Accelerating Analytics at the Edge

Ada Gavrilovska

Georgia Tech

Computer Science in 6/10ths of a second

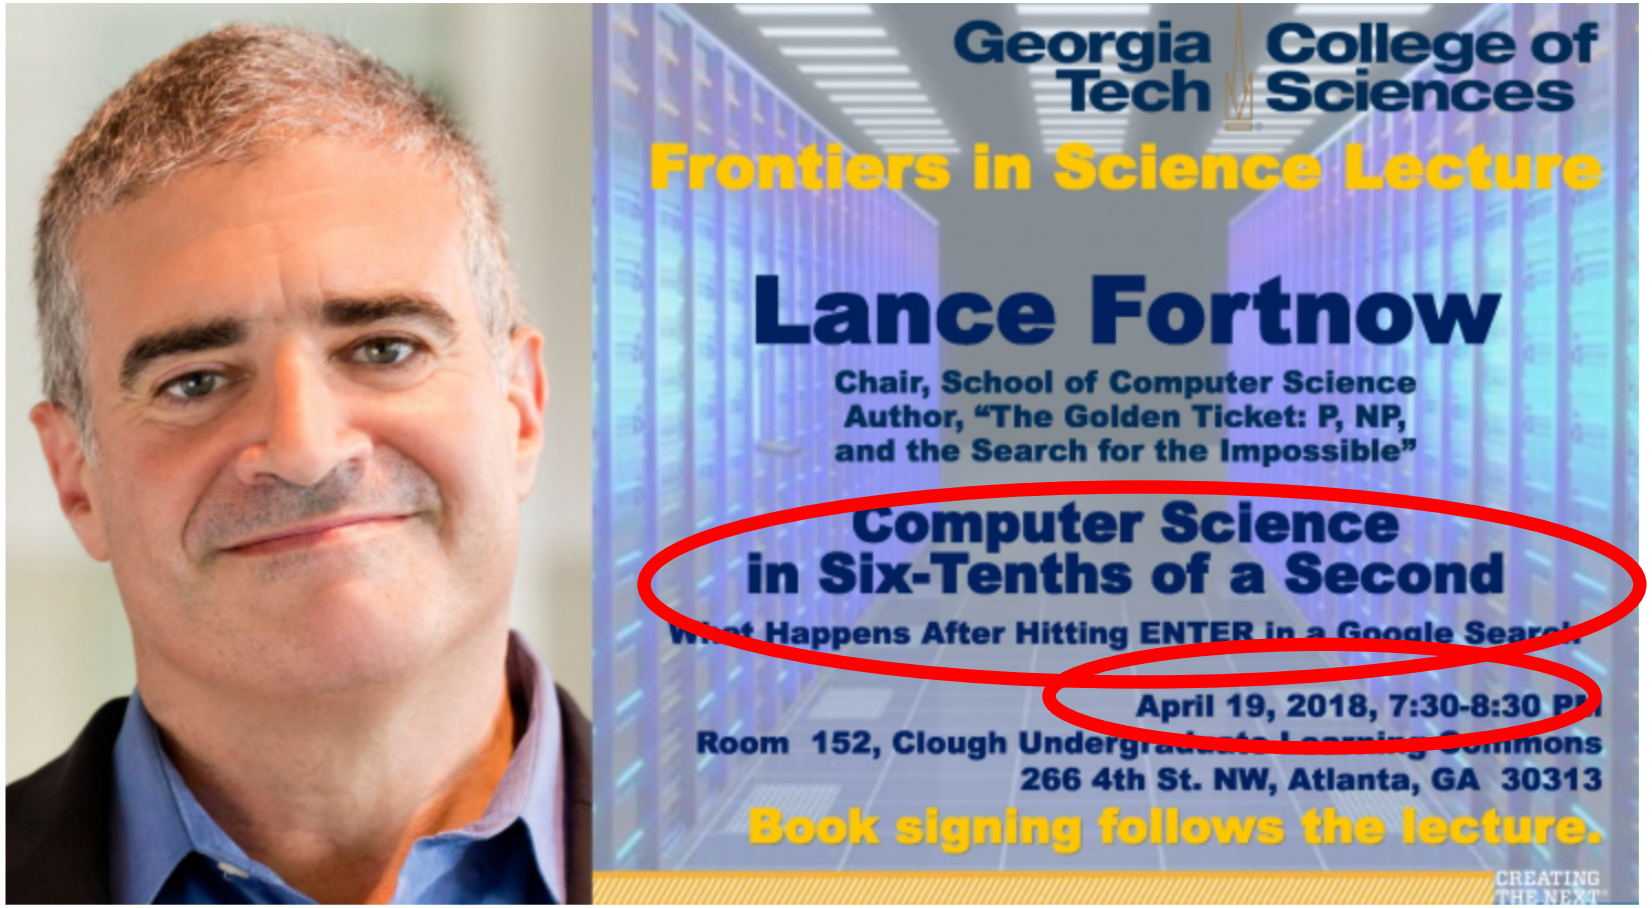

few milliseconds

Computer Science in ~~6/10ths of a second~~

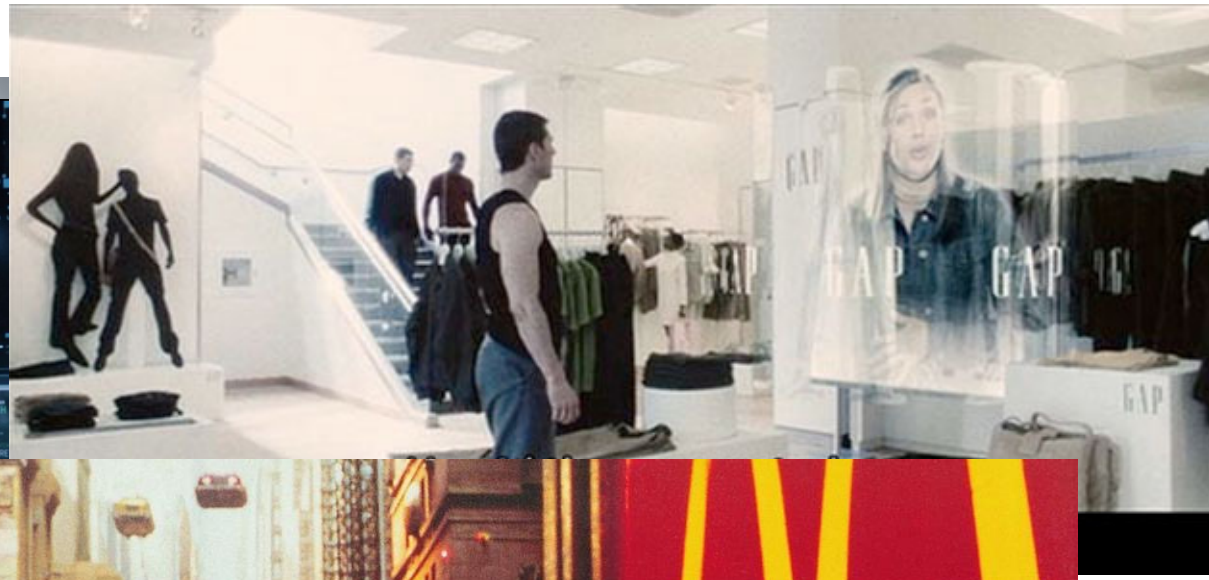

Growth in data movement demand

- New bandwidth-intensive and latency-sensitive workloads

high definition video

AR/VR

SmartCity, automation

Growth in data movement demand

- New bandwidth-intensive and latency-sensitive workloads
- Demand shifts due to new “normal”; connectivity == basic services

high definition video

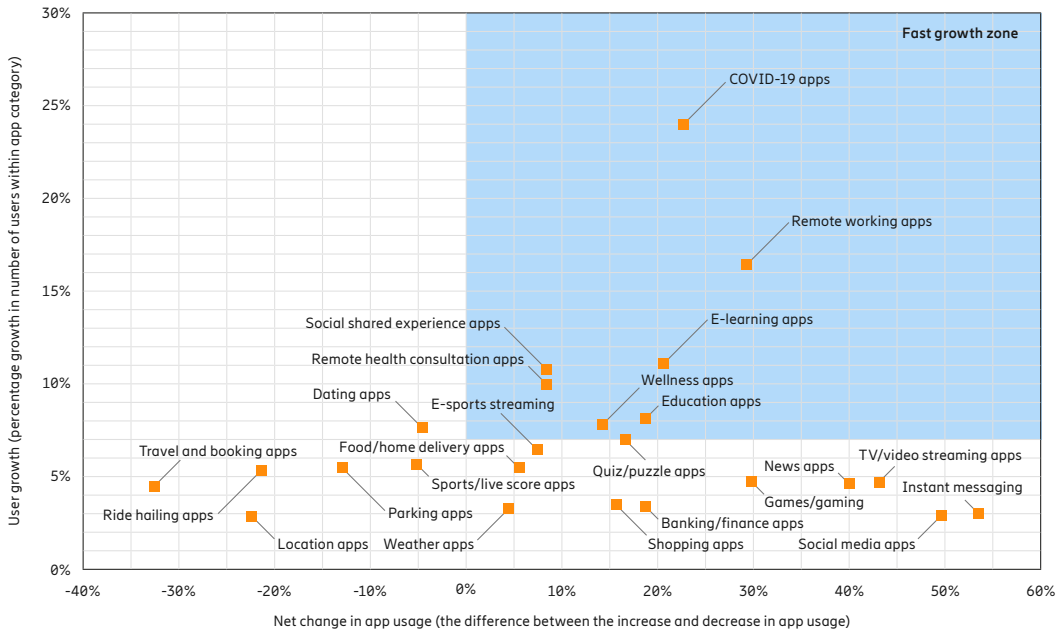

SmartCity, automation

Growth in data movement demand

- Increase in traffic volume, number of devices, wireless

10% CAGR
2017–2022

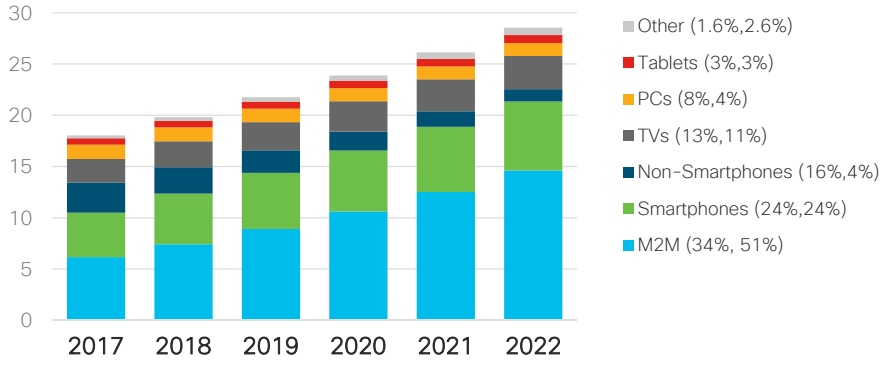

devices

26% CAGR
2017–2022

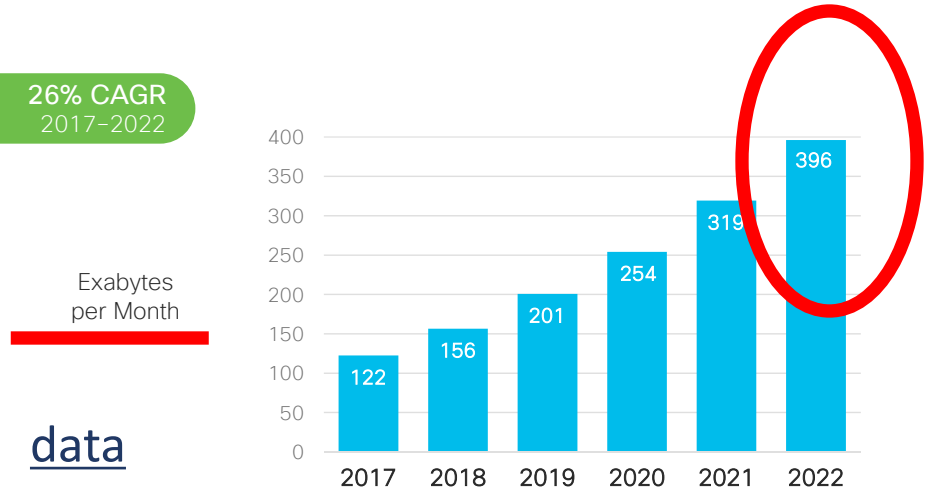

data

26% CAGR
2017–2022

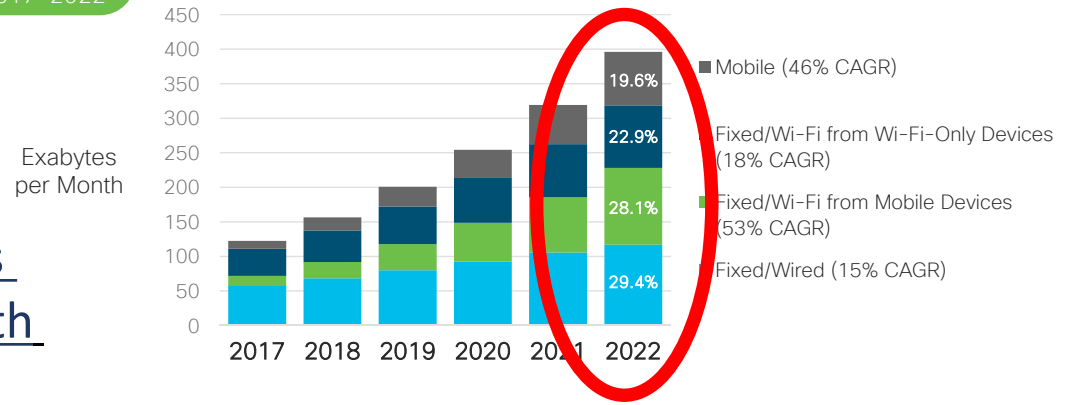

wireless bandwidth

Source: Cisco Virtual Network Index

What does this mean? Just latency and bandwith?

- Past and recent datapoints:
  - 70 TWh to run the Internet, LBNL, 06/2016
  - 50 TWh to run China's mobile network, Huawei, 07/2020
- Updated traffic predictions – no slowdown!
- EB/month cost?
  - wide range based on factors: technology, distance, system scope, … *
  - **1.8 TWh /EB**
  - => **1.2 million tons of CO2** (EPA calculator)
  - **per EB**

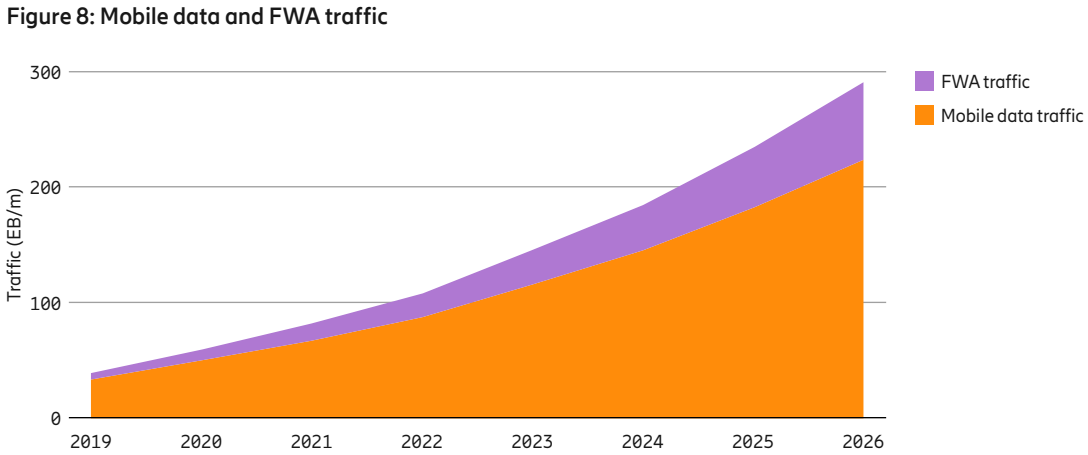

Ericsson Mobility Report (11/2020)

* https://www.wholegraindigital.com/blog/website-energy-consumption/

What does this mean?

