DLPS: Dynamic Laser Power Scaling for Optical Network-on-Chip

Fan Lan[†], Rui Wu[§], Chong Zhang[§], Yun Pan[†], and K.-T. Tim Cheng[§*]

[†] Zhejiang University

[§] University of California, Santa Barbara

[*] Hong Kong University of Science and Technology

Outline

- Background of optical interconnect
  - Components in an optical link
  - Power consumption of a laser
  - Power consumption of an optical link
- Laser power control strategy
  - Four operation modes of a laser in an optical link
  - Dynamic power scaling strategy
- Experiments
  - Case study – an exemplar optical NoC architecture
  - Case study – benchmark FFT-1024
  - Public benchmark suite
- Conclusion

Photonic integration for optical interconnects

- Fully integrated WDM transceiver network for inter- and intra- optical link
- Compact footprint with 300+ active units (48 DFB, 93 EAM, 67 PD) and 120+ passive units (7 AWG, 1X15 MZI switch) on single chip;
- Total 2.56 Tbps (8 × 8 × 40 Gbps) capacity of the photonic NoC chip.
- This technique is scalable and potential for low energy consumption communications (~pJ/bit), dominated by the laser power at the photonic circuit side (regardless the CMOS driver design)

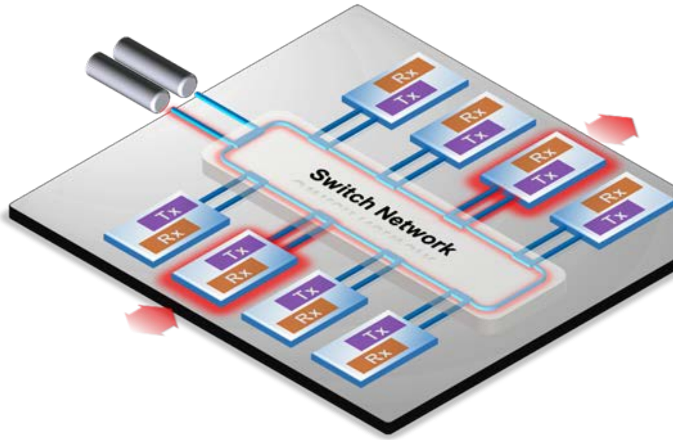

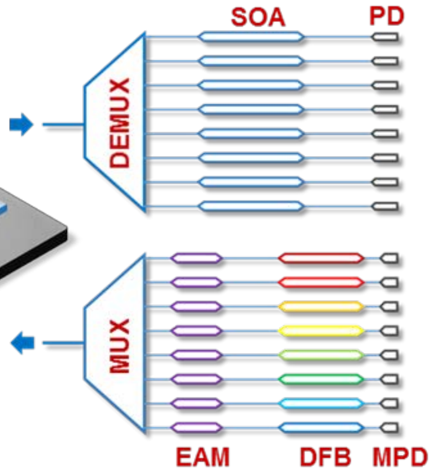

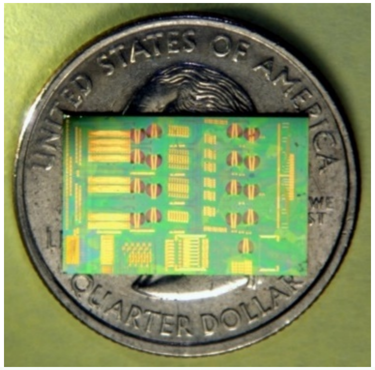

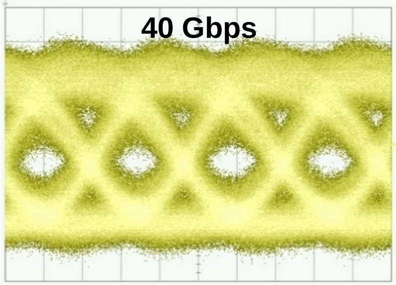

Laser: pJ/bit, PD/MOD ~fJ/bit

C. Zhang, et al., Optica 3(7), 2016

Components in an Optical Link

- On-chip lasers (> 10 mW [1])
- SerDes (~2.2 mW [2])
- Photodetectors
- TIA (~1.3 mW [3])
- Modulator and filter
  - Microrings-based
  - Need tuning
- Tuning (~10 uW [4])

