
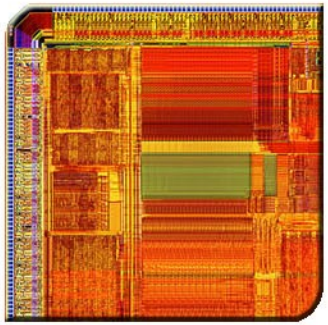
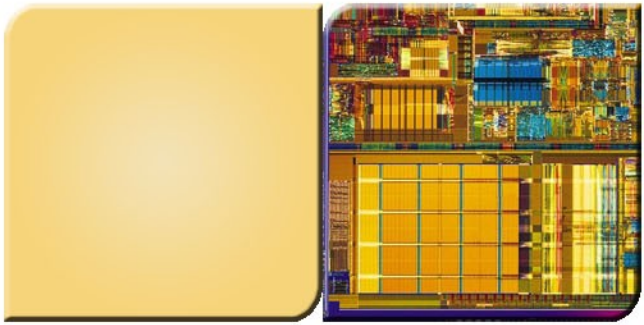

High Performance Non-blocking Switch Design in 3D Die-Stacking Technology

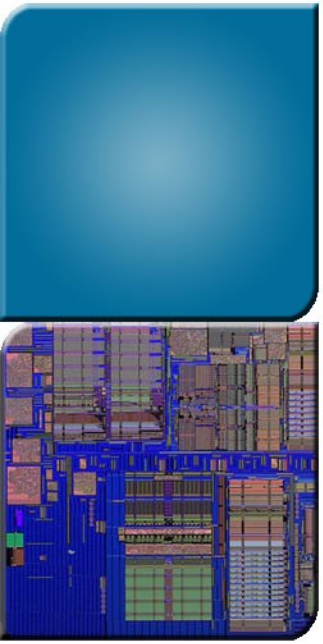

Dean L. Lewis, Sudhakar Yalamanchili,
Hsien-Hsin S. Lee

IEEE Computer Society Annual Symposium on VLSI
Tampa, Fl, 2009

http://arch.ece.gatech.edu/mars.html

Georgia Tech MARS

3D Integration

- Multiple layers of silicon

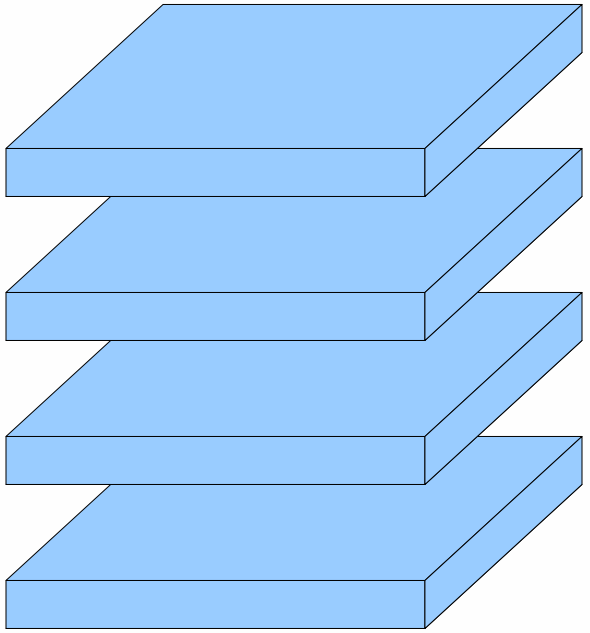

3D Integration

- Multiple layers of silicon
- Interconnected with TSVs
  - Etched through thinned wafers

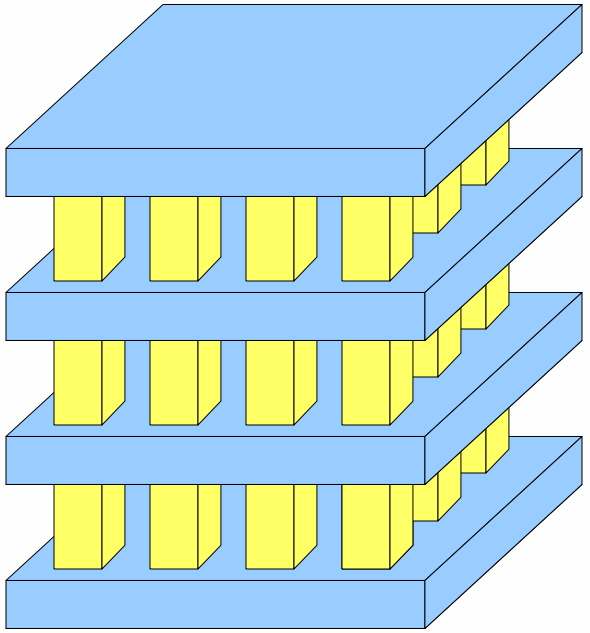

3D Integration

- Multiple layers of silicon
- Interconnected with TSVs
  - Etched through thinned wafers
- Several integration options
  - Technology
    - Different processing technologies are tightly integrated

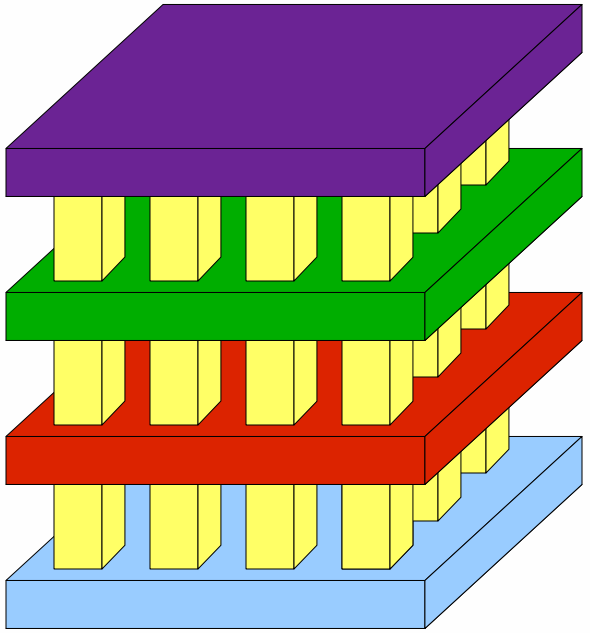

3D Integration

- Multiple layers of silicon
- Interconnected with TSVs
  - Etched through thinned wafers
- Several integration options
  - Technology
    - Different processing technologies are tightly integrated
  - Architecture
    - Blocks split across layers

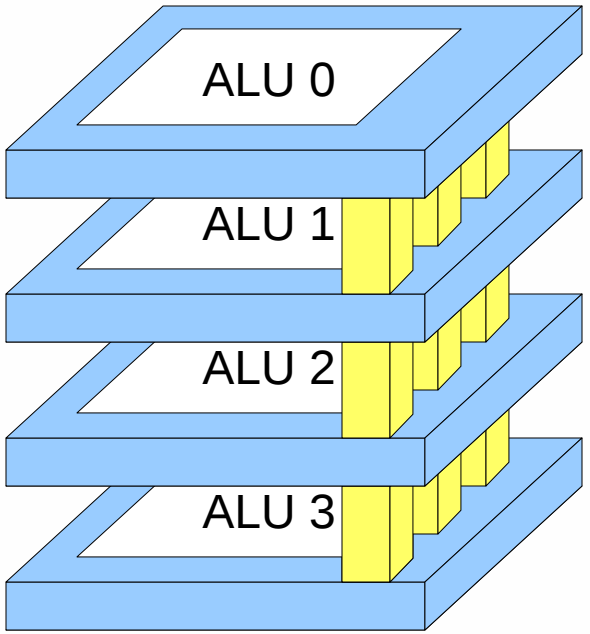

3D Integration

- Multiple layers of silicon
- Interconnected with TSVs
  - Etched through thinned wafers
- Several integration options
  - Technology
    - Different processing technologies are tightly integrated
  - Architecture
    - Blocks split across layers
  - Circuit
    - Transistors split across

