A Smorgasbord of Embedded and Pervasive Computing Research

Kishore Ramachandran

(part of systems group which includes Ada Gavrilovska, Taesoo Kim, Ling Liu, Calton Pu, and Alexey Tumanov)

Georgia Tech College of Computing

- Current PhD inmates!


Enrique Saurez


Harshit Gupta


Zhuangdi Xu


Hall


Ashish Bijlani

- Recently escaped!


Tyler Landle

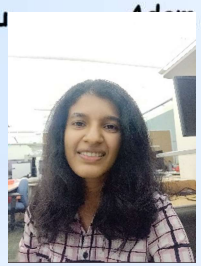
Manasvini Sethuraman

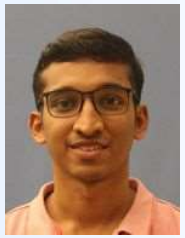
Anirudh Satma

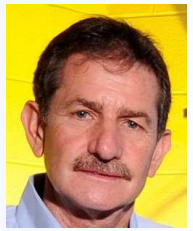
Alan Nussbaum

  - Hyojun Kim (IBM Almaden then Startup and now Google); Lateef Yusuf (Amazon then Google); Mungyung Ryu (Google); Kirak Hong (Google and now CTRL-labs); Dave Lillethun (Seattle U.); Dushmanta Mohapatra (Oracle); Wonhee Cho (Microsoft); Beate Ottenwalder (Bosch); Ruben Mayer (TU Munich), Ashish Bijlani (Startup)
- Plus a number of MS and UGs

Georgia Tech College of Computing

Embedded Pervasive Lab

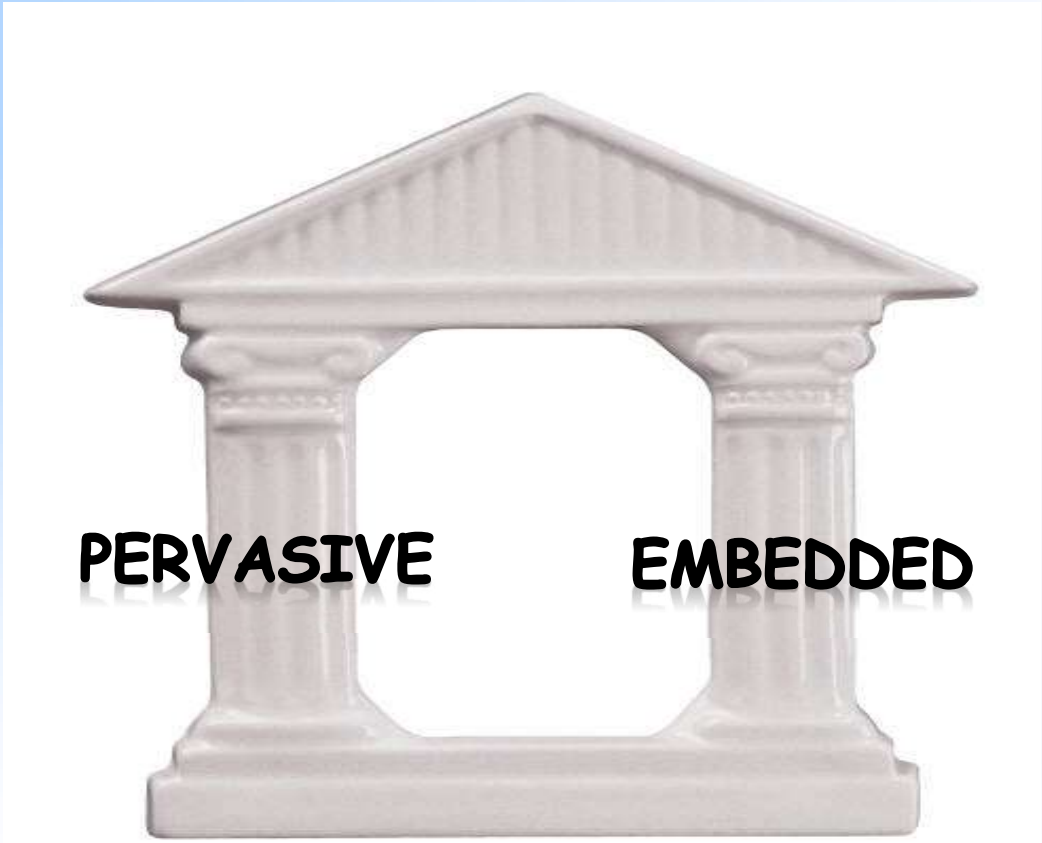

Georgia Tech College of Computing

Pervasive side of the house

- Embedded devices treated as black boxes
- System Support for IoT
  - Fog/Edge computing

