Challenges in the implementation of DWDM optical interconnects using *resonant* silicon photonics

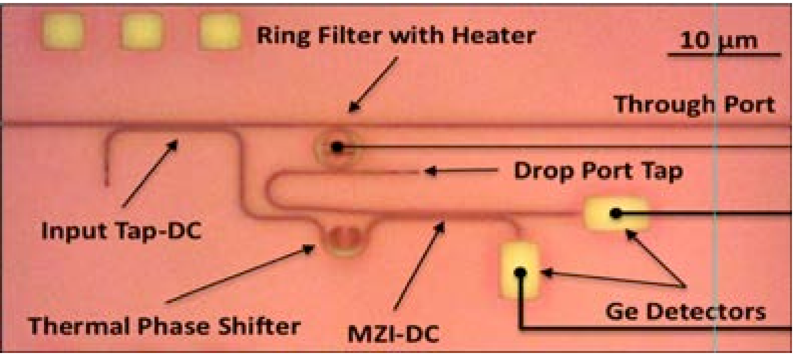

Anthony L. Lentine and Christopher T. DeRose

Sandia National Laboratories, Albuquerque NM

alentine@sandia.gov

Optical 'short reach' Interconnects

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**

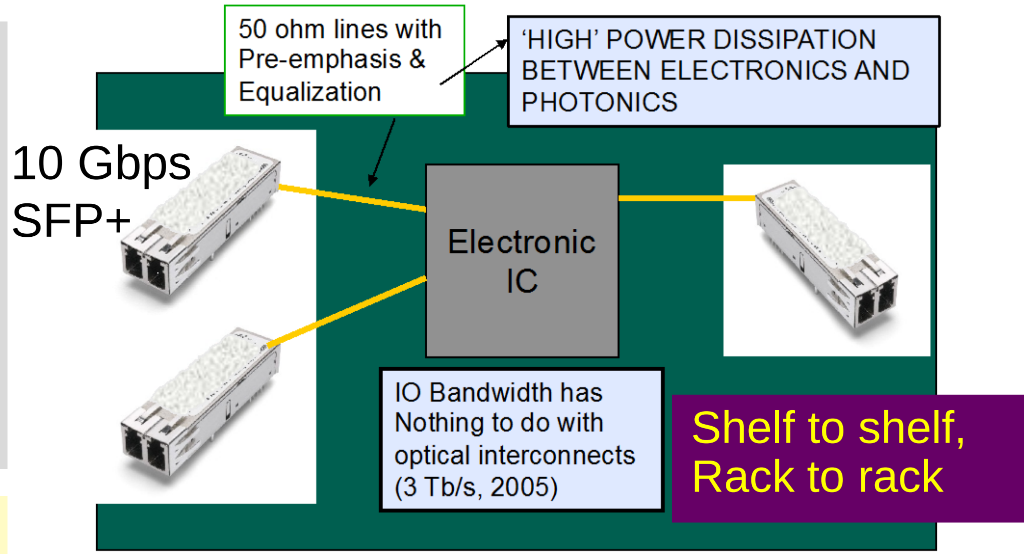

OPTICS FOR DISTANCE

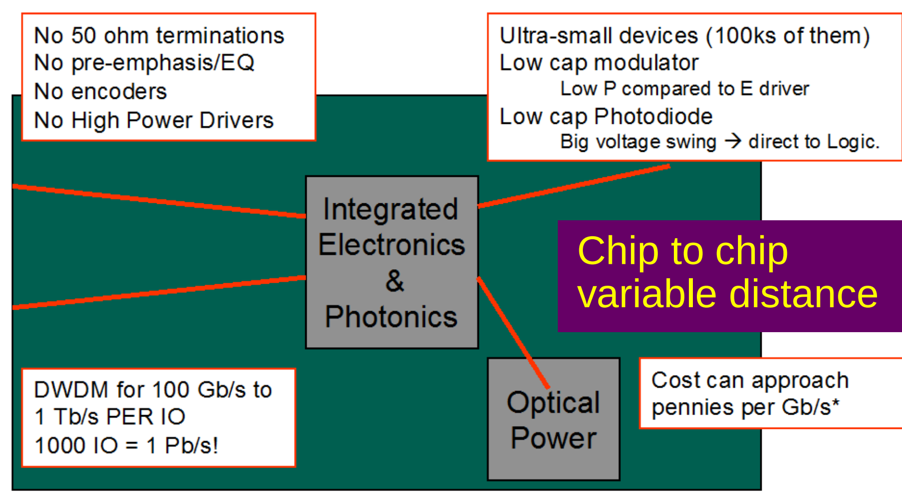

Not addressing on-chip interconnects

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

Si Photonics for DWDM

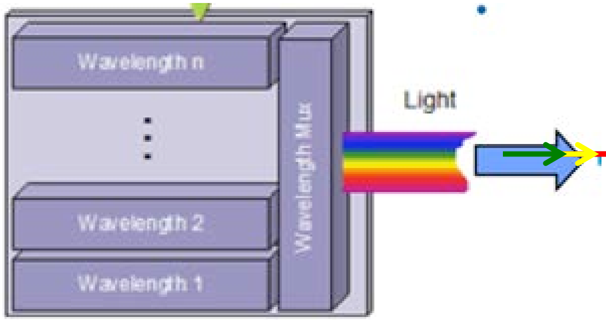

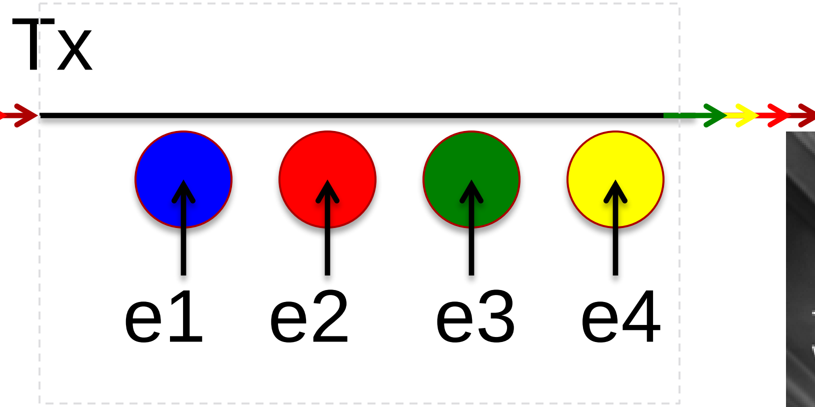

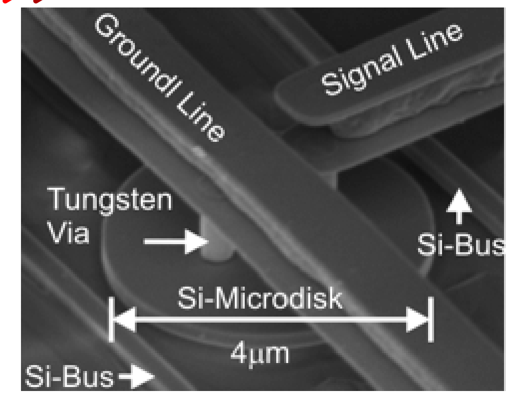

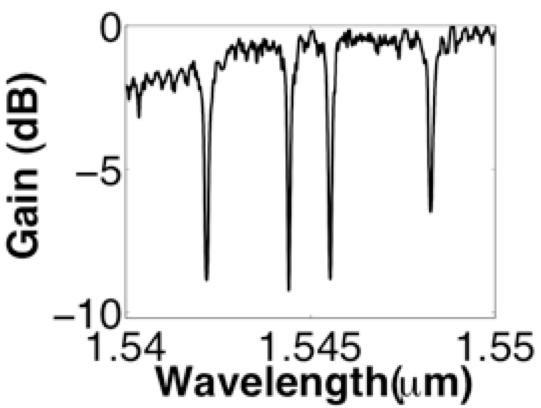

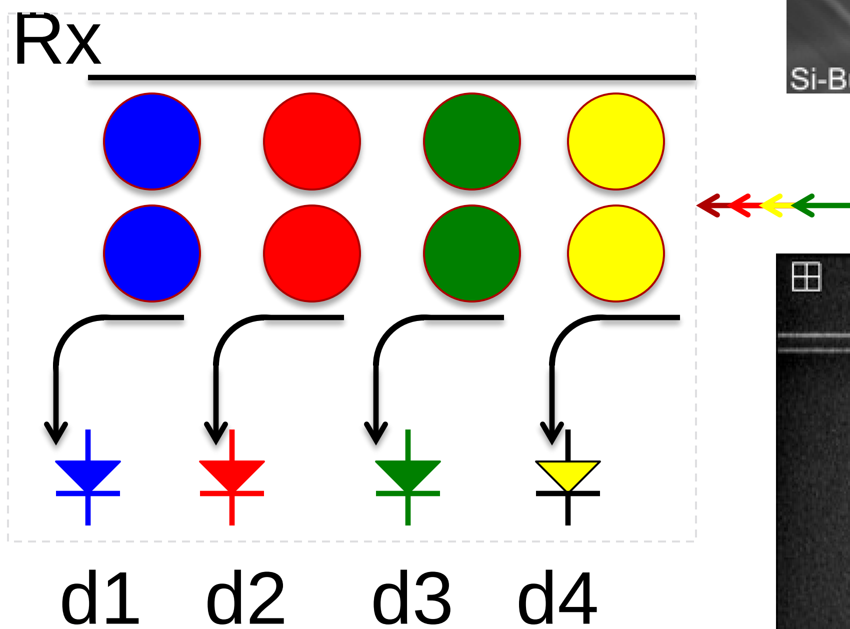

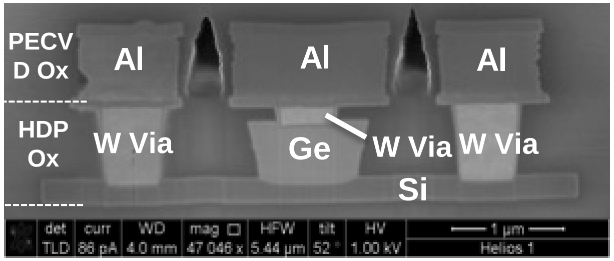

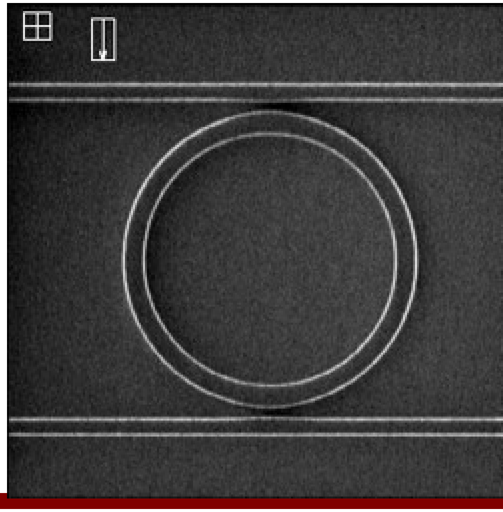

Silicon Photonics Layer Structure

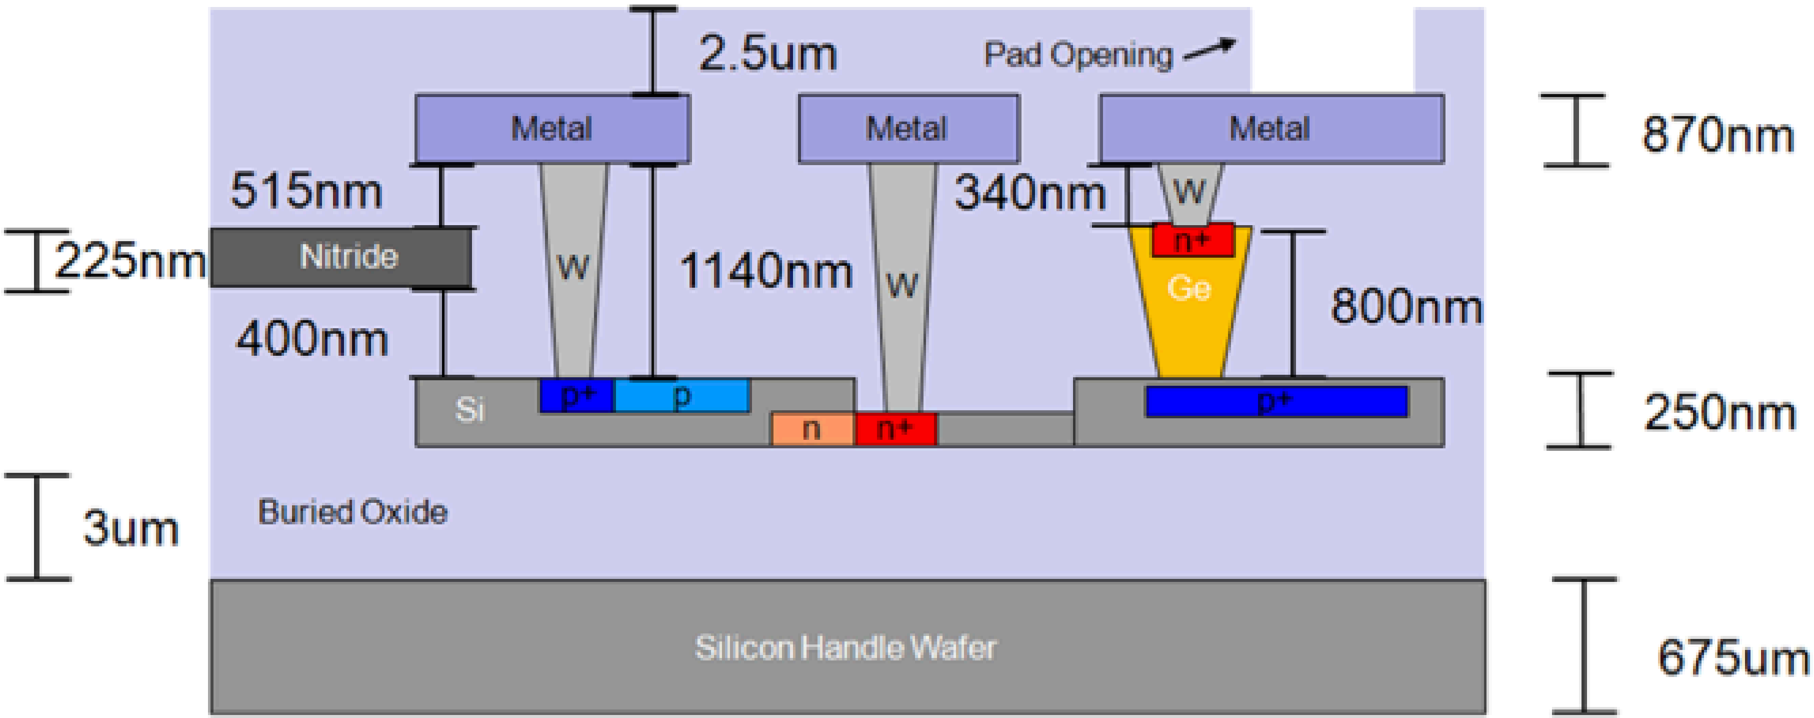

See e. g. A.L. Lentine, C. T. DeRose, P. S. Davids, N. J. D. Martinez, W. A. Zortman, J. A. Cox, A. Jones, D.C. Trotter, A. T. Pomerene, A. L. Starbuck, D. J. Savignon, T. Bauer, M. Wiwi, and P. B. Chu, “Silicon Photonics Platform for National Security Applications, “in 2015 IEEE *Aerospace Conference,*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

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)

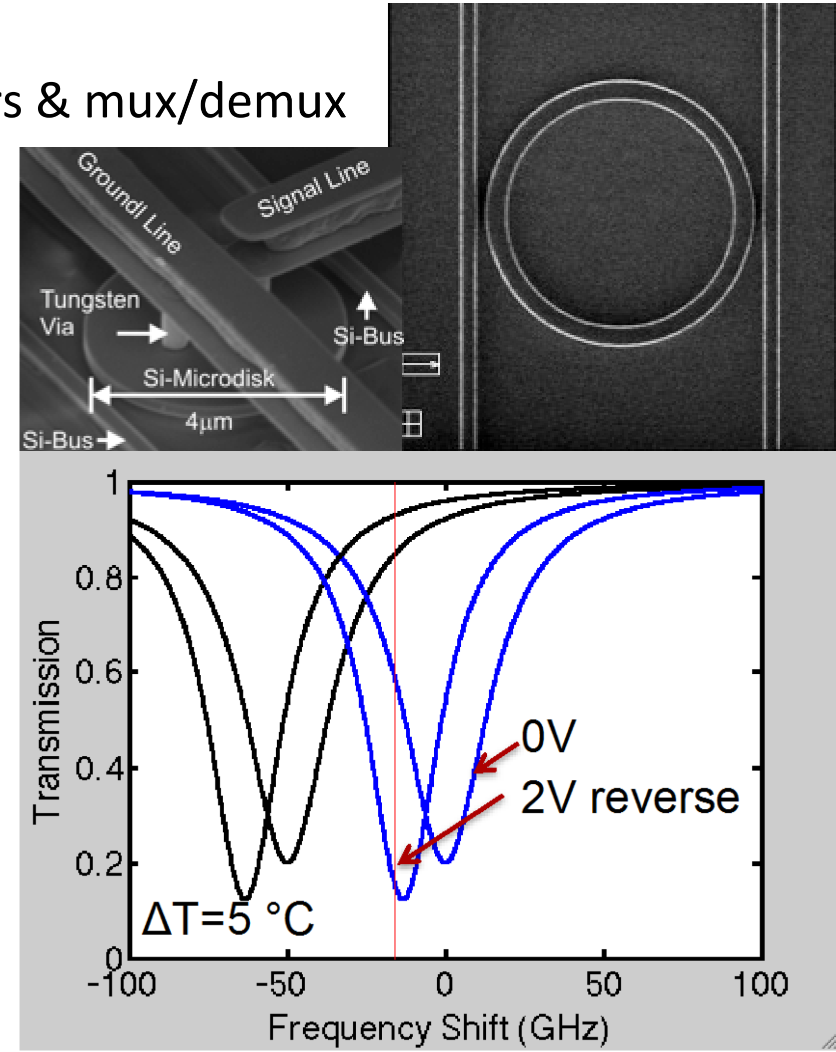

Why DWDM vs. high speed and multi-level formats?

- 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***

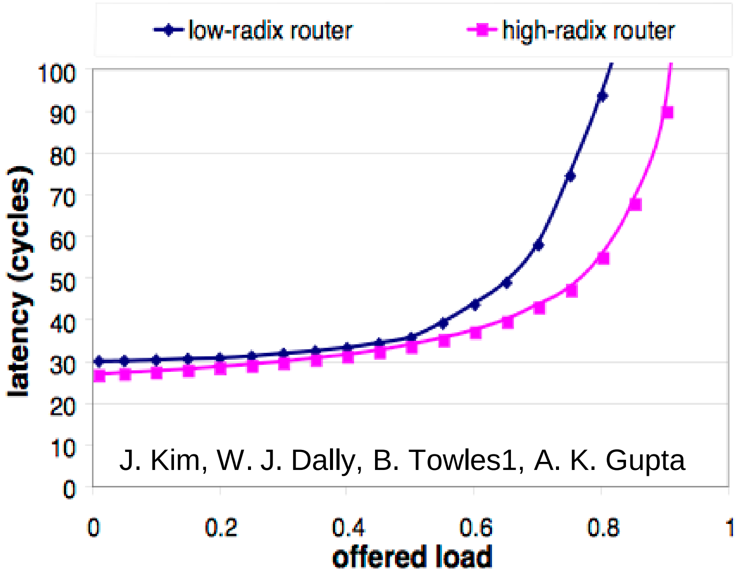

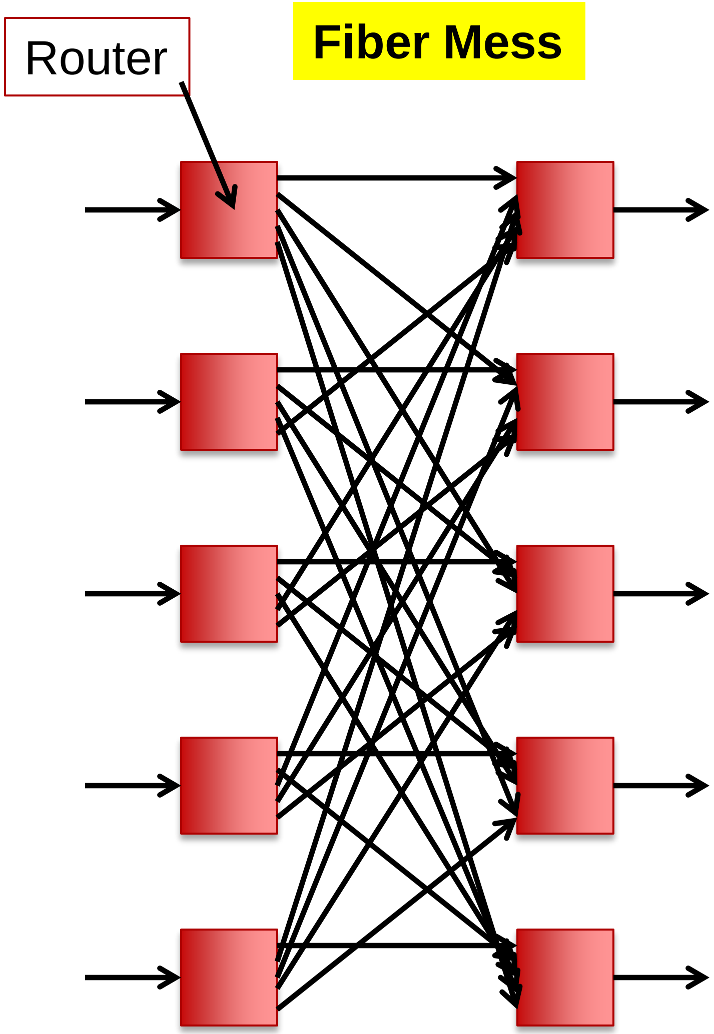

High radix and DWDM?

