

The Digital Computer as a Musical Instrument

A computer can be programmed to play "instrumental" music, to aid the composer, or to compose unaided.

M. V. Mathews

With the aid of suitable output equipment, the numbers which a modern digital computer generates can be directly converted to sound waves. The process is completely general, and any perceivable sound can be so produced. This potentiality of the computer has been of considerable use at the Bell Telephone Laboratories in generating stimuli for experiments in the field of hearing, and for generating speech sounds and connected speech in investigations of the factors which contribute to the intelligibility and naturalness of speech.

The quality of sound is of great importance in two fields—that of speech and communication and that of music. Our studies at the Bell Laboratories in the first of these fields have led us, over the past few years, to related studies in the production of musical sounds and their organization into musical compositions. I believe that this by-product of our work on speech and hearing may be of considerable value in the world of music, and that further work in this direction will be of substantial value in furthering our understanding of psychoacoustics.

There are no theoretical limitations to the performance of the computer as a source of musical sounds, in contrast to the performance of ordinary instruments. At present, the range of computer music is limited principally by cost and by our knowledge of psychoacoustics. These limits are rapidly receding.

In addition to generating sound, the computer can also function as a ma-

chine for composing music. It can either compose pieces based entirely on random numbers generated by itself or it can cooperate with a human composer. It can play its own compositions.

Here I first describe the process for converting numbers to sounds, then I describe a program for playing music. Next I consider a psychoacoustic problem which is typical of those posed in attempts to make more interesting sounds. Finally, I look to the future, to the time when the computer is itself the composer.

Sound from Numbers

How can the numbers with which a computer deals be converted into sounds the ear can hear? The most general conversion is based upon the use of the numbers as samples of the sound pressure wave. A schematic diagram of this process is shown in Fig. 1. Here a sequence of numbers from the computer is put into an analog-to-digital converter, which generates a sequence of electric pulses whose amplitudes are proportional to the numbers. These pulses are smoothed with a filter and then converted to a sound wave by means of an ordinary loudspeaker. Intuitively, we feel that if a high enough pulse rate is used and the amplitudes of the pulses are generated with sufficient precision, then any sound wave can be closely approximated by this process. Mathematically, it has been established (1) that this conclusion is correct. A sound wave with frequencies from 0 to B cycles per second can be generated from a sequence of two B pulses per second. Thus, for

example, by running our computer at a rate of 30,000 numbers per second, we can generate sound waves with frequencies from 0 to 15,000 cycles per second. Waves in this frequency range are about the only ones the human ear can perceive.

The signal-to-quantizing-noise ratio of the sound wave depends on the accuracy with which the amplitudes of the pulses are represented. Computers deal with a finite number of digits and, hence, have limited accuracy. However, the computer limits are more than sufficient acoustically. For example, amplitudes represented by four-digit decimal numbers, are accurate to within 1 part in 10,000, an accuracy which represents a signal-to-noise ratio of 80 decibels; this is less noise than the ear can hear, and less noise than would be introduced by any audio equipment, such as the best tape recorder.

The sampling process just described is theoretically unrestricted, but the generation of sound signals requires very high sampling rates. The question should immediately be asked, "Are computers of the type now available capable of generating numbers at these rates?" The answer is "Yes," with some qualifications. A high-speed machine such as the I.B.M. 7090, using the programs described later in this article, can compute only about 5000 numbers per second when generating a reasonably complex sound. However, the numbers can be temporarily stored on one of the computer's digital magnetic tapes, and this tape can subsequently be replayed at rates up to 30,000 numbers per second (each number being a 12-bit binary number). Thus, the computer is capable of generating wideband musical sounds. Because of the cost of computer time, we often limit our studies to those for which the computer is run at lower rates, such as 10,000 numbers per second—a rate which yields a bandwidth of 5000 cycles per second.

