POMI2020 Mobility

Nick McKeown
nickm@stanford.edu

The Stanford Clean Slate Program
http://cleanslate.stanford.edu

Technology competition?

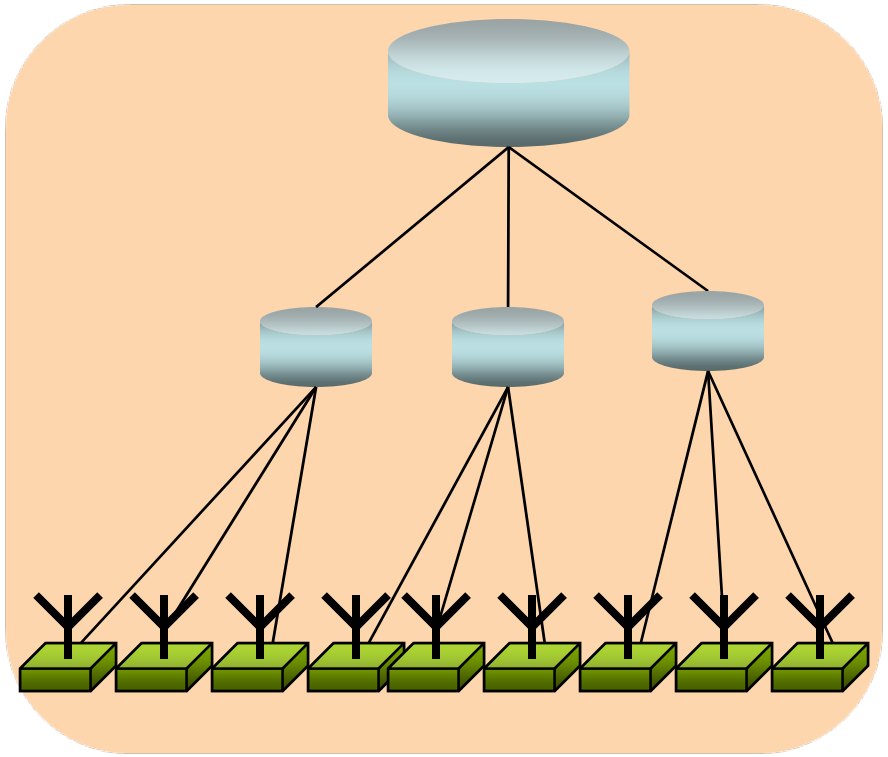

Cellular providers

3G, Wimax, LTE, 3GPP, …

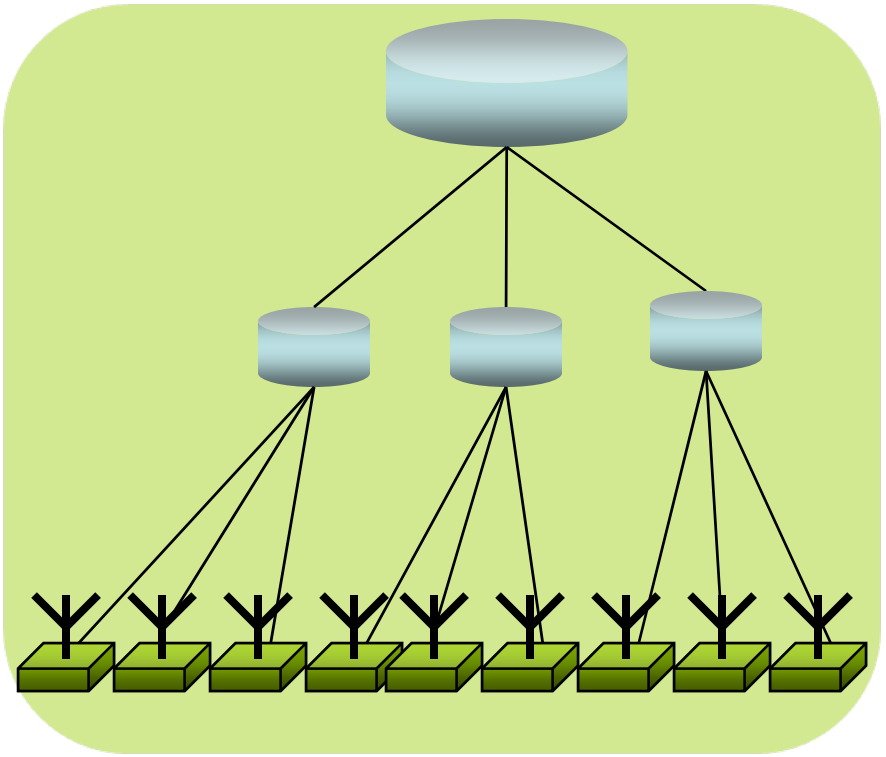

WiFi infrastructure

Employer, city, home, neighbor, …

High investment, desire to keep closed

Intertwined radio/network, specialized network

Low investment

Open/closed?

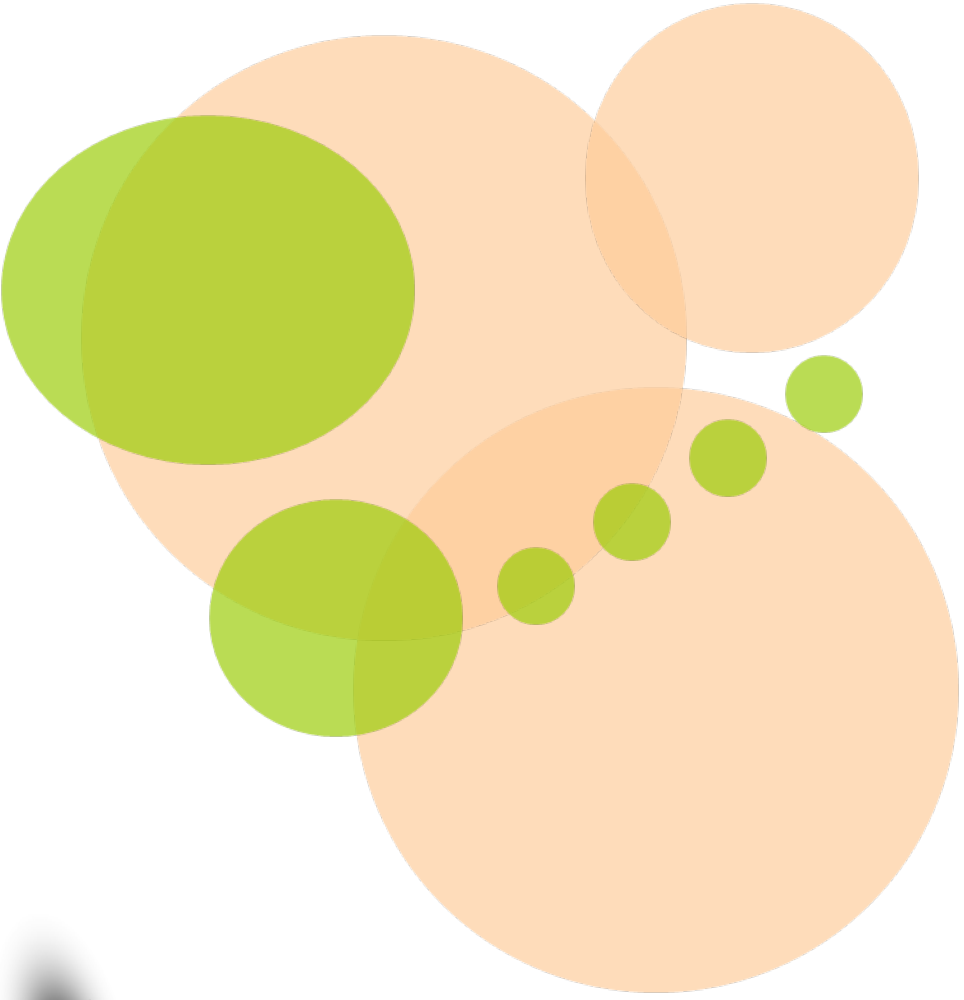

- Expect a rich combination of both.
- Both will evolve.
- More a question of ownership than technology.

