Code Together Podcast Episode 25: ***Inspiring the Next Generation of Scientists & Coders***
June 7, 2021
Host: Nicole Huesman, Intel
Guests: Katherine Riley, Argonne Leadership Computing Facility; Joe Curley, Intel

---

Nicole Huesman: Welcome to Code Together, an interview series exploring the possibilities across architecture development with those who live it. I'm your host, Nicole Houstman. Today's discussion focuses on inspiring the next generation of scientists and coders. And it's my sincere pleasure to welcome two guests to today's program. Katherine Riley is the director of science for Argonne's Leadership Computing Facility, or ALCF. Where she leads a fantastic team of experts in computational science, performance engineering, visualizations, and data sciences. She has helped guide the facility's strategic vision, assisting with the design and development of its super computers by identifying scientific requirements, applying her career long focus on how scientific application architecture impacts performance, scalability, and extensibility. So great to have you with us today Katherine.

Katherine Riley: Oh, thank you. It is my pleasure to be here, excited for our conversation.

Nicole: And Joe Curley. He serves as senior director of oneAPI products, solutions and ecosystem. His primary responsibilities include supporting the oneAPI industry initiative and its developer ecosystem as well as product management of Intel's oneAPI product implementation. Over the last 14 years at Intel, Joe has held multiple other strategic planning, ecosystem development and business leadership roles. Thanks so much for joining the program, Joe.

Joe Curley: It's great to be here, Nicole.

Nicole: Katherine, can you kick off today's discussion by talking about the work that you do at ALCF?

Katherine: Yeah, I'd love to. So, the starter point, right, is kind of what all of those words mean. What we are is a national user facility. Why we exist is to build the biggest computers in the world in order to tackle the biggest scientific and engineering challenges, ones that you can't do, unless we really build a giant supercomputer to tackle your problem. And in the end, right in designing a supercomputer like that, they're always kind of bleeding edge, right? They're always really pushing technology forward and to help scientists with that, we actually have a pretty large team of scientists at the facility and computer scientists at the facility where we collaborate with researchers who think they need to use a supercomputer who are trying to use a supercomputer to help them move it forward, right. To help them understand what their requirements are and how we can get them tackling their science.

So, the other way to describe this, as a national user facility, we're open to everybody, right? Everybody can apply to use the ALCF and it's a peer review process. One of the things that we view as really important is making sure that there is a pipeline of people, right? That there's new scientists coming up and engineers coming up who have some idea about how they would like to use high-performance computing to tackle their research problems. And that's one of our other big components of the work that we do is really not just, "Hey, we have this big machine, do you want to use it?" But also making sure we've got people coming in who are thinking about that, in their future research plans in the next 5, 10, 15 years.

Nicole: Thanks so much, Katherine, for explaining really how ALCF works. Brick and motor labs have really changed. Joe, can you talk a little bit more about this?

Joe: When you get to be of a certain age, which is what all of us who are getting older start to say is we grew up in an era of test tubes, laboratories. If you wanted to test how a laser would end up working, you tried to build a laser in the basement of your college lab and you tried to fire it and you'd try to see what would happen through various different mediums. That's expensive, it's time-consuming, it's error-prone and the scientific method starts with a theory and testing the theory out and gathering results, and then continuing to iterate through all of that. The lab of the present has changed, brick and mortar labs do that same exact experiment and trying to look at how one might build that laser can be done entirely in simulation. And there are some reasons why you'd want to do that.

One of the first ones is it's going to be quicker, less expensive. You can take a very, very complex experiment, model it, get a result. And when you get your errors, you can iterate very quickly on top of that experimentation before actually going to a physical prototype. And so, we've moved from an era, 30 years ago, of paper design to digital design and from digital design to digital simulation. And now we're in digital prototyping. And so, the idea that supercomputers can take on some of the most massive challenges of the world. We're reading a lot about COVID. But trying to figure out how the protein hits the cell wall and how the bio mechanisms work that can be done iteratively over years with lots of loss of life through human experimentation, or it could be done in a supercomputer, in a relatively shorter time. And that's what's happened, is that supercomputers have replaced kind of the brick and mortar lab environment, broadly for some of the smallest challenges of science. I don't know, Katherine, would you agree with that assertion?