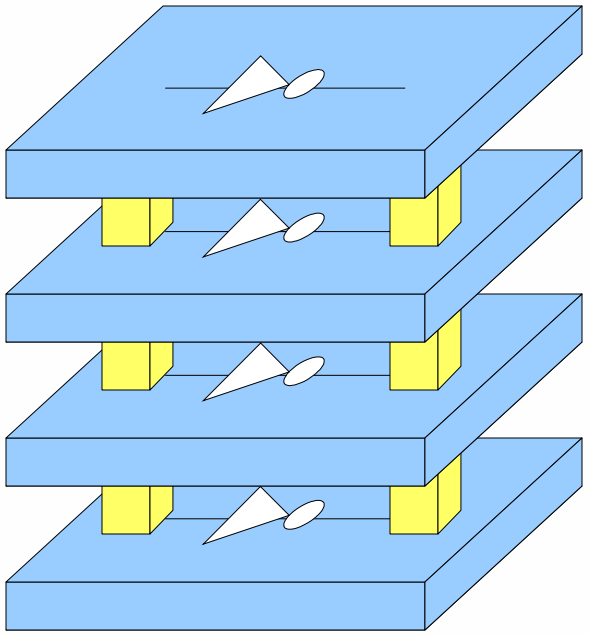

Switch Design

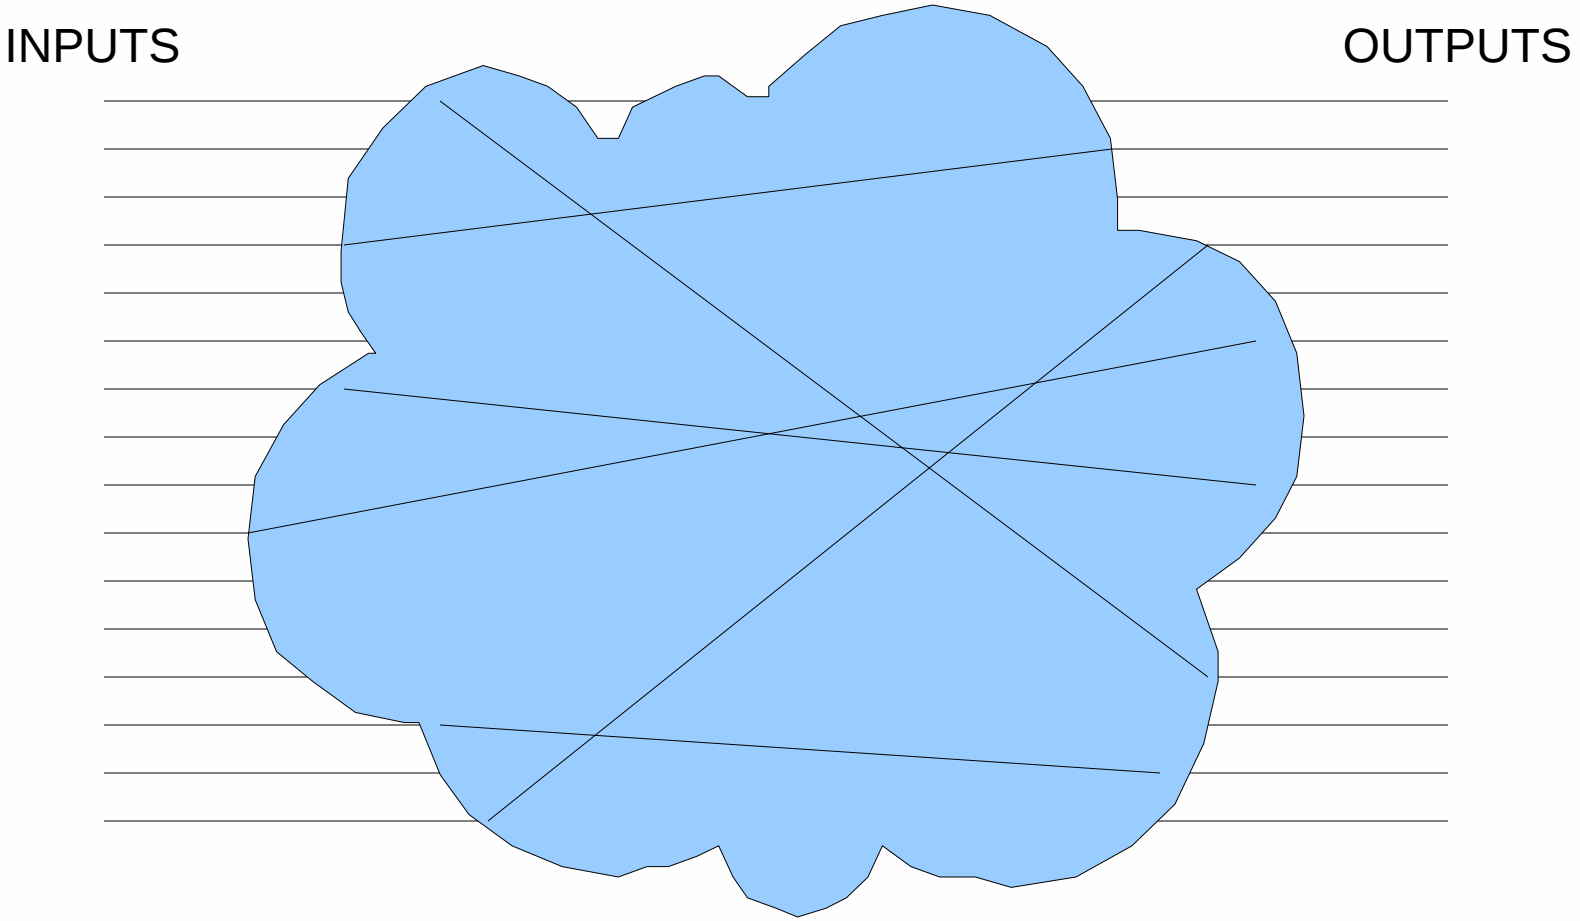

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections

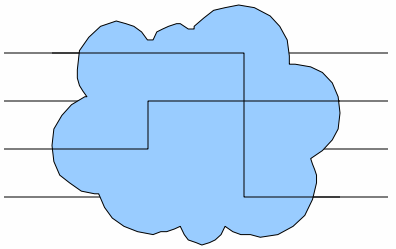

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections

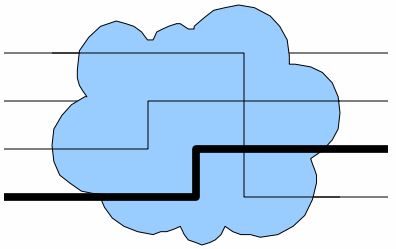

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged

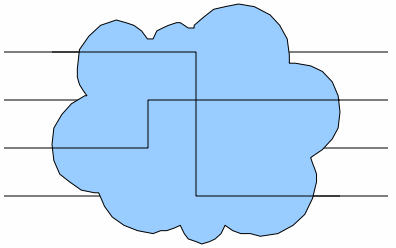

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged

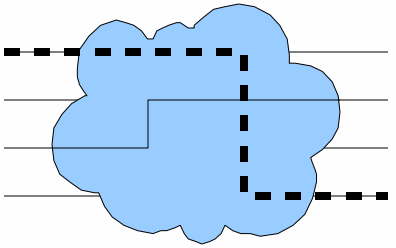

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged

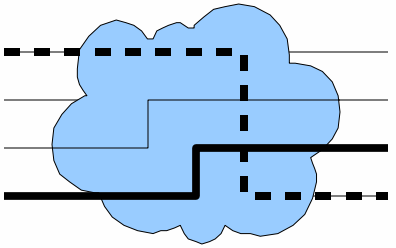

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged
- Blocking
  - Making a new connection may disconnect an existing

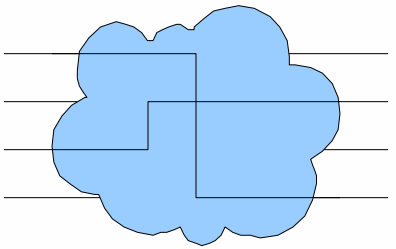

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged
- Blocking
  - Making a new connection may disconnect an existing

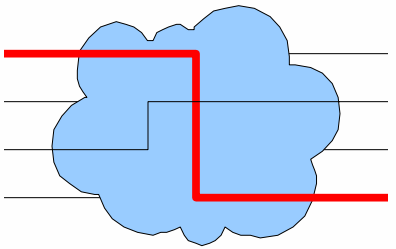

Connect(3,2)

Switch Design

- Strictly non-blocking
  - Any free input/output pair can be connected
  - No effect on existing connections
- Weakly non-blocking
  - Any free input/output pair can be connected
  - Existing connections may need to be rearranged
- Blocking
  - Making a new connection may disconnect an existing

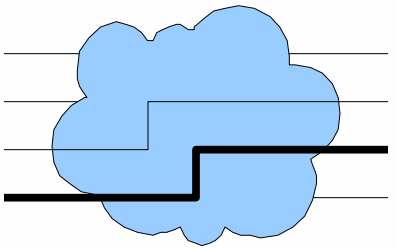

Connect(3,2)

