

Lab #3 due tonight,
LPSet 8 Thurs 10/11

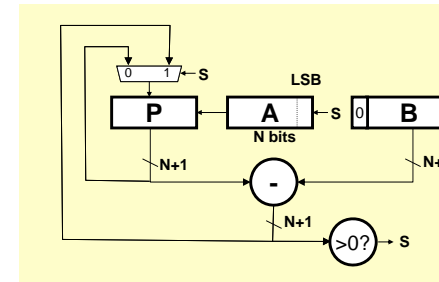


Pipelining & Verilog

- Latency & Throughput
- Pipelining to increase throughput
- Retiming
- Verilog Math Functions
- Debugging Hints

Sequential Divider

Assume the Dividend (A) and the divisor (B) have N bits. If we only want to invest in a single N-bit adder, we can build a sequential circuit that processes a single subtraction at a time and then cycle the circuit N times. This circuit works on unsigned operands; for signed operands one can remember the signs, make operands positive, then correct sign of result.



```

Init: P ← 0, load A and B
Repeat N times {
  shift P/A left one bit
  temp = P - B
  if (temp > 0)
    {P ← temp, ALSB ← 1}
  else ALSB ← 0
}
Done: Q in A, R in P
  
```

Verilog divider.v

```

// The divider module divides one number by another. It
// produces a signal named "ready" when the quotient output
// is ready, and takes a signal named "start" to indicate
// the input dividend and divisor is ready.
// sign = 0 for unsigned, 1 for two's complement

// It uses a simple restoring divide algorithm.
// http://en.wikibooks.org/wiki/Divisions_of_Integer#/restoring_division
module divider #(parameter WIDTH = 8) (ready, start, quotient,
    remainder, dividend,
    divisor, sign, clk);

    input clk;
    input sign;
    input start;
    input [WIDTH-1:0] dividend, divisor;
    output [WIDTH-1:0] quotient, remainder;
    output

    reg [WIDTH-1:0] quotient, quotient_temp;
    reg [WIDTH-2:1:0] dividend_copy, divisor_copy, diff;
    reg negative_output;

    wire [WIDTH-1:0] remainder = ({negative_output} ?
        dividend_copy[31:0] :
        "dividend_copy[31:0] + 1'b1");

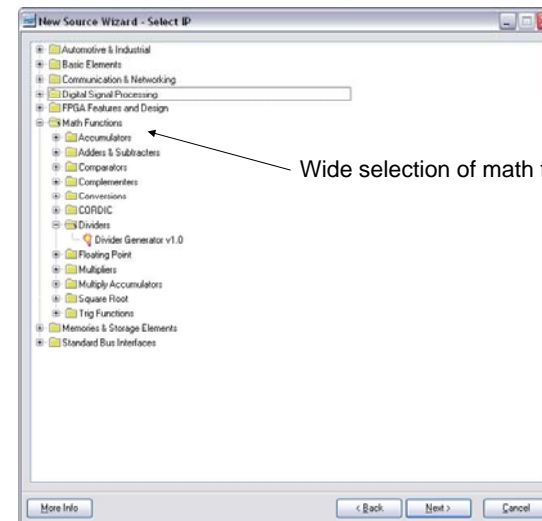
    reg [7:0] bit;
    reg del_ready = 1;
    wire ready = (!bit) & "del_ready";

    wire [WIDTH-2:0] zero = 0;
    initial bit = 0;
    initial negative_output = 0;

    always @(posedge clk) begin
        del_ready <= !bit;
        if (start) begin
            bit = WIDTH;
            quotient = 0;
            quotient_temp = 0;
            dividend_copy = ({sign} ? {dividend[WIDTH-1]} ?
                {1'b0, zero, dividend + 1'b1};
            divisor_copy = {sign} & {divisor[WIDTH-1]} ?
                {1'b0, divisor + 1'b1, zero};
            negative_output = sign &&
                ((divisor[WIDTH-1] && {dividend[WIDTH-1]}) ?
                    !({divisor[WIDTH-1] && {dividend[WIDTH-1]}}));
            diff = dividend_copy - divisor_copy;
            quotient_temp <= 0;
            if (!diff[WIDTH-2:1]) begin
                dividend_copy = diff;
                quotient_temp[0] = 1'd1;
            end
            quotient = ({negative_output} ?
                quotient_temp :
                quotient_temp + 1'b1);
            divisor_copy = divisor_copy >> 1;
            bit = bit - 1'b1;
        end
    end
endmodule
  
