Planes, Trains, and Data Centers

Mendel Rosenblum

PLATFORMLAB

Talk Agenda

- What will be new applications that need data centers?
    - Controlling autonomous entities
- Some brainstorming about controlling autonomous devices
    - A funding proposal
- What will the data center be doing in the future?
    - Big Control
- Summary of "Infrastructure for Collaborative Device Swarms"
    - NSF Expedition Pre-Proposal

Disclaimer: Most material in talk from someone else

- NSF Pre-Proposal: Infrastructure for Collaborative Device Swarms
    - John Ousterhout
    - Balaji Prabhakar
    - Pat Hanrahan
    - Mac Schwager
    - Guru Parulkar
    - Dan Boneh
    - Sachin Katti
    - Mykel Kochenderfer
    - Christos Kozyrakis
    - Phil Levis
    - Keith Winstein

Background: Search for next-gen data center apps

- Existing data center driving applications
  - High Performance Computing (HPC) - Parallel computing
    - Mainly simulations
  - Communication
    - Messaging, social networks
  - Big Data
    - Search, machine learning with neural networks
  - Cloud Computing
    - All non-embedded computing
- Wild and crazy idea: All embedded computing
  - **Big Control**
    - Autonomous vehicles, drones, Internet of Things, ….

Focus: Centralized versus distributed control

- Some previously implemented distributed system have been move to more centralized implementations
- Internet
  - Early vision: Data center in middle of the country
  - First app: Distributed email system
  - Current incarnation: Gmail running in a data center in Council Bluffs**,** Iowa
- Network packet routing
  - Version 1: Border gateway protocol (BGP) - Fully distributed route state computation
  - Current hot idea: Software-Defined Networking (SDN) - Centralized control
- What about autonomous entities?
  - Can we can get examples from nature?

What about nature?

- Most believe there is not centralized control of life
  - Notwithstanding Intelligent Design
- Kirstin Petersen's bio-inspired robots
  - **Robot Collectives** - Control based on social insects
    - Termites - Simple control systems but good at construction tasks
      - Highly fault tolerant and scalable (e.g. Deborah Gordon's Ants and TCP slow start)
      - Fully distributed
    - Robots have sensors and actuators
    - Simple control without communication
- Argument against:
  - Impact communication has had on humans
    - Language, books, telephone, Internet

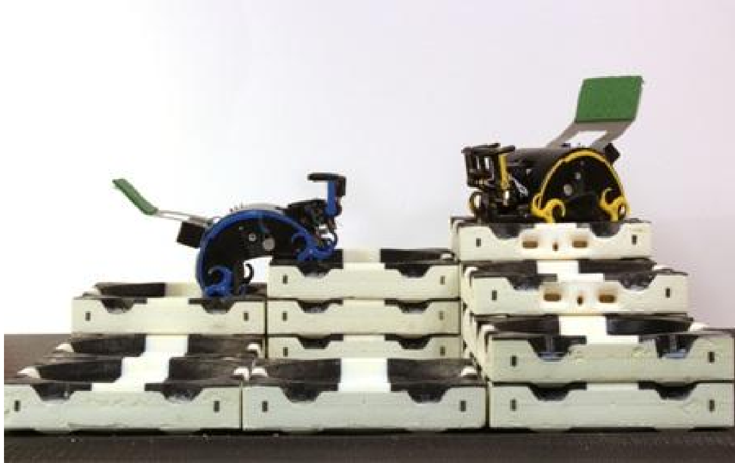

http://www.eecs.harvard.edu/~kirstin/Pics/Robots/Kali&Isis.jpg

Robot to Robot Communication

- Mac Schwager's **robot swarms**
  - Communicate to self-organize to do something
    - Example: Drones searching a disaster area
  - Robots build shared knowledge of system-wide state
    - Used to guide local decision making
    - Wonderful CS problem: Distributed consensus algorithms
  - Ideal disk communication model
    - Can use radio to communication with nearby drones
- Low power radio not a requirement - Could communicate to a controller
  - Pat Hanrahan's observation: Motors used way more power than electronics
    - Radios/CPU: Milliwatts
    - Motors: Watts

http://sites.bu.edu/msl/files/2013/07/OilCropForestFig.jpg

Communication latency and reaction times

- Some times:
  - Current Internet round-trip: 100s of milliseconds
  - Speed of light to Iowa and back: 10s of milliseconds
  - Human reaction times: ~1 sec
  - Anti-lock brakes reaction time: a few milliseconds
  - Anti-collision systems: a few milliseconds
- Max's drones:
  - Safety systems kept drones one meter apart
- Speed of light argues that not all control can be centralized
  - But much control can be centralized

Centralized versus distributed control tradeoff

- Distinction in proposal:
  - Reflex behaviors - Control decision made by the device itself
  - Planned behaviors - Control decision from centralized control

- Reflex behaviors
  - Safety systems and ???

