2017 Winter Review

2017 Winter Review: Lab Overview and Update

John Ousterhout
Faculty Director

Thank You, Sponsors!

EQUINIX

Special Thanks To...

Platform Lab Vision

- **New platforms enable new applications:**
  - Relational databases → Enterprise applications
  - HTTP + HTML + JavaScript → Internet commerce
  - GFS + MapReduce → Big Data (large-scale analytics)
  - Smart phones + GPS → Google Maps, Uber, …
- **Mission:**
  - Define the next generation of platforms
  - Stimulate new classes of applications
  - One or two flagship projects at any given time
- **Current focus: platforms for large-scale control**
  - Big Control Platform
  - Granular Computing Platform

What is a Platform?

- **General-purpose hardware or software substrate**
- **Simplifies construction of a class of applications (or higher-level platforms)**
  - Solves common problems
  - Usually introduces (simplifying) restrictions
- **Example: MapReduce**
  - Applications: large-scale analytics
  - Problems solved: hides latency, handles slow/crashed servers
  - Simplifying restrictions: two-phase decomposition, large sequential accesses

Platform Lab Faculty

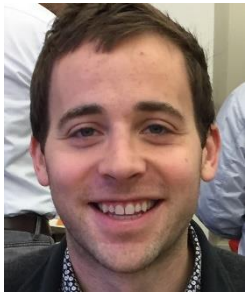

Peter Bailis

Bill Dally

Sachin Katti

Christos Kozyrakis

Phil Levis

Nick McKeown

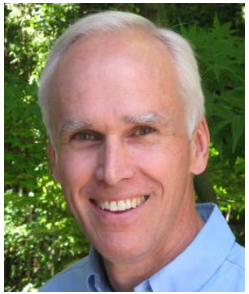

John Ousterhout (Fac. Director)

Guru Parulkar (Exec. Director)

Balaji Prabhakar

Mendel Rosenblum

Keith Winstein

Matei Zaharia

Platform Lab News

- **Additional faculty:**
  - Peter Bailis (Big Data, databases)
  - Balaji Prabhakar (Networking)
  - Matei Zaharia (Big Data, systems, Spark creator)
- **NSF Expedition proposal on Big Control**
  - Promoted to second round
- **Definition of granular computing platform**
- **Promotions and awards:**
  - Sachin Katti: tenure
  - Christos Kozyrakis: ACM Fellow
  - Dinesh Bharadia: MIT TR35, Marconi Young Scholar
  - Best paper awards: ISCA, MOBICOM, Sensys (runner-up)

Recent/Soon-To-Be Graduates

Ankita Kejriwal	Secondary indexes for RAMCloud	Google
Camilo Moreno	Communication in many-core chips	Intel Labs
Kanthi Nagaraj	Programmable network fabrics	
Hang Qu	Task scheduling for cloud analytics	

Big Control

- **Enormous swarms of devices:**
  - Collaborative
  - Centrally controlled
- **The morning commute**
  - 1M+ self-driving cars
- **Large distribution center**
  - 10,000+ indoor drones
- **Disaster recovery**
  - 1000+ drones, automated ground-based vehicles
  - Coordinated mapping and search
  - Data integration
  - Mobilized response

Big Control, cont'd

- **Interesting properties:**
  - Scale
  - Collaboration
  - Latency
- **Control will become more centralized:**
  - Easier application development
  - More powerful features (e.g., integrate back-end datasets)
  - More robust!
- **Lab goal: define and enable the Big Control paradigm**
- **Create two new platforms:**
  - Big Control Platform (BCP)
  - Granular Computing Platform

