About the instructor

- **Instructor:** Dr. Walter W. Schilling, Jr.
- **Office:** Walter Schroeder Library 335
- **Office Hours:**
  - **MTF 9:00-10:00**
  - **W 10:00-11:00**
  - **R 15:00-16:00**
  - While I post office hours, I keep an open door policy. If I am in my office and the door is open, please feel free to stop in.
- **Telephone:** 414 277 7370
- **E-mail:** schilling@msoe.edu
  - Best method to contact me during non-class days
  - Please prefix subject with SE3821.
- **Course Web Page:**
  - http://www.walterschilling.us/msoe/fall20132014se498/fall2013se498.php

About the Instructor (Continued)

- Ohio Northern University graduate in Electrical Engineering
  - Computer Science Minor
- Masters and PhD. from University of Toledo
  - Specialized in Computer Systems Design and Software Reliability
- Worked in Automotive Industry for approximately 6 years
  - Audio Software Engineer – Embedded Systems Design
    - US Patent 6,707,768
    - “Randomized Playback of Tracks in a Multimedia Player”

- Personal Website: http://www.walterschilling.org

MSOE

Catalog Description

- This course provides students with an introductory overview of parallel and high performance computing. In class, students will be exposed to basic parallel algorithms and parallel thinking, as well as given an overview of parallel computing architectures. Lab activities include students implementing a set of parallel algorithms using OpenMP, MPI, OpenAcc, and other appropriate technologies. As part of the lab activities, students will experimentally analyze the impact of parallelization. A small independent project will complete the course. Limited domain exposure to matrix math calculations, image processing, and DNA sequencing will also be provided.

Outcomes

- 1. Justify the need for parallel program and algorithm development
- 2. Explain the different parallel computing architectures
- 3. Calculate performance metrics related to the parallelization of elementary programs
- 4. Construct simple applications using openMP, MPI, and openACC technology

Prerequisites

- CS3841
- Junior Standing

C/C++ programming

Textbook

- *Peter S. Pacheco, An Introduction to Parallel Programming*
- Textbook will occasionally be supplemented with articles and other handouts.

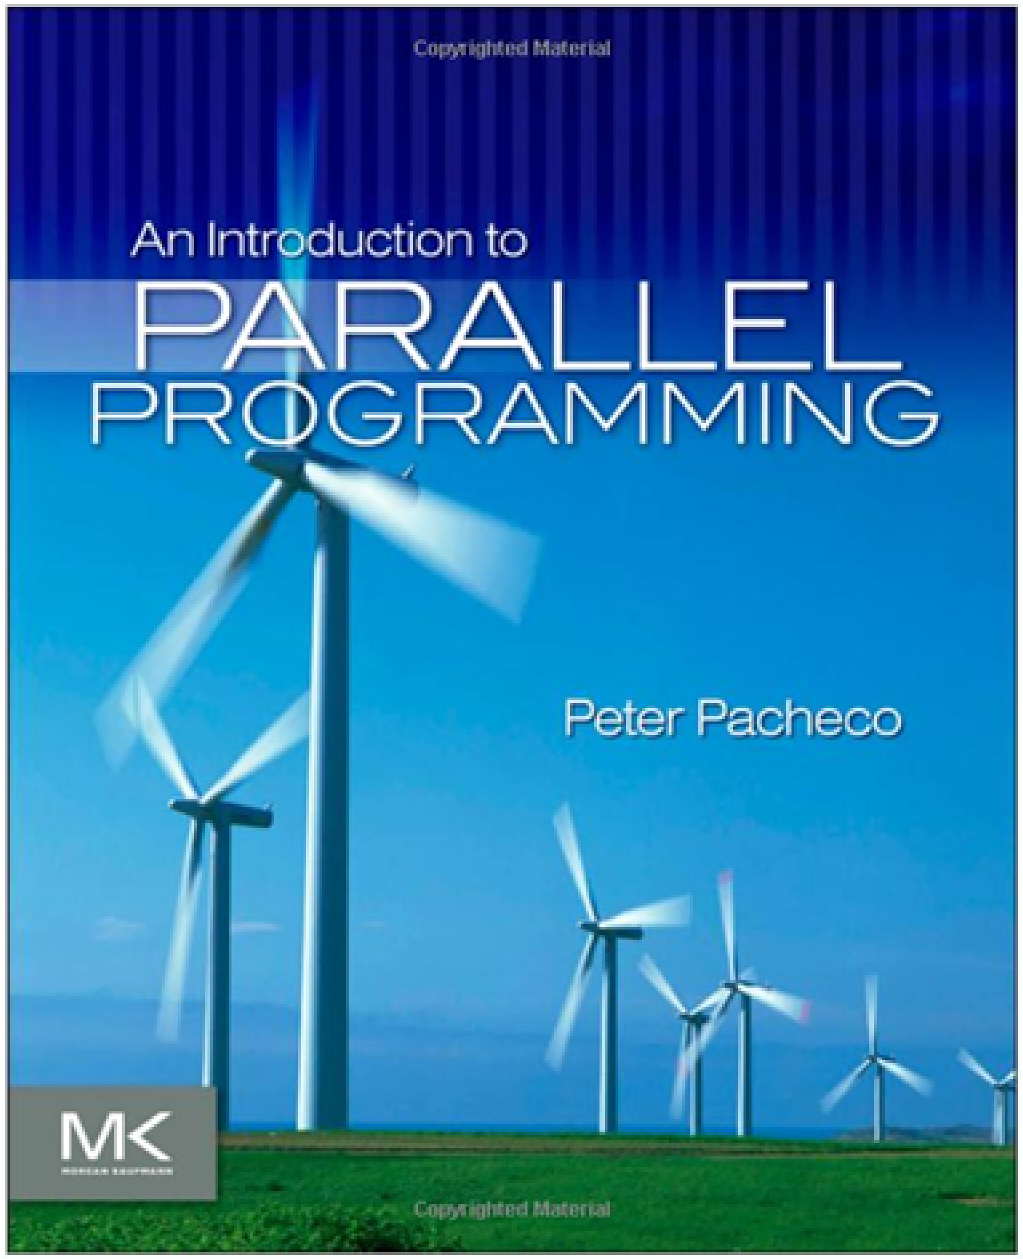

Class Materials

- Textbook
- Laptop Computer with standard MSOE image
  - You are responsible for completing the assignments on time, even if your computer malfunctions. No extensions will be granted due to computer problems.

Student Integrity

- All students are expected to abide by MSOE's policy on student integrity. If at any point in the semester you have a question about an assignment, please come discuss it with me.
- Violations of this policy will be dealt with seriously, and may result in significant penalty, up to and including failure of the course.
- You are expected to abide by the terms of the usage agreement with the Ohio Supercomputer Center for ethical usage of the resources.

Lecture Notes / Handouts

- Lecture notes and handouts may be made available on my course website.
  - These are for your own personal usage and are not to be circulated outside of the MSOE domain.
- Lecture notes and handouts are subject to copyright law.

General Notes

- **Laptop Usage**
  - Laptops may be used in class to view the lecture presentations live in Ubiquitous Presenter.
  - Random responses may be required at various times using Ubiquitous Presenter
- **Class Participation and Activities**
  - Class attendance is mandatory.
  - Attendance is required at all lab sessions.
    - You must stay for the entire lab period unless the current assignment is complete and turned in.

Etiquette

- Cell Phones
  - Please turn off all cell phones and pagers during class.
    - Cell phones can be disruptive to the professor as well as fellow students.
  - No text messaging in class!

- If you must miss a class or lab for an acceptable reason, please let the professor know in advance and follow-up afterwards.

Course Coverage

- See Syllabus

SE498 Parallel Computing
Lecture 1: Introduction

Lecture Objectives:

1) Explain the concept of Moore's law
2) Define high performance computing
3) Calculate the power for a given microprocessor at a given voltage and a given frequency
4) Explain the relationship between voltage, frequency, and dissipated power.

