Optical Network on Chip in 3D Architectures

Sébastien Le Beux

Ecole Centrale de Lyon, France
contact : Sebastien.Le-Beux@ec-lyon.fr

October 20th, 2011

Outline

MPSoC Design Trends

- Multiple layers → multiple technologies (3D integration + heterogeneity) :
  - computation → electrical layer
  - communication → optical layer

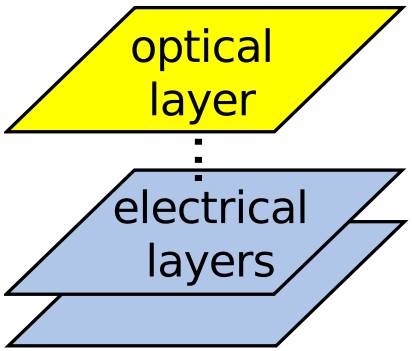

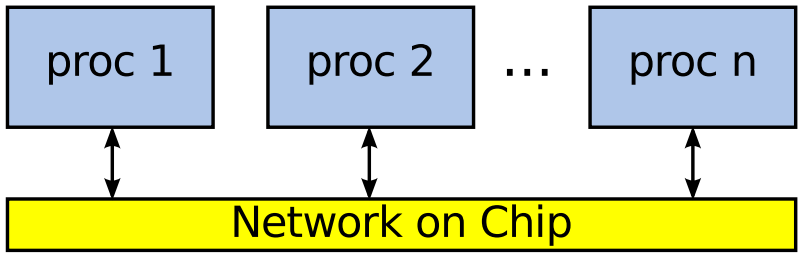

MPSoC Design Trends

- Multiple layers → multiple technologies (3D integration + heterogeneity) :
  - computation → electrical layer
  - communication → optical layer

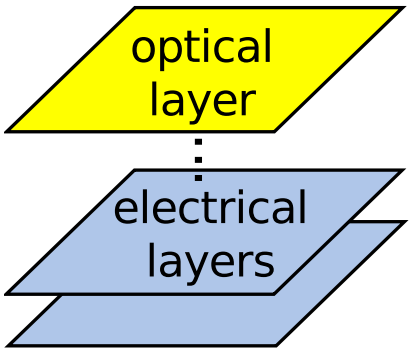

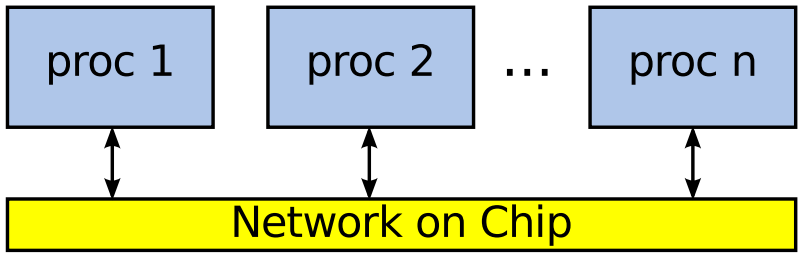

- Optical Network on Chip (ONoC) characteristics
  - high throughput : Wavelength Division Multiplexing, WDM
  - long range communications (chip scale)
  - low latency

Related Work and Contribution

- Mesh [4], multistage [2], fat-tree [5], torus [12]
  - waveguide crossings $\rightarrow$ losses (about 0.3dB for the multistage)

Related Work and Contribution

- Mesh [4], multistage [2], fat-tree [5], torus [12]
  - waveguide crossings $\rightarrow$ losses (about 0.3dB for the multistage)
- Ring topology : Corona [1], Firefly [10], FlexiShare [9], ATAC [11]
  - no waveguide crossings and simpler layout

Related Work and Contribution

- Mesh [4], multistage [2], fat-tree [5], torus [12]
  - waveguide crossings → losses (about 0.3dB for the multistage)
- Ring topology : Corona [1], Firefly [10], FlexiShare [9], ATAC [11]
  - no waveguide crossings and simpler layout

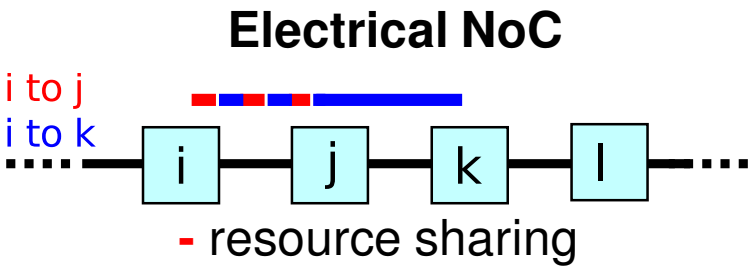

- resource sharing

+ number of resources

Related Work and Contribution

- Mesh [4], multistage [2], fat-tree [5], torus [12]
  - waveguide crossings → losses (about 0.3dB for the multistage)
- Ring topology : Corona [1], Firefly [10], FlexiShare [9], ATAC [11]
  - no waveguide crossings and simpler layout

Electrical NoC

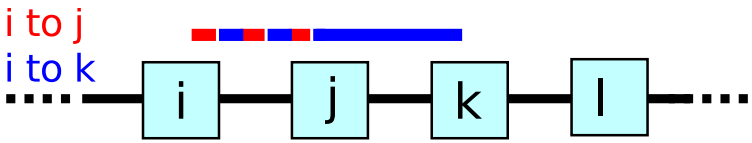

- resource sharing

+ number of resources

Optical NoC

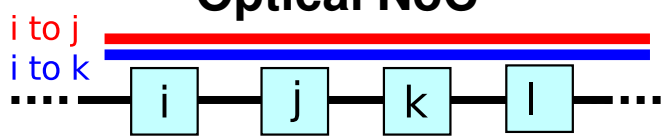

+ no contention in ONoC

- shared access, number of resources

Related Work and Contribution

- Mesh [4], multistage [2], fat-tree [5], torus [12]
  - waveguide crossings → losses (about 0.3dB for the multistage)
- Ring topology : Corona [1], Firefly [10], FlexiShare [9], ATAC [11]
  - no waveguide crossings and simpler layout

Electrical NoC

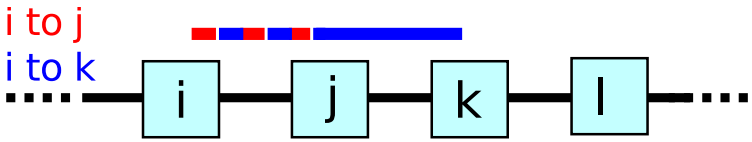

- resource sharing

+ number of resources

Optical NoC

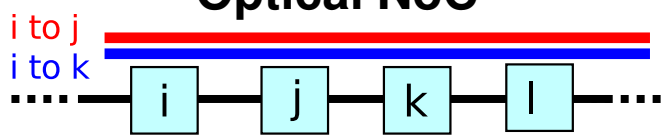

+ no contention in ONoC

- shared access, number of resources

ORNoC : Optical Ring Network on Chip

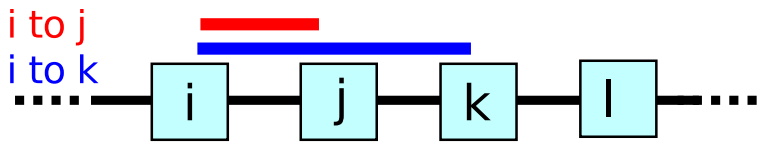

+ zero contention (WDM + virtual waveguide partitioning)

+ optimal number of resources

Optical Network Interface (ONI)

- Operation mode :
  - ejection
  - pass through
  - injection

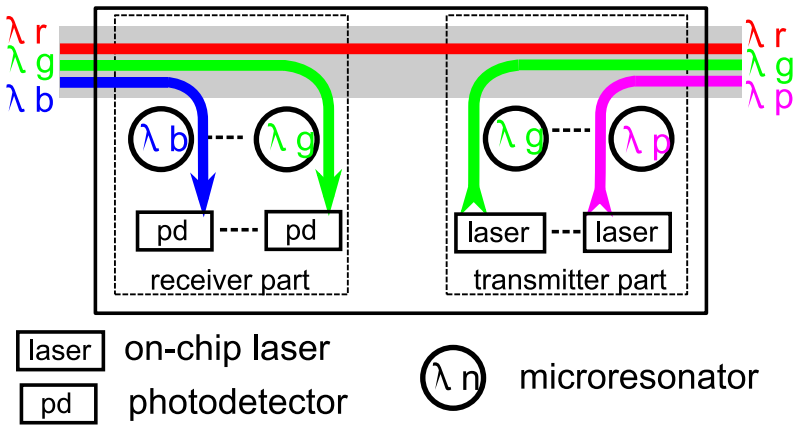

Optical Network Interface (ONI)

- Operation mode :
  - ejection
  - pass through
  - injection

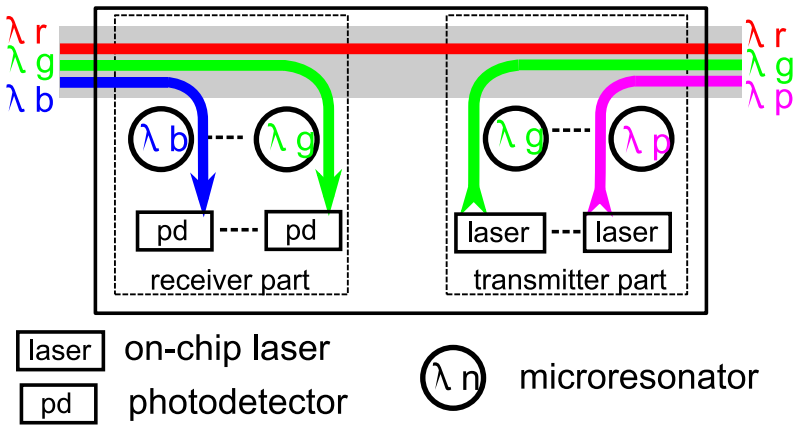

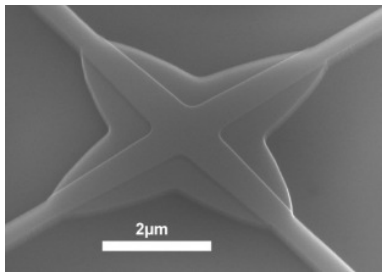

W. Bogaerts et al.,
IEEE JSTQE, 16(1), 33(2010)

Optical Network Interface (ONI)

- Operation mode :
  - ejection
  - pass through
  - injection

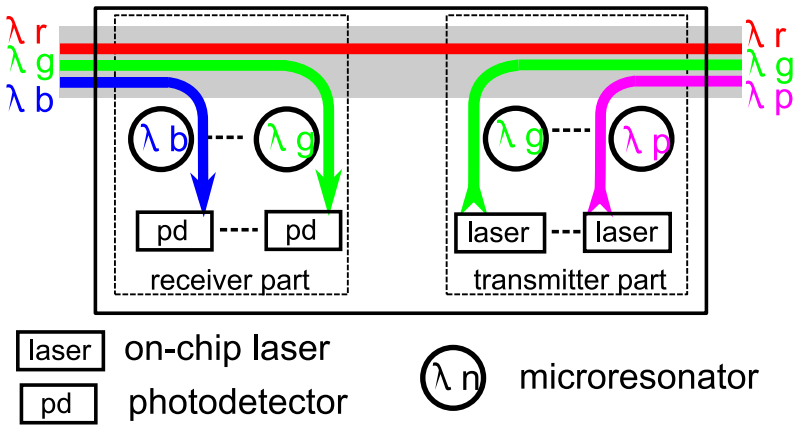

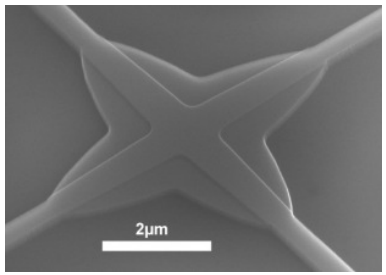

W. Bogaerts et al.,
IEEE JSTQE, 16(1), 33(2010)

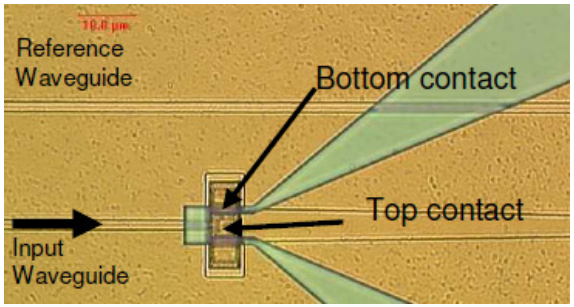

L. Vivien et al.,
Optics Express, 17, 6252 (2009)

Optical Network Interface (ONI)

