

# The Chatterbox

A simple speech synthesizer for demonstration and amusement

by Ian H. Witten, M.A., M.Sc., Ph.D., M.I.E.E. and Peter H. C. Madams, B.Sc., M.Sc.

Department of Electrical Engineering Science, University of Essex.

The device described is a hand-controlled, electronic model of the acoustic properties of the vocal tract, and was built to illustrate the physiological and acoustic nature of speech. Although designed as a portable demonstration and lecturing tool it makes a fascinating toy for adults and children alike, and has been used as a stimulus for retarded and autistic children. After discussing the nature of speech and the mechanism of electronic speech synthesis, the authors explain the design principles of the Chatterbox and in a later article will give further circuit details and instructions on how to make it talk.

PEOPLE SPEAK by using their vocal chords as a sound source, and making rapid gestures of the articulatory organs (lips, tongue, mouth, etc.). The resulting changes in shape of the vocal tract allow the production of the different sounds that we know as the vowels and consonants of ordinary language. For several years it has been possible to simulate the action of the vocal tract electrically, using a device similar to an electronic organ to produce sounds with the same character as

those of human speech. The first of these "speech synthesizers," built in the early 1950s, comprised many racks of equipment, consumed a lot of power, and cost a great deal of money. Now, however, with the advent of cheap integrated circuits, it has become possible to build simple, compact, and quite inexpensive synthesizers, without sacrificing the ability to produce the full range of speech sounds.

Of course, to make the ever-changing patterns of speech, a synthesizer needs some form of continuously varying control, and just as there are many vocal tract organs involved simultaneously in speaking, so it is necessary to control several parameters of the synthesizer at once. Most speech research laboratories nowadays use a digital computer to manipulate the control signals for their synthesizers. However, for the purposes of informal experiments with speech or just to learn about the sounds we make, a pair of hands will suffice – with the added

advantage that the operator can use his long-standing experience with real speech to mould the sounds into voice-like ones.

By way of illustration of these points, we have built a small, manually controlled speech synthesizer, suitable for home construction – the "Chatterbox" (Fig. 1). In experienced hands it can be encouraged to utter recognizable words and phrases ("hello," "how are you," etc.), while even a complete novice can make it generate a great variety of astonishingly different noises, all of which are immediately recognizable as speech-like. The Chatterbox was originally designed as a portable demonstration and lecturing tool for illustrating the different sounds of speech and how they can be synthesized, but we quickly found that the fascination of artificial speech makes it a successful and compelling toy for adults and children alike. As a hand-controlled, electronic model of the acoustic properties of the vocal tract, it provides a natural feel for the growing science of phonetics – the study of what people do when they are talking and when they are listening to speech.


This first article discusses the nature of speech and the mechanisms of electronic synthesis. The Chatterbox design is described later, with some circuit details, as a concrete example of the implementation of a speech synthesis system.

## The anatomy of speech

The so-called "voiced" sounds of speech – like the sound you make when you say "aaah" – are produced by passing air up from the lungs through the larynx or voicebox, which is situated just behind the Adam's apple. The vocal tract from the larynx to the lips acts as a resonant cavity, amplifying certain frequencies and attenuating others.

The waveform generated by the larynx, however, is not simply sinusoidal. (If it were, the effect of the vocal tract resonances would merely be to give a sine wave of the same frequency but amplified or attenuated according to how close it was to the nearest resonance.) The larynx contains two folds of skin – the vocal cords – which blow apart and flap together again in each cycle of the pitch period. The pitch of a male voice in speech varies from as

Fig. 1. The Chatterbox, showing the controls for hand operation.



low as 20Hz to perhaps 250Hz, with a typical median value of 100Hz. For a female voice, of course, the range is correspondingly higher. The flapping action of the vocal cords gives a waveform which can be approximated by the triangular pulse of Fig. 2. This has a rich spectrum of harmonics, decaying at around 12dB/octave, and each harmonic is affected by the vocal tract resonances.

A simple model of the vocal tract is an organ-pipe-like cylindrical tube with a sound source at one end (the larynx) and open at the other (the lips), as shown in Fig. 3. This has resonances at wavelengths  $4L, 4L/3, 4L/5, \dots$ , where  $L$  is the length of the tube; and these correspond to frequencies  $c/4L, 3c/4L, 5c/4L, \dots$  Hz, where  $c$  is the speed of sound in air. Calculating these frequencies, using a typical figure for the distance between larynx and lips of 17cm, and  $c = 340$  m/sec for the speed of sound, leads to resonances at approximately 500Hz, 1500Hz, 2500Hz, ...

When excited by the harmonic-rich waveform of the larynx, the vocal tract resonances produce energy peaks known as *formants* in the energy spectrum of the speech wave (Fig. 4). The lowest formant, called formant 1, varies from around 200Hz to 1000Hz during speech, the exact range depending on the size of the vocal tract. Formant 2 varies from around 500 to 2500Hz, and formant 3 from around 1500 to 3500Hz.

Of course, speech is not a static phenomenon. The organ-pipe model describes the speech spectrum during a continuously held vowel with the mouth in a neutral position such as for "aaah." But in real speech the tongue and lips are in continuous motion, altering the shape of the vocal tract and hence the positions of the resonances. It is as if the organ-pipe were being squeezed and expanded in different places all the time. Say "ee" as in "heed" and notice how close your tongue is to the roof of your mouth, causing a constriction near the front of the vocal cavity.

Linguists and speech engineers use a frequency analyser called a sound spectrograph to make a three-parameter plot of the variation of the speech energy spectrum with time. Fig. 5 shows a spectrogram of the utterance "go away." Frequency is given on the vertical axis, and bands are shown at the beginning to indicate the scale. Time is plotted horizontally, and energy is given by the darkness of any particular area. The lower few formants can be seen as dark bands extending horizontally, and they are in continuous motion. Notice that in the neutral first vowel of "away," the formant frequencies approximate the 500Hz, 1500Hz and 2500Hz that we calculated earlier. (In fact, formant 2 is around 1250Hz and formant 3 around 2300Hz.) The fine vertical striations in the spectrogram correspond to single openings of the




Fig. 2. Approximate waveform produced by the larynx.




