Fog Computing:
On the Road to an Open Infrastructure for the Internet of Things

Keynote Talk KuVS Fachgespräche, TU Darmstadt

Ruben Mayer, University of Stuttgart

08th of March, 2018

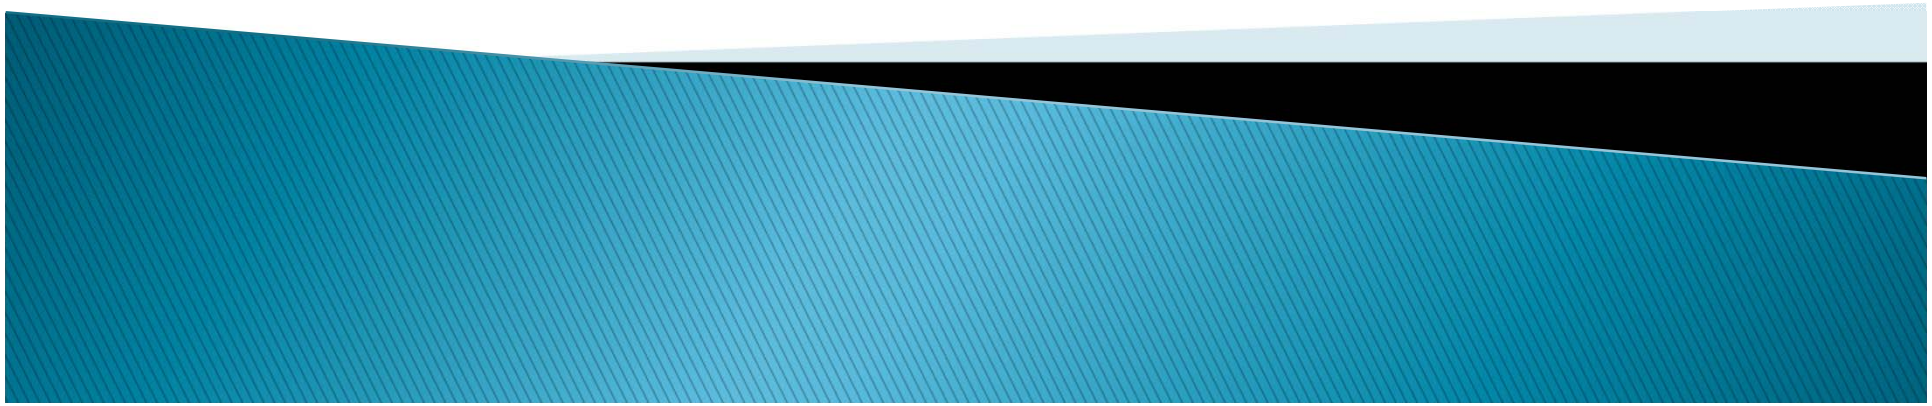

Overview

- Motivation
- Introduction
- Openness and standardization
- Research
- Conclusion and outlook

Takeaways

- You will be able to answer three questions

1. What is fog? *(definition)*
2. What can fog do for me? *(benefits)*
3. What can I do for fog? *(research)*

The Data Surge...

- The Internet of Things (IoT) rapidly evolves
  - 100 M IoT-connected cars (2022)
  - 800 M smart meters (2020)
  - Heavy sensor streams (1 GB/sec in a car)
- Data needs to be analyzed and stored

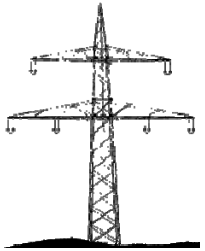

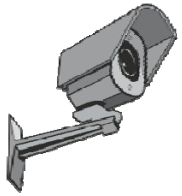

Cloud Computing

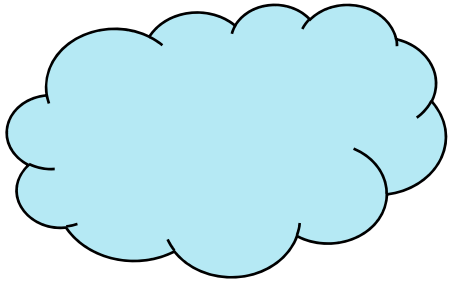

- Centralization of computing in data centers
- Virtualization allows for elasticity and multitenancy
  - Efficient resource usage
  - Cutting costs
  - Pay-as-you-go

Does Cloud Computing Suit IoT?

Cloud has shortcomings

1. Bandwidth constraints
   - Bandwidth is expensive
2. High latency
   - > 20 ms latency
3. Dependent on Internet Connectivity
   - No Internet – No Cloud!