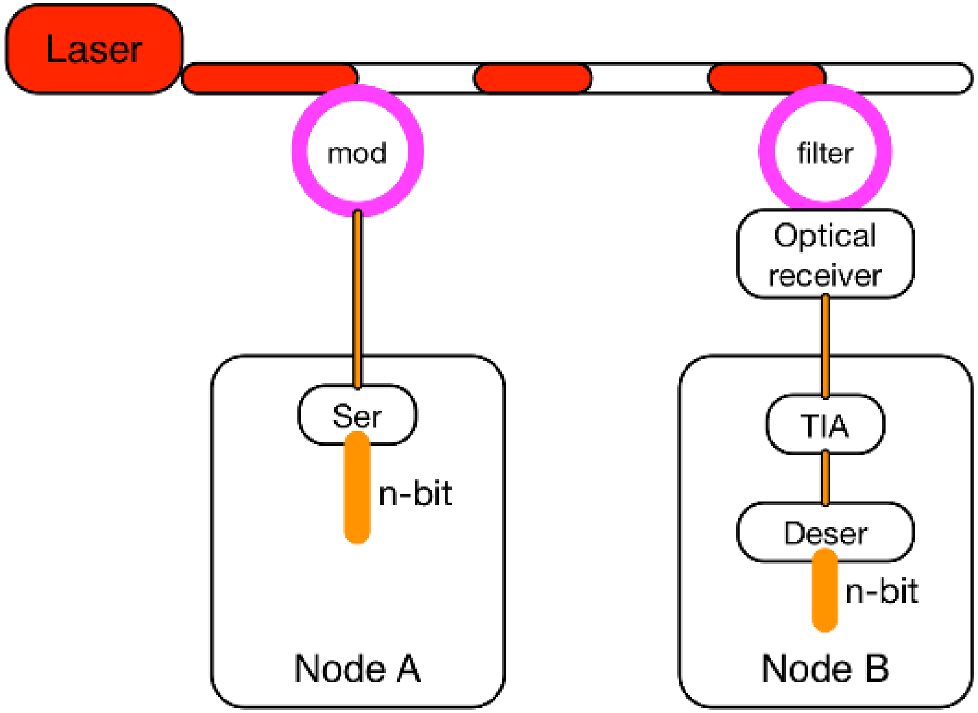

- *Lasers consume most power in an optical link*

[1] C. Zhang, et al., Optical Express, 2014
[2] A. Carpenter, et al., Int. Symp. Comput. Archit., 2011
[3] C. Li, et al., ISSCC, 2013
[4] Y. Zheng, et al., DATE, 2014

Power Consumption of a Laser

- Electric power consumption $P_{in} = I^2 R_S + IV_d$
  - Laser diode injected current – I

- Optical power output $P_{out} = \frac{h\nu}{q} \eta_d' \left(I - I_{th}' \right)$
  - Efficiency – η_d'
  - Threshold current – I_{th}'

- η_d' and I_{th}' decrease as I increases
  - Due to increased temperature
  - Leading to the roll off curve

