

ESE680-002 (ESE534): Computer Organization

Day 17: March 19, 2007
Interconnect 5: Meshes




Previously

- Saw
 - need to exploit locality/structure in interconnect
 - a mesh might be useful
 - **Question:** how does w grow?
 - Rent's Rule as a way to characterize structure

Today

- Mesh:
 - Channel width bounds
 - Linear population
 - Switch requirements
 - Routability
 - Segmentation
 - Clusters
 - Commercial


Mesh



Mesh Channels


- Lower Bound on w ?
- Bisection Bandwidth
 - $BW \propto N^p$
 - $N^{0.5}$ channels in bisection

$$W \propto \frac{N^p}{\sqrt{N}} = N^{(p-0.5)}$$




Straight-forward Switching Requirements

- Switching Delay?
- Total Switches?



Switch Delay

- Switching Delay: $2\sqrt{N_{\text{subarray}}}$
 - worst case: $N_{\text{subarray}} = N$




Penn ESE680-002 Spring2007 -- DeHon

7

Total Switches

- Switches per switchbox:
 - $4 \times (3w \times w) / 2 = 6w^2$
 - Bidirectional switches
 - $(N \rightarrow W \text{ same as } W \rightarrow N)$
 - double count




Penn ESE680-002 Spring2007 -- DeHon

8

Total Switches

- Switches per switchbox:
 - $6w^2$
- Switches into network:
 - $(K+1)w$
- Switches per PE:
 - $6w^2 + (K+1)w$
 - $w = cN^{p-0.5}$
 - Total $\propto N^{2p-1}$
- Total Switches: $N^*(Sw/PE) \propto N^{2p}$



Penn ESE680-002 Spring2007 -- DeHon

9

Routability?


- Asking if you can route in a given channel width is:
 - NP-complete

Penn ESE680-002 Spring2007 -- DeHon

10

Traditional Mesh Population


- Switchbox contains only a linear number of switches in channel width



Penn ESE680-002 Spring2007 -- DeHon

Linear Mesh Switchbox

- Each entering channel connect to:
 - One channel on each remaining side (3)
 - 4 sides
 - W wires
 - Bidirectional switches
 - $(N \rightarrow W \text{ same as } W \rightarrow N)$
 - double count
 - $3 \times 4 \times W / 2 = 6W$ switches
 - vs. $6w^2$ for full population




Penn ESE680-002 Spring2007 -- DeHon

12

Total Switches

- Switches per switchbox:
 - $6w$
- Switches into network:
 - $(K+1)w$
- Switches per PE:
 - $6w + (K+1)w$
 - $w = cN^{p-0.5}$
 - Total $\propto N^{p-0.5}$
- Total Switches: $N^*(Sw/PE) \propto N^{p+0.5} > N$




Penn ESE680-002 Spring2007 -- DeHon

13

Total Switches (linear population)

- Total Switches
 - $\propto N^{p+0.5}$
 - $N < N^{p+0.5} < N^{2p}$
- Switches grow faster than nodes
- Wires grow faster than switches




Penn ESE680-002 Spring2007 -- DeHon

14

Checking Constants

When do linear population designs become wire dominated?

- Wire pitch = 8λ
- switch area = $2500 \lambda^2$
- wire area: $(8w)^2$
- switch area: $6 \times 2500 w^2$
- crossover
 - $w=234$?
 - (practice smaller)




Penn ESE680-002 Spring2007 -- DeHon

15

Total Switches (linear population)

- Total Switches
 - $\propto N^{p+0.5}$
 - $N < N^{p+0.5} < N^{2p}$
- Switches grow faster than nodes
- Wires grow faster than switches



When wire dominated, want to minimize use of wires
...switches not matter.


Penn ESE680-002 Spring2007 -- DeHon

16

Checking Constants: Full Population

Does full population really use all the wire physical tracks?

