

Bruke Kifle: This is ACM ByteCast, a podcast series from the Association for Computing Machinery, the world's largest education and scientific computing society. We talk to researchers, practitioners and innovators who are at the intersection of computing research and practice. They share their experiences, the lessons they've learned and their own visions for the future of computing. I am your host, Bruke Kifle.

The intersection of computer science and music has led to remarkable innovations in how we create, perform, and experience music and the arts. Advancements in real-time interactive systems, AI and digital audio processing have opened up new possibilities for musicians and composers. At the forefront of this exciting field is our next guest, Dr. Roger Dannenberg, who has been pioneering groundbreaking technologies that bridge the gap between computer science and music for over four decades. Dr. Roger Dannenberg is a Professor Emeritus of Computer Science, Art & Music at Carnegie Mellon University. He received his PhD in computer science from Carnegie Mellon in 1982 and has since made significant contributions to the field of computer music.

Dr. Dannenberg is internationally renowned for his research in computer music, particularly in the areas of programming language design, real-time interactive systems and AI music. Throughout his career, he has developed innovative technologies that have revolutionized the music industry and is known for creating Audacity, the widely known and used audio editor. Surprisingly, in addition to his academic work, Dr. Dannenberg is an accomplished trumpet player and composer having performed in prestigious venues around the world. Dr. Roger Dannenberg, welcome to ByteCast.

Dr. Roger Danne...: Thanks for having me. It's good to be here.

Bruke Kifle: Very, very excited to have you. You have such a wide range of interests and experiences, and you've been able to bridge this personal passion and love for music with a very accomplished academic career. Could you describe some key inflection points in your personal and professional journey that have inspired you to pursue a career in computing, and specifically within your area of research at the intersection of computer science, arts and music?

Dr. Roger Danne...: Sure. I think that one thing that might be good to understand is that I've always been really attracted to mathematics and science and engineering. When I was a kid, I liked to build things. My father had a workshop and taught me a lot about using tools, and so that's just kind of part of my nature. And I was also really attracted to music early on and pursued music through trumpet playing and writing and arranging. And so I kind of grew up with both of these things as very separate directions. And one inflection that comes to mind is in high school I went to a music store that had an analog synthesizer on display, and these days that would be nothing or it'd be hard to imagine discovering synthesizers that late in life. But for me, these were very new things and really fascinating and I just really took to it.

I had a chance to just play around with this thing in the music store and being able to turn knobs and parametrically adjust sounds was something that just was really awakening and it sort of tied everything that I knew about mathematics and functions and physics of sound to the musical ideas that I had in my head. And in some sense, looking back on that, I see that that was really an awakening and there was no going back. So I did try to pursue computer science as a profession, a solid way to make a living and enjoy it, and at the same time, practicing and developing as a musician, which I thought I would, through computer science, I would have the resources and the freedom, make the kind of music I wanted to make. And it wasn't really until after graduate school in computer science that I started really putting the two together and doing research in the computer music area.

And I guess one other thing I would mention is that after I finished my PhD, there were people around the computer science department at Carnegie Mellon that were very encouraging. I think Allen Newell and Herb Simon were among those, these are great names and great people in computer science, and their attitude was that computer science was so all encompassing that the important thing was to do good research and not worry about anything else, and that attitude was really pervasive in the department. There were just all kinds of things going on. And so I felt invited and welcome to really pursue my passion and that was tremendously helpful.

Bruke Kifle: That's amazing. I think going from having a personal passion and then seeing that as independent or somewhat exclusive of your academic pursuits, but finally getting to a point where you were able to combine these two pursuits. And so being in a position where not only is music something that you can pursue in your free time as an exercise of enjoyment, but rather it becomes the focal point of your career, and so I think that's certainly something that a lot of people will take great inspiration from. I think as you described it, you've had this long career, not just as an academic, but also as an accomplished trumpet player and composer, how do you think your musical background and performance in sort of different cultural spaces had influenced your approach to teaching and scholarship more broadly? Do you find that there are parallels between music and your academic work, or do you find that it's shaped your perspective as a researcher?

Dr. Roger Danne...: I would say the most important thing that music has done for me is giving me a real basis for asking important questions and really forming research questions. So it's through the practice of music... Well, let me give just one example. One of the early things that I worked on was getting computers to listen to live performing musicians and play along with them. And that was sort of inspired by a talk that I heard where someone was building a sensor for a conductor's baton. And their idea was that if you could sense baton motions, then someone could actually conduct a computer through processing sensor data to figure out what the beat was and what the tempo was, and that would be a good way to achieve expression through computer music.

