

Torsten Hoefler

Full Professor of Computer Science, Swiss Federal Institute of Technology (ETH Zürich) and
Chief Architect for Machine Learning, Swiss National Supercomputing Center (CSCS)

Education

- 2005–2008 **Ph.D., Computer Science (Dr. rer. nat.)** **Indiana University
Bloomington, IN, USA**
GPA: 4.0/4.0 (“summa cum laude”)
Committee: Andrew Lumsdaine, Randall Bramley, Jack Dongarra, Richard Graham, Minaxi Gupta
IU Young Alumni Award 2014 and **Luddy Distinguished Alumni Award 2022**
- 2000–2004 **Diplom, Informatik (Master of CS)** **Chemnitz University of Technology
Chemnitz, Germany**
Grade: sehr gut (“very good”)
Universitätspreis 2005 (best student in class), Advisor: Wolfgang Rehm
- 1993–1999 **Gymnasium (Academic High School)** **Gymnasium Oelsnitz
Oelsnitz, Germany**
Graduated top of class (3rd best among 90 students)

Research Interests

My research interests revolve around the central topic of "Performance as a Science" in the context of High-Performance Computing (HPC) and Large-Scale Artificial Intelligence (AI). There are three main sub-branches in my work: (1) **high-performance parallel programmings**, (2) **high-performance networking**, and (3) **climate simulations and AI training**.

Awards and Honors (career awards underlined)

- 2023 **Best Student Paper Award at SC23** **Denver, CO, USA**
best technical student paper at ACM/IEEE Supercomputing 2023; selected by a committee during the conference out of a set of four finalists (out of 376 submissions), \$1,000
- 2023 **Best Reproducibility Advancement Award at SC23** **Denver, CO, USA**
best paper in terms of reproducibility at ACM/IEEE Supercomputing 2023; selected by a committee from all 90 accepted papers), \$300
- 2023 **Jack Dongarra Early Career Award** **Hamburg, Germany**
“... to an upcoming researcher [...], who has been a catalyst for scientific progress through exceptional work in their field.”, inaugural award at the 2023 International Supercomputing Conference, €5,000
- 2023 **ACM Fellow, class of 2022** **New York, NY, USA**
“The ACM Fellows program recognizes the top 1% of ACM Members”; **citation**: “for foundational contributions to High-Performance Computing and the application of HPC techniques to machine learning”
- 2023 **Who is Who in Zurich** **Zurich, Switzerland**
“The 200 most prominent persons in Zurich and Basel.” (population of ≈2M in the metro areas)
- 2022 **IEEE CS Sidney Fernbach Award** **Los Alamitos, CA, USA**
“Established in 1992 in memory of high-performance computing pioneer Sidney Fernbach, the Sidney Fernbach Memorial Award recognizes outstanding contributions in the application of high-performance computers using innovative approaches.” (oldest and one of the most significant career awards in HPC)
- 2022 **Best Paper Award at SC22** **Dallas, TX, USA**
best technical paper at ACM/IEEE Supercomputing 2022; selected by a committee during the conference out of a set of seven finalists (out of 320 submissions), \$1,000
- 2022 **Best Reproducibility Advancement Award at SC22** **Dallas, TX, USA**
best paper in terms of reproducibility at ACM/IEEE Supercomputing 2022; selected by a committee during the conference out of a set of two finalists (out of 320 submissions), \$300
- 2022 **Luddy Distinguished Alumni Award, IU Luddy School of Informatics** **Indianapolis, IN, USA**
“The Luddy Distinguished Alumni Award recognizes outstanding contributions to computer science, informatics, and engineering over the course of a career [...]” (the school had ≈1,800 students)

- 2022 **IEEE Fellow, class of 2021** **Zurich, Switzerland**
 "following a rigorous evaluation procedure, the IEEE Fellow Committee recommends a select group of recipients for elevation to IEEE Fellow. Less than 0.1% of voting members are selected annually."
- 2022 **SIGHPC Certificate of Appreciation** **Dallas, TX, USA**
 As recognition for reproducible methods in and as an SC22 student cluster competition task. A single paper selected from SC21's proceedings out of 98 papers that appeared at the conference.
- 2021 **HPCWire Person to Watch** **San Diego, USA**
 "You have been selected for [...] the impact you will have on the HPC community and ecosystem globally. Your leadership [...] will shape not only the future of our industry but will make a difference for the quality of life on our planet today and generations to come."
- 2021 **Distinguished Paper Award at ACM OOPSLA'21** **Chicago, USA**
 designated as distinguished paper of the ACM Symposium on Object-oriented Programming, Systems, Languages, and Applications (six out of 71 accepted papers)
- 2020 **ERC Consolidator Grant** **Brussels, Europe**
 "ERC Consolidator Grants are awarded to outstanding researchers [... with ...] a scientific track record showing great promise and an excellent research proposal." (13% acceptance). EUR 2M
- 2020 **BenchCouncil Rising Star Award** **Qingdao, China**
 "The BenchCouncil Rising Star Award is given to a young scholar who has made outstanding contributions to the field of benchmarking, measurement and optimization. The winner of this award will be automatically elected as a Senior Member of the ISTC [...]" (\$1000 award and keynote at the Bench'20 conference, digitally due to COVID-19)
- 2019 **ACM Gordon Bell Prize** **Denver, CO, USA**
 "[...] recognize outstanding achievement in high-performance computing. The purpose of the award is to track the progress over time of parallel computing, with particular emphasis on rewarding innovation in applying high-performance computing to applications in science, engineering, and large-scale data analytics." awarded in the categories "Sustained Application Performance" and "Novelty of Programming Approach", \$10,000
- 2019 **Best Student Paper Award at SC19** **Denver, CO, USA**
 advisor on the best student paper at ACM/IEEE Supercomputing 2019; selected by a committee during the conference out of a set of twelve candidates (out of 387 submissions), \$1,000
- 2019 **IEEE TCSC Award for Excellence in Scalable Computing (MCR)** **Zhangjiajie, China**
 "for contributions on all aspects of large-scale scalable parallel processing systems and supercomputers"
- 2017 **Student Teaching Award "Best Interaction"** **Zurich, Switzerland**
 elected democratically by all computer science students
- 2017 **Best Student Paper Award at IEEE HOTI'17** **Santa Clara, CA, USA**
 collaborator on the best student paper at IEEE Hot Interconnects 2017, \$250
- 2016 **Outstanding Paper Award at ACM OOPSLA'16** **Amsterdam, Netherlands**
 designated as outstanding paper of the ACM Symposium on Object-oriented Programming, Systems, Languages, and Applications (four out of 52 accepted papers (203 submissions))
- 2016 **Best Student Paper Award at IEEE HOTI'16** **Santa Clara, CA, USA**
 advisor on the best student paper at IEEE Hot Interconnects 2016, \$250
- 2016 **Karsten Schwan Best Paper Award at ACM HPDC'16** **Kyoto, Japan**
 designated as best paper of the ACM Symposium on High-Performance Parallel and Distributed Computing (out of 20 accepted papers (129 submissions))
- 2015 **Latsis Prize of ETH Zürich** **Zürich, Switzerland**
 "The purpose of the Latsis Prize is to recognize and reward scientific work of particular excellence from all fields of research undertaken at the ETH Zurich" (one award per year to one scientist across all disciplines), CHF 25,000
- 2015 **ERC Starting Grant** **Brussels, Europe**
 "ERC Starting Grants aim to support up-and-coming research leaders who are about to establish a proper research team and to start conducting independent research [...]" (8% acceptance). EUR 1.5M
- 2015 **Best Student Paper Award at IEEE HOTI'15** **Santa Clara, CA, USA**
 advisor on the best student paper at IEEE Hot Interconnects 2015, \$250

- 2015 **Best Paper Award at ACM HPDC'15** **Portland, OR, USA**
designated as best paper of the ACM Symposium on High-Performance Parallel and Distributed Computing (out of 19 accepted papers (116 submissions))
- 2015 **Best Paper Award at IEEE Intl. Parallel & Distr. Processing Symposium** **Hyderabad, India**
designated as best paper of the software track at IPDPS'15 (four tracks, one award each, plenary presentation, of 108 accepted papers (496 submissions))
- 2014 **Best Student Paper Award at SC14** **New Orleans, LA, USA**
advisor on the best student paper at ACM/IEEE Supercomputing 2014; selected by a committee during the conference out of a set of seven finalists (out of 394 submissions), \$1,000
- 2014 **Young Alumni Award, Indiana University School of Informatics** **Indianapolis, IN, USA**
"in recognition of outstanding early career achievement that brings acclaim and recognition to the field of informatics, and honor and distinction to Indiana University." (the school had \approx 1,800 students)
- 2013 **Best Paper Award at SC13** **Denver, CO, USA**
designated as best paper at ACM/IEEE Supercomputing 2013; selected by a committee during the conference out of a set of thirteen finalists (out of 457 submissions)
- 2013 **IEEE TCSC Young Achiever in Scalable Computing** **Denver, CO, USA**
"Awarded to individuals who have made outstanding, influential, and potentially long-lasting contributions in the field of scalable computing within 5 years of receiving their PhD."
- 2013 **IBM Faculty Award** **Yorktown Heights, NY, USA**
"To qualify for this internationally competitive award [...] candidates must have an outstanding reputation for contributions in their field or, in the case of junior faculty, show unusual promise.", \$30,000
- 2013 **Best Paper Award at EuroMPI'13** **Madrid, Spain**
designated as best paper of EuroMPI 2013 after a two-round review process, \approx \$3,400 gift
- 2012 **SIAM SIAG/SC Junior Scientist Prize** **Savannah, GA, USA**
"awarded to an outstanding junior researcher in the field of algorithms research and development for parallel scientific and engineering computing", \approx \$2,000 travel funds
- 2011 **Best Poster Award PPOPP'11** **San Antonio, TX, USA**
designated as best poster at the 2011 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming; selected by a committee during the poster session.
- 2010 **Best Paper Award at SC10** **New Orleans, LA, USA**
designated as best paper at ACM/IEEE Supercomputing 2010; selected by a committee during the conference out of a set of nine finalists (out of 257 submissions) \$1,000
- 2010 **Best Paper Award LSAP'10** **Chicago, IL, USA**
designated as best paper at the 2010 ACM Workshop on Large-Scale System and Application Performance; selected by a committee
- 2009 **Best Paper Award LCI'09** **Boulder, CO, USA**
best student paper at the Linux Cluster Institute Conference 2009; selected by a committee; \$500
- 2008 **Cluster Challenge Champion SC'08** **Austin, TX, USA**
co-advised the winning team at IEEE/ACM SC08's Cluster Challenge; a competition involving seven international teams of undergraduate students running HPC applications on a self-made cluster computer
- 2008 **Travel Award CCGrid'08** **Lyon, France**
IEEE/TCSC Doctoral Symposium for Cluster Computing and the Grid 2008, \$2,000
- 2005 **State Fellowship for Doctoral Studies** **Chemnitz, Germany**
Saxon Ministry of Science and the Fine Arts (Sächsisches Ministerium für Wissenschaft und Kunst), one of four reputable fellowships at TU Chemnitz, €1,400/month; extension declined after one year.
- 2005 **Universitätspreis 2005 (Best Student Award)** **Chemnitz, Germany**
Chemnitz University of Technology, €2,000
- 2005 **GI/PARS Nachwuchspreis 2005 (PARS Junior Researcher Award)** **Lübeck, Germany**
Group of Parallel Algorithms, Computer Architectures and System Software in the German Computer Society (Gesellschaft für Informatik, GI), €500
- 2005 **HPC Europa, Scientific Highlight** **Caseleccio di Reno, Italy**
Selected as outstanding visitor of the HPC Europe scientific exchange program.

Awards of Mentees (received while being mentored by me)

2024	ISC'24 Cluster Competition ETH's Cluster Challenge Team won the Third Prize overall	Hamburg, Germany
2023	SC23 Cluster Competition ETH's Cluster Challenge Team won the Overall Student Cluster Competition	Denver, CO, USA
2023	Informatics Europe Best Dissertation Award Salvatore Di Girolamo's dissertation was Informatics Europe Best Dissertation Award Runner Up	Edinburgh, UK
2023	ISC'23 Cluster Competition ETH's Cluster Challenge Team won the Highest Linpack Award and the Third Prize overall	Hamburg, Germany
2022	SIGHPC Dissertation award Maciej Besta received the SIGHPC dissertation award for his PhD thesis under my supervision	Dallas, TX, USA
2022	ACM Student Research Competition Marcin Copik received the ACM SRC Gold Medal (graduate)	Dallas, TX, USA
2022	ACM-IEEE CS George Michael Memorial HPC Fellowship Marcin Copik received the prestigious ACM/IEEE-CS George Michael Memorial HPC Fellowship	Dallas, TX, USA
2022	ETH Medal for MSc Thesis Patrick Iff received the ETH Medal for his MSc thesis supervised by Maciej Besta and myself	Zurich, Switzerland
2022	ISC'22 Cluster Competition ETH's Cluster Challenge Team won the Silver Medal	Hamburg, Germany
2022	Kaivalya Dixit dissertation award (runner up) Maciej Besta received the Dixit award (runner up) for his PhD thesis under my supervision	Beijing, China
2022	ETH Medal for PhD Thesis Maciej Besta received the ETH Medal for his PhD thesis under my supervision	Zurich, Switzerland
2020	Xilinx Open Hardware Design Contest Johannes de Fine Licht won in the Compute Acceleration category	Virtual/COVID-19
2020	Xilinx Open Hardware Design Contest Manuel Burger won in the PhD Student Category	Virtual/COVID-19
2019	ACM Student Research Competition Marcin Copik received the ACM SRC Gold Medal (graduate)	Denver, CO, USA
2019	ISC'19 Cluster Competition ETH's Cluster Challenge Team won the highest Linpack and overall Bronze Medal (2/5 awards)	Frankfurt, Germany
2018	ACM Student Research Competition Salvatore Di Girolamo received the ACM SRC Bronze Medal (graduate)	Dallas, TX, USA
2018	SNSF Ambizione Fellowship Tal Ben-Nun received the prestigious Ambizione Fellowship to build a research group	Bern, Switzerland
2018	ETH Medal for MSc Thesis Cedric Renggli received the ETH Medal for his MSc thesis under Dan Alistarh's and my supervision	Zurich, Switzerland
2015	ACM-IEEE CS George Michael Memorial HPC Fellowship Maciej Besta received the prestigious ACM/IEEE-CS George Michael Memorial HPC Fellowship	Austin, TX, USA
2013	ACM Student Research Competition Robert Gerstenberger received the ACM SRC Bronze Medal (undergraduate)	Denver, CO, USA
2016	SNSF Ambizione Fellowship Tobias Grosser received the prestigious Ambizione Fellowship to build a research group	Bern, Switzerland

