

# Ayar Labs

## Deeply integrated photonic I/O

---

Mark Wade | 11/13/2016

EEHPC Workshop 2016



# The exascale gap

	Optical link energy use		Interconnect cost	
Today's HPC systems	50 pJ/bit		\$5/Gbps	
Exascale projection	12 MW (on 25 MW total system target)		\$1.2 Billion	
Assumptions:	Assumptions:		Assumptions:	
<ul style="list-style-type: none"><li>• <math>10^{18}</math> Flops</li><li>• 0.01 Bytes/Flop injection bandwidth</li><li>• <math>10^{16}</math> Bps total injection bandwidth</li></ul>	<ul style="list-style-type: none"><li>• 50 pJ/bit <math>\times</math> 8 = 400 pJ/B</li><li>• 400 mW/GBps</li><li>• 4 MW injection bandwidth</li><li>• <math>\times</math> 3 for network bisection</li></ul>		<ul style="list-style-type: none"><li>• \$5/Gbps <math>\times</math> 8 = \$40/GBps</li><li>• <math>\times</math> <math>10^{16}</math> Bps = \$400M</li><li>• <math>\times</math> 3 for network bisection</li></ul>	
Exascale target	<1.2 MW	<5 pJ/b	\$60 Million	\$0.10/Gbps



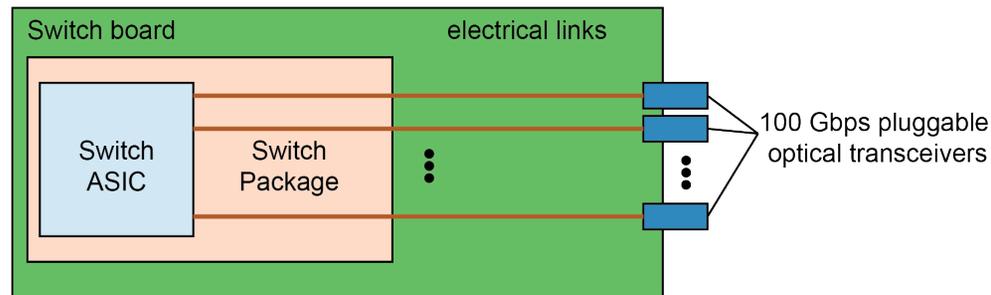
# The problem with electrical SerDes

	Optical link energy use	Interconnect cost
Overall targets	<5 pJ/b	\$0.10/Gbps

**The problem**      **Electrical SerDes**

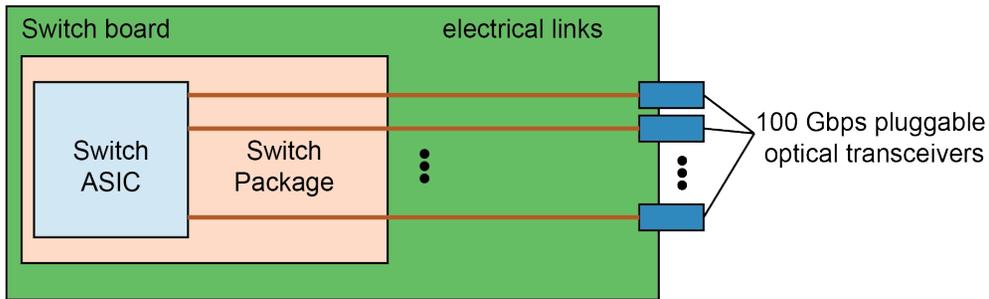
~ 10 pJ/b

- Switch ASICs consume ~ 33% of their power on SerDes
- Optical transceivers consume ~ 50% of their power on SerDes





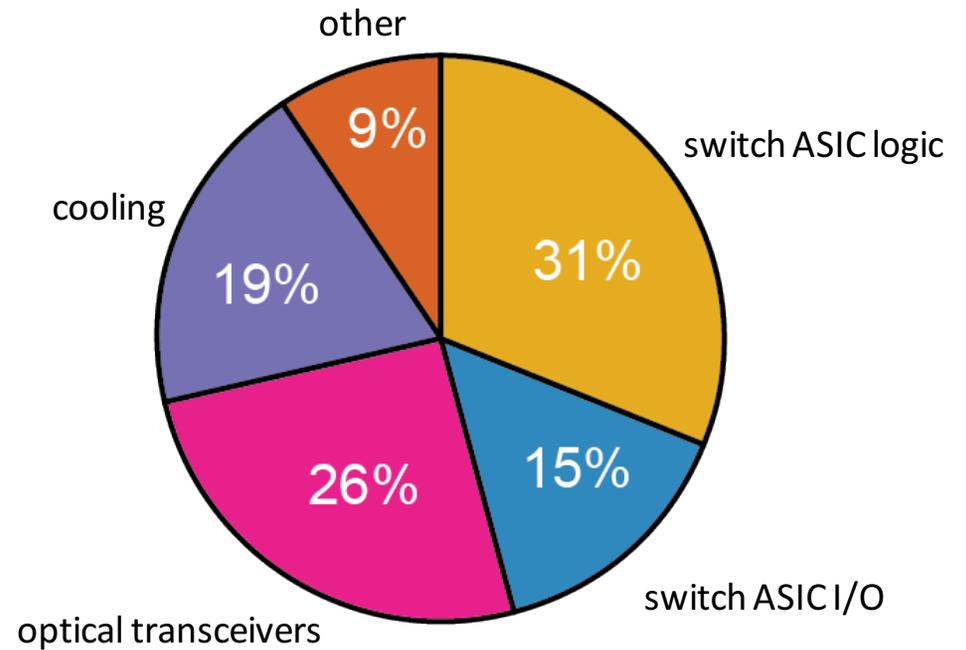
# Dealing with SerDes power: Today's optical links



32x100G Switch Baseline

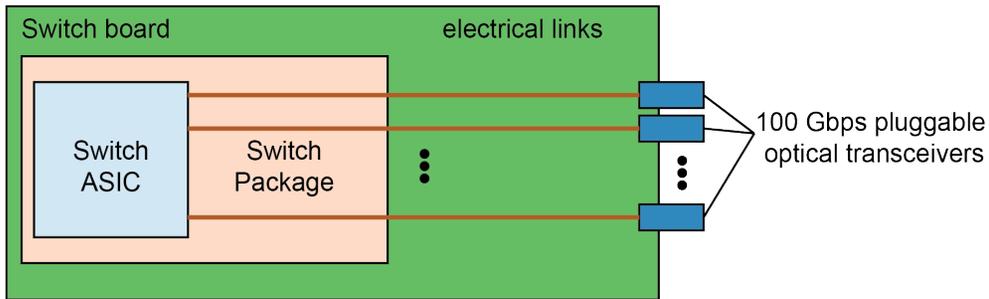
436W power

100% (baseline)





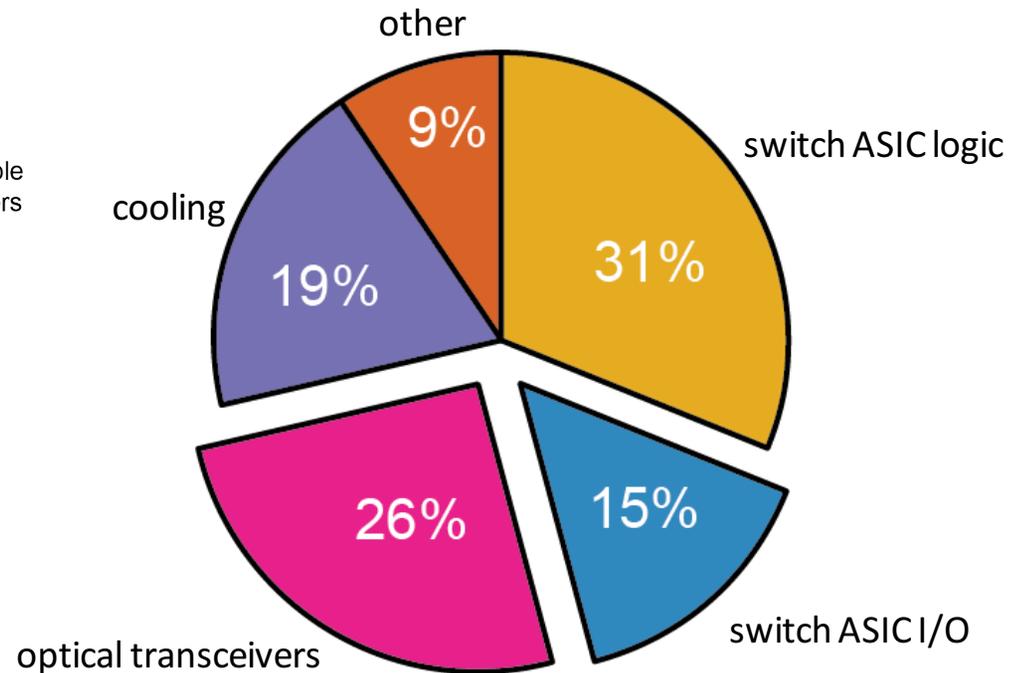
# Dealing with SerDes power: Today's optical links



32x100G Switch Baseline

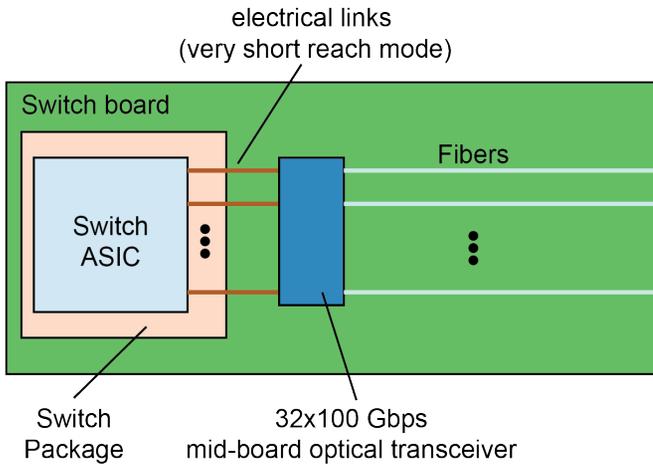
436W power

100% (baseline)