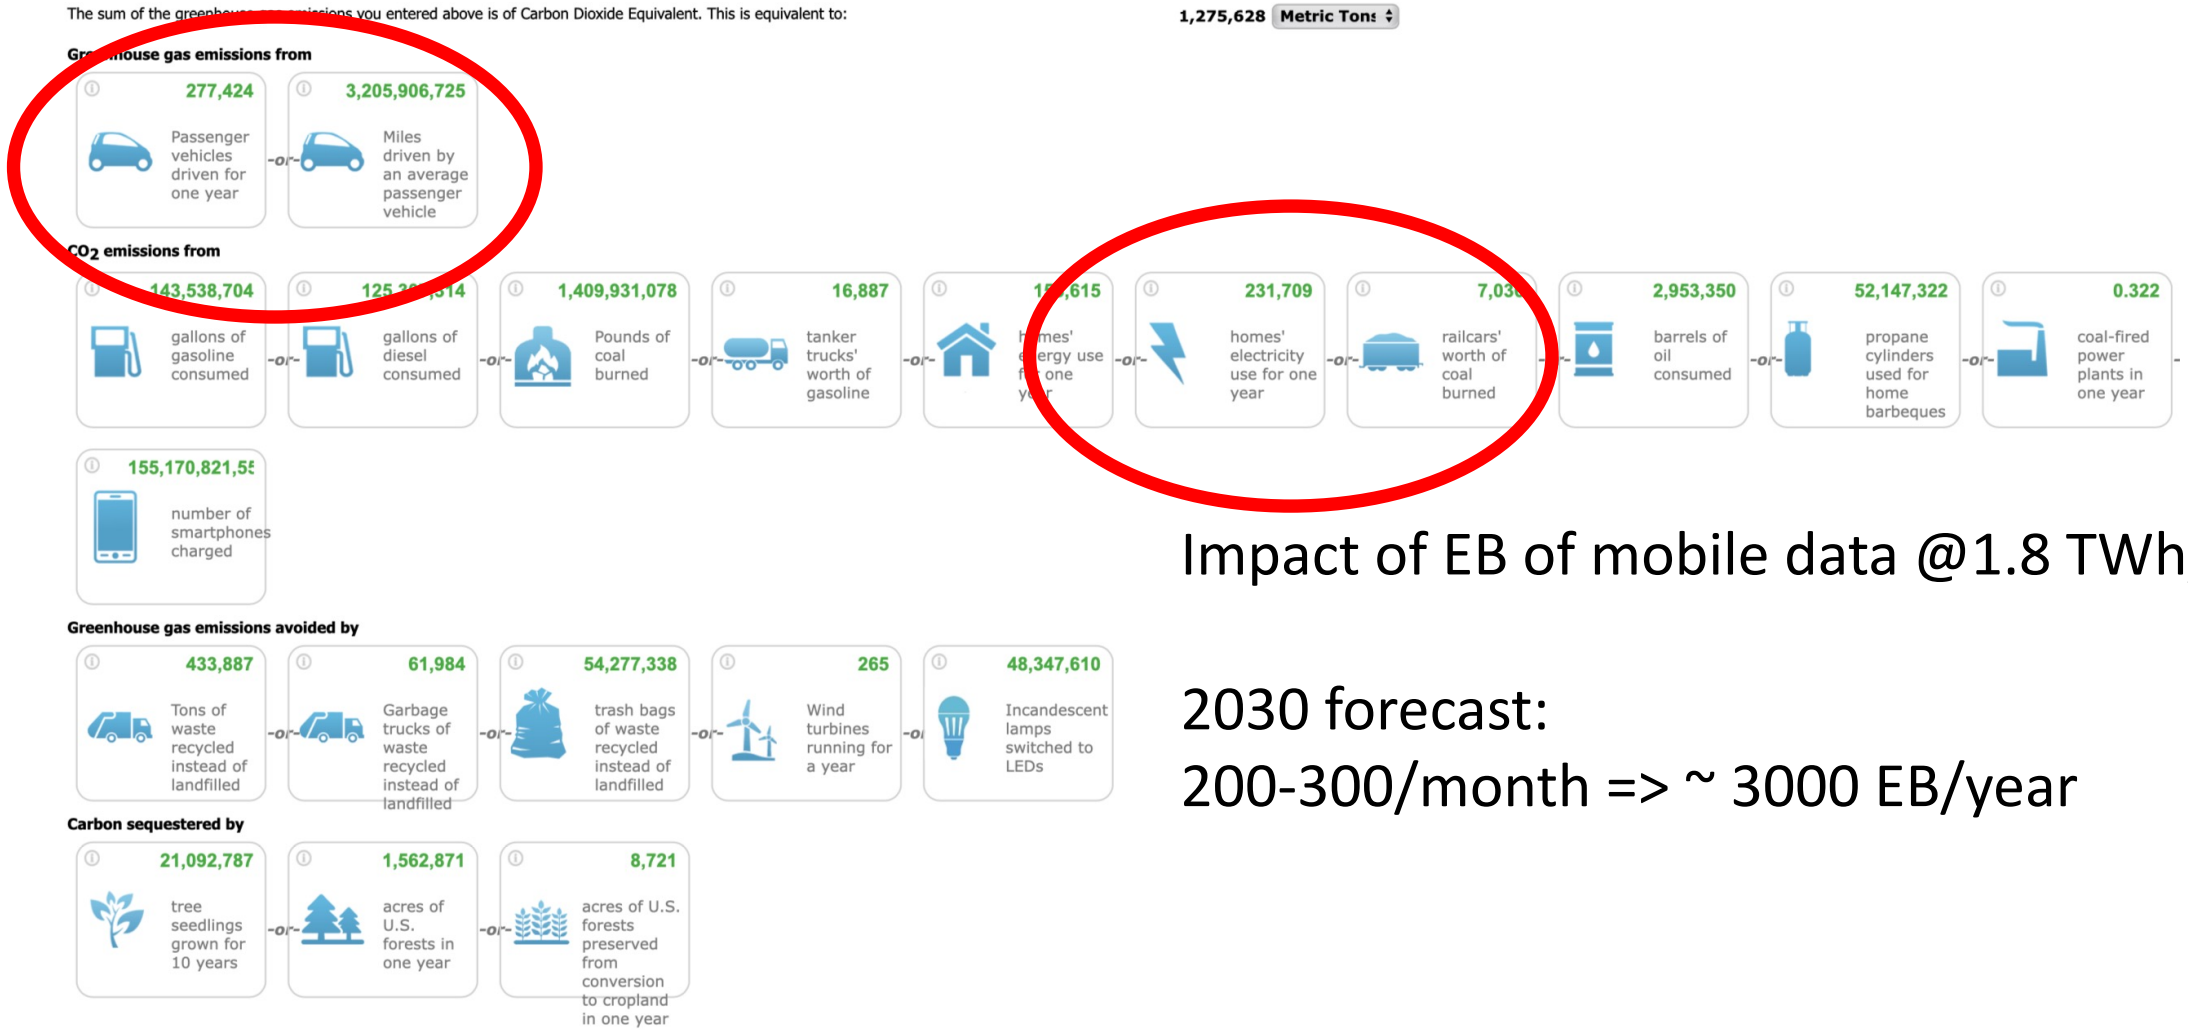

Impact of EB of mobile data @1.8 TWh/EB

2030 forecast:
200-300/month => ~ 3000 EB/year

Edge Computing Opportunities

- Reduce/remove data movement
- 5+G/6+WiFi/…, => new energy-efficient technologies
- Virtualization, software functions/server, …
- Enabler for new applications
- Aligned with UN SDG, Exponential Energy Roadmap

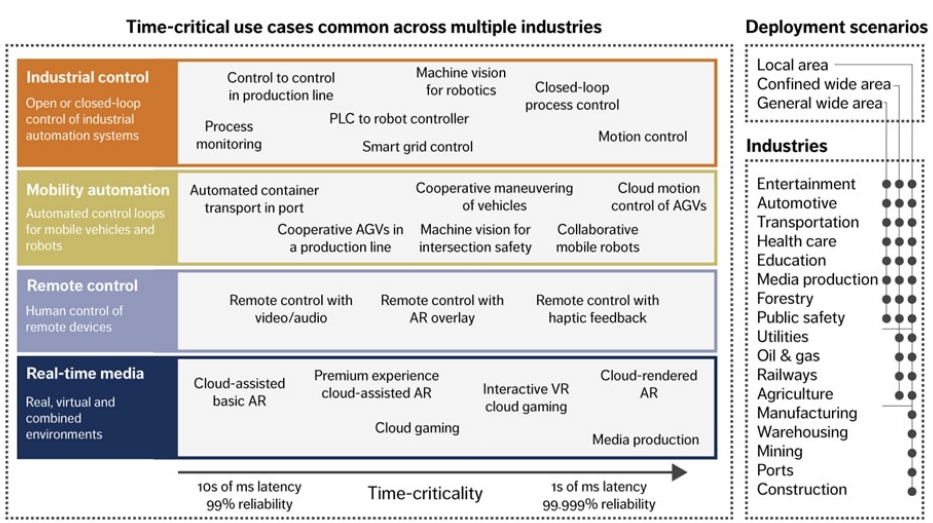

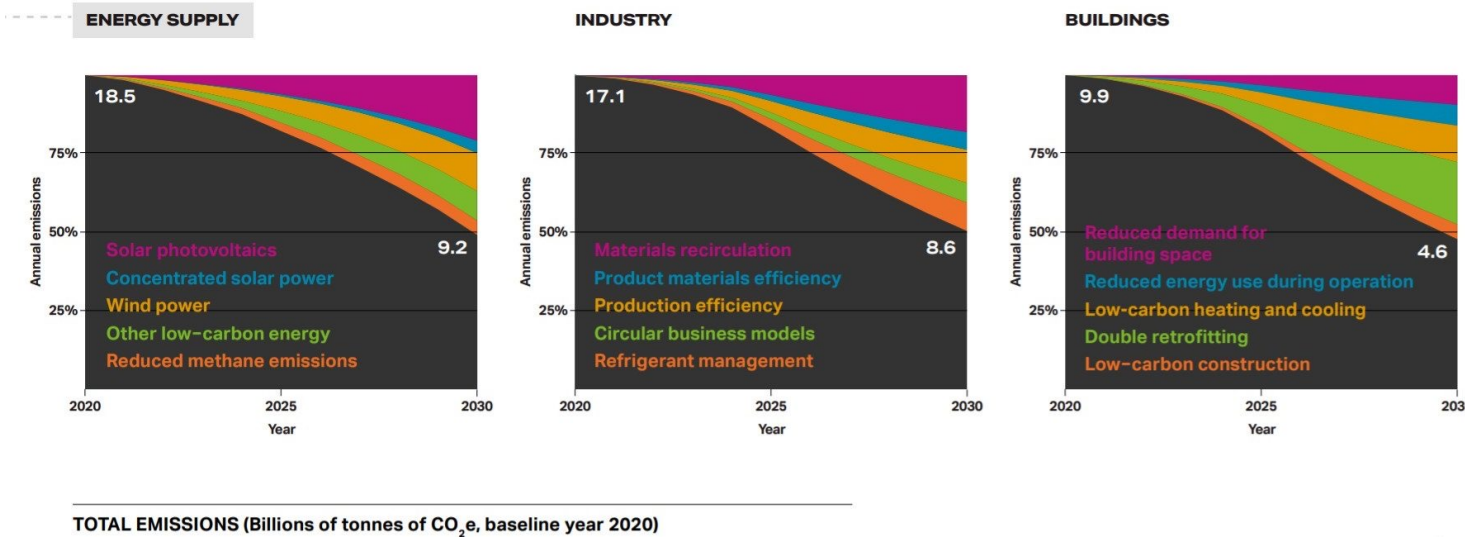

Emerging Systems Lab


Ada Gavrilovska


Greg Eisenhauer


Ketan Bhardwaj


Pradeep Fernando


Thaleia Doudali


Ranjan S. Venkatesh


Harshit Daga

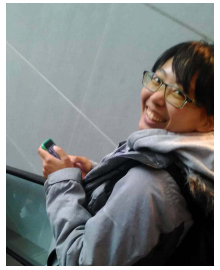
Carol Hsu


Rafael Oliveira


Tony Mason


Jim Choncholas


Jin Heo


Misun Park


Daniel Zahka


Vaibhav Bhasole

Systems Software for Edge Computing

Edge-native technologies

Edge ≠ Cloud

- Lightweight containers for the edge [HotEdge'20, TECHCON'20]
- Security and privacy [HotEdge'18, HotEdge'20, NSF'19-22]
- Systems support for distributed analytics and learning @edge [SEC'19, SOCC'19]

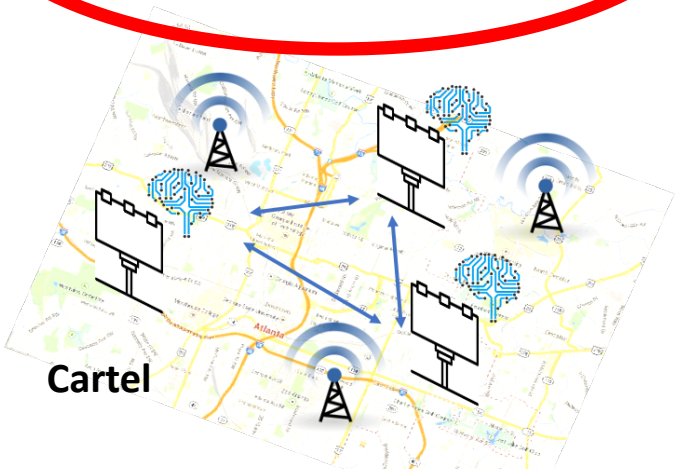

Edge use cases

Applications and platforms

- CDN, AR/VR, Video360, IoT analytics, Visual computing
- Edge and emerging hardware: in-network and in-storage accelerators, NVM, …
- Mostly under SRC JUMP ADA center [SRC'19-23]

MEC and 5+G

Are MNO stacks edge ready?

