2110412 Parallel Comp Arch Performance and Benchmarking

Natawut Nupairoj, Ph.D.
Department of Computer Engineering, Chulalongkorn University

Performance Questions

- How to characterize the performance of applications and systems?
- User's requirements in performance and cost?
- How about performance measurement?
- How will system perform when having more resources or more workload?

Important Keywords

- Peak Performance
  - Theoretical performance.
  - Typically, peak of single CPU * n
- Sustained Performance
  - The maximal achievable performance by running a benchmark.

Performance Metrics

- Indicators of how good the systems are.
- To evaluate correctly, we must consider:
  - What is the metric (or metrics) ?
  - What is its definition ?
  - How to measure it ? Benchmark algorithm ?
  - What is the evaluating environment ?
    - Configuration.
    - Workload.

Popular Metrics

- Time - Execution Time
- Rate - Throughput and Processing Speed
- Resource – Utilization
- Ratio - Cost Effectiveness
- Reliability – Error Rate
- Availability – Mean Time To Failure (MTTF)

Execution Time

- Aka. Wall clock time, elapsed time, delay.
- CPU time + I/O + user + …
- The lower, the better.
- Factors
  - Algorithm.
  - Data structure.
  - Input.
  - Hardware/Software/OS.
  - Language.

Definition of Time

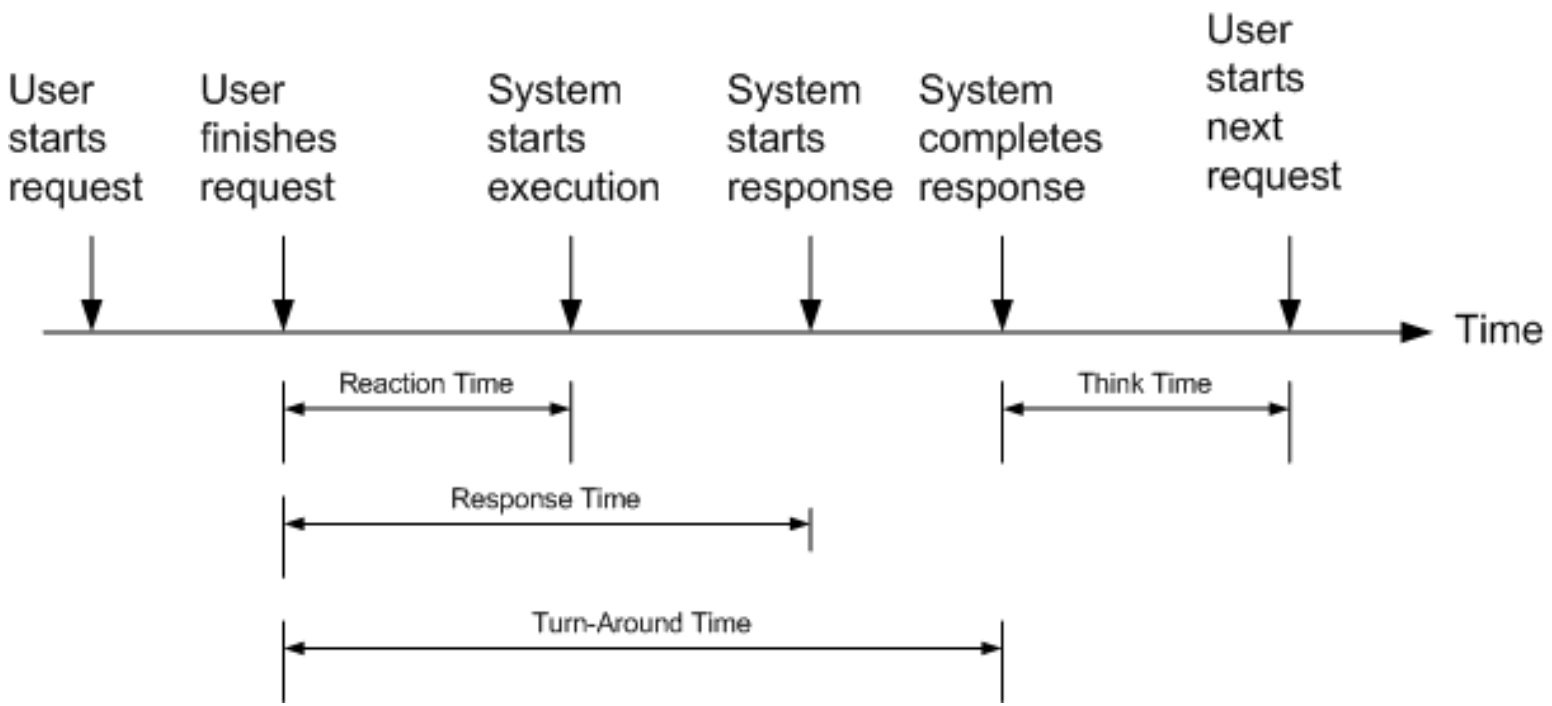

Analysis of Time

- Let's try "time" command for Unix

```
90.7u 12.9s 2:39 65%
```

- User time = 90.7 secs
- System time = 12.9 secs
- Elapsed time = 2 mins 39 secs = 159 secs
- (90.7 + 12.9) / 159 = 65%
- Meaning?

Processing Speed

- How fast can the system execute ?
- MIPS, MFLOPS.
- The more, the better.
- Can be very misleading !!!

```
k = m + n;
k = m + n;
k = m + n;
k = m + n;
...
```

```
for j=0 to x
    k = m + n;
```

```
for j=0 to x/4
    k = m + n;
    k = m + n;
    k = m + n;
    k = m + n;
```

Moore's Law (1965)

