

TECH[NOCULTURE

The artistic gesture: Research at CCRMA

Episode 40

Full transcript

Guest: John Chowning [John]

Host: Federica Bressan [Federica]

[Federica]: Welcome to a new episode of Technoculture. I am Federica Bressan, and today my guest is John Chowning, composer, researcher, and founding director of CCRMA, the Center for Computer Research in Music and Acoustics at Stanford University. Welcome to Technoculture, John.

[John]: Thank you.

[Federica]: Thank you so much for having me at CCRMA today. I'm thrilled to be here because music and electronic music are closer to my experience than other topics that I've had on the show. And so to prepare for this interview, I have actually reached out to friends and colleagues in this field. And they have submitted to me some questions that they would like to ask - to someone who has made history like yourself - I would like to begin, though, by asking you about the very centre CCRMA. Now, the name as well as the centre, is a couple of decades old. The name that you chose then, Centre for Computer Research in Music and Acoustics, do you think that it still reflects what is being done today?

[John]: Yeah, I think it does. We do work with computers, as you can see. We do a lot of work with acoustics, which can be in range from physical acoustics to room reverberation to psychoacoustics. So, I think these are broad enough names that characterize fairly what we do here. The actual work began in 1964, when I was a graduate student. I read the article by Max Mathews in Science magazine, and I had no experience in electronic music or electronics. I had no experience in computers. I was a musician, trained and lived and loved music, had come from Paris to study with Nadia Boulanger in 1959 until 1962. And while there, I had

heard music by European composers, Boulez, Stockhausen, Pousseur, Berio... and the idea of composing for loudspeakers was very interesting to me, but I had no means, no qualifications. So, when I came to Stanford as a graduate student composer I thought, there's no opportunity to do that here. And then I read Max's article in my second year of graduate work - someone gave it to me, I read it and something snapped in my head, which was, if I could learn to program a computer, I can, in theory, generate anything that my mind can conceive as long as the source is a loudspeaker. Because the only equipment that's required in Max's initial article was a computer, digital-analog converter and loudspeaker. It's exactly what we have in every device. It's still in our iPhones or iPads and this and that: we have DACs, computers and loudspeakers. All of which have increased in quality and speed. But the idea of this abstract notion that, without having to learn how to patch or deal with electronic instruments, I could go directly from musical thought through programming to music making... that was the attraction. So, I could just say that in 1964 I began work in spacialization, having been struck by Stockhausen's Kontakte, for example, in four channels, and that was my goal. Then other graduate students became interested in what I was doing, joined in, I finished my degree and began teaching. And we had a nice team by 1972: myself, Lauren Rush, another composer Andy Moorer, a computer scientist, John Grey, a psycho-acoustician, or perceptual sciences and psychology. And so, in 1974 we formed a Centre for Computer Research in Music and Acoustics: CCRMA.

[Federica]: And you're still active here. What makes you curious today?

[John]: Oh, well, I still compose.... slowly, but that's my way of working. I like the time, the unconstrained, well, composing without pressure. I've only had one commission in my life. Others I've turned away. So, I work at making music, and it always involves research, usually in some perceptual domain. At the moment, we've just submitted a project to the scientists at Chauvet in the Chauvet Caves in Southern France, doing a project with them, which is quite compelling and quite unusual.

[Federica]: Do you want to talk about it a little bit?

[John]: So, Chauvet Caves, the wall paintings are 32,000 and greater years old, and some of the most beautiful of all the wall paintings that have been found anywhere on this earth - and some of the oldest. And it turns out that these caves were discovered in 1964, and the French prehistoric scientists and researchers said, "We can't let what happened to Moscow happen to Chauvet." So they closed it to the public. Now, there's been a theory since the late 1980s that some of the wall paintings in other caves like the skull, or others, quite a few in Spain and elsewhere... their position and subject was determined by the acoustic response at a certain point in the cave. In other words, some of the famous images of Chauvet, for example, are the overlaid horses, horses heads that seem to be in motion. And the question is, what if you click

two rocks together at that point in front of those images, might you have heard an acoustic response of this echoes? Crump, crump, like horses moving on a plain. What an interesting idea, right? But there's no way really to test that. Because these caves - since the beginning of the end of the Ice Age, from 19,000 years ago, to about 11,000 years ago - the de-glaciation has caused stalactites and stalagmites to form and other calcite deposits that have kind of filled the caves now with all these pieces of rock. So, you can't really do experiments in the cave because they've been, in a sense, "polluted" by the natural consequence of de-glaciation. So, I had this idea a couple years ago of teaming up with geologists, and have them make estimations of the amount of growth of these stalactites, since 29,000 years ago, when the cave was absolutely closed because of a natural rock fall. And using contemporary image processing, subtracting out all the accretions from the image then we, as scientists or engineers and musicians who worked with simulation of reverberation, as we did with the Hagia Sophia and others in Chavín de Huántar in Peru, recreate the acoustics of the world at that time, so that we hear what the creators of these images heard. So, I contacted this geologist whom I had found in the literature, and asked him if that would be interesting. And we just last week received a very positive response for us to come and concretize this idea in France at the end of March. So

[Federica]: So you're going to take a 3D image of the actual space?

