WINLAB Research Overview
March 2022

WINLAB

Rutgers, The State University of New Jersey
www.winlab.rutgers.edu
Contact: Prof. D. Raychaudhuri, Director
ray@winlab.rutgers.edu

WINLAB at Rutgers University:

University-Industry Research Center Specializing in Wireless Technology

Founded in 1989

WINLAB Tech Center Facility, North Brunswick, NJ

~25 Faculty & Staff
~45 PhD Students
~10 sponsor companies

Center's research portfolio spans: information theory, radio technology, wireless systems, mobile networks and computing

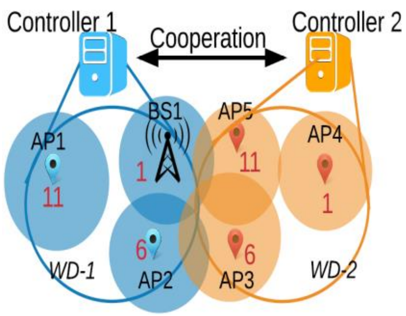

Dynamic Spectrum

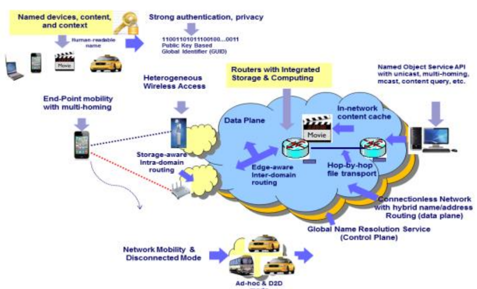

Future Internet Arch

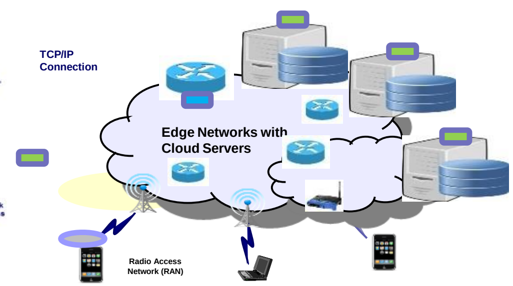

Edge Cloud

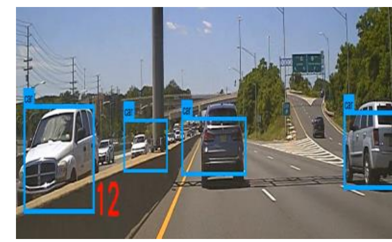

Connected Vehicles

Extensive experimental research infrastructure including ORBIT & GENI testbeds, SDR, …

Low Power IoT device

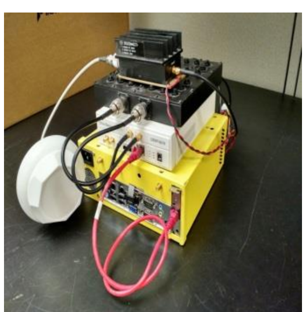

SDR unit

Massive MIMO

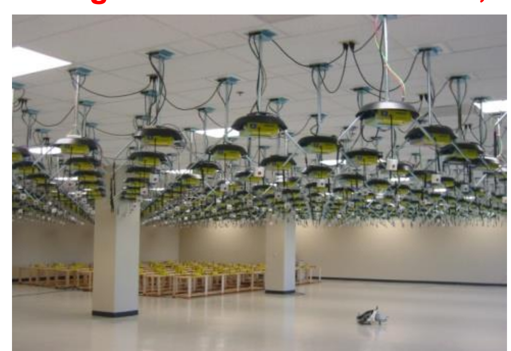

ORBIT Radio Grid

GENI Cloud Rack

WINLAB Research Topics Map (~2020-22)

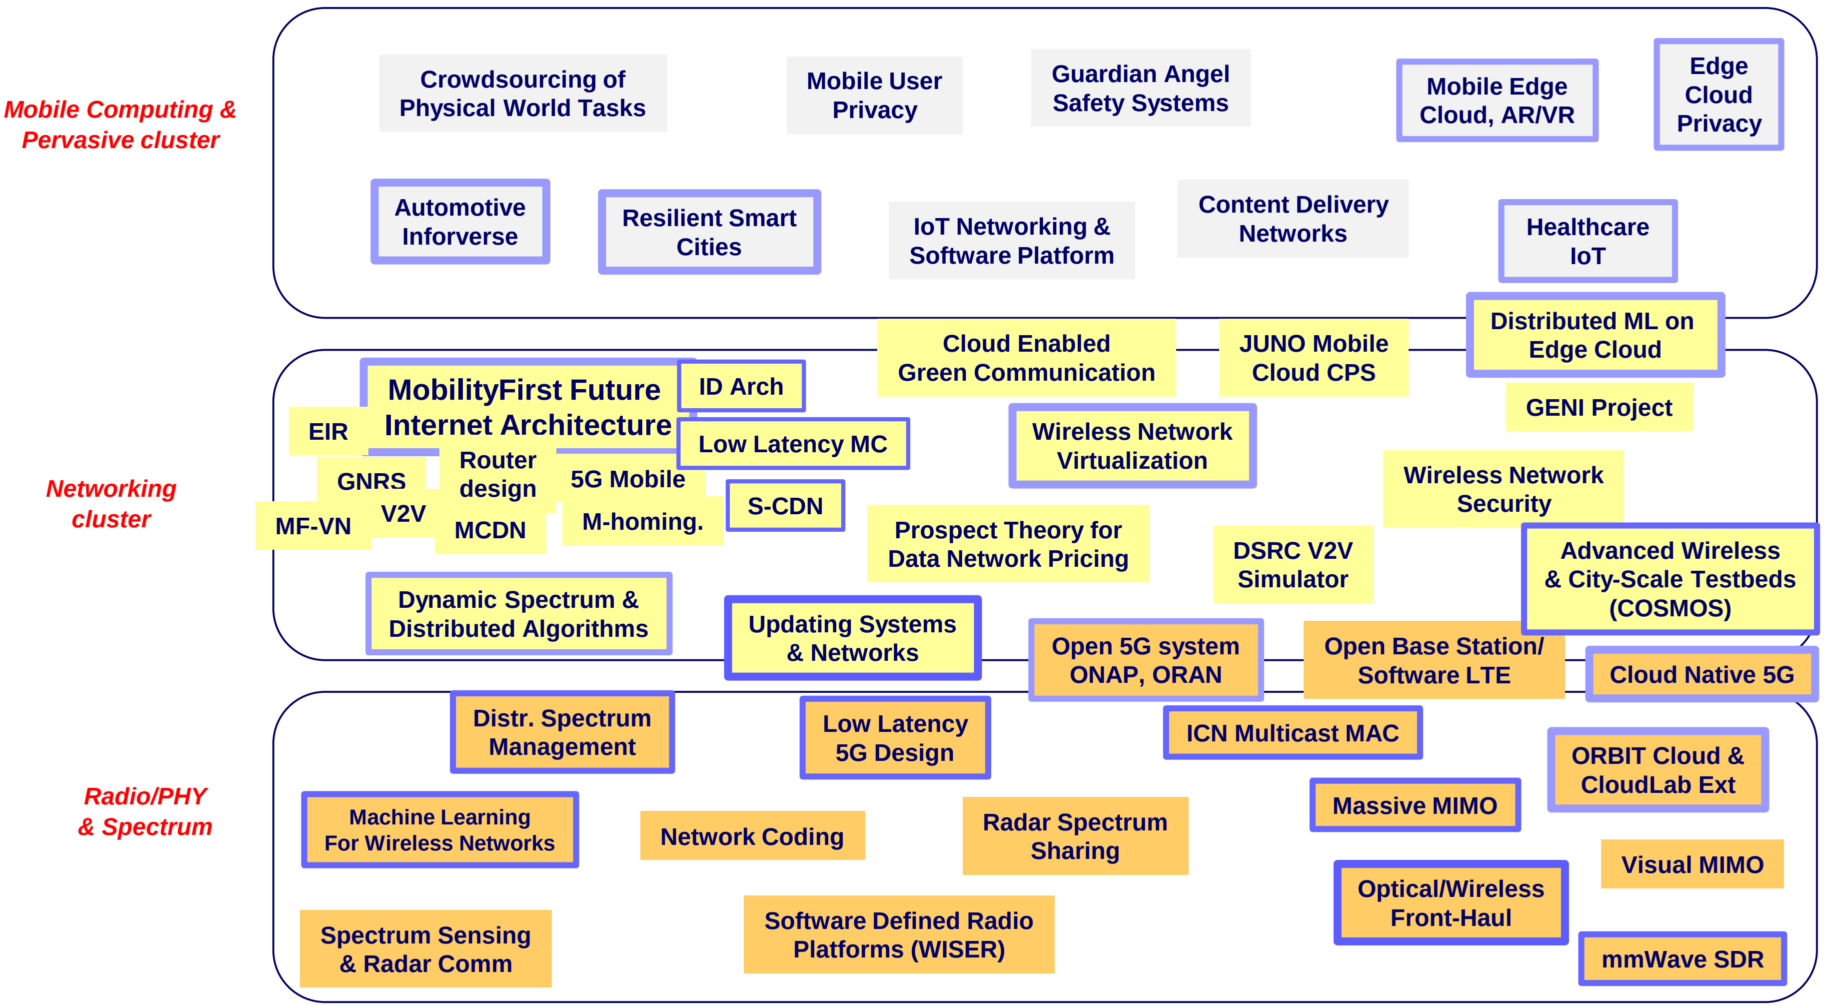

WINLAB System Research Themes

- Next-gen mobile/Internet – FIA/5G/NGMN, …
- Decentralized spectrum architecture (SII, SCM)
- Edge cloud for mobile real-time CPS services (AoI, distributed ML)
- Software defined wireless networks (O-RAN, O-NAP)
- Research testbeds: ORBIT, CloudLab, COSMOS, COSMIC
- Automotive infoverse, Cloud-assisted autonomous vehicles
- Smart Building, Smart City, …(ERC collaboration with Columbia)
- Health IT systems and applications
- Resilient architecture; disaster recovery networks (NSF, NIST)

Assistant Professor

WINLAB Summary: People

Dipankar Raychaudhuri

Roy Yates

Narayan Mandayam

Yinying Chen

Wade Trappe

Predrag Spasojevic

Yanyong Zhang

Marco Gruteser

Ivan Seskar

Yicheng Lu

Athina Petropulu

Waheed Bajwa

Dick Frenkiel

Rich Howard

Richard Martin

Anand Sarwate

Bo Yuan

Thu Nguyen

Vivek Singh

Lisa Musso

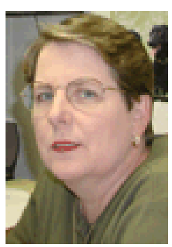

Noreen DeCarlo

Janice Campanella

Jakub Kolodziejski

Elisa Servito

Prashanthi Madala

~30-PhD & MS
(see www.winlab.rutgers.edu for photos)

WINLAB Summary: Industry Sponsors & Partners (Current & Recent)

verizonwireless

ERICSSON
TAKING YOU FORWARD

BAE SYSTEMS

Cisco Systems

InterDigital

NTT docomo

SES

accenture

NEC Empowered by Innovation

QUALCOMM

Raytheon
BBN Technologies

NiCT National Institute of Information and Communications Technology *

TOYOTA InfoTechnology Center *

NOKIA *

ARUBA networks *

InPoint *
Igolgi
Semandex

***Research Partners**

Selected Research Projects

ORBIT Testbed: Massive-MIMO

Ivan Seskar,

- 32 USRP X310s
  - Available FPGA resources:

Resource Type	Number
DSP48 Blocks	58K
Block Rams (18 kB)	14K
Logic Cells	7.2M
Slices (LUTs)	1.5M

  - RF 2 x UBX-160 (10 MHz - 6 GHz RF, 160 MHz BB BW)
  - 2 x 10G Ethernet for fronthaul/interconnect
  - Four corner movable mini-racks (4 x 16 x 16 -> 1 x 64 x 64)
- > 500+ GPP Cores/CloudLab Rack
- Number of GPU platforms
- 32x40G SDN aggregation switch

