PVM MPI 12

Sistemas Operativos y Distribuidos

Prof. Javier Echaiz

D.C.I.C. – U.N.S.

http://cs.uns.edu.ar/~jechaiz

je@cs.uns.edu.ar

Outline

- Motivation
- Modern scientific method
- Evolution of supercomputing
- Modern parallel computers
- Seeking concurrency
- Programming parallel computers
- **PVM**

What is Parallel and Distributed computing?

- Solving a single problem faster using multiple CPUs
- Parallel = Shared Memory among all CPUs
- Distributed = Local Memory/CPU
- Common Issues: Partition, Synchronization, Dependencies

Why Parallel and Distributed Computing?

- Grand Challenge Problems
  - Weather Forecasting; Global Warming
  - Materials Design – Superconducting material at room temperature; nano-devices; spaceships.
  - Organ Modeling; Drug Discovery
- Physical Limitations of Circuits
  - heat and light effect
  - Superconducting material to counter heat effect
  - Speed of light effect – no solution!

Why Parallel and Distributed Computing?

- VLSI – Effect of Integration
  - 1 M transistor enough for full functionality
  - Rest must go into multiple CPUs/chip
- Cost – Multitudes of average CPUs give better FLPOS/$ compared to traditional supercomputers
- Idling workstations should be utilized
- Everyday Reasons
  - Solve compute-intensive problems faster
    - Make infeasible problems feasible
    - Reduce design time
  - Solve larger problems in same amount of time
    - Improve answer's precision
    - Reduce design time
  - Gain competitive advantage

Definitions

- Parallel computer
  - Multiple-processor system supporting parallel programming
- Parallel programming
  - Programming in a language that supports concurrency explicitly

Why MPI and PVM?

- MPI = “Message Passing Interface”
- PVM = “Parallel Virtual Machine”
- Standard specification for message-passing libraries
- Libraries available on virtually all parallel computers
- Free libraries also available for networks of workstations or commodity clusters

Why Shared Memory programming?

- Easier conceptual environment
- Programmers typically familiar with concurrent **threads** and **processes** sharing address space
- OpenMP an application programming interface (API) for shared-memory systems
  - Supports higher performance parallel programming of symmetrical multiprocessors

