

# RFNoC Crossbar Architectures

---

Background

---

# Networking 101

## Main Components of a Network

### 1. Network Topology

- The arrangement of network components (terminals, routers, links, etc)

### 2. Routing Algorithm

- The path selection for packets traversing the network

### 3. Flow Control


- The process of managing the rate of transmission between components

### 4. Microarchitecture

- The organization and implementation of router components

# Definitions

- The **throughput** of a network is the data rate (bps) that the network accepts per input port
- The **injection bandwidth** is the max throughput for a given channel
- The **channel load** is the ratio between the bandwidth on a channel to the injection bandwidth
- The **latency** is the amount of time it takes to traverse the network




# Load vs Latency Graph

Used to evaluate network performance under various load conditions. Traces represent percentile latency for various implementations. Conveys:

- Statistical spread of latency
- Maximum load tolerated by network

Examples:

1. Blue: High latency network with high load capacity
2. Orange: Low latency network with low capacity



# Traffic Patterns

Network performance depends on the traffic pattern i.e. the pattern of the source to destination paths.

Examples:

- **Uniform:** A node sends to all nodes with equal probability
- **Uniform Others:** A node sends to all other nodes with equal probability
- **Neighbor:** A node only sends traffic to their neighbor
- **Bit Complement:** A node only sends traffic to the diametrically opposite node
- **Sequential:** A node sends to all nodes sequentially
- **Loopback:** A node sends to itself

Uniform Random  
(with self)

0.25	0.25	0.25	0.25
0.25	0.25	0.25	0.25
0.25	0.25	0.25	0.25
0.25	0.25	0.25	0.25

Uniform Random  
(without self)

0.00	0.33	0.33	0.33
0.33	0.00	0.33	0.33
0.33	0.33	0.00	0.33
0.33	0.33	0.33	0.00

Neighbor  
 $dest=(src+1)\%N$

0.00	1.00	0.00	0.00
0.00	0.00	1.00	0.00
0.00	0.00	0.00	1.00
1.00	0.00	0.00	0.00

Bit Complement  
 $dest=(\sim src)\%N$

0.00	0.00	0.00	1.00
0.00	0.00	1.00	0.00
0.00	1.00	0.00	0.00
1.00	0.00	0.00	0.00

Random  
Permutation

1.00	0.00	0.00	0.00
0.00	0.00	0.00	1.00
0.00	1.00	0.00	0.00
0.00	0.00	1.00	0.00

Other


0.20	0.40	0.30	0.10
1.00	0.00	0.00	0.00
0.05	0.70	0.15	0.10
0.00	0.50	0.50	0.00

# RFNoC Crossbar Implementations

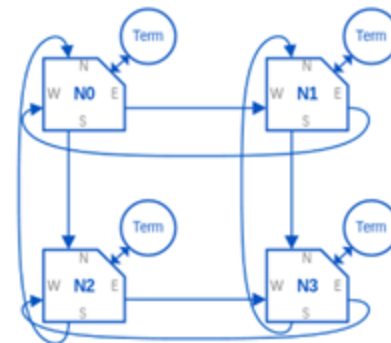
---

# axis\_ctrl\_crossbar\_2d\_mesh

- **Topology:** Bidirectional Mesh or Unidirectional Torus
- **Routing:** Wormhole or Store-and-Fwd
  - Wormhole: A packet can be in several routers at a time.
  - Store-and-Fwd: A packet is completely buffered in one router before moving to the next one.
- **Flow Control:** Packet buffer with cut-through
  - Packet buffer: Entire packet is buffered in router
  - Cut-through: Only flits (words) are backpressured



2x2 Mesh




2x2 Torus




# axis\_ctrl\_crossbar\_2d\_mesh

## Microarchitecture

This crossbar has been optimized for low-throughput control and can scale to a large number of ports. The underlying implementation uses a 2-dimensional mesh topology which can be configured either as a bidirectional mesh or a unidirectional torus. The crossbar is not deadlock free by design but there are various features implemented that reduce the possibility of deadlock. In the event of a deadlock the crossbar will self-recover by dropping all the packets in flight. This behaviour makes the crossbar lossy. The underlying mesh topology can only be a square so the number of ports supported are  $N^2$  for  $N > 1$ . All unused ports need to be terminated using the `axis_port_terminator` module.