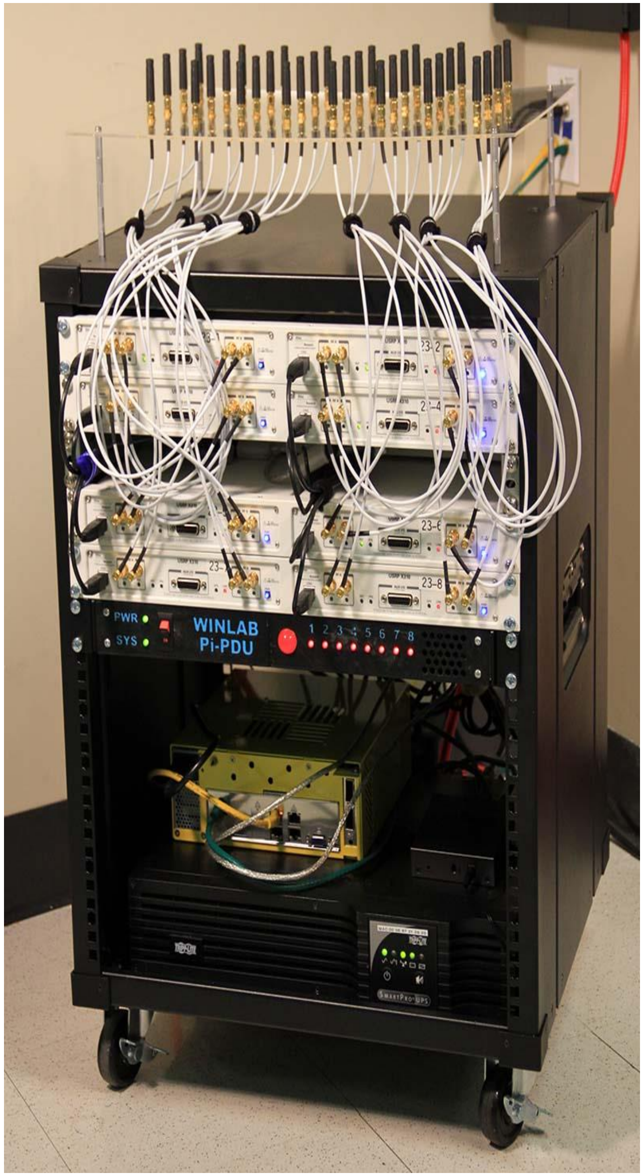

Massive MIMO Measurements

Ivan Seskar,

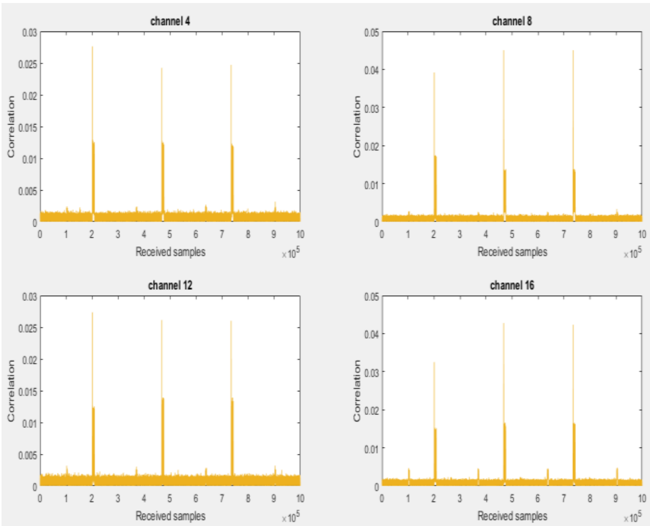

OFDM packet detection and time synchronization

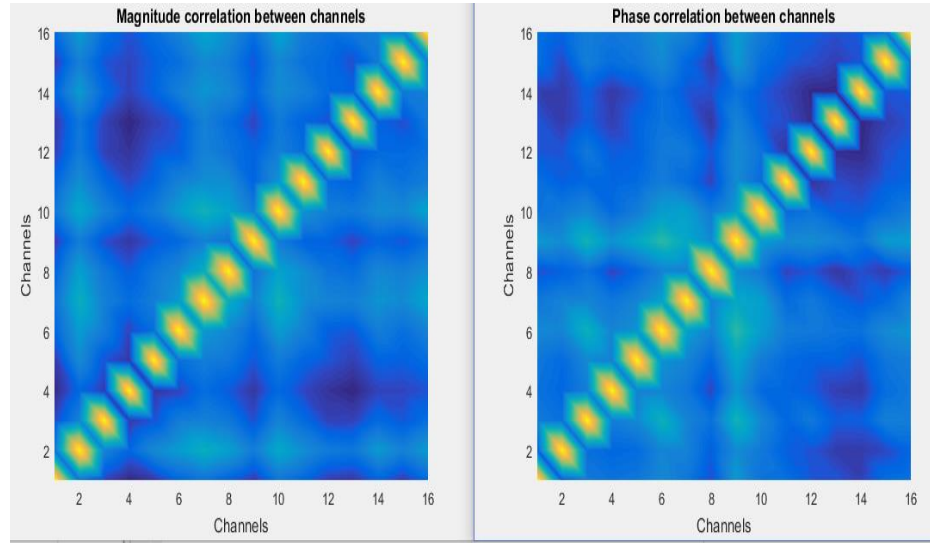

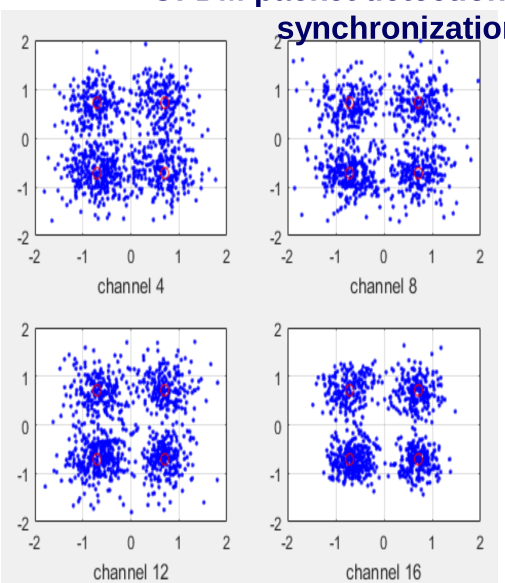

QPSK Decoding

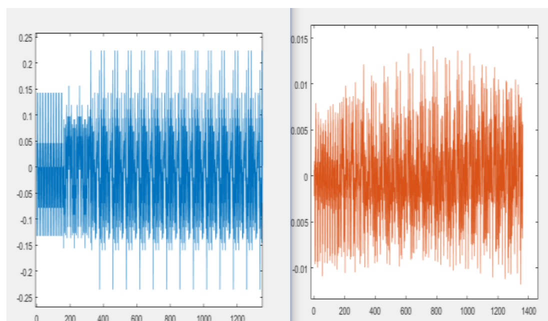

Transmitted Signal Received signal

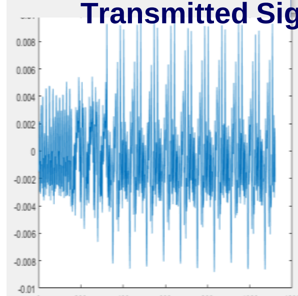

Received signal (with external reference)

