Silicon Photonic Super-Switch

Ming C. Wu (PI), Vladimir Stojanovic, and Eli Yablonovitch, UC Berkeley
Paraskevas Bakopoulos, Elad Mentovich, Mellanox

ELECTRONICS RESURGENCE INITIATIVE SUMMIT & MTO Symposium 2020 | Seattle, WA | August 18-20

PIPES — ***Heterogeneous 3D***

Background

Ever-Demanding Computing Power for AI and Machine-Learning Workloads

- The computing power required in AI and machining-learning workloads has increased by 300,000 times in the past 7 years (doubling in 3.5 months)
- However, the hardware improvement as predicted by Moore's Law is only 12 times in the same period (doubling every 18 months)
- Fundamental changes in how we construct high performance computing clusters are needed to keep up with the demand
- Traditional electrically switched networks are facing fundamental limits in power consumption
- For example, commercial switch chips with 12.8 Tb/s throughput consumes 400W
- Optically switched network could enable much larger cluster HPC

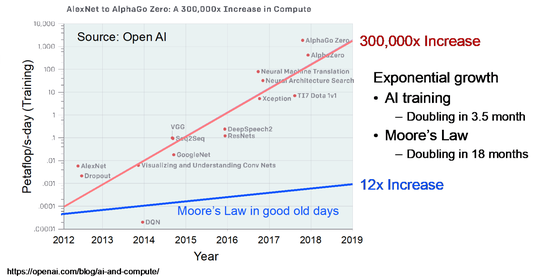

"Power Wall" in Interconnect Network of Large-Scale Compute Clusters

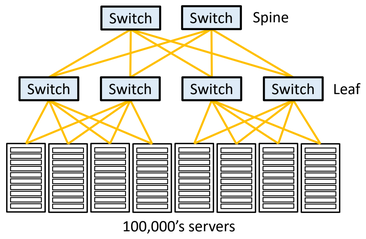

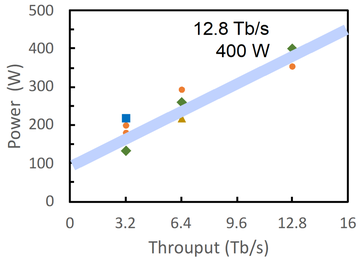

Scaling of AI Servers with Optically Switched Network

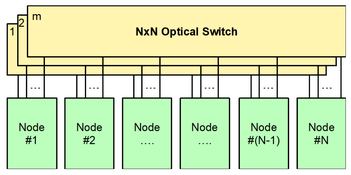

Advantages of optical switches:
- "Unlimited" bandwidth
- Future proof (transparent to data rates, formats)
- Low power consumption (100x lower than electrical switches)
- Low PHY latency – speed of light

Challenges:
- Currently, there is no large-scale optical switch with microsecond switching time
- Efficient network control and management need to be developed

Goal of this research is to develop optical switches with port counts approaching (exceeding) those of free-space switches but with 100,000x faster switching speed. Below is a comparison of switch technologies:

Commercial 3D MEMS-based Free-Space Switches

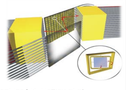

[Calient Inc.] R. Helkey *et al. Optics & Photonics News*, May 2002

Switch Type	3D MEMS[1]	Planar Lightwave Circuits (PLC)[2]	Si Photonic, Thermo-Optic[3]	Si Photonic MEMS[4]	Super-Switch (This Program)
Construction	Free Space	Integrated	Integrated	Integrated	Integrated
Size	~ 30x45x50 cm^3	Chip: 10x10 cm^2	Chip: 2.6x1 cm^2	Chip: 4x4 cm^2	Chip: 10x10 cm^2
Cost	High	Medium	Low	Low	Low
Port Count	320x320	32x32	32x32	240x240	1000x1000
Optical Loss	< 3dB	6dB	10dB	< 10dB	< 3dB
Switch Time	25 ms	> 1 ms	> 10μs	< 1μs	< 1μs
Crosstalk	< -65dB	~ -20dB	~ -20dB	< -60dB	< -60dB
Bandwidth	> 500nm	30nm	< 15nm	> 300nm	> 100nm
Power Consp	50W	10W	2W	< 1W	< 10W