What is in best interest of user?

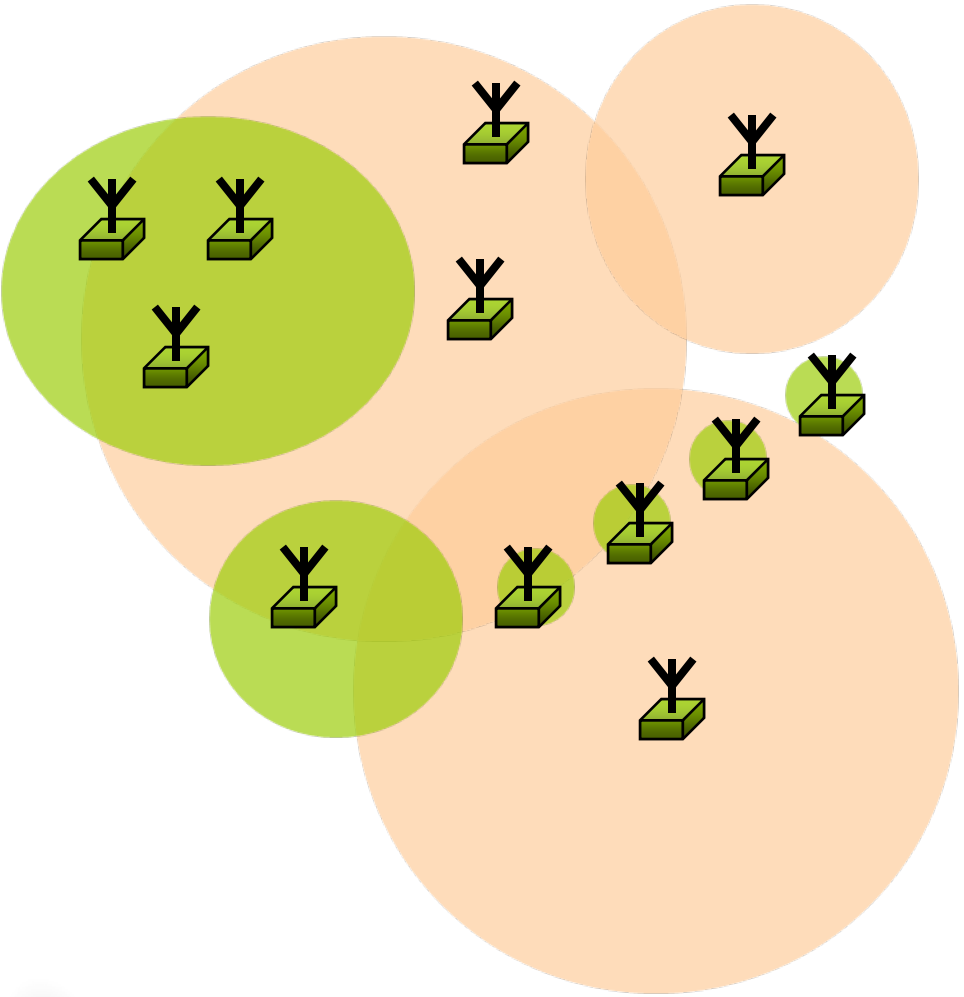

- Today: many cellular networks visible (5-7 common), many wifi networks visible (10-15 common).
- But not practically available to me – closed infrastructures.

How can I use of all the infrastructure around me?

Goal

Maximize choice for the user

Therefore

- Assume rich deployment of radios
- Be radio technology neutral
- Minimize cost of switchover and handover

Problem

- How to help maximize choice

Technical goals

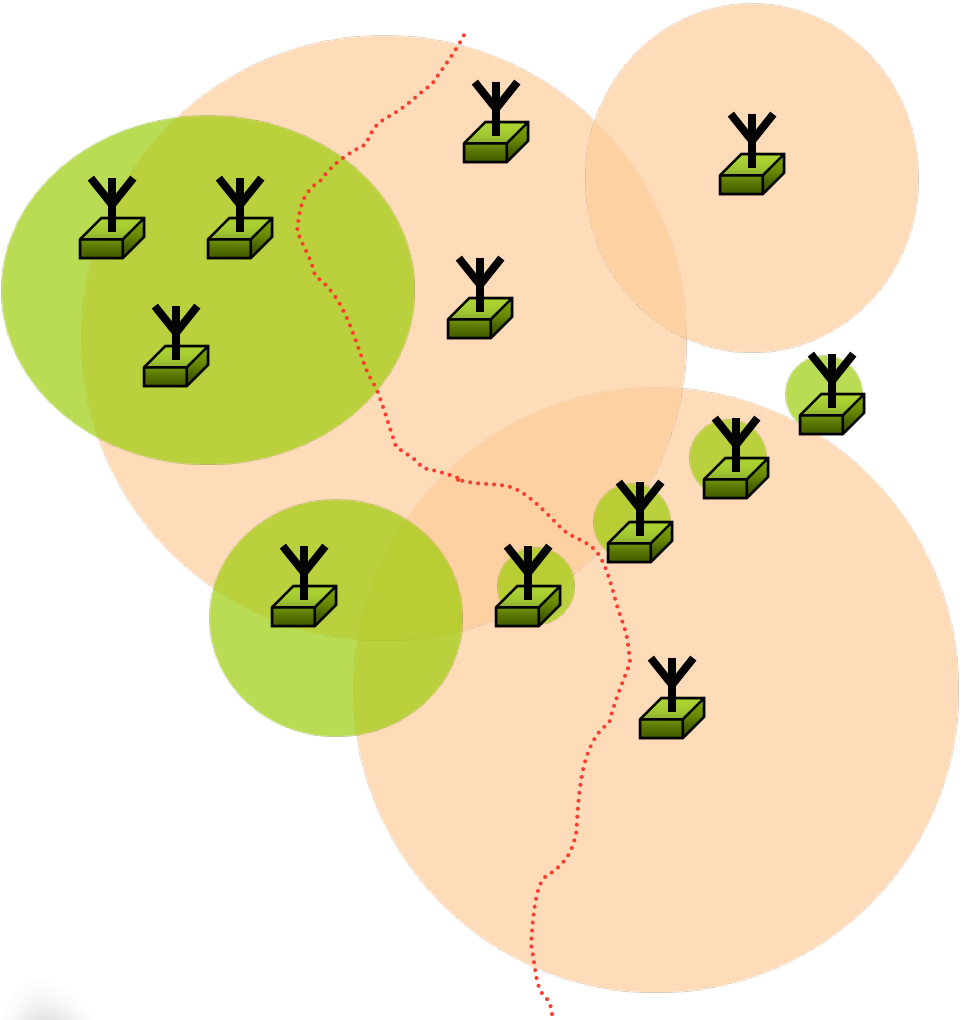

- Access to all infrastructure
- Continued connectivity as I move
- User choice
  - Radio
  - Handoff
- Allow innovation
  - Handoff mechanisms
  - AAA, billing, …