Power

$$W = V \times I$$

$$W = V^2 / R$$

W=> Power (Watts)
V=> Voltage
I=> Current
R=> Resistance

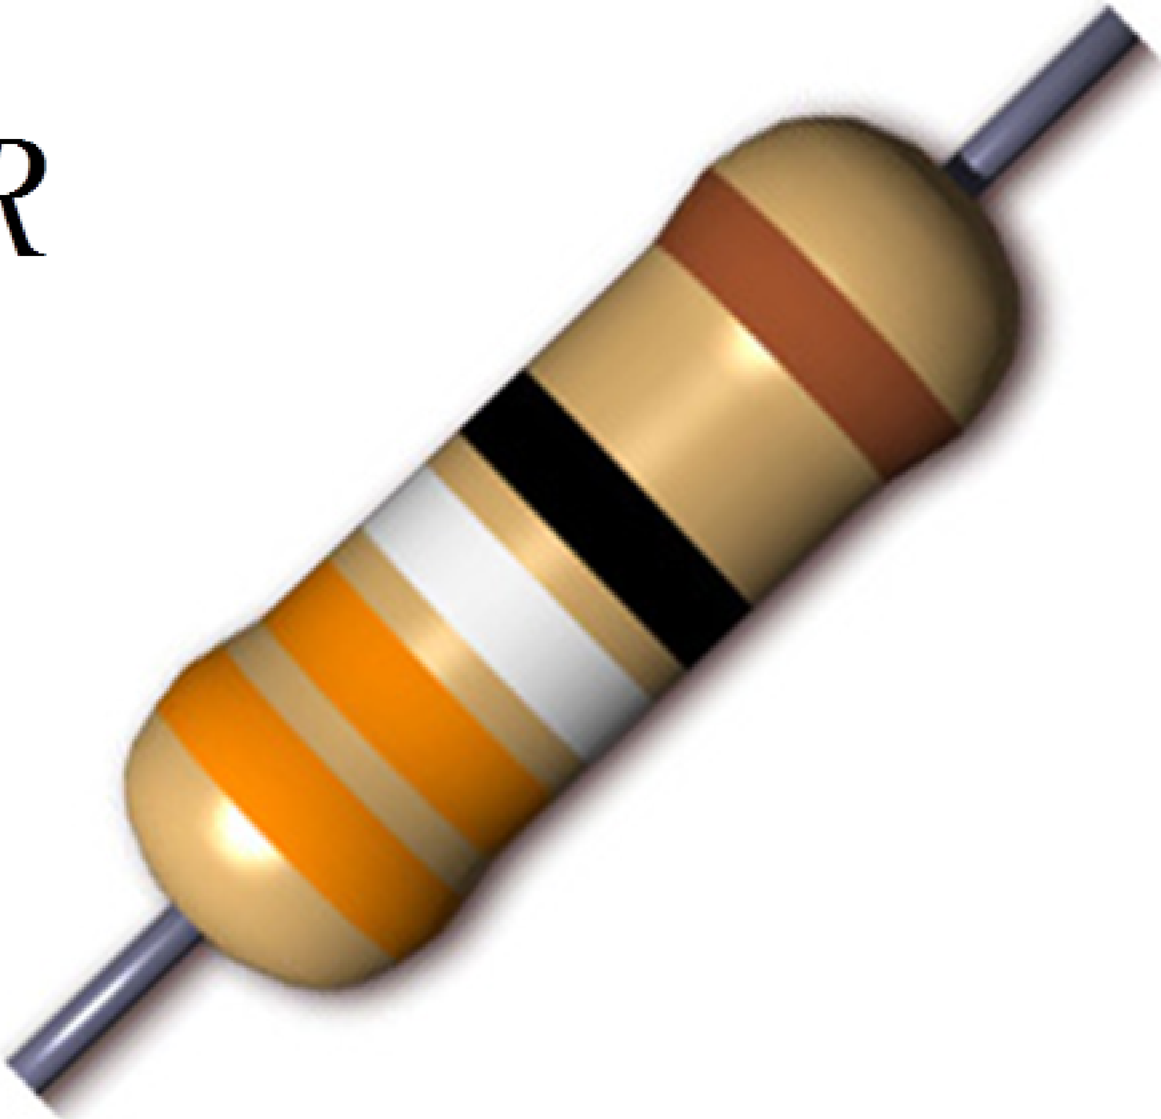

What is power

- Power is the energy used by an electronic device to do work
  - For computers, it's byproduct is heat.

Demo part 2

- Using a modified SunROM board, measure the power used to run the board at 1Mhz, 4 Mhz, and 16 MhZ and between 2 and 5 volts.
- Sample code for the program is available.

