Analyzing Multicore Characteristics for a Suite of Applications on an XT5 System

Courtenay T. Vaughan and Douglas W. Doerfler
Sandia National Laboratories

Cray User Group
May 2010

Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin company, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Motivation

- Study applications on multicore multisocket machine
- Preparation for Cielo
- For details see the talk by Sudip Dosanjh (SNL) and John Morrison (LANL), “An Alliance for Computing at the Extreme Scale” on Thursday at 11:30

Sandia National Laboratories

XT5 Details

- Cray XT5 with 160 compute nodes
- dual socket with 6 core AMD Istanbul processors
- 2.4 GHz processors
- 32 GB of 800 MHz DDR2 Memory per node
- 6 x 4 x 8 3D torus with SeaStar 2.2
- CNL 2.2.41
- PGI 9.0.2
- Panasas filesystem

Sandia National Laboratories

Applications

- CTH - shock hydrodynamics
- Charon - semiconductor device modeling
- SAGE - hydrodynamics with AMR
- xNOBEL - hydrodynamics with adaption and high-explosive burn
- AMG - 3D, deterministic, multigroup, photon transport code for unstructured meshes
- UMT - parallel algebraic multigrid solver

Sandia National Laboratories

Comparison with Red Storm

- Comparison with quad-core nodes of Red Storm
  - 2.2 GHz with 8 MB of 800 MHz DDR2 memory per node
  - Catamount (CNW 2.1.56.1)
  - SeaStar 2.2
  - PGI 10.0.0
- Comparison using CTH

Sandia National Laboratories

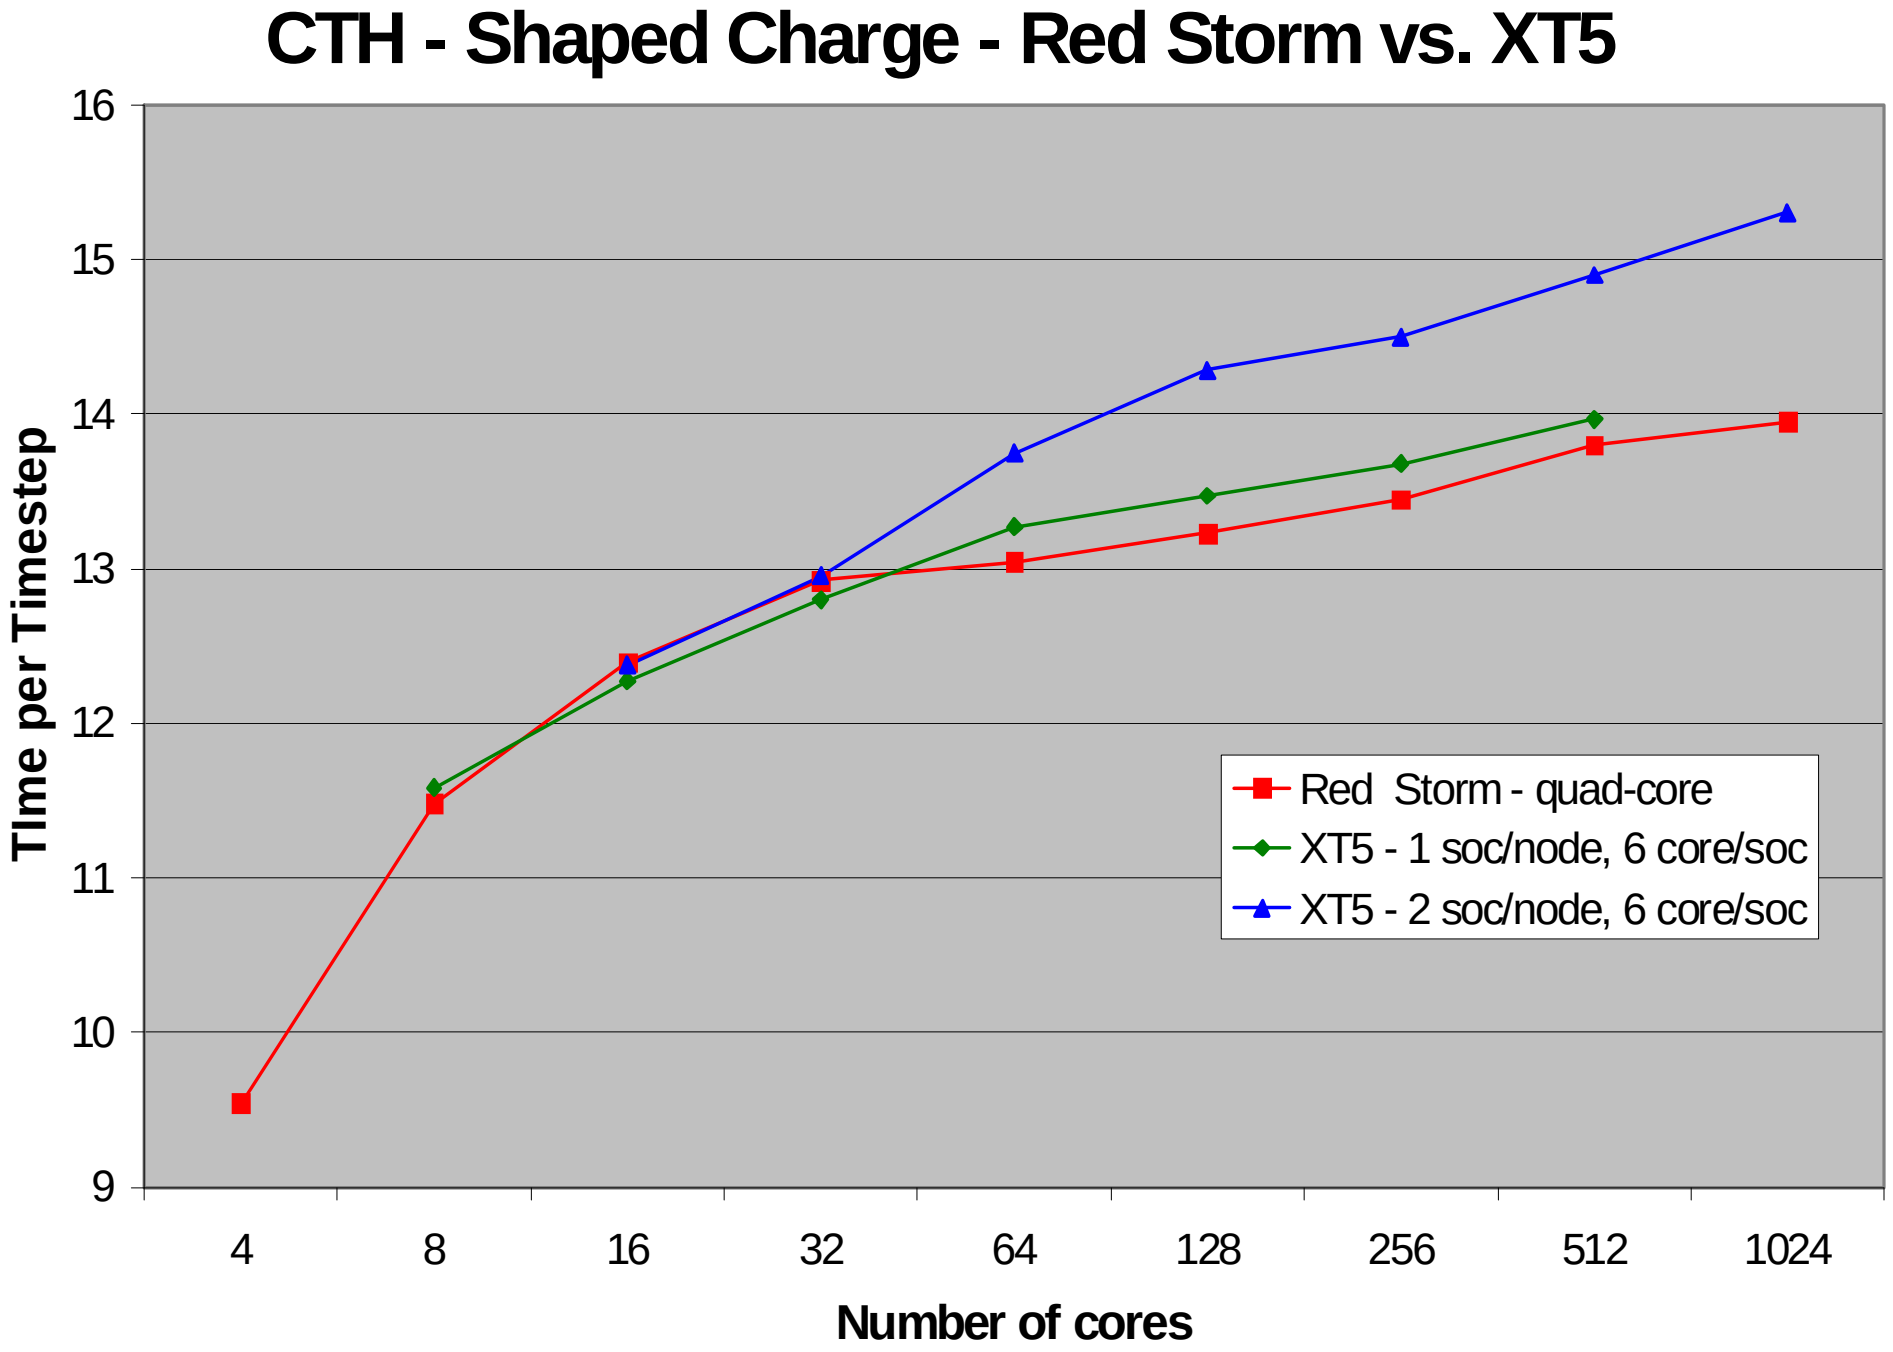
CTH - Shaped Charge - Red Storm vs. XT5
Time per Timestep
16
15
14
13
12
11
10
9
4
8
16
32
64
128
256
512
1024
Number of cores
Red Storm - quad-core
XT5 - 1 soc/node, 6 core/soc
XT5 - 2 soc/node, 6 core/soc

