

- Bandwidth or throughput
  - Total work done in a given time
  - 10,000-25,000X improvement for processors
  - 300-1200X improvement for memory and disks
- Latency or response time
  - Time between start and completion of an event
  - 30-80X improvement for processors
  - 6-8X improvement for memory and disks

## Performance

Comparing Machines/Systems

- Response Time (latency) =  $T$ 
  - How long does it take for my job to run?
  - How long does it take to execute a job?
  - How long must I wait for the database query?
- Throughput =  $\gamma$ 
  - How many jobs can the machine run at once?
  - What is the average execution rate?
  - How many queries per minute?

avg/best case/worst case

What do we "really" want to know?

--- Which system works best in our larger system?

--- What costs can be traded off?

avg

Time?

- Elapsed Time
  - Counts everything (disk and memory accesses, I/O, etc.)
  - A useful number, but often not good for comparison purposes
    - E.g., OS & multiprogramming time make it difficult to compare CPUs


→ "wall clock"

Depends on load, disk layout, ...

more abstract

- CPU time (CPU = Central Processing Unit = processor)
  - Doesn't count I/O or time spent running other programs
  - Can be broken up into system time, and user time

→ user cpu time



- Our focus: user CPU time
  - Time spent executing the lines of code that are "in" our program
  - Includes arithmetic, memory, and control instructions...

Time CPU used for our job (+ overhead)

unix ⇒ localhost > time myjob

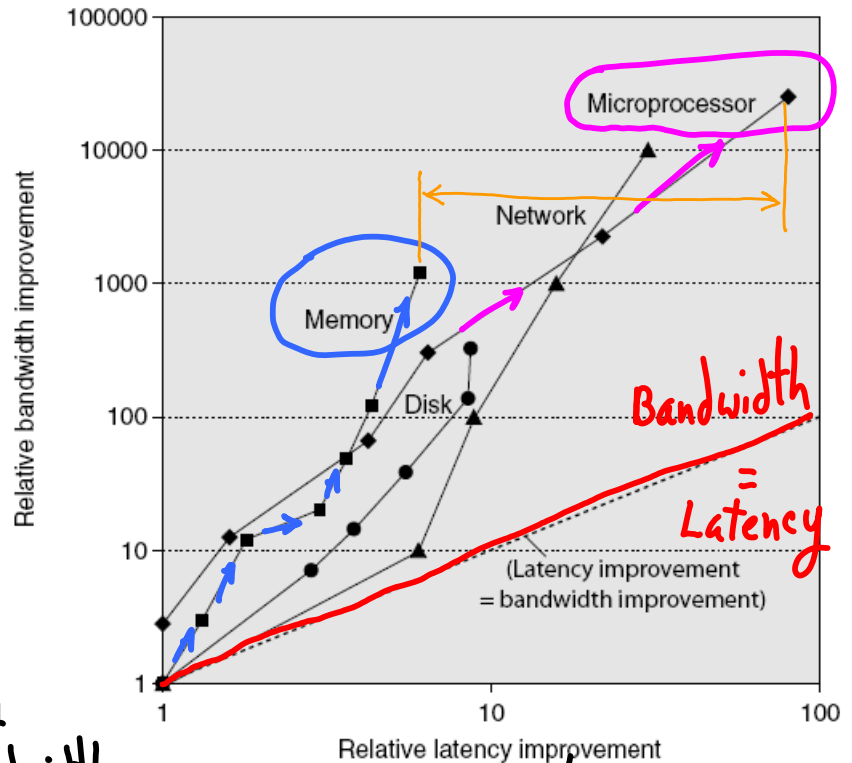


# relative performance

Latency vs. Bandwidth


Increasing complexity/density/speed

1.  $V / L$  gets worse for each component.
2.  $V_{proc} / L_{mem}$  gets worse even faster



↑  
Log Bandwidth  
V

→ Log Latency L



- Dynamic energy
  - Transistor switch from 0 → 1 or 1 → 0
  - $\frac{1}{2} \times \text{Capacitive load} \times \text{Voltage}^2$
- Dynamic power
  - $\frac{1}{2} \times \text{Capacitive load} \times \text{Voltage}^2 \times \text{Frequency switched}$
- Reducing clock rate reduces power, not energy

$$f_{max} = \propto V$$

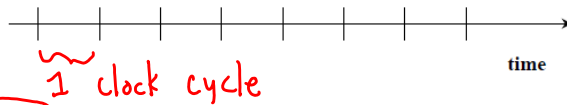
cpu Clock Cycles  $\Rightarrow$  cpu time

- Instead of reporting execution time in seconds, we often use cycles

$$\text{CPU Time} = \left( \frac{\text{seconds}}{\text{program}} \right) = \left( \frac{\text{cycles}}{\text{program}} \right) \times \left( \frac{\text{seconds}}{\text{cycle}} \right)$$

$$T_{\text{cycle}} = \left( \frac{\text{seconds}}{\text{Tick}} \right)$$

- Clock "ticks" indicate when to start activities:



$$\text{Freq} = \left( \frac{\text{Ticks}}{\text{sec}} \right) = \frac{1}{T_{\text{cycle}}}$$