ATMEGA 32 Current versus Crystal Frequency

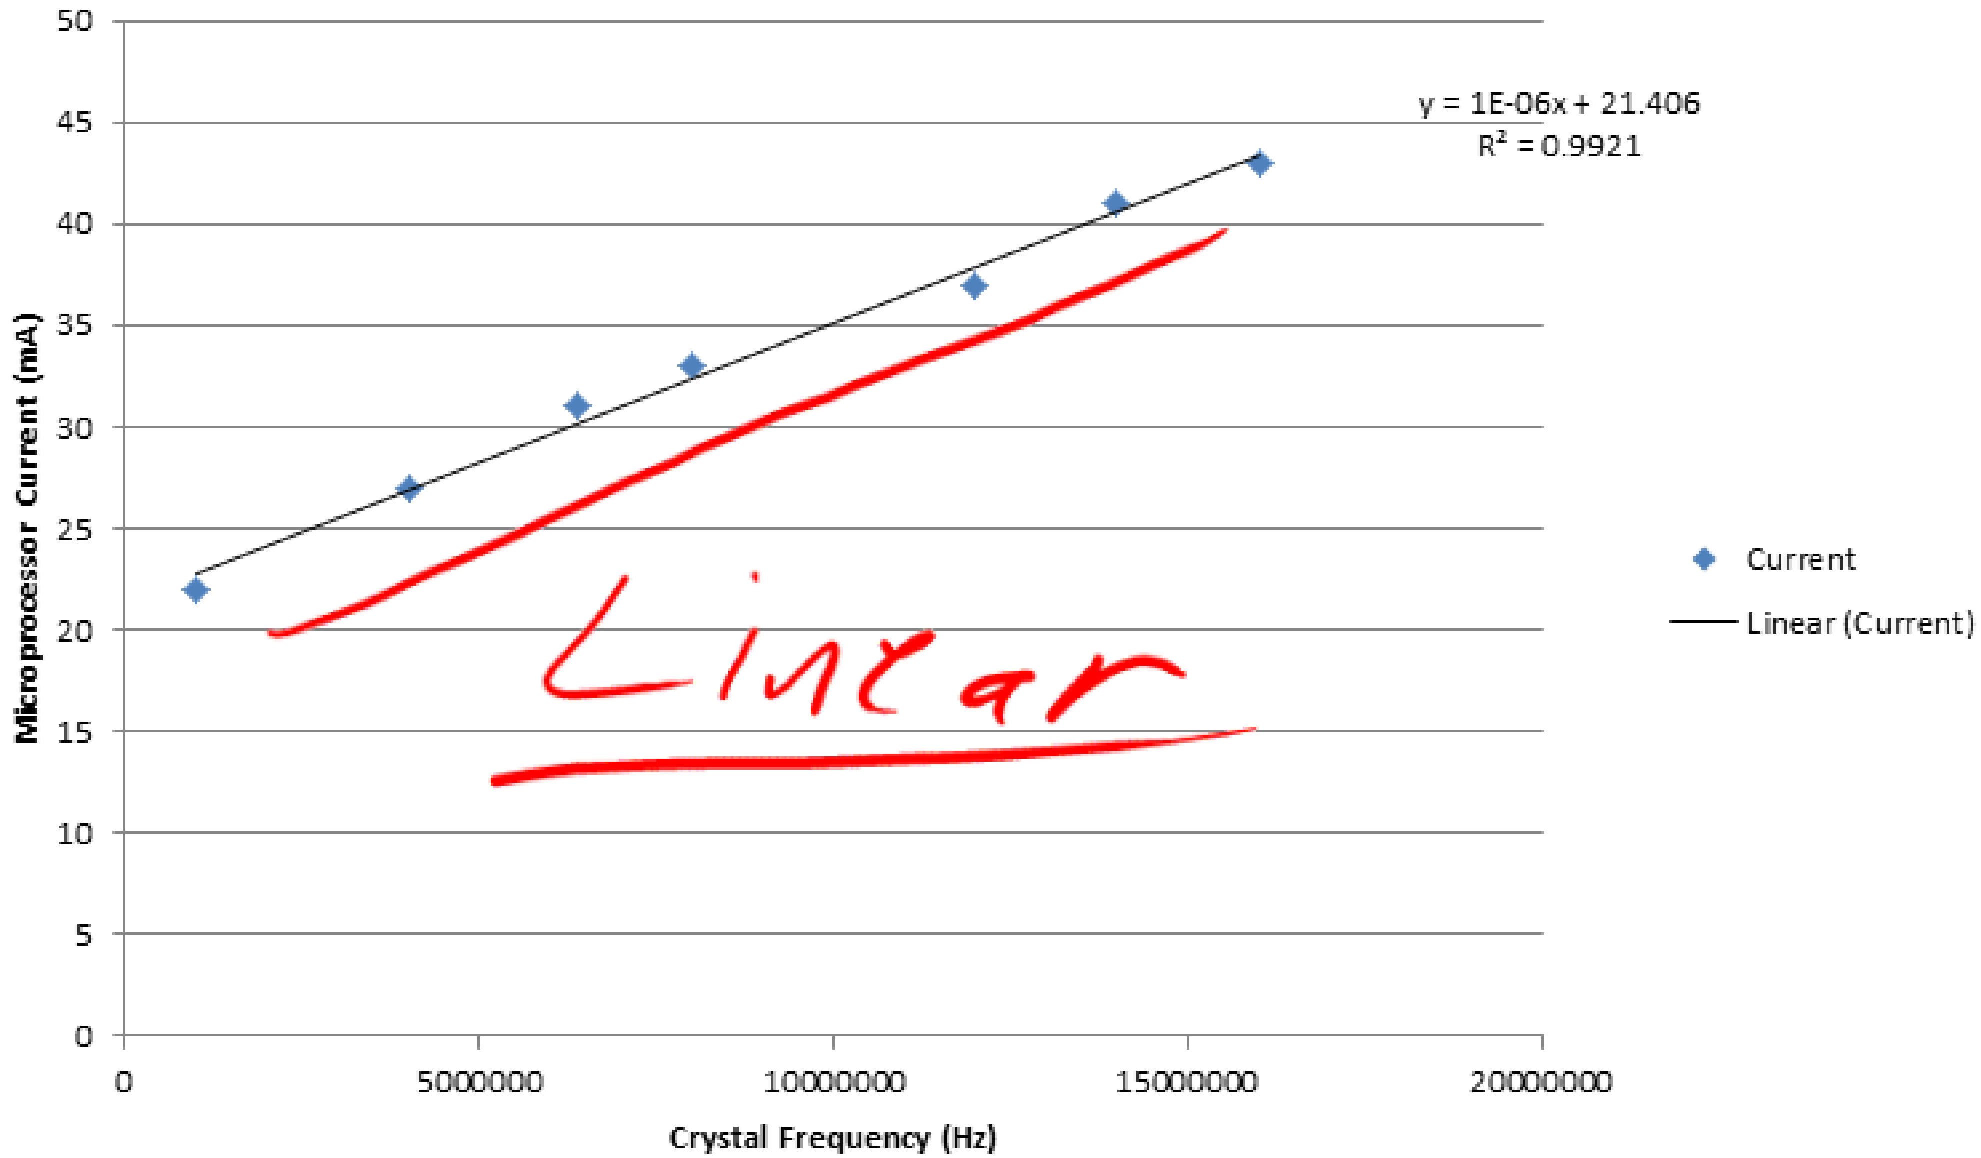

Instructor's Data (11/14/2012)

Increase freq

Instructor's Data (11/14/2012)

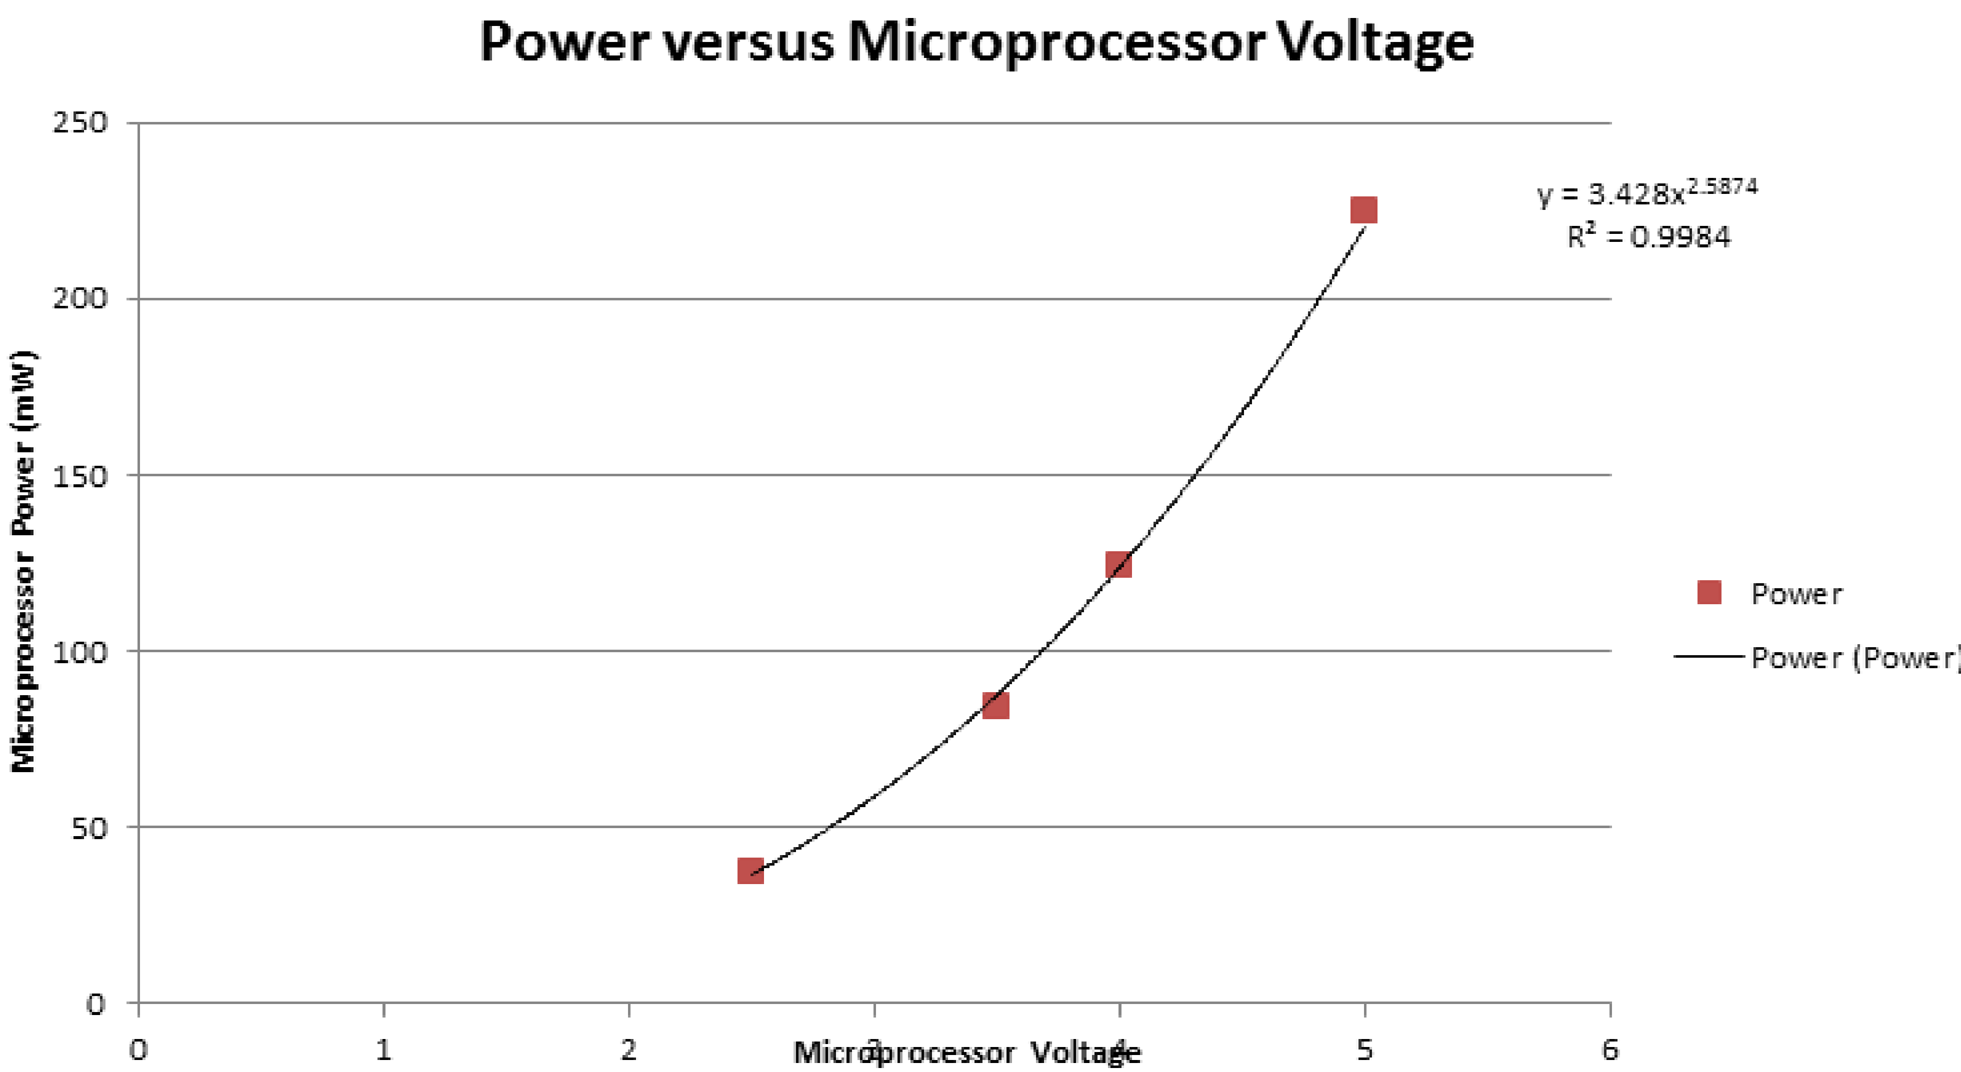

Exponential

Power Calculation

$$Power = CapacitiveLoad \times Voltage^2 \times ClockFrequency$$

- What will cause power to go up?

