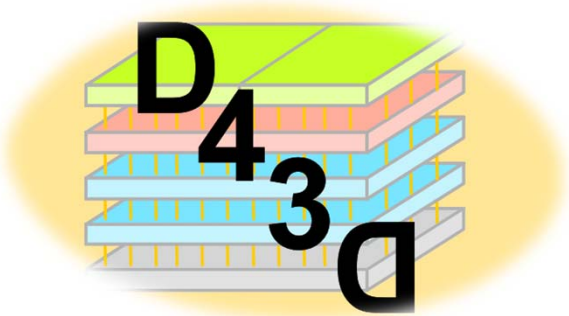

Toward a Top-Down Synthesis Methodology for 3D-Stacked Wavelength-Routed Optical Networks-on-Chip

Davide Bertozzi
University of Ferrara (Italy)

Collaboration with prof. Ulf Schlichtmann (TU Munich), with prof. Milos Krstic (IHP Microelectronics and Univ. Potsdam), and with prof. Maddalena Nonato (Univ. Ferrara)

OPTICAL NETWORKS-ON-CHIP

Optical interconnection networks stand out as the most promising emerging technology to work around on-chip communication bottlenecks

EMERGING CHALLENGES:

- ***Latency sensitivity***
- ***Bandwidth criticalities***
- ***The power overhead***
- ***Memory bottleneck***

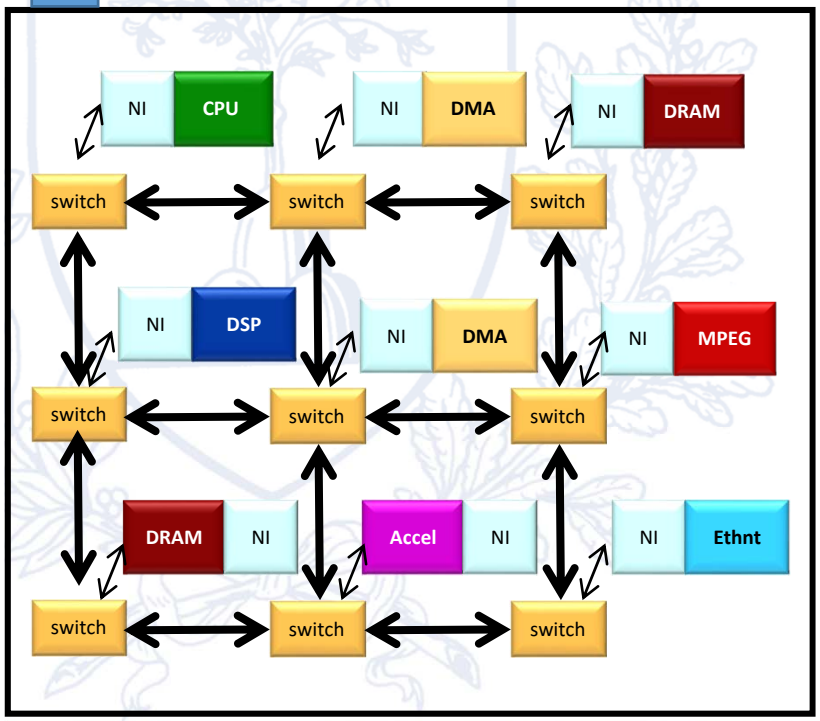

Laser sources

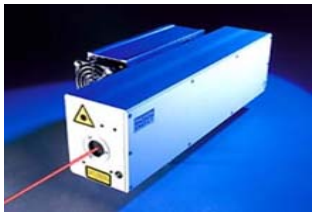

Optical modulation

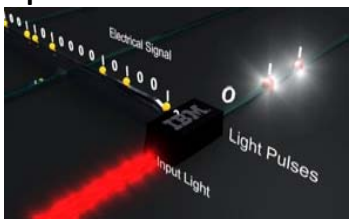

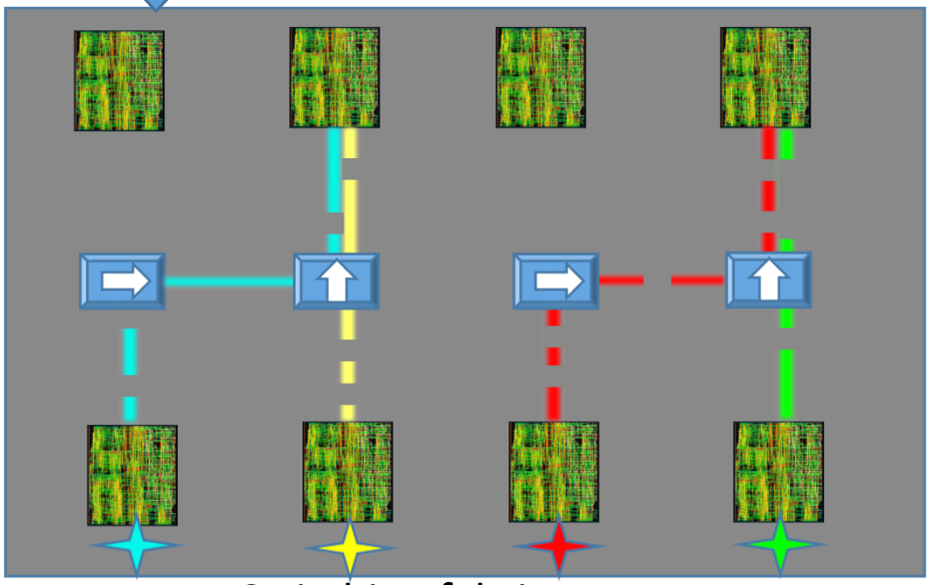

Switching fabric

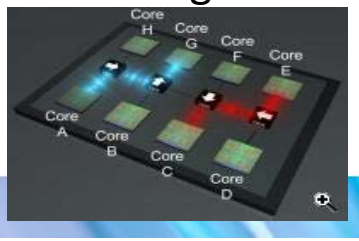

Photodetector

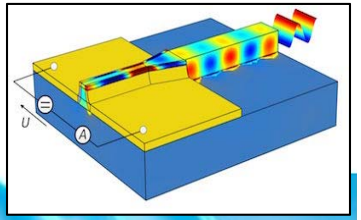

OPTICAL NETWORKS-ON-CHIP

Optical interconnection networks stand out as the most promising emerging technology to work around on-chip communication bottlenecks

EMERGING CHALLENGES:

- ***Latency sensitivity***
- ***Bandwidth criticalities***
- ***The power overhead***
- ***Memory bottleneck***

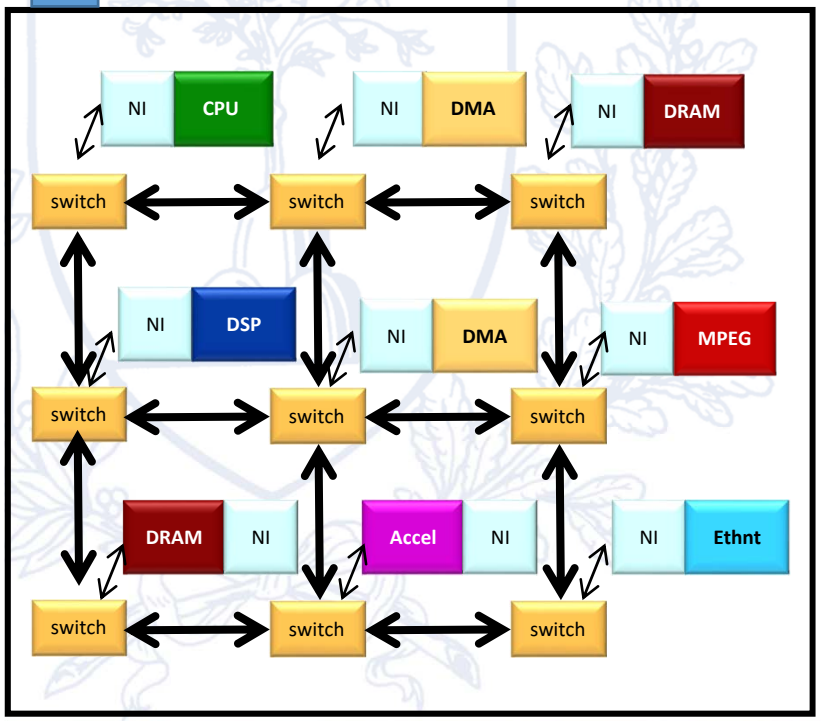

Laser sources

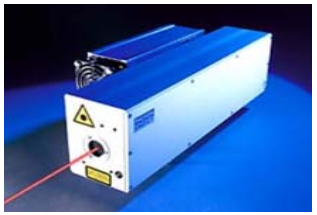

Optical modulation

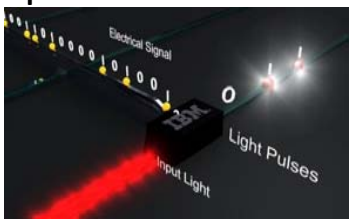

Still far away from a technology of practical relevance. Why?