Some separation happening

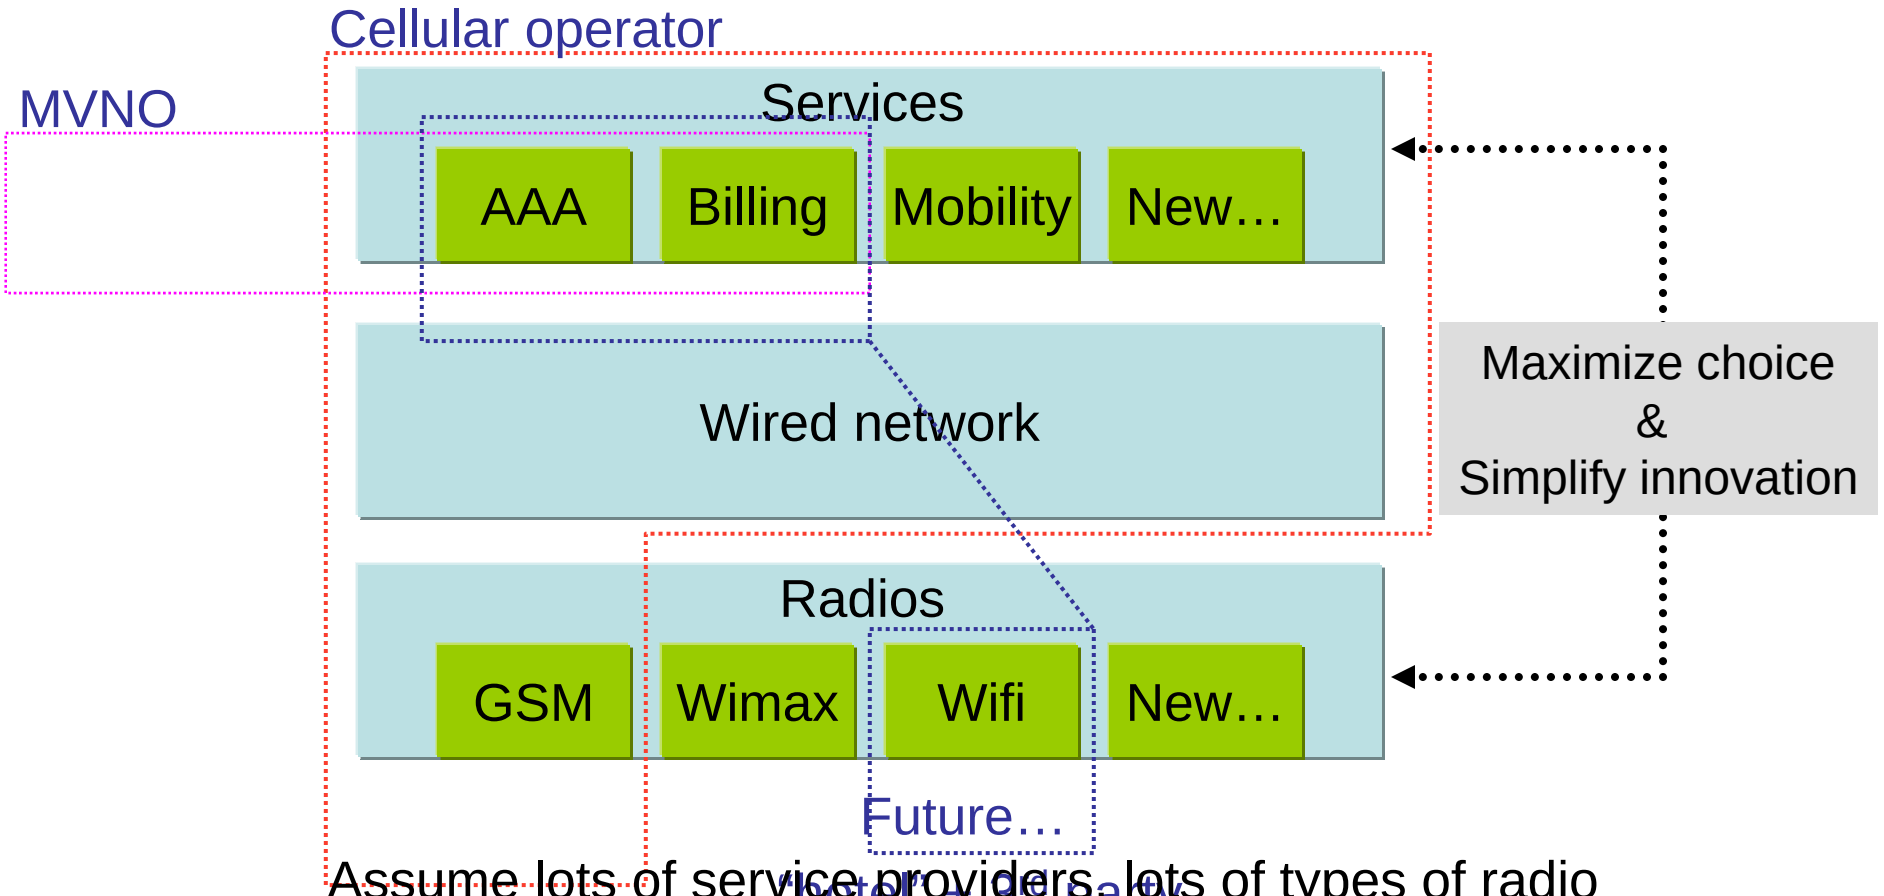

Assume lots of service providers, lots of types of radio

Assume lots of diversity of space, channels, multiple radios, APs, …

Assume always make-before-break

Implications on mobility

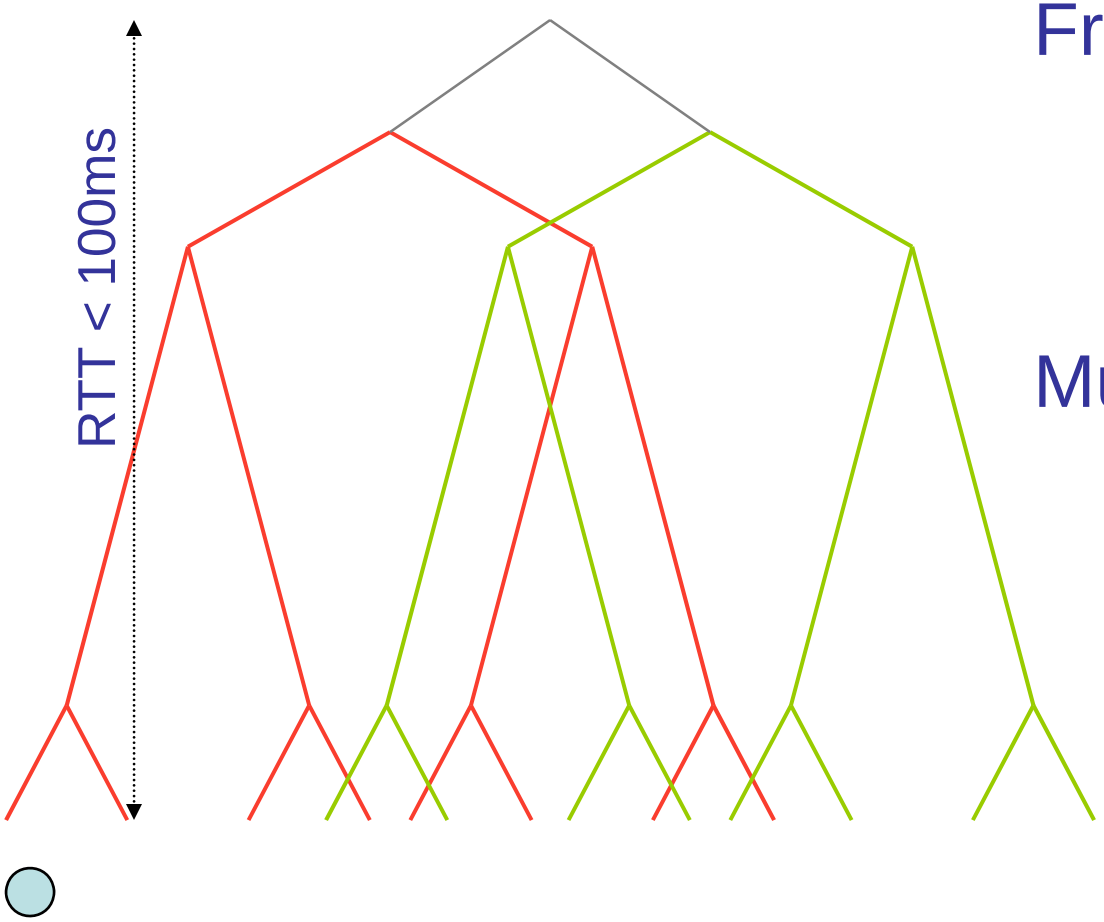

Frequency of handoff

- "Cell" size
- Speed of motion
- Signal degradation

Must finish one handoff before start next.

