# ECE 590
# COMPSI 590

Special Topics: Edge Computing

August 27th, 2018





# About the Instructor (1/2)

- Started at Duke in July
- Previously: Associate Research Scholar, Princeton University, Electrical Engineering

- Ph.D. Columbia University, Electrical Engineering
- M.Sc., B.Sc. University of Ottawa, Canada




# About the Instructor (2/2)

- Worked in industry before, during, and after all degrees

IBM

Disney Research SCIENCE & PLAY

D E Shaw Research

Canadian Intellectual Property Office

Telcordia Technologies

DEFENCE R&D DÉFENSE

SiriusXM SATELLITE RADIO

# Introductions



# In this Lecture

- Introduction to edge computing, part 1
- Course structure and syllabus
- Introduction to edge computing, part 2

# Cloud: Computing in Datacenters

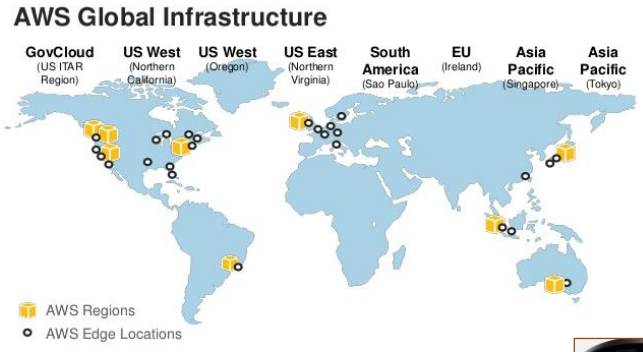


- AWS: 44 locations worldwide, MS Azure: 30
- For emerging applications: **fundamental limitations** in **latency**, **bandwidth**

7



# Edge/Fog: Computing Closer to the Users

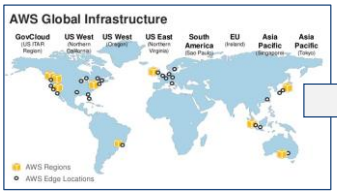


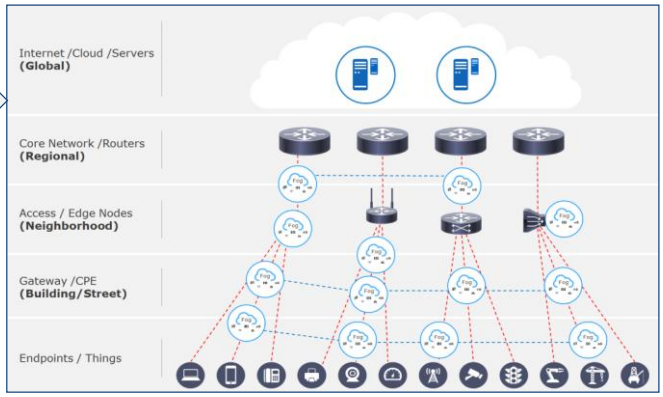


- Data processing, business logic, decision-making at multiple points in the hierarchy

Smart city IoT deployments: computing in buildings, neighborhoods, zip codes

8

# Edge Computing: Cloud to Fog

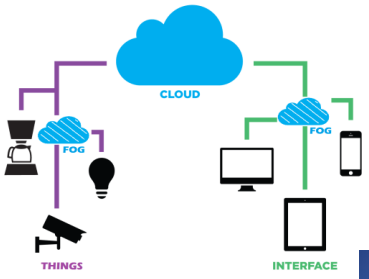


- Distribute computation, storage, communication and control services immersively closer to end-users **along the Cloud-to-Things (C2T) continuum**

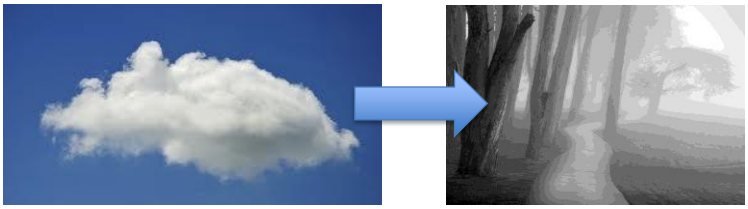

9



# Edge Computing: Core to Edge

- An architecture that uses **one** or a **collaborative multitude** of **end-user clients** or **near-user edge devices** to carry out a **substantial amount** of computation, storage, communication, and control
- **Core → Edge**

