CS184a: Computer Architecture (Structure and Organization)

Day 15: February 12, 2003
Interconnect 5: Meshes

Previous

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

Today

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

Mesh

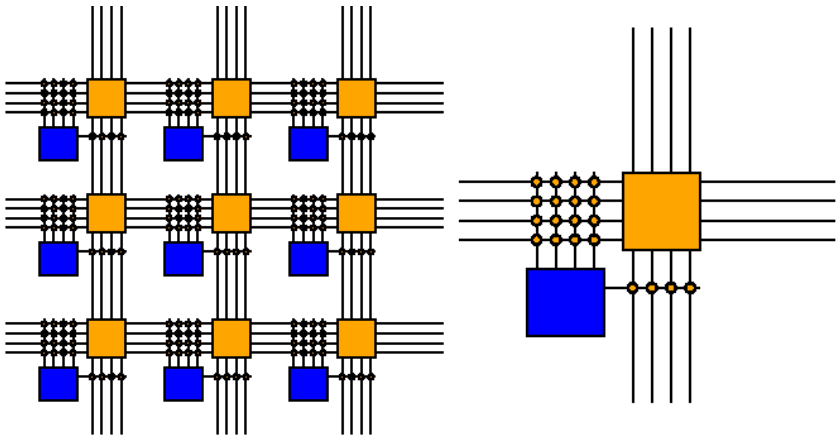

Mesh Channels

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

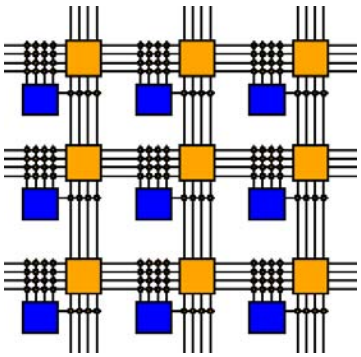

$$W \propto N^p / \sqrt{N} = N^{(p-0.5)}$$

Straight-forward Switching Requirements

- Total Switches?
- Switching Delay?

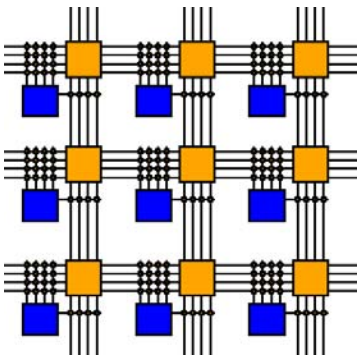

Switch Delay

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

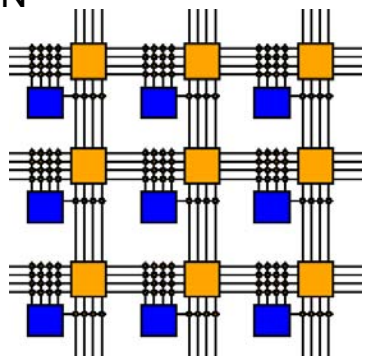

Total Switches

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

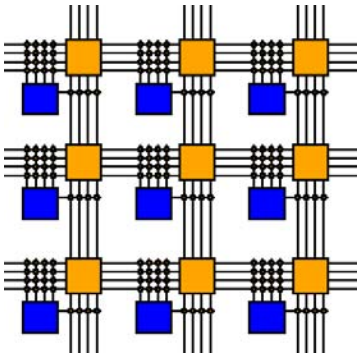

Total Switches

- Switches per switchbox:
  - 4 3w×w / 2 = $6w^2$
- 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) $\propto N^{2p}$

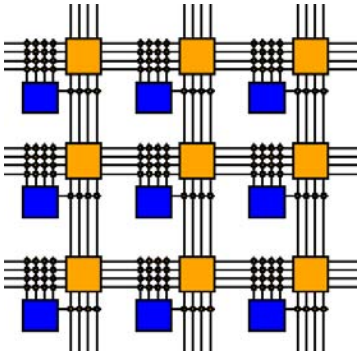

Routability?