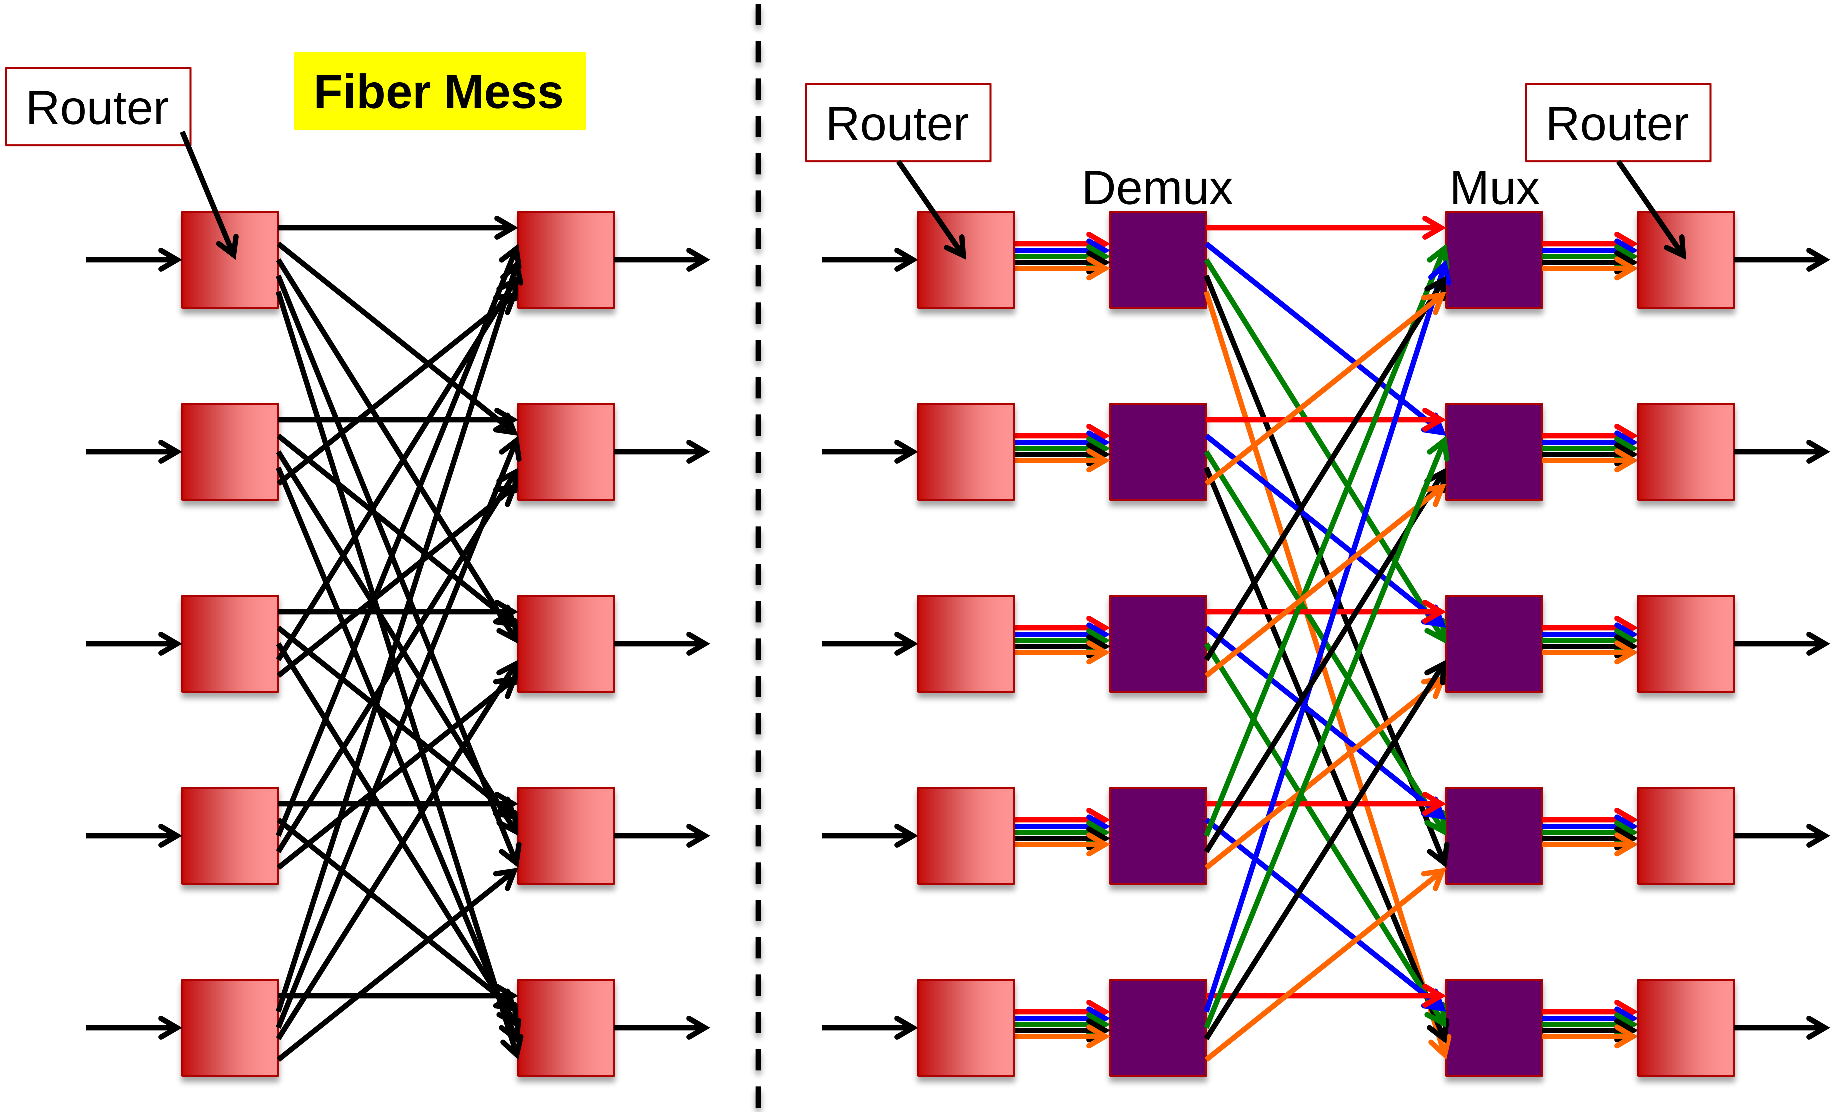

High radix and DWDM?

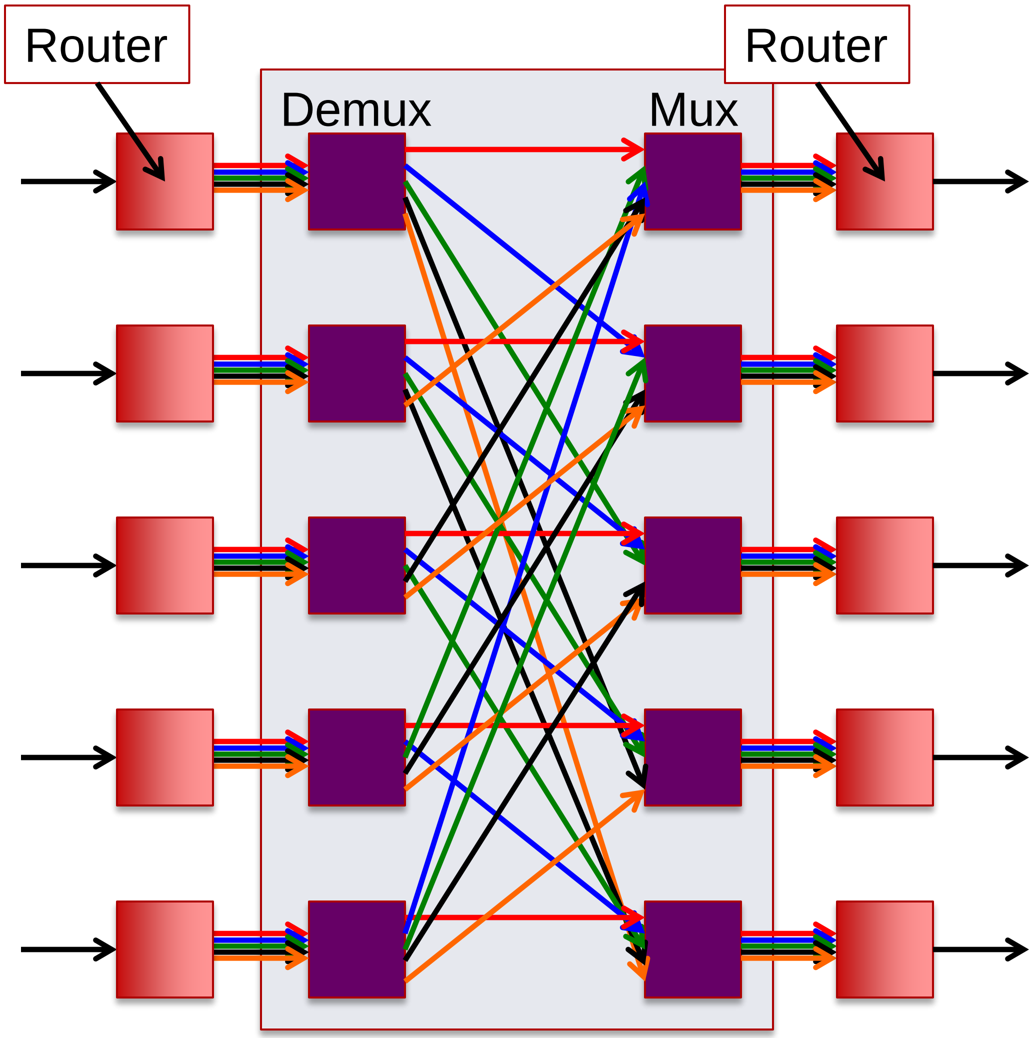

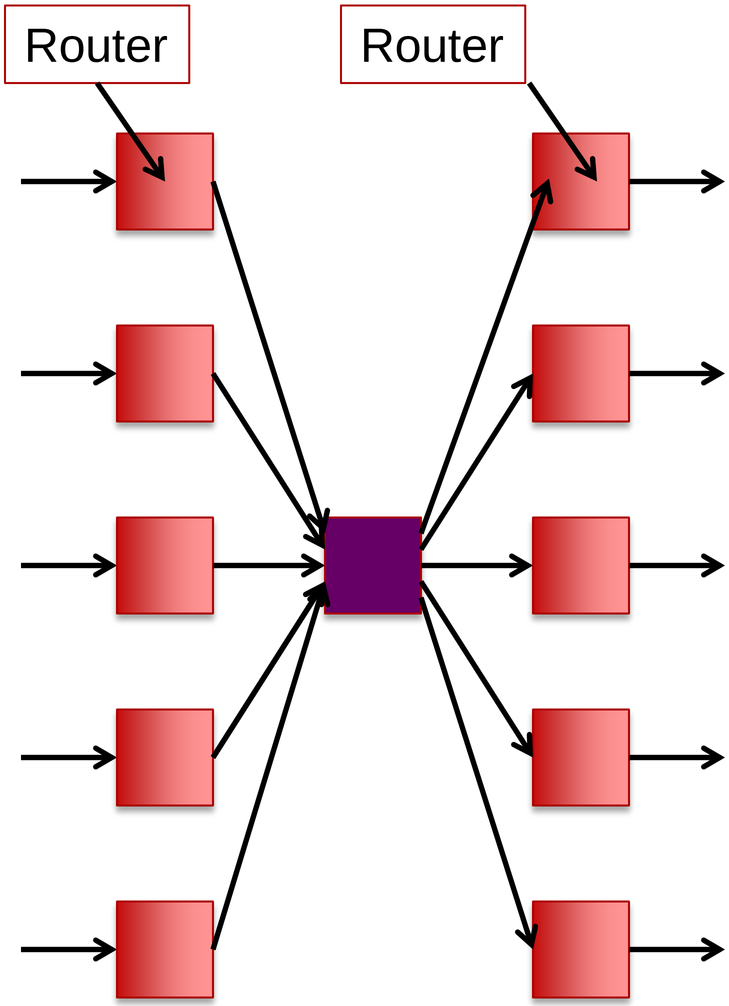

- DWDM reduces fiber cost

High radix and chip-scale DWDM?

Sandia National Laboratories

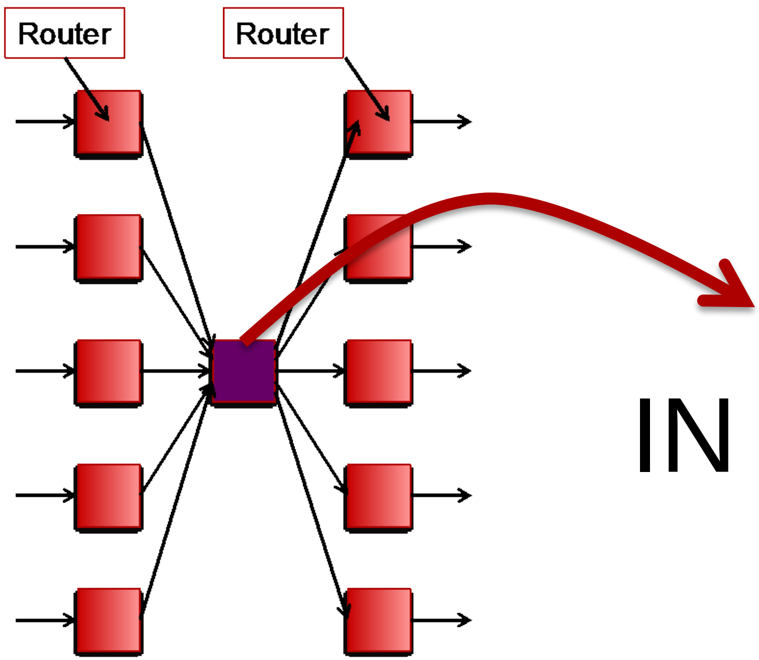

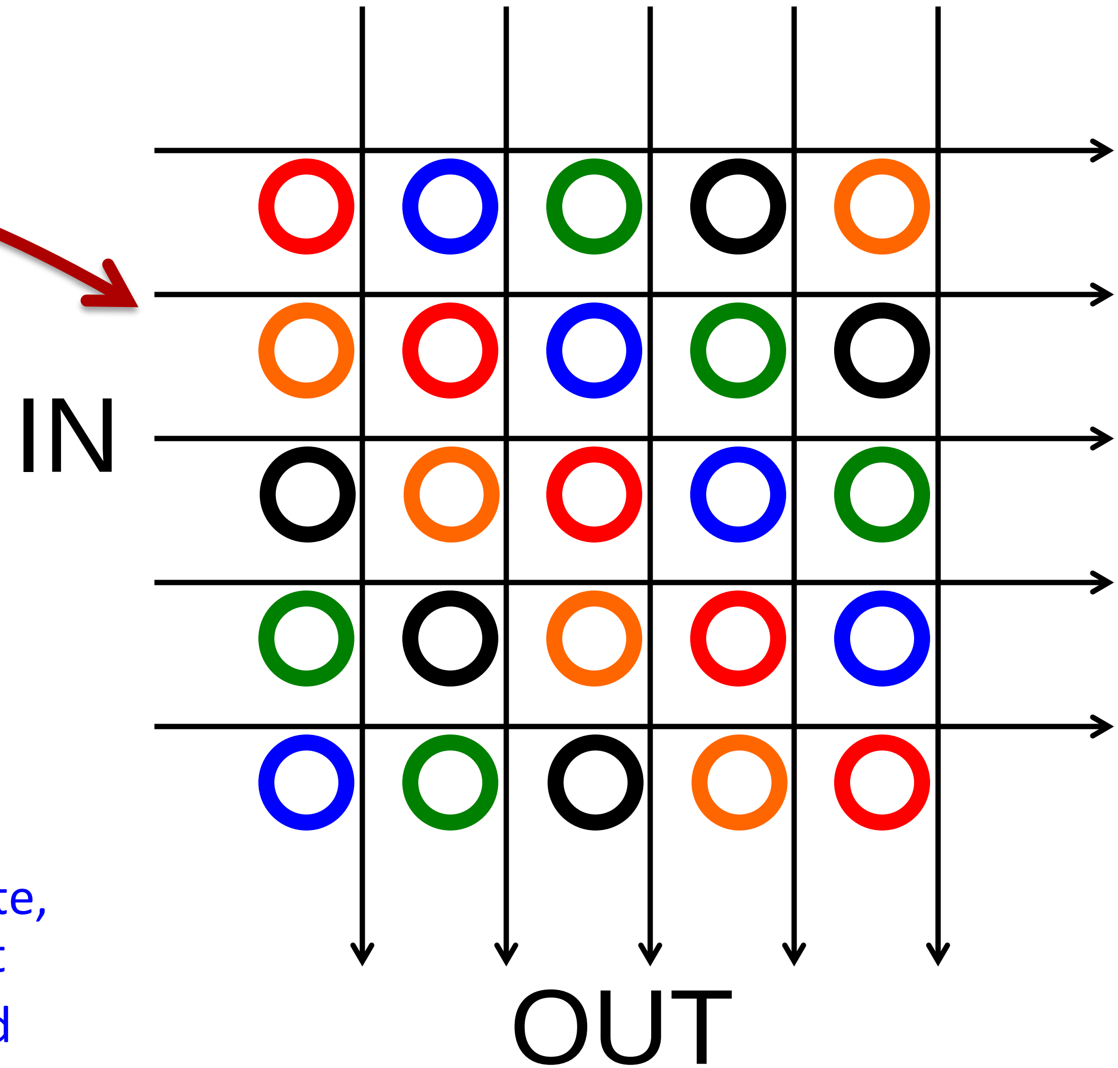

- Chip-scale solution
- Low loss (pass most resonators)
- Impossible with high-data rate, advanced modulation format
- Saves fiber routing compared to parallel interconnect
- Can be ***reconfigurable***

Technology Challenges (40 λ @25 Gbps)

- 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*

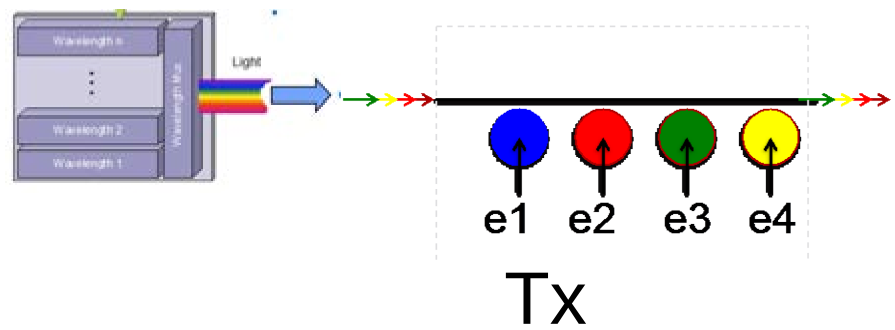

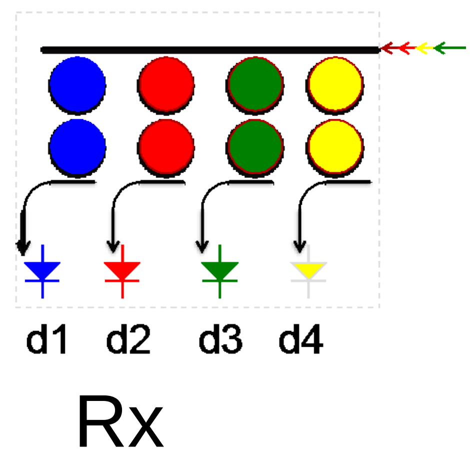

Technology Challenges

- 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*

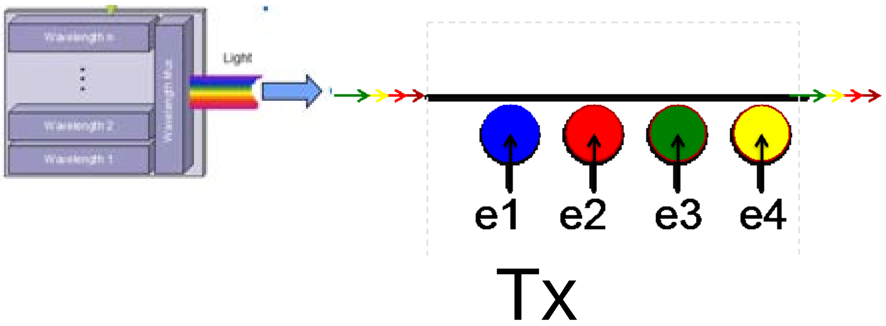

- Silicon Photonics
  - ***Efficient Laser source***
  - *Modulator and optical filter resonant wavelength stability and uniformity*
  - *Filter shape, coupling variations*
  - *Low energy receivers*

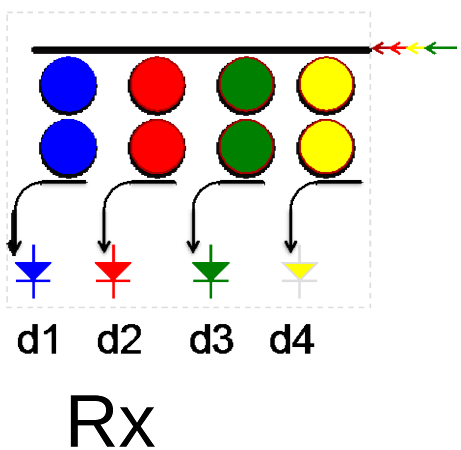

- Interface Electronics
  - *Efficient clock and data recovery*
  - *Data TDM multiplexing (SERDES)*
  - *Efficient Data encoding and error correction*

Laser Source Architecture

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

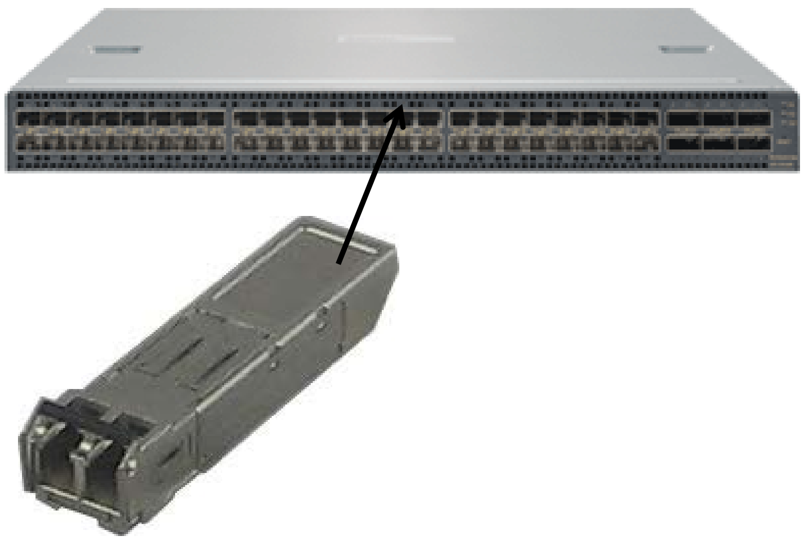

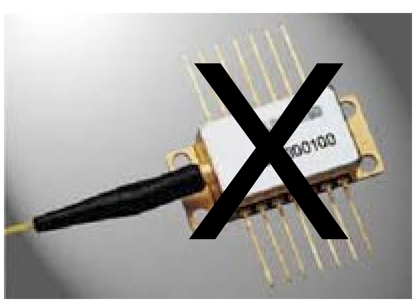

- In the *long term*, routing the laser in an optical waveguide from an on-board laser is analogous to routing electrical power.