Crossbar

- Strictly Non-blocking
- Lots of switch points
- Little wiring

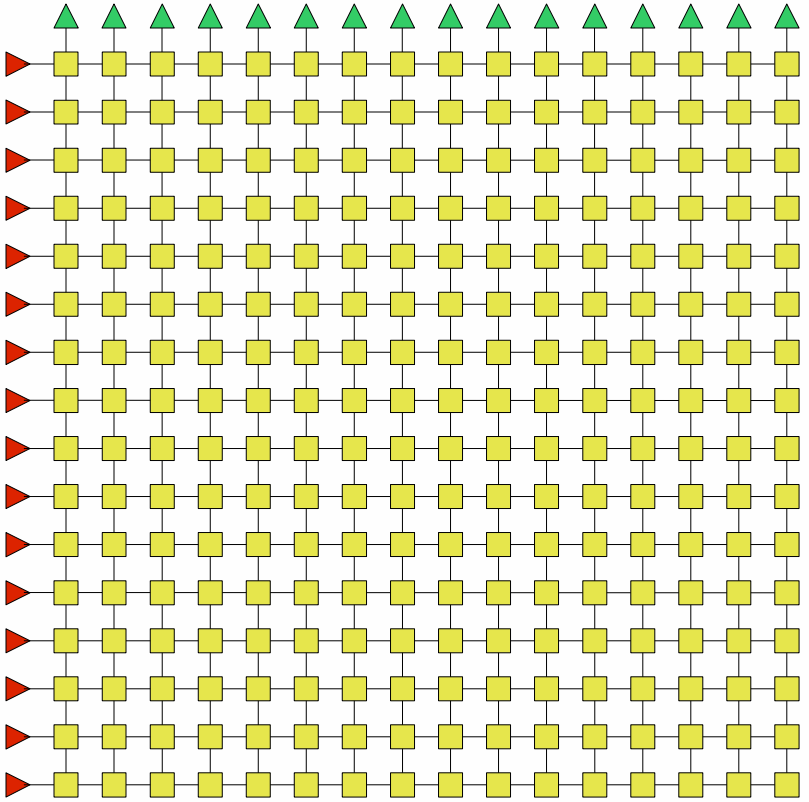

Crossbar

- Strictly Non-blocking
- Lots of switch points
- Little wiring
- Very long critical path

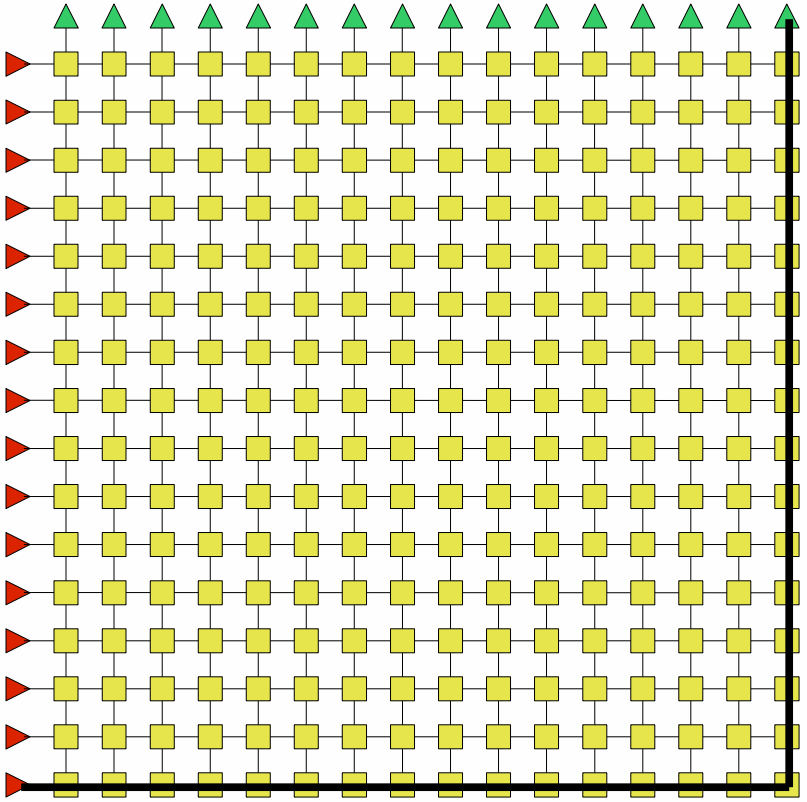

Connect(15,15)

Row-Split 3D Crossbar

- Cannot target wiring directly
  - Too short and direct
- Instead, change critical path using 3D design
- Split rows in half, placing each half on a different layer
- Add multiplexers to bypass unused half
- Critical path wirelength is reduced by 25%

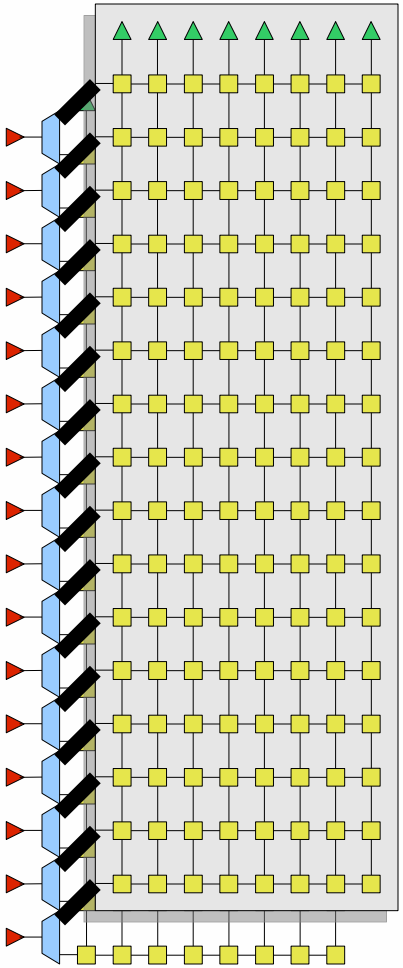

Row-Split 3D Crossbar

- Cannot target wiring directly
  - Too short and direct
- Instead, change critical path using 3D design
- Split rows in half, placing each half on a different layer
- Add multiplexers to bypass unused half

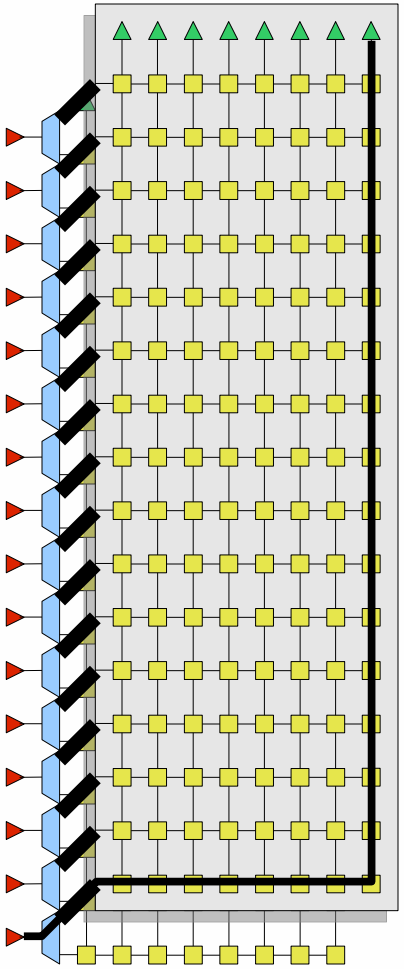

Connect(15,15)