Switching fabric

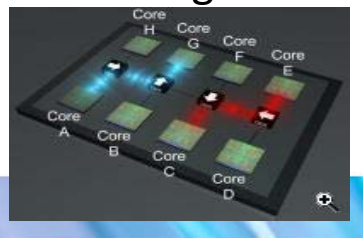

Photodetector

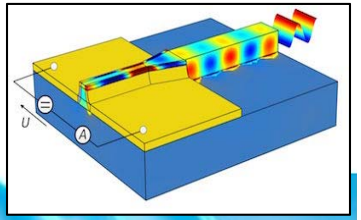

Mind the Gap

The gap between system-level designers and technology developers is currently still huge!

Architecture design points stem directly from designers' intuition

Descriptive information at different abstraction layers are mixed

Designs are difficult to compare with one another

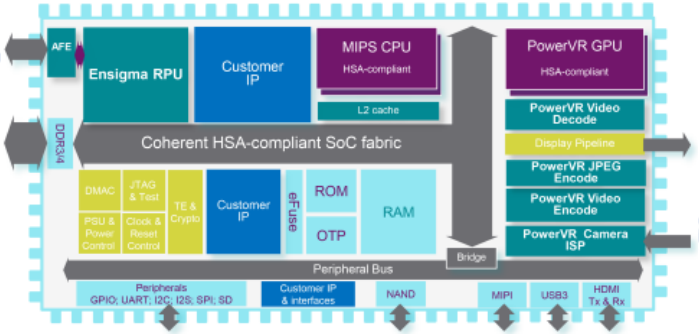

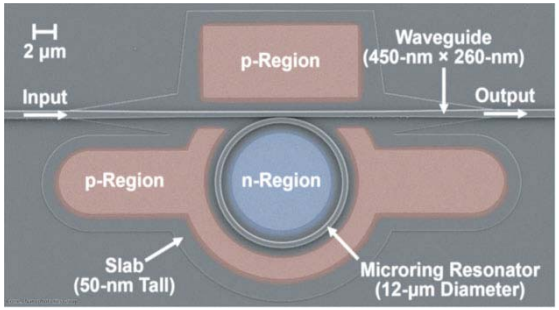

The application of well-known optimization techniques is difficult

No consistent methodologies to explore the design space

Technology-aware design not possible

Requirement targeted by this work:

starting from a high-level description, operate on abstractions and refine them into an actual implementation with components from a technology library.

EDA BEYOND ITS «E-ROOTS»

Can we extend EDA paradigms and methodologies to ONoC topology synthesis?

Digital VLSI design

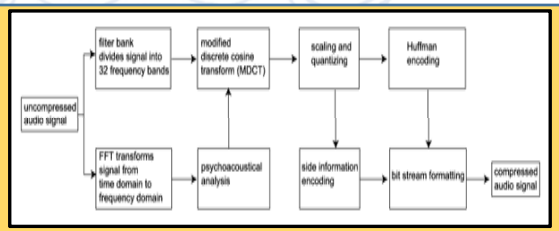

High-level specification

Technology-independent Logic Library

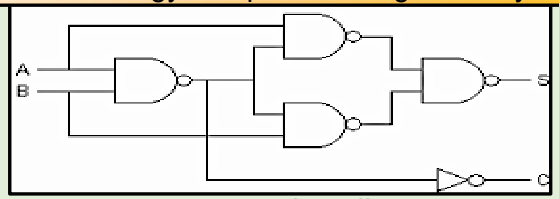

Gate-Level Netlist

Technology Library

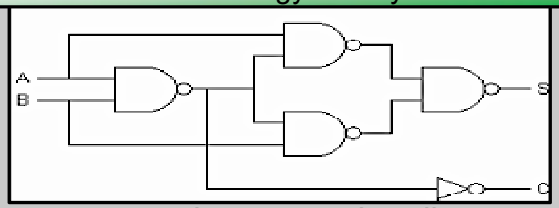

Mapped Gate-Level Netlist

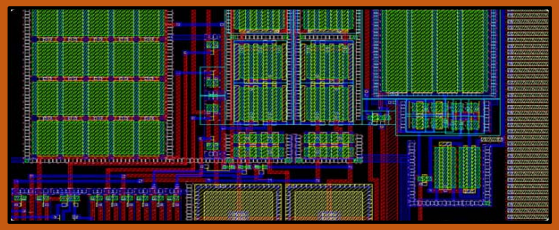

Planar geometric shapes

ONoC topology synthesis

0. Routing Protocol Selection
- I. Network Interface Design
- II. Switching Primitives Representation
- III. Technology Mapping
- IV. Assignment of carriers wavelengths
- V. Device Parameter Selection
- VI. Placement and routing
- VII. Physical Design

WAVELENGTH-ROUTED OPTICAL NOCS (WRONOCS)

Our synthesis methodology currently targets the Wavelength-selective routing protocol:

Benefits
Predictable all-optical interconnect solution (concurrent contention-less all-to-all connectivity)

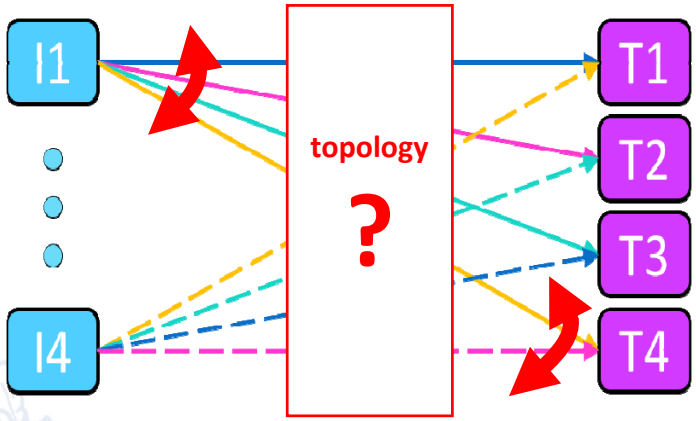

Limitation
Scalability concern

Requirement
Same-wavelength channels never overlap in the topology

16x16 server blade

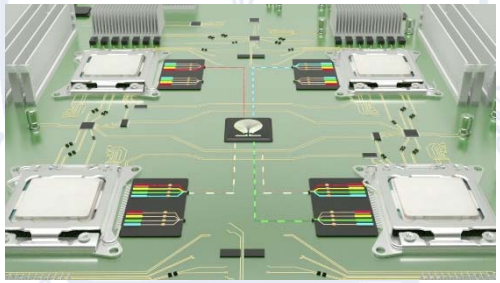

ICT-STREAMS multi-socket interconnect

Hi-LION at UC Davis

100000s of nodes

Key intuition: the basic abstract primitive for the construction of each topology is the **1x2 Drop Filter (DF)**

On-resonance signal λi
Off-resonance signal λj

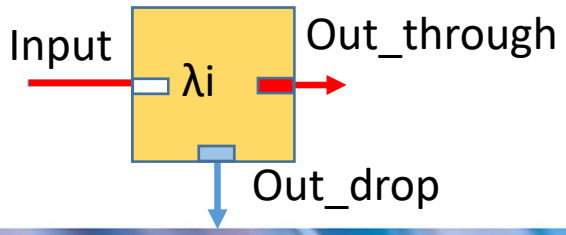

Premature!

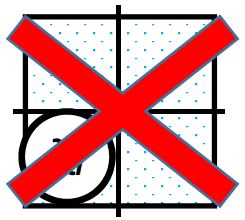

FRONT-END REFINEMENT METHODOLOGY

NxN CONNECTIVITY REQUIREMENT | **Future: NxM, bandwidth requirements,…**

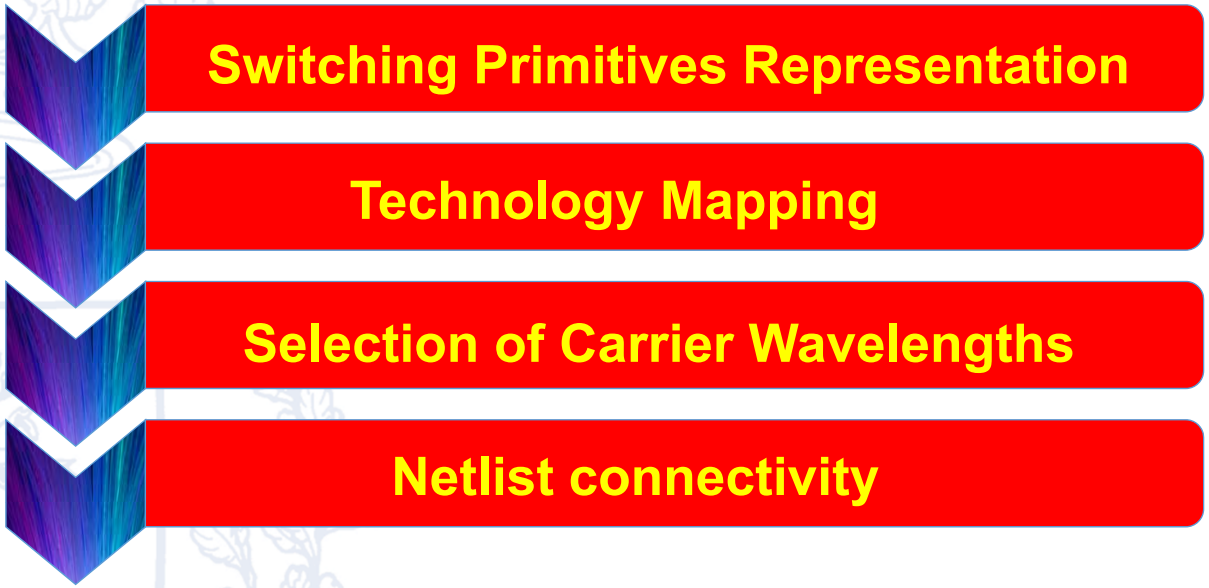

COMPLIANT LOGIC TOPOLOGY | **Future: Pruning of the design space**

How does the design space of NxN WRONoC topologies look like?

