ELECTRONOTES 110

NEWSLETTER OF THE MUSICAL ENGINEERING GROUP

1 PHEASANT LANE

ITHACA, NEW YORK 14850

VOLUME 12, NUMBER 110 FEBRUARY 1980

GROUP ANNOUNCEMENTS:

CONTENTS OF EN#109

This issue, like the previous one, has a major report that involves a lot of math. The reader not inclined to mathematics should read the PWM report lightly, mostly looking at the resulting spectra. Pay particular attention to the patches on page 14, as they show what is perhaps the most useful thing to come out of the math. Students taking "engineering math" courses will perhaps find a lot of applied math in this report that will help motivate them. The construction project at the end is a new type of contour or envelope generator which works quite nicely, but applications are still being developed.

A SHORT LIST OF "BUZZ WORDS" FOR THE IC BUYER AND USER:

Many of our readers are just starting to use IC's and/or ordering them in large(r) quantities. We often get questions as to what certain terms mean. The following list will help you when looking over data sheets, and when talking to salesmen. At the very least, in the manner of all "buzz word" lists, they will help you to sound like you know what you are talking about, even when you don't.

Absolute Max Electrical vaules of an IC which, if exceeded, may well destroy it.

Anti-Stat Anti-static electricity. Protective shipping for sensitive FET IC's

Bonding Wires Wires inside the IC, usually gold, which connect the external leads to the internal silicon structures.

Bug IC's often look like bugs, and are called the same.

Can A type of IC package made of metal, and often round, usually with 8 or 10 wire leads out the base. The can is about 3/8" round and 1/4" high for the usual package, but there is also a much larger can for voltage regulators, etc. Also called "metal can," "metal," or just "round package."

Chip Technically just the inside silicon part of the IC, but often used for whole IC.

Date Code A four-digit number on an IC, the first two being the last two of the year, and the last two being the week the IC was made. Example, 7914 would be the 14th week of 1979.

Delivery Time   Just what it says - and don't forget to ask.

DIP   Dual Inline Plastic (or Package). The common IC package with two rows of legs on either side of a plastic body. Typically there are 8, 14, 16, 24, 28, or 40 legs total, although others are also found.

Expediter   When you call about an order, and don't remember who you placed it with, ask for an expediter. This is a person who can find your lost order by the fourth try nearly every time.

Family   The basic group to which a certain IC belongs. Typical family names would be TTL, CMOS, Linear, BiFET, and so on, although the term is not totally defined.

Gold Adder   Not really a precious metal snake, but can be dangerous. This is a new term, an added cost for the gold content of IC's as the gold market goes up and down. Typically a few cents per IC, but not always stated ahead of time. Ask.

House Number   Not your house, but a manufacturer wants his part number on an IC rather than the one the manufacturer would normally put on. They do this to aid in their assembly effort. If such parts get outside the house, they can be difficult to identify as well, for better or worse.

Mil-Spec   An especially reliable part (Military Specifications). You don't need it or the added cost. You need a "commercial" grade usually.

Mini-DIP   The 8-pin DIP package. Often written mDIP. The usual op-amp package.

OEM   Original Equipment Manufacturer. You probably aren't one, but you may be buying in a quantity where you can get OEM prices.

Package   The "package" of the IC such as DIP or Can. Another common package is "ceramic" which is often a DIP type, and a better grade device.

Pin-out   A listing of the functions of the various pins or leads of an IC, and often with some indication of the physical orientation of these leads.

Process   The IC fabrication process, usually related to the family.

Purchase Order   Often just "PO." An offer to buy without sending cash with the order. Usually the seller wants payment within 30 days after he delivers. Usually for well known companies ("rated firms").

Rail   This is the funny shaped plastic or metal tubes, about two feet long, which DIP IC's come in end to end, perhaps 40 or 50 mini-DIP's to the rail. These can be ordinary plastic, or anit-stat plastic or metal for IC's that can't tolerate static electricity.

Second Source   This means that the part is available from another manufacturer. This is, strangely enough, something manufacturers are glad to tell you. Why? Because if you are making a product, you want to know you can get the same part from another source if the first one has delays for any reason.

Stock   Or "in-stock." This does not always mean that the parts are on the shelf behind the salesman. It does mean that the parts are available to him without going to the manufacturer (or that his computer says they are), usually through another branch of his distributership.

Type Number   The "type number" of an IC usually consists of three parts. The prefex is often one, two, or three letters identifying the manufacturer (although not always related to his name). The body is the actual number, and all or part of this number may be related to the number of another manufacturer who makes the same or a similar product. The suffex is usually one or two letters identifying the "grade" of the product and/or the packaging. The grade is very important in that better grades cost a lot more, and you probably want the cheapest.

Velofoam   A black "sponge rubber" type of material that is conductave, used as an anit-stat shipping and storage medium.

A NEW LOOK AT PULSE WIDTH MODULATION - PART 2:

-by Bernie Hutchins and Lester Ludwig

INTRODUCTION: In part 1, we looked at a "quasi-static" case of pulse-width modulation (PWM). Here we will begin to look at the case where the quasi-static case does not apply any longer. A simple analogy with frequency modulation (FM) will serve to show the direction we are now moving. With a slowly varying frequency, the ear can easily follow the changes up and down, but as is well known[1] *, somewhere in the range of 7 Hz to 15 Hz, the ear gives up following the change of frequency and begins to resort to an overall impression, which can be quantified by the mathematical "sideband" structure of the FM process. When we get well above 15 Hz, the sideband process is quite effective in describing the resulting sound. In the case of PWM, it is timbre rather than pitch that is actually modulated. Slow changes of pulse-width (the quasi-static case) can be followed by the ear. Again, once these changes get faster than about 7 Hz to 15 Hz, the actual change gives way to an overall impression. This is the area we want to get into now. As expected, the mathematics and interpretation gets more complicated, but final results (useful patches) need not increase by the same degree of complexity.

Central to our investigation is the determination of the spectrum of a PWM signal. We will illustrate several useful methods of arriving at the spectrum (or at least a good notion of what the spectral features are). Taking the overall Fourier Series of a harmonic spectrum will prove the most useful in the long run, but other methods (direct expansion in Bessel functions, convolution of inferred multiplying waveforms, reconstruction of the PWM waveform from known waveforms) will also tie things down and provide intuitive understanding.

We will find that there is an important difference between PWM signals resulting from different driving waveforms and different reference waveforms, some cases of which will cause a shifting of spectral energy and a resulting apparent pitch shift. These cases appear mainly where both the reference and the drive are audio frequency waveforms, and not in quasi-static cases or near-quasi-static cases, where the same patches may give an annoying "pop" in the output sound.

PWM belongs to a class of modulations that are useful for activating[2] otherwise static waveforms, but which do not alter the amplitude of the waveform. Under the right conditions (rapid rate of modulation), nearly any modulation method, including AM and balanced modulation, can enter this class, but the principal members are the various phase modulations (including FM), and the various pulse modulations (including PWM). A comparison of FM and PWM shows that PWM implementation is simpler, involving mainly a comparator. At the same time, the use of FM in electronic music is much more extensively explored. Consequently we will be looking at procedures familiar from FM work to see if we can apply them to PWM. In particular, dynamic depth PWM will be explored.

BASICS AND A SAMPLE RESULT:

A basic PWM system involves two waveforms and a comparator, in the manner of Fig. 1 where one waveform is called the reference while the other is called the drive. Most electronic music VCO's already contain this stage as their pulse producing circuit, in which case the drive is provided by any convenient VCO waveform, and the reference is obtained from a pulse width control circuit which is typically a manually set voltage (inital pulse width) and a variable voltage-control input. It is often the case that the reference in such a VCO is an external envelope. Note however that the choice between reference and drive is somewhat arbitrary.

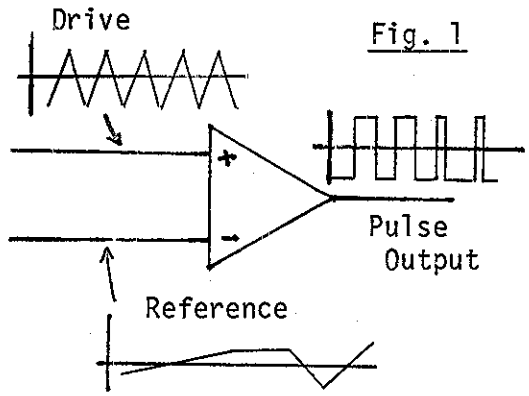

Fig. 1

*all footnotes at end of report

Let's assume that the drive input to the comparator is some periodic waveform, and any waveform such as a saw, triangle, or sine will do. Now, if the reference is a fixed voltage (and within the range of the periodic waveform), the output of the comparator is a periodic pulse train of fixed pulse width. If the pulse width changes very slowly as a result of slow changes of reference voltage, we have the quasi-static case we discussed in part 1. Here we want to consider the case where the reference waveform is also periodic and is of an audio frequency. To make things a bit more concrete, let's say you have one VCO with PWM input, and this is initially at an audio frequency. You then connect to the PWM control input an audio signal from a second VCO. The first VCO provides, internally, the drive, while the second VCO provides the reference.

The computation of the spectrum of the PWM output in the case where the reference frequency and the drive frequency are arbitrary is rather complicated, but there are cases where we can use a simple procedure. In the case where the reference and the drive are harmonically related, the PWM output is periodic, taking on the period of the drive or the reference, according to whichever of the two has a longer period. A typical example of this harmonic case is seen in Fig. 2 where both drive and reference are upward going sawtooth waves, with the drive having eight times the frequency of the reference, as shown in the upper portion of the figure. The middle portion of the figure shows the resulting pulse train[3], where we consider the pulse to be high when the drive is higher than the reference. Since the two sawtooth waves are harmonically related (f and 8f), the relationship in the top portion of the figure, and the pulse train in the middle, can be thought of as repeating periodically. The bottom portion of the figure shows the spectrum, which we will get to shortly.

To get a better idea of what we are getting at, Fig. 3 shows the resulting pulse train over a more extended period of time, a little over three times as long. in this case, we can consider the PWM output to be periodic waveform with period 1/f,

