

COSMOS Tutorial: Programmable Optical x-Haul Network

Craig Gutterman

Electrical Engineering, Columbia University

ACM SenSys 2019

Nov. 10, 2019


Joint work with Tingjun Chen, Artur Minakhmetov, Jiakai Yu, Michael Sherman,
Shengxiang Zhu, Ivan Seskar, Dipankar Raychaudhuri, Daniel Kilper, Gil Zussman

COSMOS Team: Rutgers, Columbia, and NYU in partnership with New York City, IBM, Silicon Harlem, City College of New York, U. Arizona



NSF: Platforms for Advanced Wireless Research (**PAWR**)

- Wireless networking research
 - In 2012, the number of cellular users exceeded the number of toothbrush users
 - Unparalleled growth in the number of devices, data rates, and traffic volume
 - Services evolving from high-speed data and video towards AR/VR and IoT
- Emerging wireless communications paradigms




Platforms for Advanced
Wireless Research



Millimeter
Waves



Small Cell



Massive
MIMO



Beamforming



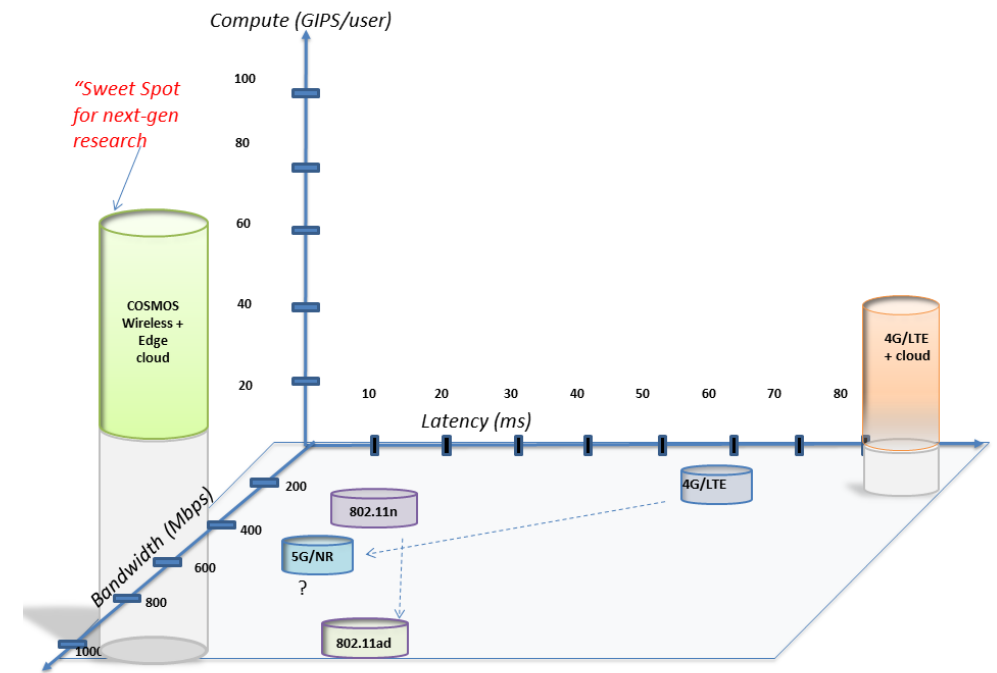
Full Duplex

Figure source: A. Nordrum and K. Clark, "Everything you need to know about 5G," *IEEE Spectrum*, 2017.

- **PAWR: 4 city-scale testbeds for future wireless technologies – COSMOS testbed in NYC**

COSMOS – Project Vision

- COSMOS = Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment
- **Latency** and **compute power** are the two new dimensions for characterizing wireless access
- **Edge computing** is an enabler for real-time services
- Latency for 4G cellular >50 ms, while targets for 5G are <10 ms
- COSMOS will enable researchers to investigate ultra-high bandwidth (~Gbps), low latency (<5 ms), and edge computing (~10–100 GIPS)




Optical Networking – Backbone of COSMOS



320x320 Optical Space Switch



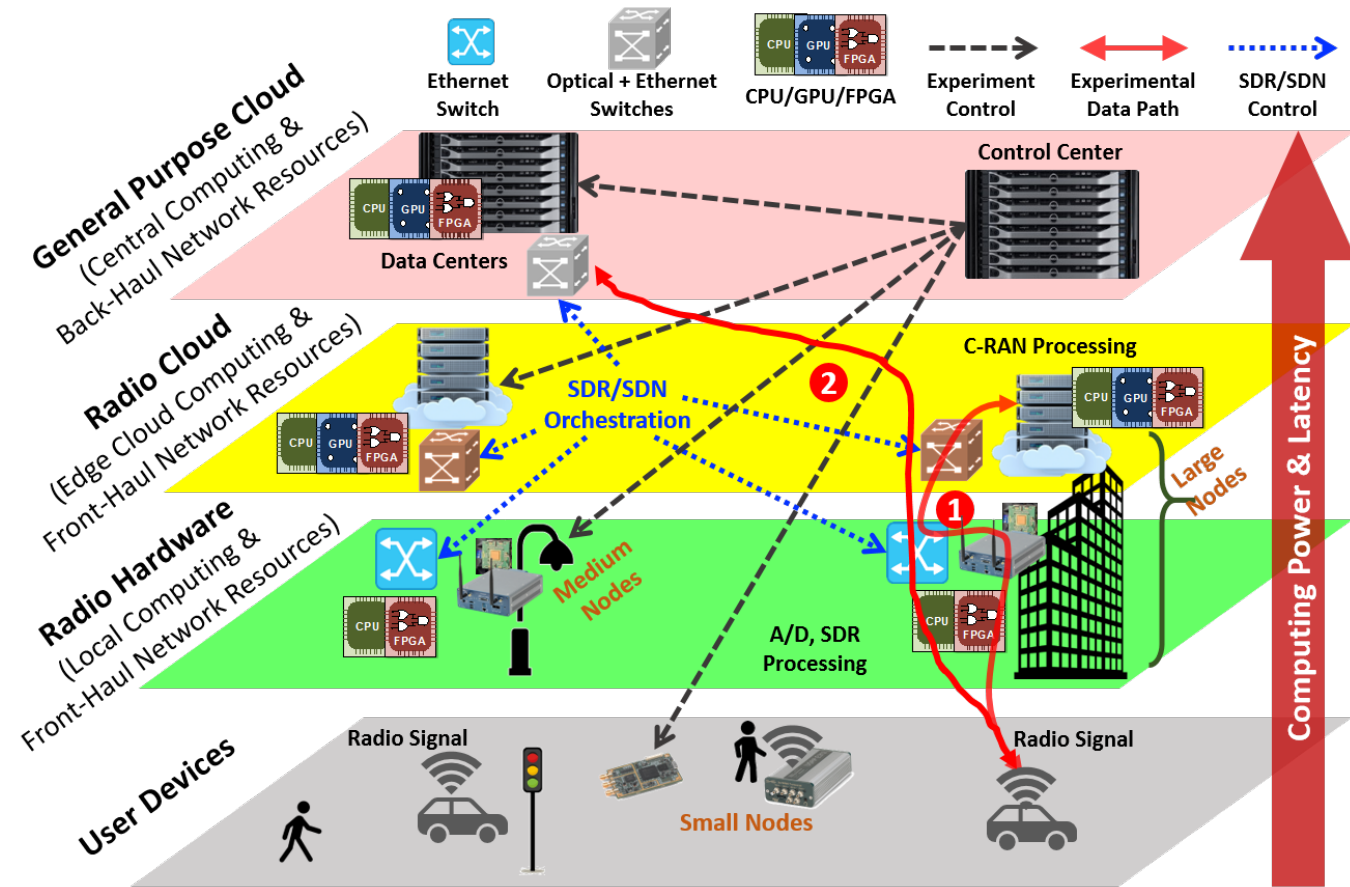
1x20 ROADMs



7-mile Fiber Connections

COSMOS – System Architecture

- Key design challenge: Gbps performance + full programmability at the radio level
- Fully programmable multi-layered computing architecture for flexible experimentation
- **Key technologies:**
 - Software-define radio (SDR)
 - mmWave
 - **Optical networks**
 - **Software-defined networking and cloud**
 - Control and management software



COSMOS's multi-layered computing architecture with different data paths of example experiments with local/remote computing