- Kubernetes-based mobile networks and MEC stack
- Latency-centric orchestration in multi-tenant MEC
- Infrastructure services for edge (MEC-L-DNS, MEC-CDN, …) [HotNets'20]

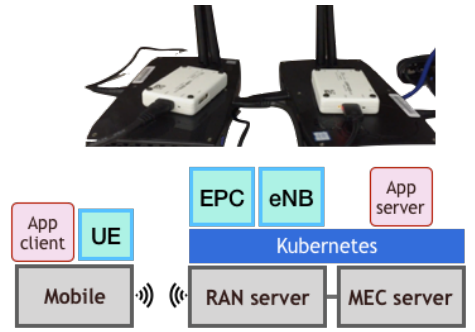

Demo at: https://tinyurl.com/mec-in-a-box

Collaborative Learning for the Edge

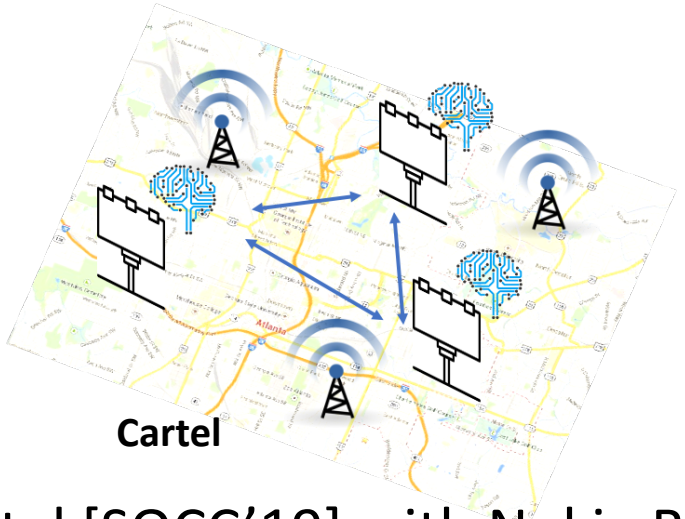

with

Harshit Daga

Cartel [SOCC'19] with Nokia Bell Labs

vmware

Hewlett Packard Enterprise

cisco

Online learning over data from edge

Centralized System

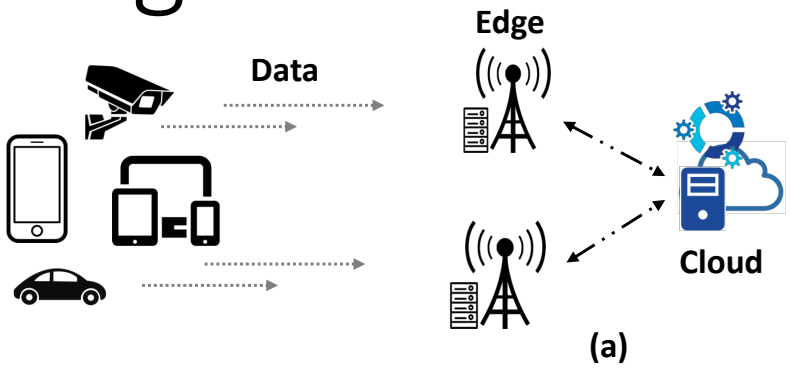

Problems

- **Data movement** is time consuming and uses a lot of backhaul network bandwidth.
- **Distributed ML** across geo-distributed data can **slow down** the execution up to 53X[1].
- **Regulatory constraints** (GDPR)
- Even federated learning requires continuous model updates to be aggregated and propagated

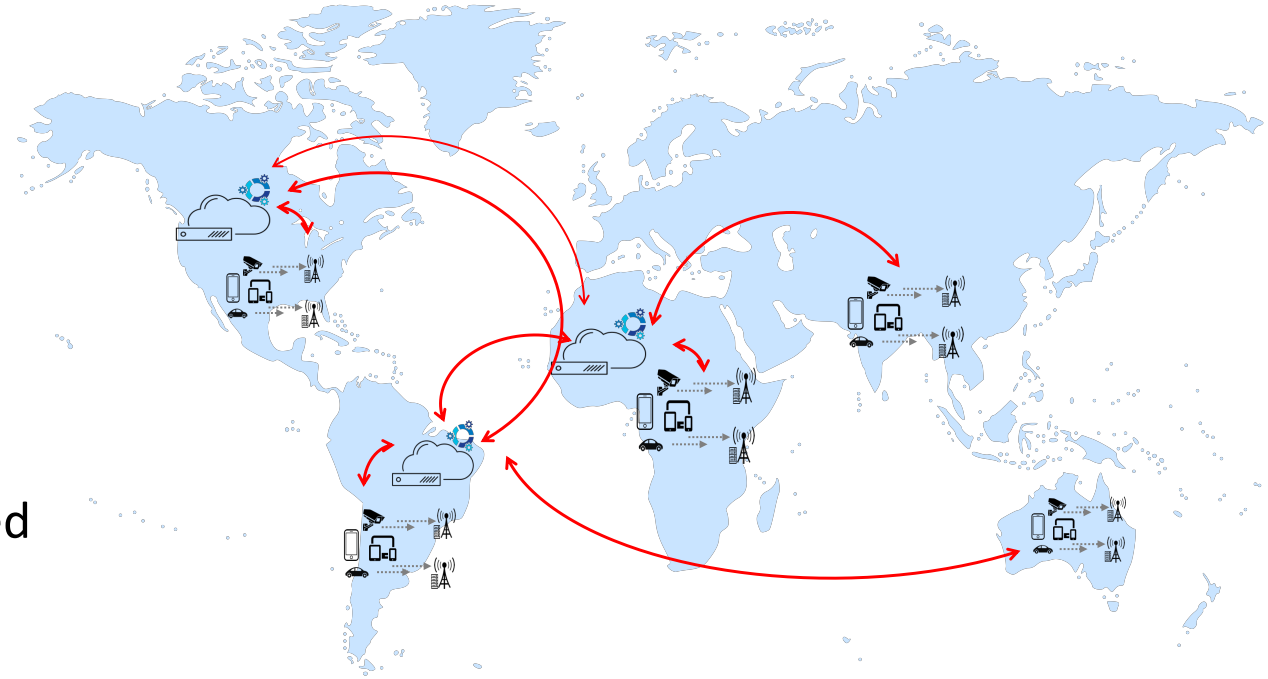

[1] Kevin et al. Gaia: Geo-Distributed Machine Learning Approaching LAN Speeds.

An Alternative Approach

Isolated System

- Train machine learning models independently at each edge, in isolation from other edge nodes.
- The isolated model performance gets heavily impacted in scenarios where there is a need to **adapt** to changing workload.

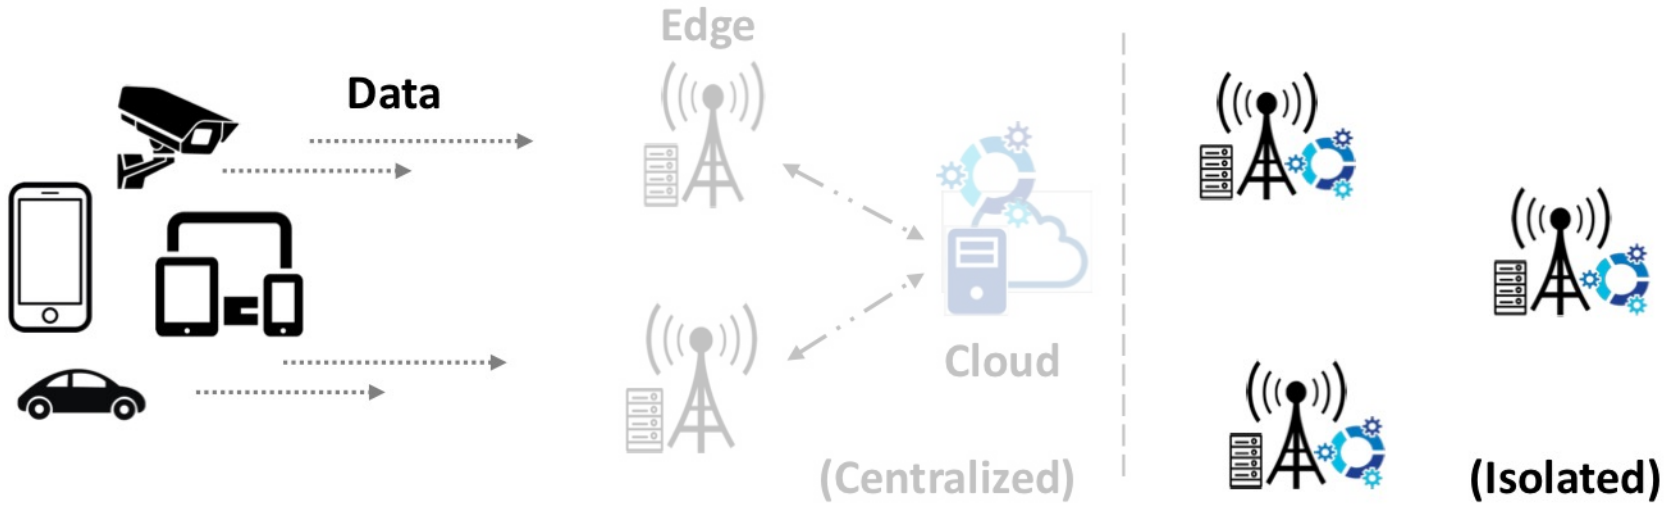

Solution Overview

Cartel : A System for Collaborative Transfer Learning at the Edge

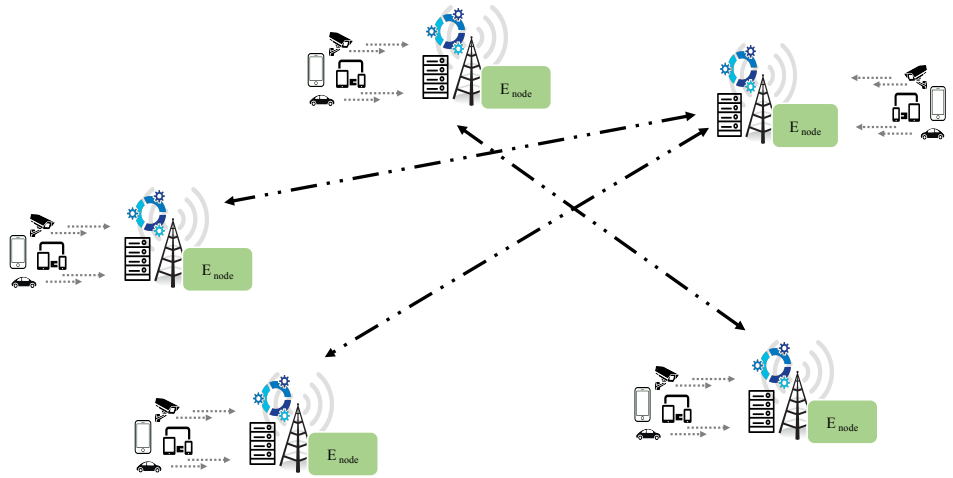

	Centralized	Isolated	Cartel
Lightweight Models	x	✓	✓
Data Transfer	↑	↓	↓
Online Training Time	↑	↓	↓
High Model accuracy	✓	x	✓

- Cartel maintains **small customized models** at each edge node.
- ***When there is change*** in the environment or variations in workload patterns, Cartel provides a jump start to **adapt** to these changes by ***transferring knowledge*** from other ***edge(s) where similar patterns*** have been observed.

Key Challenges & Idea

C1 : When to request for model transfer?

C2 : Which node (logical neighbor) to contact?

C3 : How to transfer knowledge to the target edge node?

Metadata Server (MdS)

- Do not share raw data between any edge nodes or with the cloud.
- => Use **metadata**
  - Statistics about the network
  - Software configuration
  - Active user distribution by segments
  - Estimates of class priors (probability of certain classes), etc.

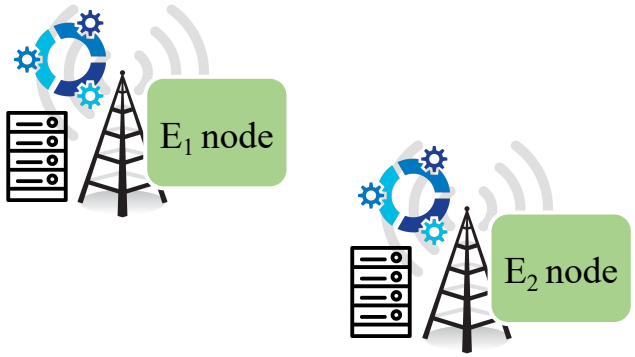

Evaluation

Adaptability to Change in the Workload

Introduction Workload

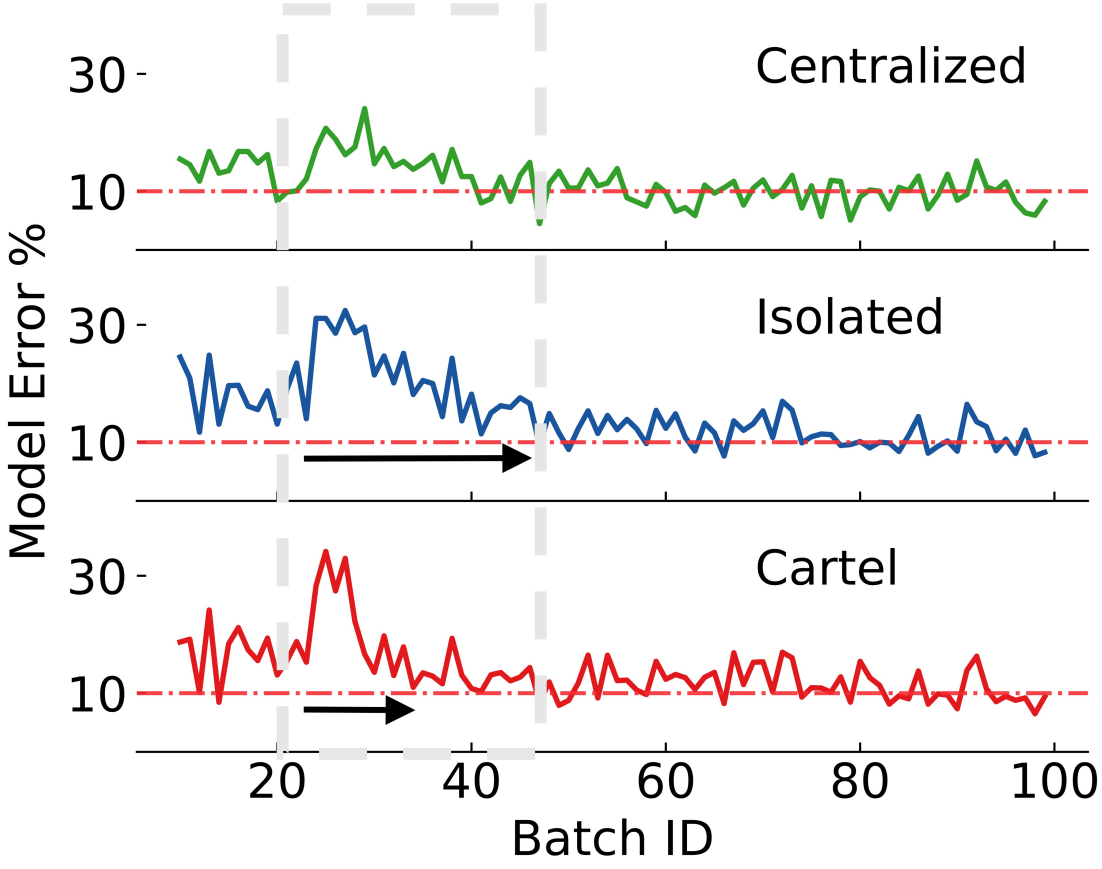

Online Random Forest (ORF)

Evaluation

Adaptability to Change in the Workload

Benefits of Cartel:

- **Adapts quickly to changes** in workload (up to 8x faster compared to an isolated system).
- **Reduces total data transfer costs** significantly (1500x ↓ compared to a centralized system).
- Enables **use of smaller models** (3x ↓) at an edge node **leading to faster training** (5.7x ↓) when compared to a centralized system.

