LA-UR- 04-0445

Approved for public release; distribution is unlimited.

Title: An ASCI User Perspective

Author(s): John M. Scott

Submitted to: ASCI S&CS Requirements and Planning Meeting
January 27-29, 2004
Santa Rosa, CA

An ASCI User Perspective

John M. Scott

Technical Staff Member
Los Alamos National Laboratory
X-2, Thermonuclear Applications
jmscott@lanl.gov

Abstract

- The ASCI project supports the hardware and software projects that ultimately are put to use by the nuclear weapon design community in order to solve problems relevant to the stockpile. This talk describes a group of ASCI code project users and an example problem from one of those users in terms of its computing resources. Significant discussion focuses on important considerations as a user when interacting with ASCI hardware and software.

Outline

- Crestone user background
  - Who are they? What do they do?
- Example parameters from Crestone simulation of a stockpile issue
- As a user, what do I think are important considerations.

Crestone Users

- The bulk of Crestone users reside in X-2, a group of about 70 people.
  - Approximately 50% to 75% have used Crestone project codes.
  - Applications vary from weapon or NTS simulations, radiation flow, shock tube, inertial confinement fusion, astrophysics, asteroid impact, etc.
  - Types of problems
    - 1-,2-,3-D problems
    - complex geometries
    - 10^2 to ~10^9 cells
    - rich in physics
- Crestone users are currently the bulk of the usage for LANL secure machines.

Crestone Project codes are applied to relevant stockpile problems

- Users have been quick to adopt Crestone project codes and apply them to real problems involving the stockpile.
  - A majority of usage has been on specific NTS events for validation .
    - V&V milestone and weapon system baselining.
  - There has been application to resolution of significant finding investigations.
- Strong, positive interaction between developers and users has led to project successes.

Parameters from a recent weapons application

- This problem was inherently a 3-D problem, but was distilled into 2-D.
  - 1-D simulations were used to assess differences between legacy and ASCI code calculations.

- Over the course of about 8 months...
  - ~100 1-D runs, hour to overnight runtime, 100s to 1000s of cells, 100s MB information per simulation
  - ~100 2-D runs, 1 wk to 1 month runtime, $5x10^5$ to $4x10^6$ cells using 48 to 120 processors, 10s GB information per simulation
  - Majority of runs on Compaq Q and C platforms
    - Early 1-D runs on Bluemountain
    - Started 2-D runs on QB and migrated to CA, CB, and CC
  - Some visualization with Ensight
  - Analysis of HDF and text data files performed with Igor.

What is important for the user?

- Throughput--the amount of time it takes to solve a problem of interest
  - The ability to run a large, complex problem without significant human intervention is important.
    - Mean time to hardware failure; more hardware means more things to go wrong. This is more of a concern as larger jobs become the norm.
  - Analysis of 3-D simulation output is difficult and requires a team of people. There have been impressive results.
  - Setup can be easy to very difficult depending on the nature of the code.

What is important for the user?

- Speed--the amount of time of wall clock time spent by a machine on a simulation.
- Issues that the design community are being asked to address are becoming increasingly complex.
  - We are being asked and are calculating scenarios that could not have been simulated in the past.
  - Many of these scenarios can really only be addressed in a 3-D environment.
  - The surveillance program is identifying issues that have not been seen before.
- ASCI codes are expected to contain better first principles physics models.
- These are two of many reasons that drive the need for larger and more complex simulations that require faster machines.

What is important for the user?

- Capacity--how many and what size jobs can be run at the same time?
  - Generally, it is not the size of a single job but the number of jobs of a specific size that can be completed in a given amount of time that is important.
  - Parameter studies are important in order to understand the sensitivity of the results to physics models that are not always adequate. This requires a lot of simulations that are very similar in nature.

What is important for the user?

- Ease of use--How much does the user need to understand about the system hardware, system software, and the application in order to accomplish his/her task?
  - the many flavors of UNIX (shell environment and command lexicon)
  - the different file system structures and scratch file systems
  - bit/byte ordering, IEEE vs. Cray
  - hpss (psi and xpsi)
  - scratch purging

Concerns as a user...

- Is the end-user experience important?
  - Can more machine debugging occur before machine is open for general use?
    - How can this process be accelerated by users?
  - Uptime and reliability
    - If a user goes away for two weeks, will a job go to completion without any help?
  - What can be done to make the interface more consistent across platforms?
    - issues from previous slide
  - Simple, understandable documentation
    - Help the user understand the system

Summary

- A relatively small issue in comparison to other stockpile issues required:
  - 30-50% of a user's time over 8 mos.
  - ~ 10^6 hrs of single processor time
  - ~ TB of storage
- As a user, one wants to get to the right or wrong answer as quickly as possible.
- Speed, capacity, and ease of use all factor into throughput.

I am one in a sea of many.

- I am just one user from a group of 70.
- That group is only the secondary design group.

- My opinion is not meant to represent a consensus or majority view of ASCI code project users.