Application Characteristics and Performance on a Cray XE6

Courtenay T. Vaughan
Sandia National Laboratories

Cray User Group
May 2011

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.

Cielo

- Cray XE6 with 6654 compute nodes
- dual-socket oct-core AMD Magny-Cours nodes
- clocked at 2.4 GHz
- 32 GB of 1.333 GHz DDR3 memory per node
- 3D torus with Gemini interconnect
- have large machine and smaller machines

- were configured briefly as XT6 with same nodes and SeaStar interconnect

Sandia National Laboratories

XT5

- 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

Sandia National Laboratories

XE6 node

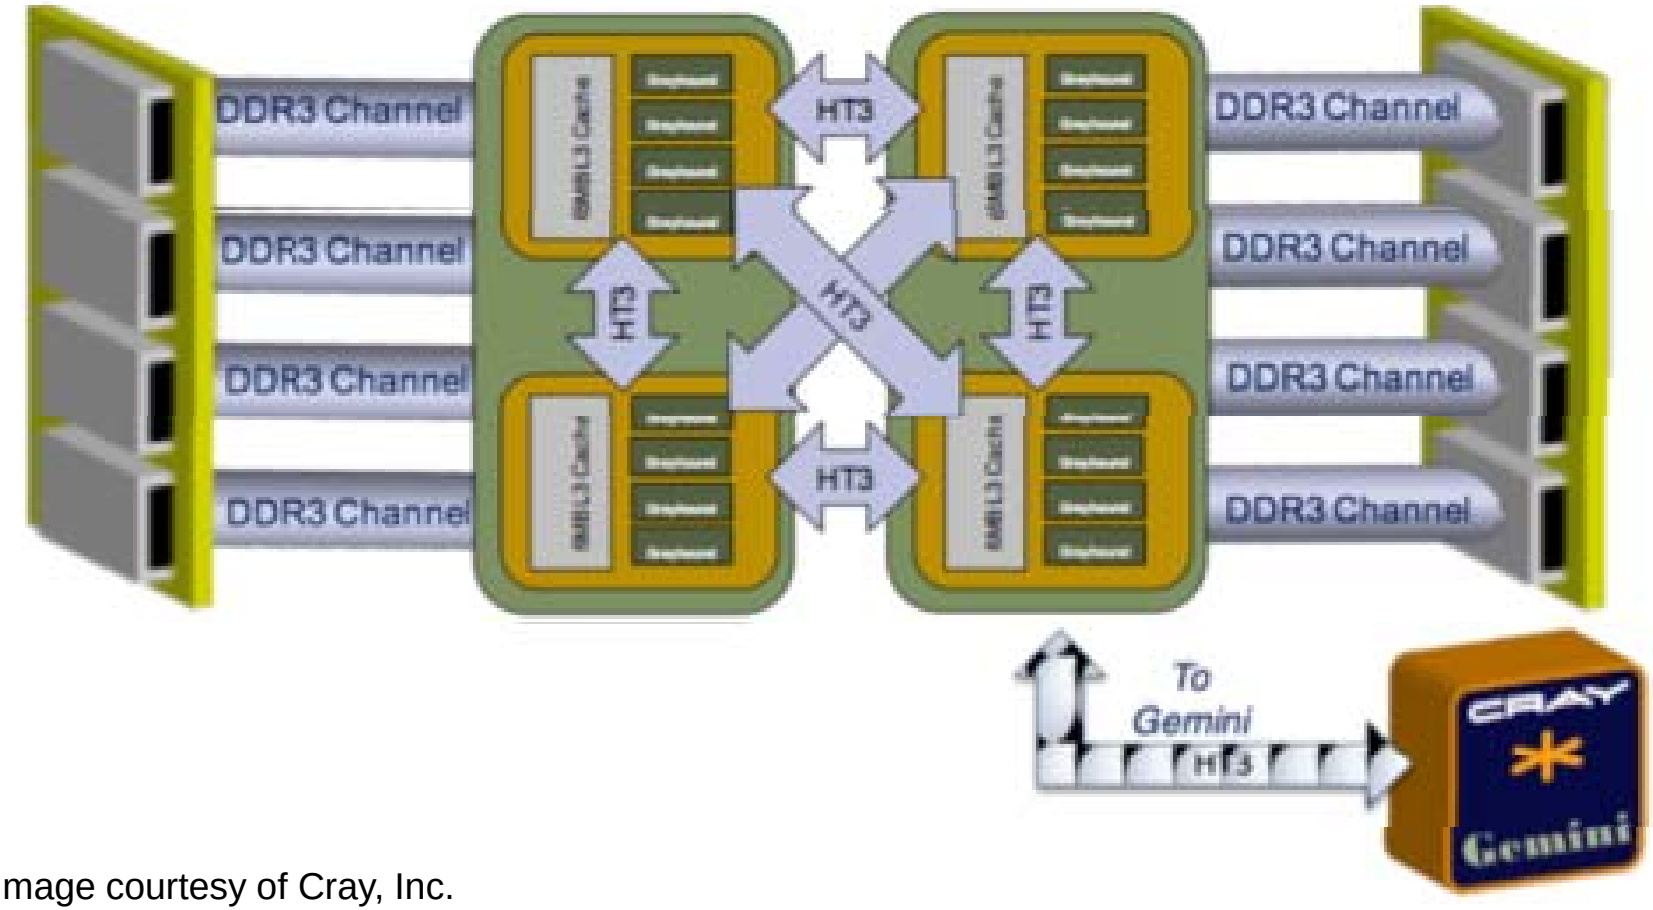

Image courtesy of Cray, Inc.

Sandia National Laboratories

CTH

- Three-dimensional shock hydrodynamics code
- Ran in flat mesh mode - no AMR (Automatic Mesh Refinement)
- Several points in each timestep where each processor sends a few large messages to up to six neighbors
- Messages are aggregated from several variables per cell
- Code is mostly FORTRAN with a little C

Sandia National Laboratories

CTH Problems

- explosively formed Shaped-Charge problem with 4 materials, high explosives, and 90 x 216 x 90 cells/processor in weak scaling mode
  - Messages aggregate 40 variables per cell and average 5.2 MB
- impact Meso-Scale problem with 11 materials and 80 x 80 x 275 cells/processor in weak scaling mode
  - Messages aggregate 75 variables per cell and average 10.4 MB