The direct conversion of numbers to sound is only one of the ways in which the computer can generate sounds. An alternate procedure is to use the numbers from the computer to control electronic apparatus such as oscillators and filters, which, in turn, generate the sounds. These processes have been carried out by the Radio Corporation of America music synthesizer (2) and by a machine constructed at the University of Illinois (3). This procedure has the advantage that a much lower rate

The author is director of the Behavioral Research Laboratory, Bell Telephone Laboratories, Murray Hill, N.J.

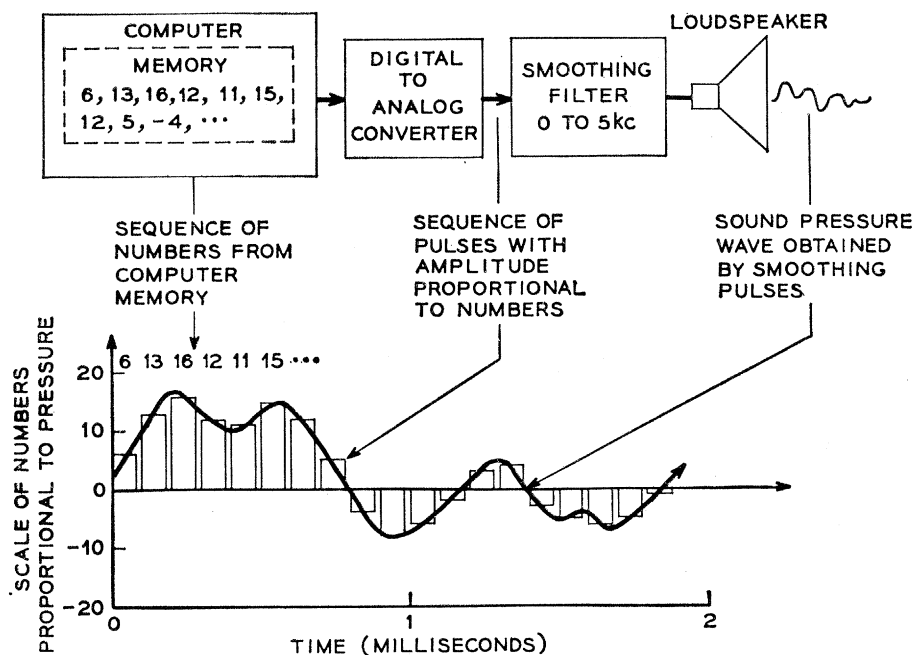


Fig. 1. Schematic diagram depicting the conversion of a sequence of numbers stored in a computer memory to a sound pressure wave form. The sampling rate is 10,000 numbers per second to yield a bandwidth of 5000 cycles per second for the sound wave.

of output is required from the computer, and hence the computation costs are less. The disadvantage is that the only sounds that can be generated are those produced by the particular electronic apparatus employed, and hence the generality of the sampling process is not attainable.

The work described here, which was done at the Bell Telephone Laboratories, is based entirely on the sampling method.

Playing a Computer

To specify individually 10,000 to 30,000 numbers for each second of music is inconceivable. Hence, the numbers-to-sound conversion is useless musically unless a suitable program (set of computer instructions) can be

devised for computing the samples from a simple set of parameters. The central contribution of the Bell Telephone Laboratories to computer music is a program for computing the many samples in a note from the few parameters characterizing it. The details of the program determine the limits of the sounds now obtainable (from the standpoint of practicality) with a computer. The program represents a compromise between a general procedure, through which any sound could be produced but which would require an inordinate amount of work on the part of the composer, and a very simple procedure, which would too greatly limit the range of musical sounds obtainable. In order to give the composer flexibility between these two extremes, the program is divided into two parts. In the first part the composer specifies,

in computer language, the characteristics of a set of musical instruments. The program unit that represents the instrument (or the "instrument unit") may be as simple or as complex as he desires. He then prepares a score consisting of a list of notes to be played on the instrument-units he has created. The samples of sound wave are generated by putting the score, in a form the machine can read, into the computer, together with the instrument-units, and turning on the computer. The numerical output is recorded on a digital magnetic tape for subsequent conversion to acoustic form.