Objective: Take it Outside

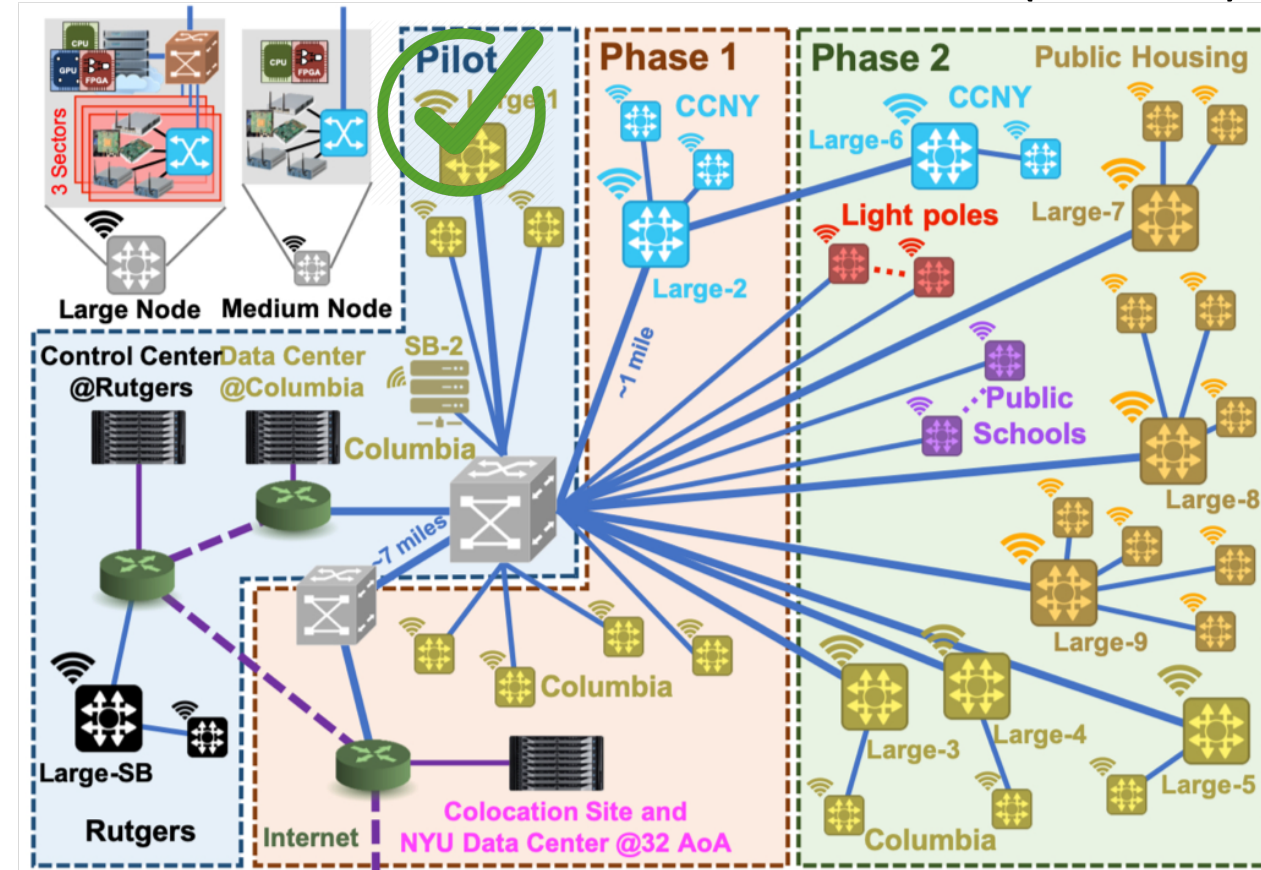



Objective: Take it Outside



COSMOS Deployment – A Phased Approach

West Harlem w/ an area of ~1 sq. Mile + Fiber connection to NYU Data Center (32 AoA)