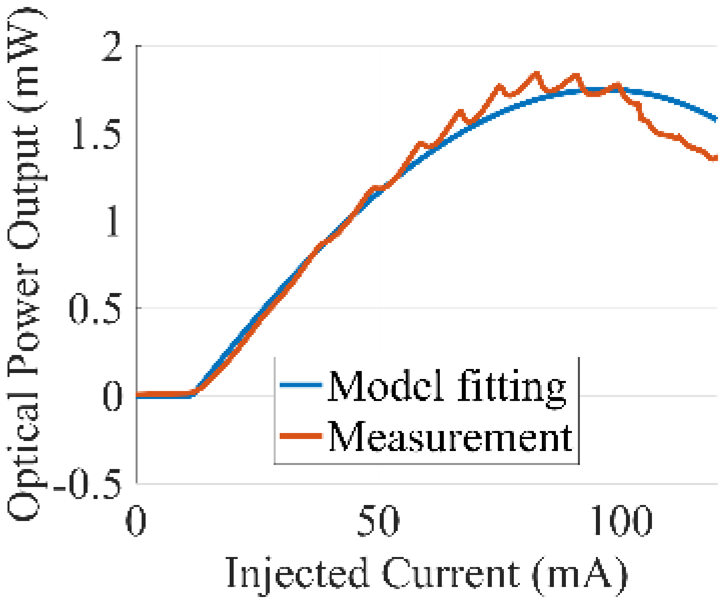

Power Consumption of an Optical Link

- Optical power vs. data rate for an optical receiver*
  - At a fixed BER=10^{-12}
  - P_{out} ≥ optical power
- Total power consumption of an optical link:

$$P_{link} = P_{in} + P_{TIA} + 2P_{tuning} + P_{SerDes}$$

  - $P_{link} = P_{in} + P_{TIA} + 2P_{tuning} + P_{SerDes}$ power of TIA
  - $P_{link} = P_{in} + P_{TIA} + 2P_{tuning} + P_{SerDes}$ power of tuning
  - $P_{link} = P_{in} + P_{TIA} + 2P_{tuning} + P_{SerDes}$ power of SerDes

