



The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes¹
August 4, 2020, 12-2 pm ET

Virtual

Participants

Alison Derbenwick Miller (Oracle)	Kathy Austin (TTU)
Anshu Dubey (ANL)	De Kaushik (UTA)
Bill Spotz (DOE)	Keith Beattie (LBL)
Dan Gunter (LBL)	Mallory Hinks (NCO)
David Martin (ANL)	Mike Heroux (Sandia National Lab)
Dhruv Chakravorty (Texas A&M)	Miron Livny (OSG)
Graham Lopez (NVIDIA)	Richard Carlson (DOE-SC)
H Birali Runesha (U of Chicago)	Sharon Broude Geva (U of Michigan)
Hal Finkel (DOE/SC)	Terence Sterba (GA Tech)
Ian Cosden (Princeton)	Terrill Frantz (Harrisburg Univ)
Jack Wells (NVIDIA)	Tevfik Kosar (NSF)
Jeff Larkin (NVIDIA)	Tim Costa (NVIDIA)
Jonathan Bentz (NVIDIA)	Valerie Taylor (ANL)
Kate Evans (ORNL)	

Introductions: This meeting was chaired by Richard Carlson (DOE/SC) and Tevfik Kosar (NSF)

Sustainable Software Speaker Series: Best practices in designing and developing sustainable research software

Toward Improved Sustainability of the Exascale Computing Project Software Stack

Mike Heroux, Senior Scientist at Sandia National Laboratories and Director of Software Technology for the US DOE Exascale Computing Project

- Noted that the reason to sustain is just as important as the how to sustain
- Mentioned as the “why” for the Exascale Computing Project: they work on products people use and on emerging platforms
- Described the key themes, software categories, and some example products
- SLATE port to AMD and Intel Platforms
 - SLATE is a distributed, GPU-accelerated, dense linear algebra library, intended to replace ScaLAPACK
 - Covers parallel BLAS, linear system solvers, least squares, eigensolvers, and the SVD

¹ Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.

- Accomplishments include:
 - Refactored SLATE to use BLAS++ as portability layer
 - Ported BLAS++ to AMD rocBLAS and Intel oneMKL
- Broader Community Engagement
 - Extreme-scale Scientific Software Stack (E4S) Forum – September 24, 2020
- How to sustain
 - Software architecture
 - Quality expectations
 - Access and understanding
 - Plan, execute, deliver
 - Better, faster, cheaper
- Extreme-scale Scientific Software Stack (E4S)
 - E4S: HPC Software Ecosystem – a curated software portfolio
 - Spack-based distribution of software tested for interoperability and portability to multiple architectures
 - Available from source, containers, cloud, binary caches
 - Not a commercial product – an open resource for all
- xSDK: Software Development Kits: Primary delivery mechanism for ECP math libraries continual advancements toward predictive science
- Delivering an open, hierarchical software ecosystem
 - More than a collection of individual production
 - ST Products > SDKs > E4S
- E4S DocPortal
 - Single point of access
 - All E4S products
 - Rendered daily from repos
- ECP ST Planning Process: Hierarchical, 3-phase, cyclical
- ECP ST Lifecycle summary

Q&A

- Dhruv Chakravorty (Texas A&M): Mentioned that they are heavily invested in EasyBuild at Texas A&M. He asked how Mike saw the path forward for them.
 - Mike said he thinks there may be some opportunities to look at leveraging the two different build systems and maybe some sort of migration path.
 - Mike mentioned that one of the things that has been hampering them has been a lack of support for Windows. He said that's one of the things they're looking at trying to establish.
 - Mike said he didn't want to recommend dropping EasyBuild. He said Spack is an up and coming meta build tool and there's a tremendous investment in it within the HPC community, so he thinks it's worth tracking.
 - Dhruv said they would be very interested in advice for migrating.