Current-Generation Applications

Cloud computing's utility model commoditized hardware…

Enabling large-scale apps from centralized data centers…

facebook NETFLIX

Next-Generation Applications

Future apps will be data-driven, model-driven, machine-in-the-loop, and far more demanding…

Autonomous Vehicles (AV)

Augmented / Virtual Reality (AR/VR)

Latency <20 ms

Bandwidth >50 Mbps

Per device

Smart Cities

Smart Camera Networks

Next-Generation Applications

- Sense -> Process -> Actuate
- Common Characteristics
  - Dealing with real-world data streams
  - Real-time interaction among mobile devices
  - Wide-area analytics
- Requirements
  - Dynamic scalability
  - Low-latency communication
  - Efficient in-network processing

Cloud Computing

- Good for web apps at human perception speeds
  - Throughput oriented web apps with human in the loop
- Not good for many latency-sensitive IoT apps at computational perception speeds
  - sense -> process -> actuate
- Other considerations
  - Limited by backhaul bandwidth for transporting plethora of 24x7 sensor streams
  - Not all sensor streams meaningful
    => Quench the streams at the source
  - Privacy and regulatory requirements

Fog/Edge Computing

- Extending the cloud utility computing to the edge
- Provide utility computing using resources that are
  - Hierarchical
  - Geo-distributed

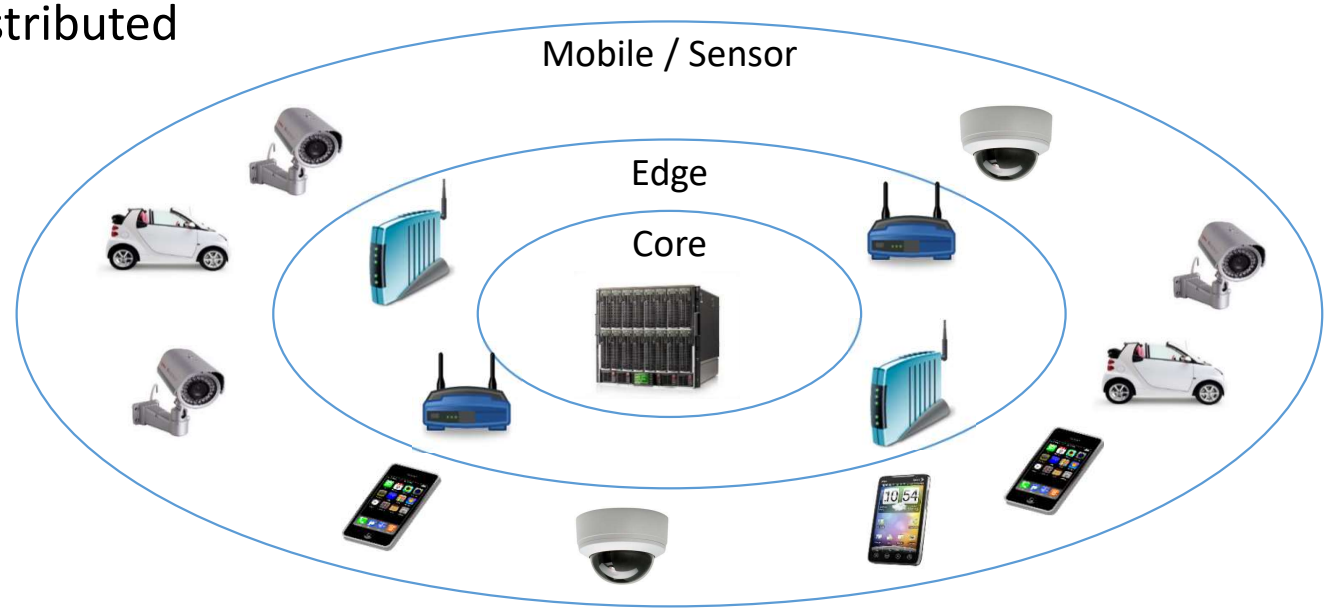

Fog/edge computing today

- Edge is slave of the Cloud
  - Platforms: IoT Azure Edge, CISCO Iox, Intel FRD,
- Mobile apps beholden to the Cloud

