A COMPARISON OF ELECTRONIC SOUND (MUSIC)
SYNTHESIS SYSTEMS FOR MUSICAL COMPOSITION
IN THE UNITED STATES AND EUROPE

by J.W. Beauchamp

April 10, 1973.

During the past decade a number of techniques for electronic synthesis of sounds have been developed to enable composers to directly create sounds and succession of sounds without necessarily resorting to the use of tape editing techniques. These techniques may be divided into two general categories: 1) real time control or performance, and 2) non-real time score processing or computation, although mixtures of these two techniques are possible. I will review those systems with which I am familiar either by personal observation and use or by review of documentation (articles and tapes) and through conversations with those who have used the equipment.

Real time control: Commercial Synthesizers

Several electronic music synthesizers are being produced in the U.S. and England, the most famous of which are a) Moog (Moog Music Inc.,Buffalo, N.Y.), Arp (Newton, Mass.), Buchla (Buchla and Associates, Berkeley, Calif.), and Putney (also Synthi and Digitan, EMS, Putney, London). All of these machines work on basically the same principles for sound synthesis although there are differences in detailed configurations and human engineering.

The fundamental idea behind these synthesizers is to identify useful musical/acoustical parameters and provide means for programming changes in these parameters in order to provide for a continual change in the basic psychoacoustical parameters pitch, loudness, timbre, and duration. One powerful characteristic of the so-called "analog" synthesizer is the ability to connect the output of one "module" with the input of another. Thus the synthesizer is really just a collection of analog processors which take input data (in the form of control voltages) and produce output data (again in the form of voltages which may be either used controls for other modules or as final output signals (sounds). A module is generally provided to control each specific sound attribute which is desired to be controlled. Examples of typical modules are as follows: 1) voltage-controlled frequency module, 2) voltage-controlled gain module (or 2-quadrant analog multiplier), 3) voltage-controlled filter, 4) envelope generator, 5) random voltage generator, 6) voltage-controlled reverberator, 7) sequenser (which produces sequences of preset voltages at constant or variable tempos). Other auxiliary modules which do not refer to any specific parameters but which are used to combine signals are 8) adders (mixers), 9) multipliers (ring modulators), 10) voltage shifters, 11) exponential and logarithmic converters, 12) nonlinear processors (general type).

Although there is now a tendency toward inclusion of digital memories in synthesizers, most synthesizers generate sounds on the basis of preset values of potentiometers (a rudimentary form of memory) as well as real time tactile input from a performer/composer. As long as real time performance is the goal, the quality of composition (or performance) depends strictly on the flexibility and facility provided by the machine and the skill of the performer. For listeners who can tolerate redundancy (or even thrive on it) a synthesizer can provide a ready-made self-composition box without even the requirement of human performance. On the other hand, if a voltage-controlled synthesizer is used to generate single sounds or short successions of sounds and is used in combination with conventional tape editing and montage techniques, music composition by this method should surely be a great improvement over the method combining manual generation of sounds and tape manipulation which characterized most of the first work in electronic music in the early 1950's. Unfortunately, most composers have not taken advantage of this improvement in capability for timbre synthesis. This syndrome may be explained as follows: a) Many very good composers have become bored with electronic tape music as a medium and have turned to real time electronic synthesis, to mixed media (light, sculpture, dance, drama, etc.), or to instrumental music. Most of these composers lacked vision or interest in the possibilities of timbre synthesis, became completely disenchanted with the amount of tedious work involved in tape production, and simultaneously have decided that the performance of tapes has a small impact on the public (or more appropriately, the musically enlightened public) compared to personalized performances. b) On the other hand, many others presuming to be composers have decided to take the easy way out and let synthesizers control their compositions to such a degree that the frequencies of clichés becomes overwhelming.

