

⚡ OpenOptics ⚡

MAX PLANCK INSTITUTE  
FOR INFORMATICS



# OpenOptics: Enabling Open Research and Implementation of Optical Data Center Networks

Yiming Lei<sup>1</sup>, **Federico De Marchi**<sup>1</sup>, Jialong Li<sup>2</sup>, Raj Joshi<sup>3</sup>, Shu-Ting Wang<sup>4</sup>, Xiaoqi Chen<sup>5</sup>,  
Balakrishnan Chandrasekaran<sup>6</sup>, Yiting Xia<sup>1</sup>

Max Planck Institute for Informatics<sup>1</sup>

Shenzhen University of Advanced Technology<sup>2</sup>, Red Hat / Harvard University<sup>3</sup>,

UC San Diego<sup>4</sup>, Purdue University<sup>5</sup>, Vrije Universiteit Amsterdam<sup>6</sup>

# Ever Growing Cloud Traffic



 PyTorch

**Model Training**



**Gemini**



**Claude**



**ChatGPT**

**AI Workloads**



**Cluster Management**



**Streaming**



**Hugging Face**




**GitHub**




**Community Platforms**

# Moore's Law for Electrical Switches

- Electrical switches double their bandwidth every two years at the same power and cost.




# Optical Data Center Networks




- Optical circuit switches
  - ▶ Dedicated links per configuration
  - ▶ An accurate transmission “schedule”
- Bufferless
  - ▶ No queuing
  - ▶ Need to buffer packets elsewhere

# Optical Data Center Networks


- **Power efficient**
  - ▶ Passive optical switches
- **Low cost**
  - ▶ Lower per-port cost than electrical switches
- **Future proof**
  - ▶ High bandwidth limit on optical switches
- **Flexible network topology**
  - ▶ Bring data and computation closer




c-Through  
[SIGCOMM 2010]




Helios  
[SIGCOMM 2010]




OSA  
[NSDI 2012]




Mordia  
[SIGCOMM 2013]




RotorNet  
[SIGCOMM 2017]




Sirius  
[SIGCOMM 2020]



Jupiter  
[SIGCOMM 2022]



Shale  
[SIGCOMM 2024]



Realizing RotorNet  
[SIGCOMM 2024]

# High Barrier of Entry

- **Closed ecosystem**

- ▶ Specialized optical hardware + customized software stack
- ▶ Home-grown simulation + small-scale testbed


- **Siloed research and innovation**

- ▶ Software systems tied to the underlying hardware
- ▶ Lack of a unified platform for implementation and evaluation

# High Barrier of Entry

- Long Development Cycle

## RotorNet



[JLT'16] A scalable, partially configurable optical switch for data center networks

[SIGCOMM'17] RotorNet: A Scalable, Low-complexity, Optical Datacenter Network



[NSDI'20] Expanding across time to deliver bandwidth efficiency and low latency

[SIGCOMM'24] Realizing RotorNet: Toward Practical Microsecond Scale Optical Networking


# High Barrier of Entry

- Architecture-Specific Design

---

	Jupiter	Sirius
Company		 Microsoft
OCS Type	MEMS	AWGR
Switch Interval	Minutes	Nanoseconds
Scheduling	Traffic-aware	Traffic-oblivious

---




⚡ *OpenOptics* ⚡

Decouple Software from Hardware → Evolve Independently

# What is in OpenOptics

- ▶ Optical DCN architecture plug-and-play
- ▶ A simple abstraction + user API
- ▶ Multi-backend support: Testbed, Simulation, Emulation
- ▶ To enable easy design and testing



# Who is OpenOptics for

- **With physical hardware**
  - ▶ Shorten development cycle with built-in software
  - ▶ Software-stack evaluation with emulated optical hardware
- **Without physical hardware**
  - ▶ Single-laptop emulation or simulation
  - ▶ For initial research and exploration, for classroom education