```

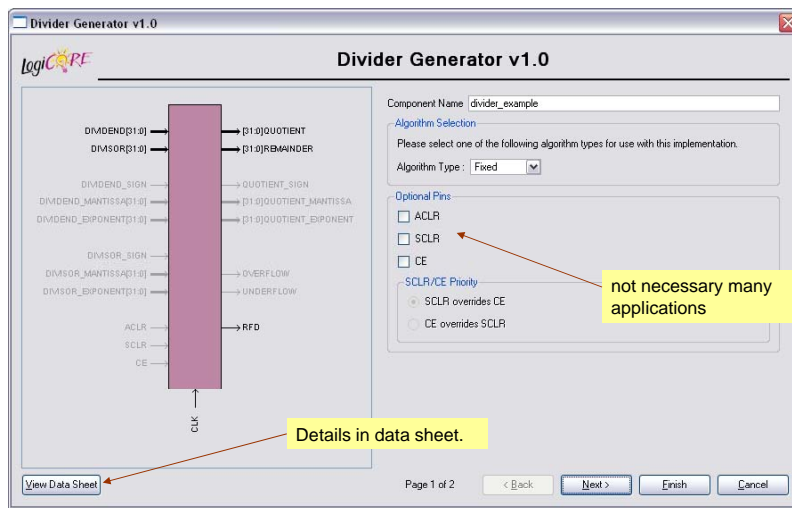
L. Williams MIT '13

Math Functions in Coregen

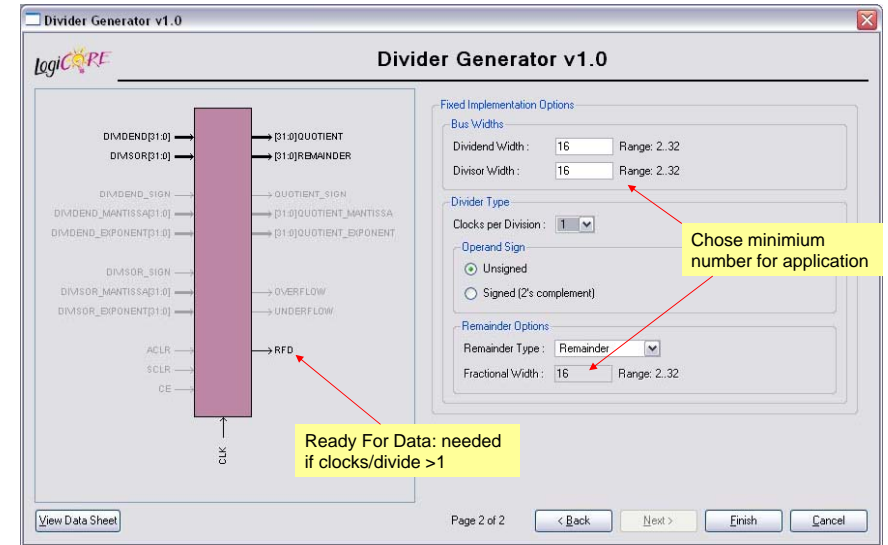


Wide selection of math functions available

Coregen Divider



Coregen Divider



Performance Metrics for Circuits

Circuit **Latency** (L): time between arrival of new input and generation of corresponding output.

For combinational circuits this is just t_{PD} .

Circuit **Throughput** (T): Rate at which new outputs appear.

For combinational circuits this is just $1/t_{PD}$ or $1/L$.

Coregen Divider Latency




Figure 2: Latency Example (Clocks per Division = 4)

Table 4: Latency of Fixed-point Solution Based on Divider Parameters


Signed	Fractional	Ckls/Div	Latency
False	False	1	M+2
False	False	>1	M+3
False	True	1	M+F+2
False	True	>1	M+F+3
True	False	1	M+4
True	False	>1	M+5
True	True	1	M+F+4
True	True	>1	M+F+5

Note: M=dividend width, F=fractional remainder width.

The divclk_sel parameter allows a range of choices of throughput versus area. With divclk_sel = 1, the core is fully pipelined, so it will have maximal throughput of one division per clock cycle, but will occupy the most area. The divclk_sel selections of 2, 4 and 8 reduce the throughput by those respective factors for smaller core sizes.