For the systematically inclined composer who prefers to work from a score the synthesizer offers little hope. The principle difficulty is that although there are typically large numbers of parameters available for controlling musical structures there are no systematic techniques for programming successions of these structures from musical scores. The closest thing to programming capability rests with the sequencer, a device which suffers from limited memory and inability to read from any convenient permanent representations of scores, such as paper, cards, or their magnetic equivalents.

I should mention in passing that many artists who have become disillusioned with the restrictions of commercial synthesizers have decided that the best solution to the problem of producing electronic music lies in their undertaking the design and construction of their own custom equipment. There are many examples of this phenomena in the United States, one of the most notable being Gordan Mumma. In a sense, the machine is the composition. This of course, makes it almost impossible to perform a composer's composition without the presence of both the composer and his machine at the place of performance. My own personal observation is that in the majority of cases the research involved in preparing the machine becomes the principal activity, and that the musical activity becomes secondary by comparison.

Composition by Instruction to Computer: Total computational synthesis

The first use of a general-purpose computer for musical sound synthesis was demonstrated by Max Mathews at the Bell Telephone Laboratories (Murray Hill, N.J.) about 1960. The basic technique has remained fundamentally unchanged since this time although the computer program has been expanded and refined over the years culminating in the Music 4 and Music 5 programs. A feature of any computer system is the inherent capability for the utmost precision in combining musical events from direct or algorithmically computed score. Unlike most analog apparatus, repeatibility is insured. One of the most attractive features is the generality of the method: Conceptually speaking, any sound which can be described mathematically can be generated by a computer system. The method consists of the computation of large strings of numbers which are fed continuously to one or more digital-to-analog converters at rates of typically 20,000 numbers per second. While it is possible to consider the generation of numbers in real time, any reasonably complex sound synthesis example will exceed the capabilities of even the fastest serial computers available today (e.g. the IBM 360/91). Therefore one must recognize at the outset that non-real time operation is necessary and intermediate storage of the number strings is necessary. At 20,000 numbers per second, 12 bits per number only about six minutes of sound (monaural) can be stored on a typical full reel of digital tape. Thus, the storage of sounds in digital form is quite expensive compared to conventional audio storage. However, the most insurmountable problems at present for becoming involved in computer music are 1) access to a general purpose computer, 2) cost of computation time, 3) computer turnaround time, and 4) availability of the digital-to-analog conversion system, which usually requires an off-line digital-to-analog converter coupled to a high speed magnetic tape unit. Because of these problems very few musical compositions have been competed by total computer synthesis as compared to the use of manual tape and synthesizer techniques.

Probably the most efficient users of the computer music technique are members of a group at Stanford University, Leland Smith and John Chowning (composer/programmers) and David Poole (system programmer/musician), who thanks to the good graces of Artificial Intelligence Project (directed by Prof. John McCarthy) have been able to make extensive use of the A.I.P.'s very powerful computer system. After six years or so of concentrated experimentation they have developed programs and techniques allowing them to create by 1971 several quite lengthy and very sonourously sophisticated compositions. Having worked at the A.I. lab during 1968-69 (on speech recognition techniques) I feel that I was able to observe the work of this group in some detail. The A.I. computer system consisting of one PDP/10 and one PDP/16 computer which share a 200K core memory, a Lieberscope swapping disk, a large IBM2314 disk system (roughly 500x106 bits storage), 8 III keyboard/

displays, 40 or so small keyboard/displays (one in each office), very high rate A/D input for TV monitor as well as audio, multi-channel D/A output for sound and for control of artificial arms and TV servos, etc.(These figures have been quoted from memory and should not be considered as accurate.) Extensive programming by several very good systems programmers have been made for the development of a system which allows time-shared editing of files, interactive computation using graphic/alphanumeric displays, and use of special compiler languages, while at the same time providing for high rate continuous analog input and output. It is probably the only system of this type in the world. Since David Poole is the principal systems programmer and engineer of the A.I. Project and has from the outset been involved with the computer music project, Smith and Chowning have been in a particular good position to capitalize on all the capabilities of the A.I. computer system.