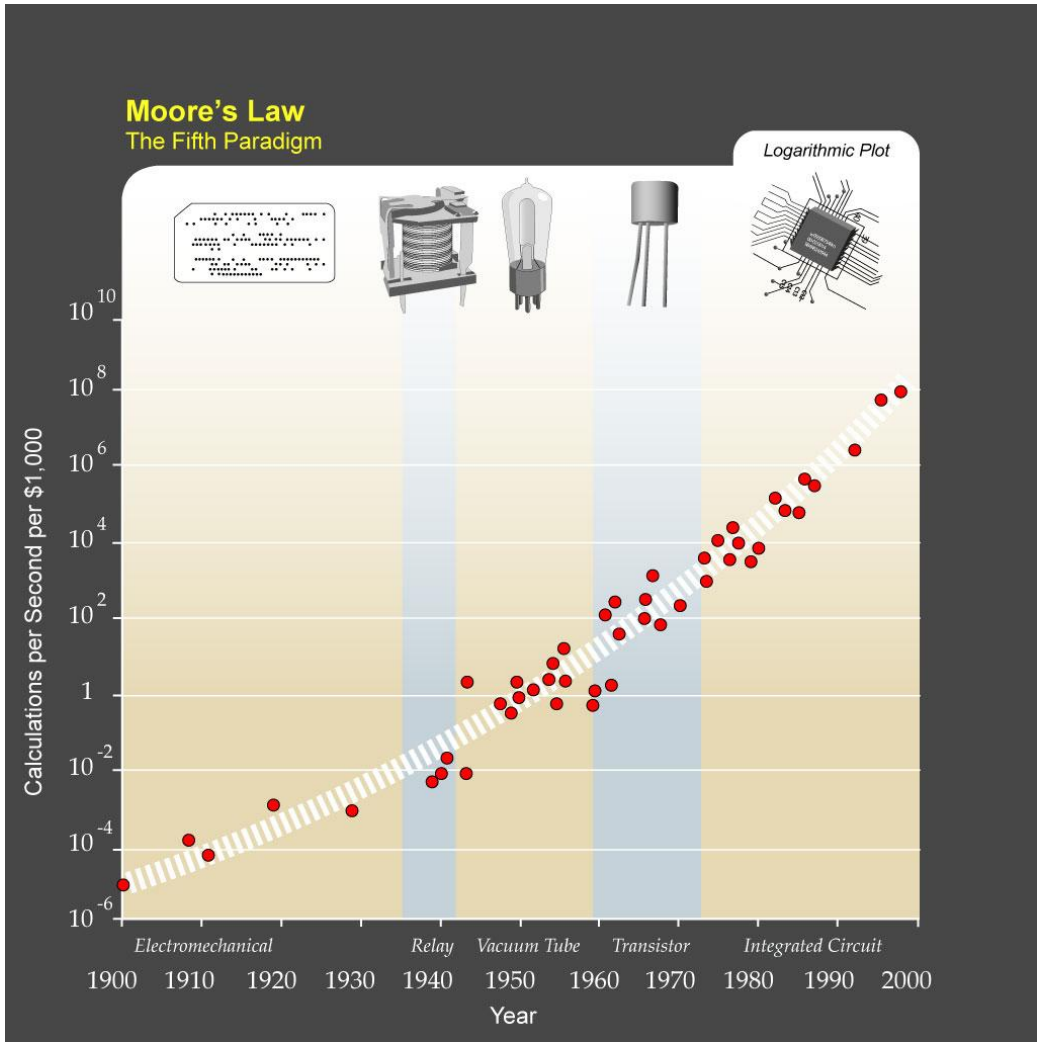

Kurzweil: The Law of Accelerating Returns

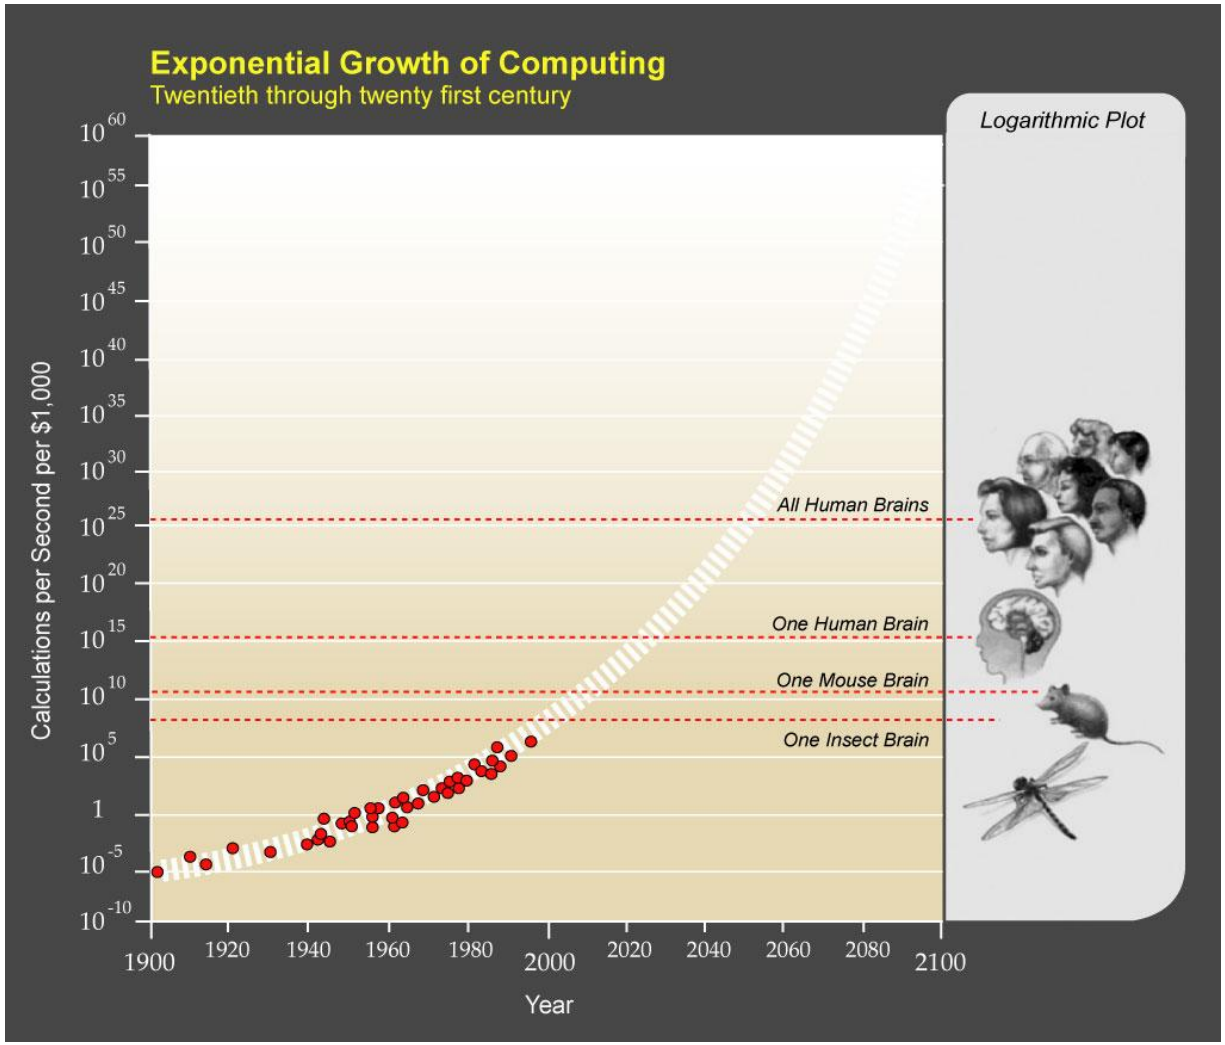

Throughput

- Number of jobs that can be processed in a unit time.
- Aka. Bandwidth (in communication).
- The more, the better.
- High throughput does not necessary mean low execution time.
  - Pipeline.
  - Multiple execution units.

Utilization

- The percentage of resources being used
- Ratio of
  - busy time vs. total time
  - sustained speed vs. peak speed
- The more the better?
  - True for manager
  - But may be not for user/customer
- Resource with highest utilization is the "bottleneck"

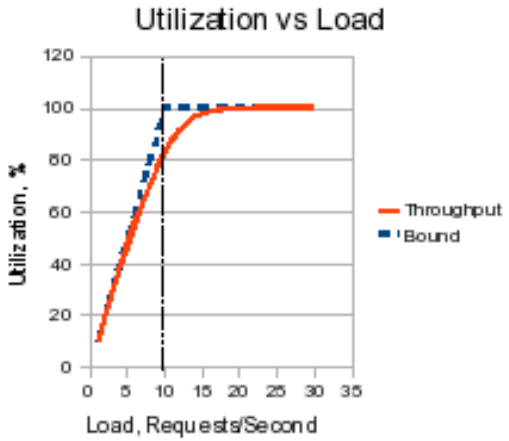

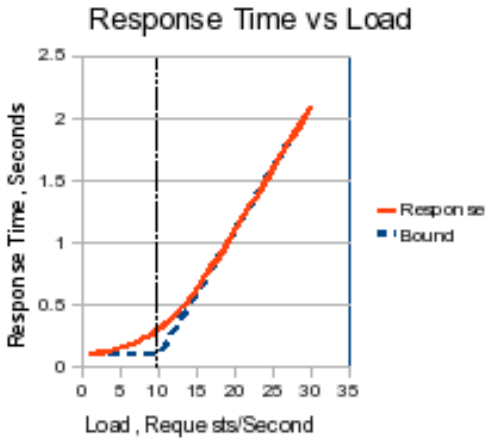

Typical Utilization when Running Program

- sustained speed vs. peak speed
- Sequential: 5-40%
  - Stalled Pipe.
  - I/O.
- Parallel: 1-35%
  - Low degree of parallelism.
  - Overheads: communication, I/O, OS, etc.