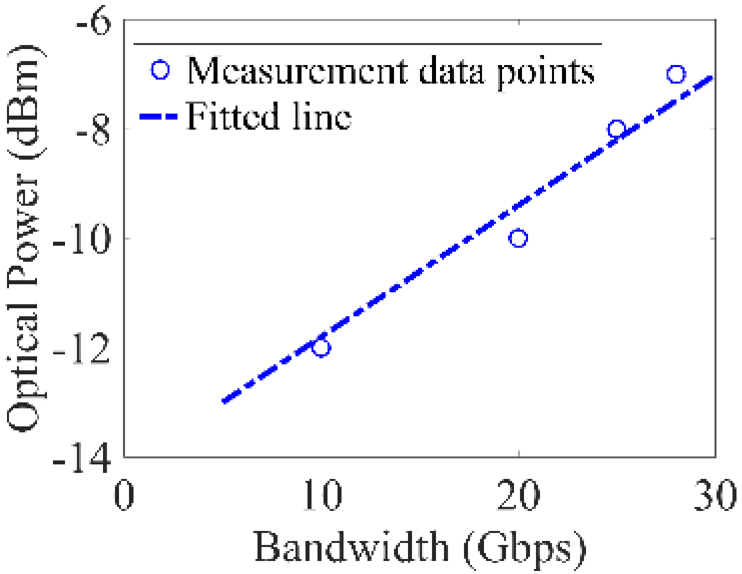

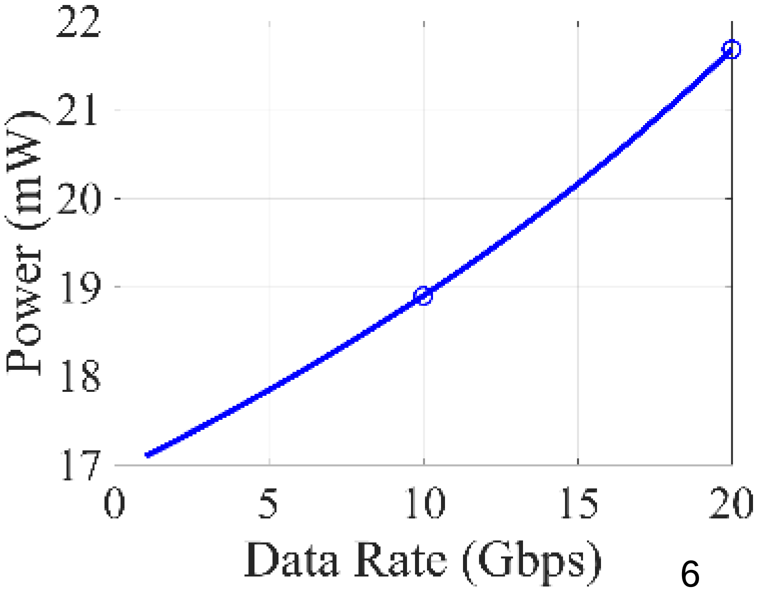

* H. Pan, et al., Optical Express, 2012

Four Operation Modes of a Laser in an Optical Link

- Full-on mode
  - Fast, power-consuming
- Off mode
  - Cannot transfer data
  - Zero power
- Standby mode
  - Cannot transfer data
  - Small power
  - Can be turned-on quickly
- Intermediate data rate mode
  - Data rate adjustable
  - Power varies as data rate changes

