The Routability of Multiprocessor Network Topologies in FPGAs

by

Manuel Saldaña, Lesley Shannon and Paul Chow

8th International Workshop on
System Level Interconnect Prediction (SLIP06)
March 4-5, 2006
Munich, Germany

Overview

- Motivation
- Infrastructure
- Experiments and Results
- Conclusions
- Future Work

Motivation

- FPGAs have a **fixed set of resources** that must be used not wasted
- How do well-known network topologies map to a FPGA routing fabric?

Context:
Mapping Topologies to FPGAs

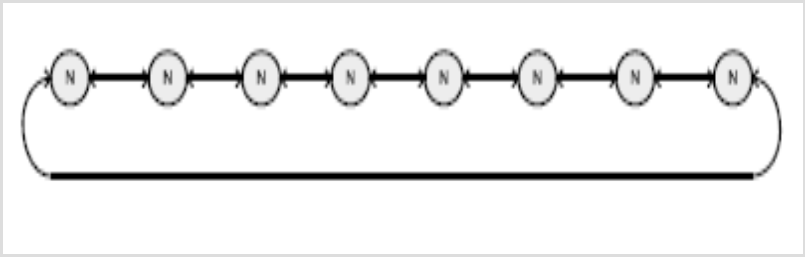

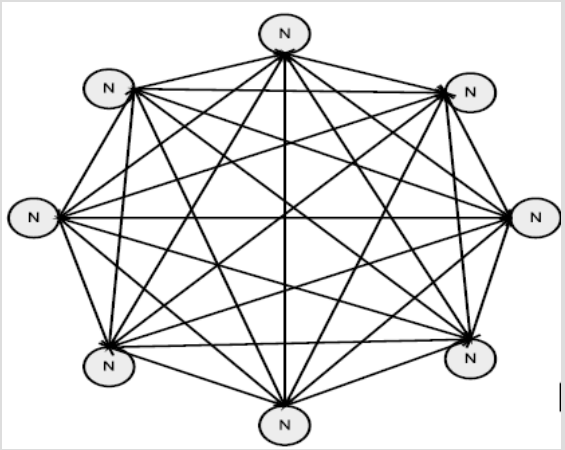

Synthesize
Map
PAR

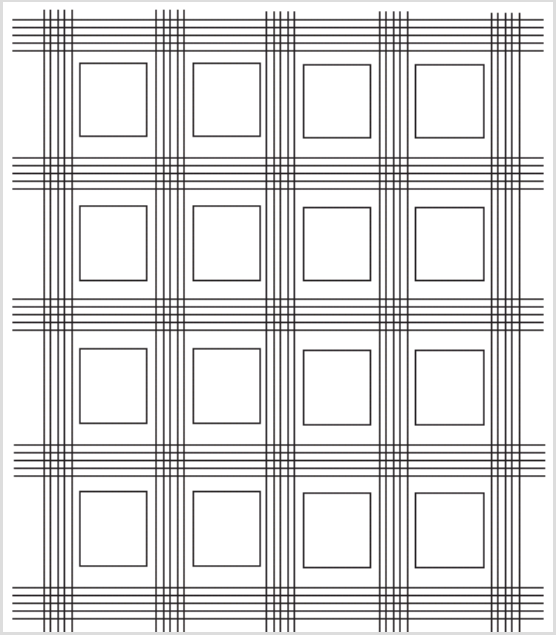

FPGA

Infrastructure: FPGA Structure

Routing channels

Xilinx Virtex 2 and Virtex 4 LX

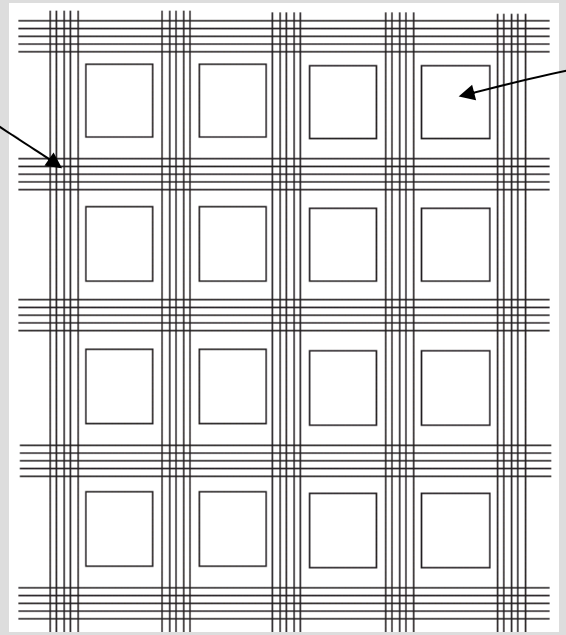

Block of logic

Routing Fabric

Logic Element

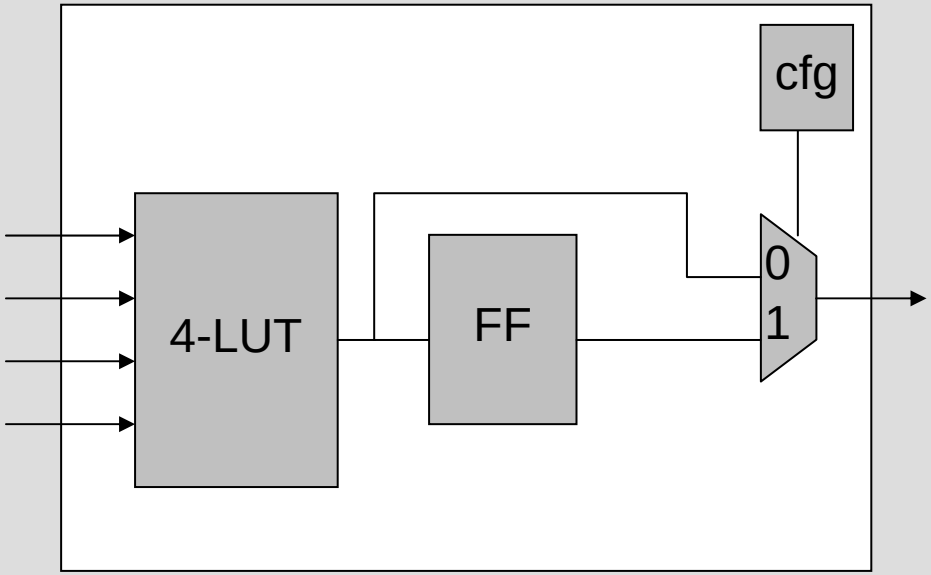

Infrastructure: Topologies

Ring

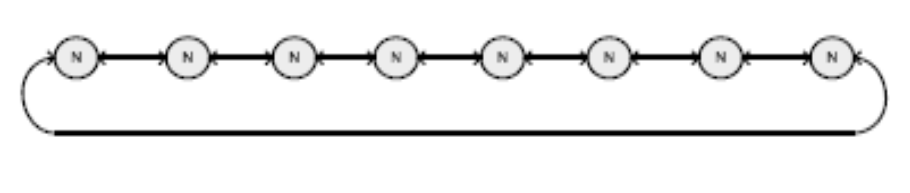

Star

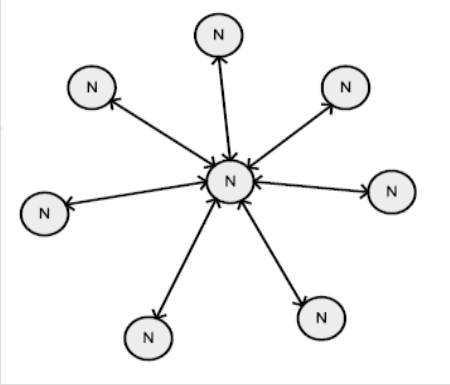