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

Traditional Mesh Population

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

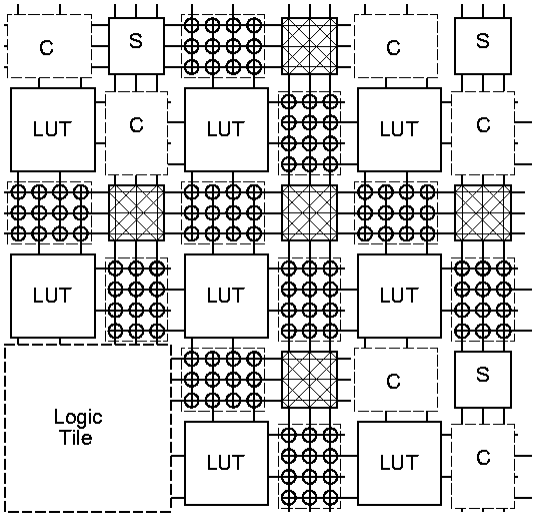

Linear Mesh Switchbox

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

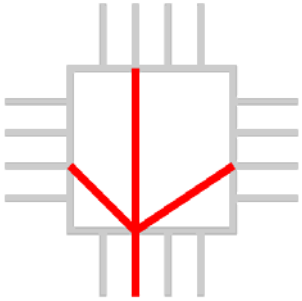

Total Switches

- 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: $N*(Sw/PE) \propto N^{p+0.5} > N$

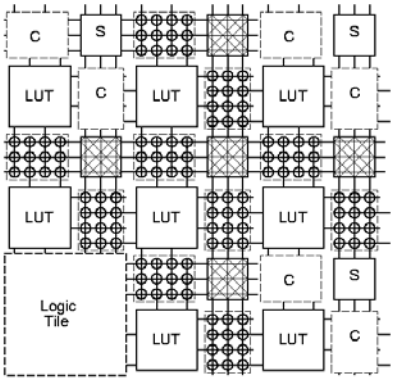

Total Switches

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

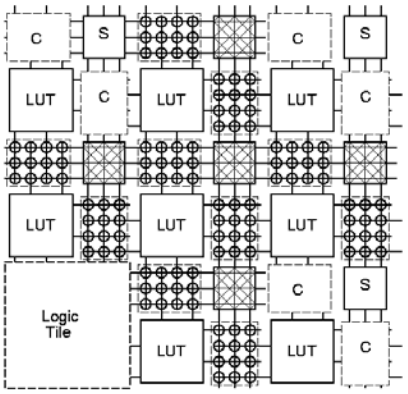

Checking Constants

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

Checking Constants: Full Population

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

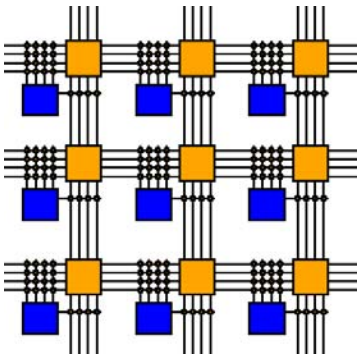

Practical

- Just showed:
  - would take 15× 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

Diamond Switch

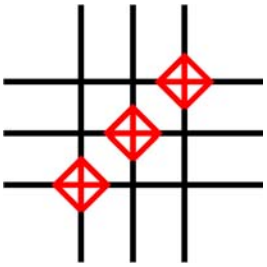

- Typical switchbox pattern:
  - Used by Xilinx
- Many less switches, but cannot guarantee will be able to use all the wires
  - may need more wires than implied by Rent, since cannot use all wires
  - this was already true…now more so

Domain Structure

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

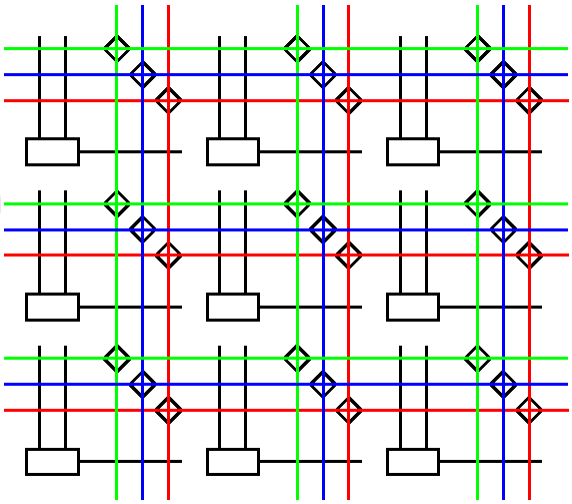

Universal SwitchBox

- 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

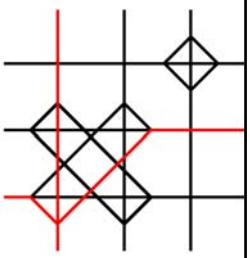

Diamond vs. Universal?

