

ELECTRONOTES 102

NEWSLETTER OF THE MUSICAL ENGINEERING GROUP

1 PHEASANT LANE

ITHACA, NEW YORK 14850

VOLUME 11, NUMBER 102

JUNE 1979

GROUP ANNOUNCEMENTS:

CONTENTS OF EN#102

- Page 2 A Vocal-Effects Waveform Animator:
 The Synthesis of "Animated" Sounds - Part 4: -by Bernie Hutchins
- Page 10 Reader's Questions

We remind readers that this issue is the second of the approximately half sized "Summer" issues that we have scheduled this year. These smaller issues will continue through September, at which time they will resume full size for the rest of the volume.

In this issue, we present another of the "animator" features and this one worked out quite well. Applications will be discussed in EN#103. If you haven't built anything for a while, and are thinking about it, this is a good one.

It seems that there is some confusion about the Electronotes Builder's Guide and Preferred Circuits Collection (EBG&PCC) despite our (feeble) efforts to get the point across. Let's start back at the beginning with the basic logic. First, a broad coverage thing like the Musical Engineer's Handbook (MEH) is just fine, but it takes too long to produce to keep up with the rapid changes in the circuitry. Thus we have made a separate publication just for the circuitry, and this is the EBG&PCC, formerly just the PCC. This is a quarterly publication. Each quarterly edition is the same in scope - it covers the building procedures, and preferred circuits for the full range of synthesizer modules (VCO's, VCA's, VCF's etc.) Each edition is different from the previous one because we update whatever we feel needs updating. Think of it like an encyclopedia. Each new edition of an encyclopedia has updated material (the moons of Jupiter for example) and material not likely to be updated (the Civil War for example). You don't need each new edition of an encyclopedia because you read much of the new material in the newspaper. It is the same with the EBG&PCC. While each one is an update, if you read the newsletter, you can easily glide through and miss several editions. Now, we don't like to waste paper and effort, so we only print as many copies of each edition as we have orders for. This is why orders must be in by a deadline. This is also why we generally offer several editions at the same time - if you come in too late for one, we will hold you over for the next one. I doubt many readers really need more than one edition each year unless they like to have home and work copies, or something like that. Note however, that in changing from the PCC to EBG&PCC, we are making a major revision. The building information is better organized. Some circuit description is added in a complete form. We are also offering it to non-subscribers for the first time (at a higher price). If you have ordered two or more consecutive editions, and would rather spread them out a little, just drop us a note and we will shove the later ones down the line. It's easy to do.

A VOCAL-EFFECTS WAVEFORM ANIMATOR:

THE SYNTHESIS OF "ANIMATED" SOUNDS - PART 4:

-by Bernie Hutchins

INTRODUCTION: This report continues our series on animators. This time we have what amounts to a filter bank with a small number of filters, but which is dynamic in that the filter characteristics, frequencies and Q's, can be made to vary. At the same time, the unit was designed with a vocal sound imitation capability in mind. It was useful (perhaps necessary) to have this vocal sound imitation in mind because it allows us to make some design decisions that would otherwise have been virtual guesses. There are a lot of variables in this sort of scheme, and when it was necessary to make a decision about a given variable (such as how many filters), we look at the vocal sound as a guide (usually three variable formants).

THE BASIC SETUP IDEAS:

A standard filter bank is composed of a set of perhaps 20 to 40 fixed filters with spacing staggered with respect to harmonic positions. If we look at the vowels of English speech (or singing), we find that they are produced to some approximation by an excitation source and three or four dynamically varying filters. To actually "voice" a given vowel sound, these filters must vary in a non-arbitrary way. The present animator is composed of three VCF bandpass filters simulating speech formants, one fixed bandpass for an upper formant, and one variable notch for a "nasal" sound. The VCF BP filters are also subject to a voltage-controlled Q.

No attempt is being made to produce a given vowel sound at any one time, nor are we trying to produce undetermined, but standard vowel sounds. Instead, the filters are selected so that they could take on the formant positions of standard vowels. Put another way, the ranges of the VCF bandpass filters are set to cover the general ranges of the first, second, and third formants occurring in English vowels, yet no attempt is made to have them in corresponding positions for well studied standard vowels.

Early speech synthesizers were somewhat like this device. They supplied control voltage curves to VCF's and these curves were adjusted for the desired speech sound. Here we will not be using envelope control curves, but rather, summed Low-Frequency Oscillator (LFO) outputs. We have looked at the process of summing LFO's, either directly, or by controlling some sort of processing with LFO's, and then summing the processors. We have found that a mix of three or more LFO's may be sufficient to hide any obvious repetitions, and yet not sound totally random. The LFO's used in the current animator are actually simple VCO's, so the LFO's can all be moved in parallel to raise or lower the average frequency (roughly analogous to talking faster). The VCF's can also be raised or lowered in parallel.

To summarize the basic idea, we are looking at a device which is roughly an analog of the vocal sound system. We will excite this device with an externally chosen waveform. The device will then process this waveform in a dynamic way, where the dynamic control voltages are summed LFO's - basically out of the control of the user except for certain average settings. The idea is to eventually end up with an animator capable of producing vocal-like sounds, and as many others as possible.