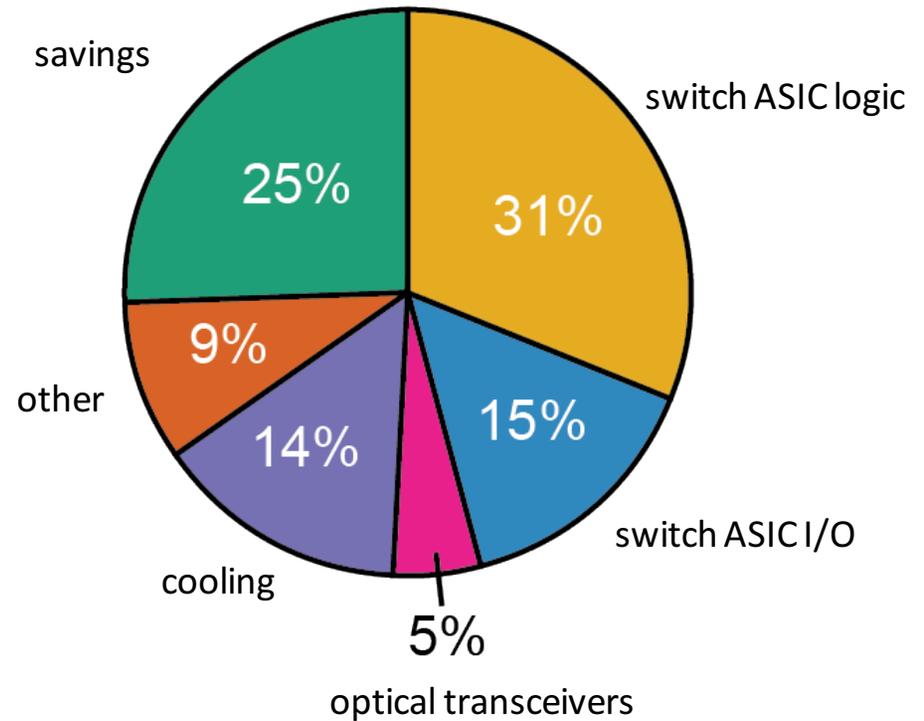
# Dealing with SerDes power: Mid-board optical links



32x100G Switch  
with mid-board optics

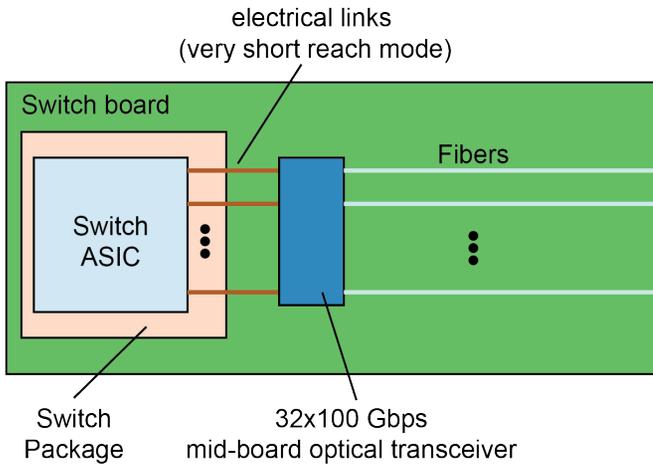
326W power

75% energy use





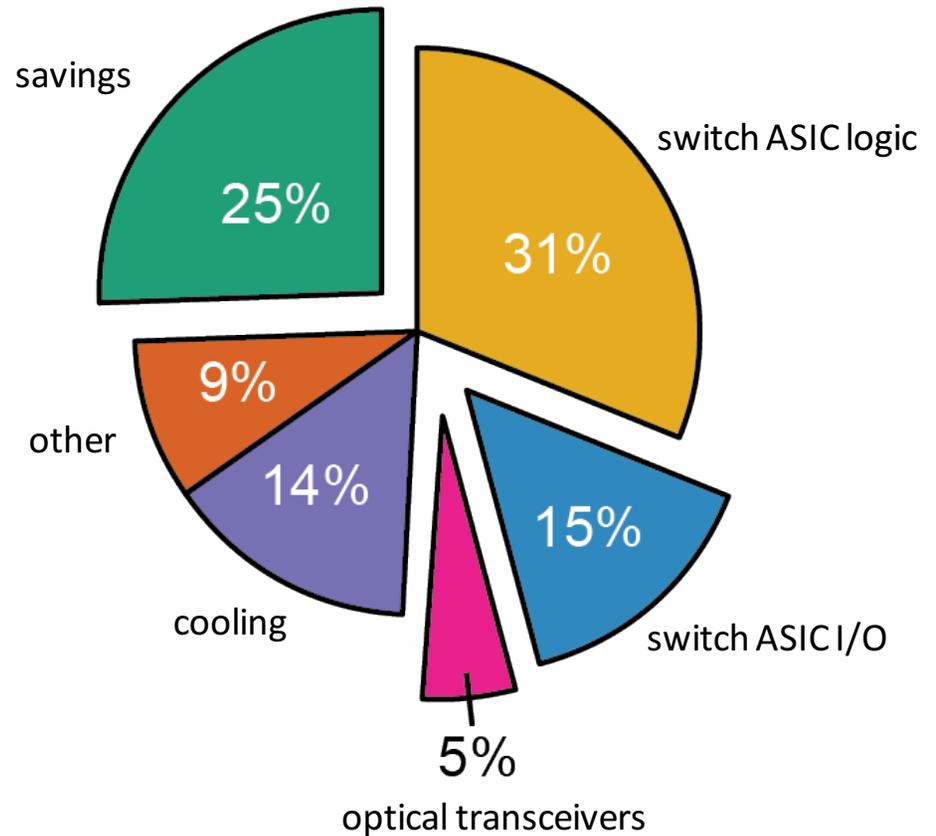
# Dealing with SerDes power: Mid-board optical links



32x100G Switch  
with mid-board optics

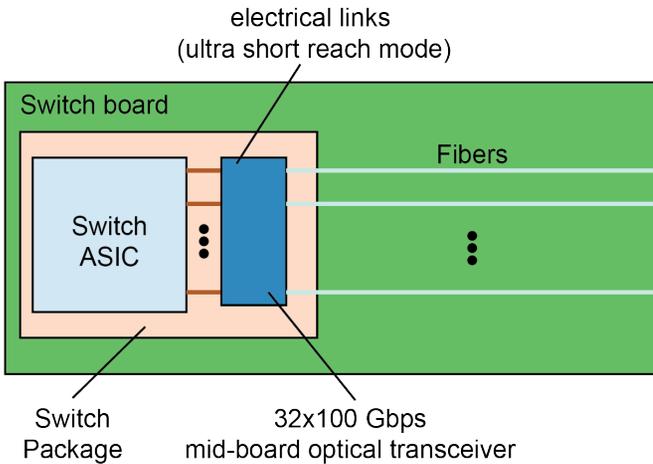
326W power

75% energy use





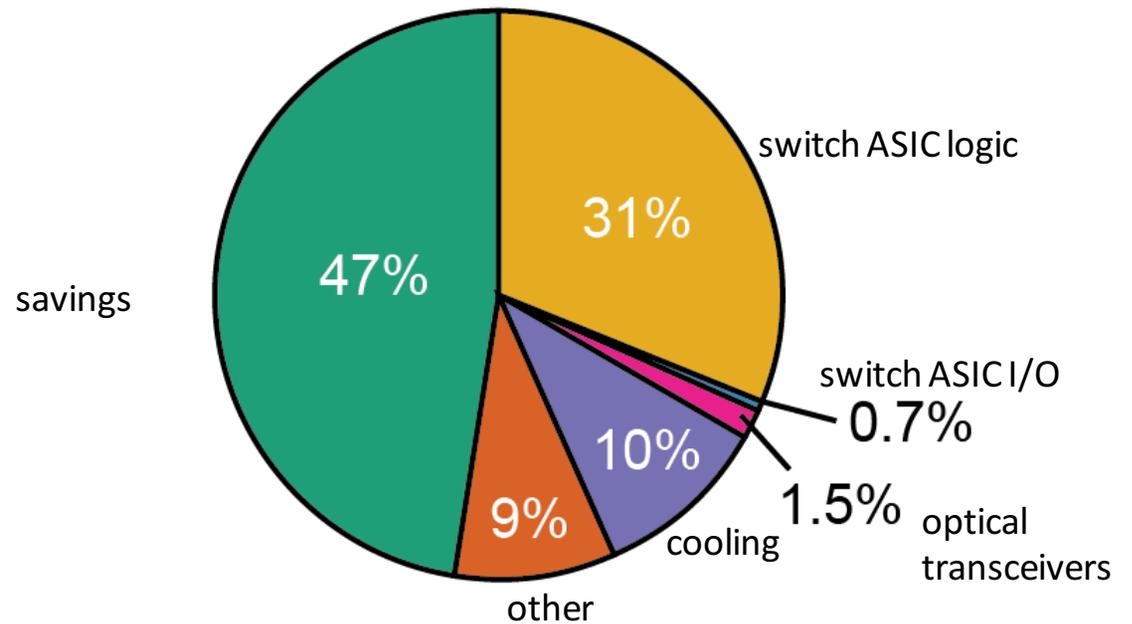
# Dealing with SerDes power: In-package optical links