2x2 Mesh



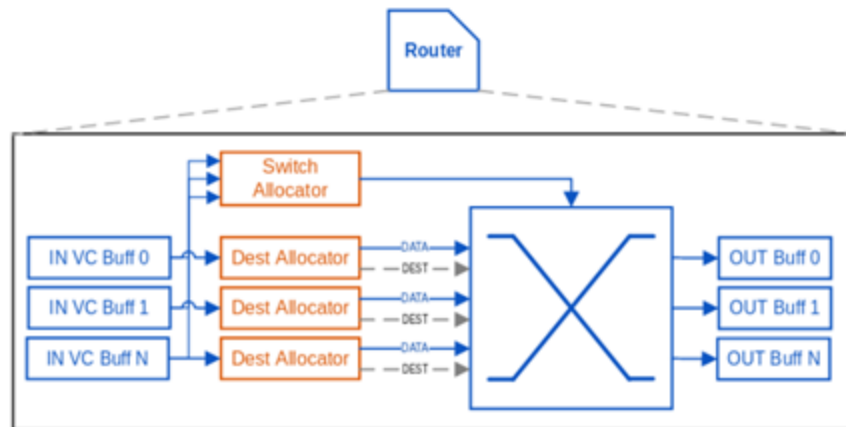
2x2 Torus

# axis\_ctrl\_crossbar\_2d\_mesh

## Router Architecture

Each node in the crossbar is comprised of a terminal and a router. A terminal is the interface to client logic and the network of routers performs packet switching. It consists of:


- A switch to do the routing
- Ingress buffers with one or more virtual channels
- A switch allocator to choose an input port to drive the switch
- A destination selector to choose the destination port
- Egress buffers



# axis\_ctrl\_crossbar\_2d\_mesh

## Scaling

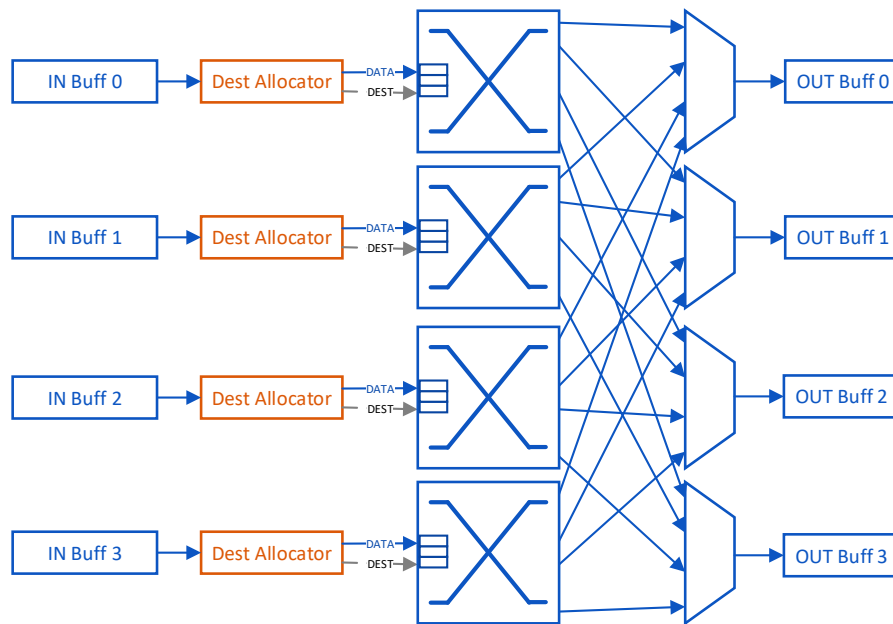
The mesh can scale to an arbitrary number of ports.



