- 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?

avg/best case/worst case

- Throughput = ν
  - How many jobs can the machine run at once?
  - What is the average execution rate?
  - How many queries per minute?

avg

What do we "really" want to know?

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

--- What costs can be traded off?

Time?

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

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

OS
I/O

- Our focus: user CPU time → Time CPU used for our job (+ overhead)
  - Time spent executing the lines of code that are "in" our program
  - Includes arithmetic, memory, and control instructions...

unix ⇒ localhost > time myjob

Latency vs. Bandwith

Increasing complexity/density/speed

1. V / L gets worse for each component.

2. Vproc / Lmem gets worse even faster

relative performance

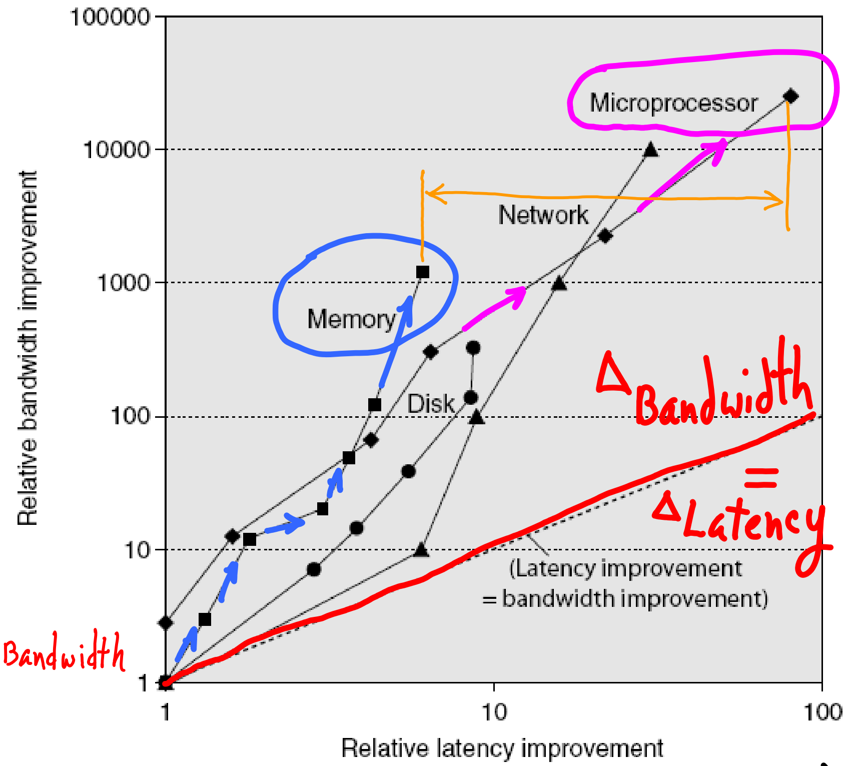

$$\log \frac{V_{new}}{V_{old}} = \Delta_{Bandwidth}$$

$$\log L_{old}/L_{new} = \Delta_{Latency}$$

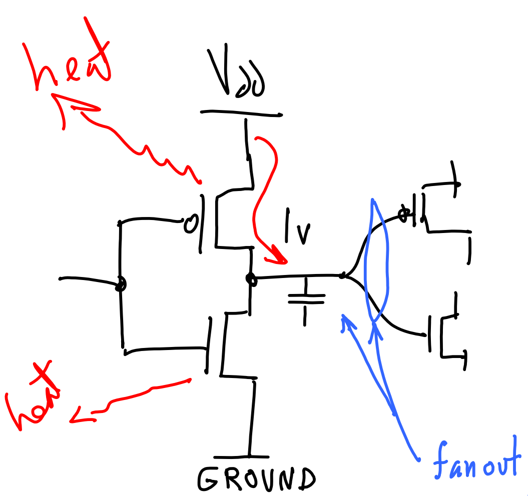

- Dynamic energy
  - Transistor switch from 0 -> 1 or 1 -> 0
  - ½ x Capacitive load x Voltage^2
- Dynamic power
  - ½ x Capacitive load x Voltage^2 x Frequency switched
- Reducing clock rate reduces power, not energy

= capacitive load, changing capacitance up to V_{dd} or draining to GROUND.

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

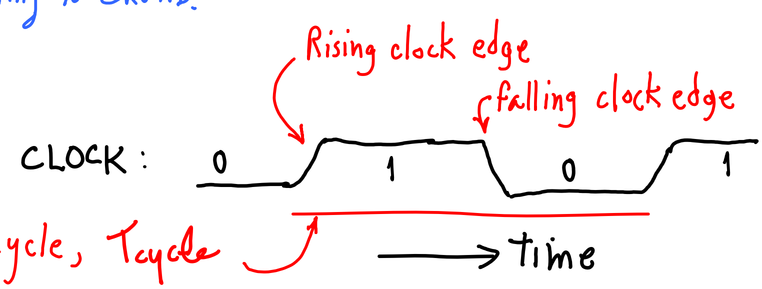

cpu Clock Cycles ⇒ cpu time

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

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

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

- Clock "ticks" indicate when to start activities:

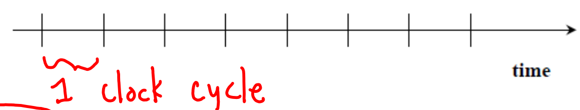

$$\text{Freq} = \left(\frac{\text{Ticks}}{\text{sec}}\right) = 1/T_{cycle}$$

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

$$2\text{ GHz clock} \Rightarrow \text{Freq} = \left(\frac{2\times10^9 \text{ ticks}}{\text{sec}}\right) \Rightarrow T_{cycle} = \left(\frac{1 \text{ sec}}{2\times10^9 \text{ Ticks}}\right)$$

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

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

$$= 500\text{ ps}$$

10^{-12} sec

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

---

rewrite

- User CPU execution time

\# cycles × T_{cycle}

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

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

$$\textit{Execution Time} = \frac{\textit{Clock Cycles for Program}}{\textit{Clock Rate}} = \#\text{cycles}\left(\frac{1}{T_{cycle}}\right)^{-1} = \frac{\#\text{cycles}}{CR}$$

- 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

$$\#\ \text{Clock Cycles} = \text{Instructions} \times \overline{\text{Clock Cycles Per Instruction}}^{\text{average}} = \#\text{instr} \times \left(\frac{\#\text{cycles}}{\text{instr}}\right) = \sum_{i=1}^{\#\text{instr}} \#\text{cycles}_i$$

each instruction is different

[lc3: ADD vs RTI]

- 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

$$CPI_{x86} \neq CPI_{mips}$$

$$\underbrace{(\#\text{instr} \times \overline{CPI})}_{} \times \underbrace{(1/CR)}_{}$$

$$\text{Time} = \#\text{cycles} \times T_{cycle}$$

Depends on
--- instruction mix,
--- system configuration,
--- data

- Drawing on the previous equation:

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

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

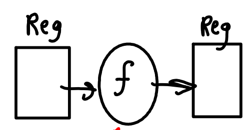

longest path limits CR

reduce Time ↓

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

$\Rightarrow$ CR ↑ $\Rightarrow$ need CPI ↓ $\Rightarrow$ but #instr ↑

shallow logic — simpler instr — more instructions

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

[look for sweet spot]