Cost Effectiveness

- Peak performance/cost ratio
- Price/performance ratio
- PCs are much better in this category than Supercomputer

Price/Performance Ratio

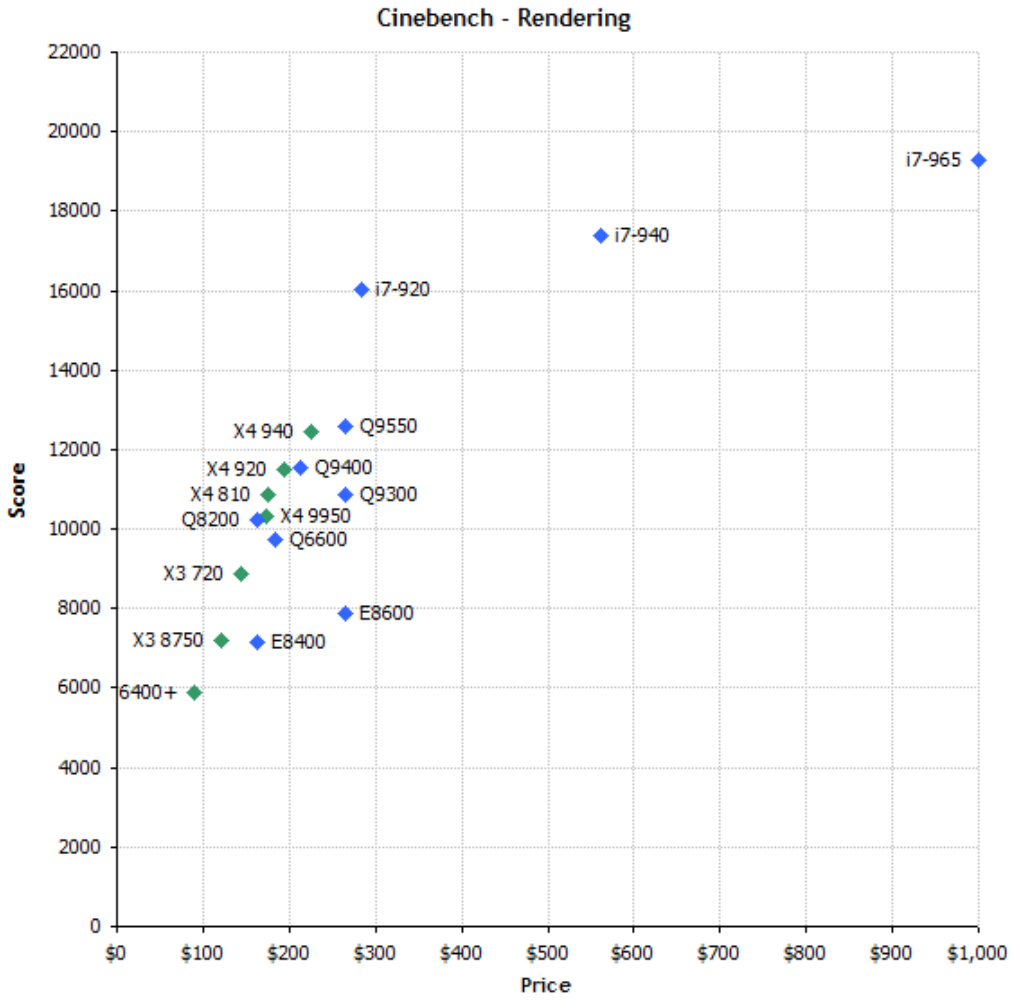

From Tom's Hardware Guide: CPU Chart 2009

Performance of Parallel Systems

- Factors
  - Components and architecture.
  - Degree of Parallelism.
  - Overheads.
- Architecture
  - CPU speed.
  - Memory size and speed.
  - Memory hierarchy.

Parallelism and Overheads

- Execution time

$$T = T_{par} + T_{seq} + T_{comm}$$

- Tpar – Time spent in Parallel
  - All nodes execute at the same time
  - Computation Time (mostly)
  - Depends on Algorithm
  - Load-imbalance (Degree of Parallelism)

Parallelism and Overheads

- Tseq – Time spent in Sequential
  - Only one node (usually master) do the job
  - Load / save data from disk
  - Critical sections
  - Usually, occurs during start and end of program
- Tcomm - Communication overhead
  - Communication between nodes
  - Data movement
  - Synchronization: barrier, lock, and critical region
  - Aggregation: reduction.

Speedup Analysis

- How good the parallel system is, when compared to the sequential system
  - Predict the scalability
- Speedup metrics
  - Amdahl's Law
  - Gustafson's Law

Execution Time Components

- Given program with Workload W:
  - Let α be the percentage of SEQUENTIAL portion in this program
  - Parallel portion = 1 - α

$$W = \alpha W + (1 - \alpha)W$$

Execution Time Components

- Suppose this program requires T time units on SINGLE processor:
  - T = Tpar + Tseq + Tcomm
  - Tpar = (1 - α)T
  - Tseq = αT
  - For simplicity ignore Tcomm

$$T = \alpha T + (1 - \alpha)T$$

Speedup Formula

$$\text{Speedup} = \frac{\text{Sequential execution time}}{\text{Parallel execution time}}$$

Amdahl's Law

- Aka. Fixed-Load (Problem) Speedup
  - Given workload W, how good it is if we have n processors (ignore communication) ?

$$S_n = \frac{\text{Time to execute W on 1 processor}}{\text{Time to execute W on n processor}}$$

$$T = \alpha T + (1-\alpha)T$$

$$S_n = \frac{T}{\alpha T + (1-\alpha)T/n} = \frac{n}{1+(n-1)\alpha} \to \frac{1}{\alpha} \text{ as } n \to \infty$$

Amdahl's Law (2)

