With these new weightings we thus arrive at the equally tempered sequence:

$$0\ 0 \rightarrow 1 \quad (= 2^{0/12})$$

$$0\ 1 \rightarrow 2^{1/12}$$

$$1\ 0 \rightarrow 2^{2/12}$$

$$1\ 1 \rightarrow 2^{3/12}$$

As successive stages are added, the equally tempered nature of the series holds. The exact resistance ratios are obtained by solving the first two equations in Fig. 3.

There you have it, a way to generate 5 octaves of control voltage to a linear oscillator from six bits of code. Further, I think that you will agree that this system behaves in <u>every way</u> as if an exponential converter were in the system somewhere. I didn't pull this out of my --ear-- sitting at the typewriter. It's the result of a lot of thought that started (according to my lab notes) over two years ago. So if I have thrown in a bunch of "obvioulsy's", take that into account.

John Simonton
PAIA Electronics
Jan. 6, 1976

* * * * * *

<u>Design and Applications Of Sample-And-Hold Modules With Slew-Limiting And Cascading Features:</u>

Bernie Hutchins, ELECTRONOTES

By now, the sample-and-hold (S&H) module is a common feature of many synthesizer systems. Their function is to capture an instantaneous voltage level and then hold this voltage until another sample command arrives. An elementary S&H system is formed as shown at the right. The switch S is closed for a brief instant during which the capacitor C charges toward the instantaneous level of V_{in}. The rate of charge is determined by the output impedance (R_{out}) of the device supplying V_{in}. If the value of R_{out} is low enough and if the sampling time is long enough, then the output of the buffer will reach the instantaneous value of the input without noticeable error, and we say we have a successful S&H design. If on the other hand the value of R_{out} is too high or the sampling time is too short, then the value of V_{out} will not reach the instantaneous value of V_{in} without noticable error. While the output tends in the right direction, it does not make it. This sort of S&H design is an engineering failure - which is not to say that it may not be useful for certain musical applications. In fact, we will look at the case where we make sure the design is slew-limited, and will find some very useful musical results.

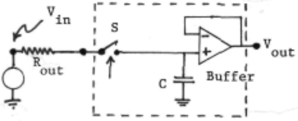

The general effect of slew-limiting of a S&H circuit is shown by the diagram at the top of the next page. We look there at the case where the device is made to respond to an input ramp. We have assumed for the purpose of making this diagram that it is the effect of R_{out} which is causing the slew limiting (hence the exponential response). In simple terms, the response is "sluggish" and lags behind the ideal response.

A very common use of the conventional S&H unit is to generate a "random" sequence of voltage levels by sampling a white noise source. It will be one of the main purposes of this report to examine how this process is changed when slew limiting is applied.

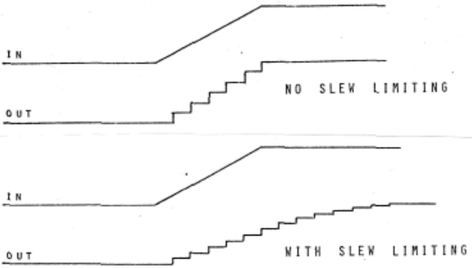

We want to look here at probability distributions of so called random samples, and also the actual structures of the sequences. First we will look at the probability distribution of the samples obtained from the usual type of white noise generator. Then we will look at a contrived example where we start with the equivalent of a flat probability distribution, and see how it is altered by slew-limiting. Moving on to sequences we will first discuss the meaning of random and non-random sequences, and then examine how slew limiting changes the note-to-note structure of sampled white noise sequences. It will be found that the major effect of slew-limiting, and the one with the most musical significance is that the tones of the sequence tend to group into what we can call "bursts" of tones where the note-to-note change is relatively small while long term structure remains random. Subjectively, this means that the sequences that we obtain will be uncontrolled, but have an overall texture that is much smoother than that obtained in the normal (non slew-limited) manner.

