SAND2009-0541C

Limits to Silicon Modulator Bandwidth and Power Consumption

Michael R. Watts

D. C. Trotter, R. W. Young, A. L. Lentine, and William A. Zortman

2009 SPIE Photonics West, San Jose CA
1/27/2009

Sandia National Labs, Albuquerque NM
Applied Photonic Microsystems

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Outline

Exascale Computing: Challenges Ahead

Modulator Bandwidth and Power Consumption

Bandpass Switches

Adiabatic Microring Resonators

Motivation: Power / Bandwidth in Supercomputers

"Based on current trends, by 2011 data center energy consumption will nearly double again, requiring the equivalent of 25 power plants. The world's data centers, according to recent study from McKinsey & Company, could well surpass the airline industry as a greenhouse gas polluter by 2020."

Quote from "Demand for Data Puts Engineers in Spotlight," New York Times, June 17, 2008

Sandia's Red Storm Supercomputer (26,569 processor cores, ~3.5MW power)

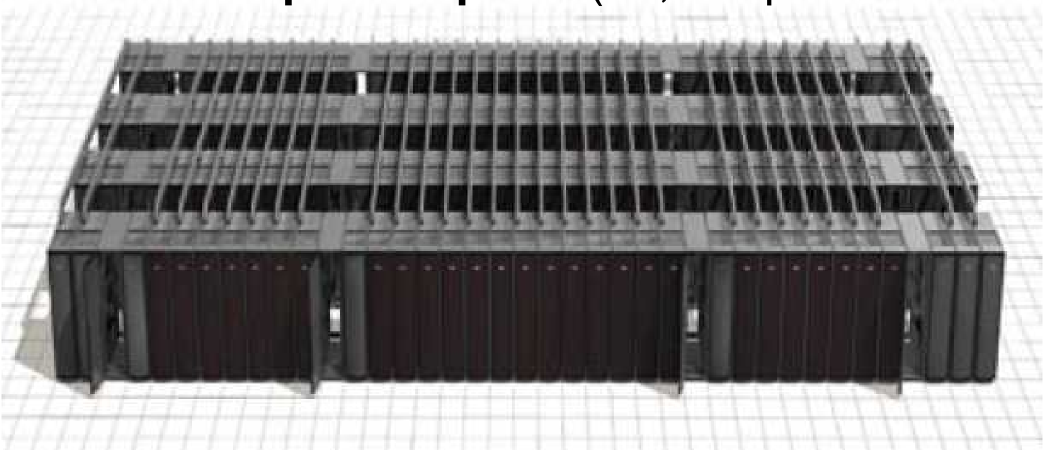

- **Question:** Now at 1PFLOPs, what does it take to get to 1-ExaFLOPs?
- **Power:** Hoover Dam provides ~1GW . . . what will future supercomputers need?
- **Bandwidth:** To scale real-world application performance, communications must scale with compute performance (i.e. bytes/FLOP ~ 1)

Sandia National Laboratories

Interconnect Requirements for Super-Computers

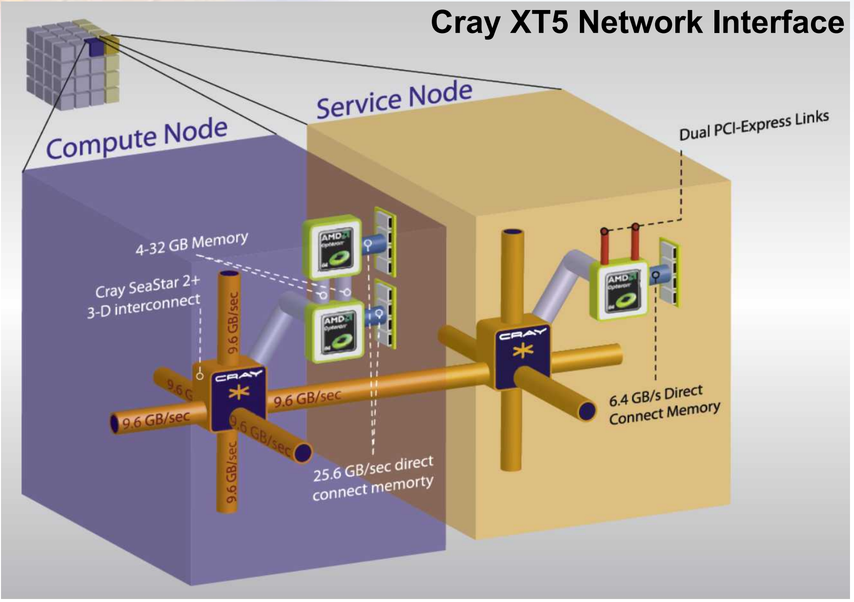

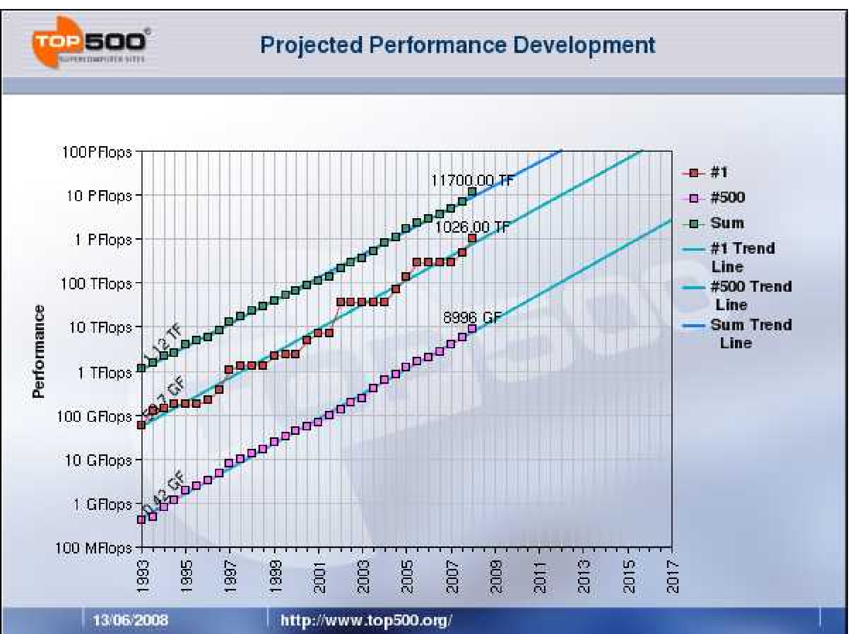

Today (2008)

- **BW/node:** 6×9.6GB/s+6.4GB/s → 0.512Tb/s/node
- **BW/core:** 0.512Tb/s (~1 Byte/FLOP orig.)
  - now BW/core is ~128Gb/s
- Top machines achieve ~1PFLOPs peak
- ~20% power is used for network comm.

Future (2018)

- Expected to reach 1-ExaFLOPs
- ITRS projects we will reach 17nm node
- Cores/chip >100, conservatively
- 100cores/chip*0.51Tb/s/core = 51Tb/s/chip

Can the bandwidth requirements be achieved with electrical signaling at reasonable power levels?

Sandia National Laboratories

Can Electronics meet BW Requirements?

IBM 65nm Process Metal

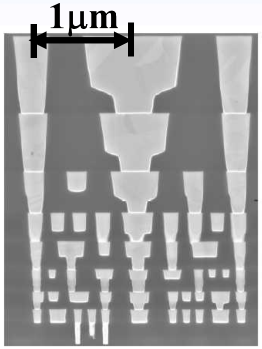

Intrachip Electrical: Charged Lines

- **Energy/Bit: $CV_0^2/4$ → ~0.5pJ/bit/cm (25W for 50Tb/s)**
- **Bandwidth:** Achieve **~1Gb/s/1μm** → 10Tb/s/cm/layer → **5 layers**
- **Conclusion:** <u>On-chip electrical signaling is troublesome in 2018</u>

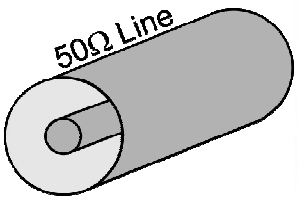

Interchip Electrical: Transmission Lines

- **Energy/Bit: $\tau V_0^2/2Z_0$ → ~10pJ/bit** (50Ω line, 1ns pulse, 1-Volt signal)
  - Chip Comm. Power: 50Tb/s × 10pJ/bit → **500W/chip**
  - System Comm. Power: 1-ExaFlOPs → **1.6GW Communications Power**
- **Bandwidth**
  - **Board: 100Gb/s/cm/layer 50 layers → 50Tb/s**
  - **Pins-to-board:** 10Gb/s/pin → **5000 pins/chip → 50Tb/s**
  - **Inter-Rack:** 5000 wires/chip, 500-chips/rack → **2.5M wires/rack**
- **Conclusion:** <u>Logistically, off-chip electrical signaling becomes impossible</u>

Receiver Power for Optical Communications

Bit-Error Rate (Fundamental)

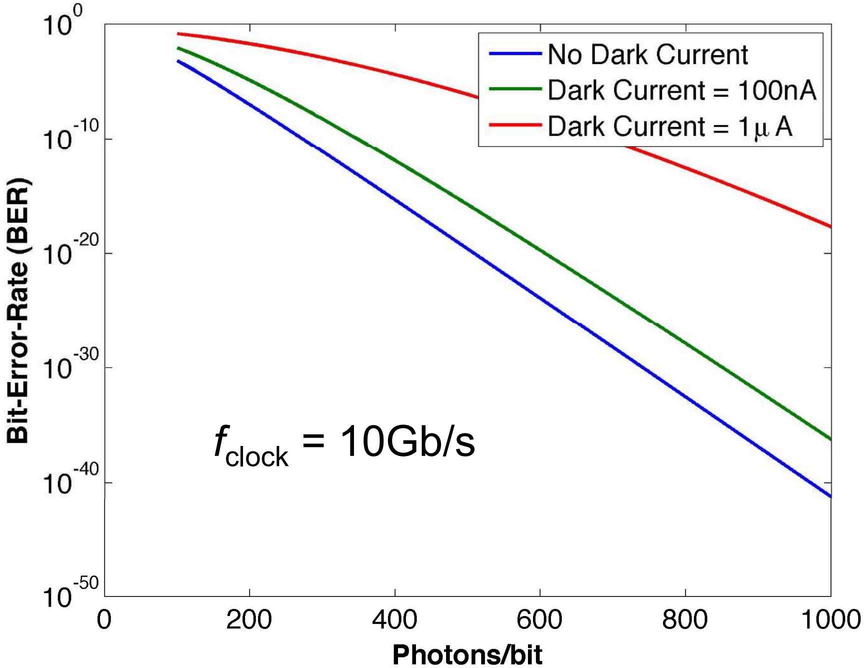

Results Based Shot, Johnson, and Dark Current Noise

- ❑ Require Dark Current < 100nA
- ❑ 1000 Photons/Bit is sufficient
- ❑ Figure 10dB loss, 10% E-O efficiency, still only 15fJ/bit

Capacitance (Technological)

Required Voltage: $V = Q/C_T$

❑ $C_T = C_g + C_d$, so capacitance can be dominated gate or photo-detector

Photodetector Capacitance

❑ $C_d = \varepsilon A/d \approx 12\varepsilon_0(2\pi \cdot 2 \cdot 0.1) \cdot 10^{-6}/0.5 \approx 0.27$ fF

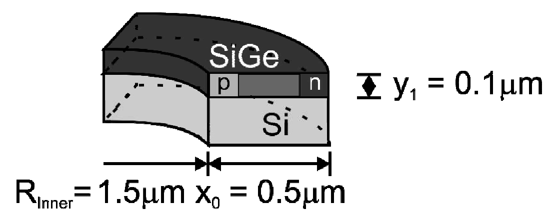

Gate Capacitance

❑ ITRS-HP Roadmap C_g ~ 0.5fF/μm

Initial Conclusion

❑ 0.3fF corresponds to ~2000 photons required or ~0.3fJ/bit at the receiver

❑ Even with 20dB penalty ~30fJ/bit

What about an Optical Network Interface?