# How to Unite Different Architectures?



DCN Applications

Transport Protocol


Routing

Packet Scheduling

Optical Architectures

# Routing in Optical DCNs

*Time Slices*




*Endpoint Node*

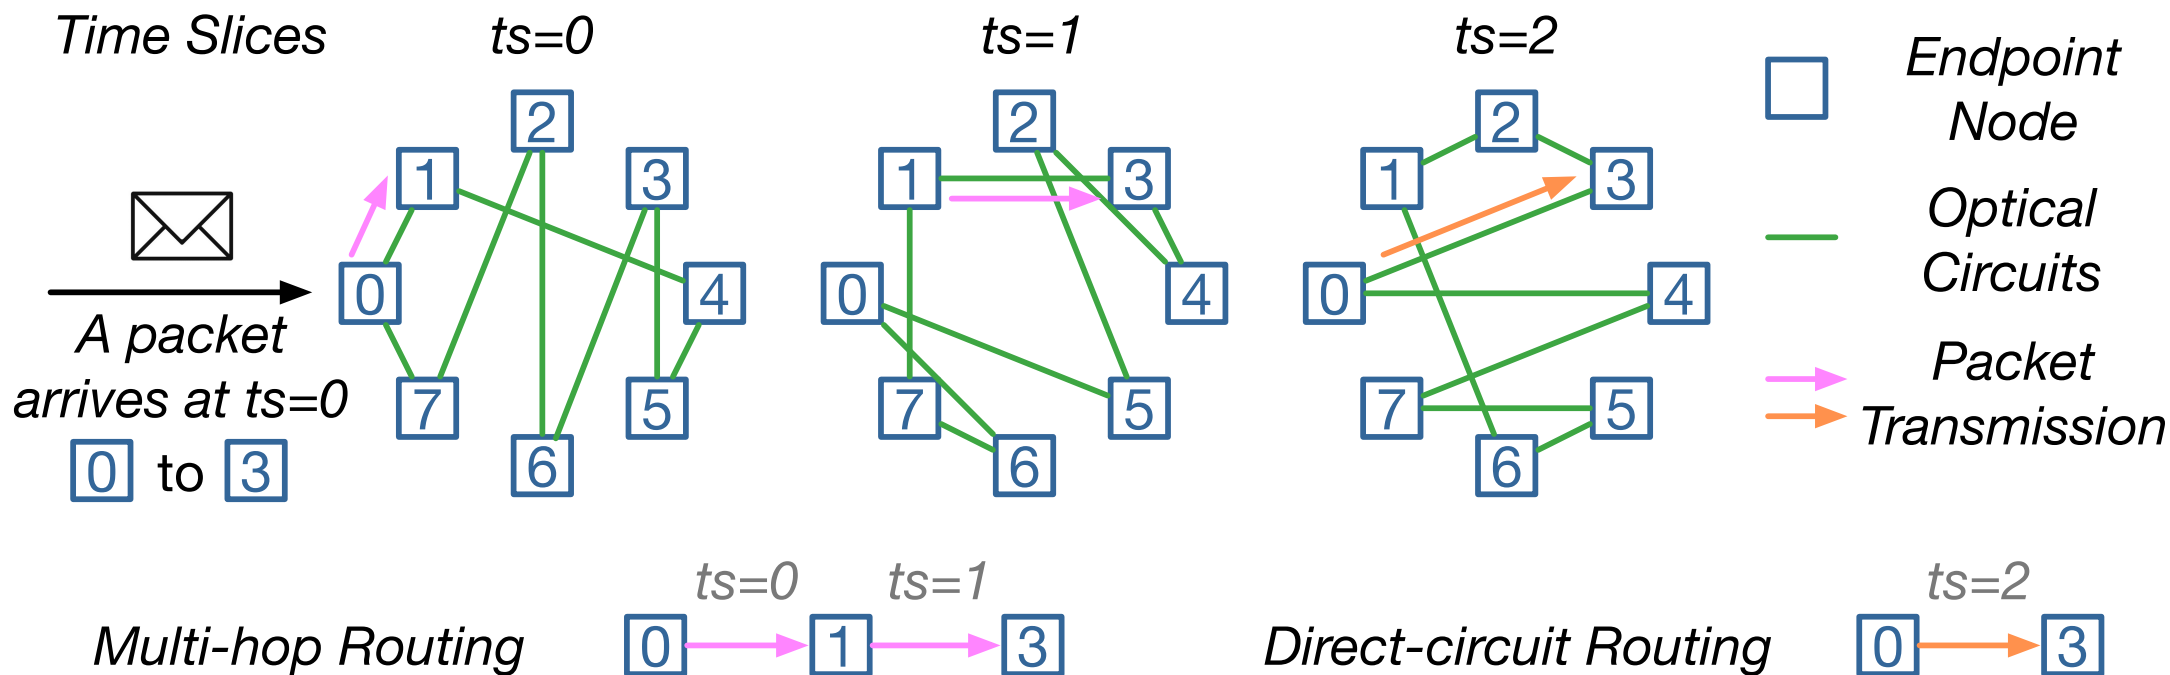


*Optical Circuits*

# Routing in Optical DCNs



# Routing in Optical DCNs



Routing is different across time slices


Packets may wait until a certain time to continue

# Time-Flow Table


*Node 0*

<i>Arrival Slice</i>	<i>Dst.</i>	<i>Egress Port</i>	<i>Departure Slice</i>
0	3	2	2

**Direct path**



# Time-Flow Table



Node 0

Arrival Slice	Dst.	Egress Port	Departure Slice
0	3	1	0

Node 1

Arrival Slice	Dst.	Egress Port	Departure Slice
0	3	2	1

Multi-hop path (per-hop lookup)

# Time-Flow Table

*Node 0*

<i>Arrival Slice</i>	<i>Dst.</i>	<i>Egress Port</i>	<i>Departure Slice</i>
<i>*</i>	<i>3</i>	<i>2</i>	<i>*</i>

Traditional flow table

*Node 0*

<i>Arrival Slice</i>	<i>Dst.</i>	<i>Source Routing Action</i> <i>Hops: &lt; Egress Port, Departure Slice &gt;</i>
<i>0</i>	<i>3</i>	<i>&lt;1,0&gt;,&lt;2,1&gt;</i>

Multi-hop path (source routing)

# Under Time-Flow Table: Time-Based Packet Scheduling

*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0
1	5	3	2







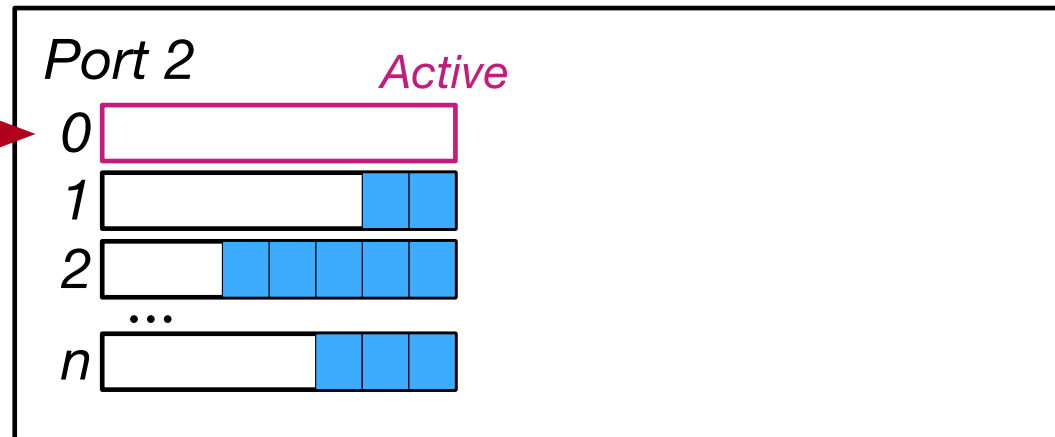
# Under Time-Flow Table: Time-Based Packet Scheduling



Time slice = 0

*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0




# Under Time-Flow Table: Time-Based Packet Scheduling

 Time slice = 1

*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0
1	5	3	2




# Under Time-Flow Table: Time-Based Packet Scheduling



Time slice = 1

*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0
1	5	3	2



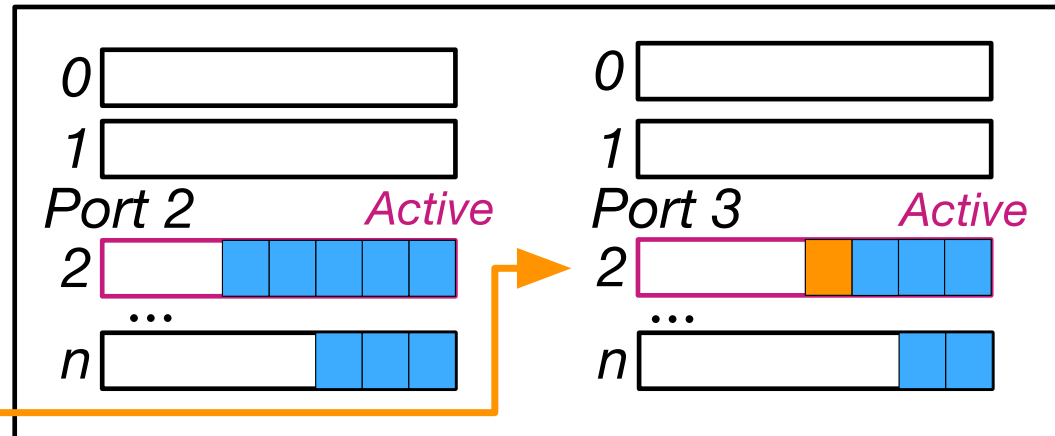
# Under Time-Flow Table: Time-Based Packet Scheduling