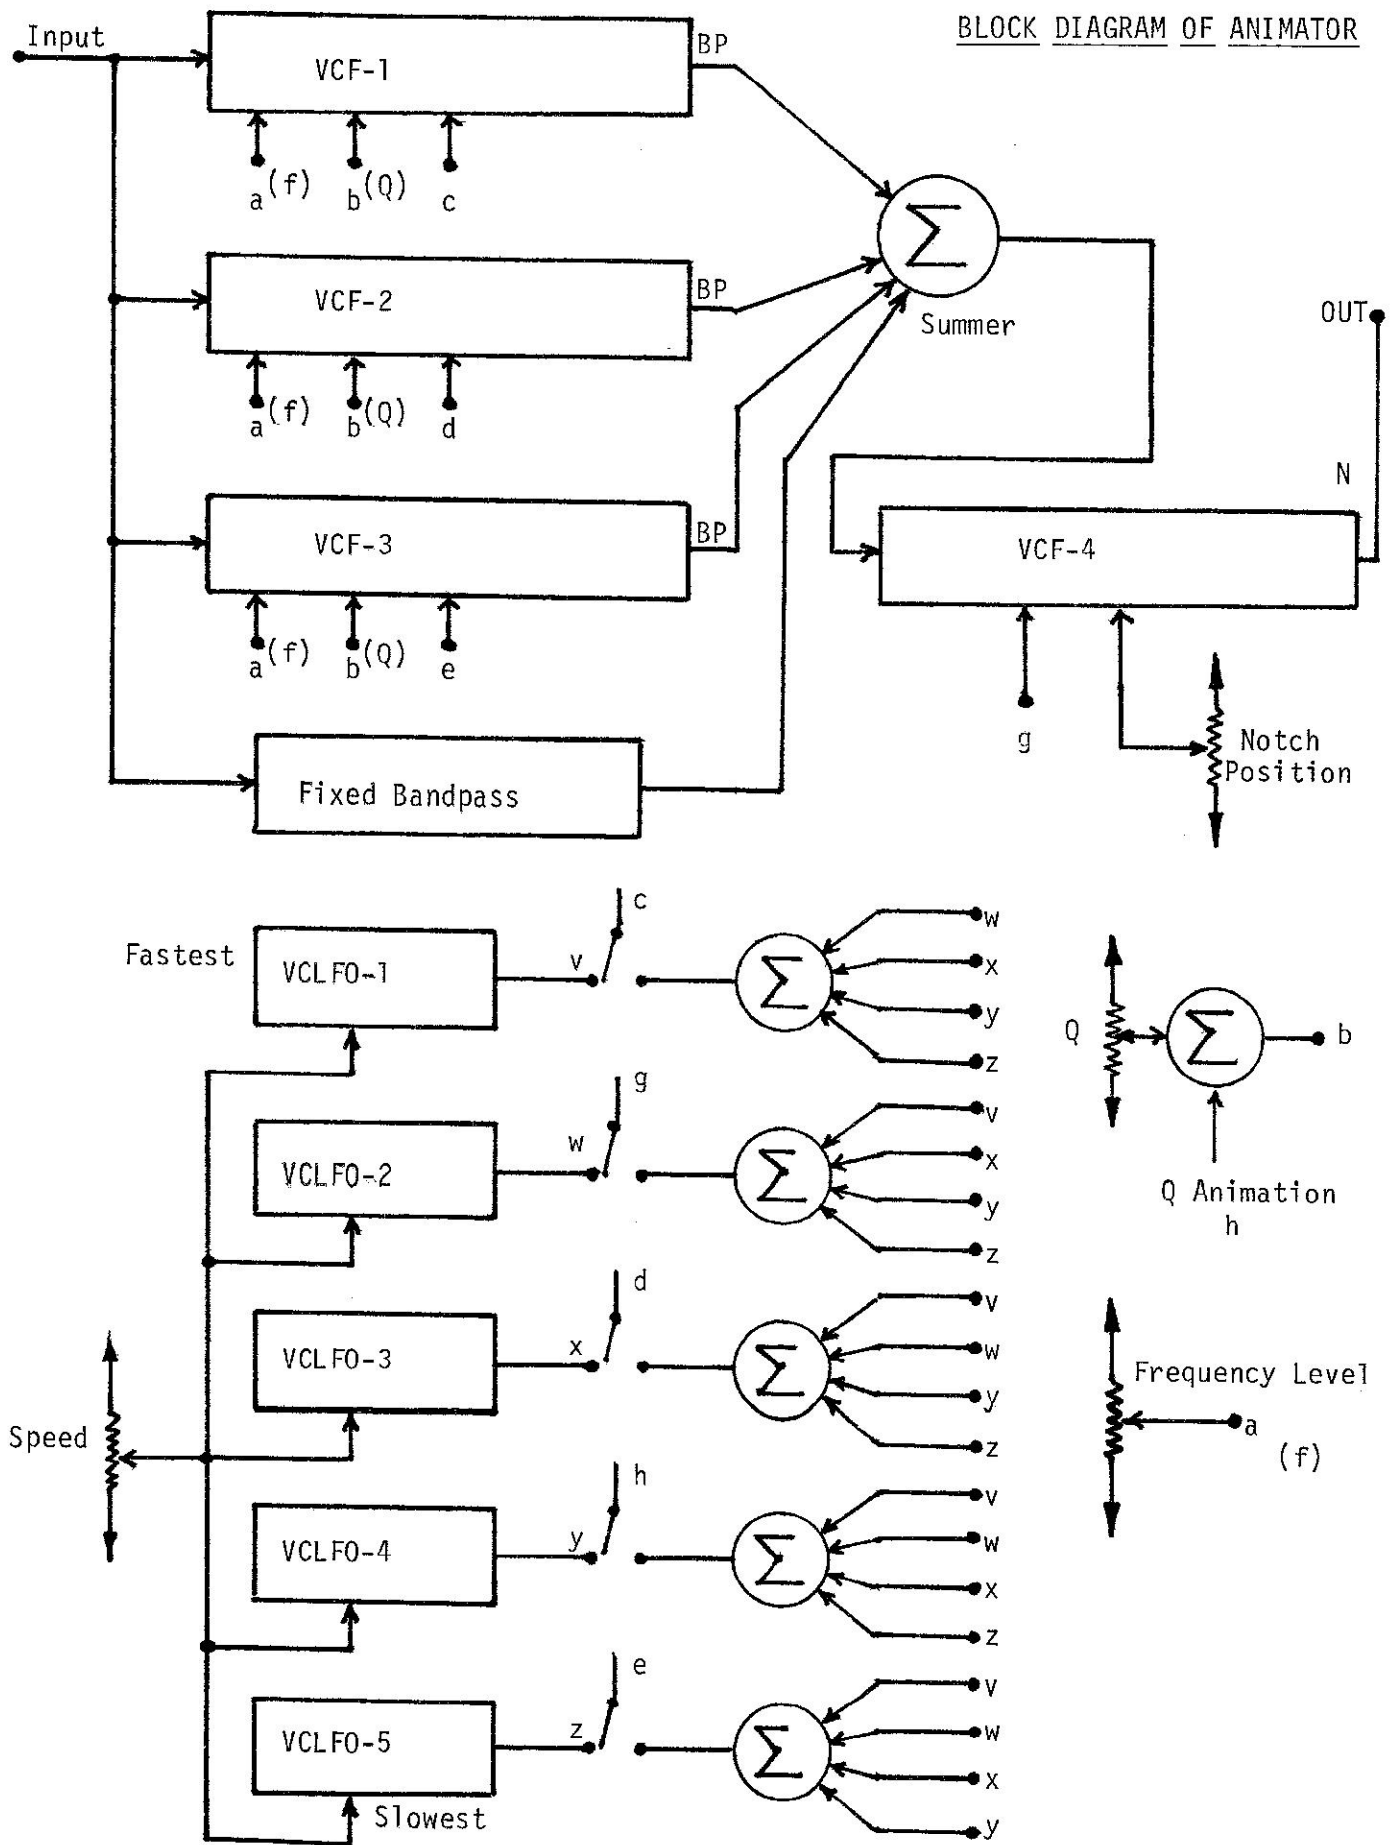
BLOCK DIAGRAM:

A block diagram of the system is shown in Fig. 1. Do not become alarmed that there are numerous VCO's and VCF's in the system, as these are simple designs and do not involve great expense or undue construction time. The total construction effort is probably about the same as that for four to five average modules. Only standard and readily available parts are used for the most part.

First, we will look at the bottom portion of the block diagram, the control

Fig. 1

BLOCK DIAGRAM OF ANIMATOR




voltage generating portion. This is comprised simply of five, voltage-controlled LFO's that are tuned in parallel. There are several ways of choosing the initial frequencies, and we will discuss this later. The five LFO's output five triangular waveforms, v, w, x, y, and z as shown. Now, a single triangular waveform is easily "tracked" when used for control purposes, so we find it useful to also make available certain summed versions of the LFO outputs. In particular, here we offer as an alternative to waveform v, a summed version of w, x, y, and z., and so on for all the other triangles. Thus, we can choose the triangle itself, or the sum of all the other triangles. This may seem arbitrary, but it seems as good as anything, and is a coherent starting point. For the moment, the reader might like to think about which set of waveforms, v, w, x, y, and z, or the corresponding "all other" sums will be the most useful for animation. Note that the triangles themselves are the most "trackable" but the sums are "correlated" to a degree.

Now we will move up to the top of the block diagram where we find four VCF's and one fixed filter. VCF-1, VCF-2, and VCF-3 correspond to the first, second, and third voice formants respectively, and are state-variable VCF's using the bandpass output. A fourth fixed formant is simulated by the fixed bandpass filter. VCF-4 is used as a notch, and serves to simulate the "nasal" sounds that results from zeros in the vocal track (sound from the nose with the mouth blocked). Note that the notch filter can be "removed" simply by moving the notch frequency up out of the way. Note that the bandpass filters are used in parallel since there is to be little overlap involved and a series combination could give large amplitude variations (relatively long "dead times").

Now, we will look at the controls to the VCF's. First, there is the main frequency level control, pot marked "f", which supplies control voltage "a" to 3 of the VCF's. Secondly, the VCF's used as bandpass filters have voltage-controlled Q, and this is controlled first by the "Q" pot, and also by a variable depth control voltage h, which is available from the LFO summer section. Finally, each of the VCF's has its own independent control (c, d, e, and g) which comes from the LFO summer section.