- Universal routes strictly more configurations

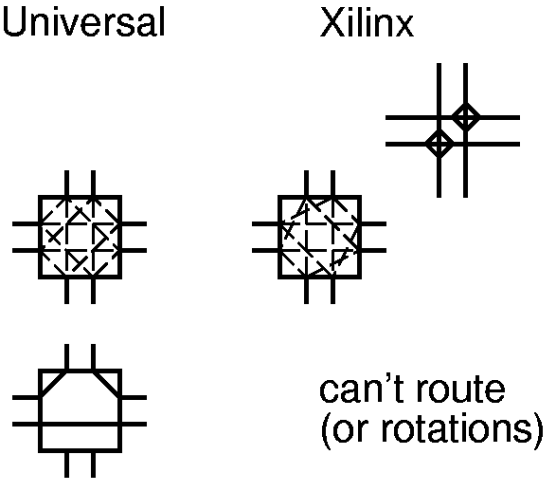

Inter-Switchbox Constraints

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

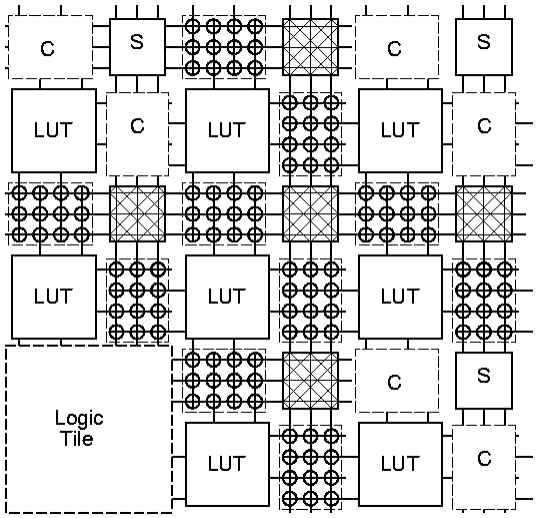

Mapping Ratio?

- How bad is it?
- How much wider do channels have to be?