Latency dependent on dividend width + fractional remainder width

Performance of Combinational Circuits




For combinational logic:

$$L = t_{PD}$$

$$T = 1/t_{PD}$$

We can't get the answer faster, but are we making effective use of our hardware at all times?




F & G are "idle", just holding their outputs stable while H performs its computation


Retiming: A very useful transform

Retiming is the action of moving registers around in the system

- Registers have to be moved from ALL inputs to ALL outputs or vice versa



Cutset retiming: A cutset intersects the edges, such that this would result in two disjoint partitions of the edges being cut. To retime, delays are moved from the ingoing to the outgoing edges or vice versa.




Benefits of retiming:

- Modify critical path delay
- Reduce total number of registers



Retiming Combinational Circuits aka "Pipelining"



$$L = 45$$

$$T = 1/45$$


Assuming ideal registers:
i.e., $t_{PD} = 0$, $t_{SETUP} = 0$

$$t_{CLK} = 25$$

$$L = 2 * t_{CLK} = 50$$

$$T = 1/t_{CLK} = 1/25$$

Pipeline diagrams



	Clock cycle			
	i	i+1	i+2	i+3
Input	X_i	X_{i+1}	X_{i+2}	X_{i+3}
F Reg		$F(X_i)$	$F(X_{i+1})$	$F(X_{i+2})$
G Reg		$G(X_i)$	$G(X_{i+1})$	$G(X_{i+2})$
H Reg			$H(X_i)$	$H(X_{i+1})$

The results associated with a particular set of input data moves *diagonally* through the diagram, progressing through one pipeline stage each clock cycle.

Pipeline Conventions

DEFINITION:

a *K-Stage Pipeline* ("K-pipeline") is an acyclic circuit having exactly K registers on *every* path from an input to an output.

a **COMBINATIONAL CIRCUIT** is thus an 0-stage pipeline.

CONVENTION:

Every pipeline stage, hence every K-Stage pipeline, has a register on its *OUTPUT* (not on its input).

ALWAYS:


The **CLOCK** common to all registers must have a period sufficient to cover propagation over combinational paths PLUS (input) register t_{PD} PLUS (output) register t_{SETUP} .

The **LATENCY** of a K-pipeline is K times the period of the clock common to all registers.

The **THROUGHPUT** of a K-pipeline is the frequency of the clock.

Ill-formed pipelines

Consider a BAD job of pipelining:



For what value of K is the following circuit a K-Pipeline? _____ none

Problem:

Successive inputs get mixed e.g., $B(A(X_{i+1}), Y_i)$. This happened because some paths from inputs to outputs have 2 registers, and some have only 1!

This **CAN'T HAPPEN** on a well-formed K pipeline!

A pipelining methodology

Step 1:

Add a register on each output.


Step 2:

Add another register on each output. Draw a cut-set contour that includes all the new registers and some part of the circuit. Retime by moving regs from all outputs to all inputs of cut-set.


Repeat until satisfied with T.

STRATEGY:

Focus your attention on placing pipelining registers around the slowest circuit elements (**BOTTLENECKS**).



Pipeline Example



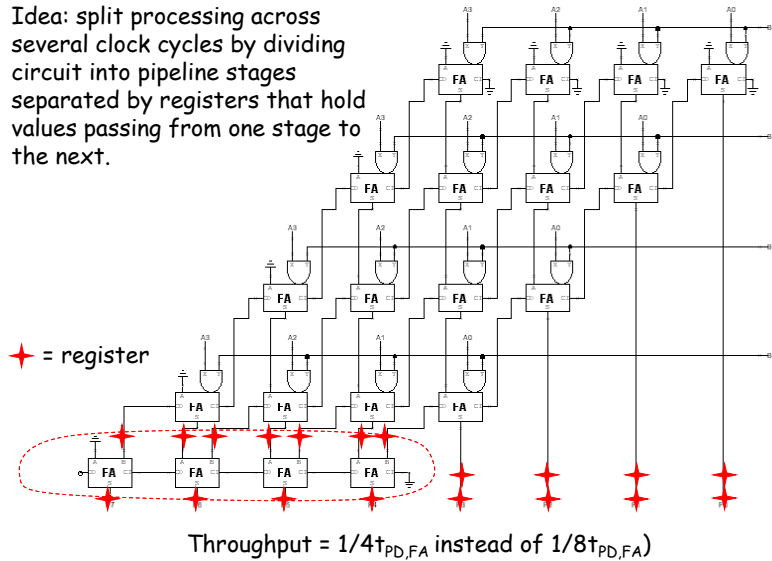
OBSERVATIONS:

- 1-pipeline improves neither L or T.
- T improved by breaking long combinational paths, allowing faster clock.
- Too many stages cost L, don't improve T.
- Back-to-back registers are often required to keep pipeline well-formed.

