Digital System Design

by
Dr. Lesley Shannon
Email: lshannon@ensc.sfu.ca
Course Website: http://www.ensc.sfu.ca/~lshannon/
Adapted from Dr. Steve Wilton's FPGA talk

Simon Fraser University

Slide Set: 4
Date: January 26, 2009

Slide Set Overview

- What's under the hood of an FPGA
  - Based on a talk originally given by Dr. Wilton, UBC
  - Highlights concepts applicable to both Altera and Xilinx FPGA architecture
    - This could affect your choice of purchase

Implementing Logic Circuits

- Different ways to implement logic:
  - Use discrete parts (eg. 7400 devices)
  - Design a custom chip
  - Use a (mask-programmed) Gate Array
  - Use a Programmable Logic Array (PLA)
  - Use a Complex Programmable Logic Device (CPLD)
  - Use a Field-Programmable Gate Array (FPGA)
- A couple slides on Discrete Parts and Custom Chip, then I'll focus on FPGAs
  - what we use in the lab (and in industry)

Using Discrete Parts

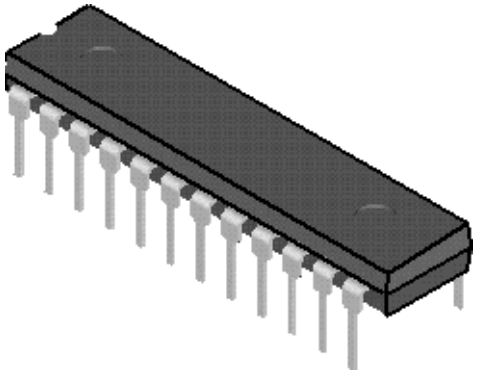

(a) Dual-inline package

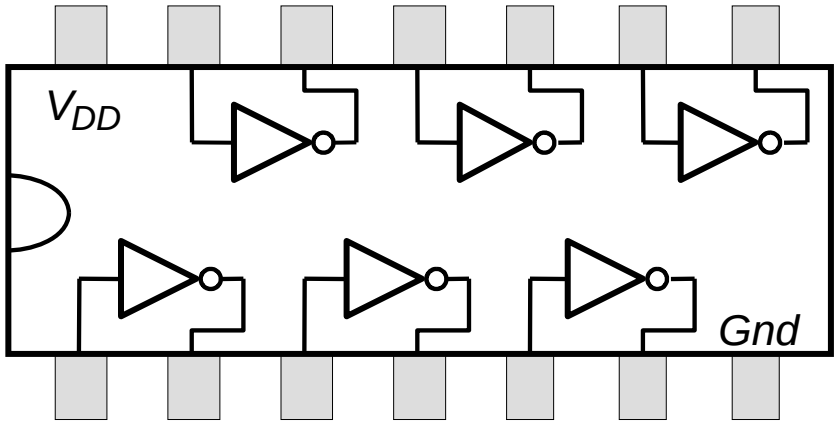

(b) Structure of 7404 chip

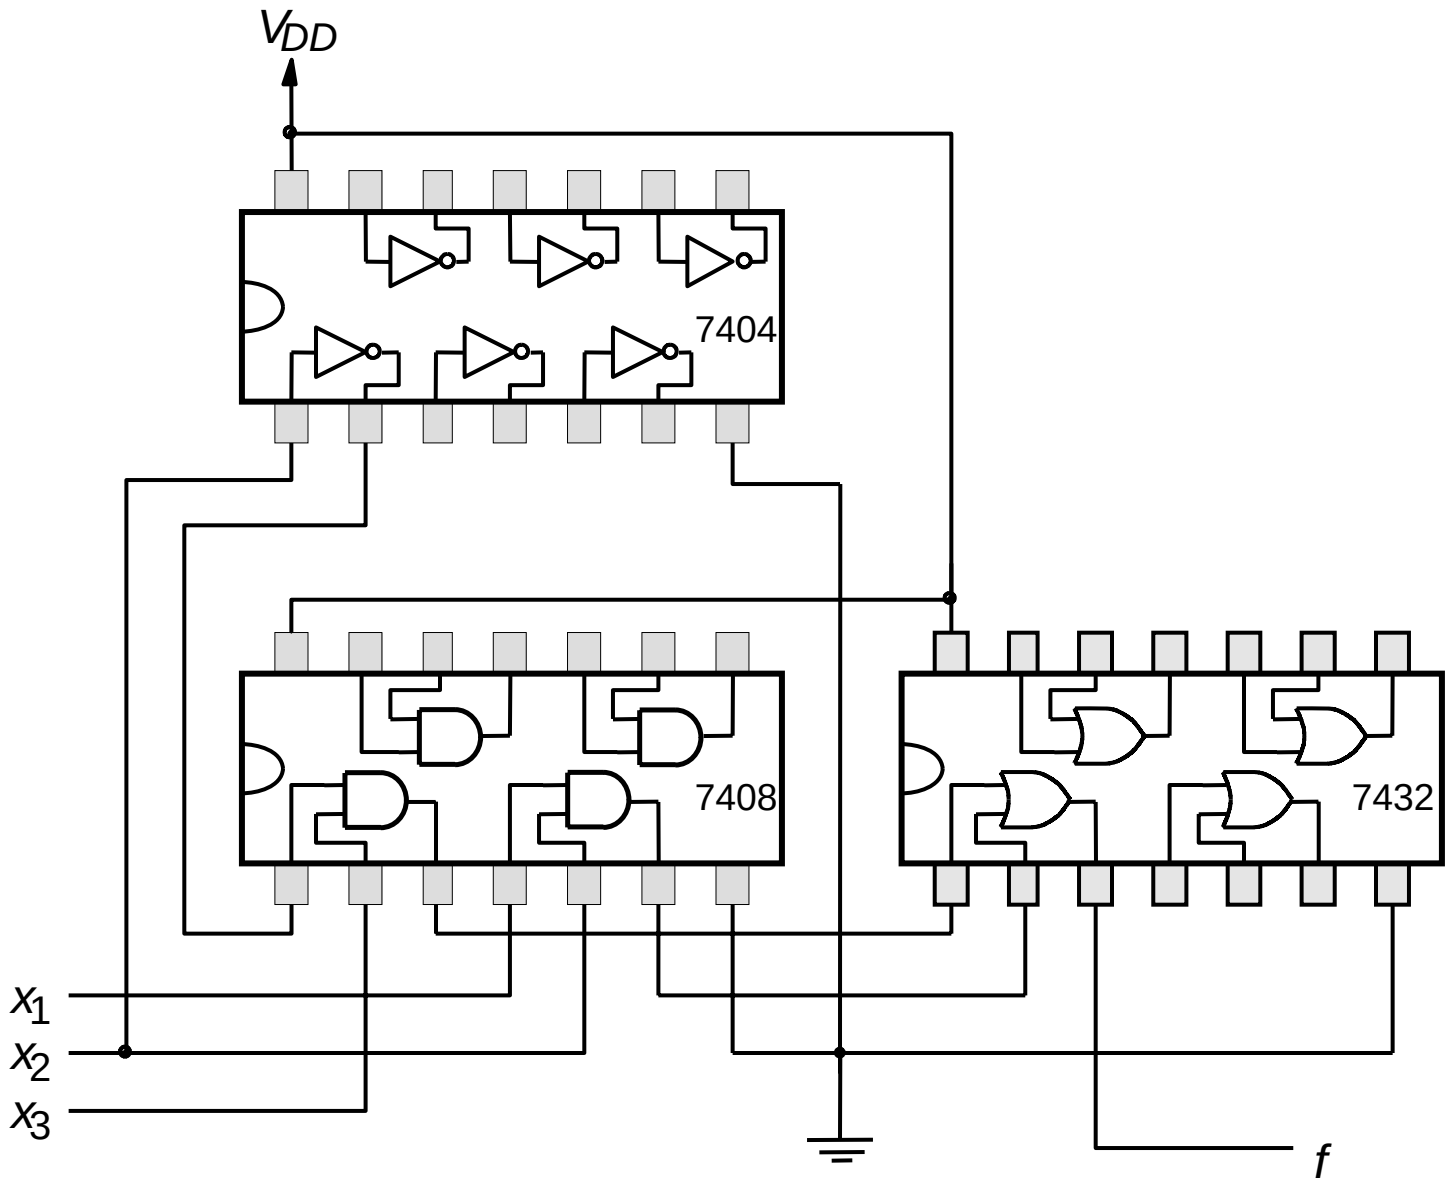

Implementation of $f = \bar{x}_2 x_3 + x_1 x_2$

Designing a Custom Chip

- Very expensive, and time consuming (> $1M just for the mask costs)
- Only used for the most high-speed or low-power applications
- More on this in ensc 450

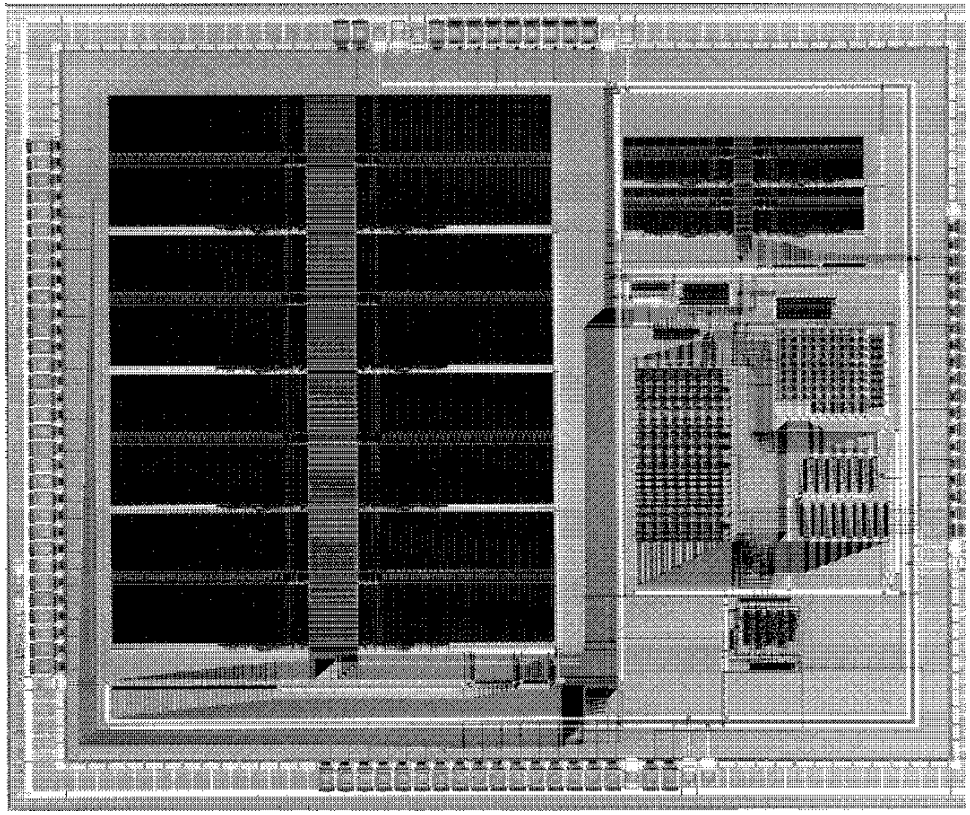

FPGAs and CPLDs

- FPGA: Field-Programmable Gate Array
- CPLD: Complex Programmable Logic Device
- **With a full custom chip:**

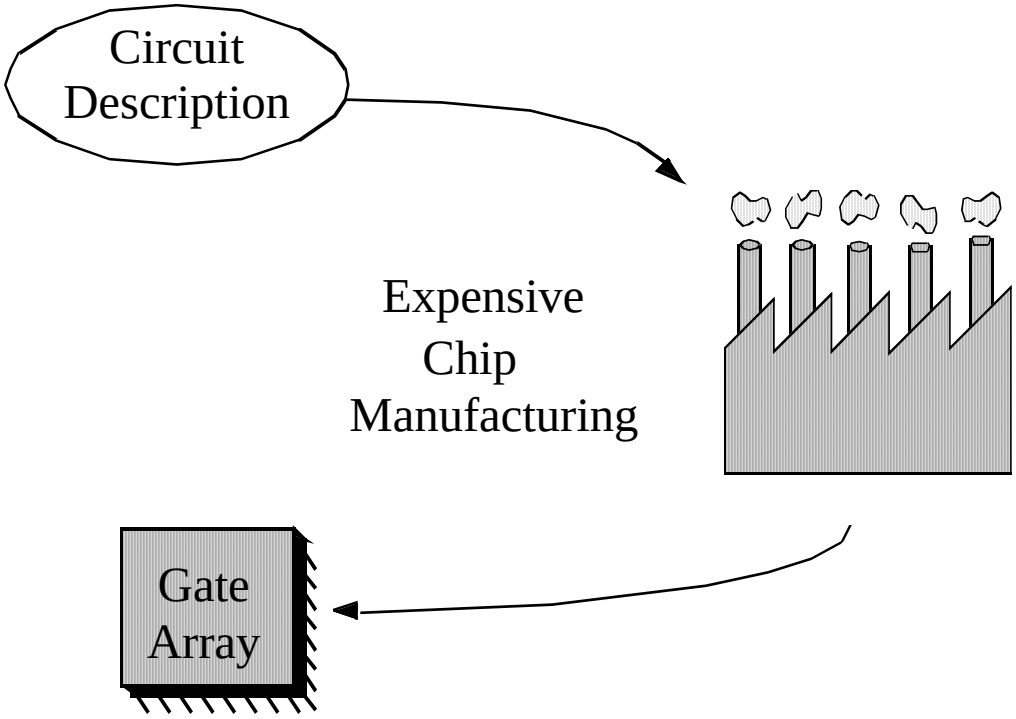

FPGAs and CPLDs

- **Field-Programmable Gate Array (FPGA):** can implement almost any digital circuit *instantly* just by reprogramming the FPGA!

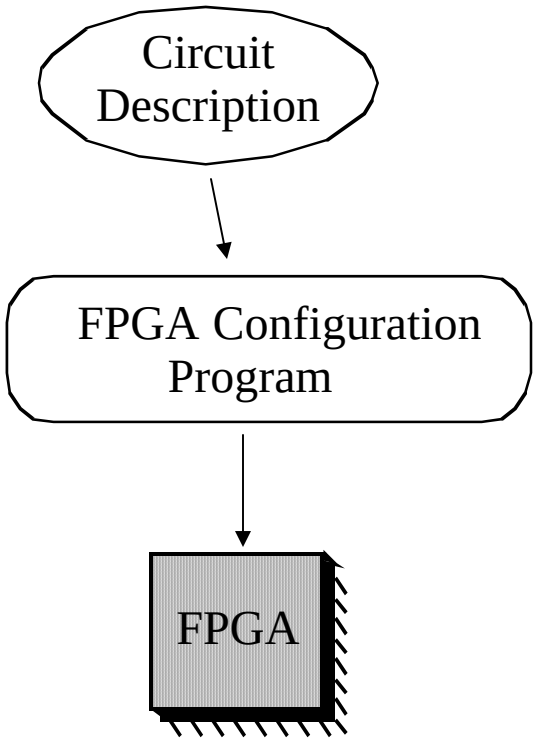

Advantages of FPGAs

1. "Instant Manufacturability": reduces time to market
2. Cheaper for small volumes because you don’t need to pay for fabrication
   - means you don’t need to be a big company to make a chip
3. Relaxes Designers -> relaxed designers live longer!

Disadvantages of FPGAs

1. Slower than gate arrays or custom chips
2. Can not get as much circuitry on a single chip
   - Today: ~ 250M gates is the best you can do
   - ~ 550 MHz is about as fast as you can get
3. For large volumes, it can be more expensive than gate arrays and custom chips

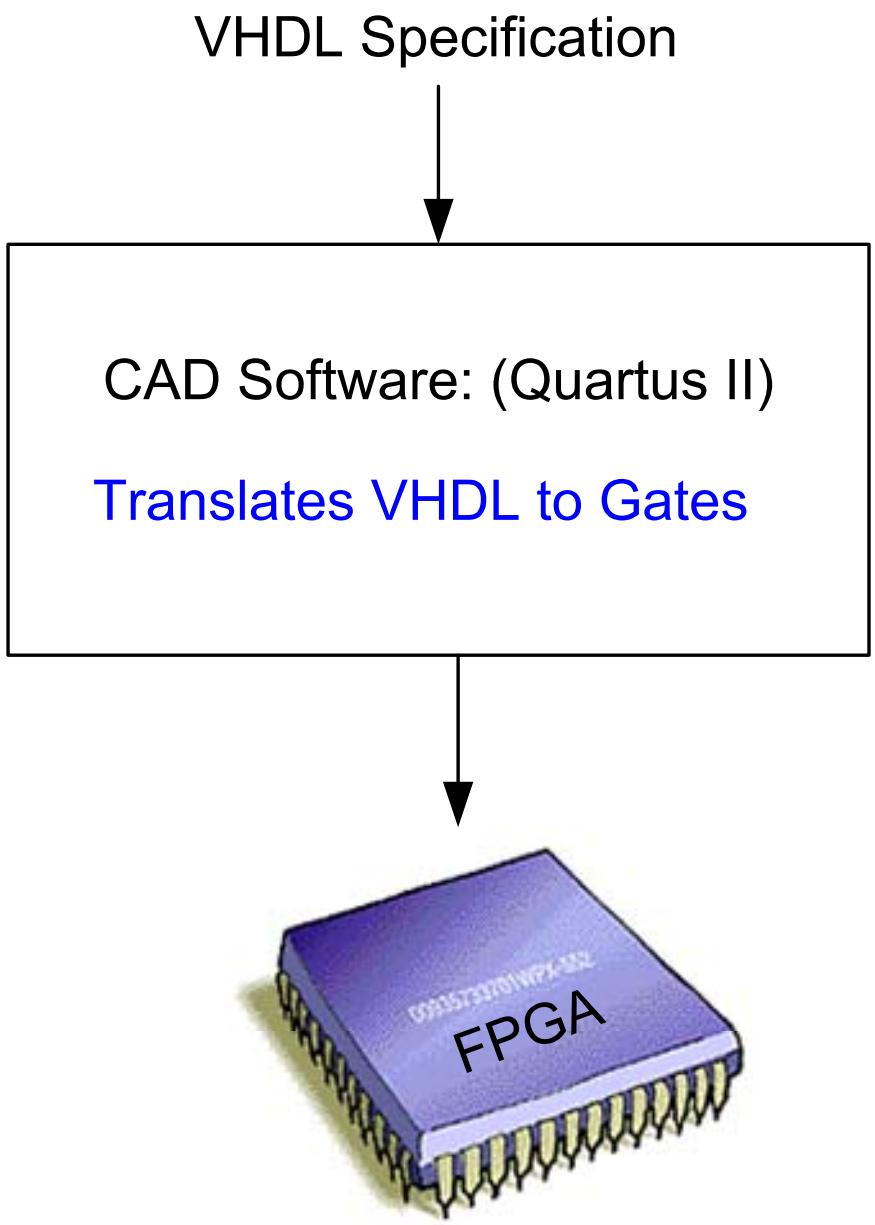
VHDL Specification
CAD Software: (Quartus II)
Translates VHDL to Gates
FPGA

Important thing to remember:
The FPGA does not "*execute*" the VHDL-
it implements ***gates***

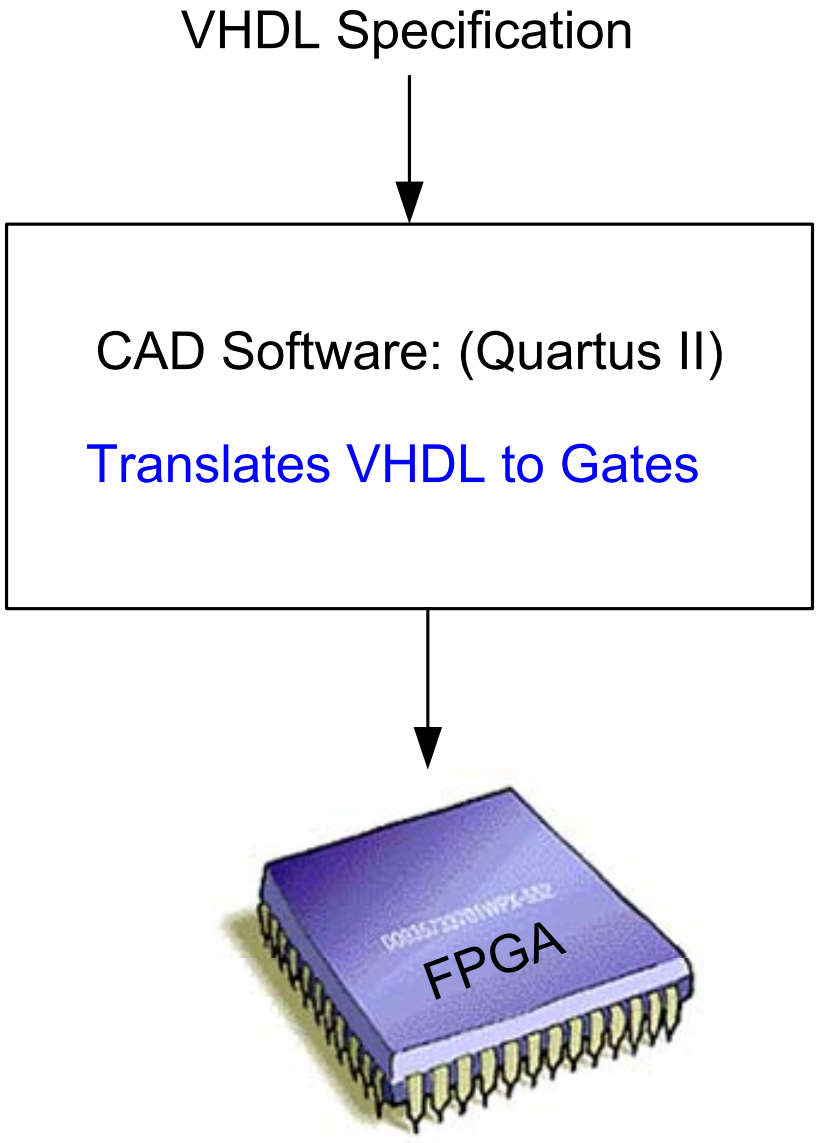

Two types of Simulation:

1. Behaviour Simulation (simulates VHDL directly)
2. Timing Simulation (simulates gates)

Timing simulation is more accurate, but behavioural simulation is faster

In this course, I recommend always using Timing Simulation

What is inside an FPGA?

- Do we care?
  - The tools shield us pretty well from the internals.