32x100G Switch  
with co-packaged optics

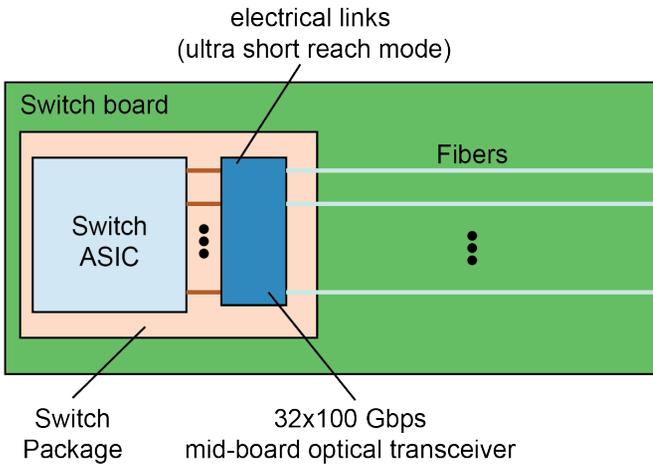
230W power

53% energy use





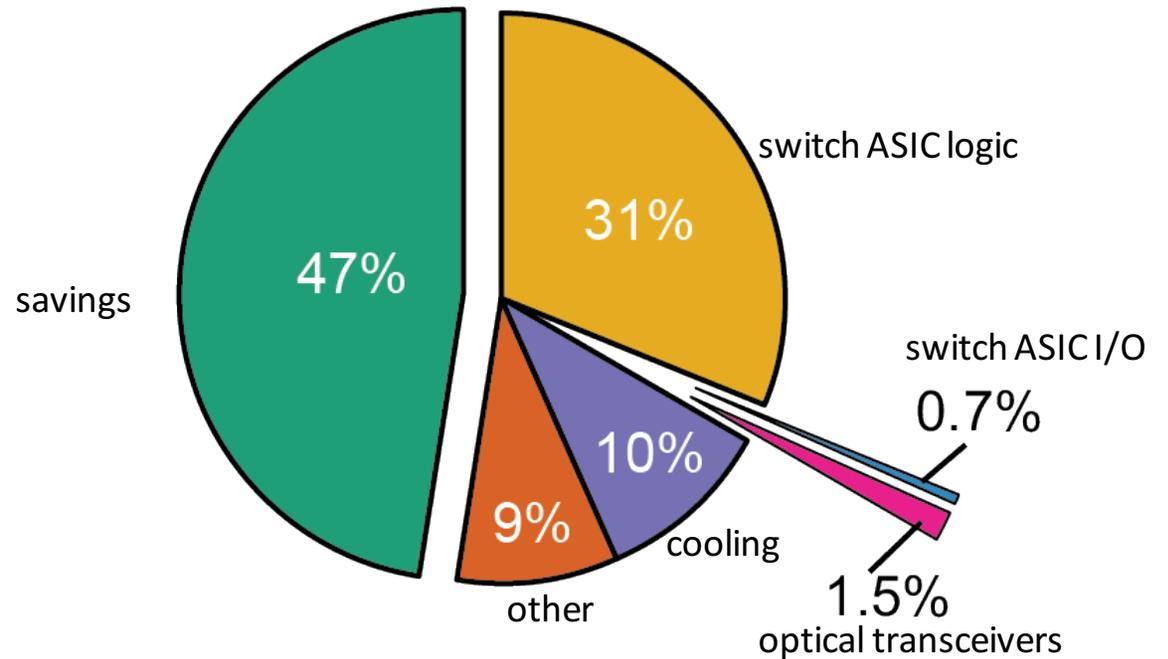
# Dealing with SerDes power: In-package optical links



32x100G Switch  
with co-packaged optics

230W power

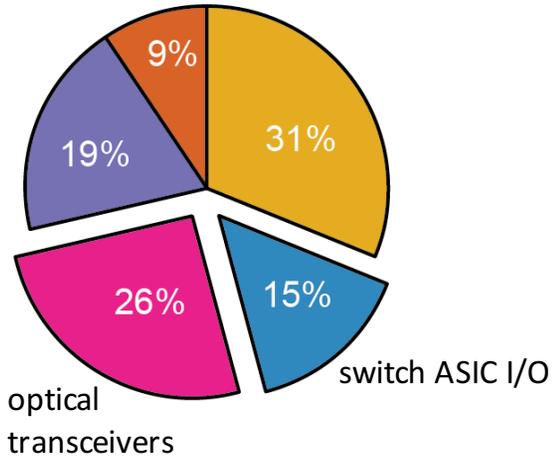
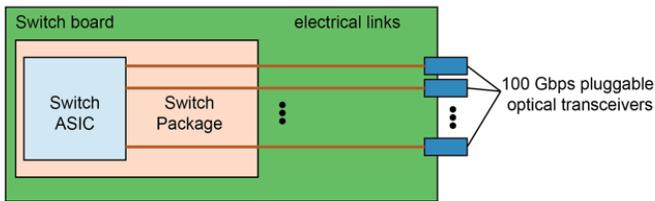
53% energy use





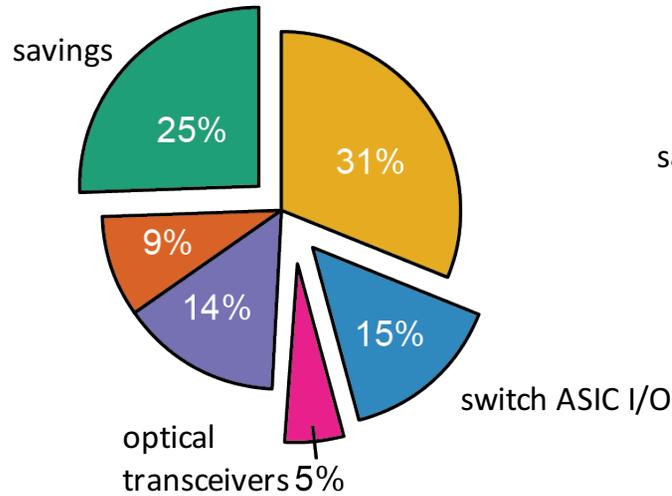
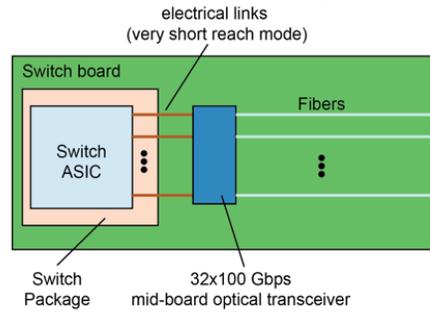
# Bringing optics into the server results in **25% - 50% system power savings**

## Baseline



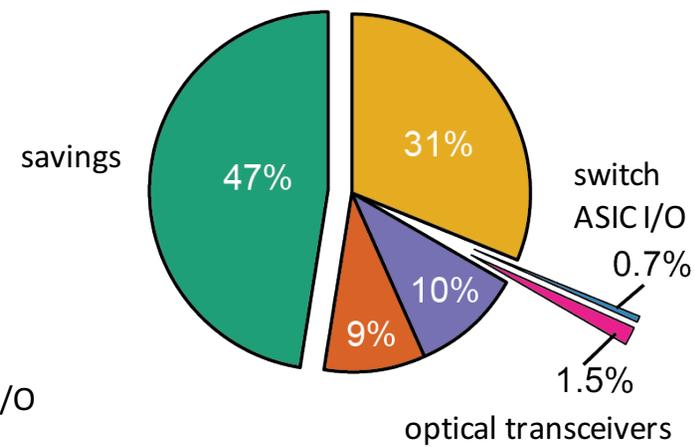
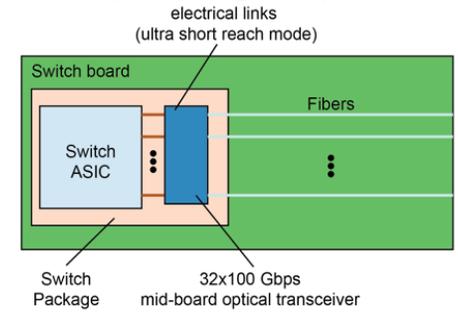
11/15/16

## Mid-board optics



Copyright Ayar Labs, Inc.

## In-package optics



10



Why can't today's optics go inside a CPU or switch package?

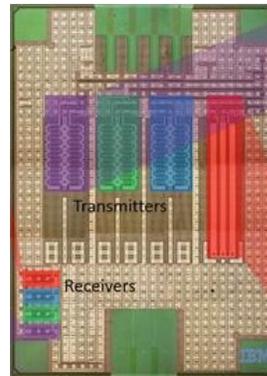
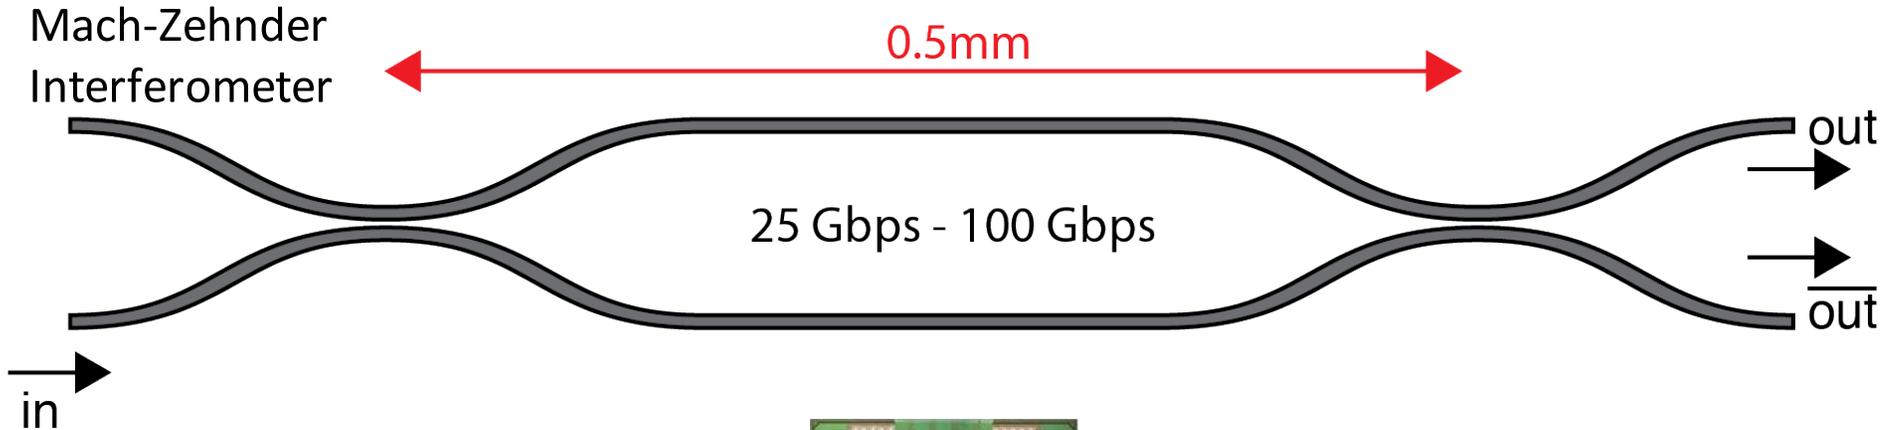
**Bandwidth density**