Sandia National Laboratories

Shaped Charge Problem

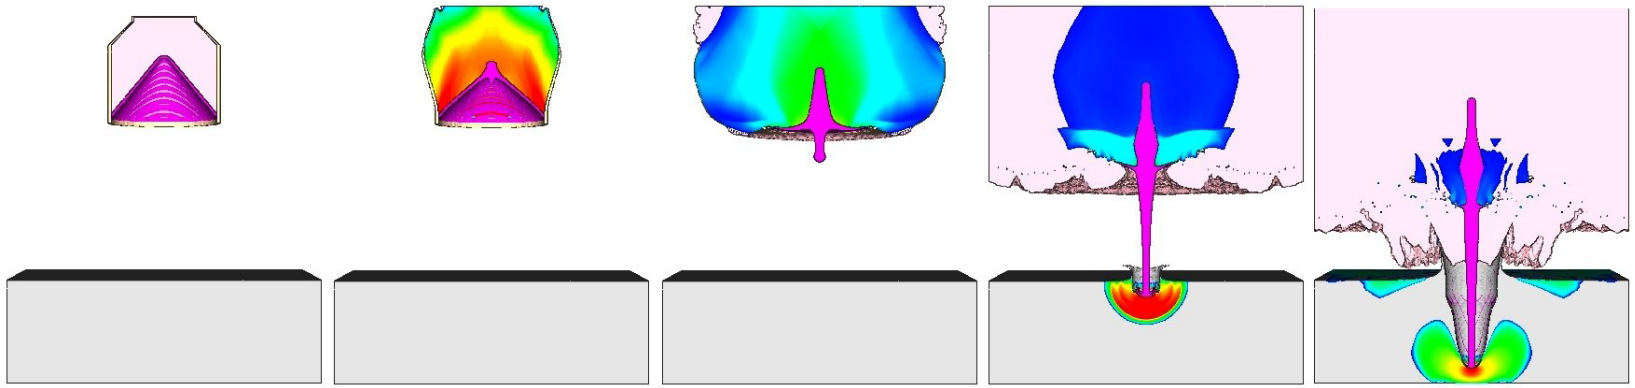

Sandia National Laboratories

CTH Communication matrices on 64 cores

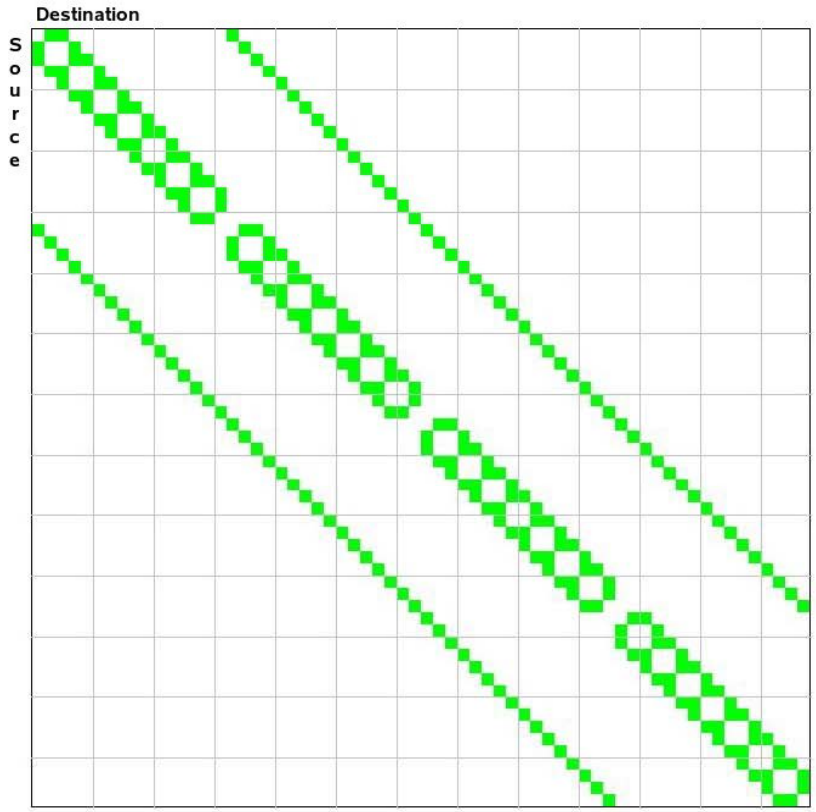

Shaped-Charge

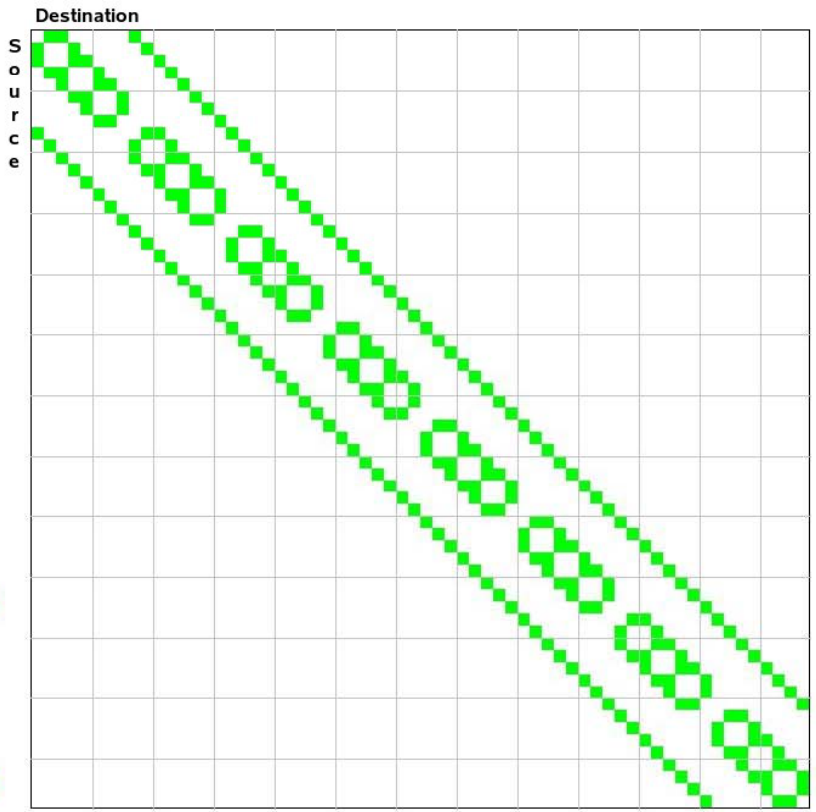

Meso-Scale

Sandia National Laboratories

CTH Communication traces from one timestep on 64 cores

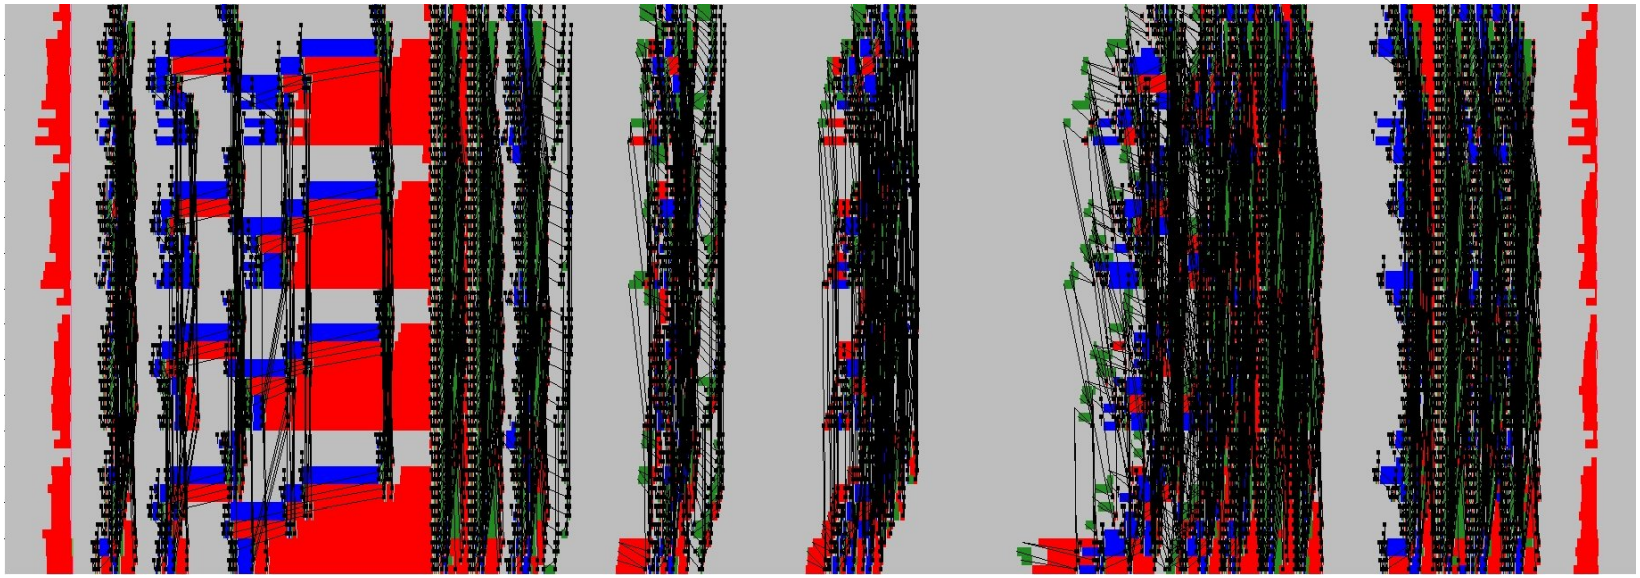

Shaped-Charge

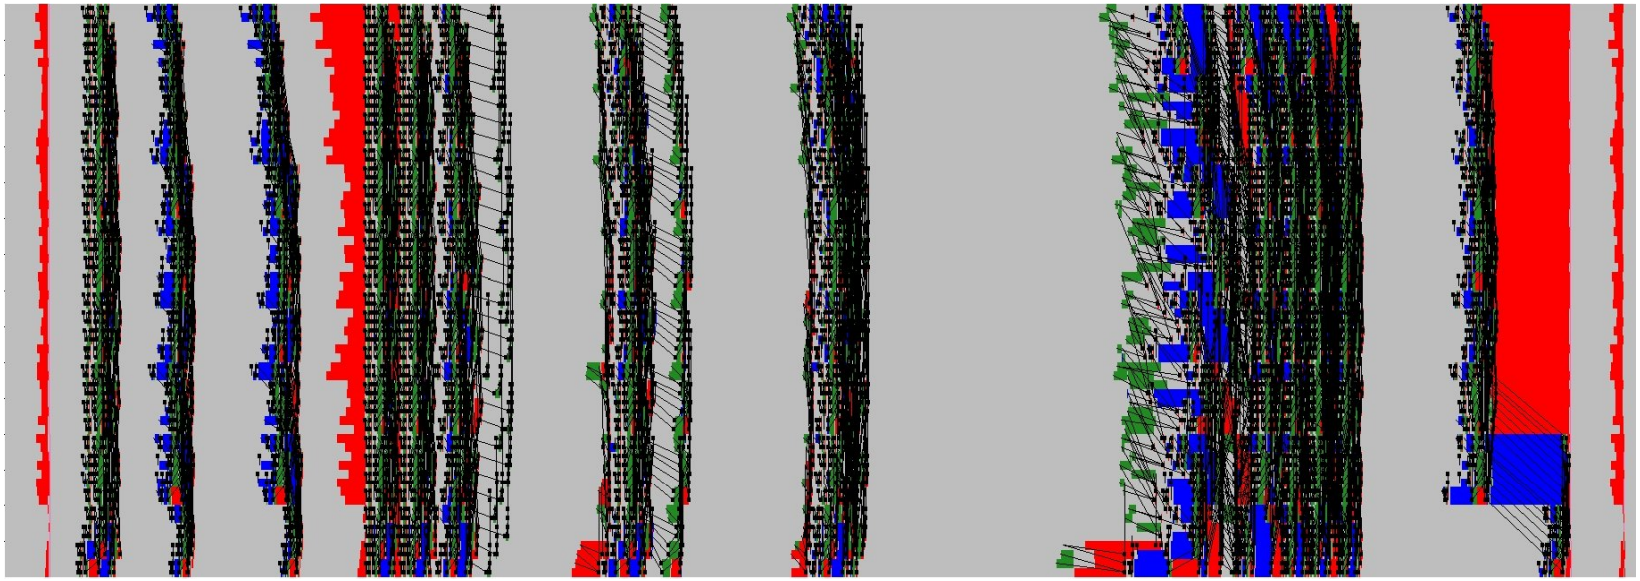

Meso-Scale

PRONTO

- Structural mechanics code with contact algorithm
- Communication for structural mechanics portion consists of boundary exchanges for single variables from static decomposition
- Contact algorithm based on dynamic secondary decomposition which changes during calculation and requires communication from and back to the primary decomposition
- Code is FORTRAN 90 with C for contact communication

Sandia National Laboratories

PRONTO Problems

- Walls problem
  - Two sets of two brick walls colliding
  - Each processor has 320 bricks each of which have 128 elements
  - All communication related to contact
- Can Crush problem
  - Cylinder crushed by block
  - Communication both for finite element and contact algorithms
  - More balanced problem