- But, it helps to understand what is going on under-the-hood:
  - You can better optimize your design if you understand how it is being implemented (smaller, faster, less power → more $$$)
  - It can be helpful during debugging
  - Important to understand how an FPGA is built when you are selecting an FPGA for a project

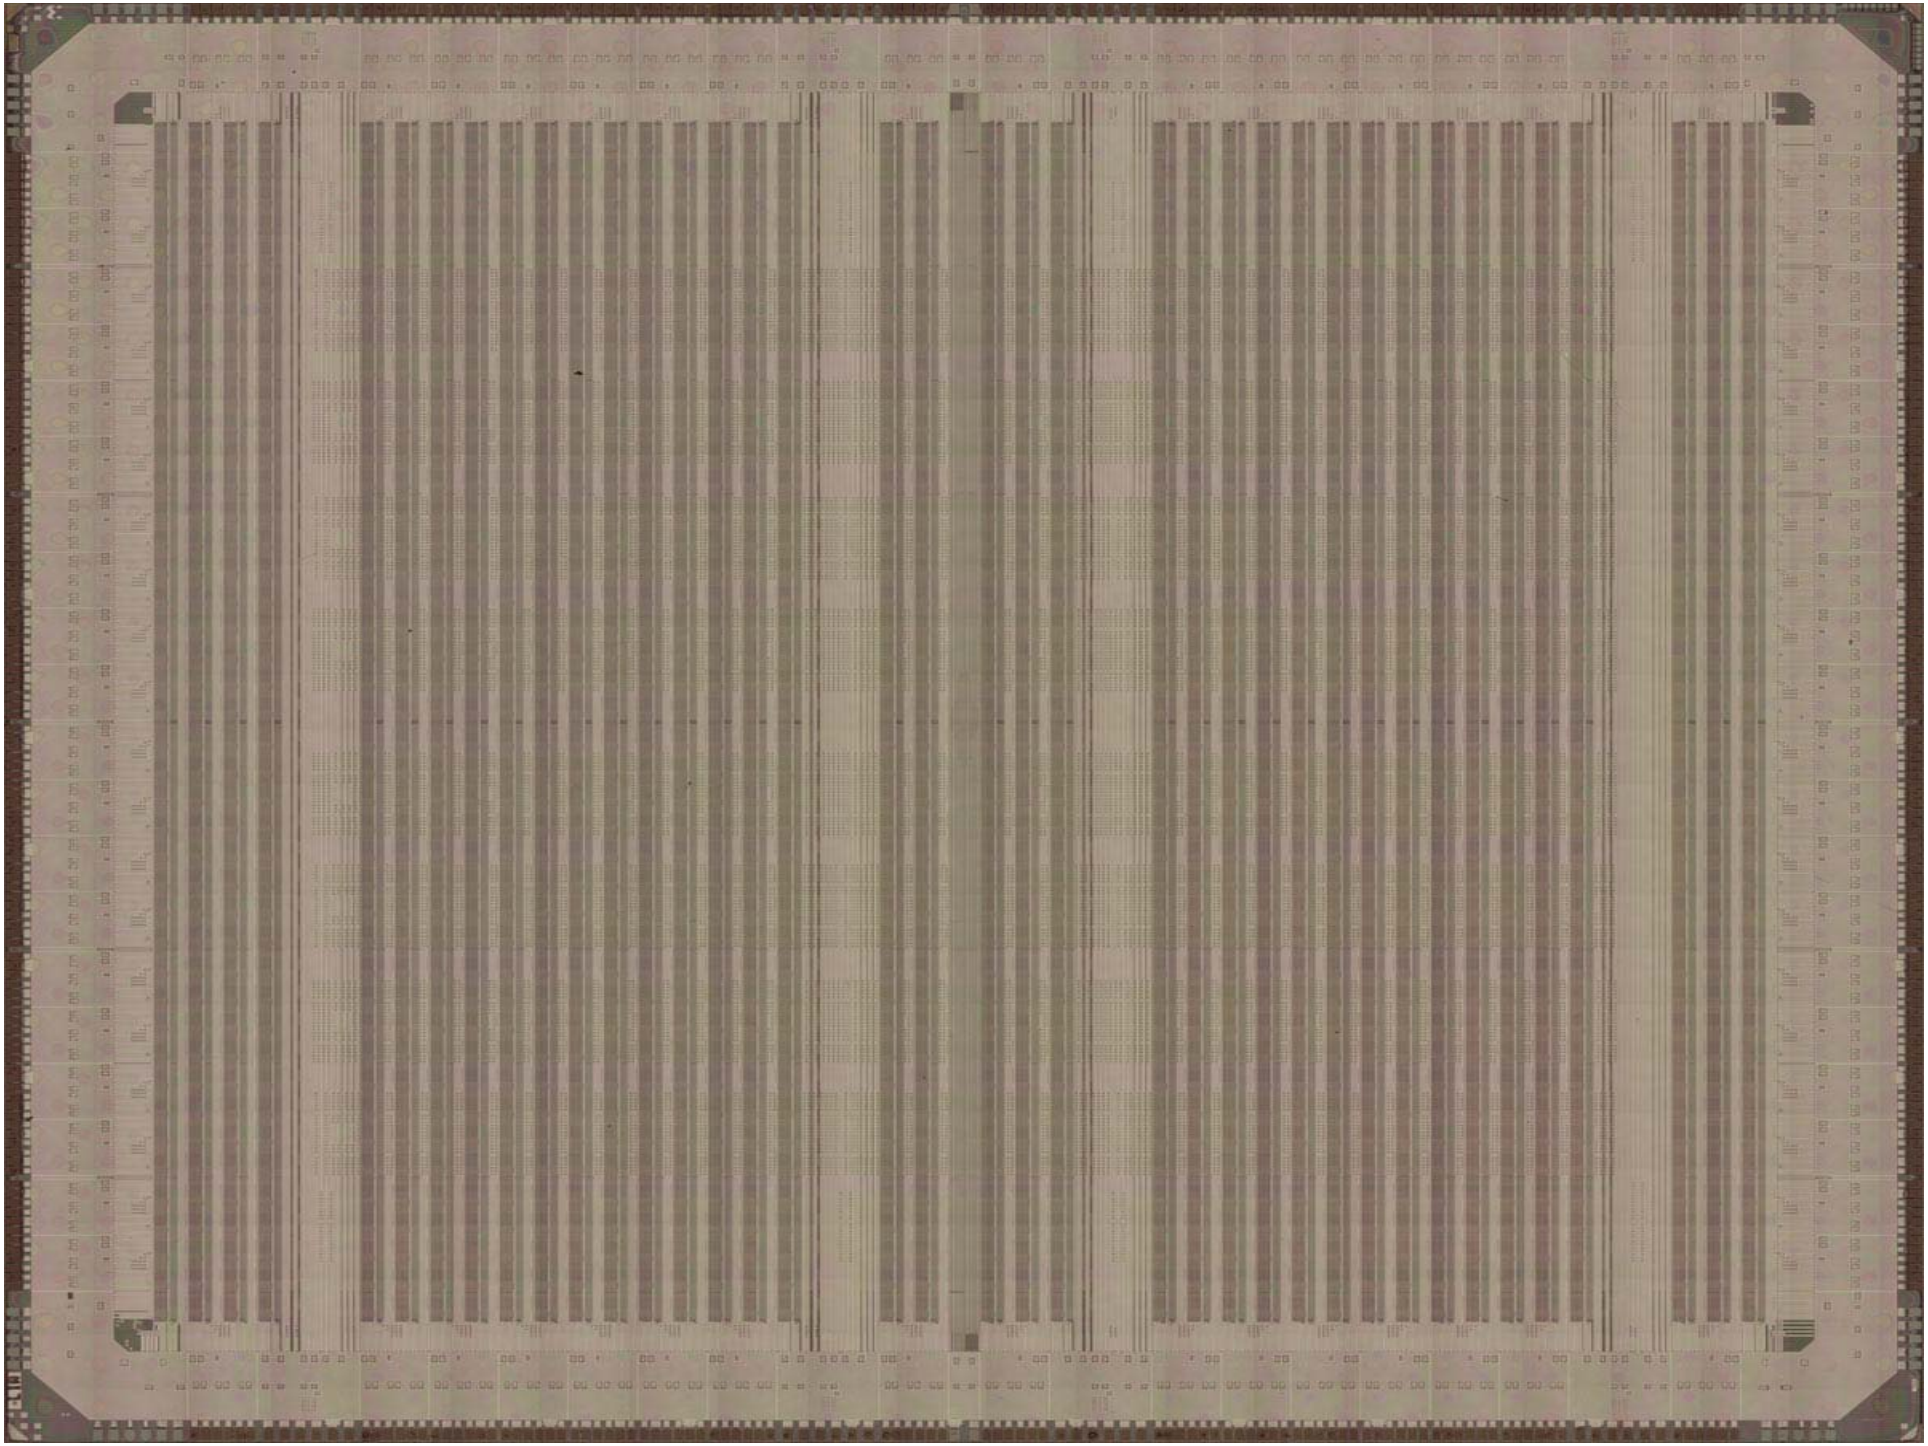

What's Inside an FPGA?

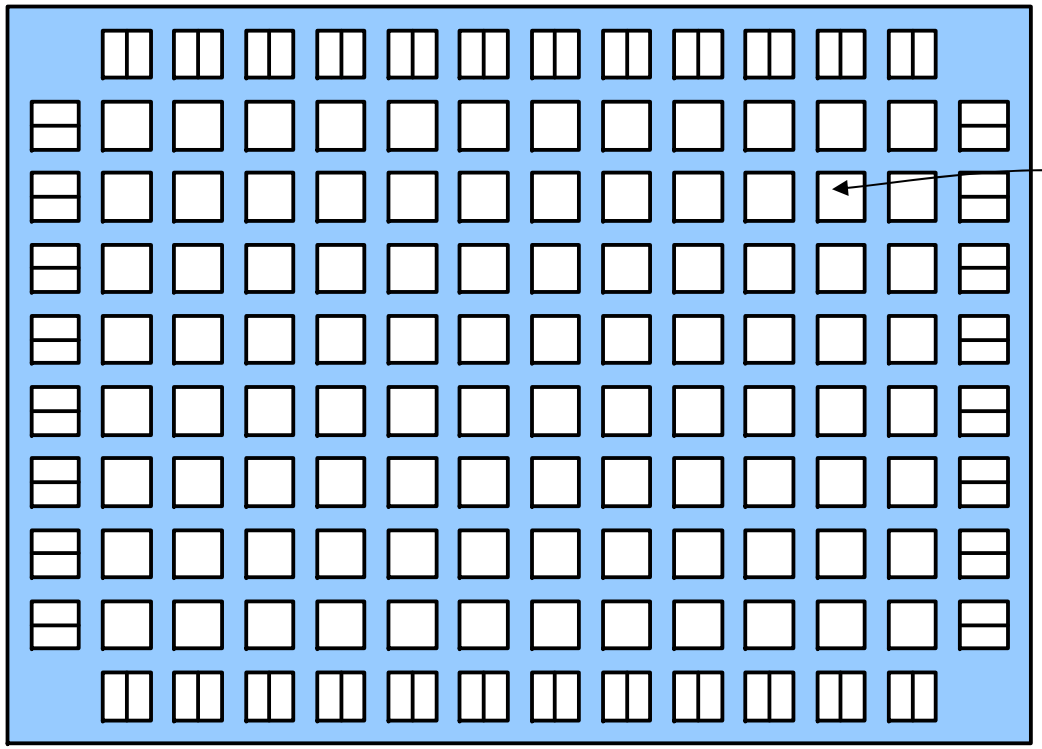

Logic Blocks

- used to implement logic
- lookup tables, LUTs and flip-flops

Altera: LABs

Xilinx: CLBs

What's Inside an FPGA?

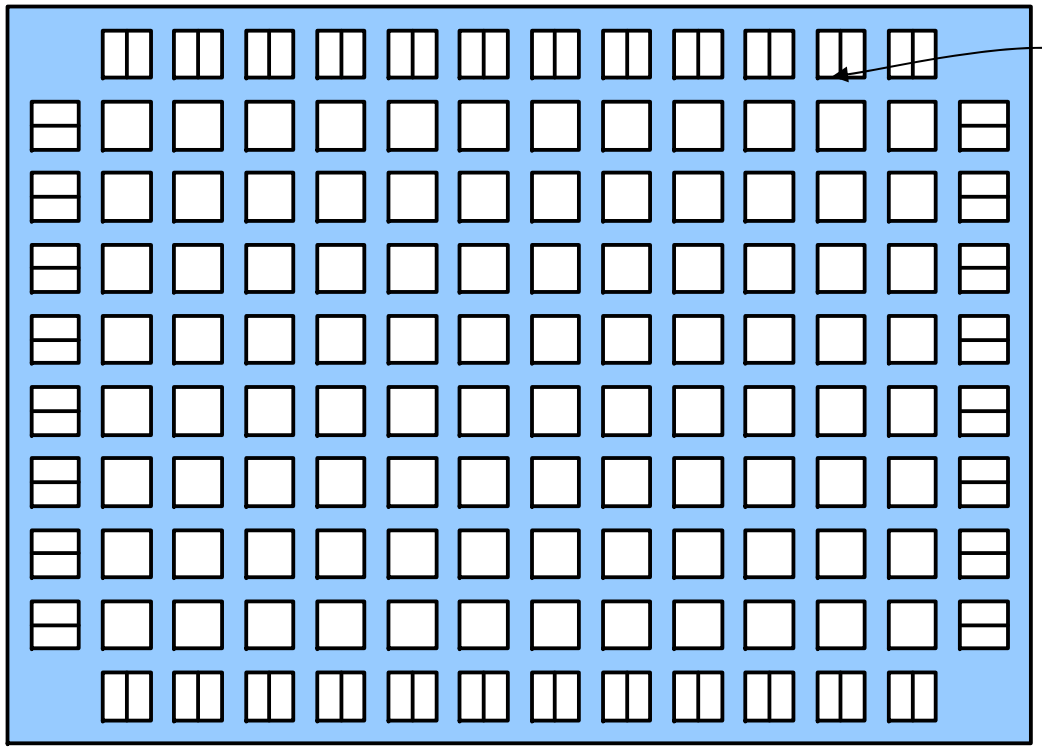

I/O Blocks

- interface off-chip
- can usually support many I/O Standards

(e.g. Virtex 2 Pro:

- 22 single ended standards
- 10 differential standards)

What's Inside an FPGA?

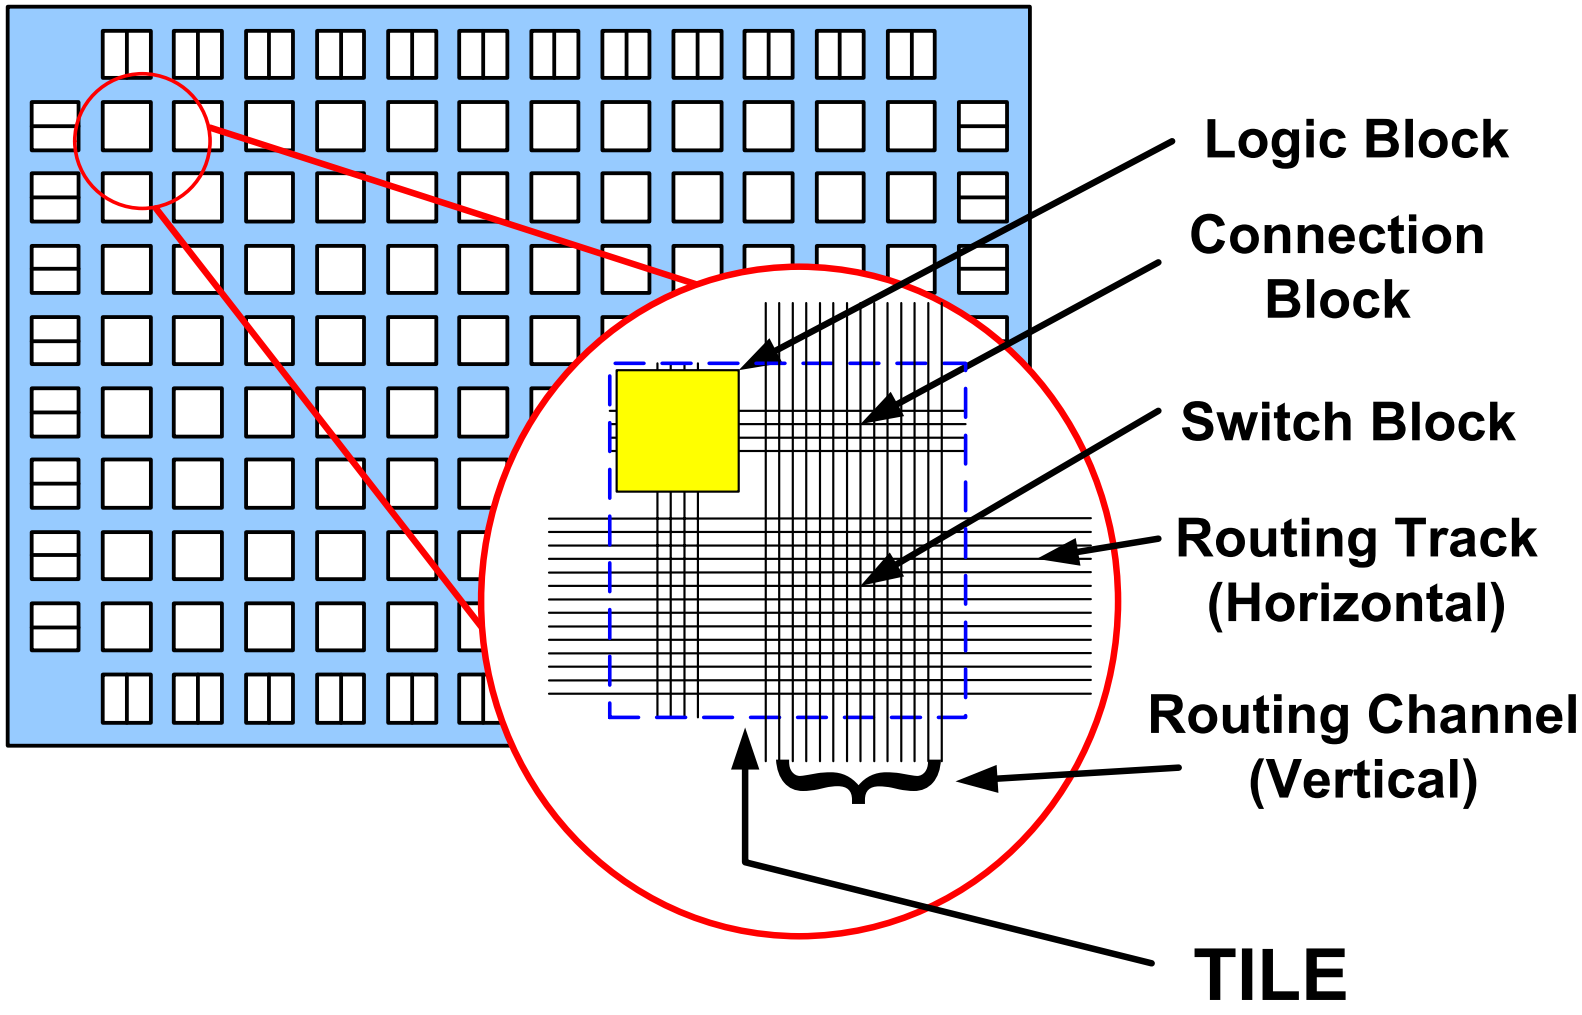

Logic Blocks implement the functionality of the circuit

Logic Block

Basic Logic Gate: Lookup-Table

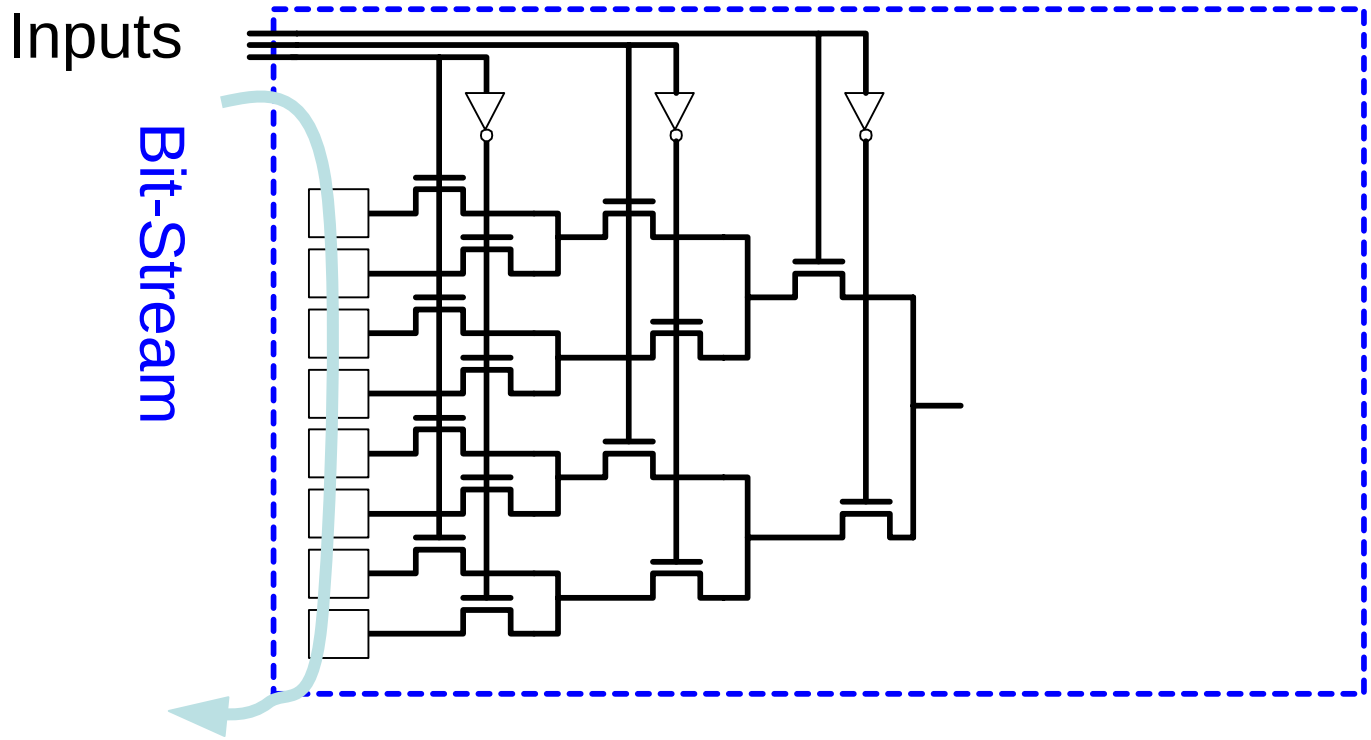

Function of each lookup table can be configured by shifting in bit-stream.

Logic Block

Quick Question: What function would this implement?

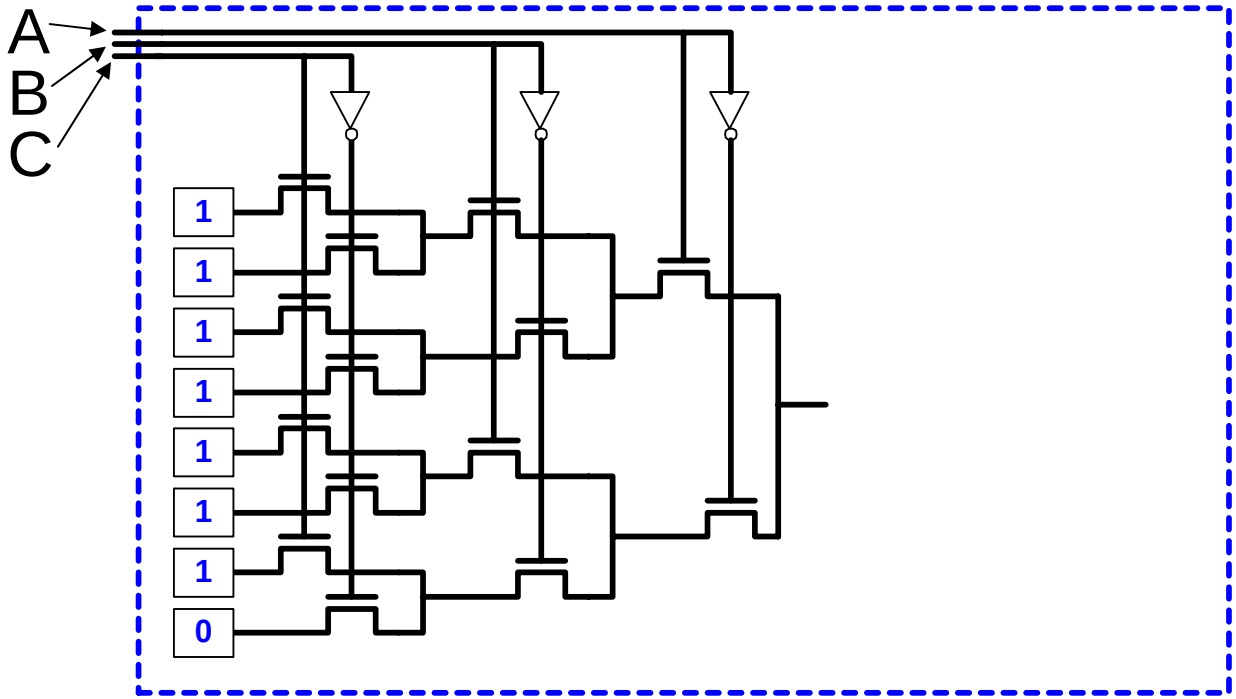

Logic Block

Quick Question: What function would this implement?