Goal of this program: Fully Integrated Si Photonic Super-Switch

1. Calient S320 switch, http://www.calient.net
2. S. Sohma, T. Watanabe, N. Ooba, M. Itoh, T. Shibata, and H. Takahashi, "Silica-based PLC Type 32 x 32 Optical Matrix Switch," ECOC 2006
3. K. Suzuki et al., "Low-Insertion-Loss and Power-Efficient 32 × 32 Silicon Photonics Switch With Extremely High-Δ Silica PLC Connector," J. Lightwave Technol., JLT, vol. 37, no. 1, pp. 116–122, Jan. 2019.
4. T. J. Seok, K. Kwon, J. Henriksson, J. Luo, and M. C. Wu, "Wafer-scale silicon photonic switches beyond die size limit," Optica, vol. 6, no. 4, pp. 490–494, Apr. 2019

Approach

Wafer-Scale Integration Enabling Unprecedented High Port Count

- We have combined several novel ideas to demonstrate an unprecendented 1024x1024 switch fully integrated on a Si chip:
  - Drastically reduce the optical loss by using MEMS-actuated adiabatic couplers, which has zero loss in the OFF state
  - Modular design with crossbar architecture. The 1024x1024 switch comprises 8x8 blocks of 128x128 switch blocks
  - Each 128x128 block fits within a lithography reticle
  - Use Wafer-scale integration to lithographically stitch the switch blocks to form large-scale switches beyond reticle size limit
  - Only 3 different reticles are needed for wafer-scale integration:

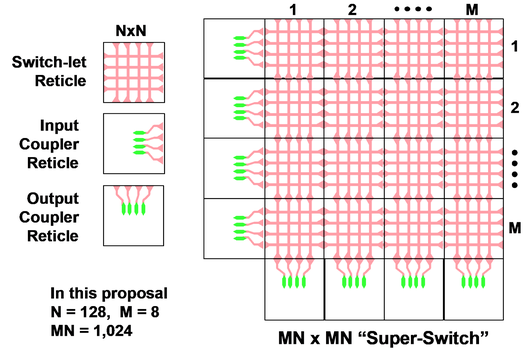

Tight CMOS Integration to Reduce Number of Electrical I/O

- CMOS control chips are directly integrated on top of the optical switches to reduce the number of electrical I/O (and hence the pin count)
- Modular CMOS design with 32x32 chiplets
- CMOS chiplets are tiled to control larger switch fabrics (e.g., 128x128, 1024x1024)
- High voltage CMOS driver capable of 70V output in 100 ns
- Fully digital addressing to select the active driver