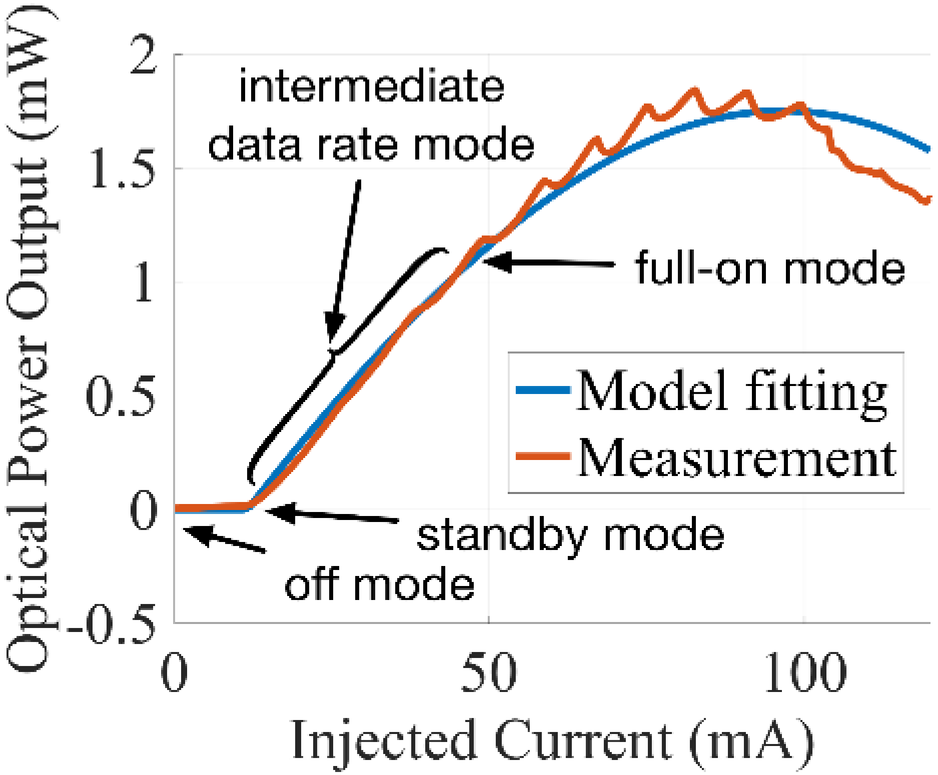

Switching Time between Laser Power Levels

40 mA DC Bias

+ 1Vpp AC Square Wave

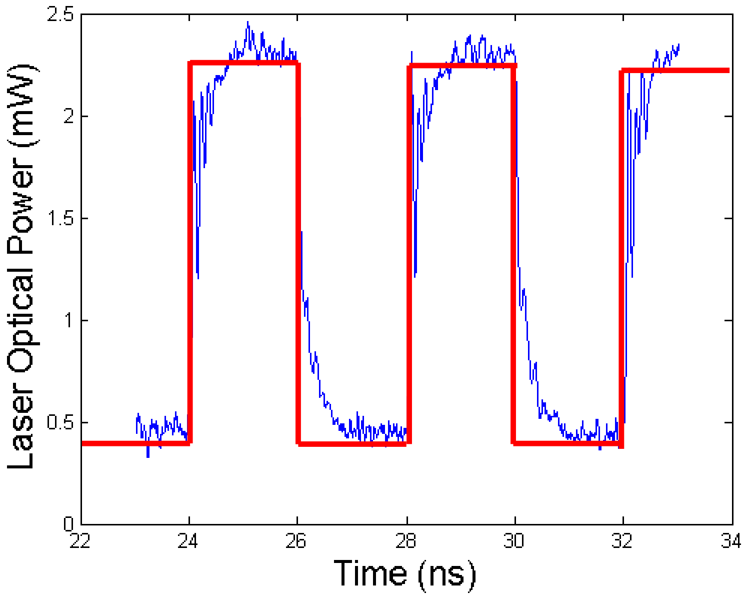

Switching time between different ON states is < 1 ns

25 mA DC Bias

+ 1Vpp AC Square Wave

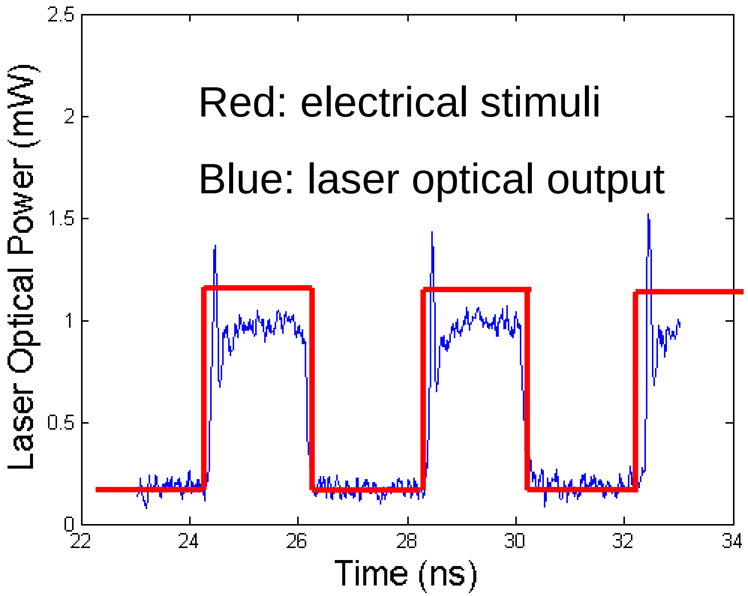

Test device parameters:
PD conversion gain 300V/W;
Total loss (fiber + coupling) ~10 dB

Typical Link Power Budget

Full-on power	21.68 mW at 20 Gbps
Intermediate power	18.90 mW at 10 Gbps
Standby power	11.61 mW
Delay from off to full-on	10 ns
Delay from standby to full-on	1 ns

- During transient period, power is estimated at the level of the target mode that the laser is switched to, e.g.
  - Off → full-on, 21.68 mW is consumed during 10 ns
  - Standby → 10Gbps, 18.90 mW is consumed during 1 ns

Dynamic Laser Power Scaling Strategy

- Data to be transmitted: switch to full-on or intermediate mode, to provide *just enough* data rate
- No data: first switch to standby mode
- Wait for *IDLE* cycles, if still no data, switch to off

