

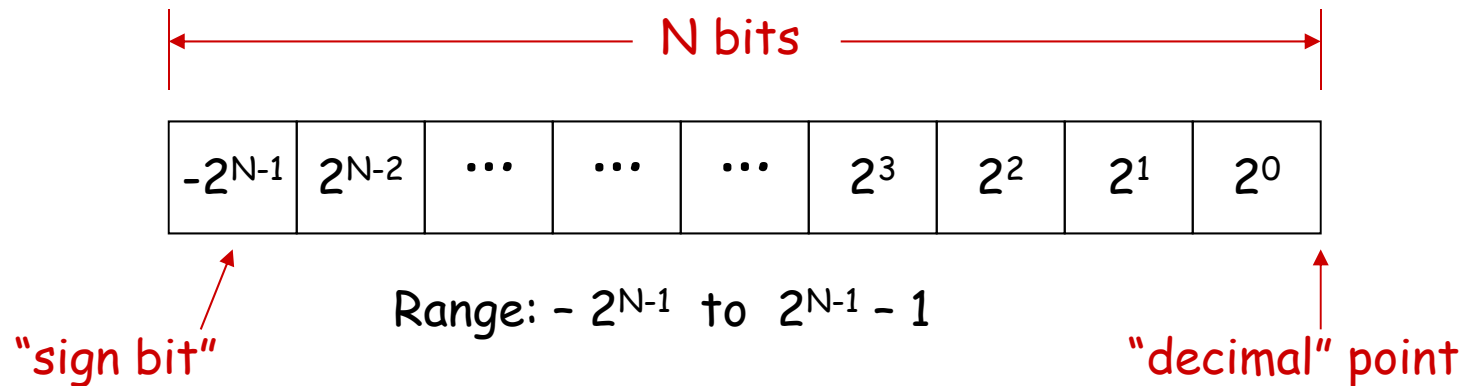


Arithmetic Circuits & Multipliers

- Addition, subtraction
- Performance issues
 - ripple carry
 - carry bypass
 - carry skip
 - carry lookahead
- Multipliers

Reminder: Lab #3 due tonight!

Signed integers: 2's complement



8-bit 2's complement example:

$$11010110 = -2^7 + 2^6 + 2^4 + 2^2 + 2^1 = -128 + 64 + 16 + 4 + 2 = -42$$

If we use a two's complement representation for signed integers, the same binary addition mod 2^n procedure will work for adding positive and negative numbers (don't need separate subtraction rules). The same procedure will also handle unsigned numbers!

By moving the implicit location of "decimal" point, we can represent fractions too:

$$1101.0110 = -2^3 + 2^2 + 2^0 + 2^{-2} + 2^{-3} = -8 + 4 + 1 + 0.25 + 0.125 = -2.625$$

Sign extension

Consider the 8-bit 2's complement representation of:

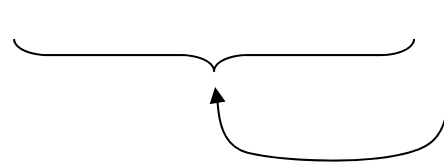
$$42 = 00101010$$

$$\begin{aligned} -5 &= \sim 00000101 + 1 \\ &= 11111010 + 1 \\ &= 11111011 \end{aligned}$$

What is their **16-bit** 2's complement representation?

$$42 = 00000000000101010$$

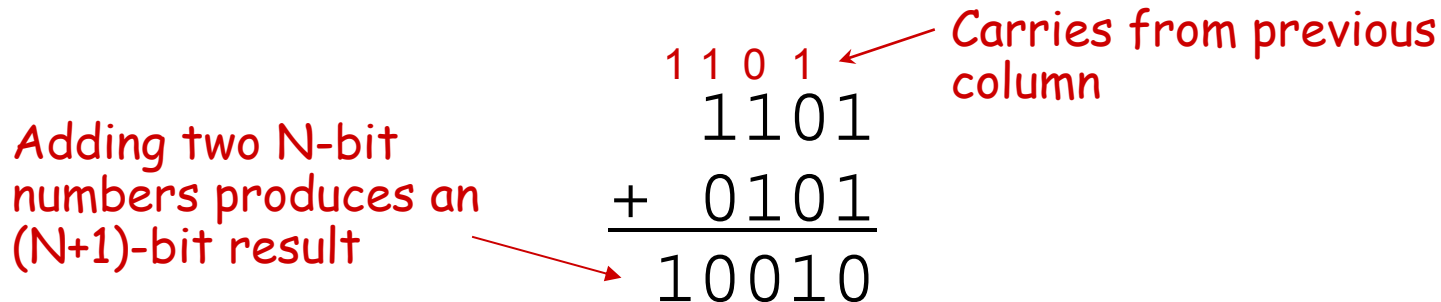
$$-5 = 1111111111111011$$



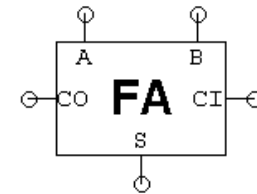
Extend the MSB (aka the "sign bit") into the higher-order bit positions

Adder: a circuit that does addition

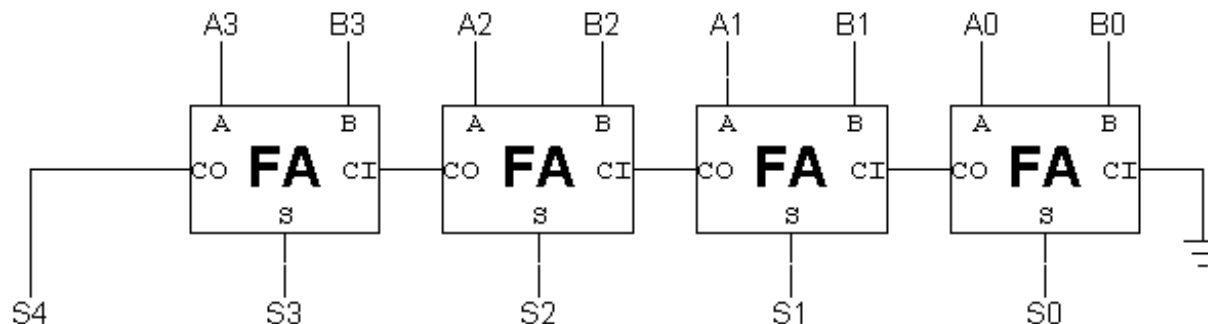
Here's an example of binary addition as one might do it by "hand":



If we build a circuit that implements one column:

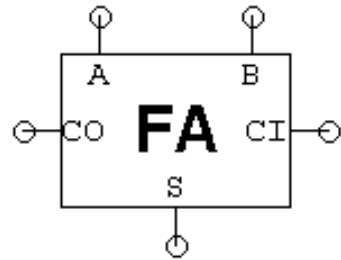


we can quickly build a circuit to add two 4-bit numbers...



“Ripple-carry adder”

“Full Adder” building block



The “half adder” circuit has only the A and B inputs



A	B	C	S	CO
0	0	0	0	0
0	0	1	1	0
0	1	0	1	0
0	1	1	0	1
1	0	0	1	0
1	0	1	0	1
1	1	0	0	1
1	1	1	1	1

$$S = A \oplus B \oplus C$$

$$\begin{aligned} CO &= \overline{A}BC + A\overline{B}C + AB\overline{C} + ABC \\ &= (\overline{A} + A)BC + (\overline{B} + B)AC + AB(\overline{C} + C) \\ &= BC + AC + AB \end{aligned}$$

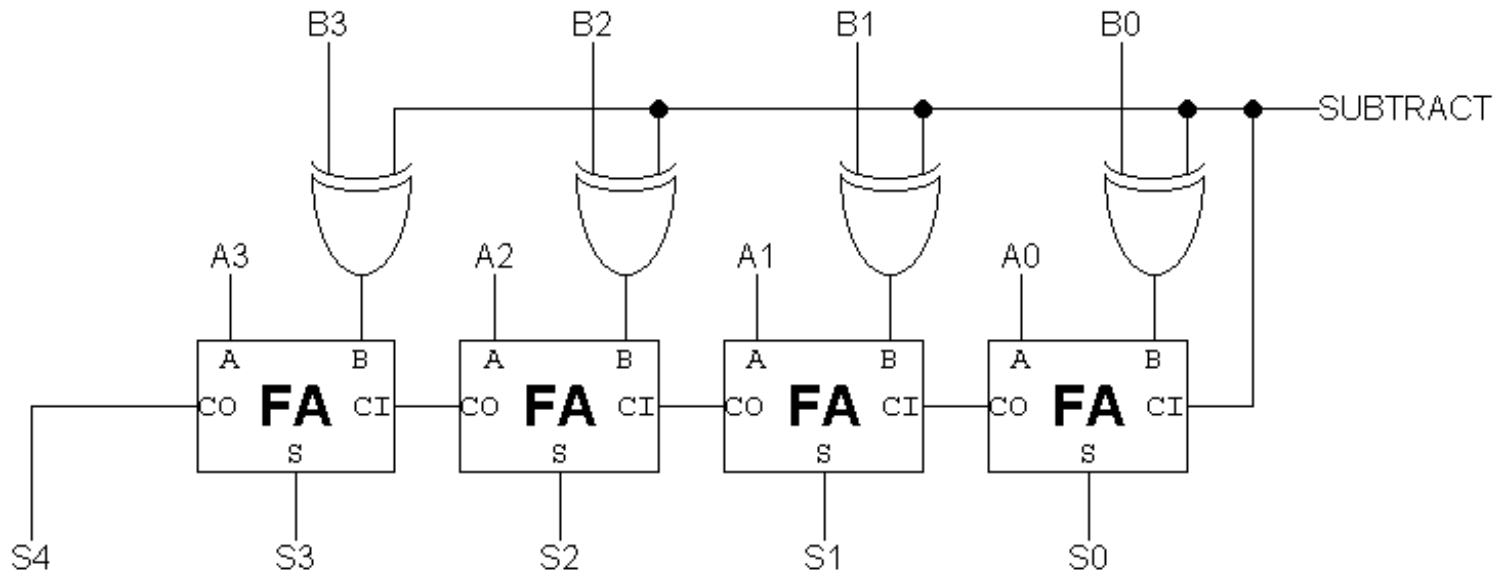
Subtraction: $A - B = A + (-B)$

Using 2's complement representation: $-B = \sim B + 1$

\sim = bit-wise complement



So let's build an arithmetic unit that does both addition and subtraction. Operation selected by *control input*.



Condition Codes

Besides the sum, one often wants four other bits of information from an arithmetic unit:

Z (zero): result is = 0 *big NOR gate*

N (negative): result is < 0 S_{N-1}

C (carry): indicates an add in the most significant position produced a carry, e.g.,
1111 + 0001 *from last FA*

V (overflow): indicates that the answer has too many bits to be represented correctly by the result width, e.g.,
0111 + 0111

$$V = A_{N-1}B_{N-1}\overline{S_{N-1}} + \overline{A_{N-1}}\overline{B_{N-1}}S_{N-1}$$

$$V = COUT_{N-1} \oplus CIN_{N-1}$$

To compare A and B, perform A-B and use condition codes:

Signed comparison:

LT $N \oplus V$

LE $Z + (N \oplus V)$

EQ Z

NE $\sim Z$

GE $\sim (N \oplus V)$

GT $\sim (Z + (N \oplus V))$

Unsigned comparison:

LTU C

LEU $C + Z$

GEU $\sim C$

GTU $\sim (C + Z)$

Condition Codes in Verilog

Z (zero): result is = 0

N (negative): result is < 0

C (carry): indicates an add in the most significant position produced a carry, e.g., 1111 + 0001

V (overflow): indicates that the answer has too many bits to be represented correctly by the result width, e.g., 0111 + 0111

```
wire signed [31:0] a,b,s;  
wire z,n,v,c;  
assign {c,s} = a + b;  
assign z = ~|s;  
assign n = s[31];  
assign v = a[31]^b[31]^s[31]^c;
```



Might be better to use sum-of-products formula for V from previous slide if using LUT implementation (only 3 variables instead of 4).

Modular Arithmetic

The Verilog arithmetic operators (+, -, *) all produce full-precision results, e.g., adding two 8-bit numbers produces a 9-bit result.

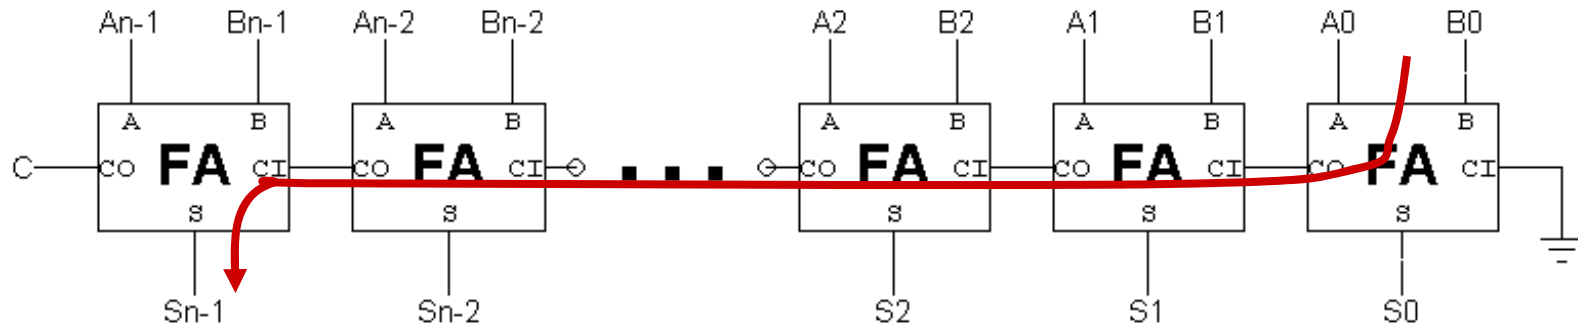
In many designs one chooses a “word size” (many computers use 32 or 64 bits) and all arithmetic results are truncated to that number of bits, i.e., arithmetic is performed modulo $2^{\text{word size}}$.

Using a fixed word size can lead to *overflow*, e.g., when the operation produces a result that's too large to fit in the word size. One can

- Avoid overflow: choose a sufficiently large word size
- Detect overflow: have the hardware remember if an operation produced an overflow - trap or check status at end
- Embrace overflow: sometimes this is exactly what you want, e.g., when doing index arithmetic for circular buffers of size 2^N .
- “Correct” overflow: replace result with most positive or most negative number as appropriate, aka *saturating arithmetic*. Good for digital signal processing.