The interconnected blocks of program which make up the instrument-unit are called unit generators. There are a number of different types of unit generators, each of which has a specific function. A typical instrument-unit is shown in Fig. 2. This instrument-unit is composed of five unit generators; three of them are oscillators, one is a random-number generator, and one is a summing circuit. Each oscillator has two inputs and one output. The upper input specifies the amplitude of the output; the lower input specifies the frequency of the output. The wave shape of the output need not be sinusoidal and can, indeed, be any one of 20 arbitrary functions stored in the computer memory. In the example given, the wave shapes are sketched on the oscillators, oscillator 1 producing a damped sinusoid, oscillator 2, a triangular attack-and-decay function, and oscillator 3, a sinusoid.

This particular instrument-unit produces notes with controlled attack and decay and with a frequency variation or vibrato. Oscillator 1 produces the main frequency, which, for example, for note A of the musical scale would be about 440 cycles per second. The amplitude of the output of oscillator 1 is modified by oscillator 2, which imposes a desired attack-and-decay function. The frequency of oscillator 2 is such that it goes through exactly 1 cycle of oscillation per note. Such low frequency operation is quite feasible for oscillators used in computer programs. The amplitude of the note is controlled by the upper input to oscillator 2, and this amplitude is one of the input parameters which the composer must specify for each note. The frequency of oscillator 1 is the sum of three components, one being the center frequency of the note (an input parameter), the other two being a periodic

Table 1. A typical computer score. The corresponding conventional score is shown in Fig. 3.

Operation code	Instrument No.	Starting time (sec)	Duration (sec)	Loudness (arbitrary units)	Frequency (cy/sec)	Periodic vibrato		Random vibrato	
						Amplitude (cy/sec)	Frequency (cy/sec)	Amplitude (cy/sec)	Bandwidth (cy/sec)
Play	1	0.0	0.25	1	466	0	0	7.0	6
Play	1	.5	.25	3	698	0	0	10.5	7
Play	1	1.0	.125	5	698	0	0	10.5	7.5
Play	1	1.5	.125	7	698	0	0	10.5	8
Play	1	2.0	.25	9	932	0	0	14.0	8.5
Play	1	2.25	.125	10	784	0	0	11.7	9
Play	2	0.5	.50	1	116.5	1.7	6	0	0
Play	2	1.5	.25	5	156	2.3	7	0	0
Play	2	2.0	.125	10	233	3.5	8	0	0

and a random vibrato, respectively. The amplitudes and frequencies of the periodic and the random vibratos can be controlled separately, and thus these amplitudes and frequencies constitute four additional input parameters. The random generator is a unit whose average amplitude is controlled by its upper input, and whose bandwidth, or average rate of variation, is controlled by its lower input.

Each of the unit generators consists of a block of a few computer instructions. By assembling particular groups of generators, almost any desired properties of the output sound wave may be achieved and controlled. The computation time is roughly proportional to the number of generators, and the number of input parameters tends to increase with the number of generators. Thus, complexity of the instrument-unit is paid for both in terms of computer time and in terms of the number of parameters the composer must supply for each note. In general, the complicated instrument-units produce the most interesting sounds, and the composer must make his own compromise between interest, cost, and work.

In supplying specifications for an instrument-unit, the composer does not have to be concerned with the computer instructions represented by each unit generator. A simple language is available by means of which he can supply the specifications in a manner

no more complicated than drawing a diagram such as that of Fig. 2. About ten different types of generators are available, and new ones may easily be defined. However, the generators most frequently used are those shown schematically in Fig. 2. Most composers wish to use a number of instrument-units, and these are distinguished simply by identifying numbers. If two or more instrument-units are to play simultaneously, the samples they generate are simply added together. This operation corresponds to the addition of the sound pressure waves from several sources in air.