Can we understand all topology design points in the context of a unified design framework?

REFINEMENT METHODOLOGY

1. Wavelength Resolution

Each channel of the WDM input signal should be resolved so to be routed to a different output

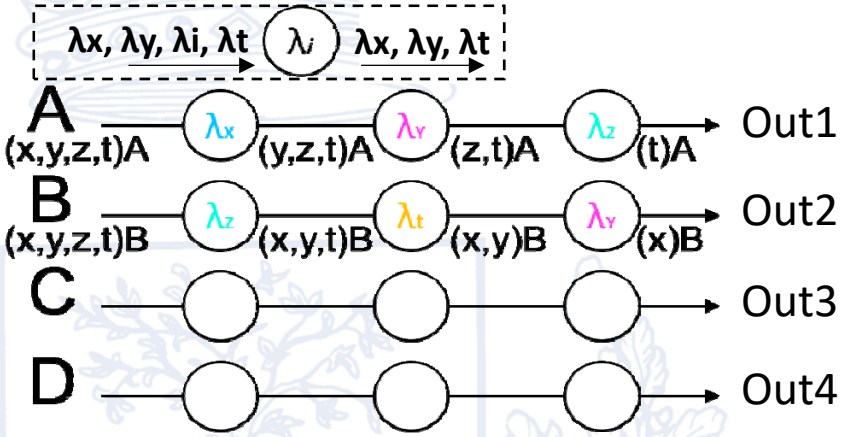

*Wavelength Resolution Graph **(WRG)** for a generic 4x4 WRONoC.*

2. Technology Mapping

E.g., Grouping the 1x2 DFs into compact 2x2 photonic switching elements (PSEs), from a technology library!

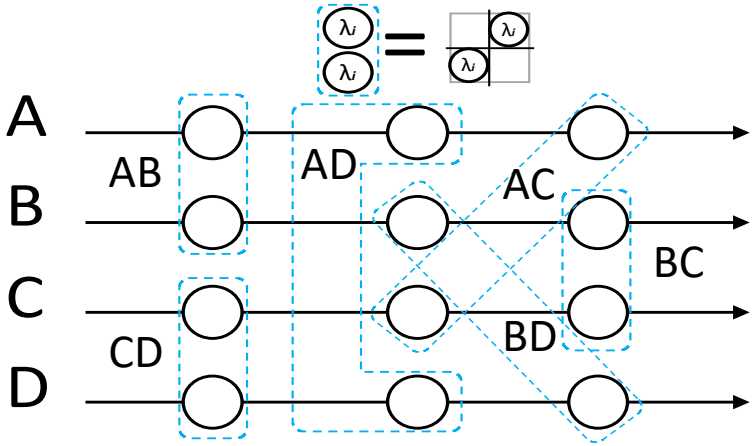

Constraint: Drop channels on rows only once!

3. Symbolic Wavelength Assignment

Assign a resonant wavelength to the MRRs

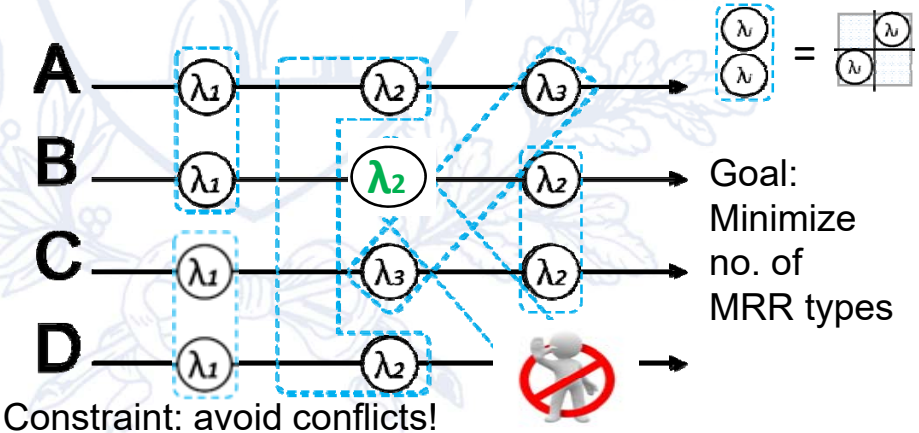

Goal: Minimize no. of MRR types

Constraint: avoid conflicts!

4. Topology Connection

Draw the topology logic scheme. It's a **λ-router!** .

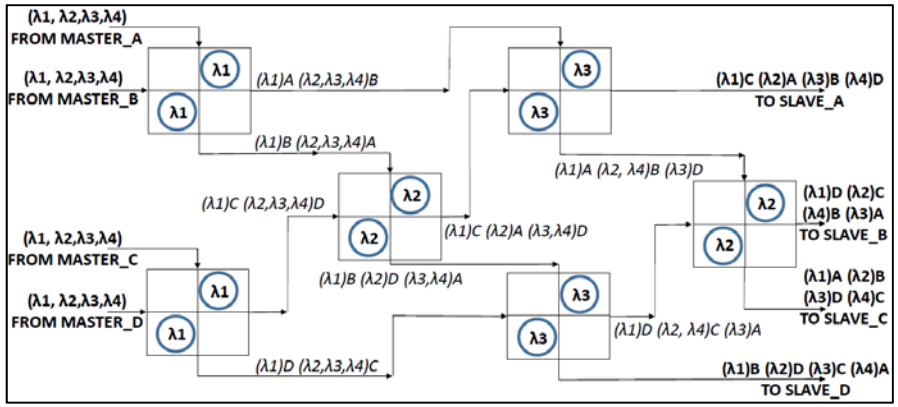

REFINEMENT METHODOLOGY

1. Wavelength Resolution

Each channel of the WDM input signal should be resolved so to be routed to a different output

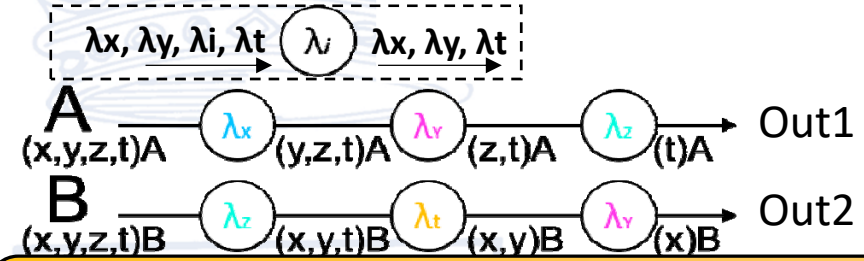

2. Technology Mapping

E.g., Grouping the 1x2 DFs into compact 2x2 photonic switching elements (PSEs), from a technology library!

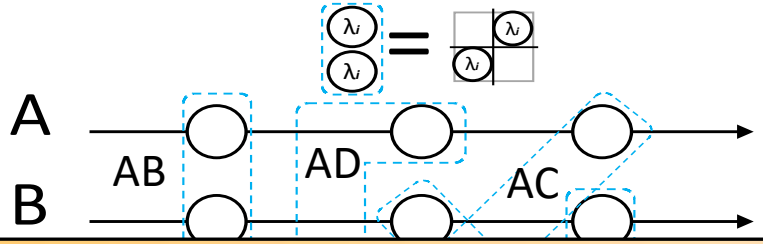

Our synthesis methodology can potentially populate the complete design space of WRONoC topologies by spanning all possible technology mappings, subject to the constraints of each stage for legal solutions.

Wavele

Only with 2x2 PSEs, the number of WRONoC topologies in the design space amounts to $[(n-1)!]^n$

3. S

Assign a resonant wavelength to the MRRs

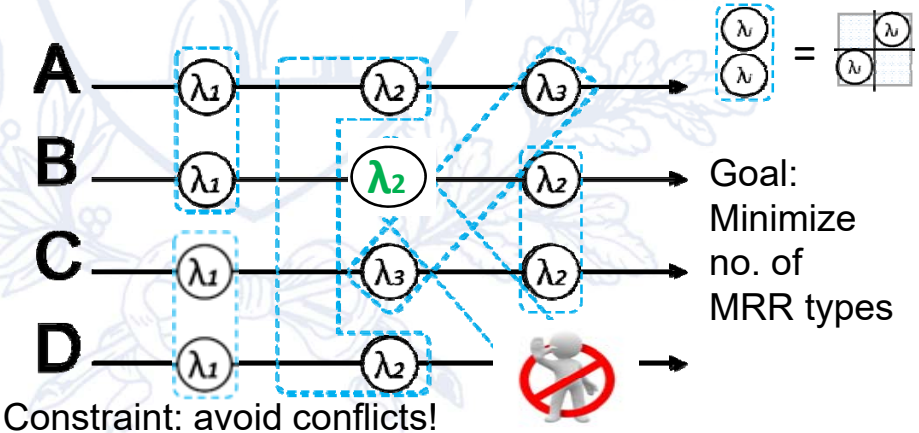

Goal: Minimize no. of MRR types

Constraint: avoid conflicts!