Electrical Network Interface

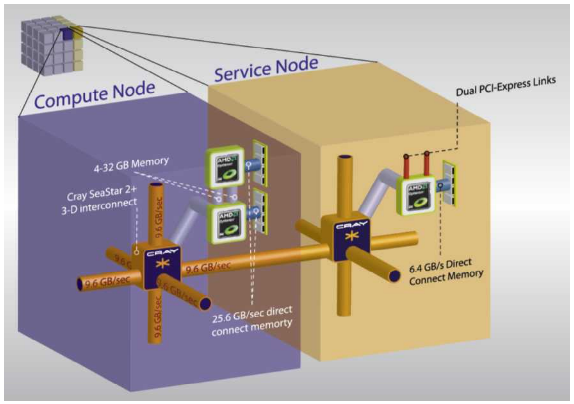

Microphotonic Network Interface

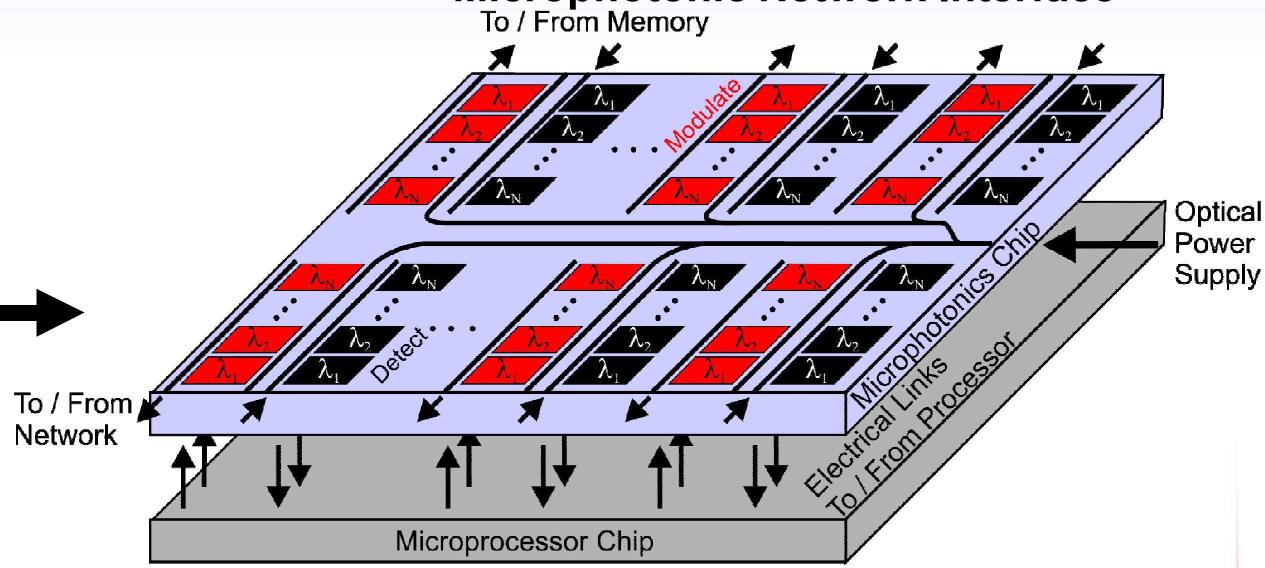

Optical Communications (Tightly Integrated with CMOS)

- **Energy/Bit = $h\nu V_0 C/(\eta q)$** (achieve <**100fJ/bit** → <**16MW @exascale**)
  - Receiver: <**1fJ/bit** required to flip a gate (Miller, 1989)
  - Modulator: Limited by capacitance (70fJ/bit power budget)
- **Bandwidth**
  - Line bandwidth: >1-Tb/s (10Gb/s @100λ's), 50GHz spacings → **5THz (40nm)**
  - Bandwidth (on-chip): > 200Gb/s/μm → **1Tb/s on a 5um pitch**
  - Bandwidth (off-chip) > ~10Tb/s/mm (**5mm → 50Tb/s, 25k fibers/rack**)
  - Routing of data could be O-E-O or in the optical domain . . .

Sandia National Laboratories

Possible Effects in Si/Ge Systems

Ge: Band-Edge Effects (Electro-Absorption $\Delta\alpha$)

Franz-Keldysh

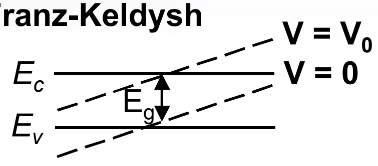

QCSE

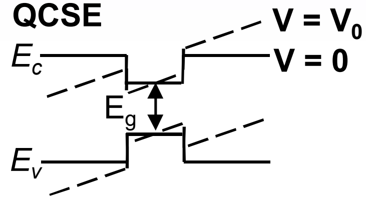

Quantum Confined Stark Effect

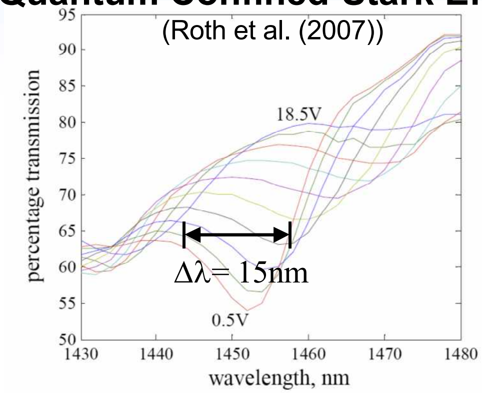

- ❑ Franz-Keldysh (Liu et al. *NP,* July 2008)
- ❑ Quantum Confined Stark (Roth et al. *OE,* April 2007)
- ❑ **Advantage:** Strong effect → small, low power devices (Liu claims 50fJ/bit)
- ❑ **Disadvantages:** → Limited Optical Bandwidth ($\Delta\lambda$ = 10nm-to-15nm) Limited Contrast ($\Delta\alpha/\alpha$ ~ 3), Intricate Fabrication

Free-Carrier Effect

(Soref and Bennett (1987))

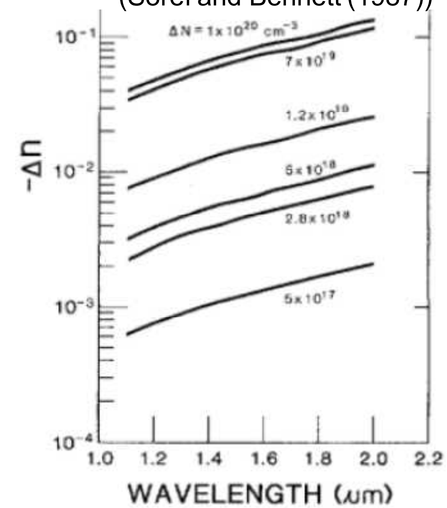

Si: Free-Carrier Effect (Electro-Refraction Δn)

- ❑ **Advantages:** Broadband effect ($\Delta\lambda$ >> 100nm), easy CMOS implementation
- ❑ **Disadvantages:** Weaker effect → need resonance for low energy/bit, but already have thousands of resonators on-chip

Magnitude of the Free-Carrier Effect

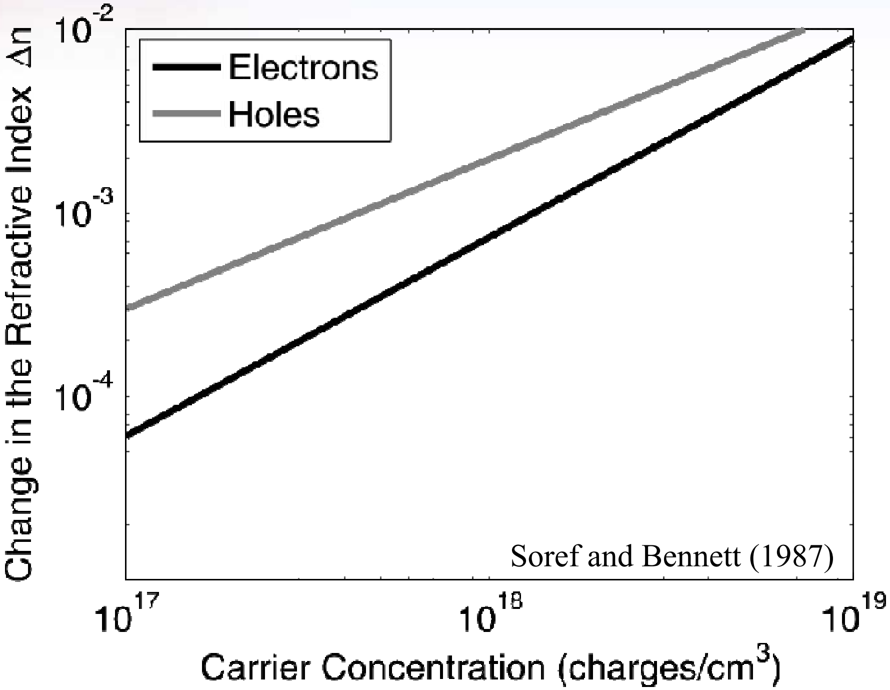

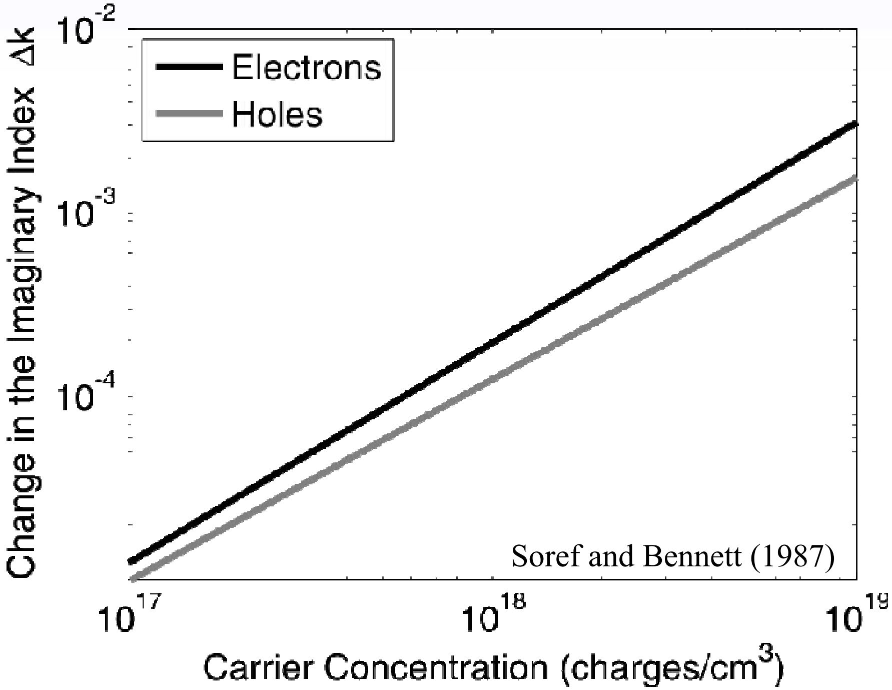

Important Points

- ❑ Effect is small
- ❑ Larger real change in index than imaginary
- ❑ Holes impact real part of the index more than electrons
- ❑ Holes induce less loss than electrons

Geometry Considerations (Depletion Approach)

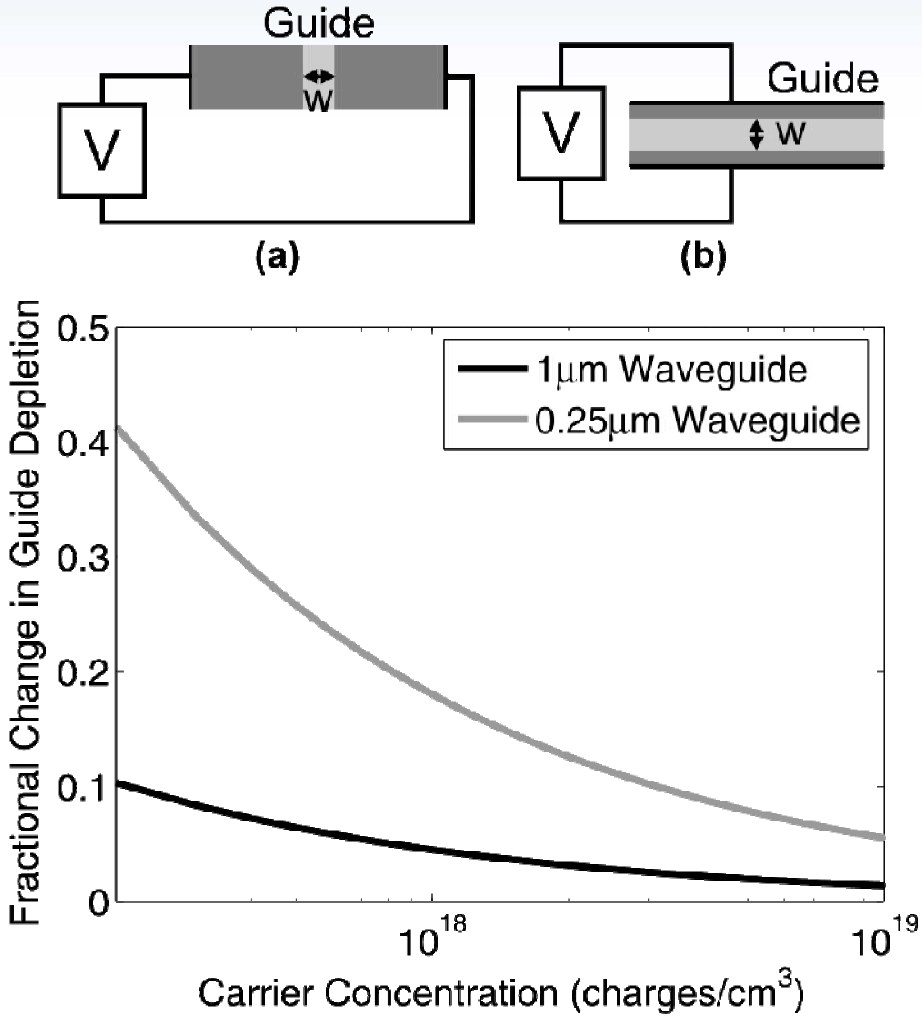

(c)

Important to Take Maximal Advantage of a Small Effect:

Since, waveguides are wide and thin for fabrication reasons, vertical junctions provide a greater overlap with depletion region