Why do we care?

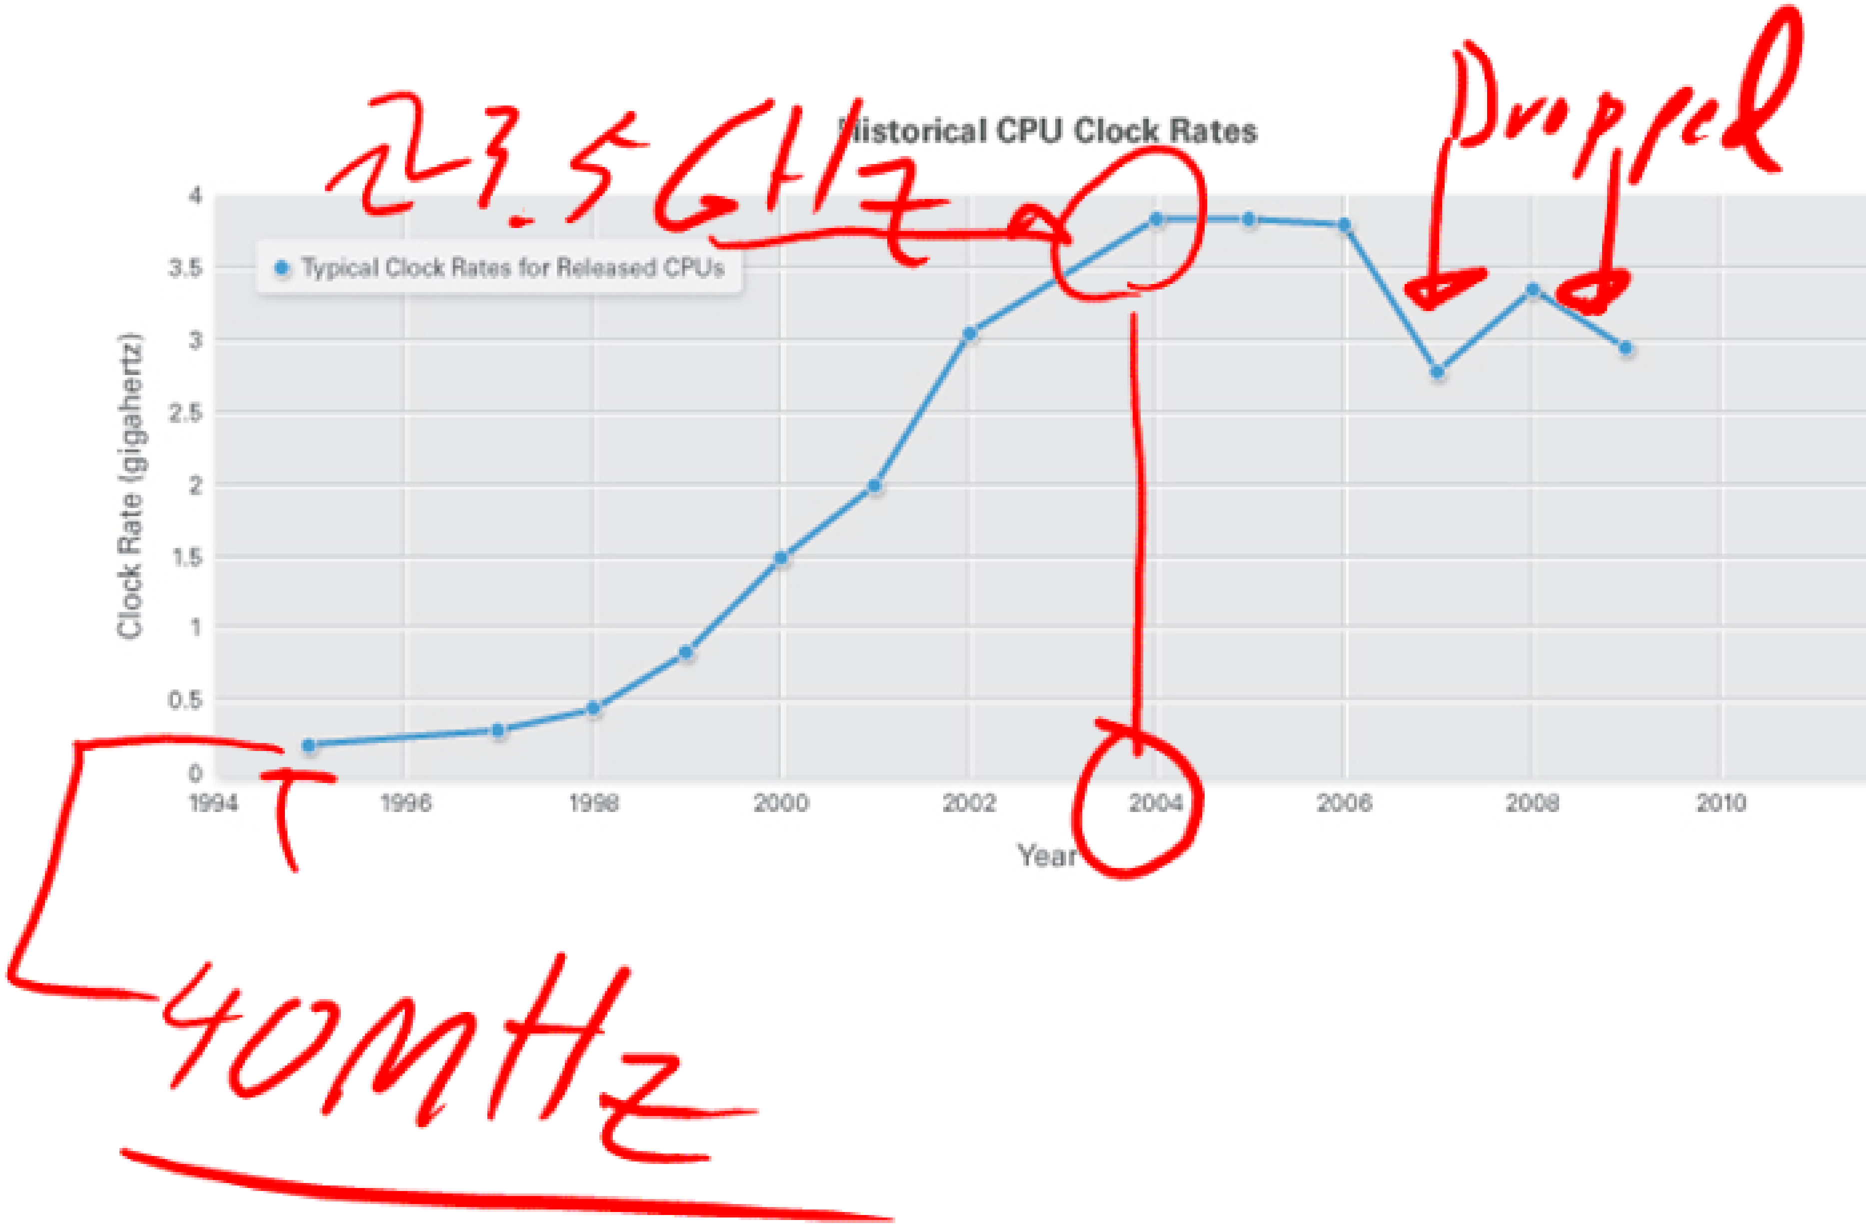

[*http://www.ni.com/white-paper/14565/en/](http://www.ni.com/white-paper/14565/en/)

Gordon Moore - Observations

- Gordon Moore worked for Fairchild Semiconductors
- He noticed a trend in IC manufacture
- Every 2 years the number of components on an area of silicon doubled*
- He published this work in 1965 – known as Moore's Law
- His predictions were for 10 years into the future
- His work predicted personal computers and fast telecommunication networks