- Planned behaviors
  - Global optimization and ???

http://roa.h-cdn.co/assets/16/19/980x490/landscape-1462981879-13123290-10153962911310289-5265986788438345559-o.jpg

Existing Big Control Example: Waze

- Popular GPS-based geographical navigation application
- Deploy sensors in many vehicle on the roadways
  - Cell phone-based sensors collect location, speed, accidents, road hazards, closures, traffic cameras, police
    - Incent users using good directions and badges
- Determine global state of roadway systems from current sensors and historical behavior
  - Make accurate prediction of travel time and best route (incentive)
- Control: Send updated navigation instructions to users
  - More of hint than a command (Reflex component uppity)
    - Slow start algorithm

Balaji's Societal Networks

- Transportation networks - Bus and trains
  - Incomplete sensor information
    - Time user swipe in at a stop
    - Time user swipe out at a stop
    - When carriages run
  - Can restructure hidden state of system
    - Can infer carriage loads, wait times, etc.

- Found basic technique to be widely applicable
  - Example: From network round trip packet times compute network switch queue depths

Experience with inference

- Erroneous input must be handled
  - Sensors can and do return bogus values (e.g. GPS position in downtown canyons)
  - Data cleaning step needed
- State reconstruction can be expensive
  - Example: Solving large set of linear equations
    - A traditional HPC workload
  - Speed of reconstruction could be limit applications
- Looking like we might need some significant computing resources

Applications for state reconstruction

- Dashboard
  - Display a visual representation of state providing the operator and users with insight
    - Use human processing to analyze the state
      - Experience: Demonstrated transit systems state previously only visible via intuition
    - Highlight events of interest
      - Notification system
- Scheduling and control
  - Optimize the system. Examples:
    - Unbunch the busses, respond to flash mobs
    - Update the routing to avoid switch queues

What if/counterfactual questions

- Can initialize a simulator with the current state and dynamic behavior
  - Explore states impractical to drive the system into
- Example:
  - What would happen if we shut down a lane this week? (e.g. Fort Lee, New Jersey)
  - What would the effect hurricane evacuation traffic?
  - What would happen if we added this new workload to the network?
- If this doesn't need a data center worth of compute, nothing does

Conclusion

- Having near infinite computing resources accessible via 10-100 ms round-trip communication time will be useful.
- And we are part of the:

PLATFORMLAB

- Just need a platform to make this easy to do...

Big Control: Infrastructure for Collaborative Device Swarms

Passé: Big Data ⇉ Exciting: Big Control

- Example applications:
  - Millions of self-driving cars - Maximize throughput, avoid congestion
  - 10,000 indoor drones tracking contents of large distribution center
  - Natural disaster response with drones scanning millions of acres
- Platform Challenges
  - Scale
  - Collaboration
  - Low latency

