


*In the name of God*

# Computer Architecture


S.Ali.Zendehbad



# SPEC CPU Benchmark

- Programs used to measure performance
  - Supposedly typical of actual workload
- Standard Performance Evaluation Corp (SPEC)
  - Develops benchmarks for CPU, I/O, Web, ...
- SPEC CPU2006
  - Elapsed time to execute a selection of programs
    - Negligible I/O, so focuses on CPU performance
  - Normalize relative to reference machine
  - Summarize as geometric mean of performance ratios
    - CINT2006 (integer) and CFP2006 (floating-point)

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$



$$\sqrt[n]{\prod_{i=1}^n \text{Execution time ratio}_i}$$

# SPEC CPU Benchmark

$$\text{Performance} = \frac{1}{\text{Execution Time}}$$

$$\frac{\text{Performance Computer A}}{\text{Performance Computer Ref}} = \frac{\text{CPU Time Ref}}{\text{CPU Time A}}$$



**Execution Time Ratio (ETR) for each Benchmark**

$$\sqrt[n]{\prod_{i=1}^n \text{Execution Time Ratio}_i}$$

# SPEC CPU Benchmark

$$\sqrt[2]{\prod_1^2 \text{Ratio } i} = \sqrt[2]{\text{Ratio } 1 \times \text{Ratio } 2} = \sqrt[2]{(CT \text{ Ref}1)/(CT \text{ A}1) \times (CT \text{ Ref}2)/(CT \text{ A}2)}$$

$$\frac{\text{Performance Computer A}}{\text{Performance Computer B}} = \frac{ETR B}{ETR A}$$

$$\sqrt[2]{\prod_1^2 \text{Ratio } i} = \frac{\sqrt[2]{(CT \text{ Ref}1)/(CT \text{ A}1) \times (CT \text{ Ref}2)/(CT \text{ A}2)}}{\sqrt[2]{(CT \text{ Ref}1)/(CT \text{ B}1) \times (CT \text{ Ref}2)/(CT \text{ B}2)}}$$

Instruction Count × CPI × Clock Cycle Time

# SPEC CPU Benchmark

## CINT2006 for Opteron X4 2356

Name	Description	IC×10 <sup>9</sup>	CPI	Tc (ns)	Exec time	Ref time	SPECratio
perl	Interpreted string processing	2,118	0.75	0.40	637	9,777	15.3
bzip2	Block-sorting compression	2,389	0.85	0.40	817	9,650	11.8
gcc	GNU C Compiler	1,050	1.72	0.47	24	8,050	11.1
mcf	Combinatorial optimization	336	10.00	0.40	1,345	9,120	6.8
go	Go game (AI)	1,658	1.09	0.40	721	10,490	14.6
hmmer	Search gene sequence	2,783	0.80	0.40	890	9,330	10.5
sjeng	Chess game (AI)	2,176	0.96	0.48	37	12,100	14.5
libquantum	Quantum computer simulation	1,623	1.61	0.40	1,047	20,720	19.8
h264avc	Video compression	3,102	0.80	0.40	993	22,130	22.3
omnetpp	Discrete event simulation	587	2.94	0.40	690	6,250	9.1
astar	Games/path finding	1,082	1.79	0.40	773	7,020	9.1
xalancbmk	XML parsing	1,058	2.70	0.40	1,143	6,900	6.0
Geometric mean							11.7



$$\sqrt[n]{\prod_{i=1}^n \text{Execution Time Ratio}_i}$$



# SPEC CPU Benchmark

## CINT2006 for Intel Core i7 920

Description	Name	Instruction Count x 10 <sup>9</sup>	CPI	Clock cycle time (seconds x 10 <sup>-9</sup> )	Execution Time (seconds)	Reference Time (seconds)	SPECratio
Interpreted string processing	perl	2252	0.60	0.376	508	9770	19.2
Block-sorting compression	bzip2	2390	0.70	0.376	629	9650	15.4
GNU C compiler	gcc	794	1.20	0.376	358	8050	22.5
Combinatorial optimization	mcf	221	2.66	0.376	221	9120	41.2
Go game (AI)	go	1274	1.10	0.376	527	10490	19.9
Search gene sequence	hmmer	2616	0.60	0.376	590	9330	15.8
Chess game (AI)	sjeng	1948	0.80	0.376	586	12100	20.7
Quantum computer simulation	libquantum	659	0.44	0.376	109	20720	190.0
Video compression	h264avc	3793	0.50	0.376	713	22130	31.0
Discrete event simulation library	omnetpp	367	2.10	0.376	290	6250	21.5
Games/path finding	astar	1250	1.00	0.376	470	7020	14.9
XML parsing	xalancbmk	1045	0.70	0.376	275	6900	25.1
Geometric mean	-	-	-	-	-	-	25.7

Float: 6.2




Integer: 6



$$\sqrt[n]{\prod_{i=1}^n \text{Execution Time Ratio}_i}$$



# Power Trends in CMOS IC



$$\text{Power} = \text{Capacitive load} \times \text{Voltage}^2 \times \text{Frequency}$$

×30

5V → 1V

×1000

# Reducing Power


- Suppose a new CPU has
  - 85% of capacitive load of old CPU
  - 15% voltage and 15% frequency reduction

$$\frac{P_{\text{new}}}{P_{\text{old}}} = \frac{(C_{\text{old}} \times 0.85) \times (V_{\text{old}} \times 0.85)^2 \times (F_{\text{old}} \times 0.85)}{C_{\text{old}} \times V_{\text{old}}^2 \times F_{\text{old}}} = 0.85^4 = 0.52$$

- The power wall
  - We can't reduce voltage further
  - We can't remove more heat
- How else can we improve performance?



