(This article was first published in *Wireless World*, October & December 1980. A brief description of sound production in the pipe organ which appeared in the original is omitted here).

Revision date: 22 Jan 07 15:40

Tone Filters for Electronic Organs

by C E Pykett B Sc, Ph D

Part 1 : Organ Tone Spectra & Source Waveforms

As the organ is a sustained-tone instrument, achieving a satisfactory imitation of the steady-state acoustic emission of organ pipes is of paramount importance. In this respect, the design of the tone-forming filters is crucial, yet there is a curious absence of definitive material dealing with filter design. This is apparently reflected in the range of commercial instruments on the market: with few exceptions, their "voicing" seems to be mainly empirical.

To derive a simple expression for the frequency response of a tone filter, consider the basic organ system, representative of a wide range of electronic instruments, shown in Fig.1. The waveforms are initially derived from a continuously running tone generator. Waveforms at various frequencies are selected by depressing keys, and envelope shaping may be applied at the instants of key attack and release to simulate the transient phenomena of organ pipes. (Whilst of considerable importance, transients are not further discussed here.) The signals are passed through various tone-forming filters depending on the stops or tone colours selected and the output from the filters is then finally amplified and fed to loudspeakers.

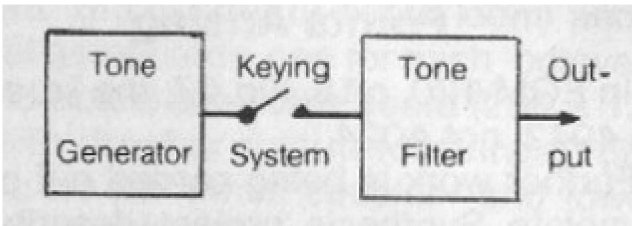

Figure 1. *Basic electronic organ system considered in this article is the subtractive kind in which an harmonically rich waveform is filtered.*

A tone filter may be thought of as an amplifier whose gain varies with frequency. The gain can therefore be explicitly written as a function of frequency, *G(f)*. Similarly, each harmonically-rich waveform from the generators is equivalent to a large number of individual sine waves of different frequencies, each sine wave having a different amplitude. This waveform can also be written as a function of frequency, say *H(f)*. Therefore the output from the tone filter, *F(f)*, is the product of the input voltage and the gain just as with any amplifier:

$$F(f) = G(f).H(f)$$

In general, the tone filter will also modify the phase as well as the amplitude of each frequency component in the input signal. As the ear is insensitive to relative phase for present purposes, this does not matter, which makes the design of tone filters much easier than it would otherwise be! It does mean, however, that the *waveform* emerging from the tone filter will not necessarily bear any resemblance to the waveform emitted by the organ pipe if both were to be viewed on an oscilloscope screen. It is only the *frequency spectra* that need to be matched as closely as possible.

If the frequency functions are expressed on a logarithmic amplitude scale then new functions are obtained that are related by addition rather than multiplication:

$$P(f) = Q(f) + R(f)$$

Rearranging this equation gives the frequency response of the tone filter, *Q(f)*, in terms of the input spectrum from the tone generator, *R(f)*, and of the output spectrum, *P(f)*:

$$Q(f) = P(f) - R(f)$$

This simple equation shows that filter design involves three basic steps. First, the logarithmic spectrum of both the tone generator waveform and of the sound to be simulated must be available. Second, the frequency response of the required filter must be derived by subtracting one from the other. Third, the response so obtained has to be realised in hardware. Subsequent sections discuss each of these stages in detail.

Acoustic Spectra of Organ Tones

Before a filter can be designed to imitate the sound of a particular type of organ pipe, the spectrum of that sound must be obtained. Following a careful search of the scientific and engineering literature extending back into the 1930's, it was discovered that very few systematic investigations into the acoustic spectra of organ tones have been reported. As this information is vital to the design of an imitative electronic instrument, three of the most useful references are appended here. (Refs 2,3 & 4). Boner's article (1938) describes one of the first attempts to use electronic techniques to analyse the sound of an organ pipe radiating in a free field (that is, away from the reverberant conditions of an auditorium) by mounting organ pipes atop a 24ft tower out of doors. From the three references quoted, spectra corresponding to the four main classes of organ tone can be extracted, viz flutes, diapasons, strings and reeds, and this goes some way to providing a framework for the design of a wide range of filters. To augment this information I have made recordings of organ sounds and analysed them. A large amount of information was obtained from a four manual instrument by Rushworth & Dreaper with some particularly fine solo stops.

Recordings were made of organ pipes *in situ* using omni-directional capacitor microphones with a frequency response from below 20Hz to about 20kHz. Two microphones were used, feeding separate channels of a tape recorder with a frequency response from 35Hz to16kHz (± 2db). The recordings were subsequently replayed monaurally into a high-resolution spectrum analysis system with a dynamic range of 60dB. The reason for using two microphones and then summing their outputs on replay was to reduce distortion of the spectrum through reflections from the surfaces in the auditorium. Because they set up standing waves, such reflections can result in a significant increase or decrease in the intensity of sound of a particular frequency at the microphone location. By using two microphones there is a reduced likelihood of an identical distorting effect occurring at both simultaneously. (A better method for averaging out the effects of reverberation would have been to use averaging in the frequency domain after phase information had been removed.) Recordings were made of four octavely-related samples from each stop on the organ, and the whole exercise has resulted in a library of some hundreds of pipe spectra.

The steady-state emission of a pipe is periodic at its fundamental frequency. This is the lowest frequency present in the spectrum in most cases and it defines the musical pitch of the pipe. Because the emitted waveform is periodic, the only other frequencies present in the spectrum are harmonics or integer multiples of the fundamental; there is virtually no acoustic energy lying between adjacent harmonics. Certain pipes, however, possess a significant noise component due to random fluctuations of the air. In other cases, the amplitudes and phases of each harmonic fluctuate randomly to a significant degree. Both of these effects produce energy that is not confined to the harmonic frequencies in the spectrum. However, it is assumed here, for simplicity, that the spectrum of an organ pipe consists only of equally-spaced lines at the fundamental and harmonic frequencies.

This structure is shown in Fig. 2, with examples of spectra corresponding to each of the four classes of tone. These have been normalised to the frequency of the fundamental so that the abscissae represent harmonic numbers (on a logarithmic frequency scale).