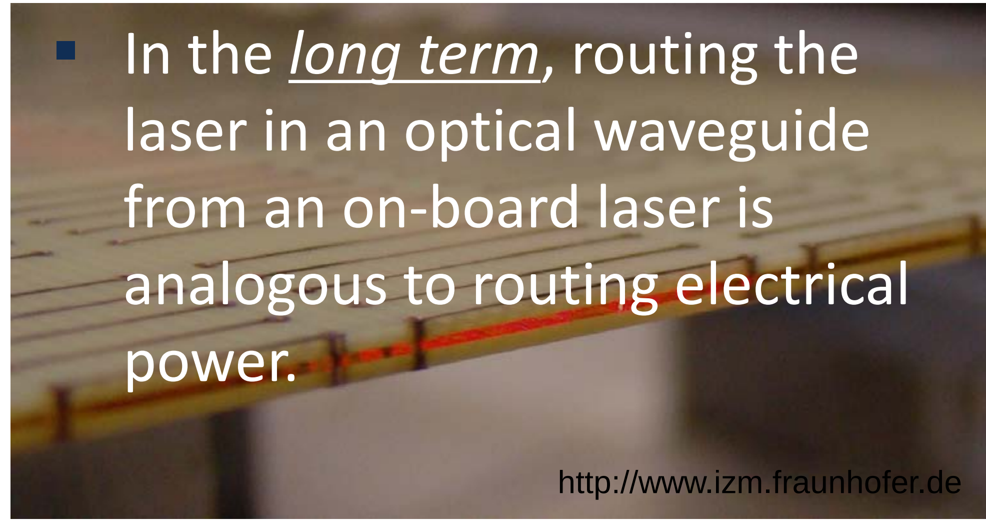

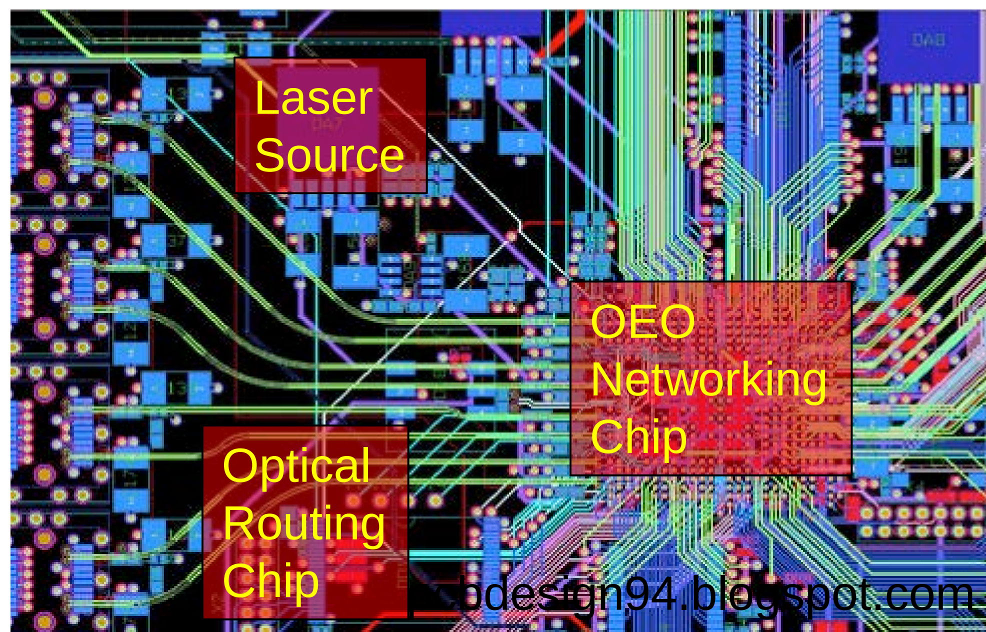

Laser Source Near-term Requirements

- Near term Requirements
  - **40 λ @ 25G = 1 Tbps;**
    - C-band (30 nm @ 100 GHz), 60 nm @ 200 GHz, 116 nm @ 400 GHz)
  - **λ stability** … depends on tracking … (not limited to ITU grid).
  - **Power per channel:** ~ 200 µW (limited by power into resonant modulators)
    - There is a tradeoff between power handling and modulation voltage (Q)
    - This low power presents a challenge for the receiver!
  - **Wall-plug efficiency:** 1% →**0.8W**/40λ (0.8 pJ/bit)
- Cost is most important!!
  - $(40λ) < $(few λ) today, we win!
  - Otherwise, parallel fibers wins.

Laser Source Technology Options

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

*** = best

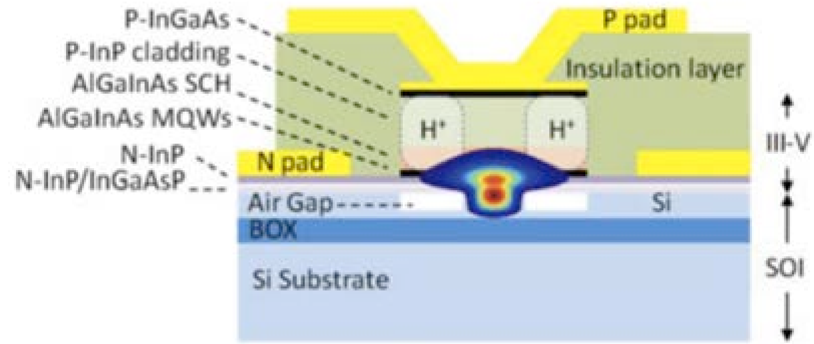

Heterogeneous: Keck et. al., JSTQE 2013

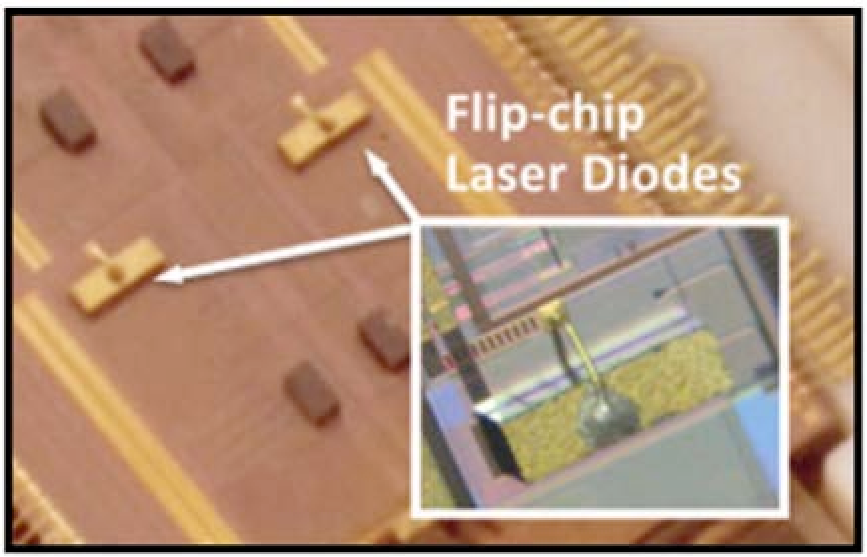

Flip-chip: Dobbelaere ECOC 2014

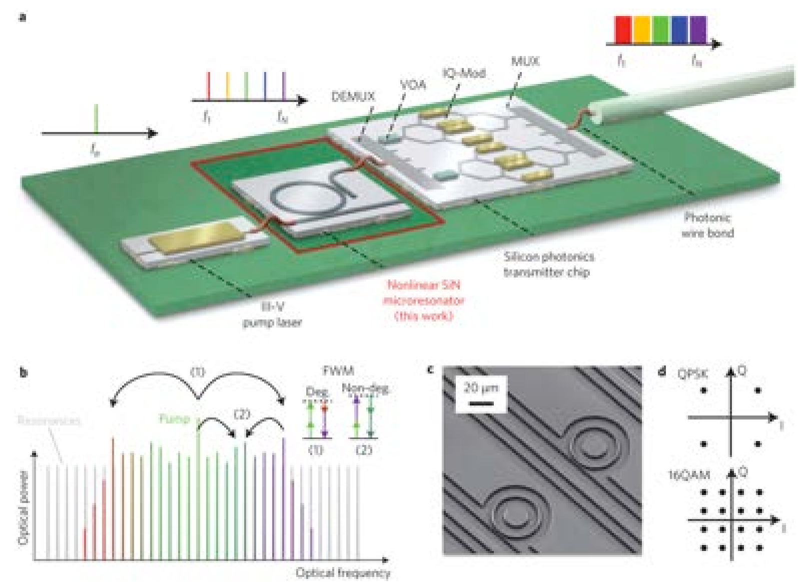

Pfeifle et. al., Nature 8 (2013)

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*

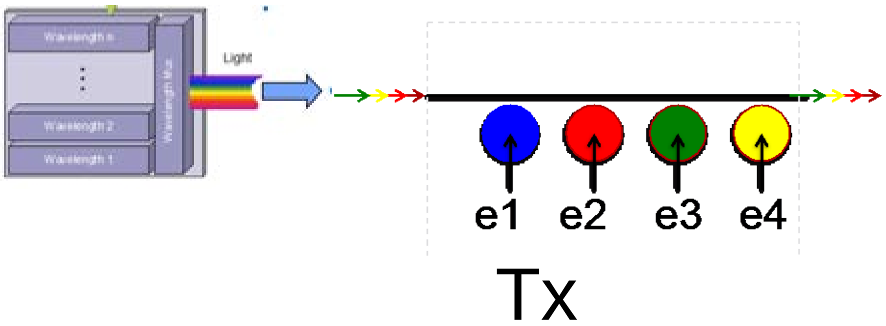

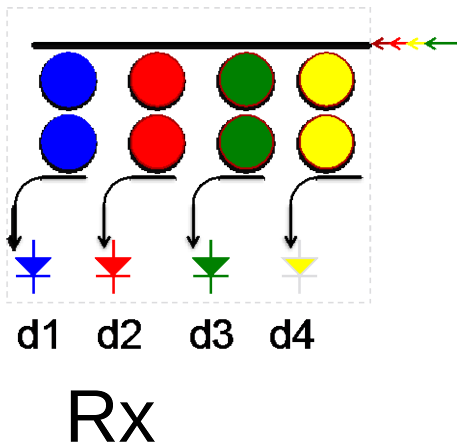

Effect of temperature on loss budget

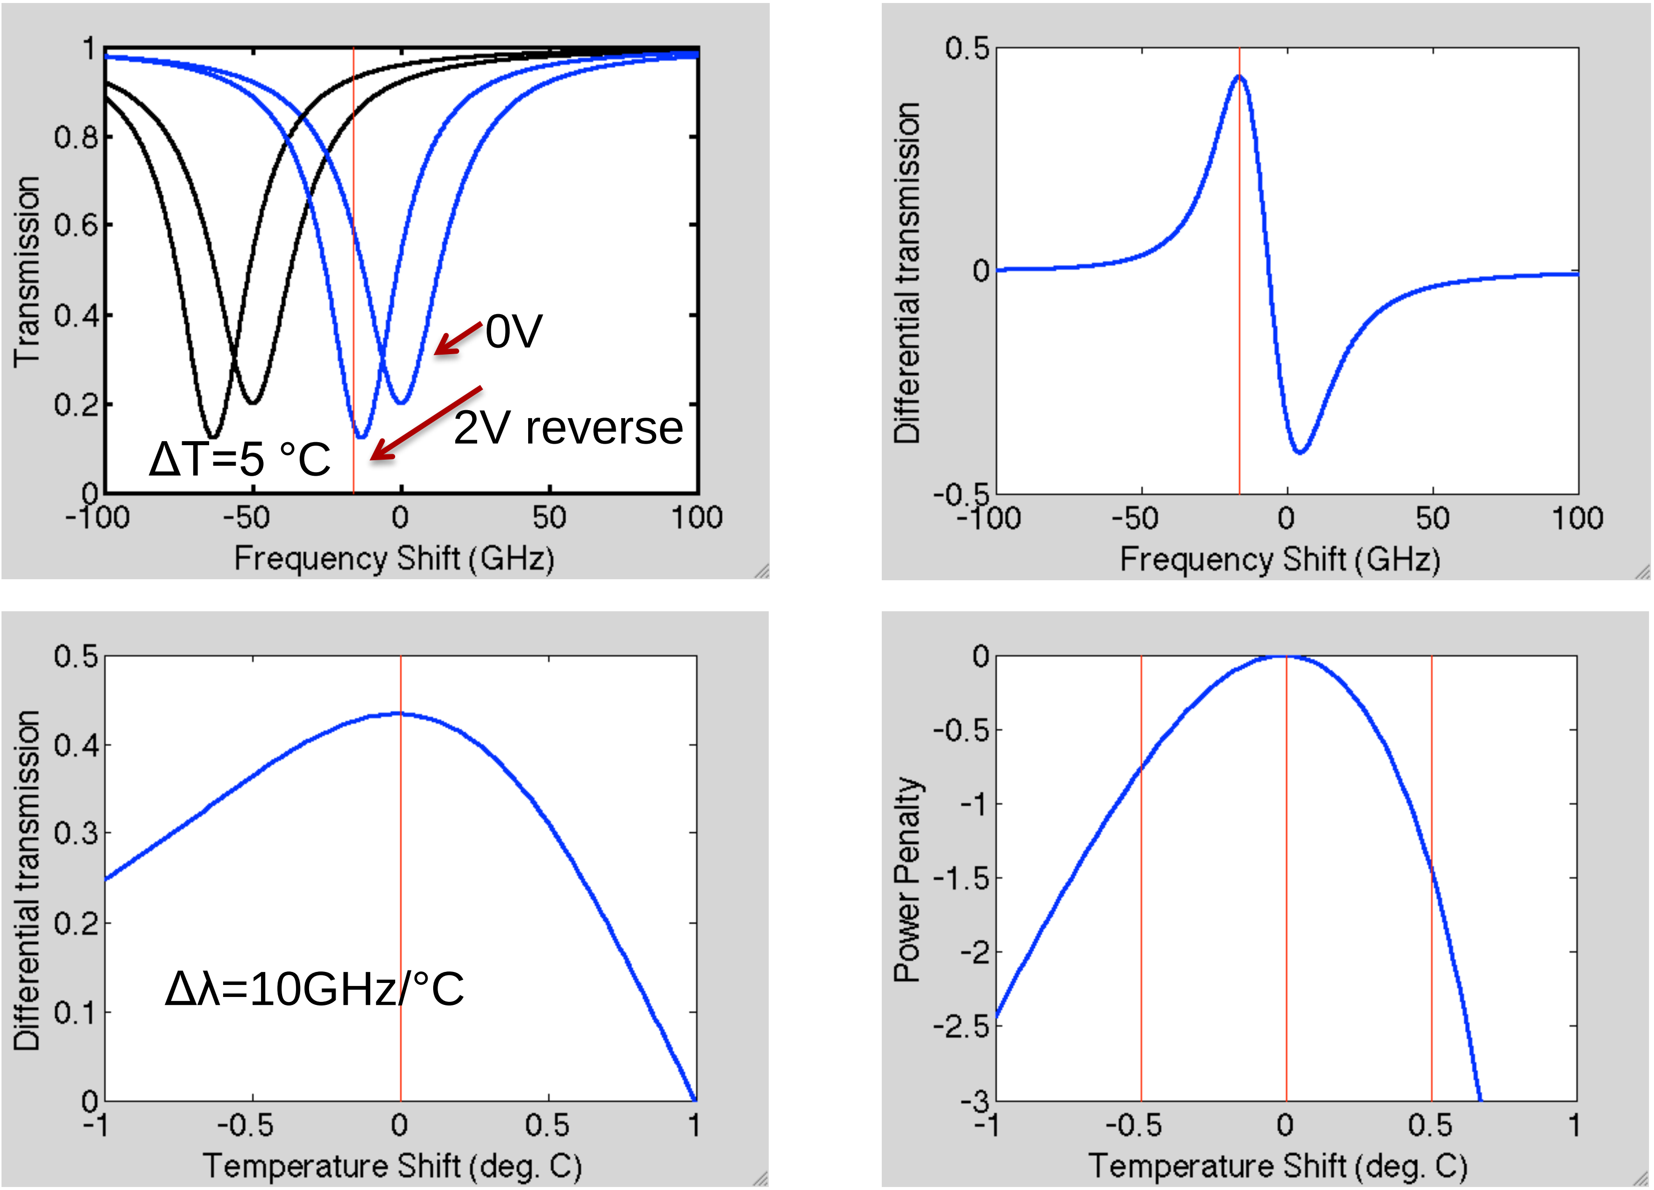

Filter allowable temperature drift.

