Jens Domke, Dr.

Curriculum Vitae

1-1-9 Minatojima,
Chuo-ku, Kobe 650-0045

Education

Oct. 2014 – Mar. 2017 **PhD student at Faculty of Computer Science**, *Technische Universität Dresden*, Dresden, Germany.
- Topic: Routing on the Channel Dependency Graph: A New Approach to Deadlock-Free, Destination-Based, High-Performance Routing for Lossless Interconnection Networks
- Final Degree: Doctor rerum naturalium (Dr. rer. nat.) awarded in June 2017

Oct. 2004 – Dec. 2010 **Diploma Student in Mathematics**, *Technische Universität Dresden*, Dresden, Germany.
- Topic: Optimized Routing for InfiniBand Networks in the field of HPC systems
- Final Degree: Diplommathematiker

Aug. 1995 – July 2003 **High School**, *Landau-Gynmasium Weißwasser*, Weißwasser, Germany.
- Qualification: High School Diploma (ger. Abitur)

Aug. 1991 – June 1995 **Primary School**, *1. Grundschule Weißwasser*, Weißwasser, Germany.

Employment History

Research

July 2022 – present **Team Leader of the Supercomputing Performance Research Team**, *RIKEN Center for Computational Science (R-CCS)*, Kobe, Japan.
- Performance modelling, performance studies, HPC co-design, and tool development
- Supervising researchers, postdoctoral researchers, and students

Apr. 2021 – June 2022 **Research Scientist of the High Performance Big Data Research Team**, *RIKEN Center for Computational Science (R-CCS)*, Kobe, Japan.
- Benchmarking processor- and network-related aspects of current HPC installation at R-CCS
- Researching alternative CPU, memory, and network architectures to steer co-design

Apr. 2019 – Apr. 2021 **Postdoctoral Researcher of the High Performance Big Data Research Team**, *RIKEN Center for Computational Science (R-CCS)*, Kobe, Japan.
- Benchmarking processor- and network-related aspects of current HPC installation at R-CCS
- Researching alternative CPU, memory, and network architectures to steer co-design

May 2019 – present **Tokyo Tech Research Fellow**, *Tokyo Institute of Technology*, Tokyo, Japan.
- Performing interconnect simulations and benchmarking newest HPC hardware

Feb. 2019 – Mar. 2019 **RIKEN R-CCS Visiting Researcher**, *RIKEN Center for Computational Science (R-CCS)*, Kobe, Japan.
- Assisting and consulting in the early stages of the HPC system procurement
- Benchmarking network-related aspects of current HPC installation at R-CCS

Apr. 2017 – Mar. 2019 **Research Staff at Matsuoka & Endo Laboratory, Global Scientific Information and Computing Center**, *Tokyo Institute of Technology*, Tokyo, Japan.

- Modifying an existing supercomputer at TokyoTech to evaluate the novel HyperX topology
- Research in the area of HPC interconnects, topology design, and simulation frameworks
- Planning and consulting activities for future AI/HPC supercomputers, such as TSUBAME4 and AIST's ABCI system
- Supervising and mentoring undergraduate and graduate students of the Matsuoka Laboratory

July 2016 – Sept. 2016 **Student Internship**, *Lawrence Livermore National Laboratory*, Livermore, United States of America.

- Applying modern network routing techniques to large-scale production HPC systems and measuring the impact
- Analyzing the disadvantages of existing routing algorithms and develop solutions to allow a more efficient and workload-specific traffic routing
- Combining two existing network simulation frameworks to improve simulation accuracy and scalability of statically routed large-scale HPC topologies

Feb. 2015 – Mar. 2017 **Tokyo Tech Research Fellow**, *Tokyo Institute of Technology*, Tokyo, Japan.

- Performing network simulations and benchmarking newest InfiniBand hardware

Oct. 2014 – Mar. 2017 **Research Associate at Faculty of Computer Science**, *Technische Universität Dresden*, Dresden, Germany.

- Research in the area of high performance interconnection networks at the Institute of Computer Engineering
- Implementation of software tools for monitoring and performance analysis
- Integration into the teaching activities of the institute and student supervision

Sept. 2013 – Sept. 2014 **Research Associate at Matsuoka Laboratory, Global Scientific Information and Computing Center**, *Tokyo Institute of Technology*, Tokyo, Japan.

