*Exceptional service in the national interest*

Sandia National Laboratories

SAND2016-10432C

# Challenges for optical interconnect for beyond Moore's law computing

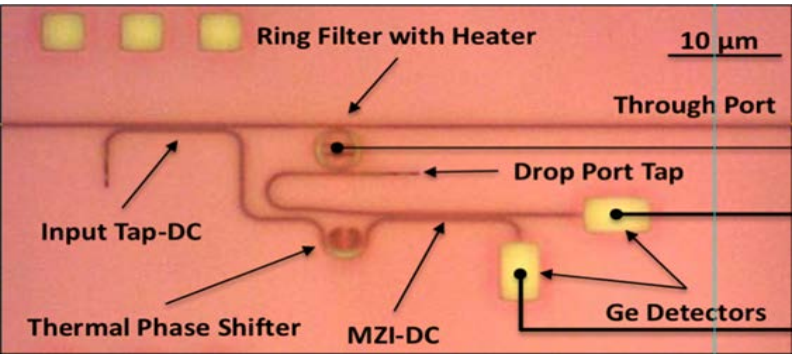



**Anthony L. Lentine and Christopher T. DeRose**

**Sandia National Laboratories, Albuquerque NM**

alentine@sandia.gov

U.S. DEPARTMENT OF ENERGY | NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

# Optics in computing and data centers

Sandia National Laboratories

- High bandwidth connections between racks of equipment
  - It's about bandwidth,
    - not power, not cost
  - 10 → 25 Gb/s per serial channel, 100 Gb/s per fiber pair or ribbon
  - 400 Gb/s on the horizon
- Just barely permeating inside of rack or boxes

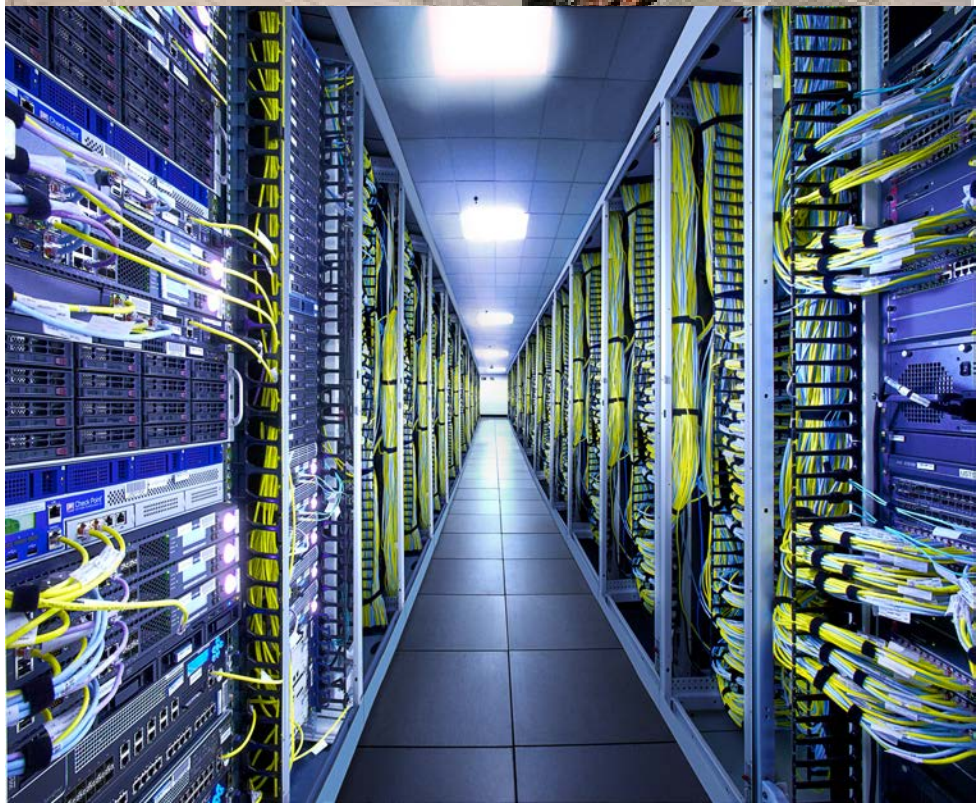

# Revolutionary vs. evolutionary optics

Sandia National Laboratories

- **Evolutionary (Modules)**
  - 100 Gbps modules available
    - Expensive, big, power hungry
  - 400 Gbps becoming available
    - Expensive, big, power hungry
  - **1000 Gbps on the horizon**

- **Revolutionary (3DI)**
  - Higher bandwidth density
  - Drastic *potential* power reduction
    - No 50 Ω lines, pre-emphasis or equalization
    - Simple low power Rx
    - Shared CDR (less delay variation.)
- **Can use this technology in TbE transceivers**

Not addressing on-chip interconnects

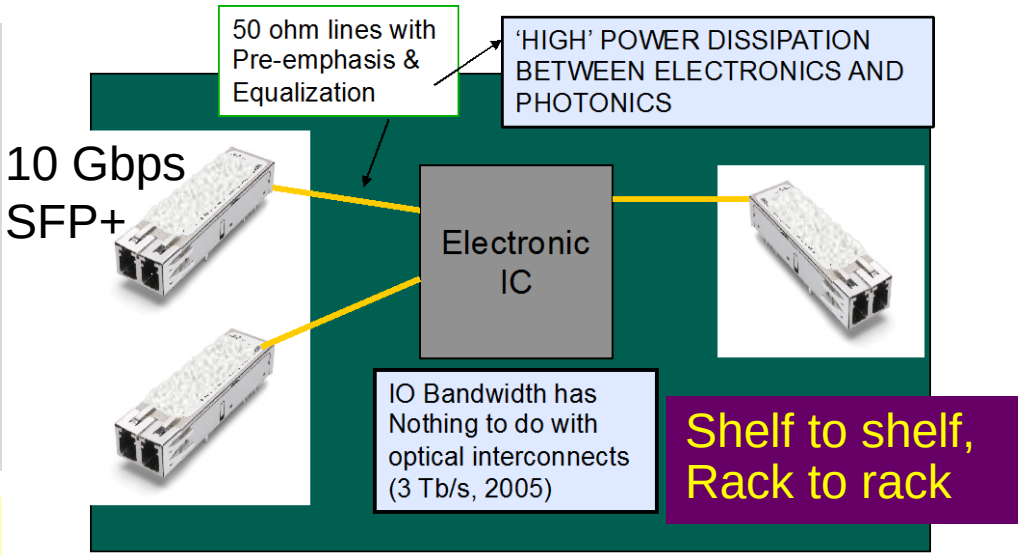


**OPTICS FOR DISTANCE**

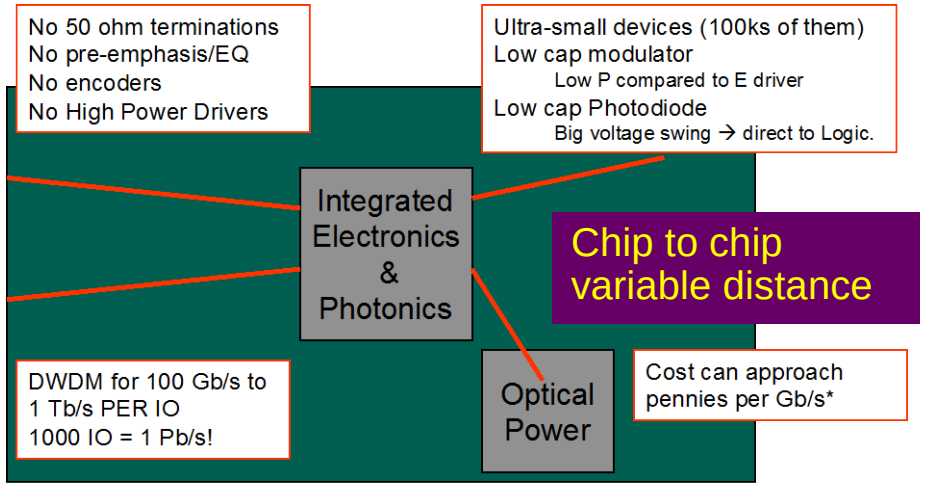


**OPTICS FOR *LOW POWER*, *HIGH BANDWIDTH DENSITY*, COST, SIZE, WEIGHT, DISTANCE**

# Intimate heterogeneous integration of photonics with high-value CMOS

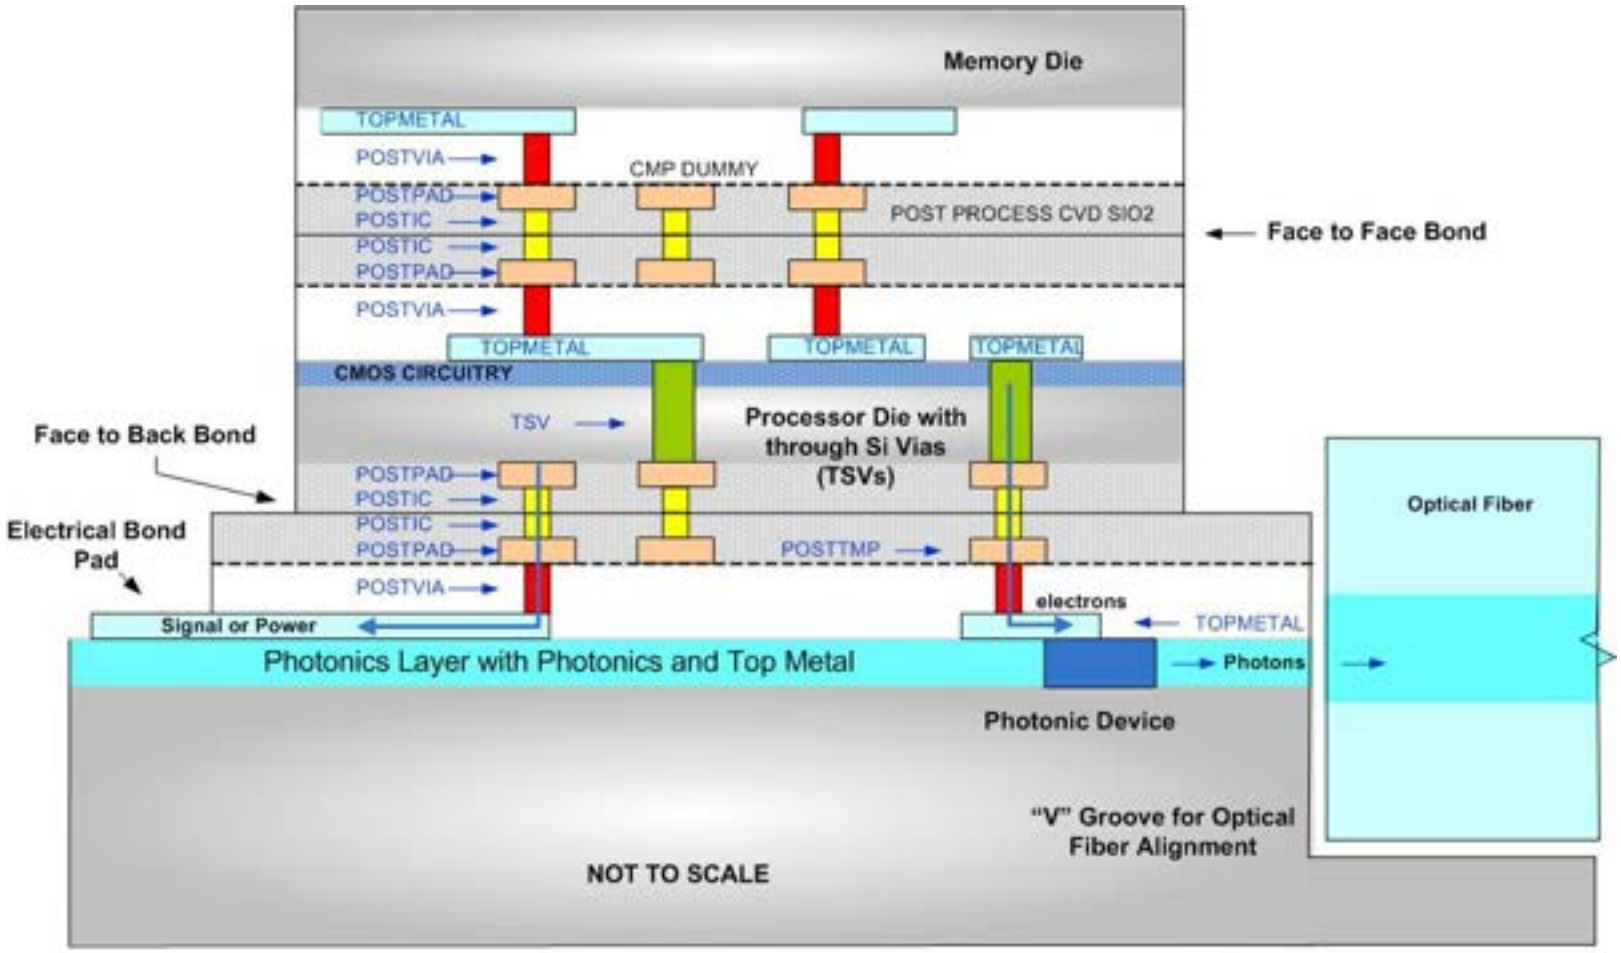


Picture courtesy of Subash Shinde

# Optical PCBs: a new growth curve?

Sandia National Laboratories

- Fiber connections will never be cheap enough to allow the use of optics to 'explode'
  - DWDM 40 λ is a first step
- Need something (old) new …

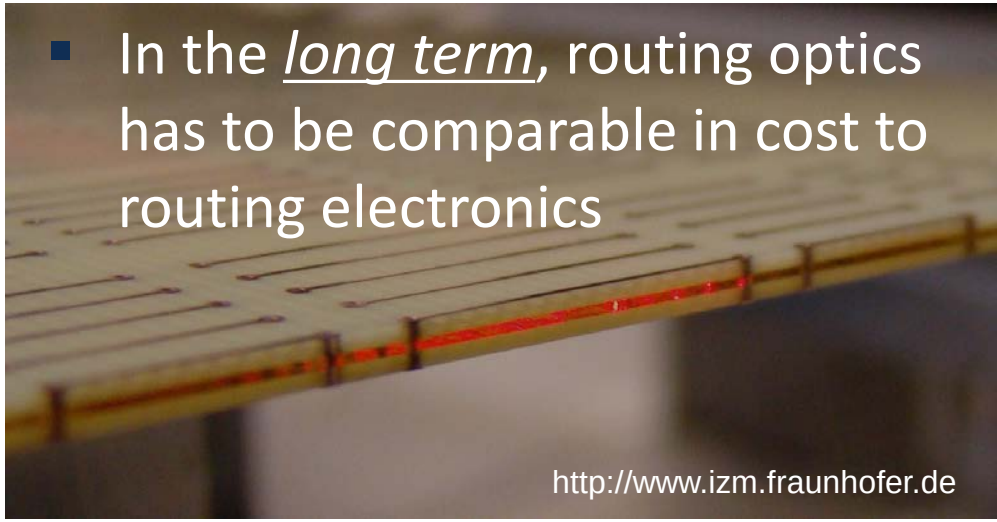


- Fiber-less on board connections
- Surface normal connections from OE chips
- Simple alignment?
- DWDM, MDM, not single wavelength multimode
- Reliable as electronics

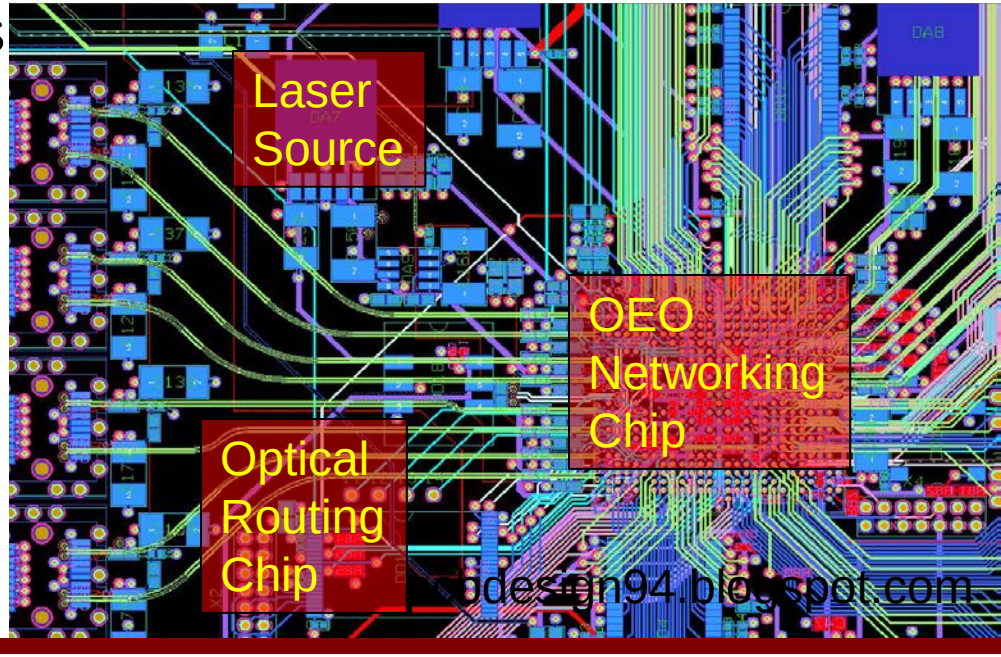

# Modulators vs. direct laser modulation

Sandia National Laboratories

**Light emitters - Pro**

- No additional optics to get the light to the output device
- Only need to turn the lasers on for the channels in use

**Light emitters - Con**

- Difficult integration with electronics
- Difficulty of wavelength control
- May require optical isolation
- May require polarization and mode control
- Relaxation oscillation limit to bit rate
- May have timing issues from turn-on delay
- All power dissipation is on-chip
- Temperature challenges
  - different shifts of bandgap and resonator wavelengths
  - decrease of laser gain with temperature

**Modulators - Pro**

- Centralized wavelength, mode, and polarization control, and optical isolation,
- Optical pulses for precise signal timing
- Only modulator drive power is on-chip
- Many approaches tolerant to high-temp.
- Can be compatible with WDM

**Modulators - Con**

- Separate light source required
- Needs optics to split and deliver the power to the many modulators
- *All illuminated modulators consume at least the optical drive power even if not driving any signals*

- Ref: D. A. B. Miller, Attojoule Opto… arXiv:1609.05510 (2016)
- Also see A. L. Lentine, IEEE LEOS, **1997**

# O/E Integration challenges

Sandia National Laboratories

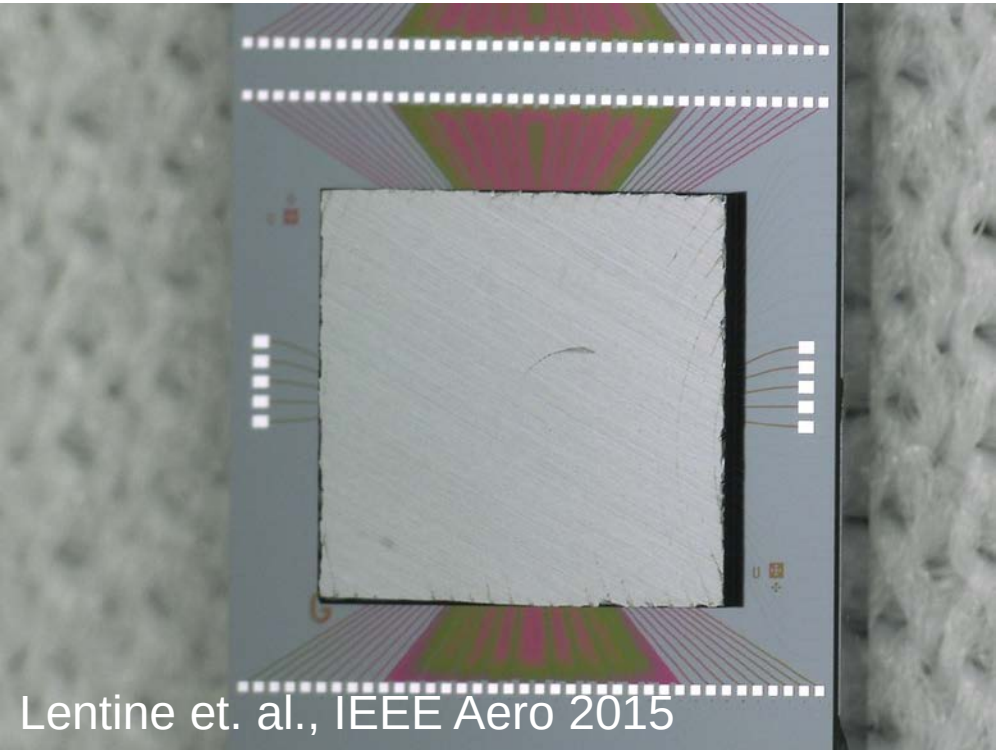


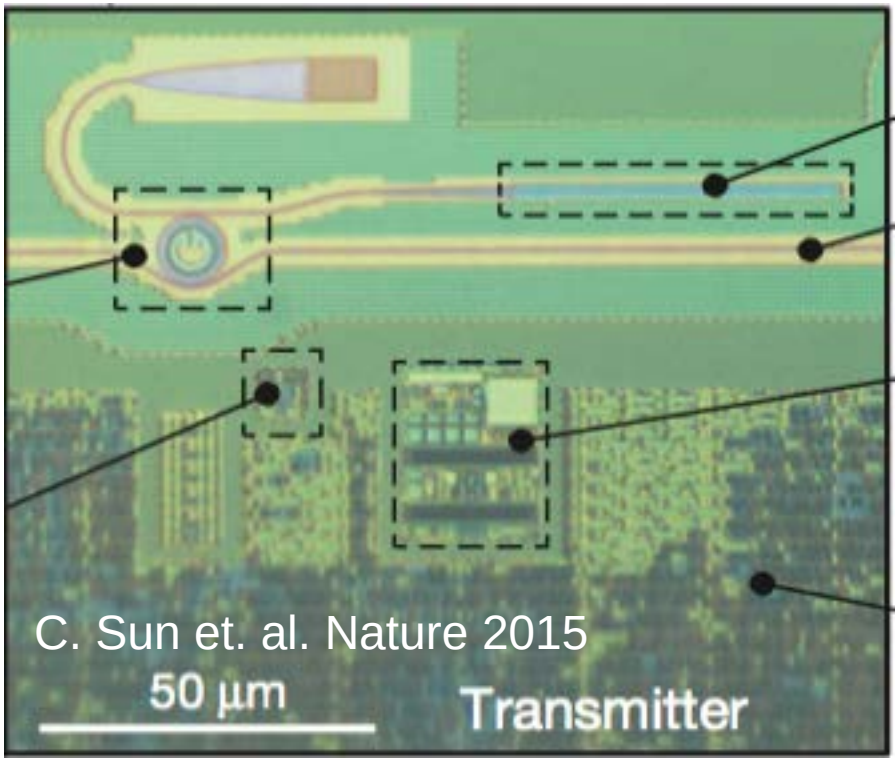


| | Heterogeneous | Monolithic |
|---|---|---|
| Connection yield | * | *** |
| Capacitance | * | *** |
| Cost/circuit density | *** | * |
| Performance | *** | * |