- Cycle time = time between ticks = seconds per cycle
- Clock rate (frequency) = cycles per second (1 Hz. = 1 cycle/sec)

$$2 \text{ GHz clock} \Rightarrow \text{Freq} = \left( \frac{2 \times 10^9 \text{ ticks}}{\text{sec}} \right) \Rightarrow T_{\text{cycle}} = \left( \frac{1 \text{ sec}}{2 \times 10^9 \text{ Ticks}} \right)$$

$$= \frac{1}{2} \text{ ns}$$

$$= \frac{1}{2} (1,000 \text{ ps})$$

$$= 500 \text{ ps}$$

$\uparrow 10^{-12} \text{ sec}$

$$\text{CR} \equiv \text{Freq}$$

- User CPU execution time  $\times$  cycles  $\times$   $T_{\text{cycle}}$

$$\text{Execution Time} = \boxed{\text{Clock Cycles for Program}} \times \boxed{\text{Clock Cycle Time}}$$

rewrite  $\downarrow$

- Since Cycle Time is 1/Clock Rate (or clock frequency)

$$\text{Execution Time} = \frac{\text{Clock Cycles for Program}}{\text{Clock Rate}} = \text{cycles} \left( \frac{1}{\text{Cycle}} \right)^{-1}$$

$$= \frac{\text{cycles}}{\text{CR}}$$

- The program should be something real people care about
    - Desktop: MS office, edit, compile
    - Server: web, e-commerce, database
    - Scientific: physics, weather forecasting
- } real jobs?  
or  
benchmarks?


## Measuring Clock Cycles

- Clock cycles/program is not an intuitive or easily determined value, so

$$\text{Clock Cycles} = \text{Instructions} \times \text{Clock Cycles Per Instruction} = \sum_{i=1}^{\#instr} \#cycles_i$$

*(avg #cycles / instr)*

- Cycles Per Instruction (CPI) used often
- CPI is an **average** since the number of cycles per instruction varies from instruction to instruction
  - Average depends on instruction mix, latency of each inst. type etc.
- CPIs can be used to compare two implementations of the same ISA, but is not useful alone for comparing different ISAs
  - An x86 add is different from a MIPS add



↑  
each instruction is different

lc3:  
ADD vs RTI

$$(\#instr \times CPI) \times \frac{1}{CR}$$


" " "  
#cycles × T<sub>cycle</sub>

Depends on instruction mix, system configuration, data

- Drawing on the previous equation:

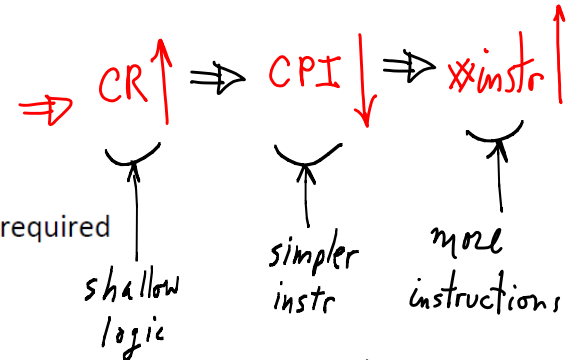
$$\text{Execution Time} = (\text{Instructions} \times \text{CPI}) \times \text{Clock Cycle Time}$$

$$\text{Execution Time} = \frac{\text{Instructions} \times \text{CPI}}{\text{Clock Rate}} = \#instr \left( \frac{CPI}{CR} \right)$$





- To improve performance (i.e., reduce execution time)
  - Increase clock rate (decrease clock cycle time) OR
  - Decrease CPI OR
  - Reduce the number of instructions

Time ↓



- Designers balance cycle time against the number of cycles required
  - Improving one factor may make the other one worse...

[look for sweet spot]



## Clock Rate ≠ Performance

- Mobile Intel Pentium 4      Vs      Intel Pentium M
  - 2.4 GHz
  - P4 is 50% faster?

$$\frac{CR_{P4}}{CR_{PM}} = \frac{2.4}{1.6} = 1.5$$

- Performance on Mobilemark with same memory and disk
  - Word, excel, photoshop, powerpoint, etc.

But - Mobile Pentium 4 is only 15% faster " Instr. Count

- What is the relative CPI? time = (Instr) (CPI/CR)
  - ExecTime = IC • CPI/Clock rate
  - ExecTime<sub>M</sub> = 1.15 ExecTime<sub>4</sub>
  - IC • CPI<sub>M</sub>/1.6 = 1.15 • IC • CPI<sub>4</sub>/2.4
  - CPI<sub>4</sub>/CPI<sub>M</sub> = 2.4/(1.15•1.6) = 1.3

$$T_{P4} = IC_{P4} \left( \frac{CPI_{P4}}{CR_{P4}} \right)$$

$$T_{PM} = IC_{PM} \left( \frac{CPI_{PM}}{CR_{PM}} \right)$$

$$T_{P4} (1.15) = T_{PM}$$

$$\left( IC_{P4} \left( \frac{CPI_{P4}}{CR_{P4}} \right) \right) (1.15) = IC_{PM} \left( \frac{CPI_{PM}}{CR_{PM}} \right)$$

Same ISA

$$IC_{P4} = IC_{PM}$$

and

$$CR_{P4} = (1.5) CR_{PM}$$

$$\frac{CPI_{P4}}{(1.5) CR_{PM}} (1.15) = \frac{CPI_{PM}}{CR_{PM}}$$

↓ rearrange

$$\frac{CPI_{P4}}{CPI_{PM}} = \frac{(1.5)}{(1.15)} = 1.304... \Rightarrow 30\% \text{ more cycles/instr on avg for P4}$$

# Break it down by classes

- Different instruction types require different numbers of cycles
- CPI is often reported for types of instructions

$$\text{Clock Cycles} = \sum_{i=1}^n (\text{CPI}_i \times \text{IC}_i)$$

\* (type i instructions)

(avg CPI for type i instructions)

- where  $\text{CPI}_i$  is the CPI for the type of instructions and  $\text{IC}_i$  is the count of that type of instruction