- Operation mode :
  - ejection
  - pass through
  - injection

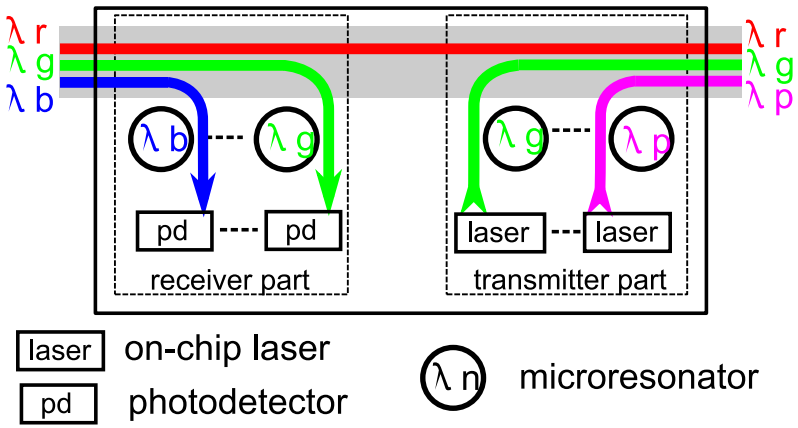

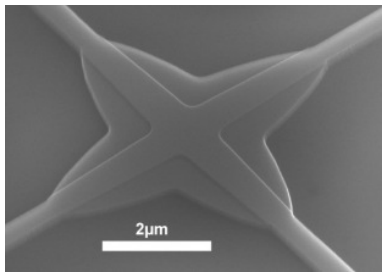

W. Bogaerts et al.,
IEEE JSTQE, 16(1), 33(2010)

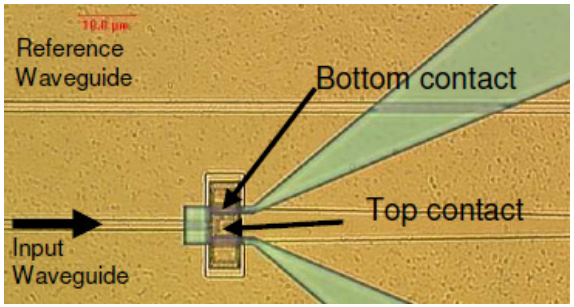

L. Vivien et al.,
Optics Express, 17, 6252 (2009)

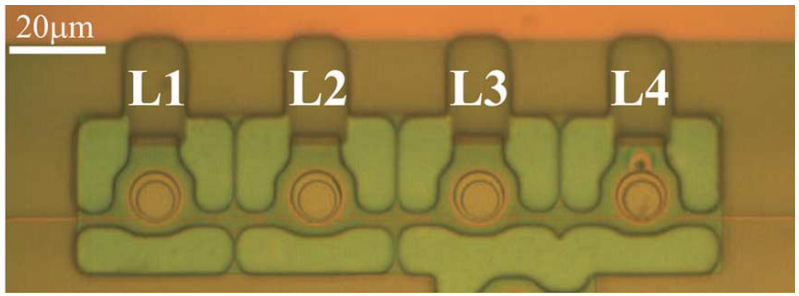

J. Van Campenhout et al., IEEE PTL, 20, 1345 (2008)

ORNoC in 3D Architecture

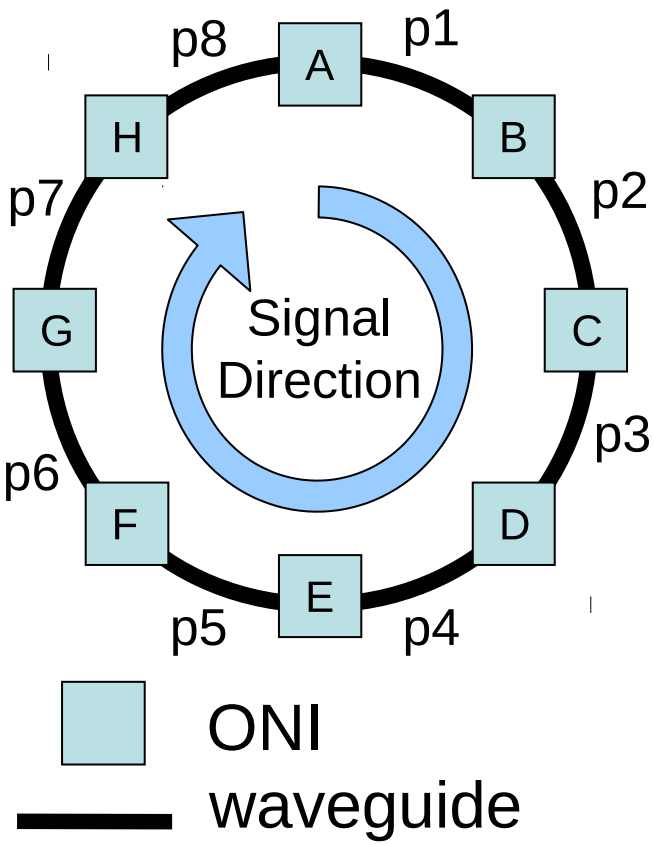

ORNoC in 3D Architecture

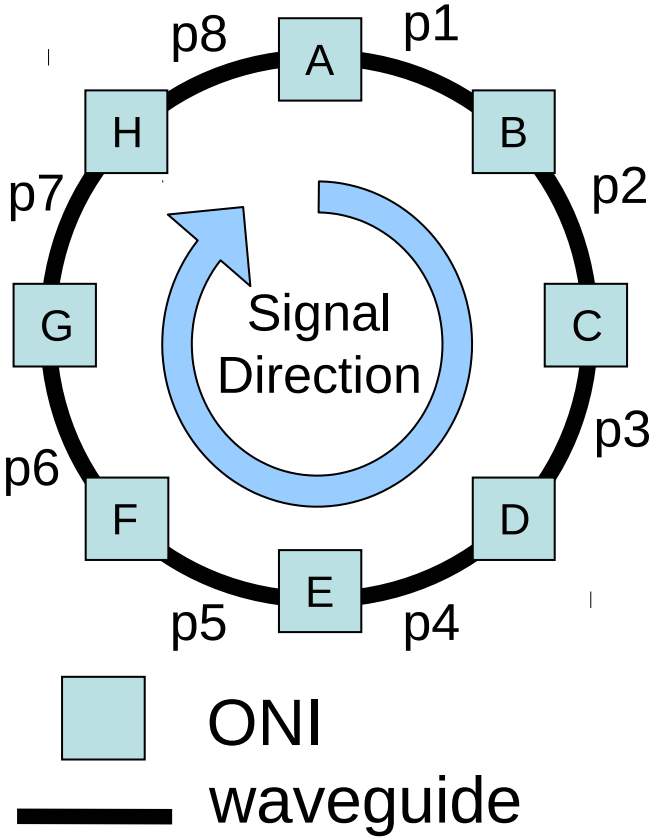

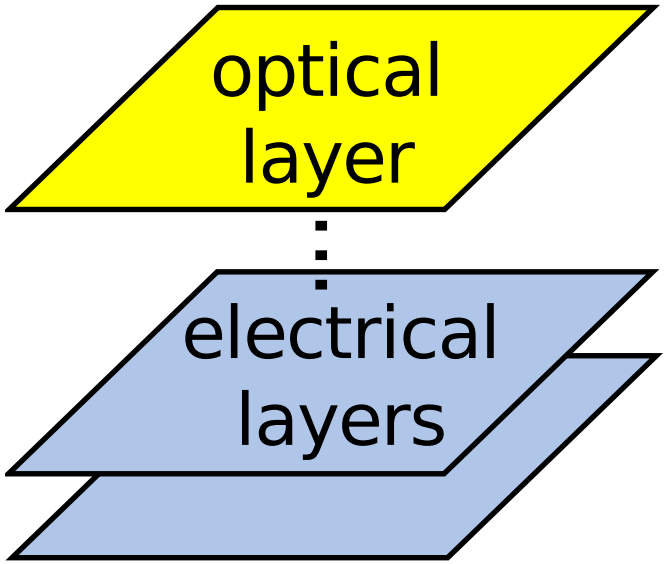

ORNoC in 3D Architecture

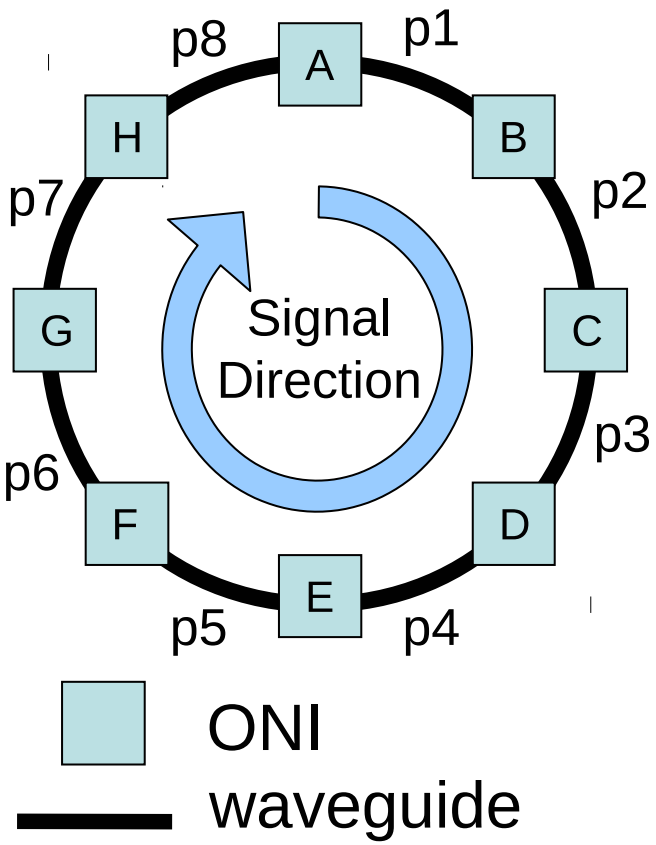

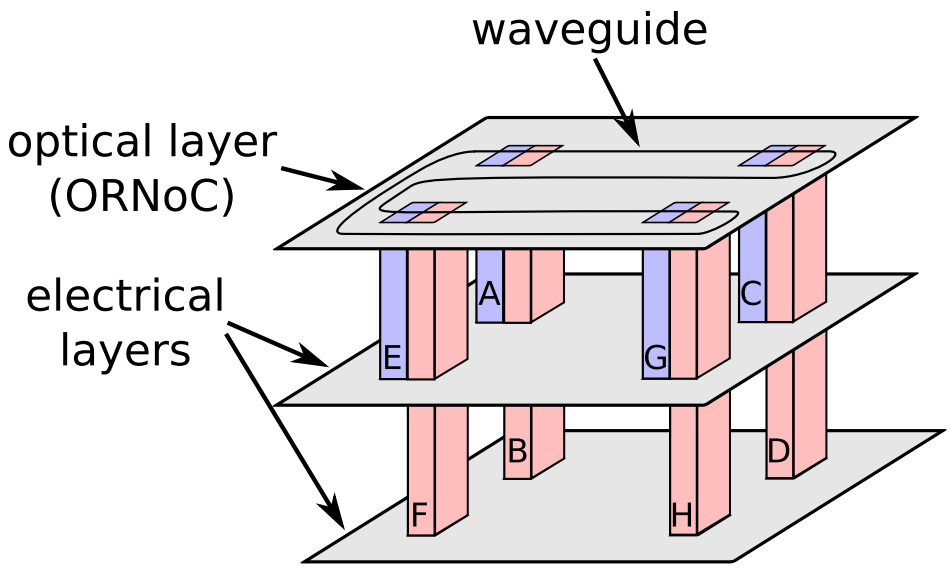

- Communication hierarchy :
  - Electrical NoC → intra-layer communication
  - ORNoC → inter-layer communication