* Sources vary regarding time period

Graph of Moore's Law

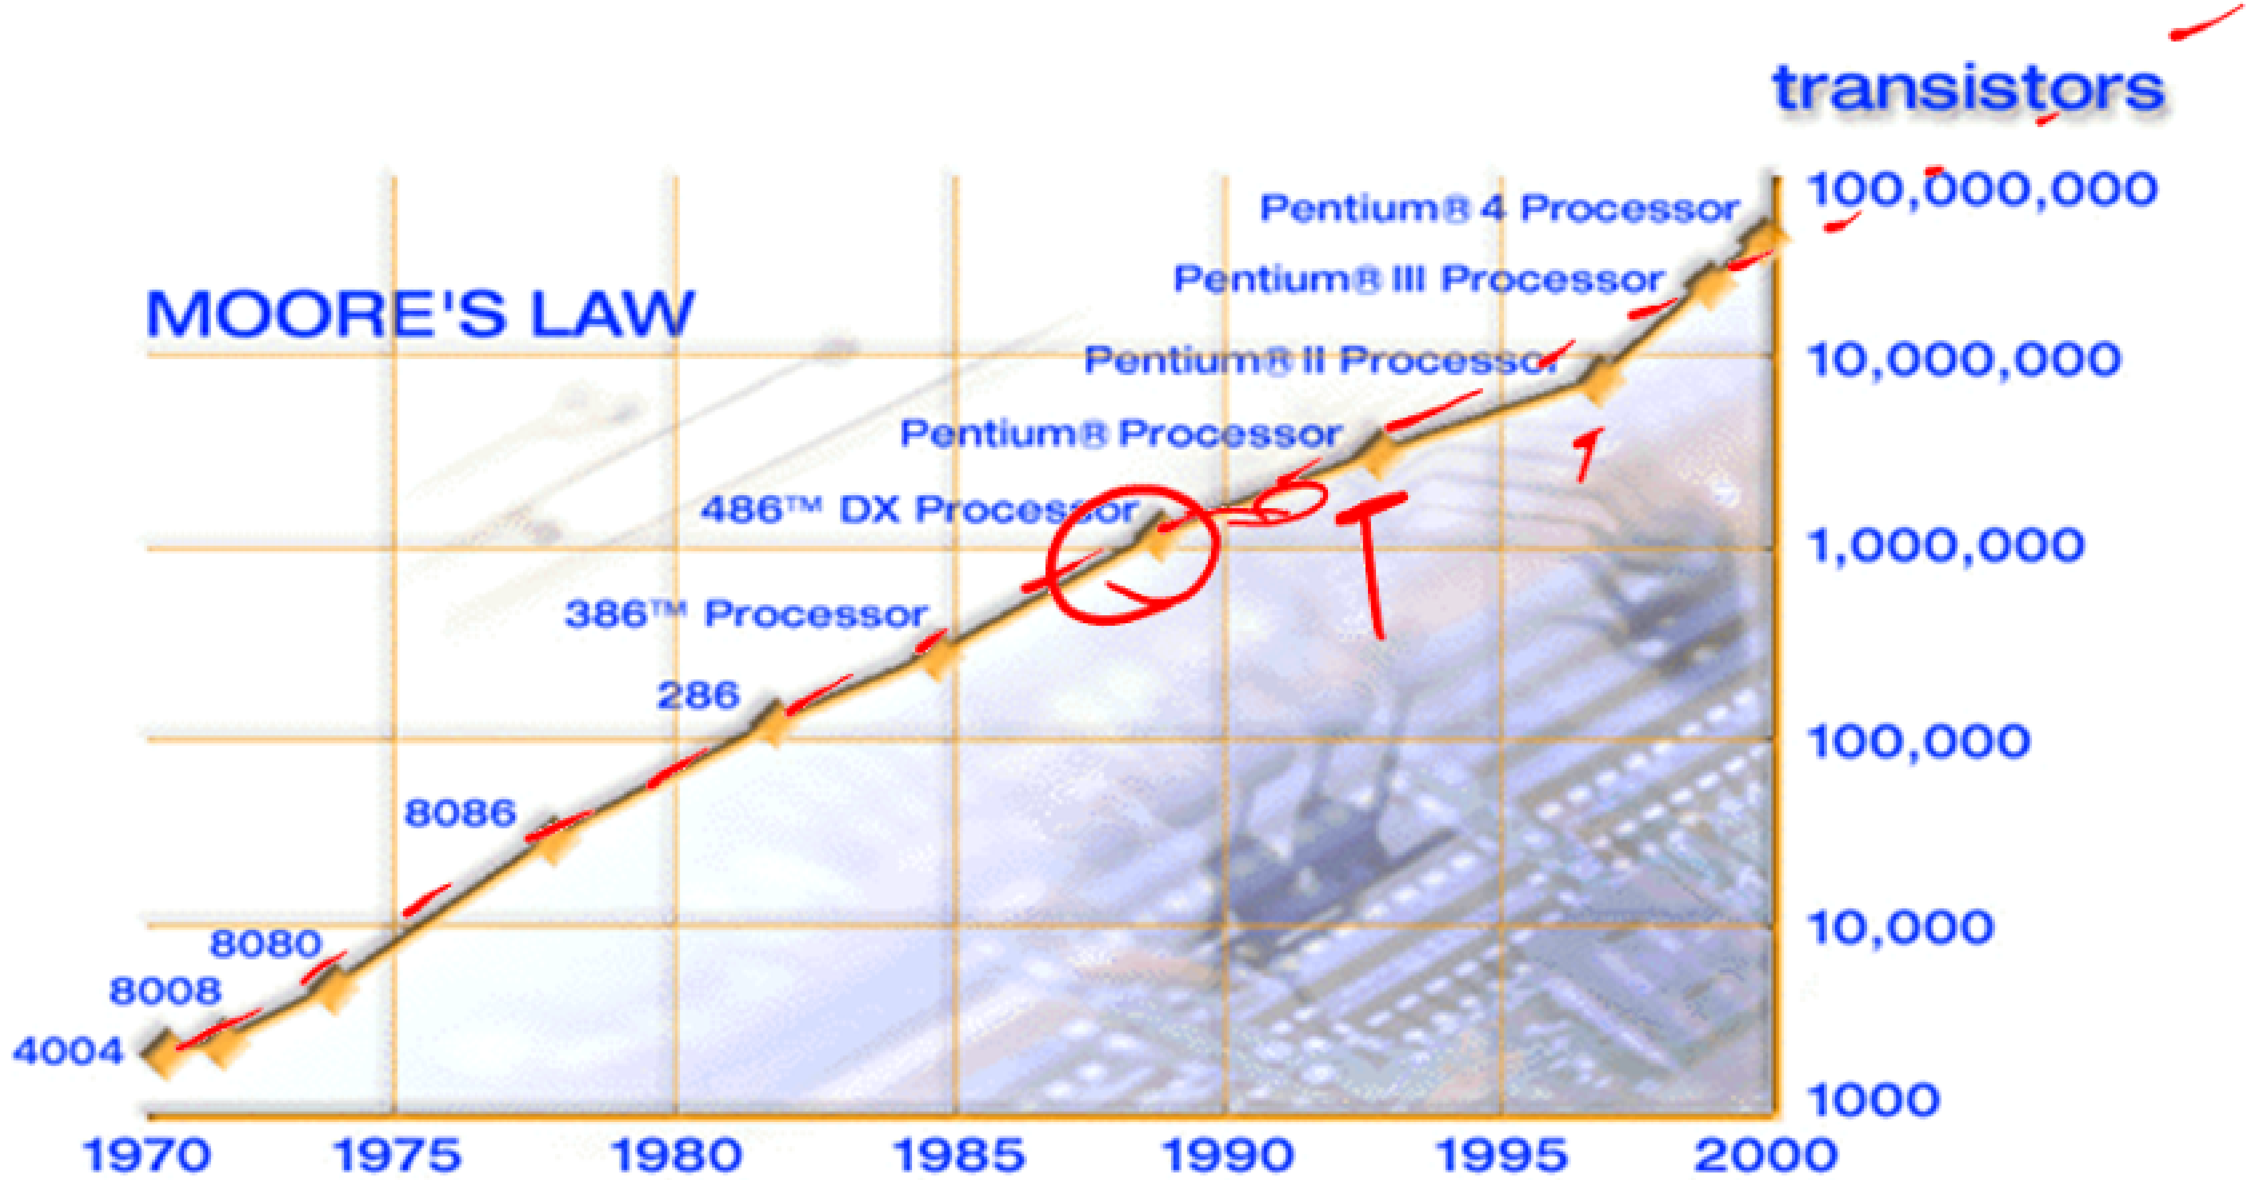

History of Supercomputing

- Seymour Cray
  - 1964: CDC 6400
  - 1976: Cray 1 (pictured)
  - 1982: Cray XMP
- 1987: Connection Machine 2 (CM-2)
  - 65,000 processors (pictured)
- New architectures
  - Vector Processing
  - Massively Parallel
  - Cluster Processing
- New technologies
  - Solid-state transistors
  - Liquid cooling (Fluorinert)
  - Alternatives to silicon

