YLEM
JOURNAL
Artists Using Science and Technology
vol. 25 - no. 6 & 8 Special CD Edition
Double Issue $12.00

CHOWNING
EVANS
GARTLAND - JONES
HIZUME
MATHEWS
SPECTOR
WARWICK

COMPUTERS AND MUSIC

SPECIAL CD EDITION

Max Mathews at Bell Labs

SSN 1057 - 2031

INTERVIEW WITH JOHN CHOWNING

by Loren Means

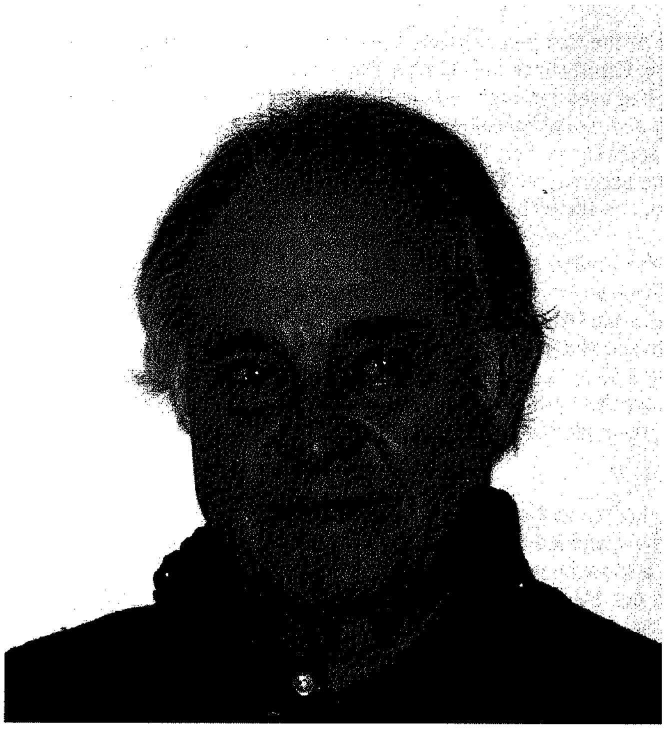

JC: With the advancement of technology, now an ordinary laptop computer has more power than probably all the computers I've ever used in the lab, plus special-purpose devices like the wonderful Systems Concepts synthesizer that was designed by Pete Samson. All of these together, maybe ten times over, exist in a little laptop. So the capacity to do fast computation and processing for music composition now is kind of unlimited. The capacity surpasses my capabilities of conception. There's more power than I can imagine using. So the whole world has changed from specialized devices, hardware synthesizers, to software. The issues now have to do with control. Kinds of controllers, like Max Mathews' Radio Baton, which is an electronic device which allows one to wave an electronic baton over a sensing surface, and this sensing surface picks up the transmitter signals from the Baton. It calculates distance from the surface and angle, so it's a three-dimensional space. That's a kind of controller that is maybe typical of what's being done today as far as technology in terms of controlling synthesis in real time. Don Buchla, who is one of our local geniuses, is also very active in this domain. His couple of machines, Thunder and Lightning, which are controllers using similar ideas...

LM: It sounds like a theremin.

JC: That's right, the theremin was the original controller. Theremin is an interesting device, but it's very limited. It's like trying to play a violin with one finger on one string. You can play certain kinds of melodic lines, but anything that takes articulation–staccato, and fast passages that have spaced intervals like thirds, fifths, et. cetera, it's very difficult. You really can't do it.

So controllers today are of two sorts, I'd say. One is the sort that Max Mathews has developed, the Radio Baton, which is at the level of the conductor. The other level is more detailed. For example, the glove, a special glove that's wired at the centers of the joints and fingertips, and so any move of the fingers, thumb, or curvature of the hand is sensed. One can control some degree of detail, certain kinds of low-level musical activity–pitches, dynamics, it can be anything that can be resolved to spatial position.

So what I've been working on is sort of in between. I have just finished a piece for solo soprano and computer synthesis, where the soprano has a microphone proximate to her mouth, and the signal that's sung is sensed by the computer. Then that is used to animate a kind of accompaniment in computer synthesis, or processing, reverberation, spatial positioning of the voice, or all of those things. That becomes very interesting, because the soprano now is in partial control of the accompaniment, and not by any means other than that with which she is intimately familiar, of course, her own voice. That's the controller. It's not the first time it's been done. Philippe Manoury in France did a piece for soprano and synthesis/processing several years ago. But it is an idea of using computers that brings it pretty close to the kinds of musical expression that we're used to experiencing with human performance. The human voice, of course, is the instrument of instruments, the first of all, and the one which probably has the greatest degree of ultimate control over the detail.