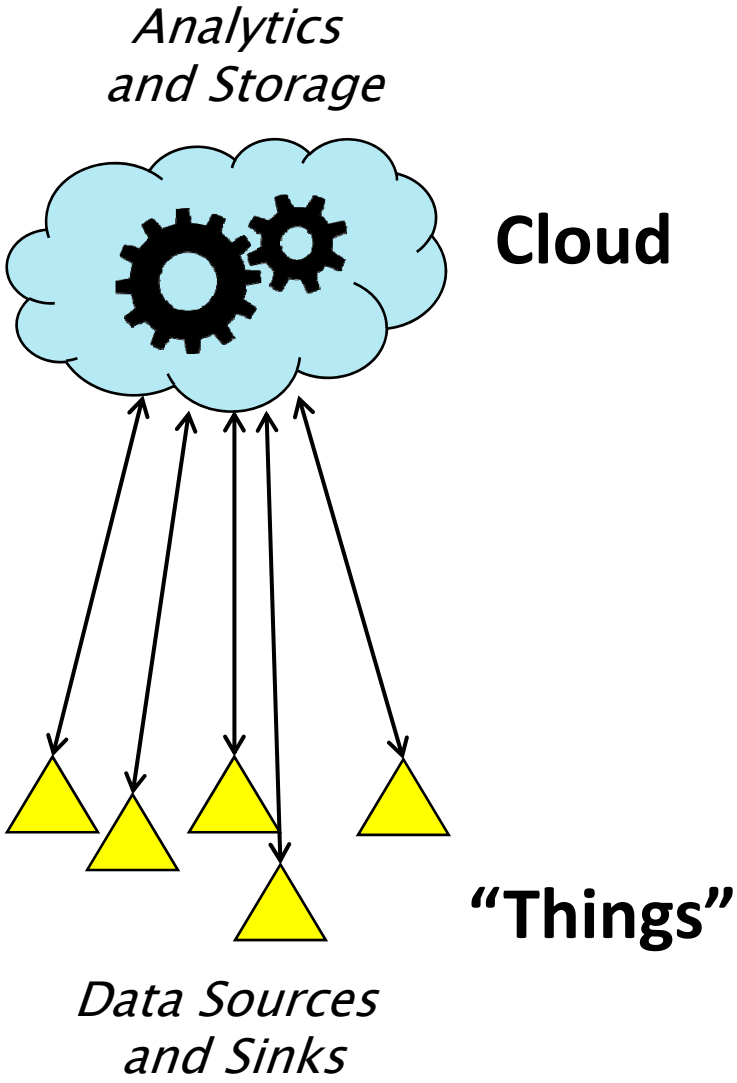

Use Case 1: Camera Networks

- Municipal security system
- Hundreds of cameras
- Real-time analytics
  - Authentication
  - Crowd behavior
- High bandwidth requirements

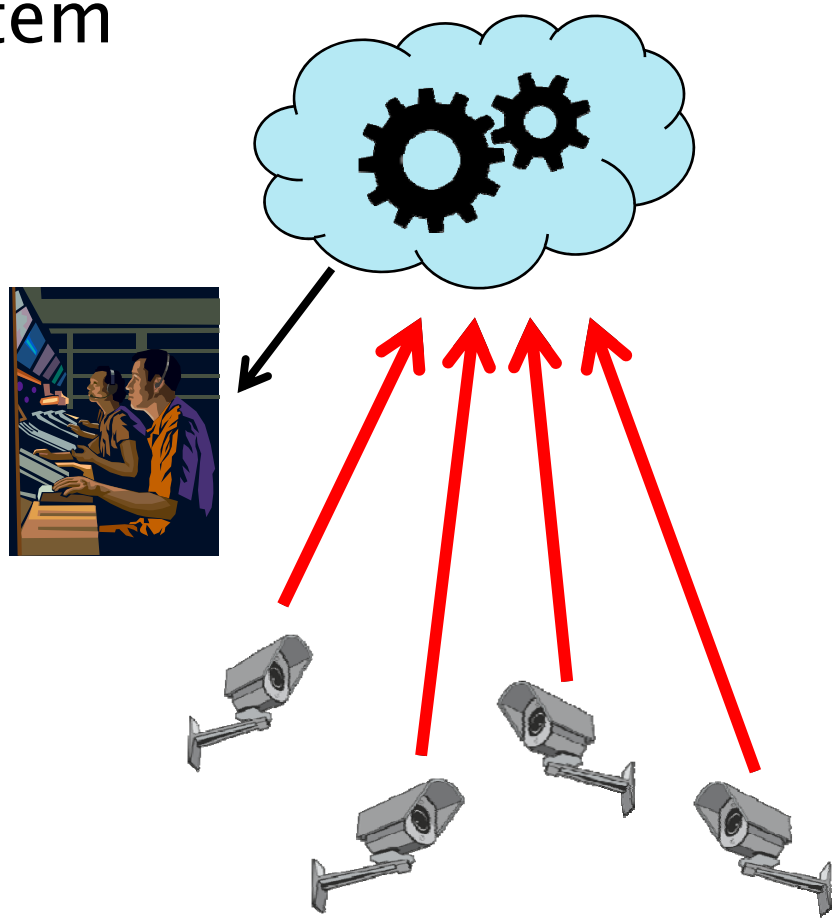

Use Case 2: Emergency Response

- Example: Natural disaster
- Supporting services
  - Safety check
  - Maps for rescue team
- Internet connectivity may be disrupted

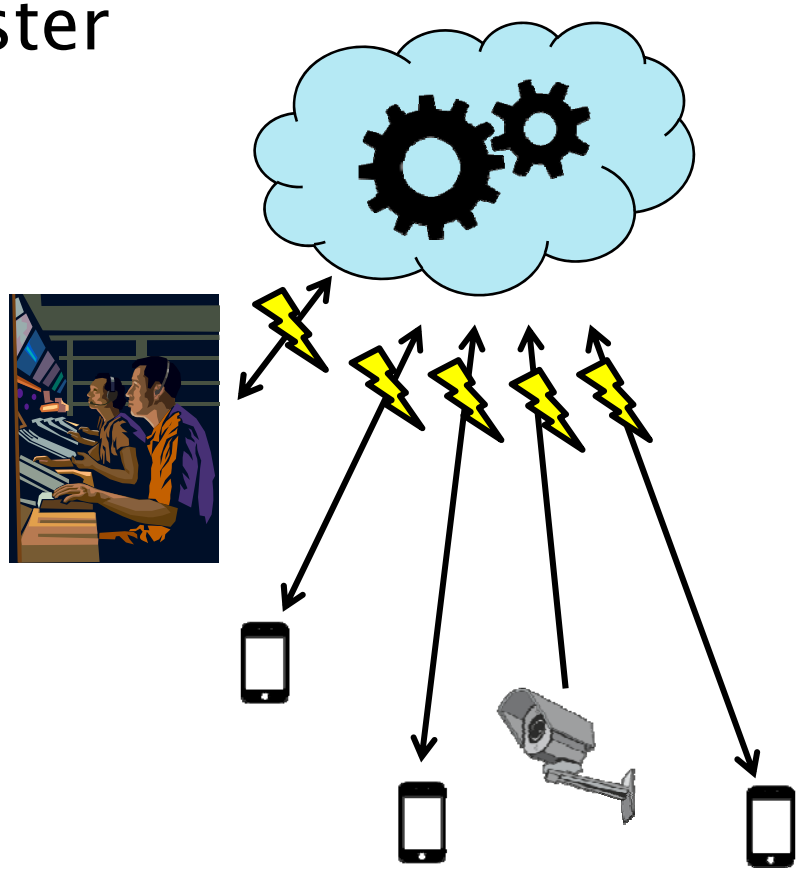

"Cloud Closer to the Ground"

- Analytics needs to happen close to the data!
- Edge computing (on devices) not enough
  - Some devices are not powerful enough
  - Collaboration between devices must be coordinated

"*We need something in between edge and cloud that makes it all work together.*"

➔ This is fog computing.