Fig. 3. Resonances in the organ-pipe model of the vocal tract.

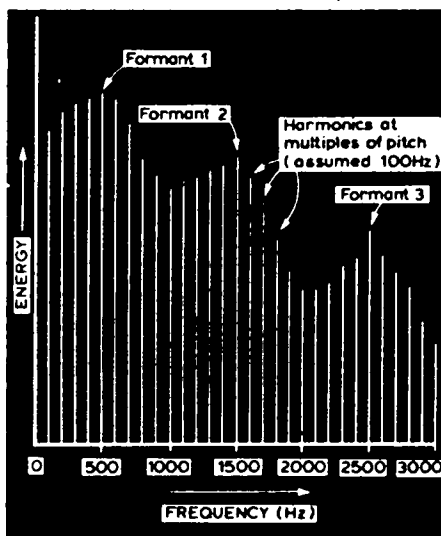


Fig. 4. The energy spectrum of speech, showing three formants.

Fig 5. Spectrogram of the utterance "go away."




Table 1. The vowels and their formant frequencies

Vowel name	Example of use	F1 (Hz)	F2 (Hz)
UH	ab(ove)	500	1500
A	bud	700	1250
E	bed	550	1950
I	bid	350	2100
O	bod	600	900
U	good	400	950
AA	bad	750	1750
EE	bead	300	2250
ER	bird	600	1400
UU	brood	300	950
AR	bard	700	1100
AW	board	450	750

vocal chords. Of course, the pitch is continuously changing throughout an utterance, and this can be seen on the spectrogram by the differences in spacing of the striations. Pitch change, or *intonation*, is singularly important in lending naturalness to speech.

On a spectrogram, a continuously held vowel shows up as a static energy spectrum. But beware - what we call a vowel in everyday language is not the same thing as a "vowel" in phonetic terms. Say "I" and feel how the tongue moves continuously while you're speaking. Technically, this is a *diphthong* or slide between two vowel positions, and not a single vowel. And there are many more phonetically different vowel sounds than the a, e, i, o and u that we normally think of. The words "hood" and "mood" have different vowels, for example, as do "head" and "mead." The principal acoustic difference between the various vowel sounds is in the frequencies of the first two formants. Table 1 gives a list of the English vowels, with a one- or two-character name for each, an example




Fig. 6. Spectrogram of "high altitude jets whizz past screaming."

word, and the two formant frequencies which characterize the sound.

Speech involves other sounds, different from the voiced ones that we have been discussing so far. When you whisper, the folds of the larynx are held slightly apart so that the air passing between them becomes turbulent, causing a noisy excitation of the resonant cavity. The formant peaks are still present, superimposed on the noise. Such "aspirated" sounds occur in the "h" of "hello," and for a very short time after the lips are opened at the beginning of "pit."

Constrictions made in the mouth produce hissy noises such as "ss," "sh," and "f." For example, in "ss" the tip of the tongue is high up, very close to the roof of the mouth. Turbulent air passing through this constriction causes a random noise excitation. For "sh," the tongue is flattened close to the roof of the mouth, in a position rather similar to that for "ee" but with a slightly narrower constriction, while "f" is produced with the upper teeth and lower lip. If the larynx is vibrating as well we get the corresponding voiced sounds "z," the "zh" in "azure," and "v." Because they are made near the front of the mouth, the resonances of the vocal tract have little effect on these hissy sounds. The complicated acoustic effects of noisy excitations in speech can be seen in the spectrogram Fig. 6 of "high altitude jets whizz past screaming."

### Speech synthesis and synthesizers

The idea of artificial speech has always fascinated man. The first genuine talking machine appears to have been demonstrated in 1791 by one Baron von Kempelen, who used bellows to inject sound into a leather tube which modeled the vocal tract, and was deformable with the hands to imitate the different vowel sounds. Progress continued sporadically (one notable achievement

being Alexander Graham Bell's encouraging his pet dog to talk by manipulating its vocal tract by hand while the dog growled), until the need for bandwidth reduction for efficient use of communication channels in the 1940s and 1950s stimulated serious research on the acoustic nature of the speech signal. Since then, the advent of widely-available real-time computing power has encouraged work on speech synthesis under computer control, and the difficult problems of pronunciation, speech rhythms, and intonation are currently being tackled to exploit this novel and effective computer output medium.

In order to simulate electrically the resonating action of the vocal tract on the sound generated by the larynx, a waveform generator and several resonant filters in cascade are needed. Varying the frequency and amplitude of the sound source simulates changes in the pitch and loudness of the speech, and different vowels can be made by adjusting the positions of the resonances appropriately.

Further analysis of the organ-pipe model reveals that simple second-order

resonators with unity d.c. gain are appropriate filters, with the slightly unusual requirement that the bandwidths should remain constant as the resonant frequencies are altered, producing sharper resonances at higher frequencies (Fig. 7). The phase response of the filters is not important. Such resonators can be achieved with simple active filter circuits.

Although vocal tracts, like organ pipes, have an indefinite number of resonances, in practice only a few filters are employed in the chain (Fig. 8). Most existing synthesizers simulate four or five formants, of which typically only the first three have controllable resonance positions. In fact, two formant filters are sufficient to generate most vowel-like speech sounds: a third is especially useful in distinguishing the "r" in "rice" from the "l" in "lice." Omitting the higher resonances means that some compensation filter needs to be introduced to give spectral lift at higher frequencies.

Whispery sounds can be synthesized by injecting noise into the chain of formant filters, instead of the harmonic-rich pulse of the waveform generator. For the sibilant sounds made at the front of the mouth, the noise should not be injected into the formant chain, but instead passed through a separate high-pass resonance whose centre frequency can be controlled to give the sounds "f," "sh," and "ss."

These considerations lead to the block diagram of Fig. 9. A synthesis system similar to this was invented in 1951 by Walter Lawrence, and he called it PAT - Parametric Artificial Talker. The eight circled numbers represent parameters of the system, and if they are varied appropriately, it can be persuaded to give a respectable imitation of almost any speech utterance. For example, the parameter tracks for "six" are shown in Fig. 10 as a set of eight graphs. You can see the onset of the




Fig. 7. Amplitude profiles of the resonant filter.

hissy sound at the beginning and end (parameter 5), and the amplitude of voicing (parameter 1) come on for the "i" and go off again before the "x." The pitch (parameter 0) is falling slowly throughout the utterance.

