Software in the Era of Extreme Heterogeneity

Joint virtual workshop of the NITRD HEC & SPSQ Interagency Working Groups

September 22 – 24, 2020

Agenda

Day I: September 22

(All times are in ET)

11:00 AM – 11:15 AM	**Welcome & Opening Remarks**	HEC/SPSQ Co-Chairs, Kamie Roberts, NITRD
11:15 AM – 11:55 AM	**Invited Speaker**	Bill Gropp, NCSA
11:55 AM – 12: 55 PM	**Session I Panel: Software Development and Sustainment Challenges in Extreme Heterogeneity** Objective: Software development and sustainment continue to be a challenge that prevents the community from taking advantage of existing and emerging processing capabilities and increasing hardware heterogeneity will only exacerbate this challenge. This session will explore how software can become more responsible for software development and sustainment decisions and efforts; what decisions and efforts can be automated and what must be performed by humans; and what opportunities and challenges lay in this space in the era of extreme heterogeneity.	Moderator: Ron Boisvert, NIST Panelists: Alex Aiken, Stanford Ira Baxter, Semantic Designs Torsten Hoefler, ETH Zurich Denys Poshyvanyk, William & Mary
12:55 PM – 1:10 PM	**OSTP Update**	Manish Parashar, OSTP
1:10 PM – 1:25 PM	**Break**	
1:25 PM – 2:25 PM	**Session I Breakout Discussions**	
	Breakout 1A: Mitigating and Accommodating the Impact of Extreme Heterogeneity on HEC Software and Sustainment	Moderator: Chris Rackauckas, MIT
	Breakout 1B: Mitigating and Accommodating the Impact of Extreme Heterogeneity on HEC Software and Sustainment	Moderator: Tal Ben-Nun, ETH Zurich
	Breakout 2A: Impact of Extreme Heterogeneity on HEC Software and on Its Development and Sustainment	Moderator: Brad Chamberlain, HPE
	Breakout 2B: Impact of Extreme Heterogeneity on HEC Software and on Its Development and Sustainment	Moderator: Hyesoon Kim, GA Tech

	Breakout 3A: HEC Software Development Decisions Humans Must Make and What They Should Know to Make Them	Moderator: Rob Schreiber, Cerebras
	Breakout 3B: HEC Software Development Decisions Humans Must Make and What They Should Know to Make Them	Moderator: Shinjae Yoo, BNL
	Breakout 4A: Automating HEC Development Decisions	Moderator: Shuvra Bhattacharyya, UMD
	Breakout 4B: Automating HEC Development Decisions	Moderator: Frederik Kjolstad, Stanford
2:25 PM – 3:10 PM	**Break/Debrief Summary Preparation Time**	
3:10 PM – 4:10 PM	**Debrief of Session I Breakout Discussions (4 min/breakout reps, Q&A remaining time)**	Moderator: Frances Hill, DoD
	Breakout 1A, 1B Reps: Jim Kirby (NRL), Tal Ben-Nun (ETH Zurich) Breakout 2A, 2B Reps: Brad Chamberlain (HPE), Hyseoon Kim (GA Tech) Breakout 3A, 3B Reps: Rob Schreiber (Cerebras), Shinjae Yoo (BNL) Breakout 4A, 4B Reps: Shuvra Bhattacharyya (UMD), Frederik Kjolstad (Stanford)	
4:10 PM – 4:15 PM	**Day I Wrap Up**	Walid Keyrouz, NIST

Day II: September 23

(All times are in ET)

