

Torsten Hoefler

Associate Professor of Computer Science, ETH Zürich

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
- 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-centric Software Development". In the context of High-Performance Computing (HPC), one can identify three sub-branches that I am actively working on: (1) **performance modeling, simulation, and optimization of large-scale parallel applications**, (2) **topologies, routing, and host interfaces of large-scale networks**, and (3) **advanced parallel programming techniques and runtime environments**.

Awards and Honors

- 2017 **Student teaching award “best interaction”** **Zurich, Switzerland**
elected democratically by computer science students
- 2017 **Best Student Paper Award at 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 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", CHF 25k
- 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 in Europe". EUR 1.5M
- 2015 **Best Student Paper Award at 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))

Universitätsstrasse 6 – 8092 Zürich

✉ htor@inf.ethz.ch • <http://htor.inf.ethz.ch/>

- 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 candidates (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 candidates (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 candidates (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 **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.

Positions and Experience

Current Significant Leadership and Service Positions

- 2013–present **ACM SIGHPC Executive Committee (2nd term)** **SIGHPC**
Member of the first and second elected Executive Committee of ACM SIGHPC, Special Interest Group on High Performance Computing. As one of two elected members-at-large, I co-represent the body of approximately 1,000 members. I was re-elected in 2016.
- 2014–present **Associate Editor (2nd term)** **IEEE TPDS**
IEEE Transactions on Parallel and Distributed Systems, I was re-appointed in 2016
- 2012–present **Associate Editor** **IJHPCA**
SAGE International Journal of High Performance Computing Applications
- 2010–now **MPI Forum WG Lead** **Message Passing Interface Forum**
I lead the MPI-3 Working Group for Collective Operations and Topology.
- 2013–now **Expert in Resilience and Software Engineering** **EESI2**
Invited member of two working groups in the European Exascale Software Initiative 2 to “provide recommendations on strategic European actions [...]”
- 2014–2016 **Scientific Advisory Board** **Simula**
Simula Research Laboratory, Norway
- 2014–now **Scientific Advisory Board** **EPIGRAM Project**
Member of the SAB of the European Project for Exascale ProGRAMming Models (EPIGRAM)

Research

- 2017–present **Associate Professor (tenured) 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.
- 2012–2017 **Assistant Professor (tenure track) 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.
- 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**
Direction des Applications Militaires (CEA-DAM)
 Bruyères-le-Châtel, France
 Parallel Quantum-Mechanical Computations with ABINIT
- Dec 2005 **Visiting Researcher** **CINECA**
CINECA Consorzio Interuniversitario
 Casalecchio di Reno, Italy
 Parallel Ab-Initio Quantum Mechanical Computations
- 2004–2006 **Research Assistant** **Chemnitz University of Technology**
 Parallel Ab-Initio Quantum Mechanical Computations, Networking Research **Chemnitz, Germany**
- Industry Experience**
- 2000–2005 **Software Engineer** **DELTA proveris AG**
 Design and Implementation of Database (Informix) and Web Applications (PHP/Perl)

Publications

total: 44 top conference papers, 12 journal papers, 57 workshop and conference papers, 2800+ citations (Google Scholar)
since 2013: 34 top conference papers, 7 journal papers, 2 workshop and conference papers, 12 best paper noms., 9 awards
 full publication list available at <http://hlor.inf.ethz.ch/publications/>

Selected Peer-reviewed Conference Publications

- 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. Gresser 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. Minkenberg, B. Prisacari, 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)*
- LCPC'13** T. Schneider, R. Gerstenberger, and T. Hoefler: Compiler Optimizations for Non-Contiguous Remote Data Movement *In Proceedings of 26th International Workshop on Languages and Compilers for Parallel Computing, LCPC'13*
- 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)*
- PPoPP'12** T. Hoefler and T. Schneider: Communication-Centric Optimizations by Dynamically Detecting Collective Operations *In Proceedings of the 17th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, PPoPP'12, pages 305–306. ACM, Feb. 2012, (poster paper) (acceptance rate (posters): 17%, 32/185)*

