VENKATRAMAN GOVINDARAJU

Amazon Web Services
2100 University Avenue
Palo Alto, CA 94303

Cell: (608) 217-6100
venkatra@cs.wisc.edu
http://cs.wisc.edu/~venkatra

Research Interests

- Architecture, compilers, runtime systems for large scale distributed systems
- Workload characterization, analysis, and application specific optimizations
- Hardware/software co-design and the interaction of hardware, compilers, and operating systems

Education

- **University of Wisconsin**—Madison, WI
  Doctor of Philosophy, Computer Science, August 2008 - August 2014
  - ◇ Ph.D., Computer Science, University of Wisconsin, Madison. August 2014
  - ◇ Dissertation Title: Energy Efficient Computing through Compiler Assisted Dynamic Specialization
- **University of Wisconsin**—Madison, WI
  Master of Science, Computer Science, August 2008 - December 2009
  - ◇ Courses: Advanced Computer Architecture I & II, VLSI Design, Construction of Compilers, Advanced Computer System Modeling, Topics in Database Management Systems
- **Anna University**—Chennai, Tamilnadu, India
  Bachelor of Engineering, Computer Science and Engineering, August 1995 - May 1999
  - ◇ Relevant Courses: Computer Architecture, Operating Systems, Compilers

Work Experience

- **Amazon Web Services**—Palo Alto, CA
  Principal Engineer - AWS Redshift, June 2021 - Present
  - ◇ Working on Amazon Redshift, a cloud data warehouse solution
  - ◇ Technical lead for distributed system and memory management aspects of AWS Redshift
  - ◇ Leading efforts on improving stability, performance and scalability of the system
- **Facebook**—Menlo Park, CA
  Performance and Capacity Engineer, July 2020 - June 2021
  - ◇ Next Generation Query Engine: Worked on Velox, a next generation SQL compliant vectorized query engine, to unify numerous query engines in Facebook. Designed and implemented spilling support, OrderBy operator and Merge Operator. Owned Memory management in Velox
  - ◇ Capacity Planning: Worked on Big Data Planner, a capacity planning service which tries to get the optimized hardware order and data placement based on demand, supply and budget
- **Amazon Web Services**—Palo Alto, CA
  Principal Engineer - AWS Redshift, May 2017 - July 2020
  - ◇ Worked on Amazon Redshift, a cloud data warehouse solution
  - ◇ Memory Management:Implemented multi-process shared memory support and an aggressive coalescing algorithm in jemalloc to support Redshift's shared memory allocation patterns. Integrated the modified jemalloc with Redshift and added monitoring hooks to track usage by allocation type, size and query. Reduced OOMs by 50%, improved performance by 15% and prevented new regressions from slipping into production.

- ◇ Cloud Wide Global Code Pool: Redshift used a local code pool to cache the compiled object file to amortize the cost of query compilation within a cluster. Designed and implemented a new micro service that caches the compiled object files across all the Redshift clusters. Reduced miss rate by $3\times$ and substantially reduced the variability in query execution time. Patent Application Submitted.
- ◇ Improving Network Performance: Improved network performance by implementing out-of-order packet handling, fast retransmits, dynamic send and receive window size by adapting to network conditions using machine learning in the inter-node communication protocol. Increased overall query performance by 10%.
- ◇ Elastic Redshift Cluster: Redshift had static cluster configurations. Designed and implemented a "membership protocol" that allow discovering new nodes and dropping existing nodes without restarting the cluster or affecting any inflight queries. This was essential for delivering elastic cluster resize, hot swapping nodes during maintenance and scaling compute resources elastically depending upon the load.
- ◇ Modernizing Codebase: Redshift was stuck with gcc-4.1.2 from 2007 for its production builds and for query code generation. Ported entire system to use C++17 and upgraded the compiler to gcc-7.3.0. This in turn enabled use of modern and updated libraries and helped deploying multiple features quickly: unload to parquet, geographic information system and allowed the use of SIMD instructions, improving performance by $2\times$. With new compiler, built the address sanitizer and undefined behavior sanitizer pipelines, which has resulted in several critical fixes.
- ◇ Code Generation for Expression: Redshift generated expression evaluation code as a monolith C++ expression. This led to increased compilation time and generated sub-optimal machine code. Revamped code generation for expression by generating simpler C++ expression and compose them to evaluate the full expression. This naturally allowed subexpression elimination and other optimizations to generate optimized code. Improved code compilation time by $3\times$ and improved expression evaluation performance by $2\times$ for large expressions.

