NERSC 5

Welcome to

Beyond Petaflops:
Specialized Architectures for Power Efficient Scientific Computing

SIAM Conference on Computational Science and Engineering
Monday, February 19, 2007
Costa Mesa California

John Shalf
LBNL/NERSC

Berkeley Lab

NERSC

Office of Science
U.S. Department of Energy

What is Going on?

- New Constraints
  - Power limits clock rates
  - Cannot squeeze more performance from ILP *(complex cores)* either!
- But Moore's Law continues!
  - What to do with all of those transistors if everything else is flat-lining?
  - Now, #cores per chip doubles every 18 months *instead* of clock frequency!
- **Power Consumption** is chief concern for system architects
- **Power-Efficiency** is the primary concern of consumers of computer systems!

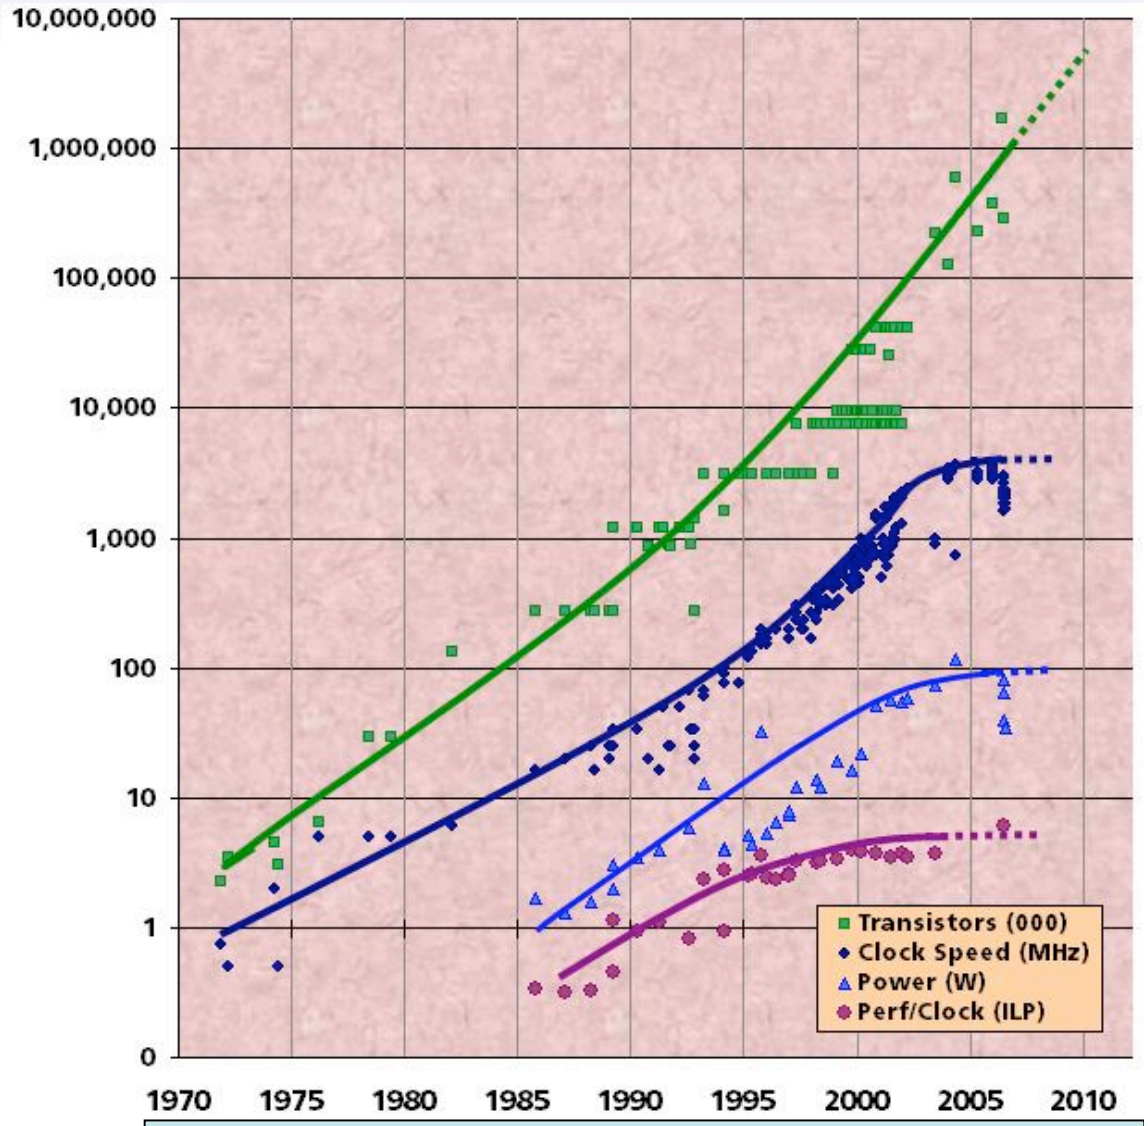