E.G. Decode time

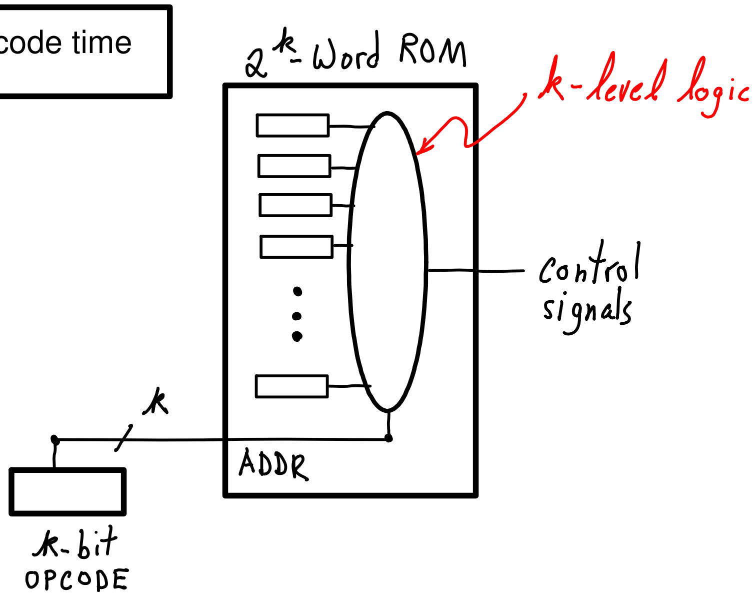

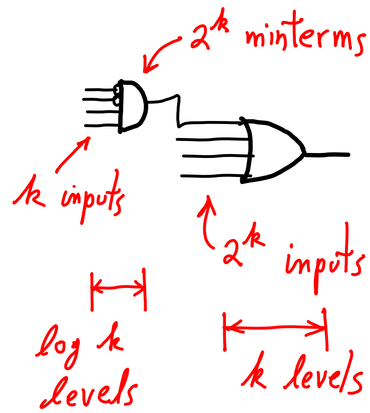

k small $\Rightarrow$ few instructions

Area under curve = Time

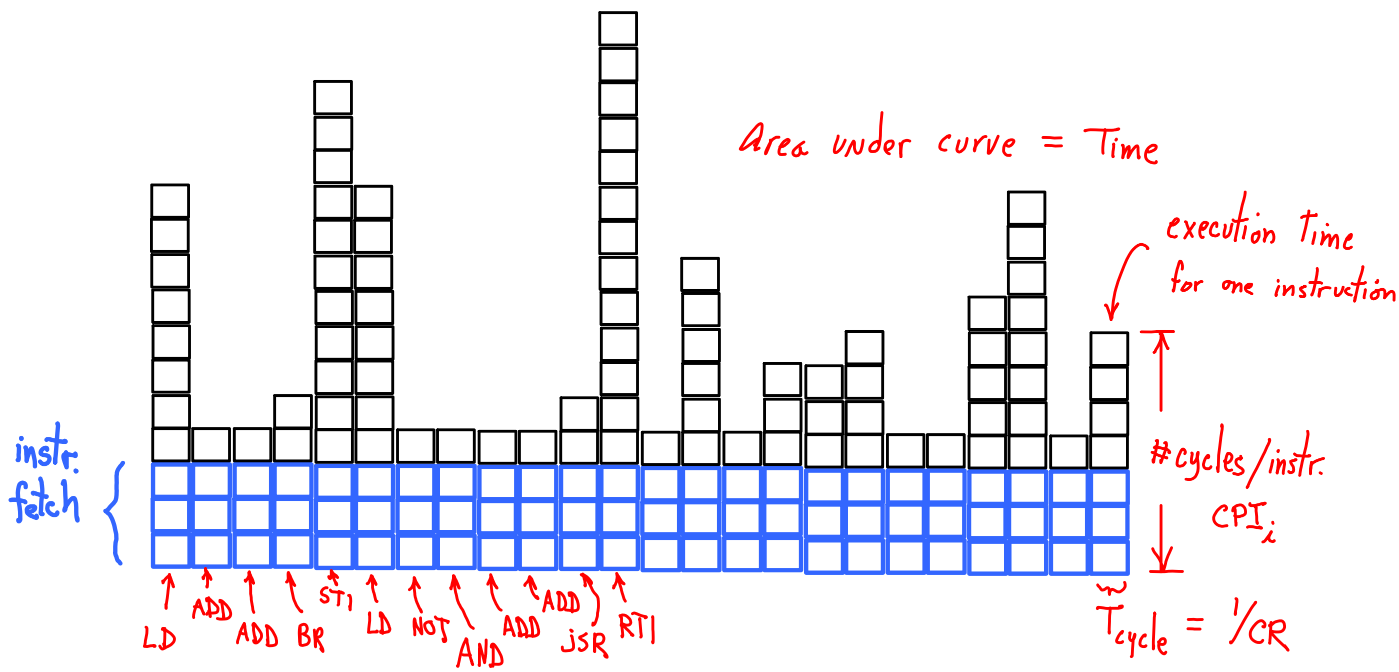

execution $\longrightarrow$ (# instructions) $\times T_{cycle}$

IC

$$\text{Time} = \text{Area} = \sum_{i}^{IC} CPI_i \times T_{cycle} = IC\ \overline{CPI}\ \left(\frac{1}{CR}\right)$$

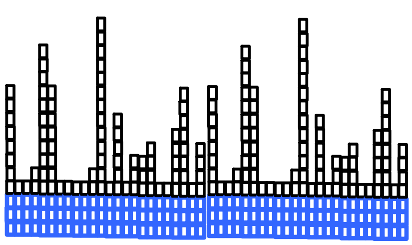

shorter T_{cycle}, more instructions

and/or

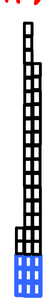

shorter cycle, more CPI_i

Instruction mix matters!