Mesh

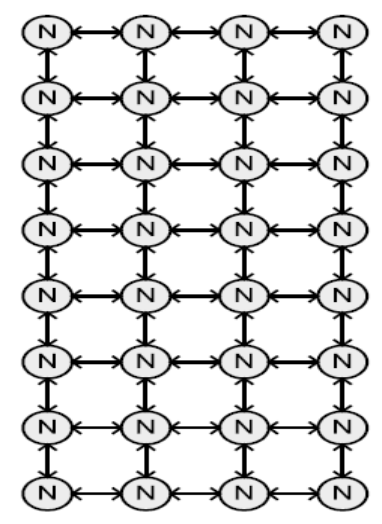

Hypercube

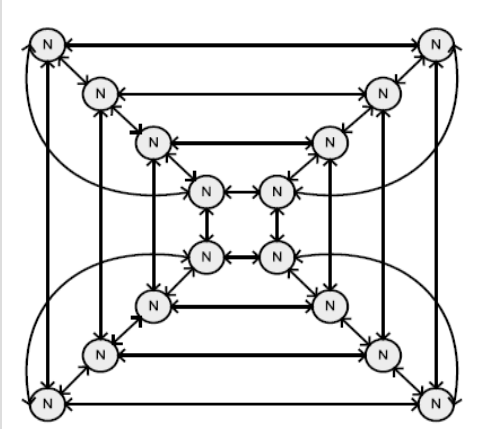

Fully-connected

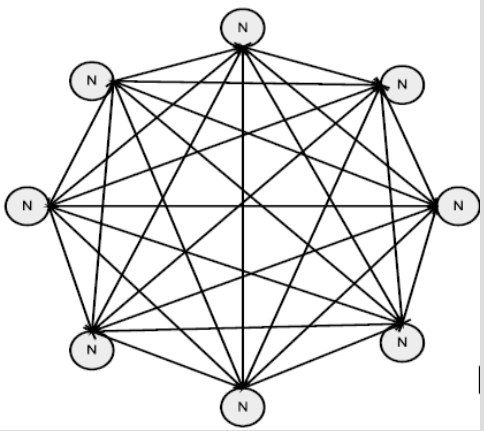

Infrastructure: Nodes

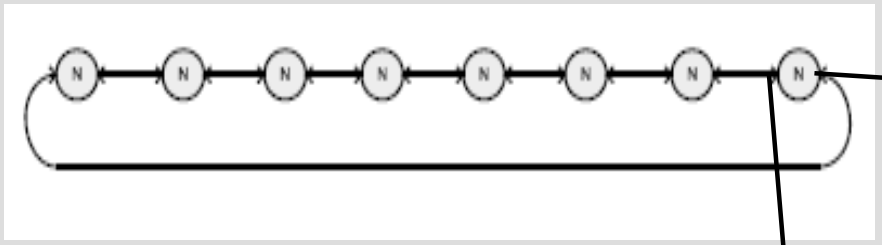

Computing node

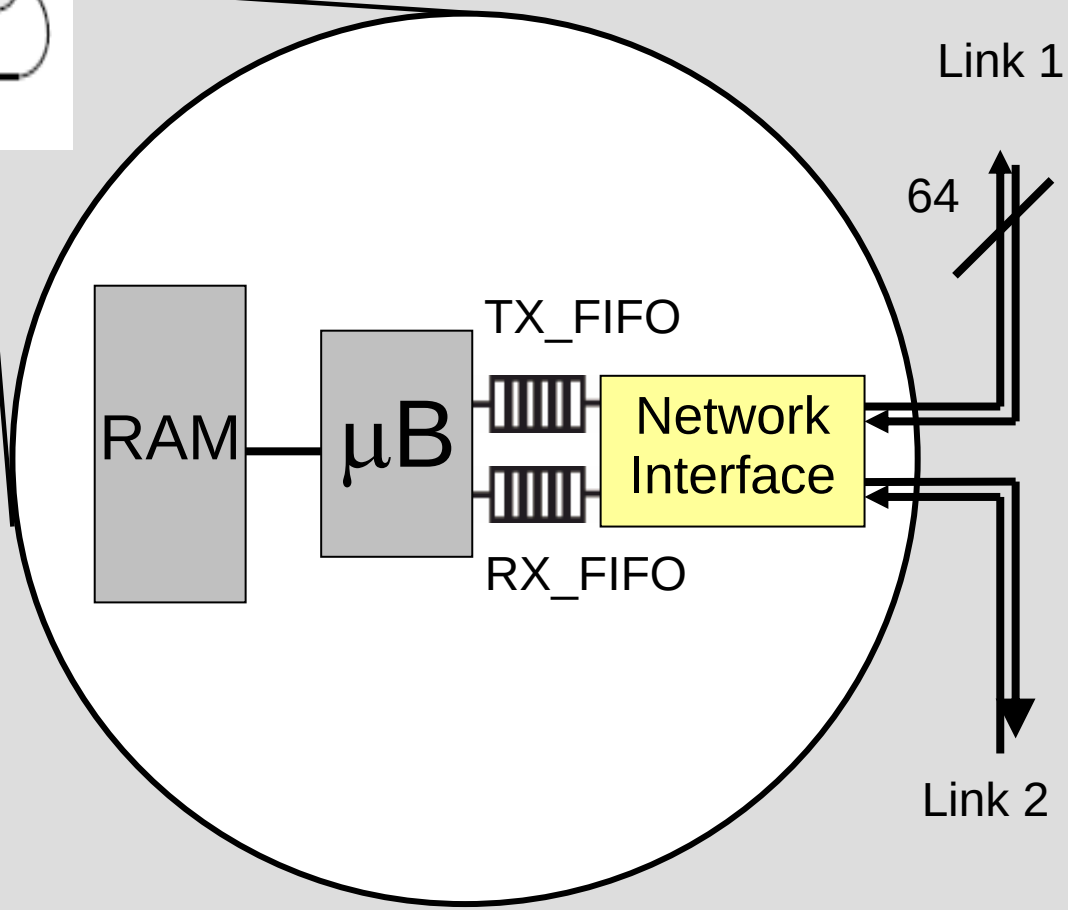

Infrastructure: Tools

Design flow used to implement a MPSoC:

- Avoiding error from interconnecting all the nodes manually
- Experimenting with different topologies by changing the input parameters.