Draw the topology logic scheme. It's a **λ-router!** .

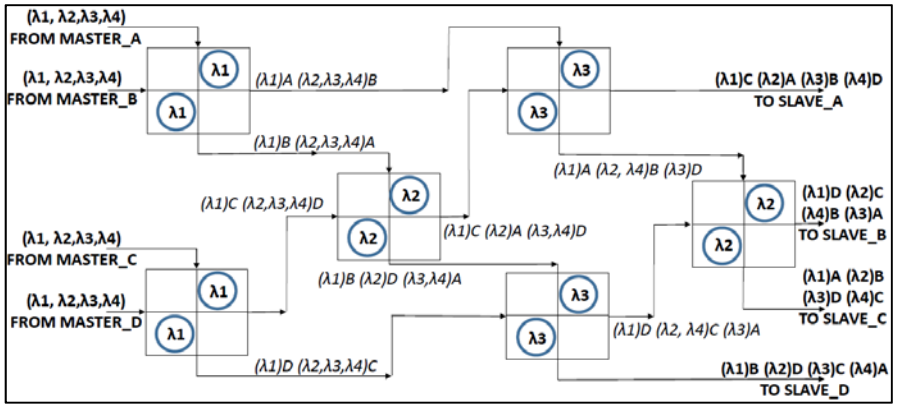

BACK-END REFINEMENT METHODOLOGY

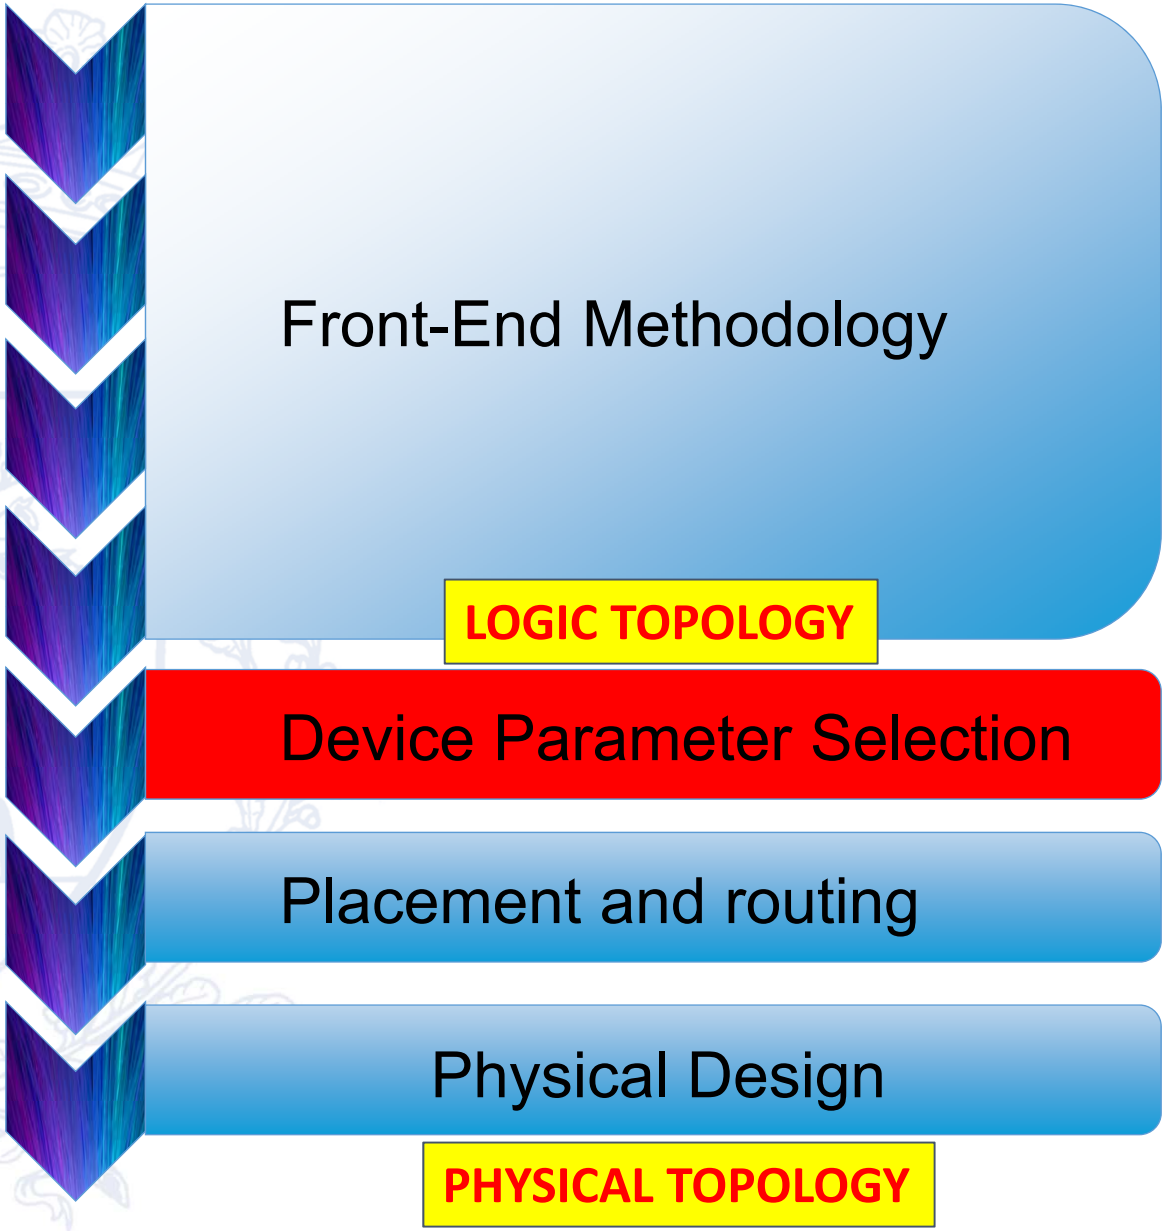

DEVICE PARAMETER SELECTION

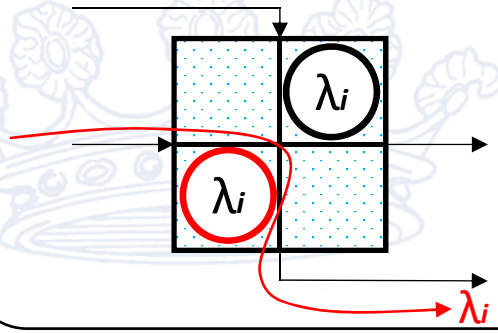

- What is the exact radius length of the MRR, typically in the range 5-20um?
- What is the exact value of the n wavelengths used by each initiator in an *n x n* wavelength-routed optical NoC?
- How much bit-level communication parallelism can be implemented on the I/O optical channel? How many wavelengths overall?

This is not just a refinement step, due to the ROUTING FAULT concern: It has implications on network-level throughput and scalability

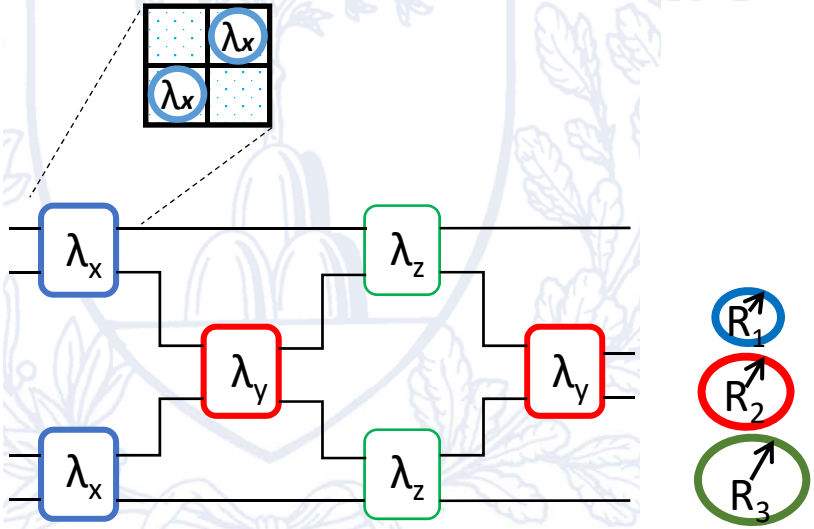

→ Available parallelism is ~~6 4~~ **3** in PSE_x
→ Available parallelism is ~~7 5~~ **3** in PSE_y
→ Available parallelism is ~~9~~ **7** in PSE_z