Ongoing extension and evaluation on NNs

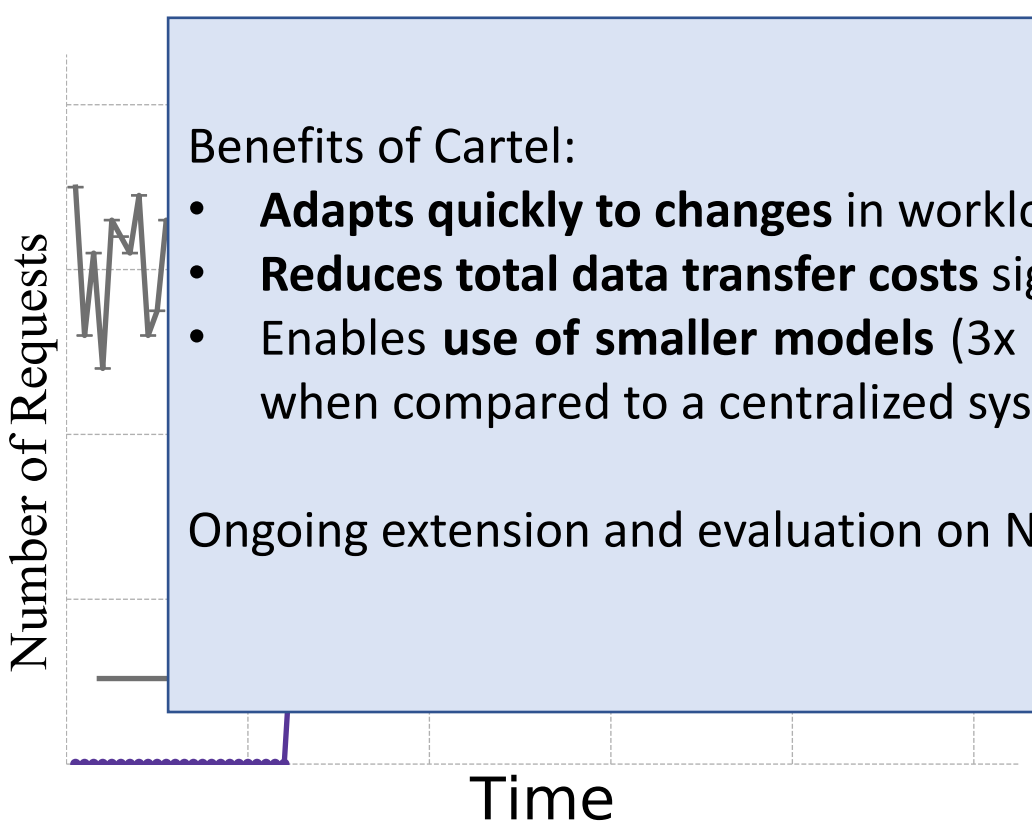

Introduction Workload

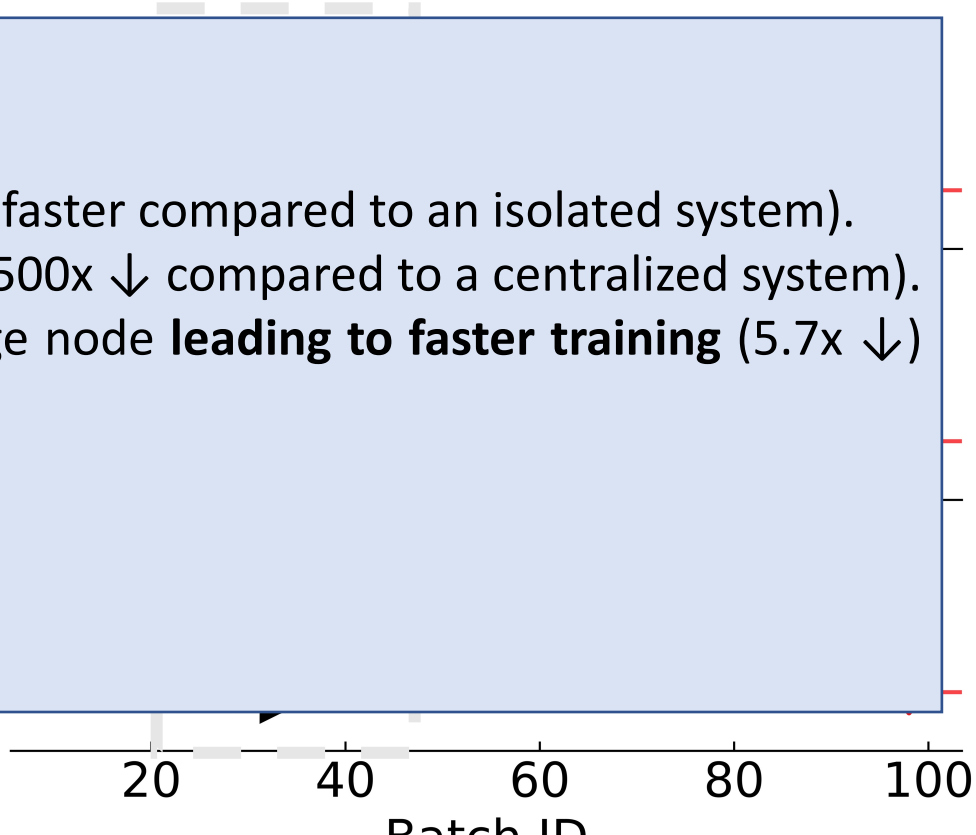

Online Random Forest (ORF)