Sandia National Laboratories

Microdisk Modulators: A Vertical P-N Junction

Vertical P-N Junction

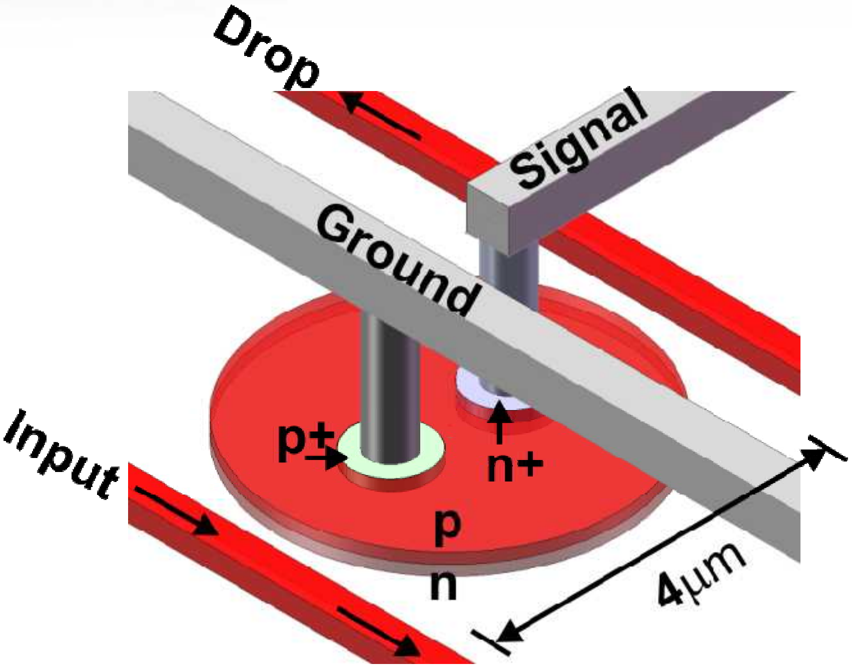

TE_{11} Cylindrical Mode

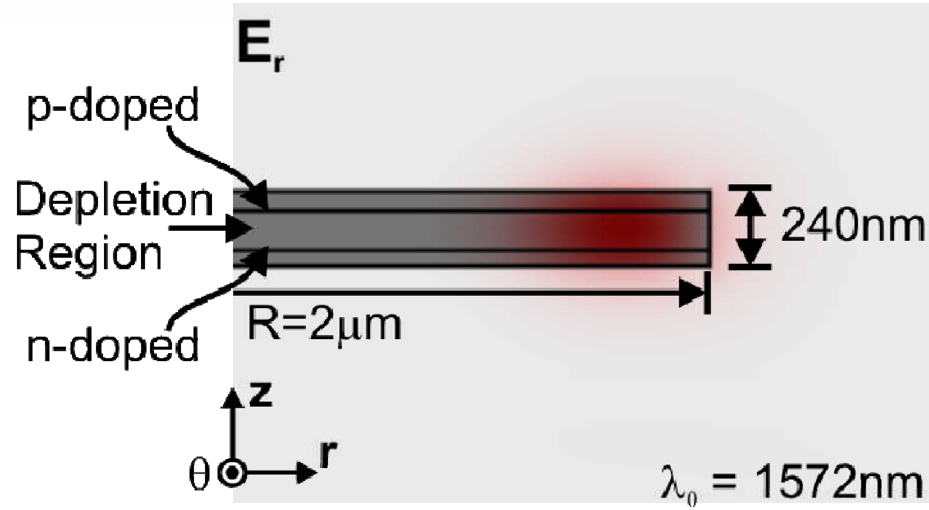

Advantages of a Vertical P-N Junction Modulator

- ❑ Vertical P-N junction enables tighter confinement and large modal overlap
- ❑ Devices as small as R = 1.5μm are possible
- ❑ Huge Free-Spectral-Range (> 9 THz possible w/ R=1.5μm)
- ❑ Smaller devices, no pre-emphasis → faster / lower power

Anticipated Optical Response

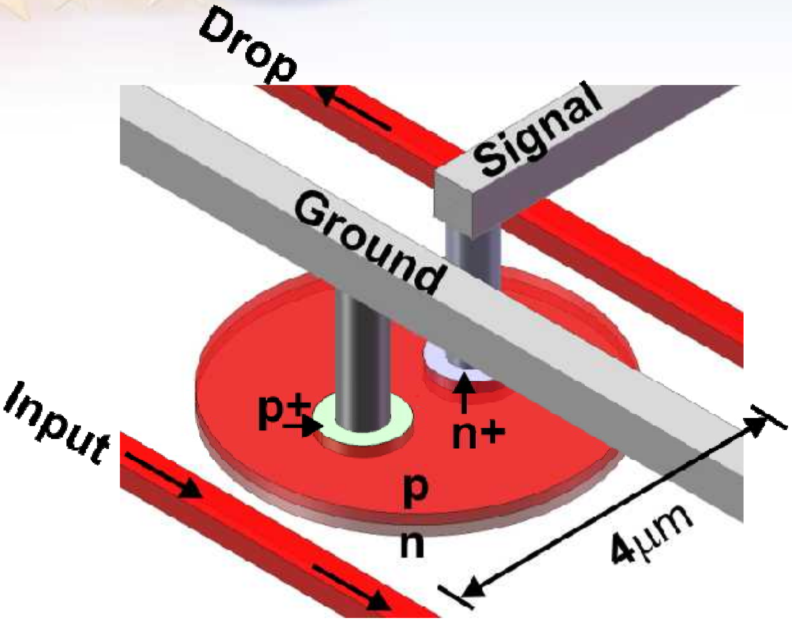

Doping Profile

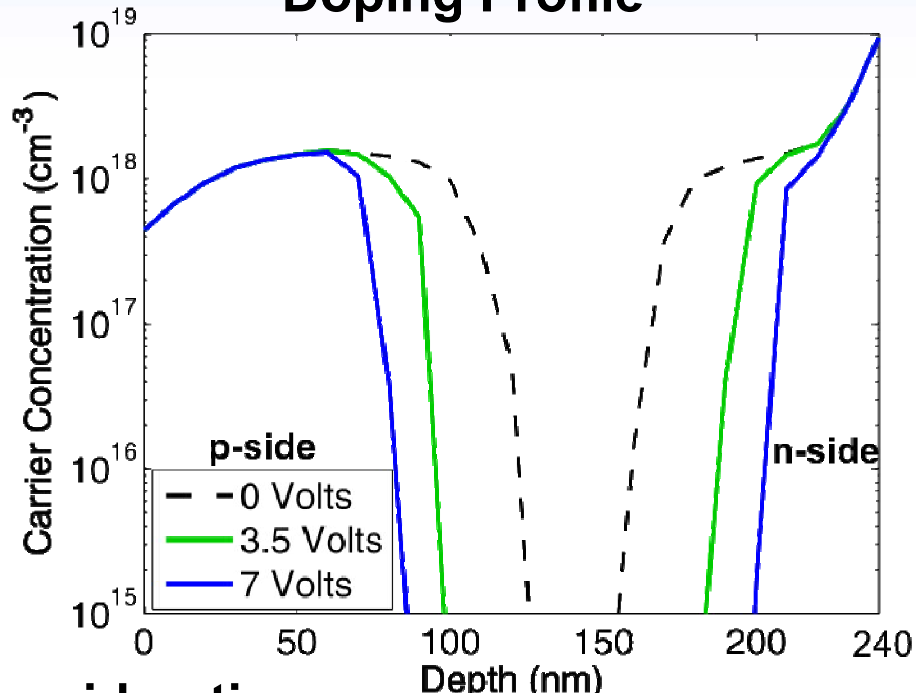

Optical Frequency Shift

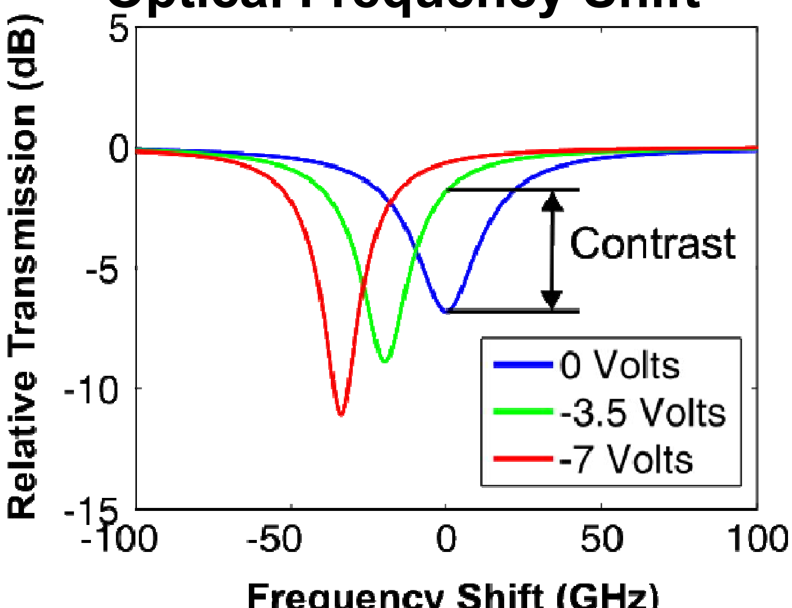

Considerations

- Depletion goes as square-root of doping, but index change is ~ linear with doping, so higher-doping ⇒ a larger change
- Higher doping levels also induce loss, limiting Q and contrast, but chosen Q must also be sufficiently low for a given BW

Sandia National Laboratories

Simple Disk: Bandwidth and Energy/Bit

Modeling Results

- **Bandwidth:** >5GHz possible at 4μm diameter
- **Energy/Bit:** ~20fJ/bit possible at 3μm diameter (2.5V) and can be reduced by going to higher-Q rings and slower modulation speeds . . . but . . . thermal control