LM: The controllers you're talking about seem to me to be working in space, whereas historically in a lot of cases, the controller of choice has been a keyboard. Didn't the first Buchla box play single lines?

JC: All the first early synthesizers were monophonic, and Buchla's, as I remember, had a ribbon, so he could create glissandi with that. It was ingenious. Don Buchla is quite an extraordinary person. We don't compare easily people like Buchla and Moog, but if I were to venture a comparison for local consumption, I'd say Buchla is somebody who has really kept his nose ahead of everything regarding technology.

LM: I remember hearing Don jamming on synthesizers with Allen Strange at 1750 Arch Street in Berkeley in the Seventies. But a real innovation you brought about in the Seventies was with moving sounds around in space. Was that your invention?

JC: It had been touched upon. I really feel an intellectual and artistic obligation to make attribution, because it's an idea that Varese had in his mind when he wrote the work for many, many, many channels at the Phillips Pavilion at the Brussels World's Fair.

LM: 1958. "Poeme Electronique." But it seems to me there are pieces by Corelli and the like where they'll put brass instruments around the stage and pass the sounds around.

JC: Gabrielli, in Venice. Those were poly-choric pieces. So the idea of using space is not a new one. In fact, I got the idea of moving sounds in space from Stockhausen, who used the idea of moving sounds to some extent in one of his early works. In "Kontacte," a piece for piano and percussion and tape, he used 4 microphones arranged in a square in the center of which a loudspeaker rotated. The microphones then sent the resultant signals to a 4-channel tape recorder. So he was able to create some motion in space.

LM: And that was in the Fifties.

JC: 1959. So when I started working with computers in '64, that idea had already been in the musical mind, in a larger sense, of electronic music. What I did was generalize it and more or less make a physical model of moving sound sources including Doppler shift and arbitrary positions in both angle and distance. So it was a big step, and it's still kind of a basic idea, used in theater sound systems and what not. It was an exciting moment in the work that I did at Stanford, because in doing that, also, I learned how to program, and realized that programming was more than simply getting a task done. It was a whole universe of possibilities that I had entered upon. I realized that this computer was no longer, in my mind, just the alternative to an analog electronic music studio, but a wholly different thing, primarily because of the power of programming languages.

LM: When you say "Doppler," you're talking about the fact that a sound appears to get louder as it moves past you...

JC: Doppler shift is really a frequency shift, so the pitch changes. It also gets louder as it comes toward you and moves away. But it's something which we're extraordinarily sensitive to in our world, because we depend upon that for our life. If we hear a bus coming, or a horn blowing, and the pitch is increasing, Doppler shift tells us that this thing is not only coming, it's coming fast, and it's close. It's a kind of cue that's had a quick evolutionary growth in human perception. After all, before the Industrial Revolution and trains, Doppler shift was not an important part of our acoustic experience.

LM: I went to a demonstration you did at Stanford in the Seventies, and people would duck as sound moved over their heads. You also could change the timbre, and make one instrument sound change into another instrument sound in the process of moving.

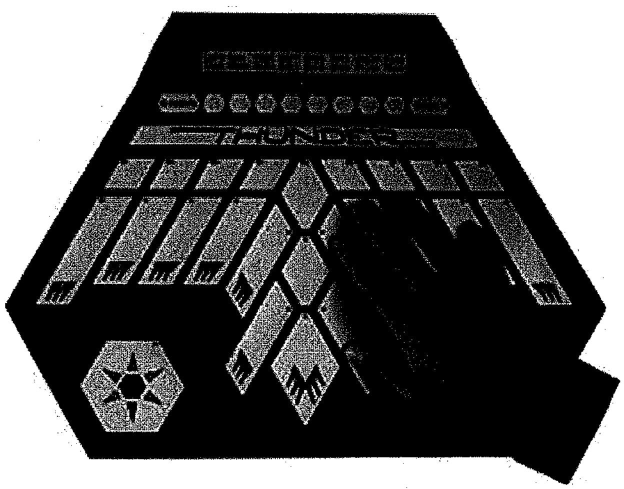

The Thunder Interface

JC: And that was using Frequency Modulation synthesis, which yielded a lot of power over timbre with a few number of controls. That developed right after I finished the spatial simulations.

LM: How were you generating the sounds in space before FM synthesis? Were you using sine waves and such?