[John]: Well, they do photograph-metric models of the caves as they are now, and from those they create a visual model, maybe a 360 hemispherical projection. And from that and the data that they have to make their original models, we can reconstruct the acoustics as the plan. Well, there's a big research that has to be done because of all the refinements of making it happen, and happen in a way that's convincing and faithful to the incredibly dedicated conditions that these French scientists at Chauvet maintained. The integrity is so great. There's a high standard that has to be met, but it's doable. I mean, the technology is here, we know how to do it, just we've never put all these disciplines together to do such a grand project. But remember that these paintings, the first year they opened the replica of the underground caves with a very careful model of the internal - it's a kind of a museum - but you walk in and you're in the cave with all these paintings. They had 600,000 visitors to this replica, so it's a big thing, and the power of these images has to do with every human being on Earth. It's independent, in particular cultures. I mean, there were local cultures, of course, the [?] people, but it's something that touches everyone, because it's so expressive, you know, this gesture. Okay, so that's one project.

[Federica]: This is early 2020 that we're talking now. Do you have timelines, you know, expect to have something to show with these projects or they're also no pressure?

[John]: Well, we do have a timeline, you know, we can only go into the caves about four weeks a year. Because the natural CO2 level is sufficiently low that during those weeks they

allow people to go in (researchers only) to reduce the amount of CO2 from exhalation, and so whether we'll get inside the cave at the end of March, we don't know yet, but the plan is if we do, the next time it's open in a year from now, we would go in and make these careful acoustic measurements. The sine sweeps and impulse responses using the latest technology, multi capsule microphones, so that we get the directional information. And at any point in the cave the acoustic response includes echoes that traveled through a gallery into another chamber and come back, and so that directional information is really important to try to faithfully reproduce what they heard.

[Federica]: I heard of some studies that involve the neurosciences and geologists, but never for a natural space. Always for like a temple, something man built or modified, not a natural cave.

[John]: Well, we did that for the Hagia Sophia, a group and team here and working with an art historian, Bissera V. Pentcheva, she's on the Stanford faculty, but Byzantine Iconography is part of what her field is, together with Jonathan Abel. But effectively, they measured the acoustics of the Hagia Sophia, in which there can no longer be performances since the early twenties, when it became a museum. And we transported the acoustics of the Hagia Sofia to the concert hall here. And they wired up the group of singers who do period music, music that was composed for the Hagia Sophia and processed their voice with a dome of speakers set up in the concert hall. Reproduce, the music's done for the Hagia Sophia in the Hagia Sofia, if you close your eyes.

[Federica]: If you close your eyes... Well, that is definitely impressive. I would like you to give a comment on all the possibilities that are there today that weren't there decades ago and on your pioneering years in the light of what followed. For example, you mentioned that a big shift happened from offline to real time...

[John]: So, that was good. But there's some consequences that result from having access to real time. And one of the consequences is you don't have to know so much, because you just pushed buttons and turn knobs and you turn things, make noises until your ear says, okay, I like that. The value that in the early days of not having real time is we had to know a lot more about what we would expect from trying to make a two second example, say, discovery of FM - discovery, not invention, because it's a gift of nature, FM and Bessel functions and all that go into it, - but to make a small test two seconds of sound, I would have to wait sometimes 10 minutes. If there were lots of users on the timeshare machine, maybe an hour, to do two seconds of music. But all the time I'm thinking. And I'm trying to understand, what will I hear? So, we learned about psycho-acoustics and perception, because the cost of computing was so great we had to inform ourselves to enrich the possible end product, the result. So, there was a lot of work in perception With real time, people just work. There's no longer the requirement