Sandia National Laboratories

Walls Problem

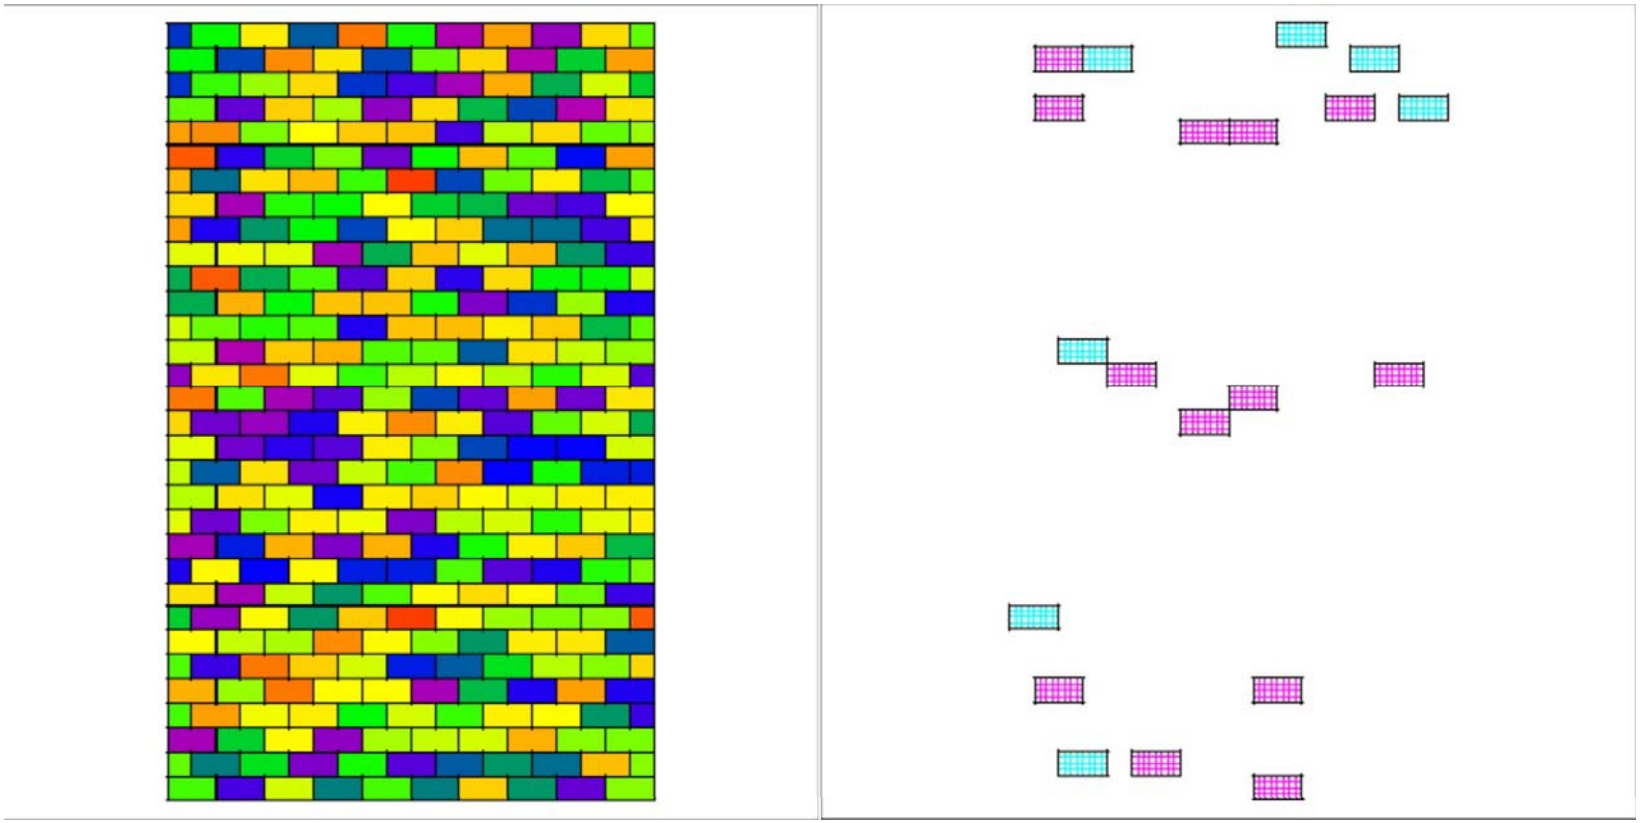

Sandia National Laboratories

Can Crush Problem

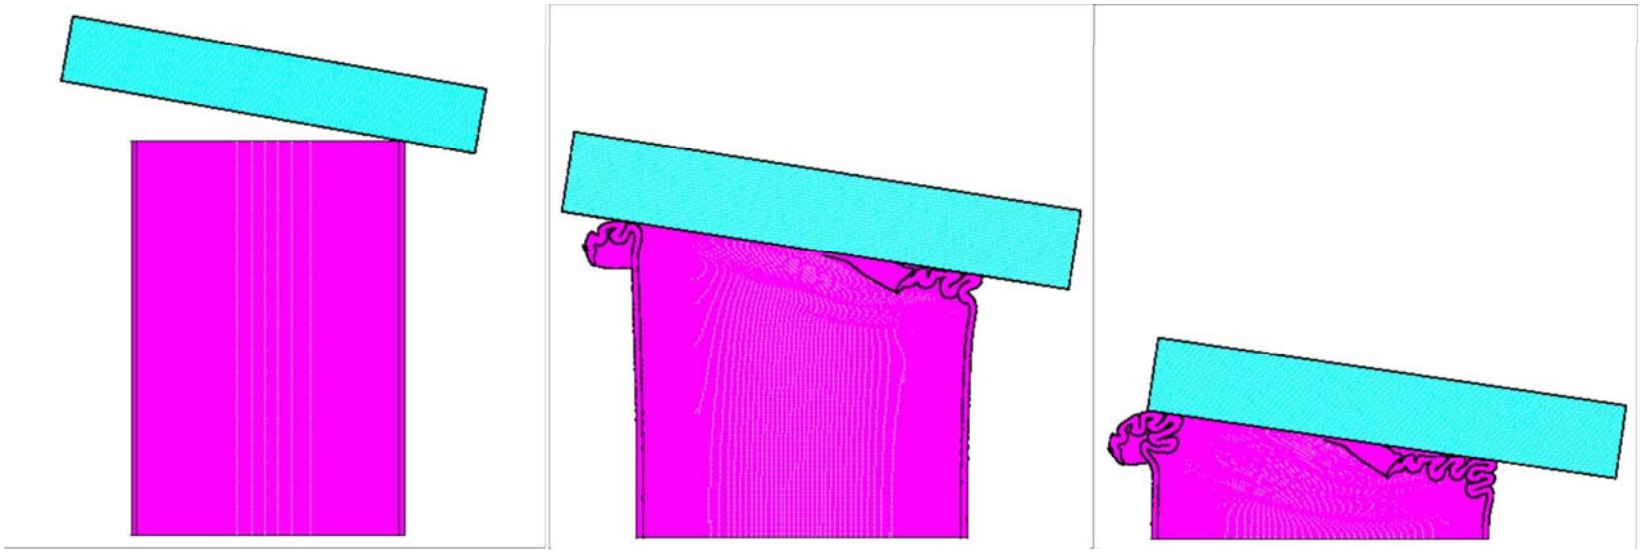

Sandia National Laboratories

PRONTO Communication matrices on 64 cores