**Laser integration**



# The problem: low bandwidth density

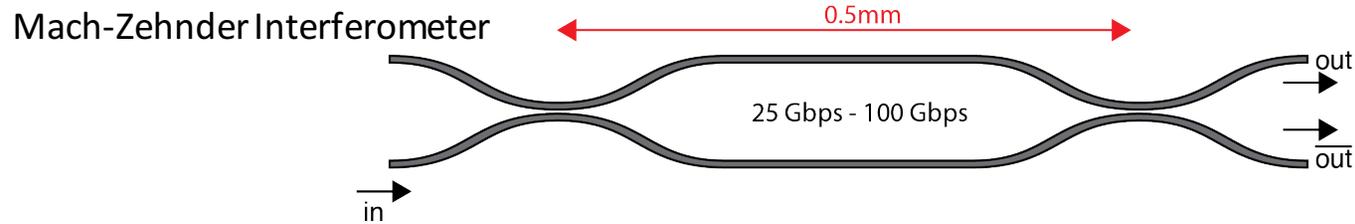
Mach-Zehnder Interferometer



100 Gbps monolithically integrated transceiver [IBM]



# The problem: low bandwidth density

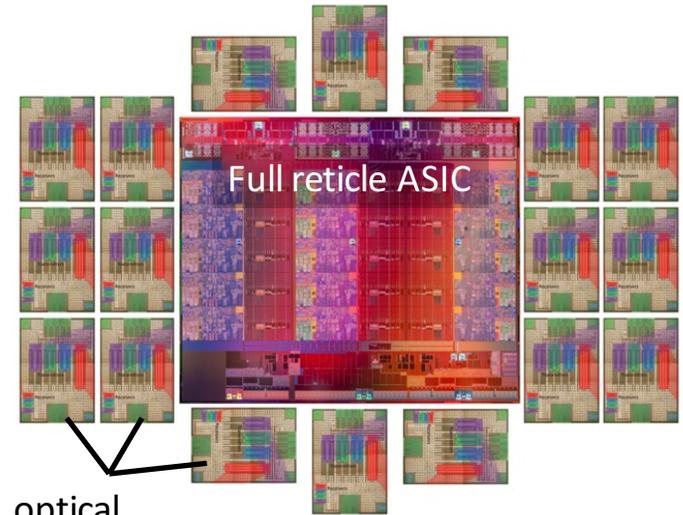


Bandwidth density:  $\sim 100\text{Gbps}/\text{die}$  ( $\sim 10\text{mm} \times 7\text{mm}$ ) =  $\sim 1.4\text{ Gbps}/\text{mm}^2$  [1]

Assume:  $400\text{Gbps}/\text{die}$  ( $\sim 10\text{mm} \times 7\text{mm}$ )

2019 era switches are targeting 6.4 Tbps

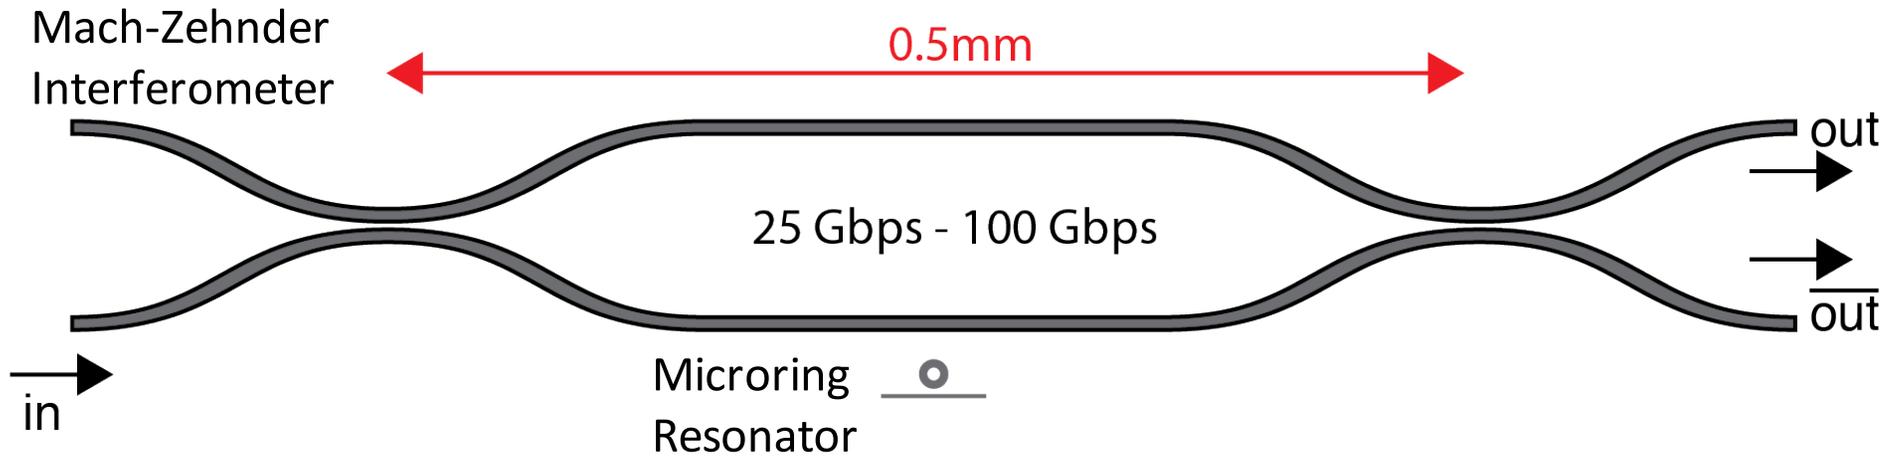
Required number of optical transceiver die: **16**