Column-Split 3D Crossbar

- Complimentary to row-split design

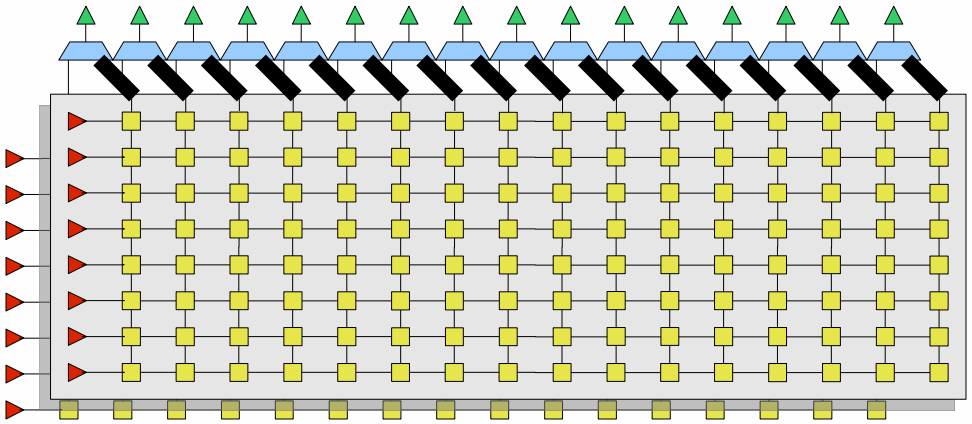

Column-Split 3D Crossbar

- Complimentary to row-split design

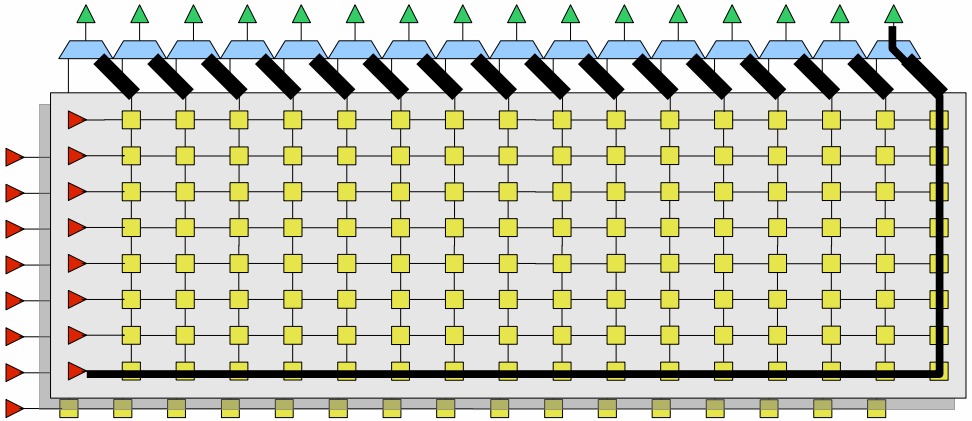

Connect(15,15)

Column-Split 3D Crossbar

- Complimentary to row-split design

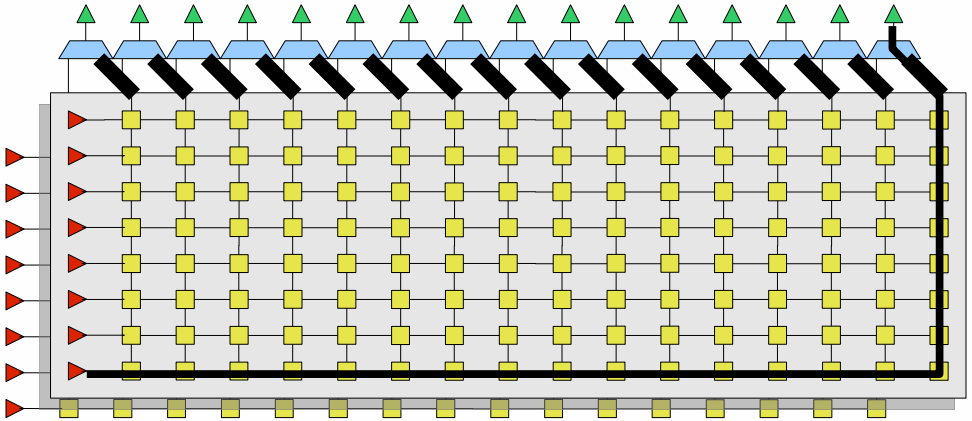

Connect(15,15)

- Row-split and column-split designs can be applied simultaneously in a four-layer stack
  - 50% reduction in critical path wirelength

Crossbar – System View

- Crossbars are great at small sizes
- This makes them good for a NoC design
  - Individual components (processors, routers, caches) are 3D
  - NoC architecture is planar (shown)
  - Of course, both components and NoC could be 3D

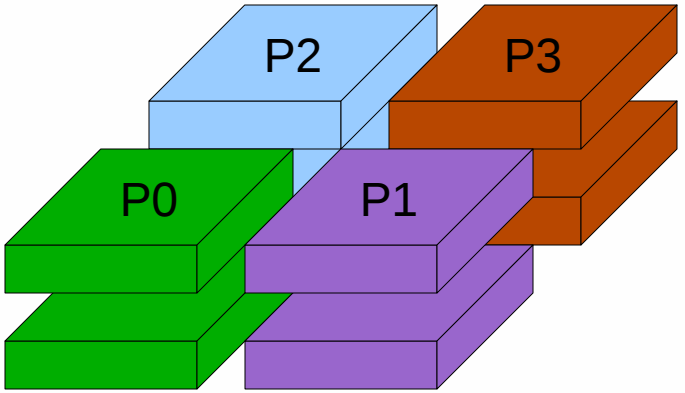

Crossbar – System View

- Crossbars are great at small sizes
- This makes them good for a NoC design
  - Individual components (processors, routers, caches) are 3D
  - NoC architecture is planar (shown)
  - Of course, both components and NoC could be 3D

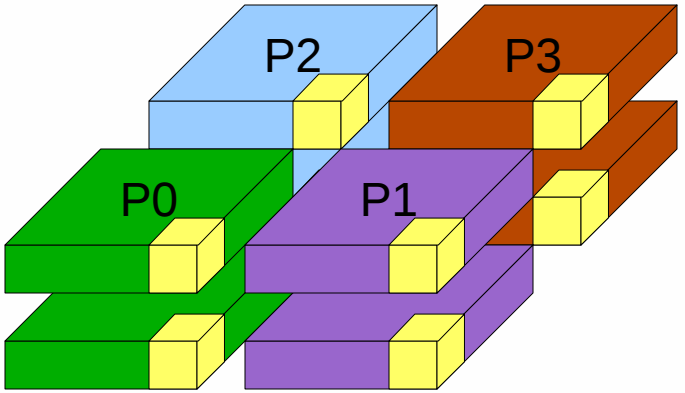

MIN

- Multistage interconnect network
- Weakly non-blocking
- Less hardware but more complex wiring than crossbars
- Usually more appropriate for larger switches
- Example system: many-input router chip

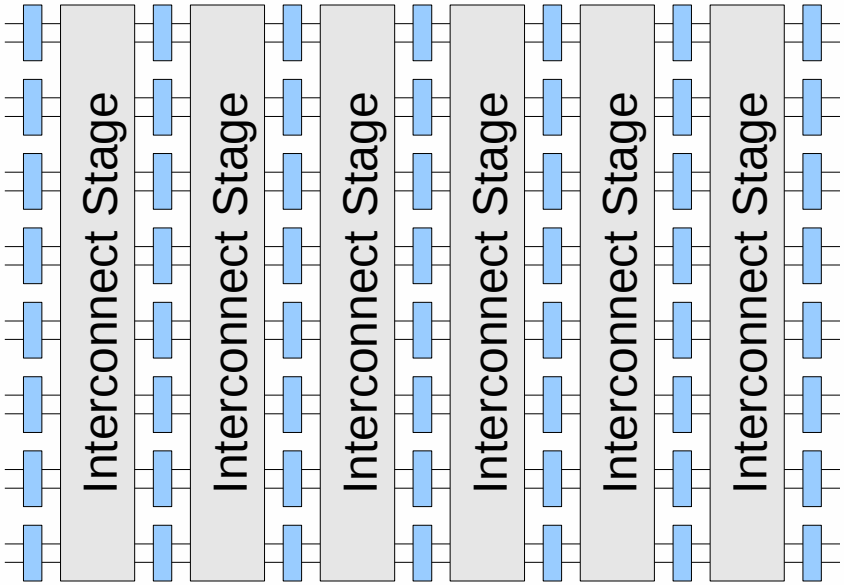

Perfect Shuffle

- Interconnect pattern – deck of cards perfectly shuffled
- Lots of overlapping wiring, so lots of wiring tracks

