

Massive Multithreading Applied to National Infrastructure and Informatics



Sandia National Laboratories
 PI: Jonathan Berry, PM: Heidi Ammerlahn,

Problem

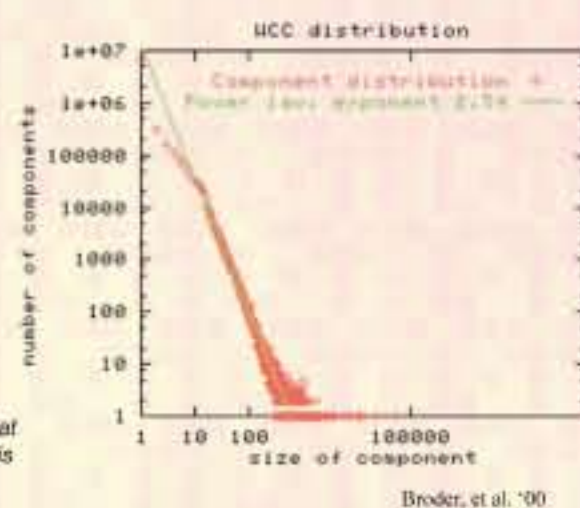
Informatics applications are memory bound. Datasets are huge. Fast CPUs are prohibitively wasteful of power. Can we use special data-driven architectures to address national informatics needs?

Informatics Datasets Are Different

Informatics: The analysis of datasets arising from "information" sources such as the WWW (not physical simulation)

Motivating Applications:

- Homeland security
- Computer security (DOE emphasis)
- Biological networks, etc.



From UCSD '08
 "One of the interesting ramifications of the fact that the PageRank calculation converges rapidly is that the web is an expander-like graph"
 Page, Brin, Motwani, Winograd, 1999

Primary HPC Implication: Any partitioning is "bad"

Informatics Problems Demand New Architectures

Distributed Memory Architectures	Massively Multithreaded Architectures	Key Issues
Fast CPU (~3GHz)	Slow CPU (~200-500MHz)	Power, concurrency
Elaborate memory hierarchy	Almost no memory hierarchy	Is cache justified?
Memory per-processor, partitioned	Global address space	Can you partition?
Operating system for threading, synchronization	Hardware for threading, synchronization	How fine-grained is your data interaction?
Programming paradigm is standardized (MPI)	Programming paradigm is machine-specific (mta-pe)	Portability, debuggability


Multithreaded architectures show promise for informatics problems, but more work is necessary...

Approach

Discover and adapt algorithms for massively multithreaded supercomputers, develop a generic software framework for these algorithms, and evaluate them on real data.

We Are Developing the MultiThreaded Graph Library

- Enable multithreaded graph algorithms
- Build upon community standard (Boost Graph Library)
- Adapters abstract data structures and other application specifics
- Hide some shared memory issues
- Preserve good multithreaded performance



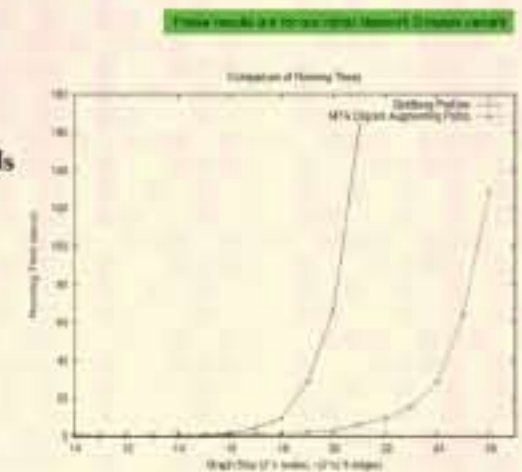
Business model: Open-source code provides a mechanism for Sandia-developed algorithms to be applied to proprietary or sensitive data structures used by customers

Results

- Scalable community detection
- Scalable detection of short cycles
- Scalable maximum flow computation
- The MultiThreaded Graph Library (MTGL)


Maximum Flow: Multithreaded Network Simplex Algorithms

- Flows through networks are a NISAC interest
- Typical runtime $O(nm)$
- State of the art: Andrew Goldberg's serial methods
- Initial results: we leverage the Cray MTA-2 to extend solvable instance size by 32X




Community Detection: The "Edge Support"

- De Facto state of the art in scalable community detection is the greedy "CNM" heuristic $O(n \log^2 n)$
- Top image shows its solution to the familiar "Zachary's karate club" [Reality: split in two]
- Our facility location based approach suggested a fractional "support" (bottom)
- Darker edges have more "support;" we can now sample an ensemble of solutions



Scalable Community Detection

- Community detection can be modeled as an "Uncapacitated Facility Location Problem" (UFL)
- The UFL has special structure that permits solution in linear time and space
- We have made the kernel of the UFL threadsafe, and demonstrated scalability out to 2560 threads (20 processors) on the Cray MTA-2
- We have a new, untested formulation that is expected to scale farther



Significance

The reputation of the MTGL has led to technical recognition (keynote at 2008 IEEE "MTAAP" workshop) and programmatic impact (new multi-FTE WFO project)

- Copyright/Berkeley open-source license paper work submitted
- Transition to SVN completed
- Nightly testing mechanism introduced
- Preparing for 1.0 release
- Synergies
 - Networks Grand Challenge uses MTGL
 - Large WFO project uses MTGL
 - Proposed "X-Caliber" architecture motivated in part by MTGL apps
 - A few MTGL algorithms are integrated with "qthreads" for SMP/CMP

Contributing Staff:
 • Jonathan Berry (01416)
 • Bruce Hendrickson (01415)
 • Randall LaViolette (05634)
 • Vilas Leung (01415)
 • Greg Mackey (06321)
 • Brad Mancke (09511)