[1] 100 Gbps transceiver, IBM, OFC 2015

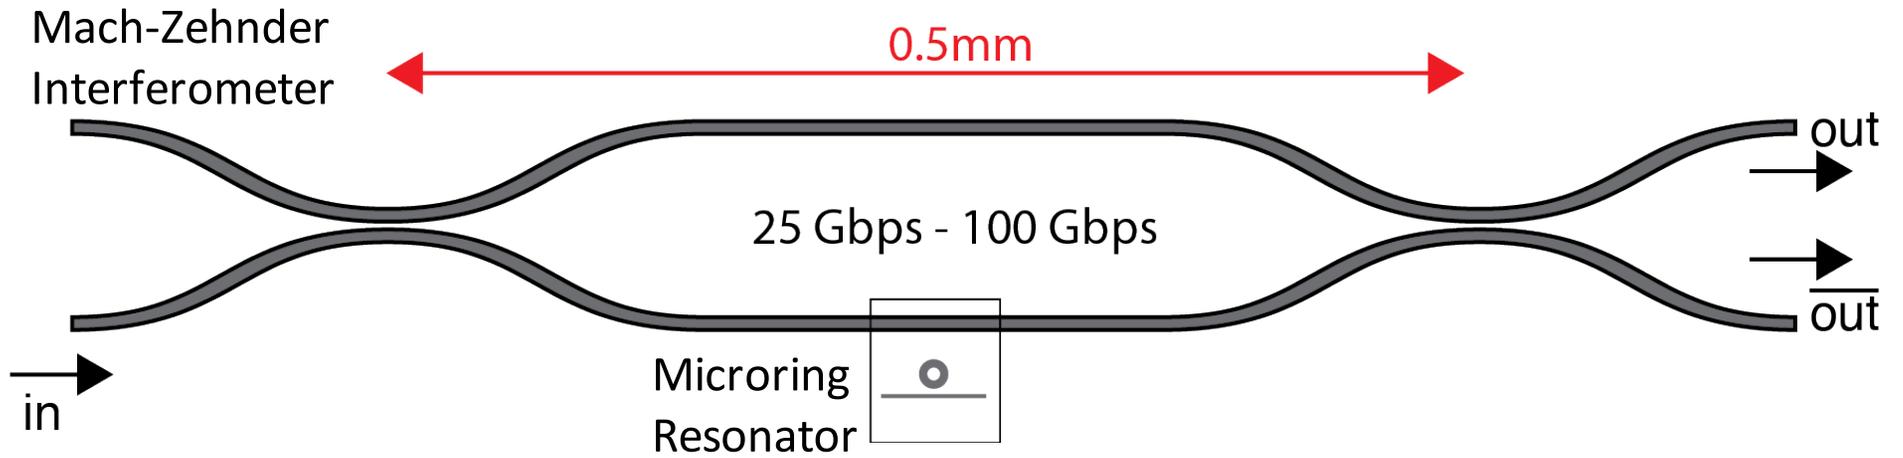


# The solution: microring resonators



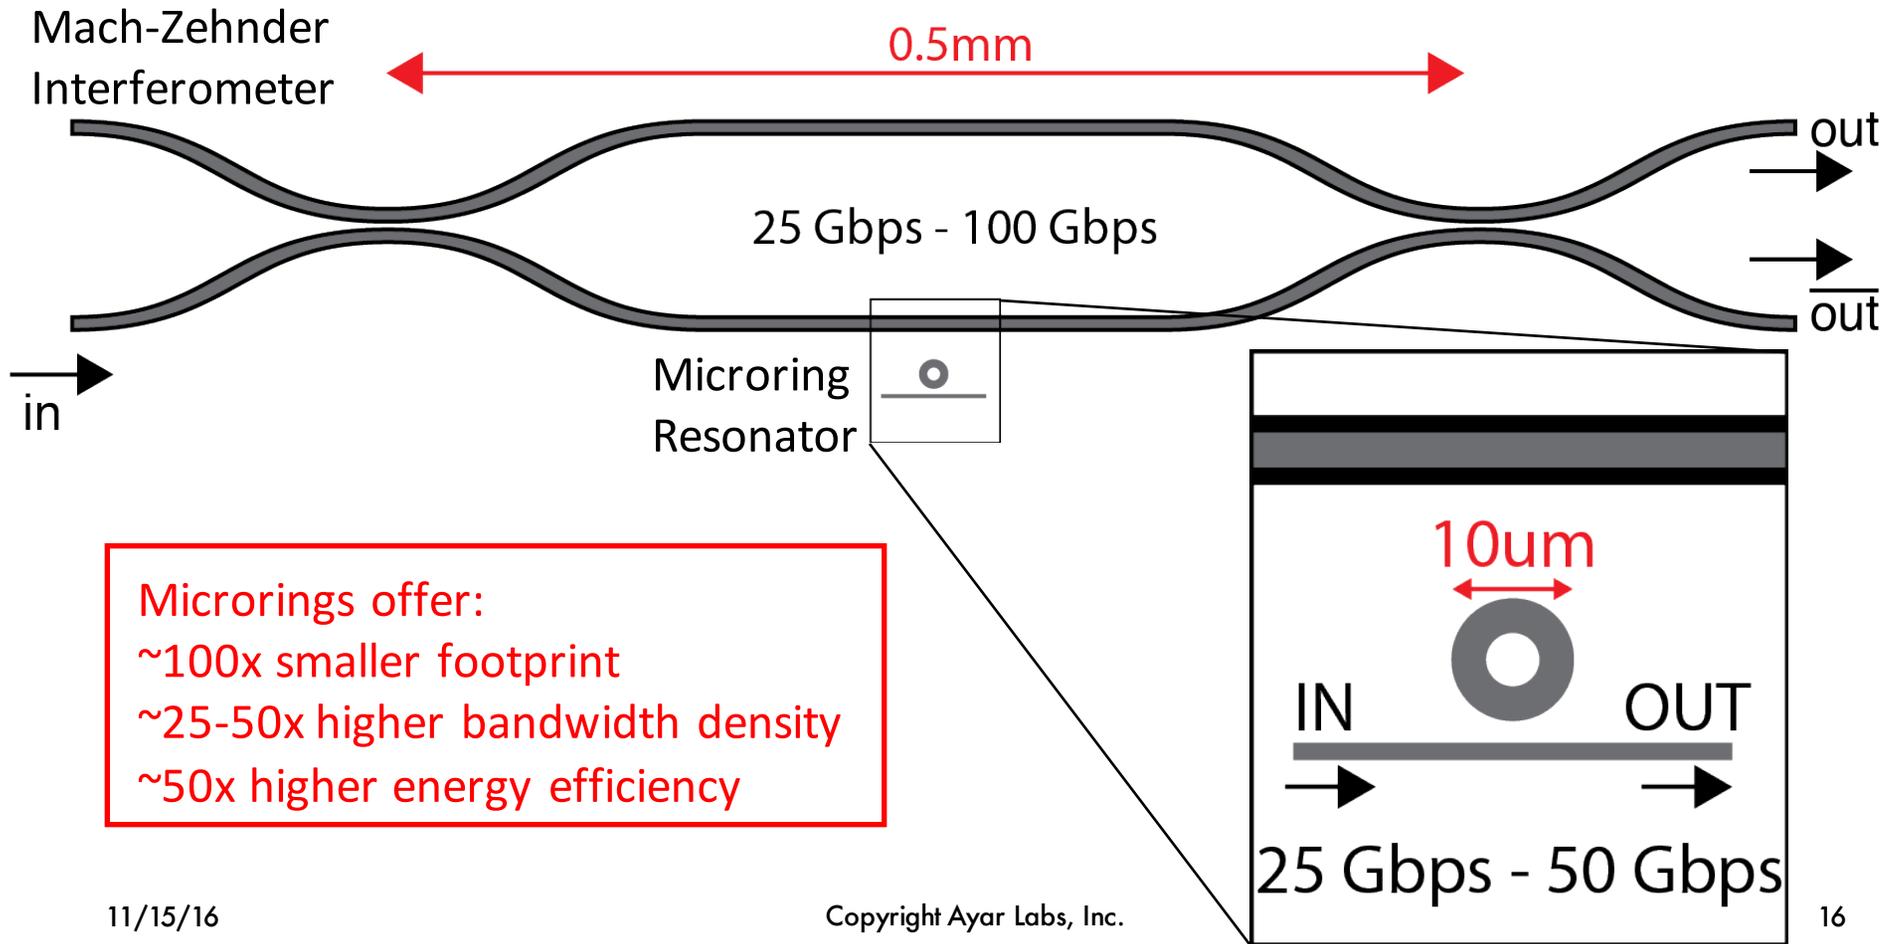


# The solution: microring resonators





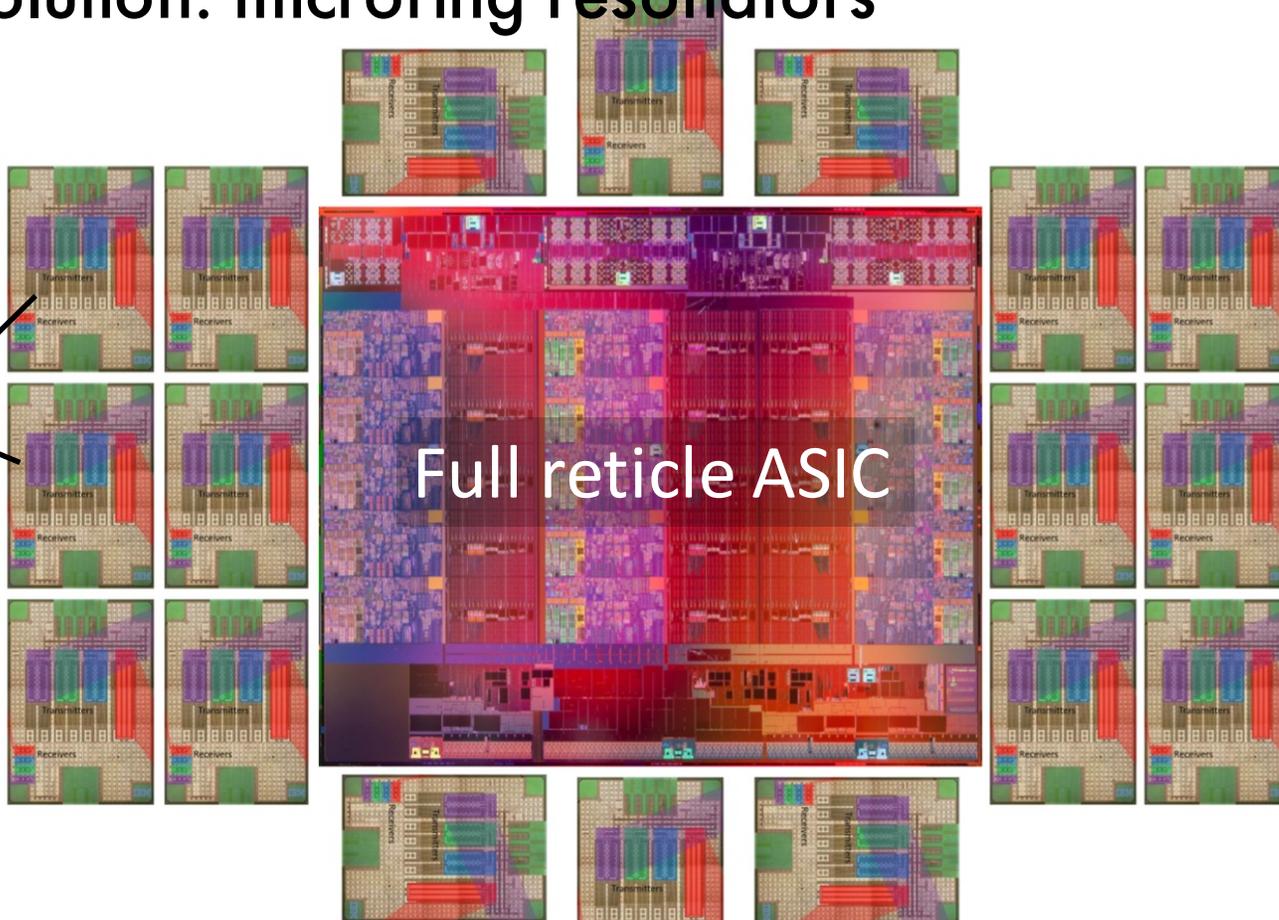
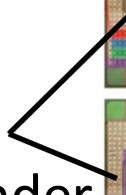
# The solution: microring resonators