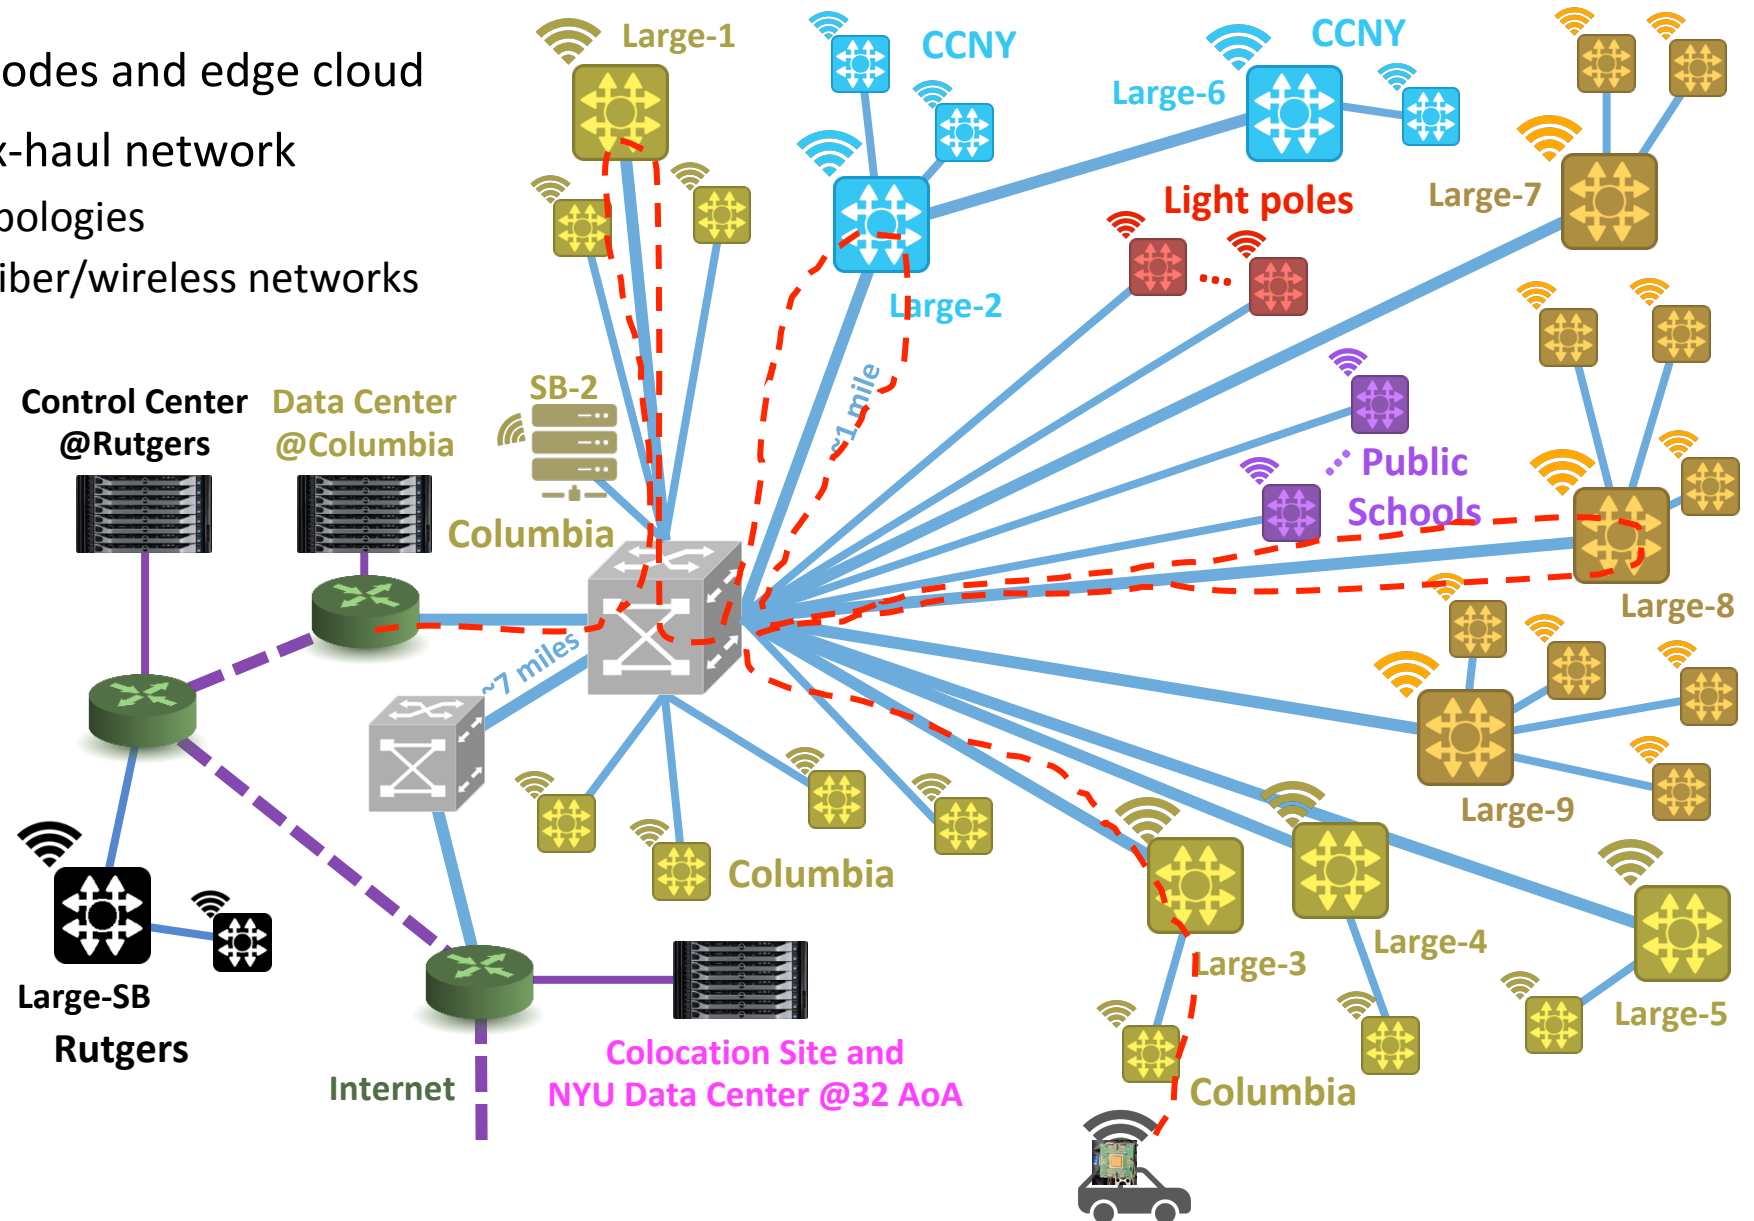
~9 **Large** sites
+ Large Sand Box

~40 Medium sites
-Building side- or light-pole-mounted

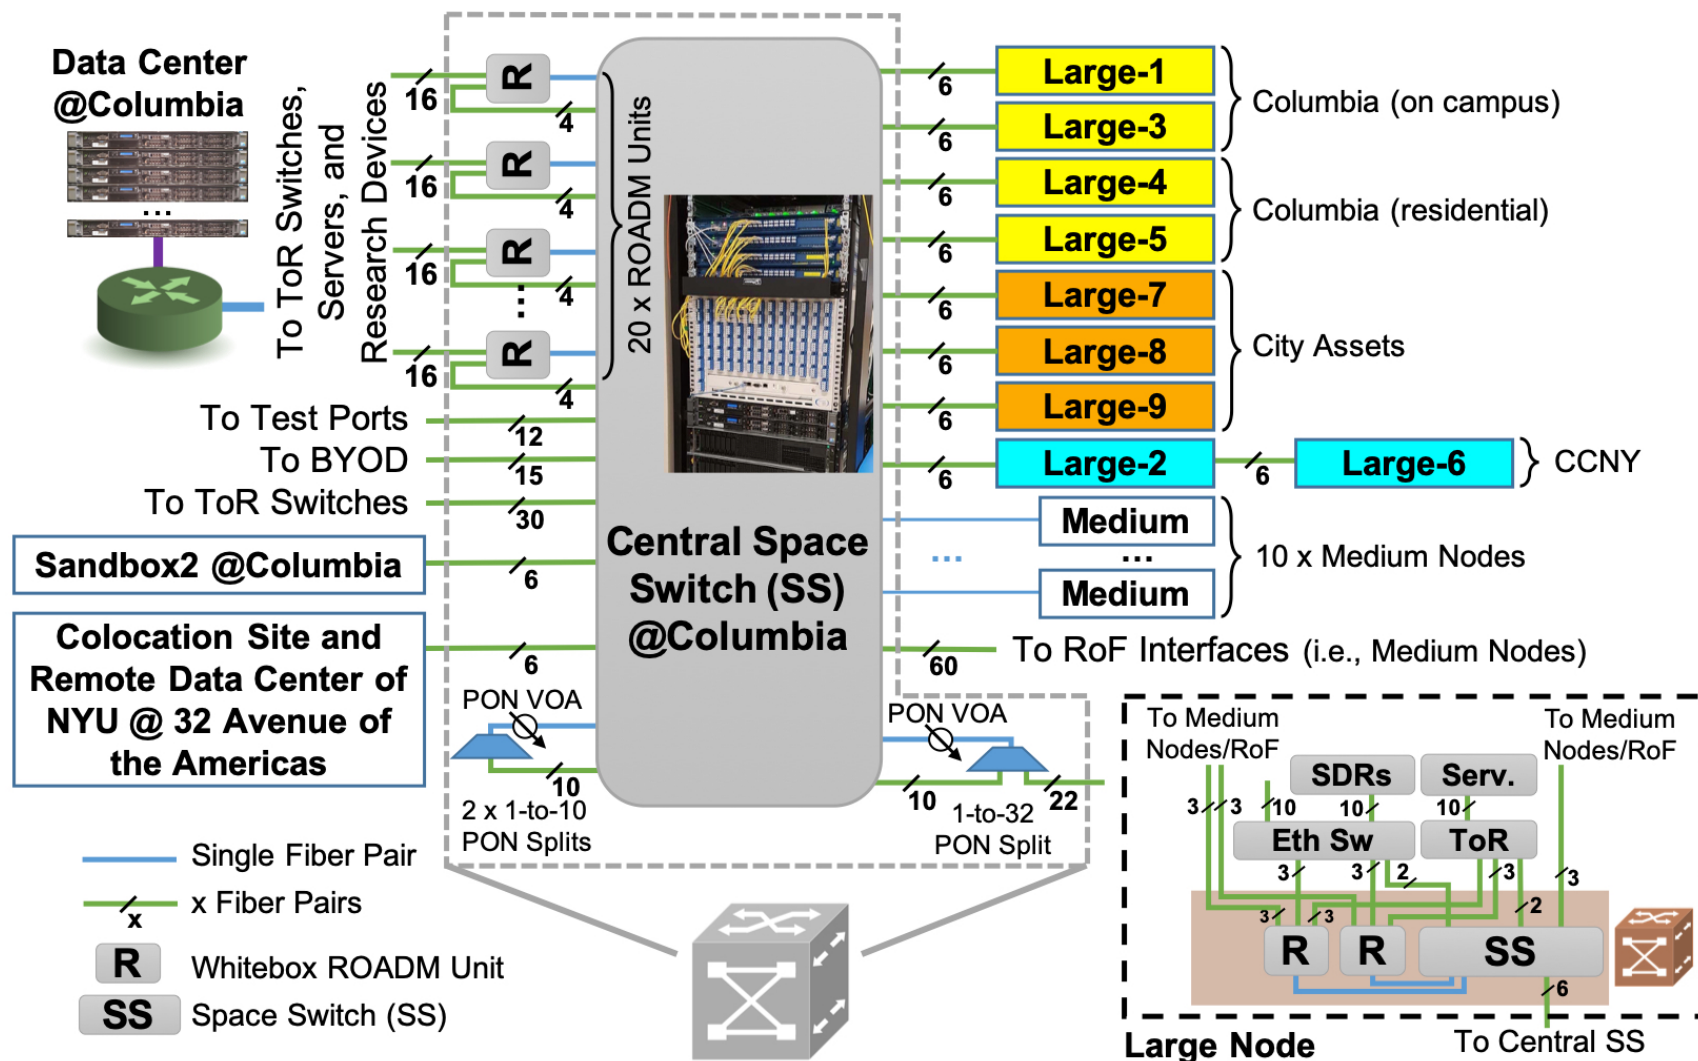
~200 Small nodes
-Including vehicular and hand-held

COSMOS Key Technologies – Optical Networking

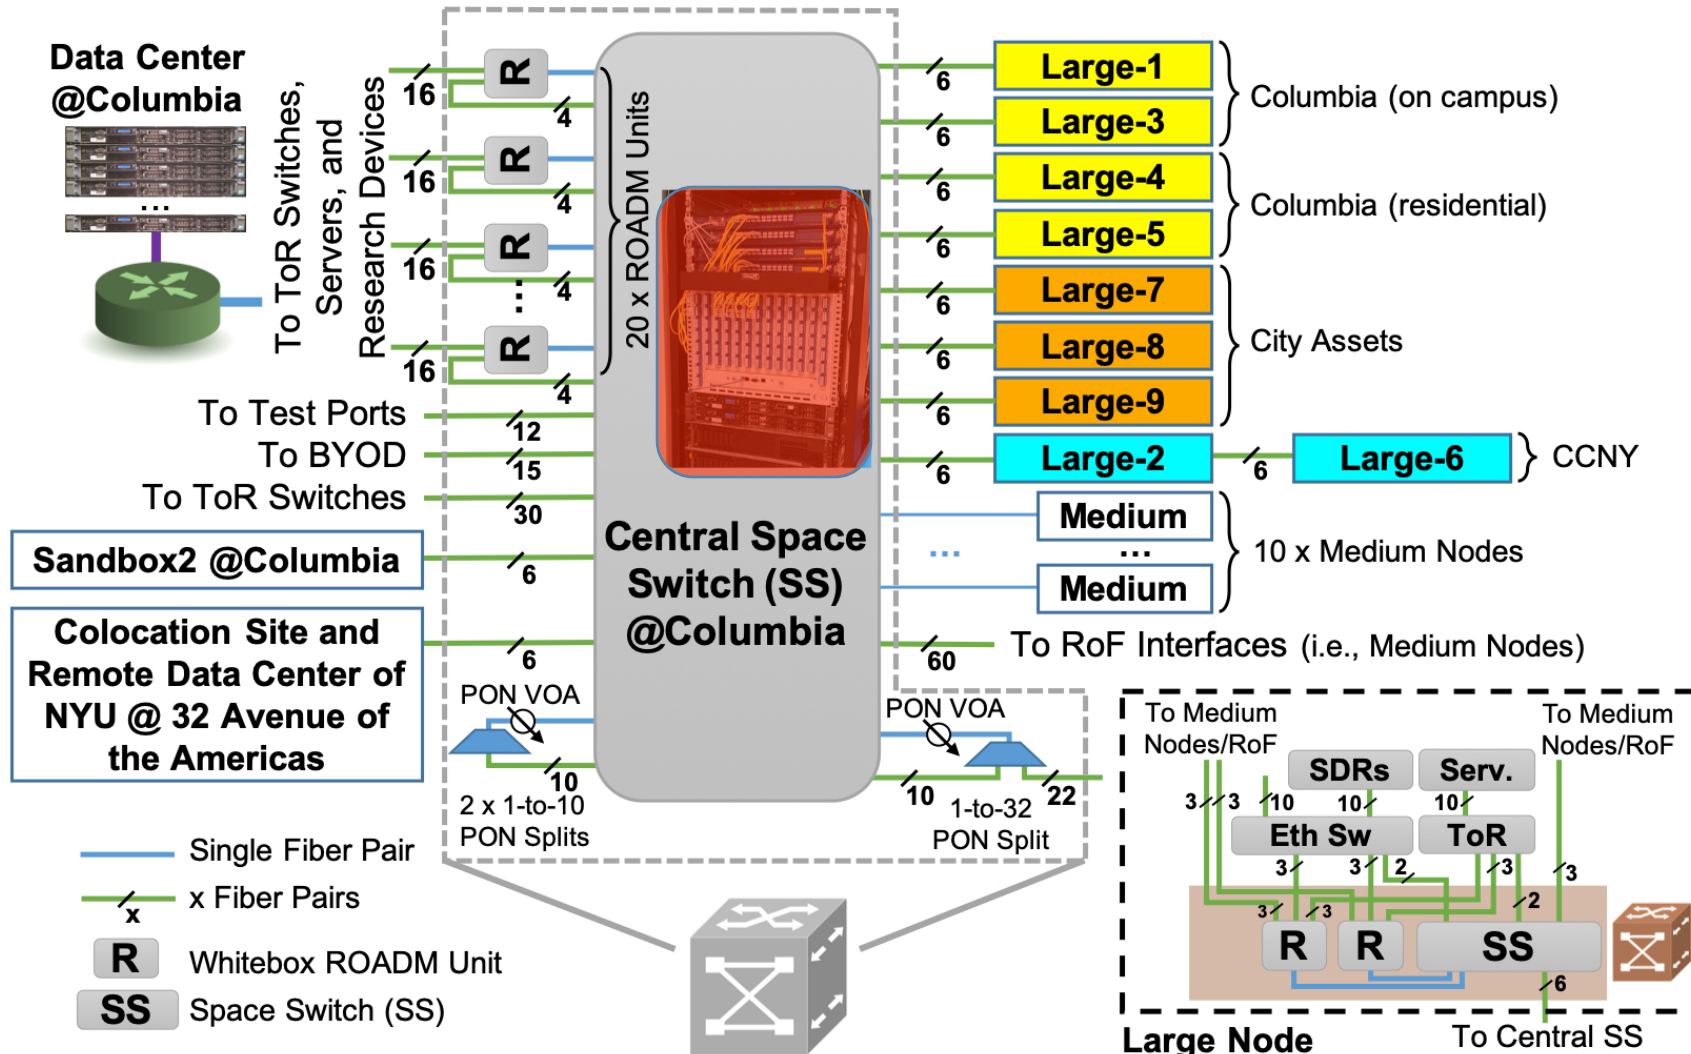
- Fast connectivity b/w radio nodes and edge cloud
- Fast and low latency optical x-haul network
 - Configure wide range of topologies
 - Experiment on converged fiber/wireless networks