Multi-Level Control Loops

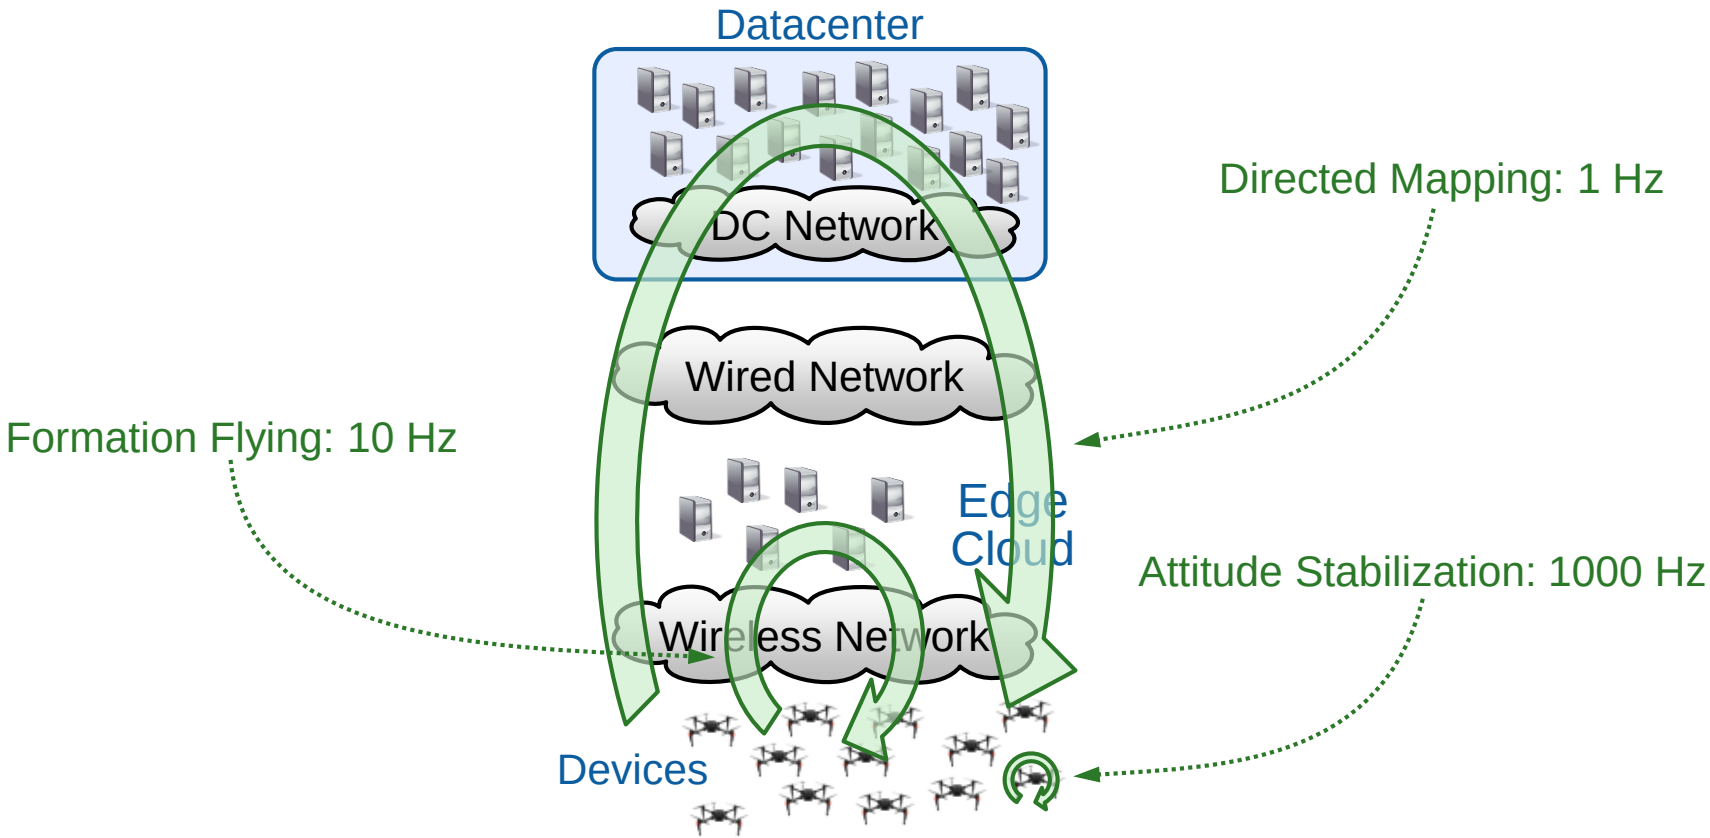

System Overview

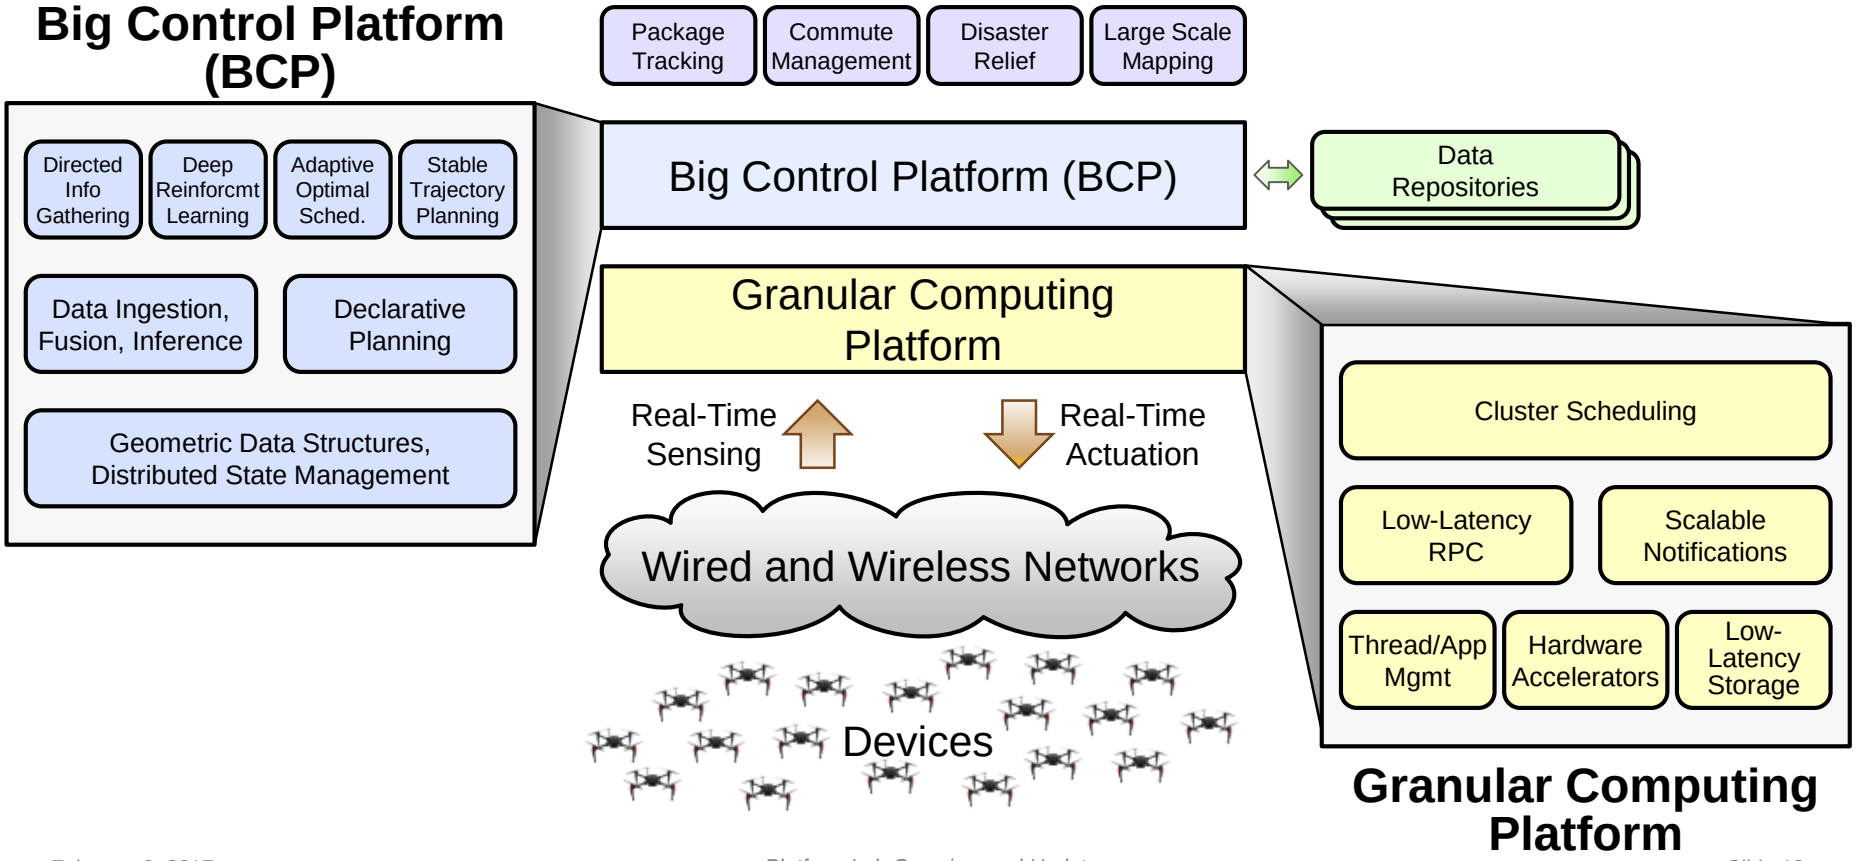

Big Control Platform

- **What is MapReduce for control?**
  - Solve common problems
  - Simple framework
- **Example: declarative planning**
  - Specify plans in high-level language (~ SQL)