- PPoPP'12** F. Kjolstad, T. Hoefler, and M. Snir: Automatic Datatype Generation and Optimization *In Proceedings of the 17th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, PPoPP'12*, pages 327–328. ACM, Feb. 2012, (poster paper) (acceptance rate (posters): 17%, 32/185)
- SC'11** 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. Chaarawi, 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
- EuroPar'11** T. Schneider, S. Eckelmann, T. Hoefler, and W. Rehm.: Kernel-Based Offload of Collective Operations - Implementation, Evaluation and Lessons Learned. *In Proceedings of the 17th international conference on Parallel processing - Volume Part II, EuroPar'11*, pages 264-275. Springer, Aug. 2011 (acceptance rate 29.9%, 81/271)
- 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).
- PADL'11** E. Holk, W. E. Byrd, J. Willcock, and T. Hoefler, A. Chauhan, and A. Lumsdaine.: Kanor – A Declarative Language for Explicit Communication. *In Proceedings of the 13th international conference on Practical aspects of declarative languages, PADL'11*, pages 190–204. Springer, Jan. 2011
- HiPC'10** N. Edmonds, T. Hoefler, and A. Lumsdaine.: A Space-Efficient Parallel Algorithm for Computing Betweenness Centrality in Distributed Memory. *In Proceedings of International Conference on High Performance Computing, HiPC'10*, pages 1–10. Dec. 2010 (acceptance rate: 19.2%)
- 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)
- HiPC'09** P. Kambadur, A. Gupta, T. Hoefler, and A. Lumsdaine.: Demand-driven Execution of Static Directed Acyclic Graphs Using Task Parallelism. In *Proceedings of International Conference on High Performance Computing, HiPC'09*, pages 284–293. Dec. 2009 (acceptance rate: 11%, 35/320)
- 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.
- EuroPar'06** F. Mietke, R. Baumgartl, R. Rex, T. Mehlan, T. Hoefler, and W. Rehm.: Analysis of the Memory Registration Process in the Mellanox InfiniBand Software Stack. *Proceedings of Euro-Par 2006 Parallel Processing*, pages 124–133. Springer-Verlag Berlin, Aug. 2006. (acceptance rate: 37.9%, 110/290)

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 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.*

Peer-reviewed Workshop Publications

- LSAP'11** T. Hoefler and M. Snir.: Performance Engineering: A Must for Petaflops and Beyond. *Proceedings of Workshop on Large-Scale System and Application Performance (LSAP 2011). Keynote Paper*
- PROPER'10** T. Hoefler.: Bridging Performance Analysis Tools and Analytic Performance Modeling for HPC. *Proceedings of Workshop on Productivity and Performance (PROPER 2010). Springer, Dec. 2010. Keynote Paper*
- LSAP'10** T. Hoefler, T. Schneider, and A. Lumsdaine.: LogGOPSim - Simulating Large-Scale Applications in the LogGOPS Model. *In Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, LSAP'10, pages 597–604. ACM, Jun. 2010 Best Paper at LSAP'10*
- HIPS'09** T. Hoefler and J. L. Traeff.: Sparse Collective Operations for MPI. *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium, HIPS'09 Workshop, May 2009.*
- LSPP'09** T. Hoefler, T. Schneider, and A. Lumsdaine.: The Impact of Network Noise at Large-Scale Communication Performance. *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium, LSPP'09 Workshop, May 2009.*
- CAC'09** C. Kaiser, T. Hoefler, B. Bierbaum, and T. Bemberl.: Implementation and Analysis of Nonblocking Collective Operations on SCI Networks. *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium, CAC'09 Workshop, May 2009.*