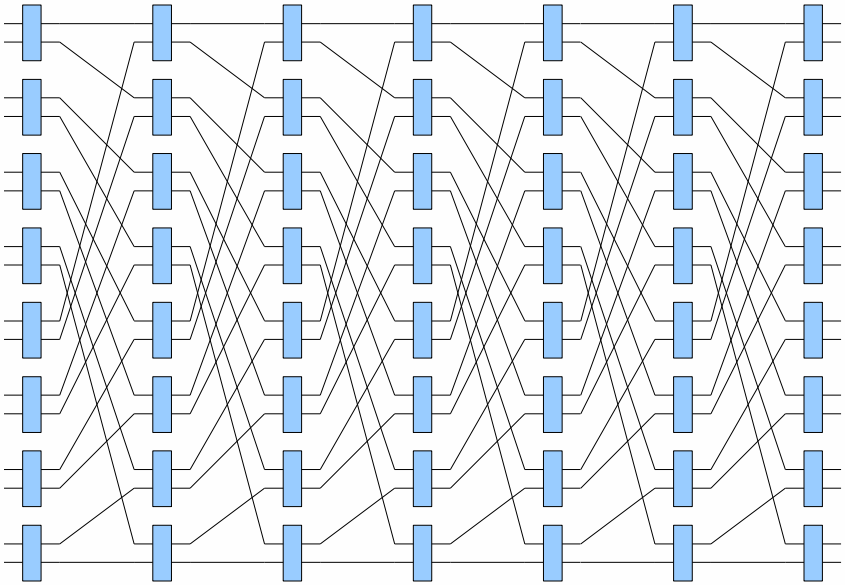

Perfect Shuffle

- Interconnect pattern – deck of cards perfectly shuffled
- Lots of overlapping wiring, so lots of wiring tracts
- Two long hops on critical path: (3,7) and (4,0)

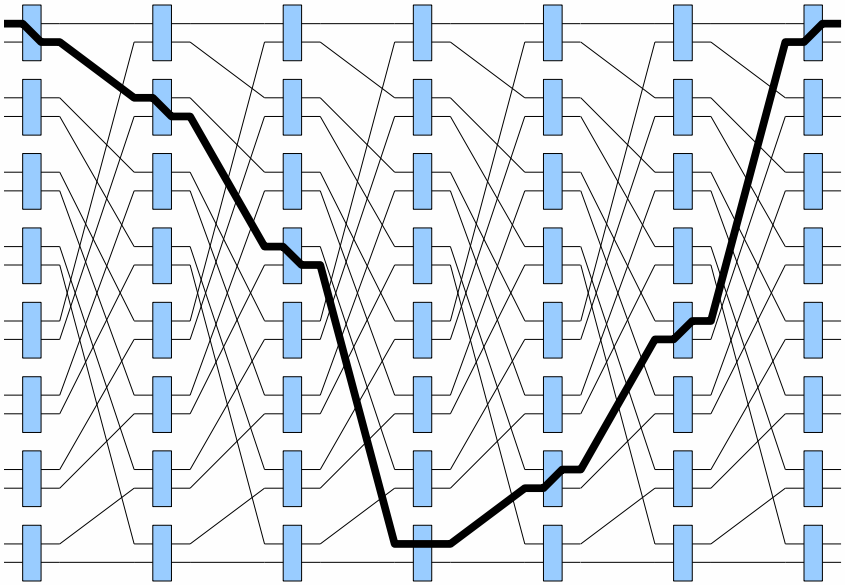

Perfect Shuffle – Wiring

- Pattern in the overlaps that 3D can target to relieve congestion

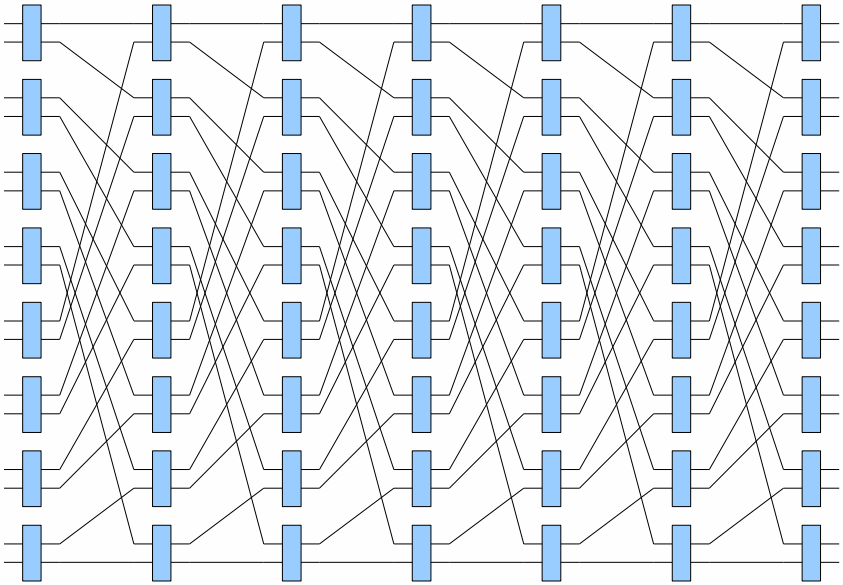

Perfect Shuffle – Wiring

- Pattern in the overlaps that 3D can target to relieve congestion
- Nodes 0 and 4 both connect to 0 and 1 in the next stage

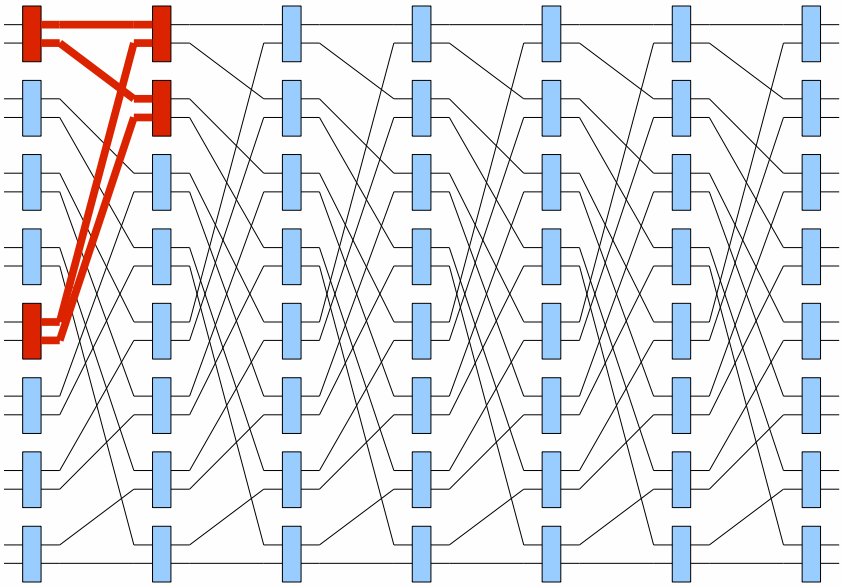

Perfect Shuffle – Wiring

- Pattern in the overlaps that 3D can target to relieve congestion
- Nodes 0 and 4 both connect to 0 and 1 in the next stage
- Similarly, nodes 3 and 7 share 6 and 7

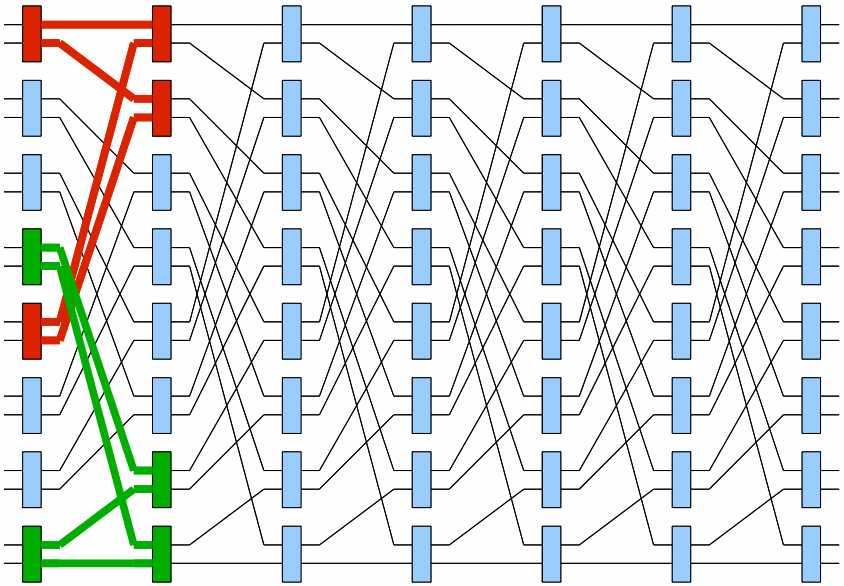

3D Perfect Shuffle

- Split MIN into bottom half and top half
- Place top half over bottom half

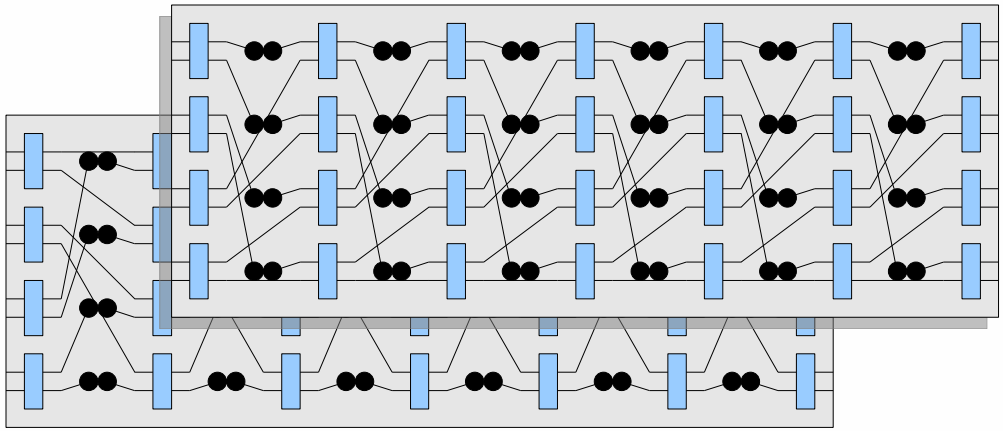

3D Perfect Shuffle

- Split MIN into bottom half and top half
- Place top half over bottom half
- Those old longest paths are now much shorter