Frequency Offset correction using external reference

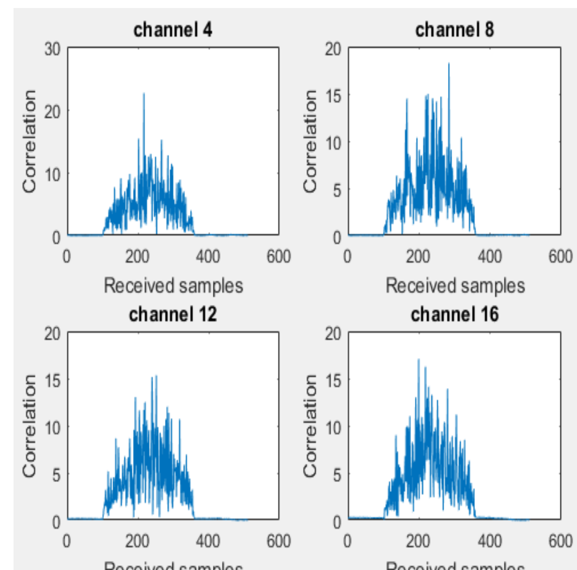

Multi PN sequence transmission and detection

Distributed Dynamic Spectrum (SMAP) Architecture

Prof. D. Raychaudhuri

Cloud-based regional aggregator level

WD controller level

Radio device level

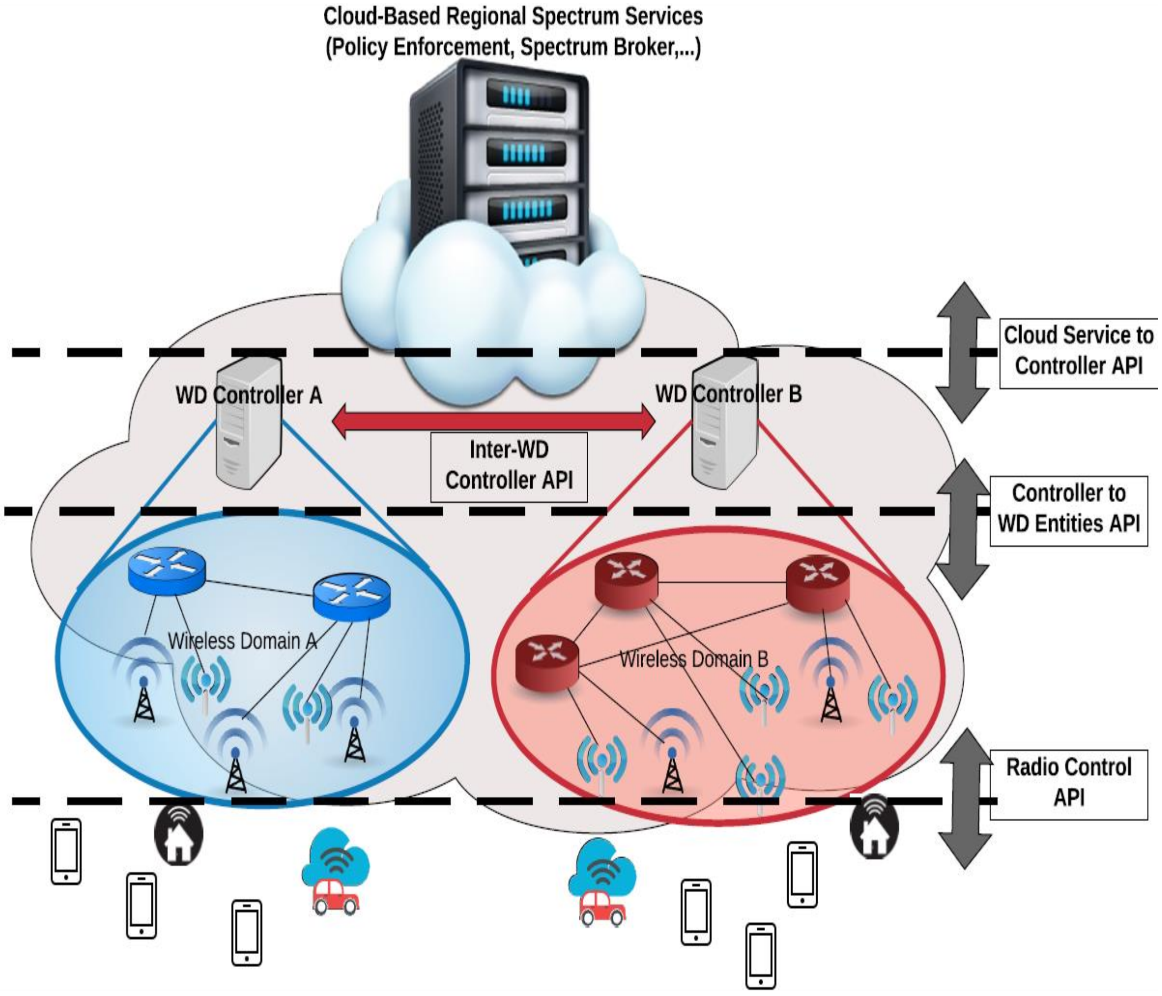

SMAP Proof of Concept using ORBIT Testbed

- Benchmark experiment setup for channel assignment algorithm

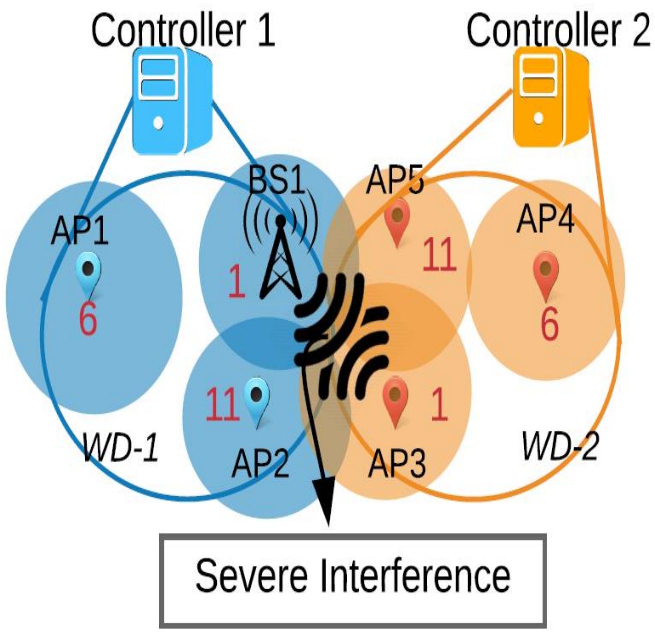

No cooperation between WD controllers:
LTE BS1 causes interference to Wi-Fi AP3 when operating on the same channel

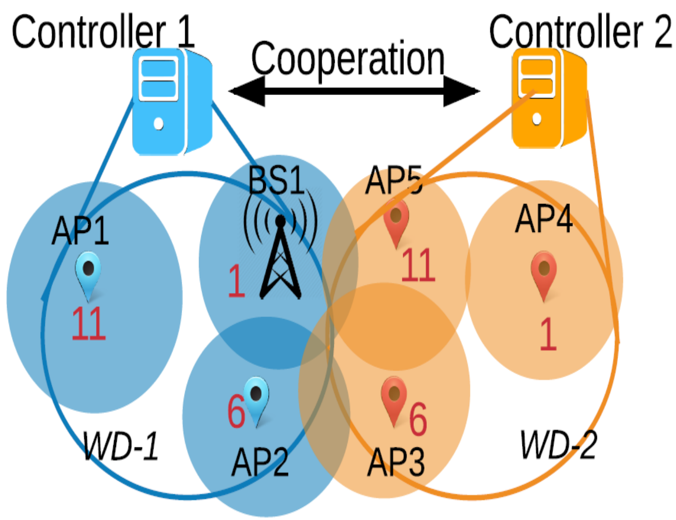

Inter-network coordination between WDs: interference avoidance between LTE and Wi-Fi

SMAP Experimental Performance: Wi-Fi/LTE Coexistence

Fairness is achieved through Inter-Network Coordination

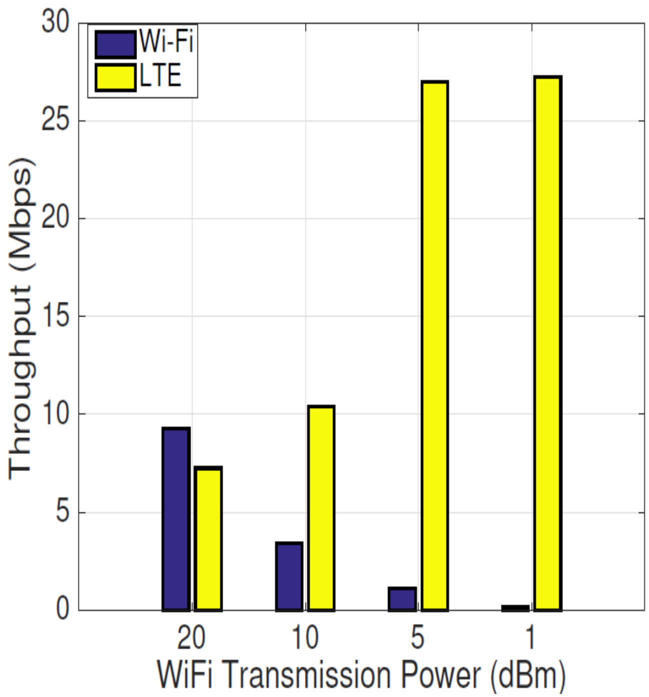

LTE-Wi-Fi coexistence: with decrease in Wi-Fi transmission power, LTE interference significantly degrades Wi-Fi throughput

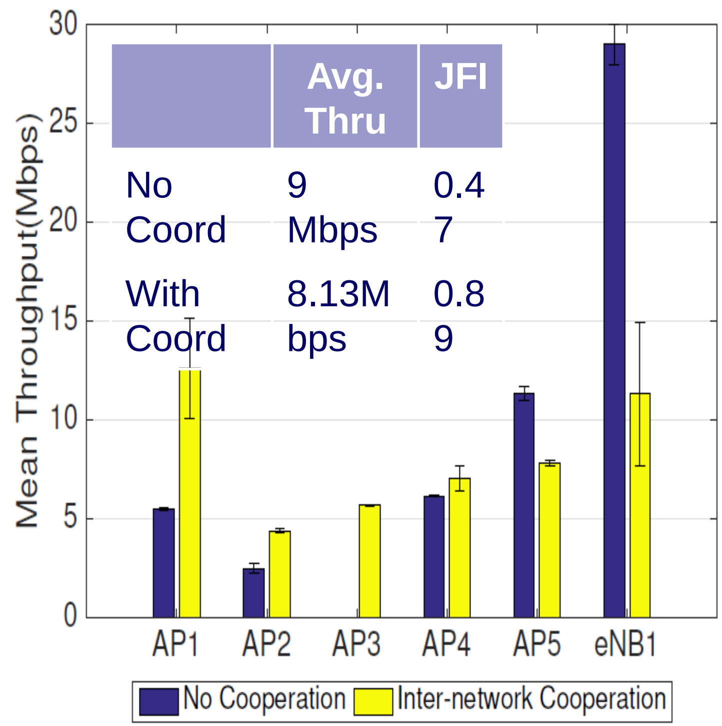

	Avg. Thru	JFI
No Coord	9 Mbps	0.47
With Coord	8.13Mbps	0.89

Throughput of Wi-Fi/LTE averaged over 5 experiment runs

Mobile Core Network: 5G vs 4G Latency

- 4G has high latency due to control plane signaling and radio framing
- 5G improvements include
  - Modular architecture (NFV)
  - Session setup and management improvements
- Further improvements will require
  - Self certifying packets (fast authentication)
  - ICN packet switched architecture with minimum signaling
  - “Flat” core network → no gateways, distributed mobility support at routers

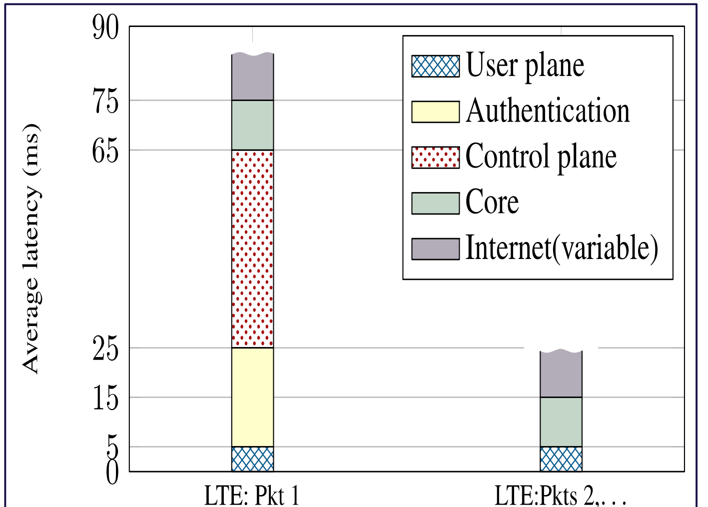

Average latency for data packets in 4G LTE §

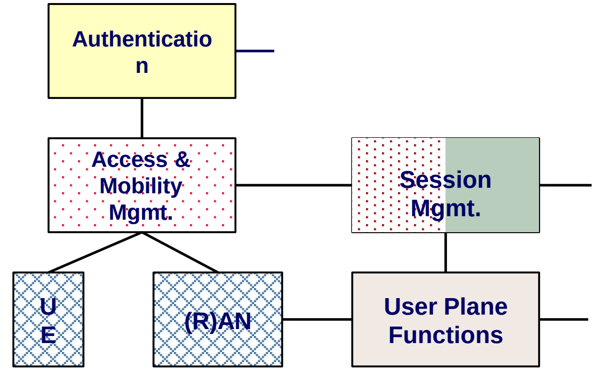

5G System Architecture in 3GPP Technical Spec. 23.501 v0.3.1 (2017-03)

Source: § J. Huang, et al., “A Close Examination of Performance and Power Characteristics of 4G LTE Networks“, in ACM MobiSys 2012

Mobile Core Network: 4G LTE Connection Establishment

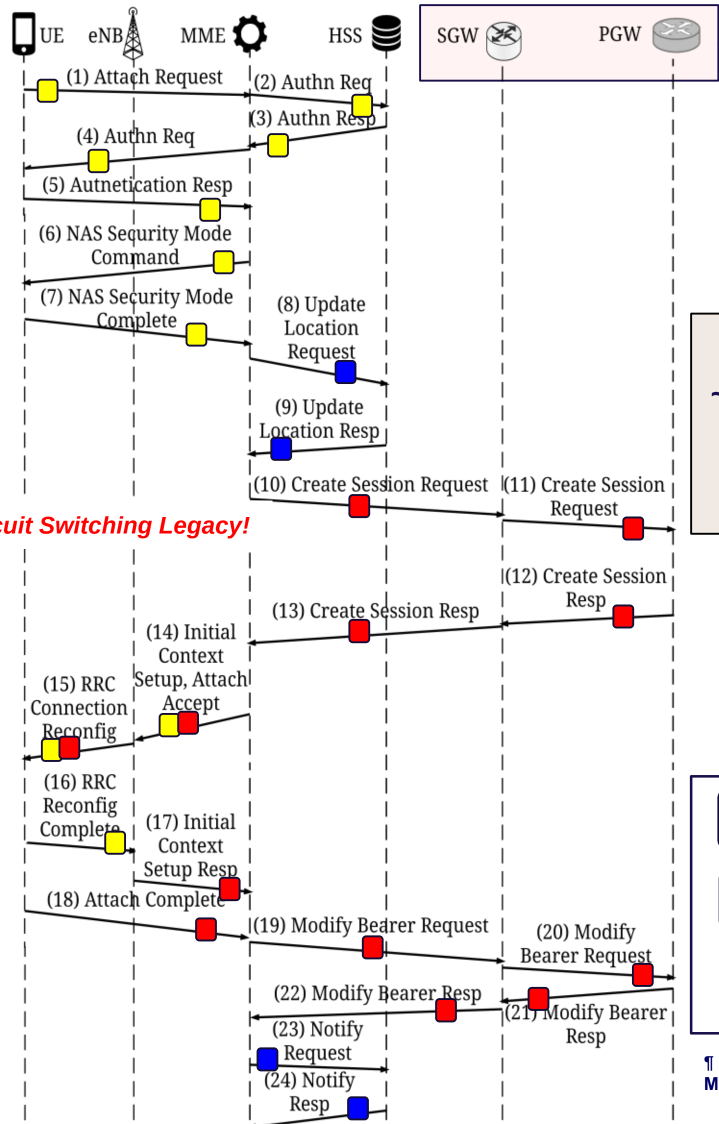

Circuit Switching Legacy!