for understanding to the extent that we had to. Well, it makes real time music very rich, of course, and there's some people who are not tempted always just to work in real time. But it was an important contribution when, for example, the DX7, the first all digital synthesizer, became available. And it connected up to a small computer, and for a couple \$1000 you could have a pretty powerful little workstation. And Jean-Claude Risset, who is my close colleague all these years, we both started work in 1964, he at Bell Labs, and I here, he started work with Max Mathews - we started working on a book, "Exploring psycho-acoustics perception using the DX7." I mean, there's a lot you can do with it. And we had some really exciting examples showing various things, like missing fundamentals and how you build attacks, and how you put which programmers and of the DX7 would call "stuff" into the sound, meaning noise at the right point... All these things that the ears tend so much to as details that are important to making a sound live. But what happened was Yamaha then had another better product and so they dropped that DX7, they didn't want to promote it, so we had no support then to write the book, because they want output that would encourage use of their newer machines. That's how they made their money, of course, so it's natural. But anyway, there's no doubt that the DX7 with a computer democratized music. Until then, we had to have large, many hundreds of \$1000s systems to be able to do this and all of a sudden, for a little bit of money, people could do wonderful things.

[Federica]: In a time where computers offered little or no support for making sound or music, what drew you to computers instead of, for example, what was more available then, for example, voltage control synthesizers?

[John]: Well, they were just available. I mean, when I started in 1964 you could get a machine from UCLA 50 miles away from here, probably for about \$10,000, maybe? But I was able to get hold of a million dollar computer for nothing from the University, because it was institutional and I was a graduate student. So, that was one thing, we didn't have the option, didn't have the money. We had no support. But there was something more important I realized quickly on, after my first project was enhancing the idea that Stockhausen had used in Kontakte with the rotary, a loudspeaker and four microphones. And I generalize that, using four loudspeakers so that I could, one could draw any path in a two dimensional space: complicated Lissajou patterns and whatever one chose. And those were stunning examples at that time, because we had distance queue based upon director reverberation signal ratios, Doppler shift. They were very compelling. When I finished that project about 1968, I realized Stockhausen had to have a team of engineers and assistants to be able to do what he did. I did what I did just because I learned how to program and got some engineers to build a four channel DAC. And without... I mean with lots of help from scientists and engineers and the Artificial Intelligence Laboratory - because I asked questions, because I had no background, - but basically by myself, without any assistant, I was able to do this and realize that the power of computers and music is that the structure of music and the sound itself became tightly

linked. So in a piece like *Stria*, which I completed 1977, done with code in a language called SAIL (Stanford Artificial Intelligence Language), which is no longer used. But it was an ALGO-like language, I created the structure that's so tightly linked to the actual sound that it's all one thing. And that was new. I mean, the idea that we could do that, we realized in the early 1960s, that this was an aspect of computers, because of the fact that programming languages represent what tens of thousands of people years thought about thought and so all of that accrues to one. It's a very special experience.

[Federica]: Are you surprised at how software has taken over the entire music production process? And do you marvel at the sheer number of tools that have emerged since the Sixties?

[John]: That's amazing. Yeah. When I started a piece for solo soprano and computer, I decided I'd make an interactive piece. I sat down and started to use Max MSP, which I had not used before. But I thought, okay, I was told that's a good programming platform for interaction. Okay, first, I'm gonna have to build a complicated reverberator, because that's what I had to do for my early pieces back when we did software synthesis. And so, I looked and, there was a beautiful reverberator already done by someone. And it was there, it was like days of work that I could just grab it, tune it to my own ear, what I wanted, and that was amazing to me. So, I wrote a little note to this programmer, just told him how meaningful it was to me that in this interim all this was done. I said, I'll give you a signed copy of my FM paper just as a... he didn't even know about FM or about me and FM. He's some guy who picked up the program, learned to use it, did this amazing thing, this fully fleshed out four channel reverberator system. That was amazing. So, I realized that this community is one of the things that has resulted from this incredible growth of software synthesis, and people talk to one another and communicate with one another. Late at night, you could be working and talk to somebody in Bulgaria, who knows where, who has the answer to your question.

[Federica]: Yes. Of course, all the innovation that we have seen since the sixties also means that when you gain something new, you lose something. So, the programming languages you were using are no longer in use, and you have to keep learning. So, as a composer, have you chased the last things for a while? Or have you chosen what works for you and stuck to these things? Because the chase can take a lot of time to learn always the new thing, right, so as a composer, how do you approach this?