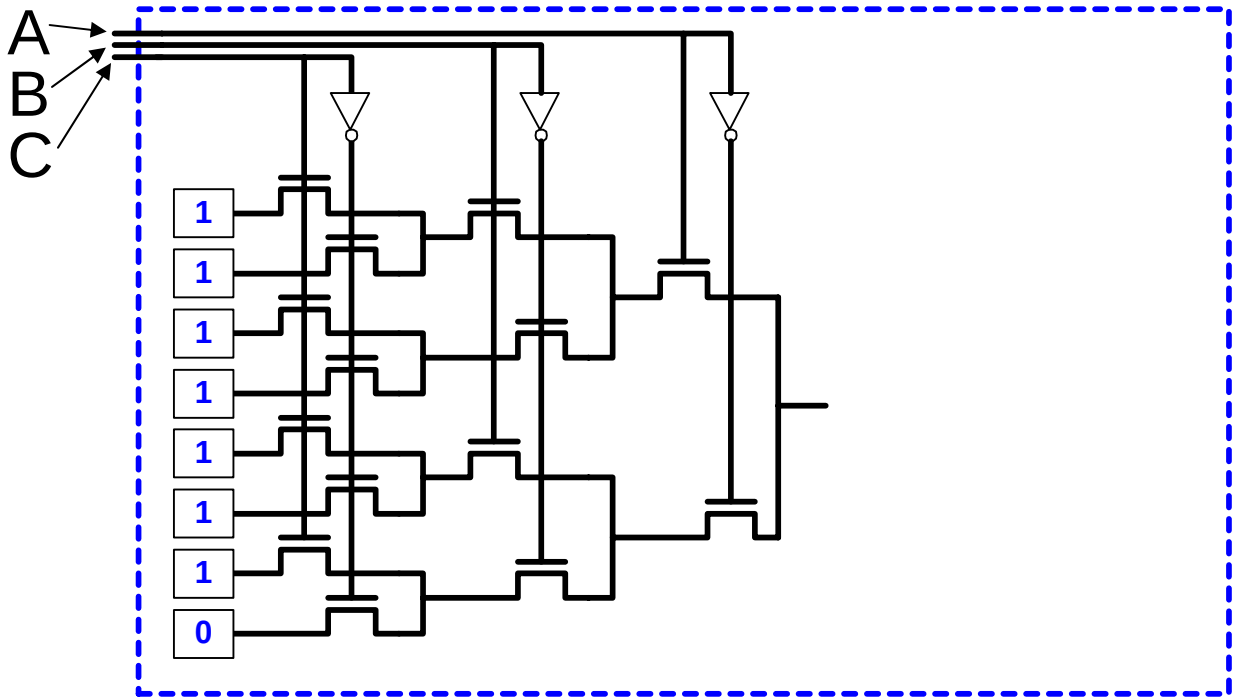

F = A + B + C

Logic Block

Basic Logic Gate: Lookup-Table

Inputs

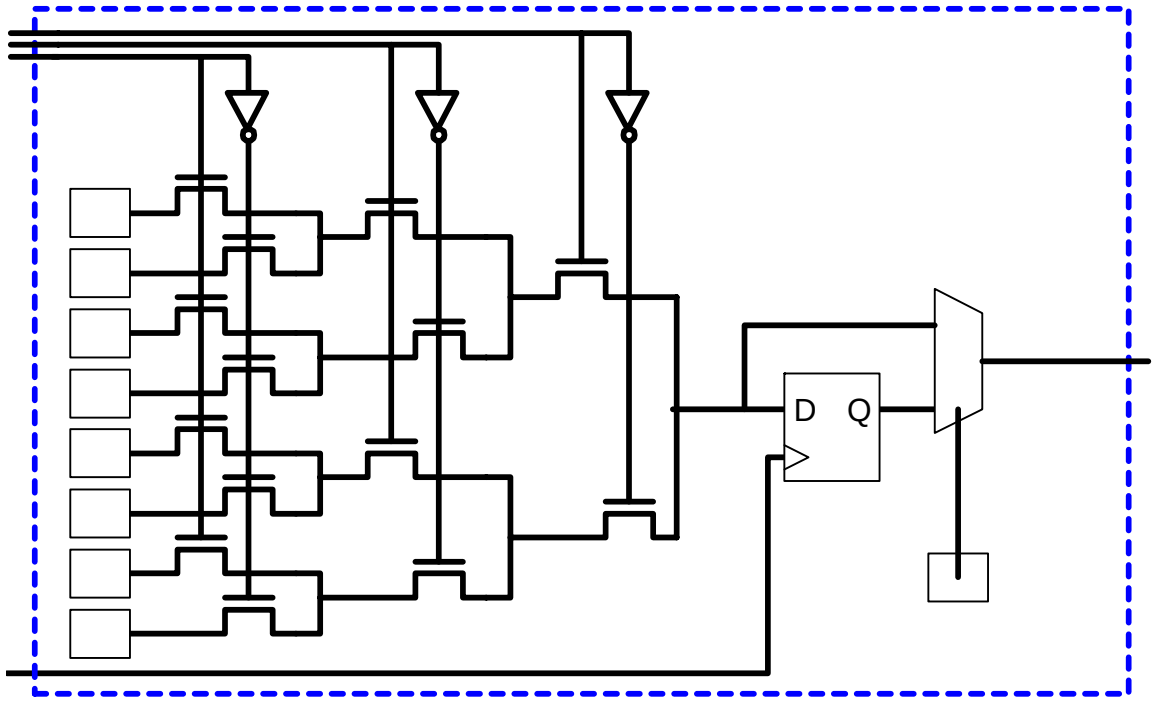

By adding a flipflop and multiplexer we can produce both registered and combinational outputs

Xilinx Virtex II Logic Block

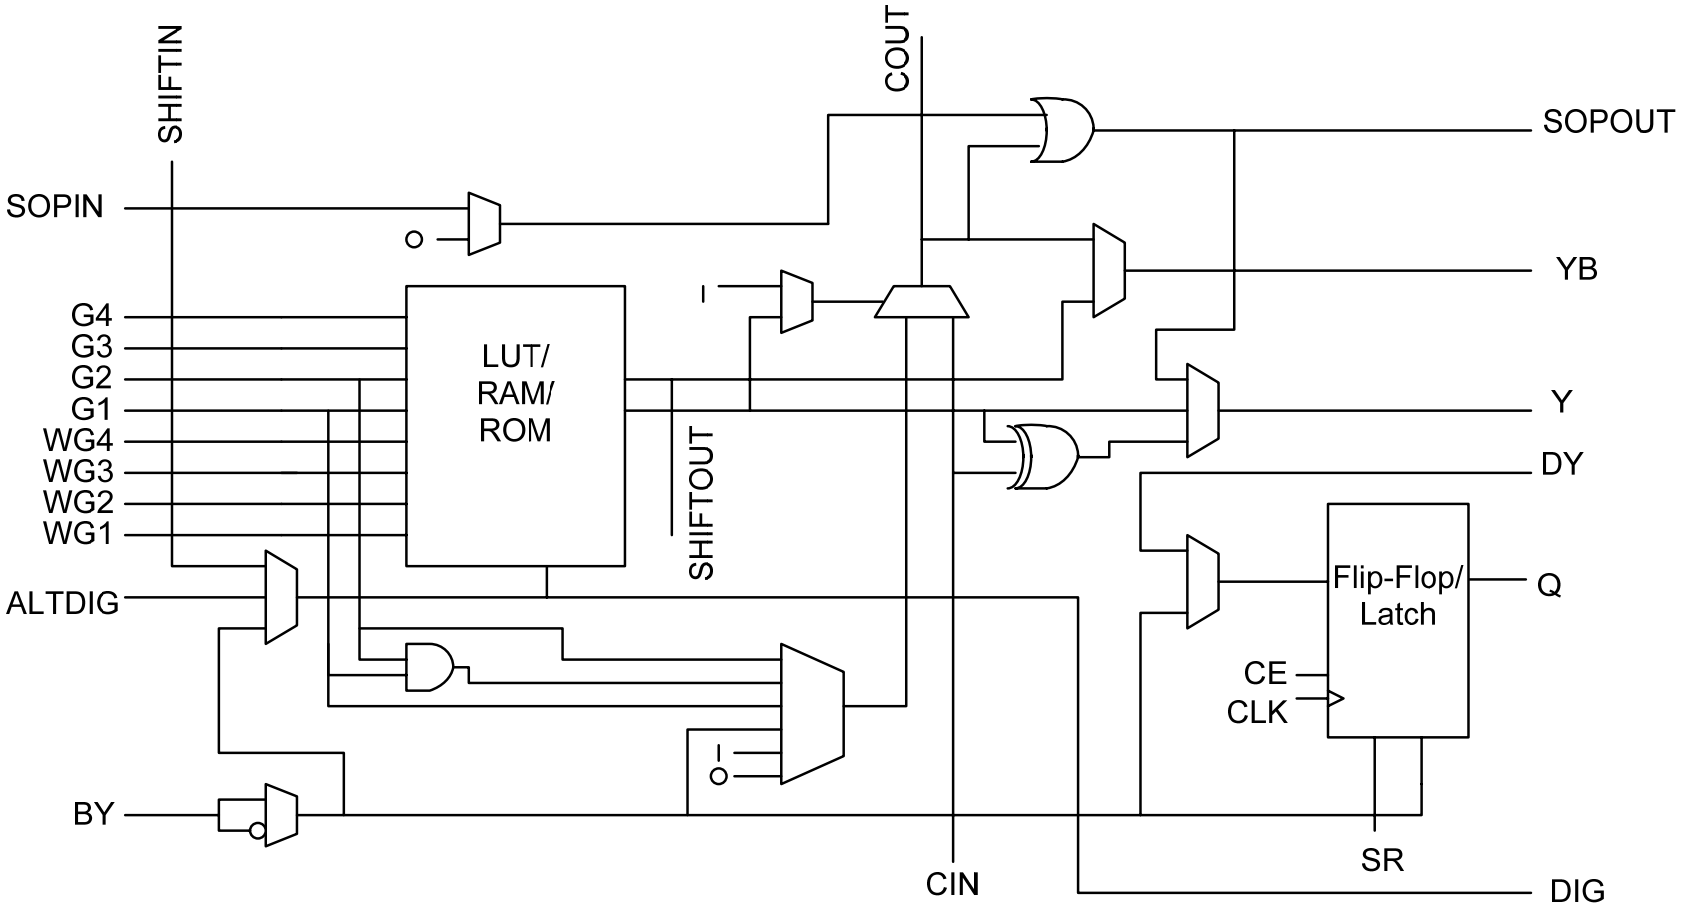

Xilinx Virtex II Logic Block

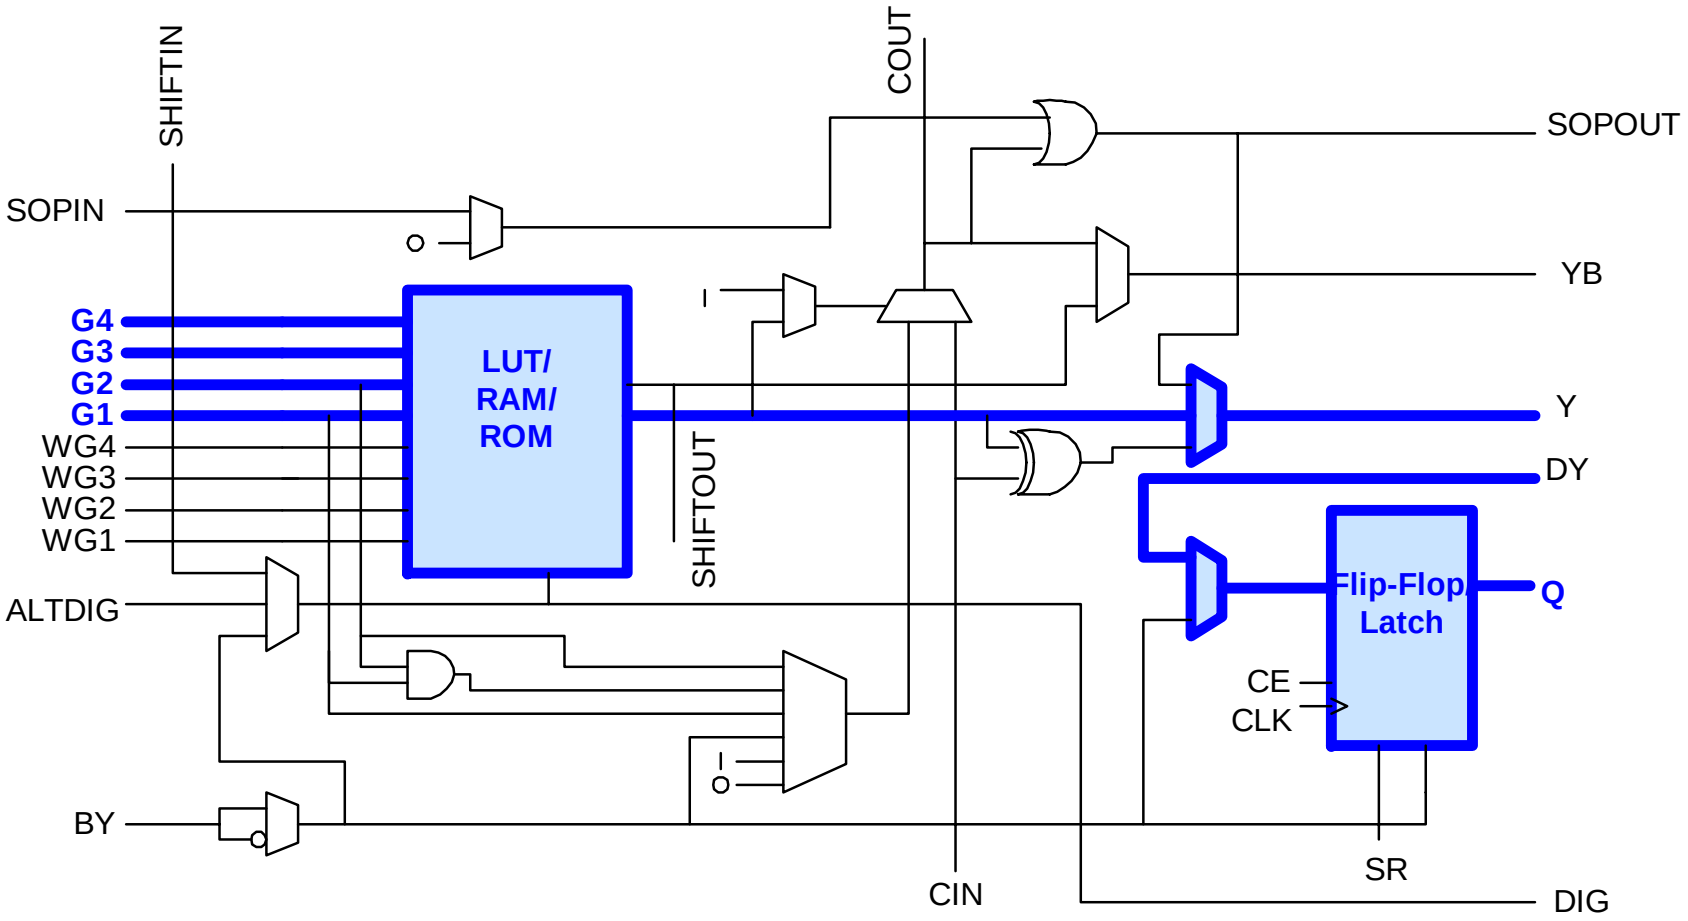

Stratix II Logic Block

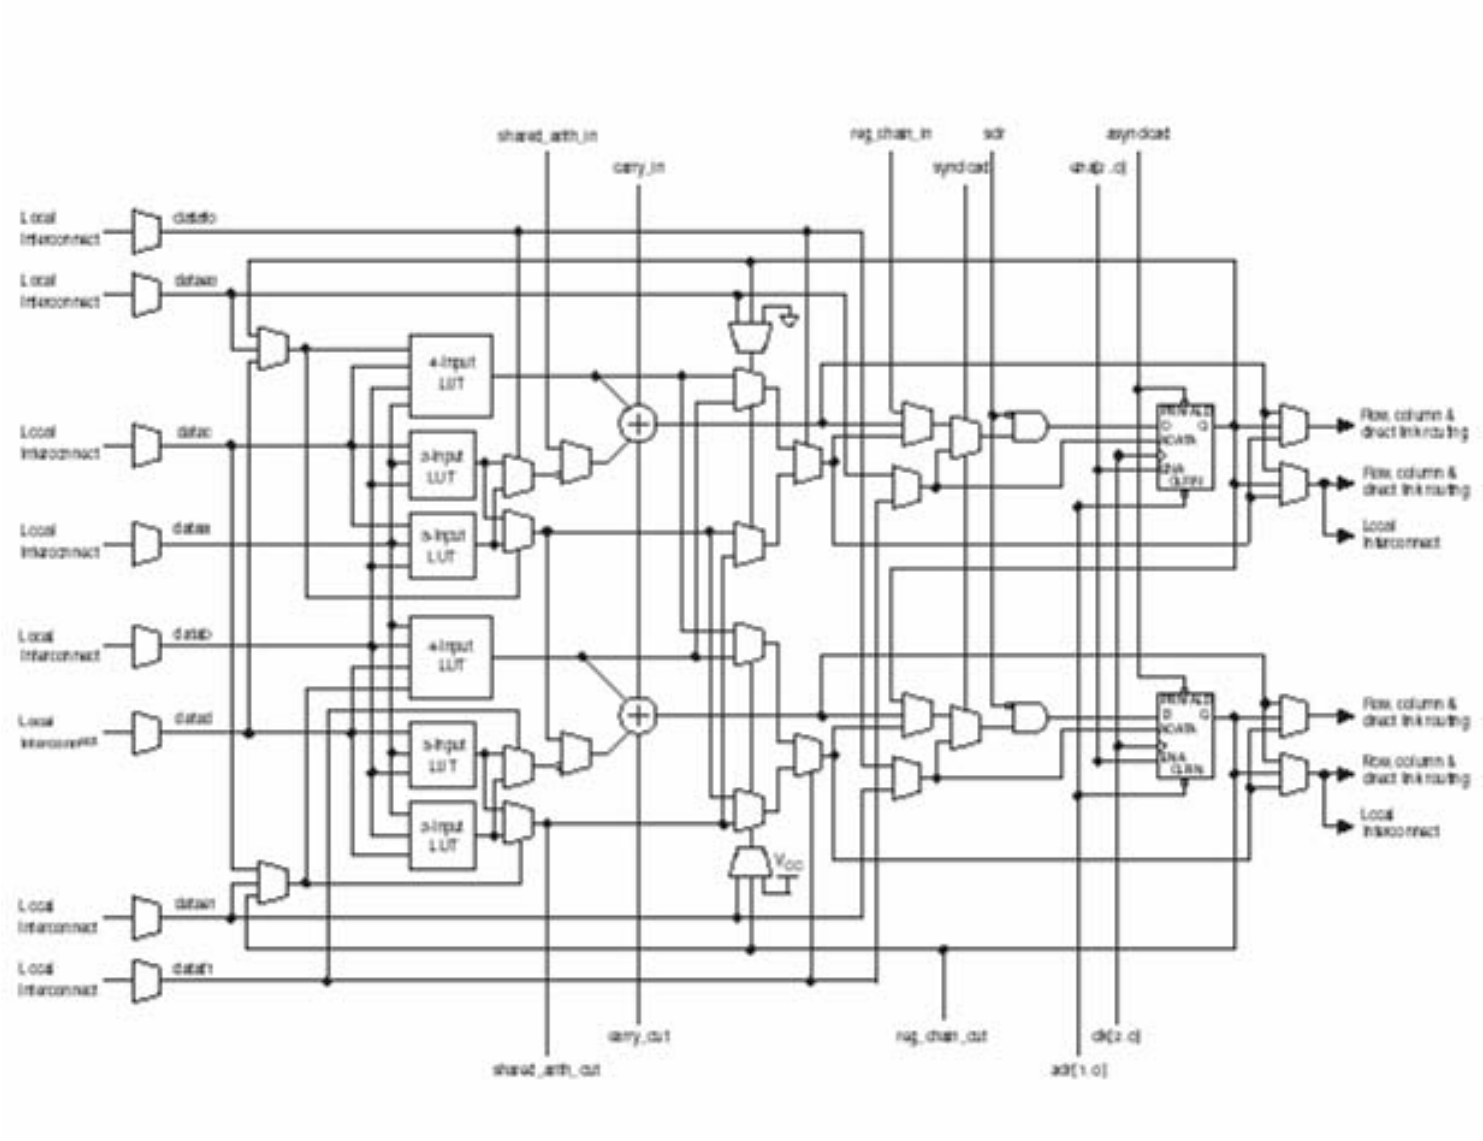

Source: Stratix II Handbook, 2005

Stratix II Logic Block

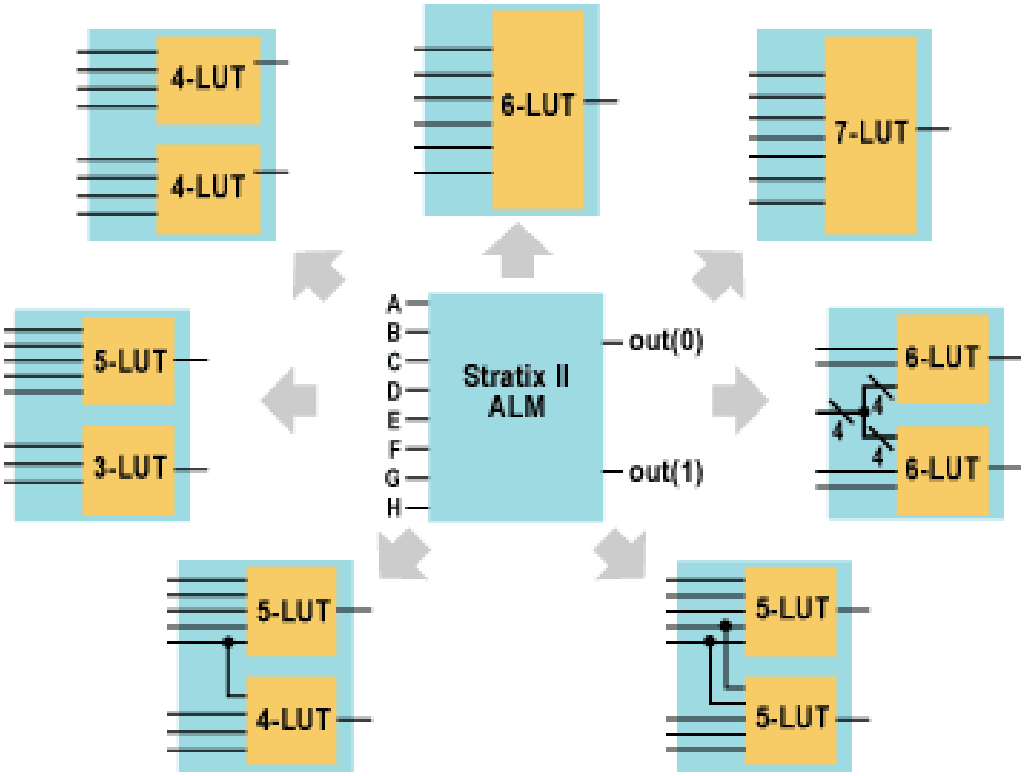

Source: Stratix II Handbook, 2005

Logic Blocks are grouped into Clusters

Logic Clusters

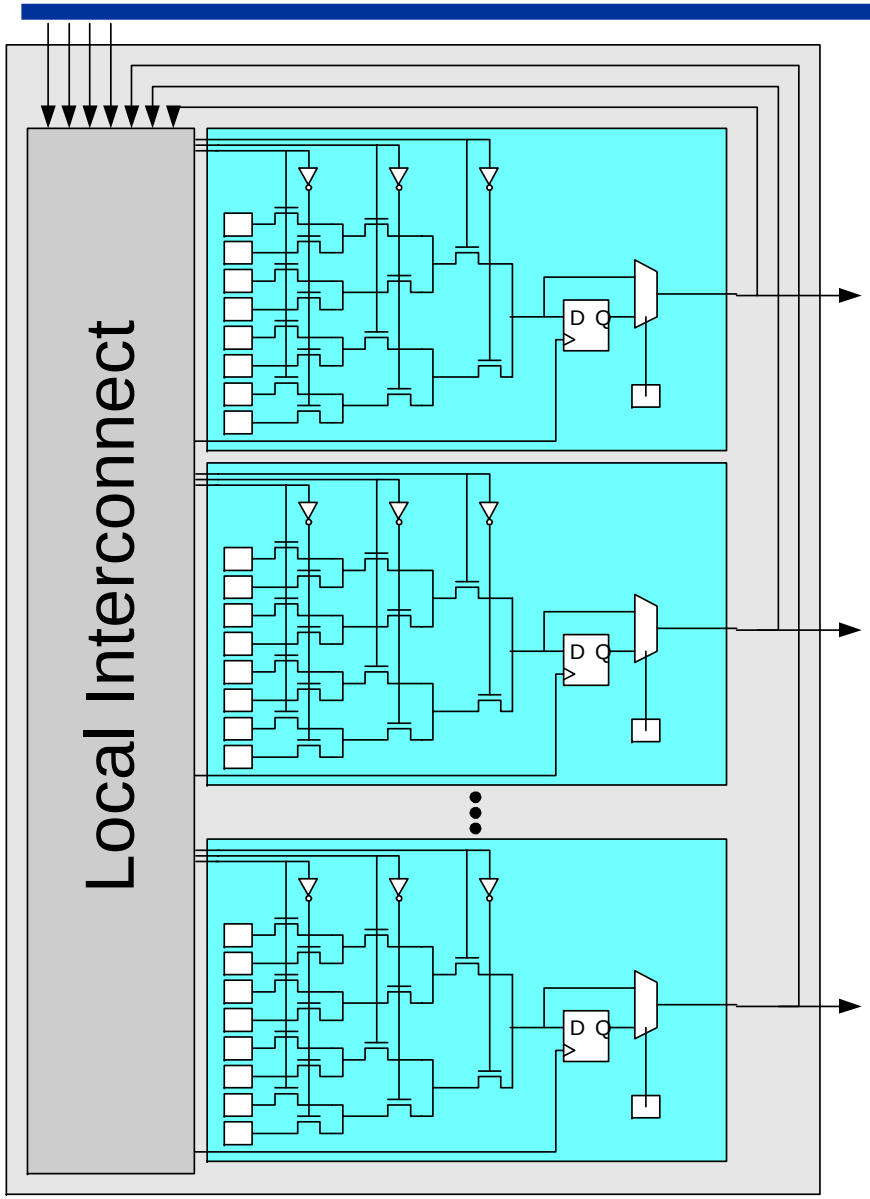

Intra-cluster connections: fast

Inter-cluster connections: slow

There is a balance:

