Designing Nanophotonic Interconnection Networks

Christopher Batten

Computer Systems Laboratory
Cornell University

Workshop on the Interaction between
Nanophotonic Devices and Systems
December 2010

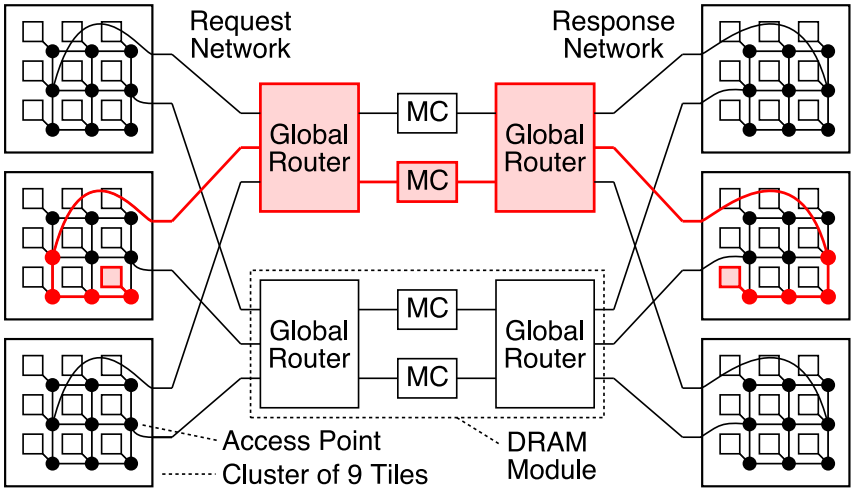

Processor-to-DRAM Network [IEEE Micro'09]

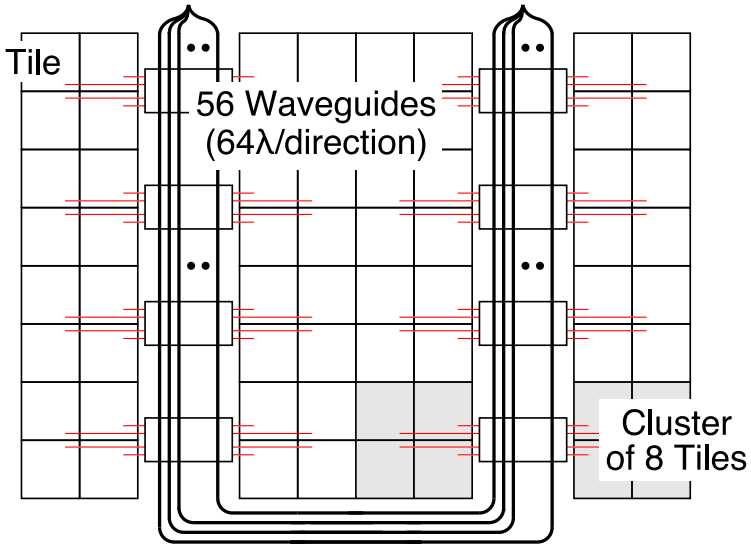

On-Chip Clos Network [NOCS'09]

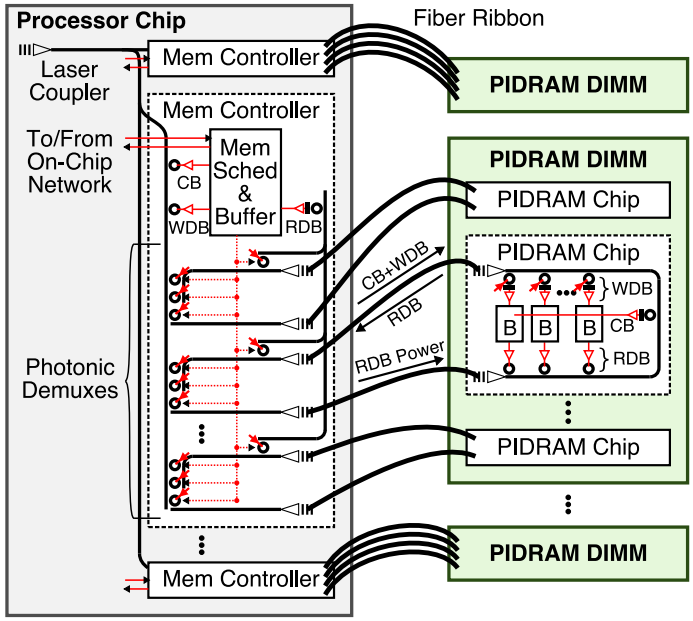

DRAM Memory Channel [ISCA'10]

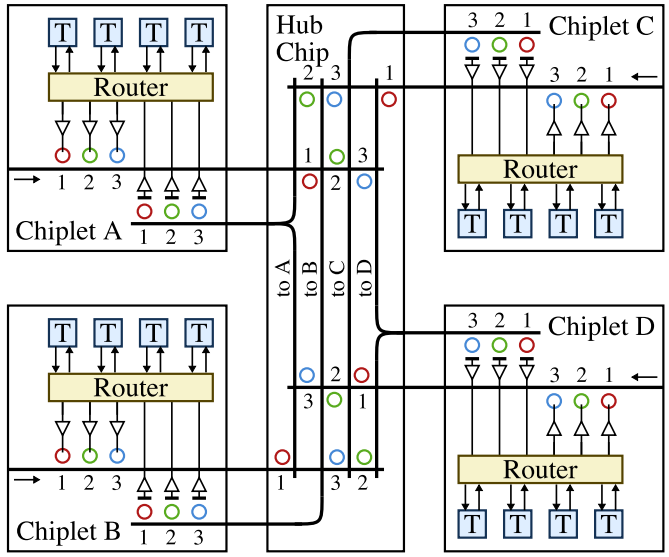

System-in-Package [WINDS'10]

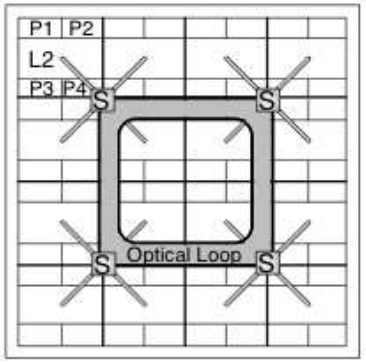

Kirman, MICRO'06

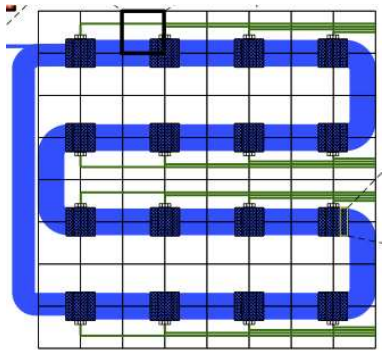

Vantrease, ISCA'08

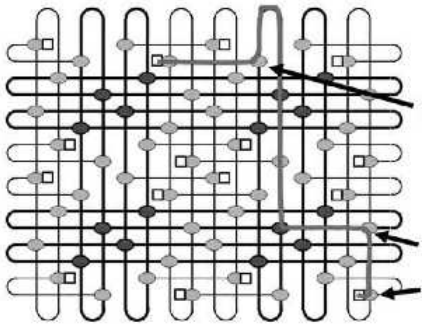

Shacham, TOC'08

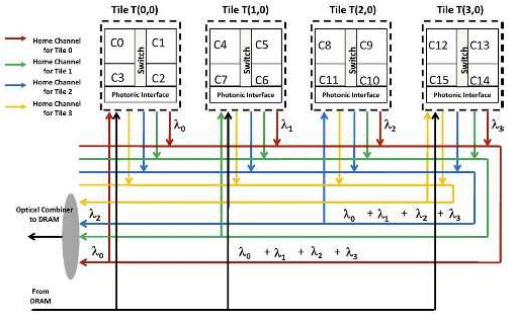

Morris, JSTQE'10

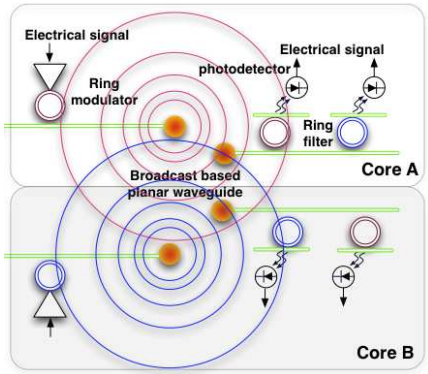

Li, DAC'09

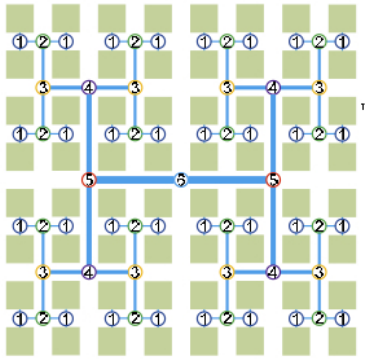

Gu, DATE'09

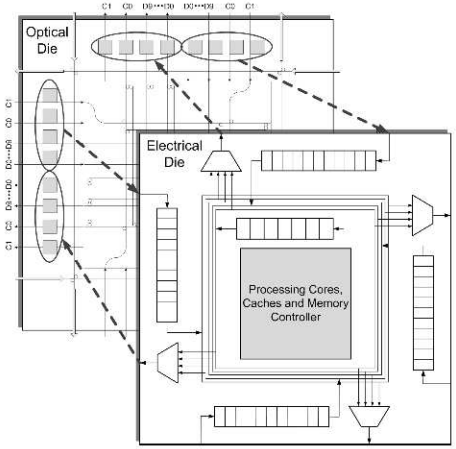

Cianchetti, ISCA'09

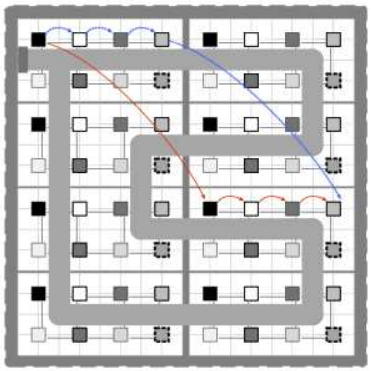

Pan, ISCA'09

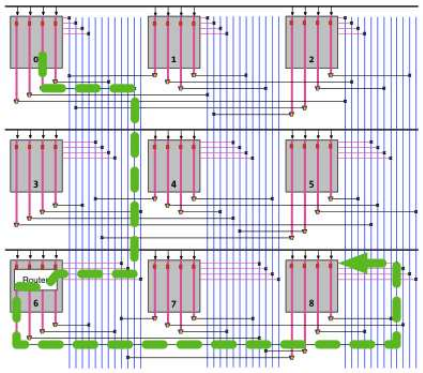

Koka, ISCA'10

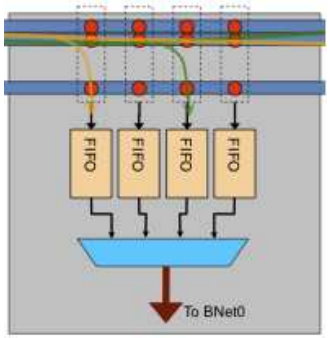

Kurian, PACT'10

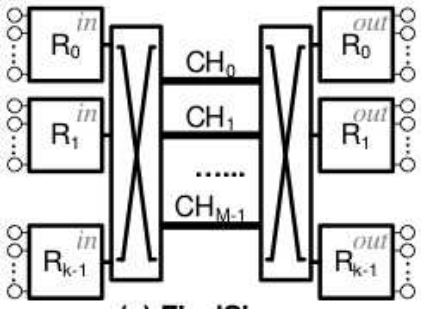

Pan, HPCA'10

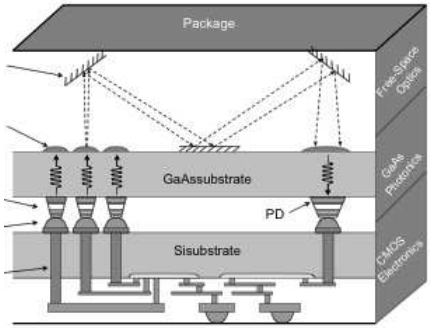

Xue, ISCA'10

Based on experiences designing nanophotonic interconnection networks and surveying the literature, can we begin to identify common design patterns and guidelines?