~4-6 gateways in the US ➔ bottleneck + non-optimal paths

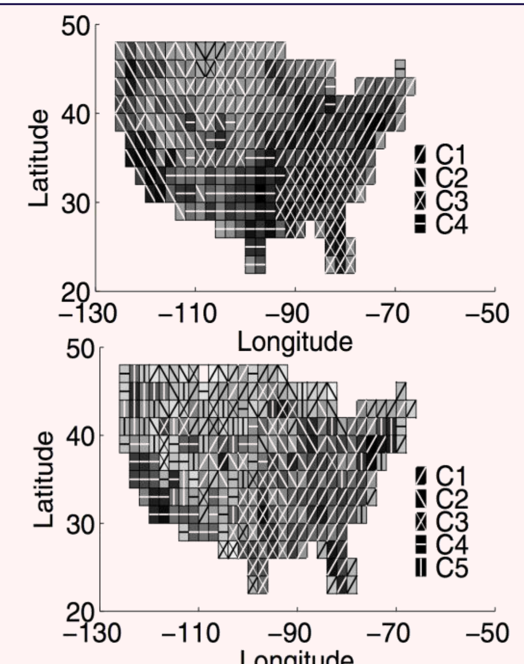

Clustering UE traffic to gateways for 2 different LTE carriers ¶

10	**Security**
12	**Session**
4	**Mobility**

~24 messages for connection setup ➔ high latency

¶ Q. Xu, et al., "Cellular Data Network Infrastructure Characterization and Implication on Mobile Content Placement ", in ACM Sigmetrics 2011

Mobile Core Network Based on Named Object MF Architecture

Control Plane

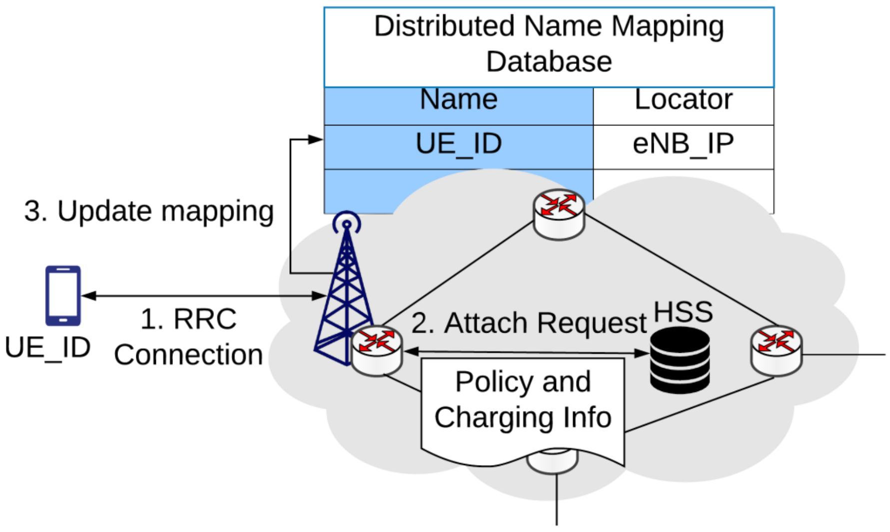

Data Plane

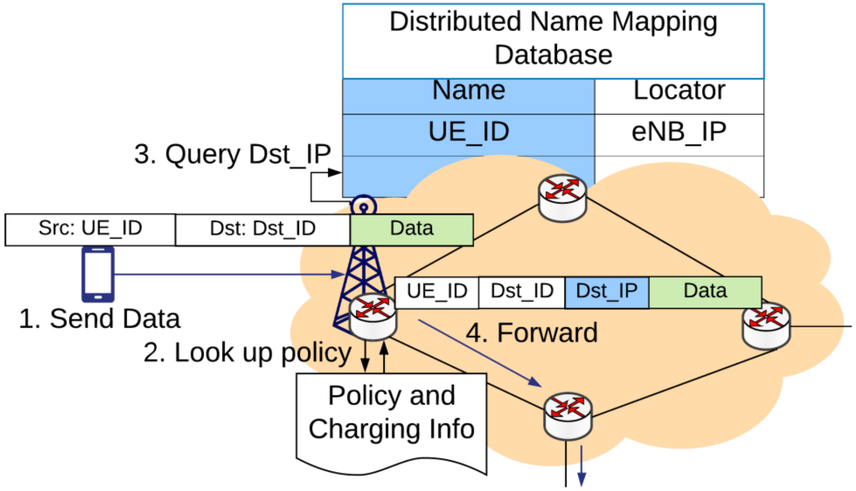

- MobilityFirst (MF) is a Future Internet Architecture based on the concept of:
  - Unique names associated with Internet-connected objects
  - Distributed hash-map based database for mapping names to addresses
- No GTPs in control plane
- Distributed forwarding in data plane

Mobile Core Network: Simplified MF-based Protocol Stack

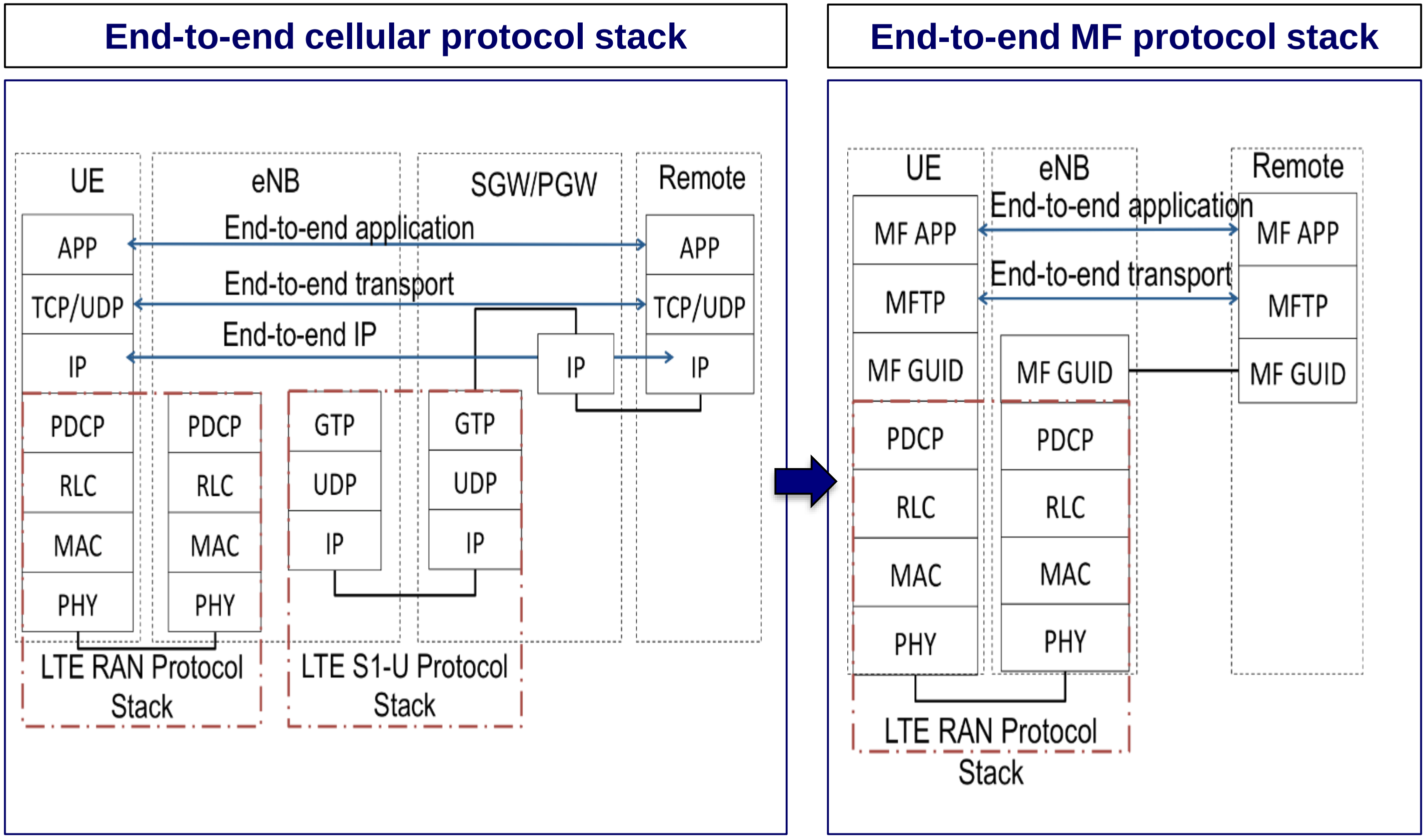

1. LTE RAN remains unchanged
2. Simplified end-to-end protocols for MobilityFirst

Mobile Core Network: Proof-of-Concept 5G Core Prototype

- MF (MobilityFirst)/ICN hoststack and Click router updated to work with OAI
- OAI modified to remove gateways and tunneling
- Preliminary ORBIT experiments on low latency

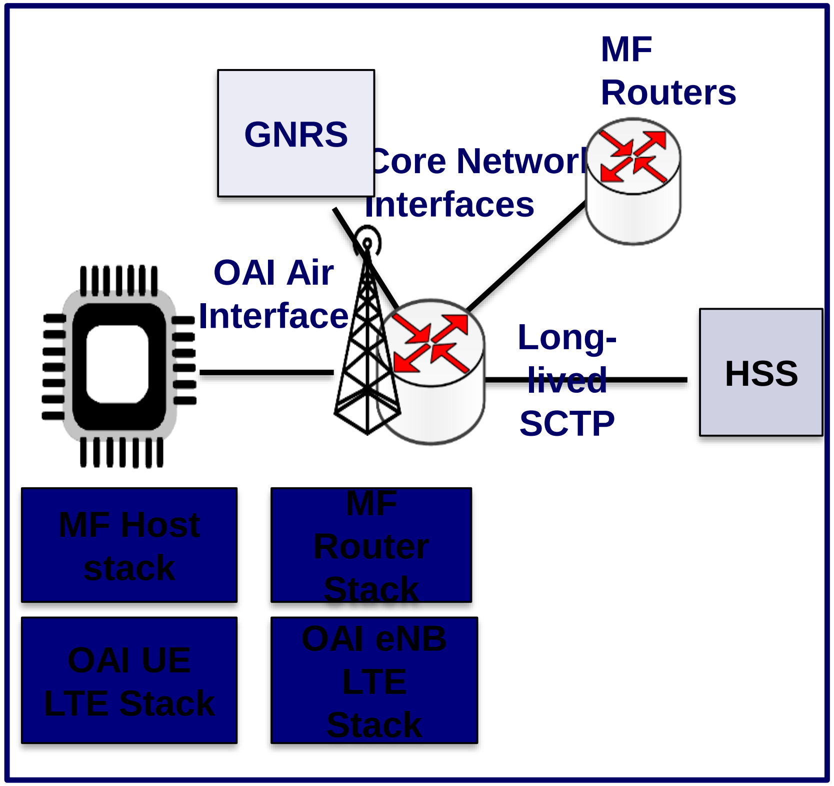

Mobile Core Network: Latency Comparison with Commercial Network

- UE attachment latency (MAC + network)
- Distributed core experiment on ORBIT
  - 50 runs of UE wakeup (i.e. control + data plane latencies)
- Commercial ISP data from MobileInsight*
  - 780 crowd-sourced smartphone logs from 2016