And it's a good idea, but actually I think it was musically uninformed because as a trumpet player who had played in orchestras and bands and worked with many different conductors, I knew that performers in the orchestra are not really following directly what comes from the baton. And if you try to synchronize visually with a baton, you're just not going to play in time. And so I started thinking about, "Well, how does that actually work? Where does time and tempo come from in an orchestra or a band?" And it comes from listening to the musicians around you, that there's this kind of collective entrainment is the technical term we use that enables people to play together, and our sense of hearing is temporally much more precise than our vision. So it's really necessary to listen.

And so that got me thinking about, "Okay, so maybe if you want to synchronize a computer to musicians, you should be listening to musicians," and I thought about how to do that, and that led to a computer accompaniment system and patents and a spinoff company. And that happened very early in my computer music career, and it's just an example of how being a musician and being concerned with real musical problems led to interesting solutions and interesting research questions.

Bruke Kifle: Well, that's amazing. So the ability to ask the right questions, but also having the personal context and knowledge yourself being a musician to pursue the right areas of research and then solutions. You touched on some of your early work, looking back from score following in the 1980s to some of your more recent projects like the Composystem. From your point of view, how has the field of interactive computer music involved since you began your work in the '80s?

Dr. Roger Danne...: One of the big changes and I would say probably the biggest change has been the speed of computing and hardware. It's hard to imagine in some ways how slow computers used to be. We're looking at machines that are a million times faster than just a few decades ago. And it's something that I thought about decades ago that, "Wow, computers are growing in speed exponentially. So what are the interesting problems going to be 10 or 20 years from now when we have so much more computational power?" And it's amazing how difficult it is to think about a world in which you have a million times more compute power. Maybe some problems, it's obvious what you would do with more power, but I think with music and audio, it wasn't so clear.

But one thing definitely that happened that changed is when I started working in the field, it was only possible really to make sound with specialized hardware. And now it's just routine to build synthesizers and very sophisticated audio processing totally in software. And that's just opened up so many possibilities for interaction. Now everyone has actually has a very sophisticated audio system in every browser in the form of Web Audio, which is just one example of how software audio signal processing has become just ubiquitous.

Bruke Kifle: Very interesting. And I think to your point, which kind of leads me to my next question, which is one of the key implications present day that we've seen of the acceleration of computing technologies is also in the field of AI. And so I'm sure as we look specifically at AI, I think maybe if you asked a lot of folks in the field 15, 20 years ago, a lot of folks saw a lot of potential for AI to enhance productivity with more routine tasks and less of some of the interesting creative tasks that we're seeing AI sort of excelling in today. And so what potential do you see for applications of some of the emerging AI capabilities that we have today in enhancing music composition and potentially live performances as well?

Dr. Roger Danne...: It's a great question. The potential is really huge. I was just thinking that back when I worked on score following and computer accompaniment, I helped a company develop a product that came out initially for the Macintosh, and at the time it was billed as the first consumer AI application. So that was considered artificial intelligence, and certainly for the time, we checked all the boxes. I think AI has been part of music research and making music along with so many other things, and we've seen an incredible acceleration of AI capabilities through deep learning. And I think we're just beginning to really understand what that can do for music and how to integrate that with music.

One of the things I'm interested in pursuing is using AI to build more interactive real-time systems. So far, most of the work has been towards less interactive systems. For example, doing music search or prompting an audio generation system with text and getting an audio composition out without a lot of interaction or real-time control. And I think that music making can be a very physical performative task and connecting some of these AI tools to real-time gesture sensors and live performance is a really interesting direction to pursue, and I think we'll see more of that in the future.

Bruke Kifle: Very interesting. And how do you generally think about how as a community we can think about balancing the role of AI or potentially AI assistance with human creativity in music while addressing sort of the importance of preserving some of these things to musicians and composers and the general ethical considerations with the role of AI in this space?

Dr. Roger Danne...: It's a very complicated situation. One thing I like to do is take a historical perspective and look at not just what's happening right now, but what's happened over the years and even over centuries that musicians have always adopted new technologies, and I think it's true of the arts in general. Even oil painting in the early days involved a lot of chemistry and processes and techniques that were being developed largely by artists themselves.

And in the music world, you look at the piano, it's just a marvel of mechanical construction, huge pieces of cast iron to withstand the tension of all those strings and intricate mechanisms in the keys. So huge amounts of engineering and technology to make a musical instrument just that could play loud and soft

with a keyboard. And so technology has always gone hand in hand with musical developments, and that followed into electronics and computers. And so it's not like AI is making fundamental changes in that sense.

