IEEE Special Issue on *Innovative R&D Toward the Exascale Era*

Sadaf R. Alam, Lois Curfman McInnes, *Member, IEEE*, and Kengo Nakajima, *Member, IEEE*

The special issue on *Innovative R&D toward the Exascale Era* explores new foundational and translational research toward enabling exascale computing for emerging scientific and societal challenges. Exascale computing is defined as the capability to perform 10^{18} operations per second. Productively harnessing such a scale of processing, storage, and networking capabilities for diverse domains—including high-performance computing (HPC) simulations, artificial intelligence (AI), and extreme data-driven computing—relies on not only revitalizing existing parallel and distributed computing technologies but also innovating new solutions. Papers in this special issue explore diverse topics in research encompassing parallel, distributed, and heterogeneous systems for exascale, including advances in applications, programming environments, runtimes, libraries, innovative algorithms, domain-specific frameworks, systems architecture, performance analysis, data processing, and networking technologies.

We thank all authors for their submissions. The selection process from the 40 submissions involved multiple stages. In the first review round, we had 13 minor revisions and 15 major revisions. In the second revision round, 18 papers were accepted, together with 8 papers under a minor revision. In the third revision round, 6 papers were accepted. In total, 24 papers were accepted for the special issue, leading to a final acceptance rate of 60 percent.

Of the accepted papers, four focus on system software and related topics, with emphasis on addressing issues in the complexity of emerging architectures:

- "Anomaly detection and anticipation in high performance computing systems," [A1]
- "Online power management for multi-cores: A reinforcement learning based approach," [A2]
- "Near-zero downtime recovery from transient-error-induced crashes," [A3] and
- "Compiler-assisted compaction/restoration of SIMD instructions." [A4]

- *Sadaf R. Alam is with the Swiss National Supercomputing Centre (CSCS), 6900 Lugano, Switzerland. E-mail: alam@cscs.ch.*
- *Lois Curfman McInnes is with the Mathematics and Computer Science Division of Argonne National Laboratory, Lemont, IL 60439 USA. E-mail: curfman@anl.gov.*
- *Kengo Nakajima is with the Supercomputing Research Division, Information Technology Center, the University of Tokyo, Kashiwa, Chiba 277-0882, Japan. E-mail: nakajima@cc.u-tokyo.ac.jp.*

Digital Object Identifier no. 10.1109/TPDS.2021.3109651

Six papers focus on programming models, runtimes, tools, and libraries that are helping users to handle the difficulties of working with multiple heterogeneous architectures while working toward performant code at extreme scale:

- "EXA2PRO: A framework for high development productivity on heterogeneous computing systems," [A5]
- "Kokkos 3: Programming model extensions for the exascale era," [A6]
- "Design and performance characterization of RADICAL-pilot on leadership-class platforms," [A7]
- "LB4OMP: A dynamic load balancing library for multithreaded applications," [A8]
- "The PetscSF scalable communication layer," [A9] and
- "An automated tool for analysis and tuning of GPU-accelerated code in HPC applications" [A10]

Four papers focus on data-intensive processing, parallel I/O, and storage systems, at scale:

- "Enabling scalable and extensible memory-mapped datastores in userspace," [A11]
- "Improving I/O performance for exascale applications through online data layout reorganization," [A12]
- "Transparent asynchronous parallel I/O using background threads," [A13] and
- "Accelerating HDF5 I/O for exascale using DAOS." [A14]

Three papers address scalable graph and network computations:

- "Characterizing performance of graph neighborhood communication patterns," [A15]
- "A parallel algorithm template for updating single-source shortest paths in large-scale dynamic networks," [A16] and
- "TianheGraph: Customizing graph search for Graph500 on the Tianhe supercomputer." [A17]

Seven papers focus on other algorithms and applications, including the use of new programming models and mixed precision computation to work toward efficient and scalable performance at exascale:

- "VPIC 2.0: Next generation particle-in-cell simulations," [A18]

- "Accelerating geostatistical modeling and prediction with mixed-precision computations: A high-productivity approach with PaRSEC," [A19]
- "libEnsemble: A library to coordinate the concurrent evaluation of dynamic ensembles of calculations," [A20]
- "Combinatorial BLAS 2.0: Scaling combinatorial algorithms on distributed-memory systems," [A21]
- "Evaluating spatial accelerator architectures with tiled matrix-matrix multiplication," [A22]
- "GSOFA: Scalable sparse LU symbolic factorization on GPUs," [A23] and
- "Accelerating restarted GMRES with mixed precision arithmetic." [A24]

We thank all members of the review committee for their hard work and insightful feedback on submissions. We also thank the community for interest in this special issue. We hope that the information in this special issue will help to further advance insights in work toward the exascale era.

APPENDIX: RELATED ARTICLES

[A1] A. Borghesi, M. Molan, M. Milano, and A. Bartolini, "Anomaly detection and anticipation in high performance computing systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 739–750, Apr. 2022.

[A2] Y. Wang, W. Zhang, M. Hao, and Z. Wang, "Online power management for multi-cores: A reinforcement learning based approach," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 751–764, Apr. 2022.

[A3] C. Chen, G. Eisenhauer, and S. Pande, "Near-zero downtime recovery from transient-error-induced crashes," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 765–778, Apr. 2022.

[A4] J. M. Cebrian *et al.*, "Compiler-assisted compaction/restoration of SIMD instructions," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 779–791, Apr. 2022.

[A5] L. Papadopoulos *et al.*, "EXA2PRO: A framework for high development productivity on heterogeneous computing systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 792–804, Apr. 2022.

[A6] Christian R. Trott *et al.*, "Kokkos 3: Programming model extensions for the exascale era," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 805–817, Apr. 2022.

