

ESE534: Computer Organization

Day 4: January 25, 2012
Sequential Logic
(FSMs, Pipelining, FSMD)



Previously

- Boolean Logic
- Gates
- Arithmetic
- Complexity of computations
 - E.g. area and delay for addition

Today

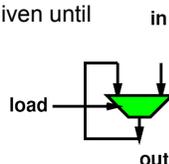
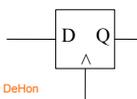
- Sequential Logic
 - Add registers, state
 - Finite-State Machines (FSM)
 - Register Transfer Level (RTL) logic
 - Datapath Reuse
 - Pipelining
 - Latency and Throughput
 - Finite-State Machines with Datapaths (FSMD)

Preclass

- Can we solve the problem entirely using Boolean logic functions?

Latches, Registers

- New element is a state element.
- Canonical instance is a register:
 - remembers the last value it was given until told to change
 - typically signaled by clock



Why Registers?

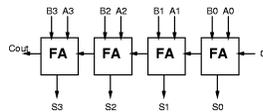
- Why do we need registers?

Reuse

- In general, we want to **reuse our components in time**

– not disposable logic

- How do we **guarantee disciplined reuse?**



Penn ESE534 Spring2012 -- DeHon

7

To Reuse Logic...

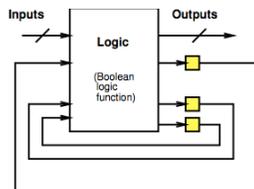
- Make sure all logic completed evaluation
 - Outputs of gates are valid
 - Meaningful to look at them
 - Gates are "finished" with work and ready to be used again
- Make sure consumers get value
 - Before being overwritten by new calculation (new inputs)

Penn ESE534 Spring2012 -- DeHon

8

Synchronous Logic Model

- Data starts
 - Inputs to circuit
 - Registers
- Perform combinational (boolean) logic
- Outputs of logic
 - Exit circuit
 - Clocked into registers
- Given long enough clock
 - Think about registers getting values updated by logic on each clock cycle



Penn ESE534 Spring2012 -- DeHon

9

Issues of Timing...

- ...many issues in detailed implementation
 - glitches and hazards in logic
 - timing discipline in clocking
 - ...
- We're going to (mostly) work above that level this term.
 - Will talk about the delay of logic between registers
- Watch for these details in ESE370/570

Penn ESE534 Spring2012 -- DeHon

10

Preclass

- How do we build an adder for arbitrary input width?

Penn ESE534 Spring2012 -- DeHon

11

Preclass

- What did the addition of state register(s) do for us?

Penn ESE534 Spring2012 -- DeHon

12

Added Power

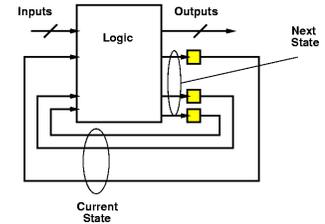
- Process *unbounded* input with finite logic
 - Ratio input:gates \rightarrow arbitrarily large
- State is a **finite** (bounded) representation of what's happened before
 - finite amount of stuff can remember to synopsize the past
- State allows behavior to depend on past (on context)

Penn ESE534 Spring2012 -- DeHon

13

Finite-State Machine (FSM) (Finite Automata)

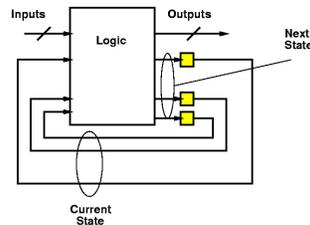
- Logic core
- Plus registers to hold state



Penn ESE534 Spring2012 -- DeHon

FSM Model

- FSM – a model of computations
- More powerful than Boolean logic functions
- Both
 - Theoretically
 - practically



Penn ESE534 Spring2012 -- DeHon

FSM Abstraction

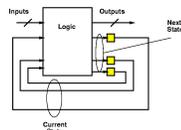
- Implementation vs. Abstraction
 - Nice to separate out
 - The abstract function want to achieve
 - The concrete implementation
 - Saw with Boolean logic
 - There are many ways to implement function
 - Want to select the concrete one that minimizes costs
- FSMs \rightarrow also separate out “desired function” from “implementation”