Time slice = 2

*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0
1	5	3	2




# Under Time-Flow Table: Time-Based Packet Scheduling

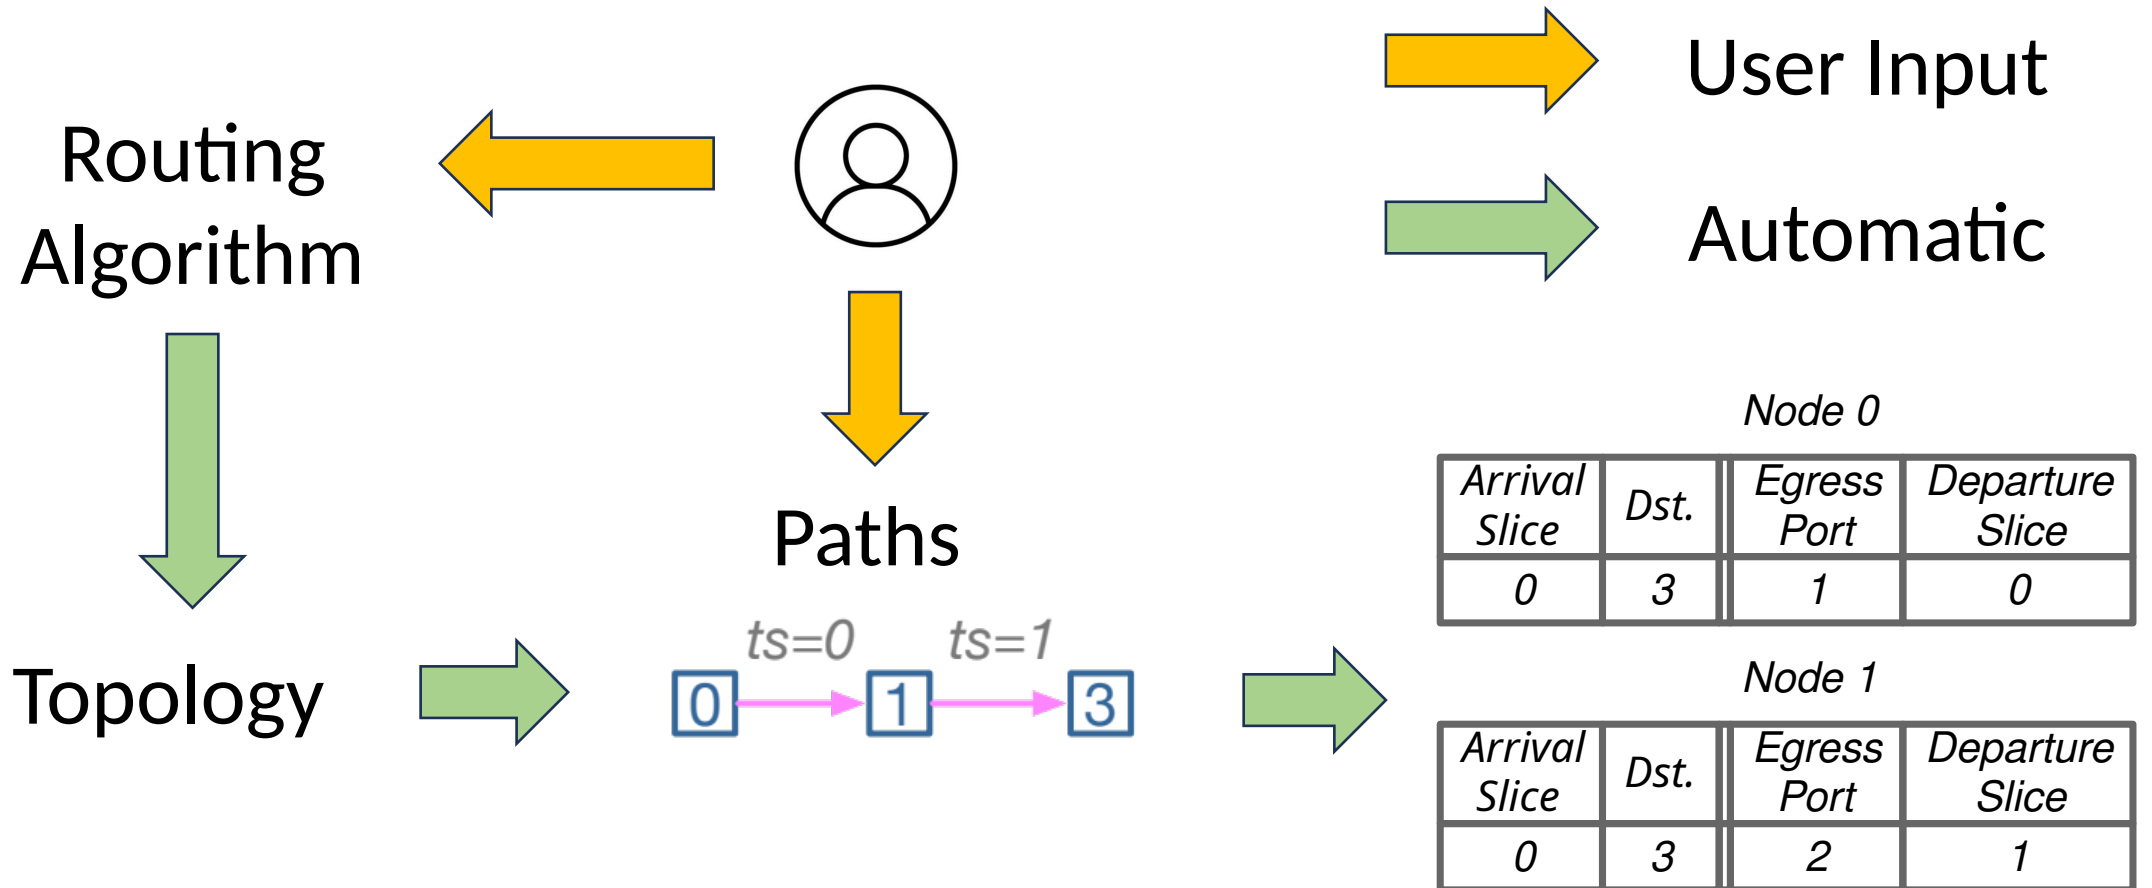


*Time-Flow Table*

<i>Arr. Slice</i>	<i>Dst.</i>	<i>Egr. Port</i>	<i>Dep. Slice</i>
0	5	2	0
1	5	3	2



# Above Time-Flow Table: Automatic Time-Flow Table Generation



# User APIs

## **Routing:**

*add\_path(Path)*

→ Registers a routing path between source and destination

# User APIs

## Routing:

*add\_path(Path)*

→ Registers a routing path between source and destination

*routing\_algo(src, dst, time\_slice, topo) -> Path*      [User-defined]