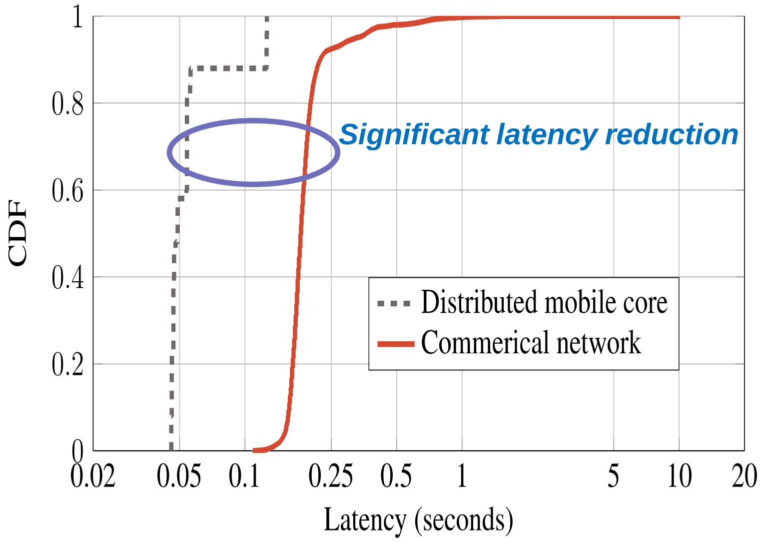

Source: Shreyasee Mukherjee, PhD thesis 2018, WINLAB

- **Average control messaging latency:**
  - **Distributed core: 49 msec**
  - **Commercial ISP: 181 msec**

Edge Cloud: Multi-tier Architecture

- Cloud service latency significant component of E2E
- Bringing the cloud closer to the edge helps:
  - Lower network RTT
  - NFV placement
- Several design challenges with edge cloud:
  - Dynamic assignment, load balancing
  - Mobility support
  - Real-time orchestration
  - Virtual network slicing

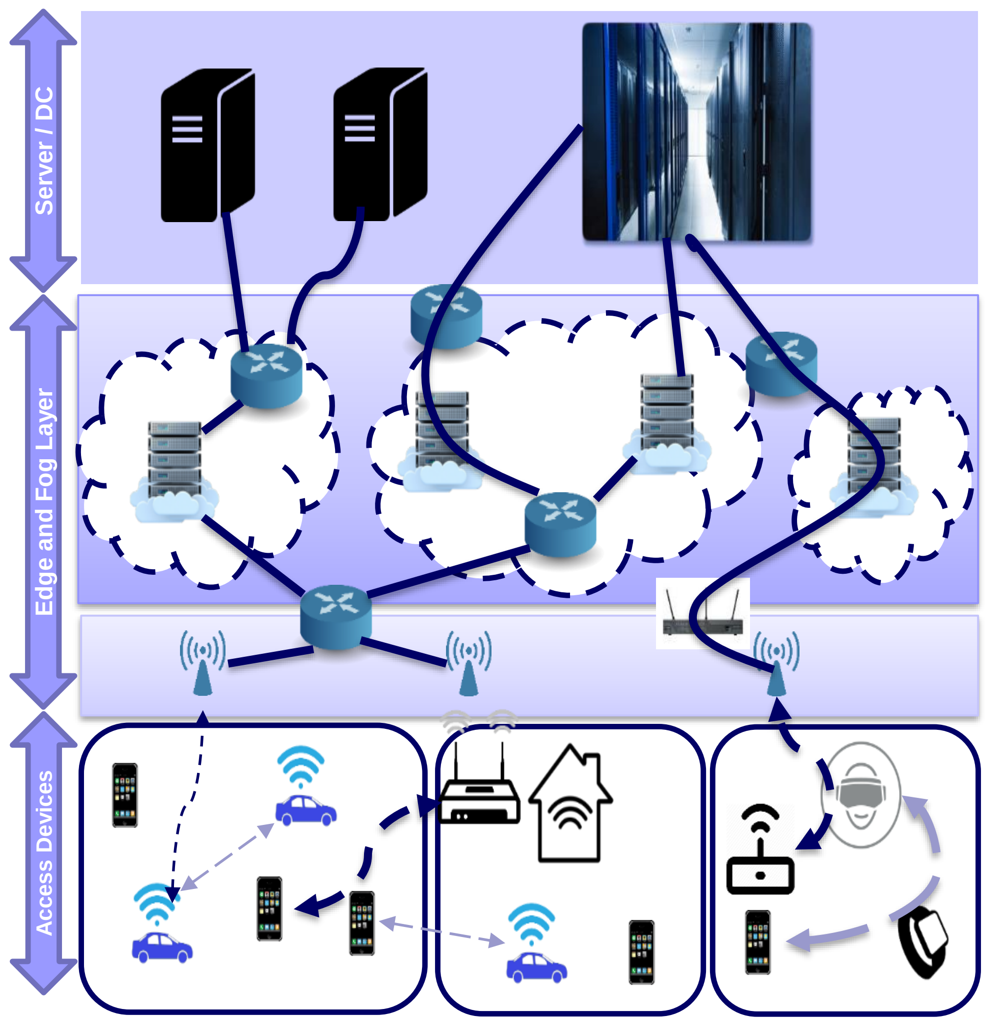

Edge Cloud: City Scale Evaluation Model

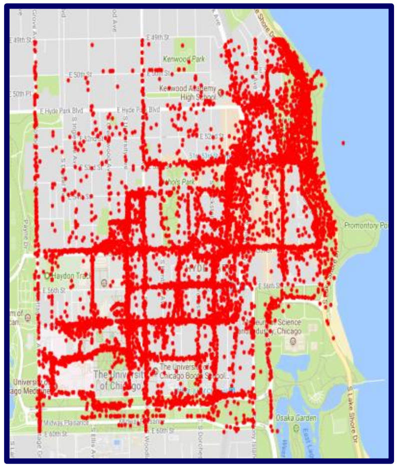

Chicago Wi-Fi AP Locations (~1.1K)

Data Source: www.wigle.net
Plot: http://hamstermap.com

Simulation Parameters

Parameter	Value
Area	5.18 km^2
Number of APs	1.1K
Number of Users	55K
Distribution of Users	Random
Bandwidth (Uplink)	27, 150 and 300 Mbps
Bandwidth (downlink)	54, 300 and 600 Mbps
Packet Size	1500 Bytes
Edge Resources (baseline)	5 Machines
α	2
β	1
γ	0.1
δ	1
ρ	0.9
w	0.5
p	10