Designing Nanophotonic Interconnection Networks

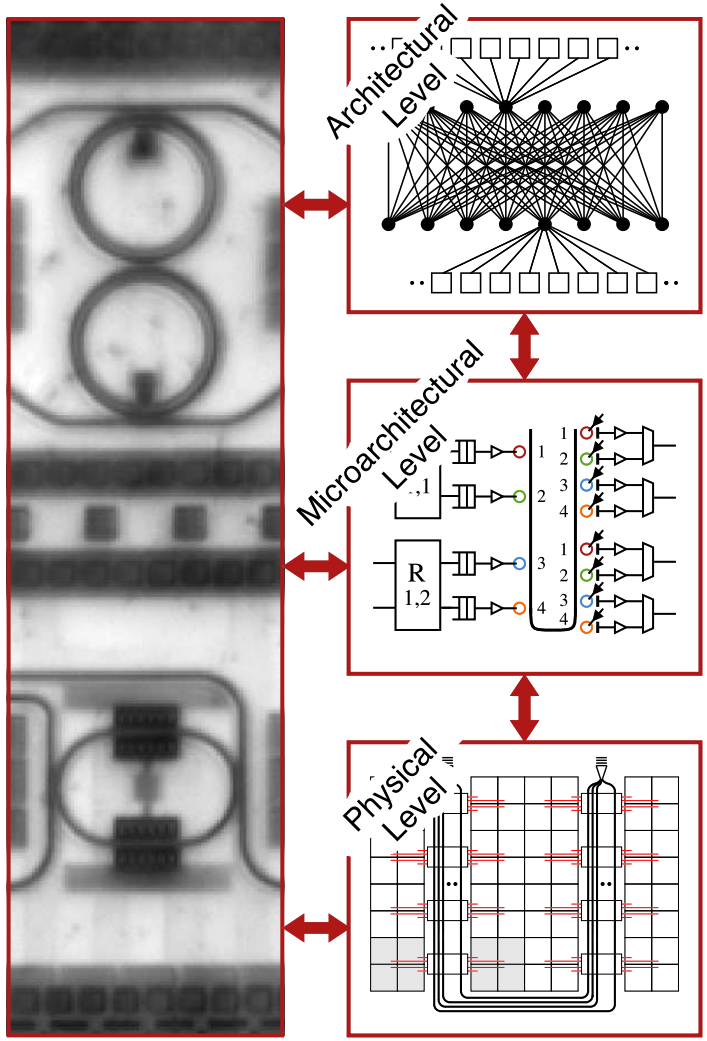

Motivation

Architectural-Level Design

Microarchitectural-Level Design

Physical-Level Design

Design Guidelines

Architectural-Level Design

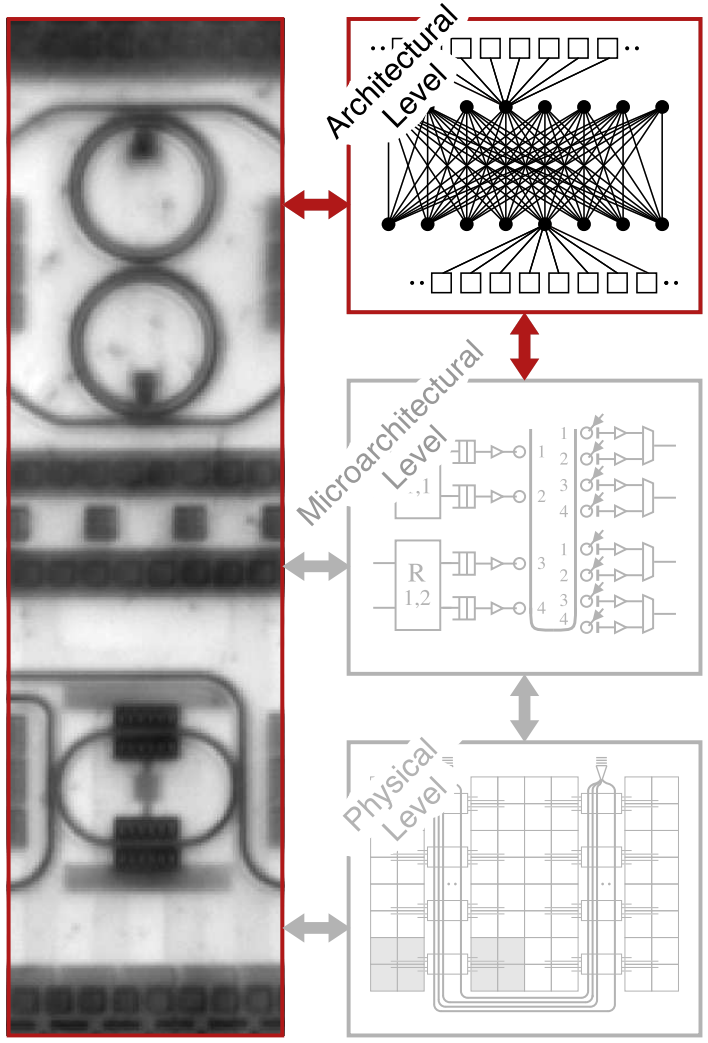

- Select logical topology which consists of input and output terminals interconnected by buses, routers, and channels
- Use topology diagram to capture design decisions
- Determine bus/channel bandwidths based on app demands
- Preliminary exploration of routing algorithms
- First-order analysis of various options to narrow design space
- Design electrical baseline network

Logical Network Topologies

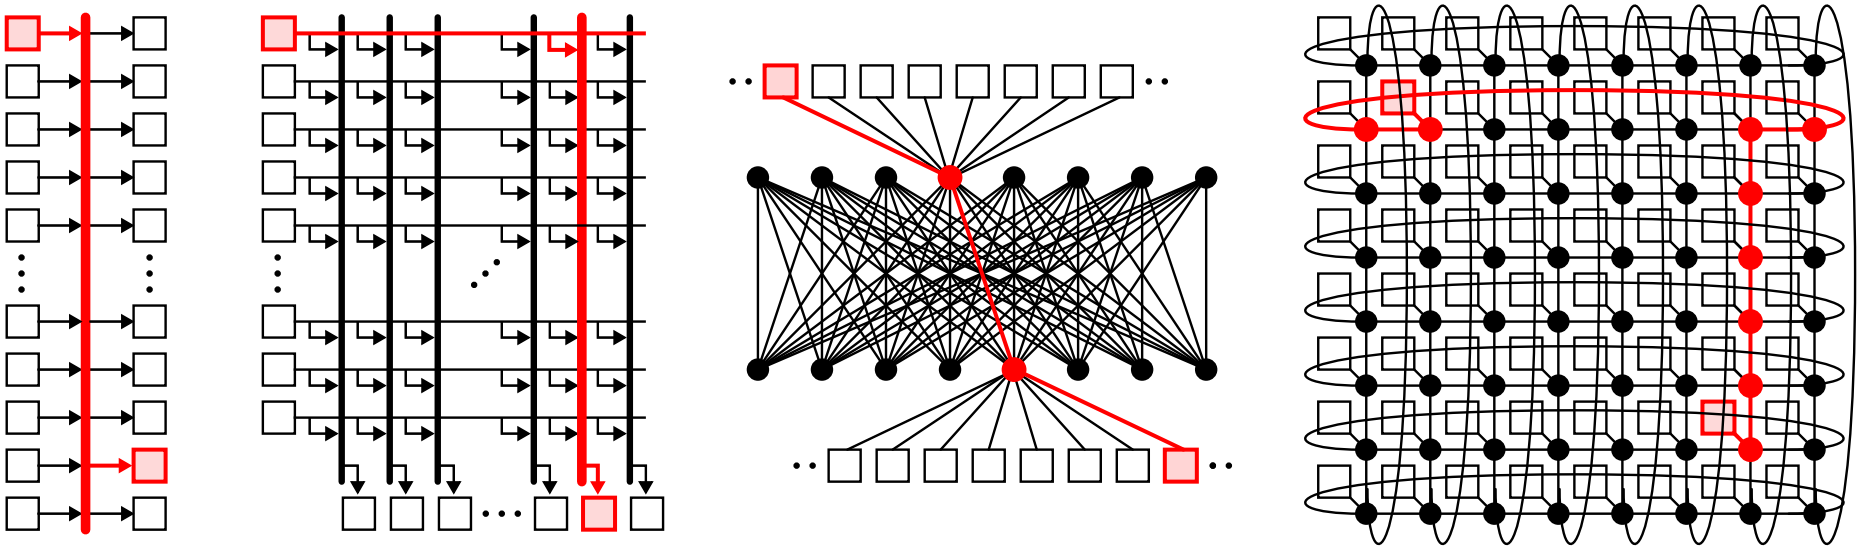

Increasing number of stages

Logical Network Topologies: Bus

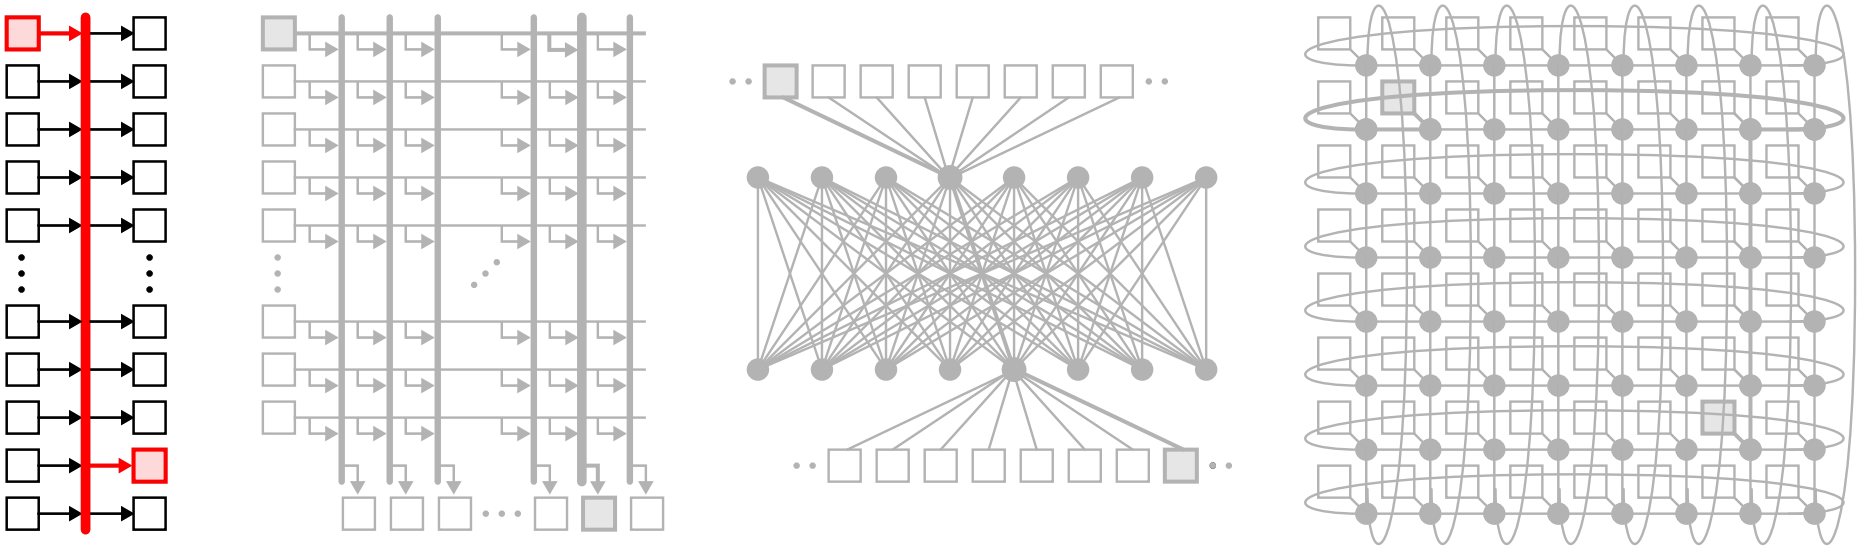

- Input terminals globally arbitrate for single shared medium
- Benefits: Simple, single stage, serialize messages, broadcast
- Challenges: Global arbitration and data transfer

Logical Network Topologies: Crossbar

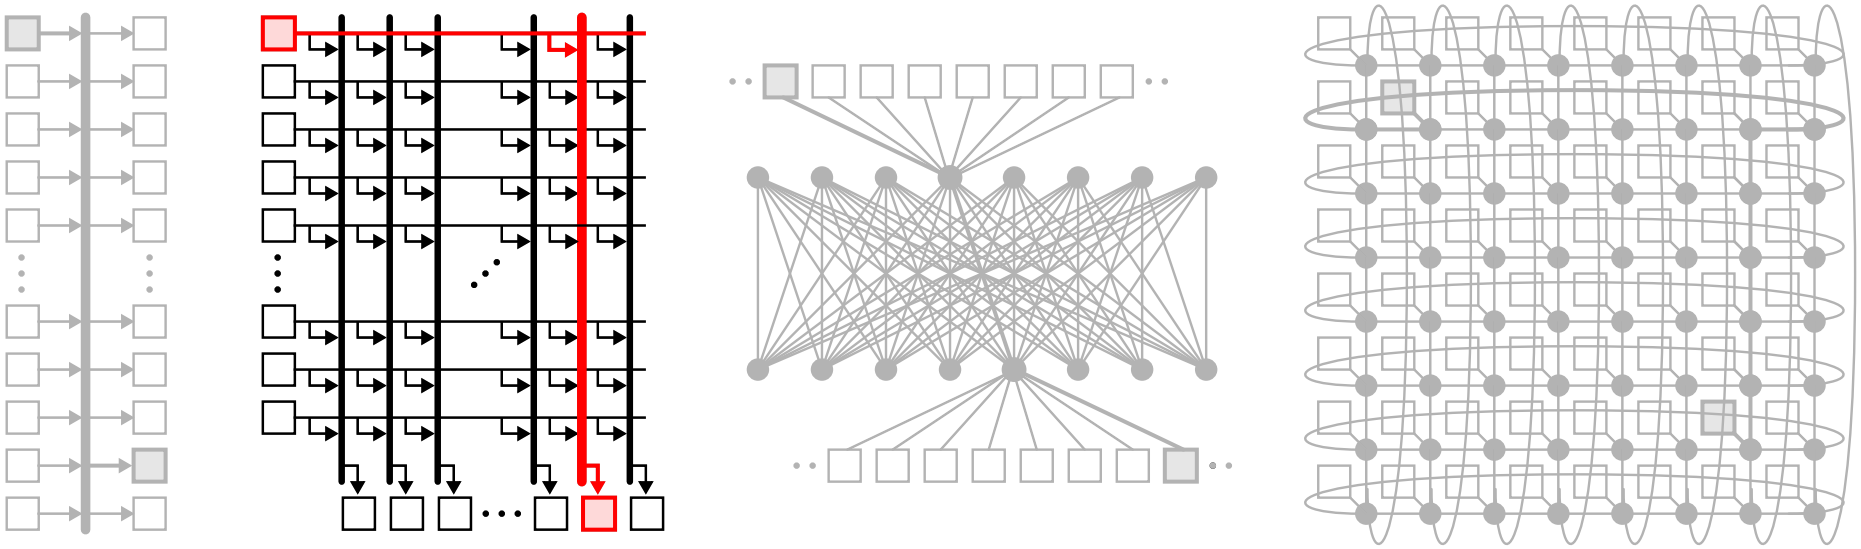

- Many parallel buses, one per terminal
- Benefits: Single stage, high throughput
- Challenges: Global arbitration and data transfer