Positions and Experience

- 2023–present **Chief Architect for Machine Learning** **Swiss National Supercomputing Center (CSCS)**
Design and architect the large-scale training and inference platform infrastructure for scientific machine learning at CSCS.
- 2020–present **Full Professor of Computer Science** **ETH Zürich**
Computer Science Department
I lead research on scalable parallel computing, advising PhD and Master students in the Scalable Parallel Computing Laboratory.
- 2020–present **Adjunct Professor of Electrical Engineering** **ETH Zürich**
Department of Information Technology and Electrical Engineering
- 2017–2020 **Associate Professor (tenured) of Computer Science** **ETH Zürich**
Computer Science Department
- 2012–2017 **Assistant Professor (tenure track) of Computer Science** **ETH Zürich**
Computer Science Department
- 2010–2013 **Adjunct Assistant Professor of Computer Science** **University of Illinois Urbana-Champaign**
Computer Science Department
I led research in high-performance computing involving CS faculty members focused on topology mapping [ICS'11] and performance modeling [SC'11]. I taught two classes on High-Performance Computing.
- 2012 **Interim Technical Program Manager Applications** **University of Illinois Urbana-Champaign**
Blue Waters Directorate, NCSA
I led the Advanced Application and User Support Group, consisting of 11 domain specialists at Masters or Ph.D. level who provide advanced scientific computing support to a small number of expert users (≈ 40) of Blue Waters in their respective domains. Also certification of application and system performance milestones during installation and bringup of Blue Waters.
- 2010–2012 **Application and System Performance Modeling and Simulation Lead** **University of Illinois Urbana-Champaign**
Blue Waters Directorate, NCSA
I performed Modeling and Simulation of Sustained Petaflop Applications for Blue Waters, MPI Forum Activities. Scientific advisors: Marc Snir, Bill Gropp.
- 2008–2010 **Postdoctoral Fellow** **Indiana University Bloomington, IN**
Open Systems Lab
Parallel Programming, Modelling and Network Research, MPI Forum Activities
Scientific advisor: Andrew Lumsdaine.
- 2006–2008 **Research Assistant** **Indiana University Bloomington, IN**
Open Systems Lab
Parallel Computing and Networking Research
- Jan 2007 **Visiting Researcher** **Commissariat à l'Énergie Atomique Bruyères-le-Châtel, France**
Direction des Applications Militaires (CEA-DAM)
Parallel Quantum-Mechanical Computations with ABINIT
- Dec 2005 **Visiting Researcher** **CINECA Casalecchio di Reno, Italy**
CINECA Consorzio Interuniversitario
Parallel Ab-Initio Quantum Mechanical Computations
- 2004–2006 **Research Assistant** **Chemnitz University of Technology Chemnitz, Germany**
Parallel Ab-Initio Quantum Mechanical Computations, Networking Research

Significant Leadership and Service Positions

2023–present	UEC Transport Working Group Co-chair Elected Co-chair of the Ultra Ethernet Transport Working Group	Ultra Ethernet Consortium
2013–present	ACM SIGHPC Executive Committee (4th term) Elected member of all Executive Committees of ACM SIGHPC 2013-2025, Special Interest Group on High Performance Computing. I am the only individual who has been on all SIGHPC executive committees since its foundation. As one of two elected members-at-large, I co-represent the body of more than 1,000 members. I was (re-)elected in 2013, 2016, 2019, and 2022.	SIGHPC
2014–2018	Associate Editor (2 terms) IEEE Transactions on Parallel and Distributed Systems, I was re-appointed in 2016	IEEE TPDS
2012–present	Associate Editor SAGE International Journal of High Performance Computing Applications	IJHPCA
2010–present	MPI Forum WG Lead I lead the MPI-3 Working Group for Collective Operations and Topology.	Message Passing Interface Forum
2013	Expert in Resilience and Software Engineering Invited member of two working groups in the European Exascale Software Initiative 2 to “ <i>provide recommendations on strategic European actions [...]</i> ”	EESI2
2014–2016	Scientific Advisory Board Simula Research Laboratory, Norway	Simula
2014–2018	Scientific Advisory Board Member of the SAB of the European Project for Exascale ProGRAMming Models (EPIGRAM)	EPIGRAM Project

Professional Memberships

ACM Fellow	Association of Computing Machinery
IEEE Fellow	Institute of Electrical and Electronics Engineers
Member Academia Europaea	European Academy of Sciences Academia Europaea
ELLIS Member	European Laboratory for Learning and Intelligent Systems

Industry Experience and Significant Consulting

2016–today	Researcher & Consultant Research and development of advanced AI and quantum architectures and systems.	Microsoft, AI and Advanced Architecture Group
2023–today	Senior Fellow Research advice in High-Performance Computing, Artificial Intelligence, and Climate topics.	Abu Dhabi Investment Authority Labs
2019	Visiting Researcher (half year sabbatical) Research and development of advanced AI and quantum architectures and systems.	Microsoft, AI and Advanced Architecture Group
2018	Visiting Researcher (summer) Research and development of quantum architectures and systems.	Microsoft, Quantum Architecture Group
2018	Scientific Advisor External advisor to the late Paul Allan to help him select groups to fund a \$500M Climate Code Modernization Initiative.	Vulcan Inc., Seattle, WA
2018	Consultant Advising on advanced networking topics.	Cray Inc.
2017	Visiting Researcher (summer) Research and development of quantum architectures and systems.	Microsoft Research, Quantum Architecture Group
2000–2005	Software Engineer Design and Implementation of Database (Informix) and Web Applications (PHP/Perl)	DELTA proveris AG

Publications

Summary: 137 A^(*) top conference, 21 journal papers, 485+ total papers, 17,000+ citations, h-index: 70 (Google Scholar) 33 papers were nominated as best (student) paper and **18 received best (student) paper awards**. See below for details. Full publication list available at <http://hlor.inf.ethz.ch/publications/>

Selected Peer-reviewed Conference Publications

- Big Data'23** Wei Qiu, Marcin Copik, Yun Wang, Alexandru Calotoiu, Torsten Hoefler: User-guided Page Merging for Memory Deduplication in Serverless Systems *In 2023 IEEE International Conference on Big Data (Big Data), Dec. 2023 (acceptance rate 17.5%, 92/526)*
- SC'23** Marcin Chrapek, Mikhail Khalilov, Torsten Hoefler: HEAR: Homomorphically Encrypted Allreduce *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'23), presented in Denver, CO, USA, Association for Computing Machinery, ISBN: 979-8-400701-09-2, Nov. 2023 (acceptance rate 23.9%, 90/376)*, **SC23 Best student paper, SC23 Reproducibility Advancement Award**
- SC'23** Maciej Besta, Paweł Renc, Robert Gerstenberger, Paolo Sylos Labini, Alexandros Ziogas, Tiancheng Chen, Lukas Gianinazzi, Florian Scheidl, Kalman Szenes, Armon Carigiet, Patrick Iff, Grzegorz Kwasniewski, Raghavendra Kanakagiri, Chio Ge, Sammy Jaeger, Jarosław Wąs, Flavio Vella, Torsten Hoefler: High-Performance and Programmable Attentional Graph Neural Networks with Global Tensor Formulations *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'23), presented in Denver, CO, USA, Association for Computing Machinery, ISBN: 979-8-400701-09-2, Nov. 2023 (acceptance rate 23.9%, 90/376)*
- SC'23** Maciej Besta, Robert Gerstenberger, Marc Fischer, Michał Podstawski, Nils Blach, Berke Egeli, Georgy Mitenkov, Wojciech Chlapek, Marek Michalewicz, Hubert Niewiadomski, Jürgen Müller, Torsten Hoefler: The Graph Database Interface: Scaling Online Transactional and Analytical Graph Workloads to Hundreds of Thousands of Cores *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'23), presented in Denver, CO, USA, Association for Computing Machinery, ISBN: 979-8-400701-09-2, Nov. 2023 (acceptance rate 23.9%, 90/376)*, **Best Paper Finalist**
- SC'23** Philipp Schaad, Timo Schneider, Tal Ben-Nun, Alexandros Nikolaos Ziogas, Alexandru Calotoiu, Torsten Hoefler: FuzzyFlow: Leveraging Dataflow To Find and Squash Program Optimization Bugs *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'23), Association for Computing Machinery, ISBN: 979-8-400701-09-2, Nov. 2023 (acceptance rate 23.9%, 90/376)*
- ACM CSUR** Maciej Besta, Robert Gerstenberger, Emanuel Peter, Marc Fischer, Michał Podstawski, Claude Barthels, Gustavo Alonso, Torsten Hoefler: Demystifying Graph Databases: Analysis and Taxonomy of Data Organization, System Designs, and Graph Queries *ACM Comput. Surv.. Vol 56, Nr. 2, Association for Computing Machinery, ISSN: 0360-0300, Sep. 2023*
- DAC'23** Patrick Iff, Maciej Besta, Matheus Cavalcante, Tim Fischer, Luca Benini, Torsten Hoefler: HexaMesh: Scaling to Hundreds of Chiplets with an Optimized Chiplet Arrangement *In Proceedings of the 60th Annual Design Automation Conference, Jul. 2023*
- ATC'23** Andrei Ivanov, Benjamin Rothenberger, Arnaud Dethise, Marco Canini, Torsten Hoefler, Adrian Perrig: SAGE: Software-based Attestation for GPU Execution *In 2023 USENIX Annual Technical Conference (USENIX ATC 23), pages 485–499, USENIX Association, ISBN: 978-1-939133-35-9, Jul. 2023*
- DAC'23** Patrick Iff, Maciej Besta, Matheus Cavalcante, Tim Fischer, Luca Benini, Torsten Hoefler: Sparse Hamming Graph: A Customizable Network-on-Chip Topology *In Proceedings of the 60th Annual Design Automation Conference, Jul. 2023*
- ICS'23** Marcin Copik, Roman Böhringer, Alexandru Calotoiu, Torsten Hoefler: FMI: Fast and Cheap Message Passing for Serverless Functions *Jun. 2023*
- IEEE TPDS** Maciej Besta, Marc Fischer, Vasiliki Kalavri, Michael Kapralov, Torsten Hoefler: Practice of Streaming Processing of Dynamic Graphs: Concepts, Models, and Systems *IEEE Transactions of Parallel and Distributed Systems. Vol 34, Nr. 6, pages 1860–1876, IEEE, Jun. 2023*
- HPDC'23** Tiziano De Matteis, Lukas Gianinazzi, Johannes de Fine Licht, Torsten Hoefler: Streaming Task Graph Scheduling for Dataflow Architectures *Jun. 2023*

- CIAC'23** Tal Ben-Nun, Lukas Gianinazzi, Torsten Hoefler, Yishai Oltchik: Maximum Flows in Parametric Graph Templates *In Algorithms and Complexity - 13th International Conference, Jun. 2023*
- ICS'23** Lukas Truemper, Tal Ben-Nun, Philipp Schaad, Alexandru Calotoiu, Torsten Hoefler: Performance Embeddings: A Similarity-Based Transfer Tuning Approach to Performance Optimization *Jun. 2023*
- ICLR'23** Langwen Huang, Torsten Hoefler: Compressing multidimensional weather and climate data into neural networks *In The Eleventh International Conference on Learning Representations, May 2023, **Notable Top 5% (Spotlight)***
- ICLR'23** Elias Frantar, Saleh Ashkboos, Torsten Hoefler, Dan Alistarh: GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers *In The Eleventh International Conference on Learning Representations, May 2023*
- CGO'23** Tal Ben-Nun, Berke Ates, Alexandru Calotoiu, Torsten Hoefler: Bridging Control-Centric and Data-Centric Optimization *In 2023 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), pages 173-185, Feb. 2023*
- NeurIPS'22** Nikoli Dryden, Torsten Hoefler: Spatial Mixture-of-Experts *In Advances in Neural Information Processing Systems 35, presented in New Orleans, Louisiana, Dec. 2022*
- SIGMETRICS** Daniele De Sensi, Tiziano De Matteis, Konstantin Taranov, Salvatore Di Girolamo, Tobias Rahn, Torsten Hoefler: Noise in the Clouds: Influence of Network Performance Variability on Application Scalability *Proc. ACM Meas. Anal. Comput. Syst.. Vol 6, Nr. 3, presented in New York, NY, USA, Association for Computing Machinery, Dec. 2022*
- NeurIPS'22** Saleh Ashkboos, Langwen Huang, Nikoli Dryden, Tal Ben-Nun, Peter Dueben, Lukas Gianinazzi, Luca Kummer, Torsten Hoefler: ENS-10: A Dataset For Post-Processing Ensemble Weather Forecasts *In Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, presented in New Orleans, Louisiana, Dec. 2022*
- SC'22** Alexandros Nikolaos Ziogas, Grzegorz Kwasniewski, Tal Ben-Nun, Timo Schneider, Torsten Hoefler: Deinsum: Practically I/O Optimal Multilinear Algebra *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), Nov. 2022*
- SC'22** Tal Ben-Nun, Linus Groner, Florian Deconinck, Tobias Wicky, Eddie Davis, Johann Dahm, Oliver Elbert, Rhea George, Jeremy McGibbon, Lukas Trümper, Elynn Wu, Oliver Fuhrer, Thomas Schulthess, Torsten Hoefler: Productive Performance Engineering for Weather and Climate Modeling with Python *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), Nov. 2022*
- SC'22** Salvatore Di Girolamo, Daniele De Sensi, Konstantin Taranov, Milos Malesevic, Maciej Besta, Timo Schneider, Severin Kistler, Torsten Hoefler: Building Blocks for Network-Accelerated Distributed File Systems *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), Nov. 2022, **Best Paper Finalist***
- CCS'22** Konstantin Taranov, Benjamin Rothenberger, Daniele De Sensi, Adrian Perrig, Torsten Hoefler: NeVerMore: Exploiting RDMA Mistakes in NVMe-oF Storage Applications *In Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22), Nov. 2022, **Best Paper Honorable Mention***
- SC'22** Philipp Schaad, Tal Ben-Nun, Torsten Hoefler: Boosting Performance Optimization with Interactive Data Movement Visualization *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), ISBN: 9784665454445, Nov. 2022*
- SC'22** Shigang Li, Kazuki Osawa, Torsten Hoefler: Efficient Quantized Sparse Matrix Operations on Tensor Cores *Nov. 2022, **Best Paper Finalist***
- SC'22** Maciej Besta, Cesare Miglioli, Paolo Sylos Labini, Jakub Tětek, Patrick Iff, Raghavendra Kanakagiri, Saleh Ashkboos, Kacper Janda, Michal Podstawski, Grzegorz Kwasniewski, Niels Gleinig, Flavio Vella, Onur Mutlu, Torsten Hoefler: ProbGraph: High-Performance and High-Accuracy Graph Mining with Probabilistic Set Representations *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), Nov. 2022, **SC22 Best Paper (1/82)***
- SC'22** Kartik Lakhotia, Maciej Besta, Laura Monroe, Kelly Isham, Patrick Iff, Torsten Hoefler, Fabrizio Petrini: PolarFly: A Cost-Effective and Flexible Low-Diameter Topology *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22), Nov. 2022*