  - Generate device-specific commands automatically (~ query optimizer)

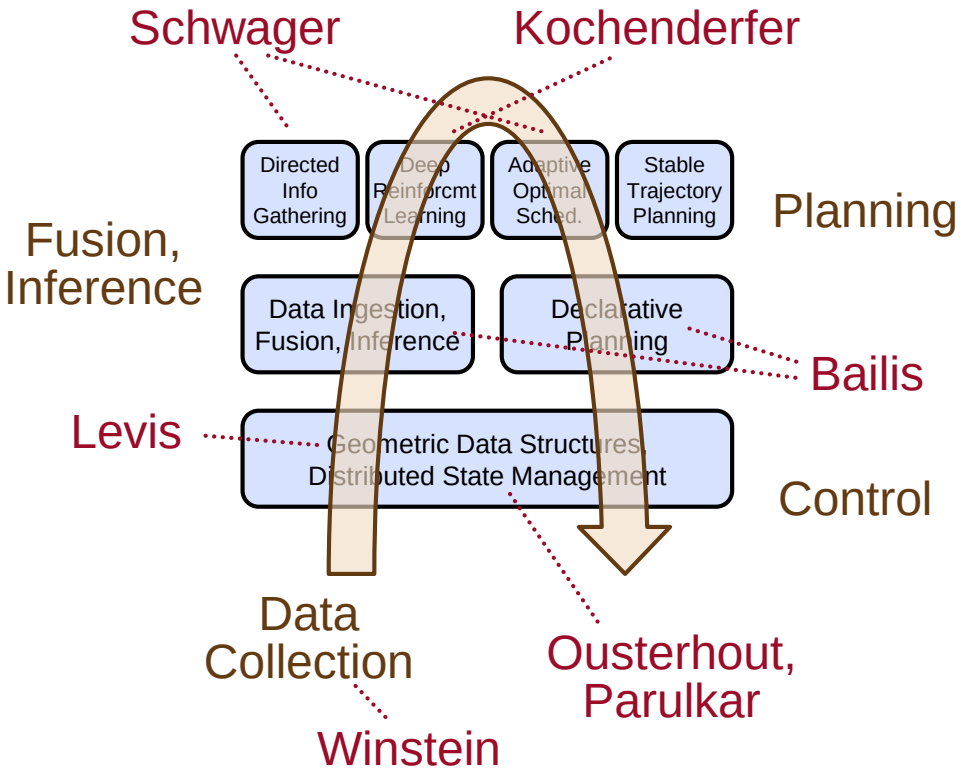

Granular Computing

- **Big Data: process data in large sequential chunks**
- **Big Control: process data in very small chunks**
  - Example: 10,000 devices, updates all-to-all every second
  - Each device notification triggers internal events for fusion, inference
- **Granular computing:**
  - Support tasks lasting 10ms → 1μs
  - Efficient instantiation, communication
  - Short duration => large numbers
  - Highly elastic
  - Must coexist with traditional large tasks
- **Related trends: micro-services, lambdas**

