SAND2012-3476C

Early Results from the ACES Interconnection Network Project

Scott Hemmert
Sandia National Laboratories

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.

Project Overview

- NNSA/ASC asked ACES to consider definition of, and technical oversight for D&E project with Cray, Inc.

- Cray Interconnect Genealogy
  - Generation 1: SeaStar, XT architecture interconnect
  - Generation 2: Gemini, XE architecture interconnect
  - Generation 3: Aries, to be deployed with “Cascade” architecture

- Final SOW signed July 2010
  - Will primarily focus on a potential future, post Aries, interconnection network, tentatively planned for the CY 2015 timeframe

- Early example of focused co-design
  - Efforts centered around impact on ASC applications

High Level Project Tasks

- NIC Studies and Analysis
  - Analyze Gemini interconnect performance
  - Look for improvements which can be enhanced in Pisces
  - Focus on NIC, with emphasis on occupancy, latency, MPI message throughput, and independent progress.

- Router & Network Studies and Analysis – Analyze Aries interconnect performance
  - (byproduct) Optimize Aries network settings for ASC applications
  - Look for possible enhancements to Pisces – Focus on network routing

- Initial architectural specification in collaboration with ACES
  - Perform a comparative study between Pisces and other state of the art interconnects, such as InfiniBand, and possibly one or two others

Early Analysis of Aries Routing using ASC Apps

- Aries implements a Dragonfly topology
  - Form of hierarchical all-to-all topology
  - Overview to follow

- Initial traces for CTH and Sage used
  - Results discussed here are from CTH traces using 1024 MPI ranks

CTH

- **CTH is a multi-material, large deformation, strong shock wave, solid mechanics code developed at Sandia National Laboratories. CTH has models for multi-phase, elastic viscoplastic, porous and explosive materials.**

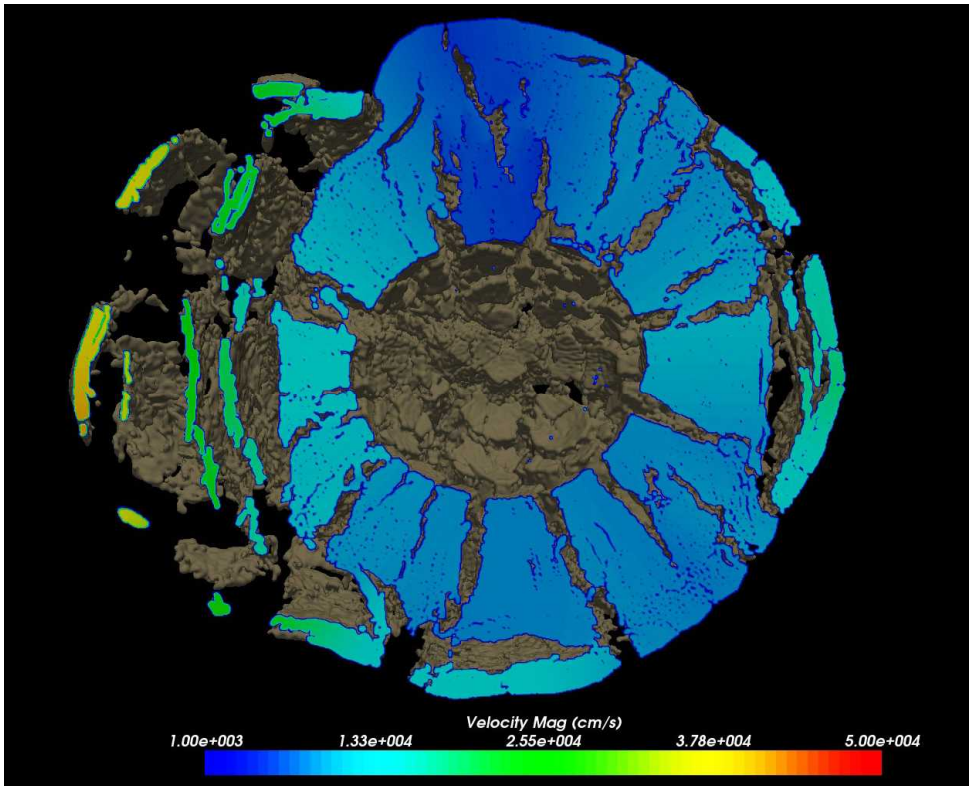

Image courtesy of ASC

Asteroid Golevka measures about 500 x 600 x 700 meters. In this CTH shock physics simulation, a 10 Megaton explosion was initiated at the center of mass. The simulation ran for about 15 hours on 7200 nodes of Red Storm and provided approximately 0.65 second of simulated time.

CTH Shaped Charge Problem

- Explicit 3d multi-material shock hydrodynamics
- Models high-speed hydrodynamic flow and the dynamic deformation of solid materials, and includes several equations of state and material strength models
- Solves the equations of mass, momentum, and energy in an Eulerian finite volume formulation
- Shaped charge problem: 40 state arrays, static mesh.