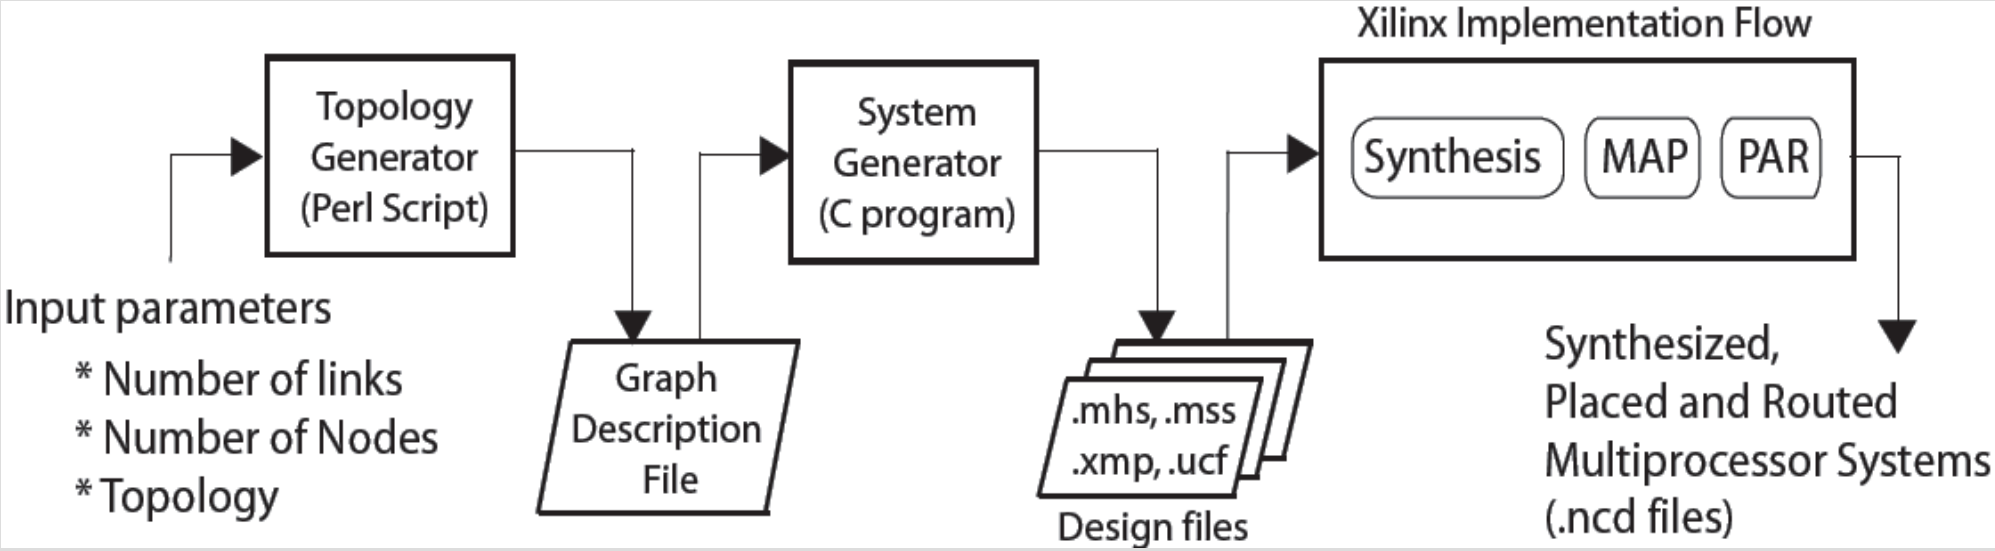

Experiments

- Routability
- Resource utilization and PAR time
- Maximum Frequency
- Cost Metric

Routability Experiment

Objective:

To determine if there are enough wires.

Routability Results

Nodes	Ring	Star	Mesh	Hypercube	Fully-connected
8	Yes	Yes	Yes	Yes	Yes
16	Yes	Yes	Yes	Yes	Yes
32	Yes	Yes	Yes	Yes	**No***

*56 nets were left unrouted!