PROBABILITY DISTRIBUTIONS

We show at the right a setup using a typical white noise generator of the type formed from a back biased semiconductor junction. The circuit described by Rossum in EN#30 (4) was the one we used for the experiment that follows. Now, the voltage level of the white noise changes very rapidly relative to human perception, and since there are no patterns that repeat, we hear only a "texture" which we describe as noise. The S&H captures a voltage level in this noise and slows it down by holding it. Typically this voltage that is being held is fed to the control input of a VCO, and is thus converted to a musical pitch. Now suppose you are a gambler. You are to bet on the interval of voltage in which an individual white noise sample will fall. You know that the noise source and S&H are supplied by a bipolar 15 volt supply so you don't expect any voltages greater than +15 or less than -15. But are you wise to suppose that a voltage in the range of say +12 to +13 volts is just as likely as a voltage between +1 and +2 volts for example. If you have ever listened to this sort of random pitch sequence, you probably know intuitively that extremely high or extremely low pitches are rare. The best bet is the interval around zero volts. The reason for this has to do with the physics of the semiconductor junction. We can do an experiment by setting up the noise source and sample-and-hold and measuring voltages. We then form a histogram by adding up all the sample voltages that fall in a certain interval. The results of such an experiment are shown on the next page. The distribution is the result of over 500 samples. We are not concerned here with the exact range of values or any individual part of the histogram, but only with the general shape, which some would say seems to be "Gaussian" or a normal "bell curve."

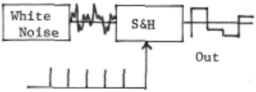

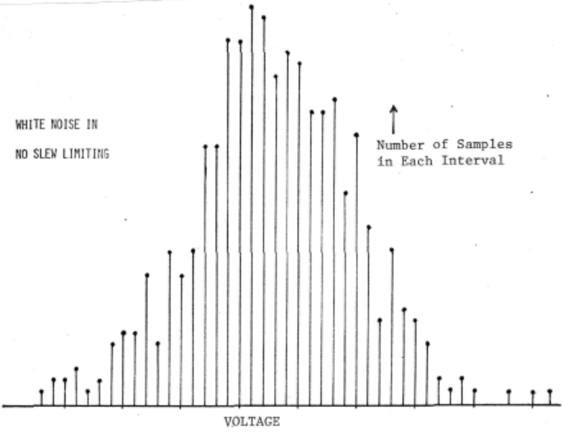

It is clear from the above histogram why we get more voltage samples near the center than near the ends. We don't really want to say too much about what the distribution means from an engineering point of view. In the first place the results show the combined action of the noise source and the S&H working together. Imperfections in the sampling process will thus alter the results, and we could not claim that we have an exact measure of the probability distribution of the noise. While the result looks Guassian, it can not be exactly Gaussian since we know for sure there will be no voltages exceeding the power supply limits, and for a true Guassian noise source we expect a finite probability beyond the supply limits. Finally, the physics of the semiconductor junction tells us that there will be a Guassian component to the noise, but there will also be other components. What we do want to say is that for the usual type of noise source, and the usual type of S&H, the above distribution will be typical of the distribution of amplitudes one will get. In short, things are crouded together toward the middle, a fact that we noted may be intuitively obvious to many synthesizer users.

Now, we want to examine the effect of slew limiting on the sampling process. One problem is that if we use a standard white noise source, the probability distribution is already "funny" and this will complicate things. We would prefer to start with a distribution that is flat on the top to begin with so that we can better see the change due to slew limiting. One way around this would be to develope a pseudo-random source and adjust the summing weights so that the probability distribution was flat. An easier way is the following: We start with some waveform that spends its time equally at all voltages (a triangle is good, or a sawtooth can be used, but not a sine or a pulse). We then sample this at "random" times. This is accomplished by using a relatively high frequency for the triangle (e.g., 1000 Hz) and use a manual trigger to the S&H. Since a human being could not time his manual actions to 1/1000 of a second (or even a shorter time to capture a certain part of the waveform), the result is the equivalent of a set of random voltage levels with a flat probability distribution. An experimental histogram of this type is shown at the top of the next page.