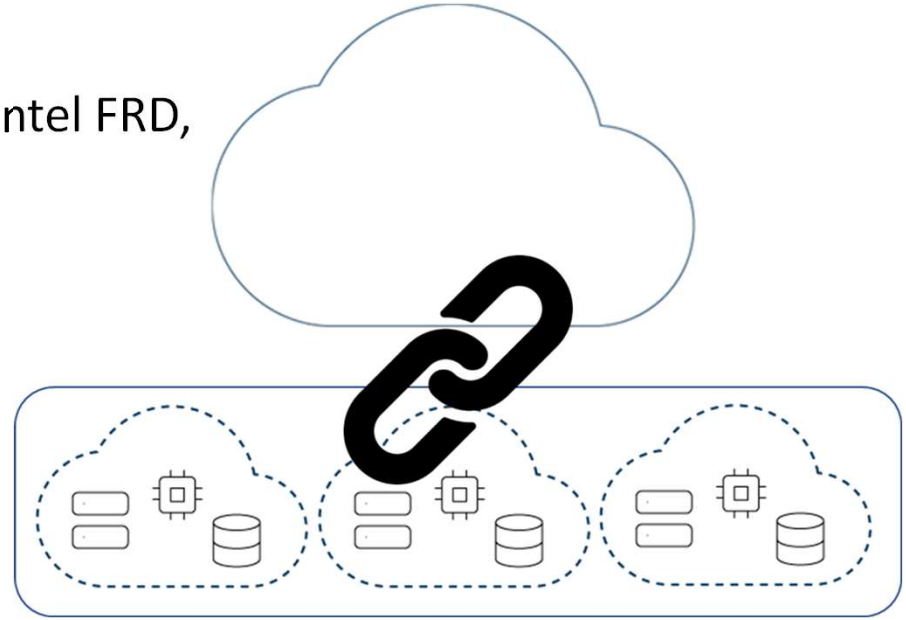

Vision for the future

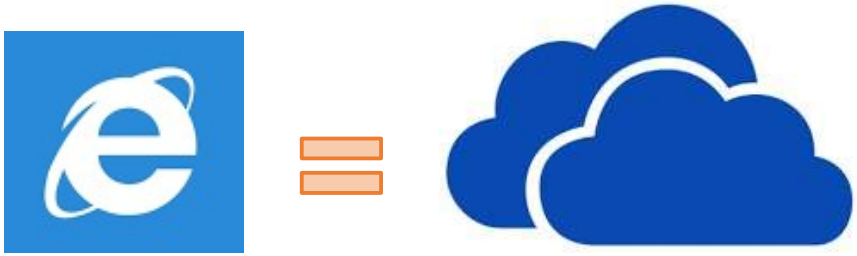

- Elevate Edge to be a peer of the Cloud
  - Prior art: Cloudlets (CMU+Microsoft), MAUI (Microsoft)
- In the limit
  - Make the Edge autonomous even if disconnected from the Cloud

Why  ?

- Interacting entities (e.g., connected vehicles) connected to different edge nodes
- Horizontal (p2p) interactions among edge nodes essential

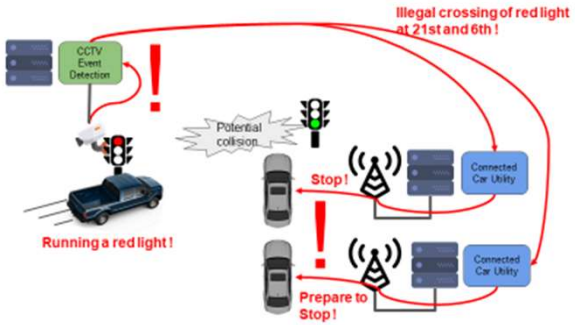

Why  ?

