Zero Buffer Optical Packet Switching Data Center Network

Shawn Shuoshuo Chen°, Weiyang Wang▪, Manya Ghobadi▪, Peter Steenkiste°, Srinivasan Seshan°

° Carnegie Mellon University School of Computer Science

▪ MIT CSAIL

Benefits, challenges & ideas

- Compared to electrical packet switching, optical networking technology promises:
  - higher bandwidth at lower costs
  - runtime reconfigurability
  - lower energy consumption per bit
- An all-optical data center network needs to address three challenges:
  - lack of data buffering in switches
  - lack of stateful routing in switches (cannot modify packet headers)
  - lack of forwarding at packet granularity

We leverage three key ideas to address the challenges accordingly.

All-optical challenges	*Proposal*
No data buffering in switches.	Schedule end host transmission globally.
No stateful routing in switches.	Source routing.
No forwarding at packet granularity.	Make cross-connects with Optical Circuit Switch (**OCS**) based on packet headers.

All-optical packet network

The goal of this project is to **build a scheduled source-routed network with all-optical packet switching**. Our design keeps data transmission *entirely in the optical domain*.

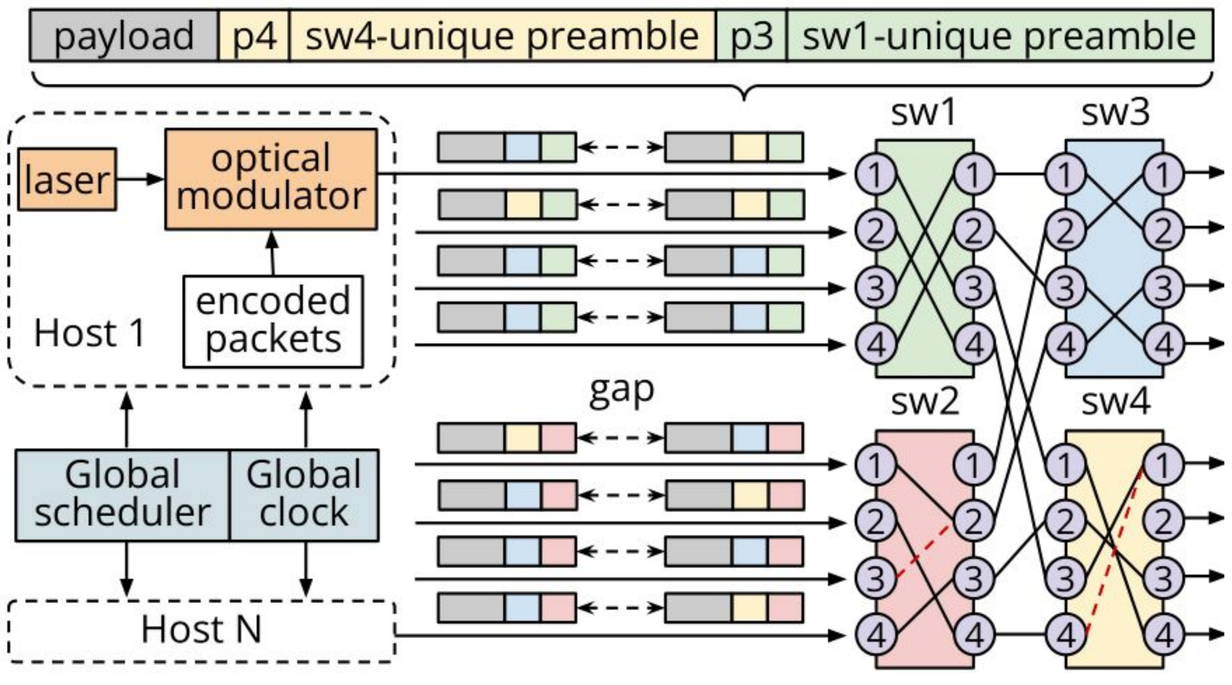

All-optical packet switch

Each switch (built on OCS) in our design:

- matches packet headers in analog optical form
- does not buffer or modify any packet

