

Thomas Richard William Scogland
1262 Cochran Dr.; Tracy, California 95377
scogland1@llnl.gov • (224) 698-7775 • <http://tom.scogland.com>

Education:

December 2014 Ph.D. Computer Science, Virginia Tech
 Thesis title: “Runtime Adaptation for Autonomic Heterogeneous Computing”
July 2012 Computer Science, Master of Science
 Graduate Index: 3.87/4.00
May 2007 B.S. Computer Science, Purdue University
 Graduate Index: 3.29/4.00

Research Interests:

- High-performance and heterogeneous computing
- Parallel languages, tools and runtimes
- Scheduling/Load-balancing
- Energy efficient high-performance computing

Employment:

- **Lawrence Livermore National Laboratory** September 2014 - Present
Center for Applied Scientific Computation, Postdoctoral scholar: Conducting research in programming models for next-generation architectures, scheduling and resource management at levels from the node to a complete site, and energy-aware scheduling.
- **Advanced Micro Devices (AMD)** June 2012 - August 2012
Summer Intern: Worked on development of the Heterogeneous Systems Architecture (HSA) platform runtime and interfaces for logically targeting multiple devices with a single workload and platform-level load-balancing.
- **Lawrence Livermore National Laboratory** May 2011 - August 2011
Center for Applied Scientific Computation, Summer Intern: Research included the design, development and analysis of a heterogeneous task scheduler to allow worksharing of parallel loops across two heterogeneous devices with distinct memory spaces.
- **EnergyWare** January 2009 - October 2009
Independent Contractor and Software Engineer: Tested and extended a dynamic voltage and frequency daemon for Linux and developed an OS X kernel module to support a port to that platform. Developed a new version of the system supporting frequency management on Android devices lacking kernel support.
- **Argonne National Laboratory** May 2008 - August 2008
Mathematics and Computer Science, Summer Intern: Helped design and implement the first version of the Hydra job launcher for MPICH2, specifically the SSH and SLURM launching engines. Hydra has since become the standard job launcher for MPICH2 and its derivatives.
- **Virginia Tech CS Department** January 2008 - September 2014
Student, M.S. and Ph.D.: Primary research focus targets automated runtime scheduling, load balancing and adaptation for heterogeneous environments, including process affinity management for MPI, worksharing across heterogeneous devices for OpenMP, and the design of policies for these systems. Other research activities include optimization analysis and evaluation for heterogeneous platforms, energy efficient computing and highly scalable parallel algorithm and data structure design.
- **Echostar Communications** July 2007 - August 2007
Software Engineering Intern: Redesigned a database and analysis platform aggregating data retrieved from set-top-boxes (STBs), reducing full-database query time from 22 hours to 15 minutes while shifting from licensed database tools to open source. Additionally developed a tool for plotting aggregate query results geographically across the US, providing the first fine-grain visualization of satellite signal strength, among other metrics.

Publications:

- [1] **Thomas R W Scogland**, Wu-chun Feng, Barry Rountree, and Bronis R de Supinski. CoreTSAR: Core Task-Size Adapting Runtime. *IEEE Transactions on Parallel and Distributed Systems*, October 2014, **preprint**.