- Larger clusters mean more intra-cluster connections
- But, larger clusters means the intra-cluster connections are not as fast
- Typically 8 to 10 LUTs per cluster

Cluster Architecture

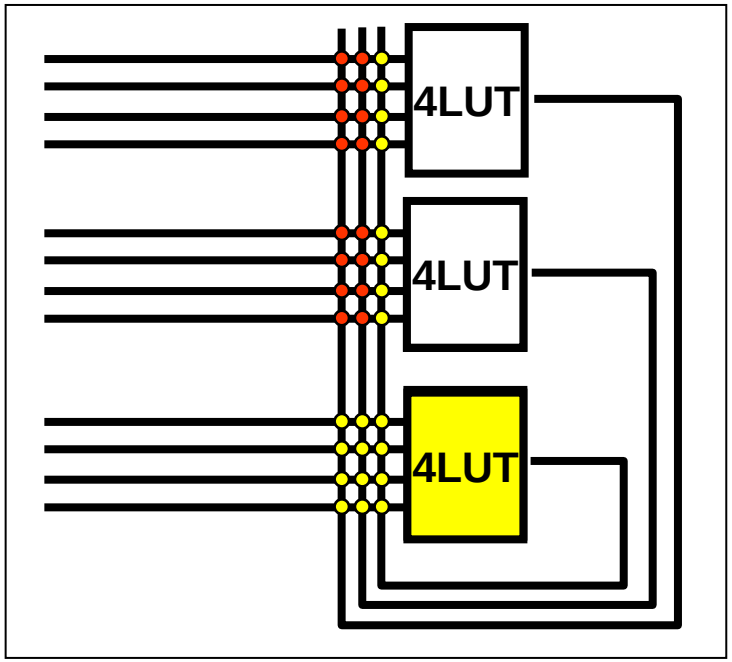

Cluster Architecture

This will significantly impact the speed and routability of the device

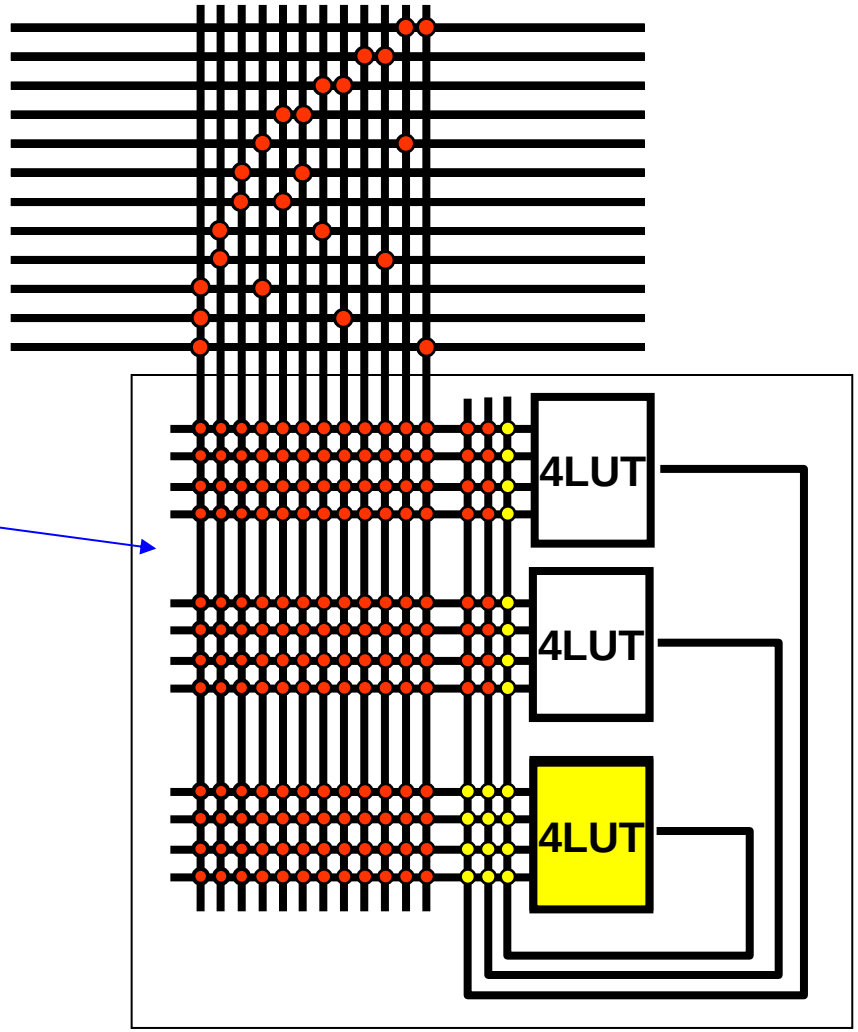

Intra-cluster routing

- Academic studies typically consider fully populated:

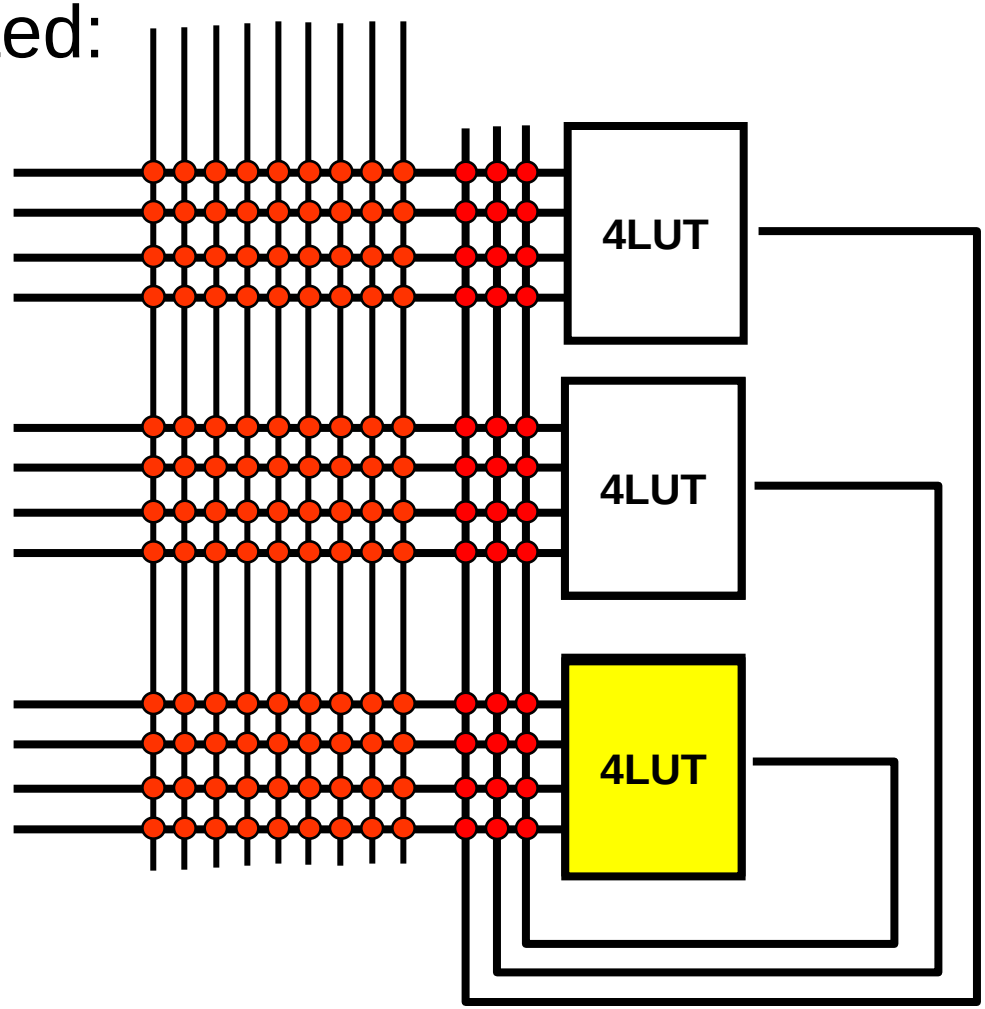

Intra-cluster routing

- Commercial parts: depopulated (this is 50%)

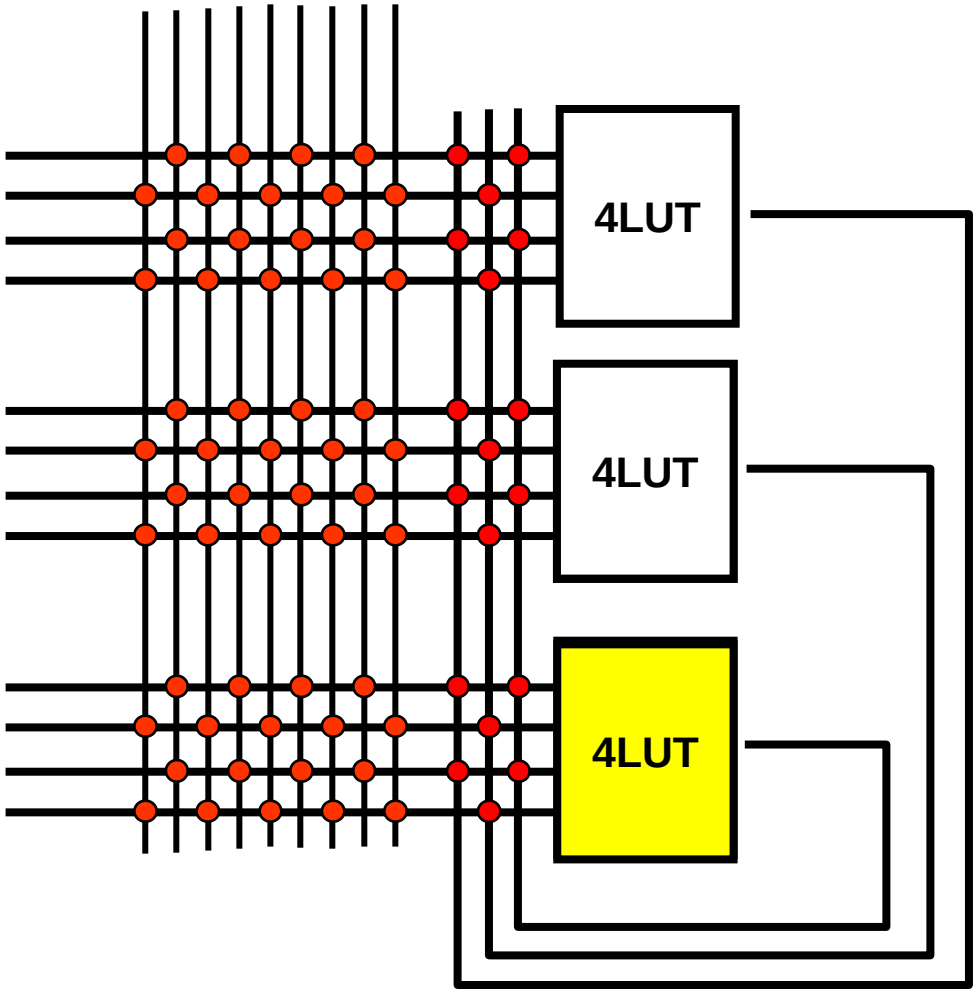

Altera Stratix LAB (Logic Array Block)

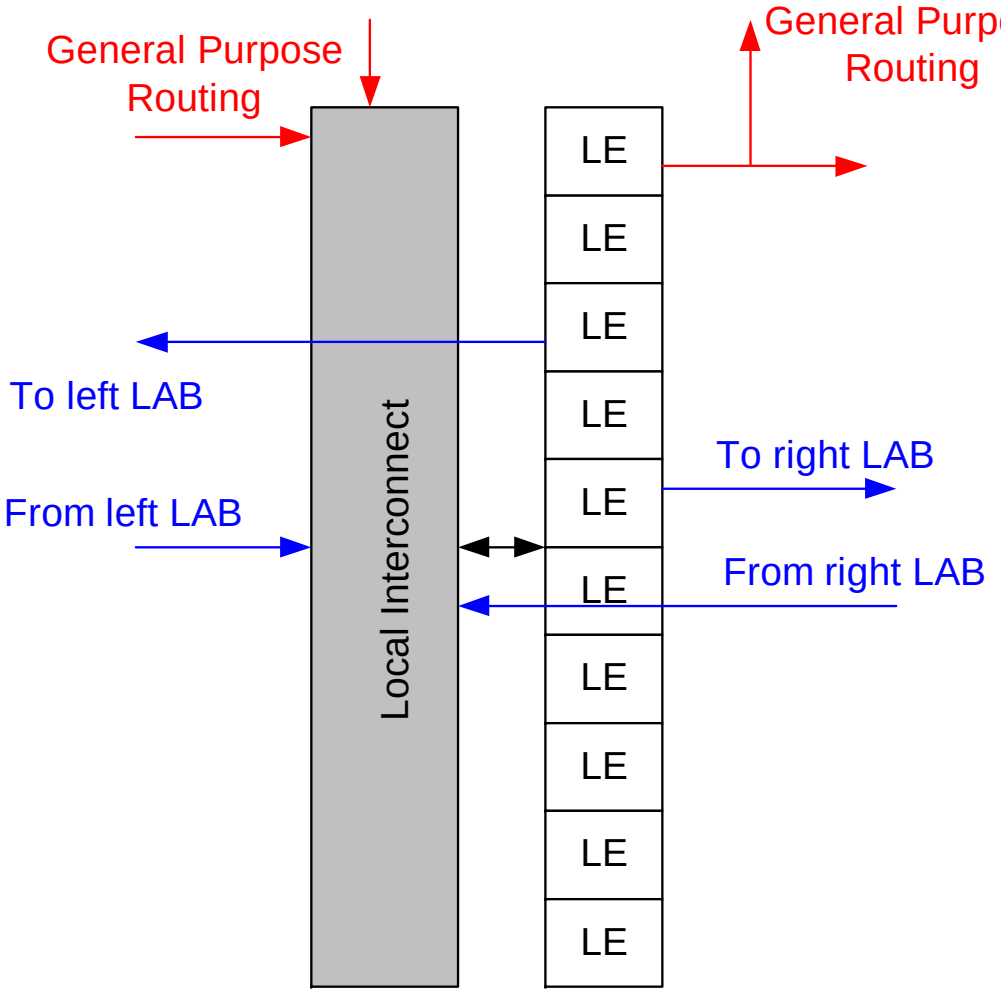

- 10 Logic Elements in each LAB
- Two carry chains through each LAB
- Connections to general purpose routing and neighbouring LABs

Altera APEX MegaLAB

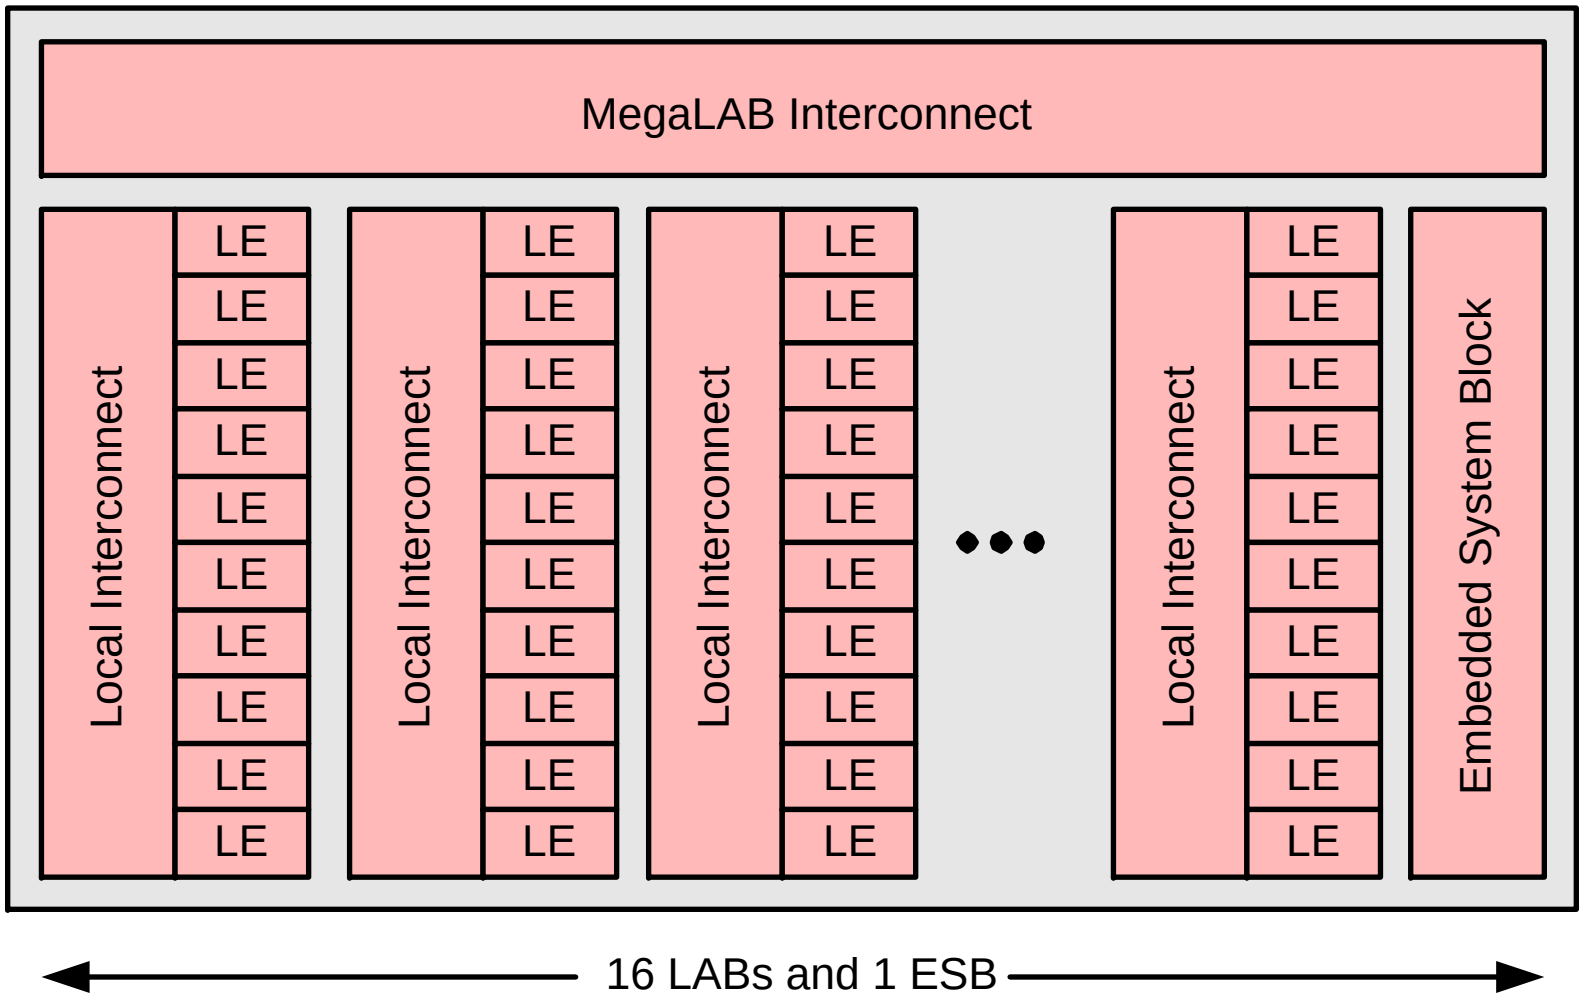

Routing Fabric

Recall: What's Inside an FPGA?

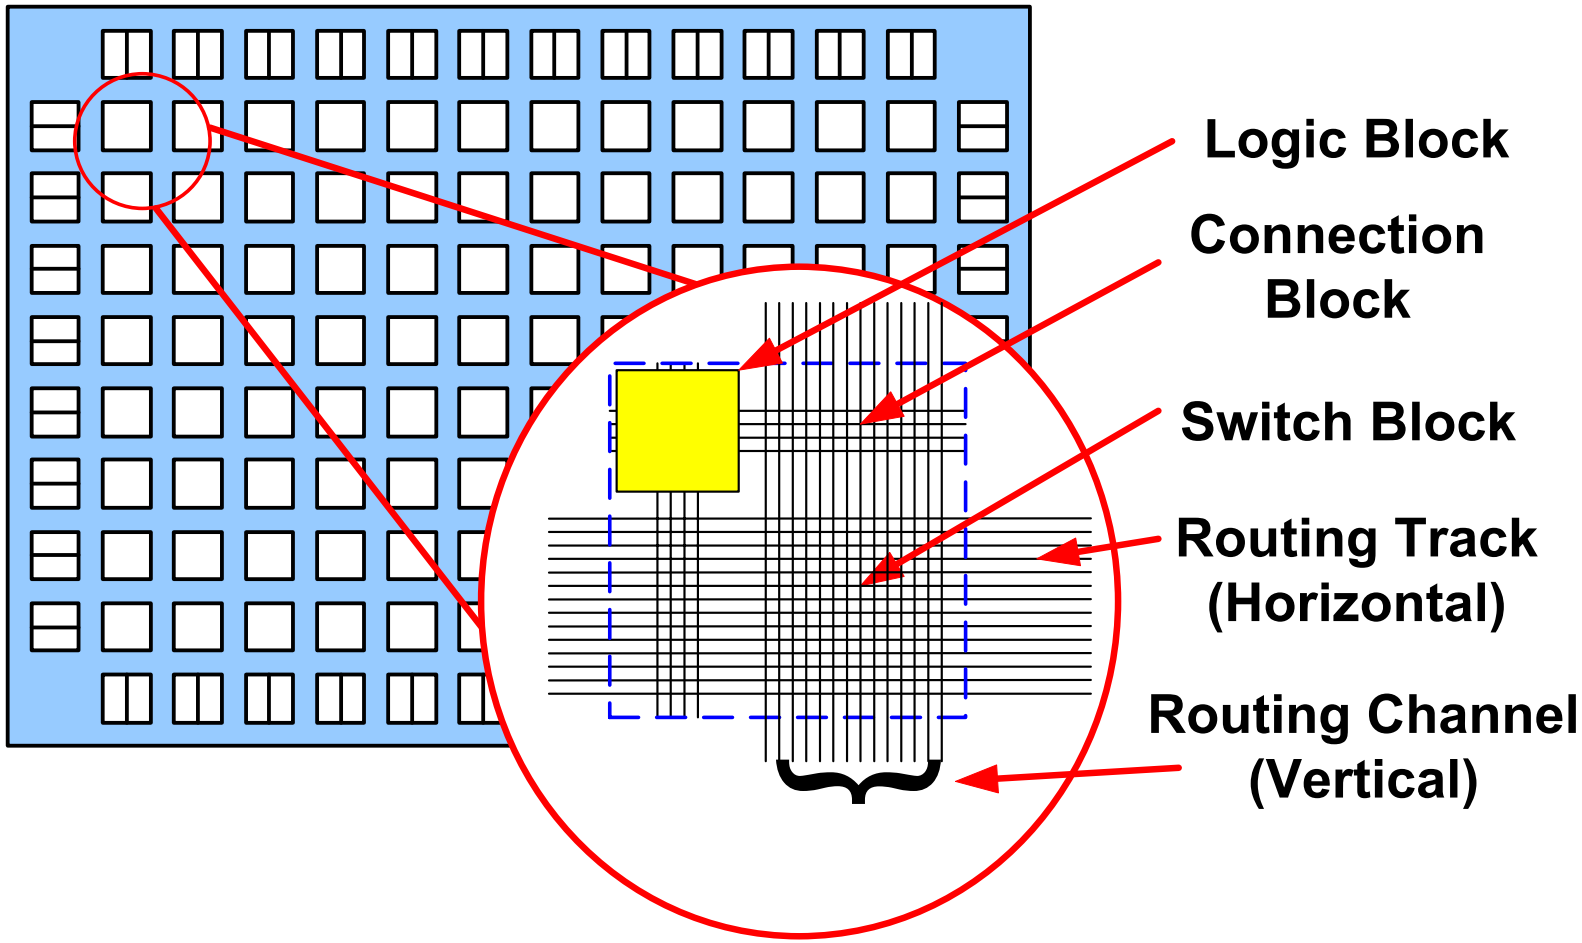

Let's look at the Routing and Switch Blocks

Routing is important!