ORNoC in 3D Architecture

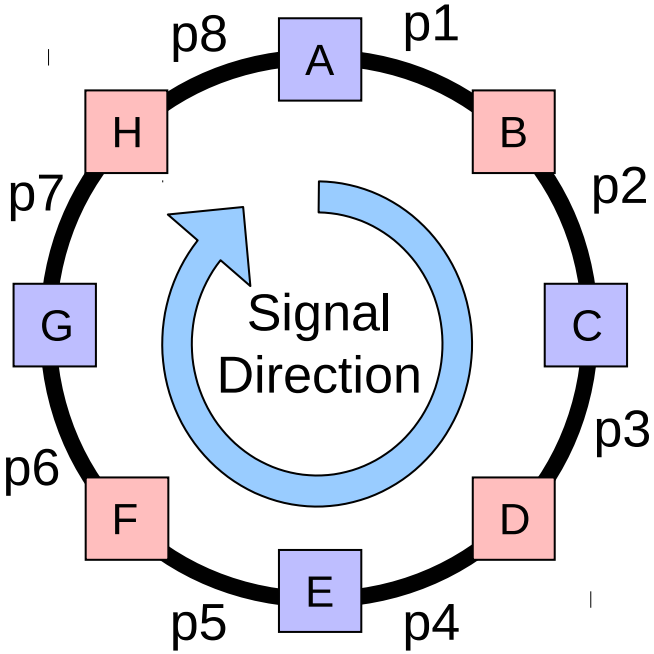

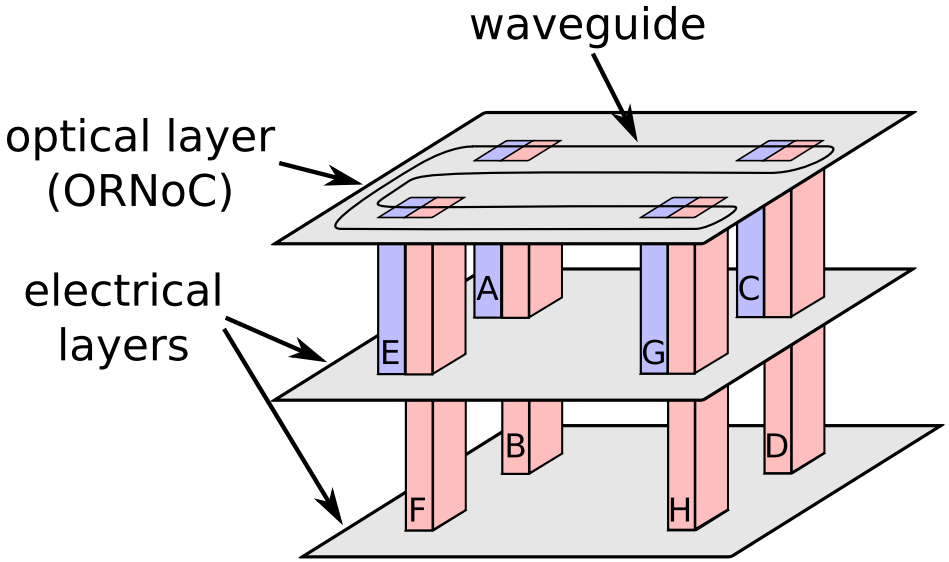

- Communication hierarchy :
  - Electrical NoC → intra-layer communication
  - ORNoC → inter-layer communication

How to configure ONI ?

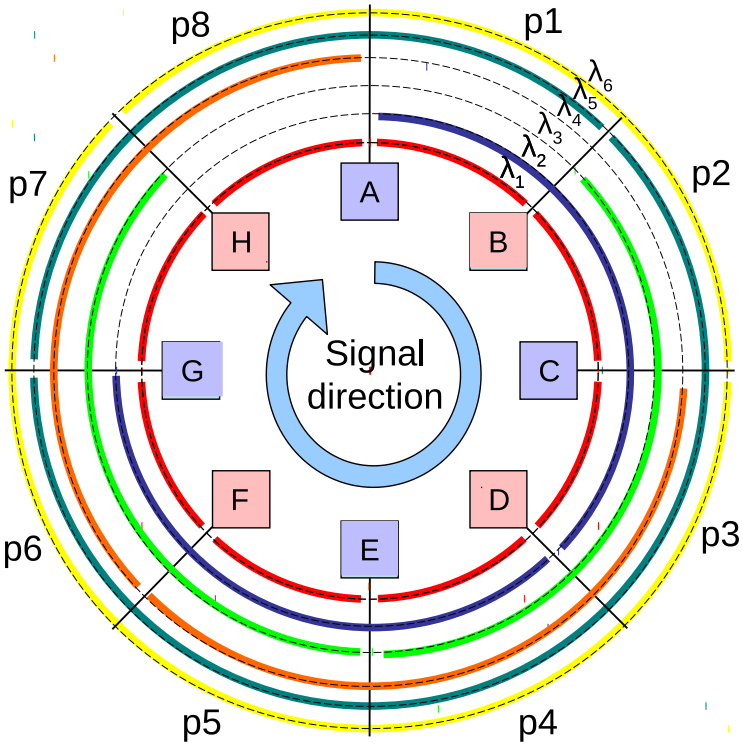

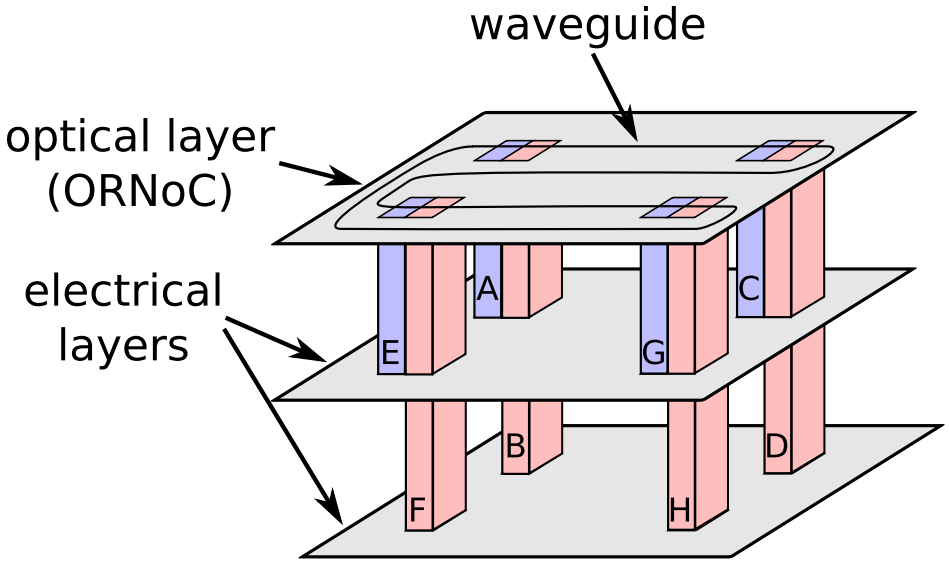

- ORNoC benefits : contention free, scalable, low power

How to configure ONI ?

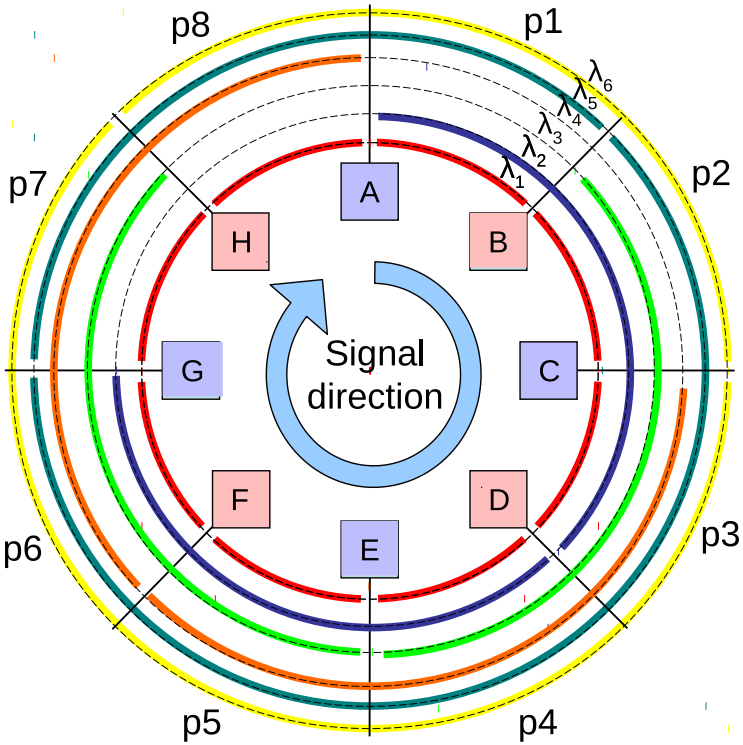

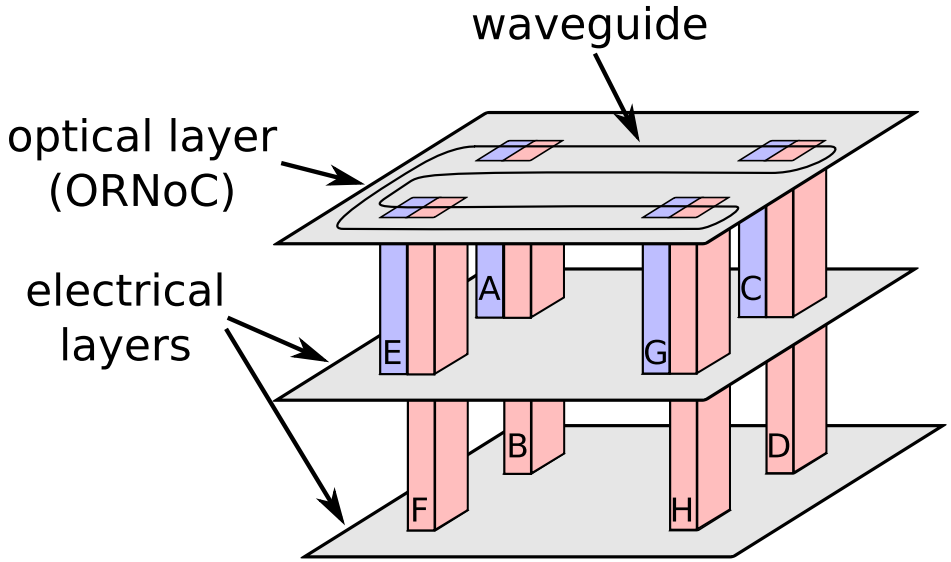

- ORNoC benefits : contention free, scalable, low power
- Need for...
  - (high level) layout guidelines
  - ONI area overhead estimation
  - regularity and reuse (x, y and z dimensions)

