Edge computing in the extreme and its applications

Suman Banerjee
CS, UW-Madison
suman@cs.wisc.edu

www.paradrop.org

Cloud computing

Centralized computing resources

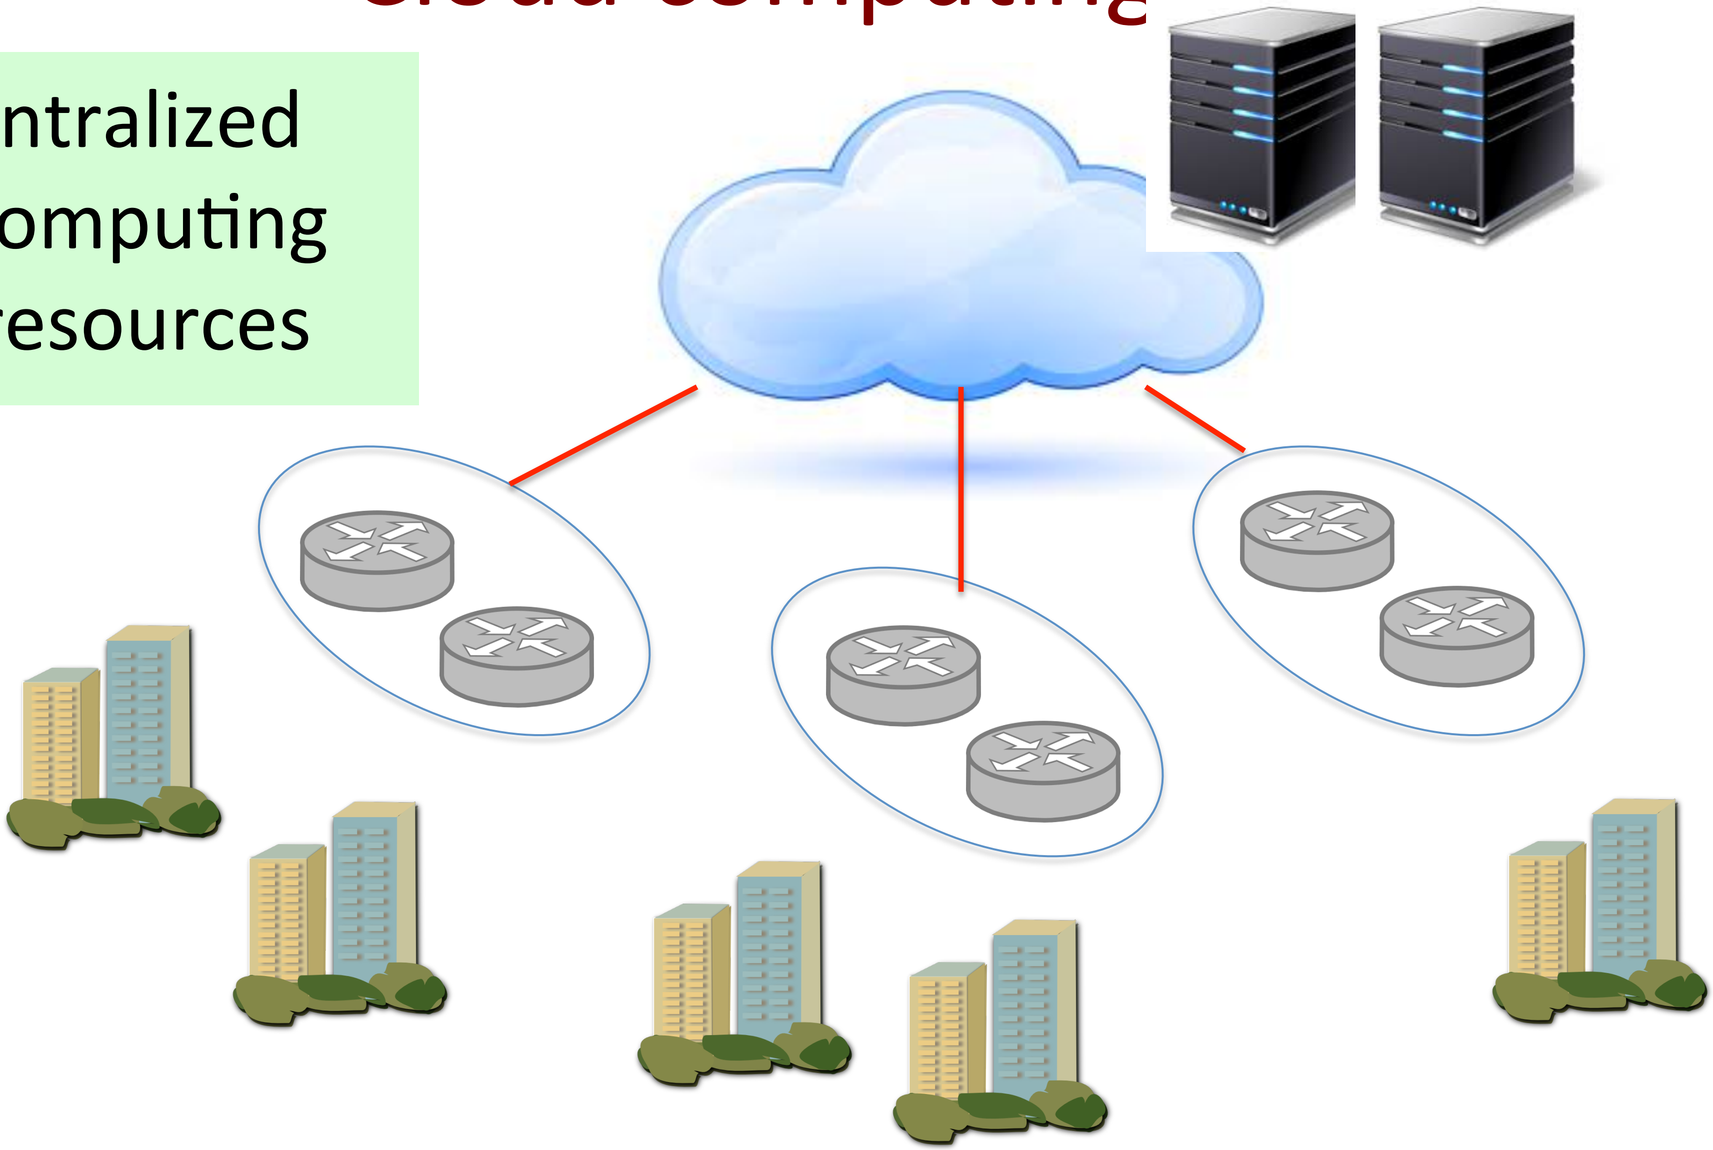

Edge computing

Computing resources in edge networks

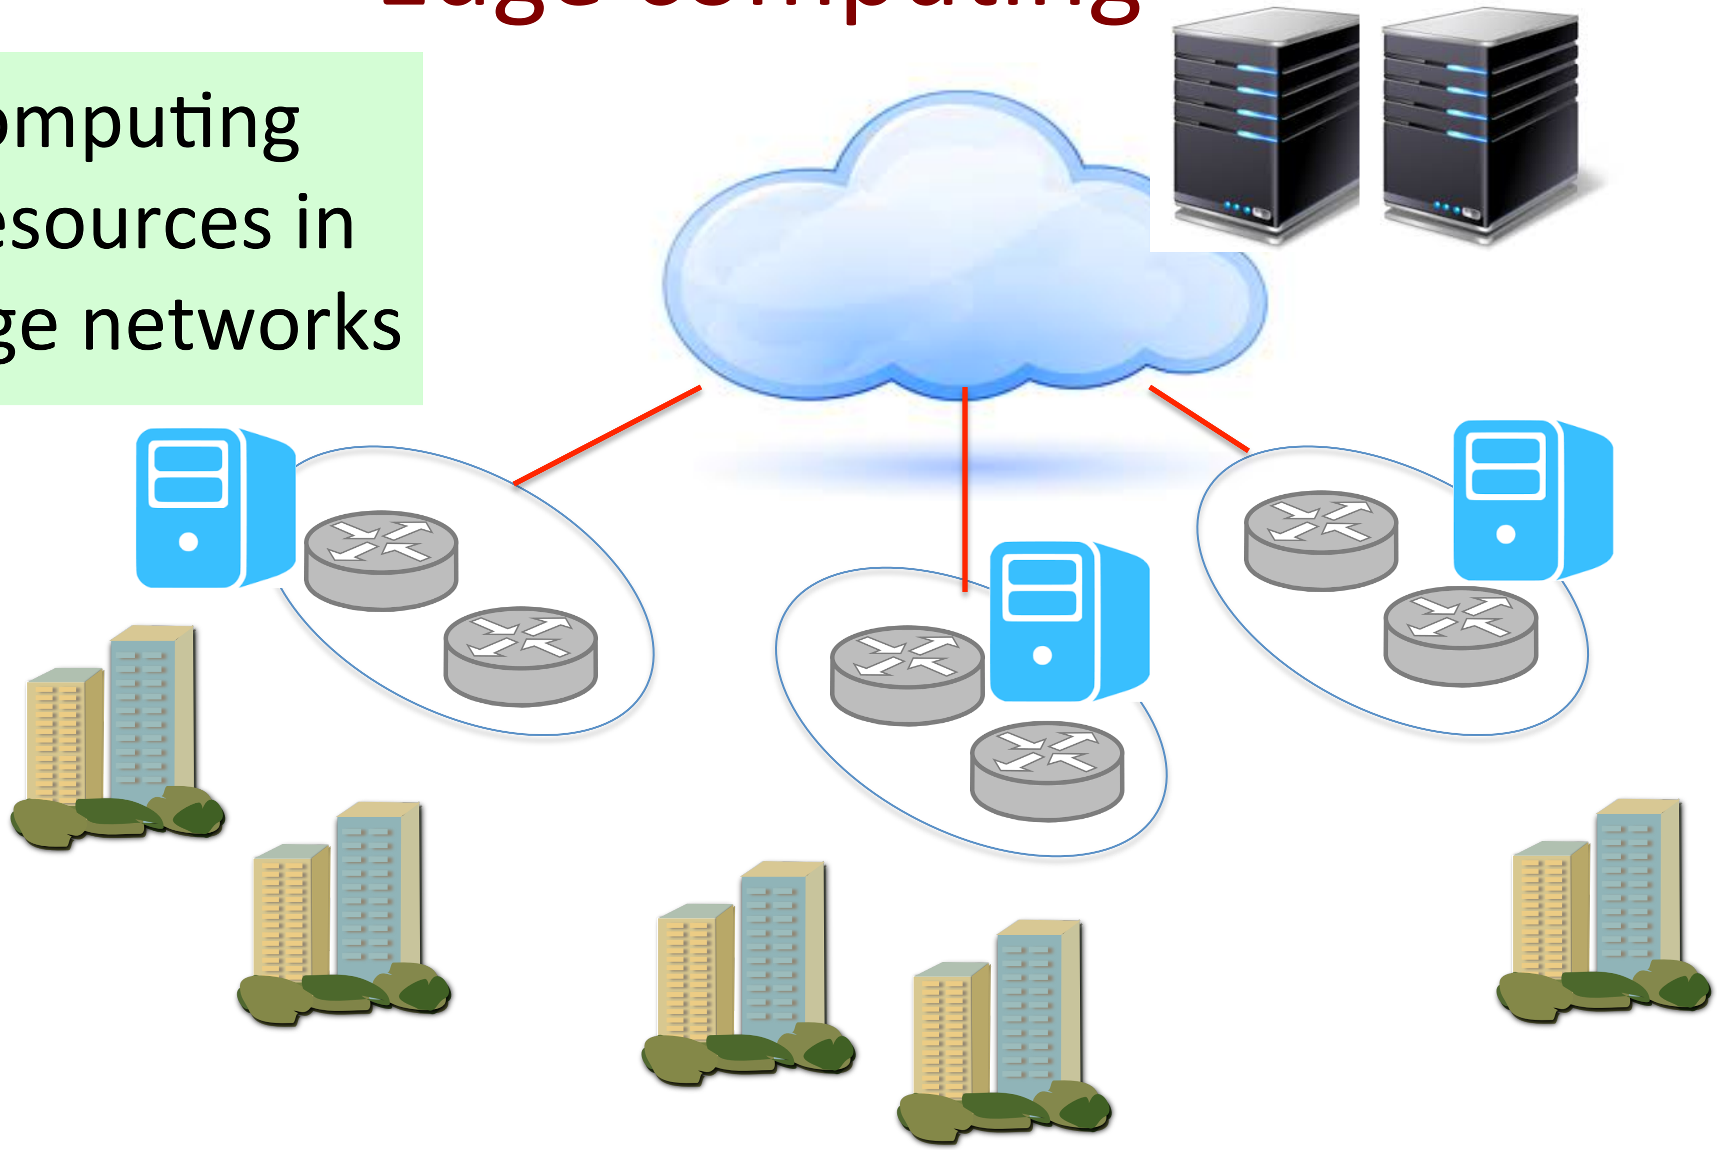

Edge or Fog

Edge: pushing computing applications, data, and services away from centralized nodes to the logical extremes of a network

Proposed years back

- Condor [Litzkow, Livny, et. al.] 1988
- Cyber foraging [Balan et. al.] 2002
- Cloudlets [Satyanarayanan et. al.] 2009
- Fog computing [Bonomi et. al.] 2012

What services?

All kinds

- Store (cache)
  - Netflix, Hulu

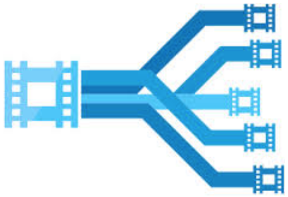

- Compute (transcode, act):
  - Live streaming, games

Talk outline

- **Edge computing in the extreme and apps**
  - ParaDrop design
    - Virtualization, RF management
  - Application 1: Sustainability
  - Application 2: Drive monitoring
- Learning through Deployments