- Wire pitch = 8λ
- switch area = $2500 \lambda^2$
- wire area: $(8w)^2$
- switch area: $6 \times 2500 w^2$
- effective wire pitch:
 - 120λ
 - ~ 15 times pitch



Penn ESE680-002 Spring2007 -- DeHon

17

Practical


- Full population is switch dominated
 - doesn't really use all the potential physical tracks
- Just showed:
 - would take $15 \times$ Mapping Ratio for linear population to take same area as full population (once crossover to wire dominated)
- Can afford to not use some wires perfectly
 - to reduce switches (area)

Penn ESE680-002 Spring2007 -- DeHon

18

Diamond Switch

- Typical switchbox pattern:
 - Used by Xilinx
- Many less switches, but cannot guarantee will be able to use all the wires
 - may need more wires than implied by Rent, since cannot use all wires
 - this was already true...now more so




Penn ESE680-002 Spring2007 -- DeHon

19

Universal SwitchBox

- Same number of switches as diamond
- **Locally:** can guarantee to satisfy any set of requests
 - request = direction through swbox
 - as long as meet channel capacities
 - and order on all channels irrelevant
 - can satisfy
- Not a global property
 - no guarantees between swboxes



Penn ESE680-002 Spring2007 -- DeHon

Diamond vs. Universal?

- Universal routes strictly more configurations
- Universal


Xilinx
- can't route (or rotations)

Penn ESE680-002 Spring2007 -- DeHon

21

Inter-Switchbox Constraints

- Channels connect switchboxes
- For valid route, must satisfy all adjacent switchboxes



Penn ESE680-002 Spring2007 -- DeHon

Mapping Ratio?

- How bad is it?
- How much wider do channels have to be?
- Mapping Ratio:
 - detail channel width required / global ch width

Penn ESE680-002 Spring2007 -- DeHon

23

Mapping Ratio


- Empirical:
 - Seems plausible, constant in practice
- Theory/provable:
 - There is no Constant Mapping Ratio
 - At least detail/global
 - can be arbitrarily large!

Penn ESE680-002 Spring2007 -- DeHon

24

Domain Structure


- Once enter network (choose color) can only switch within domain



Penn ESE680-002 Spring2007 -- DeHon

25


Detail Routing as Coloring



Penn ESE680-002 Spring2007 -- DeHon

26

Detail Routing as Coloring




- Global Route channel width = 2
- Detail Route channel width = N
 - Can make arbitrarily large difference

Penn ESE680-002 Spring2007 -- DeHon

27

Detail Routing as Coloring



Penn ESE680-002 Spring2007 -- DeHon

28

Routability

- Domain Routing is NP-Complete
 - can reduce coloring problem to domain selection
 - i.e.* map adjacent nodes to same channel
 - Previous example shows basic shape
 - (another reason routers are slow)

Penn ESE680-002 Spring2007 -- DeHon

29

Routing

- Lack of detail/global mapping ratio
 - Says detail can be arbitrarily worse than global
 - Say global not necessarily predict detail
 - Argument against decomposing mesh routing into global phase and detail phase
 - Modern FPGA routers do not

Penn ESE680-002 Spring2007 -- DeHon

30

Segmentation

- To improve speed (decrease delay)
- Allow wires to bypass switchboxes
- Maybe save switches?
- Certainly cost more wire tracks