- Mapping Ratio:
  - detail channel width required / global ch width

Mapping Ratio

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

Domain Structure

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

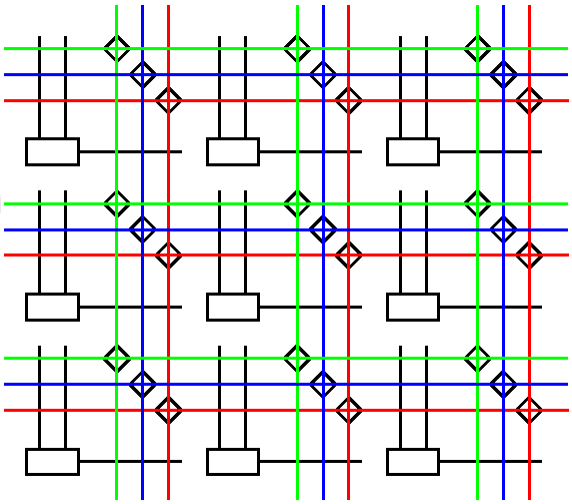

Detail Routing as Coloring

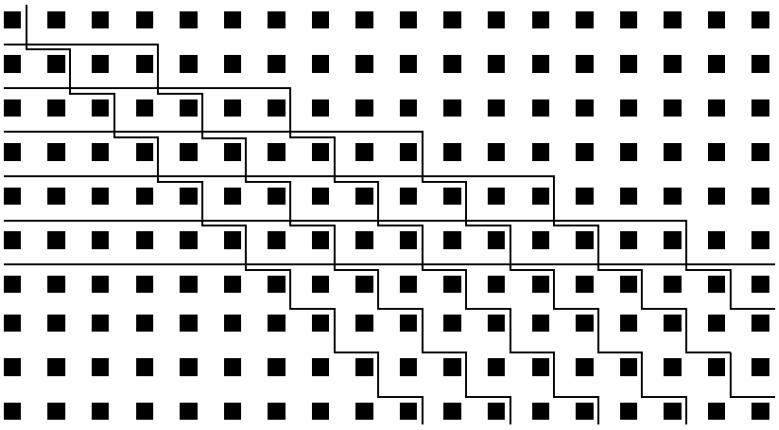

Detail Routing as Coloring

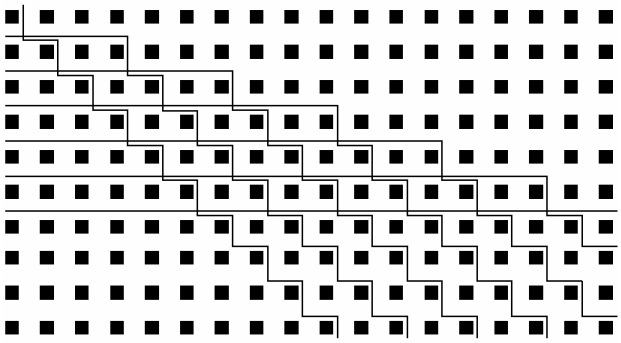

- Global Route channel width = 2
- Detail Route channel width = N
  - Can make arbitrarily large difference

Detail Routing as Coloring

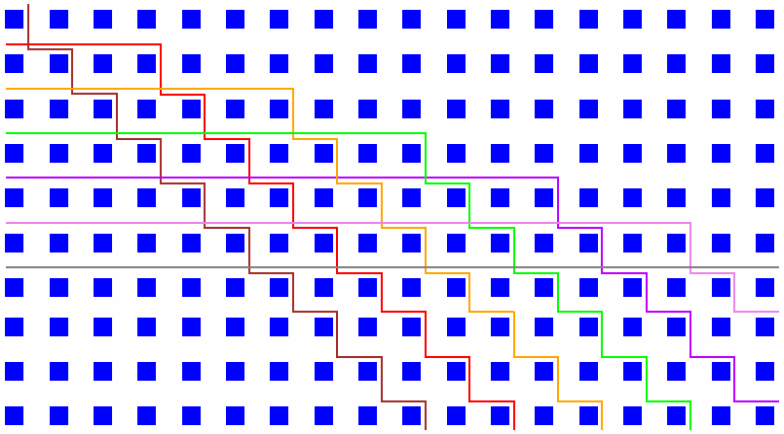

Routability

- 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)