Edge Computing in the Extreme

Edge computing in the extreme

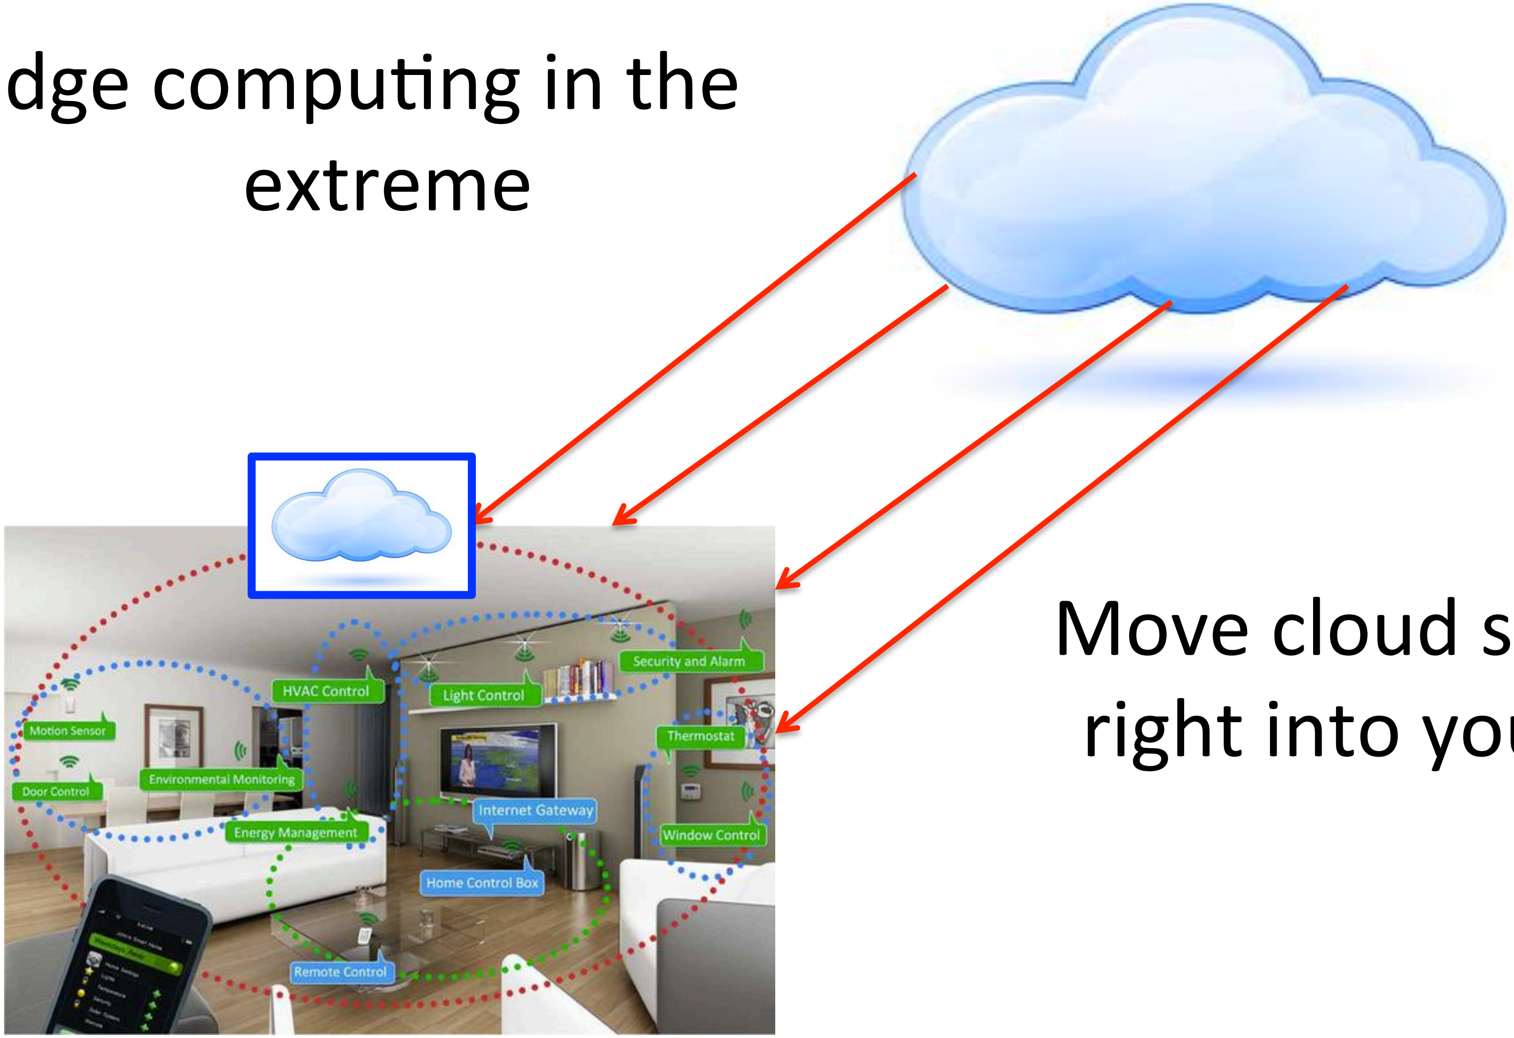

Move cloud services right into your home

Edge Computing in the Extreme

Edge computing in the extreme

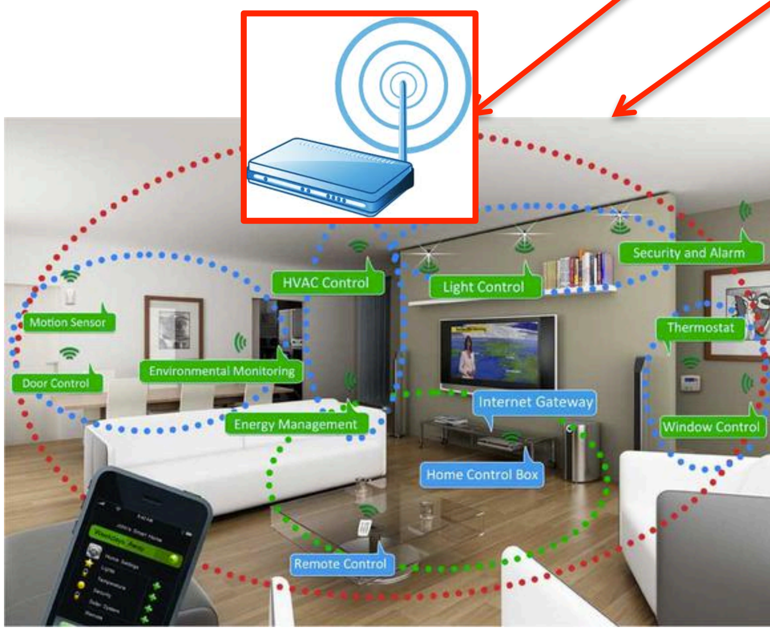

A WiFi router is always-on, low latency to home gadgets, and in the path for all home Internet traffic