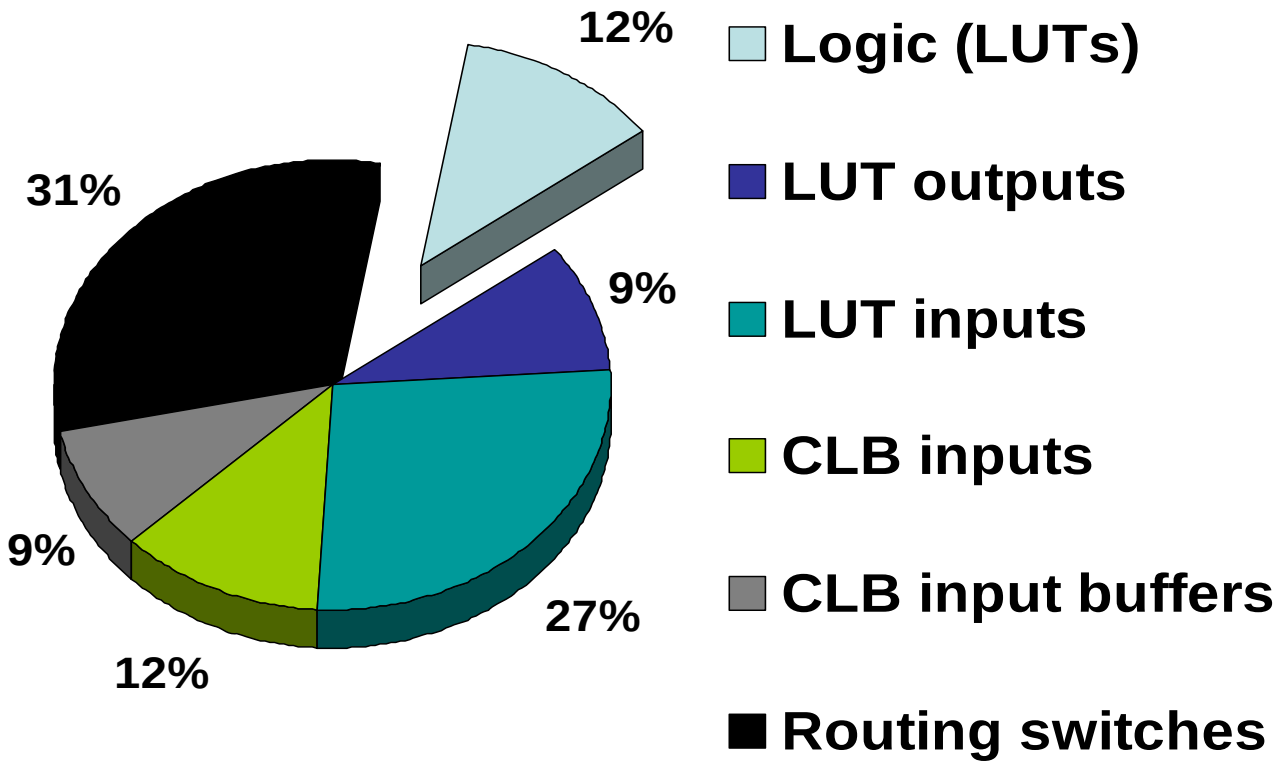

- Source: Dr. Guy Lemieux, UBC

Mesh (Island-style) FPGA

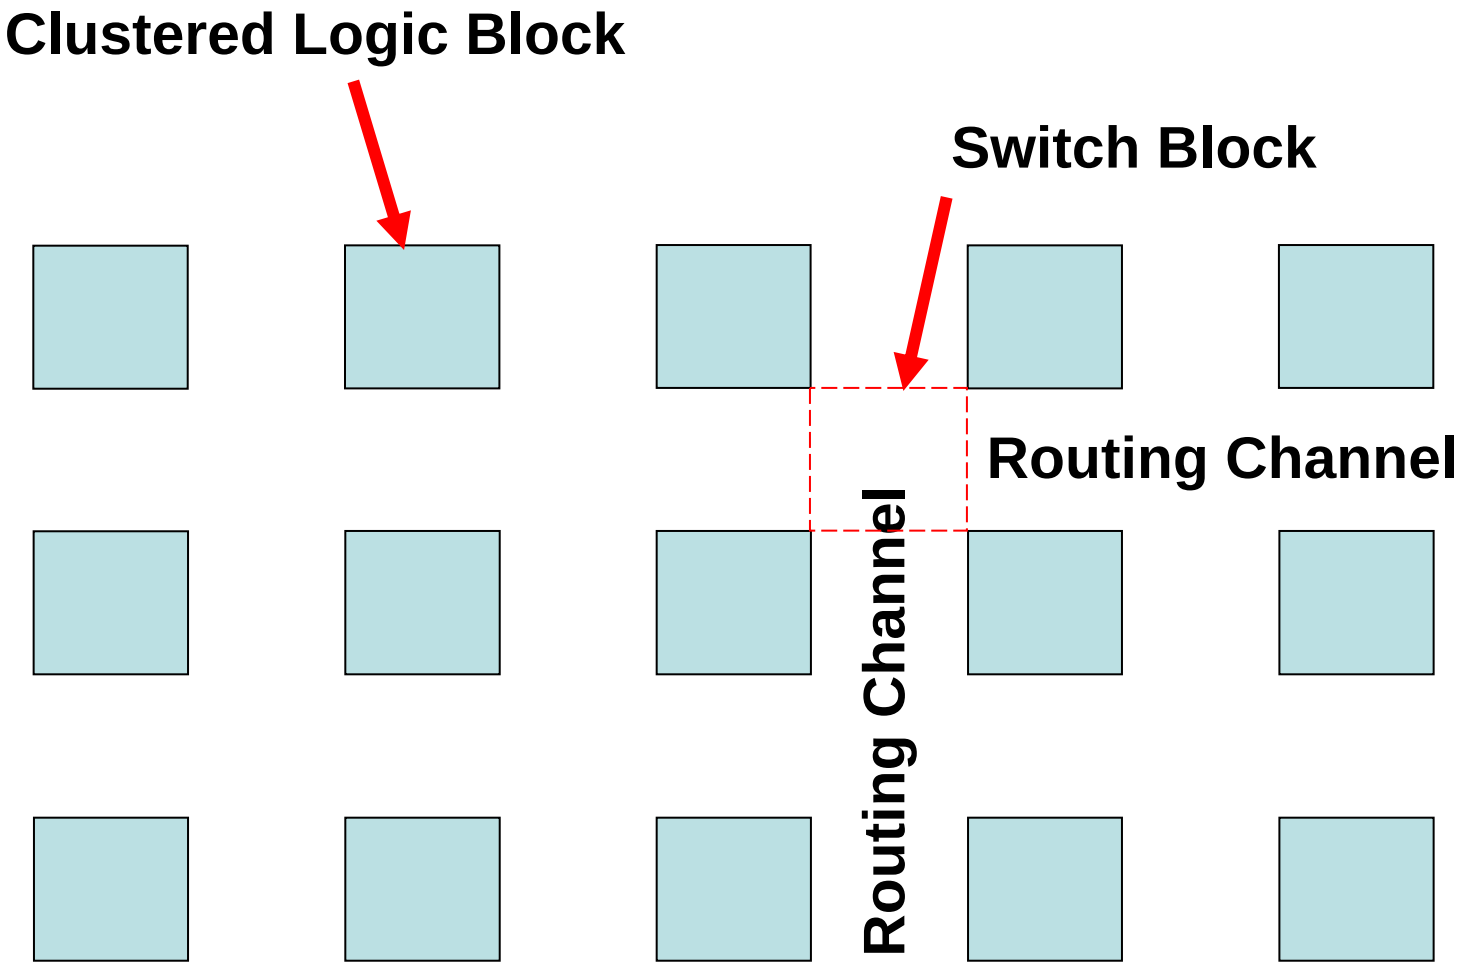

Reconfigurable Logic

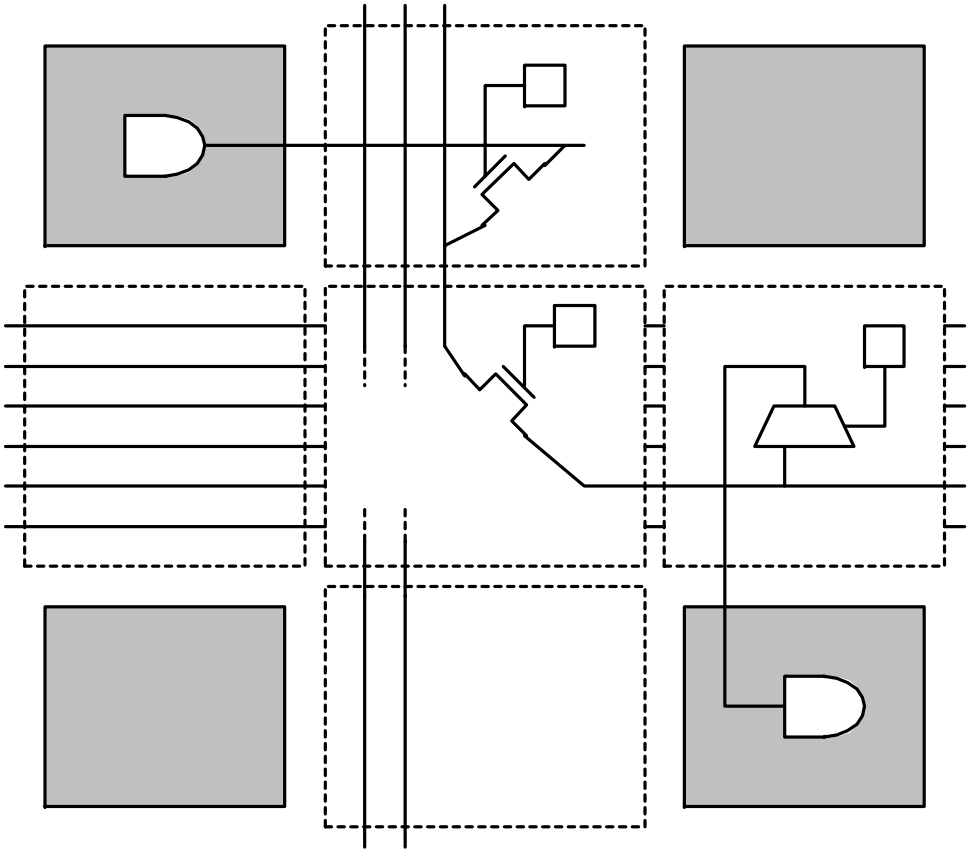

Connect Logic Blocks using Fixed Metal Tracks and Programmable Switches

Reconfigurable Logic

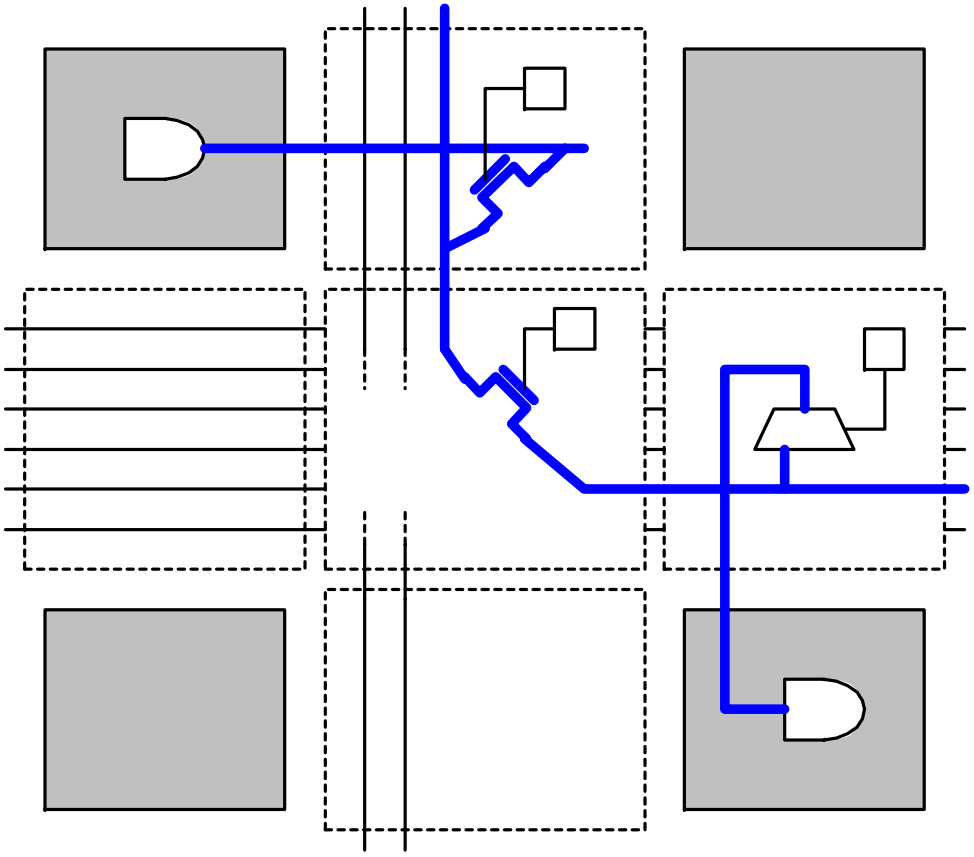

Connect Logic Blocks using Fixed Metal Tracks and Programmable Switches

Programmable Switches

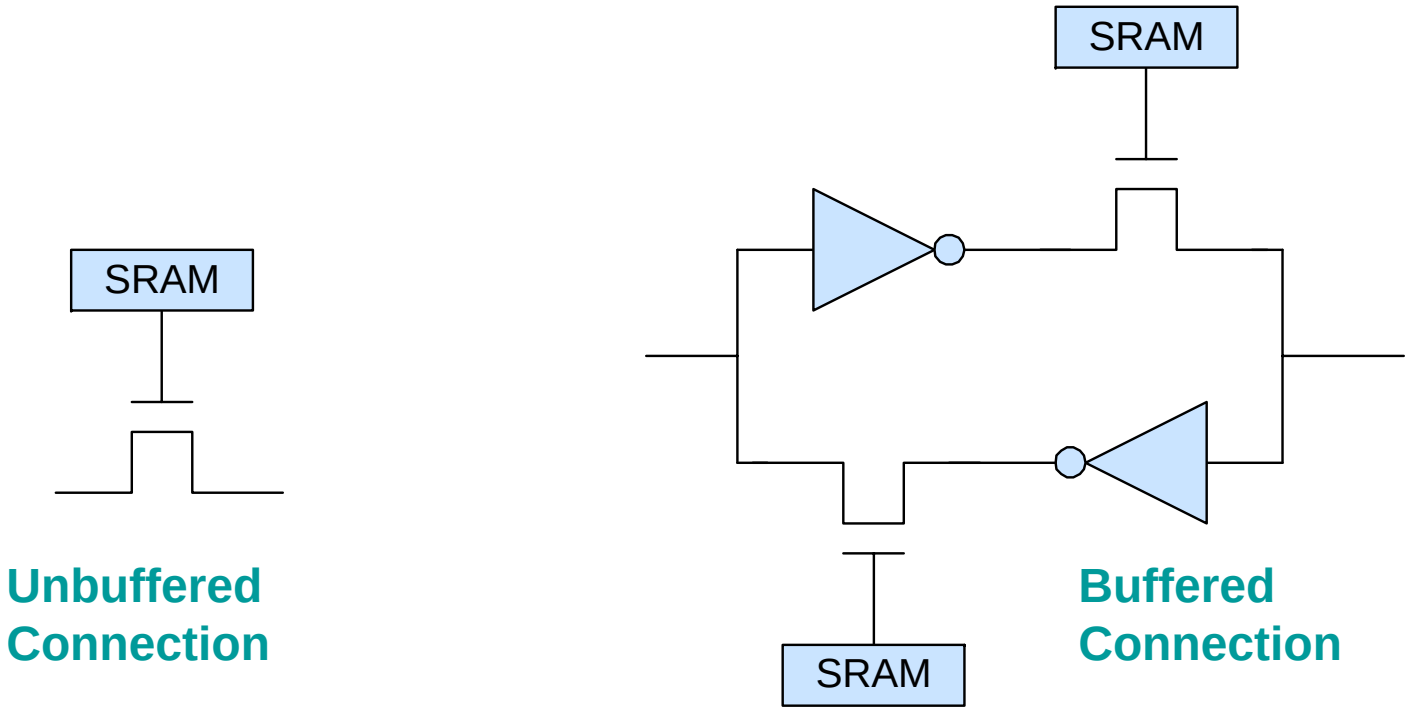

- Today, buffered connections are common

Switch Blocks

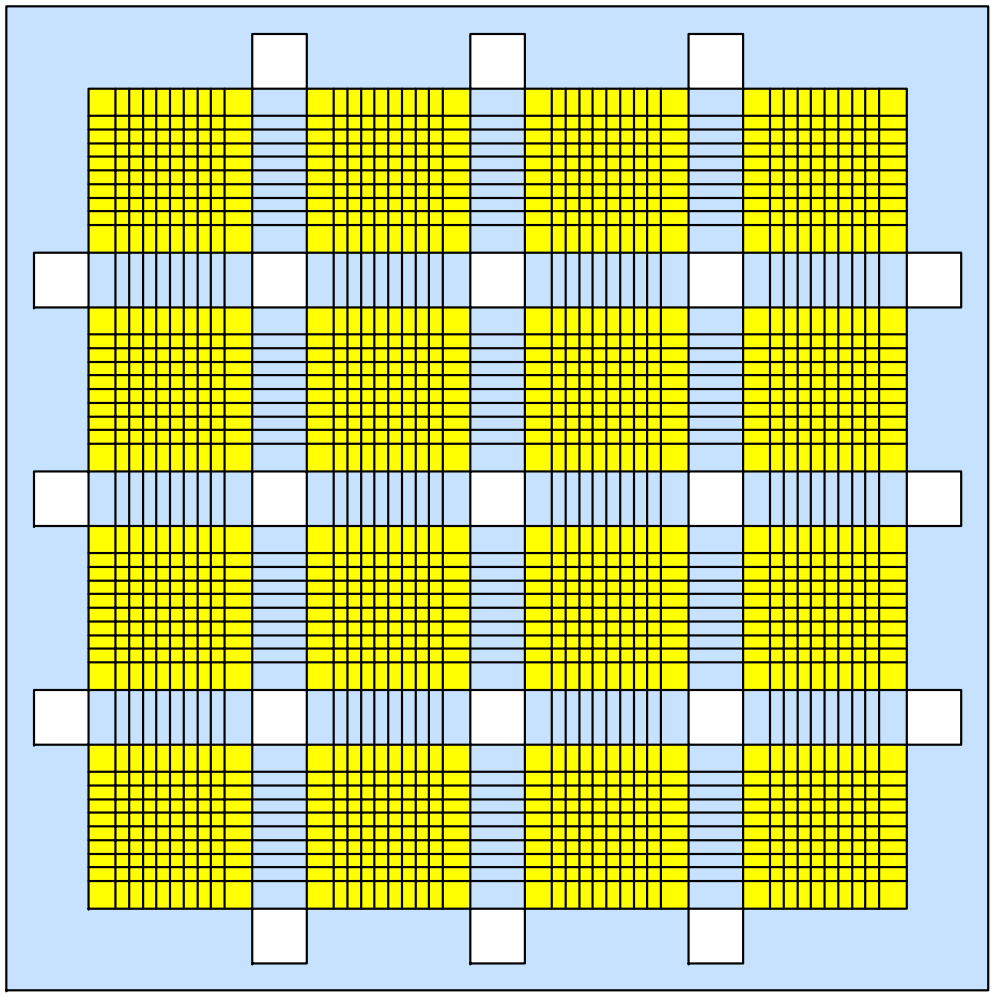

- Most of the FPGA area is due to routing
  - Fixed metal tracks arranged in horizontal and vertical channels
  - Connected to each other using switch blocks

Switch Blocks

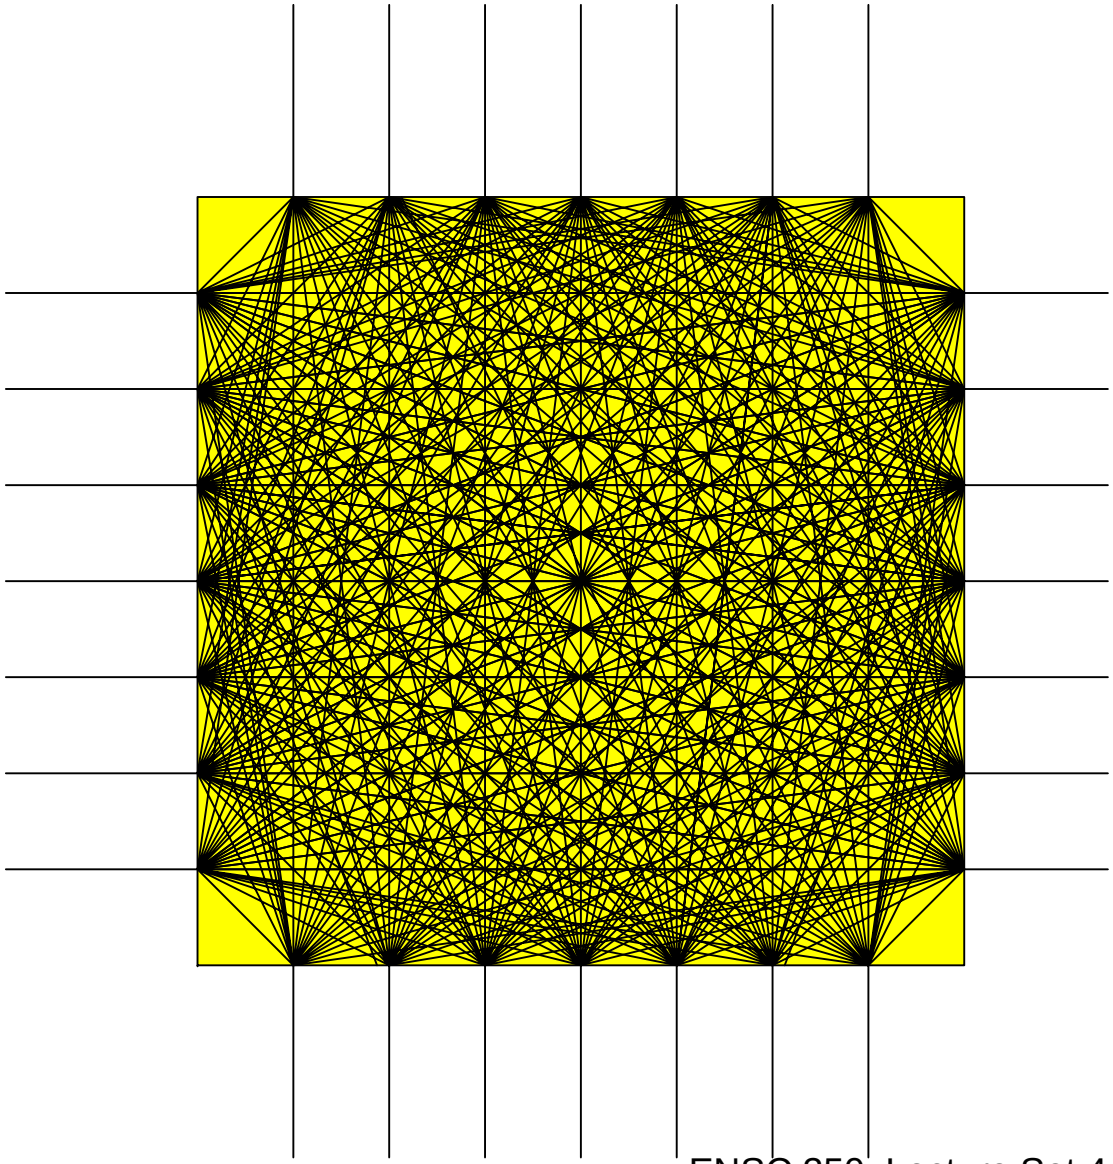

- Switch Blocks connect horizontal and vertical channels
- Every possible connection?
  - Too big
  - Too slow