The Score

Once the composer has supplied specifications for the orchestra, he must prepare a score giving the parameters of the notes he wishes played. An example of a score is given in Table 1. It corresponds to the conventional musical score of Fig. 3. Two instrument-units, of the type shown in Fig. 2, are assumed to be available; these are designated instrument 1 and instrument 2. Column 1 of Table 1 is an operation code which indicates that a note is to be played; column 2 designates the instrument-unit on which the note is to be played; column 3 gives the starting time of the note; column 4 gives the duration of the note; and columns 5 to 10 supply the input parameters re-



Fig. 3. A conventional score, corresponding to the computer score of Table 1.

quired by the instrument-unit. Both the number of these parameters and their arrangement depend entirely upon the specifications the composer has supplied.

To play the music, the computer reads a line from the score, at the appropriate time inserting the parameters given in the score into an instrument-unit and activating the instrument-unit so that it may generate numbers equivalent to the duration of the note. The entire sequencing process is automatically taken care of by the program and need not concern the composer.

Musical Examples

The best way to form an opinion about computer music is to listen to some. A commercial record is now available (4). However, a certain amount of discussion may be useful. A large number of compositions have now been produced by computer. They range from 16th-century music for the recorder to 12-tone music; from classical to popular music; from serious compositions by professional composers to acoustic experiments by psychologists; from pieces formed entirely of conventional tones to pieces formed entirely of random noise.

An outstanding advantage is the precision of the computer. Effects are exactly reproducible. Very complex effects, such as simultaneous synchronous tremolo and vibrato of two notes, can be obtained. Exact rhythmic patterns, such as the playing of seven notes in one voice against five notes in another voice, are as easy to produce as any other note sequence. Chords can be shifted in frequency with a glissando while the frequency ratios of the various voices are maintained.

The computer is also very flexible. When sufficient effort has been expended in supplying specifications for

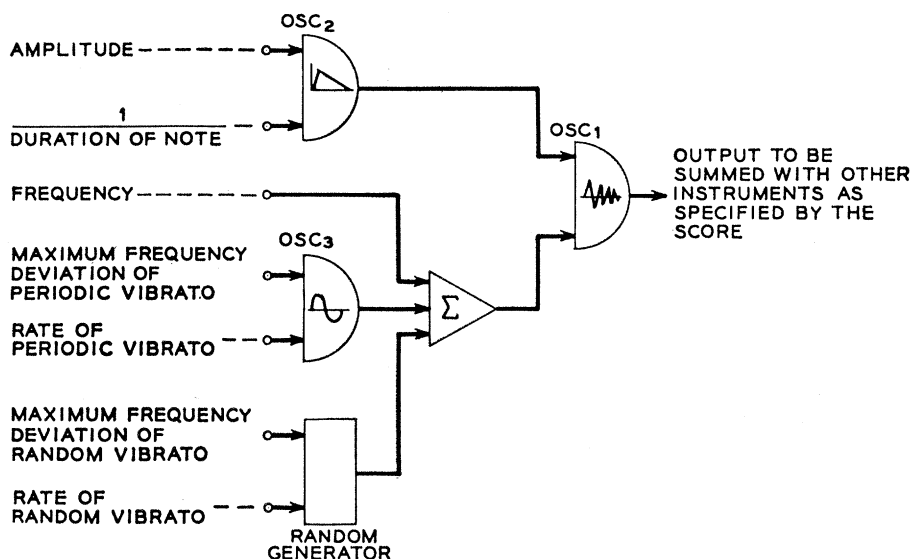


Fig. 2. Schematic diagram of a typical instrument-unit in the computer orchestra. The diagram represents a section of the computer program. In order for the computer to produce a note, numerical values for the note parameters shown at the left of the diagram are stored in the program. The program then generates samples of the sound pressure wave form.




Fig. 4. Ratios of two rise times that are just discriminable (discriminable 75 percent of the time) as a function of the shorter of the two rise times. For times greater than 5 milliseconds, the ratio of discriminability is about 3/2. For times shorter than about 5 milliseconds, the ratio corresponds to a difference of 1.5 milliseconds in rise times.

the instrument-unit, almost any sound effect can be produced, provided the wave form of the sound can be described.