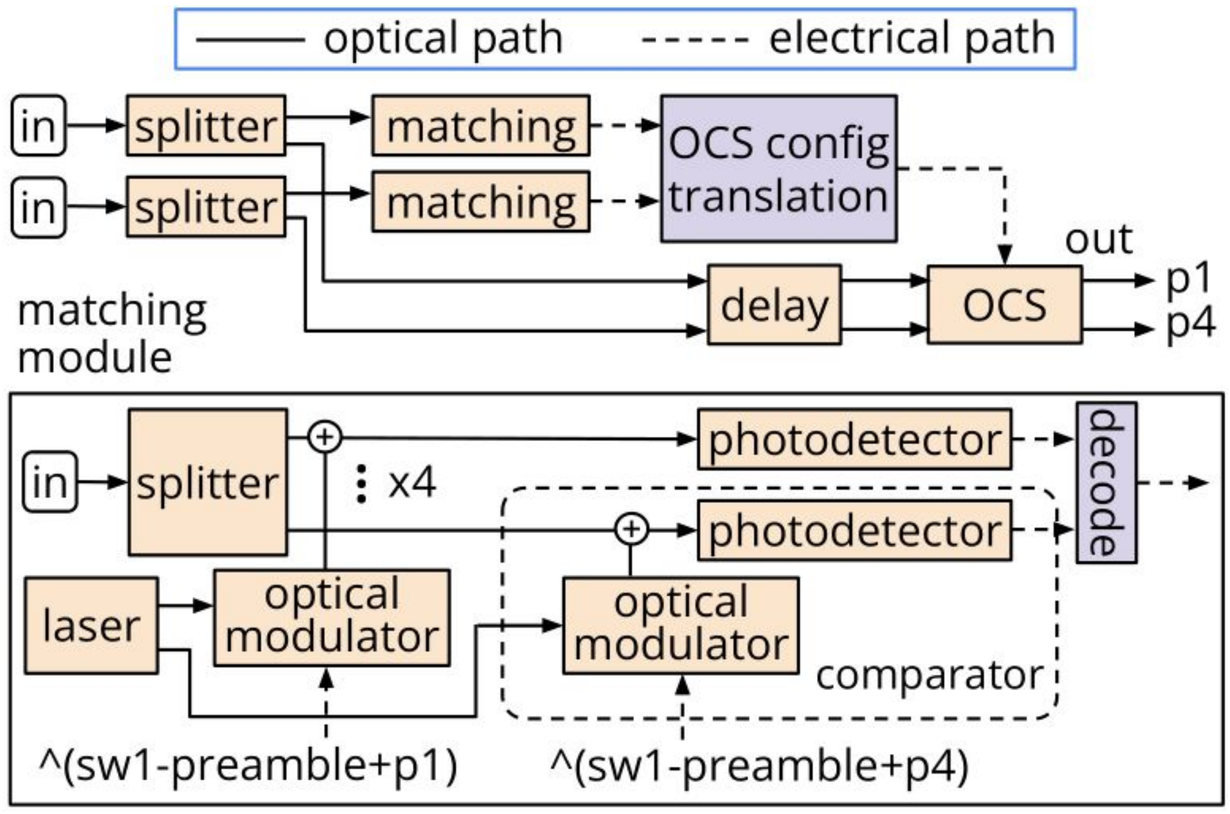

Preliminary Results

Fast OCS reconfiguration time reduces per-port throughput overhead.

- At 100ns reconfig. time, the max throughput on a 400Gbps link is 257Gbps (with 9KB packets).
- At 10ns, throughput increases to 378Gbps.

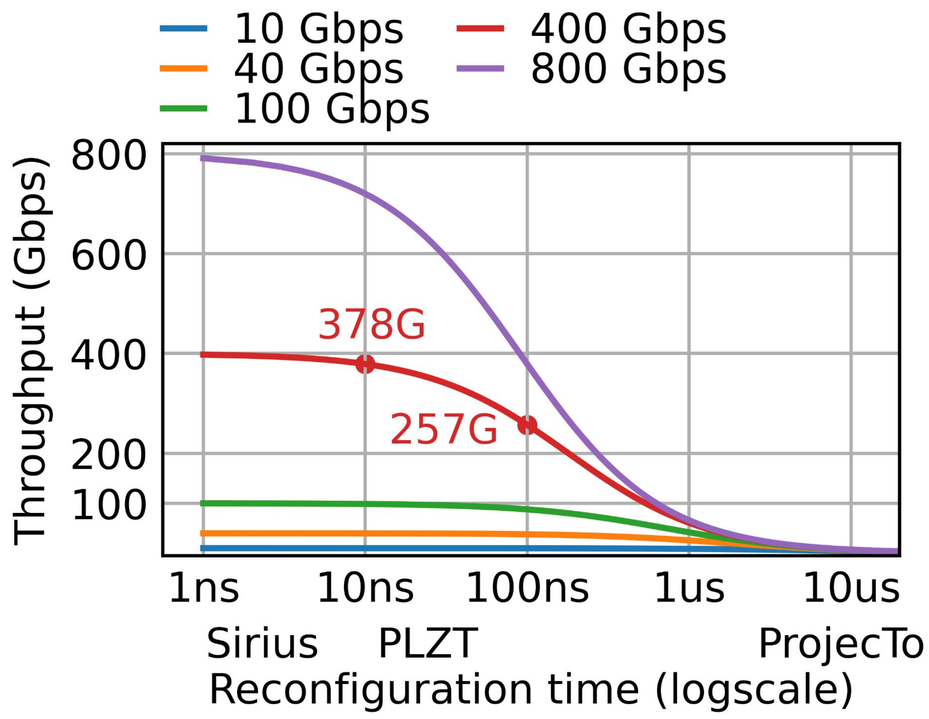

Future work

- Sub-usec clock synchronization of input port signals between hosts & switches
- Fastpass-style centralized scheduling
  - informs end hosts about when to send, how much to send
- Fast reconfigurable OCSes