COSMOS Key Technologies – Optical Networking



COSMOS Key Technologies – Optical Networking: Calient Switch



320x320 Optical Space Switch


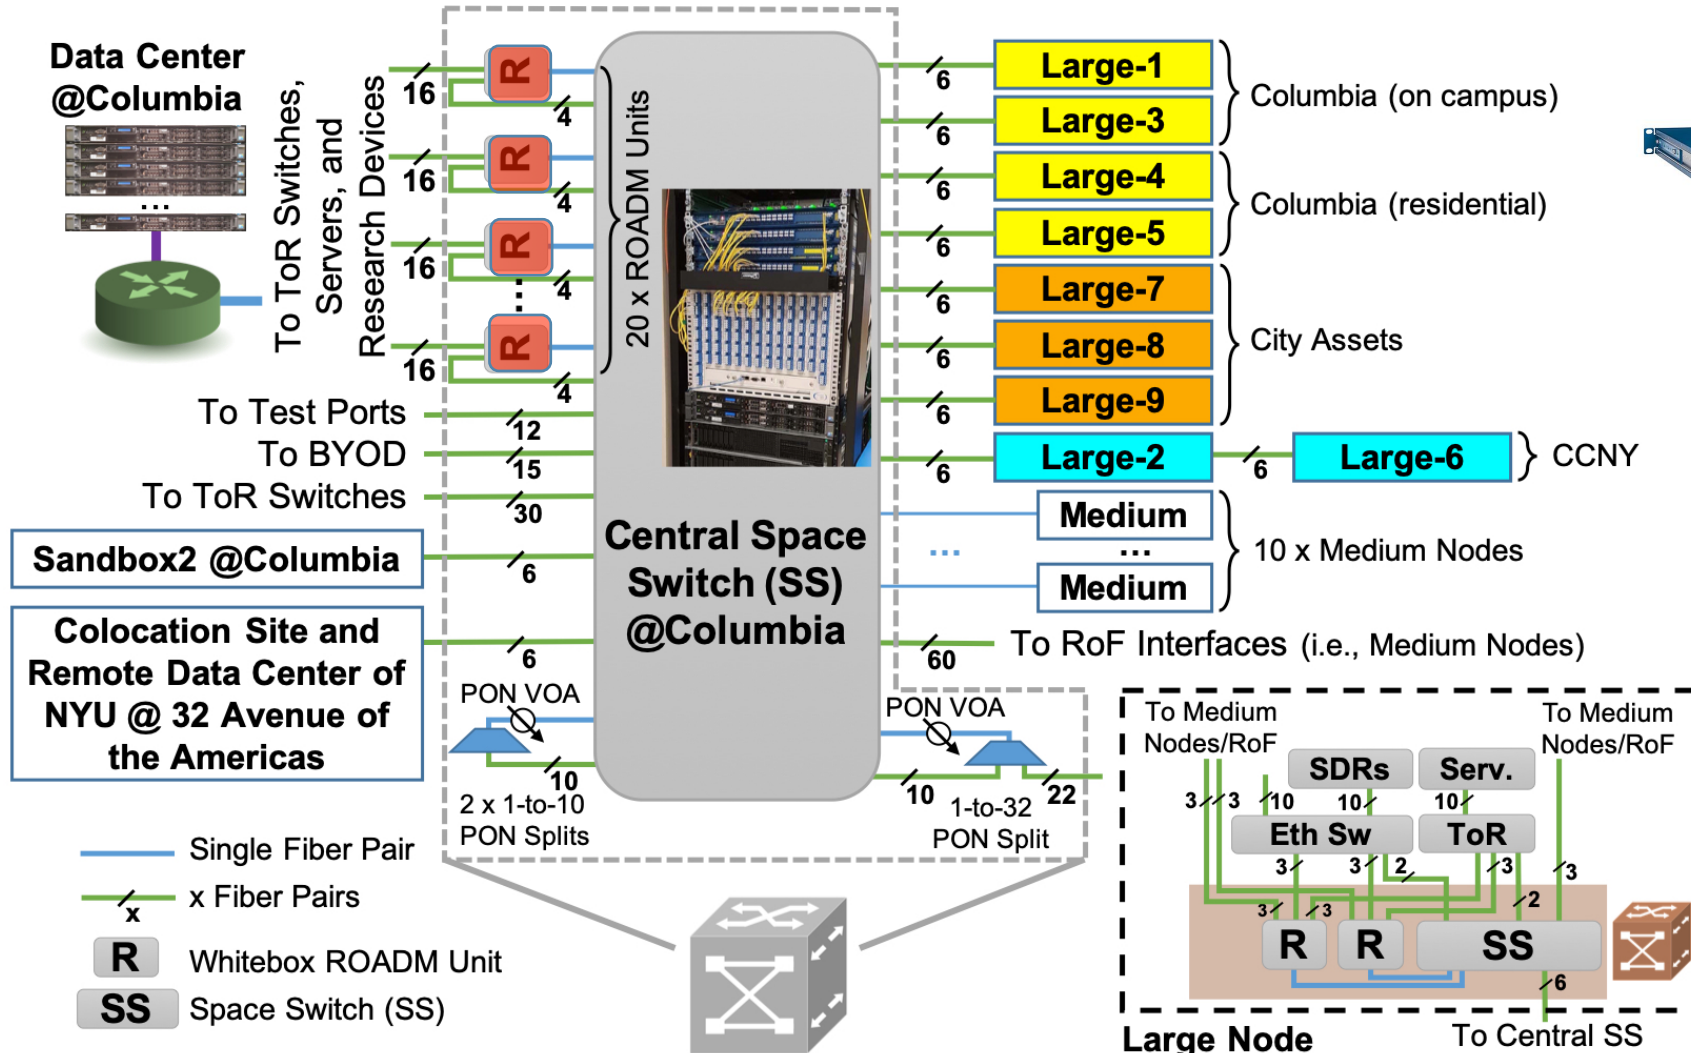
Smart “Fiber patch panel”

- Programmable
- Controllable from distance
- Dynamic changes of routes


Interconnects:

- ROADMs
- ToRs
- Nodes
- Devices

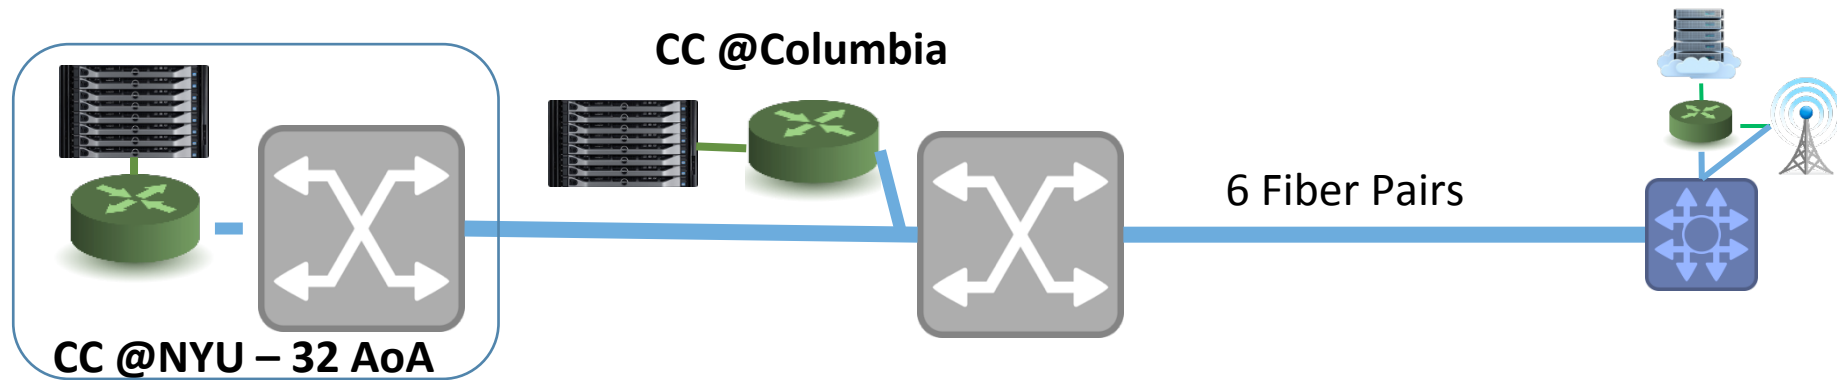
COSMOS Key Technologies – Optical Networking: ROADMs