# Si Photonics for DWDM

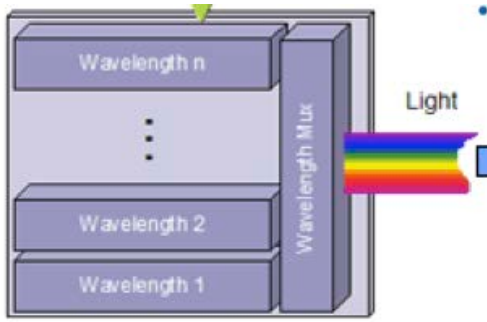


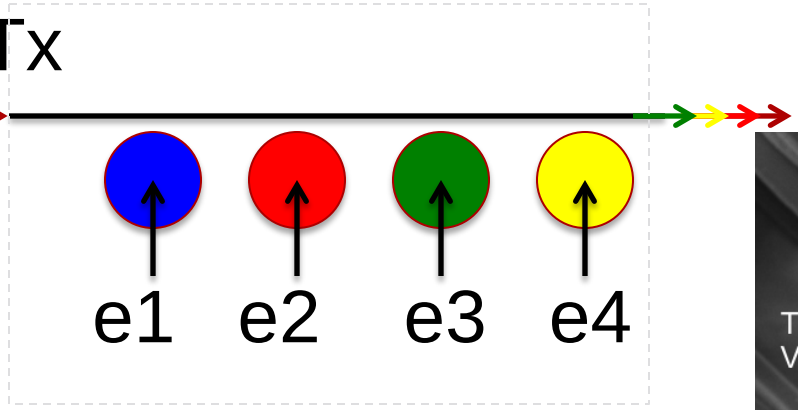


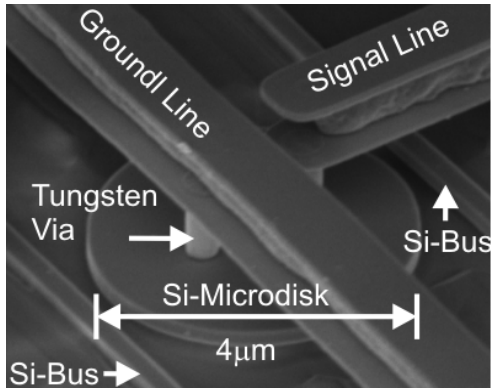


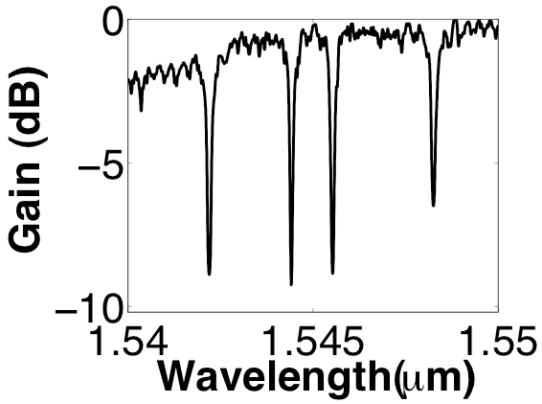


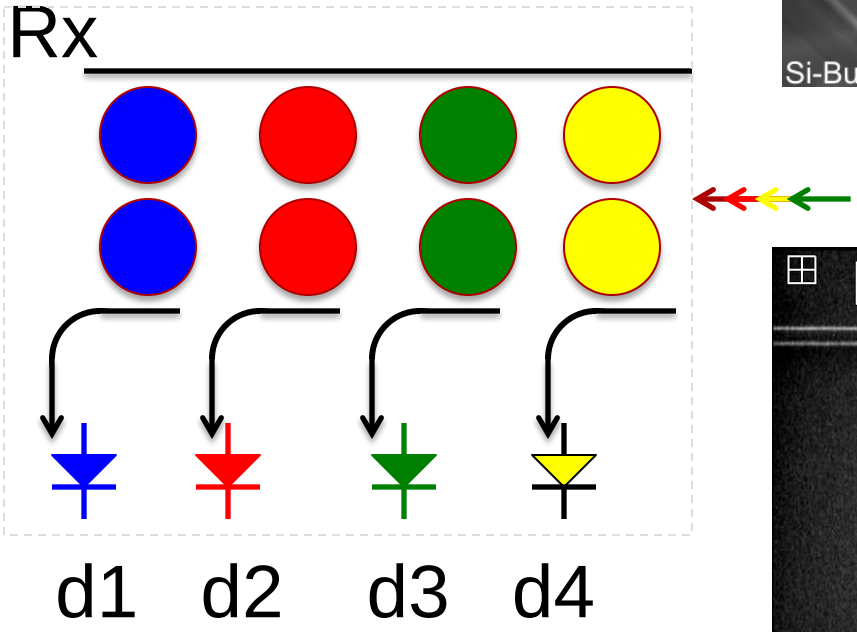


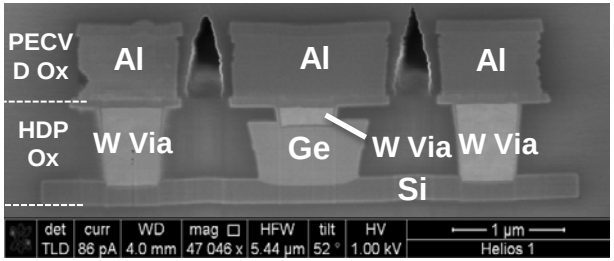


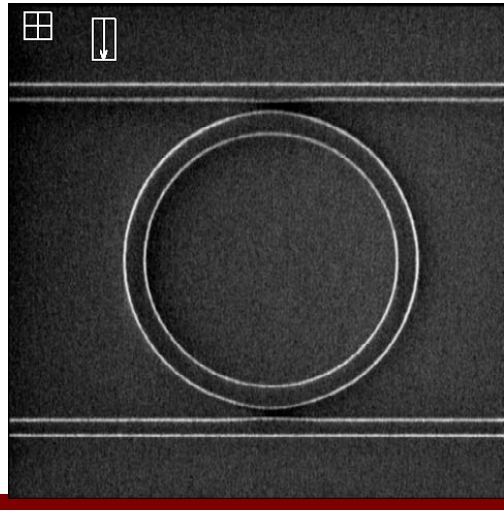

# Silicon Photonics Layer Structure

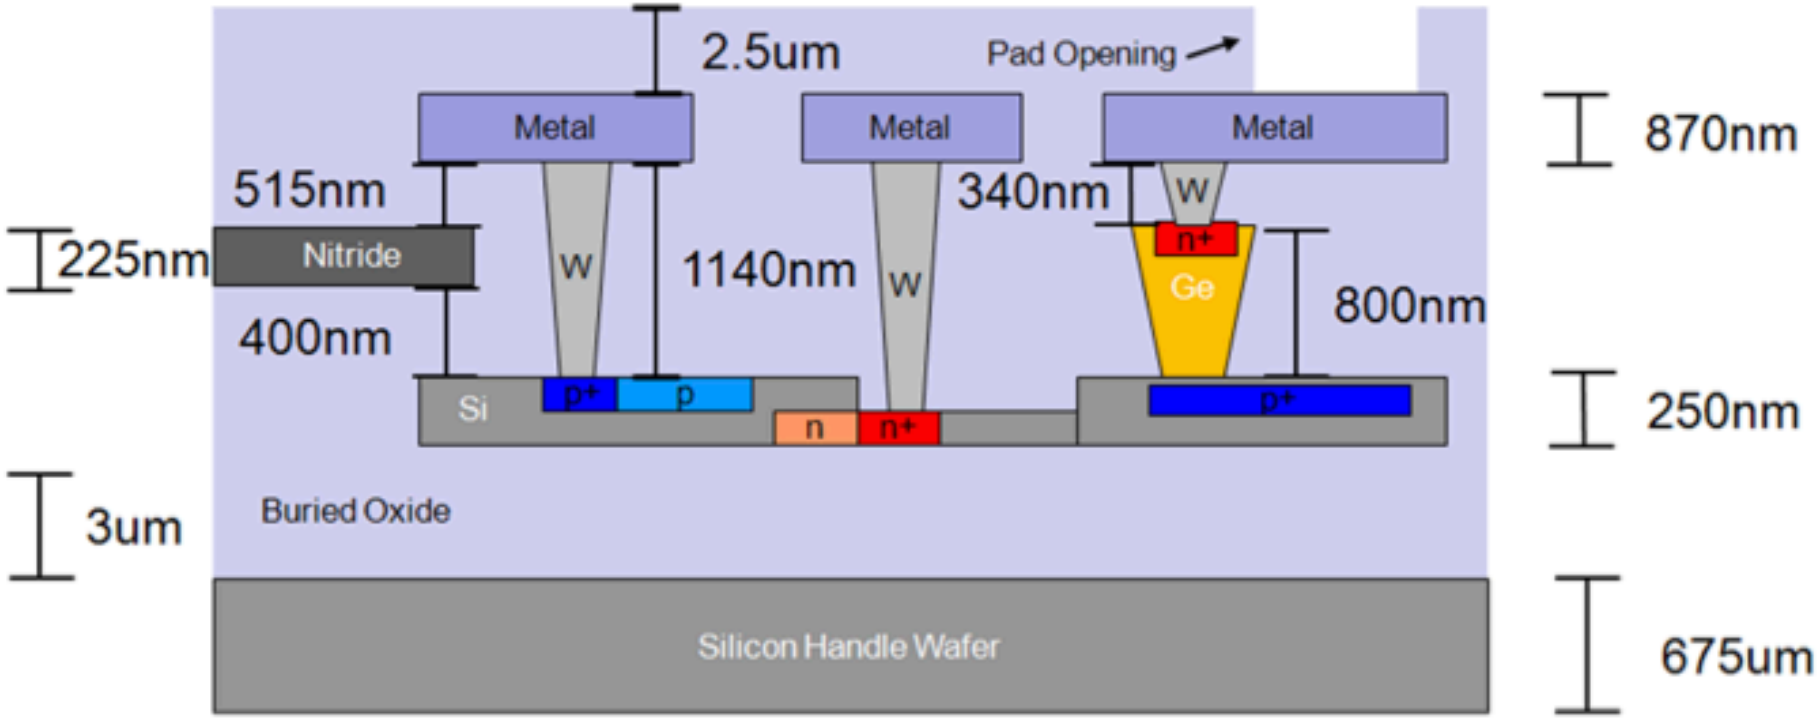


See e. g. A.L. Lentine, C. T. DeRose, P. S. Davids, N. J. D. Martinez, W. A. Zortman, J. A. Cox, A. Jones, D.C. Trotter, A. T. Pomerene, A. L. Starbuck, D. J. Savignon, T. Bauer, M. Wiwi, and P. B. Chu, "Silicon Photonics Platform for National Security Applications, "in 2015 IEEE *Aerospace Conference,* 7-14 March 2015

# Technology Challenges

- Integration
  - ***Silicon photonics integration with state of the art CMOS with low capacitance and high yield***
  - ***Cost effective, reliable packaging***
  - ***Fiber coupling and waveguide losses***
- Silicon Photonics
  - *Efficient Laser source*
  - *Modulator and optical filter resonant wavelength stability and uniformity*
  - *Filter shape, coupling variations*
  - *Low energy receivers*
- Interface Electronics
  - *Efficient clock and data recovery*
  - *Data TDM multiplexing (SERDES)*
  - *Efficient Data encoding and error correction*

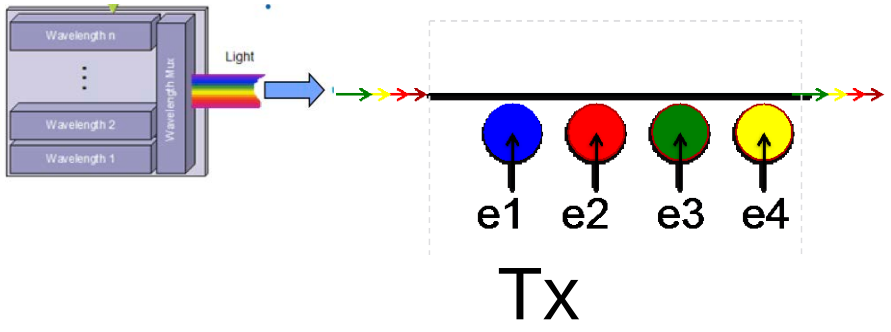


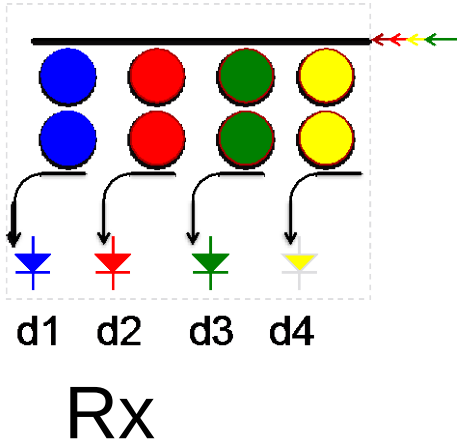

# Laser Source Technology Options

| | Flip | Hetero | Comb |
|---|---|---|---|
| Maturity | *** | ** | * |
| Efficiency | *** | ** | * |
| Size 40λ | * | ** | *** |
| Cost 40λ | * | ** | *** |

*** = best

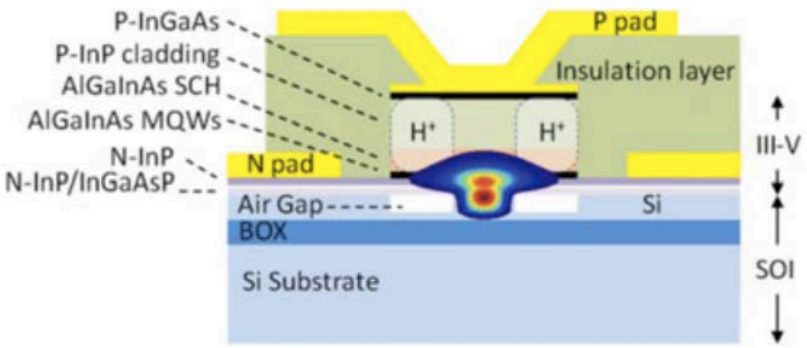


Heterogeneous: Keck et. al., JSTQE 2013

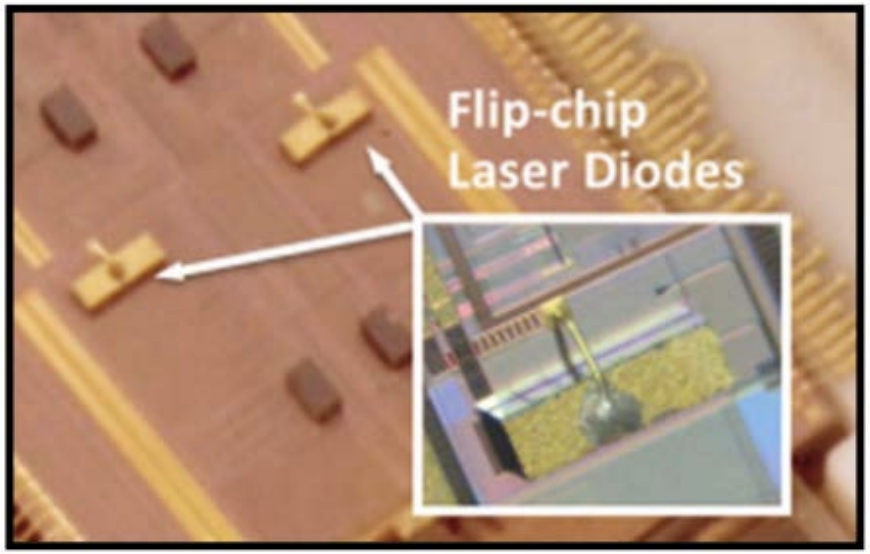


Flip-chip: Dobbelaere ECOC 2014

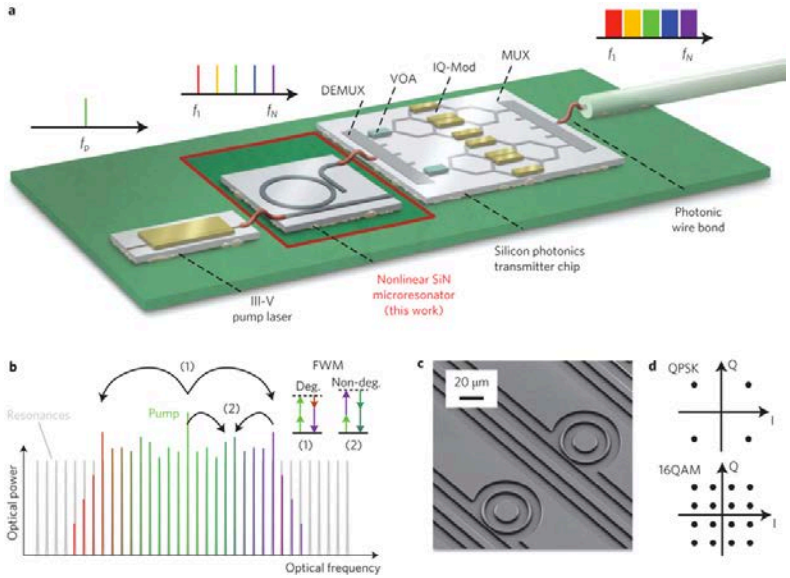


Pfeifle et. al., Nature 8 (2013)

# Simple Modulator Driver: Differential Signaling

- No pre-emphasis
- No AC coupling
- No high voltages
- CMOS logic levels

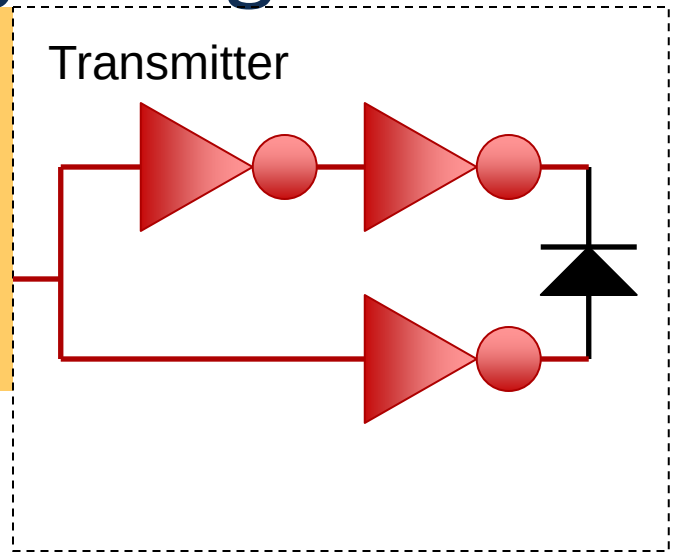


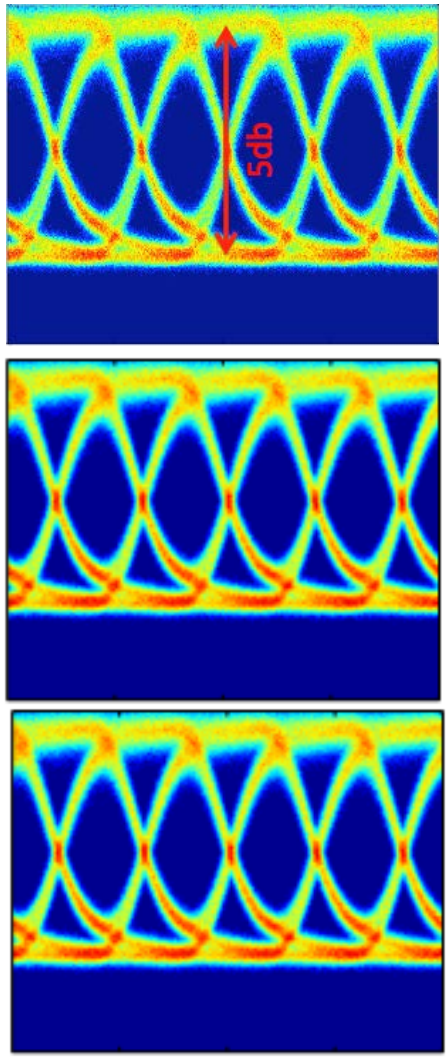


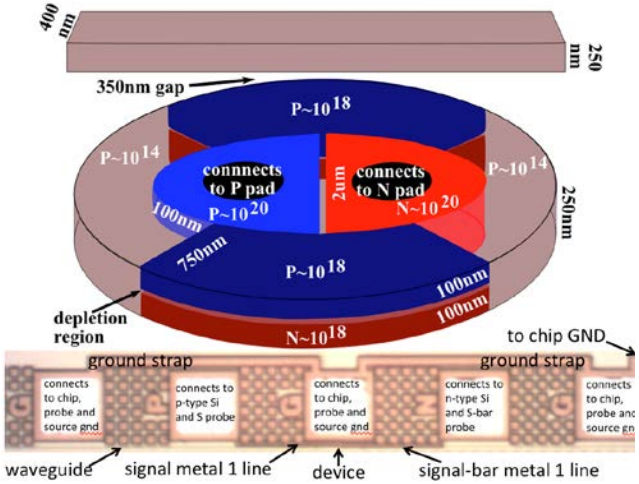


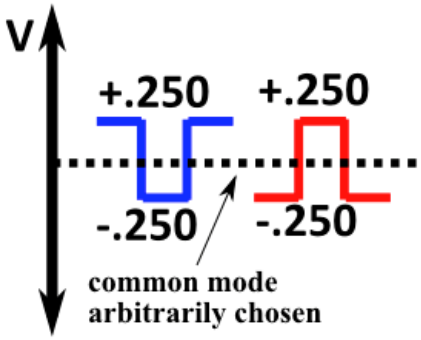


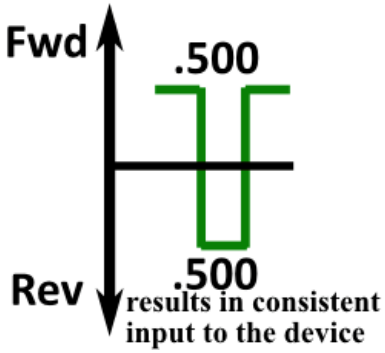


- **10 Gb/s**
- **Common Mode:**
- **.25V, .8V, 1.2V**
- **3 fJ/bit**

W. A. Zortman, A. L. Lentine, D. C. Trotter, and M. R. Watts, 'Low-voltage differentially-signaled modulators,' Opt. Express **19**, 26017-26026 (2011)

# High Transimpedance Receivers

- No AC coupling
- No data encoding
- Low input capacitance
- High trans-impedance
- Low delay

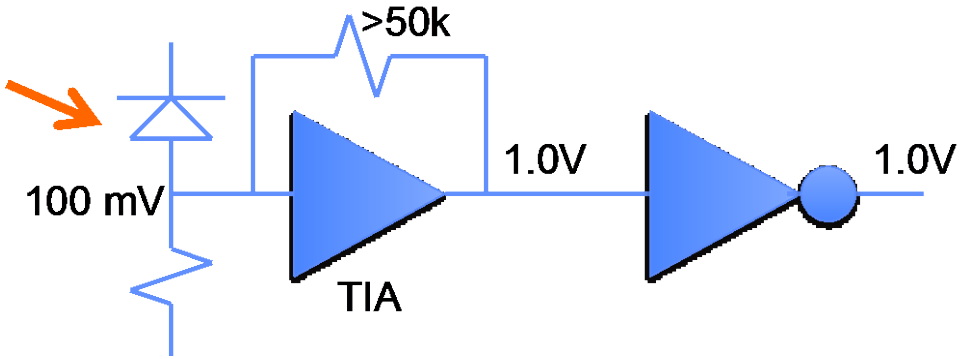


Low capacitance, high transimpedance gain
Lower noise floor, good sensitivity
Might be DC offset limited vs. noise limited