	LATENCY	THROUGHPUT
0-pipe:	4	1/4
1-pipe:	4	1/4
2-pipe:	4	1/2
3-pipe:	6	1/2

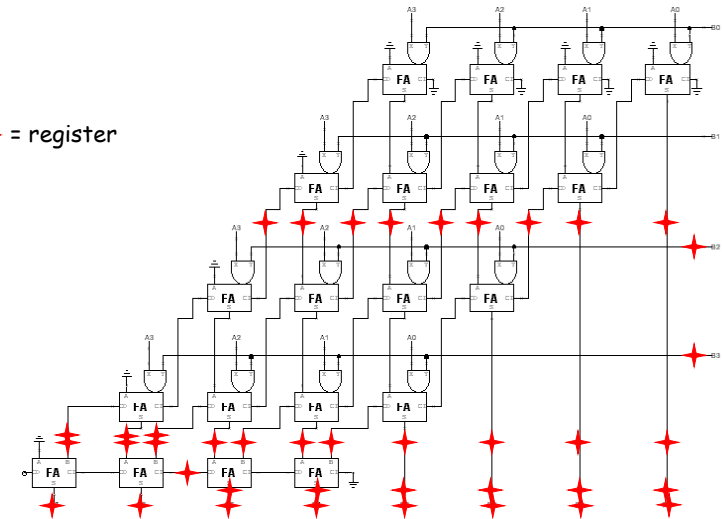
Increasing Throughput: Pipelining

Idea: split processing across several clock cycles by dividing circuit into pipeline stages separated by registers that hold values passing from one stage to the next.



How about $t_{PD} = 1/2t_{PD,FA}$?

★ = register



Timing Reports

Synthesis report

Total Propagation delay: 34.8ns

65mhz = 27mhz * 2.4

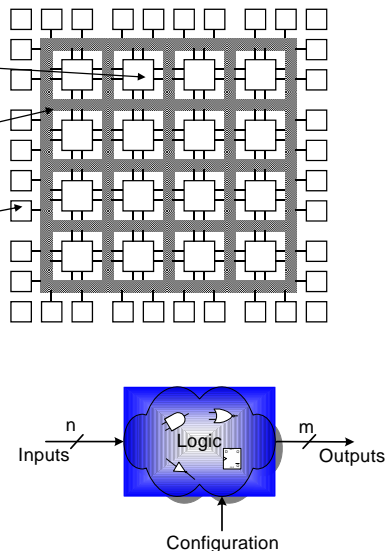
Multiple: 7.251ns

History of Computational Fabrics

- Discrete devices: relays, transistors (1940s-50s)
- Discrete logic gates (1950s-60s)
- Integrated circuits (1960s-70s)
 - e.g. TTL packages: Data Book for 100's of different parts
- Gate Arrays (IBM 1970s)
 - Transistors are pre-placed on the chip & Place and Route software puts the chip together automatically – only program the interconnect (mask programming)
- Software Based Schemes (1970's- present)
 - Run instructions on a general purpose core
- Programmable Logic (1980's to present)
 - A chip that be reprogrammed after it has been fabricated
 - Examples: **PALs**, EPROM, EEPROM, PLDs, **FPGAs**
 - Excellent support for mapping from Verilog
- ASIC Design (1980's to present)
 - Turn Verilog directly into layout using a library of standard cells
 - Effective for high-volume and efficient use of silicon area