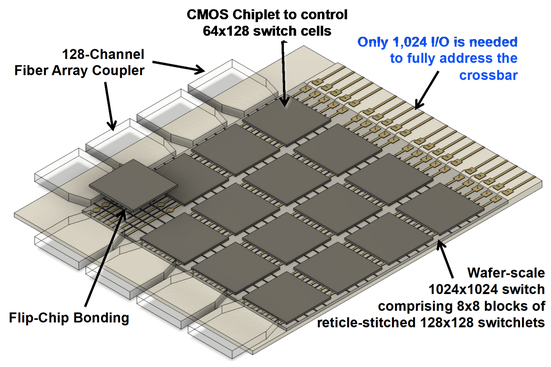

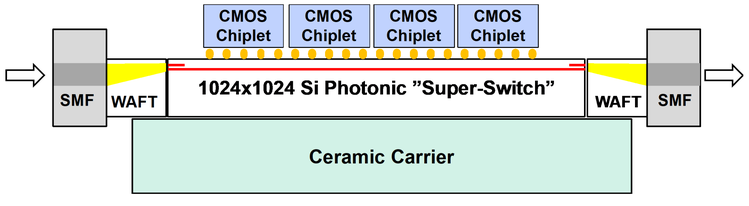

Results and Impact

Wafer-Scale Integrated 240x240 switches

Test Setup of 240x240 Switch

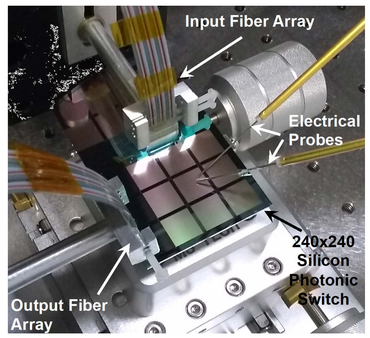

Transfer Curve

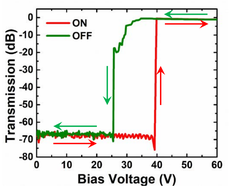

Switching Time

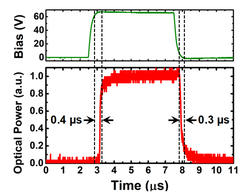

- World's largest integrated photonic switch
- Chip area: 4cm x 4 cm
- 3x3 blocks of 80x80 "switchlets"
- Digital switching with 70 dB extinction ratio
- 400 ns switching time

Phase 1 Design of Fully Integrated 128x128 Switch

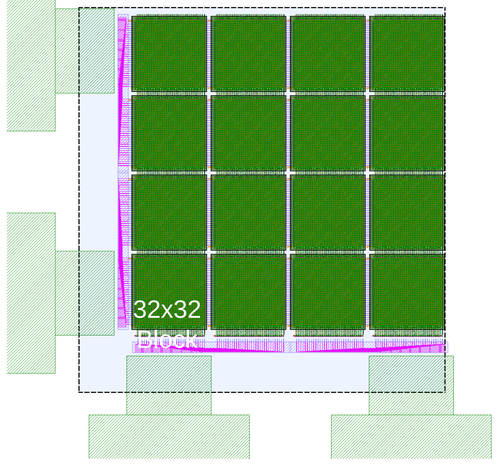

MEMS Switch Design and Simulation

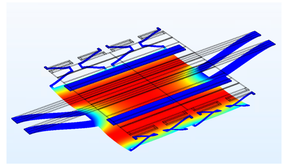

Switch Performance Simulation

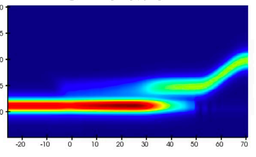

CMOS Design and Simulation
70V Swing in 100 ns

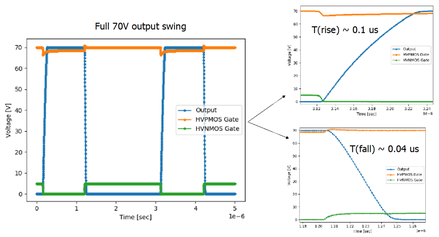

SiPh MEMS Switch Simulation
Switching time < 1 μs

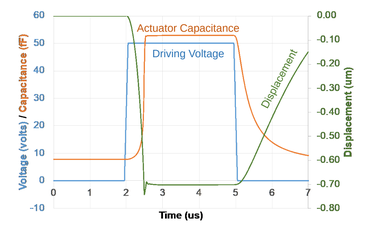

Impact

- Novel wafer-scale electronic-photonic integration enables silicon photonic switches with
  - Port counts approaching/exceeding free-space 3D MEMS switches but 100,000x faster switching time
  - Orders of magnitude lower power consumption than electrical switches
- Enable high bandwidth, low latency networks for AI servers and high-performance computers (HPC)

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).
The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.
Distribution Statement A – Approved for Public Release, Distribution Unlimited