Speed: t_{PD} of Ripple-carry Adder

$$C_O = AB + AC_I + BC_I$$



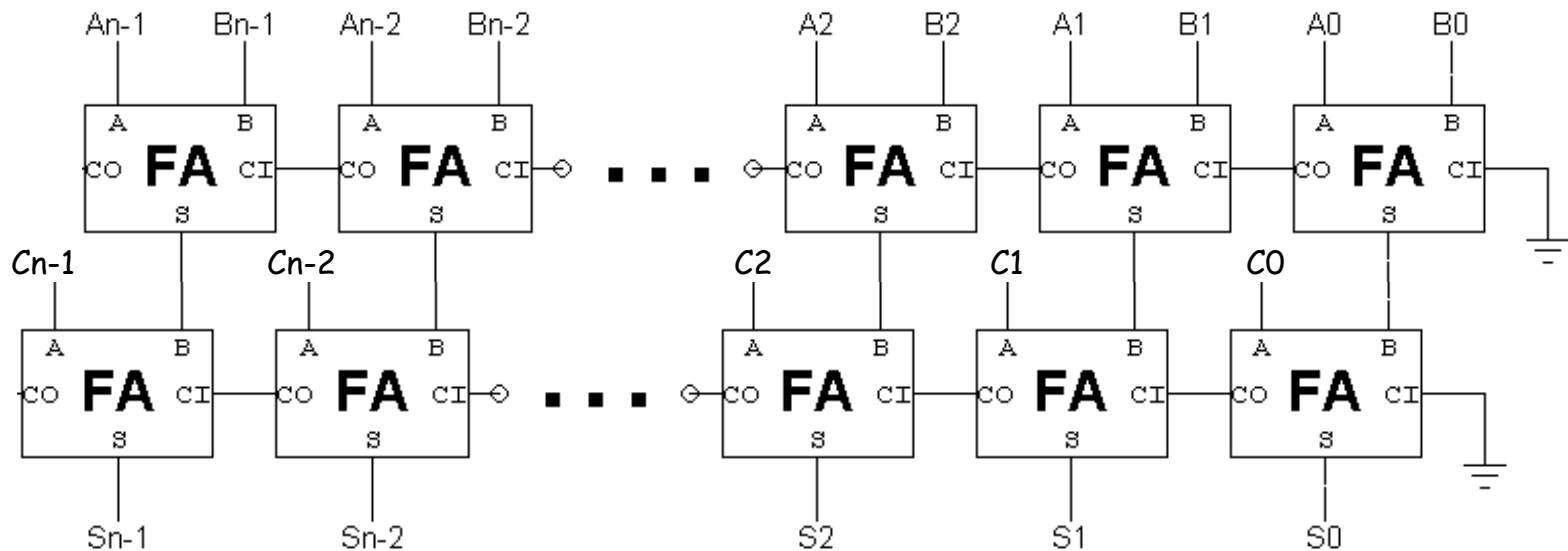
Worst-case path: carry propagation from LSB to MSB, e.g., when adding 11...111 to 00...001.

$$t_{PD} = (N-1) \underbrace{(t_{PD,OR} + t_{PD,AND})}_{CI \text{ to } CO} + \underbrace{t_{PD,XOR}}_{CI_{N-1} \text{ to } S_{N-1}} \approx \Theta(N)$$

$$t_{adder} = (N-1)t_{carry} + t_{sum}$$

$\Theta(N)$ is read “order N”: means that the latency of our adder grows at worst in proportion to the number of bits in the operands.

How about the t_{PD} of this circuit?



Is the t_{PD} of this circuit = $2 * t_{PD, N-BIT RIPPLE}$?

Nope! t_{PD} of this circuit = $t_{PD, N-BIT RIPPLE} + t_{PD, FA} !!!$

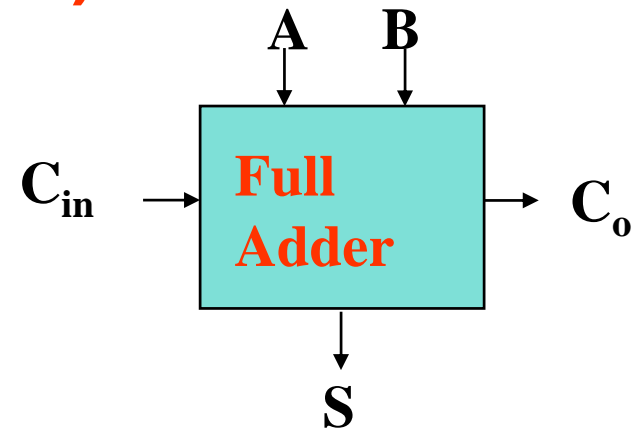


— Timing analysis is tricky!

Alternate Adder Logic Formulation

How to Speed up the Critical (Carry) Path?
(How to Build a Fast Adder?)

A	B	C_i	S	C_o	Carry status
0	0	0	0	0	delete
0	0	1	1	0	delete
0	1	0	1	0	propagate
0	1	1	0	1	propagate
1	0	0	1	0	propagate
1	0	1	0	1	propagate
1	1	0	0	1	generate
1	1	1	1	1	generate



$$\text{Generate (G)} = AB$$

$$\text{Propagate (P)} = A \oplus B$$

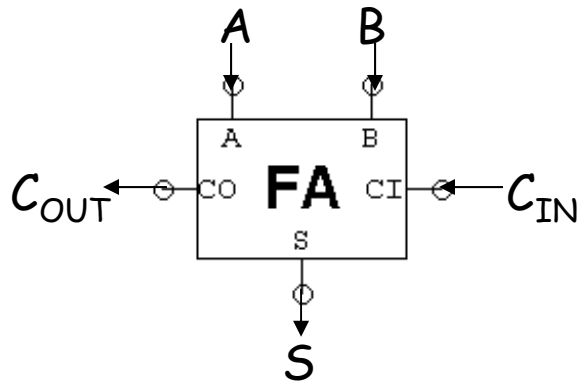
$$C_o(G, P) = G + PC_i$$

$$S(G, P) = P \oplus C_i$$

Note: can also use $P = A + B$ for C_o .

Faster carry logic

Let's see if we can improve the speed by rewriting the equations for C_{OUT} :



$$\begin{aligned}C_{OUT} &= AB + AC_{IN} + BC_{IN} \\ &= AB + (A + B)C_{IN} \\ &= G + P C_{IN}\end{aligned}$$

generate propagate

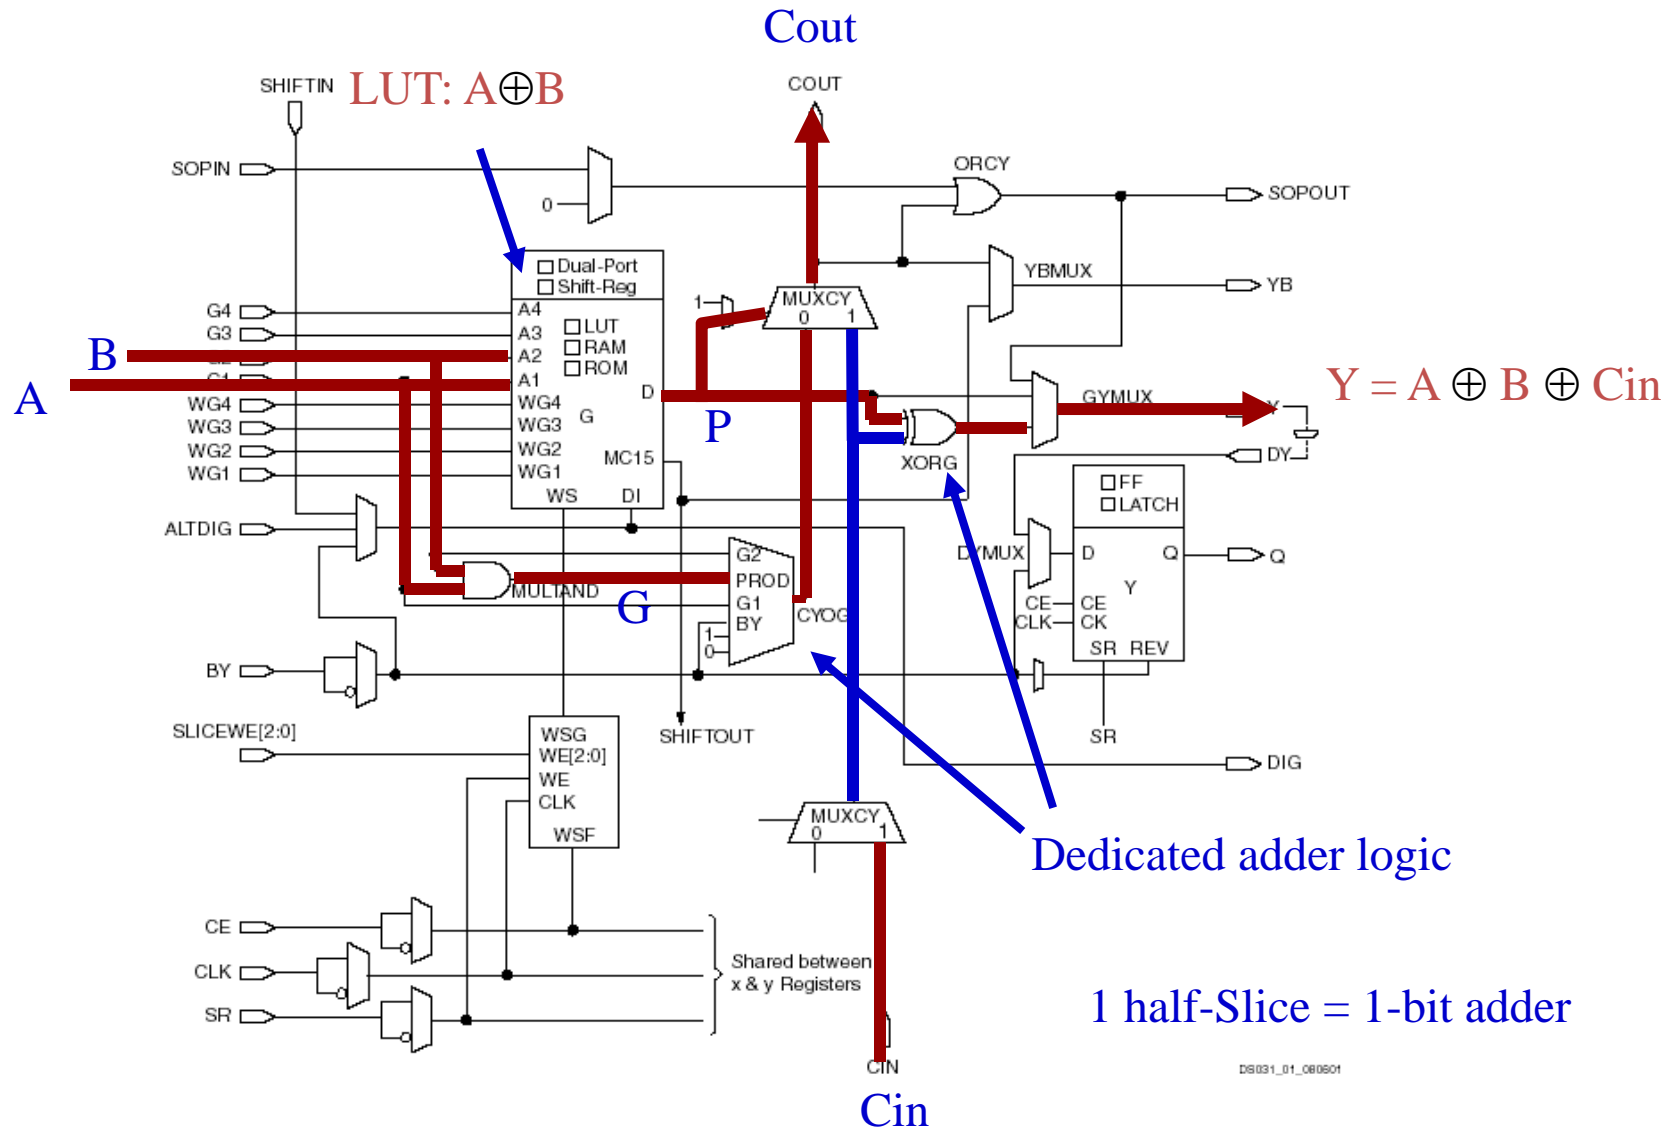
where $G = AB$
 $P = A + B$

```
module fa(input a,b,cin, output s,cout);
  wire g = a & b;
  wire p = a ^ b;
  assign s = p ^ cin;
  assign cout = g | (p & cin);
endmodule
```

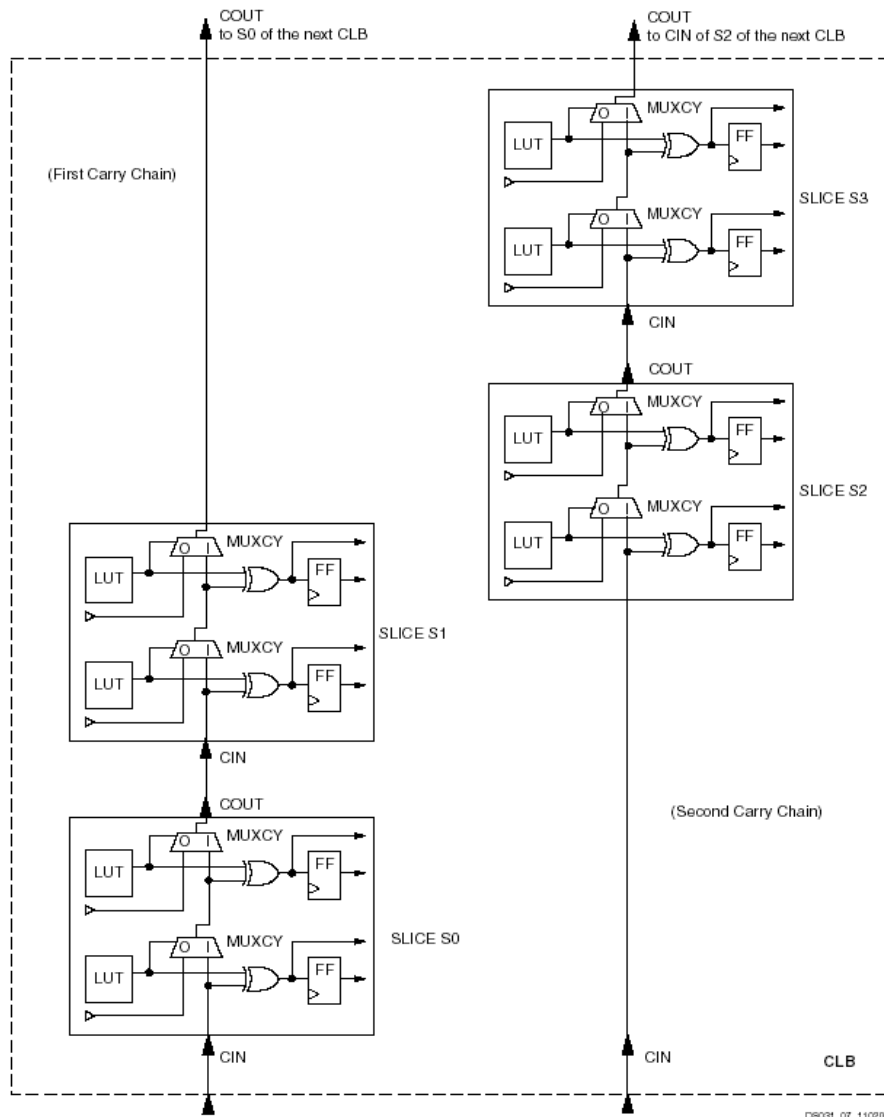
Actually, P is usually defined as $P = A \wedge B$ which won't change C_{OUT} but will allow us to express S as a simple function:

$$S = P \wedge C_{IN}$$

Virtex II Adder Implementation

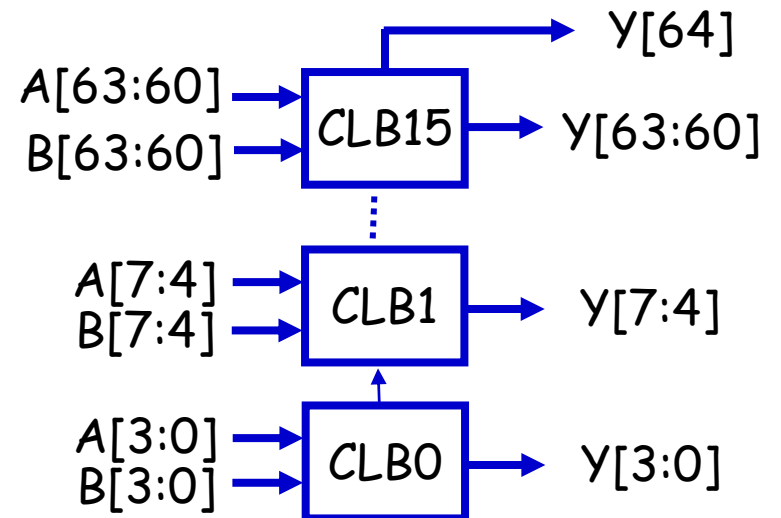
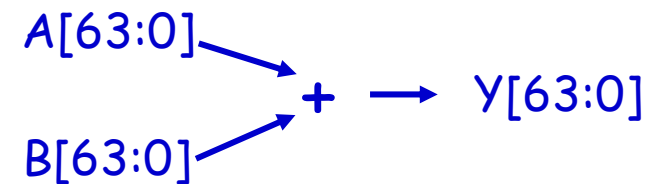