Classical Science

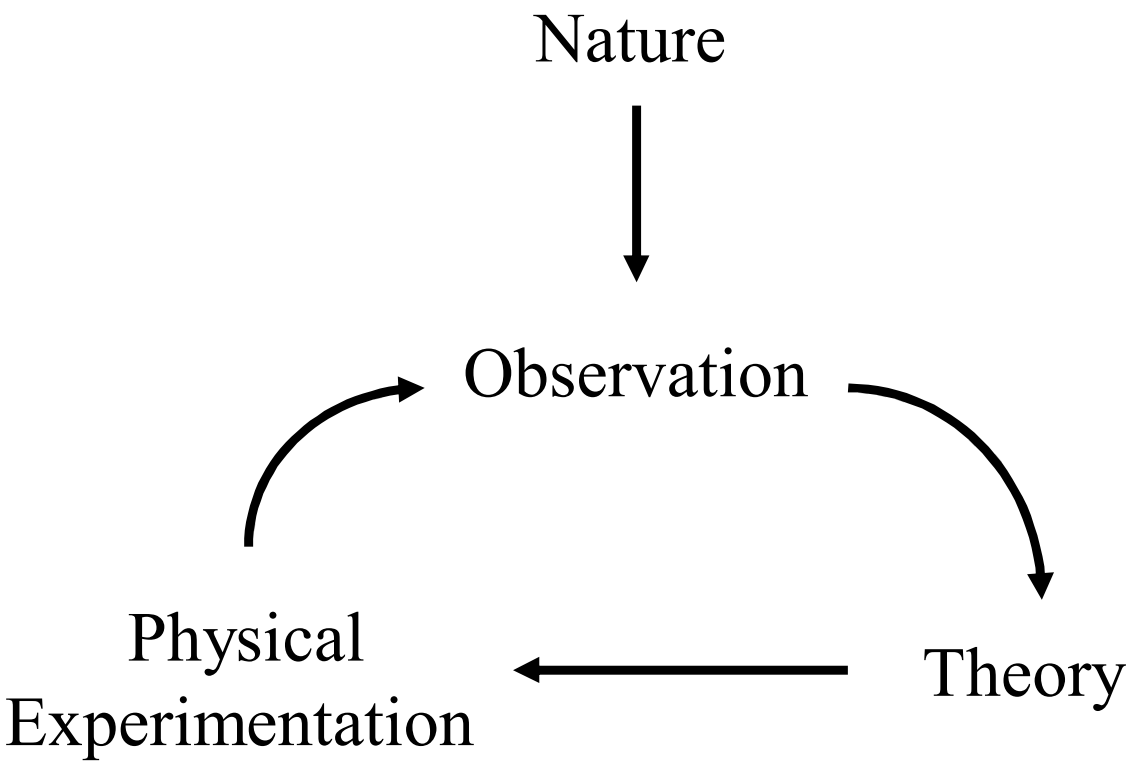

Modern Scientific Method

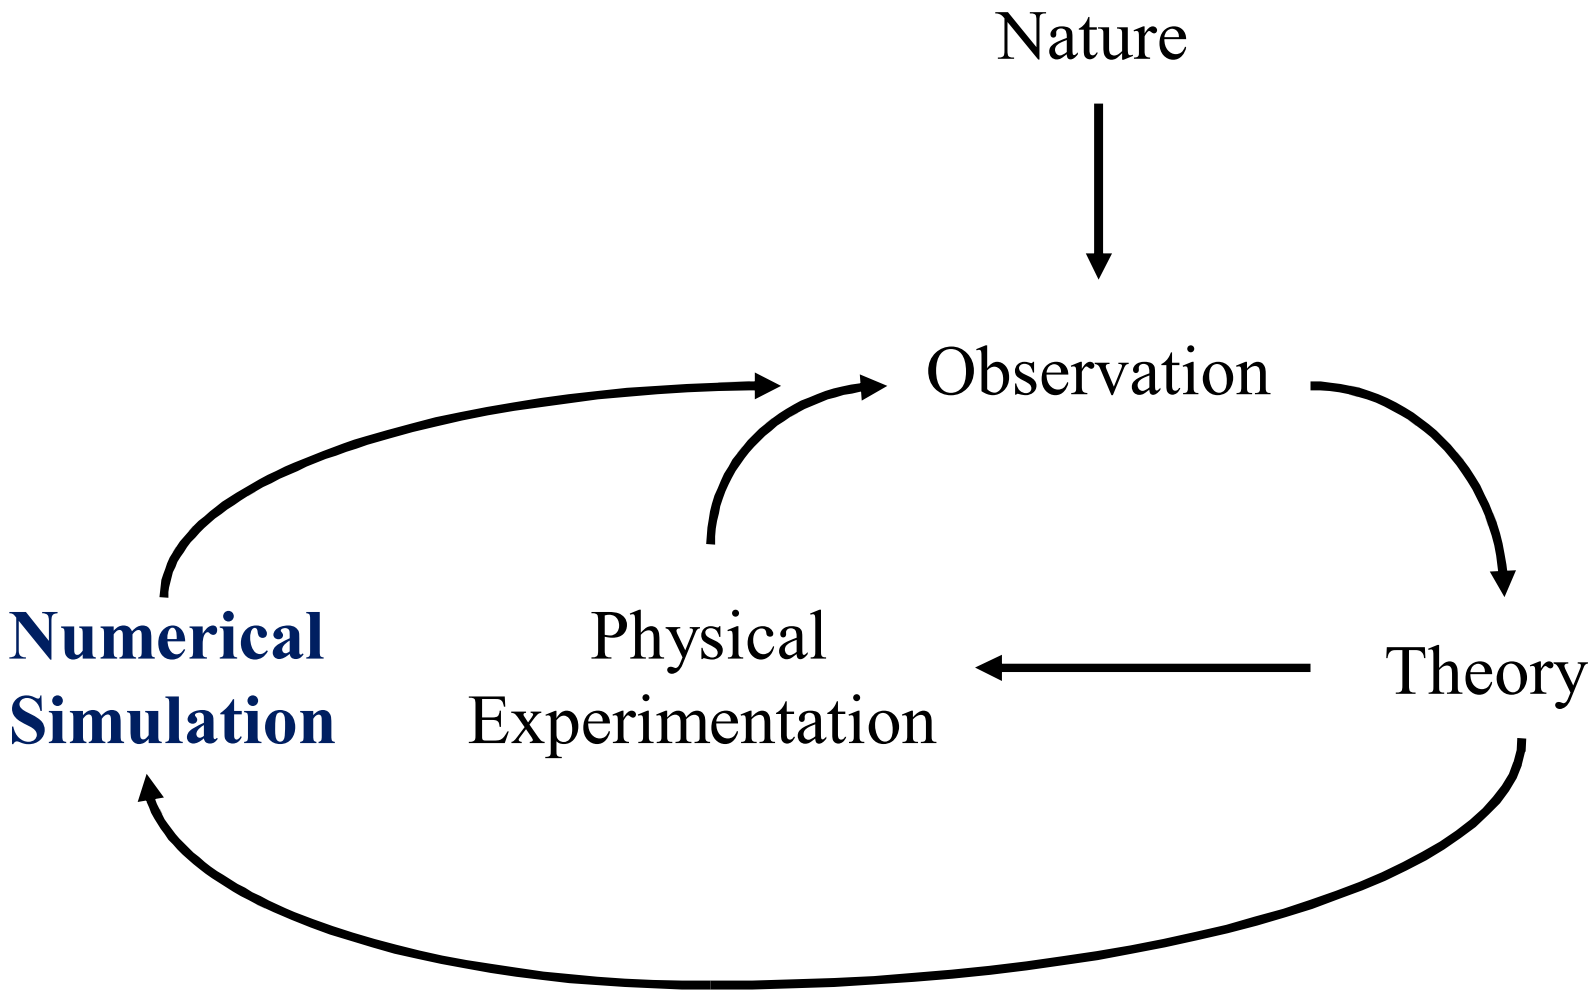

Evolution of Supercomputing

- World War II
  - Hand-computed artillery tables
  - Need to speed computations
  - ENIAC
- Cold War
  - Nuclear weapon design
  - Intelligence gathering
  - Code-breaking

Advanced Strategic Computing Initiative

- U.S. nuclear policy changes
  - Moratorium on testing
  - Production of new weapons halted
- Numerical simulations needed to maintain existing stockpile
- Five supercomputers costing up to $100 million each

ASCI White (10 teraops/sec)