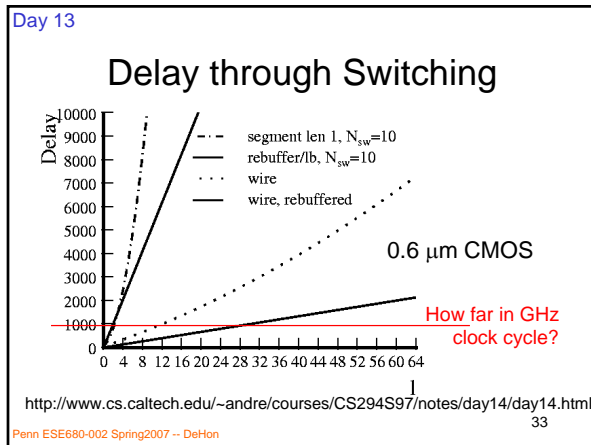
Penn ESE680-002 Spring2007 -- DeHon

Day 13

Buffered Delay

- Chip: 7mm side, 70nm sq. (45nm process)
 - 10^5 squares across chip
- $L_{\text{seg}} \approx 10^4 \text{ sq.}$ ($3.5 \times 10^3 \text{ sq.}$)
- 10 segments:
 - Each of delay $2 T_{\text{gate}}$
 - $T_{\text{cross}} = 20 \times 30\text{ps} = 600\text{ps}$ **Compare: 4ns**
 - $T_{\text{cross}} = 2 \times 30 \times 5\text{ps} = 300\text{ps}$

Penn ESE680-002 Spring2007 -- DeHon 32



Segmentation

- Segment of Length L_{seg}
 - 6 switches per switchbox visited
 - Only enters a switchbox every L_{seg}
 - SW/sbox/track of length $L_{\text{seg}} = 6/L_{\text{seg}}$

Penn ESE680-002 Spring2007 -- DeHon 34

Segmentation

- Reduces switches on path \sqrt{N}/L_{seg}
- May get fragmentation
- Another cause of unusable wires


Penn ESE680-002 Spring2007 -- DeHon 35

Segmentation: Corner Turn Option

- Can you corner turn in the middle of a segment?
- If can, need one more switch
- SW/sbox/track = $5/L_{\text{seg}} + 1$


Penn ESE680-002 Spring2007 -- DeHon 36

VPR Segment 4 Pix



Penn ESE680-002 Spring2007 -- DeHon

VPR Segment 4 Route




Penn ESE680-002 Spring2007 -- DeHon

38

C-Box Depopulation

- Not necessary for every input to connect to every channel
- Saw last time:
 - $K \times (N-K+1)$ switches
- Maybe use less?




Penn ESE680-002 Spring2007 -- DeHon

39

IO Population


- Toronto Model
 - F_c fraction of tracks which an input connects to
- IOs spread over 4 sides
- Maybe show up on multiple
 - Shown here: 2