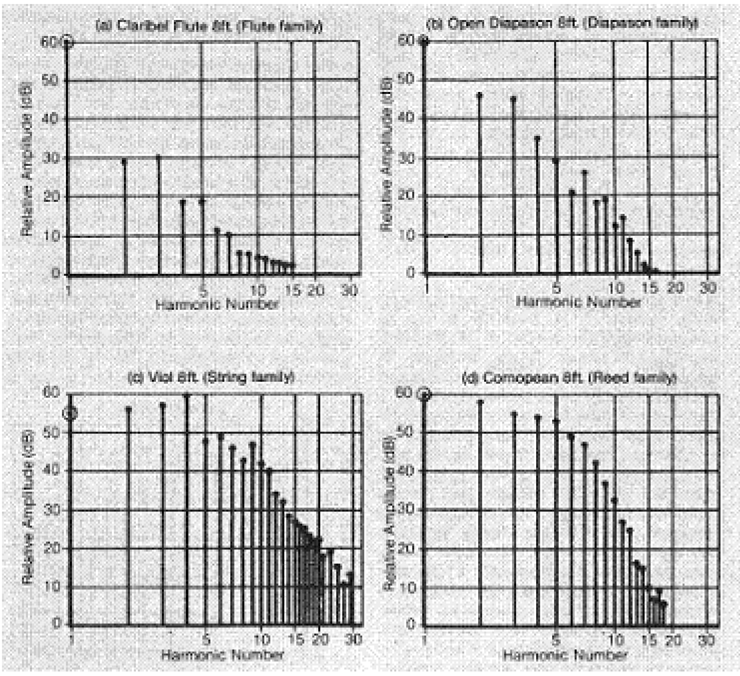

Figure 2. *Large number of harmonics in organ pipe spectra means high cost for additive instruments.*

Harmonic	Claribel Flute	Open Diapason	Viol	Cornopean
1	60	60	55	60
2	29	46	56	58
3	30	45	57	55
4	18	35	60	54
5	19	29	48	53
6	11	21	49	49
7	10	26	46	47
8	5	18	43	42
9	5	19	47	37
10	4	12	42	33
11	4	14	40	27
12	3	8	34	25
13	3	5	32	16

14	2	2	28	15
15	2	1	27	10
16	-	0	26	7
17	-	-	25	9
18	-	-	23	6
19	-	-	22	-
20	-	-	22	-
21	-	-	18	-
22	-	-	20	-
23	-	-	19	-
24	-	-	15	-
25	-	-	20	-
26	-	-	11	-
27	-	-	14	-
28	-	-	13	-
29	-	-	-	-
30	-	-	-	-

Table 1. *Harmonic amplitudes of various pipe spectra in dB, corresponding to Fig. 2*

All of these spectra contain a large number of harmonics, at least 15, within the dynamic range of 60dB. This is significant in that it clearly demonstrates that the flute is far from being a single sine wave as commonly stated. Nevertheless, as the amplitudes of the harmonics in this spectrum decrease rapidly with increasing harmonic number, it is possible to approximate to a reasonable flute tone using only a few harmonics. This is why additive sine wave instruments, which rarely have more than nine harmonics available, are able to provide good flutes, whereas their performance at synthesising almost any other type of tone leaves much to be desired. A glance at the remaining spectra in Fig. 2 shows why. For a subjectively satisfying imitation of these pipe tones, one should aim to embrace all harmonics within a dynamic range of about 60dB. Therefore even the Diapason requires about 15 harmonics and the other two spectra need more. Unless a very large number of harmonics is available in an additive instrument, the only cost-effective way to proceed is with the subtractive approach. (Whilst there are a very few additive instruments that have large numbers, perhaps in excess of one hundred, harmonics available for tonal synthesis, these are expensive experimental developments using advanced microprocessor technology and as yet they are hardly suitable for amateur construction.)

Returning briefly to the imitation of an organ flute stop of the sort illustrated by the spectrum of Fig. 2(a), this type of tone is in some ways the most difficult to simulate in spite of the apparent simplicity of the spectrum. Merely designing a filter to produce the same overall spectral features often produces a tone that seems somewhat dull and lifeless compared to the original, especially on A-B comparison using tape recordings. Ladner (ref.3) made the same point, and it seems that the role of the low-amplitude high-order harmonics is not well understood. Sumner (ref.1) reports that physical features such as the "chimney" in the flute stop of that name are responsible for subtle formant bands in the spectrum, though he does not give further details.

Passing on to the other sounds, where imitation is much easier than for flutes, consider the Diapason. The spectrum shows that the amplitude of the harmonics gradually falls off with increasing harmonic number. The viol, on the other hand, has harmonics that increase in amplitude up to the fourth, whereafter they fall. This is the result of a viol pipe being of smaller scale (narrower) than a diapason pipe of the same length.

Finally the cornopean has a spectrum in which the harmonic energy falls with frequency, though the fall is not in excess of 6dB until harmonics beyond the fifth are encountered. The relative smoothness of this curve compared to the previous three (in which more scatter is apparent) seems to be characteristic of many reed tones.

The four examples of organ pipe spectra represent the four principal categories of organ tone, and there is no reason why essentially the same spectrum should not be used to design filters for several footages, thereby producing a diapason chorus or a reed chorus, etc. The examples given here, together with others in the references cited, give a reasonably broad base of data for the construction of filters.

Electrical Waveforms

In addition to the spectrum of the sound to be simulated, we need that of the source waveform, from which the tone filters are fed. It would be a short and simple matter to present the spectra of commonly-used waveforms at this point, but several other practical problems require discussion first.

Probably the easiest waveform to generate is a square wave. With the ready availability of top-octave synthesiser, dividers and envelope-shapers in integrated-circuit form, a complete generating system of, say, 84 frequencies (seven octaves) can be contained on one card. Unfortunately, the square wave is far from ideal for tone-forming, except in a few cases, because it contains only the odd-numbered harmonics, whose amplitudes decrease at 6dB per octave (see Fig. 3(a)). A square wave cannot therefore be used to derive any of the spectra shown in Fig. 2 as these contain even harmonics. It is, however, suitable for use where tones such as a stopped diapason or a clarinet are required, in whose spectra the odd harmonics are much more prominent than the even ones.