Fig. 2

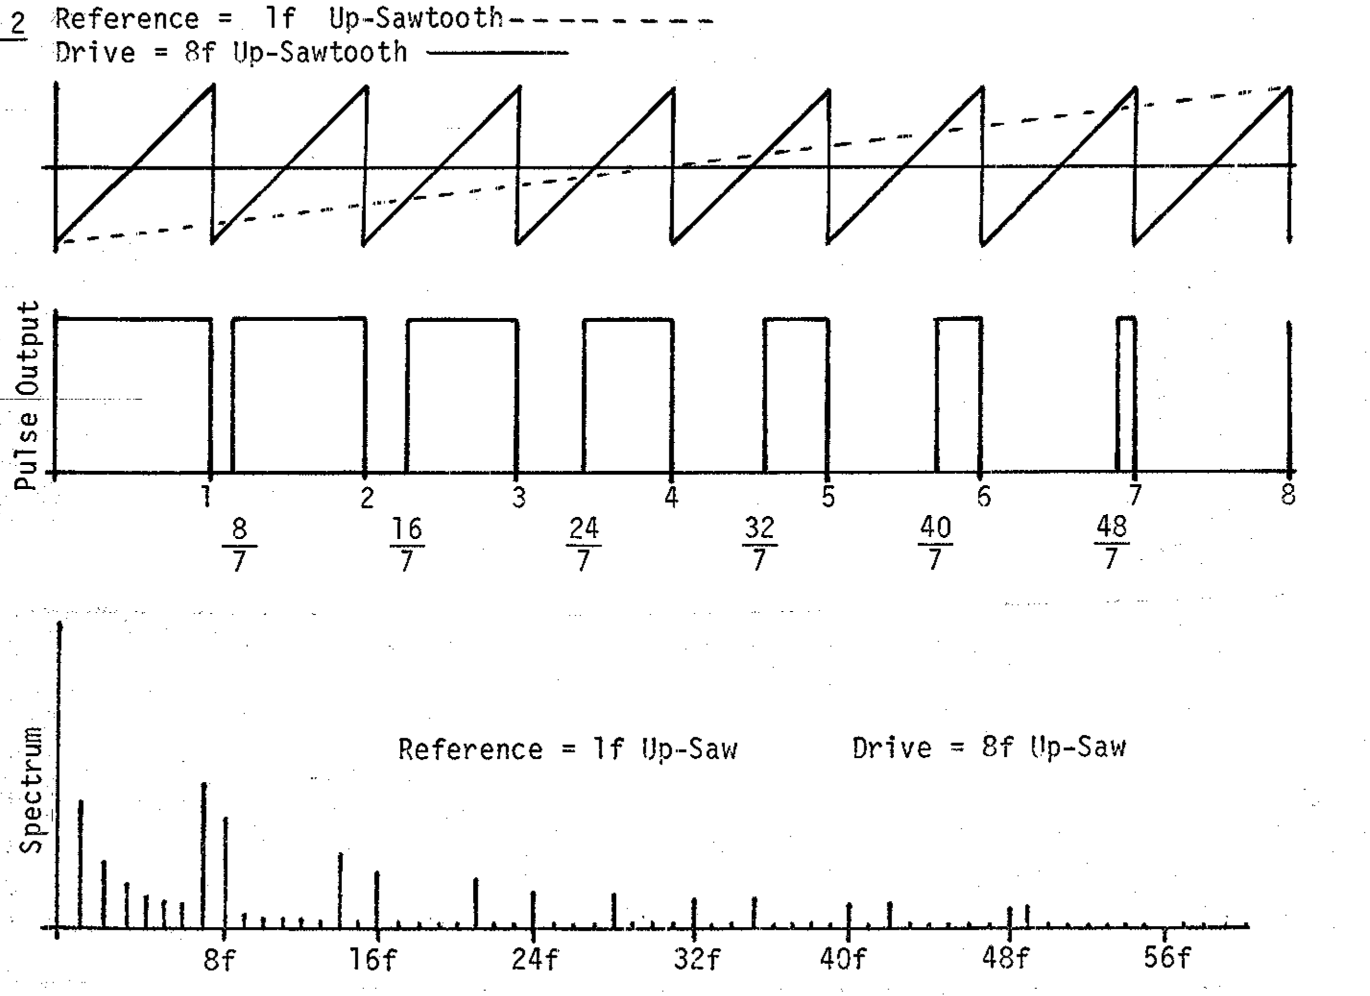

Fig. 3 Extension of middle portion of Fig. 2

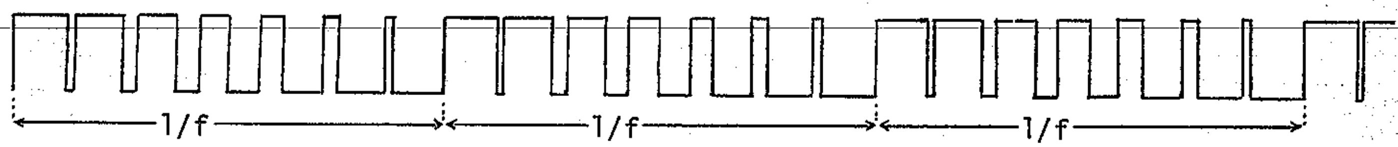

and this simplifies our spectral calculations since we can once again resort to the Fourier Series. We must keep in mind however that this applies only to the harmonic case, and later we will try to relate these results to a more general case. The calculation of the Fourier Series (FS) of the waveform of Fig. 2 or Fig. 3 is composed of three tasks. The first task is to determine the exact points in time at which the pulse rises and falls[4]. To do this, at each intersection, we write the equations of the lines representing the reference and the drive in that region, and then set the equations equal and solve. [This is of course trivial, but we can imagine that the reader would be pleased to know that the authors spent about 10 minutes fumbling around trying to get this process started!]. Once the exact points of transition are determined, it is necessary to set up the FS integrals, or an equivalent process, to determine the FS coefficients. Since the waveform here is rectangular (two levels only), integration is in fact simplified, and we can set up some simpler formulas to determine the coefficients. The details of this process are given in application note AN-160 for the interested reader. The third step is to actually run the formulas, and these come out to be summations which are easily handled by a programmable calculator or computer program. The result of such procedures for the pulse train here is shown in the bottom portion of Fig. 2[5]. Much of what we will have to say later in this report could be touched on at this point simply by discussing this spectrum. However, we must go a bit more slowly, but the reader may wish to look at the spectrum carefully at this time to see how many observations he can beat us to. For now we will just remind you that the original waveform, the drive, was of frequency 8f, and would therefore contain at best only multiples of 8f. Thus we see that the PWM process has given us nearly all harmonics of f as well, resulting in a much richer spectrum. Thus the spreading of spectral energy is one major result. Another thing to note is the significant amount of low-frequency spectral energy. Note also the difference between the larger spectral lines and the smaller ones, with fewer middle amplitude lines. Finally, note the large spectral line at 7f.

MORE DATA:

While it might be reasonable to jump in to a discussion at this point, and bring in examples as we go, here we feel it will work well to present additional data at this point. This will help the reader with some of his "what if I change....?" type of questions. The additional data given here will look at all cases of interest that involve sawtooth waves and triangle waves. Note that there is a difference between upward going sawtooth waves and downward going sawtooth waves in cases where only sawtooth waves are involved.

Fig. 4 shows the case where the drive is a up-sawtooth of frequency 8f, the same as the drive in Fig. 2, but the reference is now a downward going sawtooth. While this might at first seem like a minor difference, the pulse train (at first glance) appearing to be mainly time-reversed, the actual spectrum is quite a bit different, with spectral components originally shifted down now shifting up.

Fig. 5 shows the case where the reference is the original 1f up-sawtooth, but now the drive is replaced with an 8f triangle. Note that the spectrum here has some distinct similarities to the sawtooth cases, but also has some marked differences. In particular, note the signs of symmetry of spectral features about the harmonics of 8f.[6] Fig. 6 then gives the final case of this type, the triangle-triangle case, and you will note that this spectrum is very much different from those we have looked at so far. In particular, note that the common structure in the range from 1f to 6f that appeared in the first three cases is gone. Also the spectrum seems to have many regions where every other harmonic is somewhat weaker. Finally, for convenience, all four spectra are given on one plot in Fig. 7.

Fig. 4 Reference = 1f Down-Sawtooth - - - - - -
Drive = 8f Up-Sawtooth ————

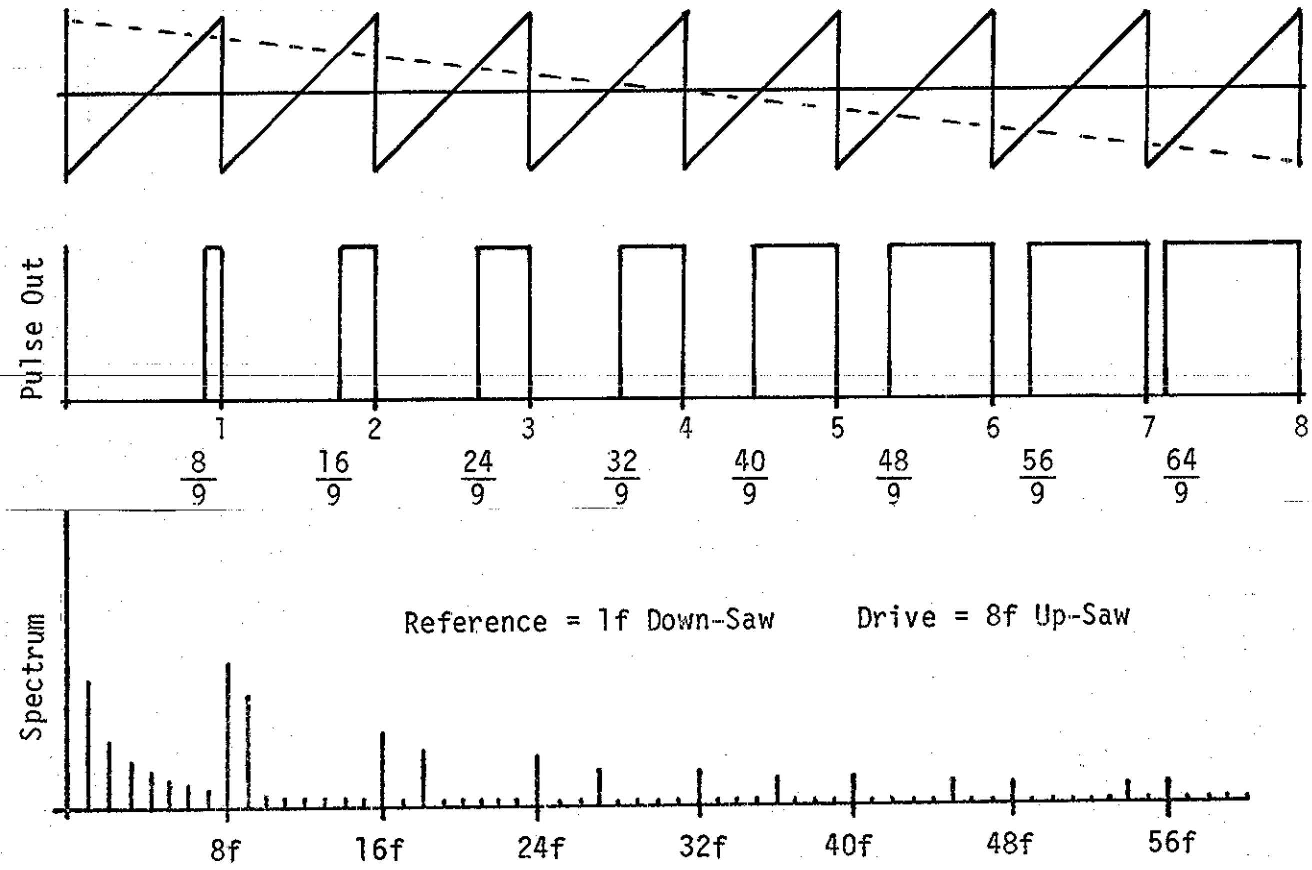

Fig. 5 Reference = 1f Up-Sawtooth - - - - - -
Drive = 8f Triangle ————

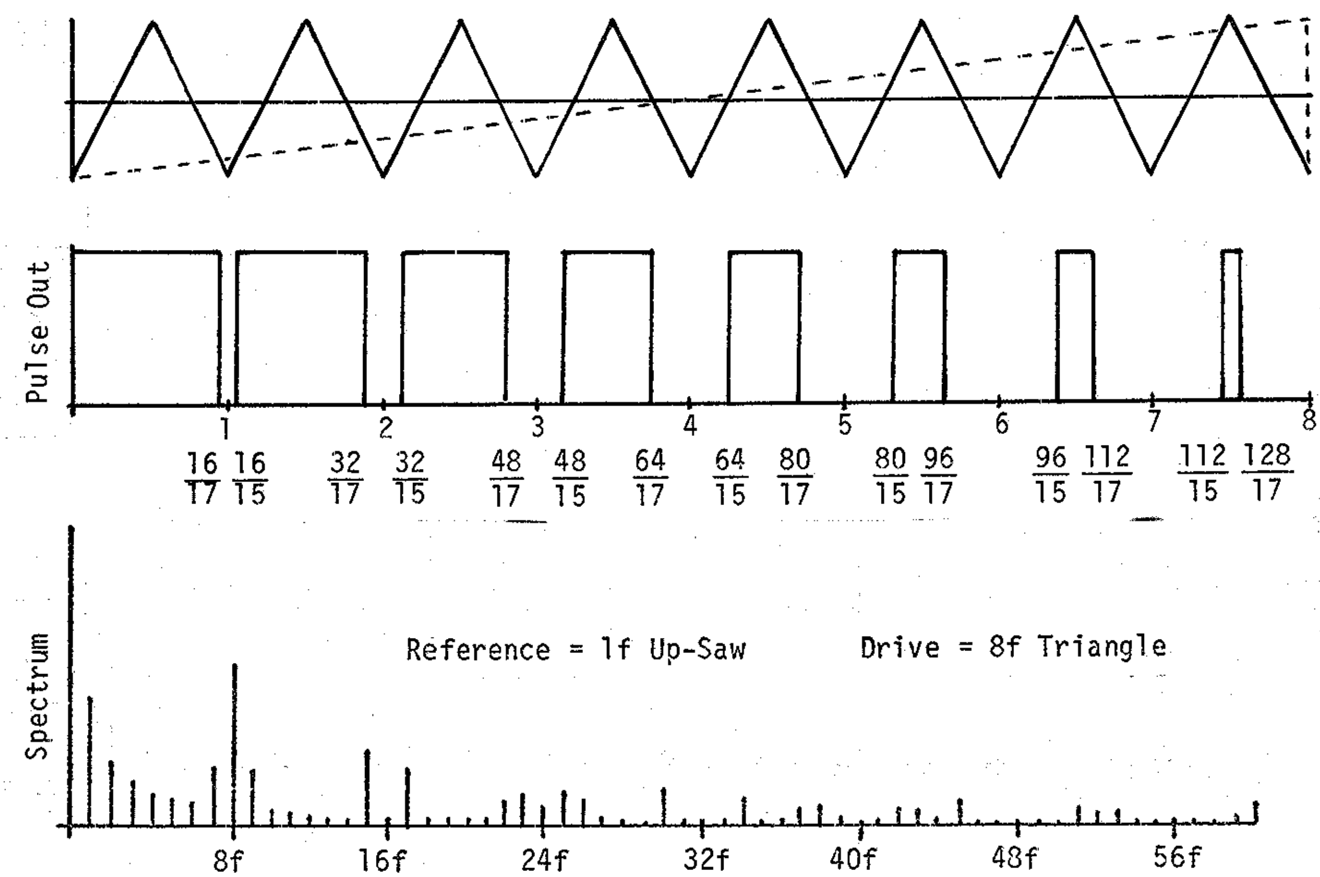

Fig. 6

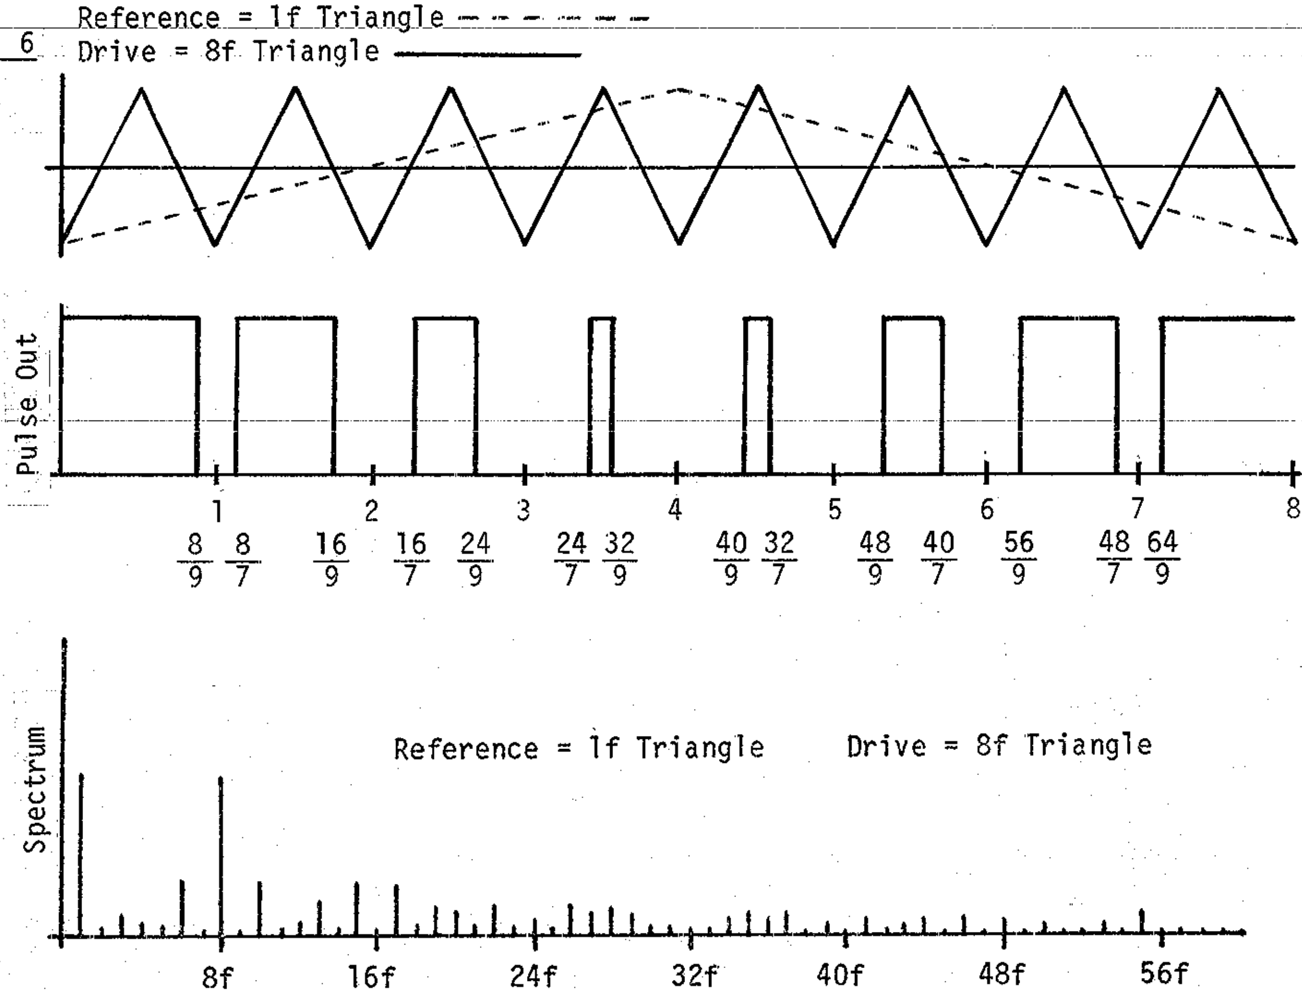

INTERPRETATIONS AND ALTERNATIVE CALCULATION METHODS:

Here we want to take a look at some alternative calculation methods for some PWM cases, and will find that they can add greatly to our understanding of the spectra calculated so far. First we will show how sawtooth-sawtooth PWM can be understood as the sum of three sawtooth waveforms. Secondly, we will look at method of forming triangle-triangle PWM from a multiplication of two known waveforms, thus making a spectral calculation based on convolution possible. Finally, we will look at the direct expansion method very briefly.

