SAND2005-1550C

System Architecture for Nuclear Power Simulations

Jim Tomkins

Presentation to SINEMA Workshop
October 21, 2004

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

Outline

- The Scale of Simulation
  - 20 years ago
  - Today
  - Tomorrow
- Parallel Computing Infrastructure
  - Programming Model
  - Numerical Methods
- Computer Architectures for Simulation
  - Today
  - Tomorrow
- Wide Area Network
- Conclusions

The Scale of Simulation – 20 Years Ago

MELPROG Code - Light Water Reactor Safety Analysis

- 2-Dimensional (R-Z)
- Multi-phase flow (4 fields); conduction, convection, and radiation heat transfer; oxidation; structural failure; debris formation
- ~100 fluid cells in Reactor vessel
- Typically 3 fuel rods plus 1 control rod modeled - a few thousand mesh points in total
- Reactor Vessel and Internal Structures
- Coupled to TRAC code to provide coolant loops, etc.

Used for extensive analysis of Three Mile Island accident

- 1 processor of XMP
- ~4 MW (32 MB)
- ~100 hours of CPU time

The Scale of Simulation – Today

Red Storm Computer System

- 20 TB of memory - over 600,000 times as much memory
- 10,368 processors - ~100,000 times faster(each processor is ~10 times faster than an XMP processor, peak is 17 times higher)

Simulation Comparison

- Full 3-d simulations are possible
- Explicit modeling of every flow channel, fuel rod, control rod, and structure in the reactor vessel
- Very high level of physics
- 250,000,000 fluid cells in the reactor vessel
- Fluid Cell size around a centimeter
- CPU time may grow as fluid cells shrink in size

Sandia National Laboratories

The Scale of Simulation – Tomorrow

PetaFlop Computer System - ~2010

- 500 TB of memory - 15,000,000 times a much memory (25)
- 50,000 processors - ~2,000,000 times faster (20) (each processor is ~40 times faster than an XMP processor, peak is 85 times higher)

Simulation Comparison

- Even more detail than with today's largest machines
- 5,000,000,000 cells in the reactor vessel
- Cell size will be a fraction of a centimeter
- CPU time may grow even longer as fluid cell size shrinks even smaller

Sandia National Laboratories

Parallel Computing Infrastructure

- Twenty years ago there wasn't any parallel computing infrastructure and most computing was done on a single processor
- Today parallel machines are common for scientific and engineering work
- Programming Model
  - Explicit Message Passing - MPI has become a standard
  - Explicit Message Passing will continue to be the programming model for simulation for the foreseeable future - MPI will evolve
- Numerical Methods
  - There has been a lot of work on parallel solvers - particularly sparse solvers
  - Sparse solvers are now scalable to thousands of processors
  - Solver libraries are readily available as open source software packages

Computer Architectures - Today

Workstations

- Appropriate for many problems
- Code development platform

Linux Clusters

- Distributed memory MIMD
- A few processors to a few thousand processors
- Appropriate for small to medium size simulations
- Scalable performance to large numbers of processors is an issue

Large Massively Parallel Processor Systems (MPPs)

- Distributed memory MIMD
- Many thousands of tightly coupled processors
- Designed to provide high parallel efficiency
- Appropriate for simulations involving thousands of processors