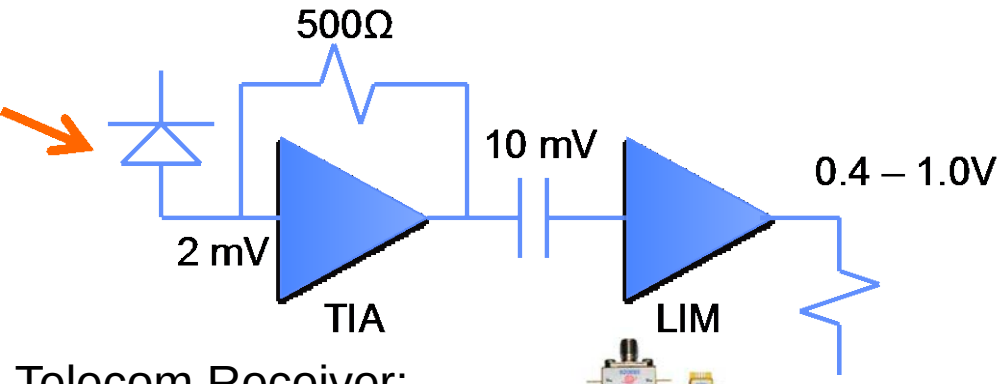


Telecom Receiver:
Multi-mode data com Rx:
Too much power, delay

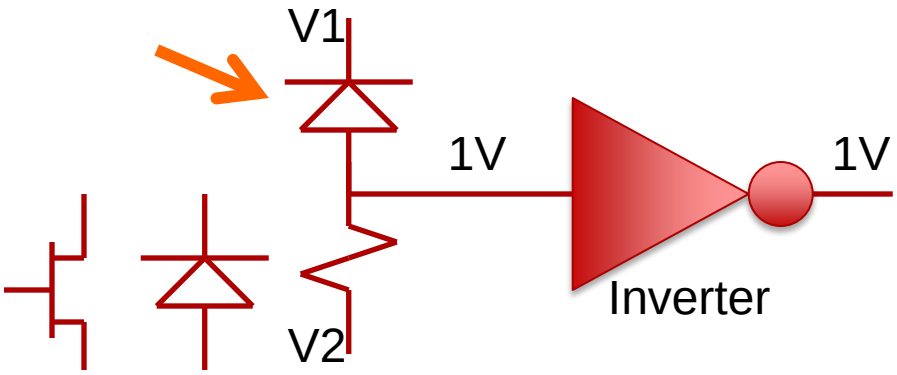


Very low energy, poor sensitivity

D. A. B. Miller et. al., PTL 1989
L. M. F. Chirovsky et. al., IEEE Int. Opt. (1994)
C. Debaes, et. al., in IEEE JSTQE (2003)

# Receiver Energies (Simulation)

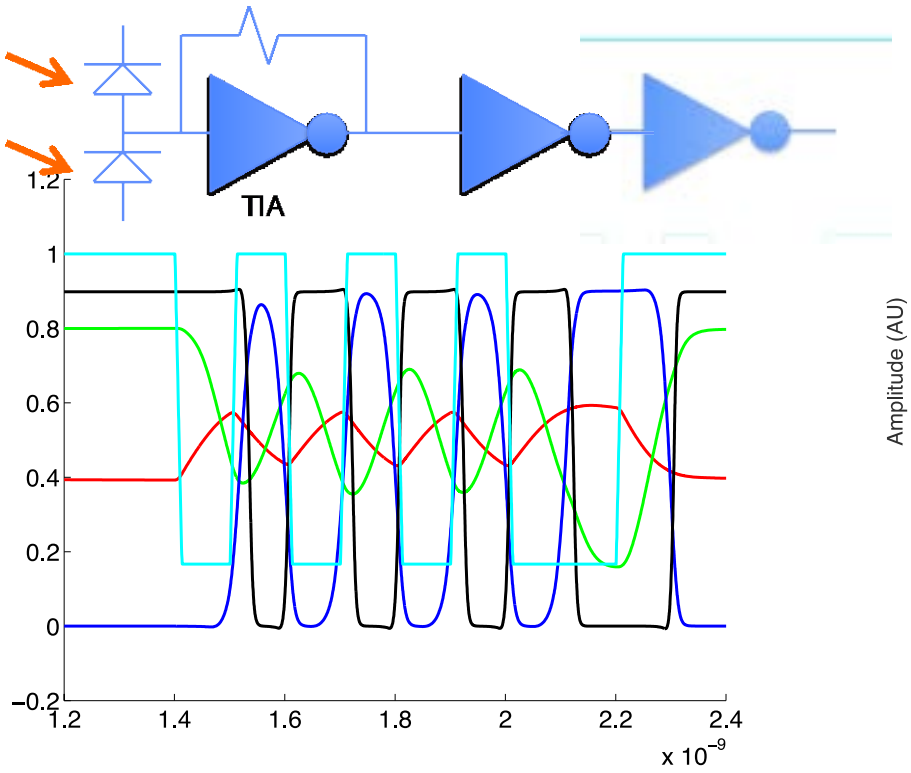


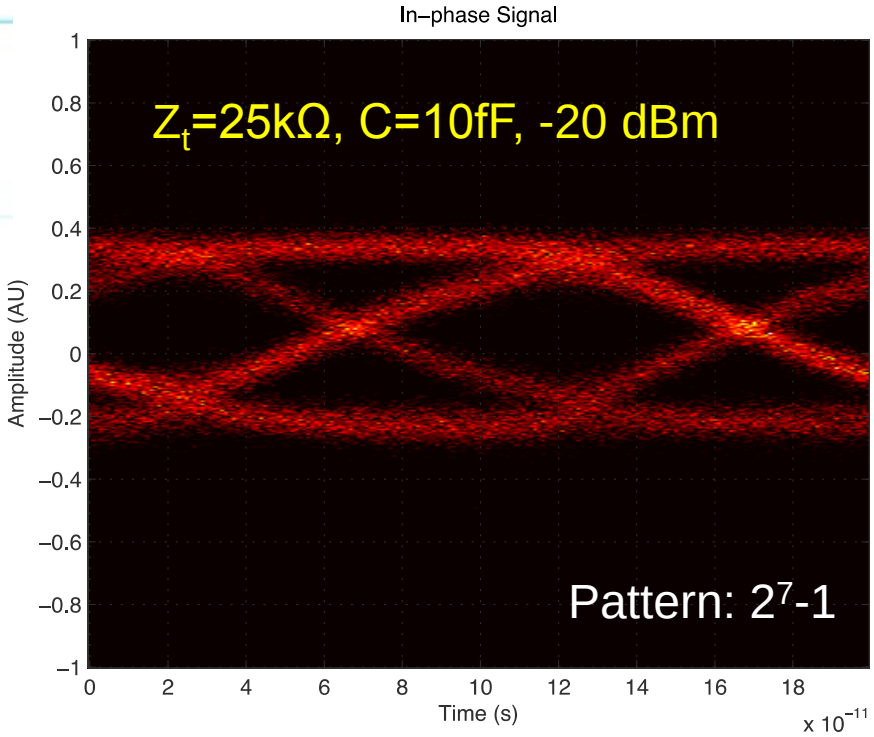


Bit Rate = 10.00Gbps, *Energy per bit = 2.74fJ*

- 45 nm technology node,
- optimum sample delay = 109 ps, total margin= 70ps,
- BER = 0.000000e+00,

# Noise limits: receiver sensitivity

- Ultra-low capacitance allows drastic increase in Zt and Rx sensitivity
- Not yet realized (why?)
- For integrated Rx:
  - achieve noise limited performance vs. DC-offset, power supply noise, etc.
- APDs should help even more
  - Optimal circuits?

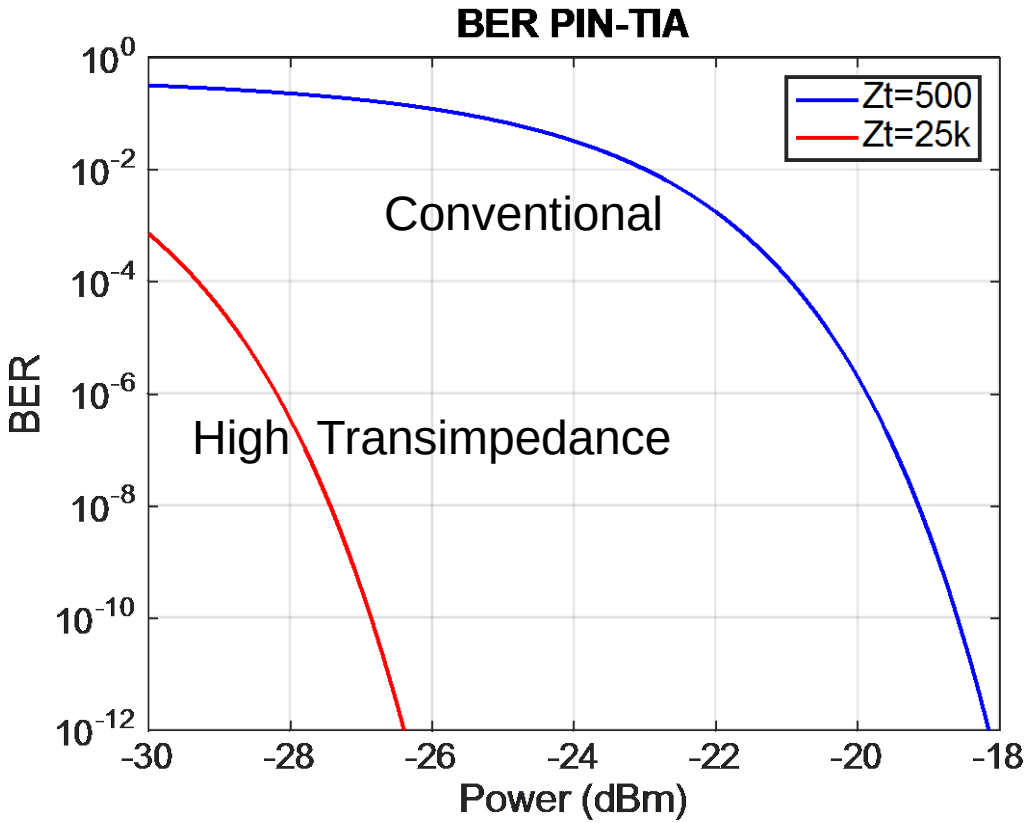


Noise Limited Performance

# Beyond Moore's Law transistors

- Examples of BMC transistors
  - TFETs,
  - SETs,
  - Superconducting J-FETs
  - …
- Characteristics of BMC transistors
  - Lower voltage
  - Poor drive
  - Sometimes cooled
  - Non-CMOS (2-terminal)
- Digital computing (above)
- Analog computing
- Optical computing
- Reversible computing
- Quantum computing

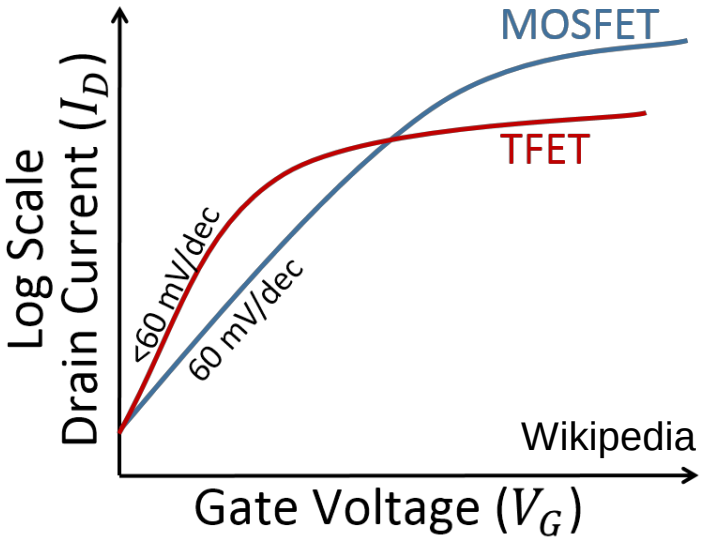


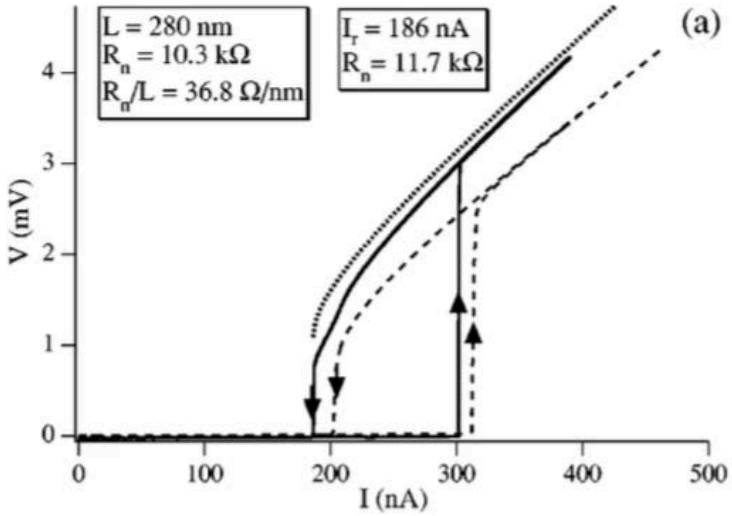


M. Tinkham, PHYSICAL REVIEW B **68**, 134515 2003

# Si Photonics for DWDM: Link budget

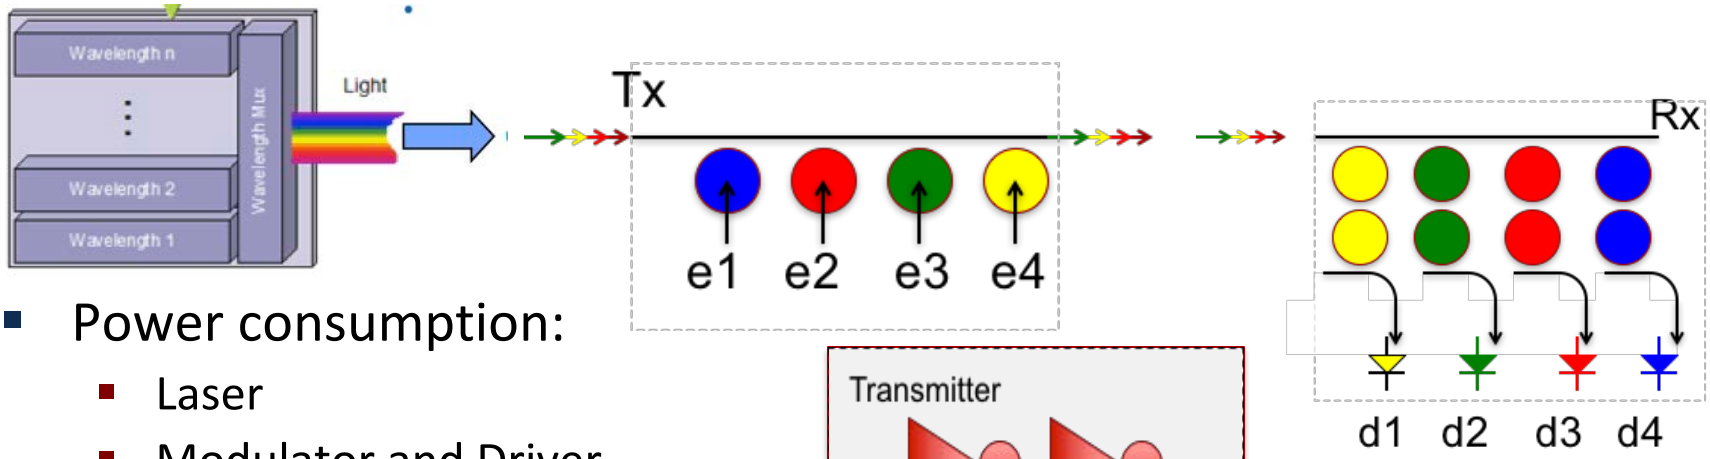


- Power consumption:
  - Laser
  - Modulator and Driver
  - Photodiode and Receiver
- Calculate receiver sensitivity[1]
  - Estimate loss in the system
    - Fiber coupling
    - Modulator (optimize difference in outputs) [2]
      - Modulation (extinction)
      - Loss
    - Waveguides
    - Photodiode responsivity
  - Work backwards to calculate laser optical power
  - Efficiency to calculate electrical laser power (dominant term)

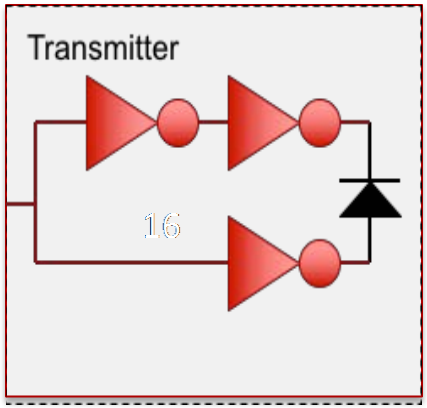


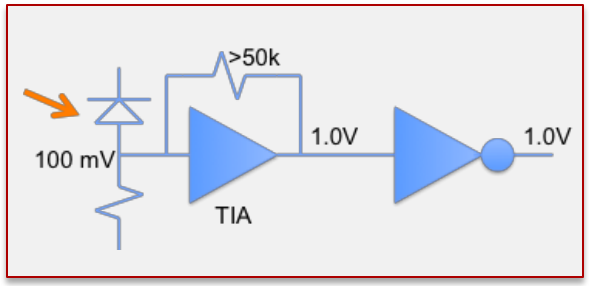


- [1] ref: Agrawal
- Bitrate → TIA gain
- → required Pin
- → noise
- → iterate for SNR

[2] ignore tuning

# Noise floor energy limits

- At very low capacitance shot noise limits the performance, not thermal noise of the feedback resistor

- Lower voltage on Rx means lower laser power.
  - Receiver easier
  - Modulator harder;

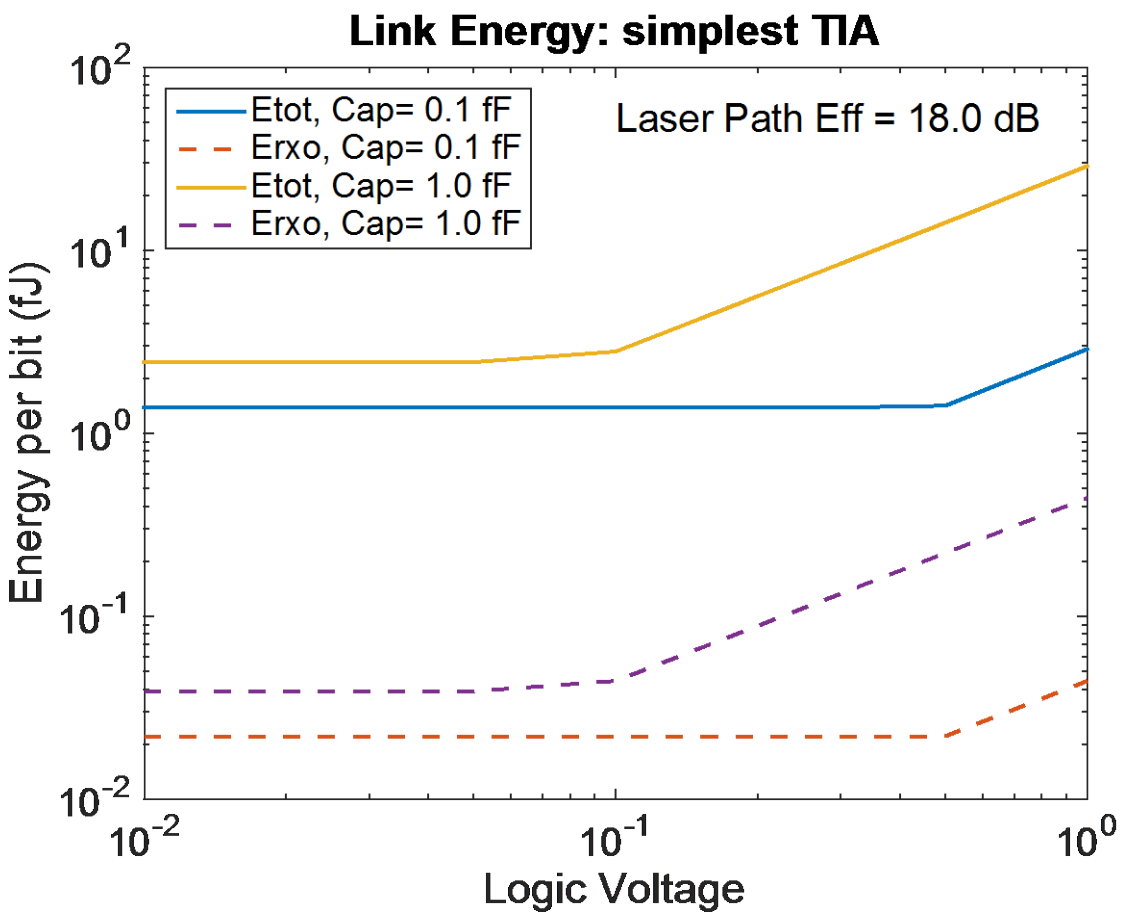

# Optical modulators

- Electro-absorption in quantum wells, III-V SiGe (QCSE)
  - Strong effect
  - Difficult integration with silicon, but possible
- Electro-optic (polymers, sol-gen, Lithium Niobate, III-V, etc.)
  - Stronger than silicon, but still weak
  - Difficult integration with silicon, but possible
- Plasma carrier effect (in silicon)
  - Weak effect, resonant or interferometric devices enhance effect
    - Manage fabrication and environmental variations
  - Stronger effect in some materials (ITO, Graphene, etc.)
  - Simple integration, especially the native silicon ones.
  - Very broadband effect (100s of nm)

# Resonant silicon micro-photonics

- Why resonant silicon photonics?
  - Small size (<4 um dia.)
  - Resonant frequency → DWDM modulators & mux/demux
- Benefits
  - Low energy
  - High bandwidth density
- Resonant Variations
  - Manufacturing Variations
  - Temperature Variations
  - Optical Power (1s density)
  - Aging?
- Requirements:
  - Resolution: +/- 0.25 ° C (depending)
  - Range: 10 – 85 ° C (depending)