→ Custom routing logic implemented by user

# User APIs

## Routing:

*add\_path(Path)*

*routing\_algo(src, dst, time\_slice) -> Path* [User-defined]

## Topology:

*connect(node1, node2, time\_slice)*

# User APIs

## Routing:

*add\_path(Path)*

*routing\_algo(src, dst, time\_slice) -> Path* [User-defined]

## Topology:

*connect(node1, node2, time\_slice)*


## Monitor:

*queue\_depth(switch, port, queue)*

*packet\_drop(switch)*

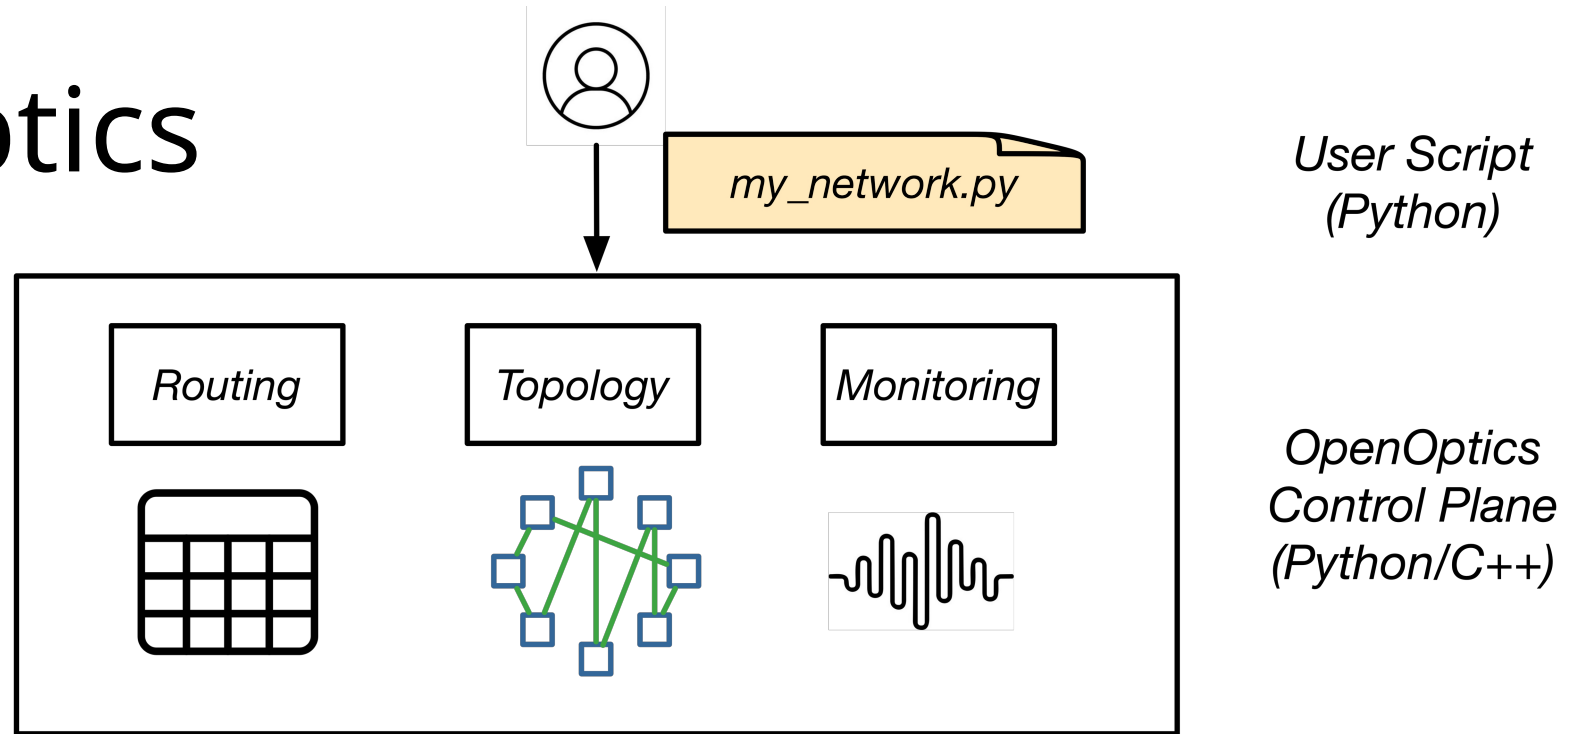
...

