MASSACHUSETTS INSTITUTE OF TECHNOLOGY
A. I. LABORATORY

Artificial Intelligence
Memo No. 213 February 1971

THE COMPUTER AS A PERFORMING INSTRUMENT

Gordon Mumma (Cunningham Dance Foundation)
Stephen Smoliar

This memo was originally presented as a Project MAC seminar on February 20, 1970.

Work reported herein was supported by the Artificial Intelligence Laboratory, an M.I.T. research program sponsored by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract number N00014-70-A-0362-0002.

Reproduction of this document, in whole or in part, is permitted for any purpose of the United States Government.

1. Background

From the outset, the computer has established two potential roles in the musical arts -- the one as a sound synthesizer and the other as a composer (or composer's assistant). The most important developments in synthesis have been due to Max Mathews at the Bell Telephone Laboratories [7]. His MUSIC V system endows a computer with most of the capabilities of the standard hardware of electronic music. Its primary advantage is that the user may specify arbitrarily complex sound sequences and achieve them with a minimum of editing effort. Its primary disadvantage is that it is not on-line, so that the user loses that critical sense of immediacy which he, as a composer, may deem valuable.

The major experiments in the area of composition of music have been due to Lejaren Hiller. Jointly with Leonard Isaacson, Hiller designed a series of computer programs whose outputs were interpreted in the standard notation of a string quartet (Illiac Suite for String Quartet, 1957) [4]. The results have been performed in concert and are currently available on a commercial recording [2]. This recording also contains Hiller's second large-scale project, Computer Cantata, which not only involved the specification of a score but the synthesis of taped sounds as well [3]. More recently, Hiller has collaborated with John Cage, when the latter was Visiting Professor of Music at the University of Illinois, in the preparation of a piece called HPSCHD, for live harpsichords and computer-generated tapes [1].

2. HPSCHD

HPSCHD was composed on a multiplicity of scales of different temperaments. A series of pitch-determination routines were applied to all divisions of the scale from five notes per octave to 56 notes per octave. Originally, 55 notes per octave was set as the smallest micro-division, because that particular scale admits of perfect fifths; but the computer had sufficient memory to accomodate one more scale.

The composition consists of the 52 computer-generated tapes (with the different divisions of the octave) and seven harpsichord solos to be played live. The sound was generated by the computer, following the principles of Max Mathews' program and making some 32 variations on harpsichord-like overtone structures. Binary decisions (such as those in the pitch-determination routines) were settled by a conventional random number generator; but when there were more than two possible choices, the decision was determined by a special routine called ICHING (based on a Chinese method of divination in which three coins are tossed six times to obtain a hexagram, i.e. number, from 1 to 64 [5]. Since then, Ed Kobrin of the University of Illinois has made another ICHING program [6]. The first ICHING program was so integrated with the other routines of HPSCHD that it couldn't be extracted; but Kobrin has made another ICHING program which not only gives random numbers from 1 to 64 in great quantity, but also is adaptable to decisions with more or less than 64 choices.

Another interesting thing to mention about HPSCHD was the physical circumstances of the performance, which was in a very large round building with the seven harpsichords in a circle leaving the audience free to

either move, or stay on the bleacher-like area around the hall, or move into the center of the space itself and be in the middle of the harpsichords. The 52 tape-machines were around the perimeter of the building and the speakers were equidistant around this circle. Then on an inner circle (like a donut) there were the seven speakers from the amplified harpsichords. Since the piece was microtonal, it was thought that the visual activity in the performance should be like telescopes. So there were 80 slide projects, 100 moving pictures, and 8,000 slides of travels through space and such things, so that visually one was seeing the world as from a telescope in space and acoustically one was hearing this fine division of the octave. Some of these 80 slide projectors were flashing images (like those one could see inside on big circular screens) on rectangular outdoor screens; so that as one came near the building one already got, visually, an impression of what was happening inside. When the performance was over (and it had lasted for five hours) the outdoor projectors were still working because it had been impossible to turn them all off.

3. Cybersonics

While HPSCHD was performed live, the role of the computer in the piece was strictly as a pre-processing device. Such has been the case for almost all applications of electronic equipment to music (with the exception of the more conventionally based electronic instruments, such as the electric organ, Ondes Martinot, and theremin). More recently there has arisen a school of thought of "live electronic music," in which the electronic musical apparatus is used in an on-line capacity rather than

just simply for the preparation of tapes prior to performance time.

