Designing VLSI Interconnects with Monolithically Integrated Silicon-Photonics

Vladimir Stojanović
MIT

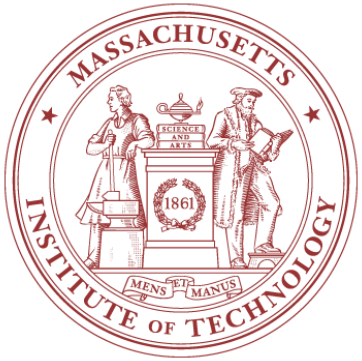

SSCS DL series – Santa Clara, CA, November, 2012

Acknowledgments

- Rajeev Ram, Henry Smith, Hanqing Li (MIT), Milos Popović (Boulder), Krste Asanović (UC Berkeley)
- Jason Orcutt, Jeffrey Shainline, Christopher Batten, Ajay Joshi, Anatoly Khilo

- Karan Mehta, Mark Wade, Erman Timurdogan, Stevan Urosevic, Jie Sun, Cheryl Sorace, Josh Wang
- Michael Georgas, Jonathan Leu, Benjamin Moss, Chen Sun
- Yong-Jin Kwon, Scott Beamer, Yunsup Lee, Andrew Waterman, Miquel Planas

- DARPA, NSF and FCRP IFC
- IBM Trusted Foundry, Solid-State Circuits Society

Chip design is going through a change

- Already have more devices than can use at once
- Limited by power density and bandwidth

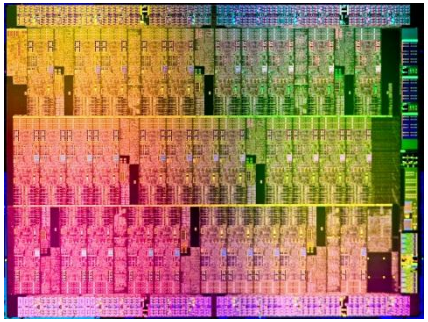

Intel Knights Corner
50 cores, 200 Threads

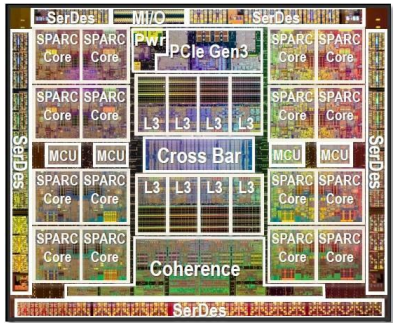

Oracle T5
16 cores, 128 Threads

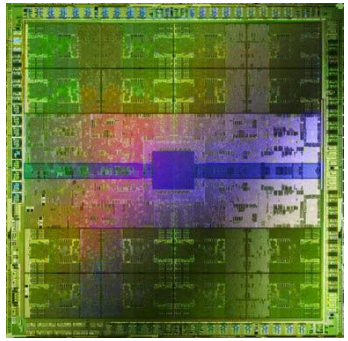

Nvidia Fermi
540 CUDA cores

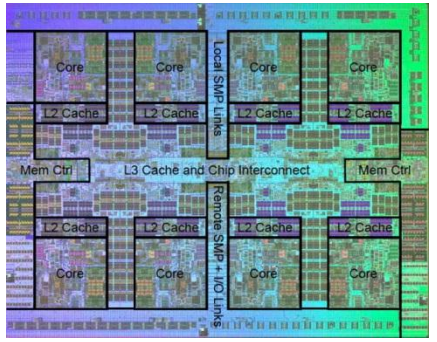

IBM Power 7
8 cores, 32 threads

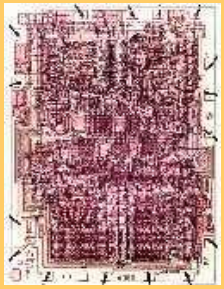

Intel 4004 (1971):
4-bit processor,
2312 transistors,
~100 KIPS,
10 micron PMOS,
11 mm^2 chip

1000s of processor cores and accelerators per die

"The Processor is the new Transistor" [Rowen]