Katherine: I would emphatically, right. I think there's also this component that people kind of lose track of which there are even things you can't do in a lab, right? Where you really need to have a supercomputer to help you. And I mean, some of these are ones that we're used to thinking about, or at least you might've heard before, which … You can't blow up a star in a lab. You can't evolve the universe in the lab, but very similarly, right, you can't do certain types of experiments on people in a lab, you just can't do that. And so, the computer allows you to really explore these questions that you wouldn't be able to explore otherwise. And this component, that I think it's a necessary part of almost all science. You don't necessarily need a supercomputer for every question you're asking, but I think you need to understand your computer as a tool for almost any science you're doing today.

Joe: One of the really fascinating things has happened over the last couple of years is how you can deal with the unknown.

Katherine: Yep.

Joe: And so, one of the things at Argonne in the leadership at the forefront of this is taking a look at combining simulation and data science, and for people who haven't walked out on this, there are equations, there are physical principles or scientific principles that are well known and you can build a model of it and simulate it. But then there's a whole bunch of things we don't know. For example, we have theories about dark matter or dark energy. So how do you do a calculation on something that you actually fundamentally don't understand? Scientists can take data from the universe and come up with ideas like a cosmological constant for something you don't understand yet that can help you get to the understanding of it. And this can only be done in a simulation environment. So, there are all kinds of things you can do and can't do because of physical reasons. But some of it just because you just don't know yet.

Katherine: And one of the ways I try to explain this to people so that they don't respond and go, "Okay, so you're just making stuff up" is this is like, think even fundamentally, to what you were learning about science in school, gravity was not understood. We didn't understand gravity. It first started with trying to understand how things were falling and observations about this. So even without a physical model for what gravity is, and that's a fundamental force, we were just starting with that, and it was just that observation about different bodies and the rate at which they fall that started that. You're exactly right. These types of questions that we have in some cases are just watching those objects fall and starting to sort of figure out maybe going backwards from there and finding out the more fundamental physics from those objects dropping.

But I also think the other one that's really an important one today is you also have scientists and researchers who are coming across other ways of doing their experiments and are being confronted with data that they can't manage anymore. And it's not because they're not good scientists and engineers and researchers. It's because you, as a human, can only process so much information. And if you're getting hundreds and hundreds of petabytes or even exabytes, say, of data from an experiment that you're doing, and that's pretty easy to do these days, you need a resource to help you understand that data. You need to do some really quite advanced and interesting analytics in order to pull off the science that you're really looking for. And that can be both more straightforward, just like I just need to crunch those numbers and figure out something, like in the way you might think of from your labs at school, but also applying completely new methods and trying to understand what's in there, right?

The reason people talk about deep learning and learning techniques today is that you can use some of those to help you find patterns. All these are is pattern matching ideas, right? You're looking for patterns in exabytes of data and that's another thing that supercomputers bring to the world of science that, honestly, is a little bit new. We did this on a different scale historically, but now we're being presented with challenges that really need incredibly large supercomputers and incredibly large data capabilities because you got to parse these exabytes of data coming off of experiments.

Joe: So, we've talked a lot about what, but the example you just brought up sparked an idea in my head of a very good friend of mine who came down with a unique, very personalized cancer. And it hit him in a profound way. And without getting into the drama of the story, he wasn't given very long to live. Thankfully, he worked in a place where there's Intel and there's the Oregon Health and Sciences University and a lot of very big computers. And they were able to actually take a look at his biology, his *particular* biology, and come up with ways to understand what was going on in the mechanisms of that cancer. And then, actually, I'm amazed to say that seven years later, and you can see his website here at Intel and otherwise, he's doing great. But science is science and the reason why people get into science and the reason why people use computers is because we want our kids to have a great planet. We want my friend to have a great life. We would like to stop putting so much carbon dioxide into the atmosphere, but I still want to go places. There are real things we want to do that science helps us do that helps us understand it. And so, the thing I guess I wanted to start with is we've got all these great computers, but if we're going to inspire people to use them, I really also just want to give a little quick plug for Argonne and others. By the way, one of the reasons for the Aurora supercomputer is to actually look in a very big way at exactly the kind of problem that Bryce went through. But there are compelling individual human- and planet-wide problems that we really want to address and we're out to get them. So, it's not just nerds working on ones and zeros, but we're out to try to make a better planet.