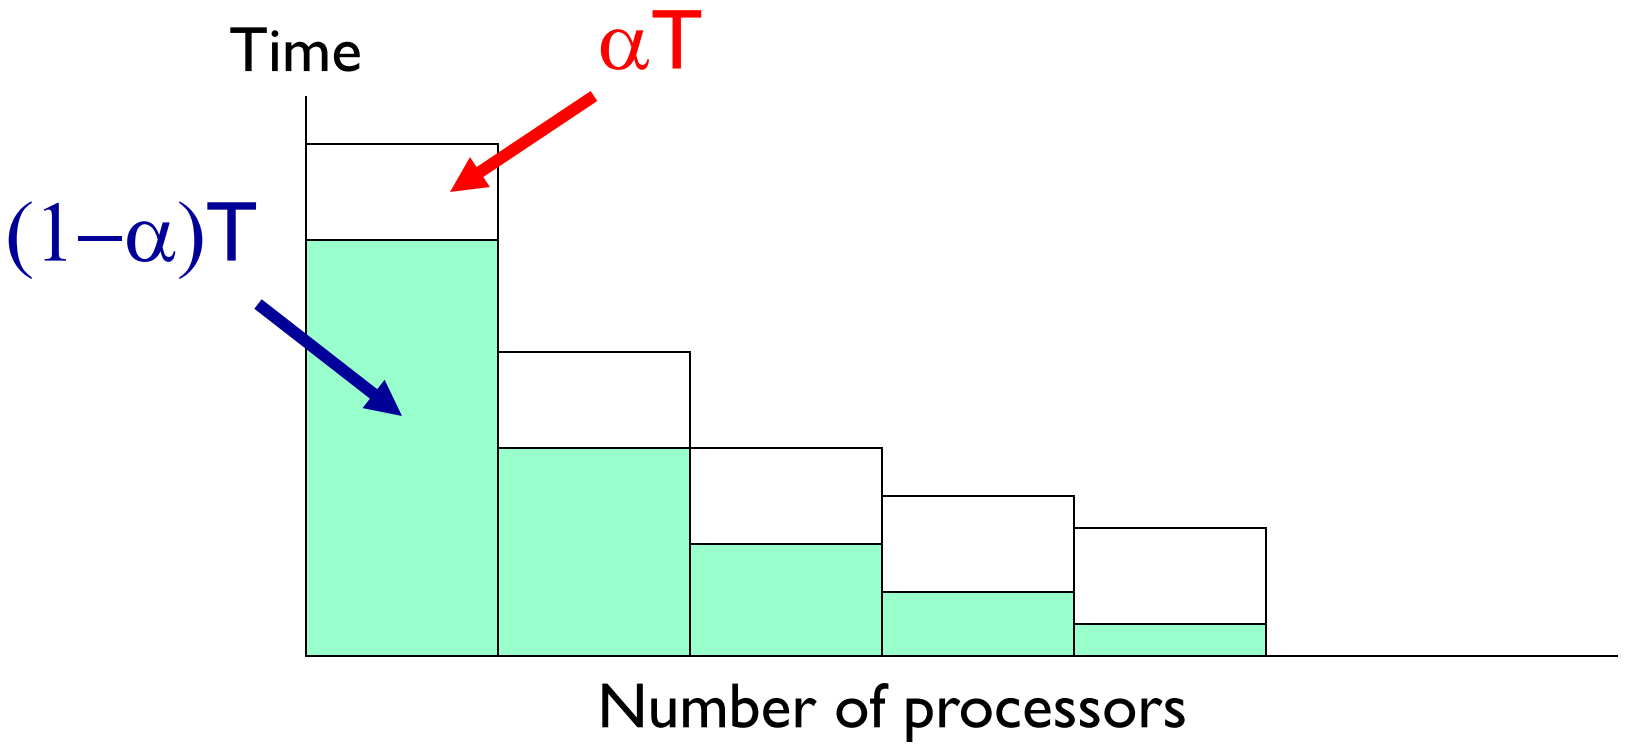

- Very popular (and also pessimistic).

Example 1

- 95% of a program's execution time occurs inside a loop that can be executed in parallel. What is the maximum speedup we should expect from a parallel version of the program executing on 8 CPUs?

Example 2

- 20% of a program's execution time is spent within inherently sequential code. What is the limit to the speedup achievable by a parallel version of the program?

Amdahl's Law (in Book)

$$\psi(n, p) \le \frac{\sigma(n) + \varphi(n)}{\sigma(n) + \varphi(n) / p + \kappa(n, p)}$$

$$\le \frac{\sigma(n) + \varphi(n)}{\sigma(n) + \varphi(n) / p}$$

Let $f = \sigma(n)/(\sigma(n) + \varphi(n))$

$$\psi \le \frac{1}{f + (1 - f) / p}$$

Limitations of Amdahl's Law

- Ignores Tcomm
  - Overestimates speedup achievable
- Very pessimistic
  - When people have bigger machines, they always run bigger programs
  - Thus, when people have more processors, they usually run bigger workloads
  - More workloads = more parallel portion
  - Workload may not be fixed, but SCALE

Problem Size and Amdahl's Law

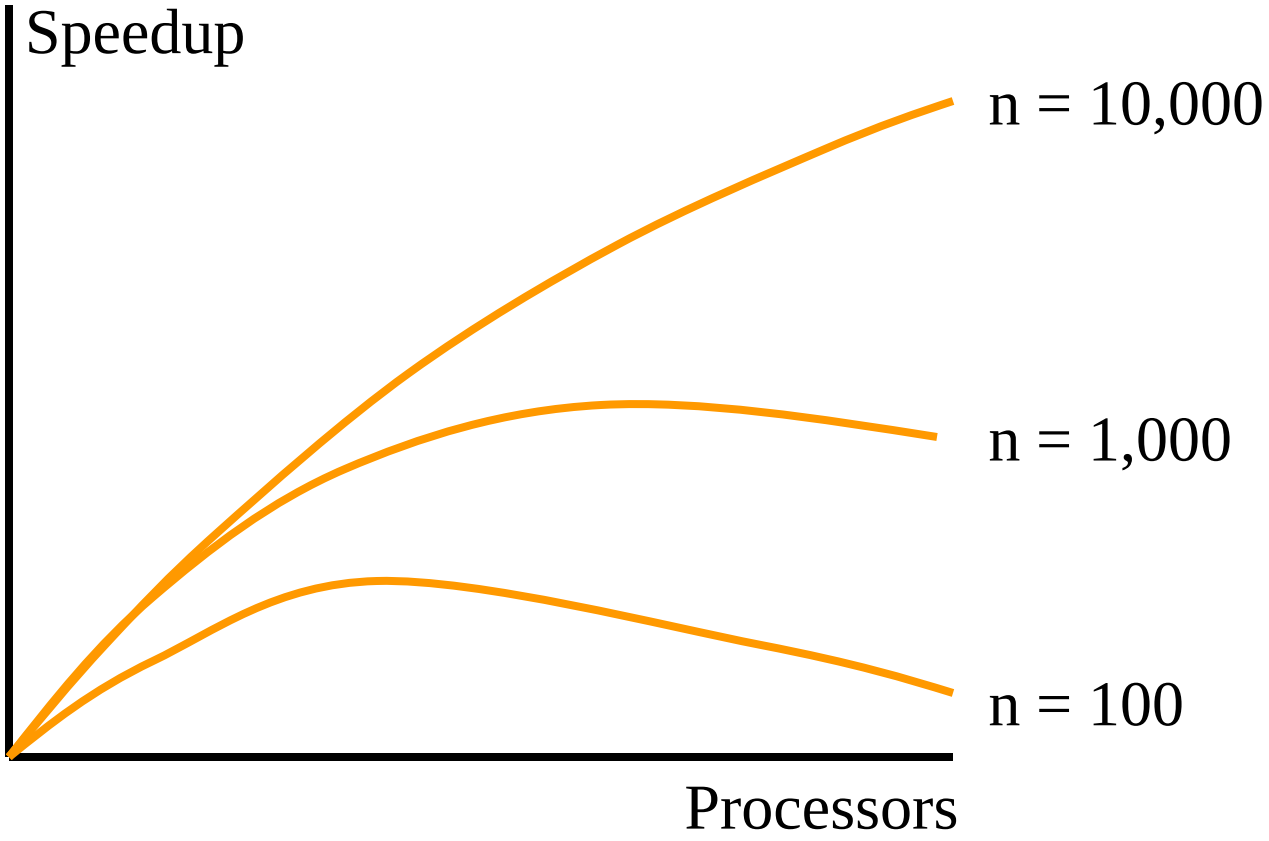

Gustafson's Law

- Aka. Fixed-Time Speedup (or Scaled-Load Speedup).
  - Given a workload W, suppose it takes time T to execute W on 1 processor.
  - With the same T, how much (workload) we can run on n processors ? Let's call it W'.
  - Assume the sequential work remains constant.

$$W = \alpha W + (1-\alpha)W \qquad W' = \alpha W + (1-\alpha)nW$$

Gustafson's Law (2)

- Fixed-Time Speedup

$$S'_n = \frac{\text{Workload size that can be executed in time T with n processors}}{\text{Workload size that can be executed in time T with 1 processors}}$$

$$S'_n = \frac{W'}{W} = \frac{\alpha W + (1-\alpha)nW}{W} = \alpha + (1-\alpha)n$$

Gustafson's Law (3)