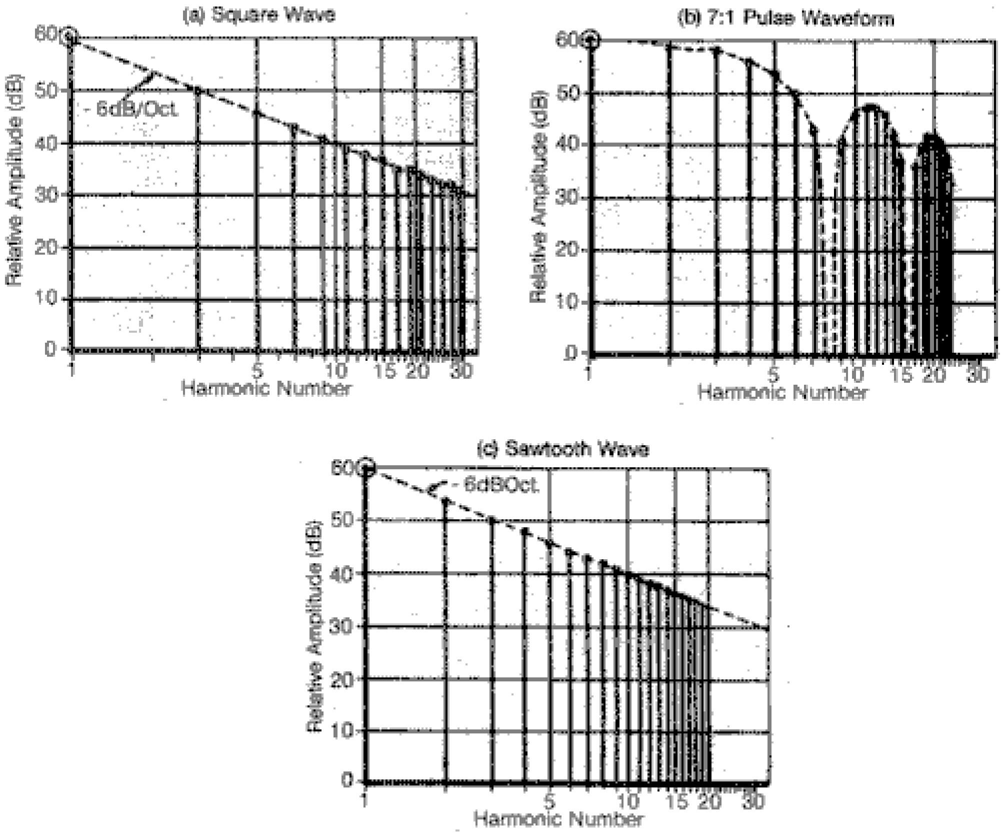

Figure 3. *Easy-keying pulse waveforms such as in (a) or (b) are deficient in harmonic content.*

Harmonic	Square	7:1 pulse	Saw tooth
1	60	60	60
2	-	59	54
3	50	58	50
4	-	56	48
5	46	54	46
6	-	50	45
7	43	43	43
8	-	-	42
9	41	41	41
10	-	46	40
11	39	47	39
12	-	47	38
13	38	46	38
14	-	42	37
15	37	37	37
16	-	-	36
17	35	36	35
18	-	40	35

19	35	42	35
10	-	42	34
21	34	41	34
22	-	38	33
23	33	33	33
24	-	-	33
25	32	33	32
26	-	37	32
27	32	39	32
28	-	39	31
29	31	38	31
30	-	36	31

Table 2. *Harmonic amplitudes of various waveforms in dB corresponding to Fig. 3*

In a square-wave multi-frequency generating system, it is relatively simple to generate pulse waveforms of different mark-space ratios. These possess, in general, both even and odd harmonics and the spectrum of a pulse waveform with a 7:1 mark-space ratio has been discussed by David Ryder (see ref. 5); this special case is of particular interest to those readers who may be building his (sine-wave) organ. The spectrum, Fig. 3(b), shows that certain harmonics are missing. This effect is always obtained with pulse waveforms, including the square wave just discussed: this is merely a "pulse" waveform with a 1:1 mark-space ratio, where the nulls happen to coincide with the even harmonics. Whilst pulse waveforms again have the desirable advantage of simple generation and keying (envelope-shaping), one possible problem concerns the low average energy of a waveform consisting of short pulses. This could give rise to noise difficulties at the output of the tone filters, as these usually introduce considerable insertion loss.

The "classical" waveform that is often used when both odd and even harmonics are required is the sawtooth. This has a spectrum containing all harmonics, whose amplitudes decrease at 6dB per octave (see Fig. 3(c)). Unfortunately, the sawtooth is not particularly economical to generate, and once generated it cannot be keyed by the simple non-linear envelope shapers commonly used for square or pulse waveforms, without introducing distortion. One way to circumvent this limitation is to generate and key pulse waveforms (i.e. square waves), and then combine them with appropriate weights so that a staircase waveform is obtained. This is a good approximation to a sawtooth.

Another approach is to generate and key a single square wave and then convert it to a sawtooth using a discharger circuit of the type shown in Fig. 4. The square wave is first converted to a series of narrow pulses (for example, by differentiation followed by rectification) which are then used to repeatedly discharge the capacitor C through the electronic switch S. In between discharges, the capacitor charges exponentially through R. A linear ramp is obtained if R is replaced by a constant-current source, though for musical purposes this would seldom be required. An exponential ramp produces little significant difference in the spectrum, even at harmonics as high as the 30^{th}. The source voltage *V* can be used to achieve envelope shaping during key attack and release.

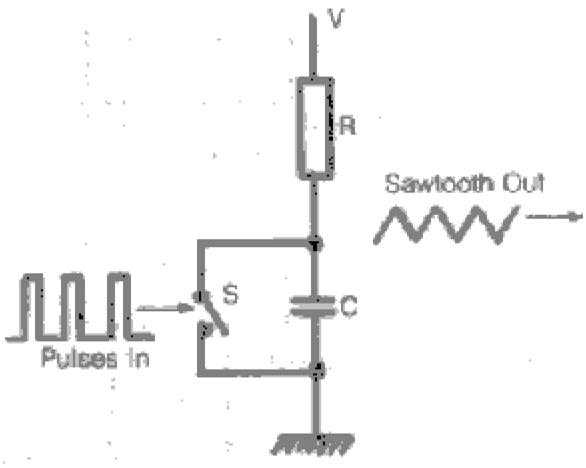

Fig. 4. *It is easier to generate and key a rectangular wave and then convert it to a sawtooth wave than to operate on the sawtooth.*

Several filters are discussed in the next article *[reproduced here as part 2 below]*, all designed assuming the availability of a sawtooth wave to feed them with. This has been chosen for the following reasons:

i) Its spectral structure is simple. Harmonic amplitudes decrease monotonically with increasing frequency rather than in the oscillatory fashion of a pulse spectrum. This results in a filter frequency response that is also much simpler than if a pulse waveform had been used. This is important because of the comparative ease with which an electrical implementation of the filter can be built.

ii) A square wave has already been rejected as being unsuitable for all but a few special tones (though in these cases it is essential).

iii) Sawtooth and square waves are available in the author's instrument. This meant that a subjective judgement could be made as to the effectiveness of a filter design. In particular, it was possible to make A-B comparisons of the electronically-generated sounds against tape recordings of the originals.