Switch Blocks

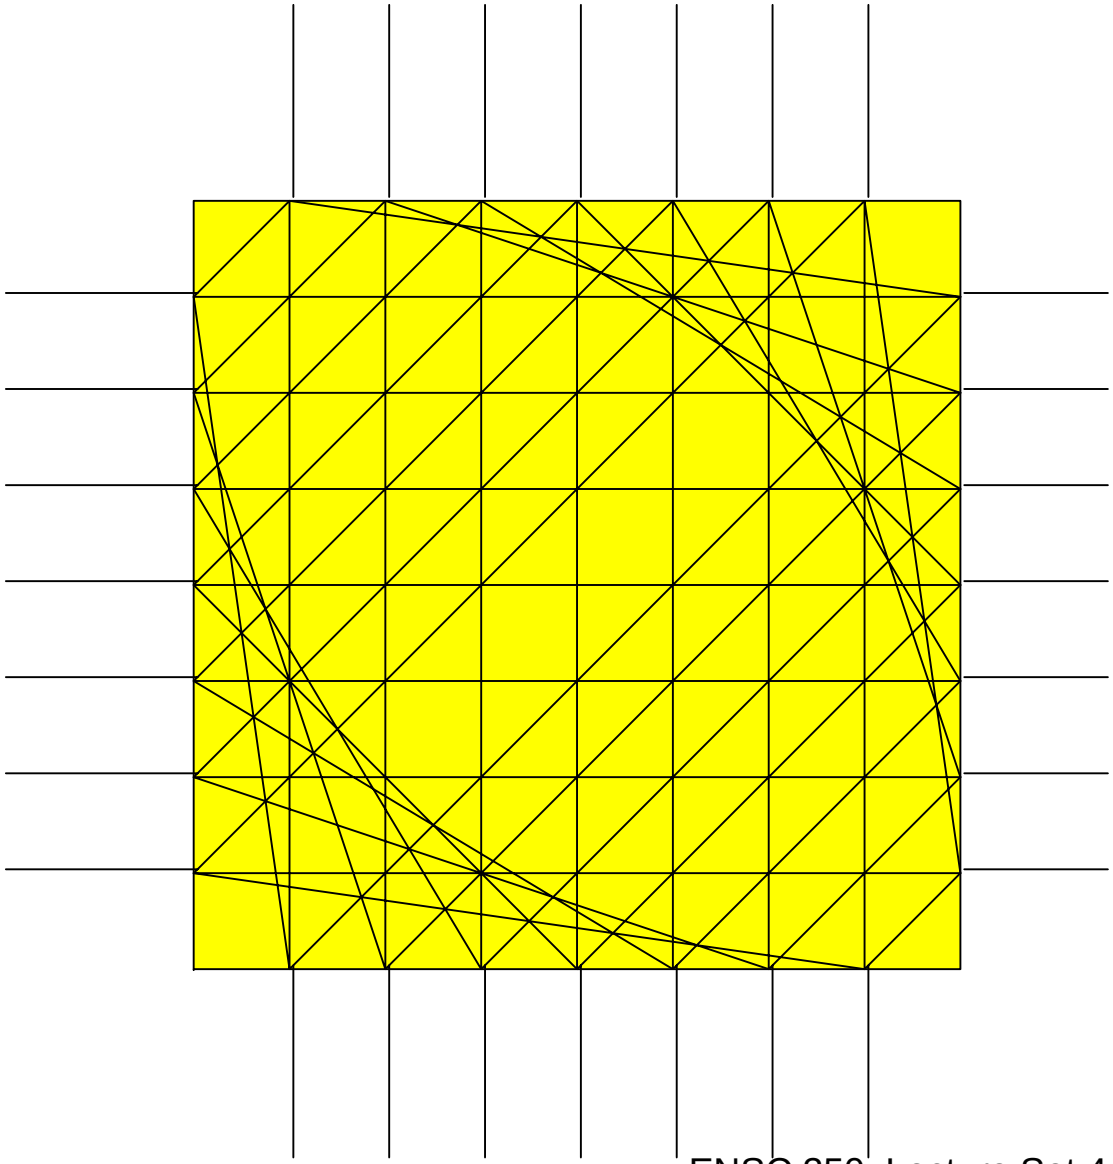

- Switch Blocks connect horizontal and vertical channels
- Every possible connection?
  - Too big
  - Too slow
- Many Topologies possible
  - Fs = 3 is common

Implementing the Switch Block

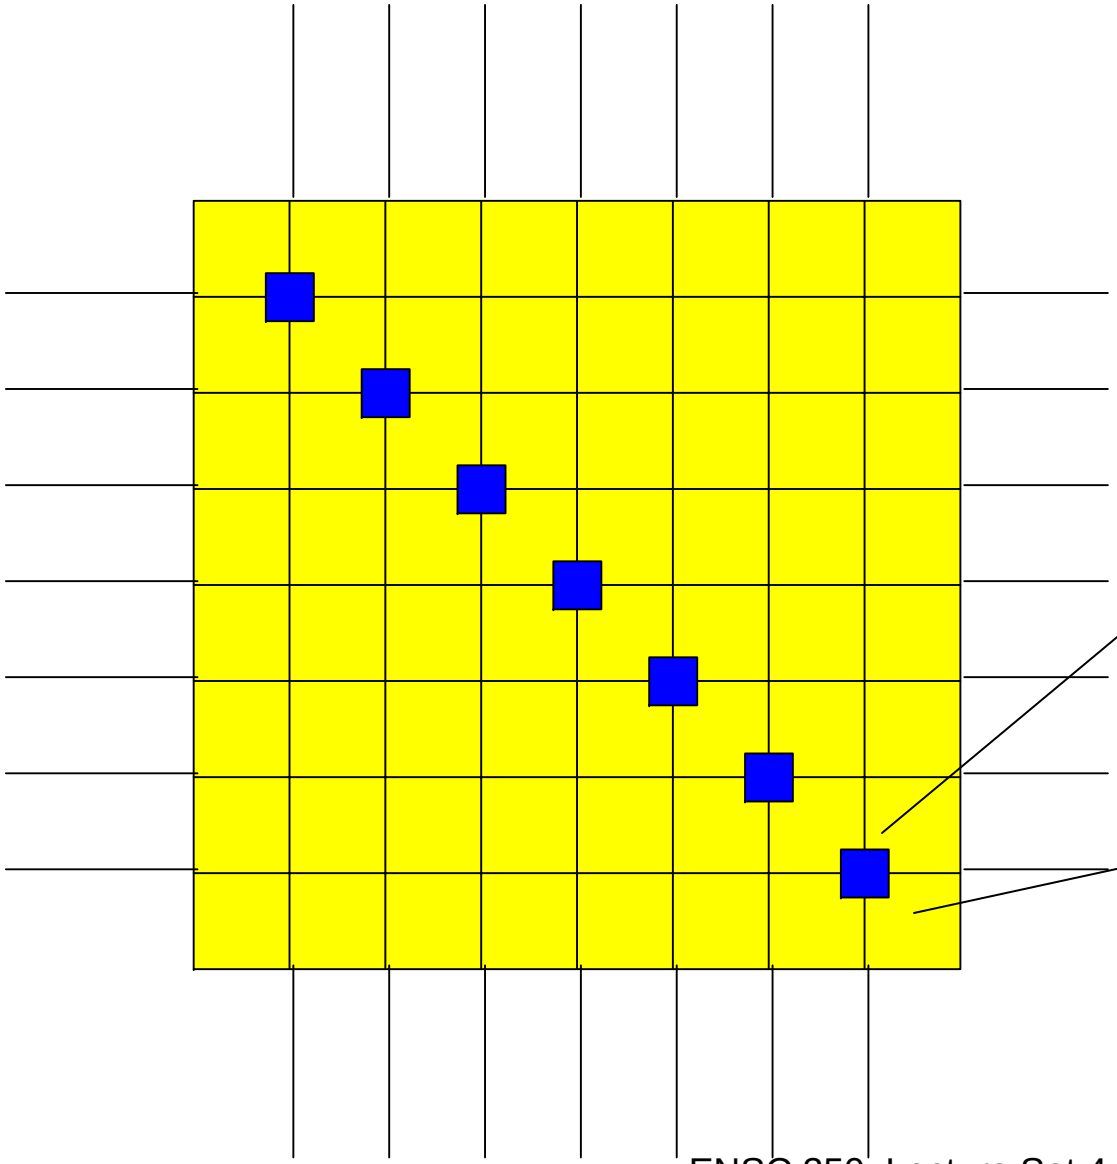

- Circuit-level design of these switch blocks will be considered later

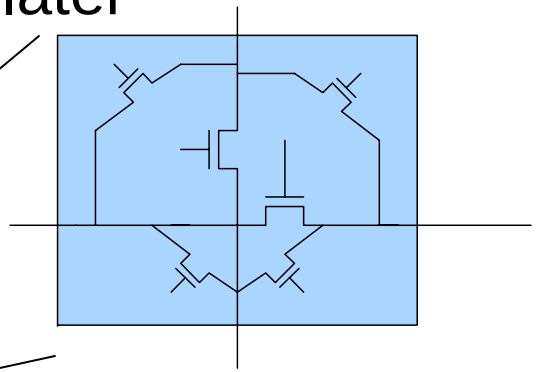

Switch Block Topologies

- 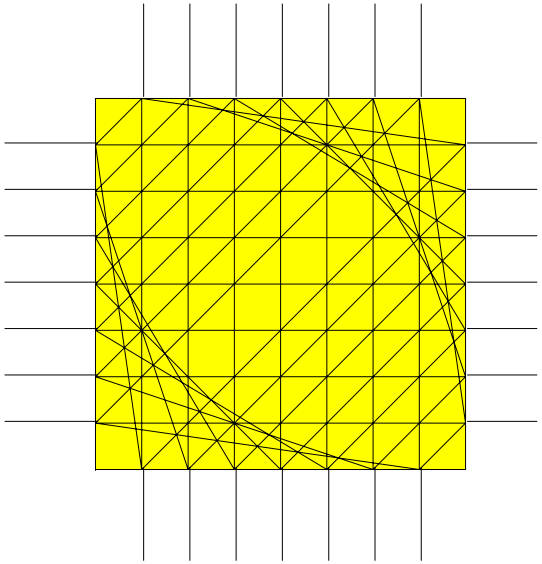

Disjoint

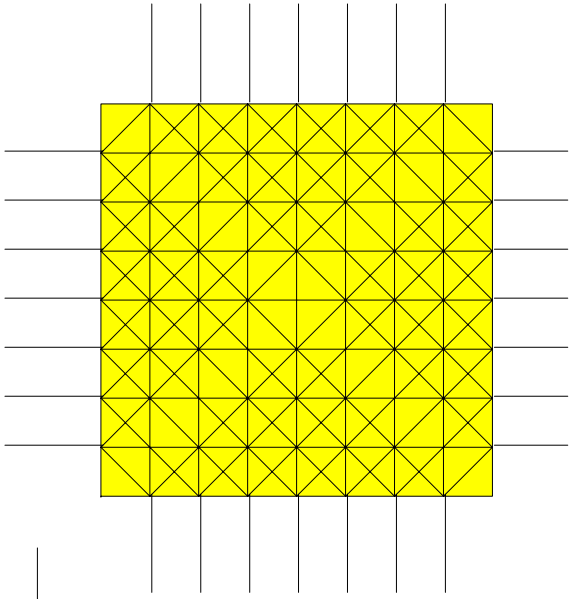

Universal

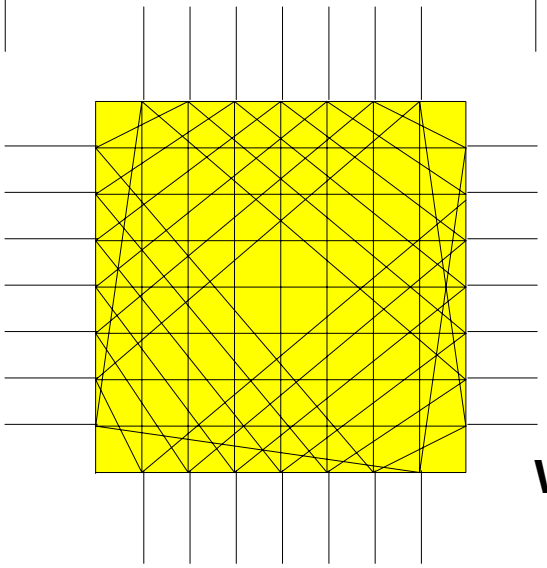

Wilton

Advantage of Wilton Switch Block

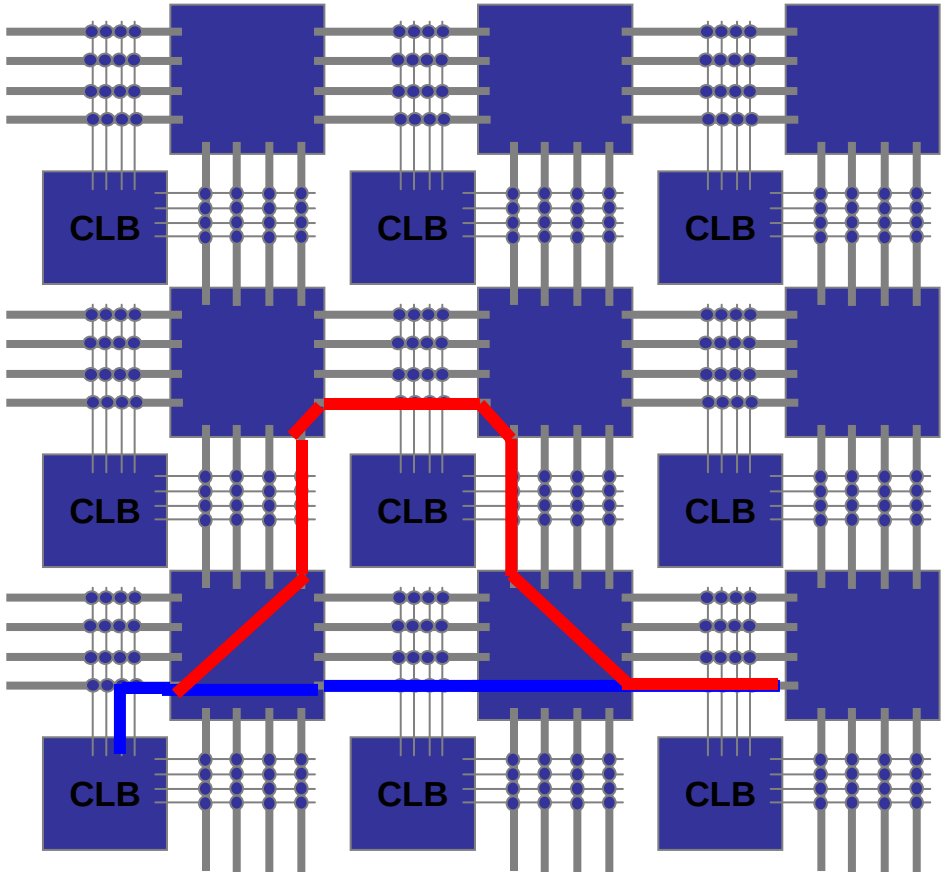

Advantage of Wilton Switch Block

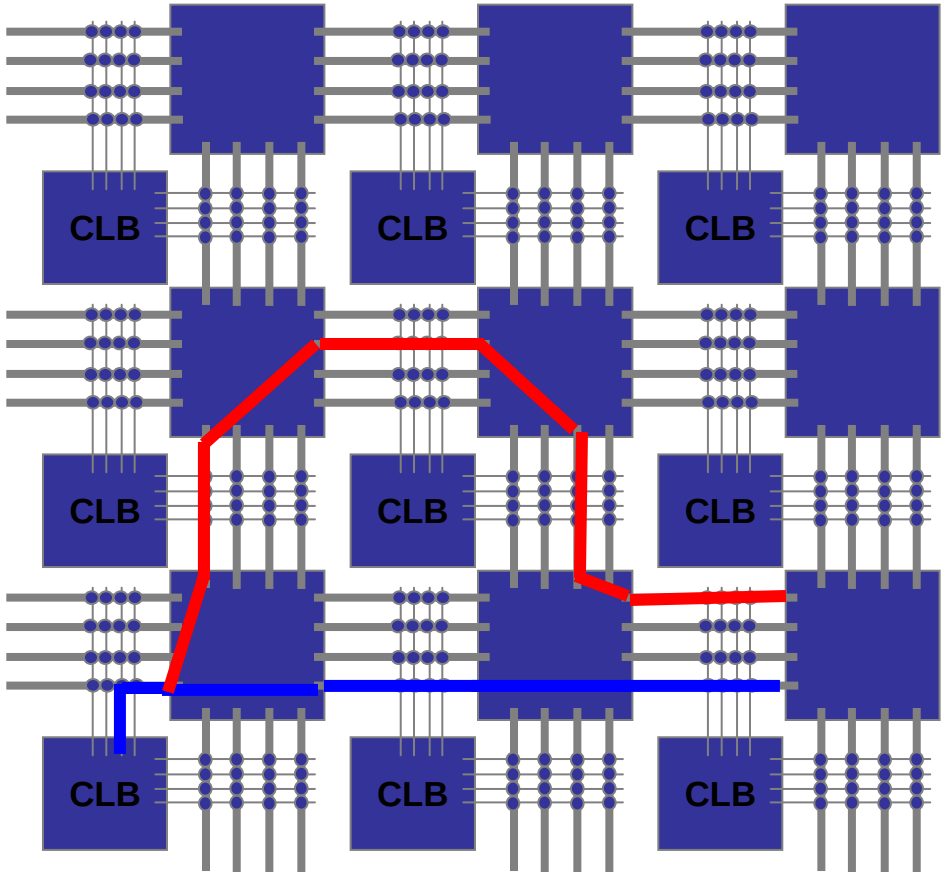

Diversity means you can “get to” more routing tracks.

It tends to provide slightly better routability. No big impact on delay.

Wiring Segments

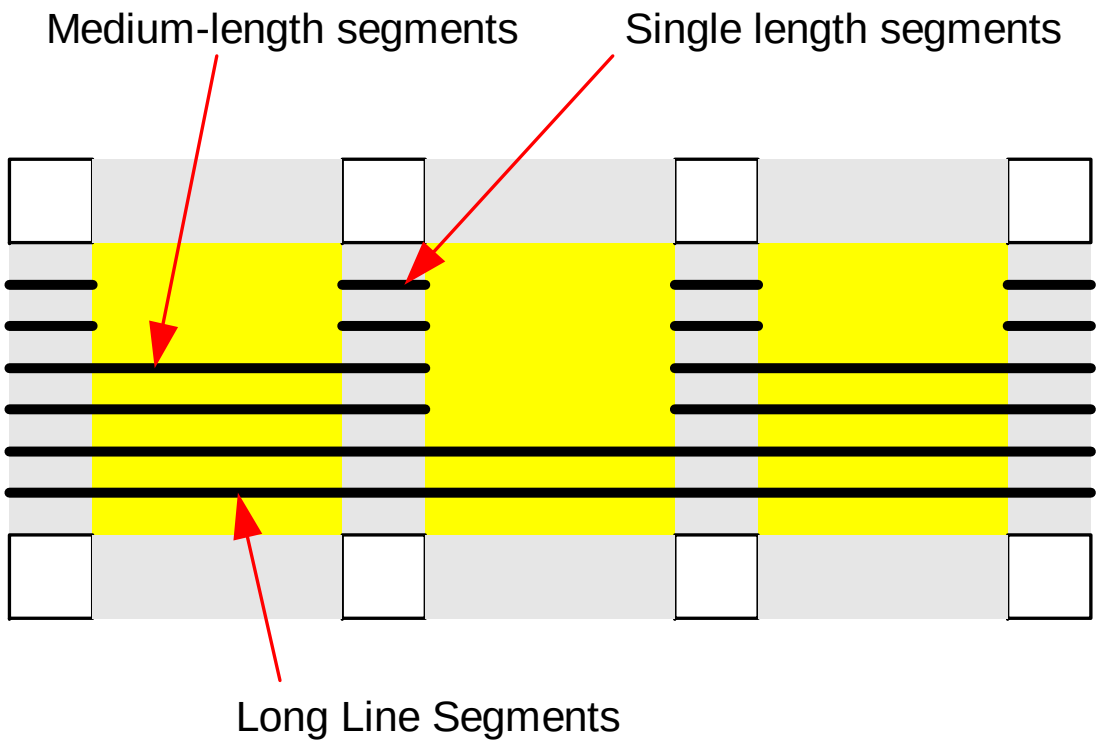

- Short segments are good for local connections
- Long segments are good for global connections
- Most FPGA's have a variety of segment lengths

Segmented Architecture

- At each switch block: some tracks end
  some tracks pass right through

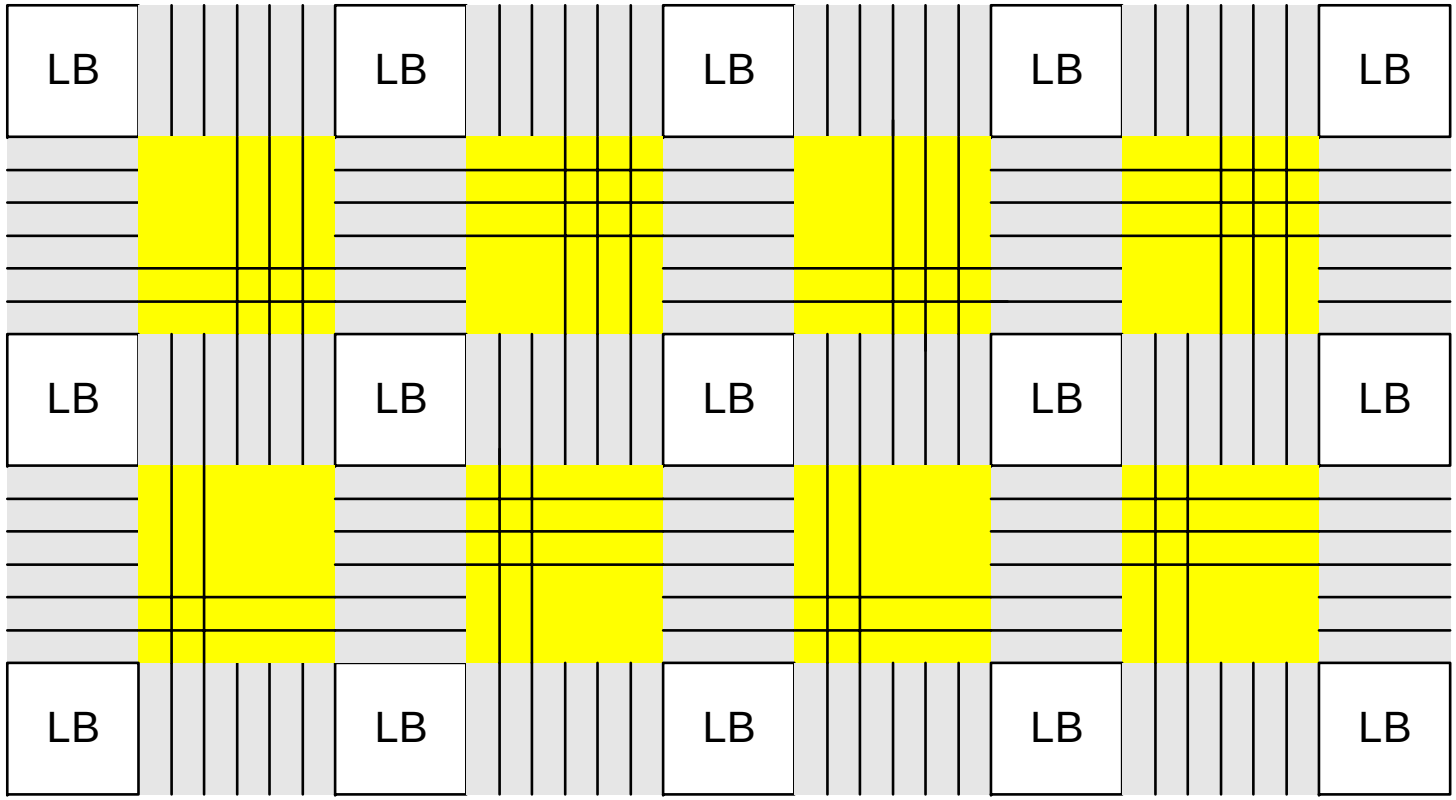

Segment Lengths

- Typically, an FPGA contains a mix of segment lengths:
  - Some wires that span only one logic block
  - Some wires that span more than one logic block
  - Some wires that span the whole chip
- If a segment is too short, must traverse many segments to reach your destination
- If a segment is too long, waste routing capacity, extra capacitance
  - Wires that span the whole chip = high capacitance
- Academic work has suggested length-4 segments