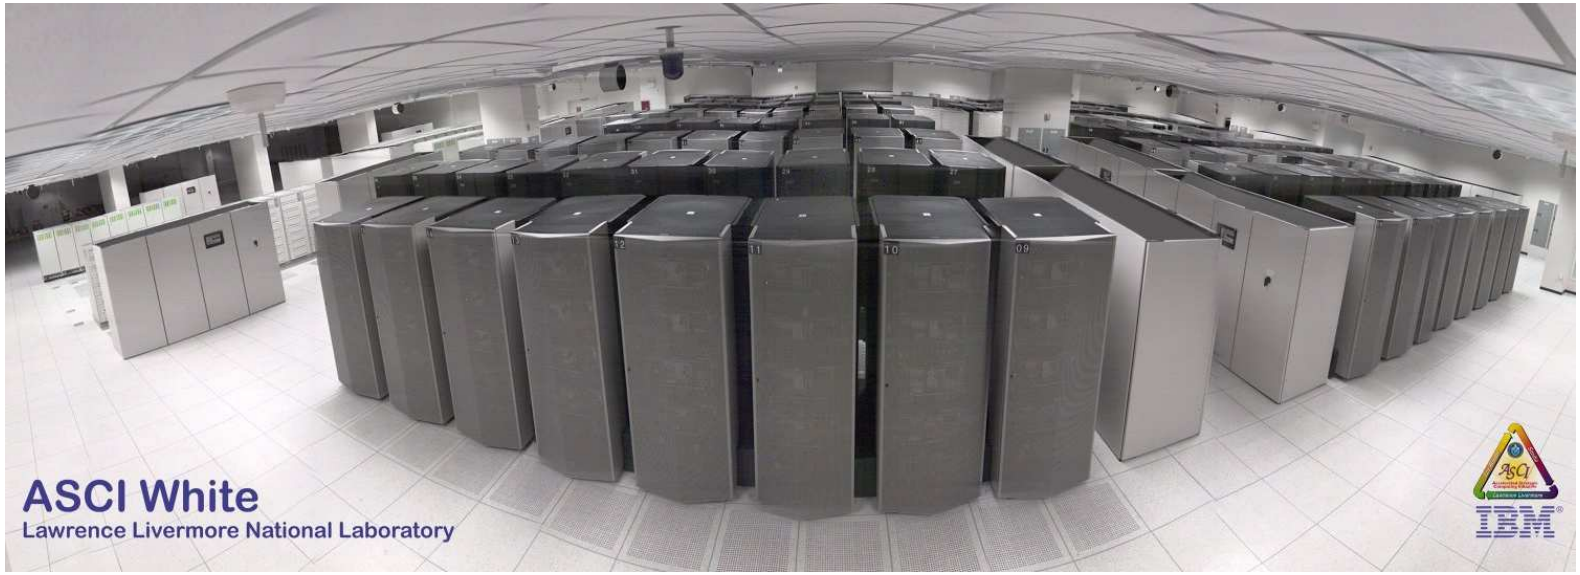

Supercomputer

- Fastest General-purpose computer
- Solves individual problems at high speeds, compared with contemporary systems
- Typically costs $10 million or more
- Traditionally found in government labs

Commercial Supercomputing

- Started in capital-intensive industries
  - Petroleum exploration
  - Automobile manufacturing
- Other companies followed suit
  - Pharmaceutical design
  - Consumer products
  - Financial Transactions