Part 2: Design Procedure and Practical Problems

This part of the article derives frequency responses of tone filters for four organ tones, whose acoustic spectra were given in part one. It completes the design procedure, discusses the number of filters needed per stop and the combining of tone colours, and various other practical points.

The frequency response of the required filter is obtained by subtracting the sawtooth spectrum from the relevant organ pipe spectrum. In practice this merely means that the numbers in Table 2, representing the individual harmonic amplitudes, are subtracted one by one from the corresponding numbers in Table 1. The resultant four series of values are presented in Table 3, and graphically in Fig. 5. In all cases the frequency response is represented on a scale that does not indicate absolute frequency but is normalised to the frequency of the first harmonic or fundamental of the original spectra. To implement a real filter circuit one needs to first convert the frequency scale back to true frequency values, which immediately begs the question of which design frequency is chosen for the filter, a subject treated later.

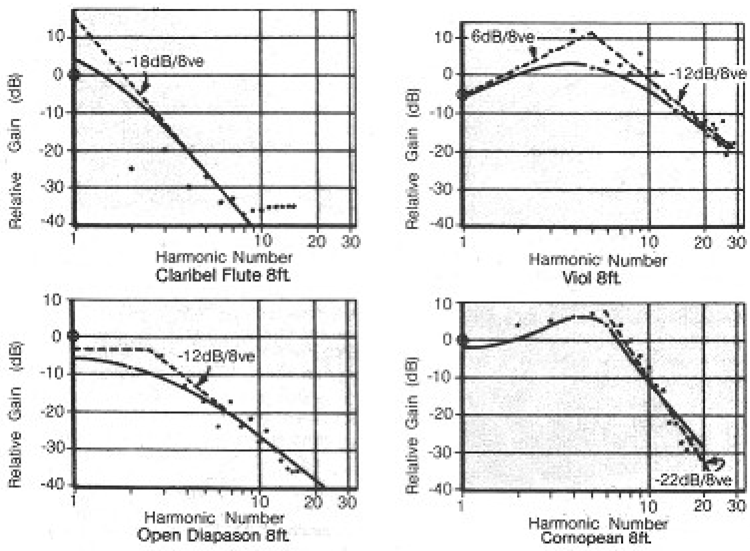

Fig. 5. *Filter frequency response curves for the tones in Fig. 2. Dots represent values of the required response at the harmonic frequencies as in Table 3. Full lines are measured frequency responses of actual filters, broken lines are Bode plots. Responses calculated assuming a sawtooth driving waveform.*

Harmonic	**Claribel Flute**	**Open Diapason**	**Viol**	**Cornopean**
1	0	0	-5	0
2	-25	-8	2	4
3	-20	-5	7	5
4	-30	-13	12	6
5	-27	-17	2	7
6	-34	-24	4	4
7	-33	-17	3	4
8	-37	-24	1	0
9	-36	-22	6	-4
10	-36	-28	2	-7
11	-35	-25	1	-12
12	-35	-30	-4	-13
13	-35	-33	-6	-22
14	-35	-35	-9	-22
15	-35	-36	-10	-27
16	-	-36	-10	-29
17	-	-	-10	-26
18	-	-	-12	-29
19	-	-	-13	-
20	-	-	-12	-
21	-	-	-16	-

22	-	-	-13	-
23	-	-	-14	-
24	-	-	-18	-
25	-	-	-12	-
26	-	-	-21	-
27	-	-	-18	-
28	-	-	-18	-
29	-	-	-	-
30	-	-	-	-

Table 3. *Normalised frequency responses in dB of tone filters for four organ tones assuming a sawtooth drive waveform corresponding to Fig. 5.*

Also shown in Fig. 5 by the full lines are the frequency responses of four actual filters intended to simulate the frequency responses suggested by the discrete points on the four graphs. (The circuit diagrams of these filters are given in Fig. 6 and they are more fully discussed later). It is, of course, permissible to draw the frequency response of a real filter as a continuous curve as the filter has a defined gain/loss at all frequencies in contrast to the experimentally derived points of Table 3, which exist at harmonic frequencies only. An additional feature of Fig. 5 is the presence of broken lines corresponding to Bode plots used in the filter design process. This is discussed later, but for the present a short qualitative discussion of the form of these responses follows as this leads naturally on to filter implementation. It is necessary that the reader is familiar with the amplitude versus frequency response of simple filter sections and (where appropriate) their equivalent Bode plot representations. Particularly important are first, second and third order passive RC networks and parallel resonant (LC) sections.

The claribel flute filter is characterised by a rapid increase in attenuation for the first six or seven harmonics, Fig. 5(a), after which the attenuation remains roughly constant at about 35 dB below the value at the fundamental frequency. After the 15th harmonic no further experimental data are available. The nature of the experimental points in this diagram shows why flutes are among the most difficult tones to emulate. It is difficult to discern a simple trend from the available information, though an interesting feature is that the attenuation of the first few even harmonics is consistently higher than at the adjacent odd harmonic frequencies. This suggests that the flute stop in question consisted of stopped pipes, though it was not possible to confirm this by an examination of the interior of the organ. Whilst a stopped construction is unusual for claribel flutes, this assumption enabled a filter response to be chosen that was based on the first four or five odd harmonic frequencies only; even harmonics were ignored. This filter consisted of a third order passive RC network whose breakpoint was the fundamental frequency. Driven with a sawtooth wave, a reasonably satisfactory flute resulted though the effect when using a square wave was not satisfactory. This is at odds with the strong suggestion from the filter response that odd harmonics ought to predominate. It seems that the proportion of odd to even harmonics is critical for flutes, and experiments with other filter configurations in which particular harmonics were selectively reinforced confirmed this. The simple filter just described makes no attempt to emulate the part of the frequency response suggested by frequencies above the tenth harmonic. Even though such high-order structure may be crucial to the production of a good flute tone as previously discussed, it was found difficult to derive a straightforward way of doing this that also yielded subjectively good results.

