IEEE FNI Testbed WG Meeting

COSMOS: An Open, Programmable, City-scale Testbed for Experimentation with Advanced Wireless Systems and Edge-computing Concepts

WINLAB

Ivan Seskar

March 26, 2021

RUTGERS

Acknowledgements

- GENI GPO
- National Science Foundatio (NSF) and Platforms for Advanced Wireless Research (PAWR) Program
- O-RAN and ONAP Alliances
- Katarzyna Keahey, Abhimanyu Gosain, Joe Breen, Andrew Buffmire, Jonathan Duerig, Kevin Dutt, Eric Eide, Mike Hibler, David Johnson, Sneha Kasera, Earl Lewis, Dustin Maas, Alex Orange, Neal Patwari, Daniel Reading, Robert Ricci, David Schurig, Leigh Stoller, Kobus Van der Merwe, Kirk Webb, Gary Wong, Ismail Guvenc, Brian Floyd, Mihail Sichitiu, Rudra Dutta, Thomas Zajkowski, Jason Proctor, Lavanya Sridharan, Vuk Marojevic, Gerard Hayes, David Love, David W. Matolak, Yufeng Xin, Mike Barts, Asokan Ram, Ilya Baldin, Anita Nikolich, Inder Monga, Jim Griffioen, KC Wang, Tom Lehman, Paul Ruth, Zongming Fei
- … and many other contributors

PAWR platforms were chosen to be geographically diverse and research focus independent

POWDER

Salt Lake City, UT

Software defined networks and massive MIMO

AVAILABLE TODAY !!

COSMOS

West Harlem, NY

Millimeter wave and backhaul research

AVAILABLE TODAY !!

AERPAW

Raleigh, NC

Unmanned aerial vehicles and mobility

Rural Broadband Platform

TBD

Coming late 2020

Colosseum – *World's largest RF emulator, located at Northeastern University in Boston*

AVAILABLE TODAY !!

POWDER Deployment

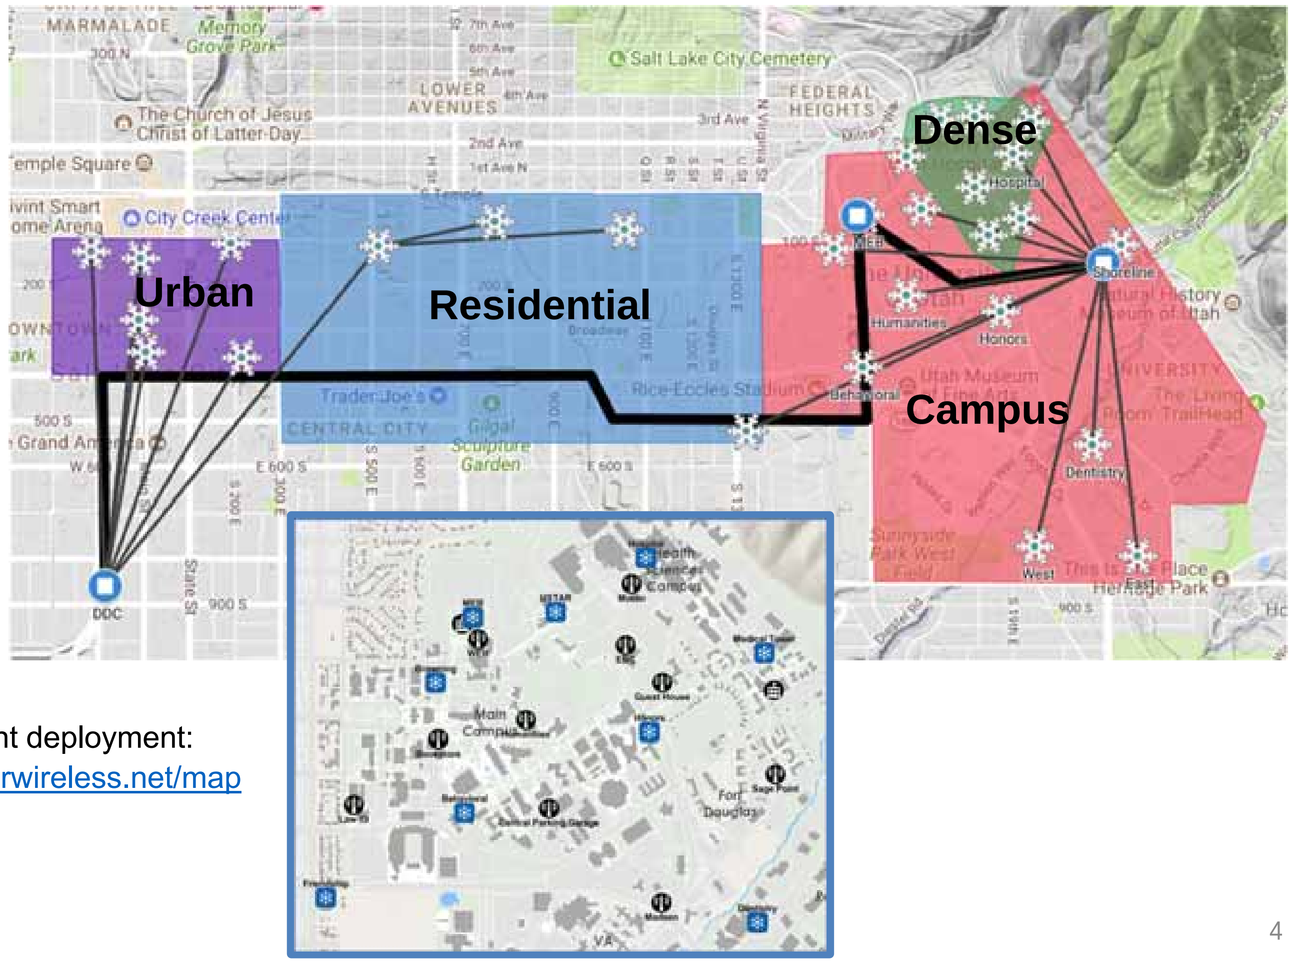

Current deployment: powderwireless.net/map

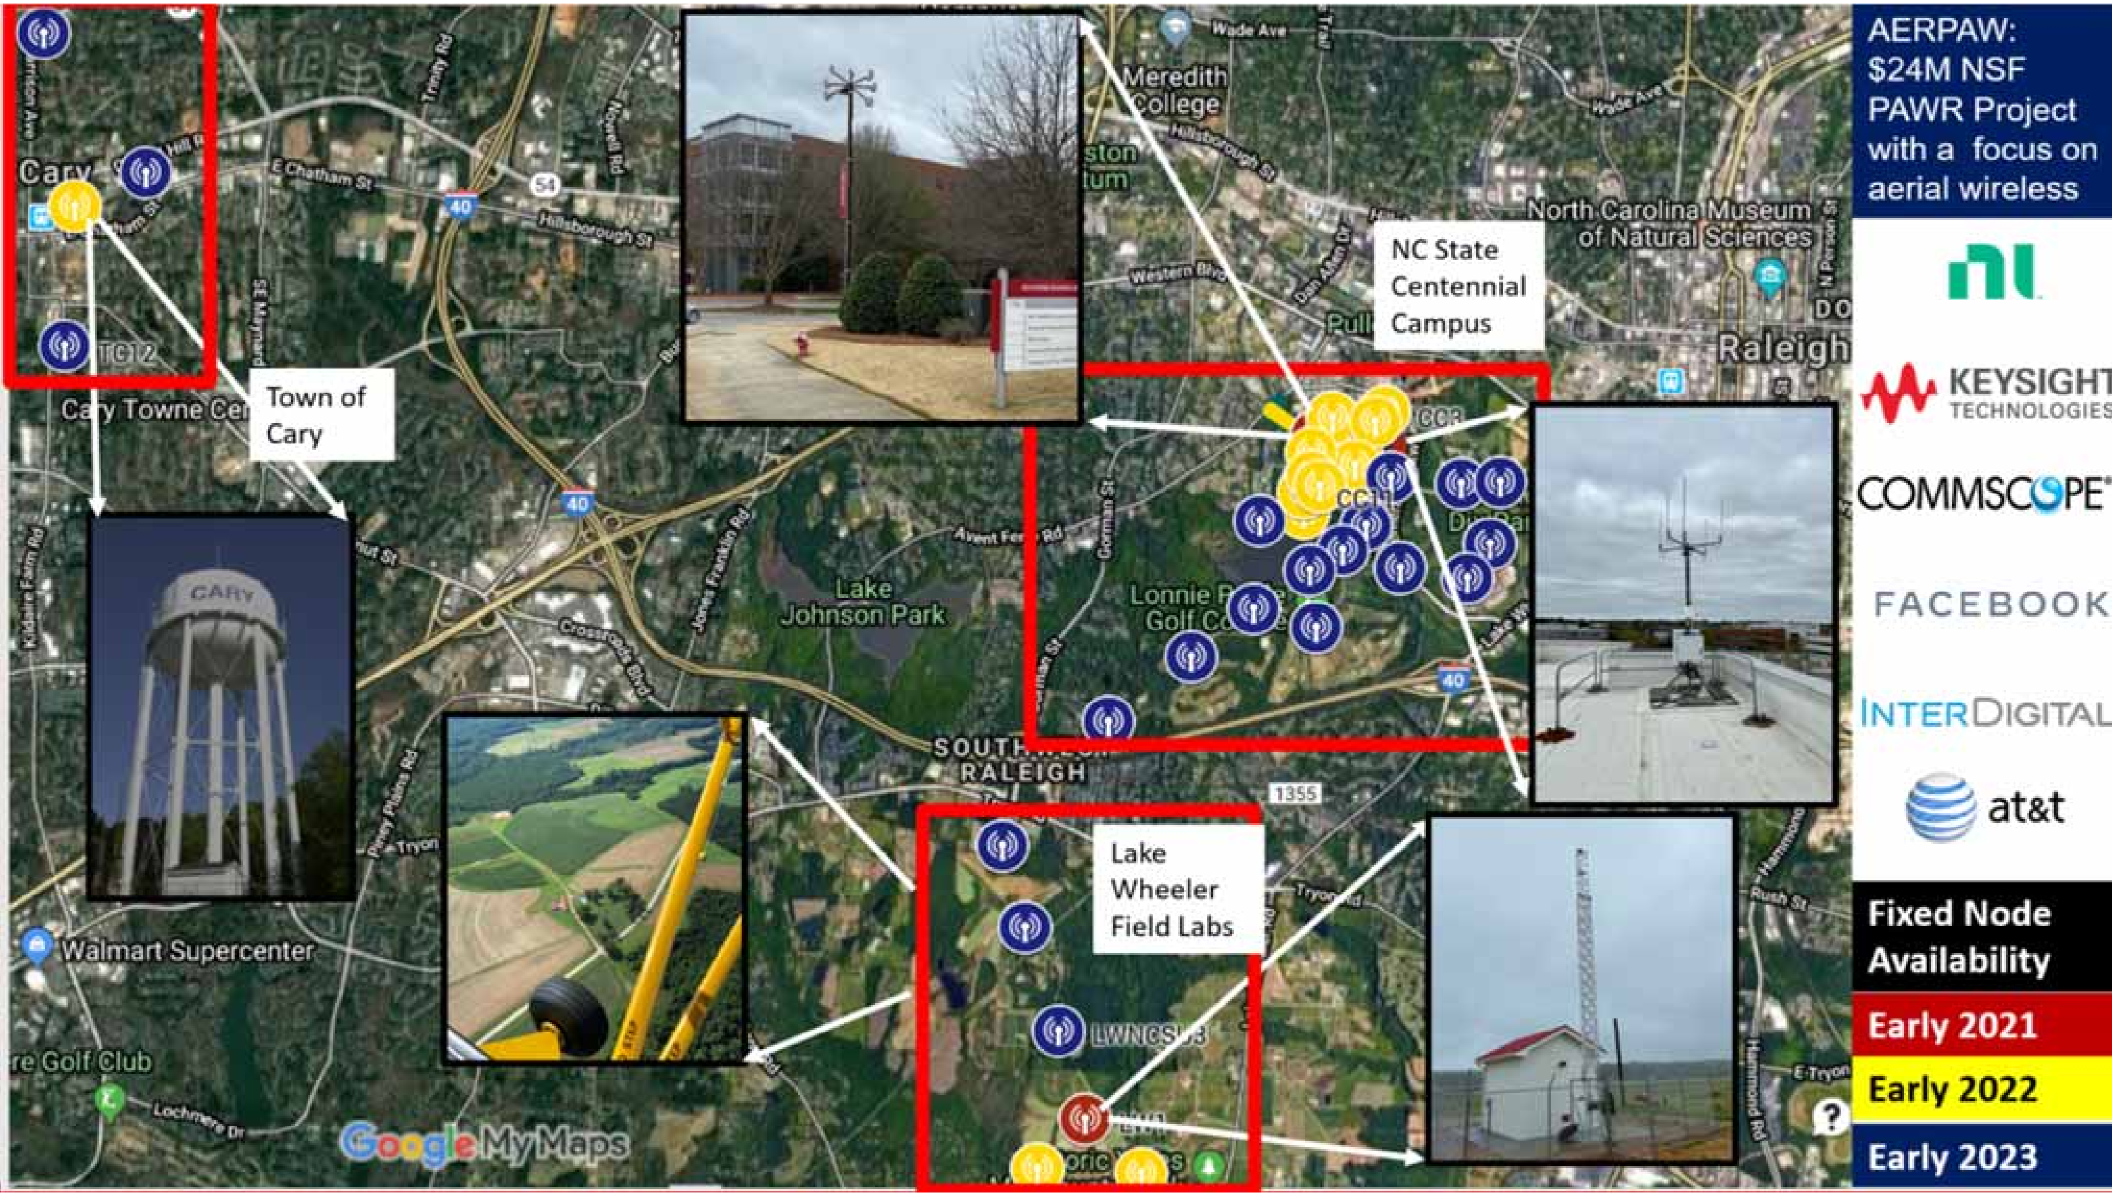

AERPAW: $24M NSF PAWR Project with a focus on aerial wireless
KEYSIGHT TECHNOLOGIES
COMMSCOPE
FACEBOOK
INTERDIGITAL
at&t
Fixed Node Availability
Early 2021
Early 2022
Early 2023
NC State Centennial Campus
Town of Cary
Lake Wheeler Field Labs
Raleigh
Cary
Lake Johnson Park
SOUTH RALEIGH
Meredith College
North Carolina Museum of Natural Sciences
Walmart Supercenter
Google My Maps

Colosseum

Colosseum is the world’s largest wireless network emulator with granularity at the RF signal level