Two years ago I visited the A.I. lab, and Leland Smith demonstrated the latest version of his SCORE program to me. With me was a composer friend from San Mateo, Calif. who had brought one of his own conventional music scores. Sitting at the display console, Smith took us through the process of realizing 4 bars of the composer's score. The score was translated into computer notation, typed in at the console and data were computed. For some reason, it took about 30 minutes before we heard the first sounds. During the next three hours we went through about 10 permutations and modifications to the basic score. Smith demonstrated some very immediate techniques for performing such operations as adding parallel thirds and octaves, producing thrills, altering pitches, and adding ornamental lines. The overall effect was very pleasing. It should be noted that we were operating on Saturday morning, a time when there was very little competition for computer time.

We might compare this to my own experience with Music 5 at the University of Illinois, a much more typical computer music system (and probably better than most). In fall 1969 we obtained the Music 5 program from Bell Laboratories via Stanford University after a programmer at Stanford had spent several months in adapting it to the IBM 360. Within 2 months we had the program running on our 360 and obtained our first output. However, there was something wrong with the quality of sound, and our programmer could not determine the cause of the problem. Also, there seemed to be a problem in the use of the D/A system (which was then performed on an on-line 1800 computer attached directly to the 360 system.).

For one year we concentrated on programmatic aspects of Music 5, improving various features of it. In 1971 the 1800 computer was removed and a hybrid CDC 1700/EAI 680 was delivered which had the capability for D/A conversion. By borrowing a D/A conversion program from Purdue University (Gary Nelson) we were able to reinstate sound output by fall 1971. A student was hired to work on the project, and he managed to generate one music example by spring 1972. Up to this time, I had not been personally involved in the project except on an advisory basis, but in spring 1972 I decided to take over the task of programming myself. During the period June, through September, 1972, I uncovered several bugs in the program which had been nagging us for some time and made several additions and improvements. Beginning in October, 1972 Music 5 was fully operational at the University of Illinois and ready for use by composers and other interested persons. During the fall of 1972 several students studied the use of Music 5 as part of our regular course in Musical Acoustics.

Operation of Music 5 at the University of Illinois

Anyone who has a valid computer problem specification number can use Music 5. The program is stored on the 360 library disk and can be accessed by any user who has the proper "job control language cards". Input consists of a) "score cards" or b) FORTRAN subroutines or c) a mixture of both. After a score is determined, the necessary cards are punched and handed to a computer operator who enters the job into the "job stream" and immediately returns the deck to the user. If he wishes he can modify his deck and run another job. Music 5 jobs are of two types: 1) Calcomp plotter output which gives a plot of the waveforms corresponding to the programmed sounds. 2) Digital tape output, which contains the numbers to be converted into sounds using the hybrid computer system. Since the hybrid computer is in a building two blocks away from the IBM 360, there is a considerable delay in obtaining sound output, and most users will use the plotter output until he is certain that his program is working properly.

Advantages of Music 5:

1- Generation of virtually any sound. The testing of any acoustical model is reduced to simple FORTRAN or SCORE programming.

2- Precision in combining several layers of sound and rhythmic precision.

3- Elegance as a demonstration of the relationship of algorithms, acoustics, and music.

Disadvantages of Music 5:

There are many. Here are four:

1. The most severe limitation is the lack of short turn-around time between coded input and sound output. Several of the bottlenecks are: a) Punching cards is a cumbersome process. b) Sheer computation time slows down the

process. c) Output is in the form of a digital tape which must be hand-carried to another building; the D/A system must be scheduled for use. Note that the Stanford people have eliminated problems a) and c) at the cost of using their computer at odd hours.

2) The input must be conceptualized in terms of a very rigid, unintelligent language which tolerates no or few mistakes. It therefore is necessary to debug the input score and program to a) make the program work at all, and to b) make it correspond to the composer's desires. This is a severe road block to the creation of music and is common to any computer system which requires alphanumeric input.