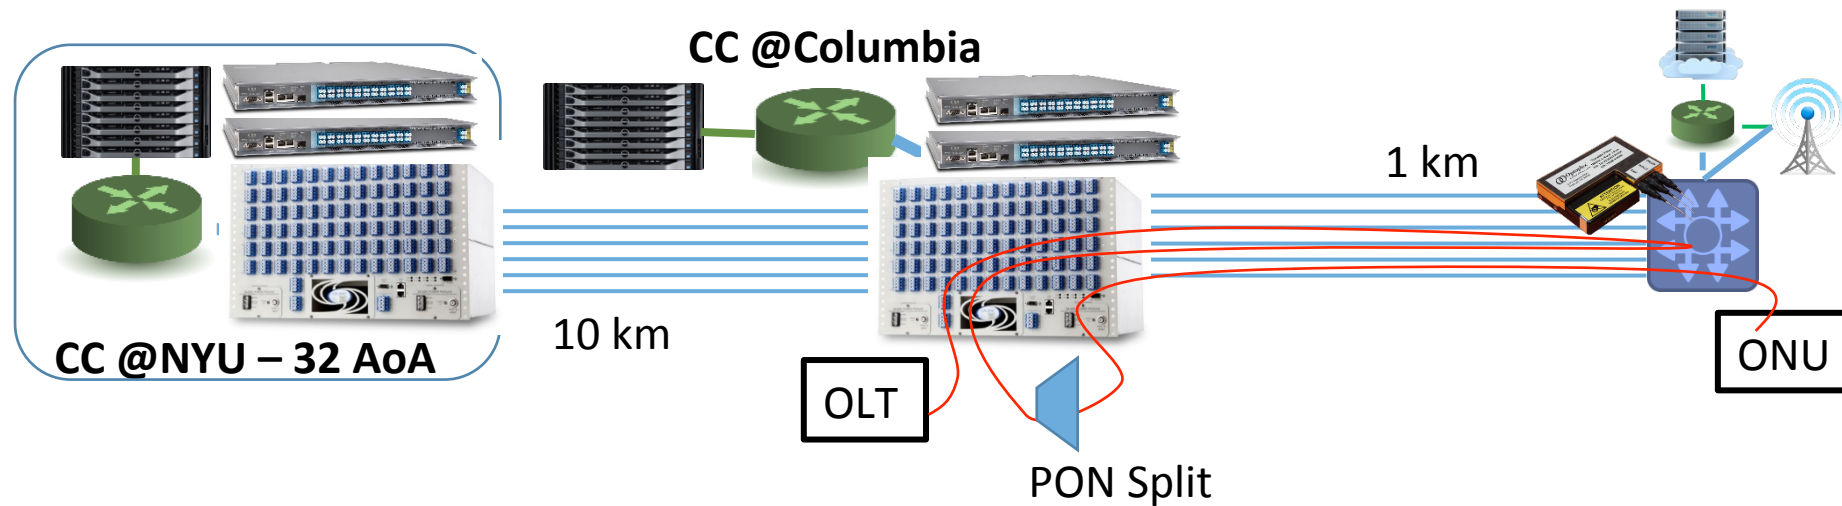
2nd degree ROADM node



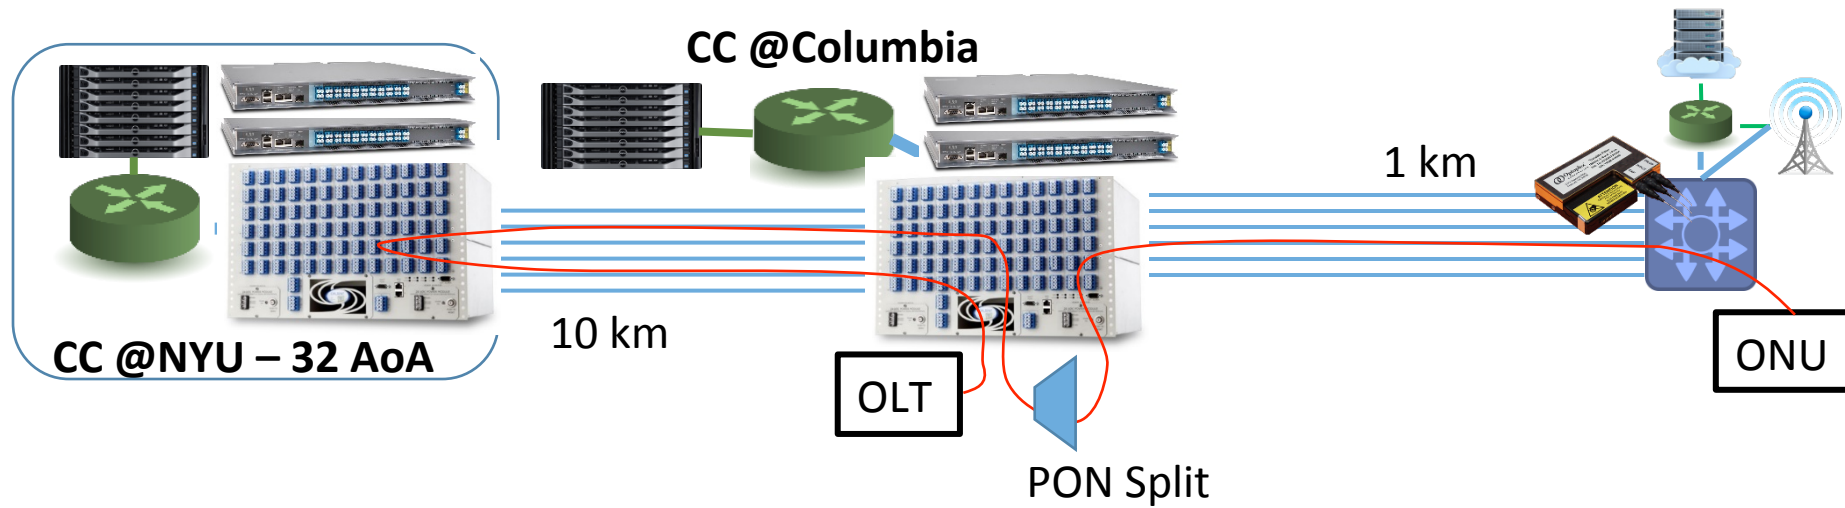
Programmable Topologies: Example



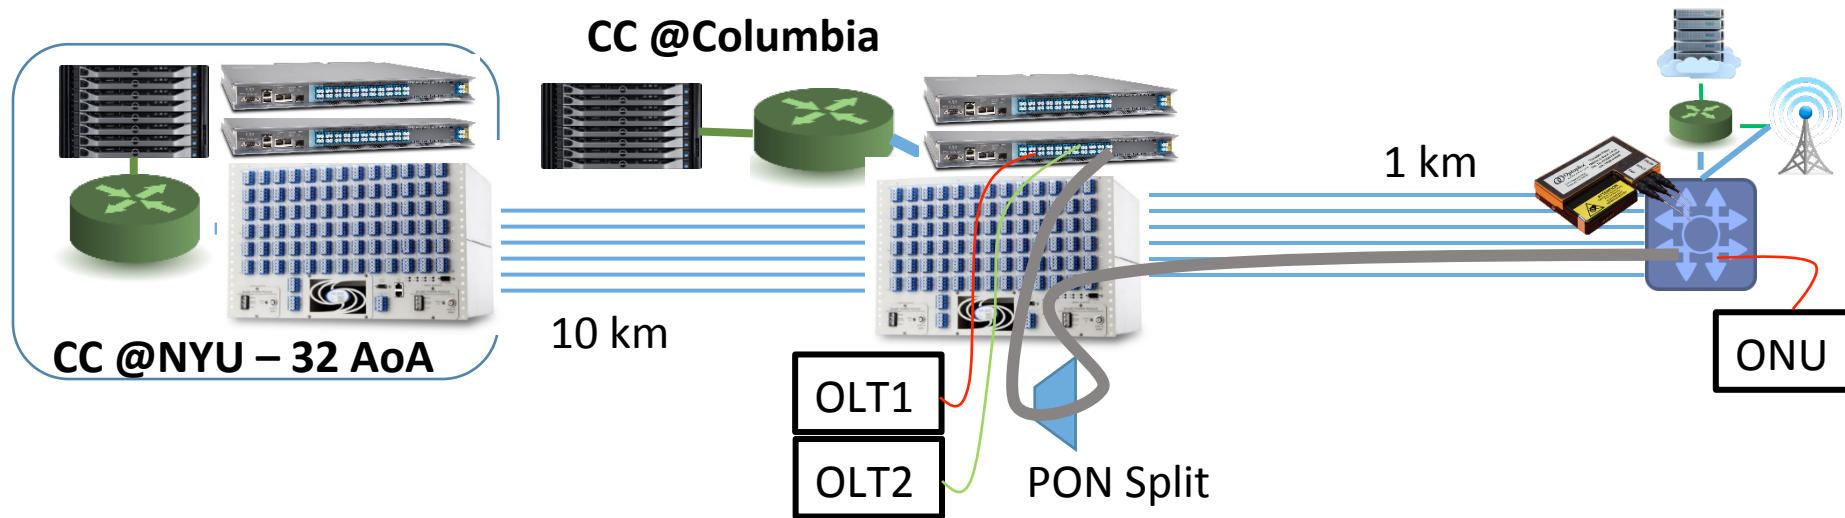
Programmable Topologies – PON



Programmable Topologies – Long Reach PON



Programmable Topologies – WDM PON



Programmable Topologies – Calient Optical Space Switch



320 x 320 bidirectional fiber interconnects :

- ROADMs
- ToRs
- Nodes
- Devices

Programmable Topologies – Calient Optical Space Switch

FROM ID: [ENTER ALIAS]

Cross Connection 1.2.1-1.5.1 Details

Connection Details

Connection ID	AS	OS	OC
1.2.1-1.5.1	CL	IS	OK
1.2.1-1.5.1	CL	IS	OK
1.5.1-1.2.1	CL	IS	OK

Group Name: 1.2.1-1.5.1

Direction: BI

LightBand: C

BoardID: admin

Port 1: 1.2.1

Port 2: 1.5.1

AutoProvision: Enabled

No Light Connect: Disabled

ProvisionTime: 02/20/19 16:14:05

Optical Power Level Details

Group Name	Connection Name	Conn ID	Port 1.2.1	Port 1.5.1	Loss
SYSTEM	1.2.1-1.5.1	1.2.1-1.5.1	-90.00	-90.00	-0.00

Active Fault Reports

Event Time	Description
------------	-------------

Fault History

Event Time	Description
------------	-------------

Buttons: Activate, Deactivate, Retry, Delete, Refresh



Home >> Summary

Alarms/Events Status Symbol Key REST API Change Password Logout

admin [05-20-2019 09:41:8]

