Ying Ye

5775 Morehouse Drive
San Diego, CA 92121

yingy@bu.edu
(857) 272-2053

Profile

6 years OS/hypervisor kernel dev; 10 years C/C++

Work Experience

Qualcomm Technologies Inc., San Diego, CA
Senior Engineer (September 2017 - current)
Working on the modem DSP architecture design, e.g., cache hierarchy, branch prediction. At the same time, working on performance optimization for modem's real-time OS kernel.

VMware Inc., Palo Alto, CA
Software Engineer Intern (May - August 2016)
Joined VMware **ESXi hypervisor** resource management team. Investigated Intel processor's new cache partitioning feature and implemented a cache management logic inside hypervisor kernel. Result showed reduced cache contention for virtual machines with low latency requirement, thus improving system's quality of service.

Cisco Systems Inc., Cambridge, MA
Software Engineer Intern (May - August 2015)
Worked in a 20+ **agile** development team building web-based application that manages Cisco security devices (routers and firewalls). Contributed most to the backend server based on Java Spring Framework. Heavily involved in the whole scrum process with weekly sprint.

Research Projects

Multicore Real-Time Operating System for Autonomous Driving
(September 2016 – April 2017)
Implemented a real-time **virtualization** system (~180K LOC). The solution allows a general purpose OS (Linux) to be running inside a VM while still being controlled by a real-time OS in another VM. Experiment showed that, while running CNN-based image classification (GPU) side by side with an object avoidance task (LIDAR), the real-time performance of the latter was strictly protected.

Predictable Resource-aware Multicore Scheduling
(September 2014 – May 2016)
Designed a real-time multicore scheduling algorithm and load balancing framework to address contention on shared caches and memory buses, while improving CPU utilization. Proposed a scheduling policy with novel performance metric to manage memory bus contention, which outperformed existing metric. Experiment showed that our system's real-time properties were greatly enhanced.

Performance Isolation via Software Cache Partitioning
(January 2013 - February 2014)

Proposed a **memory management** framework in **Linux kernel** that partitions the shared hardware cache on multicore with the use of dynamic page coloring, providing performance isolation amongst applications. Experiments showed a noticeable reduction in program performance variation up to **39%**.

Skills C/C++, assembly, x86, Java, Shell Script, Python

Awards
Intel Modern Code Developer Challenge: Trailblazer ($500) (Fall 2015)
Champion of Chinese Collegiate National Climbing (Summer 2011)

Education
Boston University, Boston, MA GPA: 3.76 / 4
PhD in Computer Science (September 2011 - August 2017)

Publications

vLibOS: Babysitting OS Evolution with a Virtualized Library OS
Ying Ye, Zhuoqun Cheng, Soham Sinha, Richard West, Technical Report, arXiv: 1801.07880, arXiv.org.

MARACAS: A Real-Time Multicore VCPU Scheduling Framework
Ying Ye, Richard West, Jingyi Zhang, Zhuoqun Cheng, in Proceedings of the 37th IEEE Real-Time Systems Symposium.

COLORIS: A Dynamic Cache Partitioning System Using Page Coloring
Ying Ye, Richard West, Zhuoqun Cheng, Ye Li, in Proceedings of the 23rd International Conference on Parallel Architectures and Compilation Techniques.

Building Real-Time Embedded Applications on QduinoMC: A Web-connected 3D Printer Case Study (**Best Student Paper**)
Zhuoqun Cheng, Richard West, **Ying Ye**, in Proceedings of the 23rd IEEE Real-Time and Embedded Technology and Applications Symposium.

Code: github.com/chrisyy/