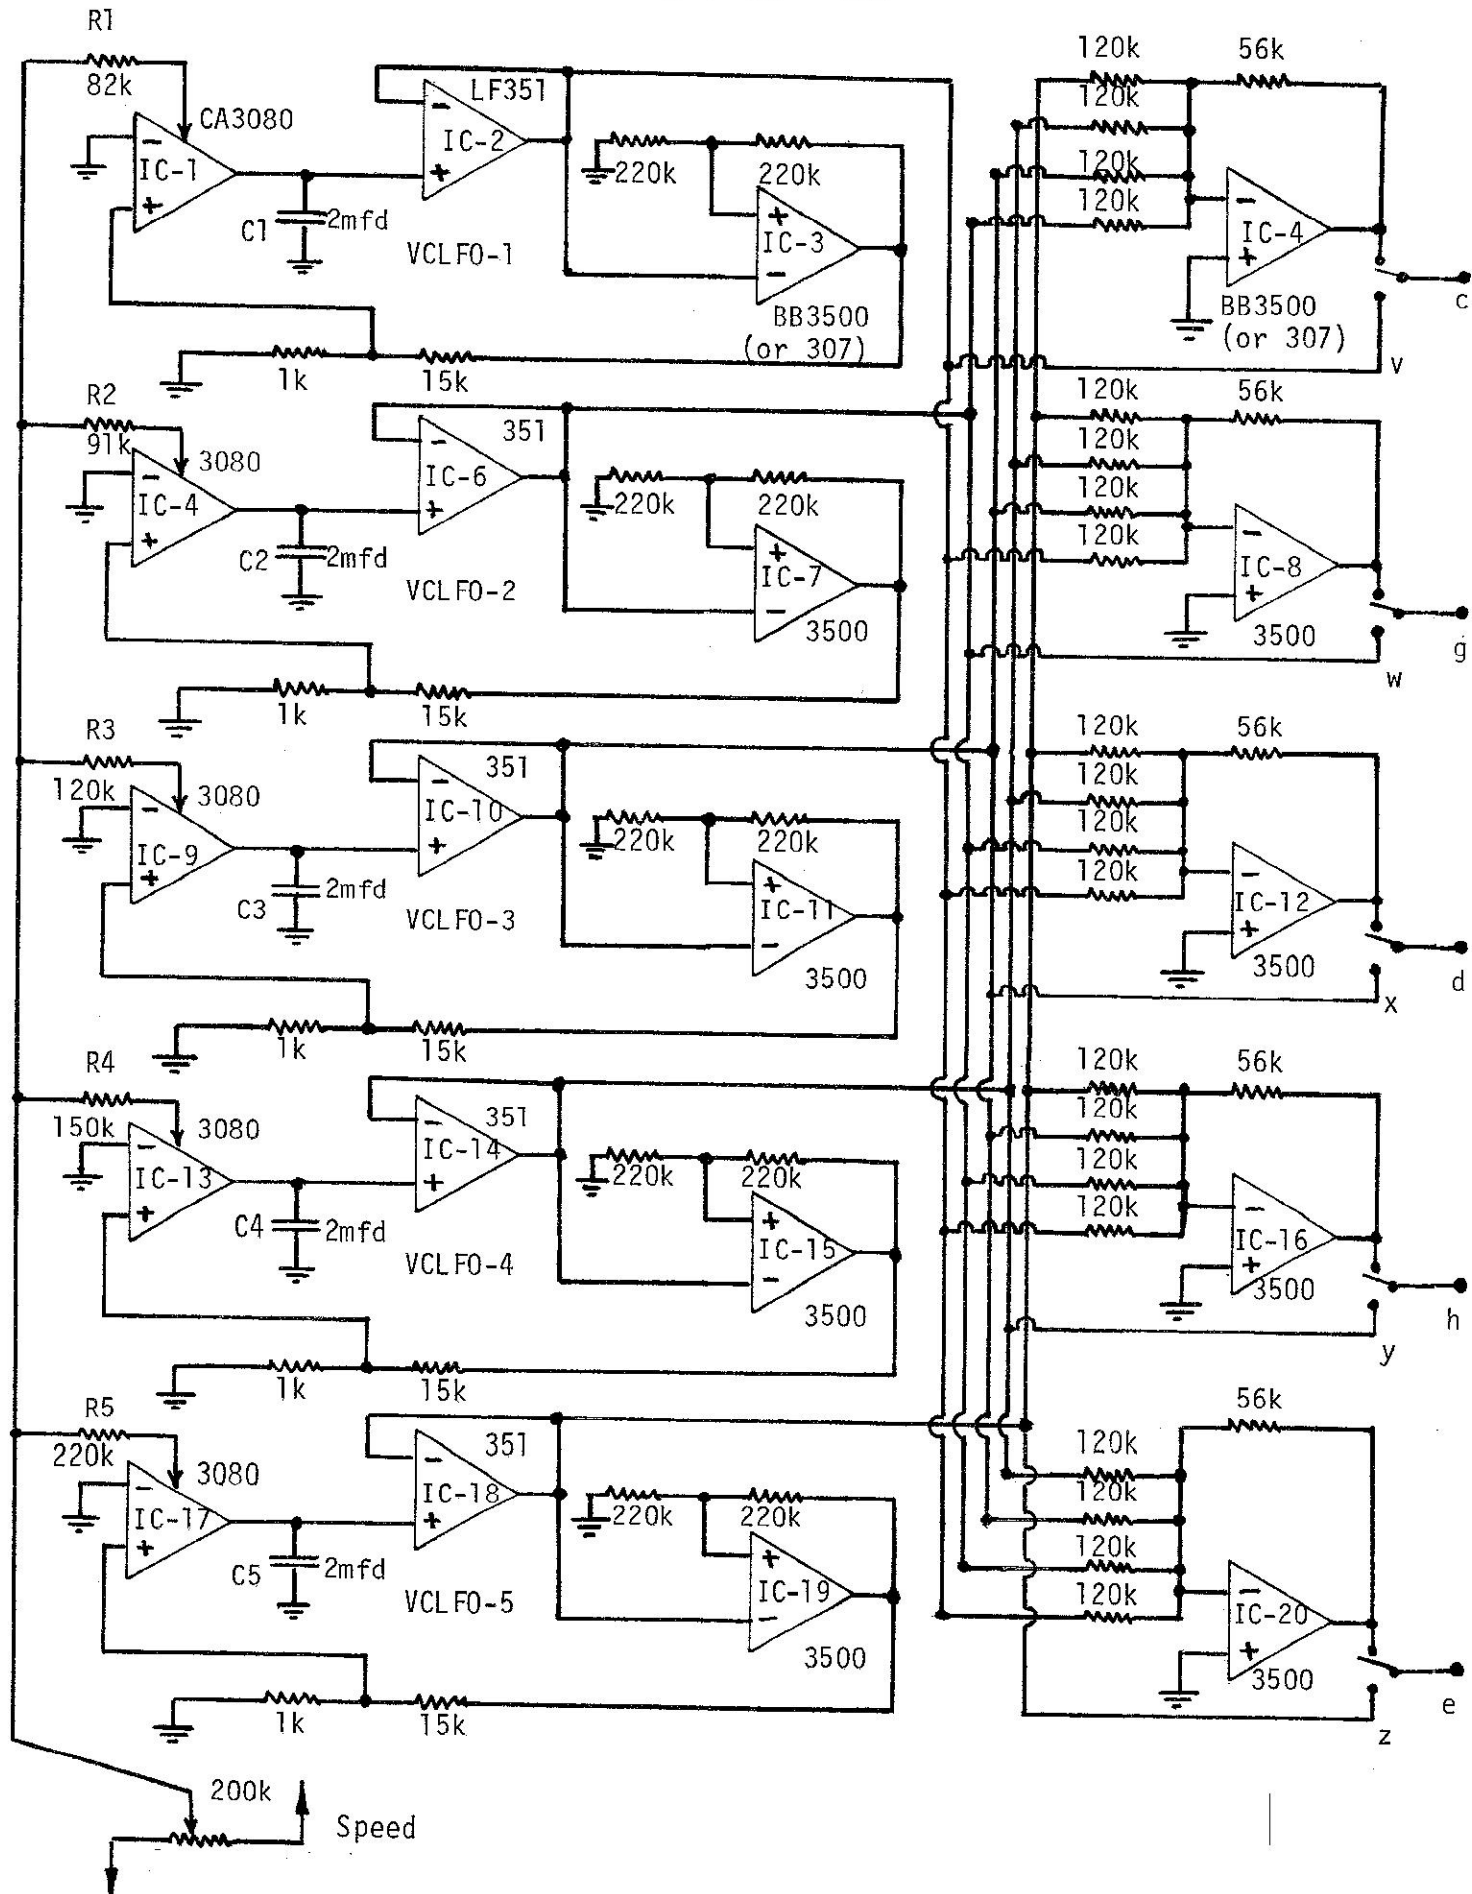
All this is pretty simple. A possible panel diagram is shown in Fig. 2 where you will note that this is basically an input-output device with a few controls. The "speed" control is used to move all LFO's up or down by a proportional amount. Note that this is not a pitch level control even though it controls VCO's. Rather, it controls the degree of "activity" in the animation. The "frequency level" control sets the initial frequency of all the VCF's, but this isn't a pitch control either because the actual pitch is determined by the external excitation. The "frequency level" control sets the region to be processed. There are two "Q" controls, one of which sets the initial Q of the filters, and another which superimposes on this a summed LFO control signal, if desired. This is followed by five switches to select the original triangle wave, or a summed LFO version for the filter and Q controls. Finally, there is a notch position (initial position) control.