Logical Network Topologies: Butterfly

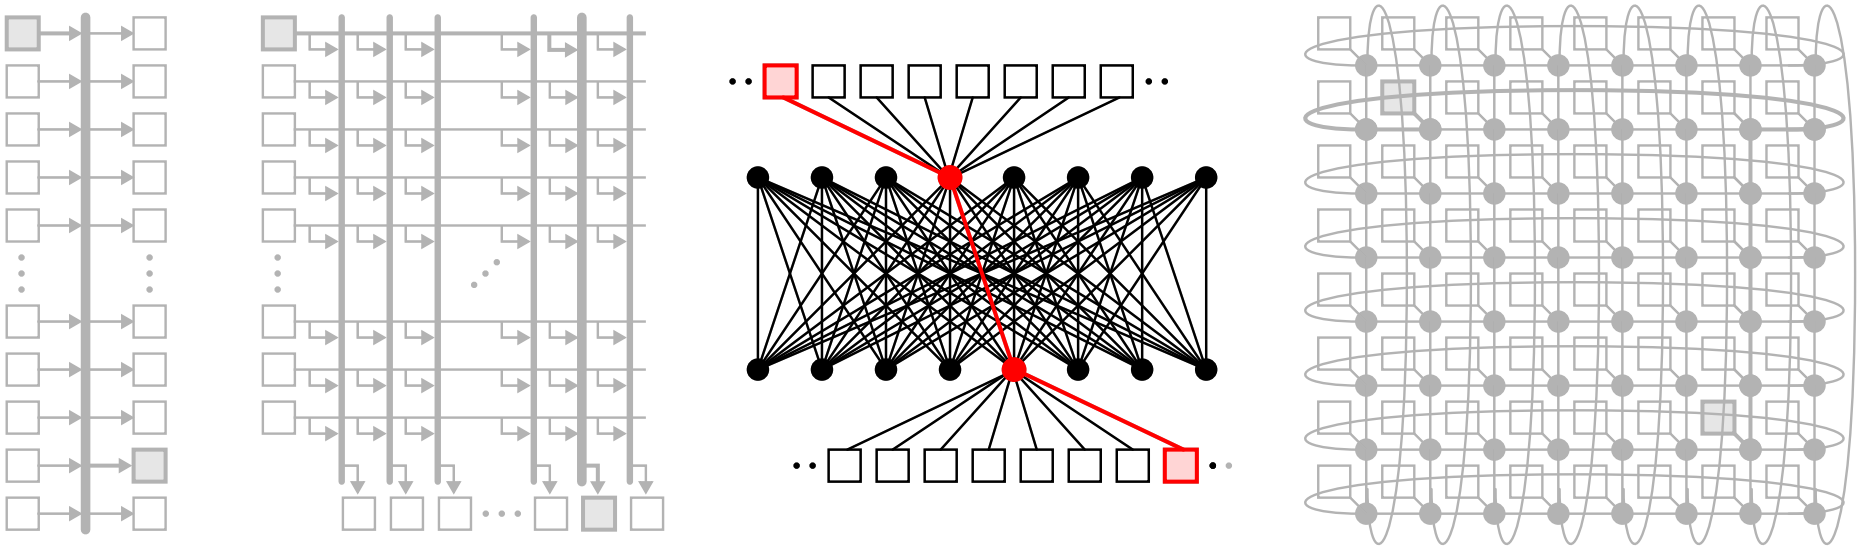

- Multiple-stages arranged in butterfly pattern
- Benefits: Distributed routing, arbitration, flow-control
- Challenges: Multiple stages of switching, global channels

Logical Network Topologies: Torus

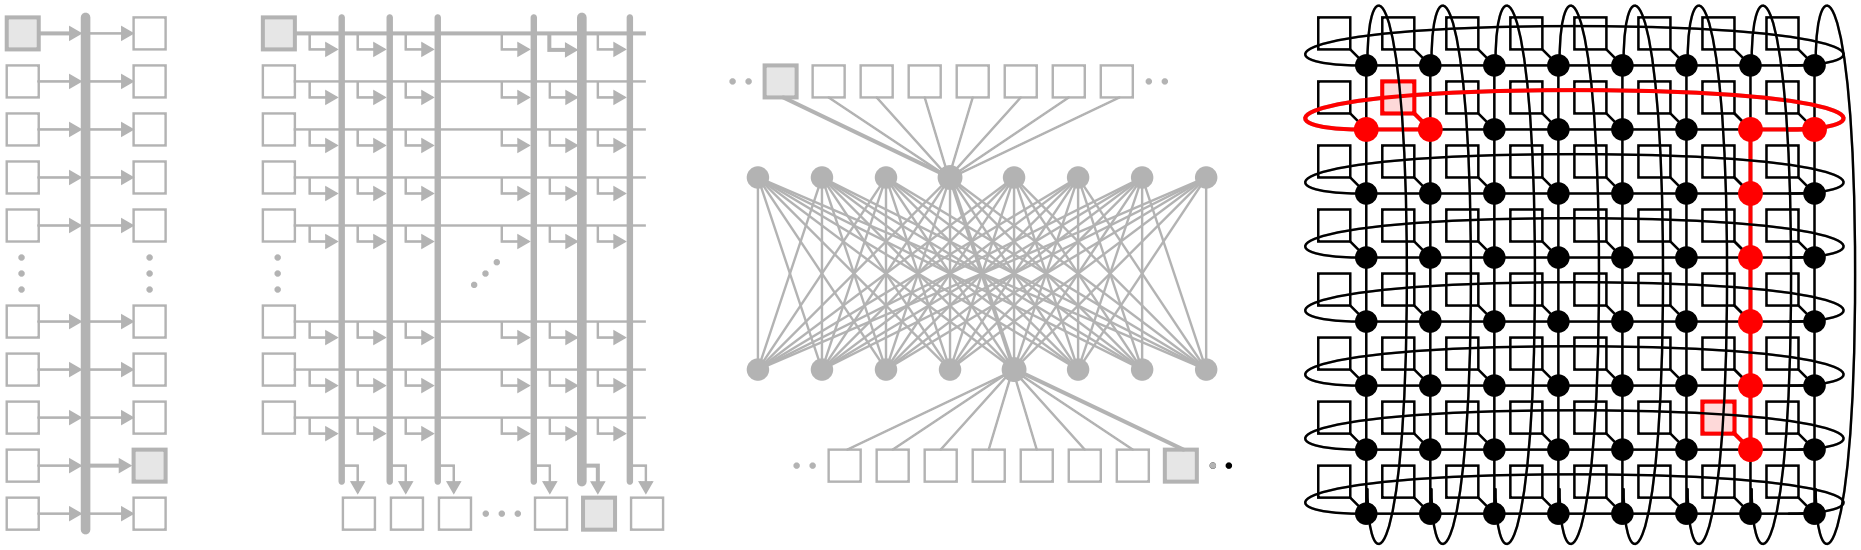

- Multiple-stages arranged in grid pattern possibly with wrap-around
- Benefits: Small localized routers, short channels (for low-dim)
- Challenges: Many stages of switching (for low-dim)

Example Architectural-Level Analysis

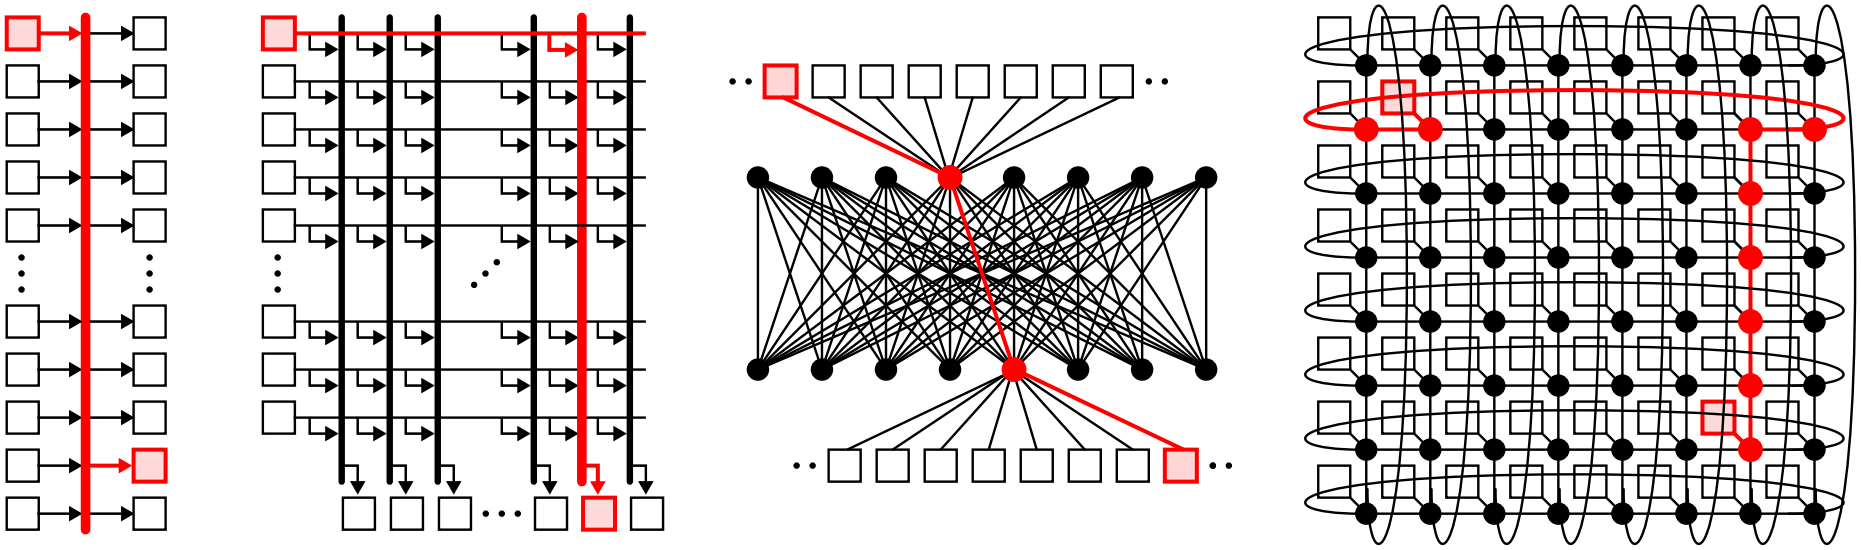

Topology		Buses & Channels				Routers		Latency				
		N_C	N_{BC}	b_C	$N_{BC} \cdot b_C$	N_R	radix	H_R	T_R	T_C	T_S	T_0
Crossbar	64×64	64	64	128	8,192	1	64×64	1	10	n/a	4	14
Butterfly	8-ary 2-stage	64	32	128	4,096	16	8×8	2	2	2-10	4	10-18
Clos	(8,8,8)	128	64	128	8,192	24	8×8	3	2	2-10	4	14-32
Torus	8-ary 2-dim	256	32	128	4,096	64	5×5	2-9	2	2	4	10-38
Mesh	8-ary 2-dim	224	16	256	4,096	64	5×5	2-15	2	1	2	7-46
CMesh	4-ary 2-dim	48	8	512	4,096	16	8×8	1-7	2	2	1	3-25