Emerging applications need DNNs

- ORF, OSVM -> knowledge transfer simple

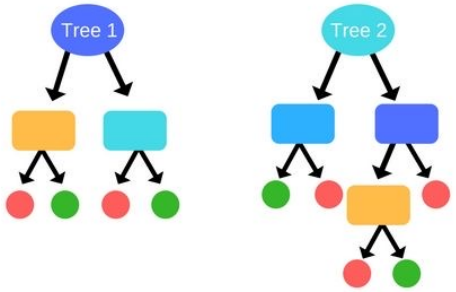

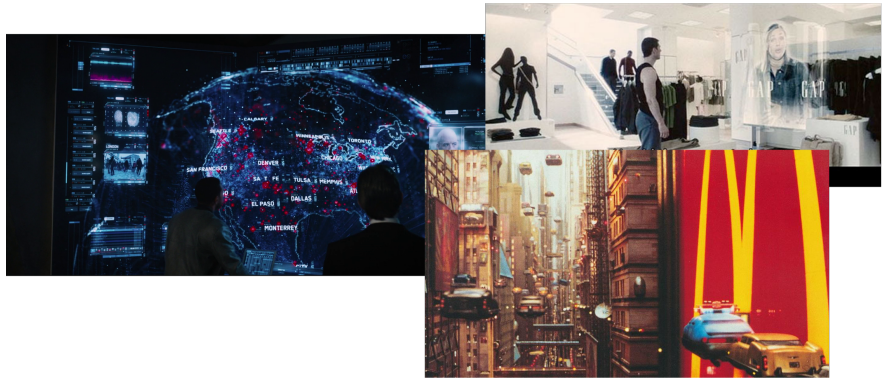

- DNNs, complex and not easily explainable

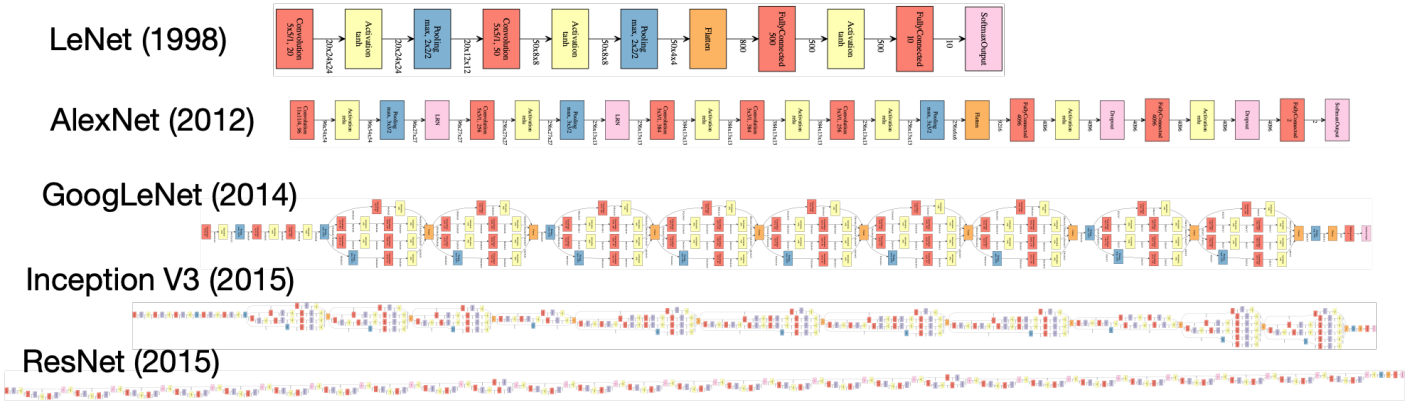

http://josephpcohen.com/w/visualizing-cnn-architectures-side-by-side-with-mxnet/

Collaborative Learning with DNNs

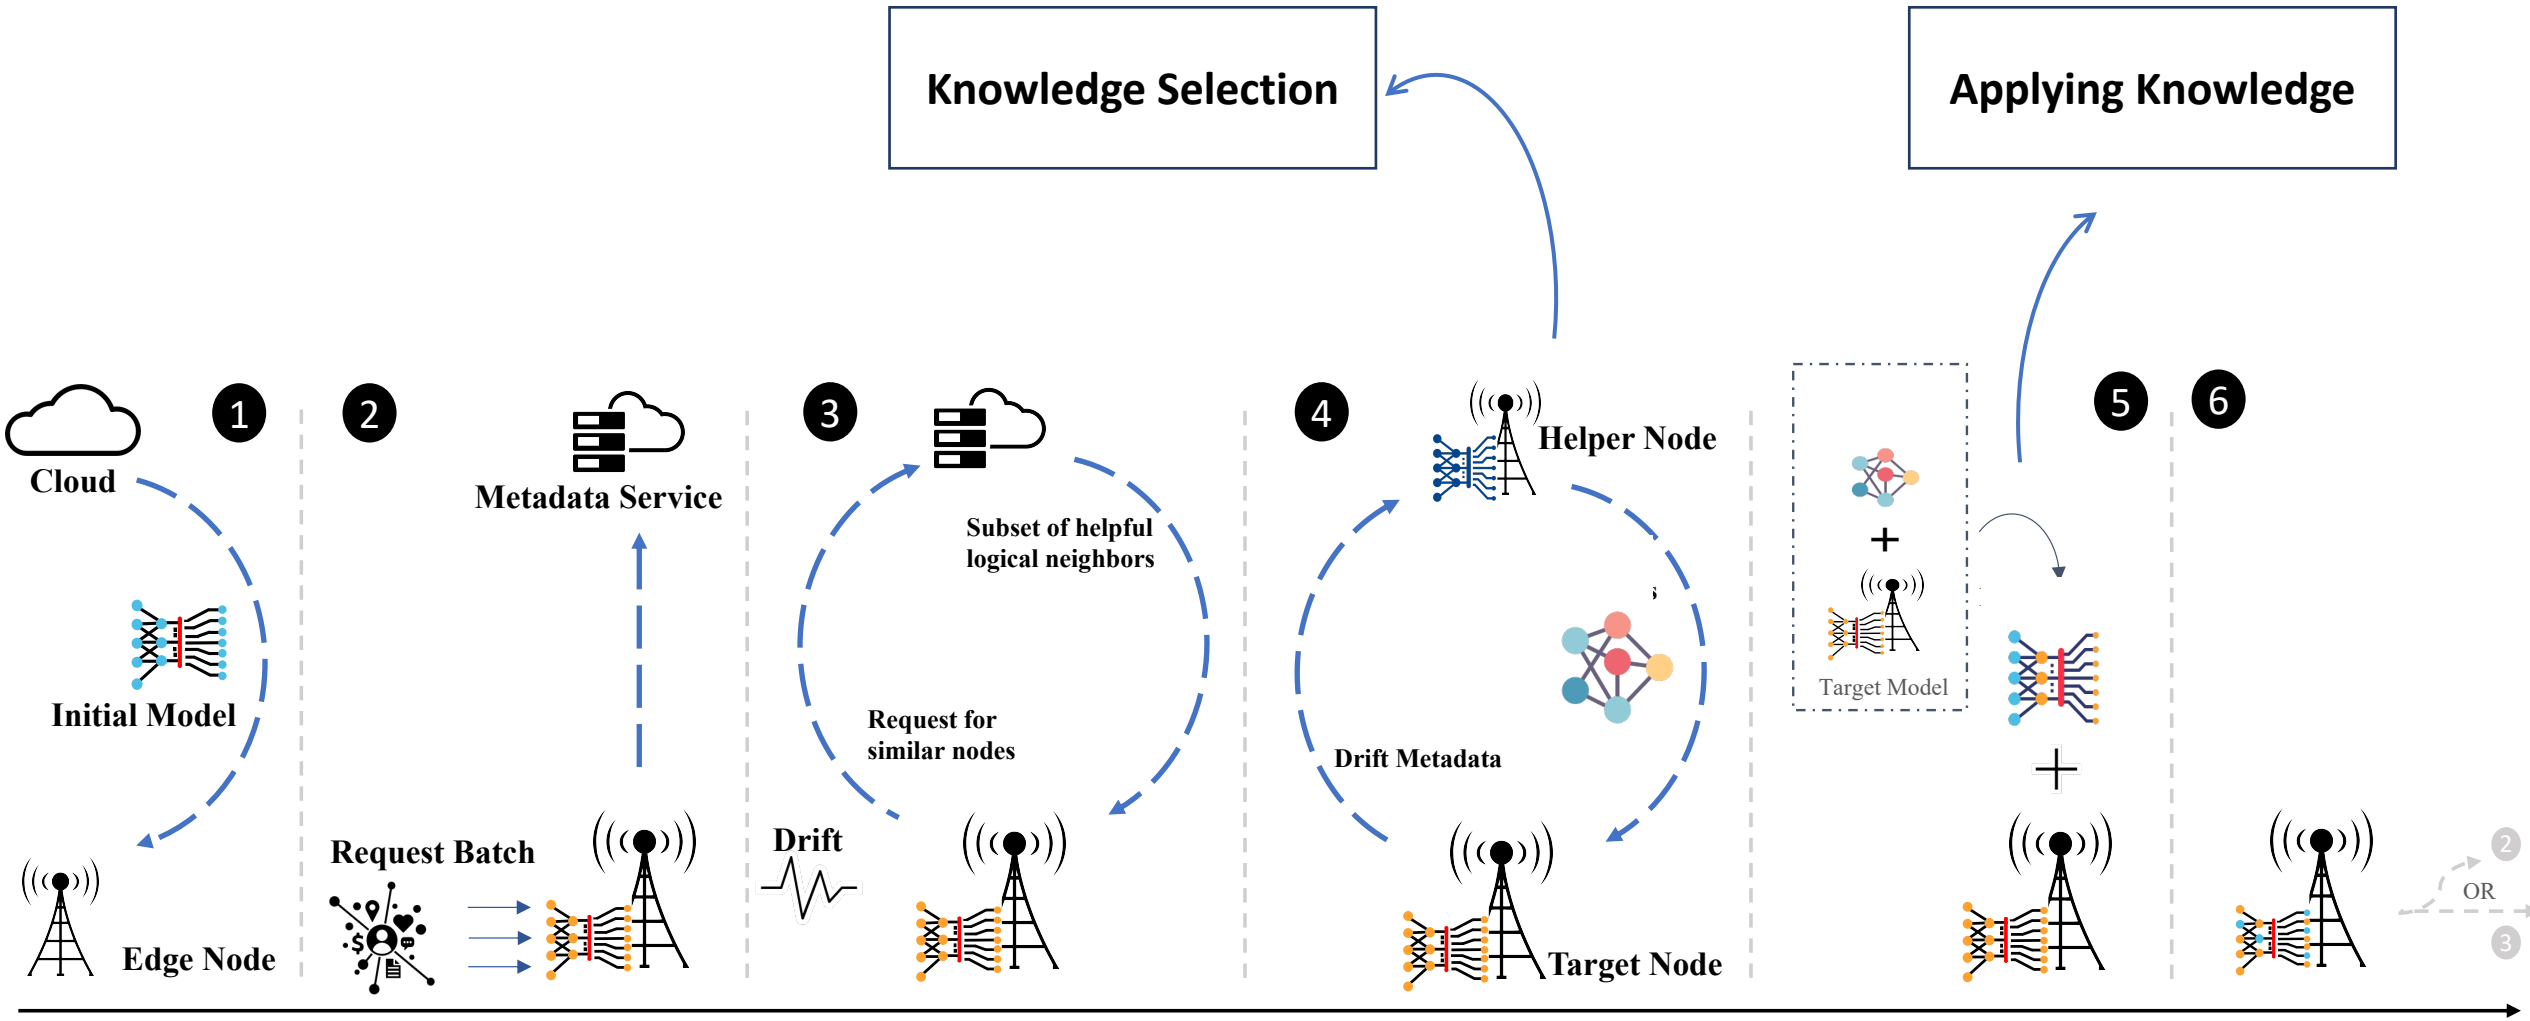

Collaborative Learning with DNNs

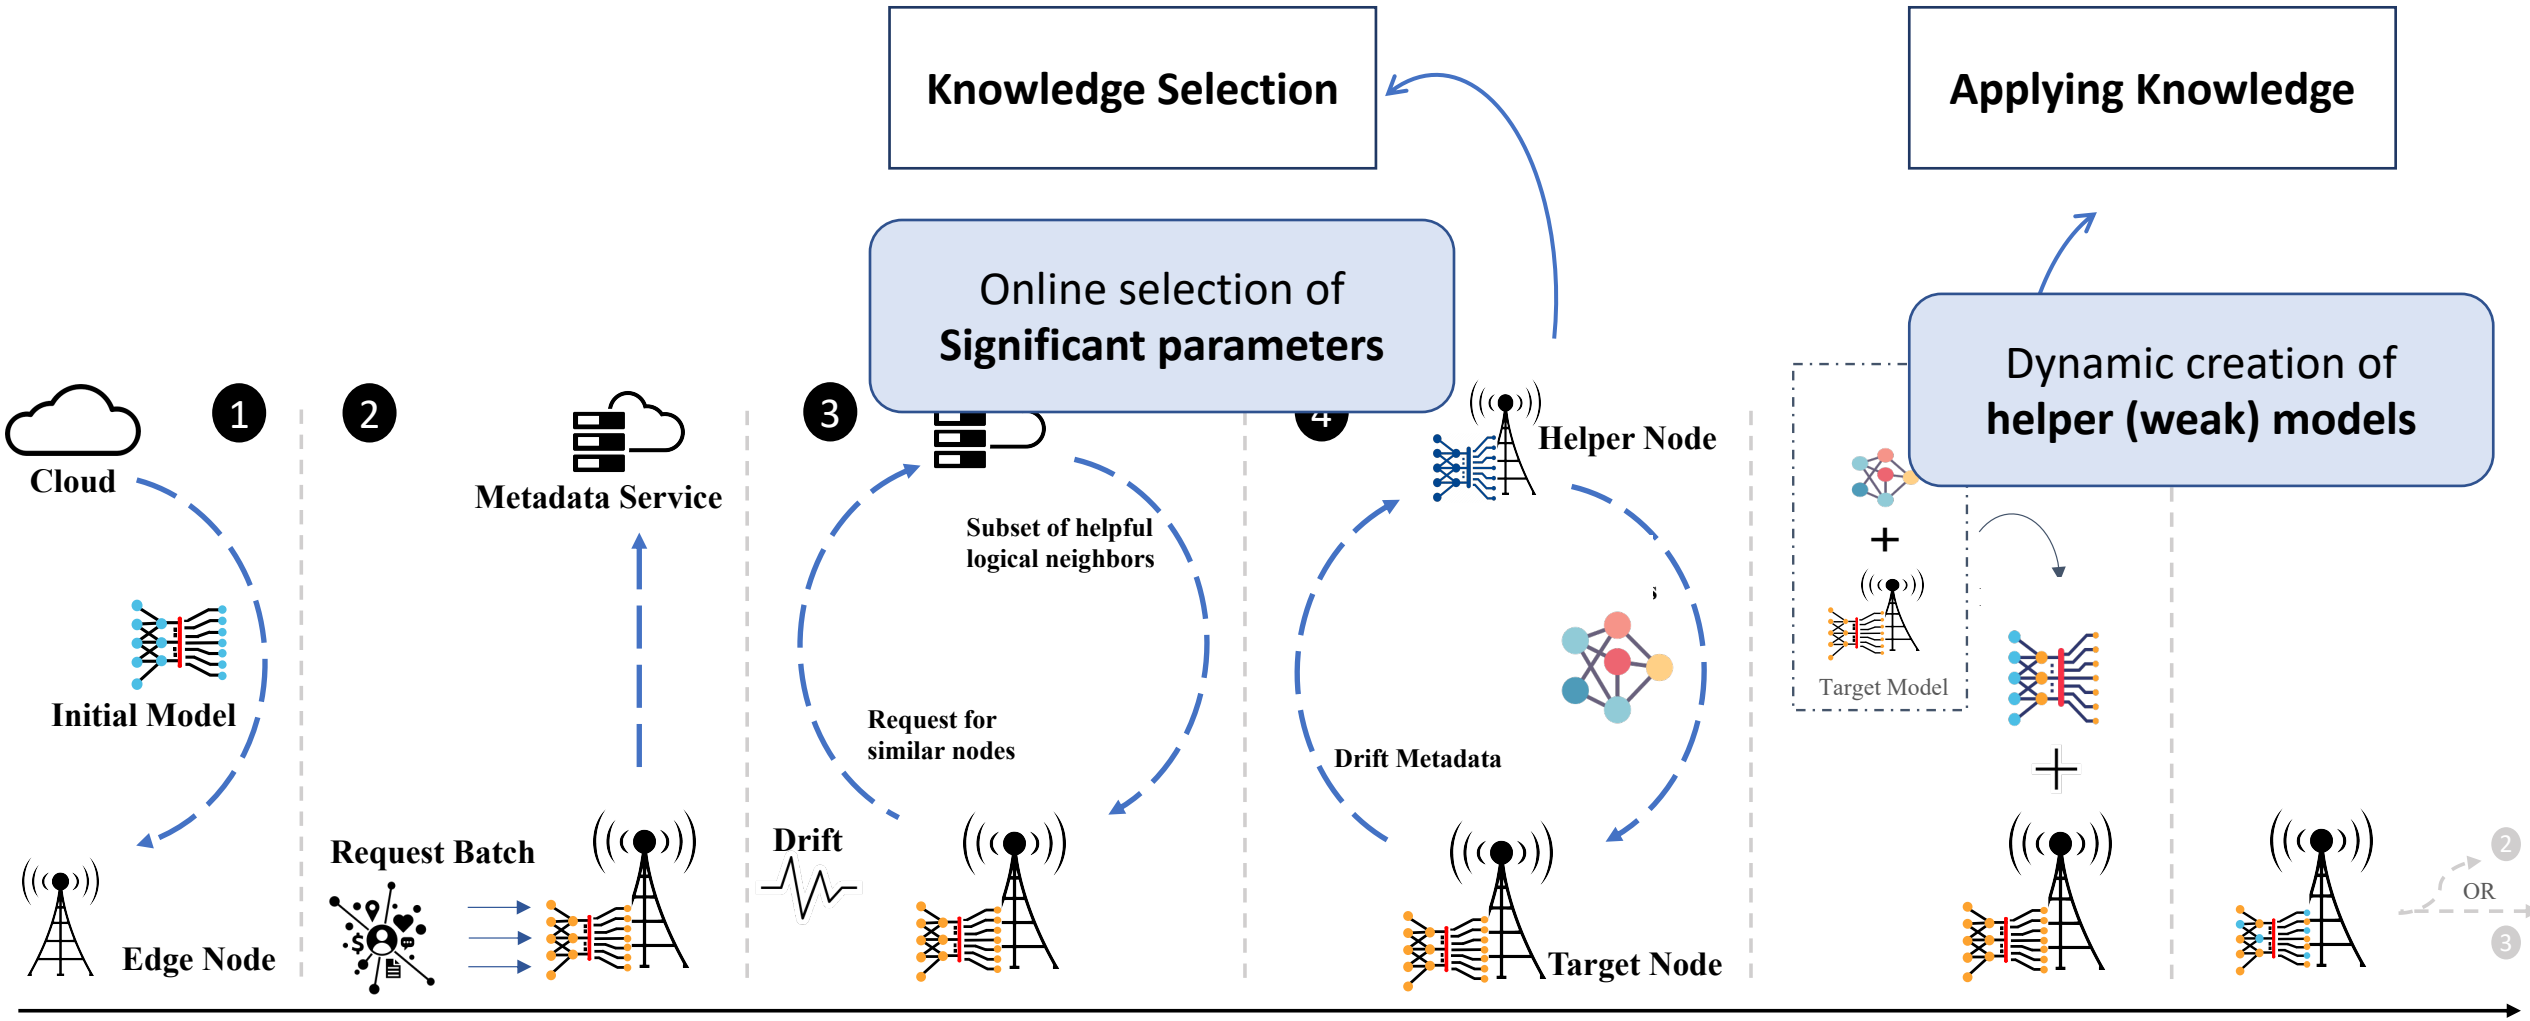

Results – Adaptability

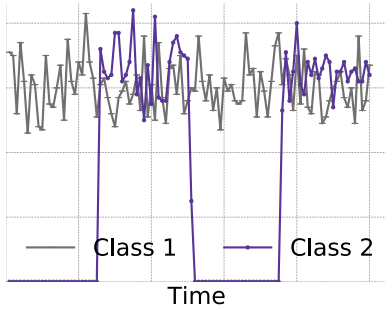

MobileNet

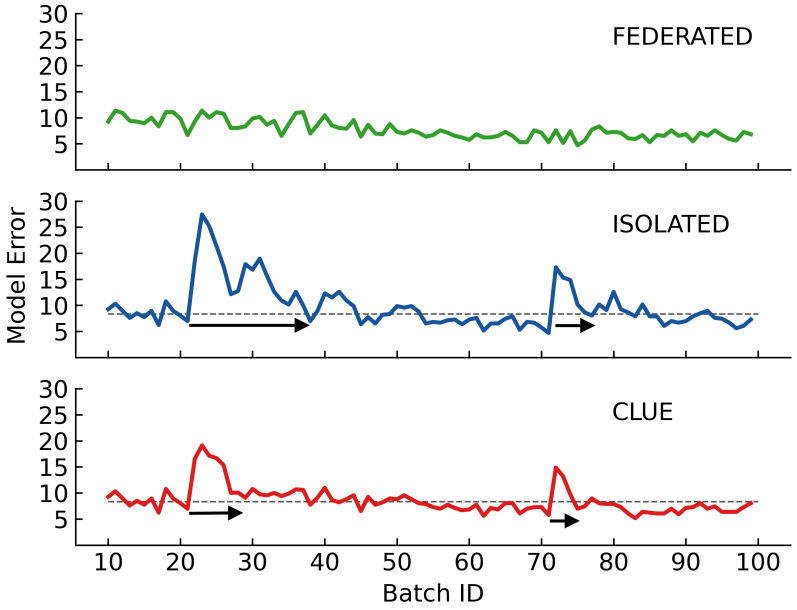

DenseNet

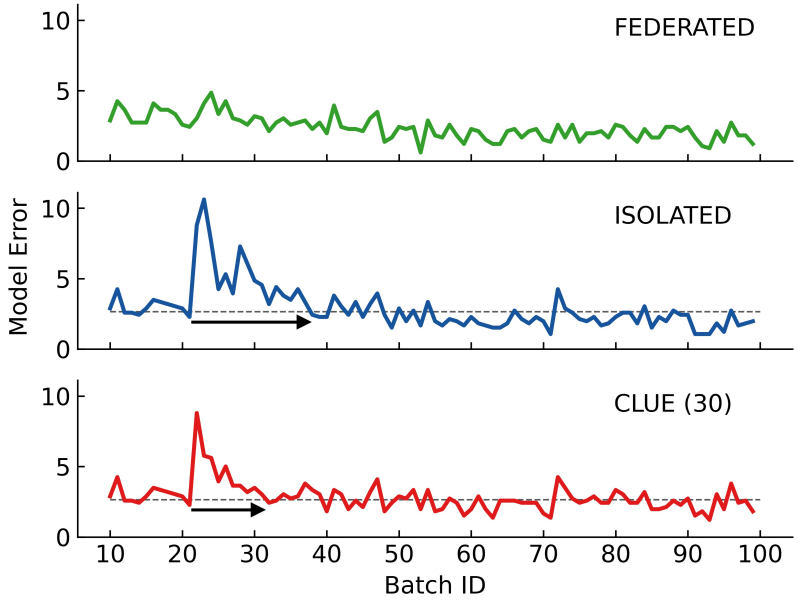

ResNet

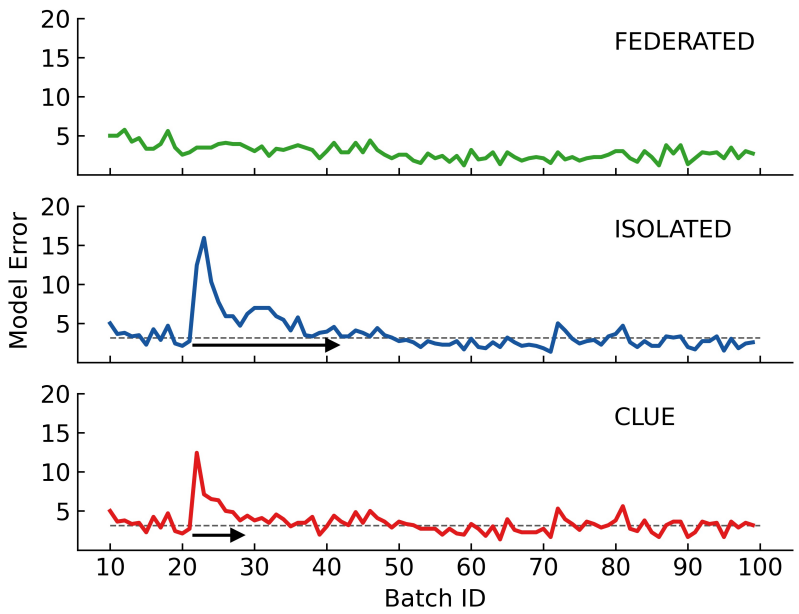

Results – Adaptability & Data Transfer Reduction

MobileNet

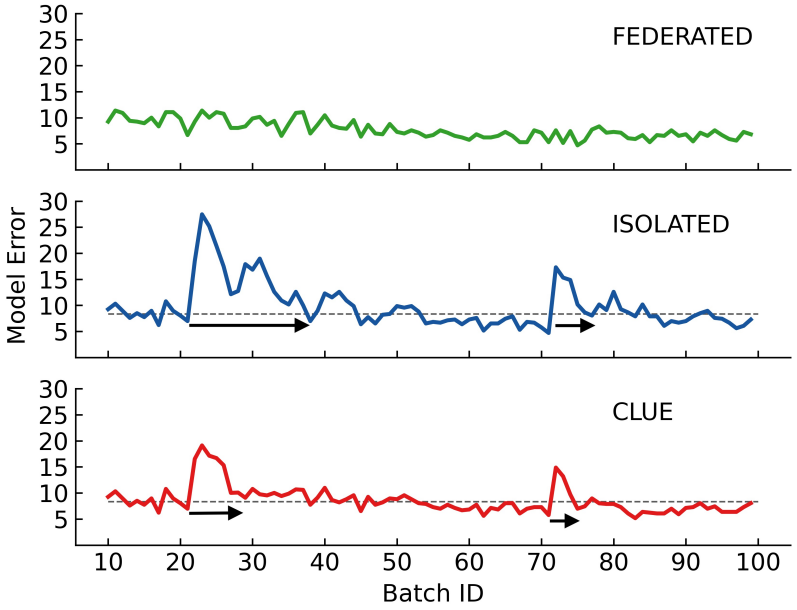

DenseNet

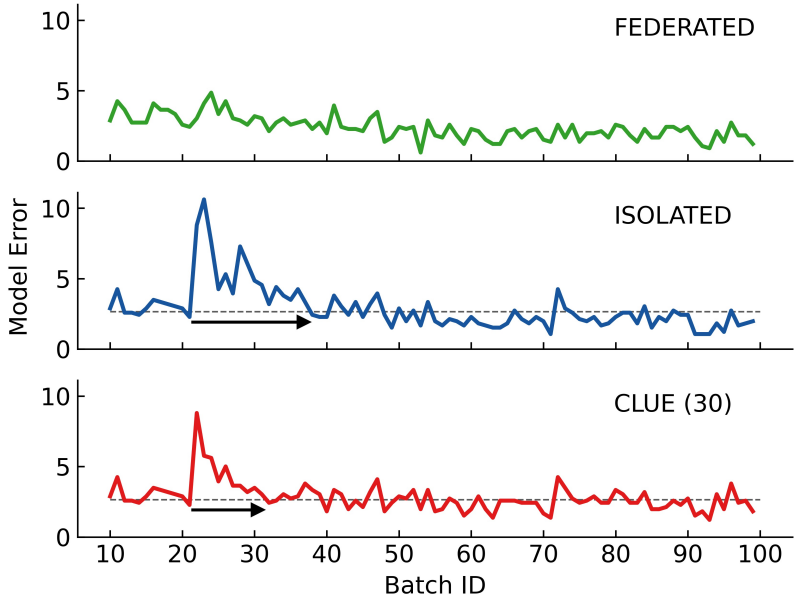

ResNet

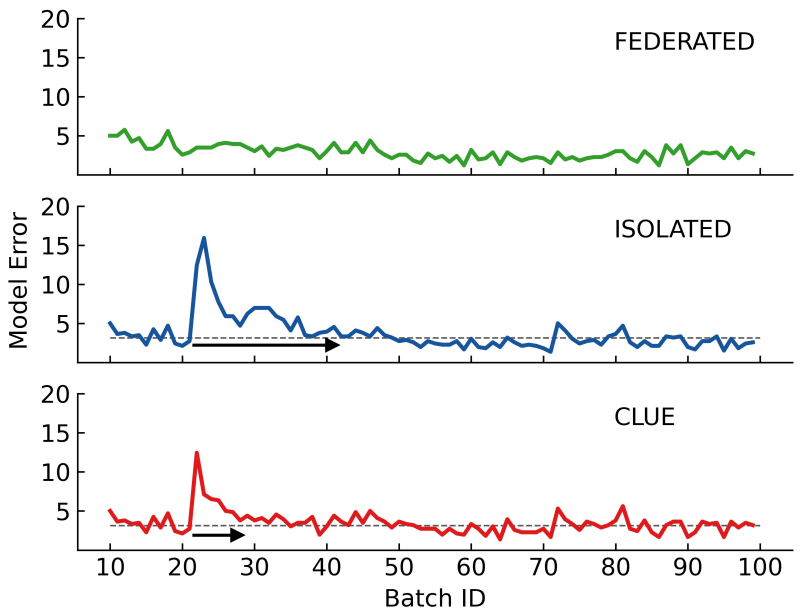

CNN Model	Million Parameters	Federated Learning		CLUE		Data Transfer Gains	
		Out	In	KT (In)	MdS (Out)	No Drift (×)	With Drift (×)
MobileNet	3.54	6.38	7.36	4.57	0.00058	23684	3
DenseNet	8.06	16.18	18.35	8.80	0.00058	59507	4
ResNet	11.69	12.19	13.75	8.55	0.00058	44713	3

Support for Multi-tenancy at the Edge

with