- 256 x 256 100 MHz RF channel emulation
- 128 Programmable Radio Nodes
- Computing resources (CPU, GPU, FPGA)
- Access control and scheduling infrastructure
- Supports remote shared access
- Colosseum supports a containerized software environment with full-stack, end-to-end experiments

northeastern.edu/colosseum

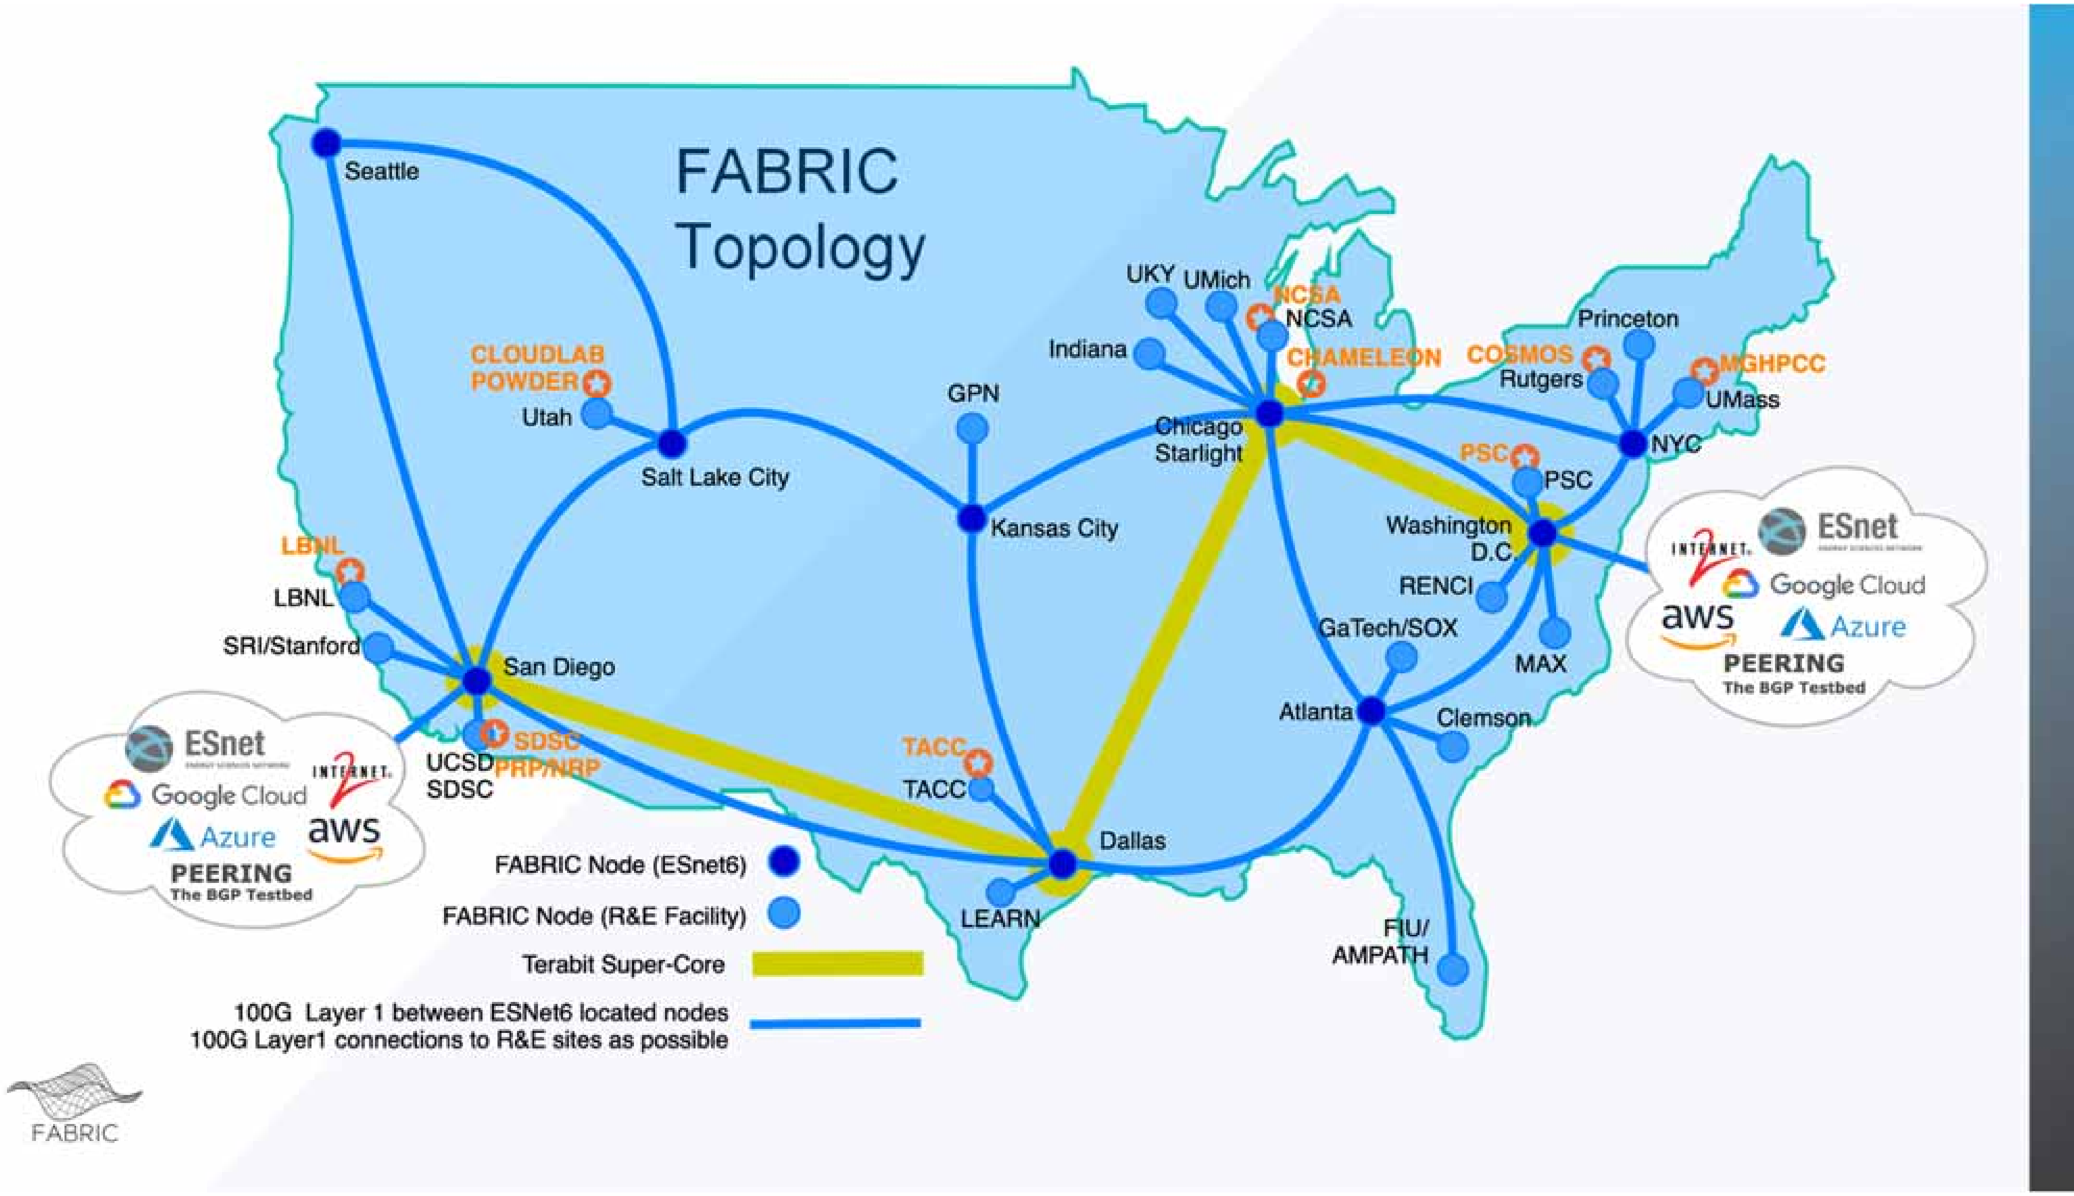
FABRIC
Topology
Seattle
CLOUDLAB
POWDER
Utah
Salt Lake City
GPN
Kansas City
UKY
UMich
NCSA
NCSA
Indiana
CHAMELEON
Chicago
Starlight
COSMOS
Rutgers
Princeton
MGHPCC
UMass
NYC
PSC
PSC
Washington
D.C.
RENCI
MAX
GaTech/SOX
Atlanta
Clemson
FIU/
AMPATH
LBNL
LBNL
SRI/Stanford
San Diego
SDSC
UCSD
PRP/NRP
SDSC
TACC
TACC
Dallas
LEARN
ESnet
INTERNET2
Google Cloud
Azure
aws
PEERING
The BGP Testbed
FABRIC Node (ESnet6)
FABRIC Node (R&E Facility)
Terabit Super-Core
100G Layer 1 between ESNet6 located nodes
100G Layer1 connections to R&E sites as possible
FABRIC

COSMOS

Acknowledgements

- WINLAB, Rutgers University: Dipankar Raychaudhuri, Jakub Kolodziejski, Michael Sherman, Prasanthi Maddala, Nilanjan Paul, Sumit Maheshwari, Newman Wilson, Janice Campanella
- Columbia University: Gil Zussman, Harish Krishnaswamy, Zoran Kostic, Tingjun Chen, Manav Kohli, Jonathan Ostrometzky, Tianwei Deng, Craig Gutterman (EE); Henning Schulzrinne (CS); Sharon Sputz (Data Science Institute); Karen Cheng, Emily Ford (Engineering Outreach); Anthony Avendano (Facilities and Operation); Flores Forbes, Victoria Mason-Ailey (Government and Community Affairs); Alan Crosswell, Daniel Gaitings, Thomas Rom (CUIT)
- New York University: Sundeep Rangan,Thanasis Korakis, Shivendra S Panwar, Panagiotis Skrimponis (ECE); Sheila Borges, Ben Esner (NYU Center for K12 STEM Education);
- University of Arizona: Dan Kilper, Jiakai Yu, Shengxiang Zhu (College of Optical Sciences)
- IBM Research: Xiaoxiong Gu , Arun Paidimarri, Bodhisatwa Sadhu, Alberto Valdes-Garcia
- New York City: Joshua Breitbart (Mayor's Office of the CTO)
- Silicon Harlem: Clayton Banks, Bruce Lincoln
- CCNY: Myung Lee, Rosemarie Wesson
- University of Thessaly: Nikos Makris
- Télécom Paris: Artur Minakhmetov
- … and many other contributors in Rutgers, Columbia, NYU, NYC, Silicon Harlem, U. Arizona, CCNY, and IBM Research.

Future Wireless Networks

Researchers' objective: design, prototype, and evaluate technologies for the wireless edge to enable novel modes of interaction between city residents and the urban environment

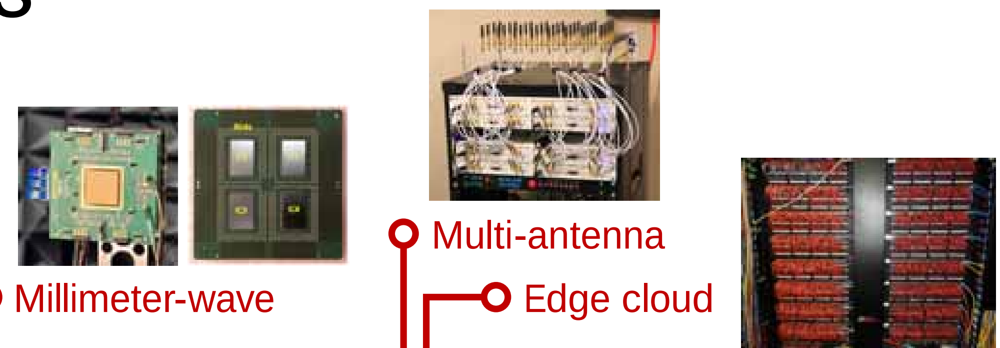

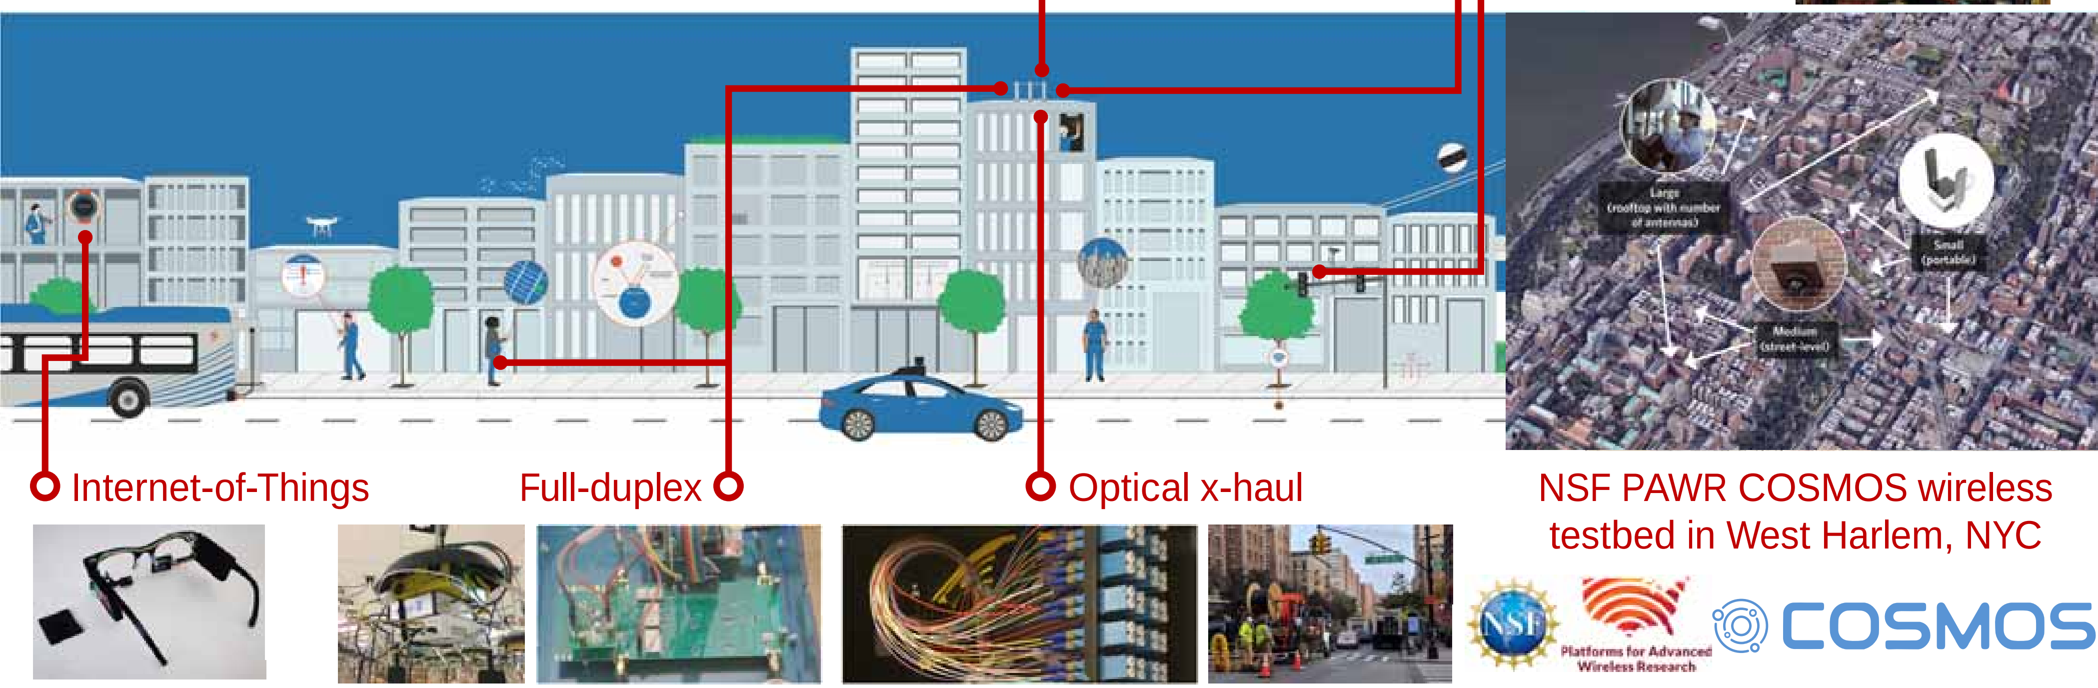

COSMOS: Project Vision

- Latency and compute power are two important dimensions and metrics
- Edge computing can enable real-time applications
- Objective: Real-world investigation or urban environments with
  - Ultra-high bandwidth (~Gbps)
  - Low latency (<5 ms)
  - Bandwidth densities (~10 Tbps/km^2)
- Enablers:
  - 10s of 64-element millimeter-wave arrays
  - 10s of miles of Manhattan dark fiber
  - B5G edge cloud base stations
  - Programmability