Scaling Data for two Sandia Engineering Codes

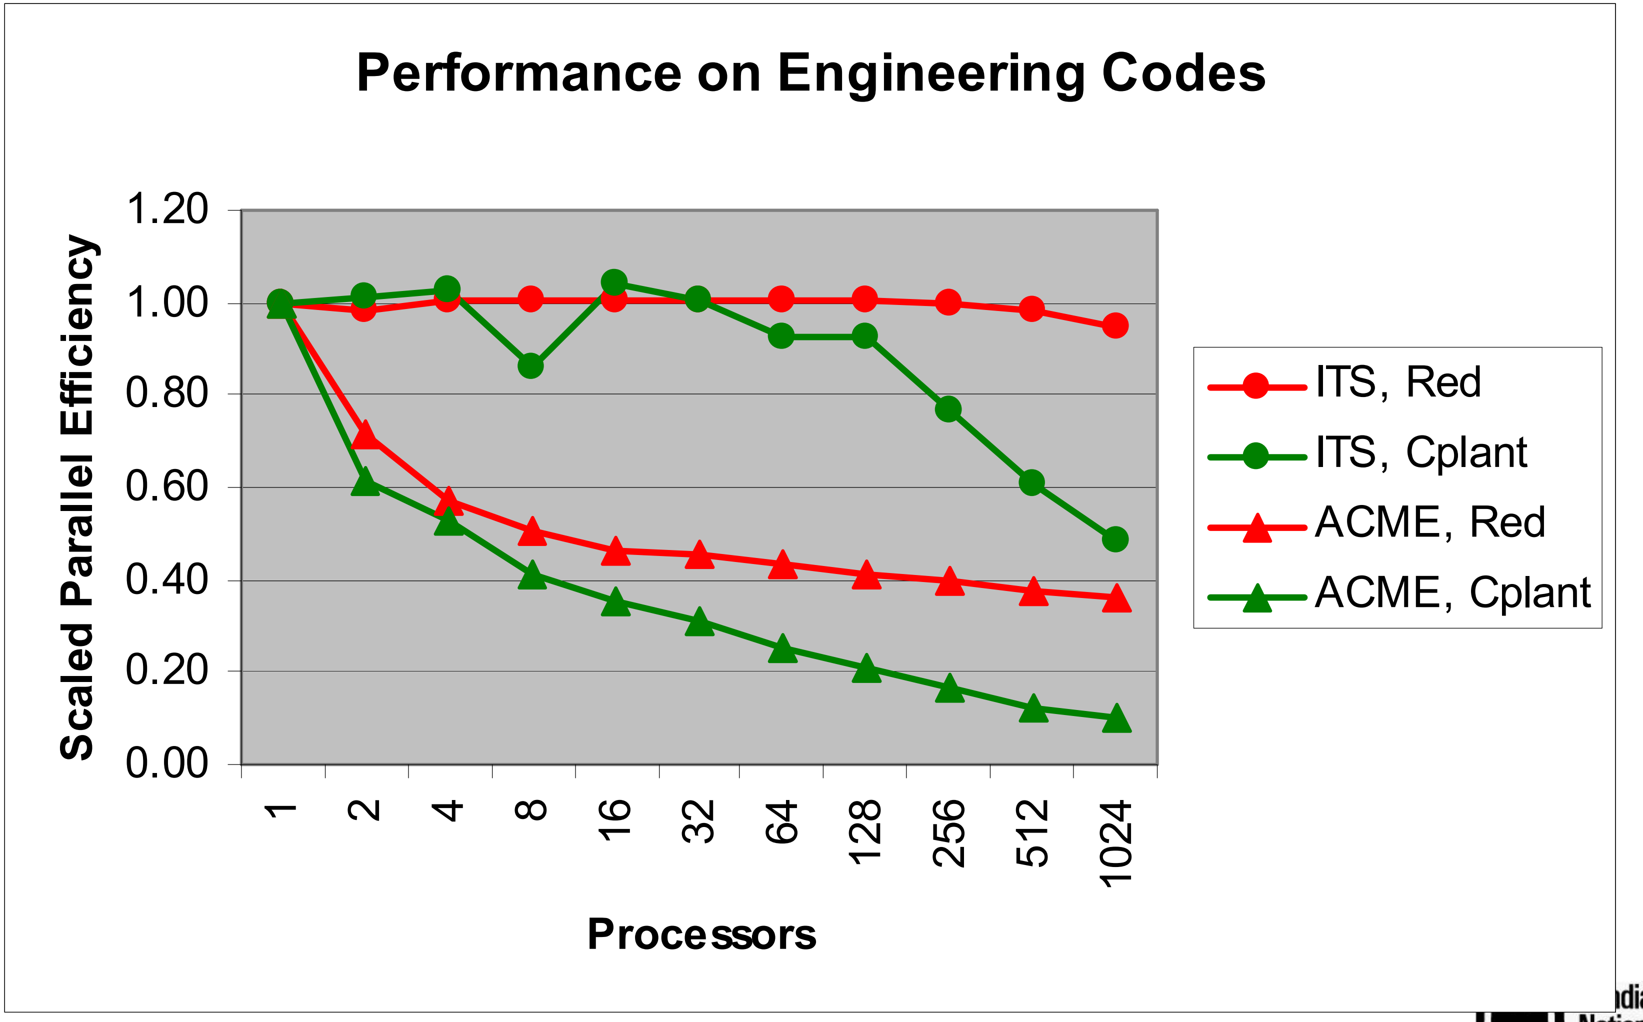

Sandia National Laboratories

Scaling Data for a Radiation Transport Code

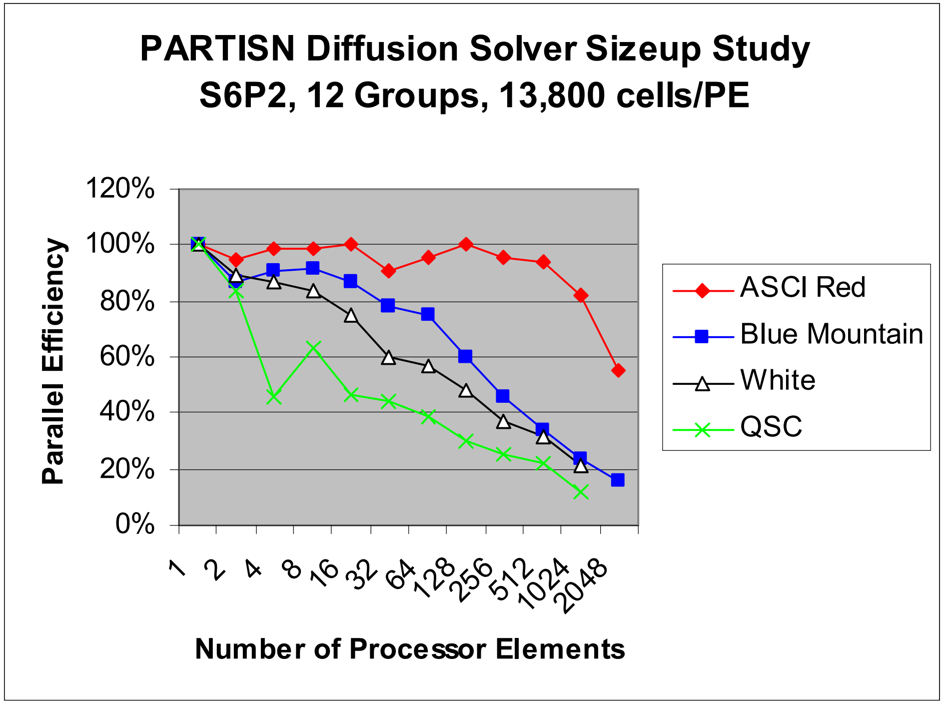

Radiation transport code

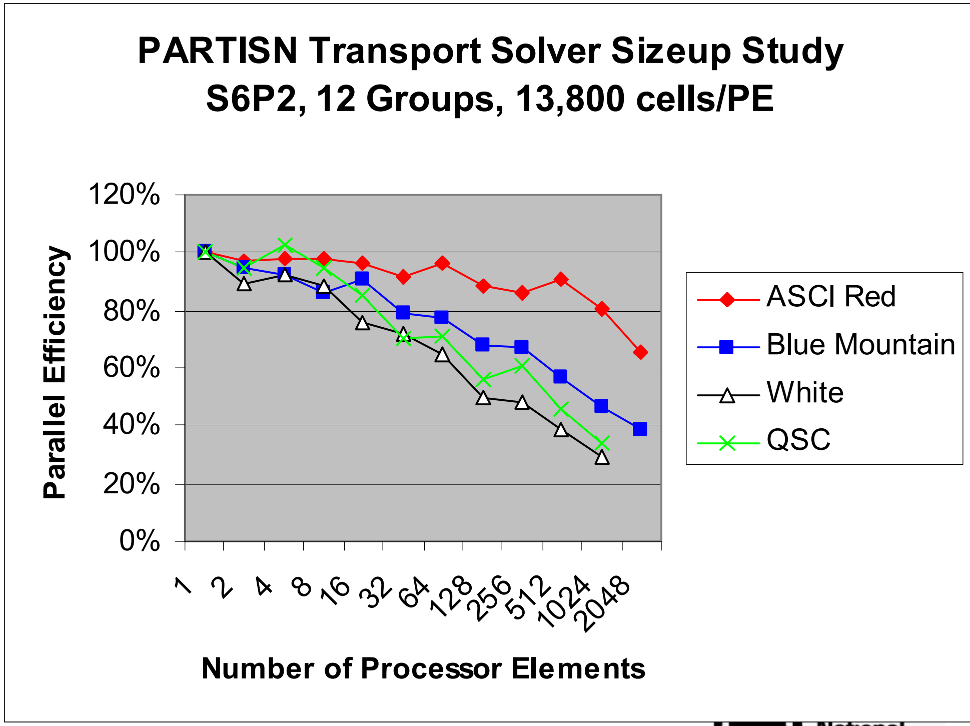

Relating Scalability and Cost

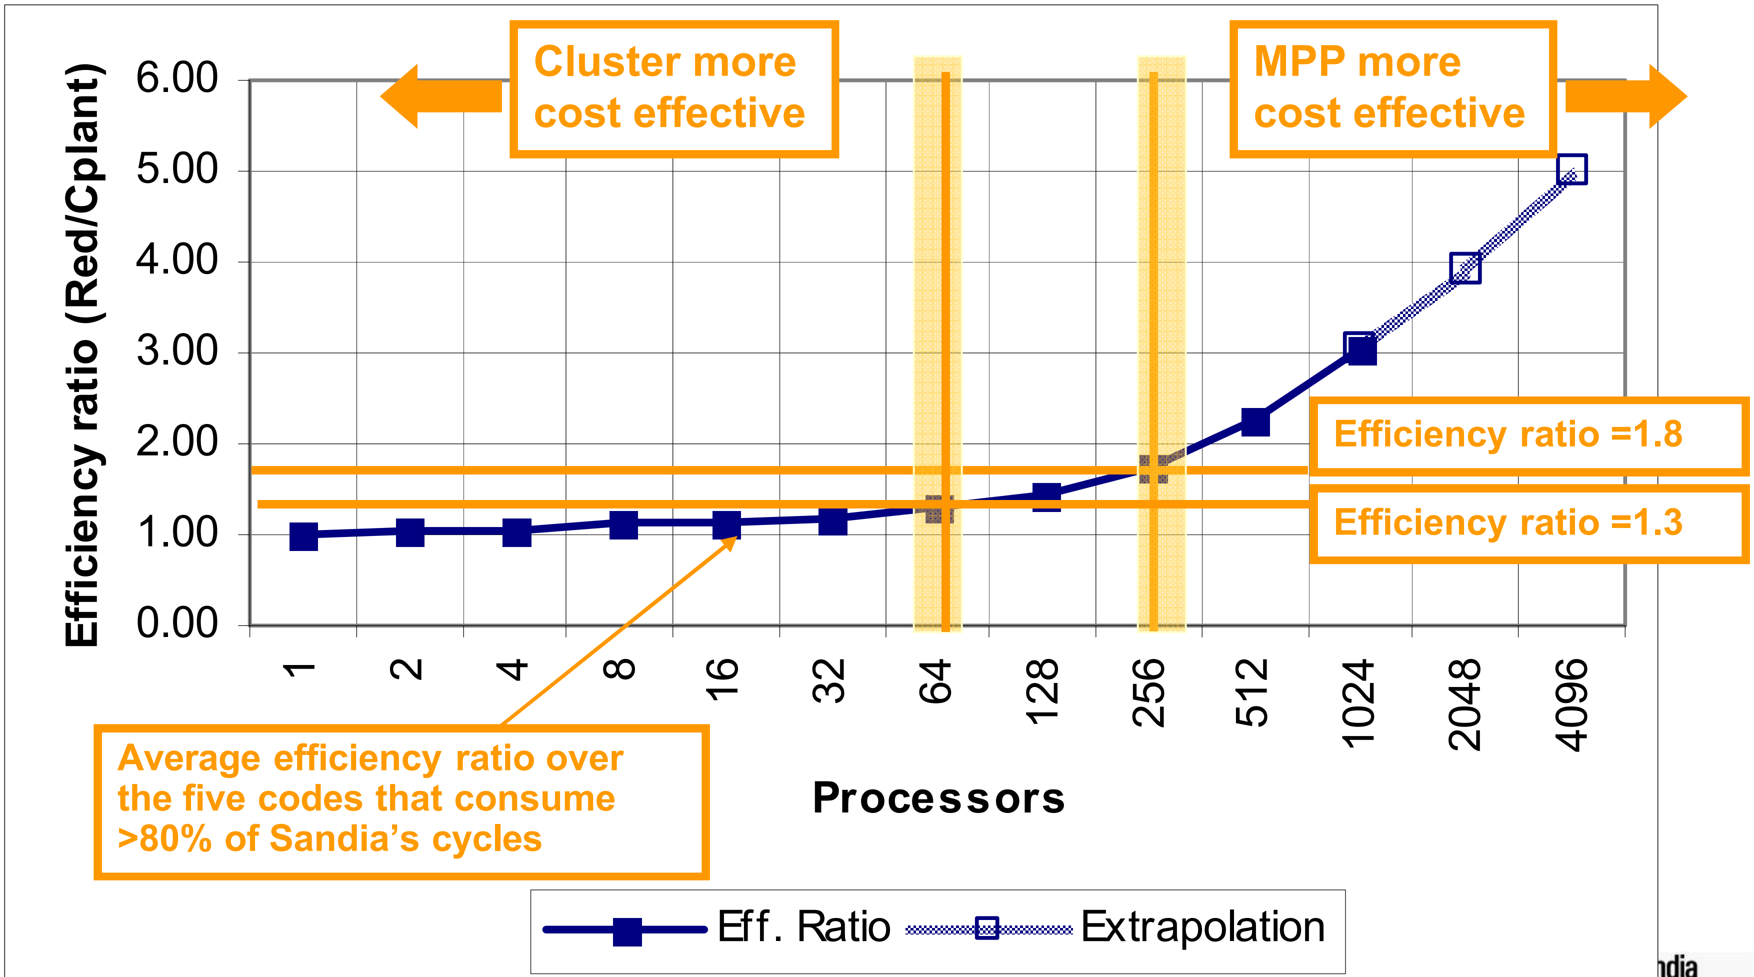