A. Saw-Saw PWM as the Sum of Saws:

A glance at the spectra for saw-saw modulation (Fig. 7a and Fig. 7b) clearly shows prominant harmonics of 7f and 8f for up-saw/up-saw, and of 8f and 9f for down-saw/up-saw. These have a 1/n-like drop off, relative to the 7f, 8f, or 9f fundamentals. In addition, both cases show a real 1/n type of behavior for the low frequency range of 1f through 6f (through 7f in the case of Fig. 7b). This is certainly a hint that the PWM waveform might be the sum of three waveforms containing all harmonics falling off as 1/n, and the sawtooth comes to mind. At this point, a bell may ring, and you might want to look back to the Musical Engineer's Handbook, pg. 2c (5) where the PWM-like nature of a sawtooth sum was noted. As believers in the relative worth of words and pictures, we now look at Fig. 8, where the PWM waveform for the up-saw/up-saw case is clearly shown to be the sum of three sawtooth waves of equal amplitude, but with frequencies of 1f, 7f, and 8f, with 1f and 7f saws inverted. The implications of this realization are quite important, as we are able to relate something quite new and poorly understood to something old and well understood - the spectrum of the saw. Note first the importance of the 1f sawtooth. This provides the harmonics from 1f-6f, 9f-13f, 15f, 17f-20f, and all others that are not harmonics of 7f or 8f, just as though

Fig. 7

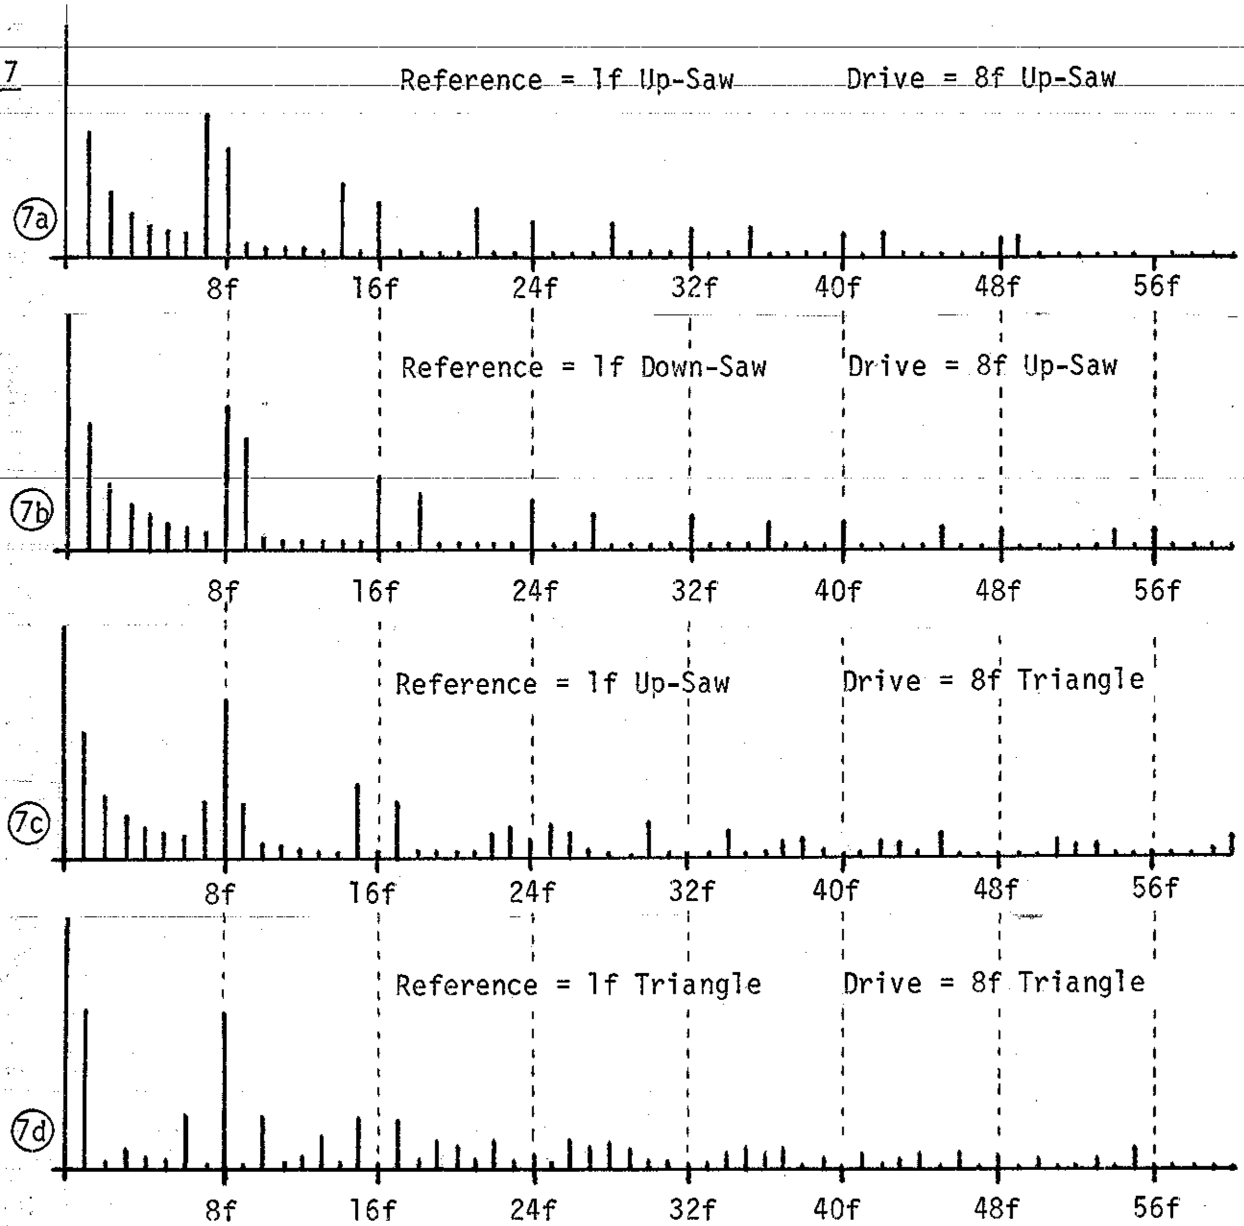

Fig. 8 PWM as the Sum of Sawtooth Waves

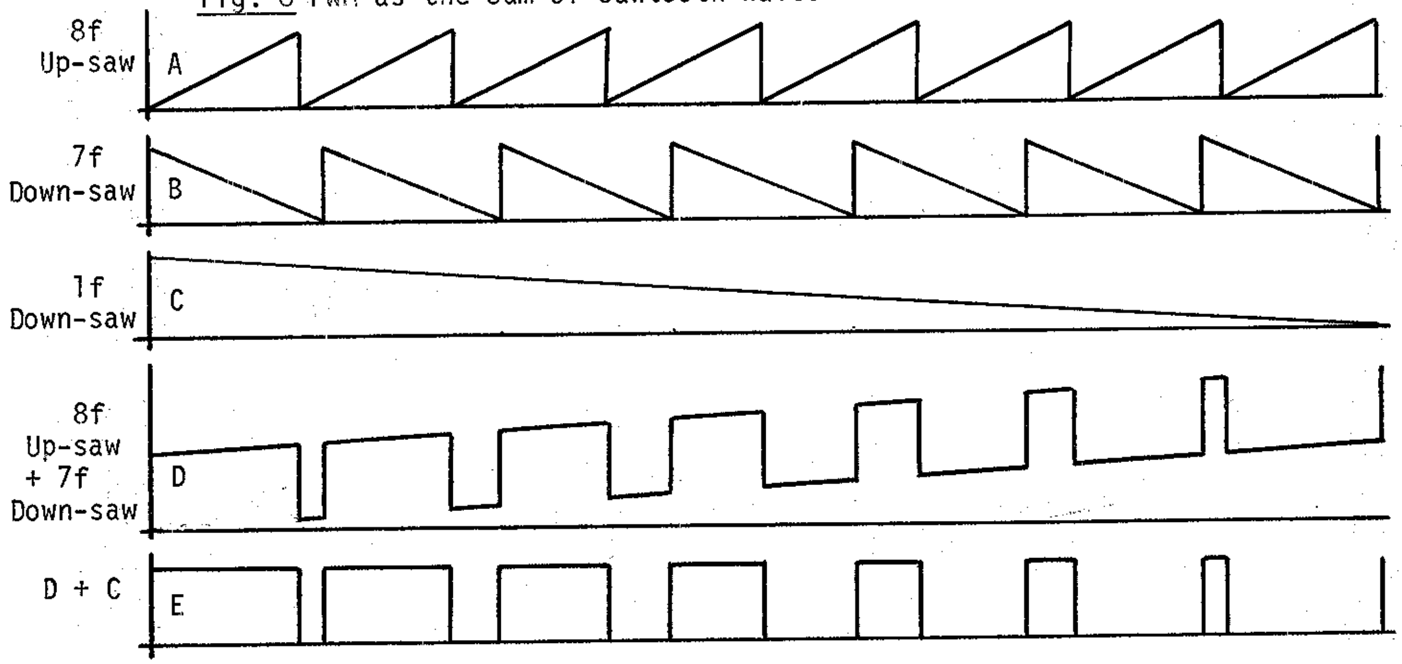

we had only this one sawtooth. Other harmonics are determined in part by the 1f sawtooth, and the 7f or 8f sawtooth waves. The component from the 1f sawtooth accounts for the departure of the harmonics of 7f and 8f from true 1/n behavior. An interesting case is the 56th harmonic of 1f, which is at 56f, which is zero. This is also the 8th harmonic of 7f, and the 7th harmonic of 8f. Thus there are three component parts of the 56f spectral line, and they must somehow cancel. Actually, this is easy to show, since we have started with all sawtooth waves in phase (or inverted) and of equal amplitude. Thus the component of the 1f sawtooth is -1/56, that of 7f is -1/8, and that of 8f is +1/7, neatly adding to zero. Note that things may change for different phase relationships between the drive and the reference, as will be shown later.

The reader can easily show the corresponding sum for the down-saw/up-saw case. It is interesting to try to arrive at a sum of waveforms for the saw-triangle case and the triangle-triangle case, but this is not simple. Perhaps some reader can suggest something reasonable. Instead, we will be looking at a different method for the triangle-triangle case.

B. Convolution of Multiplying Waveforms, Triangle-Triangle PWM:

This is a bit of a "backdoor" method, but it is interesting. We won't go into a lot of detail however, because if you know about convolution, it will be fairly obvious what we are doing, and if not, this brief description will probably give you all you need. Basically, the idea is as follows. First you find two waveforms that when multiplied together give you the PWM waveform. Since we are dealing with rectangular waveforms, the "multiplication" is really an exclusive-OR. So basically we are looking for two (fairly simple) waveforms which when exclusive ORed give you the PWM waveform. Next, the line spectra of these waveforms are determined from a Fourier Series or other suitable process. Finally, recalling that waveforms multiplied in the time domain have their spectra convolved in the frequency domain, we convolve the two spectra. The example in Fig. 9 will illustrate the idea.

To make things a bit simpler , we will change our triangle-triangle case from our original 1f/8f case to a 1f/9f case, which comes out nicer. This relationship is shown in section A of Fig. 9, and the pulse output is shown in the B part. The spectrum, determined just as we did before from the Fourier Series is in the E part of Fig. 9. Now, we have added two square waves, parts C and D of the figure. Their frequencies are 5f and 4f, and the reader can easily observe that the pulse output above them is the exclusive-OR function (one and only one) of these two. A little thought will also show that the exculsive-OR process is equivalent to multiplication if the zero level is taken to represent -1. The spectra of the square waves are well known (all odd harmonics falling off as 1/n) so if we convolve them, we get the spectrum for the PWM signal.

Convolution as a mathematical process can at times be as complex or more so than the problem it is trying to solve, so it is well to note here that the process can be straightforward. Even if you don't know what convolution is, you can understand the procedure. This is illustrated in the F and G parts of Fig. 9, which are plots of the harmonics of the square waves with the 4f square wave in the F part and the 5f square wave in the G part (shown upside-down so that we can put the baselines close together). The fact that negative frequencies are shown as well need not be of concern as it arrises from the convolution process.