Penn ESE534 Spring2012 -- DeHon

16

Formal FSM Specification (Abstract from implementation)

- An FSM is a sextuple $M = \{K, \Sigma, \delta, s, F, \Sigma_o\}$
 - K is finite set of states
 - Σ is a finite alphabet for inputs
 - $s \in K$ is the start state
 - $F \subseteq K$ is the set of final states
 - Σ_o is a finite set of output symbols
 - δ is a transition function from $K \times \Sigma$ to $K \times \Sigma_o$



Penn ESE534 Spring2012 -- DeHon

17

Finite State Machine

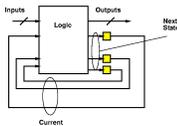
- Less formally:
 - Behavior depends not just on input
 - (as was the case for combinational logic)
 - Also depends on state
 - Can be completely different behavior in each state
 - Logic/output now depends on both
 - state and input

Penn ESE534 Spring2012 -- DeHon

18

Formal FSM Specification (Abstract from implementation)

- An FSM is a sextuple $M = \{K, \Sigma, \delta, s, F, \Sigma_o\}$
 - K is finite set of states
 - Σ is a finite alphabet for inputs
 - $s \in K$ is the start state
 - $F \subseteq K$ is the set of final states
 - Σ_o is a finite set of output symbols
 - δ is a transition function from $K \times \Sigma$ to $K \times \Sigma_o$



Relate each to concrete implementation;
identify value of abstracting.

Penn ESE534 Spring2012 -- DeHon

19

Specifying an FSM

- Logic becomes:
 - if (state=s1)
 - boolean logic for state 1
 - (including logic for calculate next state)
 - else if (state=s2)
 - boolean logic for state2
 - ...
 - if (state=sn)
 - boolean logic for state n

Penn ESE534 Spring2012 -- DeHon

20

Specifying FSM

- What's your favorite way to specify an FSM?
- Another reason we need to separate the abstract operation from the
 - Specification
 - Implementation

Penn ESE534 Spring2012 -- DeHon

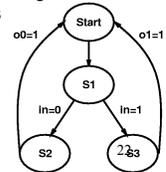
21

FSM Specification

- St1: goto St2
- St2:
 - if (I==0) goto St3
 - else goto St4
- St3:
 - output o0=1
 - goto St1
- St4:
 - output o1=1
 - goto St2

- Could be:

- behavioral language {Verilog, VHDL, Bluespec}
- computer language (C)
- state-transition graph
- extract from gates + registers



Penn ESE534 Spring2012 -- DeHon

State Encoding

- States not (necessarily) externally visible
- We have *freedom* in how to encode them
 - assign bits to states
- Usually want to exploit freedom to minimize implementation costs
 - area, delay, energy
- (there are algorithms to attack – ESE535)

Penn ESE534 Spring2012 -- DeHon

23

FSM Equivalence

- Harder than Boolean logic
- Doesn't have unique canonical form
- Consider:
 - state encoding not change behavior
 - two "equivalent" FSMs may not even have the same number of states
 - can deal with **infinite** (unbounded) input
 - ...so cannot enumerate output in all cases
 - No direct correspondence of a truth table

Penn ESE534 Spring2012 -- DeHon

24

FSM Equivalence

- What does matter?
 - What property needs to hold for two FSMs to be equivalent?

Penn ESE534 Spring2012 -- DeHon

25

FSM Equivalence

- What matters is external observability
 - FSM outputs same signals in response to every possible input sequence
- Is it possible to check equivalence over an infinite number of input sequences?
- Possible?
 - Finite state suggests there is a finite amount of checking required to verify behavior

Penn ESE534 Spring2012 -- DeHon

26

FSM Equivalence Flavor

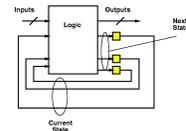
- Given two FSMs A and B
 - consider the composite FSM AB
 - Inputs wired together
 - Outputs separate
- Ask:
 - is it possible to get into a composite state in which A and B output different symbols?
- There is a literature on this

Penn ESE534 Spring2012 -- DeHon

27

Systematic FSM Design

- Start with specification
- Can compute Boolean logic for each state
 - If conversion...
 - including next state translation
 - Keep state symbolic (s1, s2...)
- Assign state encodings
- Then have combinational logic
 - has current state as part of inputs
 - produces next state as part of outputs
- Design comb. logic and add state registers