Naturally, storage of the parameter tracks presents some problems. In the earliest version of PAT, eight parameter-versus-time graphs were painted on a glass slide, which was scanned photo-electrically to read off the parameter values. Lawrence was fond of disconnecting the pitch parameter and controlling it directly with a potentiometer. One of the utterances for which he had prepared a glass slide was "What did you say before that?" and he could manipulate the pitch by hand to change the emphasis,

getting "What did you say before that?", "What did you say before that?" and so on. Even artificial singing proved possible! In fact, the earliest computer conversation on record was between PAT and a Swedish synthesizer, called "Ove." The inventors, Gunnar Fant and Walter Lawrence, stood by their

machines on a stage at an international acoustics meeting, Fant with a small transistorized table-top box, and Lawrence with several great racks of valve-based equipment. The conversation went  
Ove: "How are you?"  
PAT: "What did you say before that?"

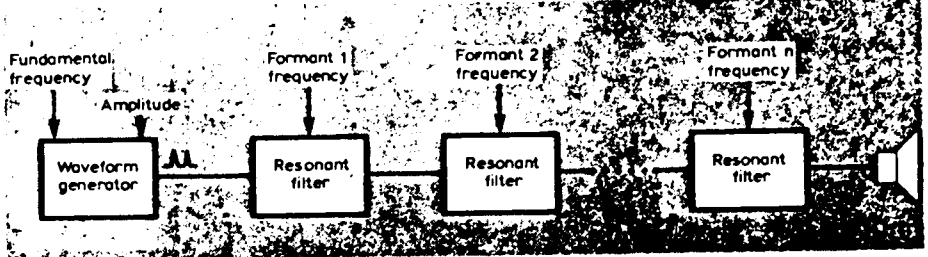


Fig. 8. Simulating the resonance action of the vocal tract.

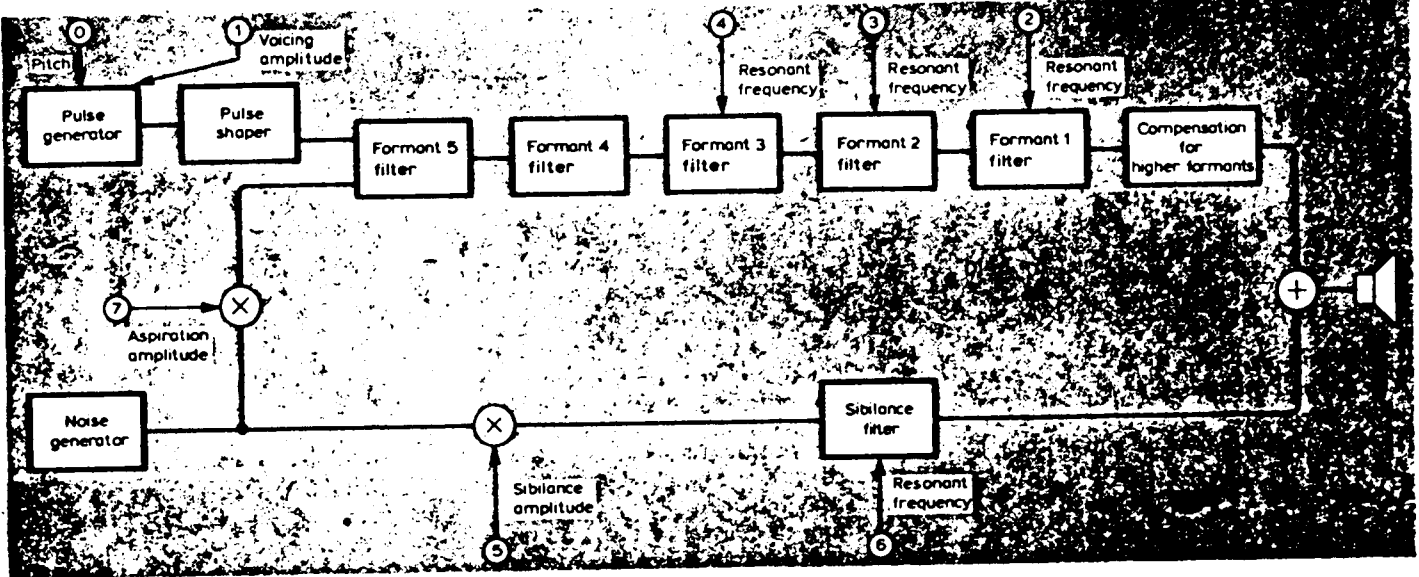



Fig. 9. Block diagram of PAT, the Parametric Artificial Talker.

Fig. 10. Parameter tracks for "six" with the Parametric Artificial Talker.



Ove: "I love you."  
PAT: "What did you say before that?"  
Ove: "I love you."  
At which point, PAT burst into song (no prizes for guessing the words!).

To obtain good speech, the best way of getting parameter tracks is to derive them from spectrograms of human utterances. Although this is a tiresome and time-consuming process, it gives the synthesizer a chance to reproduce the precise acoustic quality of the original speech. However, the parameter tracks of Fig. 10 are stylized: they don't come from a human utterance. In fact they were generated by a computer programme from the input "S I K S", a phonetic transcription of the word. This programme has direct control over the parameters of a hardware synthesizer through a computer interface, and will attempt to speak any utterance that is entered in phonetic format. In practice, the most difficult parameter to control in a convincing way is pitch. The intonation of speech is subtle, and evades classification into a form that a computer programme can handle.