Bandwidth, pin count and power scaling

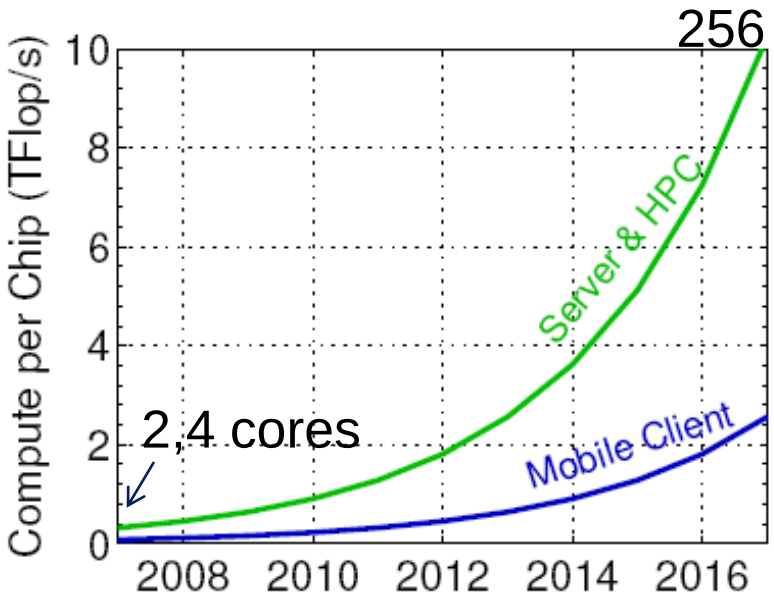

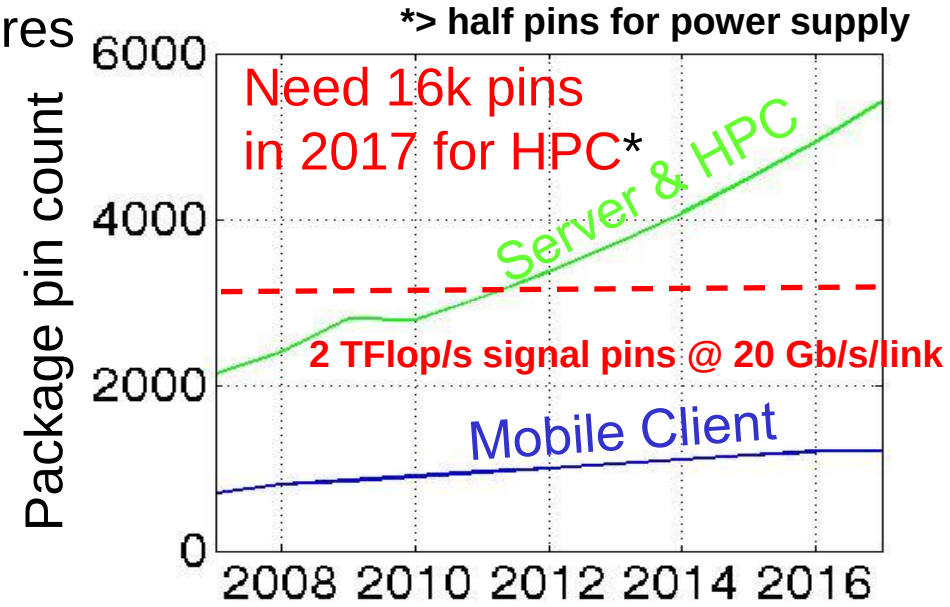

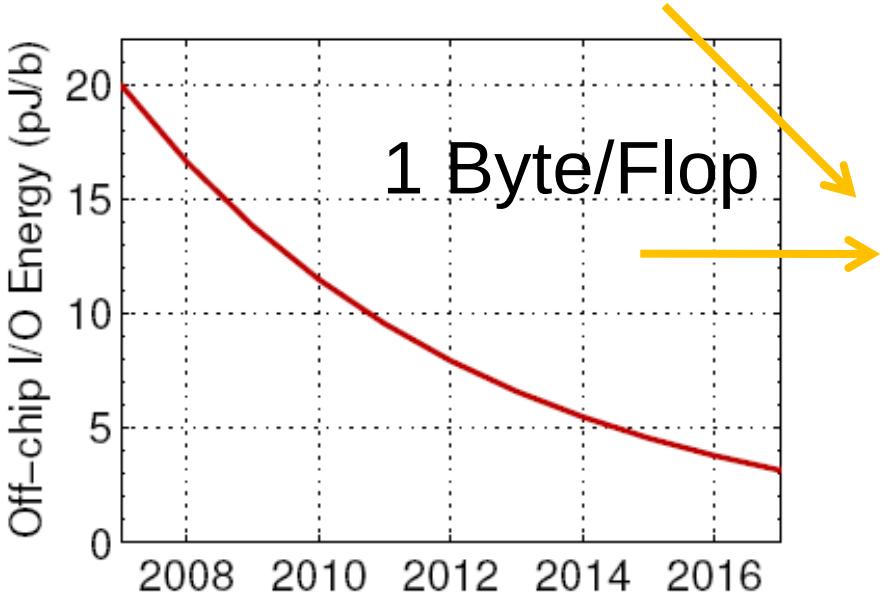

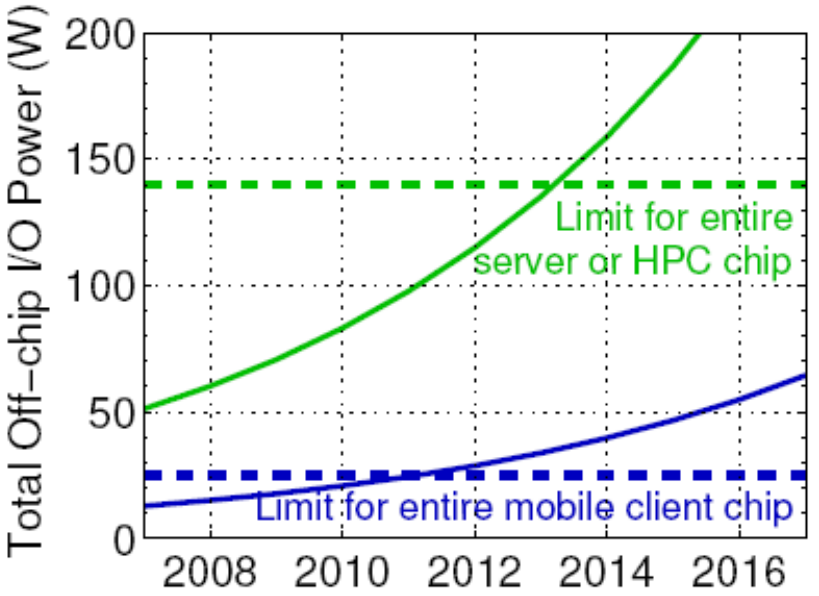

Memory interface scaling problems: Energy-cost and bandwidth density

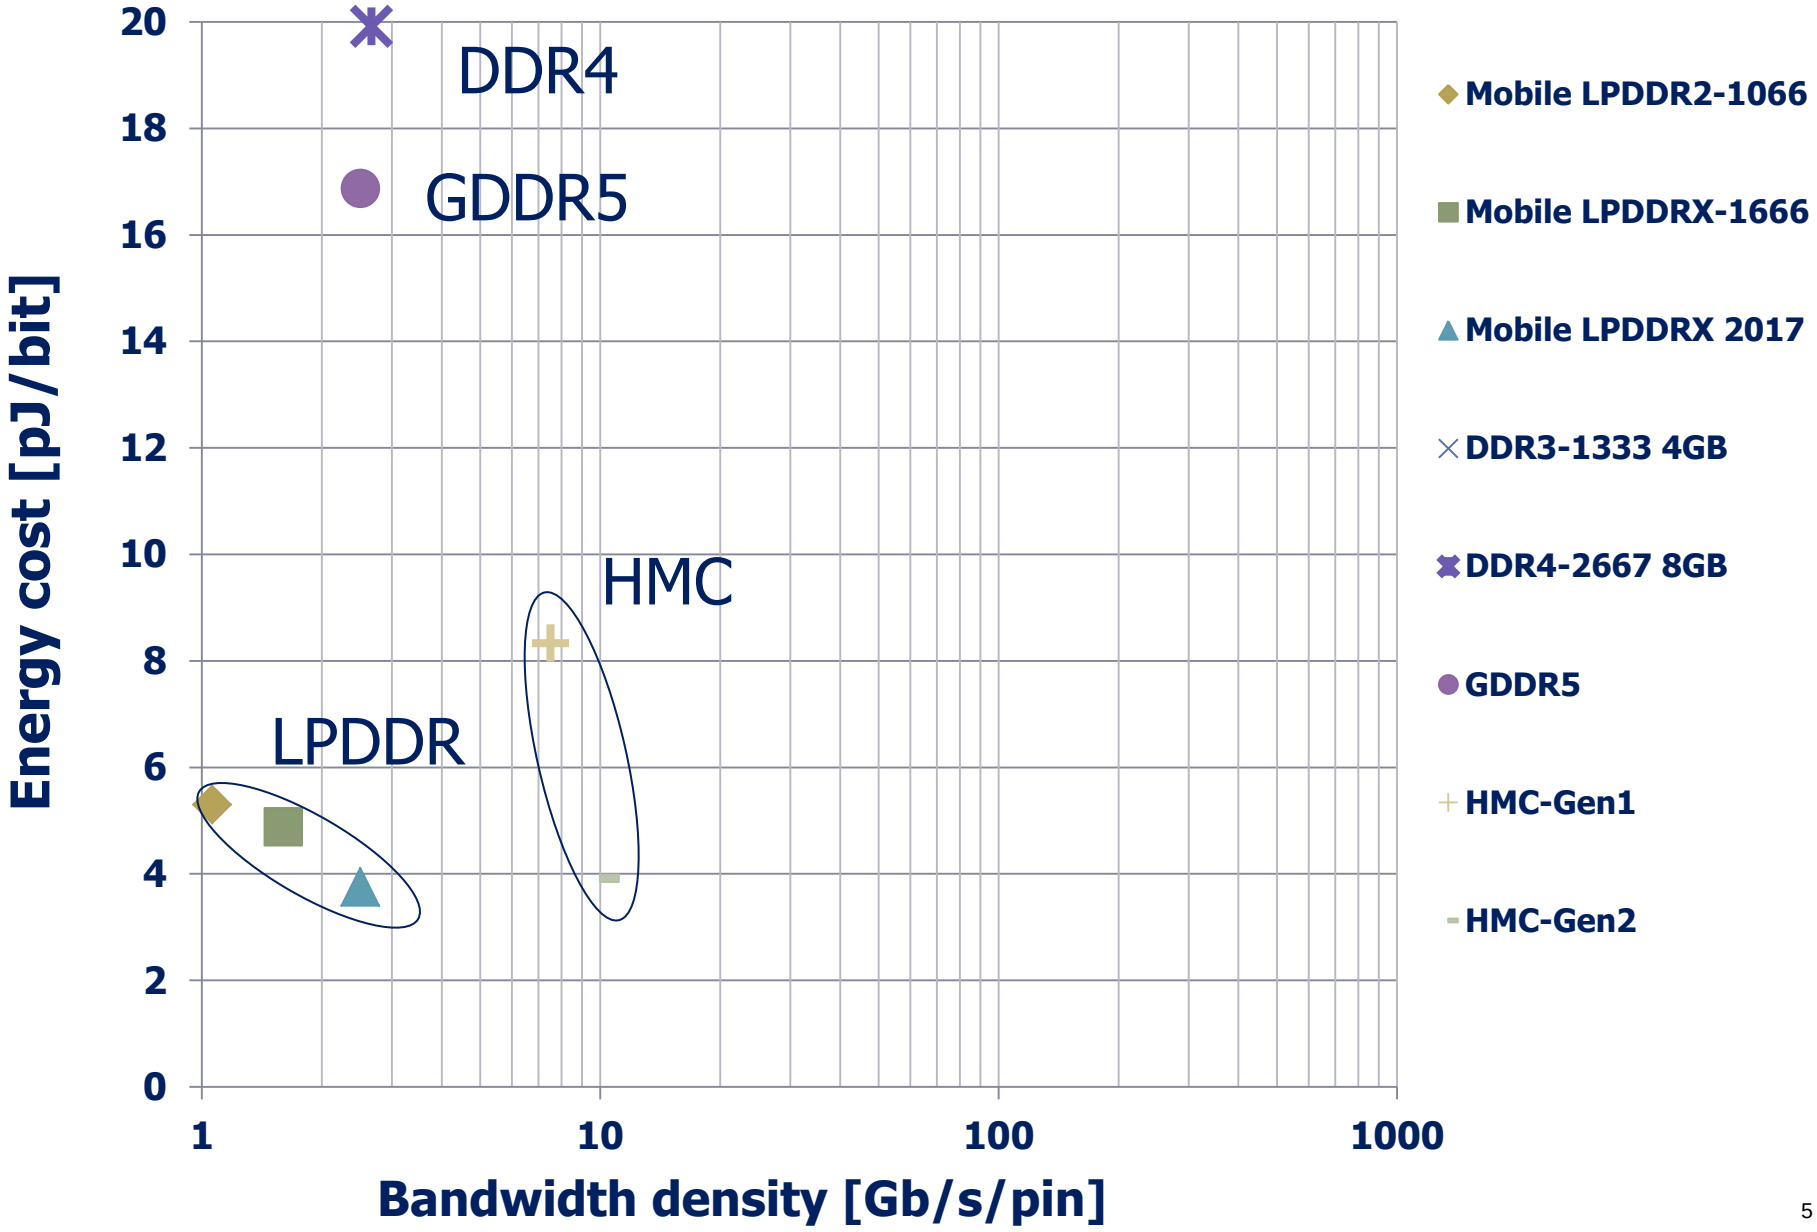

Power and pins required for 10TFlop/s

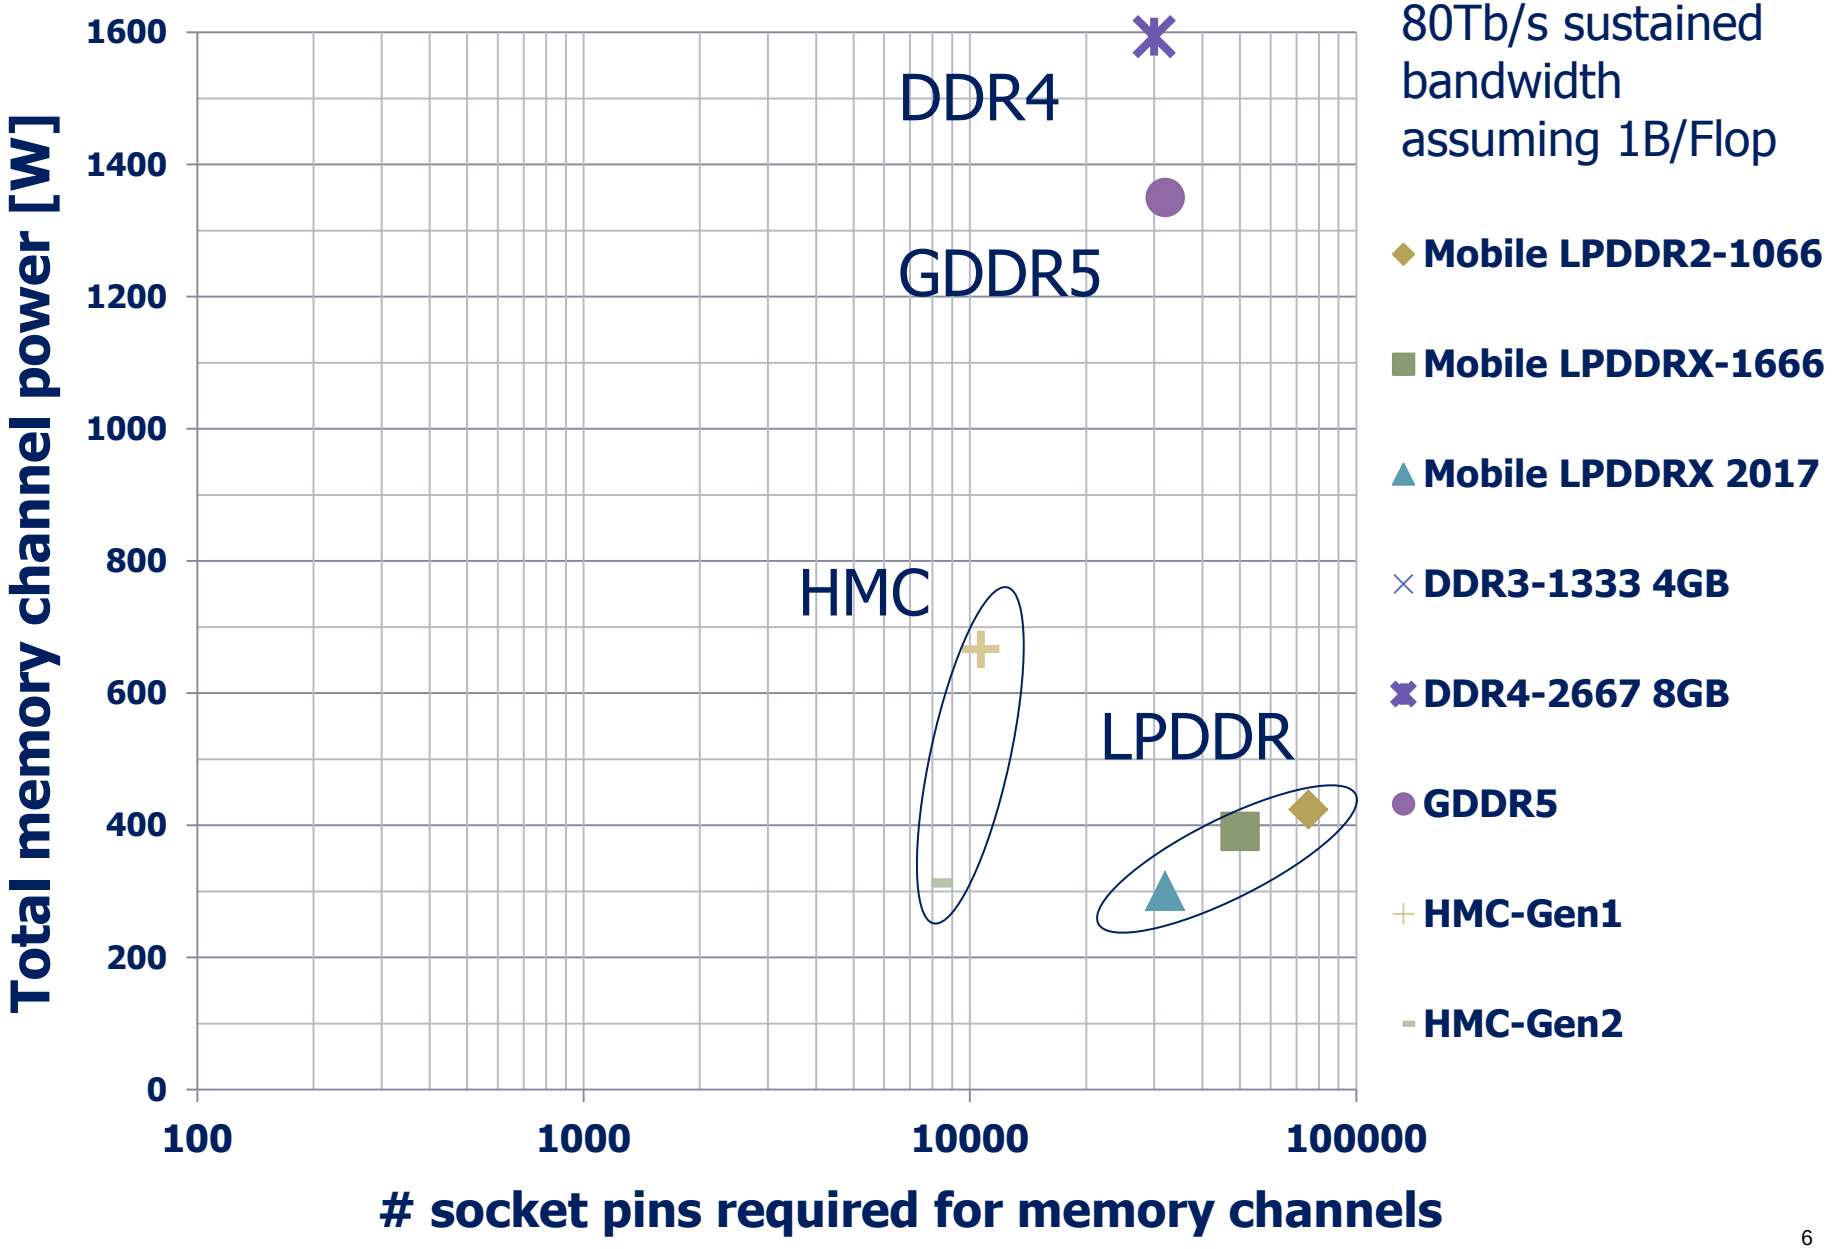

Monolithic Si-Photonics for core-to-core and core-to-DRAM networks

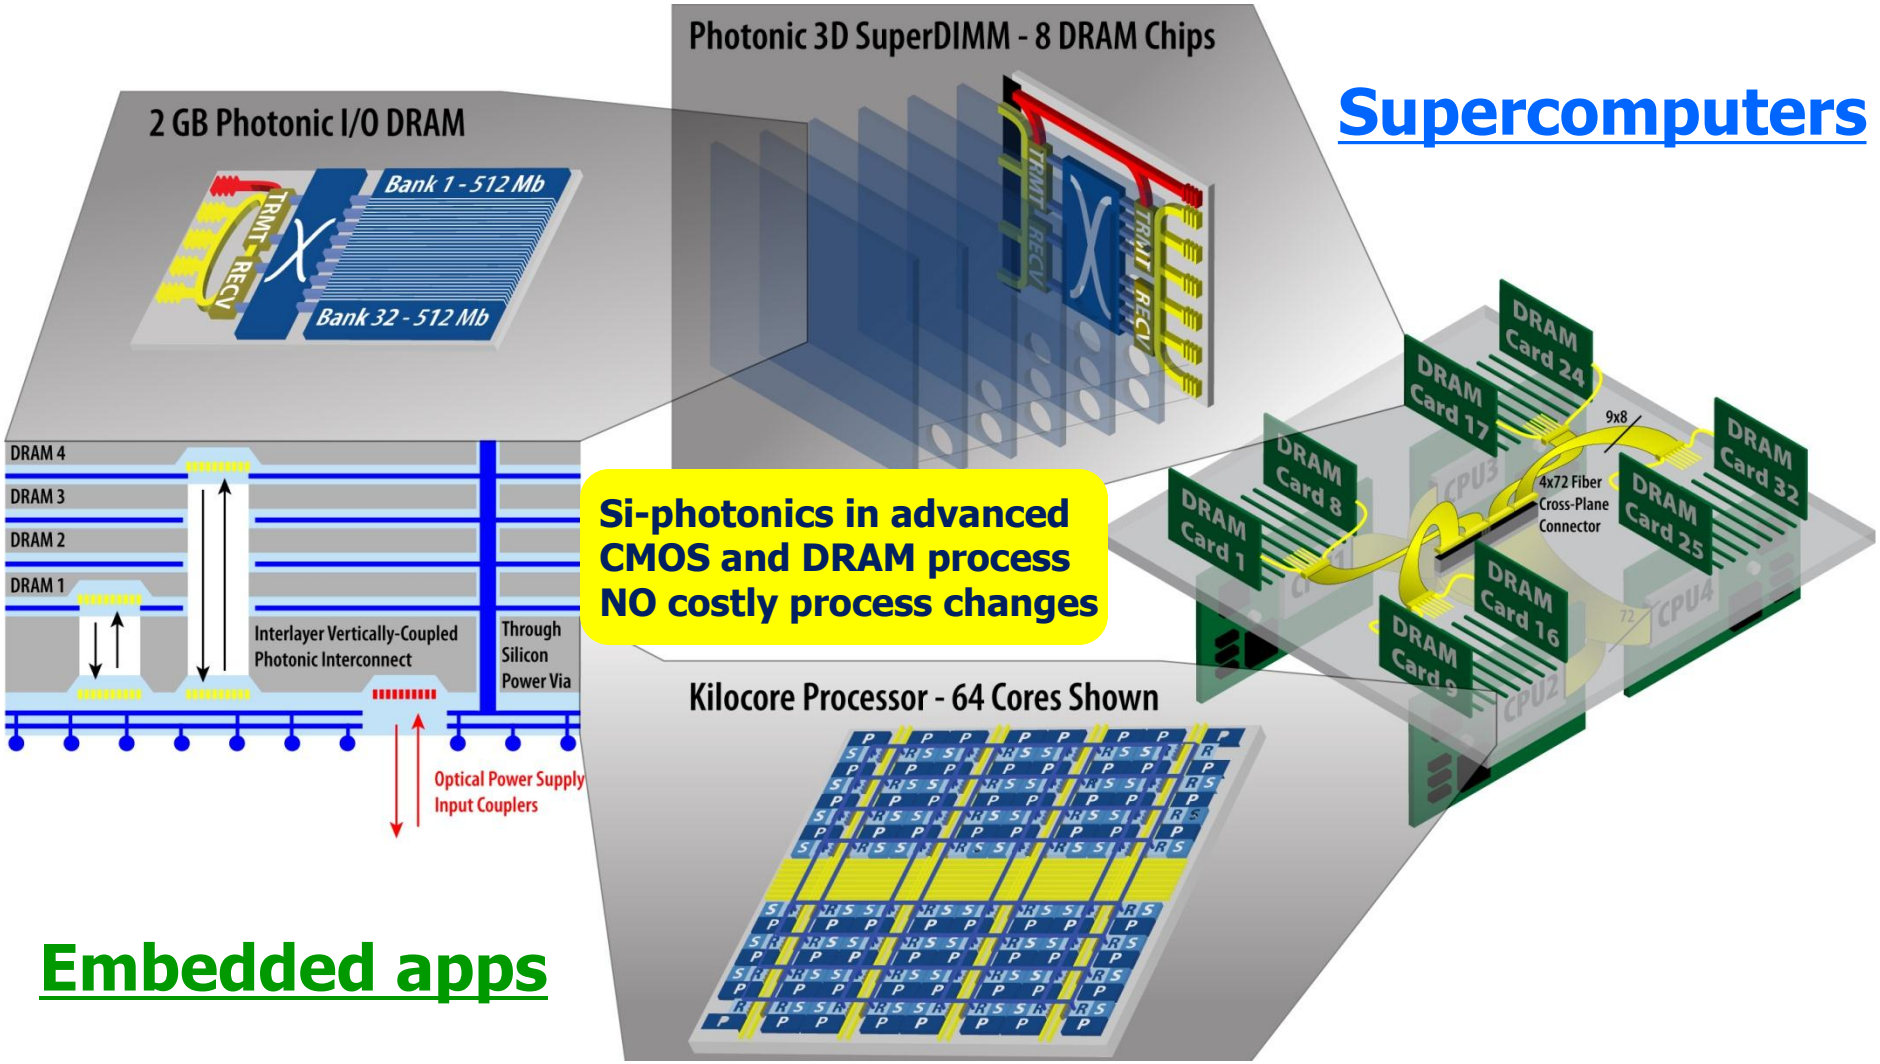

Supercomputers

Embedded apps

Bandwidth density – need dense WDM

Energy-efficiency – need monolithic integration

Integrated photonic interconnects

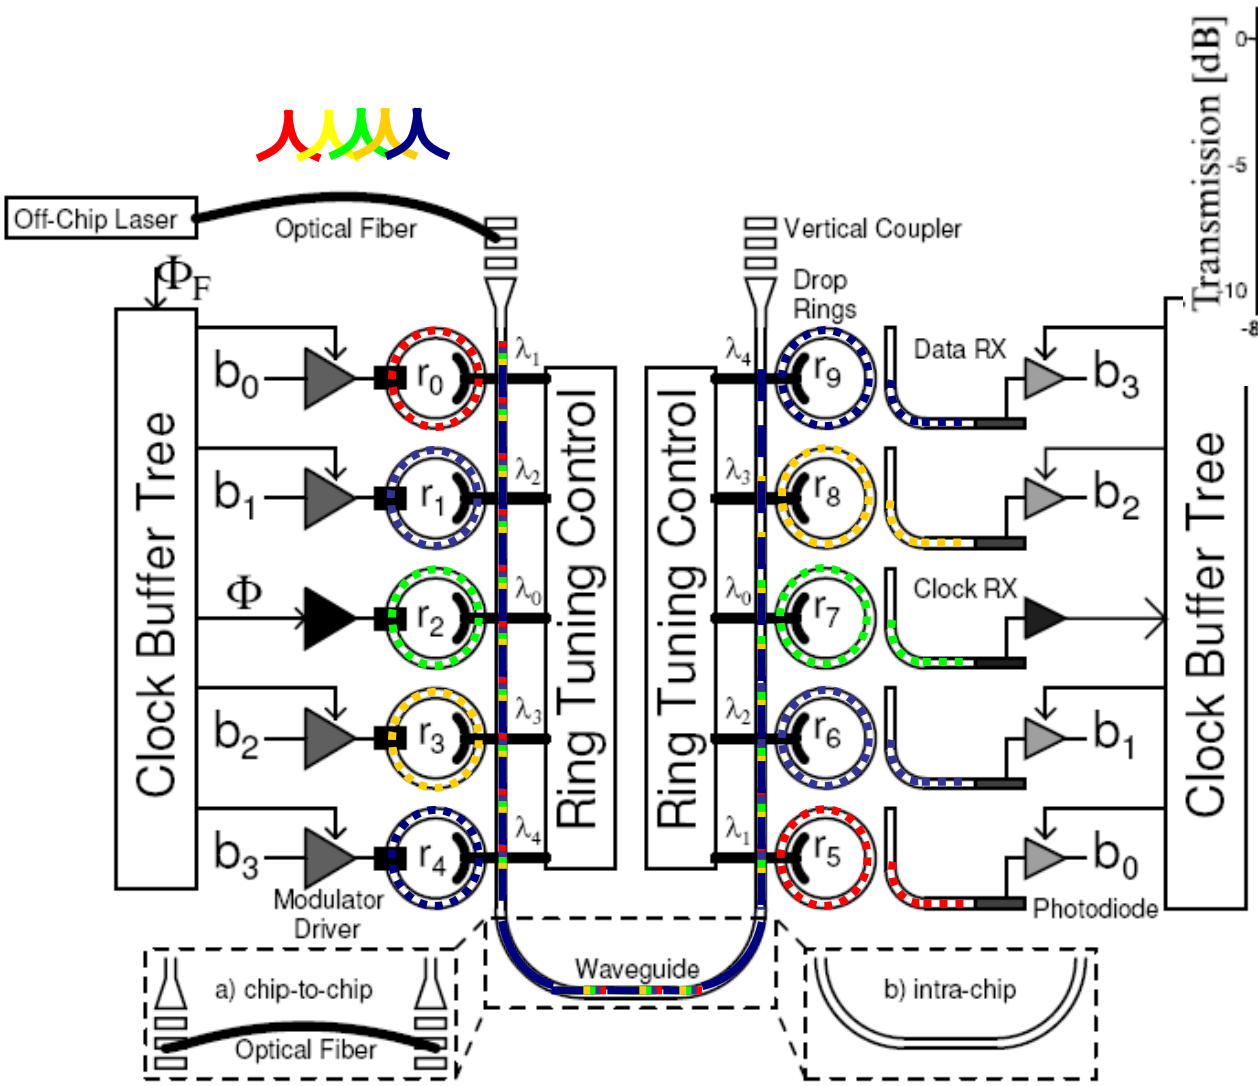

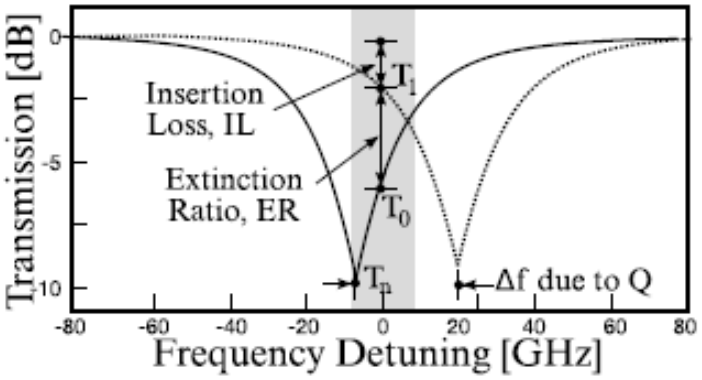

- Each λ carries one bit of data

➔Bandwidth Density achieved through DWDM

➔Energy-efficiency achieved through low-loss optical components and tight integration