Sandia National Laboratories

Scalability determines cost effectiveness

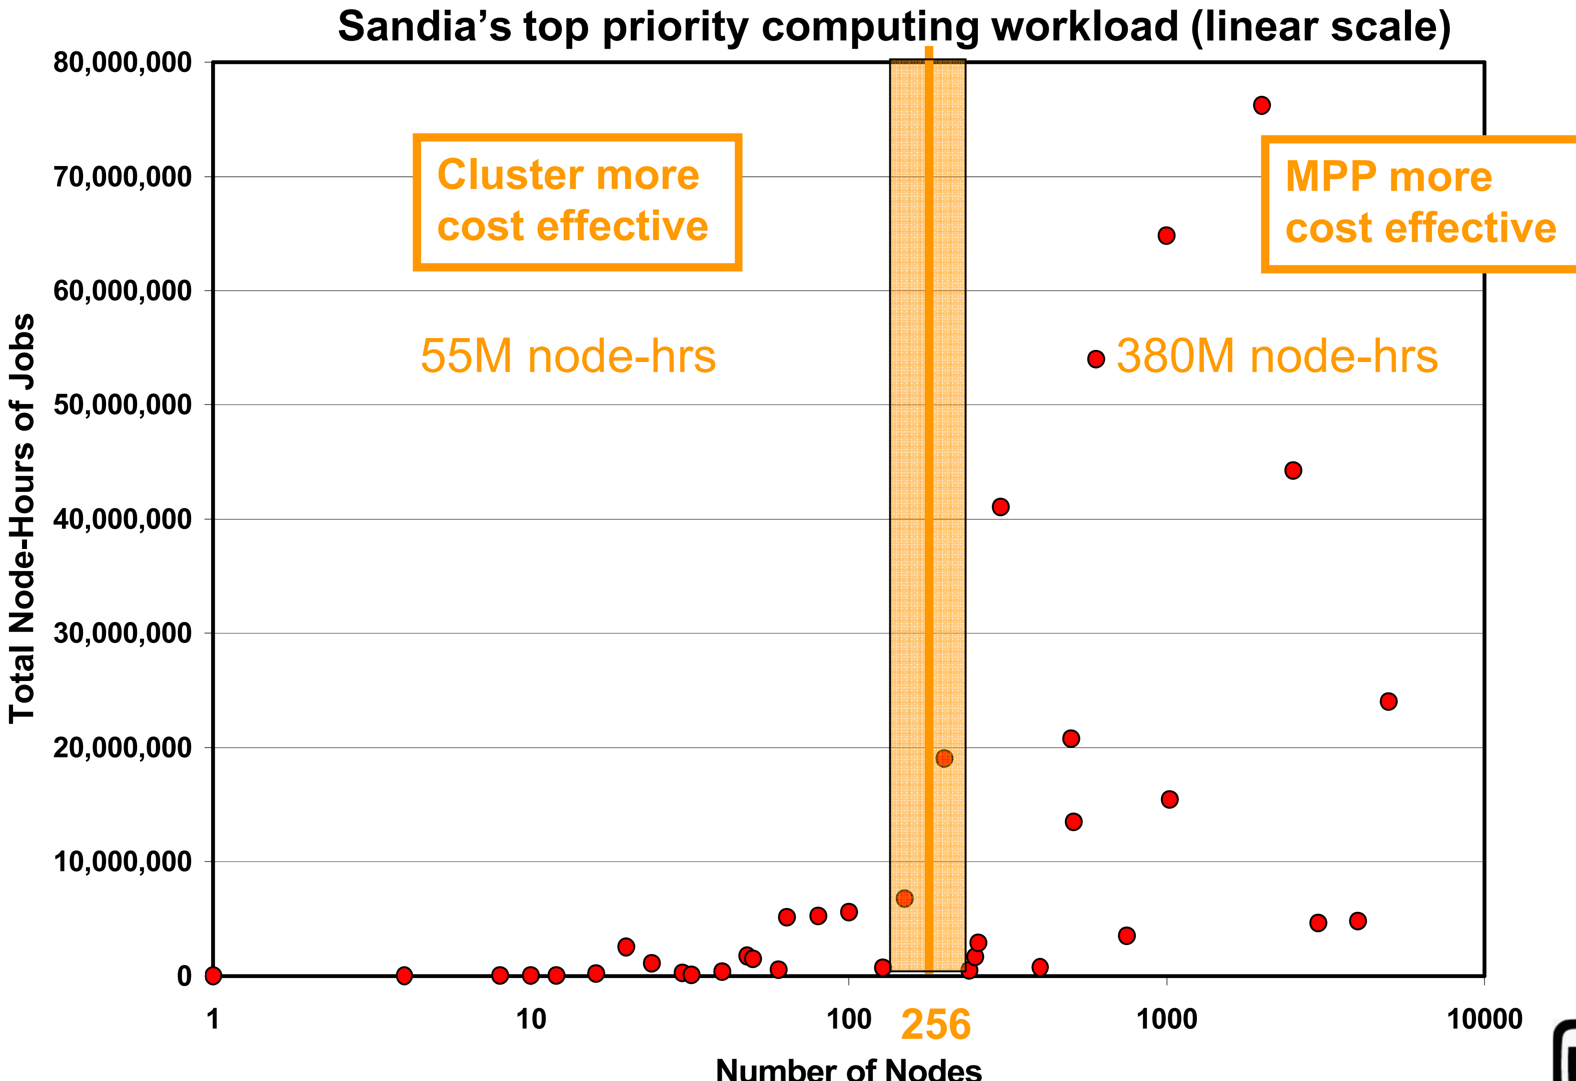

Capability computing does really work

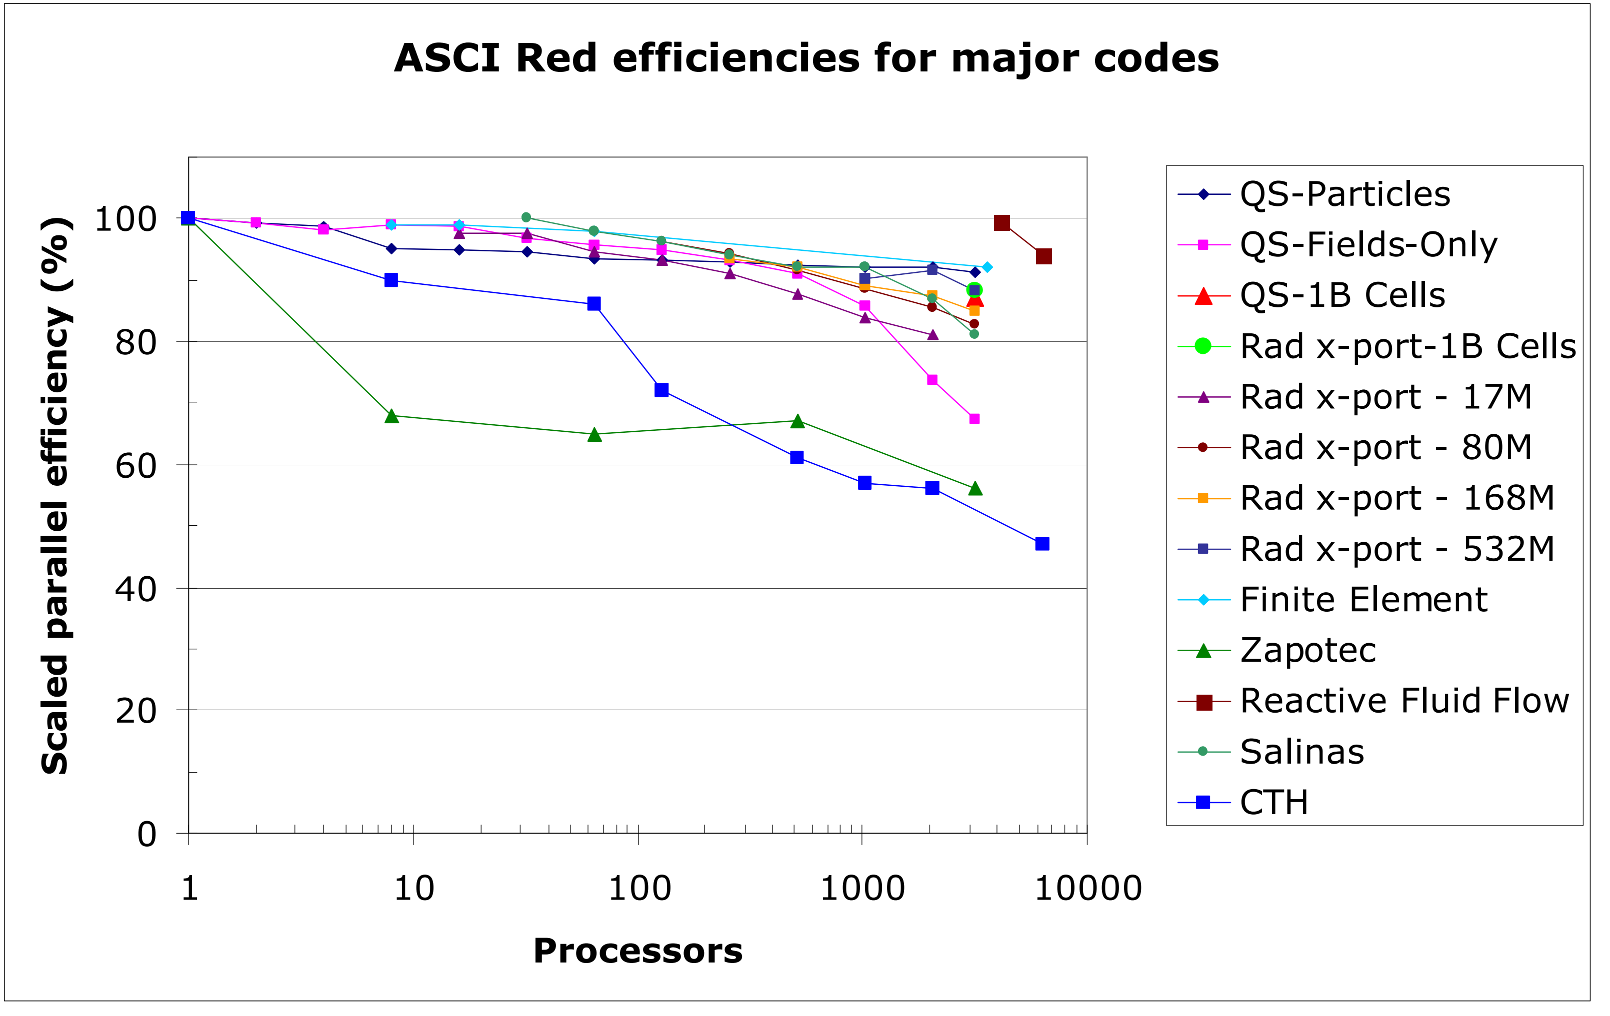

Sandia National Laboratories

Key Factors needed to Achieve High Parallel Efficiency

System Architecture - Balance
- Network Performance
  - Bandwidth