3) Acoustic limitation: For the price of precision one pays in terms of the necessity of quantization of the signals and parameters of sounds; quantization of both amplitude and time results in an approximation to the desired result in such a way as to often generate audibly significant error signals in the output. For example, an attempt to generate an ordinary sawtooth tone will produce unwanted inharmonic frequencies by virtue of the "foldover effect". Amplitude quantization can cause clipping at high levels and a sort of "modulation noise" at low levels unless a sufficient number of bits are provided.

4) Cost of computer time. Some complex sounds may cost $100 per minute, which is much more expensive than by use of a synthesizer.

Combinations of Computers and Synthesizers: The Hybrid Computer

Even though computers are very fast, their speed is limited by their serial nature. Computations must be done one at a time by a single central processing unit. The obvious solution to the problem of speed is to provide for computation in parallel so that several parameters and signals may be generated simultaneously. Many technologists and composers have been fascinated by this problem during the past several years, many systems have been proposed, and a few are actually in operation today. The number of possibilities for software programs for non-hybrid computer music systems and the number of possibilities for module design for synthesizers is essentially limitless, but at least a certain amount of standardization has taken place. However, with hybrid computers the number of possibilities for system design increases again by a large factor. Probably the most obvious step is the connection of a small computer(e.g. a PDP-8 or PDP-11) to an existing analog synthesizer.

A synthesizer may be connected to a computer by means of a suitable interface which usually consists of a set of D/A-converters, one for each parameter to be controlled, each of which may be individually addressed by the computer. The outputs of the D/A converters are analog signals which are used to controle the input parameters of the synthesizer. In addition, logic pulses can be used to time the onsets of envelope generators, sequencers, sample-hold circuits, etc. Of course, there is a very definite problem because not all parameters used to determine the state of a typical analog synthesizer can be set remotely. Most commercial synthesizers have many knobs which must be set to calibrated positions. Some parameters, such as range positions on VCO's and attack and decay times, are usually not voltage-controlled, but are set manually. Therefore, it is much more logical to consider providing a special synthesizer for the hybrid system where all saliant parameters are voltage controlled, and calibration and accuracy present no problems.

Most synthesizers use patch cords or matrix switches, to provide interconnections between the various modules contained in them. For full automatic control by computer switch networks must be provided to allow the computer have the possibility for set up the patch configuration. If every output is to be connected to every input a large number of programable switches must be provided. For example, if there are 100 outputs and 100 inputs, 10000 switches must be provided. Obviously, some compromise must be made to reduce the number of switches which would take into account these patches which are most commonly needed.

Once the hardware interface and synthesizer configuration are established the problem becomes one of providing software to translate input scores into output data within the constraints imposed by the computer system to deliver data to the interface. Whereas in total computer music real time operation is not expected, the hybrid computer system utilizing a comparatively small computer should hopefully be capable of delivering feedback in the form of sounds in a comparatively short amount of time. Therefore, a great deal of attention is necessary to decide what input language is to be used, how the language is to be translated into data, and how and what rate the information is to be delivered to the interface. The output data rate is equal to the number of parameters to be controlled times the average individual data rates to approach those of total computer systems. Obviously, the more sophisticated the synthesizer is in terms of its own internal memory and in its ability to produce transitions (frequency and amplitude curves, etc) the less difficult is the problem of the computer producing high data rates for the synthesizer.

A minimal low cost computer system typically includes a teletype unit and a high-speed paper tape unit, but one soon finds that all sorts of peripherals such as disk, cartridge tape (or DEC tape), line printer, graphic displays, and high speed tape units are needed for fast program development, computation, and production of sophisticated control signals for the parameters.

