MPSoC Research Group, **Engineering Department, University of Ferrara** (ITALY)

EDA beyond its electronic roots: toward a synthesis methodology for wavelength-routed optical networks-on-chip

Mahdi Tala, Marco Balboni, Marco Castellari, Andrea Peano, Maddalena Nonato and Davide Bertozzi

IWES 2016 Italian Workshop on Embedded Systems

Pisa, Italy 19-20 September 2016

Contact email:
mahdi.tala@unife.it

This work is a collaboration with:

- University of Zaragoza;
- TU Munich.

A New Technology Substrate for Networking

NoC-connected multi- and many-core architectures both in HPC and embedded

Intel Xeon Phi
Knights Landing
72 cores
Intel 2016

Tile-Gx72
72 cores
Tilera 2013

Kalray MPPA
256 cores
Kalray Inc. 2013

- Future system scalability trends will bring several network-on-chip (NoC) concerns to the forefront:
  - *Latency-dominated and distance-sensitive performance.*
  - *Interconnection energy does not keep up with the scalability trend of computation energy.*
  - *Extension to off-chip communication is not seamless.*

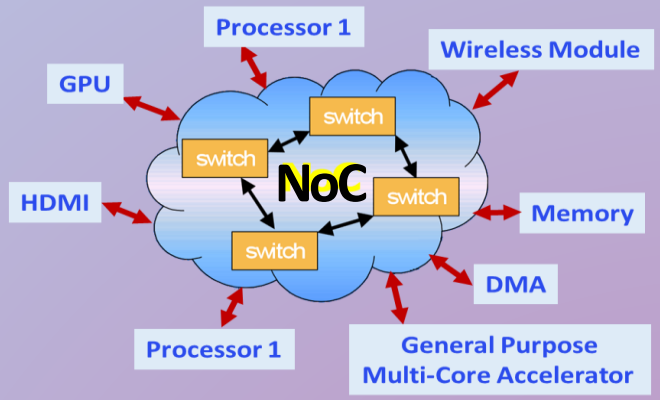

One promising solution consists of networking communication actors by means of an emerging interconnect technology:
Silicon Photonics stands out as the most promising candidate…

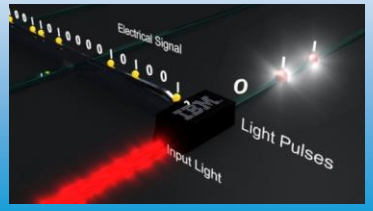

… due to intrinsic capability of the light to transport information over large distances at very high data rates and low latency with minor dynamic power dissipation.

Optical NoC Design space: largely unknown

Early stage of optical NoCs:

- Research efforts to overcome the technology barriers.
- Hardwired and loosely connected architecture design points based on designer's intuition.

Outcome: proof-of-concept ONoC architectures!

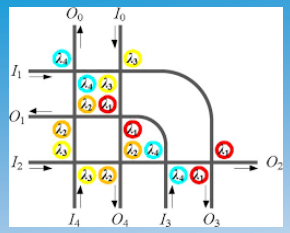

GWOR

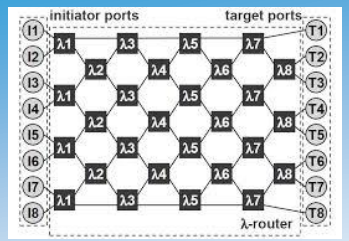

λ-Router

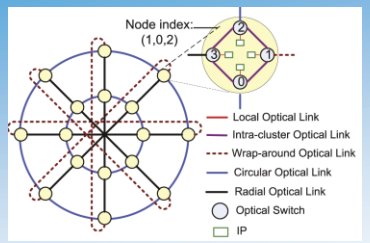

2D-HERT

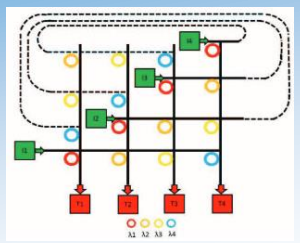

Snake

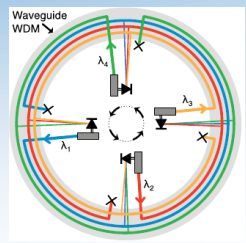

Ring

The gap with the system-level designer…

… is currently still huge!

Very few choices available for system-level architecture design. Limited design flexibility, limited room for design optimization.

How does the whole design space look like?

Major Requirement: starting from a high-level description of the interconnection system, operate on communication abstractions and refine them into an actual implementation with components from a technology library.

EDA BEYOND ITS «E-ROOTS»

This is exactly what typical EDA flows do!

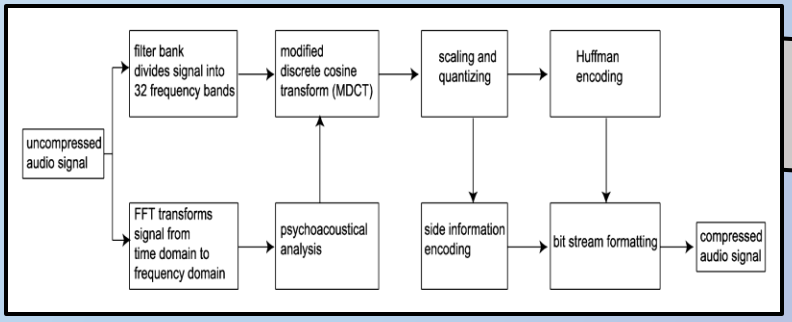

High-level specification

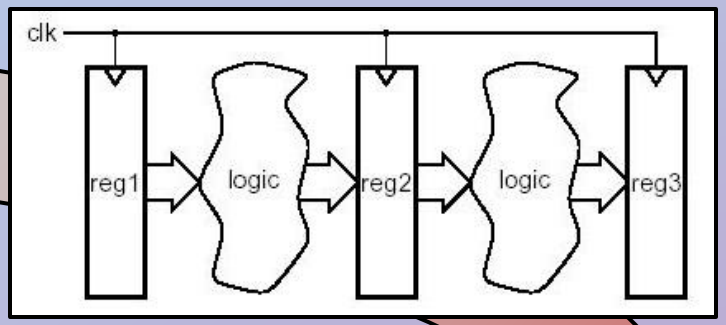

Register-Transfer Level

EDA tools can automatically transform a complex system-on-chip design from high-level functional description to a detailed geometric description across all of the intervening levels of abstraction.