The other thing to think about is that while there is some threat to musicians from artificial intelligence, generative audio algorithms for composition and synthesis and processing, it wasn't that long ago that we didn't have electronics and broadcast and transmission. And if you look at what the recording industry, a different technology did to live performance and the lives of many practicing musicians, I like to point out to people that when they talk about threats to musicians, that there aren't that many musicians left making a living at music because they've already been put out of work by the recording and the broadcast industries and all of that technology.

So things are going to change, and I think that change is inevitable, but I think that just like electronic musicians embraced the new electronics technology and created new artistic opportunities, we've got the same situation with new AI techniques. And so this will keep the arts alive and keep lots of interesting work out there for composers and creators. The downside of some of this is maybe it's going to be easier to make music coming out of black boxes and not requiring so much human labor, which has advantages and disadvantages, but it's just going to take some adjustment.

Bruke Kifle: And I think that's a really good way to put it. Of course, historically the role of technology has always existed and has always been a key enabler for further accelerating or improving the production of music, how we consume music, and so while there are of course certainly downsides in this day and age in sort of the advent of AI, we kind of have to embrace these technologies as an enabler. So I think that's well said. I would love to turn to some of your work around software design for computer music. Of course, as the co-creator of Audacity, I was the developer of the Nyquist programming language, what are some key principles that have guided your approach to designing software for computer music?

Dr. Roger Danne...: Well, one thing I can say about Audacity, which by the way I should mention that was co-creation with Dominic Mazzoni, who did a huge amount of work in the early days and was really the leader in working with people to put Audacity out there, which led to it becoming a huge open source project used by millions and millions of people. When we started Audacity, our initial goal was just to do visualization of audio waveforms and spectra to help with research because there weren't really great tools out there, and we thought it wouldn't be that hard to build something which was really the case.

And so I think one thing that would say is that the commercial stuff out there tended to be very task specific. The software was designed specifically for audio editing tasks, and I think that the designers were thinking more about taking a very narrow view of what users might want to do with something rather than

looking at editors and audio displays as a general capability that could be offered in a general form for any purpose and that's what we set out to do with Audacity.

The other thing is that the models for early audio editors treated audio files as these monolithic things, as if maybe they were pieces of tape like from a tape recorder. And so if you wanted to cut, copy paste, the early editors would copy entire sound files, which at the time, well, I mean they're still big, they're megabytes in size, and at that time, disk speeds were not so great, and copying a big multi-megabyte chunk of data took a lot of time and so these editors tended to be very slow. And we were able to look at this as computer scientists with some knowledge of data structures and really think about, "Well, where's the time going? And is that necessary? And do we have to store audio in one big monolithic file and keep it that way? Or can we break it up into chunks and reassemble chunks in ways that are quick so that we can make editing go much faster?"

And so that was early on one of the reasons I think Audacity caught on is that it was so fast for doing simple editing operations compared to all of the commercial software. And it was just because we took maybe a computer science viewpoint to simple data structures. It's not even that sophisticated, but it was obviously a good idea. And then I think similarly, if you look at Nyquist, which is a audio music composition language also used for sound synthesis, and it's actually embedded within Audacity now and so you can write audio effects in a very high level language within Audacity if you know how. So the whole point of Nyquist was to kind of revisit some existing computer music languages more from a language design and a computer science standpoint where things could be very much more unified.

Thinking about time and signals and building abstractions that could be put together in a very flexible way, and that led to a functional programming language. And interestingly, by using these abstractions, the implementation could be very specialized to different sorts of data types, and it actually ended up running much faster than some earlier designs that were more concrete. So normally we think that if you build a low level very close to the metal programming language, it's going to go faster than some abstract functional programming system. But Nyquist is at least one exception where having those nice abstractions allows you to do a lot of optimization that wouldn't happen otherwise.

Bruke Kifle: ACM Bytecast is available on Apple Podcasts, Google Podcasts, Podbean, Spotify, Stitcher and TuneIn. If you're enjoying this episode, please subscribe and leave us a review on your favorite platform. One thing that at least stood out to me is thinking about existing problems and identifying interesting or novel approaches from computer science or sort of core principles from computer science to solve those existing problems have been key to unlocking new opportunities in these spaces. How much do you see open source

development, I guess particularly in the case of Audacity, influencing both the accessibility and the widespread availability of music production?