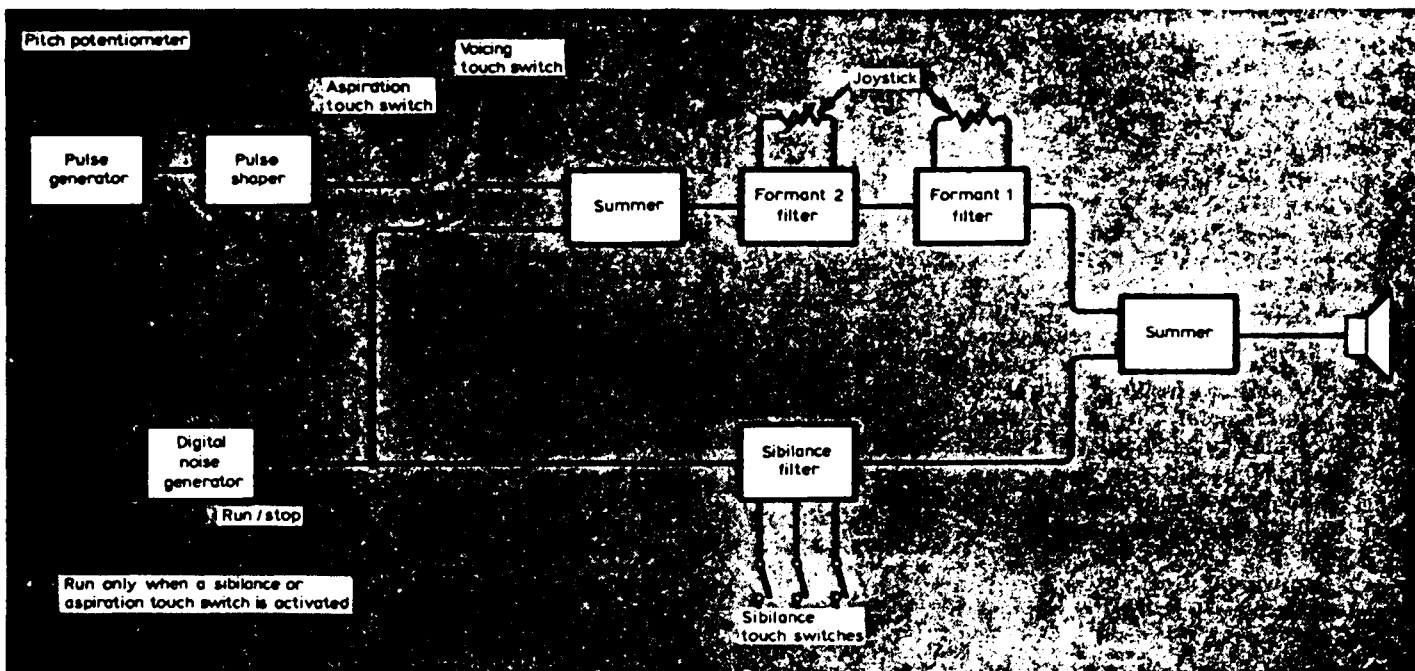


Fig. 11. Chatterbox block diagram.

The Chatterbox, however, avoids these difficulties of computer control by using a person to manipulate the parameters. Generating natural-sounding intonation is easy for people, and this turns out to be true even if they have to use their hands rather than their vocal tracts to control the pitch.

### System design of Chatterbox

**The manual controls.** Hands were never intended to speak! In your vocal tract, separate muscles control a multitude of parameters of the system simultaneously in order to produce speech sounds. Pitch is controlled by the vocal cord tension, amplitude by the lung pressure, and vowel quality by the many dimensions of movement of the tongue, teeth and lips. The greatest challenge of the Chatterbox design was the engineering of the man-machine interface: it is difficult physically to find enough degrees of freedom to control it with the hands. In fact, we even considered using arm and leg movements in addition to hands, but felt that these detracted from the neatness and compactness of the toy.

A single X-Y control is used to vary the two formant filter frequencies. Two models of Chatterbox have been designed and constructed, one with a joystick control and the other with a stylus and resistive plastic pad instead. Fortunately, the recent popularity of quadraphonic audio systems means that it is quite easy to get hold of a compact joystick assembly designed as a quad balance control. Ours had two tracking potentiometers for each direction of motion, and we took advantage of this in designing the formant filters. The alternative stylus arrangement gives a two-dimensional position indication by injecting a current from the stylus tip into a uniformly resistive plastic sheet, and monitoring the current from the sides of the sheet. Since it is fairly difficult to lay hands on a suit-

able resistive sheet for the stylus model, and the circuitry to take advantage of this is more complex anyway, we will describe only the joystick version here.

The pitch is varied by a potentiometer. We decided, after some experimentation, that linear rather than rotary control feels more natural, so a slider potentiometer is used. This is operated with one hand while the other directs the joystick.

Turning to control of the volume of the sound, it transpires that contrary to intuitive expectations, it is not important to provide variation of the amplitude of the voice source, apart from the obvious necessity to switch it on and off. Different vowel sounds do have different amplitudes, of course, because of different degrees of mouth opening (compare the vowels in "mad" and "mood", for example). However, this is taken care of by the formant filters: resonances for "mood" will naturally produce a weaker sound than in "mad" because they occur at lower frequencies, and the constant bandwidth property of the filters gives them less amplification (that is, lower Q) at lower resonance frequencies. The amplitude of the sound produced by the vocal cords corresponds more to vocal effort than to loudness, and this is more or less constant for the great majority of speech sounds. Hence we use a simple switch to turn the voicing on and off.

The hissy sounds pose the most difficult control problem. Aspiration (whispering) can be treated just like the voicing amplitude: we need only be able to turn it on and off. The joystick formant control can then be used to whisper different vowel sounds. However, constrictions in the front of the mouth must be treated separately, for here the sound type ("ss", "sh" and "f")

needs to be controlled independently of the voicing, so that the counterparts "z", "zh" and "v" can be produced simply by superimposing a normal vowel-like sound. We opted for separate switches for these three noises, instead of an analogue control which would simulate the tongue positions more accurately. These are operated by the same hand that manipulates the pitch potentiometer, as is the aspiration switch. This arrangement is not parti-

### Further reading

#### The anatomy of speech

The "source-filter model of speech production," which separates the sound source (larynx) from the filtering operations of the vocal tract, was treated most comprehensively by Gunnar Fant in *Acoustic theory of speech production*, 1960.

