Research Proposal Submitted to the National Science Foundation

Proposed Amount: $618,640. *Proposed Effective Date*: January 1, 1975 *Proposed Duration*:

Title: Computer Simulation of Music Instrument Tones in Reverberant Spaces

Principle Investigator: Leland C. Smith *Soc. Sec. No.*

Submitting Institution: Stanford University

Department: Music

Address: Stanford, Ca.
94305

Co-Principle Investigator: John M. Chowning *Soc. Sec. No.*

Make grant to: Stanford University

Endorsements:

Principle Investigators *Dept. Head* *Institutional Admin. Official*

Name:

Signature:

Title:

Telephone:

Date:

Name:

Signature:

Title:

Telephone:

Date:

COMPUTER SIMULATION OF MUSIC INSTRUMENT TONES IN REVERBERANT SPACES

B. PROPOSED FACILITY

I. HARDWARE FACILITY

The process of analysis and synthesis as described in the preceeding sections demands great quantities of computer time. The computations can take as much as 100 times the length of the produced tone. This requires great patience from the researcher. Since many of the results are empirical, many experiments have been done and must continue to be done. For this reason, we propose a special-purpose computing system. The system would initially be a satellite to the AI lab, but would serve not only to reduce the computation load of the PDP-10, but by means of special-purpose hardware, actually synthesize tones in reverberant environments in real time. The system would be powerful enough to stand alone if necessary. Since it would be extremely expensive to purchase a system with the human engineering of the AI lab facility, we propose to use the AI lab facility as an interface to the special-purpose system, thus minimizing change-over inconveniences and initial setup price. As work continues, it would be possible to upgrade the system to stand alone and eventually provide a high degree of service without the aid of the AI facility. To this end, we propose using an existing time-sharing system as the resident monitor in the special-purpose system. This saves us the trouble of having to write device controllers, memory management programs, and other system-level functions. Since we propose to begin with a time-sharing system, upward compatibility is assured. Programs will continue to run unmodified as the system is upgraded. Although the detailed hardware budget is given in section V, we shall discuss the main items here. The following is a list of the principal componants of the proposed facility:

Digital Equipment Corporation PDP-11/45 Computer
with floating-point unit
and memory management module

Digital Equipment Corporation RP03 disc drive.
Provides 20 million 16-bit words of storage.

Systems Concepts signal processor with digital
reverberation module

Digital Equipment Corporation GT-40 graphics terminal

Audio equipment, including 8-track tape recorder,
quiet booth, acoustically treated room,
8-channel amplifiers and matched speakers

The heart of the system is a PDP-11/45 with floating-point processor and memory management module. This is a powerful mini-computer capable of high-speed arithmetic and basic time-sharing operation. The purpose of such a computer is to provide a versatile test bed for research. The floating-point processor greatly aids numerical computations of the type which are so common in digital signal processing and multi-dimensional scaling. The memory management module provides for time-multiplexing of multiple tasks, so as to more effectively utilize the computer by running processes while others are idle. It also provides for cooperating parallel processes which aids decomposition of large tasks into smaller modules.

For bulk storage, an RP03 disk system is included. This is necessary for real-time operation. The bulk of the disk is to store digitized audio and control information for the signal processor. In analysing music instrument tones, we digitize recorded natural tones to a precision of 14 bits (stored in 16-bit PDP-11 words) at rates up to 50,000 samples per second. One can easily see that to store individual notes of each of the orchestral instruments in all of their characteristic playing modes would be a staggering amount of storage. The RP03 disk can provide storage for up to 400 seconds of sound. This is a compromise between price and working requirements, representing the minimum amount of storage that can be effective and useful.

An alternative to getting an RP03 disk would be to use the AI facility's disk storage. This is unwise because of the required rate of data transfer between the machines and would require a high-speed data channel for communication with the AI PDP-10. The most efficient solution is inclusion of some amount of bulk storage on the PDP-11 itself, and for real-time operation, this is the only solution.

The most important item of all is the Systems Concepts Signal Processor. This is a highly-parallel, special-purpose, programmable, digital processor designed especially for the generation and processing of audio. Together with the matching reverberation unit, it provides enough power and flexibility to synthesize all the tones we have produced to date in real time, even with spacial localization and 4-channel reverberation involved. It can even do the computations in the heterodyne filter analysis in real time. The processor is designed to be controlled by a small computer and a PDP-11 interface is a standard option. The utility of such a processor can not be overemphasized.

When one attempts to generate sounds using a new technique, often there is no good way to make an *a priori* prediction on the range of the controlling parameters. A good example of this is deciding over what range to sweep the modulation index of an FM instrument. With our current turnaround time, one must "shoot in the dark" in attempting to find the correct parameter, often wasting tremendous amounts of computer time as well as personal time in the process. The PDP-11/45 by itself offers no speedup, but combined with the Systems Concepts Signal Processor, provides the solution to the problem. It would be possible to directly connect a knob, via some PDP-11 support software, to synthesis parameters, thus allowing the experimenter to directly control the parameter as the sound is generated. This increases the efficiency of the research process immensely. It also makes possible experiments which would be otherwise impractical or even impossible. One example of this would be a two-knob experiment where one knob controls the duration of a note and one controls the loudness. To do this experiment without real-time generation of the sound would require preparation of a large number of sample sounds beforehand. With three knobs, the size of such samples exceeds even the bounds of the entire AI lab bulk storage. It is much more efficient to store a sound as the program which can synthesize it rather than the digitized waveform itself. A whole new domain of experiments immediately suggests itself, based on interactive control of complex attributes of synthesized sound.