Dr. Roger Danne...: Well, I think it's been huge. Audacity has had something like a million downloads a month for well over a decade. So you can do the math and figure out how many people are getting access to some high quality audio production tools that run on average even low end hardware. And I talked to so many young people that say that Audacity was their first exposure, and I think back to my days wandering into a music store in high school and discovering a analog synthesizer, and maybe the same thing is happening now to high school kids that download Audacity and record their guitar in their bedroom or whatever and start doing music production. It's really had a big impact there.

And at the same time, Audacity has been used... Almost every professional music producer that I talk to uses Audacity for one thing or another, just because it's maybe quicker to set up for small tasks, even though they have access to very sophisticated systems in their studios. So yeah, I think open source software, just because it's free, it encourages people to try it out and enables people to get things done.

And the other thing I would say, and this comes from Dick Moore who is a computer music guy at UCSD, and he, in talking about software and music software in particular, said that the software that academics and researchers create and put out there as open source may not always be the best software or the most sophisticated for any given task, but one thing it does is it sets a bar over which if you want to introduce some software into the market and sell it, you'd better do better than what's already out there for free. And so it's kind of establishes, has a strong influence in product design and systems design and just overall quality and capabilities that if you can't beat the stuff that researchers put out there for free, then you've got to try harder. And I think we've seen a lot of software kind of step up their game a little bit.

Bruke Kifle: Oh, that is certainly a very true point. It has indeed set a bar. And I think from your point in terms of democratizing sort of music production, but also more generally developing this new generation of folks who are just as inspired in music production and potentially the role of computers and computer science just as you were inspired by the analog synthesizer way back when I think must be an inspiring and moving driving force for you.

Dr. Roger Danne...: Right. Well, I hope so.

Bruke Kifle: So I think we touched on a couple of your contributions, and I sort of want to give space for you to highlight maybe one other project that you're most proud. Of course, looking back on your career, we talked about a couple of things with software design, programming languages, Piano Tutor, I think Audacity as well. Is there maybe another project or achievement that you're most proud of and would like to highlight?

Dr. Roger Danne...: Sure. It's kind of hard to pick something, because which of your children is your favorite? But before I move on to another thing, I just want to mention one thing I'm very proud of, hasn't been that long since I started collaborations with Jorge Sastre in writing an opera that's been translated into English and has had a few performances. And it's been a really great experience and it's a huge project that I'm very proud of. But another thing in the more technical direction, one thing that I've worked on recently is trying to understand music through studies of prediction and entropy because it's key as a music composer and music listener that music has... At a very abstract level, it's all about repetition and anticipation.

And we've tried to formalize some of that, as have many other people in computer science and music theory, but our approach has been to use especially hidden Markov models as predictors. So it's a form of sequence learning, and we've used the predictions to form estimates of how much repetition there is in music and where repetition occurs. And the repetition is how music achieves structure in many cases, maybe most cases, and once you can identify structure, you can start looking at the influence and interaction between structure and harmony and pitch and entropy or predictability.

And so I think this work is by no means finished, but it's a direction. And what I really hope that it will lead to is a deeper understanding of music, which currently there's a lot of, especially with all of the efforts to model music with artificial intelligence systems and deep learning, there's been a real focus on popular music because it's popular and it's what most people hear and listen to, and it's actually tends to be a little more formulaic than a lot of other musical forms and so I think it's easier to learn. And so what these systems are I think failing to accomplish is an understanding of what underlies all music. So not just popular music, but let's look at Beethoven, or let's look at Stravinsky, or let's look at Bebop.

How do all these forms evolve and how do they all call themselves music? What's going on that's common across all of this music? And I think a lot of it has to do with repetition and anticipation, and that's of course all tied to human perception because if we don't perceive it, then it's not going to... Doesn't really matter so much. And so anyway, I think that line of research has been really interesting, and we have a long way to go, but that's something I'd like to highlight.

Bruke Kifle: Very exciting, very exciting, and certainly an exciting area of further development. I would love to touch on some of your work sort of in the academia space as a professor and as instructor, having advised I'm sure a lot of students who've gone through the department and the field. What skills do you believe are essential for students specifically with computer music, but also entering more generally this interdisciplinary field, especially in light of some of the emerging technologies and trends that we're seeing these days?

Dr. Roger Danne...: I think for my students, the most successful ones have been the ones that really know a lot about music and have been educated musically and are really passionate about music. So I run into a lot of other researchers and lots of students who are technical and everyone loves music, and so it's common for people with a lot of technical skills and computer science interests to say, "Well, I could apply this to music," but they don't really know so much about music, and I think that makes it difficult and less successful.