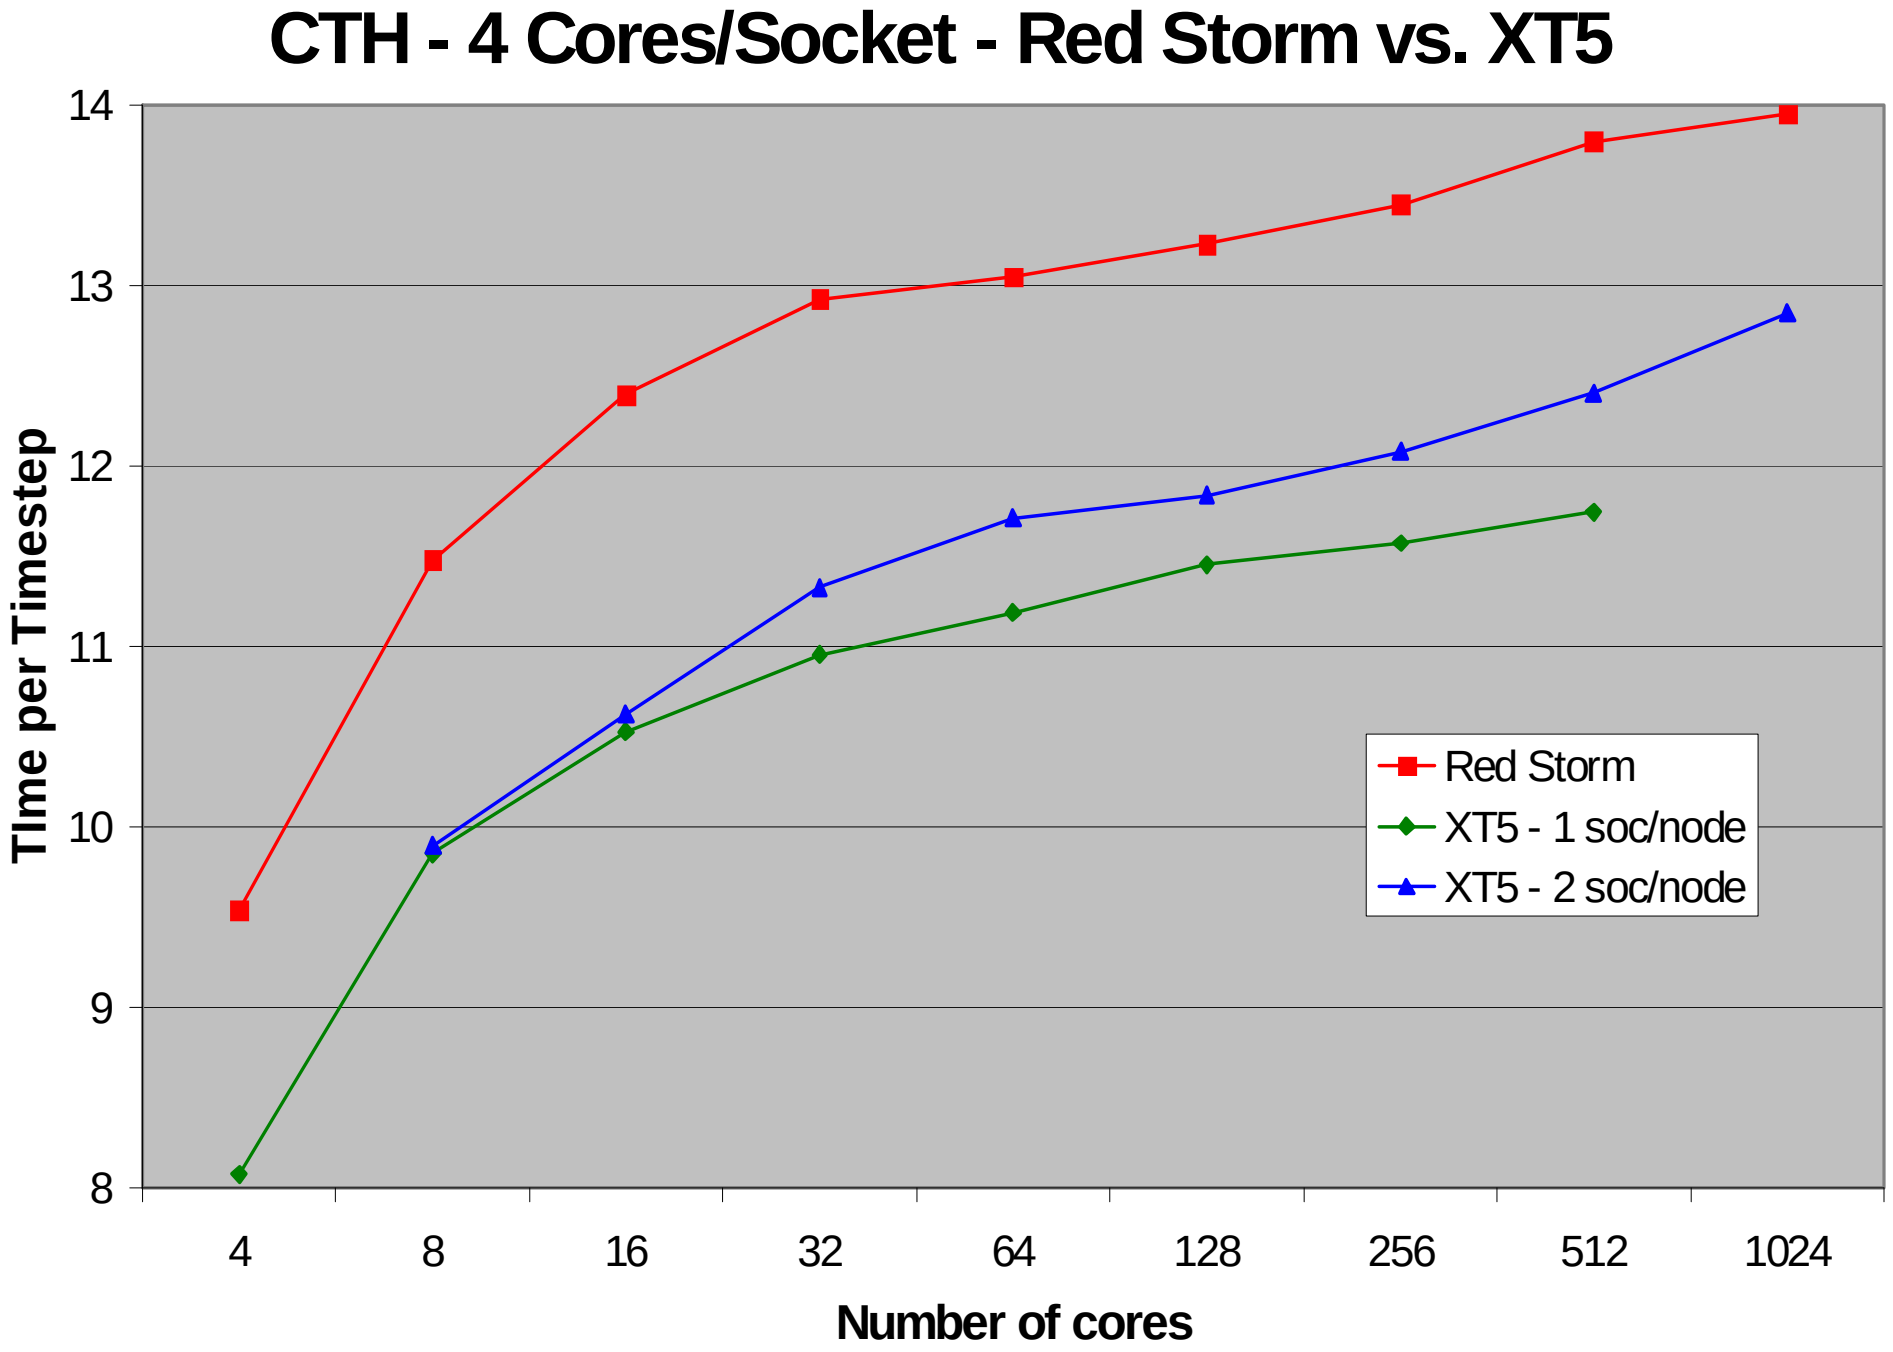
CTH - 4 Cores/Socket - Red Storm vs. XT5
TIme per Timestep
14
13
12
11
10
9
8
4
8
16
32
64
128
256
512
1024
Number of cores
Red Storm
XT5 - 1 soc/node
XT5 - 2 soc/node