- CAC'09** T. Hoefler, T. Schneider, and A. Lumsdaine.: A Power-Aware, Application-Based, Performance Study Of Modern Commodity Cluster Interconnection Networks. *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium, CAC'09 Workshop, May 2009.*
- CAC'08** T. Hoefler and A. Lumsdaine.: Optimizing non-blocking Collective Operations for InfiniBand. *Proceedings of the 22nd IEEE International Parallel & Distributed Processing Symposium, CAC'08 Workshop, Apr. 2008.*
- CAC'07** T. Hoefler, C. Siebert, and W. Rehm.: A practically constant-time MPI Broadcast Algorithm for large-scale InfiniBand Clusters with Multicast. *Proceedings of the 21st IEEE International Parallel & Distributed Processing Symposium, CAC'07 Workshop, page 232. IEEE Computer Society, Mar. 2007.*
- CAC'06** T. Hoefler, T. Mehlan, F. Mietke, and W. Rehm.: Fast Barrier Synchronization for InfiniBand. *Proceedings of the 20th IEEE International Parallel & Distributed Processing Symposium, CAC'06 Workshop, Apr. 2006.*
- PARS'05** T. Hoefler and W. Rehm.: A Communication Model for Small Messages with InfiniBand. *PARS Mitteilungen (German), pages 32–41. PARS, Jun. 2005. Received PARS Junior Researcher Award.*

Selected Invited Talks

9 invited keynotes, more than 24 invited talks, not counting normal conference or workshop presentations

- | | | |
|-----------------------------|--|-----------------------------|
| Keynote HPC
China'16 | Theory and Practice in HPC: Modeling, Programming, and Networking
Keynote at the HPC China 2016 conference | Xi'an, China |
| Keynote
Cluster'16 | Theory and Practice in HPC: Modeling, Programming, and Networking
Opening Keynote at the IEEE Cluster 2016 conference | Taipei, Taiwan |
| Keynote HIPS'15 | How fast will your application go? Static and dynamic techniques for application performance modeling.
Keynote at the HIPS'15/LSPP'15 combined workshop in conjunction with IPDPS'15 | Hyderabad, India |
| Keynote
LLVMHPC'14 | A case for runtime recompilation in HPC
Keynote at the LLVM Compiler Infrastructure in HPC workshop at SC14, Nov. 2014 | New Orleans, LA, USA |
| Keynote
ExaMPI'13 | MPI Beyond 3.0 and Towards Larger-Scale Computing
Keynote at the Workshop on Exascale MPI at SC13, Nov. 2013, ≈120 attendees | Denver, CO, USA |
| SC13 | The Second Green Graph500 List
Birds of a Feather, Nov. 2013 | Denver, CO, USA |
| Dagstuhl | Fault Tolerance for Remote Memory Access Programming Models
Invited to seminar "Resilience in Exascale Computing" | Dagstuhl, Germany |
| ISC'13 | The First Green Graph500 List
Birds of a Feather, Jun. 2013 | Leipzig, Germany |
| EASC'13 | Application-Centric Benchmarking and Modeling for Co-Design
Exascale Applications and Software Conference | Edinburgh, UK |
| Keynote MCC'12 | MPI-3.0: A Response to New Challenges in Hardware and Software
Keynote talk at Multicore Challenge Conference 2012 | Stuttgart, Germany |
| TiTech'12 | Optimized routing and process mapping for arbitrary network topologies
Invited talk at Tokio Institute of Technology | Tokio, Japan |
| Keynote
EuroMPI'11 | Writing Parallel Libraries with MPI - The Good, the Bad, and the Ugly
Keynote talk at 18th European PVM/MPI User's Group Meeting | Santorini, Greece |
| Keynote
EnA-HPC'11 | Energy-aware Software Development for Massive-Scale Systems
Keynote at the International Conference on Energy-Aware High Performance Computing | Hamburg, Germany |
| Jülich 2011 | Model-Driven HPC Software and System Design and Optimization
Jülich Supercomputing Center, Apr. 2011 | Jülich, Germany |
| Keynote
PROPER'10 | Analytical Performance Modeling and Simulation for Blue Waters
Keynote at the Workshop on Productivity and Performance in conjunction with EuroPar, Aug. 2010 | Ischia, Italy |
| Argonne Natl.
Laboratory | Nonblocking and Sparse Collective Operations on Petascale Computers
Argonne National Laboratory, Jun. 2010 | Chicago, IL, USA |

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

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

raised ≈\$3.22M in funding from government and industry.