- **Oracle Corporation**—Belmont, CA
  Principal Member of Technical Staff - Oracle Labs, July 2015 - May 2017
  - ◇ Researched on ways to use custom hardware accelerator to improve scalability and performance of large scale distributed data processing system
  - ◇ Analyzed large workloads on 1000 node clusters, identified and rectified potential performance and scalability bottlenecks
  - ◇ Worked on RAPID project, a hardware-software co-designed system targeting large scale data management and analysis
  - ◇ Designed and implemented a novel algorithm for executing top-k queries efficiently in a columnar data store system.
  - ◇ Designed and implemented the software interface to the Data Movement System (DMS) that provides hardware acceleration to data movement and partitioning operations. This helps hide the complexity of custom hardware from the upper level software stack.
  - ◇ Architected and developed a framework to test database operations and improve test coverage of the code base.
  - ◇ Architected and developed a runtime to perform dynamic instrumentation for programs running bare metal on Data Processing Unit (DPU), a non-coherent multi-core system-on-chip running bare metal execution environment.
  - ◇ Implemented Java Virtual Machine for Data Processing Unit's bare metal execution environment.
  - ◇ Publications: *A Many-core Architecture for In-Memory Data Processing* in Micro 2017 and *Big Data Processing: Scalability with Extreme Single-Node Performance* in Big Data Congress 2017

- **Oracle Corporation**—Belmont, CA
  Senior Member of Technical Staff - Oracle Labs, May 2014 - June 2015

- ◇ Worked on RAPID project, a hardware-software co-design system targeting large scale data management and analysis.
- ◇ Promoted to Principal Member of Technical Staff.

- **Department of Computer Sciences**—Madison, WI
  Research Assistant - Vertical Research Group, September 2008 - May 2014
  - ◇ Investigated energy efficient computing through compiler assisted dynamic hardware specialization.
  - ◇ Lead student architect for the Copernicus architecture, a hardware/software co-designed multicore architecture for real-time raytracing.
  - ◇ Advisor: Prof. Karthikeyan Sankaralingam
- **Oracle Corporation**—Belmont, CA
  Student/Intern - Oracle Labs, July 2012 - January 2013
  - ◇ Researched novel methods to improve energy efficiency of database operators with hardware/software co-designed dynamic specialization
- **Intel Corporation**—Santa Clara, CA
  Graduate Intern Technical - Intel Parallel Computing Lab, June 2011 - September 2011
  - ◇ Developed micro-architecture mechanisms and compiler transformations to improve performance of SIMD workloads using dynamic datapath specialization.
  - ◇ Resulted in a publication: *Unifying Functionality and Parallelism Specialization for Energy Efficient Computing* in IEEE Micro Sep/Oct 2012.
- **Intel Corporation**—Hillsboro, OR
  Graduate Intern Technical - Intel Architecture Group, May 2010 - August 2010
  - ◇ Performed graphics workload characterization on future Intel processors and assisted in performance modeling and competitive analysis of next generation graphics pipeline.
- **Epic Systems Corporation**—Madison, WI
  Software Developer: January 2000 - August 2008
  - ◇ Architected and developed a new product Resolute Reference Lab Billing
  - ◇ Directed technical operations of the company as a member of company wide Technical Operations Group
  - ◇ Designed and developed new software modules for large health care facilities
  - ◇ Designed and developed new modules for Resolute Professional Billing: Follow-up workqueue, Cash workqueue, CashDrawer, Reimbursement, Revenue management Report, Distribution/Undistribution Tracking etc.,
  - ◇ Reengineered and optimized existing modules including Payment posting, Claim-edit workqueue, Charge Review workqueue etc.,
  - ◇ Developed frameworks to improve productivity: ARControls, Formless custom billing etc.,
  - ◇ Volunteered and developed tools to improve developer productivity: ResoluteRCS, "New Item Request" tracker, Server Monitor etc.,
  - ◇ Served as Resolute Professing Billing point person for other teams including Foundations, Clarity Report, Performance, Cadence.
  - ◇ Mentored new research and development team members.
  - ◇ Conducted on-site interviews, phone screenings and helped in hiring new employees.
  - ◇ Responsibilities include system analysis, design, program analysis, programming, program reviewing, debugging, testing, quality assurance, and documentation
- **HCL Technologies LTD**—Chennai, Tamilnadu, India
  Member Technical Staff: July, 1999–December, 1999

- ◇ Developed a TL1-SNMP translation agent that translated Transaction Language 1 commands to Simple Network Management Protocol commands.