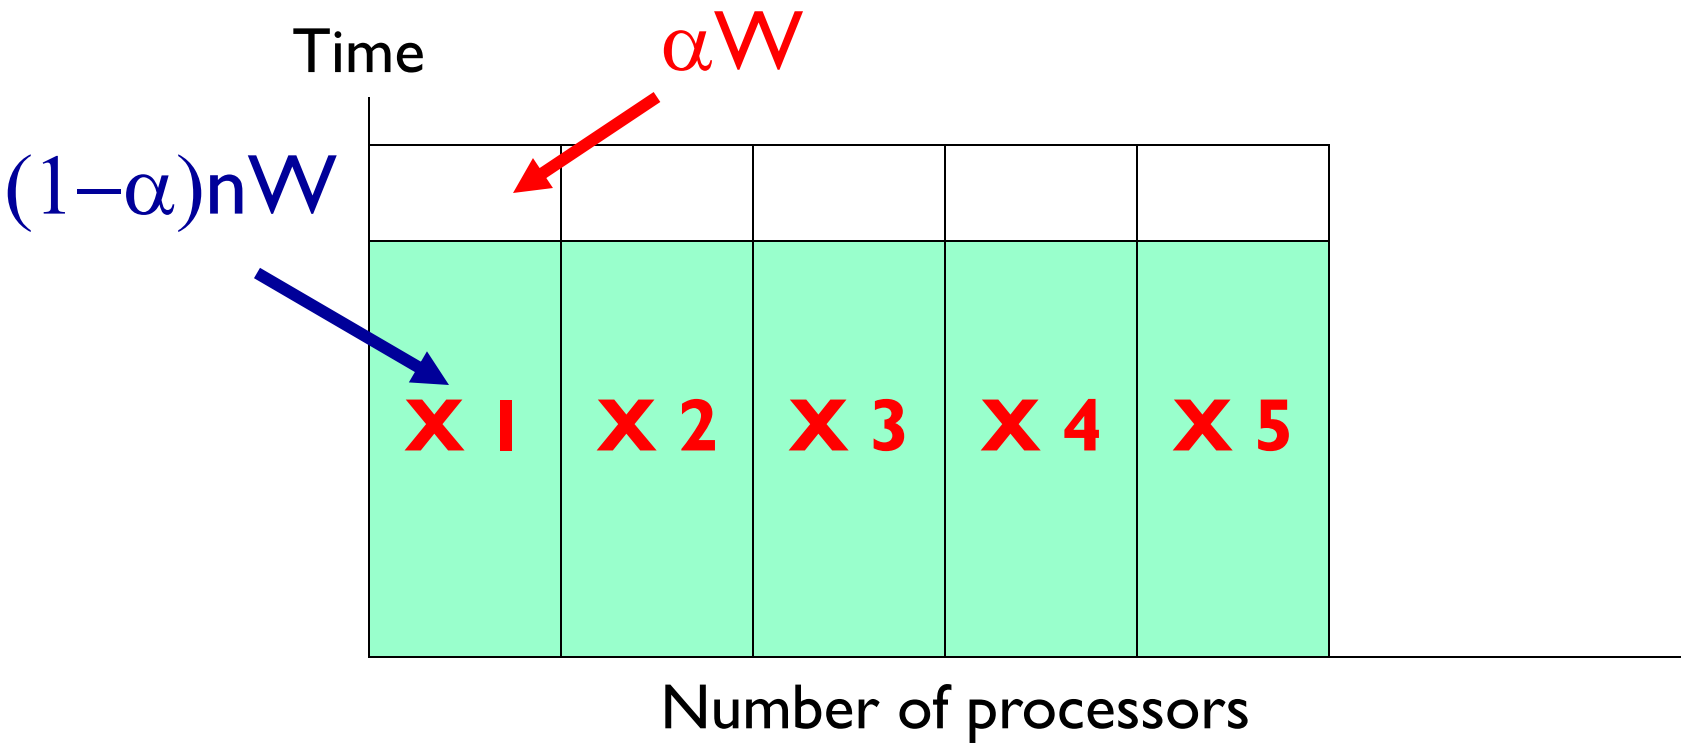

Example 1

- An application running on 10 processors spends 3% of its time in serial code. What is the scaled speedup of the application?

Example 2

- What is the maximum fraction of a program's parallel execution time that can be spent in serial code if it is to achieve a scaled speedup of 7 on 8 processors?

Performance Benchmarking

- Benchmark
  - Measure and predict the performance of a system
  - Reveal the strengths and weaknesses
- Benchmark Suite
  - A set of benchmark programs and testing conditions and procedures
- Benchmark Family
  - A set of benchmark suites

Benchmarks Classification

- By instructions
  - Full application
  - Kernel -- a set of frequently-used functions
- By workloads
  - Real programs
  - Synthetic programs

Popular Benchmark Suites

- SPEC
- TPC
- LINPACK

SPEC

- By Standard Performance Evaluation Corporation
- Using real applications
- http://www.spec.org
- SPEC CPU2006
  - Measure CPU performance
    - Raw speed of completing a single task
    - Rates of processing many tasks
  - CINT2006 - Integer performance
  - CFP2006 - Floating-point performance

CINT2006

400.perlbench	C	PERL Programming Language
401.bzip2	C	Compression
403.gcc	C	C Compiler
429.mcf	C	Combinatorial Optimization
445.gobmk	C	Artificial Intelligence: go
456.hmmer	C	Search Gene Sequence
458.sjeng	C	Artificial Intelligence: chess
462.libquantum	C	Physics: Quantum Computing
464.h264ref	C	Video Compression
471.omnetpp	C++	Discrete Event Simulation
473.astar	C++	Path-finding Algorithms
483.xalancbmk	C++	XML Processing

CFP2006

410.bwaves	Fortran	Fluid Dynamics
416.gamess	Fortran	Quantum Chemistry
433.milc	C	Physics: Quantum Chromodynamics
434.zeusmp	Fortran	Physics / CFD
435.gromacs	C/Fortran	Biochemistry/Molecular Dynamics
436.cactusADM	C/Fortran	Physics / General Relativity
437.leslie3d	Fortran	Fluid Dynamics
444.namd	C++	Biology / Molecular Dynamics
447.dealII	C++	Finite Element Analysis
450.soplex	C++	Linear Programming, Optimization
453.povray	C++	Image Ray-tracing
454.calculix	C/Fortran	Structural Mechanics
459.GemsFDTD	Fortran	Computational Electromagnetics
465.tonto	Fortran	Quantum Chemistry
470.lbm	C	Fluid Dynamics
481.wrf	C/Fortran	Weather Prediction
482.sphinx3	C	Speech recognition

Top 10 CINT2006 Speed (as of 1 Aug 2008)

System	Result	# Cores	# Chips	Cores/Chip	Processor
HP ProLiant DL160 G5 (3.4 GHz, Intel Xeon X5272)	28.4	4	2	2	Intel Xeon X5272
SGI Altix XE 250 (Intel Xeon X5272 3.4GHz)	28.4	4	2	2	Intel Xeon X5272
HP ProLiant DL380 G5 (3.16 GHz, Intel Xeon X5460)	27.7	8	2	4	Intel Xeon X5460
IBM System x 3550 (Intel Xeon X5460)	27.7	8	2	4	Intel Xeon X5460
Sun Fire X4150	27.7	8	2	4	Intel Xeon X5460
Fujitsu CELSIUS R550, Intel Xeon X5460 processor	27.6	8	2	4	Intel Xeon X5460
HP ProLiant BL480c (3.16 GHz, Intel Xeon X5460)	27.6	8	2	4	Intel Xeon X5460
HP ProLiant DL360 G5 (3.16 GHz, Intel Xeon processor X5460)	27.6	8	2	4	Intel Xeon X5460
HP ProLiant ML370 G5 (3.33 GHz, Intel Xeon processor X5260)	27.6	4	2	2	Intel Xeon X5260
IBM BladeCenter HS21 (Intel Xeon X5460)	27.6	8	2	4	Intel Xeon X5460

Top 10 CINT2006 Speed (as of 29 July 2009)

System	Result	# Cores	# Chips	Cores/Chip	Processor
Sun Blade X6275 (Intel Xeon X5570 2.93GHz)	37.4	8	2	4	Intel Xeon X5570
ASUS TS700-E6 (Z8PE-D12X) server system (Intel Xeon W5580)	37.3	8	2	4	Intel Xeon W5580
CELSIUS R670, Intel Xeon W5580	37.2	8	2	4	Intel Xeon W5580
Sun Blade X6270 (Intel Xeon X5570 2.93GHz)	36.9	8	2	4	Intel Xeon X5570
Sun Ultra 27 (Intel Xeon W3570 3.2GHz)	36.8	4	1	4	Intel Xeon W3570
Sun Fire X4170 (Intel Xeon X5570 2.93GHz)	36.8	8	2	4	Intel Xeon X5570
Sun Blade X6270 (Intel Xeon X5570 2.93GHz)	36.8	8	2	4	Intel Xeon X5570
Sun Blade X6275 (Intel Xeon X5570 2.93GHz)	36.7	8	2	4	Intel Xeon X5570
Dell Precision T7500 (Intel Xeon W5580, 3.20 GHz)	36.7	8	2	4	Intel Xeon W5580
CELSIUS M470, Intel Xeon W5580	36.6	4	1	4	Intel Xeon W5580