- To compute the overall average CPI use

$$\frac{(\text{* cycles})}{\text{IC}} \text{CPI} = \sum_{i=1}^n \left( \text{CPI}_i \times \frac{\text{Instruction Count}_i}{\text{Instruction Count}} \right)$$

$$\text{IC} = \sum_{i=1}^n \text{IC}_i$$

$$\frac{1}{\text{IC}} \left( \sum_{i=1}^n (\text{CPI}_i \times \text{IC}_i) \right)$$

$$= \sum_{i=1}^n \text{CPI}_i \times \frac{\text{IC}_i}{\text{IC}}$$

$$= \sum \text{CPI}_i \times \%_i$$

= avg CPI

$\text{CPI}_i \quad \%_i \Rightarrow \left\{ \text{CPI}_i (\%_i) \right\}$   $i^{\text{th}}$  contribution to overall avg CPI

Instruction Type	CPI	Frequency	CPI * Frequency
ALU	1	50%	0.5
Branch	2	20%	0.4
Load	2	20%	0.4
Store	2	10%	0.2

sum = 1.5 = (avg) CPI

- Given this machine, the CPI is the sum of CPI  $\times$  Frequency
- Average CPI is  $0.5 + 0.4 + 0.4 + 0.2 = 1.5$

$$= \sum_i \text{* (cycles for type i)} / \text{IC}$$

$$= \text{* cycles} / \text{IC}$$

$$= \text{avg CPI}$$

What fraction of the time for data transfer?



$$\begin{aligned}
\frac{T_{LD-ST}}{T_{Total}} &= \frac{\text{Cycles}_{LD-ST} * (\frac{1}{f})}{\text{Cycles}_{Total} * (\frac{1}{f})} = \frac{*(\text{cycles LD-ST})}{*(\text{cycles})} = \frac{(*\text{cycles LD}) + (*\text{cycles ST})}{(*\text{cycles})} = IC * CPI \\
&= \left[ \frac{(*\text{cycles LD})}{IC} + \frac{(*\text{cycles ST})}{IC} \right] \frac{1}{CPI} \\
&= \left[ \frac{CPI_{LD} * IC_{LD}}{IC} + \frac{CPI_{ST} * IC_{ST}}{IC} \right] \frac{1}{CPI} \\
&= [CPI_{LD} * \%_{LD} + CPI_{ST} * \%_{ST}] (\frac{1}{1.5}) \\
&= [0.4 + 0.2] (\frac{1}{1.5}) = 40\%
\end{aligned}$$

↑  
cycle Time

## Speedup

- Speedup allows us to compare different CPUs or optimizations

$$Speedup = \frac{CPUtimeOld}{CPUtimeNew}$$

- Example
  - Original CPU takes 2sec to run a program
  - New CPU takes 1.5sec to run a program
  - Speedup = 1.333 or speedup or 33%

$$\begin{aligned}
&\longrightarrow T_{old} = 2s \\
&T_{new} = 1.5s
\end{aligned}$$

$$S_{new-old} = \frac{T_{old}}{T_{new}} = \frac{2}{1.5} = \frac{20}{15} = 1 + \frac{5}{15} = 1.333...$$

$$\Rightarrow \frac{T_{old}}{1.3} = T_{new} \quad (?) \quad S' = \frac{1}{1.3} \quad (?)$$

Does this look like speedup? What do we mean by speedup?

$$\begin{aligned}
S' &= \frac{V_{new}}{V_{old}} = \frac{(W_{new}/T_{new})}{(W_{old}/T_{old})} = T_{old}/T_{new} \Rightarrow V_{new} = V_{old} (1.3) \\
&\Rightarrow \text{new is } 30\% \text{ faster}
\end{aligned}$$



Assuming  $W_{old} = W_{new}$

$$W = W_{parallel} + W_{sequential} = f \cdot W + (1-f)W$$

$$S = \frac{q_{old}}{q_{new}} = \frac{W/T_{old}}{W/T_{new}}$$

$$T = T_{parallel} + T_{sequential} \quad \begin{matrix} \text{parallel} = \text{improvable} \\ \text{sequential} = \text{fixed} \end{matrix}$$

$$q_p = W_p/T_p \quad q_s = W_s/T_s$$

$$q_p = q_s \quad \underline{\text{old}}$$

$$q_p = a q_s \quad \underline{\text{new}} \quad (\text{ie. } S'_p = a)$$

### Amdahl's Law

- If an optimization improves a fraction  $f$  of execution time by a factor of  $a$

$$\text{Speedup} = \frac{T_{old}}{[(1-f) + f/a] \cdot T_{old}} = \frac{1}{(1-f) + f/a}$$

Is  $f$  fixed?  $W_p = O(n)$   $W_s = O(1)$

$$T_{old} = \frac{f \cdot W}{q_s} + \frac{(1-f)W}{q_s} = W/q_s$$

- This formula is known as Amdahl's Law

- Lessons from

- If  $f \rightarrow 100\%$ , then speedup =  $a$
- If  $a \rightarrow \infty$ , the speedup =  $1/(1-f)$

- Summary

- Make the common case fast
- Watch out for the non-optimized component

target speed up to max %

$$T_{new} = \frac{W_p}{q_p} + \frac{W_s}{q_s}$$

$$= \frac{f \cdot W}{a q_s} + \frac{(1-f) \cdot W}{q_s}$$

$$= \left( \frac{f}{a} + (1-f) \right) W/q_s$$

$$S_{\text{new-old}} = \frac{T_{old}}{T_{new}} = \frac{(W/q_s)}{(W/q_s) \left( \frac{f}{a} + (1-f) \right)}$$

$$= \frac{1}{\left( \frac{f}{a} + (1-f) \right)} \xrightarrow{f=1} \frac{1}{\left( \frac{1}{a} + 0 \right)} = a = S'_{\text{all-improved}} = S'_{\text{max}}$$

