Lecture 1

Introduction to Data Centric Computation

Welcome to CSE 262!

- Your instructor is Scott B. Baden
  baden@ucsd.edu
- Office: room 3244 in EBU3B
- Office hours
  - Mon 2-4, or by appointment
- The class home page is
  http://www.cse.ucsd.edu/classes/sp09/cse262

Course Requirements

- Class participation and preparation: **20%**
  - Be prepared to discuss the assigned readings
  - Write-ups due at start of class (12)
- In class presentation of readings:**10%**
- Project: **70%**
  - Includes proposal, progress report and in class presentation, final report

Policies

- Academic Integrity
  - Do you own work
  - Plagiarism and cheating will not be tolerated
- By taking this course, you implicitly agree to abide by the following the course polices:

  http://www.cse.ucsd.edu/classes/sp09/cse262/Admin.html

Background

- Common Background
  - Programming and Data structures
- Differing backgrounds
  - Prior experience in parallel computation?
  - GPUs/Cell?
  - Numerical analysis?
  - C/C++ Java Fortran? Python? Matlab?

Background Markers

- Navier Stokes Equations
- Sparse factorization
- TLB misses
- Multithreading
- RPC
- Abstract base class
- Amdahl's law
- Relational or Spatial Join

$$\nabla \bullet u = 0$$

$$\frac{D\rho}{Dt} + \rho(\nabla \bullet v) = 0$$

$$\nabla \times u$$

$$f(a) + \frac{f'(a)}{1!}(x-a) + \frac{f''(a)}{2!}(x-a)^2 + \ldots$$

Motivation

- Technical obstacles are limiting our ability to derive knowledge from numerical simulations to understand technologically important phenomena
- Exponential growth in datasets
  - Produced by numerical simulations
  - Data generated by instruments
  - Media
- Exponential increase in hardware capability
  - Multicore processors
  - Clusters
  - Networking

Observation

- The more we know about the how the data was generated …
- … the context in which the data will be used
- … the more we can employ that knowledge to improve productivity

Syllabus

- Techniques, tools and algorithms to improve our ability to derive knowledge from data or media
- Spatial query techniques and access methods
- Feature extraction and tracking
- Data compression
- Wavelets, image processing
- Hardware acceleration
  (i.e. Cell BE and GPUs)
- Applications
  - Simulation data sets
  - Media/observational

Scheduling

- Project Symposium
  - June 4th-5th
  - No class on 6/2
- Some lectures to be rescheduled
  - Monday 5/19 → Friday 5/8
    (In class presentations of progress reports)
  - Tuesday 5/21: Friday 5/15
  - Tuesday 5/26 → 6/5

Hardware platforms

- Abe
  - Dell cluster at NCSA
  - Lincoln: attached NVIDIA Tesla
- Cell Broadband Engine (Georgia Tech)
- If you want to use Cell BE, fill out an account request form

https://cell-web.cc.gatech.edu/signup/index.pl

Motivating Applications

Computational fluid dynamics

- Direct Numerical Simulation of horizontal shear
- Sutanu Sarkar, UCSD MAE Department
- 80 Terabytes of data per simulation
  - 1024 Cray XT-4 processors for 48 hours
  - 2×10^9 unknowns, 10K time steps [4096×1024×512]
  - 1000 snapshots @ 80×10^9 bytes

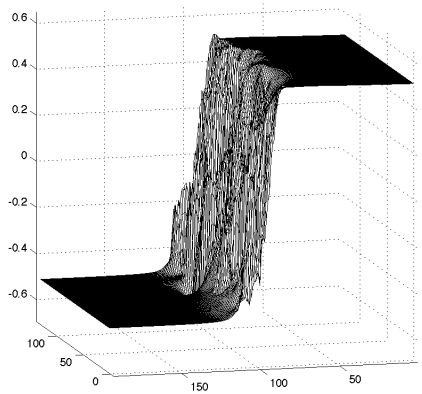

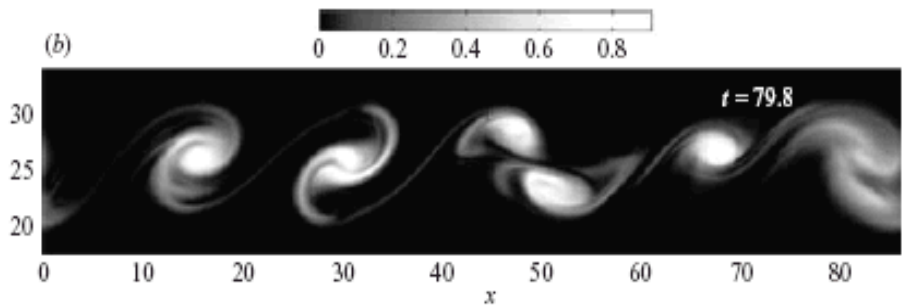

Horizontal spanwise vorticity (mag) using DNS, [Basak & Sarkar, JFM 2006]

The Input Data

- Time dependent simulation
- Each snapshot
  - 5 components: velocity (vector), density, pressure

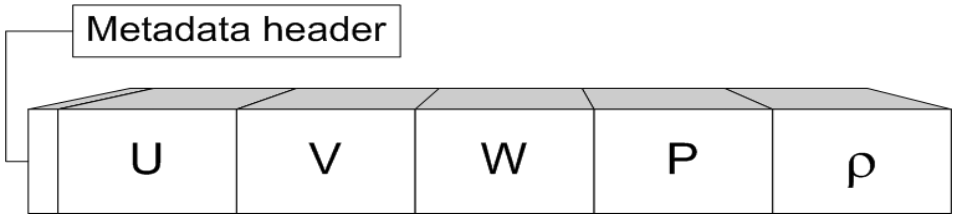

The end-to-end application

- Identify vortex cores using the "Δ criterion"