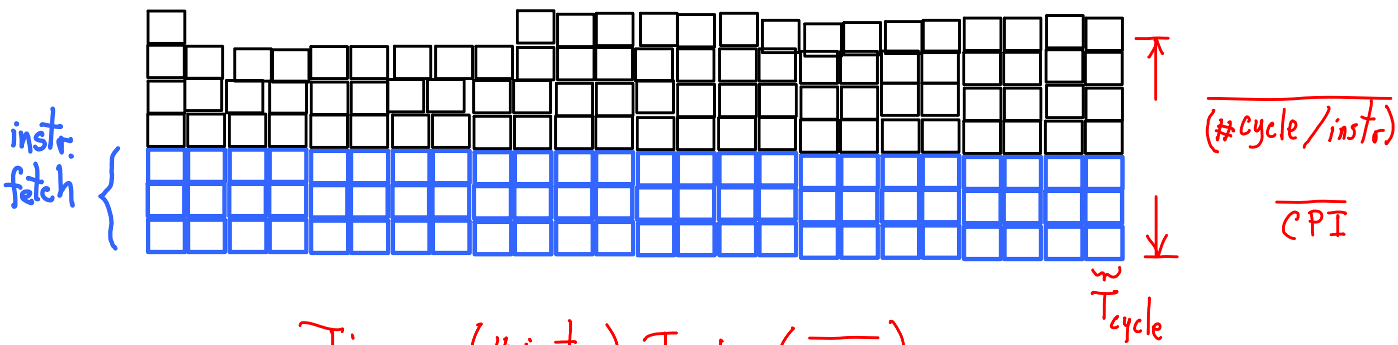

$$\text{Time} = (\#\,\text{instr}) \times T_{cycle} \times (\overline{CPI})$$

$$= IC\,(1/CR)(\overline{CPI})$$

average by classes

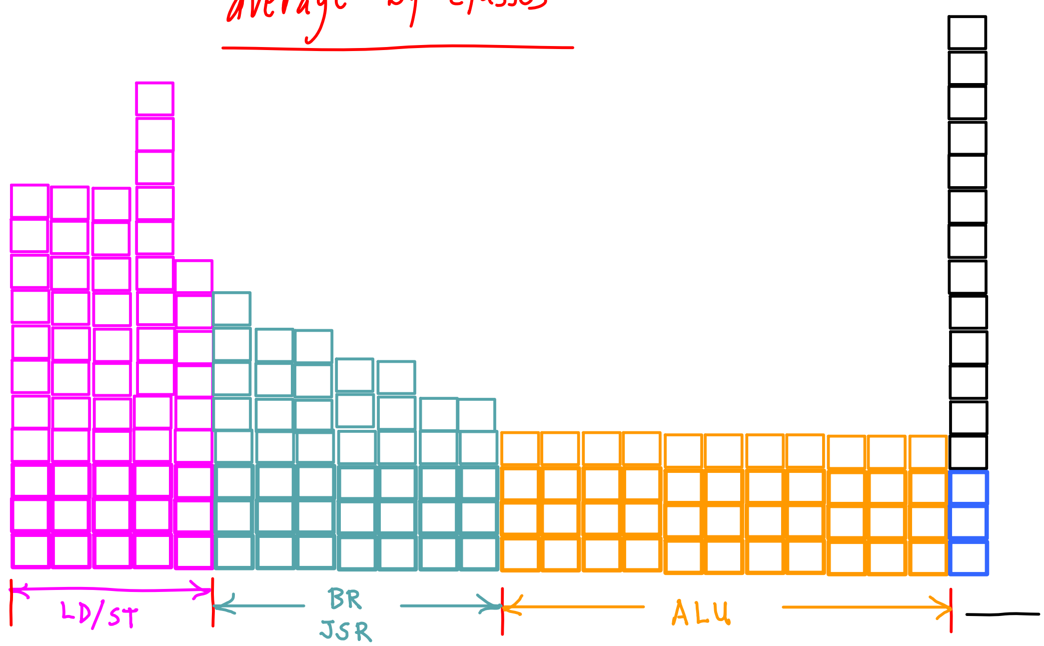

$$\text{time} = \left[(IC_1 \times \overline{CPI}_1) + (IC_2 \times \overline{CPI}_2) + (IC_3 \times \overline{CPI}_3) + (\text{OTHER}) \right] (1/CR)$$

--- Get averages by running batches of class?

--- Guessing from architecture?

Clock Rate ≠ Performance

- Mobile Intel Pentium 4 Vs Intel Pentium M
  - 2.4 GHz 1.6 GHz
  - P4 is 50% faster?
- Performance on Mobilemark with same memory and disk
  - Word, excel, photoshop, powerpoint, etc.
  - But – Mobile Pentium 4 is only 15% faster
- What is the relative CPI?
  - ExecTime = IC • CPI/Clock rate
  - $ExecTime_M = 1.15\ ExecTime_4$
  - $IC \bullet CPI_M/1.6 = 1.15 \bullet IC \bullet CPI_4/2.4$
  - $CPI_4/CPI_M = 2.4/(1.15 \bullet 1.6) = 1.3$

$$CR_{P4} = \left(\frac{2.4}{1.6} = 1.5\right) CR_{PM}$$

$$\Rightarrow T_{PM} = (1.5)\, T_{P4}\ ?$$

But, $T_{PM} = (1.15)\, T_{P4}\ !$

Is the difference $\overline{CPI}_{PM}$ versus $\overline{CPI}_{P4}$?

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

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

Same ISA

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

$$\frac{CPI_{P4}}{CPI_{PM}} = \frac{(1.5)}{(1.15)} = 1.304\ldots$$

$$CPI_{P4} = 1.304\ CPI_{PM}$$

$\Rightarrow$ 30% more cycles/instr on avg for P4

How can that be?

--- same ISA
--- same program + data
--- what is different?

Average by classes.
Average CPI?

$$\left[(IC_1 \times \overline{CPI_1}) + (IC_2 \times \overline{CPI_2}) + (IC_3 \times \overline{CPI_3}) \right] \left(\frac{1}{CR}\right) = IC\left(\overline{CPI}\right)\left(\frac{1}{CR}\right)$$

$$\Rightarrow \sum_{i=1}^{n} \left(\frac{IC_i}{IC}\right) \overline{CPI_i} = \sum_{i=1}^{n} (\%_i)\overline{CPI_i} = \overline{CPI}$$

$$\overline{CPI_i} \quad \%_i \Rightarrow \left\{ \overline{CPI_i}\,(\%_i) \right\}$$

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

$SUM = 1.5 = \overline{CPI}$

- Given this machine, the CPI is the sum of CPI ✗ Frequency
- Average CPI is 0.5 + 0.4 + 0.4 + 0.2 = 1.5
- What fraction of the time for data transfer?

$($

$$\frac{T_{LD\text{-}ST}}{T_{Total}} = \frac{Cycles_{LD\text{-}ST} * (1/f)}{Cycles_{Total} * (1/f)} = \frac{\#(cycles\ LD\text{-}ST)}{\#\ cycles} = \frac{(\#cycles\ LD) + (\#cycles\ ST)}{(\#\ cycles) = IC \times CPI}$$

$\uparrow$ cycle Time, $f = CR$

$$= \left[\frac{(\#cycles\ LD)}{IC} + \frac{(\#\ cycles\ ST)}{IC}\right]\frac{1}{CPI}$$

$$= \left[\frac{CPI_{LD} \cdot IC_{LD}}{IC} + \frac{CPI_{ST} \cdot IC_{ST}}{IC}\right]\frac{1}{CPI}$$

$$= [CPI_{LD} \cdot \%_{LD} + CPI_{ST} \cdot \%_{ST}](1/1.5)$$

$$= [0.4 + 0.2](1/1.5) = 40\%$$

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%

$\rightarrow$ $T_{old} = 2s$

$T_{new} = 1.5s$

What do we mean by speedup?

$$S_{new\text{-}old} = \frac{V_{new}}{V_{old}} = \frac{(W_{new}/T_{new})}{(W_{old}/T_{old})} = T_{old}/T_{new} = 1.3 \Rightarrow V_{new} = (1.3)\,V_{old}$$