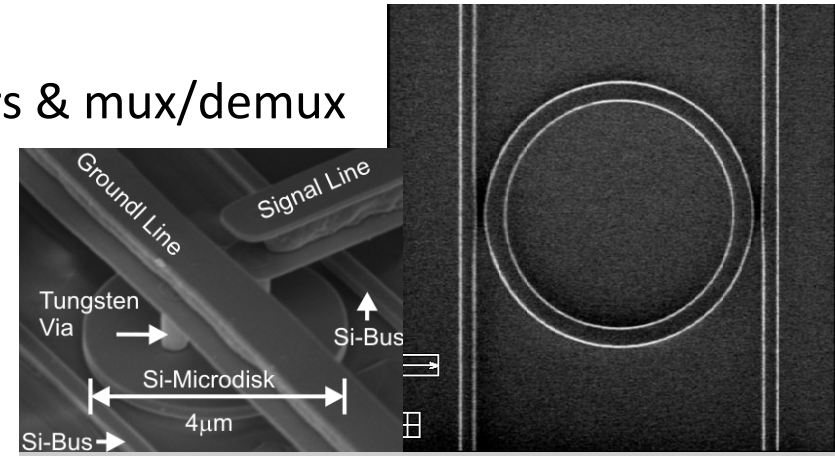


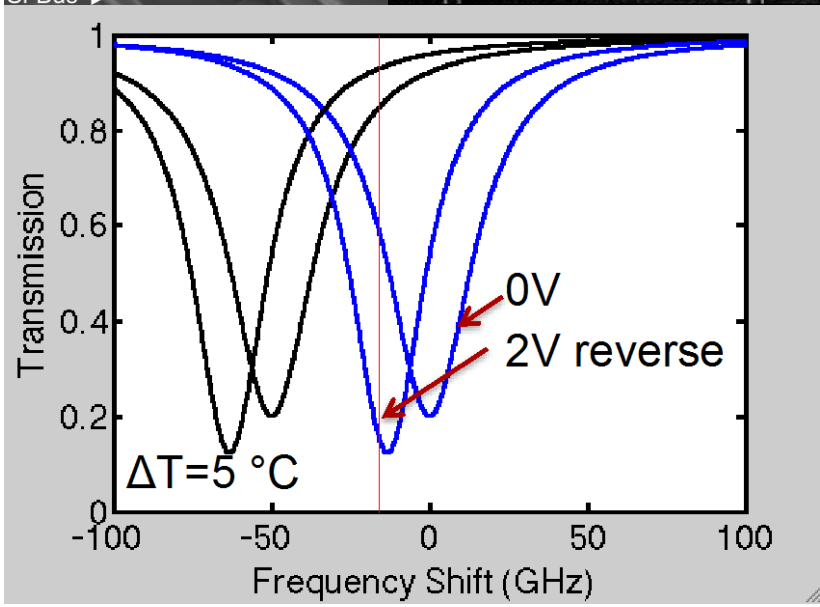

# Silicon Photonics modulator model

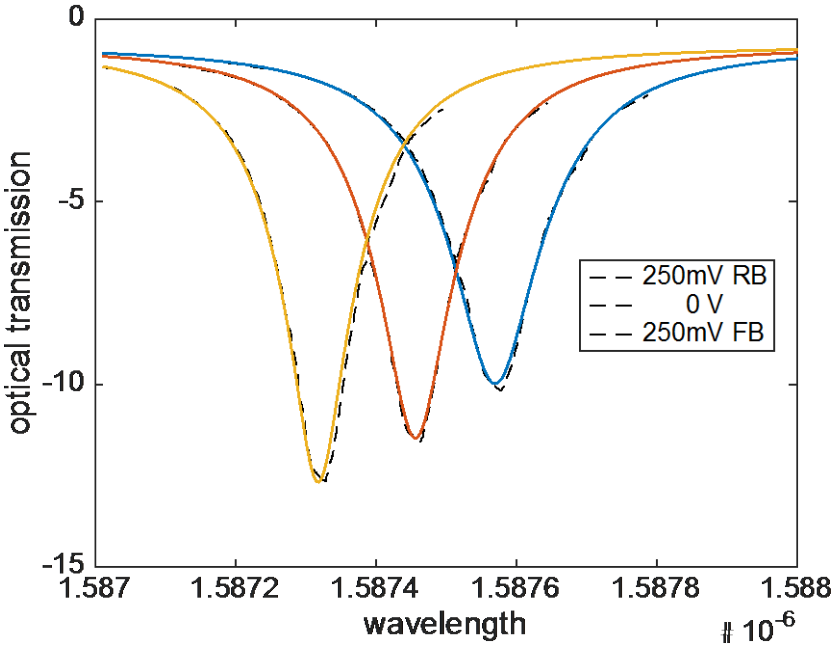


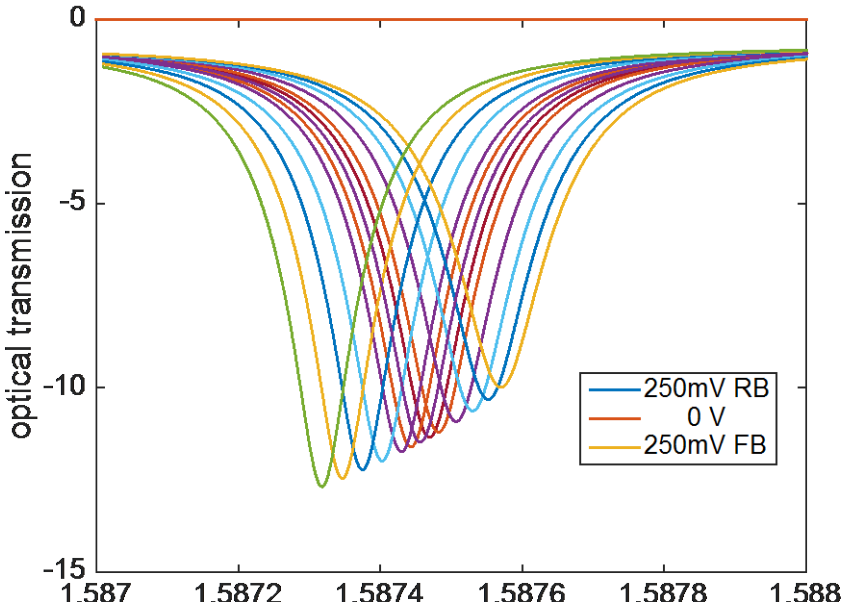


- Best published result has +/- 0.25V swing for optimal performance [1]
- At reduced voltage, there is reduced optical transmission
- Have to compensate with more power, more noise, more power.

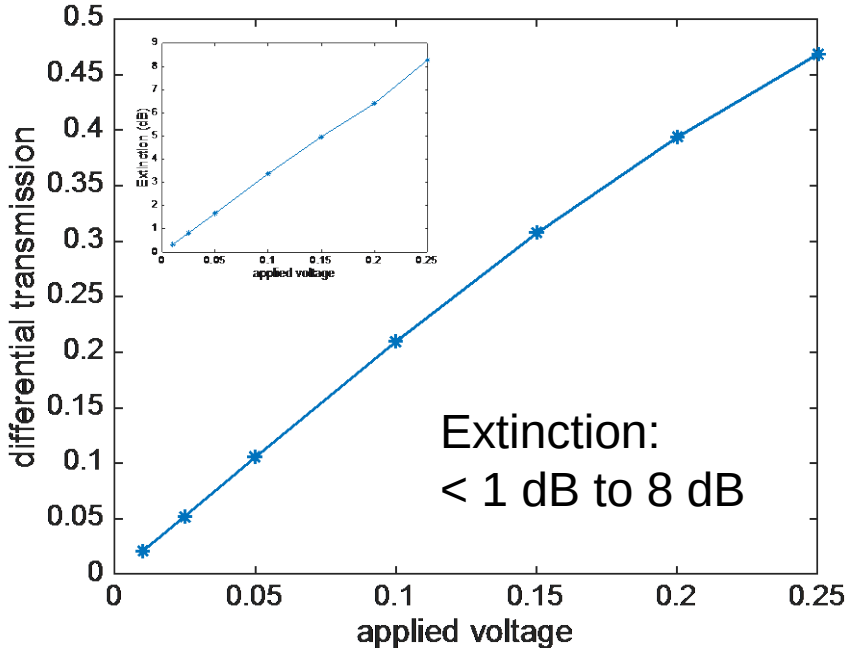


[1] Timurdogan et. al., Nature Com. 2014

# Results from an ‘actual’ modulator

- Very poor extinction at low voltages
- Increased power to compensate
- Increased shot noise
- Further increased power to compensate
- Increased energy per bit

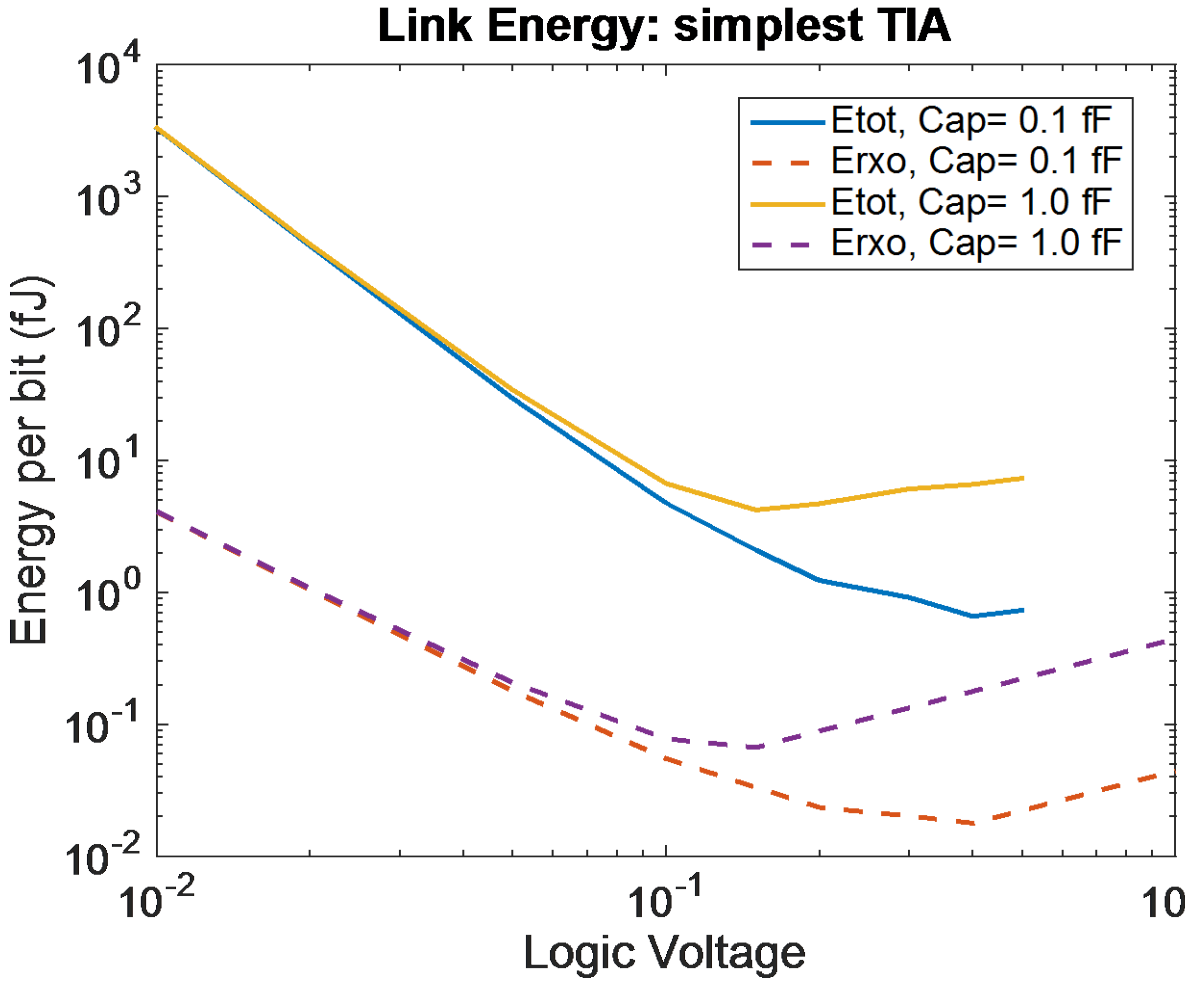


Message: BMC computing devices will need new optical technology

# Analog / multilevel interconnect

- Signal decreased in amplitude by 2X for every bit.
  - 64X @ 8bits
- Rest is shot noise

- This is interconnect
- Analog computing is different

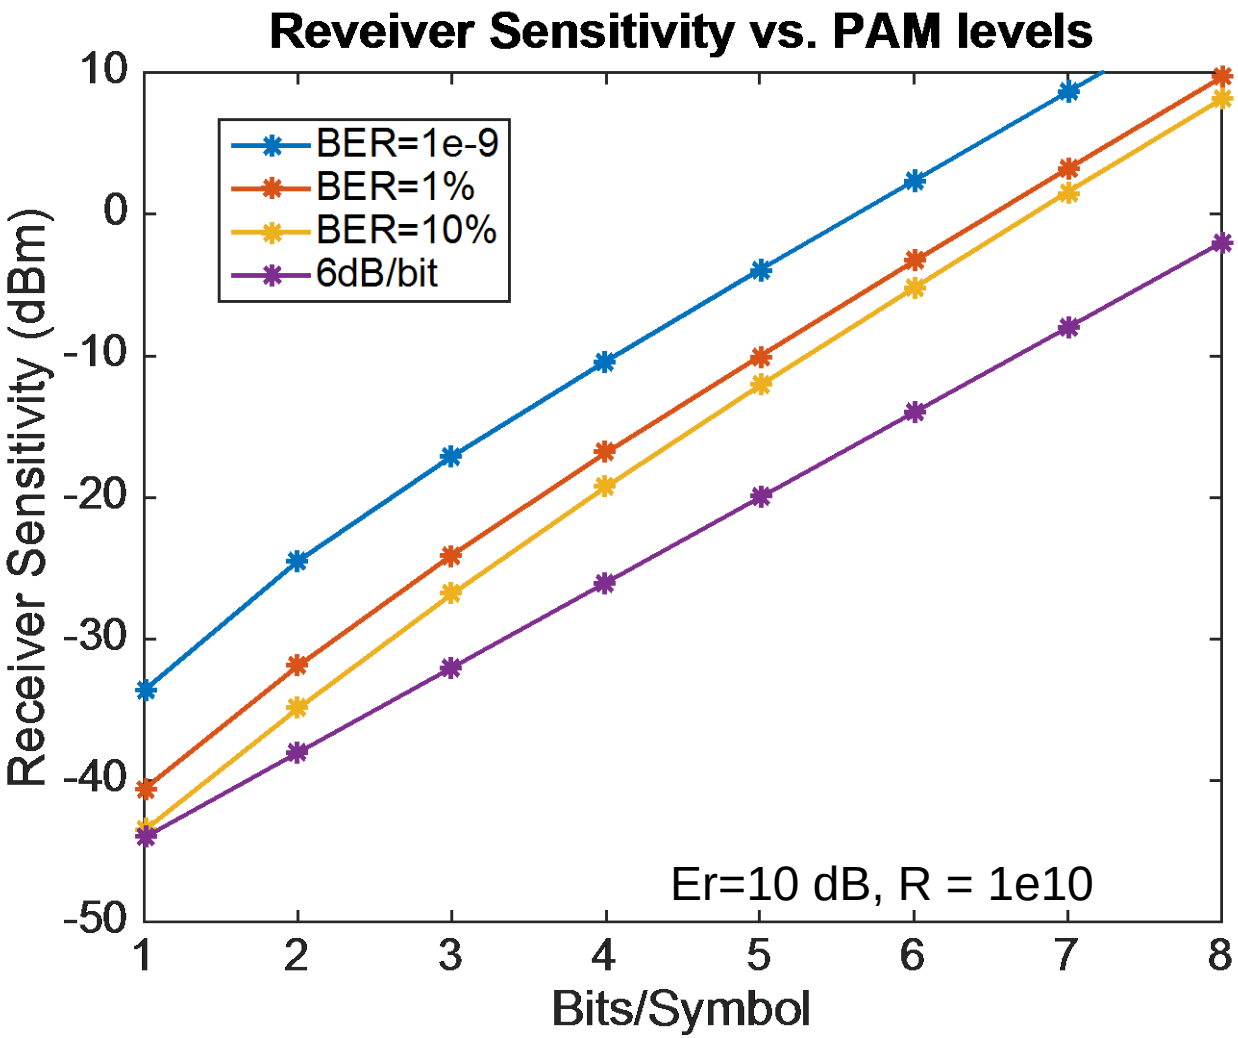

# Analog / multilevel interconnect

- Signal decreased in amplitude by 2X for every bit.
- Signal levels higher for 1/2 the bits, so more shot noise
- Need good linearity or compensation
- Best performance with non-equal levels and adaptive thresholds
- Analysis that follows is for an interconnect
- Analog optical computing may be different

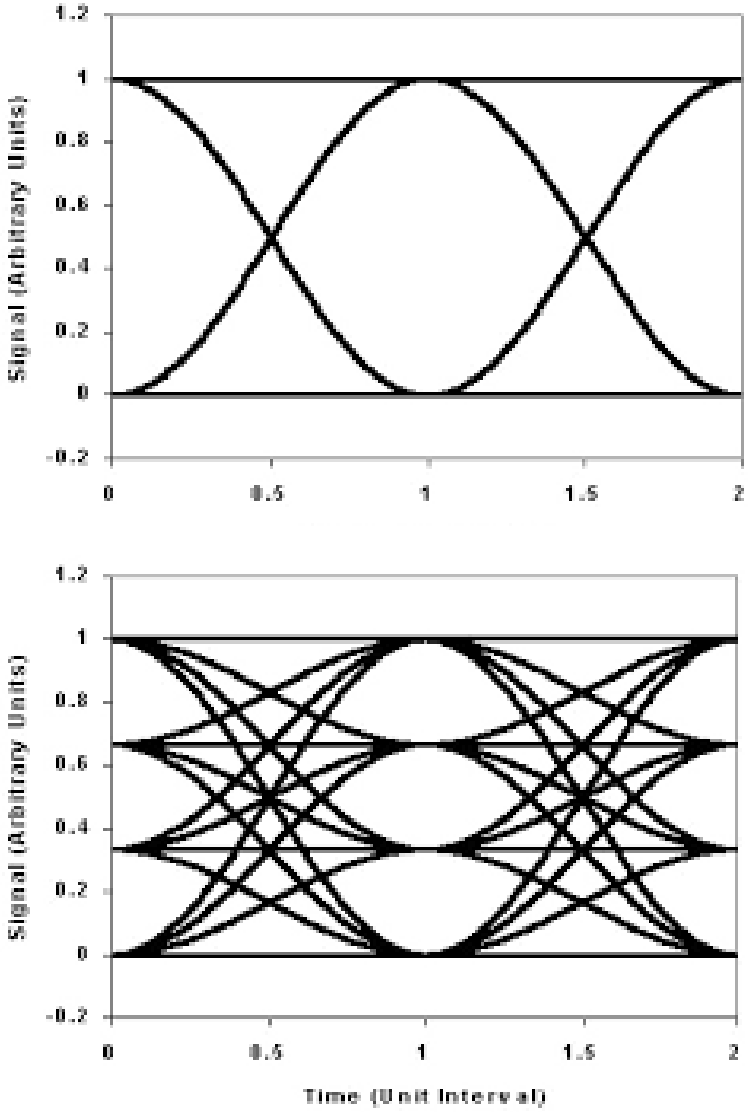

# Analog / multilevel interconnect

- Signal decreased in amplitude by 2X for every bit.
  - 64X @ 8bits
- Rest is shot noise
- This is interconnect
- Analog computing is different

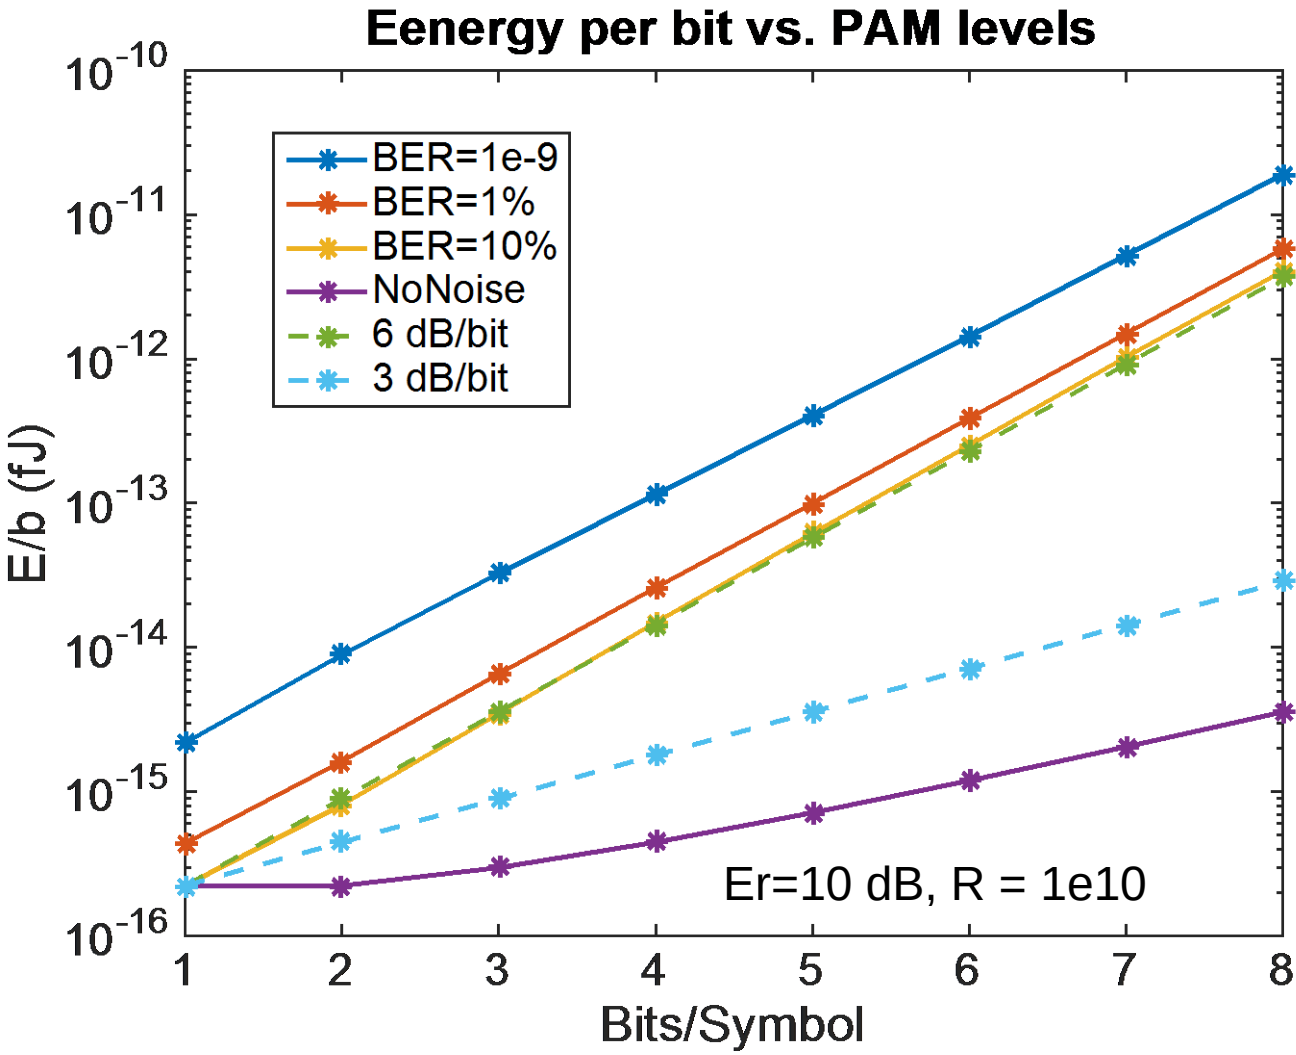

# Summary

Sandia National Laboratories

- Integrated photonics can drastically improve power consumption and bandwidth performance CMOS-based systems:
  - Need intimate integration with high-value electronics
  - Need low cost communications infrastructure (not fibers)
  - Lots and lots of challenges to getting this working (year > 2026, > $1B)

- BMC devices present additional challenges
  - Shot noise floor limits energy consumption to a fraction of 1 fJ/bit.
    - Laser efficiency, optical losses, and finite extinction >> 20 photons/bit.
  - Ultimately need better modulators for BMC low voltage, low current devices
    - Today's best devices are better suited for end of road CMOS (0.25 – 0.5V)
    - Combined optical/energy of system → even today's optics may help.
  - More work for optics to interface to non-CMOS-type logic
  - Analog interconnects maybe problematic because of SNR considerations.