What does this say about wired network and mobility?

Consequences & observations

1. If frequency of handoff > 1/RTT then we have to decentralize handoff and directories.
2. If frequency < (1/RTT + processing time) then we can choose *if* or *how* decentralized.
3. If frequency ~= (1/RTT + processing time) then probably *need* network support.
4. If frequency << 1/10s, can propagate routes using IP

   Leads to big tables, but Moore is on our side

Simple model for handoff frequency

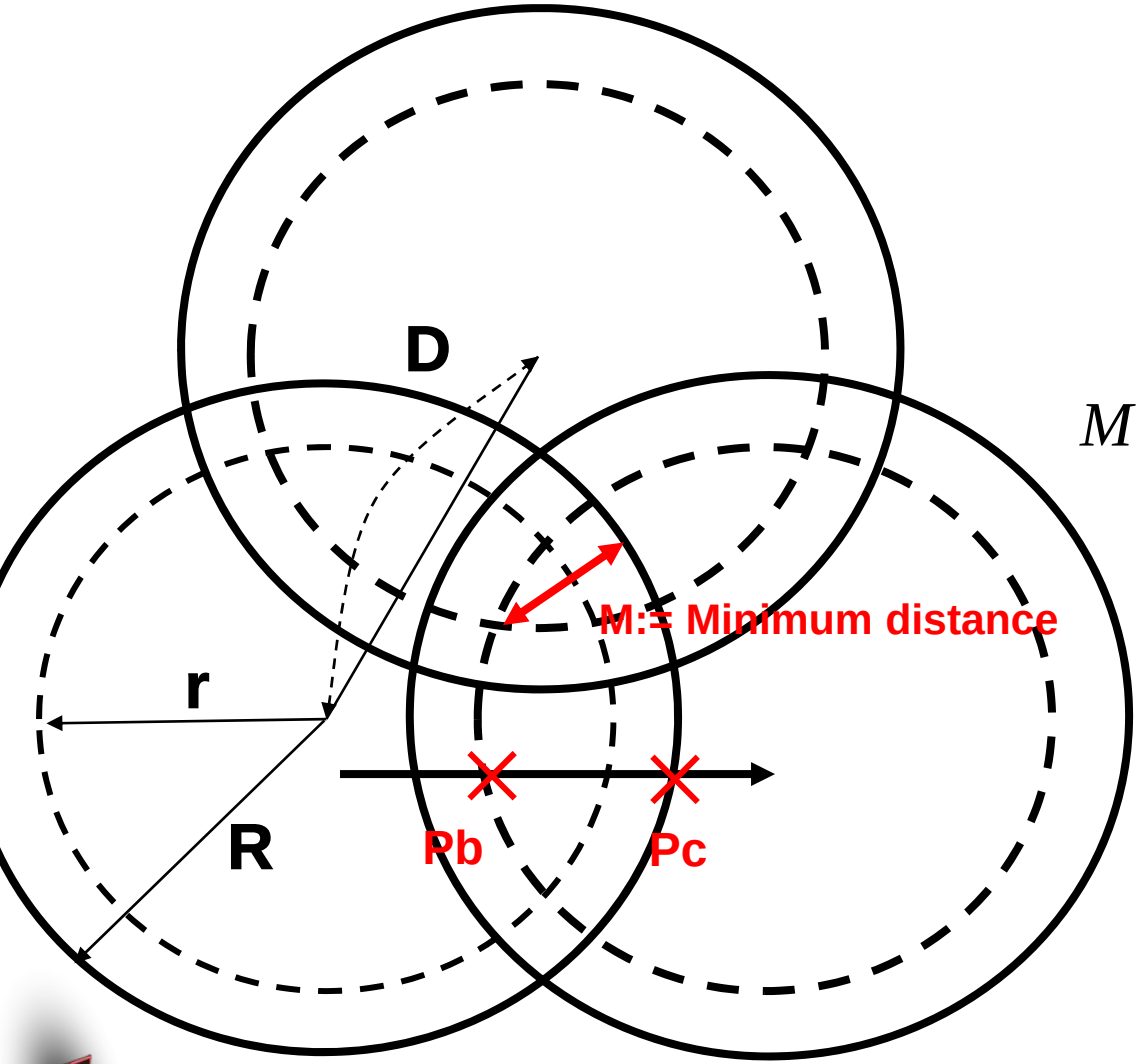

Handover time requirement
= the minimum time to
move from Pb to Pc
= (min dist. btwn Pb and Pc)/v
= M/v

$$M = \sqrt{r^2 - D^2/4} + R - (\sqrt{3}/2)D$$

[Masayoshi Kobayashi]

Handoff requirement

- Couple movement model with a variety of wireless propagation models…
- … es
- Initial

Back of the envelope

e.g. 100km/h in 100m cell → cross cell in 3s
Frequency of handoff per mobile O(1 per 1s)

  - Hard to envisage frequency > 1/100ms
- Fits well with existing standards

Conclusions on handoff frequency

- We probably do not *require* decentralization.
- (Does not mean decentralization is a bad idea).

- Lots of choice of implementation.
- Perhaps eases innovation and evolution.

Scoping the amount of information

- Directory of devices/users/location
  - Total directory O(10bn)
  - Update rate:
    - Depends on where in hierarchy
    - Back of the envelope…
      - Assume 1% of all users moving at a time and global directory event needed every 100s per user
      - $O(10^6)$ updates per second
      - Assume 10^3 bytes/update → about 10Gb/s (total)

Mobility in Networks

- Cellular network
  - O(1bn) phones
  - Multiple standards
  - Complex
  - Works
- IP
  - MobileIP (and 10^3 variants)
  - Variants of overlays and redirection
  - Slow, not scaleable … hokey
- Common
  - Mechanisms tie network, routing and policy together
  - All are closed:
    - Cellular network by design
    - IP because routing is owned by infrastructure
  - Rate of innovation is slow

Our Goal in POMI 2020 Project

1. Create an open platform/substrate suitable for innovation in mobility
2. Put into the hands of innovators
3. Stand back and watch