To enable new smart home apps

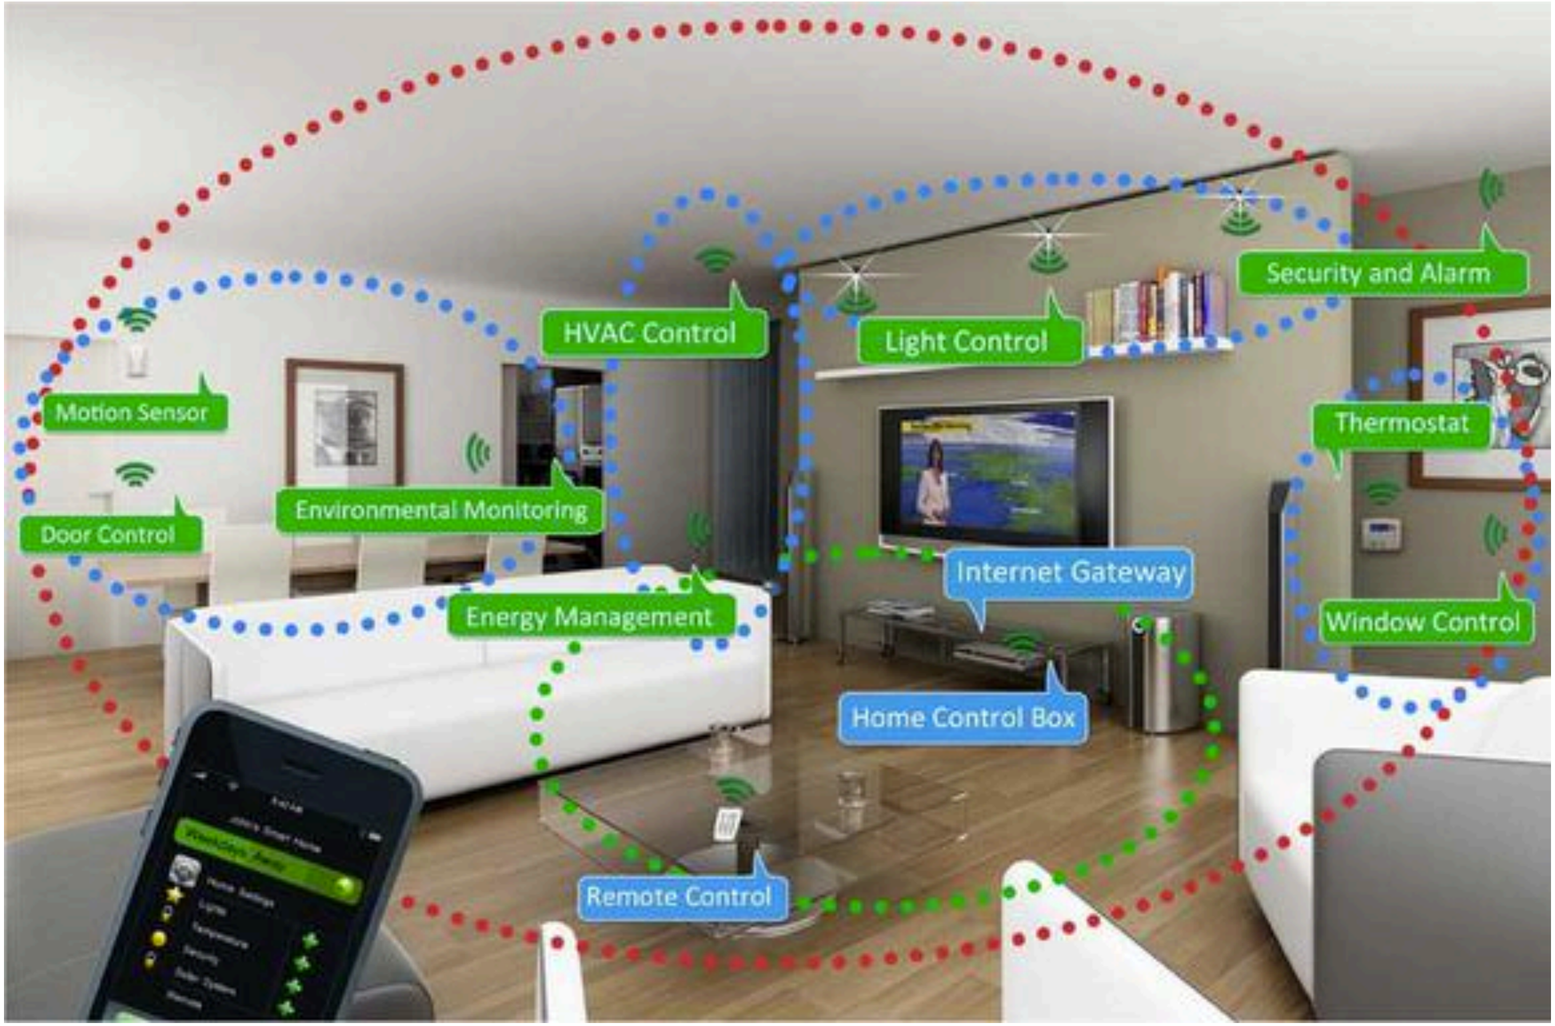

Smart home apps

Manage heating and cooling system locally

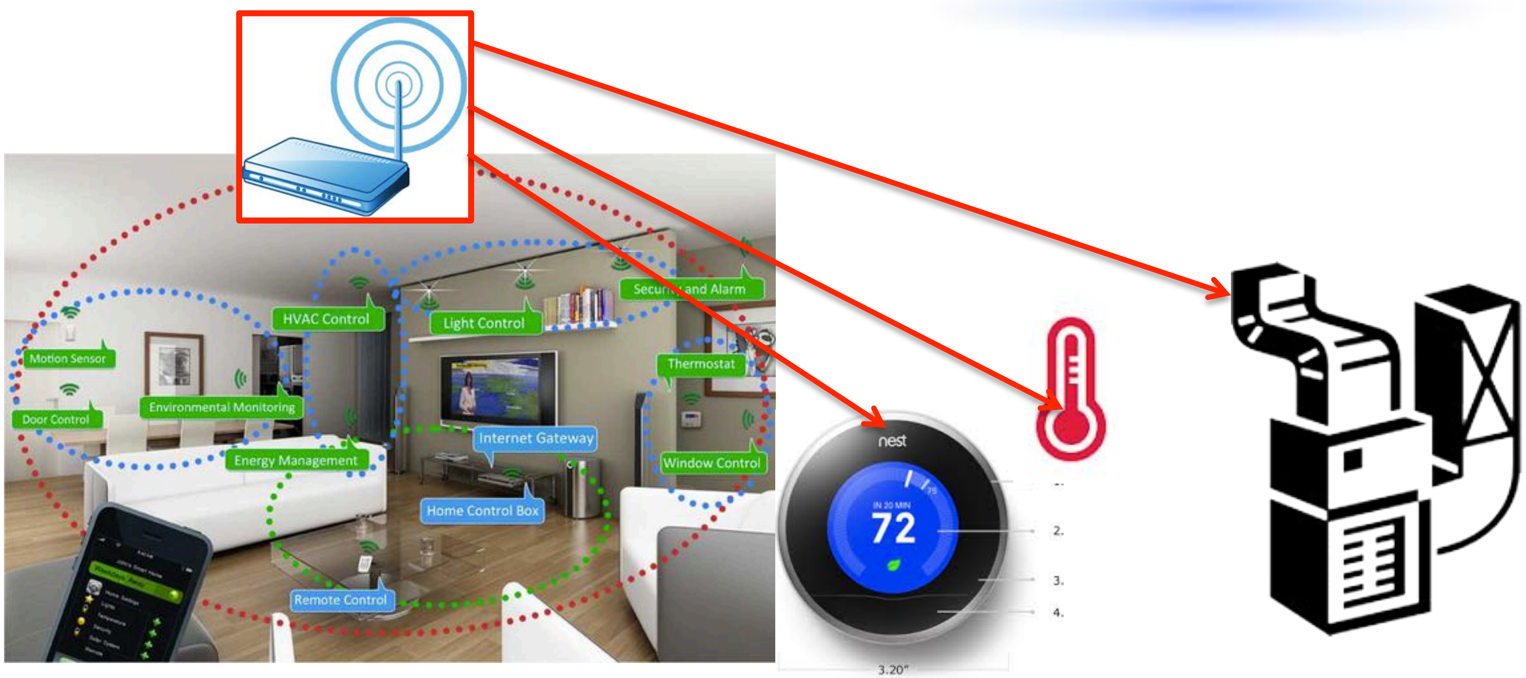

Smart home apps

Manage videos locally …

Smart home apps

Manage videos locally maybe, use it to control home environment