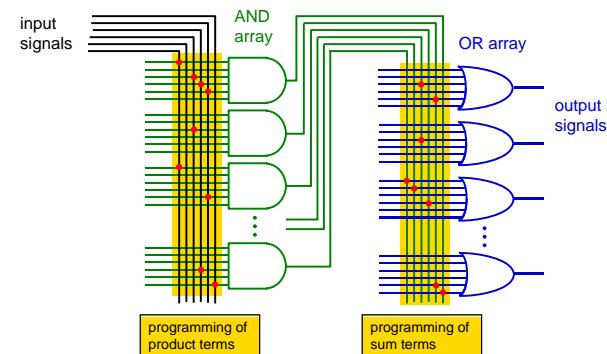
Reconfigurable Logic

- **Logic blocks**
 - To implement combinational and sequential logic
- **Interconnect**
 - Wires to connect inputs and outputs to logic blocks
- **I/O blocks**
 - Special logic blocks at periphery of device for external connections
- **Key questions:**
 - How to make logic blocks programmable? (after chip has been fabbed!)
 - What should the logic granularity be?
 - How to make the wires programmable? (after chip has been fabbed!)
 - Specialized wiring structures for local vs. long distance routes?
 - How many wires per logic block?




Programmable Array Logic (PAL)

- Based on the fact that any combinational logic can be realized as a sum-of-products
- PALs feature an array of AND-OR gates with programmable interconnect

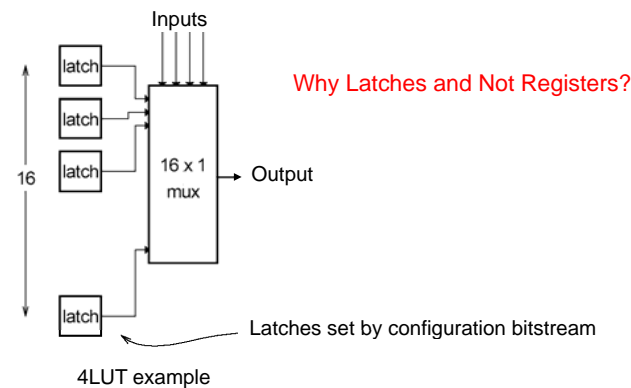


RAM Based Field Programmable Logic - Xilinx



LUT Mapping

- N-LUT direct implementation of a truth table: any function of n-inputs.
- N-LUT requires 2^N storage elements (latches)
- N-inputs select one latch location (like a memory)



Configuring the CLB as a RAM

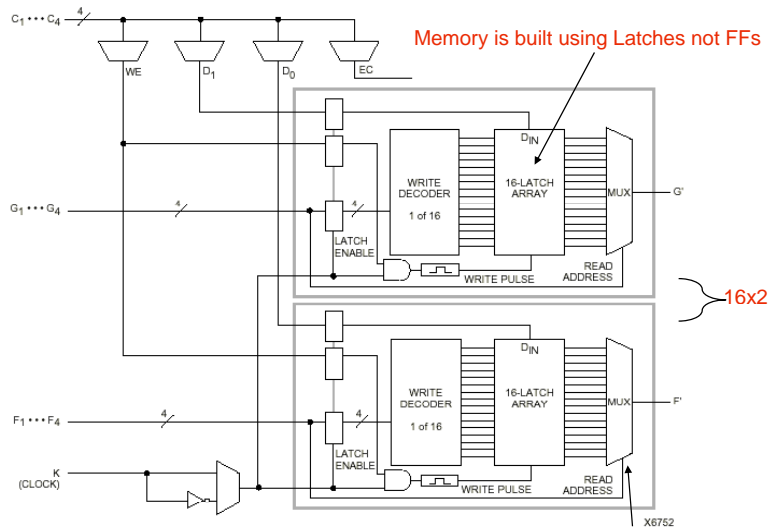


Figure 4: 16x2 (or 16x1) Edge-Triggered Single-Port RAM. Read is same a LUT Function!