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

Teaching Experience

(co)taught 3 undergraduate courses, 4 graduate courses, 6 seminars, 15 tutorials, 8 PhD committees

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
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
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 Advanced Distributed Memory Parallel Programming: MPI-2.2, MPI 3.0 and PGAS, ≈ 35 attendees	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) 6 Ph.D. students and 9 M.Sc. students.

PhD Students

ETH Zürich

Maciej Besta
Timo Schneider
Salvatore di Girolamo
Tobias Gysi
Grzegorz Kwasniewski
Bogdan Prisacari (co-advised with Cyriel Minckenberg)

Master Students

ETH Zürich

Andrea Arteaga (MS 2014), first job: MeteoSwiss
Roberto Belli (visiting MS 2014), first job: Credit Suisse
Jeremiah Bär (MS 2015)
Erik Henriksson (visiting MS 2015)

Master Students

University of Technology Chemnitz

Timo Schneider (MS 2011), first job: PhD student at ETH Zurich
Sven Eckelmann (MS 2011), first job: independent consultant
Christian Siebert (MS 2006), first job: PhD student at RWTH Aachen, Germany
Andre Lichei (MS 2006), first job: Capgemini, Munich, Germany
Mirko Reinhardt (MS 2006)

Master Students

Technical University of Dresden

Jens Domke (MS 2010), co-advised with W. Nagel, first job: Tokio Institute of Technology, Japan

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
2013	PhD Thesis Committee Member Served on the PhD committee for Omar Awile (Peter Widmayer) at D-INFK ETH Zurich	ETH Zurich
2014	PhD Thesis Committee Member Served on the PhD committee for Daniel Crisan (Lothar Thiele) at D-ITET ETH Zurich	ETH Zurich
2014	PhD Thesis Committee Member Served on the PhD committee for Ernst Gunnar Gran (Olaf Lysne) at University of Oslo	University of Oslo

- 2014 **PhD Thesis Committee Member** **University of Versailles**
Served on the PhD committee for Silvain Didelot (William Jalby) at University of Versailles
- 2014 **PhD Thesis Committee Member** **Polytechnic University of Catalonia**
Served on the PhD committee for Ana Jakanovic (Jesus Labarta) at Polytechnic University of Catalonia
- 2015 **PhD Thesis Committee Member** **University of Bordeaux**
Served on the PhD committee for Emmanuelle Saillard (Denis Barthou) at University of Bordeaux
- 2015 **PhD Thesis Committee Member** **Polytechnic University of Catalonia**
Served on the PhD committee for German Llort (Jesus Labarta) at Polytechnic University of Catalonia
- 2015 **PhD Thesis Committee Member** **ETH Zurich**
Served on the PhD committee for Animesh Trivedi (Thomas Gross) at D-INFK ETH Zurich