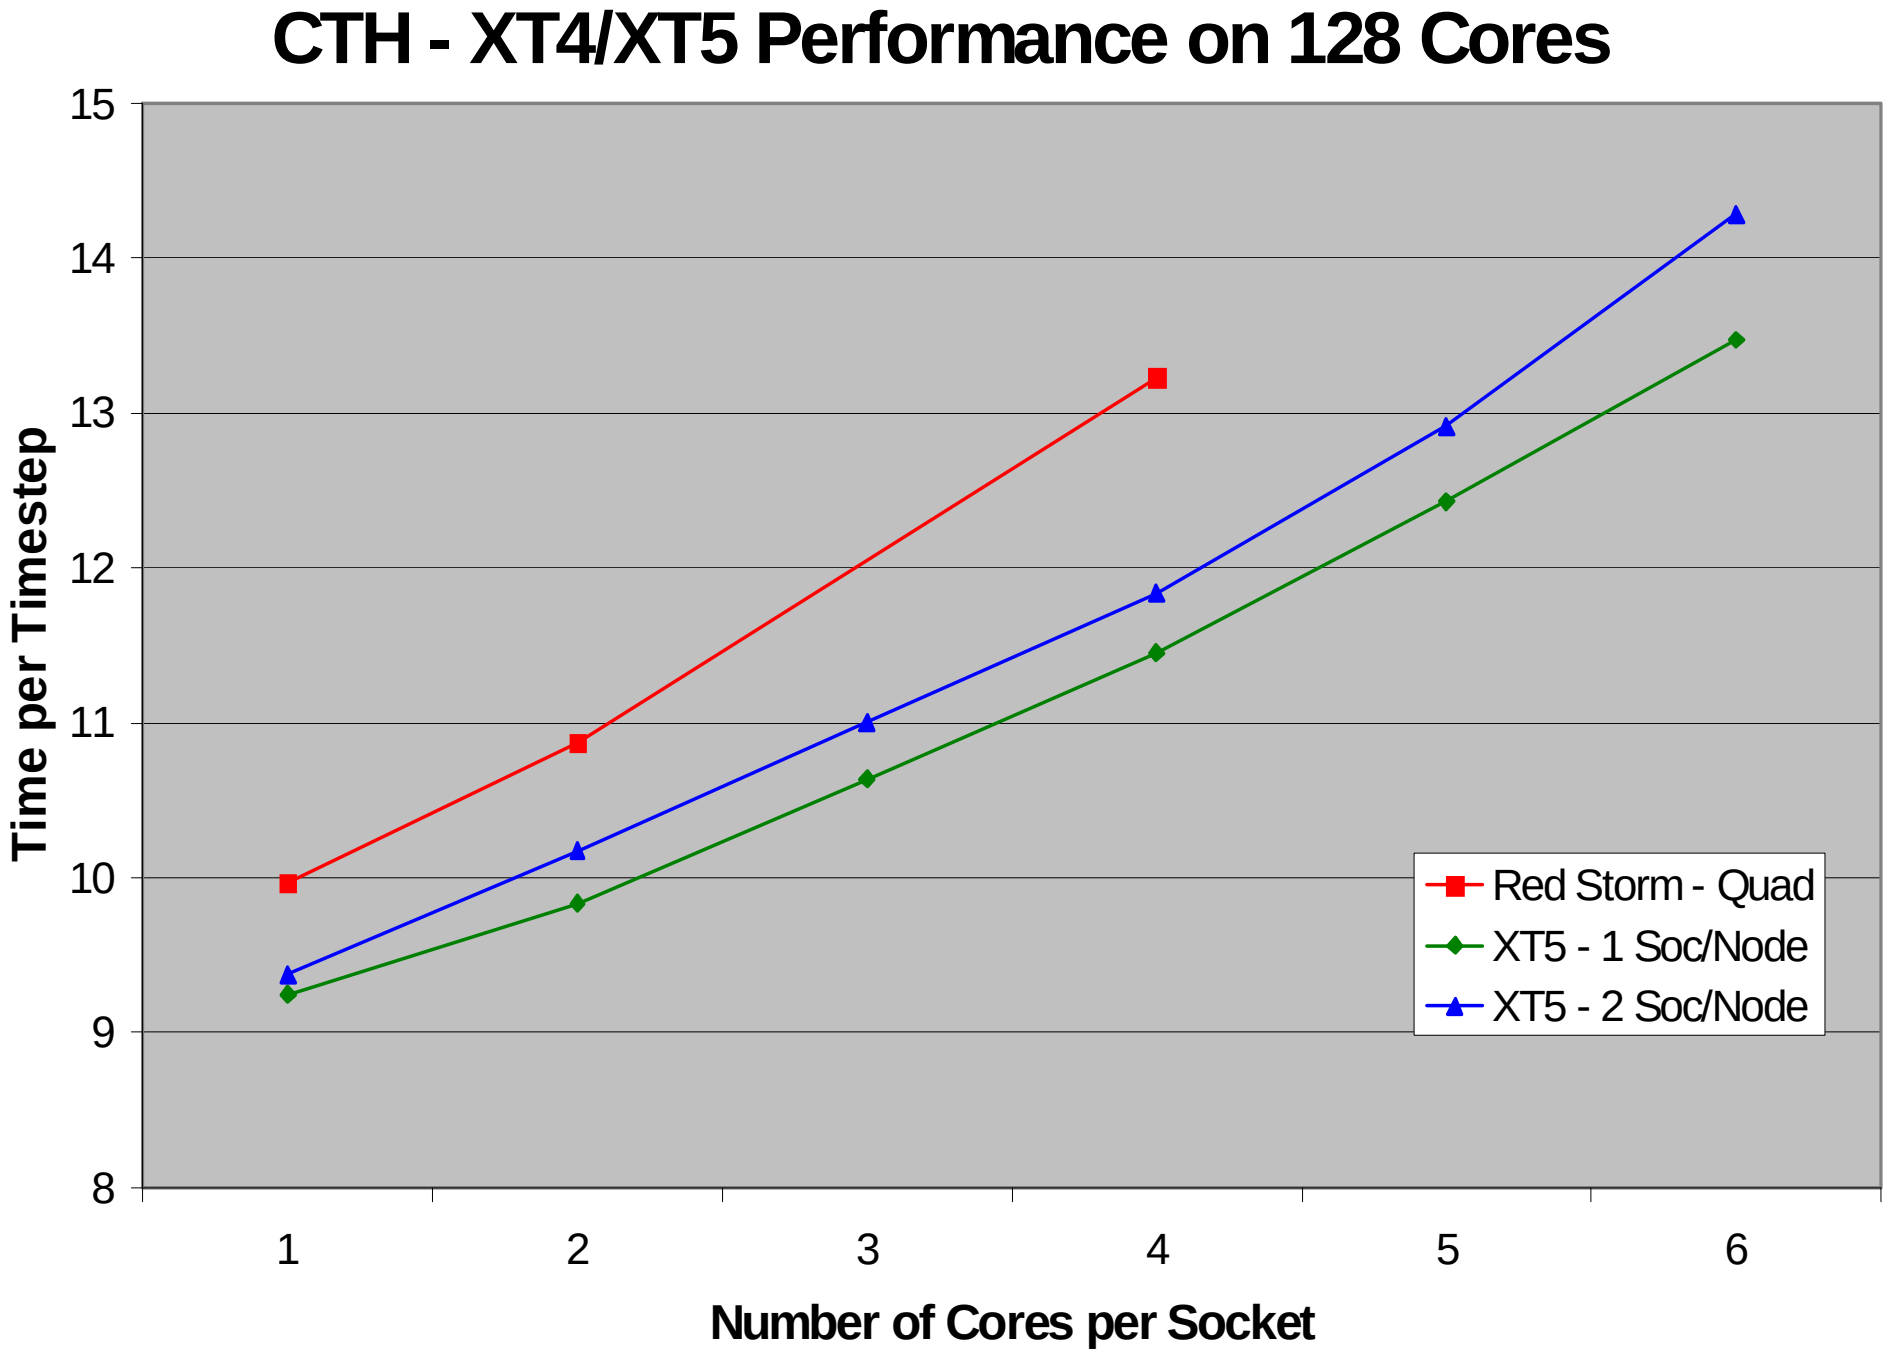
CTH - XT4/XT5 Performance on 128 Cores
Time per Timestep
15
14
13
12
11
10
9
8
1
2
3
4
5
6
Number of Cores per Socket
Red Storm - Quad
XT5 - 1 Soc/Node
XT5 - 2 Soc/Node