One facet of live electronic music is very closely related to the field of artificial intelligence and, in fact, has been given the name "cybersonics". The first piece to be so called was composed in 1963 and was a piano piece called "Medium Size Mograph 1963". To the piano was attached a microphone and the microphone was attached to an amplifier whose gain varied according to characteristics of the sound of the piano itself, so that the performer was able to make a different timbre out of the piano simply by changing its amplitude. Instead of starting loud and decreasing to a quiet sound, the amplifier has a built-in mechanism whereby it started quietly and then went through some excursions of its amplitude automatically during the course of every sound. The resulting sound was not much like a piano. That was a fairly simply application of the cybersonic concept. The major innovative feature of this concept was that one could carry a small box with which, in a relatively simple way, one could (effectively) make a "new" instrument out of a conventional one.

A more recent work along these lines is for a French horn, called "Hornpipe", in which the French horn is not fundamentally modified as the piano was in the earlier piece; instead, the acoustical properties of the space in which one performs are made use of in the piece. It is the business of the performer in this piece to determine where some of the resonances of the space exist in order to be able to make them sound through fairly simple acoustical feedback procedures which form an accompaniment for the French horn.

The performer wears a small box of circuitry attached to an amplifier and loudspeaker. The piece starts as a solo, during which the box is listening. It has within it several resonant circuits which move about on their own until they find those sounds made by the horn which are most resonant within the space. As soon as these resonant circuits have locked onto these different spatial resonances they begin to store information about the number of times that those resonances are hit. After a certain amount of time, that storage opens a gate. In effect, one or more of the several amplifiers in the box is actuated, and the space is able to respond to the loudspeakers. Further, the performer is in a situation where he can add more sounds from the French horn in such a way as to turn off the amplifiers. So the result is, effectively, a duo between the performer (and his hardware) and the space itself. There are time circuits on the amplifiers that will shut them off on their own, so the performer can turn them off in advance according to what he plays or he can extend them. He can even turn them back on again.

The hardware is realized mostly with differential amplifiers, which is possible to do in a convenient size with integrated circuits. There are eight such differential amplifiers, and they have resonance circuits in their feedback loops which are voltage-controlled as to what their resonances are. This is accomplished through the resistance between the emitter and collector of a transistor, which varies according to what's on the base. The voltage stored on the base is what tunes each of those amplifiers. (In fact, they very rarely all are operating at once. Some of them don't find a resonant frequency to lock onto.) Those amplifiers are then followed by a gated amplifier, essentially a voltage-controlled

amplifier, which turns on with a Schmidt trigger. After certain voltages are released, it turns on. These eight parallel signals are added at the output and mixed down to two outputs for stereo. The memory is capacitors' charge on the gates. There are two storage parts: first, an electrical storage to tune the resonant circuit; then, following the resonant circuit, is a voltage-controlled amplifier which has a charge added on to a second capacitor storage until it reaches the point that the gate opens. As it opens, it drains out the second capacitor after a certain period of time until it closes the gate. It can also be turned off, as was mentioned, by playing pitches that detune the previous resonant amplifier which will make that amplifier drop out by changing its resonant point.

The input consists of two lavalier microphones worn in the performer's pockets. The horn is equipped with switches for control signals to raise and lower the amplitude of the whole thing. The circuitry can sometimes get out of hand. One can get into a situation while performing where too many gates would open, in which case one can simply back down on the overall gain control on the horn. That's a safety feature.

This piece is not absolutely fool-proof. There were occasions when it would not work. The first occasion was in a recording studio for the Italian radio in Rome. They had one of those rooms that was practically anechoic; so there was nothing to resonate and we just couldn't make it work. Out-of-doors is not an ideal place to create resonances within the range of this gadget. That problem was solved simply by finding an area next to a building so that there was enough of a reflective surface behind the performer.

A recording of the piece has been made; it was done in Waltham, at the Rose Art Museum at Brandeis University. This is a big marble

building with a pool of water on the lower floor. In this space the hardware wouldn't turn off because the space was so reflective.

The title, "Hornpipe", does not refer to the old sailors' dance, but rather to the face that a horn is part of the piece and the resonant characteristics of the circuitry are very much like the resonant characteristics of pipes. There are specific things in the score which one must play on the instrument, but they are not played in a specific order; the performer has to choose from the various sequences, beginning by choosing something to play. As time goes on he has to use up the material in various ways. The specific pitches are chosen by the performer and are gradually narrowed down in the course of the performance to those things that are relevant to the space.

4. "Conspiracy 8"

The circuitry within the box for "Hornpipe" has a kind of elementary computer-like function, in that it makes decisions which are, in a very real sense, determined by the composer in advance. It stores information, in an analog rather than a digital form, and it makes decisions which affect what the performer does outside of the device itself. It's a special-purpose computer -- of no good to mankind except to make a bit of sound.

The use of a large-scale digital computer can be of interest to the composer, not only from the standpoint of sound generation but also as a "personality", if you will, in the business of the performance. Aside from being a sonic profession, music is also a social activity.