The mechanics of the process (literally) is to move one spectrum with respect to the other, and where lines pass over each other, multiply the magnitudes of these lines together, and sum all such overlaps for that particular spacing. Imagine that the G spectrum strip has been cut out so that we can slide that paper strip with respect to the F strip. As shown, we show a particular case of displacement of 7f units. This displacement, 7f in this case, is the spectral line being evaluated in the convolved spectrum. Note that two strong lines are juxtaposed by this displacement (12f and 5f) as well as two weak ones (-28f and -35f), and also there are other weak ones beyond the ends of the present strips. Thus the 7f component is $(2/\pi)(2/3\pi) + (2/7\pi)(2/7\pi) +$ smaller terms = 0.153. For other components, you just slide the strips for the

Reference = 1f Triangle
Drive = 9f Triangle

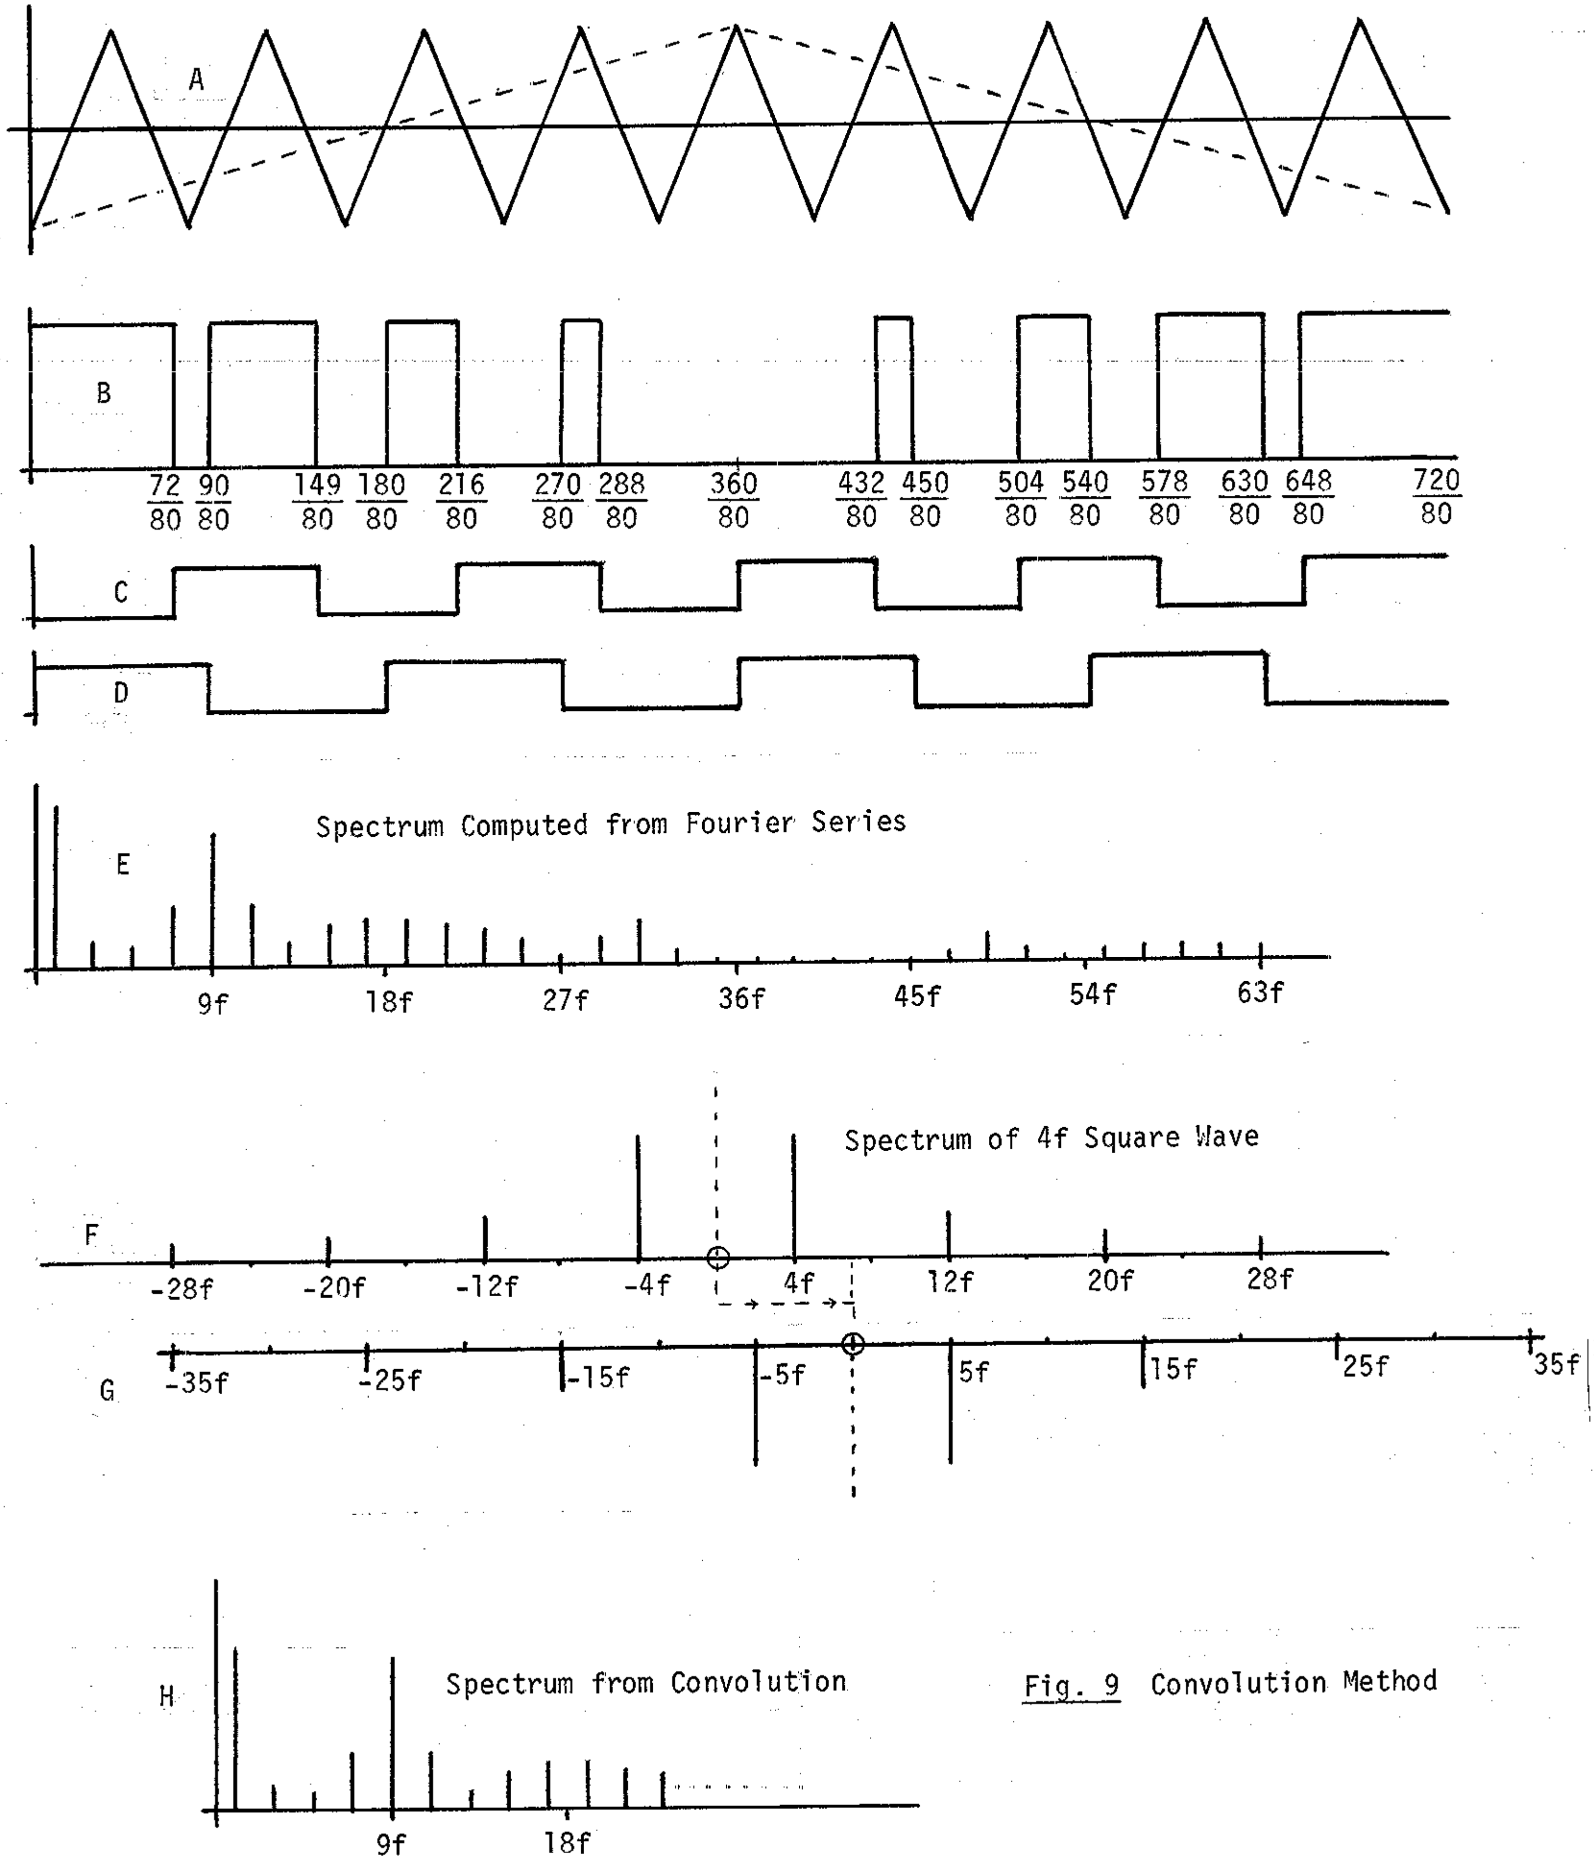

Fig. 9 Convolution Method

appropriate displacement. For example, for the 9f component, the bottom strip would move two notches to the right, lining up the major lines 4f and -5f, along with weaker juxtapositions. Note that because this is a multiplication of a line with a positive frequency by one with a negative frequency, the 4f/-5f component is of opposite sign to those that arise from like sides of the spectrum. The H part of Fig. 9 shows the result of convolution in this manner (with strips about four times as long however to pick up more of the samller terms). The spectrum in H is very close to the actual spectrum obtained from the Fourier Series in E. Any differences are due to the neglected terms off the ends. We will want to discuss things a bit more, but first we will discuss briefly the direct expansion method.

C. Direct Expansion of PWM Equation:

This method is straightforward, but of limited practical use for our purposes. It begins with the static pulse expansion as a Fourier Series as given in part 1:

$$f(x) = A/B + (2/\pi)\sum_{n=1}^{\infty}\frac{(-1)^n}{n}\,\mathrm{Sin}(n\pi A/B)\mathrm{Cos}(2n\pi x/B)$$

where A/B is the duty cycle. One form of PWM is then obtained by making A a variable such as:

$$A(y) = A_0 + A'\mathrm{Cos}\, y$$

if we plug this in above, we see that it goes in two places. First into the "DC" term A/B where it results in a "DC" level varying as Cos y. In fact, it is this term and not the summation that is usually important in communication systems that use PWM, since the PWM signal is easily demodulated by low-pass filtering for this term alone. We are more interested in the summation however, where the "harmonic coefficient" (see part 1) becomes:

$$k_n = \alpha\frac{1}{n}\,\mathrm{Sin}[n\pi A_0/B + (n\pi A'/B)\mathrm{Cos}\, y]$$

Here we see something familiar - a sinusoidal function of another sinusoidal function. We have seen this in our studies of FM, and we can feel confident that if we wish, we could use a number of trig identities and expansion of the "sine of a sine" terms in Bessel function series to get an answer. We could.

Why are we not excited about this. First, it is a lot of work compared to the somewhat easier methods we have looked at. Secondly, we can get a lot of intuitive insight into the results just be thinking about it, comparing it to the FM process, and not actually doing all the arithmetic. Thirdly, we would have to repeat the expansion for each component for a complex waveform.

D. Discussion:

We have looked at several alternative methods of calculating PWM spectra, and these are interesting in themselves, add to our intuitive understanding, and provide verification. Perhaps more importantly, our simplest method is for only the harmonic case, while the other methods are more general. Thus we can use these other methods to argue that the harmonic cases are not untypical (except for the fact that sidebands may overlap). For example, we can sum sawtooth waves that are not exact harmonics, as long as the total slope is zero, and thus arrive at a non-harmonic case of saw-saw PWM. Or we can select square waves that are not harmonics, exclusive-OR them, and arrive at a triangle-triangle case of PWM that is not harmonic. The corresponding spectral calculation methods are unchanged. Thus we can be fairly confident that the general spectral forms we have found are typical, and we can more or less put the fact that we have calculated mainly harmonic cases out of our mind.

Before going on to other things, we should make a point about sawtooth reference modulation. This type of reference puts a "click" in the PWM signal due to the sudden change of the "DC" level of the signal. The nature of this click is clear from an examination of previous figures. It is the slowest sawtooth (Line C of Fig. 8) and its spectrum is that sees in the 1f - 6f (or 7f) region of the top three spectra of Fig. 7. This click is annoying if it is subaudio, but becomes part of the timbre when its rate becomes audio.[7] We shall later see that this particular spectral region can be extremely important in the dynamic depth PWM case. So certainly this click must be considered to have positive as well as a negative aspect, depending on other factors in a particular case. Yet it is clear how to get rid of it - we simply remove the "correcting" sawtooth. Actually deriving the necessary difference frequency and generating this correcting sawtooth would be quite a chore. Instead we merely want to make the point that the click does not appear if we use Line D of Fig. 8 instead of Line E. That is, there is no click in the sum of the two sawtooth waves form Lines A and B. This is reasonable, since the linear summation of two waveforms with no low-frequency components should not generate any low-frequency ones to produce a click.[8] Nice how things really work out, isn't it?

PHASE SHIFTING OF DRIVE RELATIVE TO REFERENCE:

For a practical case, it is necessary to consider phase relationships between the reference and the drive that are not ideal such as the ones shown in Figures 2, 4, 5, 6, and 9. This is of interest first because a predetermined phase relationship may be difficult or near impossible to set up, and secondly because some change of phase can be expected in harmonic setups (near harmonic) with analog equipment. We are interested to see if the spectrum changes drastically.

We will use as our test case the up-saw up-saw PWM case of which the initial phase relationship is shown in Fig. 2. We will measure displacement of the reference in terms of a number δ, where δ runs from 0 (original) to 1 (corresponding to a displacement by one period of the drive). Fig. 10 shows a set of pulse outputs for several values of δ. You will note that most of the pulses get wider as the reference shifts to the right, except for the first pulse which gets shorter, giving up as much as all the others have gained. At $\delta = 1$, we obtain the original $\delta = 0$ case shifted by one unit of 1/8f.