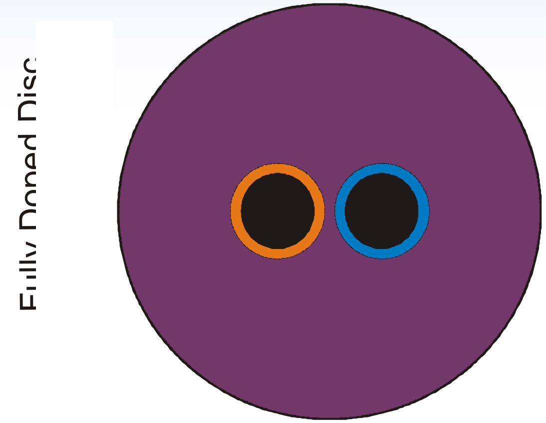

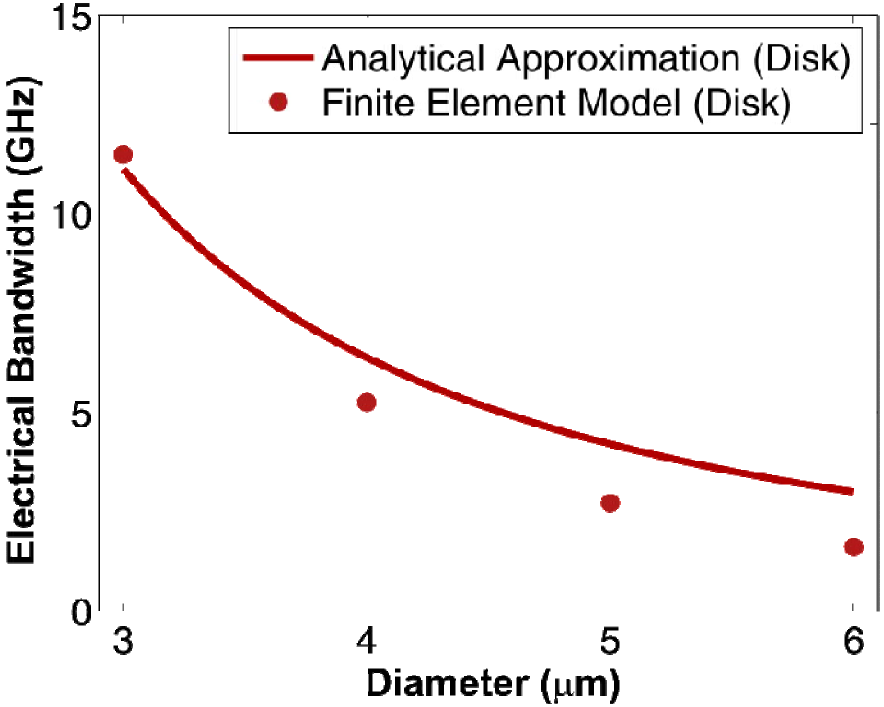

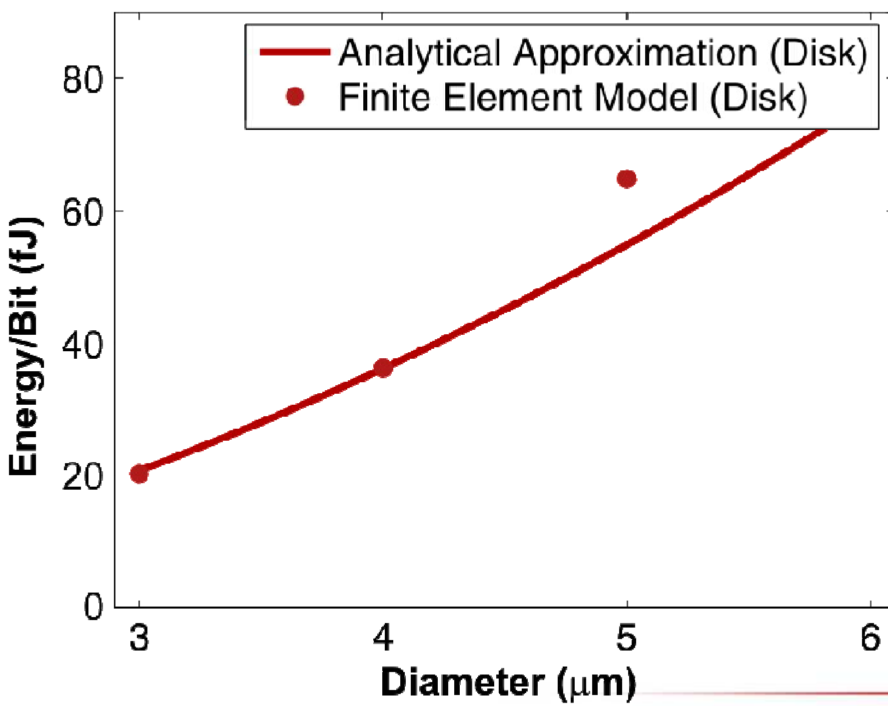

Microdisk Modulator Demonstration

SEM of the Microdisk

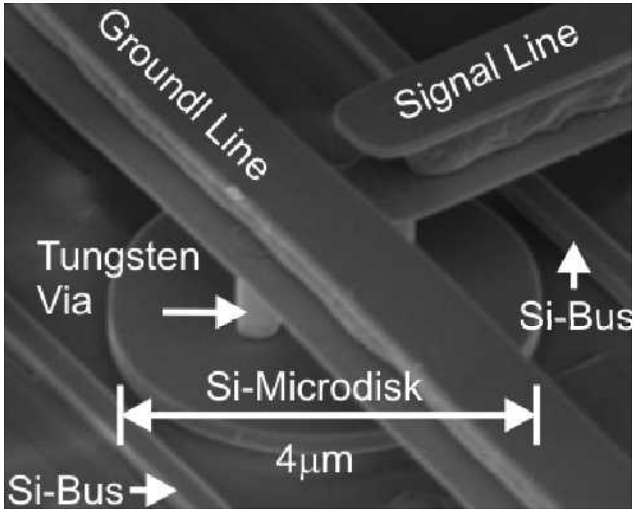

Frequency Shift vs. Reverse Bias

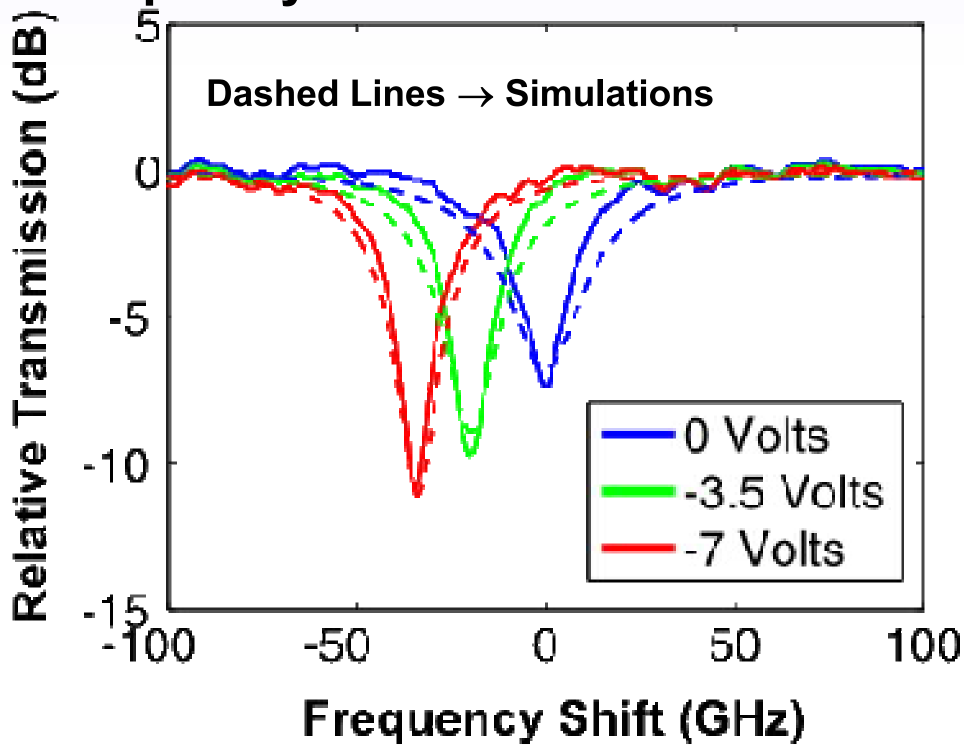

Eye Diagram (10Gb/s)

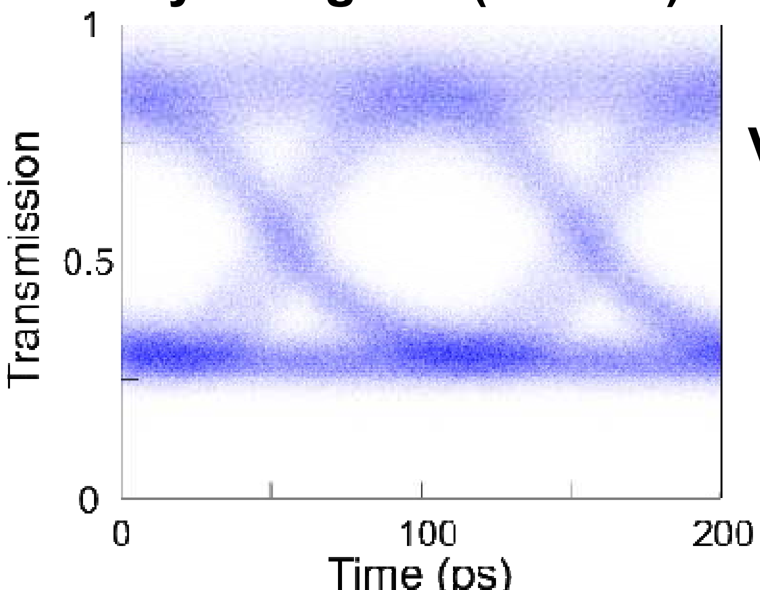

Vertical Junction Reverse-Biased Results

- 35-GHz freq. shift demo'ed, >70-GHz possible
- Achieved a BER<10^{-12} at 10Gb/s
- First resonant modulator with CMOS compatible drive (1.8V incident, ~3.5V due to reflection)

*Note: 2.5V/3.3V are avail. CMOS voltages @90nm

Sandia National Laboratories

Power Efficiency Measurement (energy/bit)

SEM of the Microdisk

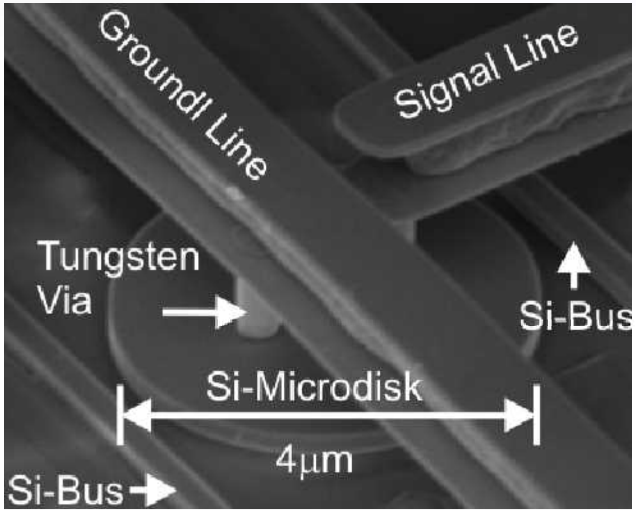

TDR Measurement vs. Simulations

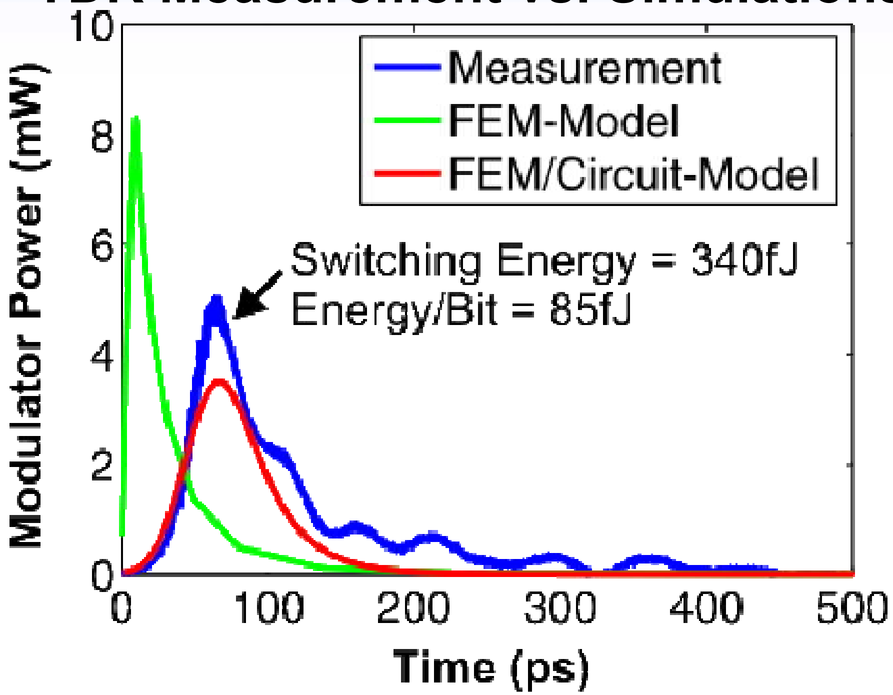

Time Domain Reflectometry Measurement Results (@3.5V)

- ❑ Depletion-based approach dissipates essentially no static power
- ❑ Switching Energy = 340fJ
- ❑ PRBS Energy/Bit = Switching Energy/4 = 85fJ **(100X less than electrical)**
- ❑ New designs indicate ~10fJ/bit is possible

Sandia National Laboratories

Can We Do Any Better?

Modeling Results

- ❑ **Bandwidth:** Reduces capacitance, increases resistance
- ❑ **Energy/Bit:** Reduced capacitance lowers energy/bit
- ❑ **Conclusions:** Do a little better . . . mostly on energy/bit

perhaps even more intricate doping schemes are necessary
. . . perhaps higher-Q devices

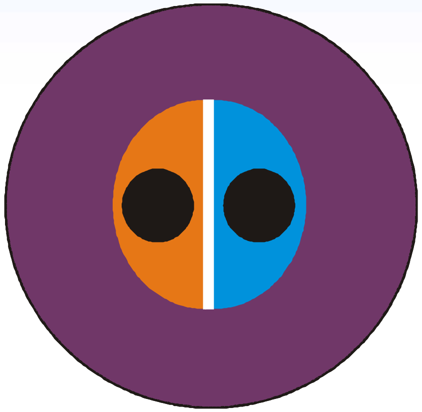

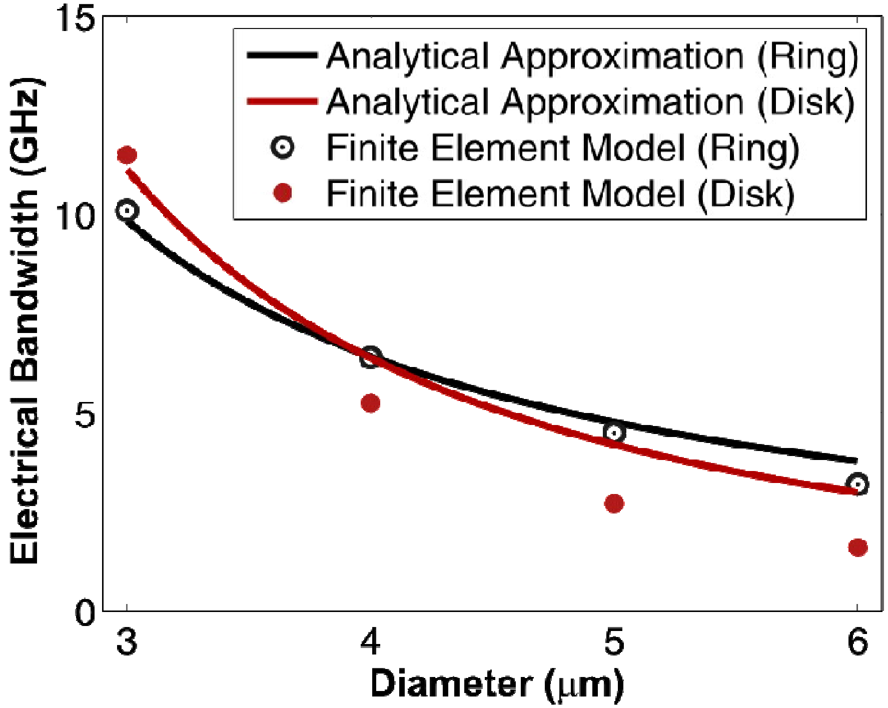

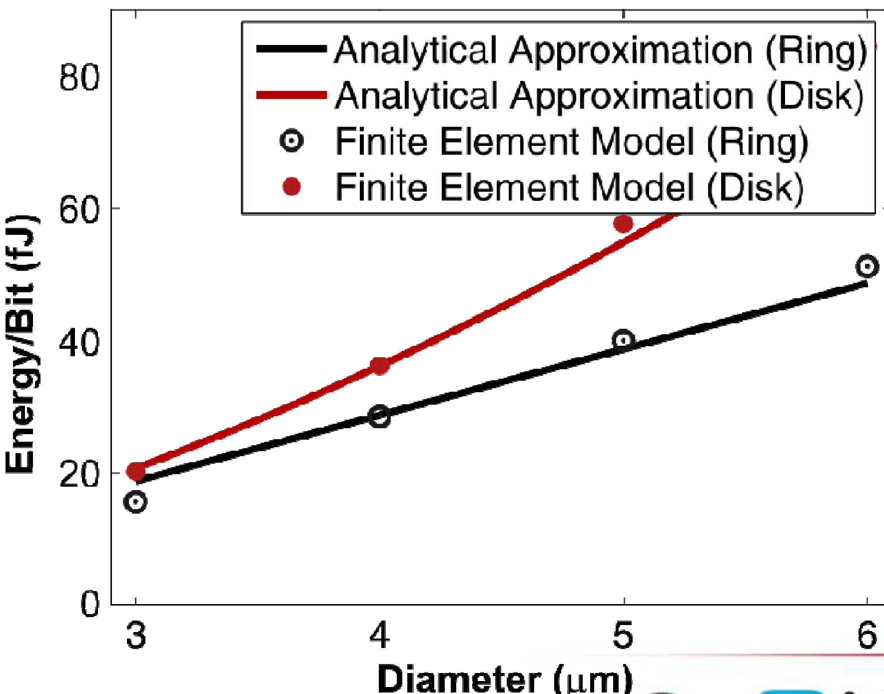

Quick Review of Modulators

Modulators

- **Potential:** Replacement for electrical interconnects with >100X increase in bandwidth density and >100X reduction in energy/bit. At exascale we would be looking at 16MW of communication power instead of 1.6GW of communications power.
- **Demonstration:** Smallest, highest speed (10Gb/s), lowest energy/bit (85fJ), and lowest voltage (3.5V) resonant silicon modulators
- **Future Work:** Improved doping profiles to reach 20fJ/bit and higher bandwidths, thermal control, and terabit/s links . . .