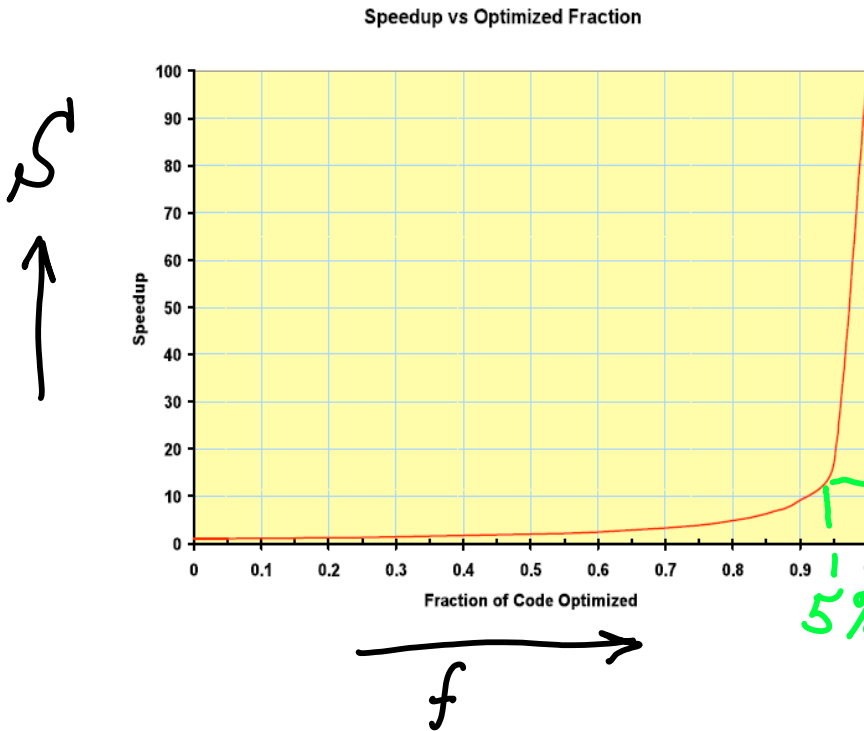
$$\xrightarrow{f=0} \frac{1}{\left( 0/a + 1 \right)} = 1 = S'_{\text{none-improved}} = S'_{\text{min}}$$

$$\xrightarrow{a \rightarrow \infty} \frac{1}{\left( \frac{f}{\infty} + (1-f) \right)} = \frac{1}{(1-f)} = S'_{\infty} \geq S'_{\text{actual}}$$

$f = 50\% \Rightarrow S'_{\infty} = \frac{1}{(1-1/2)} = 2$  max improvement for  $a \rightarrow \infty$  is twice as fast

- If  $a=100$ , what is the overall speedup as a function of  $f$ ?

e.g., use 100 CPUs  
for  $W_{parallel}$



$S = 17?$   
5% sequential  
 $cost_{new} = 100 \cdot cost_{old}$

## Amdahl's Law Example

- Suppose a program runs in 100 seconds on a machine, with multiply responsible for 80 seconds of this time. How much do we have to improve the speed of multiplication if we want the program to run 4 times faster?  $= S = 4$

$$T_{old} = T_{other} + T_{mult} = 20s + 80s = 100s$$

find  $S_p = \frac{T_{p,old}}{T_{p,new}} = ?$

$$S_p = \frac{T_{p,old}}{T_{p,new}} = \frac{80}{5} = 16$$

- How about making it 5 times faster?  $S = 5?$


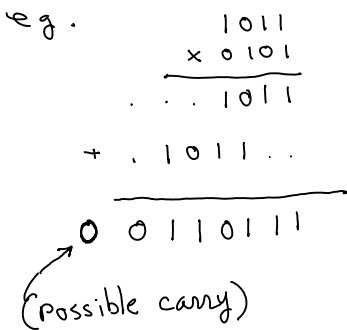

$$x \cdot 7 = x(4+2+1) = 4x + 2x + x$$

←  $y = 00\dots0111$ 
← 2 shifts (C)
 ← 1 shift (B)
 ← 0 shifts (A)

$$\begin{array}{r}
 00\dots00 = S^0 \\
 + \begin{array}{r} x_n x_{n-1} \dots x_1 x_0 \\ \hline S'_n S'_{n-1} \dots S'_1 S'_0 \end{array} = S^1 \\
 + \begin{array}{r} \text{shift} \\ x_n x_{n-1} \dots x_1 x_0 0 \\ \hline S''_{n+1} S''_n \dots S''_2 S''_1 S''_0 \end{array} = S^2 \\
 + \begin{array}{r} \text{shift} \\ x_n x_{n-1} \dots x_1 x_0 00 \\ \hline S'''_{n+2} S'''_{n+1} \dots S'''_3 S'''_2 S'''_1 S'''_0 \end{array} = S^3
 \end{array}$$

What if y has a 0 bit? Then add 0 instead of shifted x: e.g., y = 0...101 add 0, not B.

