

Process-level Power Estimation in VM-based Systems

Maxime Colmant, Mascha Kurpicz, Pascal Felber, Loïc Huertas, Romain Rouvoy, Anita Sobe



Research session 5: Virtualization, 23 April, 10:50 am.

Motivation

Problem

- Massive power consumption by data centers
- Hard to identify the largest power consumers
- Cannot attach a power meter to a VM
- Current approaches not fine-grained enough

Vision

- Process-level power estimation in VM-based systems
- Support for all CPU features
- Support for multi-tenant virtualization

Metrics

Hardware (HW) Performance Counters

- Representative and accurate metrics
- Mostly available on modern processors

Criteria selection

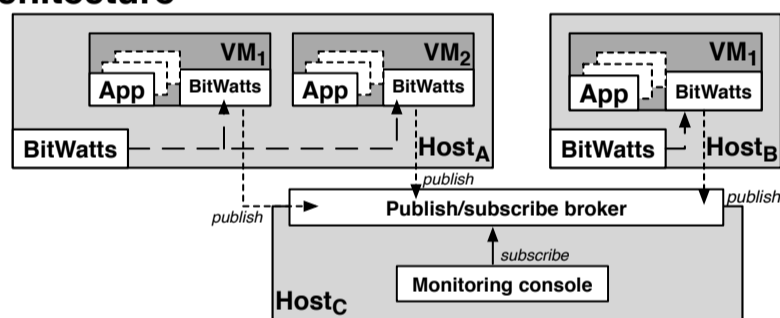
- Counter availability per CPU
- Monitoring overhead
- Best fit under several workloads

Selected HW Performance Counters

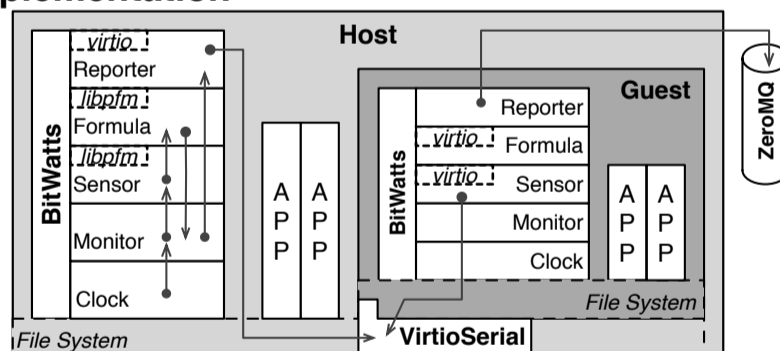
- CPU_CLK_UNHALTED:THREAD_P (uc)
- CPU_CLK_UNHALTED:REF_P

BitWatts

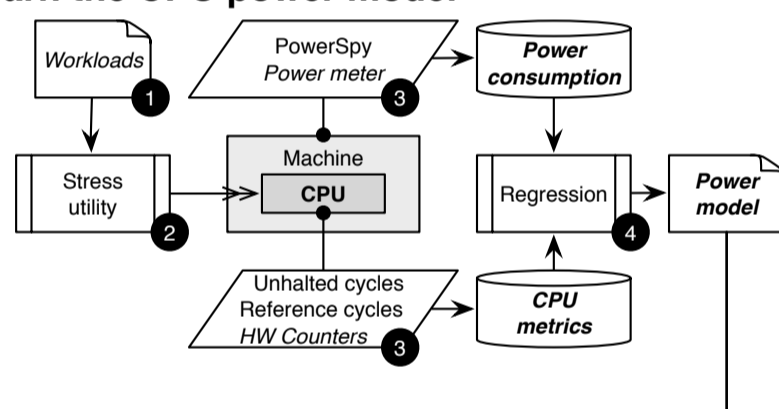
Architecture



Implementation



Learn the CPU power model

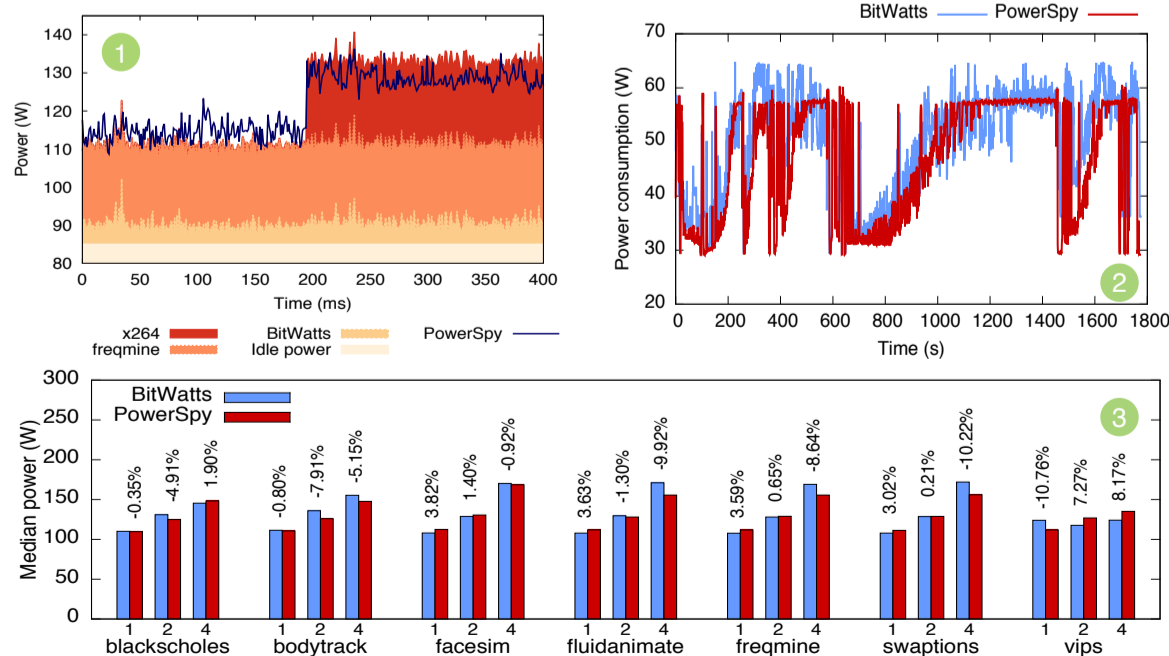


Power Model

$$P_{host}(f) = P_{idle}(f) + \sum_{pid \in PIDs} P_{cpu}(f, uc_{pid}^1 \dots uc_{pid}^N)$$

$$P_{cpu}(f, uc_{pid}^1 \dots uc_{pid}^N) = \sum_{n=1}^N P_f(uc_{pid}^n)$$

Validation



- 1 Process-level power consumption of BitWatts, x264, and freqmine on a Xeon W3520 processor
- 2 Power consumption during the execution of SPECjbb on a i3 2120 with two threads
- 3 Power consumption of the host when scaling PARSEC on multiple VMs