Monolithic CMOS photonic integration

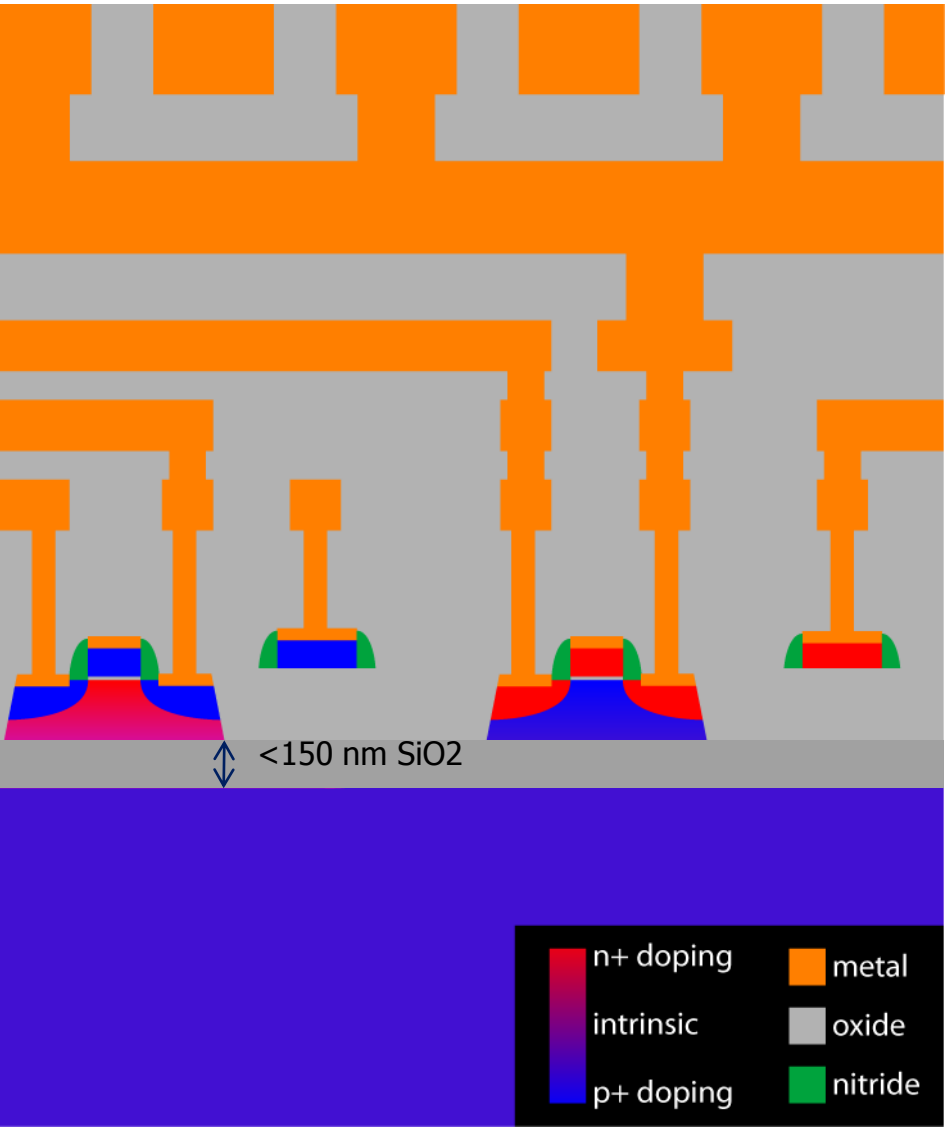

Thin BOX SOI CMOS Electronics

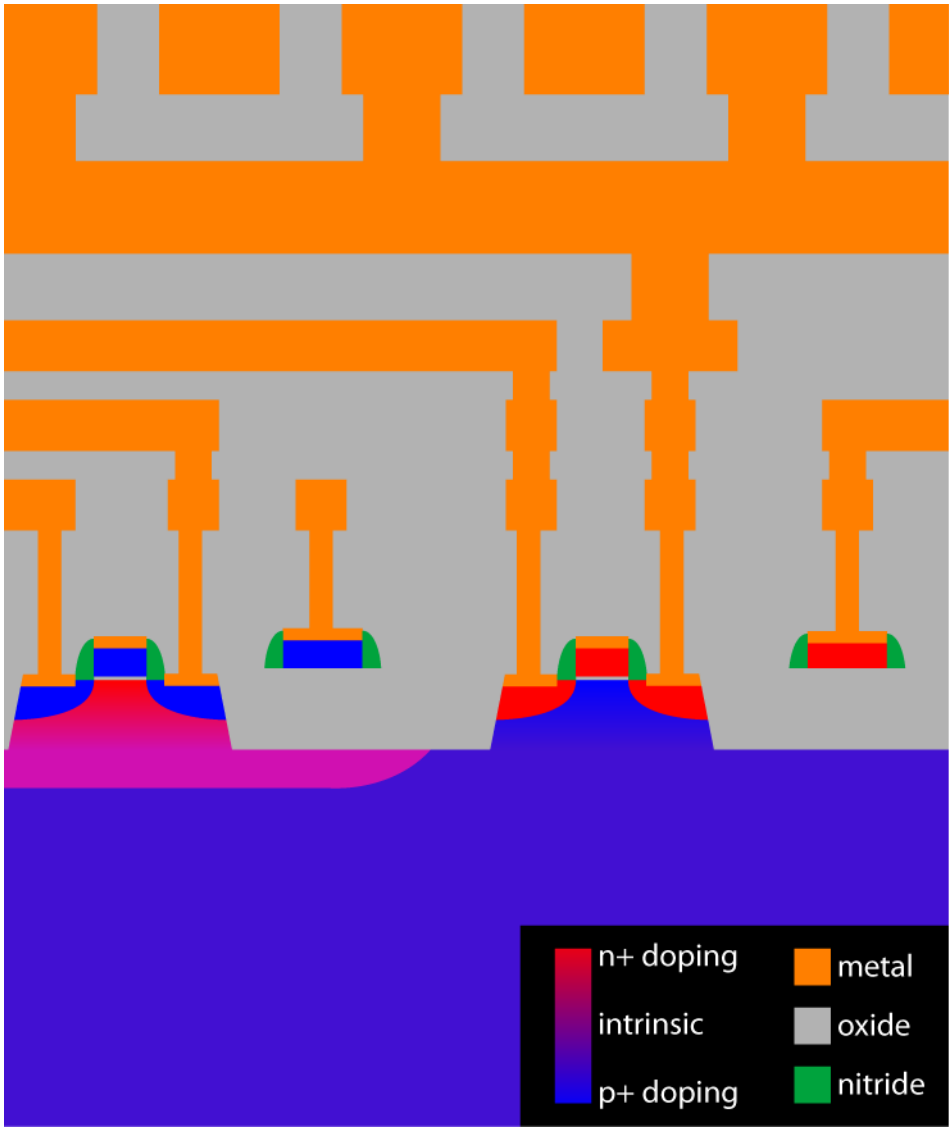

Bulk CMOS Electronics

Si and polySi waveguide formation

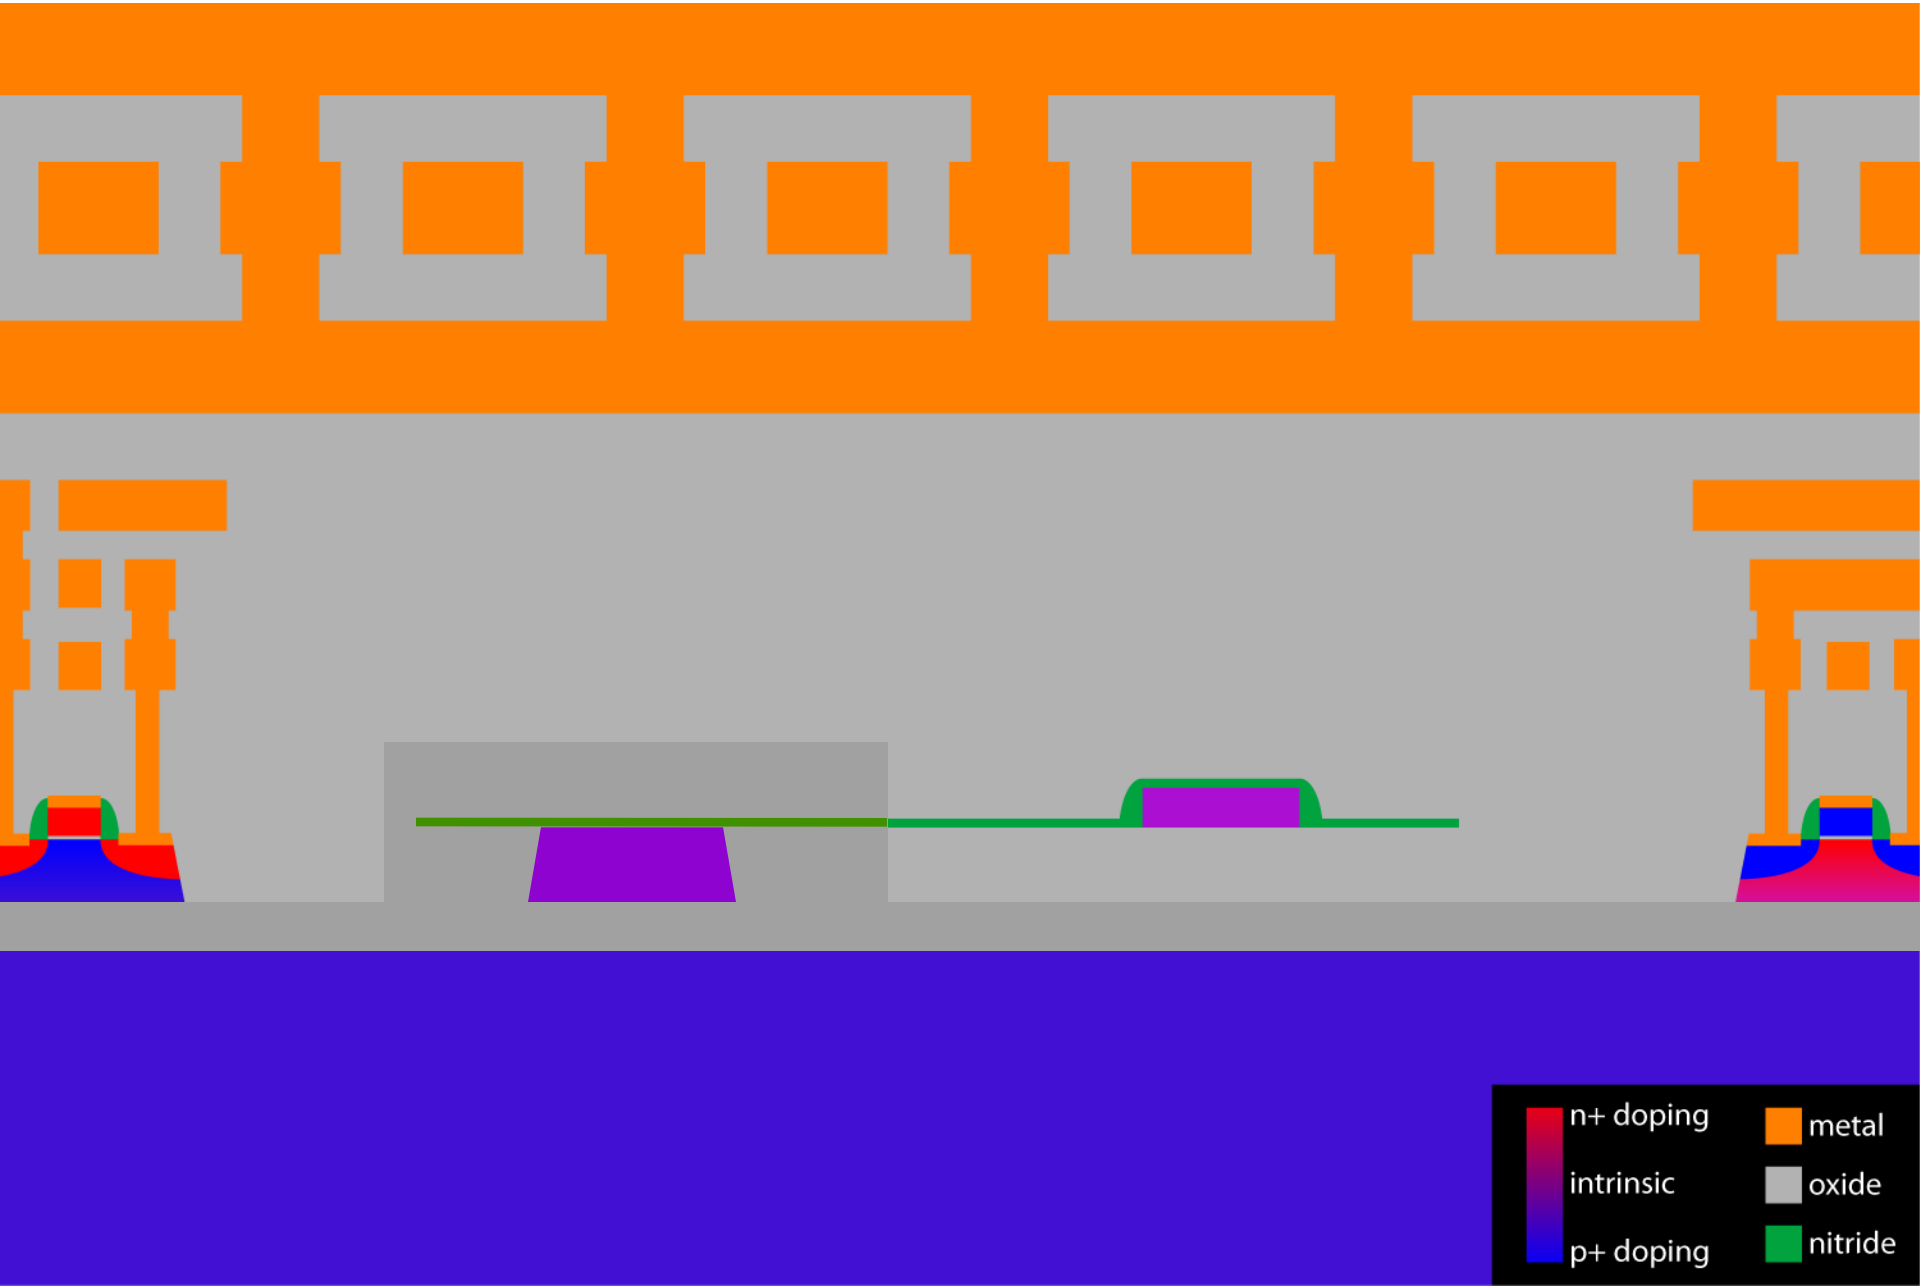

Single channel link tradeoffs

Loss

10-dB

15-dB

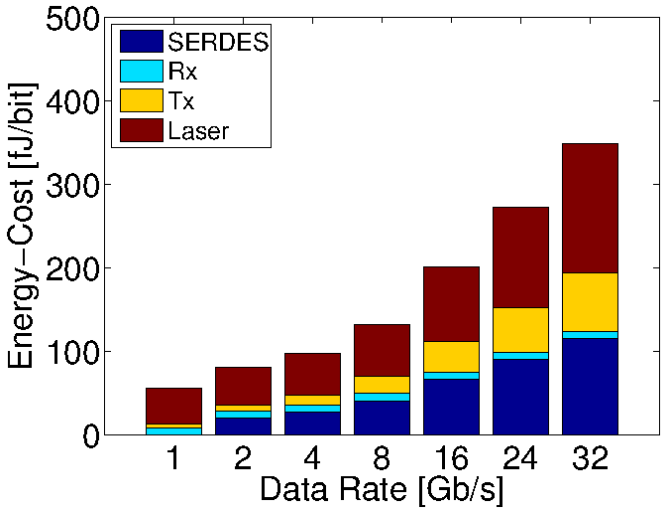

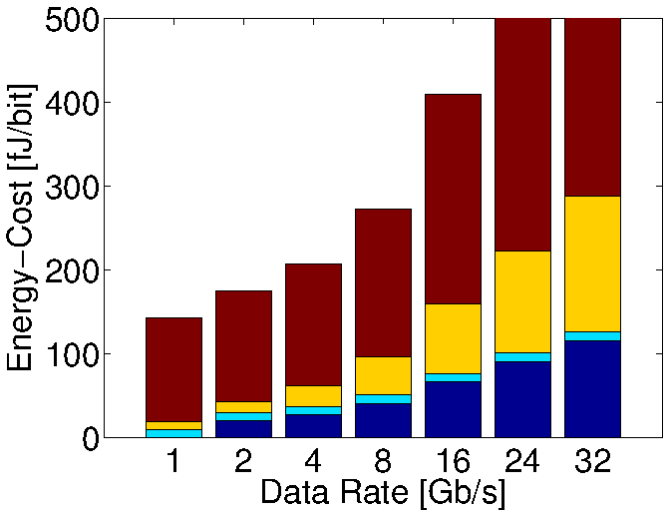

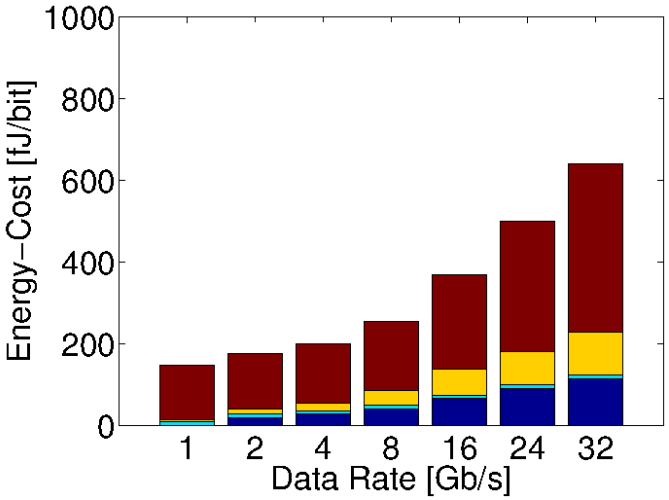

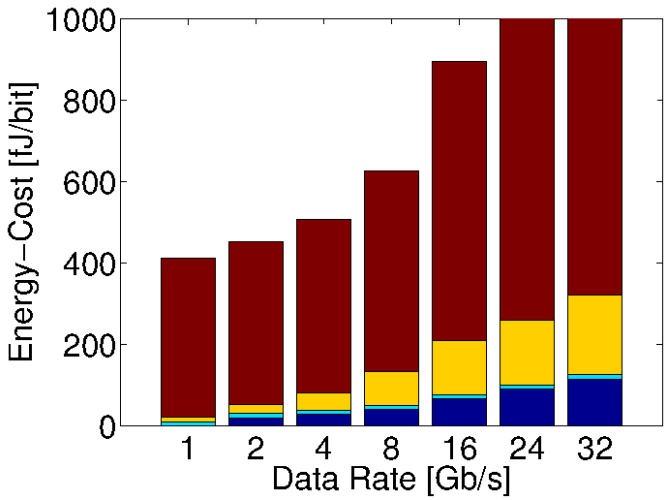

Rx Cap

5-fF

25-fF

Resonance sensitivity

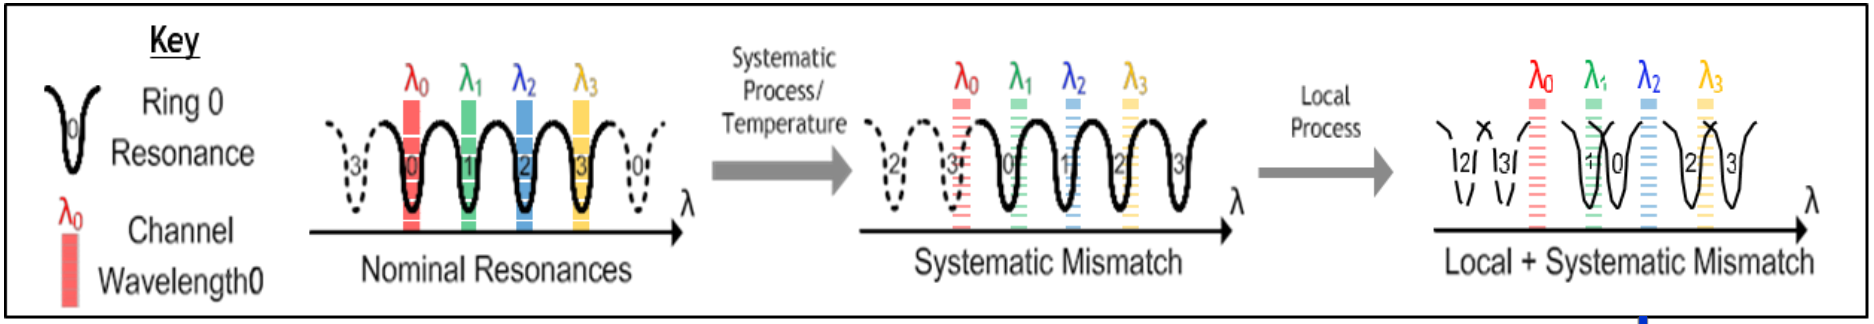

Direct thermal tuning

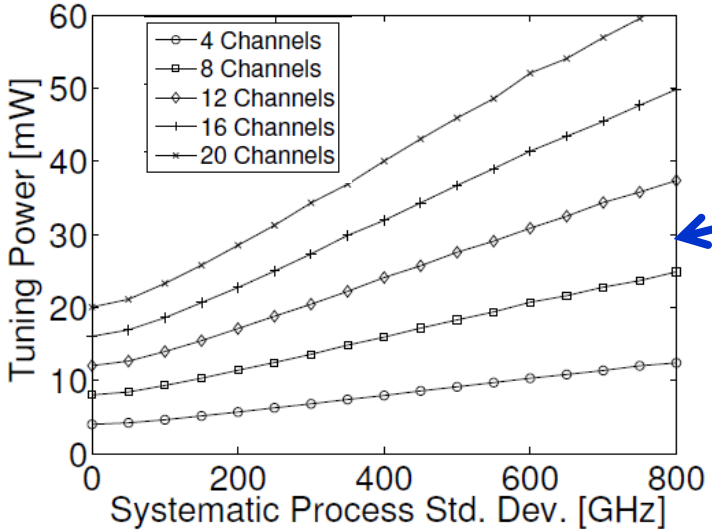

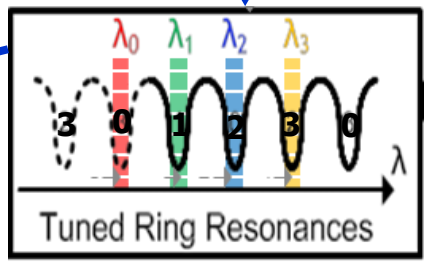

- Process and temperature shift resonances
- Direct thermal tuning cost prohibitive

Georgas CICC 2011, Sun NOCS 2012

Smarter wavelength tuning

Direct thermal tuning

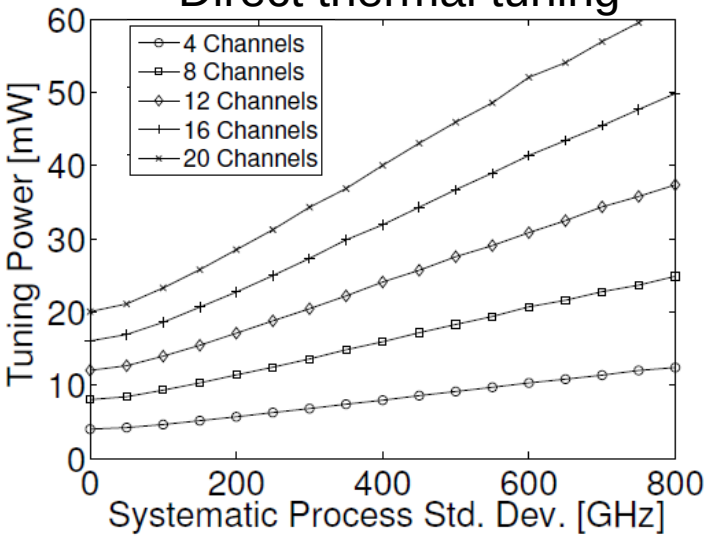

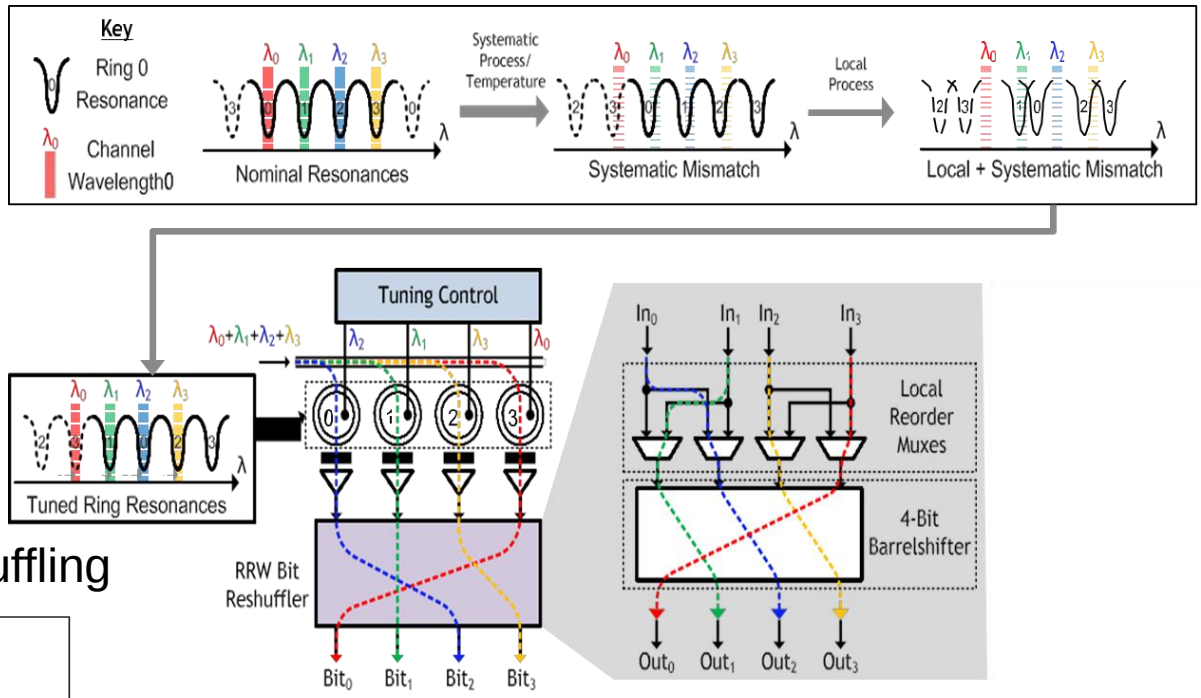

Nearest channel tuning + reshuffling

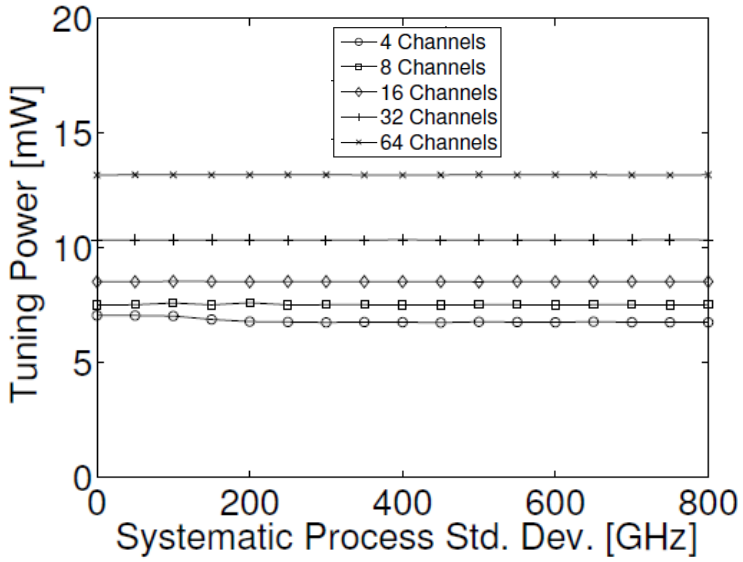

- Electrical backend enables dense WDM
  - Helps reduce tuning costs by more than 10x

Georgas CICC 2011, Sun NOCS 2012

Need to optimize carefully

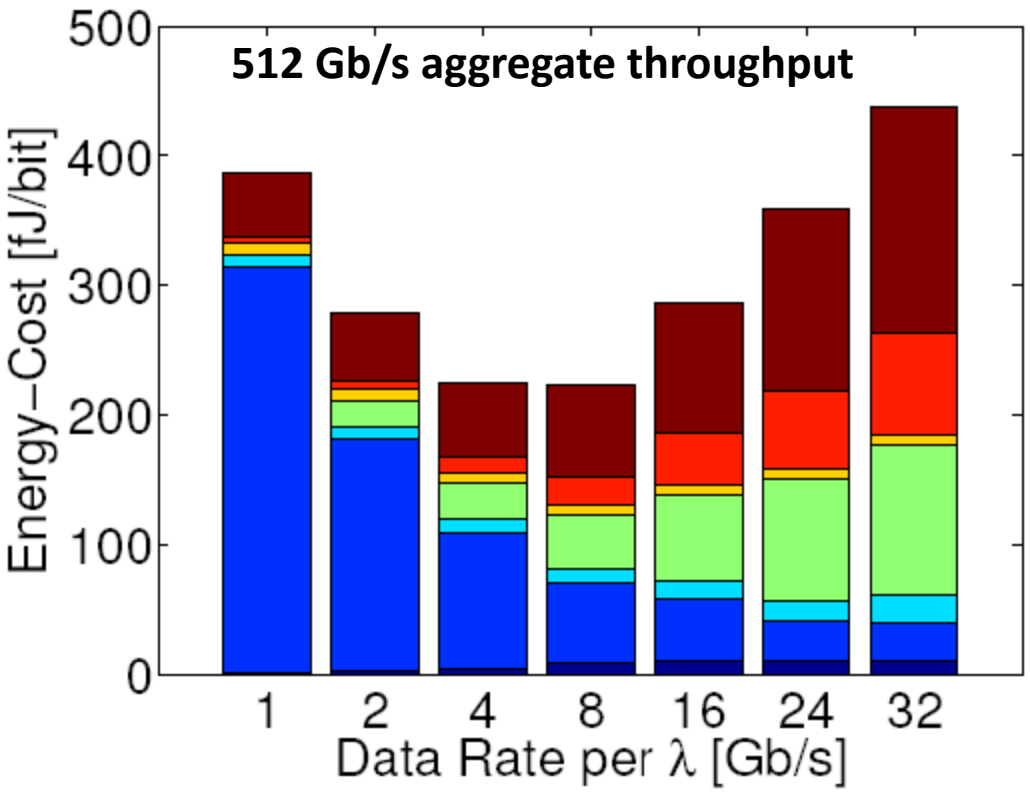

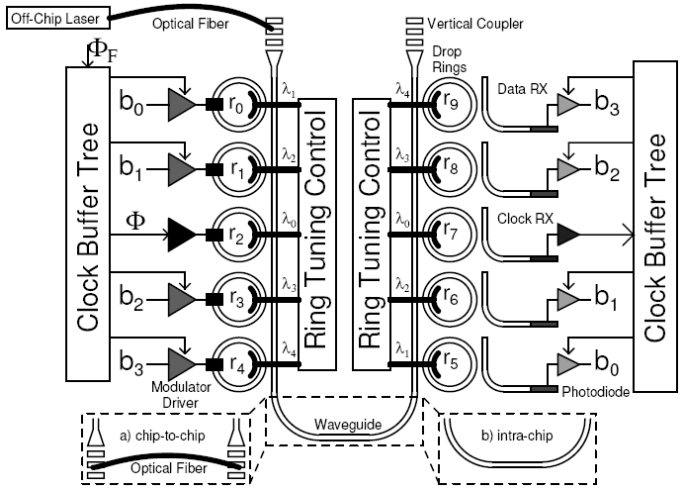

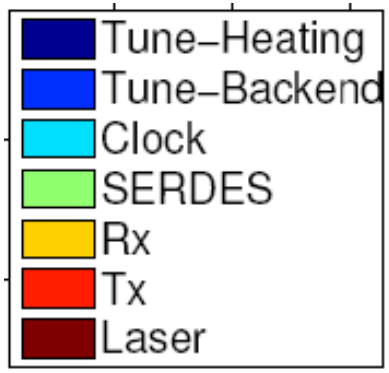

assuming 32nm CMOS

- Laser energy increases with data-rate
  - —Limited Rx sensitivity
  - —Modulation more expensive -> lower extinction ratio
- Tuning costs decrease with data-rate
- **Moderate data rates most energy-efficient**

