# Designing a multi-chiplet manycore system using the POPSTAR optical NoC architecture (invited)

Yvain Thonnart

23rd ACM/IEEE International Workshop on System-level Interconnect Pathfinding (SLIP)
Co-hosted with ACM/IEEE Intl. Conf. on Computer-Aided Design (ICCAD)
November 4, 2021

# *Designing a multi-chiplet manycore system using the POPSTAR optical NoC architecture*

*Yvain Thonnart*

*CEA-List*

# Silicon Photonics for short-range communication

- **Silicon Photonics moves forward for long distance optical wireline transceiver**
  - 100 / 400 Gigabit Ethernet

- **Large-scale electronics longs for low-latency low-energy dense communication**

- **Optical short-range communication has been a long-term target for years**
  - Needs compact optical devices to maximize bandwidth per mm²
  - ➔ **Microring optical resonators**

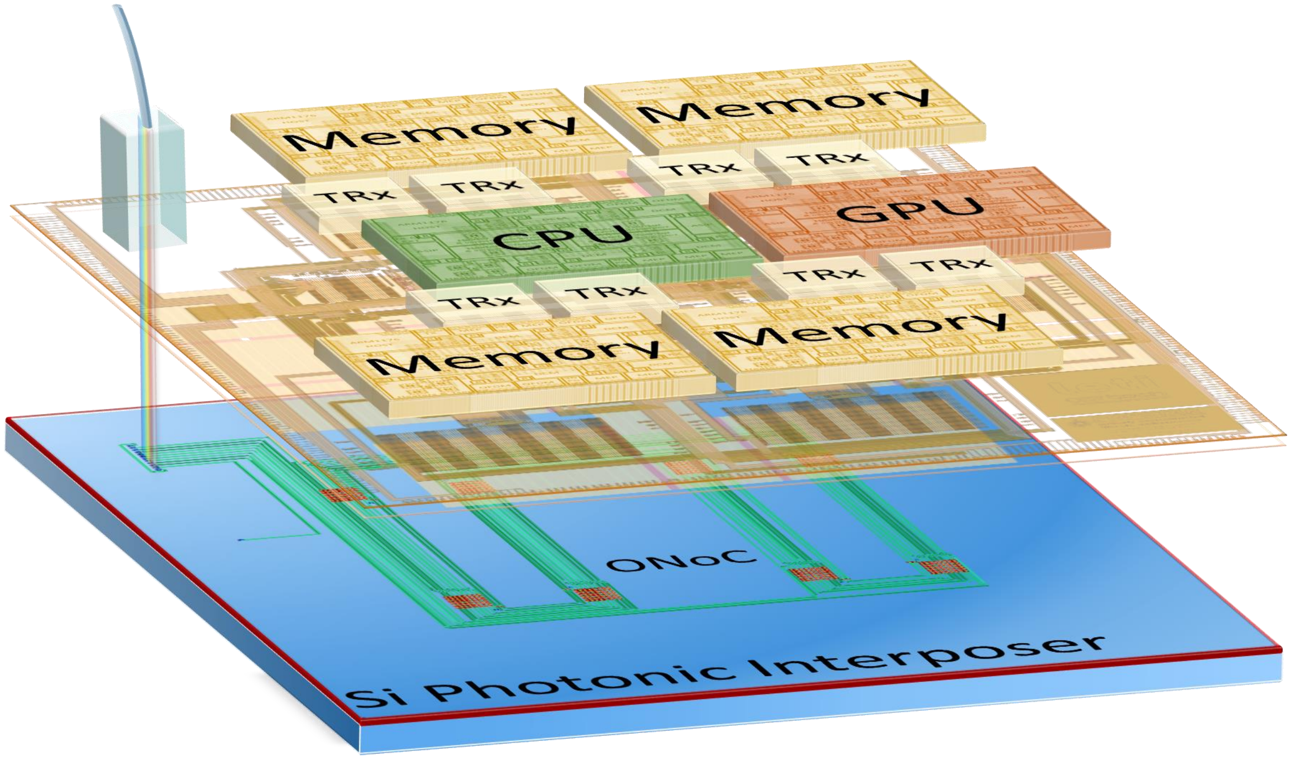


Optical Network on Chip

# Microring modulator based link

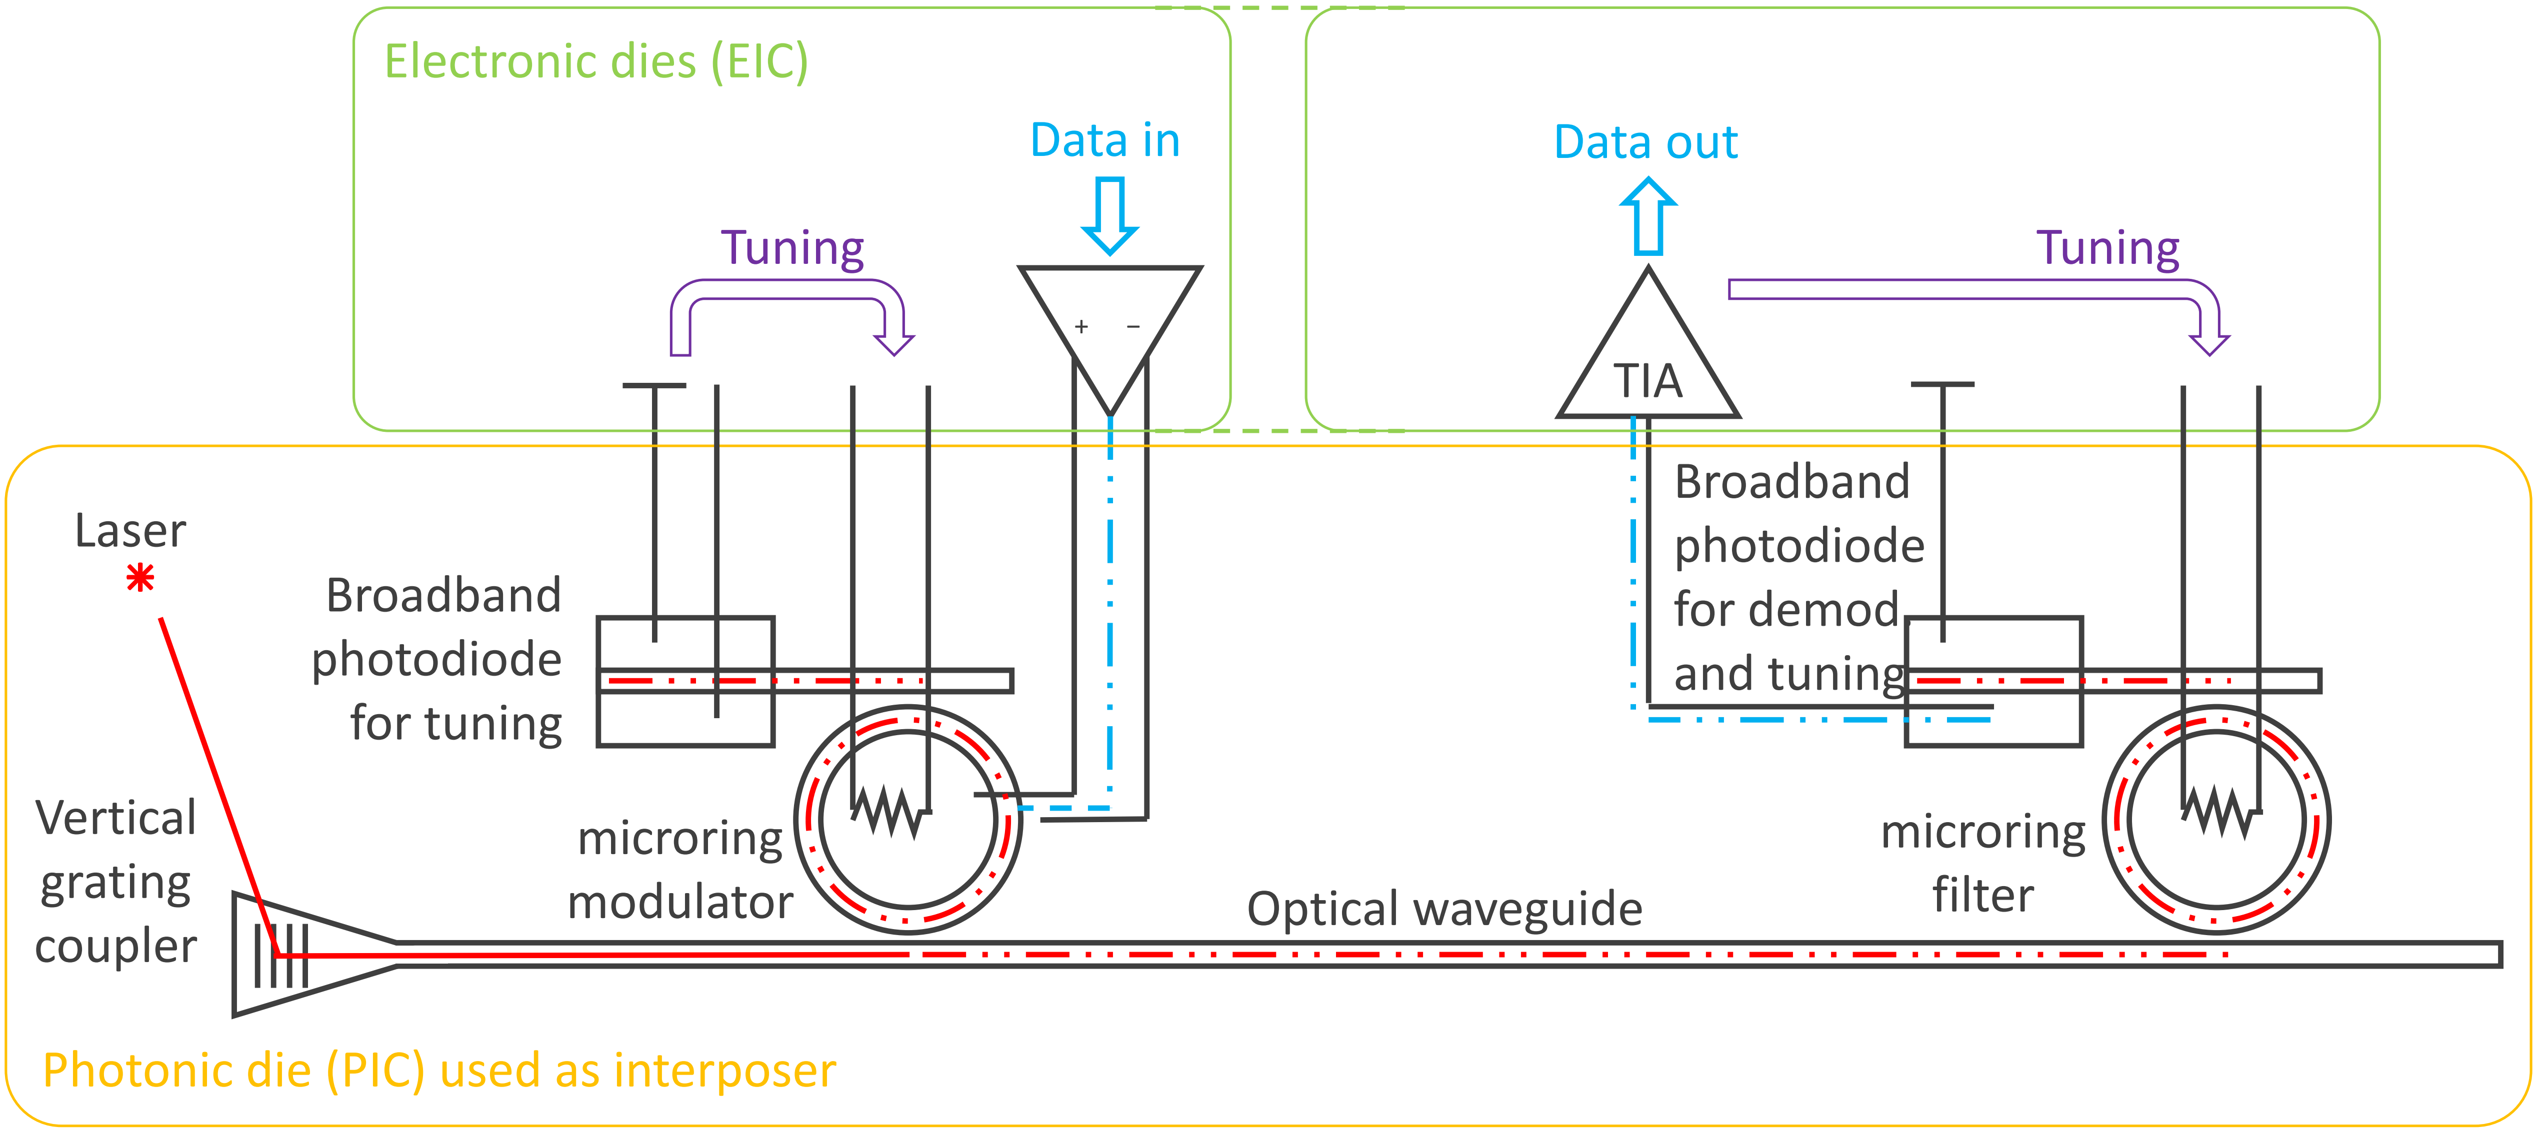

# Microring: Optical resonant cavity

- **Compact optical devices**
  - Highly resonant: Q-factor 10,000–30,000
- **Any refractive index change shifts the resonant wavelength**
- **PN or PIN diode junction can be created inside the ring for electrical control**
  - Different uses depending on diode
    - PN rings can be used as modulators (> 10 Gbps)
    - PIN rings can be used as filters (<500 MHz) for routing and wavelength demultiplexing
- **But Subject to Temperature variations**
  - ➔ Low-frequency resonance shift