  $\Delta > 1.75\times10^{-4}$

- Compute Δ directly on a compressed version of the dataset
- Compute quantities conditioned on Δ

  F(x; condition)

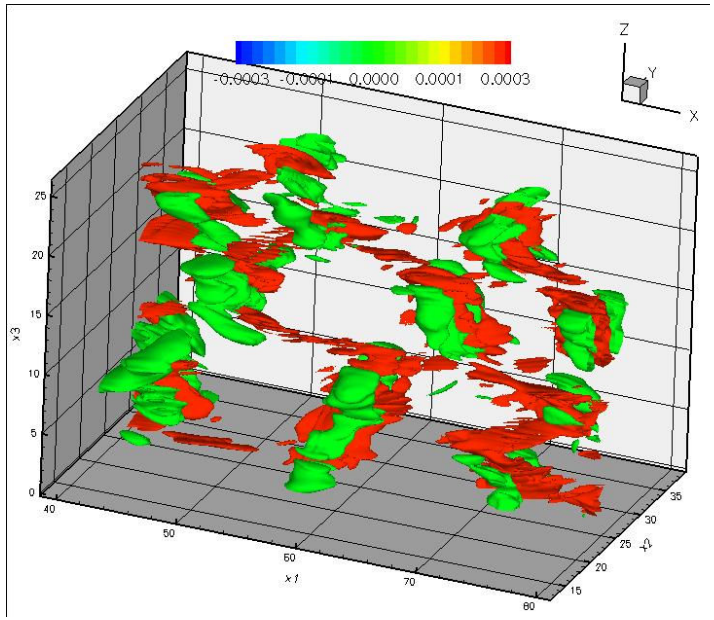

Working with sparse representations

- Controlled operations on irregular multiblock structures
- Geometric abstractions, e.g. KeLP, Chombo, Titanium
- TeraLab, Faisal Mir, KTH Sweden

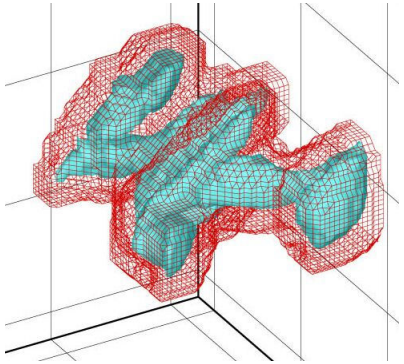

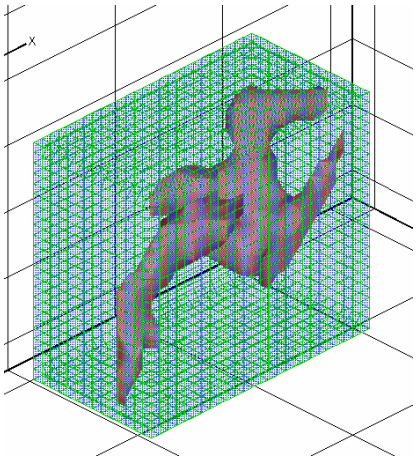

With William Kerney, Peter Diamessis, and Keiko Nomura (UCSD) [2002]

Dynamics

- Continuation: the overlap between a feature and a descendent or an ancestor
- Creations and dissipations are not continuations

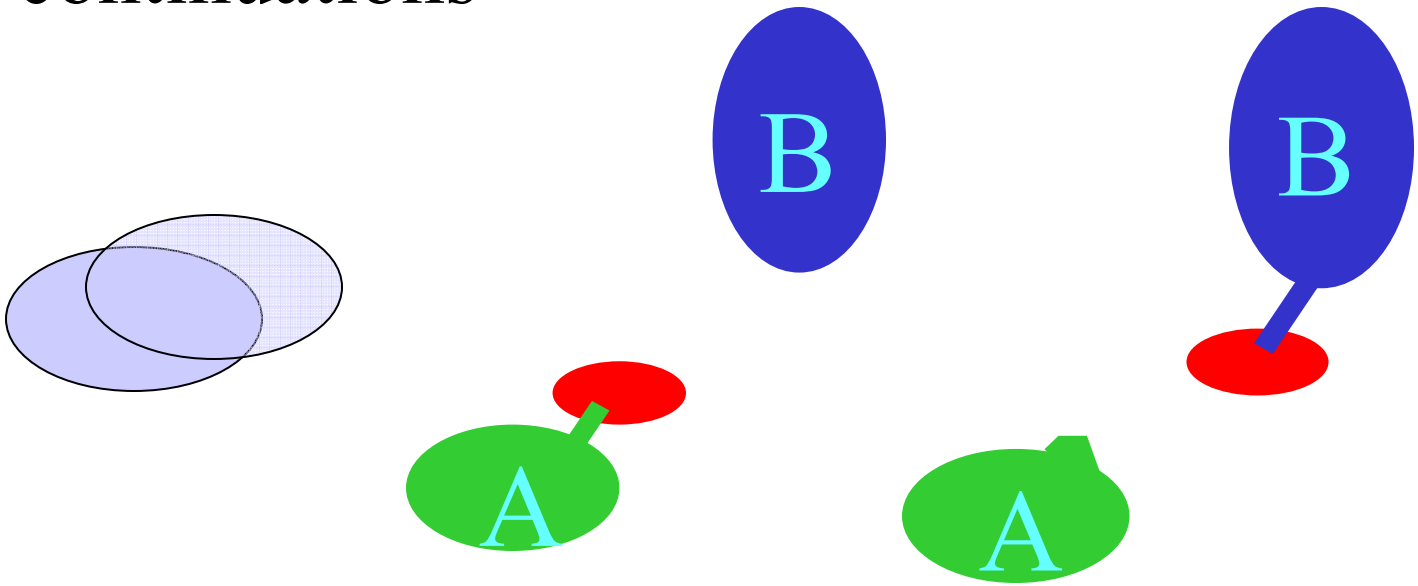

How large can a dataset get?