Fog Computing

- The definition of fog computing:

"A horizontal, system-level architecture that distributes computing, storage, control and networking functions closer to the users along a cloud-to-thing continuum."
(OpenFog Reference Architecture)

- What does this mean in practice?

Fog Architecture

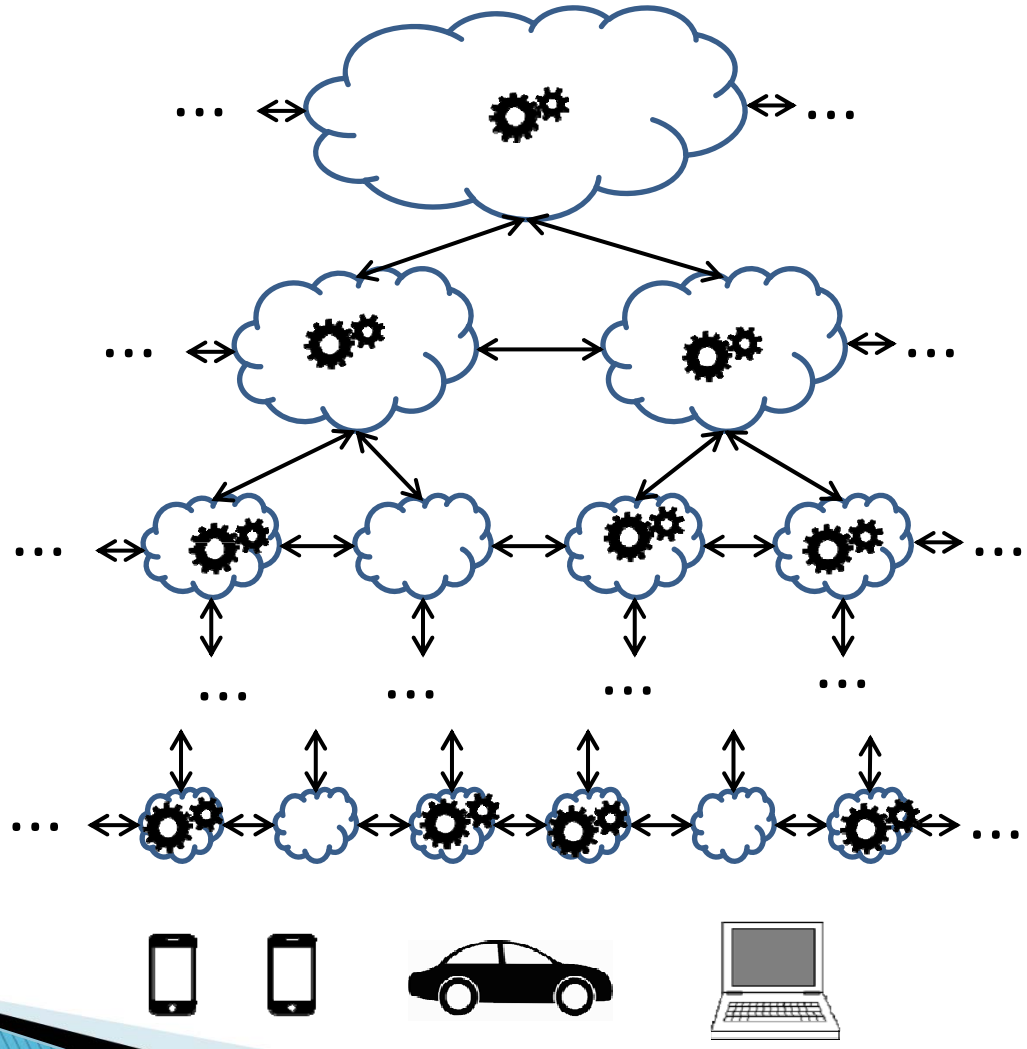

How Does Fog Computing Help?
In Camera Networks…

- Make use of fog hierarchy
  - Keep computation local
  - Global view on higher level
- Benefits
  - Low latency
  - Less network load
  - Data privacy