Application Performance on XT5

- Ran codes on 128 cores while varying the number of cores per socket and number of sockets per node
- Used Cielo benchmark “figure of merit” for the codes
  - Converted to time
- Graphed performance relative to running on one core per node
- Used only MPI level parallelism

Applications on 128 Cores

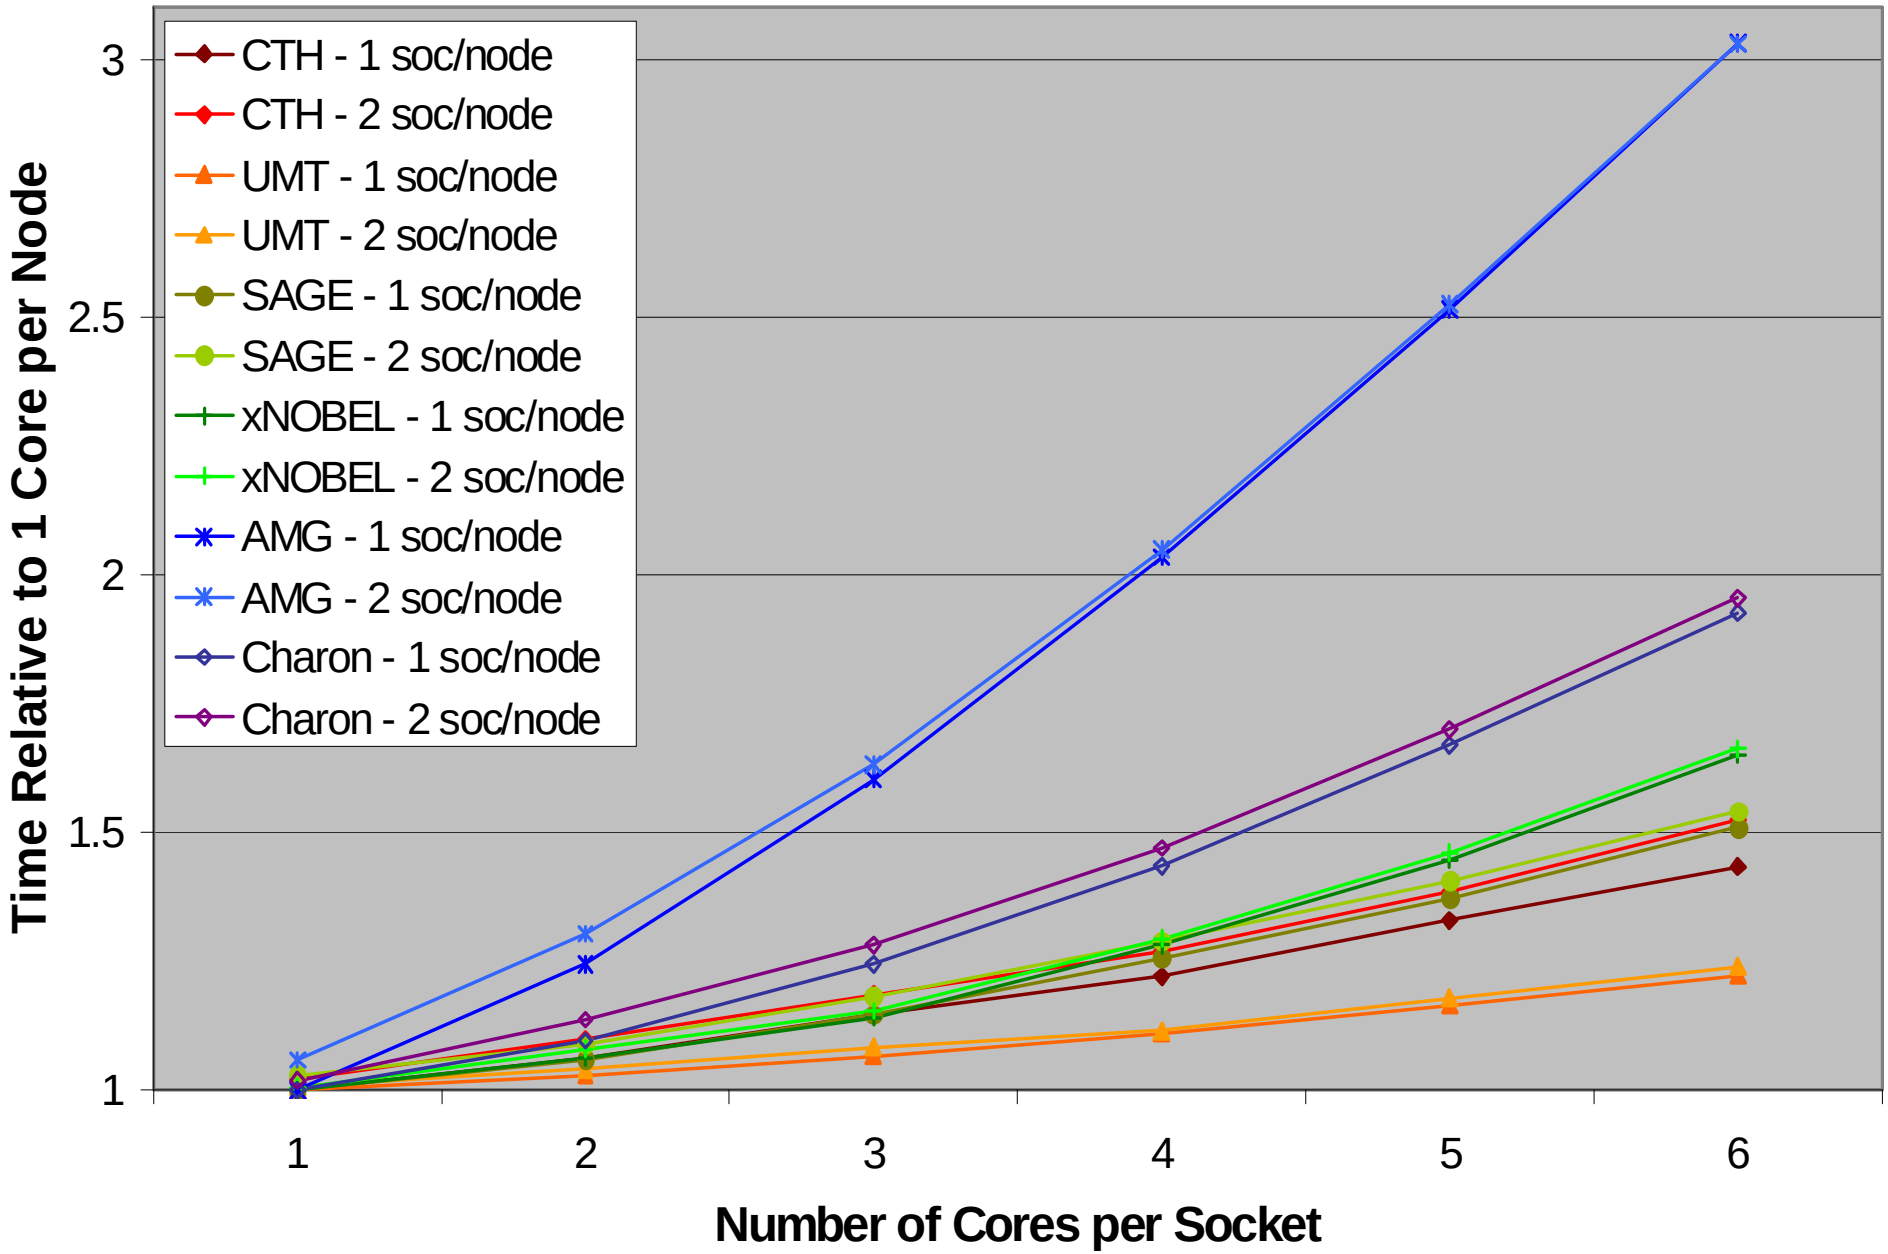

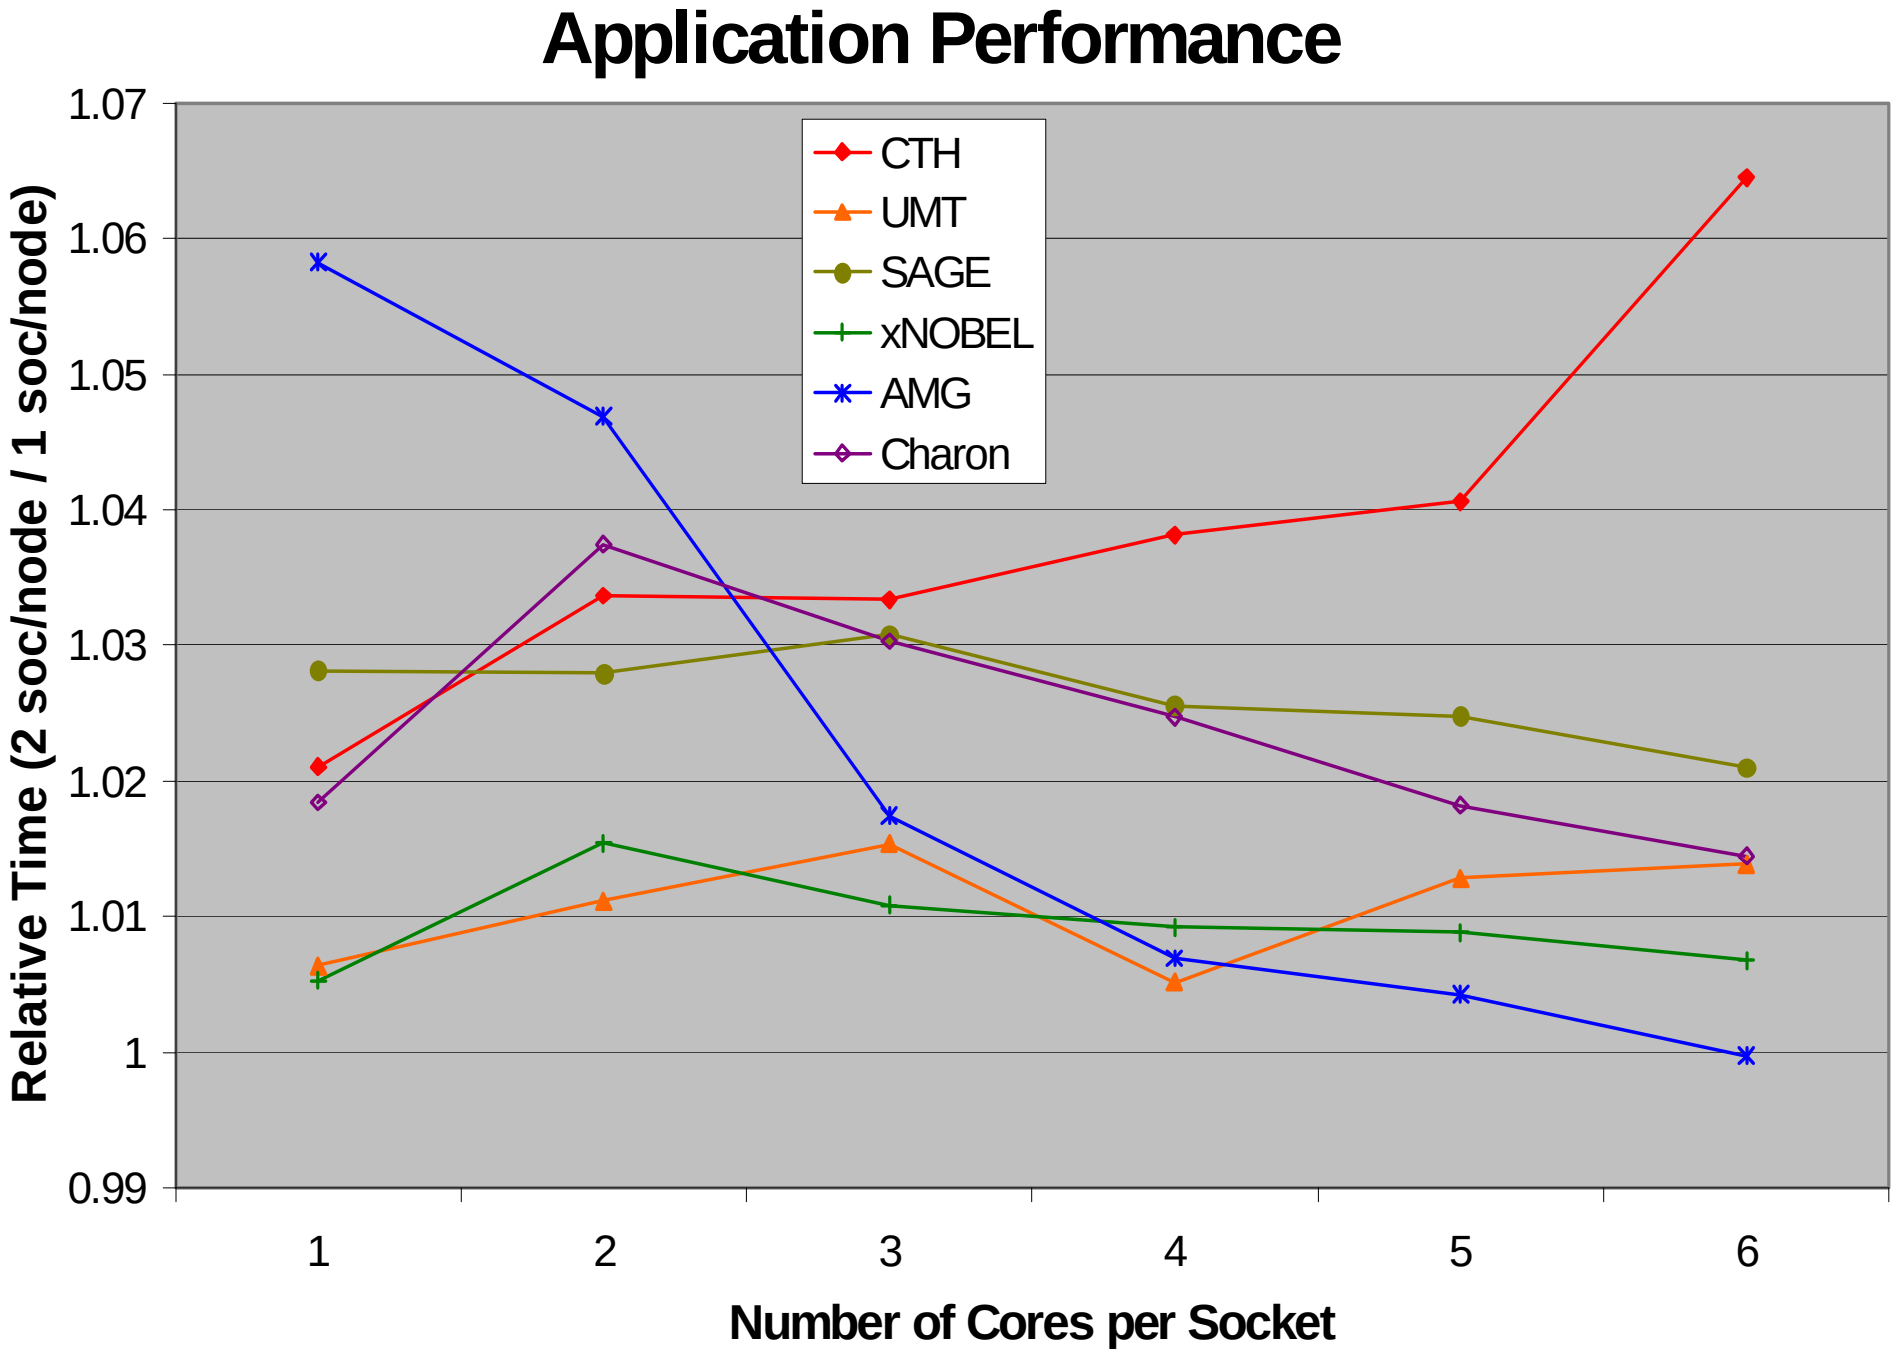
Application Performance
Relative Time (2 soc/node / 1 soc/node)
1.07
1.06
1.05
1.04
1.03
1.02
1.01
1
0.99
1
2
3
4
5
6
Number of Cores per Socket
CTH
UMT
SAGE
xNOBEL
AMG
Charon

AMG Performance

- Most sensitive of the applications to increasing core count
  - L1 cache hit 98.8%, with L2 rate of 11.9%
  - CTH L1 rate 98.9%, with L2 rate of 45.1%
  - reflects nature of memory access for multigrid
- MPI time dominated by MPI_ALLREDUCE
  - amount of time about 10% as number of cores used increases from one to six
    - Indicates optimization of MPI_ALLREDUCE on node
  - Percentage of time for MPI decreases from about 36% to about 15%

