

Abstract

Objectives

- To provide algorithm designers and performance tuners with a tool for analyzing the relationship between time, energy and power on real systems.
- Directly connect properties of an algorithm with architectural time and energy costs.
- Learn whether or how to redesign algorithms for energy efficiency.

Contributions

- Energy cost model.
- Highly optimized micro-benchmarks.
- Extensive evaluation on latest server-, desktop-, and mobile-class systems.
- New theoretical predictions for algorithm designers and architects to test experimentally.

More information

- J. Choi and R. Vuduc, "A roofline model of energy," *Proc. IEEE Int'l Parallel and Distributed Processing Symp. (IPDPS)*, Boston, MA, USA, May 2013.
- J. Choi and R. Vuduc, "A roofline model of energy," Technical Report GT-CSE-12-01, Georgia Institute of Technology, School of Computational Science and Engineering, Atlanta, GA, USA, December 2012.
- D. Bedard, M. Y. Lim, R. Fowler, and A. Porterfield, "PowerMon: Fine-grained and integrated power monitoring for commodity computer systems," *Proc. IEEE SoutheastCon 2010 (Southeast Con)*, pages 479-484, March 2010.

Experimental Setup

Powermon 2

- Designed and produced by Daniel Bedard and Robert Fowler at RENC1.
- Fine-grained measurement tool for direct current (DC) voltage and current.
- Sampling rate of up to 1024 Hz on 8 ports.

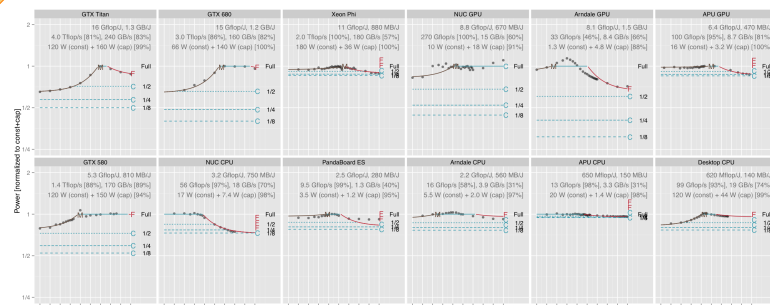
Usage

- Disconnect target device from DC power source.
- Connect the wire from the power source to the input of Powermon.
- Connect the output of Powermon to the target device
- Specify measurement configuration (e.g. sampling rate, read port).
- Read measurements via the USB port. -- OR --
- Use a grad student
- See diagram 1 for sample usage.

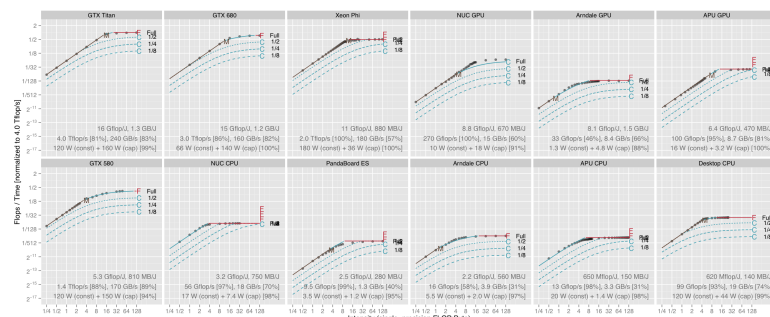


Diagram 1. Grad student connecting Powermon to Arndale development board

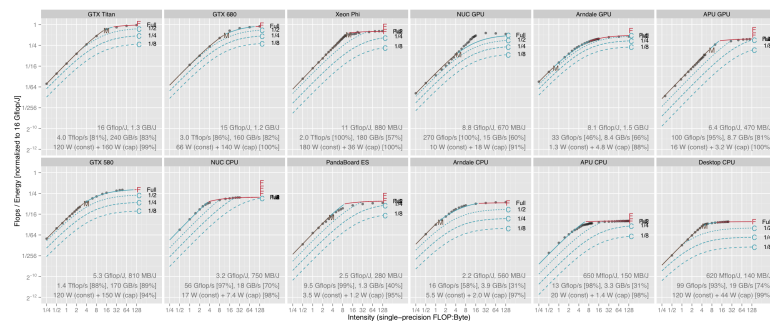
Impact of Power Cap on Time, Energy, and Power



- Hypothetical power as the usable power cap ($\Delta\pi$) decreases. Note the log-log scale, base 2. Each curve represents a power cap setting: full refers to measured max $\Delta\pi$ and $1/k$ refers to a power cap setting of $\Delta\pi/k$.



(a) Performance (Gflop/s)



(b) Energy-efficiency (Gflop/J)

- Hypothetical performance and energy-efficiency as the usable power cap ($\Delta\pi$) decreases.

Power Caps

Accounting for Power Cap

- Starting from the energy cost models from [1] and [2].
- Suppose that on top of π_0 , the system has $\Delta\pi$ additional units of *usable* power. The execution time now becomes,

$$T \equiv \max \left(W\tau_{flop}, Q\tau_{mem}, \frac{W\epsilon_{flop} + Q\epsilon_{mem}}{\Delta\pi} \right)$$

$$T(W, I) = W\tau_{flop} \max \left\{ 1, \frac{B_\tau}{I}, \frac{\pi_{flop}}{\Delta\pi} \left(1 + \frac{B_c}{I} \right) \right\}$$

- Let $B_\tau^+ \equiv B_\tau \max \left(1, \frac{\pi_{mem}}{\Delta\pi - \pi_{flop}} \right)$ and $B_\tau^- \equiv B_\tau \min \left(1, \frac{\Delta\pi - \pi_{mem}}{\pi_{flop}} \right)$

Observe that $B_\tau^- < B_\tau \leq B_\tau^+$

- When $\Delta\pi \geq \pi_{flop} + \pi_{mem}$, there is enough power to run flops and move data at their *maximum* rates
- Otherwise, $[B_\tau^-, B_\tau^+]$ defines the interval in which the system's balance can lie due to the power cap.

- We can now redefine power, P , as $\bar{P} = \bar{P}(I) = \begin{cases} \pi_{flop} + \pi_{mem} \frac{B_\tau}{I} & \text{if } I \geq B_\tau^+ \\ \pi_{flop} \frac{1}{B_\tau^-} + \pi_{mem} & \text{if } I \leq B_\tau^- \\ \Delta\pi & \text{otherwise} \end{cases}$

Variable	Description
W	# of useful compute operations
Q	# of main memory operations
I	Intensity, or W/Q (e.g., flops per byte)
τ_{flop}	Time per work (arithmetic) operation
τ_{mem}	Time per memory operation
B_τ	Balance in time, or τ_{mem}/τ_{flop} (flops per byte)
ϵ_{flop}	Energy per work (arithmetic) operation
ϵ_{mem}	Energy per memory operation
B_c	Balance in energy, or $\epsilon_{mem}/\epsilon_{flop}$ (flops per joule)
π_0	Constant power