- "TeraShake" Ground Fault Simulator
- Seismic waves generated by large earthquakes that occur on the southern San Andreas Fault
- 43 Terabytes of data
  - 240 IBM Power4 processors for 5 days
  - 2000 snapshots @ 21.1×10^9 bytes
  - Stored on tape in SDSC's Storage Research Broker (SRB) http://www.sdsc.edu/SCEC/
- Southern California Earthquake Center's Community Modeling Environment Project (SCEC/CME), NSF http://www.scec.org/cme
  - USC, Scripps Inst. for Oceanography, San Diego Supercomputer Center(SDSC), SDSU ...

TeraShake

Simulates a 7.7 earthquake along the southern San Andreas fault close to LA using seismic, geophysical, and other data from the Southern California Earthquake Center

How it works:

1. Divide up Southern California into "blocks"
2. For each block, get all the data on ground surface composition, geological structures, fault information, etc.

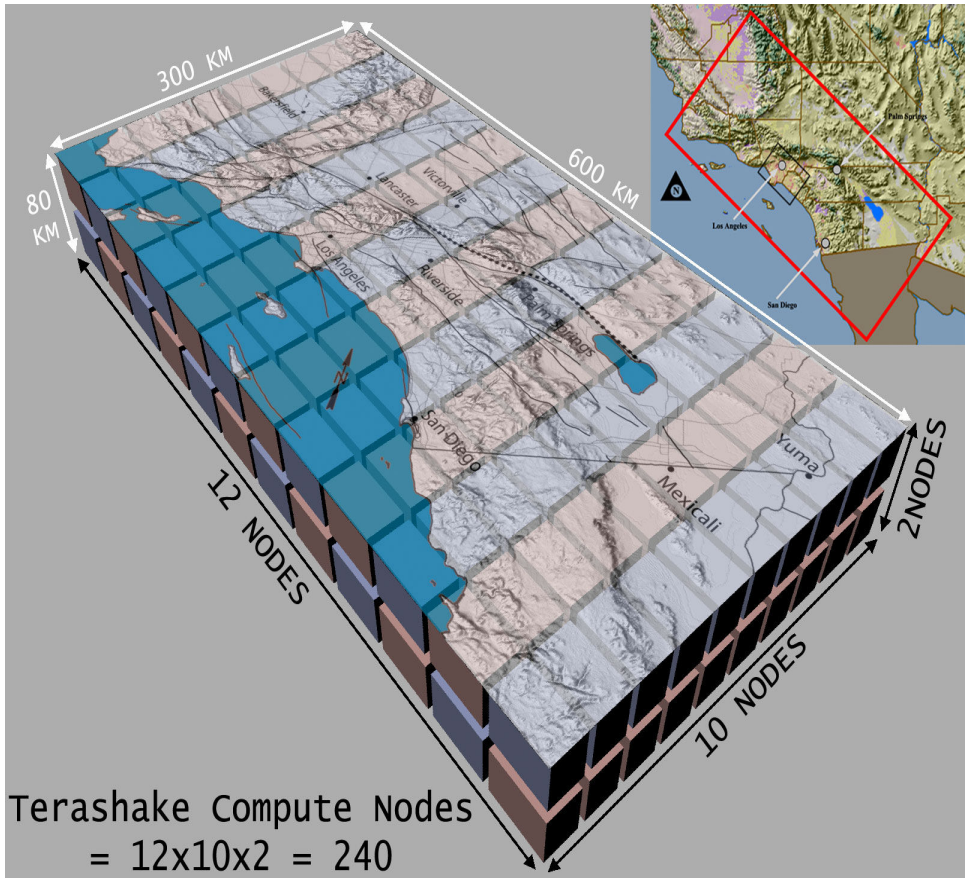

Slide Courtesy of DataCentral@SDSC

Animation

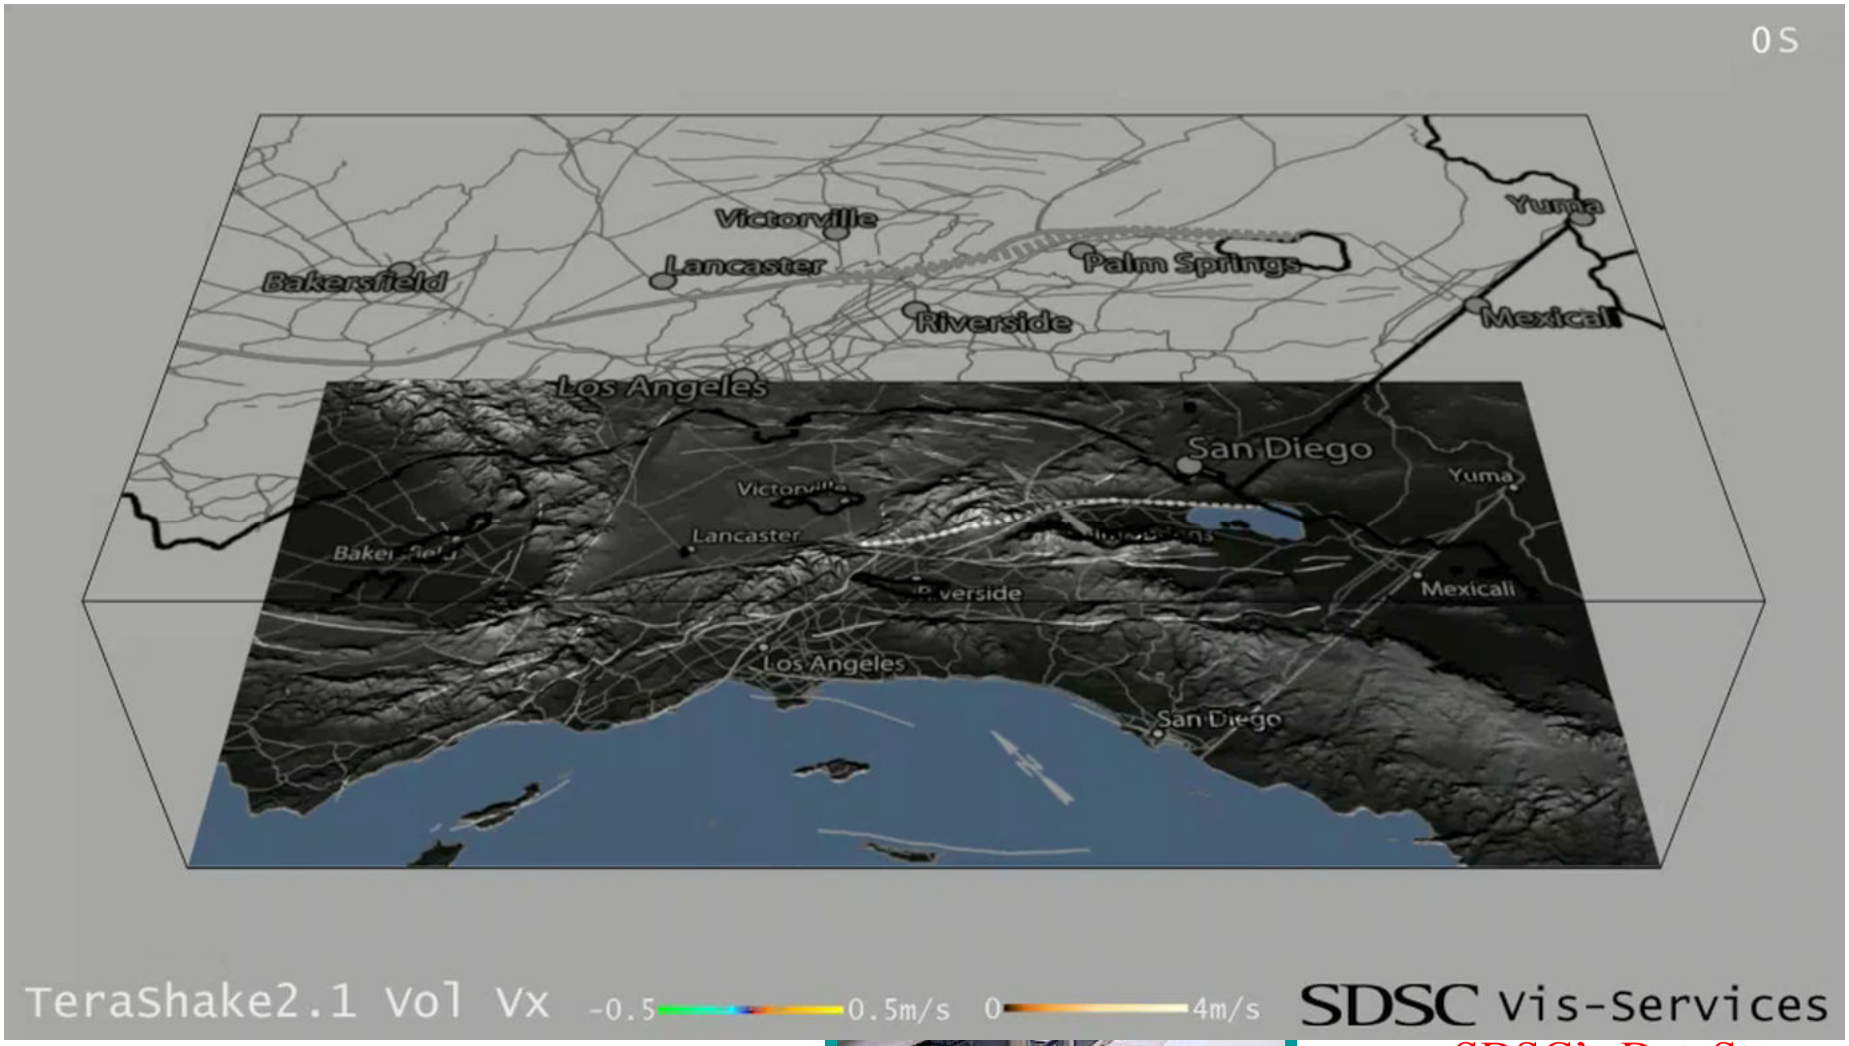

Slide Courtesy of DataCentral@SDSC

SDSC's DataStar

TeraShake Requirements

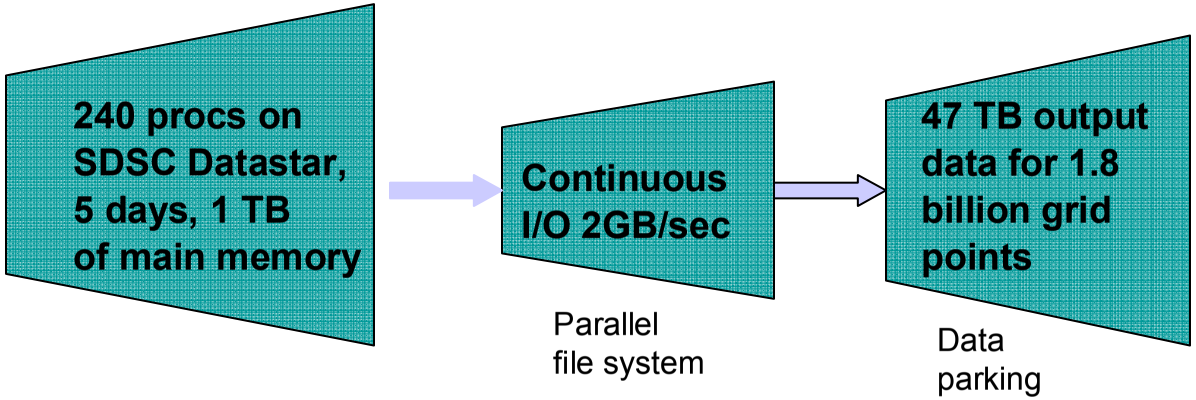

Data parking of 100s of TBs for many months

Large memory Nodes with 256 GB of DS for pre-processing and post visualization

10-20 TB data archived a day

The next generation simulation will require even more resources: Researchers plan to double the temporal/spatial resolution of TeraShake

Slide Courtesy of DataCentral@SDSC

The Role of Compression

- If the information content of a dataset is low, compression can help us …
  - increase resolution for a given storage budget
  - discover knowledge hidden within the data