Patents

- Sam Idicula, Kirtikar Kashyap, Arun Raghavan, Evangelos Vlachos, **Venkatraman Govindaraju**. Hybrid Instrumentation Framework for Multicore Low Power Processors. US Patent 11,030,073.
- Gong Zhang, **Venkatraman Govindaraju**, Sam Idicula. Tail Based top-N query Evaluation. US Patent 11,194,801.

Conference and Journal Publications

- Nikos Armenatzoglou, Sanuj Basu, Naga Bhanoori, Mengchu Cai, Naresh Chainani, Kiran Chinta, **Venkatraman Govindaraju**, Todd J. Green, Monish Gupta, Sebastian Hillig, Eric Hotinger, Yan Leshinksy, Jintian Liang, Michael McCreedy, Fabian Nagel, Ippokratis Pandis, Panos Parchas, Rahul Pathak, Orestis Polychroniou, Foyzur Rahman, Gaurav Saxena, Gokul Soundararajan, Sriram Subramanian, and Doug Terry. Amazon Redshift Re-invented. In ACM SIGMOD/PODS International Conference on Management of Data. June, 2022.
- Sandeep Agrawal, Sam Idicula, Arun Raghavan, Evangelos Vlachos, **Venkatraman Govindaraju**, Venkatanathan Vardarajan, Cagri Balkesen, Georgios Giannikis, Nipun Agarwal, and Eric Sedlar. A Many-core Architecture for In-Memory Data Processing. In Proceedings of the 50^{th} Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), October 2017.
- **Venkatraman Govindaraju**, Sam Idicula, Sandeep Agrawal, Venkatanathan Vardarajan, Arun Raghavan, Jarod Wen, Cagri Balkesen, Georgios Giannikis, Nipun Agarwal, and Eric Sedlar. Big Data Processing: Scalability with Extreme Single-Node Performance. 6^{th} IEEE International Congress on Big Data. June, 2017.
- Tony Nowatzki, **Venkatraman Govindaraju** and Karthikeyan Sankaralingam. A Graph-Based Program Representation for Analyzing Hardware Specialization Approaches. in Computer Architecture Letters(CAL), vol.14, no.2, pp.94-98, July-Dec. 1 2015. **Best of IEEE CAL**
- Chen-Han Ho, **Venkatraman Govindaraju**, Tony Nowatzki, Ranjini Nagaraju, Zachary Marzec, Preeti Agarwal, Chris Frericks, Ryan Cofell, and Karthikeyan Sankaralingam. Performance Evaluation of a DySER FPGA Prototype System Spanning the Compiler, Microarchitecture, and Hardware Implementation. 2015 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), March, 2015.
- **Venkatraman Govindaraju**, Tony Nowatzki, and Karthikeyan Sankaralingam. Breaking SIMD Shackles with an Exposed Flexible Microarchitecture and the Access Execute PDG. 22^{nd} International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2013.
- **Venkatraman Govindaraju**, Chen-Han Ho, Tony Nowatzki, Jatin Chhugani, Nadathur Satish, Karthikeyan Sankaralingam and Changkyu Kim. DySER: Unifying Functionality and Parallelism Specialization for Energy Efficient Computing. IEEE Micro Sep/Oct 2012.
- Jesse Benson, Ryan Cofell, Chris Frericks, Chen-Han Ho, **Venkatraman Govindaraju**, Tony Nowatzki, and Karthikeyan Sankaralingam. Design, Integration and Implementation of the DySER Hardware Accelerator into OpenSPARC. 18^{th} IEEE International Symposium on High Performance Computer Architecture (HPCA), February 2012.
- Shuou Nomura, Matthew D. Sinclair, Chen-Han Ho, **Venkatraman Govindaraju**, Marc de Kruijf, and Karthikeyan Sankaralingam. Sampling + DMR: Practical and Low-overhead Permanent Fault Detection, In Proceedings of the 38^{th} International Symposium on Computer Architecture (ISCA), June 2011.
- **Venkatraman Govindaraju**, Chen-Han Ho, and Karthikeyan Sankaralingam. Dynamically Specialized Datapaths for Energy Efficient Computing. 17^{th} IEEE International Symposium on High Performance Computer Architecture (HPCA), February 2011.