$\Rightarrow$ new is 30% faster

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

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

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

$$T = T_{parallel} + T_{sequential}$$

(parallel = improvable, sequential = fixed)

$$V_p = W_p / T_p \qquad V_s = W_s / T_s$$

old: $V_p = V_s$

new: $V_p = a V_s$ (ie. $S_p = a$)

Amdahl's Law

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

$$Speedup = \frac{Told}{[(1-f)+f/a]*Told} = \frac{1}{(1-f)+f/a}$$

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

- This formula is known as Amdahl's Law
- Lessons from
  - If f→100%, then speedup = a
  - If a→∞, the speedup = 1/(1-f)
- Summary
  - Make the common case fast
  - Watch out for the non-optimized component

target speed up to max %

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

$$T_{new} = W_p/V_p + W_s/V_s = \frac{f \cdot W}{a V_s} + \frac{(1-f) \cdot W}{V_s} = \left(\frac{f}{a} + (1-f)\right) W/V_s$$

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

$$= 1/\left(f/a + (1-f)\right)$$

$$\xrightarrow{f=1} 1/(1/a + 0) = a = S_{all\text{-}improved} = S_{max}$$

$$\xrightarrow{f=0} 1/(0/a + 1) = 1 = S_{none\text{-}improved} = S_{min}$$

$$\xrightarrow[f \text{ fixed}]{a \to \infty} 1/\left(f/\infty + (1-f)\right) = \frac{1}{(1-f)} = S_\infty \geq S_{actual}$$

$$f = 50\% \Rightarrow S_\infty = \frac{1}{(1-1/2)} = 2$$

max improvement for $a \to \infty$ is twice as fast

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

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

$cost_{new} = 100 \cdot Cost_{old}$

S

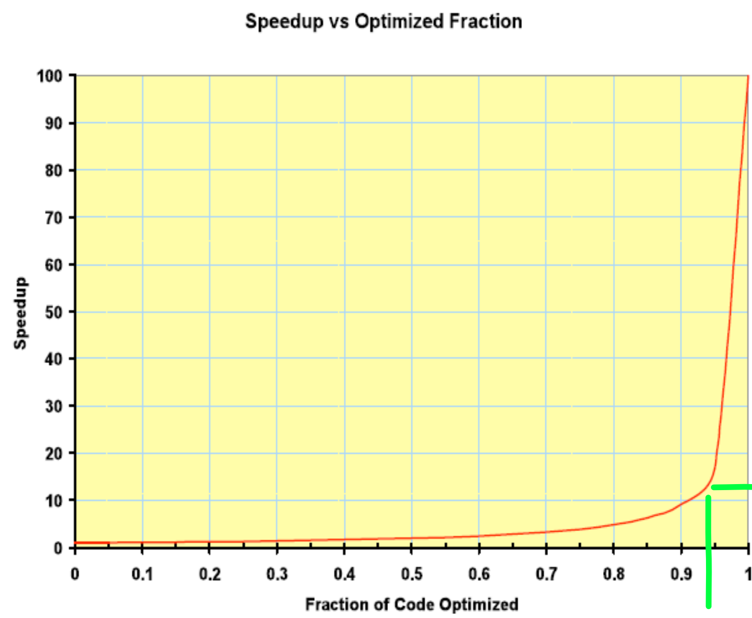

$S' = 17$?

5% sequential

f

$f = 95\%$, $cost = 100$, $S' = 17$ Hmm?

Only very large f is worth it?

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?" $= S'_{overall} = 4$

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

$S_{new\text{-}old} = 4 = T_{old}/T_{new} = \frac{100s}{20_s + 80s/S_p} \Rightarrow 80/S_p = \frac{100}{4} - 20 = 5$

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

- How about making it 5 times faster?

$S = 5$?

$$x \cdot 7 = x(4+2+1) = \underset{C}{\underset{\text{2 shifts}}{4x}} + \underset{B}{\underset{\text{1 shift}}{2x}} + \underset{A}{\underset{\text{0 shifts}}{x}}$$

$y = 00\ldots0111$

$$
\begin{array}{rcccccc|l}
 & & 0 & 0 & \cdots & 0 & 0 & = S^0 \\
+ & & x_n & x_{n-1} & \cdots & x_1 & x_0 & + (A) \\
\hline
 & & S'_n & S'_{n-1} & \cdots & S'_1 & S'_0 & S'
\end{array}
$$

shift ←

$$
\begin{array}{rccccccc|l}
+ & x_n & x_{n-1} & \cdots & x_1 & x_0 & 0 & & + (B) \\
\hline
 & S''_{n+1} & S''_n & \cdots & S''_2 & S''_1 & S''_0 & & S''
\end{array}
$$

shift ←

$$
\begin{array}{rcccccccc|l}
+ & x_n & x_{n-1} & \cdots & x_1 & x_0 & 0 & 0 & & + (C) \\
\hline
 & S'''_{n+2} & S'''_{n+1} & \cdots & S'''_3 & S'''_2 & S'''_1 & S'''_0 & & S'''
\end{array}
$$

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

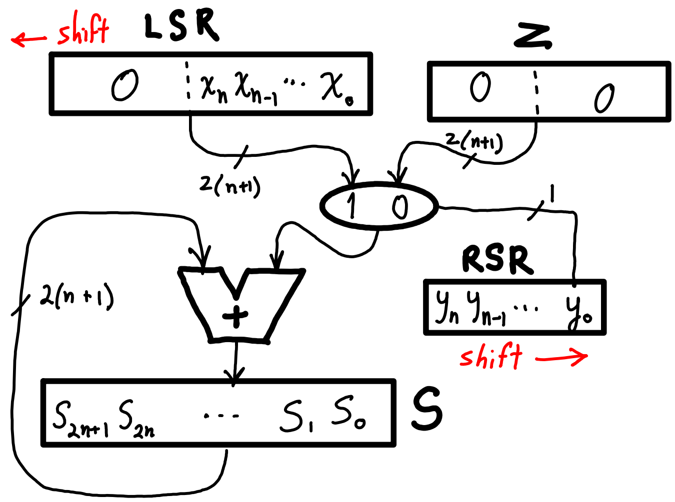

e.g.

$$
\begin{array}{r}
1011 \\
\times\ 0101 \\
\hline
1011 \\
+\ 101100 \\
\hline
00110111
\end{array}
$$

(possible carry)

rewrite ⇒

		multiplier bit
1011 × 0101		
+ 0 0 0 0 0 0 0 0	← start	
+ 0 0 0 0 1 0 1 1	← 0-shift	← $Y_0 = 1$
+ 0 0 0 0 0 0 0 0	← 1-shift	← $Y_1 = 0$
+ 0 0 1 0 1 1 0 0	← 2-shift	← $Y_2 = 1$
+ 0 0 0 0 0 0 0 0	← 3-shift	← $Y_3 = 0$
= 0 0 1 1 0 1 1 1		

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

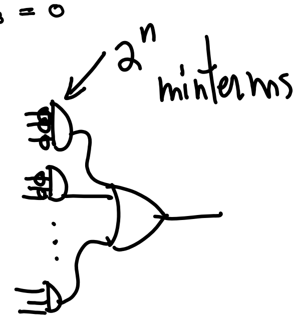

ADDER

Delay is longest path.

