Giving Research Software Engineers a Larger Stage Through the Better Scientific Software Fellowship

William F. Godoy, *Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*
Ritu Arora, *University of Texas at San Antonio, San Antonio, TX, 78249, USA*
Keith Beattie, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*
David E. Bernholdt, *Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*
Sarah E. Bratt, *University of Arizona, Tucson, AZ, 85721, USA*
Daniel S. Katz, *University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*
Ignacio Laguna, *Lawrence Livermore National Laboratory, Livermore, CA, 94551, USA*
Amiya K. Maji, *Purdue University, West Lafayette, IN, 47906, USA*
Addi Malviya-Thakur, *Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*
Rafael M. Mudafort, *National Renewable Energy Laboratory, Golden, CO, 80401, USA*
Nitin Sukhija, *Slippery Rock University of Pennsylvania, Slippery Rock, PA, 16057, USA*
Damian Rouson, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*
Cindy Rubio-González, *University of California, Davis, Davis, CA, 95616, USA*
Karan Vahi, *University of Southern California, Marina Del Rey, CA, 90292, USA*

The Better Scientific Software Fellowship (BSSwF) was launched in 2018 to foster and promote practices, processes, and tools to improve developer productivity and software sustainability of scientific codes. The BSSwF's vision is to grow the community with practitioners, leaders, mentors, and consultants to increase the visibility of scientific software. Over the last five years, many fellowship recipients and honorable mentions have identified as research software engineers (RSEs). Case studies from several of the program's participants illustrate the diverse ways the BSSwF has benefited both the RSE and scientific communities. In an environment where the contributions of RSEs are too often undervalued, we believe that programs such as the BSSwF can help recognize and encourage community members to step outside of their regular commitments and expand on their work, collaborations, and ideas for a larger audience.

In 2017, the U.S. Exascale Computing Project (ECP)[a] funded the IDEAS Productivity project[b] to support developer productivity and software sustainability and reliability for the ECP's extensive scientific software and application portfolio, including more than 100 software packages and 24 applications. The ECP encompasses roughly 1000 people across the U.S. Department of Energy's (DOE's) national laboratories, academia, and industry, including both research software engineers (RSEs) and researchers from a broad range of scientific domains as well as computer science and applied mathematics.

Given the scale and breadth of the challenge, the IDEAS project includes a wide array of "outreach"

[a]https://www.exascaleproject.org
[b]The original project title is Interoperable Design of Extreme-Scale Application Software.

1521-9615 © 2023 IEEE
Digital Object Identifier 10.1109/MCSE.2023.3253847
Date of publication 9 March 2023; date of current version 1 May 2023.

activities to complement the relatively modest level of direct engagement that the IDEAS team can sustain with ECP software teams.[1] IDEAS outreach supports the ECP project in general as well as the larger community of scientific software developers, which constitute the workforce pool that ECP and the national laboratories (in particular) draw upon, developers of dependencies of ECP software, and often additional consumers of ECP software. The three basic thrusts of the IDEAS outreach effort include 1) an online portal for the exchange of information and resources on scientific software development,[c] 2) providing seminars[d] and training[e] to help developers improve their skills, and 3) awareness and discussion of software-related issues in the community. The third element includes a great deal of event organization: creating opportunities within technical meetings for people to discuss their experiences with software development[f] as opposed to the typical focus on the *scientific* advances. But we also look for ways to support members of the community who focus on the software side of things (who often identify as RSEs), and to amplify voices in the community.

Many aspects of the IDEAS project were inspired by the work of the Software Sustainability Institute (SSI)[g] in the United Kingdom. The SSI has been sponsoring a Fellowship Program since 2012, with the goal to "improve and promote good computational practice across all research disciplines and support those who are doing this important work."[h] In the fall of 2017, the IDEAS leadership began to formulate a fellowship program of their own, bringing new ideas and new voices to the DOE/ECP scientific software community, and in 2018, the first class of the Better Scientific Software Fellowship (BSSwF) was announced.