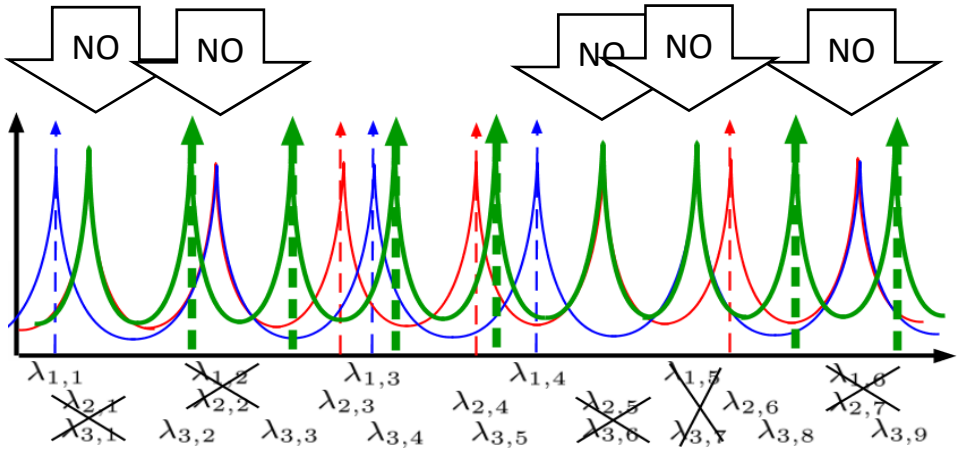

Overlapped peaks unusable

DEVICE PARAMETER SELECTION

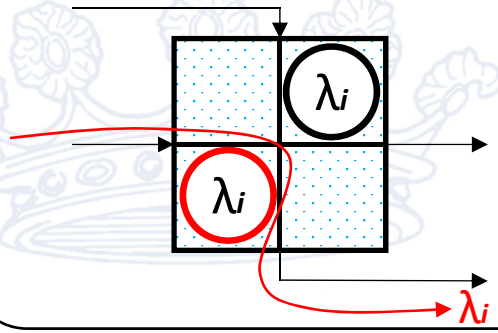

- What is the exact radius length of the MRR, typically in the range 5-20um?
- What is the exact value of the n wavelengths used by each initiator in an *n x n* wavelength-routed optical NoC?
- How much bit-level communication parallelism can be implemented on the I/O optical channel? How many wavelengths overall?

This is not just a refinement step, due to the ROUTING FAULT concern: It has implications on network-level throughput and scalability

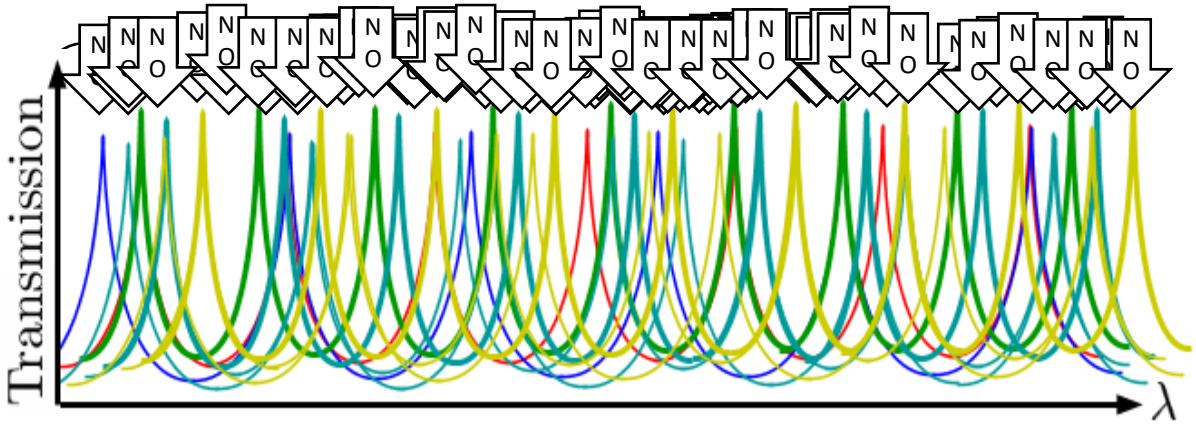

As topology size increases, the proliferation of filter types and wavelength channels may limit the availability of non-overlapped transmission peaks, which may cause the topology to be practically infeasible

Parallelism and Scalability Limitations

PARAMETER UNCERTAINTY

This is the first refinement step directly exposed to the underlying technology

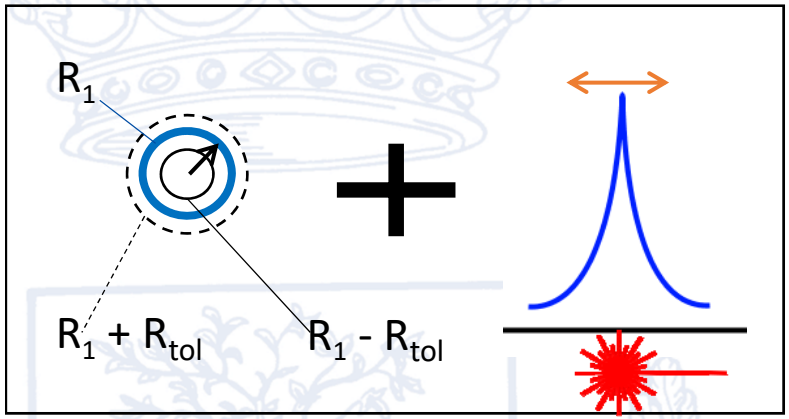

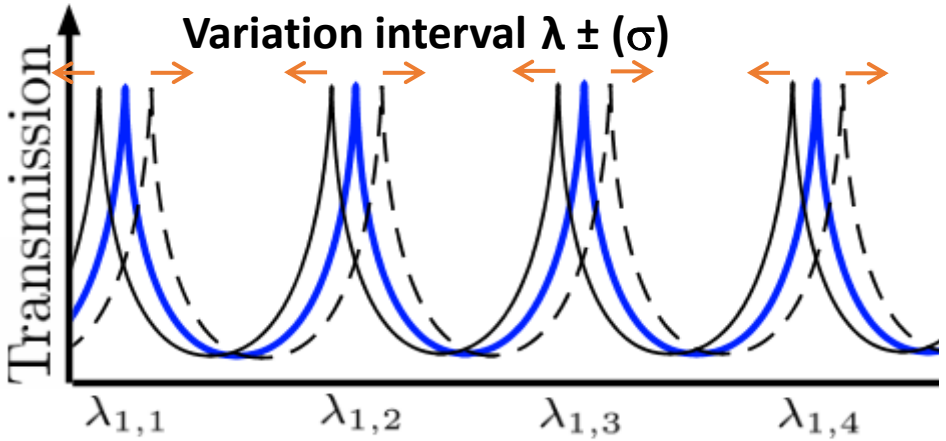

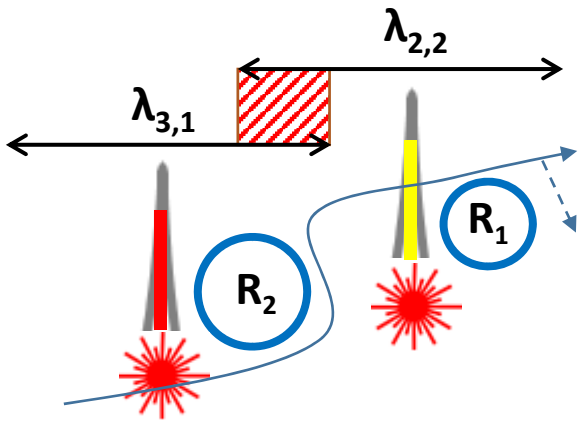

- There exists a post-fabrication variation scenario that ends up in a routing fault
- Even without overlapping, proximity raises optical crosstalk concerns

<u>Conservative/predictable design-for-reliability constraint :</u>
If the variation intervals overlap,
the transmission peaks are not usable for communication

We modelled the Ring radius/wavelength channel selection problem subject to routing fault avoidance as a Constrained Optimization Problem, and used ASP as declarative technology.