Sandia National Laboratories

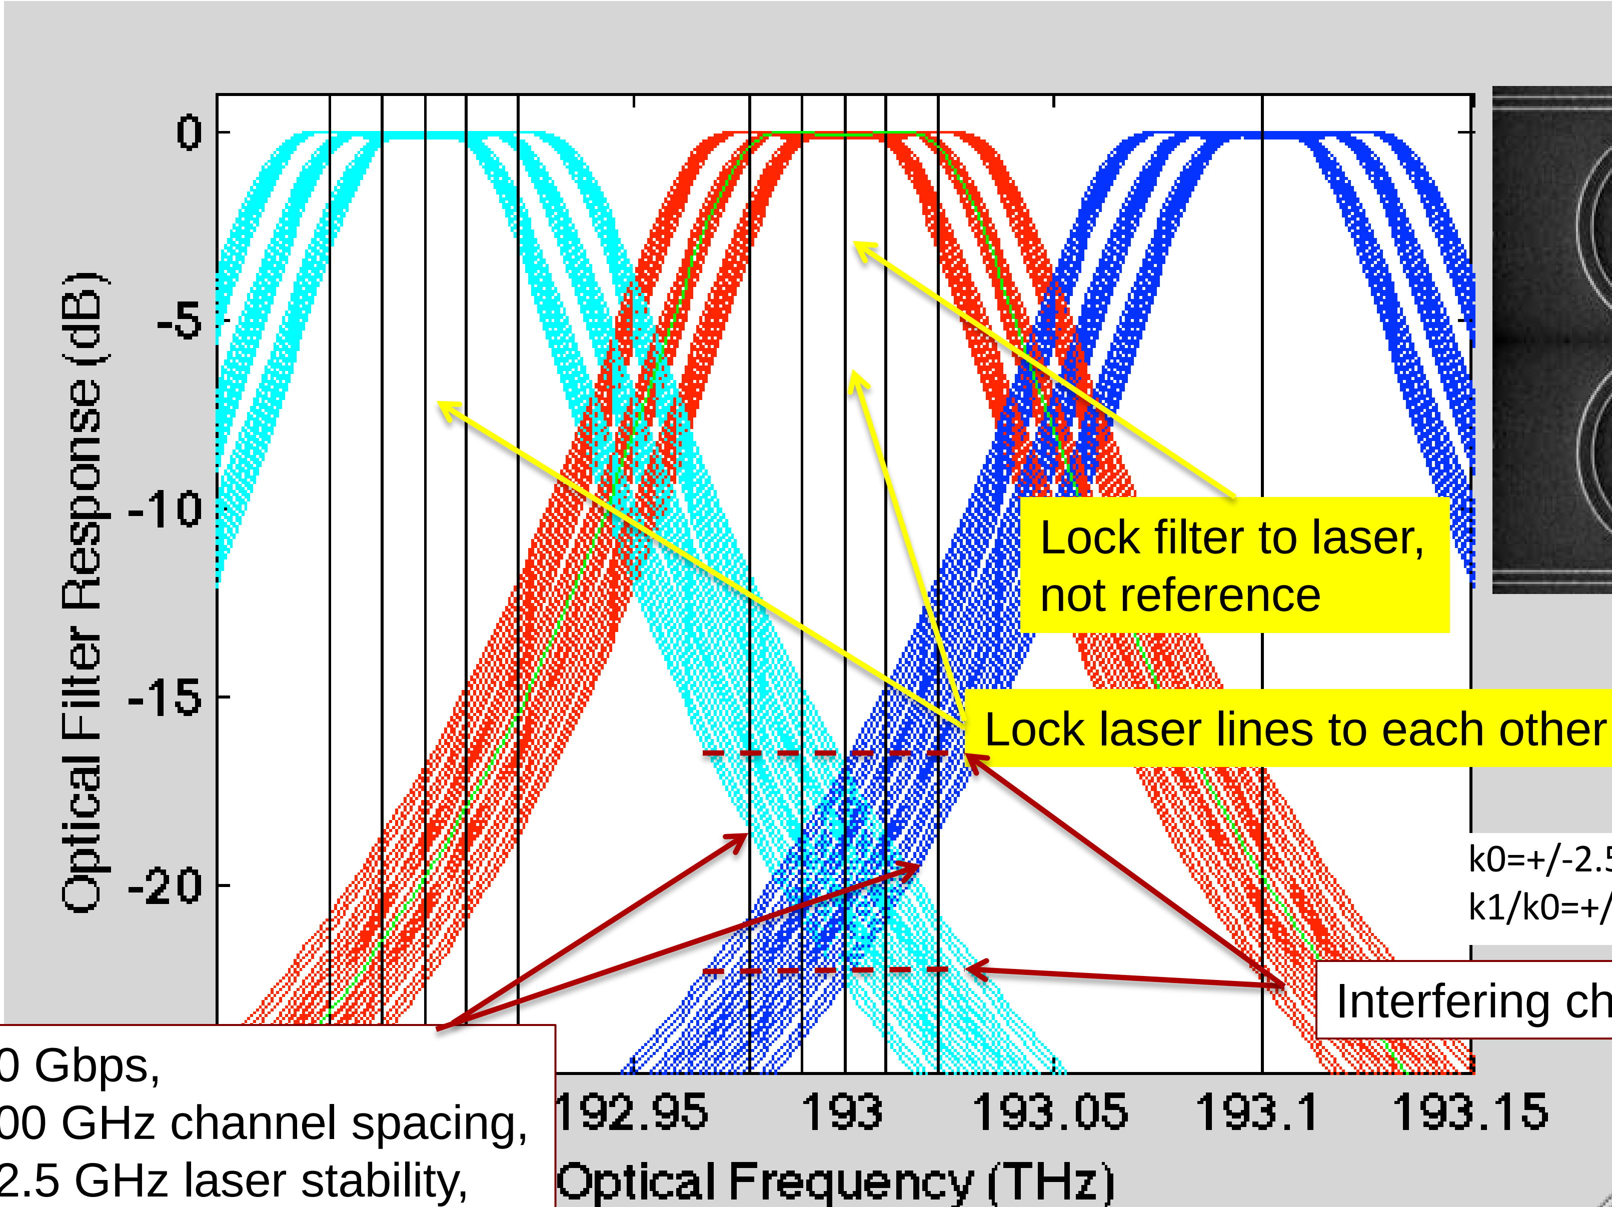

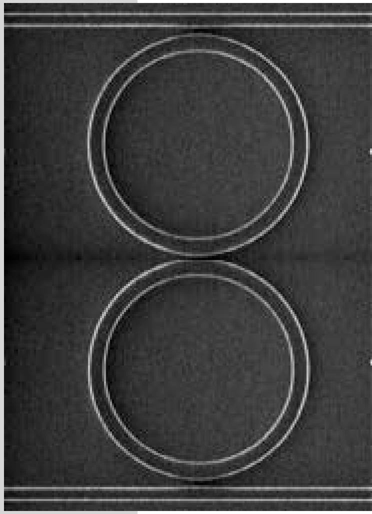

Lock filter to laser, not reference

Lock laser lines to each other

Interfering channels

k0=+/-2.5%
k1/k0=+/-2.5%

10 Gbps,
100 GHz channel spacing,
12.5 GHz laser stability,

Calculations based on Little et. al. JLT, 1997

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

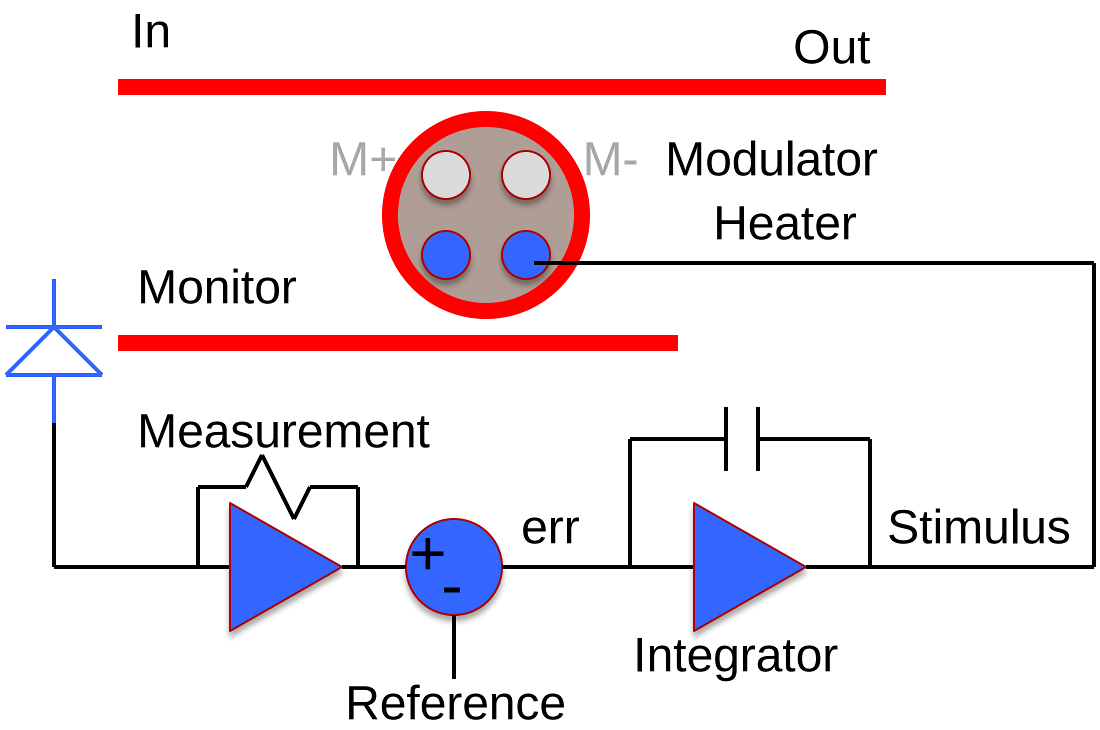

Resonant Wavelength Locking

Sandia National Laboratories

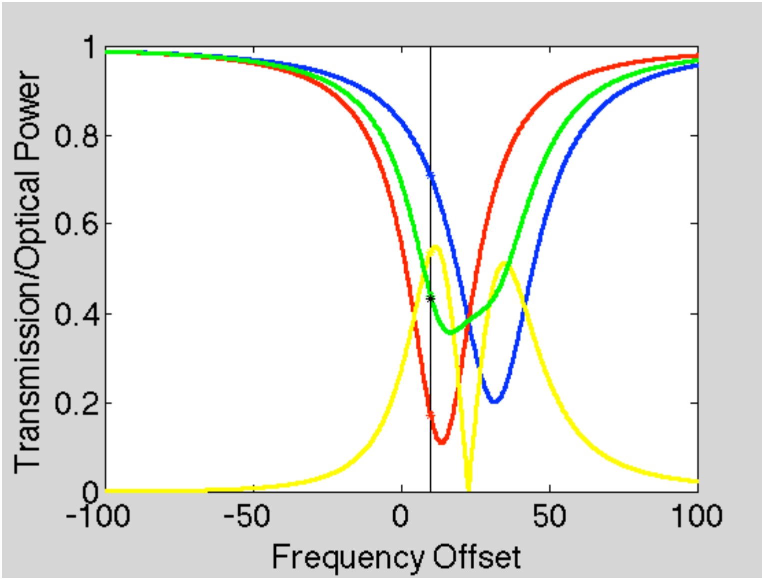

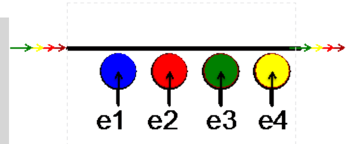

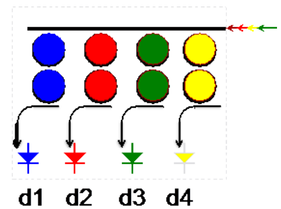

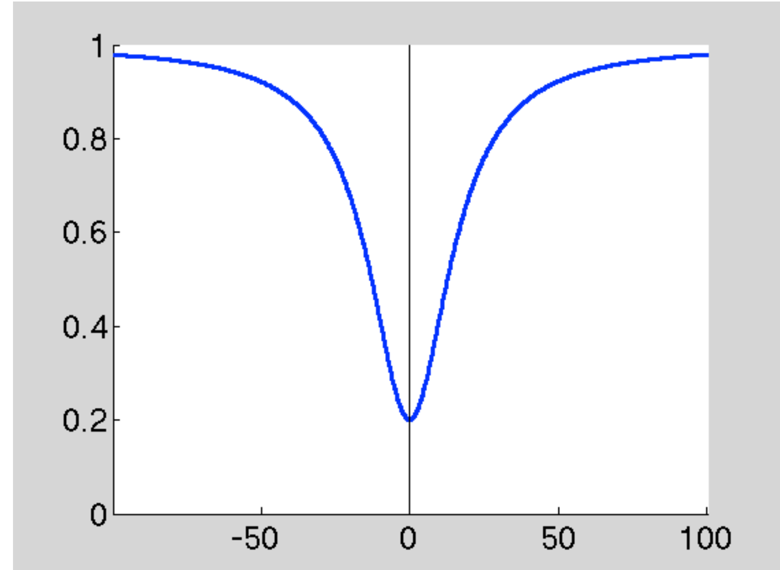

Modulator

- Lock on side of resonance

Filter (DeMux)

- Lock at minimum power

Locking using Power Sensors (MIT, Columbia, Rice, Oracle)

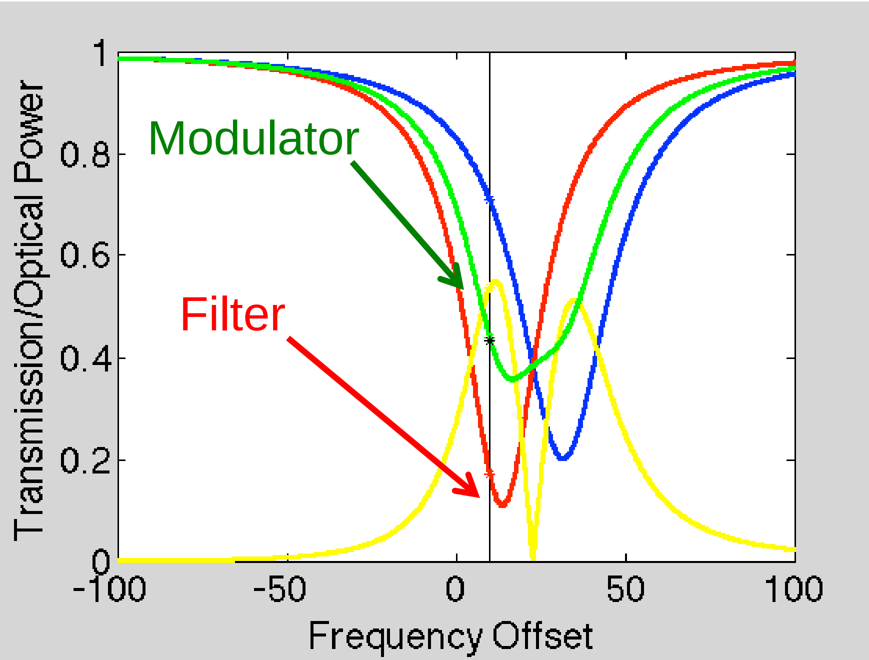

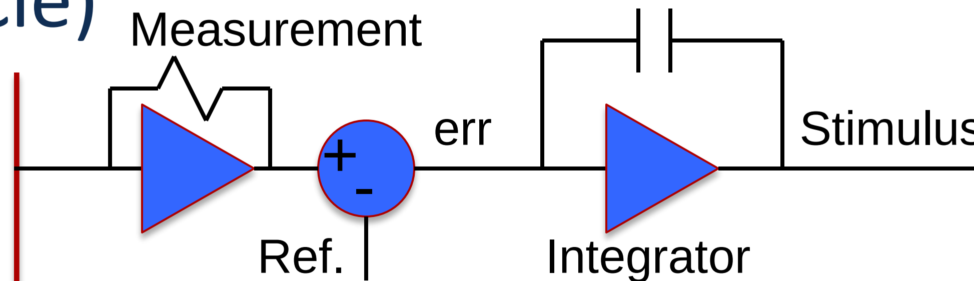

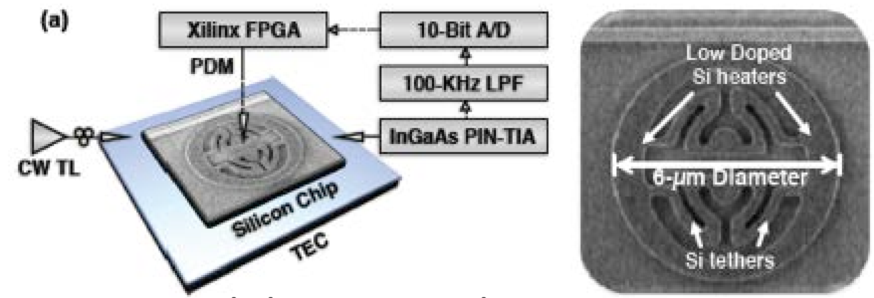

Resonator with heater, without sensor (Timurdogan et. al., CLEO 2012)

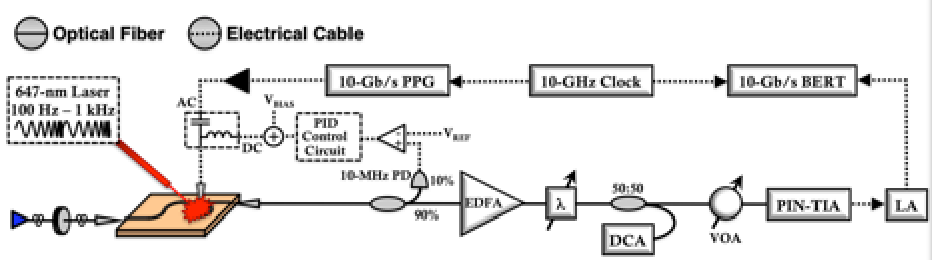

Modulator with bias induced temperature change (Padmaraju, OFC 2012)

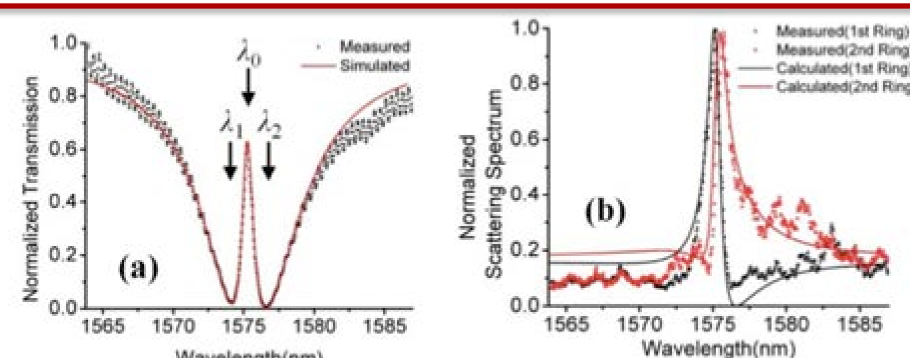

Scattering from a dual-ring modulator (Qui et. al. Opt. Exp., 2011)

Locking using a dither signal (Columbia)

- Creates a signal that is anti-symmetric (lock at zero)
- More complex electrically
- Simple optically
- Some small degradation in the optical performance with dither
- Best for filter locking

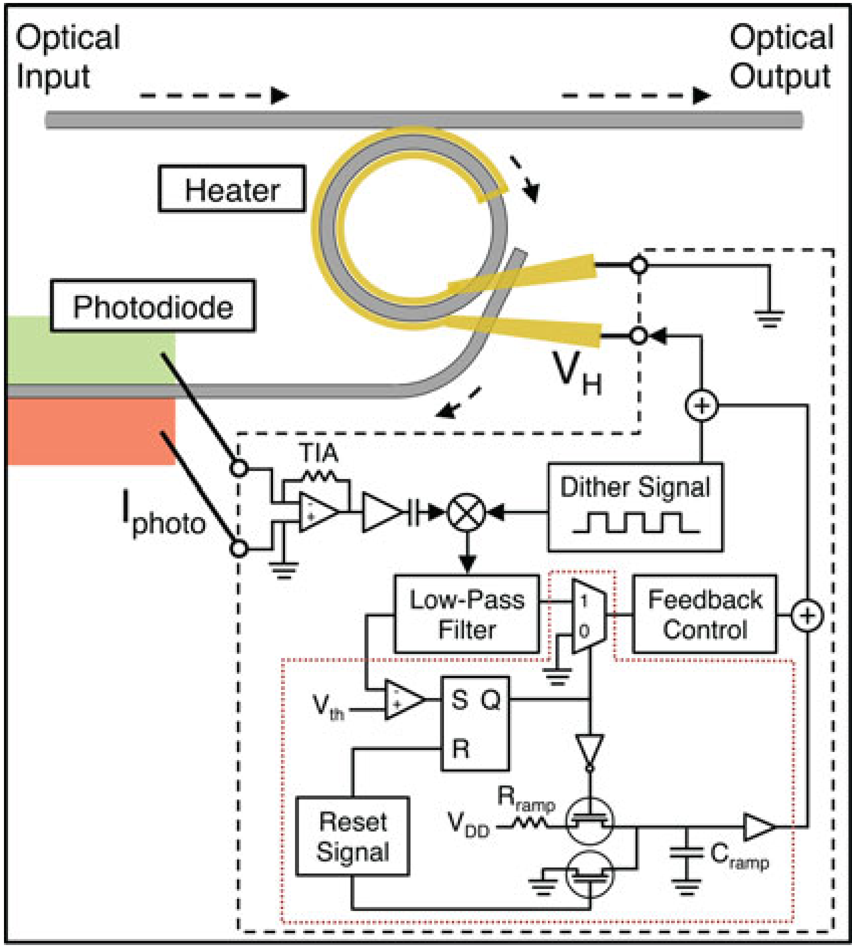

K. Padmaraju, et. al. JLT 32 (3) (2014)

Modulator wavelength stabilization using bit errors (Sandia)

- Direct measurement of the bit errors
- Requires high speed circuitry
- Most compact solution (no low pass filtering)

