JOHN R. PIERCE
931 Canon Drive
Pasadena, CA 91105

April 23, 1983

Charles S. Smith
System Development Foundation
181 Lytton Avenue, Suite 210
Palo Alto, California 94301

Dear Charlie:

I'm very sorry that I didn't get to see your before we left for Pasadena. I did talk to Nancy briefly over the phone.

First of all, Ellen and I thank you very much for having made it possible and pleasant to spend a month and a half at Stanford. We had planned to stay longer, but came back here for two reasons: the gardener reported that the water was off, and we were afraid that the garden would burn up (fortunately a kind and efficient neighbor got a plumber to fix things before we returned), and, we heard that some people who had looked at our house several years ago when it was on the market might want to buy it. They came to see the house today but said nothing definite.

I enjoyed CCRMA and Stanford very much. John Chowning seldom came to CCRMA during the first part of our stay, because of his son's illness, but the problem has been resolved and he was there full time during the rest of our stay. He and many others at Stanford were very friendly to Ellen and to me, but Ellen still feels an attachment to Pasadena. We looked at a good many houses and condominiums, off campus and on, and, of course, found nothing as nice as what we have. We did see some tolerable places. Further, John Chowning thinks that he can get an appointment as Research Professor through, and it does seem to us that it would be preferable to live on campus if we come. Right now we must see what chance there is of selling this house for a reasonable price, and how we really feel about Stanford and Palo Alto.

That's the family news, so to speak. I have a few things of a technical nature to say, and I think I'll put them on other sheets of paper so as not to get them mixed up with this letter.

Thanks again. And love from Ellen and me to Nancy as well as to you.

Yours,

John

Writing at top reads:

"Dear John - I thought you might like to see this. Thanks a lot, + love from Ellen + me.
- John"

J. R. PIERCE'S ADVENTURES AT CCRMA, ~~February~~ March 1 to April 15, 1983

(Note: I can't find my CCRMA list of personnel, so don't treat the names given below with too much confidence).

I went to CCRMA with three projects in mind:

I. I wanted to establish a working contact between Carver Mead's music synthesis chip work and potential users at CCRMA. It appears that the day-to-day work on the chip is being done by John Wawrzynek. In John Chowning's absence, Patte Wood invited him to come up for a couple of days at CCRMA expense. The meeting was arranged for April 6 and 7. I thought that it was very profitable. There was a "formal" meeting at which John gave a presentation on the nature and state of the chip to a half-dozen very interested people. Then Chris Chafe played a number of good computer pieces for him, and told him how they had been synthesized. Wawrzynek seemed particularly interested in the Chant method of voice synthesis, and Chris supplied him with an article on it. David Jaffe (who produced SILICON VALLEY BREAKDOWN using a delay line method of string synthesis) was away, but Julius Smith is very familiar with Jaffe's work, and explained that and his own violin synthesis model, which incorporates the mechanics of bowing. Julius also had an interesting mathematical analysis of a difference-equation method of generating sine waves that can be frequency modulated (the algorithm that Wawrszynek has been using doesn't allow this). Julius and Wawrszynek got on very well and exchanged both telephone numbers and ARPA electronic mail addresses. I think that further and mutually fruitful cooperation is assured.

II. I wanted to verify what Bregman and Jean-Claude had told me when I was at CCRMA last July -- that the Wessel streaming illusion is produced only when the timbres differ markedly in spectral content. I used PLA and the SAMSON BOX to synthesize some examples that convinced me and, I think, Earl Schubert of this.

III. I wanted to look into the validity of Terhardt's theory of pitch (which I am less and less inclined to believe). Terhardt gives a complicated formula for calculating the perceived pitch of any collection of sinusoidal partials. Andy Schloss, a student of Schubert's who was asked to acquaint me with the computer, was in touch with a friend at IRCAM who had a computer program (in LISP, alas) to make Terhard's computation. The program arrived at CCRMA before I left, but had not yet been translated into a runnable language; one of Earl Schubert's students was doing this. I produced a number of tones that Earl and I thought would be of interest in trying to understand pitch. We both felt that we had learned something, but further work was necessary at the time I left. Earl, who has lots of graduate students, may have one work along the lines we were pursuing.

Thus, I accomplished something of what I had thought at the start of doing. Perhaps what I hadn't thought of was more interesting.

ANALYSIS OF PERCUSSIVE SOUNDS

Andy Schloss (see above) is a percussionist with a degree in ethnomusicology from the University of Washington. At one time he made a good deal of money as a percussionist in a Broadway play. His heart is set on computer music, and he is a student of Earl Schubert's doing a thesis on drum and percussion sounds. He is quite and works the computer like mad. He was making spectral analyses of rapidly-decaying percussion sounds through a time window, and getting smoothish, not very informative spectra. I suggested that he take the spectrum of the whole sound and look for a filter-synthesis program that would find the poles. We talked with Julius and he gave Andy information on such a program (based on "the method of Prony"). Before he left on his lecture tour, Andy had found poles (and zeroes as well; this is a sort of artifact of trying to represent the drum sound by a single filter). He had not yet tried the filter as source of drum sounds, to see how well it agreed with the sounds analyzed, but he will. He seemed quite enthusiastic. I think that this method of analysis will make possible the analysis and synthesis of a host of rapidly decaying percussive sounds, including knocking on wood. I am quite enthusiastic about it. I explained it to John Chowning and Earl Schubert, and both were interested.

