} of both transistors of the pair are the same. If the drift problem disappears in this case, it is in fact one of the external voltage sources that is causing the problem. If you still have the drift, you likely have a problem with an internal component, and the main capacitor (C) of



the oscillator is a likely suspect. You can test this and other components by applying heat to them to see what happens. Body heat from fingers is good for this, but keep in mind that fingers are resistive and capacitive as well as warm, so hold the component for 20 to 30 seconds and then let go. Don't rely on any measurements you take while actually holding the component. Soldering irons are also useful for applying heat, but keep in mind that the changes of temperature you get with the soldering iron are both larger and more localized than you are likely to get in normal use. If the main oscillator capacitor is a problem, change it to a more stable type like a good quality polycarbonate type. If C is an ordinary ceramic capacitor for example, you will have substantial pitch shifts when touched with the fingers. Other internal components to suspect are shown in a general type of circuit in the figure above as R_{ic} , 400 ohms, R_{r1} and R_{r2} . R_{ic} supplies an initial current to the exponential converter, so if R_{ic} drifts, the frequency will drift as well. R_{ic} should thus be a stable metal film type or at least carbon film. R_{ri} and R_{r2} will be less of a problem because they form a voltage divider, and as long as they drift together, the ratio is maintained. So R_{r1} and R_{r2} should be of the same type. Don't mix carbon comp. and carbon film here, or in any similar voltage divider. Both could also be metal film of course.

When the control sum is not zero, the trim pot (TP) and the 400 ohm resistor come into play. You can look up or measure the temp. coeff. of the trim pot, and select a type of resistor for the 400 ohms that has a similar temp. coeff., or at least one in the same direction. Let's also look to the left of the control summing node, as this is where the control voltages are coming from. Resistors of the type R_i or R_f should be metal film if they are much less than 1 meg. The resistor R_{if} is an initial freq. setting resistor, which can often be avoided by increasing the ranges of the frequency controls for example. If it's a problem, try to get rid of it, or make it metal film. Note that this resistor is in oscillators of the type shown in the ENS-76 Option 1a circuit - its the feedback correction resistor.

What about drift due to power supply voltage drift? Often times, this is not a problem because supply voltage shifts at one point which would upset the circuit are compensated at another point. For example, if the (+) supply drops, the current through R_{ic} drops, but so does the reference from R_{r1} and R_{r2} . Shifts of one supply through a pot such as the FREQ. control will be more of a problem. Tracking regulators are a possible solution. Also, some builders use separate regulators to supply the VCO's, or just the main oscillator sections of the VCO's, leaving the waveshapers and other modules to the main supply regulators. If you expect to have an uneven load, this is worth considering.

At some point, you may have pretty well temperature stabilized the VCO, but want to improve it even more for a production model, or for your own performance requirements. At this point, you may have to rig up a make-shift "oven" to carry on your testing. For this, one of those styrofoam "picnic coolers" may work well. You place inside it a good quality laboratory thermometer, a compartment for holding some ice, and a small (e.g., 10 watt) light bulb to provide heat. You then put the VCO inside and run leads out to a frequency counter. First you put in some ice and let the VCO cool down to about 0°C or whatever stable low temperature you happen to get. Make sure the temperature is stable and has been so for an hour or so. Read the frequency for several readings and average them. Then remove the ice and let the VCO start to warm up. It may be useful to use the light bulb to bring the temperature slightly above room temperature, and then let it settle down to a stable value. You again measure the frequency. Taking data points at temperatures above room temperature is a little more difficult, but can be done. Repeat all measurements several times to get a good plot of frequency as a function of temperature. The graph you plot from the data should be either linear or well approximated by a linear function. Unless something is drastically wrong, the drift should be less than 1 part in 300 per °C, and thus can be corrected by another Q81 resistor installed in an appropriate position.

Even these steps will not assure that the VCO will not drift as a result of some sort of differential heating or cooling, but the performance should be excellent under normal conditions.

CLASSIFIEDS:

FOR SALE: 61-Note Single-Contact Electronic Music Keyboards, \$55 plus shipping, and sales tax where applicable. These are brand-new high-quality units. E. D. Cooper, 200 W. 14th St. (#2E), New York, NY 10011

FOR SALE: MOOG IIC modular synthesizer \$2900.00. Additional modules available separately for 1/2 retail price: 911's, 911A, 912, 961, 962, 1150, 959, dual S&H (\$175.00). Call David at 216-781-2207.

FOR SALE: Precision Resistors, Edge Connectors: Available early May - Corning NC55 100k 0.1% 1/4w; 1-4, \$1.00 ea; 5-9, 0.75 ea; 10-24, 0.60 ea; 25-49, 0.53 ea; 50 up, 0.49 ea. Corning NA55 100k 1% 1/4w; 1-9, 0.45 ea; 10-24, 0.20 ea; 25-49, 0.18 ea; 50 up 0.16 ea. Available now - Amphenol 143-827-1003 22 pin edge connectors for 0.055 - 0.073" boards, 0.156 centers, gold tuning fork contacts, 2 mounting holes; 1-4, 1.25 ea; 5-9, 0.50 ea; 10-19, 0.45 ea; 20 up, 0.40 ea. All prices handled, insured, postpaid. Early 100k orders greatly appreciated. Jeff Smith, 1875 Commonwealth #6, Brighton, MA 02135.

MUST SELL: Assembled ENS modules from MEH, on PC boards with sockets, tested. (# cermet pots, size 1/8" alum panel) 2-VCO (7, 6"x9") @ \$35; dual ADSR (9, 6"x6") @ \$35; 2-VCA (5, 6"x6") @ \$30; VCF#1 (6, 6"x9") @ \$35; VCF#2 (9, 6"x9") @ \$35; Bal Mod and Noise source (4, 6"x6") @ \$25; S&H (4, 6"x6") @ \$15; kybd interface (5, 6"x6") @ \$25. Also Eu 1200 with waveshapers (6, 6"x9") @ \$60. Package price available. Tom Helfrich, Box 66, RD#2, Lancaster, PA 17603 717-872-4963

FOR SALE: As a major manufacturer of custom active filters we stock a large number of 1% metal film resistors, RN55D and RN60D. Since many EN projects work better with these resistors, we're offering them to EN readers at 70 for \$10 ppd. in US, any mix of values. Certified check or MO only please. SASE brings list of values from 10 ohms to 1 Meg. TDL Electronics, Route 7, Fayetteville, ARK 72701. (501)-643-2191.

FOR SALE: ARP Odyssey, several modifications, complete service, instruction manuals, and schematics. Best offer around \$1000. Call Steven Pope, 607-273-5218

NEW RELEASE: Mother Mallard's new album "Like a Duck to Water." Draws: Waterwheel and Oleo Strut. Borden: C-A-G-E Part 2. Available from Earthquake Records, PO Box 842, Ithaca, NY 14850. \$5.00 each. NYS residents add 7% sales tax.

ELECTRONOTES, Vol. 9, No. 76 [April 1977]

Published by B. A. Hutchins, 203 Snyder Hill Rd., Ithaca, NY 14850

Send Routine Orders to: ELECTRONOTES, 213 Dryden Rd., Ithaca, NY 14850 (607)-273-8030