Yoshio Turner
Hewlett Packard Laboratories
1501 Page Mill Road, MS 1177
Palo Alto, CA 94304
Tel: (650) 236-2845
Email: yoshio.turner@hp.com
http://www.hpl.hp.com/personal/Yoshio_Turner/

RESEARCH INTERESTS

System Architecture, Networking, Operating Systems, Virtual Machines

EXPERIENCE

Senior Research Scientist — **Hewlett-Packard Laboratories**
Palo Alto, CA — May 2006 – present

Research Scientist — **Hewlett-Packard Laboratories**
Palo Alto, CA — Oct 1999 – May 2006

Linux-Based Scalable Network Platforms: (2003 – 2006)
Joint work with Intel Labs, T. Brecht, B. Lynn, G. Janakiraman. Developed and evaluated a prototype implementation of a new system software architecture supporting networking applications on TCP/IP networks. The prototype is based on Intel's ETA architecture which attempts to increase packet processing efficiency by dedicating some CPU cores of a multiprocessor to run an optimized TCP/IP stack on behalf of applications running on the remaining CPU cores. To enable applications to efficiently handle large numbers of concurrent I/Os, an API is provided which exposes to applications the asynchronous execution of network I/O operations. Specific contributions:

- Defined the research agenda and setup the experimental environment.
- Led the writing effort for our initial position paper at the 3rd Workshop on System Area Networks and presented the talk.
- Designed and developed an asynchronous I/O (AIO) layer, and extended the userver, an event-driven web server application, to use the AIO API. Identified key design issues for applications using network AIO, providing a useful case study for developers of applications using emerging AIO APIs (like Open Group Extended Sockets).
- Made significant improvements to the initial Intel ETA prototype.
- Analyzed performance results and prepared and presented final paper at EuroSys2006. Results from this project were also reported in a formal paper and presentation at HP Techcon2006. Experimental results using a SPECWeb99-like workload show modestly improved CPU efficiency with this architecture compared to a well-configured performance-tuned baseline Linux kernel. The results also quantify the impacts of the partitioning and AIO API aspects of the architecture and identify potential improvements in future systems using multi-core processors.

I/O Device Virtualization: (2004 – 2005)
Joint work with A. Menon, J. Santos, G. Janakiraman focused on understanding the performance impact of alternative approaches for network device virtualization in virtual machine environments. Specific contributions:

- Conceived the idea of developing Xenoprof, a mechanism integrated into the Xen virtual machine monitor to provide detailed insight into the performance properties of applications running in Xen. Xenoprof supports system-wide statistical performance profiling using Intel and AMD hardware performance counters.

- Guided the design and development of Xenoprof. Xenoprof is now open source, merged into the official Xen and OProfile distributions, and actively being used by developers throughout the Xen community.
- Used Xenoprof to evaluate and analyze alternatives for implementing network device virtualization in the Xen environment. This work identified and quantified major causes of significant performance degradation.
- Described the Xenoprof architecture and our experimental results and conclusions in a paper at the ACM/USENIX Virtual Execution Environments (VEE) conference in June 2005.

Application migration reducing downtime for OS upgrades: (2003 – 2005)
Joint work with D. Subhraveti, G. Janakiraman, J. Santos. This work extended the "Zap" application virtualization and migration system initiated at Columbia University. Specific contributions:

- Significantly improved the functionality of Zap to support a wider variety of applications which use several sophisticated kernel services including kernel threads, memory mapped files, and unix sockets.
- Developed mechanisms for application-transparent migration that preserves live network connections. This capability enables online upgrade of OS kernels to be performed without incurring application downtime.
- Developed a new coordinated checkpoint-restart mechanism for distributed applications which leverages the live migration capabilities. Our approach eliminates expensive complexities of prior approaches (e.g., flushing channels, continuously logging messages).
- Developed a poster and demo describing this work at HP TechCon 2005 and a paper at the Dependable Systems and Networks (DSN) 2005 conference.

Automated System Design for Availability: (2002 – 2004)
Joint work with J. Santos, G. Janakiraman. This work focused on automating the normally arduous task of designing a cluster computing infrastructure for an Enterprise or Internet service to meet a target level of service performance and service availability at minimum cost, a difficult problem faced by human designers of systems since there are many design alternatives that must be evaluated. Specific contributions:

- Contributed to develop of a proof of concept automation engine (Aved) to automatically design and configure computing infrastructure for large-scale systems.
- Defined several aspects of the modeling framework used to drive Aved's design space search procedure, enabling Aved to model a much larger variety of system infrastructure and configuration alternatives.
- Developed demonstrator, papers, and patents. An Aved description was included in an HP Techcon paper (Smart Data Center), and a paper on Aved appeared in the Dependable Systems and Networks (DSN) 2004 conference.

Congestion Control for System Area Networks: (2000 – 2003)
Joint work with J. Santos, G. Janakiraman. This work developed new mechanisms for detecting and responding to network congestion which are tailored to the unique properties of system area networks (SAN). In particular, the use of backpressure flow control and small buffer capacity input-buffered SAN switches can lead to network-wide congestion spreading and call for radically different congestion control approaches compared to TCP. Specific contributions:

- Led the design of novel switch-based explicit congestion detection and notification mechanisms, and source response mechanisms that combine rate and window control as appropriate for SAN environments.
- Evaluated performance and fairness properties of the new mechanisms using simulation to demonstrate effectiveness.
- Presented results of this work at the IEEE Infocom conference and presented to the Link-level Working Group of the InfiniBand standard body. The congestion control proposal recommended by this working group was significantly influenced by this work.