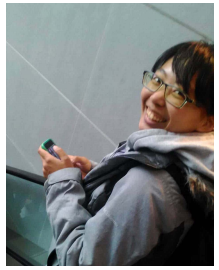
Carol Hsu


Misun Park


Ketan Bhardwaj

and with my Mom's group ☺

vmware NSF SRC ada Applications Driving Architectures

No elasticity, Maximize resource efficiency

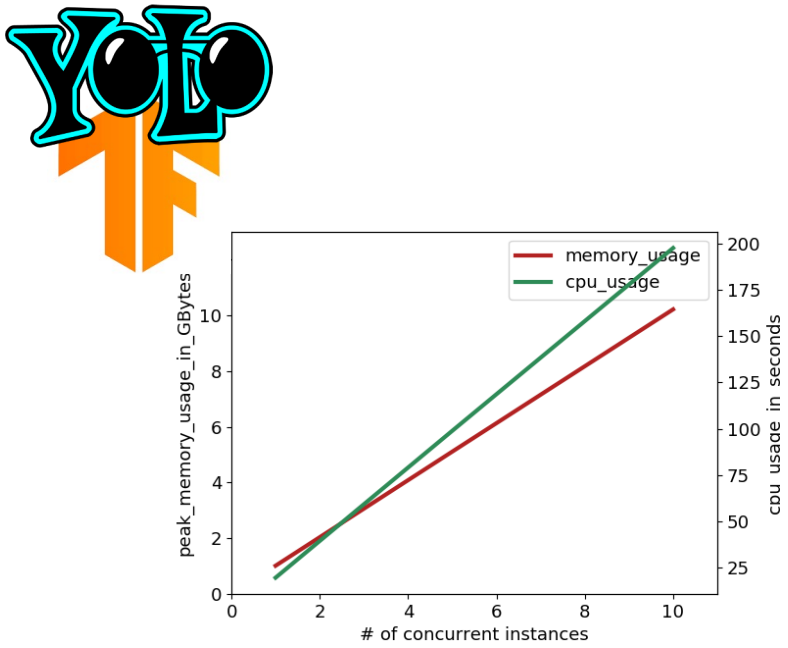

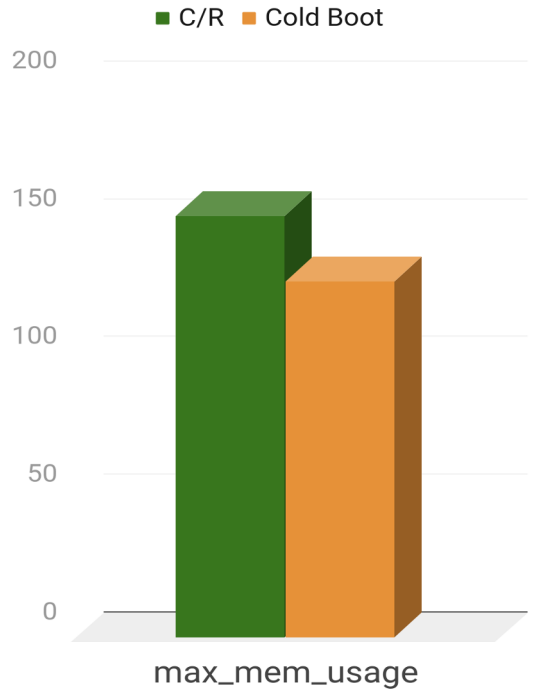

	NSDI '20 paper
Memory Limit	256mb
CPU Core	1
kernel	4.13
File system	empty
image size	0
boot time window	VM Bootup
Boot time	**200ms**

Inefficient resource sharing with cloud-native technologies

Launch-time optimizations exacerbate problems

Recent advances not designed for complex runtimes

Edge-Native Technology

- Having **shared long-running backends as runtime** for multiple applications and instances reduced

  resource pressure, application size

  improved responsiveness / warmed-up runtime

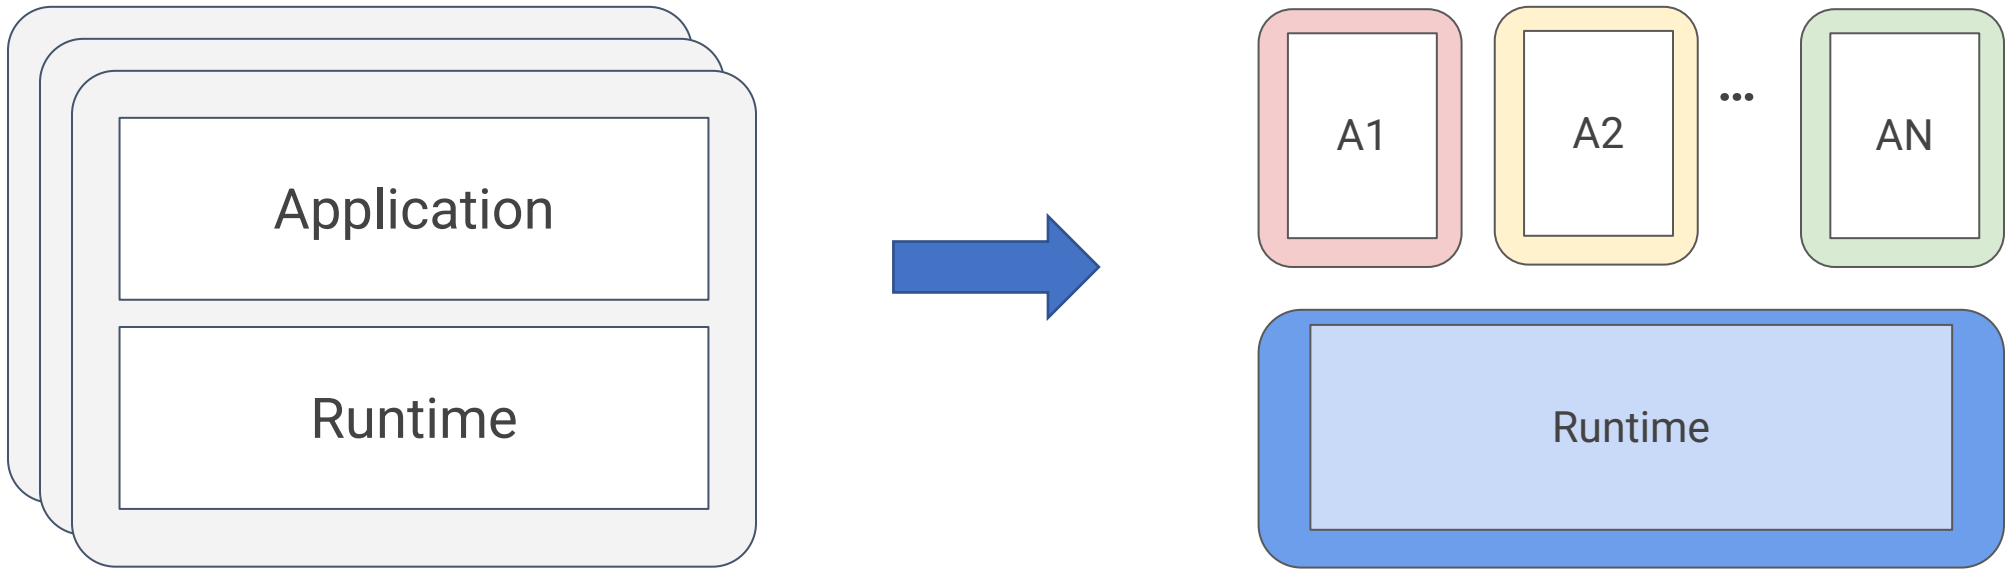

Pocket Model

Anyone say Enclaves?

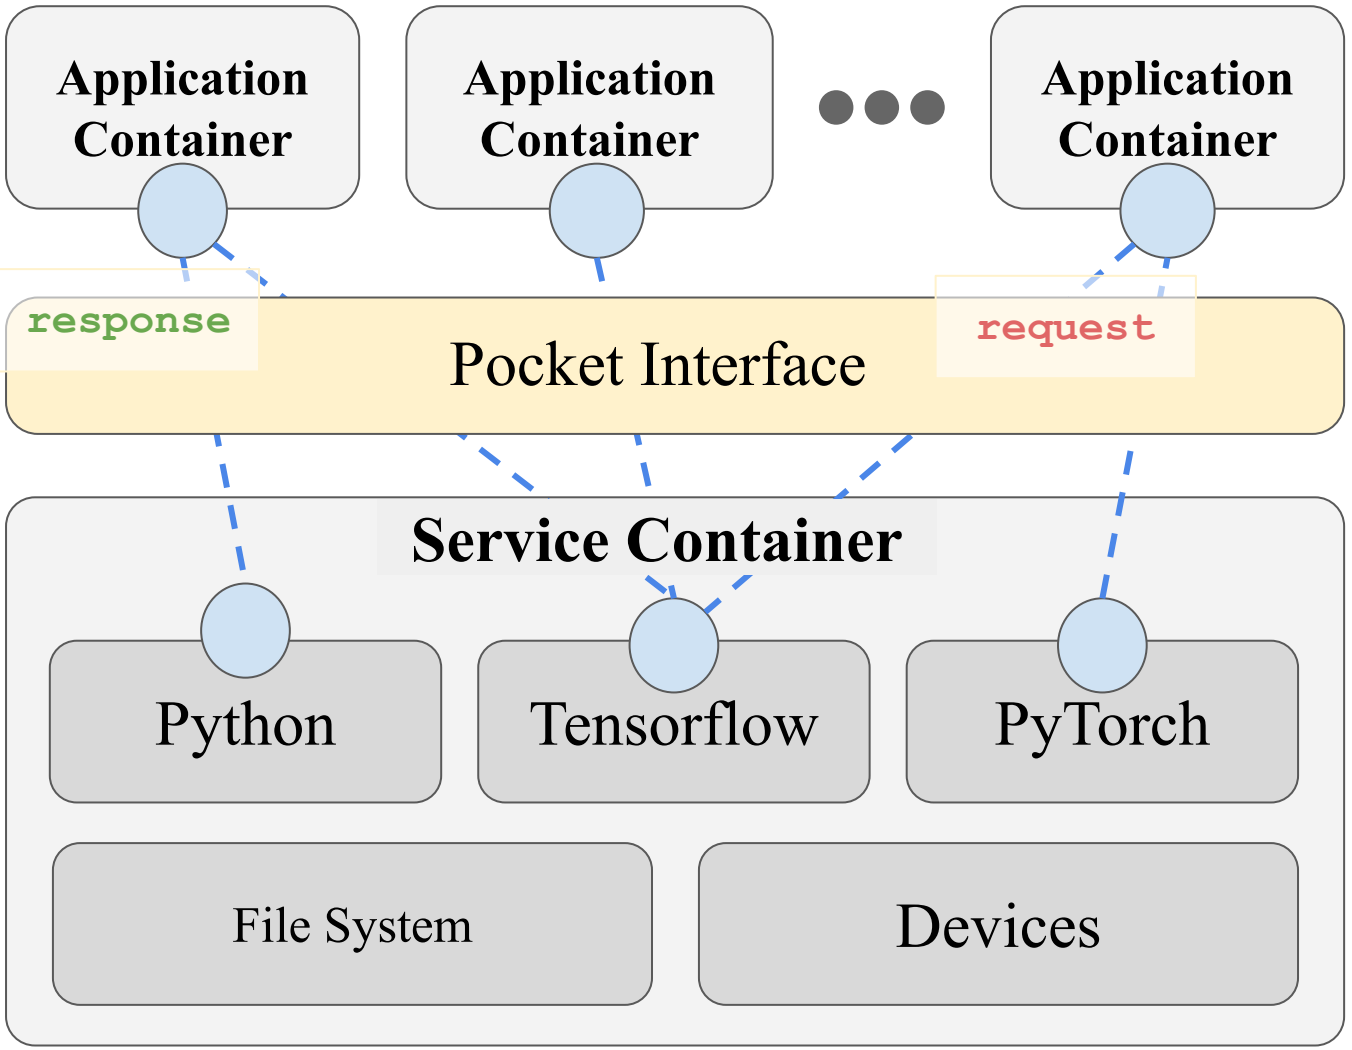

Workload Isolation

Lightweight IPC

Concurrency and Dynamic Resource Scaling in Runtime

More new problems: Latency contention problem at the edge

- Running IA and ML application together at an edge location (result is from LTE)
- IA (IoT+ROS) is at RAN, and ML (visual analytics) is moved progressively closer