# The solution: microring resonators

400 Gbps  
Mach-Zehnder  
transceivers



11/15/16

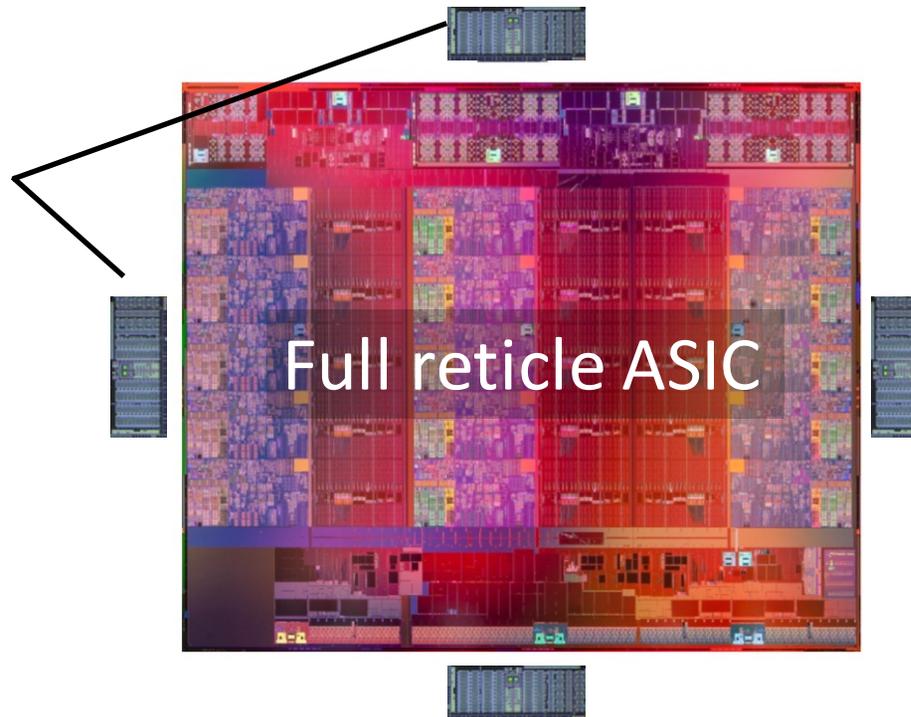
Copyright Ayar Labs, Inc.

100 Gbps transceiver, IBM, OFC 2015



# The solution: microring resonators

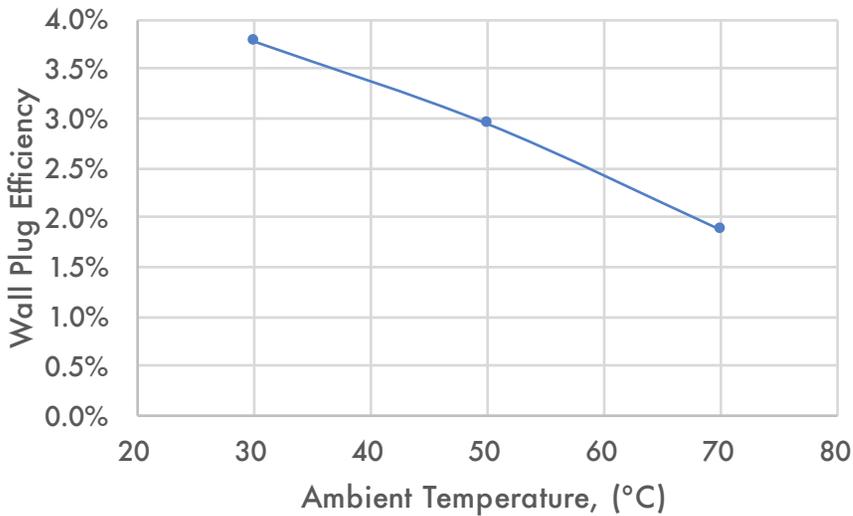
1.6 Tbps Ayar Labs  
microring-based  
transceivers



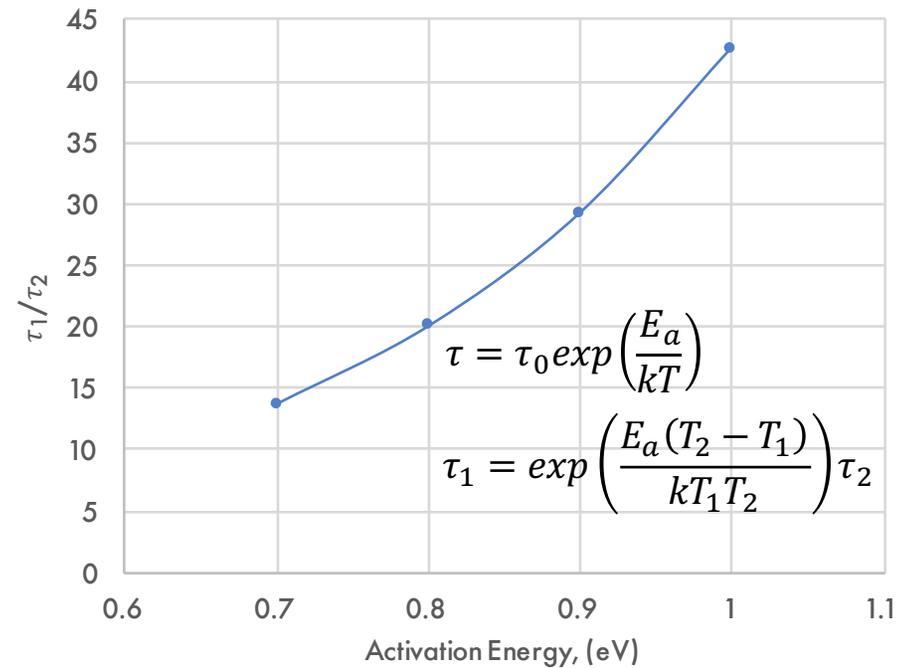


# The problem: lasers severely degrade in high temps

## Efficiency decreases



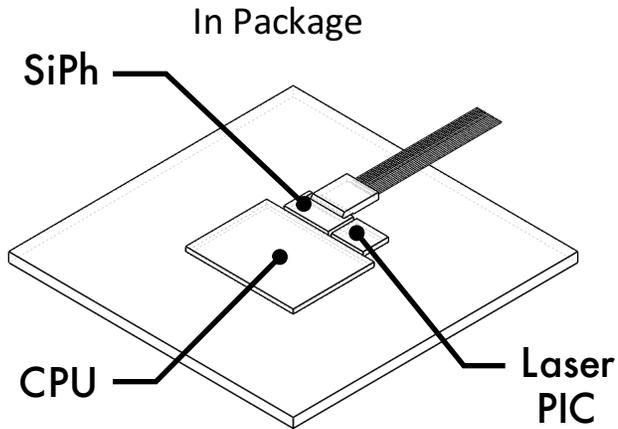
## Reliability decreases



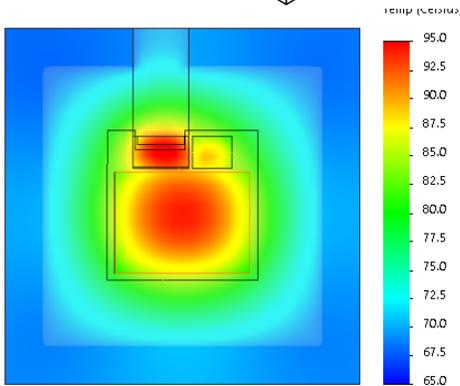
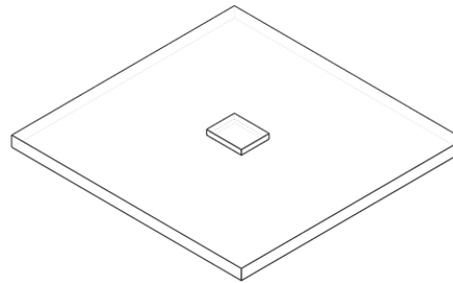
\* See JEDEC Publication No. 122F, P.77; Eyring behavior expected



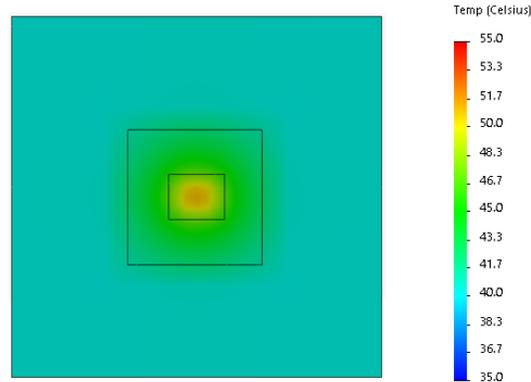
# The solution: separate the laser module



External module



50W TDP ASIC  
Max Tj ~ 89 C



$$\Delta T = 38^{\circ}\text{C}$$

~50% lower wall plug efficiency

~20x worse mean time to failure

Max Tj ~ 51 C

\* Assumptions and calculations in appendix



# What numbers are achievable with this architecture?

	Optical link energy use	Interconnect cost
Overall targets	<5 pJ/b	\$0.10/Gbps
Transceiver	Modulator Photodetector <2 pJ/b total Modulator driver circuitry Receiver circuitry Serializer/Deserializer	Modulator Photodetector Modulator driver circuitry Receiver circuitry Serializer/Deserializer
Laser	0.5 pJ/b	\$0.03/Gbps
Electrical SerDes with ultra-short reach links	1.0 pJ/b	
Total	3.5 pJ/b	~\$0.10/Gbps

Made in CMOS

Cost can scale to \$0.06/Gbps



# Where optics need to go

## Use microring resonators

- increases bandwidth density
- increases energy efficiency
- scaling vectors are aligned with efficient electronics

## Separate the laser module

- increases the reliability of the system
- increases the energy efficiency of the system
- simplifies the package



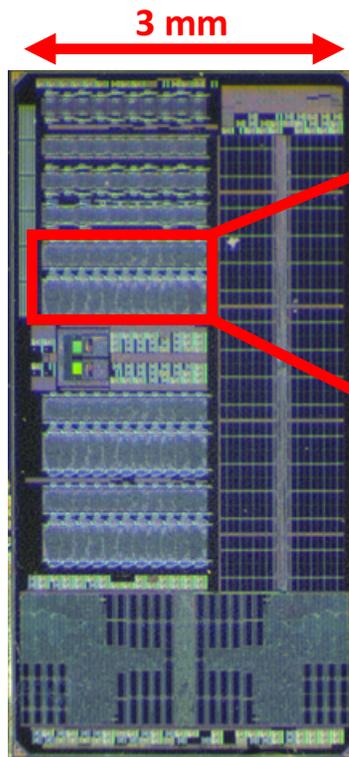
# Use existing high-volume CMOS processes to build efficient photonic I/O