Layout Guidelines

- Objective : take benefits from the regularity of ORNoC

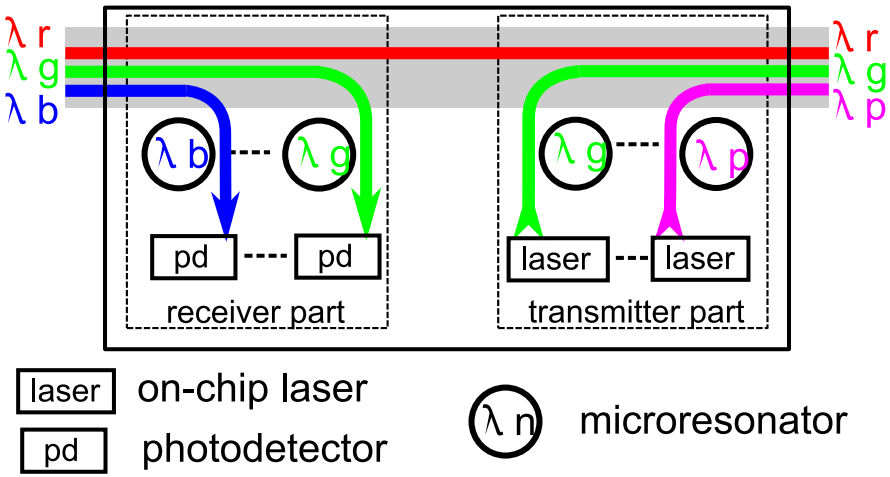

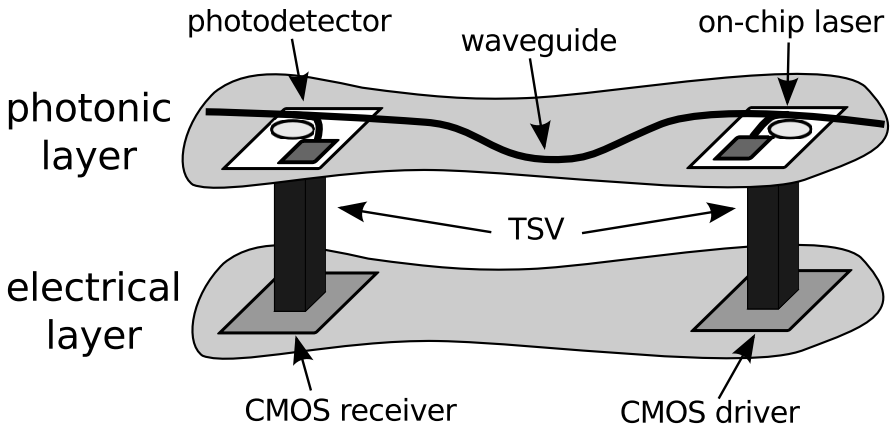

Layout Guidelines

- Objective : take benefits from the regularity of ORNoC

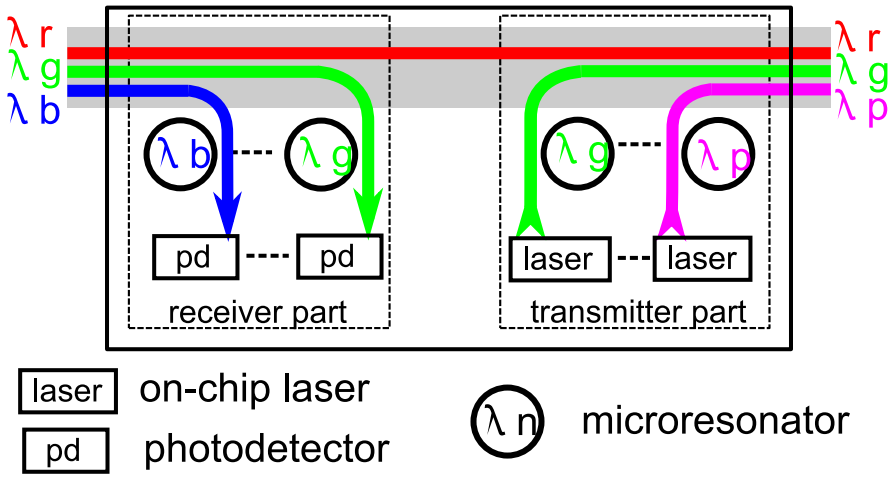

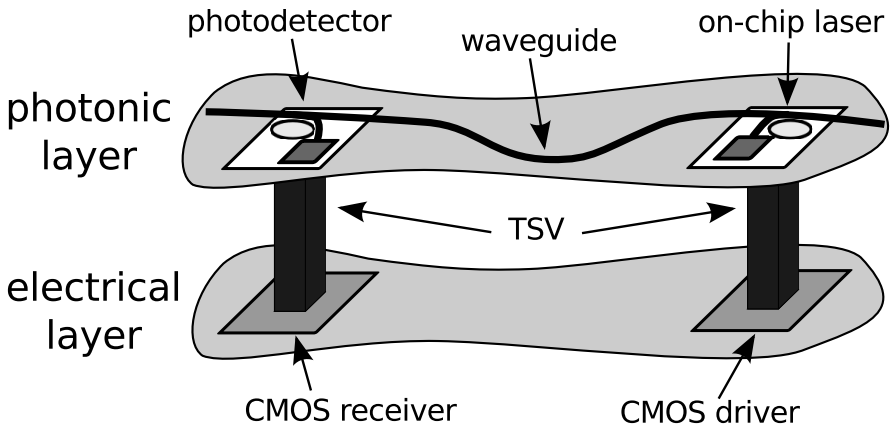

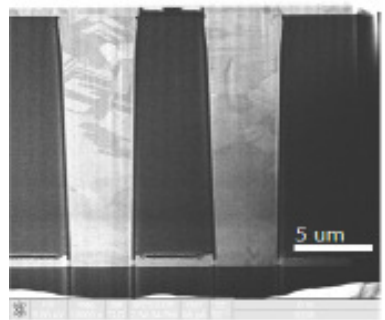

Layout Guidelines

- Objective : take benefits from the regularity of ORNoC

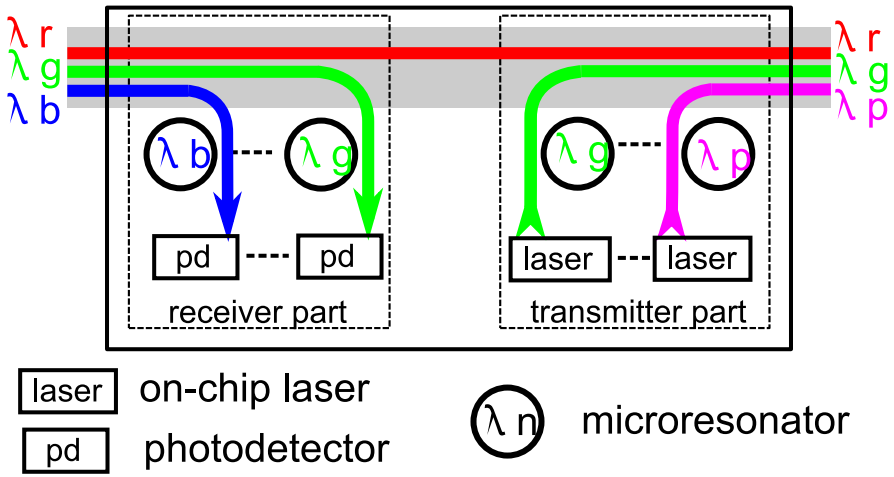

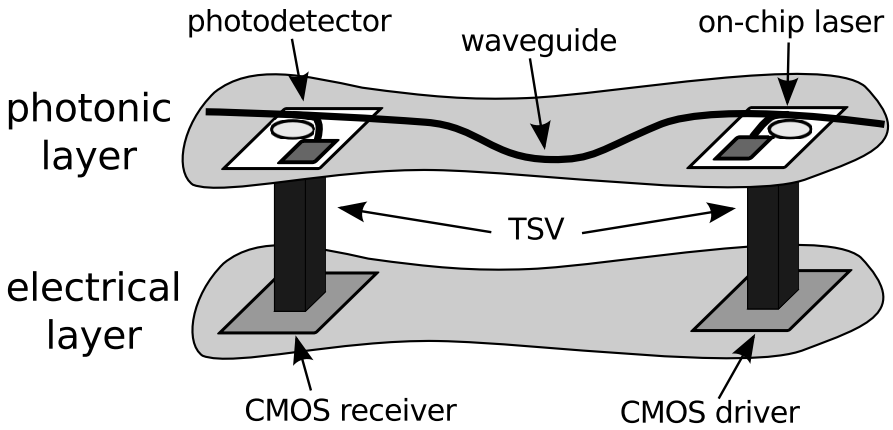

- VC_{area}
  - elementary footprint to consider (worst case)
  - highly technology dependent

