308 VINYARD AVENUE • BLACKSBURG, VA 24060
PHONE: (540) 315-1545 • E-mail: wuchun@thefengs.com • WWW: http://www.cs.vt.edu/~feng/ and http://synergy.cs.vt.edu/

WU-CHUN FENG

PERSONAL

Nickname: Wu. Gender: Male. Birth Place: Ann Arbor, Michigan, U.S.A. (U.S. Citizen)

RESEARCH INTERESTS

High-performance computing and networking, energy-efficient supercomputing, multi- and many-core computing, accelerator-based computing, cloud computing, grid computing, performance modeling and optimization, bioinformatics and the life sciences, and computer science pedagogy for K-12.

For additional information, please visit the following web sites: http://seec.cs.vt.edu/, http://synergy.cs.vt.edu/, http://www.green500.org/, http://biocomputing.cs.vt.edu/, http://www.mpiblast.org/, http://sss.cs.vt.edu/, https://ppp.cs.vt.edu/, http://psnap.cs.vt.edu/, and http://myvice.cs.vt.edu/.

EDUCATION

8/90 – 5/96 University of Illinois at Urbana-Champaign, Urbana, IL GPA: 5.00/5.00
Ph.D. in Computer Science
Thesis: *Applications and Extensions of the Imprecise-Computation Model*
Advisor: Professor Jane W.-S. Liu

8/88 – 12/89 The Pennsylvania State University, University Park, PA GPA: 4.00/4.00
M.S. in Computer Engineering
Thesis: *A Learn-by-Example, Natural-Language Processor Based on Case-Frame Instantiation*
Advisor: Professor Rangachar Kasturi

8/84 – 5/88 The Pennsylvania State University, University Park, PA GPA: 3.82/4.00
B.S. in Computer Engineering & B.S. Honors in Music

Computer Engineering
Thesis: *A Natural-Language Interface for a Paper-Based Map Information System*
Advisor: Professor Rangachar Kasturi

Music
Thesis: *Compositions in Popular Music for the Piano*
Advisor: Professor W. Bruce Trinkley

PROFESSIONAL EXPERIENCE

Employment

Virginia Tech; Blacksburg, VA, USA
January 2006 – present

1. *Site Director*; NSF Center for Space, High-Performance, & Resilient Computing (SHREC); October 2018 – present.
2. *Director*; Synergistic Environments for Experimental Computing (SEEC) Center, July 2014 – present.
3. *Professor*; Dept. of Computer Science, June 2013 – present.
4. *Faculty Member*; Wireless @ VT, May 2012 – present.
5. *Elizabeth & James E. Turner Fellow*; Dept. of Computer Science; March 2012 – April 2020.
6. *Director*; SyNeRGy Laboratory; January 2006 – present.

7. *Director* ; CUDA Research Center; May 2011 – April 2016.
8. *Site Co-Director* ; NSF Center on High-Performance Reconfigurable Computing (CHREC); January 2009 – December 2018.
9. *Associate Professor* ; Dept. of Computer Science; January 2006 – June 2013.
10. *Faculty Member* ; Center for High-End Computing Systems; January 2006 – May 2013.

Los Alamos National Laboratory; Operated by the University of California; Los Alamos, NM, USA
October 1998 – January 2006

1. *Founder, Team Leader, and Technical Staff Member;* Research & Development in Advanced Network Technology (RADIANT); Computer & Computational Sciences Division; October 2000 – January 2006.
   - Staff: 2-3 professionals, 2-3 postdocs, 3-5 students, 3-5 cross-team collaborators, 10-15 external collaborators.
   - Technical lead & principal investigator on projects in the following areas: high-performance networking and computing, network traffic characterization, monitoring tools, network protocols (OS-based and OS-bypass), computational & data grids, resource management, bioinformatics, and cyber-security.
   - Director of the Advanced Summer Curriculum for Emerging Network Technologies (ASCENT), a challenging program for students interested in learning about how to do networking-related research.
2. *Technical Staff Member (tenured in September 2000);* Network Engineering Group; Computing, Information, and Communications Division, October 1998 – October 2000.
   - Proposal writing and research in network traffic characterization, high-performance TCP, high-speed network interface cards, OS-bypass protocols, and distributed resource management.

University of Illinois at Urbana-Champaign; Urbana, IL, USA
August 1990 – May 1998

1. *Visiting Assistant Professor;* Department of Computer Science; College of Engineering; August 1996 – May 1998.
   - Teaching: computer architecture, software engineering, and computer networks.
   - Research: real-time systems, networks, and multimedia (see work at Vosaic).
2. *Graduate Research Assistant;* Department of Computer Science; College of Engineering; August 1991 – August 1996.
   - Real-time systems and networks, end-to-end scheduling, and imprecise computation.
   - Performance metering & compiler optimization of concurrent object-oriented programs.
3. *Graduate Teaching Assistant;* Department of Computer Science; College of Engineering; August 1990 – August 1991.
   - Artificial intelligence, object-oriented programming, communication networks for computers.

IBM T.J. Watson Research Center; Yorktown Heights, NY, USA
January 1990 – July 1990

1. *Applications Research Programmer;* January 1990 – July 1990.
   - Systems integration and research of speech, handwriting, and gesture recognition systems.

Penn State University; University Park, PA, USA
July 1988 – December 1989

1. *Dean's Fellow;* Department of Electrical Engineering; College of Engineering; August 1988 – December 1989.
   - Teaching: Digital design, VLSI system design, electrical circuits, power distribution.
2. *Technical Coordinator;* NSF-funded Young Scholars Academy; College of Engineering; July 1988 – August 1988.
   - Teaching, computer curriculum development, and software development & maintenance in Fortran and C.

Consulting, Courtesy, and Other Appointments (ordered by date of completion in reverse chronological order)

Adjunct Professor; Dept. of Electrical & Computer Engineering; Virginia Tech; Blacksburg, VA, USA.
January 2006 – present.

- Courtesy appointment to foster and sustain collaboration at the interface of computer hardware and software (Associate Professor: January 2006 – June 2013. Full Professor: July 2013 – present.)

Founder and Project Lead; The Green500; Blacksburg, VA, USA.
November 2007 – present.

- A grassroots movement that seeks to raise awareness in the energy efficiency of supercomputers and ensure that supercomputers only simulate climate change, not create it.

Affiliate Faculty; Dept. of Biomedical Engineering and Mechanics (BEAM); School of Biomedical Engineering and Sciences (SBES); Wake Forest University; Winston-Salem, NC, USA.
February 2017 – February 2020.

- Affiliate appointment to foster collaboration in biomedical engineering and sciences and biocomputing.
- Participation in BEAM activities, including potential teaching and committee appointments.
- Mentorship of BEAM students and postdocs.

Founder and Project Lead; The mpiBLAST Project; Blacksburg, VA, USA.
July 2002 – June 2019.

- A project that provides a freely available, open-source, parallel implementation of NCBI BLAST.

President and CEO; Abokia; Blacksburg, VA, USA.
September 2010 – January 2016.

- A company that seeks to commoditize and massively accelerate biocomputing solutions.

Adjunct Faculty Member; Virginia Bioinformatics Institute; Virginia Tech; Blacksburg, VA, USA.
April 2011 – April 2012.

- Adjunct appointment to foster and sustain collaboration in bioinformatics.

Adjunct Professor; Dept. of Cancer Biology and Translational Science Institute; Wake Forest University School of Medicine; Winston-Salem, NC, USA.
October 2010 – June 2011.

- Courtesy appointment to bridge at the interface of computer science and bioinformatics.

Chief Technology Officer; EnergyWare; Blacksburg, VA, USA.
September 2008 – December 2010.

- A company that commercialize energy-efficient software solutions for commodity hardware platforms.

Chief Scientist; Orion Multisystems; Santa Clara, CA, USA.
June 2004 – February 2006.

- Software and hardware architect for desktop cluster workstations. (Ultimately, acquired by Sun Microsystems.)

Member of the Board of Directors; Compufarm; Palo Alto, CA, USA.
December 2003 – May 2004.

- Non-profit organization established to provide ubiquitous technology infrastructure (i.e., grid of clusters) to rural as well as urban (but economically disadvantaged) areas.

Institute Fellow; Los Alamos Computer Science Institute; Los Alamos National Lab; Los Alamos, NM, USA.
November 1998 – January 2006.

- "Think tank" for fundamental and applied research in computer science.

Adjunct Assistant Professor; Dept. of Computer & Information Sciences; Ohio St. Univ.; Columbus, OH, USA.
April 2000 – April 2002 and August 2002 – August 2003.

- Courtesy appointment to facilitate collaboration between academia and government lab.

Adjunct Assistant Professor; Dept. of Electrical & Computer Engineering; Purdue Univ.; W. Lafayette, IN, USA.
October 1998 – October 2000.

- Courtesy appointment to facilitate collaboration between academia and government lab.

Research Scientist; Vosaic Corporation; Champaign, IL, USA.
February 1997 – June 1997.

- VDP: Video Datagram Protocol (streaming of audio and video over the Internet).
- AC-3™: Java-based real-time audio decoder.

Research Consultant; Spacecraft Data Systems; NASA Ames Research Center; Mountain View, CA, USA.
June 1993 – August 1993.

- Real-time scheduling on Space Station Freedom (now International Space Station).

TEACHING EXPERIENCE

Virginia Tech

Undergraduate Courses

- CS 2505: Introduction to Computer Organization I (Spring 2008, Fall 2008, Spring 2009, Fall 2009, Spring 2011)
- CS 2506: Introduction to Computer Organization II (Fall 2018, Fall 2019)
- CS 4234: Parallel Computation (Fall 2016, Fall 2017, Fall 2020)
- CS 4254: Computer Network Architecture & Programming (Spring 2010)
- CS 4504: Computer Architecture (Spring 2012, Spring 2014, Spring 2016)
- CS 4984: Accelerator-Based Parallel Computing (Spring 2009)

Graduate Courses

- CS 5234: Advanced Parallel Computation (Spring 2017, Spring 2018, Spring 2021)
- CS/ECE 5504: Computer Architecture (Spring 2012, Spring 2015, Spring 2020)
- CS/ECE 5565: Network Architecture and Protocols (Spring 2006, Spring 2007, Spring 2008, Spring 2010)
- CS 6204: Advanced Topics in High-Performance Networking (Spring 2006)
- CS 6504: Advanced Networking for High-Performance Computing (Fall 2007)
- CS 6504: Heterogeneous Multicore Computing (Spring 2011)
- CS 6504: Co-Design for Parallel Computing (Spring 2013)
- CS 6504: Accessible Parallel Computing (Spring 2015)
- CS 6504: Quantum Computing for Computer Scientists and Engineers (Spring 2019)

University of Illinois at Urbana-Champaign

Undergraduate Courses

- CS 105: Introduction to Computing (Spring 1991, Fall 1996, Spring 1997, Fall 1997, Spring 1998)
- CS 232: Computer Organization II (Fall 1996, Spring 1997, Fall 1997)

Graduate/Undergraduate Course

- CS/ECE 338: Communication Networks for Computers (Summer 1991, Spring 1998)
- CS/ECE 348: Introduction to Artificial Intelligence (Fall 1990)

Penn State University

Undergraduate Courses

- CMPEN 271: Digital System Design (Fall 1989)
- EE 220: Electrical Circuits and Power Distribution (Spring 1989)

Graduate/Undergraduate Course

- CMPEN 449: VLSI System Design (Fall 1988)

PUBLICATIONS **H-INDEX: 53 (GOOGLE SCHOLAR, DECEMBER 2020)**

In the following list of publications, an asterisk (*) denotes co-authors who were undergraduate students, graduate students, or postdoctoral research associates supervised by Feng. A double asterisk (**) indicates co-supervision by Feng.

Books, Book Chapters, and Book Articles

B1. J. A. Gomez-Pulido, B. Schmidt, W. Feng (Guest Editors), *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB), Special Issue: Accelerating Bioinformatics Applications via Emerging Parallel Computing Systems*, 12(5): 971-972, September-October 2015.

B2. W. Feng, *The Green Computing Book: Tackling Energy Efficiency at Large Scale*, Chapman and Hall / CRC Press, 2014.

B3. W. Feng (Editor), *The Green Computing Book: Tackling Energy Efficiency at Large Scale*, Chapman and Hall / CRC Press, 2014. Book Chapter by C. Hsu, W. Feng, and S. Poole, "An Adaptive Run-Time System for Improving Energy Efficiency"

B4. W. Feng, K. Cameron, and T. Scogland, "The Green500 List: A Look Back to Look Forward" in *Contemporary High-Performance Computing from Petascale to Exascale*, J. Vetter (Editor), Chapman and Hall / CRC Press, pp. 31-41, April 2013.

B5. P. Balaji, Q. Zhu, and W. Feng, "Virtualization Techniques for Graphics Processing Units" in *Scalable Computing: Theory and Practice*, John Wiley & Sons, pp. 675-697, January 2013.

B6. C. Hsu and W. Feng, "Embracing the Memory and I/O Walls for Energy-Efficient Scientific Computing" in *Energy-Aware Distributed Computing Systems*, A. Zomaya and Y. Lee (Editors), Wiley Series on Parallel and Distributed Computing, pp. 417-439, August 2012.

B7. H. Lin*, G. Fedak, and W. Feng, "Data-Intensive Computing on Desktop Grids," in *Desktop Grid Computing*, C. Cerin and G. Fedak (Editors), Chapman & Hall/CRC Numerical Analysis and Scientific Computing Series, pp. 237-256, June 2012.

B8. W. Feng and H. Lin, "From Homology to Sequence Alignment" in *Encyclopedia of Parallel Computing*, D. Padua (Editor), Springer, November 2011.

B9. W. Feng, Y. Cao, D. Patnaik*, and N. Ramakrishnan, "Temporal Data Mining for Neuroscience" in *GPU Computing Gems,* W. Hwu (Editor), Elsevier/Morgan-Kaufmann, February 2011.

B10. W. Feng and P. Balaji, "Ethernet vs. EtherNOT" in *Attaining High Performance Communication: A Vertical Approach,* A. Gavrilovska (Editor), CRC Press, 2009.

B11. W. Feng and P. Balaji (Guest Editors), *IEEE Computer, Special Issue: Tools & Environments for Multi- and Many-Core Architectures*, September 2009.

B12. W. Feng and D. Manocha (Guest Editors), *Journal of Parallel Computing, Special Issue: High-Performance Computing Using Accelerators*, Vol. 33, No. 10-11, November 2007.

B13. W. Feng and J. Duato (Editors), *2005 International Conference on Parallel Processing*, Vol. 34, Los Alamitos, CA, USA, 2005. (ISBN: 0-7695-2380-3.)

B14. A. Mezzacappa, J. Chen, P. Cummings, J. Dongarra, J. Drake, W. Feng, W. Gropp, J. Hack, R. Harrison, C. Johnson, D. Keyes, J. Mitchell, R. Mount, E. Ng, R. Ryne, N. Samatova, R. Sugar, and W. Tang (Editors), SciDAC 2005: Scientific Discovery through Advanced Computing, *Journal of Physics: Conference Series*, Vol. 16, Institute of Physics Publishing, Temple Back, Bristol, United Kingdom, 2005. (ISSN: 1742-6588.)

B15. W. Feng, "Green Destiny + mpiBLAST = Bioinfomagic," Book Chapter in *Advances in Parallel Computing* Series, Vol. 13, G.R. Joubert, W.E. Nagel, F.J. Peters, and W.V. Walter (Editors), Elsevier. 2004.

B16. R. Devon and W. Feng, *Fortran at the Keyboard*, Kendall/Hunt Publishing Company, Dubuque, IA, September 1989.

Reports & Think-Tank Meetings

R1. NSF Workshop Report on Exploiting Parallelism and Scalability (XPS), Arlington, VA, June 1-2, 2005.

R2. Institute for Computing in Sciences (ICiS) Think-Tank Workshop on *Future of the Field*, Park City, UT, August 1-6, 2012.

R3. Institute for Computing in Sciences (ICiS) Think-Tank Workshop on *Sequence Analysis Library*, Park City, UT, July 29-August 3, 2012.

R4. Institute for Computing in Sciences (ICiS) Think-Tank Workshop on *Future of the Field*, Park City, UT, July 30-August 6, 2011.

R5. Institute for Computing in Sciences (ICiS) Think-Tank Workshop on *Future of the Field*, Snowbird, UT, July 24-31, 2010.

R6. Institute for Computing in Sciences (ICiS) Think-Tank Workshop on *Computational Methods and Terabase Metagenomics*, Snowbird, UT, July 17-24, 2010.
- Meeting Report: The Terabase Metagenomics Workshop and the Visiono of an Earth Microbiome Project → Published in Standards in Genomic Sciences

R7. DOE Workshop on Cross-Cutting Technologies for Computing at the Exascale, February 2010.

R8. Opportunities in Biology at the Extreme Scale of Computing, August 2009. Contributions include
- W. Feng, "Massively Parallel Sequence Alignment: When Computational Biology Becomes I/O Biology," Session: Populations, Communities, Ecosystems and Evolutionary Dynamics: Genomics and Metagenomics.
- W. Feng, "Accelerating the Computation of Long-Range Interactions: Towards Rational Drug Design," Session: Populations, Communities, Ecosystems and Evolutionary Dynamics: Genomics and Metagenomics.

R9. Final Report on the NSF Workshop on Instrumentation Needs of Computer and Information Science and Engineering Research (held on July 12-13, 2008 in conjunction with CRA Snowbird), 24 pages. Sponsor: Rita Rodriguez, NSF.

R10. NSF Workshop Report on Petascale Computing in the Biological Sciences, August 2006. Invited Contributor.

R11. DOE Science Networking Challenge: Roadmap to 2008, August 2003. Primary Contributor and Working Group Chair. http://www.es.net/hypertext/welcome/pr/Roadmap/Roadmap%20to%202008.pdf.

R12. DOE Workshop on Ultra High-Speed Transport and Dynamic Provisioning for Large-Scale Science Applications, April 2003. Program Vice-Chair and Working Group Chair with Prof. Donald Towsley of the University of Massachusetts.

R13. NSF/DOE Large-Scale Network Security Workshop, May 2003. Primary Contributor.

Refereed Articles in Archival Journals (unless otherwise noted)

J1. Q. Al Hajri**, S. Dash**, W. Feng, H. Garner, and R. Anandakrishnan, "Identifying Multi-Hit Carcinogenic Gene Combinations: Scaling Up a Weighted Set Cover Algorithm Using Compressed Binary Matrix Representation on a GPU," *Scientific Reports: A Nature Research Journal*, 10(1):2022-2048, February 2020.

J2. S. Pumma**, M. Si, W. Feng, and P. Balaji, "Scalable Deep Learning via I/O Analysis and Optimization," *ACM Transactions on Parallel Computing (TOPC)*, 6(2): 6:1-6:34, July 2019.

J3. X. Yu*, H. Wang*, W. Feng, H. Gong, and G. Cao, "GPU-Based Iterative Medical CT Image Reconstructions," *Journal of Signal Processing Systems*, 91(3-4):321-338, March 2019.

J4. S. Dash**, N. Kinney, R. Varghese, H. Garner, W. Feng, and R. Anandakrishnan, "Differentiating between Cancer and Normal Tissue Samples using Multi-hit Combinations of Genetic Mutations," *Scientific Reports: A Nature Research Journal*, 9(1):1005-1017, January 2019.

J5. S. Banerjee**, H. Zhu*, M. Tang*, X. Wu, W. Feng, and H. Xie, "Identifying Transcriptional Regulatory Modules among Different Chromatin States in Mouse Neural Stem Cells," *Frontiers in Genetics,* 9:731-745, January 2019.

J6. K. Hou*, H. Wang*, and W. Feng, "A Framework for the Automatic Vectorization of Parallel Sort on x86-based Processors," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 29(5):958-972, May 2018.

J7. A. McCaskey, E. Dumitrescu, D. Liakh, M. Chen**, W. Feng, and T. Humble, "A Language and Hardware-Independent Approach to Quantum–Classical Computing, *SoftwareX*, 7:245-254, January-June 2018.

J8. S. Pumma**, P. Phunchongharn, S. Chapeland, and T. Achalakul, "A Runtime Estimation Framework for ALICE," *Future Generation Computer Systems*, 72:65-77, July 2017.

J9. A. Feng, M. Gardner, and W. Feng, "Parallel Programming with Pictures is a Snap!," *Journal of Parallel & Distributed Computing (JPDC)*, *Special Issue: Keeping up with Technology: Teaching Parallel, Distributed and High-Performance Computing,* 105:150-162, July 2017. January 2017 (online version).

J10. J. Zhang*, S. Misra, H. Wang*, and W. Feng, "muBLASTP: Database-Indexed Protein Sequence Search on Multicore CPUs," *BMC Bioinformatics, 17(1):443*, November 2016.

J11. X. Shu**, J. Zhang*, D. Yao, and W. Feng, "Fast Detection of Transformed Data Leaks," *IEEE Transactions on Information Forensics & Security (TIFS)*, PP(99): 1-16, March 2016.

J12. T. Scogland**, W. Feng, B. Rountree, and B. de Supinski, "CoreTSAR: Core Task-Size Adapting Runtime," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 26(11): 2970-2983, November 2015.

J13. J. Zhang*, H. Wang*, and W. Feng, "cuBLASTP: Fine-Grained Parallelization of Protein Sequence Search on CPU+GPU," *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 14(4):830-843, October 2015.

J14. K. Krommydas*, W. Feng, C. Antonopolous, and N. Bellas, "OpenDwarfs: Characterization of Dwarf-Based Benchmarks on Fixed and Reconfigurable Architectures," *Journal of Signal Processing Systems*, October 2015.

J15. A. Aji*, L. Panwar*, F. Ji**, K. Murthy, M. Chabbi, P. Balaji, K. Bisset, J. Dinan, W. Feng, J. Mellor-Crummey, X. Ma, R. Thakur, "MPI-ACC: Accelerator-Aware MPI for Scientific Applications," *IEEE Transactions on Parallel and Distributed Systems*, PP(99): 1-14, June 2015.

J16. B. P. Pickering**, C. W. Jackson, T. R. W. Scogland*, W. Feng, and C. Roy, "Directive-based GPU Programming for Computational Fluid Dynamics," *Computers and Fluids*, 114(2015): 242-253, March 2015.

J17. J. Yin**, J. Zhang**, J. Wang, and W. Feng, "SDAFT: A Novel Scalable Data Access Framework for Parallel BLAST," *Parallel Computing*, 40(10): 697-709, December 2014.

J18. E. Kolker, V. Özdemir, L. Martens, W. Hancock, E. Dumbill, N. Anderson, S. Aynacioglu, A. Baranova, S. Campagna, R. Chen, J. Choiniere, S. Dearth, W. Feng, L. Ferguson, G. Fox, R. Grossman, A. Heath, R. Higdon, M. Hutz, I. Janko, L. Jiang, S. Joshi, A. Kel, J. Kemnitz, I. Kohane, N. Kolker, D. Lancet, E. Lee, W. Li, A. Lisitsa, A. Llerena, C. MacNealy-Koch, J.-C. Marshall, P. Masuzzo, A. May, G. Mias, E. Montague, S. Mooney, A. Nesvizhskii, S. Noronha, G. Omenn, H. Rajasimha, J. Sheehan, L. Smarr, C. Smith, T. Smith, M. Snyder, R. Srikanth, S. Srivastava, L. Stanberry, E. Stewart, S. Toppo, P. Uetz, K. Verheggen, B. Voy, L. Warnich, S. Wilhelm, and G. Yandl, "Towards More Transparent and Reproducible Omics Studies through a Common Metadata Checklist and Data Publication," *OMICS: A Journal of Integrative Biology, Special Issue on BIG DATA*, January 2014.

J19. M. Gardner, P. Sathre*, W. Feng, and G. Martinez*, "Characterizing the Challenges and Evaluating the Efficacy of a CUDA-to-OpenCL Translator," *Journal of Parallel Computing, Special Issue on Programming Models, Systems Software and Tools for High-End Computing,* 39(12): 769-786, December 2013.

J20. H. Lin*, X. Ma, and W. Feng, "Reliable MapReduce Computing on Opportunistic Resources," *Journal of Cluster Computing*, 15(2): 145-161, June 2012.

J21. K. Lee*, H. Lin*, and W. Feng, "Performance Characterization of Data-Intensive Kernels on AMD Fusion Architectures," *Computer Science – Research & Development*, 10 pages, Springer-Verlag, June 2012.

J22. B. Subramaniam* and W. Feng, "GBench: Benchmarking Methodology for Evaluating the Energy Efficiency of Supercomputers," *Computer Science – Research & Development*, 10 pages, Springer-Verlag, June 2012.

J23. T. Scogland*, B. Subramaniam*, and W. Feng, "The Green500 List: Escapades to Exascale," *Computer Science – Research & Development*, 9 pages, Springer-Verlag, June 2012.

J24. M. Daga* and W. Feng, "Multi-Dimensional Characterization of Electrostatic Surface Potential Computation on Graphics Processors," *BMC Bioinformatics*, 13(Suppl 5):S4, 12 pages, April 2012.

J25. K. Bisset, A. Aji**, M. Marathe, and W. Feng, "High-Performance Biocomputing for Simulating the Spread of Contagion over Large Contact Networks (Extended Version)," *BMC Genomics*, 13(Suppl 2):S3, 12 pages, April 2012.

J26. Y. Cao, D. Patnaik**, S. Ponce**, J. Archuleta*, P. Butler**, W. Feng, and N. Ramakrishnan, "Parallel Mining of Neuronal Spike Streams on Graphics Processing Units," *International Journal of Parallel Processing*, 40(6): 605-632, July 2011.

J27. F. Ge, A. Eltawil, H. Lin*, W. Feng, C. J. Chiang, and R. Chadha, "Augment Mobile Cloud Computing with Cognitive Radio," *IEEE COMSOC MMTC E-Letter*, Vol. 6, No. 10, 4 pages, October 2011.

J28. J. Faris, E. Kolker, A. Szalay, L. Bradlow, E. Deelman, W. Feng, J. Qiu, D. Russell, E. Stewart, and E. Kolker, "Communication and Data-Intensive Science in the Beginning of the 21st Century," *OMICS: A Journal of Integrative Biology,* 15(4):213-215, April 2011.

J29. H. Lin**, X. Ma, W. Feng, and N. Samatova, "Coordinating Computation and I/O in Massively Parallel Sequence Search," *IEEE Transactions on Parallel & Distributed Systems*, 22(4): 529-543, April 2011.

J30. J. Gilbert, F. Meyer, D. Antonopoulos, P. Balaji, C. T. Brown, N. Desai, J. A. Eisen, D. Evers, D. Field, W. Feng, D. Huson, J. Jansson, R. Knight, J. Knight, E. Kolker, K. Konstantindis, J. Kostka, N. C. Kyrpides, R. Mackelprang, A. McHardy, C. Quince, J. Raes, A. Sczyrba, A. Shade, R. Stevens, "The Terabase Metagenomics Workshop and the Vision of an Earth Microbiome Project," *Standards in Genomic Sciences*, 3(3): 243-248, December 2010.

J31. P. Balaji, W. Feng, H. Lin*, J. Archuleta*, S. Matsuoka, A. Warren**, J. Setubal, E. Lusk, R. Thakur, I. Foster, D. Katz, S. Jha, K. Shinpaugh, S. Coghlan, and D. Reed, "Global-Scale Distributed I/O with ParaMEDIC," *Concurrency and Computation: Practice and Experience*, 22(16): 2266-2281, November 2010.