Technology-Independent Logic Library

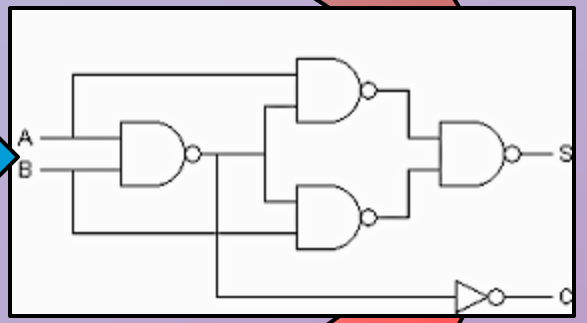

Gate-Level Netlist

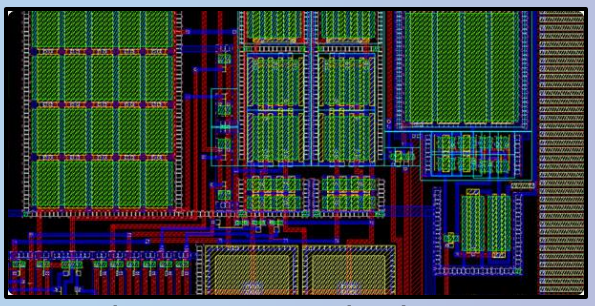

Planar geometric shapes

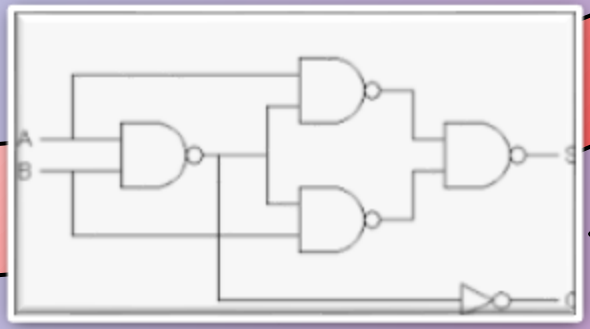

Mapped Gate-Level Netlist

Technology-Library

EDA BEYOND ITS «E-ROOTS»

Can we extend the paradigms and methodologies of EDA to the context of emerging silicon photonic interconnect?

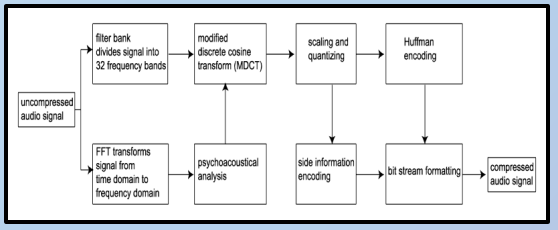

High-level specification

Technology-independent Logic Library

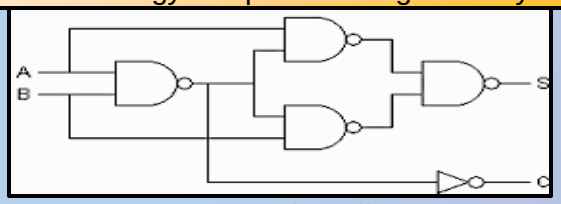

Gate-Level Netlist

Technology Library

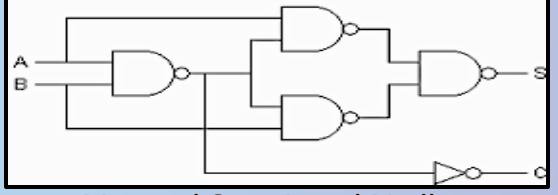

Mapped Gate-Level Netlist

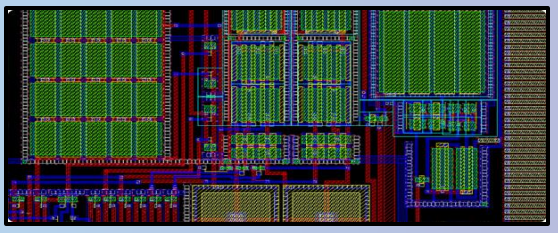

Planar geometric shapes

Can we identify equivalent steps for the synthesis of an Optical Network-on-Chip?

WAVELENGTH-ROUTED OPTICAL NOCS (WRONOCS)

The choice of this work is **Wavelength-selective Routing**, i.e., all-to-all contention-free connectivity.

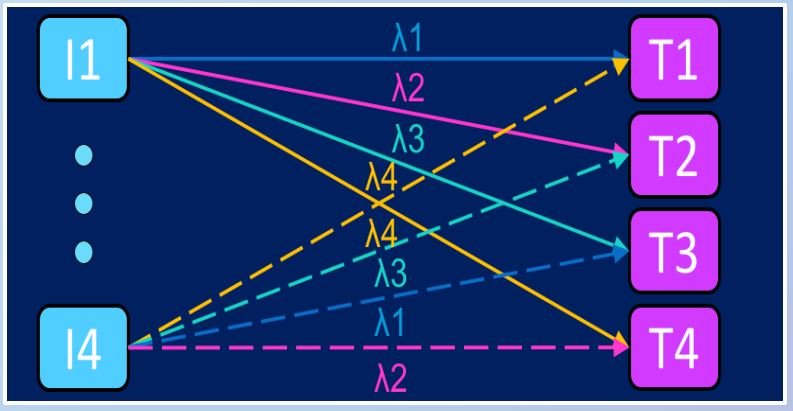

PROS

- ✓ No time is spent in routing/decoding and arbitration.
- ✓ High communication performance predictability.

CONS

- x Hard to scale to a large number of cores.

The wavelength- selective routing function is ultimately fulfilled by means of two abstract primitives:

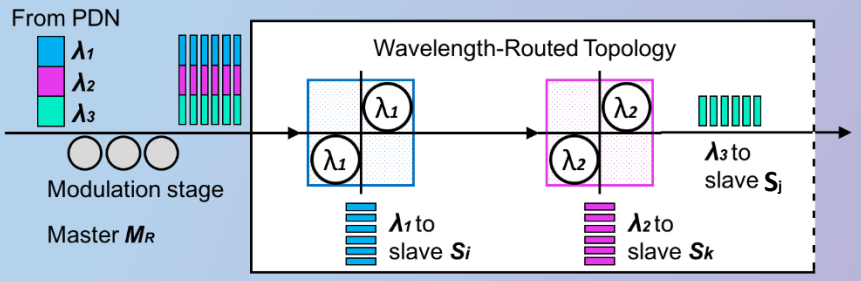

Drop function at initiator's side.

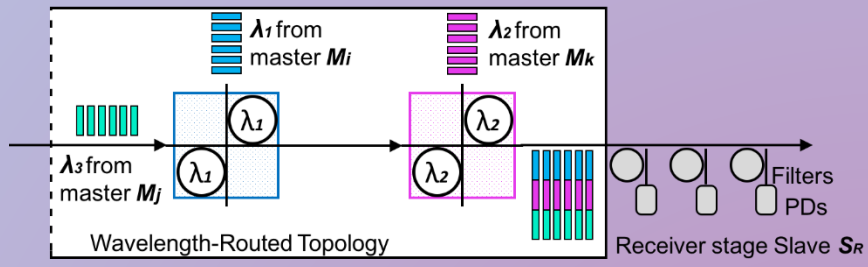

Add function at target's side.

Basic building block for the implementation of the add and of the drop function: the **1x2 ADD/DROP FILTER**

On-resonance signal

Off-resonance signal

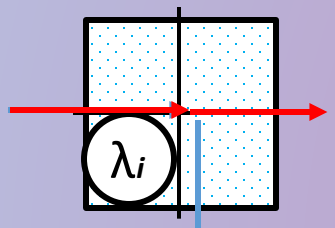

Synthesis Methodology

We come up with a six-step synthesis methodology of WRONoC topologies

I. Wavelength Resolution

II. Technology Mapping

III. Wavelength Assignment

IV. Topology Connection

V. Device Parameter Selection

VI. Placement and Routing

SYNTHESIS METHODOLOGY

1. Wavelength Resolution
   i. WDM input signal from masters should undergo the drop function by going through *n − 1* ADFs, when assuming the connectivity of *n* masters with *n* slaves.
   ii. The resolution pattern may differ in each row
   iii. On each row there must be always one wavelength channel which is not on resonance with any of the ADFs.