Categorizing Previous Proposals

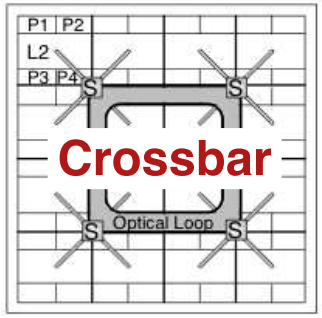

Kirman, MICRO'06

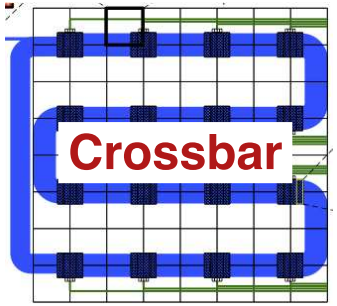

Vantrease, ISCA'08

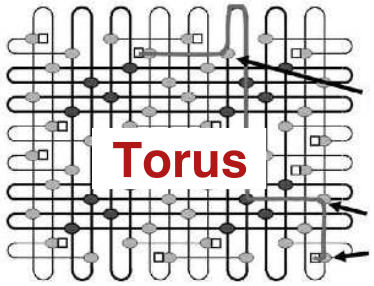

Shacham, TOC'08

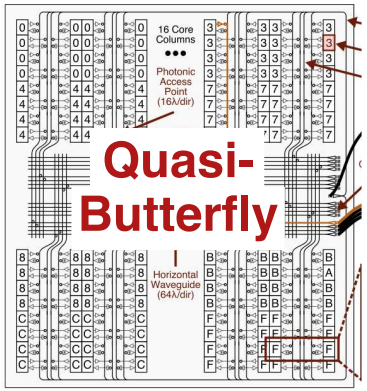

Batten, IEEE Micro'09

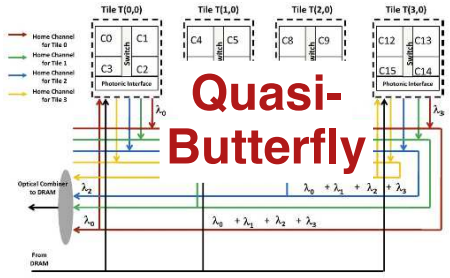

Morris, JSTQE'10

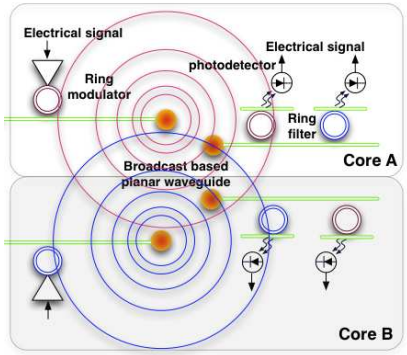

Li, DAC'09

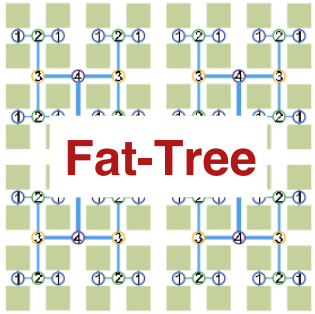

Gu, DATE'09

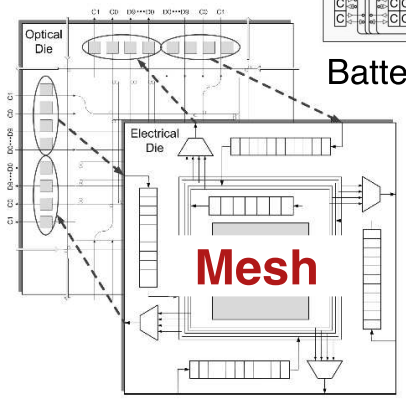

Cianchetti, ISCA'09

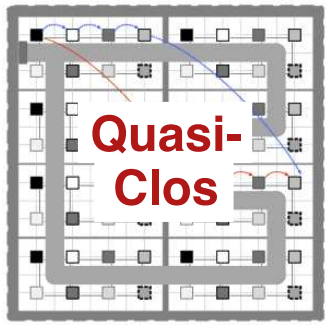

Pan, ISCA'09

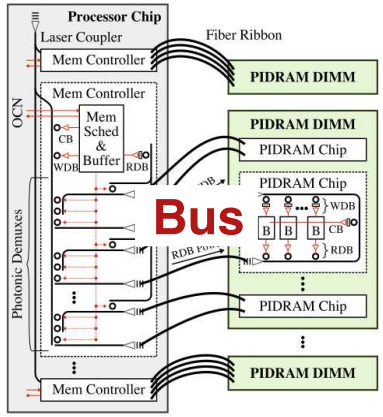

Beamer, ISCA'10

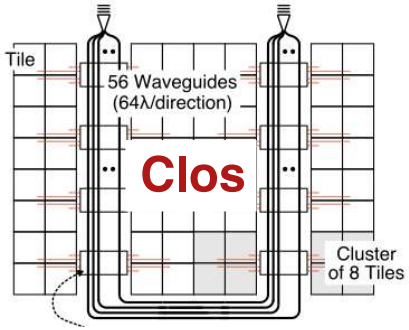

Joshi, NOCS'09

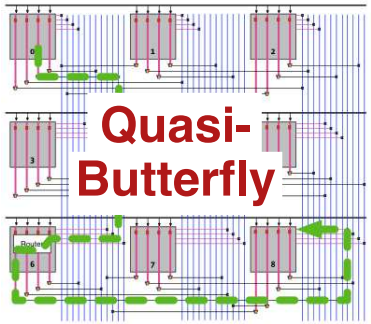

Koka, ISCA'10

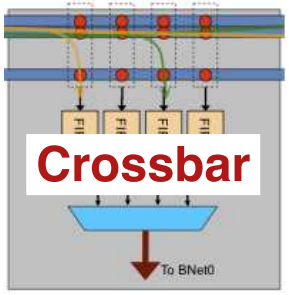

Kurian, PACT'10

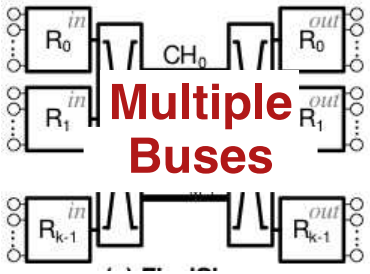

Pan, HPCA'10

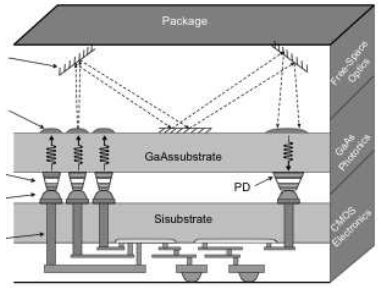

Xue, ISCA'10

Microarchitectural-Level Design

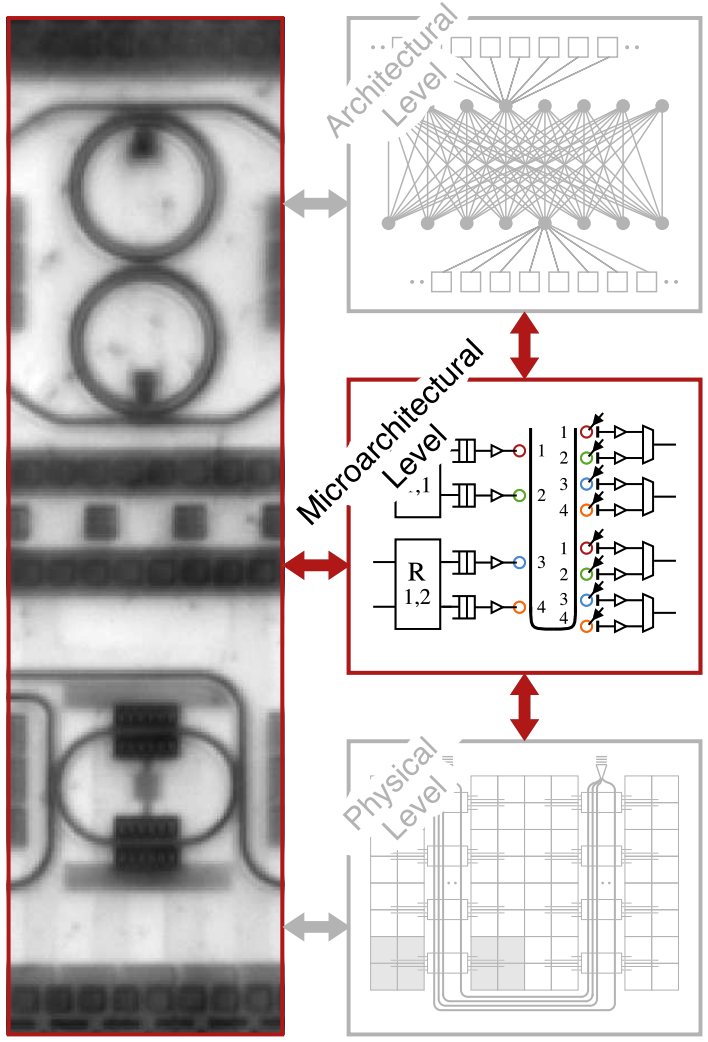

- Choose which buses, routers, and channels to implement electrically and which to implement optically
- Use nanophotonic schematic to capture design decisions
- Decide where to use transmitters, receivers, active filters
- Decide arbitration for wavelengths, manage electrical buffering
- Finalize routing algorithm
- Ignore wavelength to waveguide mapping and waveguide layout