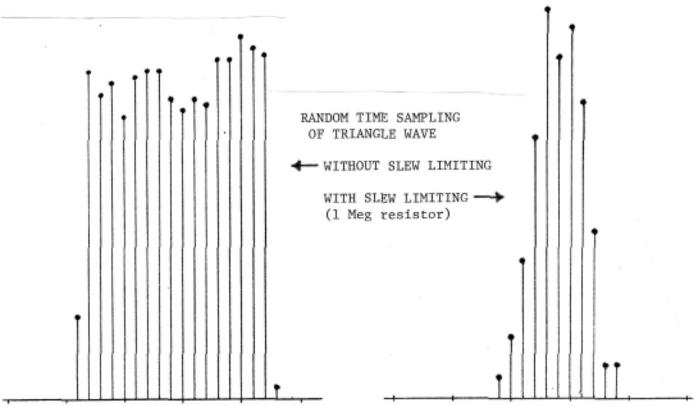

The distribution shown on the left above shows that we have come close enough to getting a flat top. The two small lines on either side are due to the fact that the triangle wave partially entered an interval but did not fully sweep it. If we had been more careful, we could have gotten rid of the side features.

We are now in a position to examine the effect of slew limiting without having to worry about funny effects in the original distribution. The distribution shown above on the right shows the effect of slew limiting (a 1 meg resistance in the circuit we will be giving later on). We see that we have obtained a rather peaked distribution rather than the flat one we started with. This is all easily understood. For the output of the slew limited S&H to reach one of the extremes of the original distribution, we would have to have several successive samples called for at the extreme voltage. Since samples are at random, this would be a rare event. Rather than go into a detailed explanation here, the reader will gain the proper understanding by spending the same amount of time just thinking about this.

Note that the dynamic range of the output is not changed by slew limiting. Although we did not observe any samples toward the extremes of the distribution, we would have if we had waited long enough. As a practical matter (musical appliactions), the dynamic range appears reduced on the time scale of interest, and it may be desirable to add some amplification to the output in the case of slew limiting. We did the above experiment on our triangle random sampling distribution rather than on white noise. What happens if we apply slew limiting to white noise? Clearly the distribution on page 6 would be further narrowed and even more samples would be concentrated toward the center of the distribution. So far we have been talking about probability distributions which are really a matter of personal preference - do you prefer your random sequences to have relatively more samples around zero or not. Musically, this is probably much less important than the effect of slew limiting on the note-by-note structure of sequences which we shall examine below.

SUBJECTIVELY RANDOM AND SUBJECTIVELY ORDERED SEQUENCES

It is easy to misuse the term "random" when applied to musical compositional processes. However, it is not our purpose here to rogorously define the term.

Rather we want to point out that a sequence may be subjectively random or selectively ordered regardless of its true nature in a rigorous mathematical sense. As an example, consider the following sequences, and consider if they will be subjectively random or subjectively ordered:

(a)	1 8 2 8 1 8 2 8
(b)	1 0 5 9 4 6 3 1
(c)	7 9 2 7 2 1 1 6
(d)	7 7 7 7 7 7

It may be of interest to actually play these sequences as musical pitches (as whole tones for example) to get a better feel for their nature.

We now want to look at the true origins of the different sequences. [The perceptive reader has probably guessed that these examples have been selected or rigged so that things are not necessarily as they first appear.] The first example (a) is just part of the expansion of the irrational number "e", the base of the natural log or 2.718281828459045... It is not random since every digit is predetermined by the fact that we say it is the number e. It also happens that this sequence (the 18281828 part) is subjectively ordered due to the repeating pattern. It would have been just as easy to select another part of the sequence (say 82845904) from just about anywhere else in the expansion for "e" and we arrive at a sequence that is not random, but which is subjectively random since the mind does not recognize any repeating patterns. Thus the 18281828 sequence is just an accident. Note however that the same sequence could have originated from a conscious effort to write down an ordered sequence consisting of the pattern 1828. In this case, the sequence would be not random and not subjectively random, a combination that might be assumed the most natural.