The “Imran” Switch Block

- At each Switch Block, some tracks terminate:

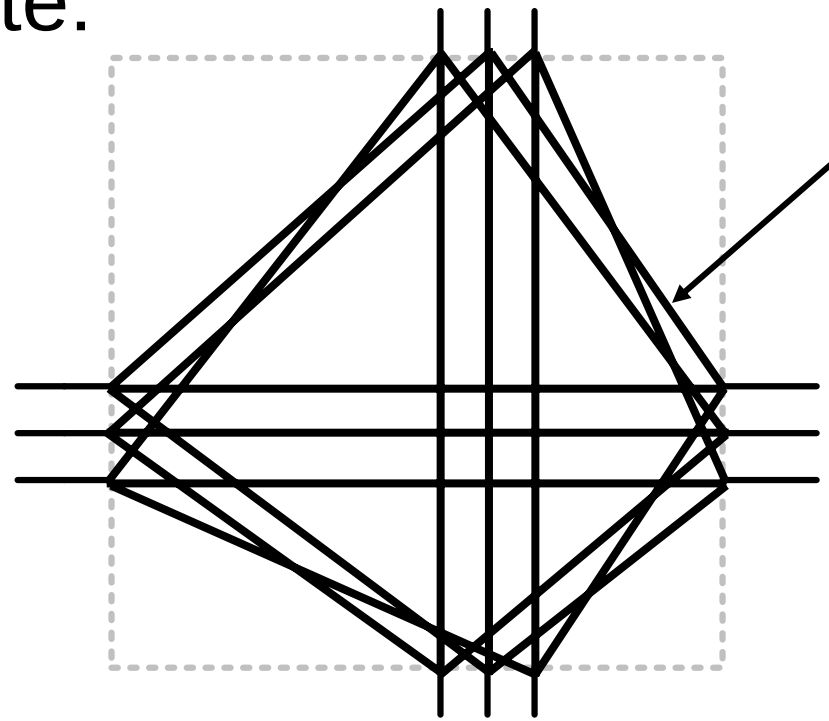

Connect using “Wilton” pattern

The “Imran” Switch Block

- Some tracks do *not* terminate:

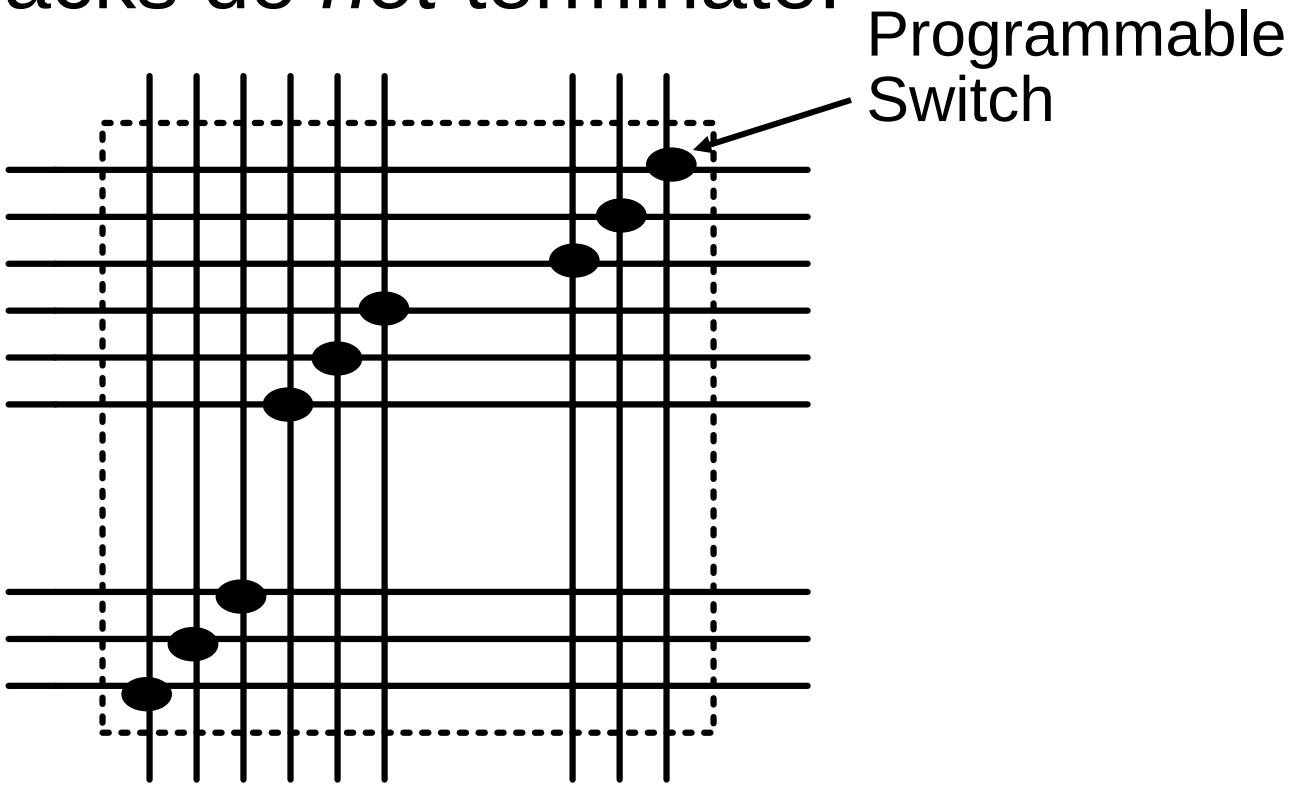

Connect using “Disjoint” pattern

The “Imran” Switch Block

- Put the two together:

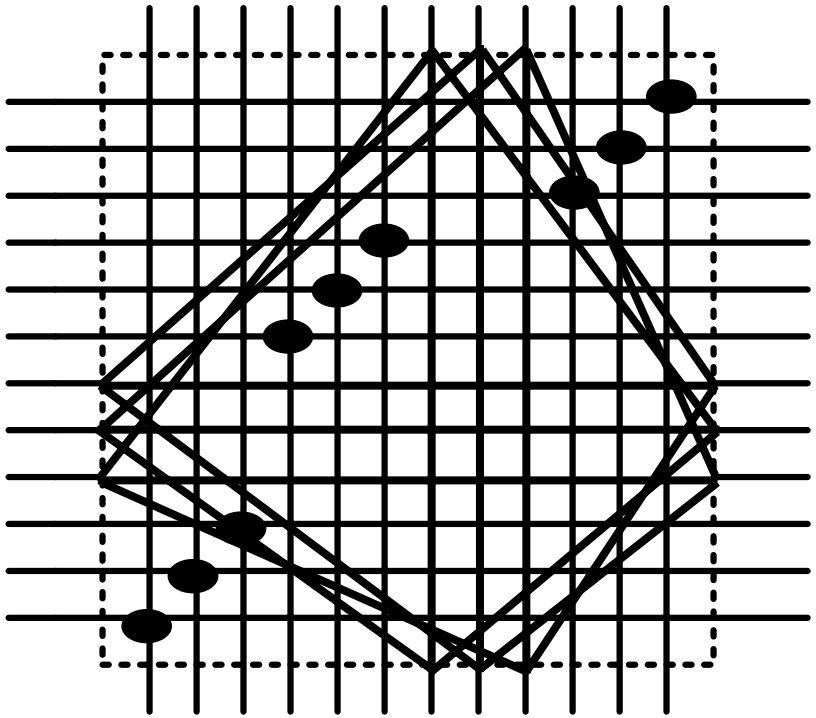

Gives good results for segmented architectures

Connection Blocks

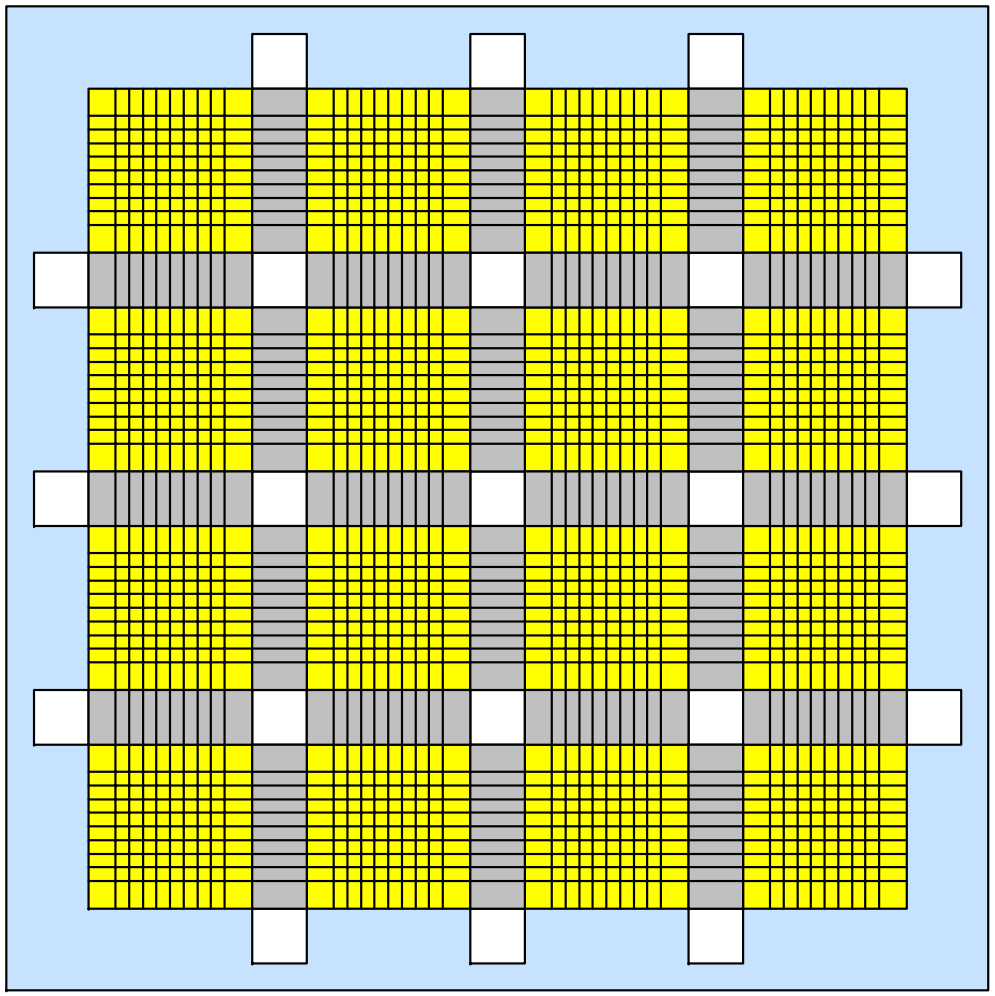

- Most of the FPGA area is due to routing
  - Fixed metal tracks arranged in horizontal and vertical channels
  - Connected to each other using switch blocks
  - Connected to logic blocks using connection blocks

A Connection Block

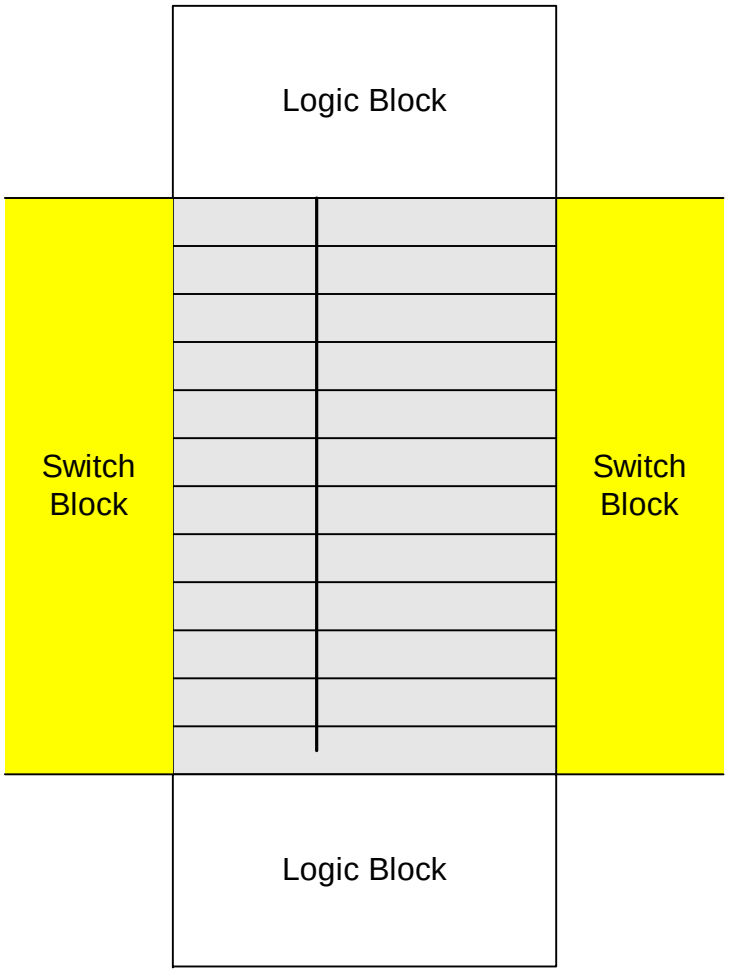

- Each pin can connect to a subset of the tracks in an adjacent channel

Detailed Routing Diagram (XC4000X)

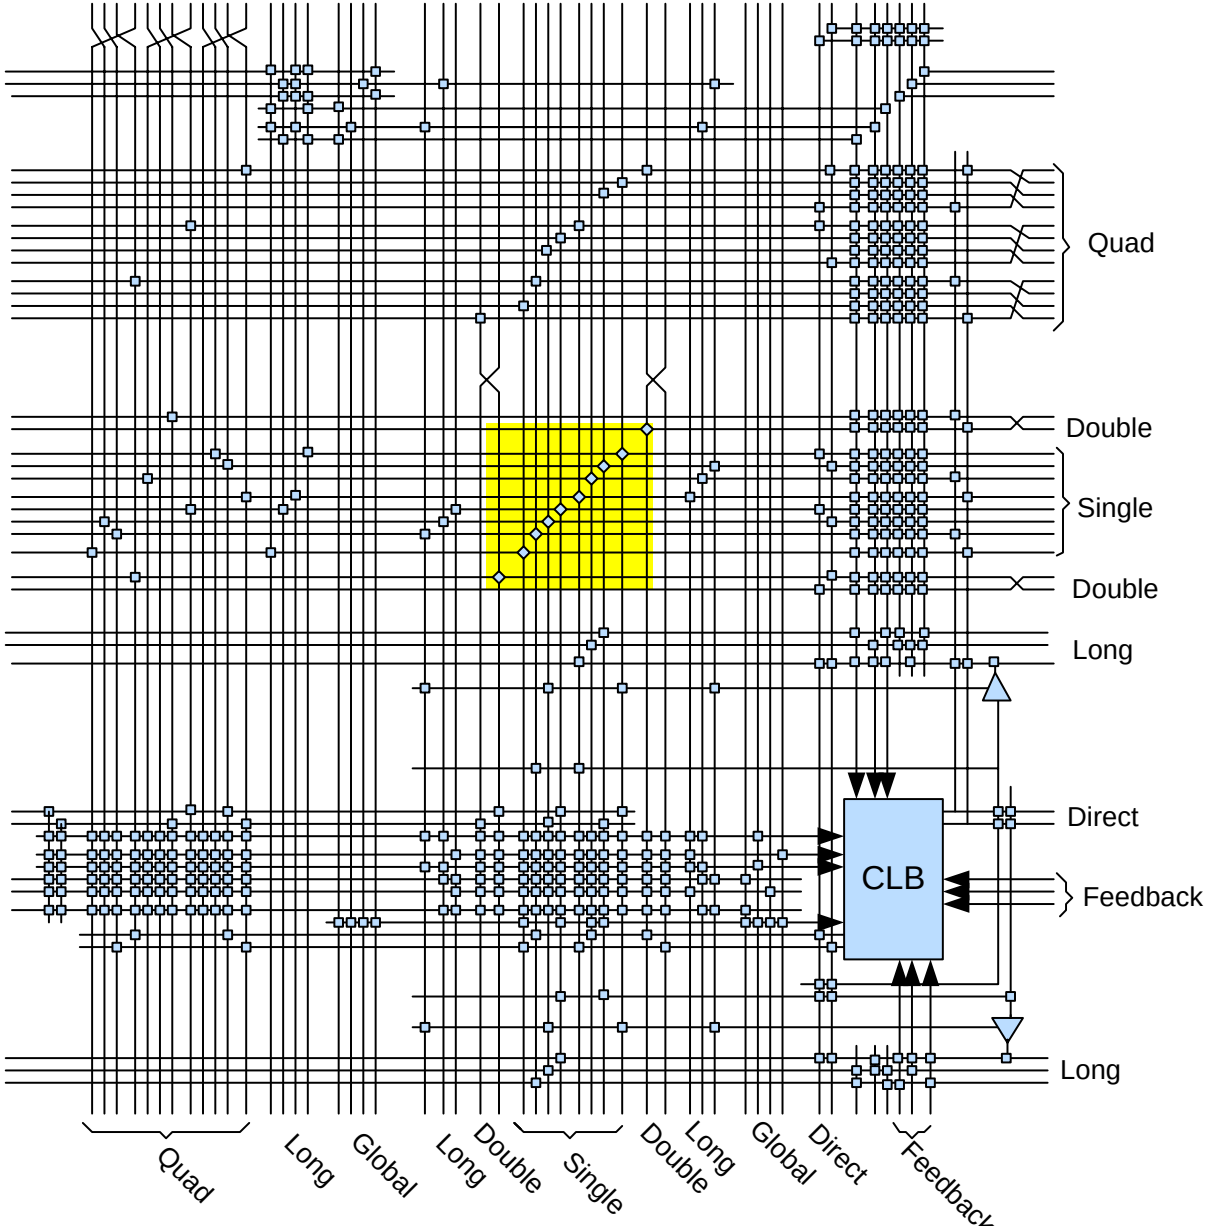

- Dots represent Programmable Connections
- Yes, this is old, but it illustrates the parts.
- Today, vendors don't publish the routing details

Altera Stratix

- Horizontal: R4 Lines, R8 Lines, R28 Lines
- Vertical: C4 Lines, C8 Lines, C16 Lines
- Local Interconnects

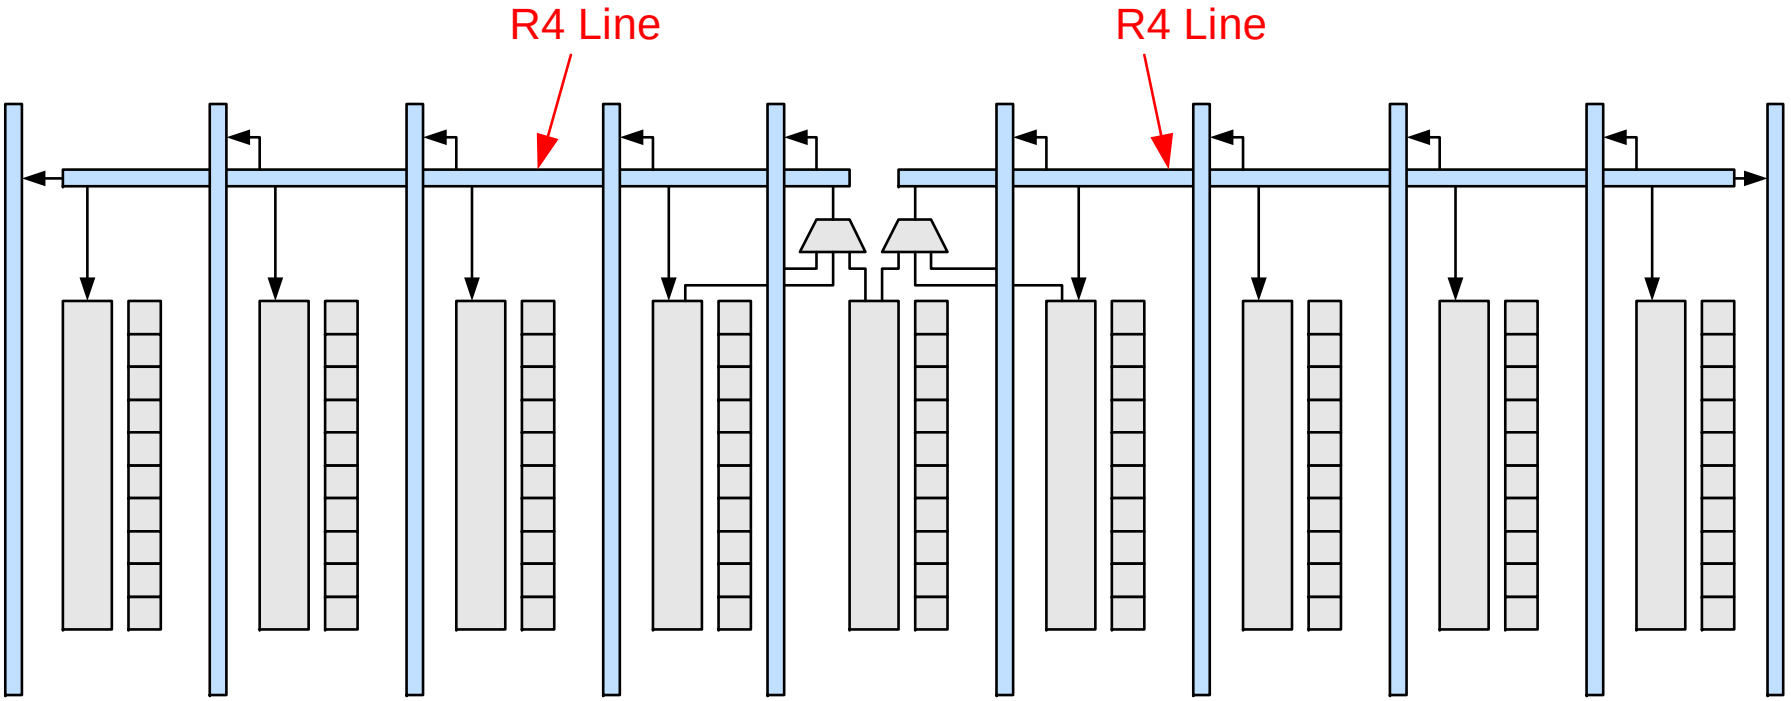

Altera Stratix II

- Horizontal: R4 Lines, R24 Lines
- Vertical: C4 Lines, C16 Lines
- Local Interconnects