Other Interesting SPECs

- SPEC MPI2007
  - Benchmark based on MPI to measure floating-point computational intensive applications on clusters and SMP
- SPEC jAppServer2004
  - Measure the performance of J2EE 1.3 application servers
- SPEC Web2009
  - Emulates users sending browser requests over broadband Internet connections to a web server
- SPECpower_ssj2008
  - Evaluates the power and performance characteristics of volume server class computers

TPC

- Transaction Processing Performance Council
- http://www.tpc.org
- TPC-C: performance of Online Transaction Processing (OLTP) system
  - tpmC: transactions per minute.
  - $/tpmC: price/performance.
- Simulate the wholesale company environment
  - N warehouses, 10 sales districts each.
  - Each district serves 3,000 customers with one terminal in each district.

TPC Transactions

- An operator can perform one of the five transactions
  - Create a new order.
  - Make a payment.
  - Check the order's status.
  - Deliver an order.
  - Examine the current stock level.
- Measure from the throughput of New-Order.
- Top 10 (Performance, Price/Performance).

Top 10 TPC-C Performance (as of 1 Aug 2008)

Rank	Company	System	tpmC	Price/tpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	IBM	IBM Power 595 Server Model 9119-FHA	6,085,166	2.81 US $	12/10/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/10/08	N
***	Bull	Bull Escala PL6460R	6,085,166	2.81 US $	12/15/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/15/08	N
2	HP	HP Integrity Superdome-Itanium2/1.6GHz/24MB iL3	4,092,799	2.93 US $	08/06/07	Oracle Database 10g R2 Enterprise Edt w/Partitioning	HP-UX 11i v3	BEA Tuxedo 8.0	02/27/07	N
3	IBM	IBM System p5 595	4,033,378	2.97 US $	01/22/07	IBM DB2 9	IBM AIX 5L V5.3	Microsoft COM+	01/22/07	N
4	IBM	IBM eServer p5 595	3,210,540	5.07 US $	05/14/05	IBM DB2 UDB 8.2	IBM AIX 5L V5.3	Microsoft COM+	11/18/04	N
5	FUJITSU	PRIMEQUEST 580 32p/64c	2,196,268	4.70 US $	04/30/08	Oracle 10g Enterprise Ed R2 w/ Partitioning	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	10/30/07	N
6	IBM	IBM System p 570	1,616,162	3.54 US $	11/21/07	IBM DB2 Enterprise 9	IBM AIX 5L V5.3	Microsoft COM+	05/21/07	N
***	Bull	Bull Escala PL1660R	1,616,162	3.54 US $	12/16/07	IBM DB2 9.1	IBM AIX 5L V5.3	Microsoft COM+	12/17/07	N
7	IBM	IBM eServer p5 595	1,601,784	5.05 US $	04/20/05	Oracle Database 10g Enterprise Edition	IBM AIX 5L V5.3	Microsoft COM+	04/20/05	N
8	NEC	NEC Express5800/1320Xf (16p/32c)	1,245,516	4.57 US $	04/30/08	Oracle Database 10g R2 Enterprise Edt w/Partitioning	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	01/21/08	N
9	FUJITSU	PRIMEQUEST 540 16p/32c	1,238,579	3.94 US $	12/15/06	Oracle Database 10g Enterprise Edition	Red Hat Enterprise Linux AS 4.0	BEA Tuxedo 8.1	11/30/06	N
10	HP	HP Integrity Superdome – Itanium2/1.6 GHz-64p/64c	1,231,433	4.82 US $	06/05/06	Microsoft SQL Server 2005 Enterprise Edt SP1	Microsoft Windows Server 2003 Datacenter Ed.(64-bit)SP1	Microsoft COM+	11/28/05	N

Top 10 TPC-C Performance (as of 29 July 2009)

Rank	Company	System	tpmC	Price/tpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	IBM	IBM Power 595 Server Model 9119-FHA	6,085,166	2.81 USD	12/10/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/10/08	N
***	Bull	Bull Escala PL6460R	6,085,166	2.81 USD	12/15/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/15/08	N
2	HP	HP Integrity Superdome-Itanium2/1.6GHz/24MB iL3	4,092,799	2.93 USD	08/06/07	Oracle Database 10g R2 Enterprise Edt w/Partitioning	HP-UX 11i v3	BEA Tuxedo 8.0	02/27/07	N
3	IBM	IBM System p5 595	4,033,378	2.97 USD	01/22/07	IBM DB2 9	IBM AIX 5L V5.3	Microsoft COM+	01/22/07	N
4	IBM	IBM eServer p5 595	3,210,540	5.07 USD	05/14/05	IBM DB2 UDB 8.2	IBM AIX 5L V5.3	Microsoft COM+	11/18/04	N
5	FUJITSU	PRIMEQUEST 580A 32p/64c	2,382,032	3.76 USD	12/04/08	Oracle Database 10g R2 Enterprise Edt w/Partitioning	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	12/04/08	N
6	FUJITSU	PRIMEQUEST 580 32p/64c	2,196,268	4.70 USD	04/30/08	Oracle 10g Enterprise Ed R2 w/ Partitioning	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	10/30/07	N
7	IBM	IBM System p 570	1,616,162	3.54 USD	11/21/07	IBM DB2 Enterprise 9	IBM AIX 5L V5.3	Microsoft COM+	05/21/07	N
***	Bull	Bull Escala PL1660R	1,616,162	3.54 USD	12/16/07	IBM DB2 9.1	IBM AIX 5L V5.3	Microsoft COM+	12/17/07	N
8	IBM	IBM eServer p5 595	1,601,784	5.05 USD	04/20/05	Oracle Database 10g Enterprise Edition	IBM AIX 5L V5.3	Microsoft COM+	04/20/05	N
9	FUJITSU	PRIMEQUEST 540A 16p/32c	1,354,086	3.25 USD	11/22/08	Oracle Database 10g release2 Enterprise Edt	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	11/22/08	N
10	NEC	NEC Express5800/1320Xf (16p/32c)	1,245,516	4.57 USD	04/30/08	Oracle Database 10g R2 Enterprise Edt w/Partitioning	Red Hat Enterprise Linux 4 AS	BEA Tuxedo 8.1	01/21/08	N

Top 10 TPC-C Price/Performance (as of 1 Aug 2008)