11:00 AM – 11:15 AM	**Welcome and Opening Remarks**	Robinson Pino, DOE/SC
11:15 AM – 11:55 AM	**Invited Speaker**	Hal Finkel, ANL
11:55 AM – 12: 55 PM	**Session II Panel: Productivity** Objective: This session examines what is scientific computing productivity, performance, and reproducibility and how it can be delineated and augmented in the emerging heterogeneous computing environments. It will also examine the impact of heterogeneity on scientific discovery.	Moderator: Jim Ang, PNNL Panelists: Saman Amarasinghe, MIT Katie Antypas, LBNL Ron Brightwell, SNL Michael Garland, NVidia Jeff Vetter, ORNL
12:55 PM – 1:25 PM	**Break**	
1:25 PM – 2:25 PM	**Session II Breakout Discussions** Breakout 1A: Productivity Challenges and Opportunities in Emerging Heterogeneous Environments, Including Performance, Reproducibility, and Portability	Moderator: Pat McCormick, ANL

	Breakout 1B: Productivity Challenges and Opportunities in Emerging Heterogeneous Environments, Including Performance, Reproducibility, and Portability	Moderator: Catherine Schuman, ORNL
	Breakout 2A: Modeling and Simulation for Predicting Heterogeneous System Performance	Moderator: Maya Gokhale, LLNL
	Breakout 2B: Modeling and Simulation for Predicting Heterogeneous System Performance	Moderator: Richard Lethin, Reservoir Labs
	Breakout 3A: Software Infrastructure and Sustainability for Heterogeneous Computing	Moderator: Jeff Vetter, ORNL
	Breakout 3B: Software Infrastructure and Sustainability for Heterogeneous Computing	Moderator: Katie Antypas, LBNL
	Breakout 4A: Maintaining and Improving Programmer Productivity in Heterogeneous Systems	Moderator: John Shalf, LBNL
	Breakout 4B: Maintaining and Improving Programmer Productivity in Heterogeneous Systems	Pete Beckman, ANL
2:25 PM – 3:10 PM	**Break/Debrief Summary Preparation Time**	
3:10 PM – 4:10 PM	**Debrief of Session II Breakout Discussions (4 min/breakout reps, Q&A remaining time)**	Moderator: Sonia Sachs, DOE/SC
	Breakout 1A, 1B Reps: Pat McCormick (ANL), Catherine Schuman (ORNL) Breakout 2A, 2B Reps: Maya Gokhale (LLNL), Richard Lethin (Reservoir Labs) Breakout 3A, 3B Reps: Jeff Vetter (ORNL), Katie Antypas (LBNL) Breakout 4A, 4B Reps: John Shalf (LBNL), Pete Beckman (ANL)	
4:10 PM – 4:15 PM	**Day II Wrap Up**	Jim Kirby, NRL

Day III: September 24

(All times are in ET)

11:00 AM – 11:15 AM	**Welcome and Opening Remarks**	Alan Sussman, NSF
11:15 AM – 11:55 AM	**Invited Speaker**	Mary Hall, U of Utah
11:55 AM – 12: 55 PM	**Session III Panel: Workforce and Tools**	
	Objective: Supporting software on high-end systems is already a challenging task, and preparing the workforce to support tools and applications on extremely heterogeneous systems	Moderator: Vivek Sarkar, GA Tech Panelists: Rudi Eigenmann, U of Delaware