The BSSwF is somewhat different in structure and goals from the SSI Fellowship. The BSSwF program "fosters and promotes practices, processes, and tools to improve developer productivity and software sustainability of scientific codes."[i] And, recognizing differences in how researchers are supported between the United Kingdom and the United States, the BSSwF supports fewer fellows with larger stipends. The fellows can pursue activities through other funding sources that may be available to them to achieve meaningful results. Each class of the BSSwF includes both fellows and honorable mentions. Honorable mentions do not receive funding but are included in many of the networking opportunities arranged by the IDEAS project to help both fellows and honorable mentions get more exposure to the DOE, ECP, and the National Science Foundation (NSF) software communities, and identify opportunities for collaboration. We encourage applications from people at all career stages, and awards have been made to graduate students, early career, midcareer, and more senior levels. Likewise, awards have been made to individuals in academia, national laboratories, and small businesses. Honorable mentions are eligible to reapply, but fellows can receive the award only once. Due to the nature of fellowship funding, the primary limitation on applicants is that they must be associated with a U.S.-based institution capable of receiving federal funds.

THE BSSwF PROGRAM "FOSTERS AND PROMOTES PRACTICES, PROCESSES, AND TOOLS TO IMPROVE DEVELOPER PRODUCTIVITY AND SOFTWARE SUSTAINABILITY OF SCIENTIFIC CODES."

The size of the fellowship classes has varied over the years based primarily on the funding available. We started with four fellows and four honorable mentions in each of the 2018 and 2019 classes, while in 2020 we supported three sets. In 2021, the NSF joined the DOE as a cosponsor of the fellowship, allowing a class of four and then six sets in 2022.

This article is coauthored by eight fellows, five honorable mentions, and one member of the fellowship's Executive Committee to illustrate the impact of the BSSwF to develop, connect, and grow a community that recognizes efforts toward better scientific software. In the remainder of the article, each recipient provides a personal account on his or her experiences around the BSSwF, we then highlight the synergies and lessons learned resulting from the fellowship activities. Finally, we outline our shared view advocating for more initiatives like the BSSwF to elevate, reward, and promote RSE stewardship efforts on mission-critical software.

[c]https://bssw.io
[d]https://ideas-productivity.org/events/hpc-best-practices-webinars
[e]https://bssw-tutorial.github.io/
[f]https://betterscientificsoftware.github.io/swe-cse-bof/
[g]https://software.ac.uk
[h]https://software.ac.uk/programmes-and-events/fellowship-programme
[i]https://bssw.io/pages/bssw-fellowship-program

EXPERIENCES

Due to the individual nature of the award, the goal of this section is to showcase the recipients' personal experiences. Each account illustrates the diverse drivers, interests, and efforts of people from different backgrounds and career stages contributing toward building a better scientific software community that elevates the work of RSEs.

FELLOWS

Daniel S. Katz

At the time of his fellowship (2018), Katz led the scientific software and applications division at National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign. He was concerned about career paths and recognition for RSEs. The BSSwF allowed him to fund and focus on activities that would recognize software as a valid scientific metric through the software citation principles initiative.[2] He co-led the FORCE11 Software Citation Implementation Working Group alongside astronomers, physicists, and geoscientists to produce the principles and promote practices for software citation to recognize the RSE community behind these efforts. The more significant outcome from Katz's fellowship experience is that software citation is now becoming widely accepted in several major conference venues, while software developers and maintainers are encouraged to provide citable sources for their work. The FORCE11 Group has continued this work (with Katz still co-leading it) as well as focusing on other groups, such as publishers, funders, and repositories and registries, to institutionalize citation practices among scientific software stakeholders.

Ignacio Laguna

As a part of his fellowship (2019), Laguna organized a "Tutorial on Floating-Point Analysis Tools for Scientific Software"[j] at venues such as IEEE/ACM Supercomputing (SC) and the Practice and Experience in Advanced Research Computing (PEARC) conferences. The tutorial's goal was to demo tools that allow programmers to gain insight about how different aspects of floating-point arithmetic affect their code and how to fix potential bugs. It focused on three specific areas of floating-point analysis: 1) error bounds and re-evaluation of expressions to reduce the error, 2) compiler-induced error due to optimizations, and 3) floating-point exceptions in accelerators, e.g., GPUs. Developing accurate and reliable scientific software is notoriously difficult, and efforts such as Laguna's enable researchers to ensure numerical reproducibility. The latter becomes an even bigger challenge with the heterogeneity of high-performance computing (HPC) systems, thus posing a very difficult problem going from compiler optimizations to different precision arithmetic, which can significantly affect the final numerical results. The BSSwF helped him expand his network and made his tools more visible to HPC developers. As a fellow, Laguna participated at several ECP annual meetings, which provided a venue to meet new people, socialize his technical ideas in the DOE HPC community, and come up with novel research ideas to address important problems.