- Development of new deadlock-free oblivious routing algorithms

Oct. 2012 – Sept. 2013 **Visiting Researcher at Matsuoka Laboratory, Global Scientific Information and Computing Center**, *Tokyo Institute of Technology*, Tokyo, Japan.

- Enhancing the deadlock-free oblivious routing algorithm and focusing on research related to routing algorithms, network topologies and network faults
- Providing knowledge and support in the field of performance analysis tools, like Vampir and Scalasca, for an application scalability study

Aug. 2011 – Aug. 2012 **Research Associate at Joint Institute for Computational Sciences (JICS) – University of Tennessee and Oak Ridge National Laboratory (ORNL)**, *University of Tennessee*, Knoxville, United States of America.

- Enhancing the functionality of the Vampir framework, from the Technische Universität Dresden, to support running in a mixed environment of X86 processors and General-Purpose computation on Graphics Processing Units processors
- Supporting the analysis process of large scale applications on Jaguar (Jaguar was the third fastest HPC system in the Top500 list, Nov. 2011)

Feb. 2011 – Aug. 2011 **Research Associate at Center for Information Services and High Performance Computing (ZIH)**, *Technische Universität Dresden*, Dresden, Germany.

- Extending the Vampir/VampirTrace tool set within an European–Russian project, called HOPSA

Nov. 2007 – Dec. 2010 **Student Research Assistant at Center for Information Services and High Performance Computing (ZIH)**, *Technische Universität Dresden*, Dresden, Germany.

- Implementation/parallelization of different sparse linear algebra algorithms for Cell BE
- BenchIT Project: Designing and testing of measurement kernels
- Performed measurements on a variety of HPC installations: SGI Altix 4700, Linux Networx Evolocity II PC-Farm, IBM BladeCenter QS21, NEC SX6 (ZIH, Dresden), NEC SX8 (HLRS, Stuttgart)
- Winner (Team) of the Cluster Challenge at Supercomputing 08 (Austin, Texas)

Vocational Internship

Aug. 2007 – Sept. 2007 **Usability investigation of Cell BE based console for efficient calculation of sparse matricies**, *Center for Information Services and High Performance Computing at Technische Universität Dresden*, Dresden, Germany.

- Build and tuned a cluster of Playstation 3 consoles and implemented different sparse linear algebra algorithms

German Armed Forces

Oct. 2003 – June 2004 **Officer Candidate**, *German Air Force*, Bayreuth, Fürstenfeldbruck, Germany.

- In early 2004, resigned as officer candidate and moved to basic military service

Computer skills

Languages:	C, Python, FORTRAN, Bourne Shell, MATLAB, C++, Java	Programming:	InfiniBand, MPI, OpenMP, x86, Cell BE
Platforms:	GNU/Linux, Mac OS X	Math Packages:	MATLAB, Maple, SPSS
Performance Analysis:	Vampir/VampirTrace, Score-P, BenchIT, GNU gprof, TAU, PAPI	Publishing:	LaTeX

Languages

German:	Fluent	*native language*
English:	Upper-Intermediate	*ten years at school, two semester at university, living in the US for 1 year*

Publications

J. Mueller, T. Schneider, J. Domke, R. Geyer, M. Haesing, T. Hoefler, S. Hoehlig, G. Juckeland, A. Lumsdaine, M. Mueller, and W. Nagel, "Cluster Challenge 2008: Optimizing Cluster Configuration and Applications to Maximize Power Efficiency," in *Proceedings of the 10th LCI International Conference on High-Performance Clustered Computing*, Mar. 2009. LCI'09 2nd Best Paper Award.

J. Domke, T. Hoefler, and W. E. Nagel, "Deadlock-Free Oblivious Routing for Arbitrary Topologies," in *Proceedings of the 25th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, (Washington, DC, USA), pp. 613–624, IEEE Computer Society, May 2011.

J. Domke and D. Wang, "Runtime Tracing of the Community Earth System Model: Feasibility Study and Benefits," *Procedia Computer Science*, vol. 9, no. 0, pp. 1950–1958, 2012. Proceedings of the International Conference on Computational Science, ICCS 2012.

J. Domke, T. Hoefler, and S. Matsuoka, "Fail-in-Place Network Design: Interaction between Topology, Routing Algorithm and Failures," in *Proceedings of the IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC14)*, SC '14, (New Orleans, LA, USA), pp. 597–608, IEEE Press, Nov. 2014.