Xilinx 4000 Interconnect

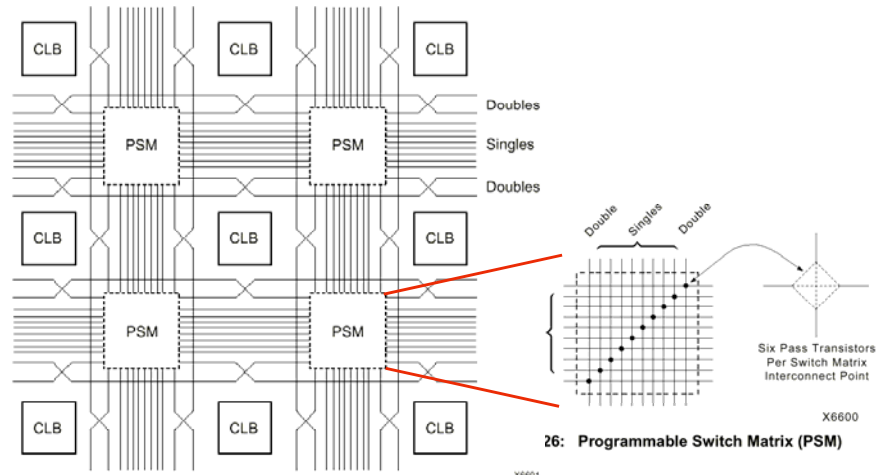
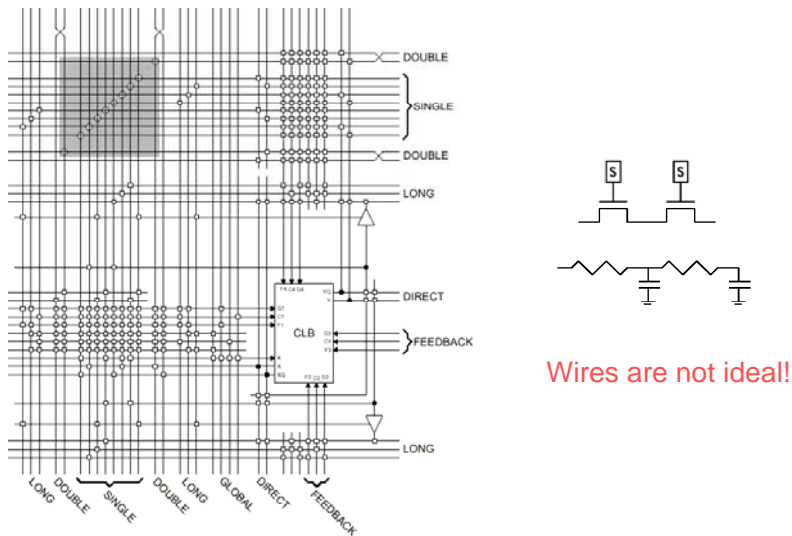



Figure 28: Single- and Double-Length Lines, with Programmable Switch Matrices (PSMs)

Xilinx 4000 Interconnect Details

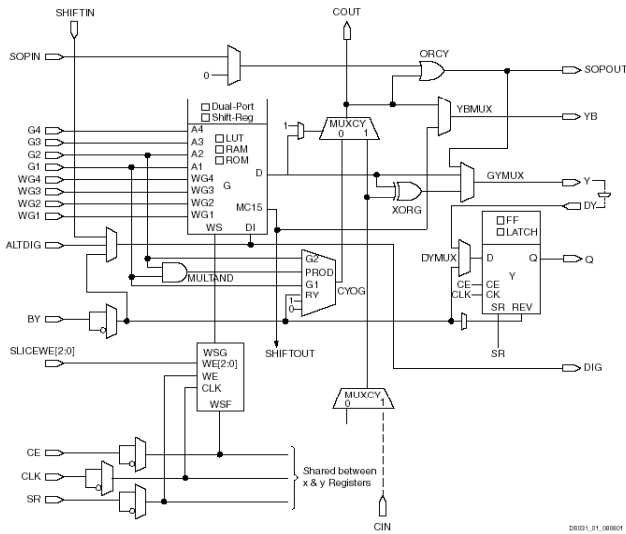


Add Bells & Whistles

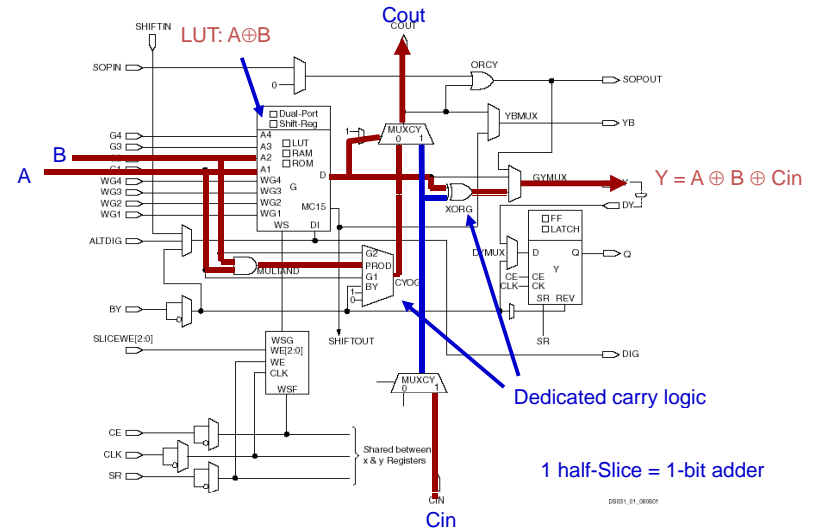


Courtesy of David B. ParLOUR, ISSCC 2004 Tutorial, "The Reality and Promise of Reconfigurable Computing in Digital Signal Processing"