# OpenOptics System

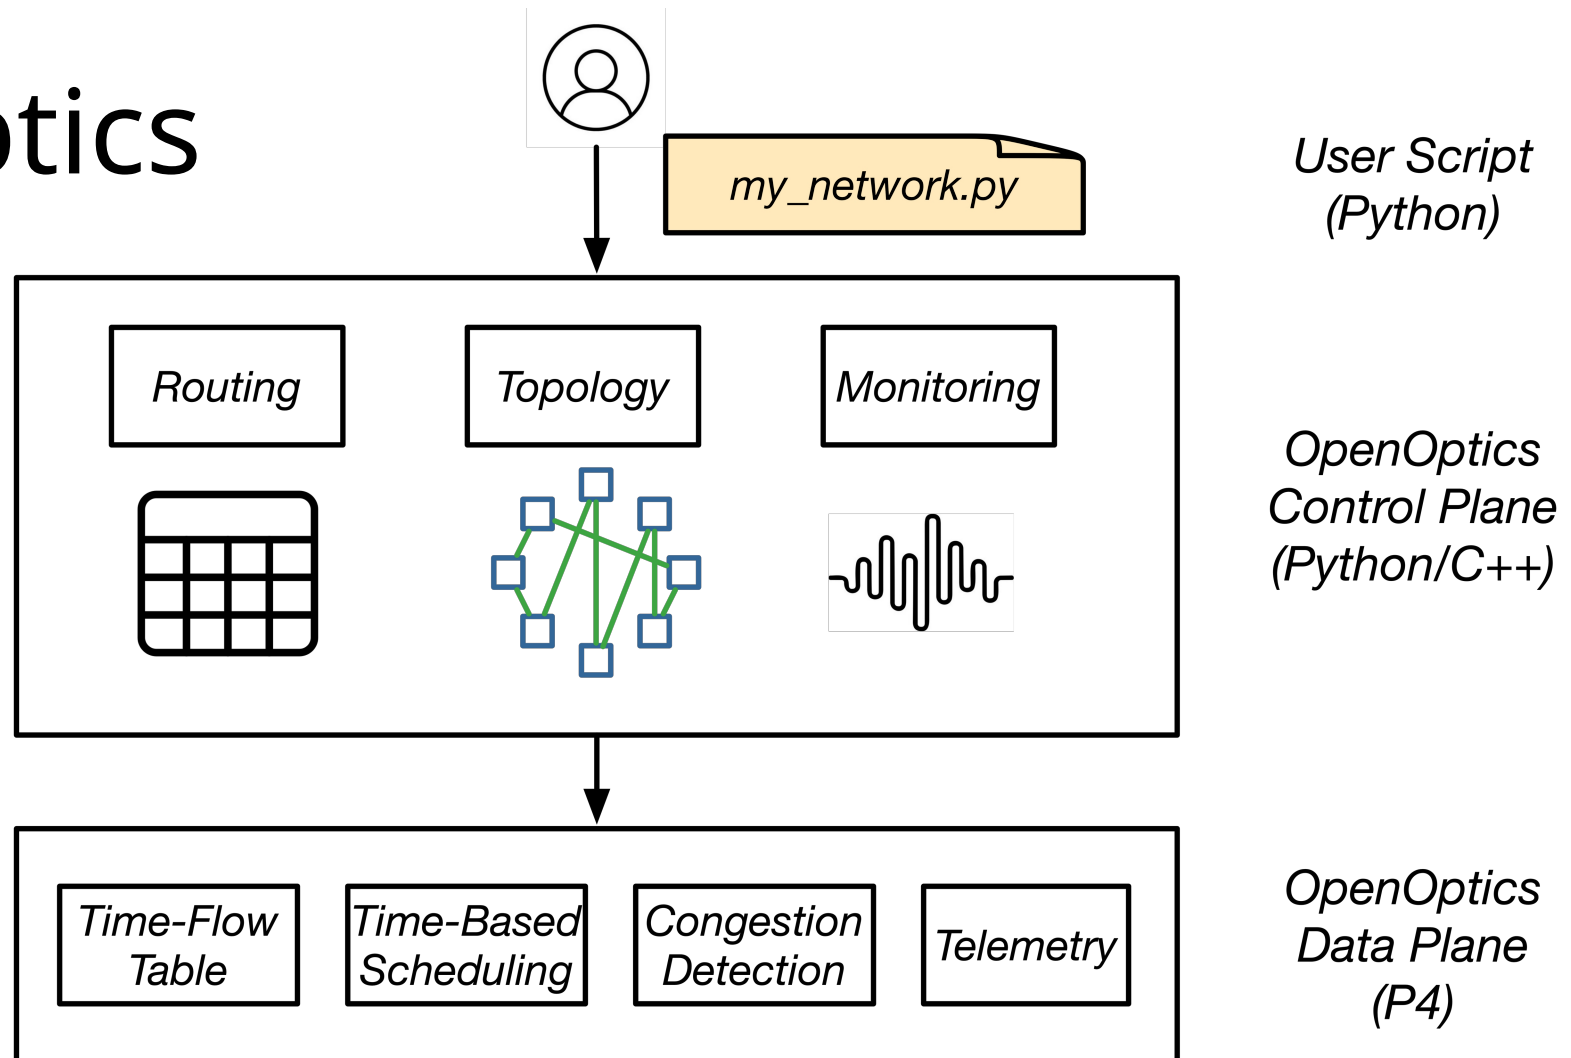


*User Script  
(Python)*

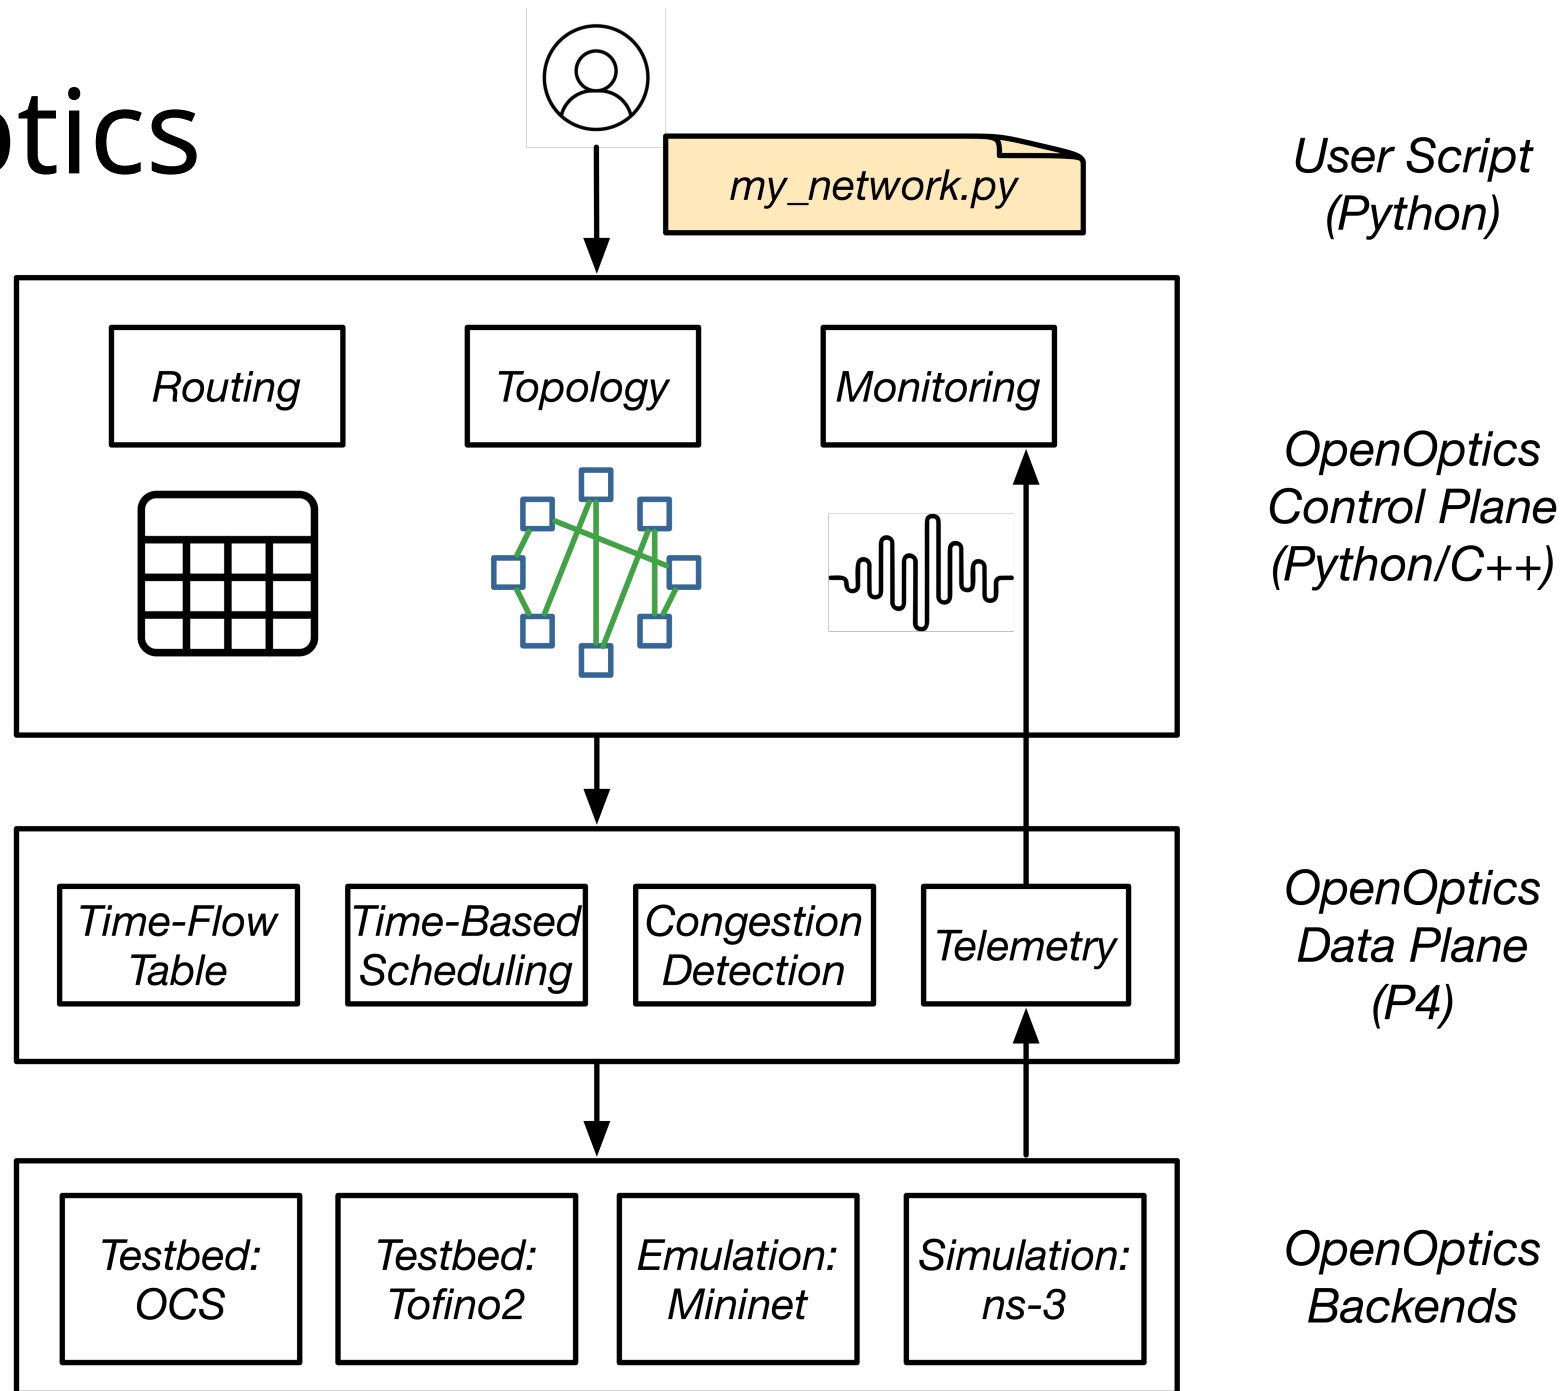
# OpenOptics System



# OpenOptics System



# OpenOptics System



# Infrastructure Services

- Congestion detection
  - Traffic push-back
  - Buffer offloading
  - Flow pausing
  - Traffic collection
- 
- Please refer to our paper

# The Question:

*How Easy Can It Be to Deploy an Optical DCN  