We wish to add to the signal processor the reverberation memory option. This device provides for a number of variable-length digital delays which are easily interfaced with the signal processor itself to provide reverberation in all the forms we have realized to date, and have enough generality to provide for any future forms of reverberation we may discover. Again, the parameters of the reverberation could be easily attached to knobs, giving the user direct real-time control over the character of the reverberation. This would greatly enhance the productivity of the researcher.

Again, such a device as the Systems Concepts Signal Processor could be interfaced directly to the AI PDP-10, but it requires real-time control of the type that is not generally available in time-sharing systems.

The GT-40 display console is a general-purpose graphics terminal. It provides a direct graphical interface to the PDP-11 and thus to the Signal processor. As was noted many times in the previous sections, the use of computer graphics is an essential piece of human engineering that we take advantage of constantly. The AI facility has made graphics highly available and easily used. It thus has crept into many programs as a debugging aid as well as a research aid.

The audio equipment is the final stage of the sound production. The waveform must be converted to sound by an audio system whose quality matches the extreme quality of digital synthesis. We propose to place 8 speakers in an acoustically treated room for listening tests. The number 8 is somewhat of a compromise with cost, as we have little *a priori* reason to believe that 8 channels will be enough. Our only real evidence is our success with 4 channels.

The importance of the listening room is great. The environment of a computer facility is very noisy. Computers require extensive cooling and air conditioning, making the computer room sound somewhat like a continuous hurricane. In addition, the particular building we are in uses forced-air cooling in each room, adding a gentle but disturbing hiss to all offices. To do perceptual testing, it is essential that the acoustical environment be controlled entirely by the investigator.

A small quiet booth is also proposed for recording and analysing sounds from real music instruments. The 8-track tape recorder provides off-line storage of acoustical data. Audio tape is still the most economical way to store large quantities of audio, at the cost of some loss of fidelity and signal-to-noise ratio.

2. PROPOSED RESEARCH SUPPORT SOFTWARE

An amount software would have to be written to integrate the proposed system into the current system. This could be done in layers, each preserving compatibility with the previous systems but each advancing the researcher's capabilities. The following discussion attempts to summarize the software which would have to be developed.

THE PDP-11 MONITOR

Although we intend to use the UNIX time-sharing system from Bell Laboratories as a base, it will need some amount of modification to deal with our specific set of peripherals: the GT-40, the RP03 disk, the PDP-10 interface, and the Signal Processor. It is not expected that this would require a great deal of effort to get a basic monitor up and running. Improvements and streamlining can always be added later as the need arises.

THE SIGNAL PROCESSOR

Writing software for the Signal Processor presents somewhat of a problem. One must provide methods not only for setting up the internal computation flow, but also for synchronization and data transfer with cooperating PDP-11 processes. One example might be interactive control of mixing natural and synthesized sounds. The PDP-11 would be reading analog to digital converter output, interactive knob settings, as well as controlling the signal processor and providing it with the parameters read from the knob and the waveform from the A/D converter all at once. The theory of cooperating parallel processes, however, is rather well developed such that straightforward methods can be applied here with little difficulty. It is a great advantage that the UNIX time-sharing system provides for parallel processes and has a highly developed inter-process comunication system. Only slight modifications would have to be made to provide for real-time interaction. The efficiency of the inter-process comunication system would have to be examined and possibly reprogrammed to assure that it does not incur prohibitive overhead.

One would write a program for the Signal Processor just as one writes a program for a computer. At first, the language would be a low-level one, corresponding to an assembler for a computer. As we learn more about controling such a device, the language could be upgraded to a higher level. One would specify communication with other processes through a system of ports. At run time, the user would be able to specify which ports should be connected to which channels. It is then that the user would specify which communication ports should be directed to which knobs, and which ports should be connected to which disk files, and so on. This maintains a level of generality which permits great flexibility in testing and repeating a run.

At first, we would merely seek to construct an interface with our existing PDP-10 synthesis programs such as to ease the change-over problems. This would allow for direct synthesis of all the sounds we have generated to date, and would allow research to proceed at a much greater rate, but would not allow for real-time interaction until the assembler for the Signal Processor was completed.

Since many of the analysis routines, as well as synthesis routines, could be executed on the Signal Processor, one would want ways to send digitized waveforms to the Signal Processor. It should be possible to come in directly from the analog-to-digital converter as well as retrieve stored waveforms from the disk. The output of the signal processor could be directed to the digital-to-analog converter as well as directed to a disk file. The system could also get at files on the AI Lab IBM 3330 disk, but only at a limited rate.

This ability to store the output of the signal processor is quite important. The most compelling reason is one of growth. As we progress, it is quite possible that we will demand computations so complex that even the Signal Processor can not complete them in real time. In this case, we must have the option of directing the output to the disk. In this manner, we could break up the computations into smaller runs and mix them in the PDP-11 for later playback.

There are also many utility routines that would have to be written for the PDP-11. We would need an implementation of the fast Fourier transform algorithm, an intermediate-level graphics package for convenience in displaying complex functions, routines for getting at the digital-analog converters, as well as ways to direct asynchronous processes in a uniform and convenient manner.

All of this is programming support that would be needed for interfacing our current research to the new system. Eventually, we would begin writing research programs directly on the PDP-11.