The Virtex II CLB (Half Slice Shown)



Adder Implementation



Virtex-6



DSP with 25x18 multiplier

Gigabit ethernet support

Part Number	LX79T	LX130T	LX190T	LX240T	LX360T	LX550T	LX760	SX315T	SX470T
Logic Cells	74.5K	126K	200K	241K	364K	555K	756K	915K	476K
CLB Flip-Flops	93.1K	160K	250K	301K	455K	687K	948K	904K	556K
Maximum Distributed RAM (Kbits)	1,045	1,740	3,040	3,650	4,130	6,500	8,580	5,000	7,640
Block RAM/PIPO w/ ECC (8Kbits each)	156	264	344	416	416	692	720	704	1,084
Total Block RAM (Kbits)	5,816	9,504	12,984	14,076	14,076	22,752	25,920	25,344	38,204
Mixed Mode Clock Managers (MMCM)	5	10	10	12	12	18	18	12	18
DSP48E1 Slices	288	480	640	768	576	864	864	1,344	2,016
PCI Express® Interface Blocks	1	2	2	2	2	2	0	2	2
10/100/1000 Ethernet MAC Blocks	4	4	4	4	4	4	0	4	4
4096 User Programmable Transceivers	12	20	30	24	24	40	0	24	40
Package	Area (Pitch)		Maximum User I/O: Select I/O Interface Pins (GTX Transceivers)						
FF484	29 x 29 mm (1.0 mm)	240 (8)	240 (8)						
FF734	29 x 29 mm (1.0 mm)	360 (12)	400 (12)	400 (12)					
FF1156	35 x 35 mm (1.0 mm)	600 (20)	600 (20)	600 (20)	600 (20)				
FF1759	42.5 x 42.5 mm (1.0 mm)			720 (24)	720 (24)	840 (28)		720 (24)	840 (28)
FF1760	42.5 x 42.5 mm (1.0 mm)					1,200 (30)	1,200 (30)		

* Preliminary product information, subject to change. Please contact your Xilinx representative for the latest information.

	CLB	Dist RAM	Block RAM	Multipliers
Virtex 2	8,448	1,056kbit	2,592kbit	144 (18x18)
Virtex 6	667,000	6,200kbit	22,752kbit	1,344 (25x18)
Spartan 3E	240	15kbit	72kbit	4 (18x18)

Design Flow - Mapping

- Technology Mapping: Schematic/HDL to Physical Logic units
- Compile functions into basic LUT-based groups (function of target architecture)




```


always @(posedge clock or negedge reset)
begin
  if (! reset)
    q <= 0;
  else
    q <= (a & b & c) | (b & d);
end
    
```


Design Flow - Placement & Route

- **Placement** - assign logic location on a particular device



- **Routing** - iterative process to connect CLB inputs/outputs and IOBs. Optimizes critical path delay - can take hours or days for large, dense designs





Iterate placement if timing not met
Satisfy timing? → Generate Bitstream to config device

Challenge! Cannot use full chip for reasonable speeds (wires are not ideal).
Typically no more than 50% utilization.


Example: Verilog to FPGA

```
module adder64 (
  input [63:0] a, b;
  output [63:0] sum);

  assign sum = a + b;
endmodule
```





64-bit Adder Example



Virtex II - XC2V2000

How are FPGAs Used?

Logic Emulation



FPGA-based Emulator
(courtesy of IKOS)

- **Prototyping**
 - Ensemble of gate arrays used to emulate a circuit to be manufactured
 - Get more/better/faster debugging done than with simulation
- **Reconfigurable hardware**
 - One hardware block used to implement more than one function
- **Special-purpose computation engines**
 - Hardware dedicated to solving one problem (or class of problems)
 - Accelerators attached to general-purpose computers (e.g., in a cell phone!)