# chdr\_crossbar\_nxn

- **Topology:** Trivial. Single router.
- **Routing:** Store-and-Fwd
  - Store-and-Fwd: A packet is completely buffered in one router before moving to the next one.
- **Flow Control:** Packet buffer with cut-through
  - Packet buffer: Entire packet is buffered in router
  - Cut-through: Only flits (words) are backpressured

\*Alternative for axi\_crossbar\*




# Performance Metrics: Control Crossbar


---

# Load vs Latency: **axis\_ctrl\_crossbar\_2d\_mesh** (TORUS, 25 nodes, 25 traffic generators)

Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: LOOPBACK, LPP: 10)




Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: SEQUENTIAL, LPP: 10)




# Load vs Latency: **axis\_ctrl\_crossbar\_2d\_mesh** (TORUS, 25 nodes, 25 traffic generators)


Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: UNIFORM, LPP: 10)



Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: UNIFORM\_OTHERS, LPP: 10)




# Load vs Latency: **axis\_ctrl\_crossbar\_2d\_mesh** (TORUS, 25 nodes, 25 traffic generators)






# Load vs Latency: **axis\_ctrl\_crossbar\_2d\_mesh** (TORUS, 25 nodes, 4 traffic generators)

Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: LOOPBACK, LPP: 10)




Load Latency Graph for axis\_ctrl\_2d\_torus  
(Traffic: SEQUENTIAL, LPP: 10)




# Load vs Latency: `axis_ctrl_crossbar_2d_mesh` (TORUS, 25 nodes, 4 traffic generators)


Load Latency Graph for `axis_ctrl_2d_torus`  
(Traffic: UNIFORM, LPP: 10)




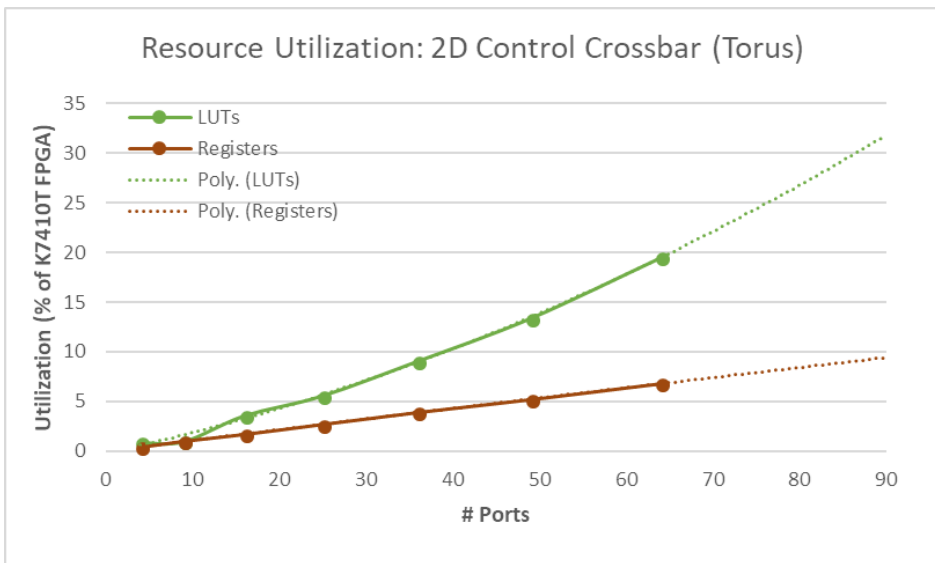
Load Latency Graph for `axis_ctrl_2d_torus`  
(Traffic: UNIFORM\_OTHERS, LPP: 10)



# Load vs Latency: **axis\_ctrl\_crossbar\_2d\_mesh** (TORUS, 25 nodes, 4 traffic generators)




# FPGA Utilization: axis\_ctrl\_crossbar\_2d\_mesh



\* Utilization percentages are relative to the X310 FPGA (Kintex7 410T Speed Grade 3)

# Timing: `axis_ctrl_crossbar_2d_mesh`



\* Design was synthesized for the X310 FPGA (Kintex7 410T Speed Grade 3)


\*\* Fmax is post-synthesis and pre-optimization

# Performance Metrics: CHDR Crossbar


---

# Load vs Latency: **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)

Load Latency Graph for chdr\_crossbar\_nxn  
(Traffic: LOOPBACK, LPP: 100)




Load Latency Graph for chdr\_crossbar\_nxn  
(Traffic: SEQUENTIAL, LPP: 100)




# Load vs Latency: **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)

Load Latency Graph for chdr\_crossbar\_nxn  
(Traffic: UNIFORM, LPP: 100)

