

Research Topics for Senior Projects

Hitoshi Oi

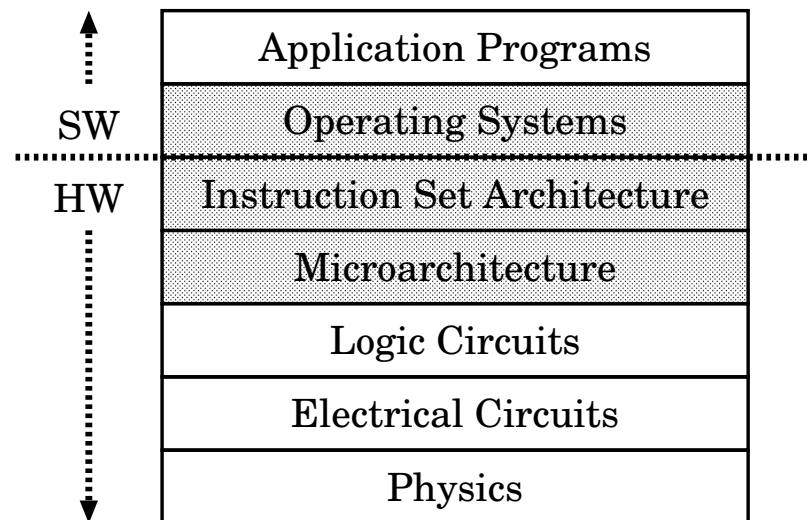
The University of Aizu

October 28, 2019



Computer Architecture and Operating Systems Group

Computer Systems Abstraction Layers



We are mostly working in the shaded layers

Research Interests in General

Hardware/Software Interaction and Co-Design

- How modern (& realistic) software accesses hardware components ?
- How modern computer systems are designed and how they can be better utilized ?

Primary Metrics

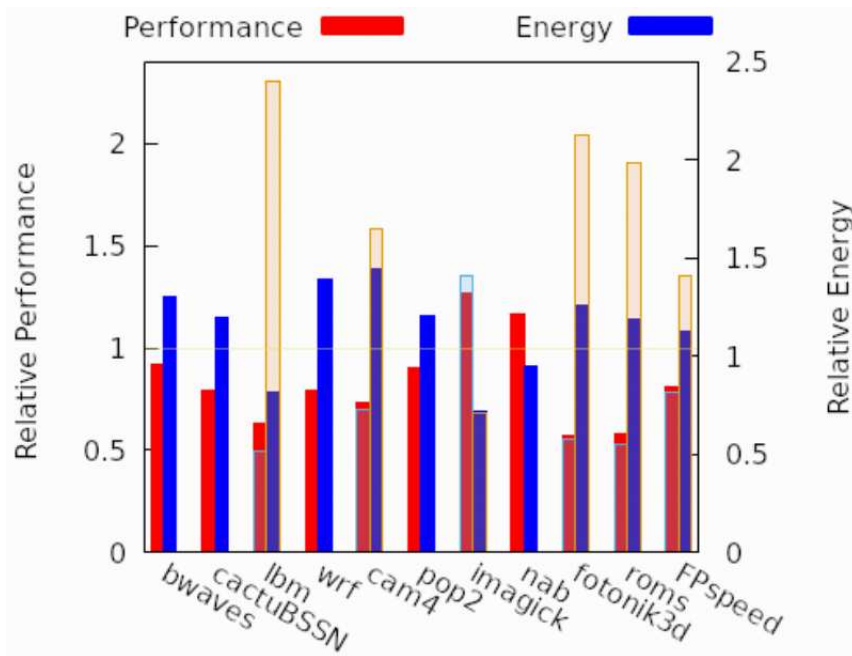
- Performance: How many tasks (or much work) can be finished in a fixed amount of time ?
- Energy-efficiency: How can we reduce the energy consumption for the same amount of work ?

Example of Recent Research

Intel Core i7 vs AMD Ryzen 5 using SPEC CPU2017

- Both are (relatively) new higher-end CPUs from competing companies
- Before Ryzen, high-end CPU market was dominated by Intel
- The particular models used in this study are Core i7-8700 and Ryzen 5 1600; both are six-core and 12-thread processors.
- Ryzen has: larger L1I, L2 and L3 caches and dedicated schedulers for INT and FP (Core has a shared scheduler).
- CPU2017 is the latest benchmark suite from SPEC. Next slide shows the comparison using FP Speed.