Penn ESE534 Spring2012 -- DeHon

28

Arbitrary Adder

- Work through design as FSM if necessary
 - Encode inputs, outputs
 - States
 - Encodings
 - Logic

Penn ESE534 Spring2012 -- DeHon

29

RTL

- Register Transfer Level description
- Registers + Boolean logic
- Most likely: what you've written in Verilog, VHDL

Penn ESE534 Spring2012 -- DeHon

30

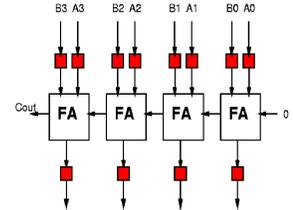
Datapath Reuse

Penn ESE534 Spring2012 -- DeHon

31

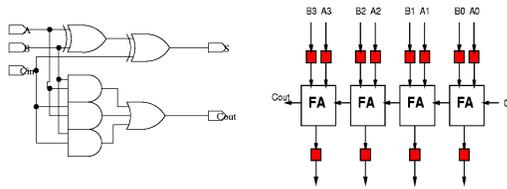
Reuse: "Waiting" Discipline

- Use registers and timing for orderly progression of data



Penn ESE534 Spring2012 -- DeHon

Example: 4b Ripple Adder



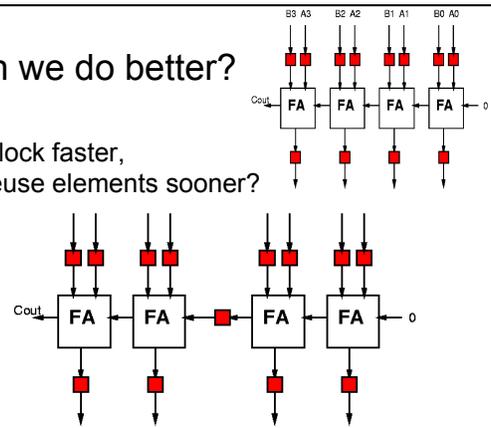
- How fast can we clock this?
- Min Clock Cycle: 8 gates A, B to S3

Penn ESE534 Spring2012 -- DeHon

33

Can we do better?

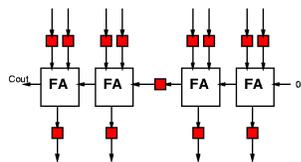
- Clock faster, reuse elements sooner?



Penn ESE5

Stagger Inputs

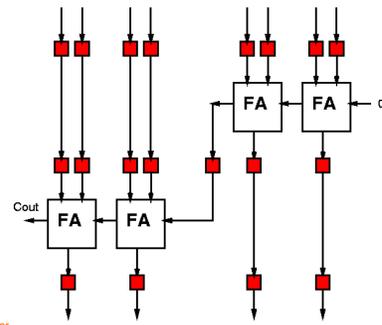
- Correct if expecting A,B[3:2] to be staggered one cycle behind A,B[1:0]
- ...and succeeding stage expects S[3:2] staggered from S[1:0]



Penn ESE534 Spring2012 -- DeHon

Align Data / Balance Paths

Good discipline to line up pipe stages in diagrams.



Penn ESE534 Spring2012 -- DeHon

Speed

How fast can we clock this?

Assuming we clock that fast, what is the delay from A,B to S3?

Penn ESE534 Spring2012 -- DeHon

Pipelining and Timing

- Once introduce pipelining
 - Clock cycle = rate of reuse
 - Is **not** the same as the delay to complete a computation

Penn ESE534 Spring2012 -- DeHon

Pipelining and Timing

- Throughput
 - How many results can the circuit produce per unit time
 - If can produce one result per cycle,
 - Reciprocal of clock period
- Throughput of this design?

Penn ESE534 Spring2012 -- DeHon

Pipelining and Timing

- Latency
 - How long does it take to produce one result
 - Product of
 - clock cycle
 - number of clocks between input and output
- Latency of this design?

Penn ESE534 Spring2012 -- DeHon

Example: 4b RA pipe 2

Latency and Throughput:

- Latency: 8 gates to S3
- Throughput: 1 result / 4 gate delays max

41

Penn ESE534 Spring2012 -- DeHon

Throughput vs. Latency

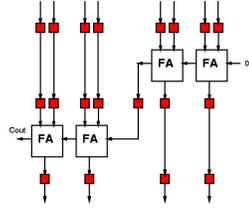
- Examples where throughput matters?
- Examples where latency matters?