Fig. 10 Pulse Output for Various Displacements of Reference

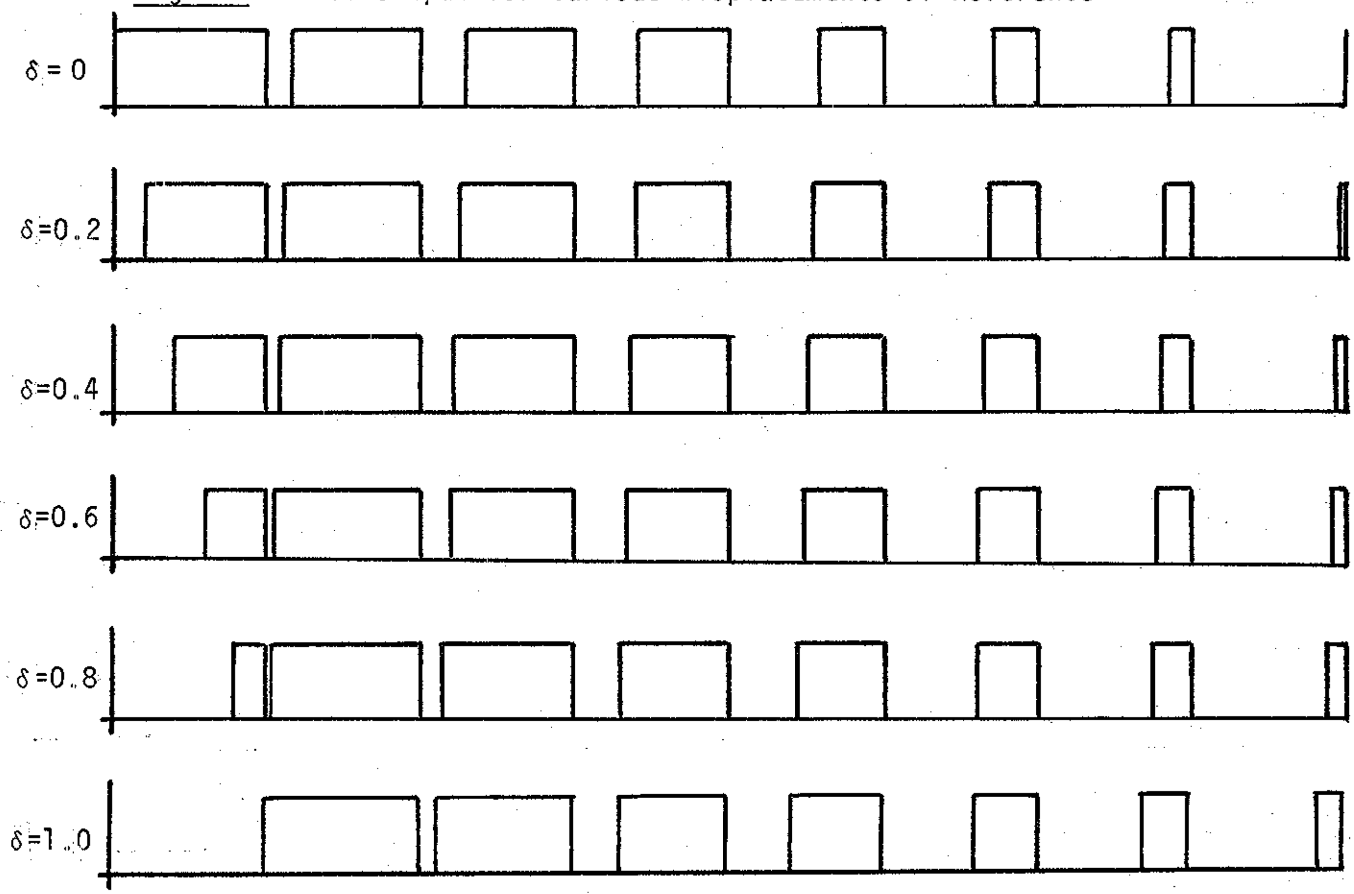

δ = displacement of reference relative to drive in units of 1/8f

It is clear that there are some substantial differences among these pulse trains, and we might wonder what sort of effect this will have on the corresponding spectra, and this is examined with the Fourier Series method. As it turns out, the shift has a suprisingly minor effect on the spectra. Most of the components stay the same, and those that do change trade energy back and forth with a close neighbor. These effects are shown in Fig. 11 and Fig. 12, which show typical data.

Fig. 11 shows the variation (or rather, the lack of variation) of coefficients c_1 and c_2 as δ goes from 0 to 1. Note that these are constant, although the actual energy shifts between its a_n (cosine) and b_n (sine) components. This constant behavior is typical of all components that are <u>not</u> harmonics of 7f or 8f. Note that the fact that the a_n and b_n components do not return to their original values at $\delta = 1$ is not a problem. What has happened is that the two sum to form a phase shift appropriate to the displacement. For example, the a_1 and b_1 components become equal in magnitude (and such that their hypotenuse is the same at all times) to obtain a 45° phase shift, appropriate to the 1/8 cycle shift of the reference.

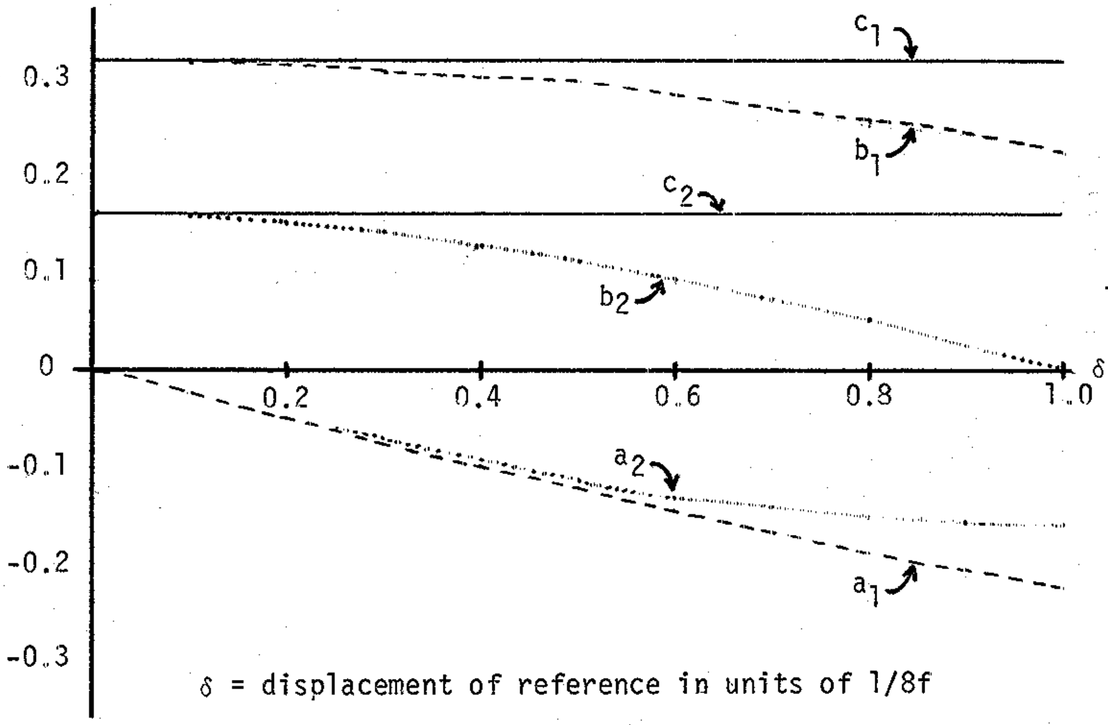

Fig. 11

Components that <u>are</u> harmonics of 7f and 8f do vary somewhat (their c_n values as well as their a_n and b_n values vary), although not drastically. Fig. 12 shows the variation for c_7, c_8, c_{14}, and c_{16}. Note the trading of energy between neighboring strong harmonics.

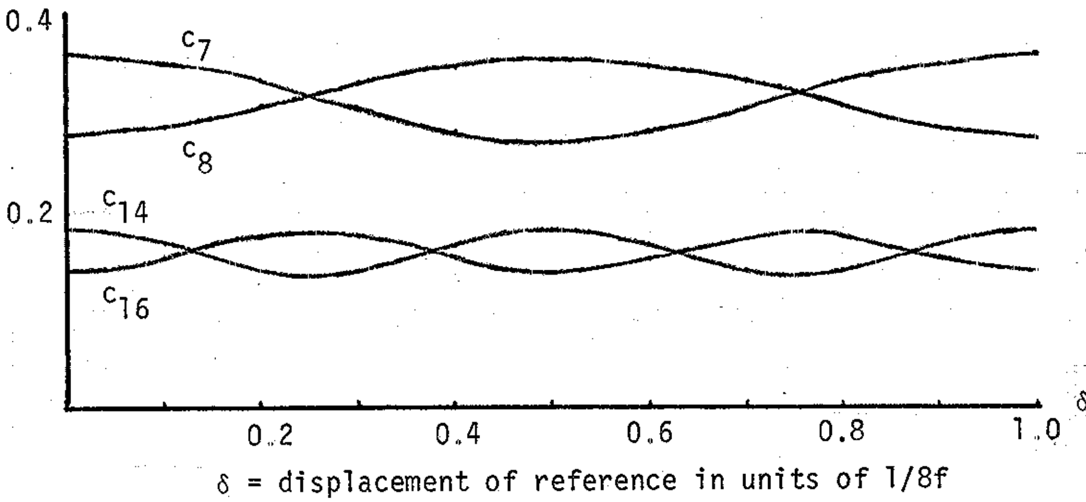

Fig. 12

DYNAMIC DEPTH PWM:

We also need to look at the PWM case where the amplitude of the reference varies with respect to the amplitude of the drive. For the moment we will assume that the reference varies from zero up to the amplitude of the drive while the drive voltage remains at constant amplitude.[9] We will use the up-saw/up-saw case. The results of some test runs are shown in Fig. 13.

The top line A shows zero modulation depth, or zero-crossing of a sawtooth wave, which results in a square wave of frequency 8f, and we find the expected spectrum of odd harmonics falling off as 1/n. The bottom line E is the full 100% modulation case, just what we saw in Fig. 2, and the difference in spectral density is quite apparent. The three inbetween cases (B, C, and D) give some idea as to how the spectrum evolves. In hardware terms, this is just what you get if you turn up the modulation depth control on a VCO. Some interesting things to note are the initial spreading of sidebands about original components (B, C), the filling in of significant components at formerly missing harmonics of 8f (C, D), and the shifting of energy among close components (D, E). However, one of the most striking things from this point of view is the appearance of the harmonics in the range 1f to 7f where none are originally. This is the "click" portion of the spectrum, and we can exploit the way in which it appears.

First however we might note that it is easy for us to set up dynamic depth PWM in a voltage-controlled manner, and an acceptable patch is shown in Fig. 14. The control element is shown as a balanced modulator, but a VCA could also be used for what we are considering at the present time. VCO-1 is the oscillator being modulated (PWM) so can be considered the drive and comparator, while VCO-2 provides the reference, with the depth of modulation (how far down Fig. 13) controlled by the envelope applied to the balanced modulator. This is an easy thing to try yourself, and the results, while interesting, are a bit harsh. Probably this is because we go from a hollow type of sound (square wave) to a very rich sound (100% modulation), while at the same time the "fundamental" jumps downward from 8f to 1f, giving us the impression that our perception was somehow all wrong in the first place. Try it.

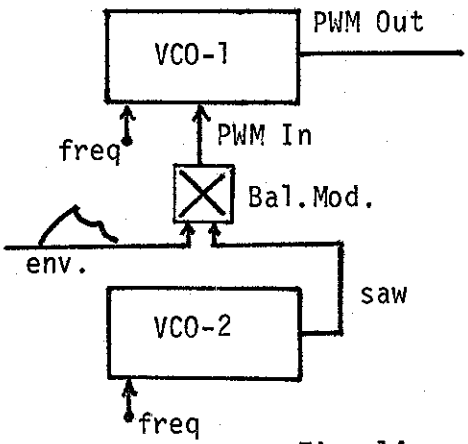

Fig. 14

One easy way around this difficulty (or perhaps we should say, one way of achieving more traditional type sounds) is to simply add a filter to the patch and listen to a more restricted portion of the resulting PWM spectrum. The patch is shown in Fig. 15, and it is not unreasonable to assume for the moment that any combination of parameters we might consider could be tried. One interesting case however is where VCO-1 is set to a very high frequency, even above the audio range, and the VCF is set in a low-pass mode to filter this out completely. In the standby state therefore, nothing comes through. Now, if we set VCO-2 to an audio frequency, and apply its signal to the PWM of VCO-1 (by raising the envelope #1 in Fig. 15 for example), then low-frequency sidebands will come up from zero frequency and will come through the filter. Thus we have a rather strange case where envelope #1 controls spectral shape and amplitude of the