- Autonomy of edge (disaster recovery)

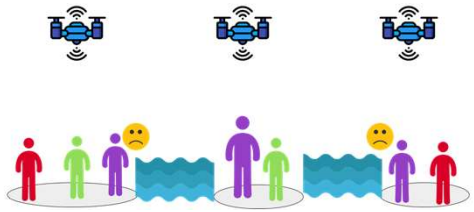

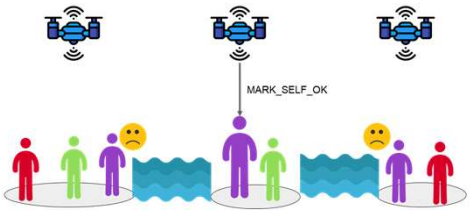

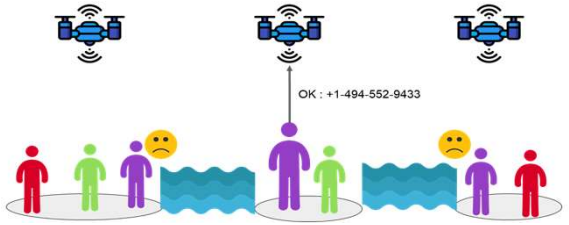

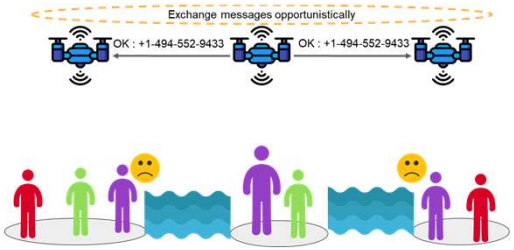

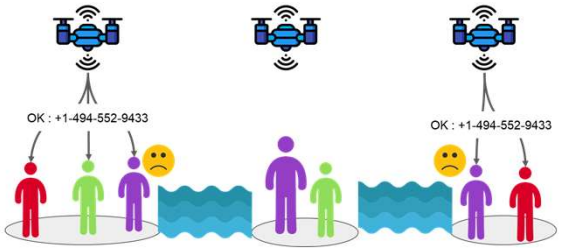

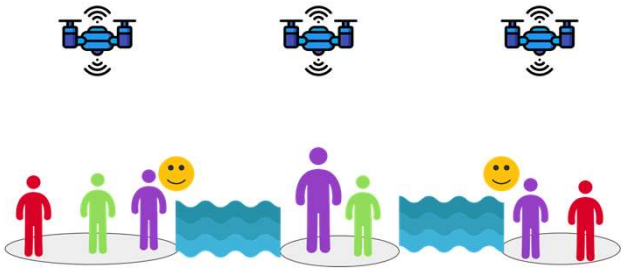

Challenges for making 

- Need for powerful frameworks akin to the Cloud at the edge
  - Programming models, storage abstractions, pub/sub systems, …
- Geo-distributed data replication and consistency models
  - Heterogeneity of network resources
  - Resilience to coordinated power failures
- Rapid deployment of application components, multi-tenancy, and elasticity at the edge
  - Cognizant of limited computational, networking, and storage resources

Thoughts on Meeting the Challenges
(https://www.cc.gatech.edu/~rama/recent_pubs.html)

Theme: Elevating the Edge to be a peer of the Cloud

- Vision papers:
  - A case for elevating the edge to be a peer of the cloud", GetMobile, 2020
  - eCloud: Vision for the Evolution of Edge-Cloud Continuum, IEEE Computer, May 2021
- Geo-distributed programming idioms for Edge/Cloud continuum
  - Foglets (ACM DEBS 2016)
  - ePulsar (ACM SEC 2021)
- Geo-distributed data replication and resource management
  - FogStore (ACM DEBS 2018)
  - DataFog (HotEdge 2018)
- Efficient Edge runtimes
  - Serverless functions using WebAssembly (ACM IoTDI 2019)
- Applications using autonomous Edge
  - Social Sensing *sans* Cloud (SocialSens 2017)
  - STTR: Space Time Trajectory Registration (ACM DEBS 2018)
  - STVT: Space-Time Vehicle Tracking (HotVideoEdge 2019)
  - Coral-Pie: Space-Time Vehicle Tracking at the Edge (ACM Middleware 2020)

Embedded side of the house

- Infinite storage for mobile devices
- Optimizing Mobile Video Downloads

Infinite Storage for Mobile Devices

- Seamlessly extend the storage on mobile to the Cloud for any app
  - User space file system
    - APSys 2018, USENIX ATC 2019, Sigmetrics 2021
- Use machine learning to build user's everyday working set and (off)load (un)wanted data
- Issues
  - Latency
  - Energy consumption
  - Security and privacy