Resource Utilization and PAR time Experiments

Objectives:

- To measure the logic and routing resources
- To determine the Place and Route (PAR) time

Routing Overhead Results

- Network Nets = Link Nets + Network Interface Nets
- Not too much network overhead

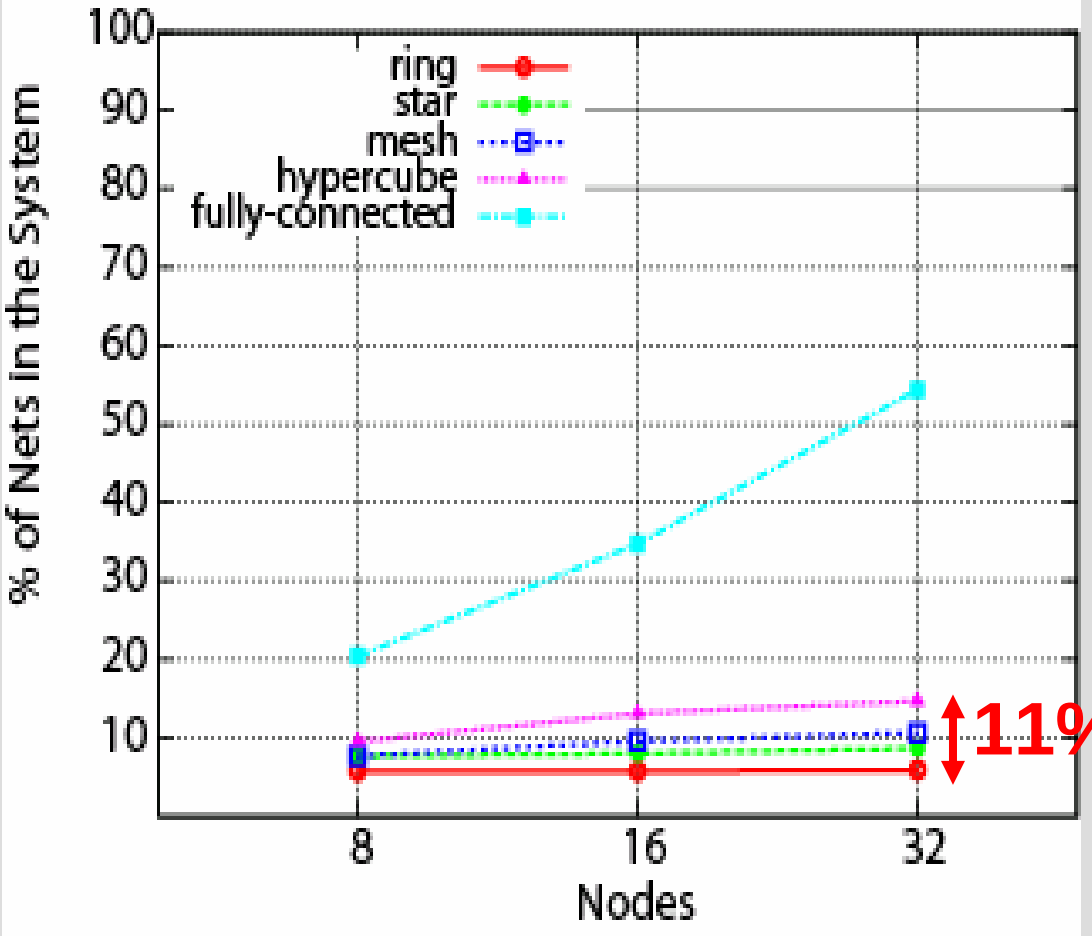

Logic Utilization Results

- Total number of LUTs in the system
- Increase in LUTs is roughly linear, except the fully-connected topology