The sound spectrograph was developed in 1946 by Koenig, Dunn, and Lacey ("The sound spectrograph," *Journal of the Acoustical Society of America*, vol. 18, pp. 19-49) and is described, with hundreds of spectrograms, by Potter, Kopp and Green (*Visible Speech*, 1947).

#### Speech synthesis and synthesizers

A classic book is *Speech analysis, synthesis and perception*, by James Flanagan of Bell Laboratories in the USA (1965, revised 1972). There is a book of collected papers on speech synthesis by Flanagan and Larry Rabiner called *Speech synthesis* (1973). A British contribution is *Speech synthesis* by John Holmes of the Government Joint Speech Research Unit (1972). Walter Lawrence wrote "The synthesis of speech from signals which have a low information rate" (in *Communication theory*, edited by W. Jackson, pp. 460-469) when he invented PAT in 1953 in the Government Signals Research and Development Establishment.

cularly easy to use, but since pitch control is unimportant during hissy sounds, it is possible to share the pitch hand satisfactorily between all these functions.

It is essential, however, that the voicing on/off switch is easy to operate while complicated pitch movements are being made, because the moment of onset and offset must be timed precisely, without disrupting the smooth flow of intonation. In the pad-and-stylus Chatterbox, it is possible to detect electrically when the stylus is in contact with the pad, and this is used to turn on the voice. A similar arrangement could be made in the joystick model if the act of grasping the joystick were detected, but this would involve modification of the joystick assembly. We opted instead to site a switch contact where it could easily be reached with the heel of the hand that operates the joystick.

All switches on the Chatterbox are touch switches, and work by detecting the skin resistance when two adjacent contacts are touched together. These are much easier to use — and cheaper to

build! — than, say, pressure operated microswitches.

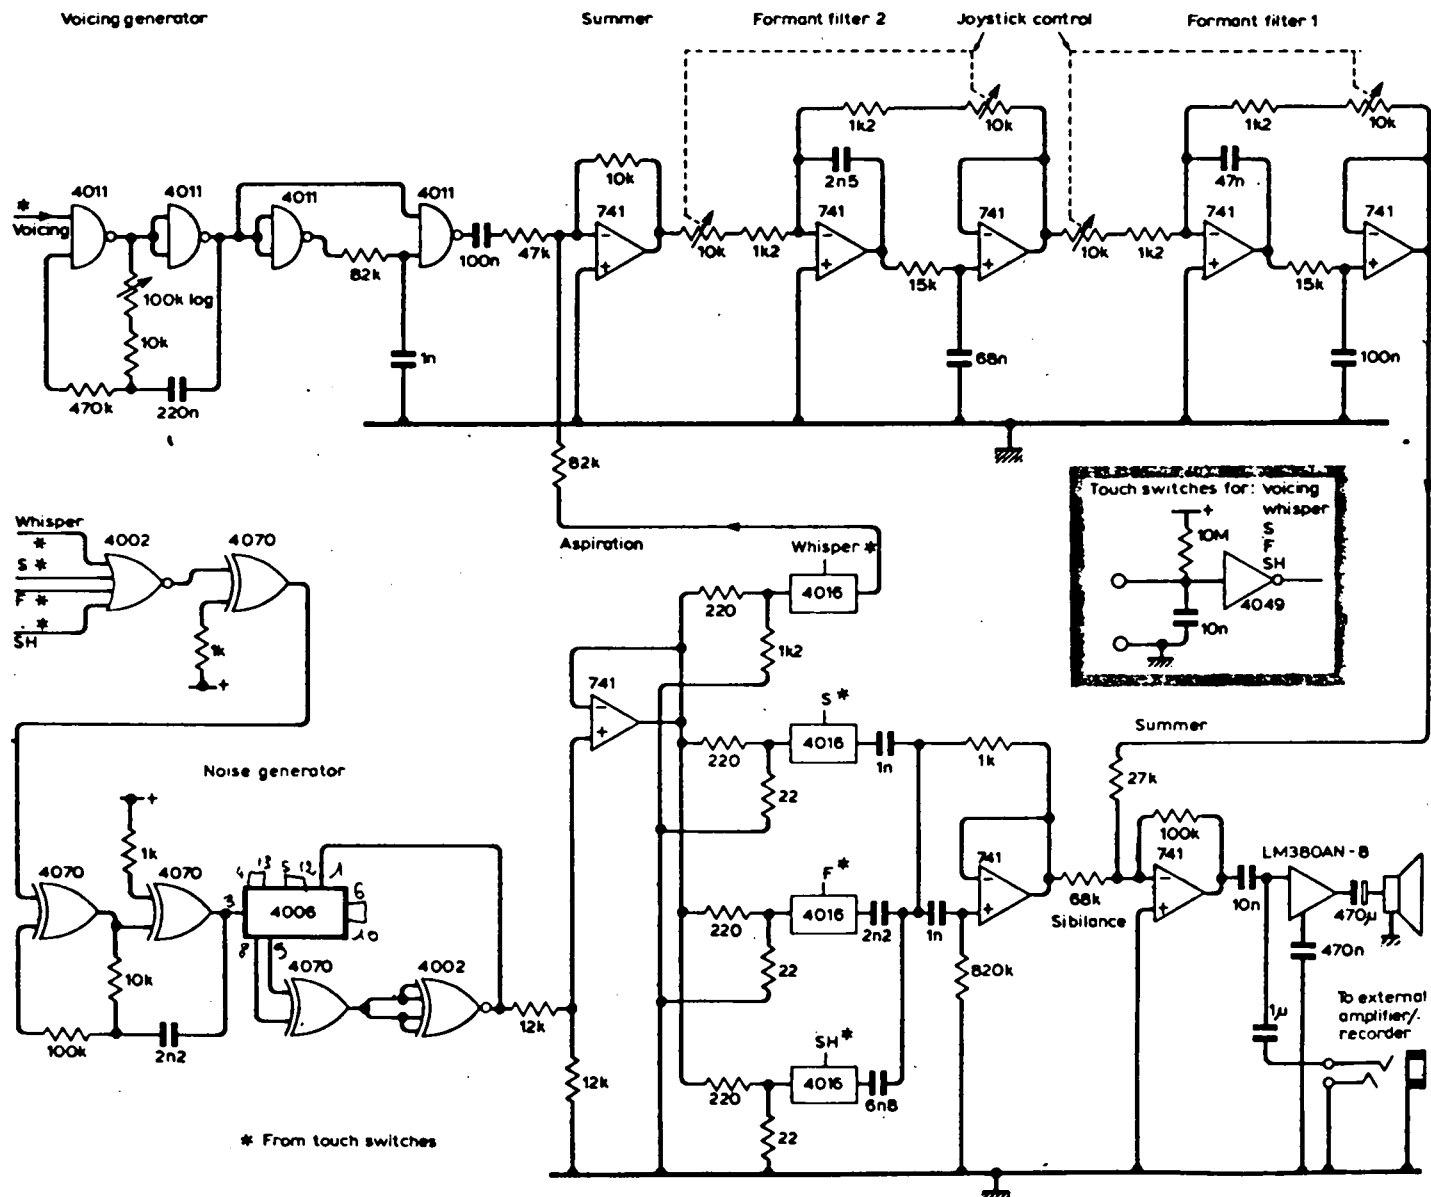
**The overall system.** The Chatterbox is essentially a simplified version of Lawrence's original PAT. As shown in Fig. 11, it consists of two parallel signal paths, a voicing/aspiration path (top of diagram) and a sibilance path (bottom). Two formant filters form the upper path, each controlled by one direction of the joystick. These can be excited by a simulated larynx pulse, produced by a variable-frequency impulse generator, or by a noise source (for aspiration). We use a digital pseudo-random generator implemented by a feedback shift register with exclusive-OR feedback. The same noise source drives the lower sibilance path, which includes a high-pass resonance to give the noise an appropriate colouration. The position of the resonance is controllable to three places by touch-switches. A full circuit diagram of the system is shown in Fig.