Optimizing Mobile Video Downloads

- Foresight (ACM MMSys 2021)
  - Bandwidth prediction across space and time for mobile users

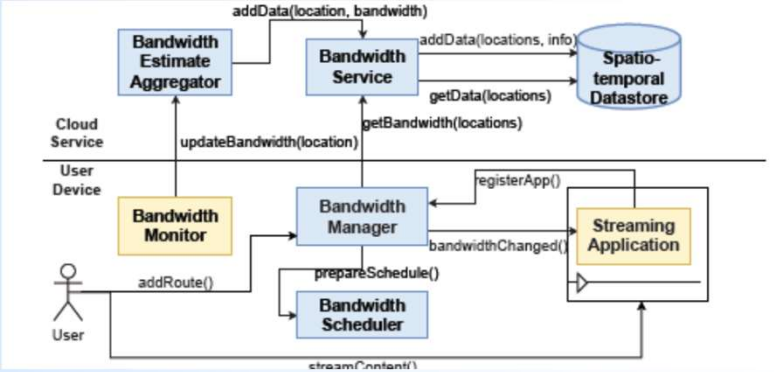

- ClairvoyantEdge
  - Short range mmWave augmentation at Edge for high bandwidth video delivery

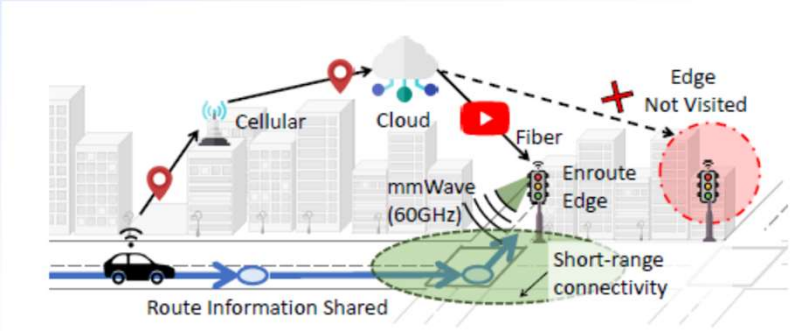

Georgia Tech College of Computing

Recap

- Infinite storage for mobile devices
- Optimizing Mobile Video Downloads
  - Foresight, ClairvoyantEdge

- Fog/Edge computing
  - eCloud
  - Foglets, OneEdge, thin virtualization for FaaS
  - Fogstore, DataFog, ePulsar, NFSlicer
  - STTR, Socialsens

Georgia Tech College of Computing

Ongoing Projects

- **eCloud**: Device-Edge-Cloud continuum
- **OneEdge**: Device/Edge/Cloud control plane using AV as exemplar
  - Scheduling edge resources, monitoring, migration
- **Foresight** and **ClairvoyantEdge**: Prescient video prefetching at the edge for AV infotainment (With Prof. **Dhekne**)
  - Use route to JIT prefetching and caching for DASH player on vehicle
  - Use mmWAVE (integrated with 5G LTE for edge node selection) to beam to passing vehicle
- Edge centric video data management systems for AV (With Prof. **Arulraj**)
  - Annotations with video for query processing, multi-tenancy, and sharing
- Nimble execution environments for the Edge
  - Analyze cold start times in containers
  - Clean slate exec environment for FaaS
- **NFSlicer**: dataplane for processing network functions (With Prof. **Daglis**)
  - Selective data movement (e.g., header vs. payload) for NF chaining
- **MicroEdge**: Low-cost edge architecture for camera processing (With Prof. **Krishna**)
- Edge computing solution for underserved communities
  - Smart information services without WAN connectivity

Pubs:
http://www.cc.gatech.edu/~rama/recent_pubs.html

E-mail: rama@cc.gatech.edu

Lab:
http://wiki.cc.gatech.edu/epl

Georgia Tech College of Computing

What does Kishore "really" do in his copious spare time when he is not teaching?

Squash anyone?

All virtual!

Soon real?

Table tennis anyone?

What you should take away?

- "Kishore" rhymes with "sea shore"
- Squash/Table-tennis
- EPL
- Fog/Edge computing
- Infinite storage on mobile/EdgeCaching

Georgia Tech College of Computing

https://www.dreamstime.com/namaste-vector-hand-drawn-symbol-yoga-design-image112101574