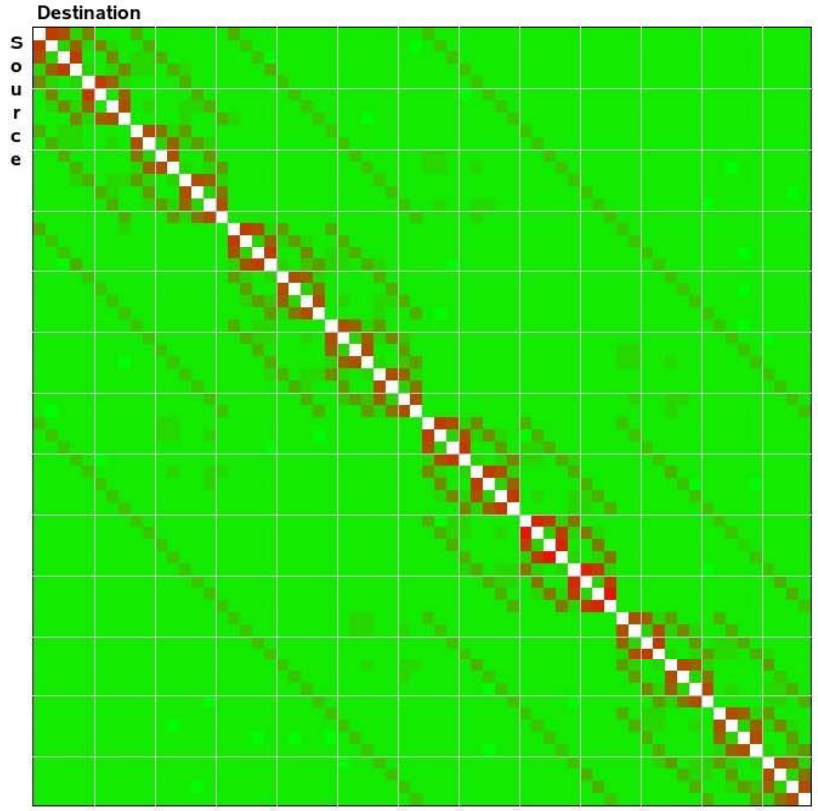

Walls

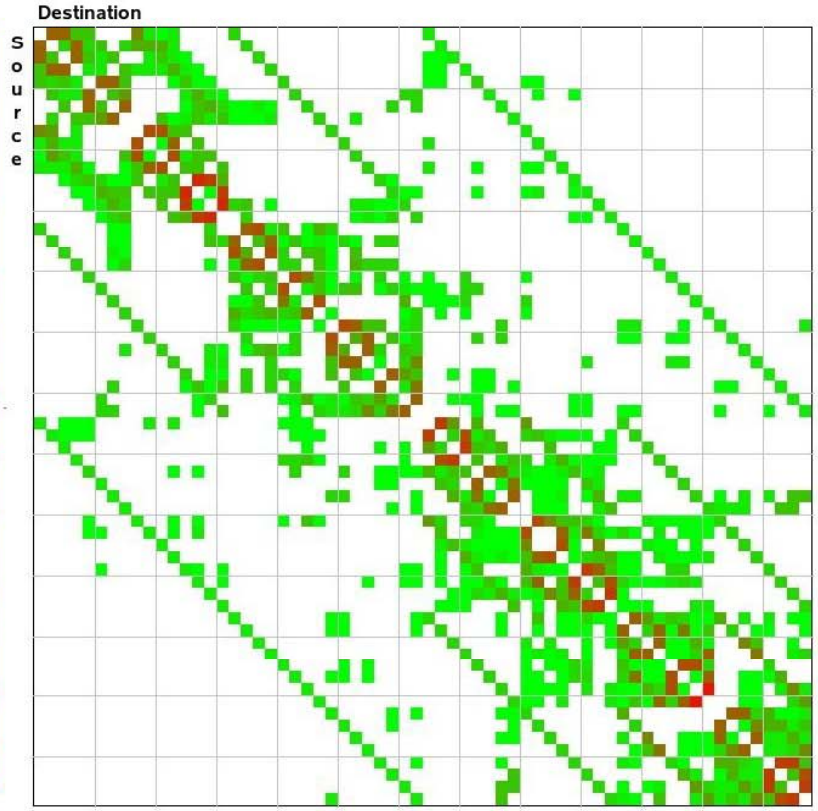

Can Crush

Sandia National Laboratories

PRONTO Communication traces on 64 cores

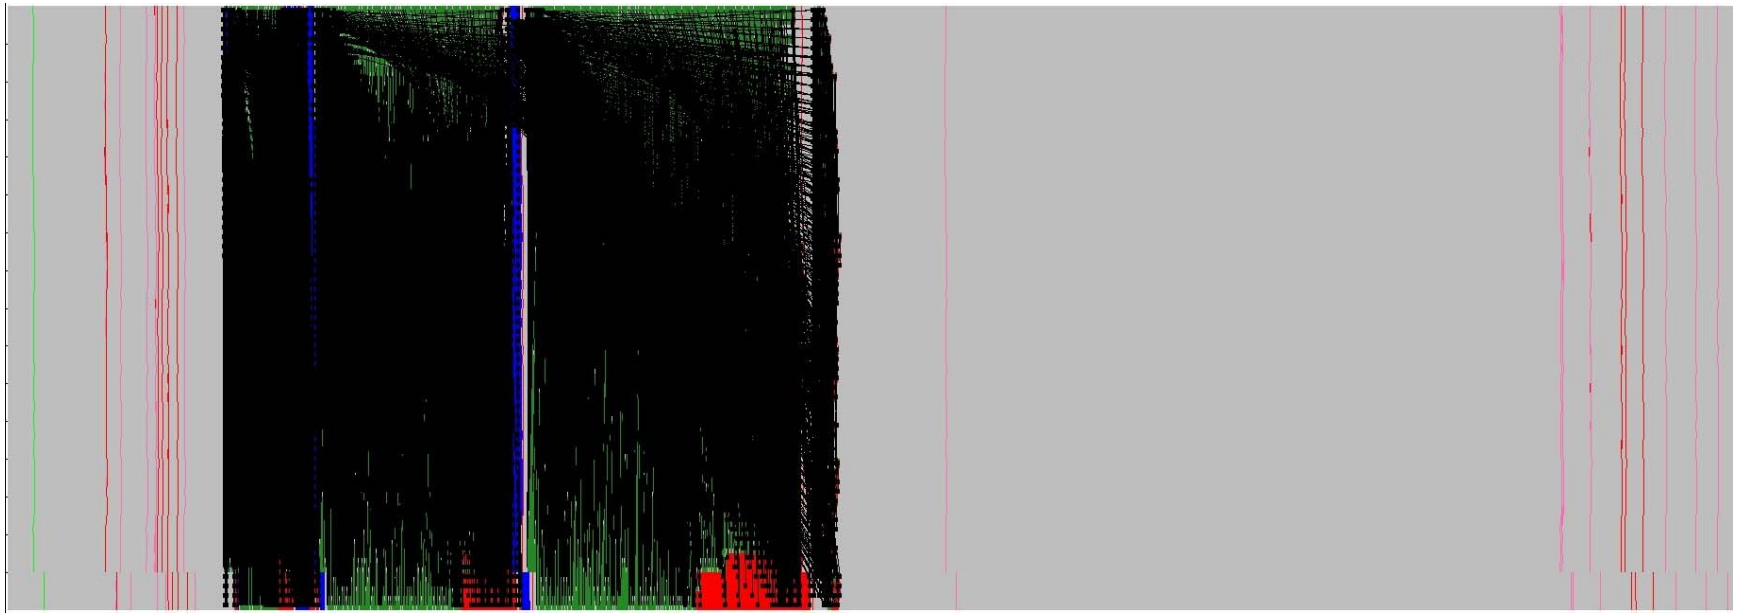

Walls

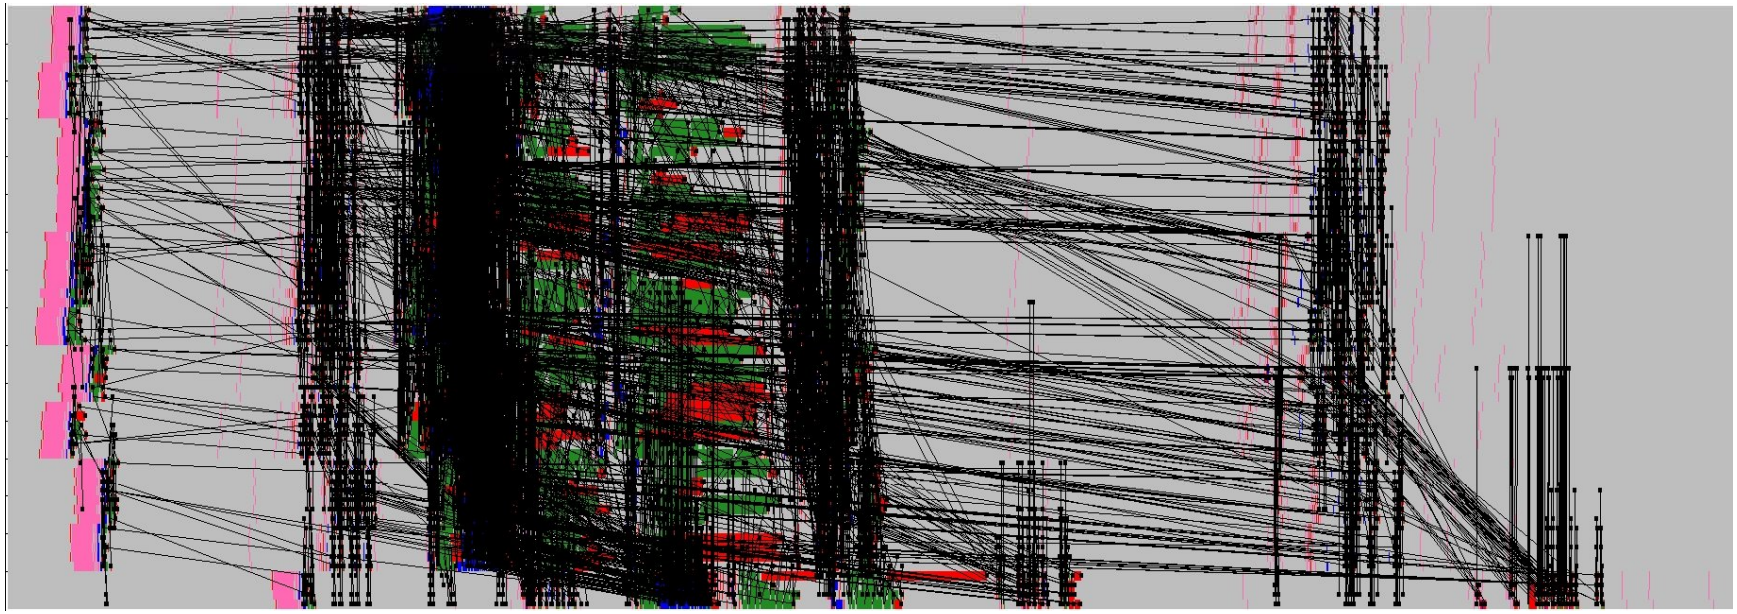