$$\lambda_{res} = \frac{P\, n_{eff}}{k}$$

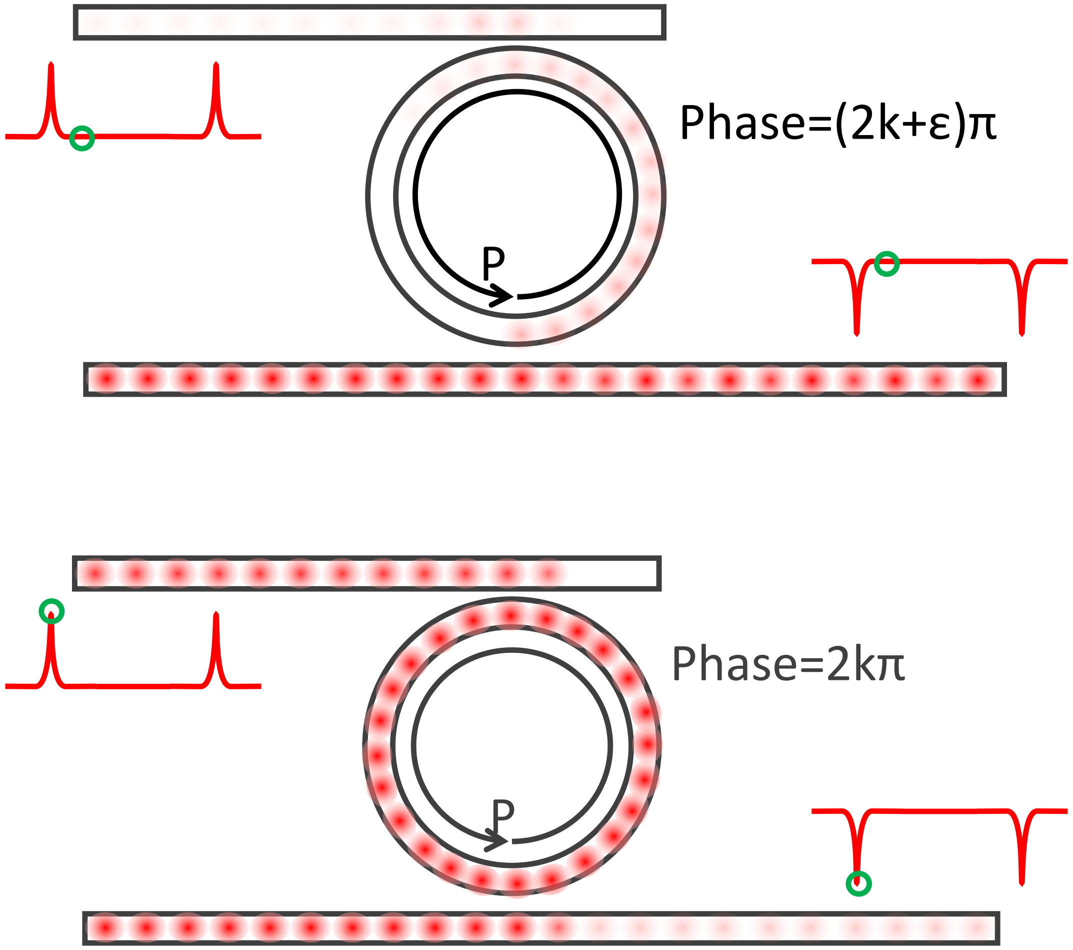

# Wavelength Division Multiplexing in a single waveguide

- **Narrow MRR resonances allow multiplexing**
  - Independent data streams
  - Paired MRR modulators and filters per wavelength
  - Additional MRR filters can be used for routing

- **Different ONoC topologies can be derived from this base scheme**

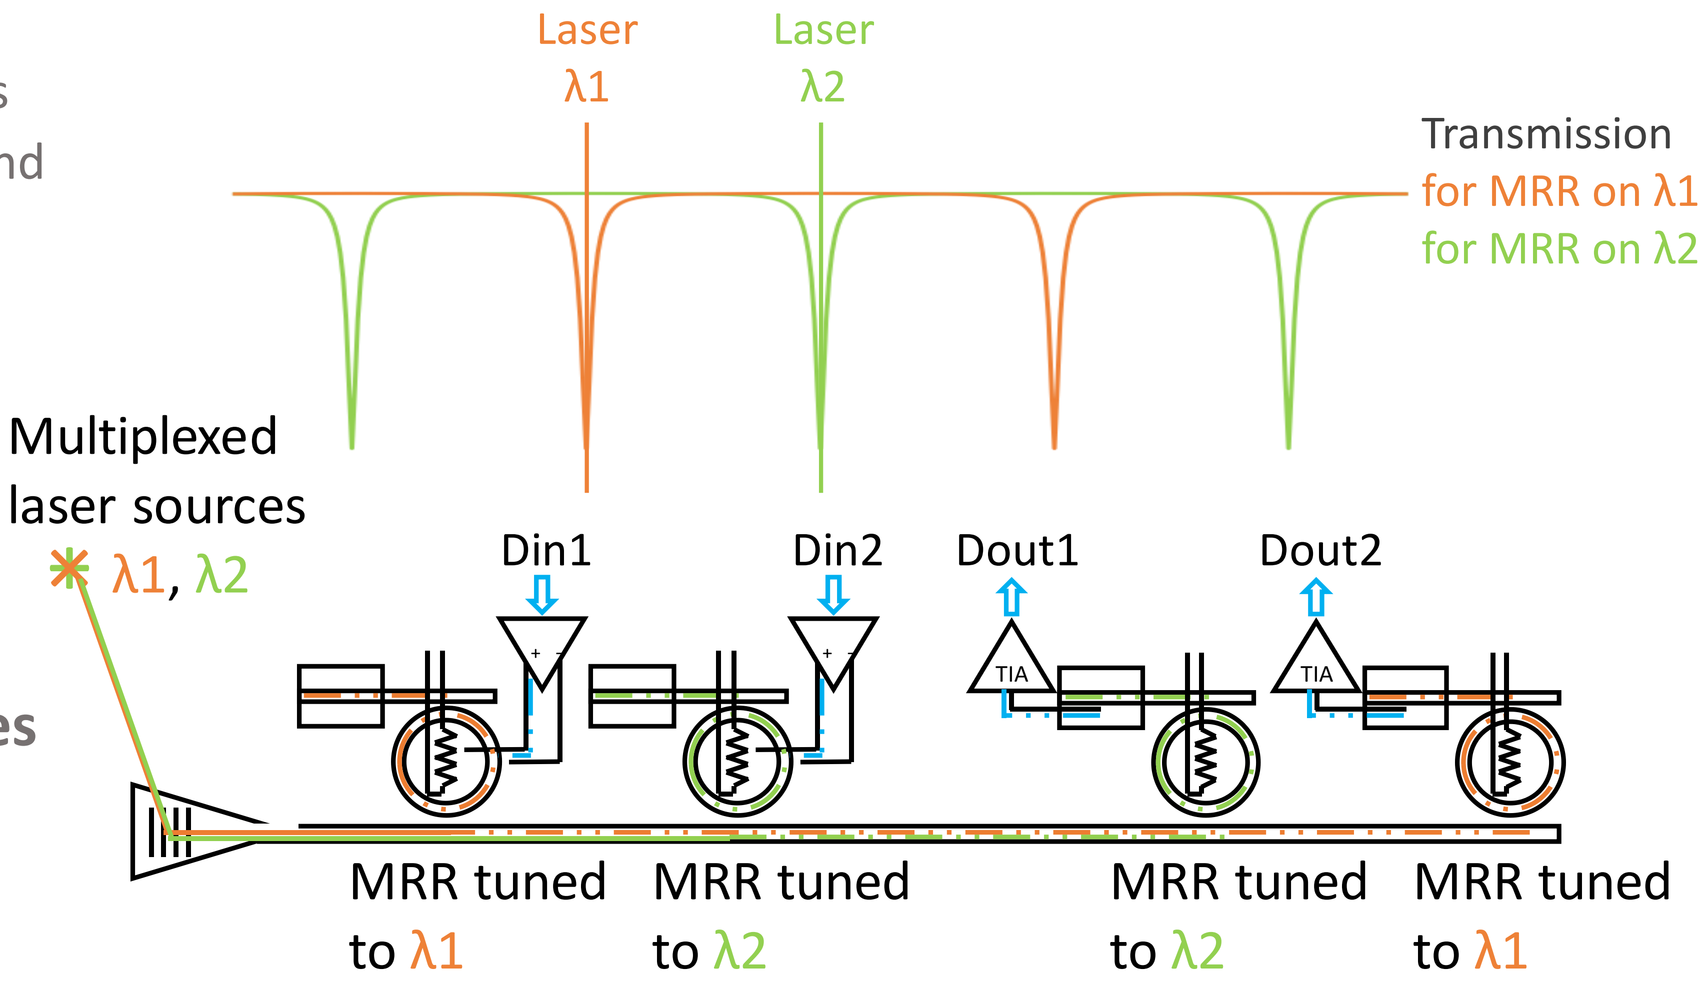

# Process variability impact on MRR resonance

- **High dependence on process variability**
  - MRR resonance can shift by about 1nm per nm of thickness
- **Geometrical variability : wafer scale characterization**
  - Identical MRR resonances characterized around 1310nm
  - FSR ~ 7.2nm
  - Variation of resonance across 5cm < 2nm in average
  - Worst-case geometrical variation: **75pm/mm**
- **Random variability : close identical rings**
  - Resonance difference of identical adjacent MRRs
  - Random variation : standard deviation: **σ=60pm**

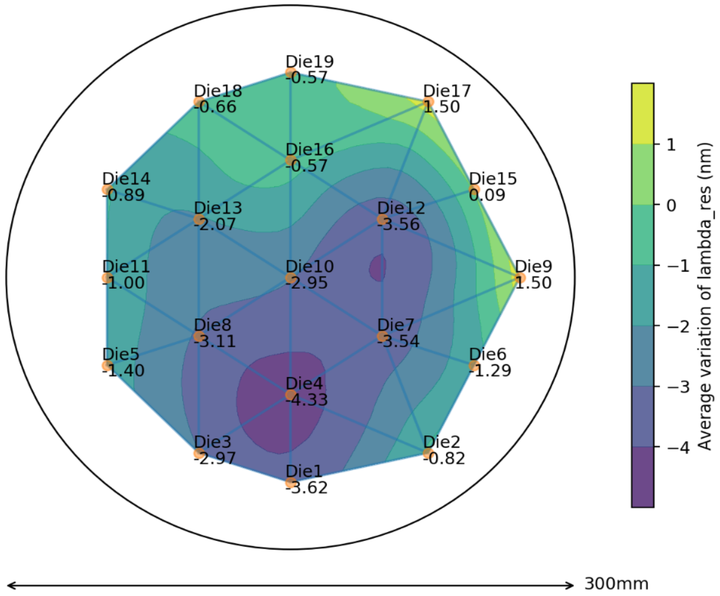


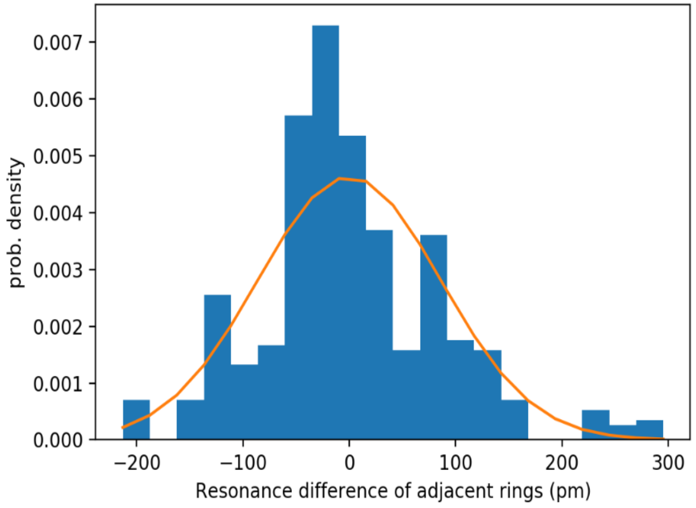

# MRR groups, WDM and crosstalk

- Q-factors > 13,000 @ 1310nm => 3dB bandwidth < 100pm
- For **~0.1dB crosstalk** => ~7x margins+3σ random variation
- up to **10-16 wavelengths** for 10nm FSR

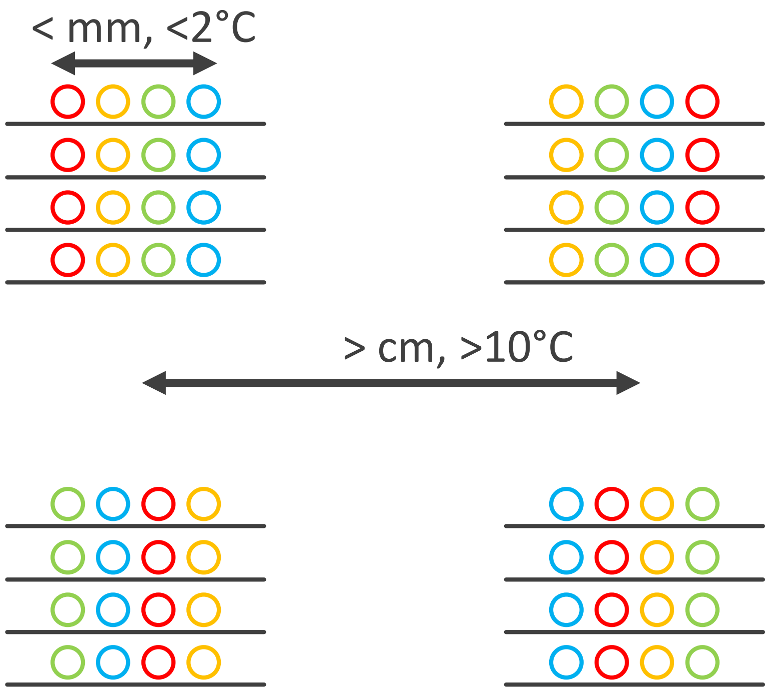


- MRR groups within 1mm distance
  - MRRs have little geometrical varability
  - Local temperature effects are smoothed & almost uniform
  - MRR groups operate consistently wrt wavelengths
- Two **MRR groups** should have **no wavelength relationship**
  - Geometrical variability becomes dominant in the cm range (>750pm)
  - Temperature effects may show large local differences.