SCALABILITY

We performed device parameter selection to assess scalability of generic topologies

Radius selection range and Incremental step

Fabrication options	R_{min}	R_{step}	R_{max}
Ropt	5µm	1µm	25µm
R'opt	5µm	1µm	30µm
R''opt	5µm	0.25µm	30µm

Tolerance: 10nm
Laser uncertainty: 0.5nm

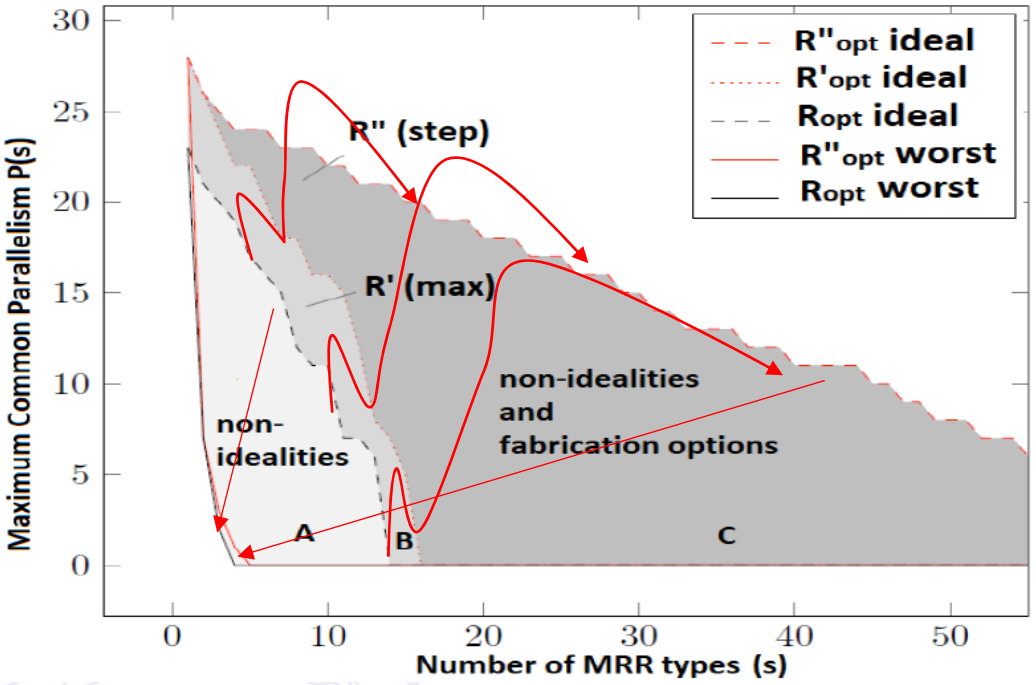

Ideal Fabrication

- With overly fine step and large rings, the upper bound is roughly a 60x60 topology, with limited parallelism though!
- Achievable parallelism most sensitive to the incremental step of MRR radii

Conservative PROCESS VARIATIONS

- Only 4x4 ONoCs are certainly feasible
- Multiple wavelength selection options useless if uncertainty ranges are not reduced accordingly

BACK-END REFINEMENT METHODOLOGY

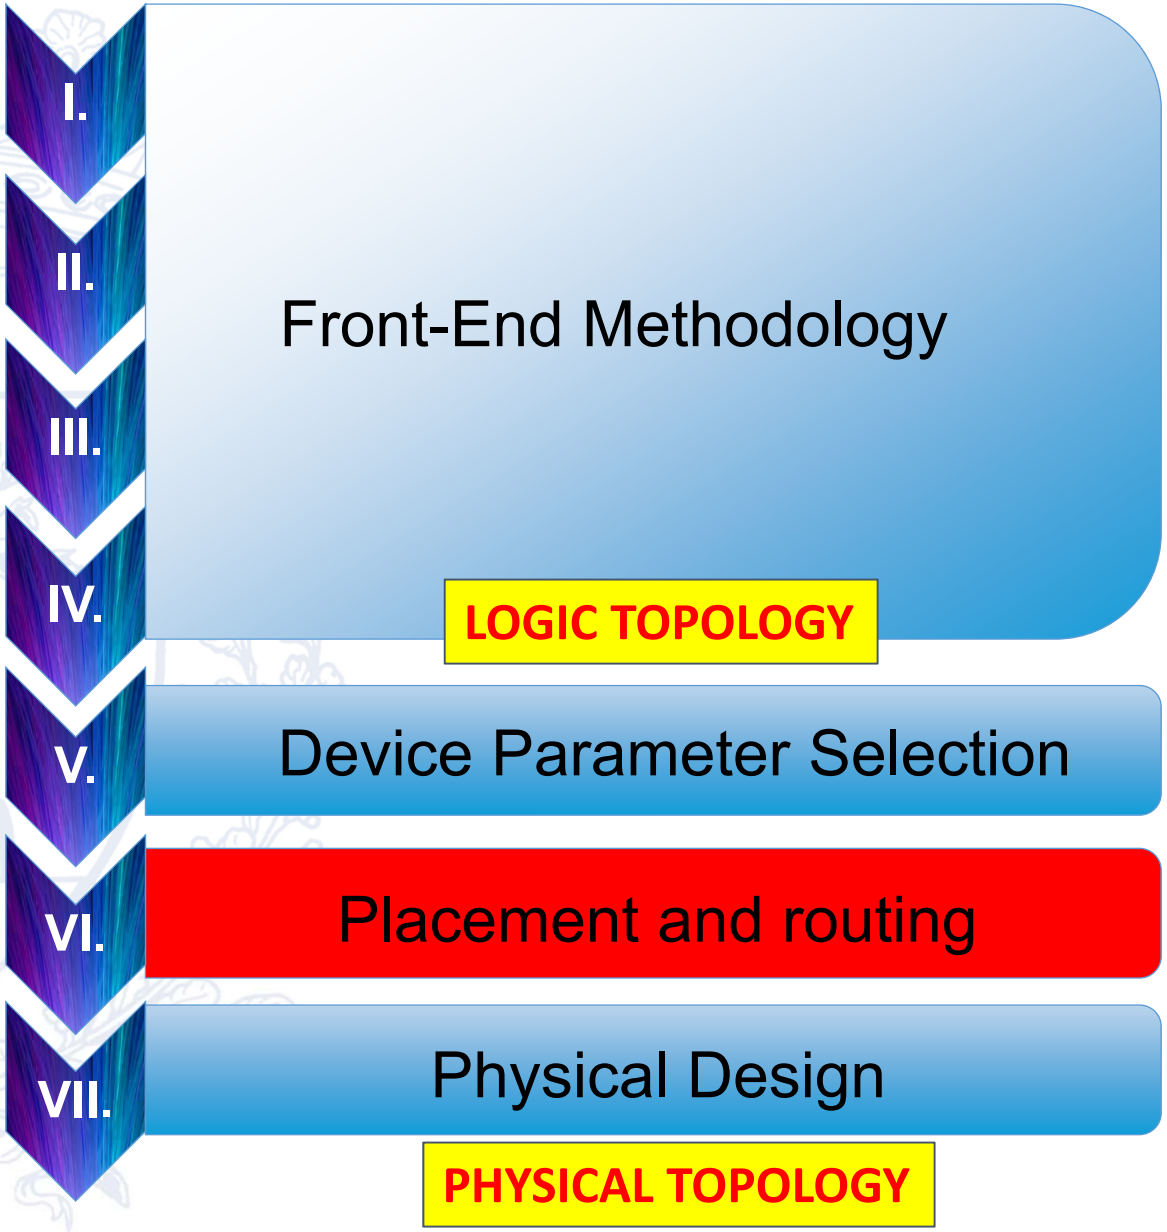

PLACEMENT AND ROUTING

There exists a design predictability gap between logic and physical topologies

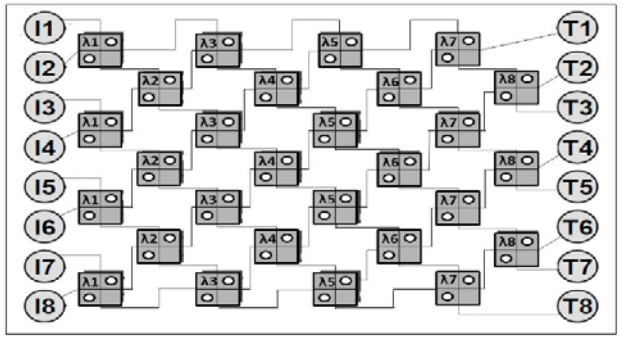

Logic Topology

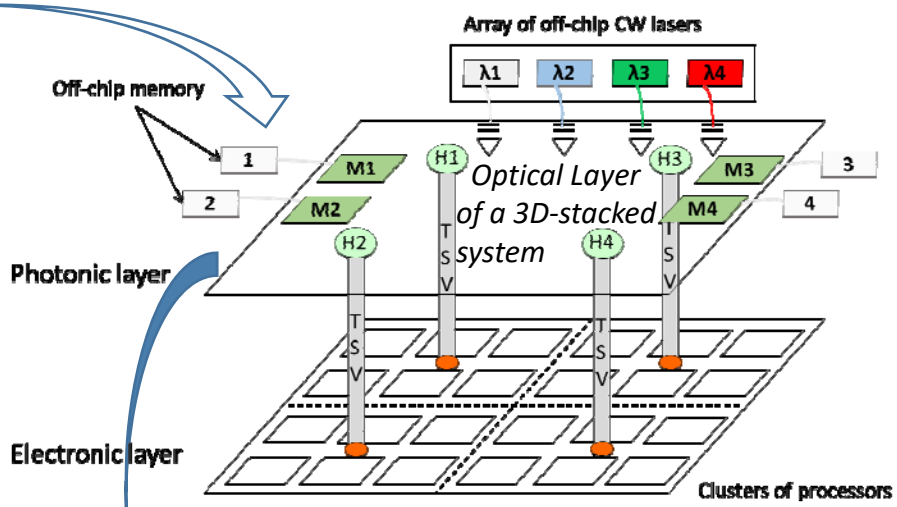

Electronic P&R tools cannot be reused here

We propose PROTON+, a tool for automatic placement and routing of ONoC topologies
(Collaboration with prof. Schlitchmann at TU Munich)

The tool tries to strike a good balance between crossing losses and propagation losses, which might be conflicting objectives