in OpenOptics?*

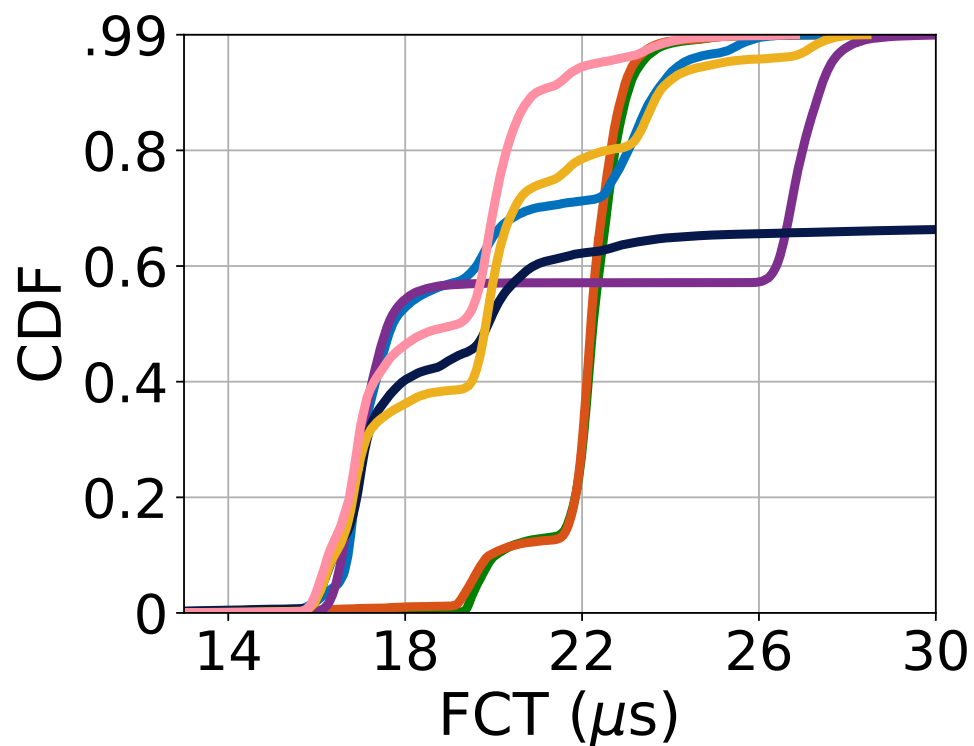
# Answer: As Little as 10 Lines of Python

```
1 from openoptics import Toolbox, OpticalTopo, OpticalRouting
2
3 net = Toolbox.BaseNetwork(nb_node=8, time_slice_duration_ms=128)
4 circuits = OpticalTopo.round_robin(nb_node=8)
5 net.deploy_topo(circuits)
6 paths = OpticalRouting.routing_hoho(net.get_topo())
7 net.deploy_routing(paths)
8 net.start()
```