Cross Connections Summary

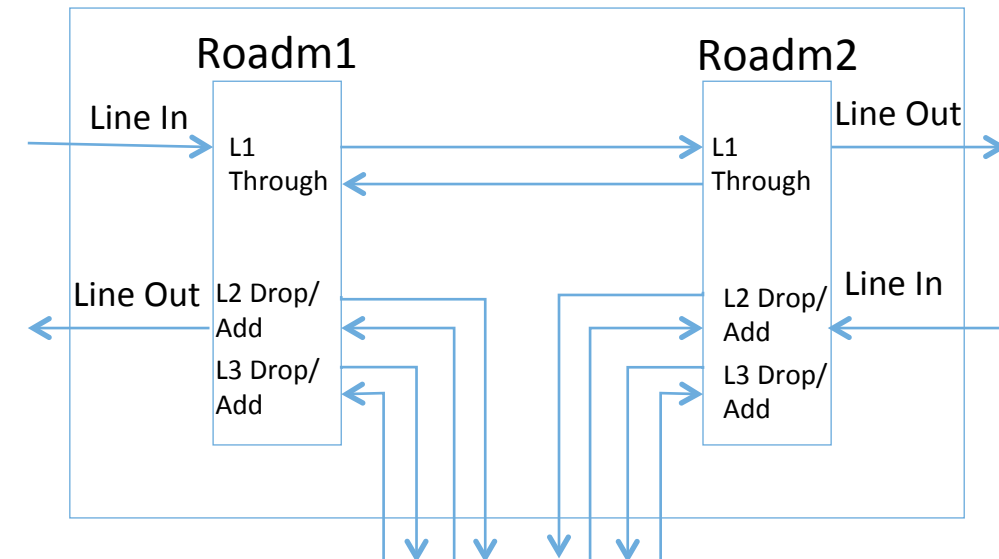
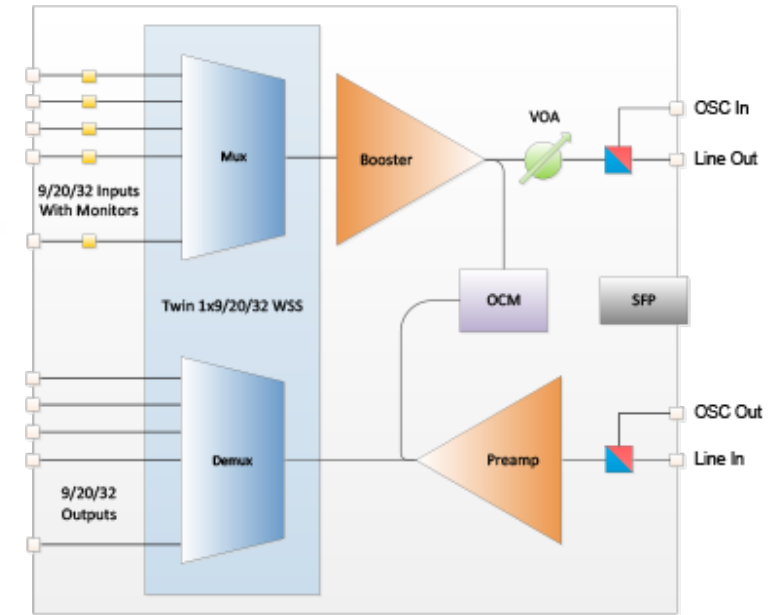
Group Filter: All

Export CSV

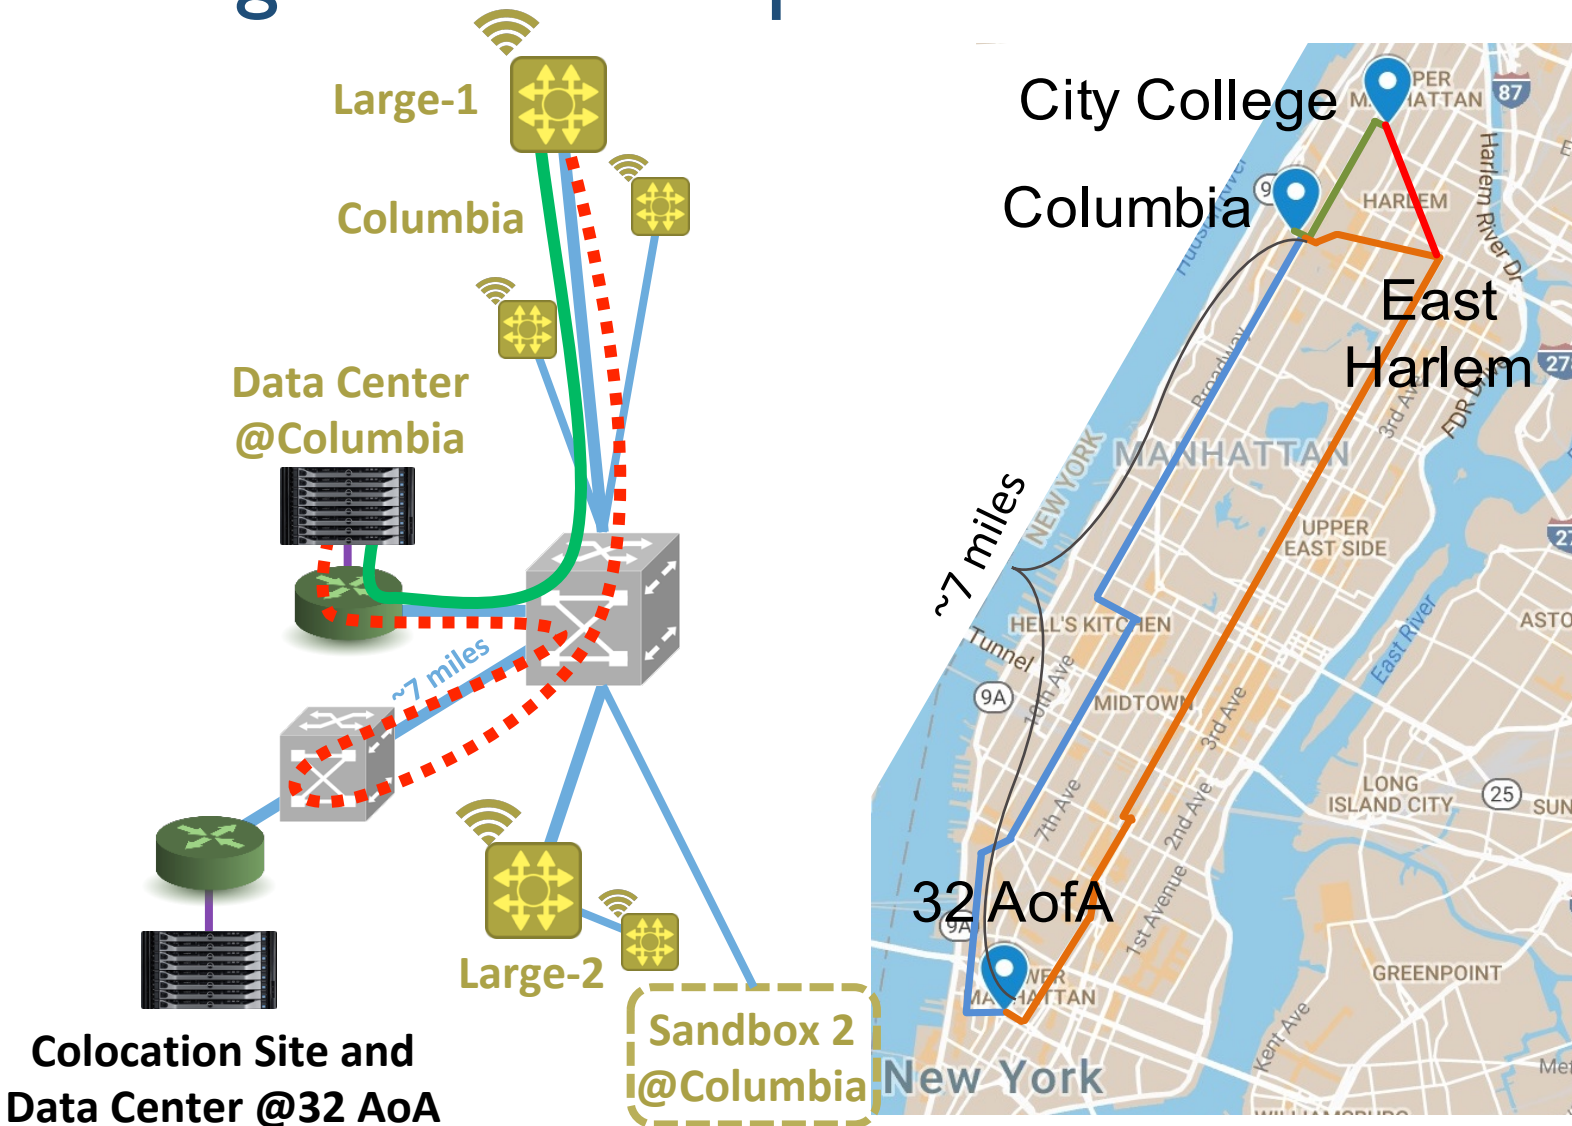
Count	Group	Connection Name	Connection ID	Dir	Band	Conn - Half	IN Power (dBm)	OUT Power (dBm)	Loss (dB)	Alarm	AS	OS	OC
1	SYSTEM	1.1.1-1.4.1	1.1.1-1.4.1	BI	CBAND	1.1.1>1.4.1	-5.60	-7.09	1.50	CL	IS	IS	OK
2	SYSTEM	1.1.2-5.7.1	1.1.2-5.7.1	BI	CBAND	1.1.2>5.7.1	-90.00	-90.00	-90.00	CL	UMA	RDY	FAIL
3	SYSTEM	1.4.8-5.8.1	1.4.8-5.8.1	BI	CBAND	1.4.8>5.8.1	-16.89	-17.98	1.09	CL	IS	IS	OK
4	SYSTEM	1.7.2-5.7.3	1.7.2-5.7.3	BI	CBAND	1.7.2>5.7.3	-8.61	-10.18	1.57	CL	IS	IS	OK
5	SYSTEM	1.7.4-5.5.1	1.7.4-5.5.1	BI	CBAND	1.7.4>5.5.1	-15.84	-18.18	2.34	CL	IS	IS	OK
6	SYSTEM	2.2.8-1.1.8	2.2.8-1.1.8	BI	CBAND	2.2.8>1.1.8	-3.49	-5.05	1.55	CL	IS	IS	OK
7	SYSTEM	5.7.4-2.2.2	5.7.4-2.2.2	BI	CBAND	5.7.4>2.2.2	-0.21	-2.31	2.10	CL	IS	IS	OK
8	SYSTEM	5.7.5-2.2.4	5.7.5-2.2.4	BI	CBAND	5.7.5>2.2.4	0.18	-1.10	1.28	CL	IS	IS	OK
9	SYSTEM	5.8.2-1.7.8	5.8.2-1.7.8	BI	CBAND	5.8.2>1.7.8	-17.20	-18.46	1.26	CL	IS	IS	OK

Programmable Topologies - ROADMs

- 3 Basic Sections
 - 96 chn MUX/DEMUX (WSS)
 - Booster Amplifier
 - Pre-Amplifier
- Single degree, bi-dir. ROADMs
 - Combine to form multi-degree
- Python scripts
 - Booster/Preamp control
 - Booster/Preamp monitor
 - WSS connection Management
 - WSS connection monitor
- RYU SDN Controller