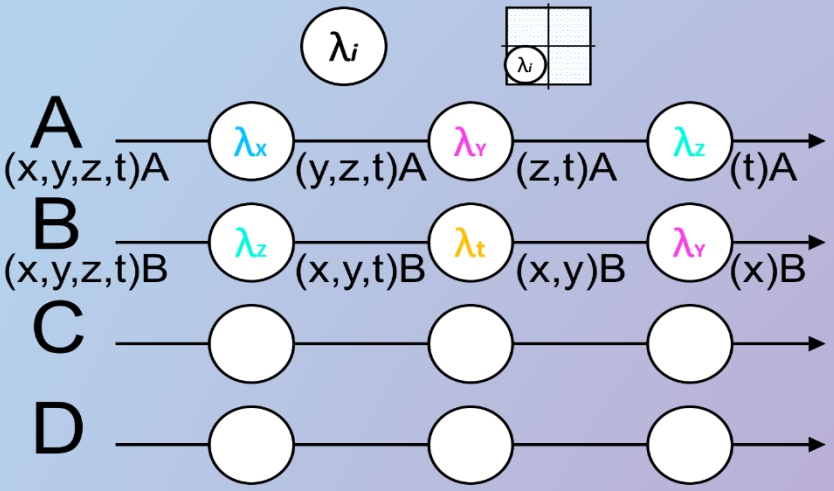

Wavelength Resolution Graph ***(WRG)*** *for a generic 4x4 WRONoC.*

Fixing the parameters of the WRG indirectly prunes the design space and biases the methodology toward a specific topology design point.

In order for the topology design point to be legal, one constraint should be enforced on the WRG:

✓ the **wavelengths of the ADFs along a single row should be different from each other!**

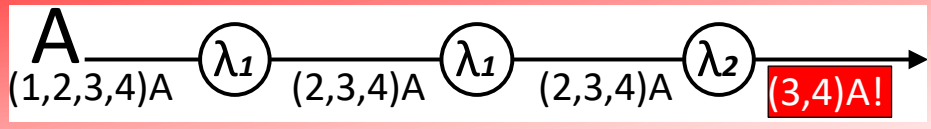

Two or more unresolved channels from the same master reach one output of the topology!

Side effect: one or more slaves will be unreachable from that master!!

SYNTHESIS METHODOLOGY

2. Technology Mapping

✓ Key Idea: a common technology mapping consists of grouping the 1x2 ADFs into compact 2x2 photonic switching elements (PSEs)!

! Mapping Constraint:

- **Legal mappings are the combinations without repetitions**

C(n,2)

for picking 2 unordered outcomes from *n* possibilities.

(…assuming 4 masters and 4 slaves, C(4,2)=6 possible mappings)

AB,AC,AD,BC,BD,CD

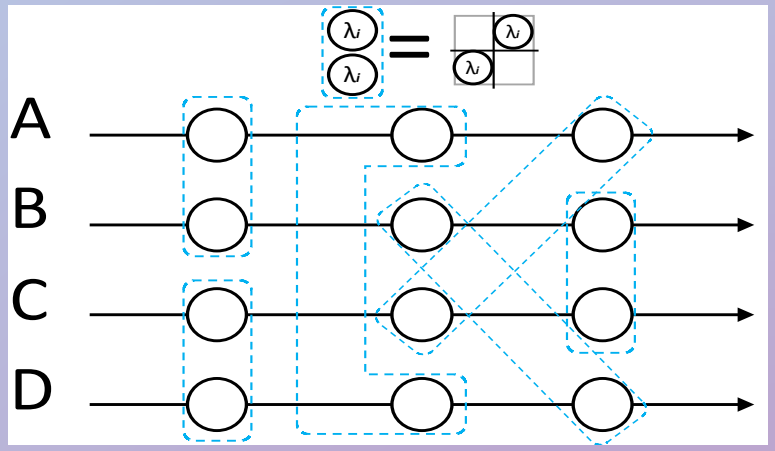

Legal technology mapping example

✓ Legal mapping:

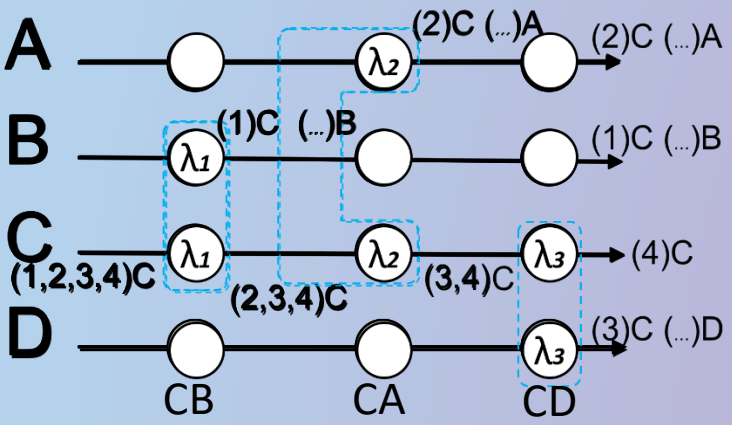

An efficient add function should couple a distinct wavelength channel onto each row.

SYNTHESIS METHODOLOGY

2. Technology Mapping

✓ Key Idea: a common technology mapping consists of grouping the 1x2 ADFs into compact 2x2 photonic switching elements (PSEs)!

There are $n - 1$ ADFs on the associated WRG row for each master.

→ $(n - 1)^2$ ways to group an ADF filter of A with an ADF filter of B.

→ only 1 should be selected for each topology design point.

! Mapping Constraint:

- **Legal mappings are the combinations without repetitions**

$C(n,2)$

for picking 2 unordered outcomes from n possibilities.

(…assuming 4 masters and 4 slaves, C(4,2)=6 possible mappings)

AB,AC,AD,BC,BD,CD

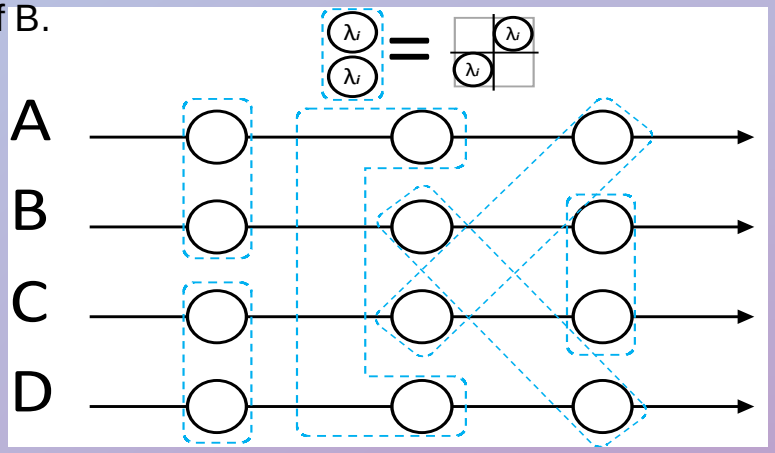

Legal technology mapping example

x Illegal mapping:

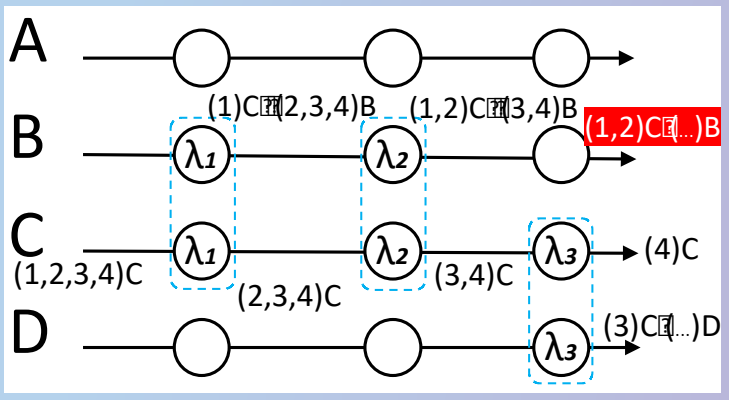

If the WDM signal from a specific master is coupled onto another row twice (despite the two PSEs are correctly tuned to different wavelengths), the dropped channels end up being recombined onto the target row.

→ A few destinations end up being unreachable

SYNTHESIS METHODOLOGY

2. Technology Mapping

✓ Key Idea: a common technology mapping consists of grouping the 1x2 ADFs into compact 2x2 photonic switching elements (PSEs)!

There are $n - 1$ ADFs on the associated WRG row for each master.

→ $(n - 1)^2$ ways to group an ADF filter of A with an ADF filter of B.

→ only 1 should be selected for each topology design point.

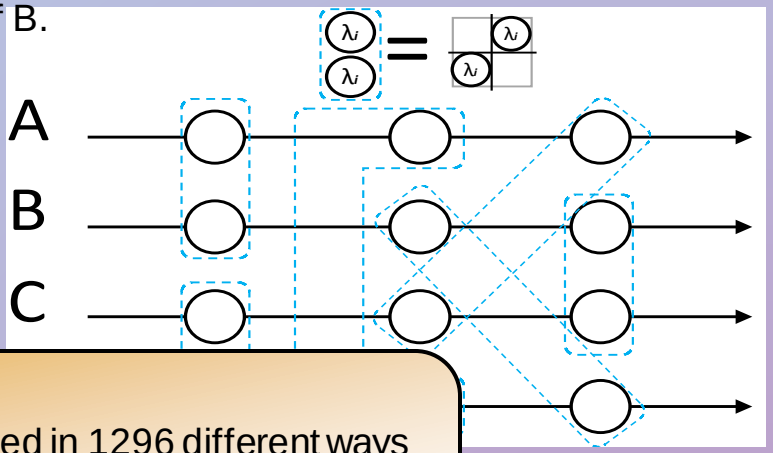

ping example

! Mapping Constraint:

- **Legal mappings are the combinations without repetitions**

$C(n,2)$

for picking 2 unordered outcomes from n possibilities.

(…assuming 4 masters and 4 slaves, C(4,2)=6 possible mappings)

A 4x4 WRONoC topology can be implemented in 1296 different ways

For n > 4, the full enumeration of all possible design points becomes computationally unaffordable.

x Illegal mappin

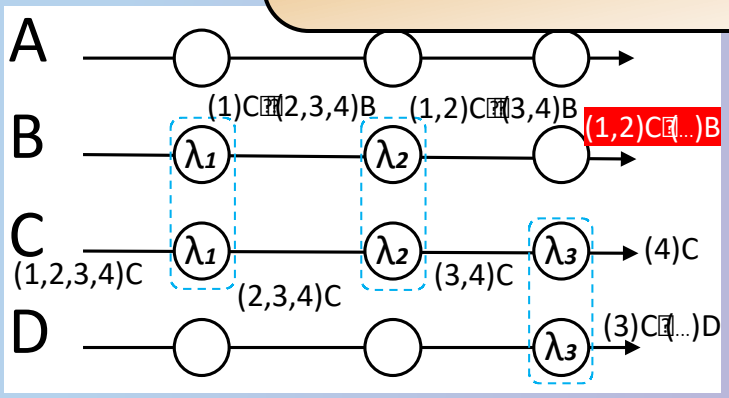

If the WDM signal from a specific master is coupled onto another row twice (despite the two PSEs are correctly tuned to different wavelengths), the dropped channels end up being recombined onto the target row.

→ A few destinations end up being unreachable

SYNTHESIS METHODOLOGY

3. Wavelength Assignment

✓ Key Idea: each master needs 1 distinct channel to communicate with each one of the *n* slaves, and such *n* wavelengths are reused across masters.

!! Assumption: we do not consider real wavelengths at this stage (e.g., 1550 nm) but symbolic ones, such as *λ1* to *λn* in an *n×n* WRONoC.

- Greedy algorithm to perform complete wavelength assignment:

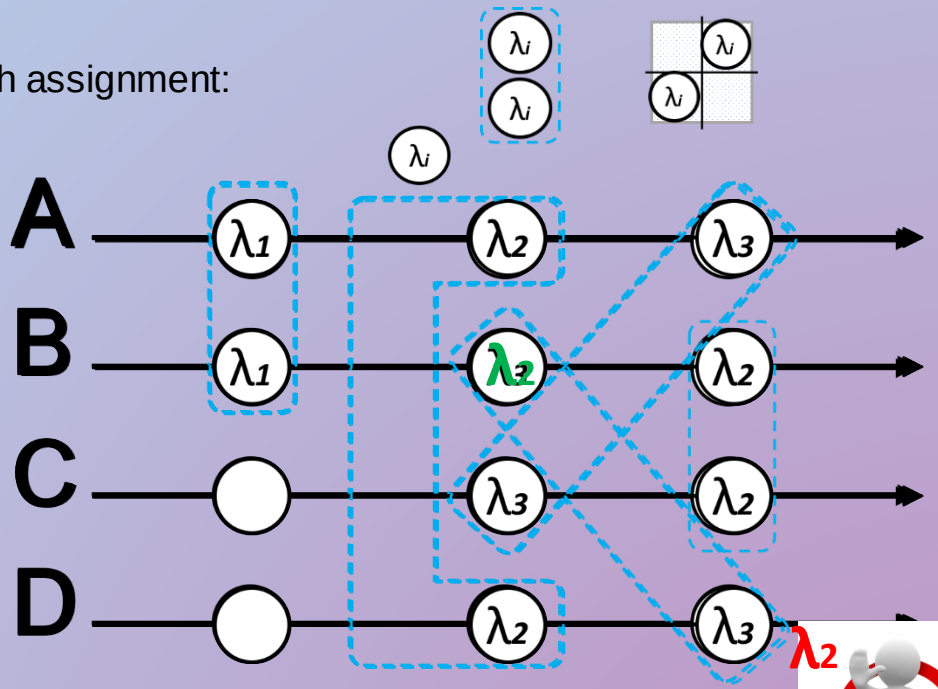

- ✓ Mutually-exclusive property of resonant wavelengths should hold on both companion rows.
- ✓ When multiple wavelength options are available for assignment, the algorithm selects the one with lowest identifier.

Under this assumptions, the algorithm typically optimizes the number of microring resonator types.