JC: I was using any sort of signal, square waves, standard waves. Then in '67 I discovered the idea of modulating a sinusoid's pitch rapidly by another sinusoid. That is a special case of Frequency Modulation that is used in radio broadcasting. The same formulas apply, only in this case we don't demodulate. It's all in the audio band. It's kind of counter-intuitive, but allowed me to produce a large number of timbres with a very small computational engine, a couple of oscillators and a couple of envelopes. So with that kind of control, it was very simple to change timbre. As sounds passed through the space or over the space, it was possible to make these metamorphoses of timbre ... like Escher does in his engravings.

LM: How did you actually find that FM worked?

JC: FM was an ear discovery, it was not an engineering discovery or a scientific discovery. I was experimenting ... toying with parameters. I was creating a vibrato over a pure tone, a sinusoid, and I kept increasing the vibrato rate and its depth. Vibrato is the effect produced by a violinist who rotates the finger...

LM: ...The pitch moves up and down around a fundamental pitch...

JC: ...And usually it's within less than a semitone, plus or minus the average pitch.

LM: Usually you'll make it wider as the pitch is sustained, to add interest to the sound.

JC: So violinists use, it singers use it, most...

LM: ...Trombonists use it, and you can see the slide movement.

JC: Oboes, all the instruments use it except French horn.

LM: The bebop revolution in jazz wouldn't allow vibrato. Miles Davis said, "Don't use no vibrato. When you get old you'll shake anyway."

JC: Anyway, this is what I was doing. I was experimenting with just a sinusoid and kept increasing the vibrato rate, so all of a sudden it didn't sound like listening to a change in pitch in time, but rather I began to hear timbral differences, sound quality differences. So the vibrato became very, very fast, hundreds of times per second, and very, very deep, as if the violinist had a different fingerboard, and the finger was whipping up and down at very high rates and very great distances. That would be sort of a physical metaphor for this. What I was hearing was timbral changes. And this was done with only three controls: the frequency around which the vibrato was occurring, the rate at which it was occurring, and the amount, what we call the depth. So with these three controls and two oscillators, I did a number of experiments and got percussive tones, brass-like tones, and woodwind-like tones. That's when I realized that there was enormous power here that was predictable. After I had done a set of experiments, I then asked an engineer, George Gucker, for some help in understanding what physically was going on. So we pulled out the engineering text by Turman, a Stanford professor, and looked at the equation in detail. It certainly predicted what my ear was telling me was happening. So that was the beginning of FM.

LM: I talk to artists and roboticists who are using principles of emergence, and they keep getting back to this delightful element of surprise.

JC: That's right. It wasn't an invention as much as it was a discovery. It was always there. And ever since, it's been full of surprises for everybody who has ever used FM. The DX-7, I think, had more surprises per user than any other instrument that's ever been on the market.

LM: That's the Yamaha synthesizer.

JC: That was their first popular implementation of FM synthesis. They had another instrument or two before, but they were not so general, and one was larger and more expensive, the GS1. Quincy Jones still has one. It was a beautiful, piano-like instrument, and it had a wonderful keyboard feel. But the DX-7 was the first one that was within the price range of synthesizer buyers. It had some attributes that made it absolutely distinctive, that Yamaha never anticipated would be so widely accepted. One of the attributes was the fact that it had velocity sensitivity that could be coupled to timbre, tone quality. That, I think, is one of its most distinctive features, because that's a feature on which really good musicians depend to make the performance expressive, like a pianist who spends hours and hours learning to control the rate at which that key is depressed in relationship to the context of the surrounding keys. Most synthesizers, and even the B-3 organ, still a popular musical instrument, did not have that attribute. It doesn't matter how hard you hit that key, it's always going to be the same loudness, depending upon the pedal. But the pedal doesn't allow you to control that note-to-note velocity which relates to loudness and timbre, an important detail, whereas the DX-7 allowed you to do that.

LM: That instrument came into being because of your work at Stanford.