# Modulator principle, Thermal sensitivity dλ/dv & dλ/dt measurements

▶ **13.6 pm/V** modulation efficiency

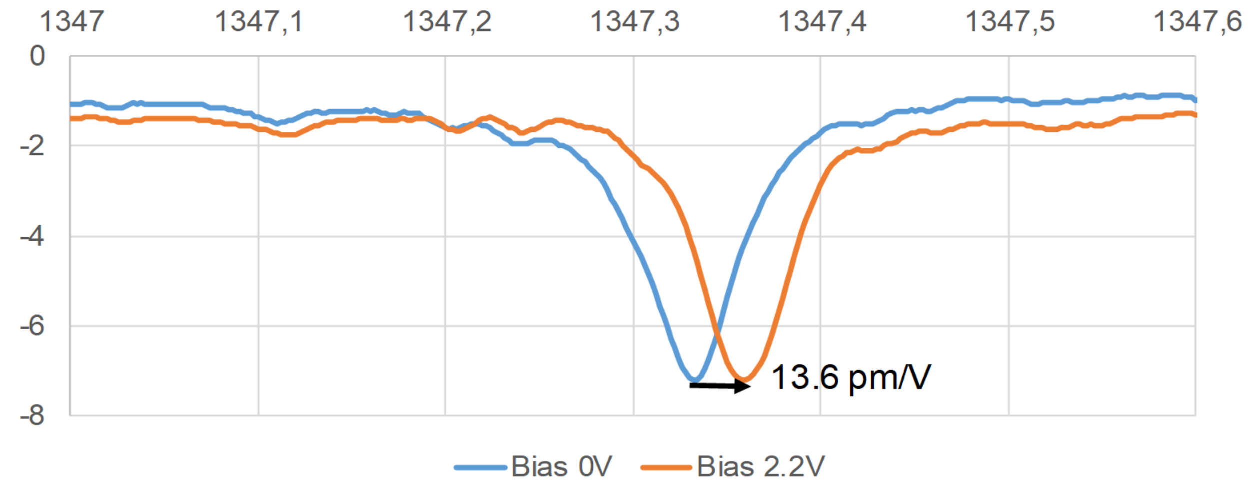


▶ **76 pm/K** thermal sensitivity

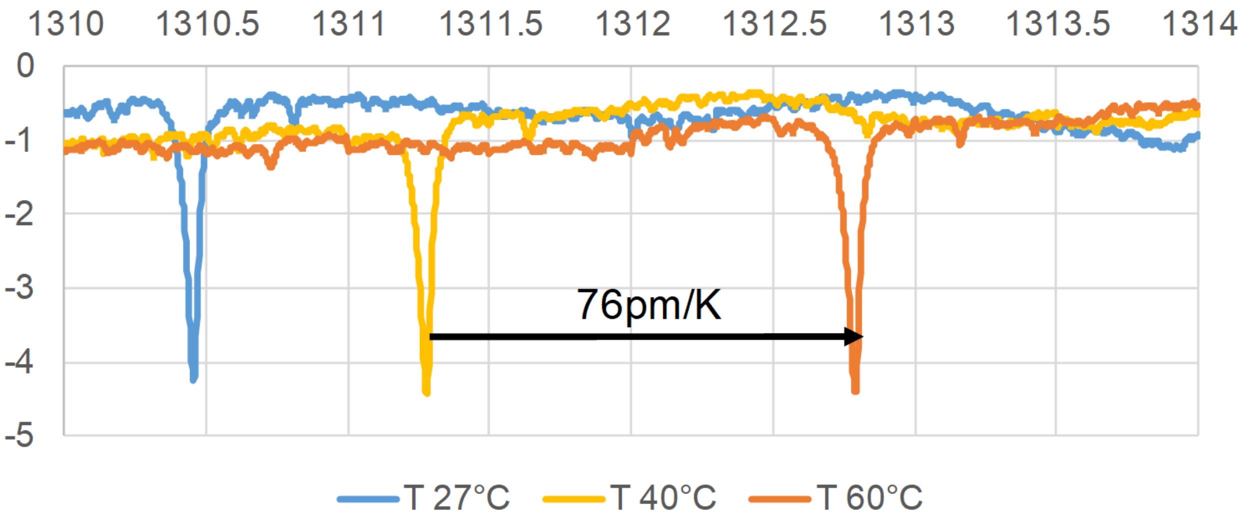


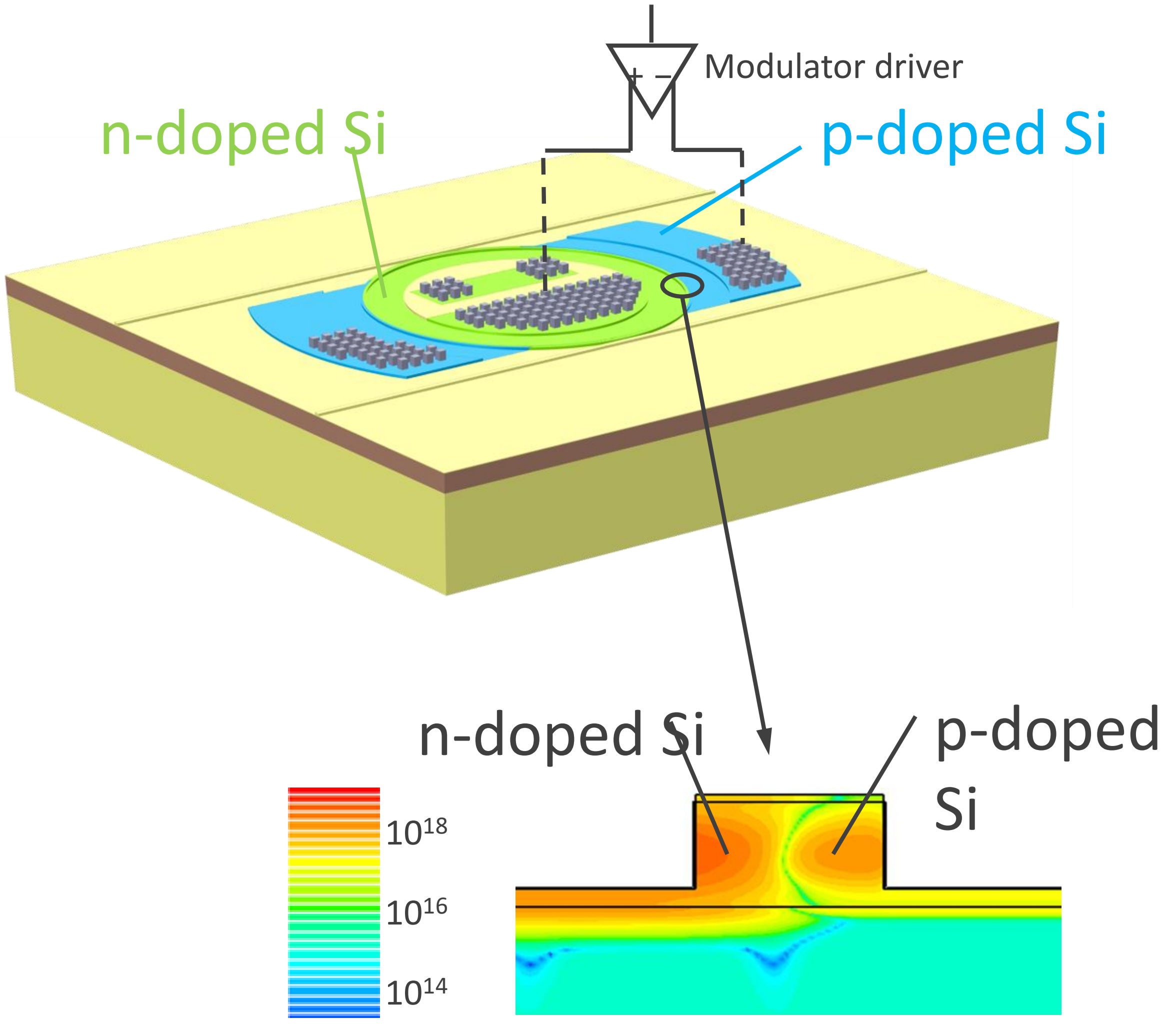

# Heater efficiency & dλ/dP measurements

- **Heating using a titanium loop resistor**
  - 120Ω Resistive path 900nm above the ring
- **Average ring temperature increase:**
  - Measurements: dλ/dP ~ **600pm/mW**