# Questions?

Sandia National Laboratories

**Anthony L. Lentine**
**Sandia National Laboratories, Albuquerque NM**
**alentine@sandia.gov**

# Common CDR among 40 λ

Sandia National Laboratories

1. Optical delay variation vs. λ is really small (~ ps)
2. Rx delay vs. optical power is likely dominant
3. FET variations ?

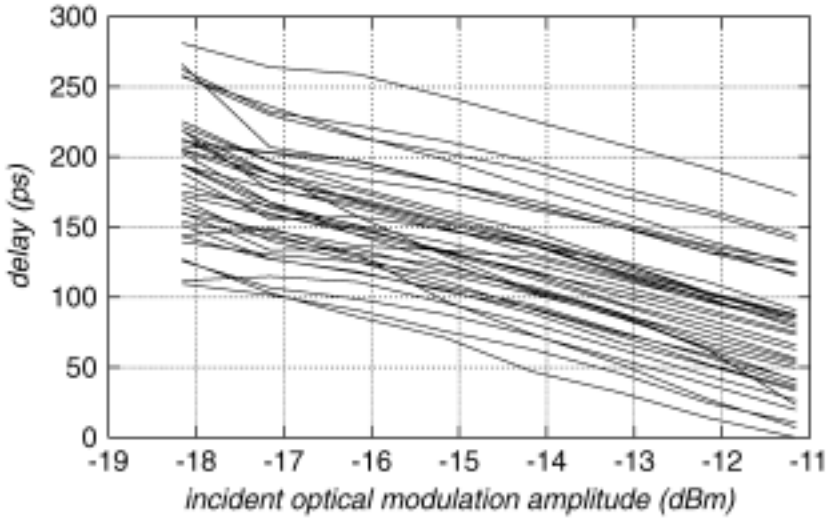


Wilde, Rits, Baets, Van Campenout
**64 ch. VCSEL Links**, JLT 2008

- **Silicon Photonics vs. VCSELs**
- High-Zt Rx – way less complex,
  - less delay variation
- Modulator vs. VCSEL
  - More likely to have uniform characteristics (drivee, vs. driver)

A similar tolerancing analysis needs to be done for DWDM silicon photonics links

# SERDES (Trends)

- Serdes must contain:
  - Mux/Demux
  - Clock multipliers and dividers
  - Receiver re-timing (phase alignment)
  - Data ordering
  - These are not necessarily inherently power hungry

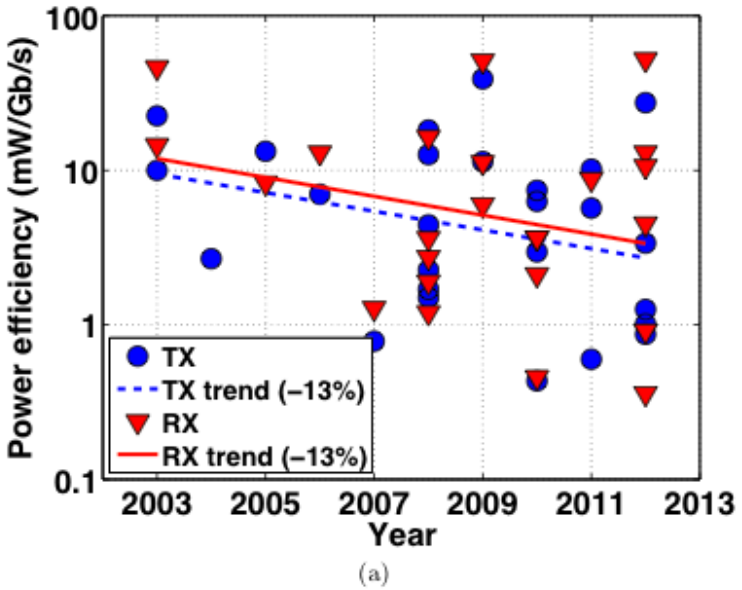


Arash Zargaran-Yazd, PhD thesis UBC

- ‘Un-necessary’ features ?
  - Electrical Line Driver (Pre-emphasis/Equalization)
  - Variable line rate
  - Clock and data recovery
  - Coding and decoding
  - Diagnostics

- Ave: 6 pJ/bit 2012 (-13%/yr)
  - 300 fJ/bit Rx (best)
  - 500 fJ/bit Tx (best)
- 0.25X in 10 years!
- Why such a spread?

Why can’t a SERDES be designed in the < 100 fJ/bit vs. ~ 1 pJ/bit?

# Near Term Packaging Costs

Sandia National Laboratories

- Optical packaging independent of the number of wavelengths
  - If we develop low cost DWDM lasers
  - If we develop robust, low overhead resonator stabilization circuits
  - If we solve other minor issues, filter shape, improve Rx sensitivity, etc.
  - Amortize the fiber connection cost by putting more data per connection
  - Cost (40λ) = Cost (1λ) for optical packaging.
- Electrical packaging independent of the number of optical λ for a given total bandwidth
  - Future 1 Tb/s ≠ today 1 Gb/s because the electrical IO is challenging.
- Need *low cost* optical package with many high speed electrical IO
- There's a lot of good work to do just that
  - But in the long run …

# Technology Challenges

- Integration
  - *Silicon photonics integration with state of the art CMOS with low capacitance and high yield*
  - *Cost effective, reliable packaging*
  - *Fiber coupling and waveguide losses*
- Silicon Photonics
  - *Efficient Laser source*
  - *Modulator and optical filter resonant wavelength stability and uniformity*
  - *Filter shape, coupling variations*
  - *Low energy receivers*
- Interface Electronics
  - ***Efficient clock and data recovery***
  - ***Data TDM multiplexing (SERDES)***
  - ***Efficient Data encoding and error correction***

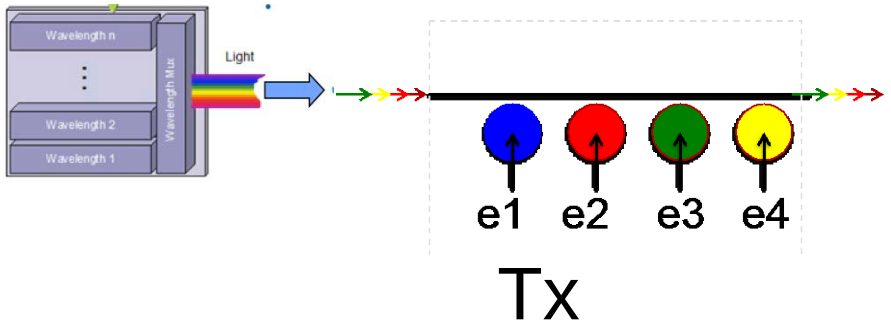


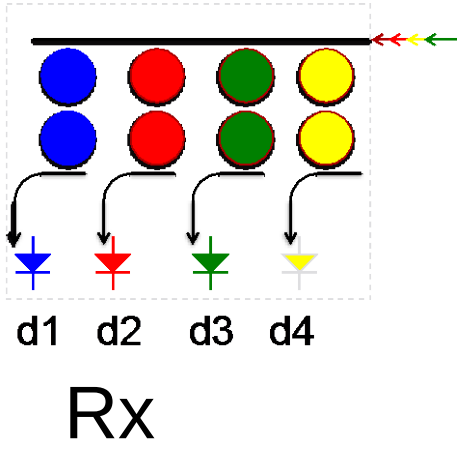

# Laser Source Architecture

- For a transceiver, ***everyone*** wants the source in the package!

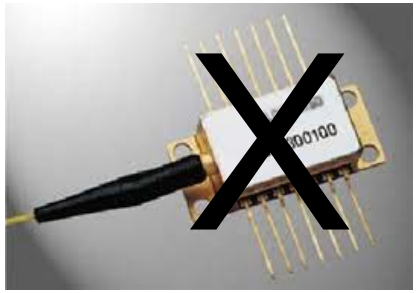

- In the *long term*, routing the laser in an optical waveguide from an on-board laser is analogous to routing electrical power.

http://www.izm.fraunhofer.de

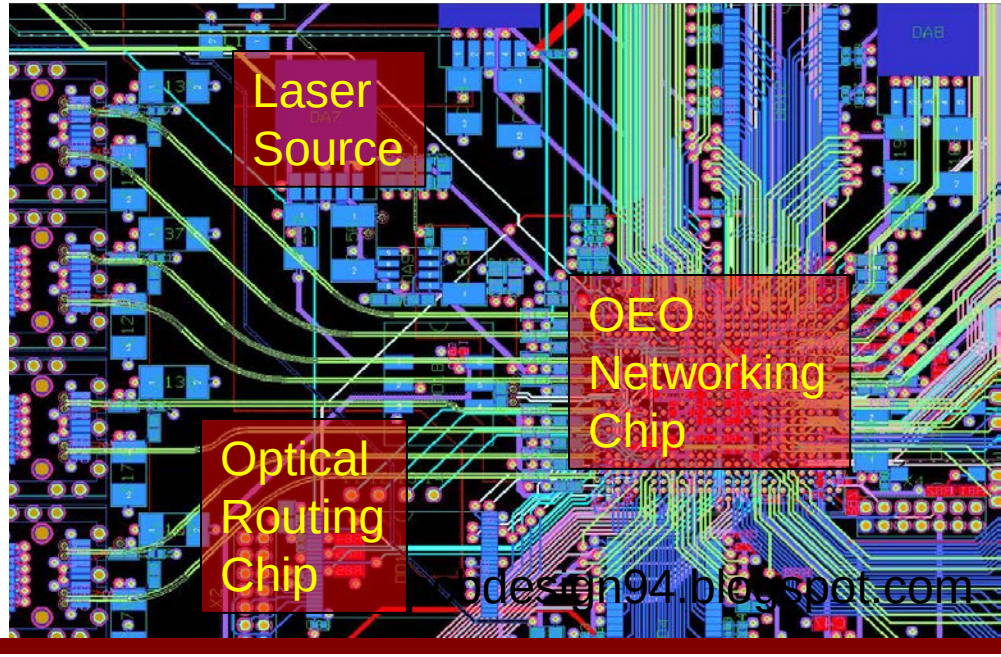

# Computing: can optics keep up?

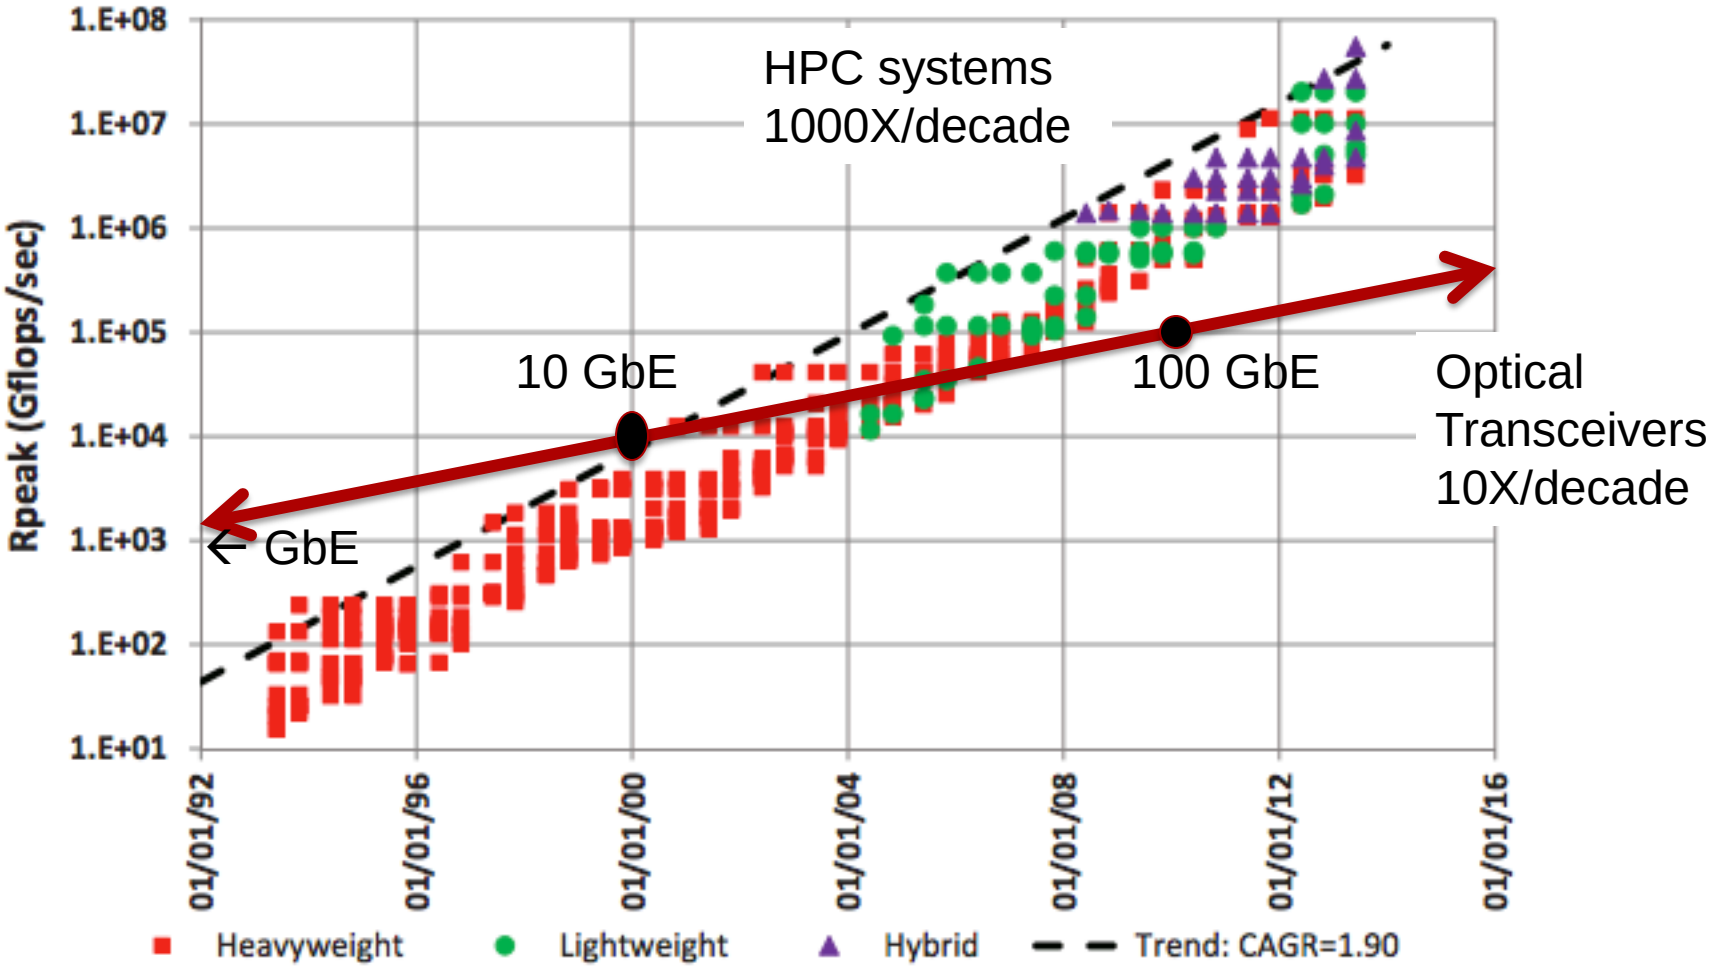


Koegge and Resnick, SANDIA Report 2013

# Can optics keep real computing on track?

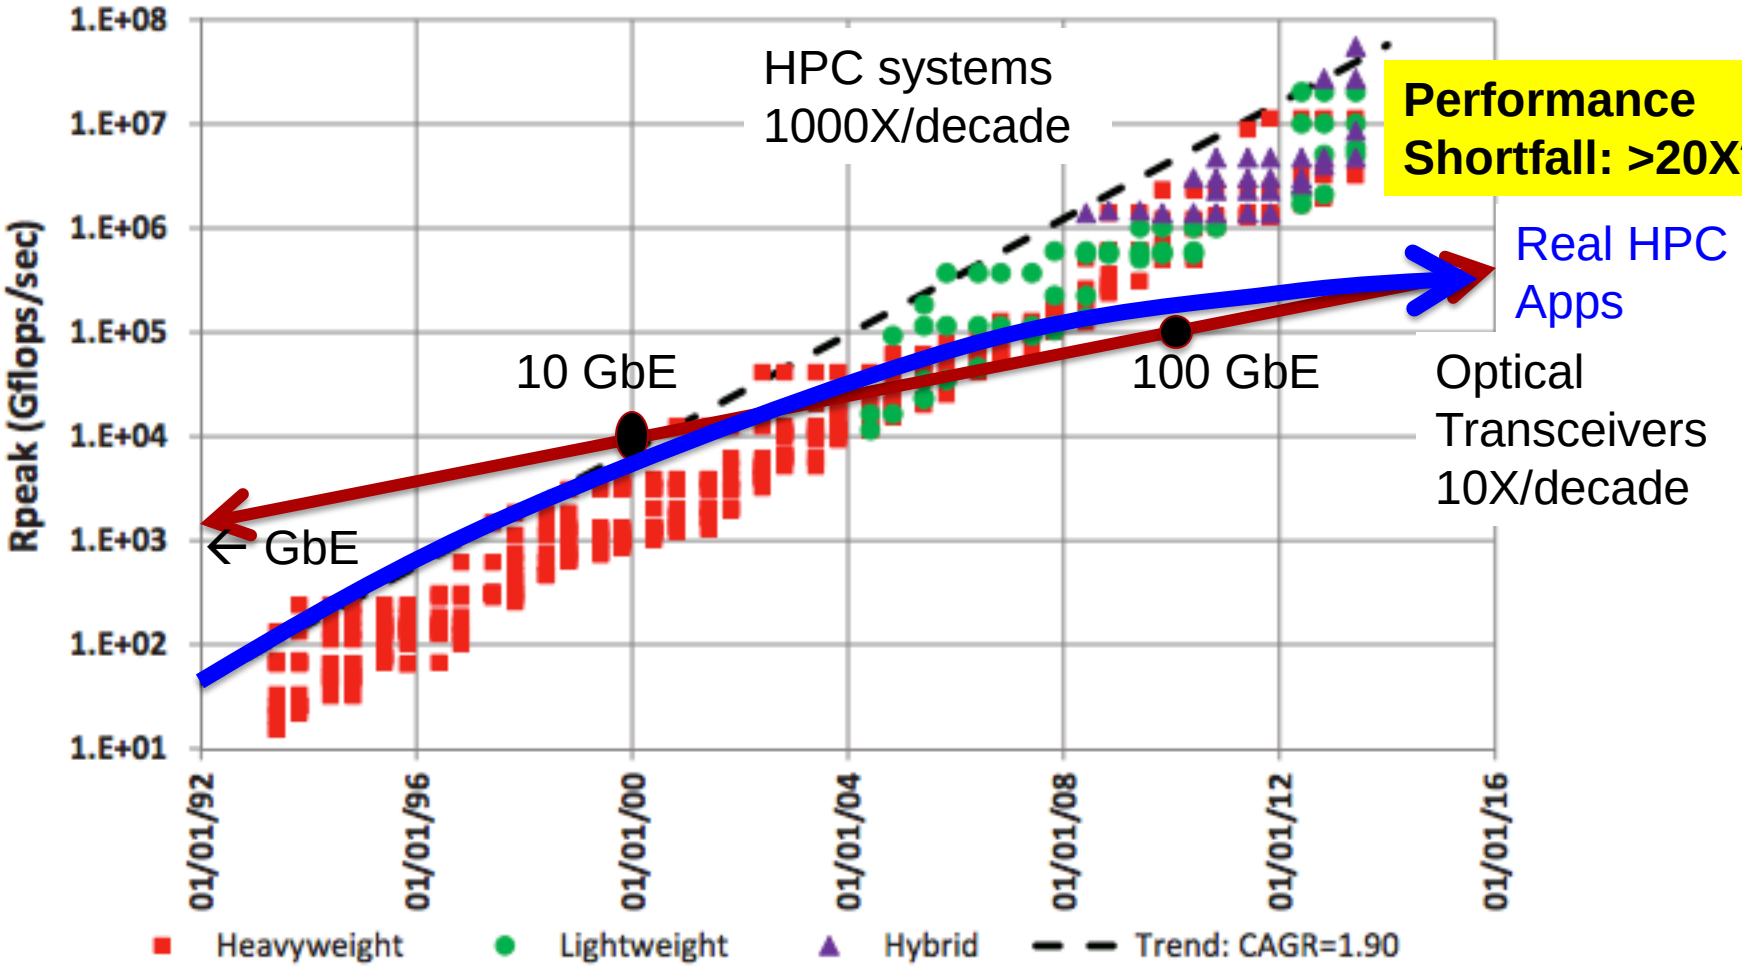


Koegge and Resnick, SANDIA Report 2013, *John Shalf, private communications

# Si Photonics 2 x 2 WSS

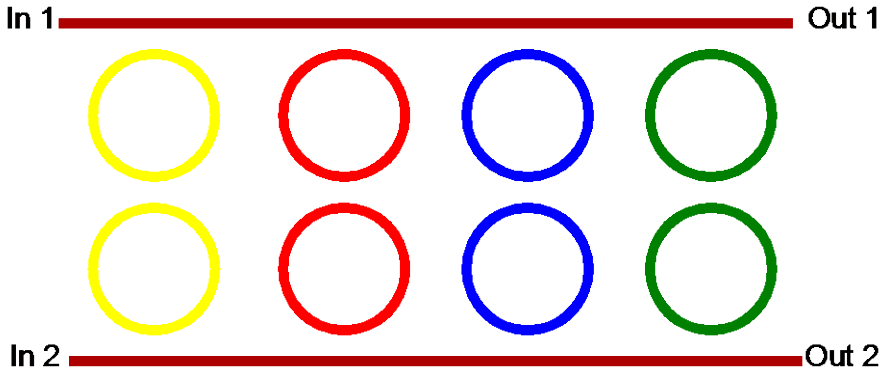


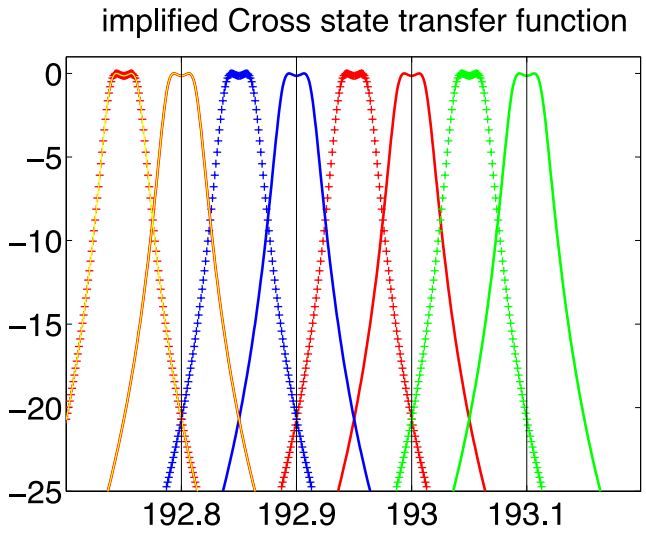


**LONG TERM (IDEAL) SPECS:**

- Ultimate Switch time < 25 ps
- Loss (cross state) 1 – 2 dB
- Loss (bar state) < 0.2 dB
- Crosstalk (15 – 30+ dB)
- Resonant wavelength stabilization

- Ring Size ~ 4 - 6 um
- Coupling gaps ~ 200 - 500 nm
- Ring to ring spacing ~ 4 – 6 um*
- Size < 12 um × λ × 10 um.