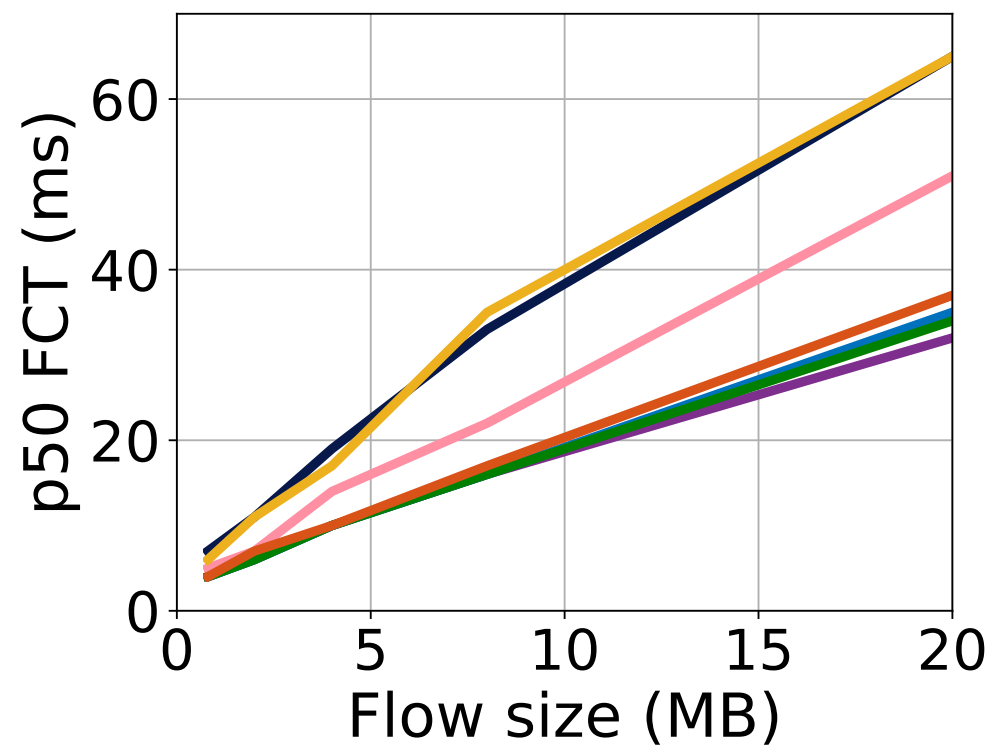
# Evaluation on the Testbed



# Cross Architecture Comparison




Memcached

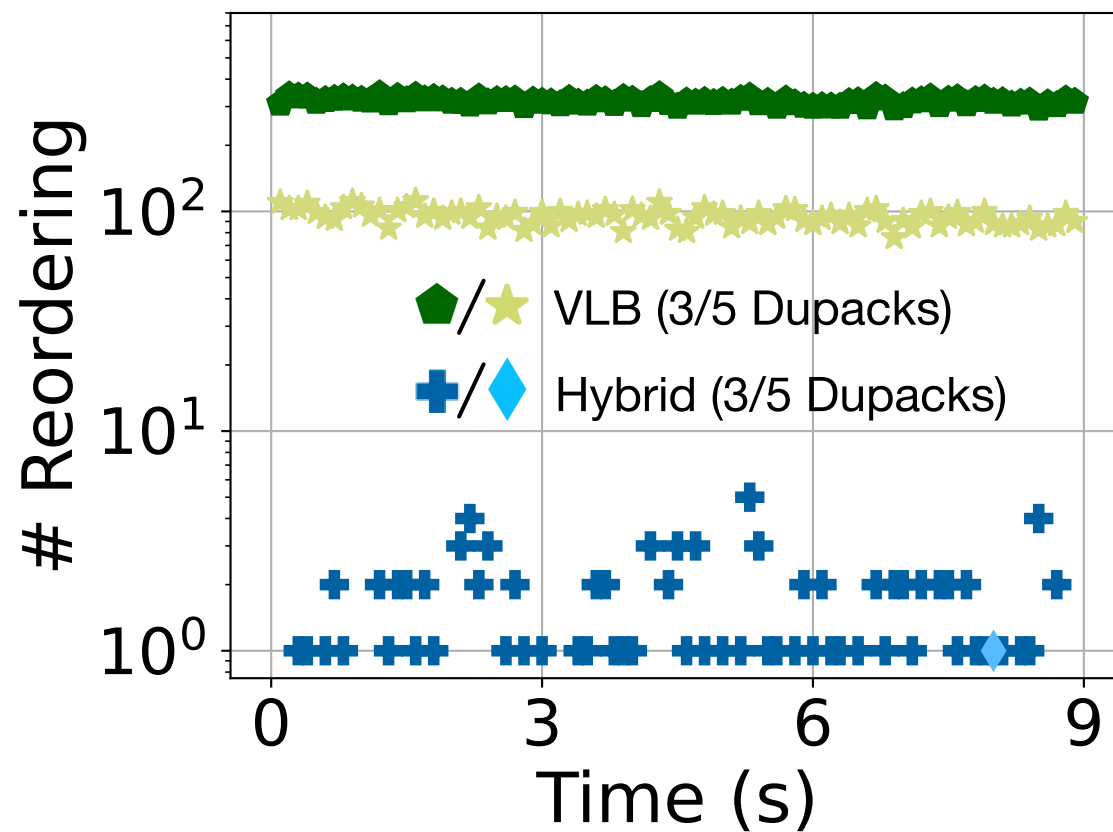
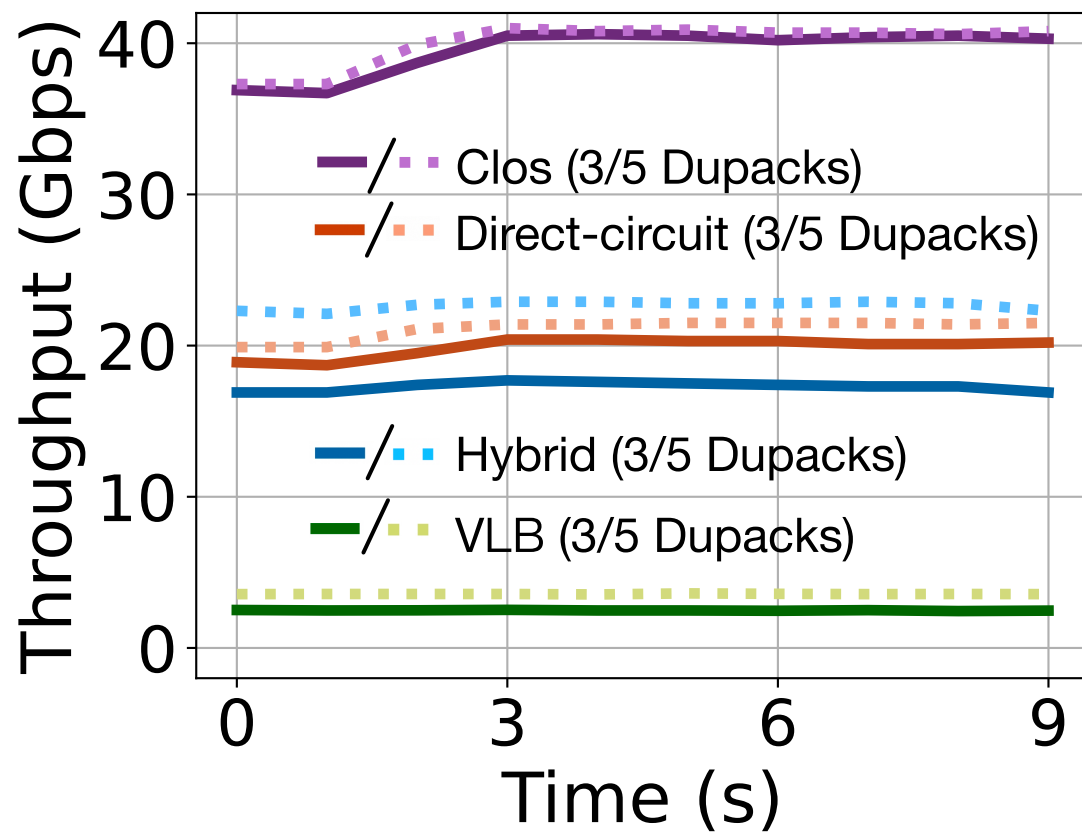