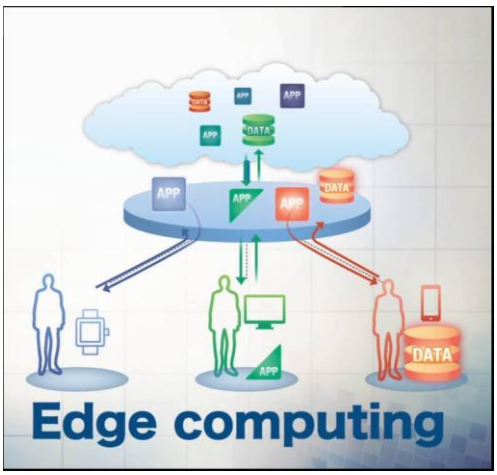


10

# End-User Clients or Near-User Edge Devices: A Range of Options

- Gateways, stationary or mobile
- Set-top boxes
- Servers, cloudlets
- Mini-datacenters

- Different properties


11



# Substantial Amount of Work Near the Users: Range of (De)centralization Alternatives

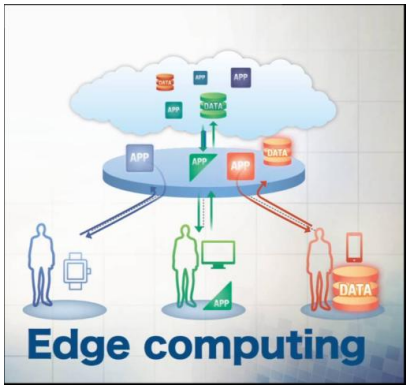


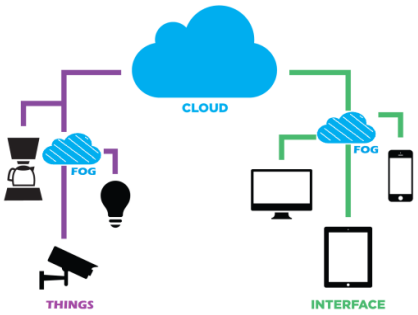


12

# What Does Edge Provide?

- Latency, bandwidth
- Cognition – advanced intelligence close to the users
- Privacy

- **Improve the performance of existing applications** and **enable new ones**

13



# Why Edge Computing? (1/2)

- "Next billion-dollar tech market" ANDREESSEN HOROWITZ
- "Most interesting part of cloud computing"
- Fundamental enabler of the pervasive computing vision



14

# Why Edge Computing? (2/2)

- Advanced intelligence close to users is cool
- Emerging field
- On the boundary of domains not usually examined together
- Requires re-thinking implicit assumptions

15

Duke UNIVERSITY

# My Life On the Edge: Background

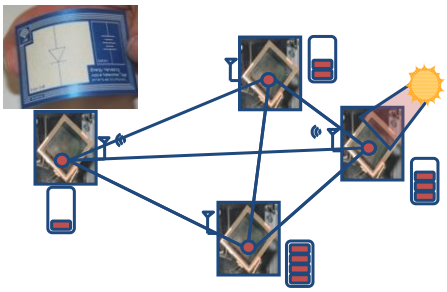

D E Shaw Research

- See edge as an enabler for the true potential of the Internet of Things
- Unusual IoT-to-cloud perspective

16

Duke UNIVERSITY

# My Life on the Edge: 2016-2018

- OpenFog Consortuim technical committee member, working group co-chair, board of directors alternate
- Contributor to the IEEE 1934 Fog Computing Standard
- Associate director of Princeton Edge Lab

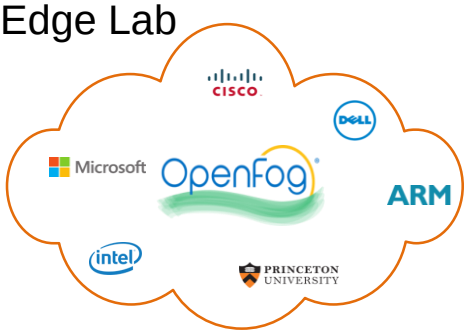


17

Duke UNIVERSITY

# In this Lecture

- Introduction to edge computing, part 1
- Course structure and syllabus
- Introduction to edge computing, part 2

18

Duke UNIVERSITY

# Course Logistics

- Lecture times: **16:40-17:55 PM**, Mondays and Wednesdays
  - No lecture October 1st. Lecture makeup date TBD.
- Office hours: Mondays **09:00-11:00** AM, 2471 CIEMAS
- Readings before every class