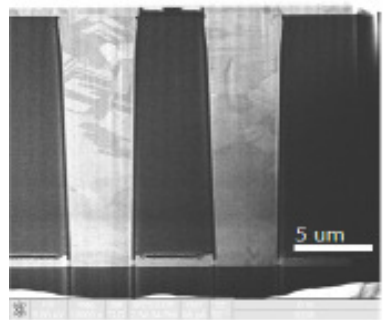

Layout Examples

WL : number of wavelengths

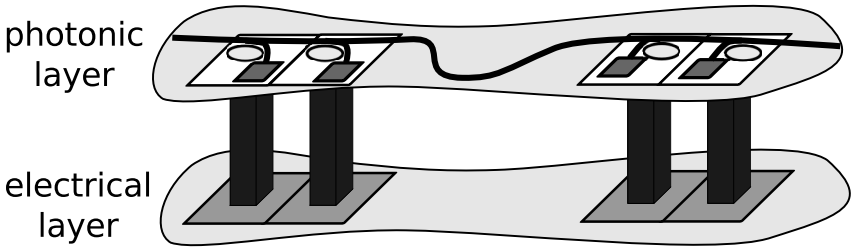

Layout Examples

WL : number of wavelengths

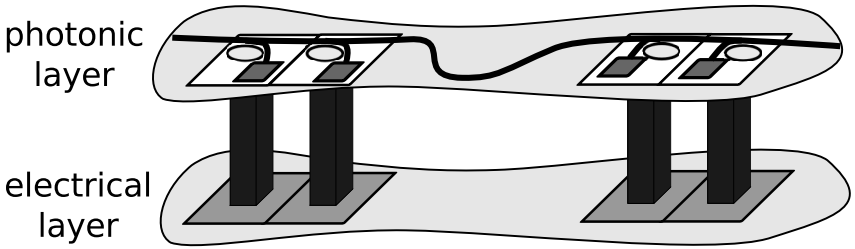

WG : number of waveguides

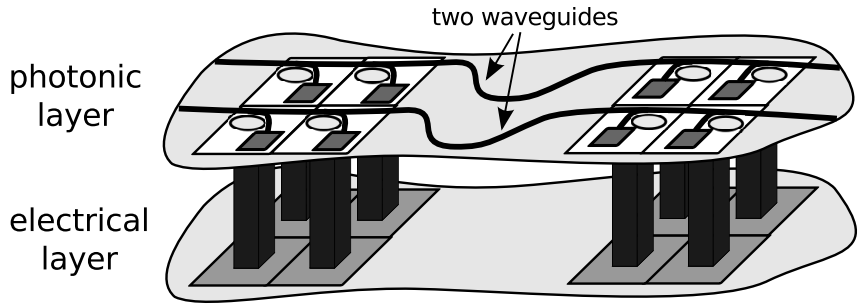

Layout Examples

WL : number of wavelengths

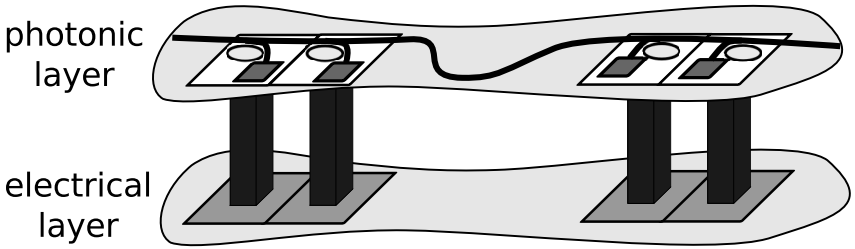

WG : number of waveguides

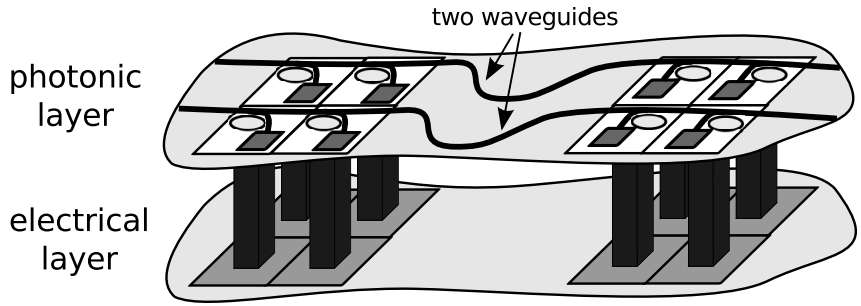

L : number of electrical layers

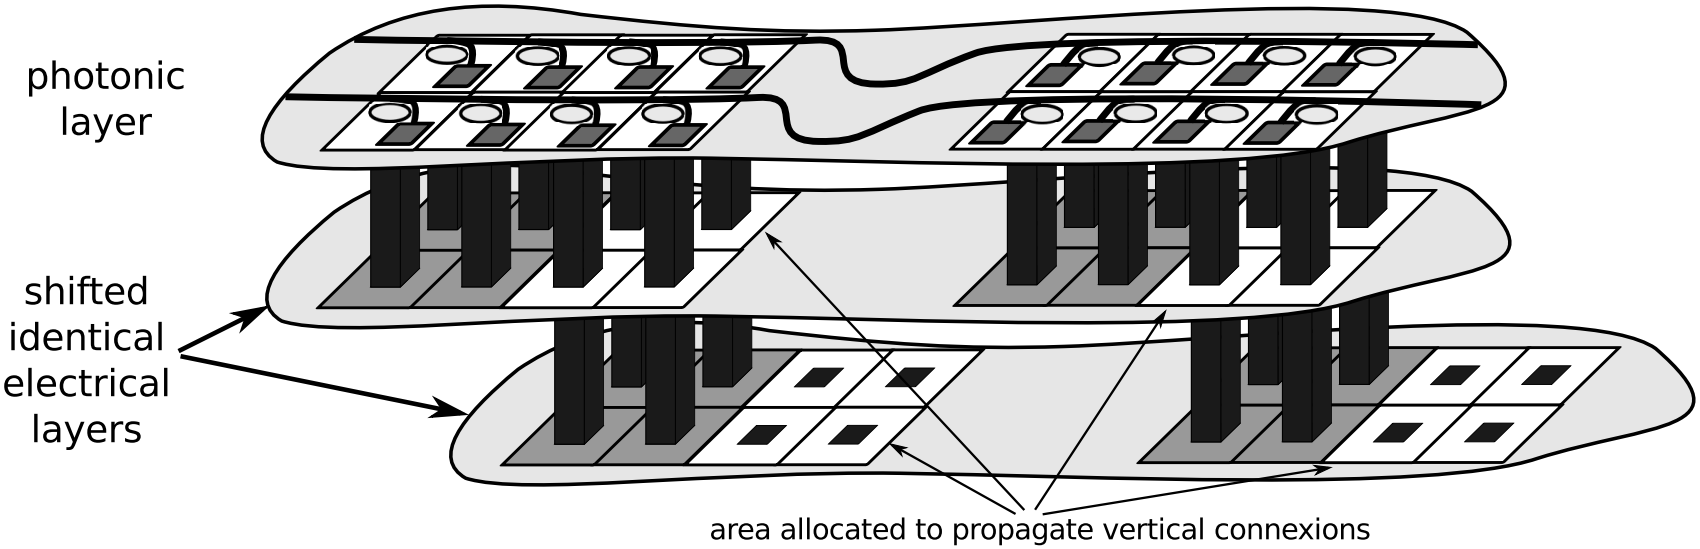

Area overhead estimation

- Area used to implement each ONI :

$$ONI_{area} = 2 \times VC_{area} \times WL \times WG$$

- Area overhead used to “propagate” ONI :

$$ONI_{area\_overhead} = ONI_{area} \times (L - 1)$$

Area overhead estimation

- Area used to implement each ONI :

$$ONI_{area} = 2 \times VC_{area} \times WL \times WG$$

- Area overhead used to “propagate” ONI :

$$ONI_{area\_overhead} = ONI_{area} \times (L - 1)$$

- Total area obtained by considering the number of ONI per layer

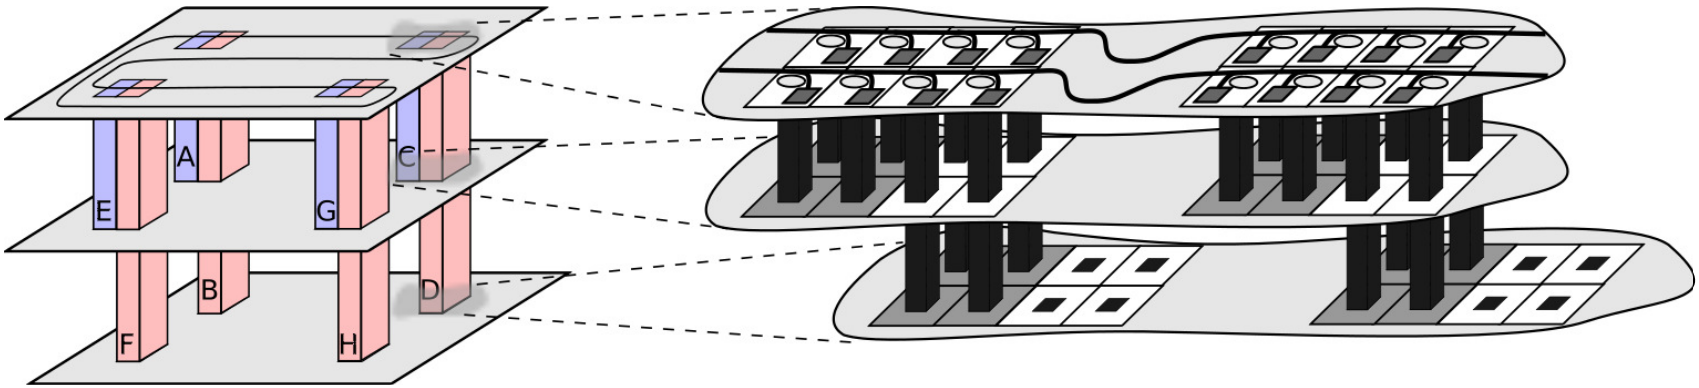

Area overhead estimation

- Assumptions :
  - VC_{area} ?
    - CMOS driver : $8.5\mu m \times 9.5\mu m$ ($350\mu A$ current threshold [13] and $0.13\mu m$ CMOS technology [7])
    - TSV : $pitch \approx 5\mu m \times 5\mu m$ [3]
    - photonic receiver : less than $20\mu m^2$ [6]
    - CMOS receiver footprint < CMOS driver footprint [7]
    - photonic transmitter : microdisk laser radius ($7.5\mu m$ [13]), microdisk resonator radius ($10\mu m$), waveguide diameter ($1\mu m$ [8])

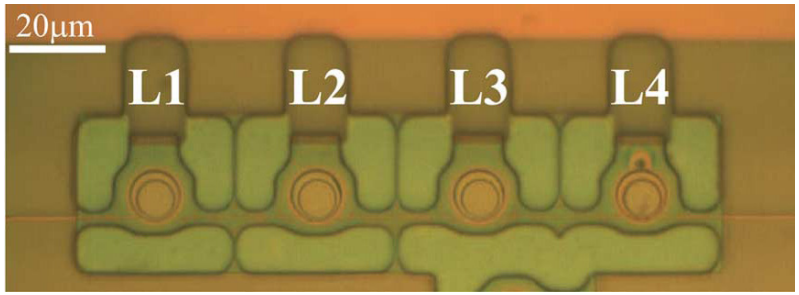

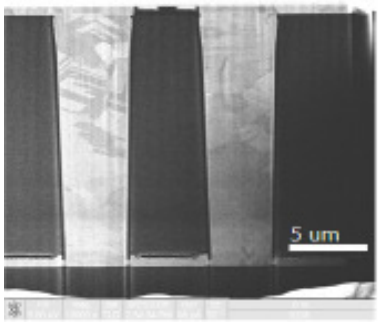

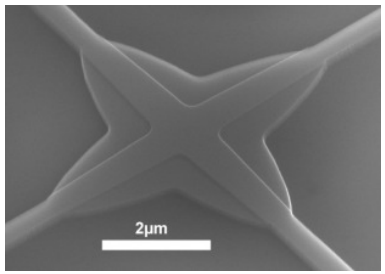

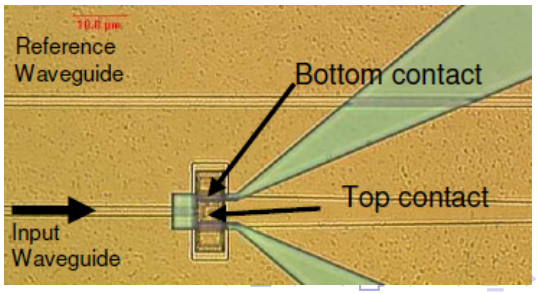

Area overhead estimation

- Assumptions :
  - $VC_{area} = 10\mu m \times 18.5\mu m$
  - initial electrical die size : $491 mm^2$ for 256 cores [11]
  - 2 electrical layers

Area overhead estimation

- Assumptions :
  - $VC_{area} = 10\mu m \times 18.5\mu m$
  - initial electrical die size : $491 mm^2$ for 256 cores [11]
  - 2 electrical layers

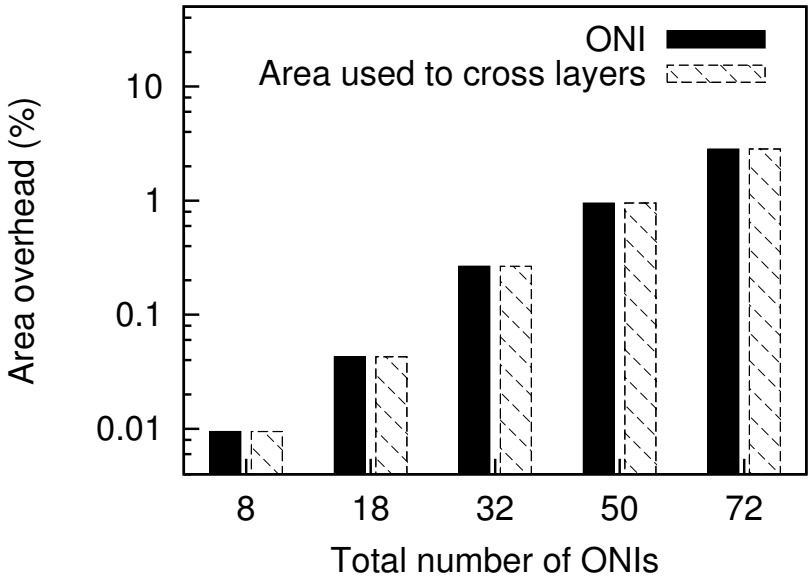

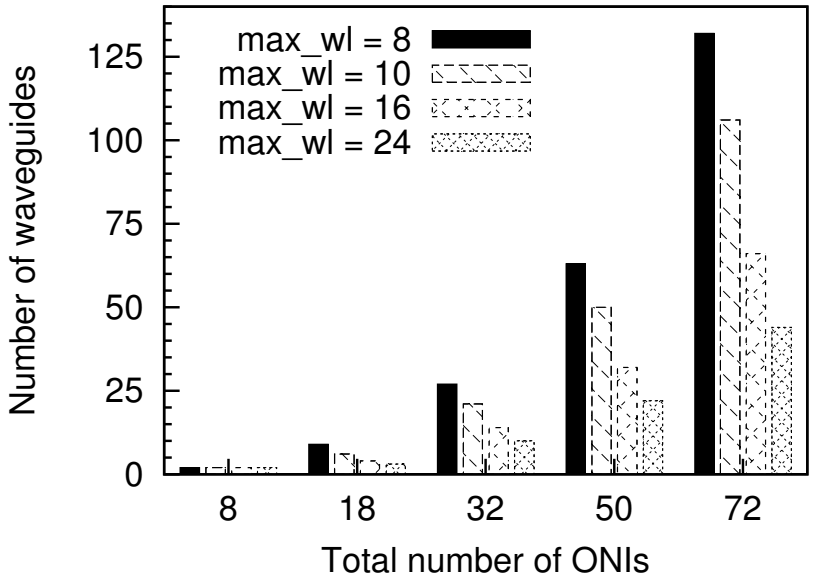

Area overhead estimation

- Assumptions :
  - $VC_{area} = 10\mu m \times 18.5\mu m$
  - initial electrical die size : $491 mm^2$ for 256 cores [11]
  - 2 electrical layers

