



Introduction of Parallel Program Execution and Architecture Models

Guang R. Gao

ACM Fellow and IEEE Fellow Endowed Distinguished Professor Electrical & Computer Engineering University of Delaware ggao@capsl.udel.edu







Outline



- A short history of parallel computing systems
- Program execution models (PXMs)
- Examples of parallel program execution models
 - SIMD and Vector processing model
 - MIMD model
 - Message-passing (e.g. MPI-style) execution model
 - Shared memory (e.g. OpenMP) execution model
 - ILP (instruction level parallelism) model
 - Superscalar model
 - VLIW model
- Evolution of (dynamic) fine-grain multithreaded program execution models.
- Summary





Recent Advances in Technologies

Memory DDR2 SDRAM 2004 1.8V vs. 2.5V for DDR, Up to 8,533MB/s DDR3 2007 1.5V or 1.35V for DDR3L, 6,400-17,066MB/s DDR4 2012 1.05-1.2V, 2133-4266MT/s

Solid State Drives

Significant reduction in price per GB since mid-2000 More expensive than HDDs of comparable capacity 2007: 320GB 100k IOPS 2009: 1TB SSD with 654MB/s write and 712MB/s read BW 2009: First SSD with 6Gbps SATA interface 2011: 2.2TB, 2.7GB/s read bandwidth 2012: 99.9% of 4KB random accesses within 0.5ms

InfiniBand

Scalable switched fabric communications technology High throughput and low latency Point-to-point bidirectional serial links Quality of service and failover features 2.5Gbps signaling rate in each direction per link in SDR speed Since June '12 InfiniBand is the dominant class of interconnects in TOP500

Ethernet

Gigabit Ethernet is the prevalent during last decade Gained the lead in TOP500 system count in June 2005 Reaches the peak of 56.6% of all TOP500 systems in June 2008 The dominance continues until June 2012 Fast Ethernet 100Mbps Gigabit Ethernet 1000Mbps 10Gigabit Ethernet 10Gbps 652.41-42

Cray Interconnects

Seastar (Red Storm, Cray XT3), 2004 3-D mesh topology with link throughput of 2.5GB/s in each direction SeaStar2 (Cray XT4), 2006 Peak bidirectional bandwidth per link: 7.6GB/s, SeaStar2+ (Cray XT5), 2009 9.6GB/s peak bidirectional bandwidth per link Gemini (Cray XE6, XK6, XK7), 2010 4 links per X and Z direction, 2 links per Y (10 total per NIC) Aries/Dragonfly (Cray XC30), 2012 High radix tiled router (48 tiles), 8 processor tiles, 4 NICs

^{1-intro} Courtesy of Prof. Thomas Sterling





World's Fastest Supercomputers



Japan, 2002: Earth Simulator 35.86 TFLOPS LINPACK



USA, 2004: Blue Gene/L 70.72 TFLOPS Linpack



USA, 2008: *Roadrunner* 1026 TFLOPS Linpack



USA, 2009: Jaguar 1759 TFLOPS Linpack



USA 2012: Sequoia 16324 TFLOPS Linpack



China 2010: Tianhe-1a 2566 TFLOPS Linpack



USA 2012: Titan 17590 TFLOPS Linpack



Japan 2011: K (京) 10510 TFLOPS Linpack



China 2013: Tianhe-2 33860 TFLOPS in HPL

Courtesy of Prof. T. Sterling





Tianhe-1A 2.566 Petaflops







Technology and Historical Perspective:

A peek of the microprocessor Evolution

Intel Pentium 4 Northwood



(10) Store Buffer (24 entries)

(14) Cache Line Read / Write Transferbuffers and 256 bit wide bus to and from L2 cache

four way set associative. 1R/1W

April 19, 2003 www.chip-architect.com

Intel Pentium 5 Prescott



April 19, 2003 www.chip-architect.com

- (14) Cache Line Read / Write Transferbuffers and 256 bit wide bus to and from L2 cache
- (12) 16 kByte Level 1 Data cache
- four way set associative. 1R/1W

(9) Load Buffer (96 entries) (10) Store Buffer (48 entries)





Technology Progress Overview

- Processor speed improvement: 2x per year (since 85). 100x in last decade.
- DRAM Memory Capacity: 2x in 2 years (since 96). 64x in last decade.
- DISK capacity: 2x per year (since 97).
 250x in last decade.



Num of Transistors x Clock Rate (Mil * MHz / 100)





Pentium M



Thermal Maps from the Pentium M obtained from simulated power density (left) and IREM measurement (right). Heat levels goes from black (lowest), red, orange, yellow and white (highest)

Figures courtesy of Dani Genossar and Nachum Shamir in their paper Intel ® Pentium ® M Processor Power Estimation, Bugdeting, Optimization and Validation published in the Intel Technical Journal, May 21, 2003





What Is Next?

- Move to "multiprocessor on a chip"?
 - cooler
 - simpler
 - cheaper

. . .





Architecture Features and Trends (Revisitd)

- core arch- simpler and simpler : RISC Core
- # of cores larger and larger : 160 cores
- on-chip memory per core smaller and smaller :
 < 32 KB/core
- On-chip bandwidth is becoming larger and larger : > 0.3 TB/sec
- Energy efficiency support more and more : 500 MHz



Outline



- A short history of parallel computing systems
- Program execution models (PXM)
- Examples of program execution models
 - Sequential execution model
 - Parallel execution model
 - SIMD and Vector processing model
 - MIMD model
 - Message-passing (e.g. MPI-style) execution model
 - Shared memory (e.g. OpenMP) execution model
 - ILP (instruction level parallelism) model
 - Superscalar model
 - VLIW model
- Evolution of (dynamic) fine-grain multithreaded program execution models.
- Summary

CAPSL

What is a Program Execution Model?