42

Penn ESE534 Spring2012 -- DeHon

Deeper?

- Can we do it again?
- What's our limit?
- Why would we stop?

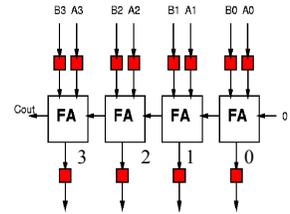


Penn ESE534 Spring2012 -- DeHon

43

More Reuse

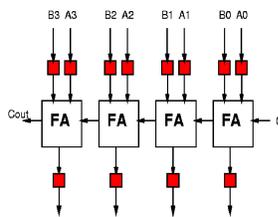
- Saw could pipeline and reuse FA more frequently
- Suggests we're **wasting** the FA part of the time in non-pipelined
 - What is FA3 doing while FA0 is computing?



Penn ESE534 Spring2012 -- DeHon

More Reuse (cont.)

- If we're willing to take 8 gate-delay units, do we need 4 FAs?



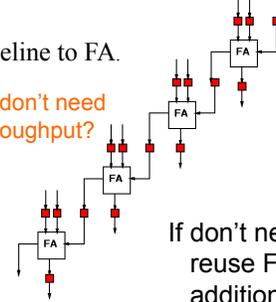
Penn ESE534 Spring2012 -- DeHon

45

Ripple Add (pipe view)

Can pipeline to FA.

What if don't need the throughput?

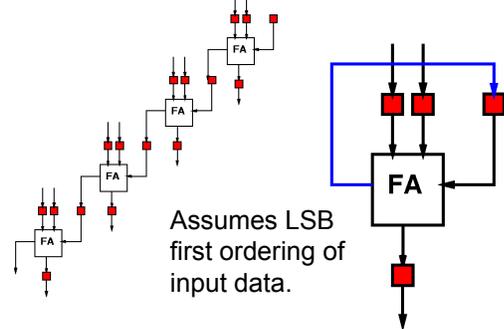


If don't need throughput, reuse FA on **SAME** addition.

Penn ESE534 Spring2012 -- DeHon

46

Bit Serial Addition



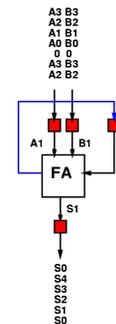
Assumes LSB first ordering of input data.

Penn ESE534 Spring2012 -- DeHon

7

Bit Serial Addition: Pipelining

- **Latency and throughput?**
- Latency: 8 gate delays
 - 10 for 5th output bit
- Throughput: 1 result / 10 gate delays
- Registers do have time overhead
 - setup, hold time, clock jitter



Penn ESE534 Spring2012 -- DeHon

48

Multiplication

- Can be defined in terms of addition
- Ask you to play with implementations and tradeoffs in homework 2

Penn ESE534 Spring2012 -- DeHon

49

Design Space for Computation

Penn ESE534 Spring2012 -- DeHon

50

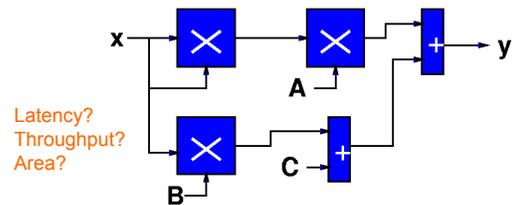
Compute Function

- Compute:
 $y = Ax^2 + Bx + C$
- Assume
 - $D(\text{Mpy}) > D(\text{Add})$
 - E.g. $D(\text{Mpy})=24, D(\text{Add})=8$
 - $A(\text{Mpy}) > A(\text{Add})$
 - E.g. $A(\text{Mpy})=64, A(\text{Add})=8$