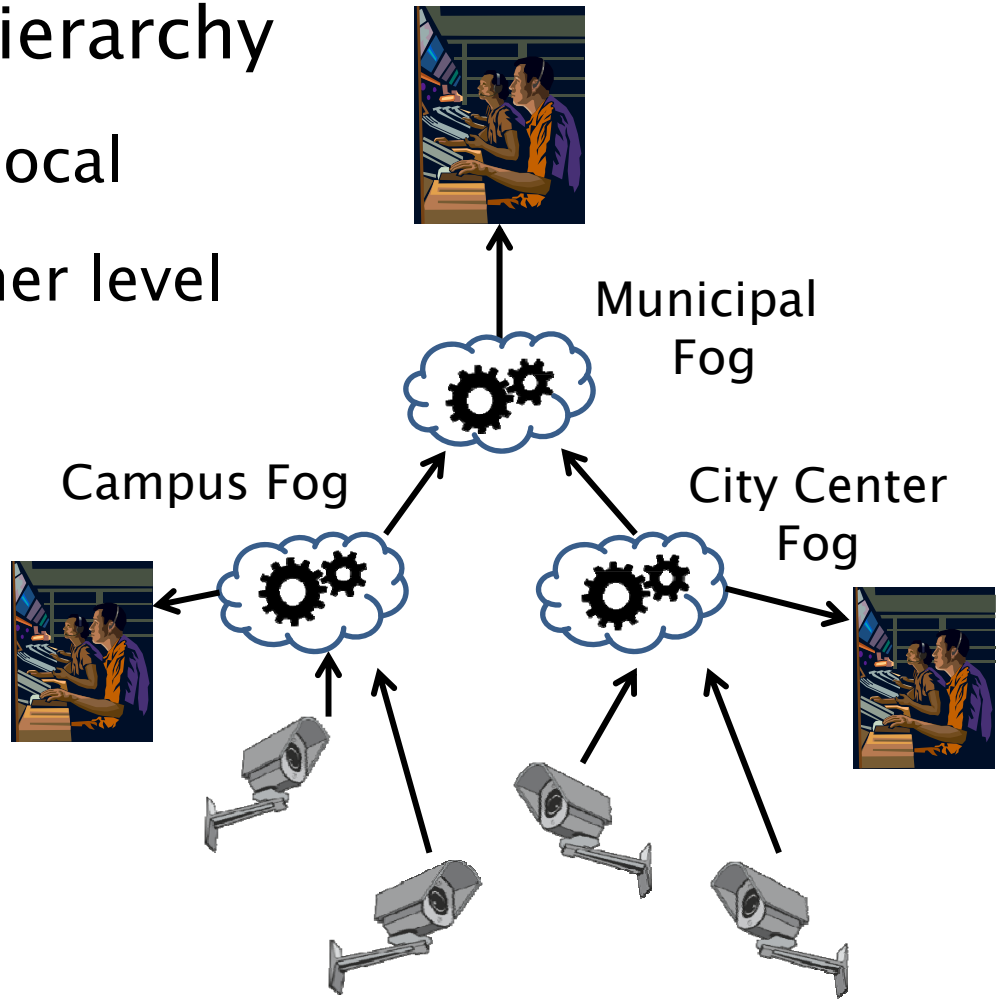

How Does Fog Computing Help?
In Emergency Response…

- Use mobile drones as fog nodes
- Benefits
  - Local connectivity
  - Works when infrastructure unavailable

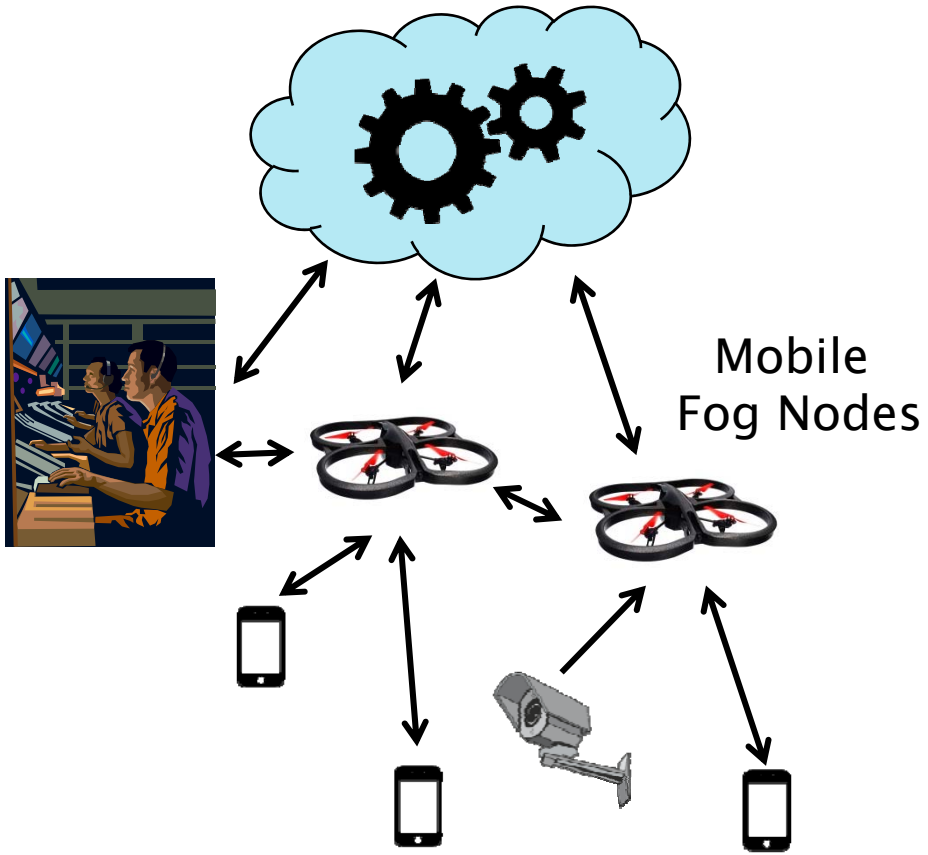

Overview

- Motivation
- Introduction
- **Openness and standardization**
- Research
- Conclusion and outlook

Silos: A Pitfall in Fog

- "Fog computing is great. Let's build our own fog infrastructure."
  - ➔ Cisco Fog, Bosch Fog, …
  - ➔ Smart city fog, traffic fog…
- Redundant developments
- Small businesses cannot develop fog solutions
- Everyone is losing revenues