INT MULTIPLY:  $S = x * y$   
 LSR: partial products, initially x.  
 S: partial sum, initially 0.  
 RSR: initially y.  
 Z: all 0s



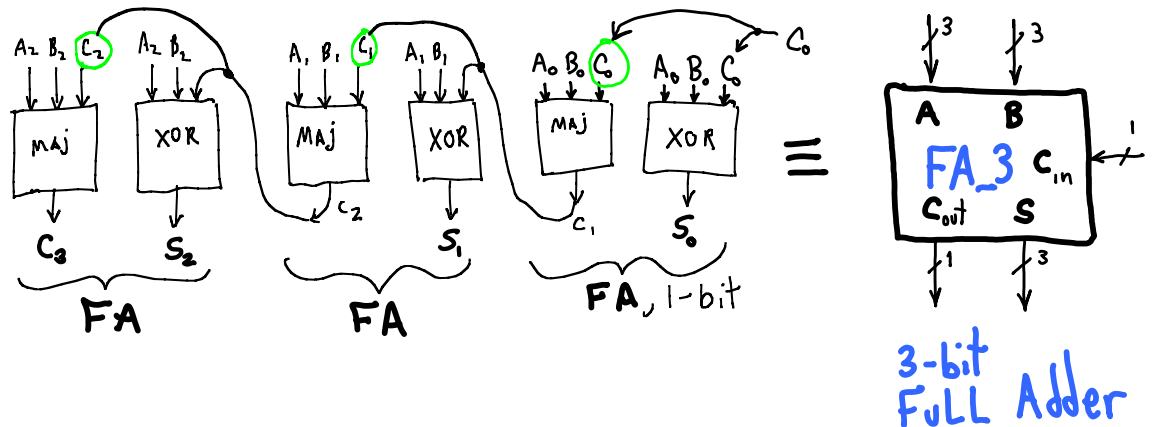
⇒ We shift left (1011) every time, but add either the shifted (1011) or all zeroes, depending on whether  $y_i$  is 1 or 0.

### ADDER

Delay is longest path.


2 levels per MAJ, n bit operands.

2n gate delays until result is ready.






3-bit FULL Adder

Shift Register Delay = 1 (all stages in Parallel)





Overall delay:  $(n \text{ shifts}) \times (2n \text{ delay}_{\text{ADDER}}) = 2n^2$

Improvement? Recursive Refinement?



(Use Recursively)



$\Rightarrow$  longest delay,  $f(n)$ , is,  
 $4n$ -bit Add +  $(n/2)$ -bit MULT delay:  
 $f(n) = 4(2n) + f(n/2)$

$$f(n) = 8n + f(n/2) = 8n + 8(n/2) + f(n/4) \Rightarrow 16n (\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots) \approx 16n$$

time improvement:  $2n^2 \Rightarrow 16n$  Area increase:  $7n \Rightarrow ?$   $A(n) = 4A(n/2) + 4(n)$

32-bit  $2(2^5)^2 = 2^{11} \Rightarrow 2^7(2^5) = 2^9$   $S' = 2^{11}/2^9 = 2^2 = 4$   $S'_{64-bit} = 8$

OUR Requirement

$$4 = S'_{overall} = \frac{q_{new}}{q_{old}} = \frac{W/T_{new}}{W/T_{old}} = \frac{T_{old}}{T_{new}} = \frac{W_s/V_{s-old} + W_p/V_{p-old}}{W_s/V_{s-new} + W_p/V_{p-new}}$$

$$= \frac{(0.2)W/V_{p-old} + (0.8)W/V_{p-old}}{(0.2)W/V_{p-old} + (0.8)W/S'_p \cdot V_{p-old}}$$

$$= \frac{1}{(0.2 + 0.8/S'_p)} = 4$$

Assume

$$V_{p-old} = V_{s-old} = V_{s-new}$$

$$V_{p-new} = S'_p V_{p-old}$$

$$\Rightarrow 0.25 = 0.2 + 0.8/S'_p$$

$$\Rightarrow 0.05 = 0.8/S'_p$$

$$\Rightarrow S'_p = \frac{0.8}{0.05} = \frac{80}{5} = 16 = S'_{128-MULT}$$


OK!

We can just make it if our data is 128-bit

$$S'_{overall} = 5? \quad 1/(0.2 + 0.8/S'_p) \stackrel{?}{=} 5 \quad \text{Can we?}$$

$$S'_{overall-\infty} = 1/(0.2 + 0.8/\infty) = 1/0.2 = 5 \quad \text{Possible? if we eliminate MULT time?}$$

what else could we try?



n-stage pipeline

1 result every shift

$$\text{delay} = 4n \quad (2n\text{-bit ADD})$$

(32-bit):


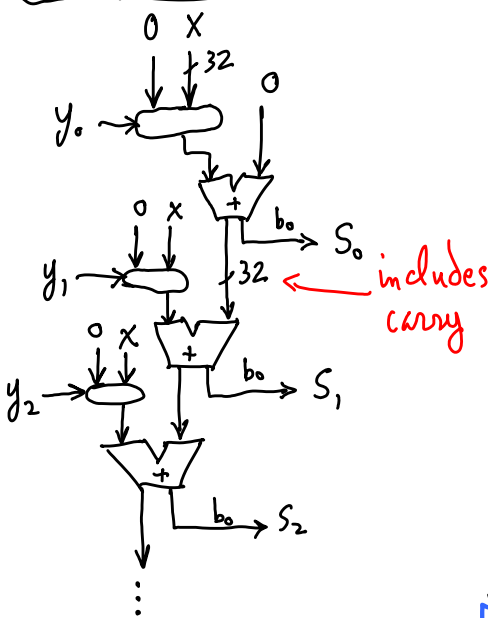
$$S'_{32-MULT} = \frac{2(2^5)^2}{4(2^5)} = 2^4 = 16$$

Can we keep it full?