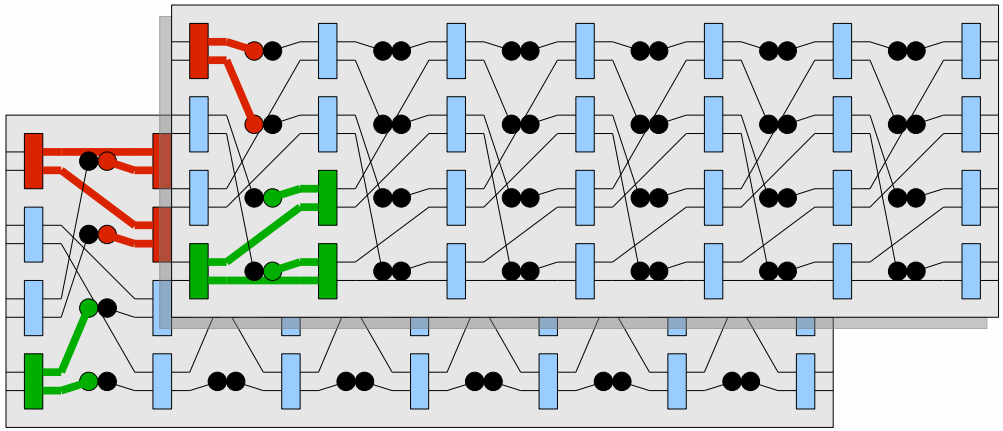

3D Perfect Shuffle

- Split MIN into bottom half and top half
- Place top half over bottom half
- Those old longest paths are now much shorter
- Consequently, critical path is reduced

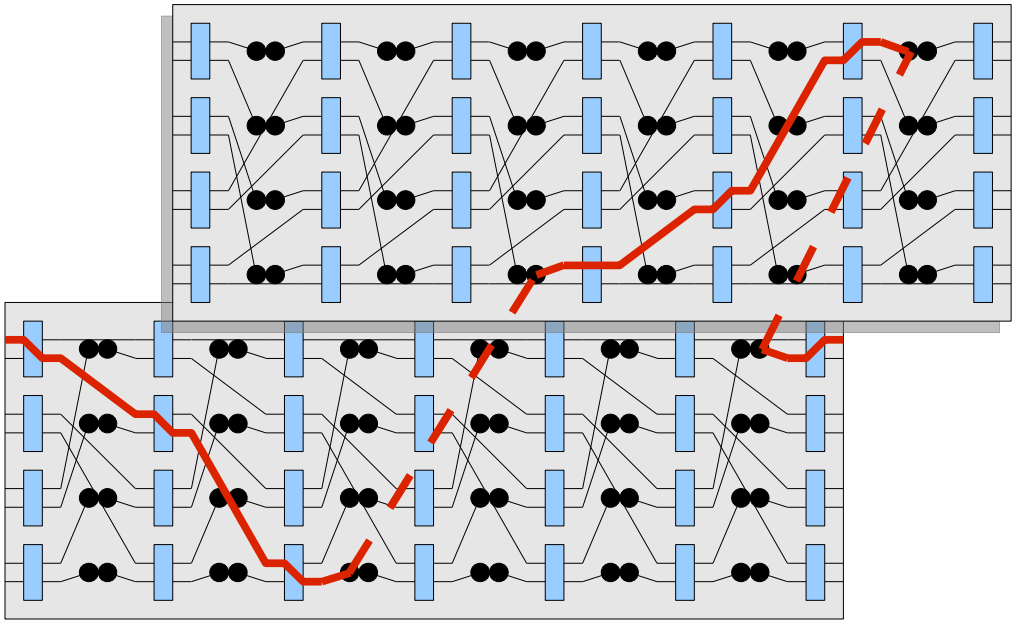

Butterfly Shuffle

- Interconnect pattern – wiring destination doubles at each stage
- Lots of overlapping wires in the outer stages

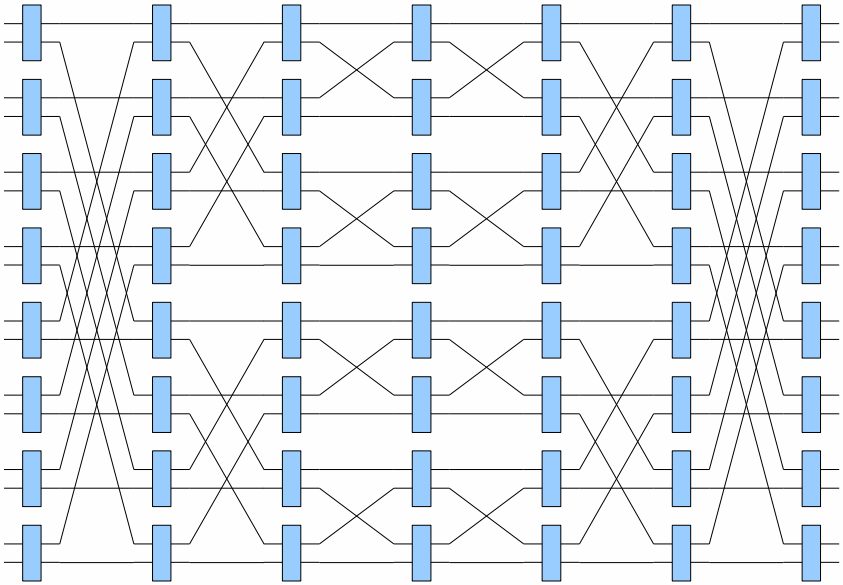

Butterfly Shuffle

- Interconnect pattern – wiring destination doubles at each stage
- Lots of overlapping wires in the outer stages

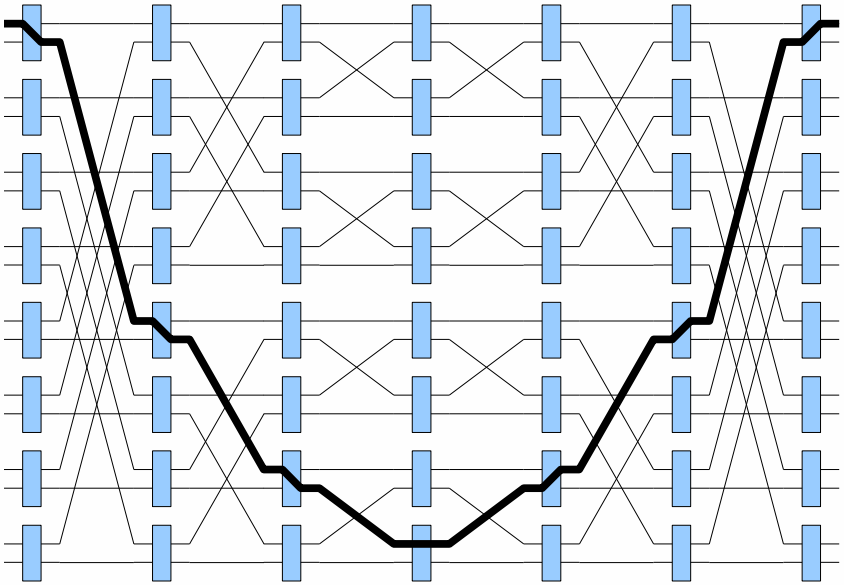

Butterfly Shuffle

- Interconnect pattern – wiring destination doubles at each stage
- Lots of overlapping wires in the outer stages
- These wires do not share switch resources