Schematic and Layout Symbols

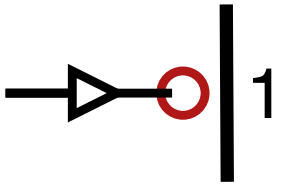

Transmitter

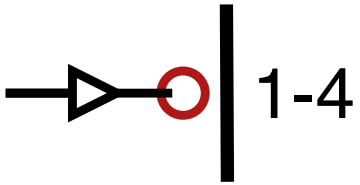

Multiple Transmitters

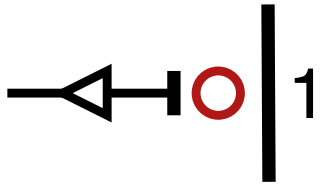

Receiver

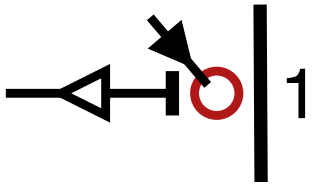

Tunable Receiver

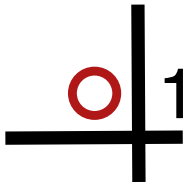

Passive Filter

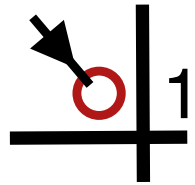

Active Filter

Bus Microarchitectural Design Patterns

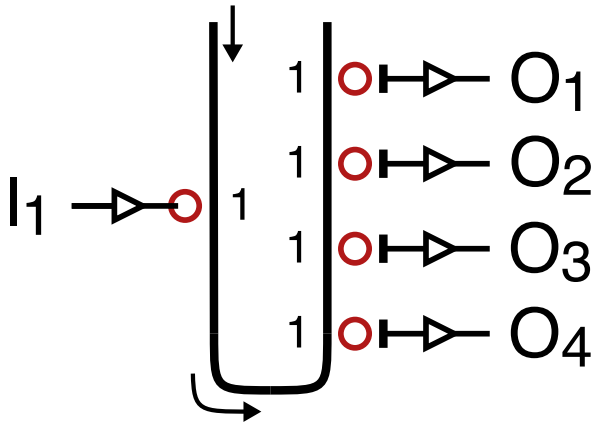

Single-Writer Broadcast-Reader

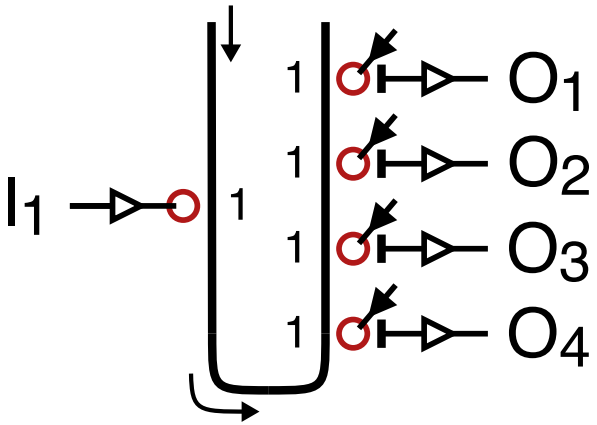

Single-Writer Multiple-Reader

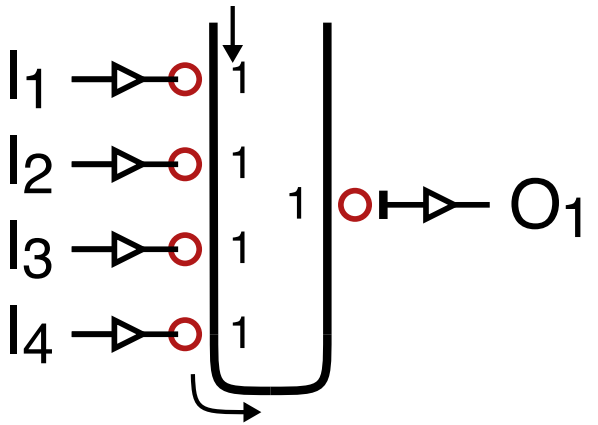

Multiple-Writer Single-Reader

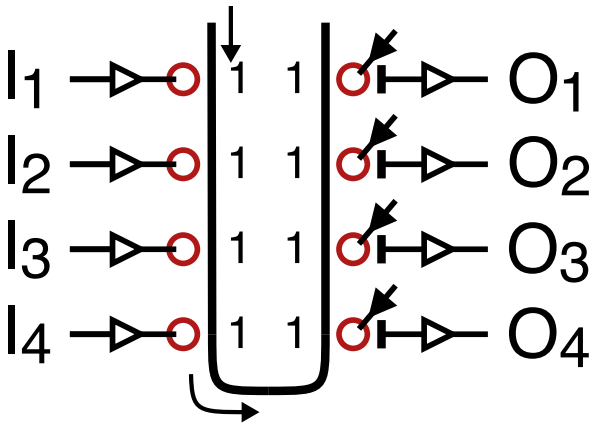

Multiple-Writer Multiple-Reader

Crossbar Microarchitectural Design Patterns

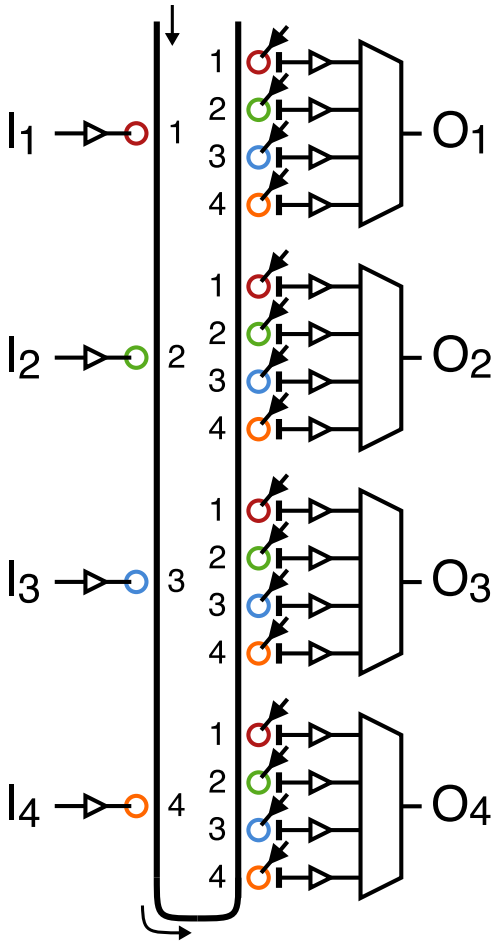

Single-Writer
Multiple-Reader

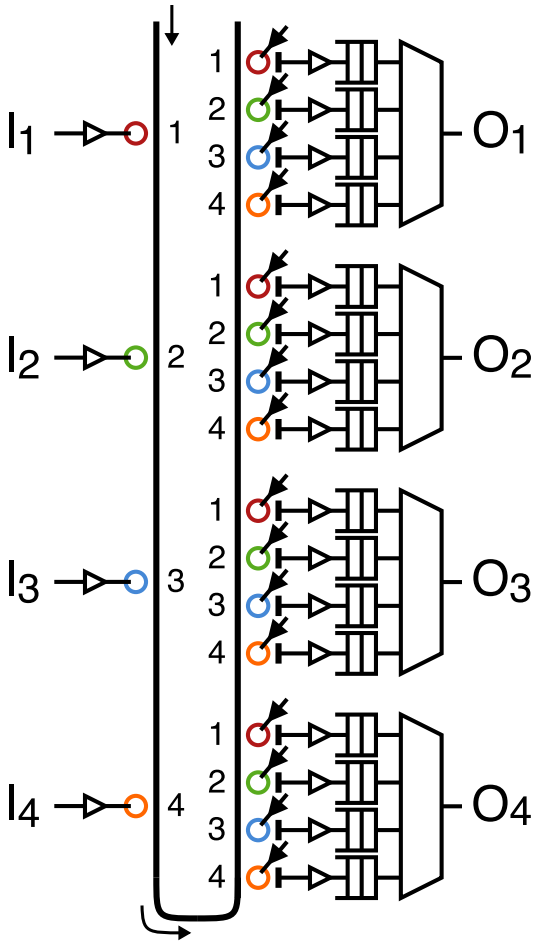

Buffered Single-Writer
Multiple-Reader

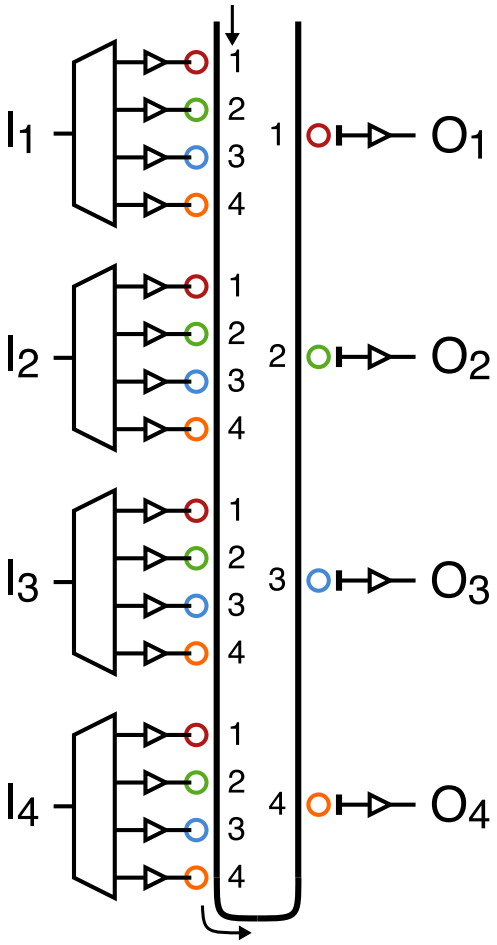

Multiple-Writer
Single-Reader

Butterfly Microarchitectural Design Patterns

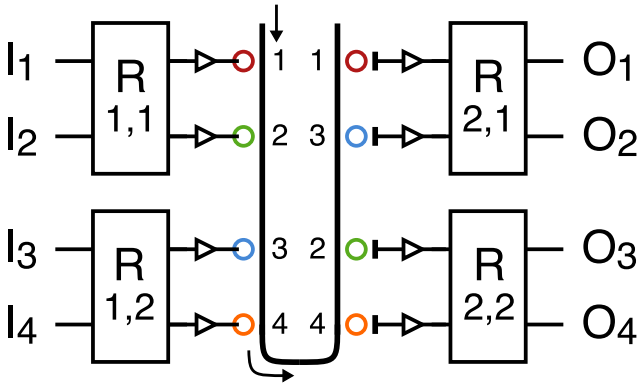

Nanophotonic Channels

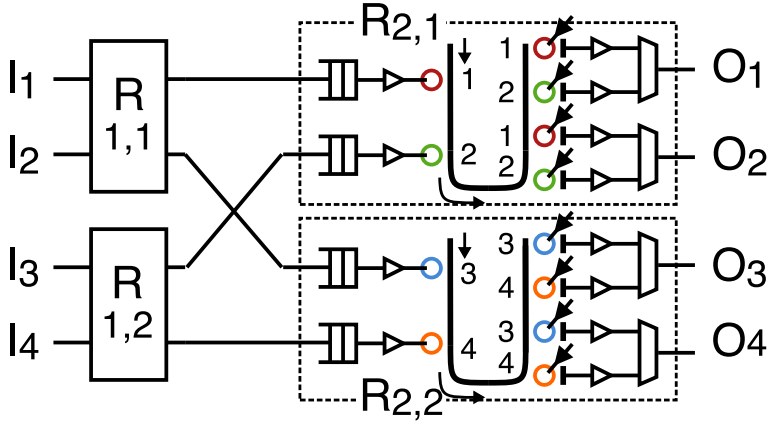

Nanophotonic Second-Stage Routers

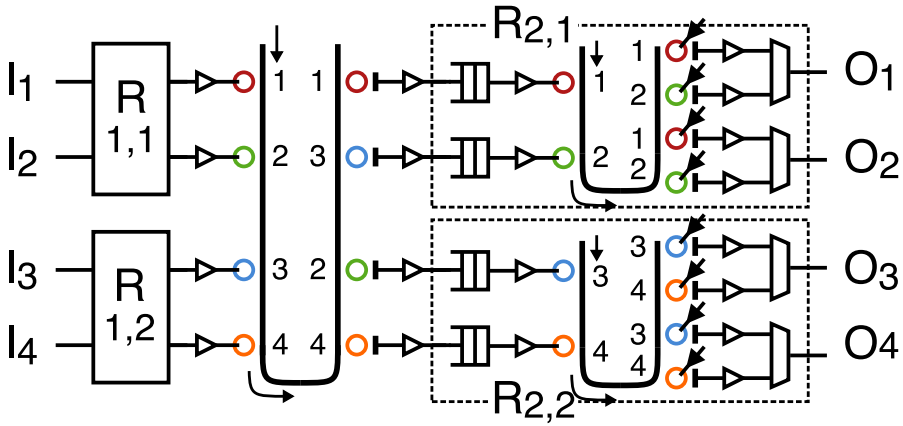

Nanophotonic Channels and Second-Stage Routers

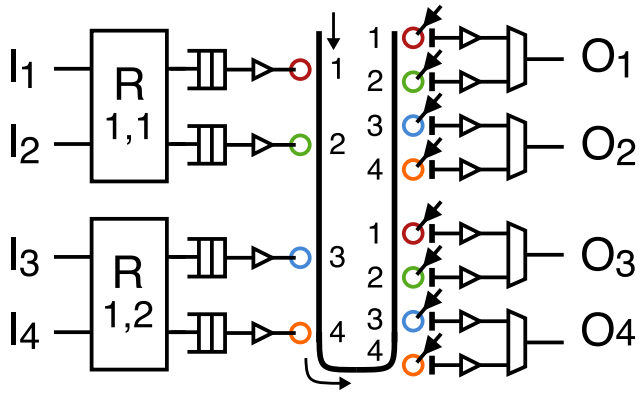

Unified Nanophotonic Channels and Second-Stage Routers

Torus Microarchitectural Design Patterns