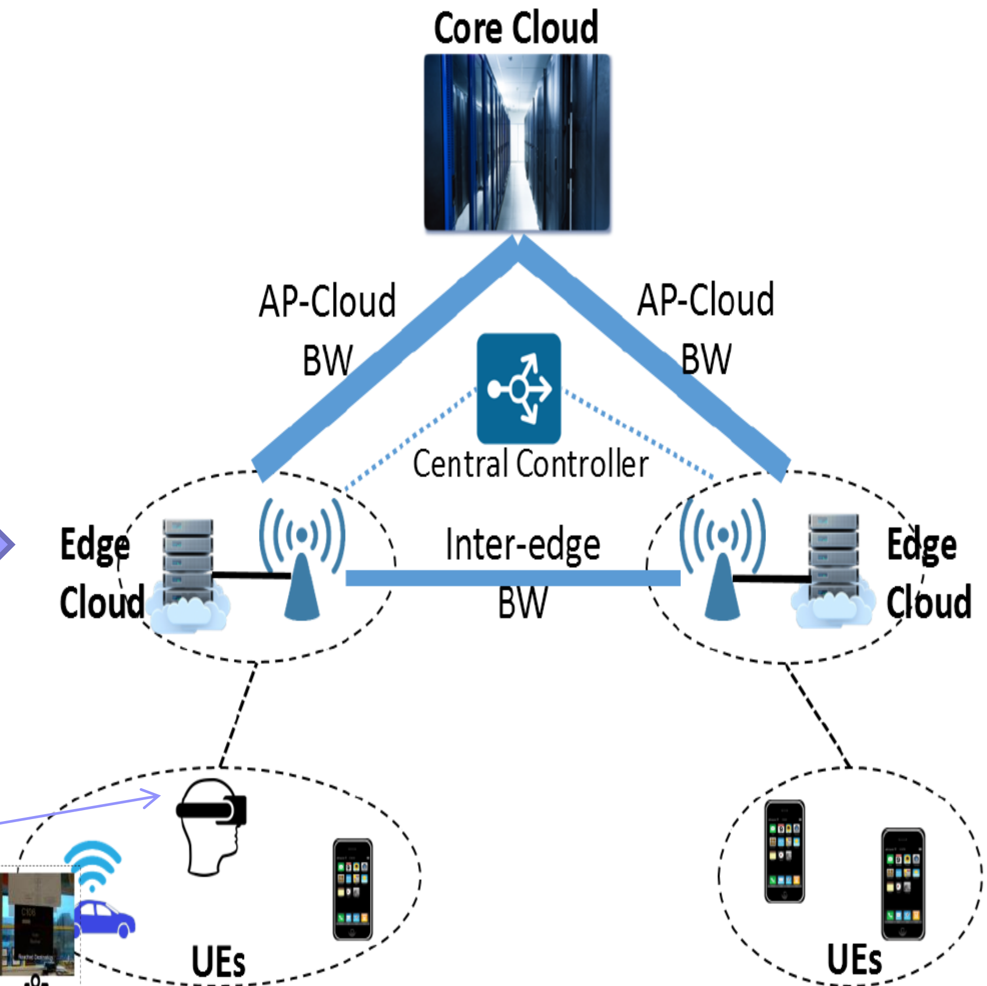

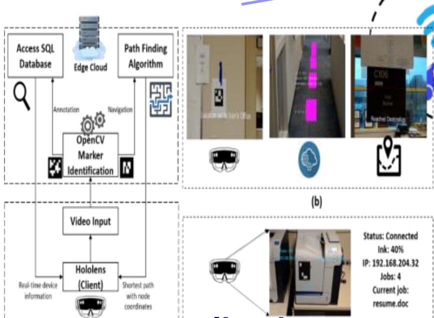

AR Application

Source: Maheshwari et al, IEEE SEC 2018

Edge Cloud: Sample Results for Hybrid MEC Evaluation

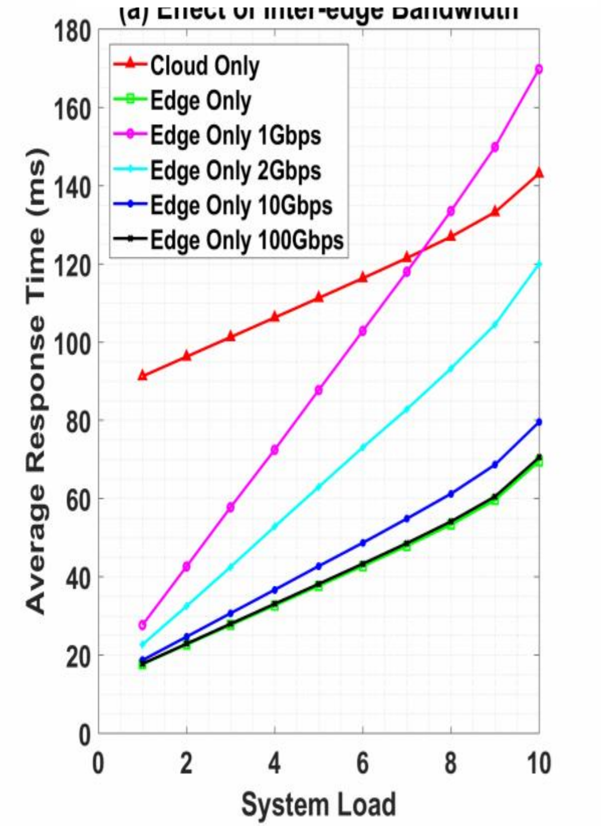

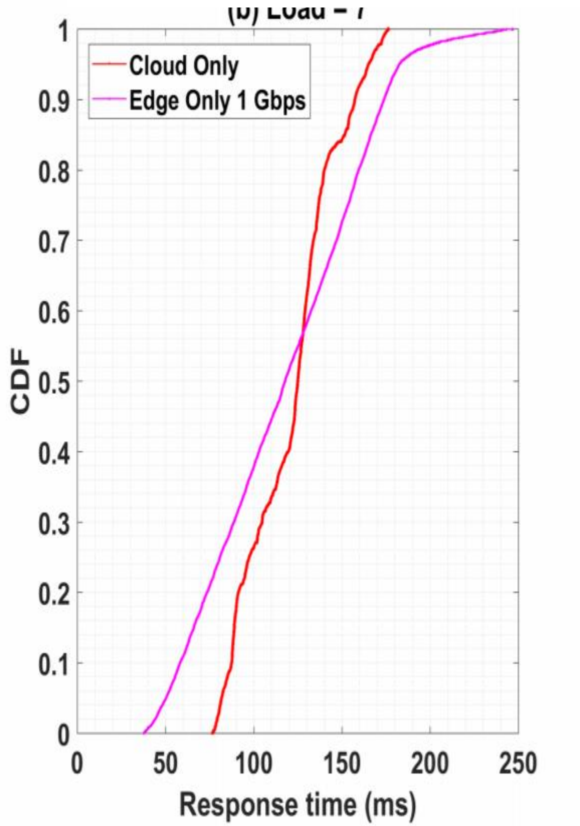

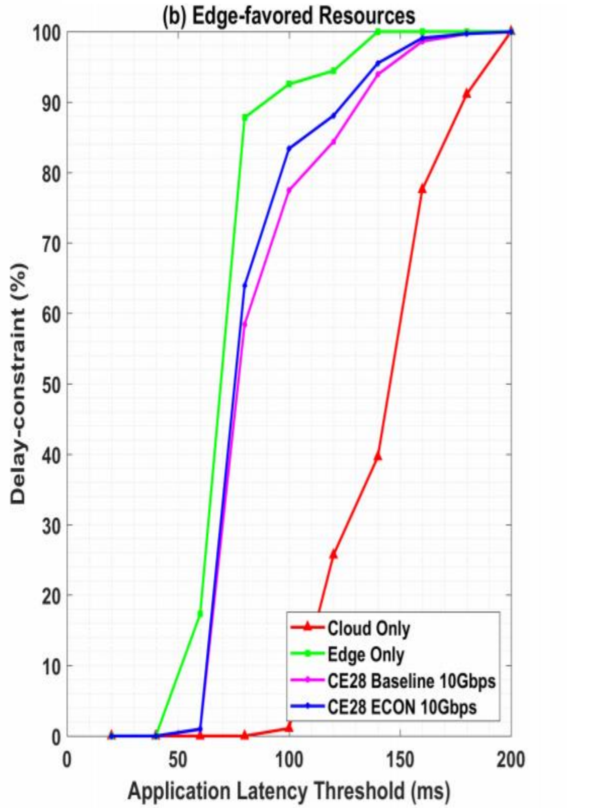

Source: Maheshwari et al, IEEE SEC 2018

- *Edge provides significant performance improvement*
- *Fast front-haul network needed to enable the edge…*

Distributed ML over Edge Clouds: System Architecture

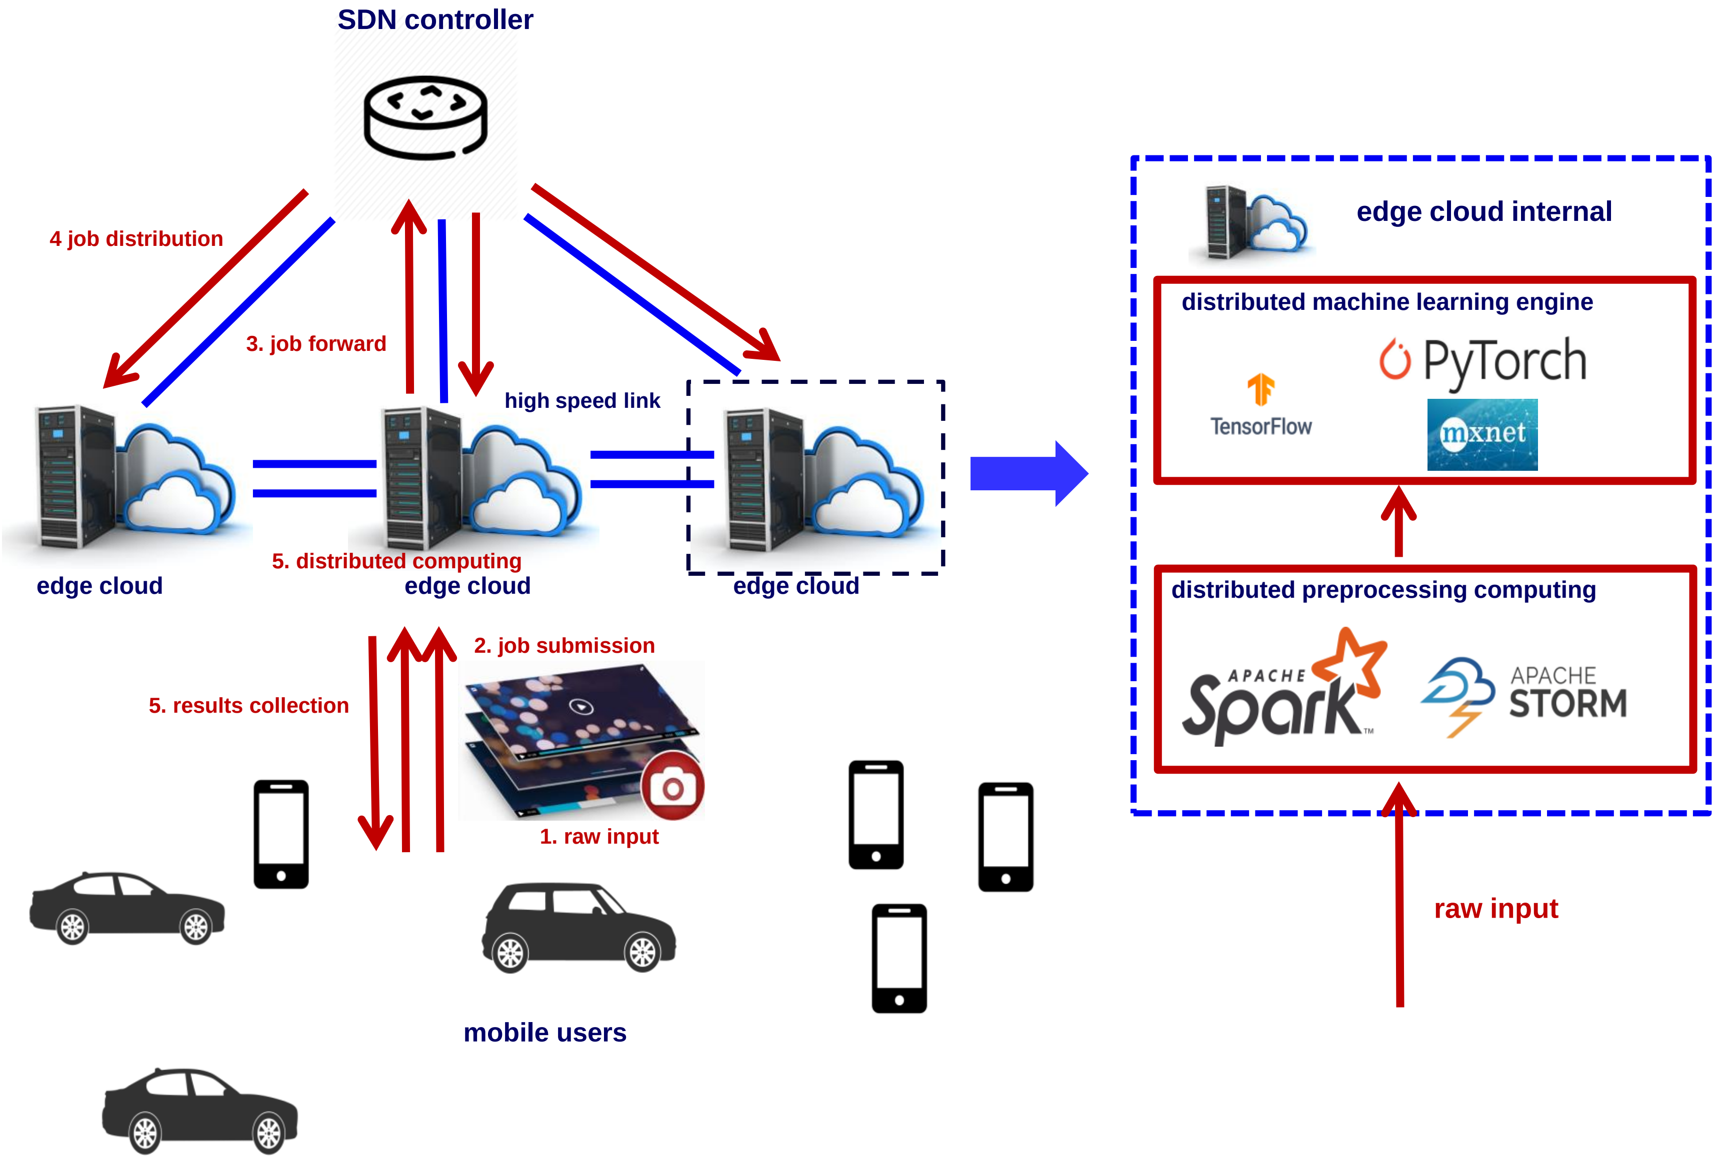

Distributed ML over Edge Clouds: Mobile Offload Performance Profiling

- **mobile device: *Jetson Tx2*; edge device: *1080Ti***
- **link: *200Mbps***
- ***network: VGG19***

Layer (type)	Output Shape	Parameters#	Size (KB)	MFLOPS	Cloud Proc (ms)	Mobile Proc (ms)	Trans (ms)	Total (ms)
Original Image	(224, 224, 3)	0	173	0	3.713	0	6.92	10.633
$\text{conv1}_1 + relu1_1$	(224, 224, 64)	1792	12544	173	3.627	0.427	501.76	505.814
$\text{conv1}_2 + relu1_2$	(224, 224, 64)	36928	12544	3699	3.362	1.743	501.76	506.865
pool1	(112, 112, 64)	0	3136	2.4	3.289	2.106	125.44	130.835
$\text{conv2}_1 + relu2_1$	(112, 112, 128)	73856	6272	1850	3.158	2.757	250.88	256.795
$\text{conv2}_2 + relu2_2$	(112, 112, 128)	147584	6272	3699	2.953	3.775	250.88	257.608
pool2	(56, 56, 128)	0	1568	1.2	2.910	3.988	62.72	69.618
$\text{conv3}_1 + relu3_1$	(56, 56, 256)	295168	3136	1850	2.774	4.664	125.44	132.878
$\text{conv3}_2 + relu3_2$	(56, 56, 256)	590080	3136	3699	2.561	5.722	125.44	133.723
$\text{conv3}_3 + relu3_3$	(56, 56, 256)	590080	3136	3699	2.345	6.794	125.44	134.579
pool3	(28, 28, 256)	0	784	0.6	2.310	6.968	31.36	40.638
$\text{conv4}_1 + relu4_1$	(28, 28, 512)	1180160	1568	1850	2.145	7.788	62.72	72.653
$\text{conv4}_2 + relu4_2$	(28, 28, 512)	2359808	1568	3699	1.874	9.134	62.72	73.728
$\text{conv4}_3 + relu4_3$	(28, 28, 512)	2359808	1568	3699	1.604	10.475	62.72	74.799
pool4	(14, 14, 512)	0	392	0.3	1.558	10.703	15.68	27.941
$\text{conv5}_1 + relu5_1$	(14, 14, 512)	2359808	392	925	1.451	11.235	15.68	28.366
$\text{conv5}_2 + relu5_2$	(14, 14, 512)	2359808	392	925	1.346	11.756	15.68	28.782
$\text{conv5}_3 + relu5_3$	(14, 14, 512)	2359808	392	925	1.239	12.288	15.68	29.207
pool5	(7, 7, 512)	0	98	0.075	1.203	12.466	3.92	17.589
fc6 + relu6	(1, 1, 4096)	102764544	16	206	0.239	17.254	0.64	18.133
fc7 + relu7	(1, 1, 4096)	16781312	16	34	0.060	18.143	0.64	18.843
fc8	(1, 1, 1000)	4097000	3.9	8	0	18.441	0.156	18.441
total		138357544	540459					

