

Maximizing The Impact of Emerging Photonic Switches At The System Level

**George Michelogiannakis⁺, Min Yeh Teh^{*},
Madeleine Glick^{*}, John Shalf⁺, Keren Bergman^{*}**

+LBNL

*Columbia University





- ★ Photonics West presentation in 2018 on simulation photonic networks at the architectural level

How many Js
per bit are
enough?

I'd rather we
focus on other
metrics

- ★ Provide a discussion of:
 - ▣ Emerging trends in HPC and datacenter systems
 - ▣ How key photonic switch properties motivate different network architectures
- ★ Using photonics as a drop-in replacement in existing networks does not promise a substantial system-wide impact
- ★ Highlight the value of co-design and multi-disciplinary collaboration



Pries, R., Jarschel, M., Schlosser, D., Klopff, M., and Tran-Gia, P., "Power consumption analysis of datacenter architectures," in [Green Communications and Networking]

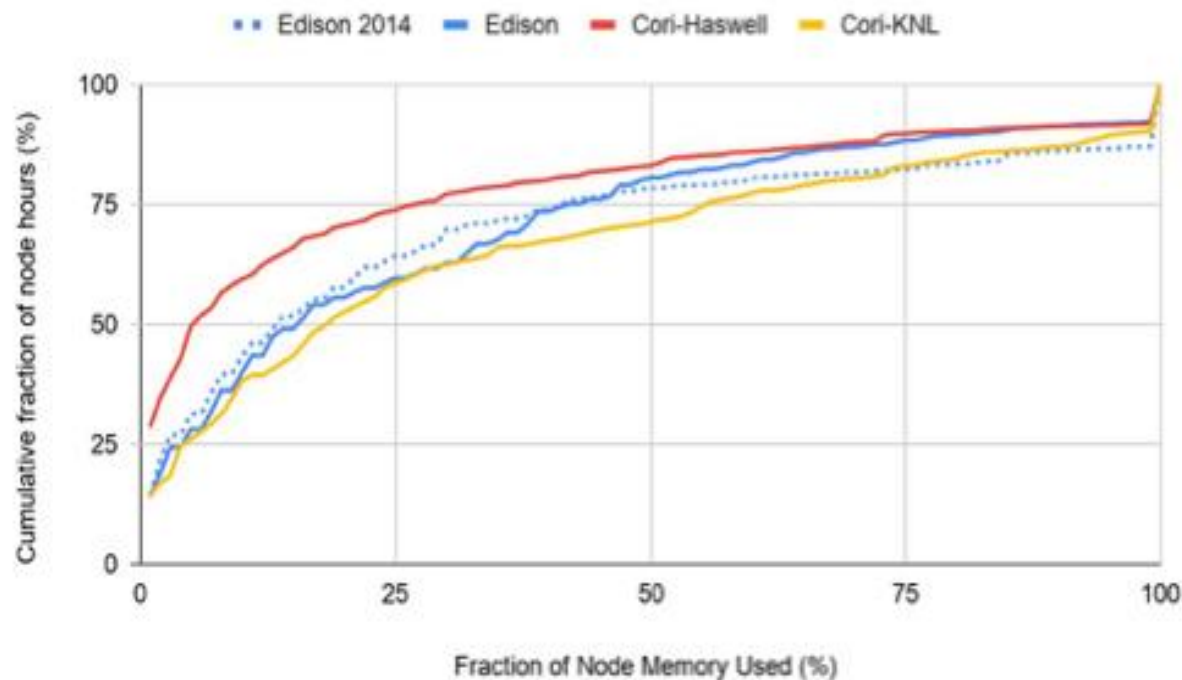
- * Resource Disaggregation
 - * Larger scales
 - * Non-volatile memory
 - * 3D memory
 - * Compute specialization
 - * Vector computing
 - * Collective operations
 - * Application data locality management
- * Adapting photonic switches to serve these trends bear significant promise