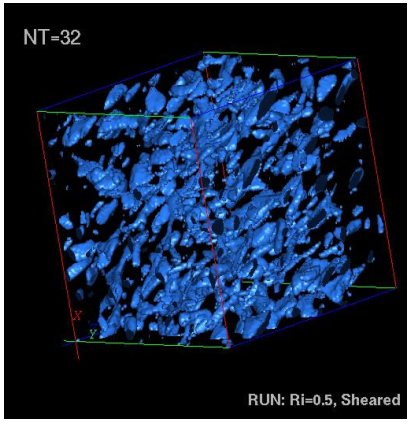

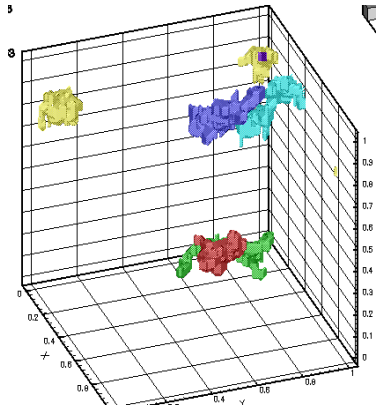

Diamessis, Kerney, Nomura, Baden [Para '02]

Data Compression

- Lossless
  - Perfect reconstruction, not more than 2:1 compression
  - RLE, Huffman coding, LZ77, DEFLATE (gzip)
- In practice: lossy
  - Wavelet Compression (WC)
  - Feature Extraction
  - Introduces artifacts into the data
  - Adaptive Compression (with Hromadka, Unat, Shafaat)

Technology

- Multicore processors may enable us to continue the trend of exponential growth in processing speeds
- Workstations or clusters with accelerators
  - Roadrunner: with STI Cell BE (IBM and Los Alamos)
  - Lincoln: with Nvidia Tesla (NCSA and Nvidia)
- Distributed computation
  - "Cloud" and "Grid" computing

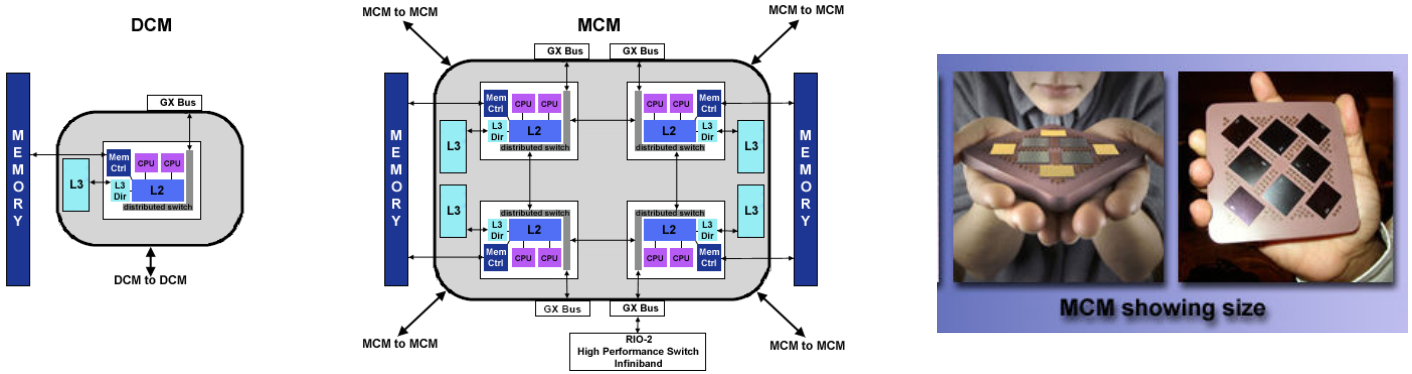

http://computing.llnl.gov/tutorials/ibm_sp (IBM)