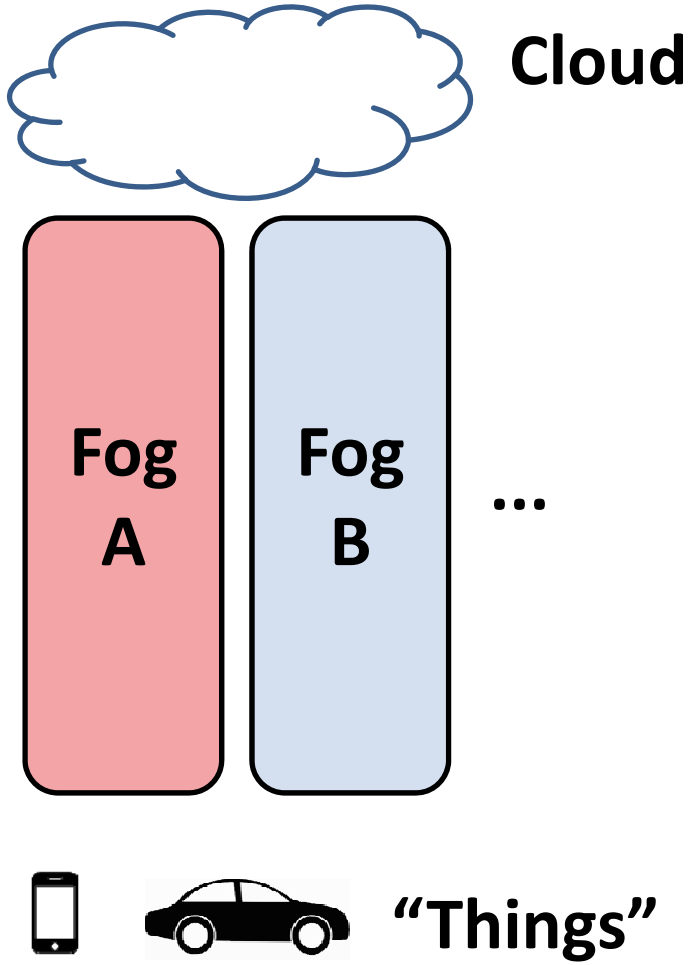

OpenFog

- Consortium of industry and academia partners

arm CISCO PRINCETON UNIVERSITY …

57 members

- Goal: Interoperable, open fog infrastructure
- "Fog as a service"
- Standardization (with IEEE)

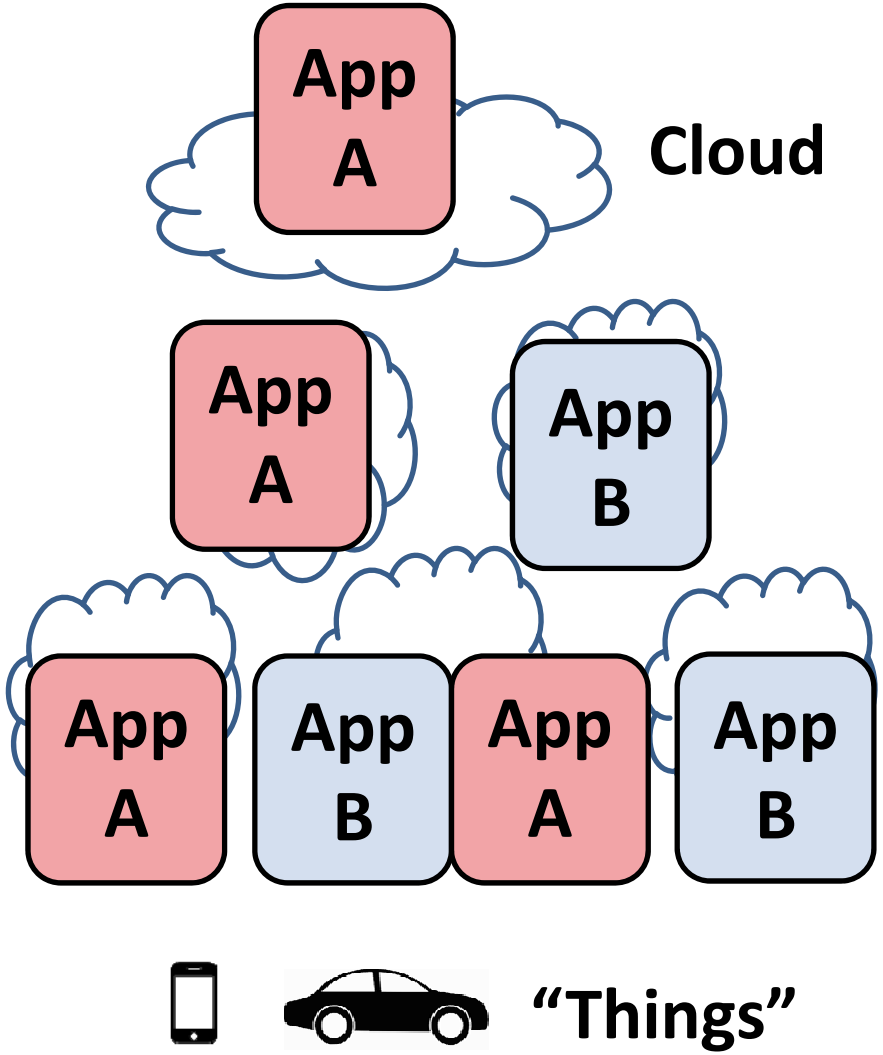

Open Fog Reference Architecture

Source: The OpenFog Consortium 2017

Overview

- Motivation
- Introduction
- Openness and standardization
- **Research**
  - FogStore – Data store for fog
  - EmuFog – Fog computing emulator
- Conclusion and outlook

Differences to Cloud

- Administrative domains:

  Fog is not in one domain
- Locality is first-class citizen in fog
- Fog nodes can be highly heterogeneous