It is very easy to use the computer in this way. The electronic equipment (computer and output equipment) has been constructed once and for all. There are no soldering irons, tape-splittings, or even knob-twistings involved, as there are with other electronic equipment for producing music. No manual dexterity is required. In-

stead, one writes down and gives the computer a sequence of numbers.

The computer has proved to be of more use to the composer who wishes to produce new compositions with new sounds than to the performer who wishes to duplicate existing music. It is difficult and expensive to copy all the acoustic details of conventional instruments well enough to produce excellent instrumental sounds. Hence, the playing of pieces already written for existing instruments seems better left to the instrumental performer.

The most apparent limitation in the field of computer music results from lack of adequate knowledge of the sound of a given pressure wave. The computer sounds are described in terms of the waveshapes produced by the unit generators in the instrument-units. This method for describing sound is quite different from the method of ordinary music, in which the sound is specified by the instrument which produces it, when certain instructions have been received by the performer. Musicians have had a great deal of experience in listening to the sounds produced by violins, oboes, horns, and other instruments and are well able to predict the contribution of these instruments to the total sound. By comparison, computer musicians have had very little experience in trying to predict the effect of a given harmonic-composition factor or a given attack-and-decay function on the timbre of a note.

An example of a psychoacoustic

surprise is the dominance of the rates of attack and decay in determining the character of a sound. These rates are a much more significant factor than the harmonic composition. Thus, a "violin," if artificially given the attack-and-decay characteristic of a piano, sounds much more like a piano than a "piano" does when it is given the attack-and-decay characteristic of a violin.

Another unexpected result is the importance of suitable random variations in almost all parameters of a note for introducing richness and interest. A sound which is otherwise quite plain can be greatly improved by introducing a random variation of up to 50 percent in its amplitude, at an average rate of something between 8 and 20 cycles per second, and by introducing a random variation of perhaps 1 percent in its frequency, at a similar rate.

Psychoacoustic Questions

Our musical studies with the computer indicate that, in this area, the major problem to be overcome by a composer concerns the relation of the physical description of the sound waves to the psychoacoustic effects which he desires. Contrary to the situation with conventional instruments, with the computer the composer himself is solely responsible for the sound. He has no conductor to interpret his composition. He himself must give careful consideration to even such a simple matter as the relative loudness of the instruments in a group.

Our experience has shown how little we now know about the relation of the quality of sound to various features of waveform. A new body of psychoacoustic data is necessary. These data should relate the properties of the acoustic waves of music to perceived qualities of sound. Part of the task of assembling these data can, of course, be given to the composer, and part of the data can be supplied by interested psychologists. An increase in knowledge in this field is bound to be of value and interest in other fields, including those of speech and hearing.

Typical of the sort of new knowledge that is needed is knowledge concerning the rate of attack of a note. As already noted, this parameter has a strong influence on the timbre, and if the composer is to make use of it he must know how small a change in rate is perceptible. Not being able to




Fig. 5. Section of a score prepared for a study by J. C. Tenney. The average values for note-duration and loudness are shown by the solid lines as functions of playing time. The allowable range of variation of these parameters is shown by the dashed lines surrounding the solid lines.

find these data in the literature, J. C. Tenney made a small study. The results are shown in Fig. 4. Here the just-detectable difference in the rise times of two tones is shown as a function of the shorter of the two rise times. These data indicate that, for rise times longer than about 5 milliseconds, a ratio of 3 to 2 is just detectable. For times shorter than 5 milliseconds, a difference of 1.5 milliseconds is necessary. Using these data, the composer can select a "scale" of attack functions which are separated by intervals that are equal in terms of the listener's ability to differentiate between the attack functions. Such data are typical of those required by a composer.