Today's laptop would have been yesterday's supercomputer

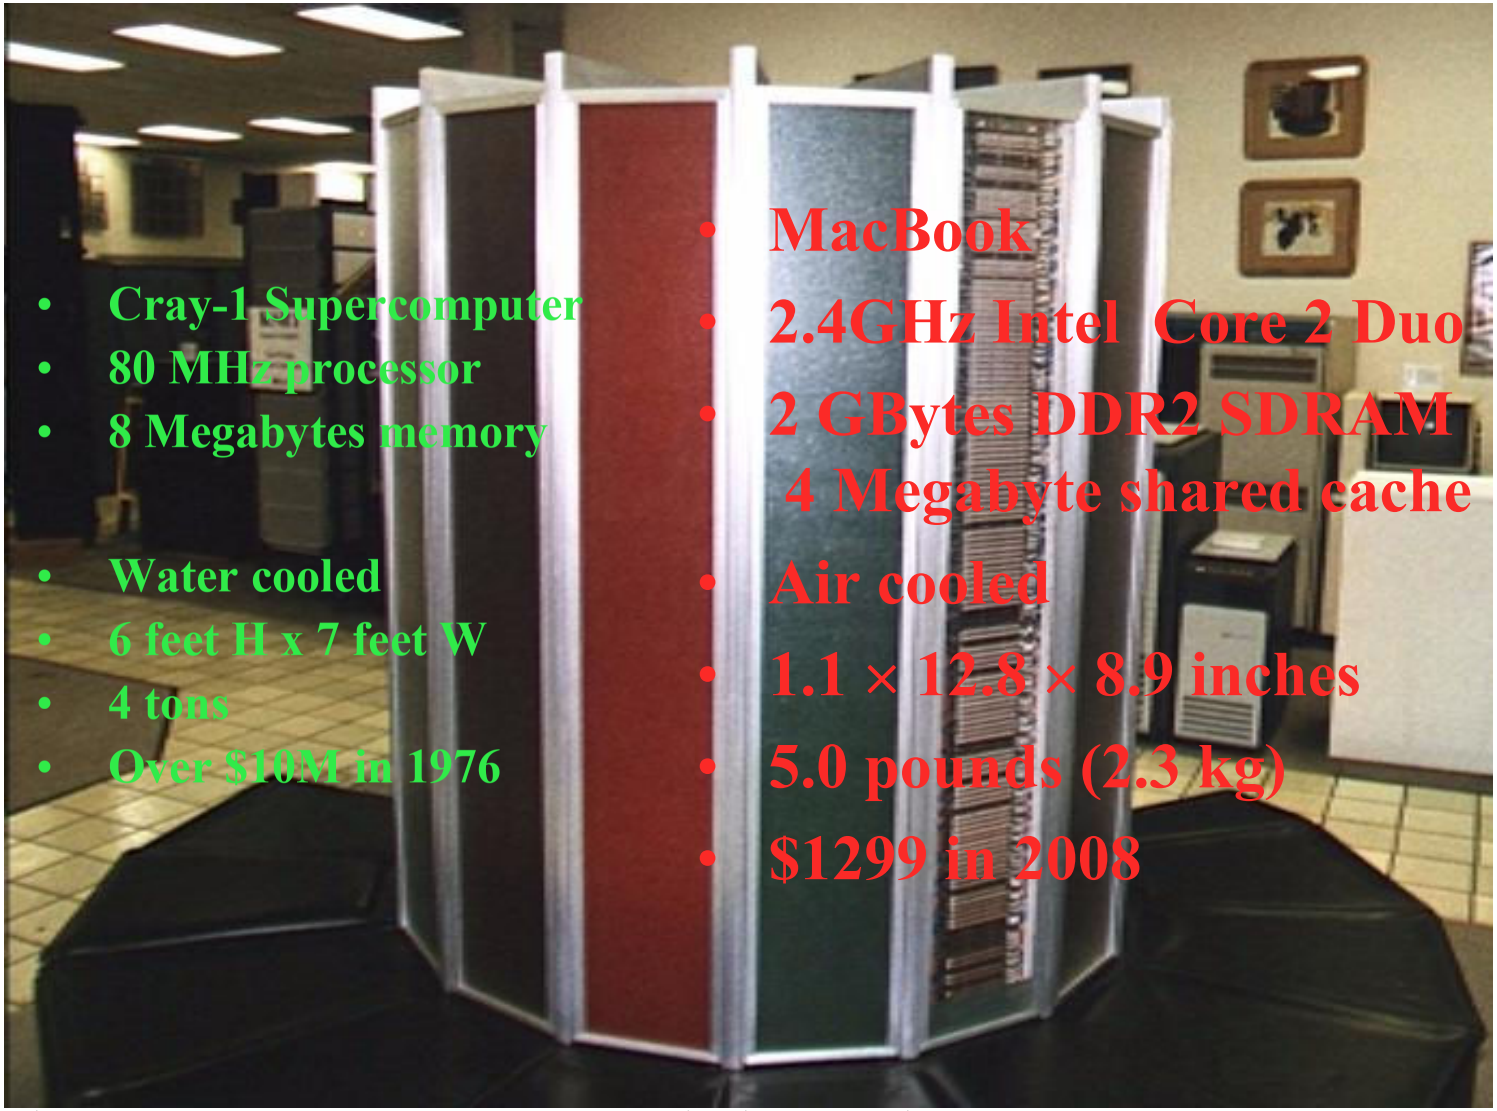

- Cray-1 Supercomputer
- 80 MHz processor
- 8 Megabytes memory

- Water cooled
- 6 feet H x 7 feet W
- 4 tons
- Over $10M in 1976

- MacBook
- 2.4GHz Intel Core 2 Duo
- 2 GBytes DDR2 SDRAM
  4 Megabyte shared cache
- Air cooled
- 1.1 × 12.8 × 8.9 inches
- 5.0 pounds (2.3 kg)
- $1299 in 2008

