

HPC User Forum Update

Interviews with Steering Committee Members: Doug Kothe, Exascale Computing Project

Steve Conway and Thomas Gerard
June 2020

IN THIS UPDATE

After the global pandemic forced Hyperion Research to cancel the April 2020 HPC User Forum planned for Princeton, New Jersey, we decided to reach out to the HPC community in another way – by publishing a series of interviews with members of the HPC User Forum Steering Committee. Our hope is that these seasoned leaders' perspectives on HPC's past, present and future will be interesting and beneficial to others. To conduct the interviews, Hyperion Research engaged Rich Brueckner, president of insideHPC Media. We welcome comments and questions addressed to Steve Conway, sconway@hyperionres.com or Earl Joseph, ejoseph@hyperionres.com.

This interview is with Douglas B. (Doug) Kothe, who has over three decades of experience in conducting and leading applied R&D in computational applications designed to simulate complex physical phenomena in the energy, defense, and manufacturing sectors. Kothe is currently the Director of the Exascale Computing Project (ECP). Prior to that, he was Deputy Associate Laboratory Director of the Computing and Computational Sciences Directorate (CCSD) at Oak Ridge National Laboratory (ORNL). Other prior positions for Kothe at ORNL, where he has been since 2006, include Director of the Consortium for Advanced Simulation of Light Water Reactors, DOE's first Energy Innovation Hub (2010-2015), and Director of Science at the National Center for Computational Sciences (2006-2010). Before coming to ORNL, Kothe spent 20 years at Los Alamos National Laboratory, where he held a number of technical, line, and program management positions, with a common theme being the development and application of modeling and simulation technologies targeting multi-physics phenomena characterized in part by the presence of compressible or incompressible interfacial fluid flow. Kothe also spent one year at Lawrence Livermore National Laboratory in the late 1980s as a physicist in defense sciences. Doug holds a Bachelors in Science in Chemical Engineering from the University of Missouri - Columbia (1983) and a Masters in Science (1986) and Doctor of Philosophy (1987) in Nuclear Engineering from Purdue University.

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. More than 75 HPC User Forum meetings have been held in the Americas, Europe and the Asia-Pacific region since the organization's founding in 2000.

DOUG KOTHE INTERVIEWED BY RICH BRUECKNER, insideHPC

Brueckner: Today my guest is Doug Kothe from the Exascale Computing Project. Doug, welcome to the show today.

Kothe: Thanks, Rich. Great to be here.

Brueckner: Well, it's great to see you through the social distancing. We see each other two or three times a year, it seems like, for HPC User Forums and the various conferences, but we want to get some more background from the Steering Committee. Let's start there, how did you get into HPC?

Kothe: I'm here in the basement of my Knoxville home. We're all working at home these days so it's good to be able to connect with you, Rich, and the HPC community. For me, it was the mid-1980's, I was doing my PhD research at Purdue University in inertial confinement fusion, so I guess I would call myself a trained physicist. It was really fortuitous. My professor, Chan Choi, was a consultant at Los Alamos National Lab at the time and took me out there in the summer of 1984, where I was exposed to computing. I sat in an inertial fusion and plasma theory group and was mentored by Jerry Brackbill, in particular, but a number of staff there that exposed me to their plasma physics codes and hydrodynamics codes. And at that time, I was logging in remotely on a 300-band modem to MFEECC, the Magnetic Fusion Energy Computer Center, which is now NERSC. My first exposure was running some Fortran, probably Fortran 66 codes, in fusion and just got hooked on being able to will the computer to do interesting science and physics. I was hooked from the summer of 1984 that way and didn't plan it, but I'm sure glad it happened

Brueckner: So, what have you been up to since, at least since you've got into this ECP role?

Kothe: I've been at Oak Ridge National Lab now for almost 15 years and before that, Los Alamos National Lab, as I mentioned, for 20 years, and Lawrence Livermore for one. I guess I'm doing the DoE lab tour these days. That said, at Oak Ridge National Lab I've had three jobs, and one of the nice things about national labs is they're so full of diverse people and diverse sets of disciplines that you find yourself getting interested in new problems and challenged by new things, and you don't have to be a careerist and plan your next move. Interesting opportunities just happen.

For me, before the Exascale Computing Project I was fortunate enough to lead the first DoE energy innovation hub, known as CASL, the Consortium for Advanced Simulation of Light Water Reactors. We built a virtual reactor that I think by all measures was very successful, known as VERA, the Virtual Environment for Reactor Applications. I did that from 2010 to 2015. Since April Fool's Day, 2015, I've been involved in trying to initially help start up the Exascale Computing Project with Paul Messina and running it the last two-and-a-half years. I first came to Oak Ridge in 2006, Thomas Zacharia and Jeff Nichols recruited me. I want to say I was the tenth or eleventh new hire into the Oak Ridge Leadership Computing Facility. They had one Leadership Computer Facility award jointly with Argonne, I think. When I say they, I mean Oak Ridge and Argonne in late 2004. They recruited me about a year later to come into the OLCF, the Oak Ridge Leadership Computing Facility, as their first director of science. That was a fantastic job that they let me define as my own and I did that job from 2006 to 2010. I'm on my third job right now, Rich, and really enjoying it here at Oak Ridge.

Brueckner: That's terrific. Well, Doug, you've certainly seen a lot of changes in computing, but what are the major changes that you've seen in HPC?