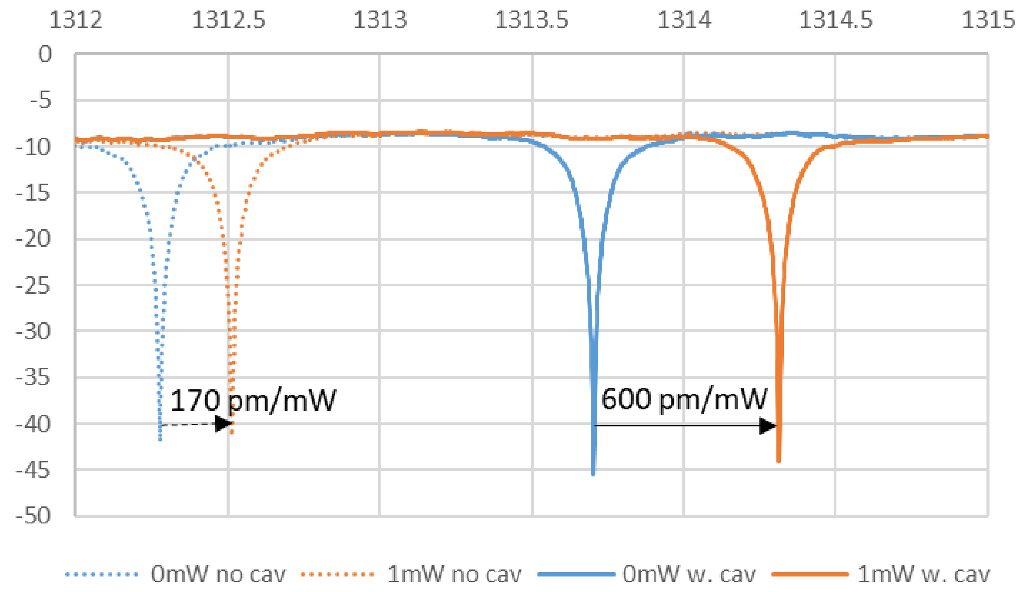


Temperature (K) and Heat Flux under 1mW heating

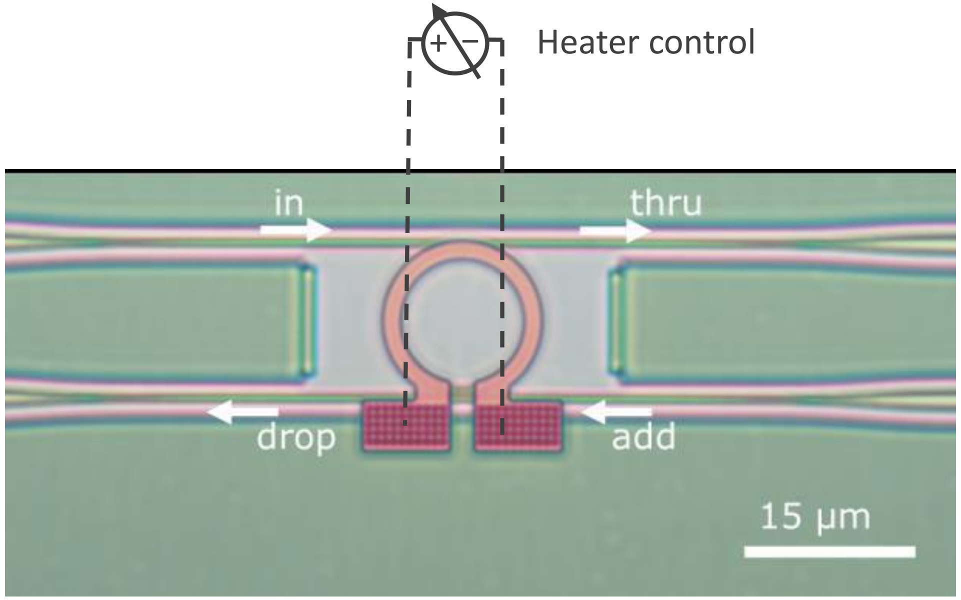


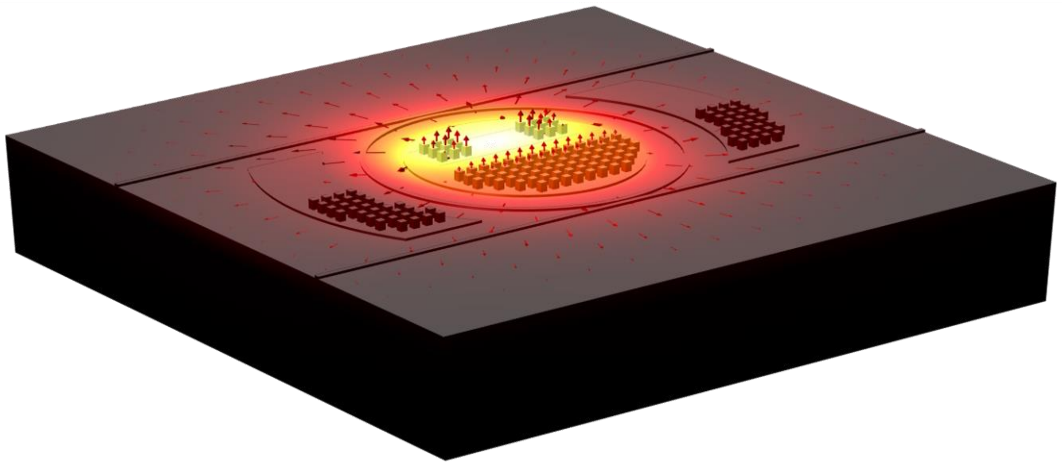

# Thermal tuning and Thermal coupling between MRR

- **Without cavity, heat flow from heater is mostly drained by substrate**
  - Efficiency limited to 170pm/mW

- **Thermal tuning range extension by back-side substrate removal**
  - Back-side cavity allows extending to ~600pm/mW,
  - i.e more than a quarter of FSR for 4mW

- **Simulations show limited thermal coupling between adjacent rings: <1pm/mW**
  - Thermal crosstalk is not an issue

A ring seen from the back side cavity

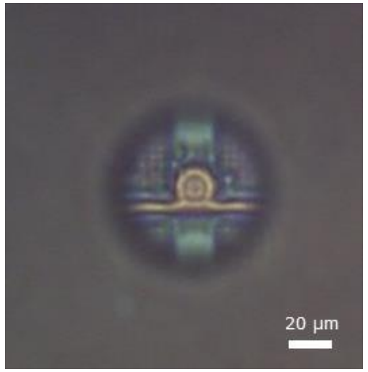


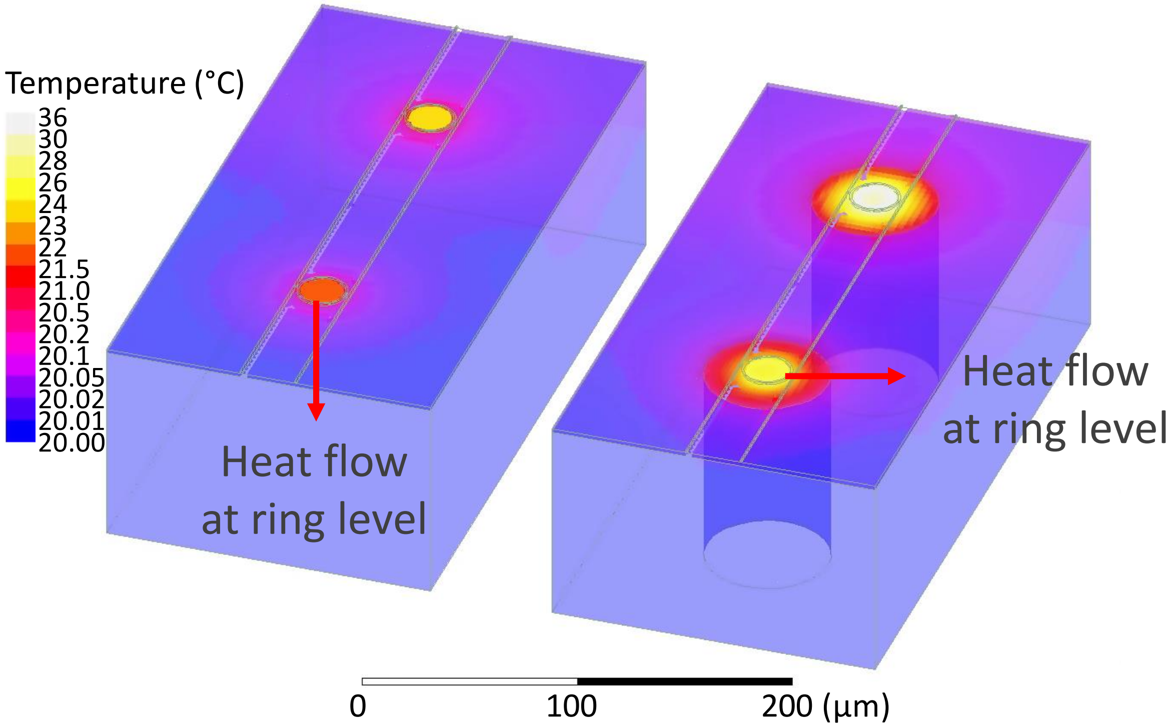

# Ring modulator operation (tuning+modulation)

- **Ring resonant wavelength unpredictible at design time**
  - 1 nm thickness variation ≈ 1 nm resonance shift
- **But finesse, free-spectral-range & amplitudes are well-controlled**
- **Thermal tuning is used to align ring resonance on laser source**
  - Low-frequency control
- **Voltage is used to modulate light**
  - High frequency modulation

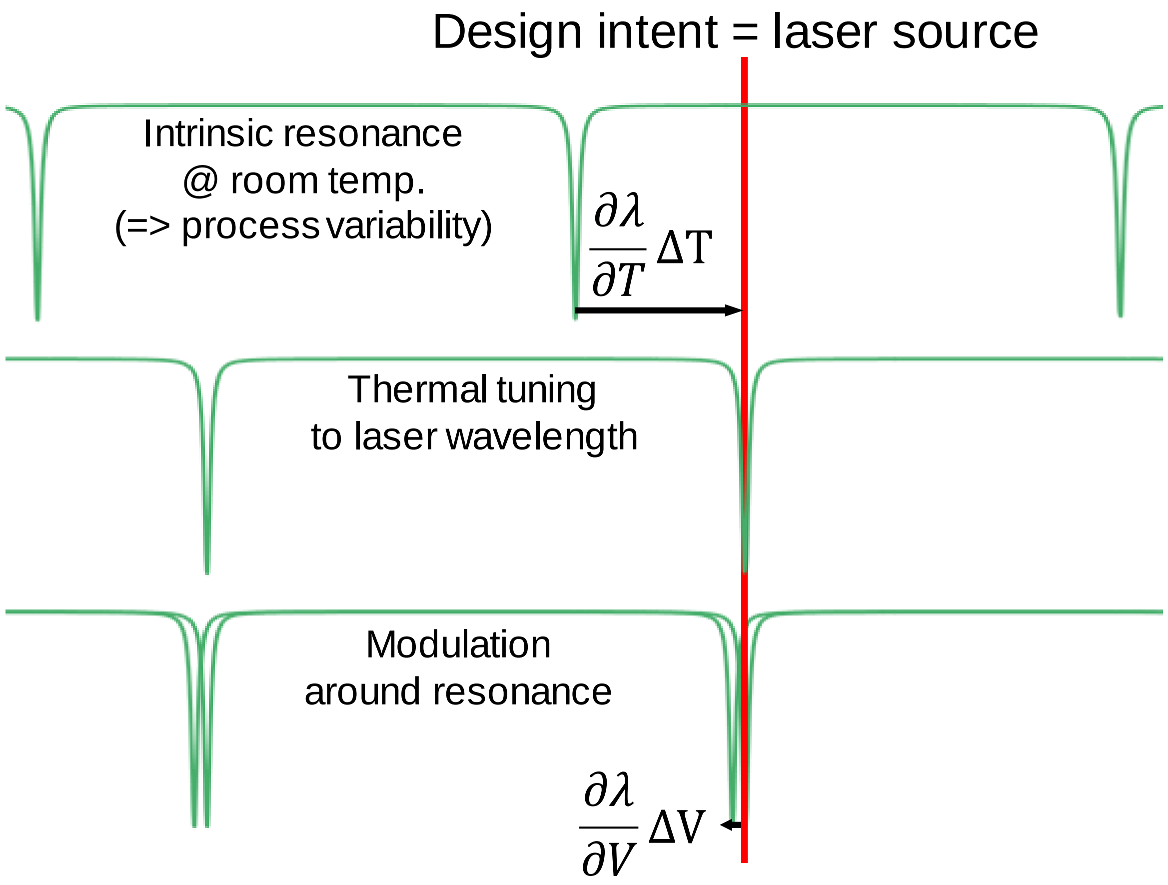

# Reduced thermal tuning cost using WDM and MRR remapping

- WDM reduces the maximum thermal shift to the closest resonance
- Closed-loop MRR tuning fixes the temperature to a stable modulation point
- Remapping occurs when ambient temperature varies too much
  - Details in [9] Thonnart et al., ISSCC, 2018.