- SC'22** Torsten Hoefler, Tommaso Bonato, Daniele De Sensi, Salvatore Di Girolamo, Shigang Li, Marco Heddes, Jon Belk, Deepak Goel, Miguel Castro, Steve Scott: HammingMesh: A Network Topology for Large-Scale Deep Learning *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'22)*, Nov. 2022, **SC22 Reproducibility Advancement Award**
- ICCAD'22** Carl-Johannes Johnsen, Tiziano De Matteis, Tal Ben-Nun, Johannes de Fine Licht, Torsten Hoefler: Temporal Vectorization: A Compiler Approach to Automatic Multi-Pumping *In 2022 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, Oct. 2022
- LOG'22** Maciej Besta, Patrick Iff, Florian Scheidl, Kazuki Osawa, Nikoli Dryden, Michal Podstawski, Tiancheng Chen, Torsten Hoefler: Neural Graph Databases *In Proceedings of the Learning on Graphs Conference (LOG'22)*, Sep. 2022
- IEEE Computer** Torsten Hoefler: Benchmarking data science: Twelve ways to lie with statistics and performance on parallel computers *IEEE Computer*. Vol 55, pages 49-56, Aug. 2022, **Cover Feature Research Reproducibility**
- KDD'22** Maciej Besta, Raphael Grob, Cesare Miglioli, Nicola Bernold, Grzegorz Kwasniewski, Gabriel Gjini, Raghavendra Kanakagiri, Saleh Ashkboos, Lukas Gianinazzi, Nikoli Dryden, Torsten Hoefler: Motif Prediction with Graph Neural Networks *In Proceedings of the 28th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'22)*, Aug. 2022
- IEEE Computer** Torsten Hoefler, Ariel Hendel, Duncan Roweth: The Convergence of Hyperscale Data Center and High-Performance Computing Networks *IEEE Computer*. Vol 55, Nr. 7, pages 29-37, Jul. 2022, **Cover Feature Technology Predictions**
- ICS'22** Oliver Rausch, Tal Ben-Nun, Nikoli Dryden, Andrei Ivanov, Shigang Li, Torsten Hoefler: A Data-Centric Optimization Framework for Machine Learning *In Proceedings of the 2022 International Conference on Supercomputing (ICS'22)*, Jul. 2022
- ICS'22** Alexandru Calotoiu, Tal Ben-Nun, Grzegorz Kwasniewski, Johannes de Fine Licht, Timo Schneider, Philipp Schaad, Torsten Hoefler: Lifting C Semantics for Dataflow Optimization *In Proceedings of the 2022 International Conference on Supercomputing (ICS'22)*, Jul. 2022
- ICS'22** Larissa Schmid, Marcin Copik, Alexandru Calotoiu, Dominik Werle, Andreas Reiter, Michael Selzer, Anne Kozirolek, Torsten Hoefler: Performance-Detective: Automatic Deduction of Cheap and Accurate Performance Models *In Proceedings of the 2022 International Conference on Supercomputing (ICS'22)*, Jul. 2022
- SIGMOD'22** Konstantin Taranov, Steve Blyan, Virendra Marathe, Torsten Hoefler: KafkaDirect: Zero-copy Data Access for Apache Kafka over RDMA Networks *In Proceedings of the 2022 ACM SIGMOD International Conference on Management of Data*, Jun. 2022
- IPDPS'22** András Strausz, Flavio Vella, Salvatore Di Girolamo, Maciej Besta, Torsten Hoefler: Asynchronous Distributed-Memory Triangle Counting and LCC with RMA Caching *In Proceedings of the 36th IEEE International Parallel and Distributed Processing Symposium*, Jun. 2022
- ICST'22** Andrei Lascu and Alastair F. Donaldson and Tobias Grosser and Torsten Hoefler: Metamorphic Fuzzing of C++ Libraries *In IEEE International Conference on Software Testing, Verification and Validation*, Jun. 2022
- IPDPS'22** Niels Gleinig, Maciej Besta, Torsten Hoefler: I/O-Optimal Cache-Oblivious Sparse Matrix-Sparse Matrix Multiplication *In Proceedings of the 36th IEEE International Parallel and Distributed Processing Symposium*, Jun. 2022
- FCCM'22** Johannes de Fine Licht, Christopher A. Pattison, Alexandros Nikolaos Ziogas, David Simmons-Duffin, Torsten Hoefler: Fast Arbitrary Precision Floating Point on FPGA *In Proceedings of the 30th IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM'22)*, May 2022
- PPoPP'22** Shigang Li, Torsten Hoefler: Near-Optimal Sparse Allreduce for Distributed Deep Learning *In Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Apr. 2022
- ICLR'22** Bryan A. Plummer, Nikoli Dryden, Julius Frost, Torsten Hoefler, Kate Saenko: Neural Parameter Allocation Search *In Tenth International Conference on Learning Representations*, Apr. 2022
- DATE** Andrea Cossettini, Konstantin Taranov, Christian Vogt, Michele Magno, Torsten Hoefler, Luca Benini: A RDMA Interface for Ultra-Fast Ultrasound Data-Streaming over an Optical Link *In Proceedings of Design, Automation, and Test in Europe (DATE)*, 2022
- DAC'21** Niels Gleinig, Torsten Hoefler: An Efficient Algorithm for Sparse Quantum State Preparation *In Proceedings of the 58th Annual Design Automation Conference, presented in San Francisco, CA, USA, ACM, Dec. 2021 (acceptance rate 23%)*

- Middleware 2021** Marcin Copik, Grzegorz Kwasniewski, Maciej Besta, Michal Podstawski, Torsten Hoefler: SeBS: A Serverless Benchmark Suite for Function-as-a-Service Computing *In Proceedings of the 22nd International Middleware Conference, presented in Québec city, Canada, ACM, ISBN: 9781450385343, Dec. 2021*
- SC21** Thomas Häner, Damian S. Steiger, Torsten Hoefler, Matthias Troyer: Distributed Quantum Computing with QMPI *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), Nov. 2021 (acceptance rate 25.9%, 98/379)*
- SC21** Daniele De Sensi, Salvatore Di Girolamo, Saleh Ashkboos, Shigang Li, Torsten Hoefler: Flare: Flexible In-Network Allreduce *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), presented in St. Louis, Missouri, ACM, Nov. 2021 (acceptance rate 25.9%, 98/379)*
- SC21** Nikoli Dryden, Roman Böhringer, Tal Ben-Nun, Torsten Hoefler: Clairvoyant Prefetching for Distributed Machine Learning I/O *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), presented in St. Louis, Missouri, ACM, Nov. 2021 (acceptance rate 25.9%, 98/379)*
- SC21** Grzegorz Kwasniewski, Marko Kabić, Tal Ben-Nun, Alexandros Nikolaos Ziogas, Jens Eirik Saethre, André Gaillard, Timo Schneider, Maciej Besta, Anton Kozhevnikov, Joost VandeVondele, Torsten Hoefler: On the Parallel I/O Optimality of Linear Algebra Kernels: Near-Optimal Matrix Factorizations *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), Nov. 2021 (acceptance rate 25.9%, 98/379)*
- SC21** Shigang Li, Torsten Hoefler: Chimera: Efficiently Training Large-Scale Neural Networks with Bidirectional Pipelines *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), presented in St. Louis, Missouri, ACM, Nov. 2021 (acceptance rate 25.9%, 98/379), Best Paper Finalist*
- SC21** Alexandros Nikolaos Ziogas, Timo Schneider, Tal Ben-Nun, Alexandru Calotiu, Tiziano De Matteis, Johannes de Fine Licht, Luca Lavarini, Torsten Hoefler: Productivity, Portability, Performance: Data-Centric Python *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), Nov. 2021 (acceptance rate 25.9%, 98/379)*
- OOPSLA'21** Arjun Pitchanathan, Christian Ulmann, Michel Weber, Torsten Hoefler, Tobias Grosser: FPL: fast Presburger arithmetic through transprecision *OOPSLA '21: Proceedings of the ACM international conference on Object oriented programming systems languages and applications. ACM, Nov. 2021, OOPSLA distinguished paper award (6/71)*
- MICRO'21** Maciej Besta, Raghavendra Kanakagiri, Grzegorz Kwasniewski, Rachata Ausavarungnirun, Jakub Beránek, Konstantinos Kanellopoulos, Kacper Janda, Zur Vonarburg-Shmaria, Lukas Gianinazzi, Ioana Stefan, Juan Gómez Luna, Marcin Copik, Lukas Kapp-Schwoerer, Salvatore Di Girolamo, Nils Blach, Marek Konieczny, Onur Mutlu, Torsten Hoefler: SISA: Set-Centric Instruction Set Architecture for Graph Mining on Processing-in-Memory Systems *In Proceedings of the 54th IEEE/ACM International Symposium on Microarchitecture (MICRO), Oct. 2021*
- JMLR** Torsten Hoefler, Dan Alistarh, Tan Ben-Nun, Nikoli Dryden, Alexandra Peste: Sparsity in Deep Learning: Pruning and growth for efficient inference and training in neural networks *Journal of Machine Learning Research. Vol 22, Nr. 241, pages 1-124, Sep. 2021*
- VLDB'21** Maciej Besta, Zur Vonarburg-Shmaria, Yannick Schaffner, Leonardo Schwarz, Grzegorz Kwasniewski, Lukas Gianinazzi, Jakub Beranek, Kacper Janda, Tobias Holenstein, Sebastian Leisinger, Peter Tatkowski, Esref Ozdemir, Adrian Balla, Marcin Copik, Philipp Lindenberger, Pavel Kalvoda, Marek Konieczny, Onur Mutlu, Torsten Hoefler: GraphMineSuite: Enabling High-Performance and Programmable Graph Mining Algorithms with Set Algebra *In Proceedings of the 47th International Conference on Very Large Data Bases (VLDB'21), Aug. 2021*
- TQC'21** David Ittah, Thomas Häner, Vadym Kliuchnikov, Torsten Hoefler: QIRO: A Static Single Assignment-Based Quantum Program Representation for Optimization *In ACM Transactions on Quantum Computing, Association for Computing Machinery, ISSN: 2643-6809, Aug. 2021*
- ICML'21** Chris Cummins, Zacharias V. Fisches, Tal Ben-Nun, Torsten Hoefler, Michael O'Boyle, Hugh Leather: ProGraML: A Graph-based Program Representation for Data Flow Analysis and Compiler Optimizations *In Thirty-eighth International Conference on Machine Learning, presented in Virtual, PMLR, Jul. 2021 (acceptance rate 21%)*
- USENIX ATC'21** Maksym Planeta, Jan Bierbaum, Leo Sahaya Daphne Antony, Torsten Hoefler, Hermann Härtig: MigrOS: Transparent Live-Migration Support for Containerised RDMA Applications *In Proceedings of the 2021 USENIX Annual Technical Conference, USENIX, Jul. 2021 (acceptance rate 18.8%, 64/341)*

- USENIX ATC'21** Konstantin Taranov, Rodrigo Bruno, Gustavo Alonso, Torsten Hoefler: Naos: Serialization-free RDMA networking in Java *In Proceedings of the 2021 USENIX Annual Technical Conference, USENIX, Jul. 2021 (acceptance rate 18.8%, 64/341)*
- SPAA'21** Grzegorz Kwasniewski, Tal Ben-Nun, Lukas Gianinazzi, Alexandru Calotoiu, Timo Schneider, Alexandros Nikolaos Ziogas, Maciej Besta, Torsten Hoefler: Pebbles, Graphs, and a Pinch of Combinatorics: Towards Tight I/O Lower Bounds for Statically Analyzable Programs *In Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'21), Jul. 2021 (acceptance rate 14.9%)*
- SPAA'21** Lukas Gianinazzi, Maciej Besta, Yannick Schaffner, Torsten Hoefler: Parallel Algorithms for Finding Large Cliques in Sparse Graphs *In Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'21), ACM, Jul. 2021*
- IEEE TPDS** Marcin Copik, Tobias Grosser, Torsten Hoefler, Paolo Bientinesi, Benjamin Berkels: Work-stealing prefix scan: Addressing load imbalance in large-scale image registration *IEEE Transactions on Parallel and Distributed Systems (TPDS). IEEE, Jul. 2021 to appear*
- ISCA'21** Salvatore Di Girolamo, Andreas Kurth, Alexandru Calotoiu, Thomas Benz, Timo Schneider, Jakub Beránek, Luca Benini, Torsten Hoefler: A RISC-V in-network accelerator for flexible high-performance low-power packet processing *In Proceedings of the 48th Annual International Symposium on Computer Architecture (ISCA'21), Jun. 2021*
- ICS'21** Alexandros Nikolaos Ziogas, Tal Ben-Nun, Timo Schneider, Torsten Hoefler: NPbench: A Benchmarking Suite for High-Performance NumPy *In Proceedings of the 2021 International Conference on Supercomputing (ICS'21), Jun. 2021*
- SIGMOD'21** Konstantin Taranov, Salvatore Di Girolamo, Torsten Hoefler: CoRM: Compactable Remote Memory over RDMA *In Proceedings of the 2021 ACM SIGMOD International Conference on Management of Data, Jun. 2021*
- IPDPS'21** Marcus Ritter, Alexander Geiss, Johannes Wehrstein, Alexandru Calotoiu, Thorsten Reimann, Torsten Hoefler, Felix Wolf: Noise-Resilient Empirical Performance Modeling with Deep Neural Networks *In IPDPS '21: Proceedings of the 35th IEEE International Parallel and Distributed Processing Symposium, May 2021*
- MLSys'21** Andrei Ivanov, Nikoli Dryden, Tal Ben-Nun, Shigang Li, Torsten Hoefler: Data Movement Is All You Need: A Case Study on Optimizing Transformers *In Proceedings of Machine Learning and Systems 3 (MLSys 2021), Apr. 2021 (acceptance rate: 23.5% (52/221)), Outstanding Paper Award (5/52)*
- Nature CompSci** Peter Bauer, Peter D. Dueben, Torsten Hoefler, Tiago Quintino, Thomas C. Schulthess, Nils P. Wedi: The digital revolution of Earth-system science *Nature Computational Science. Vol 1, Nr. 1, pages 104-113, Feb. 2021*
- RSTA** Peter Grönquist, Chengyuan Yao, Tal Ben-Nun, Nikoli Dryden, Peter Dueben, Shigang Li, Torsten Hoefler: Deep Learning for Post-Processing Ensemble Weather Forecasts *Philosophical Transactions of the Royal Society A. Vol 379, Nr. 2194, The Royal Society, Feb. 2021*
- PPoPP'21** Marcin Copik, Alexandru Calotoiu, Tobias Grosser, Nicolas Wicki, Felix Wolf, Torsten Hoefler: Extracting Clean Performance Models from Tainted Programs *In PPoPP '21: Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Feb. 2021 (acceptance rate: 21% (31/150))*
- TPDS** Maciej Besta and Jens Domke and Marcel Schneider and Marek Konieczny and Salvatore Di Girolamo and Timo Schneider and Ankit Singla and Torsten Hoefler: High-Performance Routing with Multipathing and Path Diversity in Ethernet and HPC Networks *In IEEE Transactions of Parallel and Distributed Systems, 2021*
- IEEE TPDS** Shigang Li, Tal Ben-Nun, Giorgi Nadiradze, Salvatore Di Girolamo, Nikoli Dryden, Dan Alistarh, Torsten Hoefler: Breaking (Global) Barriers in Parallel Stochastic Optimization with Wait-Avoiding Group Averaging *IEEE Transactions on Parallel and Distributed Systems (TPDS). Vol 32, Nr. 7, pages 1725-1739, IEEE, 2021*
- USENIX Security'21** Benjamin Rothenberger, Konstantin Taranov, Adrian Perrig, Torsten Hoefler: ReDMark: Bypassing RDMA Security Mechanisms *In Proceedings of the 2021 USENIX Security Symposium, USENIX, 2021*
- CGO'21** Johannes de Fine Licht, Andreas Kuster, Tiziano De Matteis, Tal Ben-Nun, Dominic Hofer, Torsten Hoefler: StencilFlow: Mapping Large Stencil Programs to Distributed Spatial Computing Systems *In Proceedings of the 19th ACM/IEEE International Symposium on Code Generation and Optimization (CGO'21), 2021*
- IEEE TPDS** Johannes de Fine Licht, Maciej Besta, Simon Meierhans, Torsten Hoefler: Transformations of High-Level Synthesis Codes for High-Performance Computing *IEEE Transactions on Parallel and Distributed Systems (TPDS). 2021*