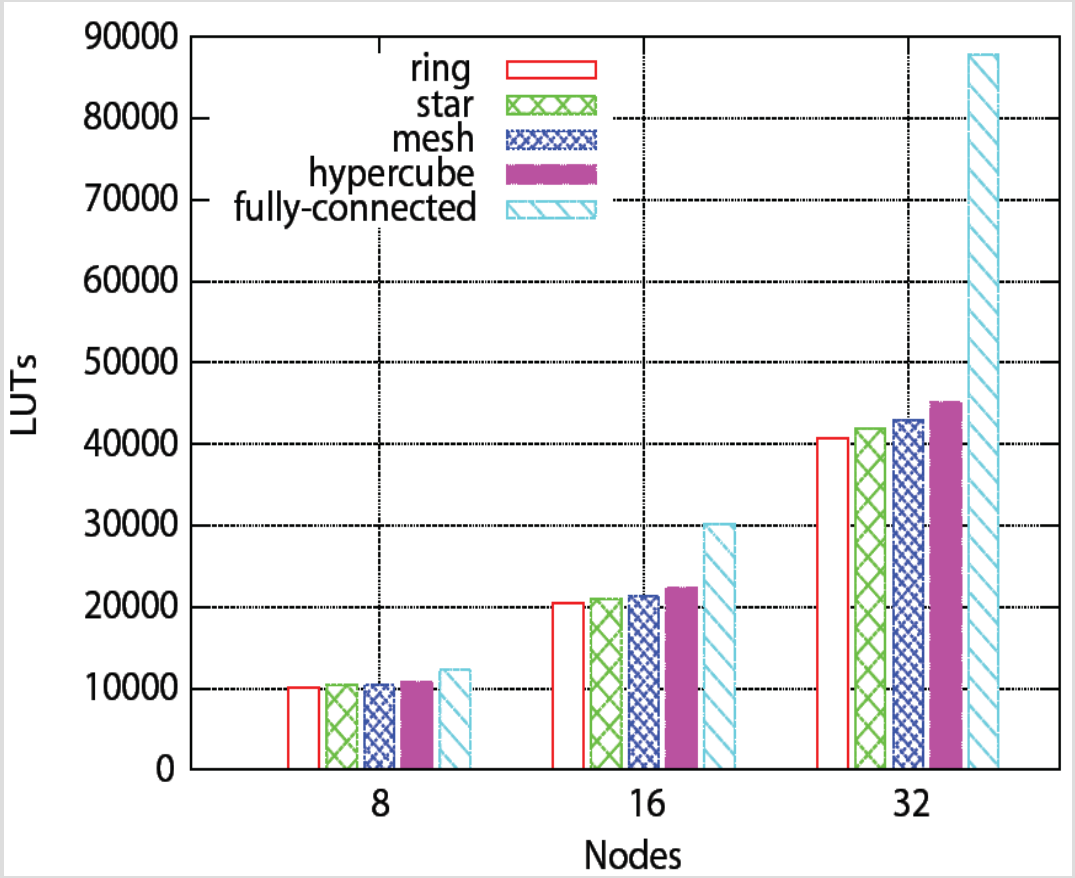

PAR Time Results

Nodes	Place and Route Time	
	Ring, Star, Mesh and Hypercube	Fully-connected
8	12 min.	15 min.
16	30 min.	12 hr.
32	5 hr.	3 days*

* Did not route completely
IBM IntelliStation Pentium 4 HT, 2.8 GHz, 2GB RAM on WindowsXP
Xilinx ISE/EDK 7.1 SP 3

Maximum Frequency Experiments

Objective:

To obtain the maximum clock frequency (using automatic optimizations).

Maximum Frequency Results

- We used the Xilinx Xplorer utility
- MicroBlaze max. freq. 180MHz at speed grade 12
- Congested designs showed a decrease in frequency

Topology	Nodes	Max Freq. (MHz)	Speed Grade	Best Run	Total Runs
fully con.	8	170	12	2	6
ring	16	180	12	1	1
star	16	180	12	4	4
mesh	16	180	12	2	2
hypercube	16	180	12	2	2
fully con.	16	126	12	5	6
ring	32	123	11	5	6

Cost Metric Experiment

Objective:

To derive an experimental metric to compare topologies by relating physical implementation factors and performance factors.

Proposed Cost Metric

- The objective is to minimize the cost metric

$$CM = \frac{D \times Routing/Node \times K^{(f_{target} - freq)} \times LUTs/Node}{Bisection\_width \times Channel\_width \times freq}$$

- Diameter (D) as an indicator of latency
- Penalize the systems that did not meet timing
- Assume single-cycle transmission for bisection bandwidth