Katherine: And I hope people have some confidence, right, that we're not just pontificating and presenting the big questions that we'll never get answers. Some really real examples of ways that supercomputers have changed your daily life is probably in your parking spot. Cars did not look like they do today, 30 years ago. And the auto industry really was changed dramatically by the introduction of supercomputers and that didn't just necessarily help them, it helped you, right? And the reason is it costs a lot of money to do crash tests, so maybe a car would be designed with five or 10 crash tests. Instead, what you can do at this point is do thousands and thousands of crash tests and understand how the car behaves. And as a result, you get a car that's a heck of a lot safer to drive. And a lot of the innovations of safety in the car come from using supercomputers.

And so, while that's a very, very specific one, right, at the same time, this is directly connected to what Joe's talking about. When we want to understand ways that we can still go places and explore the world, but also stop putting CO_2 in the atmosphere, we're looking at how you improve the design of airplanes. How can you improve our fuel sources, right? How can you change our fuel sources? These are real concrete questions we're looking at. How can you turn your windows into a solar cell, not just putting solar cells on your roof, but actually sort of have a button, and your windows become opaque, so you have shades at night, but they're always also fueling your house with electricity. These are the types of questions we're actually tackling, they're not theoretical ones, we're working on them.

Joe: If the higher calling doesn't work for you, I can just give you greed, money. About seven years ago, the Audi group demonstrated at a Supercomputing conference with us that they had gone from physical design of cars and prototypes crash testing, which Katherine talked about, but kind of the last and most expensive part of auto making was coming up with the physical prototypes, because you want to know how the car will actually look. What they were able to get just like a regular data center rack and turn it into a supercomputer and actually design, prototype and build all the way down to fit, finish, fully ray traced, visual examples of that car. And the first article that was produced was in the factory for production. Why that was important was, they took years out of their development cycle. It gave them a competitive edge of years in the marketplace with new technologies over their competitors. And if you looked at the ascendance of the brand of that particular make of automobile versus its competition, they really exploded in growth. So, your health, your happiness, make the planet better, or just make money. There's lots of reasons to do this.

Katherine: I'll also say though, that applies to scientists too, right? You want to be the first. The questions that are driving you, whatever that question might be, you probably have other people who are also looking into it as well. And there is something exciting about being the first.

Nicole: So, Katherine and Joe, as we invite newcomers, as we invite this next generation of scientists and coders to really join us, to join you, can you talk about your experiences? What inspired you in your career? What helped you along and how do we bring this next generation along with us?

Katherine: I started originally wanting to study astrophysics. This is why I had the star example earlier. It's always one I come back to. But what actually distracted me and sent me down the course was this idea of building this incredible scientific tool.

When you're looking at using a supercomputer, it's not just building that supercomputer, but it's building that scientific application that codifies what your understanding of physics is at that particular time, understanding the science, but also aspects of understanding the mathematics that go into it. It's a very multidisciplinary field, right? And that's really exciting. You have to have a lot of understanding and more importantly, you have to know where to ask questions. You have to understand where your knowledge ends and where you really have to start including other people and asking the right question. And I actually thought that multidisciplinary component, right, of using computers to tackle questions was really motivating.

It was really exciting because it's hard just to go off and do it on your own, which is a weird thing for me to say, because I like to hide in my room and get work done. But at the same time, you get more out of a collaborative effort. You get more out of people thinking and working together to tackle a problem, so that was the first thing. What really drew me into this was this idea that you're building this incredible tool, multiple tools, the computer *and* the application and there has to be relationship between them *and* the question of the science you're trying to tackle. And that's incredibly complicated and fun. But I think what really worked with me is I got lucky. I would say over and over again in my career, I got lucky. And along with that, I made the choice to sort of dive into something.

And this is something I try to encourage people to do. I was also lucky to be able to take some of the jumps that I took, and I grant that. But as I was given opportunities, I sort of had to acknowledge, this is really exciting to me right now. I know I had this whole plan, I was going to go this way, but I actually went that way, right? And I took the jump just to sort of see what would happen. And I'm really grateful that I did. And I had people around me who were more senior to me, who were very established, who could back me up in that. And that was really important.