J32. T. Scogland*, H. Lin*, and W. Feng, "A First Look at Integrated GPUs for Green High-Performance Computing," *Computer Science – Research & Development*, 25(3-4): 125-134, September 2010.

J33. R. Anandakrishnan**, T. Scogland*, A. Fenley**, J. Gordon**, W. Feng, and A. Onufriev, "Accelerating Electrostatic Surface Potential Calculation with Multiscale Approximation on Graphics Processing Units," *Journal of Molecular Graphics and Modelling*, 28(8): 904-910, June 2010.

J34. A. Warren*, J. Archuleta*, W. Feng, and J. Setubal, "Missing Genes in the Annotation of Prokaryotic Genomes," *BMC Bioinformatics*, 11:131, 12 pages, March 2010.

J35. Z. Cao**, D. Easterling**, L. Watson, D. Li, K. Cameron, and W. Feng, "Power Saving Experiments for Large Scale Global Optimization," *International Journal of Parallel, Emergent and Distributed Systems*, 25(5): 381-400, February 2010.

J36. W. Feng, X. Feng, and R. Ge, "Green Supercomputing Comes of Age," *IT Professional*, 10(1): 17-23, January/February 2008.

J37. A. Banerjee**, W. Feng, D. Ghosal, and B. Mukherjee, "Algorithms for Integrated Routing and Scheduling for Aggregating Data from Distributed Resources on a Lambda Grid," *IEEE Transactions on Parallel & Distributed Systems*, 19(1): 24-34, January 2008.

J38. W. Feng and K. Cameron, "The Green500 List: Encouraging Sustainable Supercomputing," *IEEE Computer, Special Issue: Green Computing*, 40(12): 50-55, December 2007.

J39. P. Balaji**, W. Feng, and D. Panda, "Bridging the Ethernet-Ethernot Performance Gap," *IEEE Micro, Special Issue: High-Performance Interconnects*, 26(3): 24-40, May-June 2006. (Acceptance Rate: 10%)

J40. J. Hurwitz* and W. Feng, "Analyzing MPI Performance over 10-Gigabit Ethernet," *Journal of Parallel and Distributed Computing, Special Issue: Design and Performance of Networks for Super-, Cluster-, and Grid-Computing*, 65(10): 1253-1260, October 2005.

J41. X. Zhang**, L. Bhuyan, and W. Feng, "Anatomy of UDP and M-VIA for Cluster Communication," *Journal of Parallel and Distributed Computing, Special Issue: Design and Performance of Networks for Super-, Cluster-, and Grid-Computing*, 65(10): 1290-1298, October 2005.

J42. W. Feng, "The Importance of Being Low Power in High-Performance Computing," *Cyberinfrastructure Technology Watch*, National Science Foundation, 1(3): 12-21, August 2005.

J43. C. Jin, D. Wei, S. Low, G. Buhrmaster, J. Bunn, D. Choe, R. L. Cottrell, J. C. Doyle, W. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, and S. Singh, "FAST TCP: From Theory to Experiments," *IEEE Network*, 19(1): 4-11, January-February 2005.

J44. M. Gardner, S. Thulasidasan*, and W. Feng, "User-Space Auto-Tuning for TCP Flow Control in Computational Grids," *Computer Communications, Special Issue on Network Support for Grid Computing*, 27(14): 1364-1374, September 2004.

J45. M. Veeraraghavan, X. Zheng**, W. Feng, H. Lee, E. Chong, and H. Li, "Scheduling and Transport for File Transfers on High-Speed Optical Circuits," *Journal of Grid Computing*, 1(4): 395-405, June 2004.

J46. L. Carey*, A. Darling*, and W. Feng, "Takeoff with mpiBLAST: Parallel Searching," *ClusterWorld*, 2(3), March 2004.

J47. M. Baker and W. Feng, "10-Gigabit Ethernet Helps Relieve Network Bottlenecks for Bandwidth-Intensive Applications," *Dell Power Solutions*, pp. 113-116, March 2004.

J48. J. Hurwitz* and W. Feng, "End-to-End Performance of 10-Gigabit Ethernet on Commodity Systems," *IEEE Micro*, 24(1): 10-22, January-February 2004. (A preliminary version of this paper appeared in *IEEE Hot Interconnects: A Symposium on High-Performance Interconnects* in August 2003 and was selected as one of the best papers, and hence, invited for publication in a special issue of *IEEE Micro*.)

J49. W. Feng, "Making a Case for Efficient Supercomputing," *ACM Queue*, 1(7): 54-64, October 2003.

J50. W. Feng, M. Gardner, M. Fisk*, and E. Weigle*, "Automatic Flow-Control Adaptation for Enhancing Network Performance in Computational Grids," *Journal of Grid Computing*, 1(1): 63-74, 2003.

J51. W. Feng, M. Fisk*, M. Gardner, and E. Weigle*, "Dynamic Right-Sizing: An Automated, Lightweight, and Scalable Technique for Enhancing Grid Performance," *Lecture Notes in Computer Science*, 2334: 69-83, 2002. (A preliminary version of this paper appeared in the *7th IFIP/IEEE Workshop on Protocols for High-Speed Networks*.)

J52. W. Feng, M. Gardner, and J. Hay*, "The MAGNeT Toolkit: Design, Evaluation, and Implementation," *Journal of Supercomputing*, 23(1): 67-79, August 2002.

J53. A. Feng, A. Kapadia*, W. Feng, and G. Belford, "Packet Spacing: An Enabling Mechanism for the Delivery of Multimedia Content," *Journal of Supercomputing*, 23(1): 51-66, August 2002.

J54. W. Feng, "Research & Development in Advanced Network Technology (RADIANT)," *In Brief* Contribution, *D-Lib Magazine* (http://www.dlib.org), ISSN 1082-9873, Vol. 8, No. 3, March 2002.

J55. F. Petrini*, W. Feng, A. Hoisie, S. Coll, and E. Frachtenberg, "The Quadrics Network (QsNet): High-Performance Clustering Technology" (Extended Version), *IEEE Micro*, 22(1): 46-57, January-February 2002. (A preliminary version of this paper appeared in *IEEE Hot Interconnects: A Symposium on High-Performance Interconnects* in August 2001 and was selected as one of the best papers, and hence, invited for publication in a special issue of *IEEE Micro*.)

J56. F. Petrini* and W. Feng, "Improved Resource Utilization with Buffered Coscheduling," *Journal of Parallel Algorithms & Applications* (Special Issue), 16(2-3): 123-144, 2001.

J57. F. Petrini* and W. Feng, "Time-Sharing Parallel Jobs in the Presence of Multiple Resource Requirements," *Lecture Notes in Computer Science*, 1911: 113-136, 2000. (A preliminary version of this paper appeared in the *Workshop on Job Scheduling Strategies for Parallel Processing*, 2000.)

J58. W. Feng and J. W.-S. Liu, "Algorithms for Scheduling Real-Time Tasks with Input Error and End-to-End Deadlines," *IEEE Transactions on Software Engineering*, 23(2): 93-106, February 1997.

J59. A. Chien, W. Feng, V. Karamcheti, and J. Plevyak, "Techniques for Efficient Execution of Concurrent Object-Oriented Programs," *Lecture Notes in Computer Science*, Vol. 757, 1993. (A preliminary version of this paper appeared in the *Workshop on Languages and Compilers for Parallel Computing*, 1993.)

J60. R. Kasturi, R. Fernandez, M. Amlani, and W. Feng, "Map Data Processing in Geographic Information Systems," *IEEE Computer*, December 1989.

Refereed Articles in Conferences (unless otherwise noted)

Articles listed below were refereed by three or more reviewers. When available, acceptance rates are indicated in parentheses.

C1. S. Dash**, Q. Al-Hajri**, W. Feng, H. Garner, and R. Anandakrishnan, "Scaling Out a Combinatorial Algorithm for Discovering Carcinogenic Gene Combinations to Thousands of GPUs," *35th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Portland, Oregon, USA, May 2021.

C2. P. Sathre, A. Gondhalekar*, M. Hassan*, and W. Feng, "MetaCL: Automated 'Meta' OpenCL Code Generation for High-Level Synthesis on FPGA," *24th IEEE High-Performance Extreme Computing (HPEC) Conference*, Waltham, MA, September 2020.

C3. W. Feng, D. Zhang*, J. Zhang*, K. Hou*, S. Pumma*, and H. Wang*, "A Feasibility Study for MPI over HDFS," *24th IEEE High-Performance Extreme Computing (HPEC) Conference,* Waltham, Massachusetts, September 2020.

C4. A. Gondhalekar* and W. Feng, "Exploring FPGA Optimizations in OpenCL for Breadth-First Search on Sparse Graph Datasets," *30th International Conference on Field-Programmable Logic and Applications (FPL)*, Gothenburg, Sweden, pp. 133-137, August-September 2020.

C5. K. Youssef* and W. Feng, "SparkLeBLAST: Scalable Parallelization of BLAST Sequence Alignment Using Spark," *20th IEEE/ACM International Symposium on Cluster, Cloud, and Internet Computing (CCGrid 2020)*, Melbourne, Australia, May 2020.

C6. S. Pumma*, D. Buono, F. Checconi, X. Que, and W. Feng, "Alleviating Load Imbalance in Data Processing for Large-Scale Deep Learning," *20th IEEE/ACM International Symposium on Cluster, Cloud, and Internet Computing (CCGrid 2020)*, Melbourne, Australia, May 2020.

C7. G. Abram, V. Adhinarayanan*, W. Feng, D. Rogers, and J. Ahrens, "ETH: An Architecture for Exploring the Design Space of In-Situ Scientific Visualization," *34th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, New Orleans, LA, USA, May 2020.

C8. M. Hassan**, S. Pakin, and W. Feng, "C to D-Wave: A High-level C Compilation Framework for Quantum Annealers," *IEEE High-Performance Extreme Computing (HPEC) Conference*, Waltham, Massachusetts, USA, September 2019. **Innovative Paper Award**

C9. F. Wanye*, V. Gleyzer, and W. Feng, "Fast Stochastic Block Partitioning via Sampling," *IEEE High-Performance Extreme Computing (HPEC) Conference*, Waltham, Massachusetts, USA, September 2019. **Student Innovation Award**

C10. A. Helal**, A. Aji, M. Chu, B. Beckman, and W. Feng, "Adaptive Task Aggregation for High-Performance Sparse Solvers on GPUs," *28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Seattle, Washington, USA, September 2019. (Acceptance Rate: 27% = 34/126)

C11. X. Cui* and W. Feng, "Iterative Machine Learning (IterML) for Effective Parameter Pruning and Tuning in Accelerators," *16th International Conference on Computing Frontiers,* Alghero, Sardinia, Italy, April-May 2019. (Acceptance Rate: 19/79 = 24.1%)

C12. P. Sathre*, M. Gardner, and W. Feng, "On the Portability of GPU-Accelerated Applications via Automated Source-to-Source Translation," *HPC Asia,* Guangzhou, China, January 2019.

C13. P. Sathre*, A. Helal*, and W. Feng, "A Composable Workflow for Productive Heterogeneous Computing on FPGAs via Whole-Program Analysis and Transformation," *International Conference on Reconfigurable Computing and FPGAs (ReConFig)*, Cancun, Mexico, December 2018. (Acceptance Rate: 27/68 = 39.7%)

C14. M. Hassan**, A. Helal*, P. Athanas, W. Feng and Y. Hanafy, "Exploring FPGA-Specific Optimizations for Irregular OpenCL Applications, *International Conference on Reconfigurable Computing and FPGAs (ReConFig)*, Cancun, Mexico, December 2018. (Acceptance Rate: 27/68 = 39.7%)

C15. K. Krommydas*, P. Sathre*, R. Sasanka, and W. Feng, "A Framework for Auto-Parallelization and Code Generation: An Integrative Case Study with Legacy FORTRAN Codes," *47th International Conference on Parallel Processing (ICPP)*, Eugene, Oregon, USA, August 2018. (Acceptance Rate: 88/313 = 23.2%)

C16. A. Helal*, C. Jung, W. Feng, Y. Hanafy, "CommAnalyzer: Automated Estimation of Communication Cost and Scalability on HPC Clusters from Sequential Code," *27th International Symposium on High-Performance Parallel and Distributed Computing (ACM HPDC 2018)*, Tempe, Arizona, USA, June 2018. (Acceptance Rate: 22/112 = 19.6%)

C17. J. Zhang*, A. Aji, M. Chu, H. Wang*, and W. Feng, "Taming Irregular Applications via Advanced Dynamic Parallelism on GPUs," *15th ACM International Conference on Computing Frontiers*, Ischia, Italy, May 2018. (Acceptance Rate: 23/68 = 33.8%)

C18. B. Dutta*, V. Adhinarayanan*, and W. Feng, "GPU Power Prediction via Ensemble Machine Learning for DVFS Space Exploration," *15th ACM International Conference on Computing Frontiers*, Ischia, Italy, May 2018. (Acceptance Rate: 23/68 = 33.8%)

C19. K. Hou*, H. Wang*, W. Feng, J. Vetter, and S. Lee, "Highly Efficient Compensation-based Parallelism for Wavefront Loops on GPUs," *32nd IEEE International Parallel and Distributed Processing Symposium*, Vancouver, British Columbia, May 2018. (Acceptance Rate: 113/461 = 24.5%)

C20. X. Yu*, K. Hou*, H. Wang*, and W. Feng, "Robotomata: A Framework for Approximate Pattern Matching of Big Data on an Automata Processor," *IEEE International Conference on Big Data*, Boston, Massachusetts, USA, December 2017. (Acceptance Rate: 79/437 = 18.1%)

C21. S. Dash*, A. Verma*, C. North, and W. Feng, "Portable Parallel Design of Weighted Multi-Dimensional Scaling for Real-Time Data Analysis," *IEEE International Conference on High Performance Computing and Communications (HPCC)*, Bangkok, Thailand, December 2017. **Best Paper Finalist**

C22. S. Pumma*, M. Si, W. Feng, and P. Balaji. "Towards Scalable Deep Learning via I/O Analysis and Optimization," *IEEE International Conference on High Performance Computing and Communications (HPCC)*, Bangkok, Thailand, December 2017.

C23. S. Pumma*, M. Si, W. Feng, and P. Balaji. "Parallel I/O Optimizations for Scalable Deep Learning," *IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Shenzhen, China, December 2017. (Acceptance Rate: 89/271 = 32.8%)

C24. M. Nourian**, X. Wang**, X. Yu*, W. Feng, and M. Becchi. "Demystifying Automata Processing: GPUs, FPGAs or Micron's AP?. *31st ACM International Conference on Supercomputing*, Chicago, Illinois, USA, June 2017. (Acceptance Rate: 28/177 = 15.8%)

C25. K. Hou*, W. Liu, H. Wang*, and W. Feng, "Fast Segmented Sort on GPUs," *31st ACM International Conference on Supercomputing*, Chicago, Illinois, USA, June 2017. (Acceptance Rate: 28/177 = 15.8%)

C26. A. Verma**, H. Zhou, S. Booth, R. King, J. Coole, A. Keep, J. Marshall, and W. Feng, "Developing Dynamic Profiling and Debugging Support in OpenCL for FPGAs," *54th Design Automation Conference (DAC)*, Austin, Texas, USA, June 2017. (Acceptance rate: 160/676 = 23.6%)

C27. V. Adhinarayanan*, W. Feng, D. Rogers, J. Ahrens, and S. Pakin, "Characterizing and Modeling Power and Energy for Extreme-Scale In-situ Visualization," *31st IEEE International Parallel and Distributed Processing Symposium*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C28. J. Zhang*, S. Misra, H. Wang*, and W. Feng, "Eliminating Irregularities of Protein Sequence Search on Multicore Architectures," *31st IEEE International Parallel and Distributed Processing Symposium*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C29. H. Wang*, J. Zhang*, D. Zhang*, S. Pumma*, and W. Feng, "PaPar: A Parallel Data Partitioning Framework for Big Data Applications," *31st IEEE International Parallel and Distributed Processing Symposium*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C30. V. Adhinarayanan*, W. Feng, D. Rogers, J. Ahrens, and S. Pakin, "Characterizing and Modeling Energy for Extreme-Scale In-Situ Visualization," *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C31. X. Cui*, T. Scogland, B. de Supinski, and W. Feng, "Directive-Based Partitioning and Pipelining for Graphics Processing Units," *31st IEEE International Parallel and Distributed Processing Symposium*, Orlando, Florida, USA, May 2017. (Acceptance Rate: 116/498 = 23.2%)

C32. K. Hou*, H. Wang*, and W. Feng, "GPU-UniCache: Automatic Code Generation of Spatial Blocking for Stencils on GPUs," *14th ACM International Conference on Computing Frontiers*, Siena, Italy, May 2017.

C33. X. Yu*, H. Wang*, W. Feng, H. Gong*, and G. Cao, "An Enhanced Image Reconstruction Tool for Computed Tomography on GPUs," *14th ACM International Conference on Computing Frontiers*, Siena, Italy, May 2017.

C34. U. Kalim*, M. Gardner, E. Brown, and W. Feng, "SLIM: Enabling Transparent Extensibility and Dynamic Configuration via Session-Layer Abstractions," *13th ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS '17)*, Beijing, China, May 2017. (Acceptance Rate: 15/55 = 27.3%)

C35. X. Cui**, T. Scogland, B. de Supinski, and W. Feng, "Directive-Based Partitioning and Pipelining for Graphics Processing Units," *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C36. J. Zhang**, S. Misra, H. Wang*, and W. Feng, "Eliminating Irregularities of Protein Sequence Search on Multicore Architectures," *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C37. H. Wang*, J. Zhang**, D. Zhang**, S. Pumma**, and W. Feng "PaPar: A Parallel Data Partitioning Framework for Big Data Applications," *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Orlando, Florida, USA, May-June 2017. (Acceptance Rate: 116/498 = 23.2%)

C38. S. Banerjee**, X. Chen, X. Wu, H. Xie, J. Xuan, and W. Feng, "ChIP-GMM: A Gaussian Mixture Model for Inferring Binding Regions in ChIP-seq Profiles," *9th International Conference on Bioinformatics and Computational Biology (BiCoB)*, Honolulu, Hawaii, USA, March 2017.

C39. K. Krommydas* and W. Feng, "Telescoping Architectures: Evaluating Next-Generation Heterogeneous Computing," *23rd International Conference on High-Performance Computing, Data, and Analytics (HiPC'16)*, Hyderabad, India, December 2016. (Acceptance Rate: 40/160 = 25%)

C40. A. Helal*, P. Sathre*, and W. Feng, "MetaMorph: A Library Framework for Interoperable Kernels on Multi- and Many-core Clusters," *SC|16: International Conference for High Performance Computing, Networking, Storage and Analysis*, Salt Lake City, Utah, USA, November 2016. (Acceptance Rate: 82/446 = 18.4%)

C41. V. Adhinarayanan**, I. Paul, J. Greathouse, W. Huang, A. Pattnaik, and W. Feng, "Measuring and Modeling On-Chip Interconnect Power on Real Hardware," *IEEE International Symposium on Workload Characterization (IISWC)*, Providence, Rhode Island, USA, September 2016. **Best Paper Award.** (Acceptance Rate: 21/69 = 30.4%)

C42. I. Harb* and W. Feng, "Characterizing Performance and Power Towards Efficient Synchronization of GPU Kernels" (Short Paper), *24th IEEE International Symposium on Modeling, Analysis, and Simulation on Computer and Telecommunication Systems (MASCOTS)*, London, England, September 2016.

C43. K. Krommydas*, R. Sasanka, and W. Feng, "Bridging the FPGA Programmability-Portability Gap via Automatic OpenCL Code Generation and Tuning," *27th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, London, England. July 2016. (Acceptance Rate: 18/85 = 21.2%)

C44. H. Wang*, W. Liu, K. Hou*, and W. Feng, "Parallel Transposition of Sparse Data Structures," *30th ACM International Conference on Supercomputing*, Istanbul, Turkey, June 2016. (Acceptance Rate: 43/178 = 24.2%)

C45. V. Adhinarayanan*, B. Subramaniam*, and W. Feng, "Online Power Estimation of Graphics Processing Units," *16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Cartagena, Colombia May 2016. (Acceptance Rate: 40/200 = 20%)

C46. X. Yu*, H. Wang*, W. Feng, H. Gong, and G. Cao, "cuART: Fine-Grained Algebraic Reconstruction Technique for Computed Tomography Images on GPUs," *16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Cartagena, Colombia May 2016. (Acceptance Rate: 40/200 = 20%)

C47. V. Adhinarayanan* and W. Feng, "An Automated Framework for Characterizing and Subsetting GPGPU Workloads," *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Uppsala, Sweden, April 2016. (Acceptance Rate: 27/77 = 35.1%)

C48. X. Yu**, W. Feng, D. Yao, and M. Becchi, "O3FA: A Scalable, Finite Automata-based, Pattern-Matching Engine for Out-of-Order Packet Inspection in IDS, *12th ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)* – co-located with USENIX NSDI, Santa Clara, CA, USA, March 2016. (Acceptance Rate: 12/58 = 20.6%)

C49. K. Hou*, H. Wang*, and W. Feng, "AAlign: A SIMD Framework for Pairwise Sequence Alignment on x86-based Multi- and Many-core Processors," *30th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Chicago, Illinois, USA, May 2016. (Acceptance Rate: 114/496 = 23%)

C50. D. Zhang*, H. Wang*, K. Hou*, J. Zhang*, and W. Feng, "pDindel: Accelerating Indel Detection on a Multicore CPU Architecture with SIMD," *5th IEEE International Conference on Computational Advanced in Bio and Medical Sciences (ICCABS)*, Miami, Florida, USA, October 2015.

C51. A. Aji**, A. Pena, P. Balaji, and W. Feng, "Automatic Command Queue Scheduling for Task-Parallel Workloads in OpenCL," *IEEE International Conference on Cluster Computing,* Chicago, Illinois, USA, September 2015. (Acceptance Rate: 24% = 38/157)

C52. K. Krommydas*, R. Sasanka, and W. Feng, "GLAF: A Visual Programming and Auto-Tuning Framework for Parallel Computing," *44th International Conference on Parallel Processing*, Beijing, China, September 2015. (Acceptance Rate: 32% = 99/305)

C53. K. Hou*, H. Wang*, and W. Feng, "ASPaS: A Framework for Automatic SIMDization of Parallel Sorting on x86-based Many-core Processors," *25th International Conference on Supercomputing (ICS)*, Newport Beach, California, USA, June 2015. (Acceptance Rate: 25% = 40/160)

C54. L. Luo, J. Edwards, H. Luo, F. Mueller, and W. Feng, "Optimization of a Fine-grained BILU by CUDA Inter-block Synchronization," *22nd AIAA Computational Fluid Dynamics Conference*, Dallas, Texas, USA, June 2015.

C55. T. Scogland* and W. Feng, "Design and Evaluation of Scalable Concurrent Queues for Many-Core Architectures," *ACM/SPEC International Conference on Performance Engineering*, Austin, Texas, USA, January 2015. (Acceptance Rate: 32% = 18/56)

C56. K. Krommydas*, T. Koehn, K. Kepa, W. Feng, and P. Athanas, "On the Performance and Energy Efficiency of FPGAs and GPUs for Polyphase Channelization," *International Conference on ReConFigurable Computing and FPGAs* (ReConFig), Cancun, Mexico, December 2014.

C57. N. Mohamed*, N. Maji*, J. Zhang*, N. Timoshevskaya*, and W. Feng, "Aeromancer: A Workflow Manager for Large-Scale MapReduce-Based Scientific Workflows," *3rd IEEE International Conference on Big Data Science and Engineering (BDSE)*, Beijing, China, September 2014.

C58. T. Scogland* and W. Feng, "Locality-Aware Memory Association for Multi-Target Worksharing in OpenMP" (Short Paper), *23rd International Conference on Parallel Architectures and Compilation*, Edmonton, Alberta, August 2014.

C59. J. Yin**, J. Wang, W. Feng, X. Zhang*, and J. Zhang*, "SLAM: Scalable Locality-Aware Middleware for I/O in Scientific Analysis and Visualization," *23rd International Symposium on High-Performance Parallel and Distributed Computing*, Vancouver, Canada, June 2014. (Acceptance Rate: 28% = 37/130)

C60. K. Krommydas*, W. Feng, M. Owaida, C. Antonopoulos, and N. Bellas, "On the Portability of OpenCL Dwarfs on Fixed and Reconfigurable Parallel Platforms," *25th IEEE International Conference on Application-specific Systems, Architectures and Processors*, Zurich, Switzerland, June 2014. **Best Paper Finalist.** (Acceptance Rate: 26% = 22/85)

C61. N. Timoshevskaya* and W. Feng, "SAIS-OPT: On the Characterization and Optimization of the SA-IS Algorithm for Suffix Array Construction," *4th IEEE International Conference on Computational Advanced in Bio and Medical Sciences (ICCABS)*, Miami Beach, Florida, USA, June 2014.

C62. A. Amritkar, D. Tafti, P. Sathre*, K. Hou*, S. Chivukula*, and W. Feng, "Accelerating Bio-Inspired MAV Computations using GPUs," *AIAA Aviation and Aeronautics Forum and Exposition 2014*, Atlanta, Georgia, USA, June 2014.

C63. C. del Mundo* and W. Feng, "Towards a Performance-Portable FFT Library for Heterogeneous Computing," *11th ACM International Conference on Computing Frontiers*, Cagliari, Italy, May 2014.

C64. B. Subramaniam* and W. Feng, "Enabling Efficient Power Provisioning for Enterprise Applications," *14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Chicago, Illinois, USA, May 2014. (Acceptance Rate: 19% = 54/283)

C65. J. Zhang*, H. Wang*, H. Lin, and W. Feng, "cuBLASTP: Fine-Grained Parallelization of Protein Sequence Search on a GPU," *28th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Phoenix, Arizona, USA, May 2014. (Acceptance Rate: 21% = 114/541)

C66. J. E. McClure, H. Wang*, J. F. Prins, C. T. Miller, and W. Feng, "Petascale Application of a Coupled CPU-GPU Algorithm for Simulation and Analysis of Multiphase Flow Solutions in Porous Medium Systems," *28th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Phoenix, Arizona, USA, May 2014. (Acceptance Rate: 21% = 114/541)

C67. T. Scogland*, C. Steffen, T. Wilde, F. Parent, S. Coghlan, N. Bates, W. Feng and E. Strohmaier, "A Power-Measurement Methodology for Large-Scale, High-Performance Computing, *5th ACM/SPEC International Conference on Performance Engineering (ICPE)*, Dublin, Ireland, March 2014. **Best Paper Finalist**

C68. B. Pickering**, C. Jackson**, T. Scogland*, W. Feng, C. Roy, "Directive-Based GPU Programming for Computational Fluid Dynamics," *AIAA Science and Technology Forum and Exposition (SciTech 2014)*, National Harbor, Maryland, USA, January 2014.