Fig. 2



Panel Diagram of Vocal Effects Waveform Animator

Fig. 3 VC LFO's and Summers



CIRCUIT DESCRIPTION:

The circuitry, as shown in Fig. 3 through Fig. 7 may appear quite involved, but you will also note that much of it is repetitive and this makes construction a good deal easier. Most of the basic circuits will be familiar to our readers, although there are a few new ideas used to simplify things as much as possible.

Fig. 3 shows the Low-Frequency VCO section (VCLFO). The VCLFO's are composed of a CA3080 and two op-amps, and IC-1, IC-2, and IC-3 form a typical unit. IC-1 is a current-controlled current switch, which serves to charge and discharge capacitor C1, thus producing a triangular waveform, which is in turn buffered by IC-2. Note that IC-2 is necessary because we want to use the triangle wave - if we only wanted a square wave, IC-2 could be omitted and IC-3 would receive the capacitor voltage. IC-3 is just an op-amp Schmitt-trigger of standard design. The output will change state whenever the (-) input tries to exceed +7.5 or -7.5, and hence the triangle out of IC-2 is limited to approximately a 7.5 volt amplitude. Four other VCLFO's identical to the one formed from IC-1, IC-2, and IC-3 are found directly below it, the only difference being the resistance that supplies current to pin 5 of the CA3080 involved. These are different simply because we want the frequencies of the different VCO's to be different.

Some thought should go into the spacing of these LFO frequencies. If the frequencies are all in a fairly small range, you will get a sort of "wave packet" effect with large amplitude variations in a summed output. On the other hand, a very regular and wide spacing may not be good either because it will also result in beating effects. The ratio shown, determined by the resistors $R1 = 82k$, $R2 = 91k$, $R3 = 120k$, $R4 = 150k$, and $R5 = 220k$ seems satisfactory.

This is perhaps a good place to mention that the actual success or failure of a "disguise" of the LFO sum depends on several factors, and the circuit of Fig. 3 is a good means of testing this. We have here the five summers IC-4, IC-8, IC-12, IC-16, and IC-20 which sum all combinations of the five LFO's taken four at a time. We can add to this the sum of all five (given in Fig. 7). The main purpose in making these LFO's from a VCO design is that they may all be controlled in parallel. With the device shown in Fig. 3 (and 7), we can determine several general things about the use of summed LFO's for animation control. First, the disguise of the pattern involved is better at lower LFO frequencies than at higher LFO frequencies. This fact is in line with intuition - if patterns repeat slowly, we tend to forget them. Secondly, the disguise gets better as we move from using the LFO sum as pitch control to amplitude control, and gets still better if we move to timbre control. This again seems reasonable. In fact, using a summed LFO signal for pitch control may well remind you of a fly buzzing around your head, and that is by no means subtle. In the timbre control case however, the same pattern is better hidden because timbre variations are much more subtle than either pitch or amplitude variations. Thirdly, the full device shown that in general patterns are better hidden when they first control processors, and the processors are then summed.

There may be better choices for frequency ratios, and perhaps an unequal sum would work out, but the values in Fig. 3 will probably be as good a starting point as any. Certainly some experimentation would be advisable.

Fig. 4 shows the bandpass section of the vocal animator. This consists of three bandpass VCF's (based on the state-variable technique, of which IC-21 through IC-25 forms a typical unit, with two more similar units directly below), and a fixed "Deliyannis" bandpass filter (IC-36), with IC-37 performing a summing function of all four bandpass filters.

The three bandpass VCF's may look a little different from the ones we have been using. In fact, we have generally been building them from two CA3080's and three op-amps, while this new design has three CA3080's and two op-amps. This is