Virtex II Carry Chain



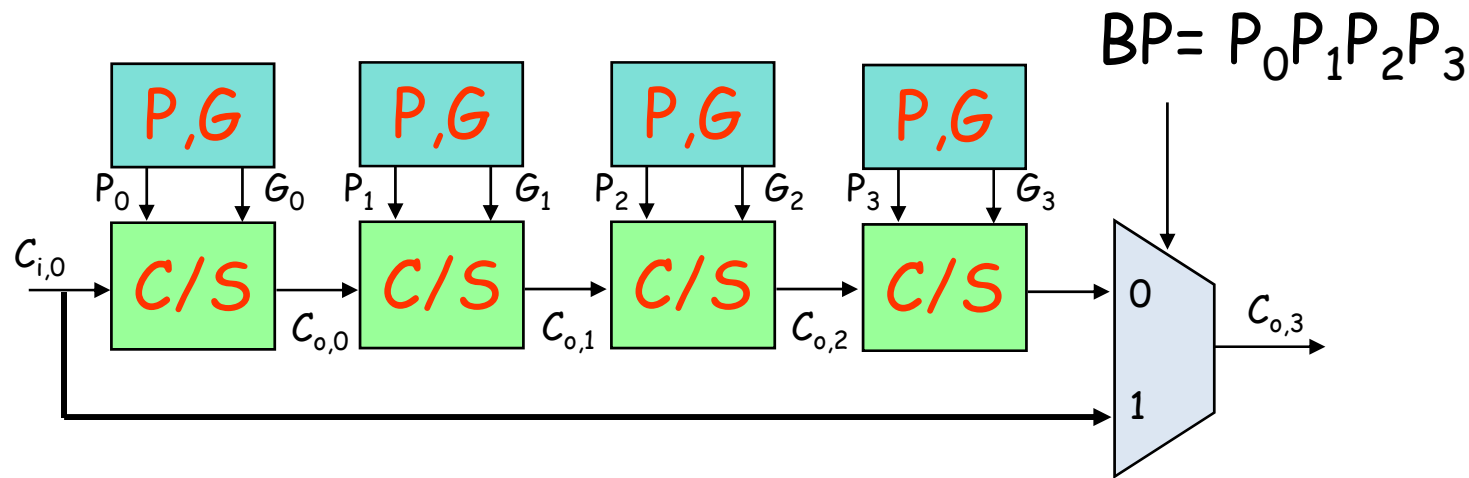
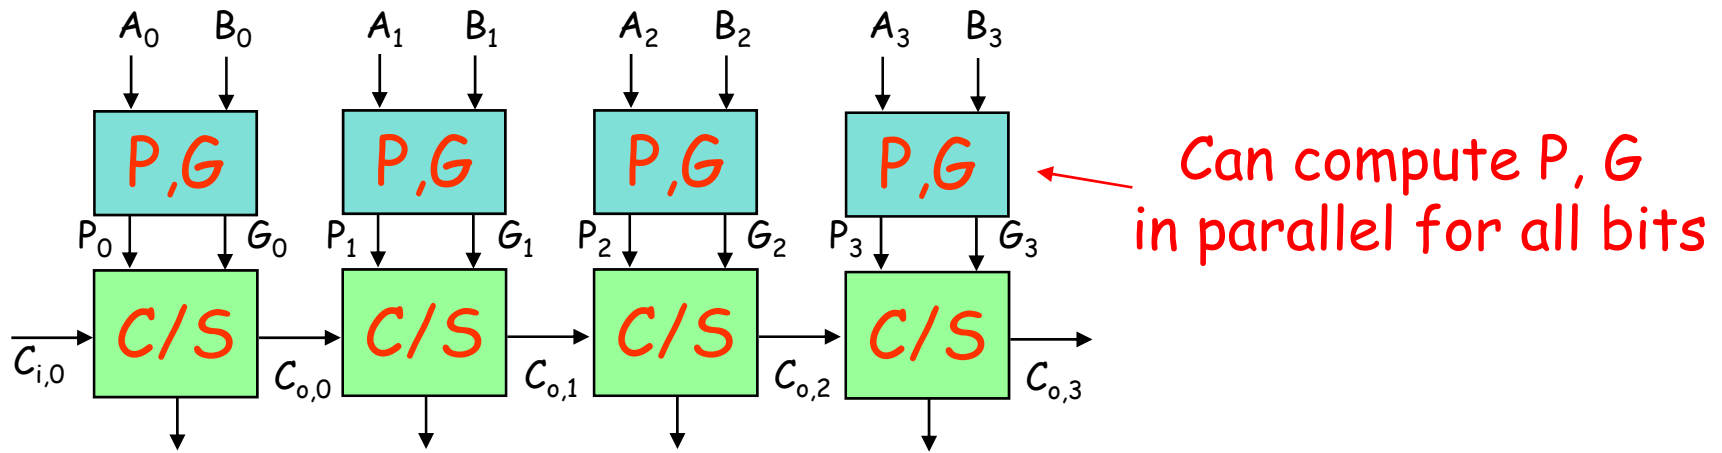
1 CLB = 4 Slices = 2, 4-bit adders

64-bit Adder: 16 CLBs



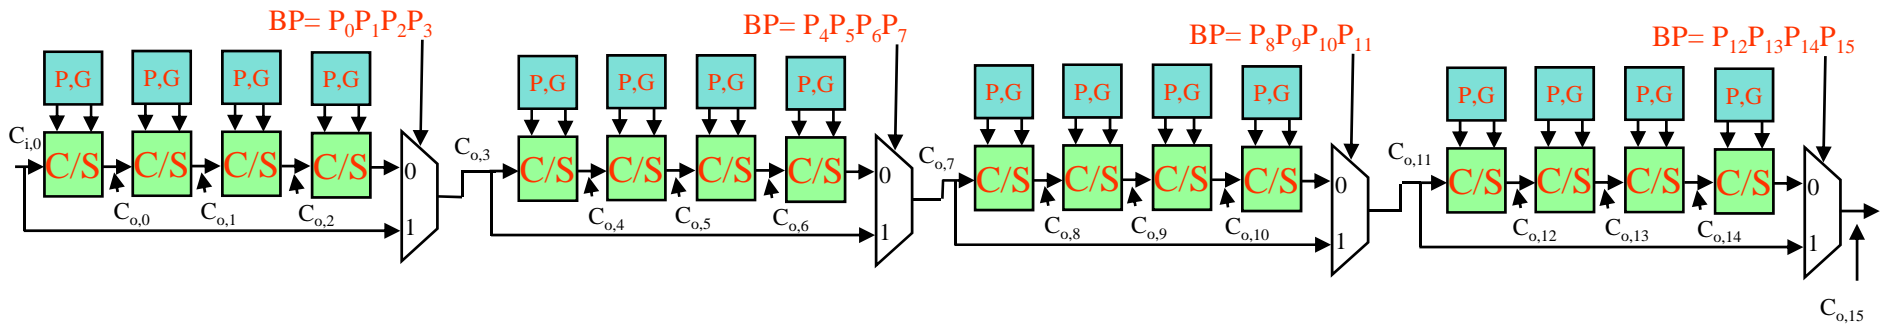
CLBs must be in same column

Carry Bypass Adder



Key Idea: if $(P_0 P_1 P_2 P_3)$ then $C_{0,3} = C_{i,0}$

16-bit Carry Bypass Adder



What is the worst case propagation delay for the 16-bit adder?

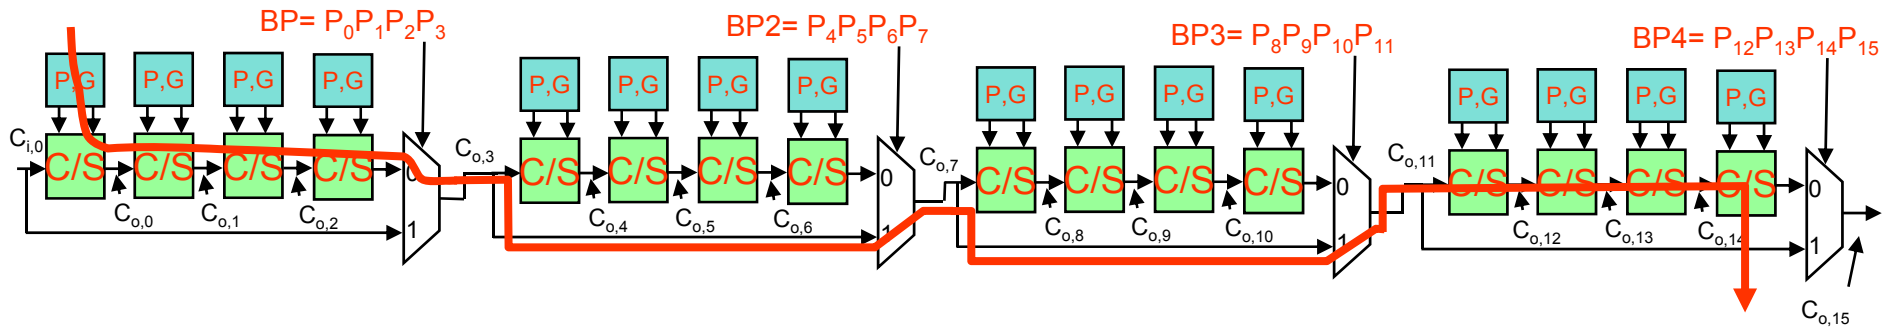
Assume the following for delay each gate:

P, G from A, B : 1 delay unit

P, G, C_i to C_o or Sum for a C/S : 1 delay unit

2:1 mux delay: 1 delay unit

Critical Path Analysis



For the second stage, is the critical path:

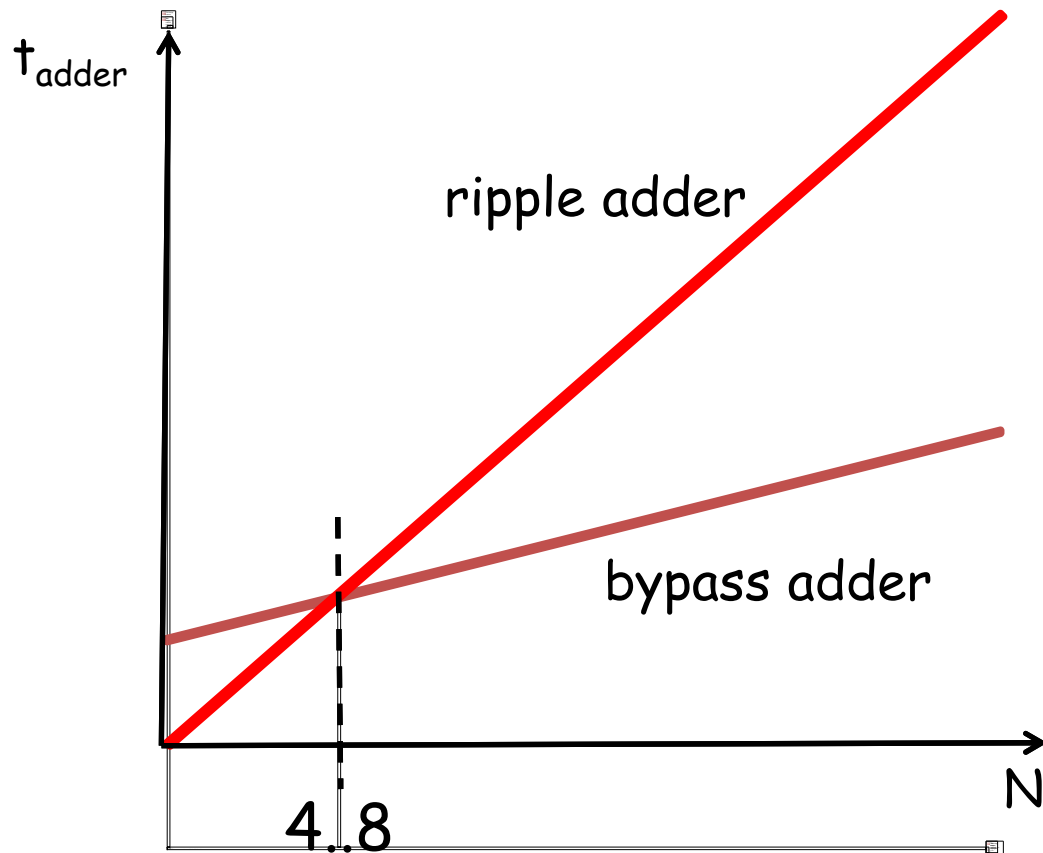
$$BP2 = 0 \quad \text{or} \quad BP2 = 1 ?$$

Message: Timing analysis is very tricky -
Must carefully consider data dependencies for false paths

Carry Bypass vs Ripple Carry

Ripple Carry: $t_{\text{adder}} = (N-1) t_{\text{carry}} + t_{\text{sum}}$

Carry Bypass: $t_{\text{adder}} = 2(M-1) t_{\text{carry}} + t_{\text{sum}} + (N/M-1) t_{\text{bypass}}$



M = bypass word size

N = number of bits being added

Carry Lookahead Adder (CLA)