Can Crush

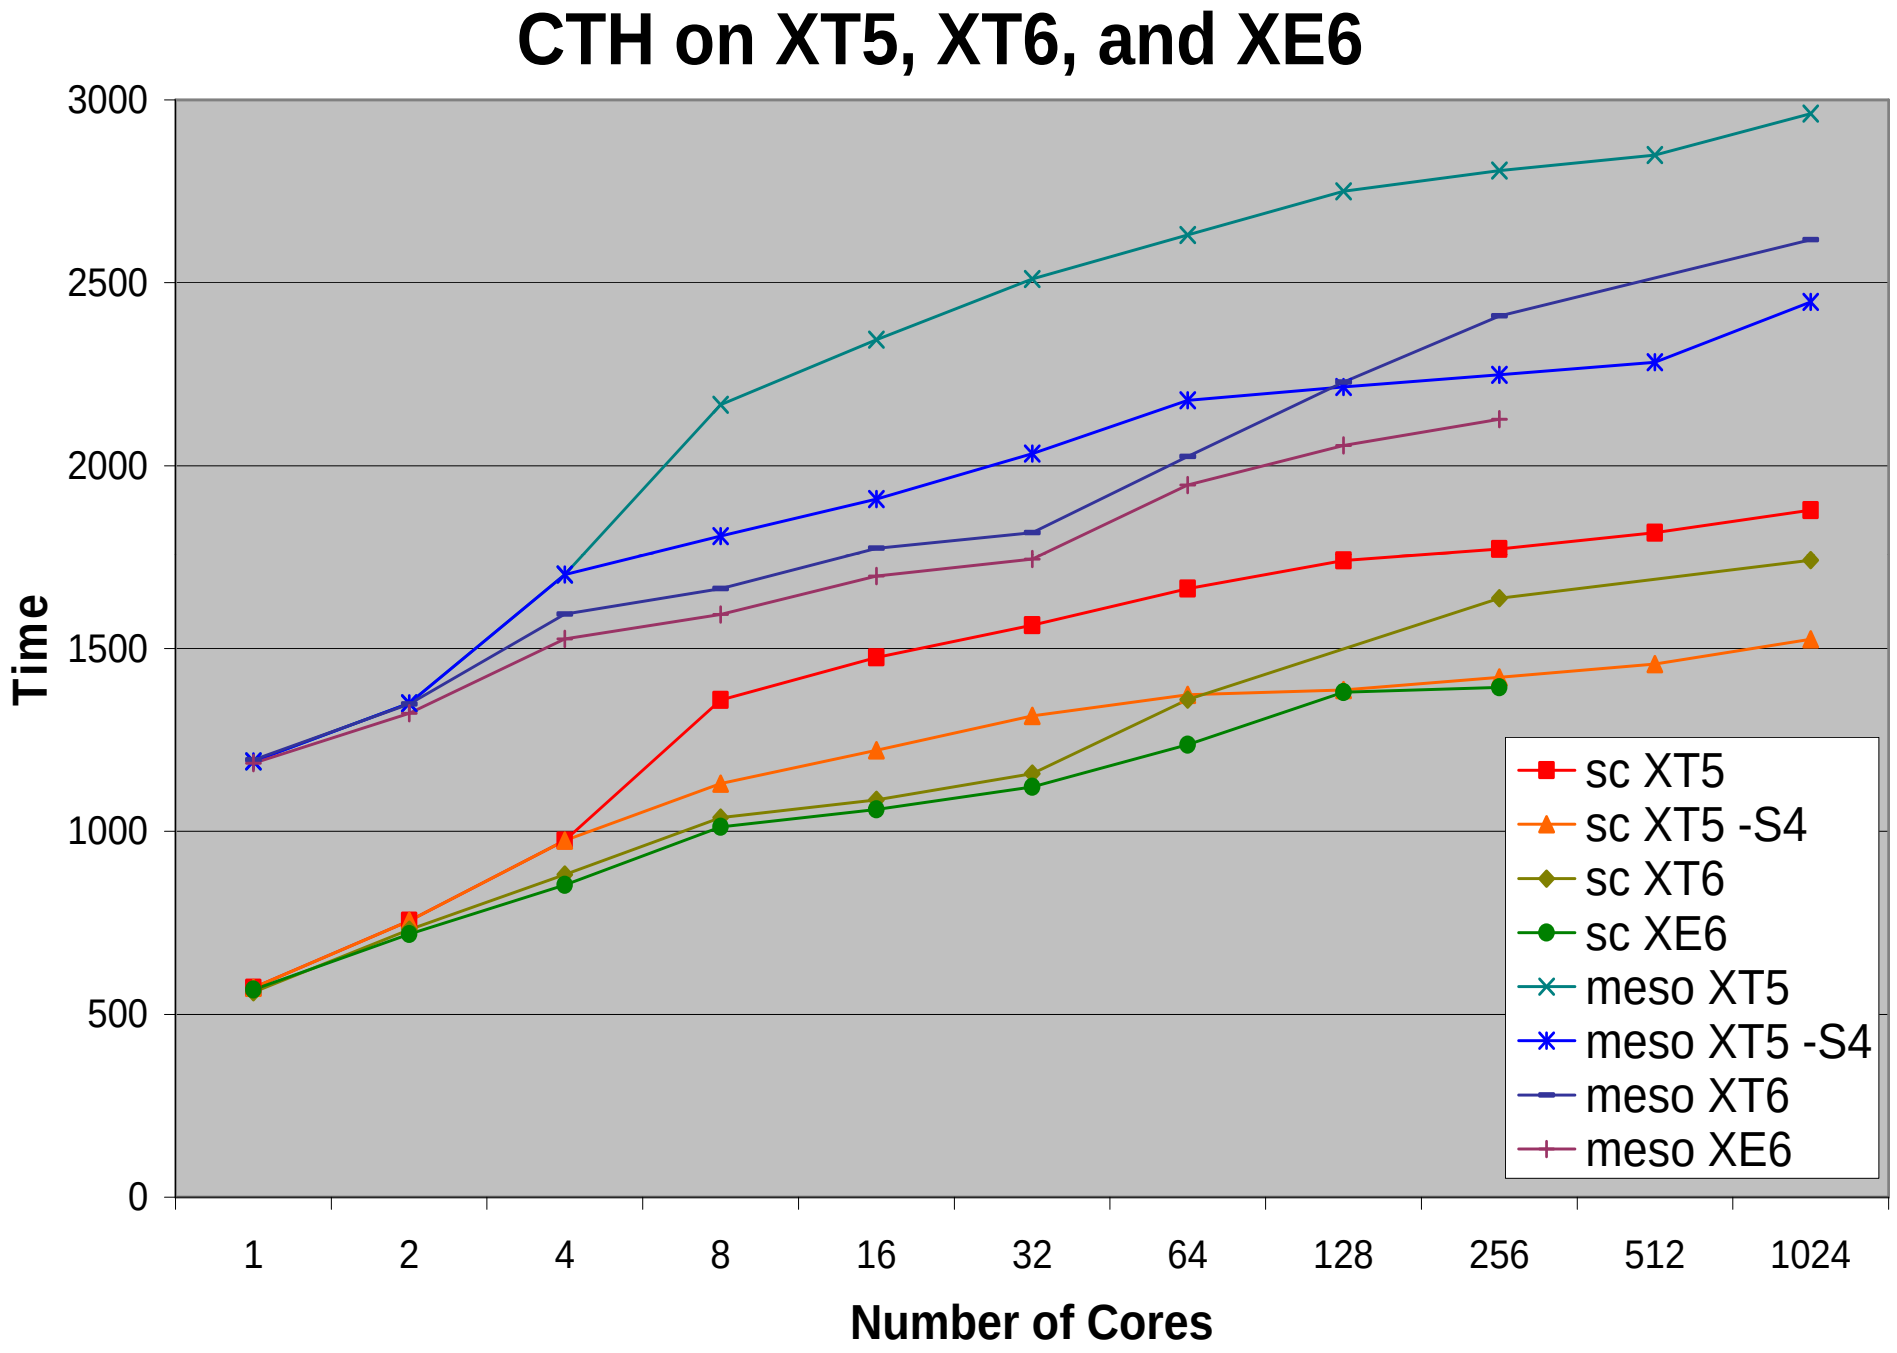
CTH on XT5, XT6, and XE6
Time
3000
2500
2000
1500
1000
500
0
1
2
4
8
16
32
64
128
256
512
1024
Number of Cores
sc XT5
sc XT5 -S4
sc XT6
sc XE6
meso XT5
meso XT5 -S4
meso XT6
meso XE6