Example Experiment: Mobility

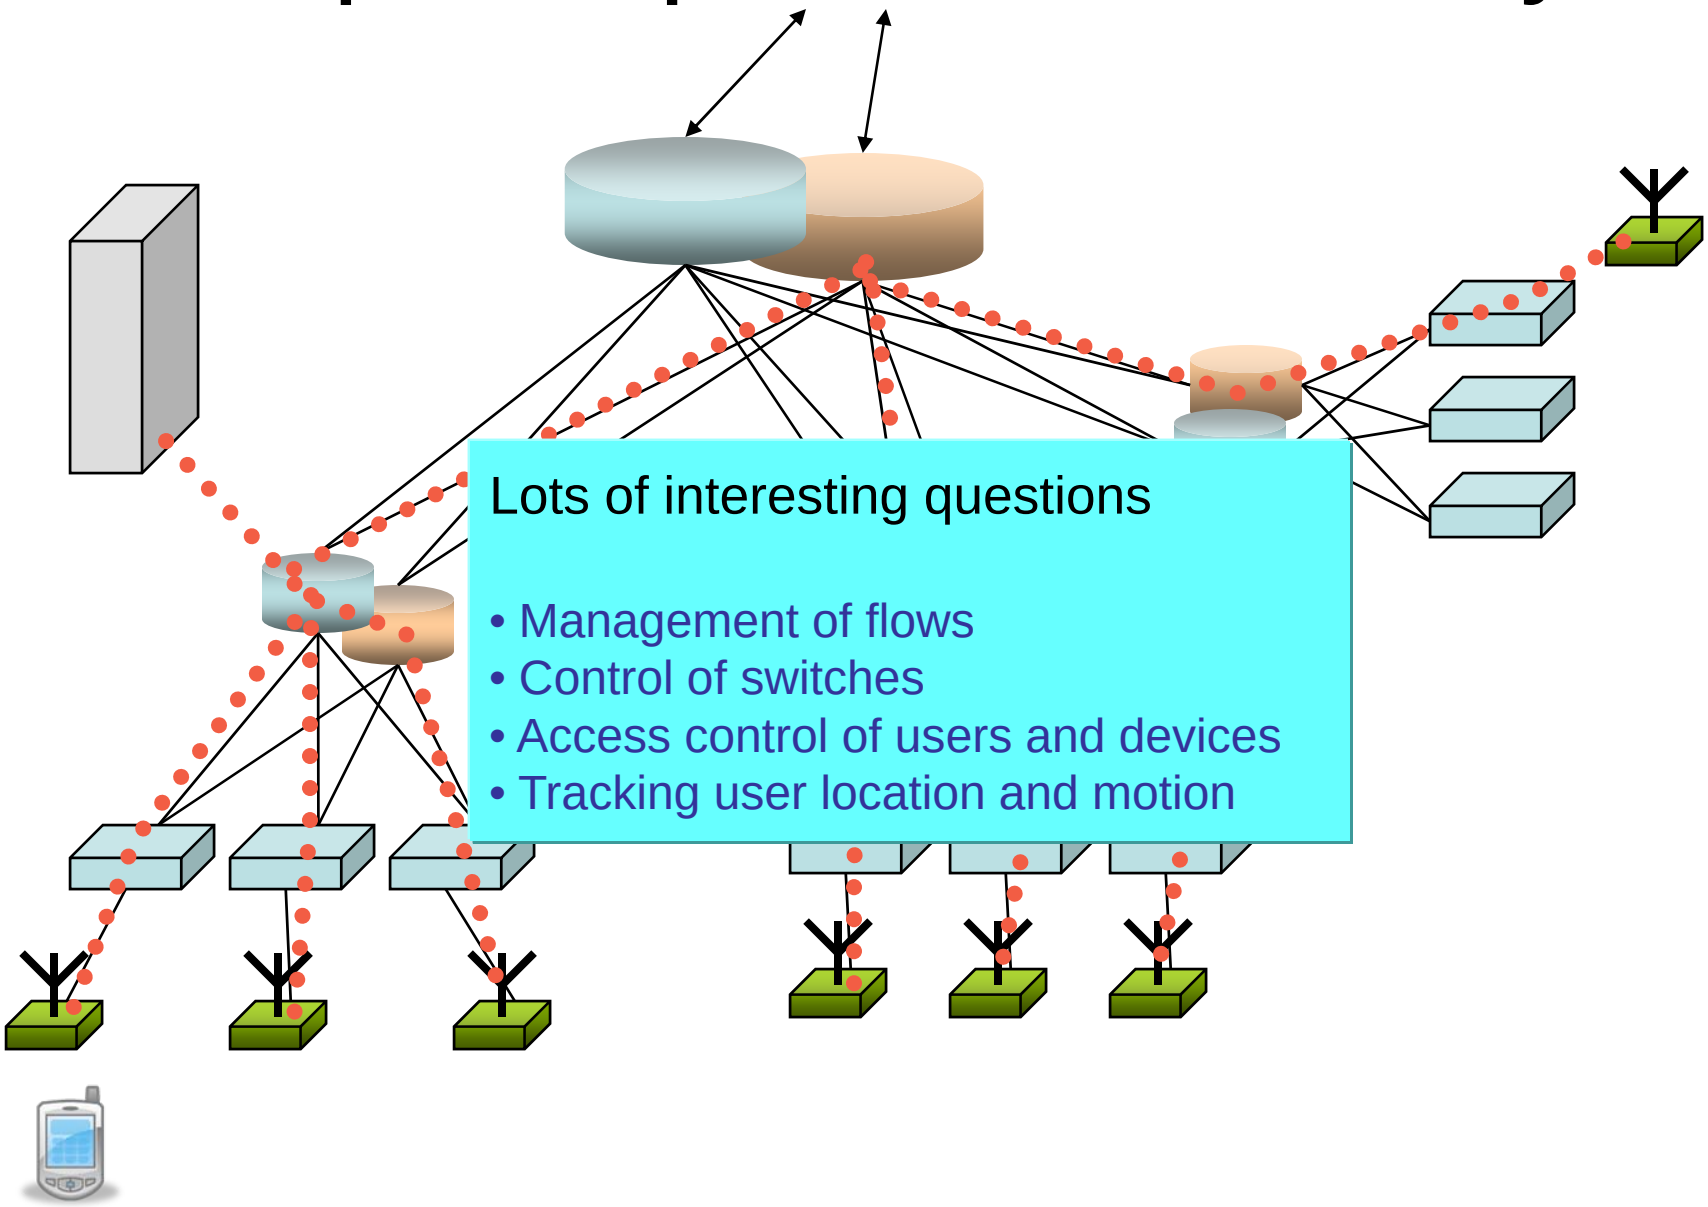

Needs

- Compatible with IP at end host, but infrastructure/routing not compelled to use IP addresses
- Possible to innovate: routing, handoff mechanism, directory service, security and access control, …
- Allow
  - Distributed or centralized control
  - Network-controlled or handset-controlled
  - Calling-plan based, free or advertising-based

OpenFlow Switching

A way to run experiments in the networks we use everyday.

A "pragmatic" compromise

Allow researchers to run experiments in their network…

…without requiring vendors to expose internal workings.

Basics

An Ethernet switch (e.g. 128-ports of 1GE)

An open protocol to remotely add/remove flow entries

Experimenter's Dream
(Vendor's Nightmare)