The Proposal Figure

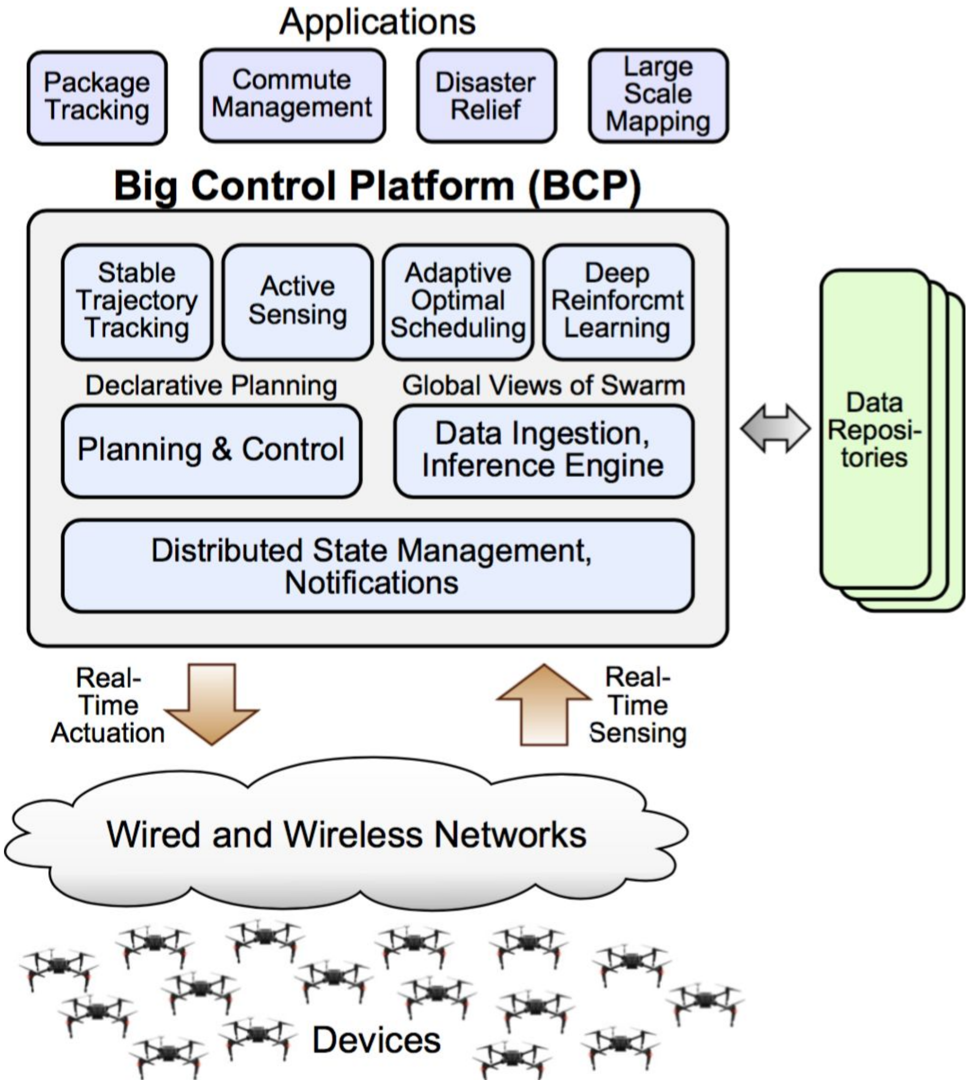

Big Control Platform

- Scalable State Management and Notifications
  - High sensor input rate
    - High performance storage
    - Notification
    - Real-time fault-tolerance
  - Low-Latency Data centers
    - Low-latency RPC system
    - Core-aware thread scheduling
    - Scheduling at data center scale
    - Low latency storage