Granular Computing Examples

- **Remote task with durable results: 20 µs**
- **Local task with volatile results: 500 ns**
- **Fanout to 100+ threads, interactive results:**
  - Real-time event-driven inference
  - Exploratory data analysis
  - Instant video encoding

Granular Computing Platform

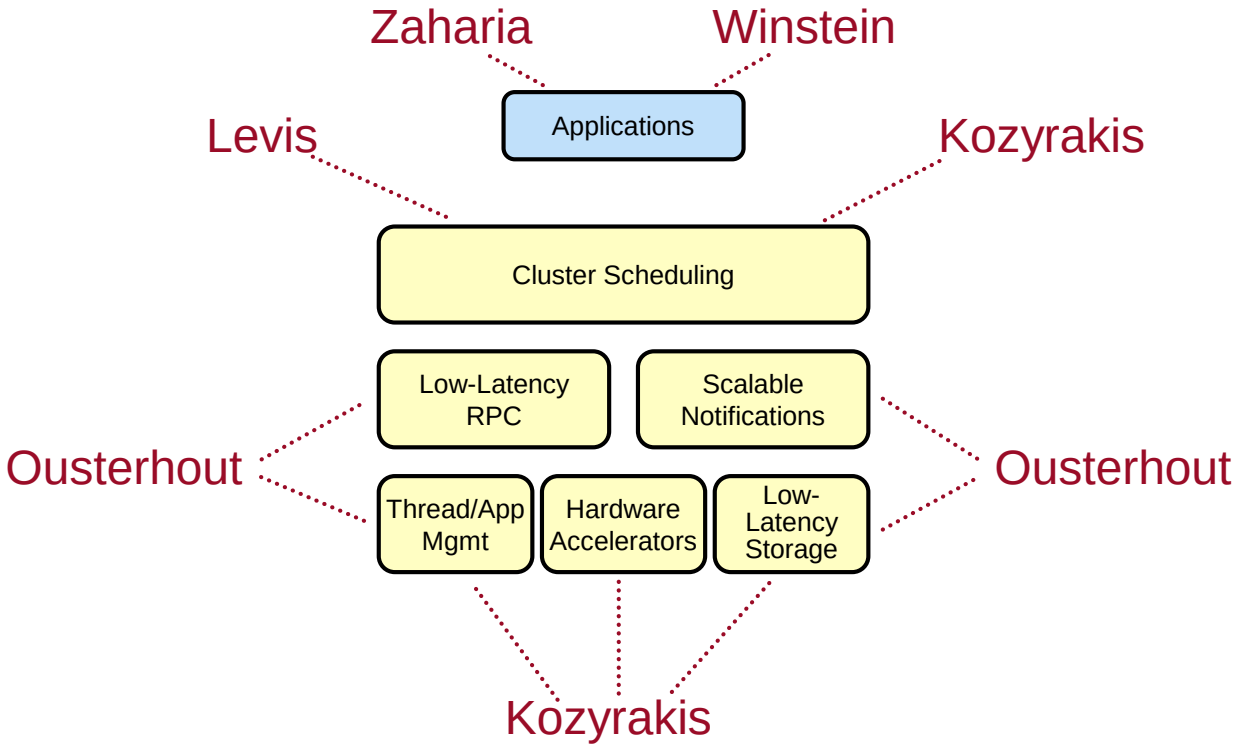

Networking Platform

- **Rearchitect the network for control applications:**
  - Ultra-reliable
  - Ultra-low and predictable latency
  - Secure, robust
- **Slicing architecture:**
  - Decouple control and data planes
  - Virtualize network substrate (multiple control/data planes)
  - New algorithms for allocating resources among competing slices