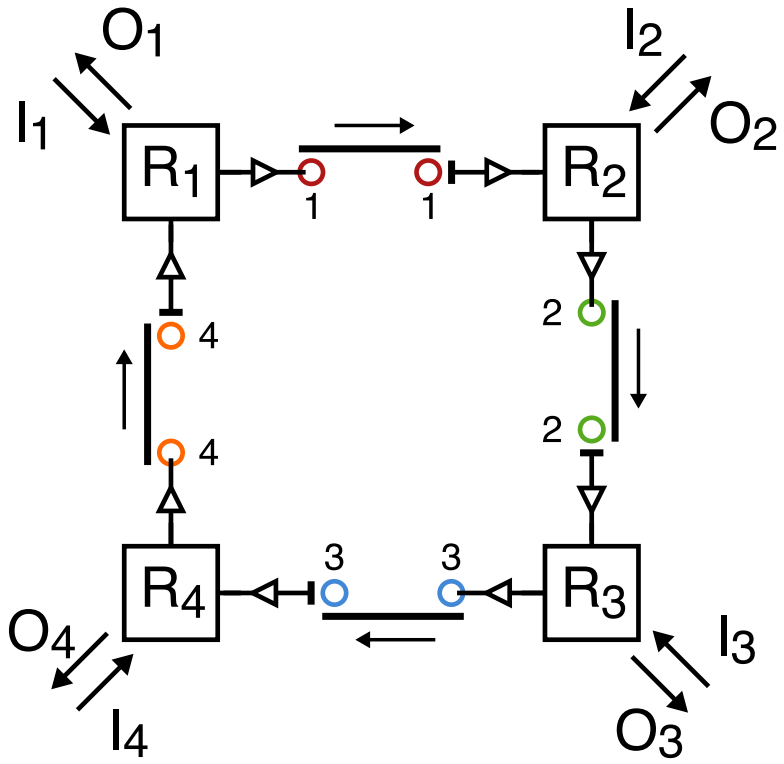

Nanophotonic Channels

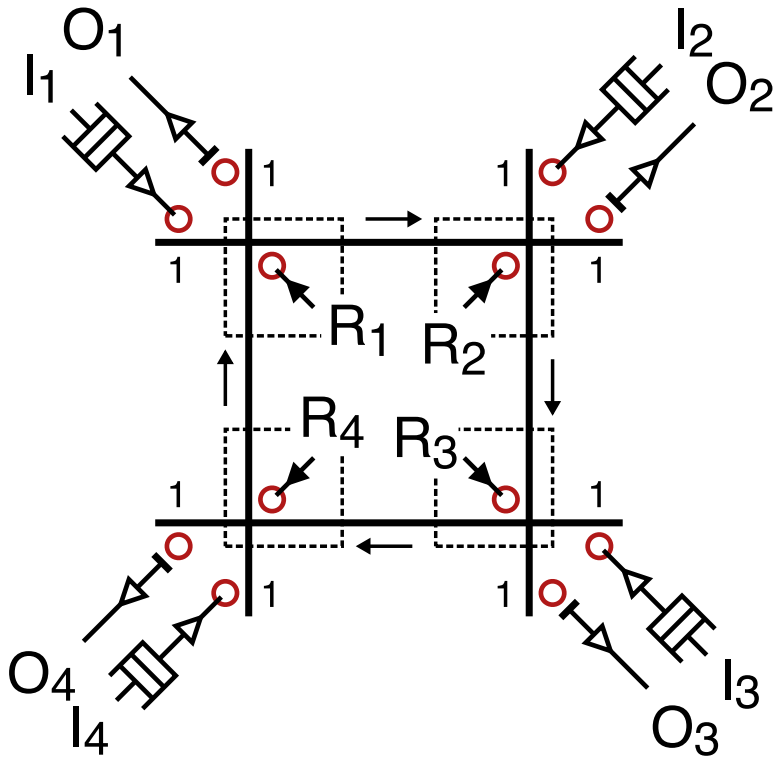

Nanophotonic Channels and Routers

Categorizing Previous Proposals

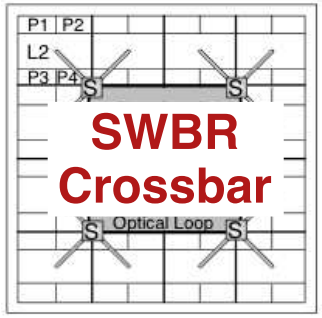

Kirman, MICRO'06

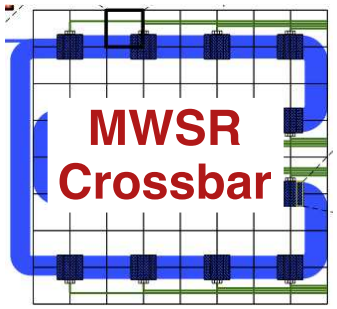

Vantrease, ISCA'08

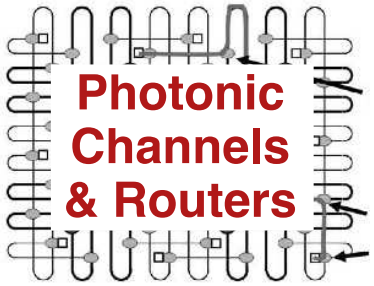

Shacham, TOC'08

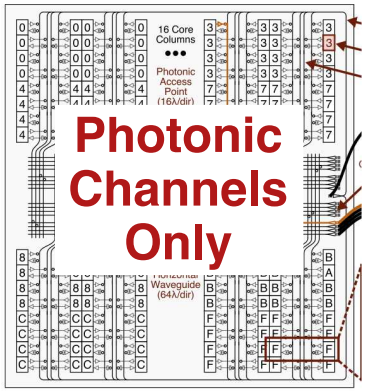

Batten, IEEE Micro'09

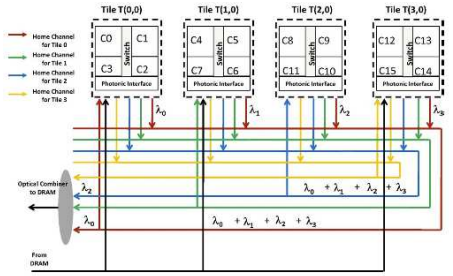

Morris, JSTQE'10

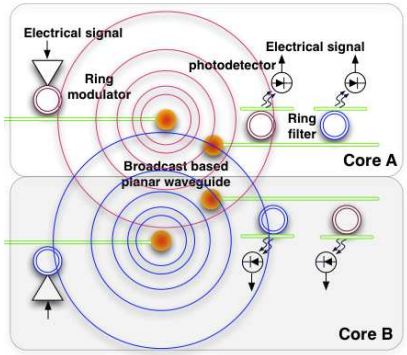

Li, DAC'09

Gu, DATE'09

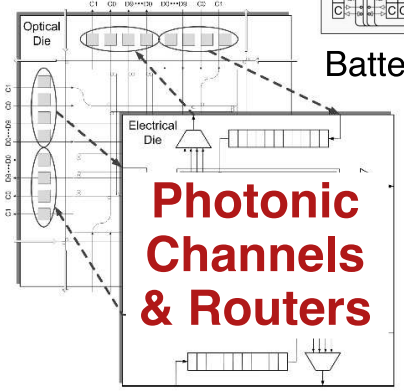

Cianchetti, ISCA'09

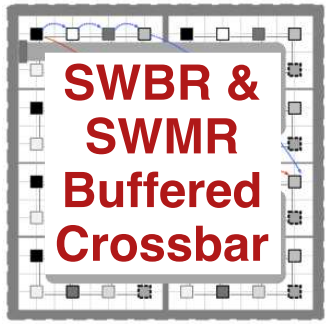

Pan, ISCA'09

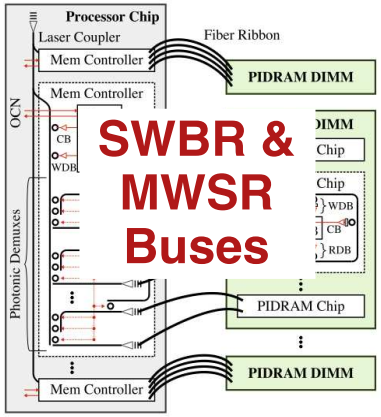

Beamer, ISCA'10

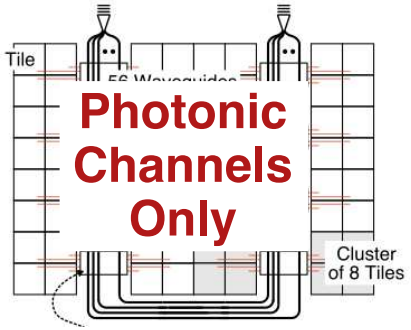

Joshi, NOCS'09

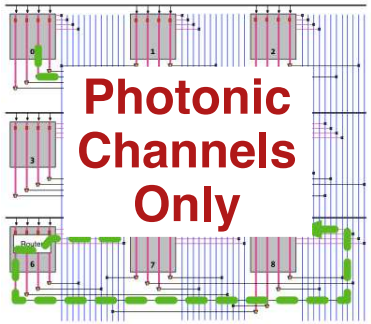

Koka, ISCA'10

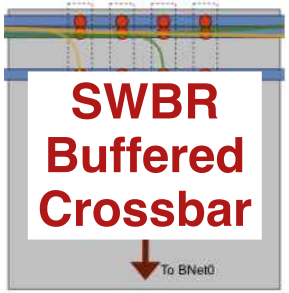

Kurian, PACT'10

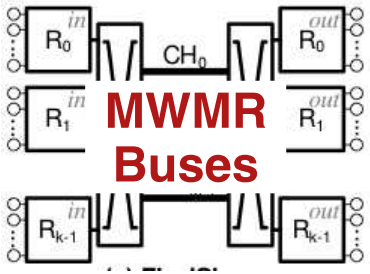

Pan, HPCA'10

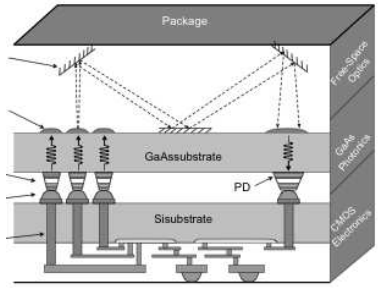

Xue, ISCA'10

Physical-Level Design

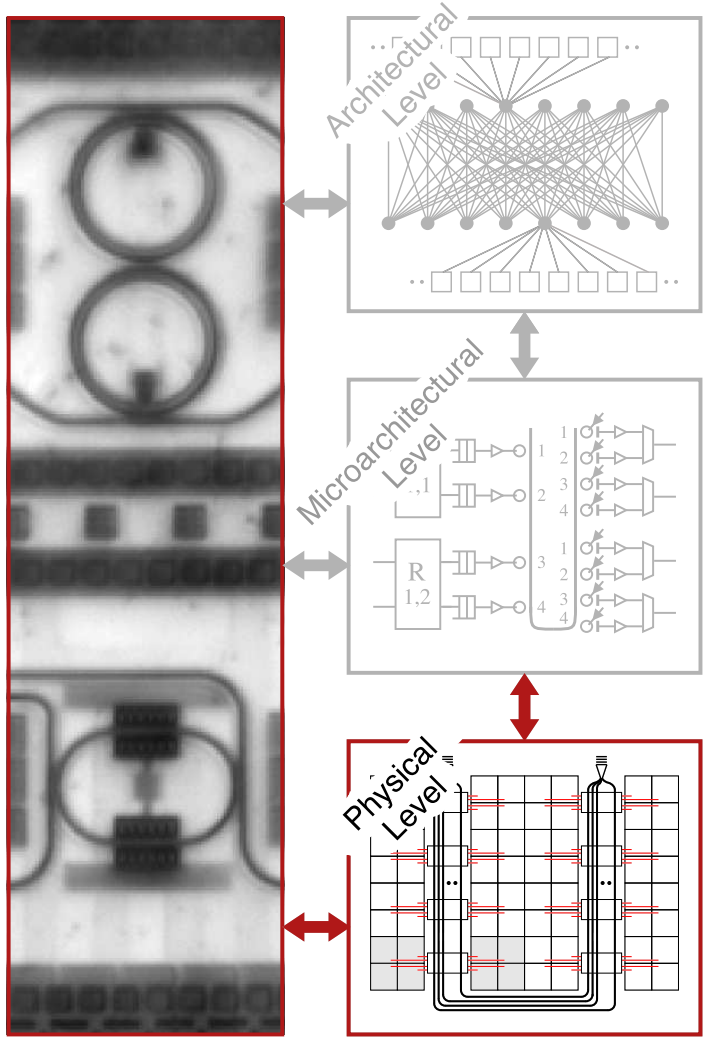

- Map microarchitectural design to physical substrate
- Use a nanophotonic abstract layout diagram to capture design decisions
- Decide how to assign wavelengths to waveguides and fibers
- Decide how to layout waveguides and organize fibers
- Decide where to place nanophotonic devices along waveguides