Rank	Company	System	tpmC	Price/tpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	DELL	Dell PowerEdge 2900	97,083	.68 US $	06/16/08	Oracle Database 11g Standard Edition One	Microsoft Windows Server 2003 Standard Ed. x64	Microsoft COM+	06/16/08	N
2	hp invent	HP ProLiant ML350G5	102,454	.73 US $	12/31/07	Oracle Database 11g Standard Edition One	Microsoft Windows Standard x64 Etd. SP1 R2	Microsoft COM+	09/12/07	N
3	hp invent	HP ProLiant ML350G5	100,926	.74 US $	06/08/07	Oracle Database 10g Standard Edition One	Oracle Enterprise Linux	Microsoft COM+	06/08/07	N
4	hp invent	HP ProLiant ML350G5	82,774	.84 US $	03/27/07	Microsoft SQL Server 2005 x64 Enterprise Edt. SP1	Microsoft Windows 2003 x64 Server Std. Ed.	Microsoft COM+	03/27/07	N
5	SYBASE iAnywhere	Dell PowerEdge 2950 III	20,705	.85 US $	08/05/08	Sybase SQL Anywhere 11	Microsoft Windows 2003 x64 Standard R2 SP2	Microsoft COM+	07/29/08	N
6	DELL	PowerEdge 2900/1 /2.33GHz/2x4M	69,564	.91 US $	03/09/07	Microsoft SQL Server 2005 Standard Ed.	Microsoft Windows 2003 Server Std Edt SP1	Microsoft COM+	03/09/07	N
7	DELL	PowerEdge 2900/3.0GHz/4M	65,833	.98 US $	06/26/06	Microsoft SQL Server 2005 Standard Ed.	Microsoft Windows 2003 Server Std Edt SP1	Microsoft COM+	06/30/06	N
8	DELL	PowerEdge 2800/1 /2.8GHz/2+2M	38,622	.99 US $	11/08/05	Microsoft SQL Server 2005 x64 Std. Ed.	Microsoft Windows 2003 x64 Server Std. Ed.	Microsoft COM+	09/26/05	N
9	hp invent	HP ProLiant DL585G5/2.5GHz	471,883	1.17 US $	07/14/08	Microsoft SQL Server 2005 x64 Enterprise Edt SP2	Microsoft Windows Server 2003 Enterprise x64 Ent. R2	Microsoft COM+	07/14/08	N
10	hp invent	HP ProLiant DL585G5/2.3GHz	402,234	1.26 US $	03/31/08	Microsoft SQL Server 2005 x64 Enterprise Edt SP2	Microsoft Windows Server 2003 Enterprise x64 Ent. R2	Microsoft COM+	03/31/08	N

Top 10 TPC-C Price/Performance (as of 29 July 2009)

Rank	Company	System	tpmC	Price/tpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	hp invent	HP ProLiant ML350 G6	232,002	.54 USD	05/21/09	Oracle Database 11g Standard Edition One	Oracle Enterprise Linux	Microsoft COM+	05/21/09	N
2	DELL	Dell PowerEdge 2900	104,492	.60 USD	02/20/09	Oracle Database 11g Standard Edition One	Microsoft Windows Server 2003 Standard Ed. x64	Microsoft COM+	02/20/09	N
3	DELL	Dell PowerEdge 2900	97,083	.68 USD	06/16/08	Oracle Database 11g Standard Edition One	Microsoft Windows Server 2003 Standard Ed. x64	Microsoft COM+	06/16/08	N
4	hp invent	HP ProLiant ML350G5	102,454	.73 USD	12/31/07	Oracle Database 11g Standard Edition One	Microsoft Windows Standard x64 Etd. SP1 R2	Microsoft COM+	09/12/07	N
5	hp invent	HP ProLiant ML350G5	100,926	.74 USD	06/08/07	Oracle Database 10g Standard Edition One	Oracle Enterprise Linux	Microsoft COM+	06/08/07	N
6	hp invent	HP ProLiant ML350G5	82,774	.84 USD	03/27/07	Microsoft SQL Server 2005 x64 Enterprise Edt. SP1	Microsoft Windows 2003 x64 Server Std. Ed.	Microsoft COM+	03/27/07	N
7	SYBASE iAnywhere	Dell PowerEdge 2950 III	20,705	.85 USD	08/05/08	Sybase SQL Anywhere 11.0	Microsoft Windows 2003 x64 Standard R2 SP2	Microsoft COM+	07/29/08	N
8	DELL	PowerEdge 2900/1/2.33GHz/2x4M	69,564	.91 USD	03/09/07	Microsoft SQL Server 2005 Standard Ed.	Microsoft Windows 2003 Server Std Edt SP1	Microsoft COM+	03/09/07	N
9	hp invent	HP ProLiant DL585G5/2.7GHz	579,814	.96 USD	11/17/08	Microsoft SQL Server 2005 x64 Enterprise Edt SP2	Microsoft Windows Server 2003 Enterprise x64 Ent. R2	Microsoft COM+	11/17/08	N
10	hp invent	HP ProLiant DL580G5	639,253	.97 USD	01/26/09	Oracle Database 11g Standard Edition	Oracle Enterprise Linux TP	Microsoft COM+	01/16/09	N

LINPACK

- Linear Algebra Package
- By Jack Dongarra at University of Tennessee
- http://www.top500.org
- Collection of FORTRAN subroutines
  - Solve linear equations
  - Numerical, Micro, Kernel, Synthetic
  - Used in Top-500 list

LINPACK

- Metrics and parameters
  - R(max) - sustained maximal speed achieved.
  - N(max) - problem size when R(max) is achieved.
  - N(1/2) - problem size when half of R(max).
  - R(peak) - theoretical peak speed of the system measured.
- Top-500 list
  - See results.