50 Years of Speed Increases

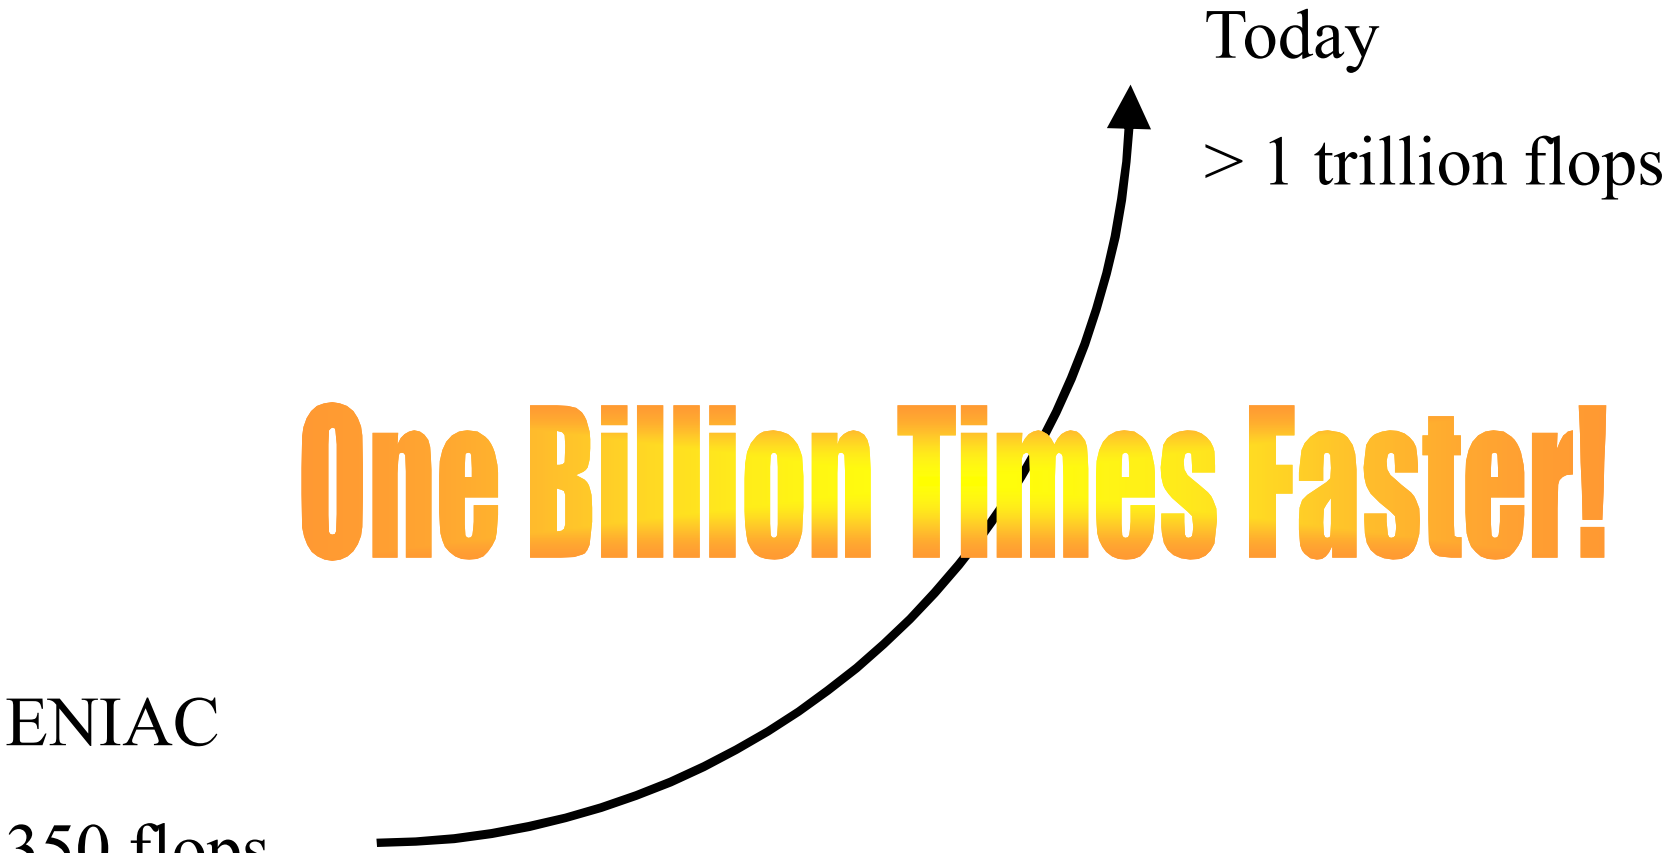

CPUs 1 Million Times Faster

- Faster clock speeds
- Greater system concurrency
  - Multiple functional units
  - Concurrent instruction execution
  - Speculative instruction execution

Systems 1 Billion Times Faster

- Processors are 1 million times faster
- Combine thousands of processors
- Parallel computer
  - Multiple processors
  - Supports parallel programming
- Parallel computing = Using a parallel computer to execute a program faster