- Parameters:
  - IDLE = 5
  - 3 transmissions

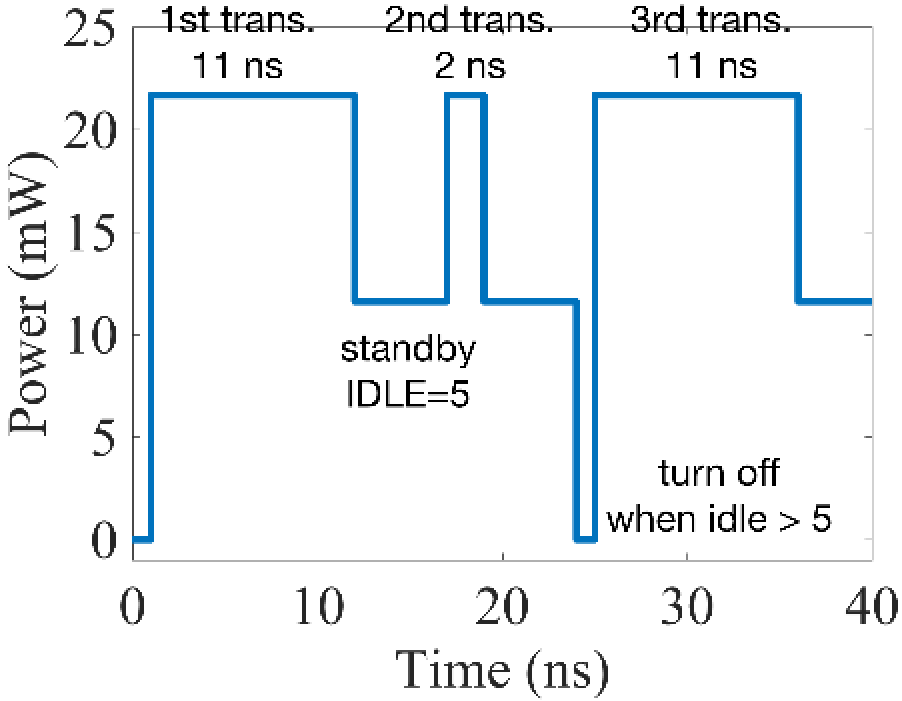

Case Study – An Exemplar Optical Network-on-Chip Architecture

- A Single-Writer-Multiple-Reader (SMWR) architecture
  - N waveguides for N nodes
  - In each waveguide
    - One writer, N-1 readers
  - Point-to-point communication

- Example - 4 nodes SMWR

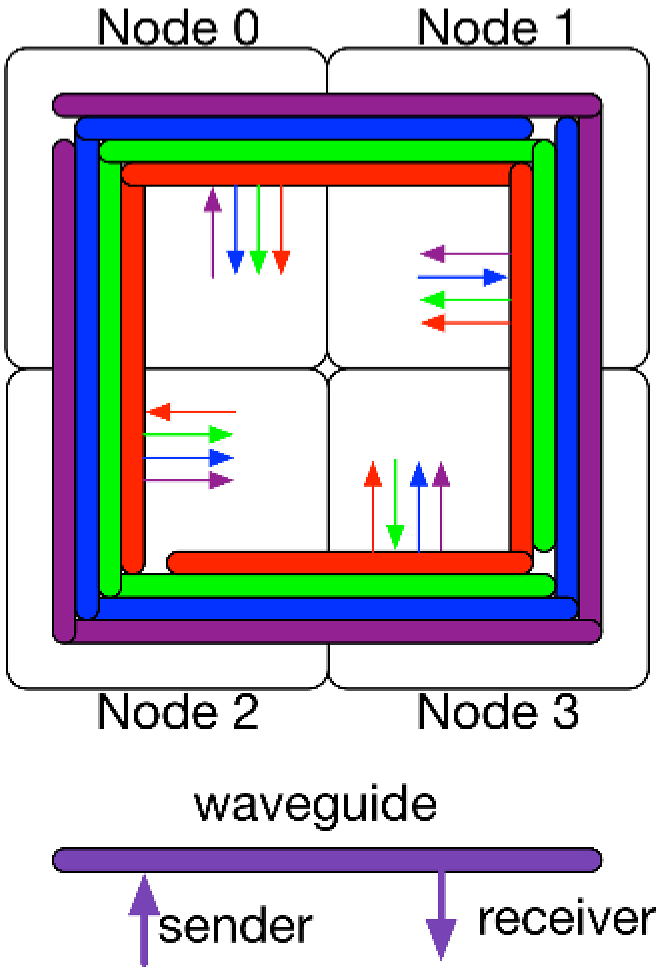

Case Study – Benchmark FFT-1024

- Use the traffic trace in a benchmark suite

- Assume no turn-on delays → limit
  - Min time (performance)
  - Min power (energy)