- Recall that $C_{OUT} = G + P C_{IN}$ where $G = A \& B$ and $P = A \hat{=} B$
- For adding two N-bit numbers:

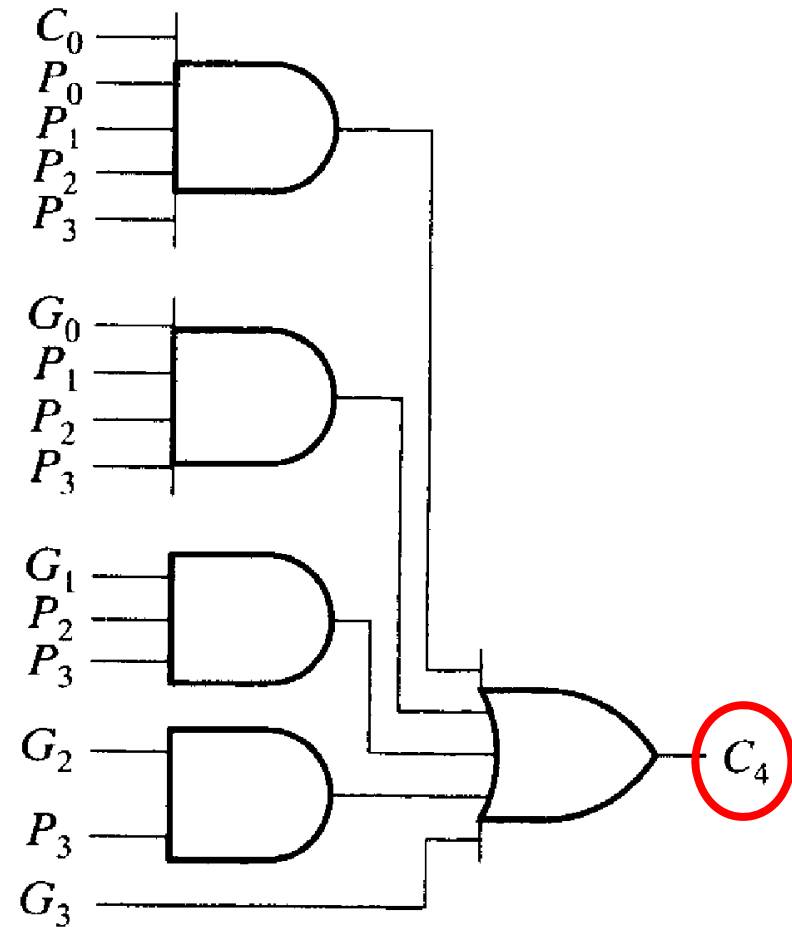
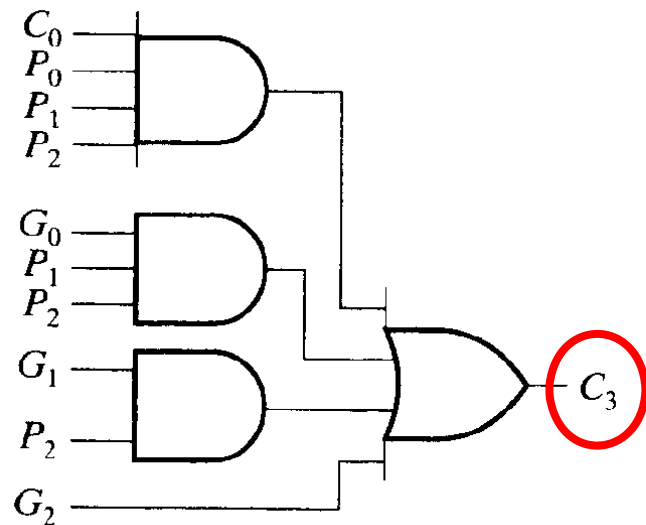
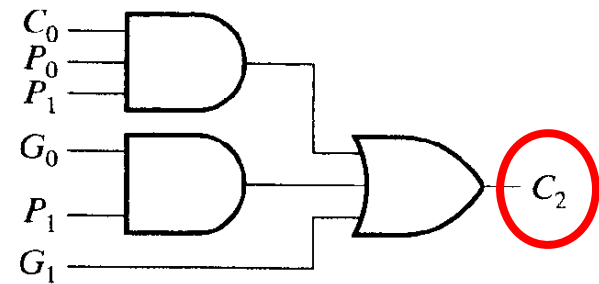
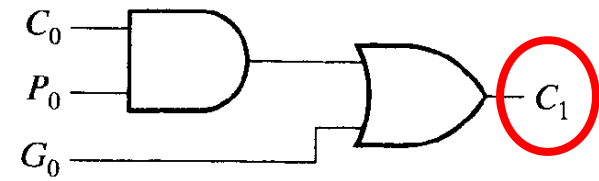
$$\begin{aligned} C_N &= G_{N-1} + P_{N-1} C_{N-1} \\ &= G_{N-1} + P_{N-1} G_{N-2} + P_{N-1} P_{N-2} C_{N-2} \\ &= G_{N-1} + P_{N-1} G_{N-2} + P_{N-1} P_{N-2} G_{N-3} + \dots + P_{N-1} \dots P_0 C_{IN} \end{aligned}$$

C_N in only 3 gate delays* :
1 for P/G generation, 1 for ANDs, 1 for final OR

*assuming gates with N inputs

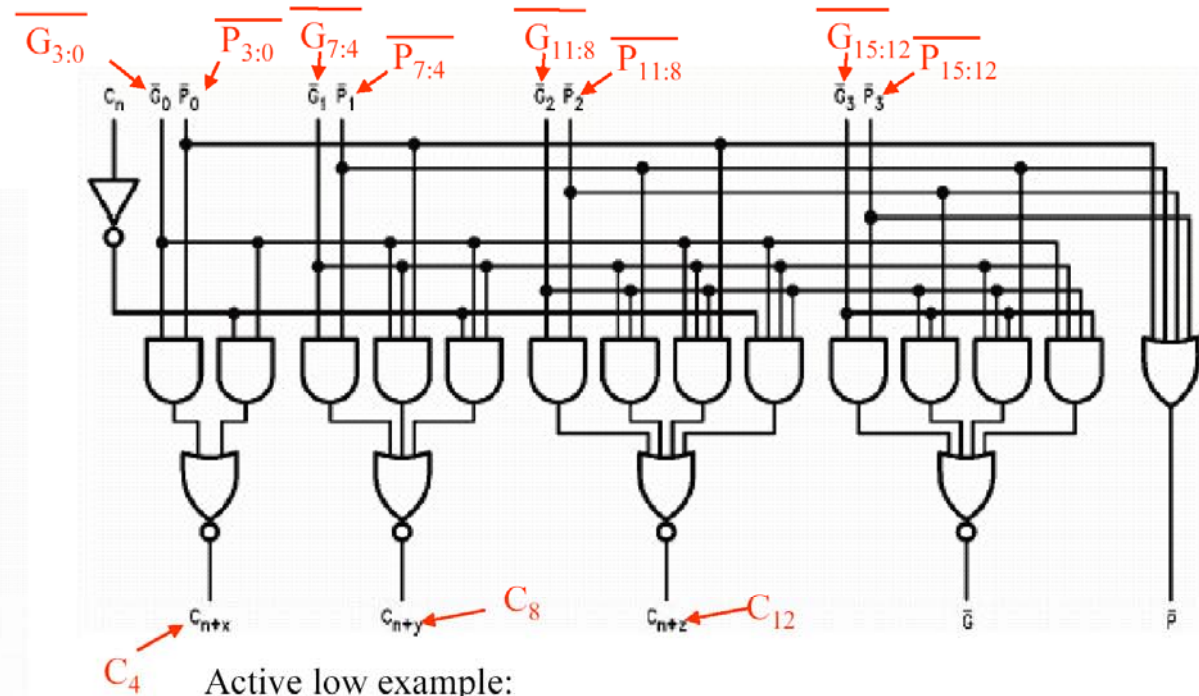
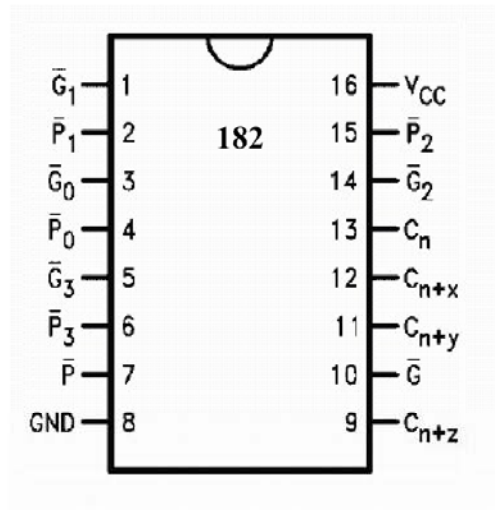
- Idea: pre-compute all carry bits as $f(Gs, Ps, C_{IN})$

Carry Lookahead Circuits



The 74182 Carry Lookahead Unit

74182 carry lookahead unit



Active low example:

$$\begin{aligned} C_{n+x} &= \overline{\overline{G_0} \cdot \overline{P_0}} + \overline{\overline{G_0} \cdot \overline{C_n}} \\ &= \overline{\overline{G_0} \cdot \overline{P_0} \cdot \overline{G_0} \cdot \overline{C_n}} \\ &= (G_0 + P_0) \cdot (G_0 + C_n) = G_0 + P_0 C_n \end{aligned}$$

$$\Rightarrow C_4 = G_{3:0} + P_{3:0} C_n$$

$$C_{n+y} = C_8 = G_{7:4} + P_{7:4} G_{3:0} + P_{7:4} P_{3:0} C_{i,0} = G_{7:0} + P_{7:0} C_n$$

$$\begin{aligned} C_{n+z} = C_{12} &= G_{11:8} + P_{11:8} G_{7:4} + P_{11:8} P_{7:4} G_{3:0} + P_{11:8} P_{7:4} P_{3:0} C_n \\ &= G_{11:0} + P_{11:0} C_n \end{aligned}$$

- high speed carry lookahead generator
- used with 74181 to extend carry lookahead beyond 4 bits
- correctly handles the carry polarity of the 181

Block Generate and Propagate

G and P can be computed for groups of bits (instead of just for individual bits). This allows us to choose the maximum fan-in we want for our logic gates and then build a hierarchical carry chain using these equations:

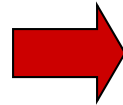
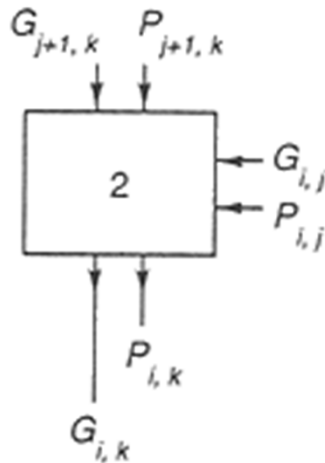
$$C_{J+1} = G_{IJ} + P_{IJ}C_I$$

$$G_{IK} = G_{J+1,K} + P_{J+1,K}G_{IJ}$$

$$P_{IK} = P_{IJ}P_{J+1,K}$$

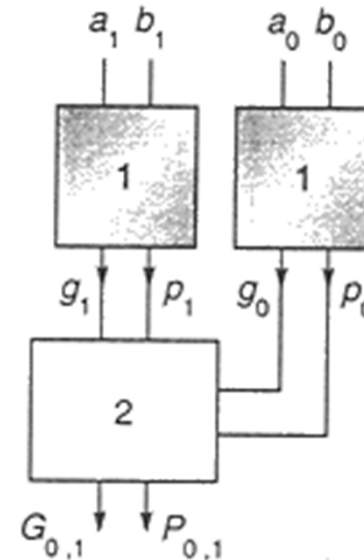
“generate a carry from bits I thru K if it is generated in the high-order ($J+1,K$) part of the block or if it is generated in the low-order (I,J) part of the block and then propagated thru the high part”