Kothe: Oh, man. That's a tough one because we've seen so many changes. Every change seems disruptive and scary. But, in retrospect, at least in my experience, they have not been that much. I think it's because, as a community, we're fast moving, we're agile, we're always learning. I think we're really on the cutting edge of our craft. I'm going to say this not because I've been in this community my whole career. I kind of have and kind of haven't, meaning I've done stints in nuclear energy and with fusion and other areas, and I think other domain sciences don't move as quickly as high-performance-computing, so every change seems scary and disruptive and in retrospect I think we forge our way through. I can remember working on the [Thinking Machines] Connection Machine and really loving the SIMD approach and programming in Connection Machine Fortran, CMS, and remembering Fortran 90 and really feeling like PVM and MPI were going to be super disruptive and scary. You're thinking, is this really going to work? Over a period of a few months, we ported our code from CMF, Connection Machine Fortran (and I still have a copy of that code, by the way) to MPI and it was surprisingly not easy, but we pulled it off and it was fun and it worked and we looked back going, wow, this actually wasn't as bad as we thought.

Well, here we are now, especially in ECP, where we're worried about accelerators. We've got different types of GPUs and, frankly, at Oak Ridge, I was part of standing up the first large system of GPUs, although I would argue that the cell-based system, certainly Roadrunner at Los Alamos, was the same situation. Sure enough, I think we found that it was disruptive and scary, but we pulled it off. We're seeing the same thing in ECP, Rich. We're figuring things out and knocking down barriers and exciting things are happening to make accelerators really more demystified and accessible – everything from interesting abstraction layers that creative people are working on, like RAJA and Kokkos, to pushing the standards of OpenMP to moving toward what I would hope to be a more unified programming model for accelerators. This has been a big, big change. It's hard for me, Rich, being an old person here, to not look back and say, you know, every one of these things was somewhat disruptive. I think the talent of the community and the agility of the community really is a testament to the fact that we collectively pull it off.

Brueckner: Well, you've kind of touched on this but where do you think HPC is headed? Is it heterogenous accelerated computing? You're at the forefront of this with your team.

Kothe: Certainly, I do think that is the case, Rich. I think that we would've failed in ECP if we developed a one-off boutique set of codes and programming models. We're not really forcing, nor should we, one programming model. I do think what we're seeing is, we are collectively figuring out how to adapt to and program on very heterogenous architectures. I think they're here to stay, certainly in my career, but I think, frankly, for decades. If you view an accelerator, and I view it this way, as hardware that is tuned to accelerate certain operations, let's just take the Volta V100. It really does make its matrix multiply accumulate 4x4 tensor stream. That's a very specific hardware acceleration of a specific algorithm. By that definition I think they're here to stay. We're going to see accelerators that are tuned to certain operations and certainly AI/machine learning is driving that roadmap. I think we're going to see all kinds of interesting possibilities that are based on computational motifs or domain-based languages. Even ASICs, certainly with neuromorphic and quantum, one could step back and say, this is an accelerator. I think what we need to do in ECP is develop the underpinning and the groundwork to be able to exploit those accelerators in a way that is not too, back to your term, disruptive to what we've done already to our applications. The proof will be in the pudding, but I feel like we're on a good path.

Brueckner: Are there any HPC trends that have you particularly excited or concerned?

Kothe: Mostly excited, Rich. Let me just be a little retrospective and then try to look forward a little bit. Relative to 10 or 15 years ago, Rich, the development environment, the IDE, for computing is really very mature. I'm really continually amazed at the tools we have at our disposal to do configuration management, to do issue tracking, to do continuous integration and testing. We're really fortunate that we don't have to essentially roll our own a lot like we used to, roll our own testing system, our build system. These things are black holes for people's time and talent and funding, frankly. Right now, we're all trying to figure out, and maybe we won't, what's the best container? What's the best virtualization? While we are not thrashing around, there's this tremendous technology.

I think I am concerned that the community won't move in a direction that is sort of unified with regard to containerization. Here's where the HPC community needs to sit more closely with the cloud community and say, we're one community, we need to learn from each other. I think the cloud community is really, frankly, ahead of the more traditional, science-based HPC community. I wouldn't call that a concern, Rich, as much as something we need to proactively manage, which is this more productized, containerized virtualization where we are bringing in, let's talk heterogeneity now, in terms of multiple machines running together in a coordinated fashion for a larger workflow which makes sense say, for, machine learning-based augmented workflows relative to science-based. So, now, instead of a particular node being tuned for a particular aspect of the workflow, it's one system sitting in one state and another system sitting in another state and we're working together in a much more federated, cohesive fashion. Those are challenges moving ahead. I think the containerization aspect is one key technology that's going to be enabling that.

Let me just make a comment about the time that we're living in, I think that we're seeing, with this pandemic, the merits and efficacy and importance of computational biology. My own personal view is the science community and the HPC community really have stepped up to the plate and shown that what we are doing really is substantial for the quality of our life and society in general. At ECP, we have very talented staff and we're doing everything we can do to help knock out this pandemic and be ready to prevent or at least mitigate the next one if it does come along.

Brueckner: Well, terrific, Doug. I want to really thank you once again for your time today.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA

612.812.5798

www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2020 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com or www.hpcuserforum.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.