Service

Leadership Service

- 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–now **Platform for Advanced Scientific Computing Conference**, *Steering Committee*
- 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**, *Area Co-Chair*
- ACM PASC'16 **ACM Platform for Advanced Scientific Computing Conference**, *Program Co-Chair*
- HOTI'14 **Hot Interconnects**, *Tutorials Co-Chair*
- SIAM PP'14 **SIAM Parallel Processing**, *Member of the Organizing Committee*
- HOTI'13 **Hot Interconnects**, *General Co-Chair*
- EuroPar'13 **European Conference on Parallel Processing**, *Local Topic Chair for High-Performance Networks and Communication*
- ROSS'13 **International Workshop on Runtime & Operating Systems for Supercomputers**, *General Co-Chair*
- HOTI'12 **Hot Interconnects**, *Program Chair*
- SC'12 **IEEE/ACM Supercomputing**, *Technical Posters Chair*
- ROSS'12 **International Workshop on Runtime & Operating Systems for Supercomputers**, *General Co-Chair*
- HIPS'11 **16th International Workshop on High-Level Parallel Programming Models and Supportive Environments**, *General Chair*
- ROSS'11 **International Workshop on Runtime & Operating Systems for Supercomputers**, *General Co-Chair*
- HOTI'11 **Hot Interconnects**, *Program Co-Chair*
- HOTI'10 **Hot Interconnects**, *Tutorials Chair*

Standardization Committees

- 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*

Advisory Service

- 2014 **NSF Reviewer (Big Data call)** **NSF, Washington, DC**
(reviewing grant proposals asking for ≈\$1M per proposal)

2013	Proposal Reviewer (reviewing grant proposals asking for ≈\$400.000 per proposal)	SNF, Switzerland
2012	DOE Exploratory Panel Performance Modeling Early Explorations (basis for future funding decisions, resulted in DOE workshop and program)	DOE, Washington, DC
2011	NSF Review Panel (reviewing grant proposals asking for ≈\$500.000 per proposal)	NSF, Washington, DC
2010	Scientific Software Innovation Institute for Quantum Chemistry Exploratory Workshop	NSF, Washington, DC
2010	Scientific Software Innovation Institute for Parallel Tools Exploratory Workshop	NSF, Washington, DC

Journal Editorial Boards

2014–present	Subject Area Editor Supercomputing Frontiers and Innovations	SuperFri
2014–present	Associate Editor IEEE Transactions on Parallel and Distributed Systems	IEEE TPDS
2013–present	Associate Editor Elsevier Parallel Computing Journal	PARCO
2012–present	Associate Editor SAGE International Journal of High Performance Computing Applications	IJHPCA

Technical Program Committee Member (alphabetically)

AMP/PLDI	Advances in Message Passing, 2010
Big Data	IEEE Conference on Big Data, 2013
CACHES/ICS	Characterizing Applications for Heterogeneous Exascale Systems, 2011
AsHES/IPDPS	Accelerators and Heterogeneous Exascale Systems, 2012
CASS/IPDPS	Comm. Architecture for Scalable Systems (formerly CAC), 2010, 2011, 2012, 2013
CCGrid	IEEE Symposium on Cluster Computing and the Grid, 2009, 2010, 2011, 2012, 2013, 2014
Cluster	IEEE Conference on Cluster Computing, 2010, 2012
E2SC/SC	International Workshop on Energy Efficient Supercomputing, 2015
ESPAS/HiPEAC	Extreme Scale Parallel Architectures and Systems, 2012
EuroMPI	former EuroPVM/MPI, 2009, 2010, 2011, 2012, 2013
HiPC	High Performance Computing Conference, 2011, 2012, 2013
HIPS/IPDPS	High-Level Par. Programming Models and Supportive Environments, 2011, 2012, 2013
HotI	IEEE Hot Interconnects, 2009, 2010, 2011, 2012, 2013
HPCC	High Performance Computing and Communications, 2011, 2012
ICPP	Intl. Conference on Parallel Processing, 2012, 2013
ICS	ACM International Conference on Supercomputing, 2011, 2014
IPDPS	IEEE Intl. Parallel & Distributed Processing Symposium, 2010, 2013, 2015
ISC	International Supercomputing Conference, 2012, 2014 (tutorials)
IWAPT	International Workshop on Automatic Performance Tuning, 2014, 2017
PACT	Parallel Architectures and Compilation Techniques, 2015
LLVM/SC	Workshop on the LLVM Compiler Infrastructure in HPC, 2015, 2016
NRE/SC	Workshop on Numerical reproducibility in high-performance computing, 2015, 2016
PARCO	Parallel Computing, 2014
P2S2/ICPP	Par. Progr. Models and Systems Software for High-End Comp., 2010, 2011, 2012
PGAS	Partitioned Global Address Space Programming Models, 2014