Maintainable Software through Standard Language Parallelism

Jeff Larkin, Principal HPC Application Architect, NVIDIA

- Mentioned that NVIDIA is best known for GPUs, but they also have an interconnect division
- Jeff said that NVIDIA is a platform company, and they need to provide software and ways for programmers to address concerns
- They start with a foundation of accelerated libraries (core libraries).

- They believe that the best place for most programmers is in the accelerated standard languages
- Noted that if your goal is absolute performance, platform specialization is an option
- Future of programming the NVIDIA platform:
 - Accelerated standard languages will make up a larger portion
 - ISO C++
 - ISO Fortran
 - Python
- Shared a code snippet from a code called Lulesh, which is out of Lawrence Livermore National Lab
 - Mini app
 - Worked with the developers to support standard C++ using parallel execution policies
 - Showed a simplified code that was the result
- Shared an example of a full app: STLBM – many-core Lattice Boltzmann with C++ Parallel Algorithms
- Gave an example for ISO Fortran
- Legate: Automatic python acceleration and scalability
 - Product still in early stages, but they've made an alpha release available on GitHub
- Future of maintainable parallel programming
 - Parallel-first programs using standard programming languages

Q&A

- Anshu Dubey: Where all of these systems have a limitation and therefore lack adoption by real multi-physics applications that run on large platforms is that parallelism needs to be hierarchical. A simple concurrent kind of parallelism doesn't cut it because as soon as you have heterogeneity within the application all of these solutions become impossible to use. Have you given any thought to that aspect of HPC workloads?
 - Jeff: Yes, and the place where we are working most closely on this is in C++ today. We're working to define ways to specify hierarchical parallelism and to support more task-based approaches through C++ executors.
 - Anshu: How would that be different from other solutions?
 - Jeff: The primary benefit is that it is in the standard language so you can rely on it being available and supported by multiple vendors and multiple compilers.
 - Mike Heroux: I want to emphasize that this is an ongoing process. The objective is to get into the language standards and to get into community ecosystems. But then it's also a pipeline where the front end doesn't dry up. You continue to innovate, you continue to adapt, you continue to deliver to these standard ecosystems.
 - Anshu: I do appreciate that you included Fortran in this because it gets left out of conversations mostly.

Discussion

- Tefvik asked the group what other aspects of designing and developing sustainable software should we be discussing or focusing on?
- Anshu said software design is a critical issue which is not given enough attention and focus.
- Dan Gunter said he would agree with Anshu. The process of learning what the software needs to do is very important and sometimes skipped over as obvious.
- Mike Heroux said public published APIs are very important. The more we can get published APIs and multiple implementations underneath it, the easier it is to sustain software.
- Mike Heroux said long-term sustained funding
- Tefvik said that is already on their list

