

CS184a: Computer Architecture (Structure and Organization)


Day 18: February 21, 2003
Retiming 2:
Structures and Balance



Caltech CS184 Winter2003 -- DeHon

Last Time

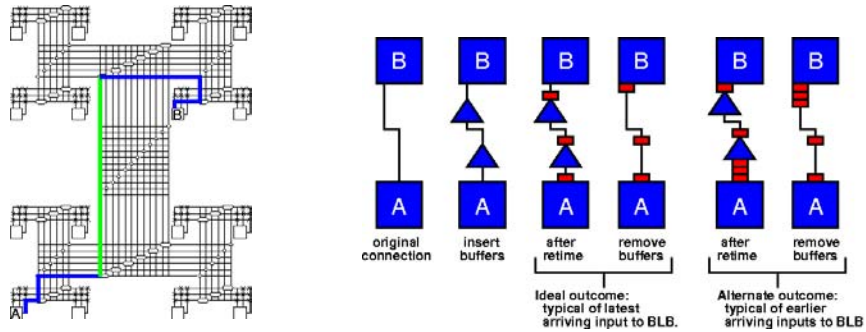
- Saw how to formulate and automate retiming:
 - start with network
 - calculate minimum achievable c
 - c = cycle delay (clock cycle)
 - make c -slow if want/need to make $c=1$
 - calculate new register placements and move



Caltech CS184 Winter2003 -- DeHon

Last Time

- Systematic transformation for retiming
 - “justify” mandatory registers in design