- [2] **Thomas R W Scogland** and Wu-chun Feng. Design and Evaluation of Scalable Concurrent Queues for Many-Core Architectures. In *ACM/SPEC International Conference on Performance Engineering (ICPE)*, January 2015.
- [3] **Thomas R W Scogland**, Barry Rountree, Wu-chun Feng, and Bronis R de Supinski. CoreTSAR: Adaptive Worksharing for Heterogeneous Systems. In *International Supercomputing Conference*, Leipzig, June 2014.
- [4] **Thomas R W Scogland**, Craig P Steffen, Torsten Wilde, Florent Parent, Susan Coghlan, Natalie Bates, Wu-chun Feng, and Erich Strohmaier. A Power-Measurement Methodology for Large-Scale, High-Performance Computing. In *ACM/SPEC International Conference on Performance Engineering (ICPE)*, March 2014.
- [5] Brent P Pickering, Jackson W Charles, **Thomas R W Scogland**, Wu-chun Feng, and Christopher J Roy. Directive-Based GPU Programming for Computational Fluid Dynamics. In *52nd Aerospace Sciences Meeting*, Reston, Virginia, January 2014. American Institute of Aeronautics and Astronautics.
- [6] Konstantinos Krommydas, **Thomas R W Scogland**, and Wu-chun Feng. On the Programmability and Performance of Heterogeneous Platforms. In *International Conference on Parallel and Distributed Systems*, Gungam, October 2013.
- [7] Wu-chun Feng, Kirk Cameron, and **Thomas R W Scogland**. The Green500 List: A look back to look forward. In Jeffrey S Vetter, editor, *Contemporary High Performance Computing From Petascale toward Exascale*, pages 31–41. Chapman and Hall/CRC, January 2013.
- [8] Balaji Subramaniam, Winston Saunders, **Thomas R W Scogland**, and Wu-chun Feng. Trends in energy-efficient computing: A perspective from the Green500. In *International Green Computing Conference (IGCC)*, pages 1–8, 2013.
- [9] **Thomas R W Scogland**, Barry Rountree, Wu-chun Feng, and Bronis R de Supinski. Heterogeneous Task Scheduling for Accelerated OpenMP. In *International Parallel and Distributed Processing Symposium*, pages 144–155. IEEE Computer Society, May 2012.
- [10] Wu-chun Feng, Heshan Lin, **Thomas R W Scogland**, and Jing Zhang. OpenCL and the 13 Dwarfs: a Work in Progress. In *Proceedings of the third joint WOSP/SIPEW international conference on Performance Engineering*, Boston, MA, April 2012. ACM.
- [11] **Thomas R W Scogland**, Balaji Subramaniam, and Wu-chun Feng. The Green500 List: Escapades to Exascale. *Computer Science-Research and Development*, pages 1–9, 2012.
- [12] **Thomas R W Scogland**, Balaji Subramaniam, and Wu-chun Feng. Emerging Trends on the Evolving Green500: Year Three. In *Workshop on High-Performance, Power-Aware Computing (IPDPSW: HPPAC)*, pages 822–828. IEEE Computer Society, May 2011.
- [13] Mayank Daga, Wu-chun Feng, and **Thomas R W Scogland**. Towards accelerating molecular modeling via multi-scale approximation on a GPU. In *International Conference on Computational Advances in Bio and Medical Sciences (ICCABS)*, pages 75–80, 2011.
- [14] Mayank Daga, **Thomas R W Scogland**, and Wu-chun Feng. Architecture-Aware Mapping and Optimization on a 1600-Core GPU. In *International Conference on Parallel and Distributed Systems*, pages 316–323, Tainan, Taiwan, 2011. IEEE Computer Society.
- [15] Marwa Elteir, Heshan Lin, Wu-chun Feng, and **Thomas R W Scogland**. StreamMR: An Optimized MapReduce Framework for AMD GPUs. In *International Conference on Parallel and Distributed Systems*, pages 364–371, 2011.
- [16] **Thomas R W Scogland**, Heshan Lin, and Wu-chun Feng. A First Look at Integrated GPUs for Green High-performance Computing. *Computer Science-Research and Development*, 25(3-4):125–134, August 2010.
- [17] Ramu Anandkrishnan, **Thomas R W Scogland**, Andrew T Fenley, John C Gordon, Wu-chun Feng, and Alexey V Onufriev. Accelerating Electrostatic Surface Potential Calculation with Multi-scale Approximation on Graphics Processing Units. *Journal of Molecular Graphics and Modelling*, 28(8):904–910, June 2010.
- [18] Mark K Gardner, Adam Herr, David Mazary, Heshan Lin, **Thomas R W Scogland**, and Wu-chun Feng. A Flexible and Extensible Framework for Delivering Designer Images for K-12 Pedagogy. In *3rd International Conference on the Virtual Computing Initiative (ICVCI)*, Research Triangle Park, NC, October 2009.
- [19] Wu-chun Feng and **Thomas R W Scogland**. The Green500 List: Year One. In *Workshop on High-Performance, Power-Aware Computing (IPDPSW: HPPAC)*, Rome, Italy, May 2009. IEEE Computer Society.