  - Latency
  - Message Throughput
- Minimizing Operating System Noise
- Avoiding Serial Bottlenecks
  - Collectives
  - I/O

Application Codes
- Keep parallelism that is in underlying problem
- Maximize computation to communication ratio

Computer Architectures - Tomorrow

Workstations
- Faster Processors
- Multiple-processors on a chip

Linux Clusters
- Distributed memory MIMD
- A few tens of processors to a few thousand processors
- Multiple-processors on a chip
- Scalable performance to large processor counts will still be an issue

Large Massively Parallel Processor Systems (MPPs)
- Distributed memory MIMD
- 4 to 5 times as many processors as today
- Designed to provide high parallel efficiency
- Appropriate for simulations involving thousands of processors

Wide Area Networks

Today

- ASC has a 10 Gbit network connecting LANL, LLNL, and Sandia
- Bandwidth is not sufficient to move large amounts of data around - Need to keep most data with the machine
- Provides sufficient performance to move processed results
- Price is controlling factor - electronics not fiber

Tomorrow

- 100 Gbit may be possible
- Bandwidth will still be to low to support large scale data movement
- Price will still be the controlling factor - electronics not fiber

Sandia National Laboratories

Conclusions

- Large scale parallel computing has been demonstrated to work for performing complex, high resolution scientific and engineering simulations
- The explicit message passing programming model will remain the programming model for at least the next several years.
  - MPI has become a de-facto standard because it works and MPI codes are portable
  - There is a huge investment in explicit message passing codes