Penn ESE534 Spring2012 -- DeHon

51

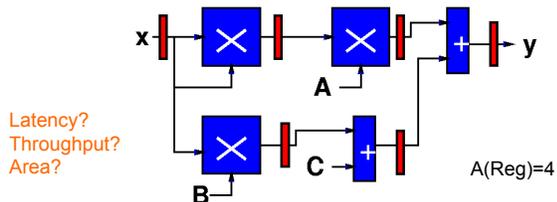
Spatial Quadratic



- $D(\text{Quad}) = 2 * D(\text{Mpy}) + D(\text{Add}) = 56$
- Throughput $1 / (2 * D(\text{Mpy}) + D(\text{Add})) = 1/56$
- $A(\text{Quad}) = 3 * A(\text{Mpy}) + 2 * A(\text{Add}) = 208$ ⁵²

Penn ESE534 Spring2012 -- DeHon

Pipelined Spatial Quadratic



- $D(\text{Quad}) = 3 * D(\text{Mpy}) = 72$
- Throughput $1 / D(\text{Mpy}) = 1/24$
- $A(\text{Quad}) = 3 * A(\text{Mpy}) + 2 * A(\text{Add}) + 6A(\text{Reg}) = 232$ ⁵³

Penn ESE534 Spring2012 -- DeHon

53

Quadratic with Single Multiplier and Adder?

- We've seen reuse to perform the **same** operation
 - pipelining
 - bit-serial, homogeneous datapath
- We can also reuse a resource in time to perform a **different** role.

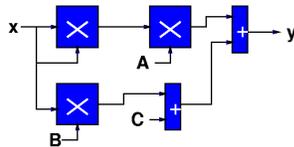
Penn ESE534 Spring2012 -- DeHon

54

Repeated Operations

- What operations occur multiple times in this datapath?

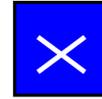
- x^2 , $A(x^2)$, Bx
- $(Bx)+c$, $(A^2x^2)+(Bx+c)$



Penn ESE534 Spring2012 -- DeHon

Quadratic Datapath

- Start with one of each operation
- (alternatives where build multiply from adds...e.g. homework)



Penn ESE534 Spring2012 -- DeHon

56

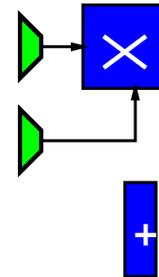
Lecture Ended Here

Penn ESE534 Spring2012 -- DeHon

57

Quadratic Datapath

- Multiplier serves multiple roles
 - x^2
 - $A(x^2)$
 - Bx
- Will need to be able to steer data (switch interconnections)

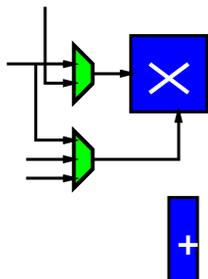


Penn ESE534 Spring2012 -- DeHon

58

Quadratic Datapath

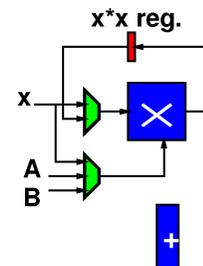
- Multiplier serves multiple roles
 - x^2
 - $A(x^2)$
 - Bx
- Inputs
 - a) x , x^2
 - b) x, A, B



Penn ESE534 Spring2012 -- DeHon

Quadratic Datapath

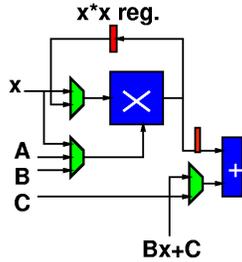
- Multiplier serves multiple roles
 - x^2
 - $A(x^2)$
 - Bx
- Inputs
 - a) x , x^2
 - b) x, A, B



Penn ESE534 Spring2012 -- DeHon

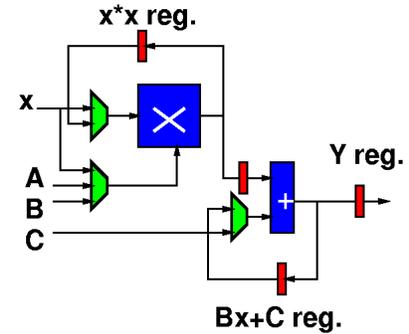
Quadratic Datapath

- Adder serves multiple roles
 - $(Bx)+c$
 - $(A*x*x)+(Bx+c)$
- Inputs
 - one always mpy output
 - $C, Bx+C$