Microprocessor Revolution

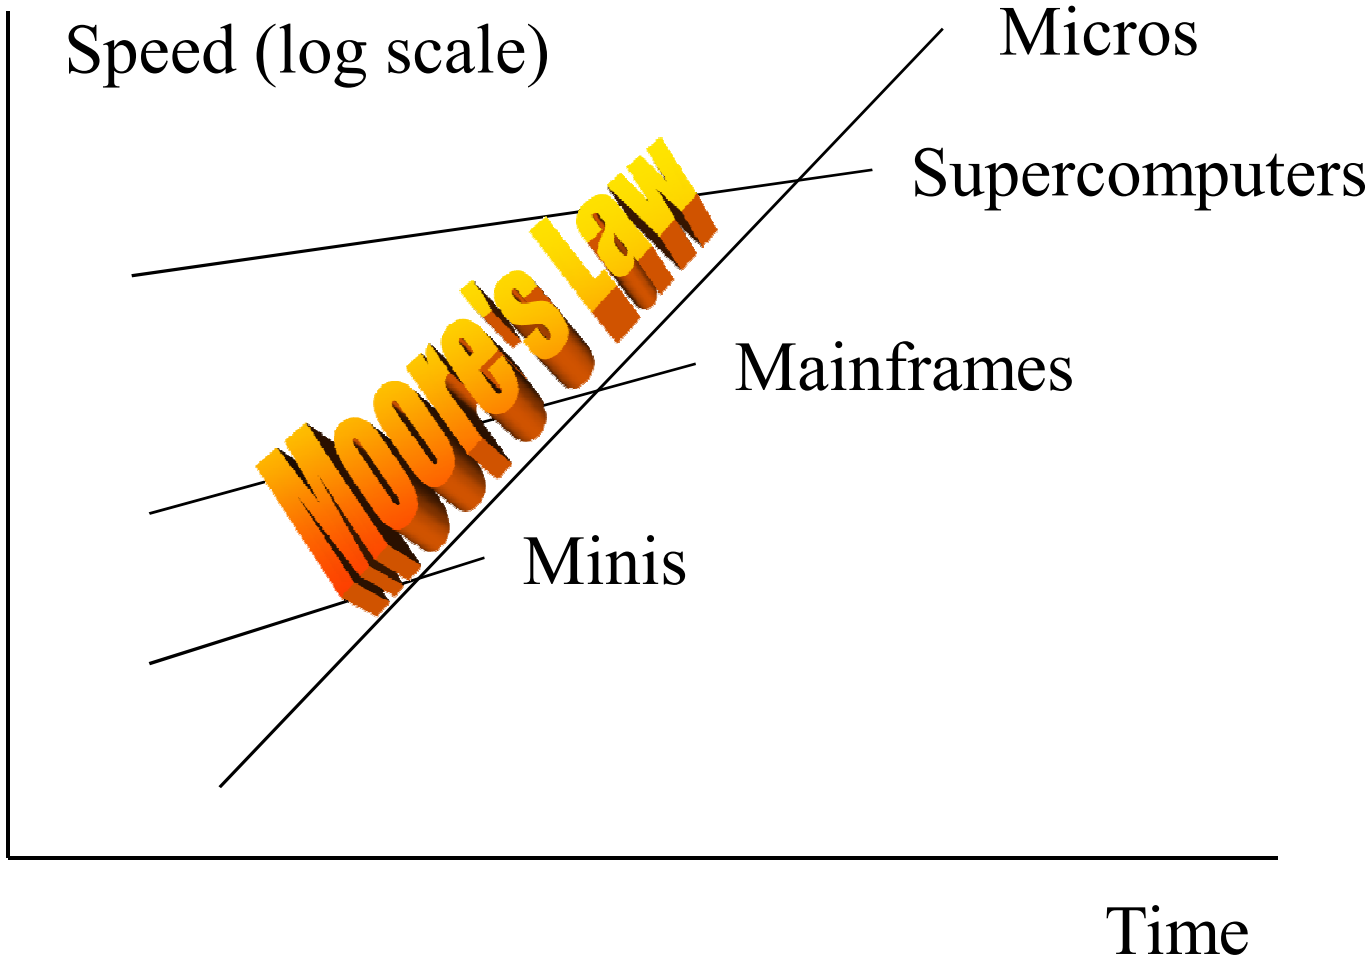

Modern Parallel Computers

- Caltech's Cosmic Cube (Seitz and Fox)
- Commercial copy-cats
  - nCUBE Corporation (512 CPUs)
  - Intel's Supercomputer Systems
    - iPSC1, iPSC2, Intel Paragon (512 CPUs)
  - Lots more
- Thinking Machines Corporation
  - CM2 (65K 4-bit CPUs) – 12-dimensional hypercube - SIMD
  - CM5 – fat-tree interconnect - MIMD

Copy-cat Strategy

- Microprocessor
  - 1% speed of supercomputer
  - 0.1% cost of supercomputer
- Parallel computer = 1000 microprocessors
  - 10 x speed of traditional supercomputer
  - Same cost as supercomputer

Why Didn't Everybody Buy One?

- Supercomputer $\neq \Sigma$ CPUs
  - Computation rate $\neq$ throughput (#jobs/time)
  - Slow Interconnect
  - Inadequate I/O
  - Customized Compute and Communication processors meant inertia in adopting the fastest commodity chip with least cost and effort
  - Focus on high end computation meant no sales volume to reduce cost
- Software
  - Inadequate operating systems
  - Inadequate programming environments

Commercial Parallel Systems

- Relatively costly per processor
- Primitive programming environments
- Focus on commercial sales
- Scientists looked for alternative

Beowulf Concept

PCPD

- NASA (Sterling and Becker)
- Commodity processors
- Commodity interconnect
- Linux operating system
- MPI/PVM library
- High performance/$ for *certain* applications

Seeking Concurrency

- Data dependence graphs
- Data parallelism
- Functional parallelism
- Pipelining