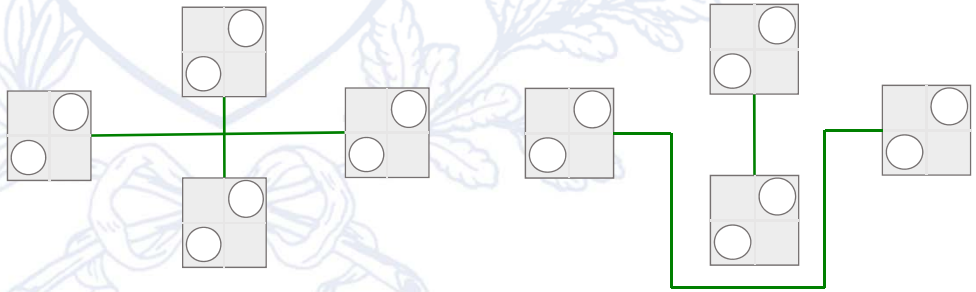

Minimize waveguide length Minimize no. of crossings

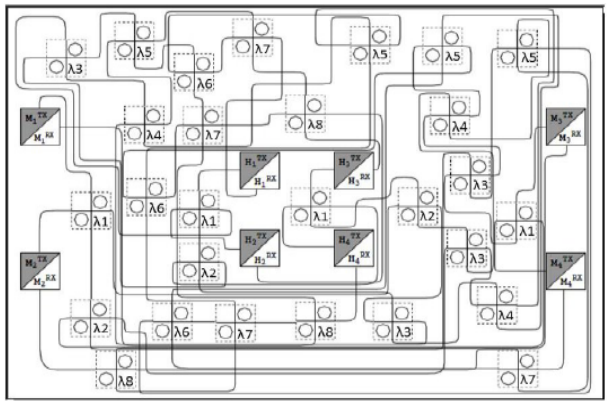

Lots of unexpected waveguide crossings (which burden on the static power budget)

Physical Design Space Exploration

The objective function should reflect a trade-off between accuracy in capturing power loss mechanisms and computation efficiency

$$\min_{\mathbf{x},\mathbf{r}} \max_{p \in P} \alpha \cdot L_p(\mathbf{x},\mathbf{r}) + \beta \cdot C_p(\mathbf{x},\mathbf{r})$$

Our objective functions minimizes the insertion loss across the lossiest path. This indirectly limits total laser power.

Where Lp and Cp are approximate functions of path lengths and no. of crossings

Placement: Non-linear optimization problem solved with an IPM

Routing: adaptation of the Lee's algorithm «Maze Router»

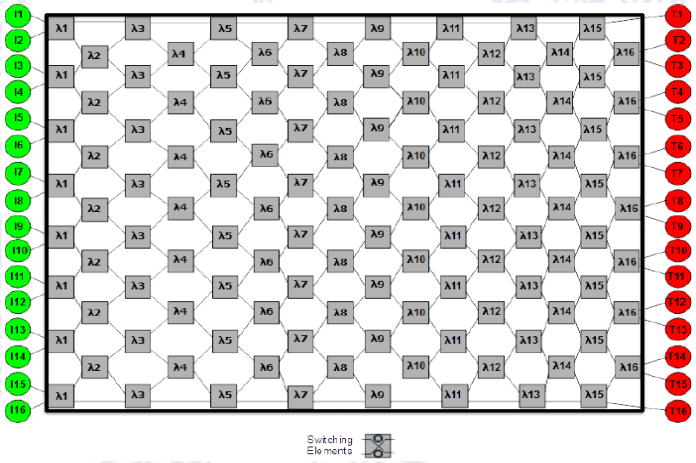

16x16 λ-Router

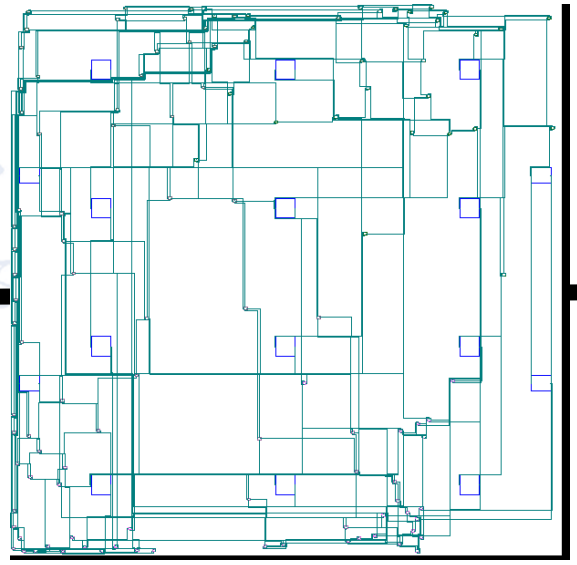

PROTON+: 44dB, 25k secs

Force-directed placement

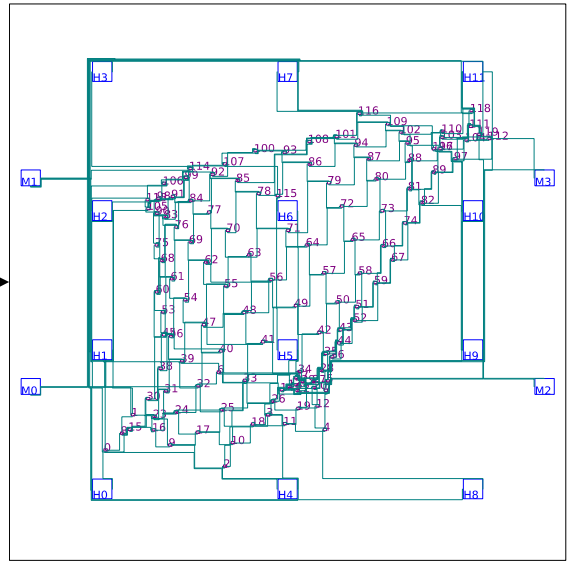

PROTON+ v2.0: 22dB, 64 secs

VERTICAL INTEGRATION

We exhaustively generated all 4x4 WRONoC topologies and mapped them with Proton+

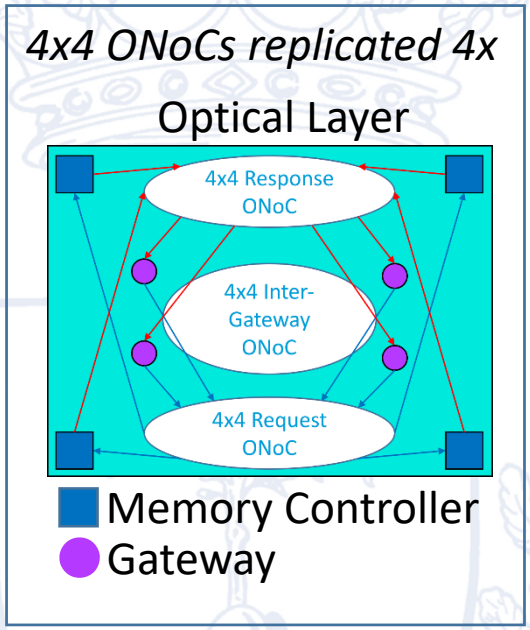

Physical Design with Proton+

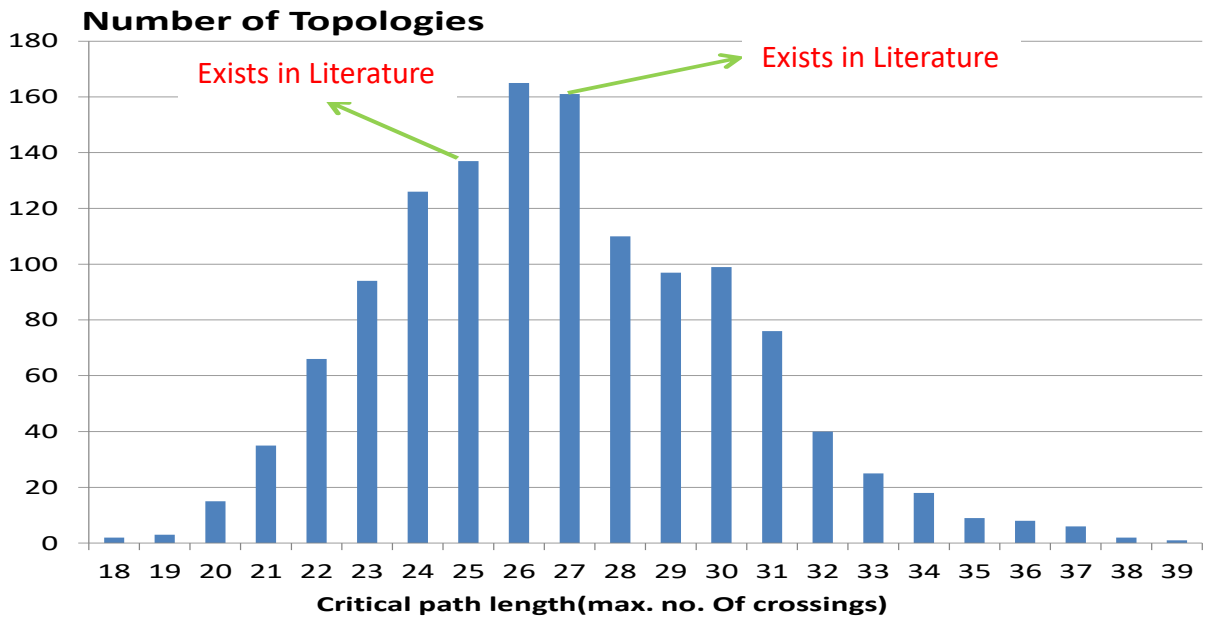

Distribution of the critical path after physical mapping.