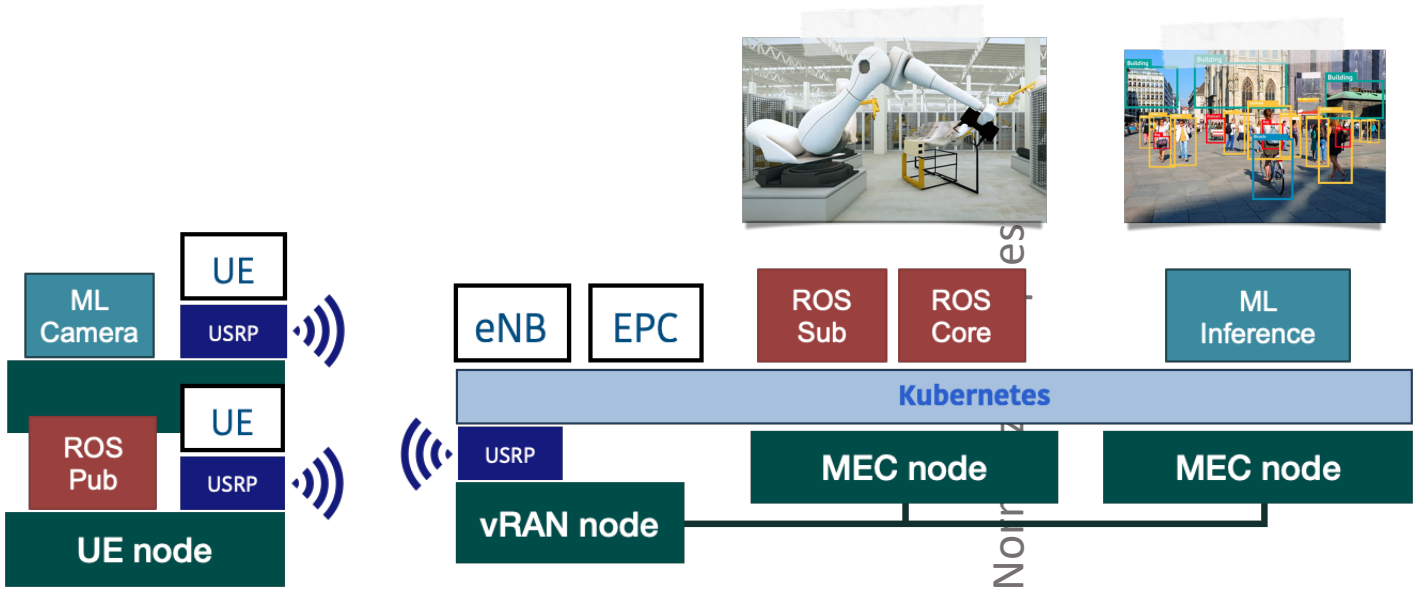

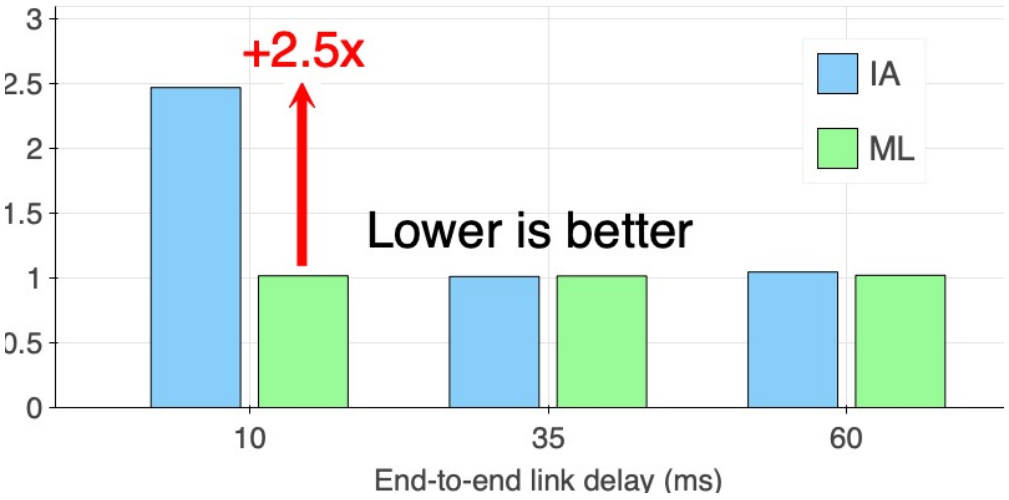

More new problems: Latency contention problem in MEC

- Transport-level flow control engines clocked by RTT lead to longer and denser packet bursts
- Need a solution independent of endpoint protocols and in-network QoS support

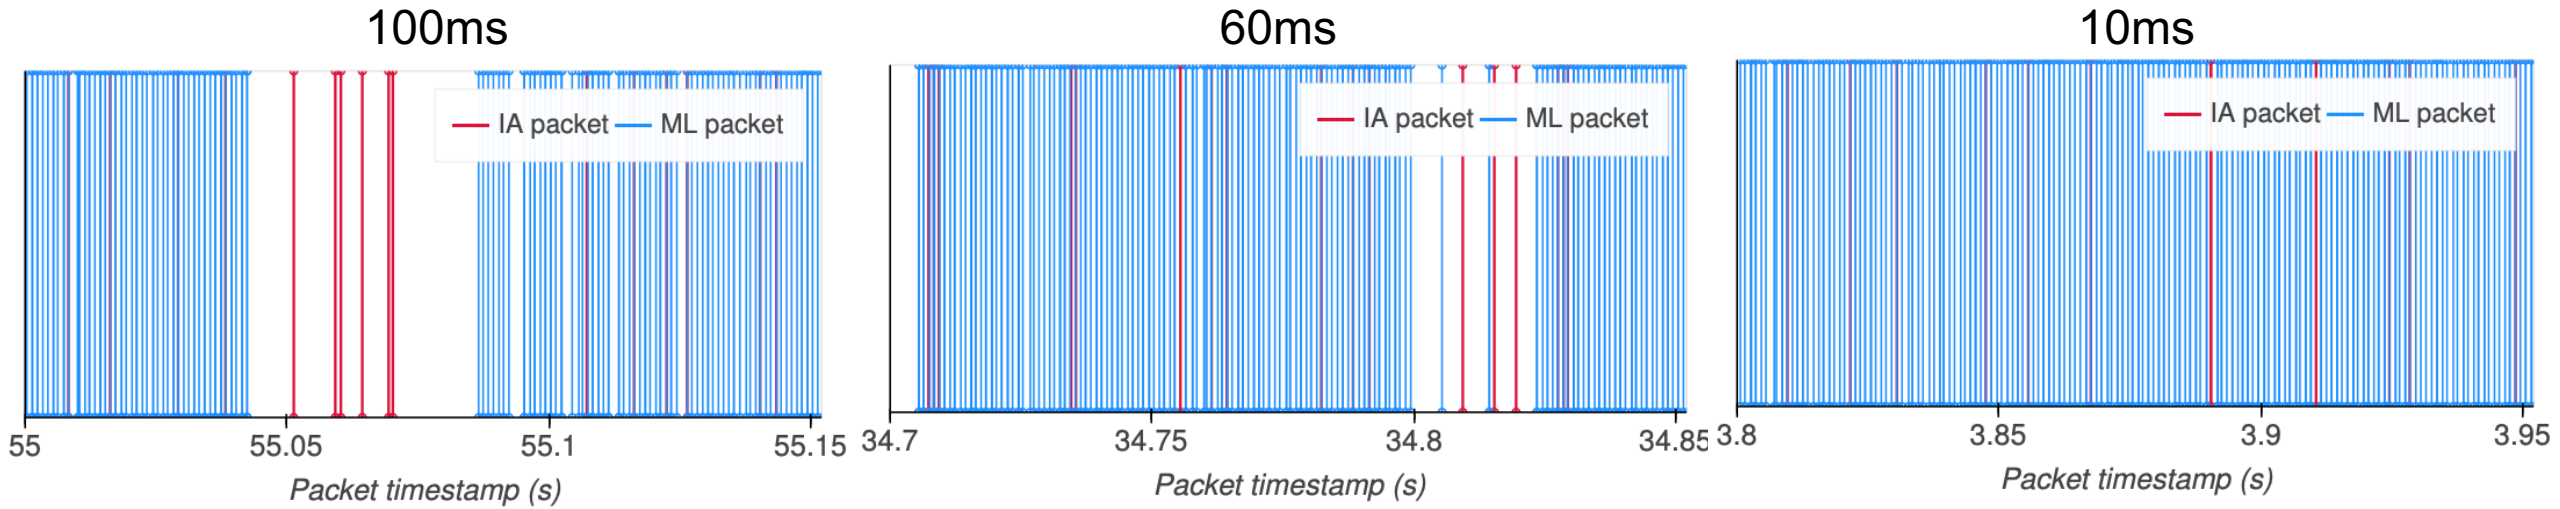

- **MEC-based collocated workloads trigger latency contention that can obviate the edge benefits for LC applications**

Need new fine-grained latency-aware packet burst management for a multi-tenant edge

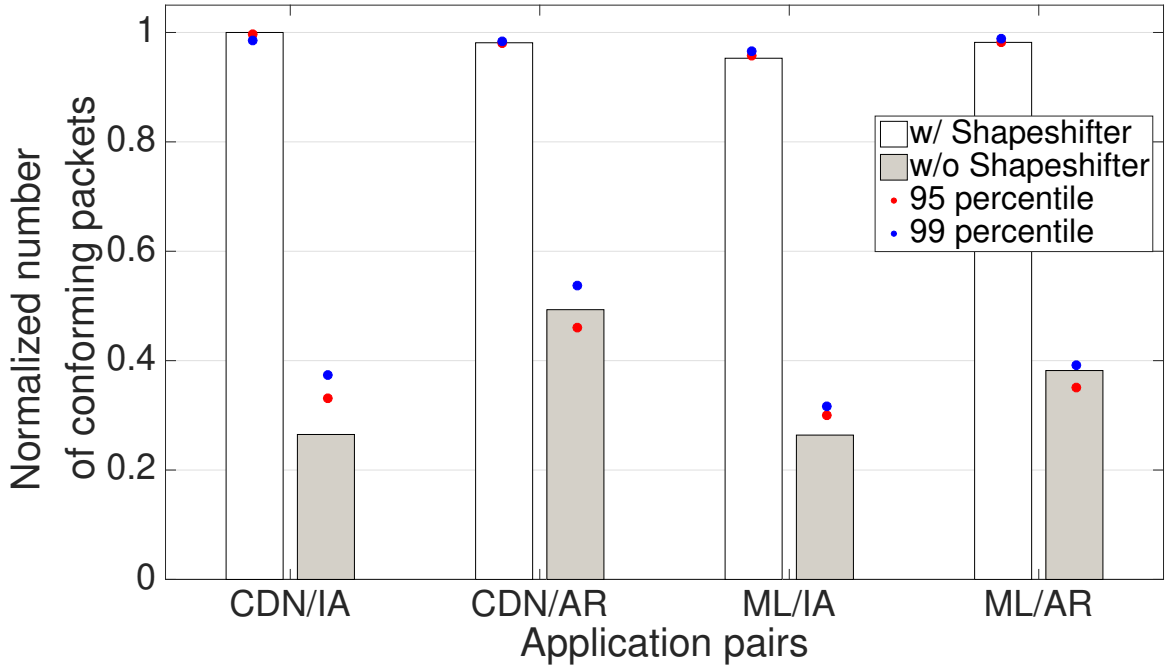

ShapeShifter can support latency-centric QoS

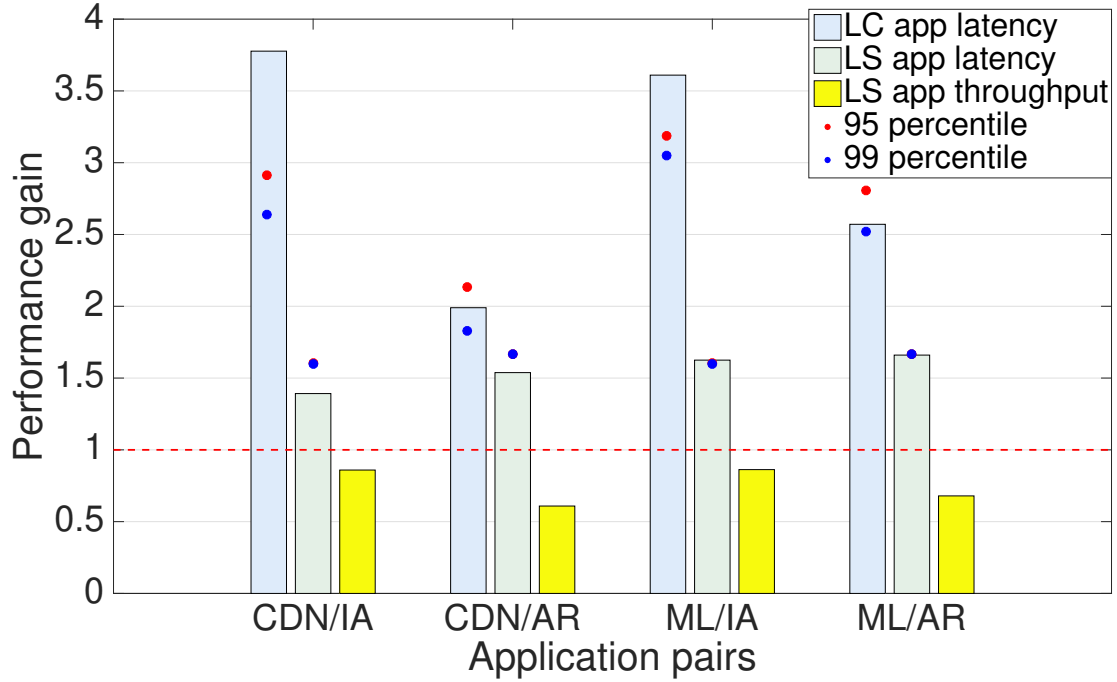

ShapeShifter provides benefits to latency-critical and latency-sensitive applications