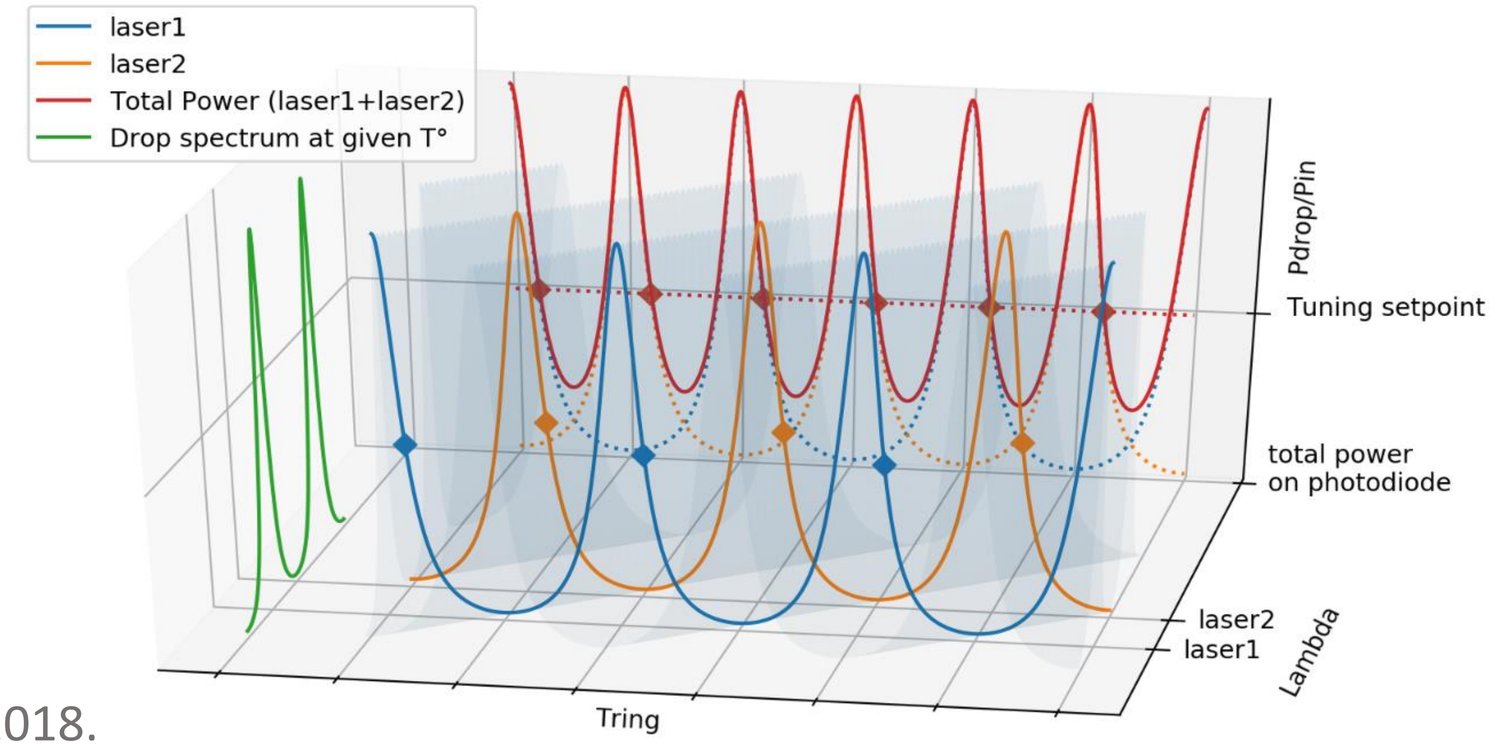


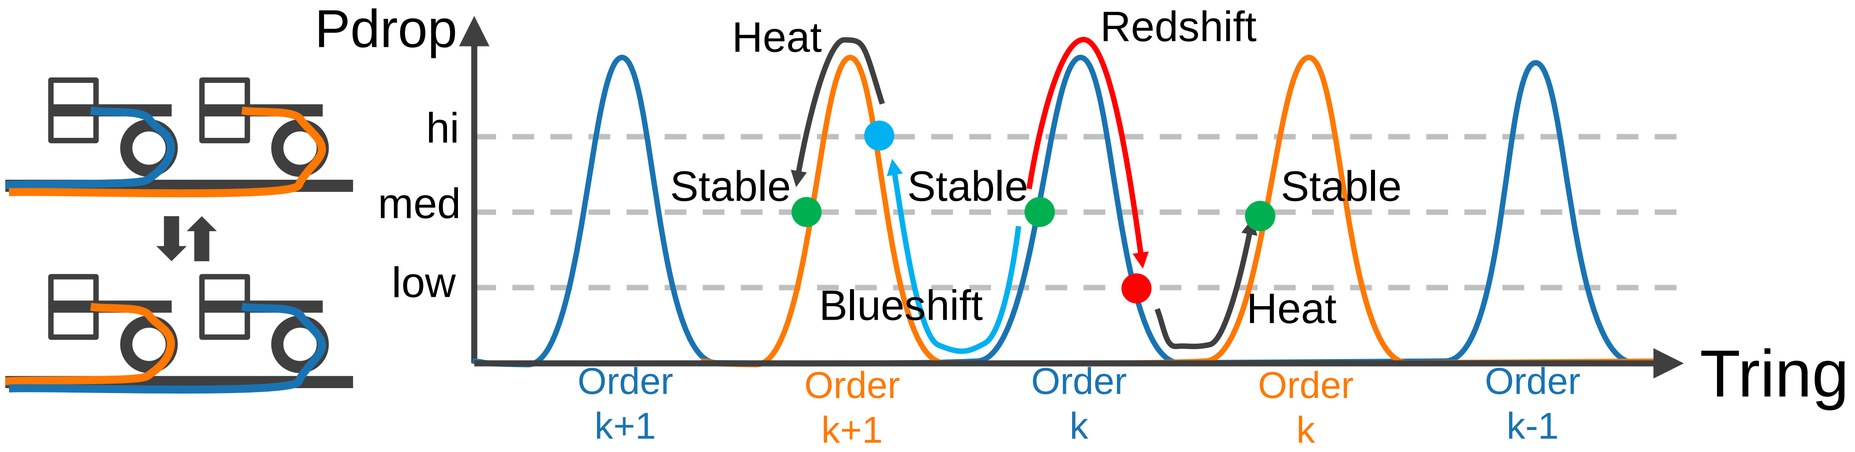

# Complete WDM link operation vs temperature

- ▶ Activity load in compute chiplets creates local temperature increase
- ▶ Closed-loop control regulates Heater power for each MRR to lock to closest laser wavelength & maintain a constant MRR temperature
- ▶ When Heater power becomes too small, MRRs are remapped to lower λ
- ▶ When heater power becomes too high, MRRs are remapped to higher λ

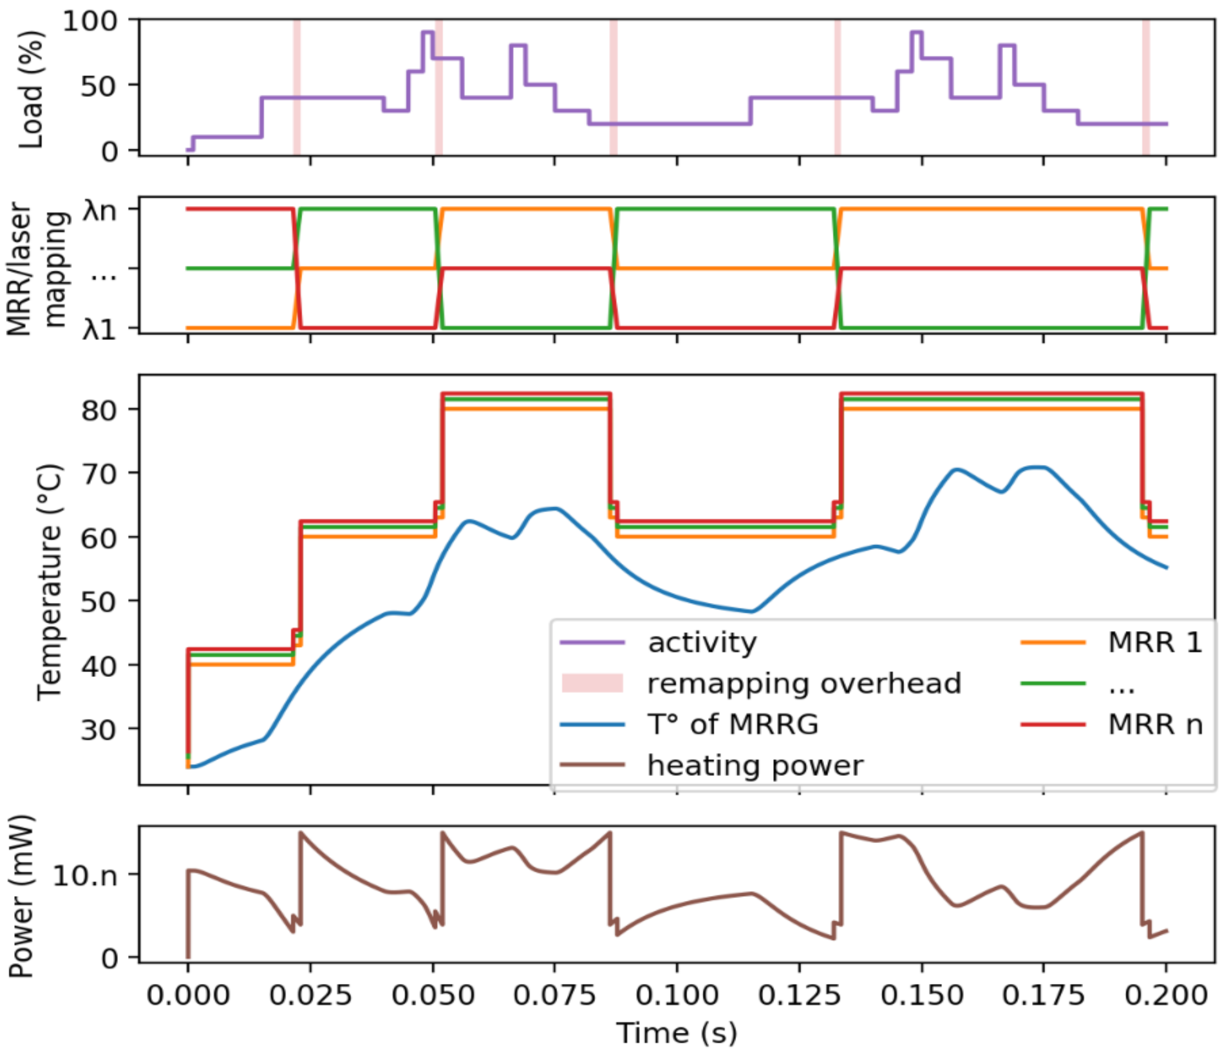

# PN MRR model for WDM modulation

- **Rings with PN junction have a low switching time**
  - Suited for data modulation > 10 Gbps
- **But modulation efficiency remains low**
  - About 20% of bandwidth per volt
- **Tradeoff to find between ‘off’ losses and extinction ratio**
- **Drop transmission correlated to peak extinction ratio**
  - Sufficient power required for tuning

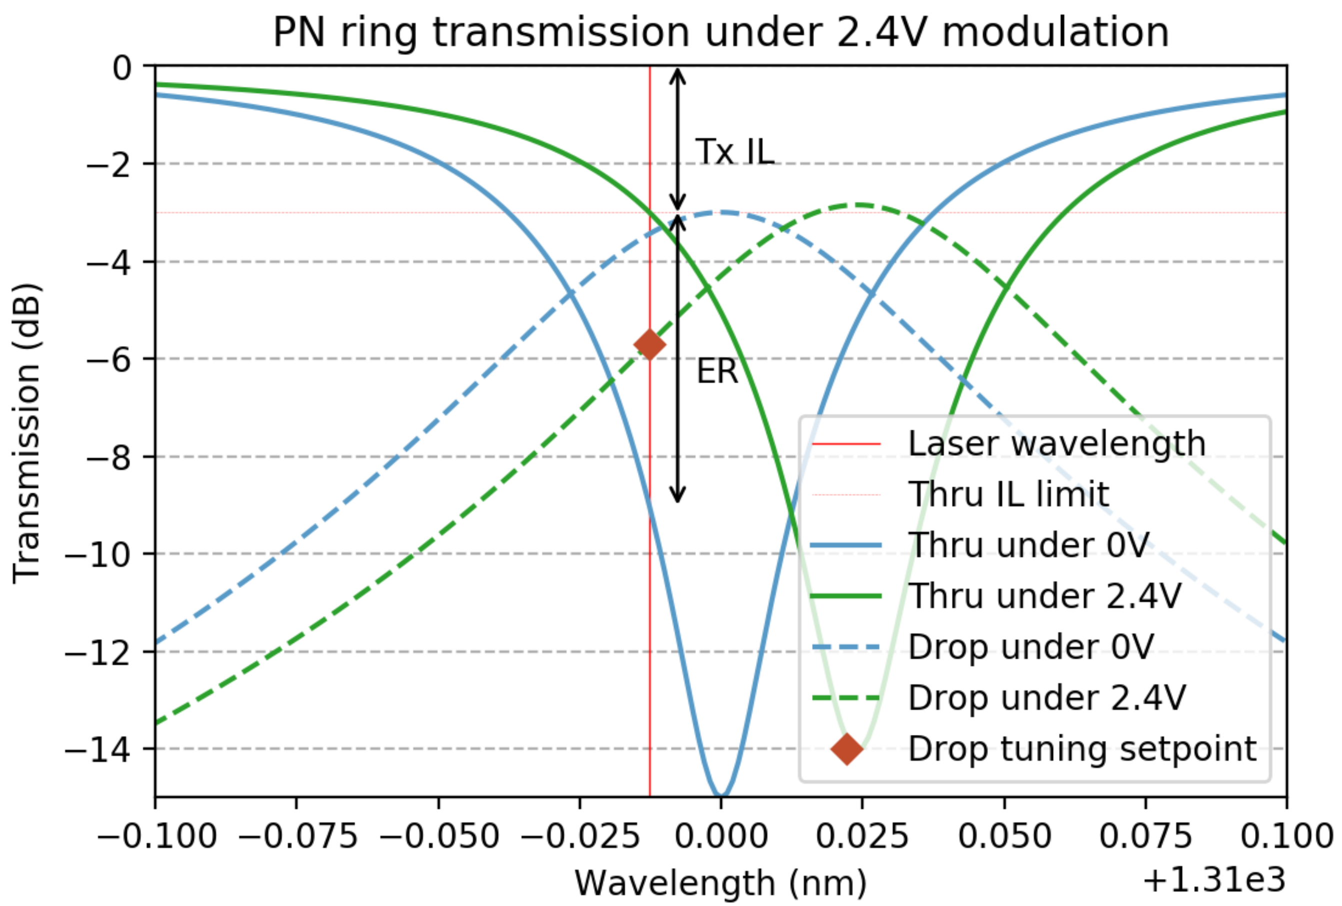

# PIN MRR model for WDM filtering and SWMR routing

- **Rings with PIN junction have a higher modulation efficiency and higher drop transmission**
  - Suited for data filtering
- **But switching speed is limited**
  - About 200 MHz
  - Higher with pre-emphasis
- **Tradeoff to find between 'off' losses and 'off' drop transmission power**
  - Sufficient power required for tuning