- TACO21** Tobias Gysi, Christoph Müller, Oleksandr Zinenko, Stephan Herhut, Eddie Davis, Tobias Wicky, Oliver Fuhrer, Torsten Hoefler, Tobias Grosser: Domain-Specific Multi-Level IR Rewriting for GPU: The Open Earth Compiler for GPU-Accelerated Climate Simulation *ACM Trans. Archit. Code Optim.*. Vol 18, Nr. 4, Association for Computing Machinery, ISSN: 1544-3566, 2021
- DATE** Paul Scheffler, Florian Zaruba, Fabian Schuiki, Torsten Hoefler, Luca Benini: Indirection Stream Semantic Register Architecture for Efficient Sparse-Dense Linear Algebra *In Proceedings of Design, Automation, and Test in Europe (DATE)*, 2021
- SC20** Yuyang Jin, Haojie Wang, Teng Yu, Xiongchao Tang, Torsten Hoefler, Xu Liu, Jidong Zhai: SCALANA: Automating Scaling Loss Detection with Graph Analysis *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, Nov. 2020 (acceptance rate 25.1% (95/378))
- SC20** Daniele De Sensi and Salvatore Di Girolamo and Kim H. McMahon and Duncan Roweth and Torsten Hoefler: An In-Depth Analysis of the Slingshot Interconnect *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, Nov. 2020 (acceptance rate: 25.1% (95/378))
- IEEE TCAD** Asif Ali Khan, Hauke Mewes, Tobias Grosser, Torsten Hoefler, Jeronimo Castrillon: Polyhedral Compilation for Racetrack Memories *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. Vol 39, Nr. 11, IEEE, Nov. 2020
- OOPSLA'20** Tobias Grosser, Theodoros Theodoridis, Maxmilian Falkenstein, Arjun Pitchanathan, Michael Kruse, Manuel Rigger, Zhendong Su, Torsten Hoefler: Fast Linear Programming through Transprecision Computing on Small and Sparse Data *OOPSLA'20: Proceedings of the ACM international conference on Object oriented programming systems languages and applications*. ACM, Nov. 2020
- OOPSLA'20** Thomas Häner, Matthias Troyer, Torsten Hoefler: Assertion-based optimization of quantum programs *OOPSLA '20: Proceedings of the ACM international conference on Object oriented programming systems languages and applications*. ACM, Nov. 2020
- SC20** Maciej Besta and Armon Carigiet and Kacper Janda and Zur Vonarburg-Shmaria and Lukas Gianinazzi and Torsten Hoefler: High-Performance Parallel Graph Coloring with Strong Guarantees on Work, Depth, and Quality *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, Nov. 2020 (acceptance rate: 25.1% (95/378))
- SC20** Maciej Besta and Marcel Schneider and Marek Konieczny and Karolina Cynk and Erik Henriksson and Salvatore Di Girolamo and Ankit Singla and Torsten Hoefler: FatPaths: Routing in Supercomputers and Data Centers when Shortest Paths Fall Short *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, Nov. 2020 (acceptance rate: 25.1% (95/378))
- SC20** Tiziano De Matteis and Johannes de Fine Licht and Torsten Hoefler: FBLAS: Streaming Linear Algebra on FPGA *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, Nov. 2020 (acceptance rate: 25.1% (95/378))
- IEEE TOC** Florian Zaruba, Fabian Schuiki, Torsten Hoefler, Luca Benini: Snitch: A tiny Pseudo Dual-Issue Processor for Area and Energy Efficient Execution of Floating-Point Intensive Workloads *IEEE Transactions on Computers (TOC)*. IEEE, Sep. 2020, **Featured Paper in November 2021 issue**
- VLDB'20** Claude Barthels, Ingo Müller, Konstantin Taranov, Torsten Hoefler, Gustavo Alonso: Strong consistency is not hard to get: TwoPhase Locking and TwoPhase Commit on Thousands of Cores *In Proceedings of the VLDB Endowment*, Vol. 12, No. 13, VLDB Endowment, Sep. 2020
- USENIX ATC'20** Konstantin Taranov, Benjamin Rothenberger, Adrian Perrig, Torsten Hoefler: sRDMA – Efficient NIC-based Authentication and Encryption *In Proceedings of the 2020 USENIX Annual Technical Conference, USENIX, Jul. 2020*, (acceptance rate 18.6%, 65/348)
- SPAA'20** Lukas Gianinazzi, Torsten Hoefler: Parallel Planar Subgraph Isomorphism and Vertex Connectivity *In Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'20)*, ACM, Jul. 2020
- DAC'20** Andreas Kurth, Samuel Riedel, Florian Zaruba, Torsten Hoefler, Luca Benini: ATUNs: Modular and Scalable Support for Atomic Operations in a Shared Memory Multiprocessor *In Proceedings of the 57th Annual Design Automation Conference, ACM, Jun. 2020*, **Best Paper Finalist (6/228)**
- CVPR'20** Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefler, Daniel Soudry: Increasing batch size through instance repetition improves generalization *In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020

- IPDPS'20** Marcus Ritter, Alexandru Calotoiu, Thorsten Reimann, Torsten Hoefler, Felix Wolf: Performance Modeling at a Discount *presented in New Orleans, LA, USA, IEEE, May 2020, Accepted at the 34th IEEE International Parallel and Distributed Processing Symposium (IPDPS'20)*
- IPDPS'20** Maciej Besta, Raghavendra Kanakagiri, Harun Mustafa, Mikhail Karasikov, Gunnar Rättsch, Torsten Hoefler, Edgar Solomonik: Communication-Efficient Jaccard Similarity for High-Performance Distributed Genome Comparisons *May 2020, In Proceedings of the 34th IEEE International Parallel and Distributed Processing Symposium*
- FPGA'20** Johannes de Fine Licht, Grzegorz Kwasniewski, Torsten Hoefler: Flexible Communication Avoiding Matrix Multiplication on FPGA with High-Level Synthesis *Feb. 2020, In Proceedings of the 28th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*
- PPoPP'20** Shigang Li, Tal Ben-Nun, Salvatore Di Girolamo, Dan Alistarh, Torsten Hoefler: Taming Unbalanced Training Workloads in Deep Learning with Partial Collective Operations *In Proceedings of the 25th Symposium on Principles and Practice of Parallel Programming (PPoPP'20), Feb. 2020, (acceptance rate: 23.1% (28/121))*
- SC19** Cedric Renggli, Dan Alistarh, Mehdi Aghagolzadeh, Torsten Hoefler: SparCML: High-Performance Sparse Communication for Machine Learning *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Tiziano De Matteis, Johannes de Fine Licht, Jakub Beránek, Torsten Hoefler: Streaming Message Interface: High-Performance DistributedMemory Programming on Reconfigurable Hardware *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, Torsten Hoefler: Optimizing the Data Movement in Quantum Transport Simulations via Data-Centric Parallel Programming *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, Torsten Hoefler: A Data-Centric Approach to Extreme-Scale Ab initio Dissipative Quantum Transport Simulations *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019, **Gordon Bell Prize Winner***
- SC19** Salvatore Di Girolamo, Konstantin Taranov, Andreas Kurth, Michael Schaffner, Timo Schneider, Jakub Beranek, Maciej Besta, Luca Benini, Duncan Roweth, Torsten Hoefler: Network-Accelerated Non-Contiguous Memory Transfers *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Daniele De Sensi, Salvatore Di Girolamo, Torsten Hoefler: Mitigating Network Noise on Dragonfly Networks through Application-Aware Routing *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Tal Ben-Nun, Johannes de Fine Licht, Alexandros Nikolaos Ziogas, Timo Schneider, Torsten Hoefler: Stateful Dataflow Multigraphs: A Data-Centric Model for Performance Portability on Heterogeneous Architectures *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Maciej Besta, Simon Weber, Lukas Gianinazzi, Robert Gerstenberger, Andrey Ivanov, Yishai Oltchik, Torsten Hoefler: Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344))*
- SC19** Grzegorz Kwasniewski and Marko Kabić and Maciej Besta and Joost VandeVondele and Raffaele Solcà and Torsten Hoefler: Red-Blue Pebbling Revisited: Near Optimal Parallel Matrix-Matrix Multiplication *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), presented in Denver, CO, USA, ACM, Nov. 2019 (acceptance rate: 22.7% (78/344)), **SC19 Best Student Paper***
- PASC'19** Felix Thaler, Stefan Moosbrugger, Carlos Osuna, Mauro Bianco, Hannes Vogt, Anton Afanasyev, Lukas Mosimann, Oliver Fuhrer, Thomas Schulthess, Torsten Hoefler: Porting the COSMO Weather Model to Intel KNL *presented in Zurich, Switzerland, ACM, Jun. 2019, Accepted at the ACM Platform for Advanced Scientific Computing Conference (PASC19)*

- DAC'19** Niels Gleinig and Frances Ann Hubis and Torsten Hoefler: Embedding Functions Into Reversible Circuits: A Probabilistic Approach to the Number of Lines *Proceedings of the 56th Annual Design Automation Conference, presented in Las Vegas, NV, USA, ACM, ISBN: 978-1-4503-6725-7/19/06, Jun. 2019*
- PLDI'19** T. Gysi, T. Grosser, L. Brandner, T. Hoefler: A Fast Analytical Model of Fully Associative Caches *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, presented in Phoenix, AZ, USA, pages 816–829, ACM, ISBN: 978-1-4503-6712-7, Jun. 2019*
- ICS'19** Paul R. Eller, Torsten Hoefler, William Gropp: Using Performance Models to Understand Scalable Krylov Solver Performance at Scale for Structured Grid Problems *Proceedings of the 2019 ACM International Conference on Supercomputing (ICS'19), presented in Phoenix, AZ, ACM, Jun. 2019*
- IPDPS'19** S. Di Girolamo, P. Schmid, T. Schulthess, T. Hoefler: SimFS: A Simulation Data Virtualizing File System Interface *IEEE, May 2019, Accepted at the 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS'19)*
- IPDPS'19** T. Ben-Nun, M. Besta, S. Huber, A. N. Ziogas, D. Peter, T. Hoefler: A Modular Benchmarking Infrastructure for High-Performance and Reproducible Deep Learning *IEEE, May 2019, Accepted at the 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS'19)*
- PPoPP'19** Martin Kuettler, Maksym Planeta, Jan Bierbaum, Carsten Weinhold, Hermann Haertig, Amnon Barak, Torsten Hoefler: Corrected Trees for Reliable Group Communication *Feb. 2019, Accepted at The ACM Conference Principles and Practice of Parallel Programming 2019 (PPoPP'19) (acceptance rate: 19% (29/152))*
- FPGA'19** Maciej Besta, Marc Fischer, Tal Ben-Nun, Johannes De Fine Licht, Torsten Hoefler: Substream-Centric Maximum Matchings on FPGA *Feb. 2019, In Proceedings of the 27th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (acceptance rate: 23%), Best Paper Finalist (4/30)*
- NIPS'18** Tal Ben-Nun, Alice Shoshana Jakobovits, Torsten Hoefler: Neural Code Comprehension: A Learnable Representation of Code Semantics *In Advances in Neural Information Processing Systems 31, presented in Montreal, Canada, Curran Associates, Inc., Dec. 2018*
- NIPS'18** Dan Alistarh, Torsten Hoefler, Mikael Johansson, Sarit Khirirat, Nikola Konstantinov, Cedric Renggli: The Convergence of Sparsified Gradient Methods *In Advances in Neural Information Processing Systems 31, presented in Montreal, Canada, Curran Associates, Inc., Dec. 2018*
- PACT'18** M. Besta, D. Stanojevic, T. Zivic, J. Singh, M. Hoerold, T. Hoefler: Log(Graph): A Near-Optimal High-Performance Graph Representation *Limassol, Cyprus, ACM, Nov. 2018, Accepted at the 27th International Conference on Parallel Architectures and Compilation (PACT'18)*
- SC18** Heng Lin, Xiaowei Zhu, Bowen Yu, Xiongchao Tang, Wei Xue, Wenguang Chen, Lufei Zhang, Torsten Hoefler, Xiaosong Ma, Xin Liu, Weimin Zheng, Jingfang Xu: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC18), presented in Denver, CO, USA, ACM, Nov. 2018, Gordon Bell Award Finalist*
- Cluster'18** Y. Oyama, T. Ben-Nun, T. Hoefler, S. Matsuoka: Accelerating Deep Learning Frameworks with Micro-batches *In IEEE International Conference on Cluster Computing, CLUSTER 2018, Belfast, UK, September 10-13, 2018, presented in Belfast, UK, IEEE, ISBN: 978-1-5386-8319-4, Sep. 2018, (28% (44/154))*
- Cluster'18** Alexandru Calotiu, Alexander Graf, Torsten Hoefler, Daniel Lorenz, Sebastian Rinke, Felix Wolf: Lightweight Requirements Engineering for Exascale Co-design *In IEEE International Conference on Cluster Computing, CLUSTER 2018, Belfast, UK, September 10-13, 2018, presented in Belfast, UK, IEEE, ISBN: 978-1-5386-8319-4, Sep. 2018, (28% (44/154))*
- EuroSys'18** K. Taranov, G. Alonso, T. Hoefler: Fast and strongly-consistent per-item resilience in key-value stores *Apr. 2018, EuroSys '18: Thirteenth EuroSys Conference 2018, April 23–26, 2018, Porto, Portugal (acceptance rate: 16% (43/262))*
- ASPLOS'18** M. Besta, S. M. Hassan, S. Yalamanchili, R. Ausavarungrun, O. Mutlu, T. Hoefler: Slim NoC: A Low-Diameter On-Chip Network Topology for High Energy Efficiency and Scalability *Mar. 2018, Accepted at the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'18)*
- PPoPP'18** Lukas Gianinazzi, Pavel Kalvoda, Alessandro De Palma, Maciej Besta, Torsten Hoefler: Communication-Avoiding Parallel Minimum Cuts and Connected Components *In Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, presented in Vienna, Austria, pages 403–404, ACM, ISBN: 978-1-4503-4982-6, Feb. 2018*

