HPC Application Performance Analysis and Prediction

LDRD
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

Sandia National Laboratories

Michael A. Heroux (PI), S. Scott Collis (PM), H. Carter Edwards, Alan Williams

Problems

- The arrival of multi-core computer processors is a 'step-change' in computer technology, which requires different programming approaches.
- We need improved decision-making for selecting next-generation computer systems.
- How will we develop parallel application codes that effectively utilize multi-core processors?
  - Will "pure MPI" be good enough? (1 MPI process per core)
  - Will MPI implementations do "the right thing" for multi-core?
  - Will explicit hybrid programming (MPI+threads) be necessary?

Approach

Mantevo* Project

* Greek: Augur, Guess, Predict, Presage

- Multi-faceted project to develop benchmarks and simulators for studying application performance.
- Three types of packages:
  - Microapps: Small, self-contained programs.
    - HPCCG: unstructured implicit FEM/FVM.
    - pHPCCG: parameterized scalar/int, SpMV kernel.
    - phdMesh: explicit FEM, contact detection.
    - MD: Parameterized from simple to bio molecules.
  - Microdrivers: Wrappers around Trilinos packages.
    - Beam: Intrepid+FEI+Trilinos solvers.
    - Epetra Benchmark Tests: Core Epetra kernels.
  - Motif framework: Collection of "dwarves."
    - Prolego: Composable fragment collection to mimic real apps.
- Open Source: Fosters external collaboration.
- Staffing: Application & Library developers.

Node Architectures: Key Focus Area

- **Multicore:**
  - New HPC systems axis.
  - First Mantevo analysis focus.
- **Quantitative results:**
  - Confirm, sharpen intuitive sense.
  - Sometimes counter intuition.

Intel Clovertown

AMD Barcelona

Sun Niagara2

Multicore: First focus area for Mantevo micro apps/drivers.

phdMesh

- Portable compact application with parallel heterogeneous dynamic unstructured mesh.
- Geometric proximity search is a performance constraining kernel for contact detection and multiphysics loose-coupling.
- Driver problem – six 3D counter rotating "gears" with continually changing contact surfaces.

Beam: 3-D FEM "beam" of Hex-8 elements

- Mimics important computational characteristicsof implicit finite-element applications.
- Heavily exercises Trilinos (linear solvers andinfrastructure) and FEI (sparse matrix assembly).
- Shown scaling to 2 Billion equations on Red Storm
- Beam test program is portable, parallel, and freely available

Prolego: compact performance-prediction application

- Collection of "kernels" that exhibit the performance characteristics of a real application kernel.
- Application performance can be modeled by a calibrated collection of "basis kernels."
- Prolego infrastructure has been developed:
  - XML input file parsing (run-time selection and specification of kernels)
  - Some kernels are in place:
    - BLAS operations (vector axpy, dot operations, matrix-vector, matrix-matrix)
    - Sparse matrix-vector multiply
    - MPI communication operations
- Near-term future work:
  - add more kernels representative of Sandia target applications
  - collect and compare predicted performance data to refine the concept of generating performance bases.

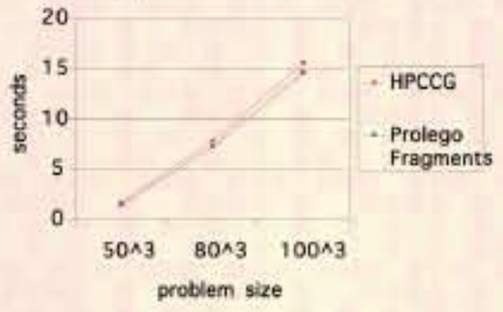

Library Efforts in Response to Node Architecture Trends

- Block Krylov Methods (Belos & Anasazi)
- Specialized sparse matrix data structures
- Templated Kernel Libraries (Tpetra & Tifpack)
- Shared memory node-only algorithms
- Kokkos Node class: Support for Intel TBB, OpenMP, Pthreads, ...
- MPI-only+MPI/PNAS

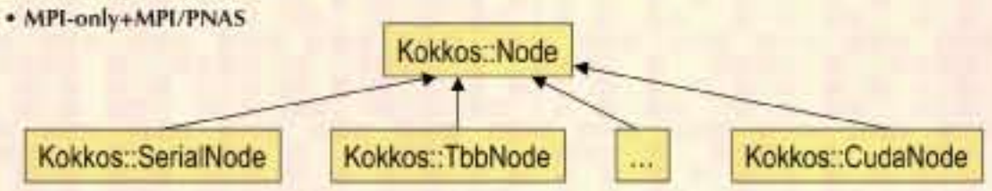

- Trilinos/Kokkos: Trilinos compute node package.
- Abstraction definition in progress: Will look a lot like TBB.

Results

A Few Programming Model Results

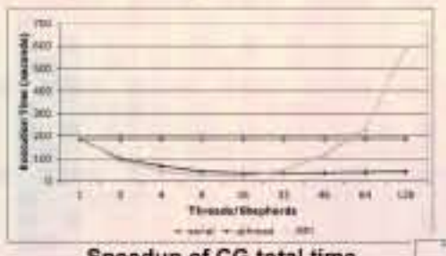

- Model HPCCG rewritten:
  - Qthreads: Massively threaded library.
  - BEC: Bundle-Exchange-Compute.
- MPI & MPI+threads.
  - App: MPI-only
  - Solver: MPI+threads

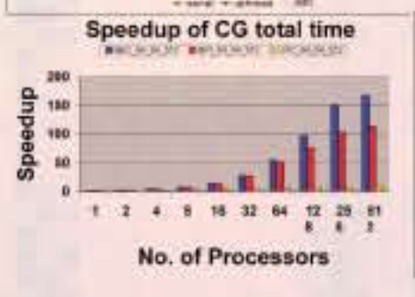

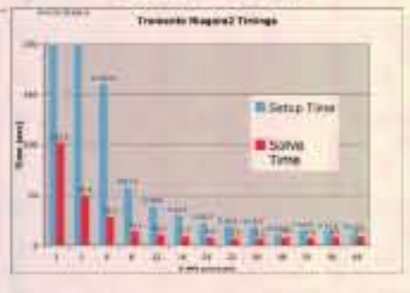

Significance of Accomplishments

- Micro-drivers and Compact Apps:
  - Provide new information source for system designers.
  - Provide new test-bed for experimental programming models.
  - Quantify the performance advantages of different programming models.
  - Establish dialogue between systems and applications staff.
- External visibility:
  - Enhances Sandia's voice in the external community.
  - Provides valuable insight from external experts.

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.

SAND 2008-5464C CreativeArts 98429.LW.8.08

NNSA National Nuclear Security Administration

Sandia National Laboratories