[j]https://fpanalysistools.org

Damian Rouson

At the time of his fellowship (2020), Rouson worked part time for the Sustainable Horizons Institute (SHI), where he collaborated with SHI President Mary Ann Leung (also a BSSwF recipient), and archaeologic engineer Brad Richardson on a series of workshops that introduced agile software development training. A novel aspect of these workshops was that there was no programming language prerequisite. Instead, the focus was on the collaborative aspects of source control using Git and GitHub workflows, including pull requests and continuous integration using a simple proxy project, such as markdown document creation. The participants were also exposed to concepts like pair programming, unit testing, and test-driven development. One of the important goals of Rouson's fellowship project was to attract students and scientists across a broad range of disciplines, programming backgrounds, career stages, and demographics, with an emphasis on reaching out to groups that are underrepresented in computational science and engineering. The BSSwF enabled Rouson to contribute to lowering the barriers to entry into projects related to scientific software. During this experience, he discovered that most of what one needs to learn to employ agile practices effectively does not require doing software development.

Cindy Rubio-González

During her fellowship (2020), Rubio-González focused on disseminating techniques for automated floating point mixed-precision tuning and software testing of numerical software. Mixed precision balances the tradeoff between accuracy and cost of full numerical representations, as used heavily by scientific applications. As a part of her BSSwF work, Rubio-González presented state-of-the-art efforts at various venues and developed a graduate seminar that focuses on the use

and evaluation of techniques that improve the reliability and performance of numerical software, with an emphasis on the reproducibility of experiments affected by numerical precision. The fellowship provided Rubio-González with a unique opportunity to connect, share, and learn from others who are also passionate about building better scientific software. Being a part of the BSSw community has inspired Rubio-González to continue her work designing and developing tools that target a significant topic in scientific computing: writing better, more efficient, and reproducible scientific software.

Ritu Arora

During her fellowship (2022), Arora developed training material on how to optimize input–output (I/O) in scientific applications targeting artificial intelligence (AI)/machine learning (ML) workflows. The topics covered include 1) optimizing I/O for data analysis and checkpointing in serial and parallel applications written in C, C++, Fortran, Python, R, message passing interface, OpenMPI, and CUDA; 2) optimizing I/O and checkpointing AI/ML models and workflows; and 3) techniques for leveraging the features in the underlying hardware and filesystems (e.g., Lustre) for optimizing applications' I/O while being aware of portability issues. The material will be distributed through accessible online resources such as YouTube, GitHub-hosted repositories, LinkedIn, and bssw.io blog articles. The BSSwF program has been immensely valuable for Arora for expanding her network of colleagues, especially in DOE laboratories, and exploring opportunities for collaboration. The program has provided visibility to the cause that she is very passionate about, which is "optimizing I/O" from serial and parallel applications to efficiently take advantage of the underlying HPC hardware as more data-driven AI/ML drive scientific discovery.

Nitin Sukhija

During his fellowship (2022), Sukhija organized a tutorial on "Best Practices and Tools for Secure Scientific Software Development" at the Society for Industrial and Applied Mathematics Conference of Mathematical Data Science. This tutorial includes components for evaluating design practices and processes for creating and managing secure software, threat modeling, and quality-assurance testing using both static and dynamic analysis tools. Participants from industry and academia learned techniques to mitigate threats to prevent losing business and sensitive information due to a variety of potential vulnerabilities from unsecured scientific software. The materials were distributed to the community at large via open source collaboration platforms. The fellowship provided Sukhija with an opportunity to collaborate with distinguished researchers through the ECP project. These interactions helped him and his students learn more about HPC and cybersecurity challenges, which are available in his department/center at the Slippery Rock University of Pennsylvania.