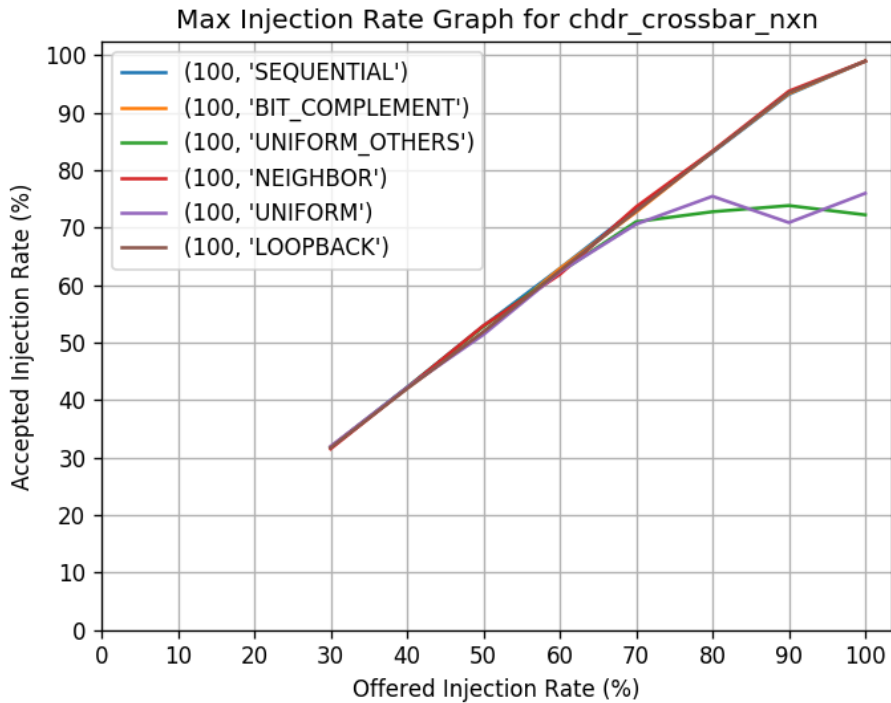


Load Latency Graph for chdr\_crossbar\_nxn  
(Traffic: UNIFORM\_OTHERS, LPP: 100)






# Load vs Latency: **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)




# FPGA Utilization: chdr\_crossbar\_nxn

## Resource Utilization: CHDR Crossbar (64 bit)




## Resource Utilization: CHDR Crossbar (256 bit)



\* Utilization percentages are relative to the X310 FPGA (Kintex7 410T Speed Grade 3)

# Timing: **chdr\_crossbar\_nxn**



\* Design was synthesized for the X310 FPGA (Kintex7 410T Speed Grade 3)


\*\* Fmax is post-synthesis and pre-optimization

# Performance Comparison: Old vs New CHDR Crossbar


---

# Load vs Latency: **axi\_crossbar** vs **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)

Load Latency Graph for **axi\_crossbar**  
(Traffic: SEQUENTIAL, LPP: 100)




Load Latency Graph for **chdr\_crossbar\_nxn**  
(Traffic: SEQUENTIAL, LPP: 100)




# Load vs Latency: **axi\_crossbar** vs **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)

Load Latency Graph for **axi\_crossbar**  
(Traffic: NEIGHBOR, LPP: 100)




Load Latency Graph for **chdr\_crossbar\_nxn**  
(Traffic: NEIGHBOR, LPP: 100)




# Load vs Latency: **axi\_crossbar** vs **chdr\_crossbar\_nxn** (12 nodes, 12 traffic generators)

Load Latency Graph for **axi\_crossbar**  
(Traffic: UNIFORM, LPP: 100)




Load Latency Graph for **chdr\_crossbar\_nxn**  
(Traffic: UNIFORM, LPP: 100)



# Load vs Latency: axi\_crossbar vs chdr\_crossbar\_nxn (12 nodes, 12 traffic generators)

Load Latency Graph for axi\_crossbar  
(Traffic: UNIFORM\_OTHERS, LPP: 100)



Load Latency Graph for chdr\_crossbar\_nxn  
(Traffic: UNIFORM\_OTHERS, LPP: 100)