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

Segmentation

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

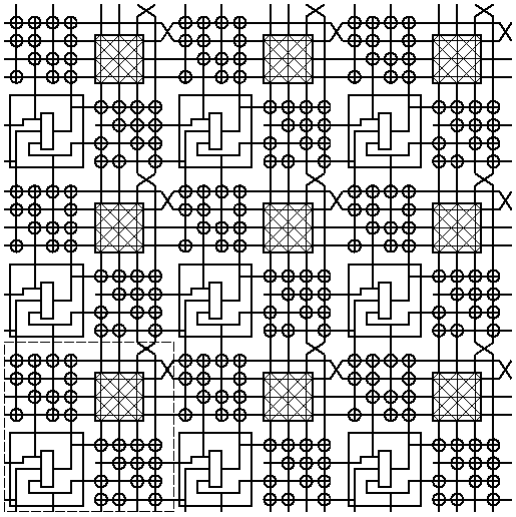

Segmentation

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

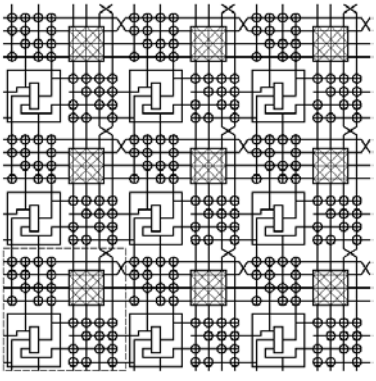

Segmentation

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

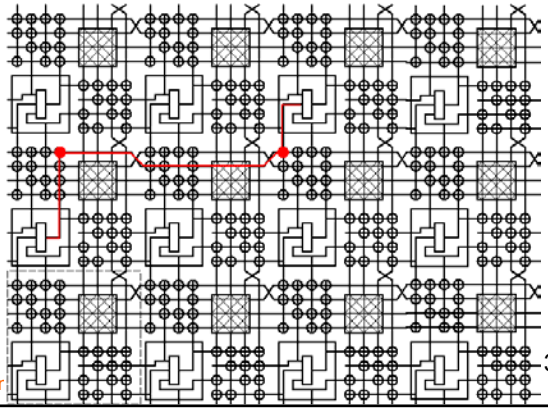

Segmentation: Corner Turn Option

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

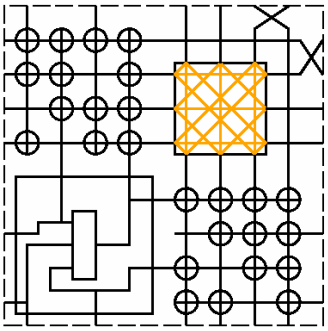

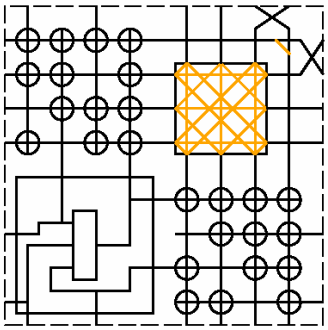

VPR Segment 4 Pix

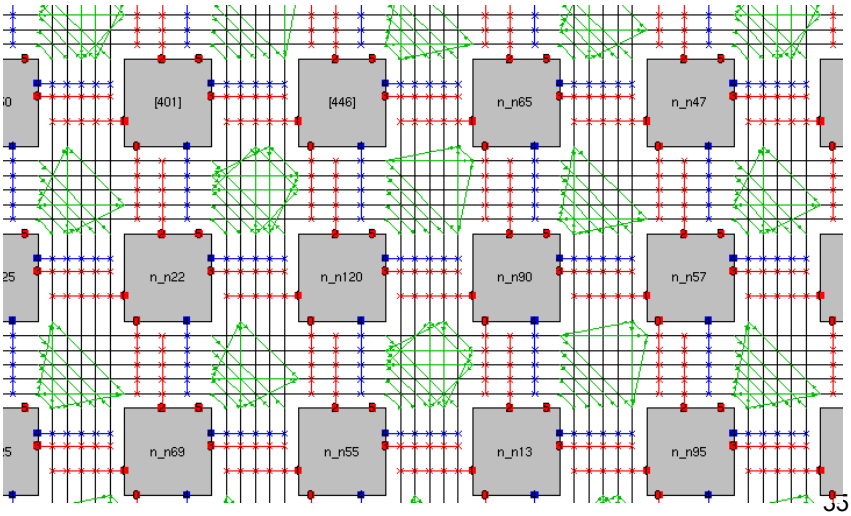