So there's a lot said about interdisciplinary studies and interdisciplinary teams with the assumption or just observation that not everyone's an expert in everything, and maybe the way to make progress across disciplines is you have disciplinary specialists that talk to each other a lot, and you get some kind of benefit from getting experts together. And I think that's certainly true and there are lots of good case studies and examples of that. But personally, I've always worked more on personal projects or projects with me and a student and not so much in these broader interdisciplinary team sorts of approaches. And I think it's just really valuable to be really solidly grounded in music or whatever it is that you're interested in pursuing.

Bruke Kifle: Yeah, yeah, very much so. And so I guess maybe as we sort of wrap up this conversation, I'd love to get some of your thoughts on future directions for the field of computer music. So from your point of view, as somebody who's worked in this field for a number of decades, where do you see the field heading in say maybe the next 5 to 10 years, particularly as we think about some of the role with advancements in computing, some of the rapid progress that we're seeing with AI integration and interactive systems? So where do you see some of the exciting areas of research?

Dr. Roger Danne...: Well, it seems to me that artificial intelligence has become, in a sense, the new computer science. So just as in the early days when computers were becoming available, people were looking around saying, "Oh, I can apply computing to this and I can apply computing to that." We're seeing that with artificial intelligence that, "Oh, now I can apply machine learning to this and to that." And just as computers and computer programming became pervasive, I think intelligence systems are going to be pervasive in almost everything that we do, including music.

And so I think for the future, the big question out there is how do we make intelligence systems more interactive and collaborative and working with humans? So as I said earlier, I think real-time performance is a great area to think about applying AI and machine learning, even though machine learning is traditionally been a batch-oriented, take a whole bunch of data, try to distill it into some knowledge and build an inference engine. So that's kind of the paradigm that's developing, but I think that it has its shortcomings and thinking about real time and interaction is going to open up a lot of new possibilities. So I think that's a future thing to work on.

And otherwise, just generally interaction and control. We have amazing systems that are capable of taking prompts and creating stuff from that. What we don't have are systems that after they create something, allowing artists to go in and modify them, perfect them, teach the system something it didn't know that would make everything better or would push everything in a certain direction. People do that in a very limited way with prompting, but I think we have a lot to learn about how to really control systems, especially for artistic purposes, where the standards are so high.

Bruke Kifle: Very promising and exciting areas of work. And I guess maybe just to wrap up, of course, we have a lot of folks who tune into ByteCast, who are excited about computing, who are potentially interesting in establishing a career in the computing space, and much like you potentially at your interdisciplinary field of study. And I know you provided some interesting nuggets or perspectives on your work in with music and some of the models of students that you found to have great success. But more generally, what advice would you give to young professionals, young researchers, students who are interested in computing more broadly and potentially interdisciplinary field of study?

Dr. Roger Danne...: Well, I think the best advice I can give, and it's advice that many others have given so it's nothing new, but pursuing your passions is really important, that you will be tempted to do what people tell you to do or do what your parents tell you to do or do whatever there is out there that's offering the most money and the most income. And ultimately, I think that the thing that people should prioritize is having fun and enjoying themselves and being proud of their work. And you can achieve that by following what you're passionate about, and at least not having regrets of leaving opportunities behind that you really wish you had pursued.

And I think it's a difficult world to do that, and so it's difficult to give that advice to people because it doesn't guarantee that you'll be successful. But I think it's still a good thing to aspire to and a way to provide some guidance and look for the right opportunities and make the right decisions.

Bruke Kifle: I think that is a great piece of advice, to pursue your passions, to find things that provide you joy and fun and allow you to live a sense of purpose. And I think coming from you as somebody who's been able to live that advice and live that mission, I'm sure a lot of our listeners will take great inspiration in your journey and not just your advice. So Dr. Dannenberg, thank you so much for joining us on ByteCast, and we're very excited to see the future of the computer music field.

Dr. Roger Danne...: Oh, thank you.

Bruke Kifle: ACM ByteCast is a production of the Association for Computing Machinery's Practitioner Board. To learn more about ACM and its activities, visit acm.org. For

This transcript was exported on Dec 09, 2024 - view latest version [here](#).

more information about this and other episodes, please visit our website at learning.acm.O-R-G/B-Y-T-E-C-A-S-T. That's learning.acm.org/bytecast.