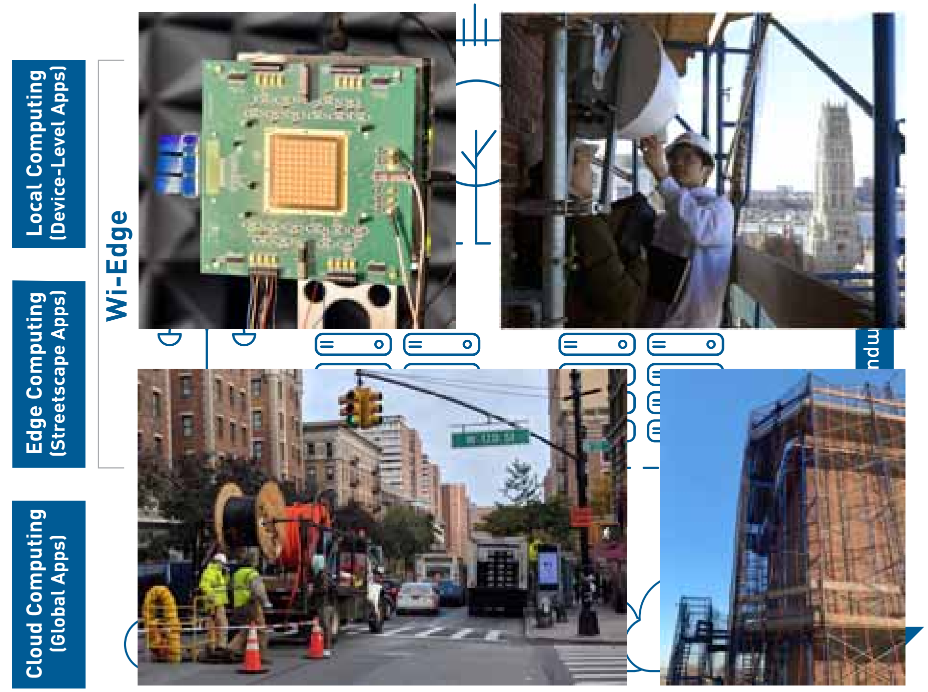

Ultra-high bandwidth, low latency, and powerful edge computing will enable new classes of real-time applications. Domains including AR/VR, connected cars, smart city (with high-bandwidth sensing), and industrial control

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC IBM

Objective: Take it Outside

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC The City College of New York IBM

Objective: Take it Outside

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC The City College of New York IBM

COSMOS: Envisioned Deployment

- West Harlem with an area of ~1 sq. mile
  - ~15 city blocks and ~5 city avenues

- ~9 Large sites
  - Rooftop base stations
- ~40 Medium sites
  - Building side- or lightpole-mounted
- ~200 Small nodes
  - Including vehicular and hand-held

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC IBM

COSMOS: Envisioned Deployment

- A phased approach:
  - May 2019: Pilot completion
  - Sept. 2019: FCC Innovation Zone
  - June 2020: General Available
  - During 2021*: Phase 1 completion

    *Deployments affected by the COVID-19 pandemic

- Fiber optic connection from most sites
- Fiber connection to Rutgers, NYU Data Center (at 32 Ave. of Americas), GENI, and Internet2, etc.

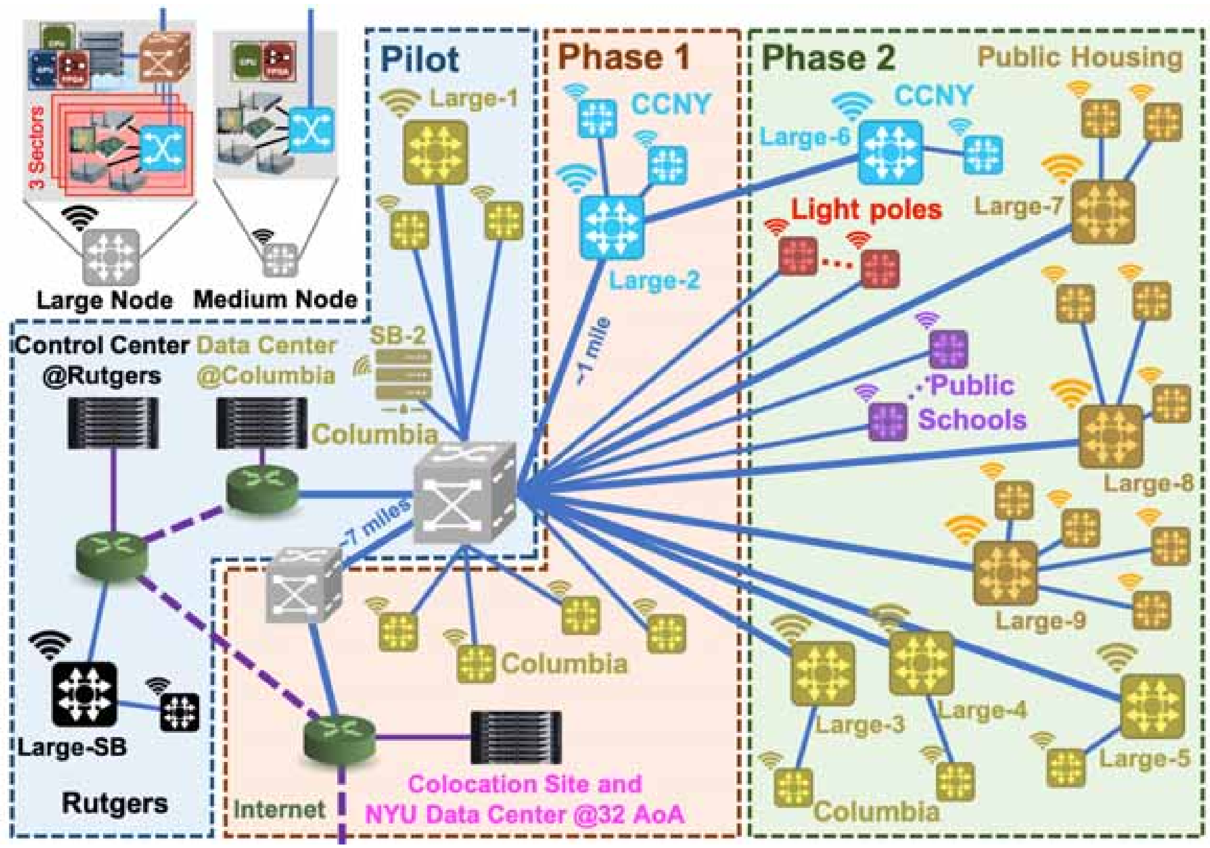

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC IBM

COSMOS (PAWR) Experimental Licenses

FCC Innovation Zone: “The New York City Innovation Zone will encompass area bounded by W 123rd Street on the south, Amsterdam Avenue to the east, W 134th Street to the north and Broadway to the west”

Frequency Band	Type of operation	Allocation	Maximum EIRP (dBm)
2500-2690 MHz	Fixed	Non-federal	20*
3700-4200 MHz	Mobile	Non-federal	20*
5850-5925 MHz	Mobile	Shared	20*
5925-7125 MHz	Fixed & Mobile	Non-federal	20*
27.5-28.35 GHz	Fixed	Non-federal	20*
38.6-40.0 GHz	Fixed	Non-federal	20*

Program Experimental License: at Rutgers, Columbia and CCNY campuses

COSMOS: Envisioned Deployment

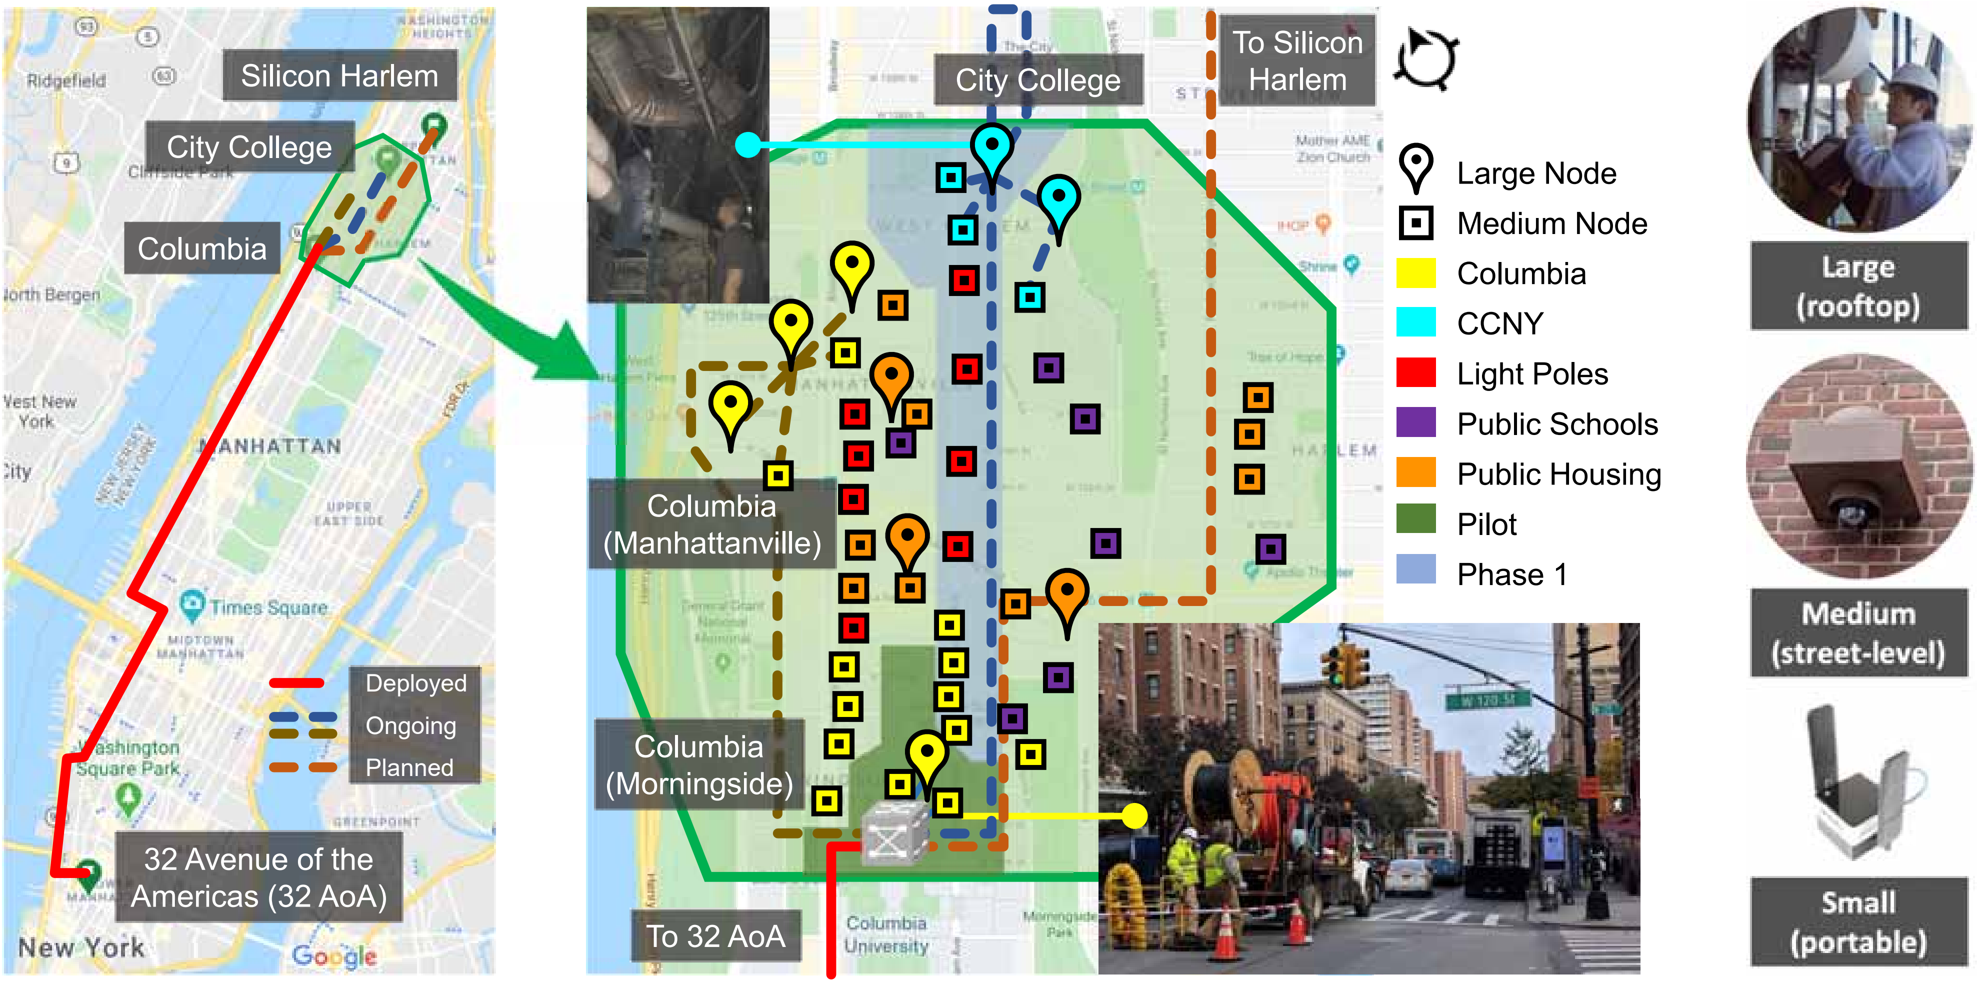

COSMOS: Design and Architecture

- Key design challenge: Gbps+ performance and full programmability at the radio level
  - Fully programmable multi-layered computing architecture for flexible experimentation
- Key technologies:
  - Software-define radios (SDRs)
  - Millimeter-wave (mmWave) radios
  - Optical x-haul networks
  - Software-defined networking and cloud
  - Control and management software
- Pilot experiments:
  - Open-access full-duplex wireless 1 2
  - Optical-wireless x-haul networking 3 4

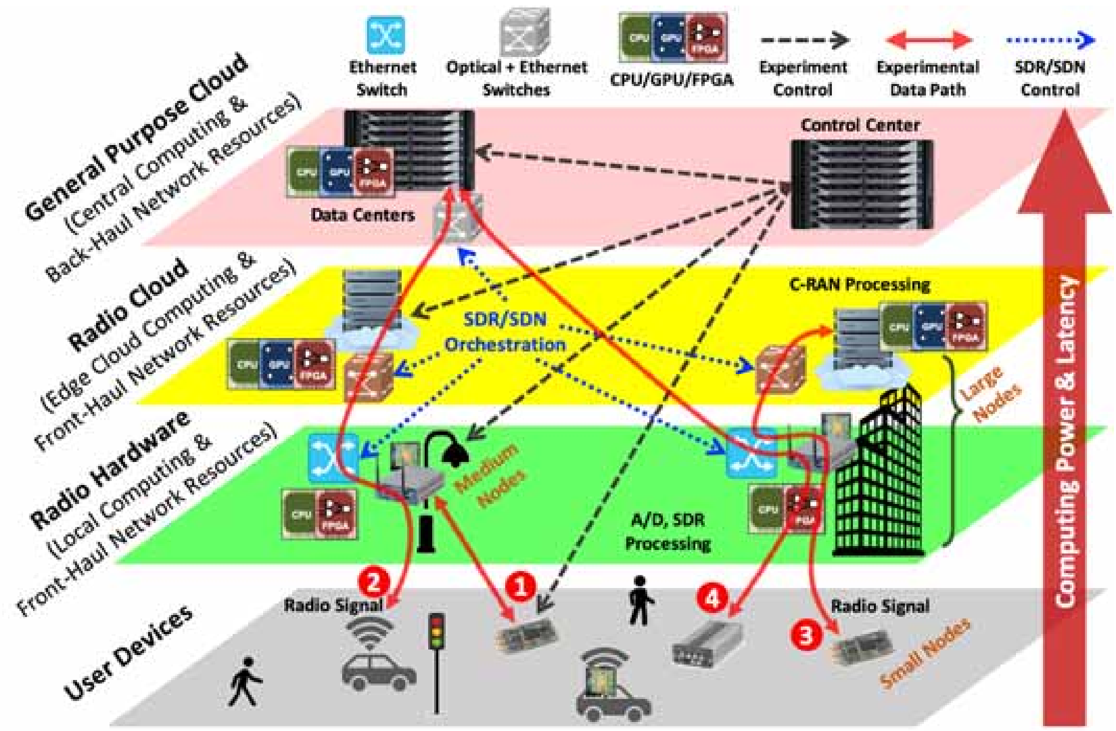

COSMOS's multi-layered computing architecture

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC IBM

Key Technology: Software-Define Radios

- Software-defined radio (SDR) nodes at various performance levels and form factors:
  - (i) 0.4–6 GHz bands, (ii) 28/60 GHz bands (with up to ~500 MHz bw), Gbps
- Signal processing can be spread between radio node & edge cloud RAN