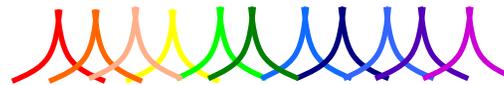
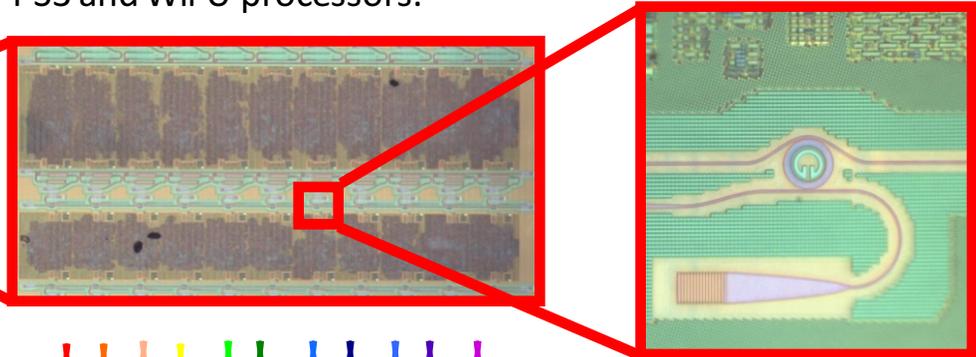
10x less energy use

40x higher bandwidth density

compared to copper I/O



Made in the foundry that builds IBM's flagship processors. Made the processors in 2 of the world's top 10 supercomputers, PS3 and Wii U processors.



Dense Wavelength Division Multiplexing (DWDM)

10 um

Our modulator (microring resonator)



# Team

- Founded 2015
- \$2.5M of equity funding led by  **FOUNDERS FUND**
- Several grants from DOE and NSF. Winners of \$275k MIT Clean Energy Prize

## Founding team



Chen Sun, CTO



Mark Wade,  
Chief Scientist



Alex Wright, CEO



## Recent hires

Roy Meade, VP of Manufacturing



- Formerly senior business line manager at Micron
- Two decades in semiconductor and optics



John Fini, Principal Engineer for Optical Devices

- Led next-gen receiver design at Lumentum (JDSU)
- Over a decade in optical device design

Evelina Yeung, High Speed Electrical Links Lead



- Led high speed links engineering team at Intel
- Two decades in top-tier semiconductor companies

## Founding Technology Principals



Prof. Rajeev Ram

- PI at MIT
- Former ARPA-E Program Director
- Founder of Pharyx

11/15/16



Prof. Vladimir Stojanović

- PI at UC Berkeley
- Founder of NanoSemi



Copyright Ayar Labs, Inc.



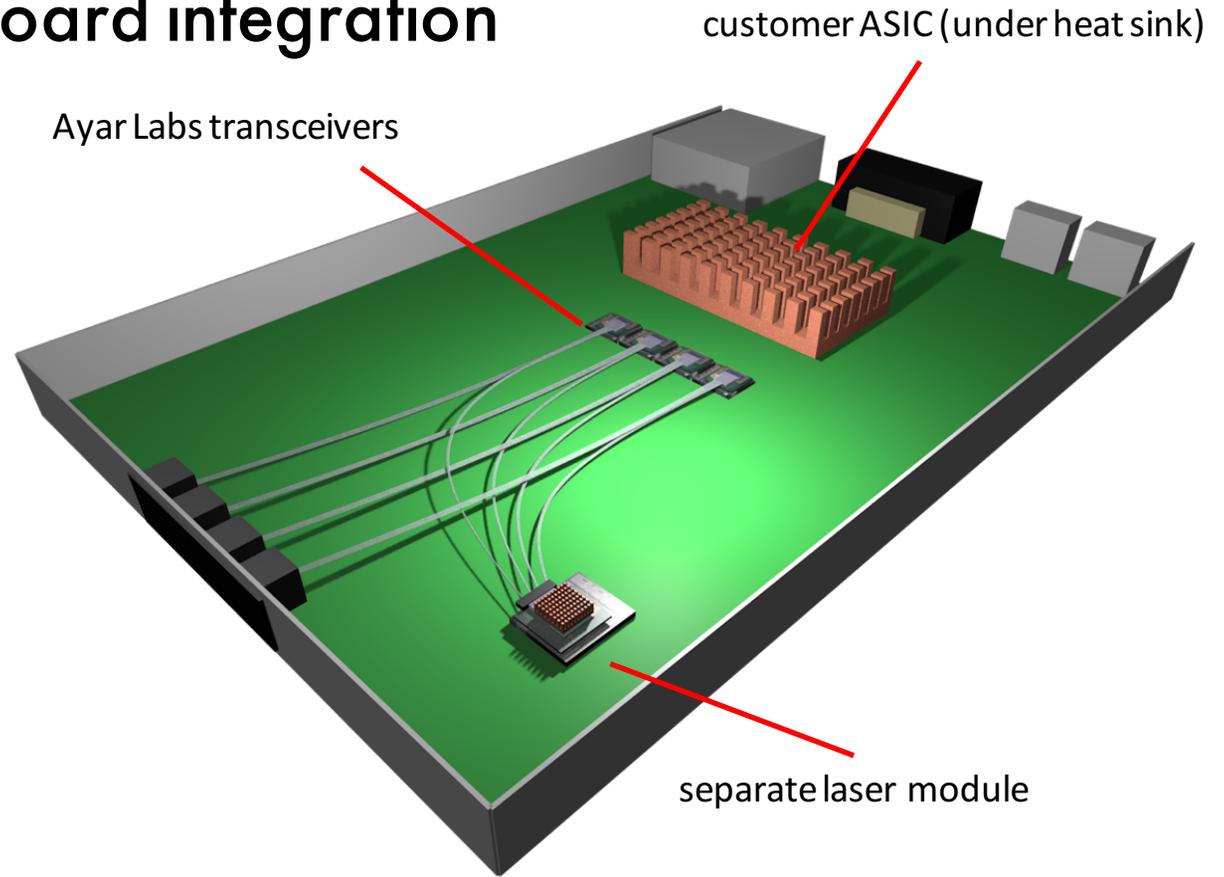
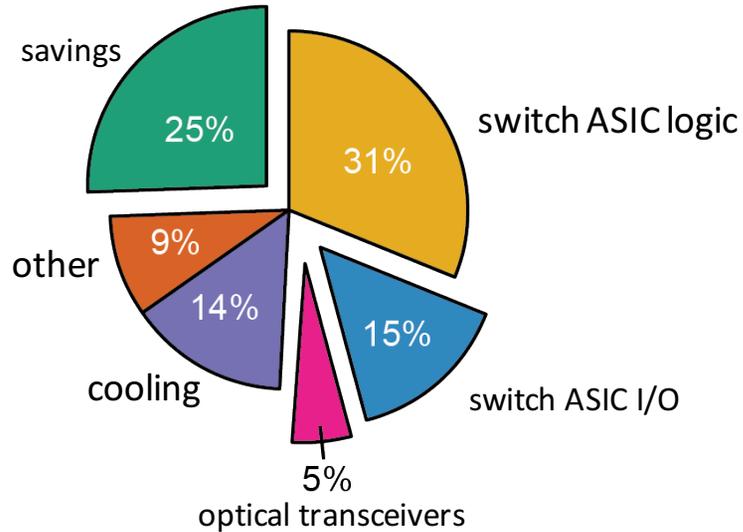
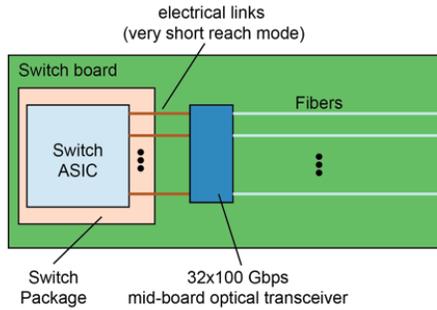
Prof. Miloš Popović