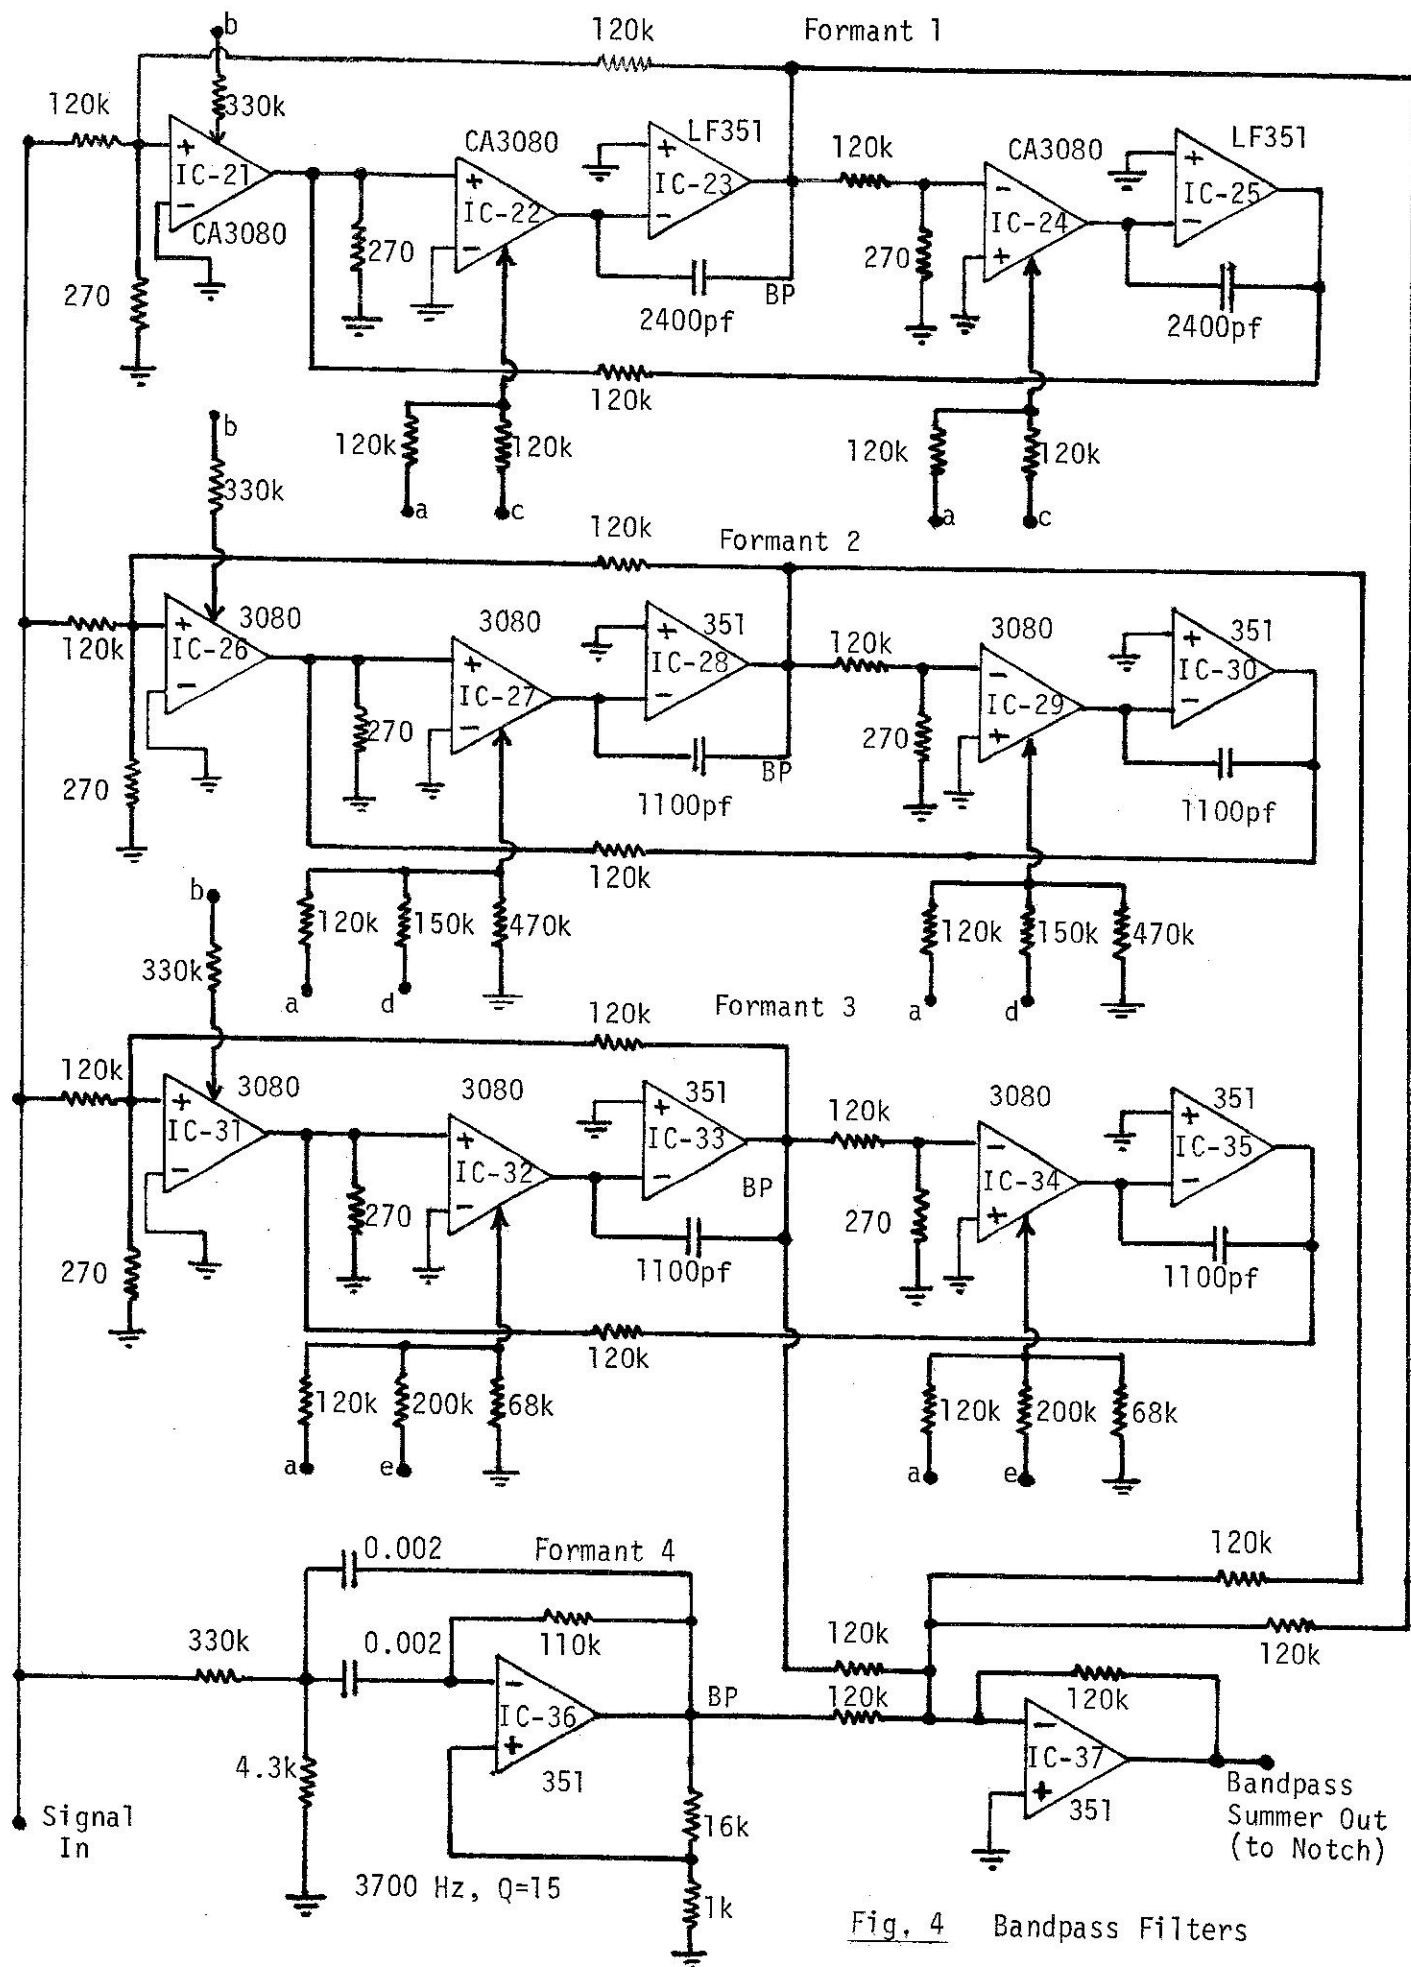


Fig. 4 Bandpass Filters

not a step backward, since the third CA3080 is for Q-control, and is an added feature, while one op-amp has really been done away with. We will describe the circuit in terms of the top VCF (IC-21 through IC-25) and the same description applies to the other two (IC-26 - IC-30, and IC-31 - IC-35).