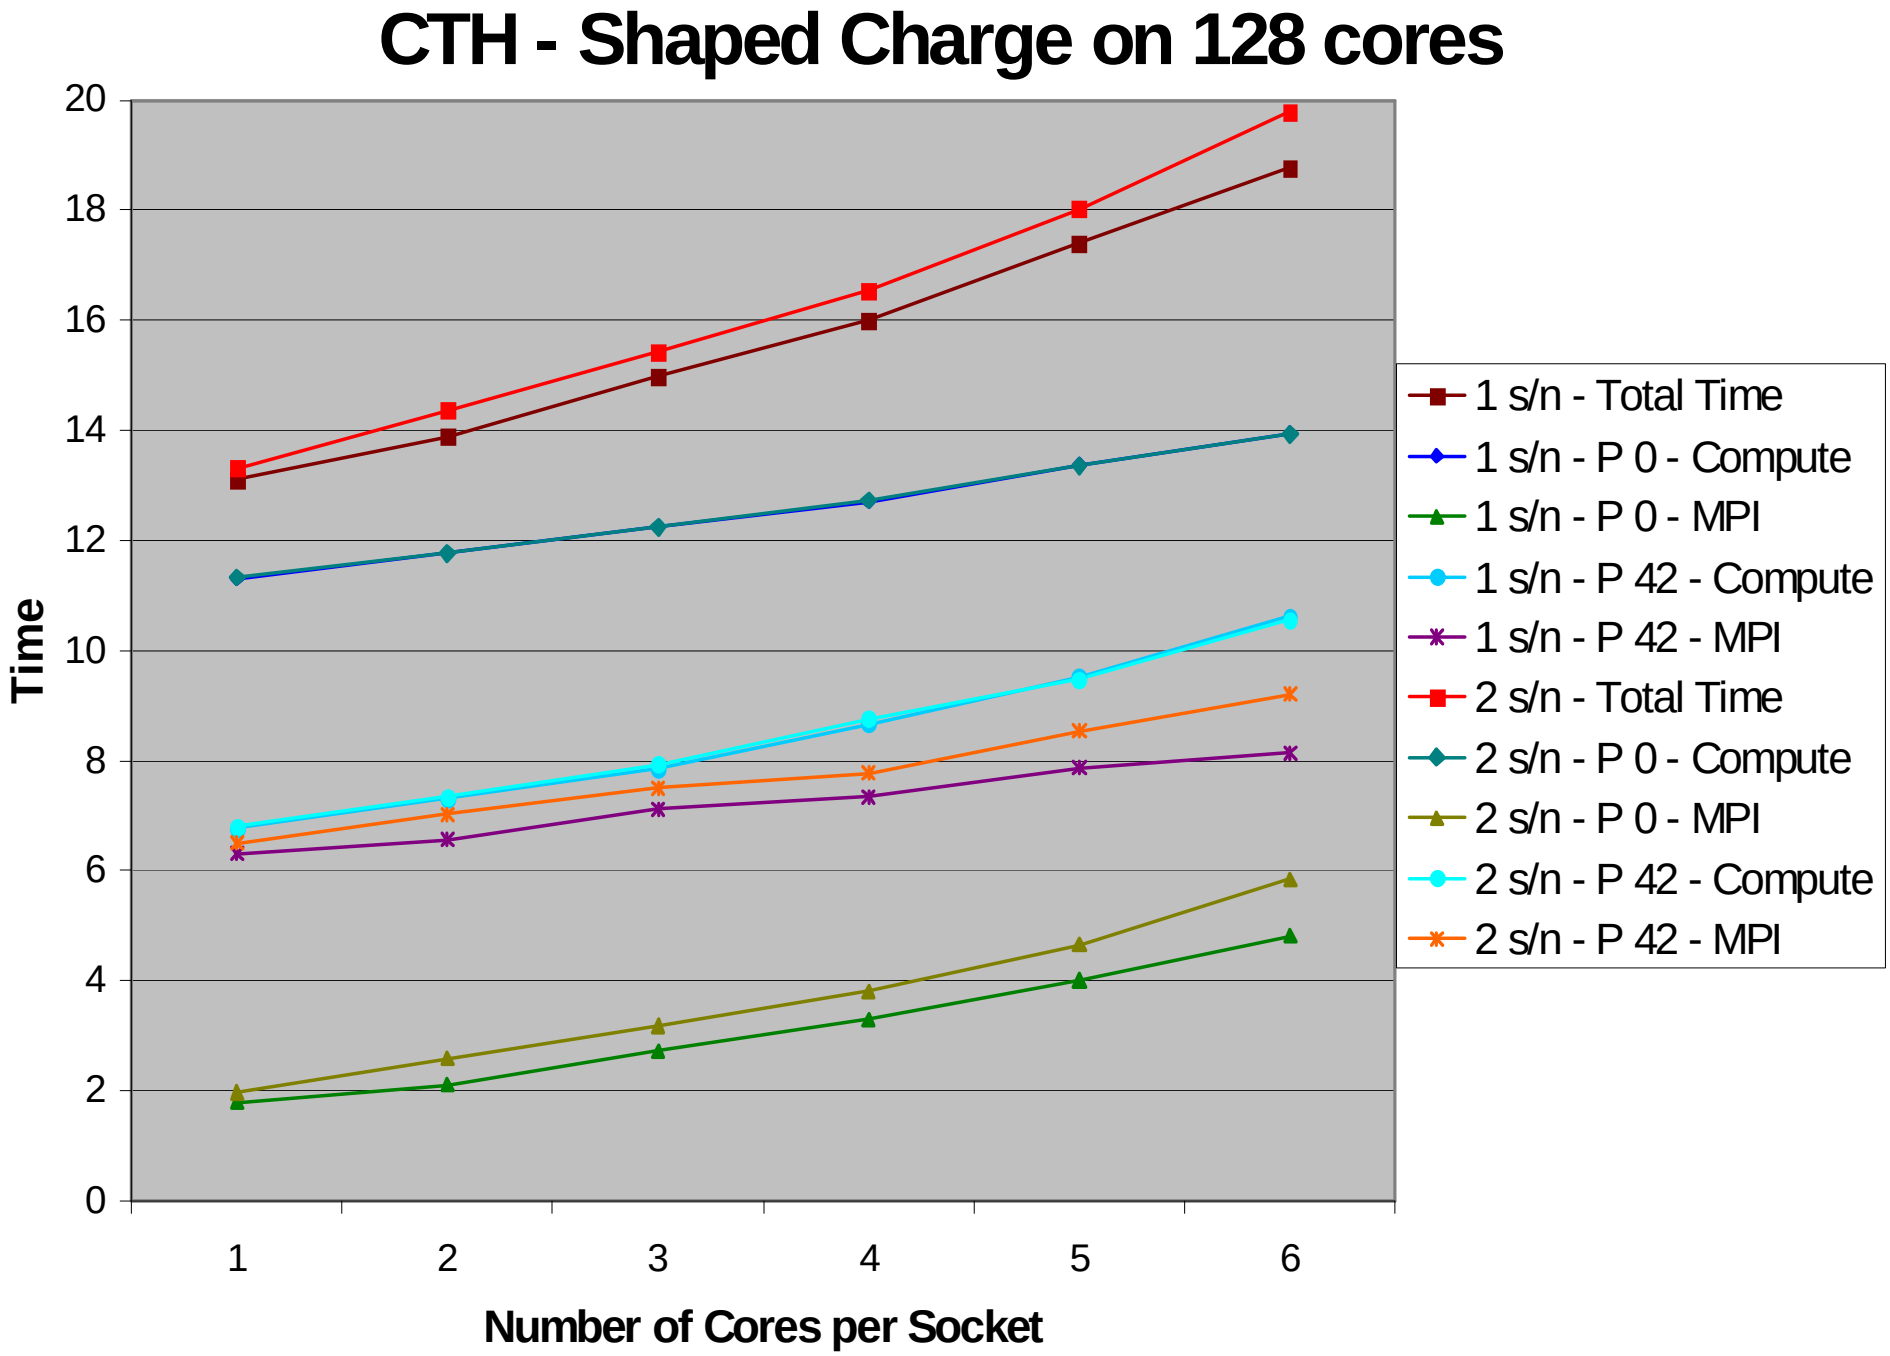
CTH - Shaped Charge on 128 cores
Time
0
2
4
6
8
10
12
14
16
18
20
1
2
3
4
5
6
Number of Cores per Socket
1 s/n - Total Time
1 s/n - P 0 - Compute
1 s/n - P 0 - MPI
1 s/n - P 42 - Compute
1 s/n - P 42 - MPI
2 s/n - Total Time
2 s/n - P 0 - Compute
2 s/n - P 0 - MPI
2 s/n - P 42 - Compute
2 s/n - P 42 - MPI

