Matthew E. Tolentino

8226 Steamboat Island Rd NW Olympia, WA 98502,
253-371-9254 (office), 360-951-8607 (mobile), 360-878-9853 (home)
matthew.e.tolentino@intel.com, matthew.tolentino@gmail.com

Education

Virginia Polytechnic Institute and State University, Blacksburg, VA ***2009***
Doctor of Philosophy in Computer Science
Advisor: Professor Kirk W. Cameron
Committee: Gary Beihl, Professors Ali Butt, Dimitris Nikolopoulos, Calvin Ribbens
Ph.D. Thesis: "Managing Memory for Power, Performance, and Thermal Efficiency"

University of Washington ***2004***
Master of Science in Computing & Software Systems
Advisor: Professor George Mobus
M.S. Thesis: "Flexible Operating System Structure for Dynamic Memory Management"

University of Maryland, College Park, MD ***2000***
Bachelor of Science in Computer Science

Research Interests

Computer architecture, high-performance computing, parallel and distributed systems, throughput computing, operating systems, power and performance modeling and analysis, and data analytics/scientific computing

Academic Experience

UNIVERSITY OF WASHINGTON ***2011 – Present***
Tacoma, WA

Affiliate Assistant Professor/Industry Fellow ***2011- Present***

- Classes instructed/co-instructed:
  - TCSS/CE 371 - Intro to Computer Architecture & Machine Organization
  - TCSS/CE 372 - Advanced Computer Architecture
  - TCSS/CE 422 - Operating Systems
  - TCSS 590 - Big Data Architecture
  - TCSS 558 - Applied Distributed Systems

Professional Experience

INTEL CORPORATION ***2000 – Present***

Research Scientist ***2010- Current***

- Leading cross-divisional (DCSG, SSG, TMG), micro-architectural I/O characterization and modeling initiative for server workloads including Facebook (MemcacheD/Tao), SPC-2, Hadoop/UCB Spark-based workloads, SPECWeb, and HPC-Challenge benchmarks to determine future Silicon I/O landing zone requirements for 2018+ server platforms across all server segments. Characterization includes bandwidth details, locality given varied on-die interconnect topologies, frequency and impact of conflicts, packet size distribution. Developed network, on-die interconnect, fabric, and storage analytical models to analyze the performance impact of architectural alternatives at component, node, rack, and data center granularities using key workloads from Industry to determine the direction of the Intel roadmap for server CPU, chipset, fabric, and storage technologies through 2020.
- Researching performance impact of tiered storage models and shared rack-scale fabric interconnects for Intel proposed Rack Scale Architecture (RSA). Extending component and subsystem-level modeling to develop a cluster-level analytical model (CPU/memory/network/storage) to reason about potential impact of on-die integration of IO components, fabric bandwidth requirements, ASIC requirements for fabric switches, and performance impact of increased latency. Paper detailing methodology and initial results to be submitted to ccGrid 2014.
- Co-developed a customer-representative Hadoop benchmark suite (including an ETL proxy, CDR-based queries, and compute-bound iterative solvers) to enable first pass architectural analysis of representative Big Data workloads in multi-tenant data center environments. Results published at IISWC 2013.
- Analyzed memory architectural options (1LM vs. 2LM composition) based on MemcacheD, SIR, and mWeb memory utilization and traffic to determine pre-roadmap, Denverton micro-server CPU architecture definition. Leveraged integrated performance counters along with DDR3 DIMM traces in an 8Core/8Cbo configuration as well as scaled 2Core/2Cbo configuration on current SNB-based JKT Xeon servers to empirically determine Core CPI, LLC MPI, and blocking factors to drive analytical modeling for 2LM architectural performance input to Denverton program. Results showed minimal performance benefit for 2LM.
- Developed trace-based simulation methodology to identify MPI-based HPC workload sensitivity to interconnect bandwidth and latency for upcoming Stormlake fabric architecture to be integrated on-die into Intel 2017 Xeon Phi CPUs (KNH, KNP). Extended simulator to model impact of increases in per-node fabric bandwidth under various latency scenarios and evaluated over 25 HPC kernels using between 64 MPI ranks to 4096 MPI ranks. Analysis narrowed scope of further fabric modeling and analysis efforts to 6 key HPC workloads for further large-scale fabric simulation as part of KNH path-finding.
- Co-PI for research grant from National Security Agency (via Intel Federal) for Power-Performance Architectural Analysis for Future Exascale Systems. Led micro-architecture power and performance characterization team for Intel Exascale research

contract with US Government (NSA). The goal of this research was to analyze the power characteristics of on-die interconnect traffic patterns in many-core CPUs using HPC codes and propose architectural changes to improve the energy and performance efficiency for 2017+ Exascale CPUs and systems. Set up an experimental cluster configuration to capture current power characteristics of each power plane within current, pre-production Xeon server CPUs for HPC-Challenge benchmarks, SSCA-2, and Graph-500. Defined analytical modeling techniques to assess scaling micro-architectural power management techniques to future many-core, Exascale processors (Intel Knight's family) in the 2020 timeframe. Delivered technical report with findings and results and presented results at several DoE National Labs. Analysis and methodology integrated into Intel proposal (currently in process) for the Exascale CORAL RFP from DoE.