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

Microprocessors: Up Against the Wall(s)

From Joe Gebis: UCB/LBNL

- Microprocessors are hitting a power wall
  - Higher clock rates and greater leakage increasing power consumption
- Reaching the limits of what non-heroic heat solutions can handle
- Newer technology becoming more difficult to produce, removing the previous trend of “free” power improvement

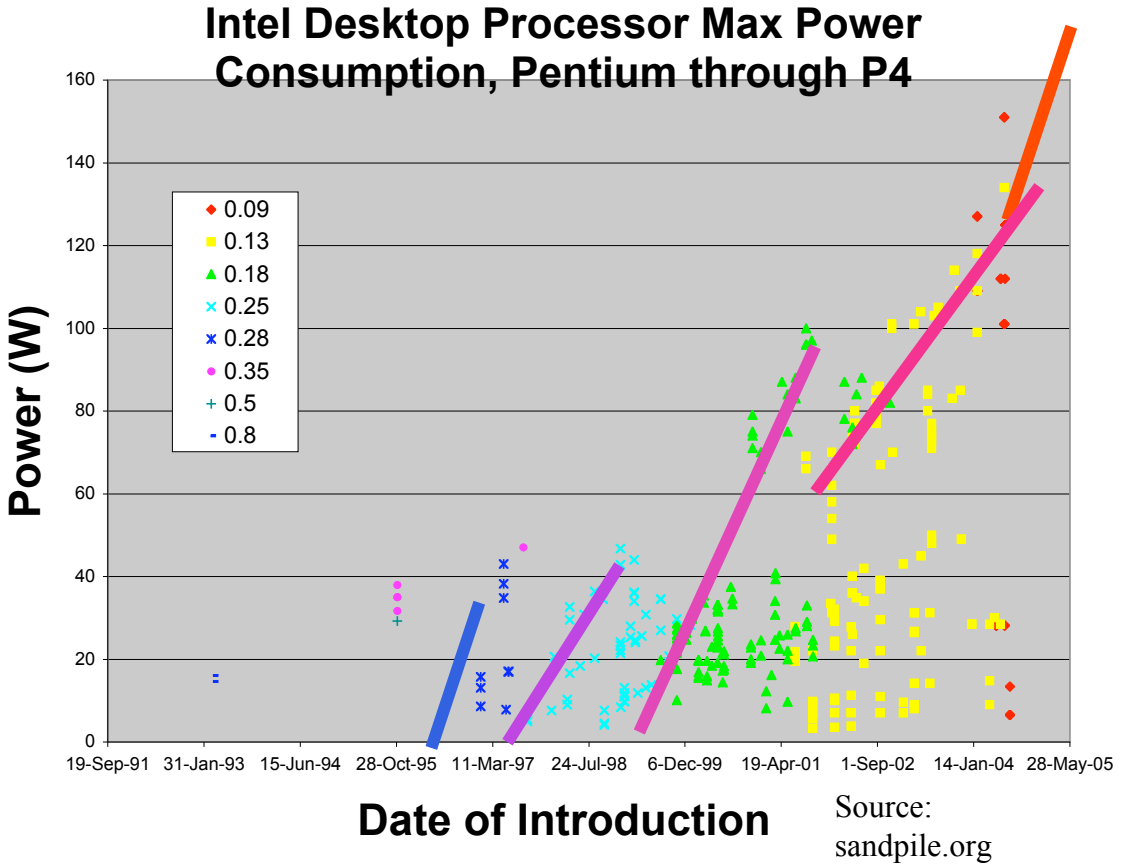

Source: sandpile.org

Example: ORNL Power Cost Projection 2006 - 2011

- Immediate need to add 8 MW to prepare for 2007 installs of new systems
- NLCF petascale system could require an additional 10 MW by 2008
- Need total of 40-50 MW for projected systems by 2011
- Numbers just for computers: add 75% for cooling
- Cooling will require 12,000 – 15,000 tons of chiller capacity