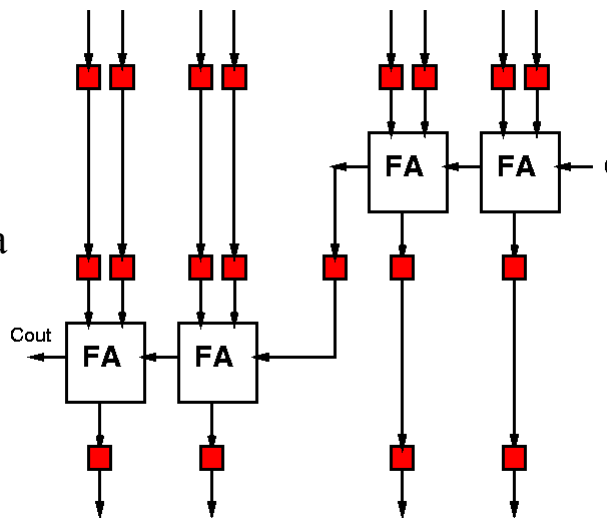
Today

- Retiming in the Large
- Retiming Requirements
- Retiming Structures

Retiming in the Large

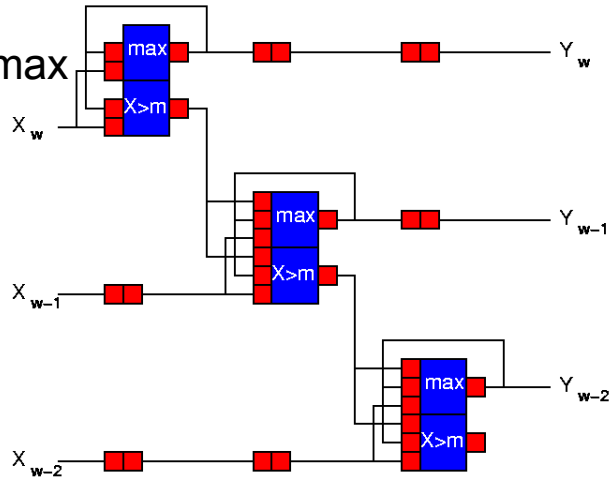
Align Data / Balance Paths

Day3:
registers
to align data



Systolic Data Alignment


- Bit-level max



Caltech CS184 Winter2003 -- DeHon

Serialization

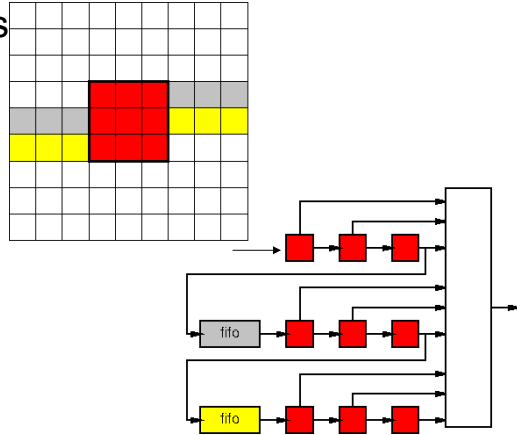
- Serialization
 - greater serialization → deeper retiming
 - **total:** same **per compute:** larger



Caltech CS184

Data Alignment

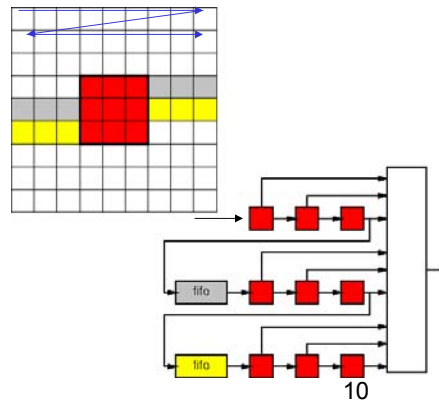
- For video (2D) processing
 - often work on local windows
 - retime scan lines
- E.g.
 - edge detect
 - smoothing
 - motion est.



Caltech CS184 Winter2003 -- DeHon

Image Processing


- See Data in raster scan order
 - adjacent, horizontal bits easy
 - adjacent, vertical bits
 - scan line apart




Caltech CS184 Winter2003 -- DeHon

Wavelet

- Data stream for horizontal transform



- Data stream for vertical transform
– N=image width



Caltech CS184 Winter2003 -- DeHon

Retiming in the Large

- Aside from the local retiming for cycle optimization (last time)
- Many intrinsic needs to retime data for correct use of compute engine
 - some very deep
 - often arise from serialization

Caltech CS184 Winter2003 -- DeHon

Reminder: Temporal Interconnect

- Retiming \equiv Temporal Interconnect
- Function of *data* memory
 - perform retiming

Requirements not Unique

- Retiming requirements are not unique to the problem
- Depends on algorithm/implementation
- Behavioral transformations can alter significantly

Requirements Example

$$Q = A * B + C * D + E * F$$

- For $l \leftarrow 1$ to N
 - $t1[l] \leftarrow A[l] * B[l]$
- For $l \leftarrow 1$ to N
 - $t2[l] \leftarrow C[l] * D[l]$
- For $l \leftarrow 1$ to N
 - $t3[l] \leftarrow E[l] * F[l]$
- For $l \leftarrow 1$ to N
 - $t2[l] \leftarrow t1[l] + t2[l]$
- For $l \leftarrow 1$ to N
 - $Q[l] \leftarrow t2[l] + t3[l]$
- For $l \leftarrow 1$ to N
 - $t1 \leftarrow A[l] * B[l]$
 - $t2 \leftarrow C[l] * D[l]$
 - $t1 \leftarrow t1 + t2$
 - $t2 \leftarrow E[l] * F[l]$
 - $Q[l] \leftarrow t1 + t2$
- left \Rightarrow $3N$ regs
- right \Rightarrow 2 regs


Retiming Structure and Requirements

Structures

- How do we implement programmable retiming?
- Concerns:
 - Area: λ^2/bit
 - Throughput: bandwidth (bits/time)
 - Latency important when do not know when we will need data item again

Just Logic Blocks


- Most primitive
 - build flip-flop out of logic blocks
 - $I \leftarrow D^*/\text{Clk} + I^*\text{Clk}$
 - $Q \leftarrow Q^*/\text{Clk} + I^*\text{Clk}$



- Area: 2 LUTs (800K \rightarrow 1M λ^2/LUT each)
- Bandwidth: 1b/cycle

Optional Output


- Real flip-flop (optionally) on output



- flip-flop: $4-5K\lambda^2$
- Switch to select: $\sim 5K\lambda^2$
- Area: 1 LUT ($800K \rightarrow 1M\lambda^2/\text{LUT}$)
- Bandwidth: 1b/cycle

