System Utilization of Processor Hardware Counters

Kai Shen

Dept. of Computer Science, University of Rochester

Processor Hardware Counters

Hardware counter metrics (Example: Intel Xeon processors)

NONHALT_TICKS	Num. of ticks that CPU is in non-halt state
INSTRCTN_RTD	Num. of retired instructions
UOPS_RETIRED	Num. of retired micro-ops
L1_MISS_RTD	Num. of L1 cache misses due to retired accesses
L2_MISS_RTD	Num. of L2 cache misses due to retired accesses
L2_MISS	Num. of L2 cache misses
L2_REFERENCE	Num. of L2 cache references
DTLB_MISS_RTD	Num. of data TLB misses due to retired accesses
PGWKMISS_DTLB	Num. of page walks that page miss handler performs due to data TLB misses
DELIVER_MODE	Num. of cycles in trace cache deliver/build modes
TRACECACHE_MS	Num. of trace cache lookup misses
PGWKMISS_ITLB	Num. of page walks that page miss handler performs due to instruction TLB misses
FSB_DATAREADY	Num. of Data Ready and Data Busy events that occur on the front side bus
BUSACCES_CHIP	Num. of transactions on the bus issued by chip
BUSACCES_ALL	Num. of transactions on the bus (may be issued by chip or DMA)
MACH_CL_COUNT	Num. of entire pipeline being cleared
X87_FP_UOP	Num. of X87 float point micro-ops
MEM_CANCEL	Num. of canceled requests in the Data Cache Address Control Unit
UOPQ_W	Num. of valid micro-ops written to the micro-op queue
RES_STALL	Num. of stalls in the Allocator
MISPRED_BRANCH	Num. of mis-predicted branches
RTD_MISPRED_BRANCH	Num. of retired mis-predicted branches
BRANCH	Num. of branches
FRONT_END_EVENT	Num. of load/store micro-ops

Processor Hardware Counters

- Counter registers
  - metrics must map to limited number (2, 4, … 18) of counter registers for observation
- How to access them?
  - counter register access instructions are privileged (at least on Intel Xeon processors)
  - but OS may choose to expose statistics through system calls

Utilization as Performance Indicators

- Context (hardware resource sharing):
  - on multi-core and multi-threading processors, simultaneous processes compete for shared resources (memory bus, cache, floating point unit)
- Hardware counters indicate execution performance:
  - Case #1: too much competition on floating point unit when processes A/B run simultaneously
    ⇒ schedule them not to run simultaneously.
  - Case #2: inefficient/unfair use of cache – process A misses a lot while process B doesn't miss at all
    ⇒ adjust cache size allocation between them.

Utilization as Workload Signatures

- Identifying requests for server workloads

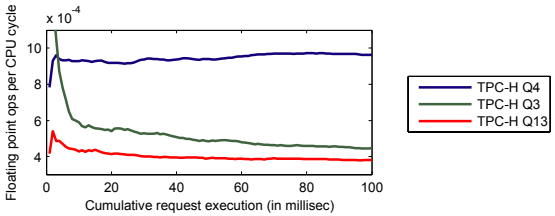

- On-the-fly: identify a request while it still executes
- Utilizations:
  - Predicting request properties to guide OS adaptations
  - Classifying requests on-the-fly to detect anomalies

Utilization: Shortest-Job-First Scheduling

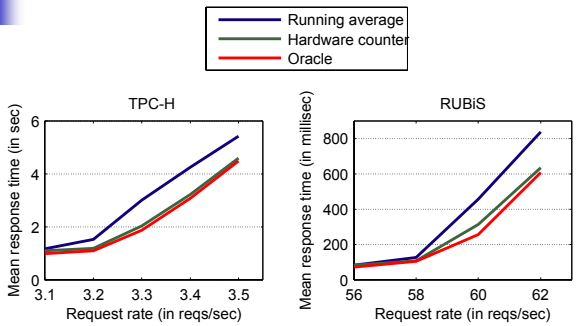

- 15-27% shorter response time than running average
- perform similar to oracle

Utilization: Request Classification and Anomaly Detection

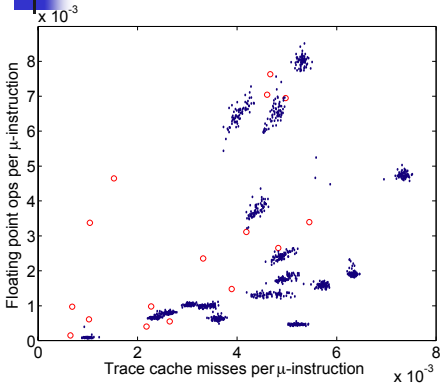

- Dots are normal TPC-H requests
- Circles are anomalies (SQL injection attacks)

Security Risk: Stealing RSA key in OpenSSL

- In OpenSSL/RSA, X^d is decomposed into a series of square and multiply operations.
  - e.g., $X^{11} = ((X^2)^2*X)^2*X$
- Execution sequence of squares and multiplys can help infer the RSA key d.
- Distinguish square and multiply through hardware counters.

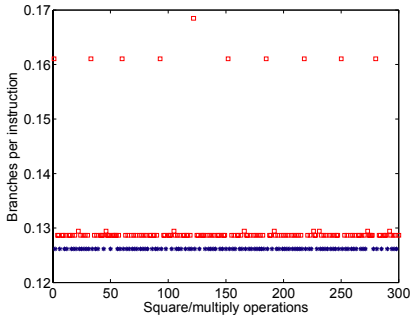

Security Issues

- Side-channel attacks
  - attacks based on machine hardware information collected during the physical execution of a cryptosystem
- Covert-channel attacks
  - a program encodes information into a series of executions with easily recognizable hardware metrics that the collaborator can learn
- Solution: prevent attackers from reading hardware metrics of targets
  - some metrics are reported in a combined fashion on a multi-processor