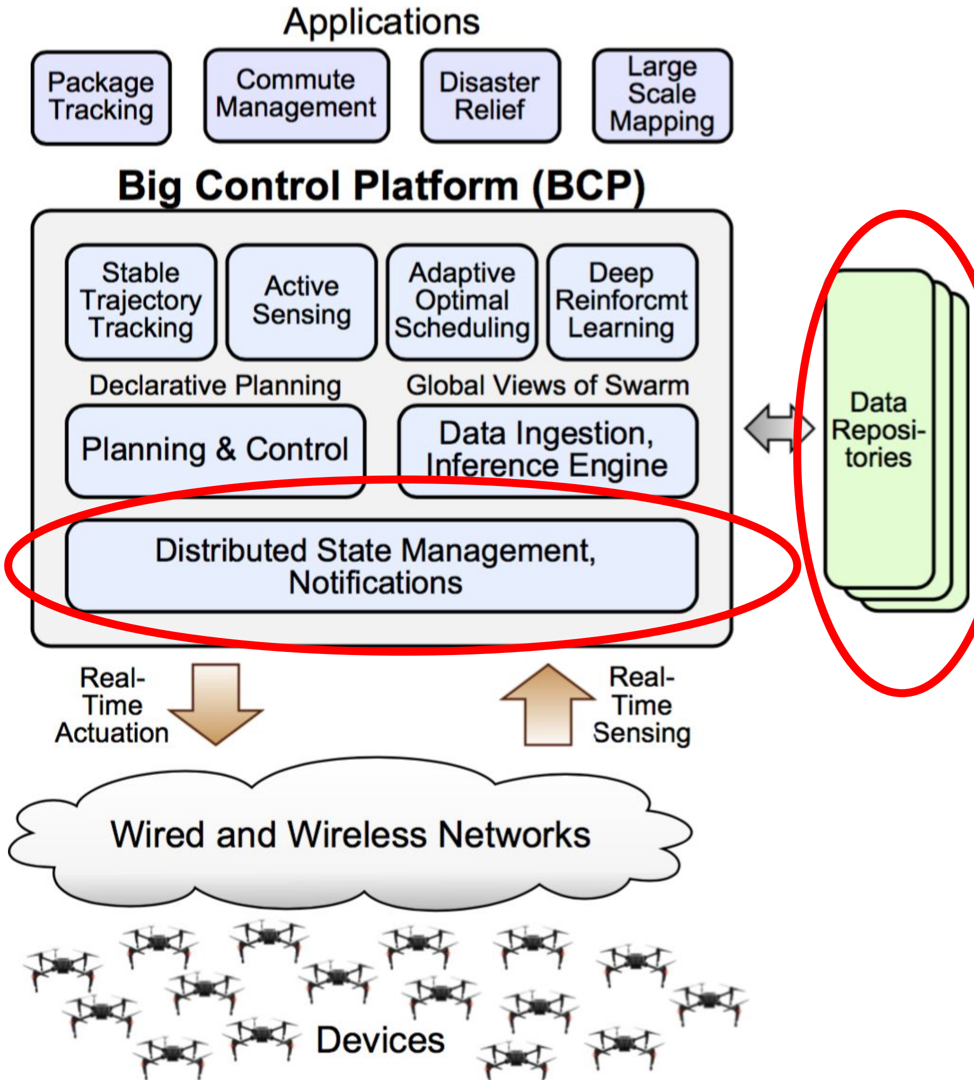

Big Control Platform

- Data Ingestion and Inference
    - Capture state of device swarm
    - Example: Large, octree point cloud-like