LINPACK - Results Interpretation

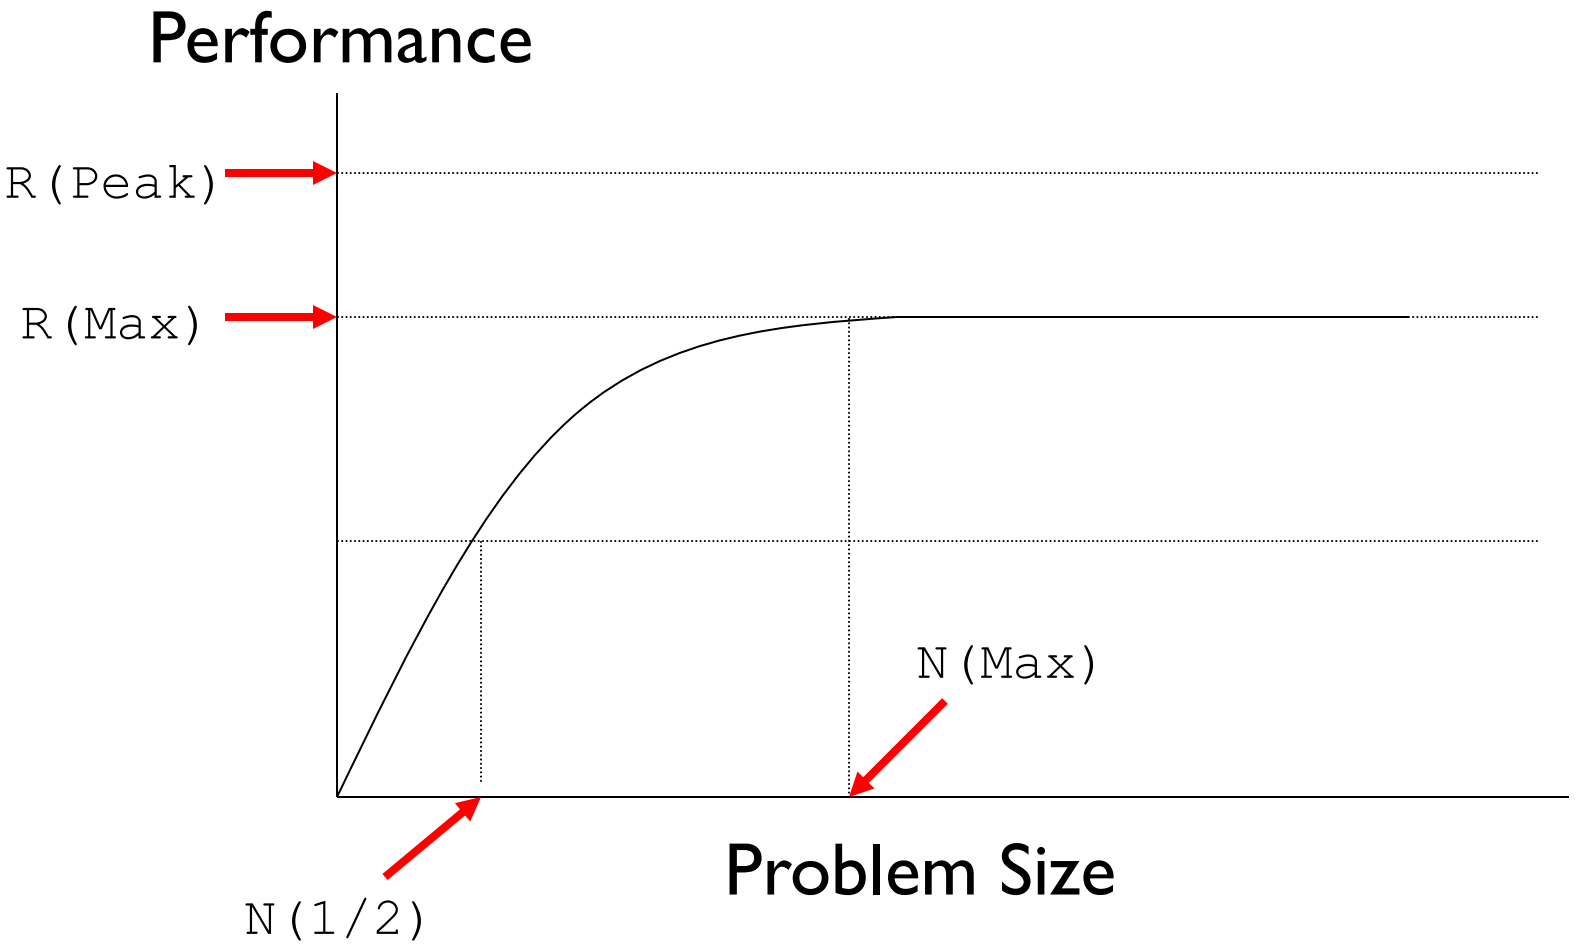

Top 10 of Top 500 Performance (as of June 2008)

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband / 2008 IBM	122400	1026.00	1375.78	2345.50
2	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60
3	Argonne National Laboratory United States	Blue Gene/P Solution / 2007 IBM	163840	450.30	557.06	1260.00
4	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron Quad 2Ghz, Infiniband / 2008 Sun Microsystems	62976	326.00	503.81	2000.00
5	DOE/Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz / 2008 Cray Inc.	30976	205.00	260.20	1580.71
6	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution / 2007 IBM	65536	180.00	222.82	504.00
7	New Mexico Computing Applications Center (NMCAC) United States	Encanto - SGI Altix ICE 8200, Xeon quad core 3.0 GHz / 2007 SGI	14336	133.20	172.03	861.63
8	Computational Research Laboratories, TATA SONS India	EKA - Cluster Platform 3000 BL460c, Xeon 53xx 3GHz, Infiniband / 2008 Hewlett-Packard	14384	132.80	172.61	786.00
9	IDRIS France	Blue Gene/P Solution / 2008 IBM	40960	112.50	139.26	315.00
10	Total Exploration Production France	SGI Altix ICE 8200EX, Xeon quad core 3.0 GHz / 2008 SGI	10240	106.10	122.88	442.00

Top 10 of Top 500 Performance (as of June 2009)

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2008 IBM	129600	1105.00	1456.70	2483.47
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT5 QC 2.3 GHz / 2008 Cray Inc.	150152	1059.00	1381.40	6950.60
3	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution / 2009 IBM	294912	825.50	1002.70	2268.00
4	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz / 2008 SGI	51200	487.01	608.83	2090.00
5	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60
6	National Institute for Computational Sciences/University of Tennessee United States	Kraken XT5 - Cray XT5 QC 2.3 GHz / 2008 Cray Inc.	66000	463.30	607.20	
7	Argonne National Laboratory United States	Blue Gene/P Solution / 2007 IBM	163840	458.61	557.06	1260.00
8	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband / 2008 Sun Microsystems	62976	433.20	579.38	2000.00
9	DOE/NNSA/LLNL United States	Dawn - Blue Gene/P Solution / 2009 IBM	147456	415.70	501.35	1134.00
10	Forschungszentrum Juelich (FZJ) Germany	JUROPA - Sun Constellation, NovaScale R422-E2, Intel Xeon X5570, 2.93 GHz, Sun M9/Mellanox QDR Infiniband/Partec Parastation / 2009 Bull SA	26304	274.80	308.28	1549.00

Top 500 – Projected Performance (as of June 2009)

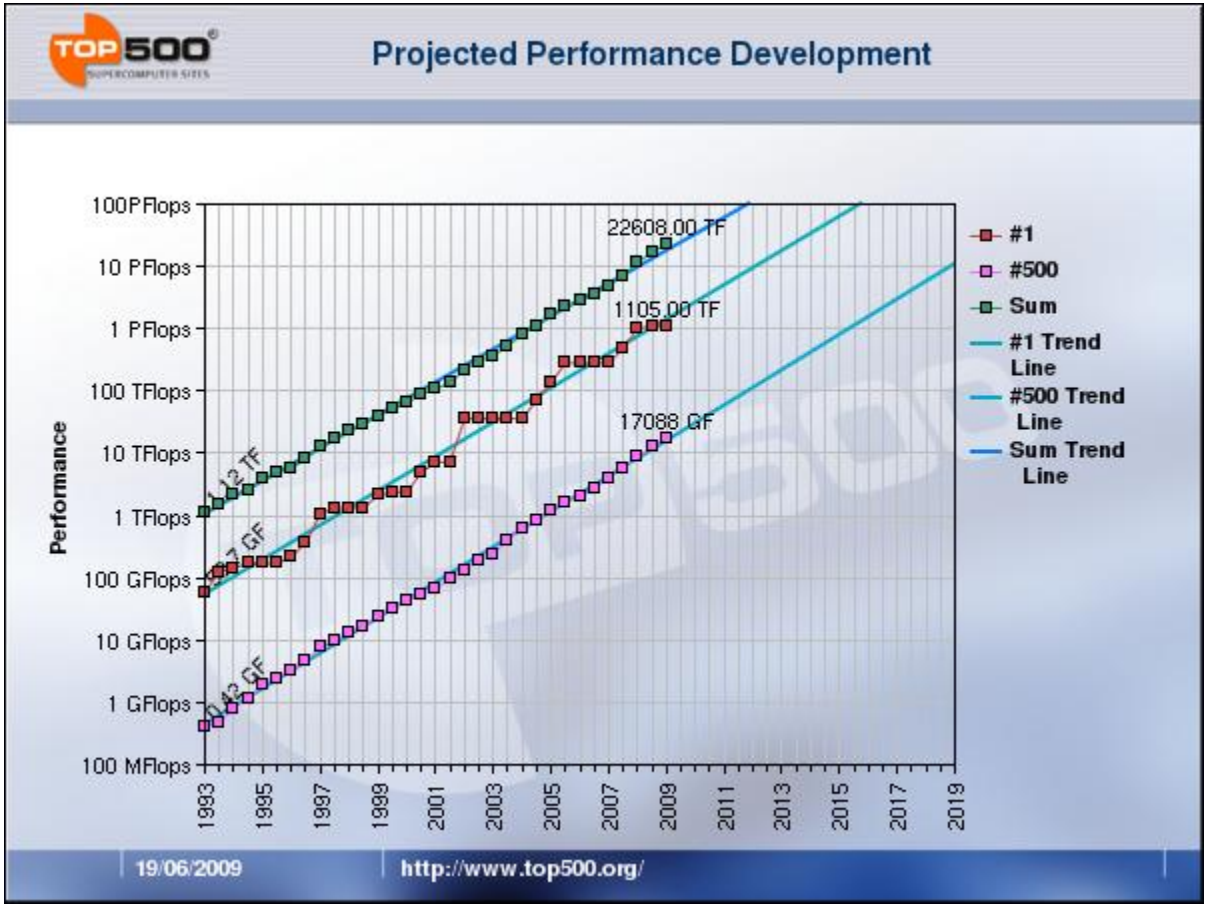

Top 500 – Architecture Distribution (as of June 2009)