

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-

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

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