So, What Else?

- Try coupling some disks together . . .

Reconfigurable Networks: Bandpass Switch

Demonstration

- ❑ 1st demonstration of high-speed, electrically active, silicon bandpass switches
- ❑ Completely shift resonator bandpass out of the channel

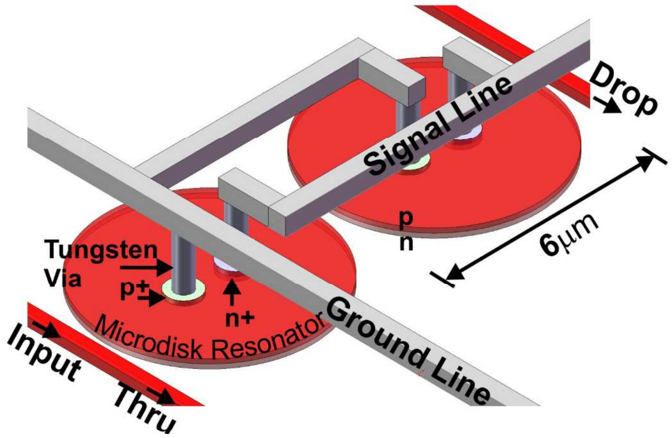

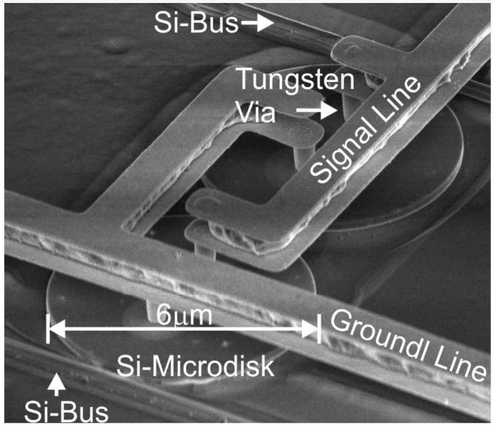

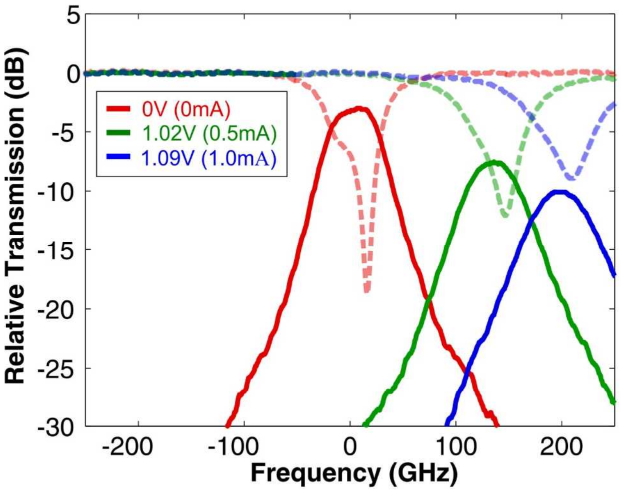

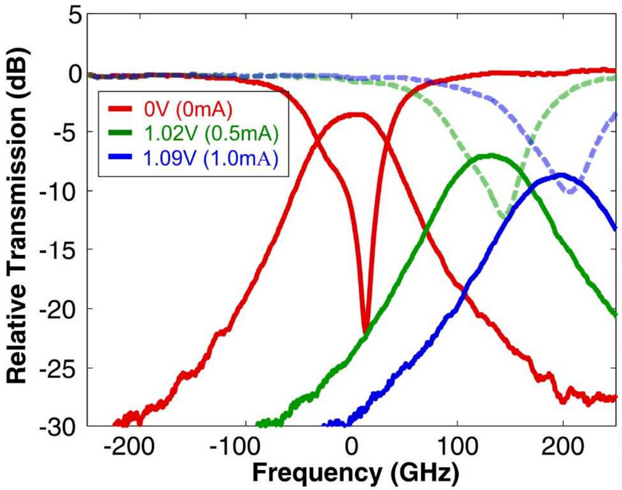

Benefits

- ❑ Potential for ultralow power switching in the optical domain (to avoid OEO conversions)

M. R. Watts et al., OFC Postdeadline Presentation (Feb. 2008)

Sandia National Laboratories

Data Routing with Bandpass Switch

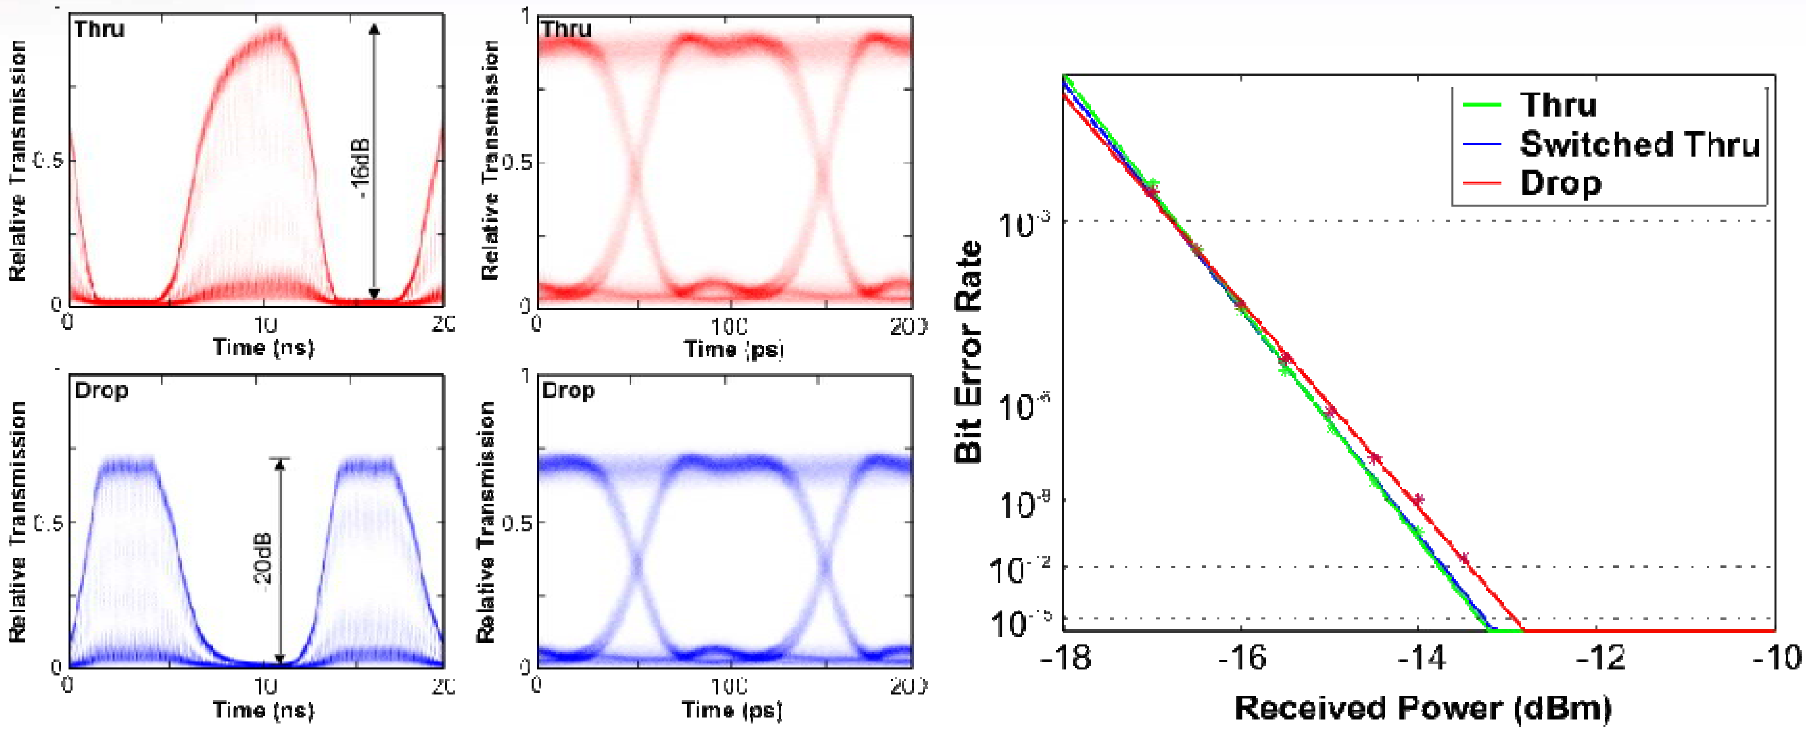

Switch Results

- Data switched error-free (BER<10^{-12}) with little power with ~2ns rise time
- Power penalty measured to be <0.4dB in Drop Port and <0.1dB in Thru Port
- Driven with ~0.6V (~1V due to reflection), so CMOS compatible
- <u>Still, there are some problems with microdisks . . .</u>

Problem 1: Coupling Losses

TE Coupling: R = 3μm

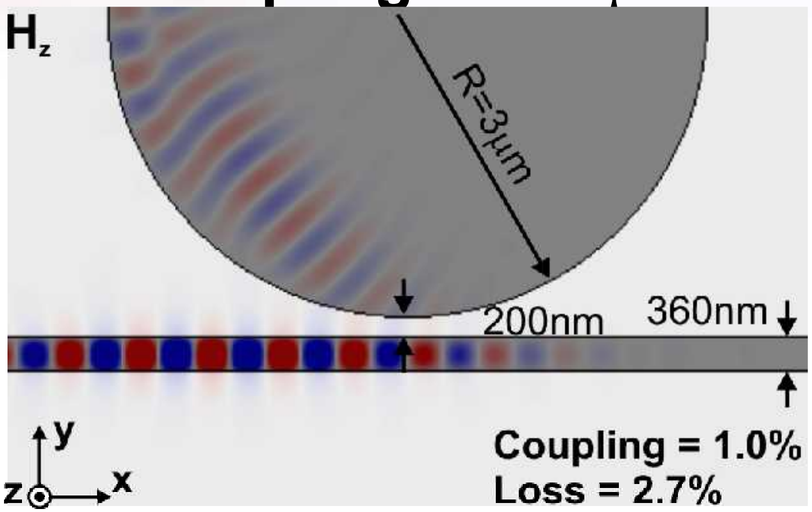

TE_{21} Mode: R = 3μm

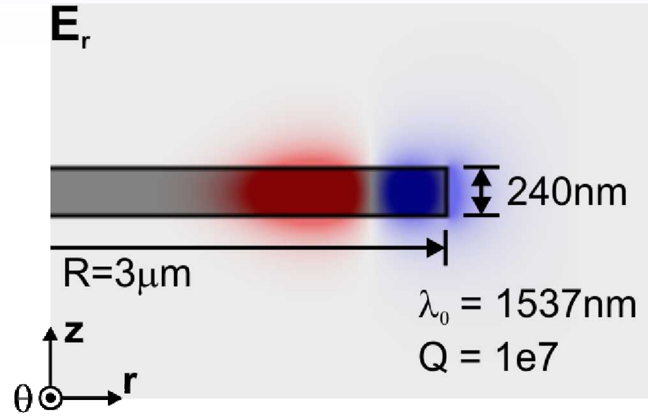

TE Coupling: R = 1.5μm

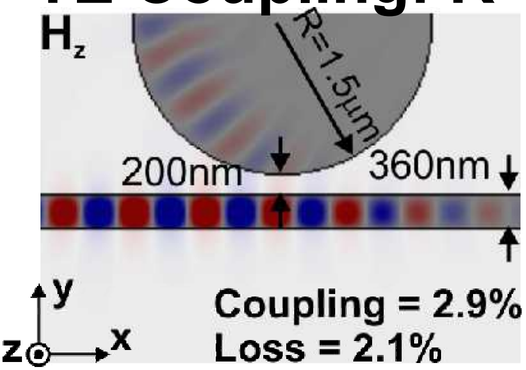

TE_{21} Mode: R = 1.5μm

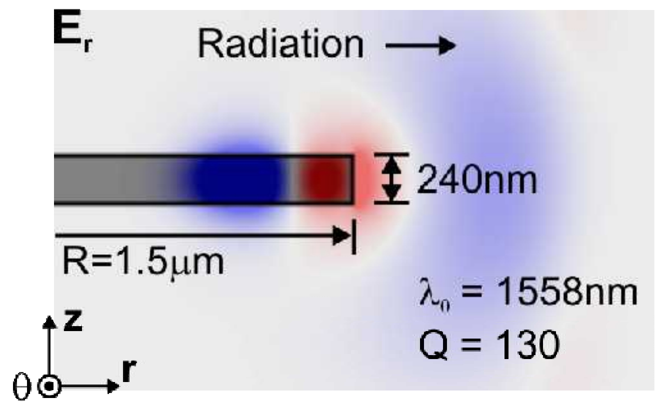

Important Considerations

- Smaller disk ⇒ Lower coupling loss? Why?
- Secondary mode can be cut-off, but requires very small disks