- IDLE=0, on-off control
- IDLE=+∞, on-standby

- Optimal IDLE=10

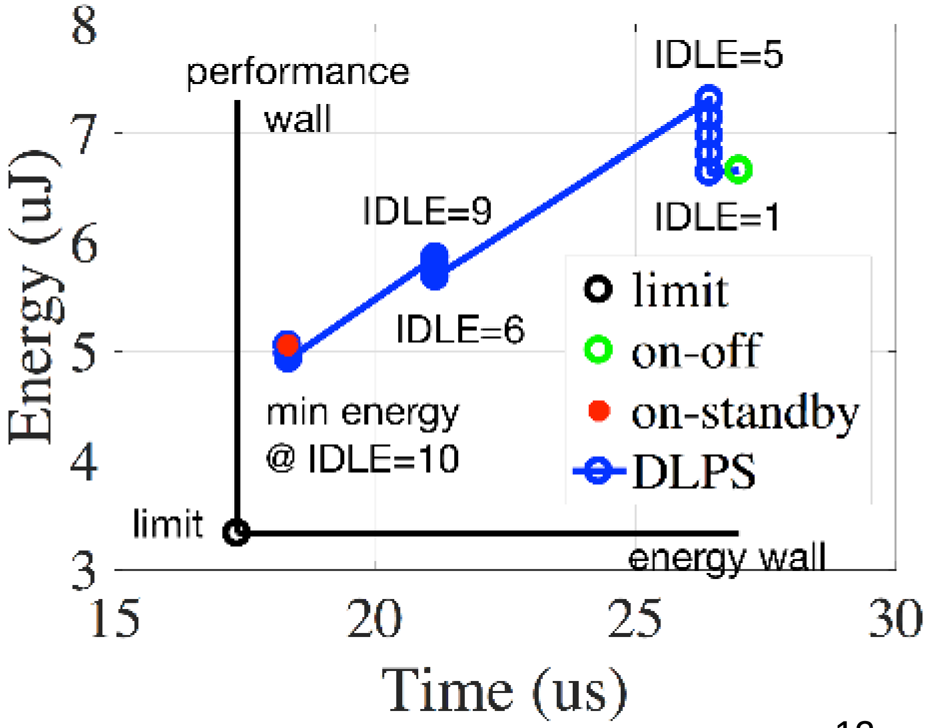

Determining the Optimal *IDLE*

- The optimal IDLE (min energy) should guarantee:
  - Energy in standby mode ≤ Turn-on energy from off mode

$$IDLE_{opt} \cdot P_{standby} + D_{standby} \cdot P_{on} \leq D_{off} \cdot P_{on}$$

  - $IDLE_{opt}$ ≥ the max interval between the transmission sessions

- FFT-1024 case study
  - $0 \leq IDLE_{opt} \leq 16$
  - $IDLE_{opt} \geq 10$
  - Thus: $IDLE_{opt} = 10$