C69. L. Panwar*, A. Aji*, J. Meng, P. Balaji and W. Feng, "Online Performance Projection for Clusters with Heterogeneous GPUs," *19th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Seoul, Korea, December 2013. (Acceptance Rate: 30% = 52/170)

C70. V. Adhinarayanan* and W. Feng, "Wideband Channelization for Software-Defined Radio via Mobile Graphics Processors," *19th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Seoul, Korea, December 2013. (Acceptance Rate: 30% = 52/170)

C71. K. Krommydas*, T. Scogland*, and W. Feng, "On the Programmability and Performance of Heterogeneous Platforms," *19th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Seoul, Korea, December 2013. (Acceptance Rate: 30% = 52/170)

C72. J. Zhang*, H. Wang*, H. Lin*, and W. Feng, "Consolidating Applications for Energy Efficiency in Heterogeneous Computing Systems," *15th IEEE International Conference on High-Performance Computing and Communications (HPCC)*, Zhangjiajie, China, November 2013. (Acceptance Rate: < 30%)

C73. P. Lama, Y. Li, A. Aji*, P. Balaji, J. Dinan, S. Xiao*, Y. Zhang, W. Feng, R. Thakur, and X. Zhou, "pVOCL: Power-Aware Dynamic Placement and Migration in Virtualized GPU Environments," *33rd IEEE International Conference on Distributed Computing Systems*, Philadelphia, Pennsylvania, USA, July 2013. (Acceptance Rate: 13% = 61/464)

C74. B. Subramaniam*, W. Saunders, T. Scogland*, and W. Feng, "*Trends in Energy-Efficient Computing: A Perspective from the Green500*. In *Proceedings of the 4th International Green Computing Conference*, Arlington, Virginia, USA, June 2013.

C75. A. Aji*, L. Singh*, W. Feng, P. Balaji, J. Dinan, R. Thakur, F. Ji, X. Ma, M. Chabbi, K. Murthy, J. Mellor-Crummey, and K Bisset, "On the Efficacy of GPU-Integrated MPI for Scientific Applications," *22nd ACM International Symposium on High-Performance Distributed Computing*, New York City, New York, USA, June 2013. (Acceptance Rate: 15% = 20/131)

C76. C. del Mundo*, V. Adhinarayanan*, and W. Feng, "Accelerating Fast Fourier Transform for Wideband Channelization," *IEEE International Conference on Communications (ICC)*, Budapest, Hungary, June 2013. (Acceptance Rate: 39% = 948/2422)

C77. J. Zhang*, H. Lin*, and W. Feng, "Optimizing Burrows-Wheeler Transform-Based Sequence Alignment on Multicore Architectures," *13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Delft, Netherlands, May 2013. (Acceptance Rate: 22% = 57/257)

C78. B. Subramaniam* and W. Feng, "Towards Energy-Proportional Computing for Enterprise-Class Server Workloads," *4th ACM/SPEC International Conference on Performance Engineering*, Prague, Czech Republic, April 2013. **Best Paper Award.**

C79. N. Mohamed*, H. Lin*, and W. Feng, "Accelerating Data-Intensive Genome Analysis in the Cloud," *5th International Conference on Bioinformatics and Computational Biology (BiCoB)*, Honolulu, Hawaii, USA, March 2013.

C80. A. Aji*, D. Buntinas, J. Dinan, P. Balaji, W. Feng, K. Bisset, R. Thakur, "Efficient Inter-node Communication in High-Performance GPU Systems," *14th IEEE International Conference on High-Performance Computing and Communications (HPCC)*, Liverpool, United Kingdom, June 2012. (Acceptance Rate: < 30%)

C81. F. Ji**, A. Aji*, J. Dinan, D. Buntinas, P. Balaji, R. Thakur, W. Feng, X. Ma, "DMA-Assisted, Intranode Communication in GPU-Accelerated Systems," *14th IEEE International Conference on High-Performance Computing and Communications (HPCC)*, Liverpool, United Kingdom, June 2012. (Acceptance Rate: < 30%)

C82. K. Lee*, H. Lin*, and W. Feng, "Performance Characterization of Data-Intensive Kernels on AMD Fusion Architectures," *27th International Supercomputing Conference (ISC)*, Hamburg, Germany, June 2012. (Acceptance Rate: 46% = 16/35)

C83. B. Subramaniam* and W. Feng, "GBench: Benchmarking Methodology for Evaluating the Energy Efficiency of Supercomputers," *27th International Supercomputing Conference (ISC)*, Hamburg, Germany, June 2012. (Acceptance Rate: 46% = 16/35)

C84. T. Scogland*, B. Subramaniam*, and W. Feng, "The Green500 List: Escapades to Exascale," *27th International Supercomputing Conference (ISC)*, Hamburg, Germany, June 2012. (Acceptance Rate: 46% = 16/35)

C85. T. Scogland*, B. Rountree, W. Feng, and B. de Supinski, "Heterogeneous Task Scheduling for Accelerated OpenMP," *26th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Shanghai, China, May 2012. (Acceptance Rate: 21% = 118/569)

C86. S. Xiao*, P. Balaji, J. Dinan, Q. Zhu, R. Thakur, S. Coghlan, H. Lin*, G. Wen, J. Hong, and W. Feng, "Transparent Accelerator Migration in a Virtualized GPU Environment," *12th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid)*, Ottawa, Canada, May 2012. (Acceptance Rate: 27% = 83/302)

C87. S. Xiao*, P. Balaji, Q. Zhu, R. Thakur, S. Coghlan, H. Lin*, G. Wen, J. Hong, and W. Feng, "VOCL: An Optimized Environment for Transparent Virtualization of Graphics Processing Units," *IEEE Innovative Parallel Computing*, San Jose, California, USA, May 2012. (Acceptance Rate: 40% = 25/62)

C88. R. Braithwaite*, W. Feng, P. McCormick, "Automatic NUMA Characterization using Cbench," *3rd ACM/SPEC International Conference on Performance Engineering*, Boston, Massachusetts, USA, April 2012.

C89. W. Feng, H. Lin*, T. Scogland*, and J. Zhang*, "OpenCL and the 13 Dwarfs: A Work in Progress," *3rd ACM/SPEC International Conference on Performance Engineering*, Boston, Massachusetts, USA, April 2012.

C90. M. Daga*, T. Scogland*, and W. Feng, "Architecture-Aware Mapping and Optimization on a 1600-Core GPU," *17th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Tainan, Taiwan, December 2011. (Acceptance Rate: 27% = 96/352)

C91. M. Elteir*, H. Lin*, and W. Feng, "StreamMR: An Optimized MapReduce Framework for AMD GPUs," *17th IEEE International Conference on Parallel and Distributed Systems (ICPADS),* Tainan, Taiwan, December 2011. (Acceptance Rate: 27% = 96/352)

C92. G. Martinez*, M. Gardner, and W. Feng, "CU2CL: A CUDA-to-OpenCL Translator for Multi- and Many-Core Architectures," *17th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Tainan, Taiwan, December 2011. (Acceptance Rate: 27% = 96/352)

C93. C. Wu, J. Ke, H. Lin*, and W. Feng, "Optimizing Dynamic Programming on Graphics Processing Units via Adaptive Thread-Level Parallelism," *17th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, Tainan, Taiwan, December 2011. (Acceptance Rate: 27% = 96/352)

C94. K. Pereira**, P. Athanas, H. Lin*, and W. Feng, "Spectral Method Characterization on FPGA and GPU Accelerators," *International Conference on Reconfigurable Computing and FPGAs (ReConFig)*, Cancun, Mexico, November-December 2011.

C95. M. Elteir*, H. Lin*, and W. Feng, "Performance Characterization and Optimization of Atomic Operations on AMD GPUs," *IEEE Cluster 2011,* Austin, Texas, USA, September 2011. (Acceptance Rate: 27% = 38/140)

C96. E. Brown, M. Gardner, U. Kalim*, and W. Feng, "Restoring End-to-End Resilience in the Presence of Middleboxes," *IEEE International Conference on Computer Communications and Networking*, Maui, Hawaii, USA, July-August 2011. (Acceptance Rate: 29% = 144/490)

C97. M. Daga*, A. Aji*, and W. Feng, "On the Efficacy of a Fused CPU+GPU Processor for Parallel Computing," *Symposium on Application Accelerators in High-Performance Computing (SAAHPC),* Knoxville, Tennessee, USA, July 2011.

C98. K. Krommydas*, C. Antonopoulos, N. Bellas, and W. Feng, "AVS Video Decoder on Multicore Systems: Optimizations and Tradeoffs" (Short Paper), *IEEE International Conference on Multimedia and Expo*, Barcelona, Spain, July 2011.

C99. S. Xiao*, H. Lin*, and W. Feng, "Accelerating Protein Sequence Search in a Heterogeneous Computing System," *25th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Anchorage, Alaska, USA, May 2011. (Acceptance Rate: 20% = 112/571)

C100. A. Aji*, M. Daga*, and W. Feng, "Bounding the Effect of Partition Camping in GPU Kernels" *8th ACM International Conference on Computing Frontiers*, Ischia, Italy, May 2011. (Acceptance Rate: 22% = 22/101)

C101. M. Daga*, T. Scogland*, and W. Feng, "Towards Accelerating Molecular Modeling via Multi-Scale Approximation on a GPU," *1st IEEE International Conference on Computational Advances in Bio and Medical Sciences (ICCABS)*, Orlando, Florida, USA, February 2011. (Acceptance Rate: 43% = 38/89)

C102. K. Bisset, A. Aji*, M. Marathe, and W. Feng, "High-Performance Biocomputing for Simulating the Spread of Contagion over Large Contact Networks," *1st IEEE International Conference on Computational Advances in Bio and Medical Sciences (ICCABS)*, Orlando, Florida, USA, February 2011. (Acceptance Rate: 38/89 = 42.7%)

C103. A. Aji*, L. Zhang, and W. Feng, "GPU-RMAP: Accelerating Short-Read Mapping on Graphics Processors," *13th IEEE International Conference on Computational Science and Engineering*, Hong Kong, China, December 2010. (Acceptance Rate: 17% = 37/213)

C104. W. Feng and B. Subramaniam*, "Understanding Power Measurement Implications in the Green500 List," *IEEE International Conference on Green Computing and Communications*, Hangzhou, China, December 2010. (Acceptance Rate: 18% = 41/233)

C105. Y. Jiao*, H. Lin*, and W. Feng, "Power and Performance Characterization of Computational Kernels on the GPU," *IEEE International Conference on Green Computing and Communications*, Hangzhou, China, December 2010. (Acceptance Rate: 41/233 = 17.6%)

C106. B. Subramaniam* and W. Feng, "Statistical Power and Performance Modeling for Optimizing the Energy Efficiency of Scientific Computing," *IEEE International Conference on Green Computing and Communications*, Hangzhou, China, December 2010. (Acceptance Rate: 41/233 = 17.6%)

C107. M. Elteir*, H. Lin*, and W. Feng, "Enhancing MapReduce via Asynchronous Data Processing," *16th International Conference on Parallel and Distributed Systems (ICPADS)*, Shanghai, China, December 2010. (Acceptance Rate: 77/188 = 29.6%)

C108. F. Ge**, H. Lin*, A. Khajeh, J. Chiang, A. Eltawil, C. Bostian, W. Feng, and R. Chadha, "Cognitive Radio Rides on the Cloud," *IEEE Military Communications Conference (MILCOM)*, San Jose, California, USA, October-November 2010.

C109. T. Scogland*, H. Lin*, and W. Feng, "A First Look at Integrated GPUs for Green High-Performance Computing," *1st International Conference on Energy-Aware High-Performance Computing (EnA-HPC)*, Hamburg, Germany, September 2010.

C110. S. Sharma*, D. Gillies, and W. Feng, "On the Goodput of TCP in Mobile Networks," *19th International Conference on Computer Communications and Networks (ICCCN)*, Zurich, Switzerland, August 2010.

C111. M. Gardner and W. Feng, "Broadening Accessibility to Computer Science for K-12 Education," *15th Conference on Innovation and Technology in Computer Science Education*, Ankara, Turkey, June 2010. (European version of ACM SIGCSE.)

C112. H. Lin**, X. Ma, J. Archuleta*, W. Feng, M. Gardner, and Z. Zhang, "MOON: MapReduce On Opportunistic eNvironments," *19th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC'10)*, Chicago, Illinois, USA, June 2010. (Acceptance Rate: 25% = 23/91) **Addendum: Selected Amongst "Best 20 Papers over the Past 20 Years" in HPDC, April 2012.**

C113. W. Feng and S. Xiao*, "To GPU Synchronize or Not GPU Synchronize?" *IEEE International Symposium on Circuits and Systems*, Paris, France, May 2010.

C114. Y. Cao, D. Patnaik**, S. Ponce**, J. Archuleta*, P. Butler**, W. Feng, and N. Ramakrishnan, "Towards Chip-on-Chip Neuroscience: Fast Mining of Neuronal Spike Streams Using Graphics Hardware," *7th ACM International Conference on Computing Frontiers*, Bertinoro, Italy, May 2010. (Acceptance Rate: 27% = 30/113)

C115. S. Xiao* and W. Feng, "Inter-Block GPU Communication via Fast Barrier Synchronization," *24th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Atlanta, Georgia, USA, April 2010. (Acceptance Rate: 24% = 127/527)

C116. S. Xiao*, A. Aji*, and W. Feng, "On the Robust Mapping of Dynamic Programming onto a Graphics Processing Unit," *15th International Conference on Parallel and Distributed Systems (ICPADS)*, Shenzhen, China, December 2009. **Best Paper Nominee.**

C117. A. Singh*, P. Balaji, and W. Feng, "GePSeA: A General-Purpose Software Acceleration Framework for Lightweight Task Offloading," *38th International Conference on Parallel Processing*, Vienna, Austria, September 2009. (Acceptance Rate: 32% = 71/220)

C118. J. Archuleta*, Y. Cao, T. Scogland*, and W. Feng, "Multi-Dimensional Characterization of Temporal Data Mining on Graphics Processors," *23rd IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Rome, Italy, May 2009. (Acceptance Rate: 23% = 100/440)

C119. S. Huang* and W. Feng, "Energy-Efficient Cluster Computing via Accurate Workload Characterization," *7th IEEE International Symposium on Cluster Computing and the Grid (CCGrid)*, Shanghai, China, May 2009. (Acceptance Rate: 21% = 57/271)

C120. M. Kumar*, V. Chaube*, P. Balaji, W. Feng, and H. Jin, "Making a Case for Proactive Flow Control in Optical Circuit-Switched Networks," *15th International Conference on High-Performance Computing (HiPC)*, Bangalore, India, December 2008. (Acceptance Rate: 14% = 46/319)

C121. H. Lin**, P. Balaji, R. Poole, C. Sosa, X. Ma, and W. Feng, "Massively Parallel Genomic Sequence Search on the Blue Gene/P Architecture," *ACM/IEEE SC|08: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Austin, Texas, USA, November 2008. (Acceptance Rate: 21% = 59/277)

C122. T. Scogland*, P. Balaji, W. Feng, and G. Narayanaswamy*, "Asymmetric Interactions in Symmetric Multi-core Systems: Analysis, Enhancements, and Evaluation," *ACM/IEEE SC|08: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Austin, Texas, USA, November 2008. (Acceptance Rate: 21% = 59/277)

C123. A. Aji* and W. Feng, "Optimizing Performance, Cost, and Sensitivity in Pairwise Sequence Search on a Cluster of PlayStations," *IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*, Athens, Greece, October 2008.

C124. G. Narayanaswamy*, P. Balaji, W. Feng, "Impact of Network Sharing in Multi-core Architectures," *17th International Conference on Computer Communications and Networks (ICCCN)*, St. Thomas, U.S. Virgin Islands, August 2008. (Acceptance Rate: 26% = 104/399)

C125. P. Balaji, W. Feng, and H. Lin*, "Semantic-Based Distributed I/O with the ParaMEDIC Framework," *17th ACM/IEEE International Symposium on High-Performance Distributed Computing (HPDC'08)*, Boston, Massachusetts, USA, June 2008. (Acceptance Rate: 17% = 18/103)

C126. P. Balaji, W. Feng, H. Lin*, J. Archuleta*, S. Matsuoka, A. Warren, J. Setubal, E. Lusk, R. Thakur, I. Foster, D. Katz, S. Jha, K. Shinpaugh, S. Coghlan, and D. Reed, "Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer," *23rd International Supercomputing Conference (ISC '08)*, Dresden, Germany, June 2008. **Distinguished Paper Award.**

C127. M. Gardner and W. Feng, "Towards a Virtual Ecosystem for K-8 Education," *2nd International Conference on the Virtual Computing Initiative (ICVCI'08)*, Research Triangle Park, North Carolina, USA, May 2008.

C128. A. Aji*, W. Feng, F. Blagojevic, and D. Nikolopoulos, "Cell-SWat: Modeling and Scheduling Wavefront Computations on the Cell Broadband Engine," *5th ACM International Conference on Computing Frontiers*, Ischia, Italy, May 2008. (Acceptance Rate: 27% = 30/110)

C129. P. Balaji, W. Feng, J. Archuleta*, H. Lin**, R. Kettimuthu, R. Thakur, and X. Ma, "Semantics-Based Distributed I/O for mpiBLAST," *13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Salt Lake City, Utah, USA, February 2008. (Acceptance Rate: 34% = 22/65)

C130. M. Arisoylu* and W. Feng, "Achieving Edge-Based Fairness in a Multi-Hop Environment," *5th IEEE Consumer Communications and Networking Conference*, Las Vegas, Nevada, USA, January 2008.

C131. P. Balaji, W. Feng, S. Bhagvat, D. Panda, R. Thakur, and W. Gropp, "Analyzing the Impact of Supporting Out-of-Order Communication on In-Order Performance with iWARP," *ACM/IEEE SC|07:* The International Conference on High-Performance Computing, Networking, Storage, and Analysis, Reno, Nevada, USA, November 2007. (Acceptance Rate: 20% = 54/268)

C132. J. Archuleta**, E. Tilevich, and W. Feng, "A Maintainable Software Architecture for Fast and Modular Bioinformatics Sequence Search," *23rd IEEE International Conference on Software Maintenance*, Paris, France, October 2007. (Acceptance Rate: 21% = 46/214)

C133. R. Ge**, X. Feng**, W. Feng, and K. Cameron, "CPU MISER: A Performance-Directed, Run-Time System for Power-Aware Clusters," *36th International Conference on Parallel Processing*, Xian, China, September 2007.

C134. G. Narayanaswamy**, P. Balaji, and W. Feng, "An Analysis of 10-Gigabit Ethernet Protocol Stacks in Multicore Environments," *15th Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2007.

C135. J. Archuleta**, W. Feng, and E. Tilevich, "A Pluggable Framework for Parallel Pairwise Sequence Search," *29th International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Lyon, France, August 2007.

C136. P. Datta*, W. Feng, and S. Sushant*, "End-System Aware, Rate-Adaptive Protocol for Network Transport in LambdaGrid Environments," *ACM/IEEE SC|06: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Tampa, Florida, USA, November 2006. (Acceptance Rate: 23% = 54/239)

C137. M. Gardner, W. Feng, J. Archuleta*, H. Lin**, and X. Ma, "Parallel Genomic Sequence-Searching on an Ad-Hoc Grid: Experiences, Lessons Learned, and Implications," *ACM/IEEE SC|06: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Tampa, Florida, USA, November 2006. (Acceptance Rate: 23% = 54/239) **Best Paper Finalist**

C138. A. Ching**, W. Feng, H. Lin**, X. Ma, and A. Choudhary, "Exploring I/O Strategies for Parallel Sequence Database Search Tools with S3aSim," *15th IEEE International Symposium on High-Performance Distributed Computing (HPDC'06)*, Paris, France, June 2006. (Acceptance Rate: 15% = 24/157)

C139. P. Datta*, S. Sharma*, and W. Feng, "A Feedback Mechanism for Network Scheduling in LambdaGrids," *6th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid'06)*, Singapore, May 2006. (Acceptance Rate: 24% = 61/257)

C140. A. Banerjee**, W. Feng, B. Mukherjee, and D. Ghosal, "RAPID: An End-System Aware Protocol for Intelligent Data-Transfer over LambdaGrids," *20th International Parallel & Distributed Processing Symposium (IPDPS'06)*, Rhodes, Greece, April 2006. (Acceptance Rate: 24% = 125/531)

C141. H. Lin**, X. Ma, W. Feng, A. Geist, and N. Samatova, "Efficient Data Handling in Comparative Genome Analysis Applications," *12th SIAM Conference on Parallel Processing for Scientific Computing*, San Francisco, California, USA February 2006.

C142. J. Gans**, W. Feng, and M. Wolinsky, "Whole Genome, Physics-Based Sequence Alignment for Pathogen Signature Design," *12th SIAM Conference on Parallel Processing for Scientific Computing*, San Francisco, California, USA, February 2006.

C143. C. Hsu* and W. Feng, "A Power-Aware Run-Time System for High-Performance Computing," *ACM/IEEE SC2005: The International Conference on High-Performance Computing, Networking, and Storage*, Seattle, Washington, USA, November 2005. (Acceptance Rate: 24% = 63/260)

C144. C. Hsu* and W. Feng, "A Feasibility Analysis of Power Awareness in Commodity-Based High-Performance Clusters," *7th IEEE International Conference on Cluster Computing (IEEE Cluster)*, Boston, Massachusetts, USA, September 2005. (Acceptance Rate: 33% = 45/138)

C145. P. Balaji*, W. Feng, Q. Gao, R. Noronha, W. Yu, and D. Panda, "Head-to-TOE Evaluation of High-Performance Sockets over Protocol Offload Engines," 7th *International Conference on Cluster Computing (IEEE Cluster)*, Boston, Massachusetts, USA, September 2005. (Acceptance Rate: 33% = 45/138.)

C146. W. Feng, P. Balaji*, C. Baron*, L. Bhuyan, and D. Panda, "Performance Characterization of a 10-Gigabit Ethernet TOE," *13th IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2005.

C147. C. Hsu* and W. Feng, "Reducing Overheating-Induced Failures via Performance-Aware CPU Power Management," *The 6th International Conference on Linux Clusters: The HPC Revolution 2005,* Chapel Hill, North Carolina, USA, April 2005.

C148. S. Ayyorgun*, S. Vanichpun*, and W. Feng, "Q-Composer and CpR: A Probabilistic Synthesizer and Regulator of Traffic," *24th IEEE INFOCOM*, Miami, Florida, USA, March 2005. (Acceptance Rate: 17% = 244/1419)

C149. S. Ayyorgun* and W. Feng, "A Systematic Approach for Providing End-to-End Probabilistic QoS Guarantees," *13th IEEE International Conference on Computer Communications and Networks (IC3N'04)*, Chicago, Illinois, USA, October 2004. (Acceptance Rate: 35%)

C150. W. Feng, "A Multimodal Interface for the Immediate Transcription of Radiology Dictation," *17th IEEE Symposium on Computer-Based Medical Systems (CBMS'04)*, Bethesda, Maryland, USA, June 2004.

C151. W. Feng and C. Hsu*, "Green Destiny and Its Evolving Parts," *19th International Supercomputer Conference*, Heidelberg, Germany, June 2004. **Innovative Supercomputer Architecture Award**

C152. A. Engelhart*, M. Gardner, and W. Feng, "Re-Architecting Flow-Control Adaptation for Grid Environments," *18th IEEE International Parallel & Distributed Processing Symposium (IPDPS'04)*, Santa Fe, New Mexico, USA, April 2004. (Acceptance Rate: 32% = 142/447)

C153. W. Feng and C. Hsu*, "The Origin and Evolution of Green Destiny," *IEEE Cool Chips VII: An International Symposium on Low-Power and High-Speed Chips*, Yokohama, Japan, April 2004.

C154. S. Ayyorgun* and W. Feng, "A Deterministic Characterization of Network Traffic for Average Performance Guarantees," *38th Annual Conference on Information Sciences and Systems (CISS'04)*, Princeton, New Jersey, USA, March 2004.

C155. M. Gardner, W. Deng, T. S. Markham*, C. Mendes, W. Feng, and D. Reed, "A High-Fidelity Software Oscilloscope for Globus," *GlobusWORLD 2004*, San Francisco, California, USA, January 2004.

C156. W. Feng, J. Hurwitz*, H. Newman, S. Ravot, L. Cottrell, O. Martin, F. Coccetti, C. Jin, D. Wei, and S. Low, "Optimizing 10-Gigabit Ethernet for Networks of Workstations, Clusters, and Grids: A Case Study," *ACM/IEEE SC 2003: High-Performance Networking and Computing Conference*, Phoenix, Arizona, USA, November 2003. (Acceptance Rate: 29% = 60/207)

C157. M. Veeraraghavan, X. Zheng**, H. Lee, M. Gardner, and W. Feng, "CHEETAH: Circuit-Switched High-Speed End-to-End Transport ArcHitecture," *SPIE/IEEE Optical Networking and Computer Communications Conference (OptiComm)*, Dallas, Texas, USA, October 2003. (Acceptance Rate: 35%) **Best Paper Award**

C158. W. Feng, "Green Destiny + mpiBLAST = Bioinfomagic," *10th International Conference on Parallel Computing 2003 (ParCo'03)*, Dresden, Germany, September 2003.

C159. J. Hurwitz* and W. Feng, "Initial End-to-End Performance Evaluation of 10-Gigabit Ethernet," *11th IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2003.

C160. A. Darling*, L. Carey*, and W. Feng, "The Design, Implementation, and Evaluation of mpiBLAST," *4th International Conference on Linux Clusters: The HPC Revolution 2003* in conjunction with the *ClusterWorld Conference & Expo 2003*, San Jose, California, USA, June 2003. (Acceptance Rate: ~40%) **Best Paper: Applications Track**

C161. S. Thulasidasan*, W. Feng, and M. Gardner, "Optimizing GridFTP Through Dynamic Right-Sizing," *12th IEEE Symposium on High-Performance Distributed Computing (HPDC'03)*, Seattle, Washington, USA, June 2003. (Acceptance Rate: 20% = 25/125.)

C162. M. Gardner, W. Feng, M. Broxton*, J. Hurwitz*, and A. Engelhart*, "MAGNET: A Tool for Debugging, Analysis, and Adaptation in Computing Systems," *3rd IEEE/ACM Symposium on Cluster Computing and the Grid (CCGrid'03)*, Tokyo, Japan, May 2003. (Acceptance Rate: 34% = 39/114)

C163. M. Gardner, M. Broxton*, A. Engelhart*, and W. Feng, "MUSE: A Software Oscilloscope for Clusters and Grids," *17th IEEE International Parallel & Distributed Processing Symposium*, Nice, France, April 2003. (Acceptance Rate: 29% = 119/407)

C164. W. Feng, and S. Vanichpun*, "Ensuring Compatibility Between TCP Reno and TCP Vegas," *3rd IEEE Symposium on Applications and the Internet (SAINT'03)*, Orlando, Florida, USA, January 2003. (Acceptance Rate: 29% = 43/149)

C165. W. Feng, "An Integrated Multimedia Environment for Speech Recognition Using Handwriting and Written Gestures," *36th Hawai'i International Conference on System Sciences (HICSS-36)*, Big Island, Hawaii, USA, January 2003.