Sandia National Laboratories

Problem 2: Limited Free-Spectral Range

Goal: Large channel count, large bandwidth → Large FSR

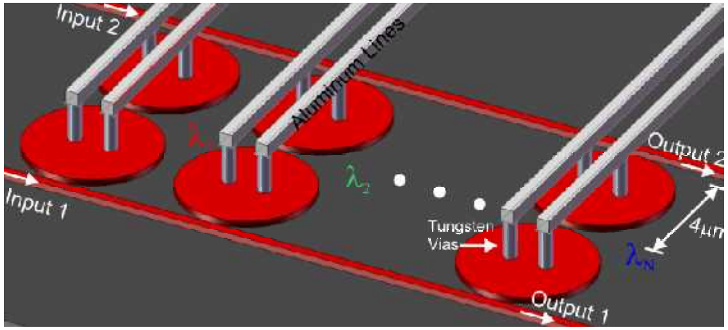

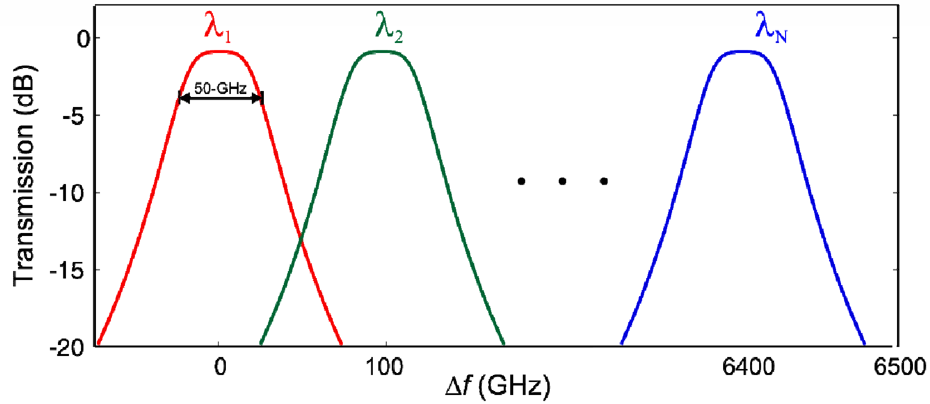

Problem: Microdisks propagate multiple spatial modes, corrupting FSR

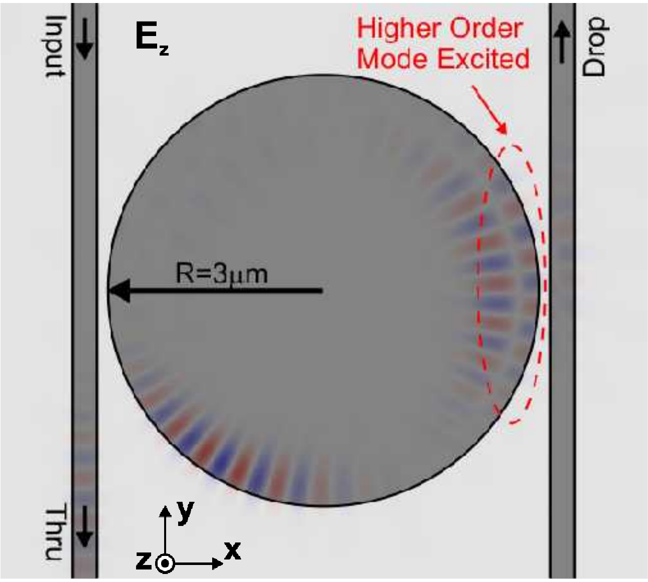

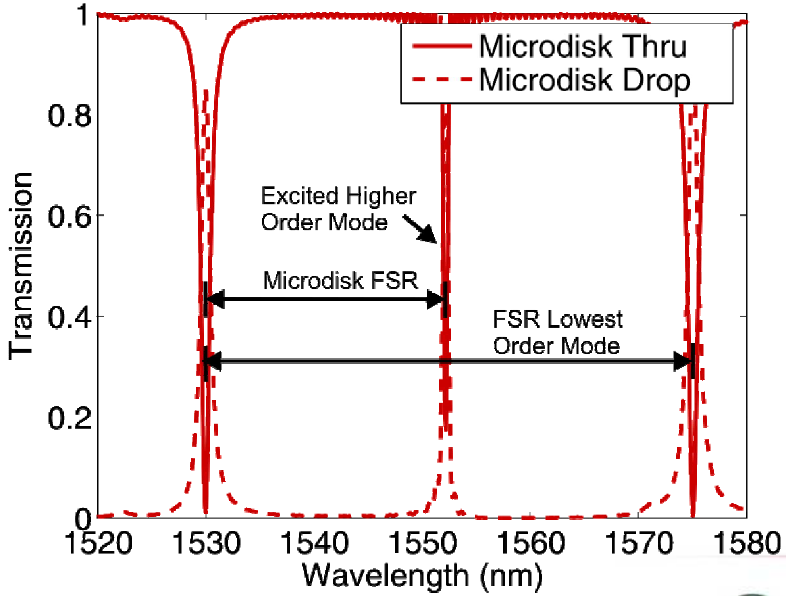

Sandia National Laboratories

Experimentally Observed Microdisk FSR

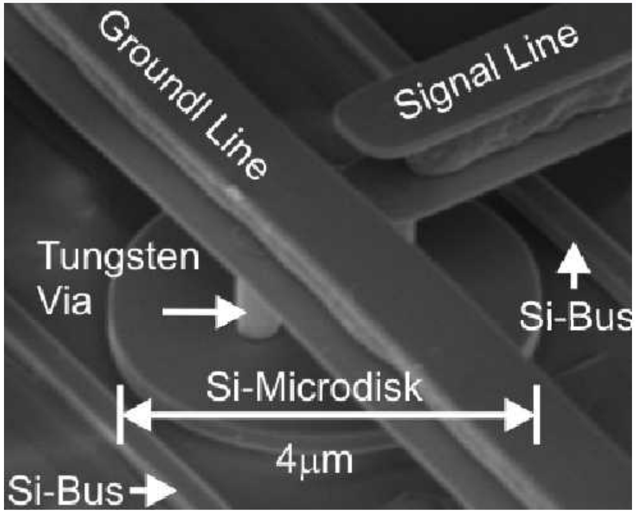

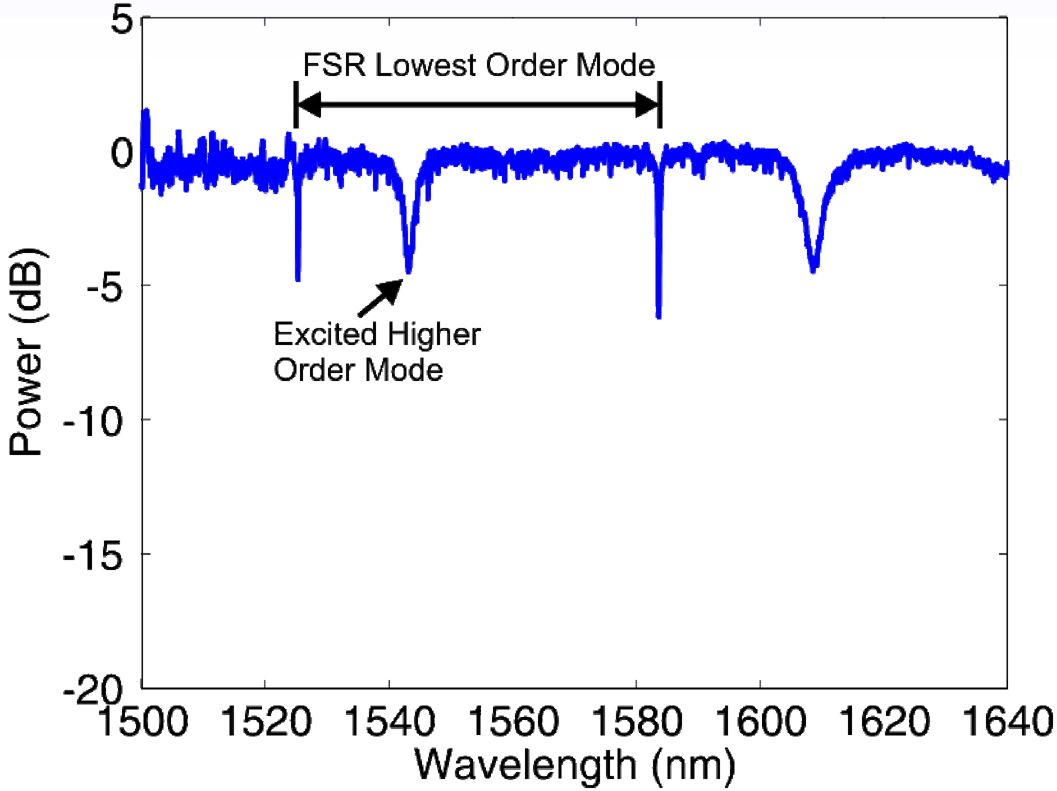

Results

- ❑ As expected, microdisk propagates higher order modes corrupting FSR and limiting the available line bandwidth
- ❑ Can be fixed with a microring, but how do you make electrical contact?

What about Directly Contacting a Microring?

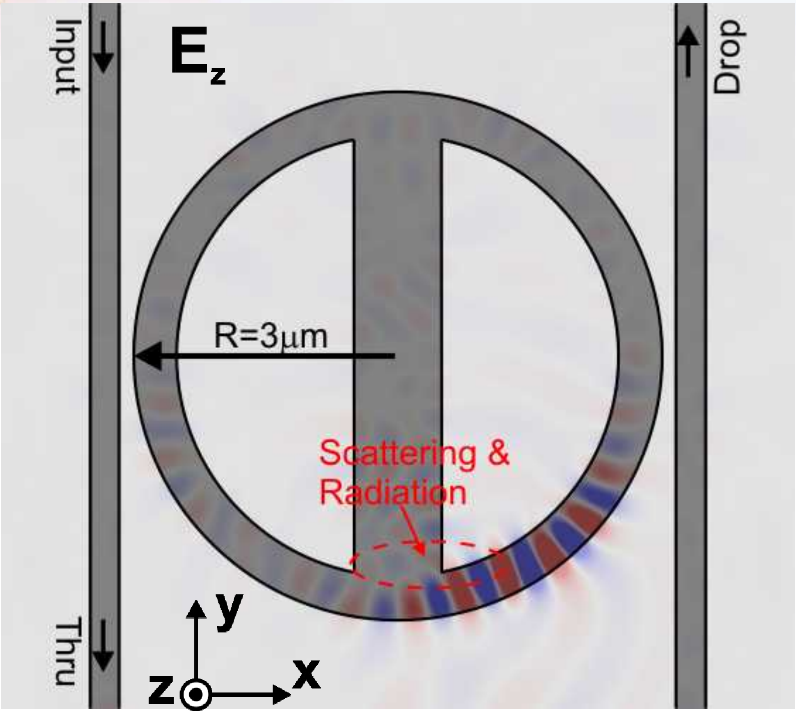

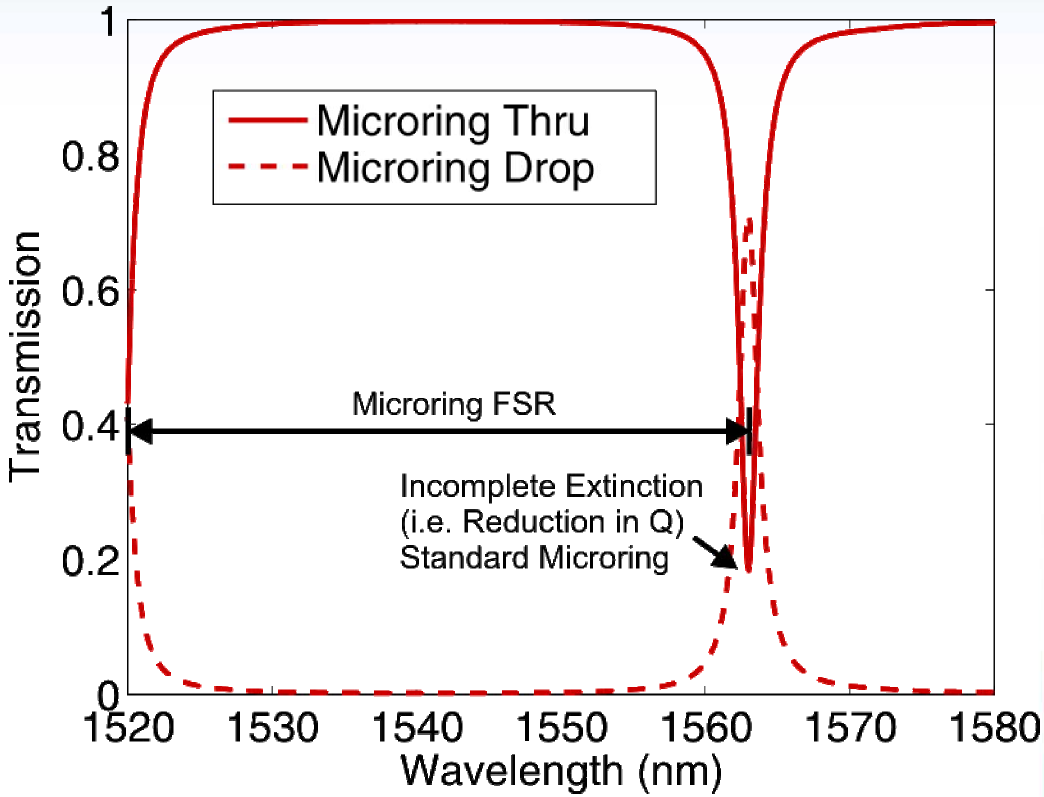

Results

- ❑ Microrings enable a recovery of the full Free Spectral Range
- ❑ However, the contact leads to scattering and a reduction in Q
- ❑ Can we modify the ring geometry to enable contact without loss?