Turning now to the open diapason, the response fits a second order Bode plot very nicely, with the break point occurring at a frequency equal to 2.6 times the fundamental. The actual response of such a filter (full curve) fits the experimental points well, with only a few reaching a maximum divergence of 6 dB. Subjectively this simple diapason filter produced entirely acceptable and realistic sounds that were "hard and bright" rather than "dull and woofy". A complete diapason chorus, from a 16 foot double diapason to a three rank mixture, was built up using a total of 32 such filters and the effect had something of the tonal excitement of a similar flue chorus on a pipe organ.

The experimental points for the viol filter suggest a bandpass characteristic, and they are again well approximated by the Bode plot illustrated in the diagram. This consists of a 6 dB/octave rise changing to a 12 dB/octave fall, the transition between the two being at the fifth harmonic of the fundamental. Such a filter has the true response illustrated by the full curve. The subjective verdict on this filter was again favourable, though it was too "stringy" for some tastes. This is possibly due to the fact that this filter was derived from Boner's data (ref. 2) in which measurements were made in a free field with the microphone close to the pipe. In an organ, a viol rank would be placed well inside the organ case and almost certainly inside a swell box. Therefore significant high frequency attenuation would result, with the tone of the pipe sounding less "stringy" to a listener in the auditorium.

Finally, the cornopean data are again strongly suggestive of a bandpass characteristic. In this case the filter was implemented using a parallel resonant circuit tuned to the fifth harmonic with a Q of about 2. To achieve the asymmetry of the response, which rapidly falls off above resonance, a third order RC filter was also used breaking at the eighth harmonic. The reasons for using this particular bandpass filter configuration instead of one akin to the viol are given in the next section. For the present the actual response is seen to fit the experimental values closely. The effect of this filter was a convincing bright reed tone, definitely typical of a cornopean or trumpet rather than of a close-toned tromba or tuba. Again, a family of such filters was built with worthwhile results. The unique tone of an organ reed pipe seems, in part at least, to be due to an harmonic structure that is relatively constant in amplitude up to an harmonic order between the fifth and tenth, depending on the particular tone. After this frequency the amplitude falls off rapidly; this falling characteristic is reflected in the filter response. It is therefore essential to copy the "asymmetrical resonance curve" of the filter, as without the rapid attenuation above resonance the effect is completely synthetic and quite unlike the original.

Hardware Realisation

Filter responses need not be matched exactly to the calculated values at each harmonic frequency of the driving waveform. These points originate from experimental measurements in which a large number of variables, most of them uncontrollable, affect the results such that divergences of a few dB can be neglected provided they are random rather than noticeably systematic.

Flue pipe tones can nearly always be well approximated by the use of a simple passive RC filter:

- Flutes generally need a third order low pass system
- Diapasons generally need a second order low pass system
- Strings generally need a bandpass system

Circuit examples of these types of filter are given in Fig. 6(a), (b) and (c).

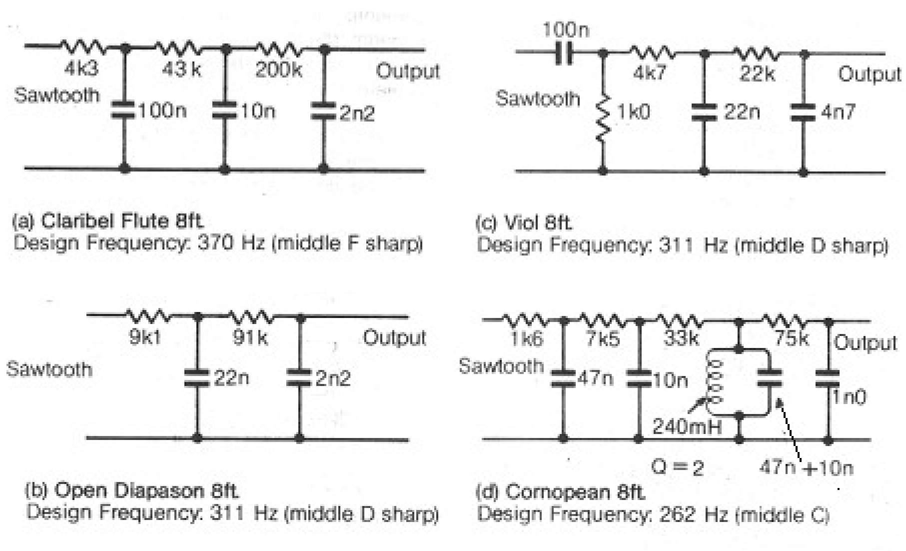

Fig. 6. *Filter circuits giving the required frequency responses of Fig. 5. Inductor in (d) can be realised electronically.*

Reeds can nearly always be well approximated by implementing the asymmetrical bandpass characteristic previously described. It is usually found that the Q of the hump in this bandpass is significantly greater than unity for reeds, whereas for strings (which also require a bandpass) the Q tends to be less than this. Therefore, whilst a simple RC passive bandpass filter can be used for strings as noted above, a resonant circuit or its equivalent is usually necessary for reeds. If a parallel LC circuit is used, as in the example in Fig. 6 (d), the rapid roll-off on the high frequency side of the resonant peak can be achieved by using an additional RC network. In Fig. 6 (d) this network is of third order.

The majority of organ tones are best derived from a sawtooth wave, or one that has both odd and even harmonics. However, there are some important exceptions where a waveform containing only the odd harmonics (e.g. a square wave) is preferable if not actually essential. A partial list of stops where odd harmonics predominate might have names such as stopped diapason, lieblich gedackt, bourdon (all stopped flue pipes), and clarinet, vox humana, cromorne (reed pipes with cylindrical resonators).

These design guidelines just given apply to the filter circuits in Fig. 6. For flue pipe tones, the Bode plot of an appropriate passive network is first matched to the experimental points and then the corresponding filter is implemented. This procedure requires a certain amount of experience and judgement; for the first example turn to the open diapason frequency response in Fig. 5(b). The Bode plot best suited to the experimental data appeared to be a second order system in which there is first a horizontal line (zero slope) followed by a line of slope -12 dB/octave. The breakpoint is the frequency at the point of intersection of the two line segments. The -12 dB/octave part of the response was drawn so that it fitted the slope of the experimental data as well as possible as judged by eye, then the breakpoint was adjusted bearing in mind that the actual response at this frequency will be 6 dB less in amplitude. A breakpoint of 2.6 times the fundamental frequency resulted. The frequency response of the filter is given by the full line in Fig. 5(b) and Fig. 6(b) gives the circuit. This corresponds to the particular form of the Bode plot in that the two sections have the same time constant (RC product) and they are arranged such that they do not mutually load each other. (It is usually possible to avoid buffer amplifiers by choosing the component values to avoid mutual interaction). The circuit was designed for a fundamental sawtooth frequency of 311 Hz, so that each section has a time constant of

$$RC = 10^6 / (2\pi \times 311 \times 2.6)$$

where R is in kohm and C in nF. The question of how to choose the design frequency of the filter is deferred until later as it raises some important practical issues.