- Led path finding analysis for newly proposed C-state for Skylake/KNL (Intel CPUs scheduled for 2016) for SAF. This idle power state reduces active power C0 residency and entry/exit latencies by avoiding the performance-penalty of flushing and reloading the mid-level (e.g. L2) cache during core C0 to C3 transitions. This enables the performance attributes of core-C1 with core-C3 power characteristics, improving energy efficiency for all workloads with idle periods by autonomically reducing CPU power. Architectural proposal accepted and integrated into Skylake (2016) CPU design.
- Analyzed QPI and PCI-e traces over the SPECweb load line to identify opportunities to dynamically scale QPI and PCI-e bus widths (e.g. L0p and DLW) to reduce power while preserving latency and bandwidth. Co-developed QPI dynamic link width algorithm proposed for FSM implementation within uncore of 2016 Skylake processor.
- Analyzed power and performance characteristics of SPECweb and SPECpower workloads on EP and EX segment server systems to project the impact of power-related micro-architectural modifications on future platforms.

Software Architect ***2005 - 2010***

- Led cross-divisional (NBI, ECG, & DCG) technical investigation of power and performance capabilities of General Vision Cognimem processor, a high-speed neural network in silicon. Co-designed PCI-Express board to analyze performance of HPC and data-mining streaming workloads, architected and built system accelerator software stack and libraries for prototype USB and PCI-Express devices, and benchmarked prototype devices against current Xeon servers and Nvidia CUDA-based GPGPUs. Provided performance analysis as input to DCG 2010 PLBP and ECG TRP processes for inclusion in DCG and ECG roadmaps.
- Processor Simulation Technologies: Led development of CPU and platform simulator for 2012 Itanium Processor for pre-Silicon validation of system software stack.
- Intel QuickPath Interconnect (QPI) Technology: Designed, developed, and delivered BIOS Reference Software for enabling the high-speed, serial interconnect to OEMs productizing systems based on Intel i7 (Nehalem) processors and supporting chipsets (Tylersburg/Boxboro). Led QPI BIOS power-on of first-ever QPI-enabled processors and chipset based platforms. Extended QPI Reference Code for Multi-processor QPI link-based platforms and led several additional board power-ons (EP and EX lines). Developed and delivered QPI BIOS training collateral to OEMs.

- Green Server Initiative: Served as firmware and operating system consultant within cross organizational working group to deliver an energy-efficient server prototype using a Nehalem-based ESPD system.

Chipset & OS Software Engineer ***2002 - 2005***

- Memory RAS: Led initiative within Intel to develop memory capacity-add, capacity-remove, and capacity-replace capabilities within the Linux operating system for high-end enterprise platforms. Collaborated with engineers from IBM, SGI, HP, Fujitsu, VA Linux, as well as the open source community to incorporate support into Linux. Delivered solution for hot-add and prototype support for hot-remove memory in support of the Intel Twincastle chipset (the first Intel chipset to support hardware-removable DDR-based memory), shipped within the Harwich platform.
- EFI/UEFI Firmware: Led Linux operating system enabling initiative for UEFI firmware standard that has been incorporated into all Intel-based platforms.
- AGP 3.0 Enabling: Developed Linux driver support for AGP 3.0 compliant chipsets and merged into mainstream Linux kernel. Developed and released chipset specific drivers for several Intel chipsets including E7505, E7205, and 875i.

Operating Systems Engineer ***2000 - 2002***

- Led operating system power-on team (Windows, Linux, and AIX) during power-on of the first dual-processor Itanium platform (Tiger2).
- Researched OS I/O scalability limitations discovered during hardware validation testing of first generation Itanium platforms. Authored internal technical report, *Linux I/O Scalability on Itanium I Platforms*, with findings and evaluation of prototype solution. Collaborated with IBM and Intel SSG division to drive scalable solution into Linux kernel.
- Developed platform validation toolkit framework to induce, detect, and report processor, chipset, and system software defects discovered during platform development prior to OEM deployment. Developed implementations for the Itanium and IA-32 architectures that are currently used in the Intel Enterprise Platform Validation Laboratories
- Ported SCSI, Fibre Channel, Ethernet device drivers for the Linux, AIX, and Unixware operating systems to the Itanium architecture during the initial development of the first-ever Itanium processor-based platform (Merced on Lion platform).