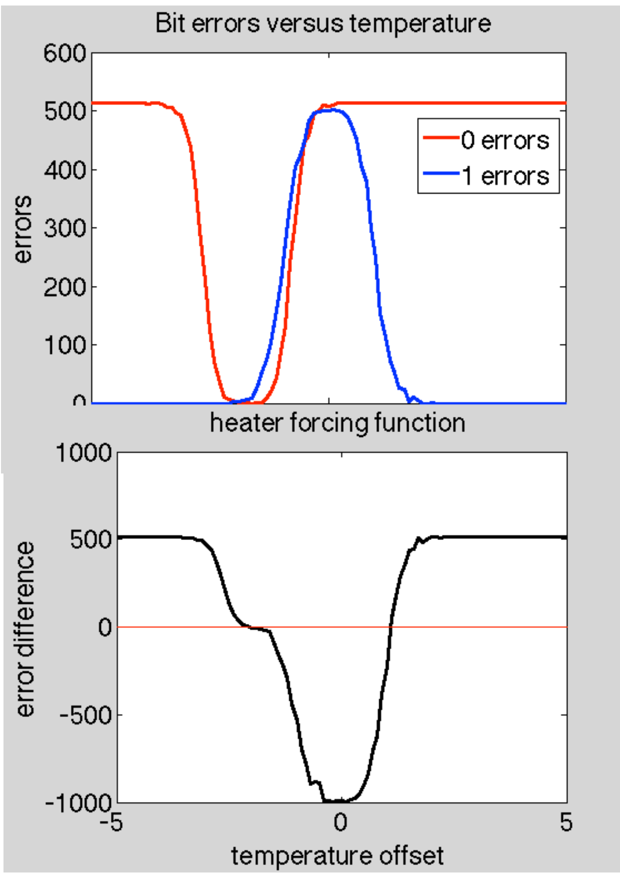

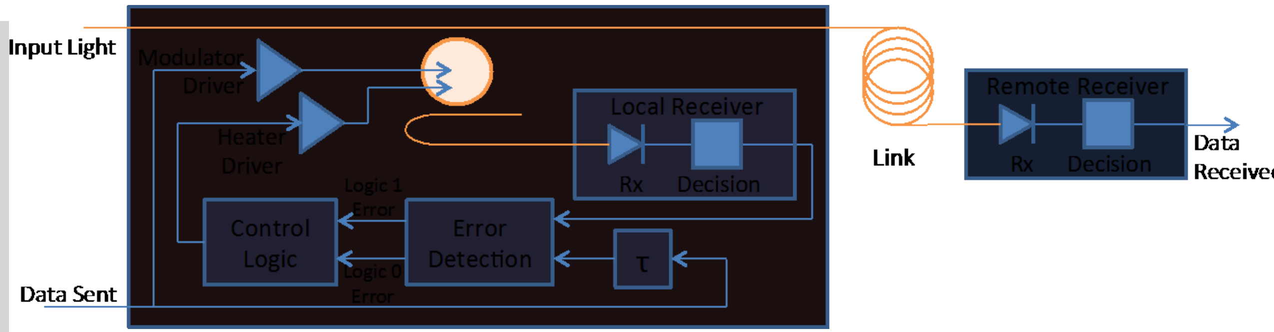

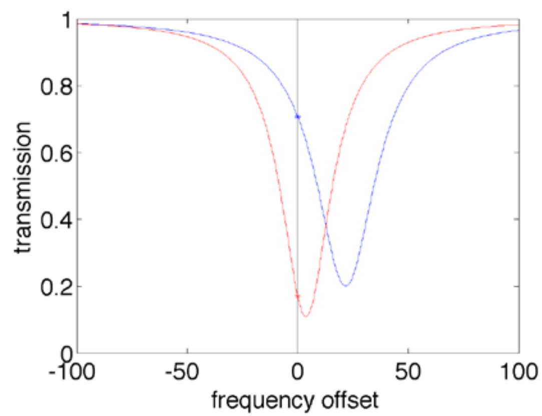

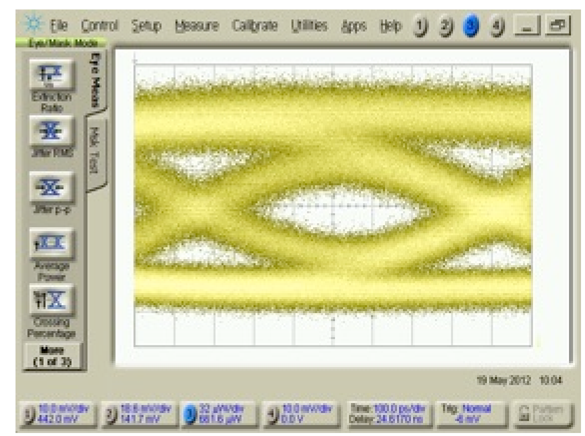

W. A. Zortman et. al.,IEEE Micro (2012)

Balanced Homodyne Detection

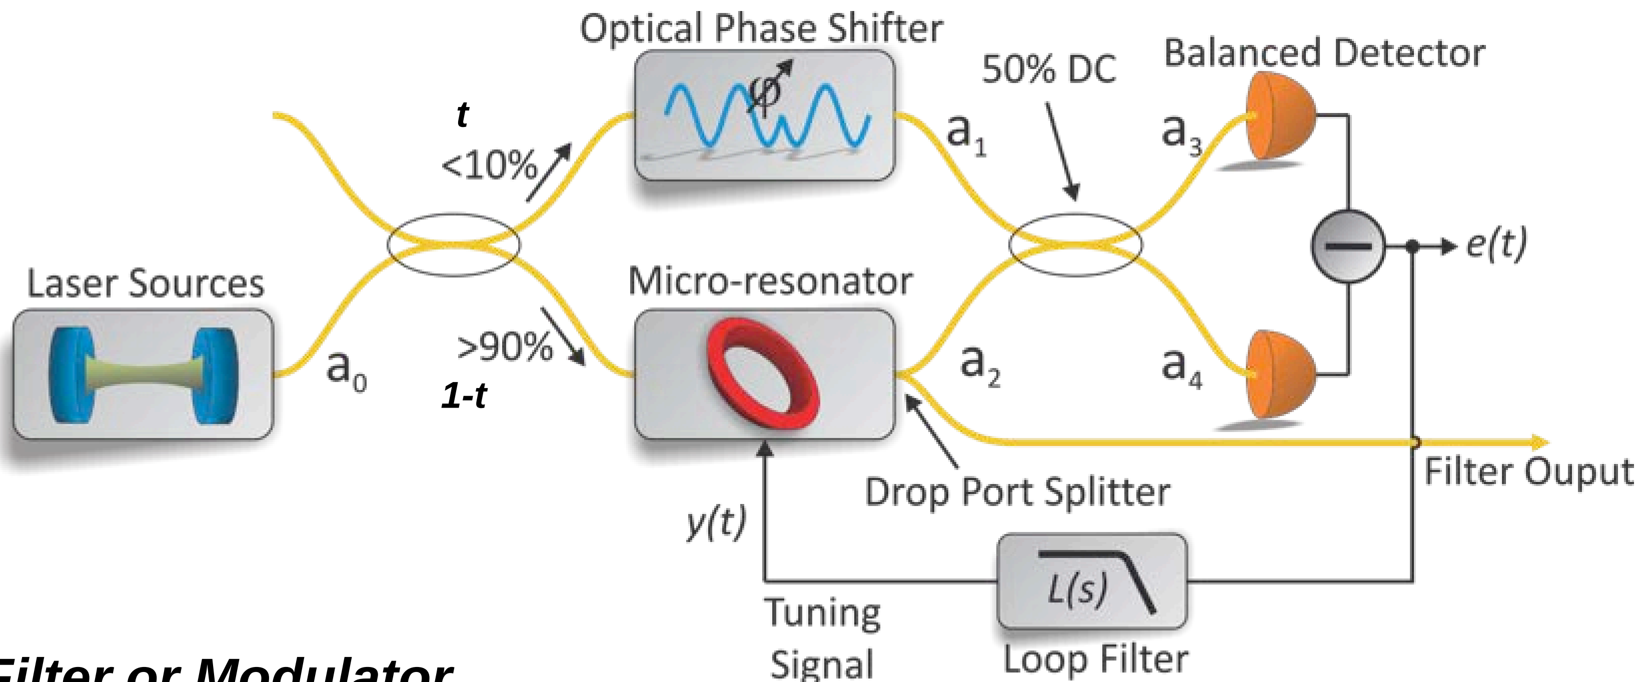

- ***Filter or Modulator***
- ***Lock to zero:*** No calibration or reference level needed for locking
- ***Amplitude insensitive:*** Locking point not influenced by optical intensity
- ***Precision locking:*** Resonator is not disturbed
- ***Minimum circuit complexity:*** Power and area consumption of control electronics is minimized

J.A. Cox, A.L. Lentine, D.C. Trotter and A.L. Starbuck, "Control of integrated micro-resonator wavelength via balanced homodyne locking," *Opt. Express* Vol. 22(9) (2014)

Locking multiple channels: 40-λ

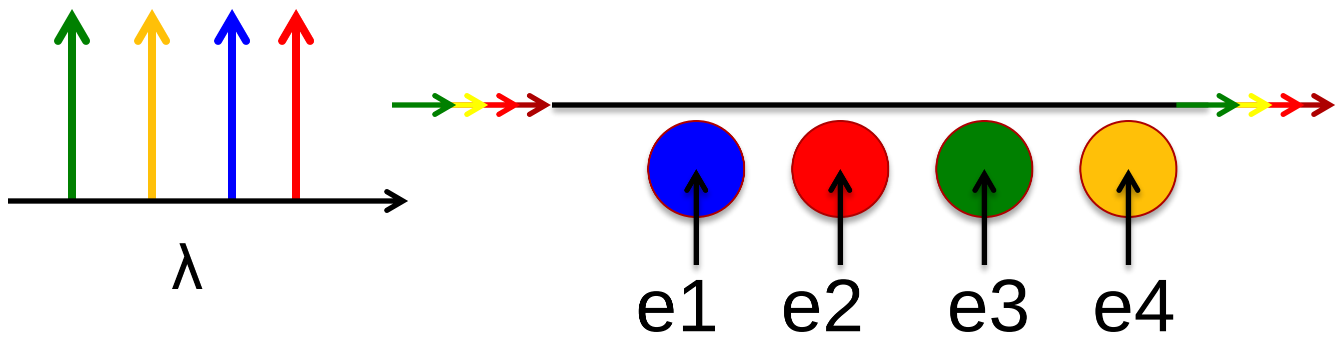

- 40 independent control loops
- Must ensure each resonator locks to an independent line
  - Not necessarily ordered
- Maintain lock while operational (training periods ≠ good)
- Independent locking … don't mess up the others.

Locking multiple channels

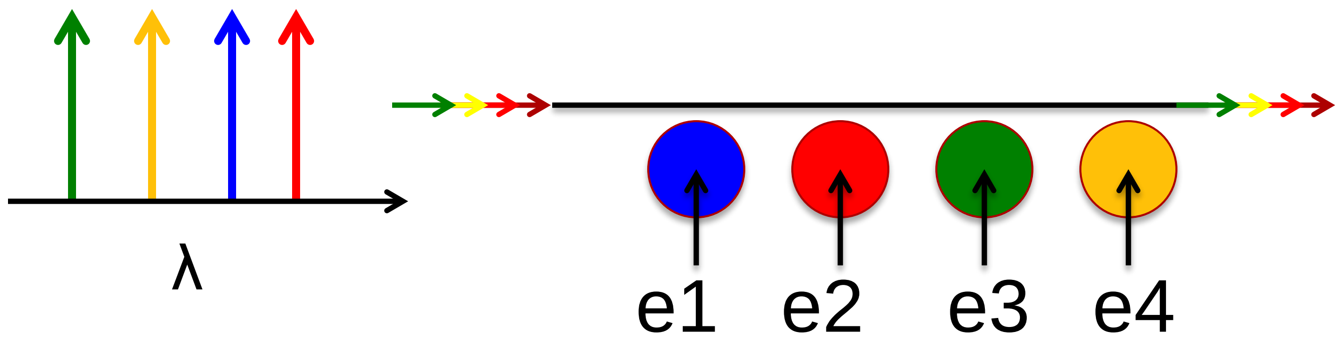

- Must have distinguishing characteristics (filter mechanism)
  - Don't see how a simple power measurement will work, especially on modulators
- Orthogonal frequency control
  - Dither, Pound-Drever-Hall
- Time sequential control
  - BER (code), Balanced Homodyne detection (tuned interferometer)
- Other techniques

Locking second order filters

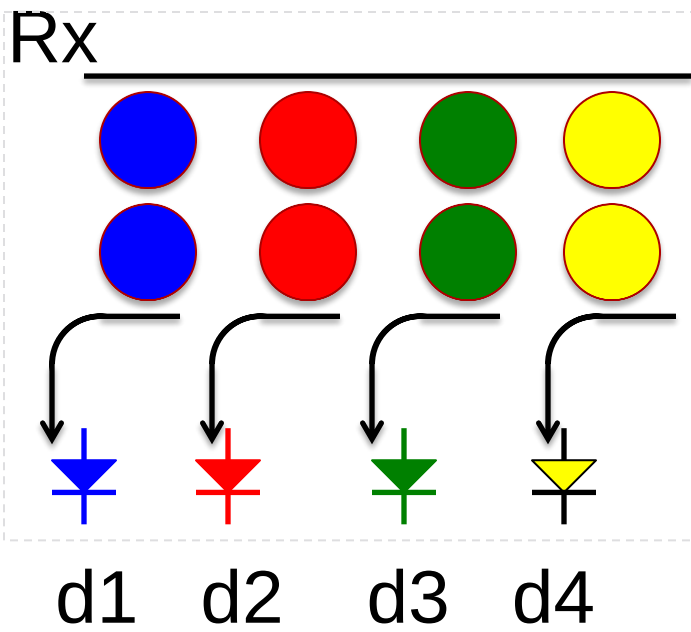

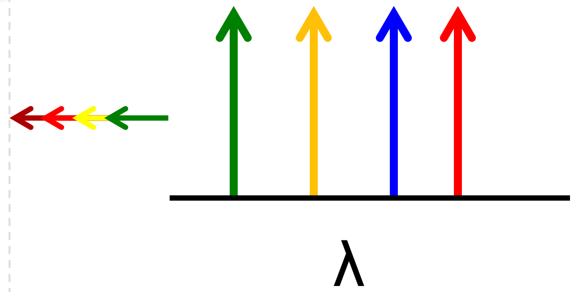

- Same problem as modulators and first order filters PLUS
- How to sense 'errors' in each ring ?
- How to tune them independently ?
  - Slight I-V variations

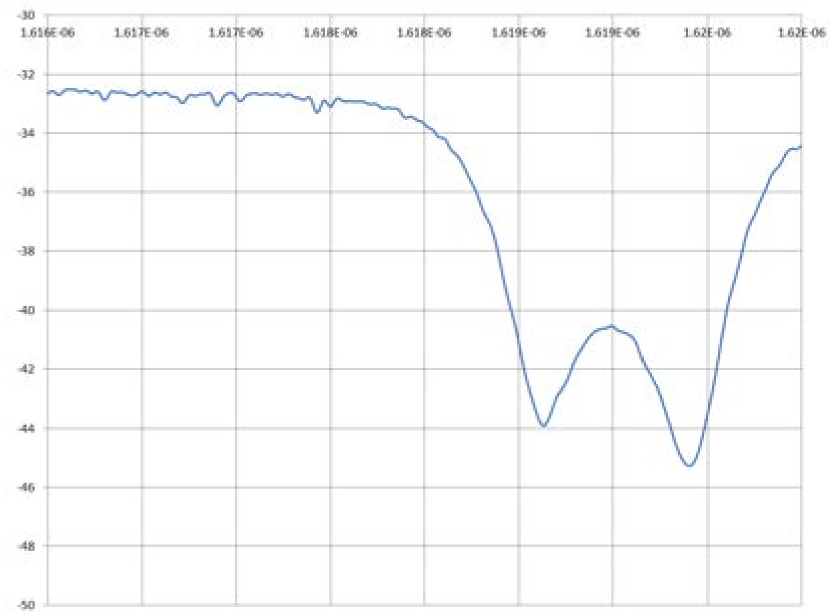

λ Stability: ITU grid or ΔT?

- 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

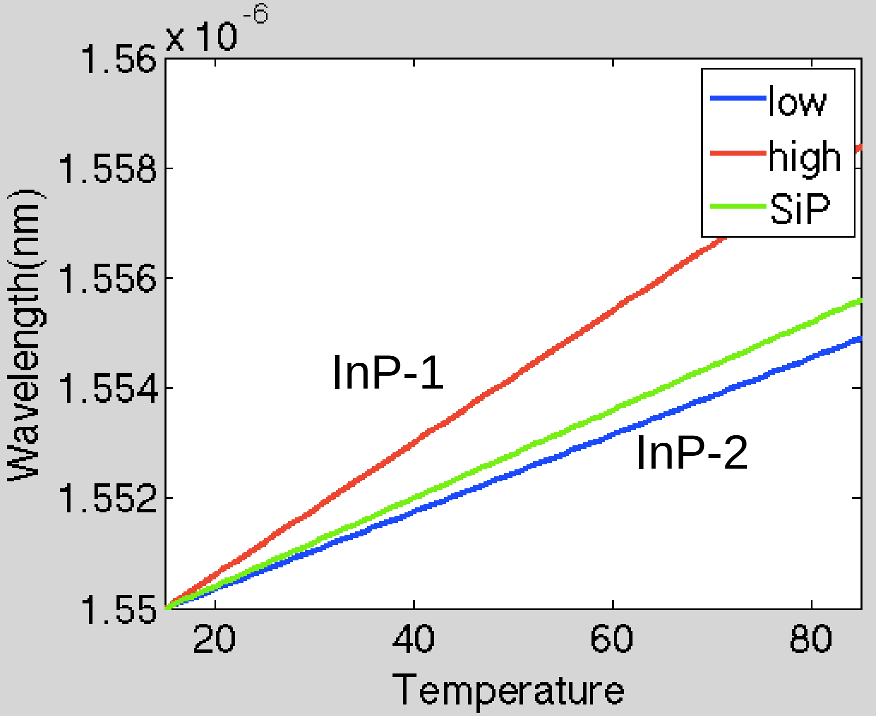

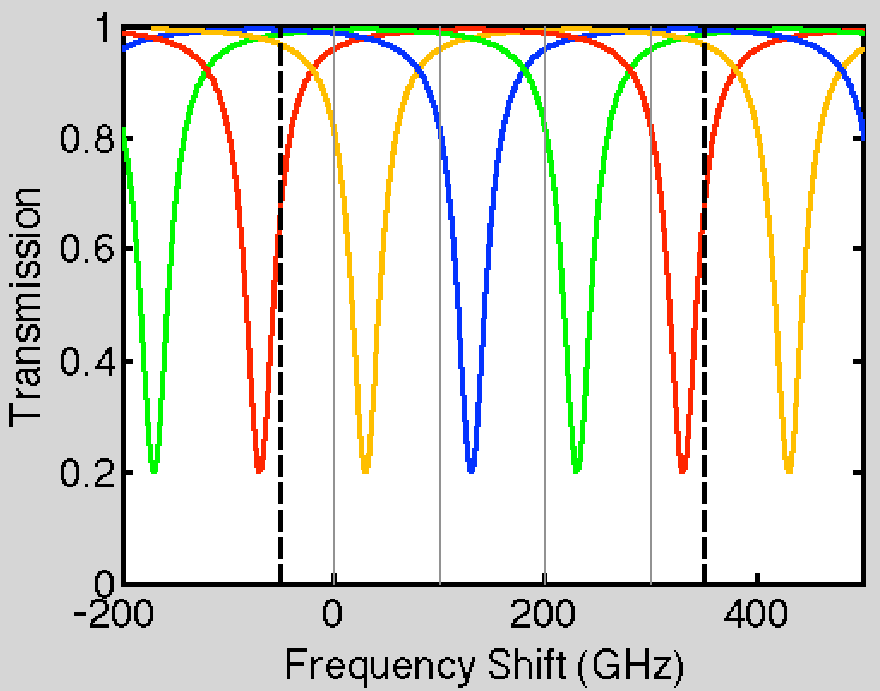

Technology Challenges

- 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*

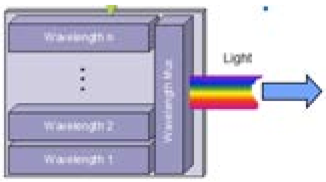

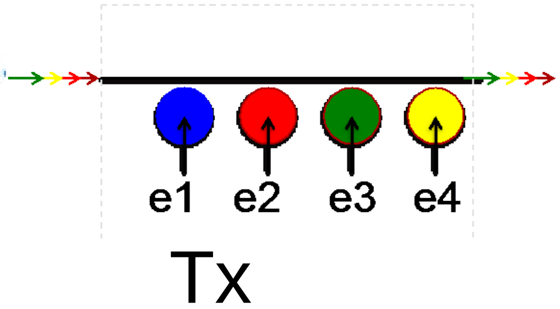

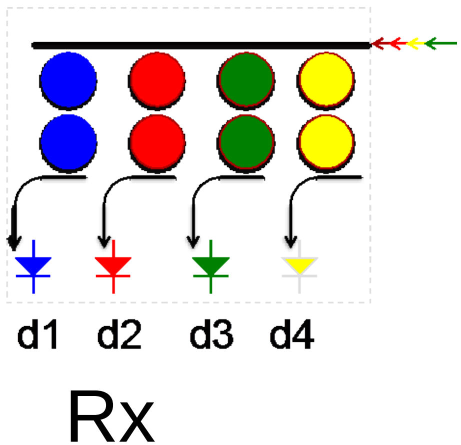

High Transimpedance Receivers

Sandia National Laboratories

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

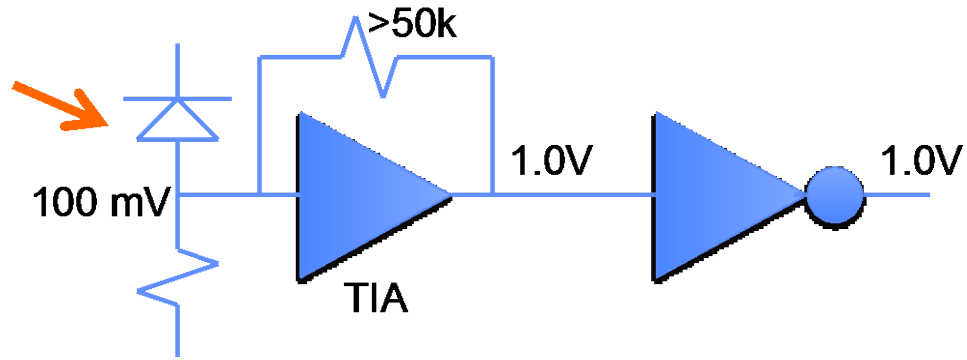

Low capacitance, high transimpedance gain
Lower noise floor, good sensitivity
Might be DC offset limited vs. noise limited

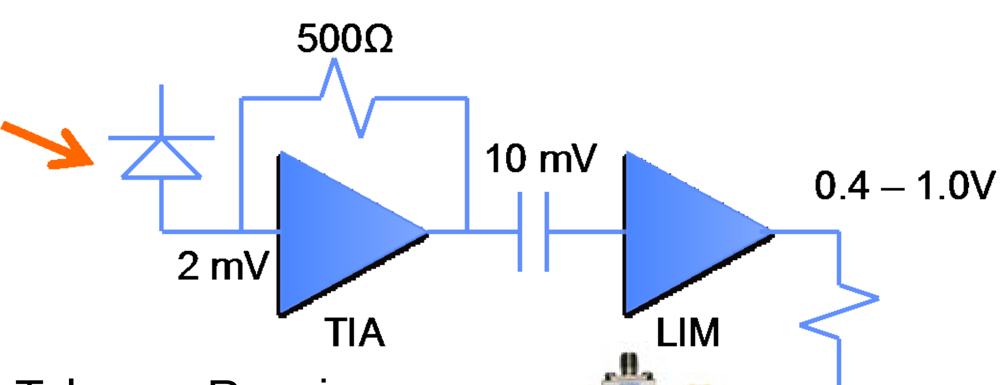

Telecom Receiver:
Multi-mode data com Rx:
Too much power, delay

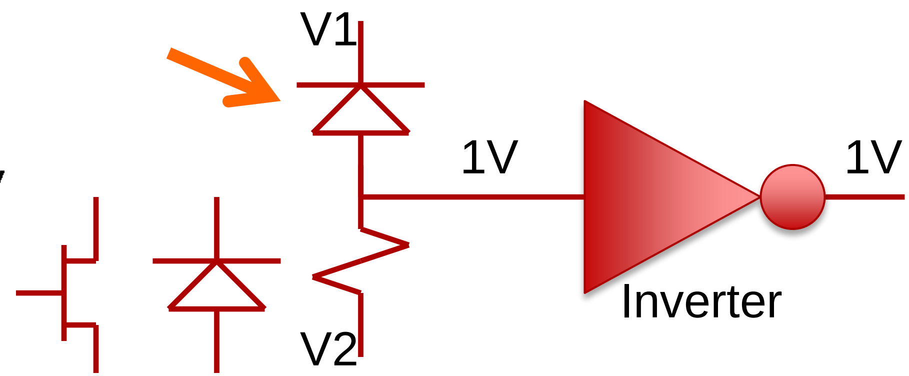

Very low energy, poor sensitivity

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)

Differential Optical Signaling?