Honestly, I think that that encouragement, it might not have been an everyday pat on the back, but it was like, no, go do this, this sounds exciting. And I think that was very important because it's a multidisciplinary thing. You have to go learn a lot of things that you're not going to go learn in the classroom. You have to be willing to sort of pull out the books, do it on your own and experiment and learn on your own and explore. And so, you need that support from the people around you to do it.

Joe: I also feel really lucky. I lucked into high-performance computing. When I talked to my kids a few years ago, gosh, almost 10 years ago, we were working on the very first teraflop computing chip. It was a big deal getting it out, but trying to launch a semiconductor of that magnitude, I mean, it was backbreaking work that was really stressful, and my son asked me why I did it. And what I found is that for the current generation of kids that are coming up, this is like the question, why are you doing any of this stuff? I mean, it's work.

And I think that he actually had exactly the right question and what we do is really important, but why you're doing it ends up actually being more. Later that chip helped professor Hawking and his team take a look at the universe, visualized in the state that it was in just after the universe was cool enough to actually visualize, which to me was really wonderful. And hearing back from them that they got those results. Or, more personally, when my phone rings and it's my credit card company saying, are you currently in Eastern Europe ringing up charges? No, I'm not, please don't process those.

Why we do things is probably more important than sometimes what we do. And what I would really do is take a focus on what do you want to do with your life? If you want better medicine, if you want breathable air, if you want clean water, if you want your bank account safe, if you want to drive a really cool car, how do you contribute to some of that? And you start with those things that you want to do.

You know, Katherine’s job is hard. We're lucky to have them [our jobs], they're really cool. But it starts with an act of co-mission when you drive to work, if you drive to work, when you put your key in the ignition of saying, I actually want to go do this. So, I think the first thing is find your calling, understand what it is that you want to do in life, and then all the rest of it, the luck kind of finds you. What Katherine was talking about, ‘I was real lucky I got this opportunity.’ She came in and started with a general area and a passion and then it evolved into something here.

The second thing is, once you've got your passion, learn. The thing about computing, it's the best programmers, the best parallel programmers I run across, oddly enough, they tend to be physicists, chemists and whatever else that had a problem, that was big enough that they really had to master the tool. They couldn't just script up a program and hope that they could get a term back. They actually had to gain control, like an artist gets control with their palette or a mathematician controls their numbers. So, I guess the first thing is, start with your passion and then two, once you do it really be great at it. Katherine, would you disagree?

Katherine: I would definitely not disagree. That's where you find the innovations, that's where you find the nuggets of discovery is when you do throw yourself into it. And that includes, for example, understanding your tools, right? And just trying to explore being creative. I think there's also less within the group of scientists and engineers, but often this sort of stereotype that the scientist isn't creative, we're just crunching numbers. It's an old stereotype, but it's still there. You have to be really creative. It's not necessarily the creative, like, I want to go paint a beautiful piece of art. The act of doing science and research and writing your code to execute the kind of research you want to do, or even if it's writing your code to support a research effort, right, and it's not necessarily yours, that's a creative process and it's a fun process. And the more you understand the tool, honestly, I feel like the more creative you can be. You see that in artists as well, the more they explore a medium that they're using, the better they can do it. It's exact same thing. It's just that our medium is code on the computer. And I can't undervalue that 'be good at it' component and 'want to learn' because otherwise, no matter what job I think you're doing, you're going to get bored. Inherently, people are creative, people are intelligent. And if we just stagnate, you kind of stop. And so, for me, life is interesting, and work, therefore, is interesting.

Joe: And learning is so critical. But the other one is a little bit of resilience because along the way, generally speaking, if it's really worthwhile to do, it's not easy. And there are a lot of barriers that will crop up: It's not cool. It's not for you. I hate to say this, but I mean, Katherine, you're probably in a much better position to talk about this than I am, but there are a lot of social barriers that pop up that tell you that you don't want to be smart. You don't want to code. You don't want to do these things. Don't listen. If it's worth doing, it's worth doing, do you know what I'm getting at with that?