Amiya K. Maji

During his fellowship (2022), Maji worked to simplify scientific Python package installation by streamlining environment management, dependency tracking, and runtime customizations through easy-to-use tools developed for scientific end users. Managing Python applications is especially challenging as packages can be installed via multiple package managers and have incomplete dependencies, and there is a growing need for providing consistency across traditional batch workloads and interactive notebooks in HPC environments. He implemented best practices for package management in the development of the `python-env-mod` tool. `Python-env-mod` helps users manage their Python environments more efficiently and load runtime configurations through the familiar abstraction of environment modules. His work improved scientific productivity by automating the process of environment management to avoid user errors. The fellowship provided Maji with the ability to network with HPC scientists who provided early feedback and suggestions about `python-env-mod`. In addition, he was able to hire and mentor undergraduate students about the HPC software ecosystem.

Karan Vahi

During Vahi's fellowship (2022), he developed training materials that examine the workflow lifecycle and challenges associated with the interdependency of various steps, such as creation, execution, monitoring, and debugging in simulation and data analysis pipelines. The overall goal is to bring the use of workflows to the wider scientific community. His work walks users through the process for how to model existing simulation pipelines into workflows, package application code in containers, and execute the workflow on HPC resources and distributed computing infrastructure such as Open Science Grid. He used interactive Jupyter notebooks to guide users on how to develop workflows that leverage application containers for running jobs using Pegasus.[3] These notebooks can be used as self-guided material, in the classroom, or as virtual training material. Dissemination of the work targeting a scientific audience was done as tutorials at PEARC 2022 and at the 2022 IEEE 18th International

Conference on e-Science. The BSSwF recognition also enabled Vahi to engage with National Energy Research Scientific Computing Center staff members on how to best expand his work to their HPC resources.

HONORABLE MENTIONS

Keith Beattie

The BSSwF recognition (2021) has allowed Beattie to be a part of the RSE community and contribute to the discussion and plans for improving not only the state of scientific software but also as an advocate for the career development and advancement of RSEs. He continues to apply the tools, processes, and rigor learned in industry to produce quality scientific software that impacts the overall scientific community. Beattie considers the BSSwF as one of several relatively new initiatives creating a community to address these needs.

Addi Malviya-Thakur

The BSSwF recognition (2021) encouraged Malviya-Thakur to continue developing methodologies and guidelines for systematic software design and development that consider underlying hardware and data architecture and environment semantics. Her focus on design patterns for scientific software identifies a very important aspect to be improved due to the unpredictable expansion of scientific software. Her efforts benefited novel scientific applications being written as a part of the DOE's ECP and the Oak Ridge National Laboratory's (ORNL's) Spallation Neutron Source to adopt design patterns in their data-reduction software. Her work raised awareness for code reliability, maintainability, coherency, extensibility, and usability. She hopes her efforts will advance the science of software development for the foreseeable future.

Sarah E. Bratt

The BSSwF recognition (2022) enabled Bratt to identify opportunities to shift the narrative and revise organizational structures to better support and accelerate future RSE work, such as integration in nontraditional fields. Her research uses user-experience design approaches to 1) develop a model of how genetics and genomics biologists have institutionalized data deposit, 2) assist emerging disciplines that are institutionalizing data deposit (e.g., the social and political sciences), and 3) develop HPC workflows.[4] The BSSwF and the ECP annual conference have created an avenue for piloting a project in workflows for GenBank metadata with collaborators that support her vision to connect library and information science research with HPC solutions. The fellowship allowed Bratt to connect with a community of diverse scholars whose expertise and advocacy for better scientific software leadership was very valuable for their own research.

William F. Godoy

The BSSwF (2022) gave Godoy recognition for the RSE work performed during his career in contributing to DOE-funded HPC frameworks, aside from regular publications. The BSSwF empowered him to continue fostering a community that considers better scientific software practices as essential investments. He increased his participation in related activities, such as offering tutorials at ORNL on software practices, being a part of initiatives such as the Supercomputing Reproducibility Committee, and promoting publication of RSE work as an author and a reviewer. Godoy also serves as a mentor in internship programs targeting underrepresented minorities: the ECP Sustainable Research Pathways and the National GEM fellowship to contribute to the next generation of RSEs and HPC practitioners.