VPR Segment 4 Route

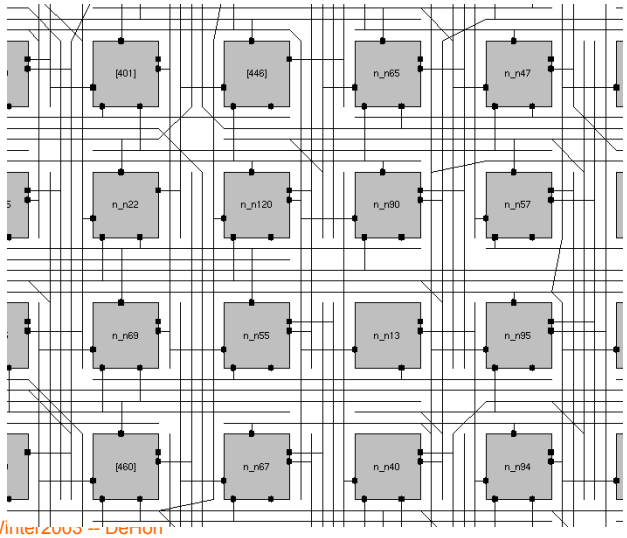

C-Box Depopulation

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

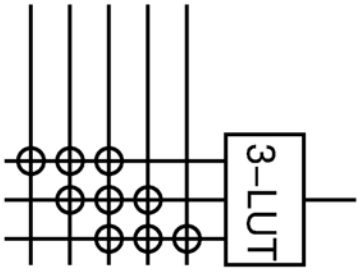

IO Population

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

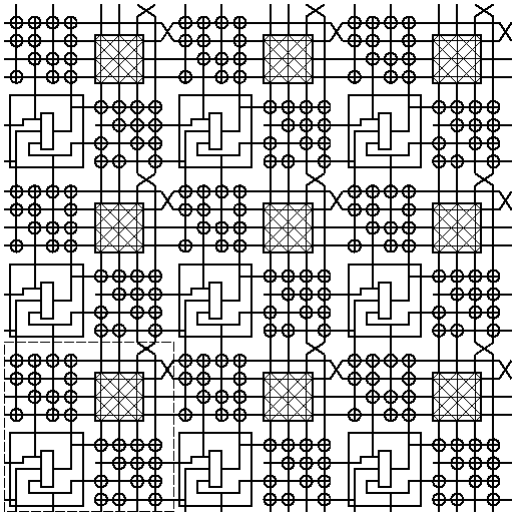

IO Population

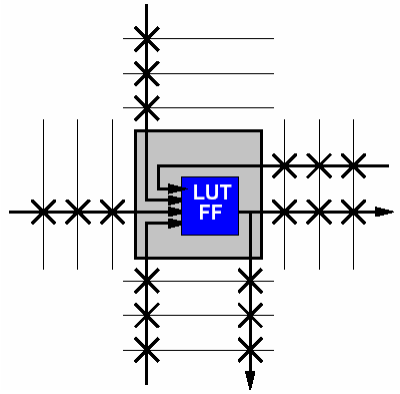

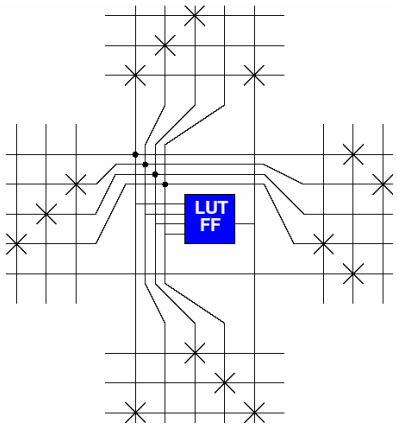

Leaves Not LUTs

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

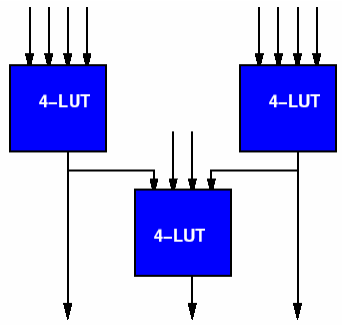

Logic Block

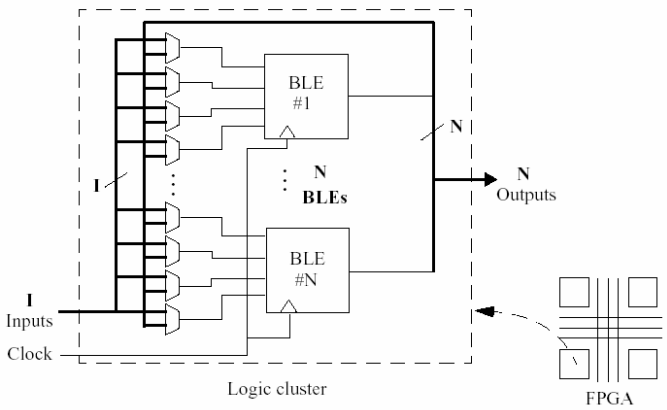

Figure 3: Logic cluster structure.

