

Lecture 1.3

Moore's Law and Dennard Scaling

EN 600.320/420

Instructor: Randal Burns


29 January 2018



Department of Computer Science, *Johns Hopkins University*

We've been duped!

- Moore's law
 - The number of transistors that can be *inexpensively placed* on an integrated circuit is increasing exponentially, doubling approximately every two years.
 - The observation has held for half a century
- It's true, but not helpful:
 - More transistors has become more cores (independent processing units on the same chip)
- It's true, but not helpful:
 - Pipelined multicore (N k-flop cores) are not as useful as a big (Nk flop) processor
- But the chip vendors tell us we have faster processors
 - So we (the programmers) must write parallel code to make software faster on cores with the same clock speed and number of transistors



<https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/>



Moore's Law

Microprocessor Transistor Counts 1971-2011 & Moore's Law




Image from http://en.wikipedia.org/wiki/File:Transistor_Count_and_Moore%27s_Law_-_2011.svg



Dennard Scaling


As transistors get smaller their power density stays constant so that power is in proportion with area

- voltage and current scale downward
- Performance per watt increases exponentially
 - smaller transistors lead to faster clock rates
- Dennard scaling ended in 2006
 - But Moore's law still alive
 - Turn to multicore processors



Dennard scaling breakdown

- In 2006
- Current leakage and heating



Moore's Law Post Dennard Scaling

- Moore's law -> parallelism -> parallel programming
 - This has already happened and software is just catching up
- Is Moore's law dead?
 - Dark silicon, 5nm
 - <http://www.extremetech.com/computing/165331-intels-former-chief-architect-moores-law-will-be-dead-within-a-decade>
- Moore's law is dead
 - Yes, likely, but not relevant to parallel programmers
 - Even if scaling does not continue, trend toward parallelism will



Discussion

- What are the hardware trends of note?
 - Ubiquitous GPU acceleration
 - Heterogeneous/reconfigurable processing
 - 3-d lithography
 - Reduced precision processors (Google TPP)
- What is the difference between cloud computing and supercomputing?
- What is “exascale” and why do I care?