where $I < J$ and $J+1 < K$



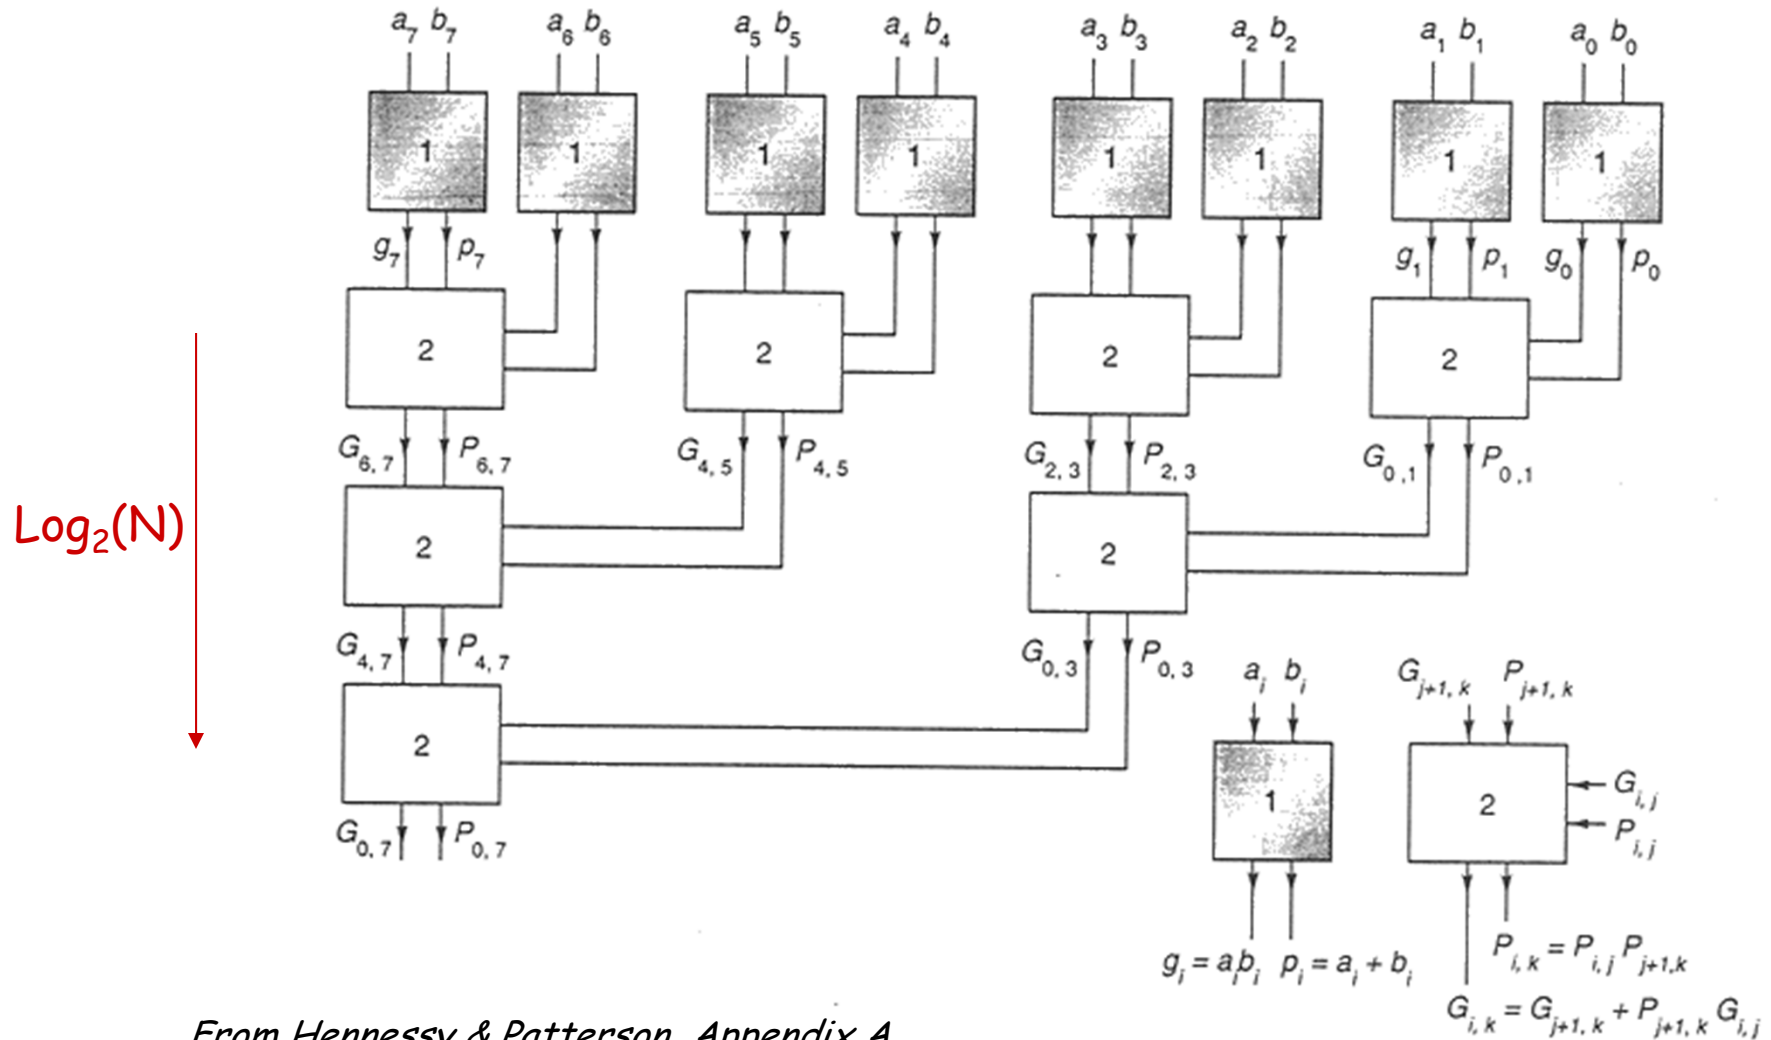
P/G generation

1st level of lookahead

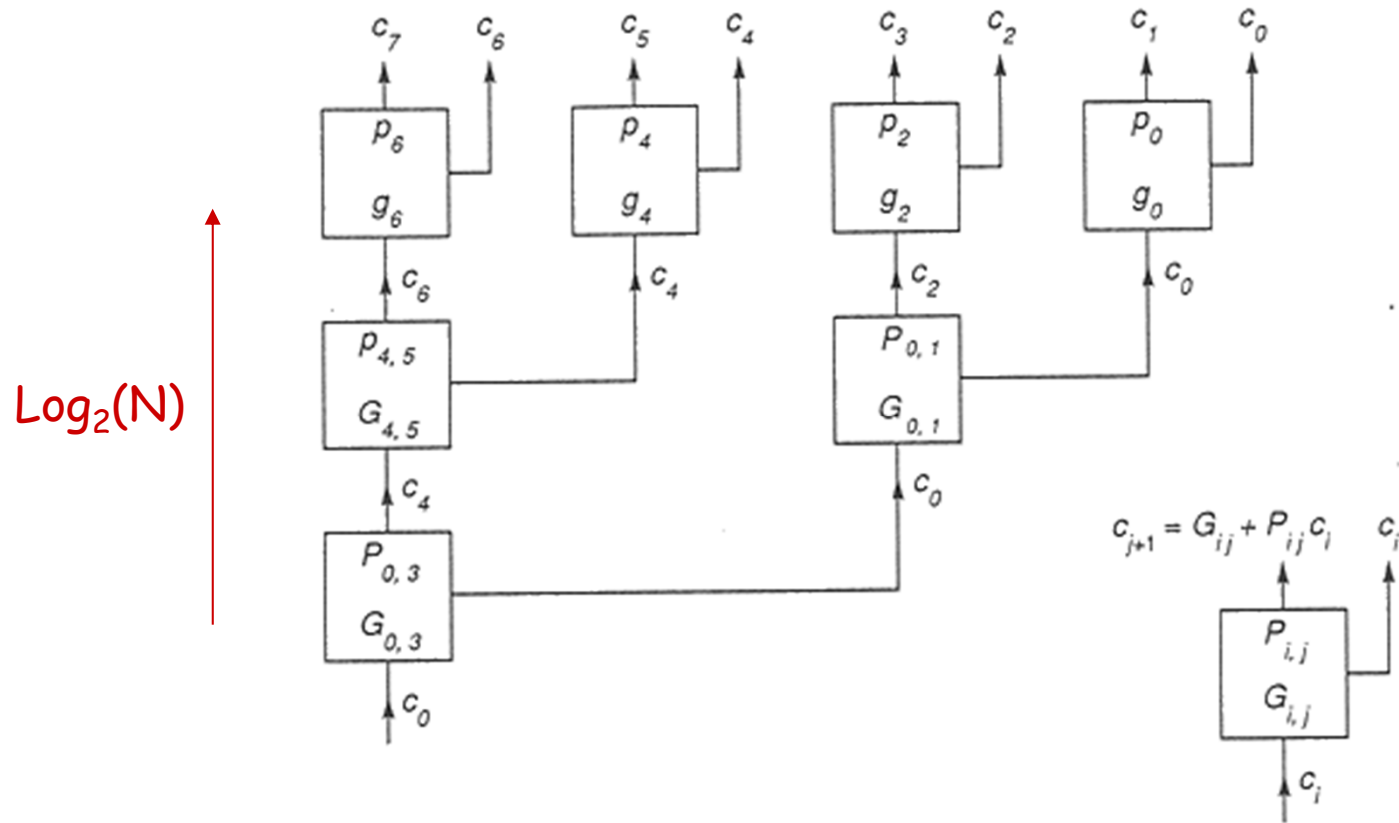


Hierarchical building block

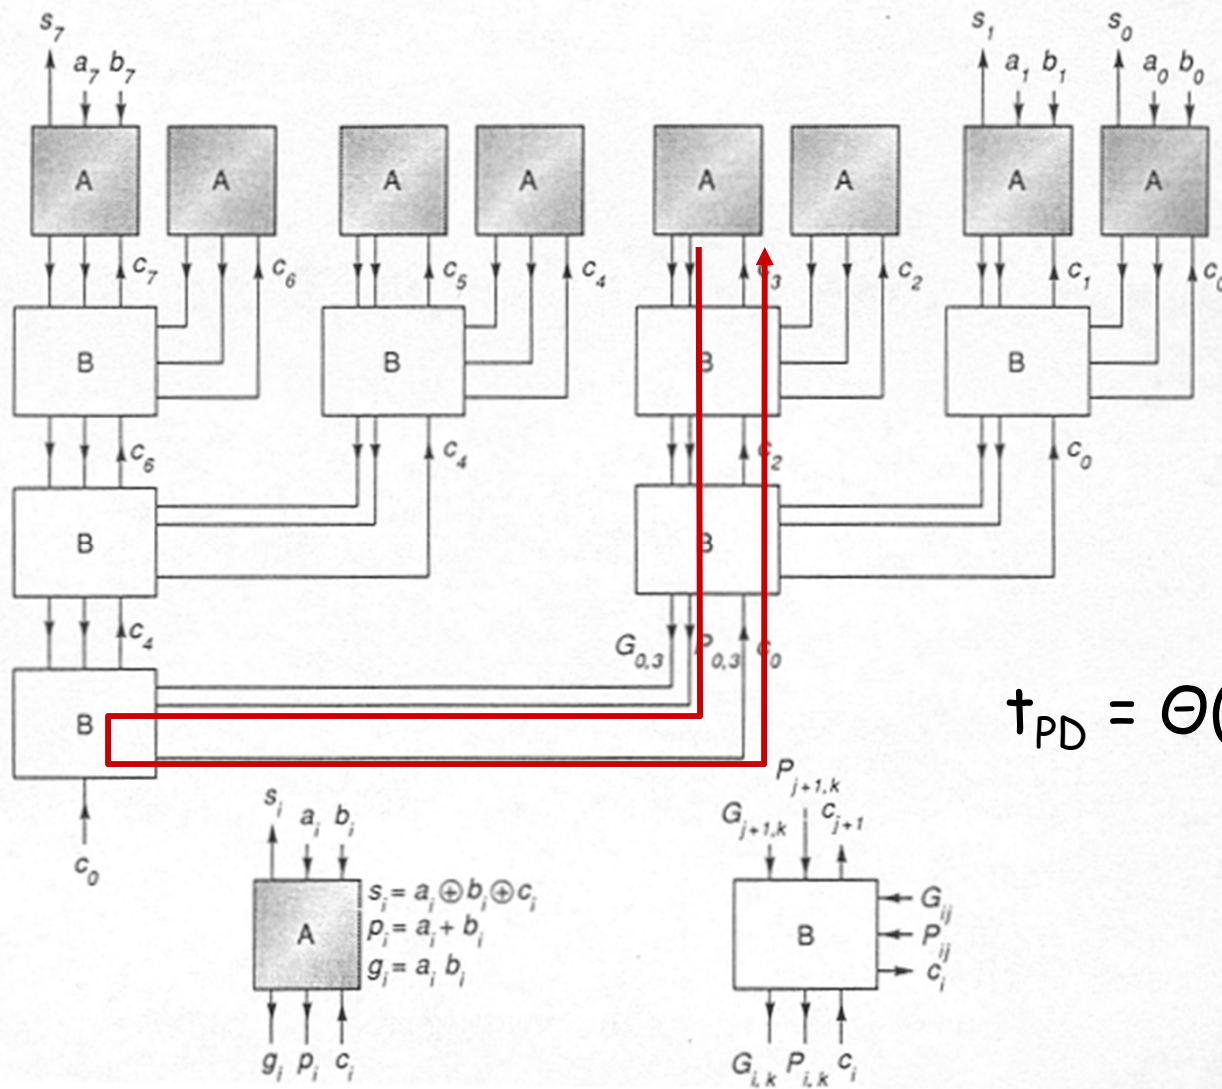
8-bit CLA (P/G generation)



8-bit CLA (carry generation)



8-bit CLA (complete)



$$t_{PD} = \Theta(\log(N))$$

Unsigned Multiplication

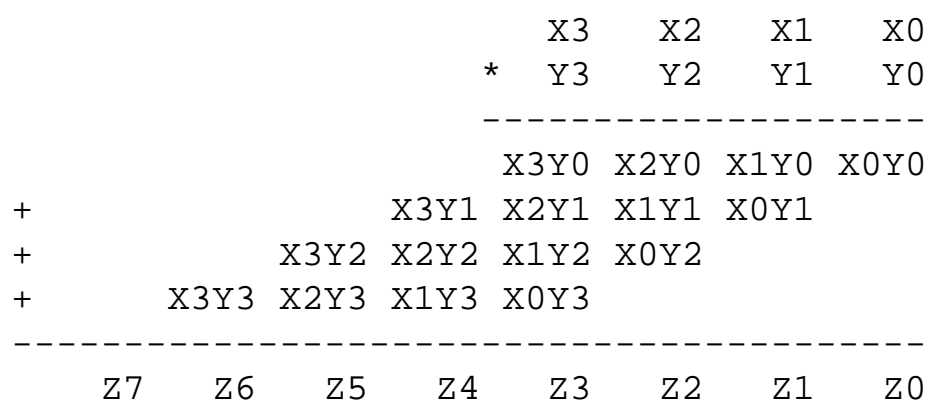
$$\begin{array}{r}
 A_3 A_1 \\
 \times B_3 B_1 \\
 \hline
 A_3 B_0 A_1 B_0 \\
 A_3 B_1 A_1 B_1 \\
 A_3 B_2 A_1 B_2 \\
 + A_3 B_3 A_1 B_3 \\
 \hline
 \end{array}$$