2 levels per MAJ, n bit operands.

2n gate delays until result is ready.

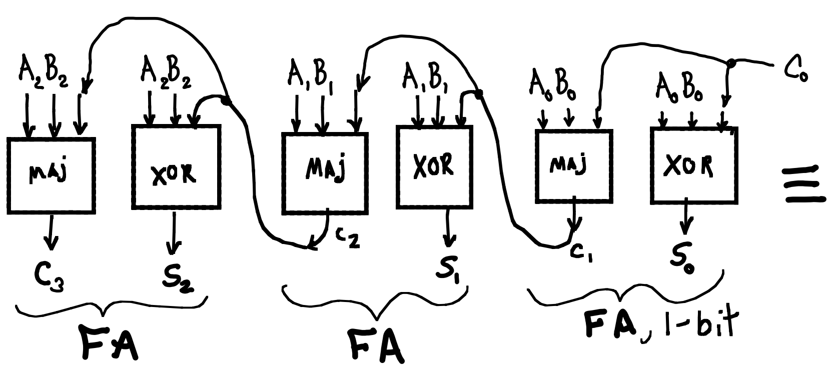

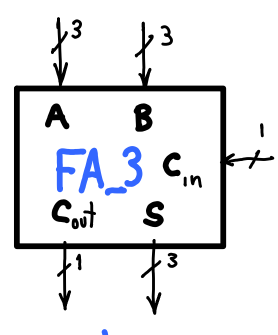

Shift Register Delay $= 1$ (all stages in Parallel)

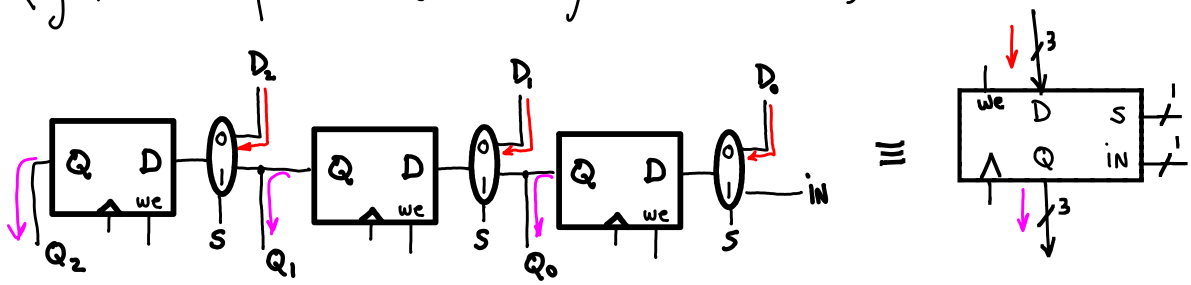

Overall delay: $(n \text{ shifts}) \times (2n \text{ delay}_{ADD}) = 2n^2$

Improvement? | Recursive Refinement?

(1-bit MULT)

$$\begin{array}{r}0\\ \times 0\\ \hline 0\end{array} \quad \begin{array}{r}0\\ \times 1\\ \hline 0\end{array} \quad \begin{array}{r}1\\ \times 0\\ \hline 0\end{array} \quad \begin{array}{r}1\\ \times 1\\ \hline 1\end{array} \Rightarrow$$

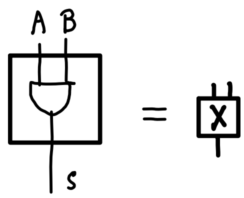

(2-bit MULT)

$$\begin{array}{rcccc} & & & (x_1 & x_0) \\ & & \times & (y_1 & y_0) \\ \hline & 0 & 0 & & x_0y_0 \\ + & 0 & x_1y_0 & & 0 \\ + & 0 & y_1x_0 & & 0 \\ + & x_1y_1 & 0 & & 0 \\ \hline & S_3 & S_2 & S_1 & S_0 \end{array}$$

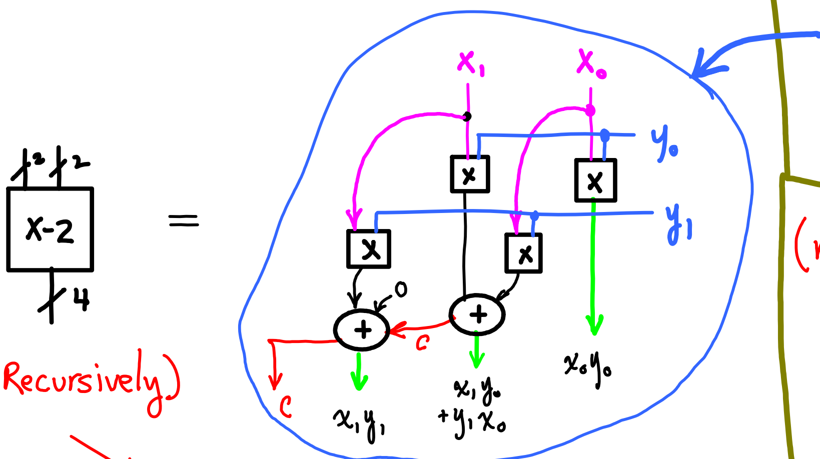

(Use Recursively)

(n-bit)

$$\begin{array}{ccccc} & & & \boxed{x_3\,x_2} & \boxed{x_1\,x_0} \\ & & \times & \boxed{y_3\,y_2} & \boxed{y_1\,y_0} \\ \hline & 0 & 0 & \boxed{X_L \cdot Y_L} & \\ & 0 & \boxed{X_H \cdot Y_L} & 0 & \\ & 0 & \boxed{X_L \cdot Y_H} & 0 & \\ + & \boxed{X_H \cdot Y_H} & 0 & 0 & \\ \hline & S_{HH} & S_{LH} & S_{HL} & S_{LL} \end{array}$$

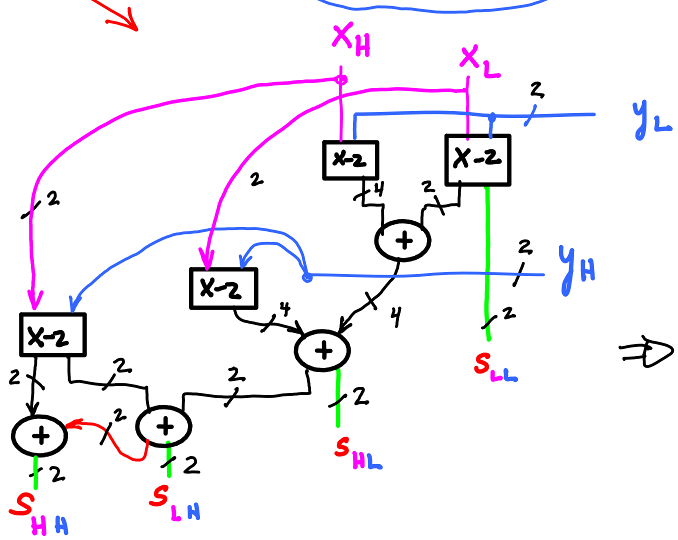

$\Rightarrow$ longest delay, $f(n)$, is,

4n-bit Add + $(n/2)$-bit MULT delay:

$$f(n) = 4(2n) + f(n/2)$$

$$f(n) = 8n + f(n/2) = 8n + 8(n/2) + f(n/4) \Rightarrow 16n\left(\tfrac{1}{2}+\tfrac{1}{4}+\tfrac{1}{8}+\cdots\right) \approx 16n$$