- This challenges existing made-for-cloud systems

Data Stores in Fog

- Data stores move to fog

  ➔ for same reasons as compute moves

- Design challenges
  - Replica placement
  - Consistency guarantees

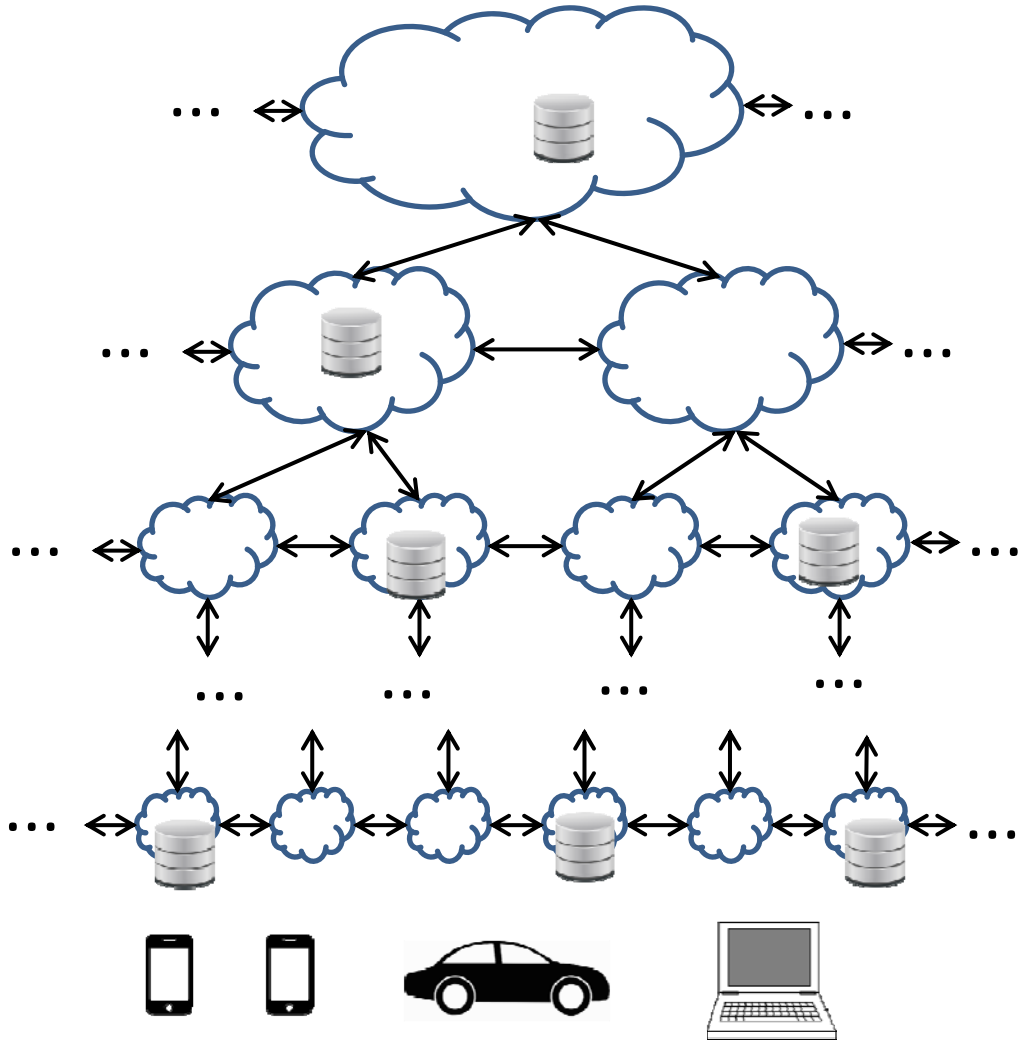

Cloud Data Stores

- Current data stores are designed for data centers
  - Replica placement → not place in same server rack
  - Consistency management → indifferent of location
  - Apache Cassandra, MongoDB, …
- Fog is different
  - Might not have server racks
  - Locality of sources and sinks is important (➔ low latency)

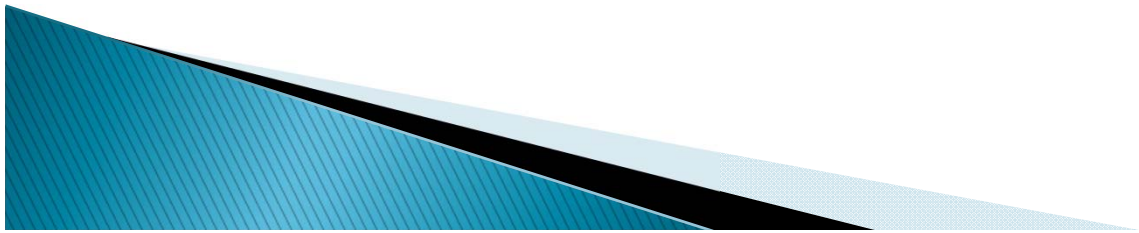

Example: Autonomous Cars

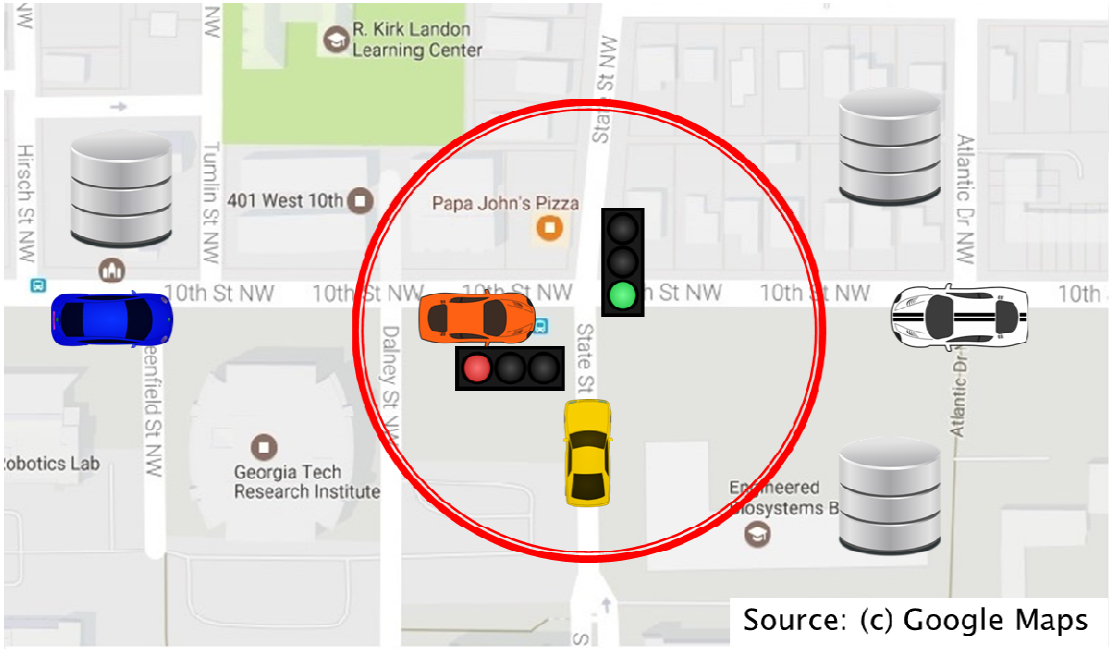

- Cars in red circle need consistent state of traffic light
  - Not both see green at the same time!
- Others can read different states for non-critical tasks
  - E.g., adapt speed to save energy