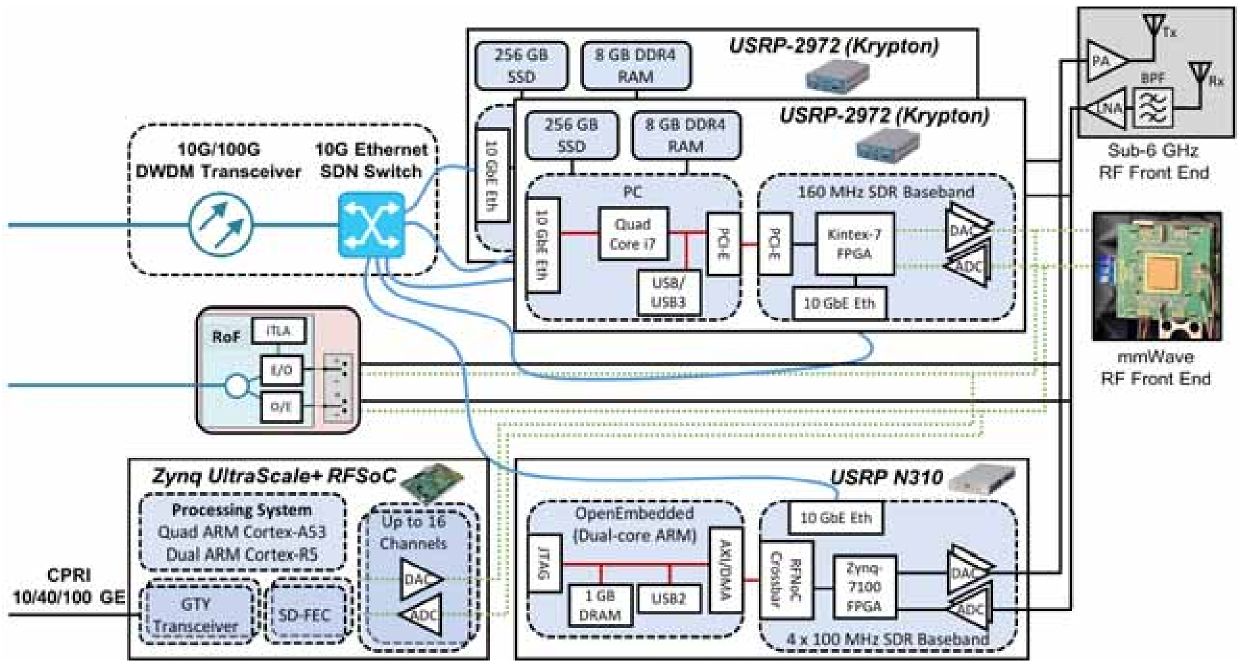

Key Technology: Software-Define Radios

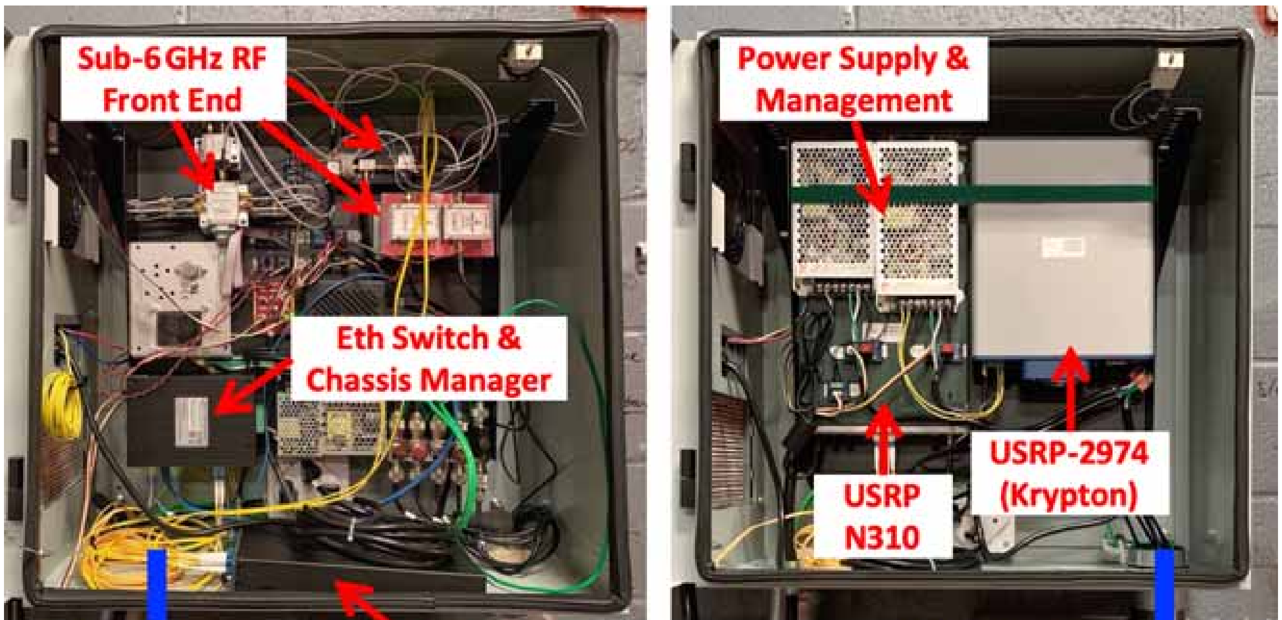

A large node sector or a medium node

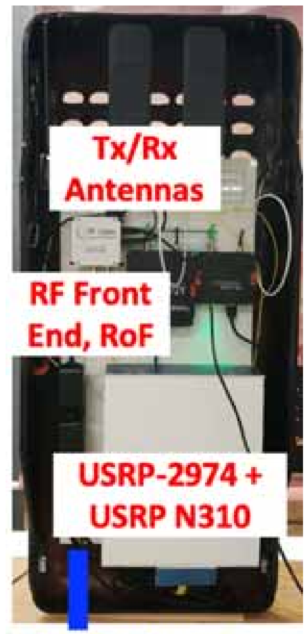

Medium-light node (lightpole-mounted)

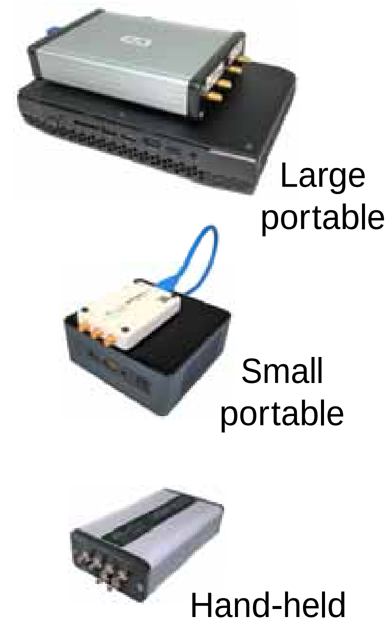

Small mobile node

COSMOS RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC <sh> The City College of New York IBM

Key Technology: mmWave

- IBM 28 GHz programmable phased array antenna modules (PAAMs)
  - Integration of the β-version in Sandbox 2
  - Up to ~500 MHz BW using the Zynq UltraScale+ RFSoC platform
  - Experiment with adaptive beamforming and MIMO

USRP-2974 (Krypton)

Zynq RFSoC

- Facebook Terragraph 60 GHz radios

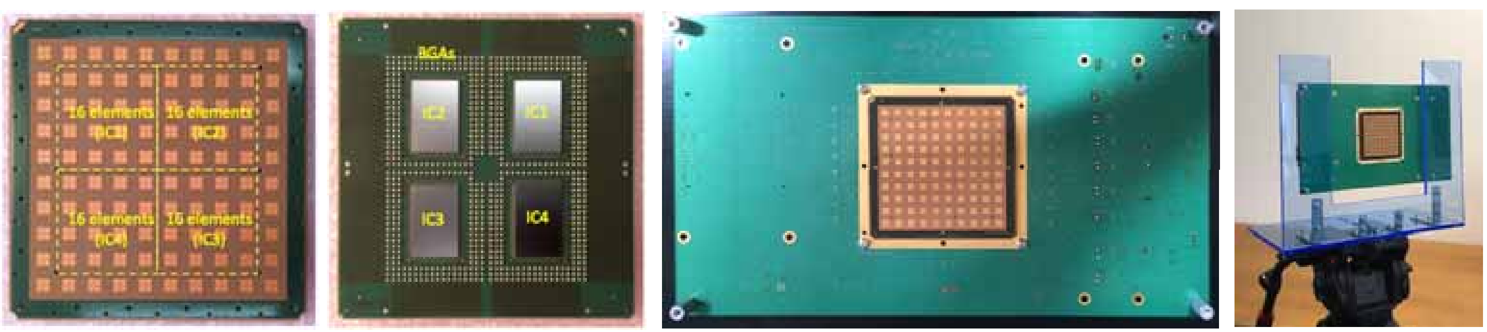

IBM 28 GHz phased array antenna module
(in collaboration with Ericsson)

Facebook Terragraph
60 GHz radios

Key Technology: mmWave

Sivers IMA 60 GHz Radios

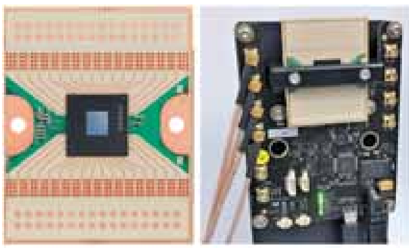

Zynq UltraScale+ RFSoC

InterDigital EdgeLink (60 GHz)

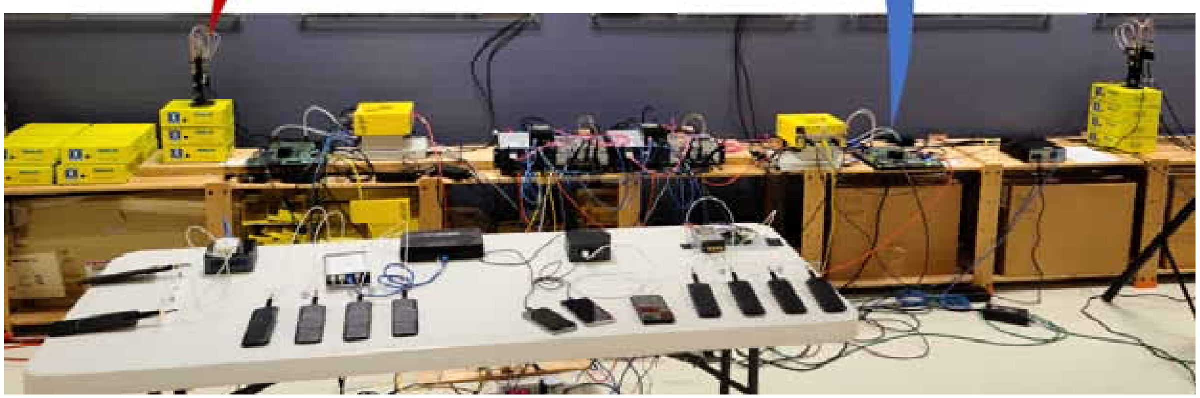

Sivers IMA Radios (60 GHz) with RFSoC

InterDigital 5G NR Platform (28 GHz)

Key Technology: Optical Networking

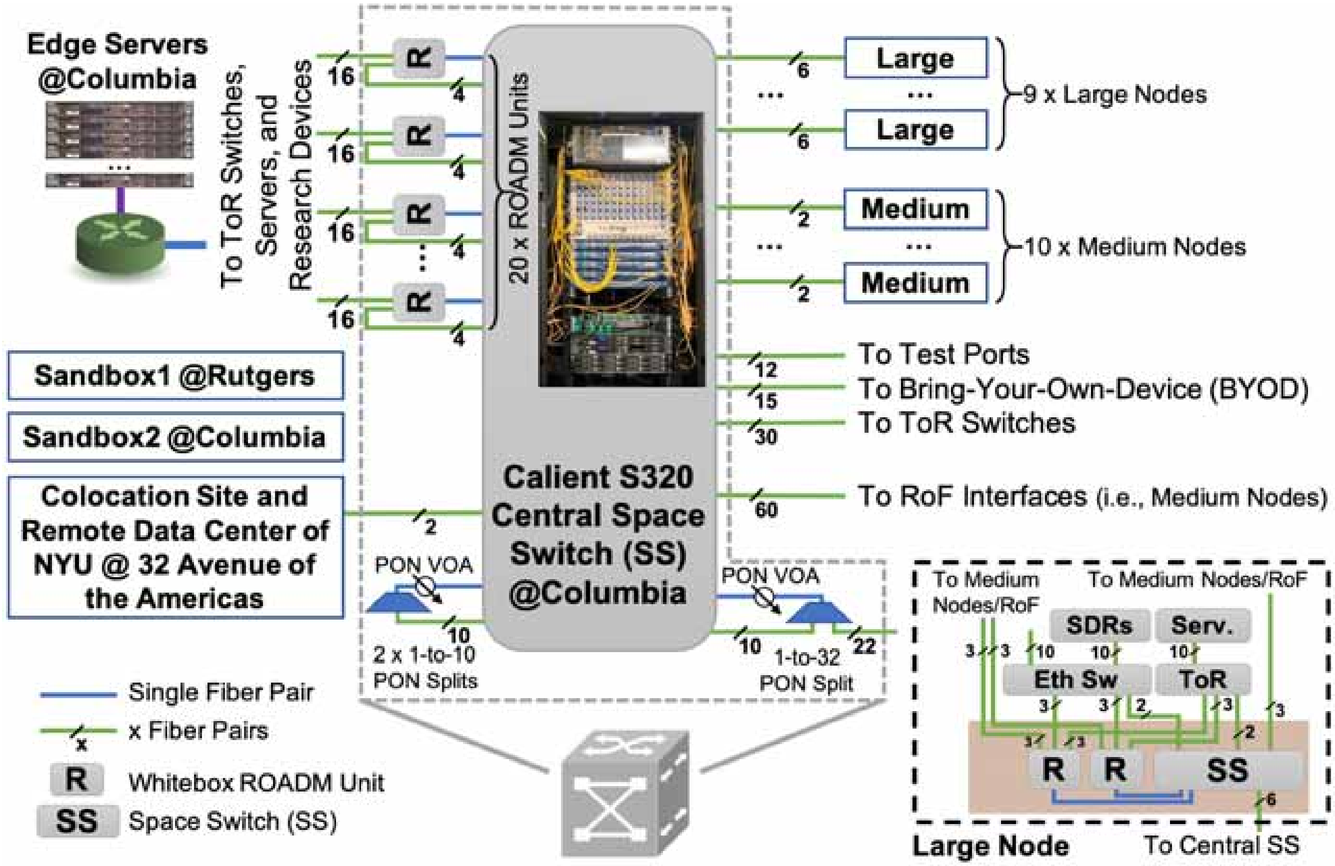

- J. Yu, C. Gutterman, A. Minakhmetov, M. Sherman, T. Chen, S. Zhu, G. Zussman, I. Seskar, and D. Kilper, “Dual use SDN controller for management and experimentation in a field deployed testbed,” in *Proc. OSA OFC’20, T3J.3*, 2020.
- J. Yu, T. Chen, C. Gutterman, S. Zhu, G. Zussman, I. Seskar, and D. Kilper, “COSMOS: Optical architecture and prototyping,” in *Proc. OSA OFC’19, M3G.3 (invited)*, 2019.