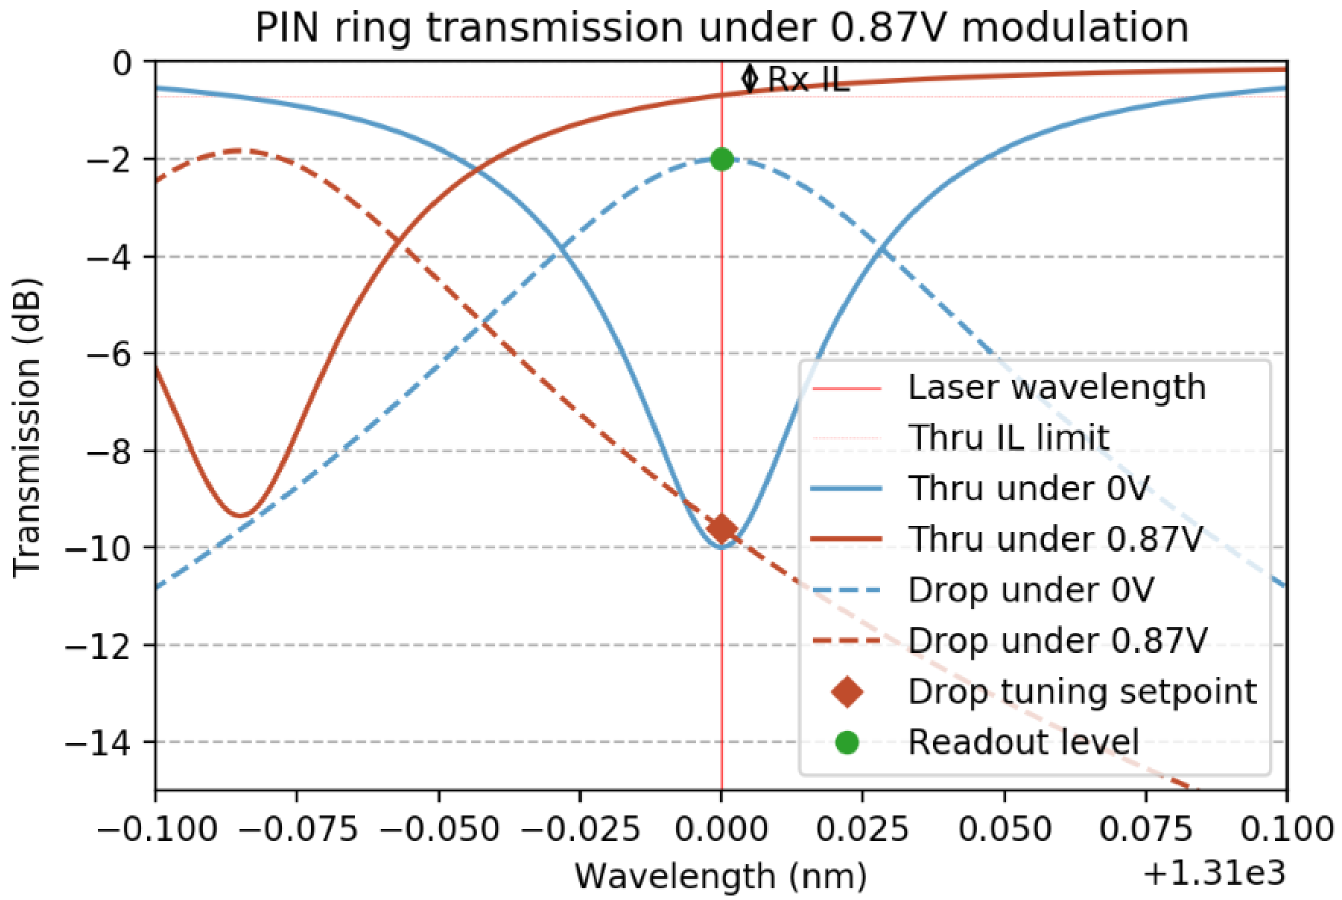

# Device Characterization from CEA-Leti Silicon Photonic platform

| Device | Parameter | Value |
|---|---|---|
| PN MRR modulator | Thru IL off-res. ("1") | -3dB |
| | Thru IL on-res ("0") | -9dB (ER 6dB) |
| | Drop IL tuning | -6dB |
| PIN MRR filter | Thru IL off-res. (deselected) | -0.7dB |
| | Drop IL on-res (selected) | -2dB |
| | Drop IL tuning | -10dB |
| Waveguide | Straight losses | -0.11dB/cm |
| | Critical radius (lossless) | 20µm |
| | Crossings (1x1 MMI) | -0.25dB |
| Grating coupler | IL | -2dB |
| Laser power | Max power in MRR | 3dBm |
| O/E sensitivity | Demod. sensitivity (10Gbps) | -15dBm |
| | Tuning sensitivity | -18dBm |

# POPSTAR architecture motivation

- **Limit the number of rings on an optical path**
  - Due to insertion losses
- **Limit as much as possible the number of « crossings », i.e. drop paths**
  - Due to drop losses
  - Favor single waveguide transmission
- **Favor PIN over PN on link budget**
  - Because of lower insertion losses
- **Consequently, use SWMR topology**
  - PN rings at Tx, PIN rings at Rxes

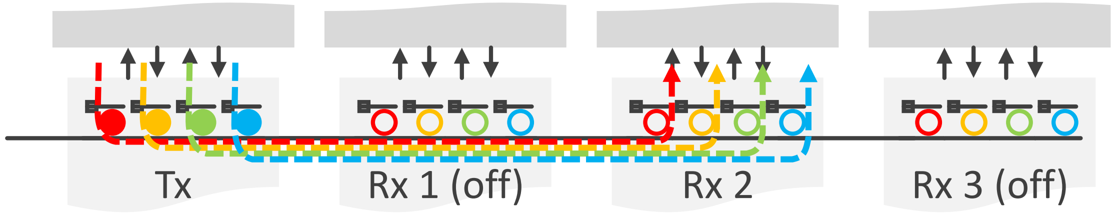


- **Use all WDM wavelengths as a single data bus**
  - No wavevength routing
  - To limit the remapping overhead
  - To avoid global communication synchronization
  - To enable decentralized local arbitration

# POPSTAR photonic interposer electro-optical architecture

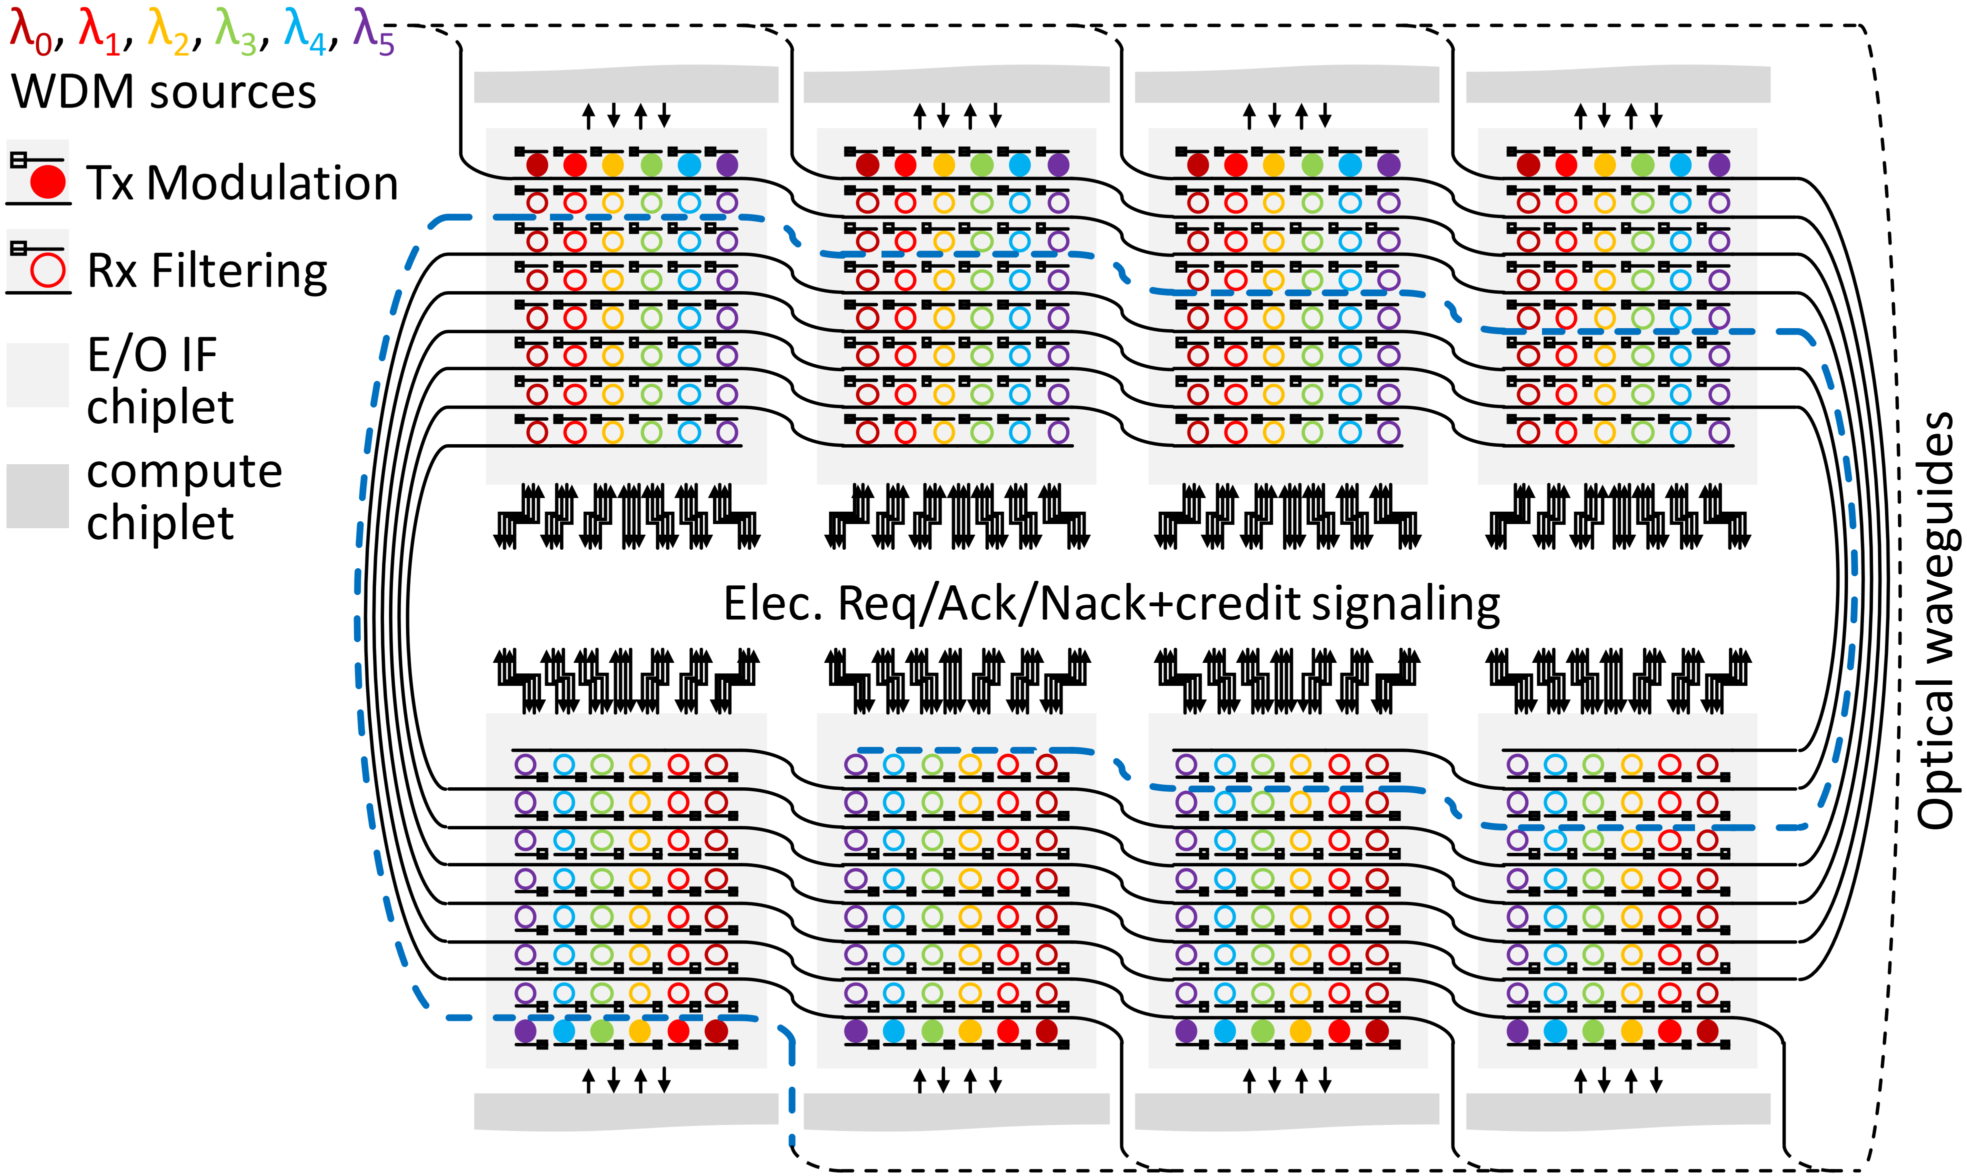

# POPSTAR photonic interposer electro-optical architecture

- **An 8-port POPSTAR instance on a folded ring**
  - 6 wavelengths are used for each channel
  - 384 microrings in total
  - Up to 1Tbyte/s @ 12GBaud aggregate bandwidth