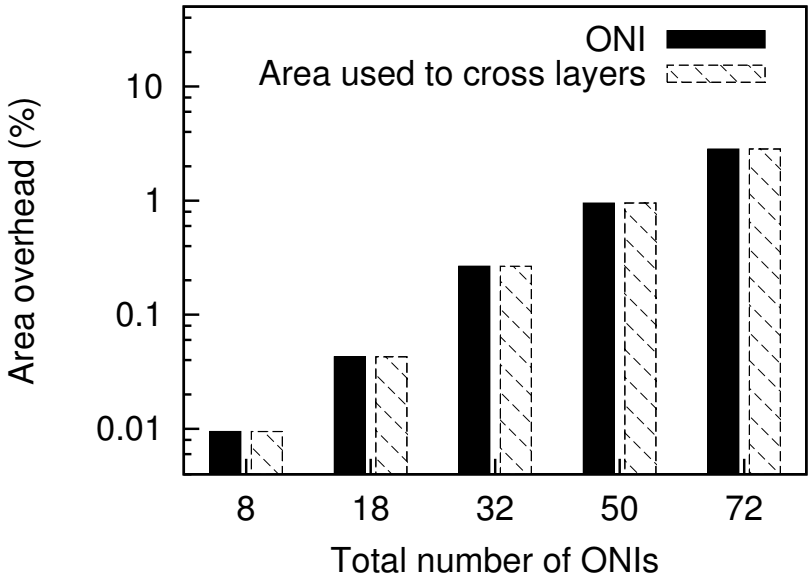

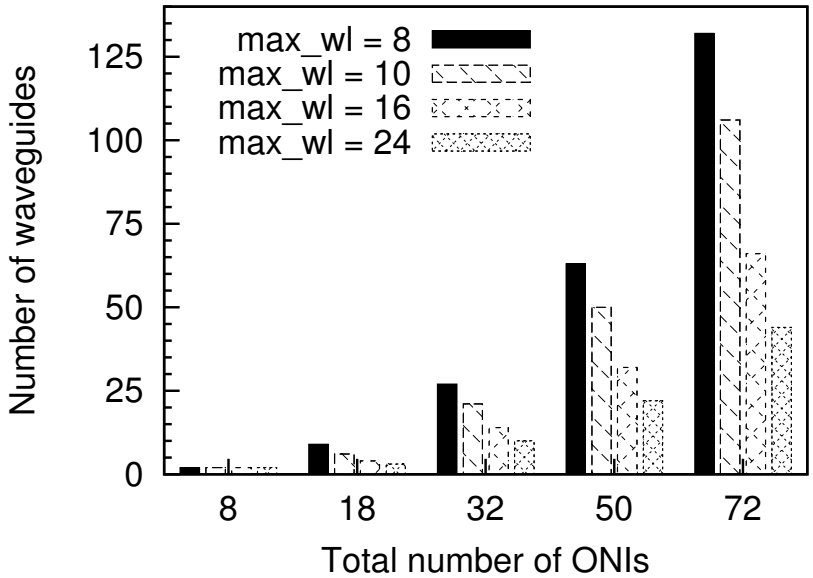

- 32 ONIs $\rightarrow$ 4×4 ONIs per electrical layer
  - each ONI is shared by 16 cores
  - 0.5% area overhead
  - reliability-complexity design tradeoff

Conclusion and Future Work

- ORNoC
  - suitable for 3D architectures (communication hierarchy)
  - layout guidelines → regularity in x, y and z dimensions
  - ONI area overhead estimation
    - e.g. : 36 ONIs, 2 electrical layers → 0.5% area overhead

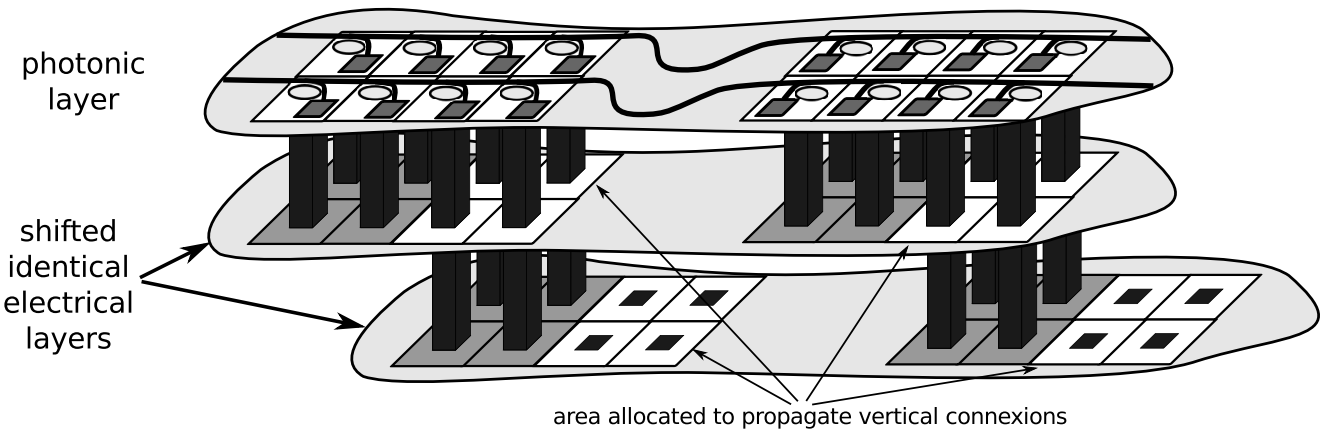

- Future work
  - loss modeling and estimation
  - laser power control
  - SERDES footprint

References

R. G. Beausoleil, J. Ahn, N. Binkert, A. Davis, D. Fattal, M. Fiorentino, N. P. Jouppi, M. McLaren, C. M. Santori, R. S. Schreiber, S. M. Spillane, D. Vantrease, and Q. Xu.
A Nanophotonic Interconnect for High-Performance Many-Core Computation.
In *Proceedings of the 16th IEEE Symposium on High Performance Interconnects*, pages 182–189, 2008.

M. Brière, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, and I. O'Connor.
System Level Assessment of an Optical NoC in an MPSoC Platform.
In *Proceedings of the conference on Design, Automation and Test in Europe*, DATE, pages 1084–1089, 2007.

Jason Cong and Yan Zhang.
Thermal-driven multilevel routing for 3-D ICs.
In *Proceedings of the Asia and South Pacific Design Automation Conference*, ASP-DAC, pages 121–126, 2005.

Huaxi Gu, Jiang Xu, and Zheng Wang.
A Novel Optical Mesh Network-on-Chip for Gigascale Systems-on-Chip.
In *Proceedings of the IEEE Asia Pacific Conference on Circuits and Systems*, pages 1728–1731, December 2008.

Huaxi Gu, Jiang Xu, and Wei Zhang.
A Low-Power fat Tree-Based Optical Network-on-Chip for Multiprocessor System-on-Chip.
In *Proceedings of the conference on Design, Automation and Test in Europe*, DATE, pages 3–8, 2009.

George Kurian, Jason E. Miller, James Psota, Jonathan Eastep, Jifeng Liu, Jurgen Michel, Lionel C. Kimerling, and Anant Agarwal.
ATAC : a 1000-Core Cache-Coherent Processor with on-Chip Optical Network.
In *Proceedings of the 19th international conference on Parallel Architectures and Compilation Techniques*, PACT, pages 477–488, 2010.

David Navarro, Matthieu Briere, Ian O'Connor, Fabien Mieyeville, Frédéric Gaffiot, and Laurent Carrel.
Quantitative Study of Area and Power Consumption Costs for 3 Gbits/s Optical Communications in a 0.13μm CMOS Circuit.

Optical Network Interface (ONI)

- Operation mode :
  - ejection
  - pass through
  - injection

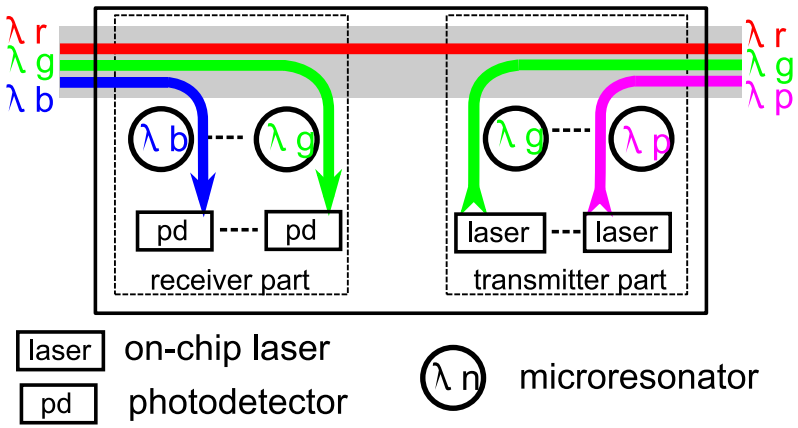

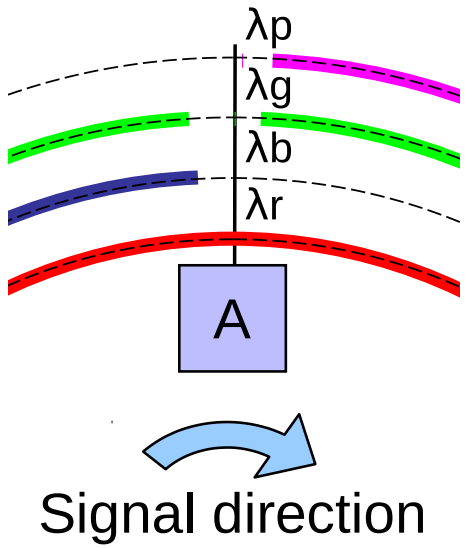

How to configure ONI ?

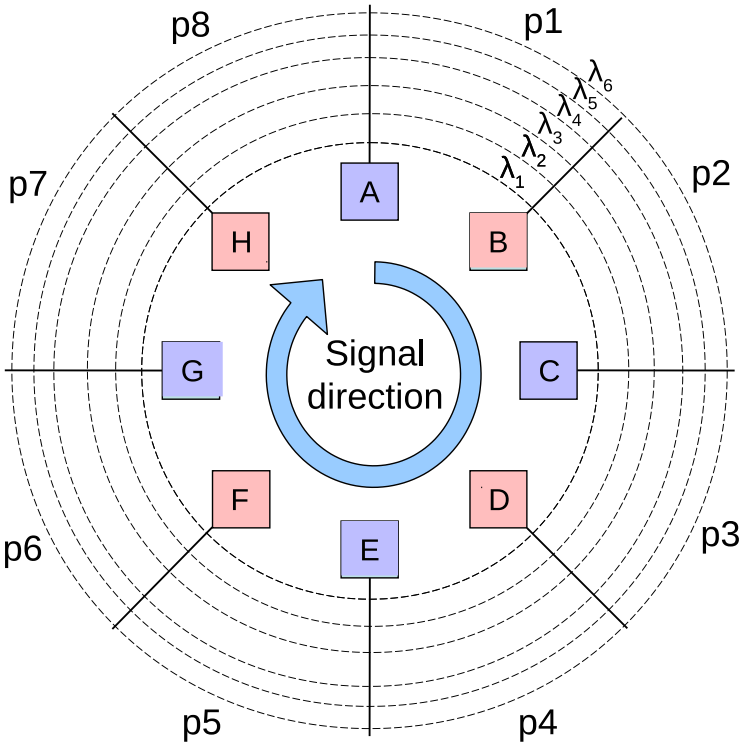

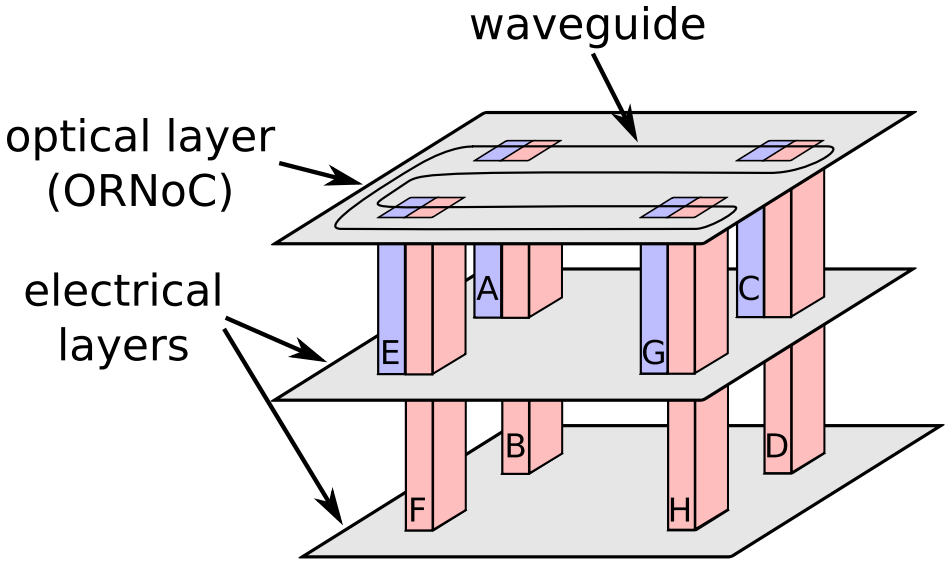

- Which waveguide partition ?
- Which wavelength ?