Key Technology: Cloud Architecture

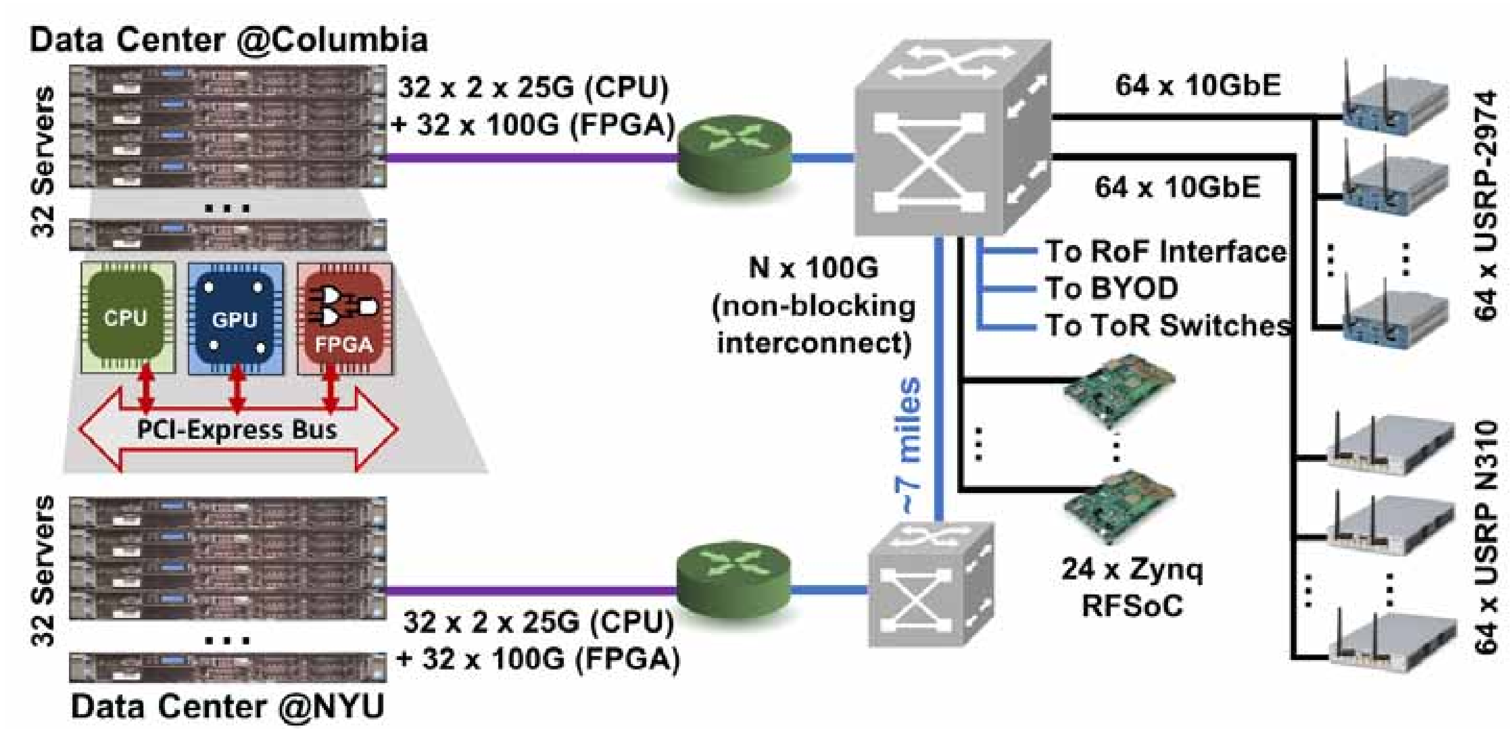

- COSMOS serves as an Open Test and Integration Center during the O-RAN North America plugfest and proof-of-concept demonstration

Pilot Experiment: Full-Duplex Wireless

- Open-access wideband full-duplex radios integrated in the COSMOS sandbox2 with open-sourced hardware, software, and example experiments

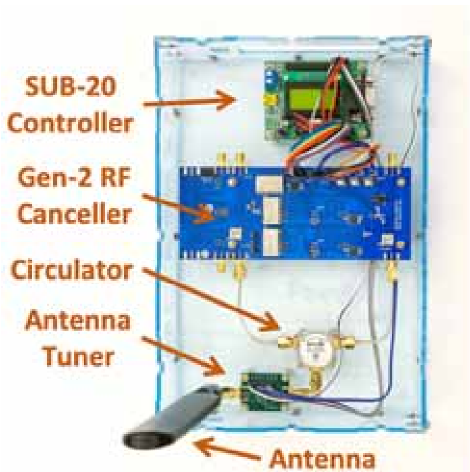

Real-time visualization of data packets, signal spectrum, and digital self-interference cancellation

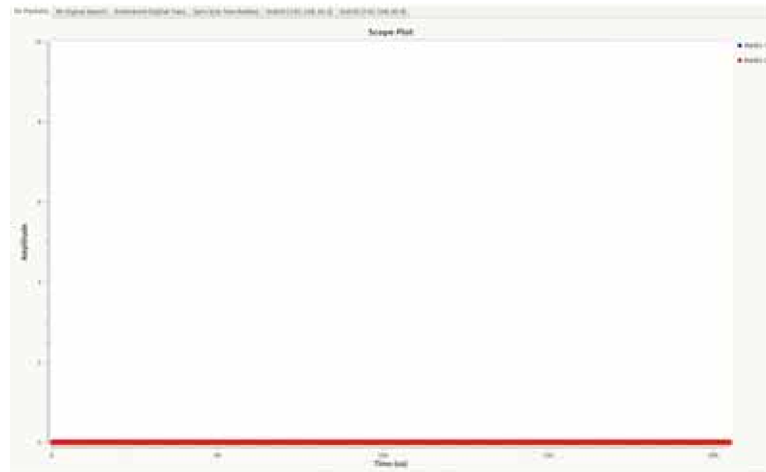

- M. Kohli, T. Chen, M. Baraani Dastjerdi, J. Welles, I. Seskar, H. Krishnaswamy, and G. Zussman, “Open-access full-duplex wireless in the ORBIT and COSMOS testbeds,” in *Proc. ACM MobiCom’20 Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH)*, 2020.
- T. Chen, M. Baraani Dastjerdi, J. Zhou, H. Krishnaswamy, and G. Zussman, “Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation,” in *Proc. ACM MobiCom’19*, 2019. **ACM MobiCom’19 Student Research Competition (SRC) Winner – First Place.**
- Tutorial available at https://wiki.cosmos-lab.org/wiki/tutorials/full_duplex, code available at https://github.com/Wimnet/flexicon_orbit.

Pilot Experiment: Remote-Processing

- Full-duplex radio integrated with COSMOS' dark fiber-based optical x-haul network

- Local RF self-interference cancellation at the full-duplex radio
- Remote digital self-interference cancellation at the server (~14 miles away from the radio)

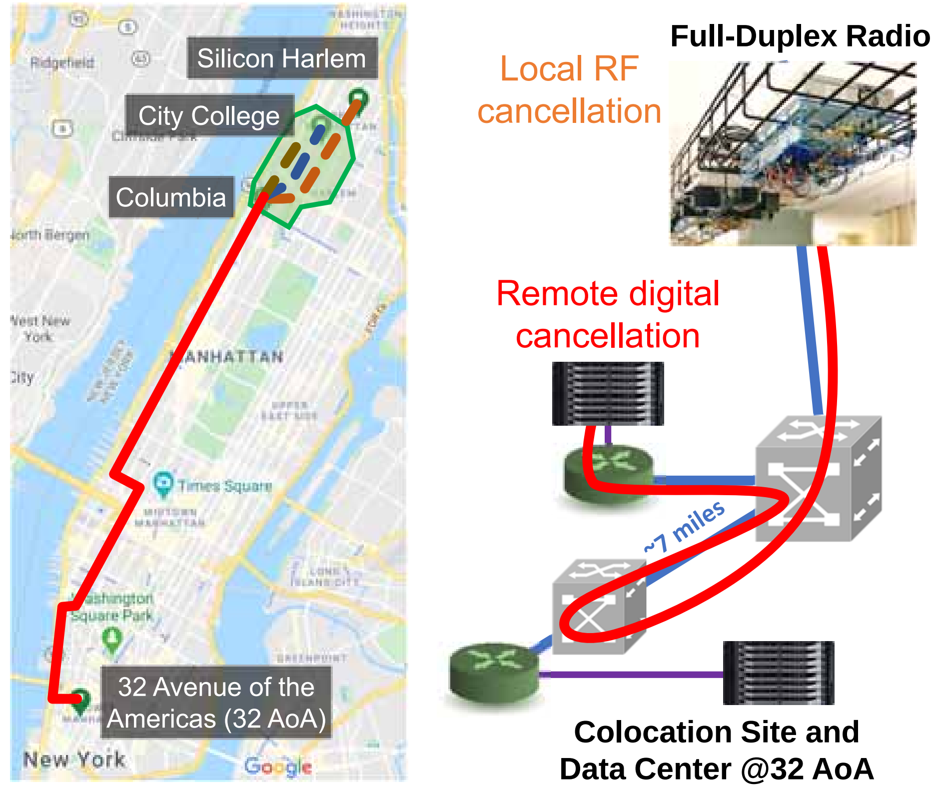

COSMOS dark fiber deployment and the supported Cloud-RAN applications

- J. Yu, T. Chen, C. Gutterman, S. Zhu, G. Zussman, I. Seskar, and D. Kilper, "COSMOS: Optical architecture and prototyping," in *Proc. OSA OFC'19, M3G.3 (invited)*, 2019.

Pilot Experiment: Wireless Handover

- SDN-based optical switching to support high bandwidth links with deterministic delay
- A vehicle taking a turn at an intersection receives services from two remote radio heads (RRHs) through dynamic optical switching and wavelength re-allocation.

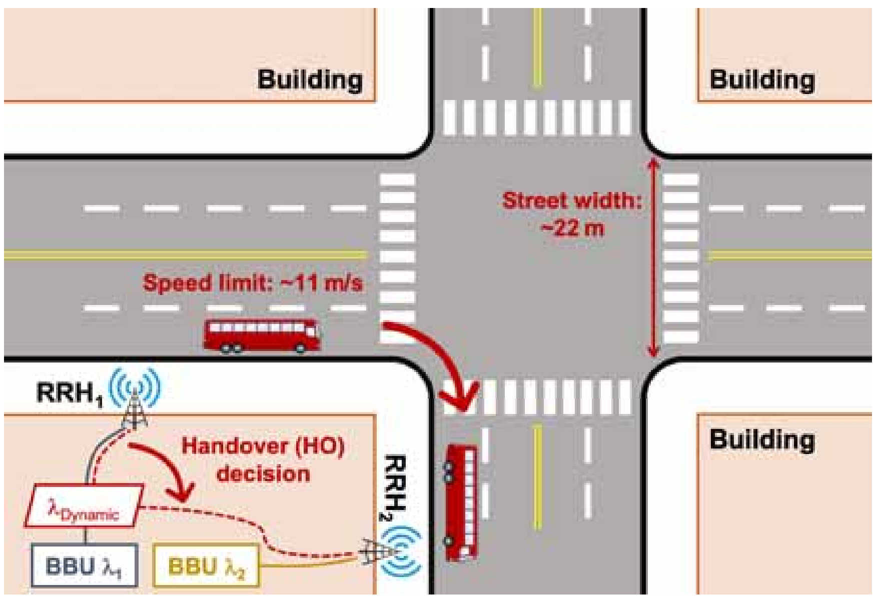

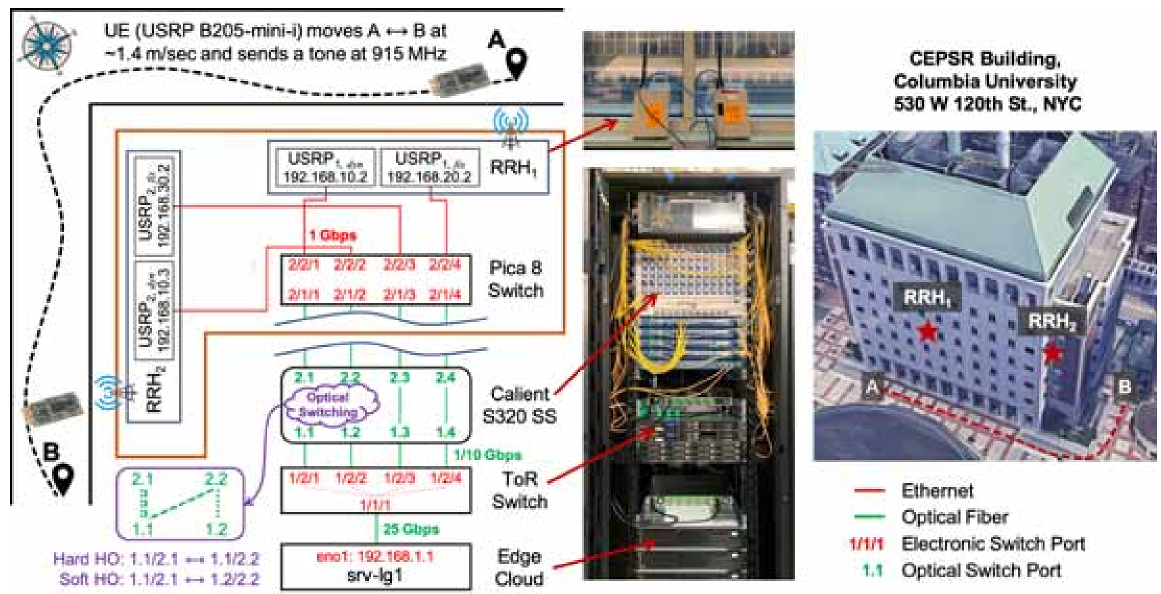

- A. Minakhmetov, C. Gutterman, T. Chen, J. Yu, C. Ware, L. Iannone, D. Kilper, and G. Zussman, “Experiments on cloud-RAN wireless handover using optical switching in a dense urban testbed,” in *Proc. OSA OFC’20, Th2A.25*, 2020.

Pilot Experiment: Smart Intersection

- Smart intersection as a core smart city asset
  - Low latency, high bandwidth wireless links, sensor data acquisition
  - Edge cloud computing and machine intelligence for interaction with pedestrians
- Real-time (latency) – something useful for traffic interaction/management
  - Vehicle speed: 10 km/h → ~3 m/s → ~0.1 m in 1 frame of a video (@30 fps)
  - Arguably useful to prevent accidents, target round-trip latency = 1/30 second

COSMOS pilot site

Bird's eye videos fed into the COSMOS edge node for vehicles/pedestrians detection and classification

- S. Yang, E. Bailey, Z. Yang, J. Ostrometzky, G. Zussman, I. Seskar, and Z. Kostic, "COSMOS smart intersection: Edge compute and communications for bird's eye object tracking," in *Proc. 4th International Workshop on Smart Edge Computing and Networking (SmartEdge'20)*, 2020.

Autonomous Vehicles in a Metropolis?

Manhattan

City roads are complicated

Some important data is not accessible to individual vehicle's sensors.

Pedestrians are "unruly".

Can infrastructure "help"?

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC <sh> The City College of New York

Smart Intersection - Research Goals

Detect and track vehicles and pedestrians

Using COSMOS resources

- **Infrastructural sensors: cameras, IoT**
- **Wireless data harvesting**
- **Data aggregation in COSMOS nodes**
- **Computing and AI intelligence of nodes**
- **Feedback to participants in the intersection**

Harvesting data from participants inside the traffic intersection

- **In-vehicle sensors: cameras, lidars, GPS, other**
- **Pedestrian phones**

Provide feedback in real time

Evaluate performance

Explore algorithms: improve accuracy and real time performance.

Smart Intersection Objectives

- Detect and track vehicles and pedestrians
- Provide feedback to participants in the intersection in real time

What is real-time (latency)?

- Something useful for traffic interaction/management

10 kilometers per hour (km/h)

- 6.2 miles per hour
- 2.778 meters per second
- ~ 3 m/s

How far does a vehicle move in (1/30) of a second (1 frame of a video), at 10km/h

- 3m/s / 30 = 0.1m
- Arguably useful to prevent accidents

Target round-trip latency = (1/30) s

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC <sh> The City College of New York

Smart Intersection - Experimentation

COSMOS pilot site:
Amsterdam Avenue and 120th St.