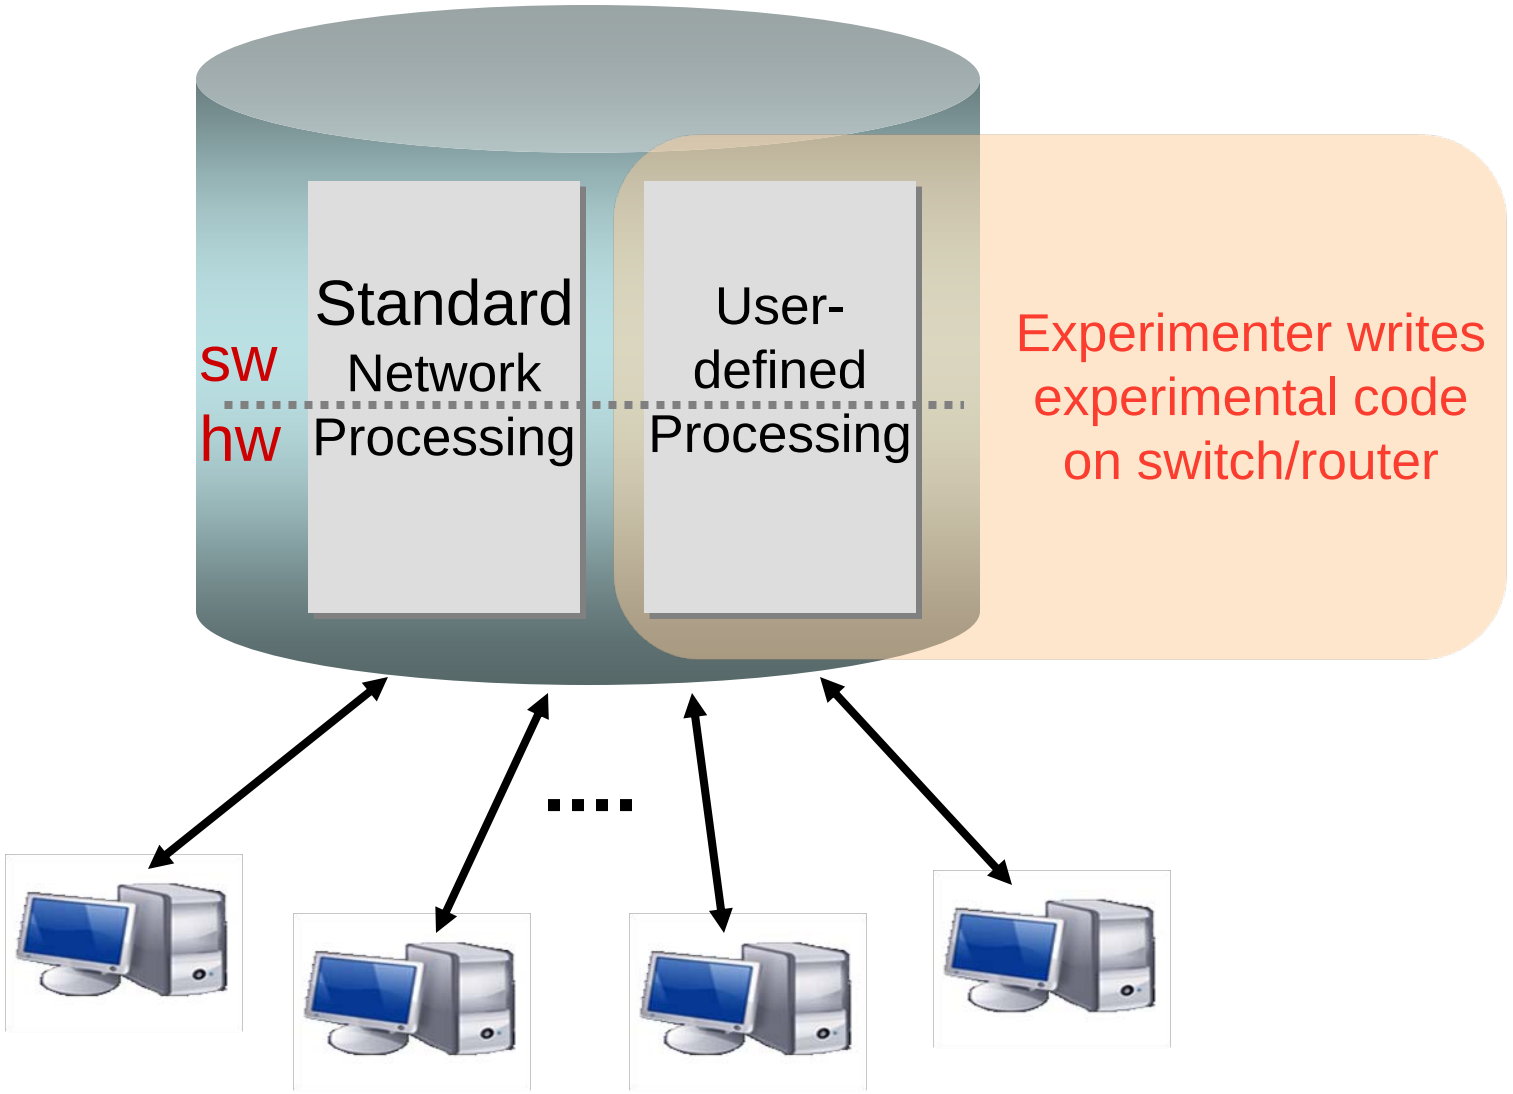

No obvious way

Commercial vendor not ready to open software and hardware development environment

- Complexity of support
- Market protection and barrier to entry

Hard to build my own

- Prototypes are flakey
- Software only: Too slow
- Hardware/software: Fanout too small (need >100 ports for wiring closet)