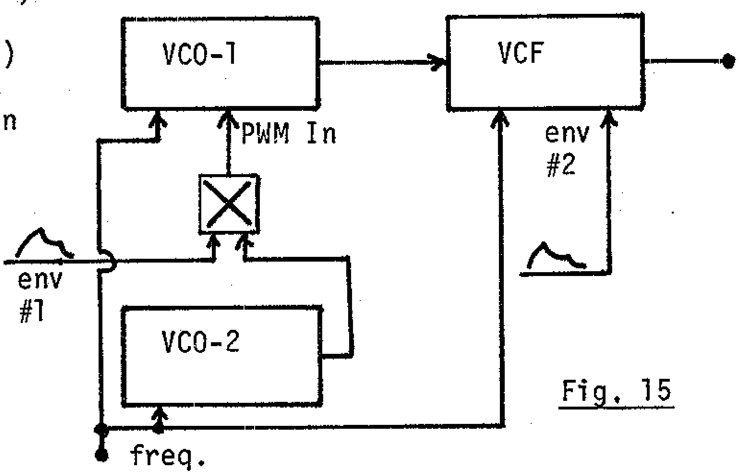

Fig. 15

Fig. 13

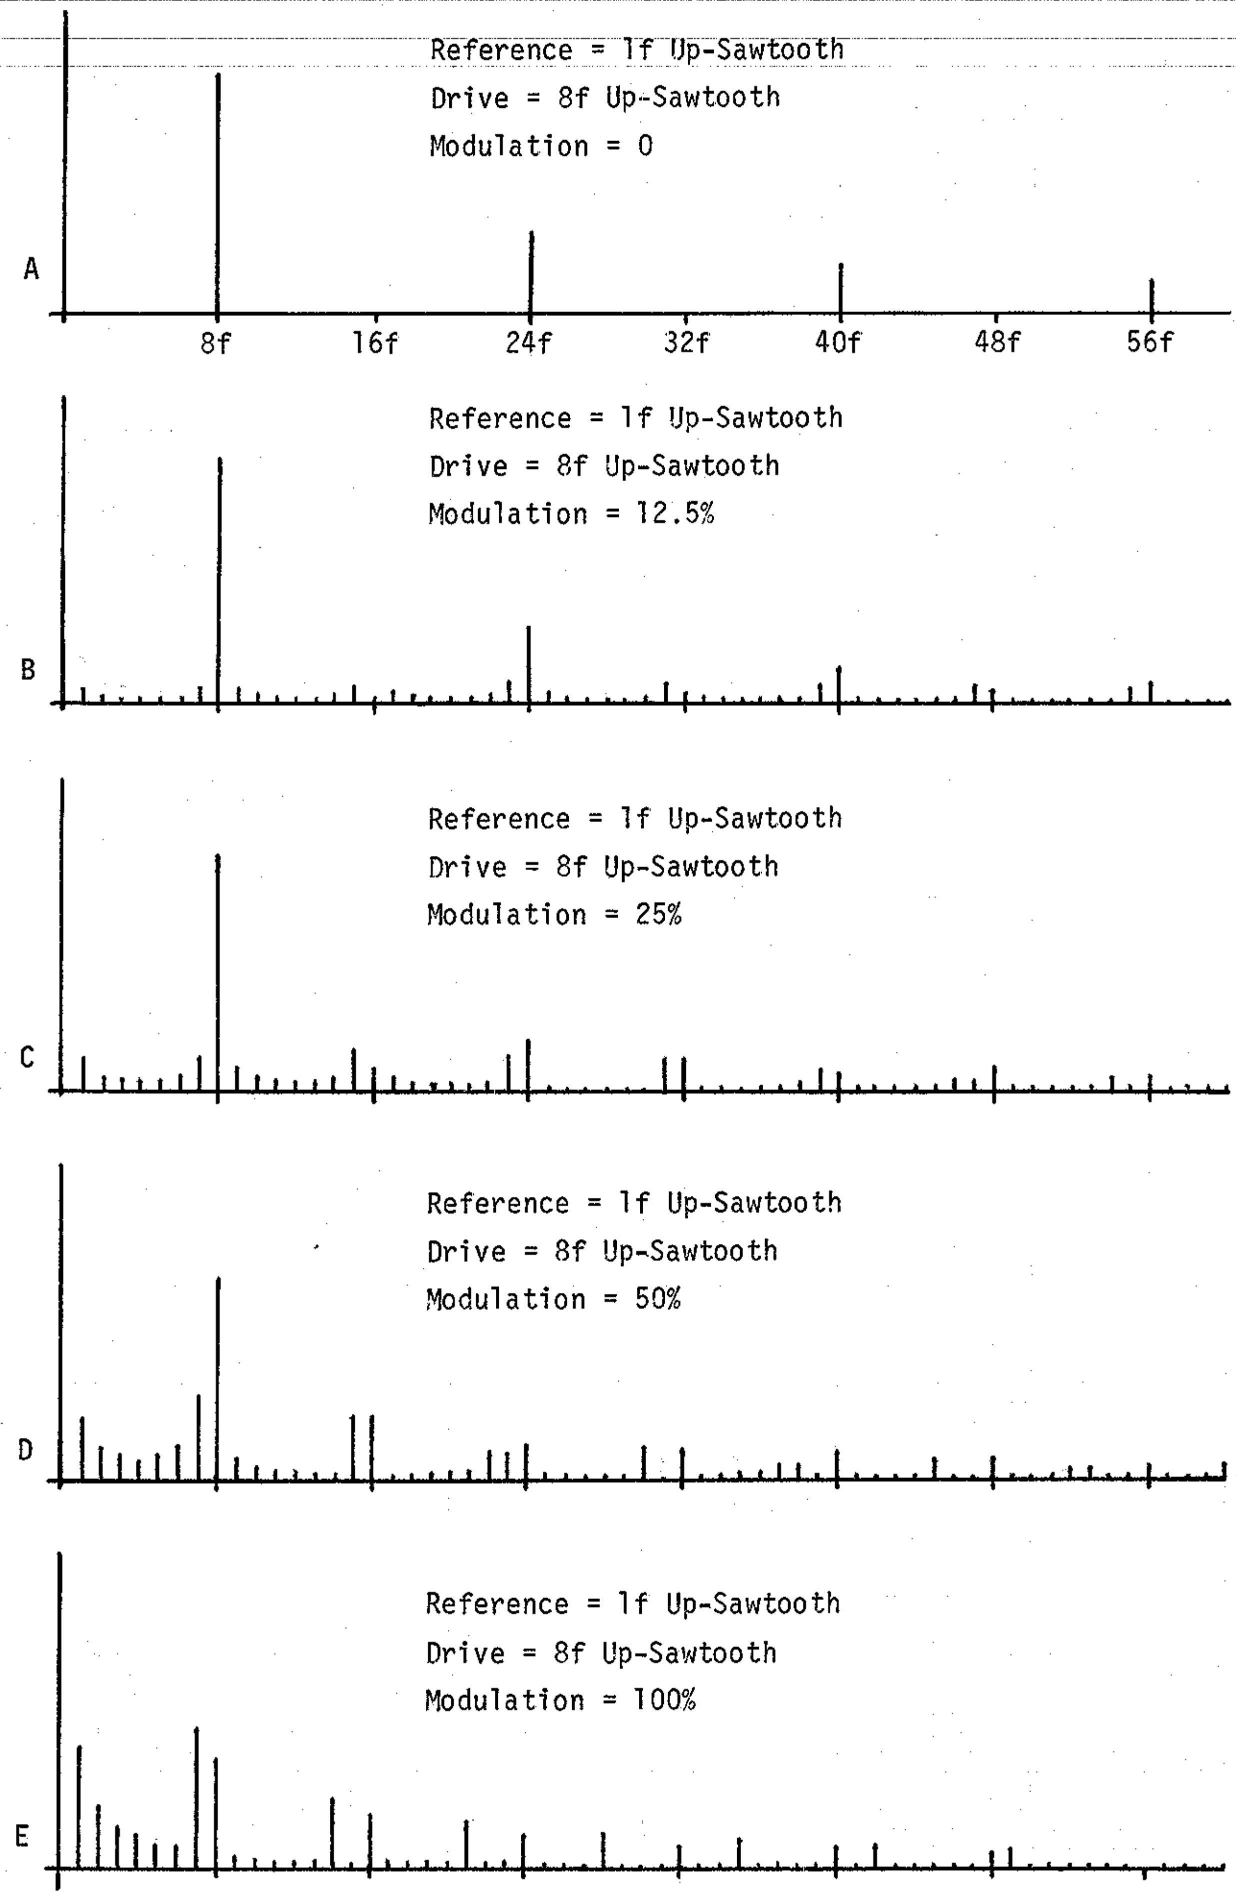
Reference = 1f Up-Sawtooth
Drive = 8f Up-Sawtooth
Modulation = 0
A
8f
16f
24f
32f
40f
48f
56f
Reference = 1f Up-Sawtooth
Drive = 8f Up-Sawtooth
Modulation = 12.5%
B
Reference = 1f Up-Sawtooth
Drive = 8f Up-Sawtooth
Modulation = 25%
C
Reference = 1f Up-Sawtooth
Drive = 8f Up-Sawtooth
Modulation = 50%
D
Reference = 1f Up-Sawtooth
Drive = 8f Up-Sawtooth
Modulation = 100%
E

output all at one time. The modulation process actually creates harmonics that can pass through the filter. The sounds produced in this way are somewhat similar to those produced for example with a fixed pulse width and a moving filter (i.e., traditional subtractive synthesis).

Things become a lot more interesting when we consider the full potential of the patch of Fig. 15. First, we can use envelope #2 as well, to add a second control of harmonic structure. Thus the VCF is used in a normal manner to control a spectrum that is already dynamic through the PWM process. This is one step above the normal additive synthesis process. However, the really interesting thing is that VCO-1 is really acting as a "carrier" for the modulation provided by VCO-2. Thus, it's frequency is not important to the pitch, while VCO-2 provides the pitch. Now if we lower the frequency of VCO-1, the large sidebands around its frequency will start to get through the filter. We can also at the same time send the filter up further to get these if we wish. Now these upper sidebands are interesting in that they will be non-harmonics in general, and through a phase shift type effect due to frequency adjustment errors, they will be somewhat modulated in amplitude (see Fig. 12). Also, since we are using low-pass filtering, these upper components may well be somewhat attenuated, even though they start out strong. A better idea of what is going on can be obtained by studying Fig. 16.

Fig. 16

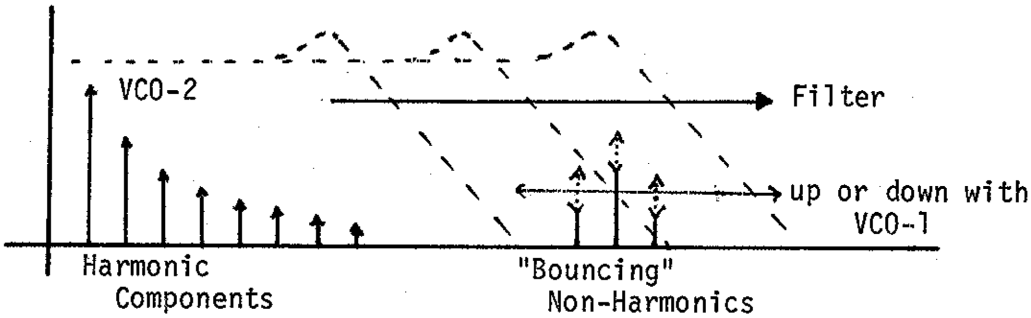

In Fig. 16, we assume that the patch of Fig. 15 is being used, and that the note being played is "attacking." We see at the bottom that the harmonic components are coming up (compare with the 1f through 6f components of Fig. 13). The filter passes over the upper harmonic components and reaches a set of "bouncing non-harmonics." These non-harmonics correspond to approximately 7f, 8f, and 9f in the harmonic case of Fig. 13. The spacing of these harmonics is the same as the modulating frequency supplied by VCO-2 (thus the same as the pitch) but the exact placement relative to zero frequency is determined by VCO-1. The result is that non-harmonics in the upper range can be added to a relatively strong harmonic base, resulting in some very interesting bell-like and other unique sounds. Many sounds are similar to those obtined with the FM technique. Variations on this technique could involve such things as small amounts of FM for either or both of the VCO's.

Another interesting case is to use Fig. 15 in a high-pass mode, and set the filter frequency to exclude the low frequency region. This avoids the problem of having the spectral balance suddenly shift low as the modulation depth is increased. However, this case still gives relatively untraditional sounds, for good or bad according to your point of view. The sounds are much harsher than in the low-pass case.

Before closing out this part, we should mention that we can consider a modulation through zero depth for this system, which is easily implemented by replacing envelope #1 with a bipolar waveform, something that is easily done since we are using a balanced modulator. While this has not been totally explored, initial tests show that it is not a spectacular effect. Since there are three frequency variables in such a case however, ther may well be room for additional exploration. Also, this is just one of several possible interpretations of "through-zero PWM" and additional results will be discussed in a later part. We would be interested in hearing from readers who have perhaps used these patches in the past, most likely on a trial and error basis, or from readers who use them now and find them useful.

FOOTNOTES:

[1] This is related to the minimum audible frequency of approximately 15 Hz. Nearly all slow processes will blur or become difficult to follow around this frequency.

[2] Here we use the word "activate" instead of "animate." By activate we mean a process that perturbes a static waveform. Typically an activated waveform will be an order of magnitude more difficult (or more) to properly display and trigger on a scope, relative to a static one. An animated waveform is one which is virtually impossible to trigger and display, and one for which we claim the ear will have more inherent interest.

[3] We will be ignoring all DC levels in waveforms, and simply using a unipolar or bipolar representation according to which is more useful, or more traditional.

[4] That is, the points such as 8/7, 16/7, etc. in Fig. 2, and also the obvious one such as 1, 2, 3,... etc. Note that the fractional numbers are shown down lower to avoid too much crowding.

[5] What we are actually plotting is the total coefficient $c_n = \sqrt{a_n^2 + b_n^2}$, where a_n and b_n are the cosine and sine components respectively. We are also ignoring the DC component, which is 0.5 in nearly all cases (except for the triangle where it can vary a small amount).

[6] It is tempting to suggest that this summetry is a result of the fact that the triangle can be thought of as composed of portions of an up-saw and portions of a down-saw. The up-saw produces components on one side, while the down-saw produces components on the other side. This could probably be demonstrated in a number of ways, but we have not tried this yet.

[7] It is interesting that this "click" does not become the pitch, unless the click rate is quite high. Instead, the pitch seems to be determined by the main spectral lines around 7f, 8f, or 9f. We have not done an accurate test on this pitch shift, but the pitch is down in the case of the up-saw / up-saw (Fig. 2) and up in the case of the down-saw / up-saw (Fig. 4), in qualitative agreement with the corresponding spectra. This can also be understood by observing that whatever "happens" will "happen" a bit later each "cycle" in the case of the up-saw reference, while it will "happen" a bit earlier each "cycle" in the case of the down-saw reference. By "happen" we could mean the occurance of a rising or falling edges of the pulse output. Another useful interpretation of "happening" is the occurance of a "center or mass" of the pulses, which shifts to the right each cycle for Fig. 2 (lower pitch) and to the left each cycle for Fig. 4 (higher pitch). This displacement of the center of mass then "snaps back" at the end with a frequency corresponding to the reference (thus to the "click").

[8] This comment is basically just a definition of a linear process, and the simple addition of two signals is such a linear process. One could suggest however that it is possible to achieve a lower-frequency periodic sensation even with a linear process. The well known "beat" in the sum of two sinewaves is (or can be) such a case.

[9] One should not assume that modulation in excess of 100% is not allowed in music synthesis in general, and in PWM in particular. In PWM, modulation exceeding 100% corresponds to a periodic drop-out, annoying at low frequency, but part of the timbre at high frequency. Note also that this effect is similar to a resonator behavior where a particular response is excited, decays away, and then can be re-excited later. Thus PWM in excess of 100% may fit in with what we have called a "generalized resonator" in the past.

— — — — — — — — — — — —

In later parts of this series, we will be looking at other aspects of PWM, and at various interpretations of through-zero PWM in particular. Any general comments, or specific comments from readers who follow up on some of the ideas presented here will as always, be welcome.