Georgas CICC 2011

DWDM link efficiency optimization

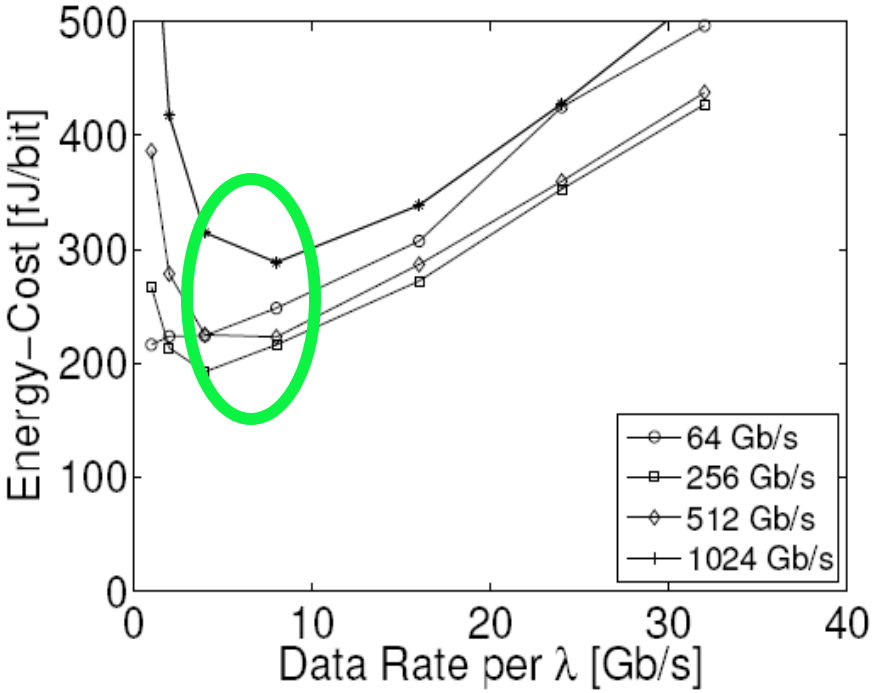

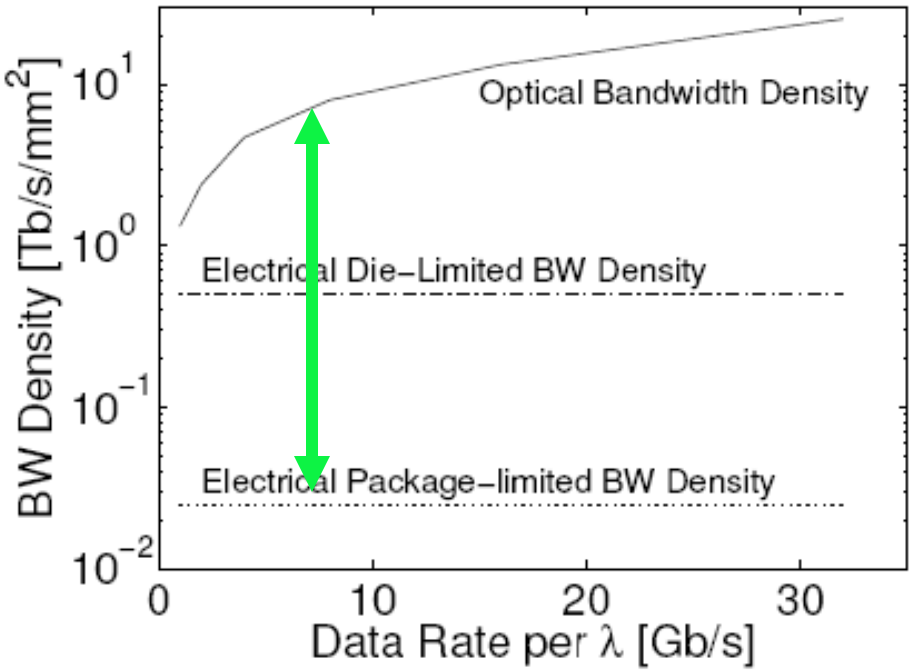

Optimize for min energy-cost

Bandwidth density dominated by circuit and photonics area (not coupler pitch)

- 10x better than electrical bump limited
- 200x better than electrical package pin limit

Photonic memory interface – leveraging optical bandwidth density

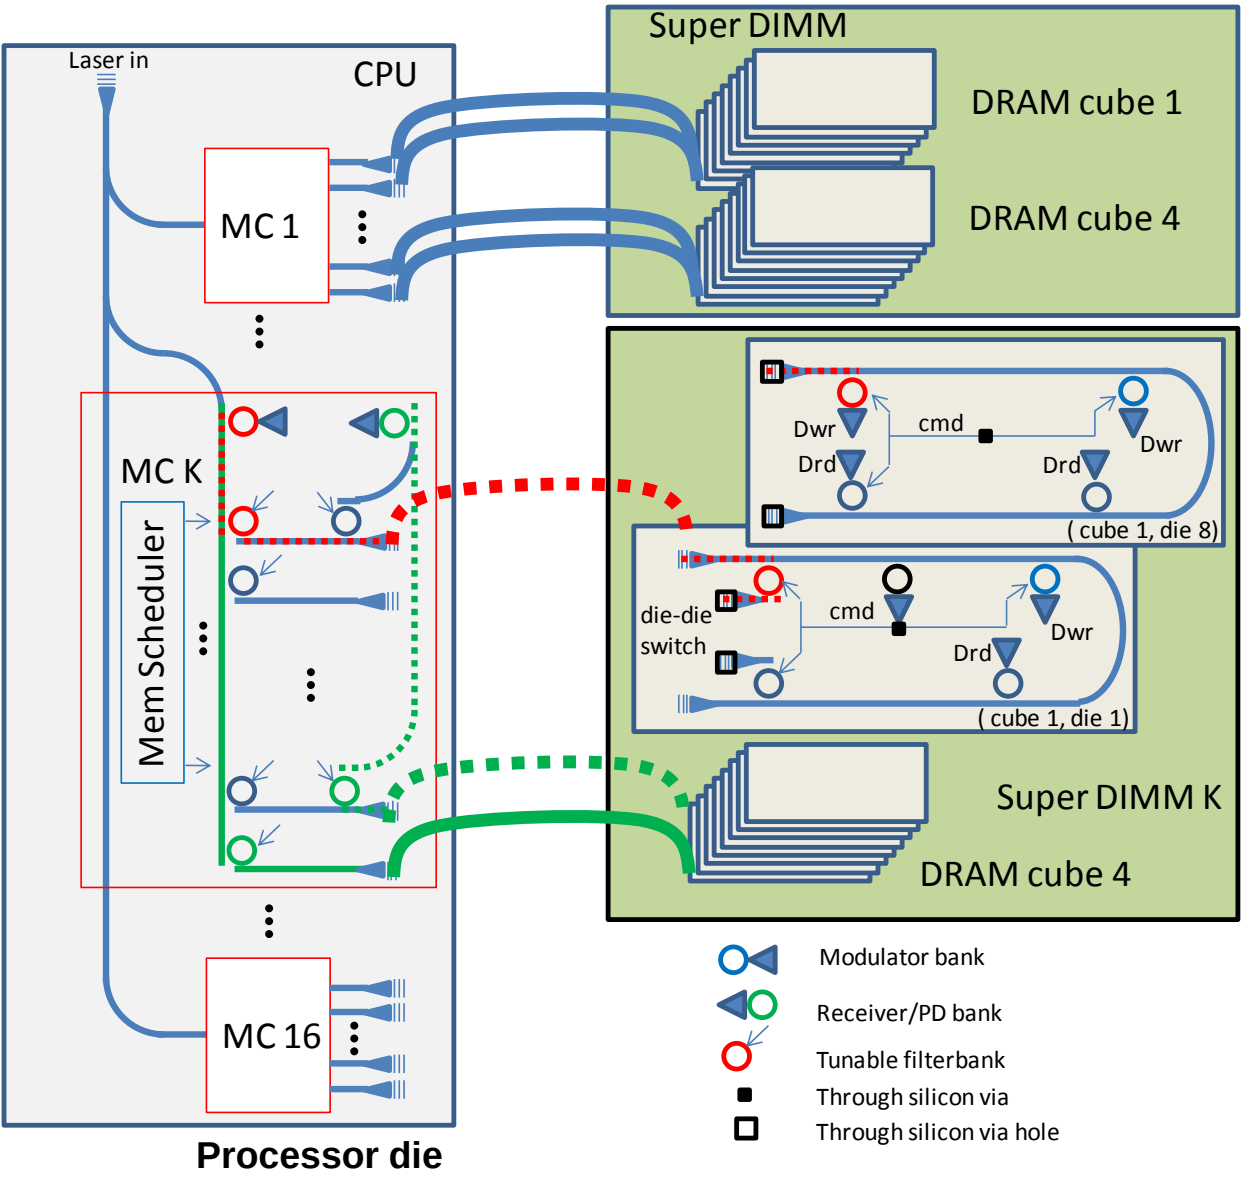

Important Concepts

- Power/message switching (only to active DRAM chip in DRAM cube/super DIMM)

- Vertical die-to-die coupling (minimizes cabling - 8 dies per DRAM cube)

-Command distributed electrically (broadcast)
- Data photonic (single writer multiple readers)

Enables energy-efficient throughput and capacity scaling per memory channel

Beamer ISCA 2010

Laser Power Guiding Effectiveness

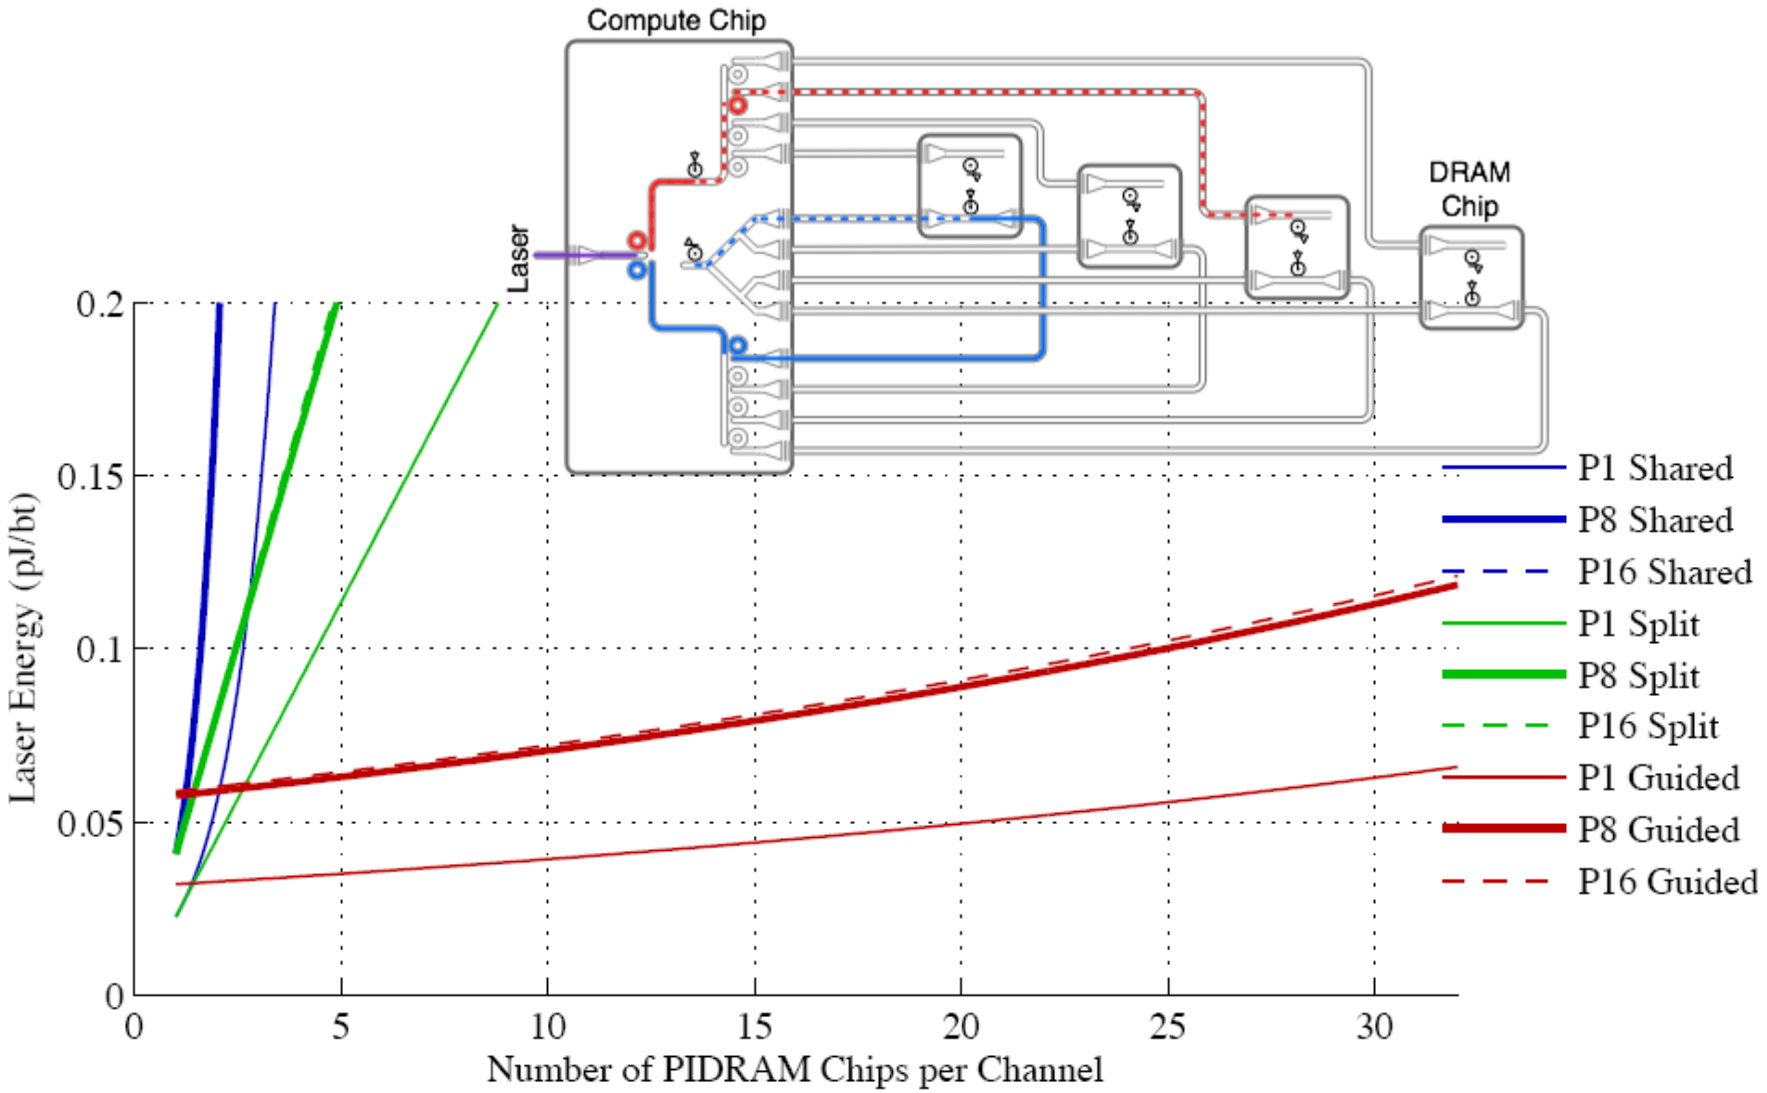

Enables capacity scaling per channel and significant savings in laser energy

Beamer ISCA 2010

Optimizing DRAM with photonics

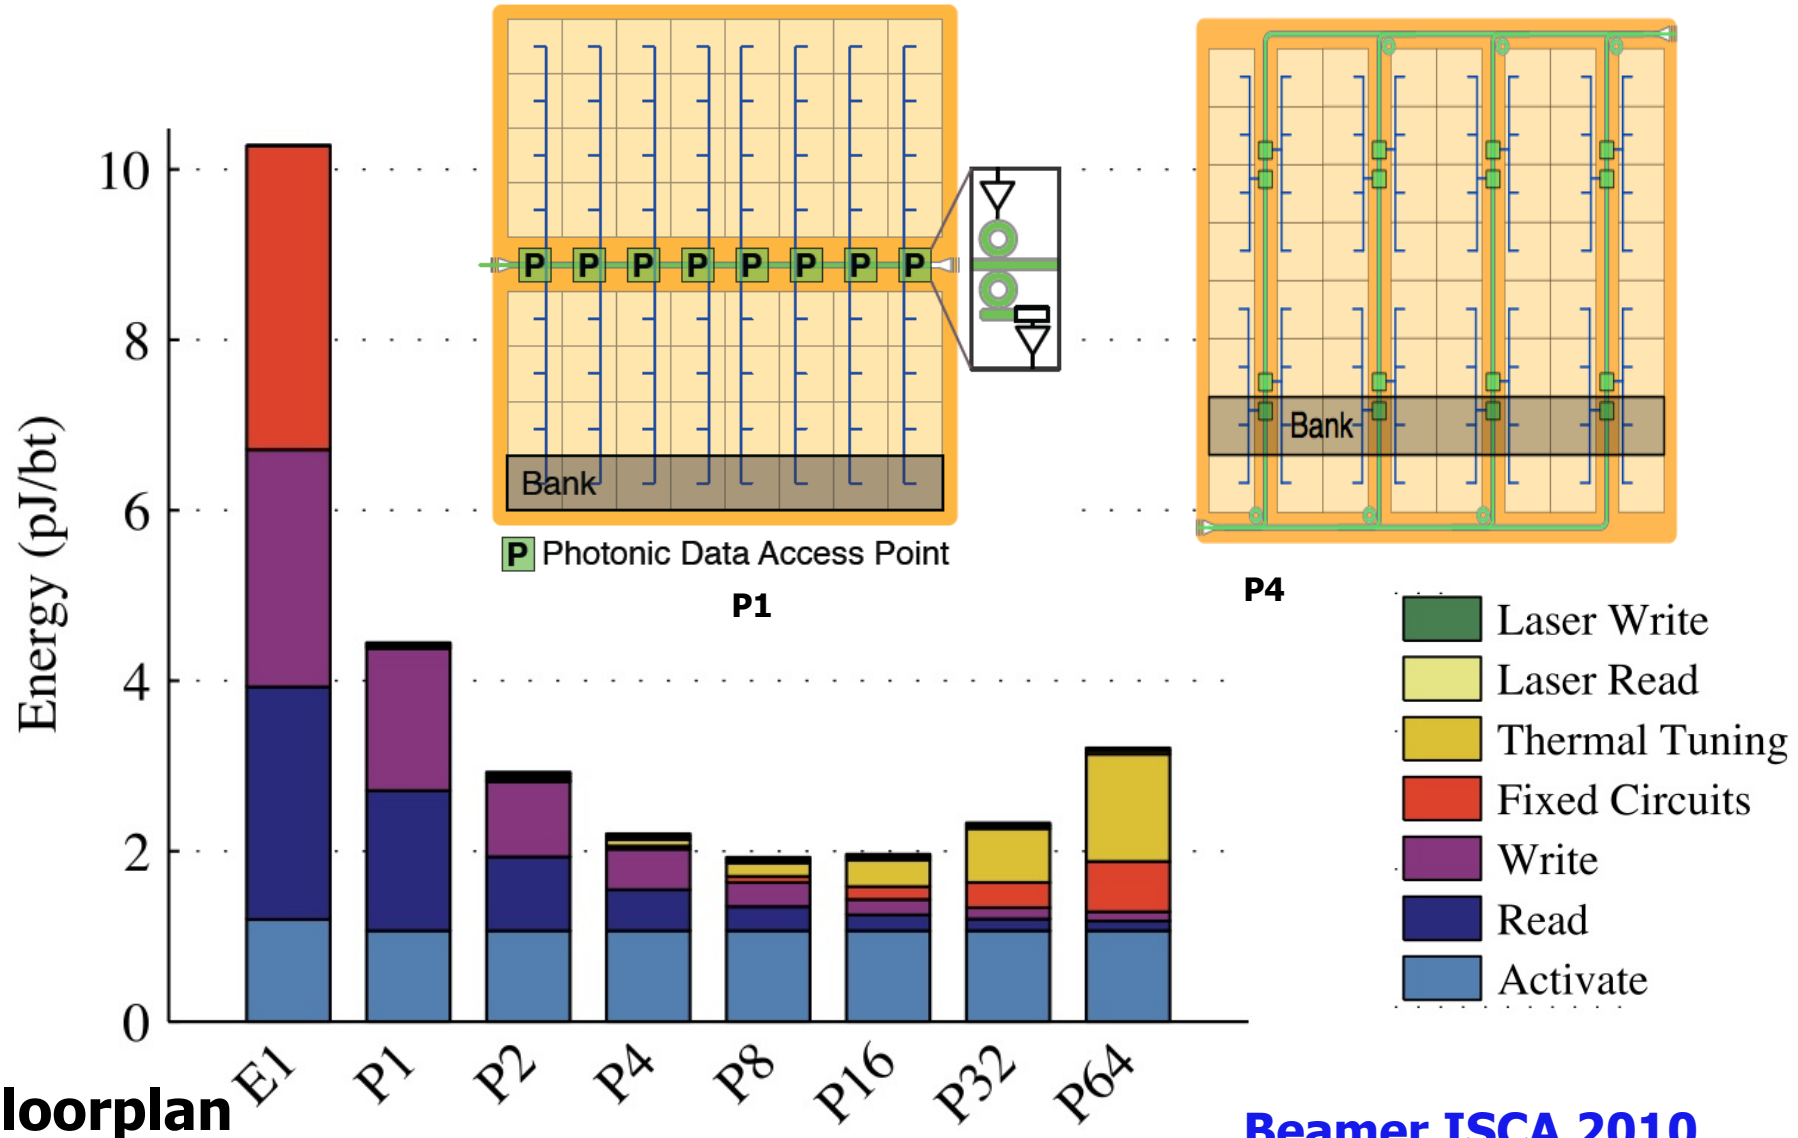

Floorplan

Beamer ISCA 2010

Design Space Exploration of Networks Tool

DSENT – A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks Modeling

Chen NOCS 2012

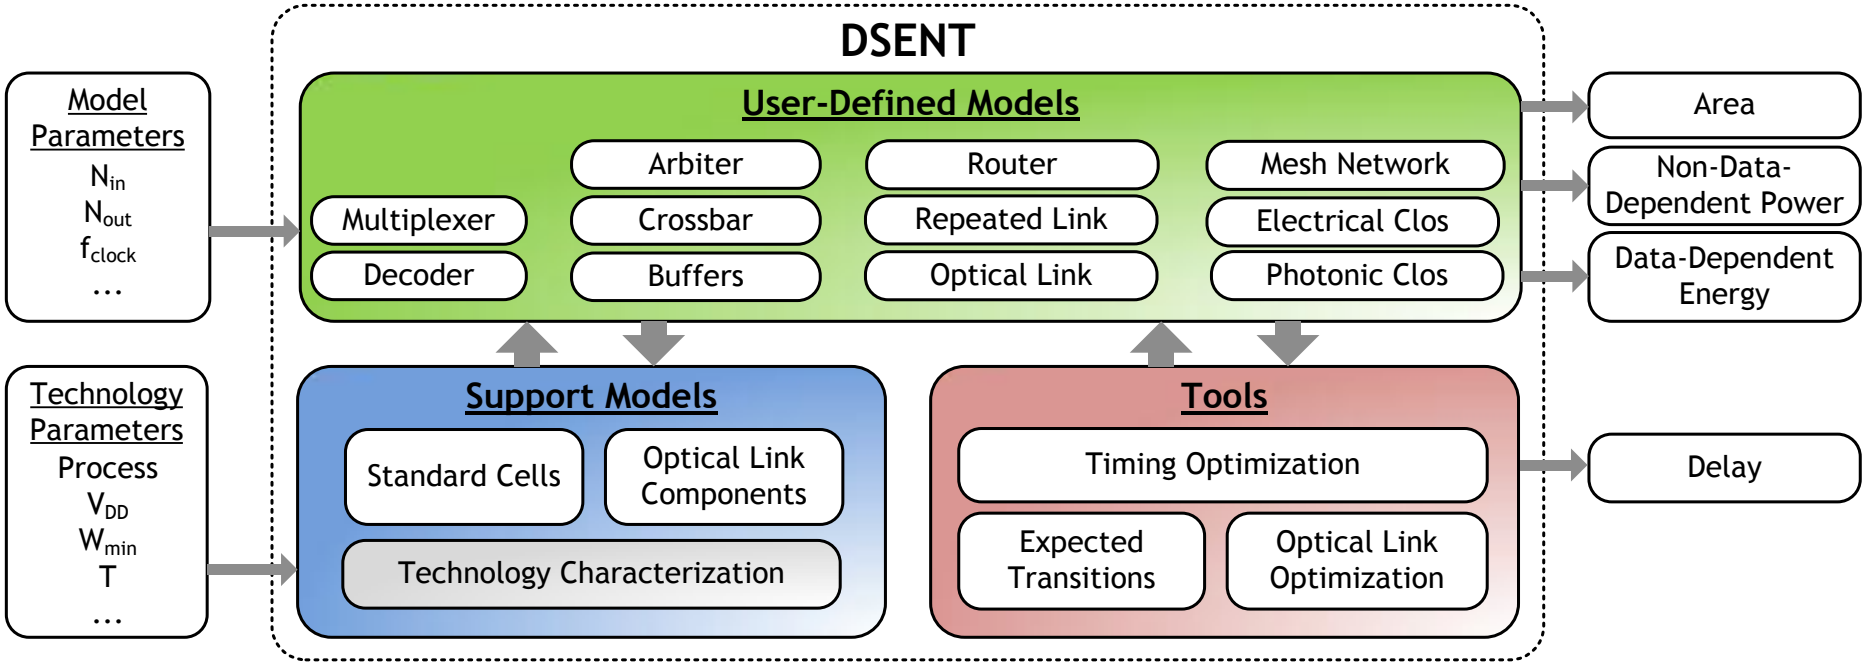

Available for download at:

https://sites.google.com/site/mitdsent/

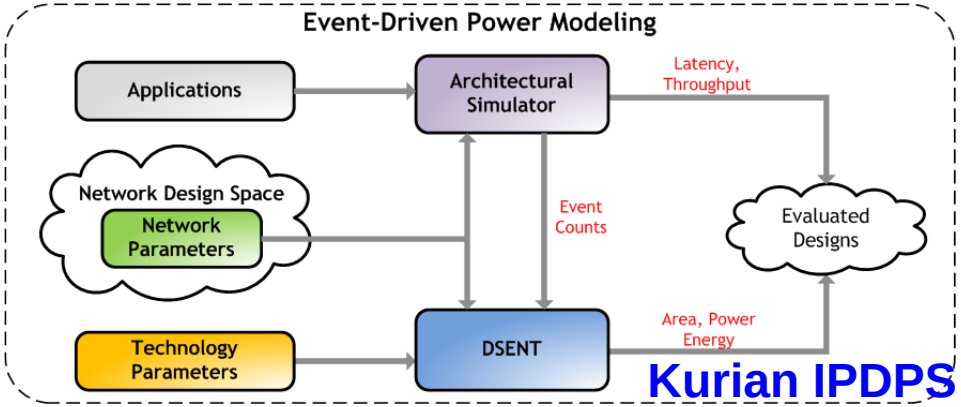

Kurian IPDPS 2012

Significant integration activity, but hybrid and older processes …

130nm
thick BOX SOI

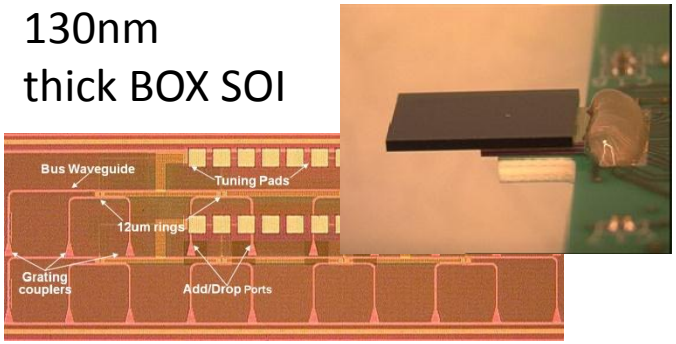

[Luxtera/Oracle/Kotura]