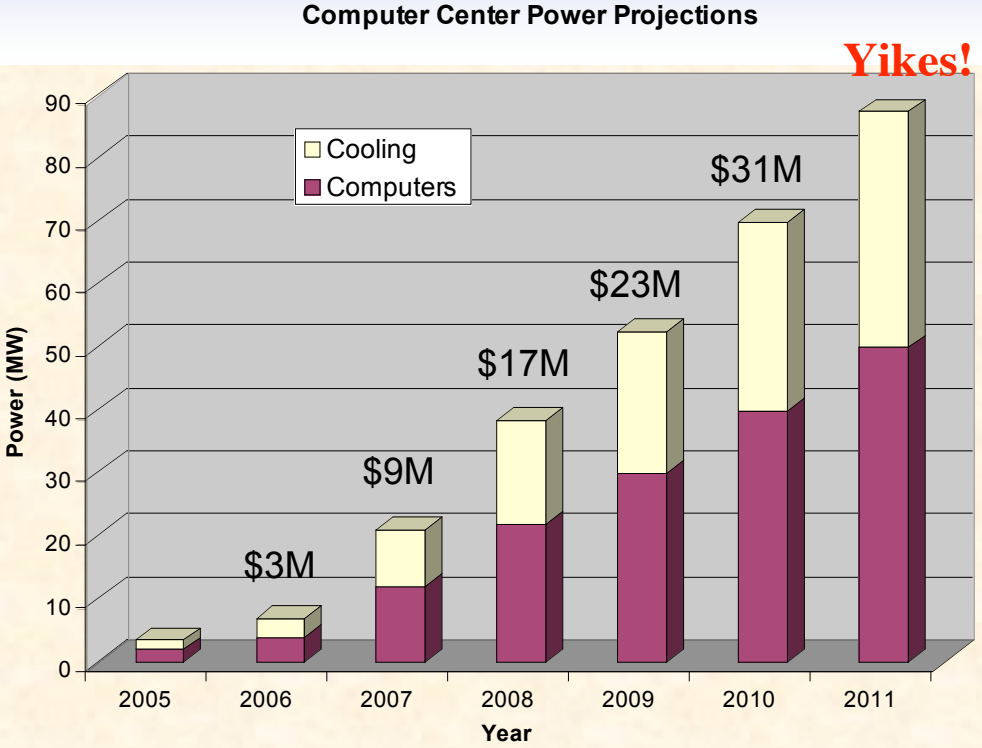

Cost estimates based on $0.05 kW/hr

Annual Average Electrical Power Rates $/MWh

Site	FY 2005	FY 2006	FY 2007	FY 2008	FY 2009	FY 2010
LBNL	43.70	50.23	53.43	57.51	58.20	56.40 *
ANL	44.92	53.01				
ORNL	46.34	51.33				
PNNL	49.82	N/A				

Data taken from Energy Management System-4 (EMS4). EMS4 is the DOE corporate system for collecting energy information from the sites. EMS4 is a web-based system that collects energy consumption and cost information for all energy sources used at each DOE site. Information is entered into EMS4 by the site and reviewed at Headquarters for accuracy.

Common Ground for MS02 Speakers

- When power costs more than capital procurement costs
  - It is cost-effective to explore more customized solutions
  - It is more power-efficient to specialize to the application
  - It is more power-efficient to employ massive parallelism
    - ***See http://view.eecs.berkeley.edu/***
  - Embedded computing market has been doing this for years!
- Start with **Science Requirements** and then design a machine to conform to requirements
  - This as opposed to building a petaflop machine and trying to find things that will run on it
- Using this approach, you can reach aggressive performance goals . . .
  - Far more rapidly than a conventional approach would allow
  - Using a reasonable power budget
  - Using a reasonable development budget