SYNTHESIS METHODOLOGY

4. Topology Connection

✓ Key Idea: drawing the connectivity pattern of the topology, as well as the features of its PSEs, from the annotated WRG.

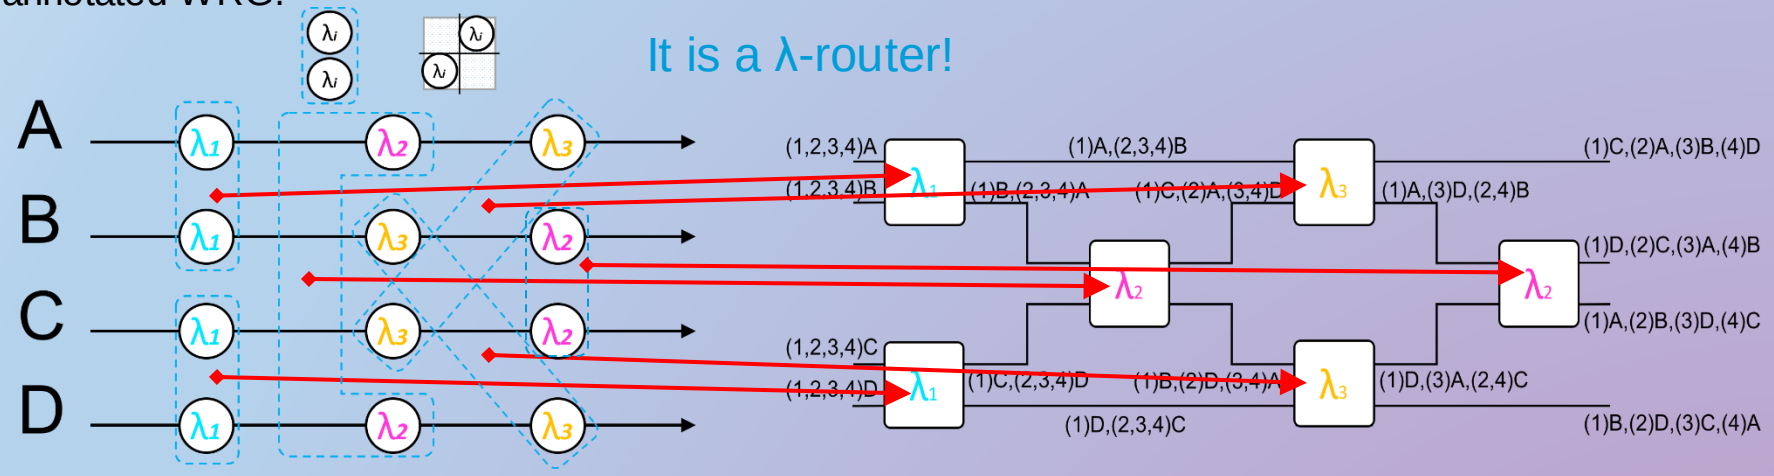

The exact connectivity pattern of the λ-router is provided!!

There is only a minor (yet relevant) difference

- With our wavelength assignment strategy, the last stage is tuned to λ2 because of its preferred choice for minimum wavelength channel identifiers.
- Baseline λ-router: the last stage is tuned to λ4

→ usage of 4 resonator types instead of 3.

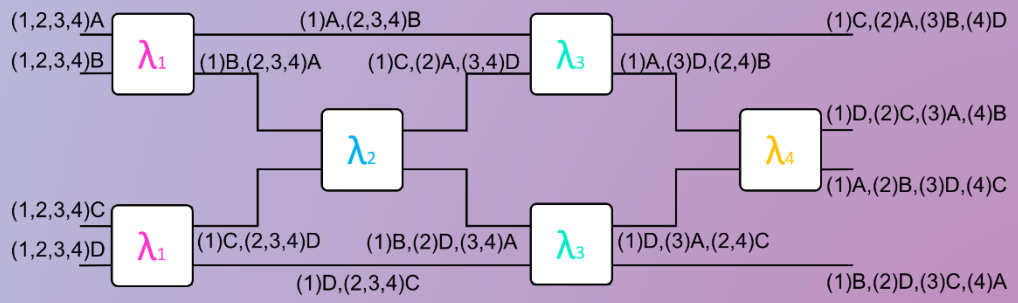

SYNTHESIS METHODOLOGY

4. Topology Connection

✓ Key Idea: drawing the connectivity pattern of the topology, as well as the features of its PSEs, from the annotated WRG.

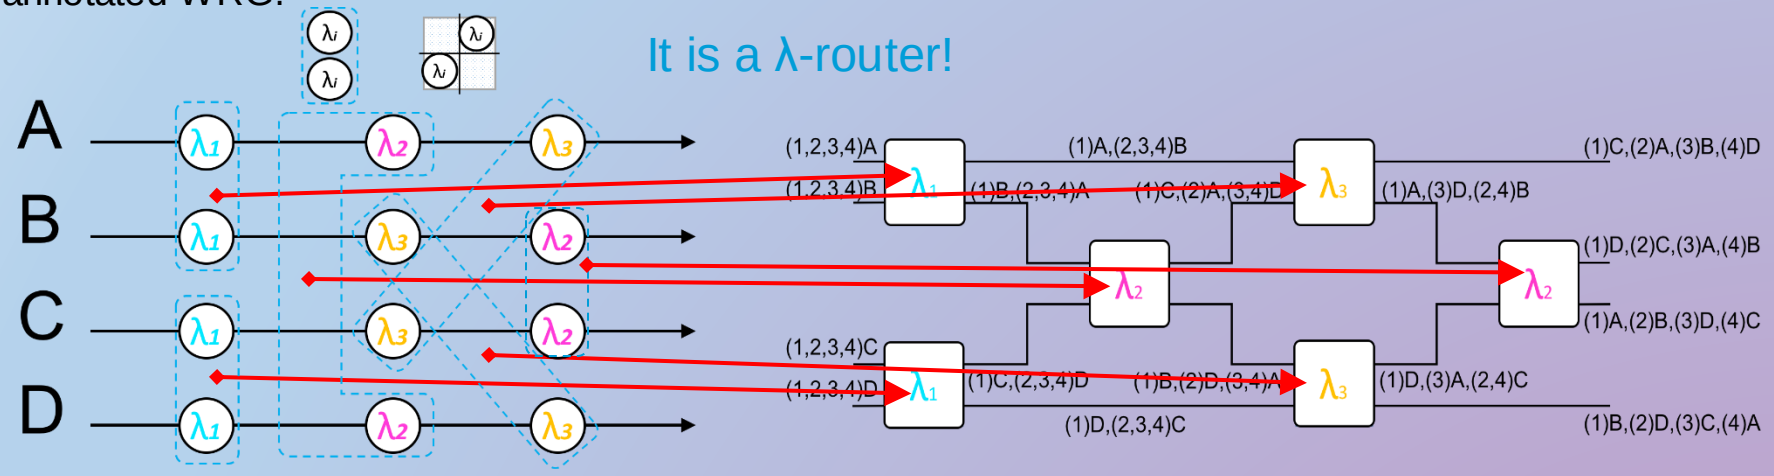

The exact connectivity pattern of the λ-router is provided!!