The second sequence (b), 10594631 is part of the expansion for the 12th root of two which is familiar as the basis of the twelve-tone equally-tempered scale. It is not random, but is subjectively random. This is the same case as you get by staying away from the "freak" (18281828) portion of the expansion of e in sequence (a).

The third sequence (c), 79272116 is a random number obtained by selecting some pages at random and writing down their numbers in series. It is random and is also subjectively random. This is another of what we might call expected cases.

The final sequence (d), 777777 is obviously subjectively ordered (not subjectively random). It is however a random number since it is obtained from a random number table.* This would be roughly equivalent to the case where white noise was being sampled and the same voltage happened to come up six times in a row.

\- - - - - - - - - -

* This is from A Million Random Digits with 100,000 Normal Deviates prepared by the Rand Corp. and published by the Free Press in 1955. The page of interest is published by Martin Gardner in his "Mathematical Games" column in Scientific American for July 1968. This is an excellent column on the meaning of randomness. The current author is indebted to the anonymous person who circled the sequence of 7's in the copy of Gardner's article at Cornell University's Clark Hall Library. (The same graditude perhaps can not be assumed on the part of the librarian). One is tempted to ask if this series of six sevens is to be expected. I have done a few calculations but am not sure the results are right. I would like to hear from readers with answers to the following questions: What is the probability of six sevens in a table of random digits? What is the probability that such a sequence is in the Rand Tables (1,000,000 digits so almost 1,000,000 sequences of six digits)? I concluded that the sequence is expected somewhere in the table. However,it is not expected that the sequence should appear on any one page selected for reprint. Gardner does not make any reference to the series of sevens in the text of his column.

We have thus seen examples of all four possible cases as shown in the table below:

SEQUENCE	RANDOM?	SUBJECTIVELY RANDOM?
1 8 2 8 1 8 2 8	NO	NO
1 0 5 9 4 6 3 1	NO	YES
7 0 2 7 2 1 1 6	YES	YES
7 7 7 7 7 7	YES	NO

It is the idea of "subjective ordering", or lack of subjective randomness, that we are interested in, and which we suggest will be induced by slew limiting. Due to the fact that common white noise sources do not have a flat probability distribution, there is a slight amount of subjective ordering to start with, but this is nowhere near as strong as when slew limiting is applied.

THE EFFECT OF SLEW-LIMITING ON THE STRUCTURE OF SAMPLED SEQUENCES

The major effect of slew-limiting on sampled sequence structure can be seen by considering the example below. The input is a square wave (around 1000 Hz) and the sampling was done manually to obtain the equivalent of a random selection between two possible levels. With no slew limiting, the sequence emerges as indicated by the top diagram below - you get one level or the other. [The lines between sample points have been put in to aid the reader in following changes. They do not represent voltage levels. All voltage changes in the output sequence are discrete jumps.] When slew-limiting is imposed, we obtain a sequence that has many possible levels and which in general stays away from the extreme voltages. Note that no two voltages in succession ever remain the same in this case, and in general all jumps are roughly the same order of magnitude (the jumps as seen are influenced by the exponential rate of change of the slew limiting process and by a rounding process used to simplify data recording and plotting). The reader should bear in mind that this sort of sequence can only be obtained with manual (forced random) sampling. If a periodic trigger is used, the sequence will be highly ordered. If the input square wave frequency is held up around 1000 Hz, and manual sampling is done by a simple pushbutton, manual sampling will in general give a subjectively random output. This is because even if the person pushing the button does his best to achieve an exact tempo, it will not be accurate enough to order the sequence.