AN "ARTICULATING" CONTOUR GENERATOR:

-by Bernie Hutchins

INTRODUCTION: This device lies somewhere between a sequencer and an envelope generator. We have decided to call it an "Articulating Contour Generator" (ACG) because it will generally be used to add a sort of character or "signature" to a tones attack (thus an articulation) and because it is not a traditional envelope generator, we call it a contour generator. It is actually just a counter, clocked by its own clock, and outputting a sequence of +1's, 0's, and -1's, user selected by switches. This sequence is then double integrated for smootening purposes. Typically it will be used to give a small pitch bend at the start of a note, or to manipulate a filter frequency, or to mix in a small bit of noise.

THEORY OF OPERATION:

The idea of using a control-voltage sequencer as an envelope generator is not new, but it is seldom used. Probably the reason is that if the sequencer has enough steps for adequate resolution, it has too many steps to adjust to be worth the trouble. In the present design, there are only eight steps, each of which has a weight of only +1, 0, or -1. Here, smooth envelopes and a variety of envelopes are obtained through the use of several integrators. In an attempt to show the variety of contours you can expect from such a device, we will show a number of calculated curves. The reader should be aware that these sketches are for arbitrarily chosen time constants and gains, since here we need only show the shapes. Also, keep in mind that the integrals in general will end at non-zero values. The first integral will hold its last value when the sequence ends, but the second one would keep on going, unless the first integral happened to end at zero. Thus, it is desirable in the actual hardware to add a means of returning the integrators to their initial zero values. This is done with a controlled reset, and in part by the natural damping that is a part of any non-ideal integrator design. We do not show these resets in our sketches, but the reader can imagine them returning to zero by some sort of exponential decay process.

Fig. 1 shows our first example for the sequence +1, 0, +1, -1, -1, 0, -1, +1. The first integral is composed of straight line segments, but the second one is much smoother, and with our supposed exponential decay added at the end, the similarity of this to a more standard (ADSR) envelope is notable. In an effort to show the variety of contours

Fig. 1

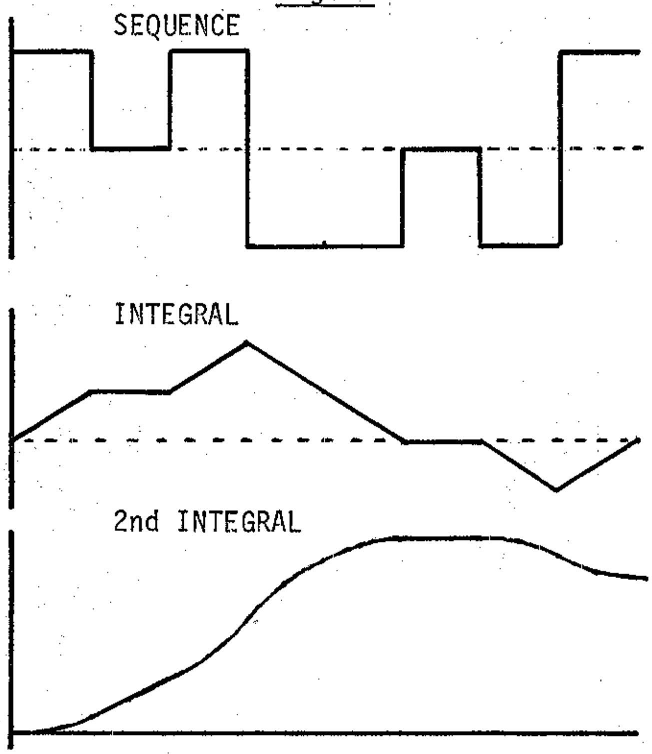

Fig. 2

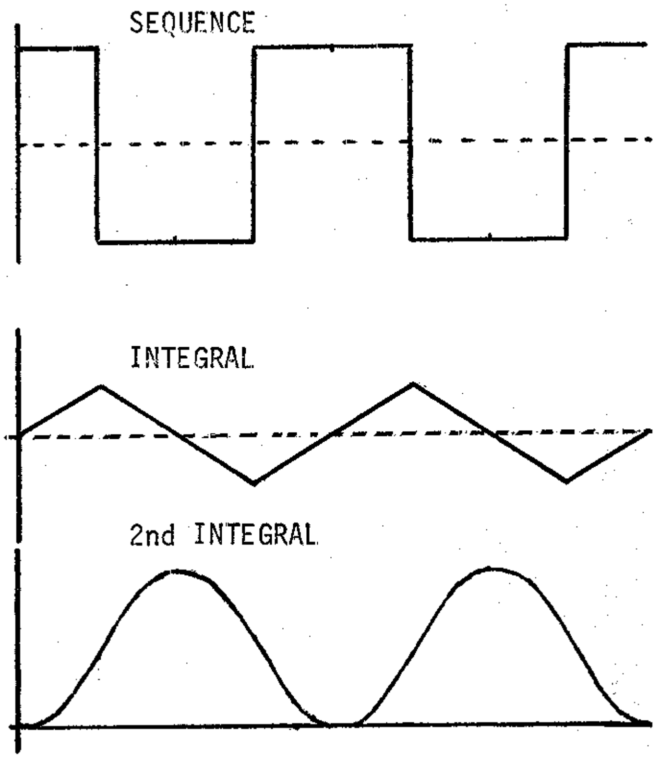

that are possible, we show in Fig. 2 the special case where both integrals end at zero (the only such case, except for half the sequence and displacements of this half). Note that the 2nd integral is a pair of parabolic (sinusoidal-like) humps. Note however that if we change the phasing of the input sequence, so that we obtain the sequence of Fig. 3, we get quite a different contour. Thus we can begin to see that even with just the sequence selection, we can get a good variety of contours.

Another feature that is easy to add to the circuit is a DC offset for the sequence. This has the effect of bending the contour up or down, as can be seen in the example of Fig. 4. Fig. 5 shows another set of contours for a sequence which has zero for its initial step, and thus the offset is more important at the beginning that it was in the case of Fig. 4.

Fig. 6 shows the case of a constant sequence, resulting in a sharp parabolic rise in the second integral. If we imagine this followed by our supposed exponential decay, it is clear that we achieve an AD type of envelope where the A as well as the D is concave. In the past, we needed a feedback technique with a voltage-controlled envelope generator to achieve this particular shape.

Fig. 7 shows the actual circuit of the ACG, and the design is relatively conventional, consisting of convenitonal sequence generating digital logic followed by standard summers and integrators. The only added circuitry is the reset structure based around the 4016 analog switch. Note also that IC-2, an op-amp, is used as a NOR gate to avoid the need for a larger IC package 3/4 of which would be wasted.

The basis clock of the system is formed around IC-6e and IC-6f. In the actual run mode, the clock signal passes through NOR gate IC-2 to the counter. The counter used here is a CMOS type 4024 which is convenient

Fig. 3

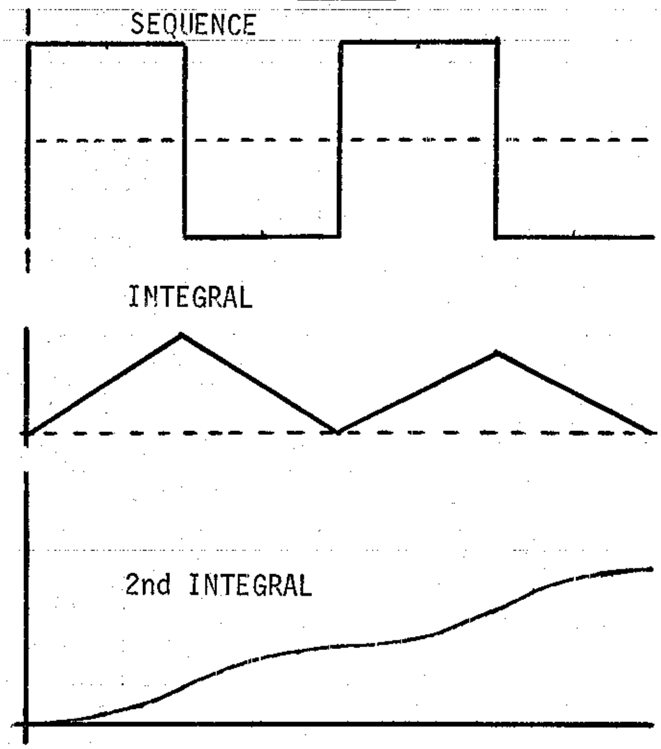

Fig. 4

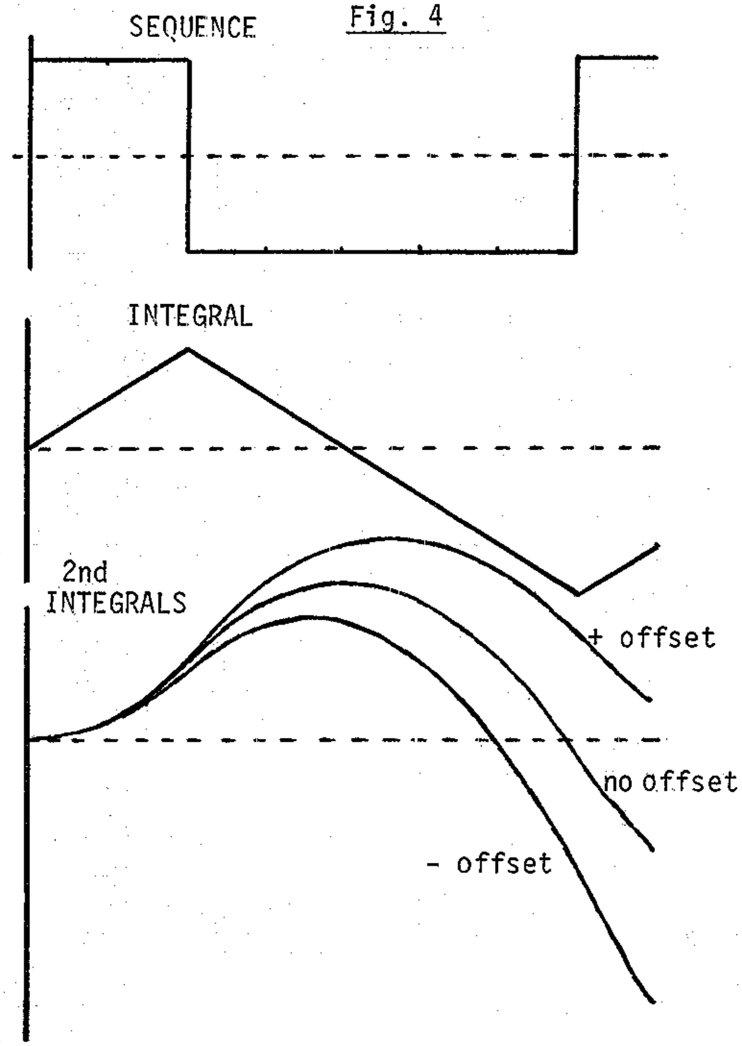

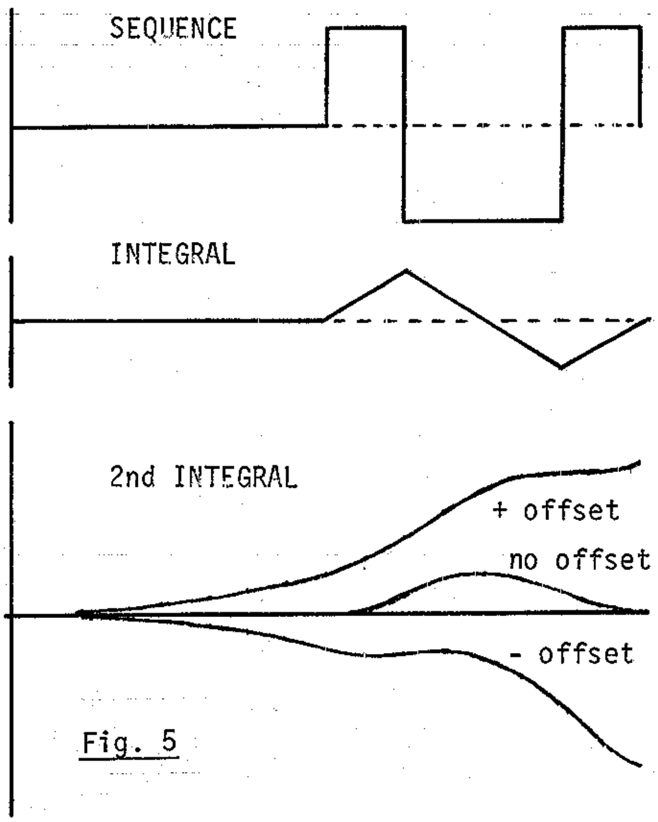

Fig. 5

Fig. 6

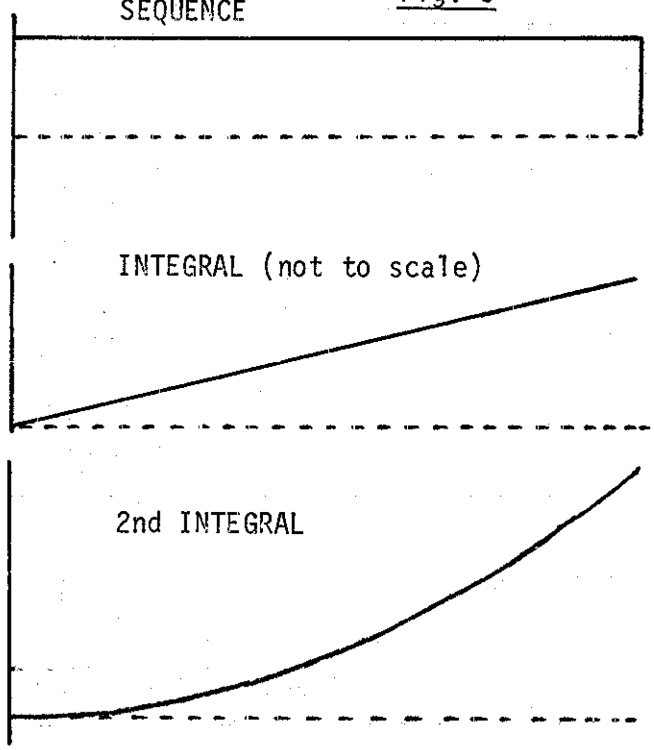

for use as a four-stage binary counter (three additional stages are not used). Other counters such as the 74C93 could probably be used. IC-3 drives the decoder IC-4, which is a 74C154 four line (A,B,C,D) to 16 line decoder. The decoded output corresponds to counts 0, 1, 2,....15, corresponding to pins 1, 2, 3,..11, 13,..17). We are using only counts 0 through 9 (pins 1 - 10) in our design. Note that count 9 (pin 10) is inverted through IC-6c and fed back to NOR gate IC-2, blocking the clock. Thus when the counter reaches count 9, it is stopped by this feedback loop. It thus produces the desired sequence and then stops - until it is reset by trigger stage IC-7, at which time the sequence runs again. Note that the sequence is formed with output counts 1 through 8 (pins 2 - 9), rather than with counts 0 through 7 (pins 1 - 8). The reason for this is that once the counter is reset, it will advance as soon as a clock pulse arrives, and random chance says that this first one could occur at any time, so the length of the zero state is irregular. To avoid this, we start on one and go one extra on the end. Note that a single positive going pulse is obtained through the use of inverters IC-5a through 5f, and IC-6a and 6b, for stages 1 - 8 respectively.