So Let's Consider Mode-Evolution

Thought: Mode-evolution is commonly used to transform modes

Principle: Slow perturbations ($\Delta\beta/\kappa$ >>1) do not result in coupled power

$$P(L) \propto \left|\frac{\kappa}{\Delta\beta}\right|^2 \left[1 - \cos(\Delta\beta L)\right]$$

An Example: A Polarization Rotator

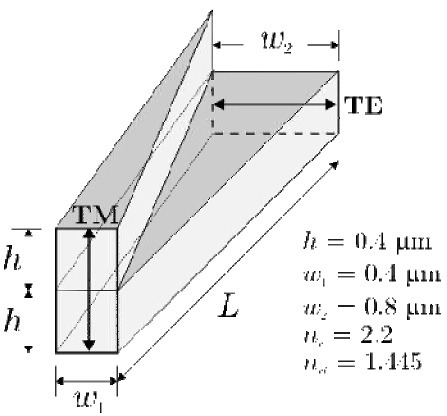

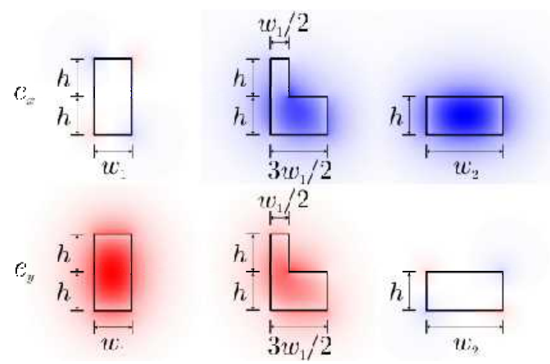

M. R. Watts and H. A. Haus, **30**, *Optics Letters*, 2005

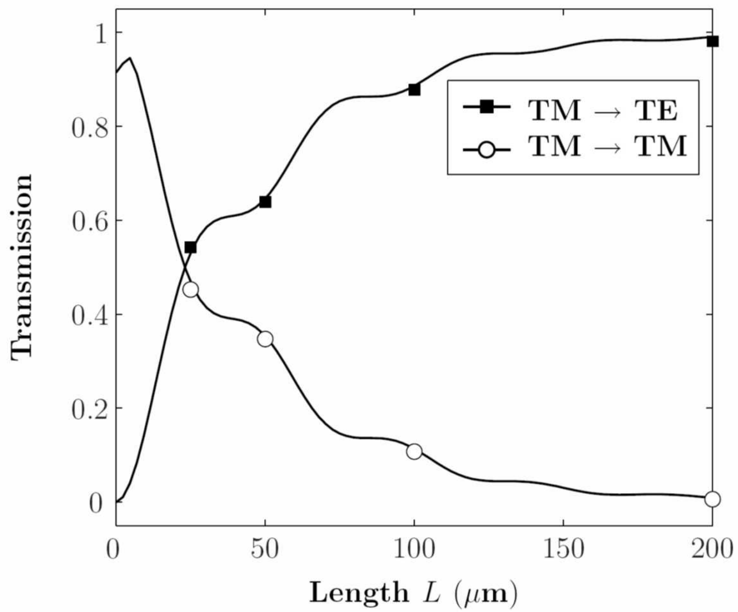

So, what if we introduce an adiabatic taper into a microring to enable a contact without inducing radiation?

Sandia National Laboratories

Simple Contact vs. Adiabatic Bend

Simple Contact

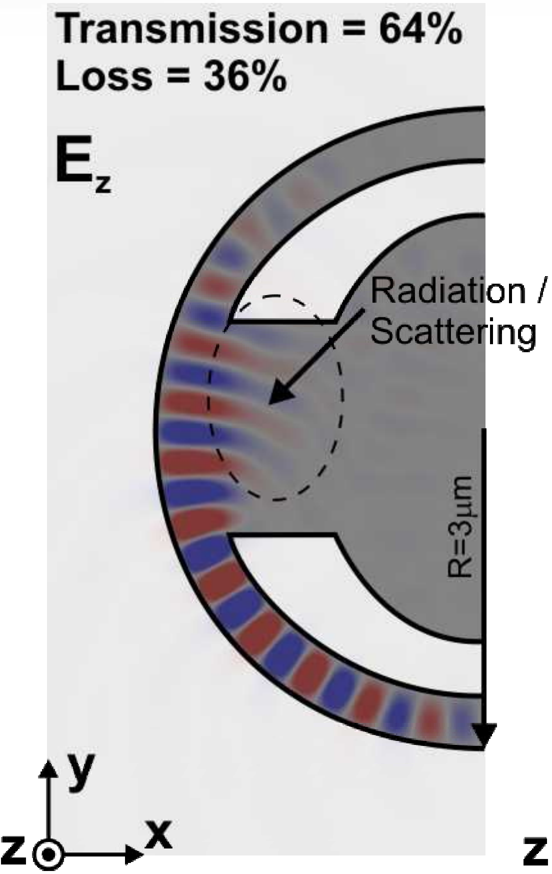

Adiabatic Bend

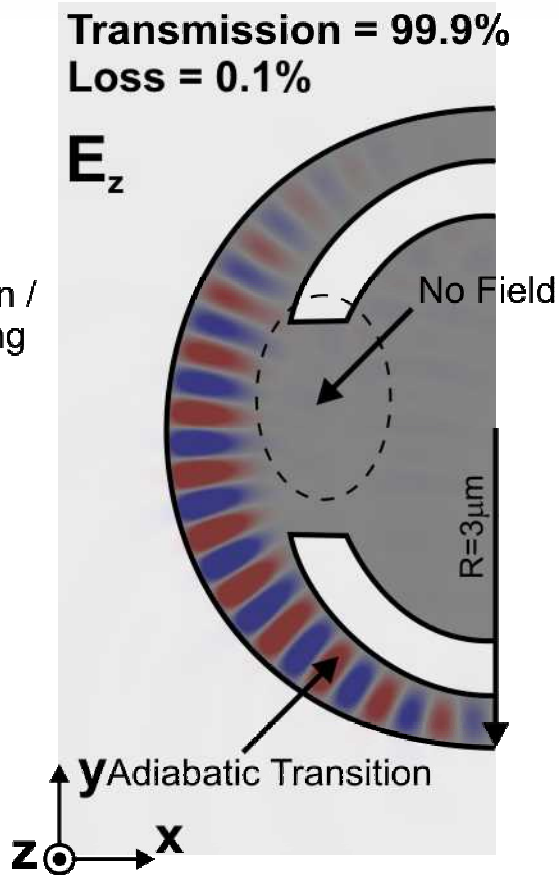

Results

- ❑ Adiabatic bend reduces losses from 36% to 0.1%
- ❑ Loss of 0.1% enables a Q exceeding 10^4
- ❑ Further optimization should enable lower losses
- ❑ Surprising how rapidly "adiabatic" transition is made (enabled by large difference in propagation rates between modes)

Adiabatic Microring-Resonator Q

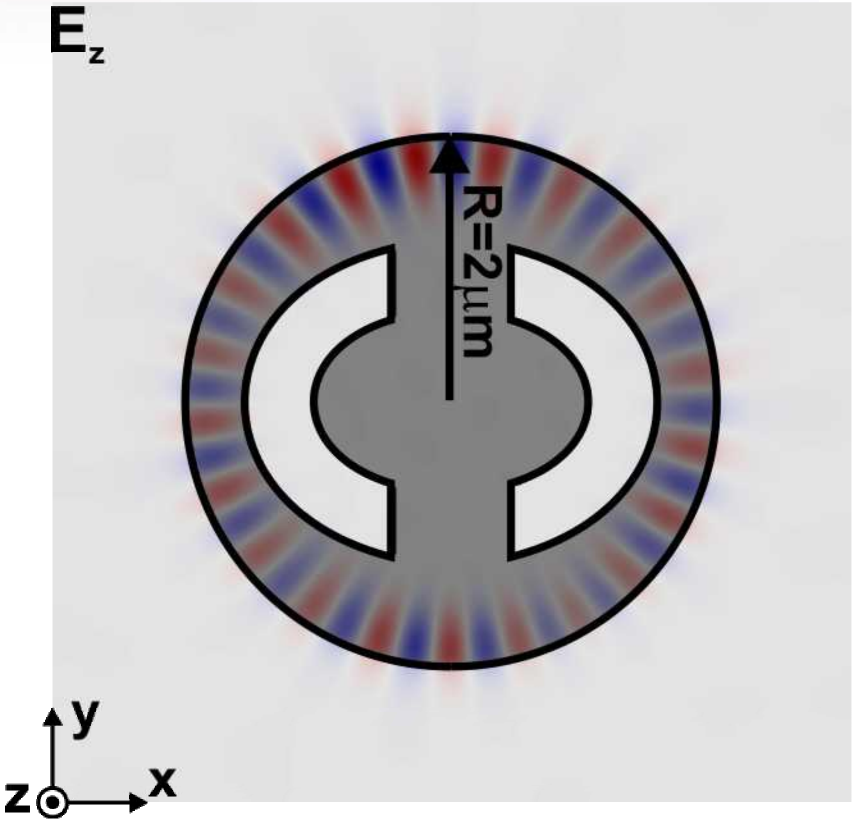

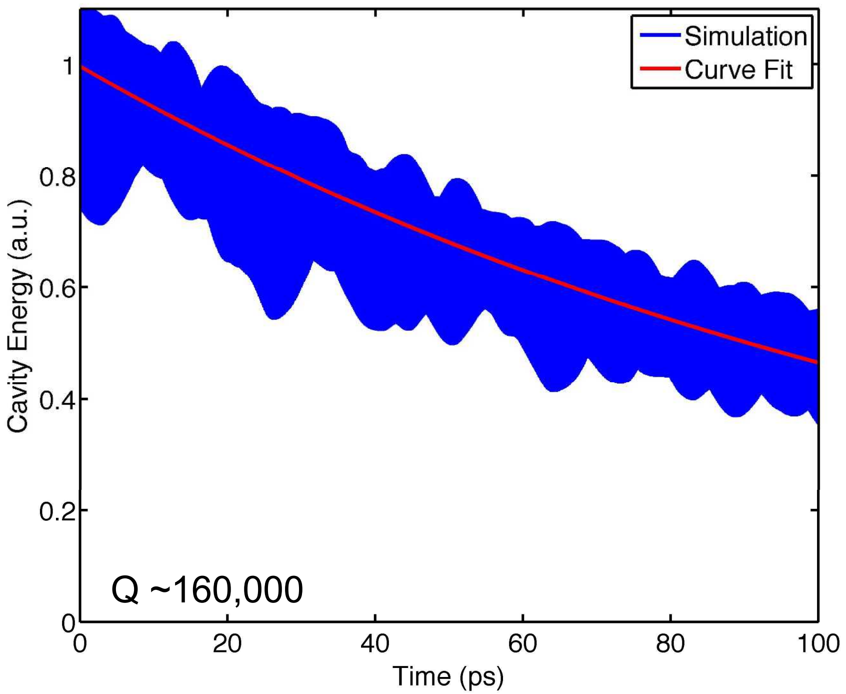

Results

- Internal Q's exceeding 10^5 are possible in small adiabatic microrings
- Q's exceeding 10^6 are likely by iterating on designs

Sandia National Laboratories

Coupling Losses: Essentially Eliminated

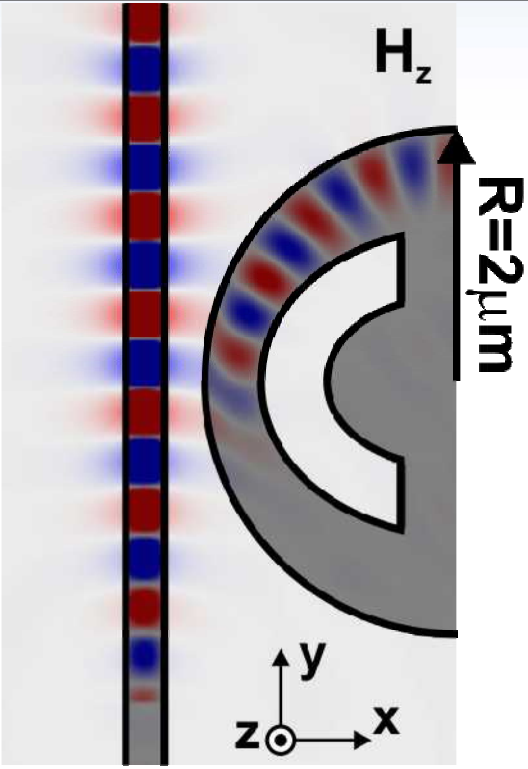

Coupling = 1.5%
Loss = 0.16%

Approach / Results

- ❑ Higher order modes are eliminated in coupling region → lossless coupling
- ❑ Narrower bus than ring waveguide used to mode-match / suppress coupling to lossy supermodes with the structure*,**

*M. R. Watts, PhD MIT Thesis (2005)

M.A. Popović et al., Opt. Lett., **31, pp. 2571-2573 (2006)