C166. M. Warren, E. Weigle*, and W. Feng, "High-Density Computing: A 240-Processor Beowulf in One Cubic Meter," *IEEE/ACM SC 2002: High-Performance Networking and Computing Conference*, Baltimore, Maryland, USA, November 2002. (Acceptance Rate: 29% = 67/230)

C167. W. Feng, A. Kapadia*, and S. Thulasidasan*, "GREEN: Proactive Queue Management over a Best-Effort Network," *IEEE GLOBECOM*, Taipei, Taiwan, R.O.C., November 2002. (Acceptance Rate: 31% = 606/1980)

C168. S. Vanichpun* and W. Feng, "On the Transient Behavior of TCP Vegas," *11th IEEE International Conference on Computer Communications and Networks (IC3N'02)*, Miami, Florida, USA, October 2002. (Acceptance Rate: 29%)

C169. W. Feng, "Securing Wireless Communication in Heterogeneous Environments," *IEEE MILCOM*, Anaheim, California, USA, October 2002.

C170. W. Feng, M. Warren, and E. Weigle*, "The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs," *4th IEEE International Conference on Cluster Computing (IEEE Cluster)*, Chicago, Illinois, USA, September 2002. (Acceptance Rate: 39% = 45/116)

C171. W. Feng and J. Al-Muhtadi*, "A General Security Infrastructure for Wireless Communication," *International Conference on Networks (ICN'02)*, Atlanta, Georgia, USA, August 2002.

C172. N. Rao and W. Feng, "Performance Tradeoffs of TCP Adaptation Methods," *IEEE International Conference on Networks (ICN'02)*, Atlanta, Georgia, USA, August 2002.

C173. W. Feng, M. Warren, and E. Weigle*, "Honey, I Shrunk the Beowulf!" *31st International Conference on Parallel Processing (ICPP'02)*, Vancouver, Canada, August 2002. (Acceptance Rate: 36% = 67/188)

C174. M. Gardner, W. Feng, and M. Fisk*, "Dynamic Right-Sizing in FTP (drsFTP): An Automatic Technique for Enhancing Grid Performance," *IEEE Symposium on High-Performance Distributed Computing (HPDC'02)*, Edinburgh, Scotland, July 2002. (Acceptance Rate: 31%)

C175. E. Weigle* and W. Feng, "A Comparison of TCP Automatic-Tuning Techniques for Distributed Computing," *IEEE Symposium on High-Performance Distributed Computing (HPDC'02)*, Edinburgh, Scotland, July 2002. (Acceptance Rate: 31%)

C176. J. Hay, W. Feng, and M. Gardner, "Capturing Network Traffic with a MAGNeT," *5th USENIX Annual Linux Showcase & Conference (ALS'01)*, Oakland, California, USA, November 2001.

C177. E. Weigle* and W. Feng, "Dynamic Right-Sizing in TCP: A Simulation Study," *IEEE International Conference on Computer Communications and Networks (IC3N'01)*, Scottsdale, Arizona, USA, October 2001. (Acceptance Rate: 31%)

C178. W. Feng, J. Hay*, and M. Gardner, "MAGNeT: Monitor for Application-Generated Network Traffic," *IEEE International Conference on Computer Communications and Networks (IC3N'01)*, Scottsdale, Arizona, USA, October 2001. (Acceptance Rate: 31%)

C179. J. Al-Muhtadi*, W. Feng, and M. Fisk*, "An Inter-Realm, Cyber-Security Infrastructure for Virtual Supercomputing," *Los Alamos Computer Science Institute Symposium (LACSI'01)*, Santa Fe, New Mexico, USA, October 2001.

C180. A. Feng**, W. Feng, and G. Belford, "Packet Spacing: An Enabling Mechanism for the Delivery of Multimedia Content," *Los Alamos Computer Science Institute Symposium (LACSI'01)*, Santa Fe, New Mexico, USA, October 2001.

C181. M. Fisk* and W. Feng, "Dynamic Right-Sizing in TCP," *Los Alamos Computer Science Institute Symposium (LACSI'01)*, Santa Fe, New Mexico, USA, October 2001.

C182. J. Hay, W. Feng, and M. Gardner, "The Design, Implementation, and Evaluation of MAGNeT," *Los Alamos Computer Science Institute Symposium (LACSI'01)*, Santa Fe, New Mexico, USA, October 2001.

C183. F. Petrini*, W. Feng, A. Hoisie, S. Coll, and E. Frachtenberg, "The Quadrics Network (QsNet): High-Performance Clustering Technology," *9th IEEE International Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2001.

C184. E. Weigle* and W. Feng, "A Case for TCP Vegas in High-Performance Computational Grids," *IEEE International Symposium on High-Performance Distributed Computing (HPDC'01)*, San Francisco, California, USA, August 2001. (Acceptance Rate: 39%)

C185. A. Kapadia*, A. Feng, and W. Feng, "The Effects of Inter-Packet Spacing on the Delivery of Multimedia Content," *IEEE International Conference on Distributed Computing Systems (ICDCS'01)*, Phoenix, Arizona, USA, April 2001. (Acceptance Rate: 31%)

C186. W. Feng and P. Tinnakornsrisuphap*, "The Failure of TCP in High-Performance Computational Grids," *IEEE/ACM SC 2000: High-Performance Networking and Computing Conference*, Dallas, Texas, November 2000. (Acceptance Rate: 29%)

C187. W. Feng, "Network Traffic Characterization of TCP," *IEEE MILCOM 2000*, Los Angeles, CA, October 2000.

C188. W. Feng and P. Tinnakornsrisuphap*, "The Adverse Impact of the TCP Congestion-Control Mechanism in Heterogeneous Computing Systems," *International Conference on Parallel Processing (ICPP'00)*, Toronto, Canada, August 2000. (Acceptance Rate: 38%)

C189. F. Petrini* and W. Feng, "Buffered Co-Scheduling: A New Methodology for Multitasking Parallel Jobs on Distributed Systems," *14th IEEE International Parallel & Distributed Processing Symposium (IPDPS'00)*, Cancun, Mexico, May 2000. (Acceptance Rate: 35%)

C190. P. Tinnakornsrisuphap*, W. Feng, and I. Philp, "On the Burstiness of the TCP Congestion-Control Mechanism in a Distributed Computing System," *IEEE International Conference on Distributed Computing Systems (ICDCS'00),* Taipei, Taiwan, R.O.C., April 2000.

C191. F. Petrini* and W. Feng, "Scheduling with Global Information in Distributed Systems," *IEEE International Conference on Distributed Computing Systems (ICDCS'00)*, Taipei, Taiwan, R.O.C., April 2000.

C192. F. Petrini and W. Feng, "Efficient Resource Utilization on a Massively Parallel System," *7th International Conference on Advanced Computing and Communications (ADCOM'99)*, Roorkee, India, December 1999.

C193. D. Tolmie, T. M. Boorman, A. DuBois, D. DuBois, W. Feng, and I. Philp, "From HiPPI-800 to HiPPI-6400: A Changing of the Guard and Gateway to the Future," *6th International Conference on Parallel Interconnects (PI '99)*, Anchorage, Alaska, USA, October 1999.

C194. W. Feng, "Dynamic Client-Side Scheduling in a Real-Time CORBA System," *IEEE International Computer Software and Applications Conference (COMPSAC 99)*, Phoenix, Arizona, USA, October 1999.

C195. D. Hull, W. Feng, and J. W.-S. Liu, "Operating System Support for Imprecise Computation," *AAAI Fall Symposium on Flexible Computation*, Cambridge, Massachusetts, USA, November 1996.

C196. D. L. Hull, W. Feng, and J. W.-S. Liu, "Enhancing the Performance and Dependability of Real-Time Systems," *IEEE International Computer Performance and Dependability Symposium*, Erlangen, Germany, April 1995.

C197. V. Lopez-Millan, W. Feng, and J. W.-S. Liu, "Using the Imprecise-Computation Technique for Congestion Control on a Real-Time Traffic Switching Element," *IEEE International Conference on Parallel and Distributed Systems*, Hsinchu, Taiwan, R.O.C., December 1994.

C198. W. Feng, "Parallel Spinodal Decomposition," *26th Annual Summer Computer Simulation Conference*, San Diego, California, USA, July 1994.

C199. W. Feng, "Using Handwriting and Gesture Recognition to Correct Speech-Recognition Errors," **Best Paper Award**, *10th International Conference on Advanced Science and Technology*, Chicago, Illinois, USA, March 1994.

C200. W. Feng, "A Natural Language Interface to Paper-Based Maps," *ACM International Conference on Human-Computer Interaction*, Boston, Massachusetts, USA, September 1989.

Refereed Articles in Workshops

Articles listed below were refereed by two or more reviewers, unless otherwise noted.

W1. U. Kalim*, M. Gardner, and W. Feng, "A Non-Invasive Approach for Realizing Resilience in MPI," *7th Workshop on Fault Toerlance for HPC at eXtreme Scale (FTXS), in conjunction with the 26th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Washington, DC, June 2017.

W2. V. Adhinarayanan*, B. Dutta*, and W. Feng, "Making a Case for Green High-Performance Visualization via Embedded Graphics Processors," *14th IEEE Workshop on High-Performance, Power-Aware Computing (HPPAC)*, in conjunction with the *32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Vancouver, British Columbia, USA, May 2016.

W3. K. Hou**, S. Che, and W. Feng, "Auto-Tuning Strategies for Parallelizing Sparse Matrix-Vector (SpMV) Multiplication on Multi- and Many-Core Processors," *7th International Workshop on Accelerators and Hybrid Exascale Systems (in conjunction with the 31st IEEE International Parallel & Distributed Processing Symposium)*, Orlando, FL, USA, May 2017.

W4. C. Hsu and W. Feng, "The Right Metric for Efficient Supercomputing: A Ten-Year Retrospective," *12th IEEE Workshop on High-Performance, Power-Aware Computing (HPPAC)*, in conjunction with the *30th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Chicago, Illinois, USA, May 2016.

W5. A. Feng and W. Feng, "Parallel Programming with Pictures in a Snap!," *6th NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar)*, in conjunction with the *30th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Chicago, Illinois, USA, May 2016.

W6. A. Feng**, E. Tilevich, and W. Feng, "Block-based Programming Abstractions for Explicit Parallel Computing," *IEEE Blocks and Beyond Workshop*, in conjunction with the *IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, Atlanta, Georgia, USA, October 2015.

W7. R. Kalidas*, M. Daga, K. Krommydas, and W. Feng, "On the Performance, Energy, and Power of Data-Access Methods in Heterogeneous Computing Systems," *11th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 29th IEEE International Parallel & Distributed Processing Symposium)*, Hyderabad, India, May 2015.

W8. V. Adhinarayanan*, W. Feng, J. Woodring, D. Rogers, and J. Ahrens, "On the Greenness of In-Situ and Post-Processing Visualization Pipelines," *11th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 29th IEEE International Parallel & Distributed Processing Symposium)*, Hyderabad, India, May 2015.

W9. X. Shu**, J. Zhang*, D. Yao, and W. Feng, "Rapid and Parallel Content Screening for Detecting Transformed Data Exposure," *IEEE International Workshop on Security and Privacy in Big Data (BigSecurity)*, co-located with *IEEE INFOCOM*, Hong Kong, China, April 2015. (Acceptance Rate: 26%)

W10. B. Subramaniam* and W. Feng, "On the Energy Proportionality of Distributed NoSQL Data Stores," *5th International Workshop on High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation (PMBS)*, in conjunction with *ACM/IEEE SC|14: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, New Orleans, Louisiana, USA, November 2014. (Acceptance Rate: 20%)

W11. K. Hou*, H. Wang*, and W. Feng, "Delivering Parallel Programmability to the Masses via the Intel MIC Ecosystem: A Case Study," *7th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, in conjunction with the *43rd International Conference on Parallel Processing,* Minneapolis, Minnesota, USA, September 2014.

W12. J. Yin**, J. Zhang**, J. Wang, and W. Feng, "SDAFT: A Novel Scalable Data Access Framework for Parallel BLAST," *ACM/IEEE International Workshop on Data-Intensive Scalable Computing Systems (DISCS)* held in conjunction with the *ACM/IEEE SC|13: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Denver, Colorado, USA, November 2013.

W13. A Aji*, P. Balaji, J. Dinan, W. Feng, and R. Thakur, "Synchronization and Ordering Semantics in Hybrid MPI+GPU Programming," *3rd IEEE International Workshop on Accelerators and Hybrid Exascale Systems* (in conjunction with the *27th IEEE International Parallel & Distributed Processing Symposium*) Boston, Massachusetts, USA, May 2013.

W14. P. Sathre*, W. Feng, and M. Gardner, "Lost in Translation: Challenges in Automating CUDA-to-OpenCL Translation," *5th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) (in conjunction with the 41st International Conference on Parallel Processing)*, Pittsburgh, Pennsylvania, USA, September 2012.

W15. B. Subramaniam* and W. Feng, "The Green Index: A Metric for Evaluating System-Wide Energy Efficiency in HPC System," *8th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 26th IEEE International Parallel & Distributed Processing Symposium)*, Shanghai, China, May 2012.

W16. F. Ji**, A. Aji*, J. Dinan, D. Buntinas, P. Balaji, W. Feng, and X. Ma, "Efficient Intranode Communication in GPU-Accelerated Systems", *2nd IEEE International Workshop on Accelerators and Hybrid Exascale Systems (AsHES) (in conjunction with the 26th IEEE International Parallel & Distributed Processing Symposium)*, Shanghai, China, May 2012.

W17. T. Scogland*, B. Subramaniam*, and W. Feng, "Emerging Trends on the Evolving Green500: Year Three," *7th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 25th IEEE International Parallel & Distributed Processing Symposium)*, Anchorage, Alaska, USA, May 2011.

W18. U. Kalim*, E. Brown*, M. Gardner, and W. Feng, "Enabling Renewed Innovation in TCP by Establishing an Isolation Boundary," *8th International Workshop on Protocols for Future, Large-Scale and Diverse Network Transports*, Lancaster, Pennsylvania, USA, November 2010.

W19. W. Feng and H. Lin*, "The Green500 List: Year Two," *6th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 24th IEEE International Parallel & Distributed Processing Symposium)*, Atlanta, Georgia, USA, May 2010.

W20. S. Huang*, S. Xiao*, and W. Feng, "On the Energy Efficiency of Graphics Processing Units for Scientific Computing," *5th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 23rd IEEE International Parallel & Distributed Processing Symposium)*, Rome, Italy, May 2009.

W21. W. Feng and T. Scogland*, "The Green500 List: Year One," *5th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 23rd IEEE International Parallel & Distributed Processing Symposium)*, Rome, Italy, May 2009.

W22. W. Feng, P. Balaji, and A. Singh*, "Network Interface Cards as First-Class Citizens," *1st Workshop on the Influence of I/O on Microprocessor Architecture (in conjunction with the 15th IEEE International Symposium on Computer Architecture)*, Raleigh, North Carolina, USA, February 2009.

W23. W. Feng, A. Ching**, and C. Hsu*, "Green Supercomputing in a Desktop Box," *3rd IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 21st IEEE International Parallel & Distributed Processing Symposium)*, Long Beach, California, USA, March 2007.

W24. W. Feng, "Real-Time Transcription of Radiology Dictation: A Case Study for Multimedia TabletPCs," *Microsoft eScience Workshop*, Baltimore, Maryland, USA, October 2006.

W25. W. Feng, "Global Climate Warming? Yes … In the Machine Room," *Clusters and Computational Grids for Scientific Computing*, Flat Rock, North Carolina, USA, September 2006. (Invited Talk)

W26. S. Sharma*, C. Hsu*, and W. Feng, "Making a Case for a Green500 List," *2nd IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 20th IEEE International Parallel & Distributed Processing Symposium)*, Rhodes, Greece, April 2006.

W27. A. Banerjee**, W. Feng, B. Mukherjee, and D. Ghosal, "End-System Performance-Aware Transport over Optical Circuit-Switched Connections," *IEEE INFOCOM High-Speed Networking Workshop: The Terabits Challenge (in conjunction with the 25th IEEE INFOCOM)*, Barcelona, Spain, April 2006.

W28. V. Vishwanath**, P. Balaji**, W. Feng, J. Leigh, and D. Panda, "A Case for UDP Offload Engines in LambdaGrids," *4th International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet'06)*, Nara, Japan, February 2006.

W29. C. Hsu*, W. Feng, and J. Archuleta*, "Towards Efficient Supercomputing: A Quest for the Right Metric," *1st IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 19th IEEE International Parallel & Distributed Processing Symposium)*, Denver, Colorado, USA, April 2005.

W30. A. Banerjee**, W. Feng, D. Ghosal, and B. Mukherjee, "Routing and Scheduling Large File Transfers over Lambda Grids," *3rd International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet'05)*, Lyon, France, February 2005.

W31. C. Hsu* and W. Feng, "Effective Dynamic Voltage Scaling through CPU-Boundedness Detection," *4th IEEE/ACM Workshop on Power-Aware Computer Systems (in conjunction with the 37th IEEE/ACM International Symposium on Microarchitecture)*, Portland, Oregon, USA, December 2004.

W32. M. Veeraraghavan, X. Zheng**, W. Feng, H. Lee, E. Chong, and H. Li, "Scheduling and Transport for File Transfers on High-speed Optical Circuits," *2nd International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet'04)*, Argonne, Illinois, USA, February 2004.

W33. W. Feng, M. Fisk*, M. Gardner, and E. Weigle*, "Dynamic Right-Sizing," *7th IFIP/IEEE Workshop on Protocols for High-Speed Networks*, Berlin, Germany, April 2002.

W34. M. Gardner, W. Feng, and J. Hay*, "Monitoring Protocol Traffic with a MAGNeT," *3rd Passive & Active Measurement Workshop*, Ft. Collins, Colorado, USA, March 2002.

W35. E. Weigle* and W. Feng, "TICKETing High-Speed Traffic with Commodity Hardware and Software," *3rd Passive & Active Measurement Workshop*, Ft. Collins, Colorado, USA, March 2002.

W36. E. Frachtenberg, F. Petrini*, S. Coll, and W. Feng, "Gang Scheduling with Lightweight User-Level Communication," *Workshop on Scheduling and Resource Management for Cluster Computing (in conjunction with the International Conference on Parallel Processing),* Valencia, Spain, September 2001.

W37. F. Petrini*, A. Hoisie, W. Feng, and R. Graham, "Performance Evaluation of the Quadrics Interconnection Network," *IEEE Workshop on Communication Architectures for Clusters (in conjunction with the IEEE International Parallel & Distributed Processing Symposium)*, San Francisco, California, USA, April 2001.

W38. E. Weigle*, W. Feng, and M. Gardner, "Why TCP Will Not Scale for the Next-Generation Internet," *IEEE Workshop on Local and Metropolitan Area Networks (LANMAN 2001)*, Boulder, Colorado, USA, March 2001.

W39. W. Feng, "The Future of High-Performance Networking," *Workshop on New Visions for Large-Scale Networks: Research & Applications*, Invited Paper, Vienna, Virginia, USA, March 2001. (Sponsors: Federal Large-Scale Networking Working Group, DARPA, DOE, NASA, NIST, NLM, and NSF.)

W40. W. Feng, U. Syyid*, and J. W.-S. Liu, "Providing for an Open Real-Time CORBA," *IEEE Workshop on Middleware for Distributed Real-Time Systems and Services (in conjunction with the IEEE Real-Time Systems Symposium)*, Raleigh-Durham, North Carolina, USA, December 1997.

W41. W. Feng and J. W.-S. Liu, "Time-Constrained Speech Processing and Generation," *IEEE Workshop on Real-Time Applications*, New York, New York, USA, May 1993.

W42. A. Chien and W. Feng, "Efficient Implementation of Concurrent Object-Oriented Programs," *Workshop on Languages and Compilers for Parallel Computing*, New Haven, Connecticut, USA, May 1992.

W43. A. Chien and W. Feng, "GST: Grain-Size Tuning for Efficient Execution of Symbolic Programs," *Workshop on Compilation of Symbolic Languages for Parallel Computers*, San Diego, California, USA, October 1991.

Refereed Posters in Conferences (unless otherwise noted)

Articles listed below were refereed by two or more reviewers. Acceptance rates are reported, when available.

P1. K. Youssef**, K. Iwabuchi, W. Feng, M. Gokhale, and R. Pearce, "Towards Optimizing Memory Mapping of Persistent Memory in UMap," *SC|20: The International Conference for High Performance Computing, Networking, Storage and Analysis*, Atlanta, Georgia, USA, November 2020.

P2. V. Adhinarayanan* and W. Feng, "Approximate Pattern Matching for On-Chip Interconnect Traffic Prediction," *29th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Atlanta, Georgia, USA, October 2020.

P3. I. Herrera*, P. Mandke*, W. Feng, and G. Cao, "A Deep-Learning Neural Network Based Reconstruction Algorithm for Sparse-View CT," *2020 Joint AAPM/COMP Meeting*, Vancouver, British Columbia, Canada, July 2020.

P4. X. Cui**, T. Scogland, B. de Supinski, and W. Feng, "Performance Evaluation of the NVIDIA Tesla V100: Block-Level Pipelining vs. Kernel-Level Pipelining," SC|18: The International Conference for High Performance Computing, Networking, Storage and Analysis, Denver, Colorado, USA, November 2018. (Acceptance Rate: 93/165 = 56%.)

P5. S. Pumma**, M. Si, W. Feng, and P. Balaji, "I/O Bottleneck Investigation in Deep Learning Systems," *47th International Conference on Parallel Processing*, Eugene, Oregon, USA, August 2018. **Best Student Poster Award.**

P6. S. Dash*, R. Anandakrishnan, and W. Feng, "Identifying Carcinogenic Multi-hit Combinations using Weighted Set Cover Algorithm," *47th International Conference on Parallel Processing*, Eugene, Oregon, USA, August 2018.

P7. X. Cui**, T. Scogland, B. de Supinski, and W. Feng, "Performance Evaluation of the NVIDIA Tesla P100: Our Directive-Based Partitioning and Pipelining vs. NVIDIA's Unified Memory," SC|17: The International Conference for High Performance Computing, Networking, Storage and Analysis, Denver, Colorado, USA, November 2017. (Acceptance Rate: 98/169 = 58%.)

P8. X. Yu*, K. Hou*, H. Wang*, and W. Feng, "Hierarchical Automata Construction for Approximate Pattern Matching on Automata Processors," IEEE International Symposium on Workload Characterization (IISWC), Seattle, Washington, October 2017.

P9. A. Verma* and W. Feng, "Accelerating 3D-Structured Grid on FPGAs via OpenCL: A Case Study with OpenDwarfs," ACM Student Research Competition, *ACM International Symposium on Code Generation and Optimization*, Austin, Texas, USA, February 2017.

P10. K. Krommydas*, A. Helal*, A. Verma*, and W. Feng, "Bridging the Performance-Programmability Gap for FPGAs via OpenCL: A Case Study with OpenDwarfs," *24th IEEE International Symposium on Field-Programmable Custom Computing Machines*, Washington DC, USA, May 2016.

P11. V. Adhinarayanan*, S. Pakin, D. Rogers, W. Feng, and J. Ahrens, "Performance, Power, and Energy of In-Situ and Post-Processing Visualization: A Case Study in Climate Simulation," *ACM/IEEE SC|15: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC),* Austin, Texas, USA, November 2015. **Best Poster Finalist.** (Acceptance Rate: 138/318 = 43%.)

P12. X. Cui*, T. Scogland, B. de Supinski, and W. Feng, "Directive-Based Pipelining Extension for OpenMP," *ACM/IEEE SC|15: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC),* Austin, Texas, USA, November 2015. (Acceptance Rate: 138/318 = 43%.)

P13. A. Feng**, W. Feng, and E. Tilevich, "PPP: Parallel Programming with Pictures," *ACM/IEEE SC|15: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC),* Austin, Texas, USA, November 2015. (Acceptance Rate: 138/318 = 43%.)

P14. X. Shu, J. Zhang, D. Yao, and W. Feng, "Rapid Screening of Transformed Data Leaks with Efficient Algorithms and Parallel Computing," *5th ACM Conference on Data and Application Security ad Privacy (CODASPY)*, San Antonio, Texas, USA, March 2015. **Best Poster Award.**

P15. P. Sathre* and W. Feng, "MetaMorph: A Modular Library for Democratizing the Acceleration of Parallel Computing Across Heterogeneous Devices," *ACM/IEEE SC|14: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC)*, Denver, Colorado, USA, November 2014. (Acceptance Rate: 76/193 = 39%)

P16. K. Krommydas, M. Owaida, C. Antonopoulos, N. Bellas, and W. Feng, "On the Portability of the OpenCL Dwarfs on Fixed and Reconfigurable Parallel Platforms," *19th IEE International Conference on Parallel and Distributed Systems (ICPADS)*, Seoul, Korea, December 2013.

P17. C. del Mundo* and W. Feng, "Enabling Efficient Intra-Warp Communication for Fourier Transforms in a Many-Core Architecture," ACM Student Research Competition, *ACM/IEEE SC|13: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC)*, Denver, Colorado, USA, November 2013. **Best Undergraduate Poster Award** from the ACM Student Research Competition. (Acceptance Rate: 83/211 = 39%)

P18. X. Shu, J. Zhang*, D. Yao, and W. Feng, "High-Performance, Data-Leak Detection," *IEEE Symposium on Security & Privacy*, San Francisco, California, USA, May 2013.

P19. U. Kalim*, M. Gardner, E. Brown, and W. Feng, "Cascaded TCP: Big Throughput for Big Data Applications in Distributed HPC" (poster), *ACM/IEEE SC|12: The International Conference for High-Performance Computing, Networking, Storage and Analysis (SC)*, Salt Lake City, Utah, USA, November 2012. (Acceptance Rate: 82/179 = 46%)

P20. B. Subramaniam* and W. Feng, "Characterizing the Performance and Energy Efficiency of Simultaneous Multithreading in Multicore Environments," *41st International Conference on Parallel Processing*, Pittsburgh, Pennsylvania, USA, September 2012.

P21. L. Thibério, L. Rangel, J. Setubal, W. Feng, H. Lin, and N. Mohamed, "Analysis of Proteobacterial Genomes Using Phylogenomic Networks," *Brazilian Symposium on Bioinformatics*, 2012.

P22. K. Lee*, H. Lin*, and W. Feng, "Characterizing the Impact of Memory-Access Techniques on AMD Fusion," *ACM/IEEE SC|11: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Seattle, Washington, USA, November 2011. (Acceptance Rate: 35%)

P23. M. Daga* and W. Feng, "Accelerating Molecular Modeling Using GPUs," *2nd GPU Technology Conference*, San Jose, California, USA, September 2010.

P24. J. Addesa*, J. Archuleta*, and W. Feng, "On the Efficacy of Haskell for High-Performance Computational Biology" (poster), *ACM/IEEE SC|09: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Austin, Texas, USA, November 2009. **Best Undergraduate Poster Award.**

P25. G. Martinez*, M. Gardner, and W. Feng, "Characterizing and Optimizing Virtualization Overhead for Portable High-Performance Networking" (poster), *ACM/IEEE SC|08: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Austin, Texas, USA, November 2008. **Best Undergraduate Poster Award.**

P26. F. Moraes* and W. Feng, "Analyzing and Enhancing Parallel I/O in mpiBLAST" (poster), *ACM/IEEE SC 2003: High-Performance Networking and Computing Conference*, Phoenix, Arizona, USA, November 2003. (Acceptance Rate: 35%)

P27. W. Feng and F. Moraes*, "mpiBLAST: A Tool for Interactive Bioinformatics" (poster), *NIH Symposium on Bioinformatics and Computational Biology – Digital Biology: The Emerging Paradigm*, Bethesda, Maryland, USA, November 2003.