Infrastructure support:
radios, cameras, edge computing node (CPUs, GPUs and FPGAs).

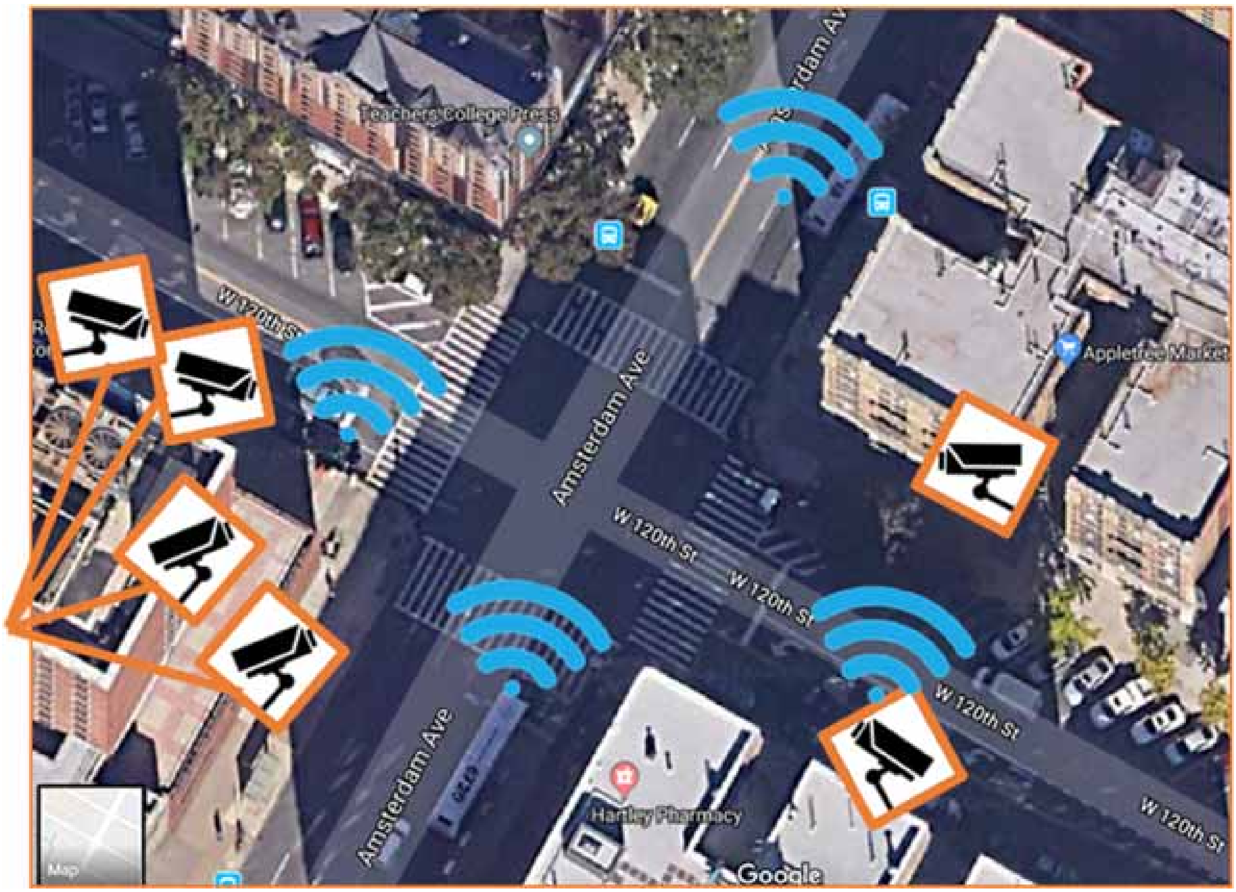

Vehicle/Pedestrian Detection & Tracking

Set of cameras in COSMOS nodes:

- **Street level (medium nodes)**
- **Rooftop level (large node)**

Bird's eye videos
(friendly for privacy and security - IRB issue)

- **Different light and weather conditions**
- **Calibrated to 90 degree bird's eye view**
- **Manually annotate for training**

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC <sh> The City College of New York

Processing Flow

Evaluating the quality and compute latencies of detection and tracking:

- Custom DL models
- DL training and inference

Target: From bird's eye video to radar map broadcast, using deep learning (DL).