How to configure ONI ?

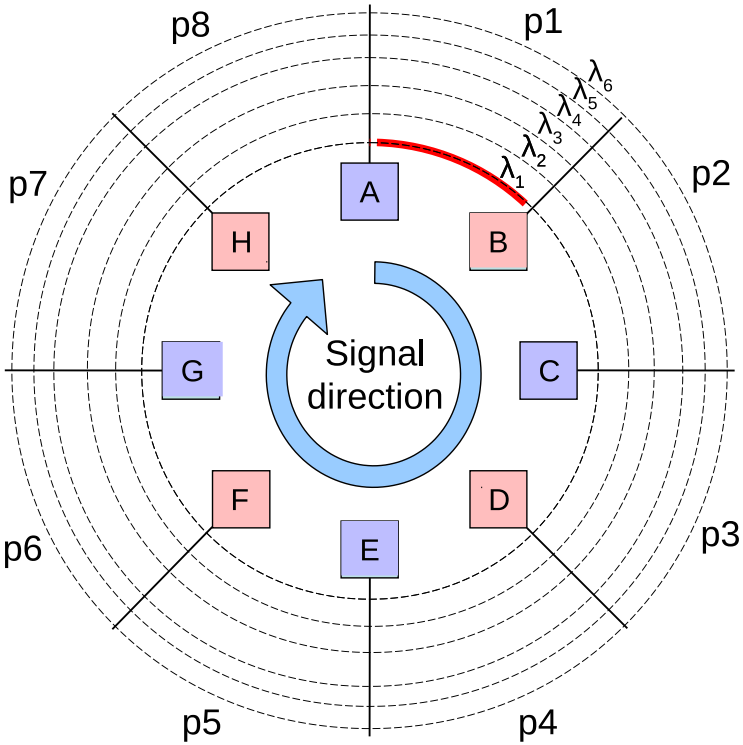

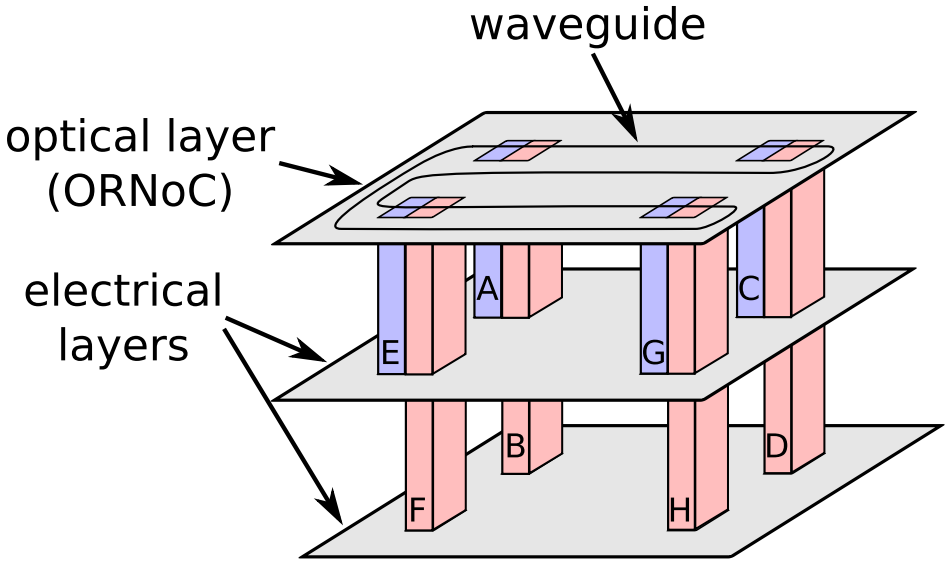

- Which waveguide partition ?
- Which wavelength ?

How to configure ONI ?

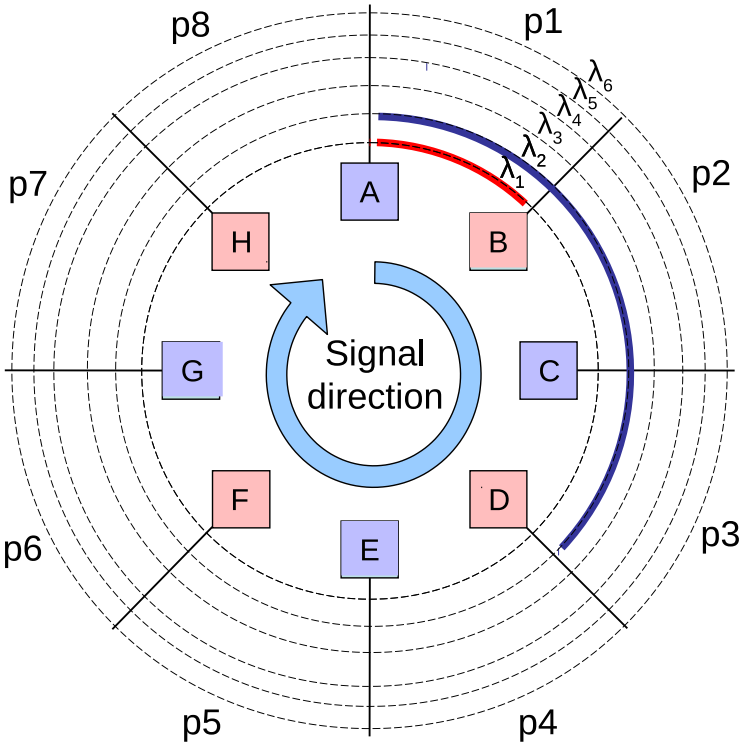

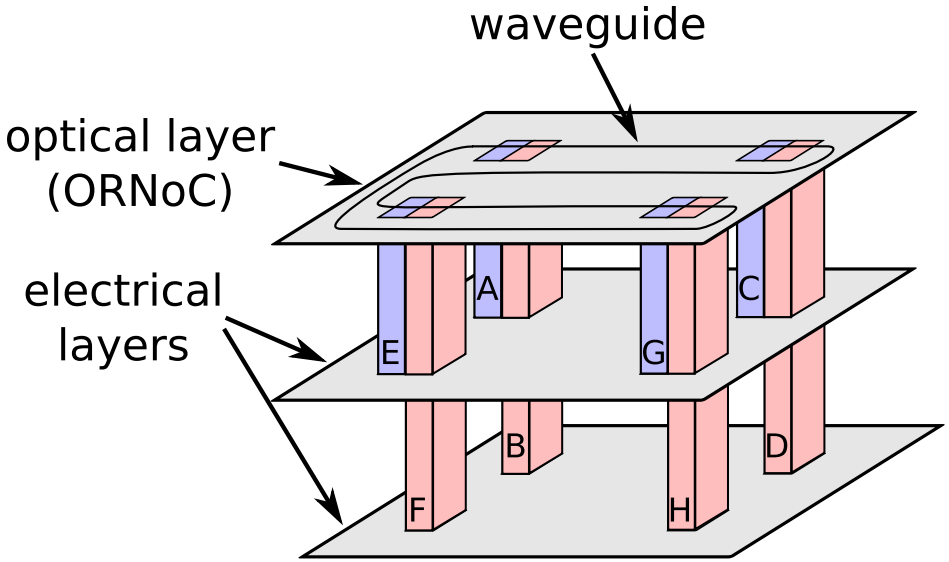

- Which waveguide partition ?
- Which wavelength ?

How to configure ONI ?

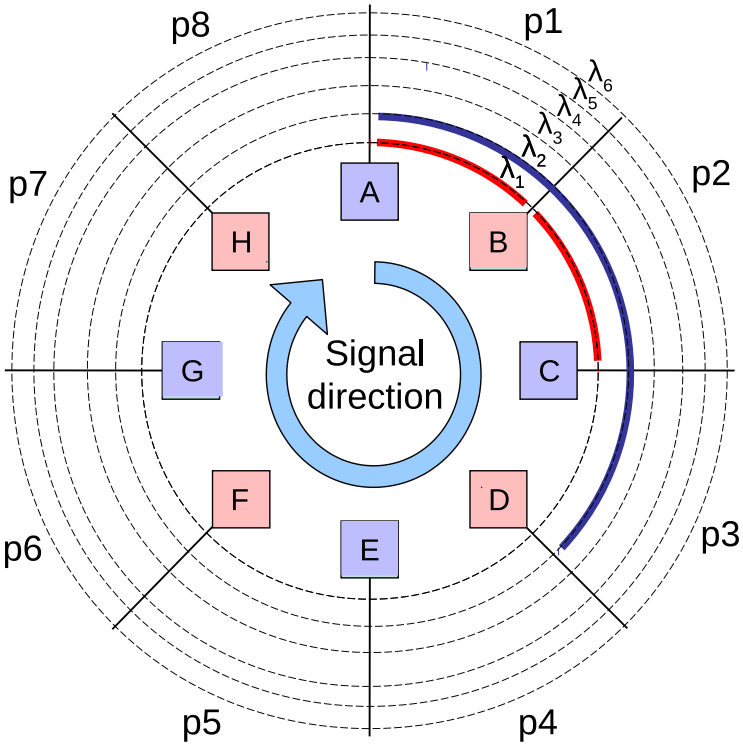

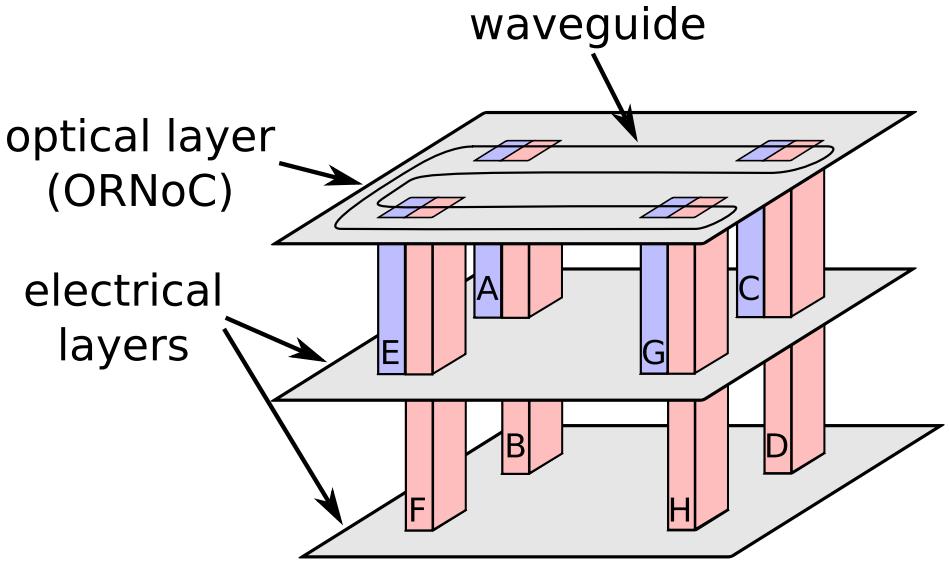

- Which waveguide partition ?
- Which wavelength ?