Programmable Optical x-Haul Network: Demo Context

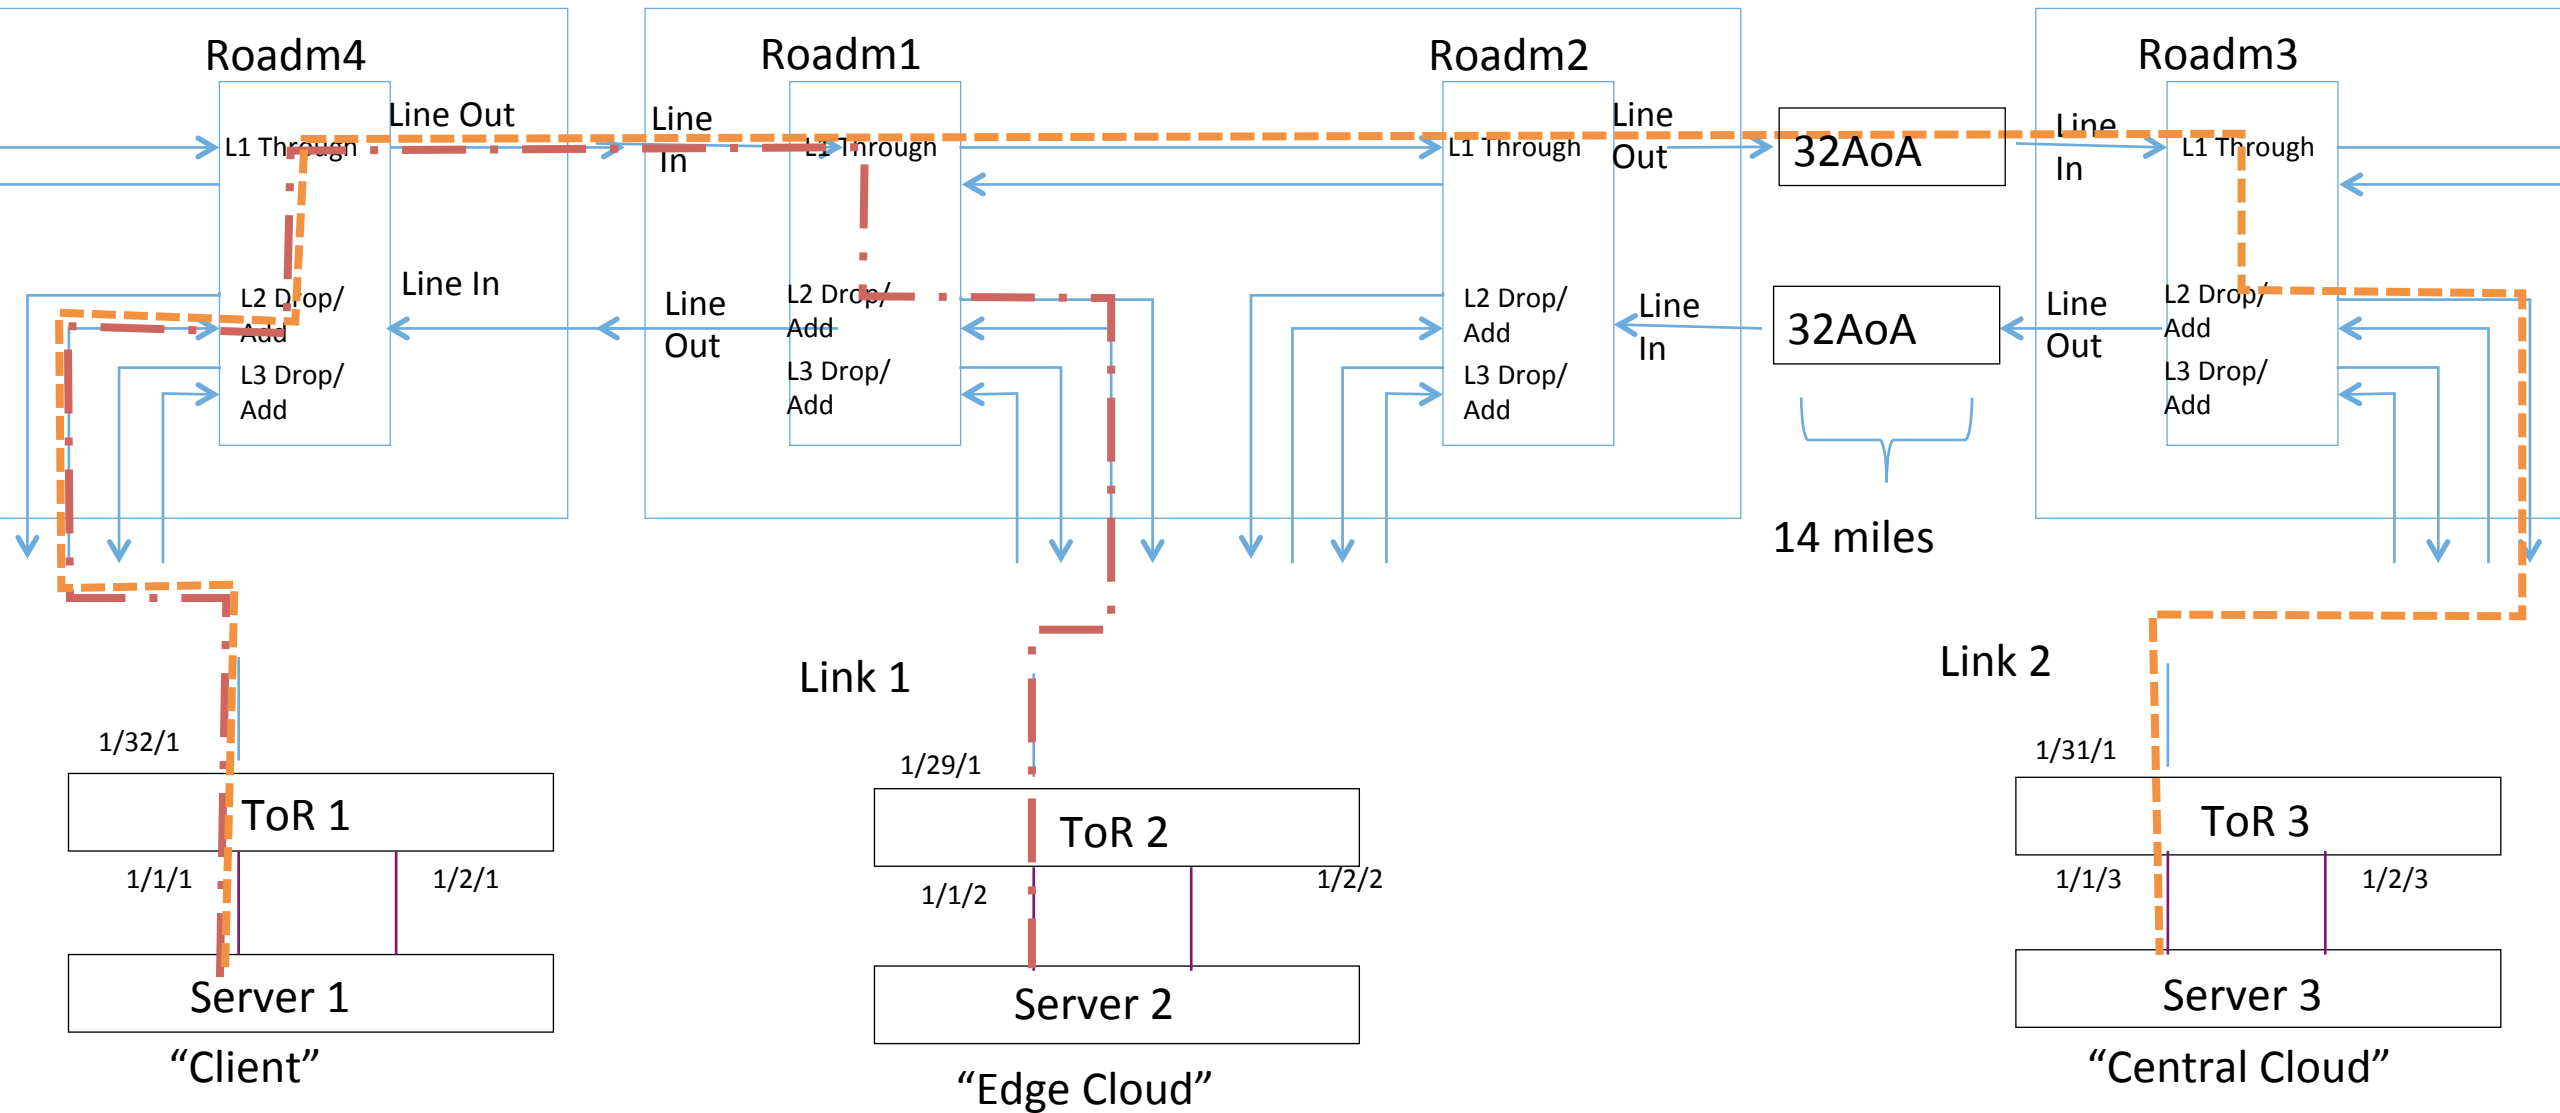


Fiber to 32 AoA, NYC:
facilitated by the City and ZenFi



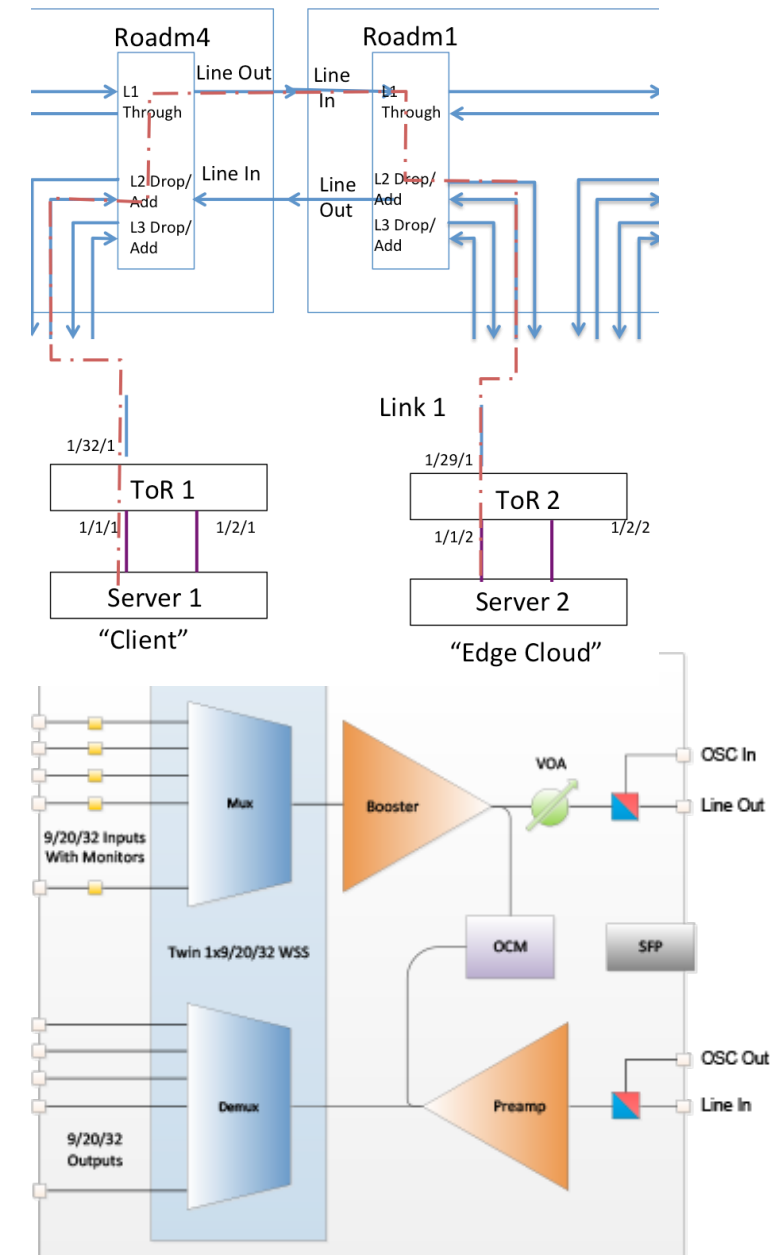
- C. Gutterman, A. Minakhmetov, J. Yu, M. Sherman, T. Chen, S. Zhu, I. Seskar, D. Raychaudhuri, D. Kilper, and G. Zussman, "Programmable optical x-haul network in the COSMOS testbed," in *Proc. IEEE ICNP'19 Workshop Midscale Education and Research Infrastructure and Tools (MERIT)*, 2019.

Programmable Optical x-Haul Network: Demo




Demo: Establish Link 1

- Preliminary Steps
 - Connect line ports of ROADM4 and ROADM1 using the Calient Switch
 - Connect line Ports of ROADM2 and ROADM3 using the Calient Switch
- Steps
 - Add MUX/DEMUX connection from ROADM4 to TOR1
 - Add MUX/DEMUX connection from ROADM1 to TOR2
- Example code
 - `python add_connection.py 10.104.1.4 1 10 in-service false 4102 4201 192950 193050 0 Exp1-FromTor1`
 - `python add_connection.py 10.104.1.4 2 10 in-service false 5101 5202 192950 193050 0 Exp1-TorwardTor1`
 - `python add_connection.py 10.104.1.1 1 10 in-service false 4102 4201 192950 193050 0 Exp1-FromTor2`
 - `python add_connection.py 10.104.1.1 2 10 in-service false 5101 5202 192950 193050 0 Exp1-TorwardTor2`



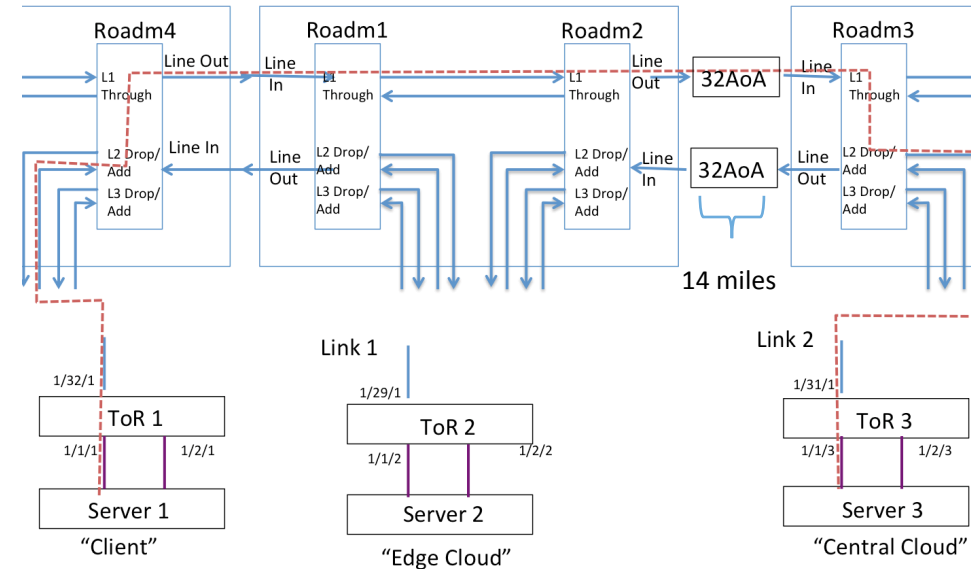
Demo: Establish Link I



Demo: Establish Link 2

- Steps

- Add MUX/DEMUX connection from ROADM4 to TOR1
- Add MUX/DEMUX connection from ROADM1 to ROADM2
- Add MUX/DEMUX connection from ROADM2 to ROADM1
- Add MUX/DEMUX connection from ROADM3 to TOR3



- Example code

- `python add_connection.py 10.104.1.4 1 10 in-service false 4102 4201 192950 193050 0 Exp1-FromTor1`
- `python add_connection.py 10.104.1.4 2 10 in-service false 5101 5202 192950 193050 0 Exp1-TorwardTor1`
- `python add_connection.py 10.104.1.1 1 10 in-service false 4101 4201 192950 193050 0 Exp1-ROADM2`
- `python add_connection.py 10.104.1.1 2 10 in-service false 5101 5201 192950 193050 0 Exp1-ROADM2`
- `python add_connection.py 10.104.1.2 1 10 in-service false 4101 4201 192950 193050 0 Exp1-ROADM1`
- `python add_connection.py 10.104.1.2 2 10 in-service false 5101 5201 192950 193050 0 Exp1-ROADM1`
- `python add_connection.py 10.104.1.3 1 10 in-service false 4102 4201 192950 193050 0 Exp1-FromTor3`
- `python add_connection.py 10.104.1.3 2 10 in-service false 5101 5202 192950 193050 0 Exp1-TorwardTor3`

Demo: Establish Link 2

```
Terminal
native@optical-linux: ~/test/craig_test
native@optical-linux:~/test/craig_test$ sh expl_stepA_add.sh
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
native@optical-linux:~/test/craig_test$

native@srv1-lg1: ~
native@srv1-lg1:~$ ping 192.168.1.2
PING 192.168.1.2 (192.168.1.2) 56(84) bytes of data.
From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
From 192.168.1.1 icmp_seq=2 Destination Host Unreachable