CAPSL

Features a User Program Depends On

Features expressed within a Programming language

But that's not all !!

Features expressed Outside a (typical) programming language

- Procedures; call/return
- Access to parameters and variables
- Use of data structures (static and dynamic)
- File creation, naming and access
- Object directories
- Communication: networks and peripherals
- Concurrency: coordination; scheduling



Developments in the 1960s, 1970s





Contour Model: Algorithm; Nested Blocks and Contours

- Johnston, 1971





Idea: A Common Base Language

This is a report on the work of the Computation Structures Group of Project MAC toward the design of a common base language for programs and information structures. We envision that the meanings of programs expressed in practical source languages will be defined by rules of translation into the base language.

The meanings of programs in the base language is fixed by rules of interpretation which constitute a transition system called the interpreter for the base language.

We view the base language as the functional specification of a computer system in which emphasis is placed on programming generality -- the ability of users to build complex programs by combining independently written program modules.

- Dennis, 1972





Terminology Clarification

- Parallel Model of Computation
 - Parallel Models for Algorithm Designers
 - Parallel Models for System Designers
 - Parallel Programming Models
 - Parallel Execution Models
 - Parallel Architecture Models





What Does Program Execution Model (PXM) Mean ?

• The notion of PXM

The program execution model (PXM) is the basic low-level abstraction of the underlying system architecture upon which our programming model, compilation strategy, runtime system, and other software components are developed.

• The PXM (and its API) serves as an interface between the architecture and the software.





Program Execution Model (PXM) – Cont'd

Unlike an instruction set architecture (ISA) specification, which usually focuses on lower level details (such as instruction encoding and organization of registers for a specific processor), the PXM refers to machine organization at a higher level for a whole class of high-end machines as view by the users

Gao, et. al., 2000



Execution Model and Abstract Machines





Abstract Machine Models May Be Heterogeneous!





Execution Model and Abstract Machines



Outline



- A short history of parallel computing systems
- Program execution models (PXM)
- Examples of program execution models
 - Sequential execution model
 - Parallel execution model
 - SIMD and Vector processing model
 - MIMD model
 - Message-passing (e.g. MPI-style) execution model
 - Shared memory (e.g. OpenMP) execution model
 - ILP (instruction level parallelism) model
 - Superscalar model
 - VLIW model
- Evolution of (dynamic) fine-grain multithreaded program execution models.
- Summary





What is your "Favorite" Program Execution Model?





Course Grain Execution Models

The Single Instruction Multiple Data (SIMD) Model



The Single Program Multiple Data (SPMD) Model



The Data Parallel Model







Data Parallel Model

Difficult to write unstructured programs

Convenient only for problems with regular structured parallelism.

Limited composability!

Inherent limitation of coarse-grain multithreading







Programming Models for Multi-Processor Systems

- Message Passing
 Model
 - Multiple address spaces
 - Communication can only be achieved through *"messages"*



- Shared Memory Model
 - Memory address space is accessible to all
 - Communication is achieved through memory







Comparison

Message Passing

- + Less Contention
- + Highly Scalable
- + Simplified Synch
 - Message Passing → Sync + Comm.
 - But does not mean highly programmable
- Load Balancing
- Deadlock prone
- Overhead of small messages

Shared Memory

- + global shared address space
- + Easy to program (?)
- + No (explicit) message passing (e.g. communication through memory put/get operations)
- Synchronization (memory consistency models, cache models)
- Scalability





Comment on OS impact?

- Should compiler be OS-Aware too ? If so, how ?
- Or other alternatives ? Compiler-controlled runtime, of compiler-aware kernels, etc.
- Example: software pipelining ...

Gao, ECCD Workshop, Washington D.C., Nov. 2007



Outline



- A short history of parallel computing systems
- Examples of program execution models
 - Sequential execution model
 - Parallel execution model
 - SIMD and Vector processing model
 - MIMD model
 - Message-passing (e.g. MPI-style) execution model
 - Shared memory (e.g. OpenMP) execution model
 - ILP (instruction level parallelism) model
 - Superscalar model
 - VLIW model
- Evolution of (dynamic) fine-grain multithreaded program execution models.
- What is program execution model anyway ?
- Summary





A Quiz: Have you heard the following terms ?

Actors (dataflow) ?

strand ?

fiber ? codelet ?







Coarse-Grain vs. Fine-Grain Multithreading

[Gao: invited talk at Fran Allen's Retirement Workshop, 07/2002]

652-14F-PXM-intro







Fine-Grain *non-preemptive* thread-The "hotel" model

Coarse-Grain vs. Fine-Grain Multithreading

[Gao: invited talk at Fran Allen's Retirement Workshop, 07/2002]

652-14F-PXM-intro



Evolution of Multithreaded Execution and Architecture Models









Outline



- A short history of parallel computing systems
- Examples of program execution models
 - Sequential execution model
 - Parallel execution model
 - SIMD and Vector processing model
 - MIMD model
 - Message-passing (e.g. MPI-style) execution model
 - Shared memory (e.g. OpenMP) execution model
 - ILP (instruction level parallelism) model
 - Superscalar model
 - VLIW model
- Evolution of (dynamic) fine-grain multithreaded program execution models.
- What is program execution model anyway ?
- Summary