

Multi-Core + Multi-Tasking = Multi-Opportunity?

Karl A. Nyberg

Overview

- Background
- Motivation
- Industry Trends
- Technical Activity
- Market Opportunity
- Questions and Answers

Background

- **Large Integers**
 - Personal Fascination
 - Perfect numbers
 - Mersenne primes
 - Fibonacci sequences
 - Implementations
 - BASIC, FORTRAN, Pascal, PL/I, Gypsy, C, Ada
- **Computer Security**
 - DES Algorithm
 - RSA Factoring

Initial Motivation

- **RSA Factoring Challenge (now withdrawn)**
 - Determine prime factors of large integers
 - Intended to show difficulty of problem, security of encryption approach and encourage research
 - Receive money \$20K - \$200K (and fame)
- **Sun Open Performance Contest (expired)**
 - Part of their “Try-and-Buy” Program to get potential customers to test multi-core system
 - Intended to have real-world evaluations
 - Receive test system (and fame, again)

One out of two is not bad! 😊

- **Sun Open Performance Contest Winner**
 - 60 day evaluation using “A Constructive Approach to Integer Factorization”
 - Available at <http://www.grebyn.com/t1000/>
 - Compared against Pentium IV & Athlon64 systems
 - Done in Ada
- **Awarded Evaluation System**
 - 8 1.0GHz cores x 4 “strands”, 16 GB memory, 2 x 73GB 10K RPM SAS drives, 4 Gigabit Ethernets
 - List price \$14,445

Future Motivation???

- **Netflix Prize (not yet expired or won 😊)**
 - Predict Movie Recommendations
 - Intended to encourage research into improvement of recommendation system
 - Receive \$1M (and that fame thing, again)
- **Eternity II (not yet expired or won either 😊)**
 - 256 piece jigsaw puzzle
 - Intended to make the inventor wealthy
 - Receive \$2M (fame, glory, notwithstanding...)

The Machine

- 8 x 1 GHz UltraSparc V9 CPU cores (Niagara 1), 4 way hardware “strands”
- 16 GB memory
- 4 x Gigabit Ethernets
- 2 x 73 GB SAS 10000 RPM disk drives
- Solaris 10 (Linux available)
- 1U Physical size
- ***Draws approximately 142 watts***
- **LOUD!** And heavy

Heart of the New Machine



Back To The Present...

- What to do with a ~\$15K 32-way computer?
 - Additional multi-tasking “research”
 - Deploy Java application
 - EBay! ☹

Research Approach

- Create Multi-Tasking Test Applications
- Execute
- Analyze Results

Test Applications

- Simple UNIX Utilities Implemented
 - wc – Word Count
 - sum – Check Sum Calculation
- Possibly More Complicated Applications Considered
 - Astronomy
 - File Compression
 - ???

Results

- WC
 - Elapsed time decreased with increased tasks
 - Not quite linear in relative improvement
 - Limited by file I/O
- sum
 - Elapsed time again decreased
 - More linear, as more CPU, less I/O
 - Tasking overhead seems to have limited total throughput to about 50%

Industry Trends in Performance

- Serial Approach
 - Programming Languages
 - Applications
 - Programmers
- Concurrent Approach
 - Programming Languages (and Libraries)
 - Applications
 - Programmers

Industry Environment

- Hardware Multi-Core Paths
 - Intel, AMD, IBM, Tileria also creating multi-core
- Software (System / Language) Support
 - OpenMP, MPI
 - Parallelizing C, C++, FORTRAN
 - Java?
- Industry expects YEARS to leverage multi-core
- University starting to fund training parallelism
- Ada is ready **TODAY!**

Opportunities

- Embedded Systems
 - Communications gear (routers, firewalls, etc.)
 - Set-top boxes (unlikely due to \$\$\$???)
- High Performance Computing
 - Using Ada's "natural" parallelism rather than other languages "add on" features
 - Physics (including astrophysics), weather, video, data mining
 - Simulations
- Commodity Applications
 - Build once, sell **LOTS**

Challenge

- Ada community too cloistered
 - Many of the same speakers and topics
 - Dwindling numbers
- Limited focus
 - Academic, contractor and tools community
 - Niche market not allowing for growth
 - Few non-Ada focused product companies
- Constrained market view
 - Build onesy, twosey, maybe 26, maybe 200
 - Not build one, deploy thousands, millions, ...

Questions & Answers

- Is multicore necessarily slower clock speed?
 - Yes, heat dissipation and power consumption issues
- Why didn't you use GNAT GPL?
 - Obsolete version
 - Restrictive licensing
 - GCCFSS opportunity
- How about comparing Ada to other languages?
 - Offers of funding will be considered... 😊
- How does multicore differ from SMP or network distributed?
 - Essentially interconnect speeds
 - But beware the “8 Fallacies of Distributed Computing”