- **Venkatraman Govindaraju**, Peter Djeu, Karthikeyan Sankaralingam, Mary Vernon, and Bill Mark. Toward A Multicore Architecture for Real-time Raytracing. In Proceedings of the 41^{st} Annual International Symposium on Microarchitecture (MICRO), November 2008.

Technical Reports

- Tony Nowatzki, **Venkatraman Govindaraju** and Karthikeyan Sankaralingam. Studying Hybrid Von-Neumann/Dataflow Execution Models. University of Wisconsin Computer Sciences Technical Report. TR-1820, July 2015.
- **Venkatraman Govindaraju**, Chen-Han Ho, Tony Nowatzki and Karthikeyan Sankaralingam. Mechanisms for Parallelism Specialization for the DySER Architecture. University of Wisconsin Computer Sciences Technical Report. TR-1773, June 2012.
- **Venkatraman Govindaraju**, Chen-Han Ho, and Karthikeyan Sankaralingam. Design and Evaluation of Dynamically Specialized Datapaths with the DySER Architecture. University of Wisconsin Computer Sciences Technical Report. CS-TR-2010-1683, November 2010.

Posters

- Jesse Benson, Ryan Cofell, Chris Frericks, **Venkatraman Govindaraju**, Chen-Han Ho, Zachary Marzec, Tony Nowatzki, and Karthikeyan Sankaralingam. Prototyping the DySER Specialization Architecture with OpenSPARC. Hot Chips 24, August 2012.

Invited and Conference Talks

- Breaking SIMD Compiler Shackles. 17^{th} Annual Wisconsin Architecture Industrial Affiliates Meeting, Madison, WI, October 2013.
- Automatic Data-Parallel Acceleration with DySER. Google, Madison, WI, October 2013.
- Breaking SIMD Shackles with an Exposed Flexible Microarchitecture and the Access Execute PDG. 22^{nd} International Conference on Parallel Architectures and Compilation Techniques (PACT), Edinburgh, Scotland, September 2013.
- Dynamic Hardware Specialization with DySER. 15^{th} Annual Wisconsin Architecture Industrial Affiliates Meeting, Madison, WI, October 2011.
- Toward A Multicore Architecture for Real-time Ray-tracing. 41^{st} International Symposium on Microarchitecture (MICRO), Lake Como, Italy, November 2008.
- Toward A Multicore Architecture for Real-time Ray-tracing. 12^{th} Annual Wisconsin Architecture Industrial Affiliates Meeting, Madison, WI, October 2008.

Open Source Contributions

- Contributions to LLVM:
  - ◇ Official code owner of SPARC backend.
  - ◇ Implemented JIT (Just In Time compilation), leaf procedure optimization, delay slot filler, PIC code generation, long double support in SPARC code generator.
  - ◇ Improved SPARC backend to generate ABI compliant code.
- Slicer - A publicly available compiler for DySER
  - ◇ Owner and maintainer of DySER Compiler
  - ◇ Available at http://research.cs.wisc.edu/vertical/dyser-compiler/

Awards and Honors

- **Best of IEEE Computer Architecture Letters Award**, 2015, for "A Graph-Based Program Representation For Analyzing Hardware Specialization Approaches"
- **Wisconsin Computer Science Departmental Summer Fellowship**, May 2009.
- **Best Student Presentation Award**, 41^{st} International Symposium on Microarchitecture (MICRO), November 2008, "Toward A Multicore Architecture for Real-time Raytracing".

Professional Activities

- Reviewer: Computer Architecture Letters (CAL) 2009, ISPASS 2010, JPDC, FCCM, Scientific Programming
- ACM member, IEEE Student member

Skills

- **Computer Languages:** C, C++, Java, Python, Perl, x86 and SPARC Assembly, Visual Basic, Unix Script
- **Software:** LLVM, GCC Internals, gem5 Simulator, Simics, Pin Tools, Jikes RVM
- **Operating Systems:** Unix variants, Windows