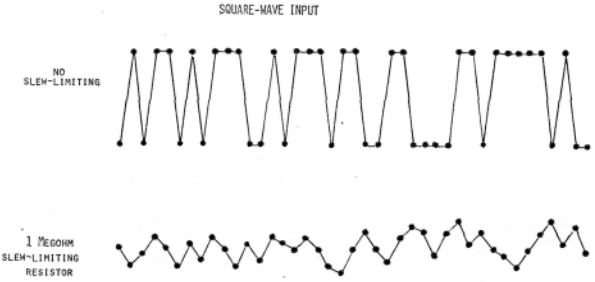

We move now to a case that is more easily related to the usual application of the S&H for sequence generation. In this case, it is the usual practive to trigger the S&H from an oscillator which is periodic and use a white noise source as the input source. For our experiment we will be using the "pseudo flat distribution trick" of manually sampling a triangle as was done for the figures on page 7. The two examples below correspond exactly to the two examples on page 7. With no slew limiting, successive samples are purely at random, and you will note that they are well scattered. In the lower example slew-limiting has been applied so that successive samples are to a degree correlated and a subjective feeling of ordering results. This is easily seen in terms of long term trends. For illustration we have added a dotted line which could be considered a threshold below which the output is blanked, and above which the output is allowed. It can be seen in the case of the imposed threshold that with no slew limiting the samples tend to occur so that only one or two emerge at one time. In the case of slew limiting, they tend to occur in "bursts" of many voltages, or tones, if we think of the sequence as driving a VCO as is a usual case. The illustration below shows what is essentially a single burst for the case of slew limiting. The essential difference between the two is that in the case of no slew-limiting the tones penetrate in rapidly occuring groups of one or two, while in the case of slew limiting the tones occur as bursts of something like 10 tones, and the tones in the burst are to a degree correlated.

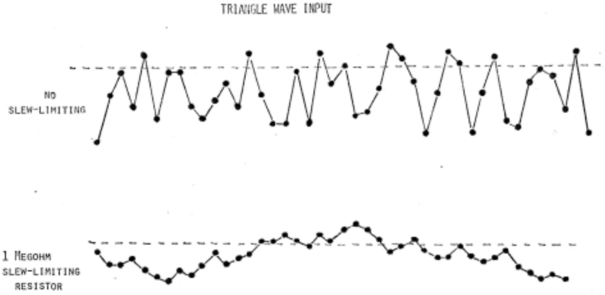

The situation using an actual white noise source and a slew-limited S&H is very similar to the case illustrated above. We found that when no slew limiting was used, and a threshold was used, it was the case that there was never a time as long as a second when no tone emerged. With the slew-limiting, silent periods of as many as 8 seconds occured and tone bursts of the same relative duration occured. The above results were obtained with a sampling rate of 3-4 Hz. Also, with slew-limiting the tones in the burst moved in what might be described as a more purposful manner as compared to the scatter of the more rapid tones with no slew limiting.

THE SAMPLE-AND-HOLD CIRCUIT

The circuit for the S&H used for these experiments is shown at the top of the next page. This basic S&H mechanism has been used in several applications and works well. Note that the control part of the circuitry in connected between ground and -15, and will provide internal, external, or single pushbutton sampling. Inputs should be in the range of -5 to +5 volts. For -10 to +10 volt operation the gain of the input amplifier should be changed to -1/2 rather than -1, and the output should be scaled to a gain of two. Note that the actual sampled voltage is stored as the negative of the actual voltage that