Cost Metric Results

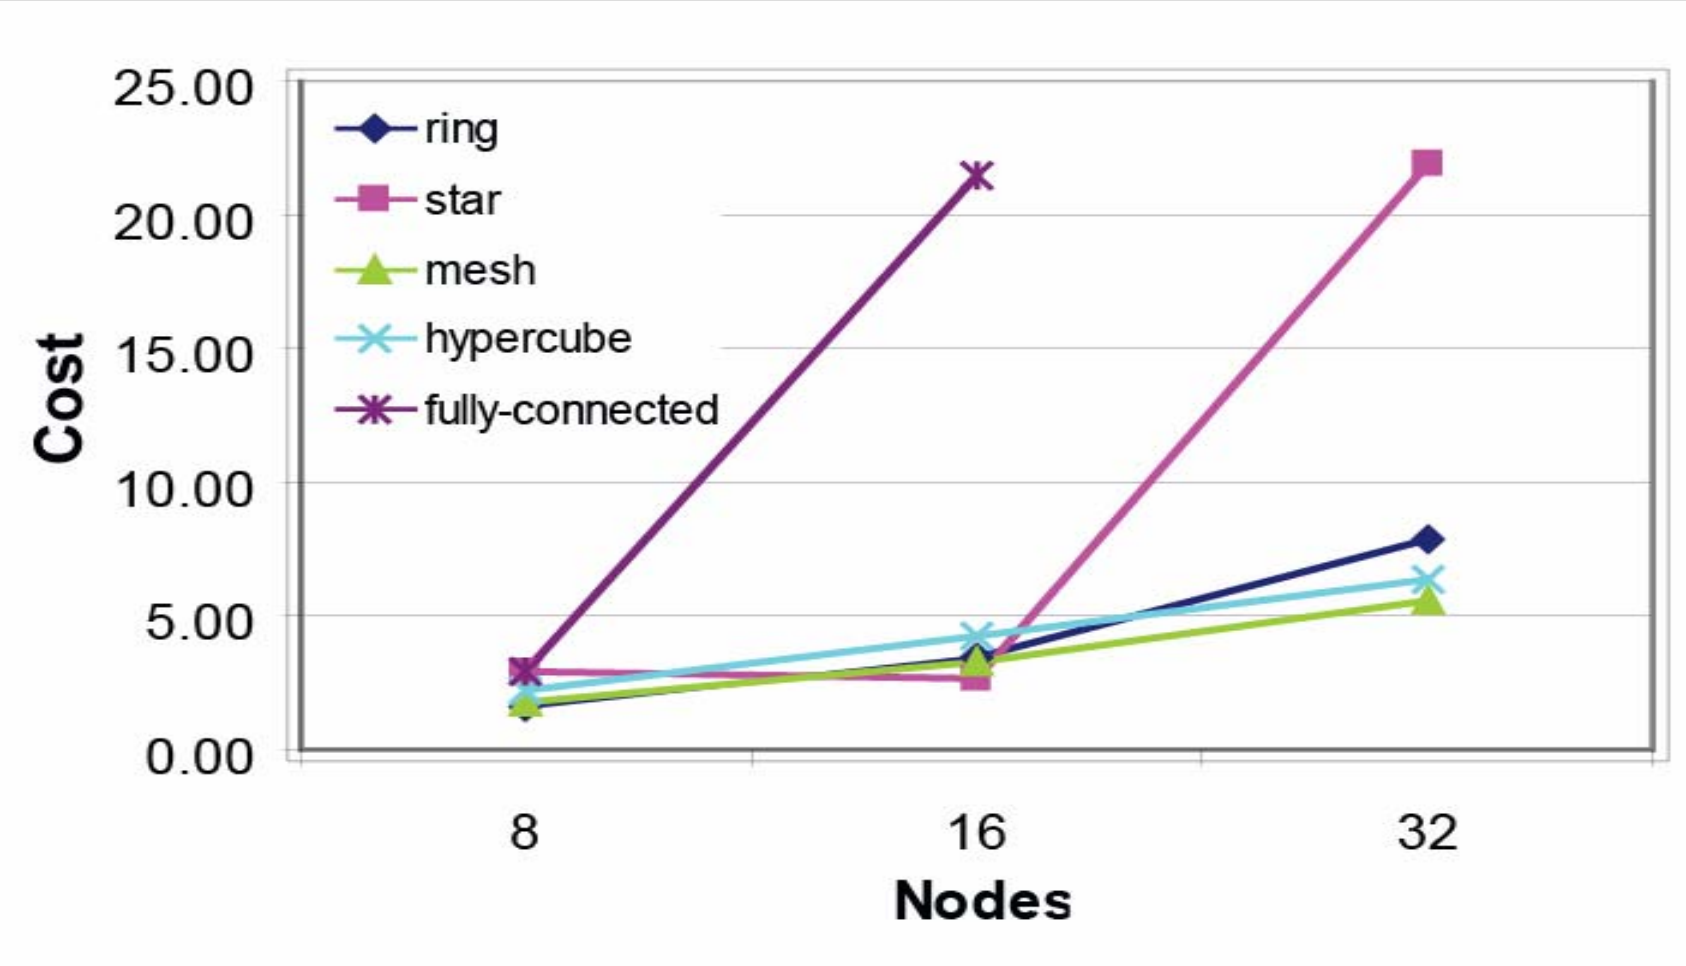

Conclusions

- FPGAs have a fixed set of resources that exist whether they are used or not.
- No need to limit the connectivity if there are resources available.
- For systems up to 16 nodes, fully-connected topologies are practical to implement.
- A 32-node fully-connected topology exceeds the routing resources on the FPGA, but there are still logic resources available to implement more nodes.

Future Work

- Include more topologies (3D-Torus, Trees)
- Predict routability of application-specific topologies

Dankeschön
(Thank you)

Experiments

Minimum Area

- To investigate the minimum area required by each system.