FogStore: Contributions

- Allows applications to specify a context-of-interest (coi) as per its semantics
  - Coi-aware replica placement for low latency access to replicas
  - Higher tolerance to geo-correlated failures
- Coi-aware per-query consistency mapping
  - Serializability to queries within coi
  - Eventual consistency to queries outside coi (focus on low latency)

FogStore: Architecture

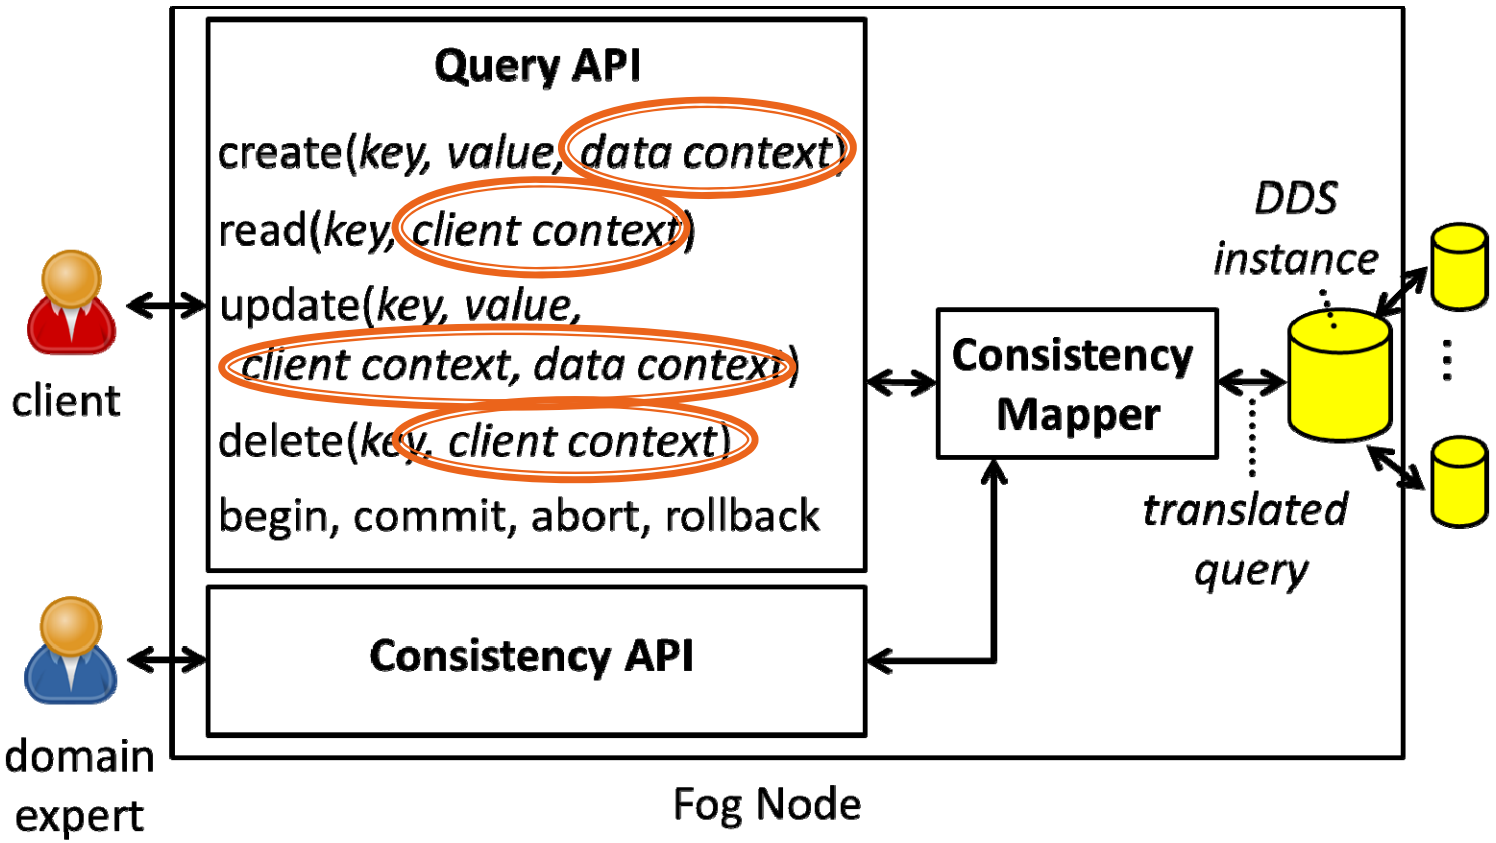

Ruben Mayer, Harshit Gupta, Enrique Saurez, and Umakishore Ramachandran. 2017. FogStore: Toward a Distributed Data Store for Fog Computing. In Proceedings of the 1st IEEE Fog World Congress (FWC '17). 6 pages.

Wanted: Test Platform for Fog

- **Fog architects:**
  "How to layout the fog architecture for a specific application workload?"
- **Application developers:**
  "How does my application behave in a specific fog architecture?"
- No real fog publicly available
  ➔ Need for a test environment

Background

- Three possibilities:

  (1) real-world test beds, (2) simulation, (3) (network) emulation

- Real-world testbeds not available
- Simulation has shortcomings:
  - Based on many assumptions that need to be modeled
  - Applications need to be adapted/modeled in simulation framework
- Network emulation is suitable:
  - Run original applications
  - Only network is emulated
  - Realistic performance measurements of applications

Shortcomings of Existing Emulators

- Existing emulation frameworks

  lack support for fog

  - Allow to load a network topology and emulate it
  - Missing part (the hard one): placing fog nodes and apps

    ➔ This has to be done “by hand”

- Goal: Provide better support!