130nm
thick BOX SOI

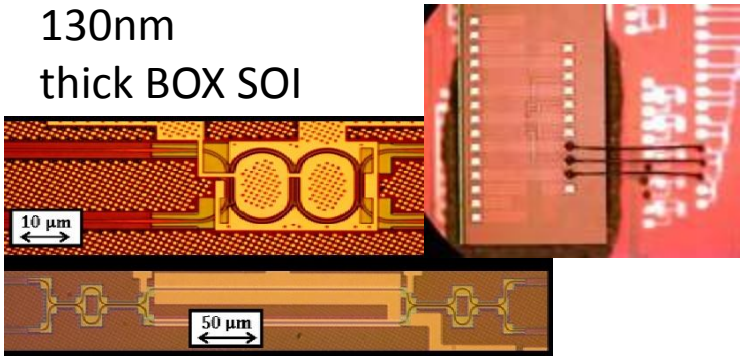

[IBM]

[Many schools]

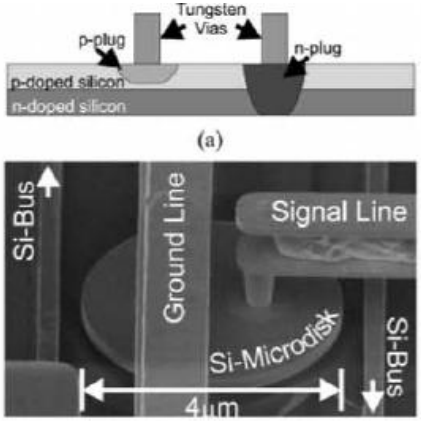

[Watts/Sandia/MIT]

[Lipson/Cornell]

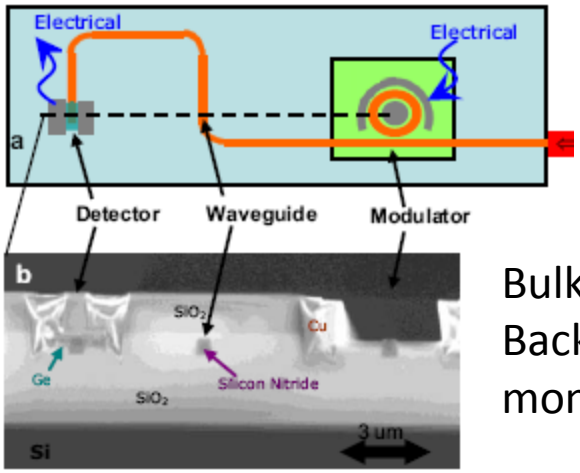

Bulk CMOS
Backend
monolithic

[Intel]

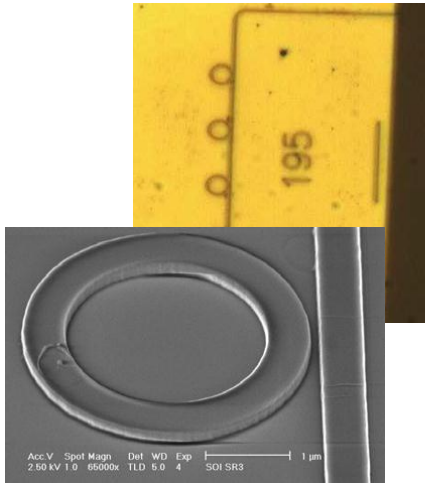

[HP]

[Kimerling/MIT]

20

EOS Platform for Monolithic CMOS photonic integration

2007

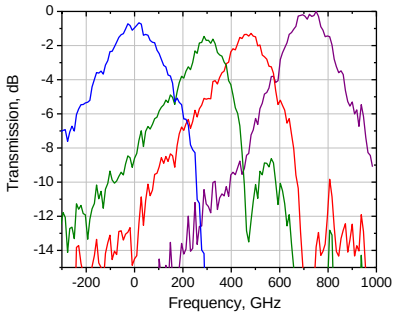

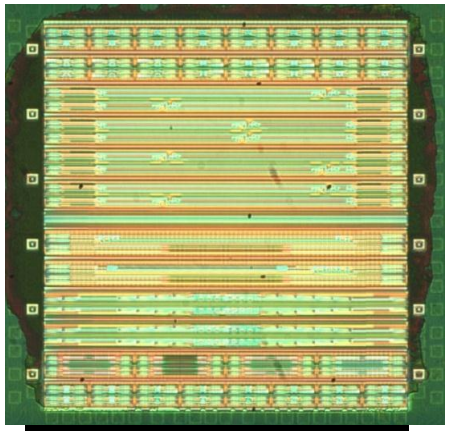

65 nm bulk CMOS
Texas Instruments

Toward 1M transistors and 100's of photonic devices

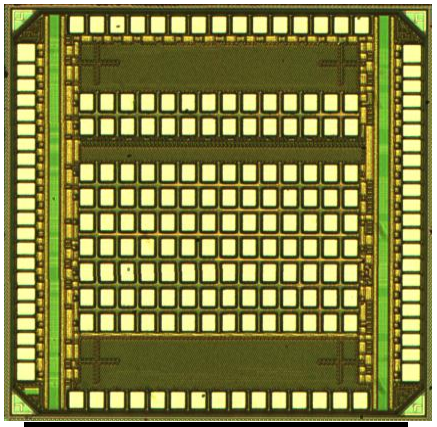

90 nm bulk CMOS
IBM cmos9sf

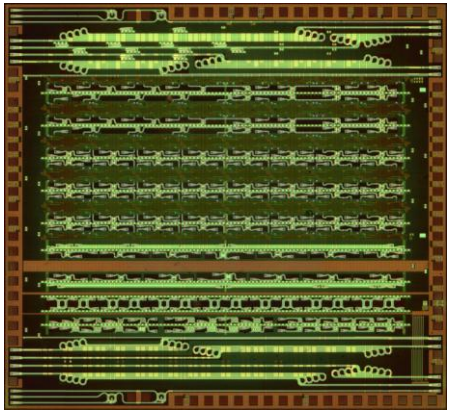

32 nm bulk CMOS
Texas Instruments

2011

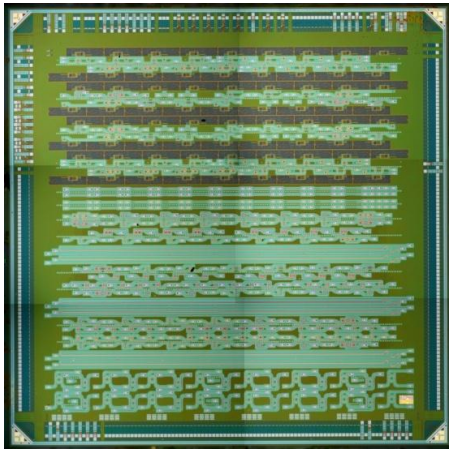

45 nm SOI CMOS
IBM 12SOIs0

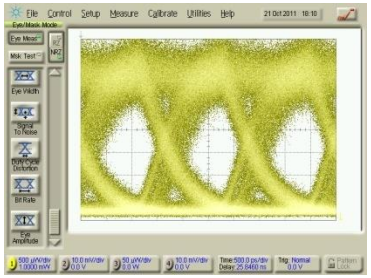

Create integration platform to accelerate technology development and adoption

EOS Platform: EOS8 fabricated in IBM12SOI

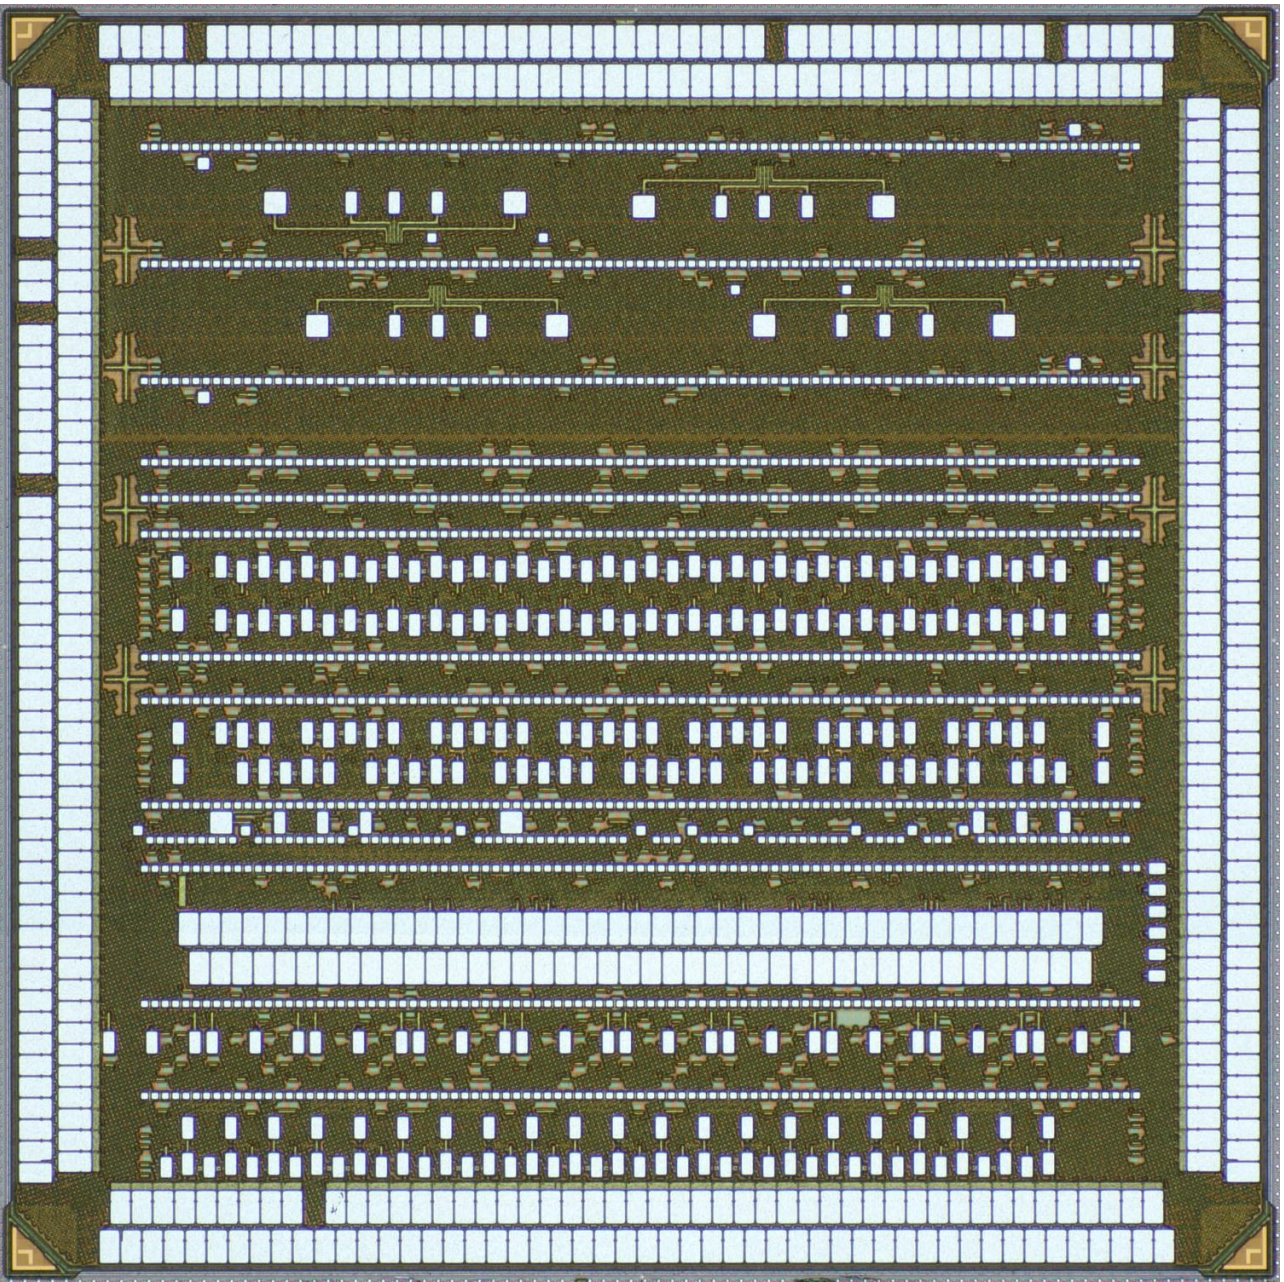

Orcutt et al,
Optics Express, 2012

3 x 3 mm die

45nm Thin Box SOI Technology
(used for Power 7 and Cell processors)

3M Transistors

400 Pads

ARM Standard Cells and
custom link circuits

EOS8 performance summary

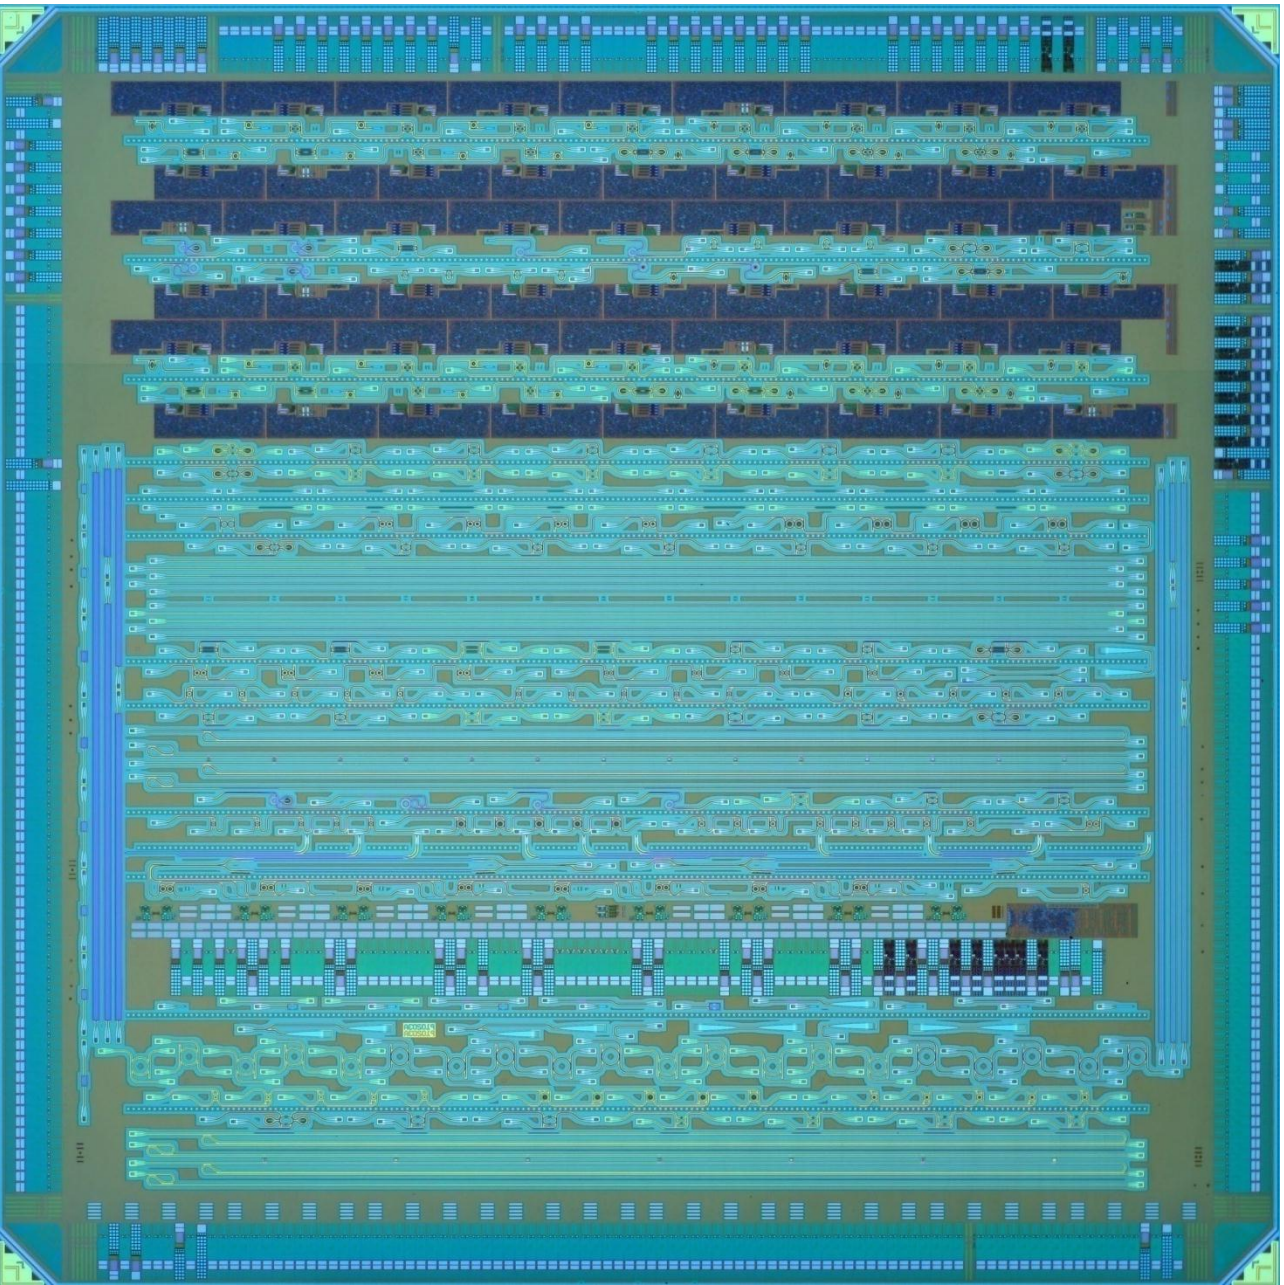

Fiber-to-chip grating couplers with 3.5 dB insertion loss

Waveguides under 4dB/cm propagation loss

10 dB extinction optical modulators

8 channel wavelength division multiplexing filter bank with <-20 dB cross talk

>20 GHz SiGe photodetectors

All integrated with electronic circuits

Full integration of photonics into VLSI tools

Layout of photonics → abstract → LEF

Layout of Circuit blocks → abstract → LEF

LEF of standard cells, I/O pads

Chip-level verilog (instantiation of .LEF macros and connectivity)

Floorplan (macro placement, power grid, routing Constraints)