How to configure ONI ?

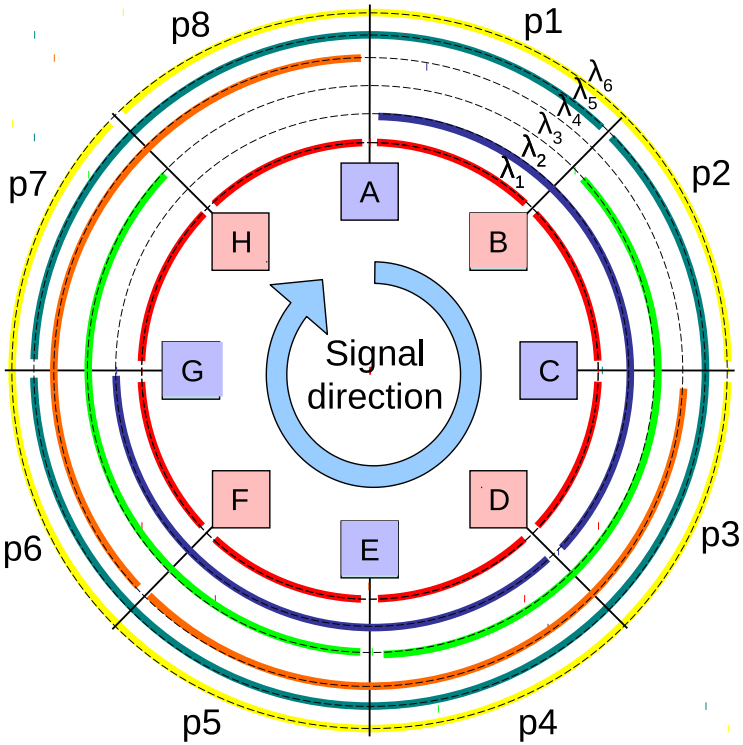

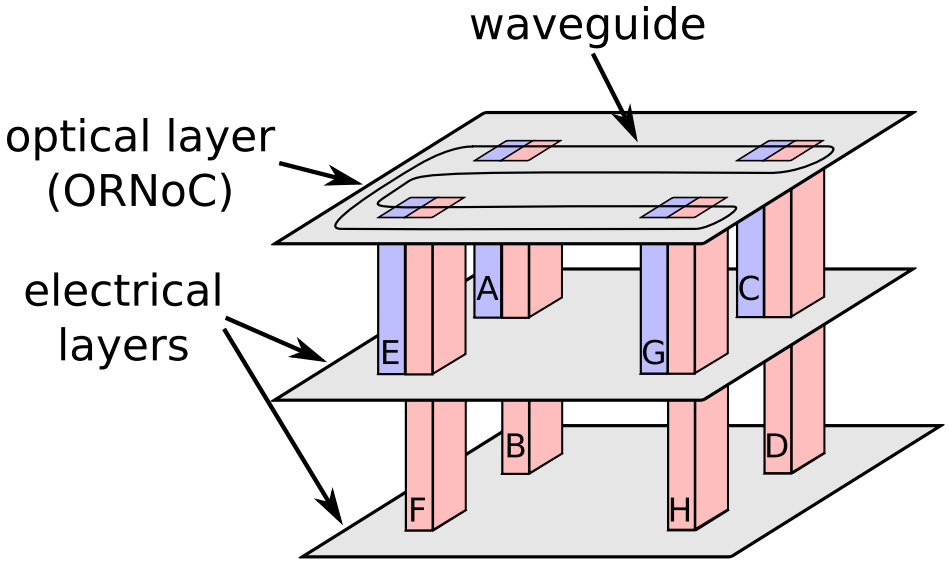

- ORNoC benefits : contention free, scalable, low power

How to configure ONI ?

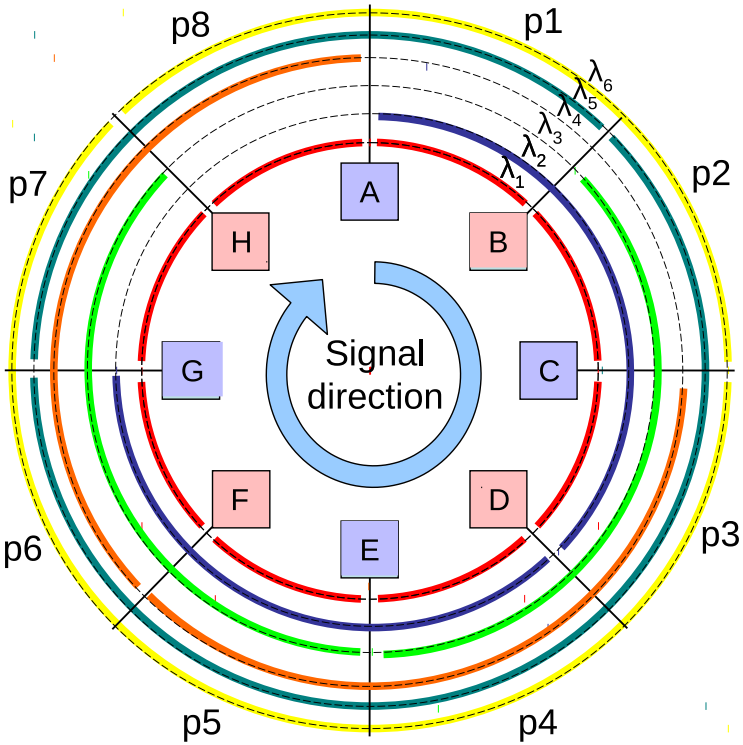

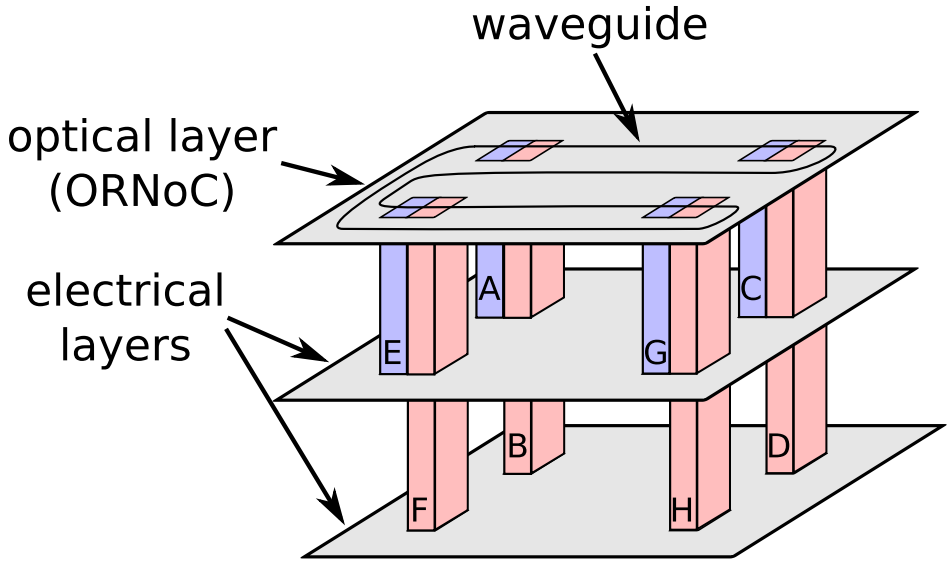

- ORNoC benefits : contention free, scalable, low power
- Need for...
  - (high level) layout guidelines
  - ONI area overhead estimation
  - regularity and reuse (x, y and z dimensions)

ORNoC in 2D Architecture

- Communication hierarchy :
  - Electrical NoC → intra-cluster communication
  - ORNoC → inter-cluster communication

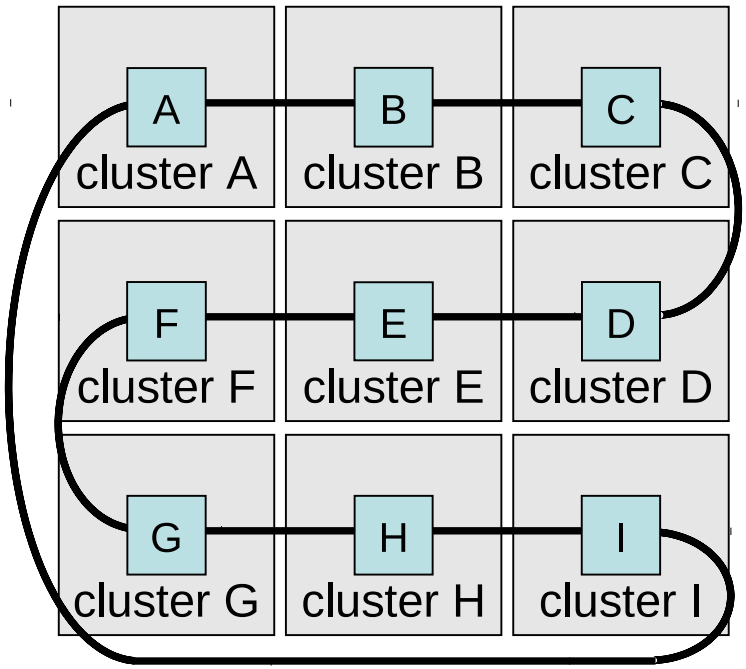

2D architecture

ORNoC in 2D Architecture

- Communication hierarchy :
  - Electrical NoC → intra-cluster communication
  - ORNoC → inter-cluster communication

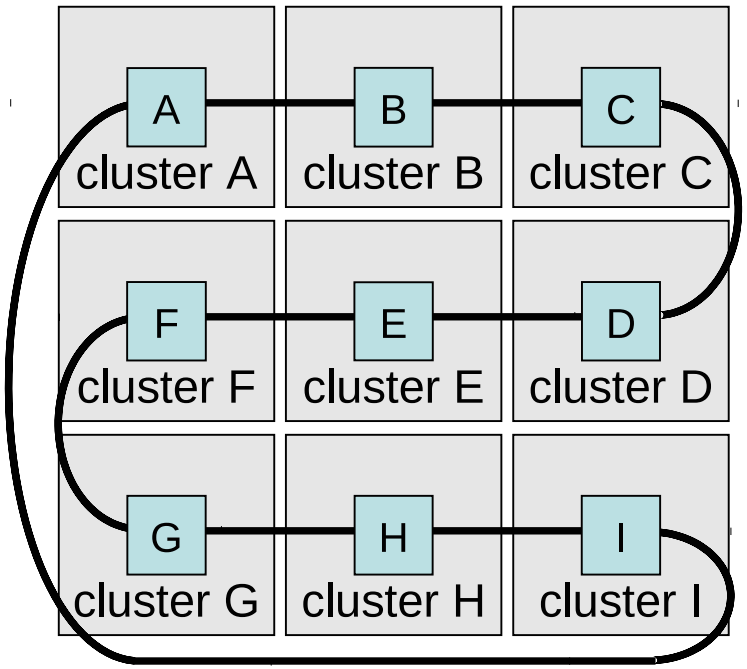

2D architecture

	A	B	C	D	E	F	G	H	I
A	0	1	1	1	1	1	1	1	1
B	1	0	1	1	1	1	1	1	1
C	1	1	0	1	1	1	1	1	1
D	1	1	1	0	1	1	1	1	1
E	1	1	1	1	0	1	1	1	1
F	1	1	1	1	1	0	1	1	1
G	1	1	1	1	1	1	0	1	1
H	1	1	1	1	1	1	1	0	1
I	1	1	1	1	1	1	1	1	0

Connectivity Matrix :
0 → no communication
1 → communication required