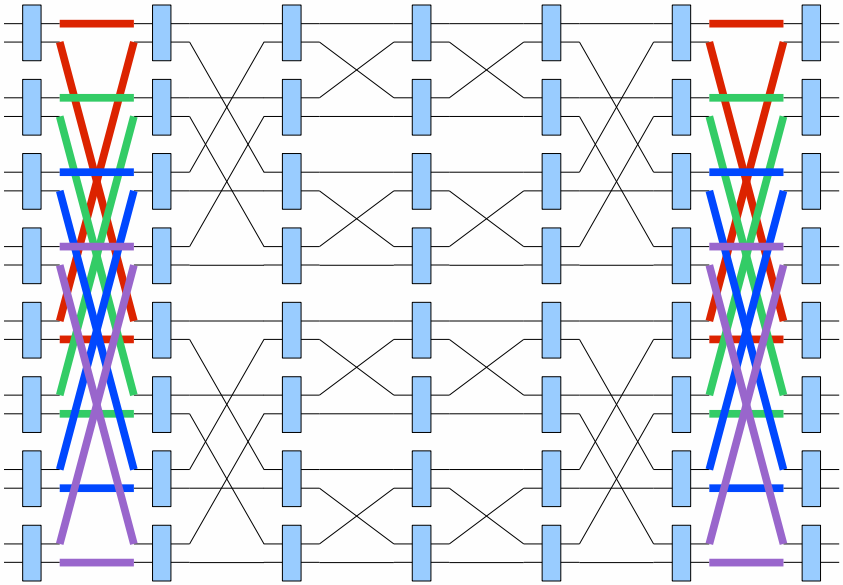

Butterfly Shuffle

- Interconnect pattern – wiring destination doubles at each stage
- Lots of overlapping wires in the outer stages
- These wires do not share switch resources

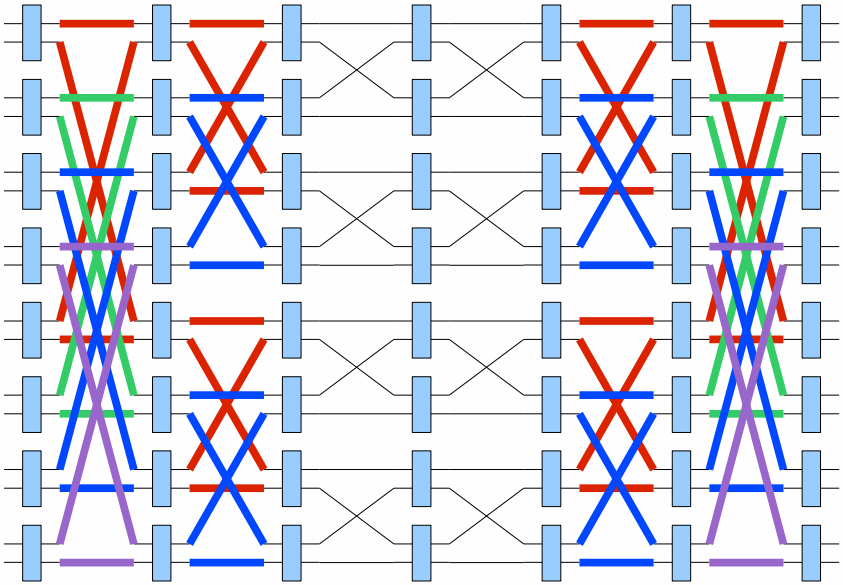

3D Butterfly Shuffle

- Use a “bit-split” partition
- This removes half of overlapping wires from **each** outer stage

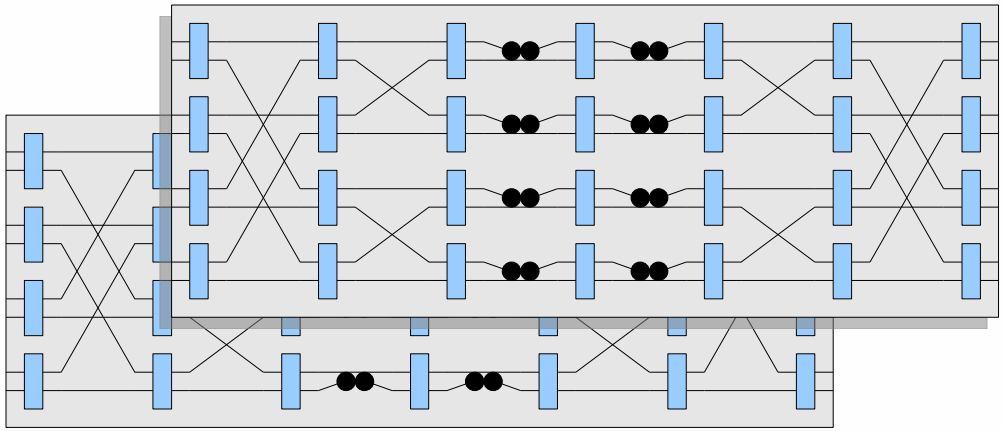

3D Butterfly Shuffle

- Use a “bit-split” partition
- This removes half of overlapping wires from **each** outer stage
- Now many fewer wiring tracts required per stage

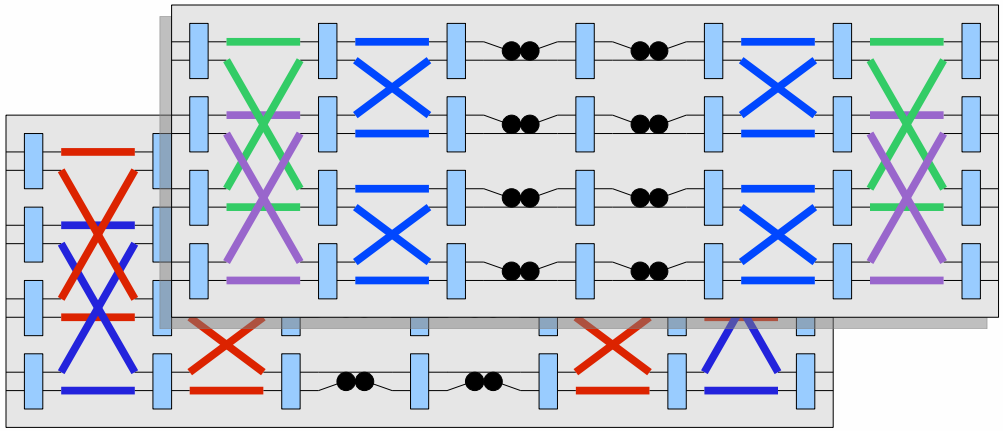

3D Butterfly Shuffle

- Use a “bit-split” partition
- This removes half of overlapping wires from **each** outer stage
- Now many fewer wiring tracts required per stage
- Critical path is significantly reduced

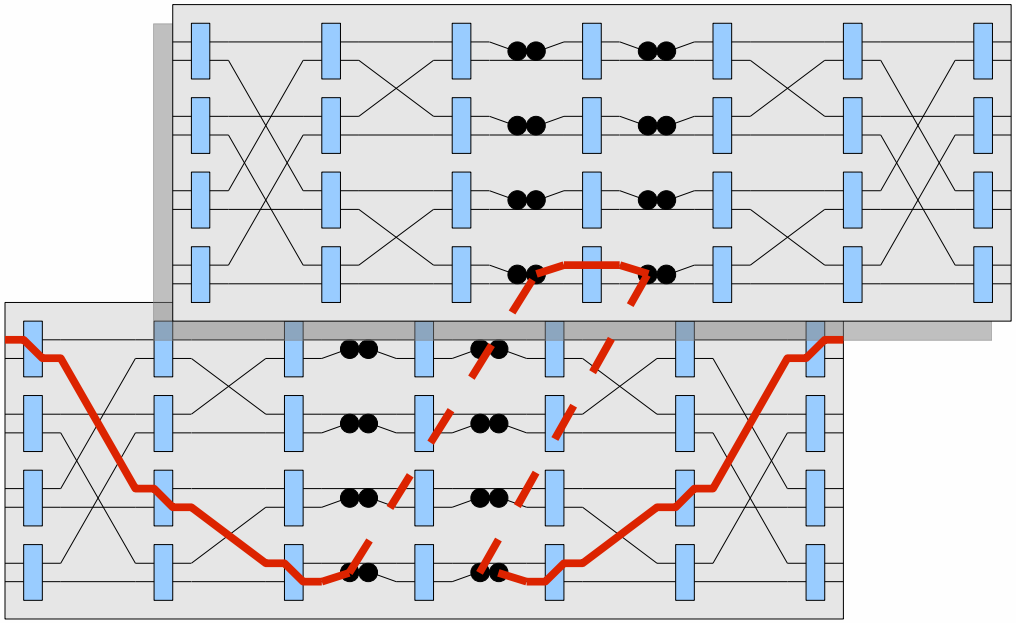

Experimental Setup

- Switch cell
  - 130nm technology
- TSV
  - Tezzaron 3D process
  - 10um x 1.7um x 1.7um
- HSPICE simulation
  - Lvl49 130nm transistor
  - Stick diagrams for wiring stages
  - Intel 130nm wire parasitics