- Valerie Taylor said it would be good to also hear beyond HPC software, maybe from experimental facilities in terms of that software space and some of the needs there.
- Anshu seconded Valerie's suggestion. There has been a group of people who have been informally talking about the need for software. And in that process, they've talked to a lot of people from facilities and they discovered that they consume as much software as we do, but their challenges are very different.
- Dhruv Chakravorty mentioned that this is something he is personally facing. He noted that he his software developers were hired away by Google, Nintendo, and Sonos. Now there is a gap where the new students have to come and learn what they did. The other challenge is that when students moved to the private landscape they signed NDAs and they are worried that by helping their old project, by just explaining what they did and what is the current way of doing things, may be a violation.
- Anshu said that another thing that is critical is that a lot of the information about the software is in people's head. How do you sensibly get it out of their heads and onto some form of documentation that is useful for the rest of the team.
- Jeff Larkin pointed out that there will always be a flow of students. He said that part of what makes what they are advocating for really attractive is that students are not needing to learn new frameworks or new abstractions.
- Anshu said that the entire software infrastructure that underlies a lot of science in today's world has reached a level of sophistication and maturity that it's time we started thinking of it as a scientific instrument and treated it as something that needs to be properly invested in.
- Valerie said that a good discussion for this group could be metrics for sustainability.
- Jeff said that you have to treat simulation as core to the scientific process. It's just like buying a new instrument. Developing a strong source of simulation software is often critical in this process.
- Anshu commented that because the software is developed over a period of time by people, the amount of investment is not immediately obvious in the way that the investment is obvious in a machine. But it usually ends up being comparable. In the past, many such investments have been allowed to drop by the wayside, which is counterproductive.
- Rich said that several months ago, we had a discussion in this group on the various kinds of presentations and topics we would like to see. Email Mallory if you would like to add anything to the list (topics and speakers).
- Mike Heroux said that a bigger picture topic is the business model, which includes funding, but it's a sensibility around what is the beginning middle and sustainable delivery of capabilities.
- Miron Livny questioned the idea of a business model for federally funded projects like Large Hadron Collider or telescopes.
- Miron also said that there is more of a culture of how to retain people in biological labs. He said we are not paying enough attention to what will make an individual stay and work on software for decade. He said he doesn't believe that it is a question of money. Sometimes it is intellectual challenge.
- Mike Heroux suggested that Miron may have interpreted his mention of a business model differently that was intended. He said that a business model is simply an understanding of where your efforts fit within the broader ecosystem of activities and making sure that as you do your work it is the most important work. And if at some point it's better for that work to be done by someone else, perhaps in industry, that you figure out a way to transfer that activity to them so that you can continue to do leading edge work.

- Mike said that he agreed with Miron regarding investing. It's often the case that the way we keep people and the way we sustain our scientific teams is to make sure they understand the mission and what their life is about as a scientist and have them buy into that.
- Anshu commented that the question of people leaving is not as simple as money vs intellectual challenge and scientific interest. There are many things that come into play, including the fact that the value of software developers is not recognized. People start to think that if I'm being thought of as a mere technician and making a small amount of money, what's the point? Let me at least do similar work and make more money. We need to garner respect for the work people do in these kinds of software.
- Dhruv added that he tried to hire his students, but the overwhelming reason they left was to guarantee H1B sponsorships.
- Anshu said that in the academic environment, there is no permanent career path for somebody who'd not doing fundamental research.
- Alison Derbenwick Miller said that one of the things that they've seen in working with researchers is that software is often treated in research labs as a disposable asset. We don't see a lot of imposition of best practices about software development go on in these labs. It might be beneficial for the community to think through if we were to have some of those best practices documented and imposed. Would that help with the value perceptions of the work?
- Mike Heroux said that with the Exascale project, these policies have elevated the skills of the people who noted this kind of work. It's very important. There's a cause and effect there that we shouldn't ignore.
- Miron asked how do mobilize somebody to define this as a grand challenge.
- Sharon Broude Geva said there is a huge divide between faculty and staff. Orchid has now started expanding the credit that you can give. Anyone who has the power to do that when they have a project should be doing it. That includes all of the people who are considered invisible and ancillary. She suggested having the agencies stress that and maybe require that as a part of funding might be very useful.
- Rich said if we need more formal workshops or some other venue to make sure that this issue is raised to an awareness inside the federal government and the agencies, DOE is certainly open to doing something. Co-sponsoring with NSF would be even better.
- Sharon shared a link to the Orchid giving credit where credit is due page: <https://info.orcid.org/credit-for-research-contribution/>

Round Table

- Mike Heroux pitched a workshop that they are planning for the second week of December on the science of scientific software. RSE-HPC-2021 Workshop at SC21 <https://us-rse.org/2021-07-16-SC21-workshop-cfp/>
- Hal Finkel pointed out that they had their federal advisory committee meeting last week and Barb Helland shared some slides that outline their Software Stewardship Task Force is working on in this regard. He shared the slides so everyone could see. He said the slides are available on their website.
 - Mentioned that the taskforce is planning to put out an RFI, so he will keep everyone updated on that.

Next Meeting

September 1 (12 pm ET)