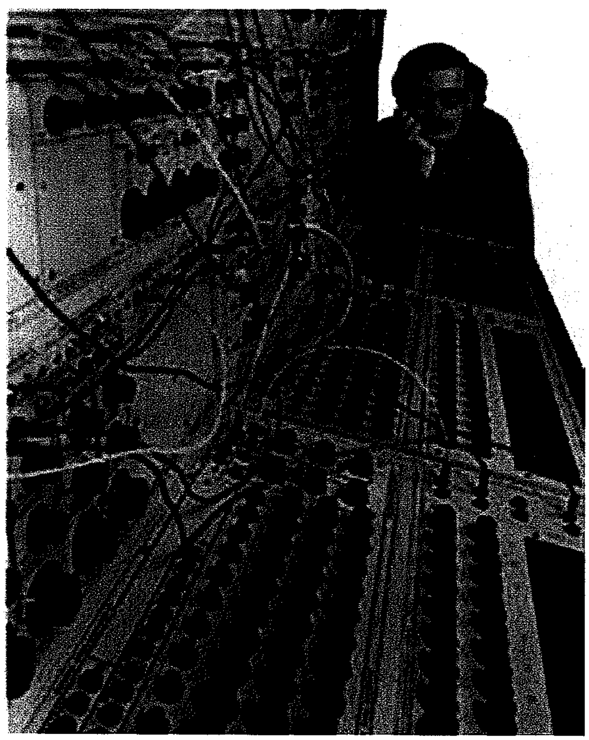

Donald Buchla

JC: I discovered the FM synthesis in '67, and then did some work with it in the intervening years, but it wasn't really until '71 that I rethought an article that Max Mathews and Jean-Claude Risset had published in '65 about an analysis of the trumpet tone that they had done at Bell Labs. So in '71 I thought about that article and a discussion I'd had with Jean-Claude, who had done most of the research on that particular aspect of the work. I realized that there was a way to use FM that was extremely powerful, coupling this idea of velocity or force or breath pressure or effort to timbre and loudness or intensity. That was a key insight. At that point, Stanford contacted a number of instrument companies to see if they were interested in producing instruments using this technology. None of them understood, because they had no idea about the digital domain at that time, except for Yamaha. They sent an engineer to Stanford, and he understood immediately what the issue was, and the problems it solved. They were already interested in the process of developing digital synthesis based upon other technology. So they took an option on a license and finally signed an agreement with the Office of Technology Licensing at Stanford, to whom I had assigned the rights.

LM: You had been fired by Stanford when this happened?

JC: Yeah. About 1973 I lost my job because I didn't get promoted from Assistant Professor to Associate Professor. So I was let go. That's the break point in the ladder of academia. If you make it from Assistant Professor to Associate Professor and you get tenure, you have a job. But I didn't. It was a curious coincidence of things. In that very year, Yamaha had taken their option. In the same year, Pierre Boulez had contacted me through Jean-Claude Risset, and asked me to help in founding his studio in Paris, IRCAM [Institut de Recherche et Coordination Acoustique/Musique]. That was his big project in the Seventies. So I lost my job, and then at the same time, all of a sudden, the work that I and people like Andy Moorer, John Gray, Loren Rush, and some other people had been doing at Stanford became of interest to one of the most renowned musicians of our day.

LM: You had already formed CCRMA [Center for Computer Research in Music and Acoustics] at this time..

JC: Yeah, in '75, actually, although the work had been going on since '64. I was in Berlin for a year, then came back as a Research Associate, because I had no other place to work. There was no other computer facility which was equivalent. And of course we had invested all this energy in the software that was there. So we kept working, and got some grants from the National Science Foundation and from the National Endowment for the Arts.

LM: So you were working at Stanford even when you didn't have a job.

JC: Well, I was a Research Associate, so it was dependent upon outside money.

LM: And then they hired you back as a professor.

JC: Then the UC system offered me a full professorship, so Stanford did, as well.

LM: And this is unprecedented, for a university to hire somebody back that they had let go?

JC: It is truly rare. I think a dean once told me that it was the first case they can remember.

LM: It's like Stanford shot themselves in the foot and then noticed they were bleeding.

JC: It's not surprising. At that time computer music was weird. To most humanists, the idea of not just using computers to analyze music, but actually to make the sound, seemed like a pretty large jump for a university to make. Listening to some of those early sounds that were made back in the Sixties, they were generally pretty raw. Had I not had any deeper understanding, I probably would have thought twice before thinking that this idea of computer music has a future. But in '64, '65, I realized that with this program, with some modest ability at programming, I could create this magical sound space with just programming, without having to know how to solder, without knowing anything about electronics, without any intervening engineer. I realized that there was a power here that had an enormous future, that we had just begun to touch upon.

LM: You've composed using the Golden Mean.