More History: Cluster Processing

- 1991: Eugene Brooks at LLNL writes "Attack of the Killer Micros"
  - Microprocessors (386, 486) became as fast as supercomputers of the day!
- Cluster: a collection of computers used to solve the same problem
- Clusters allowed supercomputers to scale out
  - 1995-2010: A good 15-year run of increasing performance through ever-increasing cluster size
  - Now, constrained by cost, size, power

The Future of Supercomputing

- 2004: Fastest desktops clocked near 3 GHz
- 2013: Fastest desktops clocked near 3 GHz
  - Power consumption keeps us from raising clock rates
- 2005-2010: Multicore (2, 4, 6, 8-cores/chip)
- 2010-on: Manycore (GPUs and accelerators)

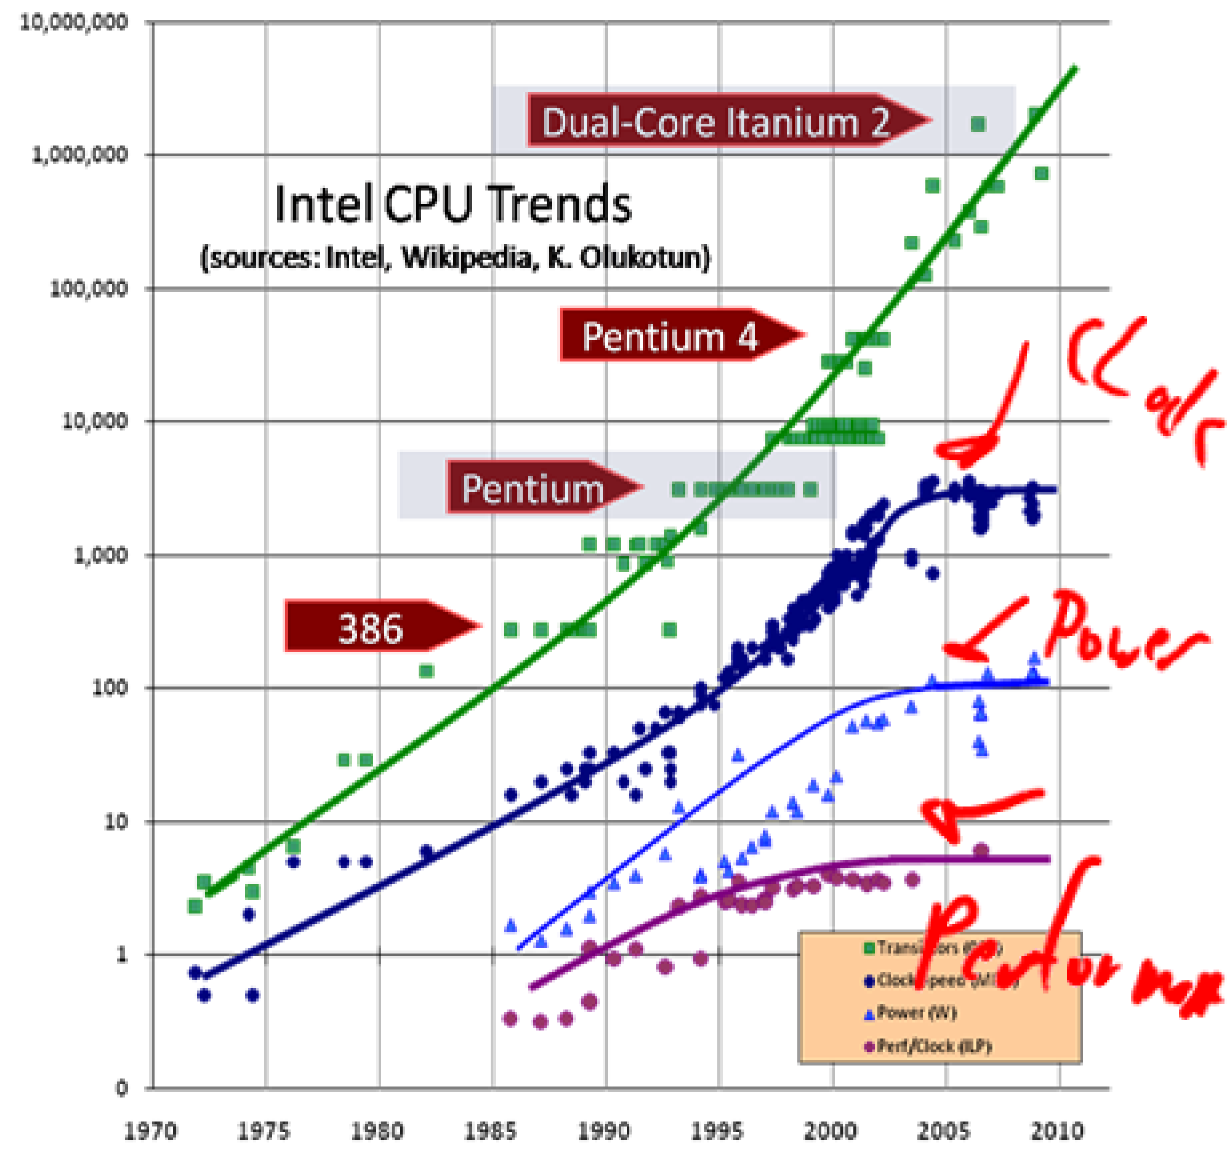

http://www.gotw.ca/publications/concurrency-ddj.htm

Top 500 Machines from top500.org

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	National University of Defense Technology	NUDT	Tianhe-2 NUDT TH-IVB-FEP, Xeon 12C 2.2GHz, IntelXeon Phi	China	3,120,000	33.9	17.8
2	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Opteron 16C 2.2GHz, Gemini, NVIDIA K20x	USA	560,640	17.6	8.21
3	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	17.2	7.89
4	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	795,024	10.5	12.7
5	Argonne National Laboratory	IBM	Mira BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	786,432	8.59	3.95
6	Texas Advanced Computing Center/UT	Dell	Stampede PowerEdge C8220, Xeon E5 8C 2.7GHz, Intel Xeon Phi	USA	462,462	5.17	4.51
7	Forschungszentrum Juelich (FZJ)	IBM	JuQUEEN BlueGene/Q, Power BQC 16C 1.6GHz, Custom	Germany	458,752	5.01	2.30
8	Lawrence Livermore National Laboratory	IBM	Vulcan BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	393,216	4.29	1.97
9	Leibniz Rechenzentrum	IBM	SuperMUC iDataPlex DX360M4, Xeon E5 8C 2.7GHz, Infiniband FDR	Germany	147,456	2.90	3.52
10	National SuperComputer Center in Tianjin	NUDT	Tianhe-1A NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2.57	4.04

Top 500 by country (November 2012)