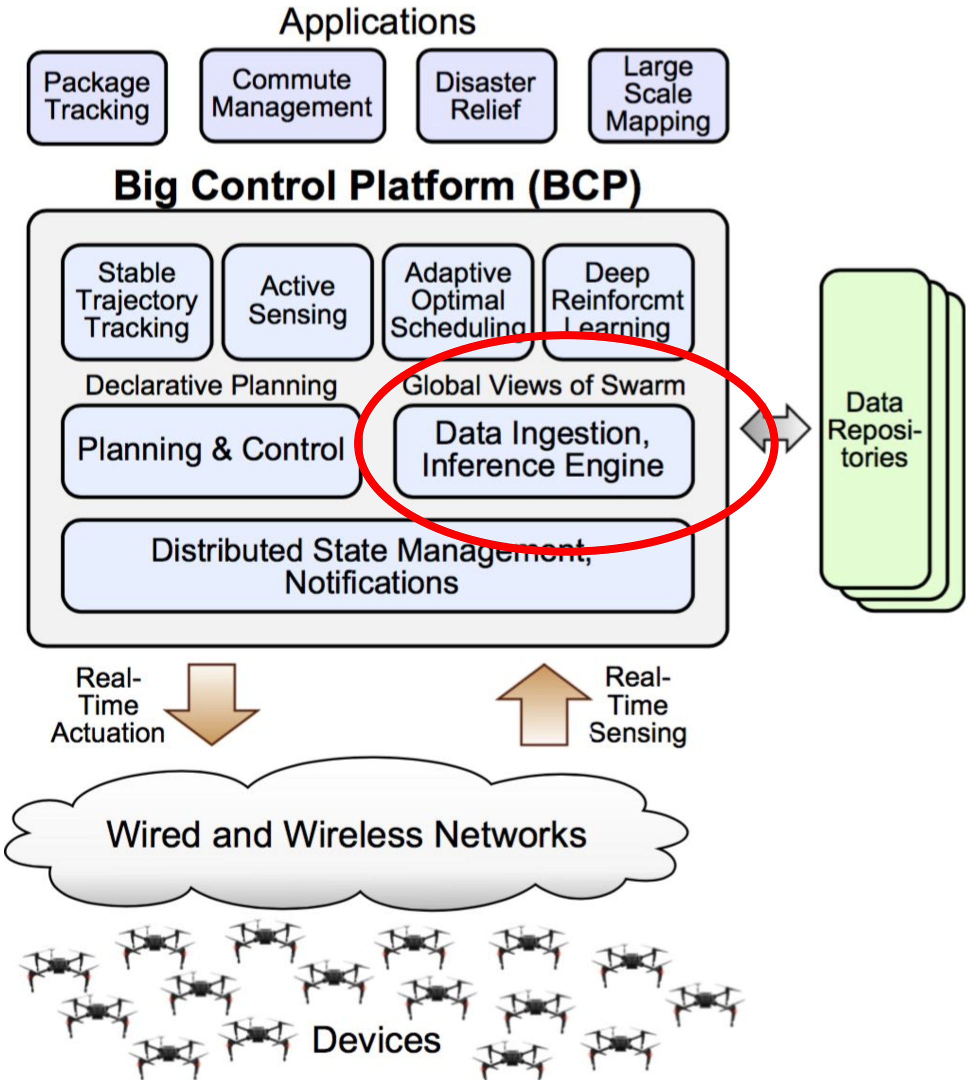

Big Control Platform

- Declarative Planning
  - Higher level declarative interface

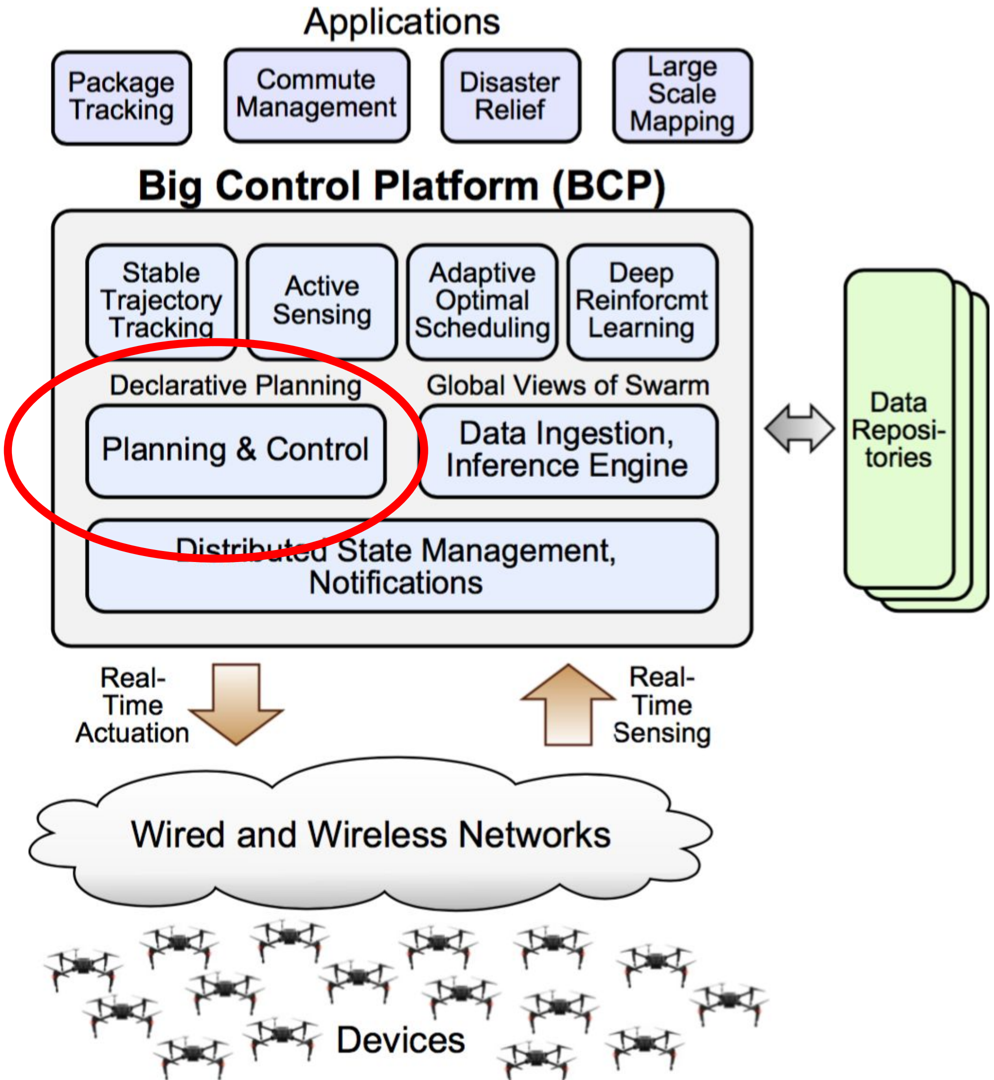

Big Control Platform

Predictable Low Latency Wide Area Network

- Decouple the control and data planes of the network
- Virtualize the hardware network
- 20ms => 1ms