AB_i called a "partial product" \longrightarrow

Multiplying N-bit number by M-bit number gives (N+M)-bit result

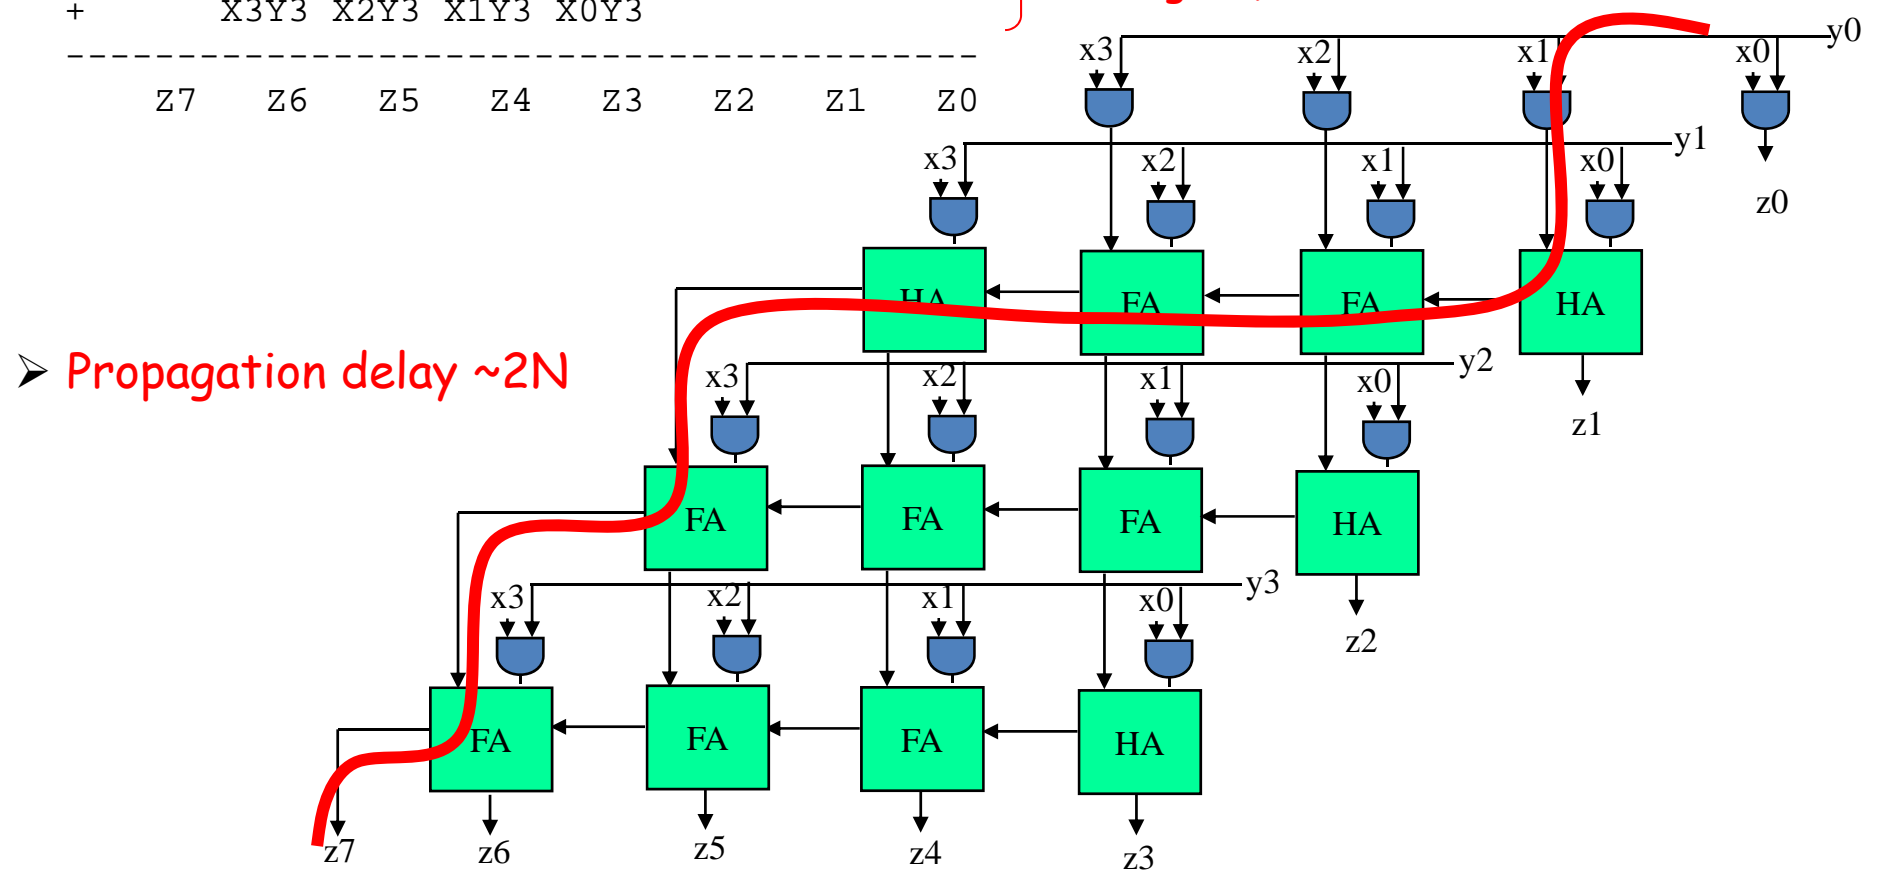
Easy part: forming partial products
 (just an AND gate since B_i is either 0 or 1)
 Hard part: adding M N-bit partial products

Combinational Multiplier (unsigned)



← multiplicand
 ← multiplier

Partial products, one for each bit in multiplier (each bit needs just one AND gate)



➤ Propagation delay $\sim 2N$

Baugh Wooley Formulation -The Math

no insight required

Assuming X and Y are 4-bit twos complement numbers:

$$X = -2^3x_3 + \sum_{i=0}^2 x_i 2^i \quad Y = -2^3y_3 + \sum_{i=0}^2 y_i 2^i$$

The product of X and Y is:

$$XY = x_3y_32^6 - \sum_{i=0}^2 x_iy_32^{i+3} - \sum_{j=0}^2 x_3y_j2^{j+3} + \sum_{i=0}^2 \sum_{j=0}^2 x_iy_j2^{i+j}$$

For twos complement, the following is true:

$$-\sum_{i=0}^3 x_i 2^i = -2^4 + \sum_{i=0}^3 x_i 2^i + 1$$

The product then becomes:

$$\begin{aligned} XY &= x_3y_32^6 + \sum_{i=0}^2 \overline{x_i}y_32^{i+3} + 2^3 - 2^6 + \sum_{j=0}^2 \overline{x_3}y_j2^{j+3} + 2^3 - 2^6 + \sum_{i=0}^2 \sum_{j=0}^2 x_iy_j2^{i+j} \\ &= x_3y_32^6 + \sum_{i=0}^2 \overline{x_i}y_32^{i+3} + \sum_{j=0}^2 \overline{x_3}y_j2^{j+3} + \sum_{i=0}^2 \sum_{j=0}^2 x_iy_j2^{i+j} + 2^4 - 2^7 \\ &= -2^7 + x_3y_32^6 + (\overline{x_2}y_3 + \overline{x_3}y_2)2^5 + (\overline{x_1}y_3 + \overline{x_3}y_1 + x_2y_2 + 1)2^4 \\ &\quad + (\overline{x_0}y_3 + \overline{x_3}y_0 + x_1y_2 + x_2y_1)2^3 + (x_0y_2 + x_1y_1 + x_2y_0)2^2 + 1 \\ &\quad + (x_0y_1 + x_1y_0)2^1 + (x_0y_0)2^0 \end{aligned}$$

Multiplication in Verilog

You can use the "*" operator to multiply two numbers:

```
wire [9:0] a,b;  
wire [19:0] result = a*b; // unsigned multiplication!
```

If you want Verilog to treat your operands as signed two's complement numbers, add the keyword `signed` to your `wire` or `reg` declaration:

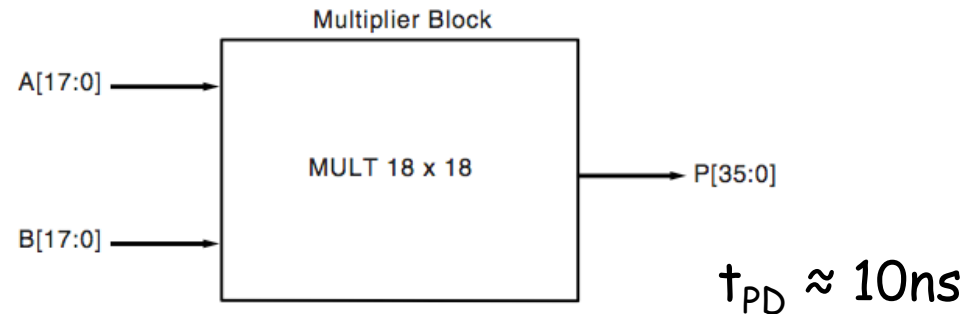
```
wire signed [9:0] a,b;  
wire signed [19:0] result = a*b; // signed multiplication!
```

Remember: unlike addition and subtraction, you need different circuitry if your multiplication operands are signed vs. unsigned. Same is true of the `>>>` (arithmetic right shift) operator. To get signed operations all operands must be signed.

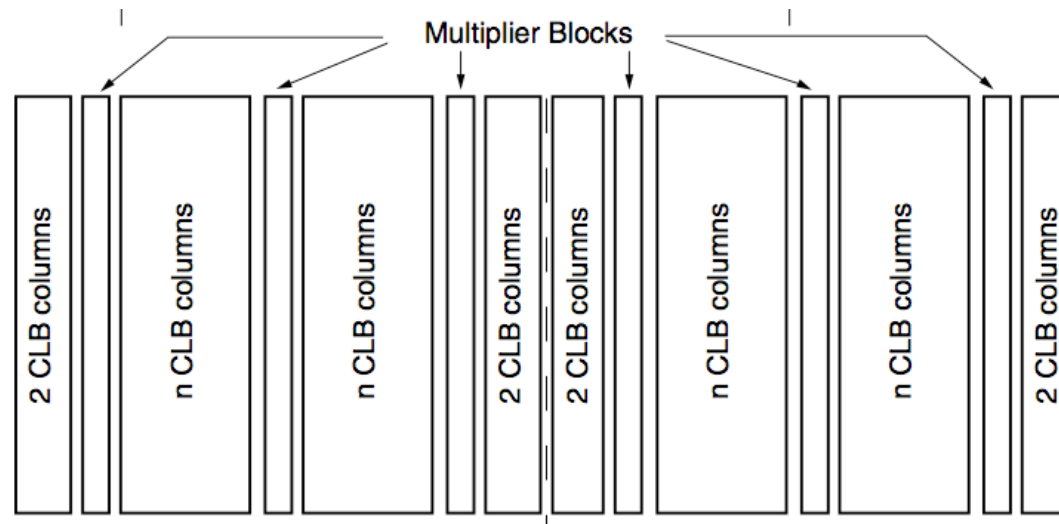
To make a signed constant: `10'sh37C`

Multiplication on the FPGA

Hardware multiplier block: two 18-bit twos complement (signed) operands

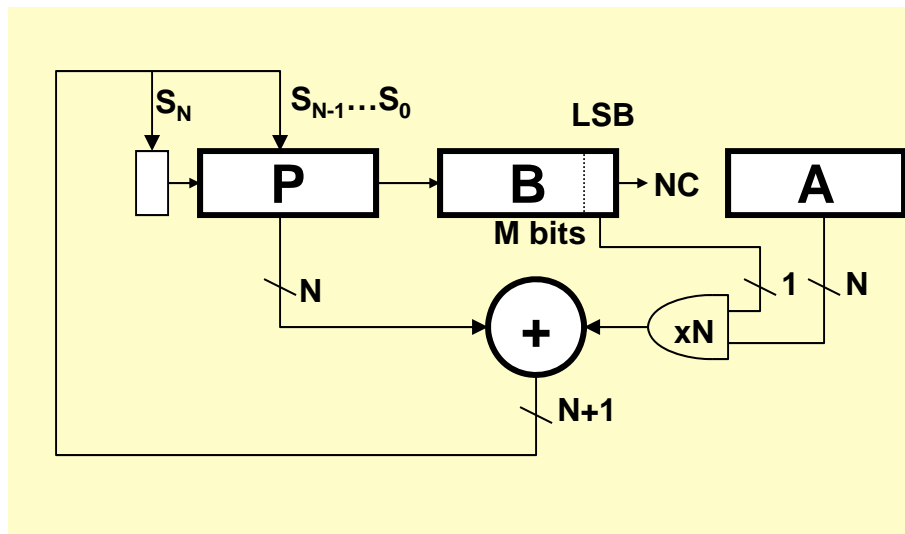


In the XC2V6000: 6 columns of mults, 24 in each column = 144 mults



Sequential Multiplier

Assume the multiplicand (A) has N bits and the multiplier (B) has M bits. If we only want to invest in a single N-bit adder, we can build a sequential circuit that processes a single partial product at a time and then cycle the circuit M times:

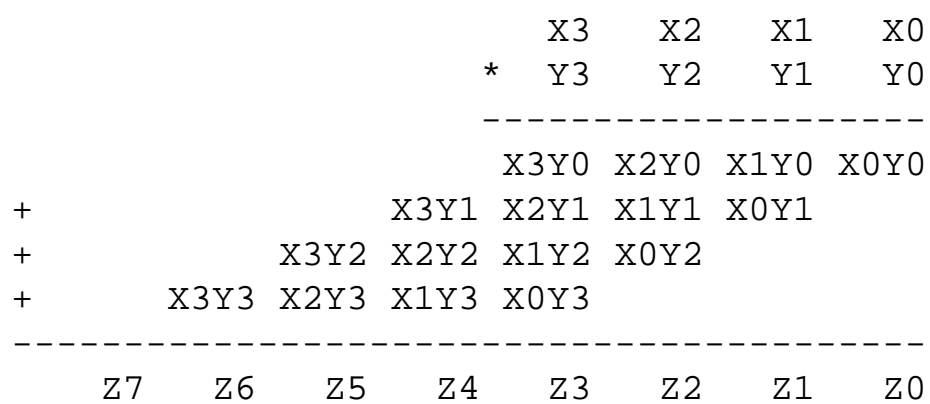


Init: $P \leftarrow 0$, load A and B

```
Repeat M times {  
     $P \leftarrow P + (B_{\text{LSB}} == 1 ? A : 0)$   
    shift P/B right one bit  
}
```

Done: (N+M)-bit result in P/B

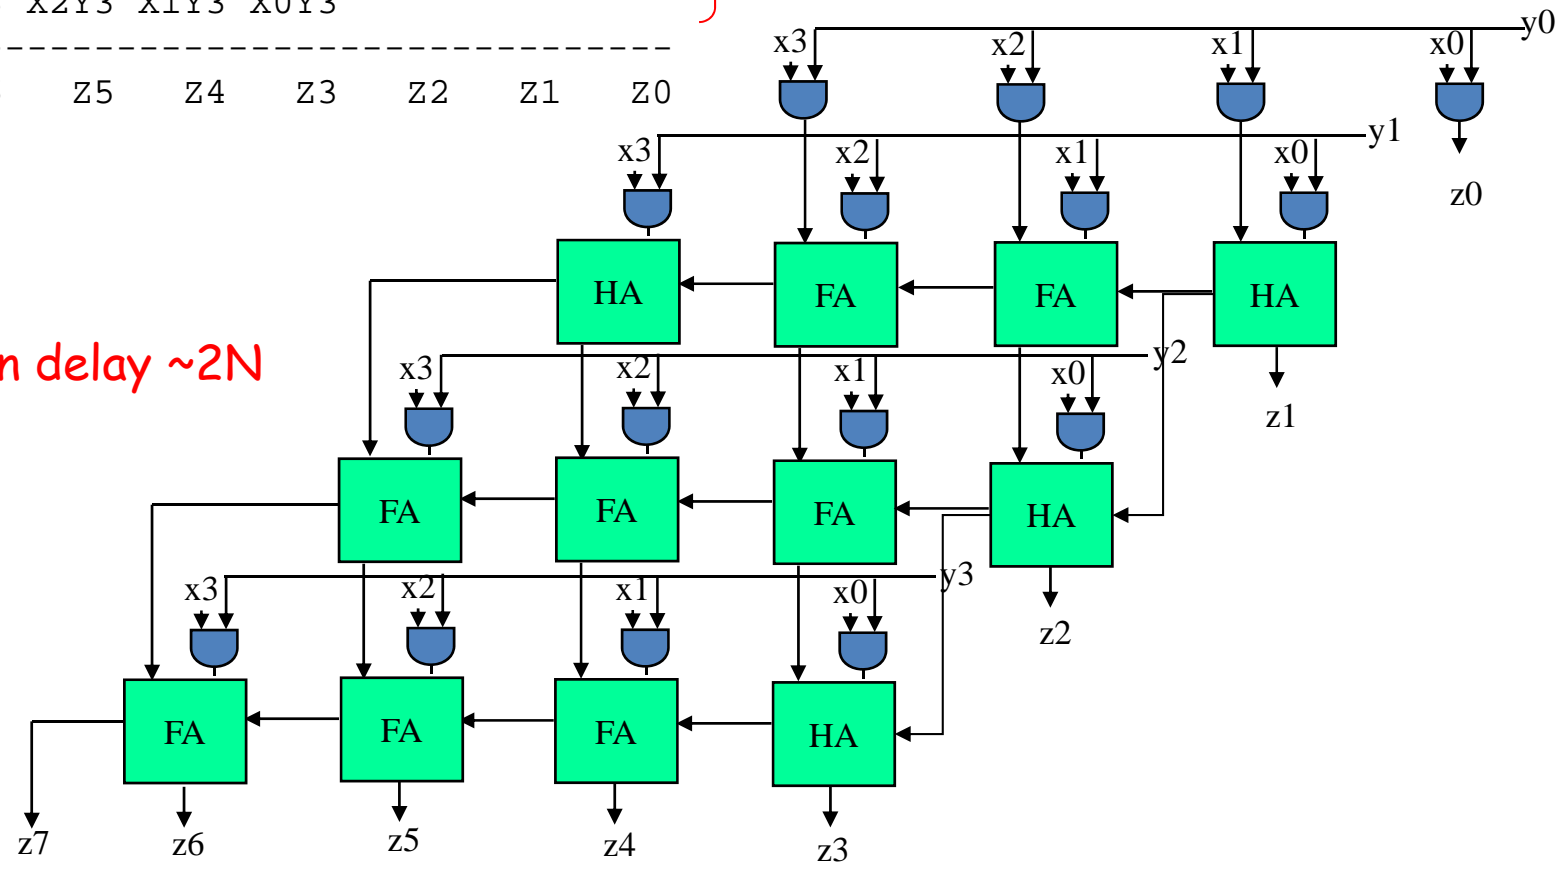
Combinational Multiplier (unsigned)