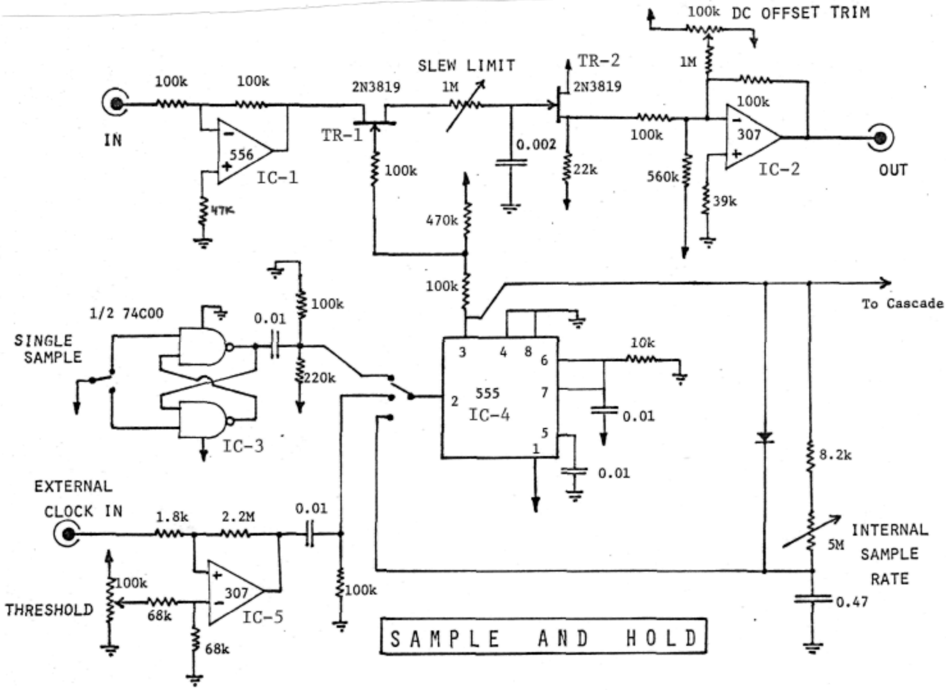

appeared at the input. This really makes no difference and the inversion is corrected by the output amplifier. It is just more convenient to handle the input and output stages as inverters.

Slew limition has been added to the circuit by means of the 1 Meg "SLEW LIMIT" pot. With this pot set to zero resistance, we have a normal S&H action. Slew limiting can then be added by increasing the resistance of the pot. The cascade feature of the circuit will be discussed later on in this report.

IC-1 serves as an input inverter/buffer. TR-1 serves as the sampling switch while TR-2 forms a source-follower buffer for the voltage held on the 0.002 mfd sampling capacitor. IC-2 serves to invert the held voltage and adjust the DC level of the output. IC-4 is a 555 timer which operates in either an astable mode or a monostable mode. In either case it provides pulses (from a -15 level to ground and back to -15) of about 100 microseconds that serve to trigger the FET switch TR-1. This pulse appears at pin 3 of the 555 and the attached resistor network gives the proper voltage at the gate of TR-1 to turn it off (pin 3 at -15) or on (pin 3 at ground). IC-4 can be triggered by IC-3 which forms a bounceless pushbutton for single manual samples. IC-4 can also be clocked externally from IC-5. Note that a variable threshold can be set for IC-5. This is very useful for triggering from a noise source so that the sample times are random. In this case the threshold can be set high enough so that it does not sample too fast to be followed, as would be the case if the zero crossings of the white noise were used. With the threshold raised, triggering can occur only on the peaks in the white noise, and

these occur much less frequently. To avoid the use of an external oscillator when this is not needed, an internal rate control circuit is also available.

The inclusion of the slew-limiting resistor is just the simplest way of achieving the desired effect. It is something that can probably be added simply to other S&H designs already in existence. At least two other means of slew limiting are possible: First one could simply adjust the sampling time so that it reached values too short for full sampling. Secondly, certain types of S&H designs (such as the one in the CA3080 application notes) function by charging a capacitor with a constant current until it reaches the input level. This would give linear slew, and the slew could be adjusted (and thus limited) by changing the charging current. Note however that it is not possible to rig a slew limit with an external slew limiting module. Adding a slew limiter to the input of the S&H or the output will give different effects. To be properly done, the slew limiting must appear at the switch. One other thing that was not done (for lack of a dual pot with different valued sections) is to adjust the gain of the output to higher values as the slew limiting is increased. As we mentioned, the slew limiting decreases the expected short term dynamic range of the output relative to the input. For musical purposes, it might be useful to have this relatively constant. The proper dual pot and some experimenting might show that this control would be useful.