Motivation of Resource Disaggregation




- * 75% of node hours use 25% or less of node memory

Memory pressure at NERSC, 2018




Resource Disaggregation




Current server




Disaggregated rack



Current rack



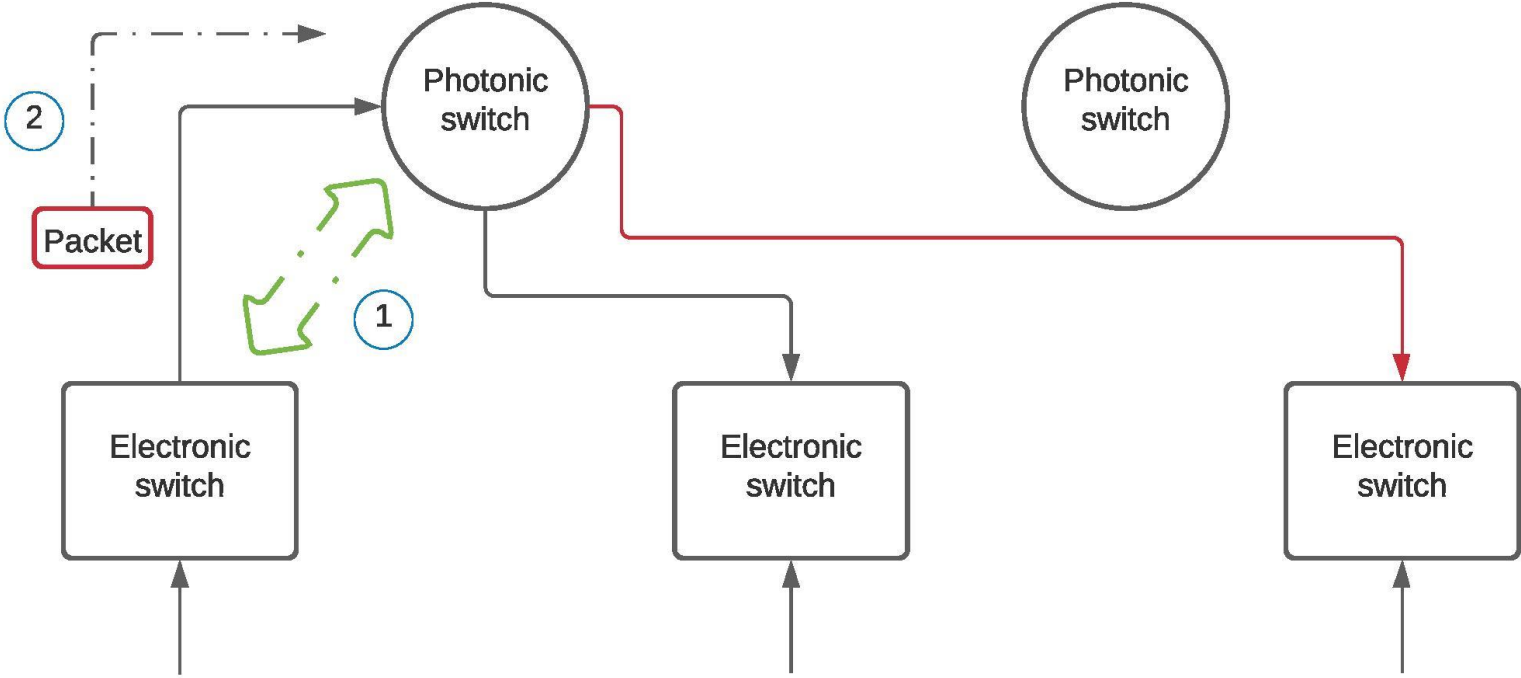
Pool and compose



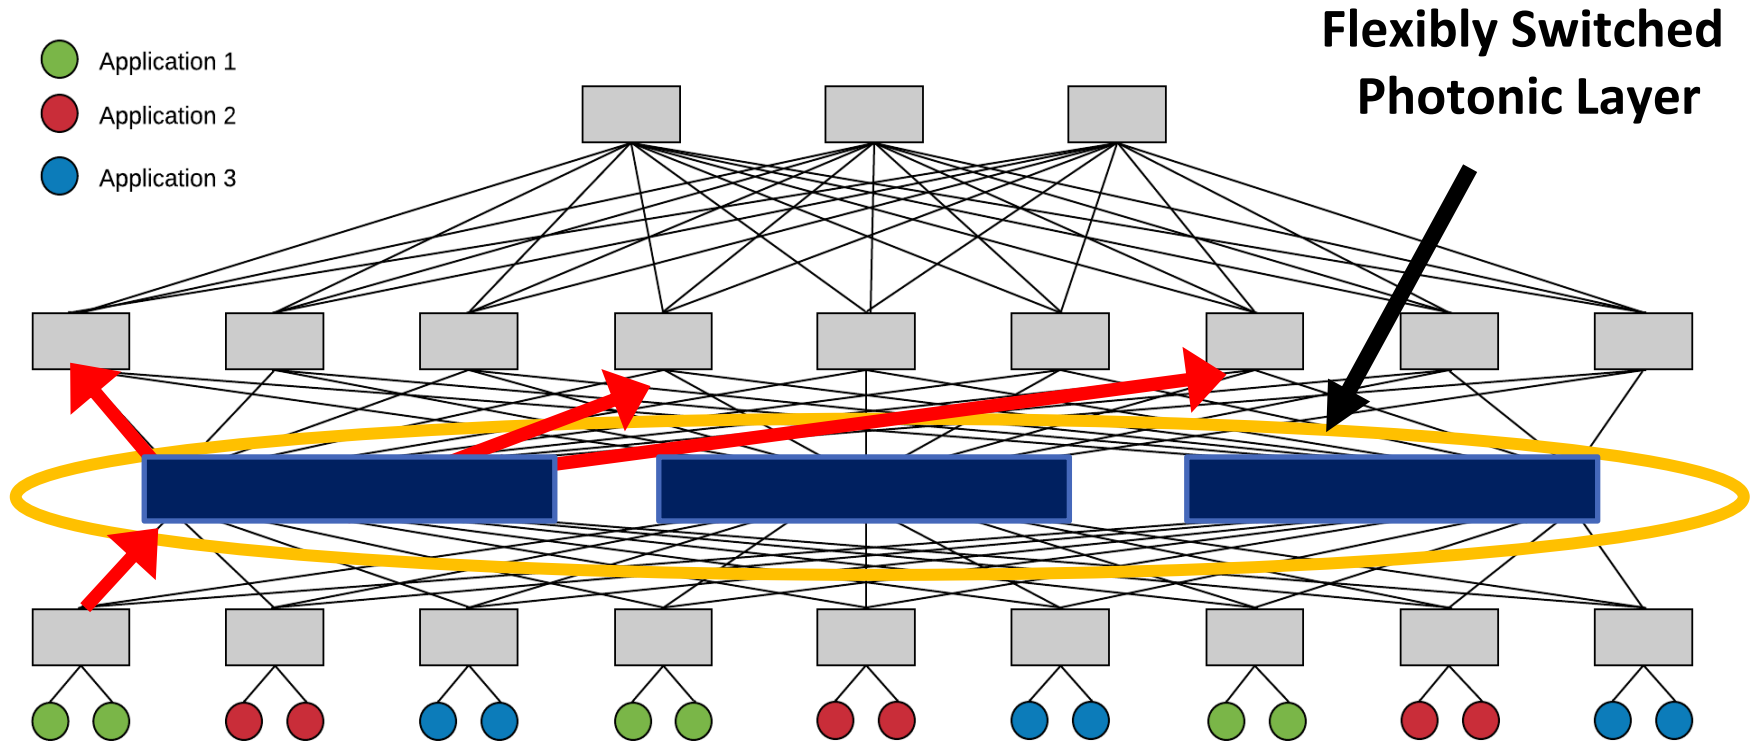
Reconfiguration Latency



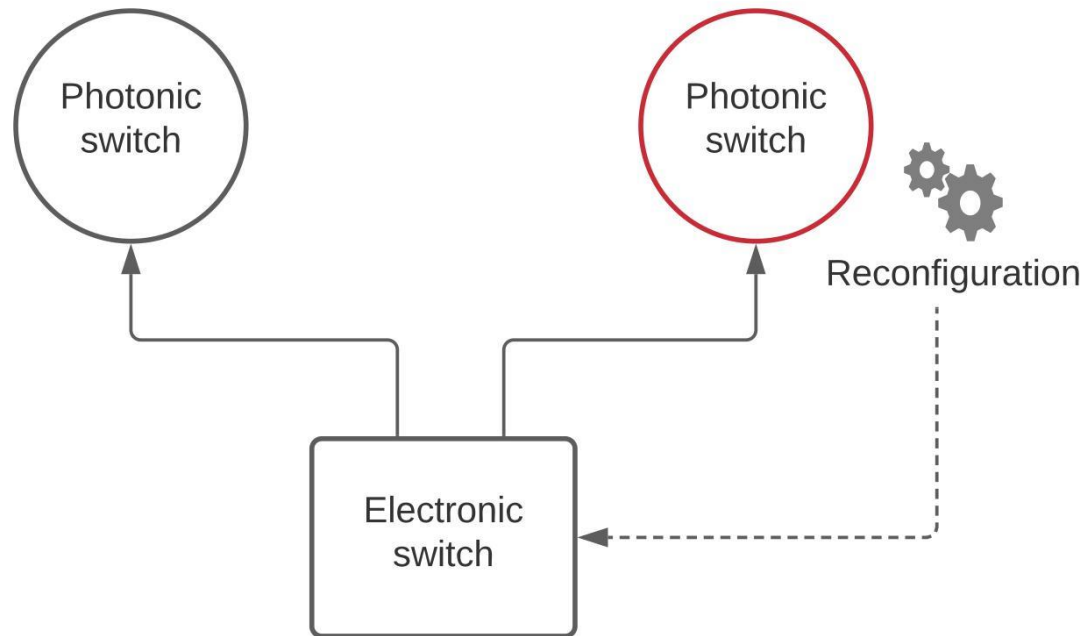
★ Nanosecond-scale reconfiguration latencies motivate per-packet reconfiguration