Bus Physical Design Patterns

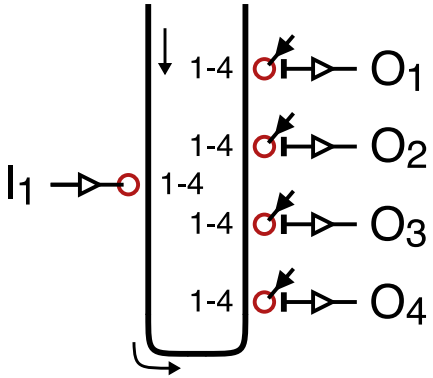

Single Waveguide

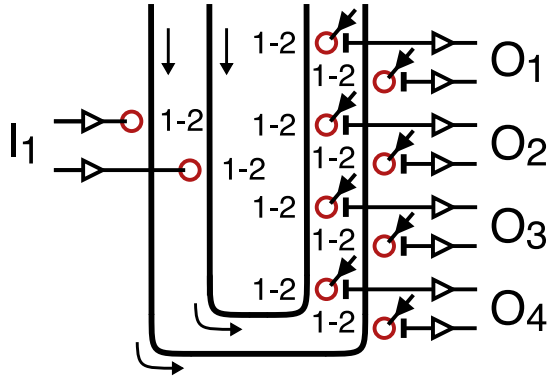

Bit Slicing

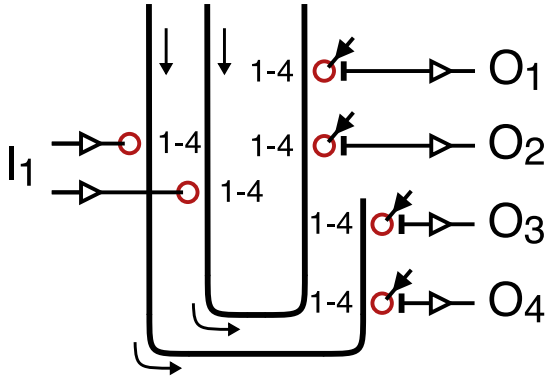

Reader Slicing with Redundant Transmitters

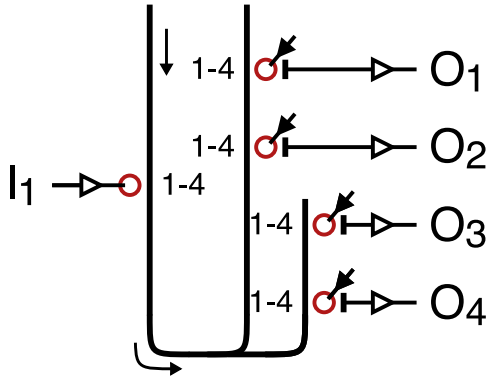

Reader Slicing with Optical Power Splitting

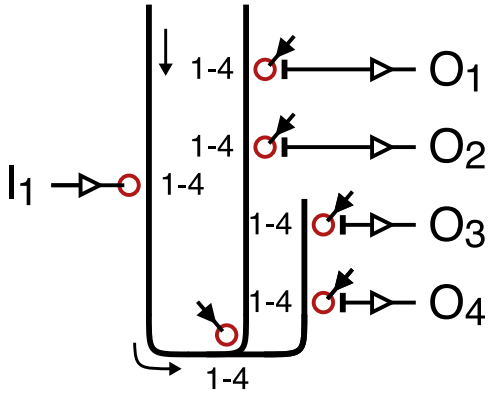

Reader Slicing with Optical Power Guiding

Crossbar Physical Design Patterns

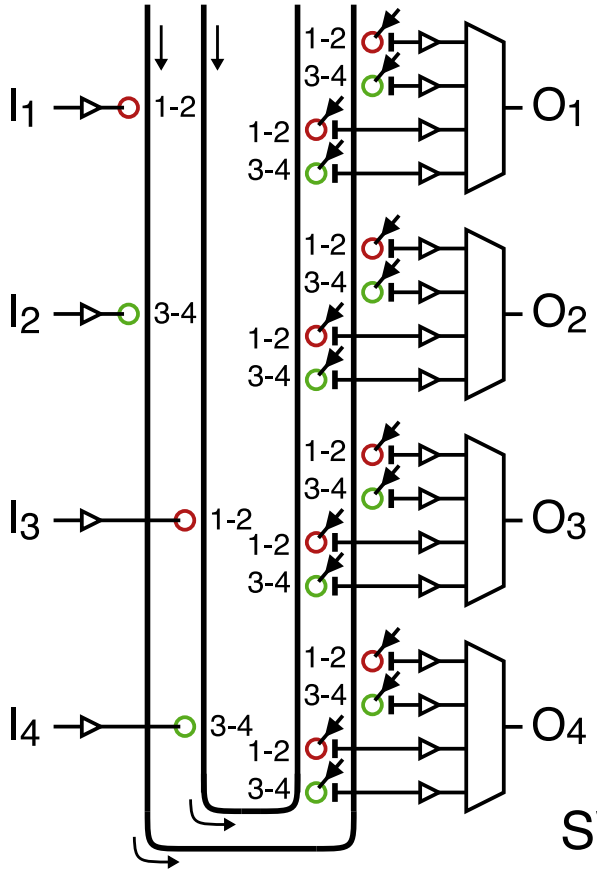

SWMR Crossbar with Bus Slicing

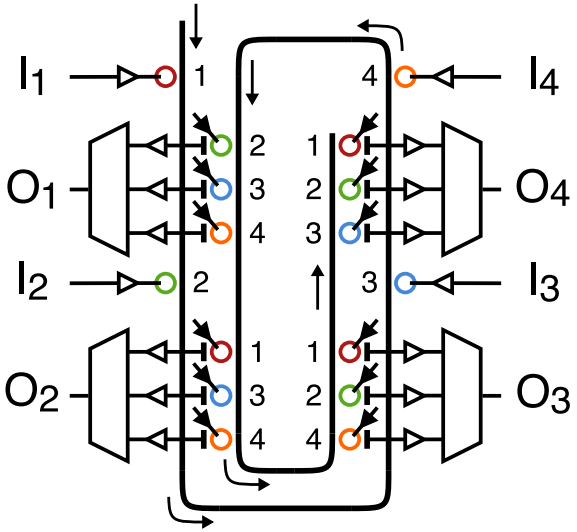

SWMR Crossbar with Double-Serpentine Layout

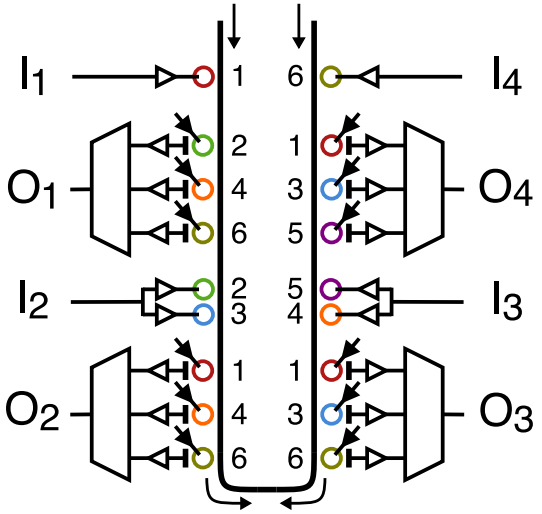

SWMR Crossbar with Single-Serpentine Layout

Point-to-Point Channels Physical Design Patterns

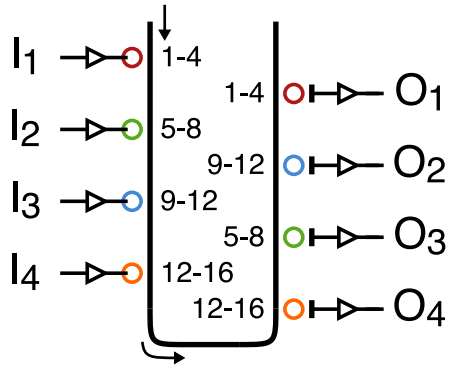

Single Waveguide and Serpentine Layout

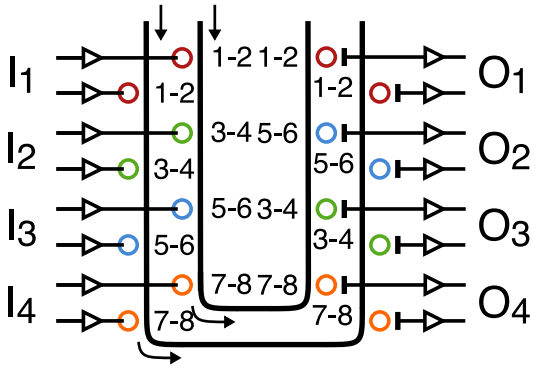

Bit Slicing and Serpentine Layout

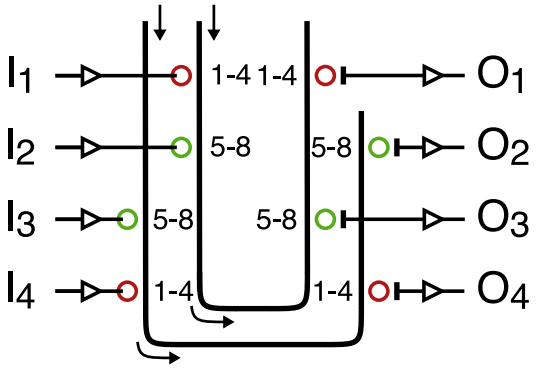

Channel Slicing and Serpentine Layout

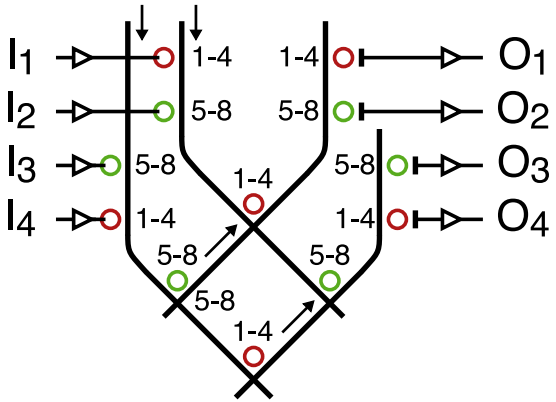

Channel Slicing and Ring-Filter Matrix Layout

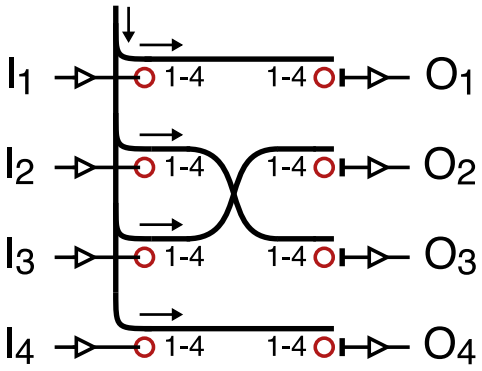

Channel Slicing and Point-to-Point Layout

Abstract Physical Layouts for Crossbar

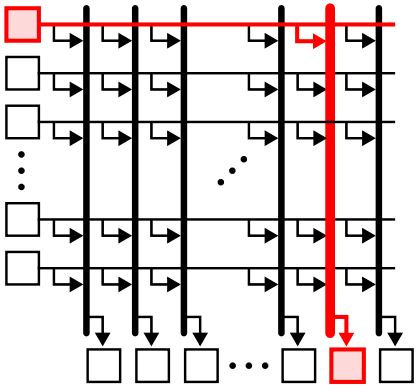

Architectural-Level
Global 64x64 Crossbar

Microarchitectural-Level
Single-Writer, Multiple-Reader

Physical-Level
Long vs Short Serpentine Layout

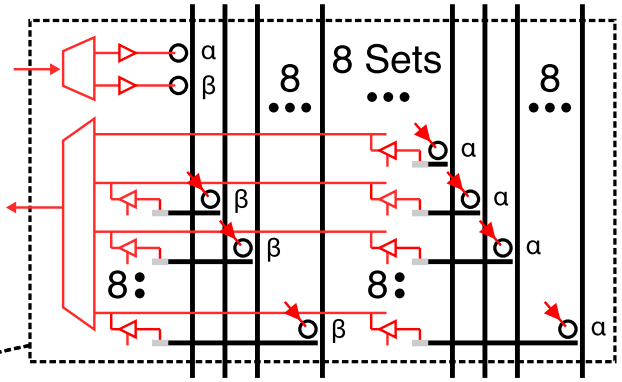

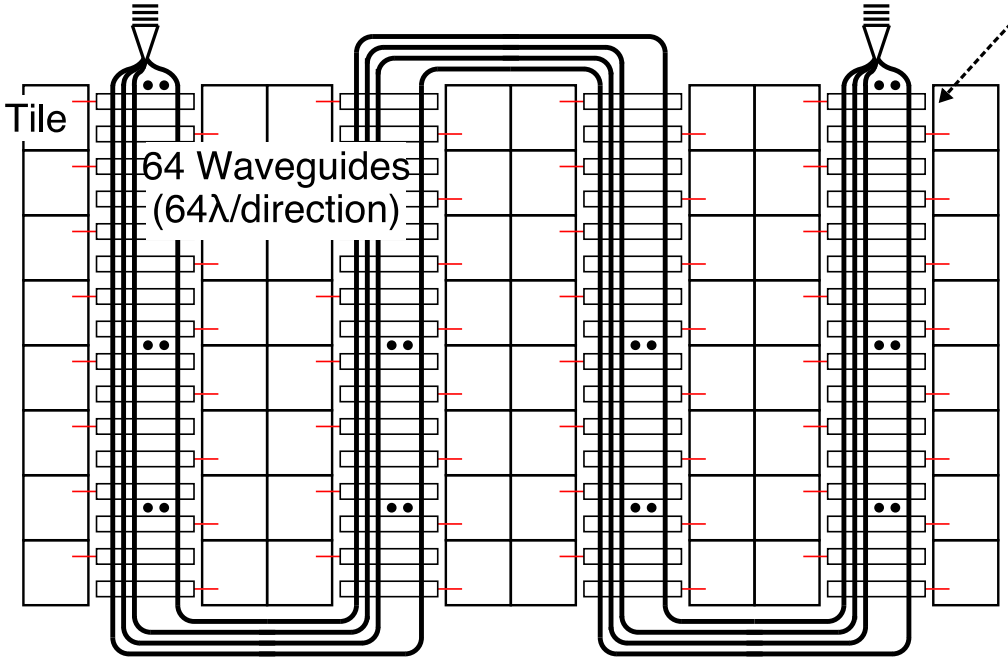

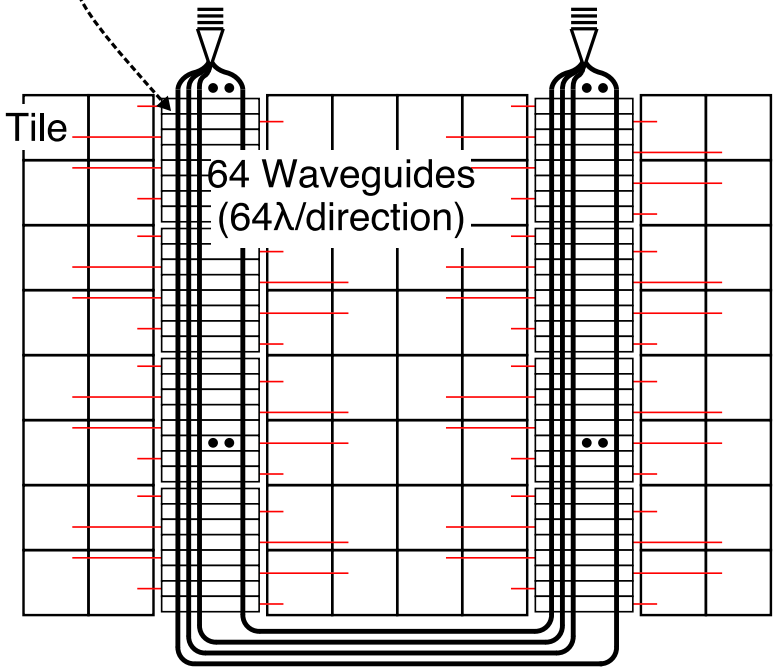

Abstract Physical Layouts for Butterfly

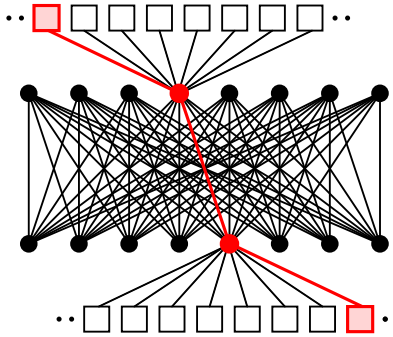

Architectural-Level
8-ary 2-fly Butterfly Topology

Microarchitectural-Level
Nanophotonic Channels, Electrical Routers

Physical-Level
Grid vs Serpentine Layout

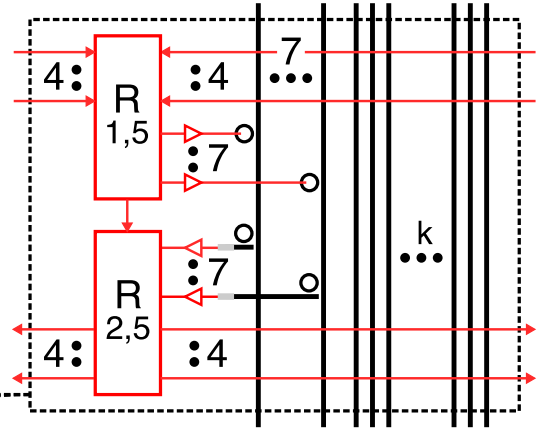