K. Brown, J. Domke, and S. Matsuoka, "Hardware-Centric Analysis of Network Performance for MPI Applications," in *2015 21th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, (Melbourne, Australia), p. 8, IEEE Press, Dec. 2015.

J. Domke, T. Hoefler, and S. Matsuoka, "Routing on the Dependency Graph: A New Approach to Deadlock-Free High-Performance Routing," in *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '16, (New York, NY, USA), pp. 3–14, ACM, 2016.

J. Domke and T. Hoefler, "Scheduling-Aware Routing for Supercomputers," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '16, (Piscataway, NJ, USA), pp. 13:1–13:12, IEEE Press, 2016.

D. Wang, J. Domke, J. Mao, X. Shi, and D. M. Ricciuto, "A Scalable Framework for the Global Offline Community Land Model Ensemble Simulation," *Int. J. Comput. Sci. Eng.*, vol. 12, pp. 73–85, Feb. 2016.

N. Wolfe, M. Mubarak, N. Jain, J. Domke, A. Bhatele, C. D. Carothers, and R. B. Ross, "Preliminary Performance Analysis of Multi-rail Fat-tree Networks," in *17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, CCGrid '17, (Madrid, Spain), pp. 258–261, IEEE Press, May 2017. Short paper.

J. Domke, *Routing on the Channel Dependency Graph: A New Approach to Deadlock-Free, Destination-Based, High-Performance Routing for Lossless Interconnection Networks*. Dissertation, Technische Universität Dresden, Fakultät Informatik, Professur für Rechnerarchitektur, 2017.

M. Mubarak, N. Jain, J. Domke, N. Wolfe, C. Ross, K. Li, A. Bhatele, C. D. Carothers, K.-L. Ma, and R. B. Ross, "Toward Reliable Validation of HPC Interconnect Simulations," in *Proceedings of the 2017 Winter Simulation Conference*, WSC '17, (Las Vegas, NV, USA), pp. 659–674, IEEE Press, Dec. 2017.

H. Bhatia, N. Jain, A. Bhatele, Y. Livnat, J. Domke, V. Pascucci, and P.-T. Bremer, "Interactive Investigation of Traffic Congestion on Fat-Tree Networks Using TREESCOPE," *Computer Graphics Forum*, vol. 37, no. 3, pp. 561–572, 2018.

S. A. Smith, C. E. Cromey, D. K. Lowenthal, J. Domke, N. Jain, J. J. Thiagarajan, and A. Bhatele, "Mitigating Inter-Job Interference Using Adaptive Flow-Aware Routing," in *Proceedings of the International Conference for High Performance Computing,*

Networking, Storage and Analysis, SC '18, (Piscataway, NJ, USA), pp. 27:1–27:15, IEEE Press, Nov. 2018. Note: Best student paper finalist.

J. Domke, K. Matsumura, M. Wahib, H. Zhang, K. Yashima, T. Tsuchikawa, Y. Tsuji, A. Podobas, and S. Matsuoka, "Double-precision FPUs in High-Performance Computing: an Embarrassment of Riches?," in *2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2019, Rio de Janeiro, Brazil, May 20-24, 2019*, (Rio de Janeiro, Brazil), May 2019.

J. Domke, S. Matsuoka, I. R. Ivanov, Y. Tsushima, T. Yuki, A. Nomura, S. Miura, N. McDonald, D. L. Floyd, and N. Dubé, "The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?," HOTI26, (Piscataway, NJ, USA), p. 4, IEEE Press, Aug. 2019. Note: peer-reviewed short paper.

J. Domke, S. Matsuoka, I. R. Ivanov, Y. Tsushima, T. Yuki, A. Nomura, S. Miura, N. McDonald, D. L. Floyd, and N. Dubé, "HyperX Topology: First At-Scale Implementation and Comparison to the Fat-Tree," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '19, (New York, NY, USA), pp. 40:1–40:23, ACM, Nov. 2019.

T. Dey, K. Sato, B. Nicolae, J. Guo, J. Domke, W. Yu, F. Cappello, and K. Mohror, "Optimizing Asynchronous Multi-Level Checkpoint/Restart Configurations with Machine Learning," in *IEEE International Workshop on High-Performance Storage*, HPS '20, (New Orleans, LA, USA), May 2020.