Composing with the Computer

So far I have described use of the computer solely as a musical instrument. The composer writes one line of parameters for each note he wishes played and hence has complete control of the note. He is omnipotent, except for lack of control over the noise produced by the random-number unit generators. Here a minor liberty is allowed the computer.

However, composing-programs are a reasonable area of computation, and work in this direction has already been done by Hiller (5), Olson (6), Brooks (7), and others. A number of different approaches can be taken toward composition by computer. At one extreme, the computer can be given a set of rules, plus a random-number generator, and can simply be turned on to generate any amount of music. Hiller's work is perhaps closest to this extreme. In the opposite direction, the human composer can maintain close control of the music, using the computer merely to avoid some of the repetitious and tedious work involved in representing his musical ideas. Once a theme with many notes has been written, a program can be devised for repeating the theme by means of a single instruction.

Furthermore, the theme can be modified in simple ways: it can be transposed to another pitch range or played upon a different instrument; its tempo can be changed or its loudness modified. Harmonization of the theme according to simple rules is possible. Other means of modifying or developing a theme in interesting ways may be forthcoming. The composer could, perhaps, form a composition from a set of thematic material, which he supplied, and a set of fixed transformations.

At present, the music-playing program has been modified so as to make transformational development of a theme possible. Certain of the simplest transformations have been programmed. These include all those mentioned above, with the exception of harmonization. As yet not enough music has been generated to assess the significance of this approach.

A slightly different method has been tried by Tenney. His approach is a compromise between a purely random and a completely specified composition. The parameters of the individual notes of the composition are generated as a sequence of independent random numbers by a random-number routine. However, the average value and the variance of these parameters are specified by the composer as functions which change slowly throughout the composition. The "score" of a section of one of Tenney's works is shown in graphic form in Fig. 5. The means and variances of the note-durations, loudness, and other parameters of the various voices are controlled. Indeed, the number of voices playing at a given time is controlled. By this relatively simple algorithm, a long-range structure which can be clearly recognized by the ear is imposed on the composition. Thus, one of the characteristic shortcomings of random compositions—a lack of long-range structure—can be overcome.

The use of computers as an aid in composition is still very new. We hope that by this means the composer can

avoid having to write out all the individual notes in a piece of music in order to express his ideas—that he will be able, rather, to write directly in parameters that are much more closely related to his musical objectives, letting the machine generate the individual notes. Whether this objective can be attained remains to be seen.

The Future of Computer Music

I have indicated how almost any sound can be produced by treating the numbers generated by a computer as samples of the sound pressure wave. A very high sampling rate is required, and, if this process is to be useful musically, programs for generating samples from the parameters of notes must be written. A broad set of these programs is now available and has been used for playing, on an experimental basis, a wide range of music. Additionally, studies are being made on possible uses of the computer as an aid in composing. In such studies the computer usually plays its own compositions and constitutes a composer-player team.

Computer music appears to be very promising technically. However, the method will become significant only if it is used by serious composers. At present, our goal is to interest and educate such musicians in its use. We believe that competent work in the field can benefit not only music but the whole field of psychoacoustics.

References and Notes

1. C. E. Shannon, *Bell System Tech. J.* **27**, 379 (1948).
2. H. Olson, H. Belar, J. Timmens, *J. Acoust. Soc. Am.* **32**, 311 (1960).
3. A musical output for the csx-1 computer in the Coordinated Science Laboratory at the University of Illinois was constructed by J. Divilbis. No written description is as yet available.
4. "Music from Mathematics." Decca record DL 9103 Monaural or DL 79103 Binaural.
5. L. Hiller, Jr., and L. Isaacson, *Experimental Music* (McGraw-Hill, New York, 1959).
6. H. Olson and H. Belar, *J. Acoust. Soc. Am.* **33**, 1163 (1961).
7. F. Brooks, Jr., A. Hopkins, Jr., P. Newmann, W. Wright, *IRE (Inst. Radio Engrs.) Trans. Electron. Computers* **EC-6**, 175 (1957).