# Wavelength switching networks

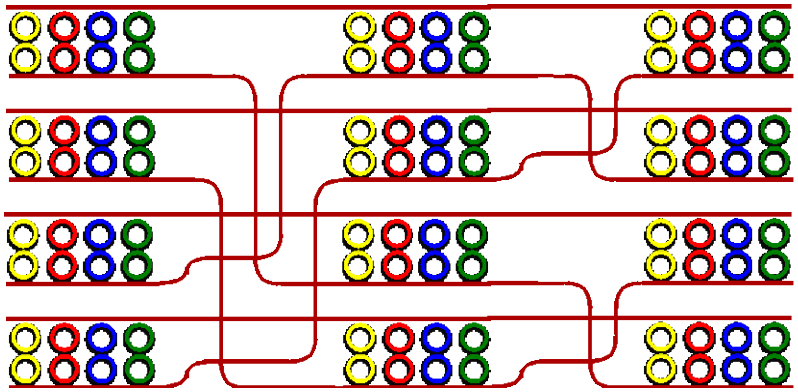

...

Chip scale 256 x 256 @ 32λ

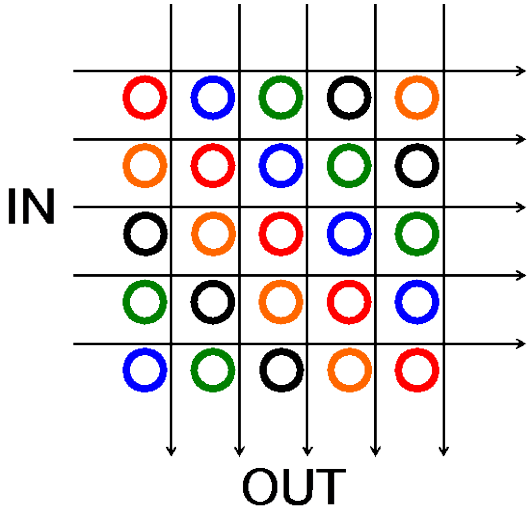


- **Networks may lag interconnect**
  - Routing,
  - Path-hunt, electronics
  - Buffering, electronics
- Low loss
- Pass band shape (larger base elements)
- Low power (non-thermal)
- *Fast* switches
  - Traditional WSS/MEMs competition for slow ones

# Silicon Photonics Challenges

Sandia National Laboratories

- **DWDM Silicon Photonics is 'inevitable'**
  - Today technology is too immature
    - Many technology and cost challenges
  - Tomorrow: *lowest power, lowest cost* solution for 1 TbE transceivers

- **2020-2030 Optics be integrated with high-value ICs?**
  - Re-awakening of optical PCBs (but single mode)
  - New design tools/teams need to be developed.
  - DWDM & potentially mode division mux to maximize IC throughput

  - **HPC and data center interconnection needs *will* require it!**
    - **When?**

# Sandia 2 x 2 silicon photonics switches

Sandia National Laboratories

- Fast (< 100ps)
- Broadband
- 1pJ/switching event
- No static power
- 1 mm size

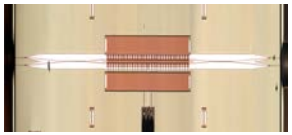

**MZ – free carrier effect**

- Fast (< 100ps)
- Wavelength selective*
- 1fJ/switching event
- No static power
- < 10 um size

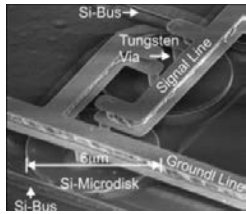


**Ring – free carrier effect**

- Slow (10 us)
- Broadband
- ~ 15 mW/2π
- Static power in one state
- < 10 um size + coupler

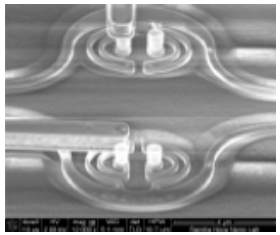

**MZ – thermo-optic**

- Slow (10 us)
- Wavelength selective
- ~ 4 uW/GHz (200uW)
- Static power in one state
- < 10 um size

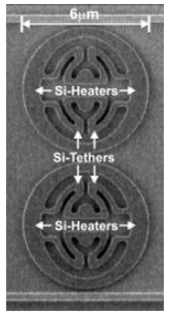


**Ring – thermo-optic**

*Can also switch all channels at once if free spectral range = channel spacing

# **Why DWDM** vs. high speed and multi-level formats?

Sandia National Laboratories

- Networks with high-radix switches
  - ***Greater connectivity leads to greater network efficiency***
- Energy Consumption and delay
  - ***Many lower speed channels vs. higher speed multi-level ones***
  - ***No error correction***

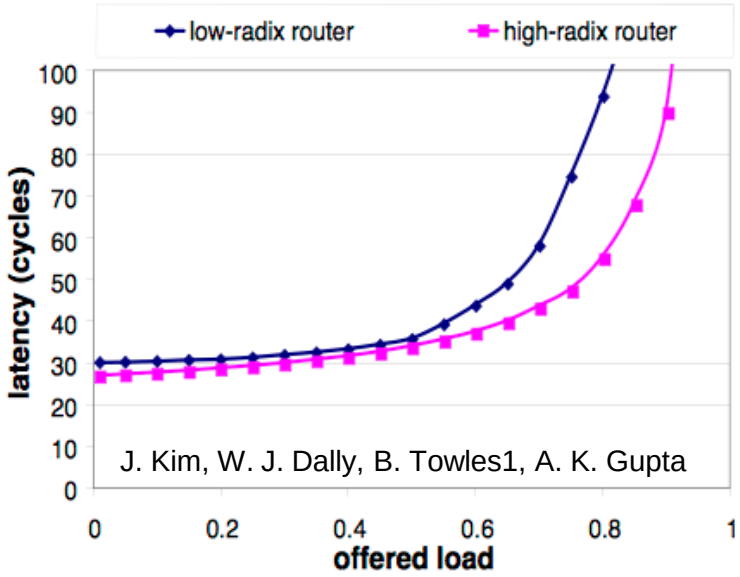


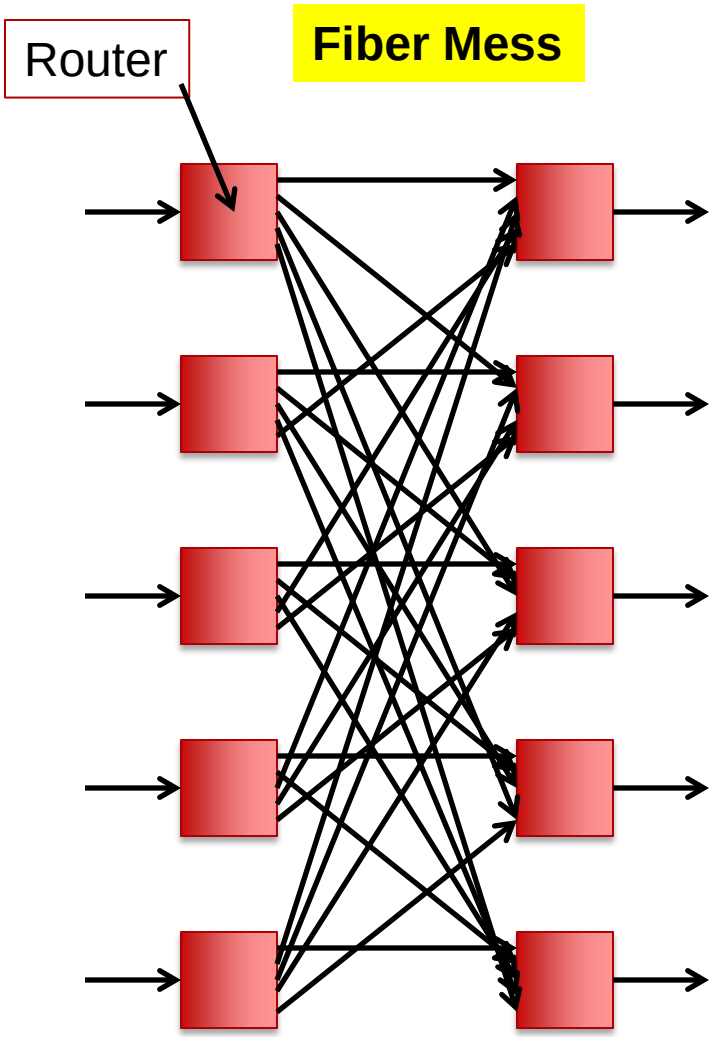

# High radix and DWDM?

Sandia National Laboratories

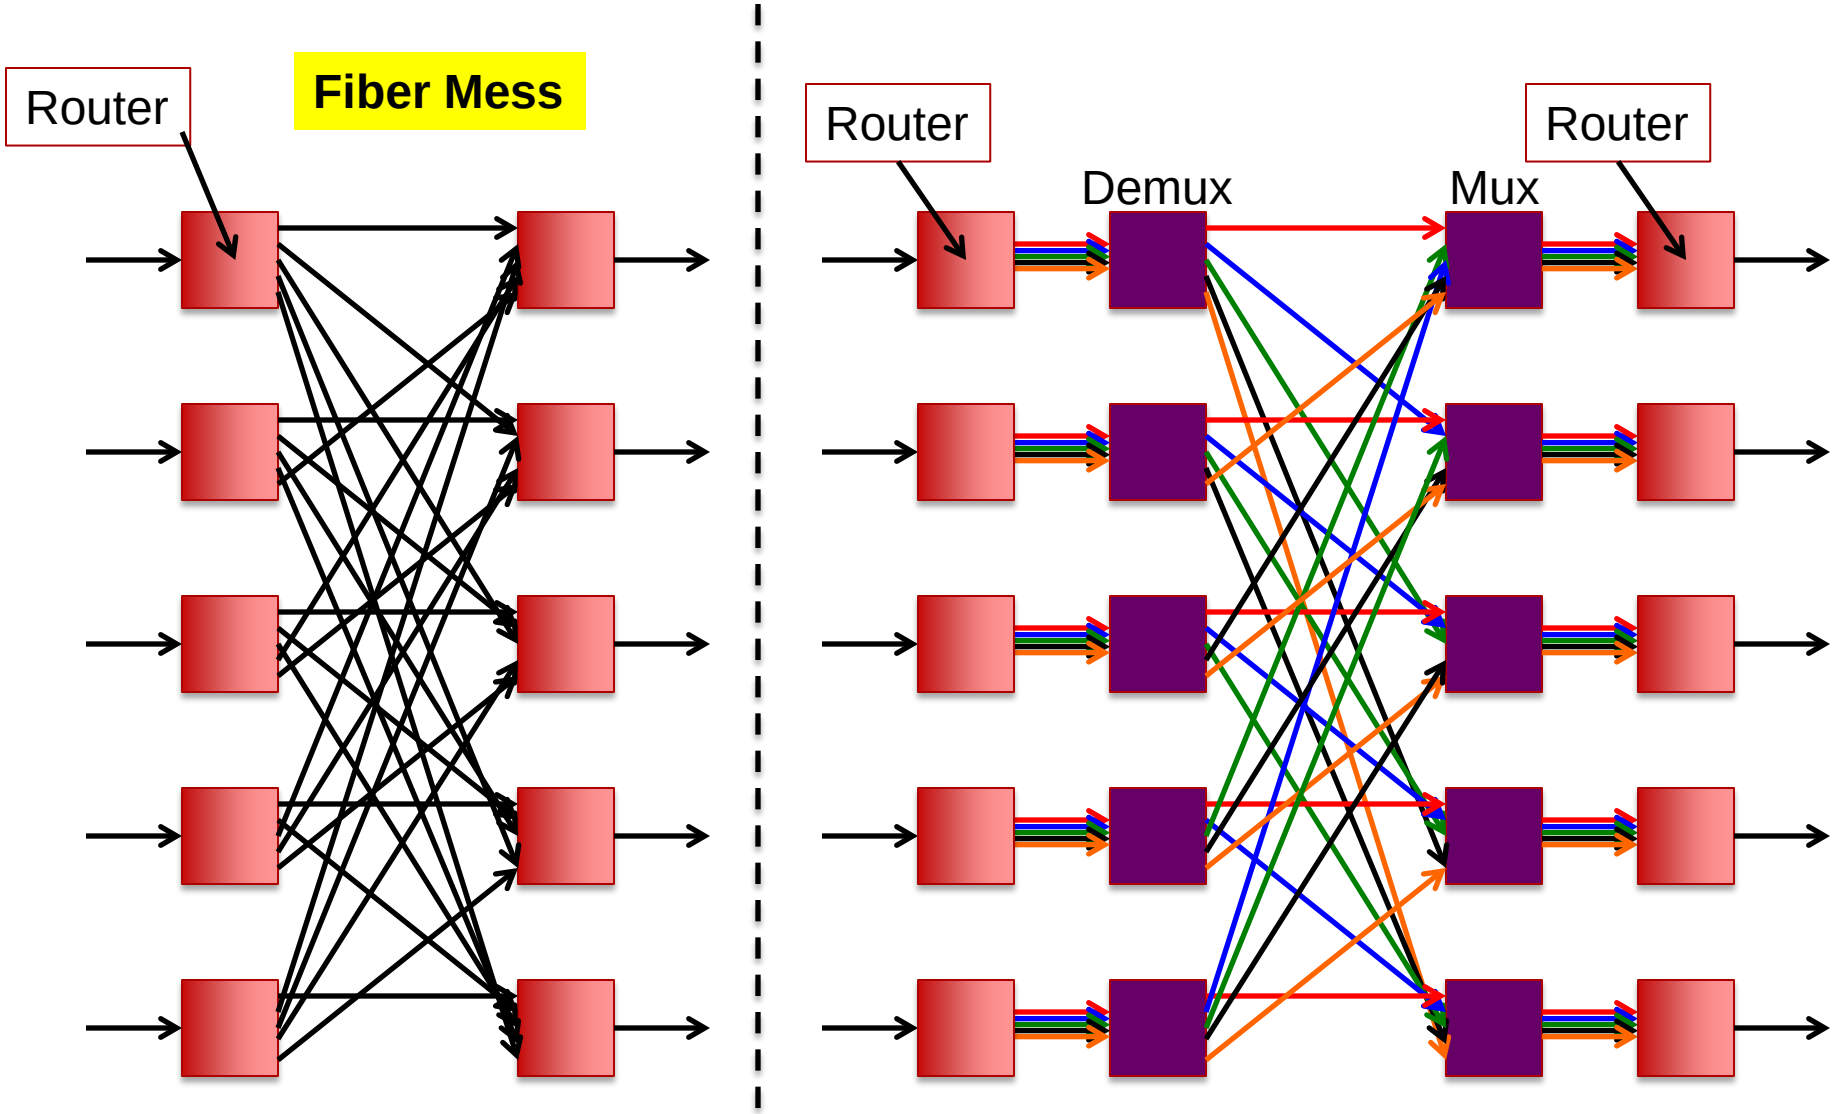

# High radix and DWDM?

Sandia National Laboratories

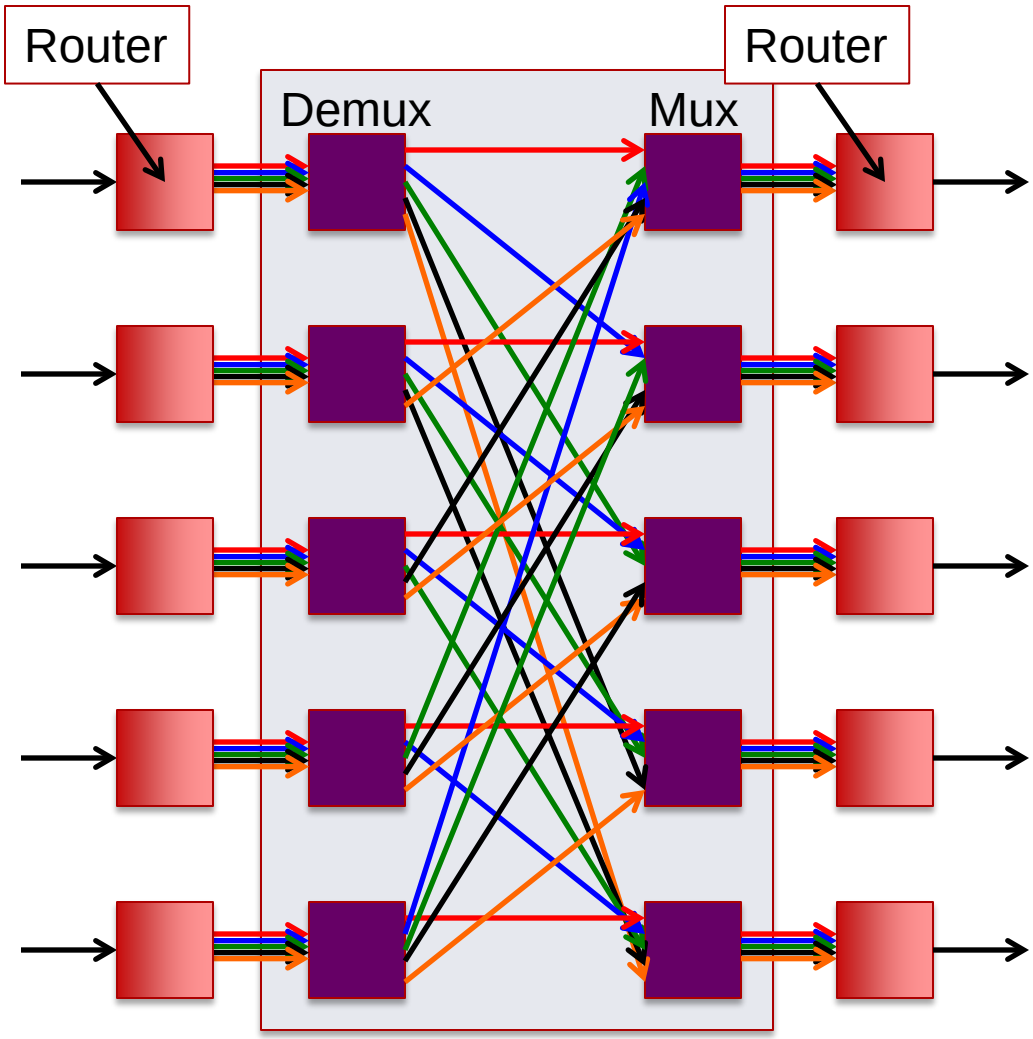


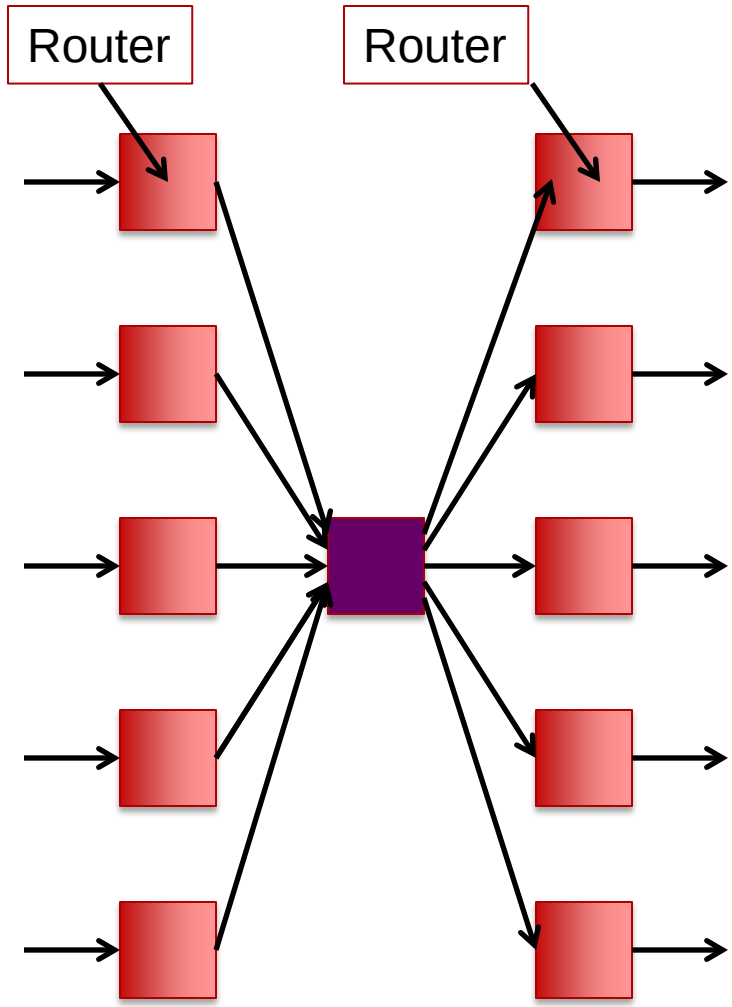


- DWDM reduces fiber cost

# High radix and chip-scale DWDM?

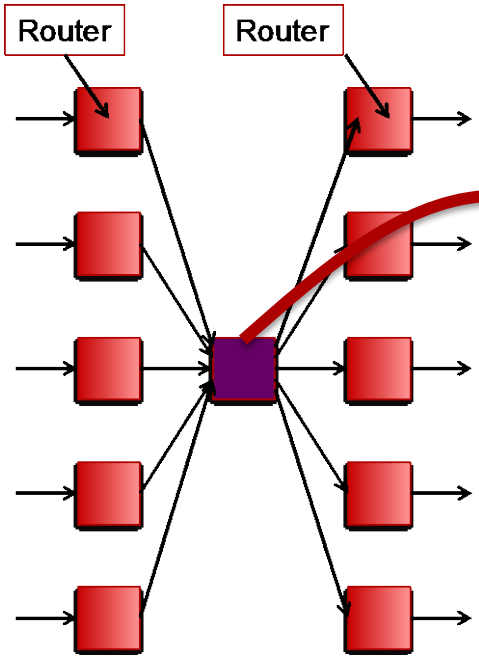


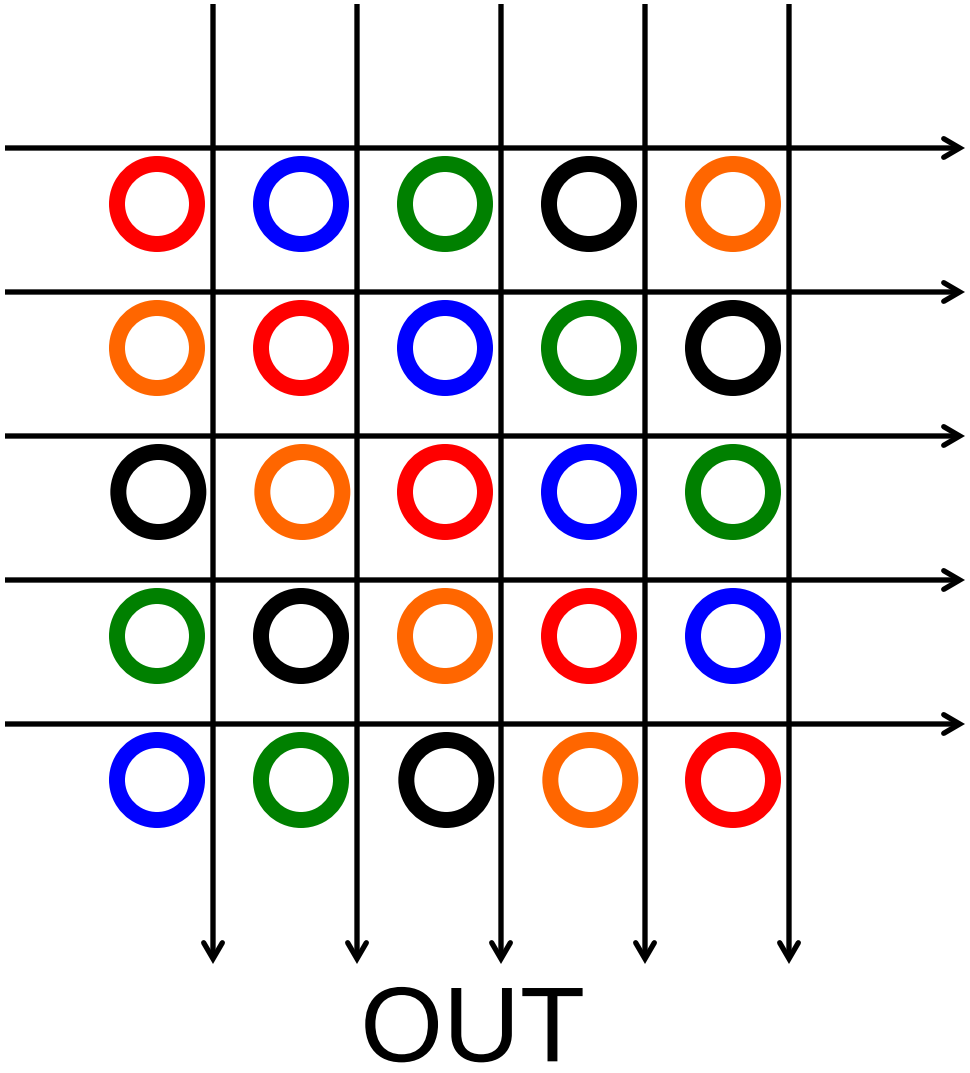


- Chip-scale solution
- Low loss (pass most resonators)
- Impossible with high-data rate, advanced modulation format
- Saves fiber routing compared to parallel interconnect
- Can be ***reconfigurable***