IC-21 is the Q-control CA3080, and thus controls the amount of signal being fed back from the bandpass output to the input. At the same time, the input is passed through the same control device (meaning IC-21), so as the Q goes up, the input is attenuated, resulting in a constant peak response as the Q varies. See application note AN-121 for more details on this scheme.

Ordinarily there is an op-amp summer next in the line, the one that provides the high-pass output. Here however we are not using the high-pass, and have no need to make it available as a low-impedance, standard level signal. Instead, we can do the necessary summing right on the input of IC-22 by using current summing. That is, the 270 ohm resistor is much smaller than the 120k resistor, and thus the 120k resistor looks like a current source, while the output of IC-21 is a current source. This pretty much takes care of the circuit, since the remainder is just two voltage-controlled integrators completing the state variable structure (IC-22 and IC-23, IC-24 and IC-25).

Now it is the time to say something about the current controlling resistors that are on pin 5 of all the CA3080's. The placement of the frequencies and the sweep ranges started out with a brief look at the literature regarding formants in speech. It is a little difficult to say things "in general" about speech, but data for the F1, F2, and F3 formants were chosen from a table of formants for vowels. [The source was Peterson and Barney, "Control Methods Used in a Study of the Vowels," J. Acoust. Soc. Amer., Vol. 24, No. 2, pp. 38-58, March 1952, as excerpted by L. R. Rabiner and R. W. Schafer in their book Digital Processing of Speech Signals, Prentice Hall (1978), pg. 45. From this, we get F1 in the approximate range of 300 - 800 Hz, F2 in the approximate range of 800 Hz to 2200 Hz, and F3 in the approximate range of 2200 Hz to 3000 Hz. With some care we could have set these to be the exact ranges swept, but in the interest of simplicity, a slightly different approach was used. This was desirable because we want to just drive the VCF's with current determining resistors, with pin 5 of the CA3080's sitting at about -14.3. This means that our zero-centered triangles are going to already supply half the possible current. Instead, we chose to set all the initial frequencies at zero triangle voltage to about 60% the desired value, and these could then be pulled up with the frequency level control. In this way, the frequency level can be lowered as well as raised, a desirable feature. Experiment tends to bear this out as well. The final step is to adjust the range. All of the triangles have a range of ± 7.5 volts, and this is fine for sweeping the lowest formant (in fact, this is how we chose 7.5 volts). The upper formants have somewhat less range however, and to simulate this, we split the normal resistor feeding in the triangle into two resistors with the same parallel equivalent. Thus, part of the triangle was supplied by the triangle, and part by ground (a zero amplitude triangle!). For the second formant, 470k is run to ground, relative to 120k to triangle, while F3 has much more shunted away, 68k to ground relative to 200k to triangle.

The reduction of formant range is evident and comes to its ultimate conclusion with a fixed fourth formant, F4, achieved with the simple bandpass filter around IC-36. This is set for a Q of 15 and a center frequency of 3700 Hz. The "Deliyannis" structure was used because initially it was not known what Q would be needed, and the Deliyannis structure is capable of high as well as low Q's.

Note that the bandpass filters are simply summed. This can be justified because we have little overlap between their ranges, and since at times we are using a high Q, series structure would often result in very little signal.