All-reduce

# Compare with Original Results



# Case Study: Congestion Control



# Observability: Dashboard

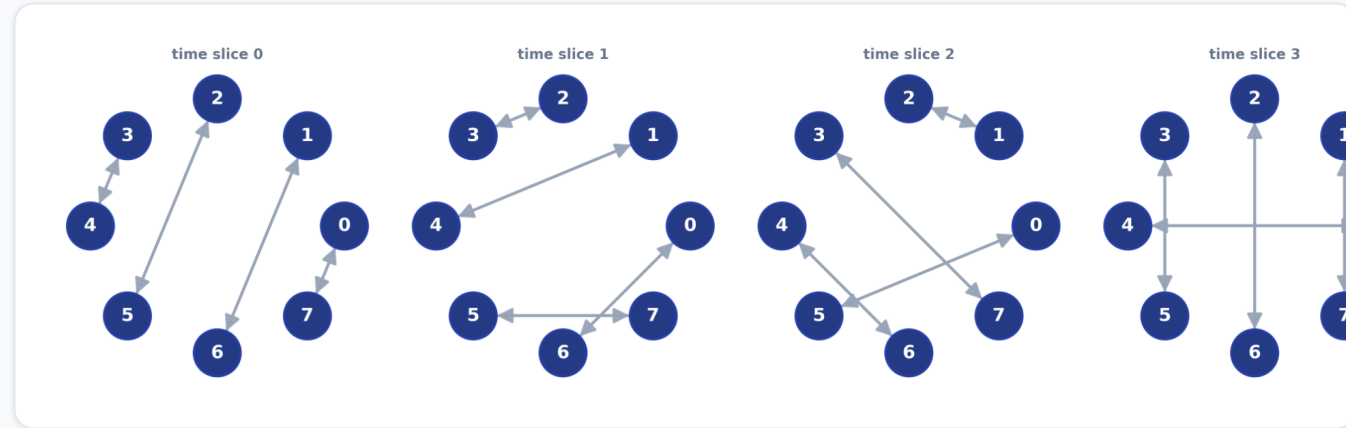


## OpenOptics Dashboard

24-04-2026 10:40:18


24-04-2026 10:41:32

24-04-2026 10:43:25




### Queue Depth (packets)

#### network




- switch=tor0
- switch=tor1
- switch=tor2
- switch=tor3
- switch=tor4
- switch=tor5
- switch=tor6
- switch=tor7

#### tor0



- port=0, queue=0
- port=0, queue=1
- port=0, queue=2
- port=0, queue=3
- port=0, queue=4
- port=0, queue=5
- port=0, queue=6

# Open Source



⚡ OpenOptics ⚡

🔍 Search

⌘ + K

Quick Start

Installation

Tofino Backend

Examples

Topology Primitive

Routing Primitive

Direct Routing

Fastest Path Routing

Source Routing



## OpenOptics: Democratizing Optical DCNs

Easy design, testing, and deployment of optical DCNs for everyone.


### Note

- The OpenOptics [paper](#) has been accepted by NSDI'26!


OpenOptics is a general framework for realizing different optical data center network architectures in a plug-and-play manner. With OpenOptics, users can deploy customized optical data center networks on the testbed, emulation, or simulation with ~10 lines of Python code. Under the hood, user configurations are translated into table entries and control plane programs, which are then deployed to the underlying optical and P4-programmable switches.

```
from openoptics import Toolbox, OpticalTopo, OpticalRouting


net = Toolbox.BaseNetwork(name="demo", backend="Mininet",
                          nb_node=8, time_slice_duration_ms=512)
circuits = OpticalTopo.round_robin(nb_node=nb_node)
net.deploy_topo(circuits)
paths = OpticalRouting.routing_direct(net.get_topo())
net.deploy_routing(paths, routing_mode="Per-hop")
```



GitHub



⚡ *OpenOptics* ⚡



Website

Easy design, testing, and deployment of optical DCNs for everyone.



Yiming Lei

ylei@mpi-inf.mpg.de