[Betz+Rose/IEEE D&T 1998]

Cluster Size

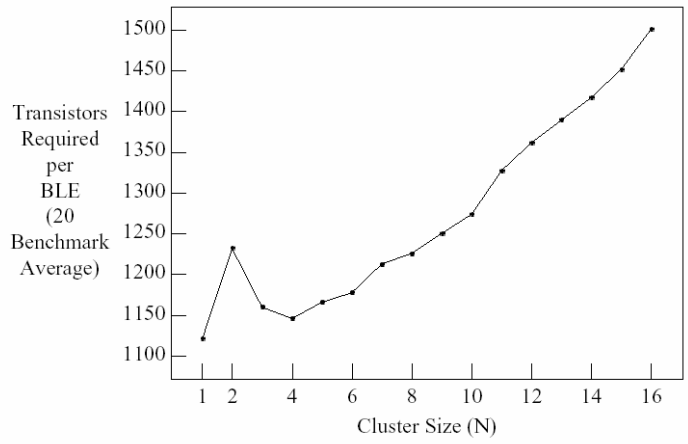

Figure 8: Area-efficiency versus cluster size.

[Betz+Rose/IEEE D&T 1998]

Inputs Required per Cluster

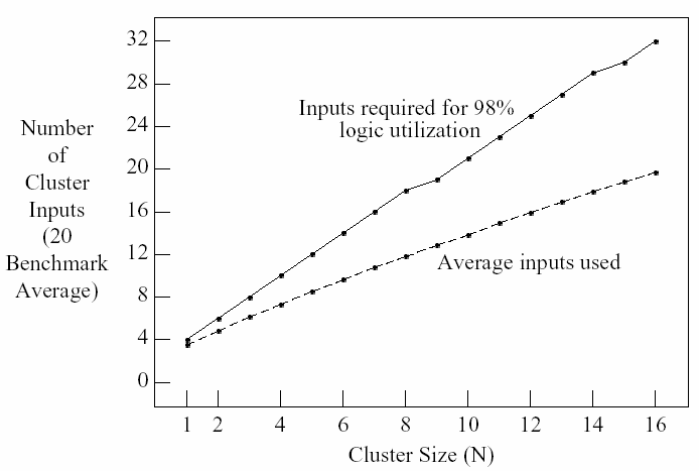

Should it be linear?

[Betz+Rose/IEEE D&T 1998]

Mesh Design Parameters

- Cluster Size
  - Internal organization
- LB IO (Fc, sides)
- Switchbox Population and Topology
- Segment length distribution
- Switch rebuffering