- There is a large variability of the design space: from 18 to 39 crossings!
  - This raises the issue of placement-aware logic topology synthesis, completely new discipline for optical NoCs.
- λ-Router and snake proposed in literature are not the best topologies from the critical path length viewpoint!
  - **Design automation helps to get the most out of a technology**

CONCLUSIONS

- Optical networks-on-chip soon sinking into quicksand if the gap between technology developers and system designers is not bridged.
- Design automation for emerging technologies: from supporting technologies to enabling them
- This work proposes an early-stage cross-layer synthesis methodology refining abstract interconnect solutions into the intents for their layout design.
- Milestones:
  - ✓ All topology design points have been understood in the context of a unified design framework.
  - ✓ Device parameter selection problem correlated to network-level aggregate bandwidth, topology scalability and reliability margins.
  - ✓ Place&route methods have been differentiated for the most relevant ONoC topology families
  - ✓ By vertically integrating refinement steps and tools, we are shedding light into the characteristics of the design space.
  - ✓ Horizontal integration addressed through bridge design with electronics

Current work: WRONoC Compiler

1- not just populating and characterizing the design space, but pruning it!

2- vertical tool integration for technology/placement-aware topology synthesis

THANK YOU

BACKUP CONTENT:

HOW TO DRIVE THE SYNTHESIZED WAVELENGTH-ROUTED OPTICAL NoC TOPOLOGIES?

Bridge Architecture

Data rate adaptation, (De-)Serialization, Flow control
Clock Resynch., Message-dependent deadlock avoidance

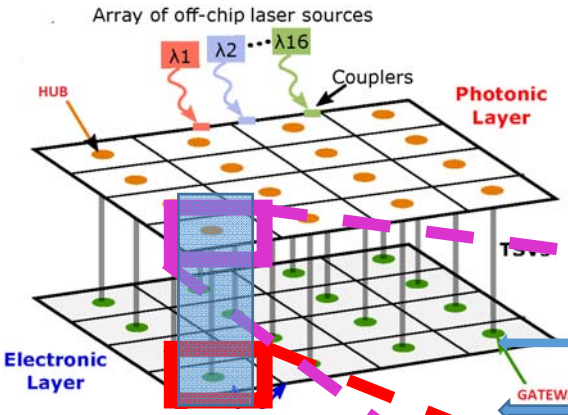

Gateway

Ind. 40nm ultra-low power CMOS

Ind. 40nm ultra-low power CMOS

IHP 130nm SiGe BiCMOS

Transmitter side

Receiver side

TRANSMITTER SIDE

RECEIVER SIDE

ENoC

VC dec

Demux 1x3

DC-FIFO (req)

DC-FIFO (reply)

dec

1x15

Mux 3x1

SER

PLL

Drivers

Modulators

λ11

λ12

λ13

λ1c

Arbiter

Credit counter 1

Credit counter 15

DC-FIFO

Mesochronous Synchronizer

Clock Divider

Filters

Comp

TIA

PD

DESER

Mux (15x2)x1

Bridge Configuration Space

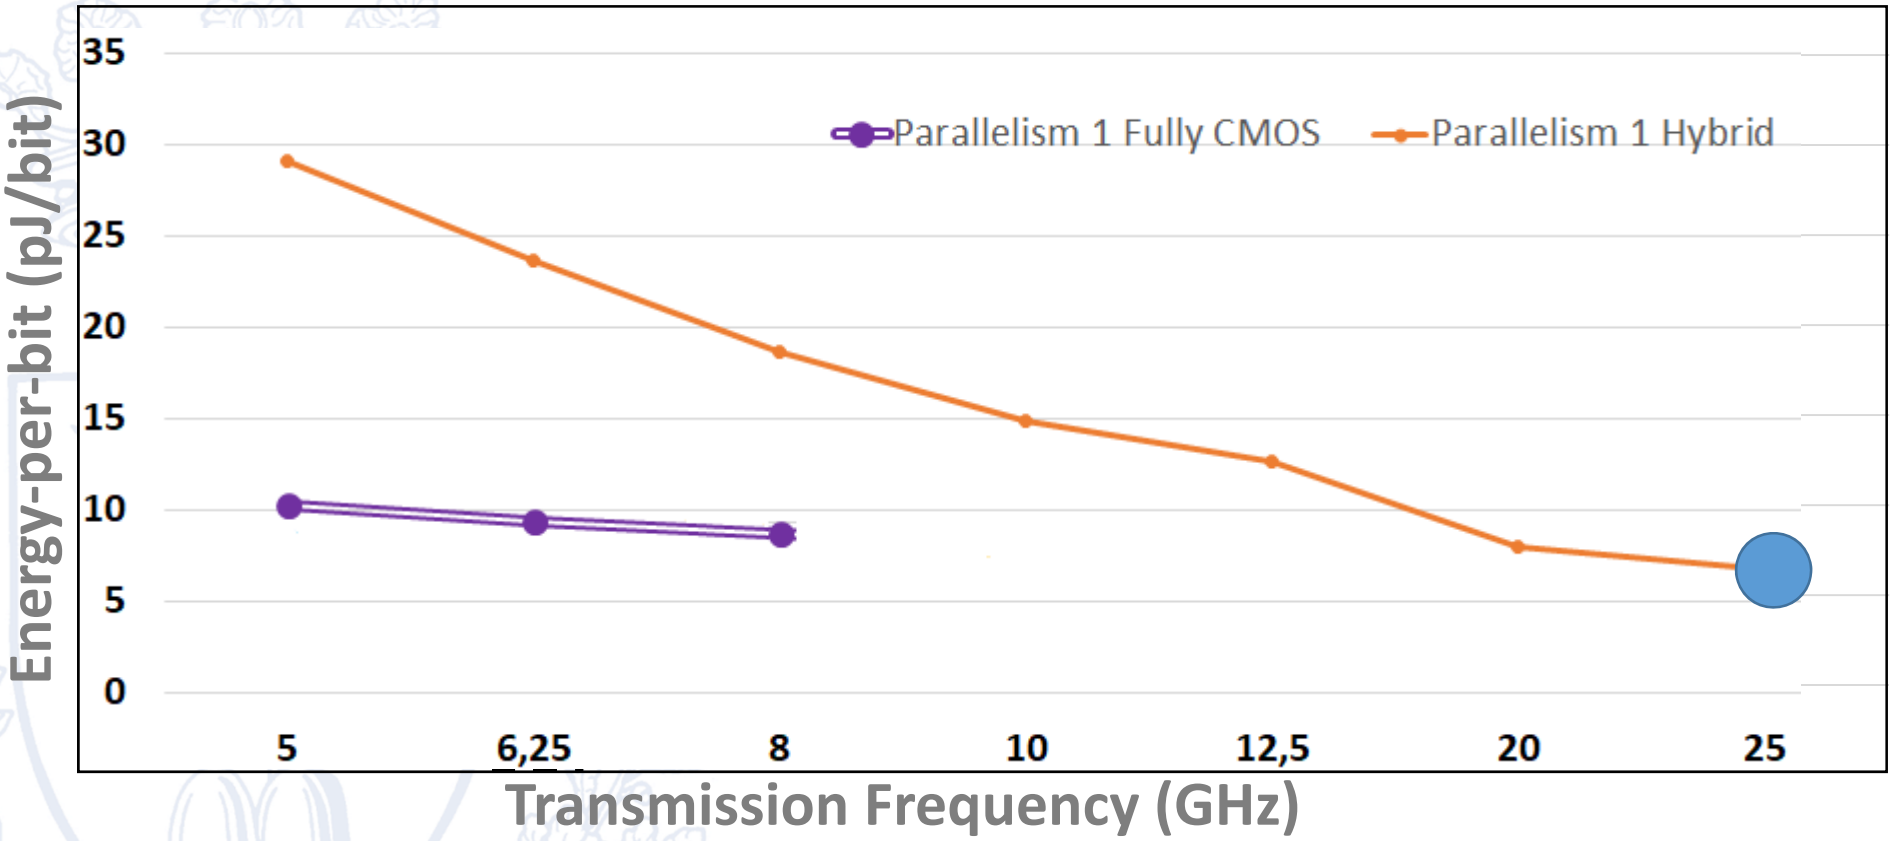

Hybrid solution achieves two goals at the same time:

Higher bit rate (**25 Gbit/s** vs. **8 Gbit/s**)

Lower energy-per-bit (**6.8 pJ/bit** vs. **8.7 pJ/bit**)

- **The E/bit at a certain rate decreases by increasing the bit-level parallelism**
- **A target throughput can be achieved in two ways, which are also energy-equivalent**

Higher frequency and **lower bit-level parallelism** **represent the best choice**