Free Spectral Range: Fully Recovered

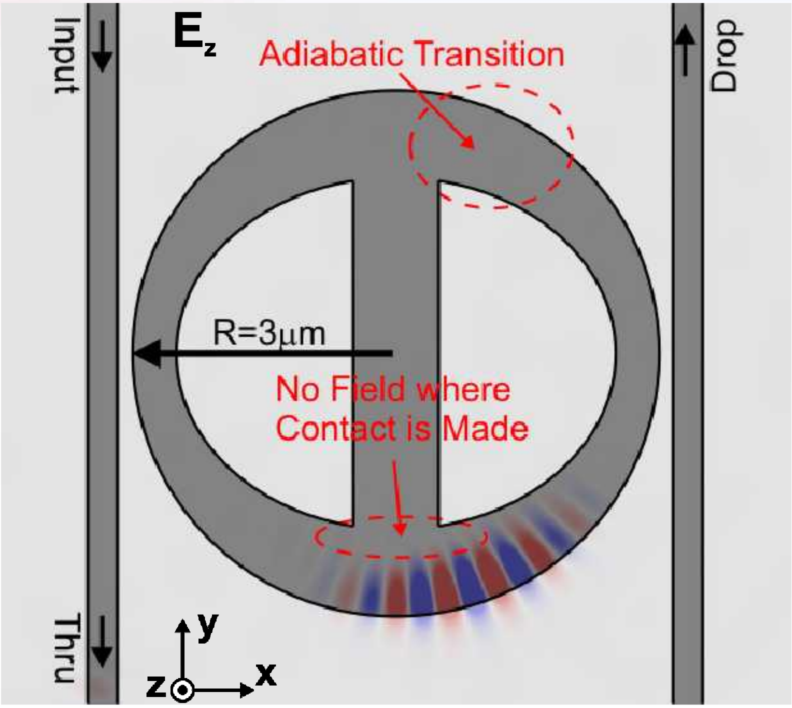

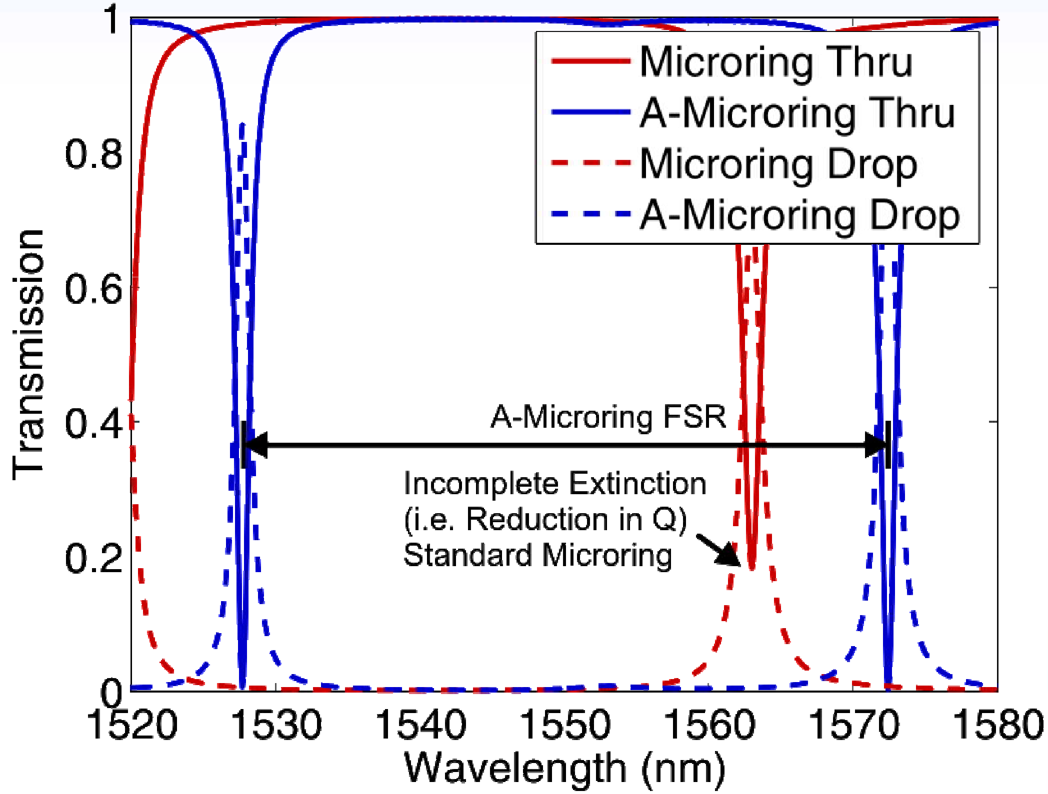

Results

- ❑ Adiabatic microrings enable a recovery of the full Free Spectral Range without inducing scattering and loss

Adiabatic Microring: Experimental Results

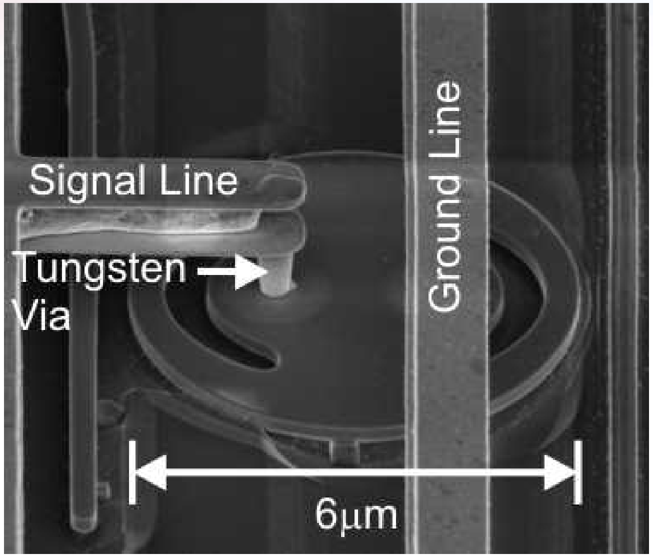

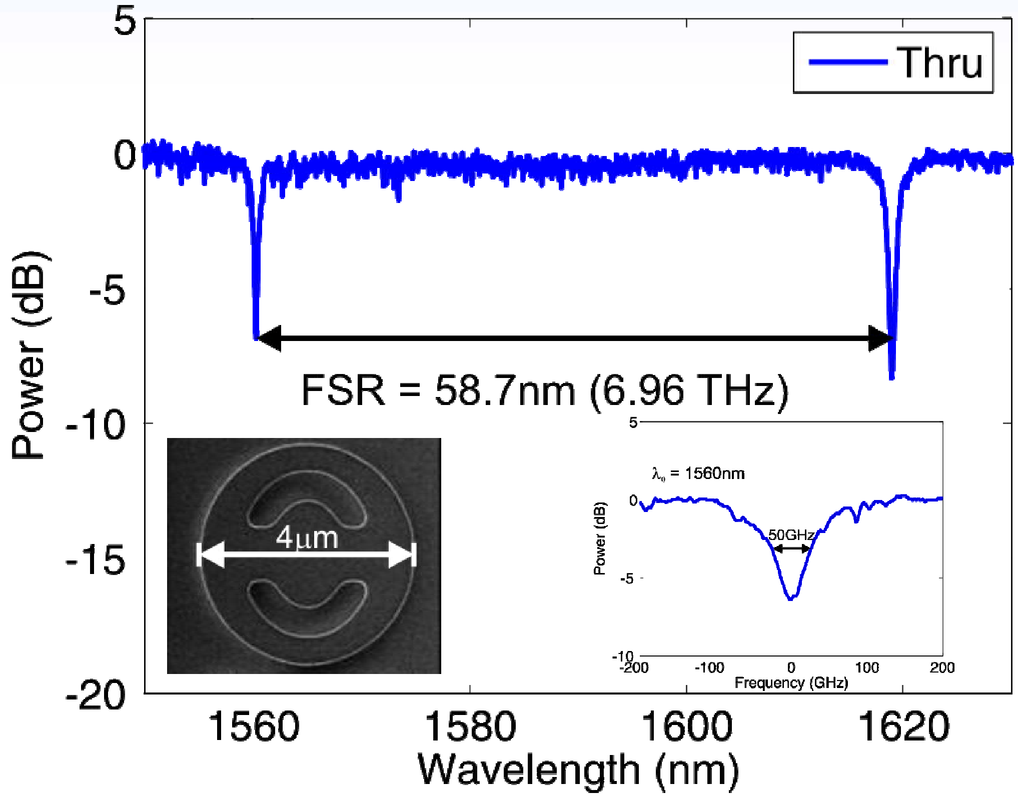

Results

- ❑ Uncorrupted 6.96THz Free-Spectral-Range
- ❑ Eliminated higher order modes without significant loss (Q_{ext}~4000)
- ❑ Slight reduction in Q due to fabrication bias (i.e. loss of ~100nm)
- ❑ Electrically active, 4μm rings currently have insufficient Q, 6μm rings have high-Q, testing now . . .

Sandia National Laboratories

Summary and Future Outlook

Modulators

- ❑ Demonstrated error-free, 10Gb/s NRZ data transmission (no pre-emphasis)
- ❑ Communications efficiency of 85fJ/bit **(100X less than electrical interchip)**
- ❑ Path to ~10fJ/bit with <2.5V drive is highly probable . . . **enabling 1-Tb/s @10mW**

1×N Silicon Modulators

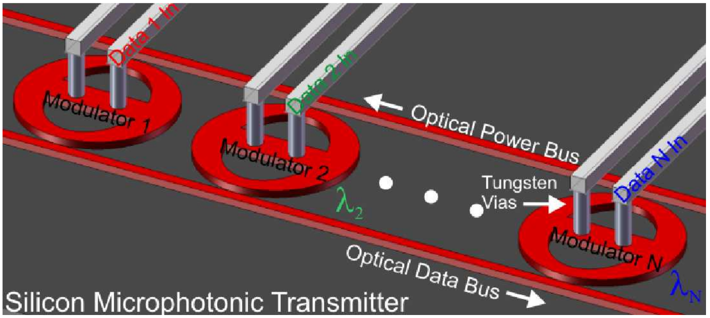

1×N Bandpass Switches

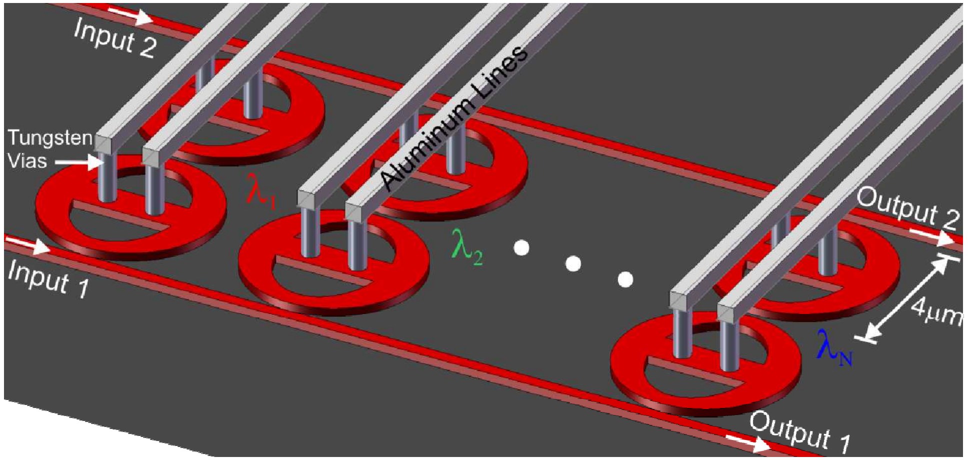

Bandpass Switches

- ❑ First demonstrated electrically active high-speed (~2ns) silicon bandpass switch
- ❑ Can be driven with ~1V drive
- ❑ Did not require any post-fabrication trimming

Challenges Ahead

- ❑ Temperature control, fabrication tolerances, dense integration

Sandia National Laboratories

Acknowledgements

FDTD Code: Christina Manolatou

Cylindrical Modesolver: Milos Popovic

Funding: Sandia LDRD and DARPA MTO (M. Haney and J. Shah)

Sandia National Laboratories

Energy/Bit in a Capacitive Modulator

Definition: $Energy/Bit = \frac{1}{BR \cdot T}\int_0^T P(t)dt,\quad where\ BR = Bit\ Rate$

Total Bits

Energy

0-to-1 Transition

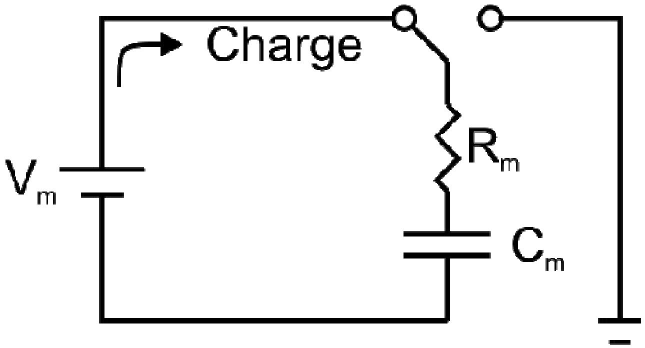

CV^2 Energy Supplied
($CV^2/2$ Dissipated, $CV^2/2$ Stored)

1-to-0 Transition

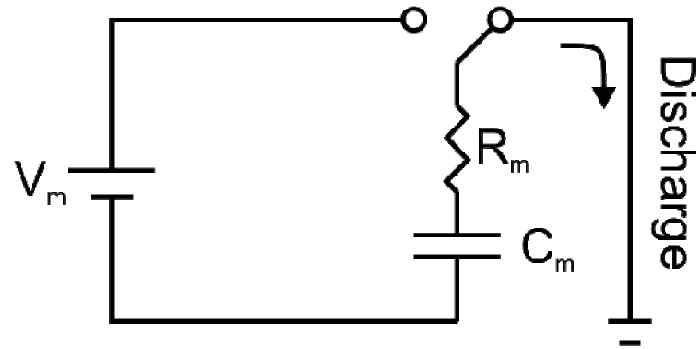

0 Energy Supplied
($CV^2/2$ Dissipated)

- **Switching Energy:** *$E_s = CV^2$ ($CV^2/2$ dissipated, $CV^2/2$ stored)*
- **Clocking Energy/"Bit":** *$CV^2/2$ (because 0-1-0-1-0-1, etc.)*
- **E(energy/bit|0)** = P(0-1|0-state)*$energy_{diss}$(0-1)+P(0-0|0-state)*$energy_{diss}$(0-0)
  = 0.5*CV^2/2+0.5*0 = $CV^2/4$
- **E(energy/bit|1)** = P(1-0|1-state)*$energy_{diss}$(1-0)+P(1-1|1-state)*$energy_{diss}$(1-1)
  = 0.5*CV^2/2+0.5*0 = $CV^2/4$
- **PRBS E(energy/bit)** = P(0)*E(energy/bit|0)+P(1)*E(energy/bit|1) = $CV^2/4$

Sandia National Laboratories