P28. K. Pattabiraman**, W. Feng, and D. Reed, "Profile-Based Dynamic Voltage Scaling for I/O-Intensive Codes" (poster), *Los Alamos Computer Science Institute Symposium (LACSI'03)*, Santa Fe, New Mexico, USA, October 2003.

P29. A. Darling*, A. Engelhart*, and W. Feng, "An Open-Source Parallelization of BLAST" (poster), *O'Reilly Bioinformatics Technology Conference (BioCon'03)*, San Diego, California, USA, February 2003.

P30. A. Darling*, L. Carey*, and W. Feng, "mpiBLAST: Delivering Super-Linear Speedup with an Open-Source Parallelization of BLAST" (poster), *8th Pacific Symposium on Biocomputing (PSB'03)*, Lihue, Hawaii, USA, January 2003.

P31. A. Darling* and W. Feng, "mpiBLAST: Parallelization of BLAST for Computational Clusters" (poster), *IEEE/ACM SC 2002: High-Performance Networking and Computing Conference*, Baltimore, Maryland, USA, November 2002. (Acceptance Rate: 35%)

P32. A. Darling* and W. Feng, "BLASTing Off with Green Destiny" (poster), *IEEE Computer Society Bioinformatics Conference (CSB'02),* now known as the *IEEE Computational Systems Bioinformatics Conference*, Palo Alto, California, USA, August 2002.

P33. M. Fisk* and W. Feng, "Dynamic Right-Sizing: TCP Flow-Control Adaptation" (poster), *SC 2001: High-Performance Networking and Computing Conference*, Denver, Colorado, USA, November 2001. (Acceptance Rate: 35%)

(Highly Cited) Technical Reports

T1. A. McCaskey, E. Dumitrescu, D. Liakh, M. Chen, W. Feng, and T. Humble, "Extreme-Scale Programming Model for Quantnum Acceleration within High-Performance Computing," *arXiv preprint arXiv:1710.01794*, October 2017.

T2. K. Krommydas, A. Helal, A. Verma, and W. Feng, "Bridging the Performance-Programmability Gap for FPGAs via OpenCL: A Case Study with OpenDwarfs," *Virginia Tech Technical Report TR-16-03*, May 2016.

T3. A. Verma, A. Helal, K. Krommydas, and W. Feng, "Accelerating Workloads on FPGAs via OpenCL: A Case Study with OpenDwarfs," *Virginia Tech Technical Report TR-16-04*, May 2016.

T4. V. Adhinarayanan and W. Feng, "An Automated Framework for Characterizing and Subsetting GPGPU Workloads," *Virginia Tech Technical Report TR-15-06*, December 2015.

T5. A. Aji, M. Daga, and W. Feng, "CampProf: A Visual Performance Analysis Tool for Memory-Bound GPU Kernels," *Virginia Tech Technical Report TR-10-10*, October 2010.

T6. R. Anandakrishnan, T. Scogland, A. Fenley, J. Gordon, W. Feng, and A. Onufriev, "Accelerating Electrostatic Surface Potential Calculation with Multiscale Approximation on Graphics Processing Units," *Virginia Tech Technical Report TR-09-15*, 2009.

T7. A. Aji and W. Feng, "Accelerating Data-Serial Applications on Data-Parallel GPGPUs: A Systems Approach," *Virginia Tech Technical Report TR-08-24*, 2008.

PATENTS

1. Adaptive Real-Time Methodology for Optimizing Energy-Efficient Computing, 2011. U.S. Patent No. 7,971,073. Issued June 28, 2011.
2. Parallel Metadata Apparatus and Method for Distributed Computation, Communication, and Input-Output, 2009. U.S. Patent Application No. 12/492,794.
3. Off-Chip Access Workload Characterization Methodology for Optimizing Computing Efficiency, 2009. U.S. Patent Application No.: 12/372,286.
4. Buffered Co-Scheduling: A New Methodology for Multitasking Parallel Jobs on Distributed Systems, 2006. U.S. Patent No. 6,993,764. Issued January 31, 2006.

PANELS

1. Supercomputing – What Does the Future Hold?, *National High-Performance Computing and Communications Council (HPCC)*, Newport, Rhode Island, USA, March 2014.
2. Crossfire, *National High-Performance Computing and Communications Council (HPCC)*, Newport, Rhode Island, USA, March 2014.
3. Fusion Processors and HPC, *AMD Fusion Developer Summit*, Bellevue, Washington, June 2011.

4. On the Three P's of Heterogeneous Computing: Performance, Power and Programmability, *ACM/IEEE SC 2010: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, New Orleans, Louisiana, USA, November 2010.
5. Accelerators: Fad, Fashion, or Future?, *39th International Conference on Parallel Processing*, San Diego, California, USA, September 2010.
6. Alternative Processors Technology, *HPC User Forum*, Roanoke, Virginia, USA, April 2009.
7. Green IT Issues: Power, Cooling, and Facility Space, *HPC User Forum*, Tucson, Arizona, USA, September 2008. Panel Moderator.
8. User Experiences with Alternative Processors*, HPC User Forum*, Tucson, Arizona, USA, September 2008.
9. Is HPC Going Green?, *23rd International Supercomputing Conference*, Dresden, Germany, June 2008.
10. When Optical Networking Meets Grid Computing?, *IEEE International Conference on Computer Communications and Networks*, Arlington, Virginia, USA, October 2006. Panel Moderator.
11. The Six-Million Processor System, *ACM/IEEE SC2005: The International Conference on High-Performance Computing, Networking, and Storage*, Seattle, Washington, USA, November 2005. Panel Moderator.
12. Tour de HPCycles, *ACM/IEEE SC2005: The International Conference on High-Performance Computing, Networking, and Storage*, Seattle, Washington, USA, November 2005. Panel Moderator.
13. Ethernet vs. EtherNot, *13th IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2005.
14. How Would You Architect a 100,000-Processor Cluster?, *19th International Parallel & Distributed Processing Symposium Workshop on Communication Architecture for Clusters*, Denver, Colorado, USA, April 2005.
15. Optical Networking and Grid Computing, *1st IEEE/ACM International Conference on Broadband Networks*, San Jose, California, USA, October 2004.
16. Grids: Hype, Substance, or Renaissance, *33rd International Conference on Parallel Processing*, Montreal, Quebec, Canada, August 2004. Panel Moderator.
17. Why Should Users Care About Blades?, *Server Blade Summit*, San Jose, California, USA, March 2004.
18. Battle of the Network Stars!, *ACM/IEEE SC 2003: High-Performance Networking and Computing Conference*, Phoenix, Arizona, USA, November 2003. Panel Moderator.
19. Practical Supercomputing, *ACM/IEEE SC 2003: High-Performance Networking and Computing Conference (Birds-of-a-Feather Panel)*, Phoenix, Arizona, USA, November 2003.
20. High-Performance Interconnection Networks for Cluster Computing, *11th IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)*, Palo Alto, California, USA, August 2003.
21. What is the Future of TCP?, *INET 2002*, Arlington, Virginia, USA, June 2002. Panel Moderator.
22. Applications of Server Blades, *Server Blade Summit*, San Jose, California, USA, May 2002.
23. The Adequacy of TCP for High-Performance Computing, *IEEE/ACM SC2000: High-Performance Networking & Computing Conference*, Dallas, Texas, USA, November 2000.
24. Real-Time CORBA, IEEE *Real-Time Systems Symposium (RTSS)*, San Francisco, California, USA, December 1997.

SELECTED INVITED TALKS & COLLOQUIA

1. *At the Synergistic Intersection of Parallel Computing, Data Analytics, and Machine Learning*, MIT Lincoln Laboratory, Lexington, Massachusetts, USA, September 2019. Invited Talk.
2. *CoreTSAR: Core Task-Size Adapting Runtime*, Barcelona Supercomputing Center (BSC), Barcelona, Spain, July 2019. Invited Talk.
3. *Confessions of an Accidental Greenie: From Green Destiny to Green500,* Universitat Politècnica de València, Valencia, Spain, July 2019. Invited Talk.
4. *On the Three P's of High-Performance Computing: Programmability, Portability, and Performance,* General Electric (GE) Research, Niskayuna, New York, USA, February 2019. Invited Talk.
5. *Forget About the Clouds, Shoot For the MOON,* Grid'5000+FIT Conf.*,* INRIA Sophia Antipolis, April 2018. Invited Talk.
6. *Confessions of an Accidental Greenie: From Green Destiny to Green500,* SC|17 Workshop on Energy-Efficient Supercomputing (E2SC), Denver, CO, November 2017. Invited Talk.

7. *The Center for High-Performance Reconfigurable Computing (CHREC): A Ten-Year Odyssey*, 14th ACM International Conference on Computing Frontiers, May 2017. Invited Talk.
8. *An Ecosystem for the New HPC: Heterogeneous Parallel Computing,* Frontiers in Computing and Data Science, Michigan State University, Lansing, Michigan, USA, October 2016. Invited Talk.
9. *Confessions of an Accidental Greenie: From Green Destiny to Green500*, International Workshop on HPC Architecture, Software, and Application at an Extreme Scale, Wuxi Supercomputing Center, Wuxi, China, September 2015. Invited Talk.
10. *Synergistic Computing: A Renaissance in Computing,* Student Transition Engineering Program (STEP), Center for the Enhancement of Engineering Diversity (CEED) at Virginia Tech, Blacksburg, Virginia, July 2015. Keynote Talk.
11. *Confessions of an Accidental Greenie: From Green Destiny to Green500*, Distinguished Lecture Series, George Mason University, Fairfax, Virginia, April 2015. Invited Talk.
12. *Is Energy-Efficient Exascale an Oxymoron?* ExtremeGreen 2014: Extreme Green & Energy Efficiency in Large Scale Distributed Systems in conjunction with the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Chicago, Illinois, USA, May 2014. Invited Talk.
13. *DNA Sequencing in the Cloud,* PASS Business Analytics Conference, San Jose, California, USA, May 2014. Invited Talk.
14. *Accelerating Data-Intensive Genomic Analysis in the Cloud,* eScience in the Cloud, Redmond, Washington, USA, April 2014. Invited Talk.
15. *Trends in Energy-Efficient Computing: A Perspective from the Green500*, High-Performance Computing and Communications USA (HPCC USA), Newport, Rhode Island, USA, March 2014. Endnote Talk.
16. *An Ecosystem for the New HPC: Heterogeneous Parallel Computing*, Naval Surface Warfare Center – Dahlgren Division, Dahlgren, Virginia, January 2014. Distinguished Lecture.
17. *On the Road to Exascale: A Perspective from the Green500*, HPC China, Changsha, China, May 2013. Invited Talk.
18. *Synergistic Co-Design for BIG DATA*, The White House, Washington, D.C., USA, May 2013. Invited Talk.
19. *Computing in the Sky: MOON Computing + Cloud Computing*, BioIT World Conference & Expo, Boston, Massachusetts, USA, April 2013. Invited Talk.
20. *Towards an Ecosystem for Heterogeneous Parallel Computing*, Yale University, New Haven, Connecticut, USA, March 2013. Invited Talk.
21. *Forget About the Clouds, Shoot For the MOON*, BioIT World Cloud Computing Summit, San Francisco, California, USA, September 2012.
22. *An Ecosystem for Heterogeneous Parallel Computing*, DoD Advanced Computing Systems (ACS), Catonsville, Maryland, MD, July 2012.
23. *CU2CL: An Automated CUDA-to-OpenCL Source-to-Source Translator*, AMD Fusion Developer Summit, Bellevue, Washington, USA, June 2012.
24. *Towards Computing the Cure for Cancer*, GPU Technology Conference, San Jose, California, USA, May 2012. Endnote Talk.
25. *An Ecosystem for the New HPC: Heterogeneous Parallel Computing*, 2nd International Workshop on Accelerators and Hybrid Exascale Systems (AsHES) in conjunction with the 26th International Parallel & Distributed Processing Symposium, Shanghai, China, May 2012.
26. *An Ecosystem for the New HPC: Heterogeneous Parallel Computing*, Wake Forest University, Winston-Salem, North Carolina, USA, April 2012.
27. *Will Computers Replace Humans?*, Kids Tech University, Blacksburg, Virginia, USA, February 2012.
28. *On the Evolution of the Green500 to Exascale*, SIAM Conference on Parallel Processing for Scientific Computing, Savannah, Georgia, USA, February 2012.
29. *An Ecosystem for Heterogeneous Parallel Computing*, University of Central Florida, Orlando, Florida, USA, December 2011.
30. *How to Run Your CUDA Program Anywhere*, NVIDIA Theater, ACM/IEEE International Conference on High-Performance Computing, Networking, Storage, and Analysis (SC), Seattle, Washington, USA, November 2011. http://nvidia.fullviewmedia.com/fb/nv-sc11/tabscontent/archive/309-wed-feng.html.
31. *Research Challenges for Applications & Algorithms in Green HPC*, Workshop on Green High-Performance Computing (HPC), Piscataway, New Jersey, USA, October 2011.
32. *An Ecosystem for Heterogeneous Parallel Computing*, AMD Research, Austin, Texas, USA, September 2011.
33. *On the Three P's of Heterogeneous Computing*, Workshop on Parallel Programming on Accelerator Clusters (PPAC) at IEEE Cluster, Austin, Texas, USA, September 2011. Keynote Talk.

34. *An Ecosystem for the New HPC – Heterogeneous Parallel Computing*, Fall Creek Falls Conference by Oak Ridge National Laboratory, Gatlinburg, Tennessee, USA, September 2011.
35. *Energy-Efficient E-putting Everywhere*, 20th ACM International Symposium on High-Performance Distributed Computing (HPDC) in conjunction with FCRC '11: Federated Computing Research Conf., San Jose, California, USA, June 2011. Keynote Talk.
36. *OpenCL and the 13 Dwarfs*, AMD Fusion Developer Summit, Bellvue, Washington, USA, June 2011. Invited Talk.
37. *GPU Computing and the Green500*, NVIDIA Theater, ACM/IEEE International Conference on High-Performance Computing, Networking, Storage, and Analysis (SC), New Orleans, Louisiana, USA, November 2010.
38. *Energy-Efficient Exascale Computing: An Oxymoron?*, International Conference on Energy-Aware High-Performance Computing, Hamburg, Germany, September 2010. Keynote Talk.
39. *On the Power of Heterogeneous Computing in a Box*, AMD GPU Day, Boston, Massachusetts, USA, June 2010. Keynote Talk.
40. *When High-Performance Computing Meets Bioinformatics*, Wake Forest University – School of Medicine, Winston Salem, North Carolina, USA, June 2010.
41. *Green Supercomputing Comes of Age*, Center for Experimental Research in Computer Systems, Georgia Tech, Atlanta, Georgia, USA, April 2010.
42. *Green Supercomputing Comes of Age*, IEEE International Parallel & Distributed Processing Symposium (IPDPS) Workshop on High-Performance, Power-Aware Computing, Atlanta, Georgia, USA, April 2010. Keynote Talk.
43. *Green Supercomputing Comes of Age*, SIAM Conference on Parallel Processing for Scientific Computing, Seattle, Washington, USA, February 2010.
44. *Reliable Cloud Computing on Volatile Resources*, Cloud Computing for DoD & Government (Sponsored by the Institute for Defense and Government Advancement), Alexandria, Virginia, USA, February 2010.
45. *To GPU Synchronize or Not GPU Synchronize?*, NVIDIA Theater, ACM/IEEE International Conference on High-Performance Computing, Networking, Storage, and Analysis (SC), Portland, Oregon, USA, November 2009. http://nvidia.fullviewmedia.com/sc09/nvidia-sc09-wu-chun-feng-virgina-tech.html
46. *Green Supercomputing Comes of Age*, HPC China, Changsha, China, October 2009. Keynote Talk.
47. *Optimizing Efficiency in Emerging Chip Multiprocessors: Multi-Core, Call, and GPGPU*, HPC Adoption '09, Burlingame, California, USA, May 2009.
48. *Green Networking: Does It Really Matter?*, IEEE International Parallel & Distributed Processing Symposium (IPDPS) Workshop on Communication Architectures for Clusters, Rome, Italy, May 2009.
49. *Color Me Green! Color Me Fast!*, 2009 Conference on High Speed Computing, Gleneden Beach, Oregon, USA, April 2009. Keynote Talk.
50. *The Origin, Evolution, and Application of mpiBLAST for Parallel Sequence Search*, Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA, October 2008.
51. *Massively Parallel Sequence-Search with mpiBLAST: From Research to Production*, IBM, Philadelphia, Pennsylvania, USA, September 2008. Endnote Talk.
52. *Green Computing Comes of Age*, HPC User Forum, Tucson, Arizona, USA, September 2008.
53. *Parallel Sequence Search with mpiBLAST*, IBM Life Sciences Virtual Expo, April 2008.
54. *The Power in Being Low Power*, EDS Data Center Design Charrette, Plano, Texas, USA, February 2008. Dinner Keynote Talk.
55. *Ethernet vs. Ethernot: Fad or Fashion?*, 16th IEEE International Conference on Computer Communications and Networks, Honolulu, Hawaii, USA, August 2007.
56. *iSERViCE: A 21st Century Cyberinfrastructure for Commerce, Research, and Education,* Southern Virginia Higher Education Center (SVHEC) Roundtable, South Boston, Virginia, USA, September 2007.
57. *The Power in Being Low Power*, Intel Corporation, Hillsboro, Oregon, USA, April 2007.
58. *The Power in Being Low Power*, Argonne National Laboratory, Argonne, Illinois, USA, March 2007.
59. *The Evolution of Power-Aware Systems*, Merrill Lynch Corporation, New York, New York, USA, March 2007.
60. *The Evolution of Power-Aware, High-Performance Clusters: From the Datacenter to the Desktop*, APC & Microsoft InfraStruXure User Group (Inaugural Meeting), Microsoft Technology Center – Silicon Valley, Mountain View, California, USA, November 2005.
61. *The Evolution of Power-Aware, High-Performance Clusters: From the Datacenter to the Desktop*, System Research Seminar Series, North Carolina State University, Raleigh, North Carolina, USA, April 2005.

62. *The Evolution of Power-Aware, High-Performance Clusters: From the Datacenter to the Desktop*, IEEE International Parallel & Distributed Processing Symposium (IPDPS) Workshop on High-Performance, Power-Aware Computing, Denver, Colorado, USA, April 2005. Keynote Talk.
63. *Frontiers in High-End Computing*, University of Texas at Dallas, Dallas, Texas, USA, March 2005. University-Wide Talk, Hosted by Provost.
64. *Green Destiny: A 240-Node Energy-Efficient Supercomputer in Five Square Feet*, University of New Mexico, Albuquerque, New Mexico, USA, October 2004.
65. *Bridging the Disconnect Between the Network and Large-Scale Scientific Applications*, ACM SIGCOMM Workshop on Network-I/O Convergence: Experiences, Lessons, and Implications (NICELI), Karlsruhe, Germany, August 2003.
66. *Green Destiny: A "Cool" 240-Node Cluster in a Telephone Booth*, Future Computing Conference at the Royal United Services Institute for Defence and Security Studies, London, England, July 2003. Keynote Talk.
67. *Report from the High-Performance Networking Workshop*, DOE Science Network Workshop, Reston, Virginia, USA, June 2003.
68. *Green Destiny: Energy-Efficient Supercomputing*, 15th Annual E-Source Forum, Colorado Springs, Colorado, USA, November 2002.
69. *Systems & Network Research for Grids*, Argonne National Laboratory, Argonne, Illinois, USA, October 2002.
70. *Green Destiny: A 240-Node Compute Cluster in One Cubic Meter*, U. of Illinois at Urbana-Champaign (also broadcast over the Internet via the Access Grid), Urbana, Illinois, USA, October 2002.
71. *Green Destiny: A 240-Node Compute Cluster in One Cubic Meter*, Eli Lilly and Company, Indianapolis, Indiana, USA, September 2002.
72. *The Future of High-Performance Networking*, Rice U., Houston, Texas, USA, January 2002.
73. *High-Performance Networking Research: Issues for Today's & Tomorrow's High-Performance Computing Environments*, DOE Scientific Discovery through Advanced Computing, Reston, Virginia, USA, January 2002.
74. *High-Speed Measurement and Monitoring with Commodity Parts*, Rice U., Houston, Texas, USA, January 2002.
75. *High-Speed Network Monitoring and Measurement with Commodity Parts*, DARPA Next-Generation Internet, McLean, Virginia, USA, January 2002.
76. *The Software Metaphor for LAN PHY ≠ WAN PHY: Why High-Speed Networking in Clusters ≠ High-Speed Networking in Grids*, 10-Gigabit Ethernet Workshop, San Diego Supercomputing Center, San Diego, California, USA, October 2001.
77. *The Future of High-Performance Networking*, Workshop on New Visions for Large-Scale Networks: Research & Applications, Vienna, Virginia, USA, March 2001. (Sponsors: Federal Large Scale Networking Working Group, DARPA, DOE, NASA, NIST, NLM, NSF.)
78. *The Failure of TCP over High-Performance Computational Grids*, U. of Illinois at Urbana-Champaign, Urbana, Illinois, USA, January 2001.
79. *Buffered Coscheduling: A New Methodology for Multitasking Parallel Jobs on Distributed Systems*, U. of Oregon, Eugene, Oregon, USA, June 2000.
80. *Network Interface Cards as First-Class Citizens*, The Ohio State U., Columbus, Ohio, USA, November 1999.

RESEARCH GRANTS

Dr. Feng conducts theoretical and applied research of *efficient systems* across a wide breadth of areas – including high-performance networking and computing, cloud computing, grid computing, energy-efficient computing, accelerator-based computing, real-time systems, bioinformatics, and computer science pedagogy for K-12 – as well as across different layers of abstraction from architecture to system software and middleware to applications software.

External Funding: Basic Research

High-Performance Deep Learning for Dairy Science. Co-Principal Investigator.
Award: $29,819/year, October 2020 – October 2021. (Personal Share: 54%)
Program: Select Sires

A Computational Deep-Learning Approach for Fast, Accurate CT Testing and Monitoring of COVID-19. Principal Investigator.
Award: $200,000/year, June 2020 – May 2021. (Personal Share: 50%)
Program: NSF IIS

Higher Accuracy and Availability of COVID-19 Testing and Monitoring via Post-CT Image Boosting and Analysis. Principal Investigator.
Award: $150,000/year, June 2020 – May 2021. (Personal Share: 50%)

Program: NSF CCF

Productivity Tools for OpenCL on FPGA. Principal Investigator.
Award: $80,000/year, January 2020 – December 2020. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – DOD

At the Synergistic Intersection of High-Performance Computing, Data Analytics, and Machine Learning. Principal Investigator.
Award: $140,000/year, January 2020 – December 2020. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – AMD, Capital One, and MIT Lincoln Laboratory

Identifying the Multi-Hit Combinations of Inherited and Somatic Mutations in Childhood Cancers (Supplement). Principal Investigator.
Award: $6,100/year, January 2020 – May 2020. (Personal Share: 100%)
Program: Edward Via College of Osteopathic Medicine (VCOM)

Models and Evaluation of Scaling Deep Learning at Oak Ridge Leadership Computing (OLCF). Principal Investigator.
Award: $61,354/year, January 2019 – December 2020. (Personal Share: 100%)
Program: DOE Oak Ridge National Laboratory

Productivity Tools for OpenCL on FPGA. Principal Investigator.
Award: $200,000/year, January 2019 – December 2019. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – DOD

High-Performance Computing and Analytics on Heterogeneous Architectures. Principal Investigator.
Award: $120,000/year, January 2019 – December 2019. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center – Capital One, DOD, and MIT Lincoln Laboratory

Phase-I IUCRC Virginia Tech: Center for Space, High-performance, and Resilient Computing (SHREC). Principal Investigator
Award: $600,000/4 years, October 2018 – October 2022. (Personal Share: 100%)
Program: NSF Industry/University Collaborative Research Center

End-to-End Productivity for Heterogeneous Computing. Principal Investigator.
Award: $240,002/year, January 2018 – December 2018. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – DOD

Directive-Based Pipelining Extensions and Unified Memory for GPU Computing. Principal Investigator
Award: $150,668/2 years ($71,203 + $79,465), August 2017 – August 2019. (Personal Share: 100%)
Program: DOE Lawrence Livermore National Laboratory

Identifying the Multi-Hit Combinations of Inherited and Somatic Mutations in Childhood Cancers. Principal Investigator.
Award: $32,295/year, July 2017 – June 2018. (Personal Share: 100%)
Program: Edward Via College of Osteopathic Medicine (VCOM)

Performance Evaluation and Optimization of Parallel Runtime Systems for Large-Scale, Data-Analytics Applications on HPC Supercomputing Environments
Award: $216,021/year, May 2017 – May 2020. (Personal Share: 100%)
Program: U.S. Department of Energy – Exascale Computing Program

A Deep-Learning Approach Towards Auto-Tuning CFD Codes. Principal Investigator
Award: $250,000/year, March 2017 – March 2018. (Personal Share: 15%)
Program: AFOSR Computational Mathematics

Rapid Prototyping via a Grid-based Auto-tuning Framework. Principal Investigator
Award: $80,000/year, January 2017 – December 2017. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – NASA

End-to-End Productivity for Heterogeneous Computing. Principal Investigator.
Award: $240,002/year, January 2017 – December 2017. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) – DOD

HPC Development using Rapid Prototyping with GLAF. Principal Investigator
Award: $47,565/year, June 2016 – December 2016. (Personal Share: 100%)
Program: NASA via Science Systems and Applications, Inc. (SSAI)

End-to-End Productivity for Heterogeneous Computing. Principal Investigator.
Award: $80,000/year, January 2016 – December 2016. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC) -- DOD

Performance, Programmability, and Portability of Heterogeneous Computing. Principal Investigator.

Award: $120,000/year, January 2015 – December 2015. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Optimizing the Energy Usage and Cognitive Value of Extreme Scale Data Analysis Approaches. Co-Principal Investigator.
Award: $2,925,000/3 years, October 2014 – October 2017. (Personal Share: 12%)
Program: DOE Scientific Data Management, Analysis and Visualization at Extreme Scale 2

Parallel Computing for Everyone. Principal Investigator.
Award: $7,000/year, August 2014 – May 2015. (Personal Share: 100%)
Program: CRA-W / CDC Collaborative Research Experience for Undergraduates.

Performance, Programmability, and Portability of Heterogeneous Computing. Principal Investigator.
Award: $120,000/year, January 2014 – December 2014. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

XPS: SDA: Collaborative Research: A Scalable Cross-Layer Framework for Computation Intensive HPC Data Parallel Applications. Co-Principal Investigator.
Award: $750,000/4 years, September 2013 – August 2017. (Personal Share: 50%)
Program: NSF Exploiting Parallelism and Scalability (XPS).

EAGER: Collaborative Research: Democratizing the Teaching of Parallel Computing Concepts. Principal Investigator.
Award: $300,000/2 years, March 2013 – February 2015. (Personal Share: 80%)
Program: NSF EAGER.

Collaborative Research: Phase-II Renewal of CHREC Center at Virginia Tech. Co-Principal Investigator.
Award: $200,000/5 years, March 2013 – February 2018. (Personal Share: 50%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Performance, Programmability, and Portability of Heterogeneous Computing. Principal Investigator.
Award: $120,000/year, January 2013 – December 2013. (Personal Share: 83%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Genomes Galore: Core Techniques, Libraries, and Domain-Specific Languages for High-Throughput DNA Sequencing. Co-Principal Investigator.
Award: $2,000,000/3 years, January 2013 – December 2015. (Personal Share: 17.5%)
Program: NSF BIGDATA.