The medium of computer-generated sound is so elaborate that it requires far more work than many composers are presently willing to spend in programming when such sounds can be accomplished much more easily through hardware. (HPSCHD was two years in the making.) Most live electronic music employs equipment which is not unsimilar to the various commercial synthesizers -- sound generators and sound modifiers, largely voltage-controlled or current-controlled devices. Such devices can be computer-controlled, relieving the computer of the burden of sound synthesis and making it available for other possibilities.

"Conspiracy 8" (1970) does not strictly involve computer-controlled hardware, but it demonstrates the possibilities inherent in such techniques. The Chicago Conspiracy Trial, with its eight defendants, plus Judge Julius Hoffman and a lot of everybody else, was like an enormous, rather spectacular musical operatic composition. The whole thing has been one of the great theatre events of the past several years. (Consider, for example, the scene when Bobby Seale was removed from the cast of performers. They changed the name to the Chicago Seven, because the whole idea of the Conspiracy Eight had become too difficult to handle.) "Conspiracy 8" was inspired by this trial. It's not a political piece, it's a social piece. It can have up to eight or more performers and it is, literally, a sort of conspiracy. The performers bring to it whatever they have to bring -- whatever they are able to do. Somebody is a moderator, whether it's Judge Hoffman himself or the computer for which the piece was designed (a PDP-6 -- the actual computer model is irrelevant, so long as it can communicate with the performers; all judges are different).

Both judges and computers make noise when they process their material. The actual relationship between the computer "judge" and the performers is still being developed. Currently, one of the performers communicates with the computer by means of a teletype; in the future, the teletype may be circumvented by means of microphone inputs to the computer. In any event, the role of the computer is as a performer. It doesn't matter what sounds occur, that is, what the sounds are that are brought to the piece; what matters is that all performers are in a position to exercise control over what each other does. The sounds of the computer, teletype and of everything else used are altered by the electronic means, and all performers share in deciding upon alterations.

To obtain sound from the PDP-6, the rightmost six bits of a memory word specified by the test switches are monitored by six "one-bit digital-to-analog converters"; these devices simply emit a click whenever the bit to which they are attached changes state. These six channels are mixed and separated into two output channels of sound, subject to a limited degree of filtration. These two channels then serve as inputs to the elctronic apparatus. Other inputs to this apparatus come from microphones for the individual performers. (In the current version, one microphone picks up the sounds of the teletype while another amplifies the sound of a musical saw.) Only one of the two computer channels is actually subject to modification; the other is always unmodified and is only subject to amplitude control. The apparatus admits of four kinds of modification procedures: frequency modulation, amplitude modulation, formant modulation (variable filtration), and degree of reverberation. In one particular mode of operation, the

frequency modulation is dependent upon the amplitude on the input signal -- i.e. the louder the signal, the higher its pitch.

For purposes of communication, the current version uses the English program language, in as reasonable a way as possible. This was accomplished by a special purpose for Weizenbaum's ELIZA [8]. This particular version, implemented in LISP, has been given a vocabulary relevant to the circumstances of "Conspiracy 8". No particular conspiracy jargon was implemented into the script (although such would be feasible, if desired). The script was written purely in terms of the available instruments. In fact, it has been modified for a concert in New York during which an additional performer was playing a melodica.

We are also considering a project involving the PDP-14, a new computer which is made for industrial automation and control procedures. It appears to have very direct applications to live electronic music in that many procedures do not require man-machine communication. A machine which is preprogrammed for specific tasks may be able to ease the burden of working with a large configuration of electronic equipment.

There is also a composer in Illinois who has built two computer kits; they are used for educational training and computer programming procedures. He built two of them and haywired them together so that he is now able to perform on them, treating them as live instruments. He doesn't make tapes from them at all. He plugs in various kinds of sound generators and uses the computers as an articulation device. There are probably several more such projects in development.

References

1. J. Cage and L. Hiller, HPSCHD, Source 4:11 (1968).

2. Computer Music from the University of Illinois, Heliodor, HS 25053.

3. L.A. Hiller and R.A. Baker, Computer Cantata: An Investigation of Compositional Procedures, Perspectives of New Music, 3:62 (1964).

4. L.A. Hiller and L.M. Isaacson, Experimental Music, McGraw-Hill Book Company, New York, 1959.

5. The I Ching, German translation by R. Wilhalm, English translation by C.F. Baynes, Princeton University Press, Princeton, N.J., 1969.

6. E. Kobrin, I Ching, Source 8:1 (1970).

7. M.V. Mathews, The Technology of Computer Music, MIT Press, Cambridge, Mass., 1969.

8. J. Weizenbaum, ELIZA -- A Computer Program for the Study of Natural Language Communication Between Man and Machine, CACM 9:36 (1966).