Output Flip-Flop Needs

- Pipeline and C-slow to LUT cycle
- Always need an output register




Number of Registers	1	2	3	4	5	6	7	8	9	10
Percentage	72	16	4.5	2.2	1.3	0.96	1.2	0.46	0.12	0.11

Table 3: Benchmark Wide Distribution of Registers Required between LUTs
Average Regs/LUT 1.7, some designs need 2--7x

Separate Flip-Flops

- Network flip flop w/ own interconnect
 - + can deploy where needed
 - requires more interconnect



- Assume routing \propto inputs
 - 1/4 size of LUT
- Area: $200K\lambda^2$ each
- Bandwidth: 1b/cycle

21

Deeper Options


- Interconnect / Flip-Flop is expensive
- How do we avoid?

22

Deeper


- Implication
 - don't need result on every cycle
 - number of regs > bits need to see each cycle
 - → lower bandwidth acceptable
 - → less interconnect

Deeper Retiming




Output

- Single Output
 - Ok, if don't need other timings of signal
- Multiple Output
 - more routing




Input

- More registers ($K \times$)
 - $7-10K\lambda^2/\text{register}$
 - $4\text{-LUT} \Rightarrow 30-40K\lambda^2/\text{depth}$
- No more interconnect than unretimed
 - **open**: compare savings to additional reg. cost
 - Area: $1 \text{ LUT } (1M+d*40K\lambda^2)$ get Kd regs
 - $d=4, 1.2M\lambda^2$
 - Bandwidth: K/cycle
 - $1/d$ th capacity