- PPoPP **Symposium on Principles and Practice of Parallel Programming**, 2012, 2013 (ERC), 2014, 2016 (ERC)
- SC **ACM/IEEE Conference on High Performance Computing**, 2010, 2011, 2012, 2013 (stepped down), 2014 (stepped down)
- SC **ACM/IEEE SC Doctoral Showcase Committee**, 2012
- SC **ACM/IEEE SC Technical Posters Committee**, 2012, 2013
- SPAA **Symposium on Parallelism in Algorithms and Architectures**, 2013
- HUCAA/HPDC **Heterogeneous Unconventional Cluster Architectures and Applications**, 2011, 2012, 2013

Technical Program Committee Reviewer/Scientific Journal Reviewer

- AJSE **Arabian Journal for Science and Engineering**, 2011, 2012
- CACM **Communications of the ACM**, 2013
- CAC/IPDPS **Communication Architecture for Clusters**, 2008
- CCPE **Elsevier Concurrency and Computation: Practice and Experience**, 2011, 2012
- CISE **IEEE Computing in Science and Engineering**, 2010
- Cluster **IEEE Conference on Cluster Computing**, 2007, 2008, 2010
- Cluster **Springer Journal on Cluster Computing**, 2014
- Computer **IEEE Computer**, 2009
- Euro-Par **Euro-Par**, 2010, 2011, 2012
- GPCE **Intl. Conference on Generative Programming**, 2013
- ICPP **Intl. Conference on Parallel Processing**, 2010
- IJHPCA **Intl. Journal of High Performance Computing Applications**, 2009
- IJPEDS **Intl. Journal of Parallel, Emergent and Distributed Systems**, 2008
- IPDPS **IEEE Intl. Parallel & Distributed Processing Symposium**, 2008, 2009, 2010, 2011, 2012
- JPDC **Intl. Journal of Parallel and Distributed Computing**, 2011
- PGAS **Partitioned Global Address Space Conference**, 2011
- SC **ACM/IEEE Conference on High Performance Computing**, 2007, 2008, 2009, 2010, 2011, 2012
- SISC **SIAM Journal on Scientific Computing**, 2014
- SPAA **Symposium on Parallelism in Algorithms and Architectures**, 2009
- TPDS **IEEE Transactions on Parallel and Distributed Systems**, 2008, 2009, 2012
- TCOM **IEEE Transactions on Communications**, 2012
- PARCO **Elsevier Parallel Computing**, 2011

Organized Workshops

- Co-Chair of Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS'14)**, Organized in conjunction with ACM ICS'14, Munich, Germany, 2014
- Co-Chair of Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS'13)**, Organized in conjunction with ACM ICS'13, Eugene, OR, 2013
- Co-Chair of Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS'12)**, Organized in conjunction with ACM ICS'12, Venice, Italy, 2012
- Co-Chair of Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS'11)**, Organized in conjunction with ACM ICS'11, Tucson, AZ, 2011
- 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*

Professional Organizations

IEEE Computer Society, *Member*

Association for Computing Machinery (ACM), *Member*

ACM SIGHPC, *Member and Member at Large (elected)*

Significant Project Involvement

Research Projects

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*

Software Projects

2008–present **LogGOPSim**, *Network performance simulator using the LogGOPS model*

2006–present **Netgauge**, *Network performance measurement tool (open source)*

2008 **ORCS**, *Oblivious Routing Congestion Simulator (completed successfully)*

2006–2008 **LibNBC**, *Implementation of Nonblocking Collective Operations (completed successfully)*

2006–2008 **Open MPI**, *Open source MPI implementation (contributed to collectives framework)*