Comparison with SPECspeed Floating Point



- Relative Speed of Core i7 is 57 to 127% of Ryzen 5 (Red).
- Relative energy 69 to 139% (Blue).
- In 4/10 benchmarks, best performance/energy efficiency achieved with less than 12-threads.

Topics of Recent Students (1)

Security in Internet of Things

- Things that used to work stand-alone are now connected each-other and accessible over the Internet
- For example, electric appliances are connected to the Internet and collect information; you can check the stock of your fridge and make (semi)-automatic order of missing items.
- On the other hand, they can be security holes:
 - These 'things' look like ordinary electric appliances and don't seem to require security protection
 - Due to the price (& other) restrictions, resources are limited (e.g. CPU, power supply); methodologies for resource-rich devices may not be applicable (e.g. encryption).

Topics of Recent Students (2)

μ -architecture Effect on Performance and Power

- x86 (desktop and servers) and ARM are two most dominant CPU architectures.
- *Basically*, programs written for an architecture should run on any platforms on the same architecture (binary compatibility).
- However, there are many different hardware-implementations of the same architecture (microarchitectures), which result in variations in the performance and power consumption.
- A student studied the effect of microarchitectures using a standard benchmark programs.

Suggested Research Topics (1)

Linux File Systems

- File systems store various and huge amount of information, such as programs, user data, system configuration.
- In addition to the increasing capacity (amount of stored information), various requirements are emerging: speed (latency and throughput), reliability, flexibility, ...
- Study the designs of current file systems and identify the issues for further improving the file systems.
- Study the types and characteristics of the workload against file systems.

Suggested Research Topics (2)

Inter-Domain Communication in Virtualized Systems

- Multiple independent “machines” can be accommodated on a single platform (virtual machines, VMs, or domains).
- When multiple VMs form a large system (multi-tier system), communication between VMs takes place.
- Inter-domain communication goes through different paths than that of physical machines (NIC, network switch, LAN cable..)
- Investigate the inter-domain communication overhead, and relate it to the behavior of the applications and configurations of the VMs.

Suggested Research Topics (3)

Hardware Acceleration of Java Virtual Machine

- Java programs (source files) are compiled into an abstract machine instructions, Java Bytecodes.
- The, Java Bytecodes are either interpreted or compiled by the CPU of the system executing the Java application (Java Virtual Machine, or JVM).
- JVM has advantages, such as platform independence, but some operations are inefficient.
- With a programmable hardware platform (e. g. FPGA), we can design a module to which inefficient operations can be offloaded.

Suggested Research Topics (4)

Heterogeneous Multi-Core Systems

- Multi-core CPUs are ubiquitous: even your smart phones should have dual or quad-core CPUs.
- Also, in addition to the main (general-purpose) CPUs, GPUs are included for faster-graphic processing
- Another type of multi-core CPUs are emerging: heterogeneous-microarchitecture. Example of commercial product: ARM big.LITTLE
- Cores have the same ISA (\approx can run the same machine code programs), but implementations are different. The difficult (but worth investigating) part is how to assign a right job to a right core.

Suggested Research Topics (5)

Training ML Model with Edge AI Devices (Still Immature)

- AI Inference Chips/Devices for Edge Computer are available (e.g Intel NCS2 or Google Edge TPU)
- They are inexpensive and power-efficient, but can only be used for inference using pre-trained models.
- ML models are usually trained with high-performance (& power and money hungry) GPUs, but the Edge AI chips can still be used for Imprinted Weights training.
- There are several restrictions and limitations in this method: “the last layer needs to be calculated on CPU, not NPU/TPU,” or “INT \iff FP conversions.”

References

- Group Website: <http://www.oslab.biz> , follow links to Public Area for the these of past students
- Open Campus website: <http://opencampus.oslab.biz> ; a bit old but written in plain Japanese for general public.
- Research page:
<http://www.u-aizu.ac.jp/~hitoshi/RESEARCH/>
publication and other research activities.
- Posters outside the lab (Research Quadangles 241-E).



QR Codes for above pages