The flute filter of Fig. 6(a) was designed in exactly the same way, though in this case the frequency response data of Fig. 5(a) offered less precise guidance as to the form that the Bode plot should take. A third order system was used, matched to the first few odd harmonics for the reasons stated previously. The three time constants were again equal and the three RC sections were again not buffered. The breakpoint was chosen to be the fundamental frequency which in this case was 370 Hz. There would have been little point in using a breakpoint lower in frequency than the fundamental; this would merely have resulted in greater insertion loss with little effect on the tone quality.

For the viol frequency response, Fig. 5(c), there were two segments clearly indicated, forming a Bode plot with slopes -6 dB/octave and -12 dB/octave. The breakpoint turned out to be at the fifth harmonic. This is a simple bandpass filter formed from three RC sections in which one is high pass and the other two low pass. The particularly simple form of the Bode plot means, again, that the time constants are all equal and that the sections must not interact. Such a circuit is shown in Fig. 6(c) and was designed for optimum operation at 311 Hz.

Reed tones generally require bandpass characteristics with Q's not less than 1.5 and often more, which implies the use of circuits such as LC resonant sections. The higher the Q, the more "reedy" the tone and the smaller the frequency range over which the circuit is effective. A Q in excess of three or four is seldom required for the imitation of organ reeds. The cornopean frequency response in Fig. 5(d) has a clearly defined resonance peak at the fifth harmonic, and a Q of about 1.5 is implied by the locus of the experimental points below resonance. To achieve the rapid attenuation above resonance an additional roll-off of about -22 dB/octave starting at the eighth harmonic is also indicated. This result was obtained after a certain amount of juggling with ruler and pencil on the original graph points. The filter constructed used a resonant circuit with a Q of 2 rather than 1.5 because it sounded better, and a roll-off of -18 dB/octave instead of -22 dB/octave for practical reasons. A version of this circuit designed for a 262 Hz sawtooth is shown in Fig. 6 (d), and its frequency response is the full curve in Fig. 5(d). The first two and the final RC sections produce a slope of -18 dB/octave at the eighth harmonic, and the central LC section is responsible for the resonant characteristic.

A parallel tuned circuit has to be driven and terminated so that its Q is not significantly affected by the adjacent circuitry. The terminating impedance can simply be a sufficiently large resistor which in this case is also used as an element of one of the low pass sections. The source resistor feeding the resonant circuit must then be chosen according to the following criteria. It must not appreciably load the preceding RC section nor must it reduce the Q of the resonant circuit. Hence its value must be as high as possible. But the insertion loss of the complete filter is influenced by the value of this source resistor because the effective resistance of the LC section at resonance equals Q^2R where R is the equivalent resistance of the inductor. Hence the source resistor and the LC section itself form a potential divider that controls the amount of signal handed on to the rest of the circuit. For this reason the value of the source resistor should be as low as possible.

The circuit in Fig. 6(d) thus contains a certain amount of compromise, though mainly in the interests of economy. If total component cost is of no account the various sections of the filter can be buffered using active devices thereby easing the design process. Such a course seems scarcely worthwhile when it is possible to approximate the desired response as well as is indicated by Fig. 5(d).

In the interests of simplicity it has so far been implied that the resonant circuit was constructed with a wound inductor. This was not the case since an electronic inductor was synthesised using a simple circuit, Fig. 7. The advantages are that the filter can be readily adjusted until a subjectively optimum effect is produced; it is much cheaper than its wound counterpart, consisting only of two resistors, a small capacitor and a cheap operational amplifier; and it is much less bulky. Design equations are as follows:

$$L = QR_2 / 2\pi f$$

where f is the resonant frequency. L is in Henrys, R in ohms and f in Hz.

$$C = L / R_1R_2$$

C is in Farads, L in Henrys and R_1, R_2 in ohms. Suitable values for R_1 and R_2 are 82k and 1k respectively.

The value of the parallel capacitor C' required to tune the circuit to f is

$$C' = 1 / 4\pi^2 f^2 L$$

C' is in Farads, f in Hz and L in Henrys.

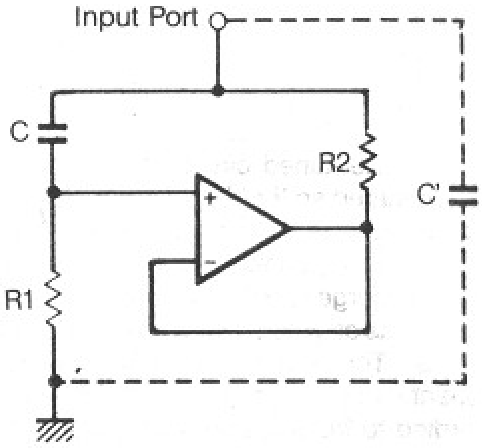

Fig. 7. *Simple electronic inductor realisation using an op-amp where C' is the tuning capacitor.*

The final version of the cornopean filter using an electronic inductor based on the above is in Fig. 8.

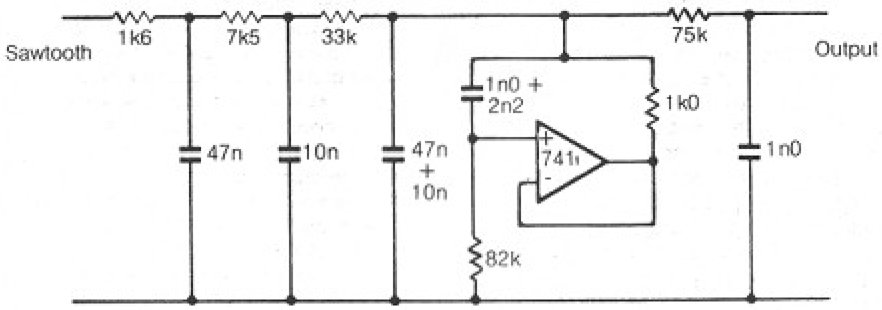

Fig. 8. *Cornopean reed filter using synthesised inductance as alternative to circuit of Fig. 6(d).*

Qualitatively at least, Fig. 5(d) is suggestive of a Q-enhanced Sallen and Key active filter response, though in practice this alone would not achieve the rate of attenuation required above resonance and additional sections would be required. Nevertheless the use of this type of circuit is a distinct possibility instead of the parallel LC circuit used here for those wishing to try it.