- [20] Jeremy Archuleta, Yong Cao, **Thomas R W Scogland**, and Wu-chun Feng. Multi-dimensional Characterization of Temporal Data Mining on Graphics Processors. In *International Parallel and Distributed Processing Symposium*, pages 1–12, 2009.
- [21] **Thomas R W Scogland**, Pavan Balaji, Wu-chun Feng, and Ganesh Narayanaswamy. Asymmetric Interactions in Symmetric Multi-Core Systems: Analysis, Enhancements and Evaluation. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing)*, 2008.

Teaching Experience:

Graduate Teaching Assistant (Virginia Tech):

- Heterogeneous Computing (CS6504), Unofficial TA, Fall 2010, 8 students. This position was part of an independent study titled "Research, Development, and Evaluation of Accessible Parallel Computing" built around the development and evaluation of a parallel computing curriculum for an active graduate course. My responsibilities included homework and project development and grading as well as lecturing on subjects including pthreads and OpenCL.
- Professionalism in Computing (CS3604), 1/3 semester Fall 2008, 30 students. I was responsible for grading of assignments.
- Network Architecture and Protocols (CS5565), Spring 2008, 10 students. My duties included creating and grading assignments and grading tests as well as lecturing on sockets and basic Internet protocols such as FTP and SMTP.

Undergraduate Teaching Assistant (Purdue):

- Operating Systems (CS354), Spring 2007, 50 students overall, 10 in my lab section. As an undergraduate TA I was responsible for grading programming assignments and part of each exam, along with two graduate TAs. I was also responsible for leading one of the three lab sections. Each lab consisted of a short lecture combined with an in-lab assignment, which I oversaw.
- Advanced Tablet Platform Applications (CS490T), Fall 2006, 17 students. As an experimental projects course in the style of a seminar, my primary responsibility was the assistance and mentoring of students through their course projects, which included aspects of computer vision, networking, compression, human computer interaction and location awareness.

Other:

- Co-Design for Parallel Computing (CS6504), Spring 2013. Gave a guest lecture on software co-design in computational fluid dynamics.
- Computer Organization / Computer Architecture (CS/ECE504), Spring 2012. Served as guest lecturer on many-core architectures and basic CUDA programming and also created a project to teach the basics of CUDA.

Service:

- Green500 List Technical Leader
- OpenMP Accelerator subcommittee member
- C Parallel EXtensions (CPLEX) committee member
- Co-chair of the first Workshop on Representative Applications (WRAP) 2015
- Program committee member for Workshop on High-Performance Power-Aware Computing (HPPAC) 2015
- Reviewer for ASHES 2013
- Member of the Energy Efficient High Performance Computing Working Group (EEHPCWG) and participating author in developing the Energy Efficient High Performance Computing Working Group power measurement methodology
- Web Chair for ICPP 2012
- Session Chair for HPPAC sessions on Energy Workloads and Power Efficient Hardware
- Session Chair for EnA-HPC Scientific Session - Power Extrapolation, Estimation and Simulation
- Former member of Engineering Projects in Community Service (EPICS) at Purdue
- Reviewer for P2S2 2010
- Student Technical Lead of the MyVICE Project, a program to deliver computer science resources and curricula to rural Virginia, 2010 through 2012

Awards:

- PACT 2014 Student Research Competition and travel grant award
- IPDPS 2014 PhD Forum
- IPDPS 2014 travel grant award
- CCGRID 2014 Doctoral Showcase
- Supercomputing 2012 Doctoral Showcase (Acceptance rate 25.5%)
- ISCR Scholar, Lawrence Livermore National Lab, Summer, 2011
- NSF Conference Travel Grant Award, Fall, 2010
- Department of Defense (DoD) National Defense Science and Engineering Graduate (NDSEG) Fellowship, August 2009 - July 2012