# Effect of temperature on loss budget

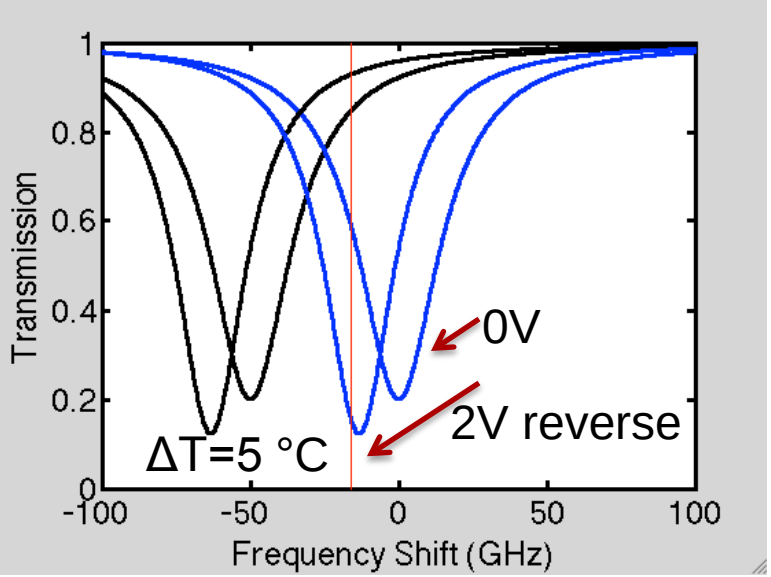


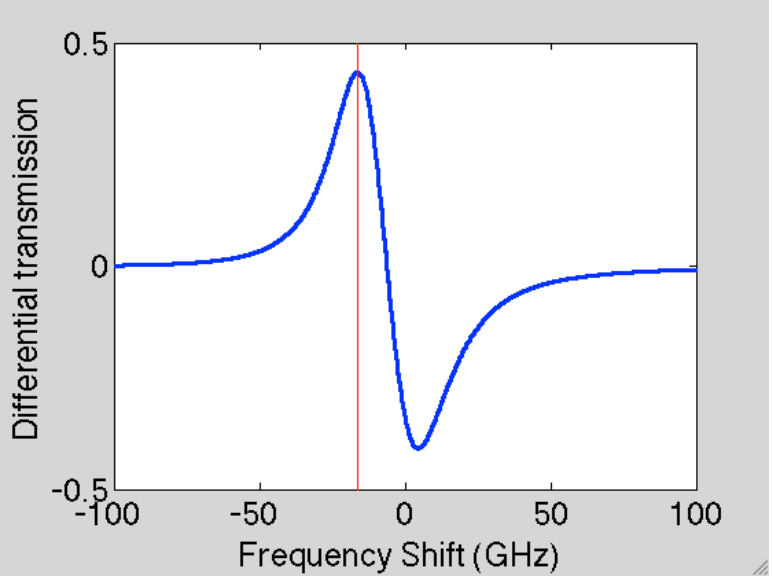


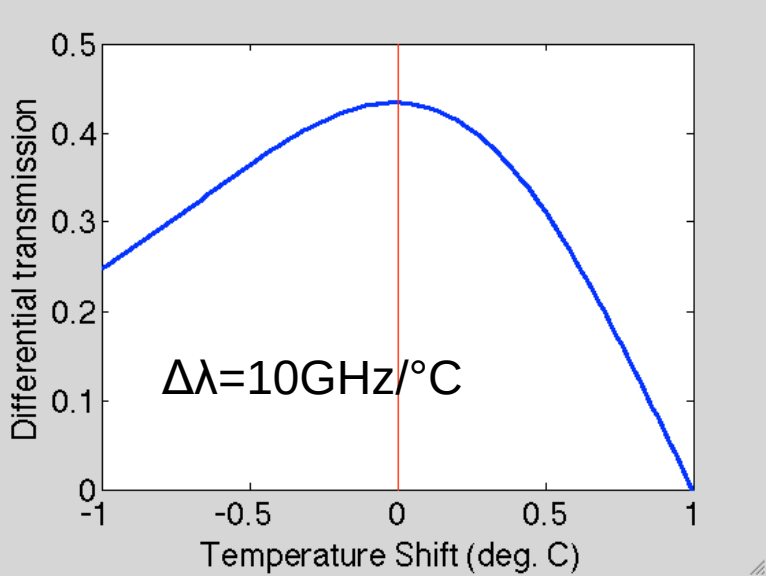


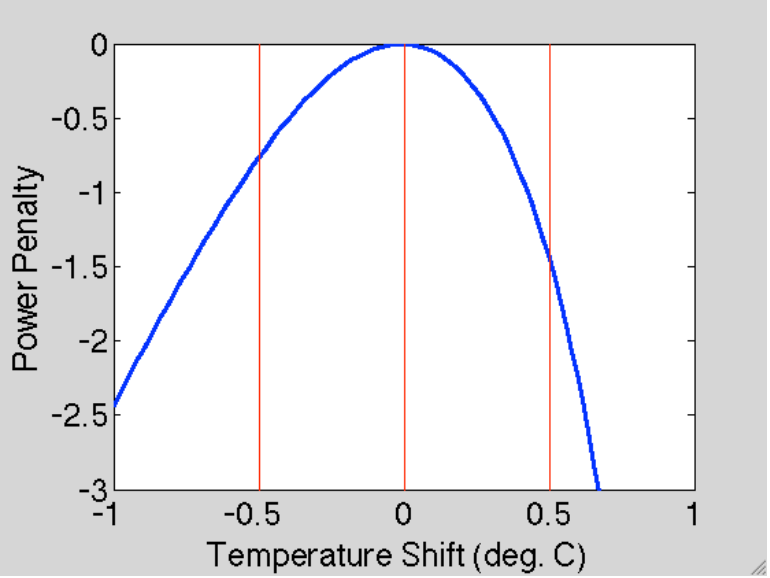

# Resonant Wavelength Closed Loop Control

- Control Loop
  - Measurement
    - Temperature
    - Power (shown)
    - Phase (BHD, PDH)
    - Bit errors
  - Integration (PI Loop)
  - Stimulus
    - Integral Heater (shown)
    - Forward bias (heater/carriers)
    - Reverse bias (carriers)
    - Strain

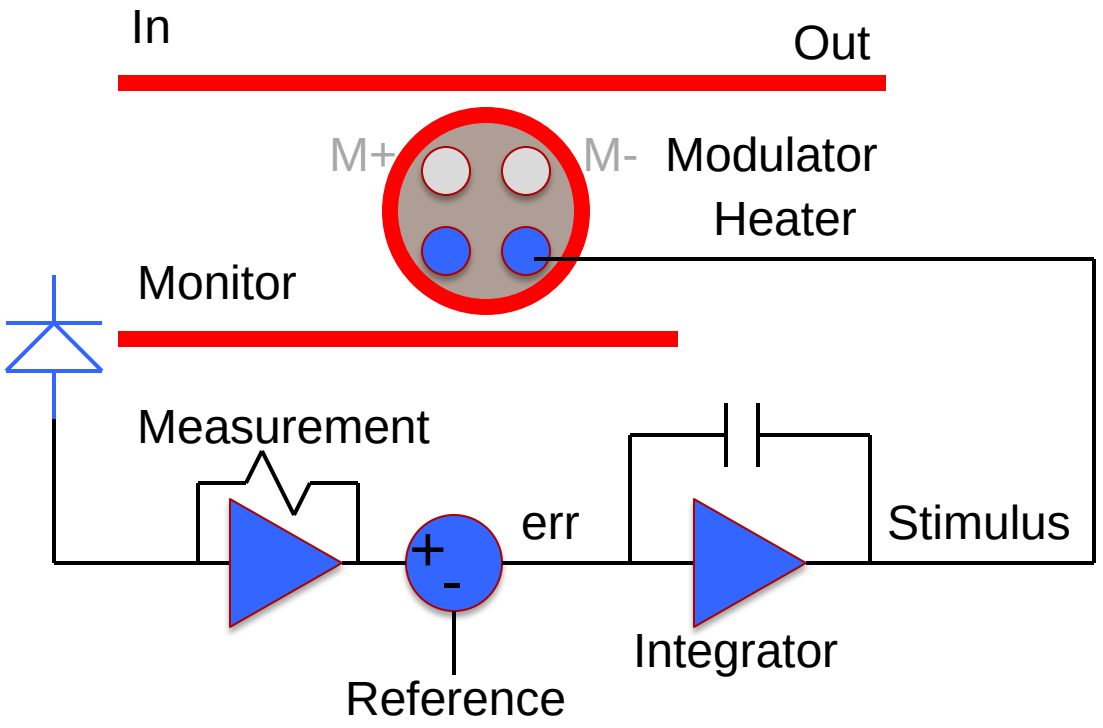

# Resonant Wavelength Locking

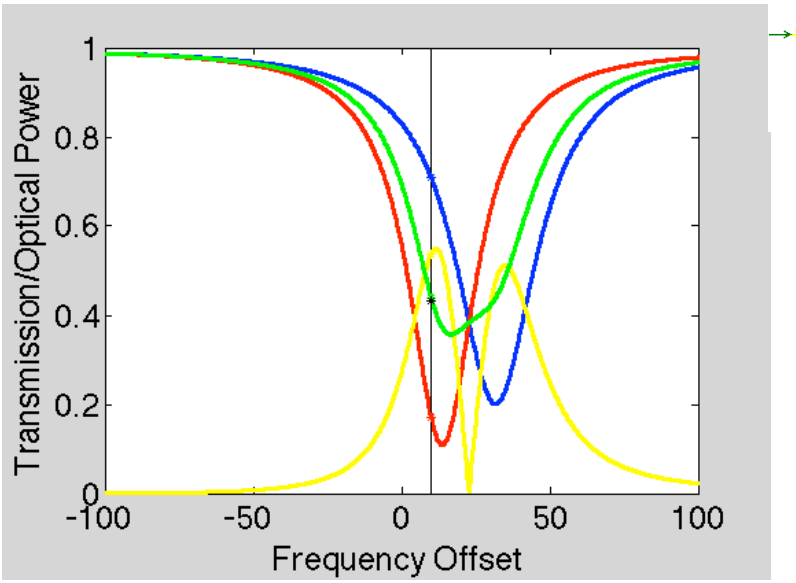


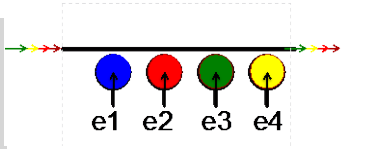


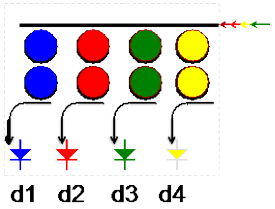


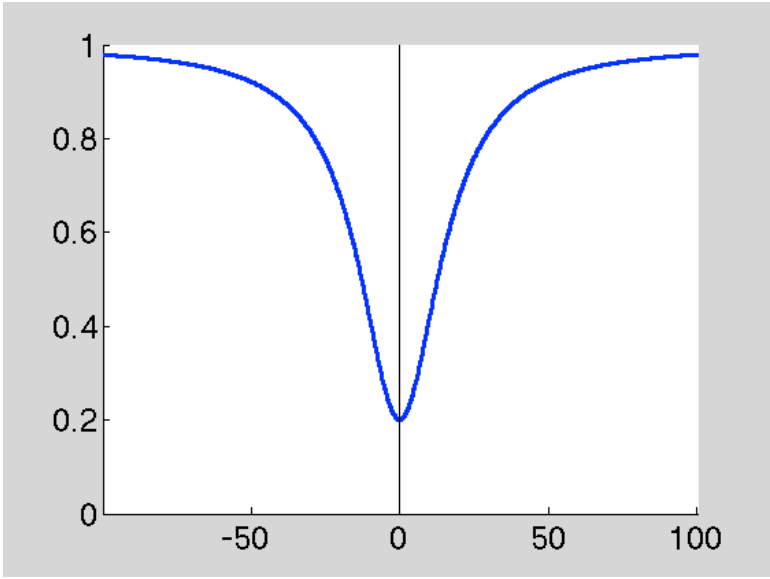


**Modulator**

- Lock on side of resonance

**Filter (DeMux)**

- Lock at minimum power

# Resonant locking of a DWDM filter

- Problem: locking on minimum power level does not lend itself to a simple control loop
- Solution: Homodyne detection with balanced detection gives optimal locking solution

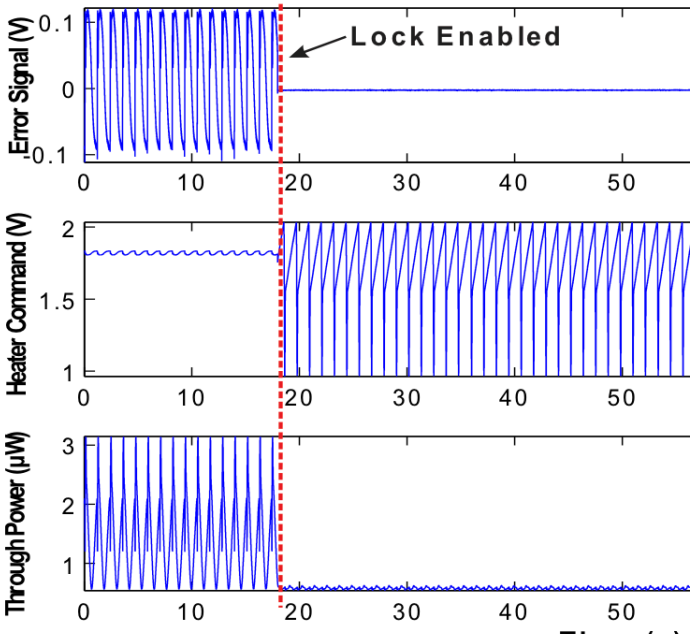


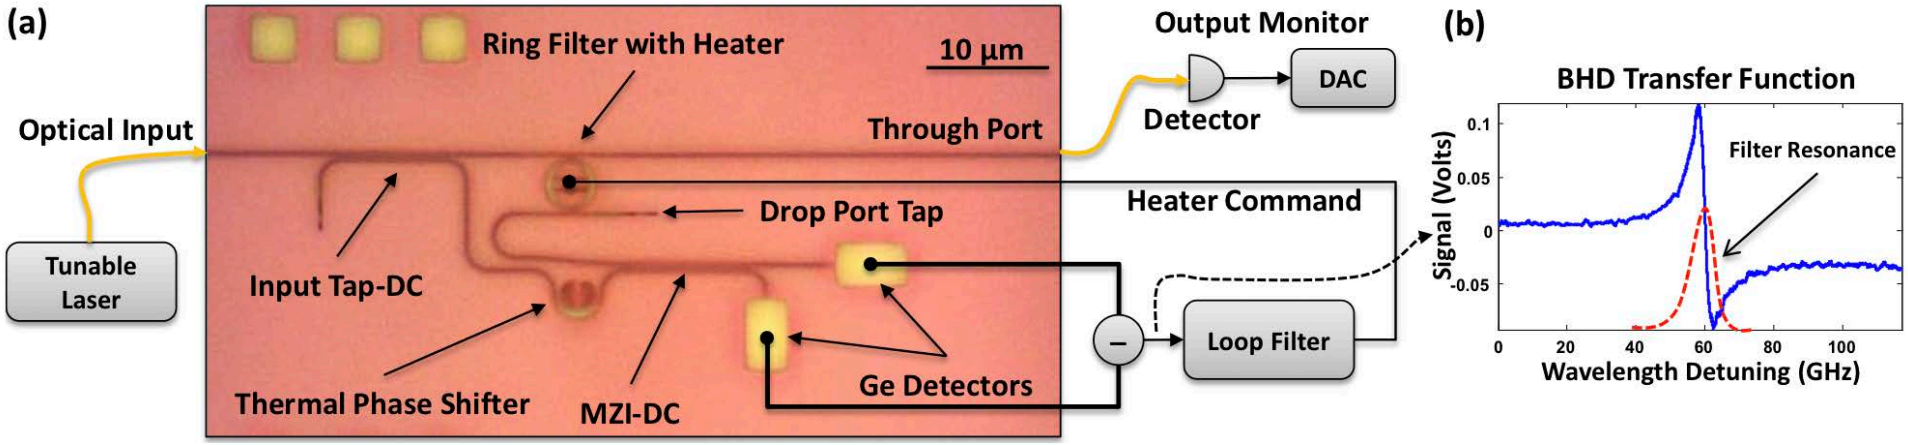


J. A. Cox et. al., IEEE Optical Interconnects 2013

# λ Stability: ITU grid or ΔT?

- Range requirements **0 – 85 ° C**
- If you laser λ wander → *difference* between Tx, Rx, Laser
- Silicon photonics frequency elements will track better

## Cyclical Channels

- Example: 4 × 100GHz channel spacing
- Max. heating = ch. spacing / df/dT
  - →100GHz/10GHz/° C = **10° C**
    - N. Binkert et al., ISCA 2011;
    - M. Gorgas et. al., IEEE CICC, 2011
    - A. Krishnamoorthy et. al., IEEE Photonics J., 2011

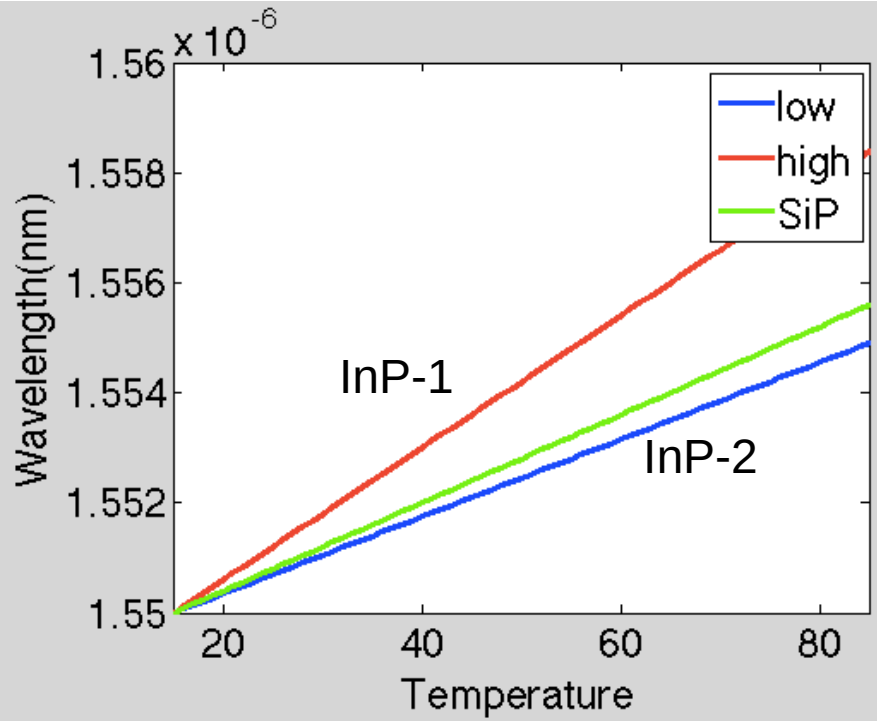


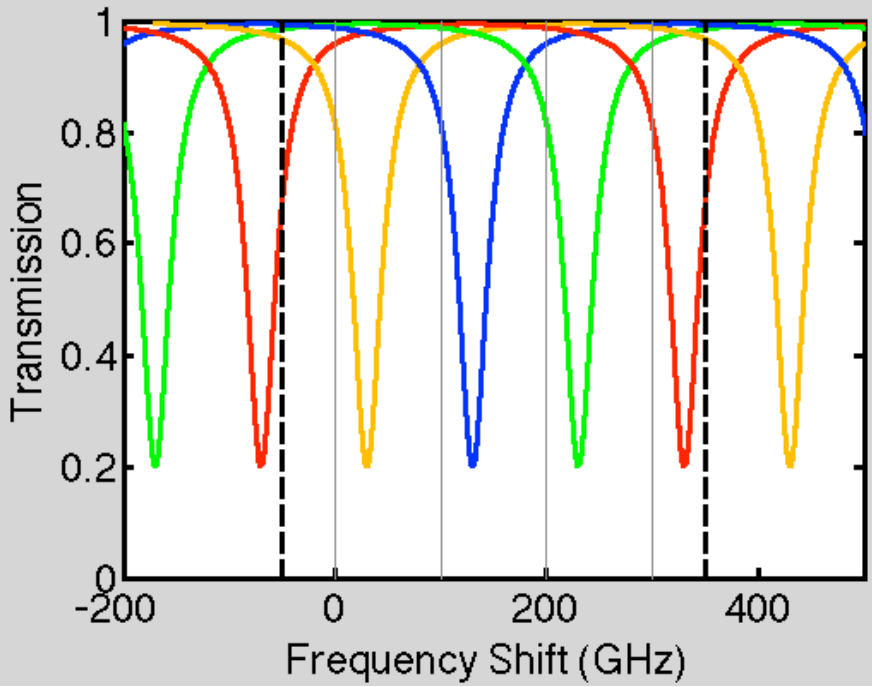

# Technology Challenges

- Integration
  - *Silicon photonics integration with state of the art CMOS with low capacitance and high yield*
  - *Cost effective, reliable packaging*
  - *Fiber coupling and waveguide losses*
- Silicon Photonics
  - *Efficient Laser source*
  - *Modulator and optical filter resonant wavelength stability and uniformity*
  - *Filter shape, coupling variations*
  - ***Low energy receivers***
- Interface Electronics
  - *Efficient clock and data recovery*
  - *Data TDM multiplexing (SERDES)*
  - *Efficient Data encoding and error correction*