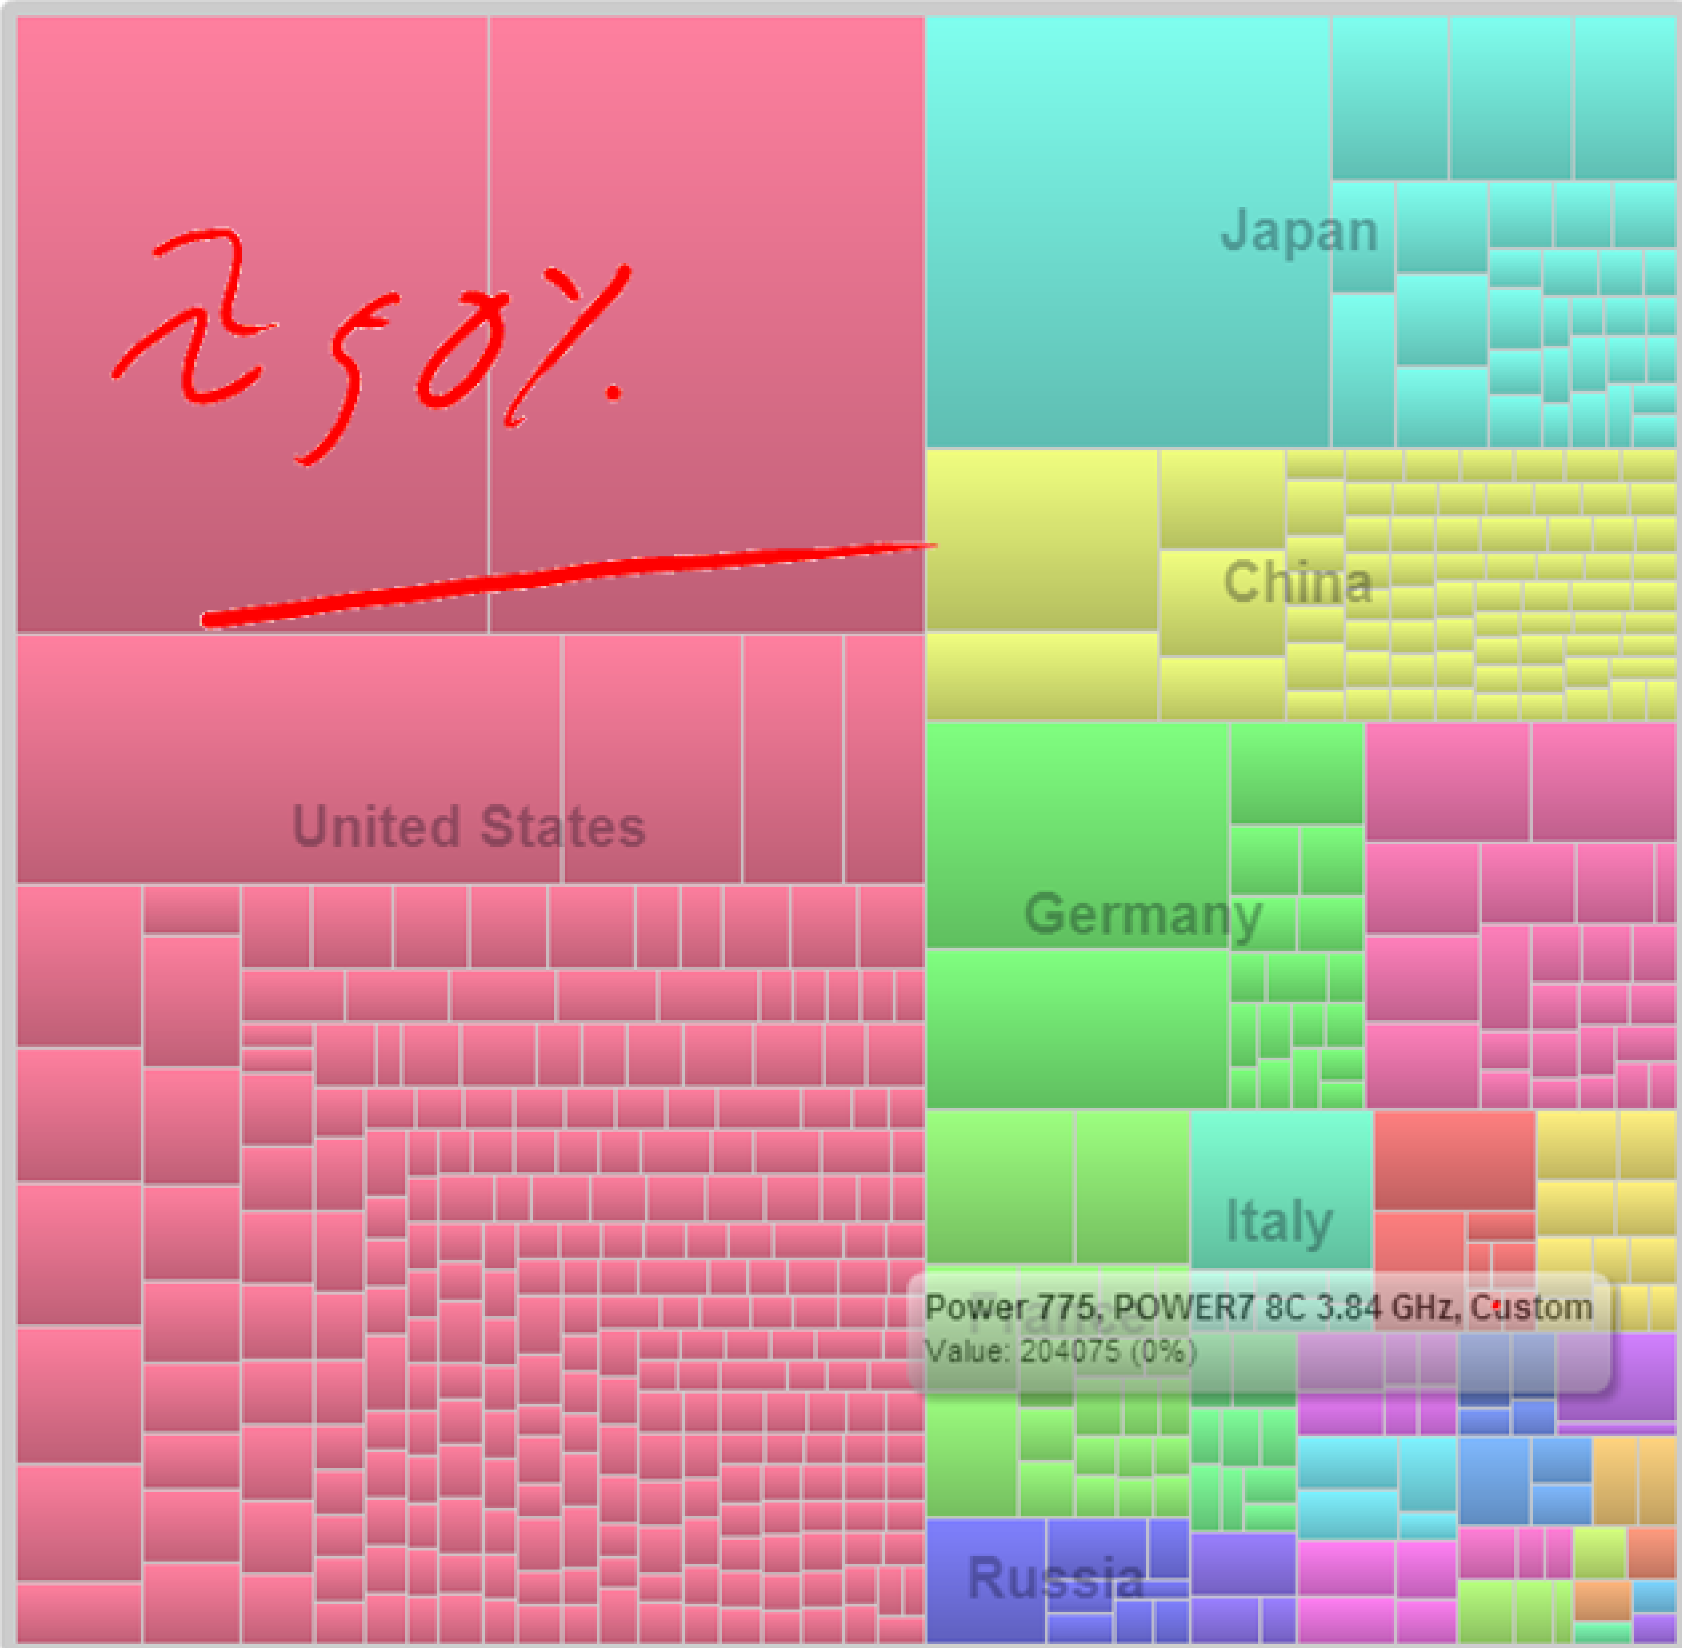

Top 500 by country (June, 2013)