Paradrop

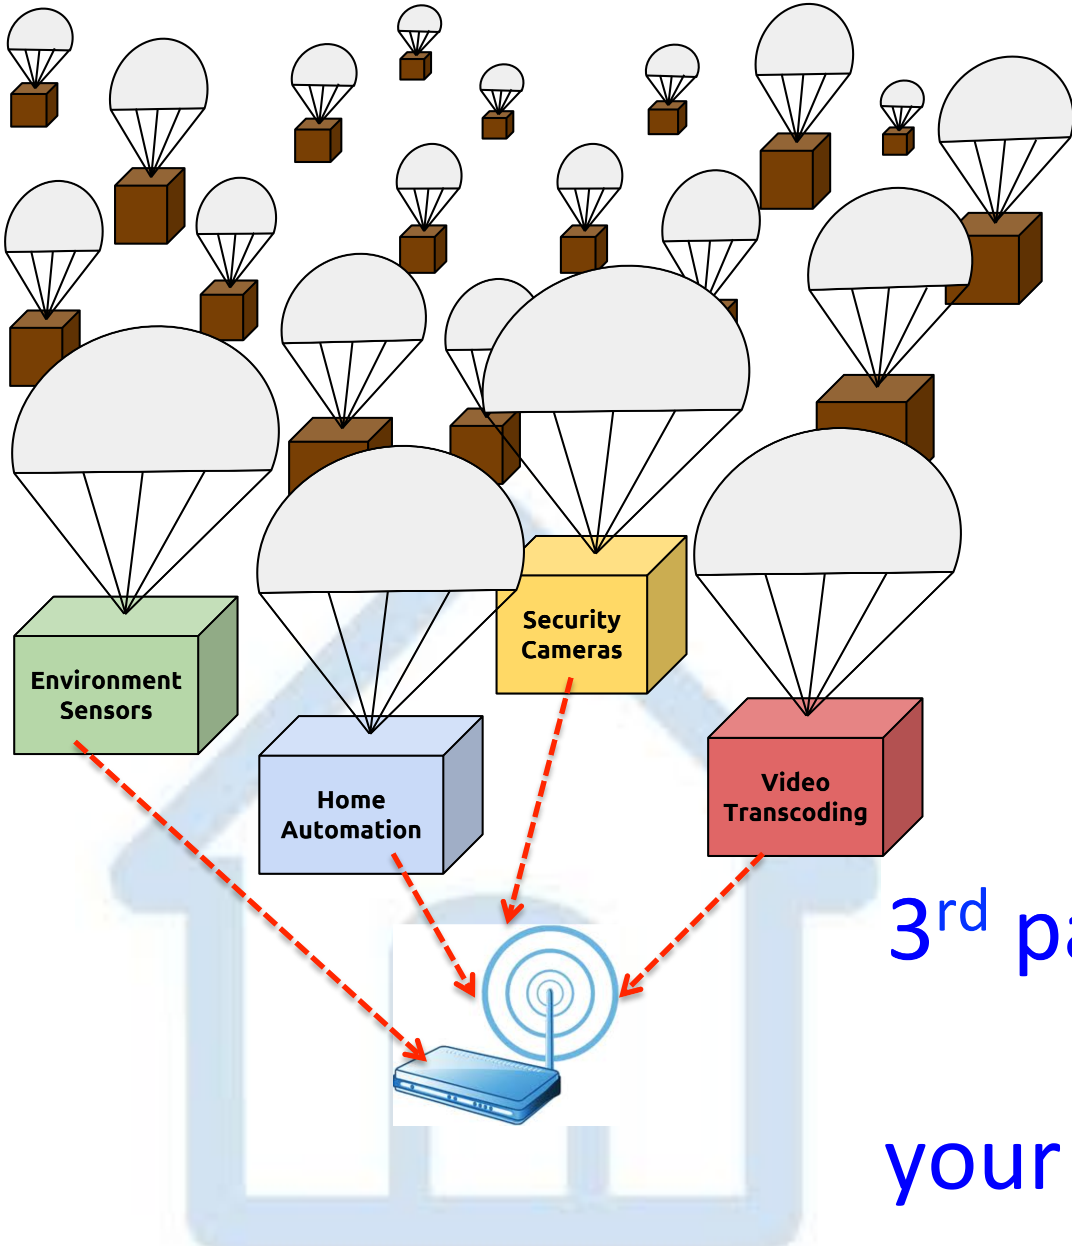

3rd party apps/services drop into your home WiFi router on-demand

Edge computing and ParaDrop

A platform that allows us to deploy new Intelligent services in WiFi routers through virtualization