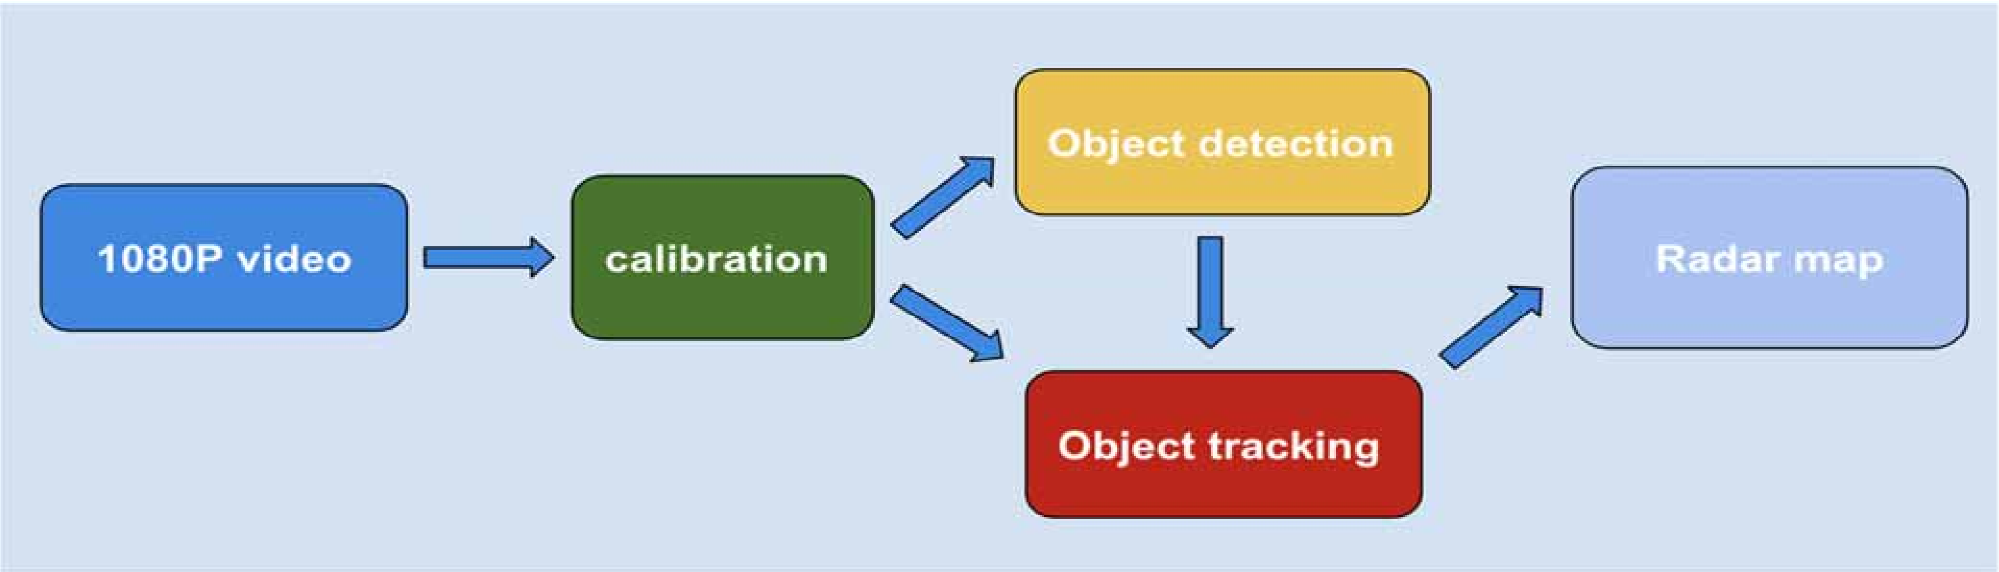

Detection Example

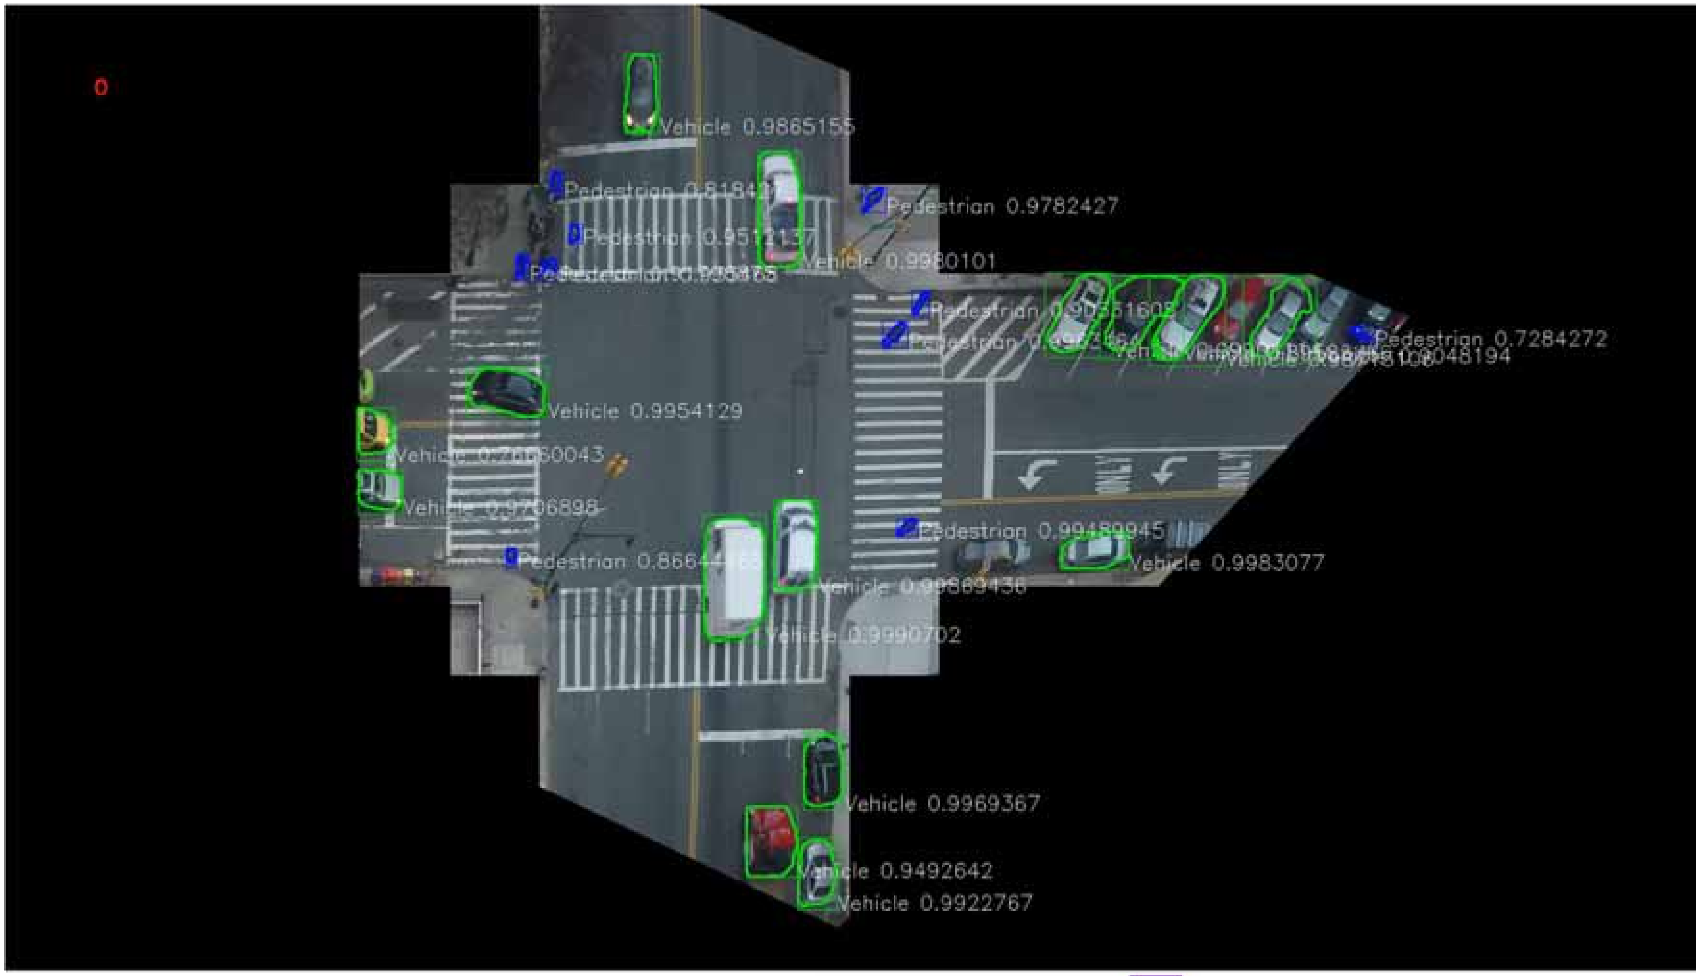

Tracking Example

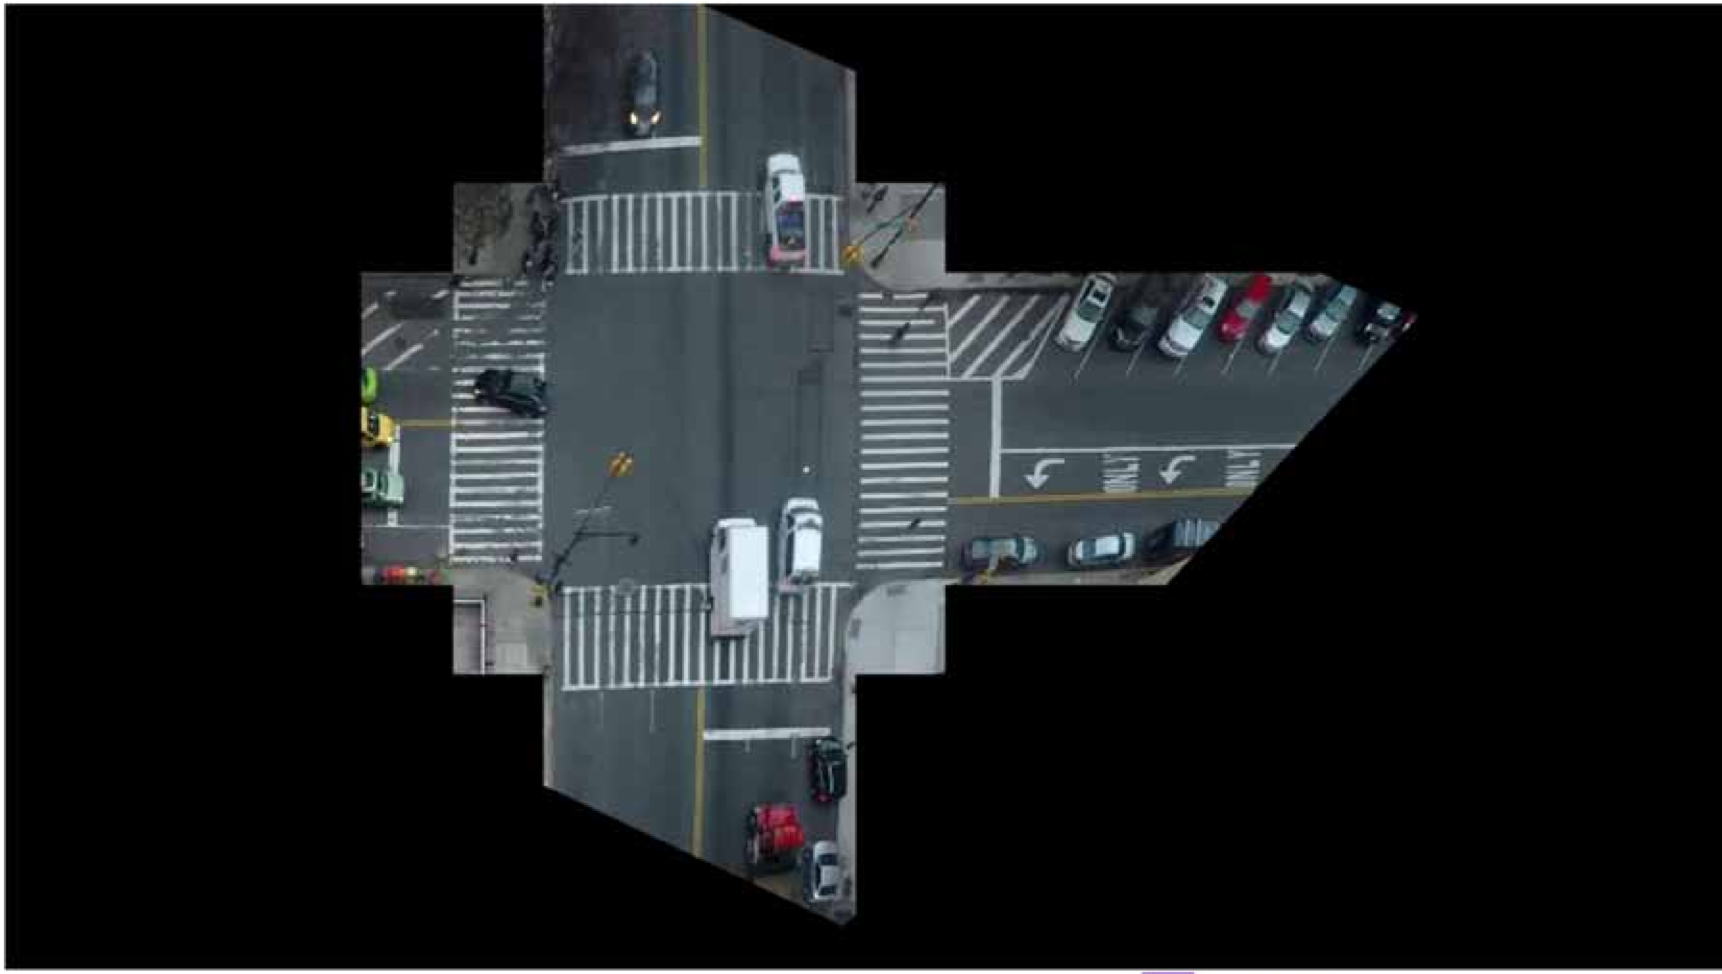

Radar Map

Broadcast to all:

From “Central” computing node to:

- all Vehicles
- all Pedestrians

...in real time.

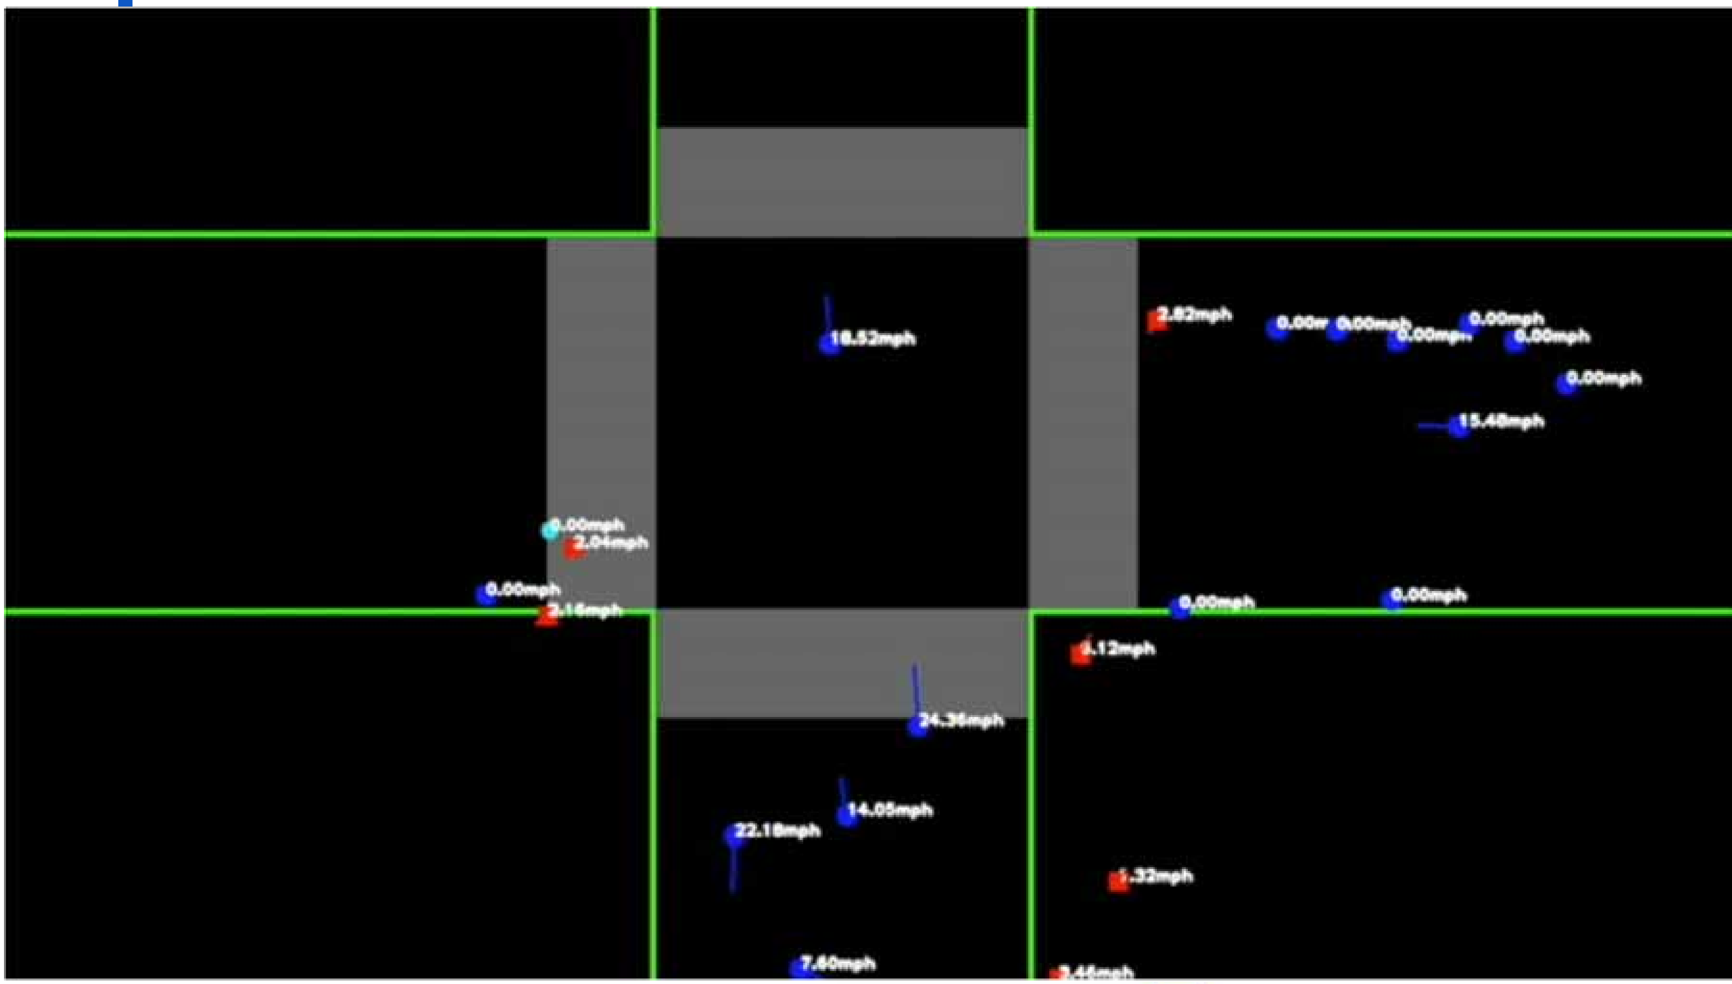

Detection & Tracking Evaluation

Significant reduction in computing time are needed - both model-wise and optimization-wise.

Detection & Tracking Methods

Object detection	Object tracking
Mask RCNN	Deep Affinity Network (DAN)
Single-Shot Detector (SSD)	Deep Simple Online and Realtime Tracking (DeepSORT)
You Only Look Once v3 (YOLOv3)	PWC net + Multicut

Performance: Detection Accuracies

Class	Detector	AP_{50}	$mIOU_{50}$	AP_{75}	$mIOU_{75}$
Veh.	MRCNN	90.7	77.3	73.3	84.4
	SSD	89.7	71.3	58.5	81
	YOLOv3	86.2	69.3	47.7	80.1
Ped.	MRCNN	62.7	65.6	37.3	80.3
	SSD	48.6	59.1	7.8	72
	YOLOv3	50.8	58.2	12.9	71.4

Performance: Frame-per-second (Nvidia P100 GPU)

Detector	FPS	Tracker	FPS
MRCNN	1	DAN	2
SSD	9.2	DSORT	3.2
YOLOv3	11.8	MCUT	0.47

Education and Outreach

- COSMOS education toolkit: A small pre-configured COSMOS node (developed in Summers 2018–2020 RET programs) offering 100+ K–12 educational labs in Math/Science/CS
- Numerous education and outreach activities

5G COVET

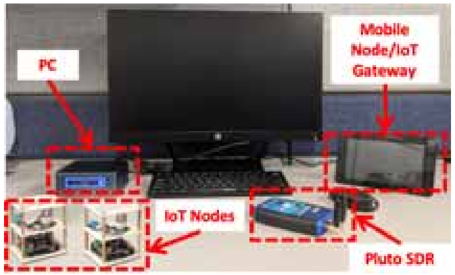

COSMOS education toolkit

COSMOS Research Experiences for Teachers (RET) program

Columbia Girls' Science Day

Kids Week at the Intrepid Museum

- P. Skrimponis, N. Makris, K. Cheng, J. Ostrometzky, Z. Kostic, G. Zussman, T. Korakis, and S. Borges Rajguru, "Evaluation: A teacher professional development program using wireless communications and NGSS to enhance STEM teaching & learning," in *Proc. ASEE Annual Conference*, 2020.
- "5G in 3D: Immersive COSMOS Education Toolkit Wins Verizon EdTech Challenge", https://www.engineering.columbia.edu/news/verizon-edtech-challenge

WINLAB Summer Internship Open House

O-RAN/ONAP (Candidate) OTIC

- Used for ONAP PoCs since 2018
- Environment
  - Entry point at console located at console.sb10.orbit-lab.org
  - 11 Ubuntu servers, all managed by OpenStack
  - Control Node and Compute Nodes
- Access Methods
  - Organizations can gain access by requesting an account – details at: https://wiki.onap.org/pages/viewpage.action?pageId=45298557
  - Tunnels to other testing and integration labs around the world

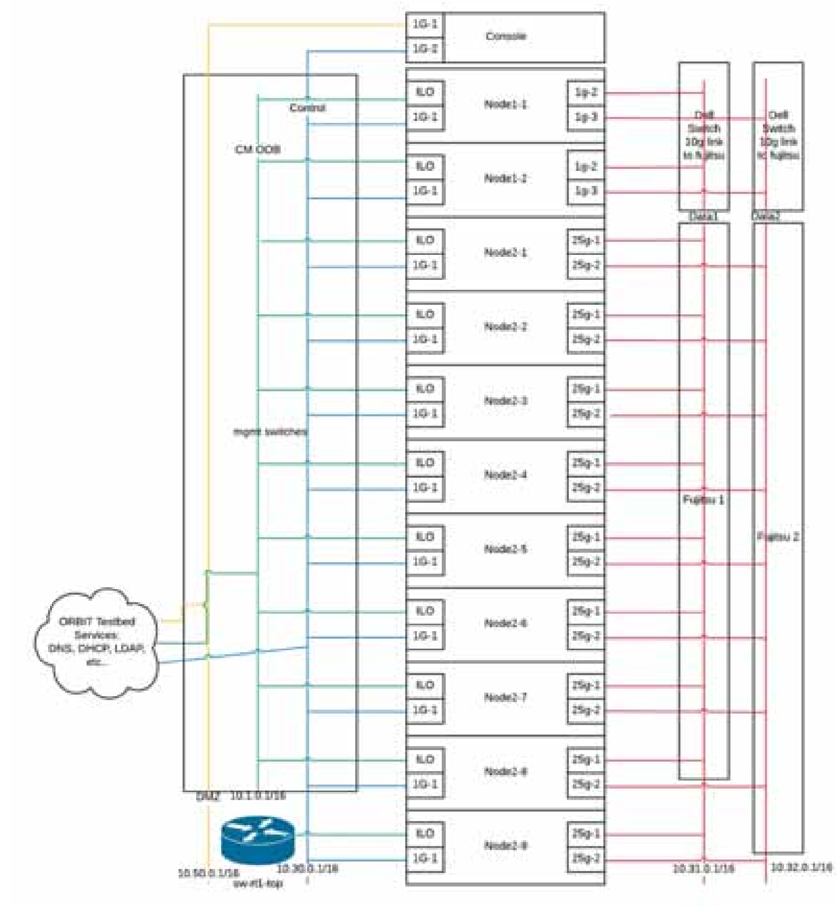

COSMOS

RUTGERS · COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK · NYU · NYC · <sh> · The City College of New York

“Typical” Wireless Experiment: O-RAN + ONAP

- **Compute Intensive**
- **Distributed**
- **Latency limited**
- **Heavily virtualized**