There is only a minor (yet relevant) difference

- With c… ice for mini…
- Base…

→ usage…

Our synthesis methodology can potentially populate the complete design space of WRONoC topologies by spanning technology mapping and wavelength assignment strategies!!

Actual Physical Channels and Micro-Ring Resonator Radii Assignment

Symbolic channels are then changed into actual physical channels by selecting a wavelength carrier in the band 1500-1600 nm

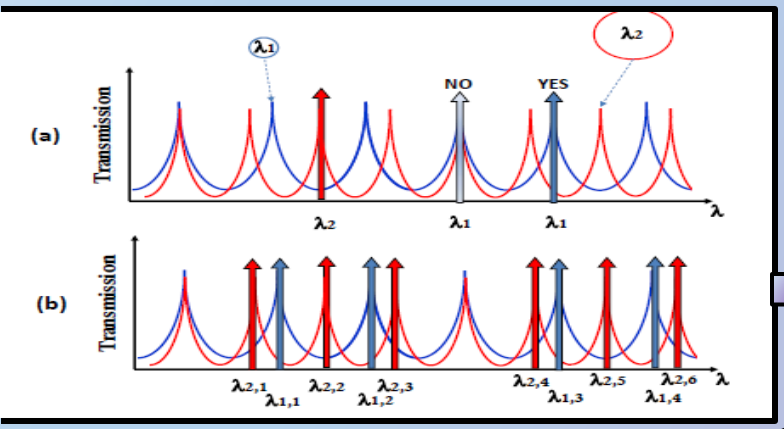

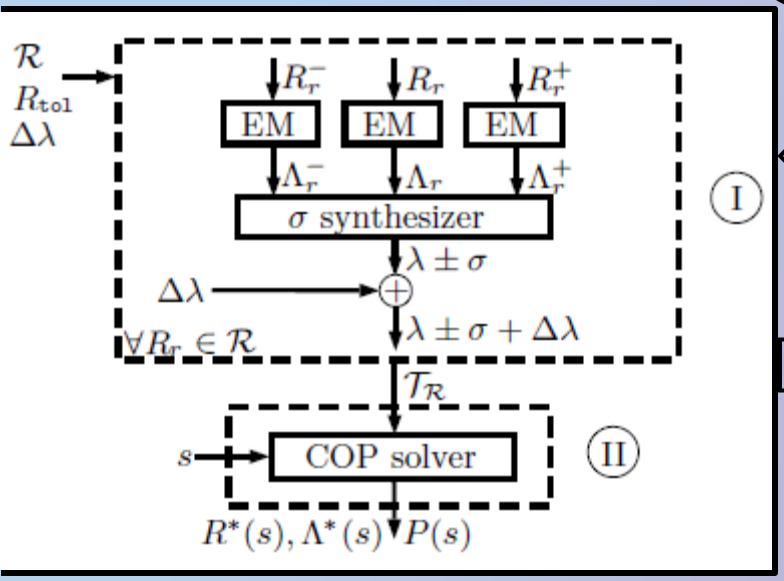

Routing-Fault Free Wavelength Selection

Maximum Parallelism of Comunication Flows

By mean of constrained optimization Problem model.

- **Legal parameter assignments (wavelength and ring diameter values).**
- **The maximum level of parallelism that can be achieved on the topology.**
  - **fabrication process uncertainty has been taken into account.**

Experimental Results

What levels of connectivity and communication parallelism can we achieve?