There are five hybrid computer systems in the world utilizing voltage-controlled synthesizers which I am aware of and have already been built. The first was a connection of special voltage-controlled apparatus to an IBM computer via a standard programmable potentiometer interface; this system was developed by James Gabura of the University of Toronto, Ontario, Canada during 1966-68. The synthesizer was essentially like the Moog except for a large number of organ filters which were switched in and out to provide a variety of tone colours. Not much music was produced with this system, although I heard a very nic Bach Fugue example (synchronized four parts) rendered by the system. The computer was disbanded about 1969, and since it was owned by the Computer Science Department there was nothing that Gabura (at that time a student) could do to keep it. In the meantime, two other hybrid computer systems were assembled starting about 1968: The GROOVE computer at the Bell Telephone Laboratories and a hybrid system built by Peter Zinovieff in Putney, London. In the last year or so two other hybrid systems have materialized: One by Edward Kobrin at the University of California at San Diego and the other by Donald Buchla of Buchla Associates in Berkely, California (in cooperation with the California Institute of the Arts, Valencia, Calif.).

Since most of these hybrid systems are in a state of flux it is difficult to give an accurate comparison between them. The GROOVE system has been documented as of 1970 ("GROOVE, A Program to Compose, Store and Edit Functions of Time" by M.V. Mathews and F.R. Moore, Bell Laboratories Internal Report (expaned version) and same title and authors in the Communications of the ACM, Vol. 13, No. 12, Dec 1970 (condensed version)). The computer used is a Honeywell DDP-224 which was purchased for speech research at Bell Labs.

The peripheral configuration consists of 16k, 24 bit core (1.7 usec cycle time), CDC 9432 disk file with removable packs, a typewriter console, high speed magnetic tape unit, 7 input A/D, and 12 8 bit and 2 12 bit D/A converters. Both input and output are utilized for the music program, which incidentally, is normally operated only after 6 PM and during weekends. Music is scored in terms of control functions which are created in real time by means of inputs from 4 potentiometers, a three dimensional wand, and a special keyboard. They may also be defined algorithmically. Functions are stored on disk and may be observed on a display screen. They also may be edited and modified with the results immediately available on the display. Functions may be combined using various simple algebriac operations. Short sections of music can be immediately synthesized, and the rate of sampling the control functions can be specified at the console. The system is fast enough that a certain amount of performance with real-time audio feedback is possible, using knobs or the keyboard to vary parameters.

While the interactive software of the GROOVE system seems very powerful (I must hedge here as I have not used the system yet.), the system at present is limited in sound synthesis capability. The sound synthesis system consists of several "electronic voices" each of which consists of a cascade connection of a voltage controlled oscillator, a voltage controlled amplifier, and a voltage controlled filter. Because they have had some trouble with the stability of their voltage-controlled oscillators, the oscillators are calibrated before each run by sending out a series of voltages to the VCO's and measuring their output frequencies by means of an A/D converter input. They are now considering switching to the use of digital oscillators. As far as I know, envelopes and frequency transitions are created in the computer memories. Several compositions have been generated by Vladimir Ussachevsky, F. Richard Moore, and Immanuel Gant. The musical textures that I have heard have been of an instrumental type, ranging from brass-like sounds to organ-like sounds to percussion sounds. I believe the limitation is basically with the synthesizer, which is actually quite simple in sonourous capabilities compared to a full-blown analog synthesizer. Another limitation is that patching must be done manually, although I understand they now have plug-in patch panels.

Another well-known hybrid system facility is directed by Peter Zinovieff in London. Although I have been in contact with his group from time-to-time, I have never heard any of the compositions generated by their computer and so can make no judgement about its success from an aesthetic point of view. An article by Zinovieff was published in 1969 ("A Computerized Electronic Music Studio", Electronic Music Reports, No 1, Utrecht, September 1969.