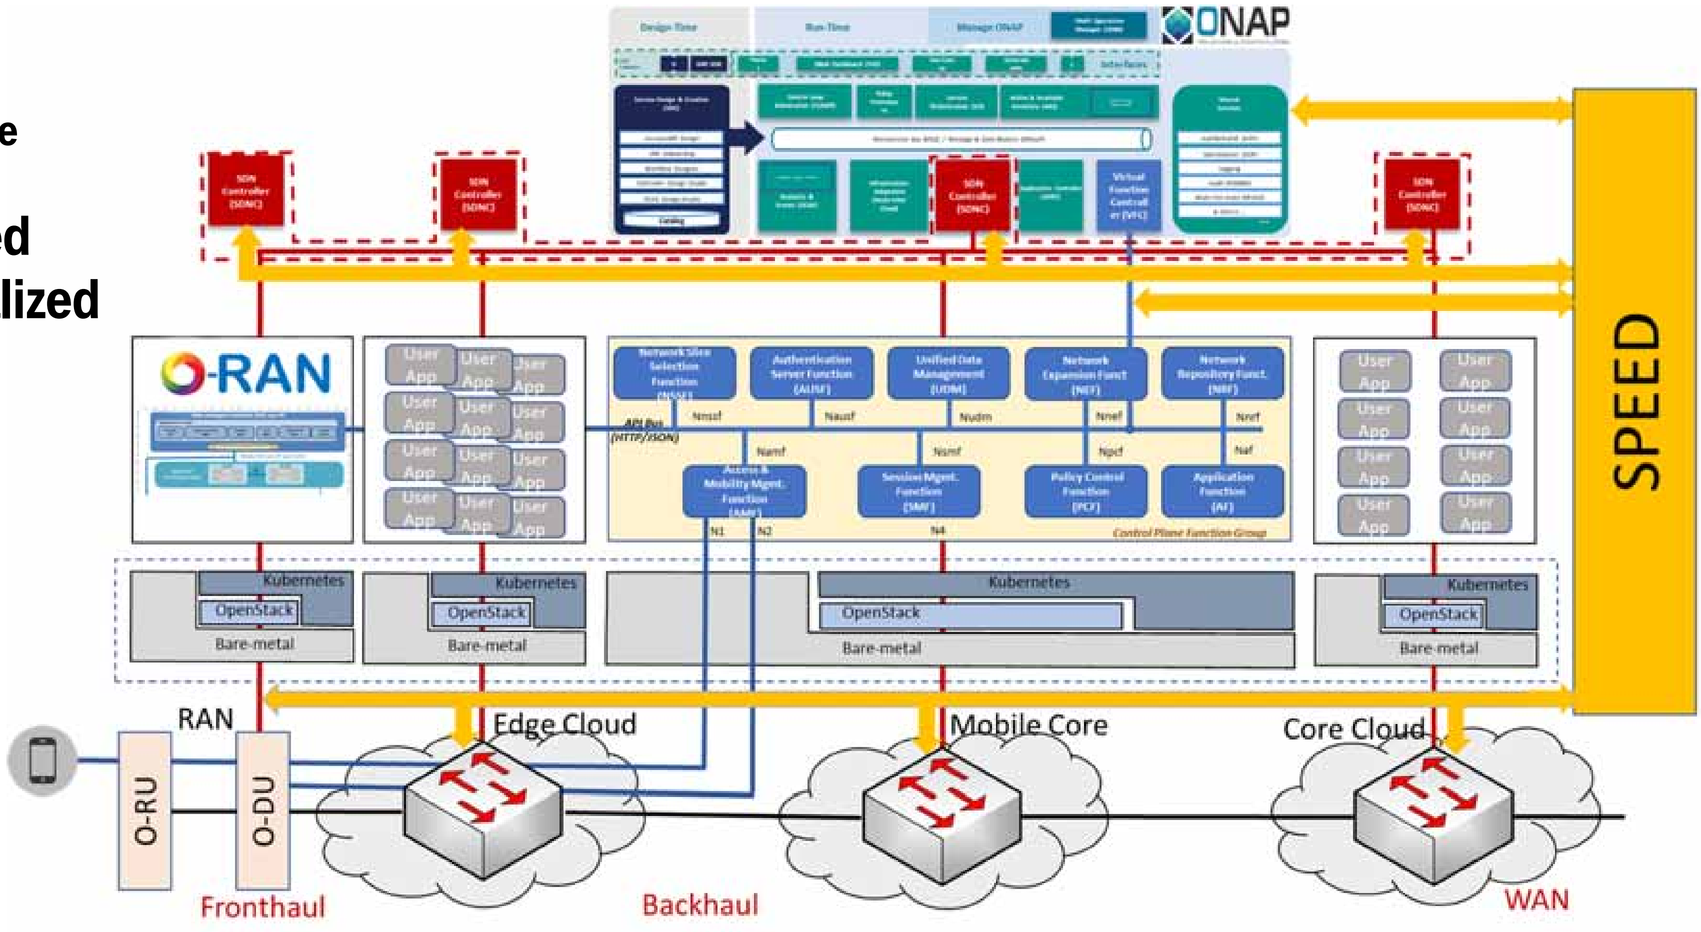

COSMOS Wide Area Connectivity

- Fully programmable (P4/SDN) wide(-er) area transport
- Low-latency high speed (100 Gbps/400 Gbps) connectivity to several (public/private) clouds
- Experimentation with reservation BW to various locations around the world

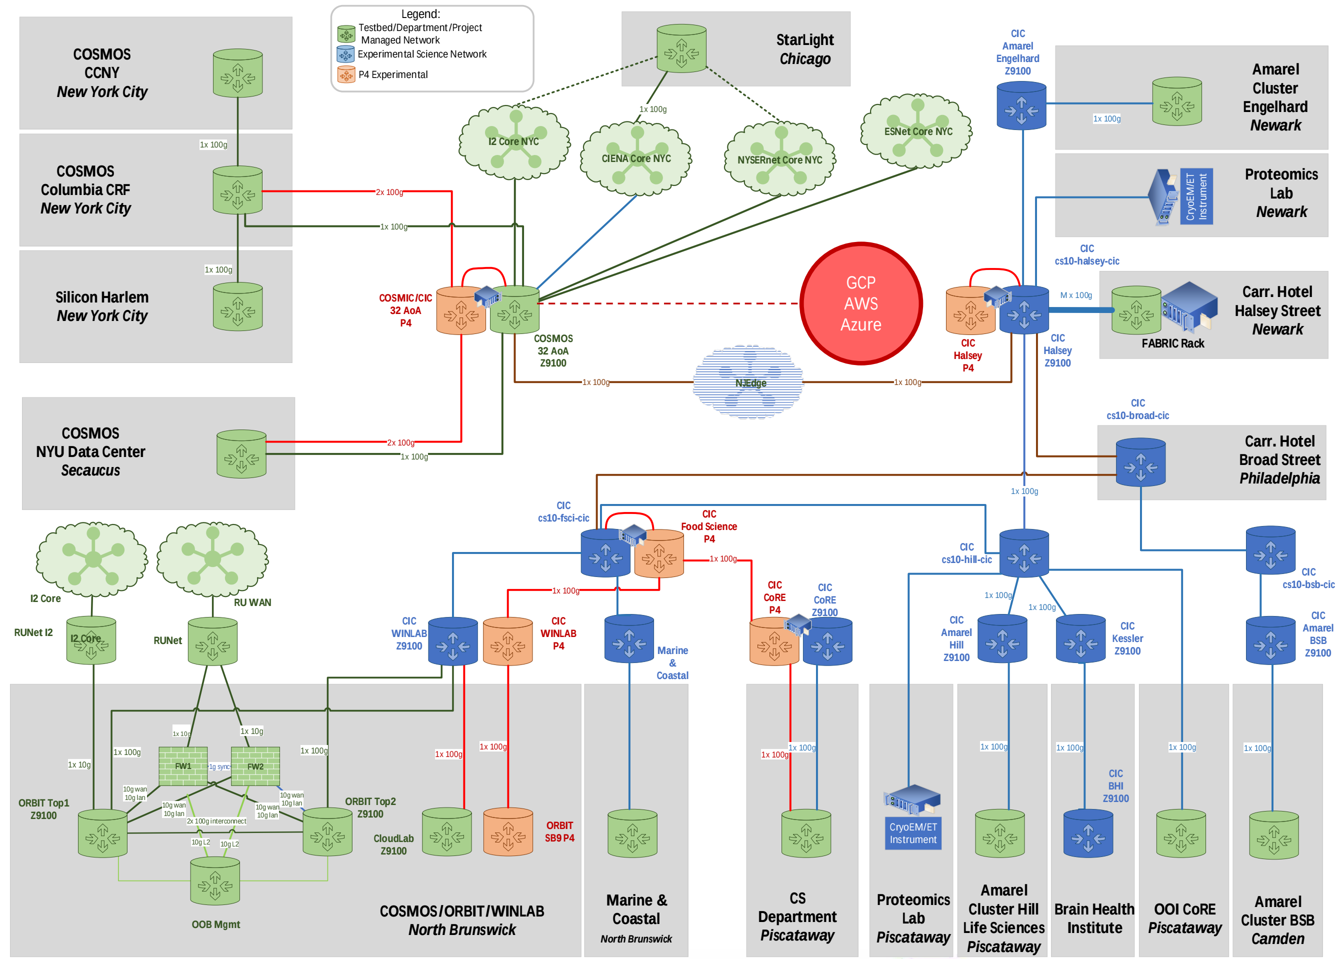

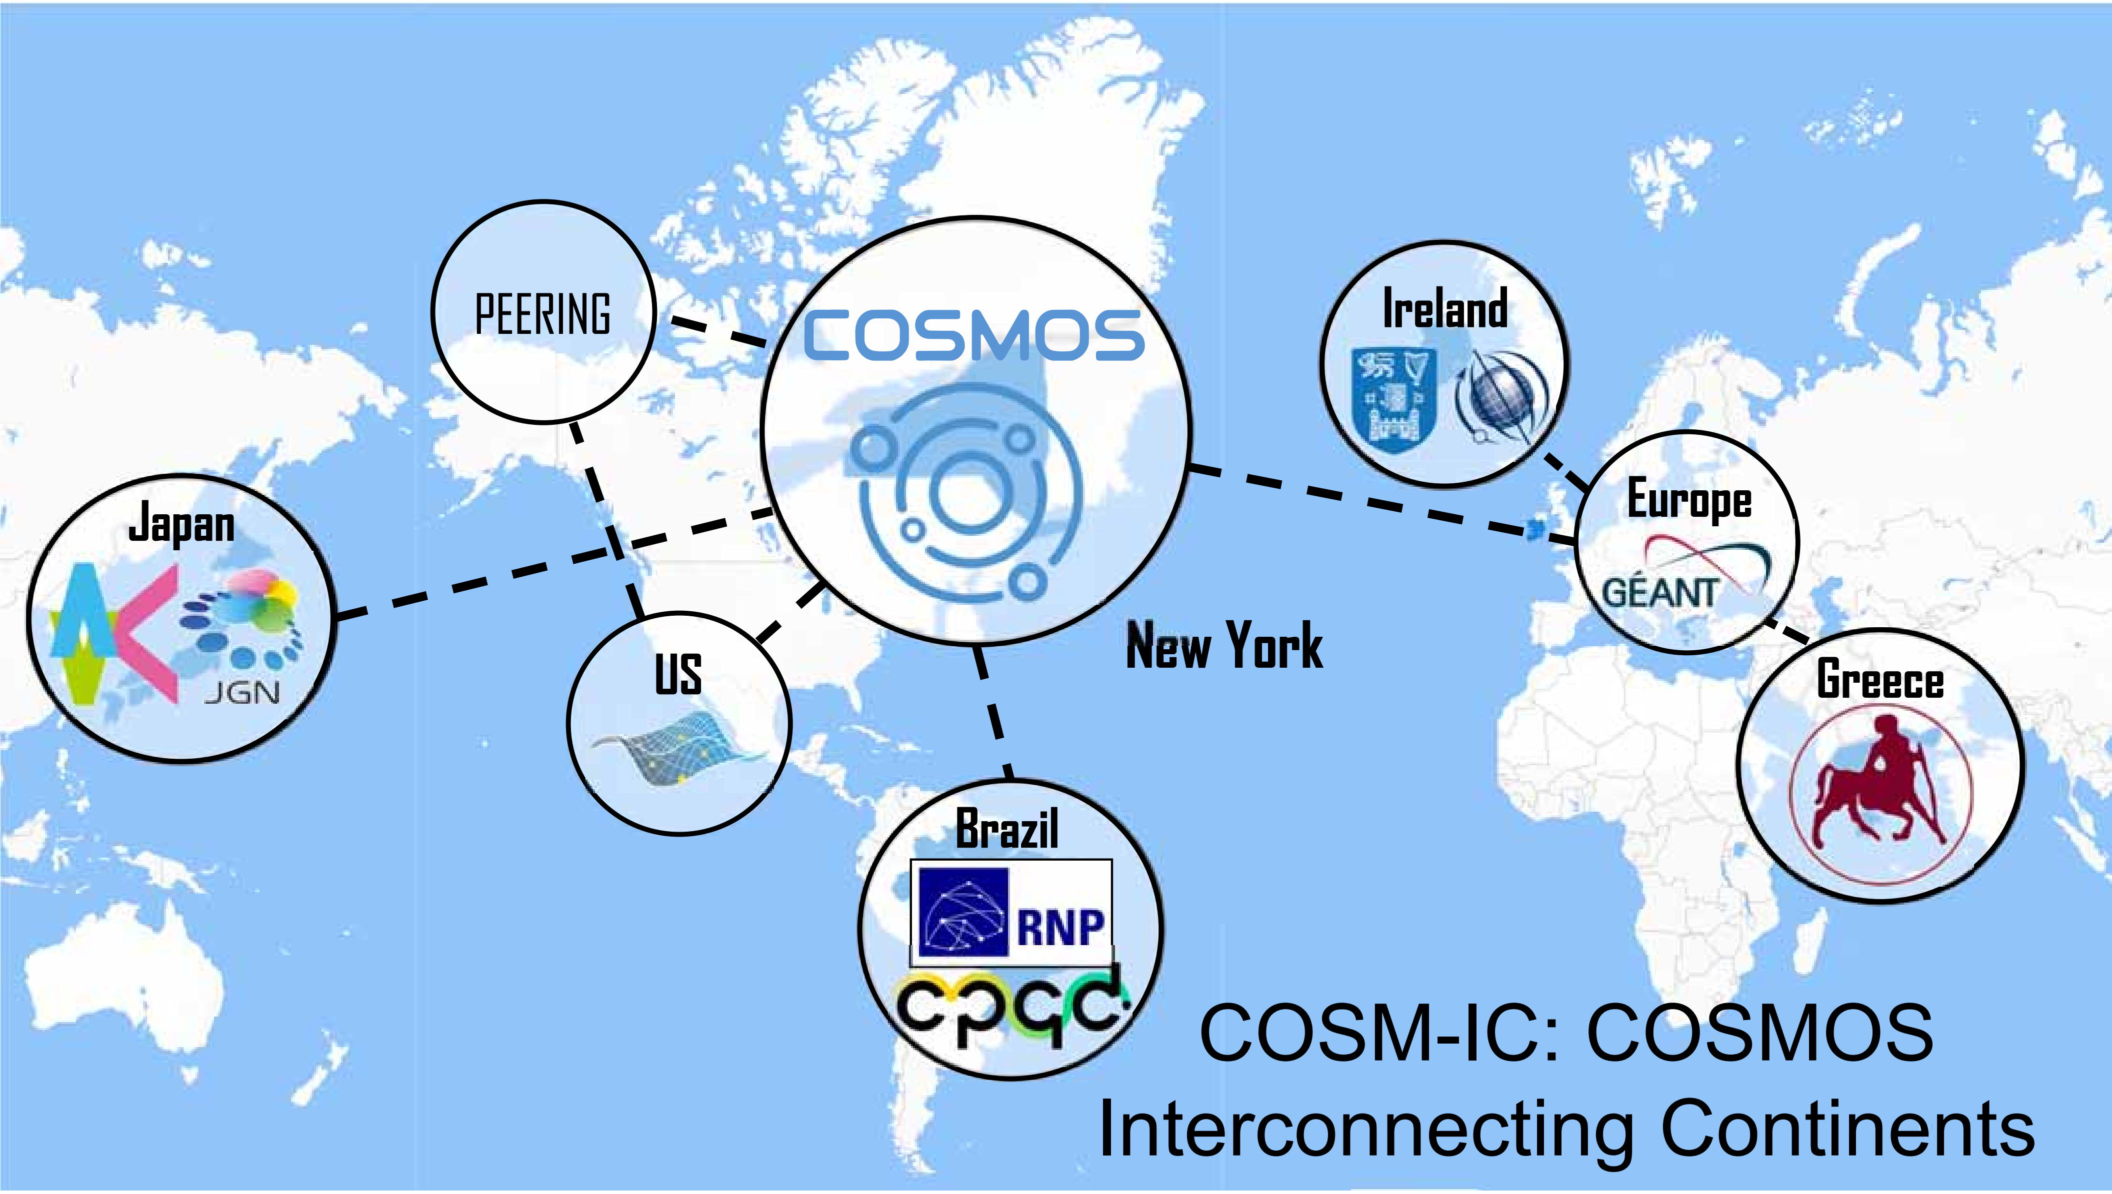
PEERING
COSMOS
Ireland
Japan
JGN
Europe
GÉANT
US
New York
Greece
Brazil
RNP
cpqd
COSM-IC: COSMOS
Interconnecting Continents

COSMOS Wireless Testbed – Summary

- Focus on ultra-high bandwidth, ultra-low latency, and edge cloud
- Open platform integrating SDRs, mmWave, and optical x-haul
- 1 sq. mile densely populated area in West Harlem
- Industry and local community outreach

COSMOS website: https://cosmos-lab.org
Tutorials: https://wiki.cosmos-lab.org/wiki/tutorials
Twitter: #pawrcosmos

Related links:

- PAWR: https://advancedwireless.org/
- ORBIT: https://www.orbit-lab.org/
- ONAP: https://wiki.onap.org
- ORAN: https://www.o-ran.org/

COSMOS

RUTGERS COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK NYU NYC <sh> The City College of New York