[A7] A. Merzky, M. Turilli, M. Titov, A. Al-Saadi, and S. Jha, "Design and performance characterization of RADICAL-pilot on leadership-class platforms," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 818–829, Apr. 2022.

[A8] J. H. Müller Korndörfer, A. Eleliemy, A. Mohammed, and F. M. Ciorba, "LB4OMP: A dynamic load balancing library for multithreaded applications," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 830–841, Apr. 2022.

[A9] J. Zhang *et al.*, "The PetscSF scalable communication layer," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 842–853, Apr. 2022.

[A10] K. Zhou, X. Meng, R. Sai, D. Grubisic, and J. Mellor-Crummey, "An automated tool for analysis and tuning of GPU-accelerated code in HPC applications," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 854–865, Apr. 2022.

[A11] I. B. Peng, M. B. Gokhale, K. Youssef, K. Iwabuchi, and R. Pearce, "Enabling scalable and extensible memory-mapped datastores in userspace," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 866–877, Apr. 2022.

[A12] L. Wan *et al.*, "Improving I/O performance for exascale applications through online data layout reorganization," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 878–890, Apr. 2022.

[A13] H. Tang, Q. Koziol, J. Ravi, and S. Byna, "Transparent asynchronous parallel I/O using background threads," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 891–902, Apr. 2022.

[A14] J. Soumagne *et al.*, "Accelerating HDF5 I/O for exascale using DAOS," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 903–914, Apr. 2022.

[A15] S. Ghosh, N. R. Tallent, and M. Halappanavar, "Characterizing performance of graph neighborhood communication patterns," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 915–928, Apr. 2022.

[A16] A. Khanda, S. Srinivasan, S. Bhowmick, B. Norris, and S. K. Das, "A parallel algorithm template for updating single-source shortest paths in large-scale dynamic networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 929–940, Apr. 2022.

[A17] X. Gan *et al.*, "TianheGraph: Customizing graph search for Graph500 on Tianhe supercomputer," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 941–951, Apr. 2022.

[A18] R. Bird, N. Tan, S. V. Luedtke, S. L. Harrell, M. Taufer, and B. Albright, "VPIC 2.0: Next generation particle-in-cell simulations," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 952–963, Apr. 2022.

[A19] S. Abdulah *et al.*, "Accelerating geostatistical modeling and prediction with mixed-precision computations: A high-productivity approach with PaRSEC," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 964–976, Apr. 2022.

[A20] S. Hudson, J. Larson, J.-L. Navarro, and S. M. Wild, "libEnsemble: A library to coordinate the concurrent evaluation of dynamic ensembles of calculations," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 977–988, Apr. 2022.

[A21] A. Azad, O. Selvitopi, M. T. Hussain, J. R. Gilbert, and A. Buluc, "Combinatorial BLAS 2.0: Scaling combinatorial algorithms on distributed-memory systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 989–1001, Apr. 2022.

[A22] G. E. Moon, H. Kwon, G. Jeong, P. Chatarasi, S. Rajamanickam, and T. Krishna, "Evaluating spatial accelerator architectures with tiled matrix-matrix multiplication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 1002–1014, Apr. 2022.

[A23] A. Gaihre, X. S. Li, and H. Liu, "GSOFA: Scalable sparse symbolic LU factorization on GPUs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 1015–1026, Apr. 2022.

[A24] N. Lindquist, P. Luszczek, and J. Dongarra, "Accelerating restarted GMRES with mixed precision arithmetic," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 4, pp. 1027–1037, Apr. 2022.

Sadaf R. Alam received the PhD degree in computer science from the University of Edinburgh, U.K., in 2004. She is currently a chief technology officer (CTO) with the Swiss National Supercomputing Centre (CSCS) in Lugano, Switzerland. Until March 2009, she was a computer scientist with the Oak Ridge National Laboratory, USA. In her role as the CTO, she ensures end-to-end integrity of HPC systems and storage solutions and leads strategic projects with the centre. She has held different roles at CSCS, including group lead of future systems, chief architect and head of operations. She is a member of ACM, ACM-W, SIGHPC and Women in HPC.

Lois Curfman McInnes (Member, IEEE) received the PhD degree in applied mathematics from the University of Virginia, in 1993. She is currently a senior computational scientist with the Mathematics and Computer Science Division of Argonne National Laboratory. Her research focuses on high-performance computational science, with emphasis on scalable numerical libraries and community collaboration toward productive and sustainable software ecosystems. She serves as deputy director of software technology with the U.S. Department of Energy's Exascale Computing Project. She is a SIAM fellow and serves on SIAM Council. She won the 2015 SIAM/ACM Prize in Computational Science and Engineering and a R&D 100 Award for PETSc in 2009 (with collaborators), and she received the 2011 Ernest Orlando Lawrence Award for outstanding contributions in R&D supporting the DOE. She is a member of ACM and SIAM.

Kengo Nakajima (Member, IEEE) received the BEng degree in aeronautics from the University of Tokyo, in 1985, the MS degree in aerospace engineering from the University of Texas at Austin, in 1993, and the PhD degree in engineering mechanics from the University of Tokyo, in 2003. He is currently a professor with the Supercomputing Research Division, Information Technology Center, University of Tokyo. Prior to joining the University of Tokyo in 2004, he spent 19 years in industry. He also has been a deputy director of RIKEN R-CCS (Center for Computational Science) since 2018. His research interests include computational mechanics, computational fluid dynamics, numerical linear algebra, parallel iterative algorithms, parallel preconditioning methods, multigrid methods, parallel programming models, adaptive mesh refinement, and parallel visualization.