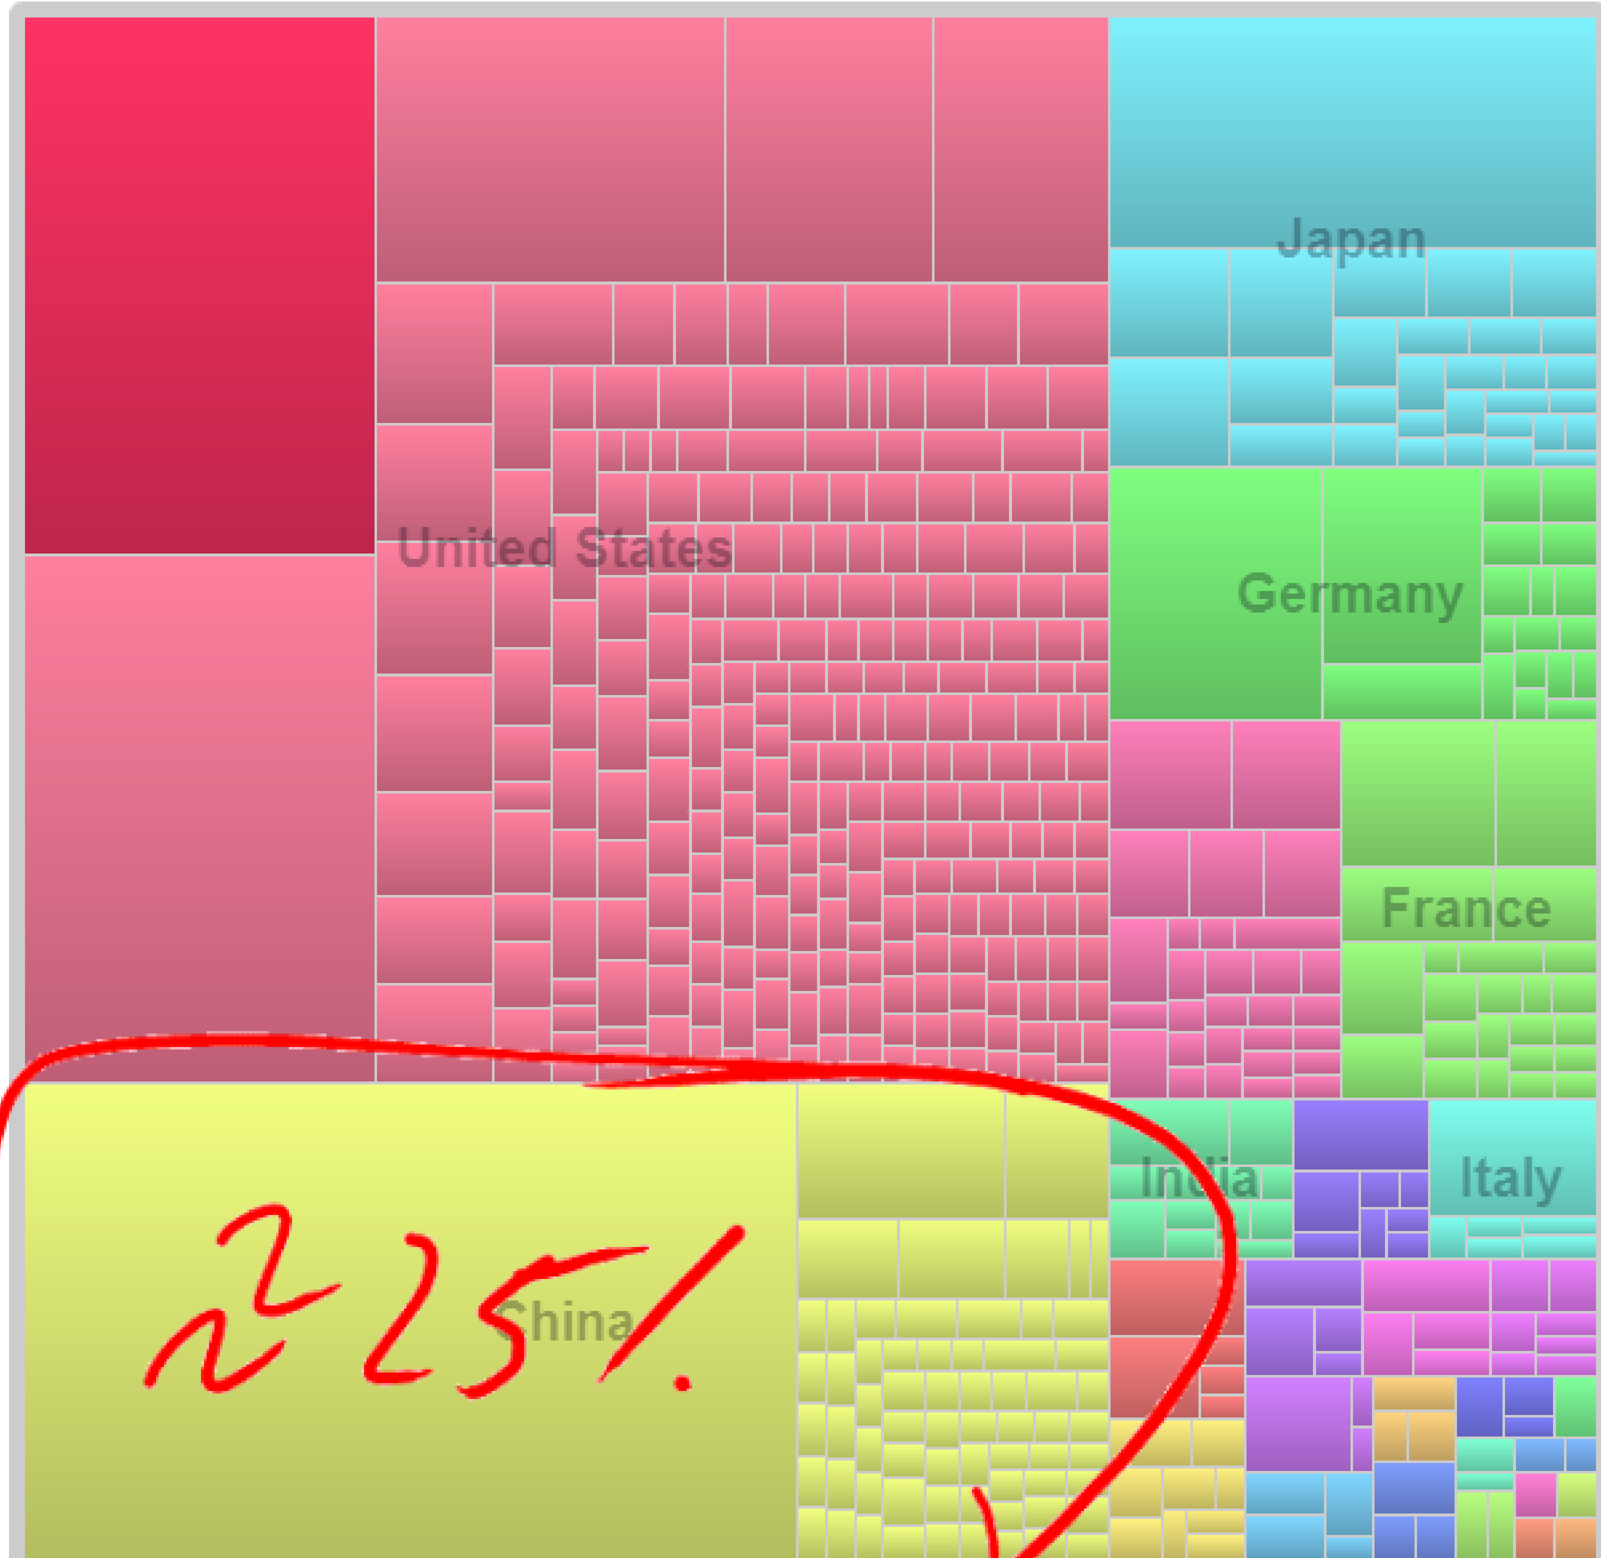

...But Are Size-Limited

Titan at Oak Ridge National Laboratory

- 8.2 MW – enough to power approximately 5000 homes or a city of ~15000 people
- 4,400 square feet – About the size of 25 MSOE dorm rooms

OSC's Newest HPC System: Oakley Cluster

Oakley & Glenn System Configurations

	Oakley System (2012)	Glenn System (Phase II, 2009)
Theoretical Peak Performance	88.6 TFs + GPU Accelerators • 65.5 TFs (DP)	53TFs + GPU Accelerators • 6 TFs (DP)
Number of Nodes	692	650
Number of CPU Sockets	1384	2632
Number of CPU Cores	8304	5300
Number / Kind of GPUs (accelerators)	128 – NVIDIA M2070s	20 – NVIDIA Quadro Plex 2000 S4's
Total Memory	~33.4 TB	24 TBs
Memory per Node	48 GBs	24 GBs
Memory per Core	4 GBs	6 GBs
Interconnect	QDR IB	DDR IB

OSC File Space Information

- Lustre – Parallel File System (Temporary Space)
  - ~570 TBs (all disk)
  - Available only on Glenn and Oakley systems
- Project Space and GPFS
  - ~1000 TBs total usable (Disk)
  - Hierarchical storage capable to tape subsystem
  - Allocated to projects in TBs, for limited time periods
  - Availability
    - Primarily Glenn and Oakley systems via NFS mounts
    - Some OSC virtual machines via NFS, limited SAMBA
- Home Directory Space / NFS
  - ~295 TBs usable (Disk)
  - Allocated to each user, 500 GB quota limit
  - Availability
    - Primarily Glenn and Oakley systems via NFS mounts
    - Some OSC virtual machines via NFS

Mass Storage Overview

- 2 Petabytes (PBs) of usable disk
- ~1 PB NFS storage
- 370 TBs GPFS storage
- 570 TBs Lustre storage
- 1.8 PBs tape

Specs: Oakley Cluster vs. Top 500 Systems in the World

Metric	Performance Ranking	Efficiency Ranking
Overall Ranking in the World	180th	37th
Overall Ranking in North America	89th	8th
Overall Academic Ranking in the World	40th	9th
Overall Academic Ranking in North America	11th	2nd

The Future of Supercomputing

- "From 1975 to 2005, our industry accomplished a phenomenal mission: In 30 years, we put a personal computer on every desk, in every home, and in every pocket."
- "In 2005, our industry undertook a new mission: to put a *personal parallel supercomputer* on every desk, in every home, and in every pocket."
- In 2011, all major form factors went parallel (multicore), including game consoles, tablets and smartphones
  - In 2012, last holdout (Nintendo Wii) went multicore

http://herbsutter.com/welcome-to-the-jungle/