Slower Reconfiguration Motivates Different Network Architecture



Synergistic Advancements



- * Electronic packet switch routing tables
- * Control plane
- * Algorithms

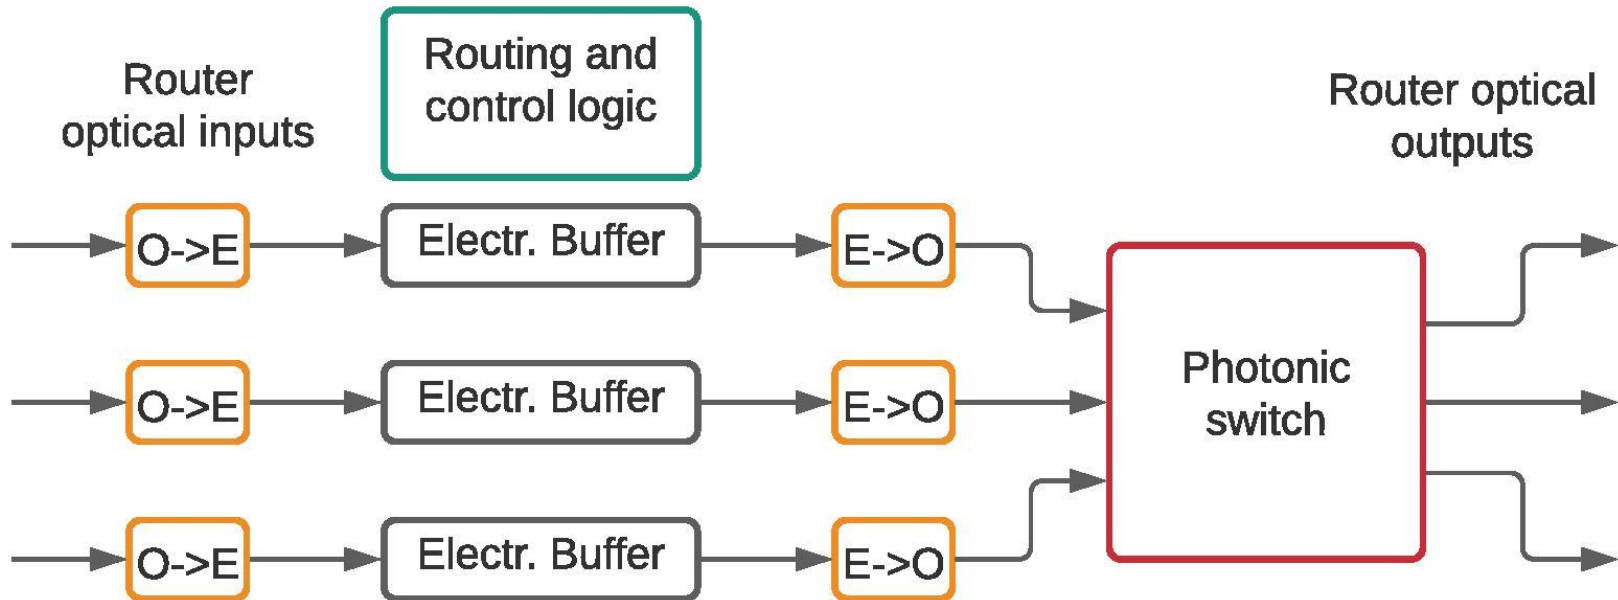
- * *Reconfiguration serves multiple trends*
- * *But it also motivates a variety of network architectures*

- ★ Currently more than one photonic switch hop is challenging due to signal loss
 - ▣ Usually have to convert back to the electrical domain
- ★ Low signal loss enables all-photonic topologies or sub-networks with multiple consecutive photonic hops
- ★ In turns, this increases the radix of efficient all-photonic paths, reducing the importance of resource placement across the system

Efficient OEO Conversions



- * OEO conversion cost affects the distance for which photonics are prudent



- ★ Optical computing
 - ▣ All-optical routers?

- ★ Multicast capability
 - ▣ Co-designed with other network components

Conclusion



- ★ Photonic switches are not best used as drop-in replacements
- ★ Instead, let's adapt photonic switches to support existing system trends
- ★ Also, understand the architectural opportunities from each photonic switch property
- ★ Co-design