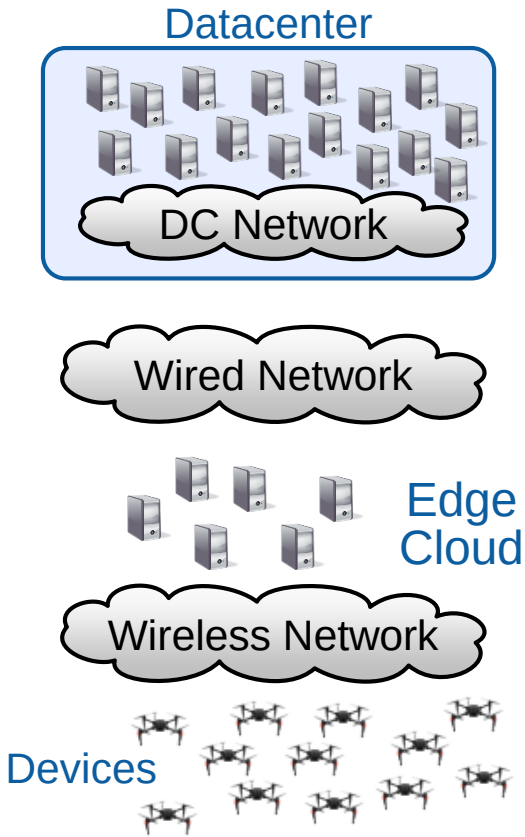

Project Plan

- **Phase 1 (Years 1-2): exploration, infrastructure**
  - Simple control applications for learning
  - Prototypes of BCP subsystems, Granular Computing Platform
  - **Milestone: ready to design BCP**
- **Phase 2 (Years 3-4): BCP version 1**
  - Integrated version of BCP
  - Revisions to Granular Computing Platform
  - Port a few applications
  - **Milestone: BCP runs a few simple applications**
- **Phase 3 (Year 5): capstone demonstration**
  - Disaster recovery demo, possibly others
  - Continued evolution of BCP, Granular Computing Platform
  - **Milestone: capstone demo**

Next Steps: More Students!

- **Long lead-time (new PhD admits commit early)**
- **Planning for heavy recruiting this spring**
- **Students rotate in 2017-2018, align in Spring 2018**
- **Big Control seminar in Fall 2017**
- **Also: funding always a challenge; hoping to win Expedition competition**

Agenda

9:15 — 10:30am **Big Control Platform (BCP) Abstractions and Services**

- Directed Information Gathering — Riccardo Spica
- Deep Reinforcement Learning for Device Control — Blake Wulfe
- Distributed Geometric Data Structures — Philip Levis

10:30am — 11:00am **Break**

11:00 — 12:15pm **Self Driving Programmable Networks**

- Data Driven Networking — Sachin Katti
- Self Driving Networks — Balaji Prabhakar
- Weld: Fast Data Analytics on Modern Hardware — Shoumik Palkar

12:15 — 12:45pm **Lightning Talks by Students**

12:45 — 2:00pm **Lunch and Poster Session**

Agenda, Cont'd

2:00 — 2:45pm **Invited Sponsor Talks**

- Network Management beyond SDN — Jeff Mogul, Google
- Potential Big Control Use Cases — Ayush Sharma, Huawei

2:45 — 3:45 pm **Panel on Granular Computing**

Keith Winstein, Christos Kozyrakis, Philip Levis, John Ousterhout

3:45 - 4:15pm **Break**

4:15pm — 5:30pm **Granular Computing Platform**

- NanoLog: A Nanosecond Scale Logger — Stephen Yang
- RAIL: Predictable, Low Tail Latency for Flash-based SSDs — Heiner Litz
- TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory — Mingyu Gao

5:30 — 5:45 pm **Wrap up**

5:45 — 7:00pm **Reception**

Conclusion

- **Platform Lab program now fully formed**
- **Time to execute**

How can we collaborate with industry for this research agenda?

Questions/Discussion