Technology files

SOC Encounter Place and route

Place&routed layout

layout

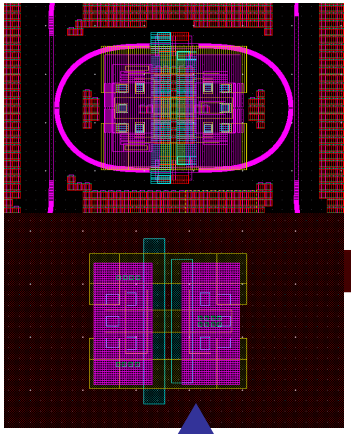

modulator.LEF

```
VERSION 5.6 ;
BUSBITCHARS "[]" ;
DIVIDERCHAR "/" ;

MACRO block_electronic_etch_row_1
  CLASS BLOCK ;
  ORIGIN -208 -1794 ;
  FOREIGN block_electronic_etch_row_1  208 1794 ;
  SIZE 2488 BY 165 ;
  SYMMETRY  X Y R90 ;
  PIN heater_a_1
    DIRECTION INOUT ;
    USE SIGNAL ;
    PORT
      LAYER ua ;
        RECT 431 1870.5 436.5 1882 ;
    END
  END heater_a_1
...
  OBS
    LAYER m1 ;
      RECT 208 1794 2696 1959 ;
...
  END
END block_electronic_etch_row_1

END LIBRARY
```

abstract

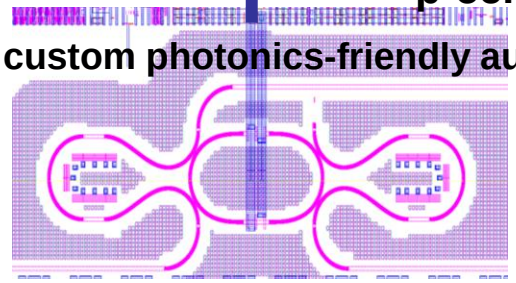

custom photonics-friendly auto-fill

Photonic device p-cell

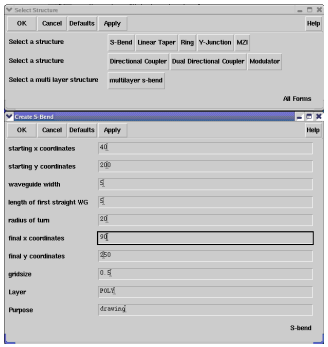

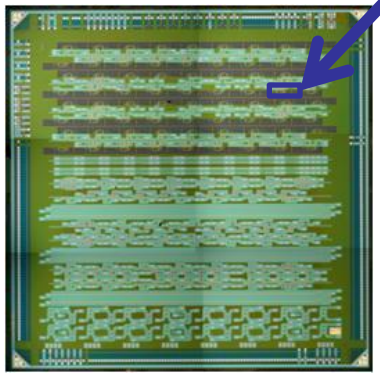

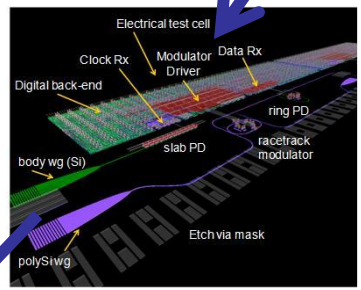

Platform Organization

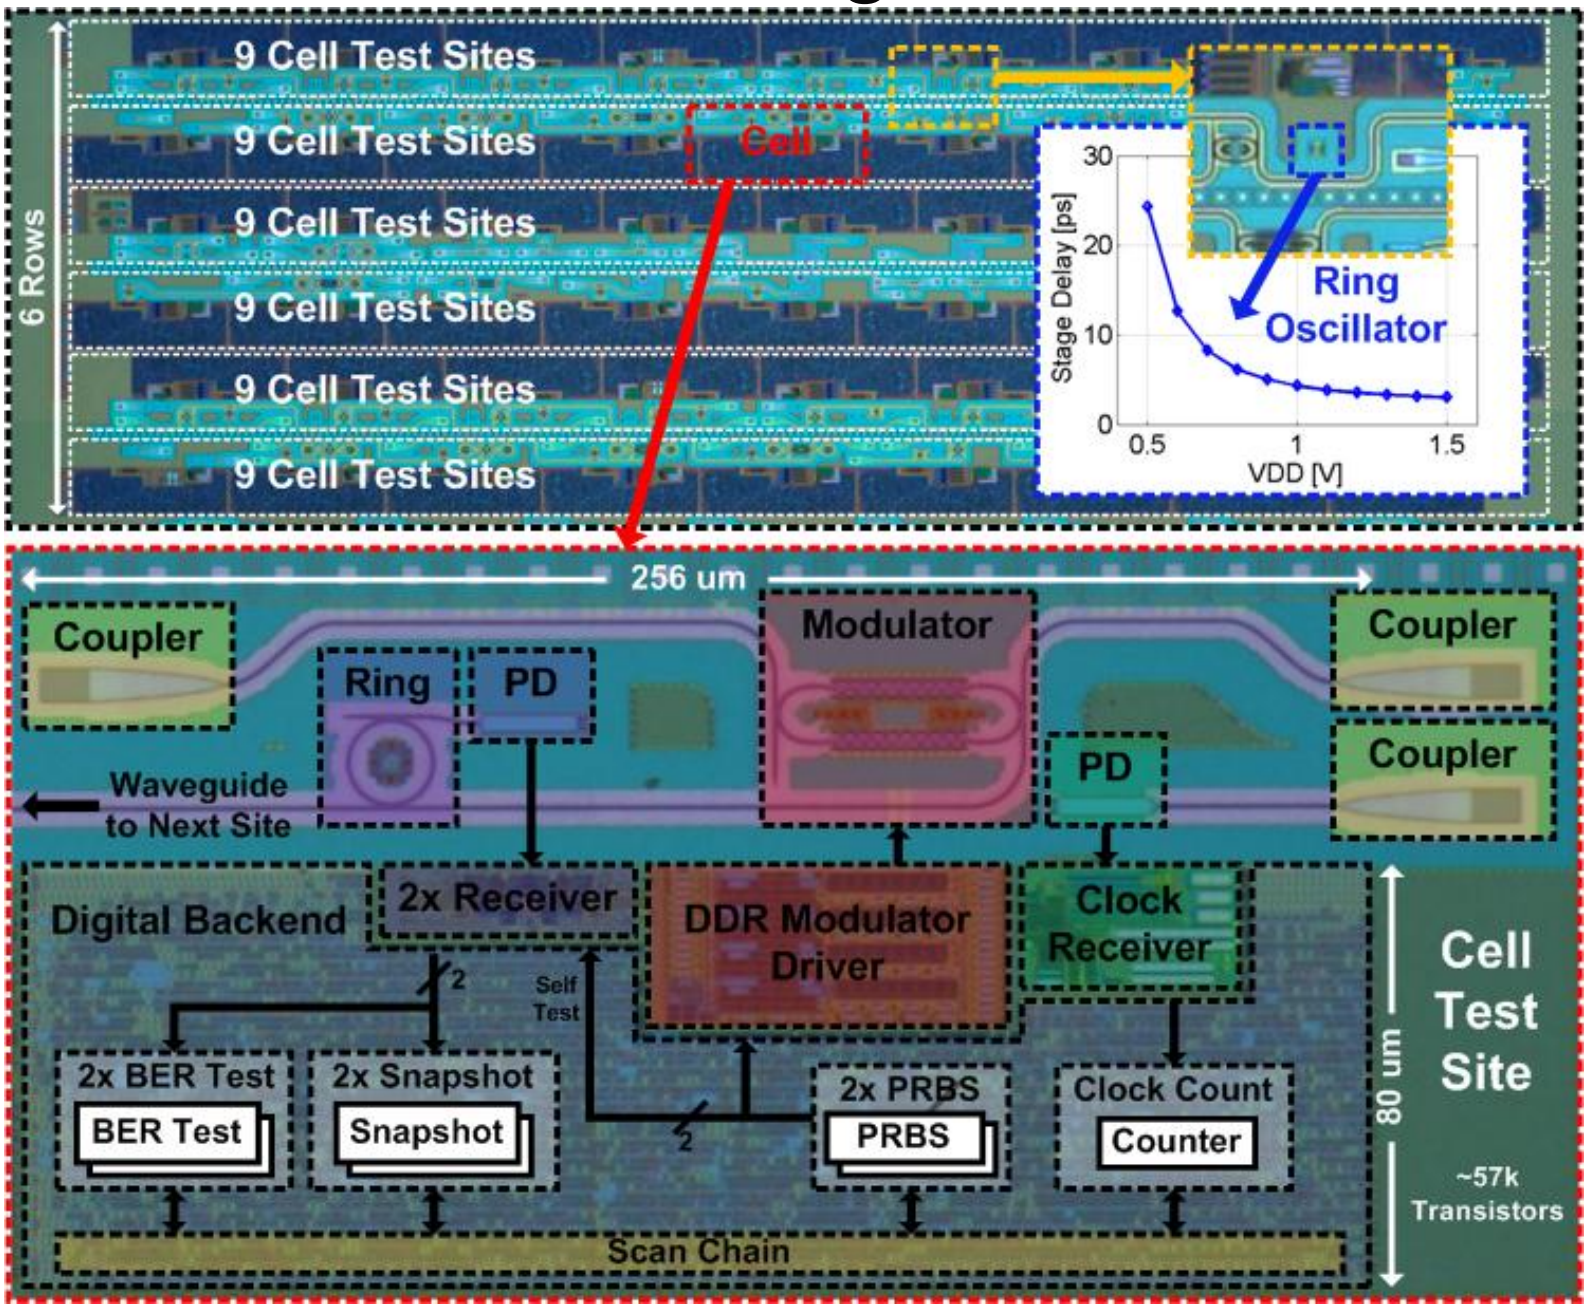

Chips fully packaged

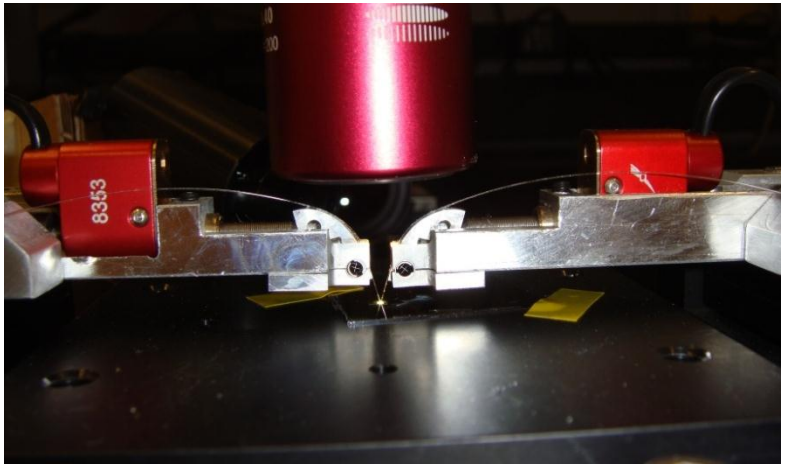

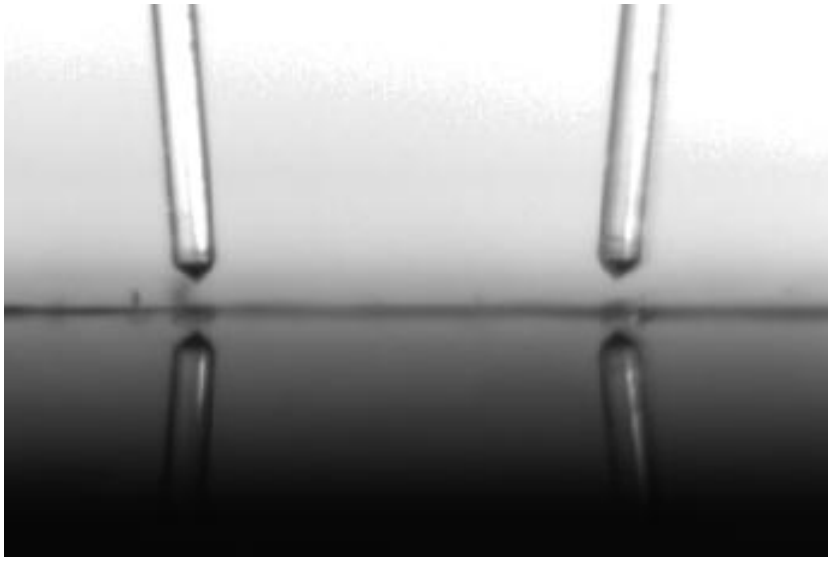

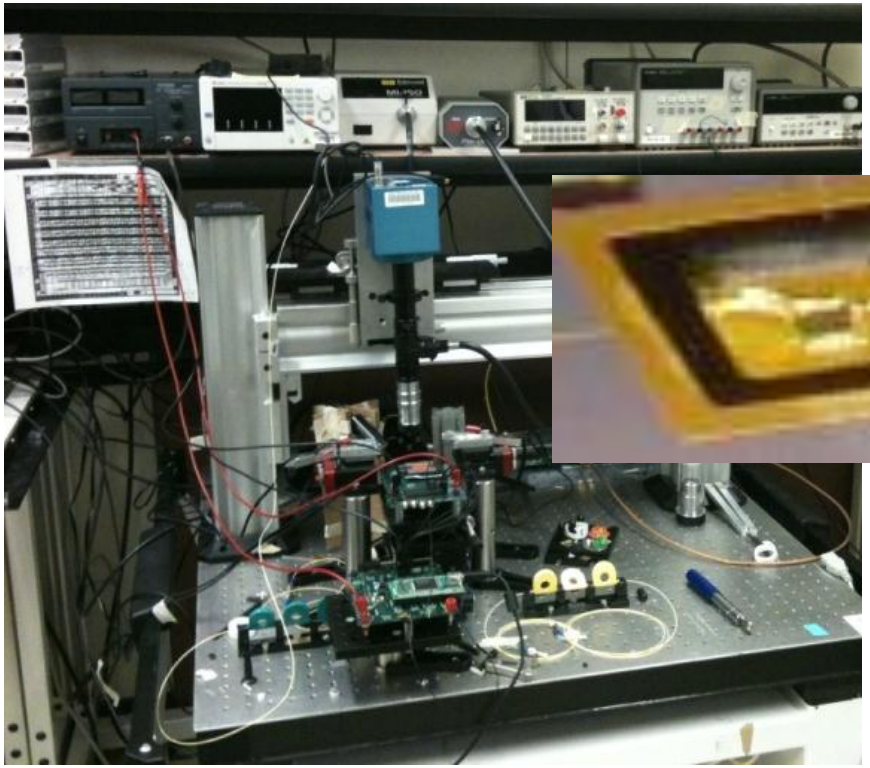

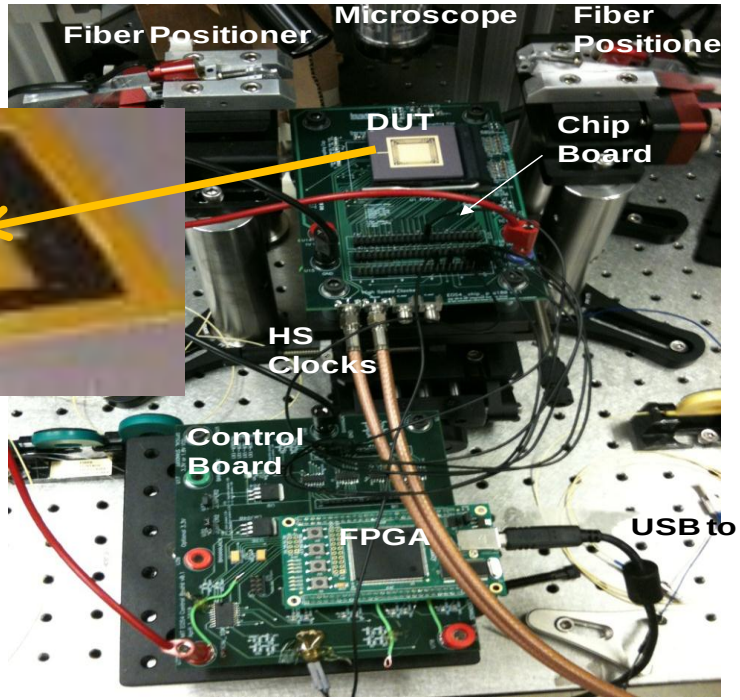

Best waveguide losses ever reported in a sub-100nm production CMOS line

- Body-Si waveguides
  - 3-4dB/cm loss
- Poly waveguides
  - 50dB/cm loss

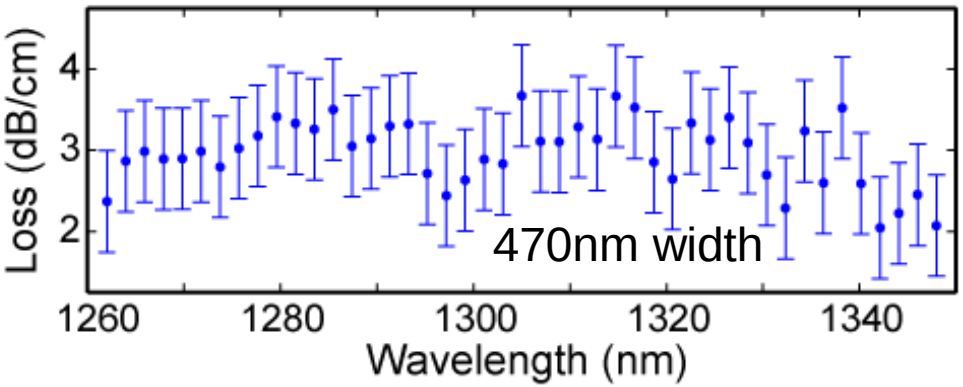

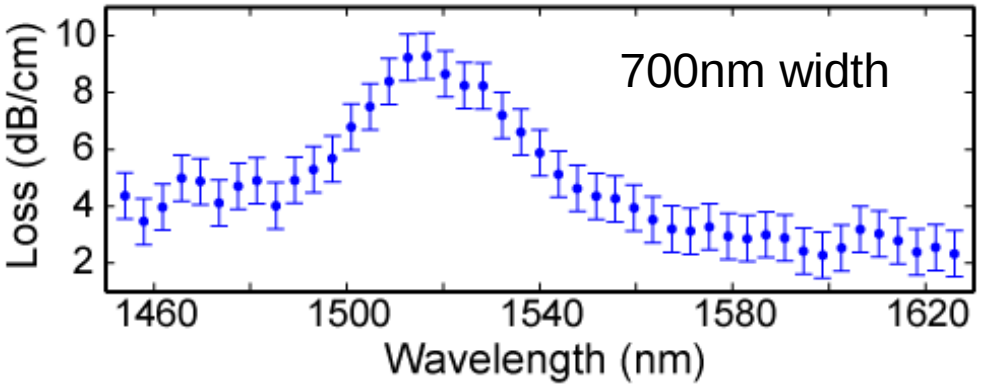

- Body-Si ring Q factor
  - 227k @ 1280nm
  - 112k @ 1550nm

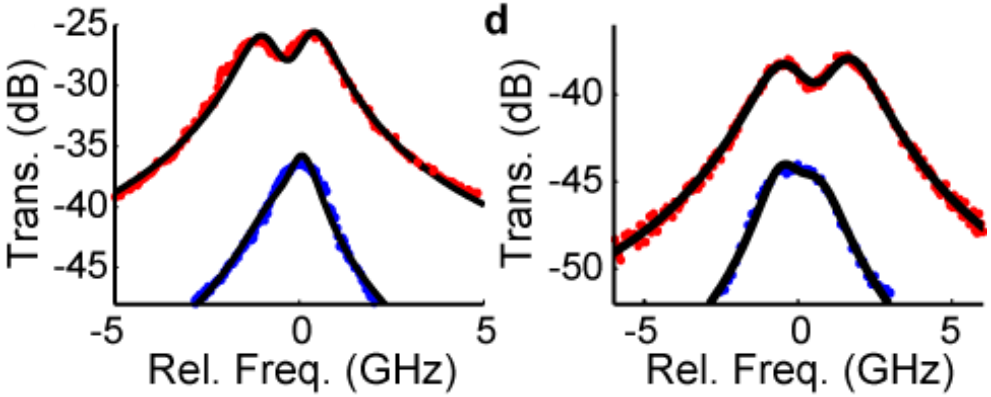

Exceptional dimensional control in 45nm node

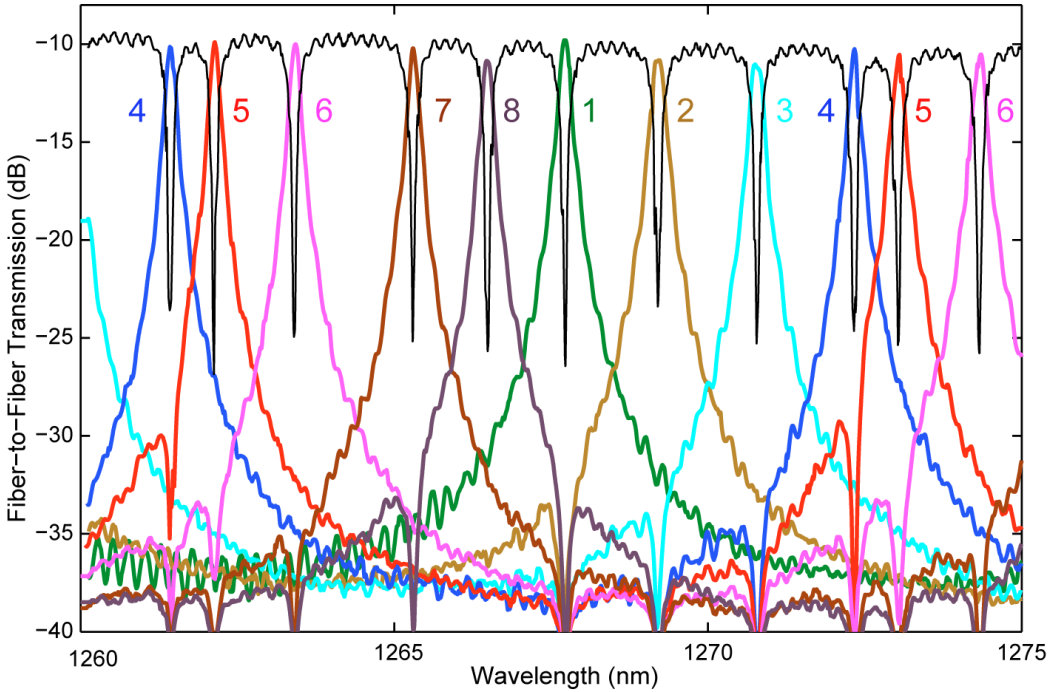

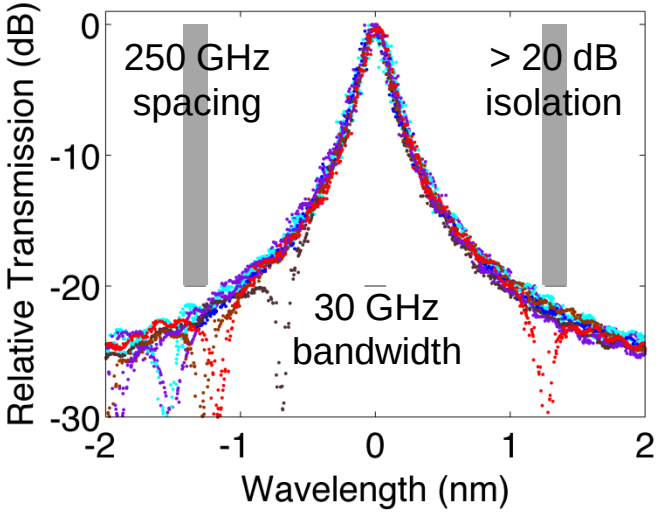

- 8-wavelength filterbank results
  - Filter channels fabricated in order
  - Less than 1nm variation
- Excellent channel isolation (>20dB at 250GHz spacing)

Integrated thermal tuning circuits

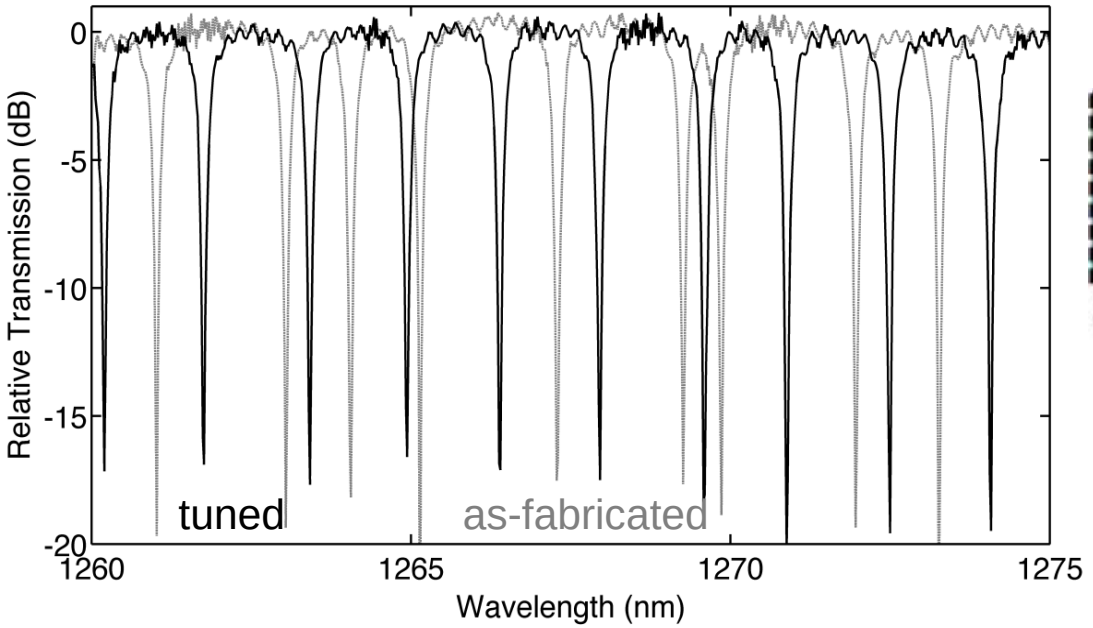

integrated digital PWM heater controller

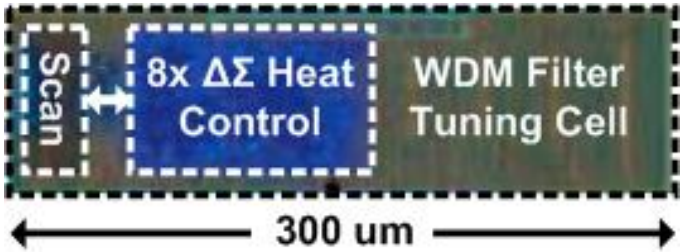

- 10mW required to retune all 8 rings
  - Negligible overhead of tuning circuits (thermal BW < 500kHz)
  - Tuning efficiency 130uW/K (32.4mW/2π) – fully substrate released chips