How Many Filters per Stop ?

A single tone filter, implemented at one design frequency, will not produce the same tonal effect across an entire keyboard which (in the case of five octaves) might represent a frequency range of 32 : 1. Yet there is evidence in favour of using single filters when cost is paramount: the single filter approach often produces subjectively reasonable results. In my experience this statement is true for flue pipe tones that are simulated using simple low pass filters (flutes and diapasons) where an effective range of three or four octaves can be obtained without difficulty. Beyond this these tones begin to sound unnaturally stringy in the bass and too characterless in the treble, and in addition there is an overall reduction in amplitude when going from low to high notes. This last problem can be mitigated by grading the isolating resistors that are nearly always found in the keying system.

There are two reasons why a single low pass filter has such a large effective frequency range. First, it is easy to show that if the filter characteristic and the source waveform both approximate to linear slopes, not necessarily identical, over a sufficiently large frequency range then the relative harmonic proportions in the output signal remain constant over this range. There is also an overall amplitude variation that can be dealt with as previously described. These approximations are valid for the claribel flute filter and the sawtooth spectrum already discussed, and also for the open diapason though to a lesser extent. The second reason why a single filter is usable in these cases is that to achieve a uniform acoustic output, the pipes in a real diapason or flute stop are scaled so that they have a relatively larger proportion of higher harmonics in the bass than in the treble. This effect is the same as that produced by driving a single flute or diapason filter over a wide frequency range.

With other tones (strings and reeds) an effective range of only two octaves or less is usual because of the more selective frequency response of the filter networks. Beyond this range the effect is artificial, particularly in the bass where the stops sound "sizzly" and thin. There is little that can be done in these cases except to use multiple filters per stop, each one designed to operate over a particular segment of the keyboard. The limiting extreme, of course, is to employ one filter per note, a tour-de-force that has certain advantages in spite of the enormous component count. The advantages stem from the ability to regulate the tone quality and loudness on a note-by-note basis, and the audible "breaks" between filters that can be troublesome when a lesser number is used do not exist. However, entirely adequate results can be achieved using different filters for each half-octave; indeed even this is usually an overkill. I have built a classical instrument of 36 speaking stops all of which employ only four filters, and the result is most satisfactory especially with regard to features such as the sound of reed choruses at the bass end of the keyboard. The method used to combine the outputs of the filters comprising one stop is illustrated in Fig. 9. Each is terminated in a resistor R' that can be used to regulate its amplitude Judicious variation of the relative amplitudes is useful in hiding the breaks between adjacent pairs of filters, yet another psychoacoustic feature of the auditory system that works in our favour. Overall gain variation is provided by making part of the negative feedback resistor R variable.

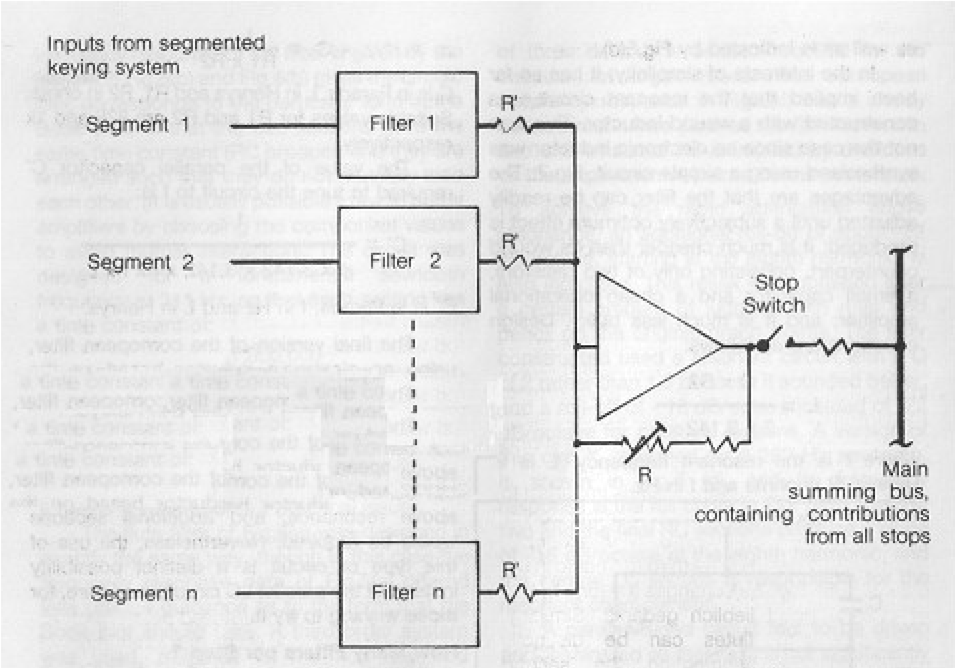

Fig. 9. *Method for combining the outputs of a number of tone filters corresponding to one stop. R' controls the regulation of the stop across the keyboard, R controls the overall amplitude of the stop.*

More Practical Points

All of the filters discussed here must be driven from a low impedance source, in practice a few tens of ohms, and terminated in a high impedance, at least five times greater than the impedances involved in the final stage of the filter. Straightforward operational amplifier techniques are suitable here.

A pronounced change can be imparted to particular tones if only one or two harmonics are selectively augmented. For example, increasing the level of the third harmonic in the claribel flute, Fig. 2(a), changes the tone to that of quite a good lieblich gedackt. Similarly, diapasons and flutes can be distinctly brightened by augmenting the second harmonic. In each case this can be done by borrowing the appropriate sawtooth wave from the multiple keying system that usually exists, in which several frequencies are switched simultaneously for each note. The additional frequencies are combined in the filter simply by providing more input resistors, as in Fig.10. This shows the claribel flute filter together with an additional input which is supplied with a sawtooth wave at the same amplitude as the existing one but at three times the frequency, i.e. at the interval of a twelfth above the note being keyed. The twelfth corresponds to 2 2/3 feet in "footage" nomenclature if the actual stop is of eight foot pitch.