The Zinovieff studio is based on two small computers, a PDP8/S and a PDP8/L which has the usually complement of tape units and disk. It has the capability for multichannel A/D input and multichannel D/A output. Editing of existing music as well as synthesis of new music can be accomplished. They also have a program for detecting the pitch of voice or instruments and recreating the idential melody at a later time with a synthetic timbre. The output interface consists of 32 12 bit storage registers, some of which are connected to switches and some of which are connected through D/A converters to various voltage-controlled apparatus. The time of events is determined by an external programmable clock which feeds back an interrupt pulse to indicate when the event has terminated so that new information can be generated to the interface. Programmable units are several voltage and digitally controlled oscillators, voltage controlled amplifiers, and voltage controlled filters, a large number of digitally controlled filters, and some programmable attack/decay circuits.

Early programs for the Zinovieff computer system were specialized for particular compositions and were written in assembly language (PAL) making extensive use of macros to provide a hierarchy of relationships to allow a sophisticated control of musical textures and relations between sounds (at least this was the intention). Later, an attempt was made to provide general purpose programs. However, it is my feeling from what I have heard about the studio that few composers have been able to gain access to it, and that program development has suffered from the lack of musical problems to be solved.

The computer system of Edward Kobrin (U.C.S.D.) is being designed with a view towards real time interactive performance. He is working with a PDP8/L in conjunction with a bank of interconnected VCO's, VCF's and VCA's, and a matric system for controlling the outputs to 16 speakers. The analog system he is using is not at all well calibrated and stable, but he claims that for the experiments he presently is involved in, such accuracy is not required. The basic technique is to provide tables of values in the computer memory which are assigned numbers or names. When the computer is placed in a performance mode, the user types in a code at a keyboard/display pertaining to the table to be used and the device to which the table is to be directed. Depressing another button initiates the transfer of data from the table to the device. Timing information may also be entered into programmable counters to control the time at which a new set of parameters will be entered. The interactive mode does not require any tape or disk units. The entire program works with 12 k of core memory. However, the output is in the form of several interacting loops which continue to perform under machine control until the computer user intervenes to change the status of the output. The program was developed at another computer installation (Argonne National Laboratories) which had an extensive peripheral configuration. It should be emphasized that the purpose of this computer system is not for score processing, but for real time performance.

In a similar vein, Salvatore Martirano of the University of Illinois has constructed a hybrid digital/analog machine, called the SAL MAR Construction which is used for real time performance. However, this machine does not utilize a digital computer to operate on data in an arithmetic sense. Digital information is entered at a console by touching some of the more than 200 light buttons available. This information is used to interact with RAM memories to set values into shift registers, and to load data into a number of special 32 segment waveform synthesizers. Waveforms are constantly changing by means of a special technique for blending old and new waveforms. The analog portion of the machine consists of 8 voltage controlled oscillators (2% accuracy, 20 to 20000 Hz range), 40 programmable percussion generators, and 8 programmable attack/duration/decay circuits. A technique called exclusive OR feedback on shift registers is used to provide finite sets of values in an analogy to serial permutation technique. The machine appears to compose by itself, providing a surprising variety of rhytmic and pitch movement. The performer controls the "flow of information" (direction and quantity) by entering new top-level information at the console. Another prominant feature of the machine is its synthesis of spatial aspects. 24 speakers are used to control the sound location of four independent programs, and reverberators may be controlled to provide distance cues. A special matrix system was designed to control the distribution of sound.

Systems involving Digitally Programmable Synthesizers

Digital synthesis modules may be of at least three types. First, there are analog modules which are programmed digitally by switched to control values of internal components. For example, a digitally-controlled oscillator where capacity or resistor values are switched. Second, there are modules which utilize digital circuitry, but the output is in analog form. Third, there are modules which utilize digital circuitry but the output is in digital form (i. e. a string of binary numbers). The first electronic example of a type one synthesizer was the wellknown RCA synthesizer. This was programmed by paper roll punched by typewriter in a binary code for the control of equal-tempered 12 tone scale, selection of waveforms and filter frequencies and attack and decay times, all as functions of time. While only four voices could be generated at once, it was possible to synchronize tape channels in such a way as to build up complex textures. One very well known piece composed on this machine in the early 1960's is "Composition for Synthesizer" by Milton Babbitt. All switching of parameters was accomplished by mechanical relays and so parameter changes were necessarily descrete. This had a definite influence on the style of music created by this machine. Timing was very precise, but smooth transitions and complex attacks were difficult to create.