UNIVERSITY OF MARYLAND, College Park, MD ***1999 - 2000***

Teaching Assistant:

- Instructed classes, graded projects, and mentored students for the following classes:
- Advanced Data Structures (CMSC 420)
- Computer System Architecture (CMSC 311)
- Introduction to C/C++ (CMIS 140)

NAVAL RESEARCH LABORATORY, Washington DC ***Summer 1998***
Undergraduate Plasma Physics Researcher:

- Researched dusty plasma theory, built and configured space simulation chambers, conducted experiments, recorded and analyzed observations and related data, and compared experimental results to previously published dusty plasma results.
- Developed analytical and graphical simulation software to model observed results.

U.S. DEPARTMENT OF THE NAVY, Pentagon, Washington DC ***1996 – 2000***
Department of the Navy Organization, Management and Infrastructure Team Staff:

- Researched inquiries from administration, congress, and general public (FOIA)
- Drafted, reviewed, and delivered Naval correspondence to DOD and cabinet-level administration officials
- Managed electronic correspondence tracking database

Deputy Under Secretary of the Navy Staff:

- Researched White House, Congressional and public inquiries; drafted, reviewed, and delivered responses to administration officials as well as House and Senate members and sub-committees.

Secretary of the Navy Staff:

- Managed imaging storage database for all internal and external correspondence
- Managed classified correspondence within secretariat

DEPARTMENT OF RECREATION ***1994 - 1996***
Director of Skating Programs

- Designed, implemented, and managed portfolio of instructional ice skating programs. Hired and managed a staff of 20 instructors, organized the annual Ice Show, directed the District Competition, represented the Department at national conventions, and managed public relations for the instructional program.

Publications

Journals:

- Matthew E. Tolentino, Kirk W. Cameron. The Optimist, the Pessimist, and the Global Race to Exascale in 20 Megawatts, *IEEE Computer, Vol. 45, Issue 1,* January, 2012.
- Matthew E. Tolentino, Joseph Turner, Kirk W. Cameron. Memory MISER: Improving Main Memory Energy Efficiency in Servers, *IEEE Transactions on Computers,* Vol. 58, Issue 3, pp. 336-350, March 2009.

Conference/Workshops:

- HcBench: Methdology, Development, and Characterization of a Customer Usage Representative Big Data/Hadoop Benchmark, *Proceedings of IEEE International Symsposium on Workload Characterization,* Portland, OR, 2013.

- Matthew E. Tolentino, Joseph Turner, Kirk W. Cameron. Memory-MISER: A Performance-Constrained Runtime System for Power-Scalable Clusters, *Proceedings of ACM International Conference on Computing Frontiers,* Ischia, Italy, 2007.
- Matthew E. Tolentino, Joseph Turner, Kirk W. Cameron. An Implementation of Page Allocation Shaping for Energy Efficiency, *Proceedings of Third Workshop on High-Performance, Power-Aware Computing (HPPAC) held in conjunction with International Parallel & Distributed Processing Symposium,* Long Beach, CA, 2007.
- Matthew E. Tolentino and Kirk W. Cameron. Improving the Energy Efficiency of High-Performance Server Systems, *Work-in-progress at the 20th Symposium on Operating System Principles (SOSP),* Brighton, UK, October 2005.
- Joel Schopp, Dave Hansen, Mike Kravetz, Hirokazu Takahashi, Iwamoto Toshihiro, Yasunori Goto, Kamezawa Hiroyuki, Matt Tolentino, Bob Picco. Hotplug Memory Redux, In *Proceedings of the Ottawa Linux Symposium,* July 2005.
- Dave Hansen, Mike Kravetz, Bradley Christiansen, Matt Tolentino. Hot-plug Memory and the Linux VM. In *Proceedings of the Ottawa Linux Symposium,* pp. 278-294, July 2003.
- Matthew E. Tolentino. Linux in a Brave New Firmware Environment. In *Proceedings of the Ottawa Linux Symposium,* pp. 433-446, July 2003.

Theses:

- Matthew E. Tolentino. Managing Memory for Power, Performance, and Thermal Efficiency, Ph.D. Thesis, Virginia Tech, May 2009.
- Matthew E. Tolentino. Flexible Operating System Structure for Dynamic Memory Management. Master's Thesis, University of Washington, Tacoma, December 2004.