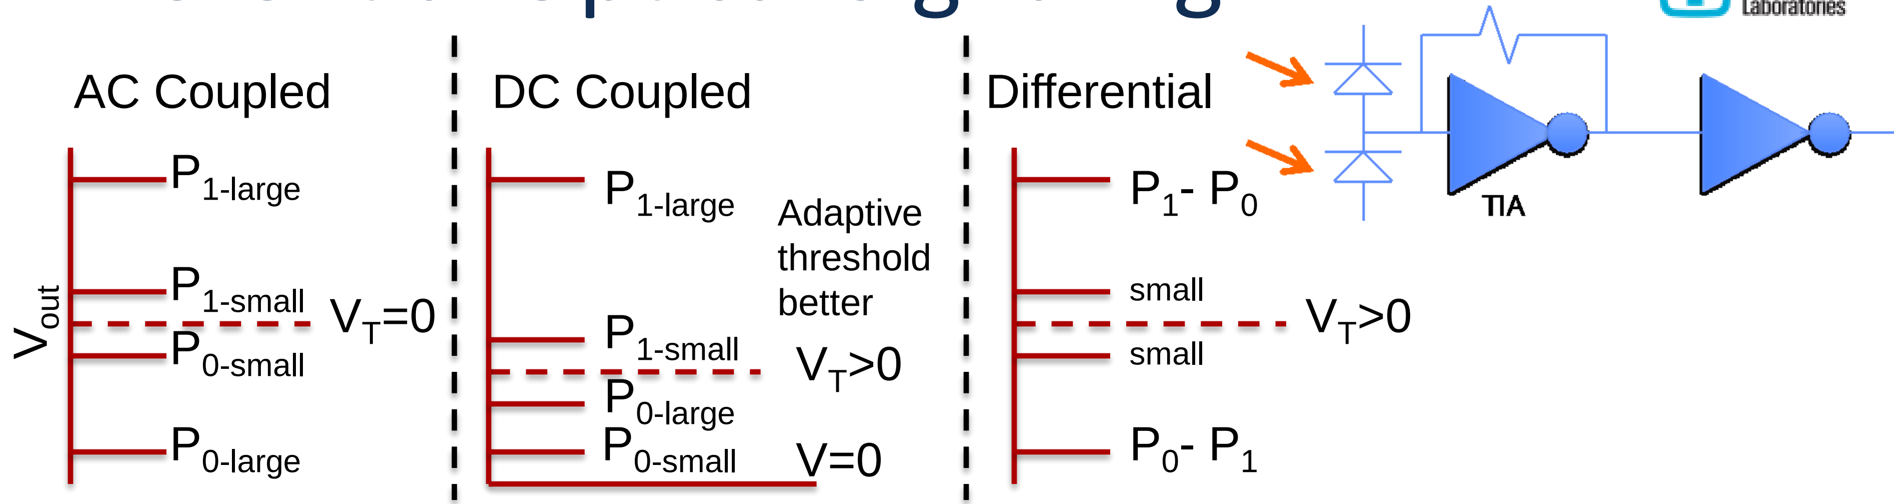

Avoid absolute optical power levels for logic 1 and logic 0

- 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

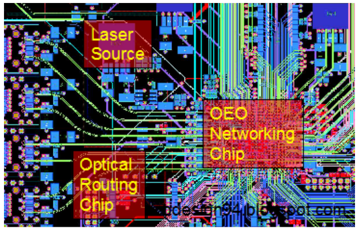

See e. g. Lentine and Miller, JQE 1993;

Receiver Energies (Simulation)

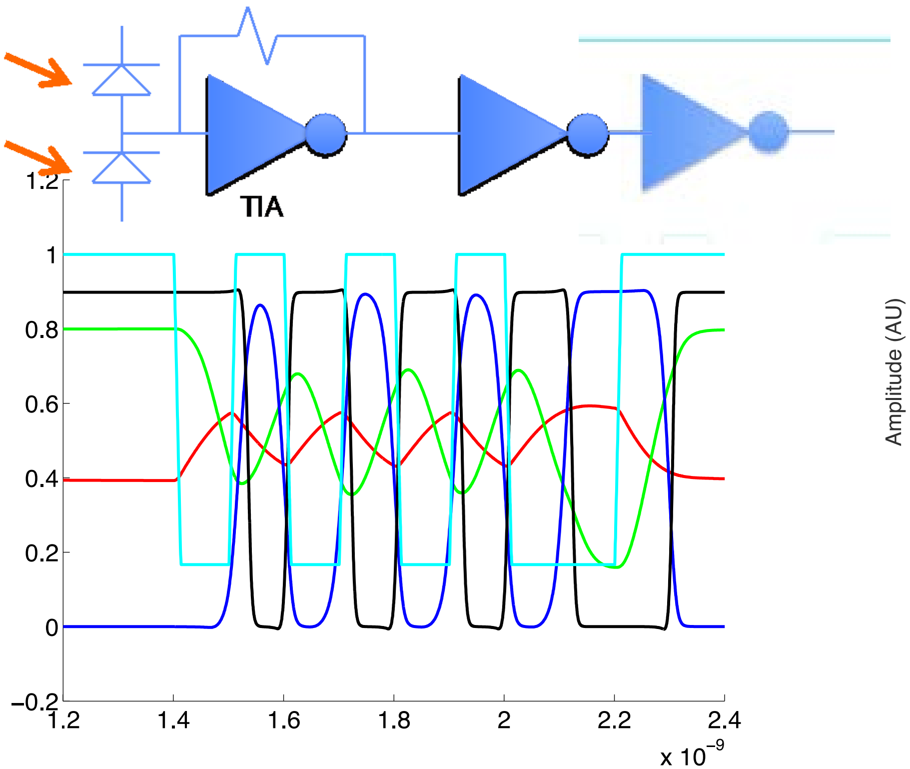

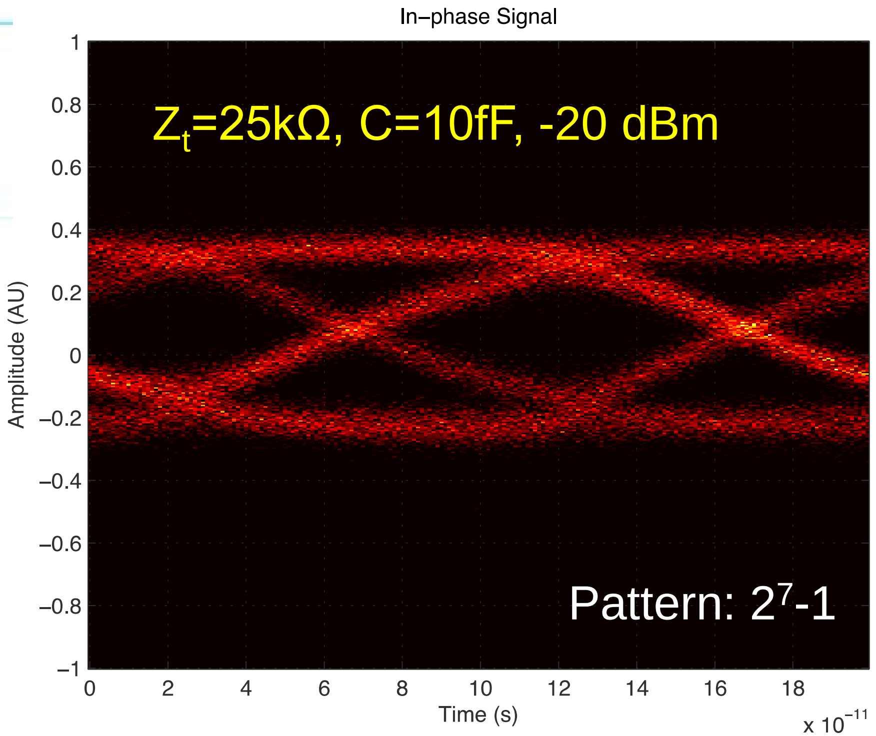

Bit Rate = 10.00Gbps, *Energy per bit = 2.74fJ*

- 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?

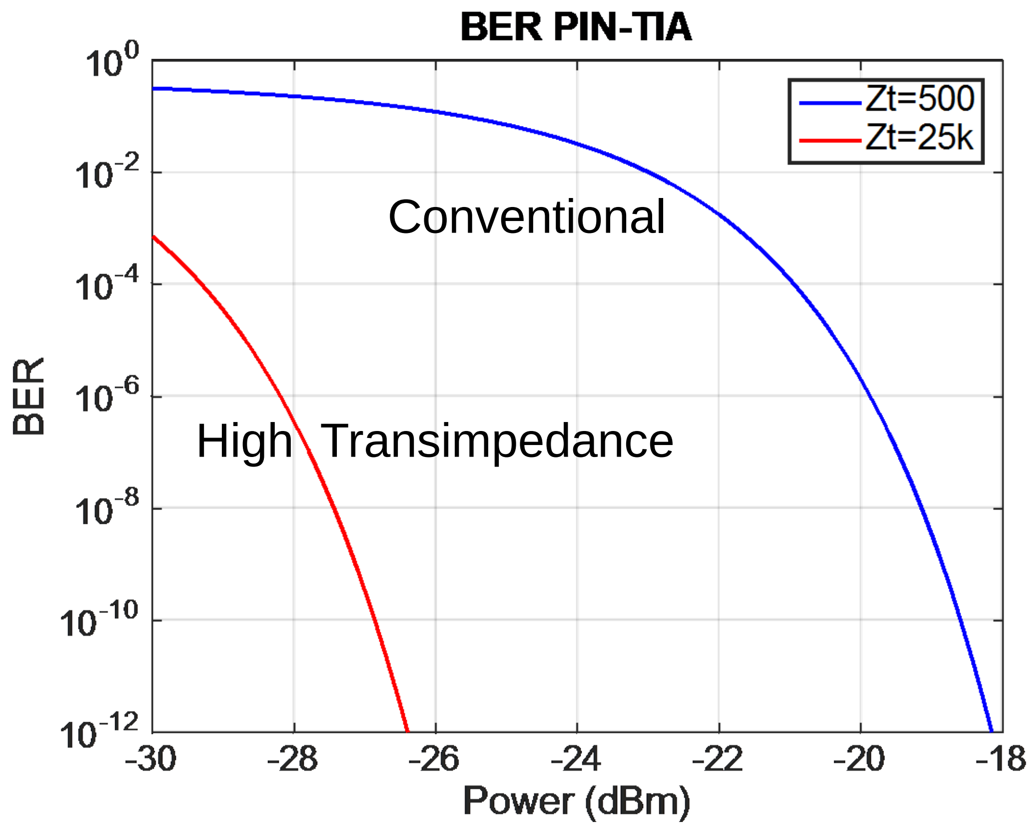

Noise Limited Performance

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*

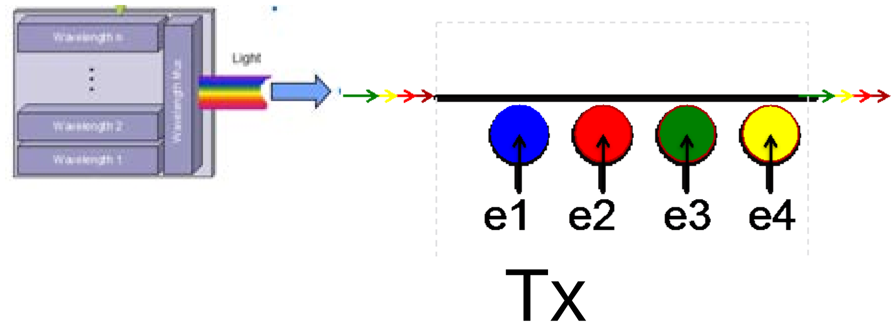

- Silicon Photonics
  - *Efficient Laser source*
  - *Modulator and optical filter resonant wavelength stability and uniformity*
  - *Filter shape, coupling variations*
  - *Low energy receivers*

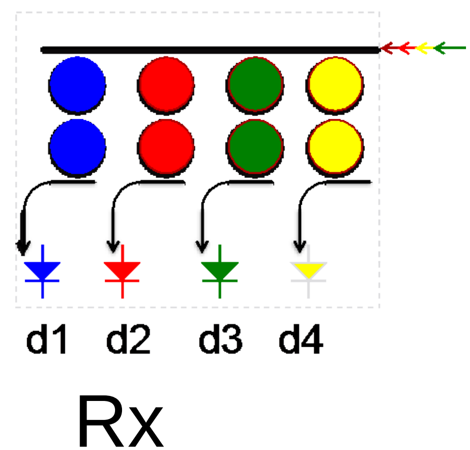

- Interface Electronics
  - ***Efficient clock and data recovery***
  - ***Data TDM multiplexing (SERDES)***
  - ***Efficient Data encoding and error correction***

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 ?

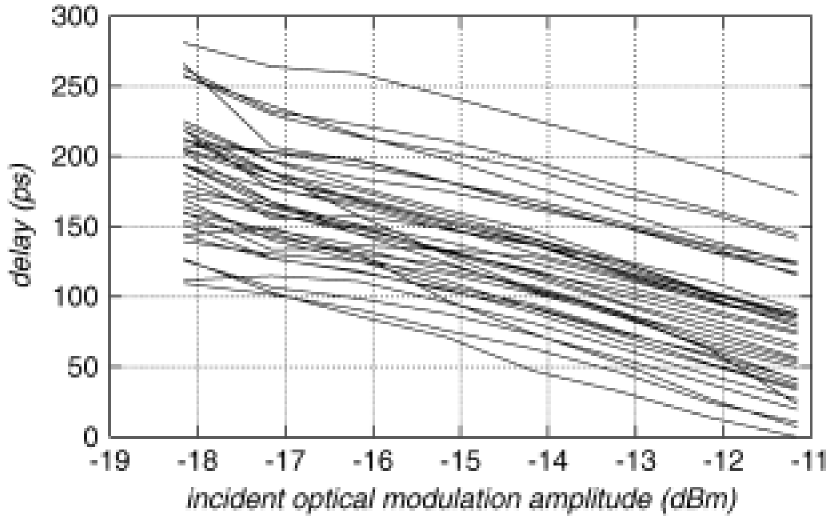

Wilde, Rits, Baets, Van Campenout
64 ch. VCSEL Links, JLT 2008

- **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

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

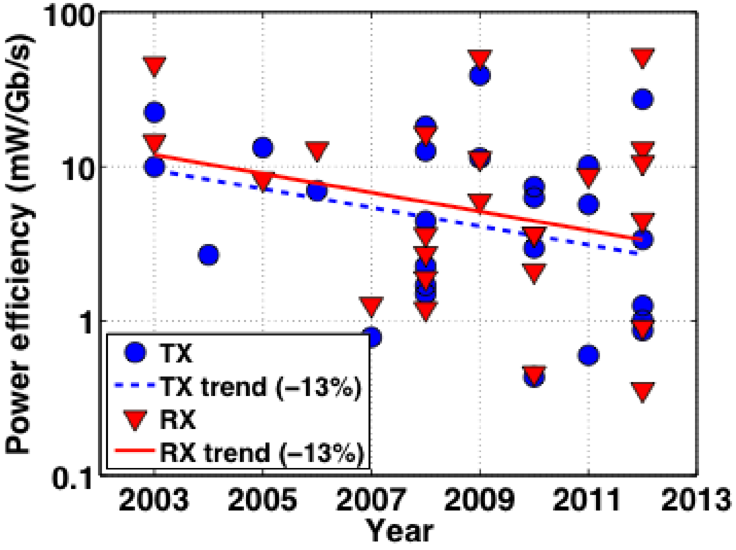

(a)

Arash Zargaran-Yazd, PhD thesis UBC

- 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?

Data encoding: necessary evil?

Sandia National Laboratories

- Why encoding is necessary
  - Power dependent characteristics
  - Continuous control loop operation
    - Gated control loops
    - Secondary measure (temperature)
  - Receiver AC coupling?

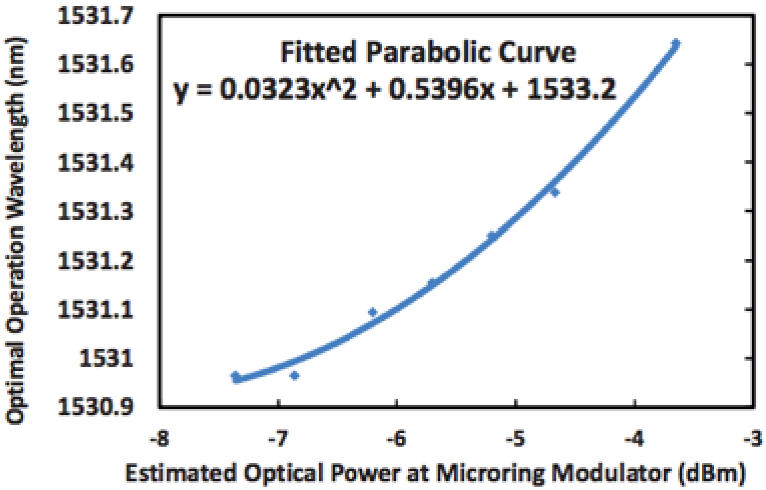

Li, Ophir, Xu, Padmaraju, Chen, Lipson, Bergman, IEEE OIC 2012

- Error correction (probably not)
  - CRC more energy efficient nominally (Koka et. al., SPIE 8265 (2012)
  - Link margin improved (6 – 12 dB)
  - Delay (bad)
  - Power (bad) → compare to optical amplification in some cases

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*

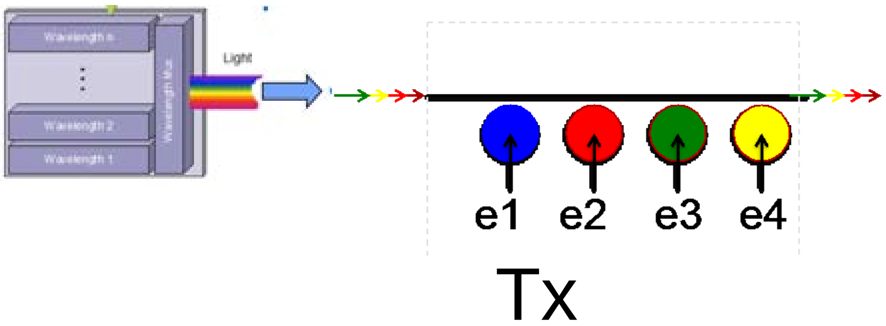

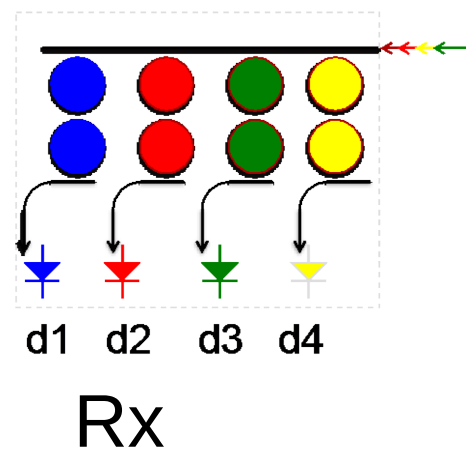

O/E Integration challenges

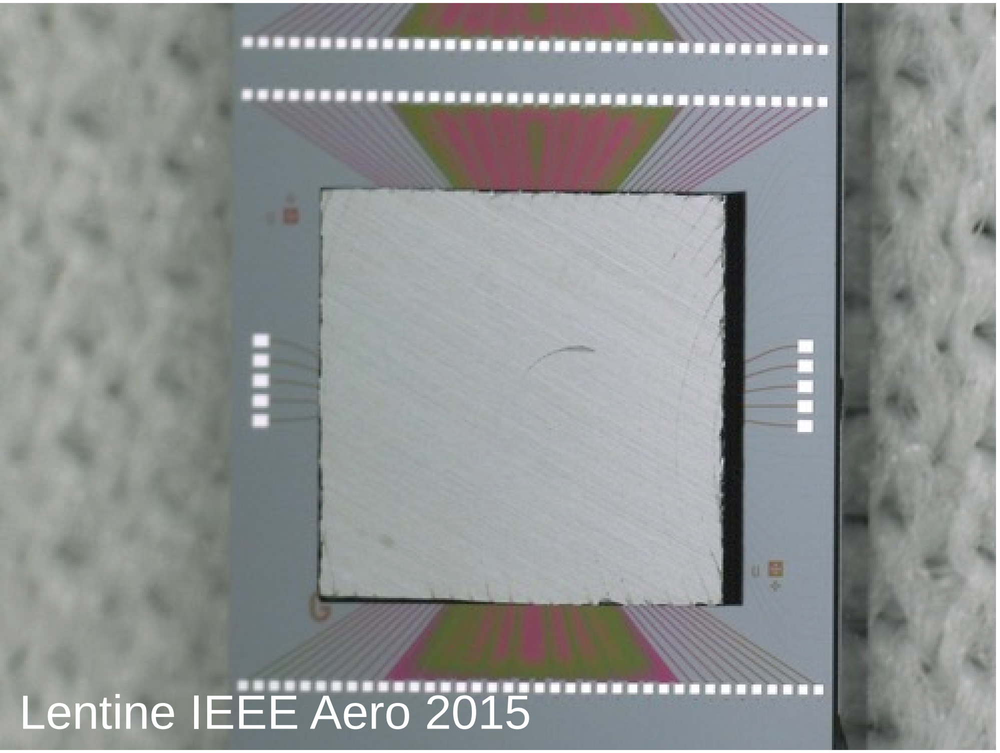

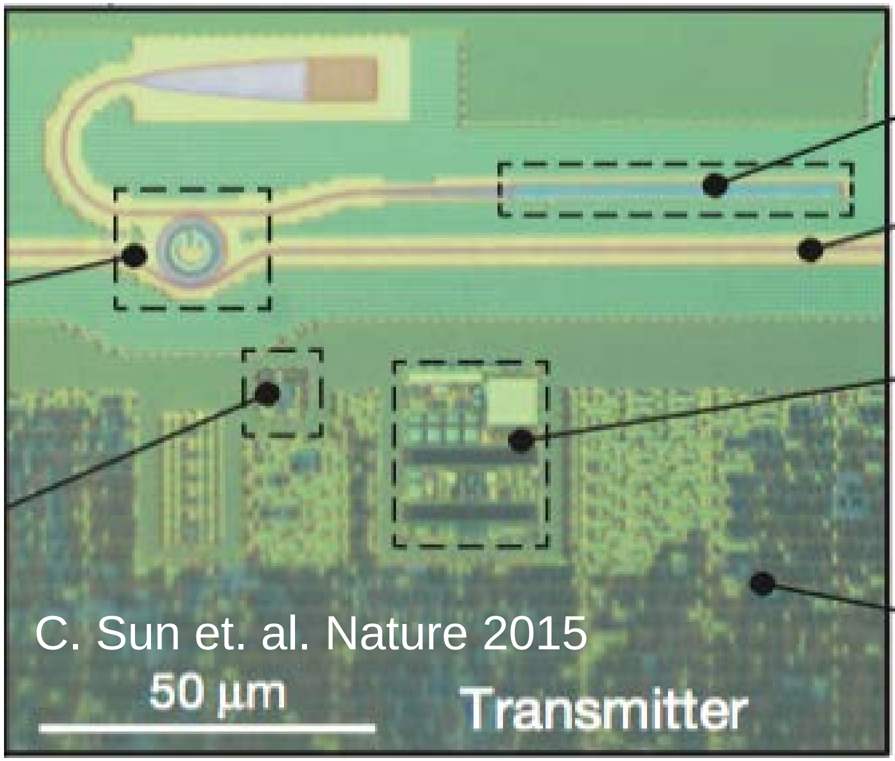

	Heterogeneous	Monolithic
Connection yield	*	***
Capacitance	*	***
Cost/circuit density	***	*
Performance	***	*

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 …

Optical PCBs: a new growth curve?