Co-Design of Hardware/Software for Predicting MAV Aerodynamics. Principal Investigator.
Award: $6,000,859/5 years, September 2012 – October 2017. (Personal Share: 71%)
Program: AFOSR Basic Research Initiative.

Performance, Programmability, and Portability of Heterogeneous Computing. Principal Investigator.
Award: $122,500/year, January 2012 – December 2012. (Personal Share: 71%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Analytical Electrostatics: Methods and Biological Applications. Co-Principal Investigator.
Award: $1,000,000/4 years, August 2011 – July 2015. (Personal Share: 20%)
Program: NIH NIGMS.

STOC: Secure, Tactical On-Demand Computing. Co-Principal Investigator.
Award: $39,498/0.75 years, May 2011 – January 2012. (Personal Share: 100%)
Program: Air Force Research Laboratory.

CIC (RDDC): Commoditizing Data-Intensive Biocomputing in the Cloud. Principal Investigator.
Award: $370,000/2 years, April 2011 – March 2013. (Personal Share: 36%)
Program: NSF Computing in the Cloud.

Performance, Power, Portability, and Programmability of Heterogeneous Computing. Principal Investigator.
Award: $192,500/year, January 2011 – December 2011. (Personal Share: 82%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Kinematic and Waveform Inversion of Microearthquake and Ambient Noise Data for Three-Dimensional Velocity Structure in Geothermal Reservoirs. Co-Principal Investigator.
Award: $708,000/3 years, August 2010 – September 2013. (Personal Share: 40%)
Program: DOE Geothermal Technologies.

A Transformative Software Toolkit for Genome Sequence Alignment. Principal Investigator.
Award: $49.776/year, August 2010 – July 2011. (Note: Award notification was sent but Virginia has not released funding.)

Program: Virginia Innovation Partnership (VIP).

On the Performance, Programmability, and Power of Accelerators. Principal Investigator.
Award: $140,000/year, January 2010 – December 2010. (Personal Share: 75%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

CSR: Small: Collaborative Research: Hybrid Opportunistic Computing for Green Clouds. Co-Principal Investigator.
Award: $470,156/3 years, September 2009 – August 2012. (Personal Share: 32%)
Program: NSF Computer Systems Research.

Characterizing and Optimizing Emerging Devices. Principal Investigator.
Award: $70,000/year, January 2009 – December 2009. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Metrics and Methodologies for High Performance System Energy Benchmarking. Co-Principal Investigator.
Award: $200,000/2 years, September 2008 – August 2010. (Personal Share: 50%)
Program: NSF Computing and Communication Foundations.

Virtual Ecosystems for K-12 Pedagogy in Science, Technology, Engineering, and Mathematics. Principal Investigator.
Award: $37,000/2 years, August 2008 – July 2010. (Personal Share: 100%)
Program: CRA-W / CDC Collaborative Research Experience for Undergraduates.

EcoDaemon: Power-Aware System Software. Principal Investigator.
Award: $10,000/year, January 2008 – August 2008. (Personal Share: 100%)
Program: Southeastern Universities Research Association (SURA) "Intellectual Property to Market" (IP2M).

Network Offloading for Genome Sequence Searching using the SmartNIC. Co-Principal Investigator with RNET Technologies.
Award: $149,704/year, January 2008 – December 2008. (Personal Share: 100%)
Program: NSF Small Business Technology Transfer (STTR).

I/UCRC Proposal for University Site at Virginia Tech: Center for High-Performance Reconfigurable Computing (CHREC). Co-Principal Investigator.
Award: $200,000/5 years, January 2008 – December 2012. (Personal Share: 50%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Process-to-Core Mapping in Advanced Architectures. Principal Investigator.
Award: $70,000/year, January 2008 – December 2008. (Personal Share: 100%)
Program: NSF Industry/University Cooperative Research Center (I/UCRC).

Earth System Grid II. Collaborator.
Award: $60,000/year, October 2004 – September 2006. (Personal Share: 100%)
Agency: DOE Office of Science.
Program: Scientific Discovery through Advanced Computing (SciDAC).
Collaborators: Argonne National Laboratory (PI: I. Foster), Lawrence Livermore National Laboratory (PI: D. Williams), National Center for Atmospheric Research (PI: D. Middleton), and a host of others.

Improvements to mpiBLAST. Principal Investigator.
Award: $10,000/year, March 2004 – June 2004. (Personal Share: 100%)
Agency: Advanced Micro Devices, Inc. (AMD).

Interface Design for High-Performance Networking. Co-Principal Investigator.
Award: $100,000/year, November 2003 – October 2004. (Personal Share: 50%)
Agency: Los Alamos National Laboratory-University of California – Cooperative Agreement on Research and Education.
Collaborators: University of California at Riverside (L. Bhuyan, PI)

An Analytical and Empirical Study of mpiBLAST for Clusters and Grids. Principal Investigator.
Award: 10,000 service units on the PSC TCS1; 50,000 service units on the NCSA TeraGrid Cluster; 50,000 service units on the U. Wisconsin Condor Flock, October 2003 – September 2004. (Personal Share: 100%)
Agency: NSF National Computational Science Alliance: Alliance Allocations Board.

Wide-Area Transport and Signaling Protocols for Genome to Life Applications. Co-Principal Investigator.
Award: $45,214/year, July 2003 – June 2004. (Personal Share: 50%)
Agency: Los Alamos National Laboratory-University of California – Cooperative Agreement on Research and Education.
Collaborators: University of California at Davis (B. Mukherjee, PI and D. Ghosal)

Software Technology to Enable Reliable High-Performance Distributed Disk Arrays. Co-Principal Investigator.
Award: $540,000/3 years, June 2002 – May 2005. (Personal Share: 50%)
Agency: NASA Applied Information Systems Research Program.

Collaborators: T-6 at Los Alamos National Laboratory (Michael S. Warren, PI and Christopher L. Fryer)

INCITE: Edge-Based Traffic Processing and Service Inference for High-Performance Networks. Co-Principal Investigator.
Award: $2,281,000/3 years, July 2002 – July 2005. (Personal Share: 40%)
Agency: DOE Office of Science.
Program: Scientific Discovery through Advanced Computing (SciDAC).
Collaborators: Rice University (R. Baraniuk, PI; R. Riedi; E. Knightly; R. Nowak), Stanford Linear Accelerator Center (L. Cottrell).

High-Performance Transport Protocols. Principal Investigator.
Award: $750,000/3 years, July 2002 – July 2005. (Personal Share: 100%)
Agency: DOE Office of Science.
Program: Base Program.

Prototyping an Earth System Grid. Co-Principal Investigator.
Award: $6,680,900/3 years, July 1999 – June 2002. (Personal Share: 18%)
Agency: DOE Office of Science.
Program: Next-Generation Internet – Applications, Network Technology & Network Testbed Partnerships.
Collaborators: Argonne National Laboratory (I. Foster), Lawrence Berkeley National Laboratory (A. Shoshani, B. Tierney), Lawrence Livermore National Laboratory (D. Williams), National Center for Atmospheric Research (S. Hammond, PI), U. Wisconsin (B. Hibbard), USC/ISI (C. Kesselman).

Network Interface Cards as First-Class Citizens. Principal Investigator.
Award: $453,489/2 years, July 1999 – June 2001. (Personal Share: 100%)
Agency: DOE Office of Science.
Program: Next-Generation Internet – Research in Basic Technologies.

External Funding: Equipment

MRI-R²: Acquisition of a Heterogeneous Supercomputing Instrument for Transformative Interdisciplinary Research. Principal Investigator.
Award: $1,992,527/3 years, July 2010 – June 2013. (Personal Share: 90%)
Program: NSF Major Research Instrumentation

MISER: A High-Performance, Power-Aware Cluster. Senior Personnel. (Personal Share: 0%. Listed as Senior Personnel.)
Award: $453,911/year, September 2007 – September 2009.
Program: NSF Computer Research Instrumentation

External Gifts and Donations

Intel CPU+FPGA Server
Award: $40,000 equipment gift, September 2018. (Personal Share: 100%)
Donor: World Wide Technology

Two AMD Radeon Vega Frontier Edition GPUs
Award: $3,000 equipment gift, June 2018. (Personal Share: 100%)
Donor: AMD

Four AMD Radeon R9 Nano GPUs (Fiji)
Award: $2,600 equipment gift, June 2016. (Personal Share: 100%)
Donor: AMD

Heterogeneous Systems Architecture Award
Award: $150,000, March 2016. (Personal Share: 100%)
Donor: Heterogeneous Systems Architecture (HSA) Foundation

Resource Allocation for the Microsoft Cloud
Award: $25,000, November 2014 (Personal Share: 100%)
Donor: Microsoft

Bittware/Altera Stratix V (S5-PCIe-HQ-D8 with Quartus 16.0)
Award: $15,000, November 2013 (Personal Share: 100%)
Donor: Bittware

NVIDIA Kepler GPUs (Two Kepler K20s at $6,398 total and One Quadro Kepler K5000 at $2,249)
Award: $8,647, December 2012 (Personal Share: 100%)
Donor: NVIDIA

AMD Build Kit with AMD Trinity APU and AMD Tahiti GPU (i.e., AMD Radeon HD 7970)
Award: $1,000, November 2012 (Personal Share: 100%)

Donor: AMD

Intel Deskside Tower with Intel Sandybridge CPU, Intel Knights Ferry MIC Card, and Two Intel Knights Corner MIC Cards
Award: $10,000, October 2012 (Personal Share: 100%)
Donor: Intel

PGI Accelerator Compiler Software
Award: $5,600 software gift, June 2012 (Personal Share: 100%. One of 45 awardees worldwide.)
Donor: PGI in cooperation with NVIDIA

Four NVIDIA Tesla Fermi C2075 Graphics Processing Units (GPUs)
Award: $10,000 equipment gift, March-April 2012 (Personal Share: 100%)
Donor: NVIDIA

Two AMD Radeon HD 7970 (Tahiti): World's First 28nm GPU
Award: $1,300 equipment gift, February-March 2012 (Personal Share: 100%)
Donor: AMD

Four AMD 16-Core Interlagos Processors
Award: $12,000 equipment gift, October 2011. (Personal Share: 100%)
Donor: AMD

AMD Fusion APU System (with AMD Llano A8-3850)
Award: $2,000 equipment gift, September 2011 (Personal Share: 100%)
Donor: AMD

IBM Faculty Award. Principal Investigator.
Award: $15,000 unrestricted gift, September 2011 (Personal Share: 100%)
Donor: IBM

CUDA Research Center. Principal Investigator.
Award: Exclusive event access; priority access to pre-release hardware and software; automatic inclusion in GPU seeding program for each major new NVIDIA GPU; live online training sessions; designated NVIDIA personnel to serve as a technical liaison; special pricing on NVIDIA equipment; joint promotion, public relations, and press activities with NVIDIA, May 2011 (Personal Share: 100%)
Donor: NVIDIA

Compute the Cure for Cancer. Principal Investigator.
Award: $100,000 unrestricted gift, April 2011 (Personal Share: 50%)
Donor: NVIDIA Foundation

AMD Research Faculty Fellowship, Principal Investigator.
Award: $75,000 unrestricted gift, January 2011 (Personal Share: 100%)
Donor: Advanced Micro Devices (AMD)

IBM Faculty Award. Principal Investigator.
Award: $16,000 unrestricted gift, August 2010 (Personal Share: 100%)
Donor: IBM

OpenCL and the 13 Dwarfs. Principal Investigator.
Award: $75,000 unrestricted gift, August 2010 (Personal Share: 100%)
Donor: AMD

IBM Faculty Award. Principal Investigator.
Award: $40,000 unrestricted gift, August 2009 (Personal Share: 100%)
Donor: IBM

NVIDIA Professor Partnership Award. Principal Investigator.
Award: $25,000 unrestricted gift, November 2008 (Personal Share: 100%)
Donor: NVIDIA

IBM Faculty Award. Principal Investigator.
Award: $35,000 unrestricted gift, August 2008 (Personal Share: 100%)
Donor: IBM

High-Performance, Power-Aware Computing. Principal Investigator.
Award: $45,000 equipment gift, September 2004 (Personal Share: 100%)
Agency: Advanced Micro Devices, Inc. (AMD)

Improvements to mpiBLAST. Principal Investigator.
Award: $100,000 equipment gift, March 2004 – June 2004 (Personal Share: 100%)

Agency: Advanced Micro Devices, Inc. (AMD)

Optimizations for 10-Gigabit Ethernet Adapters. Principal Investigator.
Award: $125,000 equipment gift, June 2002 – June 2004 (Personal Share: 100%)
Agency: Intel Corporation

Internal Competitive Funding: Basic Research

High-Performance Data Mining via Machine Learning for Analyzing Ultrasound Images. Co-Principal Investigator.
Award: $150,000/2 years, September 2019 – August 2021 (Personal Share: 25%)
Program: Institute for Critical Technologies and Applied Science (Acceptance Rate: 5/80 = 6.25%.)

SEEC: Synergistic Environments for Experimental Computing.
Award: $375,000/5 years, July 2014 – June 2019 (Personal Share: 100%)
Program: Institute for Critical Technologies and Applied Science and Dept. of Computer Science.

A Scalable Analytical Framework for Ultra-High Dimensional Neuroimaging Data. Co-Principal Investigator.
Award: $25,000/year, January 2018 – June 2018 (Personal Share: 20%)
Program: Institute for Critical Technologies and Applied Science – Data & Decisions.

An Adaptive Run-Time System for Reliable, Energy-Efficient Supercomputing. Principal Investigator.
Award: $300,000/year, October 2005 – September 2006 (Personal Share: 100%)
Program: DOE Weapons-Supported Research: Computer Science Research Foundation.

Software-Based Power-Aware Computing. Principal Investigator.
Award: $295,000/year, October 2003 – September 2005 (Personal Share: 100%)
Agency: DOE Laboratory-Directed Research & Development (Acceptance Rate: < 10%)

Smart Routers for Distributed Computational Grids. Principal Investigator.
Award: $510,000/3 years, October 2000 – September 2003 (Personal Share: 100%)
Agency: DOE Laboratory-Directed Research & Development (Acceptance Rate: < 10%)

Resource Utilization and Parallel Program Development with Buffered Coscheduling. Co-Principal Investigator.
Award: $639,000/3 years, October 2000 – September 2003 (Personal Share: 50%)
Agency: DOE Laboratory-Directed Research & Development (Acceptance Rate: < 10%)
Collaborator: CCS-3 at Los Alamos National Laboratory (F. Petrini, PI)

High-Performance TCP over the Next-Generation Internet. Principal Investigator.
Award: $330,000/3 years, October 1999 – September 2002 (Personal Share: 100%)
Agency: DOE Laboratory-Directed Research & Development (Acceptance Rate: < 10%)

Internally Competed Funding: Equipment

FPGA-Accelerated Computing with OpenCL. Principal Investigator.
Award: $24,819 equipment grant, 2019
Agency: SCHEV/ETF Fund – College of Engineering

Transforming Interdisciplinary Research and Education with Video Cards. Principal Investigator.
Award: $69,195 equipment grant, 2010
Agency: SCHEV/ETF Fund – College of Engineering

Transforming Interdisciplinary Research and Education with Video Cards. Principal Investigator.
Award: $69,195 equipment grant, 2010
Agency: SCHEV/ETF Fund – College of Engineering

Bridging the Great Parallel Divide: Research and Education for Next-Generation Computing Platforms. Co-Principal Investigator.
Award: $62,290 equipment grant, 2010
Agency: SCHEV/ETF Fund – College of Engineering

General-Purpose Computation on Graphics Processing Units. Co-Principal Investigator.
Award: $18,784 equipment grant, 2009
Agency: SCHEV/ETF Fund – Department of Computer Science

The Green500 List. Principal Investigator.
Award: $5,500 equipment grant, 2009
Agency: SCHEV/ETF Fund – Department of Computer Science

Software-Based Power-Aware Computing. Principal Investigator.
Award: $75,000 equipment grant, April 2004 (Personal Share: 100%)
Agency: DOE Laboratory-Directed Research & Development (Acceptance Rate: < 5%)

Direct Internal Funding: Basic Research

A Software Oscilloscope for Distributed Systems. Principal Investigator.
Award: $1,057,000/year, October 2004 – September 2005 (Personal Share: 100%)
Agency: Los Alamos Computer Science Institute

Deployment and Evaluation of Dynamic Right-Sizing over a Wide-Area Network. Co-Principal Investigator.
Award: $179,000/year, October 2004 – September 2005 (Personal Share: 50%)
Agency: DOE Accelerated Strategic Computing Initiative
Program: Distributed & Distance Computing

Reliable Networking in System- and Wide-Area Networks. Principal Investigator.
Award: $910,000/3 years, October 2001 – September 2004 (Personal Share: 100%)
Agency: Los Alamos Computer Science Institute

Improvements to TCP over the Wide-Area Network. Principal Investigator.
Award: $500,000/3 years, October 2001 – September 2004 (Personal Share: 100%)
Agency: DOE Accelerated Strategic Computing Initiative
Program: Distributed & Distance Computing

Research Directions in Networking and Computer Architecture. Principal Investigator.
Award: $1,035,000/3 years, October 1998 – September 2001 (Personal Share: 100%)
Agency: Los Alamos Computer Science Institute

Flow and Congestion Control over High-Speed Networks. Principal Investigator.
Award: $750,000/3 years, October 1998 – September 2001 (Personal Share: 100%)
Agency: DOE Accelerated Strategic Computing Initiative
Program: Distance & Distributed Computing

PROFESSIONAL ACTIVITIES

Selected Appointments & Memberships

Distinguished Scientist (2010-now), Member (1998-now) and Student Member (1989-1998), ACM.
Senior Member (2004-now), Member (1998-2004), and Student Member (1988-1998), IEEE and IEEE Computer Society.
Member, IDC Technical Computing Advisory Panel, 2013-now.
Distinguished Speaker, IEEE Distinguished Visitors Program, 2003-2005.

Editorships, Chairmanships, and Boards

External Advisory Board Member, *High-Performance Computing Center at Oklahoma State University*, 2020 – present.
Associate Editor, *IEEE Transactions on Parallel and Distributed Systems*, 2018 – present.
Steering Committee Chair, *International Conference on Parallel Processing*, 2016 – present.
Program Track Co-Chair: Sustainable and Green Computing, *20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, 2020.
Program Track Co-Chair: State of the Practice, *ACM/IEEE SC|19: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (formerly Supercomputing)*, 2019.
Program Track Co-Chair: High-Performance Computing and Cloud Computing, *International Green and Sustainable Computing Conference (IGSC), formerly the International Green Computing Conference (IGCC)*, 2019.
General Co-Chair, *8th International Conference on Computational Advances in Bio and Medical Sciences*, 2018.
General Co-Chair, *46st International Conference on Parallel Processing*, 2017.
Chair, Emerging Technologies, *ACM/IEEE SC|16: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, 2016.
Program Track Co-Chair: Data Centers and Big Data Computing, *23rd International Conference on Computer Communications and Networks*, 2015.
Editorial Advisory Board, *International Journal of Big Data Intelligence*, 2013 – present.
Vice Chair of Steering Committee, *International Conference on Parallel Processing*, 2013 – 2015.
General Co-Chair, *12th International Conference on Frontiers of Information Technology*, 2014.
Guest Editor, *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB), Special Issue on Advanced Parallel Computing Systems to Accelerate Bioinformatics*, 2014 – 2015.
Program Track Chair: Data Centers and Big Data Computing, *23rd International Conference on Computer Communications and Networks*, 2014.

Workshop Co-Chair, *ExtremeGreen Workshop: Extreme Green & Energy Efficiency in Large Scale Distributed Systems*, 2014. (Sponsored in part by European Action COST IC0804 on "Energy Efficiency in Large Scale Distributed Systems")
Program Co-Chair (Technical Papers), *IEEE/ACM SC (formerly Supercomputing)*, 2013.
Conflict Chair: State of the Practice, *IEEE/ACM SC (formerly Supercomputing)*, 2012.
General Chair, *41st International Conference on Parallel Processing*, 2012.
Posters Co-Chair, *41st International Conference on Parallel Processing*, 2012.
Program Co-Chair, *4th IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, 2012.
Committee Co-Chair, *Data-Enabled Life Sciences Alliance (Personalized Medicine)*, 2012.
Industrial Chair, *25th ACM International Conference on Supercomputing*, 2011.
Editorial Board, *ISRN Communications and Networking*, 2010-present.
General Co-Chair, *IEEE/ACM International Conference on Green Computing and Communications (GreenCom)*, 2010.
Guest Editor, *IEEE Computer, Special Issue on Tools and Environments for Multi- and Many-Core Architectures*, 2009.
Workshops Co-Chair, *38th International Conference on Parallel Processing (ICPP)*, 2009.
Workshops Chair, *11th IEEE International Conference on Cluster Computing*, 2009.
Program Co-Chair, *5th IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)*, 2009.
Program Co-Chair, *1st Workshop on Energy Efficiency in Grids, Clouds and Clusters (held in conjunction with the 10th IEEE/ACM International Conference on Grid Computing*, 2009.
Program Co-Chair, *4th IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)*, 2008.
Guest Editor, Journal of Parallel Computing, Special Issue on High-Performance Computing Using Accelerators, 2007.
International Liaisons Co-Chair, *36th International Conference on Parallel Processing (ICPP),* 2007.
Program Vice-Chair: Networking, *18th ACM/IEEE SC 2006: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, 2006.
Program Co-Chair, *15th IEEE International Conference on Computer Communications and Networks,* 2006.
Program Vice-Chair: Optical and High-Speed Networking, *15th IEEE International Conference on Computer Communications and Networks,* 2006.
Awards Co-Chair, *35th International Conference on Parallel Processing (ICPP),* 2006.
Program Vice-Chair: Networking, *17th ACM/IEEE SC 2005: The International Conference on High-Performance Computing, Networking, and Storage*, 2005.
Program Co-Chair, *34th International Conference on Parallel Processing (ICPP)*, 2005.
General Chair, *1st IEEE/ACM International Workshop on Provisioning and Transport for Hybrid Networks (PATHNets) at the 1st IEEE/ACM International Conference on Broadband Networks (BroadNets)*, formerly Opticomm, 2004.
Program Vice-Chair, *DOE Workshop on Ultra High-Speed Transport Protocols and Dynamic Network Provisioning for Large-Scale Scientific Applications*, May 2003.
Program Vice-Chair: Real-Time Systems, *28th International Conference on Parallel Processing (ICPP)*, 1999.

Organizing & Steering Committees

HPC Technical Computing Advisory Panel, IDC, 2014-present.
Scientific Committee, *International Conference on Computational Methods and Algorithms on HPC Platforms and Accelerators,* September 2017.
Best Paper Award Committee, *44th International Conference on Parallel Processing*, September 2015.
Scientific Computing Steering Committee, 2013 – present.
Grid5000 Scientific Advisory Committee, 2012 – present.
International Supercomputing Conference, 2012 – present.
State Council of Higher Education for Virginia (SCHEV) Outstanding Faculty Award, 2014 – present.
European Action COST IC0804 on "Energy Efficiency in Large Scale Distributed Systems," 2010 – present.
IEEE International Conference on Green Computing and Communications (GreenCom), 2010 – present.
International Conference on Parallel Processing, 2008 – present.
IEEE Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium), 2005 – present.
DOE Scientific Discovery through Advanced Computing (SciDAC) Conference, 2005.

Program Committees

2020
ACM/IEEE SC|20: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Technical Papers Committee. Track: State of the Practice.)

IEEE International Parallel & Distributed Processing Symposium (IPDPS) – Systems Software Track.
IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid) – Sustainable and Green Computing Track.
IEEE Workshop on High-Performance Computational Biology (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2019

ACM/IEEE SC|19: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Technical Papers Committee. Track: State of the Practice. Track Co-Chair.)
International Green and Sustainable Computing Conference (IGSC), formerly the International Green Computing Conference (IGCC)

2018

ACM/IEEE SC|18: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Technical Papers Committee. Track: State of the Practice.)
IEEE International Parallel & Distributed Processing Symposium (IPDPS) – Systems Software Track
International Green and Sustainable Computing Conference (IGSC), formerly the International Green Computing Conference (IGCC)

2017

ACM International Conference on Computing Frontiers (CF)
ACM/IEEE SC|17: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Technical Papers Committee. Track: State of the Practice.)
ACM/IEEE SC|17: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Posters Committee)
ACM/IEEE SC|17: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Emerging Technologies Committee)
IEEE International Parallel & Distributed Processing Symposium (IPDPS) – Multidisciplinary Track
IEEE International Symposium on Computer and Communications
XGreen: Energy Efficiency in Large Scale Distributed Systems (held in conjunction with the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing).

2016

ACM International Symposium on High-Performance Distributed Computing (HPDC)
ACM/IEEE SC|16: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Emerging Technologies Committee)
International Workshop on OpenCL (IWOCL)
IEEE Workshop on Network Science for Communication Networks (NetSciCom 2016, held in conjunction with IEEE INFOCOM)
XGreen: Energy Efficiency in Large Scale Distributed Systems (held in conjunction with the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing).