12, and the various sections of this will be explained here and next month.

**Source for voiced sounds, and mixer.** A simple circuit using c.m.o.s. gates forms the basis of the voicing waveform generator. This circuit is used because of its simplicity and low cost. The 100kΩ potentiometer varies the frequency of the oscillation. The output from the oscillator is inverted, delayed, and "ANDed" with itself. This produces a train of spikes which is the required harmonic-rich voicing waveform. The circuit is turned on and off by a logic signal from one of the touch switches shown in Fig. 12. The period can be adjusted from a very low value, 20Hz (good for sound effects and "creaky voice"), up to 200Hz. The harmonic content of the spiky waveform is very high and provides a suitable excitation for the formant filters. The 741 amplifier that follows this circuit is used both to add a signal from the noise generator to produce aspirated sounds and to adjust the amplitude of both sources.

Fig. 12. Full circuit diagram of the Chatterbox.

To be continued



# The Chatterbox — 2

Circuit details, construction and use

by Ian H. Witten M.A., Ph.D., M.I.E.E. and Peter H. C. Madams, B.Sc., M.Sc.  
Department of Electrical Engineering Science, University of Essex

Last month's article discussed the processes of electronic speech synthesis and outlined the general design principles of the Chatterbox, giving a complete circuit diagram. The authors now conclude with further description of the circuitry, notes on construction and advice on how to operate the synthesizer.

**Noise generator.** The digital feedback shift register, which is used as the noise source for the Chatterbox, is shown in Fig. 13(a). The final output from the register is exclusive-ORed with an intermediate output, and fed back to the beginning. Because it has only a finite number of states, this configuration will generate a repeating — and hence non-random — string of bits; however, if the intermediate feedback point and the number of bits, say  $N$ , in the register are carefully chosen, a maximal-length sequence of  $2^N - 1$  bits is obtained before repetition begins. We chose the 4006 c.m.o.s. 18-bit shift register chip to implement the generator because some intermediate outputs are available (see Fig. 13(b)) and it is much cheaper per bit than full parallel-output registers.

Arranging 4006 chips into a low-cost maximal-length feedback shift register presents some problems. The clock rate needs to be at least 20kHz to ensure white noise. Clocked at this rate, a 16-bit maximal-length register repeats every three seconds or so, and this repetition in the noise is unfortunately just noticeable. Although 17-bit and 18-bit maximal-length registers do exist, it is just not possible to configure the several small shift registers of the 4006 in a way which makes the appropriate bit accessible. However, a computer simulation has shown that an 18-bit register with feedback from stage 17, although not maximal length, produces 253921 bits before repeating the cycle. This is only slightly less than  $2^{18} - 1$  (262143), and so the sequence generated is very close to maximal-length. With the 22.5kHz clock that is used, the sequence repeats every 11 seconds.

It is important to start a feedback shift register in the correct state. For example, the one in Fig. 13(a) will certainly not produce noise if started with each stage containing zero, since the feedback bit will be zero also. Unfortunately, when power is turned on the c.m.o.s. shift register that we used in-





Fig. 15. Distortion of vowel positions due to square root law of filter resonances.

variably comes up in the all-zeros state. Hence a spare gate is placed in the feedback path acting as an inverter, so that the all-ones state is the one which cycles indefinitely without producing noise. The all-zeros state leads into the repetition sequence of 253921 cycles.

**Formant filters.** The formant filters form the heart of an analogue speech synthesizer like the Chatterbox. They should be constant-bandwidth, constant d.c. gain, second-order resonators with centre frequency controllable over approximately half a decade range. In order to keep the cost of the device low, we used the two operational amplifier active filter configuration of Fig. 14 instead of the more common ring-of-three design, which incidentally is




Fig. 13. (a)  $N$ -bit exclusive-OR feedback shift register; (b) shift registers in the c.m.o.s. 4006 i.c.




Fig. 14. Formant filter.

rather easier to control. The transfer function of this can be shown to be

$$H(s) = \frac{-1/C_1 R_1 C_2 R_2}{s^2 + \frac{1}{C_2 R_2} s + \frac{1}{C_1 R_1' C_2 R_2}}$$

which characterizes it as a low-pass resonance with d.c. gain of  $R_1'/R_1$ , bandwidth of  $1/2\pi C_2 R_2$  Hz, and centre frequency of  $1/2\pi \sqrt{C_1 R_1' C_2 R_2}$  Hz. By tracking  $R_1'$  with  $R_1$ , we can ensure that

the d.c. gain remains constant, and the centre frequency follows  $1/\sqrt{R_1}$ . Moreover, neither is especially sensitive to slight departures from exact tracking of  $R_1'$  with  $R_1$ .