SYNTHESIS BY SIMULATION

Julius Smith has a wonderful computer model of the violin, including the mechanism of bowing. This makes it possible to get a variety of effects by adjusting a small number of parameters, as, force on bow, velocity of bow and place of bowing. Julius has worked closely with David Jaffe, who produced worrnderful plucked string sounds by a similarly "physical" model. In these models, the vibration of the string is simulated by launching waves on a closed digital delay loop with a proper loss as a function of frequency. A fine adjustment of the "string length" can be made by varying the parameters of an all-pass filter in the loop.

To my mind this sort of digital modeling of physical processes in instruments is a powerful new method of synthesis. While I was at CCRMA a former student of mine, Paul Milenkovic, who is now an assistant professor at the University of Wisconsin, wrote me concerning a new idea for finding the shape of the vocal tract during speech. He proposes to put a tiny microphone through the nose and near the vocal cords. The signal from this together with a signal from near the lips enables one to deduce the shape of the vocal tract. This would work for singing as well as for speech.

This and the work of Julius and Dave led me to think that it might be a good idea to synthesize the singing voice by using a delay-line model of the vocal tract. Why? Chris noted that in his fm syntheses he could get fricatives right only by adding natural fricative sounds. In a vocal-tract model one can get fricatives right by injecting noise at the right point along the vocal tract. Further, in producing consonants the resonant frequencies move rapidly near closure; this can cause various

troubles in both analysis and synthesis. However, the actual shape of the vocal tract changes rather slowly. I explained this to several people interested in the synthesis of voice, including Chris, Julius, John (I can't remember his last name) and, I believe, John Strawn, and Earl Schubert. All seemed interested, and so did John Chowning. I wrote to Paul Milenkovic, urging him on and asking him to send references on vocal tract analysis. This he did and I distributed all the material he sent rather widely.

A CHEAP AND POWERFUL SYNTHESIZER

Bob Shannon, Claude Shannon's son, has a master's degree in E. E. from Stanford and has passed his qualifying examination for doctoral work. He is Allen Peterson's student (Allen Peterson is Professor of Radio Science, but much of his work, and his students', is in complicated digital hardware). At CCRMA Shannon is building a real-time digital synthesizer on one circuit board. It can be played by means of a keyboard. He has paid for the parts and for the Atari computer he is using to control it out of his own funds. He is building the synthesizer because of his own musical interests, and nothing will stop him. Some aspect, perhaps the control problem, may be acceptable to Peterson as a thesis topic.

At first I didn't know what to make of this, but I saw remarkable progress as time went on. About the time I left Bob said that the synthesizer would produce 8 simultaneous voices, each with 8 partials, and an indepent envelope on each partial. He said that the synthesizer would be more powerful than the Samson Box, and the control strategy would be better.

Shannon works in close physical proximity to Phil Gosset, who is assembling the 6 CCRMA 68,000 terminals ("SUN TERMINALS"). Phil is very able; he has worked for Funely and is a member of a partnership that is producing a very advanced fast computer graphic system. Phil believes that what Bob is doing can be done and appears to respect Bob's competence. Bob was in one of Cal Quate's classes; Cal found him eccentric but very able; he neglected home work but did very well on the final exam.

It is my guess that a cheap, powerful real-time synthesizer, suitable for use with the "SUN TERMINAL" can be made and that Bob may do this in a relatively short time. A really useful terminal will require a better A to D converter than he is now using -- his is a 10-bit device. He knows this, but he has insisted on paying for all parts himself and he hesitates to spend more than he can easily afford.

I urged Bob to interact more with the musicians at CCRMA so as to be sure that his synthesizer addressed their needs. I also sent him a copy of a memo from Barry Vercoe on control problems.

Bob seemed very agreeable and like to talk to me.

SOME GENERAL REMARKS ABOUT CCRMA

CCRMA is a constellation of extremely various but remarkably able people. The predominance of musicians (who learn a lot about computers, or at least about using them, and about software) keeps CCRMA on the musical track. The electrical engineers and the computer people, hardware and software, who seem all to be very good, make it possible to do things. The psychoacousticians (mostly Earl Schubert's students), and an occasional psychologist, add a depth to the musical work.