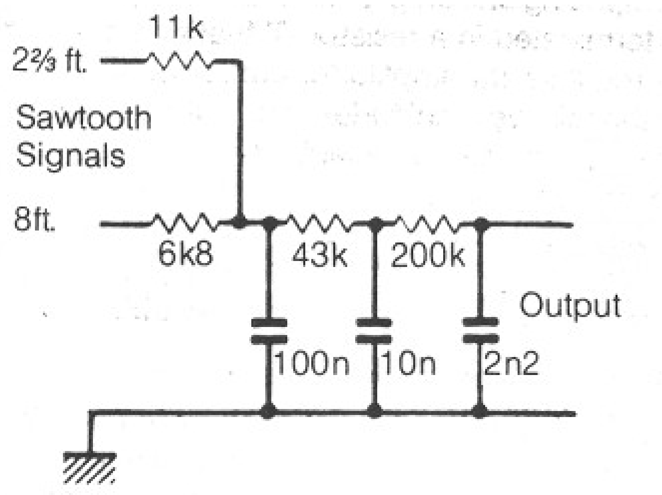

Fig. 10. *Converting the Claribel Flute into a Lieblich Gedackt by augmentation of the third harmonic.*

Three points to remember:

- It is important that the impedance of the sawtooth wave sources should be low, otherwise incorrect summation will result.
- The parallel combination of the various input resistors must approximate to the resistance calculated for the original filter.
- It is not necessary that the frequency relationships between the fundamental and the augmented harmonics be mathematically exact. This makes it possible to borrow the required harmonics from an equally-tempered tone generating system. Such borrowing can only be done to a limited extent; some intervals will be grossly out of tune though in the case of the twelfth the effect is not serious. For all octavely - related intervals, of course, this is irrelevant. A certain amount of trial and error is required to achieve the desired result by this means.

Many organs use a single generator system from which all tones are derived. This means that all stops of the same footage are fed with the same waveform when a given key is depressed, and the various signals emerging from the tone filters are then usually electronically recombined before being amplified and fed to a loudspeaker system. Take care that filters do not introduce inadvertent phase shifts due to the indiscriminate use of inverting amplifiers within the filter itself. Such amplifiers might have been used for buffering purposes. Without first designing the tone forming system as a whole and taking account of detailed points such as this, the ability to add stops one to another will be adversely affected. Buffers are therefore best implemented using non-inverting amplifiers, voltage followers for example. The problem of combining tone colours is further considered below.

The construction of analogue filter circuits for most purposes usually involves close-tolerance components, and the free use of resistors from the E24 range in these articles might imply that the same applies in this case. These values were used simply because they were available; for most purposes resistors from the 5% E12 range should be adequate. Capacitors in active filters, e.g. the synthetic inductor circuits, should be at least 5% but elsewhere 10% should prove satisfactory. The object is not to produce a highly precise scientific instrument but to reproduce musical effects in a context where 3 dB in amplitude (around 30%) is fortunately of little significance.

Combining Stops

Regardless of deliberately introduced phase inversion, filters normally produce a certain amount of phase shift, usually frequency dependent. With a common generator system, in which the same waveform is split into several paths through various filters before being recombined and amplified, there is bound to be a degree of emphasis or attenuation of particular harmonics in the final signal. This has the practical effect that the result of adding stops will be the production of a composite sound that is not necessarily the subjectively expected result of adding the individual tone colours. The effect is most noticeable for stops of the same footage, and if the problem is troublesome then various remedies can be used. The best technique is to have a multi-rank generator system in which there are as many ranks as stops that are likely to be combined. The various ranks are not phase locked to each other but must run independently. Whilst there are various technical problems inherent in this approach, not to mention cost, the chorus effect of the result can rival that of the pipe organ and it is worthwhile if economics allow. The other method, less effective but still expensive, is to retain a single generator system but only allow recombination of the filter outputs to occur acoustically through the use of a multiplicity of sound channels. Electronic "chorus" can also be judiciously applied to each channel to enhance the effect.

The combining problem is sometimes exaggerated, and a cost-effective compromise is obtainable at minimal expense simply by applying a few artistic guidelines when developing the specification of a new instrument. In normal pipe organ registration, that is, the art of selecting stops to achieve a particular tonal effect, it is preferable to minimise the number of stops of the same footage that are used. Even with the pipe organ, which has the ultimate in chorus effects owing to its huge variety of non-synchronised tone sources, it is inartistic to pile tone on tone when one or two carefully chosen stops would suffice. When major tonal build-ups are required this should be achieved by adding stops of different footages, and exactly the same guidelines apply to an electronic organ of whatever sort though particularly if it has a common generator system. In this case the addition of a 4 foot stop to an 8 foot one introduces a new harmonic series that only interferes, in the technical sense, with half as many harmonics in the basic 8 foot tone as would be the case if a second 8 foot stop had been added. The resultant tone is much more realistic in general. The only expense involved in following this principle is that the single generator rank has to be extended upwards by the appropriate number of octaves to cater for the extra upper work present in the stop list, and the keying system is made correspondingly more complex.

It might be thought that adjustable filters can be used in the filter design process to quickly arrive at a subjectively satisfactory result simply by twiddling knobs. A useful configuration, it might be argued, would be a resonance filter module as used in synthesisers in which the tuned frequency and Q are independently variable through the use of state variable techniques. This approach has been eschewed as it represents a return to the total empiricism that negates the design methodology outlined. If it is possible to calculate a frequency response then the starting point should be a filter that approximates this response in a reasonably cost-effective manner. This does not disallow small changes to the prototype circuit to secure a better result, but too much dabbling will quickly lead the ear in a false direction that becomes all too obvious if an A-B comparison with the original sound is subsequently attempted. If it is impossible to achieve a satisfactory simulation of the desired sound then the original experimental data should be suspected as being unreliable, and an attempt to obtain new data should be made.

References

1. Sumner, W L, *The Organ,* Macdonald (London), 1962.

2. Boner, C P, *Acoustic Spectra of Organ Pipes*, Journal of the Acoustical Society of America, July 1938.

3. Ladner, A W, *Analysis and synthesis of musical sounds,* Electronic Engineering, October 1949.

4. Fletcher, H et al, *Quality of organ tones,* Journal of the Acoustical Society of America, March 1963.

5. Ryder, A D, *Electronic organ tone system*, Wireless World, March 1979.