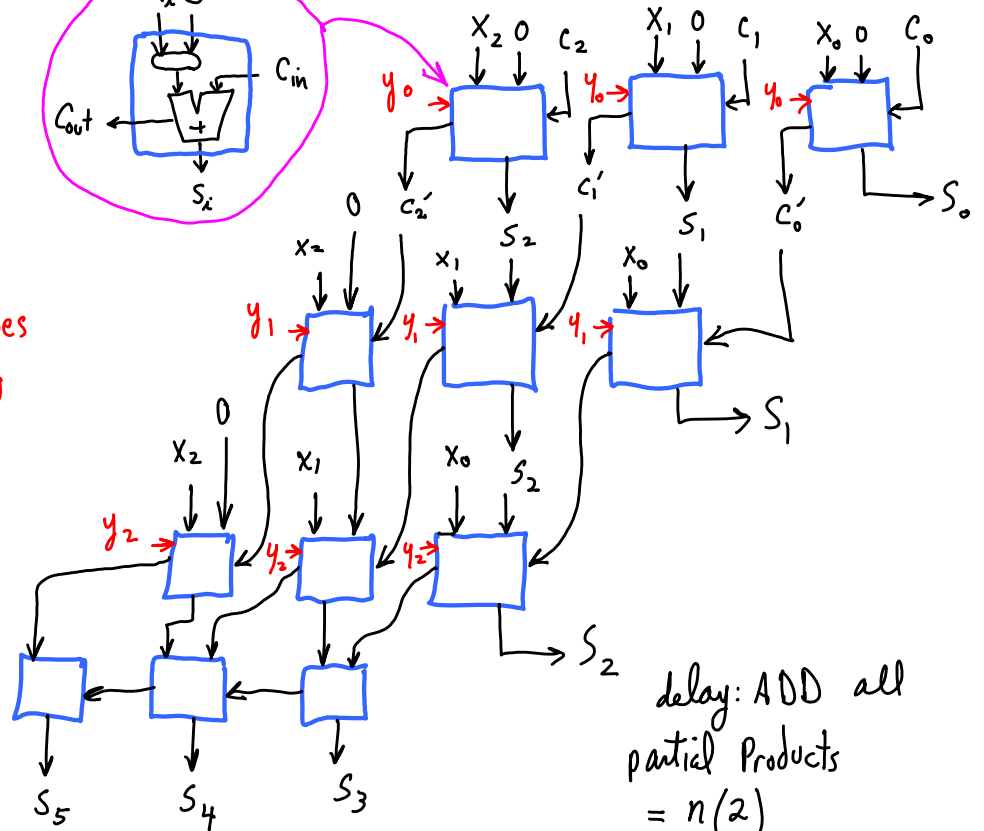
$$\text{Latency} = 4n \times (n) = 4n^2$$

# Another Approach? Back to square 1

## Adding Partial Products




## carry-save adder




delay: ADD all partial products =  $n(2)$

$$P_{32-MULT} = 2^{11} / 2^6 = 32$$

## Other Ideas?



How fast?




delay?  
area?

$$S = 1 / (0.2 + 0.8/32)$$

$$= 1 / ((2/10) + (2^3/10)/2^5)$$

$$= 10 / (2 + 1/4) = 40 / (8+1) = 40/9 = 4 + 4/9$$



Do not arith in parallel, combine results How?

Combine

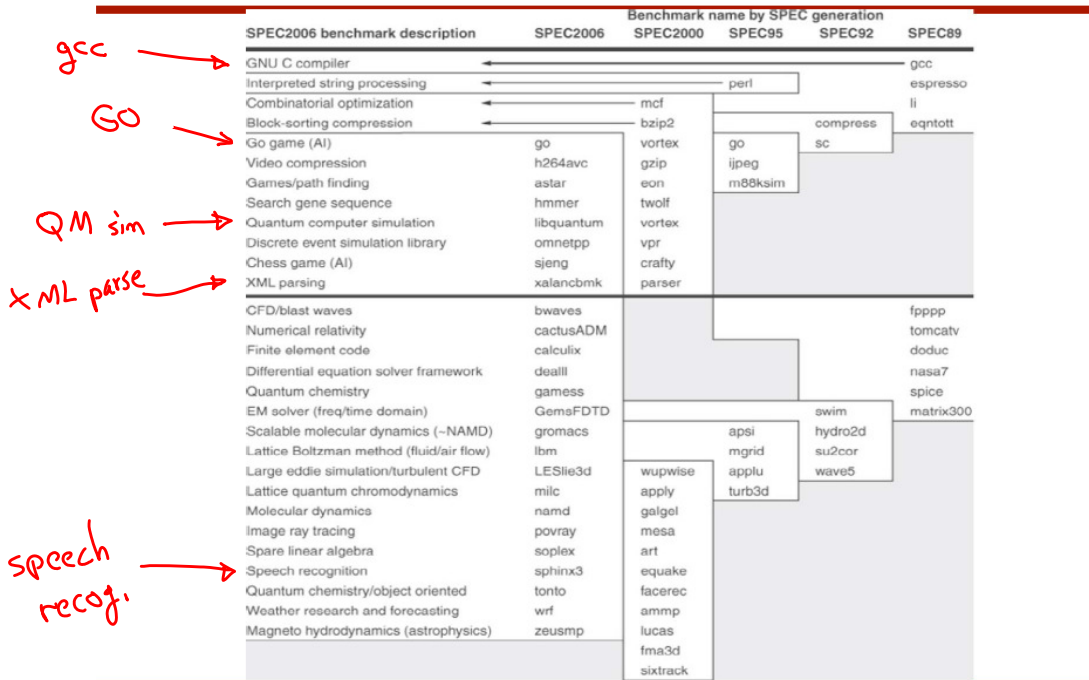
# Evaluating Performance

- Performance best determined by **running a real application**
  - Use programs **typical** of expected workload
    - e.g., compilers/editors, scientific applications, graphics, etc.
- Microbenchmarks
  - **Small programs**, synthetic or kernels from larger applications
  - Nice for architects and designers
  - **Can be misleading**
- Benchmarks
  - Collection of **real programs** that companies have agreed on
  - Components: **programs**, **inputs & outputs**, **measurements**, **rules**, **metrics**
  - Can still be abused

