

HPC User Forum Update

Interviews with Steering Committee Members: Simon Burbidge, University of Bristol

Steve Conway and Thomas Gerard
August 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 insideHPC Media. We welcome comments and questions addressed to Steve Conway, sconway@hyperionres.com or Earl Joseph, ejoseph@hyperionres.com.

This interview is with Simon Burbidge. Mr. Burbidge is the Director of the Advanced Computing Research Centre (ACRC) at the University of Bristol (UK), responsible for providing HPC, Research Data Storage and Research Software Engineering to this Russell Group University. Mr. Burbidge is a leader in the UK HPC field, participating in conferences, strategy and community groups. His career history includes leading the HPC service at Imperial College London, vendor experience and the seismic processing industry. As well as membership of the Hyperion HPC User Forum Steering Committee, Mr. Burbidge is an active participant in the UK HPC-SIG, HPE and ARM HPC communities and serves on academic committees in the UK and internationally.

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.

SIMON BURBIDGE INTERVIEWED BY DAN OLDS, insideHPC

Olds: Hello. Dan Olds here on behalf of Hyperion Research and insideHPC and I'm here with Simon Burbidge. We're going to talk about his storied career in HPC and what he sees for the future. How are you today, Simon?

Burbidge: I'm fine. Thanks, Dan. How are you?

Olds: I'm fine. So, how did you get involved with HPC in the first place?

Burbidge: Well, I think I've been involved in HPC, more or less, forever. When I was at school we started out and there was a computer hobby class that you could do when we were teenagers at the local technical college and there was an ICL computer there. We went over with our punch cards and did computing.

Olds: Starting with the cardboard.

Burbidge: I went on to do a degree in computer science and chemistry and then went into HPC straight from there at a central university computer center at University of London.

Olds: What were you doing there?

Burbidge: I started there in the QA team looking after system integration and system building. I carried on doing that for a few years, gaining more and more expertise on the CDC computers at that time and hybrid mainframes, then moving onto Cray systems.

Olds: Okay. What are you doing now?

Burbidge: So, I'm now working at the University of Bristol, which is in the west of England. I'm the director of Advanced Computing at the university, looking after the HPC computer services for the university.

Olds: Very nice. What have been some of the biggest changes that you've seen over the years in HPC?

Burbidge: I guess the biggest change was the migration from special supercomputers like the original Cray 1's and the Cray 2.

Olds: Sort of the hand-built supercomputers?

Burbidge: The hand-built ones, purpose-built, and the others that were like that, like the Convex and the MasPar and things like that, changing to use commodity microprocessors like the Intel processors, like the IBM power chips, like the SPARC chips. So, beginning the building of clusters rather than purpose-built supercomputers.

Olds: So, moving from supercomputers to dedicated Unix systems, which is kind of where my career started, but then moving into the land of entirely commodity for the most part?

Burbidge: Yeah, very much commodity at the moment. But I think the scalability of that is now becoming a challenge. We're seeing some of the architecture moving back to more dedicated and specific hardware for HPC, the advancement and the interconnects but also in the packaging, like the

cooling and the density and things like that, that are needed if you really want to scale up through the exascale to petascale and beyond.

Olds: Yes, very true. There are a lot of challenges inherent in that. A lot of people seem to think it's all on the software side, but it's on the hardware side as well.

Burbidge: It is. But I think it's always been on the software side.

Olds: Oh yes, predominantly.

Burbidge: You had to vectorize your Cray code to get the most out of it and people would spend a lot of time. Maybe that's something we've lost some of along the way. With standard processors and very reliable compilers, people just assume if their code runs it's running well. But that's not the case anymore and as you get more and more parallel, that's less and less the case. So, more work and more software engineering is required to make the most of the current computers.

Olds: So, you see hardware becoming more differentiated for HPC than it has been the last few years?

Burbidge: I think so, yes.

Olds: Okay. And, that would probably take in more exotic accelerators, FPGAs, things like that?

Burbidge: It certainly does, yes. Anything that we, as scientists, can take to make something go faster, we'll have a go at that. Anything that comes along, even if it is something that was invented to go into a teenage boy's bedroom to make the dynamics of computer games look better, we use that. That's been very successful. Let's see what's coming up from the vendors along similar or different lines and how we can use those.

Olds: Is there anything that particularly excites you in what's coming down the road in HPC? Or concerns you? Or both?

Burbidge: There are two things that excite me at the moment in the technology. The first is the advent of newer processors, in particular the ARM designed processors. We're involved in a lot of that work at Bristol, so it is of particular interest to me, and the research is here. But I think it does genuinely give some more competition to a classic Intel-based architecture.

Olds: Especially with that just being announced topping the TOP500 list at 415 petaflops. That's a big, big deal for ARM.

Burbidge: That's a big deal. And if you look at the numbers on that TOP500 score there is great improvement in the efficiency and performance of those machines. So, as they scale up, they will do better than one with a classical architecture.

Olds: Also, interestingly enough, that's an all CPU machine.

Burbidge: Exactly. So, maybe we don't need to have accelerators if you redesign your CPU's to have the capability of doing the amount of maths that you need and if you have, for example, the memory bandwidth to get those vectors and matrices in and out of the memory, then why wouldn't they be better than a GPU? They're both made of silicon. They're both made with the same kind of design, it's just how you organize that design on the chip.

Olds: But, you know, one of the systems that caught my eye was Nvidia's Selene system that most definitely is GPU-based. It is very highly efficient on a watts-per-flop basis.

Burbidge: Yes.

Olds: And that was interesting to me, given the propensity of GPUs to suck down an awful lot of electricity.

Burbidge: Well, they do. But any processor doing a lot of work is going to use a lot of power. It's kind of more about how efficiently it does that.

Olds: Exactly.

Burbidge: And I think previously, GPU-based accelerated systems have had GPUs and CPUs and the CPUs have really been a bit of an overhead on those systems. So, if you take out as many of the CPU components and host the GPUs in a much more optimized, cut down environment, then you can get those GPUs doing what they're good at and throw away the CPUs that you don't need in those systems.

Olds: And that's also, I think, a unique opportunity for ARM because, as you say, the CPUs on those systems are basically traffic cops. In these big systems. I think the Summit system at Oak Ridge, more than 90-some percent of the performance is coming from the GPUs.

Burbidge: Right.

Olds: Yes. Is there anything that has you concerned about the future in HPC?

Burbidge: I think one of the things we have to watch out for is energy consumption, total energy consumption and power efficiency. We all have a responsibility to keep our planet safe and the environment. Reducing the amount of energy we use to do computing would be a really good thing. So, if we could look at the efficiency, which is part of the nature of HPC, that would be important. It's always down to efficiency, speed, and performance. It's those same parallel, simultaneous equations that have been there for all those years. So it's just a matter of making one of those equations a bit more important than the other one and saying, "well, actually, the electricity is expensive and gee, it's not good for the environment. Let's tweak the equations a bit so we get more computing for less power."

Olds: Absolutely. And that's the way advances in computing have come, typically from HPC first.

Burbidge: I think so, yes. It's the leading edge. And some of the stuff we do never gets to see the light of day, but other things do. Things like using vectors to do faster calculations more effectively, that came from HPC in the beginning.

Olds: Sure. Well, thank you, Simon. This has been great. Really appreciate your time. Have a good rest of your day.

Burbidge: You, too.

Olds: Thank you very much. This has been another chat sponsored by Hyperion Research and InsideHPC. Thank you all for watching and listening.

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.