Results: Area

- Emulate resource depletion by restricting available area to PAR.
- RPM and low aspect ratio limit the placement (“brick effect”).
- Placement of μB does not change considerably across same size topologies.
- There are less empty spaces in Virtex2.

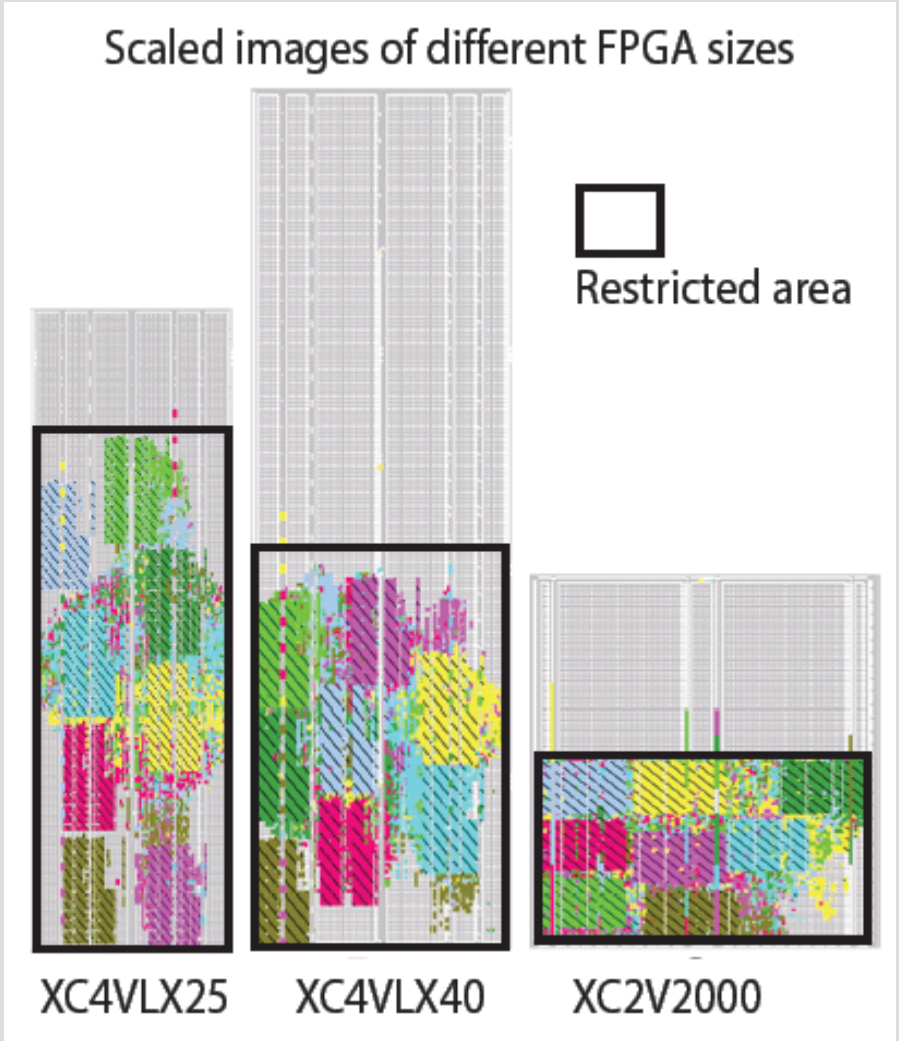

Results: Area

- μB is bigger in Virtex4 than in Virtex2.
- Topologies with more connectivity benefits from empty spaces.
- The same minimum area is required for a ring and a fully-connected.

8-node Ring MPSoC

Topology	XC2V2000		XC4VLX25		XC4VLX40	
	Min. Area Req'd (slices)	Area Utiliz. (%)	Min. Area Req'd (slices)	Area Utiliz. (%)	Min. Area Req'd (slices)	Area Utiliz. (%)
ring	5376	92.5	9352	62.1	8208	70.7
star	5568	90.8	8736	67.5	8640	68.2
mesh	5376	94.4	9240	64.0	8280	73.5
hypercube	6336	81.9	8568	70.4	8280	72.9
fully con.	6528	92.5	9688	71.2	8208	84.1
Average	5836	90.4	9116	67.0	8323	73.9

FPGA's Aspect Ratio

Chip	X_max	Y_max	Slices	Aspect Ratio
XC2V2000	95	111	10752	0.855856
XC4V25LX	55	191	10752	0.287958
XC4V40LX	71	255	18432	0.278431
XC4V60LX	103	255	26624	0.403922
XC4V200LX	231	383	89088	0.603133