*in typical environment?*

*⇒ Build compiler optimized for benchmark?  
 ⇒ "Buggy" ⇒ skips work?*

## The SPEC CPU Benchmark Suite (System Performance Evaluation Cooperative)



## Other Benchmarks

---

- Scientific computing: Linpack, SpecOMP, SpecHPC, ...
- Embedded benchmarks: EEMBC, Dhrystone, ...
- Enterprise computing
  - TCP-C, TPC-W, TPC-H
  - SpecJbb, SpecSFS, SpecMail, Streams,
- Other
  - 3Dmark, ScienceMark, Winstone, iBench, AquaMark, ...
- Watch out: your results will be as good as your benchmarks
  - Make sure you know what the benchmark is designed to measure
  - Performance is not the only metric for computing systems
    - Cost, power consumption, reliability, real-time performance, ...

## Summarizing Performance

---

- Combining results from multiple programs into 1 benchmark score
  - Sometimes misleading, always controversial...and inevitable
  - We all like quoting a single number

$$AM = \frac{1}{n} \sum_{i=1}^n (Weight_i) \cdot Time_i$$

- 3 types of means
  - Arithmetic: for times
  - Harmonic: for rates
  - Geometric: for ratios

$$HM = \frac{\cancel{n}}{\sum_{i=1}^n \frac{(Weight_i)}{Rate_i}}$$

find ratio  $\bar{r}$  s.t.  
 $r_1 \cdot r_2 \cdots r_n = (\bar{r})^n$

$$GM = \left( \prod_{i=1}^n Ratio_i \right)^{\left(\frac{1}{n}\right)}$$



$$R \Rightarrow (T_{R_1}, 10T_{R_1})$$

$$S_{A-R_1} = \frac{T_{R_1}}{100} \quad S_{A-R_2} = \frac{10T_{R_1}}{4}$$

$$S_{B-R_1} = \frac{T_{R_1}}{200} \quad S_{B-R_2} = \frac{10T_{R_1}}{1}$$

We normalize by using a reference machine R to get the speedups w.r.t. benchmark-1 and benchmark-2 (b-1, b-2).

Suppose R's time on b-2 is 10 times its time for b-1. (We can always express R's times in terms of one of its benchmark times, no matter how many b's there are.)

To combine our speedups w.r.t R, let's try getting a mean of these for A, and a mean for B, then taking the ratio of those mean speedups.

$$\begin{aligned} \bar{S}_{A-R} &= \frac{1}{2} \left( \frac{T_{R_1}}{100} + \frac{10T_{R_1}}{4} \right) \Rightarrow \frac{\bar{S}_{A-R}}{\bar{S}_{B-R}} = \frac{T_{R_1}/2 \left( \frac{4+1000}{400} \right)}{T_{R_1}/2 \left( \frac{1+2000}{200} \right)} = \left( \frac{1}{2} \right) \left( \frac{1004}{2001} \right) \approx \frac{1}{4} ? \\ \bar{S}_{B-R} &= \frac{1}{2} \left( \frac{T_{R_1}}{200} + \frac{10T_{R_1}}{1} \right) \end{aligned}$$

$\left( \frac{4+100r}{1+200r} \right) \begin{matrix} \rightarrow \frac{1}{2} & r \rightarrow \infty \\ \rightarrow 4 & r \rightarrow 0 \end{matrix}$

This doesn't seem to have worked very well. The effect of R can make all the difference: if R's time on b-2 had been one-thousandth of its b-1 time, the result would have been an overall speedup of 2. Let's try a different mean, the geometric mean.

$$\bar{S}_{A-R} = G(S_{A-R_1}, S_{A-R_2}) = \sqrt{\left( \frac{T_{R_1}}{100} \right) \left( \frac{10T_{R_1}}{4} \right)} = \sqrt{\frac{10T_{R_1}^2}{400}} = \frac{(\sqrt{10} T_{R_1})}{2 \cdot 10}$$

$$\bar{S}_{B-R} = G(S_{B-R_1}, S_{B-R_2}) = \sqrt{\left( \frac{T_{R_1}}{200} \right) \left( \frac{10T_{R_1}}{1} \right)} = \sqrt{\frac{10T_{R_1}^2}{200}} = \frac{(\sqrt{10} T_{R_1})}{\sqrt{2} \cdot 10}$$

$$\frac{\bar{S}_{A-R}}{\bar{S}_{B-R}} = \frac{(\sqrt{10} T_{R_1})}{(\sqrt{10} T_{R_1})} \left( \frac{T_{B_1} \cdot T_{B_2}}{T_{A_1} \cdot T_{A_2}} \right)^{\frac{1}{2}} \approx 0.7 \quad \left[ \text{recall } S_{A-B_1} = 2, S_{A-B_2} = \frac{1}{4} \right]$$

weights on R's times cancel.

--- (A's point of view):  $S_{avg}$  in  $[2, 0.25]$   $\implies$  30% slowdown

--- (B's point of view):  $S_{avg}$  in  $[0.5, 4]$   $\implies$  40% speedup

at least its consistent

The real world? Suppose job mix = ( $n_1$  runs of b-1) + ( $n_2$  runs of b-2)

$$S_{A-B} = \frac{n_1 T_{B_1} + n_2 T_{B_2}}{n_1 T_{A_1} + n_2 T_{A_2}} = \frac{200n_1 + n_2}{100n_1 + 4n_2} = \frac{200 + a}{100 + 4a} \quad \begin{matrix} \rightarrow \frac{1}{4}, a \rightarrow \infty \\ \rightarrow 2, a \rightarrow 0 \end{matrix}$$

( $a = n_2/n_1$ )

Sanity check: Given our result above, what **a** does GM appear to assume?

$$\frac{200 + a}{100 + 4a} \approx 3/4 \Rightarrow (200 + a)4 = (100 + 4a)3 \Rightarrow 500 = 8a$$

$$a \approx 62$$

$$\eta_2 = 62 \eta_1$$

For every short job (b1)  
we do 62 long jobs (b2)?