- 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 …

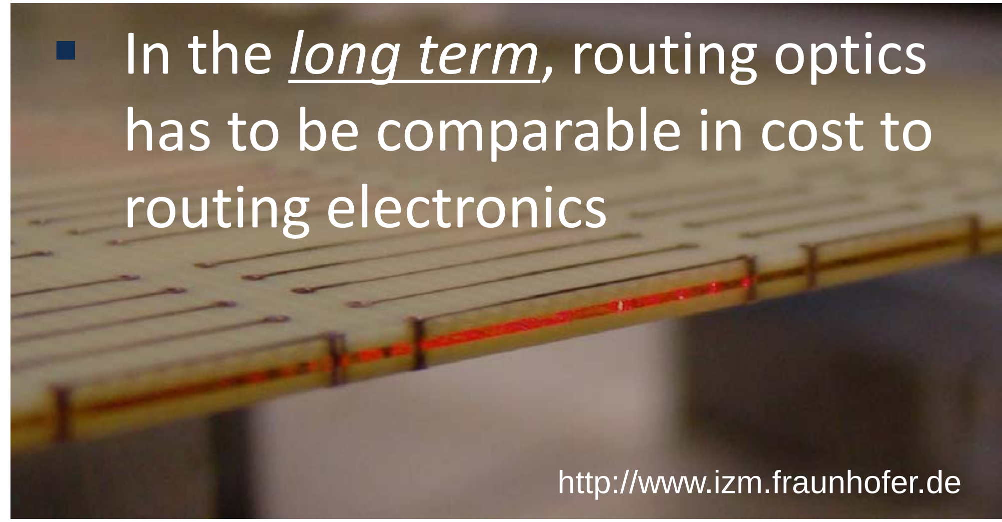

- In the *long term*, routing optics has to be comparable in cost to routing electronics

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

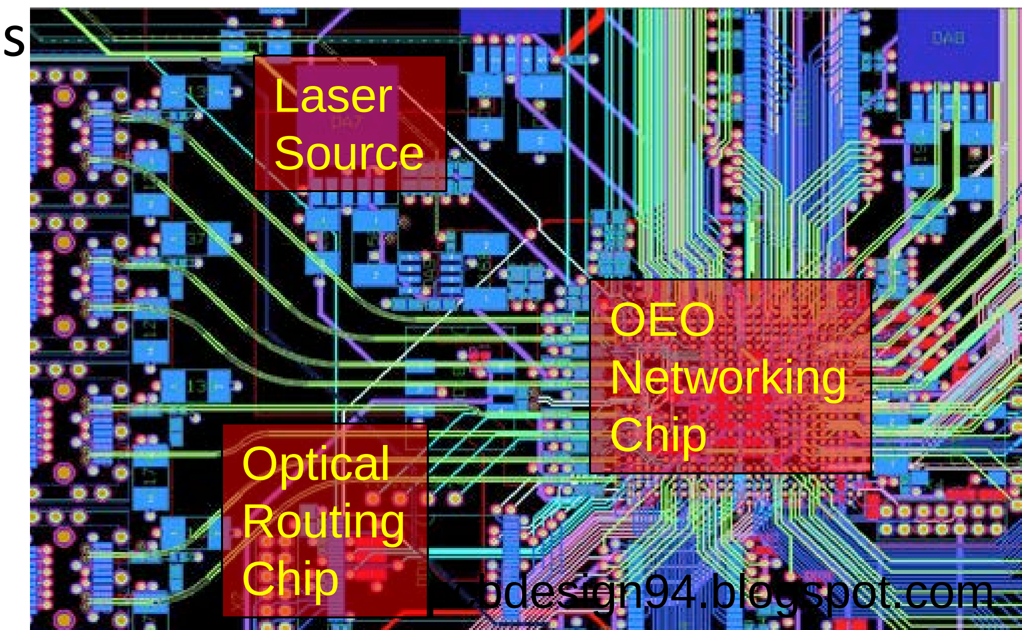

Computing: can optics keep up?

Sandia National Laboratories

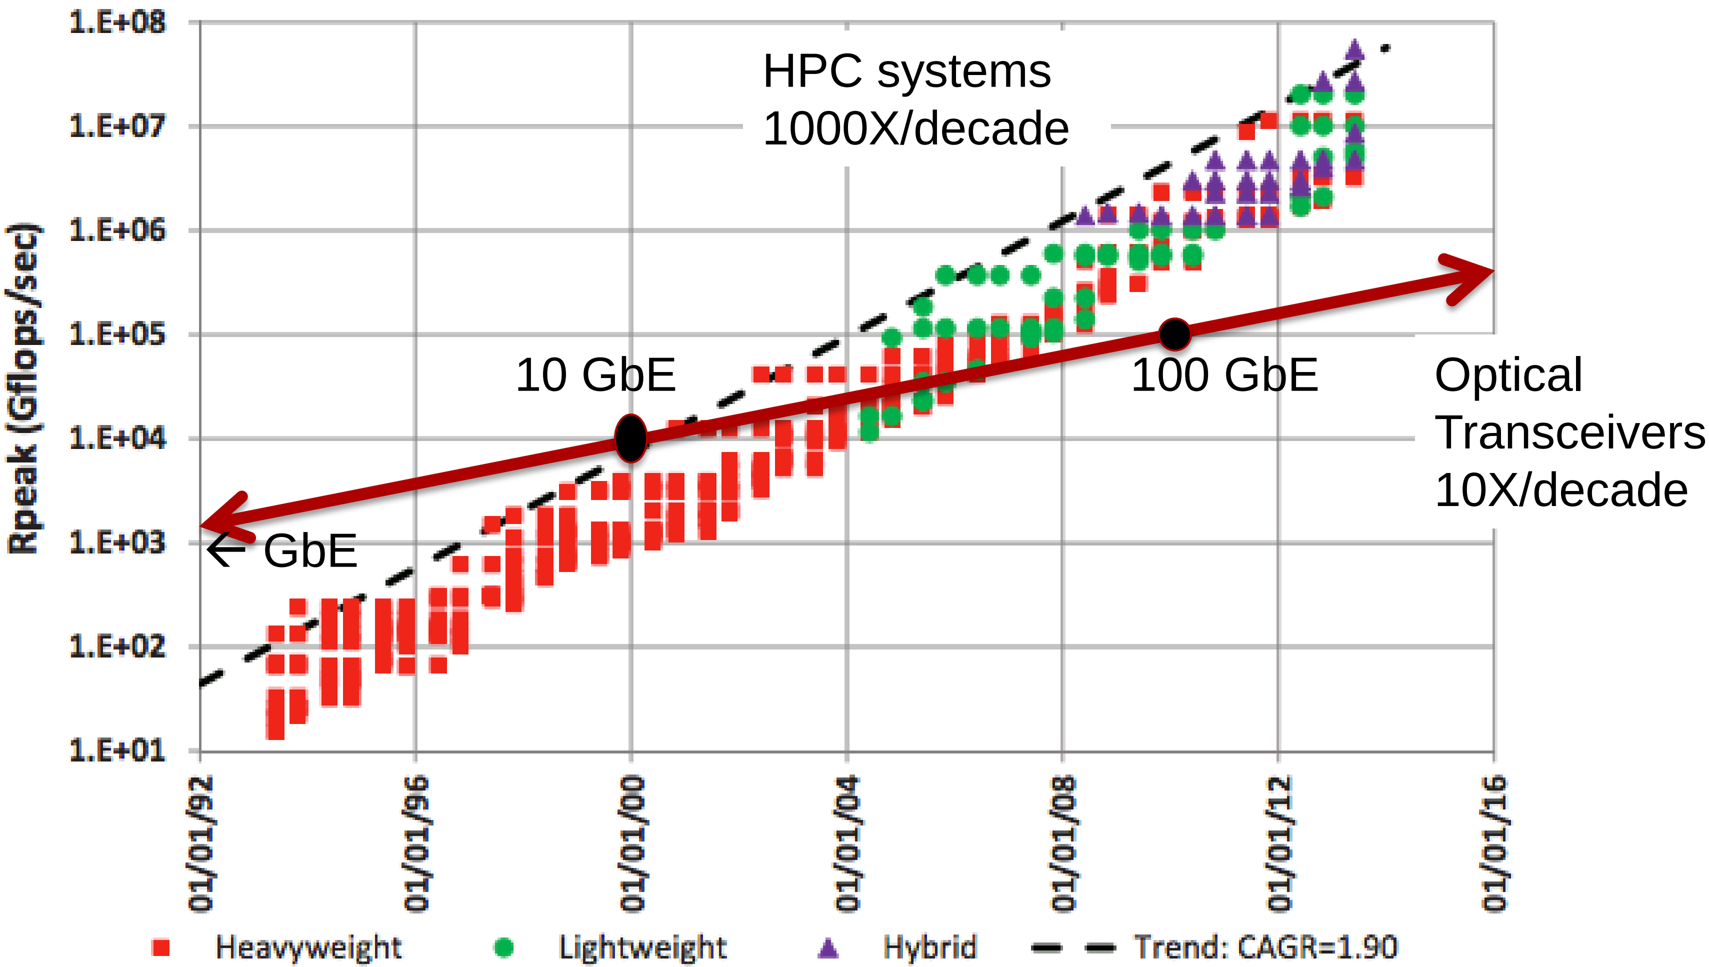

Koegge and Resnick, SANDIA Report 2013

Can optics keep real computing on track?

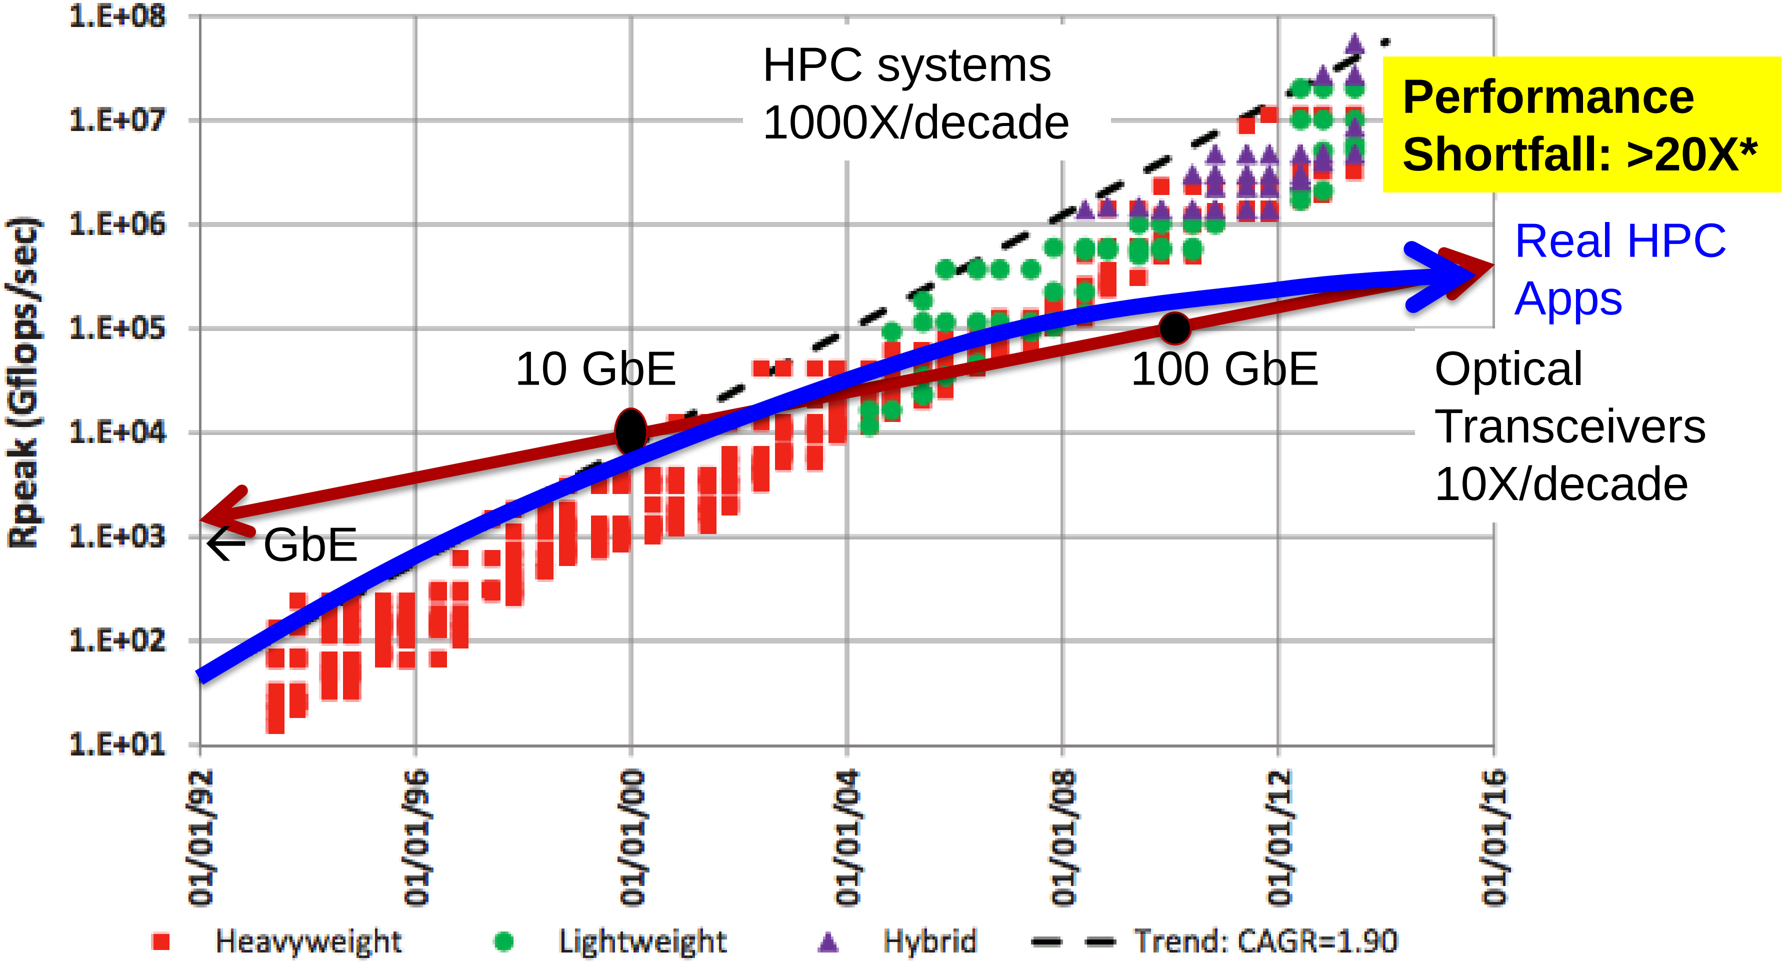

Koegge and Resnick, SANDIA Report 2013, *John Shalf, private communications

Si Photonics 2 x 2 WSS

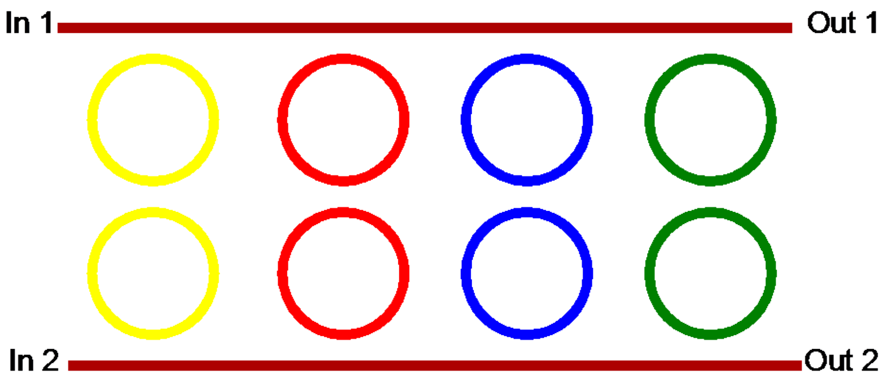

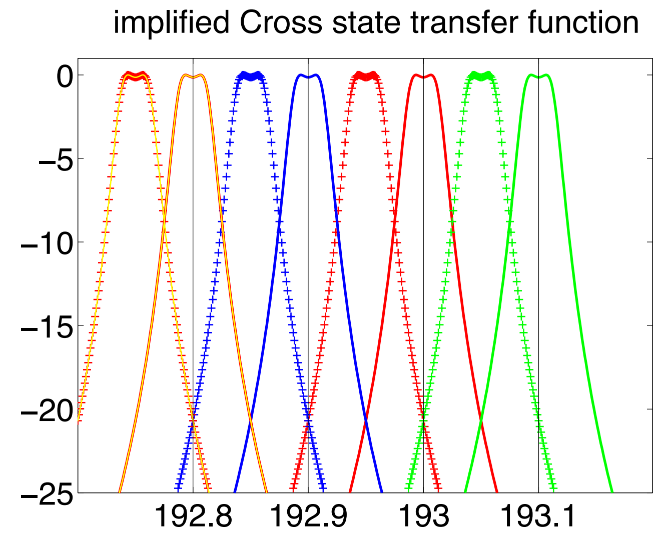

LONG TERM (IDEAL) SPECS:

- 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

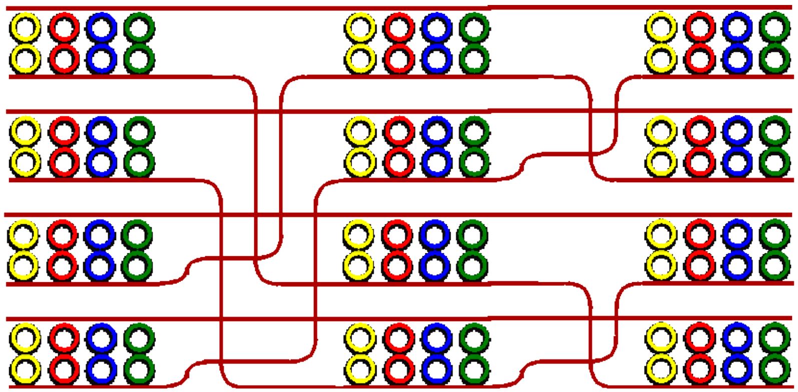

...

Chip scale 256 x 256 @ 32λ

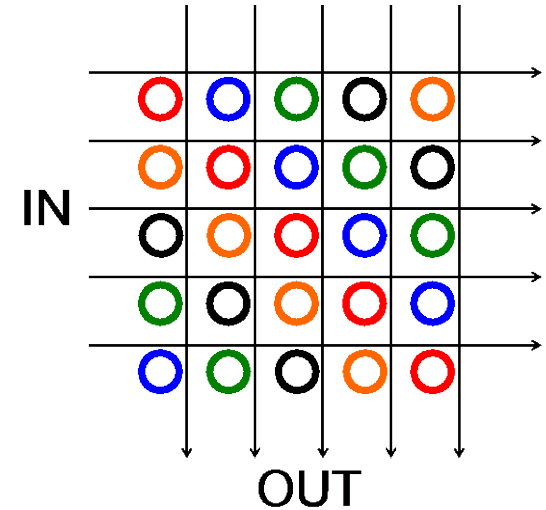

- **Networks may lag interconnect**
  - Routing,
  - Path-hunt, electronics
  - Buffering, electronics
- Low loss
- Pass band shape (larger base elements)
- Low power (non-thermal)
- *Fast* switches
  - Traditional WSS/MEMs competition for slow ones

Silicon Photonics Challenges

- **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?**

Questions?

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

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

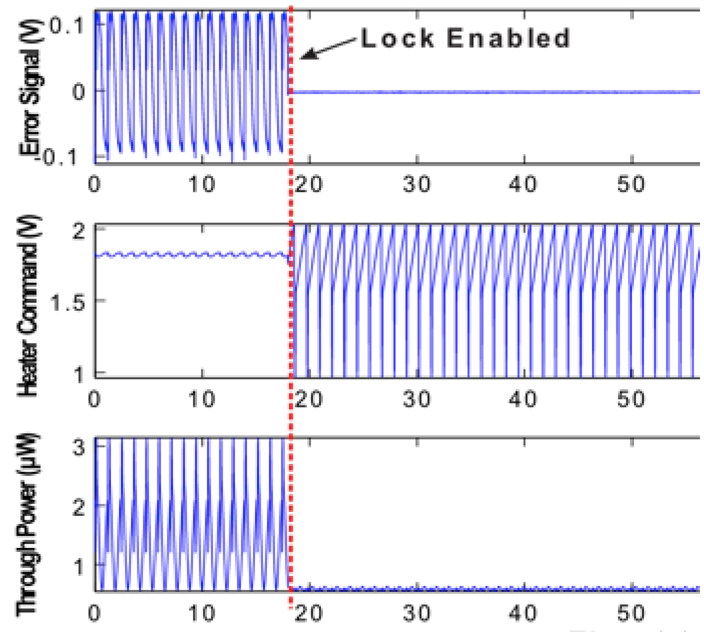

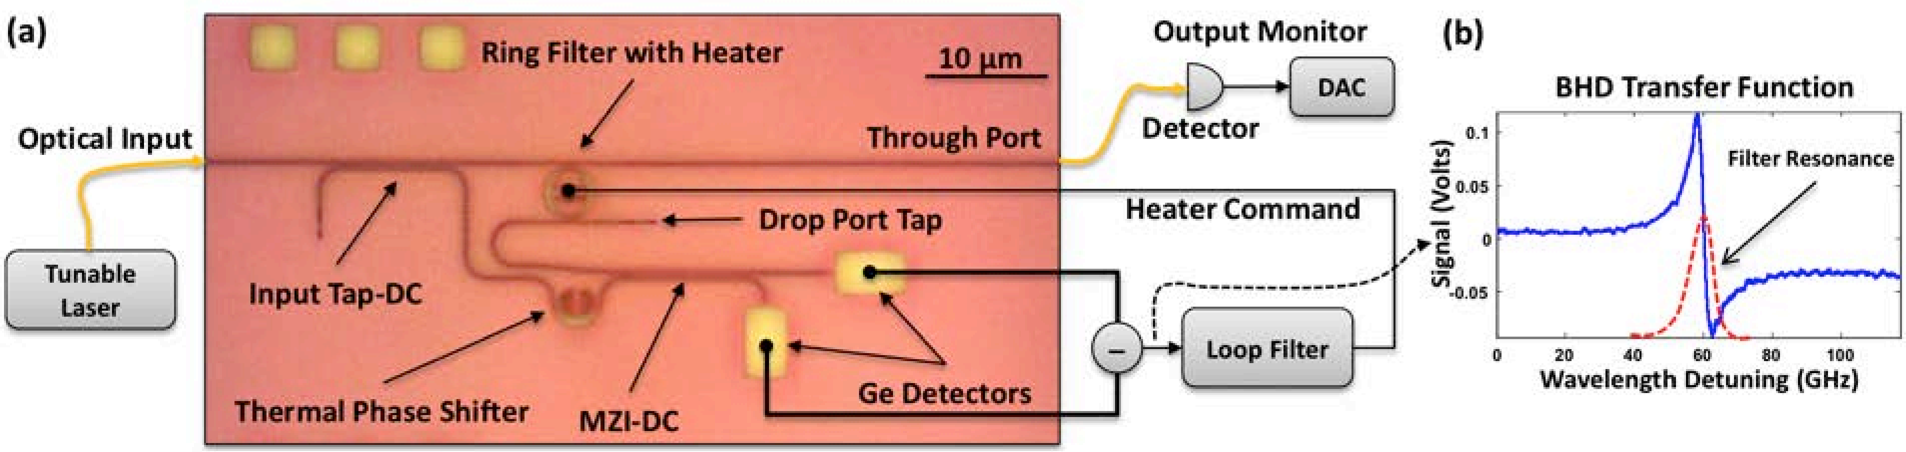

J. A. Cox et. al., IEEE Optical Interconnects 2013

First flip chip bonded SiP at Sandia

- Modulator drivers, receivers, and combinations
- IBM 45 nm CMOS
- Sandia Silicon Photonics
- 10 um bumps/14 um pads
  - Test coupons 2 um x 5 um
- 1120 connections