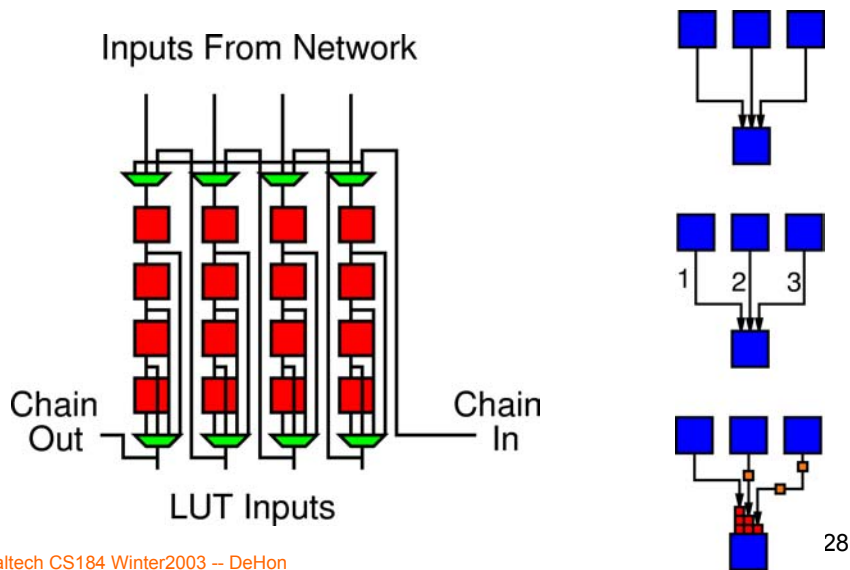


HSRA Input

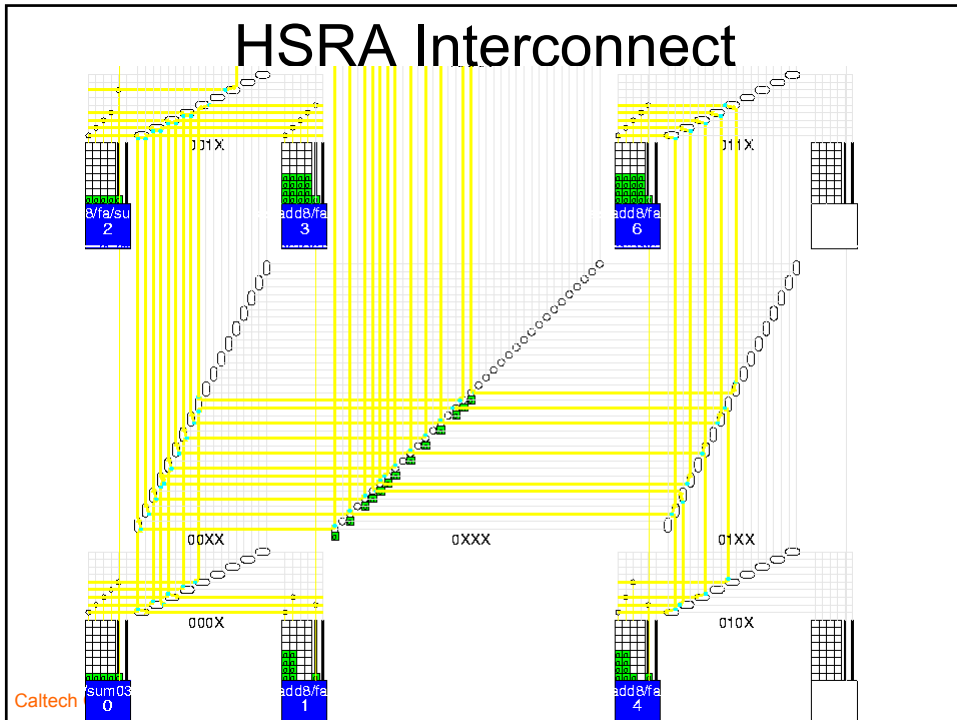


Caltech CS184 Winter2003 -- DeHon

Input Retiming



Caltech CS184 Winter2003 -- DeHon



Flop Experiment #1

- Pipeline and retime to single LUT delay per cycle
 - MCNC benchmarks to 256 4-LUTs
 - no interconnect accounting

Number of Registers	1	2	3	4	5	6	7	8	9	10
Percentage	72	16	4.5	2.2	1.3	0.96	1.2	0.46	0.12	0.11

– average 1.7 registers/LUT (some circuits 2--7)

Flop Experiment #2

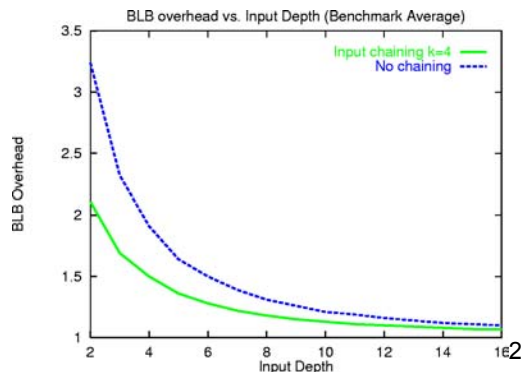

- Pipeline and retime to HSRA cycle
 - place on HSRA
 - single LUT or interconnect timing domain
 - same MCNC benchmarks

Number of Registers	1	2	3	4	5	6	7	8	9	10	>10
Percentage	60	6.9	5.9	3.8	4.3	2.7	2.6	1.9	1.5	1.2	9.2


- average 4.7 registers/LUT

Input Depth Optimization


- Real design, fixed input retiming depth
 - truncate deeper and allocate additional logic blocks



Extra Blocks (limited input depth)




Average




Worst Case Benchmark

With Chained Dual Output [can use one BLB as 2 retiming-only chains]