← multiplicand
 ← multiplier

Partial products, one for each bit in multiplier (each bit needs just one AND gate)

➤ Propagation delay ~2N

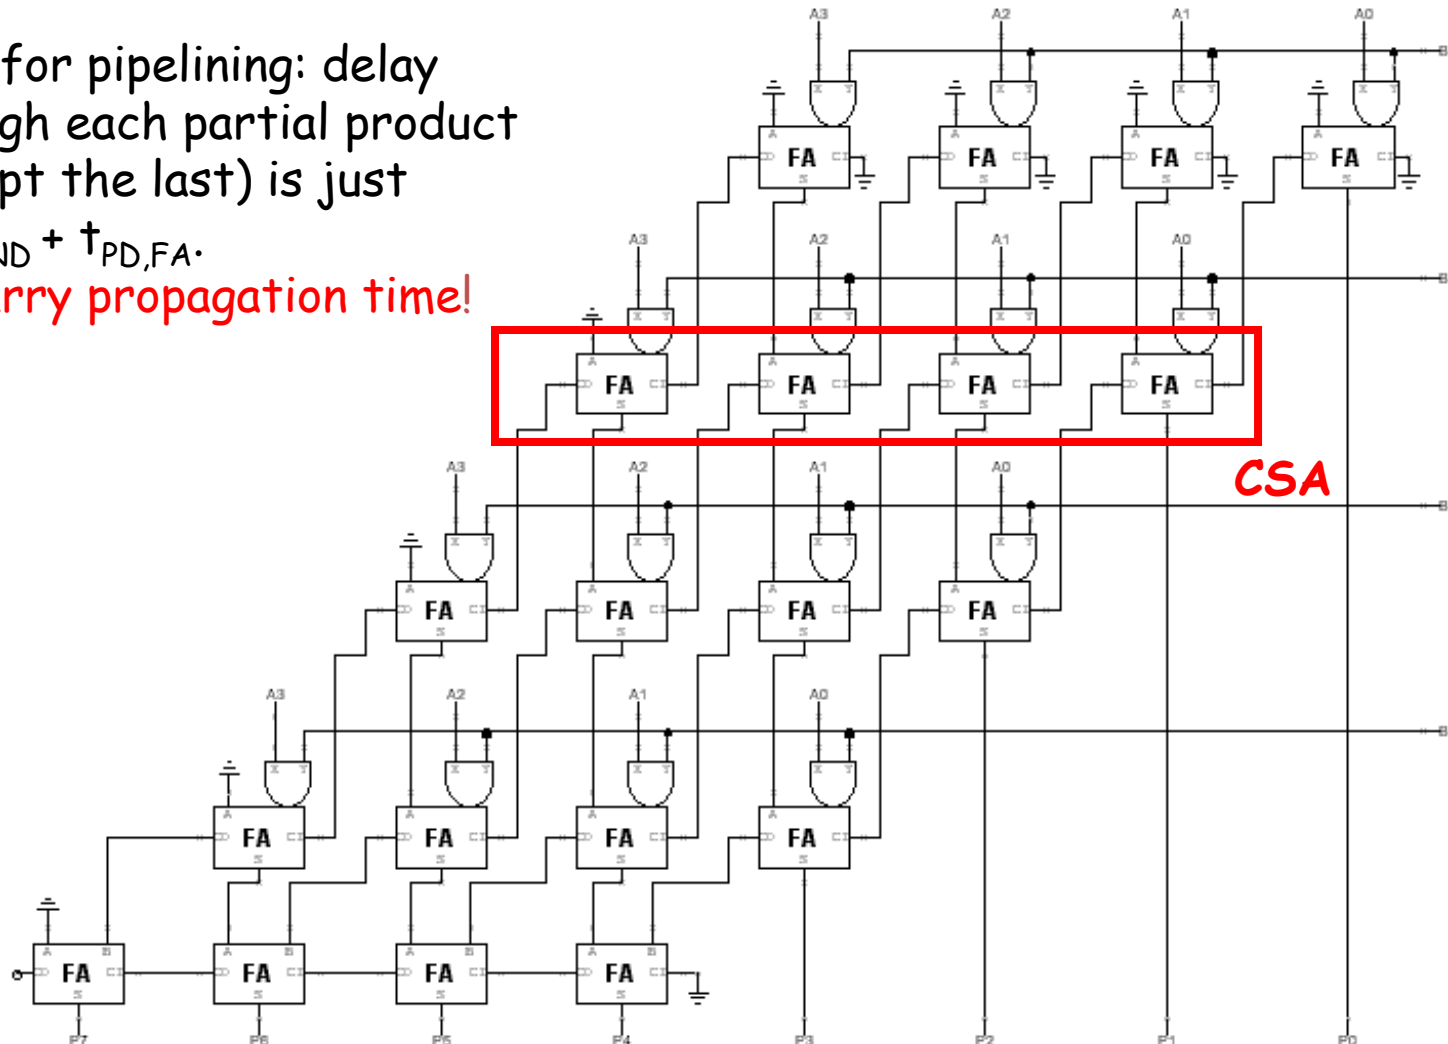


Useful building block: Carry-Save Adder

Good for pipelining: delay through each partial product (except the last) is just

$$t_{PD,AND} + t_{PD,FA}$$

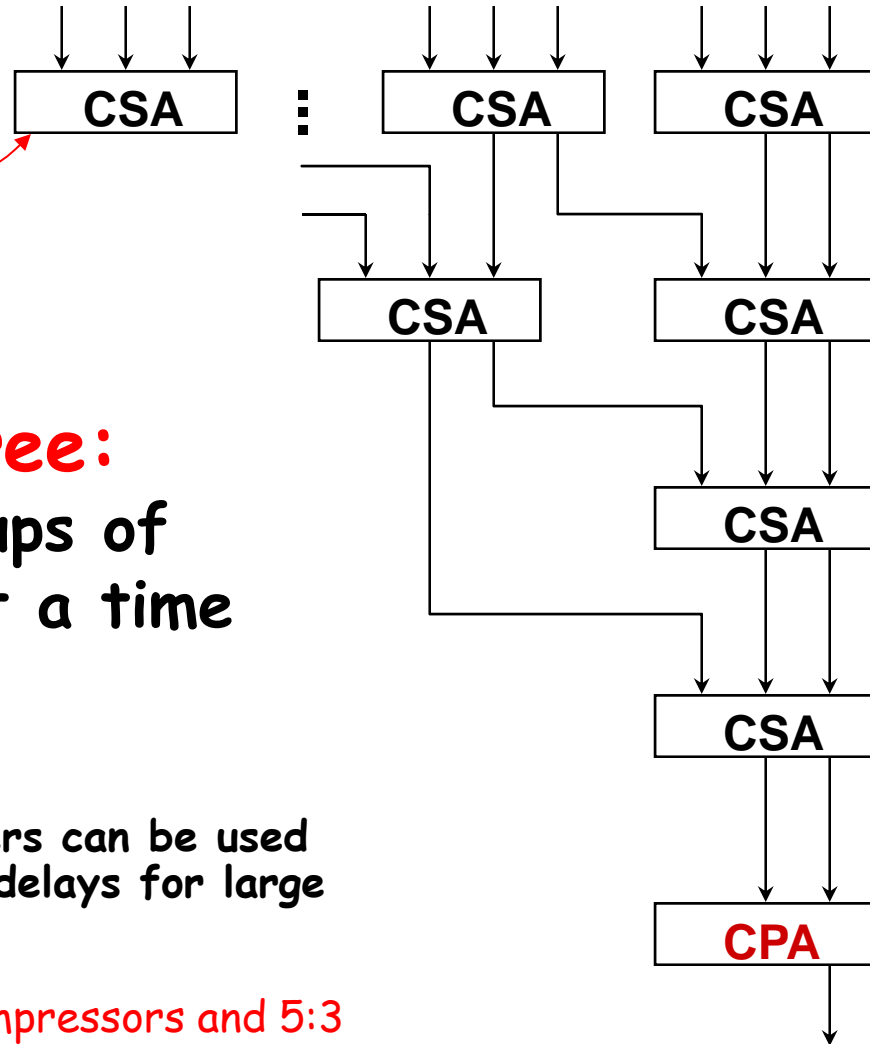
No carry propagation time!



Last stage is still a carry-propagate adder (CPA)

Wallace Tree Multiplier

This is called a 3:2 counter by multiplier hackers: counts number of 1's on the 3 inputs, outputs 2-bit result.



Wallace Tree:
Combine groups of three bits at a time

Higher fan-in adders can be used to further reduce delays for large M .

4:2 compressors and 5:3 counters are popular building blocks.

Wallace Tree * Four Bit Multiplier

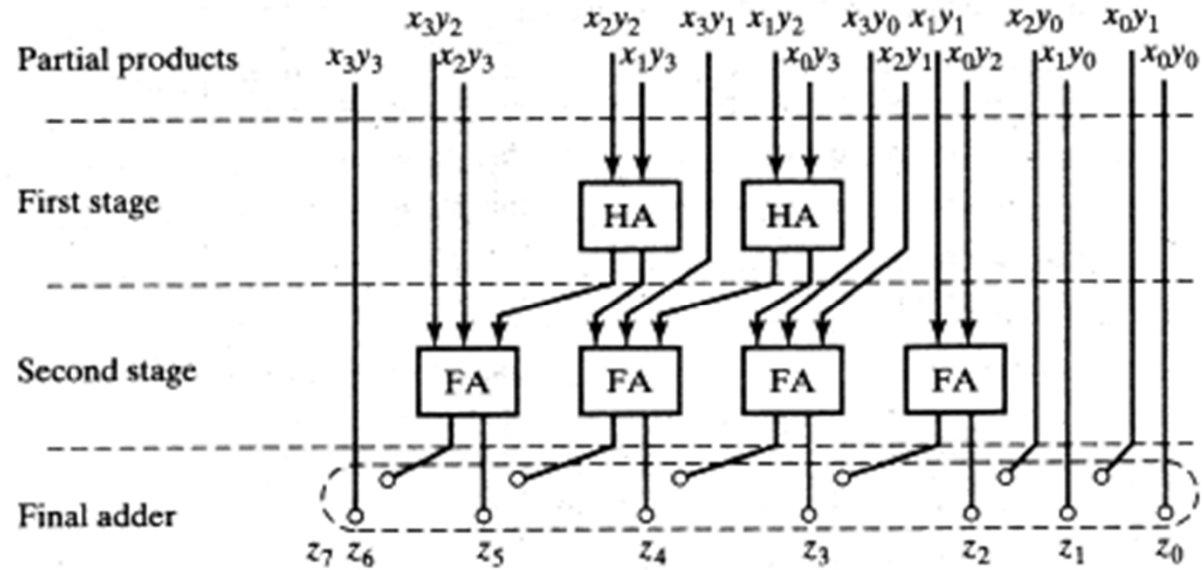


Figure 11-35 Wallace tree for four-bit multiplier.

*Digital Integrated Circuits
J Rabaey, A Chandrakasan, B Nikolic

Multiplication by a constant

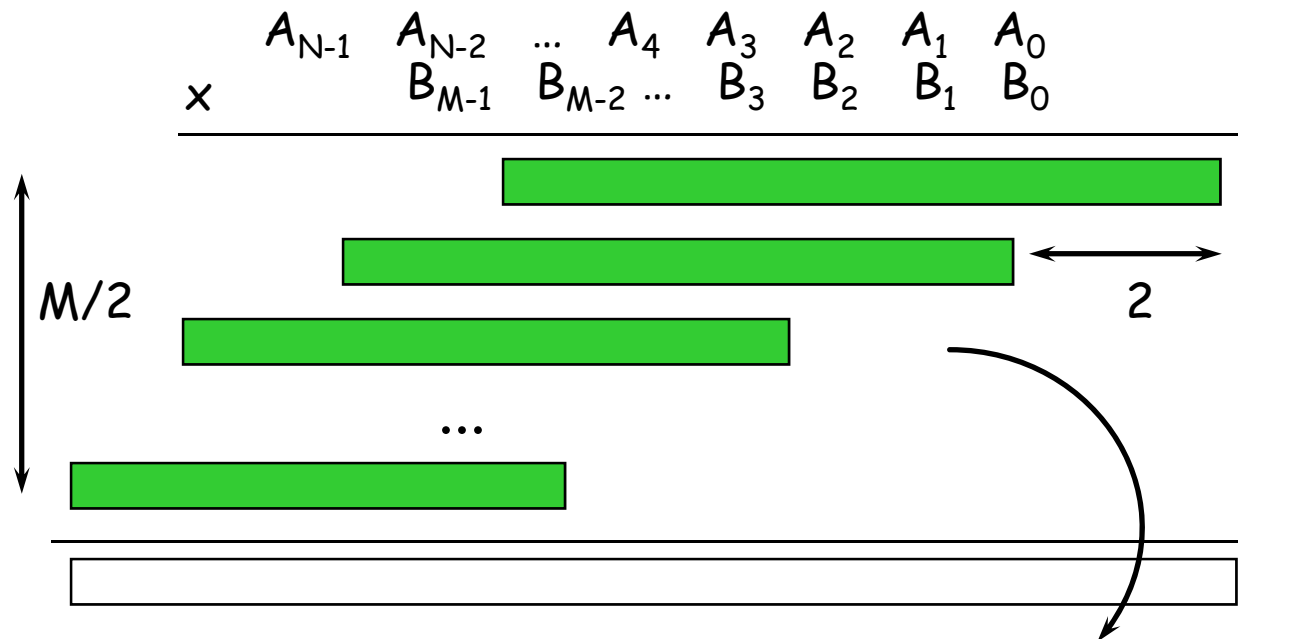
- If one of the operands is a constant, make it the multiplier (B in the earlier examples). For each "1" bit in the constant we get a partial product (PP) - may be noticeably fewer PPs than in the general case.
 - For example, in general multiplying two 4-bit operands generates four PPs (3 rows of full adders). If the multiplier is say, 12 (4'b1100), then there are only two PPs: $8*A+4*A$ (only 1 row of full adders).
 - But lots of "1"s means lots of PPs... can we improve on this?
- If we allow ourselves to subtract PPs as well as adding them (the hardware cost is virtually the same), we can re-encode arbitrarily long contiguous runs of "1" bits in the multiplier to produce just two PPs.

$$\dots 011110\dots = \dots 100000\dots - \dots 000010\dots = \dots 01000\bar{1}0\dots$$

where $\bar{1}$ indicates subtracting a PP instead of adding it. Thus we've re-encoded the multiplier using 1,0,-1 digits - aka *canonical signed digit* - greatly reducing the number of additions required.