19

Duke UNIVERSITY

# Course Structure

1. Background, core architectural principles, challenges and opportunities
   - Goal: enough "lay of the land" for the research projects
2. Domains where edge is particularly exciting
3. Notable papers in the field and techniques that need to be re-imagined for the edge

20

Duke UNIVERSITY

# Upcoming Lectures

- **8/29:** The origins and the current state of edge computing
- **9/03:** Edge helping the IoT
- **9/05:** Edge helping higher-end mobile devices
- **9/10:** Edge helping the cloud

21



# Grading

- Quizzes: 20%
- Research paper presentation: 20%
- Research project: 50%
- Participation in class discussions: 10%

22

# Please See the Syllabus for …

- Quizzes
- Research paper presentation
  - Spread over weeks of **September 12th – October 15th**
  - Sign-up Google Sheets will be available later today
- Participation in class discussions

23



# Research Project in Edge Computing

- 50% of the grade
- Teams of 1-2 people
- Research project
  - Generate and thoroughly validate a new idea
- Ideal outcome: work leading to a paper that can be published in a top venue of the field
  - But, its research – not all explorations are fruitful
  - High-risk high-rewards > incremental improvement

24

# Bonus Points: Connect Research Project to Duke


Duke LEMUR CENTER

- 2 to 5 extra points

25

Duke UNIVERSITY

# Research Project: Some Related Equipment Available

26

Duke UNIVERSITY

# Research Project: Timelines

- Teams established: Friday **September 21st**
- Proposal due: Monday **October 1st**
  - Will talk about the format next class
- Progress report due: **Friday October 26th**
- Final presentations: **weeks of November 19th and 26th**
- Final report due: **Friday December 14th**

27



# In this Lecture

- Introduction to edge computing, part 1
- Course structure and syllabus
- Introduction to edge computing, part 2

28

# Why Edge Computing?

- "Next billion-dollar tech market" ANDREESSEN HOROWITZ
- "Most interesting part of cloud computing"
- Fundamental enabler of the pervasive computing vision



29

Duke UNIVERSITY

# Next Billion Dollar Tech Market

ANDREESSEN HOROWITZ

30

Duke UNIVERSITY

# Research Themes in Edge Computing (1/5)

- What should be placed where?
  - Computing, storage, decision-making
- Restructuring applications and algorithms to fit edge/fog conditions

31



# Research Themes in Edge Computing (2/5)

- Edge in specific applications
  - Augmented reality, virtual reality
  - Networks of drones
  - Autonomous driving
  - Invited speaker: edge computing for autonomous trucking

32

# Research Themes in Edge Computing (3/5)

- Data processing that preserves privacy
- ML training on the edge
- ML inference on the edge
- Reinforcement learning on the edge



# Research Themes in Edge Computing (4/5)

- Thinking across applications, devices, platforms
- Operating across multiple computing / storage / control / decision quality levels

# Research Themes in Edge Computing (5/5)

- Multi-tenancy
- Multi-device operation
- *Uberization* of resources
- …

35

Duke UNIVERSITY

# Edge Computing at Chick-fil-A (1/2)

July 2018


36

Duke UNIVERSITY

# Edge Computing at Chick-fil-A (2/2)

Azure


AWS

Google Cloud


Cloud-fil-A


37

Duke UNIVERSITY

# Recap

- Edge computing: definitions
  - ➢Different devices
  - ➢Different degrees of application centralization
- Logistics of the class
- Research themes

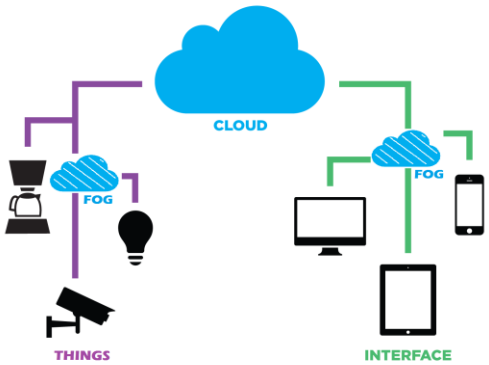


38

Duke UNIVERSITY

# Next Class and Homework (1/2)

- *"IoT meets the cloud: the origins and the current state of edge computing"*
- Reading materials for the class



# Next Class and Homework (2/2)

- Choose the dates for your presentations
- Start thinking about your research project
  - My office hours: 9-11 AM Mondays
  - Make an appointment if you need to see me sooner