Reduced Energy Decoding of MPEG Streams: (2001-2002)
With Malena Mesarina, developed an efficient and optimal application-specific scheduling and voltage scaling algorithm for trading off energy consumption and frame display rate for decoding MPEG audiovideo streams without dropping frames. In contrast to previous work, this approach factors in the limited memory resources available on portable devices. This work was selected to appear in a special issue of the ACM Multimedia Systems Journal as one of the three top papers at the ACM/SPIE Multimedia Computing and Networking 2002 conference.

System Architect **Samsung L.A. Design Center**
Cypress, California 1996 – 1998
Led the technical design aspects of a system-on-chip microcontroller ASIC and associated real-time OS and application software for the multi-function print/copy/fax machine market. Extensive experience with RTL Verilog, hardware-software co-design, RTOS, logic synthesis, FPGAs, ARM processors, and DRAM and DMA controller design.

KEY REFEREED PUBLICATIONS (* DENOTES PRESENTER)

T. Brecht, G. Janakiraman, B. Lynn, V. Saletore, Y. Turner*, "Evaluating network processing efficiency with processor partitioning and asynchronous I/O", *EuroSys2006*, Leuven, Belgium, April 2006.

G. Janakiraman*, Y. Turner, "Using Multi-core CPUs to Improve Network Processing Efficiency", *HP Techcon*, 2006.

G. Janakiraman*, J. R. Santos, D. Subhraveti, Y. Turner, "Cruz: Application Transparent Distributed Checkpoint-Restart on Standard Operating Systems," *International Conference on Dependable Systems and Networks (DSN – DCCS track)*, Yokohama, Japan, June 28 - July 1, 2005.

A. Menon, J. R. Santos, Y. Turner*, G. Janakiraman, W. Zwaenepoel, "Diagnosing Performance Overheads in the Xen Virtual Machine Environment," *First ACM/USENIX Conference on Virtual Execution Environments (VEE)*, Chicago, IL, June 11-12, 2005.

G. Janakiraman, J. R. Santos, Y. Turner, R. Badrinath, G. N. Manjunatha, D. Williams, "Zap: Application-transparent process migration," *HP Techcon*, 2005 (poster).

G. Janakiraman, J. R. Santos*, Y. Turner, "Automated System Design for Availability," *International Conference on Dependable Systems and Networks (DSN – DCCS track)*, Florence, Italy, June 28- July 1, 2004.

M. Mesarina and Y. Turner, "Reduced energy decoding of MPEG streams," *ACM/Springer Multimedia Systems Journal* (special issue with extended versions of selected papers from ACM/SPIE MMCN'02), 9(2), August 2003

C. Patel* et al, "Smart Data Center," *HP Techcon*, 2003.

J. R. Santos, Y. Turner*, G. Janakiraman, "End-to-End Congestion Control for InfiniBand," *IEEE INFOCOM 2003*, San Francisco, CA, April 1-3, 2003.

OTHER REFEREED PUBLICATIONS (* DENOTES PRESENTER)

J. Gummaraju, Y. Turner*, "Hidra: history-based dynamic resource allocation for server clusters", *International Conference on Internet Technologies and Applications (ITA-05)*, Wrexham, North Wales, September 2005.

Y. Turner*, T. Brecht, G. Regnier, V. Saletore, G. Janakiraman, B. Lynn, "Scalable networking for next-generation computing platforms," *3rd Annual Workshop on System Area Networks (SAN-3)*, Madrid, Spain, February 2004.

G. Janakiraman, J. Santos*, Y. Turner, "Automated Multi-Tier System Design for Service Availability," *First Workshop on the Design of Self-Managing Systems (at DSN-2003)*, San Francisco, CA, June 23, 2003

J. R. Santos*, Y. Turner, G. Janakiraman, "Evaluation of Congestion Detection Mechanisms for InfiniBand Switches," *High-speed Networking Symposium (IEEE GLOBECOM)*, Taiwan, November 17-21, 2002.

J. R. Santos*, K. Dasgupta, G. Janakiraman, Y. Turner, "Understanding Service Demand for Adaptive Allocation of Distributed Resources," *Global Internet Symposium (at IEEE GLOBECOM)*, Taiwan, November 17-21, 2002.

M. Mesarina* and Y. Turner, "Reduced energy decoding of MPEG streams," *ACM/SPIE Multimedia Computing and Networking (MMCN)*, San Jose, CA, January 2002

Y. Turner* and Y. Tamir, "Connection-based adaptive routing using Dynamic Virtual Circuits," *Parallel and Distributing Computing and Systems (PDCS)*, Las Vegas, NV, October 1998.

Y. Tamir and Y. Turner*, "High-performance adaptive routing in multicomputers using Dynamic Virtual Circuits," *6th Distributed Memory Computing Conference (DMCC6)*, Portland, OR, April 1991.

NON-REFEREED TECHNICAL REPORT

Y. Turner, J. R. Santos, and G. Janakiraman, "An approach for congestion control in InfiniBand," HP Labs Technical Report HPL-2001-277R1, May 2002.

EDUCATION

University of California, Los Angeles Sep 2005
Ph.D in Computer Science

University of California, Los Angeles Jun 1991
M.S in Computer Science

University of California, Los Angeles Jun 1988
B.S. in Computer Science and Engineering
summa cum laude

HONORS

Fannie and John Hertz Graduate Fellowship (1990-1995)
UCLA Chancellors Fellowship (1988-1989)