CCRMA people are extremely varied in character as well as in field. Some, including Chris Chafe, Phil Gosset and Julius Smith are extremely reasonable. Others go their own ways with various degrees of inflexibility. But, in some strange way, all interact with a sense of mutual esteem and with mutual profit. CCRMA has a mixture of chaos and cooperation which is amazing as well as productive. All are very attached to CCRMA because they are doing what they want to do, individually and collectively. Bill Schotstadt, a musican by nature, went to work for Apple in the early days, and got stock options. When the stock went up enough so that he had a modest competence, he quit Apple and returned to CCRMA as a self-supporting participant. Others stay at CCRMA by working for various small software firms.

What causes CCRMA to survive and flourish? John Chowning is an essential ingredient. He is on good terms with everyone, and he manages to get essential things done. This he certainly doesn't do by ordering people around; even the tractable ones are good enough so they aren't orderable. He enlists their cooperation. One day I found Julius Smith doing something he described as "work" (as opposed to his own very fine work, which he would classify as fun). I suppose he was paid for the "work", but chiefly, he knew that it was necessary to CCRMA's functioning.

Another important strength of CCRMA is good relations with commercial musical technology. Chowning's relation with Yamaha through the FM Synthesis patent is an example. This goes beyond royalties. Chowning visits and helps Yamaha. They are going to lend him some of their equipment for a composition he is undertaking. Sony lends or gives digital sound equipment. Loren Rush and others are in touch with people in this country who are going to produce laser-read "compact disks", and will supply computer music examples for some demonstration disks.

SOME VEXING PROBLEMS

The present music synthesis system, based on the Funely computer and the Samson Box, is necessary to the functioning and survival of CCRMA, for it is what the composers use, with great facility and skill. Yet, it is more a vestige of the past than a wave of the future. I found the text editor very powerful but with so many different commands that it was a little difficult to learn; however, all who have spent some time at CCRMA use it with great facility. I'd say the same about FLA and SAMBOX.

In the CCRMA atmosphere, the only way to go on to something better is to win over the CCRMA people with more powerful or attractive facilities that are easier to use. Part of the "more powerful" would be individual synthesizers as capable or more capable than the Samson Box. Maybe either Mead's chips or Bob Shannon's synthesizer will do this, but when? The chip is a good way off, and Bob's synthesizer might not materialize or might not insinuate itself into CCRMA.

We should note that the powerful modeling modes of synthesis that David Jaffe and Julius Smith have developed, and that might be used for voice synthesis, are quite different from the Samson Box, from Bob Shannon's synthesizer, and from anything that Carver Mead or his people have thought of in connection with the chip.

The present approach to the future seems good to me. It is to make 6 68,000-based terminals (Leland Smith will use 2 for music graphics) and to equip them with inexpensive array processors for analysis and processing. In the CCRMA view, analysis and signal processing are just as essential to computer music as synthesis is.

But, how to get cheap real-time synthesis? I've made some comments above.

I heard Barrie Vercoe and Dick Moore urge John Chowning to get an 1155 (because it is fast and cheap) and a VAX. I wonder. I think that the future of computer music lies with really cheap hardware, such as the 68,000 and special-purpose hardware for synthesis. Chowning can't afford to turn his present facilities off until CCRMA composers abandom them for something better. I don't think that software presently available elsewhere for the 1155 and the VAX would satisfy CCRMA composers without considerable augmentation. Certainly, going to an 1155 and a VAX instead of the Funely would require a lot of software work, including the conversion of many present compositional resources to a different language and a different operating system. And, there would be a considerable hardware and hardware maintenance cost and burden. All this would have to be accomplished in the CCRMA mode of operation -- consensus and persuasion.

I don't know just how CCRMA can best get into the future. Bypassing presently available resources, including the VAX, would help. That requires really cheap and flexible real-time synthesis. Are array processors a viable answer? The CCRMA people seem not to think so. Mead's chips might be. Bob Shannon's synthesizer just might be. Other, existing, real-time synthesizers seem to have their limitations. Some time ago (before my March-April visit) I suggested to Allen Peterson that it would be nice if he thought about the real-time synthesis problem. Allen has produced a remarkably powerful piece of real-time spectral analysis hardware for the SETI program. Well, he now has a graduate student (Bob Shannon) who is interested in real-time synthesis, and Allen might become interested.

Another very real concern for the future is sampling rate, and its standardization. CCRMA pieces come with almost any sampling rate. In the day of analog recording this didn't matter much. Commercial digital recording and the compact disk change things drastically. Japanese firms such as Sony are producing all sorts of very high quality digital sound equipment, and it all works at the same standard sampling rate (44,000 samples per second?) and at no other sampling rate. In supplying computer music examples for demonstration compact disks, CCRMA people face a difficult problem in high-quality conversion from one sampling rate to another. The course of computer music would be much cheaper if CCRMA and others could settle on and stick to the commercial sampling rate, or at least to some rational multiple of it (you can always omit every other sampleif frequencies above 20 KHz are avoided in synthesis). In fact, I wonder if computer music that doesn't conform to commercial digital standards has any future at all.