- **6 compute chiplets, 2 external IO interfaces**

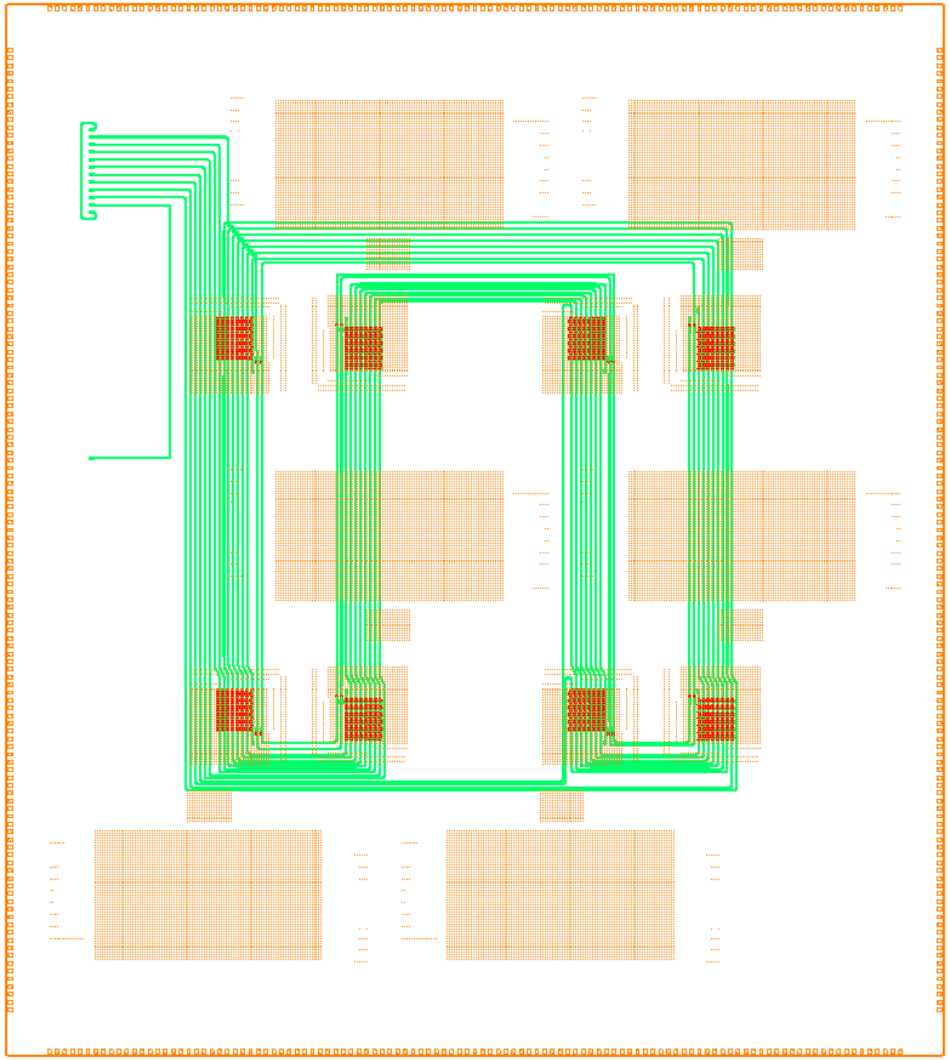

# POPSTAR photonic interposer electro-optical architecture

- **An 8-port POPSTAR instance on a folded ring**
  - 6 wavelengths are used for each channel
  - 384 microrings in total
  - Up to 1Tbyte/s @ 12GBaud aggregate bandwidth

- **6 compute chiplets, 2 external IO interfaces**

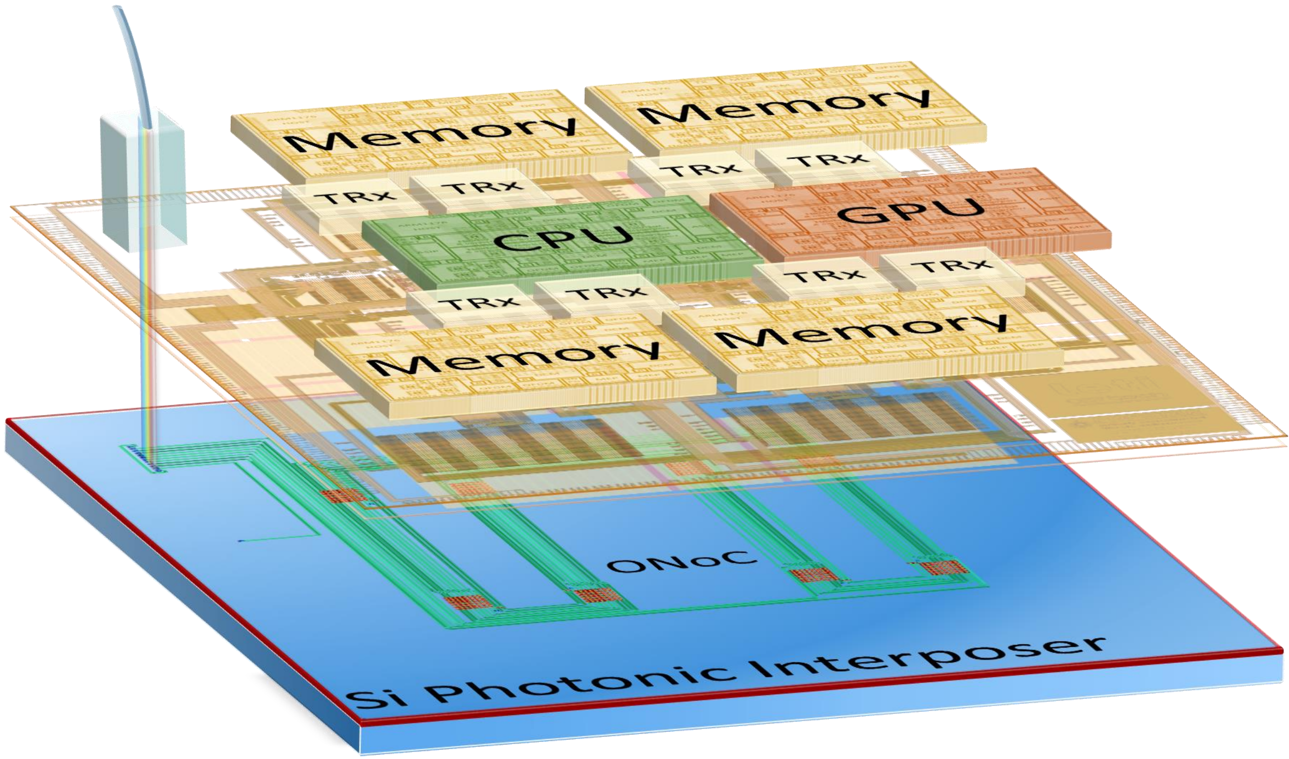

# POPSTAR Optical budget for SWMR topology

- **A single Tx IL**
  - Tuning on drop port close to resonance
- **A series of inactive Rxes**
  - Tuning for low IL
- **Low drop losses on active Rx**
  - Graph shows 0 and 1 levels based on Tx modulated data

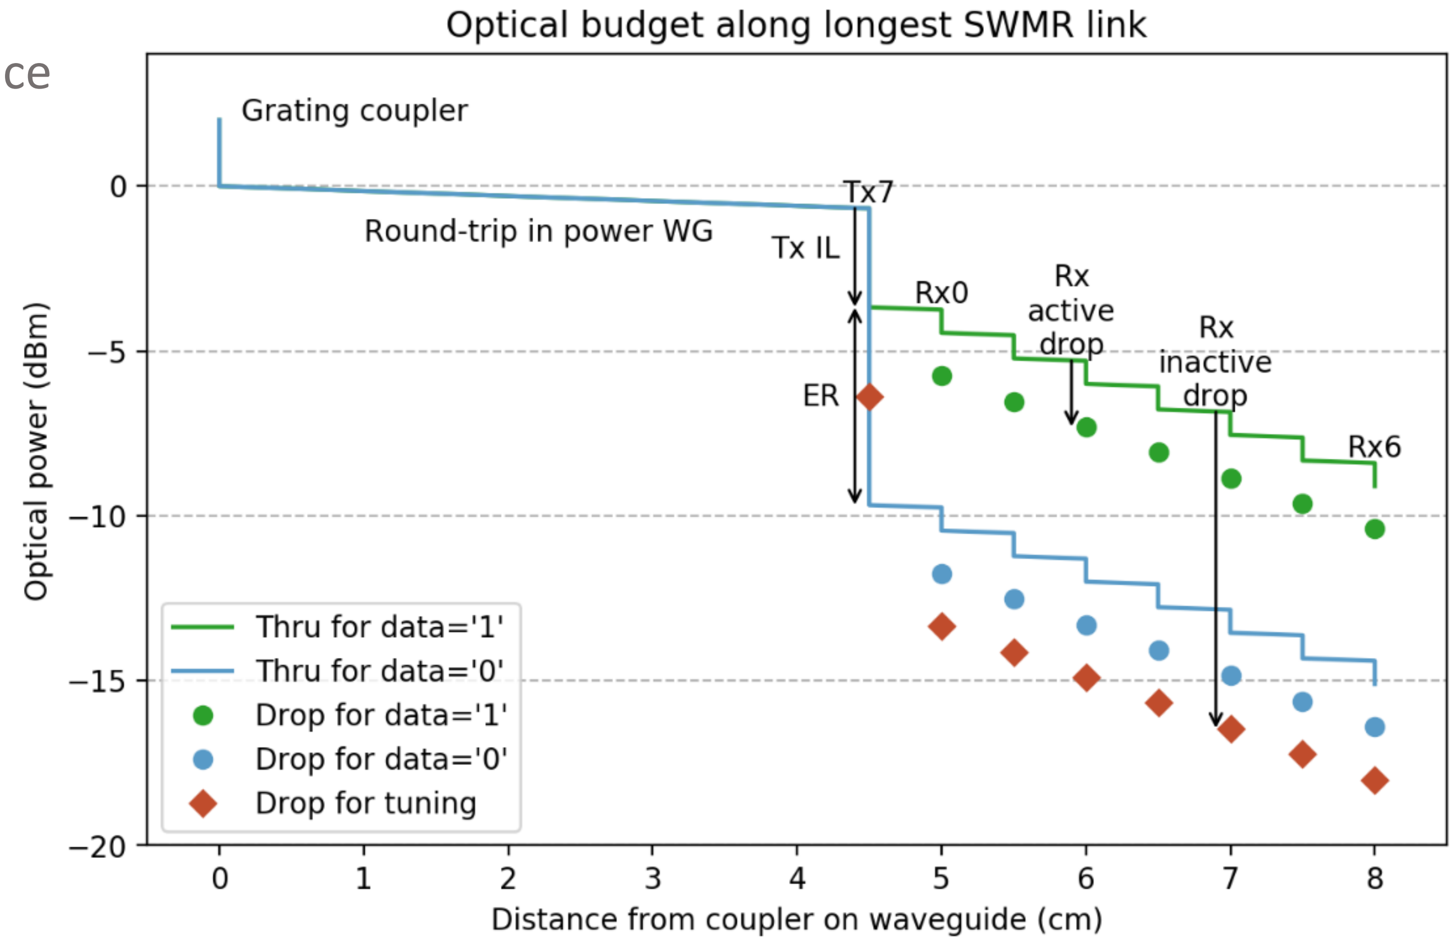

# POPSTAR communication protocol between E/O chiplets

- **Data preamble for channel setup**
  - Optical 1010101011 sequence
  - Low-level synchronization
  - Wavelength identification
- **End-to-end flow-control**
  - Using low-latency metal signaling
- **Req/Ack/Nack protocol**
  - Optical CRC encoded in the data
  - Possible retransmission
  - Store & Forward to Compute chiplets
- **Credits for low-latency**
  - Receiver always ready to receive
- **Virtualization of traffic classes**
  - Virtual channels for credits & S&F buffers