The inverse square root variation of formant frequency with  $R_1$  caused us some concern. It is important to ensure that the joystick has an audible effect on the sound quality right across its operating range, so that control is not concentrated in one corner.

Suppose we let  $k$  travel from 0 to 1 to represent the position of the joystick in one dimension. The potentiometer is linear; suppose its resistance swings from  $R_A$  to  $R_B$ . Then at position  $k$ , the resistance is  $R = (1-k)R_A + kR_B$ , and the frequency of the resonance is proportional to  $1/\sqrt{R}$ . It is easy to show from this that if the resonant frequency travels from  $f_0$  to  $f_1$ , its value at position  $k$  is

$$\frac{f_0}{1-k[(f_1/f_0)^2]}$$

Now we can plot the positions of the vowels on a two-dimensional plane, with the dimensions representing the degrees of freedom of the joystick. Fig. 15 shows a comparison between the vowel positions for a hypothetical linearly controlled filter ( $f \sim 1/(k + \text{const})$ ) – the vowels are labelled in the positions they would occupy for this – and the filter we are proposing. The distortion due to the square-root law is shown by arrows. The diagram is obtained using standard formant frequencies for a male voice, and ranges




Fig. 16. Sibilance filter.

of 200-750Hz and 750-2250Hz for formants 1 and 2 respectively.

The effect of the transformation is to crowd the vowels towards one corner of the plane, which is disadvantageous – distinguishing between vowels would clearly be easier if they were maximally separated. However, the crowding is not severe, and the simplicity of the two-amplifier filter compared with others having more suitable relationships between resonant frequency and potentiometer position was considered to outweigh the disadvantage of uneven vowel distribution.


The bandwidths of the formant filters must be chosen carefully. Published figures for formant bandwidths are surprisingly low – around 50 to 100Hz. However, low bandwidth gives a high magnification factor  $Q$ , especially in formant 2 where the resonant frequen-

Fig. 17. Inside the Chatterbox.

cies are higher, and since the gain in the first stage of the filter depends on  $Q$ , this is undesirable, because limiting will occur unless the signal levels are extremely low. In practice, we have found that the audible effect of synthesized speech does not depend critically on low formant bandwidths, and we chose to make the first formant bandwidth 110Hz and increase the second formant bandwidth slightly to 160Hz.

The component values used in the prototype Chatterbox are shown in Table 2. They give a range on formant 1 which spans slightly lower frequencies than required by the vowels of Table 1, and a formant 2 range which is rather higher than that required. This is because if the formant ranges intersect, or come close to intersection, the combined amplification of both filters can cause limiting to occur in the second filter. The problem cannot be avoided simply by turning down gains, for then the amplification when the formant frequencies are separated is so small that the signal gets lost.

**Sibilance filter.** To make the sibilant sounds "ss," "f" and "sh," the signal generated by the noise source must be filtered and attenuated. A second-order high-pass resonance is an appropriate filter, with the position of the resonance determining the type of sibilance. It is necessary also to make "f" a much weaker sound than the other two, so we sought a filter where we could change the attenuation at the same time as controlling the resonant frequency. The circuit of Fig. 16 does nicely. Capacitor





**Table 2. Component values for the formant filters**

	Formant 1	Formant 2
$R_1, R_1'$	10k $\Omega$ linear potentiometer in series with 1.2k $\Omega$	10k $\Omega$ linear potentiometer in series with 1.2k $\Omega$
$C_1$	47nF	2500pF
$R_2$	15k $\Omega$	15k $\Omega$
$C_2$	100nF	68nF
calculated bandwidth	110Hz	160Hz
calculated range of resonant frequencies	180Hz-550Hz	940Hz-2850Hz
range required by vowels of Table 1	300Hz-750Hz	750Hz-2250Hz

**Table 3. Component values for the sibilance filter**

	s	f	sh
$R_A$	220	220	220
$R_B$	22	22	22
$C_1$	1000pF	2200pF	6800pF
$R_1$	1k $\Omega$	1k $\Omega$	1k $\Omega$
$C_2$	820k $\Omega$	820k $\Omega$	820k $\Omega$
$R_2$	1000pF	1000pF	1000pF
calculated gain	1/10	1/10	1/10
calculated bandwidth	390Hz	280Hz	220Hz
calculated resonant frequency	5560Hz	3750Hz	2130Hz

$C_1$  controls the resonant frequency, and  $R_A/R_B$  controls the attenuation. The transfer function is

$$H(s) = \frac{s^2}{(1 + \frac{R_A}{R_B} + \frac{R_A}{R_2})s^2 + [(1 + \frac{R_A}{R_B}) \cdot \frac{1}{R_2} \cdot (\frac{1}{C_1} + \frac{1}{C_2}) + \frac{R_A}{R_1 C_2 R_2}]s + \frac{1 + R_A/R_B}{R_1 C_1 R_2 C_2}}$$

and if we choose  $R_B \ll R_A, R_B \ll R_2, R_2 \ll R_1$ , and  $R_B \ll R_1 C_2 / C_1$ , this has passband gain of  $R_B/R_A$ , bandwidth of  $(C_1 + C_2) / 2\pi R_2 C_1 C_2$ , and centre frequency of  $1/2\pi \sqrt{C_1 R_1 C_2 R_2}$ .

We had to experiment with the centre frequency and bandwidth to find suitable sibilant sounds. We finally settled on resonances at around 2100Hz for "sh," 3700Hz for "f" and 5600Hz for "ss," with fairly high Qs - compared to the resonant frequencies - bandwidths of 200 to 400Hz. These are on the verge of what is realizable with the circuit before capacitor values become ridiculously small and resistor values ridiculously high. Table 3 shows the component values used.