In-network Analytics

with 

Rafael Oliveira

and Hardik Sharma, Haggai Eran

Applications Driving Architectures

Accelerating Edge Analytics

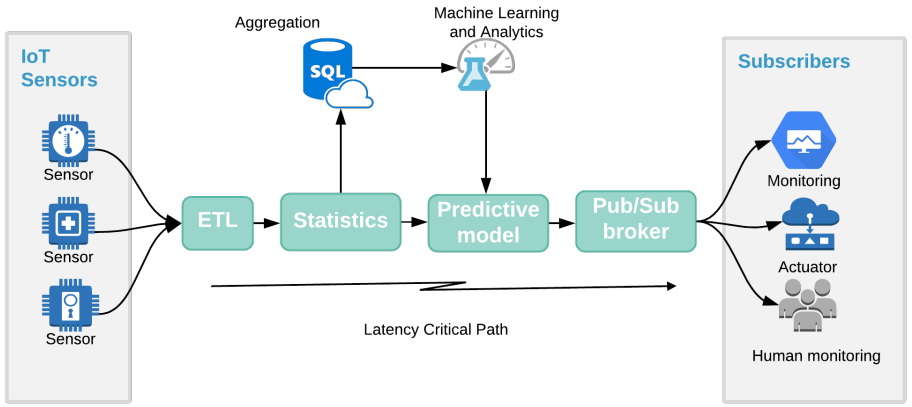

with current network and cloud => network-bound

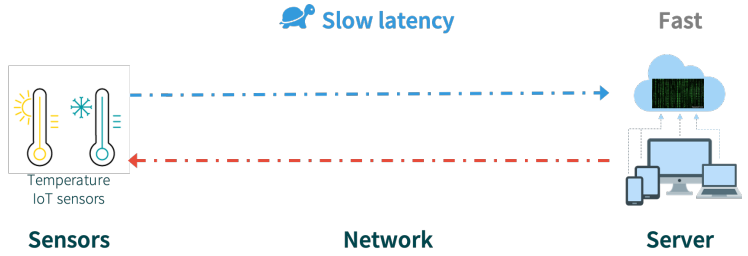

with 5G network and edge => compute-bound

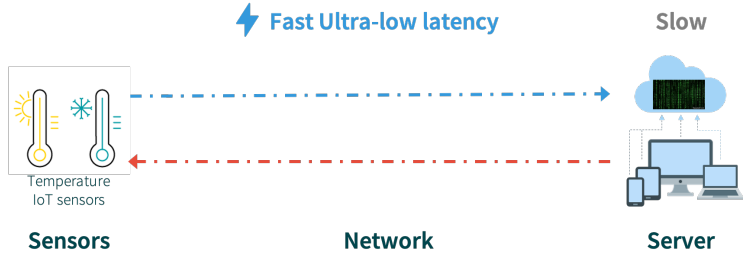

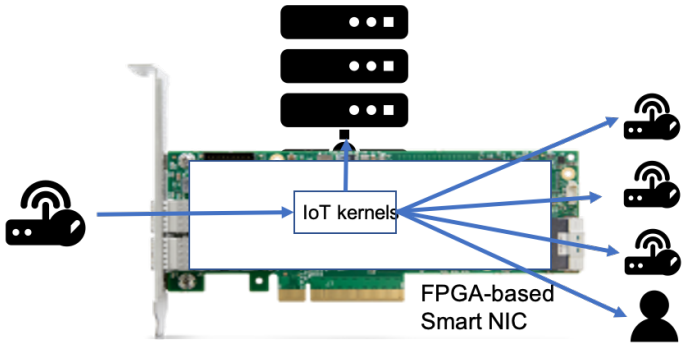

Vasado: accelerate IoT analytics with Smart NICs:

- **latency** of critical path benefits from *in-network offload engine*
- **throughput** of analytics path benefits from in-network *application-specific batching and compression*

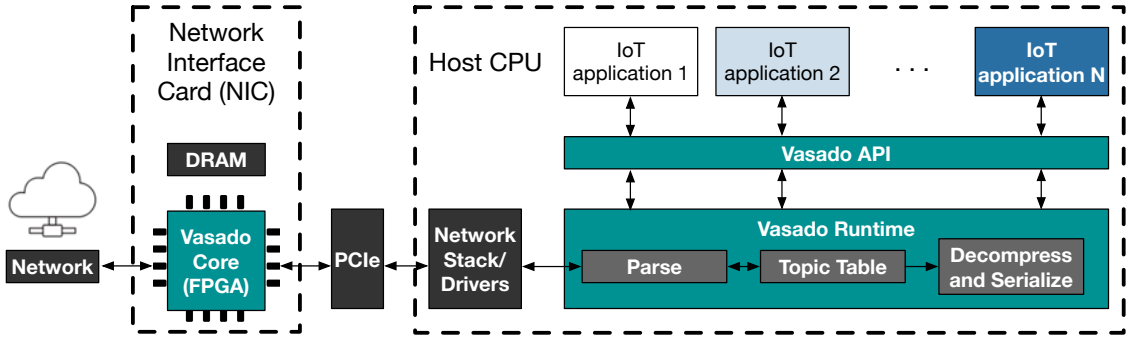

Vasado Core:

- *programmable domain-specific core*
- *fast, lightweight multiplexing (ctx_switching)*

Accelerating Edge Analytics

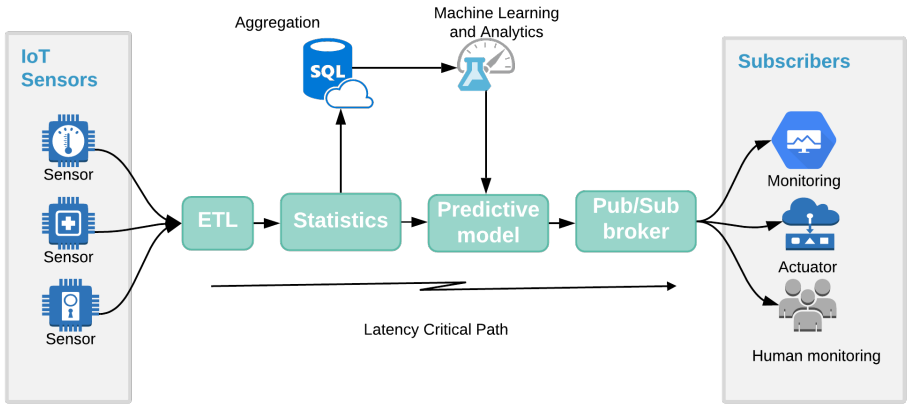

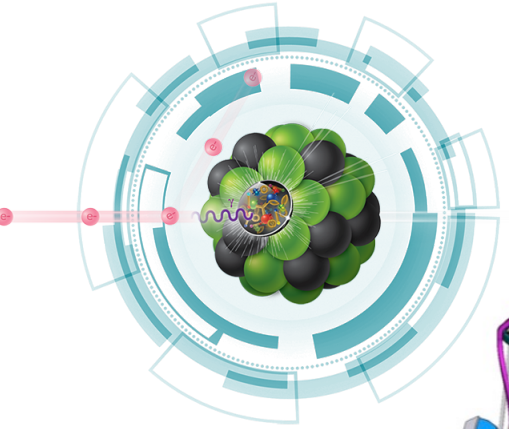

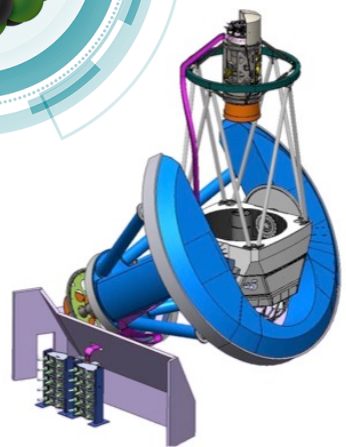

Science too!

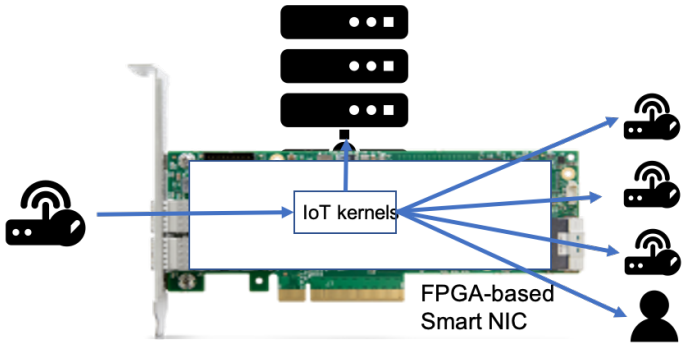

Vasado: accelerate IoT analytics with Smart NICs:

- **latency** of critical path benefits from *in-network offload engine*
- **throughput** of analytics path benefits from in-network *application-specific batching and compression*

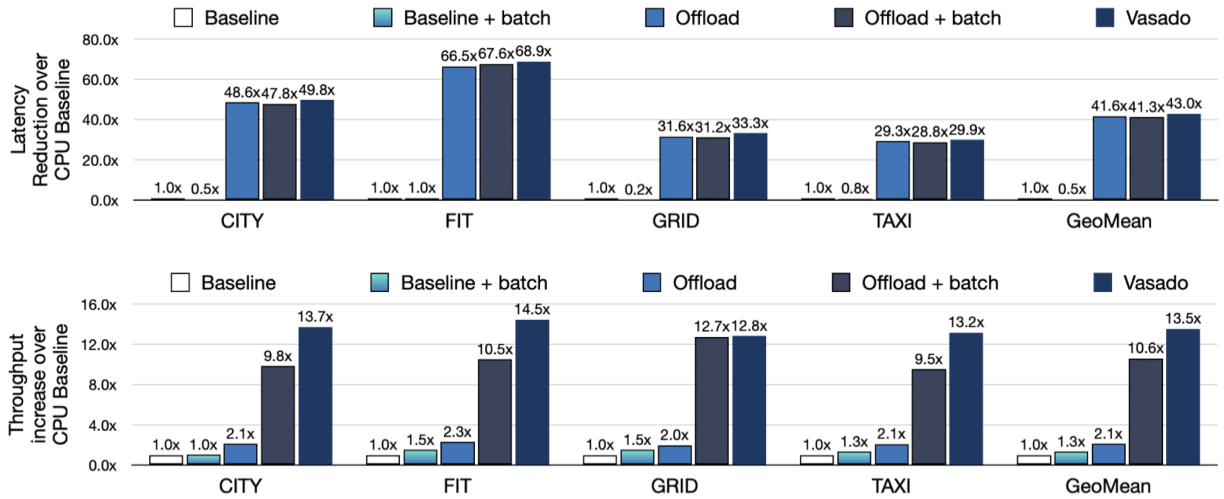

Edge Computing Challenges

- Growth in demand
  - Huawei estimate 5G transition from 50TWh to 100TWh mobile network
- Datacenter-native technologies
  - Software stack, accelerators
  - Natural cooling? PUE efficiency?
- Deployment cost, scale, and challenges
  - O(US$1000) per location
  - Densification of infrastructure, urban deployment, ensuring coverage
- Sustainability of access

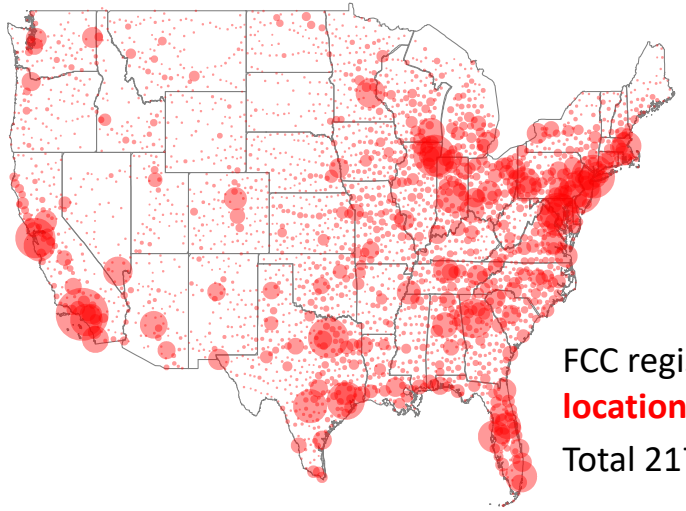

FCC registered cellular tower **locations** (Crown Castle, …)

Total 217,346, as of Mar. 2017

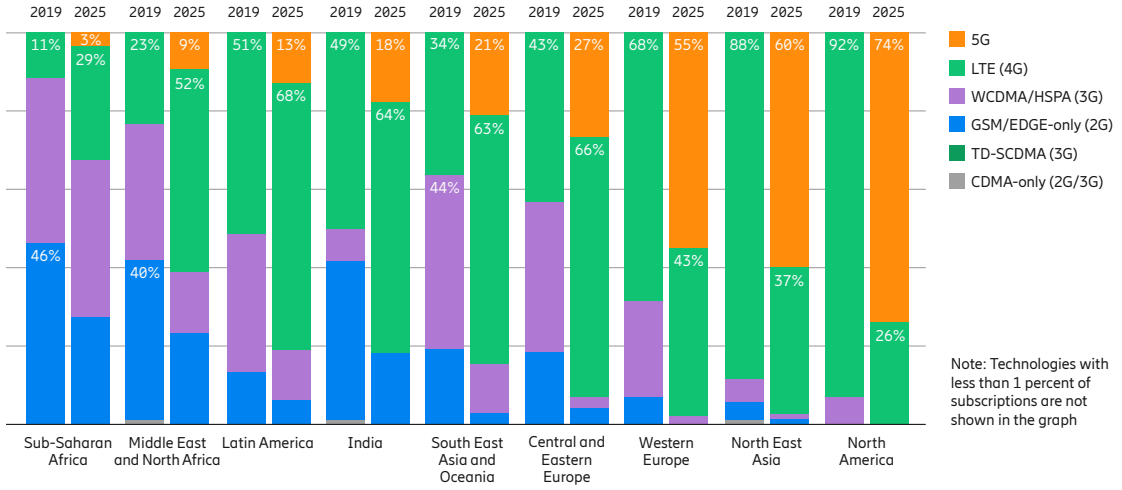

Note: Technologies with less than 1 percent of subscriptions are not shown in the graph

Thank you!

Emerging Systems Lab

Ada Gavrilovska

Greg Eisenhauer

Ketan Bhardwaj

Pradeep Fernando

Thaleia Doudali


Ranjan S. Venkatesh


Harshit Daga


Carol Hsu


Rafael Oliveira


Tony Mason


Jim Choncholas


Jin Heo


Misun Park


Daniel Zahka


Vaibhav Bhasole

www.cc.gatech.edu/~ada