The processor-memory gap

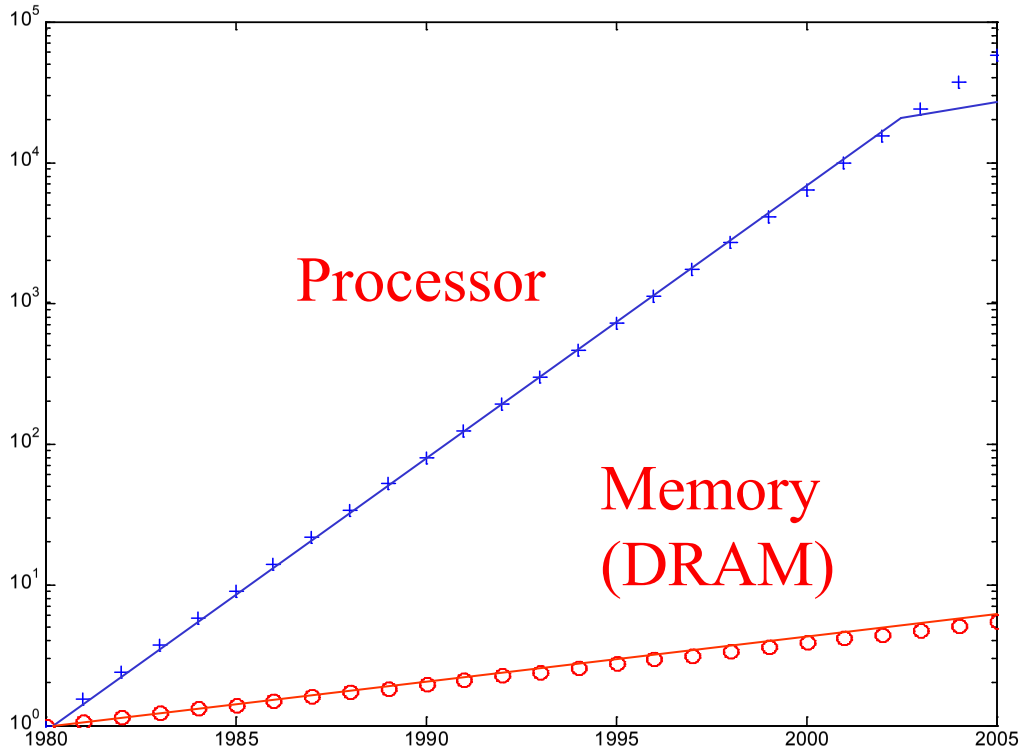

An important universal: the locality principle

- Programs generally exhibit two forms of locality when accessing memory, often involving loops
  - Temporal locality (time)
  - Spatial locality (space)
- Often involves loops

```
for t = 0 to T-1
   for i = 1 to N-2
       u[i]= (u[i-1]  + u[i+1] - 8*h*h ) / 2
```

Re-use

- Memory hierarchies rely on locality to improve memory access times through re-use
- Underlying principles for improving cache use apply to stored data, parallel memory hierarchies and hardware accelerators