They found little benefit to the length-8 lines in Stratix

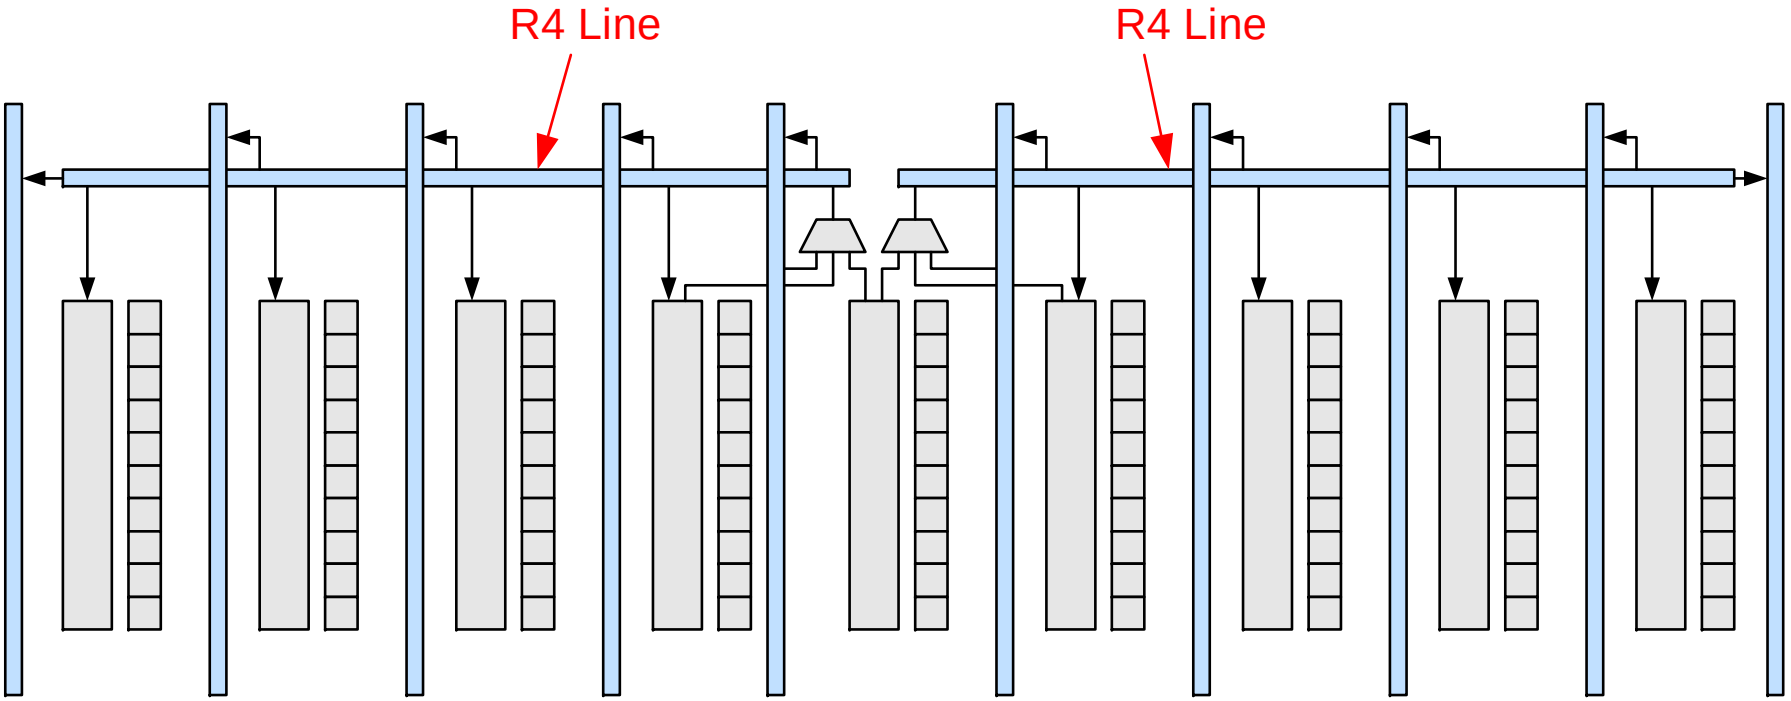

Xilinx Virtex II

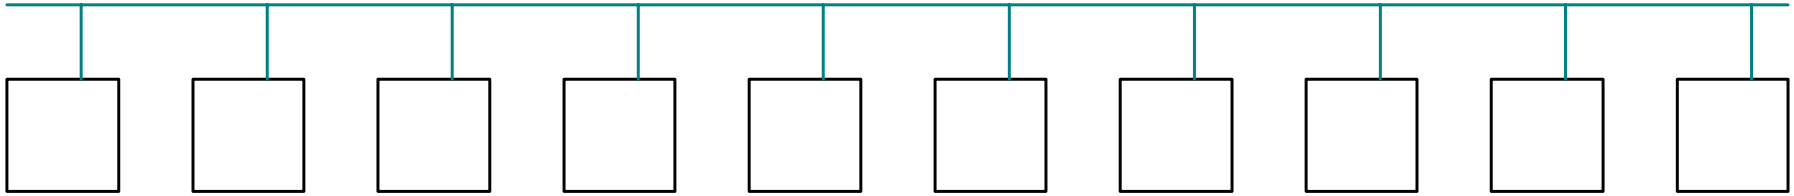

- Long Lines: Span entire chip
  - 24 in each channel (horizontal and vertical)
  - Can to connect to any logic block (actually through the neighbouring switch block)

Xilinx Virtex II

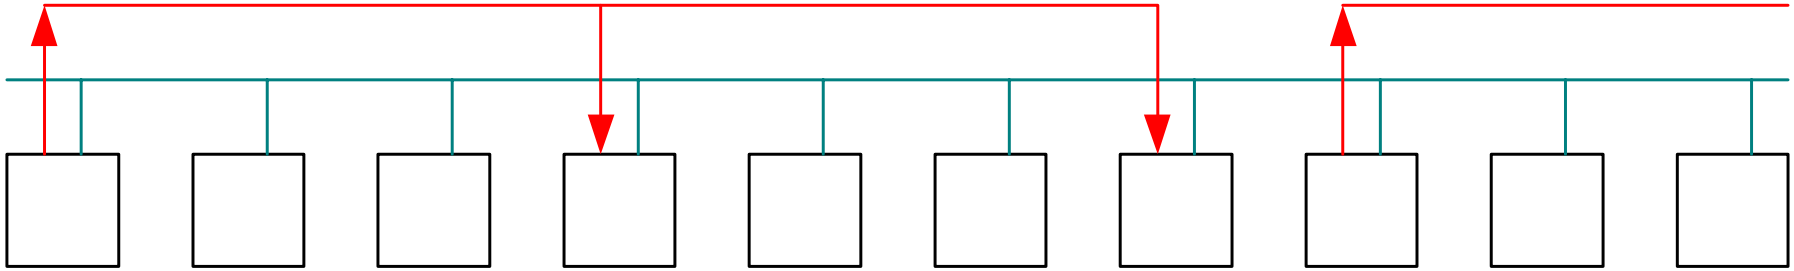

- Hex Lines:
  - 120 in each channel (horizontal and vertical)
  - Can only be driven at one end
  - Two connections to destination logic blocks

Xilinx Virtex II

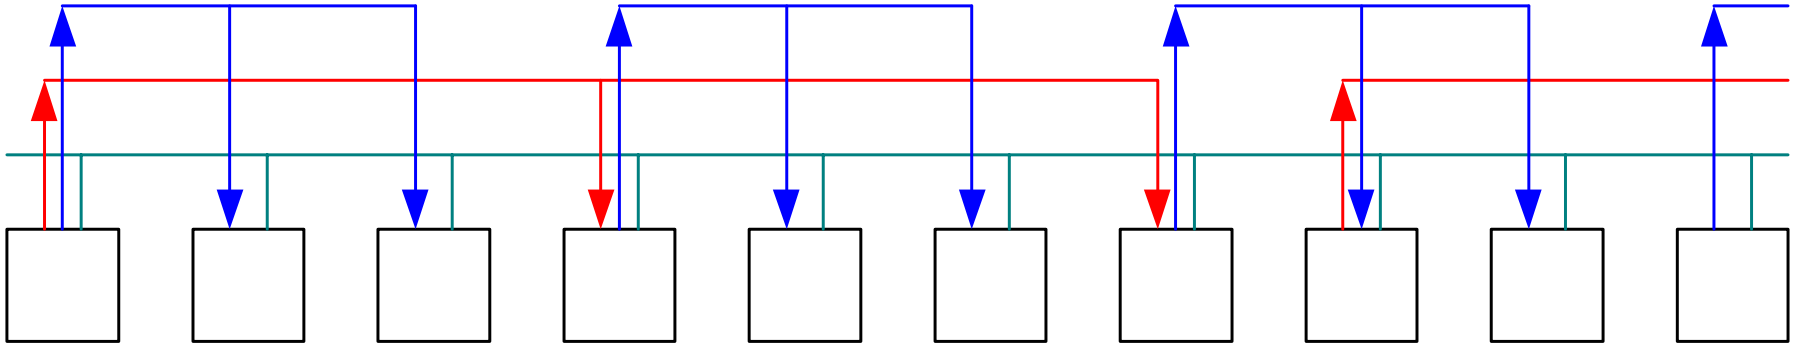

- Double Lines
  - 40 in each channel (horizontal and vertical)
  - Driven at one end

Xilinx Virtex II

- Local Interconnect between neighbouring logic blocks:

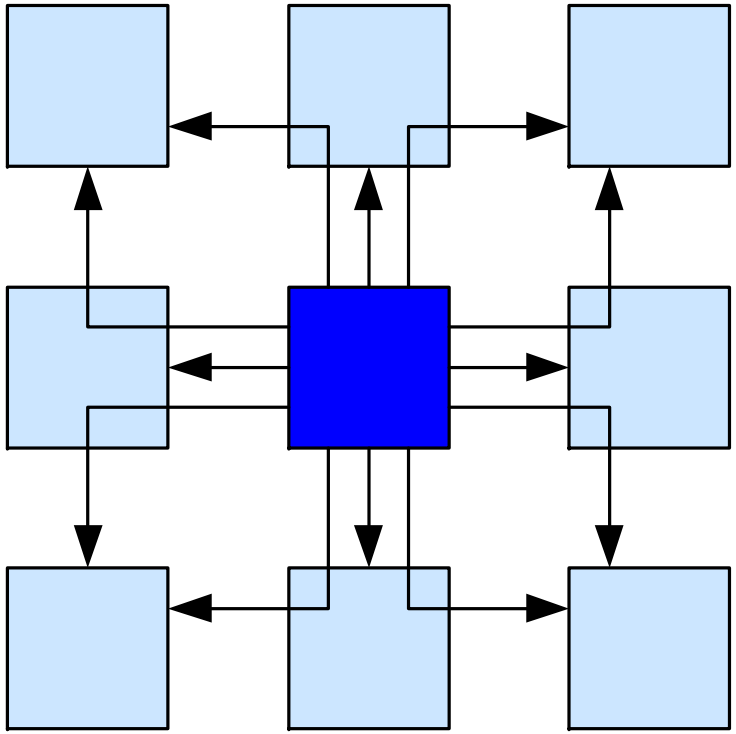

Connection Blocks

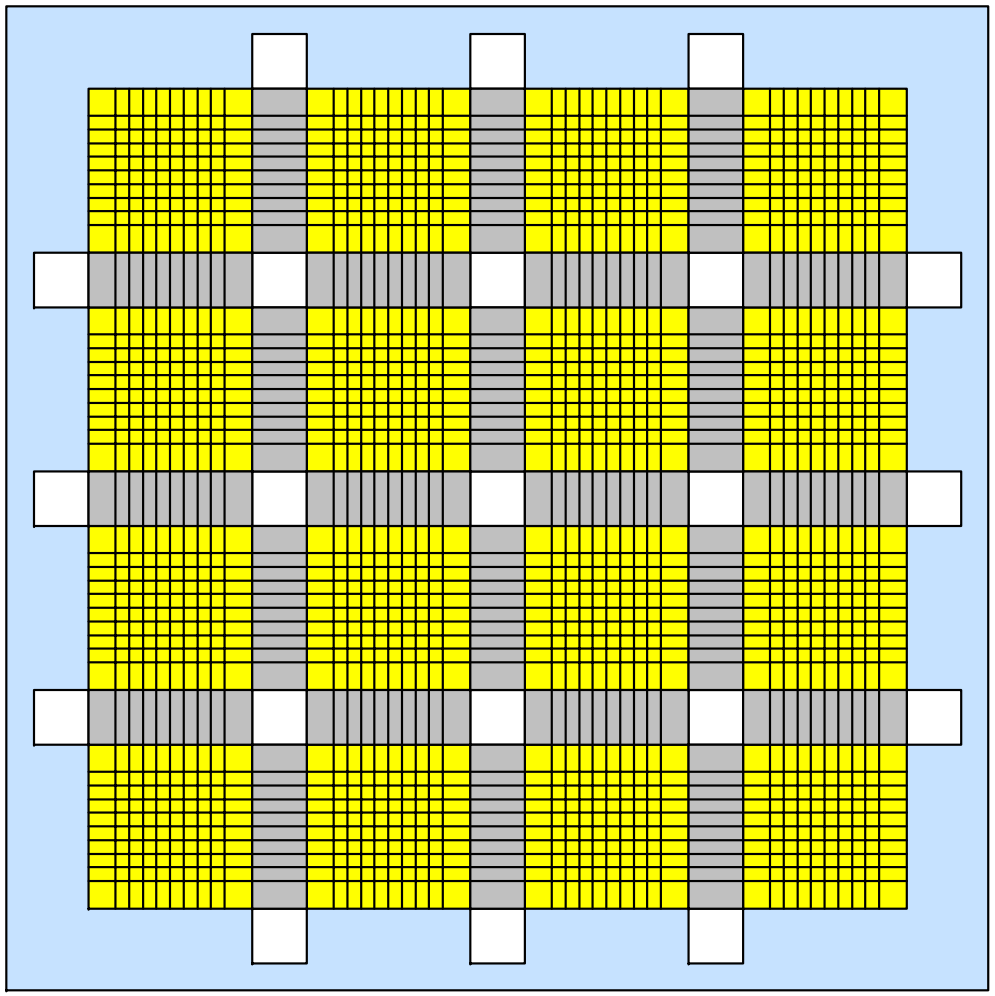

- Most of the FPGA area is due to routing
  - Fixed metal tracks arranged in horizontal and vertical channels
  - Connected to each other using switch blocks
  - Connected to logic blocks using connection blocks

Systems

Implementing Systems in an FPGA

FPGA vendors embed fixed blocks to improve speed and density:

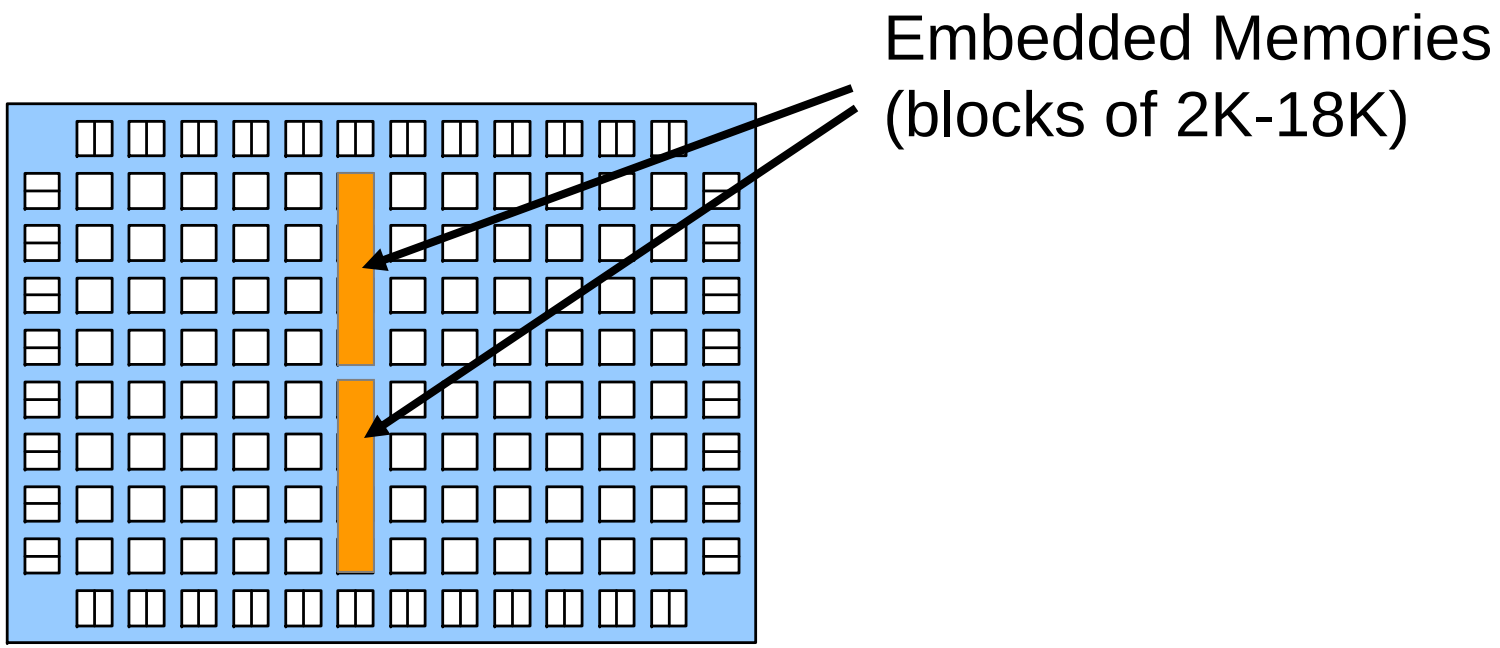

Implementing Systems in an FPGA

FPGA vendors embed fixed blocks to improve speed and density:

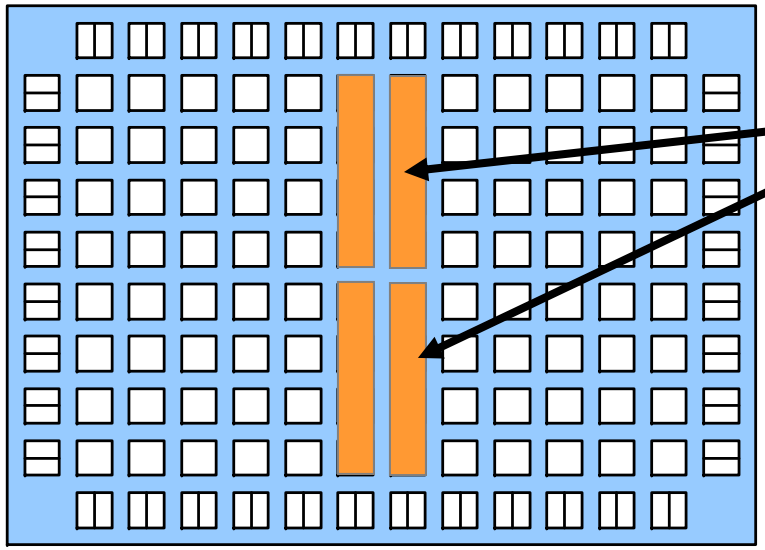

Embedded Memories (blocks of 2K-18K)

Multiplier Blocks

Implementing Systems in an FPGA

FPGA vendors embed fixed blocks to improve speed and density:

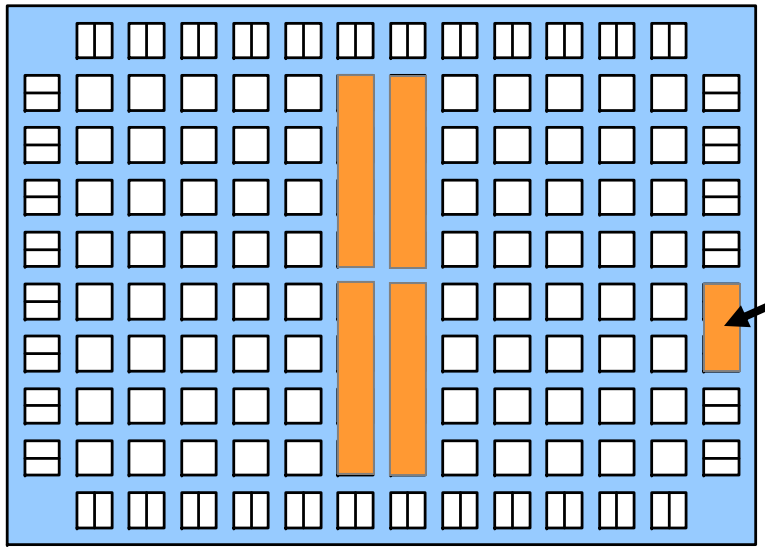

Embedded Memories (blocks of 2K-18K)

Multiplier Blocks

High-Speed I/Os

Implementing Systems in an FPGA

FPGA vendors embed fixed blocks to improve speed and density:

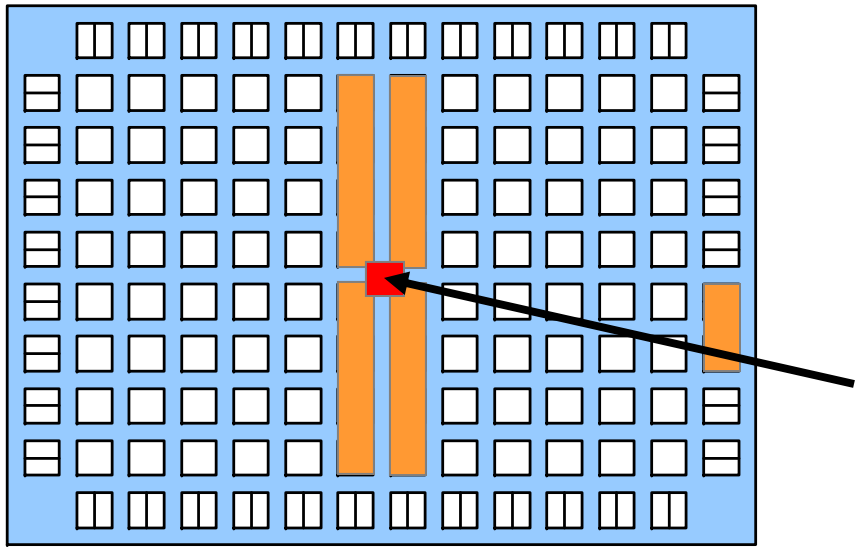

Embedded Memories (blocks of 2K-18K)

Multiplier Blocks

High-Speed I/Os

Dedicated Clock Circuitry

Implementing Systems in an FPGA

FPGA vendors embed fixed blocks to improve speed and density:

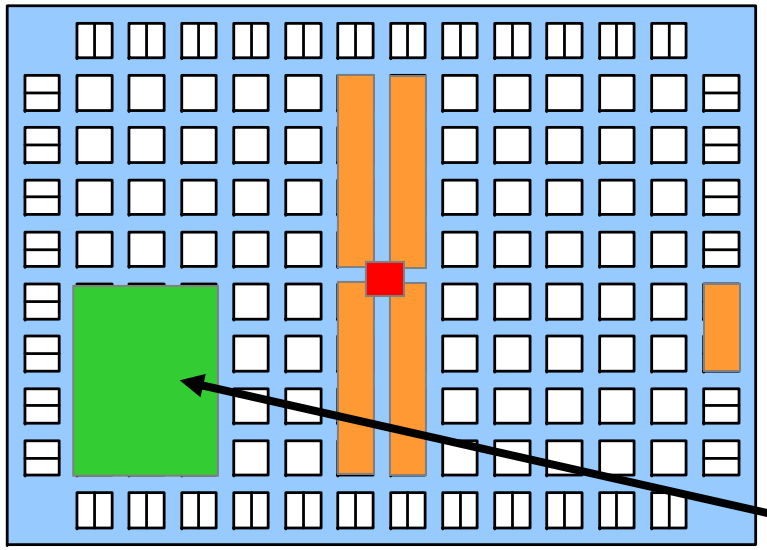

Embedded Memories (blocks of 2K-18K)

Multiplier Blocks

High-Speed I/Os

Dedicated Clock Circuitry

CPU (eg. PowerPC)

Summary

Two Sources of Flexibility in an FPGA:

1. Most FPGAs use Lookup-Tables as their basic logic resource
   - 4-LUT can implement any function of 4 inputs
   - Modern FPGAs are moving to 6-LUTs

2. Connections between logic blocks can be made using fixed metal tracks
   - these fixed tracks are connected to each other and to the logic blocks using programmable switches
   - Consume most of the area and power on the chip

Questions

- What are the fundamental components of a traditional FPGA?

- What are the newer embedded hard IP blocks in modern FPGAs?

Questions

- List 3 advantages of FPGAs.

- List 3 disadvantages of FPGAs.

Questions

- What consumes the largest portion of the FPGA fabric?

- What are the basic components of a CLB?

Questions

- What are the pros and cons of clustering?

- What are the two main components of the programmable interconnect?

Questions

- Why do FPGAs have multiple wire segment lengths? How do we measure them?

- Do clock circuits use the same routing interconnect as data I/Os?

Questions

- What's the difference between switch blocks and connection blocks?

- Be able to describe at least 2 different switch block architectures and give one pro and con of each