Distributed ML over Edge Clouds: Data Parallelism for Faster Inference

- target applications: multi-object detection / multi-face recognition

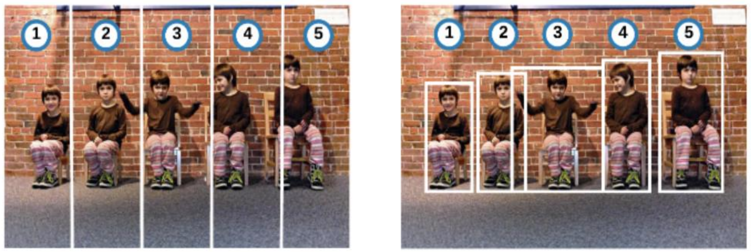

(a) image with equal partitioning (b) image with ideal partitioning

run video frame partitioning for data parallelism

no performance downgrade

lower latency

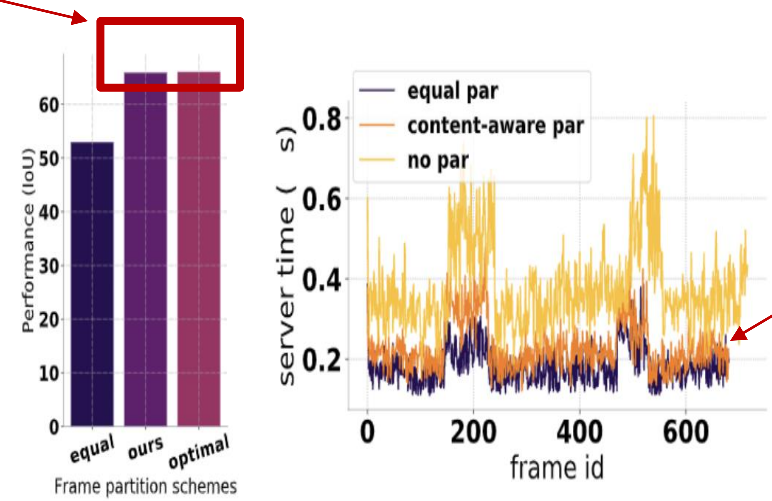

Accuracy and processing time with different frame partitioning schemes

Mobile Edge Cloud: Vehicular Applications

Cloud Assisted Driving and Self-Driving will soon become feasible due to advances in wireless network ("5G") and edge cloud technology

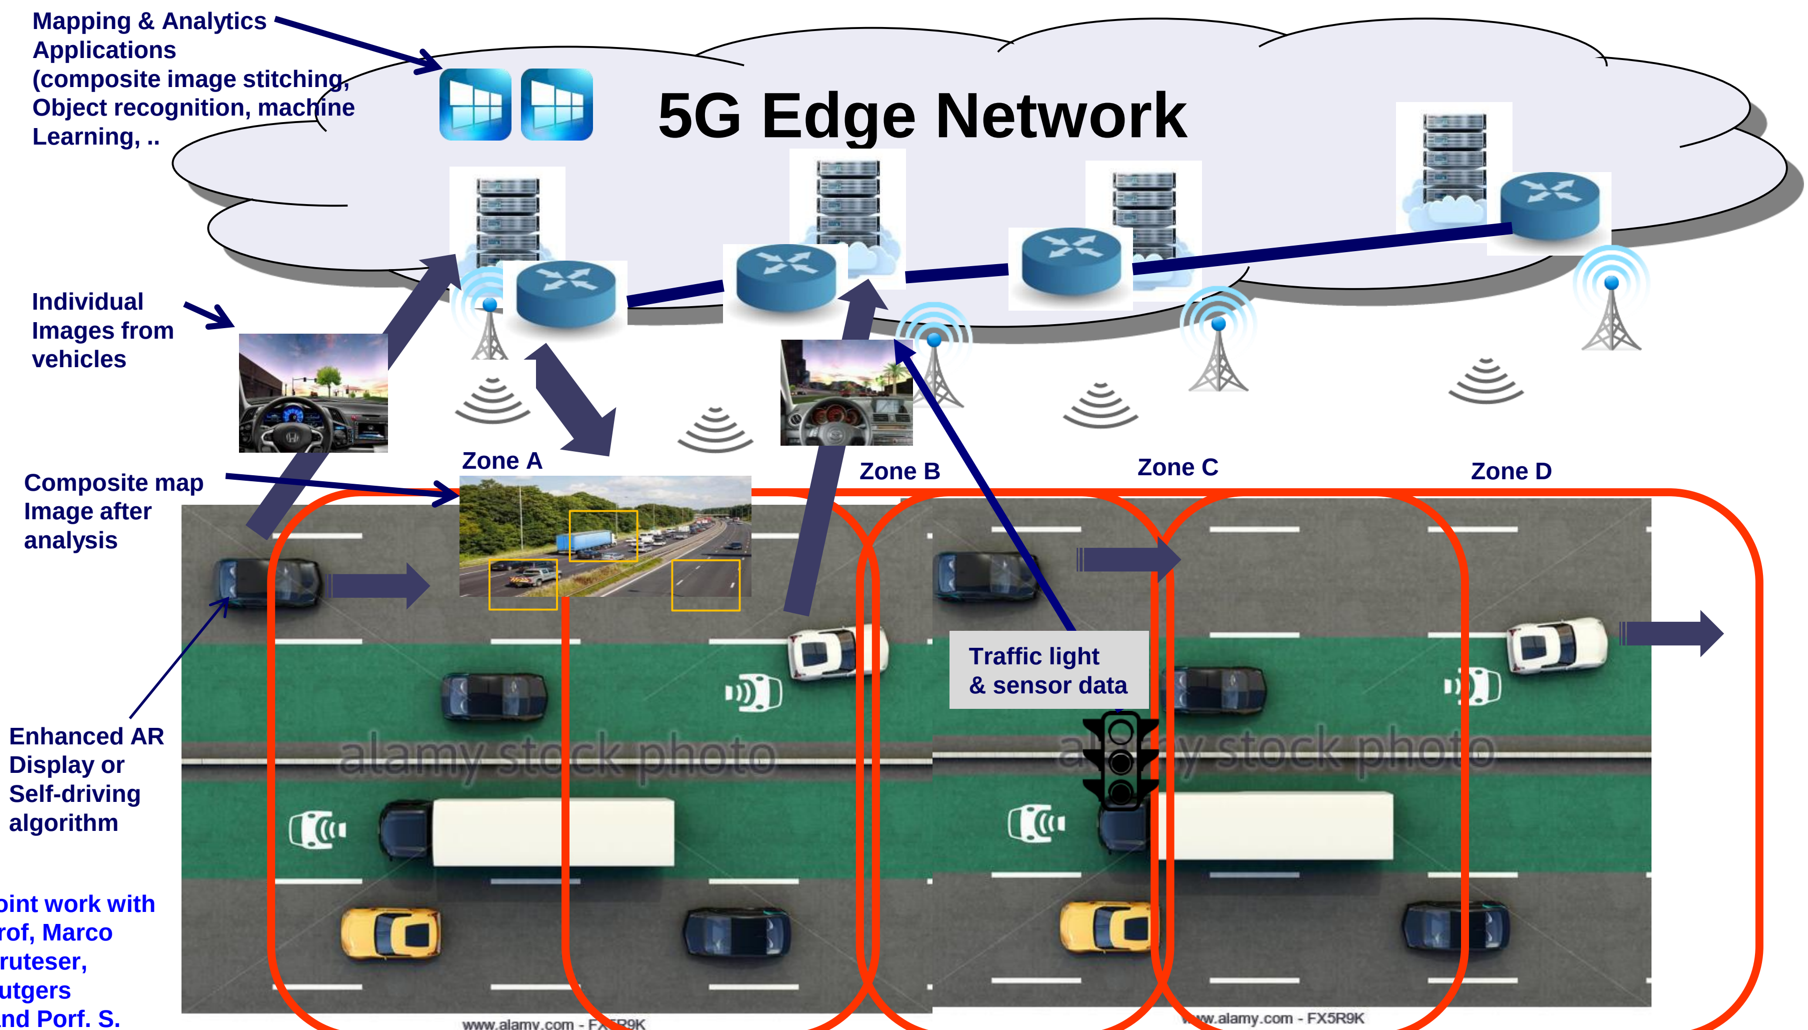

Joint work with Prof, Marco Gruteser, Rutgers
and Porf. S. Dey, UCSD

Mobile Edge Cloud: Collaborative Sensing for Autonomous Driving

Prof. Marco Gruteser

- Explore how vehicles can collaboratively sense their surroundings and gathering accurate vehicle counts and speed estimates from in-vehicle camera footage
- Developed speed estimation techniques from in-car camera video
- Applied to traffic monitoring application, which computes congestion level using vehicle counts and speeds

- We have developed a vehicle-based camera sensing platform and use a deep-learning system (YOLO*) for image detection.
- Our preliminary vehicle count accuracy is about 80% and we are in the process of refining the algorithms.

* http://pjreddie.com/darknet/yolo

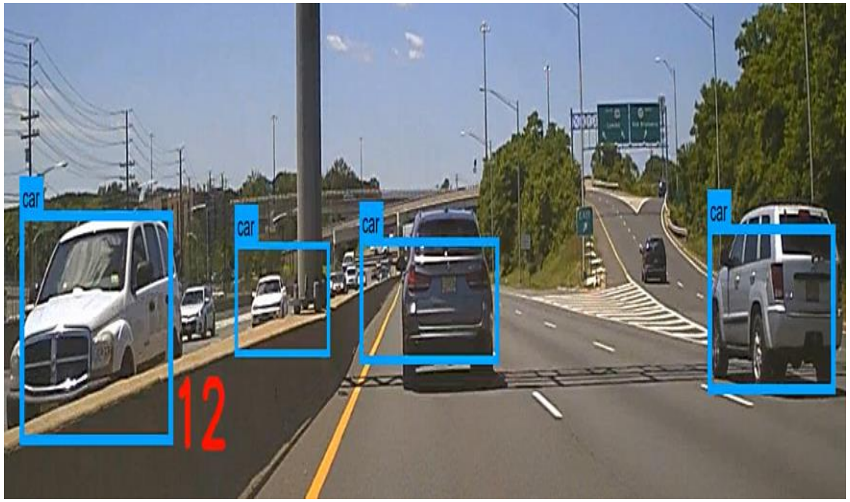

COSMOS Testbed

COSMOS: Project Vision

- Next-gen mobile expected to migrate from today's smartphones to real-time interaction with the physical world
- Application domains include AR, VR, connected car, smart city (with high-bandwidth sensing), industrial IoT, …
- Requires ultra high BW and low latency tightly coupled with edge computing

Augmented Reality

Smart City + Connected Car

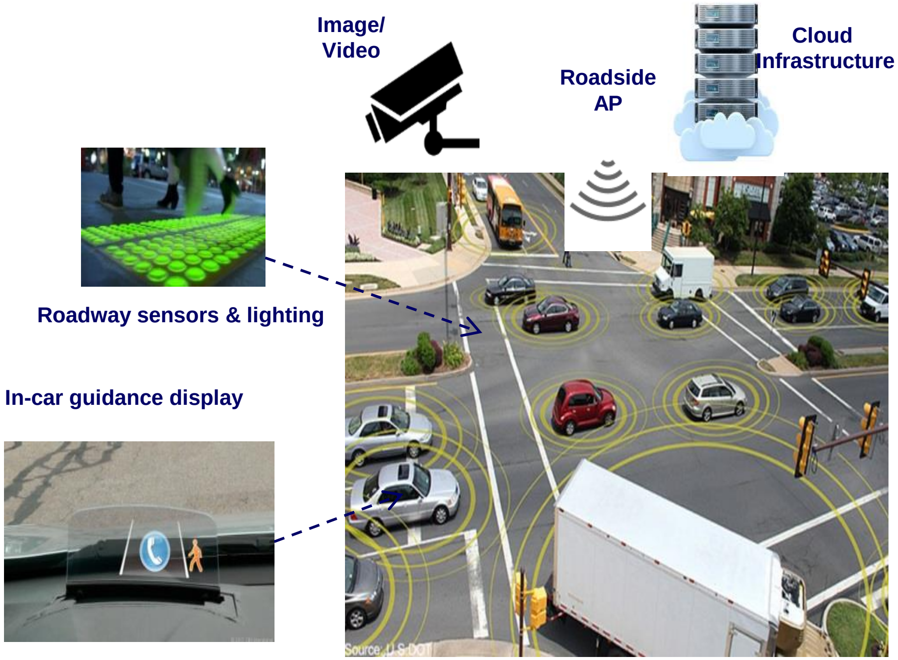

Industrial Control

COSMOS: System Architecture

- COSMOS architecture has been developed to realize ultra-high BW, low latency and tightly coupled edge computing
- Key design challenge: Gbps performance + full programmability at the radio level
- Developed a fully programmable multi-layered (i.e. radio, network and cloud) system architecture for flexible experimentation