Rafael M. Mudafort

The BSSwF recognition (2022) has enabled Mudafort to engage with the community of fellows directly and empowered him to continue advocating for the importance of software as a primary artifact of research from his position at the National Renewable Energy Laboratory (NREL). He continues to advocate among DOE and NREL stakeholders to incentivize investments to improve the quality of software artifacts. His goal is to switch the default focus from the outcomes of using software to also value how the software was constructed in a manner that ensures sustainability. To achieve this goal, Mudafort continues to promote software practices such as descriptions of design, architecture diagrams, documentation, and performance metrics to demonstrate the value of RSE work as an important component in scientific activities.

OBSERVATIONS

This section gathers some of the observations across members of the BSSwF community. These are a part of the vision for the BSSwF to continuously grow and promote the community beyond fellowship activities.

Awareness, Recognition, and Networking

These three attributes are the best motivators for BSSwF applicants. Katz, Mudafort, Malviya-Thakur, Bratt, and Beattie dedicate much of their work to

recognizing software efforts and elevating the role of RSEs in the scientific community. As such, the program has created pathways among academic, national laboratory, and private sector recipients to amplify their efforts and professional networks. Some examples are the increasing interactions between BSSwF members at conferences and project meetings like SC or ECP, which would not have happened otherwise.

Addressing Technical Challenges

Addressing technical challenges, such as developing scientific software, is a hard endeavor. Several recipients recognize the need to tackle difficult challenges in their proposed work, such as correctness (Laguna and Rubio-González), workflows (Vahi and Maji), security (Sukhija and Maji), parallel I/O (Arora and Godoy), and effective collaboration (Rouson, Malviya-Thakur, and Bratt). The BSSwF is a way to connect and enable efforts around several technical challenges in building the next generation of scientific software. Funded BSSwF project descriptions are publicly available. These also motivate applicants to provide new ideas to advance the state of scientific software from different angles.

Training

A key aspect in our scientific community is training. Much of the complexity of building scientific software is due to its dynamic nature, in which new technologies and practices can result in disruptive change. In addition, gaps due to the lack of formal software development training in the broader science and the engineering community is compensated for with tutorials offered by those with more software expertise. BSSwF members have made training a central effort in their community involvement, as reflected by many of the funded initiatives described in the "Experiences" section.

Creating Venues

Creating venues is crucial to promote and grow the community around better scientific software. Rubio-González and Laguna are the founder and chair, respectively, of the prestigious Workshop on Software Correctness for HPC Applications, running since 2017 at SC. Malviya-Thakur and Godoy are a part of the organization of the annual ORNL Software and Data Expo,[k] which provides a platform for stakeholders to celebrate RSE efforts in a national laboratory environment. Katz has spent much of his efforts in building the U.S. Research Software Engineering Association to provide a sense of community around software that is independent of a particular scientific domain.

Representation

Representation is a key aspect to advance the community. Rouson, Bratt, Malviya-Thakur, and Godoy have centered their efforts around increasing participation of underrepresented sectors of the population as well as incorporating RSE recognition similar to how scientific recognition is done traditionally (papers, awards, community, posters, and so on). The diversity of the BSSwF cohort, along the axes of gender and race, is a strength that can be leveraged to develop mentorship relationships to strengthen the retention and advancement of the diverse networks around the community. Initiatives like the BSSwF are uniquely positioned to catalyze programs to address the representation imbalance in the software engineering industry and academia through its initiatives, e.g., training, workshops, and scholarships.

Lessons

One important lesson is that the initial amount of US$10,000 given to the first class of fellows was found to be insufficient to create a meaningful impact in their project. The award was raised to US$25,000 in subsequent years to attract more applicants and increase the impact of the proposed work. Another lesson is the realization that the BSSwF brings awareness and recognition to efforts in HPC to the broader scientific community. Although HPC is a major driver for many individual projects, many BSSwF members from multiple research communities see a benefit from this exposure and interaction for their own work. The BSSwF's nature of being a cosponsored program between the NSF and the DOE allows these connections, while recognizing the wide range of RSE career paths.