- PPoPP'18** J. de Fine Licht, M. Blott, T. Hoefler: Designing scalable FPGA architectures using high-level synthesis *In Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, presented in Vienna, Austria, pages 403–404, ACM, ISBN: 978-1-4503-4982-6, Feb. 2018*
- ICDE'18** Ingo Mueller, Andrea Arteaga, Torsten Hoefler, Gustavo Alonso: Reproducible Floating-Point Aggregation in RDBMSs *Feb. 2018, In Proceedings of the 2018 IEEE 34th International Conference on Data Engineering*
- SC17** E. Solomonik, M. Besta, F. Vella, T. Hoefler: Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC17), Nov. 2017, (acceptance rate: 18% (61/327))*
- SC17** T. Hoefler, S. Di Girolamo, K. Taranov, R. E. Grant, R. Brightwell: sPIN: High-performance streaming Processing in the Network *In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC17), Nov. 2017, (acceptance rate: 18% (61/327)) **Best Paper Finalist (5/61)***
- VLDB'17** C. Barthels, T. Schneider, I. Mueller, G. Alonso, T. Hoefler: Distributed Join Algorithms on Thousands of Cores *Vol 10, Nr. 5, In Proc. VLDB Endow., presented in Munich, Germany, pages 517–528, VLDB Endowment, ISSN: 2150-8097, Aug. 2017*
- HOTI'17** P. Yebenes, J. Escudero-Sahuquillo, P. J. Garcia, F. J. Quiles, T. Hoefler: Improving Non-Minimal and Adaptive Routing Algorithms in Slim Fly Networks *In Proceedings of the 25th Annual Symposium on High-Performance Interconnects (HOTI'17), Aug. 2017, **Best Student Paper***
- HOTI'17** T. Schneider, J. Dinan, M. Flajslik, K. D. Underwood, and T. Hoefler: Fast Networks and Slow Memories: A Mechanism for Mitigating Bandwidth Mismatches *In Proceedings of the 25th Annual Symposium on High-Performance Interconnects (HOTI'17), Aug. 2017*
- HPDC'17** M. Poke, T. Hoefler, C. W. Glass: AllConcur: Leaderless Concurrent Atomic Broadcast *In Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing (HPDC'17), presented in Washington, DC, USA, ACM, Jun. 2017, (acceptance rate: 19%)*
- HPDC'17** M. Besta, M. Podstawski, L. Groner, E. Solomonik, T. Hoefler: To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations *In Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing (HPDC'17), presented in Washington, DC, USA, ACM, Jun. 2017, (acceptance rate: 19%)*
- ICCS'17** A. Arteaga, O. Fuhrer, T. Hoefler, T. Schulthess: Model-Driven Choice of Numerical Methods for the Solution of the Linear Advection Equation *In Proceedings of the International Conference on Computational Science (ICCS'17), presented in Zurich, Switzerland, Elsevier, Jun. 2017*
- SPAA'17** E. Solomonik, G. Ballard, J. Demmel, T. Hoefler: A Communication-Avoiding Parallel Algorithm for the Symmetric Eigenvalue Problem *Nr. 11, In Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'17), presented in Washington, DC, USA, pages 111–121, ACM, ISBN: 978-1-4503-4593-4, Jun. 2017*
- IPDPS'17** M. Besta, F. Marending, E. Solomonik, T. Hoefler: SlimSell: A Vectorized Graph Representation for Breadth-First Search *In Proceedings of the 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17), presented in Orlando, FL, USA, IEEE, May 2017, (acceptance rate: 22%, 116/516)*
- IPDPS'17** S. Di Girolamo, F. Vella and T. Hoefler: Transparent Caching for RMA Systems *In Proceedings of the 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17), presented in Orlando, FL, USA, IEEE, May 2017, (acceptance rate: 22%, 116/516)*
- IPDPS'17** T. Hoefler, A. Barak, A. Shiloh and Z. Drezner: Corrected Gossip Algorithms for Fast Reliable Broadcast on Unreliable Systems *In Proceedings of the 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17), presented in Orlando, FL, USA, IEEE, May 2017, (acceptance rate: 22%, 116/516)*
- IPDPS'17** T. Wicky, E. Solomonik and T. Hoefler: Communication-Avoiding Parallel Algorithms for Solving Triangular Systems of Linear Equations *In Proceedings of the 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17), presented in Orlando, FL, USA, IEEE, May 2017, (acceptance rate: 22%, 116/516)*
- IPDPS'17** S. Ramos and T. Hoefler: Capability Models for Manycore Memory Systems: A Case-Study with Xeon Phi KNL *Proceedings of the 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17), presented in Orlando, FL, USA, IEEE, May 2017, (acceptance rate: 22%, 116/516)*
- PPoPP'17** S. Shudler, A. Calotoiu, T. Hoefler, F. Wolf: Isoefficiency in Practice: Configuring and Understanding the Performance of Task-based Applications *In Proceedings of the 22nd ACM SIGPLAN symposium on Principles and practice of parallel programming, presented in College Station, TX, ACM, Feb. 2017 (acceptance rate: 21%, 29/139)*

- OOPSLA'16** Andrei Marian Dan, Patrick Lam, T. Hoefler, and Martin Vechev: Modeling and Analysis of Remote Memory Access Programming *ACM Symposium on Object-oriented Programming, Systems, Languages, and Applications OOPSLA'16* (acceptance rate: 25%, 52/203), **Outstanding Paper Award (4/52)**
- SC16** T. Gysi, J. Baer, and T. Hoefler: dCUDA: Hardware Supported Overlap of Computation and Communication *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC16)* (acceptance rate: 18%, 82/446)
- SC16** J. Domke and T. Hoefler: Scheduling-Aware Routing for Supercomputers *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC16)* (acceptance rate: 18%, 82/446)
- SC16** W. Tang, B. Wang, S. Ethier, G. Kwasniewski, T. Hoefler, K. Ibrahim, K. Madduri, S. Williams, L. Oliker, C. Rosales-Fernandez, and T. Williams: Extreme Scale Plasma Turbulence Simulations on Top Supercomputers Worldwide *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC16)* (acceptance rate: 18%, 82/446)
- SC16** M. Martinasso, G. Kwasniewski, S. Alam, T. Schulthess, and T. Hoefler: A PCIe Congestion-Aware Performance Model for Densely Populated Accelerator Servers *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC16)* (acceptance rate: 18%, 82/446)
- HOTI'16** T. Schneider, O. Bibartiu and T. Hoefler: Ensuring Deadlock-Freedom in Low-Diameter InfiniBand Networks *In Proceedings of the 24th IEEE Symposium on High-Performance Interconnects, HOTI'16*, **Best Student Paper**
- HPDC'16** 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 Symposium on High-Performance Parallel and Distributed Computing (HPDC'16)* (acceptance rate: 16%, 20/129)
- HPDC'16** P. Schmid, M. Besta, and T. Hoefler: High-Performance Distributed RMA Locks *In Proceedings of the 25th Symposium on High-Performance Parallel and Distributed Computing (HPDC'16)* (acceptance rate: 16%, 20/129), **received Karsten Schwan Best Paper Award (1/20)**
- ICS'16** T. Grosser and T. Hoefler: Polly-ACC: Transparent compilation to heterogeneous hardware *In Proceedings of the the 30th International Conference on Supercomputing (ICS'16)* (acceptance rate: 24%, 43/178)
- PACT'15** H. Schweizer, M. Besta, and T. Hoefler: Evaluating the Cost of Atomic Operations on Modern Architectures *In Proceedings of the 24th International Conference on Parallel Architectures and Compilation (PACT'15)* (acceptance rate: 21%, 38/179)
- PACT'15** A. Bhattacharyya and T. Hoefler: Using Compiler Techniques to Improve Automatic Performance Modeling *In Proceedings of the 24th International Conference on Parallel Architectures and Compilation (PACT'15)* (acceptance rate: 21%, 38/179)
- SC15** T. Hoefler and R. Belli: Scientific Benchmarking of Parallel Computing Systems *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC15)* (acceptance rate: 22%, 79/358)
- SC15** G. Kathareios, C. Minkenber, B. Priscari, G. Rodriguez, and T. Hoefler: Cost-Effective Diameter-Two Topologies: Analysis and Evaluation *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC15)* (acceptance rate: 22%, 79/358)
- HOTI'15** S. Di Girolamo, P. Jolivet, K. D. Underwood and T. Hoefler: Exploiting Offload Enabled Network Interfaces *In Proceedings of the 23rd IEEE Symposium on High-Performance Interconnects, HOTI'15*, **Best Student Paper**
- HPDC'15** M. Besta and T. Hoefler: Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages *In Proceedings of ACM Symposium on High-Performance Parallel and Distributed Computing, HPDC'15* (acceptance rate: 16% (19/116)), **HPDC'15 Best Paper (1/19)**
- ICS'15** S. Shudler, A. Calotoiu, T. Hoefler, and F. Wolf: Exascaling Your Library: Will Your Implementation Meet Your Expectations? *In Proceedings of the ACM Conference on Supercomputing, ICS'15* (acceptance rate: 25% (40/160))
- HPDC'15** M. Poke and T. Hoefler: DARE: High-Performance State Machine Replication on RDMA Networks *Accepted at ACM HPDC'15* (acceptance rate: 16% (19/116))
- ICS'15** M. Besta and T. Hoefler: Active Access: A Mechanism for High-Performance Distributed Data-Centric Computations *In Proceedings of the ACM Conference on Supercomputing, ICS'15* (acceptance rate: 25% (40/160))

- ICS'15** T. Gysi, T. Grosser, and T. Hoefler: MODESTO: Data-centric Analytic Optimization of Complex Stencil Programs on Heterogeneous Architectures *In Proceedings of the ACM Conference on Supercomputing, ICS'15* (acceptance rate: 25% (40/160))
- HPDC'15** S. Ramos and T. Hoefler: Cache Line Aware Optimizations for ccNUMA Systems *In Proceedings of ACM Symposium on High-Performance Parallel and Distributed Computing, HPDC'15* (short paper)
- IPDPS'15** R. Belli and T. Hoefler: Notified Access: Extending Remote Memory Access Programming Models for Producer-Consumer Synchronization *In Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, (acceptance rate: 21.8% (108/496)), **IPDPS'15 Best Paper (4/108)**
- SC14** M. Besta and T. Hoefler: Slim Fly: A Cost Effective Low-Diameter Network Topology *In Proceedings of IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC14)*, (acceptance rate: 21%, 82/394), **SC14 Best Student Paper (1/82)**
- SC14** J. Domke, T. Hoefler, and S. Matsuoka: Fail-in-Place Network Design: Interaction between Topology, Routing Algorithm and Failures *In Proceedings of IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC14)*, (acceptance rate: 21%, 82/394)
- SC14** K. B. Ferreira, P. Widener, S. Levy, D. Arnold, and T. Hoefler: Understanding the Effects of Communication and Coordination on Checkpointing at Scale *In Proceedings of IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC14)*, (acceptance rate: 21%, 82/394)
- PACT'14** A. Bhattacharyya and T. Hoefler: PEMOGEN: Automatic Adaptive Performance Modeling during Program Runtime *In Proceedings of 23rd Intl. Conference on Parallel Architectures and Compilation Techniques (PACT'14)*
- HPDC'14** B. Prisacari, G. Rodriguez, P. Heidelberger, D. Chen, C. Minkenberg and T. Hoefler: Efficient Task Placement and Routing in Dragonfly Networks *In Proceedings of the 23rd ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC'14)*, (acceptance rate: 16%, 21/130)
- HPDC'14** M. Besta and T. Hoefler: Fault Tolerance for Remote Memory Access Programming Models *In Proceedings of the 23rd ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC'14)*, (acceptance rate: 16%, 21/130), **Best Paper Nominee (3/21)**
- SPAA'14** T. Hoefler and G. Kwasniewski: Automatic Complexity Analysis of Explicitly Parallel Programs *In Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'14)*, (acceptance rate: 25%, 30/122)
- IPDPS'14** A. Arteaga, T. Hoefler and O. Fuhrer: Designing Bit-Reproducible Portable High-Performance Applications *In Proceedings of IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, (acceptance rate: 21.1%, 114/541)
- ACM TACO (HIPEAC)** B. Prisacari, G. Rodriguez, C. Minkenberg, and T. Hoefler: Fast Pattern-Specific Routing for Fat Tree Networks *In ACM Transactions on Architecture and Code Optimization*, and presented at the HIPEAC 2014 conference, (acceptance rate: 24%, 2011)
- SC13** A. Calotiu, T. Hoefler, M. Poke, and F. Wolf: Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, (acceptance rate: 20%, 92/457)
- SC13** R. Gerstenberger, M. Besta, and T. Hoefler: Enabling Highly-Scalable Remote Memory Access Programming with MPI-3 One Sided *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, (acceptance rate: 20%, 92/457), **SC13 Best Paper Award (1/92) and Best Student Paper Finalist (8/92)**
- SC13** A. Friedley, G. Bronevetsky, A. Lumsdaine, and T. Hoefler: Hybrid MPI: Efficient Message Passing for Multi-core Systems *In Proceedings of the IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, (acceptance rate: 20%, 92/457)
- ICPP'13** T. Schneider, R. Grant, B. Barrett, R. Brightwell, and T. Hoefler: Protocols for Fully Offloaded Collective Operations on Accelerated Network Adapters *In Proceedings of the Intl. Conference on Parallel Processing, ICPP'13*
- EuroMPI'13** T. Schneider, F. Kjolstad, and T. Hoefler: MPI Datatype Processing using Runtime Compilation *In Proceedings of ACM/SIGHPC Recent Advances in Message Passing Interface, EuroMPI'13*, **Best Paper Award (1/25)**
- ICS'13** B. Prisacari, G. Rodriguez, C. Minkenberg, and T. Hoefler: Bandwidth-optimal Alltoall Exchanges in Fat Tree Networks *In Proceedings of the 27th ACM International Conference on Supercomputing, ICS'13* (acceptance rate: 21%, 41/198)

- HPDC'13** S. Ramos Garea and T. Hoefler: Modeling Communication in Cache-Coherent SMP Systems - A Case-Study with Xeon Phi *In Proceedings of the 22nd ACM Symposium on High-Performance Parallel and Distributed Computing, HPDC'13 (acceptance rate: 15%, 20/131)*
- HPDC'13** S. Li, T. Hoefler, and M. Snir: NUMA-Aware Shared Memory Collective Communication for MPI *In Proceedings of the 22nd ACM Symposium on High-Performance Parallel and Distributed Computing, HPDC'13 (acceptance rate: 15%, 20/131), **Best Paper Nominee (3/20)***
- PPoPP'13** A. Friedley, T. Hoefler, G. Bronevetsky, and A. Lumsdaine: Ownership Passing: Efficient Distributed Memory Programming on Multi-core Systems *In Proceedings of the 18th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, PPoPP'13, pages 177–186. ACM, Feb. 2013 (acceptance rate: 21%, 100/472)*
- SC12** T. Hoefler and T. Schneider: Optimization Principles for Collective Neighborhood Communications *In Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC'12, pages 98:1–98:10. IEEE Computer Society Press, Nov. 2012, (acceptance rate: 21%, 100/472)*
- EuroMPI'12** T. Schneider, R. Gerstenberger, and T. Hoefler: Micro-Applications for Communication Data Access Patterns and MPI Datatypes *In Recent Advances in the Message Passing Interface - 19th European MPI Users' Group Meeting, EuroMPI 2012, volume 7490, pages 121-131. Springer, Sept. 2012*
- EuroMPI'12** S. Pellegrini, T. Hoefler, and T. Fahringer: Exact Dependence Analysis for Increased Communication Overlap *In Recent Advances in the Message Passing Interface - 19th European MPI Users' Group Meeting, EuroMPI 2012, volume 7490, pages 89–99. Springer, Sept. 2012*
- EuroMPI'12** T. Hoefler J. Dinan, D. Buntinas, P. Balaji, B. Barrett, R. Brightwell, W. Gropp, V. Kale, and R. Thakur: Leveraging MPI's One-Sided Communication Interface for Shared-Memory Programming *In Recent Advances in the Message Passing Interface - 19th European MPI Users' Group Meeting, EuroMPI 2012, volume 7490, pages 132–141. Springer, Sept. 2012*
- PACT'12** T. Hoefler and T. Schneider: Runtime Detection and Optimization of Collective Communication Patterns *In Proceedings of the 21st international conference on Parallel Architectures and Compilation Techniques, PACT'12, pages 263–272. ACM, Sept. 2012, (acceptance rate: 19%, 39/207)*
- SC11** T. Hoefler, W. Gropp, M. Snir, and W. Kramer: Performance Modeling for Systematic Performance Tuning *In State of the Practice Reports, SC'11, pages 6:1–6:12. ACM, Nov. 2011*
- EuroMPI'11** W. Gropp, T. Hoefler, R. Thakur, and J. L. Traeff: Performance Expectations and Guidelines for MPI Derived Datatypes *In Recent Advances in the Message Passing Interface, EuroMPI'11, volume 6960, pages 150–159. Springer, Sept. 2011*
- EuroMPI'11** V. Venkatesan, M. Charawi, E. Gabriel, and T. Hoefler.: Design and Evaluation of Nonblocking Collective I/O Operations *In Recent Advances in the Message Passing Interface, EuroMPI'11, volume 6960, pages 90–98. Springer, Sept. 2011*
- EuroMPI'11** T. Hoefler, and M. Snir.: Writing Parallel Libraries with MPI - Common Practice, Issues, and Extensions *In Recent Advances in the Message Passing Interface, EuroMPI'11, volume 6960, pages 345–355. Springer, Sept. 2011, Keynote Paper at the IMUDI session at EuroMPI 2011 Conference*
- ICS'11** T. Hoefler and M. Snir.: Generic Topology Mapping Strategies for Large-scale Parallel Architectures. *In Proceedings of the 2011 ACM International Conference on Supercomputing, ICS'11, pages 75–85. ACM, Jun. 2011 (acceptance rate 21.7%, 35/161)*
- ICS'11** J. Willcock, T. Hoefler, N. Edmonds, and A. Lumsdaine.: Active Pebbles: Parallel Programming for Data-Driven Applications. *In Proceedings of the 2011 ACM International Conference on Supercomputing, ICS'11, pages 235–245. ACM, Jun. 2011 (acceptance rate 21.7%, 35/161)*
- IPDPS'11** J. Domke, T. Hoefler, and W. Nagel.: Deadlock-Free Oblivious Routing for Arbitrary Topologies. *In Proceedings of the 25th IEEE International Parallel & Distributed Processing Symposium, IPDPS'11, pages 613–624. IEEE Computer Society, May 2011 (acceptance rate: 19.6%, 112/571)*
- PPoPP'11** J. Willcock, T. Hoefler, N. Edmonds, and A. Lumsdaine.: Active Pebbles: A Programming Model For Highly Parallel Fine-Grained Data-Driven Computations. *In Proceedings of the 16th ACM symposium on Principles and Practice of Parallel Programming, PPoPP'11, pages 305–306. ACM, Feb. 2011, **Best Poster at PPoPP'11** (acceptance rate: 25%, 26/165 papers + 16/165 poster).*
- SC'10** T. Hoefler, T. Schneider, and A. Lumsdaine.: Characterizing the Influence of System Noise on Large-Scale Applications by Simulation. *In Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC'10, pages 1–11. IEEE Computer Society, Nov. 2010, **Best Paper at SC10**, (acceptance rate: 19.8%, 50/253)*