# Uniprocessor Performance



Constrained by power, instruction-level parallelism, memory latency

# Solution




# Multiprocessing

- Multicore microprocessors
  - More than one processor per chip
- Requires explicit parallel programming
  - Compare with instruction level parallelism
    - Hardware executes multiple instructions at once
    - Hidden from the programmer
  - Hard to do
    - Programming for performance
    - Load balancing
    - Optimizing communication and synchronization



# SPEC Power Benchmark

- Power consumption of server at different workload levels
  - Performance: ssj\_ops/sec
  - Power: Watts (Joules/sec)



Utilization	Speed	Base speed:	1.20 GHz
44%	3.59 GHz	Sockets:	1
Processes	Threads	Cores:	6
227	3623	Logical processors:	8
Up time	Handles	Virtualization:	Enabled
2:23:51:15	105514	L1 cache:	544 KB
		L2 cache:	4.5 MB
		L3 cache:	10.0 MB

$$\text{Overall ssj\_ops per Watt} = \left( \sum_{i=0}^{10} \text{ssj\_ops}_i \right) / \left( \sum_{i=0}^{10} \text{power}_i \right)$$



# SPEC Power Benchmark

## SPECpower\_ssj2008 for Xeon X5650

Target Load %	Performance (ssj_ops)	Average Power (Watts)
100%	865,618	258
90%	786,688	242
80%	698,051	224
70%	607,826	204
60%	521,391	185
50%	436,757	170
40%	345,919	157
30%	262,071	146
20%	176,061	135
10%	86,784	121
0%	0	80
Overall Sum	4,787,166	1,922
$\Sigma\text{ssj\_ops}/\Sigma\text{power} =$		2,490

# Pitfall: Amdahl's Law

- Improving an aspect of a computer and expecting a proportional improvement in overall performance

$$T_{\text{improved}} = \frac{T_{\text{affected}}}{\text{improvement factor}} + T_{\text{unaffected}}$$

$$\frac{100}{\frac{80}{10} + 20} = 3.5$$

- Example: multiply accounts for 80s/100s
  - How much improvement in multiply performance to get

Speed off  5× overall?

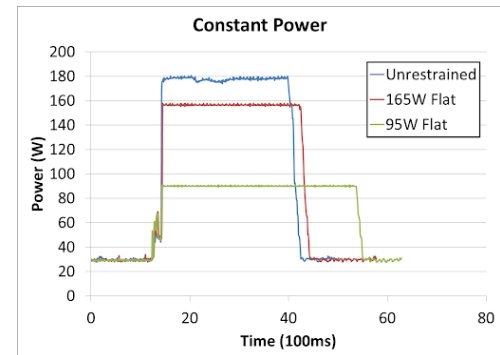

$$20 = \frac{80}{n} + 20 \quad \blacksquare \text{ Can't be done!}$$

- Corollary: make the common case fast



# Fallacy: Low Power at Idle

- Look back at i7 power benchmark
  - At 100% load: 258W
  - At 50% load: 170W (66%)
  - At 10% load: 121W (47%)
- Google data center
  - Mostly operates at 10% – 50% load
  - At 100% load less than 1% of the time
- Consider designing processors to make power proportional to load



# Pitfall: MIPS as a Performance Metric

- MIPS: Millions of Instructions Per Second

- Doesn't account for

- Differences in ISAs between computers
- Differences in complexity between instructions

**P1 100 MIPS    ARM**  
**P2 120 MIPS    Pentium**

$$\begin{aligned} \text{MIPS} &= \frac{\text{Instruction count}}{\text{Execution time} \times 10^6} \\ &= \frac{\text{Instruction count}}{\frac{\text{Instruction count} \times \text{CPI}}{\text{Clock rate}} \times 10^6} = \frac{\text{Clock rate}}{\text{CPI} \times 10^6} \end{aligned}$$

- CPI varies between programs on a given CPU

# Concluding Remark

- Cost/performance is improving
  - Due to underlying technology development
- Hierarchical layers of abstraction
  - In both hardware and software
- Instruction set architecture
  - The hardware/software interface
- Execution time: the best performance measure
- Power is a limiting factor
  - Use parallelism to improve performance

contact us

E-mail :

[Ali.zendeabad@gmail.com](mailto:Ali.zendeabad@gmail.com)

Homepage:

[Sazendeabad.ir](http://Sazendeabad.ir)