The different sibilances are obtained by switching different  $R_A, R_B$ , and  $C_1$  networks into the input of the filter. A 4016 c.m.o.s. analogue gate is employed just before  $C_1$ , controlled by the appropriate noise touch switch, and the other side of  $C_1$  in each of the three networks is commoned into the remainder of the filter. Fig. 12 shows the details of control. Because of the low input impedances of the  $R_A, R_B, C_1$  networks, the output from the noise generator is buffered with a unity-gain amplifier stage.

**Touch switches.** There are five touch switches on the Chatterbox, for voicing, aspiration, and three sibilant sounds, S, F, and SH. The voicing switch, which is operated by the heel of the hand that grasps the joystick, is replicated on both sides of the box to cater for left- and right-handed people.

The touch switches operate by detecting the skin resistance when two adjacent contacts are touched together. They are made possible by the extremely high input impedance of c.m.o.s. gates. Since only a tiny current is drawn by the gate, an extremely high pull-up resistance (10M $\Omega$ ) can be used to keep its input asserted. Then even a high skin resistance is able to overcome the pull-up resistor and bring the gate

input to ground. In fact, the conductance of people's skins varies quite a lot. Although the values shown work for

most people, if the touch-switches fail to work for you just moisten the skin a little.

The whisper and sibilant controls are "ORed" together and the result turns off the clock that drives the noise generator, so that if none of them is operated, no noise gets through to distract attention from the voiced sounds.

**Construction**

The prototype model was made in a box 19cm x 11cm x 6cm, with the joystick, touch controls, and a small loudspeaker on the top, and the pitch potentiometer control protruding from the left-hand side. Fig. 17 shows the inside of the box. A single printed circuit board, mounted on the base, contains all the components. This is joined to the controls on top by ribbon cable. Two PP3 batteries provide an internal power supply, with provision for connection to an external one via a socket. An amplifier/recorder jack output is also included.

The left-to-right motion of the joystick controls formant 2, from minimum (left) to maximum (right), and the down-and-up motion controls formant

1, from minimum (down) to maximum (up). It is important to keep the leads from the joystick potentiometer as short as possible, since any noise on them will be amplified by the filters.

**Setting up.** The component values given in Fig. 12 should provide acceptable ranges and bandwidths for the formant filters, the sibilance filter, and the sound source. However, it is necessary to adjust the relative levels of the voicing, aspiration and sibilance channels by suitable choice of resistor values.

First, adjust the levels of voicing and aspiration by substituting different values for the 47k $\Omega$  and 82k $\Omega$  resistors at the input to the first operational amplifier. It is best to examine the output of the amplifiers with an oscilloscope, so that as large a gain as possible is obtained without clipping occurring anywhere along the formant chain. The tests should be made at different positions of the joystick, because this radically affects the Qs of the filters.

Then adjust the two input resistors to the final operational amplifier to achieve a pleasing balance between sibilance and voicing. Finally, set the feedback resistor for this stage to produce maximum output voltage without clipping in the audio amplifier.

**Making the Chatterbox talk**

The best way to learn about the Chatterbox is to play with it. However, people sometimes have difficulty getting started, so here are some suggestions to help familiarize you with the controls and their effects.

First, identify the vowel positions (Fig. 18). Remember that there are far more vowels in English speech than the a, e, i, o, and u of English writing!

Now experiment with pitch variations while a steady vowel is being produced. Try a rising pitch, a falling pitch, a rise and then a fall, and a fall followed by a rise. Correct control of pitch is essential to make the speech sound natural.

Try some diphthongs next, as in "go," "toy," and "play" (Fig. 19). These are made up as a slide from one vowel position to another.

Finally, turn to the noisy sounds. One pair of touch contacts produces aspiration, and the sound is affected by the joystick position. You should be able to




Fig. 18. Approximate vowel positions.




Fig. 19. Some diphthongs.




Fig. 20. Things to say.

use it to whisper the vowels and diphthongs learned above. The other pairs of contacts produce the sounds at the beginning of "ships," "fish," and "salt." It is possible to make "t" as in "eat" with a very short burst of "s." However, it's quite difficult to integrate these noises with the vowel-like sounds to get proper words like "delicious." No wonder babies take so long to learn to talk! Fig. 20 shows some things to say.

**The Chatterbox in use**

The Chatterbox has found an exciting application as a stimulus for retarded and autistic children. (It has also been used with a young child, blind from birth.) As you may know, such children find it difficult to interact with other people and often prefer to play with machines. The Chatterbox with its relatively rich structure of controls and noises, proves an interesting device for them.

Dorinda Bath, at the University of Nottingham, has conducted some experiments with retarded and autistic children who were functioning at a mental age level of two years. They were given the toy and instructed to "play with it and see what it does". While playing with the pitch potentiometer, trying to balance the toy on its joystick, chewing it, and so on, they usually triggered one of the touch switches by accident. From then on it was a matter of discovering which of the controls worked by themselves and which operated in consort with others. Three levels of activity were defined:

- level I: making a noise by touching one of the controls;
- level II: performing two related actions

simultaneously (e.g. voice switch and joystick);

level III: performing three actions simultaneously i.e. voice switch, joystick, and pitch potentiometer).

Many of the children discovered by themselves not only level II but also level III activity. However, because the possibilities for co-ordinated control at level III usually outstripped the child's manipulative powers, because her hands were so small, she often involved the adult experimenter on her own initiative to help her ("you touch these buttons while I move this"). Thus the Chatterbox served as a catalyst for valuable interpersonal relations. This is especially important for the autistic child, who withdraws from contact with humanity and does not respond to friendly advances from other people.

As a consequence of her experiences with the Chatterbox, Dorinda suggested that it would be worthwhile to explore the possibilities of using it in child therapy, to increase auditory awareness as well as to expand attention span in children with disorders of speech and language.

**Acknowledgements.** We would like to thank all the people who contributed ideas and practical help to the Chatterbox project: Chris Corbett, Kel Fidler, Rick Jenkins, Bob Mack, Roger Moore, and especially Bob Booker, John Brazier and Richard Pope, whose hard work went well beyond the call of duty. The spectrograms were made by Linda Shockey. We are very grateful to Dorinda Bath for her enthusiasm in using the Chatterbox with emotionally retarded children.