Katherine: Yeah, I would agree. One of the things that I tell a lot of young people who are considering a challenging career, honestly, I think it applies anywhere, but especially as I know my career and I know my field, I tell them to learn how to negotiate and to learn their why. And that's almost a cliche that last bit, but both of those things. And I don't mean negotiate, not just like I want more money for my job, although that's a component of it.

Joe: Yeah.

Katherine: That's an important one. But also, this component of, like, you learn confidence from learning something like how to negotiate and how to debate. You learn to process that this person's telling me, don't do this, why would you want to do this? And you learn to call crap on it. You learn how to, in the moment, realize that, you know what, this is not relevant, I don't care if they're saying this is not something I should do, I want to keep pushing and I want to keep pursuing it.

But I will say again, it's learning. We're such social creatures, it's hard to learn that skill, and I think the earlier people are exposed to it, the better. I think every single young woman I have ever interacted with who I have talked to about going into science or engineering, I have recommended this because you have to be self-motivated and you have to learn how to edit what people are saying around you. Otherwise, frankly, unfortunately, the truth is, you're going to get inundated with a lot of things and be overwhelmed by what everybody else wants around you, and not the focus that you need.

Joe: Yeah. So, if we're talking about how to inspire the next generation, Katherine's comments are so on target. I mean, have a great why, be in this for a higher purpose. Be pragmatic, but don't listen to people that put false barriers up in front of you. And really, really dig in. There are an enormous number of tools and capabilities out there: people willing to mentor you online resources, companies like Intel that are willing to work with you early in your career and give you access to tutorials and things to see if this is a game you really want to jump into. And then, just don't give up.

Katherine: And similarly, even on the Department of Energy side, but also other government agencies like the National Science Foundation, just to choose two, for example, they have a lot of programs that start in middle school even, that allow you an opportunity to connect to real scientists and your school might be connected to them already, right, if you're a student, but you can also just go pursue them on your own, but opportunities to get out of your circle, your school, that pocket and talking to somebody else. And I think that's, as Joe was saying, which is really right that opportunity to find somebody, maybe they're not your mentor for the next 15 years, maybe it's just a year or a month of hearing that perspective and being able to sort of pull what you need to out of that and move forward.

That I cannot emphasize enough, these kinds of programs, whether you're either coming in for just a summer student status or coming in for a year are really valuable. It's not necessarily that you're going to come in and do the sexiest thing in the world, although you might, but it's true that you'll make the connections. And I think the human component, when you're young and thinking about what you might want to do and the challenges associated with it, you're going to get a better set of tools to help yourself. If you do that with exposure to people who are actually doing the job. I can't emphasize that enough.

I'll also sort of say that we try very hard at the national lab level with this is part of the Department of Energy, the national lab infrastructure, to actually give even further resources to students and to give them an opportunity to connect with somebody because we know that that's important. It's a little self-serving, we want a good pipe of scientists coming through, right? We want to see that drive and those sort of innovative ideas coming forward that you really need, you need new people coming in, new blood and diversity in order to get really the innovations that we want to see. So, it's a little self-serving because we want to see that succeed, but they are there and please pursue them.

Joe: One of the things I really loved is kind of a best-in-class experience in supercomputing. I had a team of about 15 really quite remarkable engineers, but one of my senior principal engineers, a master coder guru, and knows everything that you'd ever want to know about getting performance, was having an argument with a new hire. And was talking about how he could get this thing to work a little better and they were going back and forth and what was fun was listening to the new hire going, "Why are you wasting your time? What you're doing is chasing 3 or 4%, but if you stopped and rethought this, we could go after a hundred." And what was cool in that environment was that the senior principal engineer said, "Wait a minute, I keep forgetting, you guys are thinking out of the box."

And by the way, we did achieve the hundred X. The thing is, when you come into this, Katherine, you were saying, you want new blood in, but you want new blood because it brings in new ideas. It really does advance science and don't be intimidated by the senior principal engineer, be in there with your ideas and you'd be shocked how much you can bring to the table reasonably quickly.