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

32-bit $2(2^5)^2 = 2^{11} \Rightarrow 2^4(2^5) = 2^9 \rightarrow S_{32\text{-bit}} = 2^{11}/2^9 = 2^2 = 4$

$S_{64\text{-bit}} = 8$

$S_{128\text{-bit}} = 16$

$$4 = S_{overall} = \frac{V_{new}}{V_{old}} = \frac{W/T_{new}}{W/T_{old}} = \frac{T_{old}}{T_{new}} = \frac{W_s/V_{s\text{-}old} + W_p/V_{p\text{-}old}}{W_s/V_{s\text{-}new} + W_p/V_{p\text{-}new}}$$

OUR Requirement

$$= \frac{(0.2)W/V_{p\text{-}old} + (0.8)W/V_{p\text{-}old}}{(0.2)W/V_{p\text{-}old} + (0.8)W/S_p \cdot V_{p\text{-}old}}$$

Assume

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

$V_{p\text{-}new} = S_p V_{p\text{-}old}$

$$= 1\Big/\left(0.2 + 0.8/S_p\right) = 4$$

$$\Rightarrow 0.25 = 0.2 + 0.8/S_p$$

$$\Rightarrow 0.05 = 0.8/S_p \Rightarrow S_p = \frac{0.8}{0.05} = \frac{80}{5} = 16 = S_{128\text{-MULT}}$$

OK!

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

$S_{overall} = 5\,?$ $1/(0.2 + 0.8/S_p) \overset{?}{=} 5$ Can we?

$S_{overall-\infty} = 1/(0.2 + 0.8/\infty) = 1/0.2 = 5$ Possible? if we eliminate MULT time?

---

what else could we try?

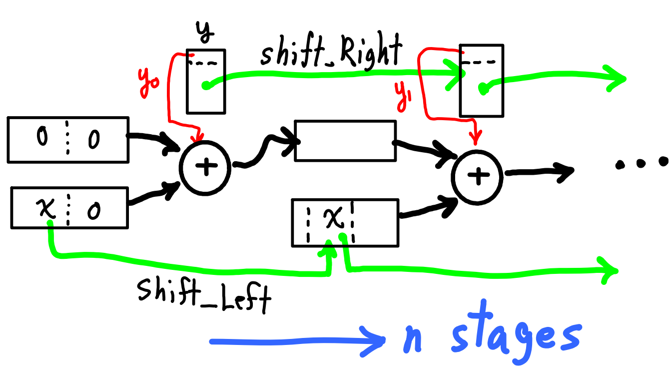

n-stage pipeline

1 result every shift

stage delay $= 4n$ (2n-bit ADD)

(32-bit):

$$S_{32\text{-MULT}} = \frac{2(2^5)^2}{4(2^5)} = 2^4 = 16$$

[Can we keep it full?]