Commercial Parts

XC4K Interconnect

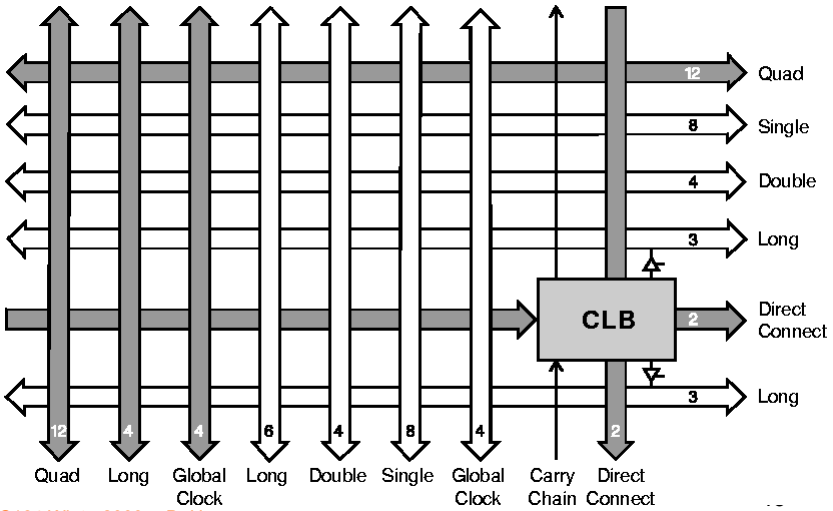

XC4K Interconnect Details

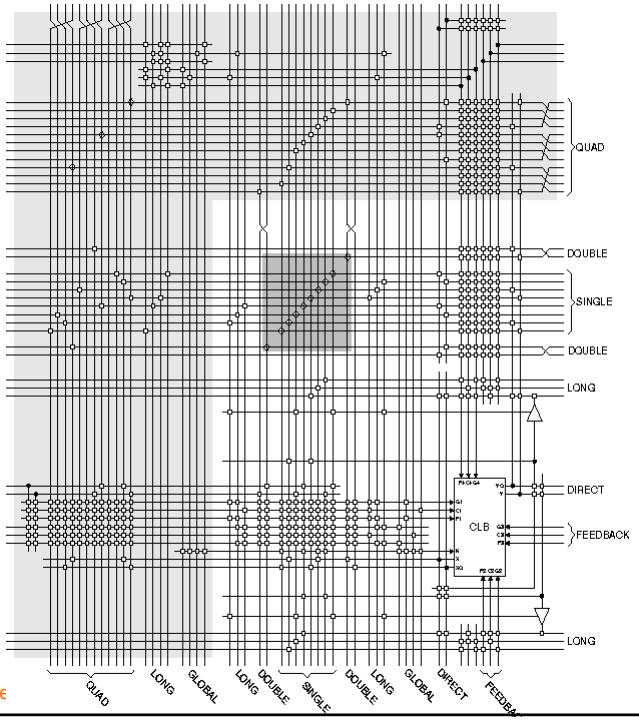

Virtex II

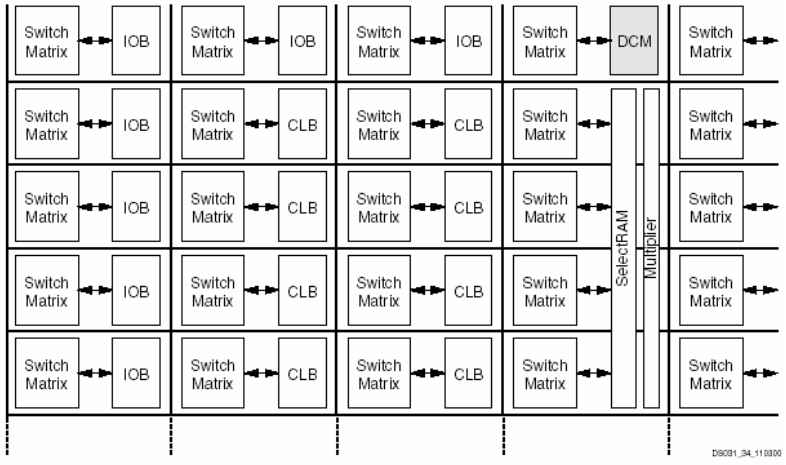

Figure 48: **Routing Resources**

Virtex II Interconnect Resources

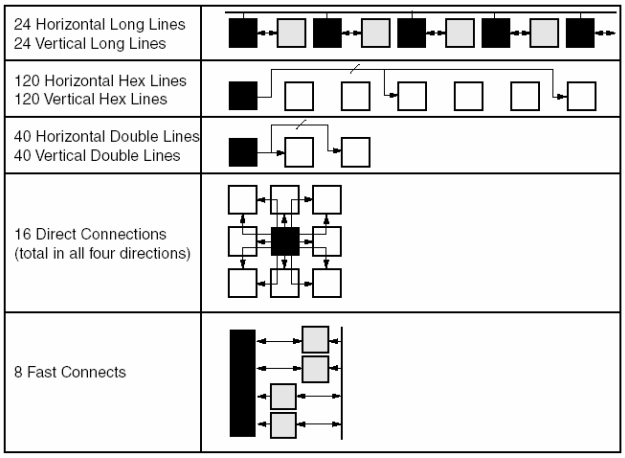

Figure 49: **Hierarchical Routing Resources**

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 $\Omega(N^{p+0.5})$
    - Unbounded global→detail mapping ratio
    - Detail routing NP-complete

Big Ideas
[MSB-1 Ideas]

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