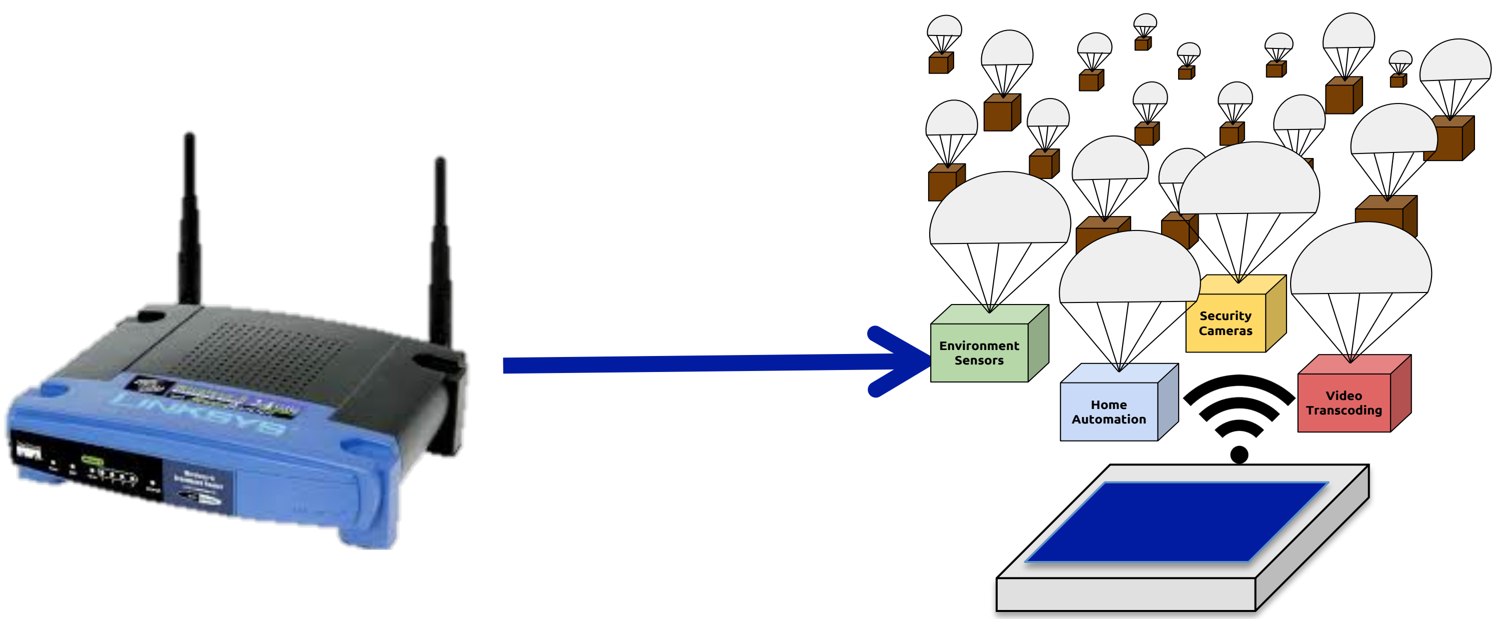

ParaDrop

- A programmable substrate
- Virtualization framework
  - “Chutes”
  - Isolated and proprietary

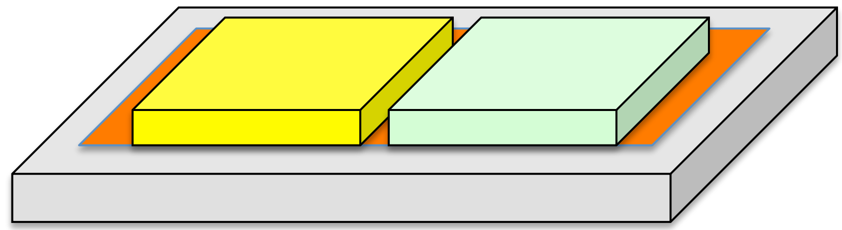

- Multiple wireless interfaces
  - WiFi, Bluetooth, ZigBee
- Wireless interference management and a context API

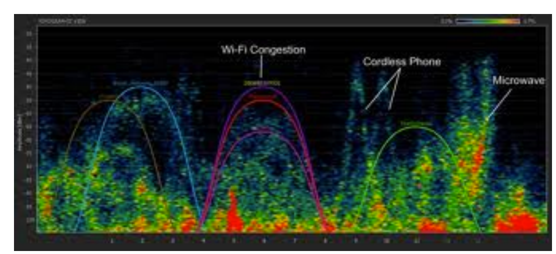

Virtualization in ParaDrop

- Uses container technology
  - LXC

OS level containers

  - Docker

App level containers

How it works?