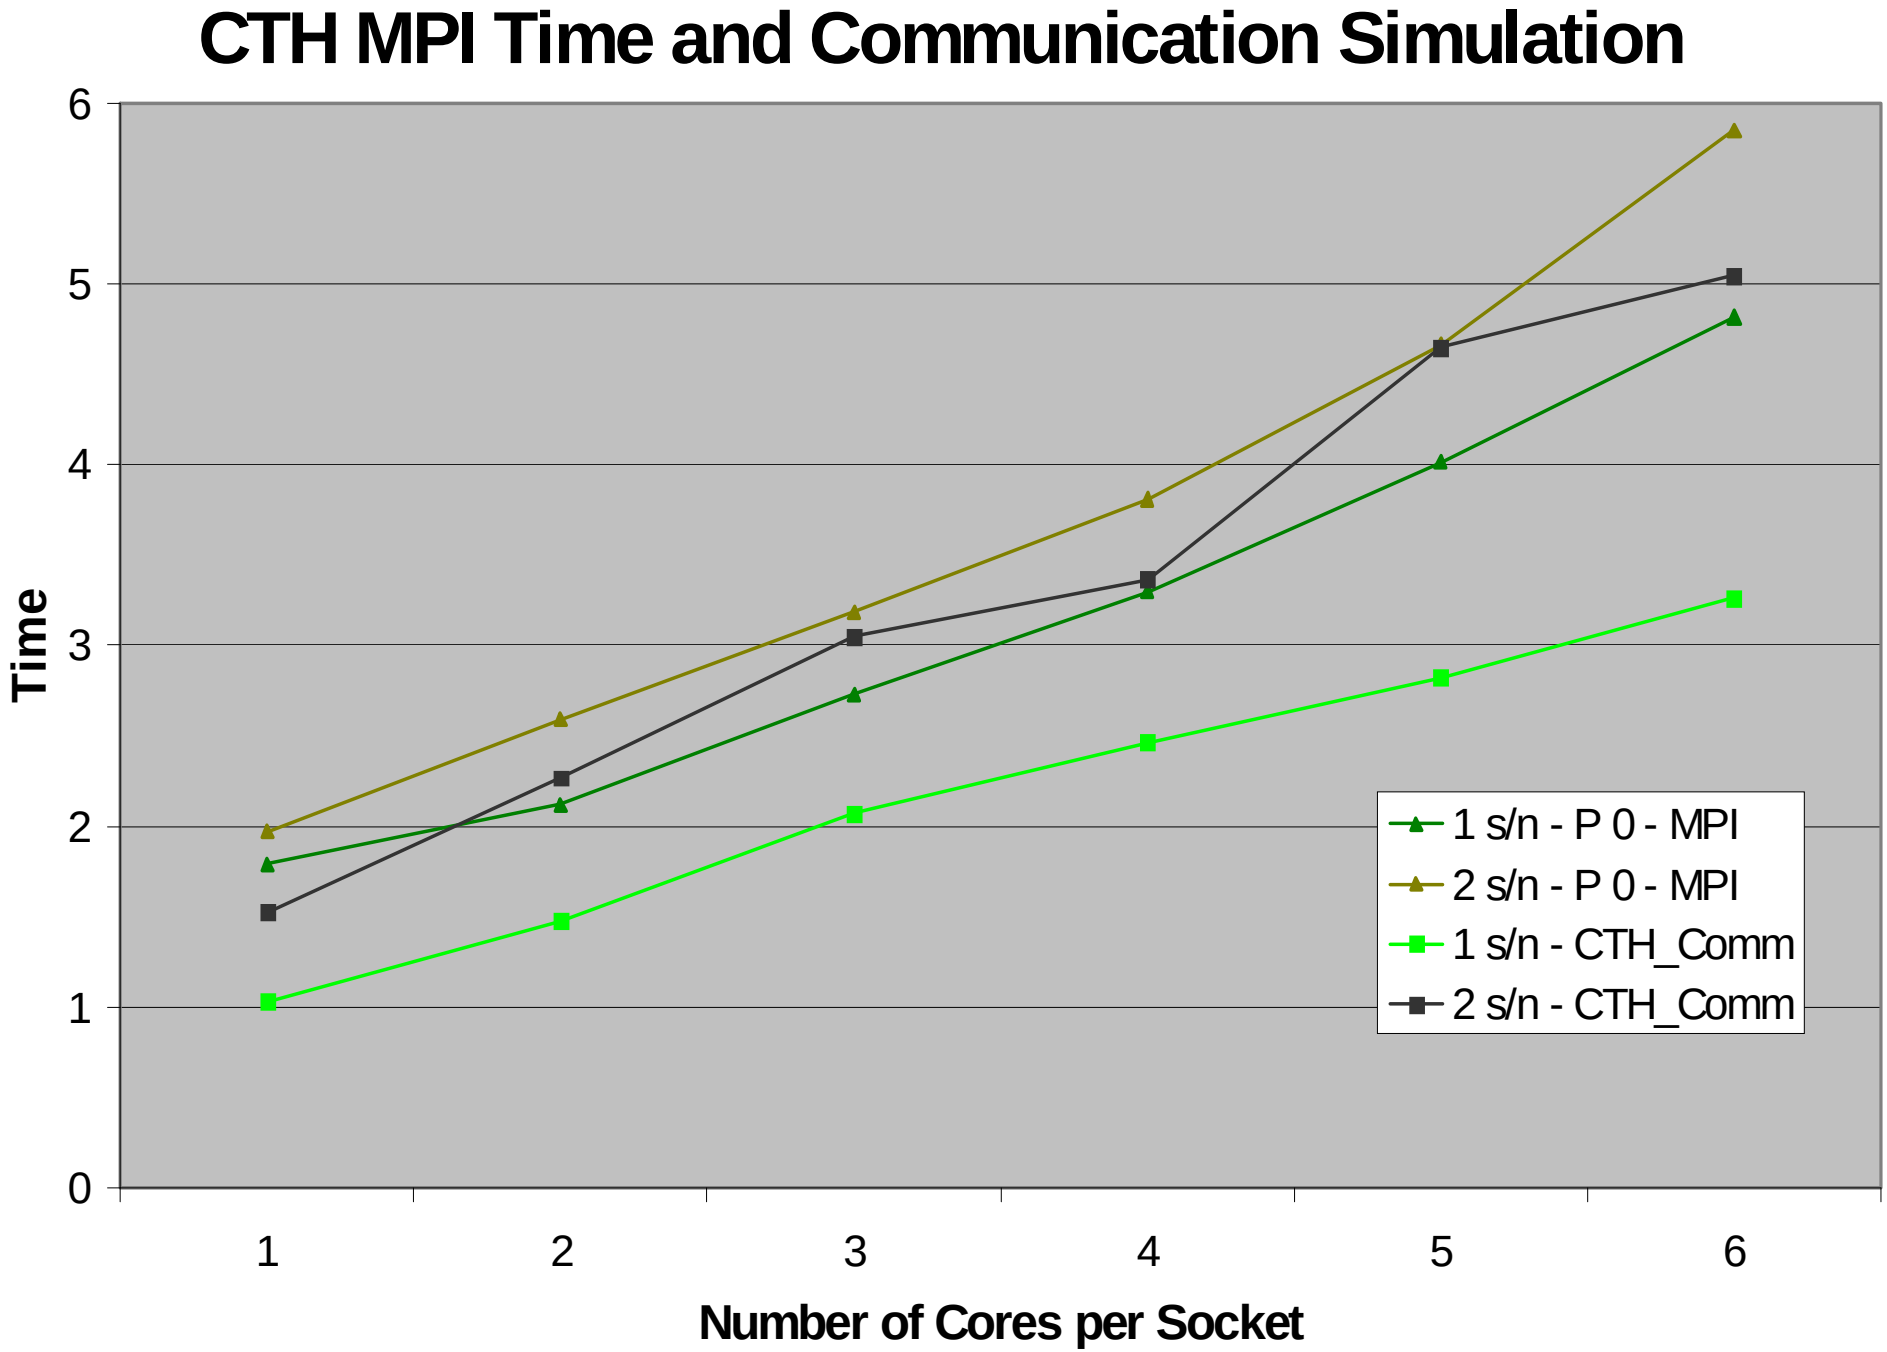
CTH MPI Time and Communication Simulation
Time
0
1
2
3
4
5
6
1
2
3
4
5
6
Number of Cores per Socket
1 s/n - P 0 - MPI
2 s/n - P 0 - MPI
1 s/n - CTH_Comm
2 s/n - CTH_Comm

Conclusions

- Scaling results for XT5 similar to Red Storm
  - As number of cores used increases, difference between single-socket and dual-socket increases
- On 128 cores, codes on average effectively used 6.6 cores out of the available 12 when using all cores on a node
  - Most of the additional time is contention on a socket for memory
  - Most codes show 1% to 3% additional slowdown due to contention for the NIC when using both nodes on a socket
  - CTH also shows additional contention for MPI when using more cores on a socket

Sandia National Laboratories

Future Work

- Use these results to model application behavior
- Use the model to predict behavior of future machines with a larger amount of node parallelism

Sandia National Laboratories