A "FUNNY" FILTER

The slew-limited sample-and-hold as described could also be called a single stage commutating filter. The fact that it is a sort of filter can be understood as follows: At low frequencies, the change between any two input samples is small since the input waveform varies slowly. The slew limiting does not restrict the response. As frequency rises, it becomes more difficult for it to catch up, and amplitude drops off. Then you run into the fact that you are working with a sample-data (discrete time) system, and frequency aliasing becomes a problem. When you start to approach any integer multiple of the sampling frequency another peak in output amplitude occurs. The output frequency however is aliased down. None the less, we can take a frequency response curve for a sinewave input, and the result as seen below shows that slew-limiting induces a sort of comb filter response curve:

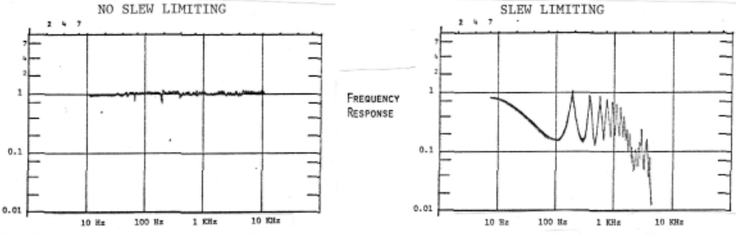

Now, this is not a comb filter by any means, but the response looks the same. We can go ahead and consider a patch of the type shown at the top of the next page. We use an envelope follower to extract the envelope of the frequency response curve, and apply the extracted amplitude to a VCA which controls the amplitude of the signal being sampled. Now, if we put a sine wave through this, we have a comb filter for sine waves. Now, if we put a general type of signal in, does it act as a comb filter? Of course not. It's not a linear system and superposition does not apply. But what do you get out? The answer quite simply is that you get some of the most annoying sounds you could hope to generate. White noise in suprisingly gives one of the most pleasent outputs. When live music or a human voice go in, you get an output whenever the waveform just happens to hit on the sampling commands for a long enough time to raise the signal to the VCA. This

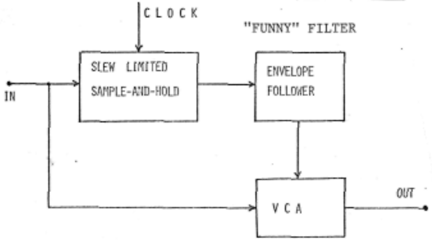

gives short and seemingly random bursts of the program material. This in general gives an unpleasent feeling and properly done, an exposure to a few minutes of this sort of sound should empty any home of unwanted guest. Is this musical? It would seem that it might be used in some cases for certain types of musical expression.

CASCADING SAMPLE-AND-HOLD DESIGNS

If before taking a new sample, one S&H unit passes its voltage on to another, the *voltage level is retained for use later. This is the basic principle of the analog* delay line. We are concerned here with an elementary analog delay line consisting of S&H modules. We form this by cascading S&H units so that voltages are passed from left to right while clocking commands pass from right to left. Thus, a given S&H *steals the voltage to its left and then tells the S&H on the left to steal the voltage* of the S&H further left, and so on. This is illustrated for three units below:

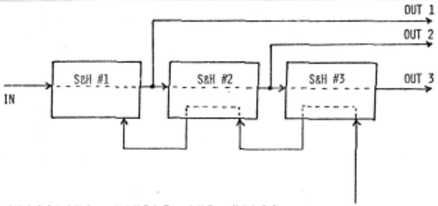

CASCADING SAMPLE-AND-HOLDS

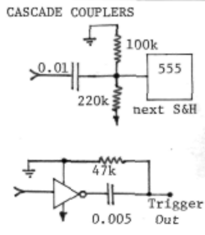

There are several useful applications for this type of device, and we shall mention three: (1) The input can be sampled white noise, and the outputs drive three *separate VCO's, the outputs of which are mixed. This gives a "Canon" structure of* random sequences. (2) The output can be looped back to give an "infinite" sequence. The term infinite is in quotes because the looped line will degrade, a fact which serves to make the system musically useful, as the degrading process is an evolution *of the sequence material. (3) If controlled by a keyboard, musical chords can be* "rolled" rapidly. This loads the notes of the chord into the S&H units and makes a versatile chord playing device.

Cascading is achieved using the "to cascade" point in the circuit on page 11. This can be done with the simple 555 cascade if the units are hard wired, or a trigger-out terminal can be formed from one of the remaining CMOS gates. This trigger is then patched to the external input of another S&H as needed. The necessary circuits are shown above.