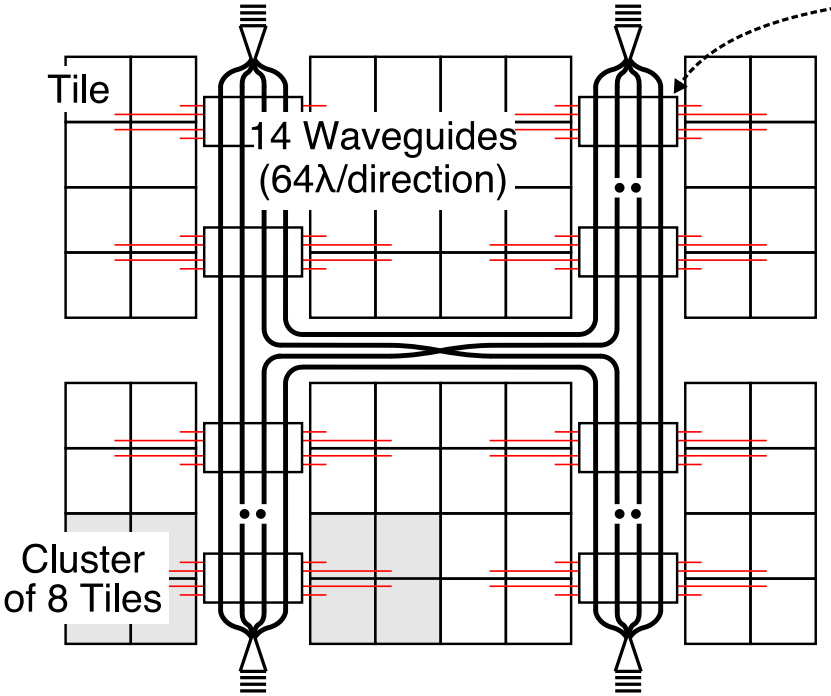

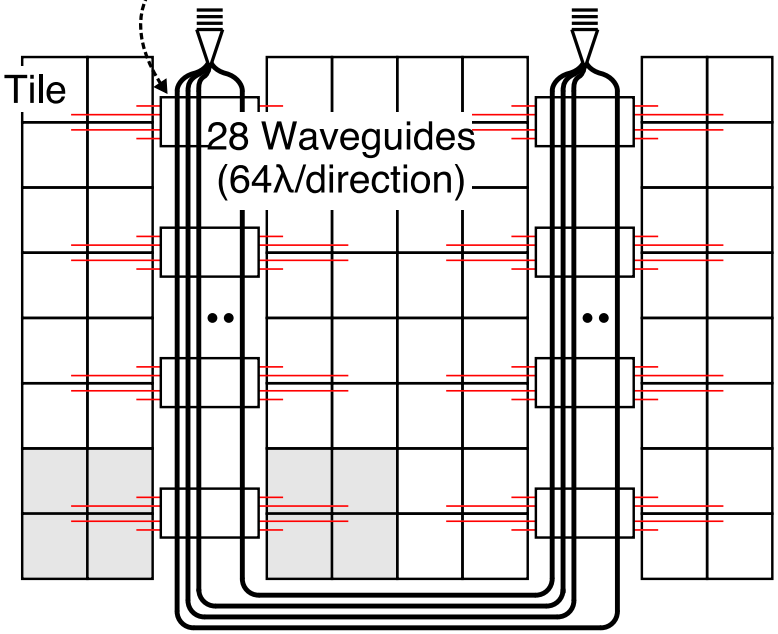

Example Physical-Level Analysis

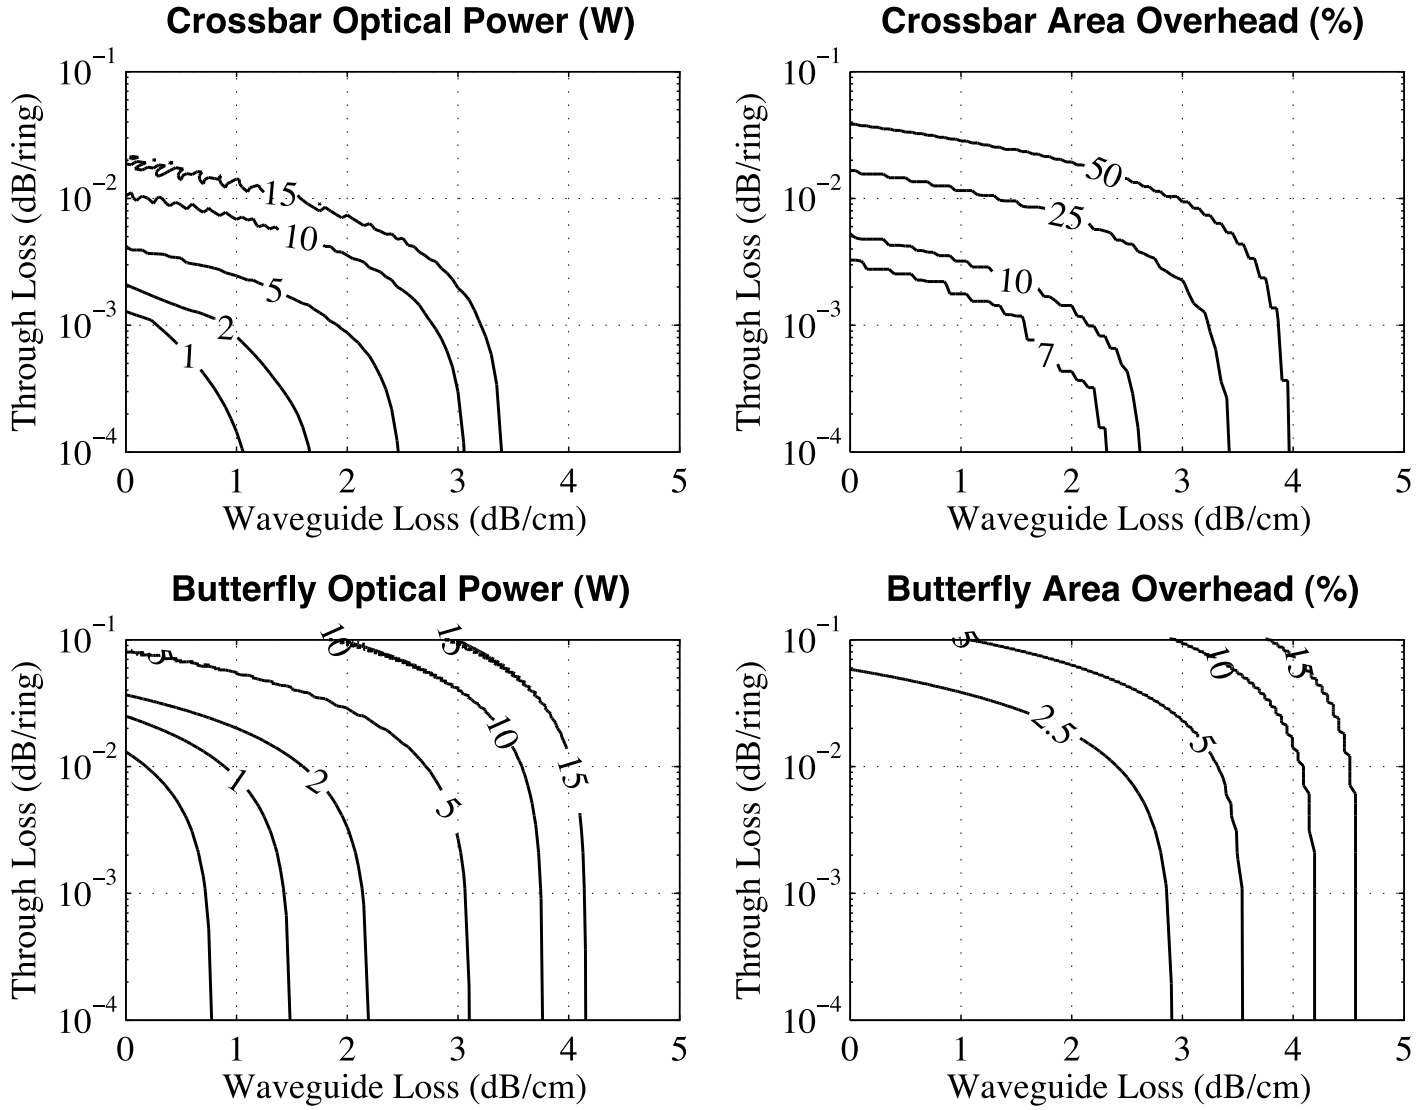

Example Network Simulation Results

Electrical Baseline ecmeshx2

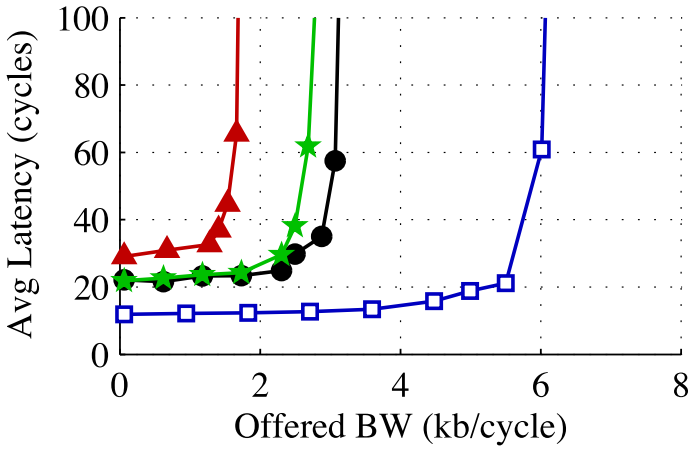

Photonic Network with Two Back-to-Back Butterfly Networks

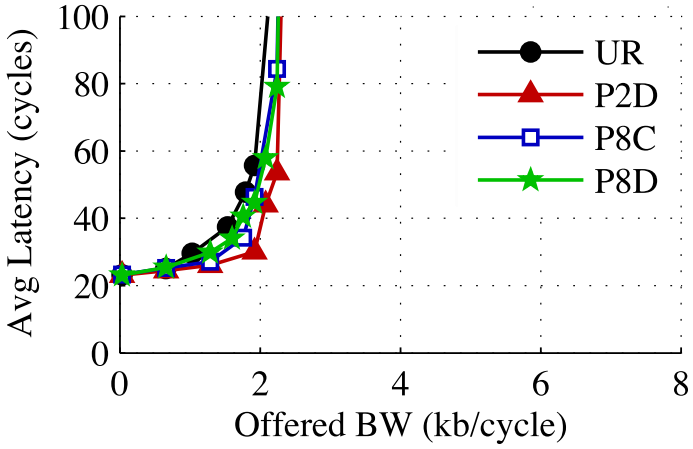

Power Breakdown for Various Networks

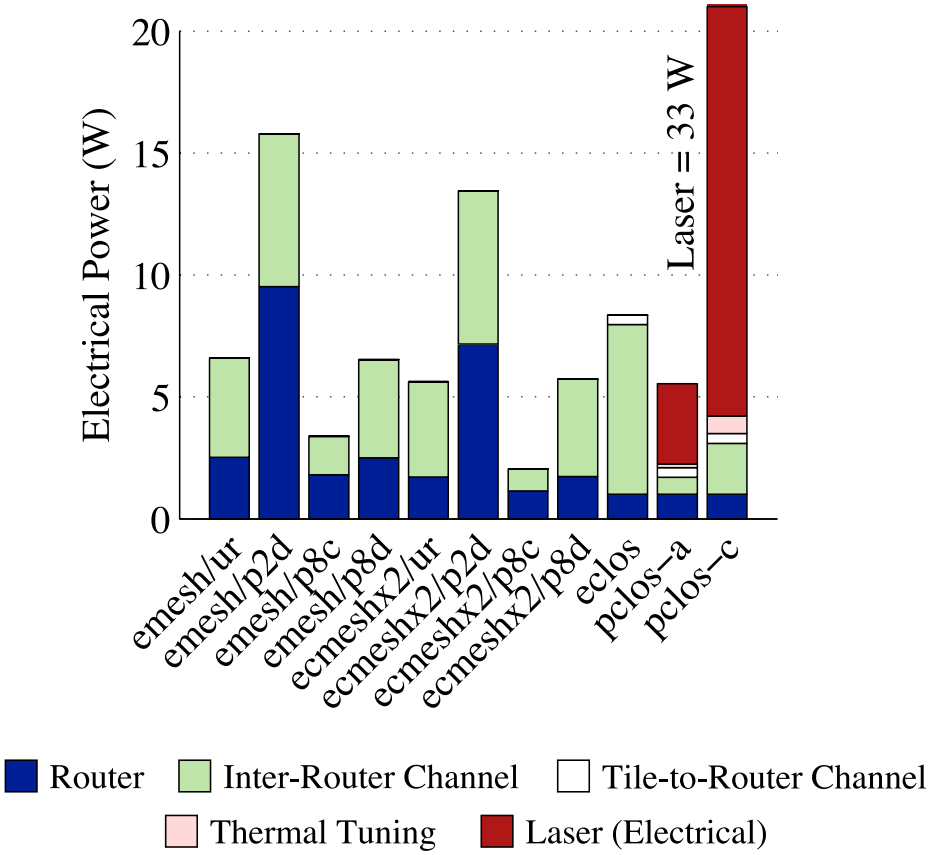

Design Guidelines

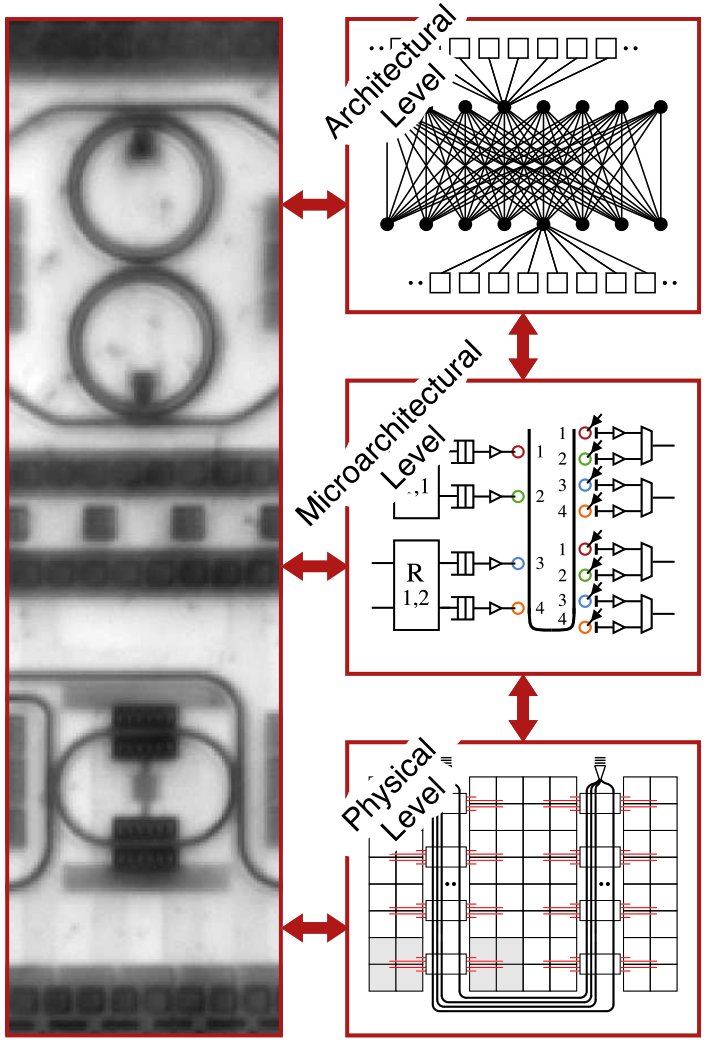

1. Clearly specify the logical topology – easy to confuse passively WDM-routed wavelengths with true network routing, helps position design relative to other proposals
2. Iterate through three design levels – many mappings at each level, iterative approach avoids over-constraining the design space
3. Use aggressive electrical baseline – is a low-dimensional mesh the best electrical comparison? are simple repeated wires really the best we can do?

Design Guidelines

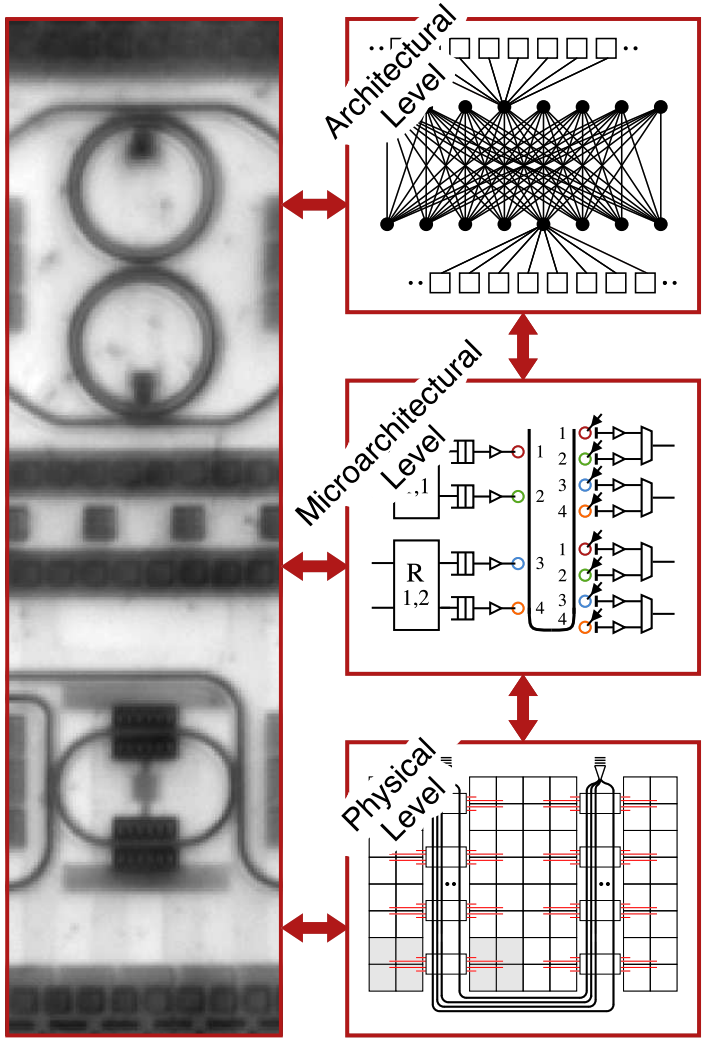

4. Broad range of device parameters – emerging technology means a single set of parameters is probably meaningless, sensitivity studies are essential
5. Consider fixed-power overheads – only studying high-utilization workloads hides the significant impact of fixed transceiver circuit power, laser power, and thermal tuning power
6. Motivate network complexity – if we should seek the simplest network architecture (and the simplest devices) that achieves our application requirements

Design Guidelines

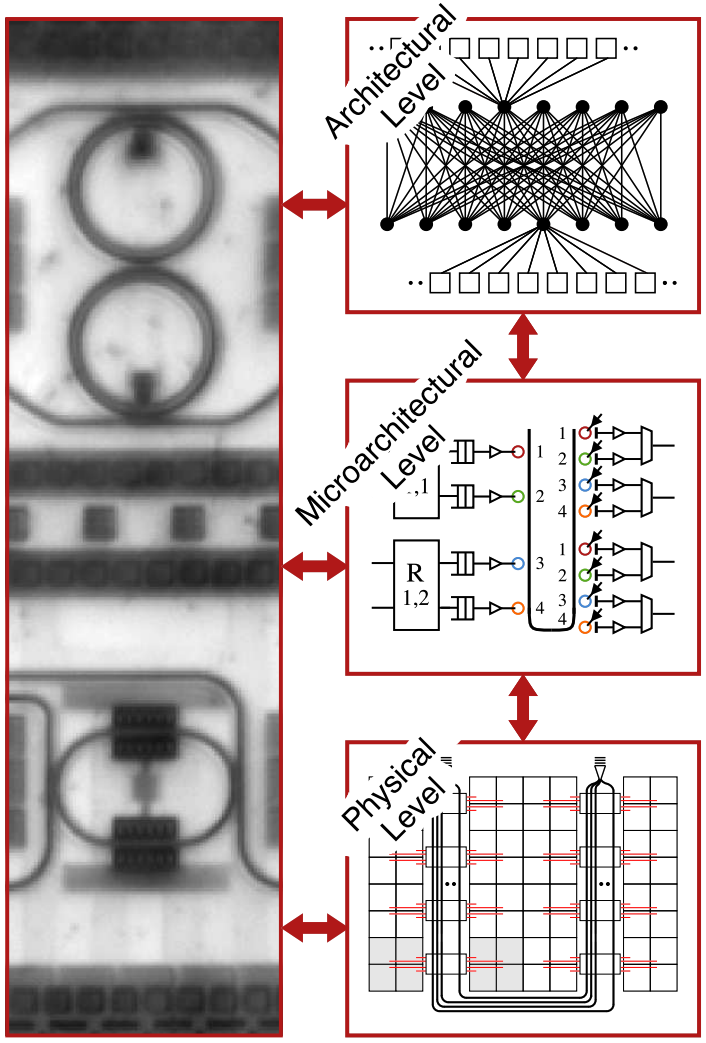

1. Clearly specify the logical topology
2. Iterate through three design levels
3. Use aggressive electrical baseline
4. Broad range of device parameters
5. Consider fixed-power overheads
6. Motivate network complexity