- PACT'10** J. Willcock, T. Hoefler, N. Edmonds, and A. Lumsdaine.: AM++: A Generalized Active Message Framework. *In Proceedings of the 19th international conference on Parallel Architectures and Compilation Techniques, PACT'10, pages 401-410. ACM, Sept. 2010 (acceptance rate: 17%, 46/266)*
- EuroMPI'10** T. Hoefler, G. Bronevetsky, B. Barrett, B. R. de Supinski, and A. Lumsdaine.: Efficient MPI Support for Advanced Hybrid Programming Models. *In Recent Advances in the Message Passing Interface, EuroMPI'10, pages 50-61, volume LNCS 6305. Springer, Sept. 2010*
- EuroMPI'10** T. Hoefler, W. Gropp, R. Thakur, and J. L. Traeff.: Toward Performance Models of MPI Implementations for Understanding Application Scaling Issues. *In Recent Advances in the Message Passing Interface, EuroMPI'10, pages 21-30, volume LNCS 6305. Springer, Sept. 2010*
- EuroMPI'10** T. Hoefler and S. Gottlieb.: Parallel Zero-Copy Algorithms for Fast Fourier Transform and Conjugate Gradient using MPI Datatypes. *In Recent Advances in the Message Passing Interface, EuroMPI'10, pages 132-141, volume LNCS 6305. Springer, Sept. 2010*
- HotI'10** B. Arimilli, R. Arimilli, V. Chung, S. Clark, W. Denzel, B. Drerup, T. Hoefler, J. Joyner, J. Lewis, J. Li, N. Ni, and R. Rajamony.: The PERCS High-Performance Interconnect. *Proceedings of 18th Symposium on High-Performance Interconnects (Hot Interconnects 2010). IEEE, Aug. 2010. (invited paper)*
- PPoPP'10** T. Hoefler, C. Siebert, and A. Lumsdaine.: Scalable Communication Protocols for Dynamic Sparse Data Exchange. *Proceedings of the 2010 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pages 159-168. ACM, Jan. 2010. (acceptance rate: 16.8%, 29/173)*
- HotI'09** T. Hoefler, T. Schneider, and A. Lumsdaine.: Optimized Routing for Large-Scale InfiniBand Networks. *17th Annual IEEE Symposium on High Performance Interconnects, HOTI'09, IEEE Computer Society, Aug. 2009. (acceptance rate: 30%, 14/47)*
- ICPP'09** T. Hoefler, C. Siebert, and A. Lumsdaine.: Group Operation Assembly Language - A Flexible Way to Express Collective Communication ICPP-2009 - *The 38th International Conference on Parallel Processing. IEEE, Sep. 2009. (acceptance rate: 32%, 71/220)*
- EuroMPI'09** T. Hoefler, A. Lumsdaine, and J. Dongarra.: Towards Efficient MapReduce Using MPI. *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 16th European PVM/MPI Users' Group Meeting, EuroPVM/MPI'09. Springer, Sep. 2009.*
- LCI'09** 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. *Proceedings of the 10th LCI International Conference on High-Performance Clustered Computing, LCI'09, Mar. 2009, **Best Student Paper at LCI'09***
- Cluster'08** T. Hoefler, T. Schneider, and A. Lumsdaine.: Multistage Switches are not Crossbars: Effects of Static Routing in High-Performance Networks. *Proceedings of the 2008 IEEE International Conference on Cluster Computing, CLUSTER'08. IEEE Computer Society, Oct. 2008. (acceptance rate: 30%, 28/92)*
- Cluster'08** T. Hoefler and A. Lumsdaine.: Message Progression in Parallel Computing - To Thread or not to Thread? *Proceedings of the 2008 IEEE International Conference on Cluster Computing, CLUSTER'08. IEEE Computer Society, Oct. 2008. (acceptance rate: 30%, 28/92)*
- HotI'08** P. Geoffray and T. Hoefler.: Adaptive Routing Strategies for Modern High Performance Networks. *16th Annual IEEE Symposium on High Performance Interconnects, HOTI'08, pages 165-172. IEEE Computer Society, Aug. 2008. (acceptance rate: 30%, 14/47)*
- SPAA'08** T. Hoefler, P. Gottschling, and A. Lumsdaine.: Leveraging Non-blocking Collective Communication in High-performance Applications. *Proceedings of the Twentieth Annual Symposium on Parallelism in Algorithms and Architectures, SPAA'08, pages 113-115. Association for Computing Machinery (ACM), Jun. 2008. (acceptance rate: 28%, 36/128)*
- SC07** T. Hoefler, A. Lumsdaine, and W. Rehm.: Implementation and Performance Analysis of Non-Blocking Collective Operations for MPI. *In proceedings of the 2007 International Conference on High Performance Computing, Networking, Storage and Analysis, SC07. IEEE Computer Society/ACM, Nov. 2007. (acceptance rate: 20%, 54/268)*
- EuroMPI'06** T. Hoefler, P. Gottschling, W. Rehm, and A. Lumsdaine.: Optimizing a Conjugate Gradient Solver with Non-Blocking Collective Operations. *Proceedings of Recent Advantages in Parallel Virtual Machine and Message Passing Interface, EuroPVM/MPI'06, pages 374-382. Springer, Sep. 2006.*

Edited Journals

- IJHPCA'13** T. Hoefler and Kamil Iskra (Editors).: Operating systems and runtime environments on supercomputers *IJHPCA*, May 2013 (vol 27 no. 2).
- IJHPCA'12** T. Hoefler and Kamil Iskra (Editors).: Issues in Large Scale Computing Environments: Heterogeneous Computing and Operating Systems *IJHPCA*, May 2012 (vol 26 no. 2).
- IEEE Micro'12** T. Hoefler, P. Geoffray, F. Petrini, J. L. Traeff (Editors).: Top Picks from Hot Interconnects 2011: Petascale Network Architectures *IEEE Micro*, Jan/Feb. 2012 (vol 32 no. 1).
- PARCO'12** T. Hoefler (Editor).: Extensions for Next-Generation Parallel Programming Models. *Elsevier Parallel Computing*, Jan/Feb. 2012.

Selected Journal Publications and Book Chapters

- IEEE TOC** Fabian Schuiki, Florian Zaruba, Torsten Hoefler, Luca Benini: Stream Semantic Registers: A Lightweight RISC-V ISA Extension Achieving Full Compute Utilization in Single-Issue Cores *IEEE Transactions on Computers (TOC)*. *IEEE*, Apr. 2020
- TRETS** Maciej Besta, Marc Fischer, Tal Ben-Nun, Dimitri Stanojevic, Johannes de Fine Licht, Torsten Hoefler: Substream-Centric Maximum Matchings on FPGA *In Proceedings of the ACM Trans. Reconfig. Technol. Systems*
- TPDS** Sergei Shudler, Yannick Berens, Alexandru Calotou, Torsten Hoefler, Alexandre Strube, Felix Wolf: Engineering Algorithms for Scalability through Continuous Validation of Performance Expectations *IEEE Transactions on Parallel and Distributed Systems (TPDS)*. Vol 30, Nr. 8, *IEEE*, Jul. 2019
- CiSE** T. Schulthess, P. Bauer, O. Fuhrer, T. Hoefler, C. Schaer, N. Wedi: Reflecting on the goal and baseline for exascale computing: a roadmap based on weather and climate simulations *Computing in Science and Engineering (CiSE)*. Vol 21, Nr. 1, *IEEE Computer Society*, ISSN: 1521-9615, Jan. 2019
- CACM** R. Gerstenberger, M. Besta, T. Hoefler: Enabling Highly-Scalable Remote Memory Access Programming with MPI-3 One Sided *In Communications of the ACM*, *ACM*, Oct. 2018, **Invited Article to Research Highlights**
- GMD'18** O. Fuhrer, T. Chadha, T. Hoefler, G. Kwasniewski, X. Lapillonne, D. Leutwyler, D. Luethi, C. Osuna, C. Schaer, T. C. Schulthess, H. Vogt: Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0 *Geoscientific Model Development*. Vol 11, Nr. 4, *Copernicus Publications*, May 2018
- IEEE TPDS'18** Shigang Li, Yunquan Zhang, Torsten Hoefler: Cache-Oblivious MPI All-to-All Communications Based on Morton Order *IEEE Transactions on Parallel and Distributed Systems (TPDS)*. Vol 29, Nr. 3, *IEEE*, Mar. 2018
- IEEE TPDS'17** Didem Unat, Anshu Dubey, Torsten Hoefler, John Shalf, Mark Abraham, Mauro Bianco, Bradford L. Chamberlain, Romain Cledat, H. Carter Edwards, Hal Finkel, Karl Fuerlinger, Frank Hannig, Emmanuel Jeannot, Amir Kamil, Jeff Keasler, Paul H J Kelly, Vitus Leung, Hatem Ltaief, Naoya Maruyama, Chris J. Newburn, and Miquel Pericas: Trends in Data Locality Abstractions for HPC Systems *IEEE Transactions on Parallel and Distributed Systems (TPDS)*. Vol 28, Nr. 10, *IEEE*, Oct. 2017
- IEEE TPDS'16** S. Ramos and T. Hoefler.: Cache Line Aware Algorithm Design for Cache-Coherent Architectures *IEEE Transactions on Parallel and Distributed Systems (TPDS)*. Vol PP, Nr. 99, *IEEE*, Jan. 2016
- IEEE MICRO'16** S. Di Girolamo, P. Jolivet, K. D. Underwood, and T. Hoefler.: Exploiting Offload Enabled Network Interfaces *IEEE MICRO*. Vol 36, Nr. 4, *IEEE*, Jul. 2016
- ACM TOPC'15** T. Hoefler, J. Dinan, R. Thakur, B. Barrett, P. Balaji, W. Gropp, K. Underwood.: Remote Memory Access Programming in MPI-3 *ACM Transactions on Parallel Computing (TOPC)*. *ACM*, Jan. 2015
- JSFI'14** T. Hoefler and D. Moor.: Energy, Memory, and Runtime Tradeoffs for Implementing Collective Communication Operations *Journal of Supercomputing Frontiers and Innovations*. Vol 1, Nr. 2, pages 58–75, Oct. 2014
- Computing'13** T. Schneider, R. Gerstenberger, T. Hoefler.: Application-oriented ping-pong benchmarking: how to assess the real communication overheads. *Journal of Computing*. *Springer*, May. 2013.
- Computing'13** T. Hoefler, J. Dinan, D. Buntinas, P. Balaji, B. Barrett, R. Brightwell, W. Gropp, V. Kale and R. Thakur.: MPI + MPI: a new hybrid approach to parallel programming with MPI plus shared memory. *Journal of Computing*. *Springer*, May. 2013.
- MPI-3.0 Standard** The MPI Forum.: MPI: A Message-Passing Interface Standard, Version 2.2. *Technical report*, *MPI Forum*, 2012. (Chapters 5 (Collective Operations), 7 (Process Topologies), and 11 (One Sided)).

- PPL'11** P. Balaji, D. Buntinas, D. Goodell, W. Gropp, T. Hoefler, S. Kumar, E. Lusk, R. Thakur, and J. L. Traeff.: MPI on Millions of Cores. *Parallel Processing Letters (PPL)*, Mar. 2011.
- CiSE'10** T. Hoefler.: Software and Hardware Techniques for Power-Efficient HPC Networking. *Computing in Science and Engineering (CiSE)*, Dec. 2010.
- CCPE'10** T. Hoefler, R. Rabenseifner, H. Ritzdorf, B. R. de Supinski, R. Thakur, , and J. L. Traeff.: The Scalable Process Topology Interface of MPI 2.2. *Concurrency and Computation: Practice and Experience*, Dec. 2010.
- MPI-2.2 Standard** The MPI Forum.: MPI: A Message-Passing Interface Standard, Version 2.2. *Technical report, MPI Forum, 2009. (Chapters 5 (Collective Operations) and 7 (Process Topologies))*.
- PPL'09** T. Hoefler, T. Schneider, and A. Lumsdaine.: The Effect of Network Noise on Large-Scale Collective Communications. *Parallel Processing Letters (PPL)*, 19(4):573–593, Aug. 2009.
- PARCO'07** T. Hoefler, P. Gottschling, A. Lumsdaine, and W. Rehm.: Optimizing a Conjugate Gradient Solver with Non-Blocking Collective Operations. *Elsevier Journal of Parallel Computing (PARCO)*, 33(9):624–633, Sep. 2007.