Penn ESE680-002 Spring2007 -- DeHon

40

IO Population




Penn ESE680-002 Spring2007 -- DeHon

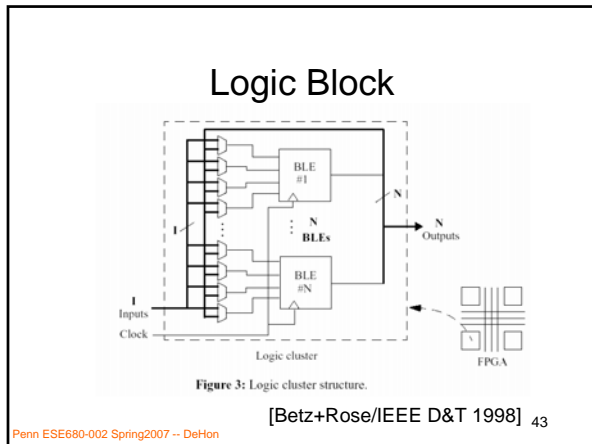
41

Leaves Not LUTs


- Recall cascaded LUTs
- Often group collection of LUTs into a Logic Block




Penn ESE680-002 Spring2007 -- DeHon



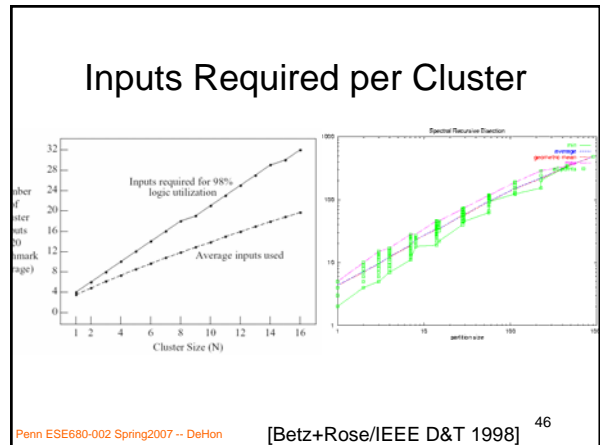
Penn ESE680-002 Spring2007 -- DeHon



Penn ESE680-002 Spring2007 -- DeHon



Penn ESE680-002 Spring2007 -- DeHon



Penn ESE680-002 Spring2007 -- DeHon

- ### Review: Mesh Design Parameters
- Cluster Size
 - Internal organization
 - LB IO (Fc, sides)
 - Switchbox Population and Topology
 - Segment length distribution
 - Switch rebuffering
- 47

Penn ESE680-002 Spring2007 -- DeHon

- ### Directional Drive
- Modern devices all directionally driven
 - *i.e.* separate channels for +x vs. -x routing
 - Faster: less parasitic stub capacitance
 - Driven
 - by scaling (saw need to buffer more often)
 - Increased demand for speed
 - Predictability
 - Harder to argue about wiring requirements
 - Not guaranteed half of wires going each direction
 - Empirically, fewer buffered switches
 - See: Lemieux et al., FPT2004
- 48

Penn ESE680-002 Spring2007 -- DeHon

Directional Drive

Total buffers same both cases.
More C-box sw per pair,
but maybe fewer pairs than tracks? [\[Lemieux.../FPT2004\]](#)

49

Penn ESE680-002 Spring2007 -- DeHon

Directional Drive

- Directional pair has same number of muxes/bufs as bidirectional track

50

[\[Lemieux.../FPT2004\]](#)

Penn ESE680-002 Spring2007 -- DeHon

Commercial Parts

51

Penn ESE680-002 Spring2007 -- DeHon

XC4K Interconnect

--

Penn ESE680-002 Spring2007 -- DeHon

XC4K Interconnect Details

53

Penn ESE680-002 S

Virtex II

Figure 48: Routing Resources

54

Penn ESE680-002 Spring2007 -- DeHon

Virtex II Interconnect Resources





Figure 49 Hierarchical Routing Resources

©2002 Xilinx, Inc.

55

Penn ESE680-002 Spring2007 -- DeHon

Stratix III



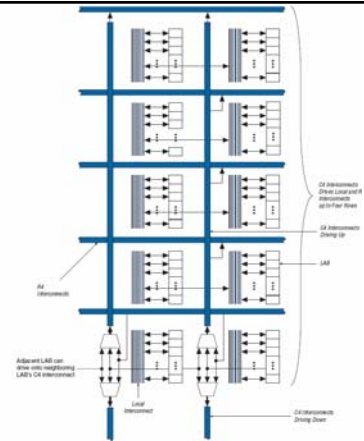
Notes to Figure 3-1

- (1) C4 and C12 interconnects can drive R4 interconnects.
- (2) This pattern is repeated for every LAB in the LAB row.

Penn ESE680-002 Spring2007 -- DeHon

56

Stratix III



Penn ESE680-002 Spring2007 -- DeHon

Admin

- Handout: reading for next Monday
- Interconnect assignment due Wed.

Penn ESE680-002 Spring2007 -- DeHon

58

Big Ideas [MSB Ideas]

- Mesh natural 2D topology
 - Channels grow as $\Omega(N^{p-0.5})$
 - Wiring grows as $\Omega(N^{2p})$
 - Linear Population:
 - Switches grow as $\Omega(N^{p+0.5})$
 - Worse than shown for hierarchical
 - Unbounded global \rightarrow detail mapping ratio
 - Detail routing NP-complete

Penn ESE680-002 Spring2007 -- DeHon

59

Big Ideas [MSB-1 Ideas]

- Segmented/bypass routes
 - can reduce switching delay
 - costs more wires (fragmentation of wires)

Penn ESE680-002 Spring2007 -- DeHon

60