What if we hadn't taken the SQRT in GM?

$$\bar{S}_{A-B} = \frac{1}{2} \Rightarrow (200 + a)2 = (100 + 4a) \Rightarrow a = 50$$

$$HM = \frac{1}{\sum r_i} \text{ where } r_1 = \frac{W_1}{T_1}, r_2 = \frac{W_2}{T_2} \dots \quad [r_i = q_i]$$

[find  $\bar{r}$  s.t., if we did all work at same rate, takes same time]

$$\frac{(W_1 + W_2 + \dots + W_n)}{\bar{r}} = (T_1 + T_2 + \dots + T_n)$$

$$\bar{r} = \frac{W}{\left(\frac{W_1}{r_1}\right) + \left(\frac{W_2}{r_2}\right) + \dots + \left(\frac{W_n}{r_n}\right)}$$

$$= \frac{1}{\sum \frac{(W_i/W)}{r_i}} \quad (W_i/W) = \omega_i$$

We have reference Times:

$$T_{R_1} = \frac{W_1}{q_{R-1}} \quad T_{R_2} = \frac{W_2}{q_{R-2}}$$

[assume  $q_{R-1} = q_{R-2} = q_R$ ]

$$W = \sum W_i$$

$$W_1 = T_{R_1} q_R \quad W_2 = T_{R_2} q_R$$

Use that to get weights:

$$\omega_i = \frac{W_i}{\sum W_i}$$

$$\omega_1 = \frac{W_1}{W_1 + W_2} = \frac{T_{R_1} q_R}{(T_{R_1} q_R + T_{R_2} q_R)} = \frac{T_{R_1}}{T_{R_1} + T_{R_2}}$$

$$= \frac{1}{(1+10)} = 1/11$$

$$\omega_2 = \frac{T_{R_2}}{T_{R_1} + T_{R_2}} = \frac{10}{1+10} = 10/11$$

$$\bar{v}_A = \frac{1}{\sum \omega_i / v_{A-i}}$$

Given our assumption that  $v_{R-1} = v_{R-2} = v_R$

$$v_{A-1} = \frac{W_1}{T_{A-1}} = \frac{W_1}{100}$$

$$v_{A-2} = \frac{W_2}{T_{A-2}} = \frac{10W_1}{4}$$

$$\begin{aligned} W_1 &= T_{R-1} v_R = \frac{T_{R-1}}{T_{R-2}} \\ W_2 &= T_{R-2} v_R = \frac{T_{R-1}}{10 T_{R-1}} \\ &= \frac{1}{10} \end{aligned}$$

or  $W_2 = 10W_1$

$$\bar{v}_A = \frac{1}{\left( \frac{(1/11)}{W_1/100} + \frac{(10/11)}{10W_1/4} \right)} = \frac{(11)}{\frac{100}{W_1} + \frac{10 \cdot 4}{10W_1}} = \frac{11W_1}{(100+4)}$$

$\uparrow$   $T_{A-1}$      $\uparrow$   $T_{A-2}$

$$\bar{v}_B = \frac{11W_1}{(200+1)}$$

$$\rightarrow \int_{A-B}^{HM} = \frac{\bar{v}_A}{\bar{v}_B} = \frac{(200+1)}{(100+4)} \approx 2$$

Suppose, again,  $n_1$  runs of  $b-1$  and  $n_2$  runs of  $b_2$ , w/  $a = n_2/n_1$ .  
Comparing our "real world" performance, what  $a$  does HM imply?

$$\int_{A-B}^{real} = \left( \frac{200+a}{100+4a} \right) = \int_{A-B}^{HM} = \frac{2001}{104}$$

$$\Rightarrow 104(200+a) = 201(100+4a)$$

$$\Rightarrow 20800 + 104a = 20100 + 804a$$

$$\Rightarrow 700 = 700a \quad \boxed{a=1}$$

# Principles of Computer Design

---

- Take Advantage of Parallelism
  - e.g. multiple processors, disks, memory banks, pipelining, multiple functional units
- Principle of Locality
  - Reuse of data and instructions
- Focus on the Common Case
  - Amdahl's Law

---

$$\text{Execution time}_{\text{new}} = \text{Execution time}_{\text{old}} \times \left( (1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}} \right)$$

---

- 1.02 #bytes per frame, time per file (cache, DRAM, ... )
- 1.03 avg CPI, CR, performance
- 1.04-05 CPI by class, CR, instr. mix,
- 1.06 compilers, avg CPI, CR, speedup, CPI by class, peak performance versus
- 1.07 Voltage scaling laws, C, power, GM, %change,
- 1.08 dynamic power, C, V
- 1.09 static and dynamic power, voltage dependence
- 1.10 multi-cores, #instructions, CPIs, execution time, power
- 1.11 die yield and cost
- 1.12 SPEC ratio from times
- 1.13 Faster clock, change ISA ==> fewer instructions executed, CPI vs CR
- 1.14 Performance measured by MFLOPS or MIPS versus overall
- 1.15 Amdahl's Law (improving only a fraction)
- 1.16 Speedup w/ communication costs