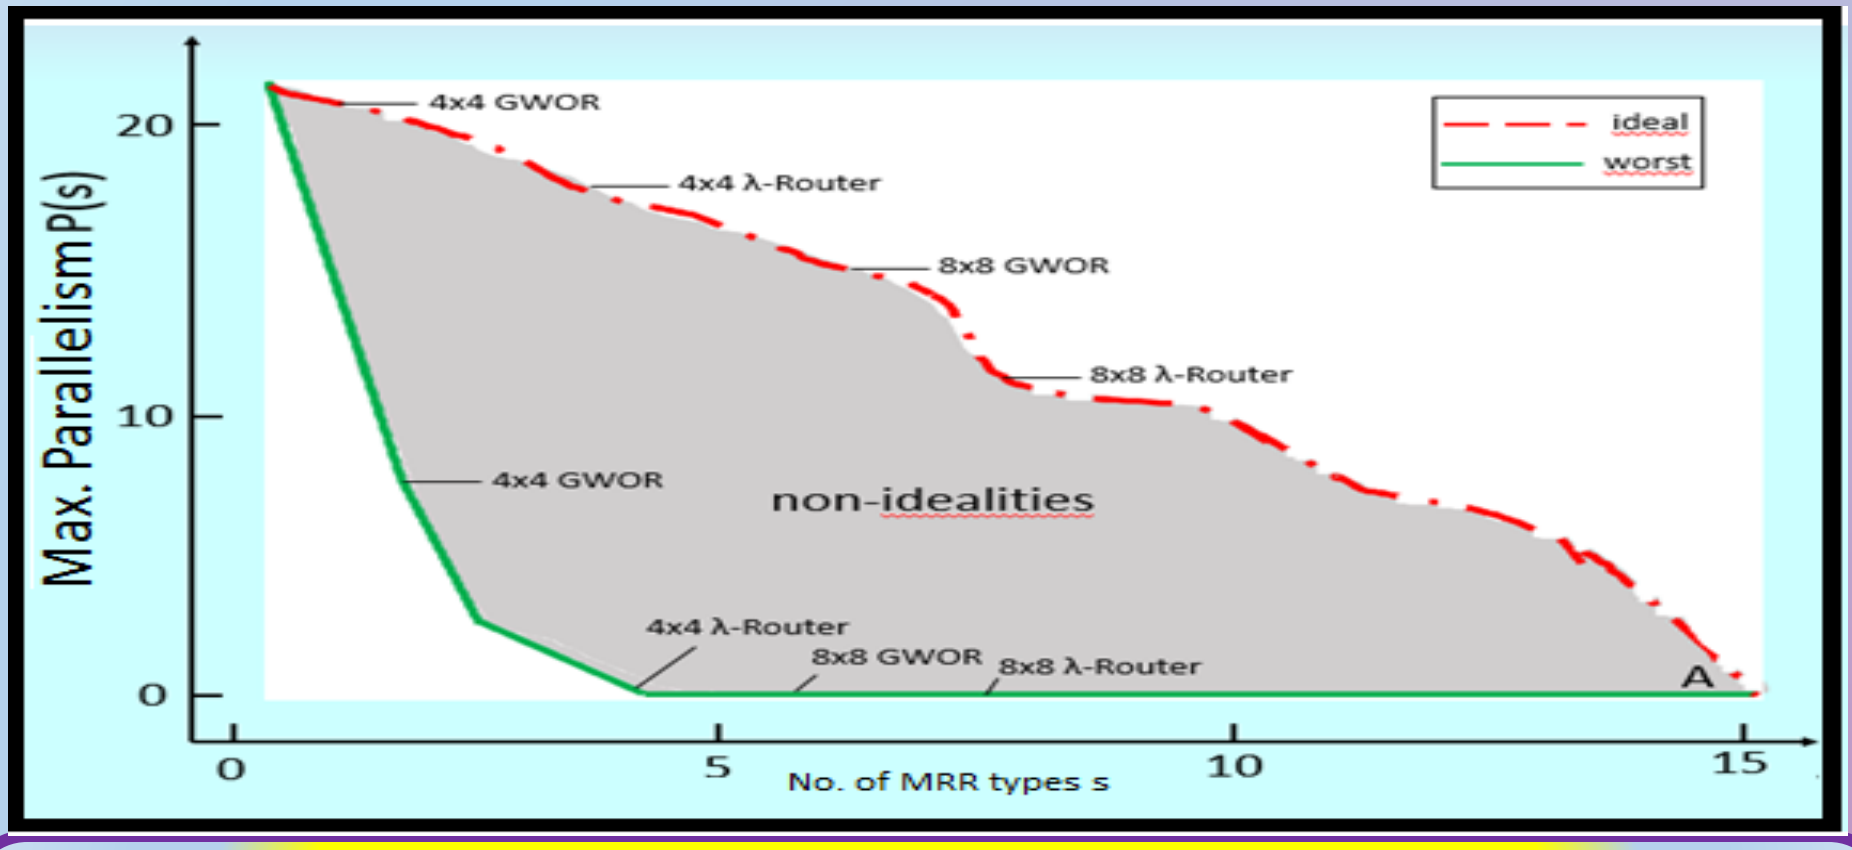

The maximum communication parallelism P(s):

(a) scales down by increasing the number of filter types in the topology at hand (e.g., GWOR has fewer types than λ-Router)

b) varies as a function of the fabrication uncertainty (tolerance on MRR radius, laser center wavelength uncertainty).

With state-of-the-art technology parameters, all 8x8 topologies turn out to be infeasible.

We assume a 3D Target Architecture

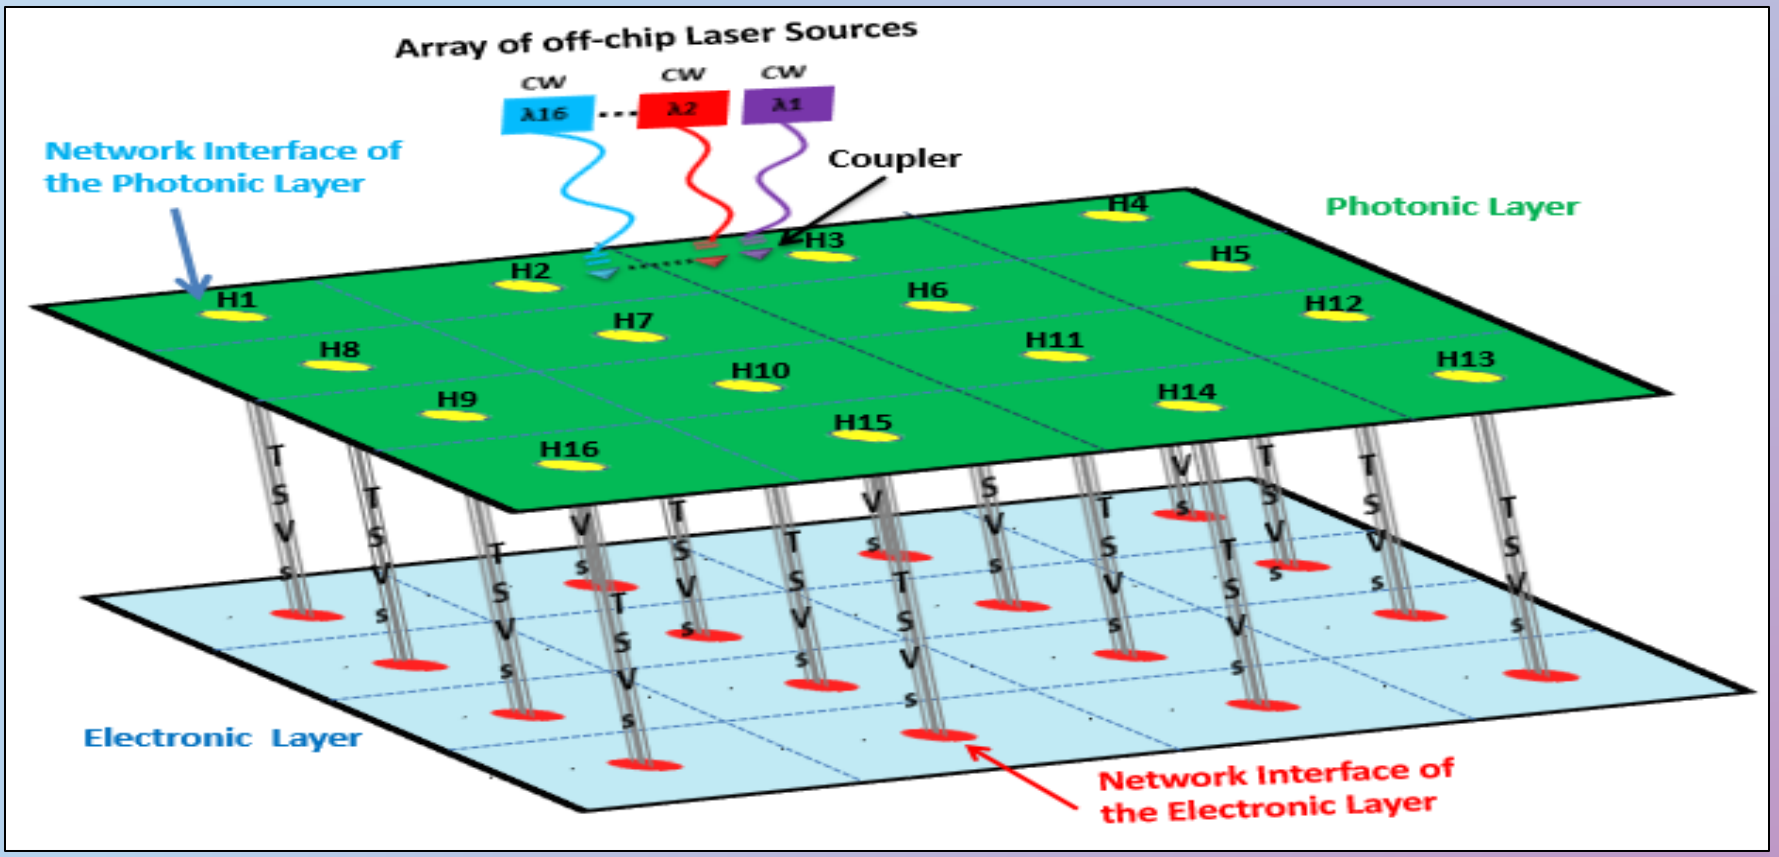

- Off-chip Laser sources.
- Optical layer placed on top of an electronic layer.
- Each layer has its Network interfaces distributed in a mesh based structure.