	is an even greater challenge. This session explores workforce requirements to support the software and what is needed to reduce the human challenge of software development, evolution, and porting. It will also explore what is needed to prepare people to develop, effectively use, and support software in the era of extreme heterogeneity.	Martin Berzins, U of Utah Bill Magro, Google David Keyes, KAUST
12:55 PM – 1:25 PM	**Break**	
1:25 PM – 2:25 PM	**Session III Breakout Discussions**	
	Breakout 1A: Creating a Flexible, Robust, and Sustainable Software Infrastructure for 21st Century Science - Long Term Challenges	Moderator: Mike Heroux, SNL
	Breakout 1B: Creating a Flexible, Robust, and Sustainable Software Infrastructure for 21st Century Science - Long Term Challenges	Moderator: Anshu Dubey, ANL
	Breakout 2A: Workforce Challenges in the Era of Extreme Heterogeneity - Short Term Responses	Moderator: Lois Curfman McInnes, ANL
	Breakout 2B: Workforce Challenges in the Era of Extreme Heterogeneity - Short Term Responses	Moderator: Kelly Gaither, TACC
	Breakout 3A: New Visions for Training the Next Generation of Cyberinfrastructure Professionals - Solving the People Problem for the Future	Moderator: Henry Neeman, U of Oklahoma
	Breakout 3B: New Visions for Training the Next Generation of Cyberinfrastructure Professionals - Solving the People Problem for the Future	Moderator: Ewa Deelman, USC/ISI
	Breakout 4A: Using AI and Machine Learning to Increase Human Productivity in Dealing with Accelerating Complexity and Rates of Change	Moderator: Andrew Chien, U Chicago/ANL
	Breakout 4B: Using AI and Machine Learning to Increase Human Productivity in Dealing with Accelerating Complexity and Rates of Change	Rajiv Ramnath, Ohio State U
2:25 PM – 3:10 PM	**Break/Debrief Summary Preparation Time**	
3:10 PM – 4:10 PM	**Debrief of Session III Breakout Discussions (4 min/breakout reps, Q&A remaining time)**	Moderator: Larry Kaplan, HPE
	Breakout 1A, 1B Reps: Mike Heroux (SNL), Anshu Dubey (ANL) Breakout 2A, 2B Reps: Lois Curfman McInnes (ANL), Kelly Gaither (TACC) Breakout 3A, 3B Reps: Henry Neeman (U of Oklahoma), Ewa Deelman (USC/ISI)	

Software in the Era of Extreme Heterogeneity

Joint virtual workshop of the NITRD HEC & SPSQ Interagency Working Groups

September 22 – 24, 2020

Breakout 4A, 4B Reps: Andrew Chien (U Chicago/ANL), Rajiv Ramnath (Ohio State U)

4:10 PM – 4:15 PM **Day III Wrap Up** Barry Schneider, NIST

Session I Breakout Discussions

- Breakout 1A/B: Mitigating and Accommodating the Impact of Extreme Heterogeneity on HEC Software Development and Sustainment

  Questions:
  - What aspects of the negative impact of extreme heterogeneity on HEC software development and sustainment are amenable to being addressed by new ideas, techniques, abstractions, and technologies?
  - What are those new ideas, techniques, abstractions, and technologies? Are new, different abstractions called for?
  - How can ML contribute?
  - How can tools and automatic code generation contribute?

- Breakout 2A/2B: Impact of Extreme Heterogeneity on HEC Software and on Its Development and Sustainment

  Questions:
  - What is the impact of extreme heterogeneity on the software that must interact with it? In particular, what exactly do we mean by computational reproducibility and accuracy under extreme heterogeneity and how does software itself meet these requirements?
  - How does heterogeneity-driven change to the software affect its development and sustainment?
  - How does heterogeneity-driven change to the software affect what is required of software developers, what they need to know, how they should best be trained?
  - How does heterogeneity-driven change to the software affect the role of software tools, what is required of them?
  - What will be the impact of evolving heterogeneity over the next 5 to 10 years on HEC software development and sustainment?

- Breakout 3A/3B: HEC Software Development Decisions Humans Must Make and What They Should Know to Make Them

  Questions:
  - What is the nature of software development decisions that humans must make? What education, training, experience, tools, abstractions, and technology is required by those who make the decisions?
  - How does extreme heterogeneity affect those decisions and what is required of those who make them? In particular, how are the computational reproducibility and accuracy of software, which are properties of the results of software execution, affected?
  - How are computational reproducibility and accuracy of software affected by human decision making and automation?

- Breakout 4A/4B: Automating HEC Development Decisions

  Questions:
  - Which of these decisions do we now know how to automate? Automating which of the decisions will require major research investments?
  - How might software reproducibility and accuracy be affected?
  - How are computational reproducibility and accuracy of software affected by human decision making and automation?