JC: That's interested me, and not from a typical metaphysical motivation. My interest in that is because after I finished "Turenas," which was the piece that exploited this moving sound source (finished in 1972) I did some experiments with inharmonic spectra. In "Turenas" I used the square root of two ratios and other common irrational numbers, and I looked for others. The Golden Mean surfaced at 1.618. It's that ratio, the Golden Proportion or the Golden Mean, that reaches back to antiquity. It's used in Greek architecture, and is found in nature. My interest began because it had a compelling kind of timbre in its simple FM synthesis formation used to create inharmonic spectra. So I started thinking about inharmonic spectra in relationship to a complementary tuning system as an analogy to the harmonic series and our common tempered system tuning system.


Donald Buchla

LM: What do you mean by "inharmonic spectra?"

JC: The non-musical acoustic world is basically chaotic and inharmonic, the relationships of frequency partials which are not in simple integers. But this chaotic world, like gongs, cymbals, wind, noises of nature, can be orderly if we have computers. With computers we can organize spectra that are inharmonic to be quite as orderly as the harmonic spectra are in nature. So I took this idea and developed it in a piece, "Stria," which is based upon the ratio of the Golden Mean, to create the actual quality of the sounds themselves. Then I structured the pitch space to be based upon a scale that is complementary to the spectra in the sense that spectral frequency components and scale step frequencies are the same or close, as octaves, fifths, 3rds, etc. are found in the harmonic spectra of musical instruments. Then the proportions of the piece, the various formal aspects of the piece, were also related to the Golden Mean. That was probably the most exciting compositional project I've ever done, because as you were mentioning earlier, so many of us find these surprises when we confront technology with artistic ideas. That piece surprised me at every instant in its composition.

I built a pretty complicated computer program to produce this piece. I continually fine-tuned the program at every compositional decision. Listening, making an adjustment, adding features, changing features, until I had this magical box which I could give directions to, which would then produce all these sounds based upon decisions I had made earlier. I knew what was in the box, because I had built everything that was inside it, but the moment when I actually used it, I didn't have to think about all this detail. I would just give the instruction to begin at such and such a time in a certain frequency space, and make it so long and with a certain kind of density, and it would just do it. Though nothing in detail was a surprise, the whole thing was kind of a huge surprise every time I implemented this procedure. So it was extremely exciting.

It was almost an organic relationship with this computer which represents tens of thousands of man-years of thought about thought ... the idea of logic and all the stuff that's inside a computer ... all that was interacting with me in this musically organic way. It was like I was being hugged by Newton and all these great people in computer science. It was like they were all part of what I was doing. It was an extremely satisfying and amazing experience.

"Stria," I thought about beginning in '73 when I was in Berlin, but I had no computer, so I did all of the calculations and planning by hand, what I would do when I could realize it. Then when I got the position as Research Associate in '75 at Stanford, I did some more thinking about it. I didn't really start programming it until the summer of '77, when Luciano Berio asked me to do a piece for the opening concert of IRCAM in Paris at the Centre Pompidou.

LM: How did you learn to program?

JC: I read Max Mathews' article that was published in *Scientific American* in November '63. It was passed to me by someone in the Stanford orchestra who knew I was interested in contemporary music and maybe I had said something to her about electronic music. Her husband was a doctor at the medical school. I didn't subscribe to *Scientific American*. I never had seen the magazine. I put the article in my pocket, and then maybe a couple of months later I found it and read it. This was Max describing the use of computers to generate sound. I had heard electronic music at the Domain Musicale series in Paris during the previous three years when I'd studied with Nadia Boulanger. So I thought, "Maybe this is a way to compose electronic music." I thought we might use computers at Stanford, as we had no electronic music studio.

So I looked in the course catalog, and there was a course for non-engineers in Algol. So I took the course. It was amazingly easy for me, because they taught this course not programming engineering problems, which was the way most every programming course was taught in those days, but rather they looked for ways to engage this population of people--it was a small group of us--in terms of solving problems that we posed. So I thought, "I'll generate a whole bunch of twelve-tone rows." Although I wasn't so much interested in that kind of music, it was a tractable problem. So I solved the problem, and in doing so I learned how to program, at least the basics of programming. Then when I really wanted to do something, the following summer, I was prepared. I got from Max Mathews at Bell Labs a box of cards with the basic program for music synthesis.

LM: I like that idea that instead of giving you the program on a disk, he gave you a box of cards. I started programming using punched cards, too. Were you influenced by Lejaren Hiller and the "Illiac Suite?" That was the first computer composition I ever heard of.