CONCLUSION

Scientific software is more crucial and complex than ever. The end of Moore's law, the avalanche of cloud services, heterogeneity of computing hardware, growth of the RSE movement, and the data tsunami from AI/ML workflows as major drivers of modern computing are disruptive to traditional software design paradigms and have henceforth put more emphasis on how we build future software. From the experiences presented in this article, we believe that the BSSwF recognizes and elevates the role of those researching, developing, and maintaining key software components in the critical path of our scientific mission. The effort to bring a diverse pool of individuals has been

[k]https://softwaredataexpo.ornl.gov

tremendously successful in enlightening the role of scientific software in innovation and discoveries. Our shared view is that the future of RSEs must incorporate initiatives like the BSSwF, which values and fosters these efforts and creates the required synergies among the scientific communities to achieve their goals.

ACKNOWLEDGMENTS

Bernholdt thanks the other members of the BSSwF Executive Committee for the honor of representing them in this publication: Varun Chandola (NSF), Mike Heroux (Sandia National Laboratories), Lois Curfman McInnes (Argonne National Laboratory), Hai Ah Nam (Lawrence Berkeley National Laboratory), and Shelly Olsan (Krell Institute). Godoy would like to acknowledge Arthur Barney Maccabe and Jay Jay Billings for being early enablers of the Software Expo and the RSE community at ORNL. This work was supported by the BSSwF program and the IDEAS subproject funded by the ECP (17-SC-20-SC), a collaborative effort of the DOE's Office of Science and the National Nuclear Security Administration; and by the NSF under Grant 2154495. This article was authored by UT-Battelle, limited liability company, under contract DE-AC05-00OR22725 with the DOE. The publisher acknowledges a U.S. government license to provide public access under the DOE's Public Access Plan.

REFERENCES

1. M. A. Heroux et al., "Advancing scientific productivity through better scientific software: Developer productivity and software sustainability report," Oak Ridge Nat. Lab., Oak Ridge, TN, USA, Tech. Rep. ORNL/TM-2020/1459, ECP-U-RPT-2020-0001, Feb. 2020. [Online]. Available: https://www.osti.gov/biblio/1606662/#main-body
2. A. M. Smith, D. S. Katz, K. E. Niemeyer, and FORCE11 Software Citation Working Group, "Software citation principles," *PeerJ Comput. Sci.*, vol. 2, Sep. 2016, Art. no. e86, doi: 10.7717/peerj-cs.86.
3. E. Deelman et al., "Pegasus, a workflow management system for science automation," *Future Gener. Comput. Syst.*, vol. 46, pp. 17–35, May 2015, doi: 10.1016/j.future.2014.10.008.
4. S. E. Bratt, J. J. Hemsley, J. Qin, and M. R. Costa, "Big data, big metadata and quantitative study of science: A workflow model for big scientometrics," *Proc. Assoc. Inf. Sci. Technol.*, vol. 54, no. 1, pp. 36–45, Oct. 2017, doi: 10.1002/pra2.2017.14505401005.

WILLIAM F. GODOY is a senior computer scientist in the Computer Science and Mathematics Division at Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA. His research interests include high-performance computing scientific software infrastructure, programming models, and performance. Godoy received his Ph.D. degree in mechanical engineering from the State University of New York at Buffalo. He is a senior member of the IEEE Computer Society and a member of Association for Computing Machinery, the American Society of Mechanical Engineers, and U.S. Research Software Engineer Association. Contact him at godoywf@ornl.gov.

RITU ARORA is the assistant vice president of research computing at the University of Texas at San Antonio, San Antonio, TX, 78249, USA. Her research interests include high-performance computing, cloud computing, large-scale data management, software engineering, and health informatics. Arora received her Ph.D. degree in computer and information science from the University of Alabama at Birmingham. Contact her at ritu.arora@utsa.edu.

KEITH BEATTIE is a computer systems engineer and a group leader in the Scientific Data Division at Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. His research interests include bringing adaptive, modern, scalable, and open source software engineering and development practices to academic and research contexts. Beattie received his MSc degree in computer science from San Francisco State University. Contact him at ksbeattie@lbl.gov.

DAVID E. BERNHOLDT is a distinguished R&D staff member in the Computer Science and Mathematics Division at Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA. His research interests include software environments for scientific computing, programming systems, software architecture, and software engineering. Bernholdt received his Ph.D. degree in chemistry from the University of Florida. He is a member of Association for Computing Machinery and the Society for Industrial and Applied Mathematics. Contact him at bernholdtde@ornl.gov.