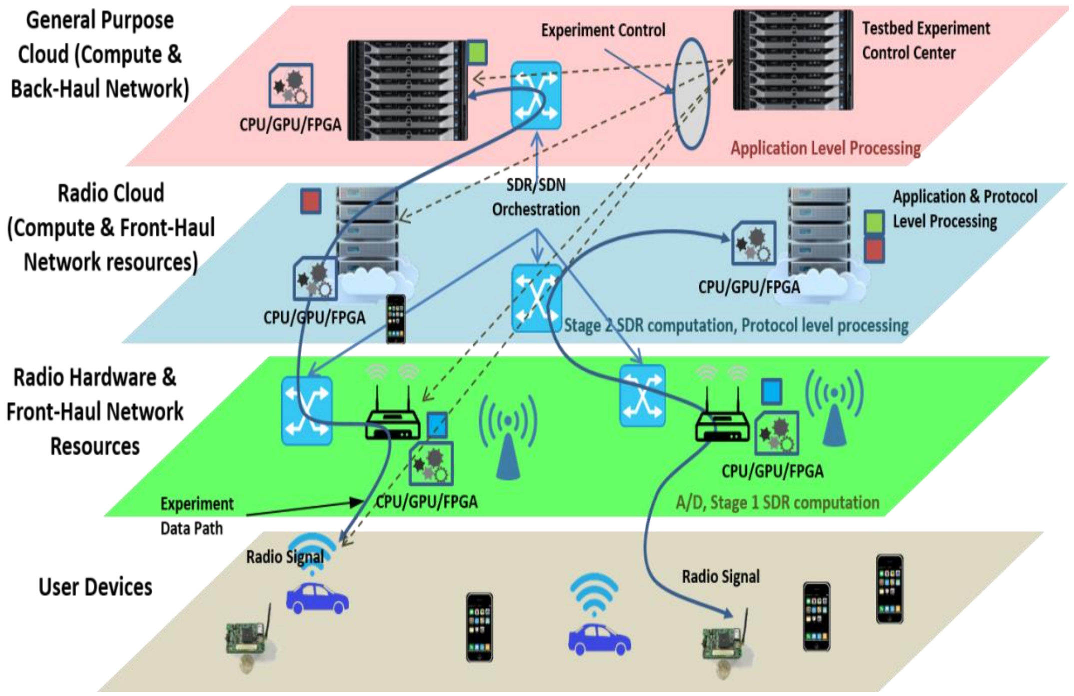

More information at: www.cosmos-lab.org

COSMOS is an open, programmable platform suitable for evaluation and testing of 5G designs and applications

COSMOS: System Implementation

- System design based on three levels of SDR radio node (S,M,L); incl mmWave
- All-optical SDN x-haul with high BW, low latency WDM switching
- Edge clouds co-located with M,L base stations
- NYC deployment with connectivity through 32 AoA PoP

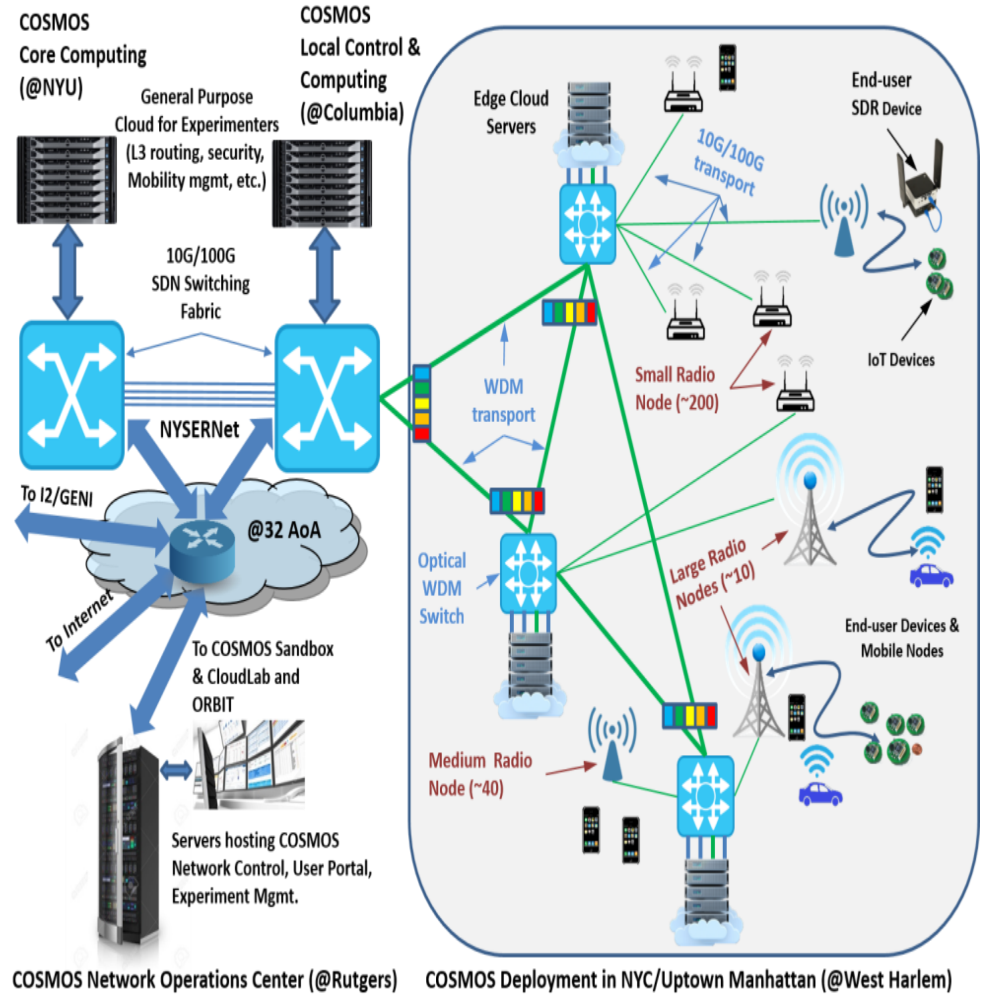

Large Node Sector

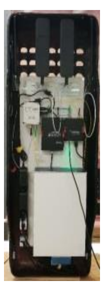

Medium Node (Pole Mount)

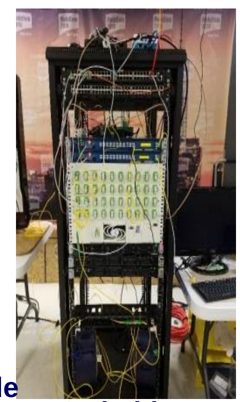

Switching Rack

COSMOS: NYC Deployment

Phased deployment plan: Pilot, Phase 1, …Phase 2

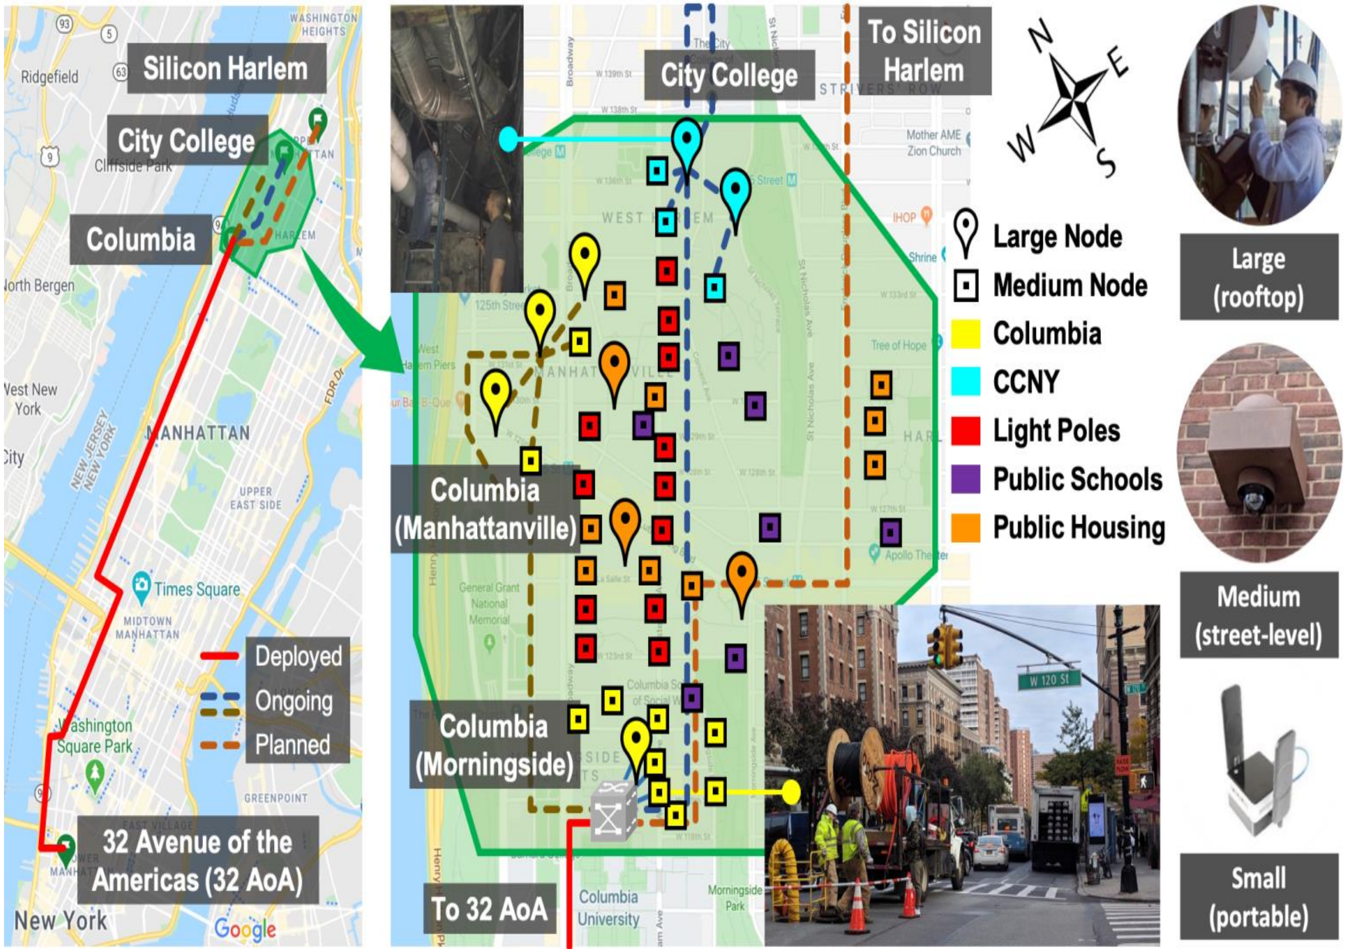

COSMOS: Smart Intersection Experiment

COSMOS Edge Cloud infrastructure deployed at 120St and Amsterdam … being used for ongoing smart intersection experiment