Session II Breakout Discussions

- Breakout 1A/1B: Productivity Challenges and Opportunities in Emerging heterogeneous Environments, Including Performance, Reproducibility, and Portability

  Questions:
  - In an era of increasingly diverse and complex computing environments, what advances in programming models, environments, and tools are required to improve the productivity of a broad range of scientific software developers?
  - What demands applications are developers facing for improved usability and productivity? Performance, Portability, and Productivity
- Breakout 2A/2B: Modeling and Simulation for Predicting Heterogeneous System Performance

  Questions:
  - Can advanced modeling and simulation efficiently and accurately predict the performance characteristics of applications running on heterogeneous technologies? The applications of this include (a) supporting the design of future systems (b) supporting static (e.g., autotuning) and dynamic (runtime) optimization processes (c) support verification of correctness (d) accurately predicting potential deadlock or livelock due to resource limitations.
  - To what extent can analytical models offer fast evaluation of the design parameter space? Are there robust methods for validating models with simulations?
  - How can modeling and simulation aid evaluation of application performance portability? of reproducibility?
- Breakout 3A/3B: Software Infrastructure and Sustainability for Heterogeneous Computing

  Questions:
  - What software infrastructure and tools will be necessary to achieve usable, productive, and sustainable scientific workflow across multiple, different and increasingly heterogeneous computing environments?
- Breakout 4A/4B: Maintaining and Improving Programmer Productivity in Heterogeneous Systems
  - Is it possible to achieve flexible, expressive, programming models and languages; intelligent, domain-aware compilers and software development tools; and composition of disparate software component content in future heterogeneous computing systems?

Software in the Era of Extreme Heterogeneity

Joint virtual workshop of the NITRD HEC & SPSQ Interagency Working Groups

September 22 – 24, 2020

- How can software supporting the future computing ecosystem balance the following: efficiency of development, debugging, verification, and validation; usability, reproducibility, manageability, extensibility, and sustainability; and performance and scalability?

Session III Breakout Discussions

- ### Breakout 1A/1B: Creating A Flexible, Robust, and Sustainable Software Infrastructure for 21st Century Science - Long Term Challenges

  Questions:

  - The current academic climate does not recognize software development as an important part of the scientific research and development cycle. For example, the need for professional software developers in many large-scale projects is well documented but there is essentially no real mechanism in academia to ensure that these individuals have permanent jobs. They rely on grants for support with not much direct help from the university. Given that they play an essential role in the science, how can that culture be changed?

- ### Breakout 2A/2B: Workforce Challenges in the Era of Extreme Heterogeneity - Short Term Responses

  Questions:

  - How do we incorporate greater flexibility into the workforce so people can transition from other jobs into CI professionals?
  - The learning curve for CI professionals is currently too steep to expect someone to have all the necessary skills when they start their positions. How can we do better to change the culture so that existing individuals interested in transitioning into CI professionals can do so with less stress?
  - What kinds of training and development resources, both online and in person, are needed to make such transitions easier?

- ### Breakout 3A/3B: New Visions for Training the Next Generation of Cyberinfrastructure Professionals - Solving the People Problem for the Future

  Questions:

  - While professional software developers are needed for many scientific projects, training, career paths, and long-term sustainability at universities and other institutions are lacking. Given that CI professionals play an essential role in the science, what can be done to better establish their place in the scientific research community?
  - Are there non-traditional sources for producing and training professional software developers with the necessary skills?

- ### Breakout 4A/4B: Using AI and Machine Learning to Increase Human Productivity in Dealing with Accelerating Complexity and Rates of Change

  Questions:

  - What can we do that would reduce the reliance on humans to write software?

For example, can deep learning and AI, or automatic code generation, help in developing software and should more effort be expended on ensuring that to happen?

- Can we build systems that will substantially reduce the efforts of people to deal with architectural changes and produce longer lived software products? The existing situation is forcing us to rewrite software on relatively short timescales.
- Since software persists much longer than hardware, can we develop tools that substantially reduce the burden on humans.