- 400 uA dynamic current draw at 1V
- 80 fJ/bit at 5 Gbps - high bond series R

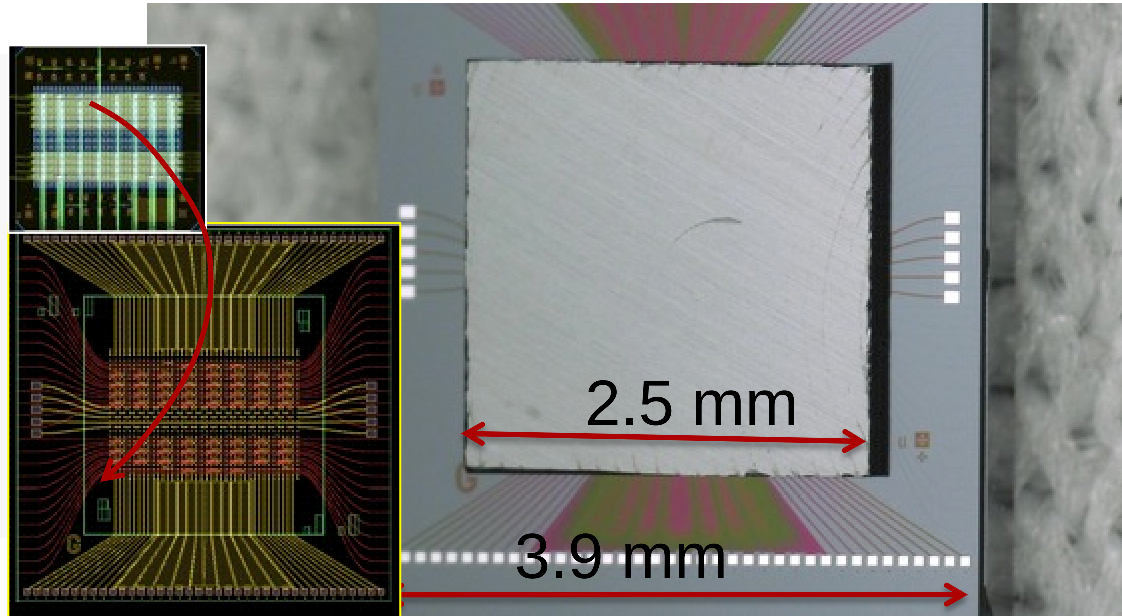

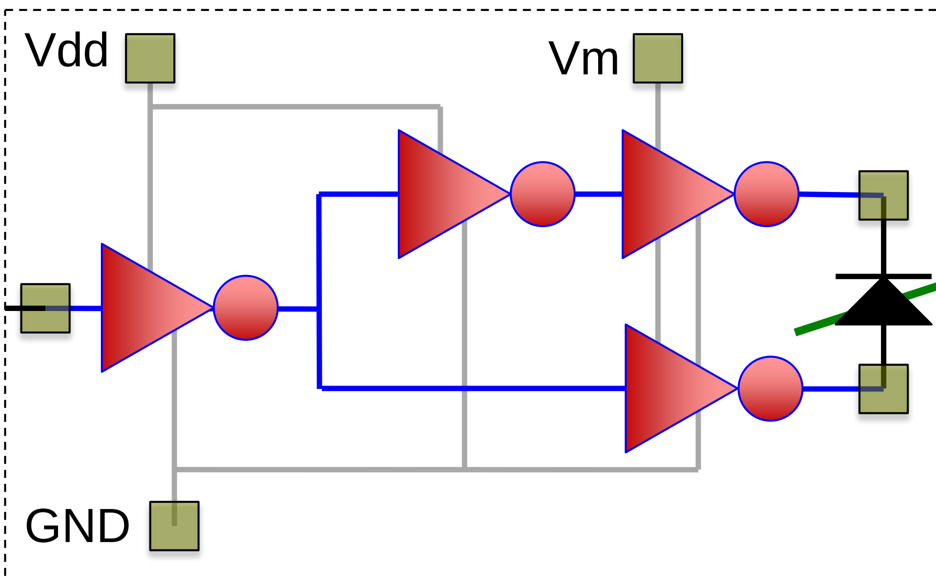

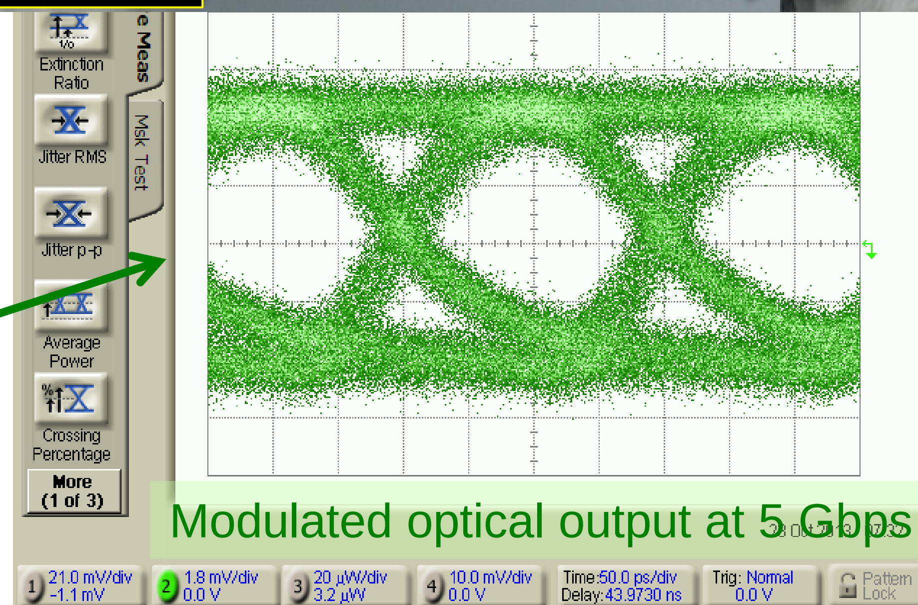

Electronic-Photonics Integration

Sandia National Laboratories

- Heterogeneous integration
  - Independent optimization of electronics & photonics
  - Need very high yields and small size

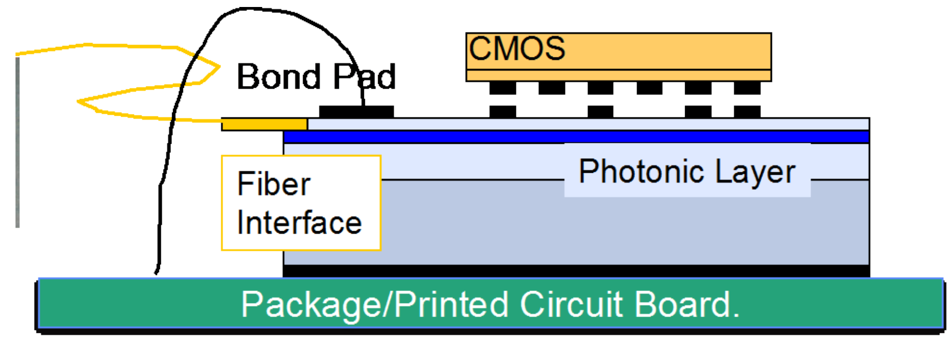

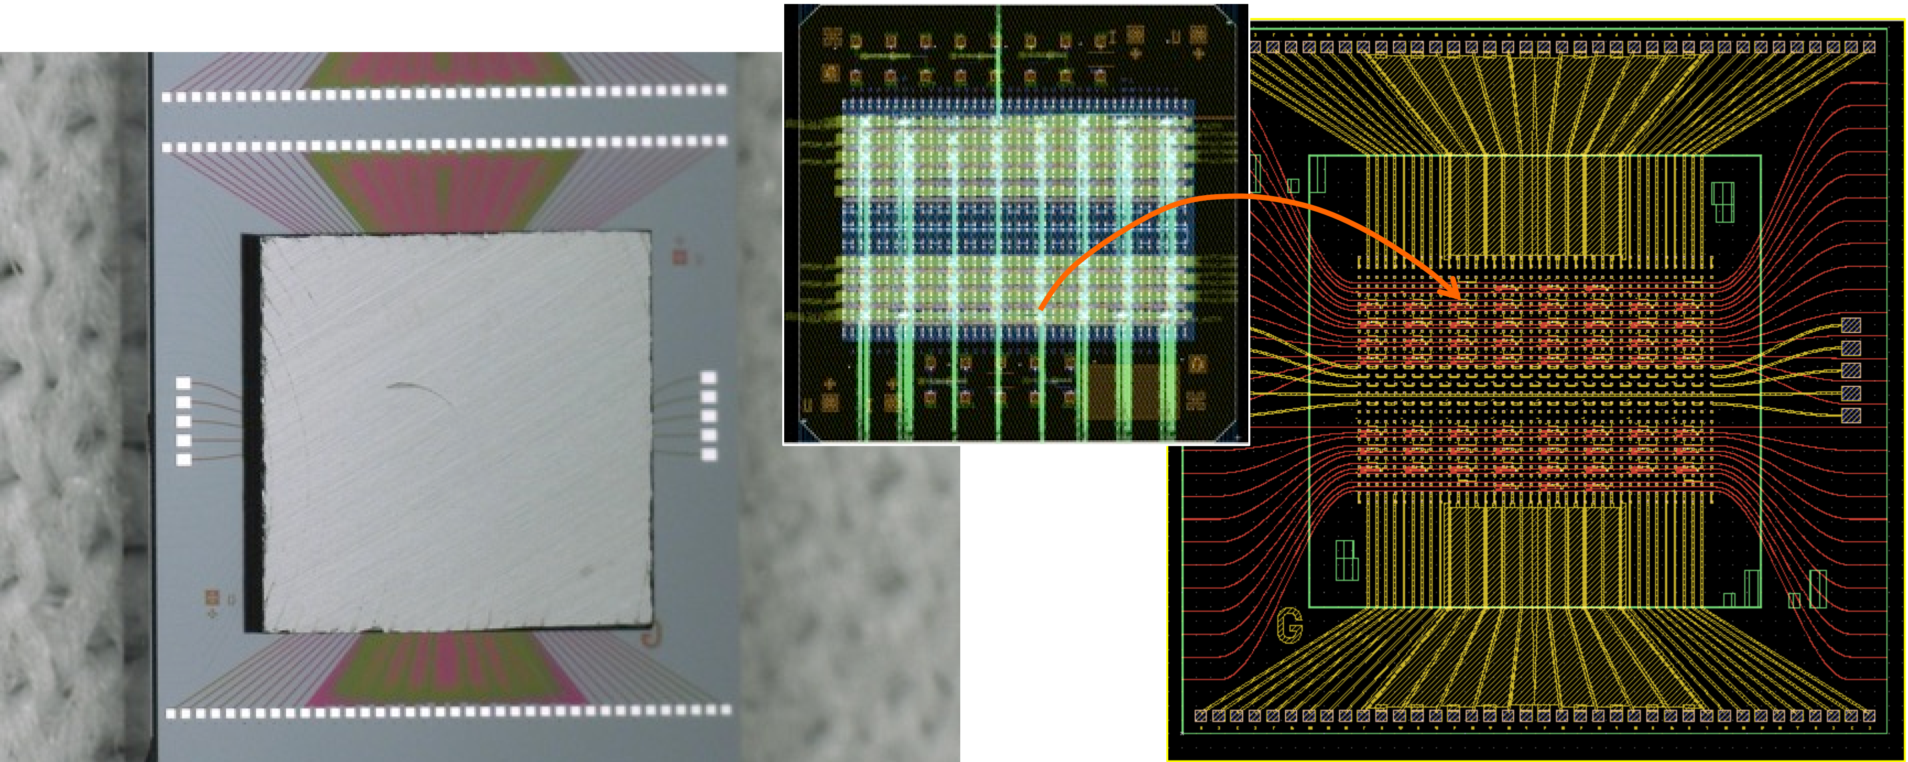

Sandia 2 x 2 silicon photonics switches

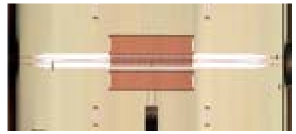

- Fast (< 100ps)
- Broadband
- 1pJ/switching event
- No static power
- 1 mm size

MZ – free carrier effect

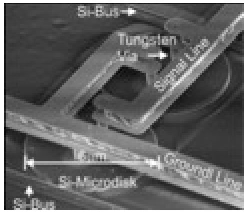

- Fast (< 100ps)
- Wavelength selective*
- 1fJ/switching event
- No static power
- < 10 um size

Ring – free carrier effect

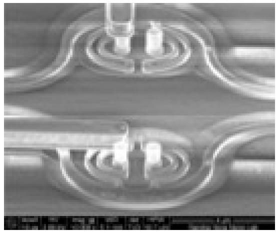

- Slow (10 us)
- Broadband
- ~ 15 mW/2π
- Static power in one state
- < 10 um size + coupler

MZ – thermo-optic

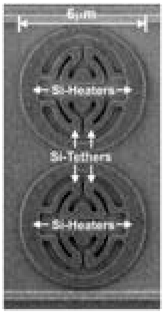

- Slow (10 us)
- Wavelength selective
- ~ 4 uW/GHz (200uW)
- Static power in one state
- < 10 um size

Ring – thermo-optic

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

Wavelength switching networks

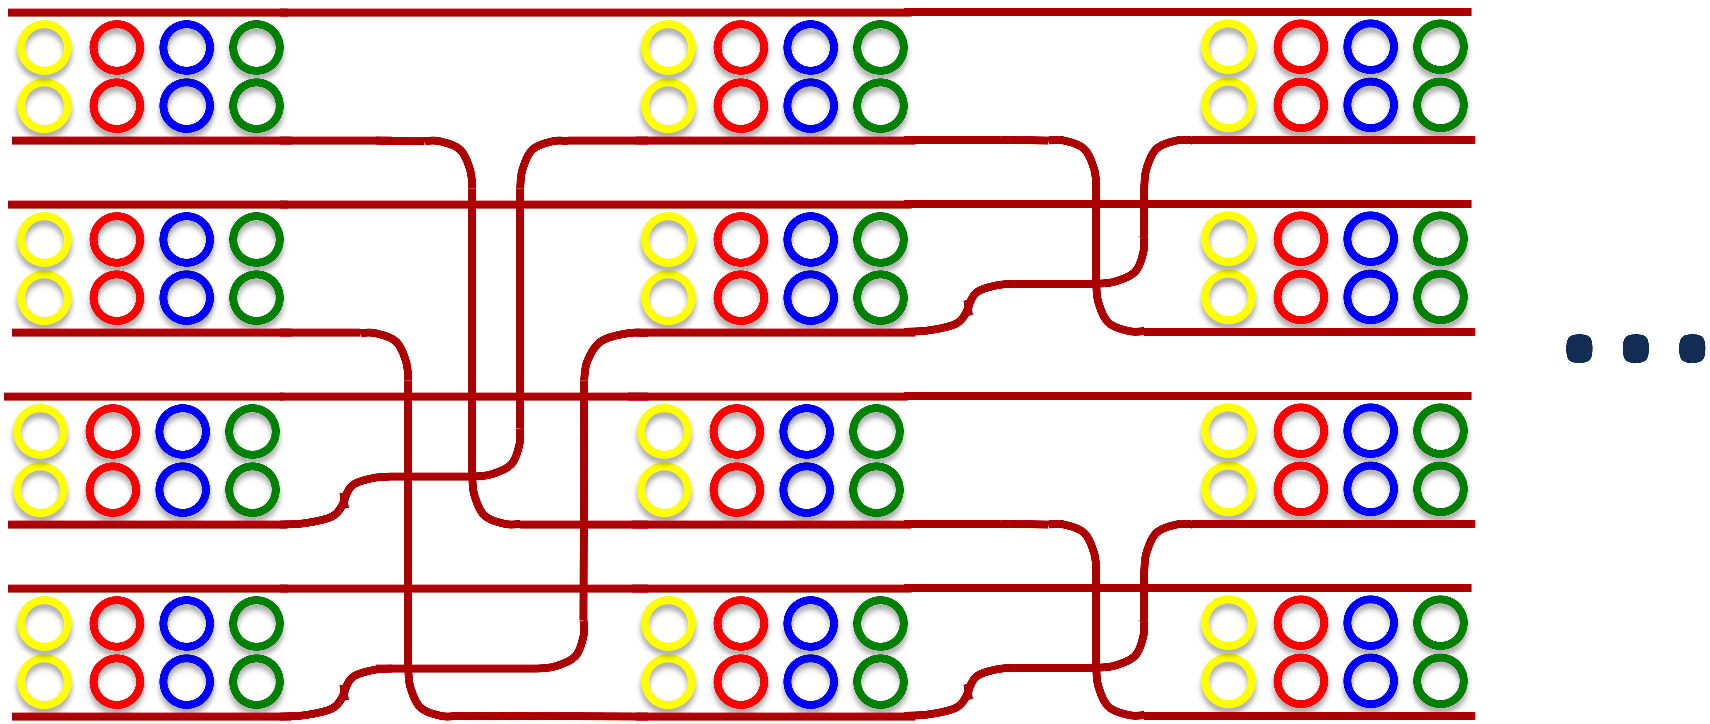

- Variety of networks from 2 x 2s
  - Squaring of crosstalk (EGS, Dilated Benes)
  - Tradeoff between initial fan-out and number of stages (EGS)
- Interconnects require planar crossings or two level optics
  - Nitride, Polysilicon: crosstalk can be very good, careful of loss

High radix and DWDM?

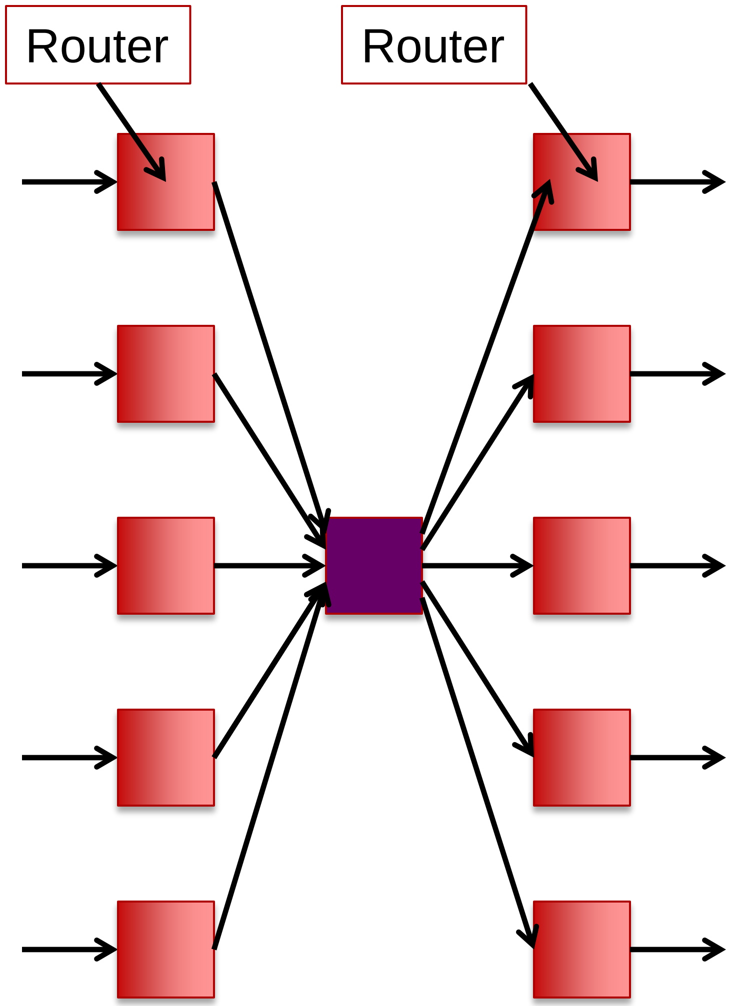

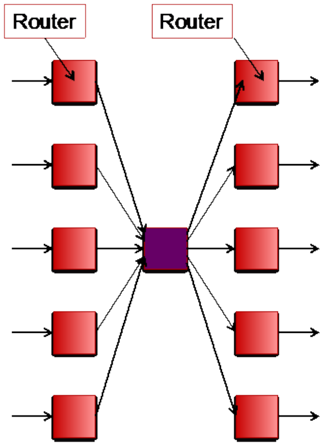

- xx

Outline

1. Title
2. Acknowledgements
3. DWDM Silicon Photonics
4. Near-term (TbE) vs. far-term
5. Why DWDM: high-radix
6. Challenge List
7. Laser-1: Options
8. Laser-2: individual lasers
9. Laser-3: hetero-integration
10. Laser-4: Comb lasers
11. Laser-5: Summary
12. Env. & man. – Oracle, Will, etc.
13. Filter coupling plot (fig.2)
14. Modulator plot
15. ITU or wander
16. Active control – 1 (general)
17. Active control – 2 (specifics)
18. Multiple filter issue (BER)
19. Multiple filter dither
20. Multiple filter BHD
21. Multiple orders
22. Filter pass band shape

Outline - II

23. Integrated receivers; simple
24. Differential signaling
25. Data encoding
26. Clock recovery
27. ‘Optical’ SERDES
28. Packaging-near-term
29. Packaging-far-term
30. Networks – challenges
31. Networks - example
32. Summary