2015

ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)
ACM/IEEE SC|15: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Technical Program Committee. Track: State of the Practice.)
ACM/IEEE SC|15: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Emerging Technologies Committee)
ACM/SPEC International Conference on Performance Engineering (ICPE)
ACM International Symposium on High-Performance Distributed Computing (HPDC)
ExtremeGreen: Extreme Green & Energy Efficiency in Large Scale Distributed Systems (held in conjunction with the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing).
International Workshop on OpenCL (IWOCL)
IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid)
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE International Conference on Green Computing and Communications (GreenCom)
IEEE International Conference on High-Performance Computing (Track: Applications)
IEEE International Conference on Parallel and Distributed Systems (ICPADS)
IEEE International Workshop on High-Performance Computational Biology (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2014

ACM International Symposium on High-Performance Distributed Computing (HPDC)
ExtremeGreen: Extreme Green & Energy Efficiency in Large Scale Distributed Systems (held in conjunction with the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing).
IEEE International Conference on Big Data Science and Engineering (BDSE)
IEEE International Conference on Communications (Track: Green Communication and Computing)
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE International Conference on Parallel and Distributed Systems (ICPADS)
IEEE International Congress on Big Data
IEEE Multicore and GPU Programming Models, Languages and Compilers Workshop (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE Workshop on High-Performance Computational Biology (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

IEEE Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2013

ACM/IEEE SC|13: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
Energy Efficiency in Large Scale Distributed Systems
ExtremeGreen: Extreme Green & Energy Efficiency in Large Scale Distributed Systems (held in conjunction with the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing).
IEEE Global Communications Conference (GLOBECOM)
IEEE International Conference on Big Data
IEEE International Conference on Communications (Track: Green Communication and Computing)
IEEE International Conference on High-Performance Computing (Track: Architecture)
IEEE International Conference on High-Performance Computing and Communications (HPCC)
IEEE International Conference on Parallel and Distributed Systems (ICPADS)
IEEE Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
International Green Computing Conference (IGCC)

2012

ACM/IEEE SC|12: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
IEEE Global Communications Conference (GLOBECOM)
IEEE International Conference on High Performance Computing and Communications (HPCC)
IEEE Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
Innovative Parallel Computing: Foundations & Applications of GPU, Manycore, and Heterogeneous Systems (InPar)
International Conference on Parallel Processing

2011

ACS/IEEE International Conference on Computer Systems and Applications (AICCSA)
IEEE International Parallel & Distributed Processing Symposium (IPDPS) – Algorithms Track
IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid)
IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom 2011)
IEEE International Conference on High-Performance Computing (Track: Architecture)
IEEE International Conference on High Performance Computing and Communications (HPCC)
IEEE Workshop on High-Performance Computational Biology (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2010

ACM International Symposium on High-Performance Distributed Computing (HPDC)
IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid)
IEEE International Parallel & Distributed Processing Symposium (IPDPS), PhD Forum
IEEE International Symposium on Cluster Computing and the Grid (CCGrid)
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)
International Conference on Cloud Computing
International Conference on Energy-Aware High-Performance Computing (EnA-HPC)
International Conference on High-Performance Computing (Track: Architecture)
International Conference on Parallel and Distributed Systems (ICPADS)
International Conference on Parallel Processing
International Conference on the Virtual Computing Initiative (ICVCI)
International Green Computing Conference (IGCC)
International Workshop on Energy-Efficient High-Performance Computing (EEHiPC)

2009

ACM/IEEE SC|09: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
IEEE INFOCOM
IEEE International Parallel & Distributed Processing Symposium (IPDPS), PhD Forum
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)
International Conference on High-Performance Computing (Track: Architecture)
International Conference on the Virtual Computing Initiative (ICVCI)
Workshop on Parallel Computational Biology

2008

ACM/IEEE SC|08: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE International Symposium on High-Performance Distributed Computing (HPDC)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)
International Conference on High-Performance Computing and Communications (HPCC)
International Conference on Parallel Processing (ICPP)
International Conference on the Virtual Computing Initiative (ICVCI)
International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet)

2007

ACM/IEEE SC|07: The International Conference on High-Performance Computing, Networking, Storage, and Analysis (Poster Committee)
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE Global Communications Conference (GLOBECOM)
IEEE International Symposium on High-Performance Distributed Computing (HPDC)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE Symposium on High-Performance Interconnects (IEEE Hot Interconnects)
International Conference on High-Performance Computing (Track: Communication Networks)
International Conference on High-Performance Computing and Communications (HPCC)
International Conference on the Virtual Computing Initiative (ICVCI)
International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet)

2006

ACM/IEEE SC|06: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE International Conference on Distributed Computing Systems (ICDCS)
IEEE International Parallel & Distributed Processing Symposium (IPDPS)
IEEE/ACM International Workshop on Cluster Security (held in conjunction with the IEEE/ACM International Symposium on Cluster Computing and the Grid)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

IEEE International Workshop on High-Performance Computational Biology (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
International Conference on High-Performance Computing and Communications (HPCC)

2005
ACM/IEEE SC|05: The International Conference on High-Performance Computing, Networking, Storage, and Analysis
IEEE/ACM International Conference on Broadband Networks (Area: Broadband Optical Networking), 2005
IEEE/ACM Workshop on Grids and Advanced Networks (held in conjunction with the IEEE/ACM International Symposium on Cluster Computing and the Grid)
IEEE International Conference on Computer Communications and Networks (IC3N)
IEEE International Symposium on High-Performance Distributed Computing (HPDC)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE International Workshop on High-Performance, Power-Aware Computing (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

IEEE Workshop on High Performance Interconnects for Distributed Computing (held in conjunction with the IEEE International Symposium on High-Performance Distributed Computing)
International Conference on Parallel Processing (ICPP)

2004
High-Performance Computing Symposium
IEEE/ACM Workshop on Grids and Advanced Networks (held in conjunction with the IEEE/ACM International Symposium on Cluster Computing and the Grid)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
IEEE International Workshop on Networks for Grid Applications (GridNets) (held in conjunction with the IEEE International Conference on Broadband Networks, formerly Opticomm)
International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet)

2003
IEEE International Conference on Local Computer Networks (LCN)
IEEE International Parallel & Distributed Processing Symposium (IPDPS)
IEEE International Symposium on High-Performance Distributed Computing (HPDC)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2002
IEEE International Conference on Local Computer Networks (LCN)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)

2001
IEEE International Conference on Local Computer Networks (LCN)
IEEE International Symposium on High-Performance Distributed Computing (HPDC)
IEEE International Workshop on Communication Architectures for Clusters (held in conjunction with the IEEE International Parallel & Distributed Processing Symposium)
International Conference on Parallel Processing (ICPP)
Workshop on Scheduling and Resource Management for Cluster Computing (held in conjunction with the 30th International Conference on Parallel Processing)

2000
ACM/IEEE SC|00: The International Conference on High-Performance Computing, Networking, Storage, and Analysis

Birds-of-a-Feather (BoF) Sessions

BoF Co-Organizer, "The Green500: Trends in Energy-Efficient Computing," *ACM/IEEE SC2020: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2020.
BoF Co-Organizer, "The Green500: Trends in Energy-Efficient Computing," *International Supercomputing Conference*, Frankfurt, Germany, June 2020.

BoF Co-Organizer, "The Green500 and its Continuing Evolution," *ACM/IEEE SC2019: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2019.
BoF Co-Organizer, "The Green500: Trends in Energy-Efficient Computing," *International Supercomputing Conference*, Frankfurt, Germany, June 2019.
BoF Co-Organizer, "The Green500 and its Continuing Evolution," *ACM/IEEE SC2018: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2018.
BoF Co-Organizer, "The Green500: Trends in Energy-Efficient Computing," *International Supercomputing Conference*, Frankfurt, Germany, June 2018.
BoF Co-Organizer, "The Green500 and its Continuing Evolution," *ACM/IEEE SC2017: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2017.
BoF Co-Organizer, "The Green500: Trends in Energy-Efficient Computing," *International Supercomputing Conference*, Frankfurt, Germany, June 2017.
BoF Co-Organizer, "The Green500 and its Continuing Evolution," *ACM/IEEE SC2016: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2016.
BoF Co-Organizer, "The Green500 List and its Continuing Evolution," *ACM/IEEE SC2015: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2015.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2014: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2014.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2013: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2013.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2012: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2012.
BoF Co-Organizer, "Setting Trends for Energy-Efficient Supercomputing," *ACM/IEEE SC2012: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Seattle, WA, November 2012.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2011: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Salt Lake City, UT, November 2011.
BoF Co-Organizer, "Setting Trends for Energy-Efficient Supercomputing," *ACM/IEEE SC2011: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Seattle, WA, November 2011.
BoF Invitee, "Critically Missing Pieces in Heterogeneous Accelerator Computing," *ACM/IEEE SC2011: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Seattle, WA, November 2011.
BoF Co-Organizer, "Setting Trends for Energy-Efficient Supercomputing," *26th International Supercomputing Conference*, Hamburg, Germany, June 2011.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2010: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, New Orleans, LA, November 2010.
BoF Co-Organizer, "Setting Trends for Energy-Efficient Supercomputing," *25th International Supercomputing Conference*, Hamburg, Germany, May 2010.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2009: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*, Portland, OR, November 2009.
BoF Co-Organizer, "The Green500 List," *ACM/IEEE SC2008: The International Conference on High-Performance Computing, Networking, Storage, and Analysis*,
BoF Co-Organizer, "The Green500 List: Past, Present, and Future," *23rd International Supercomputing Conference,* Dresden, Germany, June 2008.

Proposal Reviewing & Site Visiting

DOE Office of Science, 2002-2006, 2009.
NSF Computer & Information Science & Engineering (CISE), 1998, 2001, 2004, 2009-2012.
NSF Office of Cyberinfrastructure (OCI), 2011-2012.
DOE Laboratory-Directed Research & Development (Directed Research), Los Alamos National Laboratory, 2001-2002.
DOE Laboratory-Directed Research & Development (Exploratory Research), Computer Science & Software Engineering, Los Alamos National Laboratory, 1999.

Journal & Additional Conference Reviewing

ACM Journal of Experimental Algorithmics, 2006.
ACM Transactions on Modeling and Computer Simulation, 2010.
IEEE Computer, 2009, 2012.
IEEE Computer Graphics & Applications, 2002.
IEEE Communications Letters, 2002-2003.

IEEE International Symposium on High-Performance Computer Architecture, 2014.
IEEE Transactions on Parallel & Distributed Systems, 2000-2004, 2007-2010, 2012.
IEEE Network, 2001-2003.
IEEE Spectrum, 2005, 2009-2010.
IEEE Transactions on Computers, 1995, 2000, 2001, and 2008.
IEEE Transactions on Software Engineering, 1997. (See January 1998 issue.)
IEEE Computer, 1996. (Senior Referee Designation. See December 1996 issue.)
IEEE/ACM SC (2003), IEEE ICC (2003), IEEE INFOCOM (2002-2004), IEEE WCNC (2002), IEEE Real-Time Systems Symposium (1994-1996), IEEE Real-Time Technology & Applications Symposium (1995-1996), IEEE International Conference on Distributed Computing Systems (1995), IEEE Workshop on Real-Time Applications (1994).
International Journal of High-Performance Computing Applications (IJHPCA), 2006.
International Journal of High-Performance Computing and Networking (IJHPCN), 2005.
Journal of Microprocessors and Microsystems, 2010.
Journal of Parallel & Distributed Computing, 2004, 2009.
Journal of Parallel Computing, 2009, 2012.

Diversity and Outreach Activities

Mentor, WHPC Mentoring Programme, Women in High-Performance Computing (WHPC), Edinburgh, Scotland, United Kingdom, October 2019 – present.
Panelist, Meet-n-Greet, Society of Asian Scientists and Engineers (SASE), 30 September 2019.
Speaker, Society of Asian Scientists and Engineers (SASE), 18 March 2019.
Invited Speaker, *Synergistic Computing: A Renaissance in Computing,* Student Transition Engineering Program (STEP), Center for the Enhancement of Engineering Diversity (CEED), 16 July 2015.
Founder, Director, and Advisor for Computer Science Pedagogy in K-12, MyVICE: My Virtual Instances of Computing Environments, Virginia Tech, 2008 – present.
Founder & Director, The Green500, 2008 – present.
Keynote Speaker, Kids Tech University (for kids aged 9 to 12), Virginia Tech, February 2012.
Exhibitor, Kids Tech University (for kids aged 9 to 12), Virginia Tech, Spring Semester of 2009 – 2012.
Advisor & Mentor for Collaborative Research Experiences for Undergraduates – Females & Underrepresented Minorities, Computing Research Association for Women (CRA-W) and Coalition to Diversity Computing (CDC), Virginia Tech, 2008 – 2010, 2014 – 2015.
Dine with Faculty, Galileo and Hypatia Mentoring Program in the Center for the Enhancement of Engineering Diversity (CEED), Virginia Tech, 2010 – 2015.

Institutional Committees (Ad-Hoc, Departmental, Division, College)

Virginia Tech

- Member of the *Graduate Admissions Committee*, Dept. of Computer Science, August 2019 – August 2020.
- Member of the *Faculty Liaison to Tech Staff Committee*, Dept. of Computer Science, August 2017 – August 2020.
- Member of the *Diversity Committee*, Dept. of Computer Science, August 2016 – August 2019.
- Faculty Search Committee for *Data Analytics/Cyber Search*, Dept. of Computer Science, August 2015 – May 2016.
- Faculty Search Committee for *Human-Computer Interaction and Visualization Search*, Dept. of Computer Science, August 2013 – May 2014.
- Faculty Search Committee for *Systems/Compiler/Runtime Search*, Dept. of Computer Science, August 2012 – August 2013.
- Member of the *Committee on Entrepreneurship & Intrapreneurship*, College of Engineering, August 2011 – August 2012.
- Member of the *High-Performance Computing (HPC) Investment Committee*, College of Engineering, April 2010 – August 2012.
- Member of the *Graduate Program Committee*, Dept. of Computer Science, August 2011 – August 2013.
- Member of the *Systems & Networking Qualifier Committee*, Dept. of Computer Science, August 2009 – March 2012.
- Member of the *Diversity Committee*, Dept. of Computer Science, May 2009 – August 2015.
- Coordinator for *Graduate Preview Weekend*, Dept. of Computer Science, August 2011 – August 2014.
- Faculty Mentor for *Scieneering*, HHMI Grant, May 2011 – May 2012.
- Computer Science Representative at the *Graduate Recruiting and Professional Development Forum*, September 2011.
- Faculty Mentorship of Vicki Choi, Dept. of Computer Science, November 2009 – August 2011.
- Chair of the *Systems & Networking Qualifier Committee*, Dept. of Computer Science, August 2009 – August 2010.

- Member of the *Graduate Admissions Committee*, August 2009 – August 2011.
- Member of the *Advanced Research Computing (ARC) User Advisory Board*, Office of Information Technology, April 2008 – November 2011.
- Member of the *Executive Committee*, Dept. of Computer Science, August 2007 – August 2008.
- Member of the *Graduate Program Committee*, Dept. of Computer Science, August 2007 – August 2008.
- Member of the *Ad-Hoc Committee on Academic Year Salary Reduction*, College of Engineering, August 2007 – January 2008.
- Member of the *Department Head Search Committee*, Dept. of Computer Science, August 2007 – November 2007.
- Coordinator for an *Ad-Hoc Departmental Seminar Series*, Dept. of Computer Science, August 2006 – May 2007. (Speakers included Prof. Daniel A. Reed as a distinguished lecture; Dr. Ernst Dow; and Prof. Jack Dongarra as a distinguished lecture.)
- Faculty Presenter and Participant at *Hokie Focus – Undergraduate Recruiting*, Dept. of Computer Science, April 2006.
- Faculty Coordinator for a *Computer Organization I* course, Dept. of Computer Science, March 2006 – August 2011. (This course was formally approved as part of the CS undergraduate curriculum in the 2008-2009 academic year and listed as CS 2505.)
- Member of the *Faculty Search Committee for Systems*, Dept. of Computer Science, August 2005 – May 2006.
- Chair for *Petascale Computing Initiative*, College of Engineering, Office of Information Technology, and Office of Research, December 2005 – July 2006.

Los Alamos National Laboratory

- Member of the *Study Committee for a Computer Science Division*, January 2000 – June 2000.
- Fellow Member of the *Los Alamos Computer Science Institute*, November 1998 – January 2006.

THESIS SUPERVISION

Ph.D. Chair (or Co-Chair)

Graduated

1. Xuewen Cui, *Directive-Based Data Partitioning and Pipelining and Auto-Tuning for High-Performance GPU Computing*, Computer Science, Virginia Tech, December 2020.
2. Sajal Dash, *Exploring the Landscape of Big-Data Analytics through Domain-Aware Incremental and Approximate Algorithms*, Computer Science, Virginia Tech, August 2020.
3. Sarunya Pumma, *Scalability Analysis and Optimization for Large-Scale Deep Learning*, Computer Science, Virginia Tech, February 2020.
4. Ahmed Helal, *Automated Runtime Analysis and Adaptation for Scalable Heterogeneous Computing*, Electrical & Computer Engineering, January 2020.
5. Vignesh Adhinarayanan, *Models and Techniques for Green High-Performance Computing*, Computer Science, Virginia Tech, October 2019.
6. Kaixi Hou, *Exploring Performance Portability for Accelerators via High-level Parallel Patterns,* August 2018.
7. Jing Zhang, *Transforming and Optimizing Irregular Applications for Parallel Architectures*, Computer Science, February 2018.
8. Umar Kalim, *Cognizant Networks: A Model and Framework for Session-based Communications and Adaptive Networking*, August 2017.
9. Konstantinos Krommydas, *Towards Enhancing Performance, Programmability, and Portability in Heterogeneous Computing*, May 2017.
10. Thomas R. W. Scogland, *Runtime Adaptation for Autonomic Heterogeneous Computing,* Computer Science, Virginia Tech, 2015.
11. Ashwin M. Aji, *Programming High Performance Clusters with Heterogeneous Computing Devices*, Computer Science, Virginia Tech, 2015.
12. Balaji Subramaniam, *Metrics, Models and Methodologies for Energy-Proportional Computing*, Computer Science, Virginia Tech, 2015. **Outstanding Ph.D. Student Award (2015)** across the College of Engineering at Virginia Tech.
13. Shucai Xiao, *Generalizing the Utility of Graphics Processing Units in Large-Scale Heterogeneous Computing Systems*, Electrical & Computer Engineering, Virginia Tech, July 2013.
14. Marwa Elteir, *Towards a MapReduce Framework for Heterogeneous Computing Architectures*, Computer Science, Virginia Tech, June 2013.
15. Kumaresh Singh, *Efficient Computational Tools for Variational Data Assimilation and Information Content Estimation*, Computer Science, Virginia Tech, August 2010. (Co-chaired with Adrian Sandu.)

In Progress

1. Frank Wanye, *Scalable Graph Processing via High-Performance Computing*, Computer Science, Virginia Tech, 2022.
2. Karim Youssef, *Scalable Frameworks for Big Data Processing*, Computer Science, Virginia Tech, 2022.

M.S. Chair

Graduated

1. Anshuman Verma, *On the Programmability and Performance of OpenCL Designs for FPGA*, Electrical & Computer Engineering, Virginia Tech, March 2018.
2. Bishwajit Dutta, *Power Analysis and Prediction for Heterogeneous Computation,* Electrical & Computer Engineering, Virginia Tech, February 2018.
3. Rubasri Kalidas, *Performance and Power of Data-Access Methods in Heterogeneous Computing Systems*, Electrical & Computer Engineering, Virginia Tech, 2015
4. Lokendra Panwar, *Performance Modeling of Graphics Processing Units for Heterogeneous Computing Systems,* Computer Science, Virginia Tech, 2013. **Outstanding M.S. Student Award (2014).**
5. Paul Sathre, *On the Characterization and Optimization of a CUDA-to-OpenCL Source-to-Source Translator*, Computer Science, Virginia Tech, June 2013. **Outstanding M.S. Student Award (2013).**
6. Nabeel Mohamed, *Data-Intensive Biocomputing in the Cloud*, Computer Science, Virginia Tech, September 2013.
7. Ryan Braithwaite, *Automated Process-to-Core Mapping in Multi-Core Systems*, Computer Science, Virginia Tech, January 2012.
8. Kenneth Lee, *Characterization and Exploitation of GPU Memory Systems*, Computer Science, Virginia Tech, 2012.
9. Gabriel Martinez, *CU2CL: A CUDA-to-OpenCL Translator for Multi- and Many-Core Architectures*, Computer Science, Virginia Tech. **Outstanding M.S. Student Award (2012).**
10. Mayank Daga, *Algorithmic Mapping and Optimization of N-Body Problems onto the Graphics Processing Unit*, Virginia Tech, May 2011. **Outstanding M.S. Student Award (2011).**
11. Ajeet Singh, *Network Offloading for Genomic Sequence Searching with a SmartNIC*, Electrical & Computer Engineering, Virginia Tech, July 2009.
12. Ashwin Aji, *Mapping Biological Sequence-Search Codes to Emergent Chip Multiprocessors*, Computer Science, Virginia Tech, June 2008. **Outstanding M.S. Student Award (2009).**
13. Ganesh Narayanaswamy, *On the Interaction of High-Performance Network Protocol Stacks with Multicore Architectures*, Computer Science, Virginia Tech, May 2008. **Outstanding M.S. Student Award (2008)**, College of Engineering, Virginia Tech.
14. Christopher J. Goddard, *Analysis and Abstraction of Parallel Sequence Search*, Computer Science, Virginia Tech, September 2007.

In Progress

1. Saikat Dey, *Real-Time Anomaly Detection with High-Performance Computing*, Computer Science, Virginia Tech, 2021.
2. Sonal Jha, *High-Performance Anomaly Detection with Deep Learning*, Electrical & Computer Engineering, 2021.
3. Atharva Gondhalekar, *On the Performance and Programmability of GPUs and FPGAs for Irregular Applications*, Electrical & Computer Engineering, Virginia Tech, 2021.

THESIS COMMITTEES

Ph.D. Committees

1. Yan Huang, *Real-Time Resource Optimization for Wireless Networks*, Electrical & Computer Engineering, Virginia Tech, 2021. (Advisor: Y. Thomas Hou)
2. Steven Ross Glandon, *Time Integration Methods for Large-scale Scientific Simulations*, Computer Science, Virginia Tech, 2020. (Advisor: Adrian Sandu)
3. Moeti M. Masiane, *Towards Insight-Driven Sampling for Interactive Data-Intensive Computing*, Computer Science, Virginia Tech, 2020. (Advisor: Christopher North)
4. Haitham Abdulrahman Elmarakeby, *Deep Learning for Biological Problems*, Computer Science, Virginia Tech, 2017. (Advisor: Lenwood Heath)
5. Andrew Love, *A Modular Flow for Rapid FPGA Design Implementation*, Electrical & Computer Engineering, Virginia Tech, 2015. (Advisor: Peter Athanas)
6. Brent Pickering, *Evaluating the OpenACC API for Parallelization of CFD Applications*, Aerospace & Ocean Engineering, Virginia Tech, 2014. (Advisor: Christopher Roy)
7. Hermanpreet Singh, *Controlling Scalability in Distributed Virtual Environments*, Computer Science, Virginia Tech, 2013. (Advisor: Denis Gracanin)
8. Vedavyas Duggirala, *Open Nework Emulator: A Parallel Direct Code Execution Network Simulator*, Computer Science, Virginia Tech, 2013. (Advisors: Calvin Ribbens and Srinidhi Varadarajan)
9. Henry Monti, *An Integrated End-User Data Service for HPC Centers*, Computer Science, Virginia Tech, 2013. (Advisor: Ali Butt)

10. Tannous Frangieh, *Design Assembly Techniques for FPGA Back-End Acceleration*, Electrical & Computer Engineering, Virginia Tech, 2012. (Advisor: Peter Athanas)
11. Henry Monti, *An Integrated End-User Data Service for HPC Centers*, Computer Science, Virginia Tech, 2012. (Advisor: Henry Monti)
12. Guanying Wang, *Evaluating MapReduce System Performance: A Simulation Approach*, Computer Science, Virginia Tech, 2012. (Advisor: Ali Butt)
13. Michael Henry, *Emerging Power-Gating Techniques for Low-Power Digital Circuits*, Electrical & Computer Engineering, Virginia Tech, 2011. (Advisor: Leyla Nazhandali)
14. M. Mustafa Rafique, *An Adaptive Framework for Managing Heterogeneous Many-Core Clusters*, Computer Science, Virginia Tech, 2011. (Advisor: Ali R. Butt)
15. Dong Li, *Scalable and Energy-Efficient Execution Methods for Multicore Systems*, Computer Science, Virginia Tech, 2011. (Advisor: Kirk W. Cameron)
16. Sushant Sharma, *Cooperation in Wireless Networks*, Computer Science, Virginia Tech, 2011. (Advisor: Y. Thomas Hou)
17. Feng (Andrew) Ge, *A Network Testbed Enabling Cooperative Spectrum Sensing, Dynamic Spectrum Access, and Cross-Layer Node Adaptation*, Electrical & Computer Engineering, Virginia Tech, 2010. (Advisor: Charles W. Bostian)
18. Matthew Curtis-Maury, *Improving the Efficiency of Parallel Applications on Multithreaded and Multicore Systems*, Computer Science, Virginia Tech, 2008. (Advisor: Dimitris Nikolopoulos)
19. Filip Blagojevic, *Scheduling on Asymmetric Parallel Architectures*, Computer Science, Virginia Tech, 2008. (Advisor: Dimitris Nikolopoulos)
20. Xuan Zheng, *CHEETAH: Circuit-Switched End-to-End Transport Architecture*, Electrical & Computer Engineering, University of Virginia, 2003. (Advisor: Malathi Veeraraghavan)

M.S. Committees

1. David Kindel, *Reducing Subthreshold Leakage Power Through Hybrid MOSFET-NEMS Power Gating*, 2016. (Advisor: Leyla Nazhandali)
2. Ramakrishna Bijanapalli, *Enabling the use of Heterogeneous Computing for Bioinformatics*, Electrical & Computer Engineering, 2014. (Advisor: Peter Athanas)
3. David Uliana, *FPGA-Based Accelerator Development for Non-Engineers*, Electrical & Computer Engineering, 2013. (Advisor: Peter Athanas)
4. Ramakrishna Bijanapalli Chakri, *Enabling the Use of Heterogeneous Computing for Bioinformatics*, Electrical & Computer Engineering, 2013. (Advisor: Peter Athanas)
5. Aditya Dhoke, *On Partial Aborts and Reducing Validation Time in Fault-Tolerant Distributed Transactional Memory*, Computer Science, 2013. (Advisors: Binoy Ravindran and Eli Tilevich)
6. Iccha Sethi, *Clinician Decision Support (CDS) Dashboard: Extracting Value from Electronic Medical Records*, Computer Science, 2012. (Advisor: Harold "Skip" Garner)
7. Karl Pereira, *Characterization of FPGA-based High Performance Computers*, Electrical & Computer Engineering, 2011. (Advisor: Peter Athanas)
8. John Shalf, *Advanced System-Scale and Chip-Scale Interconnection Networks for Ultrascale Systems*, Electrical & Computer Engineering, 2010. (Advisor: Peter Athanas)
9. Sean Ponce, *Towards Algorithm Transformation for Temporal Data Mining on GPU*, Computer Science, Virginia Tech, 2009. (Advisor: Yong Cao)
10. Christian Tergino, *Cryptographic Algorithms on Low-Power Devices*, Electrical & Computer Engineering, Virginia Tech, 2009. (Advisor: Patrick Schaumont)
11. Pavan Konaki, *An Exploration of Hybrid Hard Disk Designs Using an Extensible Simulator*, Computer Science, Virginia Tech, 2008. (Advisor: Ali R. Butt)
12. Ankur Shah, *Improving the Efficiency of Parallel Applications on Multithreaded and Multicore Systems*, Computer Science, Virginia Tech, 2008. (Advisor: Dimitris Nikolopoulos)

AWARDS & RECOGNITION

1. Innovative Paper Award for *C to D-Wave: A High-Level C Compilation Framework for Quantum Annealers*, IEEE High-Performance Extreme Computing (HPEC) Conference, September 2019.
2. Student Innovation Award for *Fast Stochastic Block Partitioning via Sampling*, MIT/Amazon/IEEE Graph Challenge at the IEEE High-Performance Extreme Computing (HPEC) Conference, September 2019.
3. Best Student Poster for *I/O Bottleneck Investigation for Deep-Learning Systems*, 47th International Conference on Parallel Processing, November 2018.