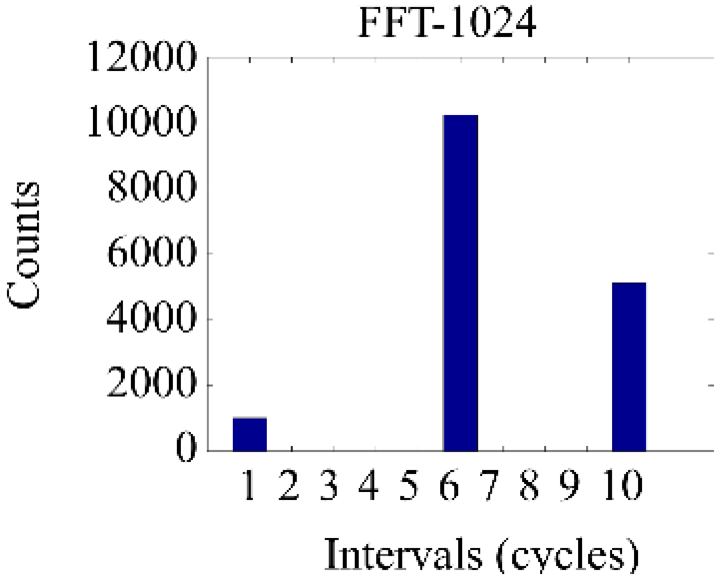

Experiments on Public Benchmark Suite

Simulation Results

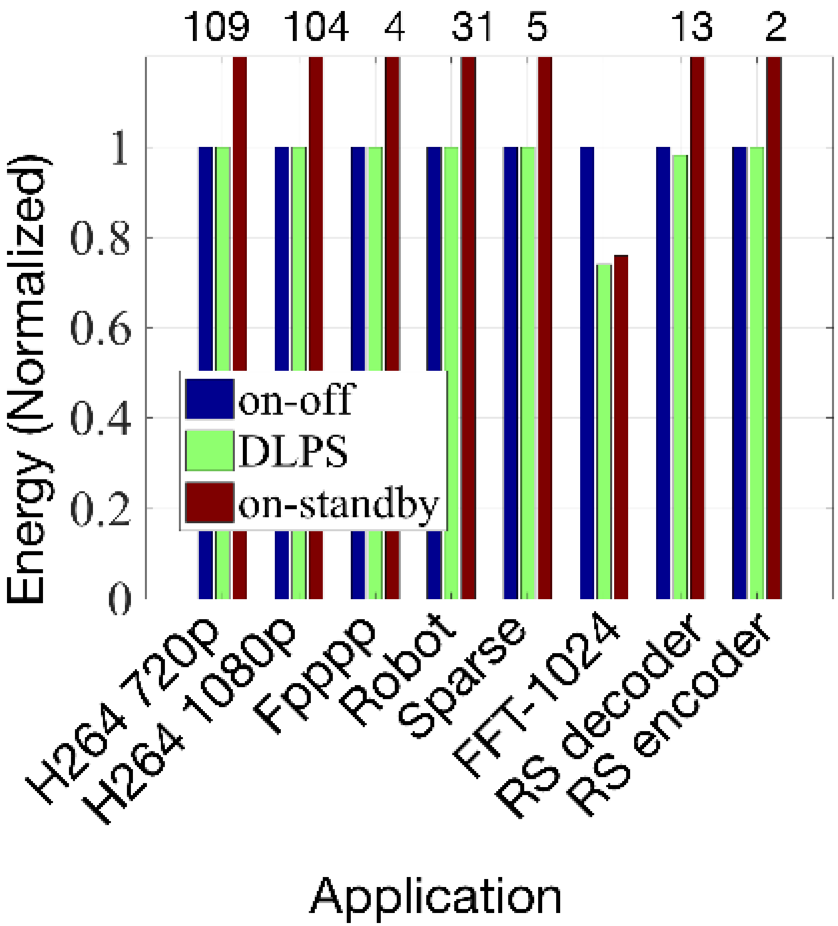

- If $\text{IDLE}_{opt} = 0$
  - DLPS rolls back to on-off
  - Energy efficiency is the same as on-off
- If $\text{IDLE}_{opt} > 0$
  - DLPS can reduce energy consumption, compared to on-off
- DLPS can best adapt to applications with different traffic patterns

Conclusion

- Propose a fine-grained control strategy: dynamic laser power scaling (DLPS)
  - Providing just-enough data rate for communication
  - Introduce two new modes: *standby* mode and *intermediate data rate* mode
  - Switch to standby mode, or even off, when no data transmissions needed
- Compared to coarse-grained on-off control strategy, our DLPS can:
  - Better adapt to applications with different traffic pattern characteristics
  - Achieve lower energy consumption under certain traffic patterns

Thank you!

Q&A