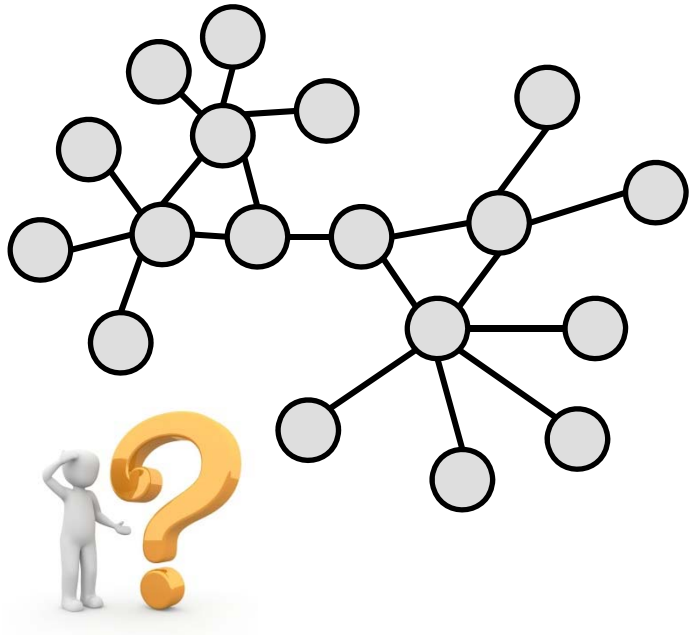

EmuFog: Contributions

EmuFog, an emulation platform for fog that supports...

- ...automatic configuration and emulation of fog infrastructure
- ...execution of natural fog applications provided as Docker containers
- ...scaling across a distributed cluster of host machines
- ...extensibility to user-defined fog node placement strategies

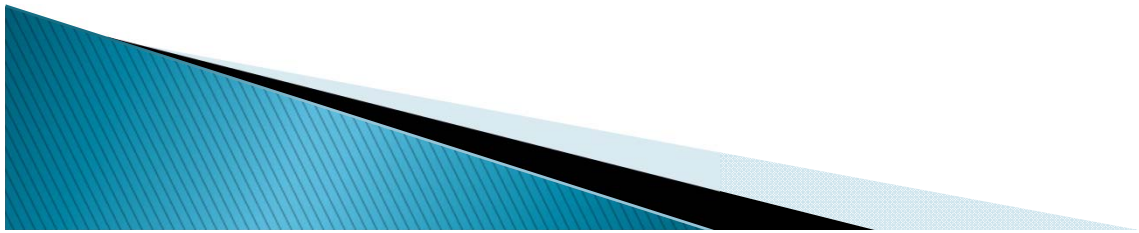

EmuFog Workflow

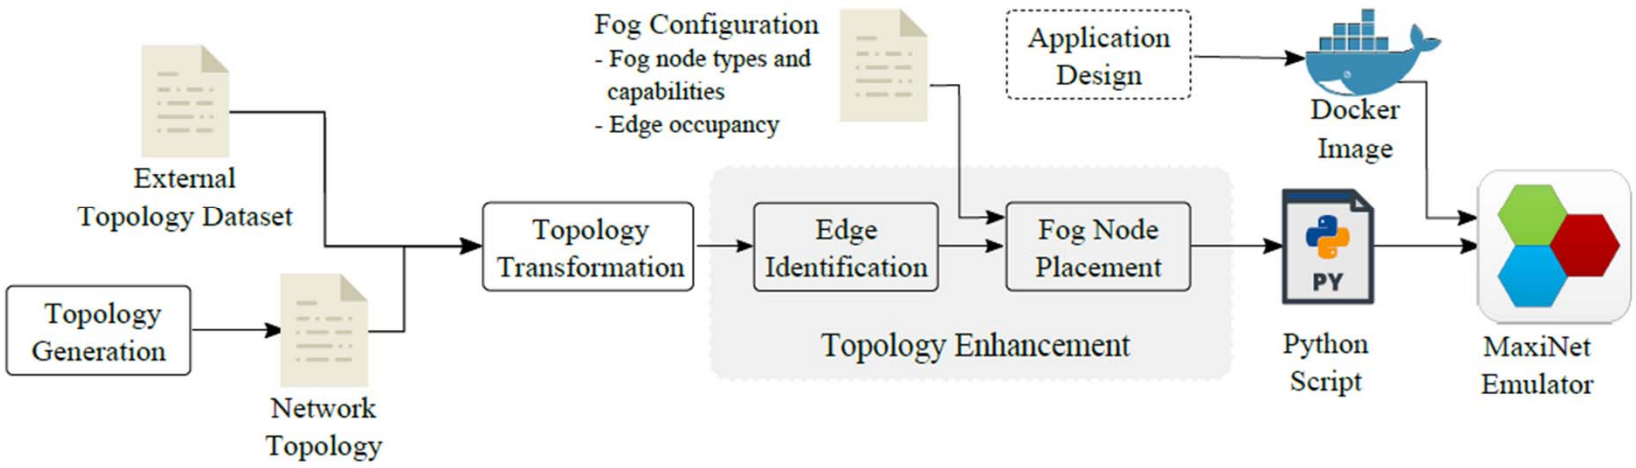

- Latest version of EmuFog available on GitHub: https://github.com/emufog/
- Paper available:
  Ruben Mayer, Leon Graser, Harshit Gupta, Enrique Saurez, and Umakishore Ramachandran. 2017. **EmuFog: Extensible and Scalable Emulation of Large-Scale Fog Computing Infrastructures**. In *Proceedings of the 1st IEEE Fog World Congress* (FWC '17). 6 pages.

Overview

- Motivation
- Introduction
- Openness and standardization
- Research
- **Conclusion and outlook**

Conclusion

1. What is fog?
2. What can fog do for me?
3. What can I do for fog?

Outlook

- Fog matches with recent trends
  - It is all about distribution and decentralization
  - Example: Blockchain
    - Decentralized consensus, cryptocurrencies, ...
    - Can it be used for accounting, provenance, etc.?
- Fog poses challenges on existing made-for-cloud systems
  - Stream processing, machine learning,...
  - Constrained resources, local processing,...

Thank you!

Time for Q & A.