The second machine of this type was created at the Siemen's studio in the late 1950's. This machine was run by coded paper tape as prepared by device which incorporated a conventional music keyboard. Apparently the system was not used extensively by composers as no well known pieces resulted from it (that I know of) and the Siemen's studio was dismantled around 1962.

The third studio to utilize digitally controlled analog synthesis apparatus is the studio at Elektronmusikstudion in Stockholm. 24 oscillators, 2 28-channels filter sets, 3 ring modulators, 2 amplitude modulators, and 100 attenuators are fully programmable by input digital information. This analog complement is under full control by a PDP 15 computer. No patching is required. The entire music console, upon which touch switches display the value of each parameter and the interpatching arrangements, is fully determined by the computer output within time resolution of 1 millisecond, althoug 20 millisecond resolution is usually used.

The PDP 15/40 computer system is probably the most powerful system in the world to be placed for the exclusive use of composers on a 24 hour basis. The software presently available is not interactive in real time although short examples may be synthesized within a few minutes using programs such as EMS-1. Longer examples take a longer period of time to compute, although it is difficult to give any real time vs. computer-time ratio figures. There is a certain sound limitation due to the nature of the synthesizer (this is true of all synthesizers, especially if they are used unimaginatively), but by means of clever programming of parameter contours, a large range of timbral variety can be achieved. The most important consideration is that the system is in operation, has the capacity for very sophisticated sound synthesis and compositional algorithms, and is available to composers who desire to work with computer synthesis techniques on a day-to-day basis. In this sense, it must be considered an experiment on a grand scale.

SUMMARY: A Discussion of Desireable Hybrid Computer Techniques

a. The Synthesizer Subsystem

One problem is that hybrid systems can be limited by the finite complexity of the synthesizer portion of the system. This is not a conceptual problem with a tape studio or with total computer music, since layers can be built up by montage in the former case, and by increased computation time in the latter case. The hybrid system, if it is to do the job it was designed for, should generate a complete composition in one shot. These means that a suitably complex orchestra should be provided. The make up of the synthesizer complex should be chosen on the basis of available psychoacoustic theory and musical experience as that which will achieve the largest variety and complexity with the minimum number of components (or parameters).

I believe in the Music 5 philosophy that all synthesis units should have inputs and outputs and that any interconnection between the units should be possible. As long as the desired accuracy is obtained, it does not matter whether the circuits are analog or digital or some mixture of both. However, I believe that the outputs and inputs should be compatible in order that arbitrary interconnections can be made. Digital circuits run the risk that descrete steps may be perceptible. (According to a psychoacoustician I talked to recently, .1 Hz (or .1%) resolution is required to prevent the perception of frequency steps and .025 db resolution is required for amplitude.) On the other hand, analog circuits run the risk of being unstable, although there has been a large improvement in this area during the past few years. As long as the ear cannot tell the difference and parameters are being programmed by a computer, the factor that should govern which type of circuitry to use should be economic.

b. The Computer Subsystem

The computer subsystem, which is a mixture of software and hardware, should be designed to make music production as convenient as possible within budget constraints. A input language for musicians free of system gobledegook is essential. For experimenting with single sounds or short musical passages, a real time interactive program is extremely desireable. A method where previously determined sound objects can be named and scored within a hierarchial framework for building up large compositions is very useful for composers. The possibility for testing, editing, and modifying sounds at various points in a composition is necessary to increase the flexibility of the system and to provide an insurance against time-consuming failures due to computational "bugs". In terms of hardware I think a display input with light pen or wand control is a very powerful tool for the composer.