Penn ESE534 Spring2012 -- DeHon

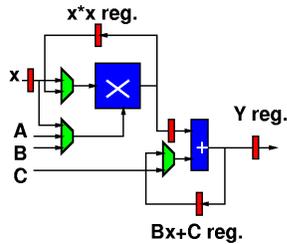
Quadratic Datapath



Penn ESE534 Spring2012 -- DeHon

Quadratic Datapath

- Add input register for x

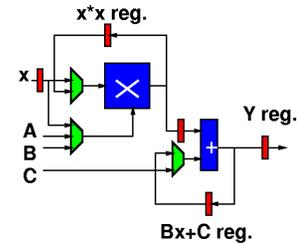


Penn ESE534 Spring2012 -- DeHon

Quadratic Control

- Now, we just need to control the datapath
- What control?

- Control:
 - LD x
 - LD $x*x$
 - MA Select
 - MB Select
 - AB Select
 - LD $Bx+C$
 - LD Y



Penn ESE534 Spring2012 -- DeHon

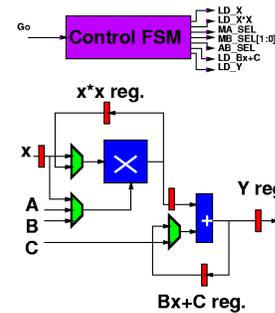
FSMD

- FSMD = FSM + Datapath
- Stylization for building controlled datapaths such as this (a **pattern**)
- Of course, an FSMD is just an FSM
 - it's often easier to think about as a datapath
 - synthesis, place and route tools have been notoriously bad about discovering/ exploiting datapath structure

Penn ESE534 Spring2012 -- DeHon

65

Quadratic FSMD



Penn ESE534 Spring2012 --

66

Quadratic FSMD Control

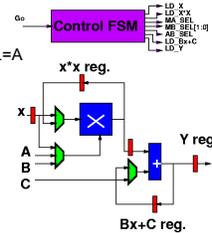
- S0: if (go) LD_X; goto S1
 - else goto S0
- S1: MA_SEL=x, MB_SEL[1:0]=x, LD_x*x
 - goto S2
- S2: MA_SEL=x, MB_SEL[1:0]=B
 - goto S3
- S3: AB_SEL=C, MA_SEL=x*x, MB_SEL=A
 - goto S4
- S4: AB_SEL=Bx+C, LD_Y
 - goto S0

Penn ESE534 Spring2012 -- DeHon

67

Quadratic FSMD Control

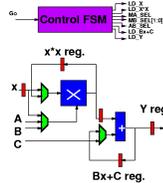
- S0: if (go) LD_X; goto S1
 - else goto S0
- S1: MA_SEL=x, MB_SEL[1:0]=x, LD_x*x
 - goto S2
- S2: MA_SEL=x, MB_SEL[1:0]=B
 - goto S3
- S3: AB_SEL=C, MA_SEL=x*x, MB_SEL=A
 - goto S4
- S4: AB_SEL=Bx+C, LD_Y
 - goto S0



Penn ESE534 Spring2012 -- DeHon

Quadratic FSM

- $D(\text{mux3}) = D(\text{mux2}) = 1$
- $A(\text{mux2}) = 2$
- $A(\text{mux3}) = 3$
- $A(\text{QFSM}) \approx 10$
- Latency/Throughput/Area?
- Latency: $5 \cdot (D(\text{MPY}) + D(\text{mux3})) = 125$
- Throughput: $1/\text{Latency} = 1/125$
- Area: $A(\text{Mpy}) + A(\text{Add}) + 5 \cdot A(\text{Reg}) + 2 \cdot A(\text{Mux2}) + A(\text{Mux3}) + A(\text{QFSM}) = 109$



Penn ESE534 Spring2012 -- DeHon

69

Admin: Reminder

- Next homework due Monday
- Reading for next week online

Penn ESE534 Spring2012 -- DeHon

70

Big Ideas [MSB Ideas]

- Registers allow us to reuse logic
- Can implement any FSM with gates and registers
- Pipelining
 - increases parallelism
 - allows reuse in time (same function)
- Control and Sequencing
 - reuse in time for different functions
- Can tradeoff Area and Time

Penn ESE534 Spring2012 -- DeHon

71

Big Ideas [MSB-1 Ideas]

- RTL specification
- FSMD idiom

Penn ESE534 Spring2012 -- DeHon

72