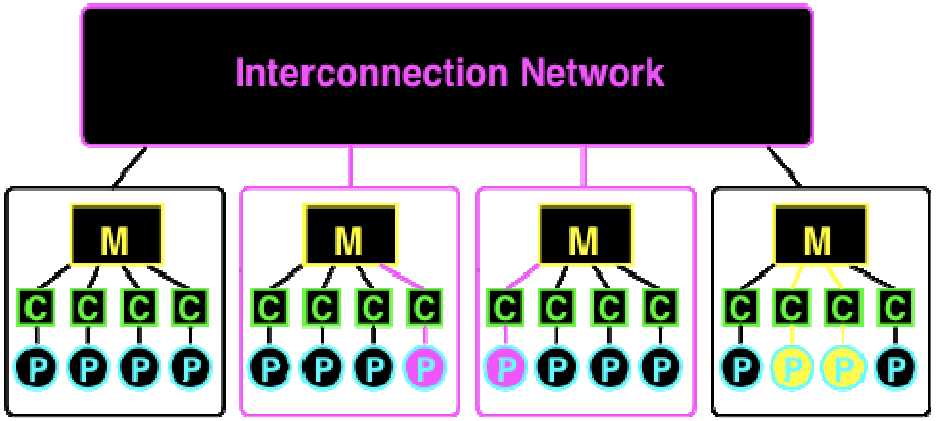

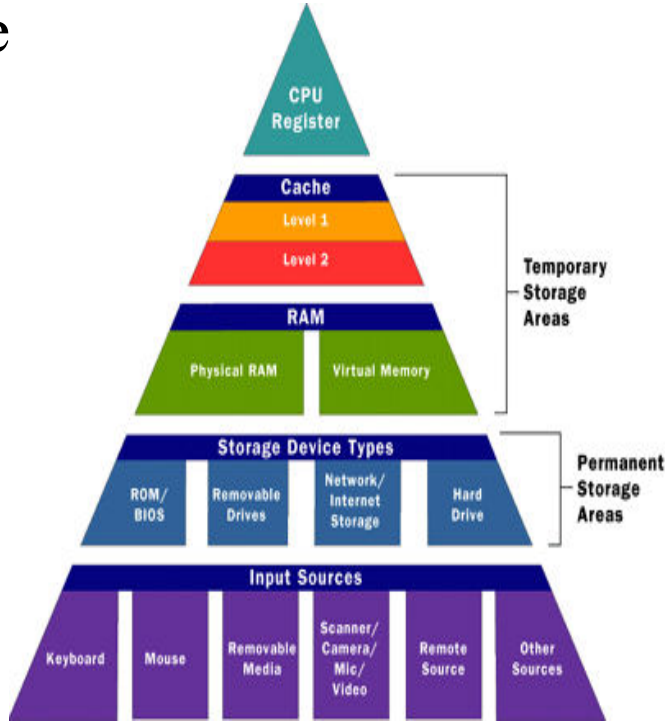

STI Cell Broadband Engine

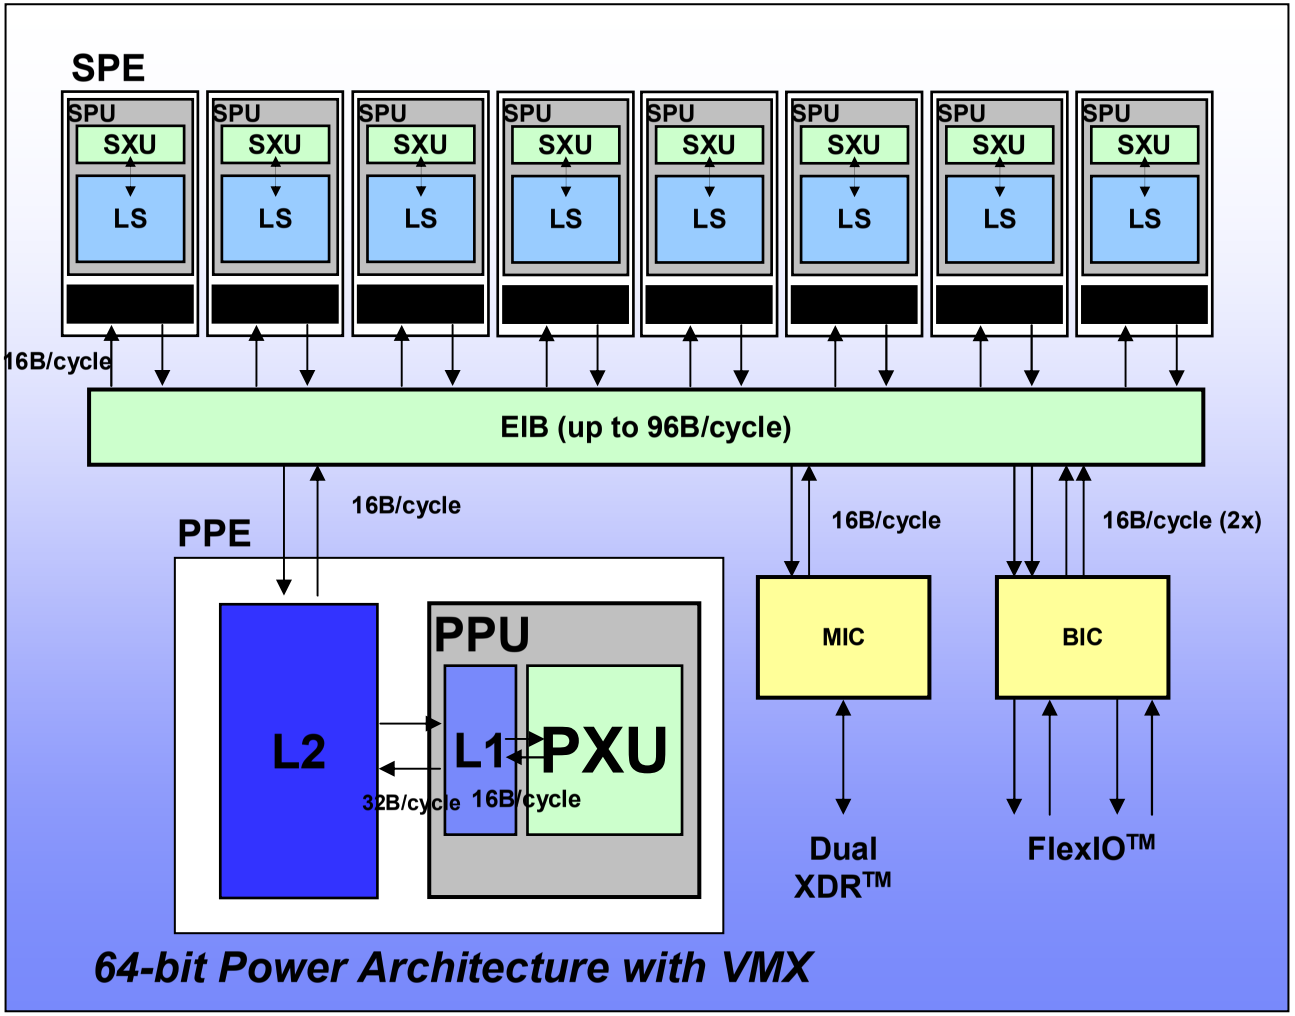

IBM

NVIDIA GeForce GTX 280

- 240 cores @ 1.296 GHz
- Parallel computing or graphic mode
- 1 GB memory
- 512 bit memory interface @ 141.7 GB/s