Booth Recoding: Higher-radix mult.

Idea: If we could use, say, 2 bits of the multiplier in generating each partial product we would **halve the number of columns and halve the latency of the multiplier!**



Booth's insight: rewrite $2*A$ and $3*A$ cases, leave $4A$ for next partial product to do!

$$\begin{aligned}
 B_{K+1,K} * A &= 0 * A \rightarrow 0 \\
 &= 1 * A \rightarrow A \\
 &= 2 * A \rightarrow 4A - 2A \\
 &= 3 * A \rightarrow 4A - A
 \end{aligned}$$

Booth recoding

On-the-fly canonical signed digit encoding!

current bit pair \swarrow \searrow \swarrow from previous bit pair

B_{K+1}	B_K	B_{K-1}	action	
0	0	0	add 0	
0	0	1	add A	
0	1	0	add A	
0	1	1	add 2^*A	
1	0	0	sub 2^*A	
1	0	1	sub A	$\leftarrow -2^*A+A$
1	1	0	sub A	
1	1	1	add 0	$\leftarrow -A+A$

A "1" in this bit means the previous stage needed to add 4^*A . Since this stage is shifted by 2 bits with respect to the previous stage, adding 4^*A in the previous stage is like adding A in this stage!

Summary

- Performance of arithmetic blocks dictate the performance of a digital system
- Architectural and logic transformations can enable significant speed up (e.g., adder delay from $O(N)$ to $O(\log_2(N))$)
- Similar concepts and formulation can be applied at the system level
- **Timing analysis is tricky**: watch out for false paths!
- Area-Delay trade-offs (serial vs. parallel implementations)

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 // the hardest module!!
 - siren.v
 - fuel_pump.v
- Run simulation on each module!
- Use hex display: show state and time
- Use logic analyzer in Vivado

Car Alarm - Inputs & Outputs

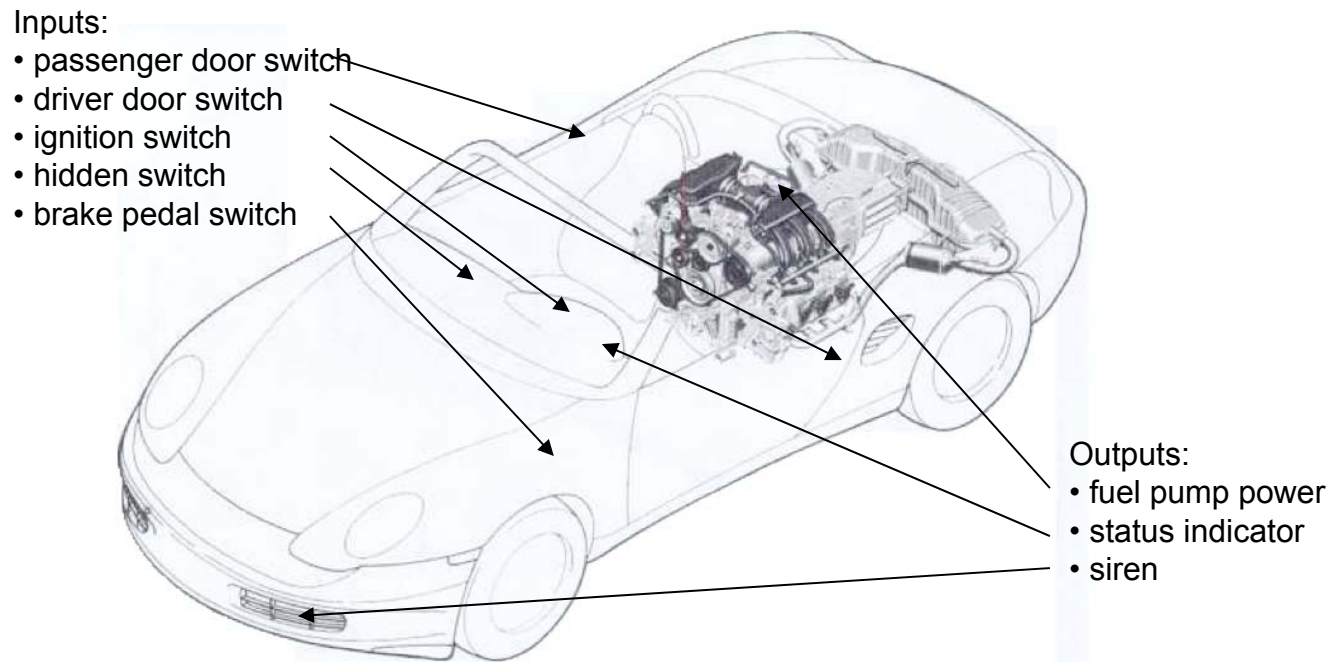
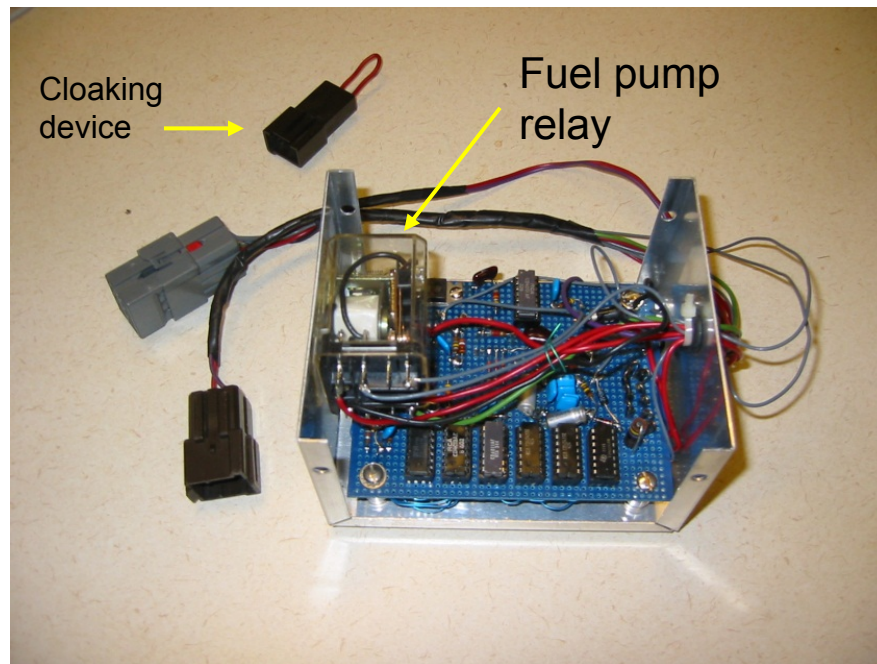


Figure 1: System diagram showing sensors (inputs) and actuators (outputs)

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
 - 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 27/25 million clock cycles. Avoids recompilations.

```
assign warp_speed = sw[6];
always @ (posedge clk) begin
    if (count == (warp_speed ? 3 : 26_999_999)) count <= 0;
    else count <= count +1;
end

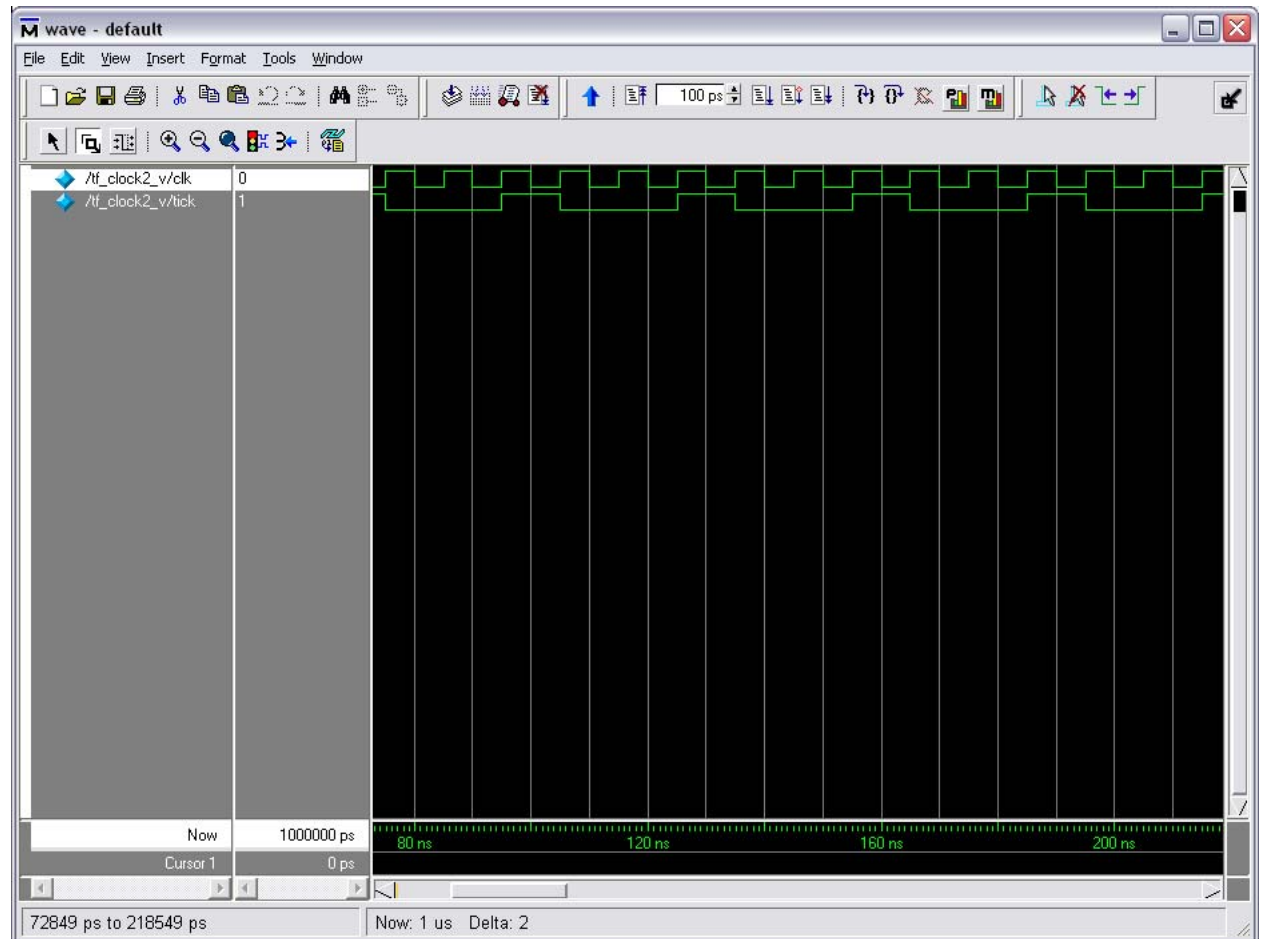
assign one_hz = (count == (warp_speed ? 3 : 26_999_999)) ;
```


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; . . .
```



For Loops, Repeat Loops in Simulation

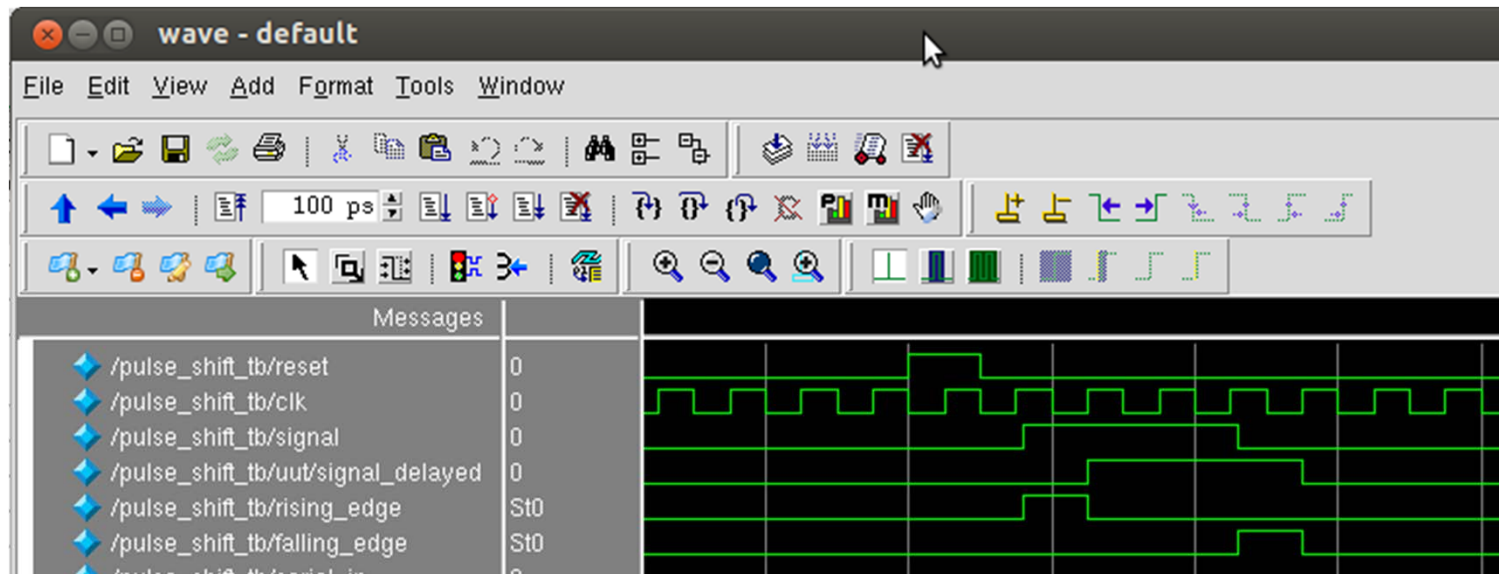
```
integer i; // index must be declared as integer
integer irepeat;

// this will just wait 10ns, repeated 32x.
// simulation only! Cannot 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

// other loops: forever, while
```

Edge Detection



```
reg signal_delayed;
```

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

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

Student Comments

- "All very reasonable except for lab 4, Car Alarm. Total pain in the ass."
- "The labs were incredibly useful, interesting, and helpful for learning. Lab 4 (car alarm) is long and difficult, but overall the labs are not unreasonable."