Switches S-1 through S-8 are used to feed the positive going pulse to either an inverting summer or a non-inverting one (IC-9 or IC-8 respectively). The summed sequence is available at the output of IC-9, and this in turn goes on to the two integration stages around IC-10 and IC-11. These are conventional inverting integrators with the 18M resistors serving to provide DC stability, and the analog switched resistor forming the discharge circuit. This discharge is controlled by IC-1, which is just a level shifted version (-6 to +6) of the output of IC-6c (0 to +15). Thus when the counter is in stage 9 (the hold and wait for a resetting trigger mode), the integrators are in discharge mode. For the second integration (IC-11) this discharge resistor is a 2M pot, and this is quite important since it is generally a "decay rate" function. The corresponding resistor in IC-10 is fixed at 100k, but could be a pot if extensive use is made of single integration. Note that IC-12, the analog switch is run on a ±6 supply obtained locally from the zener diode circuits shown at the lower right.

I used all LF13741's for the op-amps in this circuit. IC-10 and IC-11 should be a low bias current type such as the LF13741 or the LF351, etc., but the other op-amps could be just about any type on hand. The inverters used were 74C04 or 4069 hex types.

Fig. 7

Circuit Diagram of the Articulating Contour Generator (ACG)

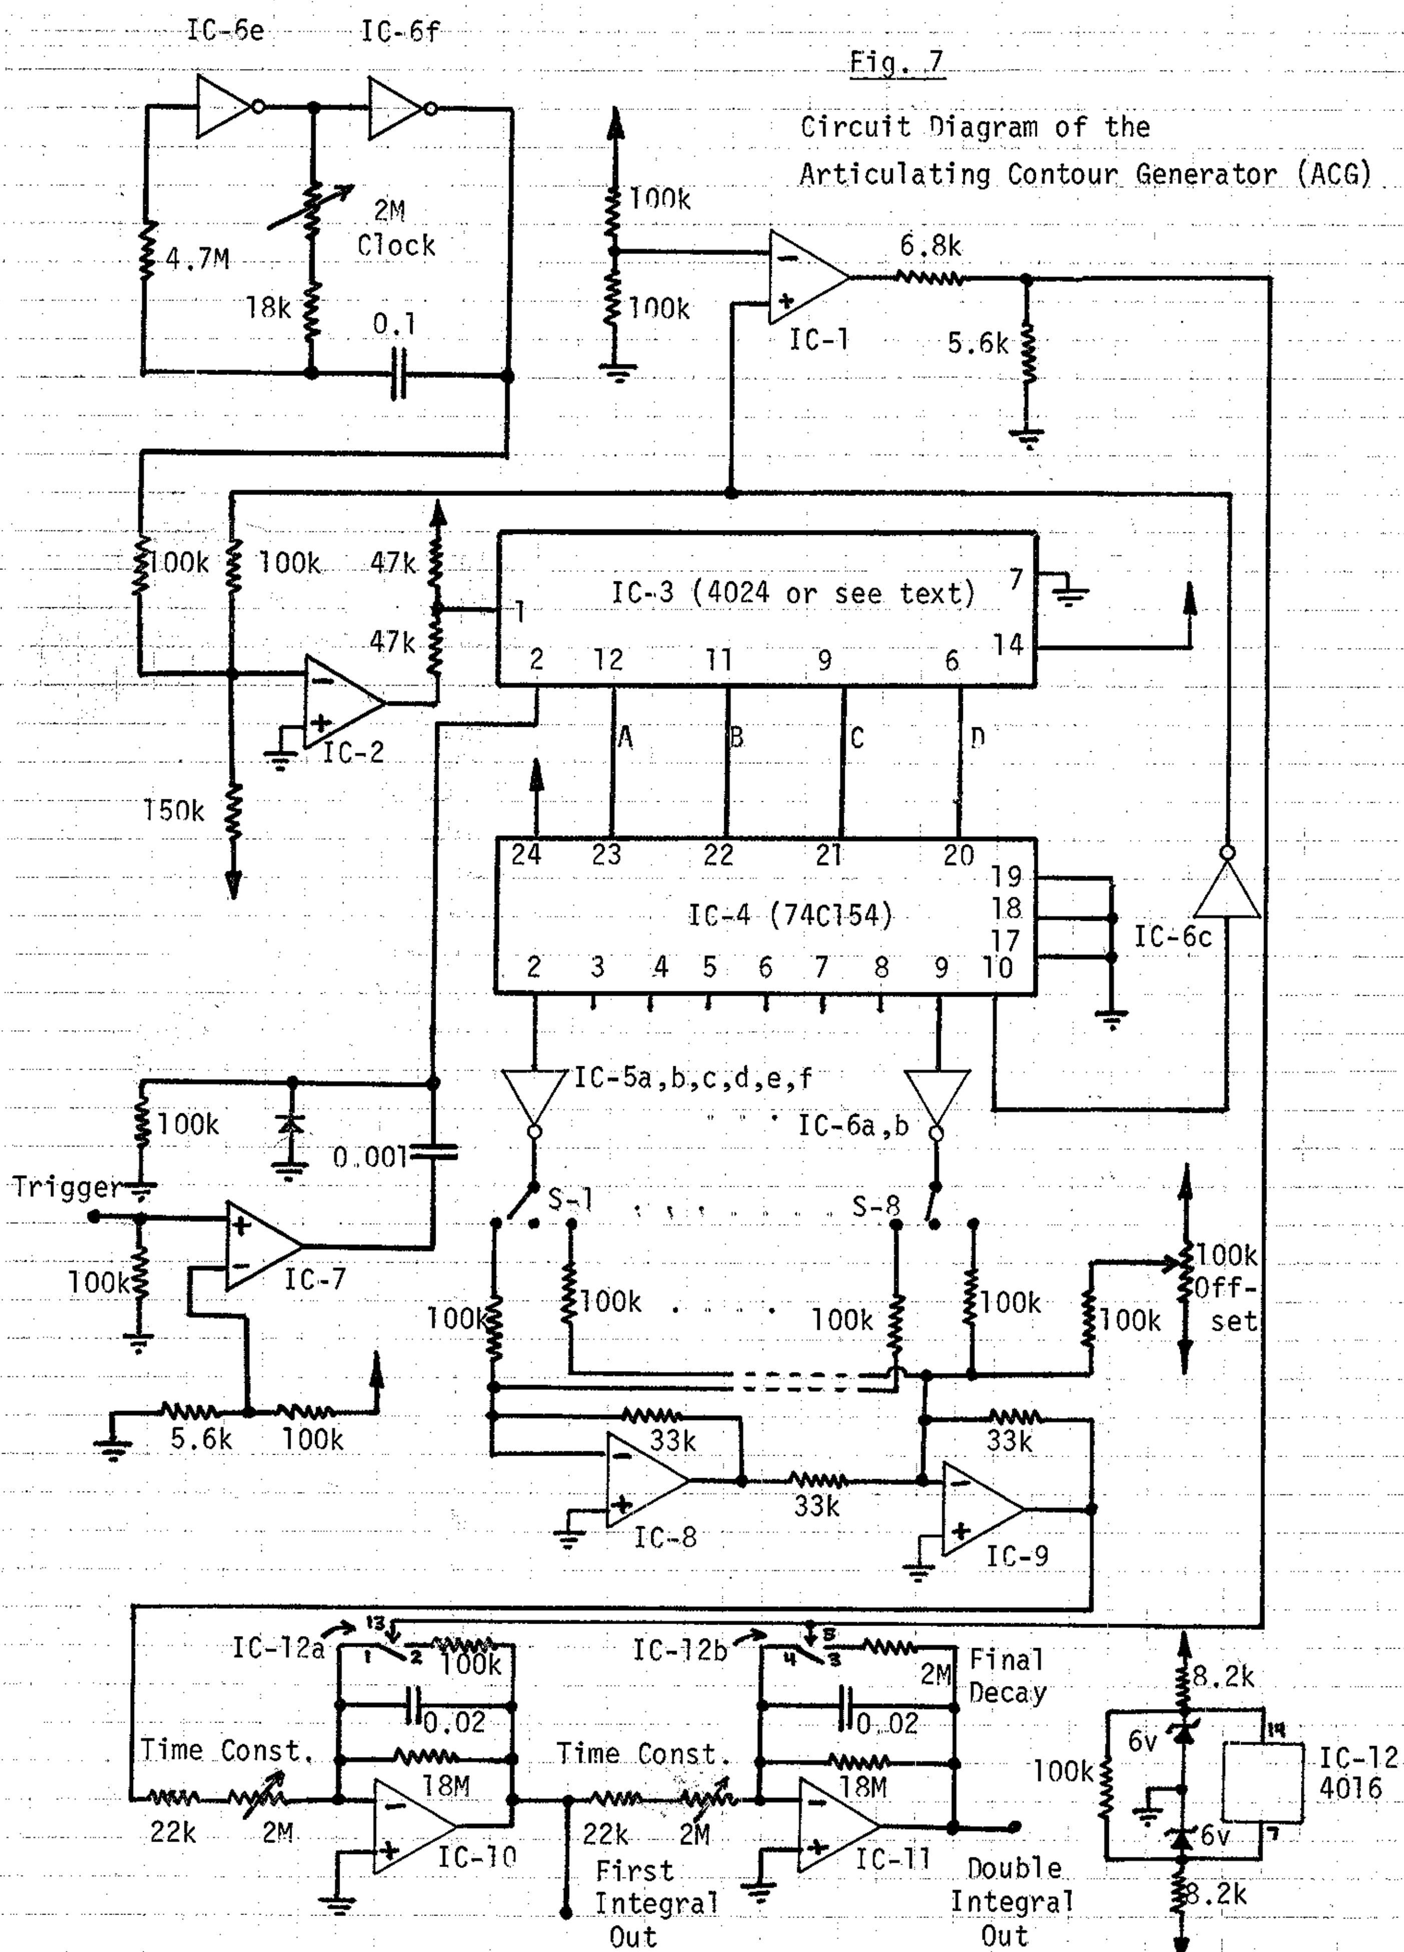

After you look the circuit over carefully, you will see that it is really a one-input, one-(or perhaps several)-output device, with numerous controls. You should think of it as a triggered envelope generator. You give it a trigger (or any convenient signal due to comparator IC-7), and it gives you a control voltage contour, and then stops waiting for another trigger. Now, the controls are very important with this unit, and have large effects on the output at times. If you had the device in front of you, you would soon see how they work. However, here we will give a brief description pointing out possibly unexpected results.

First note that the output gain is a function of several controls. If the integrator time constant controls are set to a low resistance, the integrators ramp faster and thus have a chance to reach higher amplitudes. Now, with the integrator time constants set at a given value, if the clock is slowed down, the sequence lasts longer, and again the integrators have a chance to reach a higher value. Since the clock also controls the rate at which features in the eight-step sequence eventually get out through the integrators, it makes sense to set this first, and then adjust the integrator time constants for the degree of smoothness needed, and if possible, for the correct amplitude. If not possible, try to set the amplitude higher than needed because you can almost always adjust an input attenuator on the module being controlled by the ACG.

We have discussed in the theory section the effects of choosing different sequences and different offsets. Things aren't quite as nice when it gets to the actual circuit because the integrators are not ideal, there are unwanted offsets and unequal sequence steps, and the circuits may drift. However, this is not really a practical hinderance since we have no reason to prefer any of the pictures drawn, and trial-and-error will likely be the standard mode of operation anyway. In other words, once finished, you forget the innards and just play around.

WHAT'S IT GOOD FOR?

Let's look at an example. When I finished the device, I attached it to the variable input of a VCO, and hit the trigger. The VCO chirped like a bird (I left the switches for +1, -1, +1, -1, +1, -1, +1, -1). If you max the clock rate, you get a sort of "whir" with the same patch, and if you slow the clock to minimum, you get a very slow (and very large - the second integrator clips) FM signal. The thing is it gives you four "wiggles" and then stops, unlike vibrato which keeps going. Thus, it can be a sort of "tone burst" controller (see Fig. 8).

Using less structure in the sequence gives broad contours which can be used for forming complex transients. This has not been sufficiently explored, but the results are clearly different from standard ADSR envelopes for the same functions. More results will perhaps be available in a later report.

Fig. 8

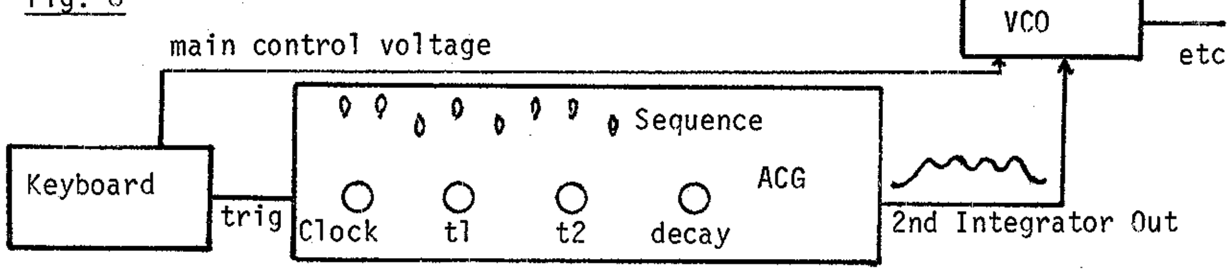

ELECTRONOTES, Vol. 12, No. 110 [February 1980](published June 1980)
Published by B. A. Hutchins, 1 Pheasant Lane, Ithaca, NY 14850
Phone (607)-273-8030
Please use above address and phone for subscription information, technical questions, and for all other business. Best phone hours - Wednesday evenings, 8 PM through 2 AM, Eastern time. Bad times to call - mornings, Tuesday and Thursday afternoons. Feel free to call any evening until midnight. To avoid the answering machine, hang up after the fifth ring if a live person does not answer.