OpenFlow Switching

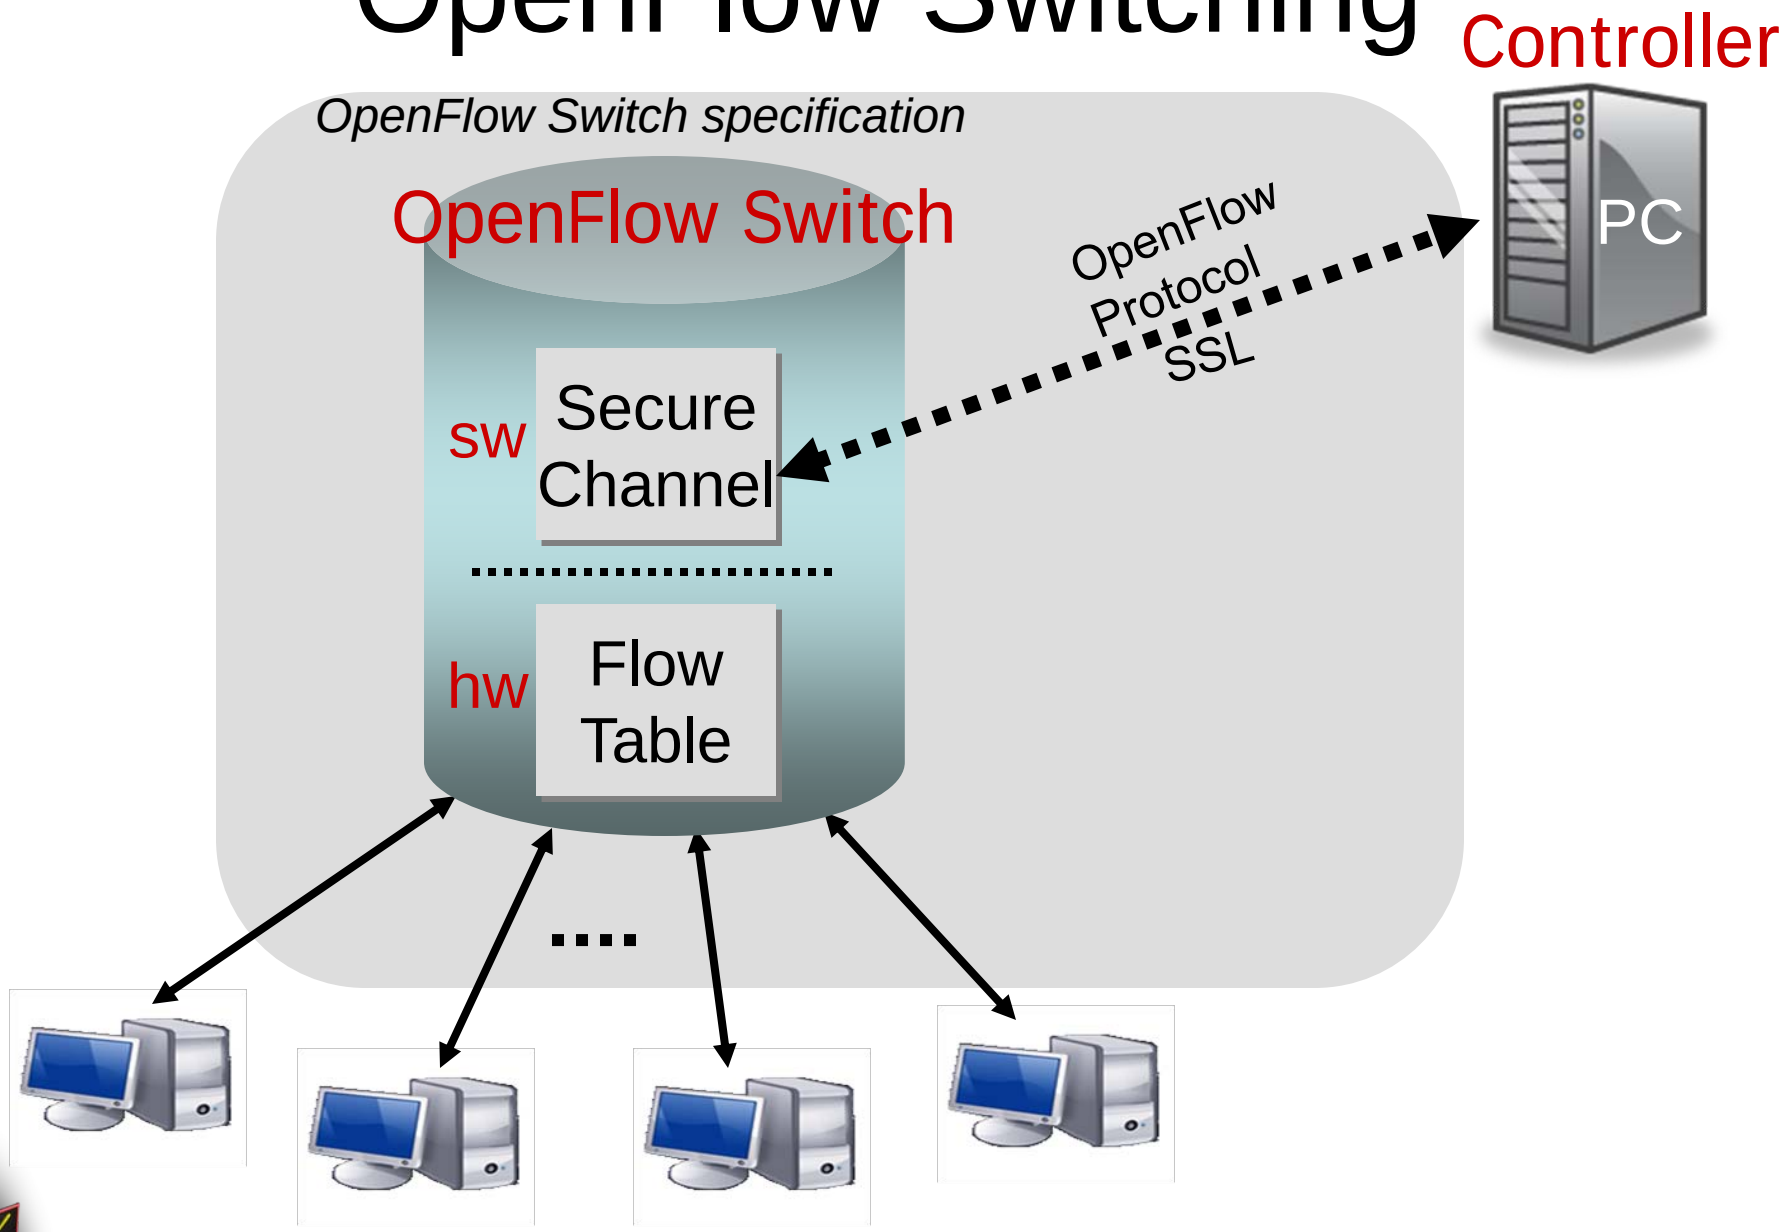

Flow Table Entry

"Type 0" OpenFlow Switch

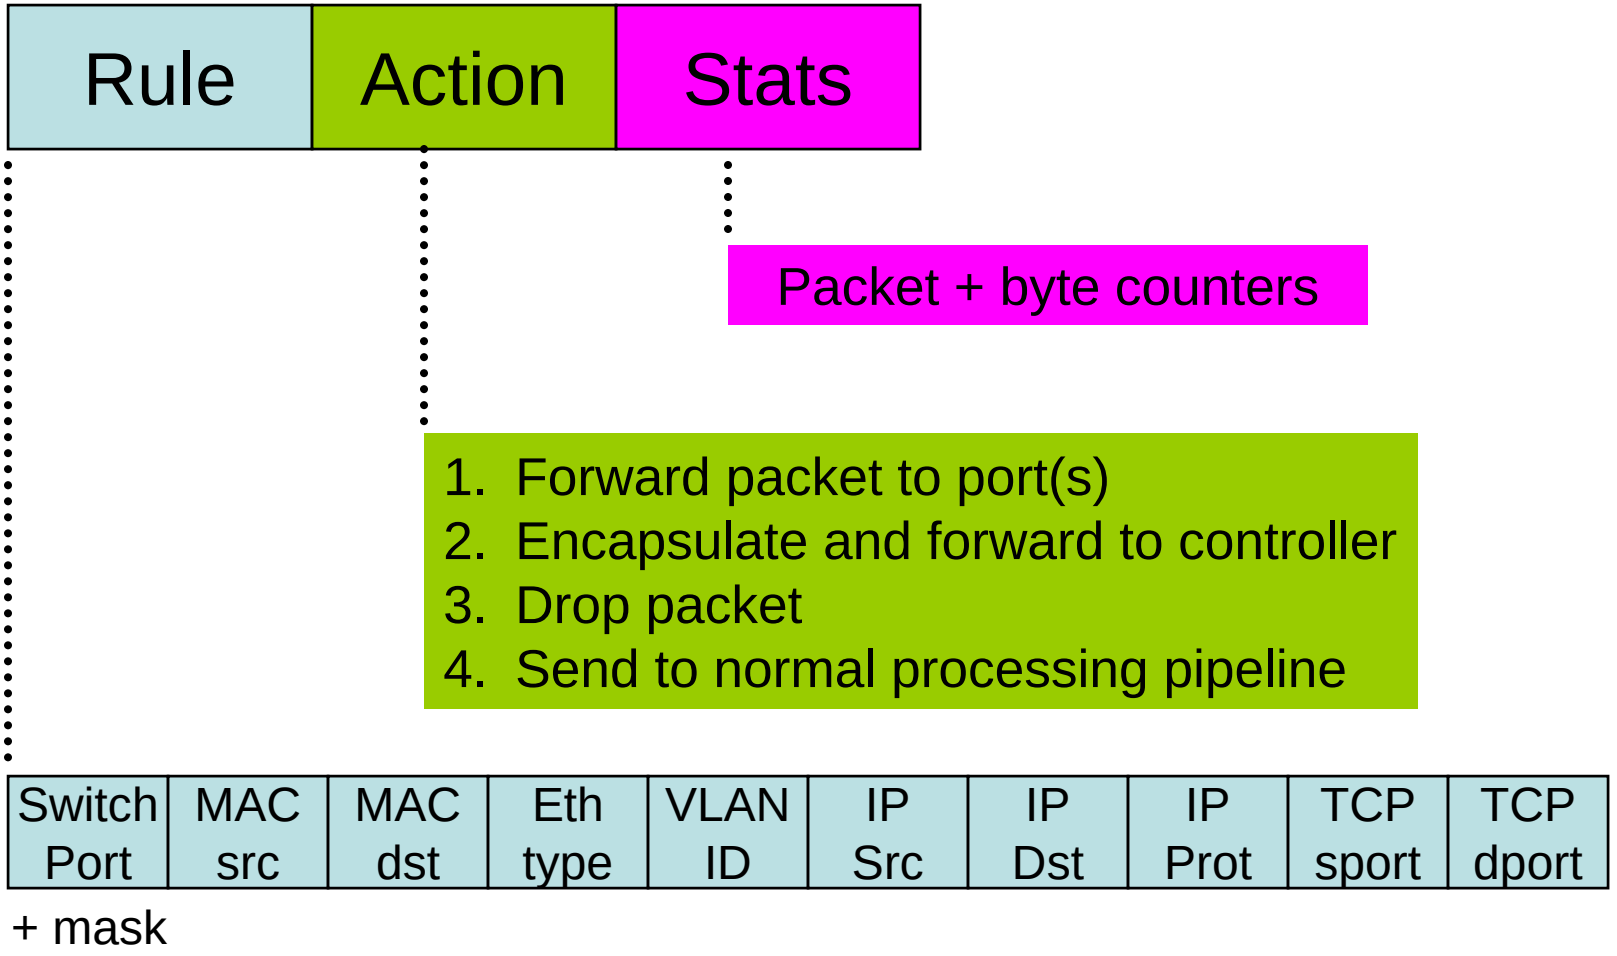

OpenFlow “Type 1”

- Definition in progress
- Additional actions
  - Rewrite headers
  - Map to queue/class
  - Encrypt
- More flexible header
  - Allow arbitrary matching of header bytes
- Support multiple controllers
  - Load-balancing and reliability