M. Wahib, H. Zhang, T. T. Nguyen, A. Drozd, J. Domke, L. Zhang, R. Takano, and S. Matsuoka, "Scaling Distributed Deep Learning Workloads beyond the Memory Capacity with KARMA," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20, IEEE Press, 2020.

M. Besta, J. Domke, M. Schneider, M. Konieczny, S. D. Girolamo, T. Schneider, A. Singla, and T. Hoefler, "High-Performance Routing With Multipathing and Path Diversity in Ethernet and HPC Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 4, pp. 943–959, 2021.

J. Domke, E. Vatai, A. Drozd, C. Peng, Y. Oyama, L. Zhang, S. Salaria, D. Mukunoki, A. Podobas, M. Wahib, and S. Matsuoka, "Matrix Engines for High Performance Computing: A Paragon of Performance or Grasping at Straws?," in *2021 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2021, Portland, Oregon, USA, May 17-21, 2021*, (Portland, Oregon, USA), p. 10, May 2021.

J. Domke, "A64FX – Your Compiler You Must Decide!," in *2021 IEEE International Conference on Cluster Computing (CLUSTER), EAHPC Workshop*, (Portland, Oregon, USA), IEEE Computer Society, 2021.

S. Farrell, M. Emani, J. Balma, L. Drescher, A. Drozd, A. Fink, G. Fox, D. Kanter, T. Kurth, P. Mattson, D. Mu, A. Ruhela, K. Sato, K. Shirahata, T. Tabaru, A. Tsaris, J. Balewski, B. Cumming, T. Danjo, J. Domke, T. Fukai, N. Fukumoto, T. Fukushi, B. Gerofi, T. Honda, T. Imamura, A. Kasagi, K. Kawakami, S. Kudo, A. Kuroda, M. Martinasso, S. Matsuoka, H. Mendonça, K. Minami, P. Ram, T. Sawada, M. Shankar, T. S. John, A. Tabuchi, V. Vishwanath, M. Wahib, M. Yamazaki, and J. Yin, "MLPerf HPC: A Holistic Benchmark Suite for Scientific Machine Learning on

HPC Systems," in *7th IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC@SC'21), St. Louis, MO, USA, November 15, 2021*, IEEE, Nov. 2021.

S. Matsuoka, J. Domke, M. Wahib, A. Drozd, R. Bair, A. A. Chien, J. S. Vetter, and J. Shalf, "Preparing for the Future—Rethinking Proxy Applications," *Computing in Science & Engineering*, vol. 24, pp. 85–90, Mar. 2022.

T. T. Nguyen, F. Trahay, J. Domke, A. Drozd, E. Vatai, J. Liao, M. Wahib, and B. Gerofi, "Why Globally Re-shuffle? Revisiting Data Shuffling in Large Scale Deep Learning," in *2022 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2022, Lyon, France, May, 2022*, (Lyon, France), May 2022.

W. S. Moses, I. R. Ivanov, J. Domke, T. Endo, J. Doerfert, and O. Zinenko, "High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs," in *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, PPoPP '23, (New York, NY, USA), pp. 119–134, Association for Computing Machinery, 2023.

S. Matsuoka, J. Domke, M. Wahib, A. Drozd, and T. Hoefler, "Myths and Legends in High-Performance Computing," *The International Journal of High Performance Computing Applications*, vol. 37, pp. 245–259, Apr. 2023.

F. Antici, K. Yamamoto, J. Domke, and Z. Kiziltan, "Augmenting ML-based Predictive Modelling with NLP to Forecast a Job's Power Consumption," in *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, SC-W '23, (New York, NY, USA), pp. 1820–1830, IEEE, Nov. 2023.

O. Pearce, A. Scott, G. Becker, R. Haque, N. Hanford, S. Brink, D. Jacobsen, H. Poxon, J. Domke, and T. Gamblin, "Towards Collaborative Continuous Benchmarking for HPC," in *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, SC-W '23, (New York, NY, USA), pp. 627–635, IEEE, Nov. 2023.

J. Domke, E. Vatai, B. Gerofi, Y. Kodama, M. Wahib, A. Podobas, S. Mittal, M. Pericàs, L. Zhang, P. Chen, A. Drozd, and S. Matsuoka, "At the Locus of Performance: Quantifying the Effects of Copious 3D-Stacked Cache on HPC Workloads," *ACM Transactions on Architecture and Code Optimization*, vol. 20, pp. 1–26, Dec. 2023.