[John]: Programming languages are sufficiently similar that the basic principles carry over. So, if you're learning C, learning Python or some other, I think that's not an issue any longer. But my own music is driven by some perceptual quirk or something of interest in the perceptual domain. So, for example, right now, I'm working on a keyboard piece based upon the complimentary relationship between tuning and spectrum. So I found this example, it was pointed out to me some years ago, an example demonstrating some perceptual phenomenon... critical

band theory, I think that was it... so, the example is a Bach chorale that we hear in four states. A Bach chorale, done with a synthesizer, where the nine harmonics, each of which is controlled by a separate oscillator, and we hear it with this synth sound Bach chorale, nine harmonics, sounds perfectly dull. Then we hear it again, where the example is they stretch the tuning system by 10%

[Federica]:. So, rather than two to the n over 12, it's 2.1 to the n over 12. But the harmonics maintain their natural structure. And of course, it sounds terrible, it sounds out of tune. Then they do the reverse. This is done in 1989, I think. And then they stretch the harmonics by 10%

[Federica]:. So, the second harmonic is closer to a whole step - not quite, but then the tuning is the same, the common practice tuning, and you hear the Bach chorale and again it sounds out of tune. Insufferably bad. But the surprise is when you stretch them both, both by 10%

[Federica]:. It sounds great! More interesting, in a way, than the original. So, I thought, this is amazing. At some point that will break, you can't keep it. So, I put together a program in Max MSP so I could test stretch and find all the combinations. Not wanting to do a Bach chorale, of course, but using that as the reference and create music based upon this play between complimentary relationship between tuning and spectra. So, it's something that's not been exploited. No one's doing it.

[Federica]: You are!

[John]: Yeah, I am. I talk about it to everyone. And I give talks, and I have this example, it shows how this works and what the math is. So, maybe there are people who are doing it, too. But I haven't heard... But, but it's exciting. So, the idea is, music based upon something that's just intrinsically interesting, the fact that our auditory system can make sense out of something that cannot be done in nature. You cannot do that with a natural cord or wind instrument. I mean, the physics binds those harmonics to the physical source, and it can't be broken. But with computers we could do, and it sounds like a Bach chorale that's better than the original, straight for the synth, perfect for...

[Federica]: The last question, before I let you go. On artificial intelligence: what do you think of algorithms and the power of machines to be musically creative? Are you following the last experimentations of Google and Spotify? That's an active field. Do you think it's interesting or not really, in that sense, artificial intelligence for musical creativity?

[John]: Well, I think it's interesting. Of course. I mean, the idea of encoding creativity is an amazing idea. I mean, actually, I don't care. I love doing the artistic gesture. So, when I look at these wall paintings at Chauvet, we don't know what they were thinking, it certainly had to do with rituals, because they didn't live in these caves. They went into the caves, made

them, then left. They certainly had something to do with belief systems and ritual. And, these pre-historians and anthropologists and art historians in France study this, trying to decipher what might they think. But there's one thing we do know, that we know a lot about, and that's what we do as artists. At the moment that creator put charcoal on the rock, he or she is thinking of some internal aesthetic that has to meet those criteria. That artistic gesture is what we all do, when we make music, whether you're writing notes, how I'm going to make the violin sound, or when I wrote a program like for Stria, which was an algorithmic program, halfway to artificial intelligence, say, but a little bit like that... writing the program for me was the same feeling that this creator of this horse's head had. It's the same thing. You don't care about the belief system at that instant in time. You care what it is that he wants to see, or I want to see, what is it do I want to hear, and that artistic gesture connects all of us together. The arts, I think. So, artificial intelligence tries to capture some of that and code it. I did that in Stria, you know, it's like a recursive procedure which started out with about 20 lines of code. By the time I finished the piece, it was about 200 lines of code or more, because every time I make a gesture, get back information both from the acoustic result and also the programming language itself, then you make a change or an addition, but it's always backward compatible, because the sounds that I first made, I want to keep... but so, it was just this big amplification of ideas, like a spiral. This interaction between my mind and the computer, and the loudspeaker output was to inform me, I informed the program by using it in different ways that came back with me suggesting ideas...

[Federica]: Feedback.

[John]: Yes feedback, that's what I was looking for. But, yeah, this kind of spiralling feedback loop, which was so enriching. For me it was such a joy. But I liked doing that. I want to write that procedure for that program. So, I don't care if they make music with... I mean, it's a great idea, it's a wonderful research project, and if there's a way to evaluate it, just as we do any kind of music, most of it is pop music, I think that their interests are not the same as mine. But nonetheless, there's a billion people who are interested in the results, and there's money to be made if Google could create music that people like that they've never heard before. Well, there are a lot of social issues and commercial issues that affect that, of course, but...

[Federica]: John, I am delighted to have met you and to share this conversation with our listeners. Thank you so much for taking the time.

[John]: My pleasure and thank you for being informed about that which we care about.

Thank you for listening to Technoculture! Check out more episodes at technoculture-podcast.com or visit our Facebook page at [technoculturepodcast](https://www.facebook.com/technoculturepodcast), and our Twitter account, hashtag [#technoculturepodcast](https://twitter.com/technoculturepodcast).