Patents

- US11076210B1** T. Hoefler, M. Heddes, D. Goel, J. Belk: Distributed processing architecture *US CN WO*
- US11720252B1** M. Heddes, A. More, N. Shah, T. Hoefler: Method and apparatus for compressing and decompressing sparse data sets *US*
- US20220291976A1** D. Goel, M. Heddes, T. Hoefler, X. Xu: Message communication between integrated computing devices *WO US TW*
- US20220244911A1** M. Heddes, T. Hoefler: Digital circuitry for normalization functions *WO US TW*
- US20220138524A1** M. Heddes, T. Hoefler, K. Colwell, A. Phanishayee: Training neural networks based on dual pipeline architectures *US WO EP*

Selected Invited Talks

Summary: 45 invited keynotes, more than 80 invited talks, not counting regular conference or workshop paper presentations

- Keynote HPCAC'24 **Scalable and Efficient AI: Federated Supercomputers and Smartphones** **Locarno, Switzerland**
Keynote at the HPC Advisory Council Conference 2024
- Keynote SCAsia'24 **High-performance Climate Simulations and AI - from Earth Virtualization to Data Compression** **Sydney, Australia**
Keynote at Supercomputing Asia Conference 2024
- Keynote ADAC'24 **Scalable and Efficient AI: Federated Supercomputers and Smartphones** **Sydney, Australia**
Keynote at 14th Accelerated Data Analytics and Computing Institute Symposium 2024
- Plenary Talk FCRC'23 **AI and HPC at a Crossroads** **Orlando, FL, USA**
Keynote at ACM Federated Computing Research Conference (FCRC), 2023
- Keynote REP'23 **Reproducing Performance - The Good, the Bad, and the Ugly** **Santa Cruz, CA, USA**
Opening keynote at 1st ACM Conference on Reproducibility and Replicability
- Keynote ICPP'23 **Scalable and Efficient AI: From Supercomputers to Smartphones** **Salt Lake City, UT, USA**
Keynote talk at the 52nd International Conference on Parallel Processing
- Keynote HPCAC'23 **Network support for High-Performance Deep Learning Systems** **Lugano, Switzerland**
Keynote at the HPC Advisory Council Conference 2023
- Keynote CCWC'23 **Distributed AI Supercomputing: Large-Scale Transformers to Graph Neural Networks** **virtual**
Keynote at the IEEE 13th Annual Computing and Communication Workshop and Conference
- Keynote HPC Asia'23 **Principles and Systems for AI Supercomputing** **Singapore**
Keynote at the HPC Asia 2023 conference, Singapore

Keynote Intel Fellows Forum	Efficient AI: From Supercomputers to Smartphones Keynote at the Intel Fellows Forum - Distributed Systems Summit, virtual	virtual
Korea AI Summit	Efficient AI: From Supercomputers to Smartphones Invited talk at the 2022 Korean AI summit in Seoul, Korea	Seoul, Korea
Keynote NTCI'22	Efficient AI: From Supercomputers to Smartphones Keynote at the 2022 International Symposium on New Trends in Computational Intelligence	virtual
Keynote IEMTRONICS'22	Sparsity in Deep Learning Keynote at the IOT, Electronics and Mechatronics Conference 2022, virtual	virtual
Keynote HPCAC'22	New trends for sPIN-based in-network computing - from sparse reductions to RISC-V Keynote as HPC Advisory Council Conference 2022, Lugano, Switzerland	virtual
Plenary HPCW'22	The Three Pillars of Large-scale Deep Learning Plenary talk at the Japanese NVIDIA HPC Weeks 2021	virtual
Keynote RSD'22	Data-Centric Python - Productivity, portability and all with high performance! Keynote at the Russian Supercomputing Days 2021	virtual
Keynote COLOC'21	Data-Centric Python - Productivity, portability and all with high performance! Keynote at the 5th workshop on data locality COLOC'21	virtual
Keynote SNN'21	Sparsity in Deep Learning Opening keynote and tutorial at the Sparsity in Neural Networks workshop (SNN'21)	virtual
Keynote DLonSC'21	High-Performance Scalable Deep Learning Keynote talk at the Deep Learning on Supercomputers Workshop in conjunction with ISC'21	virtual
Keynote HEART'21	Portable high-performance Python on CPUs, GPUs, and FPGAs Keynote talk at the International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies (HEART 2021)	virtual
Keynote CIUK'20	A data-centric approach to performance portability Keynote at the 2020 HPC conference Computing Insight UK (CIUK'20)	virtual
Keynote Bench'20	Keynote & Award Lecture (BenchCouncil Rising Star Award): Scientific Benchmarking of Parallel Computing Systems 2020 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench'20)	virtual
Keynote DISC'20	High-performance distributed memory systems from supercomputers to data centers 34th International Symposium on Distributed Computing (DISC)	virtual
Keynote HPBD&IS'20	High-Performance Communication in Machine Learning Keynote talk at the (Virtual) International Conference on High-Performance Big Data and Intelligent Systems (HPBD&IS'20)	virtual
Keynote MLHPC'19	HPC for ML and ML for HPC - Scalability, Communication, and Programming Denver, CO, USA Keynote talk at the International Machine Learning in High-Performance Computing (MLHPC'19 in conjunction with ACM/IEEE Supercomputing, SC19)	
Keynote PPAM'19	High-Performance Communication in Machine Learning Keynote talk at the 13th International Conference on Parallel Processing and Applied Mathematics (PPAM'19)	Bialystok, Poland
Keynote ParCo'19	Data-Centric Parallel Programming Keynote talk at the 18th International Parallel Computing conference (ParCo'19)	Prague, Czech Republic
Keynote EMiT'19	High-Performance Communication for Machine Learning Keynote talk at the 5th Conference on Emerging Technologies – EMiT2019	Huddersfield, UK
Keynote AsHES	Performance Portability with Data-Centric Parallel Programming Keynote talk at the 9th International Workshop on Accelerators and Hybrid Exascale Systems (AsHES)	virtual
Keynote AHPC'19	High-Performance Communication in Machine Learning Keynote at the Austrian HPC meeting 2019	Grundlsee, Austria

Keynote ExaMPI'18	RDMA, Scalable MPI-3 RMA, and Next-Generation Post-RDMA Interconnects	Dallas, TX, USA
	Keynote at ExaMPI 2018 Workshop (in conjunction with SC18)	
Tsinghua University	Performance Modeling for Future Computing Technologies	Beijing, China
	Invited talk at 60 years of CS @ Tsinghua celebration	
Keynote HPCAC	Demystifying Parallel and Distributed Deep Learning on Supercomputers: An In-Depth Concurrency Analysis	Lugano, Switzerland
	Keynote at Swiss HPC Advisory Council Conference 2018	
Keynote Multicore @ Siemens	Developing high-performance software, from modeling to programming	Nuremberg, Germany
	Invited opening presentation at the Multicore@Siemens conference	
Keynote HiPINEB'18	The three L's in modern high-performance networking: low latency, low cost, low processing load	in conjunction with HPCA'17
	Keynote at the HiPINEB workshop at HPCA'18	
Keynote EMBRACE'17	Scientific Benchmarking of Parallel Computing Systems	in conjunction with IPDPS'17
	Keynote talk at EMBRACE Workshop (IPDPS'17)	
Keynote HPC China'16	Theory and Practice in HPC: Modeling, Programming, and Networking	Xi'an, China
	Keynote at the HPC China 2016 conference	
Keynote Cluster'16	Theory and Practice in HPC: Modeling, Programming, and Networking	Taipei, Taiwan
	Opening Keynote at the IEEE Cluster 2016 conference	
Keynote HIPS'15	How fast will your application go? Static and dynamic techniques for application performance modeling.	Hyderabad, India
	Keynote at the HIPS'15/LSPP'15 combined workshop in conjunction with IPDPS'15	
Keynote LLVMHPC'14	A case for runtime recompilation in HPC	New Orleans, LA, USA
	Keynote at the LLVM Compiler Infrastructure in HPC workshop at SC14, Nov. 2014	
Keynote ExaMPI'13	MPI Beyond 3.0 and Towards Larger-Scale Computing	Denver, CO, USA
	Keynote at the Workshop on Exascale MPI at SC13, Nov. 2013, ≈120 attendees	
SC13	The Second Green Graph500 List	Denver, CO, USA
	Birds of a Feather, Nov. 2013	
Dagstuhl	Fault Tolerance for Remote Memory Access Programming Models	Dagstuhl, Germany
	Invited to seminar "Resilience in Exascale Computing"	
ISC'13	The First Green Graph500 List	Leipzig, Germany
	Birds of a Feather, Jun. 2013	
EASC'13	Application-Centric Benchmarking and Modeling for Co-Design	Edinburgh, UK
	Exascale Applications and Software Conference	
Keynote MCC'12	MPI-3.0: A Response to New Challenges in Hardware and Software	Stuttgart, Germany
	Keynote talk at Multicore Challenge Conference 2012	
TiTech'12	Optimized routing and process mapping for arbitrary network topologies	Tokio, Japan
	Invited talk at Tokio Institute of Technology	
Keynote EuroMPI'11	Writing Parallel Libraries with MPI - The Good, the Bad, and the Ugly	Santorini, Greece
	Keynote talk at 18th European PVM/MPI User's Group Meeting	
Keynote EnA-HPC'11	Energy-aware Software Development for Massive-Scale Systems	Hamburg, Germany
	Keynote at the International Conference on Energy-Aware High Performance Computing	
Jülich 2011	Model-Driven HPC Software and System Design and Optimization	Jülich, Germany
	Jülich Supercomputing Center, Apr. 2011	
Keynote PROPER'10	Analytical Performance Modeling and Simulation for Blue Waters	Ischia, Italy
	Keynote at the Workshop on Productivity and Performance in conjunction with EuroPar, Aug. 2010	
Argonne Natl. Laboratory	Nonblocking and Sparse Collective Operations on Petascale Computers	Chicago, IL, USA
	Argonne National Laboratory, Jun. 2010	

SC'09 BoF	Selected MPI-2.2 and MPI-3 Features MPICH Birds of a Feather, Nov. 2009	Portland, OR, USA
Cisco Systems	The Effects of Common Communication Patterns in Large-Scale Networks with Switch-Based Static Routing Nerd Lunch at Cisco Systems, Aug. 2008	San Jose, CA, USA
Berkeley Natl. Laboratory	Multistage Interconnection Networks are not Crossbars Lawrence Berkeley National Laboratory, Aug. 2008	Berkeley, CA, USA
Livermore Natl. Laboratory	Non-blocking Collective Operations for MPI Lawrence Livermore National Laboratory, Aug. 2008	Livermore, CA, USA
HLRS	Non-blocking Collectives for MPI-2 High Performance Computing Center Stuttgart (HLRS), Dec. 2007	Stuttgart, Germany
ABINIT Workshop	Optimization of a parallel 3d-FFT with non-blocking Collective Operations Invited to the 3rd International ABINIT Developer Workshop, Jan. 2007	Liege, Belgium
TU Munich	Fast Barrier Synchronization for InfiniBand Technical University of Munich, Sep. 2005	Munich, Germany

Impact

Microsoft Maia	Network and Programming Model Architecture Network and Programming model co-architect of a production-level AI Supercomputing System (Chip, Rack-Level, System-level co-design from ground up for AI workloads)
DFSSSP Routing	Deadlock-free Single Source Shortest Path routing The fastest routing algorithm for arbitrary topologies. Available in OpenSM (the InfiniBand subnet manager) and used at various sites. (with J. Domke)
Nonblocking Collectives	Nonblocking Collective Operations for MPI Proposed algorithms and reference implementation that are now used in virtually every MPI implementation. Drove the standardization in MPI-3.0.
Neighborhood Collectives	Neighborhood Collective Operations for MPI Proposed algorithms and reference implementation that are now used in virtually every MPI implementation. Drove the standardization in MPI-3.0.
RMA Programming	Remote Memory Access Programming Co-editor and driver of the MPI-3.0 One Sided chapter. This functionality is implemented in virtually all MPI libraries. (with W. Gropp and R. Thakur)

External Funding

2022–2026	PSAP: Productive Spatial Accelerator Programming EUR 2M; ERC Consolidator Grant	ETH Zürich
2022–2026	Green responsibLe privACy preservIng dAta operaTIONs CHF 404'980: Horizon 2021 EU Project	ETH Zürich
2022–2023	Datacenter networking and secure communication CHF 150'000: Intel donation (2x)	ETH Zürich
2021–2025	The European PILOT EUR 538'884: Euro HPC Project	ETH Zürich
2021–2024	DaFIEx: Performance portability through dataflow extraction CHF 519'906; ETH Future Computing Lab	ETH Zürich
2021–2024	DaCeMI - Harnessing future hardware using Data-Centric ML Integration CHF 440'968, Platform for Advanced Scientific Computing	ETH Zürich
2021–2024	MAELSTROM: Empowering Weather and Climate Forecast EUR 606'000, Euro HPC Project, co-funded by SBF	ETH Zürich

2021–2024	DEEP SEA: Software for Exascale Architectures EUR 297'136, Euro HPC Project, co-funded by SBF1	ETH Zürich
2021–2024	RED SEA: Network Solution for Exascale Architectures EUR 234'700, Euro HPC Project, co-funded by SBF1	ETH Zürich
2020–2021	Automatic Dataflow Modeling for HPC Applications CHF 99'992, SNF Spark (with Alexandru Calotoiu)	ETH Zürich
2018–2024	EPI: The European Processor Initiative EUR 1.82M; EU Horizon 2020, FET, (lead: Luca Benini, ETH)	ETH Zürich
2018–2021	QIRO: An Intermediate Representation for Quantum Computing \$255'000; Gift by Microsoft in the context of the MSJRC joint lab	ETH Zürich
2018–2021	PASCHA - Portability And Scalability of COSMO on Heterogeneous Architectures CHF 494'300, Platform for Advanced Scientific Computing/HPCN	ETH Zürich
2018–2021	EPIGRAM-HS: Enabling Extrame-scale Applications on Heterogeneous Hardware EUR 371'000; EU Horizon 2020, FET, (lead: Stefano Markidis, KTH)	ETH Zürich
2018–2019	Intel Parallel Computing Center \$200'000; Unrestricted gift by Intel Corp., extended into second year after evaluation	ETH Zürich
2017–2020	Portability and Scalability of COSMO on Heterogeneous Architectures CHF 494'300; Platform for Advanced Scientific Computing, led by MeteoSwiss	ETH Zürich
2017–2019	Automatic Performance Modeling of HPC Applications with Multiple Model Parameters (Catwalk 2) CHF 187,787; DFG Normal Antrag with TU Darmstadt (funded through DE-CH lead agency)	ETH Zürich
2014–2017	Data Centric Mapping \$255'000; Gift by Microsoft in the context of the MSJRC joint lab	ETH Zürich
2015–2018	Mont-Blanc 3, European scalable and power efficient HPC platform based on low-power embedded technology EUR 396,350; EU Horizon 2020 - Excellent Science FET Proactive - with the MontBlanc 3 Consortium	ETH Zürich
2015–2018	Cloud-reserving climate modeling on future supercomputing platform (crClim) CHF 176,550; Swiss National Science Fund - Sinergia with D-USYS, MeteoSwiss, CSCS	ETH Zürich
2016–2021	Data-Centric Parallel Programming (DAPP) EUR 1.5M; ERC Starting Grant	ETH Zürich
2013–2017	A Heterogeneous Compiler Platform for Scientific Codes \$649,713; Platform for Advanced Scientific Computing	ETH Zürich
2013–2016	Data-Centric Compilation Techniques for Parallel Programs \$188,171; Swiss National Science Fund	ETH Zürich
2013–2016	Google Ph.D. Fellowship for Maciej Besta \$255,000 unrestricted gift; First European Fellowship for Parallel Computing	ETH Zürich
2013	Programming Hierarchical Memory Systems for Big Data Analytics \$30,000 unrestricted gift by IBM (faculty award)	ETH Zürich
2013–2016	A Quick Development Path for Performance Models ETH's share: \$177,338; DFG Special Priority Programme SPPEXA (funded by SNF)	ETH Zürich
2011–2012	Nonblocking Collective Operations for Portals IV \$50,000 subcontract of Sandia National Laboratories, NNSA, DOE, to UIUC	University of Illinois
2010–2013	Compiled MPI: Cost-Effective Exascale Application Development UI's share: \$165,000; funded under DOE X-Stack; in Collaboration with Daniel Quinlan, Greg Bronevetsky (LLNL) and Andrew Lumsdaine (IU)	University of Illinois
2005	Quantum Mechanical Computations € 55,000; individual funding for Ph.D. studies received from AMD Saxony	Chemnitz University of Technology

Teaching Experience

(Co)taught 10 undergraduate courses, 13 graduate courses, 6 seminars, 35 tutorials, 10 PhD committees