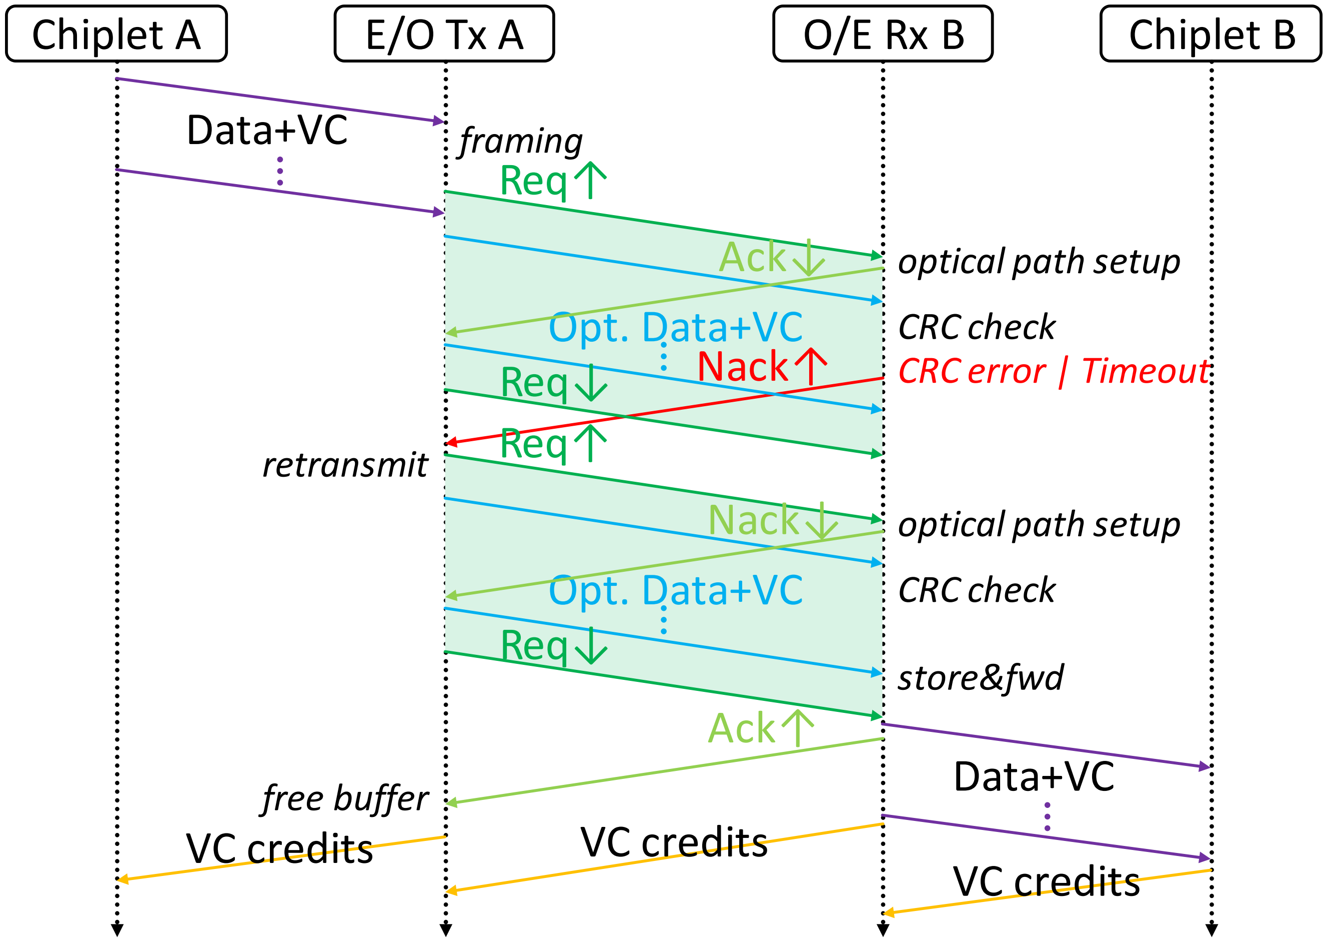

# TxRx E/O/E chiplet architecture

## ▶ Tx architecture

- VC buffering
- Flow control based on credits
- Preamble encoding
- Wavelength identification
- CRC encoding
- Serialization
- E/O driving
- Thermal tuning

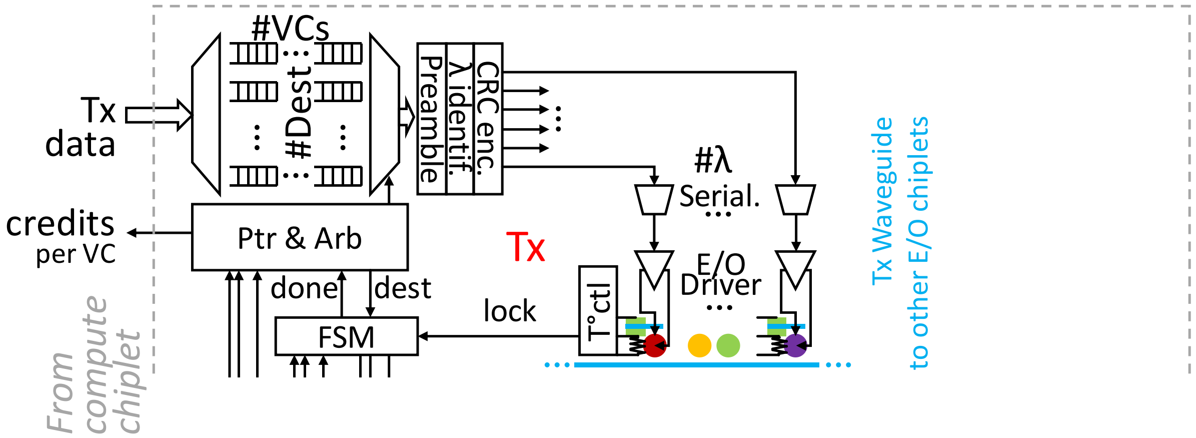

## Rx architecture

- O/E demodulation
- Thermal tuning
- Deserialization
- Resynchronization
- Preamble decoding
- Wavelength identification
- CRC decoding
- VC buffering
- Flow control based on credits

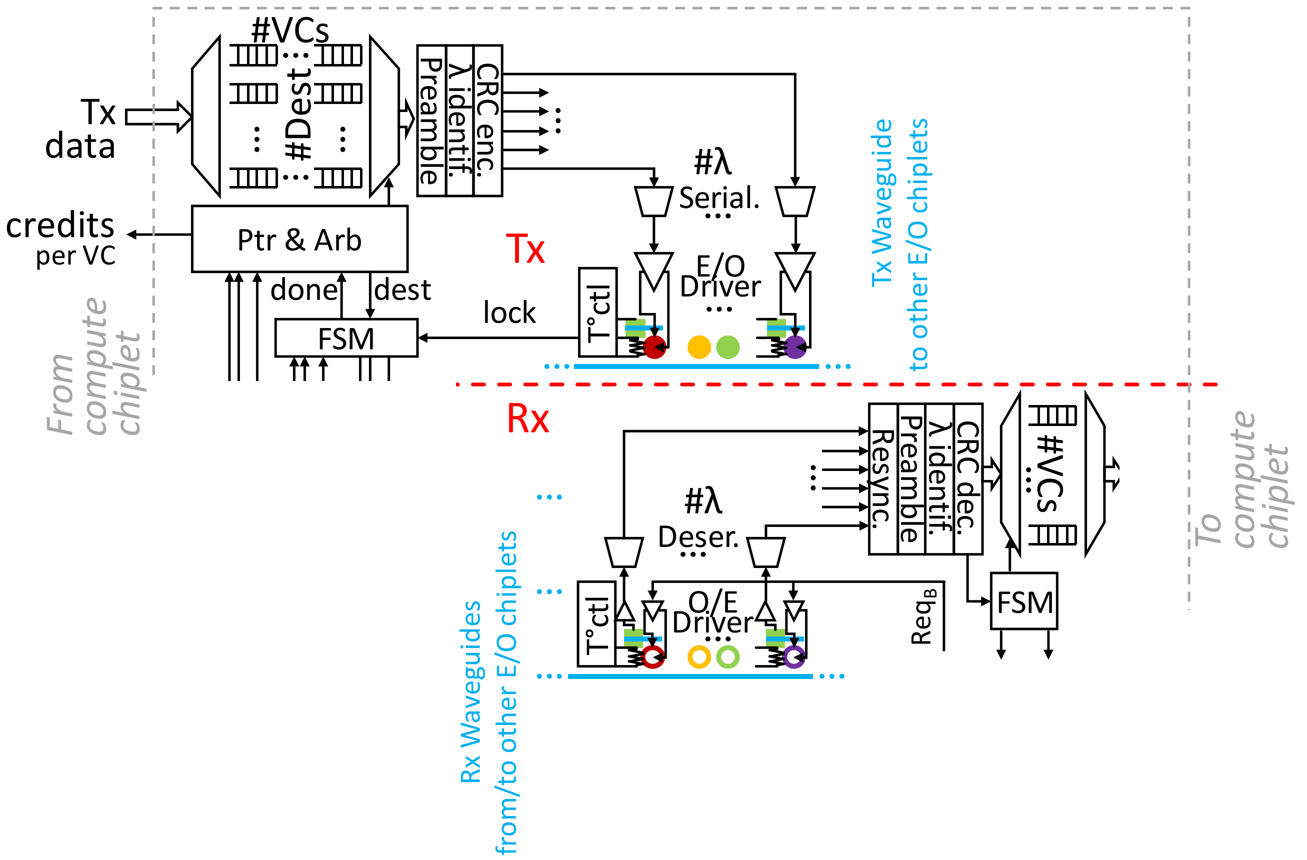

# TxRx E/O/E chiplet architecture

- Arbitration between different Txes on separate Rx channels
- Req/Ack/Nack for each pair of Tx-Rx
- Minimum end-to-end latency: **12 cycles**
  - Can be higher with S&F & queuing

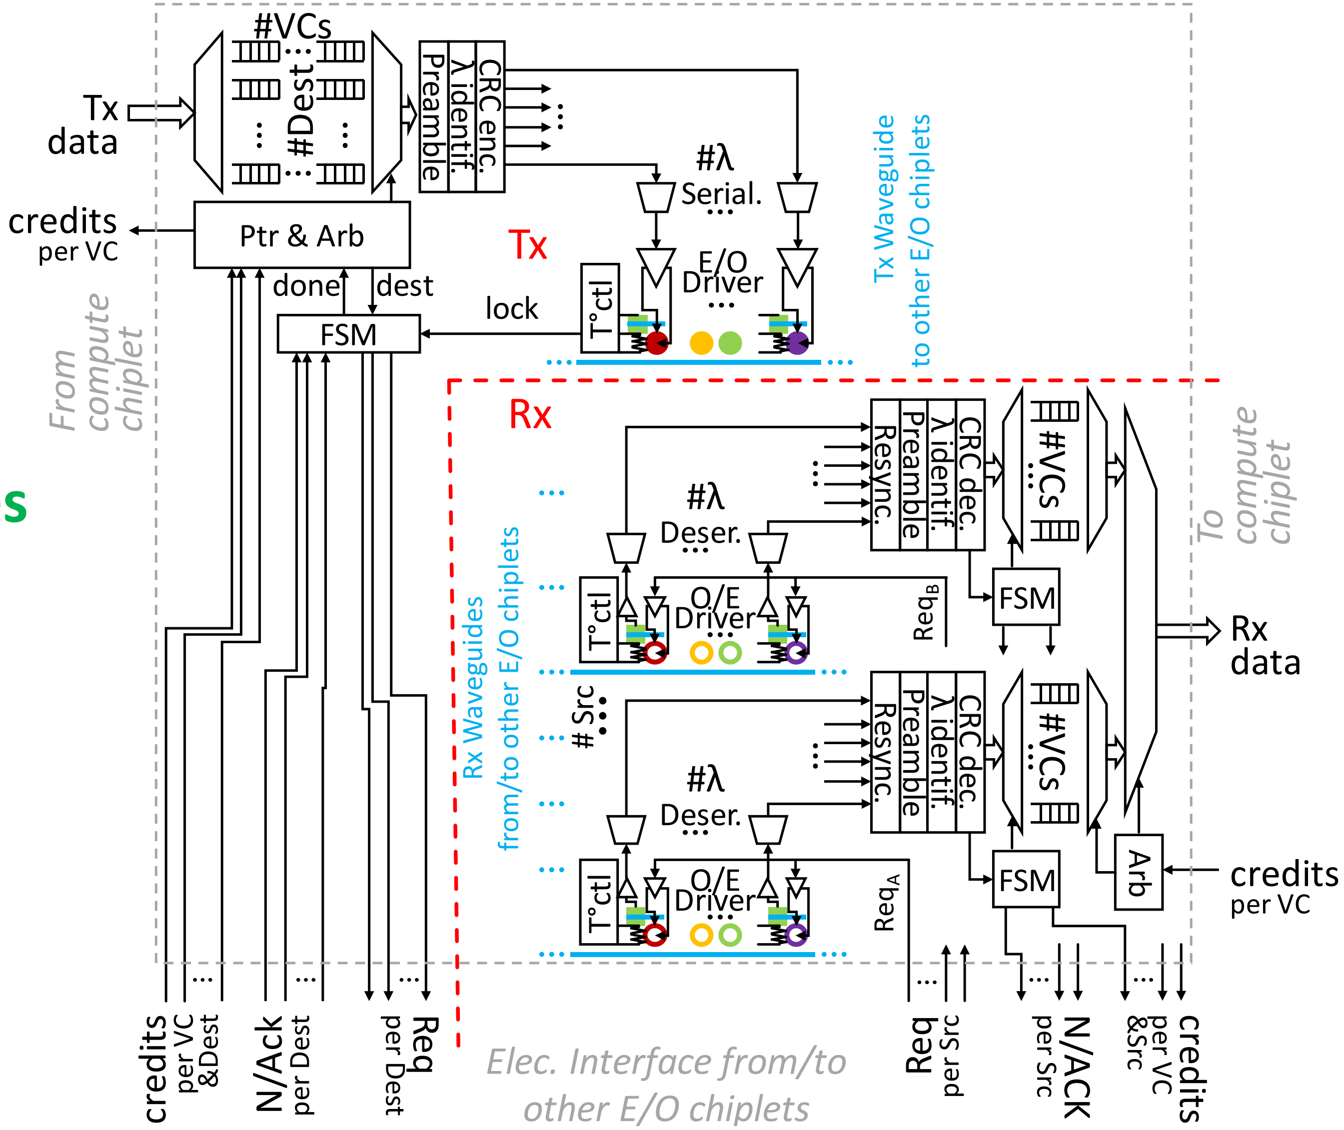

# Conclusion

- **POPSTAR: a robust modular architecture for E/O communication between chiplets on a photonic interposer**
  - Based on technological constraints
  - To cope with process and thermal variability
  - Independent SWMR links for coordinated WDM remapping
- **A standard replicable E/O chiplet to interface with compute chiplets**
  - In charge of routing, flow-control and arbitration
  - Low-latency non-blocking distributed crossbar
  - Arbitration contained in Rx chiplets without in-network contention
- **Low-latency communication for large-scale chiplet-based 3D systems**
  - New architecture opportunities for data-intensive high-performance applications