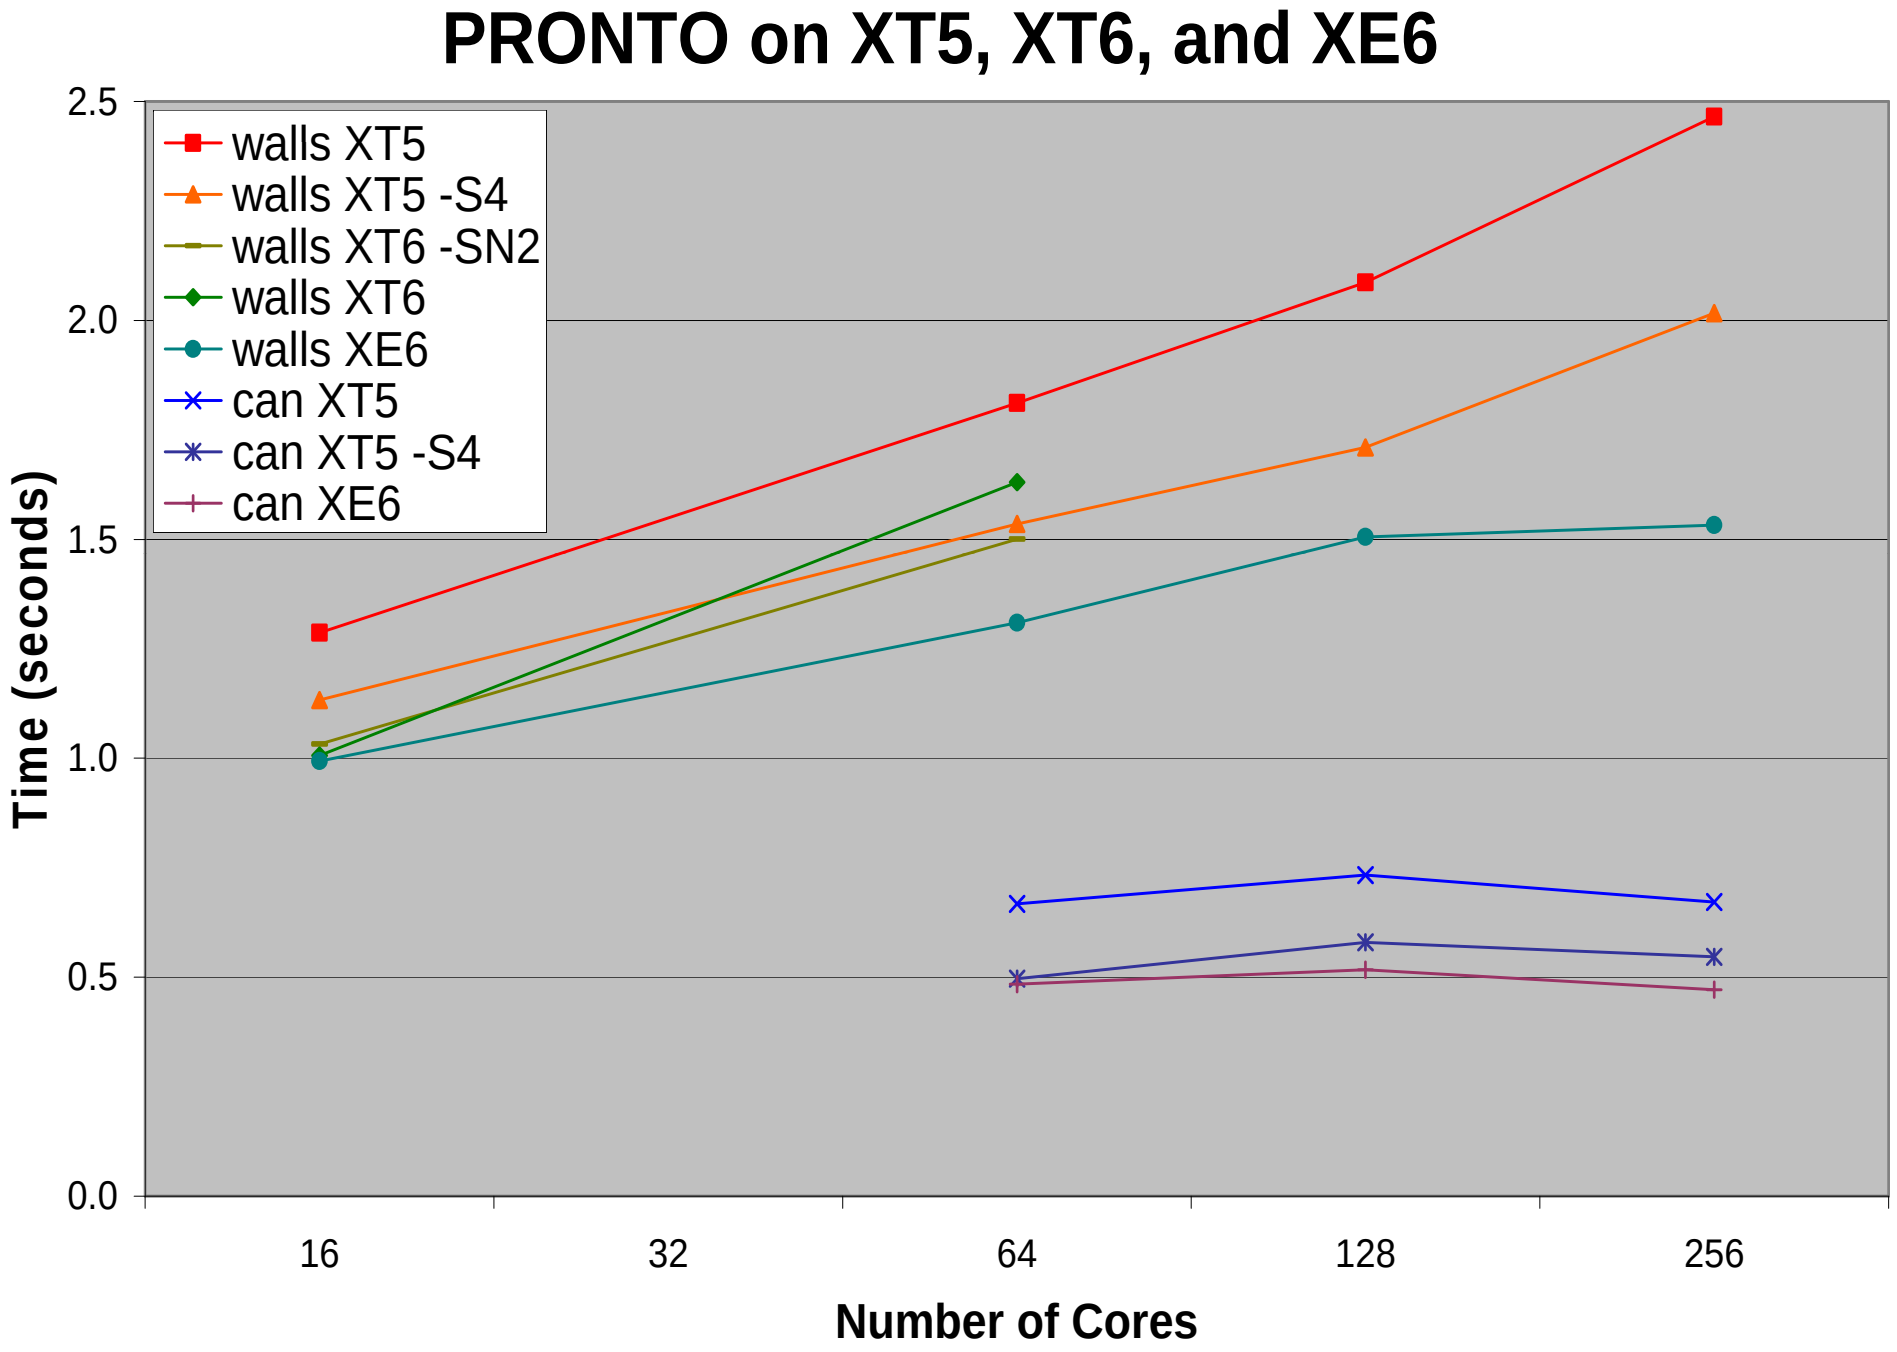
PRONTO on XT5, XT6, and XE6
walls XT5
walls XT5 -S4
walls XT6 -SN2
walls XT6
walls XE6
can XT5
can XT5 -S4
can XE6
Time (seconds)
2.5
2.0
1.5
1.0
0.5
0.0
16
32
64
128
256
Number of Cores