Low-power current-sensing optical receiver

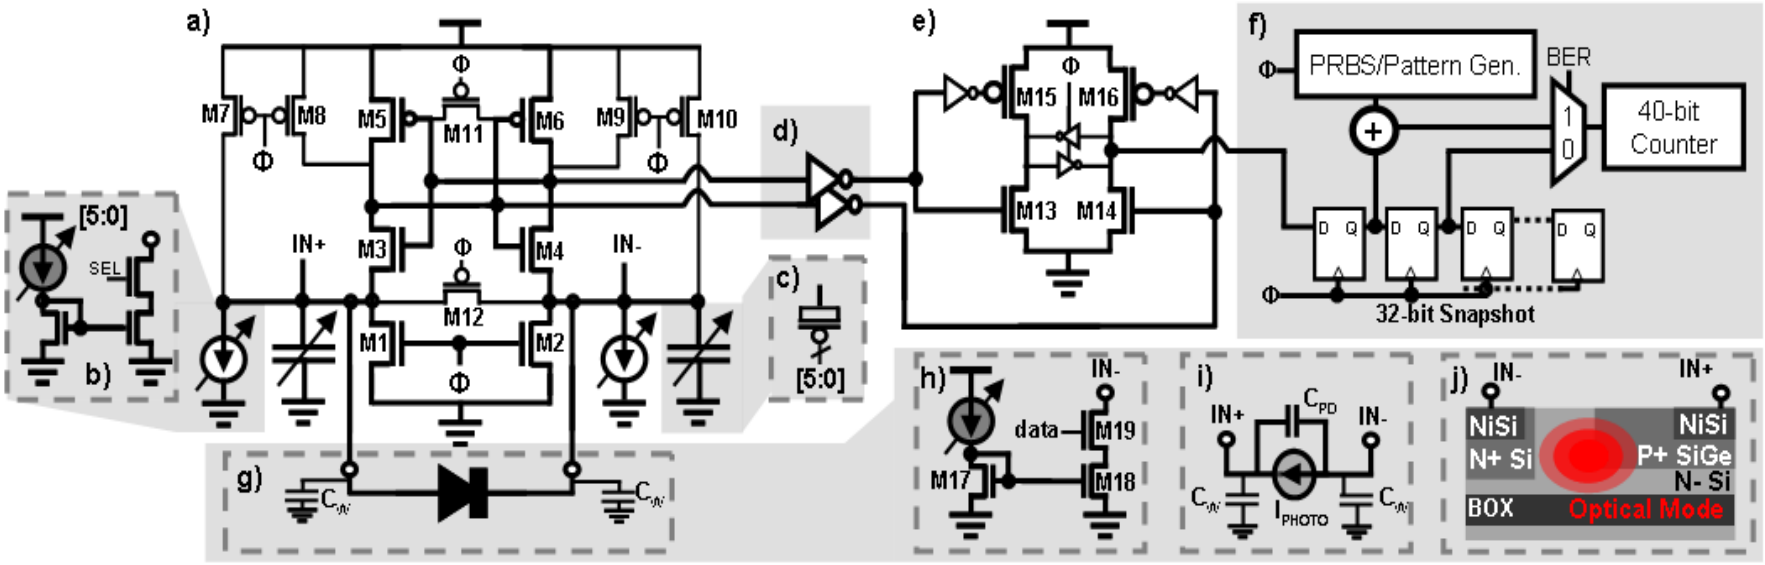

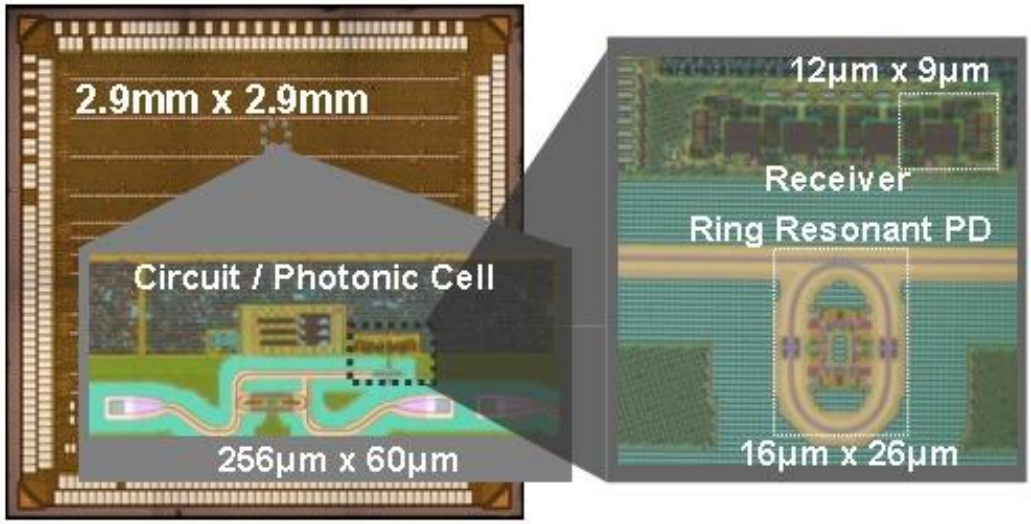

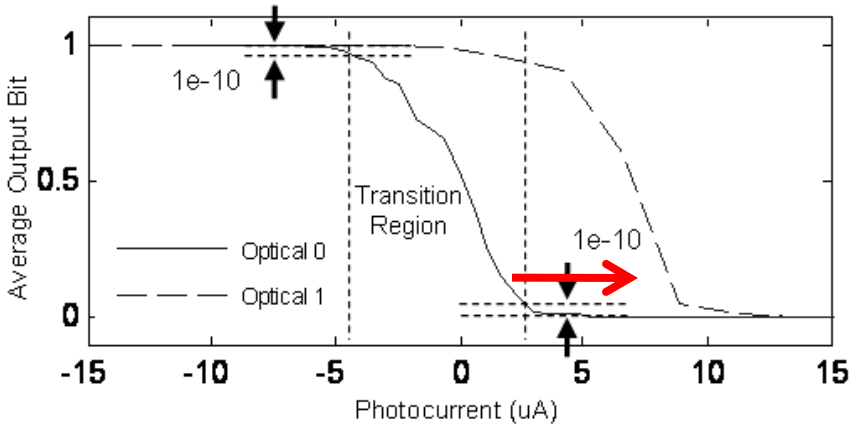

Receiver detects photo current
50fJ/b, uA sensitivities, 3-5Gb/s

Georgas ESSCIRC 2011, JSSC 2012

Optical modulator design

Shainline, Popovic

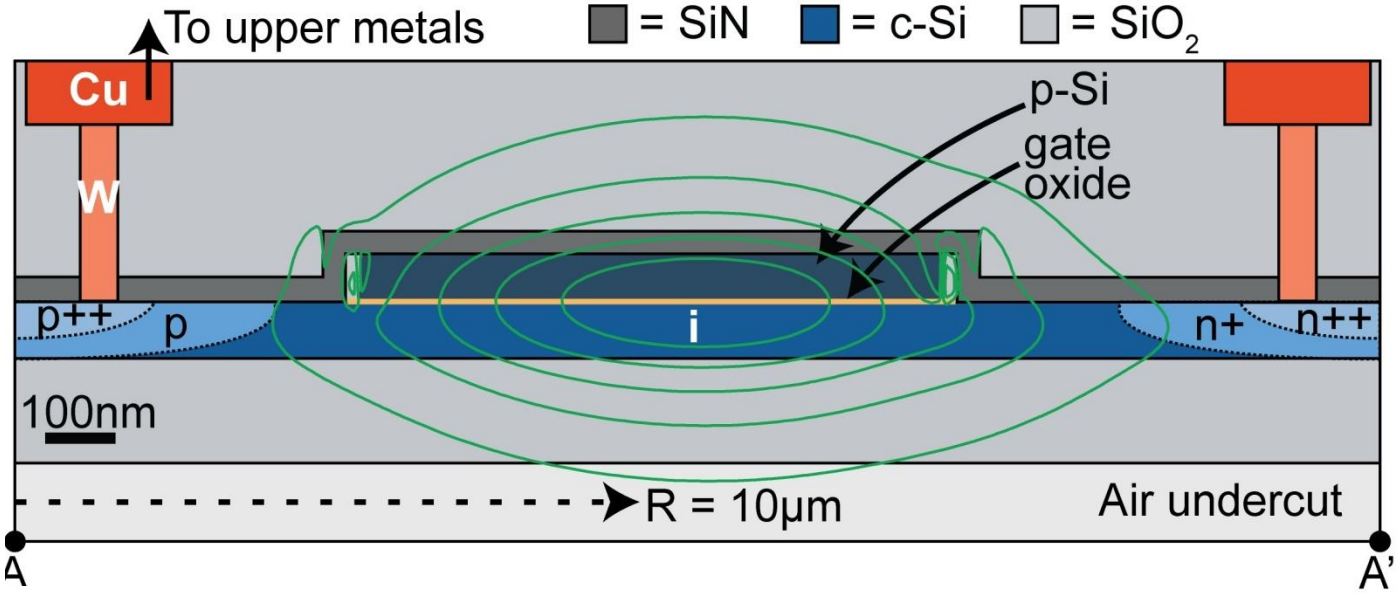

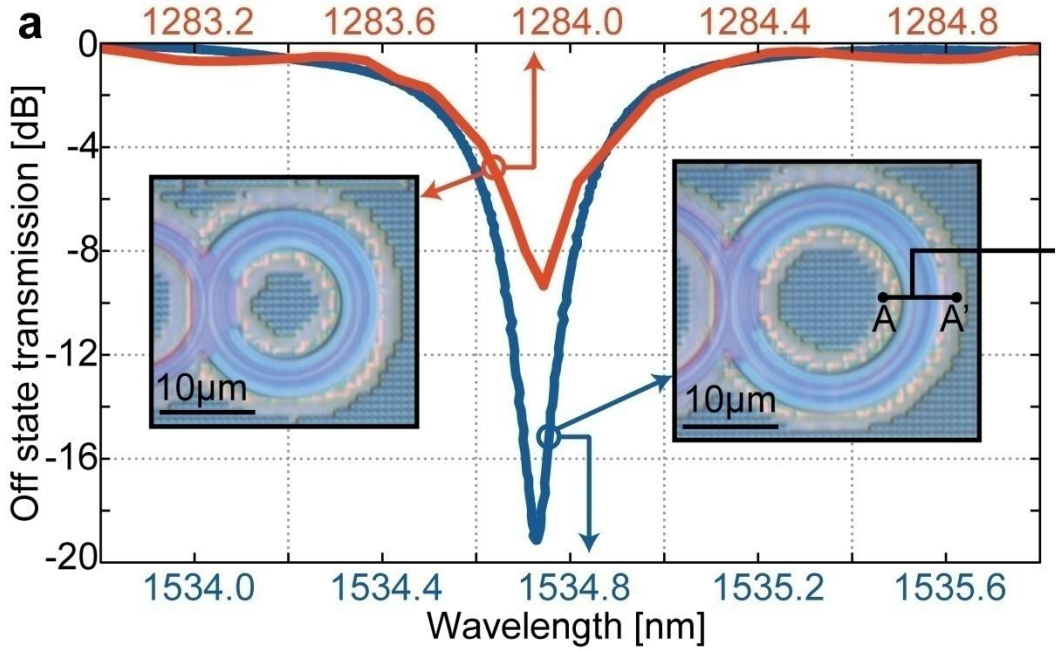

Carrier-injection device at 1550nm

- Extinction ratio 19dB
- 45GHz 3dB optical bw

at 1280nm

- Extinction ratio 9dB
- 60GHz 3dB optical bw

Optical modulator – electrical tests

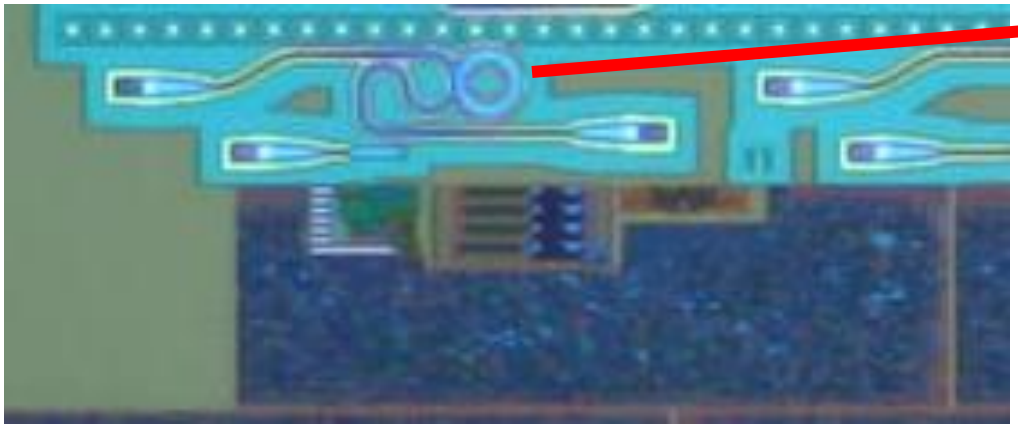

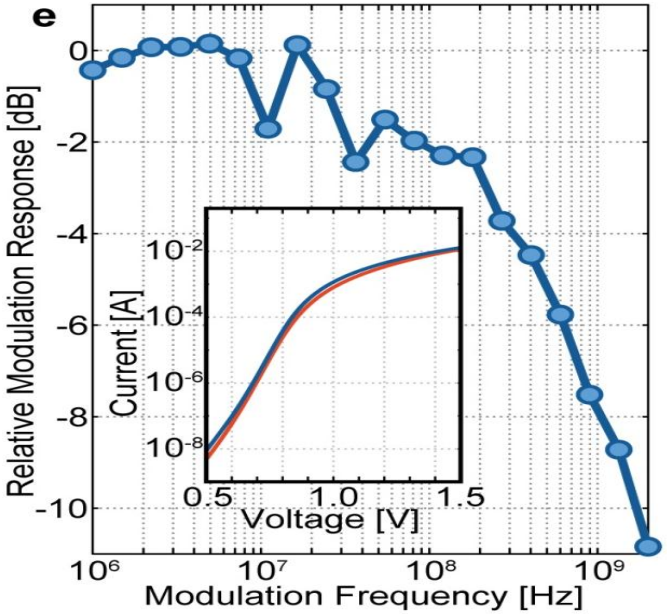

- Carrier-lifetime 2-3ns
- Diffusion time constant affected by
  - Recombination time
  - Drift conditions

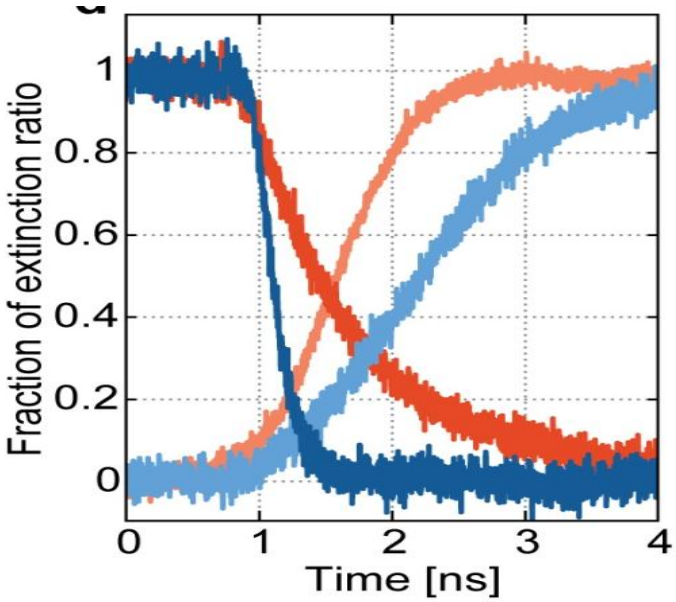

First dynamic electro-optic test in 45nm SOI

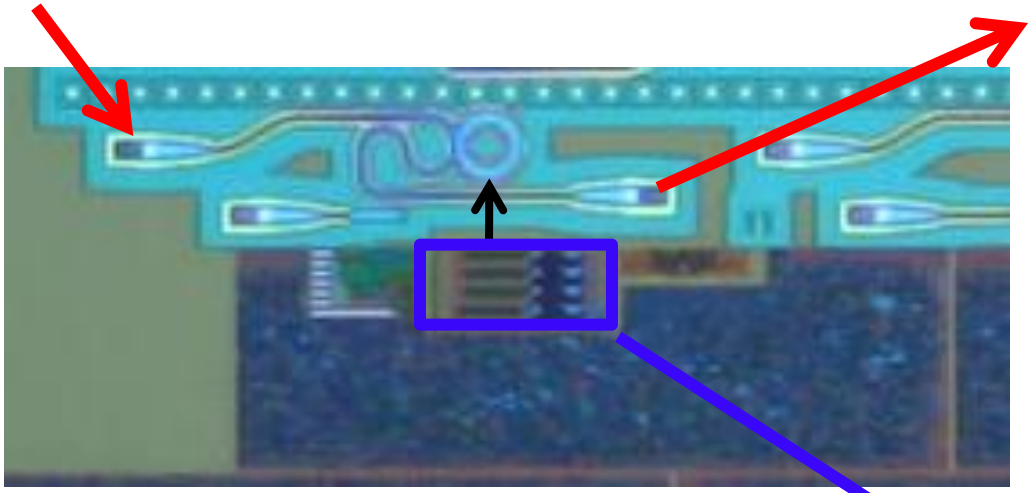

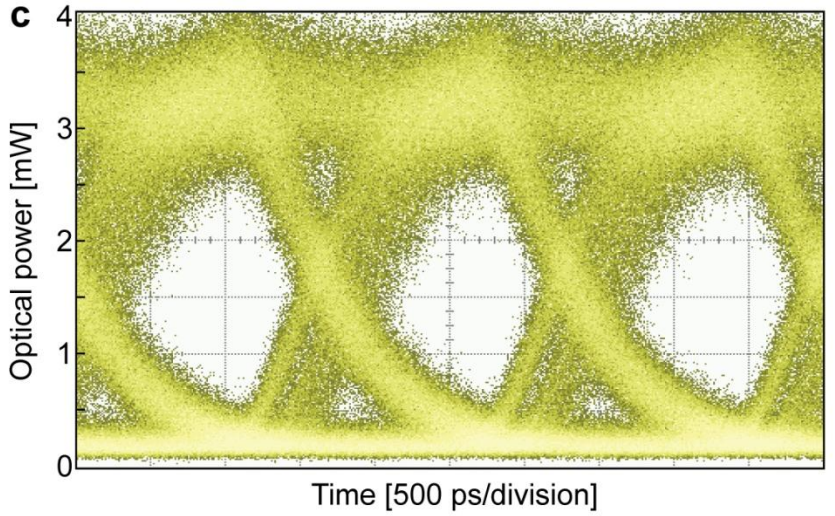

600Mb/s eyes

- Requires flexible driver
  - Split-supplies
  - Sub-bit pre-emphasis
  - 5-10Gb/s possible through waveform optimization

Transistors and Photonics can be built together in advanced CMOS!