Summary

- FPGA provide a flexible platform for implementing digital computing
- A rich set of macros and I/Os supported (multipliers, block RAMS, ROMS, high-speed I/O)
- A wide range of applications from prototyping (to validate a design before ASIC mapping) to high-performance spatial computing
- Interconnects are a major bottleneck (physical design and locality are important considerations)

Lab 4 Car Alarm Design Approach

- Read lab/specifications carefully, use reasonable interpretation
- Use modular design - don't put everything into labkit.v
- Design the FSM!
 - Define the inputs
 - Define the outputs
 - Transition rules
- Logical modules:
 - fsm.v
 - timer.v
 - siren.v
 - fuel_pump.v
- Run simulation on each module!
- Use hex display: show state and time
- Use logic analyzer

Car Alarm - Inputs & Outputs





Figure 1: System diagram showing sensors (inputs) and actuators (outputs)

Car Alarm - CMOS Implementation



- Design Specs
 - Operating voltage 8-18VDC
 - Operating temp: -40C +65C
 - Attitude: sea level
 - Shock/Vibration
- Notes
 - Protected against 24V power surges
 - CMOS implementation
 - CMOS inputs protected against 200V noise spikes
 - On state DC current <10ma
 - Include T_PASSENGER_DELAY and Fuel Pump Disable
 - First car was stolen in Cambridge
 - System disabled (cloaked) when being serviced.

Debugging Hints - Lab 4

- Implement a warp speed debug mode for the one_hz clock. This will allow for viewing signals on the logic analyzer or Modelsim without waiting for 27 million clock cycles. Avoids recompilations.

```

assign debug_on = switch[6]; // switch[6] is not used
always @ (posedge clk) begin
    if (count == (debug_on ? 3 : 26_999_999)) count <= 0;
    else count <= count +1;
end

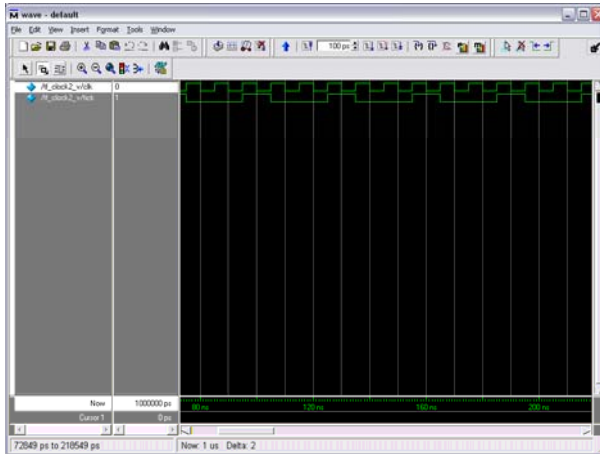
assign one_hz = (count == (debug_on ? 3 : 26_999_999)) ;
    
```

One Hz Ticks in Modelsim

To create a one hz tick, use the following in the Verilog test fixture:

```
always #5 clk=clk;
always begin
  #5 tick = 1;
  #10 tick = 0;
  #15;
end

initial begin
  // Initialize Inputs
  clk = 0;
  tick = 0; ...
end
```



For Loops, Repeat Loops in Simulation

```
integer i; // index must be declared as integer
integer irepeat;
```

```
// this will just wait 10ns, repeated 32x.
// simulation only! Can implement #10 in hardware!
irepeat = 0;
repeat(32) begin
  #10;
  irepeat = irepeat + 1;
end
```


```
// this will wait #10ns before incrementing the for loop
for (i=0; i<16; i=i+1) begin
  #10; // wait #10 before increment.
  // @(posedge clk);
  // add to index on posedge
end
```

Edge Detection

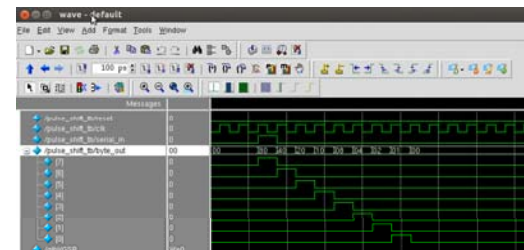
```
reg signal_delayed;
```

```
always @(posedge clk)
  signal_delayed <= signal;
```

```
assign rising_edge = signal && !signal_delayed;
assign falling_edge = !signal && signal_delayed;
```



Shift Register



```
always @(posedge clk) begin
  if (reset) byte_out <= 0;
  else byte_out <= {serial_in, byte_out[7:1]};
end
```