Swam Security

- Secure connectivity
  - Cryptographic integrity
- Prevent sensor attacks
  - Byzantine behavior

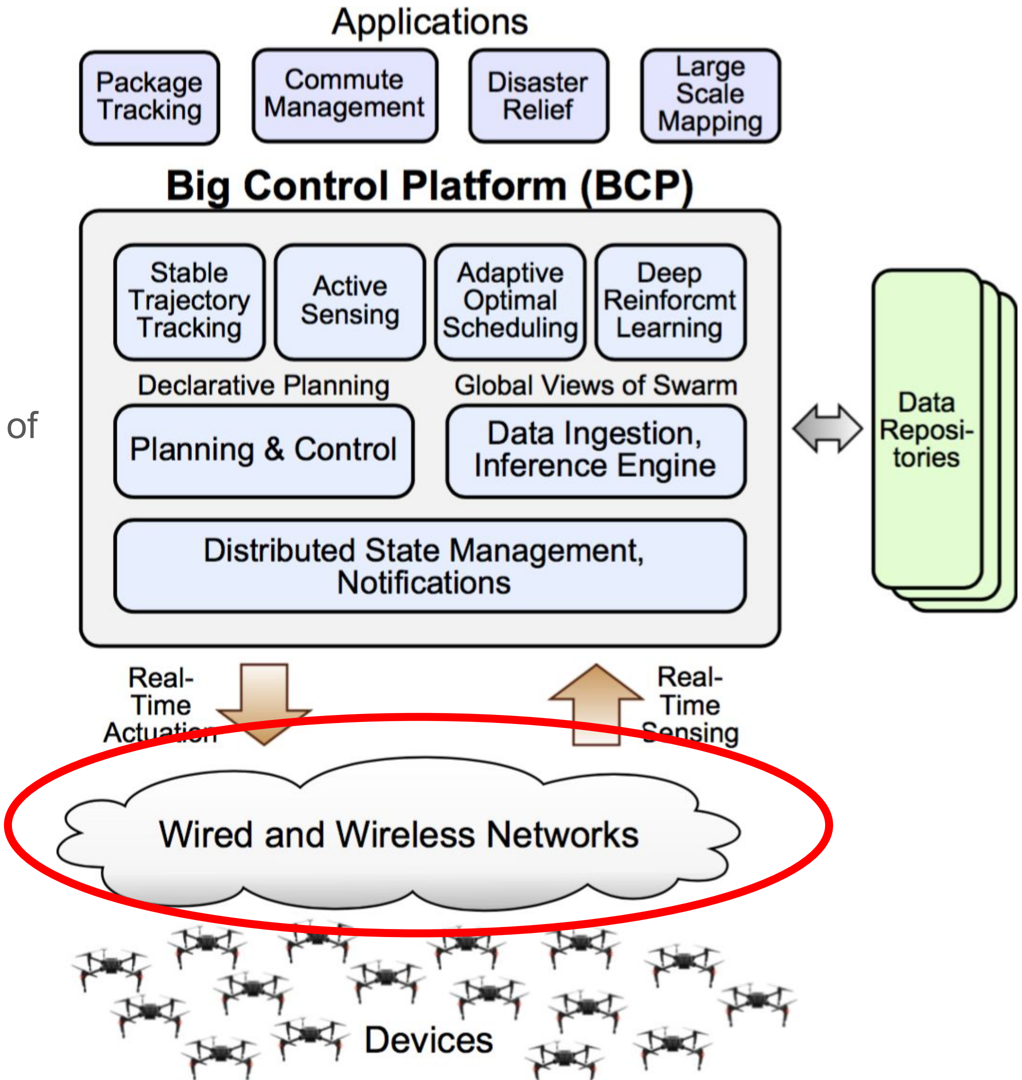

Big Control Platform

Services

Real robotics problems

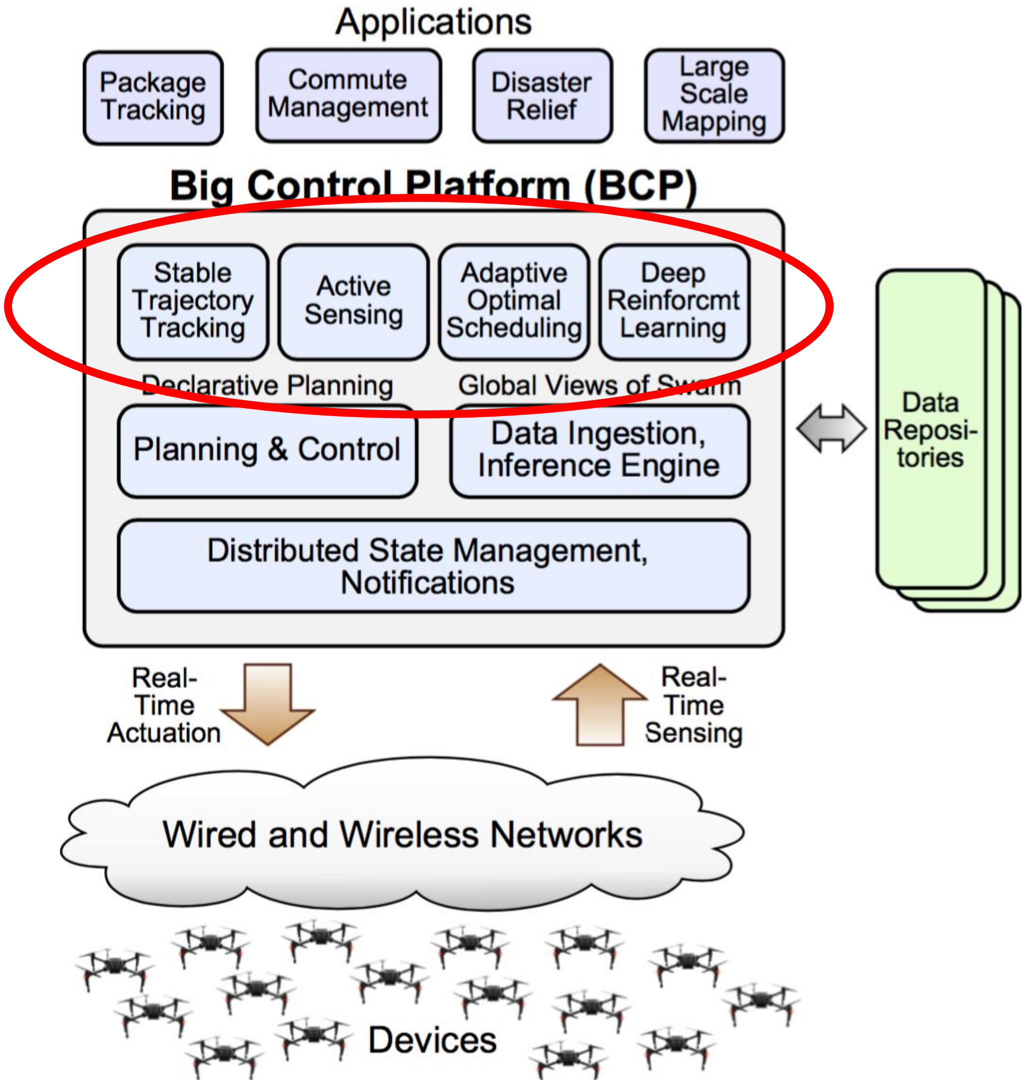

Application Examples

- 3D scanning of cities and campuses
  - Control drones with scanner to cover buildings from all necessary vantage points
- Automated agriculture
  - Drones inspect thousands of individual plants in a field
- Bus fleet management
  - Control self-driving buses based on traffic, passenger demand, weather, etc.

Grand Challenge Disaster Recovery

- After disaster send in trucks that release thousands of drones
- Area of interests
  - Coordinated search
  - Real-time sensing and inference
  - Data integration
  - Mobilizing ground vehicles
  - Small package delivery

Thank You