SARAH E. BRATT is an assistant professor at the University of Arizona School of Information, Tucson, AZ, 85721, USA. Her research interests include scientific software sustainability, neuroimaging and genomics data workflows, and science policy. Bratt received her Ph.D. degree in information science and technology from Syracuse University. Contact her at sebratt@arizona.edu.

DANIEL S. KATZ is the chief scientist at the National Center for Supercomputing Applications and a research associate

professor of computer science, electrical and computer engineering at the School of Information, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA. His research interests include advanced cyberinfrastructure, applications, algorithms, fault tolerance, and programming in parallel and distributed computing as well as policy issues such as citation and credit mechanisms for software and data. Katz received his Ph.D. degree in electrical engineering from Northwestern University. He is a Senior Member of IEEE and Association for Computing Machinery and a member of the IEEE Computer Society Board of Governors. Contact him at dskatz@illinois.edu.

IGNACIO LAGUNA is a computer scientist at the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA, 94551, USA. His research interests include high-performance computing programming models and systems, software correctness, program analysis, compilers, fault tolerance, and debugging. Laguna received his Ph.D. degree in computer engineering from Purdue University. He is a senior member of the IEEE Computer Society. Contact him at lagunaperalt1@llnl.gov.

AMIYA K. MAJI is a senior computational scientist at the Rosen Center for Advanced Computing, Purdue University, West Lafayette, IN, 47906, USA. His research interests include reliability and security of large-scale distributed systems, algorithms and tools for software testing in high-performance computing and the Internet of Things. Maji received his Ph.D. degree in computer engineering from Purdue University. He served as a fellow of Trusted CI. Contact him at amaji@purdue.edu.

ADDI MALVIYA-THAKUR is the group leader for the Software Engineering Group in the Computer Science and Mathematics Division at Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA. Her research interests include interconnected science and federated systems, research software engineering, operational workflows, software analytics, software frameworks, and ecosystems for science. Malviya-Thakur received her M.S. degree in computer science and engineering from the University of Florida. She is a member of the IEEE Computer Society and Association for Computing Machinery. Contact her at malviyaa@ornl.gov.

RAFAEL M. MUDAFORT is a senior researcher in the National Wind Technology Center at the National Renewable Energy Laboratory, Golden, CO, 80401, USA. His research interests include high-performance computing, computational engineering, multiphysics modeling, and sustainable software development. Mudafort received his MEng degree in mechanical and aerospace engineering from the University of Virginia. Contact him at rafael.mudafort@nrel.gov.

NITIN SUKHIJA is director of the Center for Cybersecurity and Advanced Computing and an associate professor of computer science at Slippery Rock University of Pennsylvania, Slippery Rock, PA, 16057, USA. His research interests include data analytics, autonomic computing, predictive control, and cyberresilience. Sukhija received his Ph.D. degree in computer science from Mississippi State University. He is a member of the IEEE Computer Society and Association for Computing Machinery (ACM) and cochair of the ACM's Special Interest Group in High-Performance Computing Education Chapter. Contact him at nitin.sukhija@sru.edu.

DAMIAN ROUSON is a staff scientist and the group leader for the Computer Languages and Systems Software Group at Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. His research interests include machine learning methods to accelerate predictions of climate change's regional impacts, programming models, and agile software development. Rouson received his Ph.D. degree in mechanical engineering from Stanford University. Contact him at rouson@lbl.gov.

CINDY RUBIO-GONZÁLEZ is an associate professor of computer science at the University of California, Davis, Davis, CA, 95616, USA. Her research interests include programming languages, software engineering, high-performance computing, reliability, and performance. Rubio-González received her Ph.D. degree in computer science from the University of Wisconsin–Madison. She is a member of the Association for Computing Machinery Special Interest Group on Programming Languages Executive Committee. Contact her at crubio@ucdavis.edu.

KARAN VAHI is a senior computer scientist in the Science Automation Technologies Group at the University of Southern California (USC) Information Sciences Institute, Marina Del Rey, CA, 90292, USA. His research interests include scientific workflows, workflow management systems, and lifecycles. Vahi received his MSc degree in computer science from USC. Contact him at vahi@isi.edu.