JC: Hiller not much. He did different things. He was actually preceded by Barbaud in Paris, and Xenakis, both of whom used computer language for higher-level music programming at the score level rather than sound synthesis. I think the real richness of this medium began with sound synthesis, and then developed into compositional algorithms which were related to synthesis. It was there that the integral nature of music, and how all these levels of control interrelate, developed out of Max Mathews' program having to do with sound synthesis. And one of his early examples is stunning in its implications, far more stunning, I think, than what Hiller did. Although the sound and the musical representation may have been not so interesting, the idea was extraordinary, which was a metamorphosis between two tunes, like some common-known folk tune and another common-known tune. He created this metamorphosis at the musical, structural level, which is really compelling.

Lissajous Curves

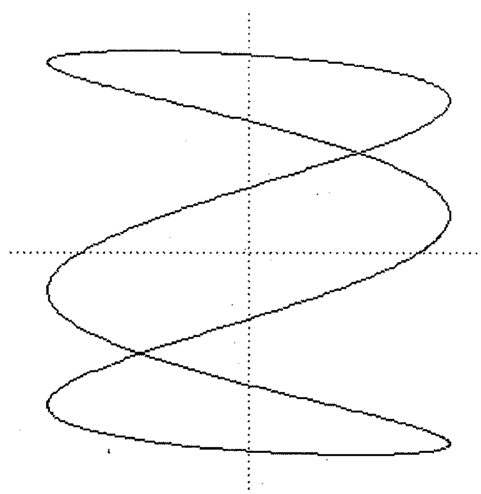

Jim Tenney was really important. He's now at Cal Arts. Jim Tenney was given by Max and John Pierce the right to do some work at Bell Labs in those days, as was Lillian Schwartz. John Pierce created this umbrella of protection for artist's work, including Max himself, because music was not his task at all. Here was this institution, which at that time was probably the greatest concentration of brain power on earth, engineering/scientific brain power. These artists were allowed to work there, all because of Pierce and Mathews. Jim Tenney was one of these people. He is a wonderfully imaginative person, and he did a piece for computer which was based upon random processes. He wrote an article for the *Yale Journal of Music Theory* describing Max's program, which was very important to my initial understanding. It was published in '63 or '64.

LM: Is there a contemporary equivalent of the Yamaha machine that is taking computer music to another level?

JC: I don't know what's really going on in the commercial synthesizer world. I never really tracked that very carefully, because it wasn't very close to me. It was a consequence rather than an interest, although I certainly spent plenty of hours in Japan helping them develop it. It's not really my world. I think that in pop music, mostly it is samplers now. Mostly what people want in popular music or commercial music, for films and TV and commercials and advertising, is instrumental sounds and natural sounds. At the time the DX-7 came out, sampling was not viable, because memory was so expensive. So most of the interest in the DX-7 was the fact that one could get somewhat realistic approximations to acoustic instruments. But then memory became so inexpensive that they just recorded sounds digitally and then played them from a keyboard or computer.

LM: My question is, why synthesize instruments that are already there instead of creating sounds that were never there before?

JC: That's because you're thinking in a different domain than somebody who is asked to give us violin sounds for a TV show, for example. It's a wholly different thing. That's why it's not very interesting to me. My whole interest is in what we can do that touches the imagination in a way that it has never been touched before. I was seeking a way of understanding some of these intricacies that existed in natural sounds. These sounds seemed to have components that are important to our perception, that are independent of whether something is natural or unnatural. Does the sound have a dynamic quality that is attractive to the perceptual system? Simulation of the vocal tract became very interesting because it's so complicated. I did some FM synthesized singing voice experiments to try to understand how the voice worked, and then that resulted in a piece, "Phone," (1981), which is based upon the metamorphoses of sounds that become voice-like just momentarily, and then transform continuously to some other timbre. So to that extent, our interest in simulating natural sounds was important, but that wasn't the end. The point was to experiment to understand, and then to do something creative with that knowledge.

LM: I'm wondering about the controversial concept of the computer doing the composing.

JC: For me it's not a controversy. There's a whole continuum of questions that relate to the degree of computer use, the degree to which it is composing. If a human writes the program, and then you put your nickel in and get your composition out, is the computer composing, or is the programmer directing the computer to compose? It's not an obvious answer to the questions regarding computer composition. In the case of "Turenas", I used Leland Smith's composing program but I wrote sub-programs to control the spatial aspects of the piece. In both "Phone" and "Stria," I wrote the entire program to automate certain processes to different degrees, but those processes were authored by me. At the moment of composition, however, those processes were independent of my specific directions. So what is an automatic composition?