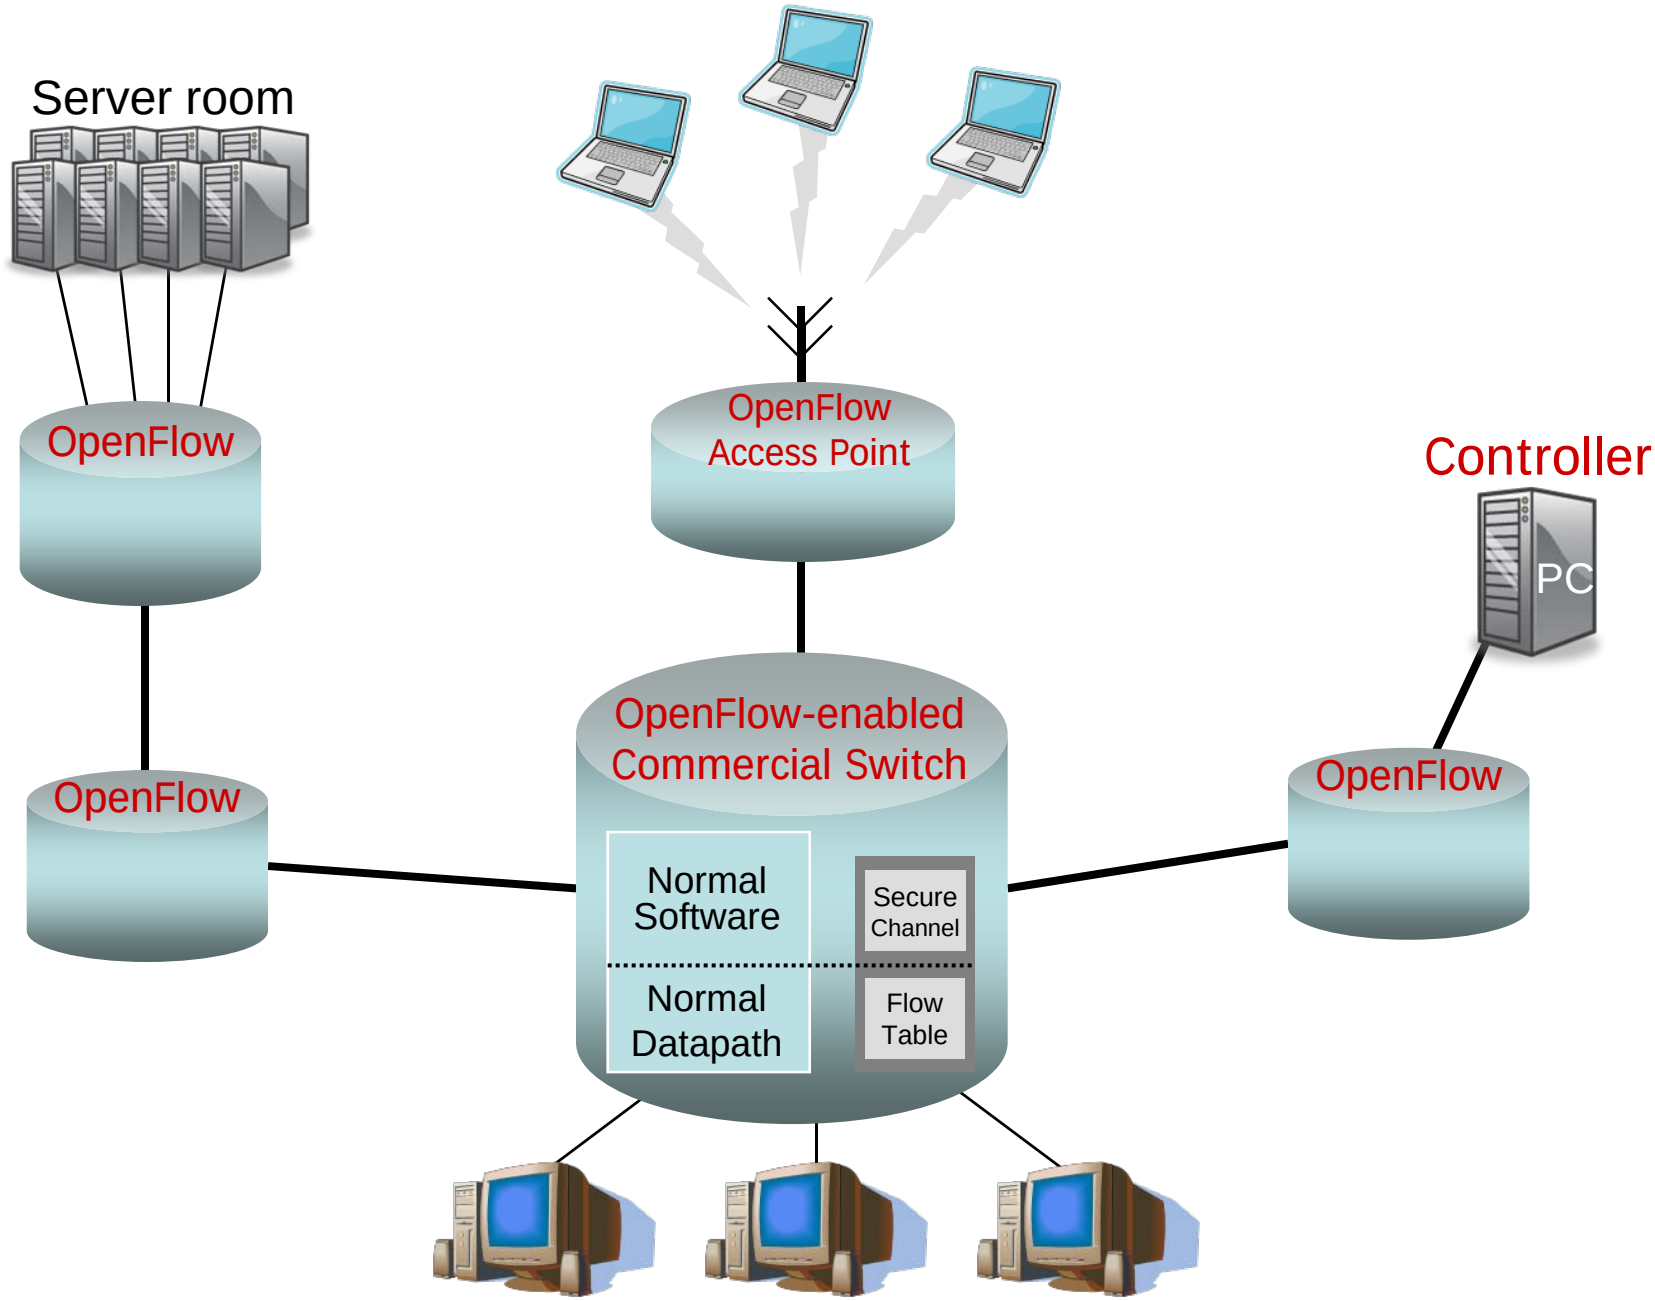
Server room
OpenFlow
OpenFlow
OpenFlow
Access Point
Controller
PC
OpenFlow-enabled
Commercial Switch
Normal
Software
Normal
Datapath
Secure
Channel
Flow
Table
OpenFlow

OpenFlow Consortium

http://OpenFlowSwitch.org

Goal: Evangelize OpenFlow to vendors

Free membership for all researchers

Whitepaper, OpenFlow Switch Specification, Reference Designs

Licensing: Free for research and commercial use

OpenFlow: Status

Commercial Ethernet switches and routers

- Working with several vendors to add to existing products
- Expect OpenFlow “Type 0” to be available in 2008-09

Reference switches

- Software: Linux and OpenWRT (for access points)
- Hardware: NetFPGA (line-rate 1GE; available soon)
- Working on low-cost 48-port 1GE switch based on Broadcom reference design

Reference controllers

- Simple test controller
- NOX controller

Deployment at Stanford

Stanford Computer Science Department

Gates Building

~1,000 network users

23 wiring closets

Stanford Center for Integrated Systems (EE)

Paul Allen Building

~200 network users

6 wiring closets

Working with HP Labs and Cisco on deployment

Experimental infrastructure

- Our goal is to deploy an OpenFlow network on campus…
- …interconnect different radio technologies.
  - WiFi and Wimax
- To enable experiments with mobility and policy mechanisms in our network.
- To understand innovation at scale.
- Then stand back and watch…