Placement and Routing Techniques

Irregular topologies with no apparent regularity pattern

These topologies are fed to ONoC-specific **place&route** tools **(PROTON, PLATON*)**, that optimize the physical mapping in terms of:

- **Number of waveguide crossings,**
- **The propagation distance.**

** A collaboration UNIFE – TU Munich*

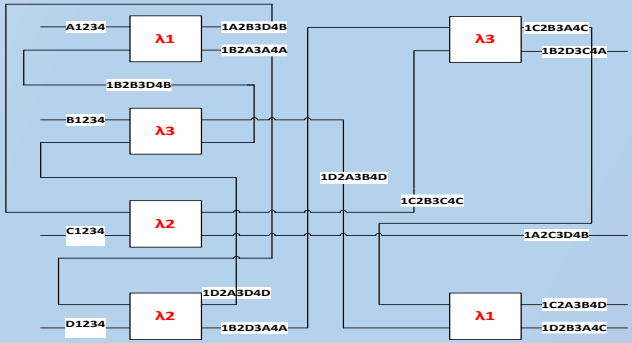

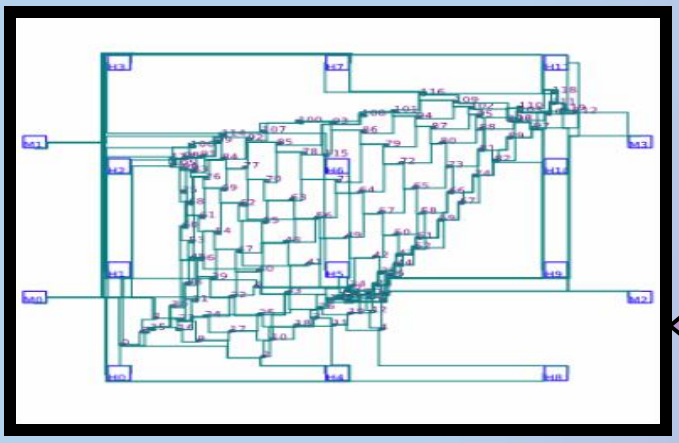

Initial Placement

Global Placement

Legalization

Routing

Layout Generation

These tools perform place&route at the Granularity of the individual photonic switches

Experimental Results

The above synthesis methodology was used to generate the entire design space of 4x4 WRONoC topologies, mapped onto the optical layer of a 3D stacked system. A power model has been applied to each solution:

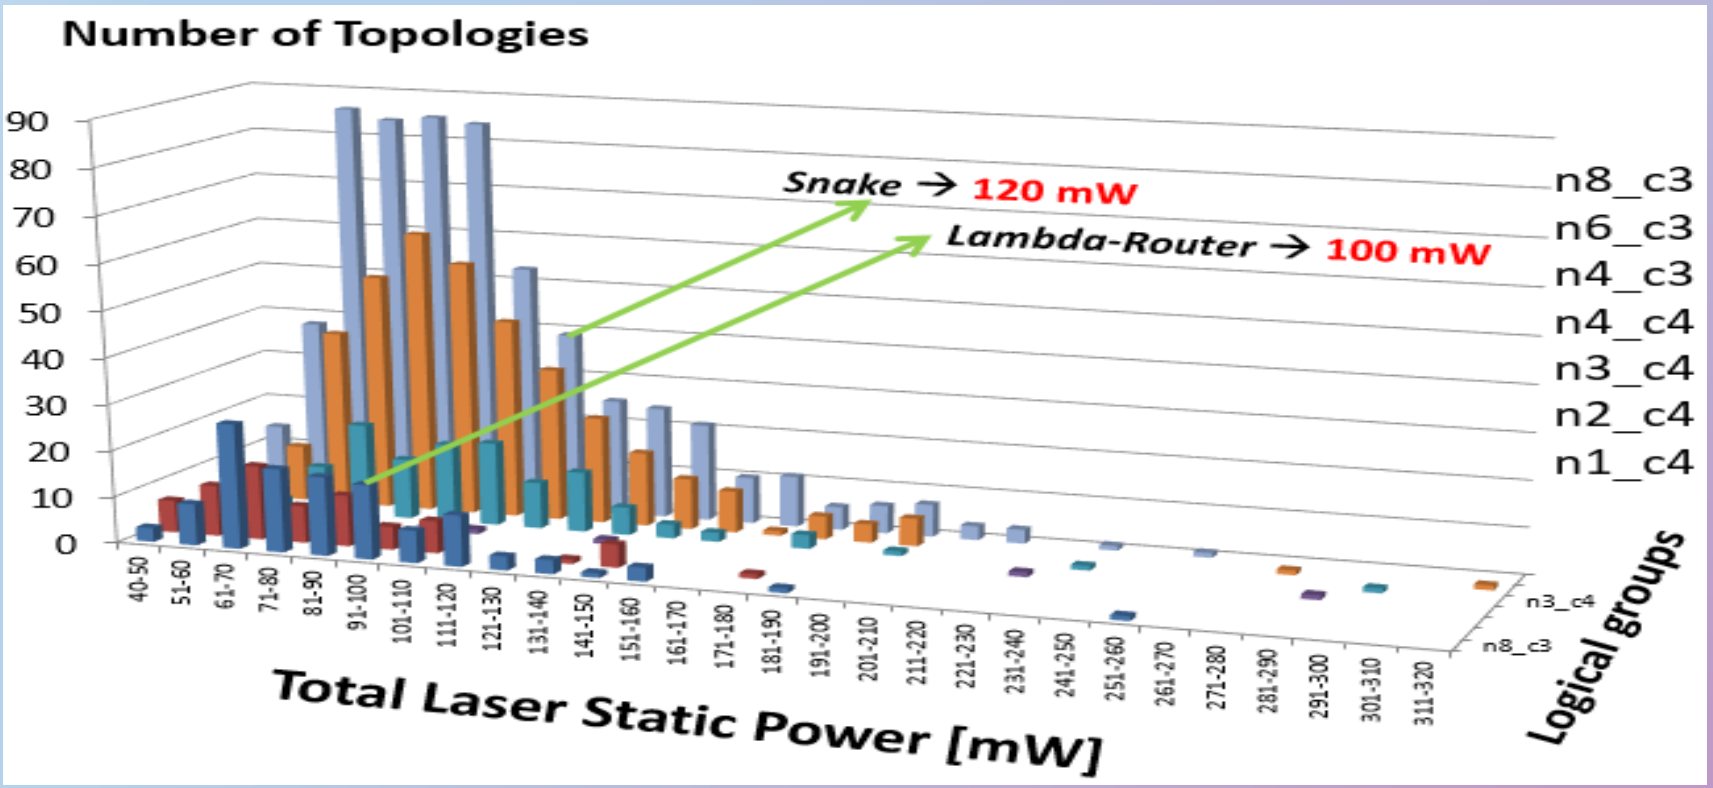

- 1296 topologies categorized as "*c* crossings on *n* critical paths."
- **The synthesis methodology has generated new topologies that are more power efficient than the best solutions reported so far in the literature** (λ-Router and Snake).

CONCLUSIONS

- This work proposes a methodology that systematically synthesizes all the points of the WRONoC topology design space.
  - ✓ Built upon the combination and aggregation of basic filtering primitives.

- The derived abstraction is quite powerful!
  - ✓ It enables to populate a design space previously limited to few design points devised by researchers' intuition.

- We demonstrate the existence of novel topologies that outperform the existing ones.

- We quantify the needed improvement of technology and process parameters to remove the current severe limitations to the connectivity scale and bit-level parallelism.

Overall, this work is a stepping stone into future work targeting automatic pruning of the design space in search for the most efficient solution that meets the requirements of the system at hand.

IWES 2016 **Italian Workshop on Embedded Systems**

Pisa, Italy 19-20 September 2016

Thanks For Your Attention