Katherine: And I will say, I think going back to connecting another point, this is where learning how to speak up for yourself is really valuable. You do have to listen, right. There are things you have to learn too, but it's exactly right. Especially as you get older, I see this happening, my box shrinks. It just does. There's some good that comes out of that, but also some bad things that come out of that. And so just realize sometimes what you're doing is you're trying to like open a box that really doesn't want to open on some folks.

Joe: Yes. I guess the thing I'd encourage as we're kind of ramping all of this down is if you're thinking about a career, one of the neat things is that not everybody's chasing this. There's a lot of opportunity. It's literally infinite. We have the best toys in the world, you can play with are supercomputers. You can model anything. Jump in. It's a chance for you to really have an impact. And I loved the point that Katherine was making about starting now and talking about NSF programs and just using the web and looking to how you can get engaged and sample it. You are the future, and the opportunity is that you can jump in and use these types of tools we're working on for things we have never imagined. And it is kind of up to you. So, we want to inspire you to go do something great. Our whole company's built on that.

Katherine: I feel like a little bit we've focused on even, very early you're maybe in high school or college thinking about these things, but this is not something like, well, I'm in an experimental lab right now getting my PhD. I don't think many people would necessarily sort of just write it off. "Okay, that's exciting." But, I will say, there's a lot of opportunities, even as an early career scientist to sort of explore this and develop capabilities again, through programs, at least in the labs that I can speak to. There's a lot of opportunities to plug in and learn what you need to do, learn how you might do it, figure out who you have to talk to, including summer schools that we have. Argonne runs a summer school every year and, basically, we fire hose you with how to do high-performance computing. But we also have programs even on the way that you use our resources, our supercomputers. We have options for early career tracks where people who are just starting out in their career and are figuring out really what they could do on large-scale resource opportunities for people to plug in there too. So, I would say that globally, it's never too late to consider it because it's just considering a new tool in your scientific tool chest, a big one that you can ask almost anything on like Joe is saying. You can do just about anything, it's just sort of the level of your commitment and creativity really.

Joe: Oh, by the way, Katherine, just to let you know, I'm registering for summer school.

Katherine: Excellent. It's awesome. I was part of the inaugural year, the person who had the job before me, Paul Messina and I really sort of brainstormed and built this and it's an awesome fire hose of information and that benefit you get of just, like, two weeks of you can think of nothing else.

Nicole: This has been such a great discussion. And this has really been such a special treat for me. I am so inspired by you both. So, thank you. As we wrap up today, Katherine, where can listeners go to learn more?

Katherine: The primary place that would launch you to everything is the Argonne Leadership Computing Facilities webpage, and that's alcf.anl.gov. And as I said, that will link you to almost everything, including even summer programs that you might be interested in. Argonne is as a whole lab, not even our facility is very connected to the Department of Energy's student programs. And so, that's all accessible there and that's an easy single place you can go that'll talk about everything that we do as a facility and I think the Department of Energy does, can be launched off of there, but also, all of our programs. If you're interested in trying to get access to a supercomputer, it'll tell you how to do that.

Nicole: Great. Thanks, Katherine. Joe, where would you point listeners to learn more?

Joe: There's so many wonderful sites, but if you just started software.intel.com, there are links to both our own sites, as well as sites in our community, around tools, languages, tutorials, downloadable samples, all kinds of things to get you started in the computer science part of the science. Additionally, if you're just a little more advanced and you're already starting a project, there's something that we call the oneAPI DevCloud. If you just google, oneAPI DevCloud, that is an environment that you can actually come into and start coding, build an account, start to code and experiment, get tutorials, work through things before you download your first set of tools or compilation environment to your computer. So, lots of stuff at software.intel.com.

Nicole: Excellent. Joe, thank you so much for being here with us today. Really appreciate it.

Joe: Oh, absolutely a pleasure. Maybe we can follow up some time and actually talk about how to code.

Katherine: I would love that. I could spend a whole series on that.

Nicole: That'd be great. We would absolutely love to have you both back on the program. Katherine, it has been so delightful to have you join the program today. Thank you so much for sharing your insights.

Katherine: Thank you. It was a pleasure. It was a lot of fun.

Nicole: And a big thank you to all of our listeners for joining us. Let's continue the conversation at oneapi.com.