Fall 2023	Design of Parallel and High-Performance Computing 125 students (capped enrollment)	ETH Zürich
Nov. 2023	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 55 attendees</i>	SC23 Denver, CO
Spring 2023	Parallel Programming <i>co-taught with Barbara Solenthaler, undergrad, 570 students</i>	ETH Zürich
Nov. 2022	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 60 attendees</i>	SC22 Dallas, TX
Nov. 2022	Half-day Tutorial: Productive Parallel Programming for FPGA with HLS <i>co-presented with J. De Fine Licht, ≈ 50 attendees</i>	SC22 virtual
Jun. 2022	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 30 attendees</i>	ISC'22 Hamburg, Germany
Fall 2022	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel, 125 students (capped enrollment)</i>	ETH Zürich
Spring 2022	Parallel Programming <i>co-taught with Barbara Solenthaler, undergrad, 550 students</i>	ETH Zürich
Nov. 2021	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 60 attendees</i>	SC21 St. Louis, MO
Nov. 2021	Half-day Tutorial: Productive Parallel Programming for FPGA with HLS <i>co-presented with J. De Fine Licht, ≈ 40 attendees</i>	SC21 virtual
Jun. 2021	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 30 attendees</i>	ISC'21 Hamburg, Germany
Fall 2021	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel, 103 students</i>	ETH Zürich
Spring 2021	Parallel Programming <i>co-taught with Barbara Solenthaler, undergrad, 500 students</i>	ETH Zürich
Nov. 2020	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 50 attendees</i>	SC20 virtual
Nov. 2020	Half-day Tutorial: Productive Parallel Programming for FPGA with HLS <i>co-presented with J. De Fine Licht, ≈ 50 attendees</i>	SC20 virtual
Jun. 2020	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 30 attendees</i>	ISC'20 virtual
Fall 2020	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel, 94 students</i>	ETH Zürich
Spring 2020	Parallel Programming <i>co-taught with Hermann Lederer, undergrad, 485 students</i>	ETH Zürich
Nov. 2019	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 50 attendees</i>	SC19 Denver, CO
Nov. 2019	Half-day Tutorial: Productive Parallel Programming for FPGA with HLS <i>co-presented with J. De Fine Licht, ≈ 50 attendees</i>	SC19 Denver, CO
Jun. 2019	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 30 attendees</i>	ISC'19 Frankfurt, Germany

Spring 2019	Parallel Programming <i>co-taught with Martin Vechev, undergrad, ≈ 400 students</i>	ETH Zürich
Fall 2018	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 60 students</i>	ETH Zürich
Nov. 2018	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 70 attendees</i>	SC18 Dallas, TX
Nov. 2018	Half-day Tutorial: Productive Parallel Programming for FPGA with HLS <i>co-presented with J. De Fine Licht, ≈ 50 attendees</i>	SC18 Dallas, TX
Jun. 2018	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 35 attendees</i>	ISC'18 Frankfurt, Germany
Spring 2018	Parallel Programming <i>co-taught with Martin Vechev, undergrad, ≈ 400 students</i>	ETH Zürich
Fall 2017	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 50 students</i>	ETH Zürich
Nov. 2017	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 70 attendees</i>	SC17 Denver, CO
Jun. 2017	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 35 attendees</i>	ISC'17 Frankfurt, Germany
Spring 2017	Parallel Programming <i>co-taught with Martin Vechev, undergrad, ≈ 400 students</i>	ETH Zürich
Fall 2016	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 40 students</i>	ETH Zürich
Nov. 2016	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 70 attendees</i>	SC16 Salt Lake City, UT
Jun. 2016	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 30 attendees</i>	ISC'16 Frankfurt, Germany
Nov. 2015	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 70 attendees</i>	SC15 Austin, TX
Nov. 2015	Half-day Tutorial: Insightful Automatic Performance Modeling <i>co-presented with A. Calotoiu, M. Schulz, F. Wolf, ≈ 30 attendees</i>	SC15 Austin, TX
Sep. 2015	Half-day Tutorial: Insightful Automatic Performance Modeling <i>co-presented with A. Calotoiu, M. Schulz, F. Wolf, ≈ 20 attendees</i>	EuroMPI'15 Bordeaux, France
Sep. 2015	Full-day Tutorial: Advanced Parallel Programming with MPI <i>≈ 20 attendees</i>	Speedup'15 Lugano, Switzerland
Fall 2015	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 40 students</i>	ETH Zürich
Jun. 2015	Half-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, ≈ 40 attendees</i>	ISC'15 Frankfurt, Germany
Spring 2015	Operating Systems and Networks <i>co-taught with Adrian Perrig, undergrad, ≈ 140 students</i>	ETH Zürich
Spring 2015	Computational Science, Seminar <i>co-taught with Peter Arbenz & Petros Koumoutsakos, ≈ 15 students</i>	ETH Zürich
Spring 2015	Research Topics in Software Engineering, Seminar <i>≈ 25 students</i>	ETH Zürich
Fall 2014	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 40 students</i>	ETH Zürich
Nov. 2014	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with P. Balaji, B. Gropp, R. Thakur, ≈ 120 attendees</i>	SC14 Denver, CO, USA

Sep. 2014	Full-day Tutorial: Advanced Parallel Programming with MPI <i>invited lecturer ≈ 50 attendees</i>	EuroMPI/Asia 2014 Kobe, Japan
Fall 2013	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 35 students</i>	ETH Zürich
Jun. 2014	Full-day Tutorial: Advanced Parallel Programming with MPI <i>co-presented with Pavan Balaji, ≈ 15 attendees</i>	ISC'13 Leipzig, Germany
Spring 2014	Operating Systems and Networks <i>co-taught with Adrian Perrig, undergrad, ≈ 130 students</i>	ETH Zürich
Spring 2014	Computational Science, Seminar <i>co-taught with Peter Arbenz & Petros Koumoutsakos, ≈ 10 students</i>	ETH Zürich
Nov. 2013	Tutorial: Advanced Parallel Programming with MPI <i>co-presented with Pavan Balaji, Rajeev Thakur, James Dinan ≈ 50 attendees</i>	SC13 Denver, CO, USA
Fall 2013	Design of Parallel and High-Performance Computing <i>co-taught with Markus Pueschel ≈ 25 students</i>	ETH Zürich
Fall 2013	Research Topics in Software Engineering, Seminar <i>co-taught with Martin Vechev ≈ 20 students</i>	ETH Zürich
Jul. 2013	MPI Programming <i>Invited lecturer at CHPC Winter School, ≈ 65 students</i>	University of Johannesburg Johannesburg, South Africa
Jun. 2013	Tutorial: Advanced Parallel Programming with MPI <i>co-presented with Pavan Balaji & Martin Schulz, ≈ 15 attendees</i>	ISC'13 Leipzig, Germany
Jun. 2013	Tutorial: Advanced Parallel Programming with MPI <i>co-presented with Pavan Balaji, ≈ 25 attendees</i>	ICS'13 Eugene, OR, USA
Spring 2013	Operating Systems and Networks <i>co-taught with Donald Kossmann, undergrad, ≈ 130 students</i>	ETH Zürich
Spring 2013	Computational Science, Seminar <i>co-taught with Peter Arbenz & Petros Koumoutsakos, ≈ 5 students</i>	ETH Zürich
Feb 24	Tutorial: MPI & Advanced Parallel Programming <i>co-presented with Pavan Balaji</i>	PPoPP'13 Shenzen, China
Fall 2012	Design of Parallel and High-Performance Computing <i>co-taught with Thomas Gross & Markus Pueschel, ≈ 25 students</i>	ETH Zürich
Fall 2012	Computational Science, Seminar <i>co-taught with Peter Arbenz & Petros Koumoutsakos, ≈ 5 students</i>	ETH Zürich
Jun 17	Tutorial: Next Generation MPI Programming <i>co-presented with Martin Schulz, ≈ 25 attendees</i>	ISC'12 Hamburg, Germany
May 23-15	Tutorial: Advanced Distributed Memory Parallel Programming <i>Advanced Distributed Memory Parallel Programming: MPI-2.2, MPI 3.0 and PGAS, ≈ 35 attendees</i>	CSCS
Spring 2011	Hot Topics in HPC: Networks and Fault tolerance, CS498 <i>(4cr grad./3cr undergrad.), co-taught with Franck Cappello, ≈ 25 students</i>	University of Illinois

Advising and Mentoring

I advise(d) 19 Ph.D. students and more than 60 M.Sc. students, of which three (Patrick Iff, Marc Fischer, Cedric Renggli) received the ETH Medal for their thesis. I mentor(ed) 13 postdocs (seven with an ETH Fellowship) and hosted two SNSF Ambizione fellows (Tobias Grosse, Tal Ben-Nun) in my group.

Postdocs (chronologically)

ETH Zürich

Sabela Ramos (ETH Fellow)
Pierre Jolivet (ETH Fellow)
Tobias Grosser
Edgar Solomonik (ETH Fellow)
Tal Ben-Nun (ETH Fellow)
Tiziano De Matteis
Nikoli Dryden (ETH Fellow)
Shigang Li
Tobias Gysi
Alexandru Calotoiu
Daniele De Sensi (ETH Fellow)
Kazuki Osawa (ETH Fellow)
Nabil Abubaker (SNSF Fellow)

PhD Students (chronologically)

ETH Zürich

Maciej Besta (graduated 2022)
Tobias Gysi (graduated 2020)
Bogdan Prisacari (graduated 2020)
Salvatore di Girolamo (graduated 2021)
Grzegorz Kwasniewski (graduated 2022)
Konstantin Taranov (graduated 2022)
Johannes de Fine Licht (graduated 2022)
Alexandros Nikolaos Ziogas (graduated 2023)
Niels Gleinig (graduated 2023)
Marcin Copik
Lukas Gianinazzi
Andrei Ivanov
Saleh Ashkboos
Philipp Schaad
Langwen Huang
Marcin Chrapek
Patrick Iff
Mikhail Khalilov
Siyuan Shen
Patrik Okanvic
Luigi Fusco

Student Cluster Challenge

I am passionate about fostering young talent and encourage all undergraduate students to participate in the student cluster challenge competition. My teams have twice been among the winners!

2020-today	Advisor ETH's Team Racklette participates in the Cluster Challenge competitions at SC and ISC Preparing the (often winning) ETH team of undergraduate students for the challenge. See section "Awards of Mentees" for details on prizes won.	ETH Zurich
2019	Advisor Cluster Challenge Preparation Preparing the (winning) ETH team of undergraduate students for the challenge at ISC'19.	ETH Zurich
2008	Co-Advisor Cluster Challenge Preparation Preparing the (winning) IU/TUD team of undergraduate students for the challenge at SC'08	Indiana University
2007	Co-Advisor Cluster Challenge Preparation Preparing the IU team of undergraduate students for the challenge at SC'07.	Indiana University

Service

Leadership Service

- 2019–now **German Helmholtz AI review board**, *External scientific advisor (yearly)*
2022 **IEEE Cluster 2022**, *Program Co-Chair*
2021 **German BMBF Quantum Centers Review Board**, *External scientific advisor*
2019–2020 **SCXY**, *IEEE/ACM Supercomputing Steering Committee*
2010–now **MPI Forum Meetings**, *MPI-3 Working Group for Collective Operations and Topology*
2012–now **Green Graph 500**, *chair the Green Graph 500 list of the greenest data analytics machines.*
2014–now **Workshop on High-Performance Interconnects in the Exascale and Big-Data**, *Steering Committee*
2014–2017 **Platform for Advanced Scientific Computing Conference**, *Steering Committee*
EuroMPI'19 **European Conference on MPI (EuroMPI)**, *General Chair*
SC18 **IEEE/ACM Supercomputing**, *Papers Chair*
ACM PASC'17 **ACM Platform for Advanced Scientific Computing Conference**, *Program Co-Chair*
ICPP'17 **International Conference on Parallel Processing**, *Area Co-Chair*
IPDPS'17 **International Parallel & Distributed Processing Symposium**, *Technical Area Chair*
ACM PASC'16 **ACM Platform for Advanced Scientific Computing Conference**, *Program Co-Chair*
SC15 **IEEE/ACM Supercomputing**, *Panels Co-Chair*
HOTI'14 **IEEE Hot Interconnects**, *Tutorials Co-Chair*
SIAM PP'14 **SIAM Parallel Processing**, *Member of the Organizing Committee*
HOTI'13 **IEEE Hot Interconnects**, *General Co-Chair*
SC13 **IEEE/ACM Supercomputing**, *Emerging Technologies Chair*
EuroPar'13 **European Conference on Parallel Processing**, *Local Topic Chair for High-Performance Networks and Communication*
HOTI'12 **IEEE Hot Interconnects**, *Program Chair*
SC12 **IEEE/ACM Supercomputing**, *Technical Posters Chair*
HIPS'11 **16th International Workshop on High-Level Parallel Programming Models and Supportive Environments**, *General Chair*
HOTI'11 **IEEE Hot Interconnects**, *Program Co-Chair*
HOTI'10 **IEEE Hot Interconnects**, *Tutorials Chair*

Standardization Committees

- 2023–present **Ultra Ethernet**, *Co-chair of the Ultra Ethernet Transport Working Group (elected)*
2012–present **MPI Forum**, *Representing ETH Zurich, Chair of the Collective Operations and Topology Working Group for MPI-3.1*
2010–2012 **MPI Forum**, *Representing University of Illinois at Urbana-Champaign, Chair of the Collective Operations and Topology Working Group for MPI-3*
2007–2010 **MPI Forum**, *Representing Indiana University, Chair of the Collective Operations Working Group, Co-Author of the Chapter 5 (Collective Communication) and Chapter 7 (Process Topologies) in MPI-2.2*

Journal Editorial Boards

- 2014–present **Subject Area Editor** **SuperFri**
Supercomputing Frontiers and Innovations
2014–present **Associate Editor** **IEEE TPDS**
IEEE Transactions on Parallel and Distributed Systems
2013–present **Associate Editor** **PARCO**
Elsevier Parallel Computing Journal

Organized Workshops

Co-Chair of Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS), Organized in conjunction with ACM ICS or ACM HPDC annually from 2011-present

16th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS'11), Organized in conjunction with IEEE IPDPS'11, Anchorage, AL, USA, 2011

1st Blue Waters Performance Modeling Workshop, Organized a performance modeling workshop with speakers from the Los Alamos National Laboratory for early users of the Blue Waters Petascale system, Urbana, IL, 2010

3rd KiCC Workshop, Co-Organized 3rd workshop on Kommunikation in Clusterrechnern und Clusterverbundsystemen, Aachen 2007

2nd KiCC Workshop, Co-Organized 2nd workshop on Kommunikation in Clusterrechnern und Clusterverbundsystemen, Chemnitz 2007

1st KiCC Workshop, Co-Organized 1st workshop on Kommunikation in Clusterrechnern und Clusterverbundsystemen, Chemnitz 2005

Significant Project Involvement

Research Projects

- 2018–today **Microsoft Maia**, AI Supercomputing System (Chip, Rack-Level, System-level co-design from ground up for AI workloads)
- 2010–2013 **NSF Blue Waters**, Sustained Petaflop Computing with the Blue Waters machine. Responsible for Modelling and Simulation of Parallel Petaflop Applications
- 2008–2010 **DOE CIFTS**, Coordinated and Improved Fault Tolerance for High Performance Computing Systems
- 2007–2010 **DOE FAST-OS II**, Forum to Address Scalable Technology for Runtime and Operating Systems
- 2005–2006 **CHiC**, Co-Design and Procurement of the Chemnitzer Hochleistungs-Linux-Cluster, project volume 2.6 + 1.7 Million Euro, 528 diskless InfiniBand nodes, 8.2 TFlop/s (73.4% HPL efficiency) #117 in Top 500 June 2007