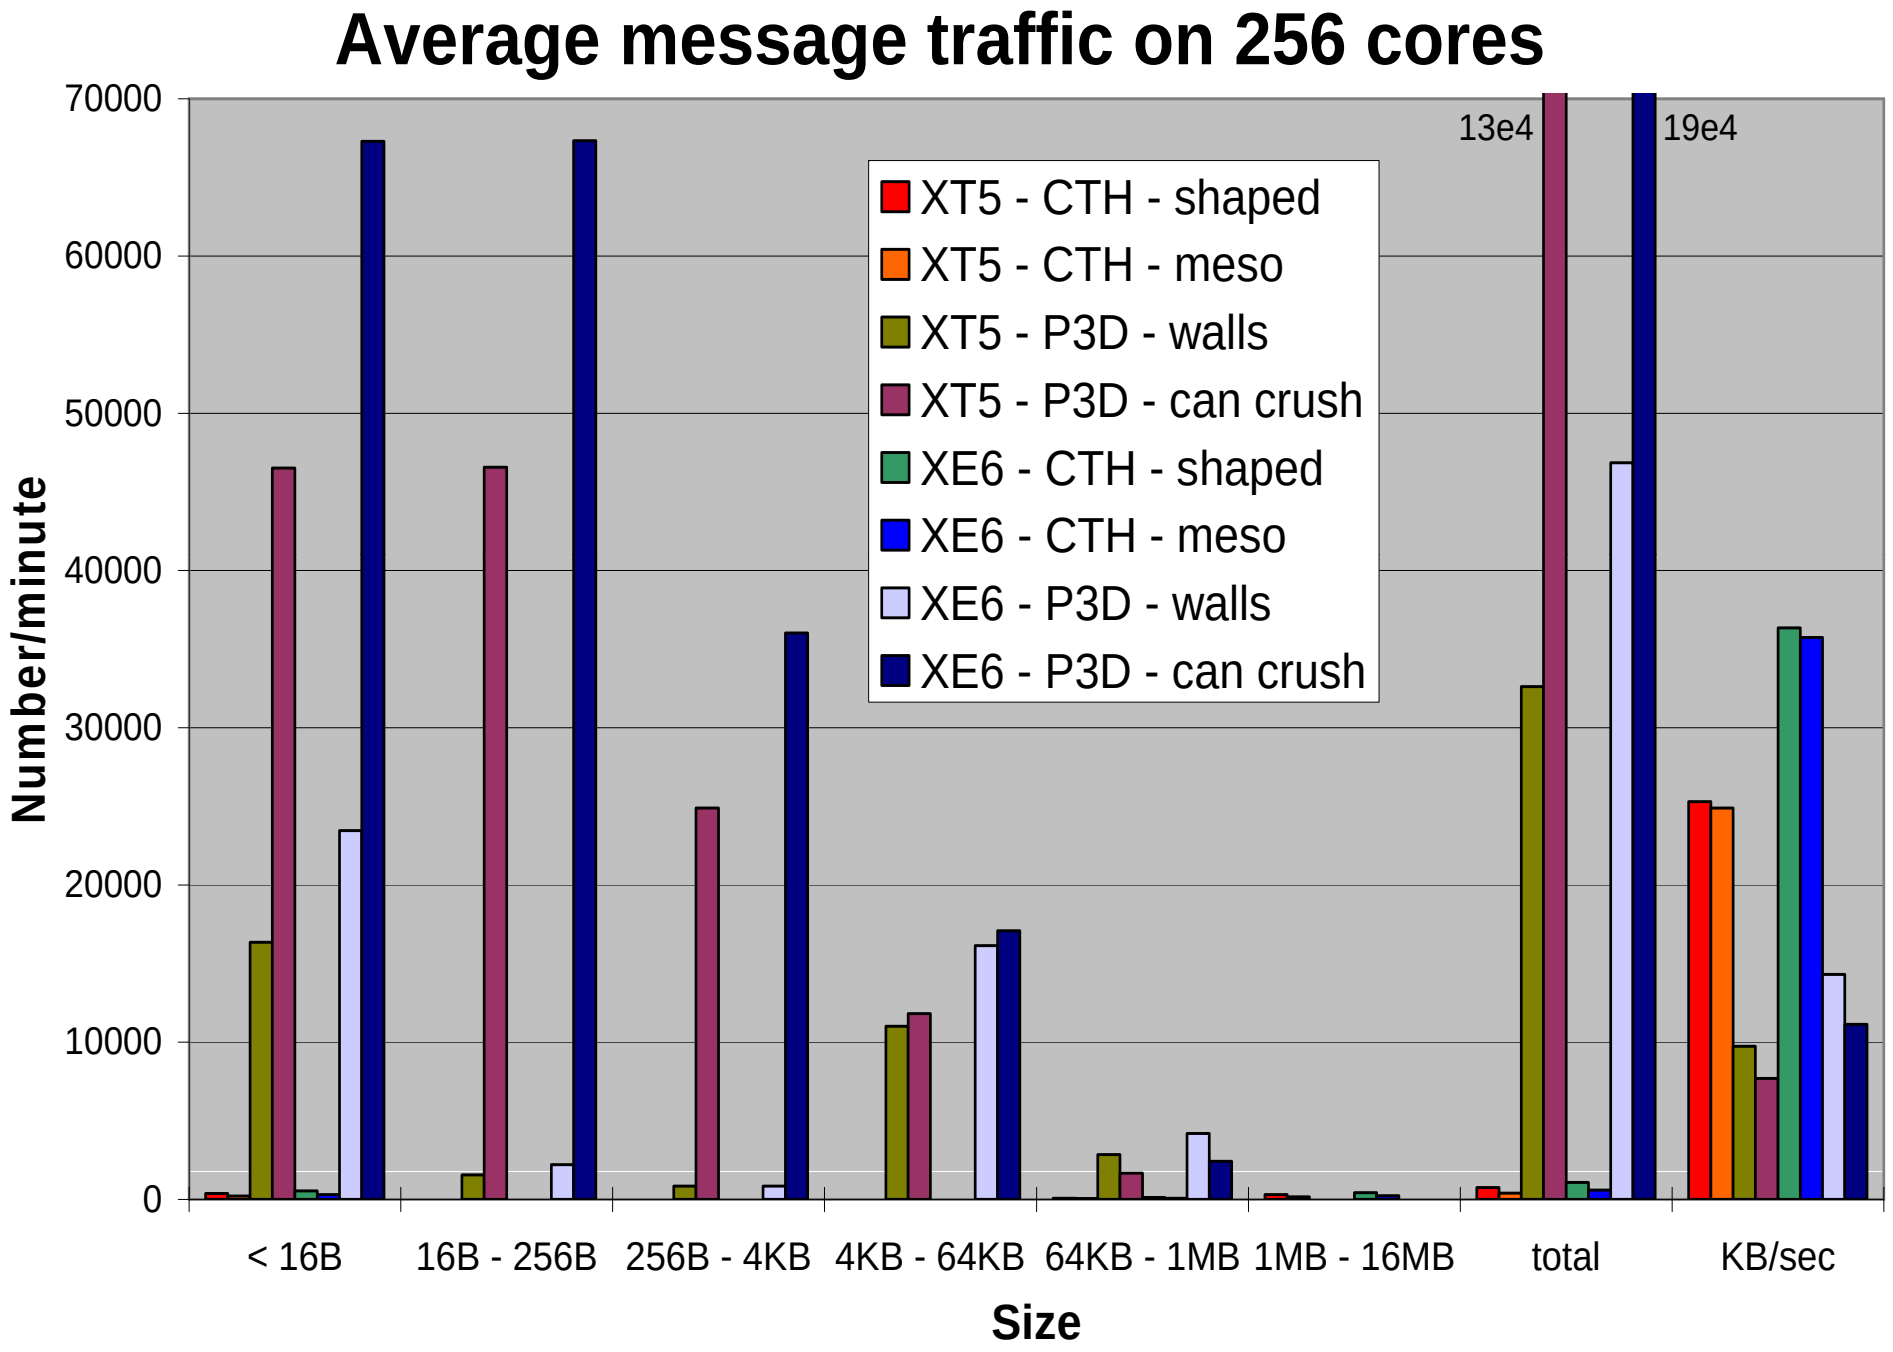
Average message traffic on 256 cores
70000
60000
50000
40000
30000
20000
10000
0
Number/minute
13e4
19e4
XT5 - CTH - shaped
XT5 - CTH - meso
XT5 - P3D - walls
XT5 - P3D - can crush
XE6 - CTH - shaped
XE6 - CTH - meso
XE6 - P3D - walls
XE6 - P3D - can crush
< 16B
16B - 256B
256B - 4KB
4KB - 64KB
64KB - 1MB
1MB - 16MB
total
KB/sec
Size

Summary of Results

- Large portion of performance difference for both codes related to memory contention on XT5 when using 6 cores per NUMA region
- CTH has large network bandwidth requirements and shows some performance improvement moving to the XE6
- PRONTO can send lots of small messages and shows more performance improvement moving to the XE6

Sandia National Laboratories

Future Work

- Extend results to larger number of processors
- Develop mini-app for CTH to see if we can take advantage of the message injection rate of the Gemini interconnect