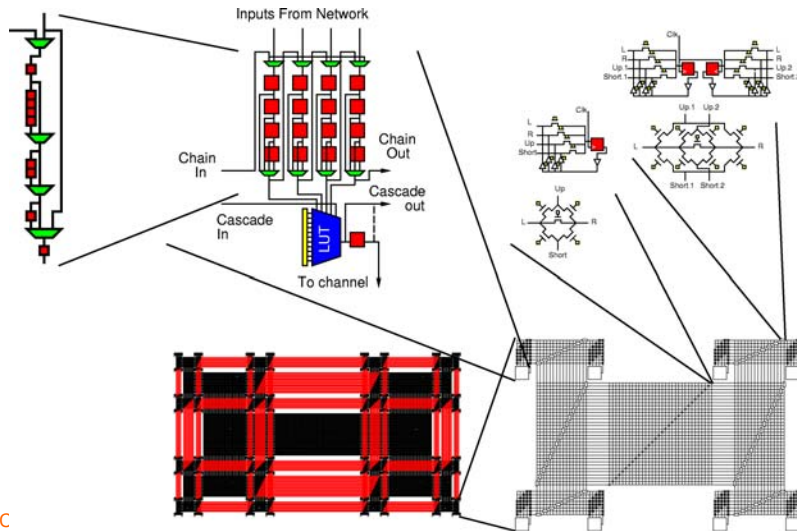


Average




Worst Case Benchmark

HSRA Architecture



Caltech C


Register File



- From MIPS-X
 - $1K\lambda^2/\text{bit} + 500\lambda^2/\text{port}$
 - $\text{Area}(\text{RF}) = (d+6)(W+6)(1K\lambda^2 + \text{ports} * 500\lambda^2)$
- $w \gg 6, d \gg 6 \quad l+o=2 \Rightarrow 2K\lambda^2/\text{bit}$
- $w=1, d \gg 6 \quad l=o=4 \Rightarrow 35K\lambda^2/\text{bit}$
 - comparable to input chain
- More efficient for wide-word cases

Xilinx CLB

- Xilinx 4K CLB
 - as memory
 - works like RF



- Area: $1/2 \text{ CLB } (640K\lambda^2)/16 \approx 40K\lambda^2/\text{bit}$
 - but need 4 CLBs to control
- Bandwidth: $1b/2 \text{ cycle } (1/2 \text{ CLB})$
 - $1/16$ th capacity

Memory Blocks

- SRAM bit $\approx 1200\lambda^2$ (large arrays)
- DRAM bit $\approx 100\lambda^2$ (large arrays)
- Bandwidth: $W \text{ bits} / 2 \text{ cycles}$
 - usually single read/write
 - $1/2^A$ th capacity

Disk Drive

- Cheaper per bit than DRAM/Flash
 - (not MOS, no λ^2)
- Bandwidth: 10-20Mb/s
 - For 4ns array cycle
 - 1b/12.5 cycles @20Mb/s

Hierarchy/Structure Summary

- “Memory Hierarchy” arises from area/bandwidth tradeoffs
 - Smaller/cheaper to store words/blocks
 - (saves routing and control)
 - Smaller/cheaper to handle long retiming in larger arrays (reduce interconnect)
 - High bandwidth out of registers/shallow memories

λ^2	DRAM	SRAM	RF bit	FF/RF	RF \times 1	XC	In FF	net FF	FF/LUT
bw/cap.	100	1200	2K	5K	40K	40K	75K	200K	800K
	$1/10^7$	$1/10^5-10^3$		$1/100$	$1/100$	$1/16$	$1/4$	$1/1$	$1/1$

Modern FPGAs

- Output Flop (depth 1)
- Use LUT as Shift Register (16)
- Embedded RAMs (16Kb)
- Interface off-chip DRAM (~0.1—1Gb)
- No retiming in interconnect
 -yet

Modern Processors

- DSPs have accumulator (depth 1)
- Inter-stage pipelines (depth 1)
 - Lots of pipelining in memory path...
- Reorder Buffer (4—32)
- Architected RF (16, 32, 128)
- Actual RF (256, 512...)
- L1 Cache (~64Kb)
- L2 Cache (~1Mb)
- L3 Cache (10-100Mb)
- Main Memory in DRAM (~10-100Gb)

Big Ideas

[MSB Ideas]

- Tasks have a wide variety of retiming distances (depths)
- Retiming requirements affected by high-level decisions/strategy in solving task
- Wide variety of retiming costs
 - $100 \lambda^2 \rightarrow 1M\lambda^2$
- Routing and I/O bandwidth
 - big factors in costs
- Gives rise to memory (retiming) hierarchy