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

Pipeling

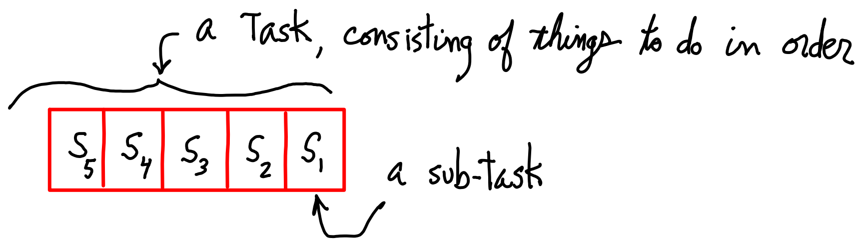

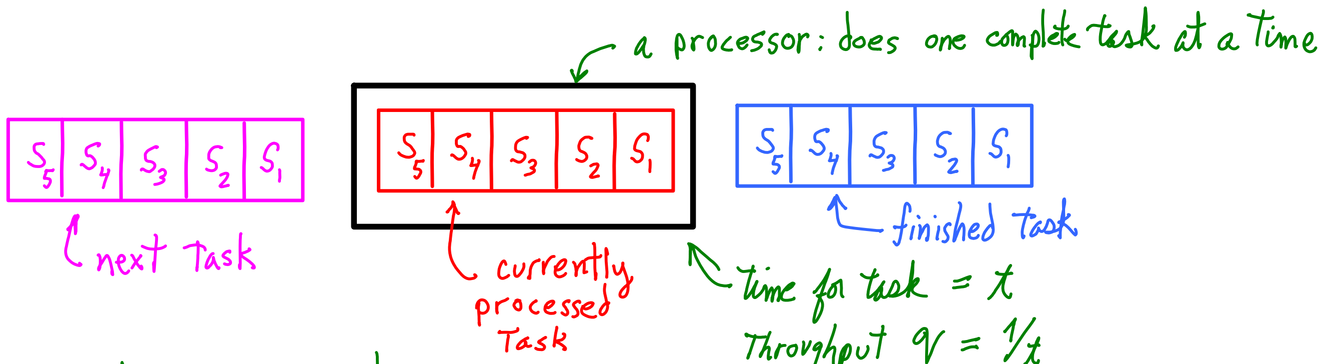

Break processor into pieces, one for each S_i

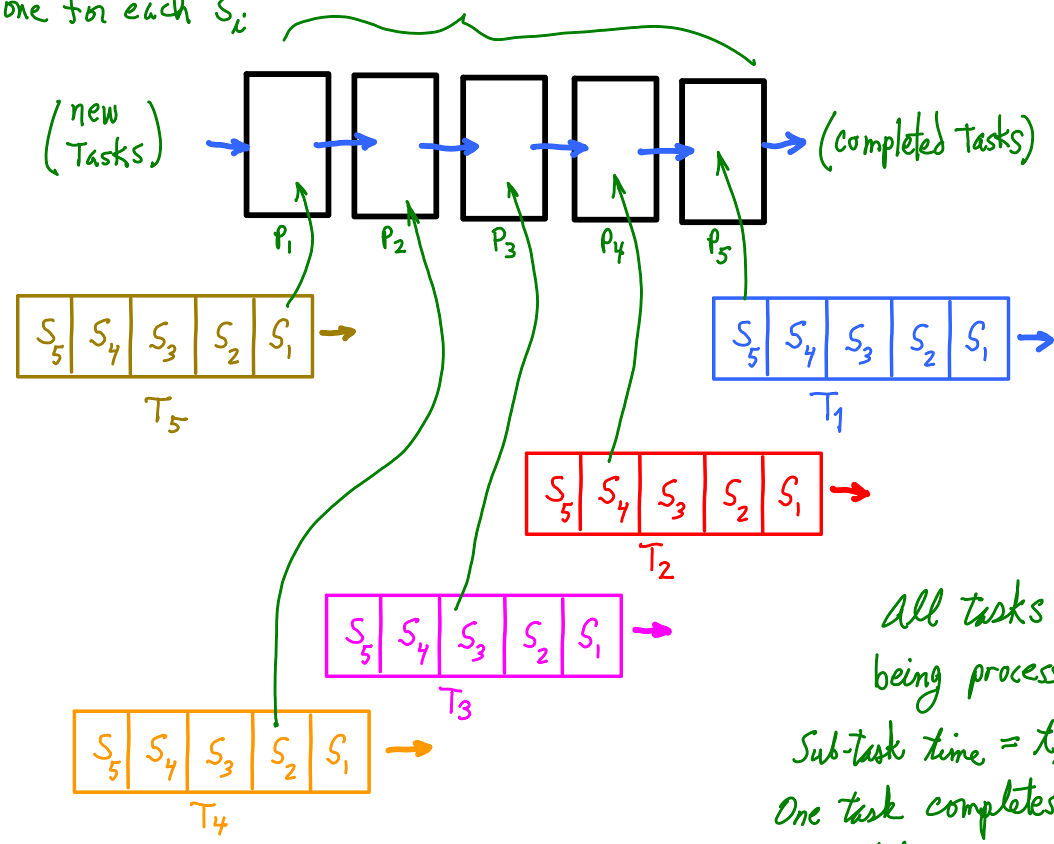

All tasks have a sub-task being processed in parallel.

Sub-task time $= t/5$.

One task completes every $t/5$.

$\nu = 1/(t/5) = 5/t$.

$S' = 5$

Another Approach? Back to square 1

Adding Partial Products

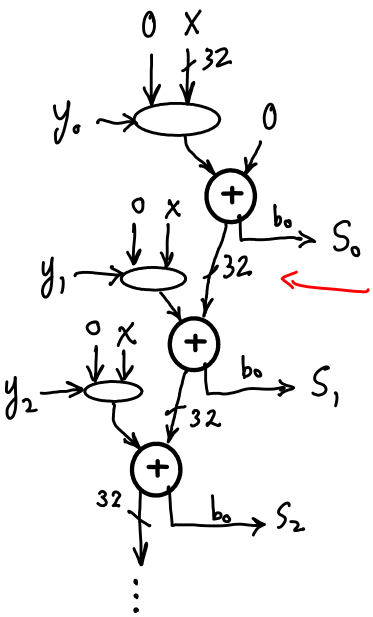

Partial sum, includes carry

carry-save adder

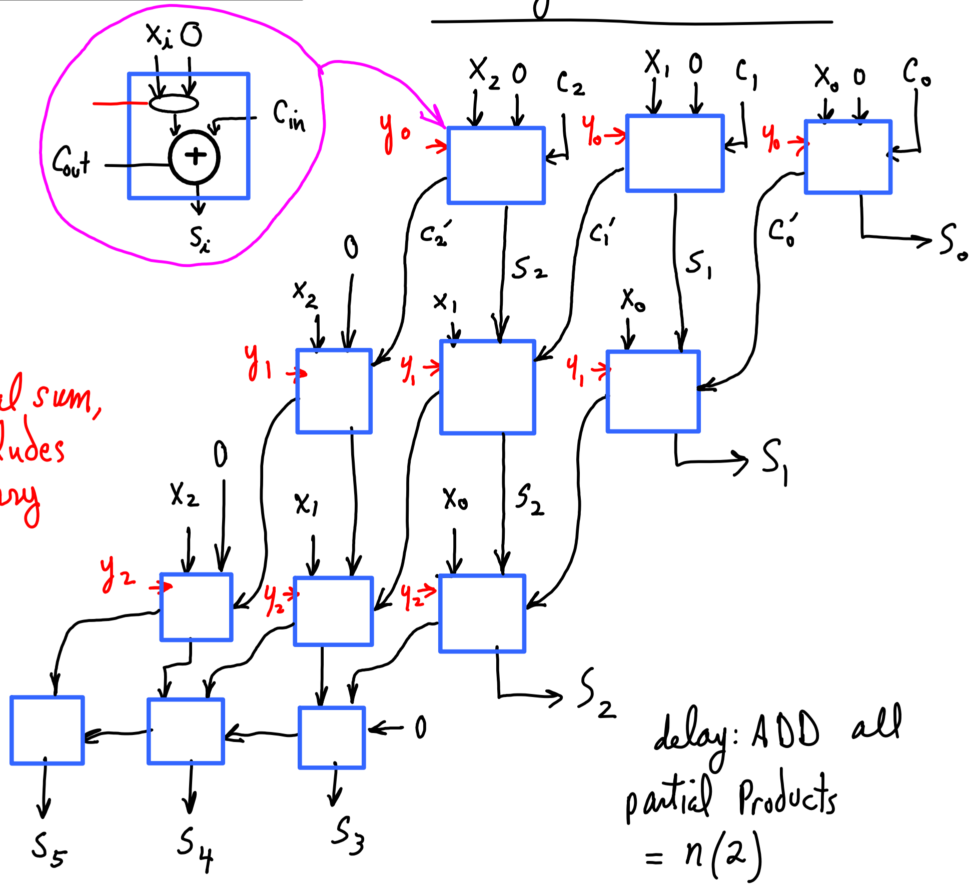

delay: ADD all partial Products $= n(2)$

$S = 2^{11}/2^3 = 2^8$

$\frac{10}{2 + 8/5} \quad \frac{10}{2 + 1/25} \quad \frac{10}{\frac{65}{32}} = \frac{320}{65} = 4.9$

$S_{32\text{-MULT}} = 2^{11}/2^6 = 32$

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

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

$= 10/(2 + 1/4) = 40/(8+1) = \frac{40}{9} = 4 + 4/9$

Other Ideas?

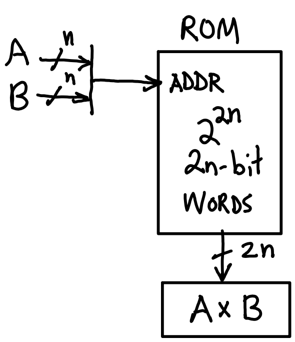

How fast?

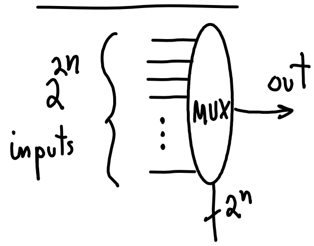

delay? 2n

area? exponential

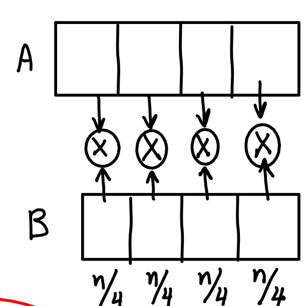

Combine

delay = n/2

Do mod arith in Parallel, combine results

How?