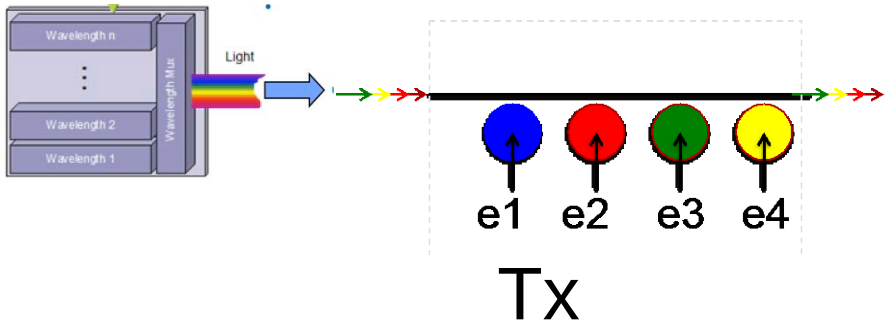


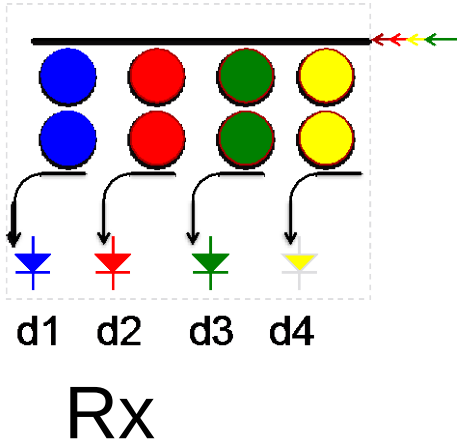

# Simple Modulator Driver: Differential Signaling *simulation*

- No pre-emphasis
- No AC coupling
- No high voltages
- CMOS logic levels

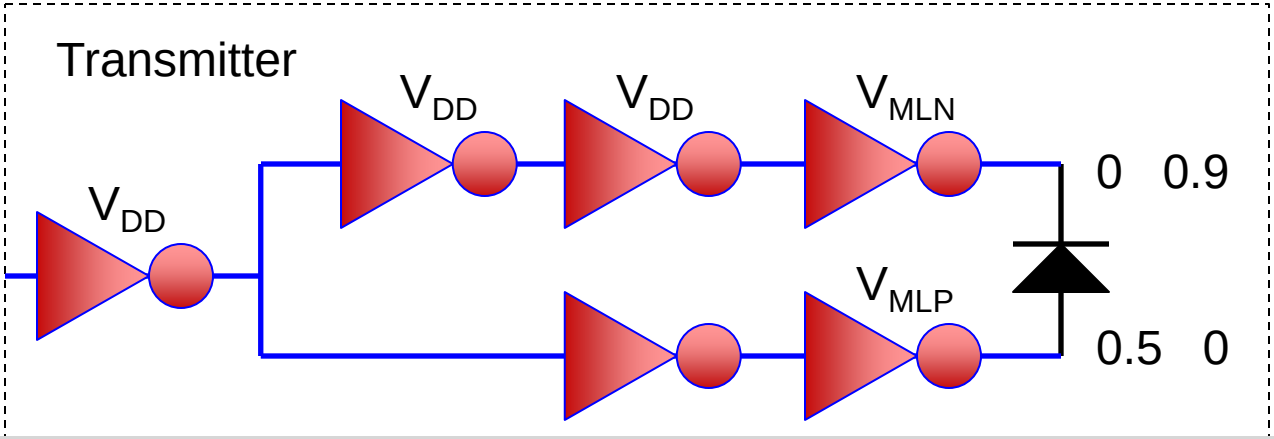


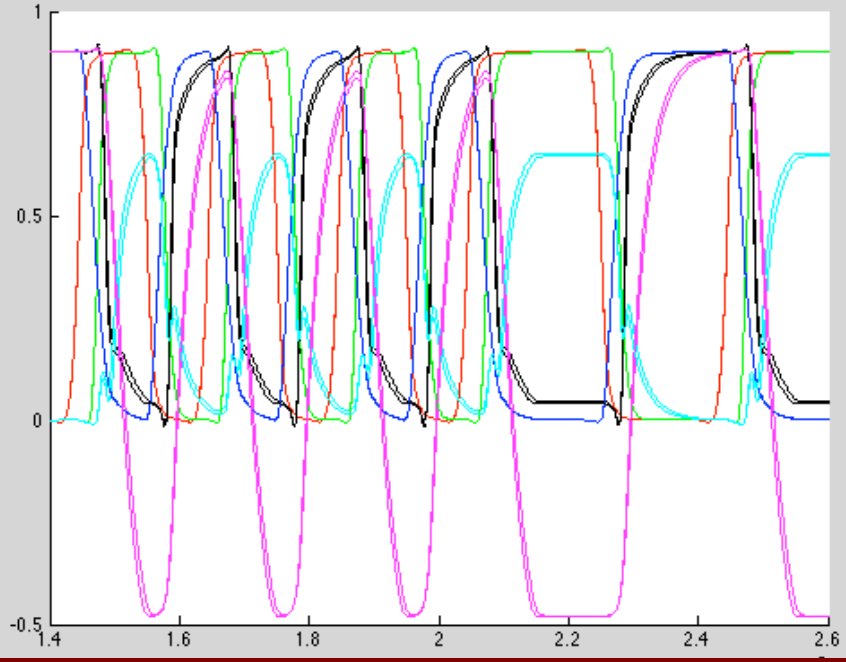


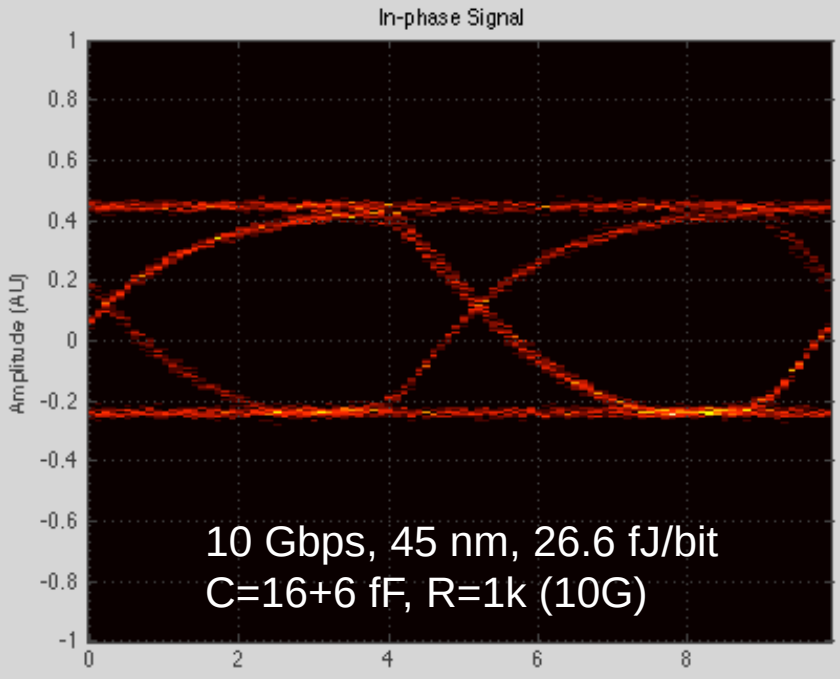

# High voltage drivers (*simulation*)

- Modulators at higher speed with low capacitance will require higher voltages (less carriers to shift)
- Simulation includes level shifter and forward bias inverters

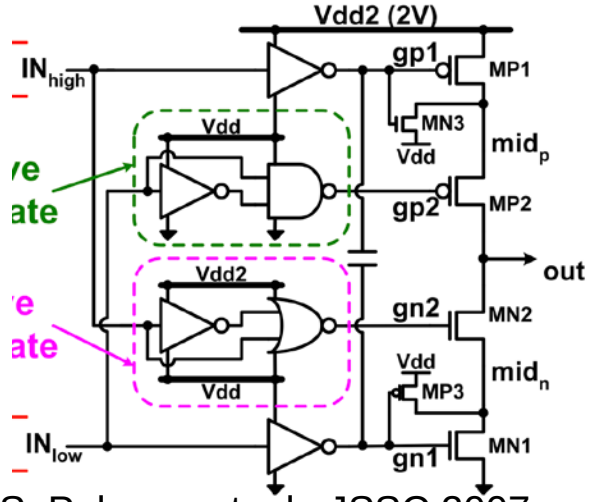


S. Palermo et. al., JSSC 2007

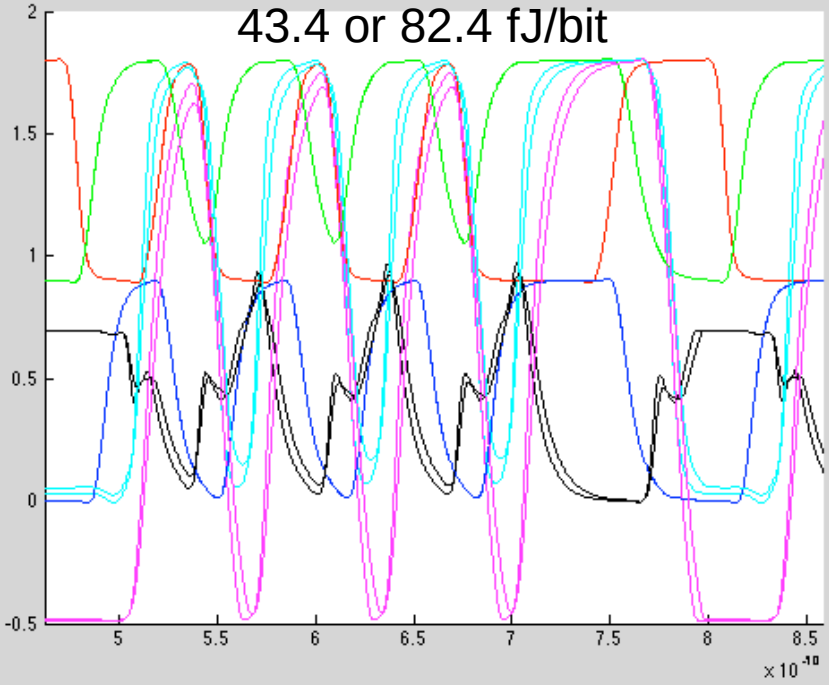


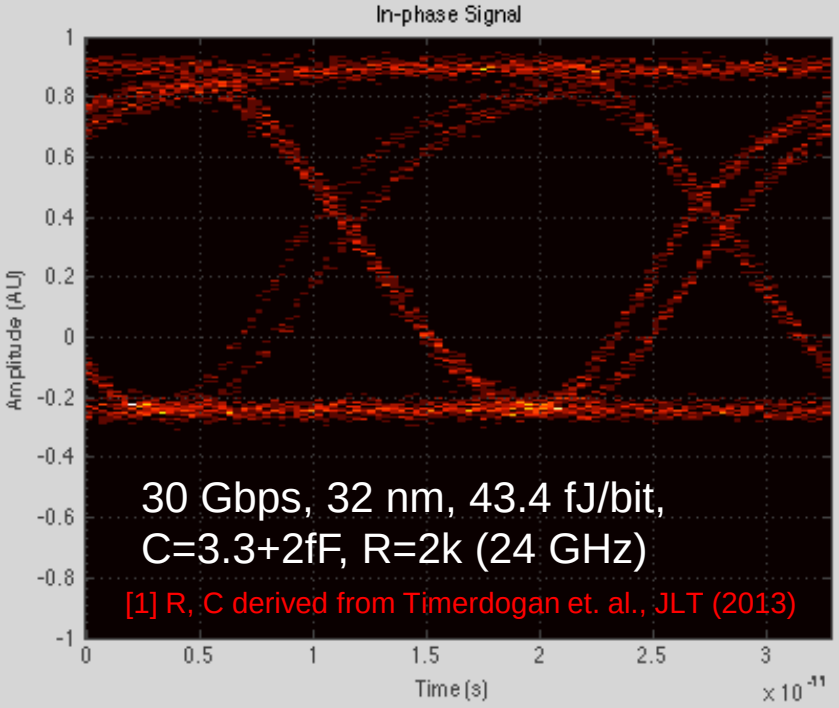

# Low Voltage Inverter Chain Driver (*simulation*)

Sandia National Laboratories

- 3 inverters in a row
- Params: MIT-HD modulator
  - C=16 + 6 fF, R=400 (25GHz),
  - 0.5V (0.7V) drive
  - E. Timerdogan et. al.," arXiv: 1312.2683 (2014)

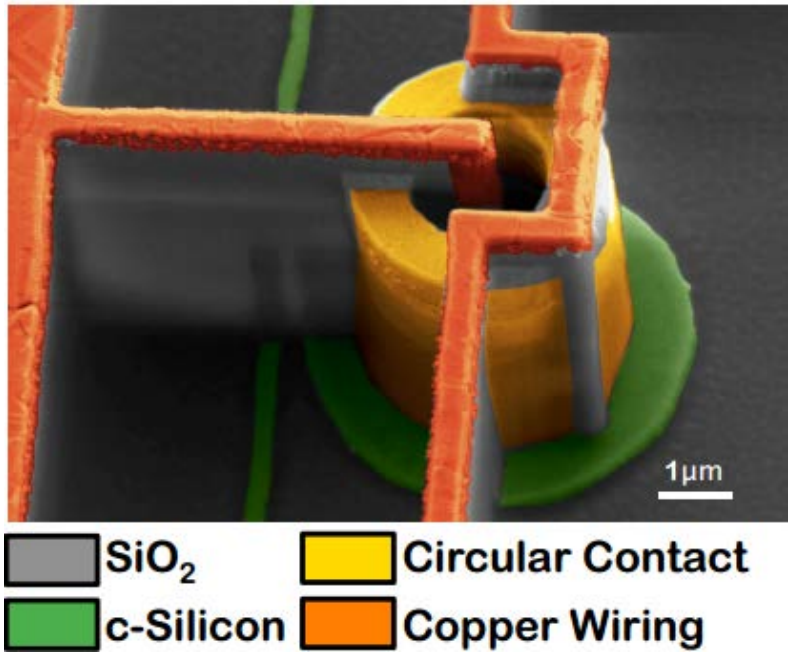


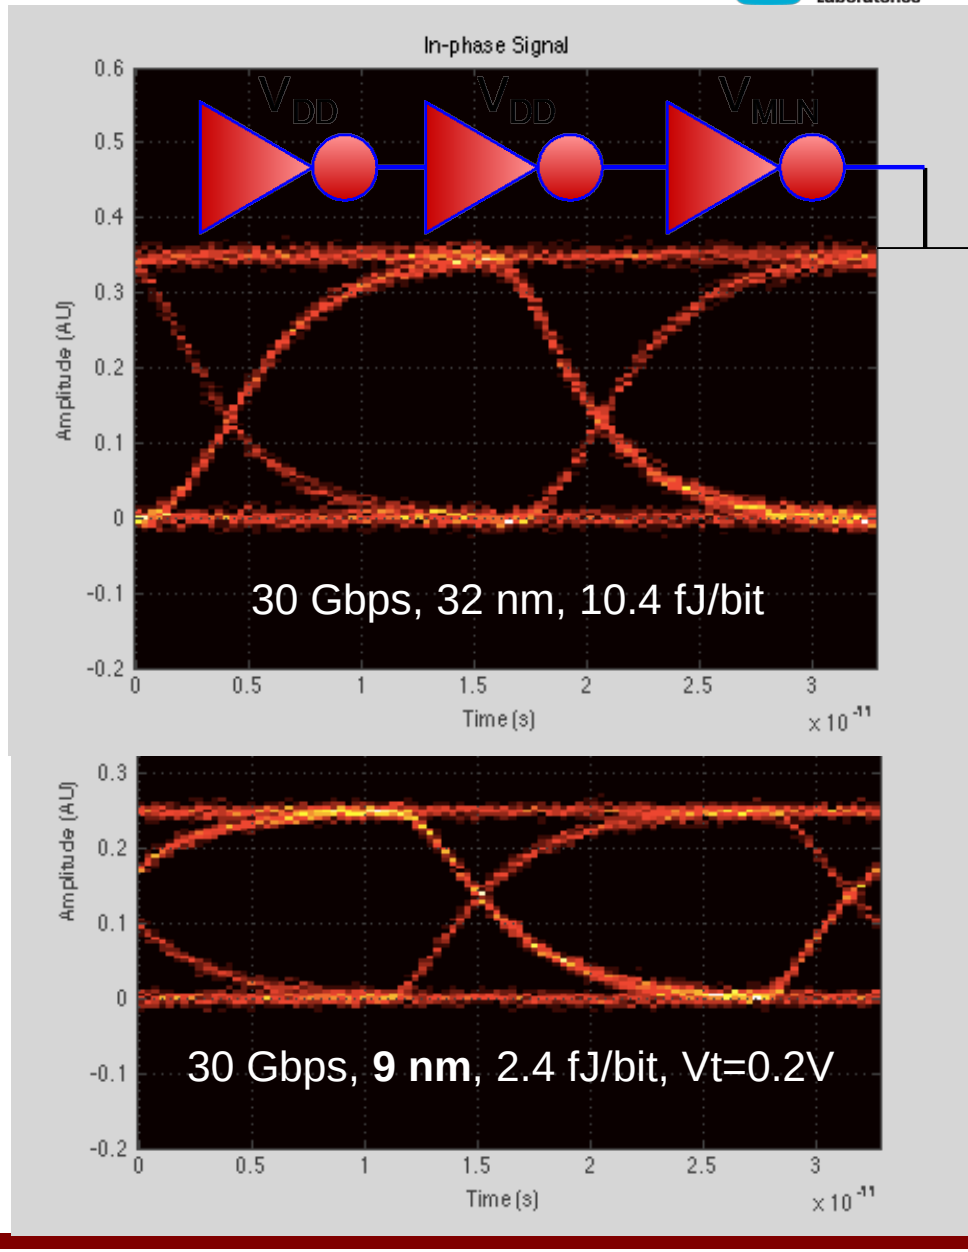

# Differential Optical Signaling?

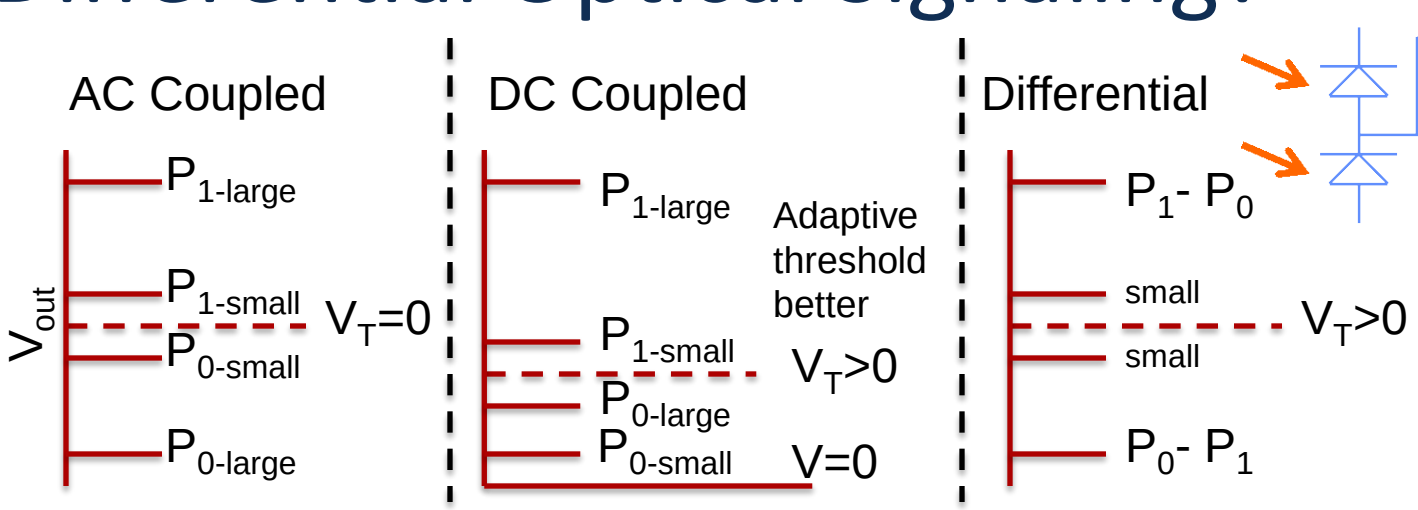


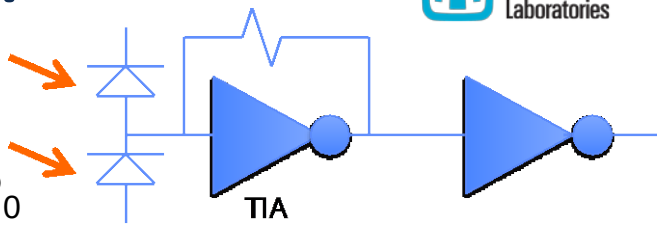


Avoid absolute optical power levels for logic 1 and logic 0

- Today's Rx work over large Rx Power because of AC-coupling
  - AC coupling requires data encoding and large capacitors
- Today's electronics is differential
  - Scalable signals, common mode rejection, 2X 'effective' power
  - Some modulators give you the signal for free
  - But … twice as many signals, equality of path, cost
  - More practical for board level than for transceivers

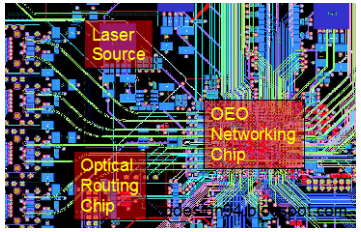


See e. g. Lentine and Miller, JQE 1993;