- PI at Boston University
- 100 publications
- Inventor on 16 patents

24



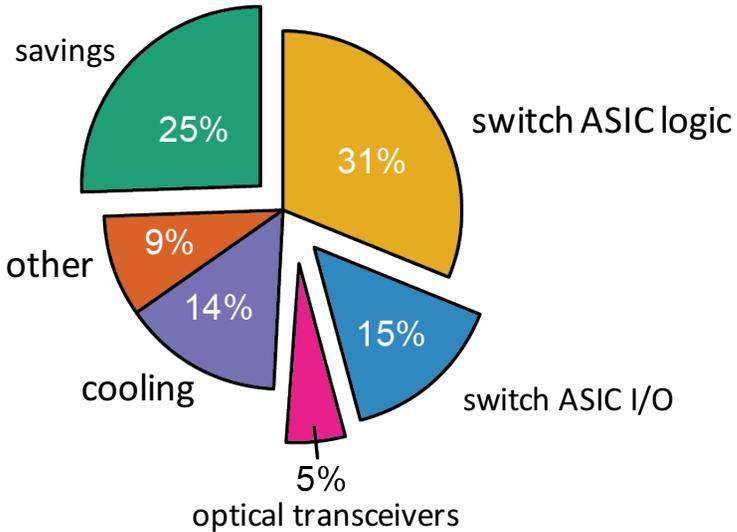
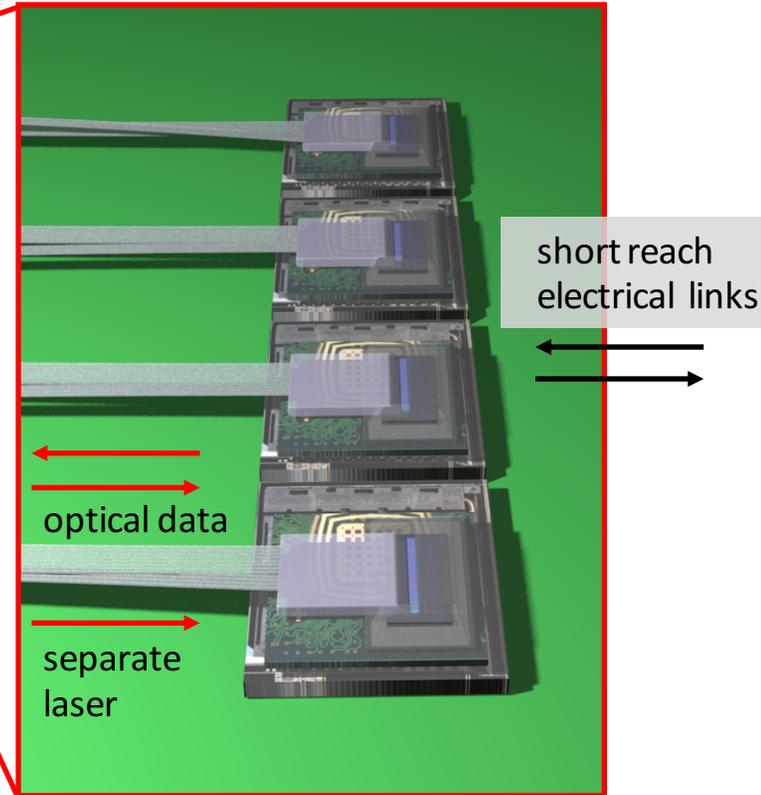
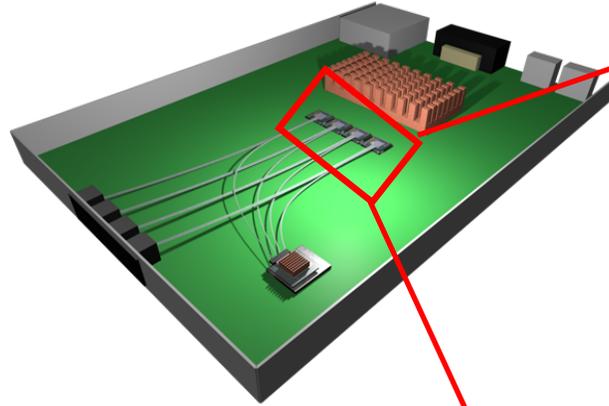
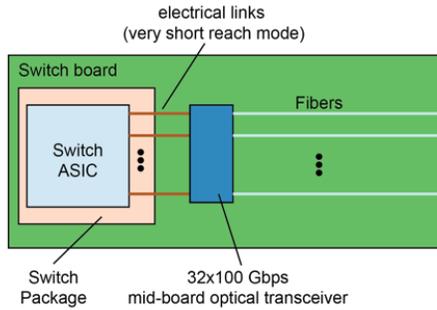
# First aim for mid-board integration





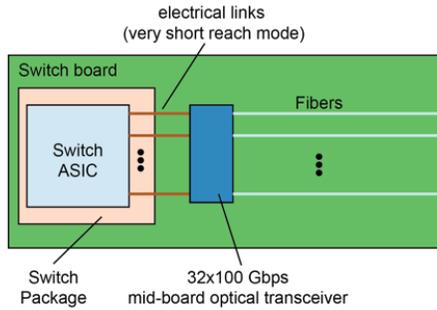
# First aim for mid-board integration

Example:  
Four 1.6 Tbps optical transceivers  
supporting 6.4 Tbps switch ASIC

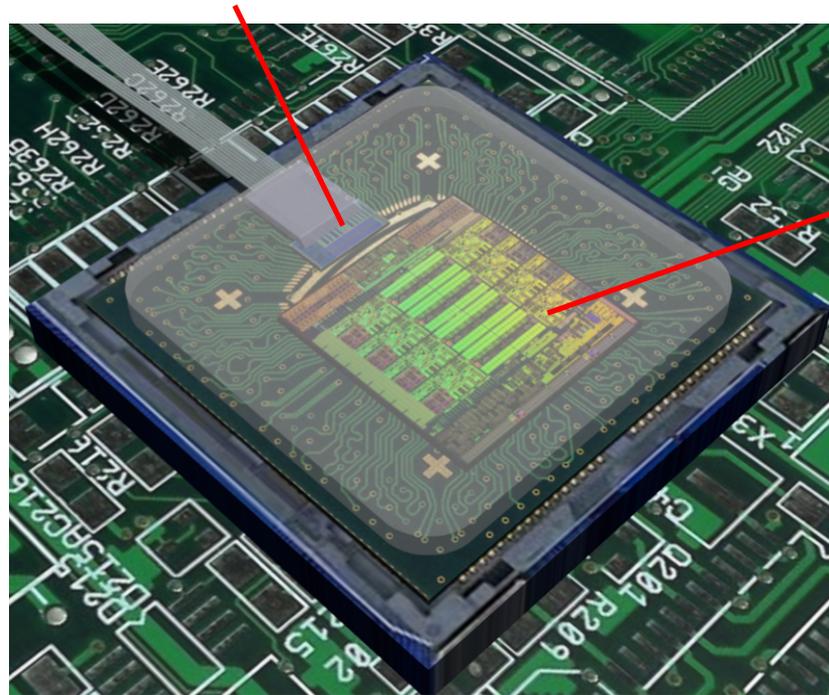




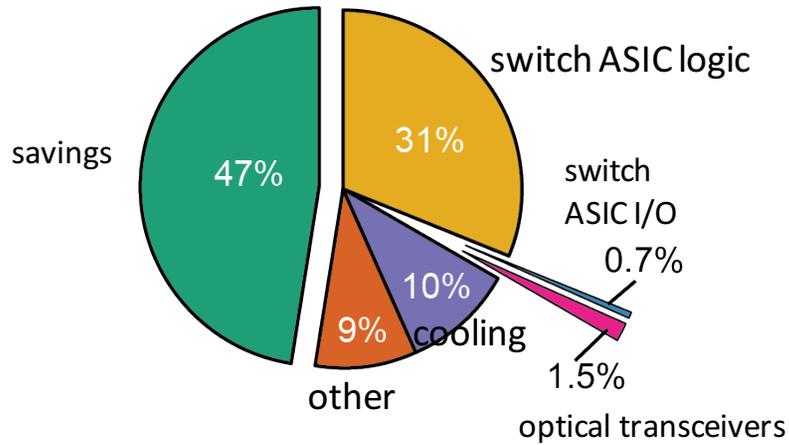
# Next transition to in-package integration



Ayar Labs transceiver chip

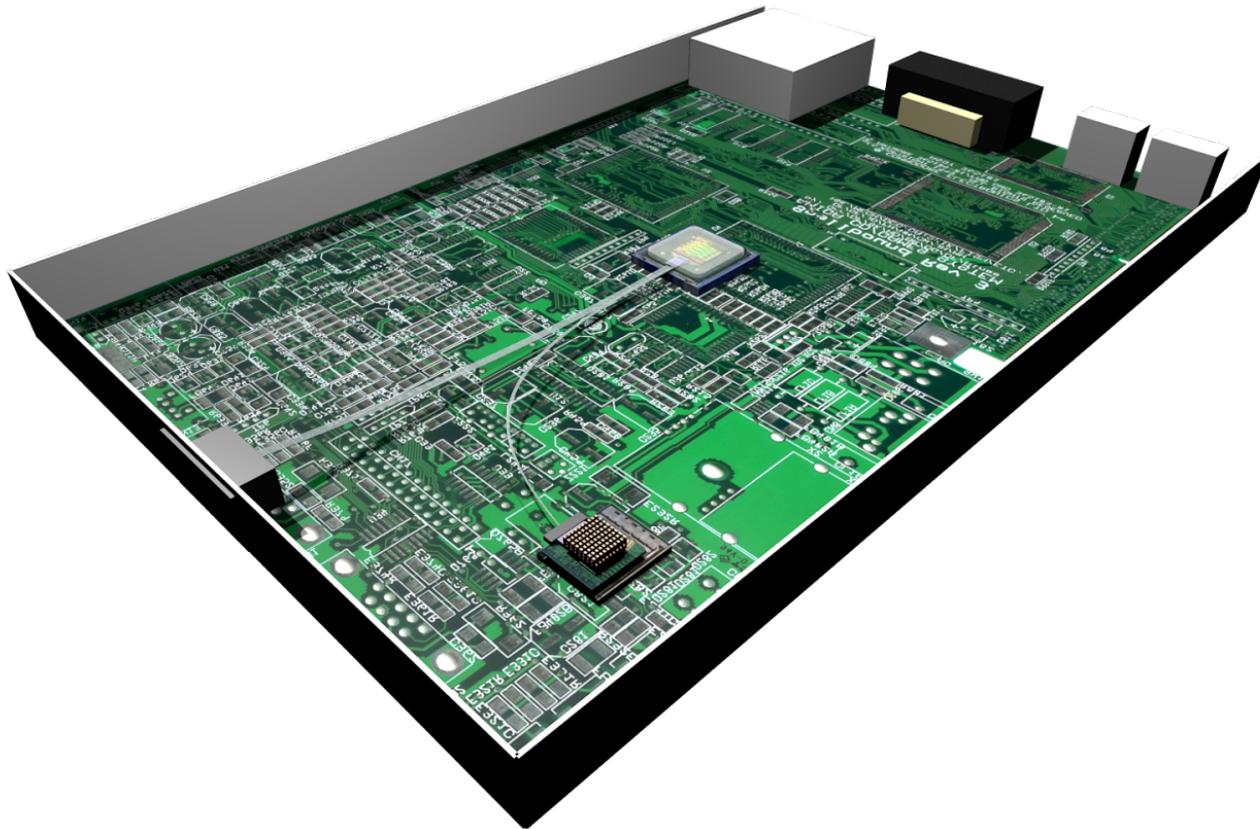


Full reticle ASIC





What would you do with  $>10\text{Tbps}$  at  $1\text{pJ/b}$  with  $2\text{km}$  reach?





**Thank you!**

---

[mark@ayarlabs.com](mailto:mark@ayarlabs.com)