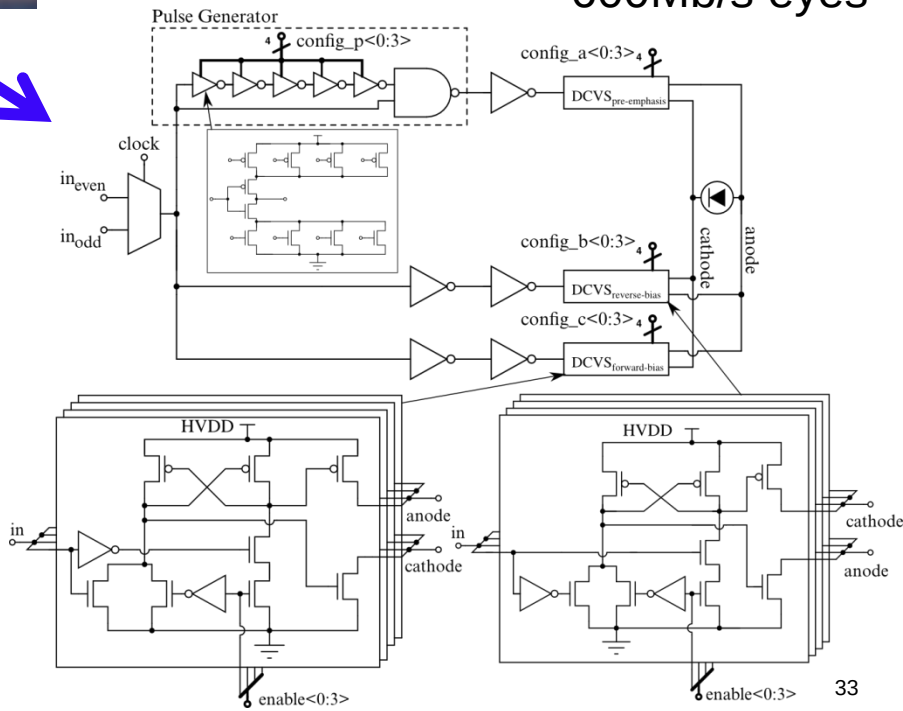

Memory interface scaling problems: Energy-cost and bandwidth density

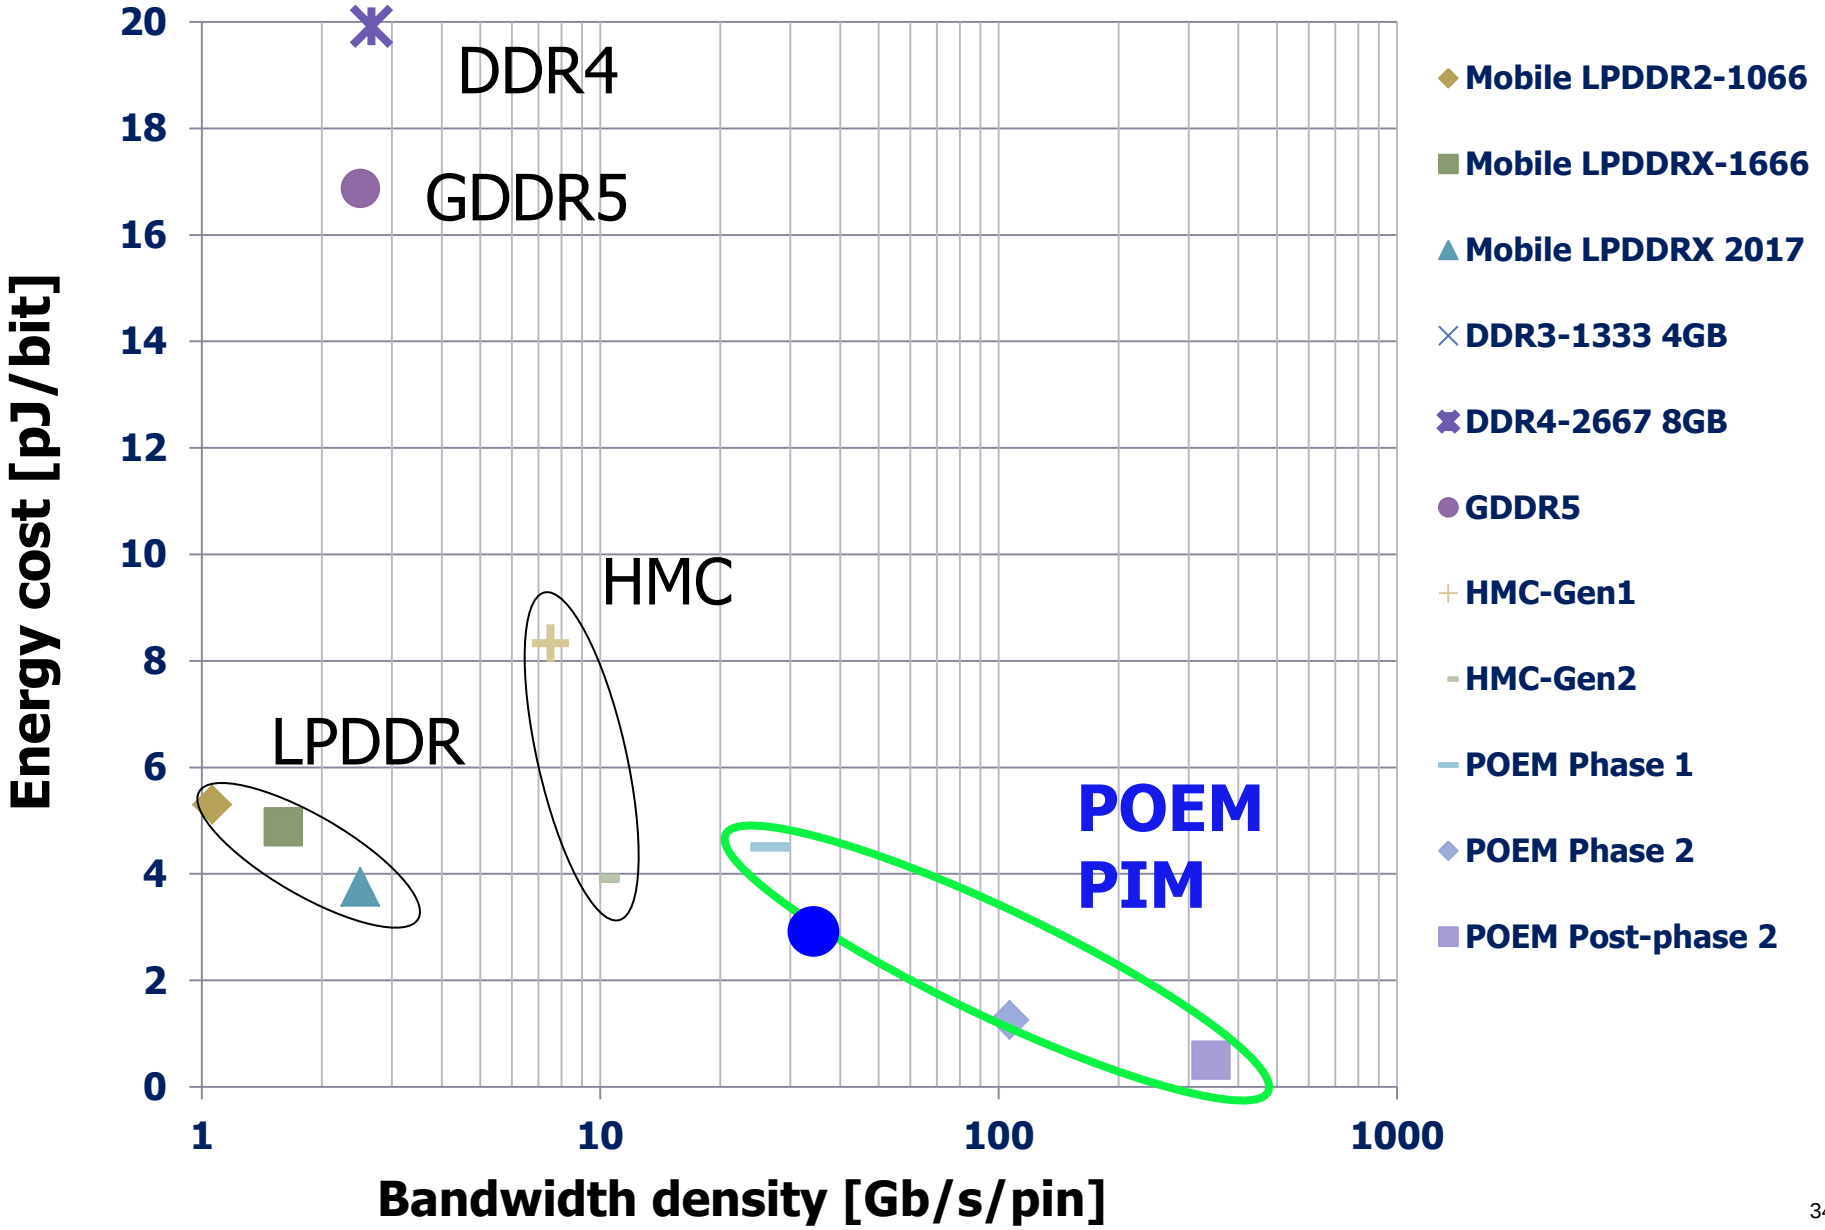

Power and pins required for 10TFlop/s

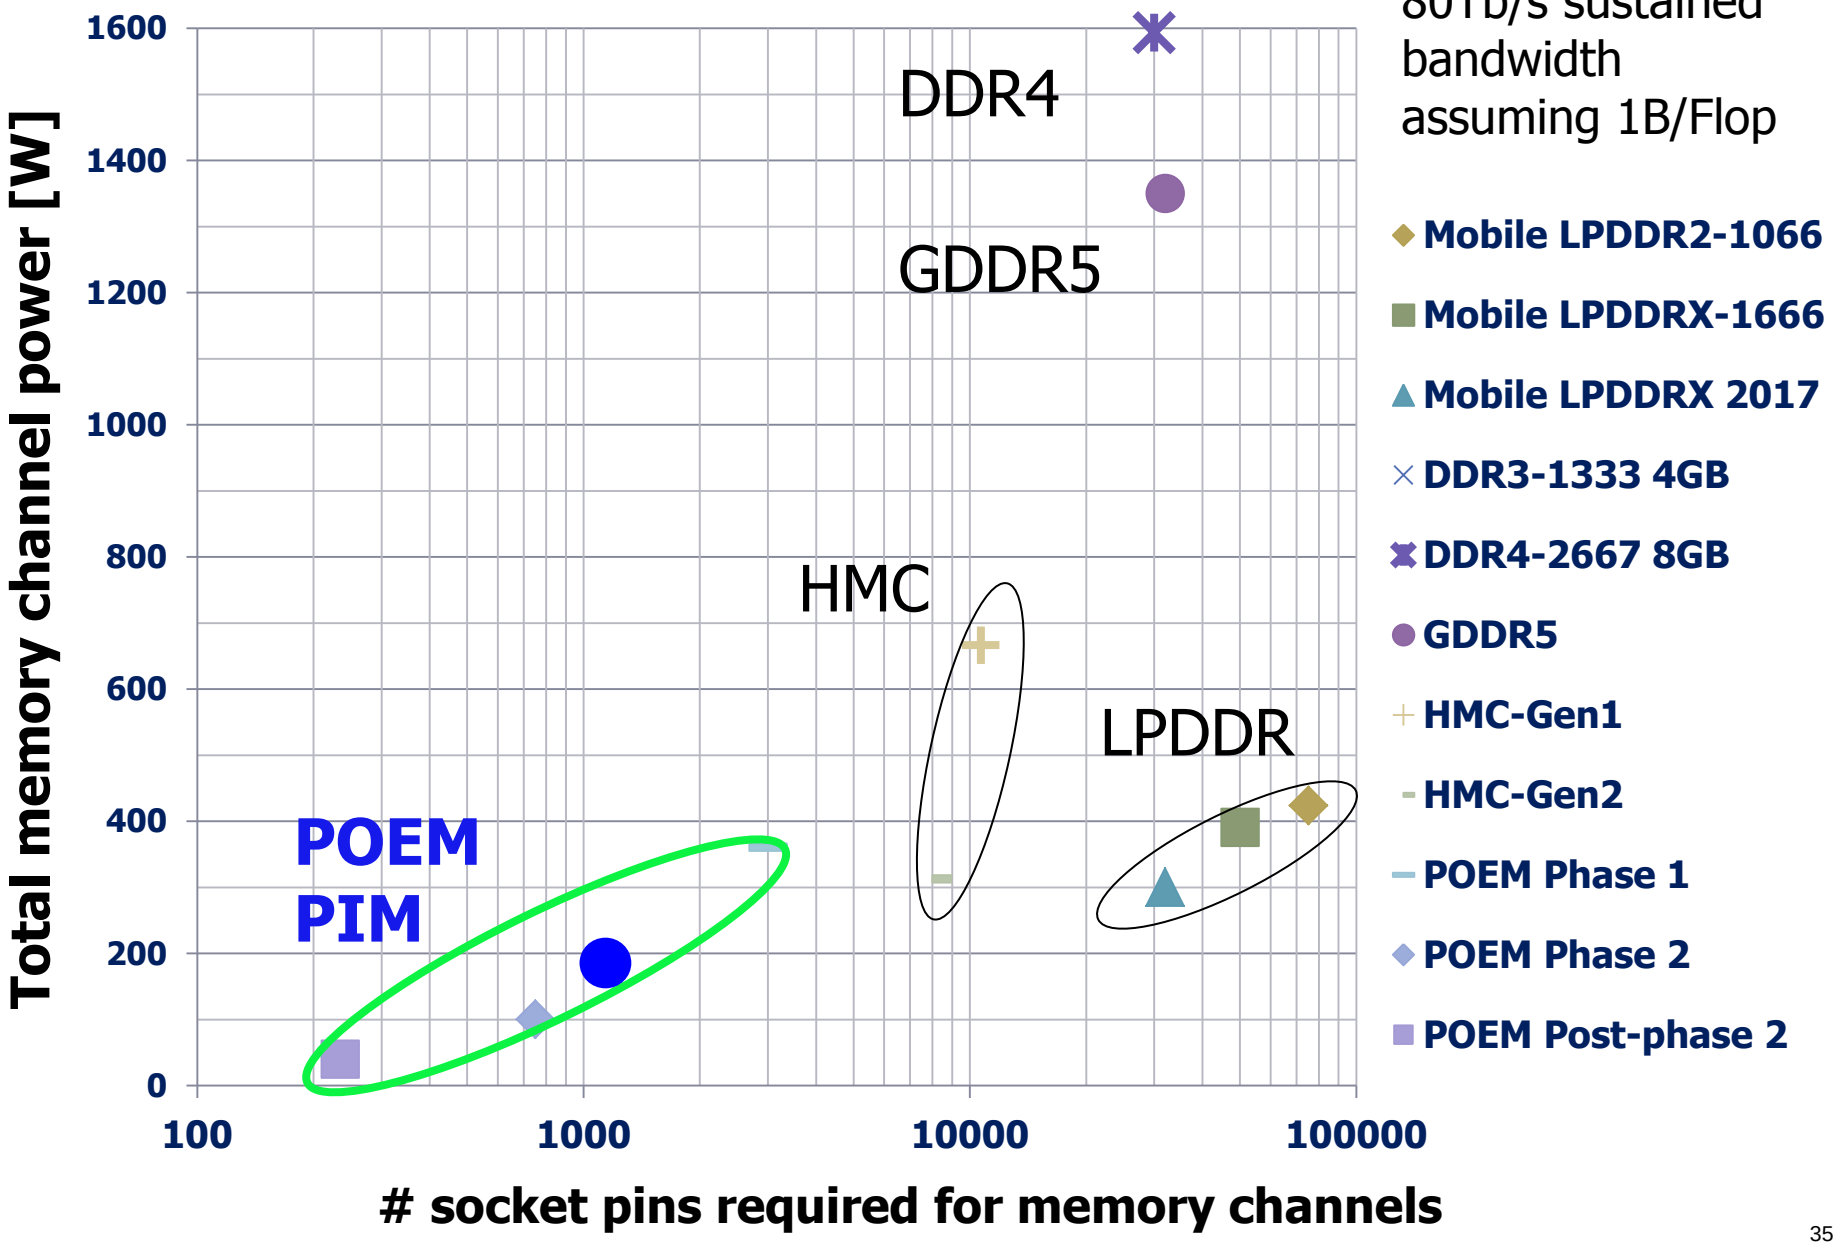

Uncooled laser sources for system efficiency

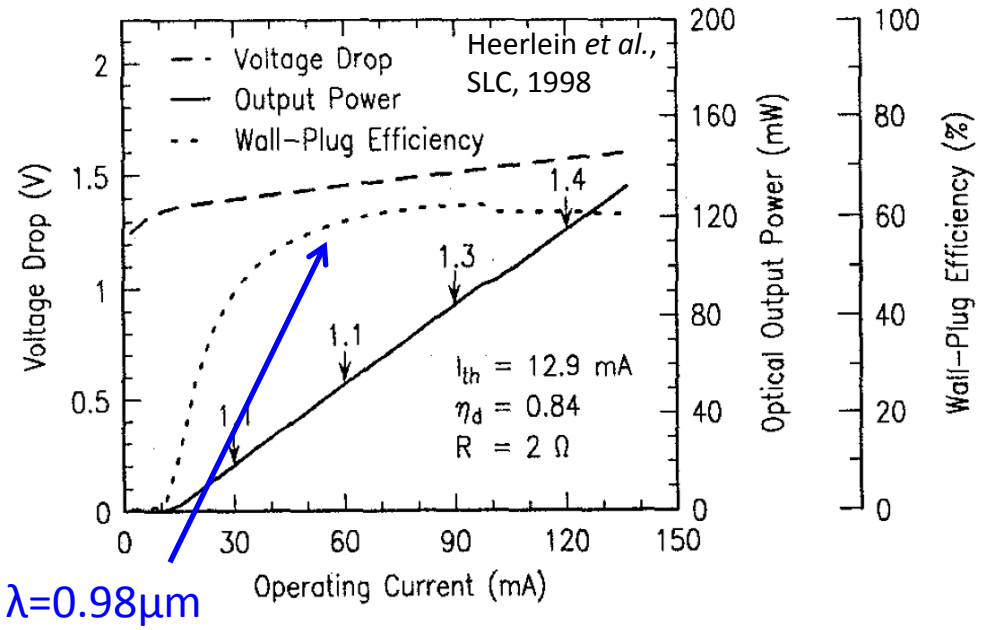

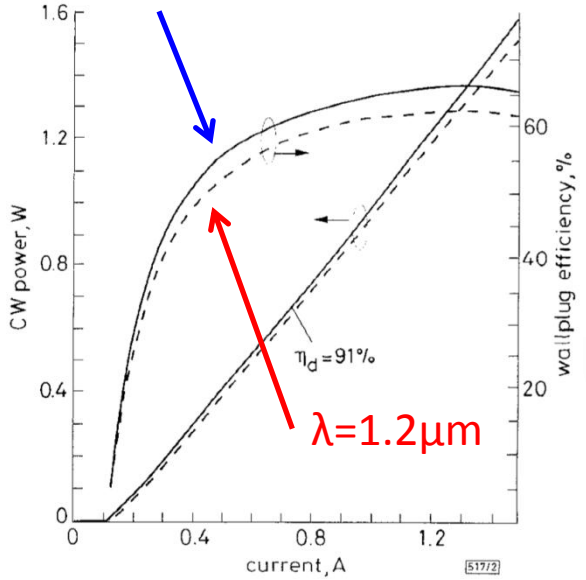

Botez *et al.*, Electronics Letters, 1996

λ=1.3µm Mitsubishi
ML7XX11 InGaAsP
Uncooled MQW DFB

35% Efficient

FTTH
Upstream

Laser Source Options (Uncooled)

- Multi-λ PIC
- FP Comb Source
- Binned DFB Bars
- Injection-Locked FP

λ = 1.2-1.3 µm Target

- Lower Laser Threshold
- Higher Published Efficiency
- Uncooled MQW Operation
- Quantum Dot Gain Media
- Larger Resonator FSR
- Smaller Optical Components

Laser reliability – Si-photonics needs fewer lasers than VCSEL links

$$FIT = \frac{\#\,Failures}{\#\,Devices} \cdot \frac{1 \times 10^{9}\ Hours}{Hours\ of\ Operation}$$

Mean Time Between Failures (MBTF) $= \frac{1 \times 10^{9}\ Hours}{FIT}$

VCSEL Laser Reliability Concerns

- Finisar 10Gb study = 2.3 FIT
- Linear data rate increases cause super-linear reliability reductions
- 100 Tbps = 10,000 VCSELs
- **MTBF = 2.3 years**
- **Intel MoBo MTBF = 19-24 years (2009-2011 Server Data)**

IBM's Blue Waters required 1M VCSELs:
Expected MTBF = 18 days

Silicon Photonics Reliability Overview

- Laser power is split for many links
- CW laser operation eliminates overdrive reliability degradation
- CyOptics 1310nm uncooled DFBs <15 FIT (200B field hour 0°C-85°C) including direct-mod. operation
- 100 Tbps = 64 DFBs (1 laser per λ)
- **MTBF @ 15 FIT/laser = 120 Years**

λ=0.98µm Pump Laser Reliability

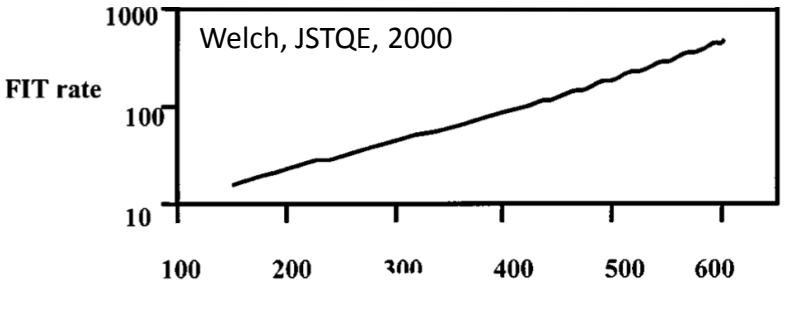

Packaging

- CPU package
  - Flip-chip <5um C4 tolerance o.k. for coupling
- DRAM package
  - Die on board
  - Connector-to-fiber alignment <2um

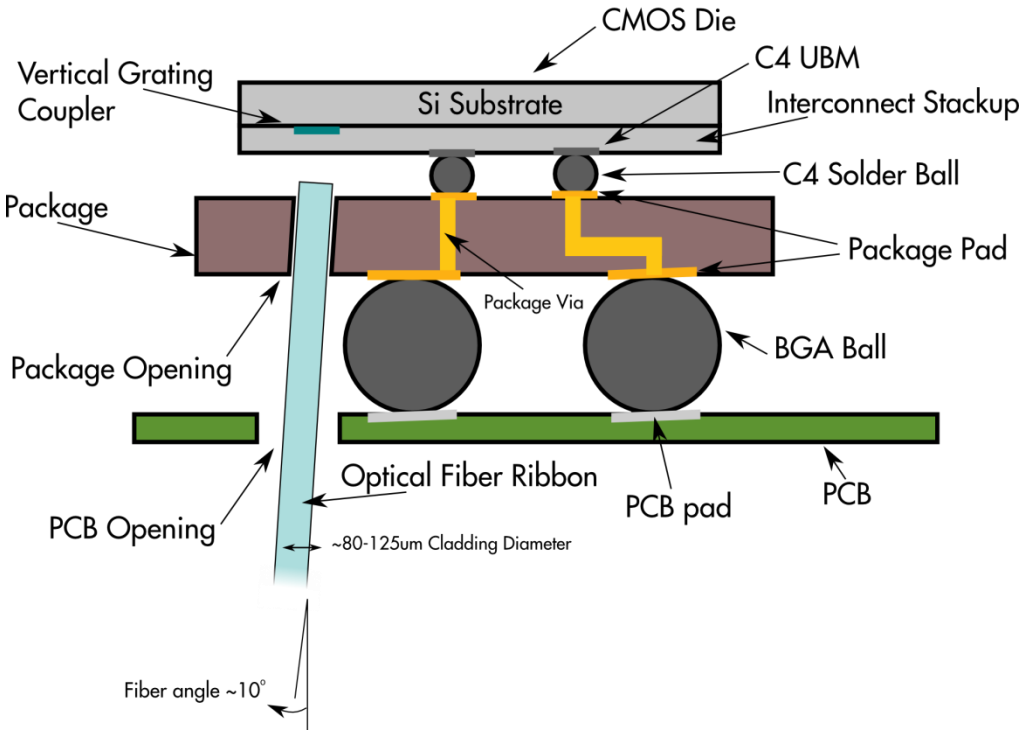

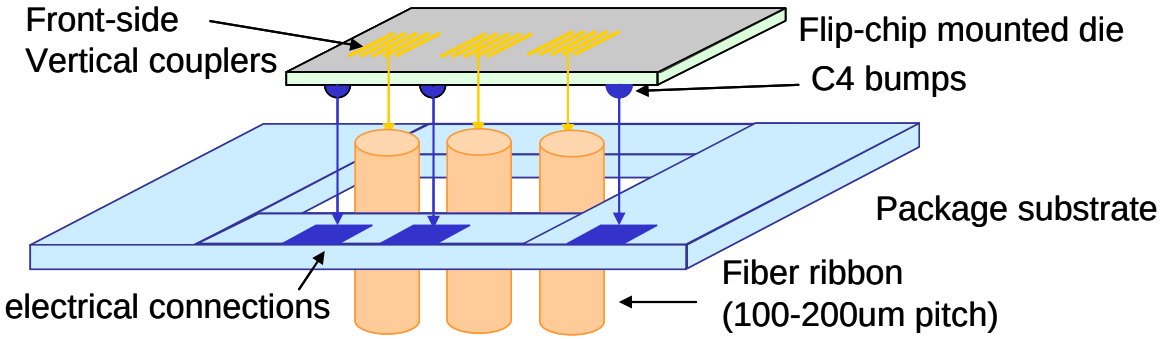

Summary

- Silicon-photonics can push both critical dimensions
  - Energy-efficiency – monolithic integration
  - Bandwidth Density - dense WDM
- Need to optimize across layers
  - Connect devices to circuits, and links to networks
- Building early technology development platforms
  - Feedback to device and circuit designers
  - Accelerated adoption
- EOS Platform designed for multi-project wafer runs
  - Best end-of-line passives in sub-100nm process (3-4dB/cm loss)
  - 50 fJ/b receivers with uA sensitivities
  - Record-high tuning efficiency with undercut ~ 25uW/K
  - First modulation demonstrated in 45nm process