Fig. 5 Notch

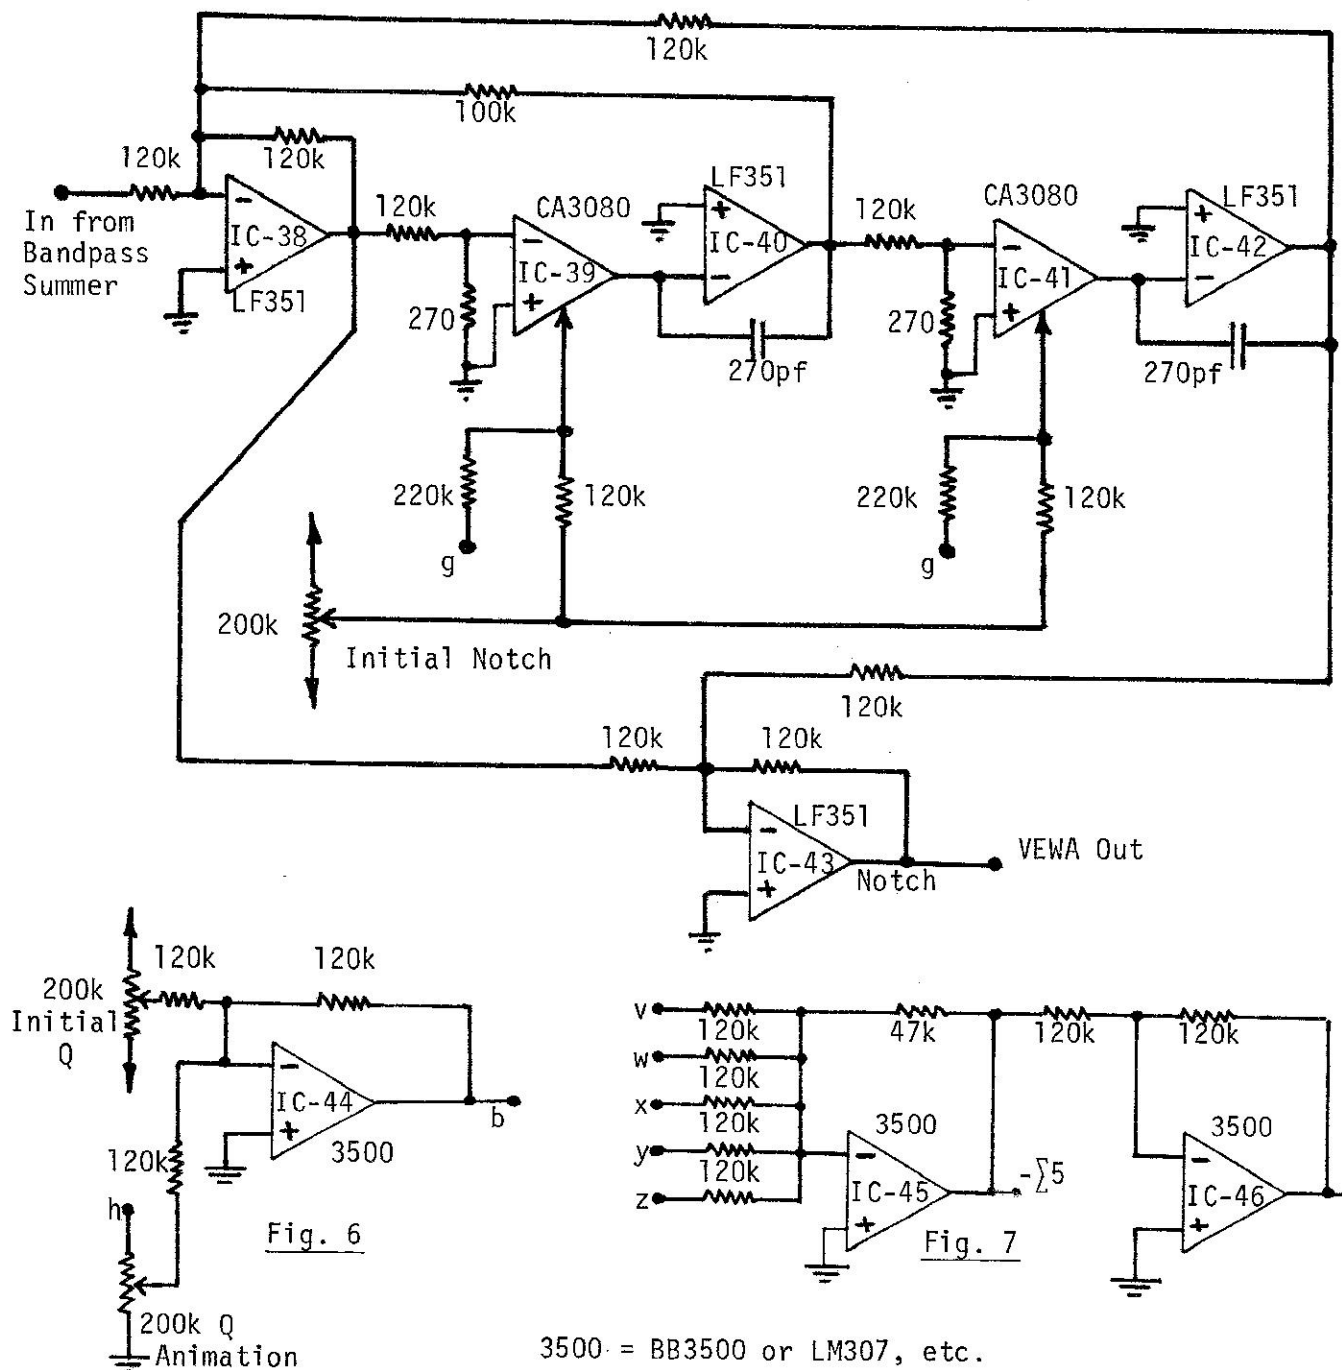


Fig. 5 shows a fourth state-variable VCF, this one without voltage-controlled Q, but with a notch feature. Here it is necessary to have a high-pass summer because we need to sum low-pass and high-pass to get the notch we desire. The result is a fairly conventional design. Note that the notch VCF follows the bandpass summer and is controlled by its own frequency level control (notch position) and by an animating voltage.

Fig. 6 and Fig. 7 are basically left-over circuits that did not fit well on the other diagrams. Fig. 6 is the Q-control summer. Note first the effect of the initial Q control. As the voltage on the wiper of this pot gets more positive, the voltage at the output of the summer gets more negative. Thus the control currents to the CA3080 Q-controls will get smaller, and this means less feedback, and therefore a higher Q. Thus, more positive voltages result in a higher Q. In addition to the initial Q, an animating voltage can be superimposed, and to a degree determined by

the 200k lower pot. In general the initial Q can have a rather dramatic effect on the sound, while the animated Q is much more subtle, or at least harder to separate from the other animation effects.

Fig. 7 is a LFO summer that goes back to Fig. 3 for its actual drive. This is a special sum of all five of the VCLFO's. It is very useful although it is not used internally. It might also be a good idea to obtain several more of these summed voltages, perhaps by an unequally weighted sum of all five. These extra LFO summed voltages are used to control such things as excitation frequency and output amplitude when the animator is used in a full patch. Fig. 7 shows both the sum and the inverted sum, and it is a good idea to bring both these out, as it is useful in at least one important patch which we shall demonstrate.

CONSTRUCTION:

A glance at the figures will perhaps show that a logical setup is to have Fig. 3 and Fig. 4 as separate boards, and the remainder of the circuitry can go on a third board. However, there is no real critical layout, and anything you work out should be just fine. It might be a good idea to leave a few additional op-amp positions on the boards somewhere in the event that you want to expand.

APPLICATIONS:

We will be looking at observations and applications in the next issue of this newsletter. We should point out here however that initial observations indicate that this is one of the more successful of the animators we have presented so far.

* * * * *

READER'S QUESTIONS:

► Q: I have a need X-Y-Z for a commercial unit. What commercial product do you recommend for this application?

A: We are not in a position to recommend commercial equipment. In the first place, we are not able to evaluate commercial equipment in a complete or useful way. We don't have it to work with, and judging criterion other than purely engineering ones are not at all clear anyway. Secondly, commercial equipment is not our main interest - we like to work with new home-built devices. The reader can rest assured however that most commercial equipment is of good quality, and with the degree of competition in the field of electronic music, any product that is below par is not going to be on the market long. We can only add that you should shop around as much as possible, and don't rush into things.

► Q: When is the new edition of the Musical Engineer's Handbook coming out?

A: This question keeps coming up - perhaps we perpetuate the idea of a 2nd edition by answering the question frequently. There is no new edition planned. Instead, the MEH will continue in its original form for some time to come. Eventually, the MEH will become superseded by a combination of books, of which the Electronotes Builder's Guide and Preferred Circuit Collection and Application Notes are perhaps early ancestors. The idea is to separate material of different sorts, particularly with regard to the need of the material for frequent update (or not). There is no real "master plan" for doing this however because it is felt that things will just fall into place as we go along.

* * * * *

ELECTRONOTES, Vol. 11, No. 102 [June 1979] (Published August 1979)

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 - any morning, all day Monday, Tuesday and Thursday afternoons. Feel free to call any evening other than Monday until 11:30 Eastern time.