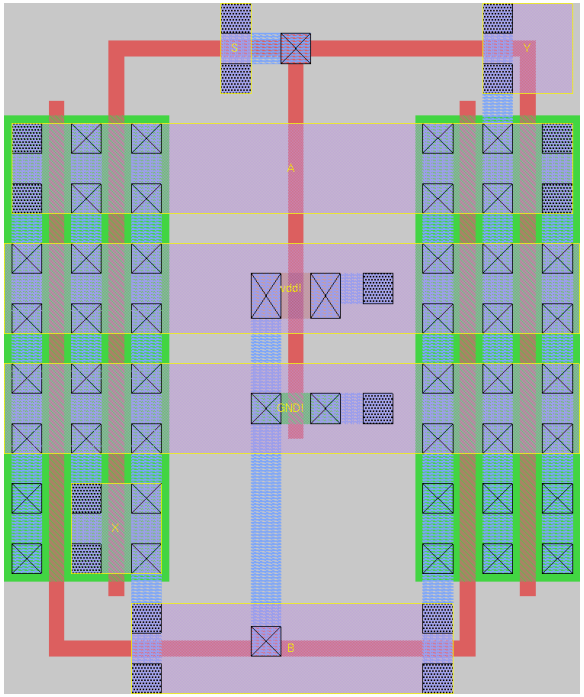

Results

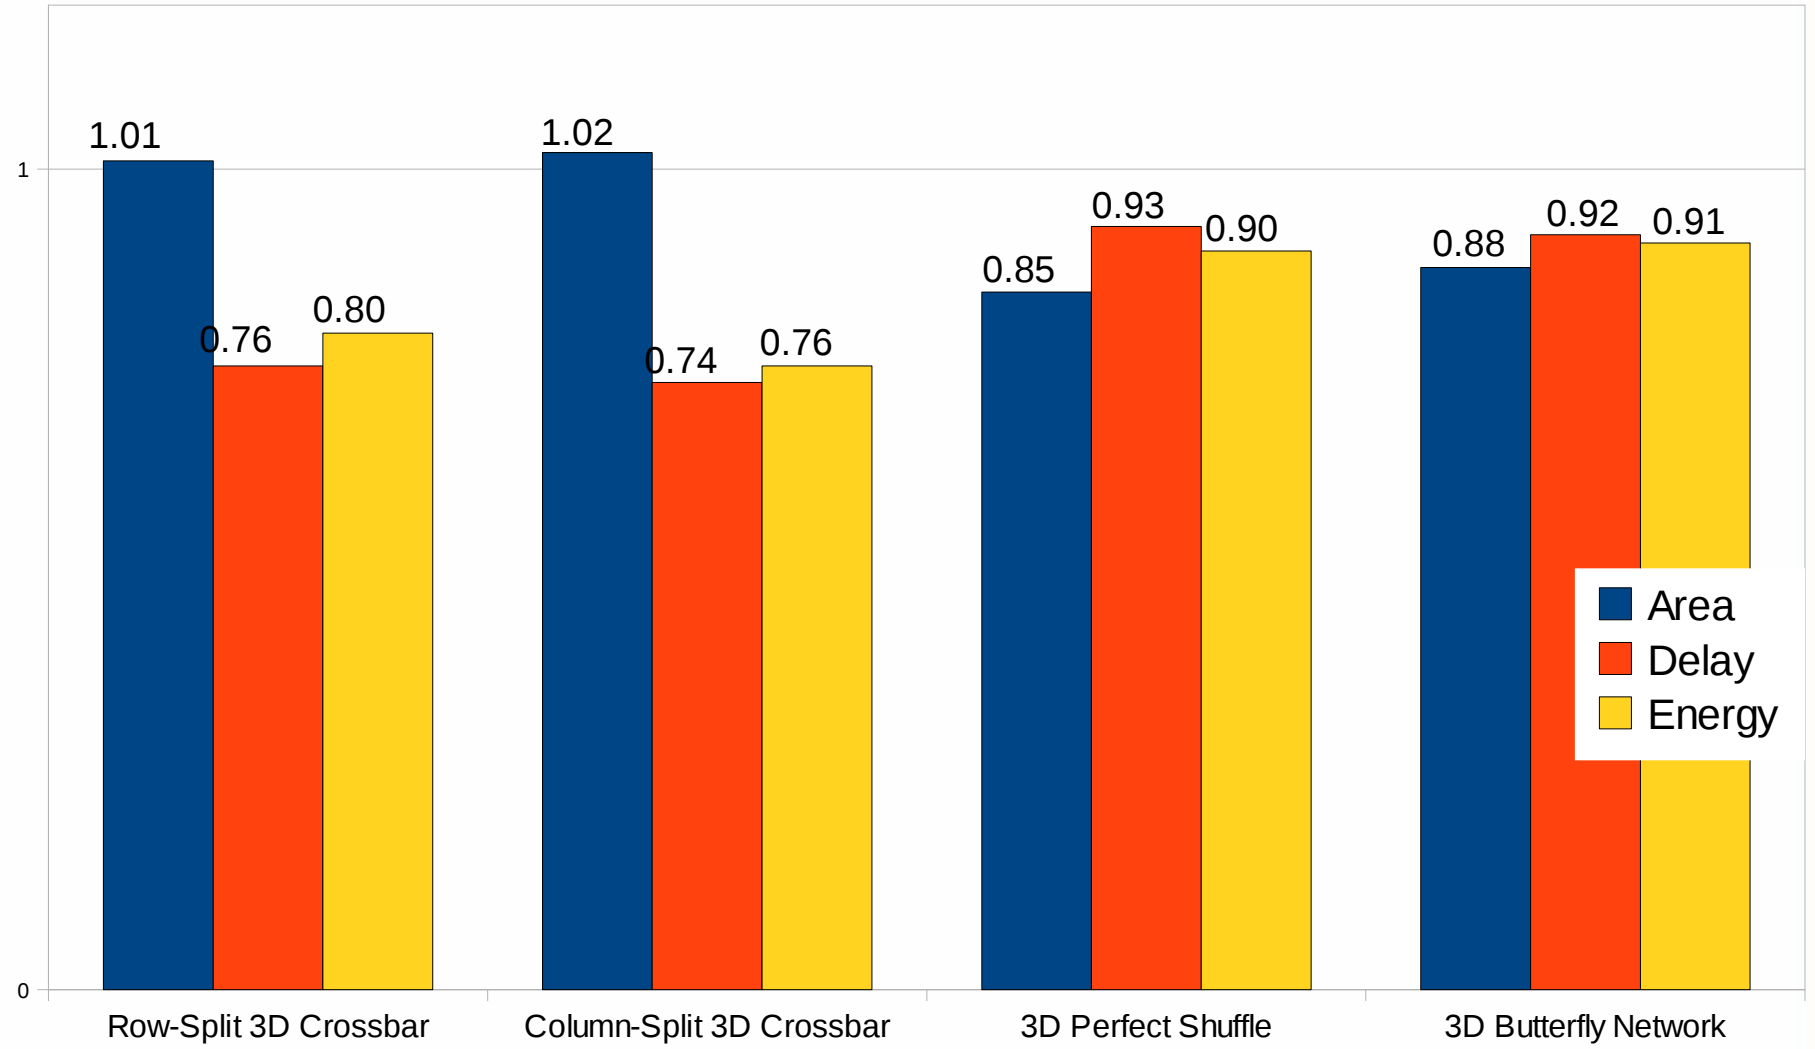
1
0
1.01
0.76
0.80
1.02
0.74
0.76
0.85
0.93
0.90
0.88
0.92
0.91
Area
Delay
Energy
Row-Split 3D Crossbar
Column-Split 3D Crossbar
3D Perfect Shuffle
3D Butterfly Network

Conclusion

- Switches are a key component of networks both large and small
- 3D integration can be used to effectively improve switch performance
- 3D circuit design effectively targets and reduces long critical paths
- Generally, 3D integration reduces area, delay, and energy **simultaneously**
- These reductions are significant
  - 8% - 26%
- 3D will be a powerful design optimization technology in near-future products



Thank you!

http://arch.ece.gatech.edu/mars.html

Georgia Tech MARS