From 192.168.1.1 icmp_seq=3 Destination Host Unreachable
From 192.168.1.1 icmp_seq=4 Destination Host Unreachable
From 192.168.1.1 icmp_seq=5 Destination Host Unreachable
From 192.168.1.1 icmp_seq=6 Destination Host Unreachable
From 192.168.1.1 icmp_seq=7 Destination Host Unreachable
From 192.168.1.1 icmp_seq=8 Destination Host Unreachable
From 192.168.1.1 icmp_seq=9 Destination Host Unreachable
From 192.168.1.1 icmp_seq=10 Destination Host Unreachable
From 192.168.1.1 icmp_seq=11 Destination Host Unreachable
From 192.168.1.1 icmp_seq=12 Destination Host Unreachable
From 192.168.1.1 icmp_seq=13 Destination Host Unreachable
From 192.168.1.1 icmp_seq=14 Destination Host Unreachable
From 192.168.1.1 icmp_seq=15 Destination Host Unreachable
From 192.168.1.1 icmp_seq=16 Destination Host Unreachable
From 192.168.1.1 icmp_seq=17 Destination Host Unreachable
From 192.168.1.1 icmp_seq=18 Destination Host Unreachable
64 bytes from 192.168.1.2: icmp_seq=19 ttl=64 time=1024 ms
64 bytes from 192.168.1.2: icmp_seq=20 ttl=64 time=0.263 ms
64 bytes from 192.168.1.2: icmp_seq=21 ttl=64 time=0.104 ms
64 bytes from 192.168.1.2: icmp_seq=22 ttl=64 time=0.105 ms
64 bytes from 192.168.1.2: icmp_seq=23 ttl=64 time=0.114 ms
64 bytes from 192.168.1.2: icmp_seq=24 ttl=64 time=0.096 ms
64 bytes from 192.168.1.2: icmp_seq=25 ttl=64 time=0.109 ms
64 bytes from 192.168.1.2: icmp_seq=26 ttl=64 time=0.099 ms
64 bytes from 192.168.1.2: icmp_seq=27 ttl=64 time=0.104 ms
^C
--- 192.168.1.2 ping statistics ---
27 packets transmitted, 9 received, +18 errors, 66% packet loss, time 26608ms
rtt min/avg/max/mdev = 0.096/113.911/1024.205/321.837 ms, pipe 4
native@srv1-lg1:~$
```


Demo: Dynamic Switching between Links 1 or Link 2

```
Terminal
native@optical-linux: ~/test/craig_test
Successfully Added Connection
native@optical-linux:~/test/craig_test$ sh exp1_stepA_delete.sh
Successfully Deleted Connection
Successfully Deleted Connection
Successfully Deleted Connection
Successfully Deleted Connection
native@optical-linux:~/test/craig_test$ sh exp1_stepB_add.sh
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
Successfully Added Connection
native@optical-linux:~/test/craig_test$
native@srv1-lg1: ~$
native@srv1-lg1: ~$
```


Summary

COSMOS has **many instruments** to work with:

- SDRs
- GPUs
- FPGAs
- mmWave
- Edge Clouds

COSMOS relies on **optical networking**:

- **Small latencies.**
- Reconfigurable Space Switch to interconnect fibers
- Reconfigurable ROADMs for different optical wavelengths
- SDN Control



City-scale PAWR COSMOS
testbed in West Harlem



COSMOS

Platforms for Advanced
Wireless Research

Thank you!

COSMOS Web-Site: cosmos-lab.org

Tutorial: <https://wiki.cosmos-lab.org/wiki/tutorials/optical-network-example>

“Programmable Optical x-Haul Network in the COSMOS Testbed”