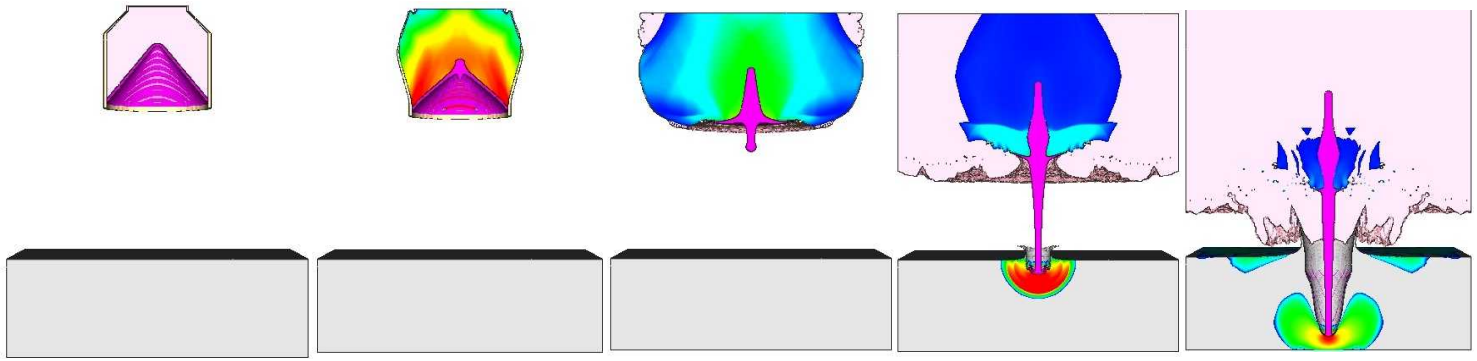

CTH Communication Characteristics

Shaped Charge Problem (weak scaling)

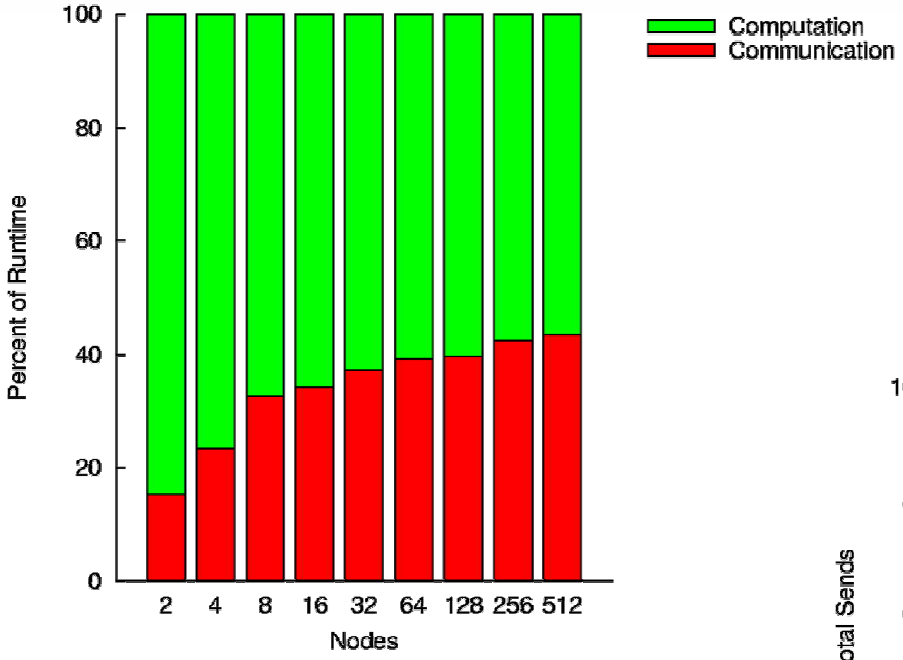

As job size increases, communication time can grow to consume around 40-50% of the runtime.

CTH communication is dominated by long messages. Fix graph ----→

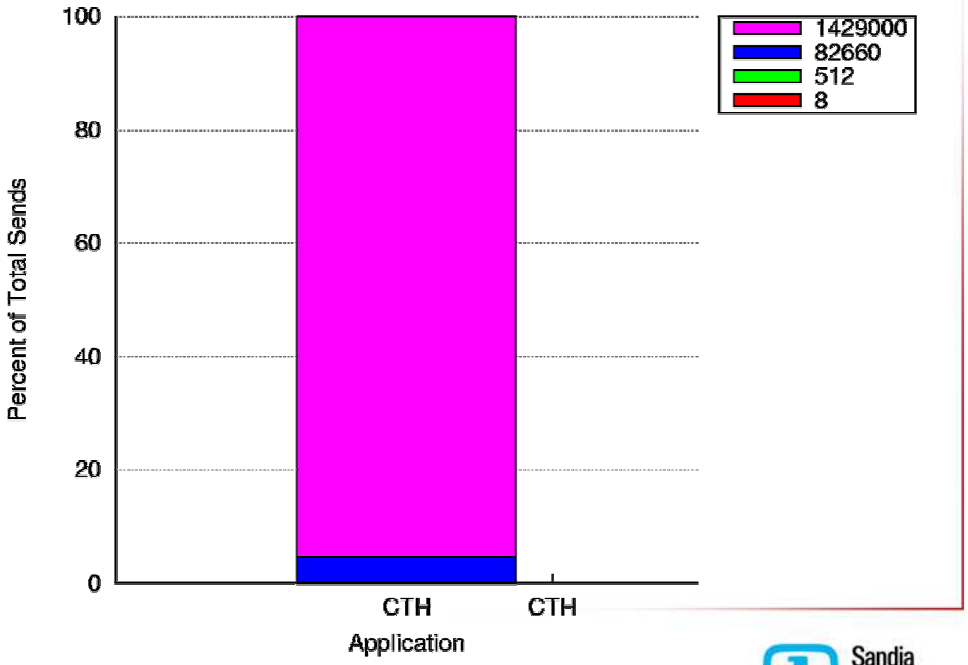

CTH Bandwidth Degradation Study

- Uses capabilities built into the Red Storm SeaStar interconnect to turn off interconnect router lanes at boot time
  - Links are made up of 4 3-bit subchannels that can be independently enabled
- Measure application performance at full and one-quarter link bandwidth
- At largest measured job size, quartering bandwidth leads to 32% longer runtime