4. "Spotlight on Transactions" Recognition in *Evolving MPI+X Towards Exascale,* IEEE Computer, August 2016. https://www.computer.org/csdl/magazine/co/2016/08/mco2016080010/13rRUwx1xMl
5. SCHEV Outstanding Faculty Award for the Commonwealth of Virginia, February 2014.
6. Best Paper Award for *Towards Energy-Proportional Computing for Enterprise-Class Server Workloads*, 4th ACM International Conference on Performance Engineering, April 2013.
7. "Best 20 Papers over the Past 20 Years" of the IEEE/ACM International Symposium on High-Performance Parallel & Distributed Computing (HPDC) for *MOON: MapReduce On Opportunistic eNvironments*, April 2012.
8. Elizabeth & James E. Turner Faculty Fellowship, March 2012 – April 2020.
9. AMD Research Faculty Fellowship, September 2011.
10. IBM Faculty Award, September 2011.
11. College of Engineering Faculty Fellow, May 2011.
12. CUDA Research Center Designation, May 2011.
13. Compute the Cure Award, NVIDIA Foundation, April 2011.
14. HPCwire's People to Watch List, January 25, 2011. http://www.hpcwire.com/specialfeatures/people_to_watch_2011.
15. ACM Distinguished Scientist, November 2010.
16. IBM Faculty Award, August 2010.
17. NVIDIA Professor Partnership Award, November 2008.
18. 1st Place, Southeastern Universities Research Association (SURA) "Intellectual Property to Market" (IP2M) Competition, October 2008.
19. IBM Faculty Award, August 2008.
20. Distinguished Paper Award for *Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer,* 23rd International Supercomputing Conference, June 2008.
21. International Storage Challenge Award for *ParaMEDIC: Parallel Metadata Environment for Distributed I/O and Computing*, ACM/IEEE SC|07: The International Conference on High Performance Computing, Networking, Storage, and Analysis, November 2007.
22. IBM Faculty Award, August 2007.
23. Best Paper Finalist for *Parallel Genomic Sequence-Searching on an Ad-Hoc Grid: Experiences, Lessons Learned, and Implications*, ACM/IEEE SC|06: The International Conference on High-Performance Computing, Networking, Storage, and Analysis, November 2006.
24. Distinguished Copyright Award, Los Alamos National Laboratory, May 2006.
25. Award for Excellence in Technology Transfer for *10-Gigabit Ethernet Adapter: Speed Really Changes Everything*, Presented by the Federal Laboratory Consortium for Technology Transfer, May 2005.
26. Award for Excellence in Technology Transfer for *Green Destiny & mpiBLAST: Hardware & Software for Super-Efficient Supercomputing*, Presented by the Federal Laboratory Consortium for Technology Transfer, May 2005.
27. R&D 100 Award for *10-Gigabit Ethernet Adapter: Speed Really Changes Everything*, October 2004. Joint entry with Intel Corp.
28. R&D 100 Award for *mpiBLAST: A High-Speed Software Catalyst for Genetic Research*, October 2004.
29. Innovative Supercomputer Architecture Award for *Green Destiny and Its Evolving Parts*, International Supercomputer Conference, June 24, 2004.
30. IEEE Outstanding Young Engineer of the Year, Achievement Award, IEEE Los Alamos Section, May 18, 2004.
31. HPCwire's Top People & Organizations to Watch List, March 18, 2004. http://www.tgc.com/hpcwire/features/topwatch04.html.
32. On the Road to a Gigabit Award: Partnership Category for "Ultralight Partnership," Sponsored by Corporation of Education Network Initiatives in California (CENIC) and California Institute for Telecommunications and Information Technology, Cal-(IT)2, March 15, 2004.
33. Asian-American Engineer of the Year Award, *National Engineers Week*, February 2004.
34. Sustained Bandwidth Award (a.k.a. "Moore's Law Move Over!" Award) for "Bandwidth Lust: Distributed Particle Physics Analysis Using Ultra High-Speed TCP on the Grid," SC2003 Bandwidth Challenge, November 2003. http://www.gridtoday.com/breaking/982.html.
35. R&D 100 Award for *Green Destiny: A 240-Processor Supercomputer in a Telephone Booth*, October 2003.
36. Best Paper Award for *CHEETAH: Circuit-Switched High-Speed End-to-End Transport ArcHitecture*, SPIE/IEEE Optical Networking and Computer Communications Conference, October 2003.
37. Distinguished Performance Award, September 2003.
38. Distinguished Mentor Performance Award, August 2003.
39. Achievement Award for the *Internet2 Land Speed Record*, Los Alamos Awards Program, July 2003.
40. Internet2 Land Speed Record in Guinness World Book of Records, Achieved February 2003, Listed Electronically July 2003. http://www.guinnessworldrecords.com/index.asp?id=58445.

41. Best Paper: Applications Track for "The Design, Implementation, and Evaluation of mpiBLAST," 4th International Conference on Linux Clusters (in conjunction with ClusterWorld Conference & Expo), June 2003.
42. On the Road to a Gigabit Award: Biggest, Fastest in the West for "High-Performance Trans-Atlantic Network Testbed," Sponsored by Corporation of Education Network Initiatives in California (CENIC) and California Institute for Telecommunications and Information Technology, Cal-(IT)2, May 7, 2003.
43. Internet2 Land Speed Record, 2.38 Gbps single-stream TCP/IP over a WAN between Sunnyvale, California and Geneva, Switzerland. Achieved: February 27, 2003. Certified: March 27, 2003. Awarded formally at Internet2 Member Meeting: April 11, 2003. (Note: This achievement also served as the multi-stream Internet2 Land Speed Record. More formally, 23,888,060,000,000,000 meters-bits/second.)
44. Achievement Award for *Green Destiny*, Los Alamos Awards Program, July 2002.
45. Certificate of Appreciation, Women's Career Development Mentoring Award, 2000.
46. Outstanding Mentor Award, 2000.
47. International Who's Who in Information Technology, 1998.
48. Senior Referee Designation, IEEE Computer Society, 1996. (See December 1996 issue of IEEE Computer)
49. Conference Travel Grant Award, Fall 1994.
50. Conference Travel Grant Award, Spring 1994.
51. Best Paper Award, 10th Annual International Conference on Advanced Science and Technology, 1994.
52. Outstanding Teaching Assistant Award, 1991.
53. The Pennsylvania State University Dean's Fellowship, 1988-1989.
54. Lance Stafford Larson Award, IEEE Computer Society, 1988.
55. Student Marshal (Magna Cum Laude) in Computer Engineering, 1988.
56. Best Student Paper, IEEE Pennsylvania Beta Chapter, 1988.
57. National Finalist in the Clara Wells Piano Competition, 1983.

SELECTED MEDIA COVERAGE

"Microsoft Cloud: Wu Feng Changing the World," TV Commercial Aired Worldwide, March 2016 – September 2016. https://www.youtube.com/watch?v=hxOcrNaQibE

"Microsoft Cloud: Empowering Cancer Research," TV Commercial Aired Worldwide, February 2015 – August 2015. https://www.ispot.tv/ad/7xoX/microsoft-cloud-empowering-cancer-research
https://vimeo.com/271732253

"Microsoft Cloud: Empowering Cancer Research," 60-Second TV Commercial, February 2015 – August 2015. https://www.dropbox.com/s/p3adn57bmua4zyq/CloudVT_HD_60_MSFT_1080_webmix.mov.mp4?dl=0

"Microsoft Cloud: Accelerating Cancer Research," 90-Second Infomercial, February 2015 – August 2015. https://www.dropbox.com/s/da4pxrys0pe2hx4/CloudVT_HD_90_MSFT_1080_webmix.mov.mp4?dl=0

"On the Verge of a Breakthrough: How the Cloud Helps Your Business Find Insights in Big Data," *TechNet Blogs*, May 15, 2014. http://blogs.technet.com/b/work/archive/2014/05/15/on-the-verge-of-a-breakthrough-how-the-cloud-helps-your-business-find-insights-in-big-data.aspx.

"Kids Tech University Takes Hard Look at Computer Future," May 4, 2014. http://www.swvatoday.com/news/washington_county/article_ba93de06-cfdb-11e3-95e2-0017a43b2370.html.

"9 Windows Azure HDInsight Big Data Success Stories," *CIO*, February 6, 2014. http://www.cio.com/slideshow/detail/139273/9-Windows-Azure-HDInsight-Big-Data-Success-Stories#slide3.

"State Council of Higher Education for Virginia honors Wu Feng with Outstanding Faculty Award," *VT News*, January 24, 2014. http://www.vtnews.vt.edu/articles/2014/01/012414-engineering-wufengstateaward.html.

"Big Data vs. Big Compute," *HPCwire,* December 10, 2013 http://www.hpcwire.com/2013/12/10/big-data-versus-big-compute/.

"New Computing Model Could Lead to Quicker Advancements in Medical Research," *HPCwire,* November 6, 2013 http://cacm.acm.org/news/169480-new-computing-model-could-lead-to-quicker-advancements-in-medical-research/fulltext.

"Tackling the Big Data Issue Behind Connected Health," *Qmed*, November 2013. http://www.qmed.com/mpmn/medtechpulse/tackling-big-data-issue-behind-connected-health.

"Computing Model Could Lead to Quicker Advancements in Medical Research," *HPCwire*, October 2013.

http://www.hpcwire.com/hpcwire/2013-11-04/computing_model_could_lead_to_quicker_advancements_in_medical_research.html.

"Microsoft's Big Data Service Available After a Year in Preview," *CIO*, October 2013. http://www.cio.com/article/742127/Microsoft_s_Big_Data_Service_Available_After_a_Year_in_Preview.

"Creator Of Most Efficient Supercomputer To Start Working On Drone-Bugs," *Popular Science*, February 2013. http://www.popsci.com/technology/article/2013-02/creator-super-efficient-computers-will-get-paid-start-working-drone-bugs.

"NSF Announces Interagency Progress on Administration's Big Data Initiative," *NSF*, October 2012. http://www.nsf.gov/news/news_summ.jsp?cntn_id=125610.

"The Unified Theory of Wu," *Communications of the ACM*, July 2012. http://cacm.acm.org/careers/153288-the-nified-theory-of-wu/fulltext.

"Project Moon: One Small Step for a PC, One Giant Leap for Data," *Wired*, May 2012. http://www.wired.com/wiredenterprise/2012/05/project_moon.

"Computer Scientists' Paper Among the Best in the History of High-Performance Distributed Computing," *VT News*, May 2012. http://www.vtnews.vt.edu/articles/2012/05/050312-engineering-cspaperbest.html.

"People in the News: Wu-chun Feng, Hugh Rienhoff, Charles Abdalian, Adam Rosenberg, and Others," *GenomeWeb*, April 2012. http://www.genomeweb.com/informatics/people-news-wu-chun-feng-hugh-rienhoff-charles-abdalian-adam-rosenberg-and-other.

"The ARRA Report: HokieSpeed," *Science360 News Service*, March 28, 2012. http://news.science360.gov/obj/video/93879471-2dce-4a46-a503-ed051b22ec2f.

"VT Nears Completion of HokieSpeed," *Engadget*, December 2011. http://www.engadget.com/2011/12/23/vt-nears-completion-of-hokiespeed-worlds-96th-most-powerful-su/.

"Prof Promises Supercomputer on Every Desktop," *Wired*, December 2011. http://www.wired.com/wiredenterprise/2011/12/vt-supercomputer/.

"Green500 Shows Continuing Trend Toward Environmentally Friendly Supercomputers," *HPCwire*, November 2011. http://www.hpcwire.com/hpcwire/2011-11-28/green500_shows_continuing_trend_toward_environmentally_friendly_supercomputers.html.

"The Next IT Revolution?: Cloud Computing Opportunities and Challenges," *Congressional Testimony to the Committee on Science, Space and Technology*, September 2011. http://science.house.gov/sites/republicans.science.house.gov/files/documents/hearings/092111_Reed.pdf. http://science.house.gov/hearing/technology-and-innovation-subcommittee-hearing-cloud-computing.

"Wu Feng Awarded AMD Fellowship for Work in Heterogeneous Computing," *HPCwire*, August 2011. http://www.hpcwire.com/hpcwire/2011-08-25/wu_feng_awarded_amd_fellowship_for_work_in_heterogeneous_computing.html.

"IBM Faculty Award To Benefit Elementary-level School Children," *The Virginia Engineer*, May 2011. http://vaeng.com/news/ibm-award-to-benefit-elementary-level-school-children.

"Virginia Tech CUDA Research Center Summary," *NVIDIA*, May 2011. http://research.nvidia.com/content/VATech-crc-summary, http://research.nvidia.com/partnerships.

"Virginia Tech Researcher Leads Effort to Compute the Cure with GPUs," *insideHPC*, May 2011. http://insidehpc.com/2011/05/04/virginia-tech-researcher-leads-effort-to-compute-the-cure-with-gpus/.

"The Sky Is No Limit: 13 Research Teams Compute in the Clouds," *NSF*, April 2011. http://www.nsf.gov/news/news_summ.jsp?cntn_id=119248.

"CUDA Spotlight: Compute the Cure," *CUDA Spotlight*, April 2011. http://developer.nvidia.com/cuda-spotlight-compute-cure.

"HPCwire Top People to Watch: Wu-chun Feng," *HPCwire*, January 2011. http://www.hpcwire.com/specialfeatures/people_to_watch_2011/Dr-Wu-chun-Feng-113294639.html.

"The Greening of HPC – A Conversation with Professor Wu Feng," *HPCwire*, November 2010. http://www.hpcwire.com/specialfeatures/sc10/features/The-Greening-of-HPC-A-Conversation-with-Professor-Wu-Feng-108983219.html.

"Accelerator Supercomputers Dominate Newest Green500 List's Top 10," *HPCwire*, November 2010.

http://www.hpcwire.com/offthewire/Accelerator-Supercomputers-Dominate-Newest-Green500-Lists-Top-10-109034899.html.

"ACM Names 47 Distinguished Members for Computing Innovations," *ACM*, November 16, 2010. http://www.acm.org/press-room/news-releases/2010/distinguished-2010.

"CPUs+GPUs=HokieSpeed," *Dr. Dobbs Journal*, August 13, 2010. http://www.ddj.com/go-parallel/article/showArticle.jhtml?articleID=226700214.

"Virginia Tech to Build New Supercomputer," *GenomeWeb*, August 2010. http://www.genomeweb.com/blog/virginia-tech-build-new-supercomputer.

"HokieSpeed Supercomputing Will Tap Both CPUs and GPUs," *Scientific Computing*, August 2010. http://www.scientificcomputing.com/news-IN-HokieSpeed-Supercomputing-will-tap-both-CPUs-and-GPUs-080910.aspx.

"Sharing the Wealth of Data," *Scientific American WorldView (SAW)*, 2010. http://www.saworldview.com/article/sharing-the-wealth-of-data.

"The Green500 List: Accelerators Raising the 'Fuel Efficiency' of Supercomputers," *HPCwire*, June 30, 2010. http://www.hpcwire.com/offthewire/97482494.html.

"mpiBLAST: Open-Source Parallel BLAST," *Dr. Dobb's Journal*, April 27, 2010. http://www.drdobbs.com/high-performance-computing/224600668;jsessionid=HYZ2HW53NKGXLQE1GHPSKHWATMY32JVN.

"High Performance Computing Reveals Missing Genes," *ScienceDaily*, April 14, 2010. http://www.sciencedaily.com/releases/2010/04/100413151911.htm.

"VBI Uses 'Ephemeral Supercomputer' to Reveal Missing Genes," *GenomeWeb*, April 13, 2010. http://www.genomeweb.com/blog/vbi-uses-ephemeral-supercomputer-reveal-missing-genes.

"Lean, Green Machines," *Communications of the ACM*, November 2009. http://cacm.acm.org/news/53018-lean-green-machines/fulltext.

"NVIDIA GPUs Speed Virginia Tech Research," *The NVIDIA Blog*, November 18, 2009. http://blogs.nvidia.com/ntersect/2009/11/nvidia-gpus-speed-virginia-tech-research.html.

"New Software Could Smooth Supercomputing Speed Bumps," *Scientific American*, October 2009. http://www.scientificamerican.com/article.cfm?id=opencl-smooths-supercomputing.

"Energy Efficiency Improvements in HPC," *insideHPC*, August 3, 2009. http://insidehpc.com/2009/08/03/energy-efficiency-improvements-in-hpc/.

"Green HPC Podcast Episode 1: Sifting Through The Hype," *insideHPC*, June 30, 2009. http://insidehpc.com/green-hpc-podcast-series/green-hpc-podcast-episode-1-sifting-through-the-hype/ http://insidehpc.com/green-hpc-podcast-series/ for the entire Green HPC Podcast series.

"Virginia Tech's Feng Preps mpiBLAST Other Search Tools for Emerging IT Architectures," *Bioinform*, August 15, 2008. http://www.genomeweb.com/informatics/virginia-tech%E2%80%99s-feng-preps-mpiblast-other-search-tools-emerging-it-architectures.

"The Greening of HPC," *HPCwire*, June 26, 2008. http://www.hpcwire.com/features/The_Greening_of_HPC.html.

"Va. Tech Researchers Part of Team Receiving Supercomputer Honors," *HPCwire*, June 16, 2008. http://www.hpcwire.com/offthewire/Va_Tech_Researchers_Part_of_Team_Receiving_Supercomputing_Honors.html.

"ParaMEDIC Enables Worldwide Supercomputer for Bioinformatics," *International Science Grid This Week*, December 12, 2007. http://www.isgtw.org/?pid=1000811.

"Argonne, Virginia Tech Win Storage Challenge," *HPCwire*, November 26, 2007. http://www.hpcwire.com/hpc/1909146.html.

"Va Tech's Green500 List Puts Supercomputing on a Diet," *HPCwire*, November 7, 2007. http://www.hpcwire.com/hpc/1881238.html.

"Green HPC Has Arrived," *Genome Technology*, November 2007. http://www.genome-technology.com/issues/2_9/webreprints/143164-1.html.

"Hippie Geeks Threaten Supercomputing Masculinity," *The Register*, November 17, 2006. http://www.theregister.co.uk/2006/11/17/green500_supercomputer/.

"New Frontiers in High-End Computing," *Government Computer News*, February 2006. http://www.gcn.com/print/25_4/38248-1.html.

"The Megascale-Processor Machine: Does It Compute?," *Genome Technology: Inside Integrated Biology*, January / February 2006. http://www.genome-technology.com. No electronic version available.

"Los Alamos Gets EnergyFit," *HPCwire*, Vol. 14, No. 35, September 2, 2005. http://news.taborcommunications.com/msgget.jsp?mid=464670&xsl=story.xsl.

"Seven Cheers for Technology," *The Scientist*, Vol. 19, No. 16, August 29, 2005. http://www.the-scientist.com/2005/08/29/21/1.

"EnergyFit Raises the Gas Mileage of Your High-Octane Computer," *LinuxHPC.org*, August 23, 2005. http://www.linuxhpc.org/mobile/news.php?Rid=05/08/23/9738054.

"A Winning Combination," Power Profile Section, *R&D Magazine*, February 2005. http://www.rdmag.com/ShowPR.aspx?PUBCODE=014&ACCT=1400000100&ISSUE=0502&RELTYPE=FE&PRODCODE=00000000&PRODLETT=AF.

"Parallel BLAST: Chopping the Database," *Genome Technology: Inside Integrated Biology*, January / February 2005. http://www.genome-technology.com. No electronic version available.

"The Wizard of Los Alamos," *TechComm Journal: The National Journal of Technology Commercialization*, Dec. 2004 / Jan. 2005. http://www.innovation-america.org/archive.php?articleID=321.

"10GbE: Enabling High-Performance Storage, Networking," *HPCwire*, Vol. 13, No. 49, December 10, 2004. http://www.tgc.com/hpcwire/hpcwireWWW/04/1210/108931.html.

"Unexpected Workstation Revival Under Way," *Bio-IT World*, September 13, 2004. http://www.bio-itworld.com/news/091304_report6017.html.

"Start-Up Introduces a Technology First: The Personal Supercomputer," *LinuxWorld*, September 7, 2004. http://www.linuxbusinessweek.com/story/46242.htm. http://www.linuxworld.com/story/46242.htm.

"New Workstations Deliver Computational Muscle," *Bio-IT World*, August 30, 2004. http://www.bio-itworld.com/news/083004_report5927.html.

"Bandwidth Challenge Teams Push Performance Envelope at SC2003," *GRIDtoday*, Vol. 2, No. 49, December 8, 2003. http://www.gridtoday.com/03/1208/102358.html. http://www.gridtoday.com/03/1208/102380.html.

"Bandwidth Challenge Teams Push Performance Envelope at ACM's SC2003 Conference," *ACM TechNews*, December 1, 2003. http://www.acm.org/technews/articles/2003-5/1201m.html#item4.

"Bandwidth Challenge Teams Push Networking Performance Envelope at SC2003 Conference – Sustained 23 Gigabits Per Second Sets New Record," *Silicon Valley Biz Ink*, December 1, 2003. http://www.prnewswire.com/cgi-bin/stories.pl?ACCT=SVBIZINK2.story&STORY=/www/story/11-25-2003/0002065579&EDATE=TUE+Nov+25+2003,+07:41+PM.

"Moore's Law for Power Consumption: Situation Critical," *The Inquirer*, November 21, 2003. http://www.theinquirer.net/default.aspx?article=12798.

"Interview with Wu-chun Feng, LANL & OSU," *HPCwire,* Vol. 10, No. 3, November 20, 2003. http://www.tgc.com/hpcwire/hpcwireWWW/03/1120/106520.html.

"Efficient Supercomputing with Green Destiny," *slashdot.org,* November 19, 2003. http://slashdot.org/articles/03/11/19/2156246.shtml?tid=126&tid=161&tid=187.

"Los Alamos National Lab Smashes Networking Records with Intel's 10-Gigabit Ethernet Server Adapter," November 1, 2003. http://www.intel.com/network/connectivity/case_studies/16832_LosAlamos_CS_r03.pdf.

"Green Destiny: A 'Cool' 240-Node Supercomputer in a Telephone Booth," *BBC News*, August 2003.

"Los Alamos Sets Internet Speed Mark in Guinness Book," *GRIDtoday*, Vol. 2, No. 31, August 4, 2003. http://www.gridtoday.com/03/0804/101764.htmll. Also *Quote of the Week* for *GRIDtoday:* "While average network speeds double every year, processor speeds are doubling only every 18 months. These trends mark the beginning of a revolution in the way we do computing."

"Los Alamos Hits The Pipe In Record Time," *IEEE Spectrum Online*, July 31, 2003. http://www.spectrum.ieee.org/WEBONLY/newslog/news07-31-03.html.

"Internet Speed Record in Guinness World Records Book," *HPCwire*, Vol. 12, No. 29, July 25, 2003. http://www.tgc.com/hpcwire/hpcwireWWW/03/0725/105593.html.

"Los Alamos TCP Pipe Hits 8.5 Gbps," *SpaceDaily*, July 25, 2003.

http://www.spacedaily.com/news/internet-03u.html.

"Los Alamos Sets Internet Speed Mark in Guinness Book," *GRIDtoday: Breaking News*, July 24, 2003. http://www.gridtoday.com/breaking/781.html.

"Internet2 Land Speed Record," Listed electronically in the *Guinness World Records*, July 7, 2003. http://www.guinnessworldrecords.com/index.asp?id=58445.

"Los Alamos Lends Open-Source Hand to Life Sciences," *The Register*, June 29, 2003. http://www.theregister.com/content/61/31471.html.

"LANL Researchers Outfit the 'Toyota Camry' of Supercomputing for Bioinformatics Tasks," *BioInform*, February 3, 2003. Electronic version only available via subscription at http://www.bioinform.com.

"Not Your Average Supercomputer," *Communications of the ACM*, Vol. 45, No. 8, pg. 9, August 2002. Electronic version only available via subscription at http://www.acm.org.

"Dr. Dobb's News & Views: 'Green Destiny' Runs Cool," *Dr. Dobb's Journal*, August 2002. Electronic version only available via subscription at http://www.ddj.com.

"At Los Alamos, Two Visions of Supercomputing," *The New York Times*, June 25, 2002. http://www.nytimes.com/2002/06/25/science/physical/25COMP.html.

"Green Destiny," *Boomernomics*, June 6, 2002. http://www.boomernomics.net/VitaminB/jun0602.htm.

"Supercomputing Coming to a Closet Near You?" *HPCwire*, Vol. 11, No. 21, May 31, 2002. http://www.tgc.com/hpc-bin/artread.pl?direction=Current&articlenumber=102814.

"Lean, 'Green' Computer Thrives: Machine Runs in Harsh Conditions," Albuquerque Journal, May 28, 2002. No web link.

"Supercomputing Coming to a Closet Near You?" *PCWorld.com*, May 27, 2002. http://www.pcworld.com/news/article/0,aid,100544,00.asp.

"National Lab Debuts Supercomputer Project," *SAP Info*, May 22, 2002. http://www.sapinfo.net/public/en/news.php4/Category-28813c6138d029be8/page/0/article/Article-114783ceb6f2789f99/en.

"Bell, Torvalds Usher Next Wave of Supercomputing," *CNN.com*, May 21, 2002. http://www.cnn.com/2002/TECH/industry/05/21/supercomputing.future.idg/index.html.

"The Green Destiny: A Cluster for the Rest of Us," *Geek.com*, May 21, 2002. http://www.geek.com/news/geeknews/2002may/gee20020521011824.htm.

"Green Destiny Draws Cheers and Jeers," *ITworld.com*, May 21, 2002. http://www.itworld.com/Tech/3494/020521greendestiny.

"Linus Launches Compact Supercomputer: Meet Green Destiny …," *Silicon.com*, May 20, 2002. http://www.silicon.com/bin/bladerunner?30REQEVENT=&REQAUTH=21046&14001REQSUB=REQINT1=53452.

Secondary Stories

"Supercomputing's Super Energy Needs and What to Do About Them," *Communications of the ACM*, September 24, 2015. http://www.wset.com/story/24782288/local-educators-honored-by-gov-mcauliffe.

"Local Educators Honored by Gov. McAuliffe," *ABC News (WSET – Channel 13)*, February 20, 2014. http://www.wset.com/story/24782288/local-educators-honored-by-gov-mcauliffe.

"In Race for Fastest Supercomputer, China Outpaces U.S.," *Newsweek*, November 28, 2011. http://www.thedailybeast.com/newsweek/2011/11/27/in-race-for-fastest-supercomputer-china-outpaces-u-s.html.

"China Wrests Supercomputer Title From U.S.," *The New York Times*, October 28, 2010. http://www.nytimes.com/2010/10/28/technology/28compute.html.

"Geeks Fight the Smelter with Embedded Processor-Based Box," *The Register*, February 2, 2008. http://www.theregister.co.uk/2008/02/02/horst_simon_cloud_computer/.

"A Change in Management at Los Alamos National Laboratory," *The Economist*, June 15, 2006. http://economist.com/science/displaystory.cfm?story_id=7055808.

"When Data Center Lose Their Cool," *Government Computer News*, May 15, 2006. http://www.gcn.com/print/25_12/40759-1.html.

"Seven Cheers for Technology," *The Scientist*, Vol. 19, No. 16, August 29, 2005. http://www.the-scientist.com/2005/08/29/21/1.

"Building a Bioinformatics Supercomputing Cluster," *Linux Journal*, April 6, 2005. http://www.linuxjournal.com/article/7936.

"The Future of Supercomputing is Its Green Destiny," *The Inquirer,* June 27, 2004.
http://www.the-inquirer.com/?article=16847.

"Five Feds Among HPCwire's People to Keep an Eye On," *Government Computer News*, Vol. 23, No. 8, April 19, 2004.
http://www.gcn.com/23_8/community/25618-1.html.

HONOR SOCIETIES

Phi Kappa Phi Honor Society
Tau Beta Pi Engineering Honor Society
Eta Kappa Nu Electrical & Computer Engineering Honor Society
Golden Key Honor Society
National Society of Professional Engineers

HOBBIES

Sports: cycling, ultimate frisbee, weightlifting, snowboarding, skiing (x-c & downhill), and running.
Music: piano – performance and composition.