Evaluating Performance

- Performance best determined by running a real application in typical environment?
  - 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

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

The SPEC CPU Benchmark Suite (System Performance Evaluation Cooperative)

	SPEC2006 benchmark description	SPEC2006	SPEC2000	SPEC95	SPEC92	SPEC89
		Benchmark name by SPEC generation				
gcc	GNU C compiler					gcc
	Interpreted string processing			perl		espresso
	Combinatorial optimization		mcf			li
	Block-sorting compression		bzip2		compress	eqntott
GO	Go game (AI)	go	vortex	go	sc	
	Video compression	h264avc	gzip	ijpeg		
	Games/path finding	astar	eon	m88ksim		
	Search gene sequence	hmmer	twolf			
QM sim	Quantum computer simulation	libquantum	vortex			
	Discrete event simulation library	omnetpp	vpr			
	Chess game (AI)	sjeng	crafty			
XML parse	XML parsing	xalancbmk	parser			
	CFD/blast waves	bwaves				fpppp
	Numerical relativity	cactusADM				tomcatv
	Finite element code	calculix				doduc
	Differential equation solver framework	dealII				nasa7
	Quantum chemistry	gamess				spice
	EM solver (freq/time domain)	GemsFDTD			swim	matrix300
	Scalable molecular dynamics (~NAMD)	gromacs		apsi	hydro2d	
	Lattice Boltzman method (fluid/air flow)	lbm		mgrid	su2cor	
	Large eddie simulation/turbulent CFD	LESlie3d	wupwise	applu	wave5	
	Lattice quantum chromodynamics	milc	apply	turb3d		
	Molecular dynamics	namd	galgel			
	Image ray tracing	povray	mesa			
	Spare linear algebra	soplex	art			
speech recog.	Speech recognition	sphinx3	equake			
	Quantum chemistry/object oriented	tonto	facerec			
	Weather research and forecasting	wrf	ammp			
	Magneto hydrodynamics (astrophysics)	zeusmp	lucas			
			fma3d			
			sixtrack			

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

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

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

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

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

find ratio $\bar{r}$ s.t.

$r_1 \cdot r_2 \cdots r_n x = (\bar{r})^n x$

$$R \Rightarrow (\overset{b_1}{T_{R_1}},\ \overset{b_2}{10\,T_{R_1}})$$

$$S_{A\text{-}R_1} = \frac{T_{R_1}}{100} \qquad S_{A\text{-}R_2} = \frac{10\,T_{R_1}}{4}$$

$$S_{B\text{-}R_1} = \frac{T_{R_1}}{200} \qquad S_{B\text{-}R_2} = \frac{10\,T_{R_1}}{1}$$

Normalize: use reference **machine R** to get **speedups w.r.t. benchmarks** (**b-1, b-2**).

(R's time on b-2) = 10 X (R's time on b-1).

Combine speedups w.r.t R:
--- Get **mean of speedups** w.r.t. R for **A**
--- Get **mean of speedups** w.r.t. R for **B**
--- Take **ratio of mean speedups.**

$$\overline{S}_{A\text{-}R} = \tfrac{1}{2}\left(\frac{T_{R_1}}{100} + \frac{10\,T_{R_1}}{4}\right)$$

$$\overline{S}_{B\text{-}R} = \tfrac{1}{2}\left(\frac{T_{R_1}}{200} + \frac{10\,T_{R_1}}{1}\right)$$

$$\Longrightarrow \frac{\overline{S}_{A\text{-}R}}{\overline{S}_{B\text{-}R}} = \frac{T_{R_1}/2\left(\frac{4+1000}{400}\right)}{T_{R_1}/2\left(\frac{1+2000}{200}\right)} = \left(\tfrac{1}{2}\right)\left(\frac{1004}{2001}\right) \cong 1/4\ ?$$

$$= \left(\tfrac{1}{2}\right)\left(\frac{4+100r}{1+200r}\right) \qquad r \to \infty : 1/4 \qquad r \to 0 : 2$$

r makes all the difference: **changing R** or **benchmarks** ===> **opposite conclusions?**

Geometric Mean

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

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

$$S_{A\text{-}B} = \frac{\overline{S}_{A\text{-}R}}{\overline{S}_{B\text{-}R}} = \frac{(\sqrt{10}\ T_{R_1})}{(\sqrt{10}\ T_{R_1})}\left(\frac{T_{B_1} \cdot T_{B_2}}{T_{A_1} \cdot T_{A_2}}\right)^{1/2} = \left(\frac{T_{B_1} \cdot T_{B_2}}{T_{A_1} \cdot T_{A_2}}\right)^{1/2} \cong 0.7$$

R cancels. Conclusion **$S_{A\text{-}B}$ = 30 % slower?** Is this fair?
---- **on b1: $S_{A\text{-}B}$ = 200/100 = 2**
---- **on b2: $S_{A\text{-}B}$ = 1/4**

job mix = (**n1 runs of b-1**) + (**n2 runs of b-2**)

$$S_{A\text{-}B} = \frac{n_1 T_{B_1} + n_2 T_{B_2}}{n_1 T_{A_1} + n_2 T_{A_2}} = \frac{200\,n_1 + n_2}{100\,n_1 + 4\,n_2} = \frac{200 + a}{100 + 4a} = \begin{cases} a \to \infty : 1/4 \\ a \to 0 : 2 \end{cases}$$

$$(a = n_2/n_1)$$

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

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

$$\implies n_2 = 62\, n_1$$ **For every short job** (b1), **62 long jobs** (b2)?

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

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

$$HM = \frac{1}{\sum w_i / r_i}$$ where $r_1 = \frac{W_1}{T_1}$, $r_2 = \frac{W_2}{T_2}$ ... $[\, r_i = V_i \,]$

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

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

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

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

$W = \sum W_i$

We have reference Times:

$$T_{R_1} = \frac{W_1}{V_{R-1}} \qquad T_{R_2} = W_2 / V_{R-2}$$ [assume $V_{R-1} = V_{R-2} = V_R$]

$$W_1 = T_{R-1} V_R \qquad W_2 = T_{R-2} V_R$$

Use that to get weights:

$$\omega_i = W_i / \sum W_i$$

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

$$= 1/(1+10) = 1/11$$

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

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

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

$$\frac{W_1}{W_2} = \frac{T_{R-1}\, v_R}{T_{R-2}\, v_R} = \frac{T_{R-1}}{T_{R-2}} = \frac{T_{R-1}}{10\,T_{R-1}} = 1/10$$

$$W_2 = 10\,W_1$$

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

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

$$\overline{v}_A = \frac{1}{\left(\frac{(1/11)}{W_1/100} + \frac{(10/11)}{10W_1/4}\right)} = \frac{(11)}{\frac{100}{W_1} + \frac{10\cdot 4}{10\,W_1}} = \frac{11\,W_1}{(\underset{T_{A-1}}{100} + \underset{T_{A-2}}{4})}$$

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

$$\longrightarrow \quad S^{HM}_{A-B} = \overline{v}_A / \overline{v}_B = \frac{(200+1)}{(100+4)} \cong 2$$

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?

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

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

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

$$\Rightarrow \; 700 = 700a$$

$$\boxed{a = 1}$$

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