Papers in Preparation/Submitted:

- Matthew E. Tolentino, Vikram Saletore. Modeling the Performance Impact of Future Storage Technologies for Data Analytics Workloads. To be submitted to IEEE/ACM International Symposium on Cluster, Cloud, & Grid Computing (ccGrid), 2014.
- Matthew E. Tolentino, Scott McMilllan, Ram Hugahalli. Latency and Bandwidth Sensitivity Analysis for Future HPC Interconnects. To be submitted.
- Matthew E. Tolentino, Vikram Saletore. Modeling I/O Architectures for Big Data Workloads – The Case for Multi-Level Tiered Storage over Shared Fabrics. To be submitted.

Talks

- Invited Panelist at SIAM Conference on Parallel Processing for Scientific Computing, Portland, OR, February 2014
- Invited Panelist, Pure Energy Session Chair, and Program Committee Member, First Usenix Workshop on Sustainable Information Technology (SustainIT '10), San Jose, CA, February 2010.
- Managing Memory for Power and Performance Efficiency, Department of Computer Science, Portland State University, OR, September 2008.

- Improving the Energy Efficiency of High-Performance Server Systems, *Work-in-progress at the 20th Symposium on Operating System Principles (SOSP)*, Brighton, UK, October 2005.
- Dynamic Physical Memory Management in Linux: An Operating System Foundation for Advanced Memory Technologies. Intel Developer Forum, San Francisco, CA, March 2005.
- Next Generation EFI Operating System: Linux & EFI. Intel Developer Forum, San Francisco, CA, September 2004.
- Next Generation EFI 32 Operating System Loader. Intel Developer Forum. San Jose, CA. September 2003

Patents & IP Disclosures

- Bruce Bahnsen, Sri Sakthivelu, Vikram Saletore, Vish Viswanathan, Matthew E. Tolentino, Gopi Govindaraju, Vincent Zimmer, "Atomic Transactions to Non-Volatile Memory", Intel Disclosure #102074, Reference #P58349, Submitted to US PTO September 2013,
- Matthew E. Tolentino, "System and method to enable quality-of-service, memory bandwidth for multiple virtually hosted operating systems," Intel Disclosure #.
- Matthew E. Tolentino, "Method to establish and dynamically control energy consumption in large-scale datacenters or IT infrastructures," US Patent Granted #US8301925B2,
- Matthew E. Tolentino, "Co-scheduling interrupt redirection with DVFS state transitions on link-based, multiprocessor systems for performance and energy efficiency," Intel Disclosure Number 56828
- Matthew E. Tolentino, Joseph Turner, Kirk W. Cameron, "Performance-centric runtime system for minimizing energy consumption in power-scalable cluster architectures," Intel Disclosure Number 52904, VTIP Disclosure Number 06.079
- Michael A. Rothman, Vincent J. Zimmer, Robert C. Swanson, Matthew E. Tolentino, "System and method to establish a peer-to-peer IT backbone," US Patent Granted #US8024477B2, September 2011.
- Robert C. Swanson, Mallik Bulusu, Vincent J. Zimmer, Michael A. Rothman, Matthew E. Tolentino, "Identifying an operating system associated with a boot path," US Patent Granted #US20060288197A1.

Professional Service

- Invited Attendee, National Academy of Engineering, 2014 Indo-American Frontiers of Engineering Symposium, Mysore, India, May 2014.
- Board Member, Advisory Board for Institute of Technology at University of Washington, Tacoma, 2010 - Present
- Technical Program Committee, International Supercomputing Conference (ISC) (Architecture Track), 2013, 2014

- Technical Program Committee, IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (ccGrid) (Architecture Track, Scheduling Track), 2013, 2014
- Technical Program Committee, IEEE/ACM Supercomputing Conference (SC) 2014
- Technical Program Committee, First Usenix Workshop on Sustainable Information Technology (Sustain IT)
- Technical Merit Reviewer/Panelist for NSF Cyber-Physical Systems Program, June 2009
- Technical Merit Reviewer, Department of Energy Smart Grid Investment Grant (DoE SGIG), August 2009
- Technical Merit Reviewer, Department of Energy Smart Grid Demonstration Program (DoE SGDP), September 2009
- Technical Merit Reviewer, Department of Energy Workforce Training for the Electric power Sector, December 2009
- Reviewer, IEEE Transactions on Computers (IEEE TC)
- Reviewer, IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS)
- Reviewer, Cluster Computing, Journal of Networks, Software Tools and Applications
- Reviewer, IEEE International Parallel & Distributed Processing Symposium (IPDPS)
- Reviewer, Parallel Architectures and Compilation Techniques (PACT)
- Reviewer, Europar
- Reviewer, High-Performance Power-Aware Computing Workshop (HPPAC)
- Invited Attendee, National Academy of Engineering Grand Challenges Summit hosted by Duke University, Durham, NC, March 2009