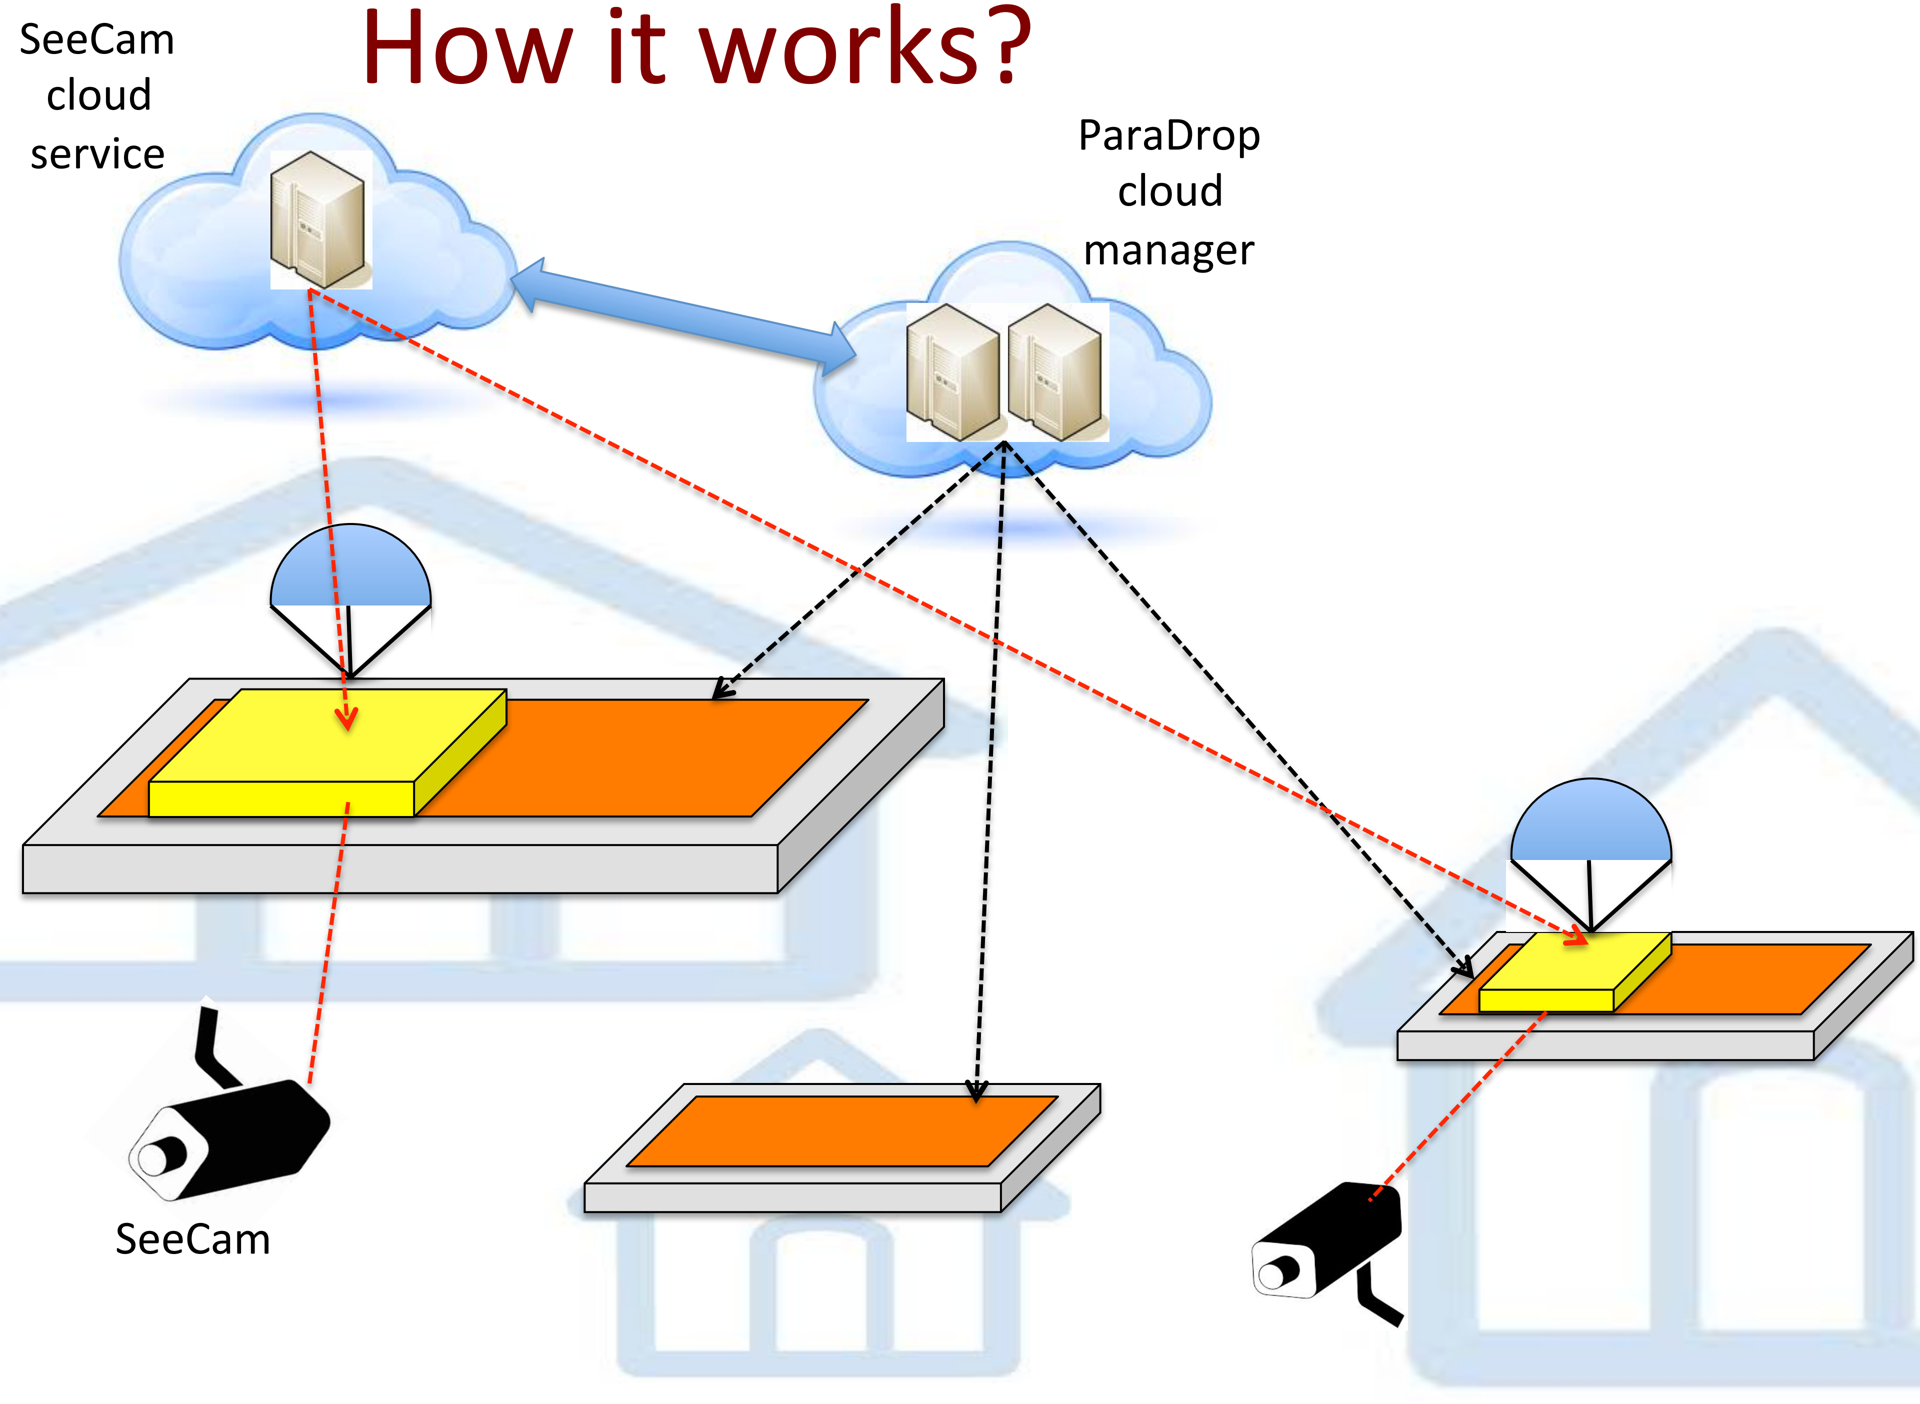

How it works?

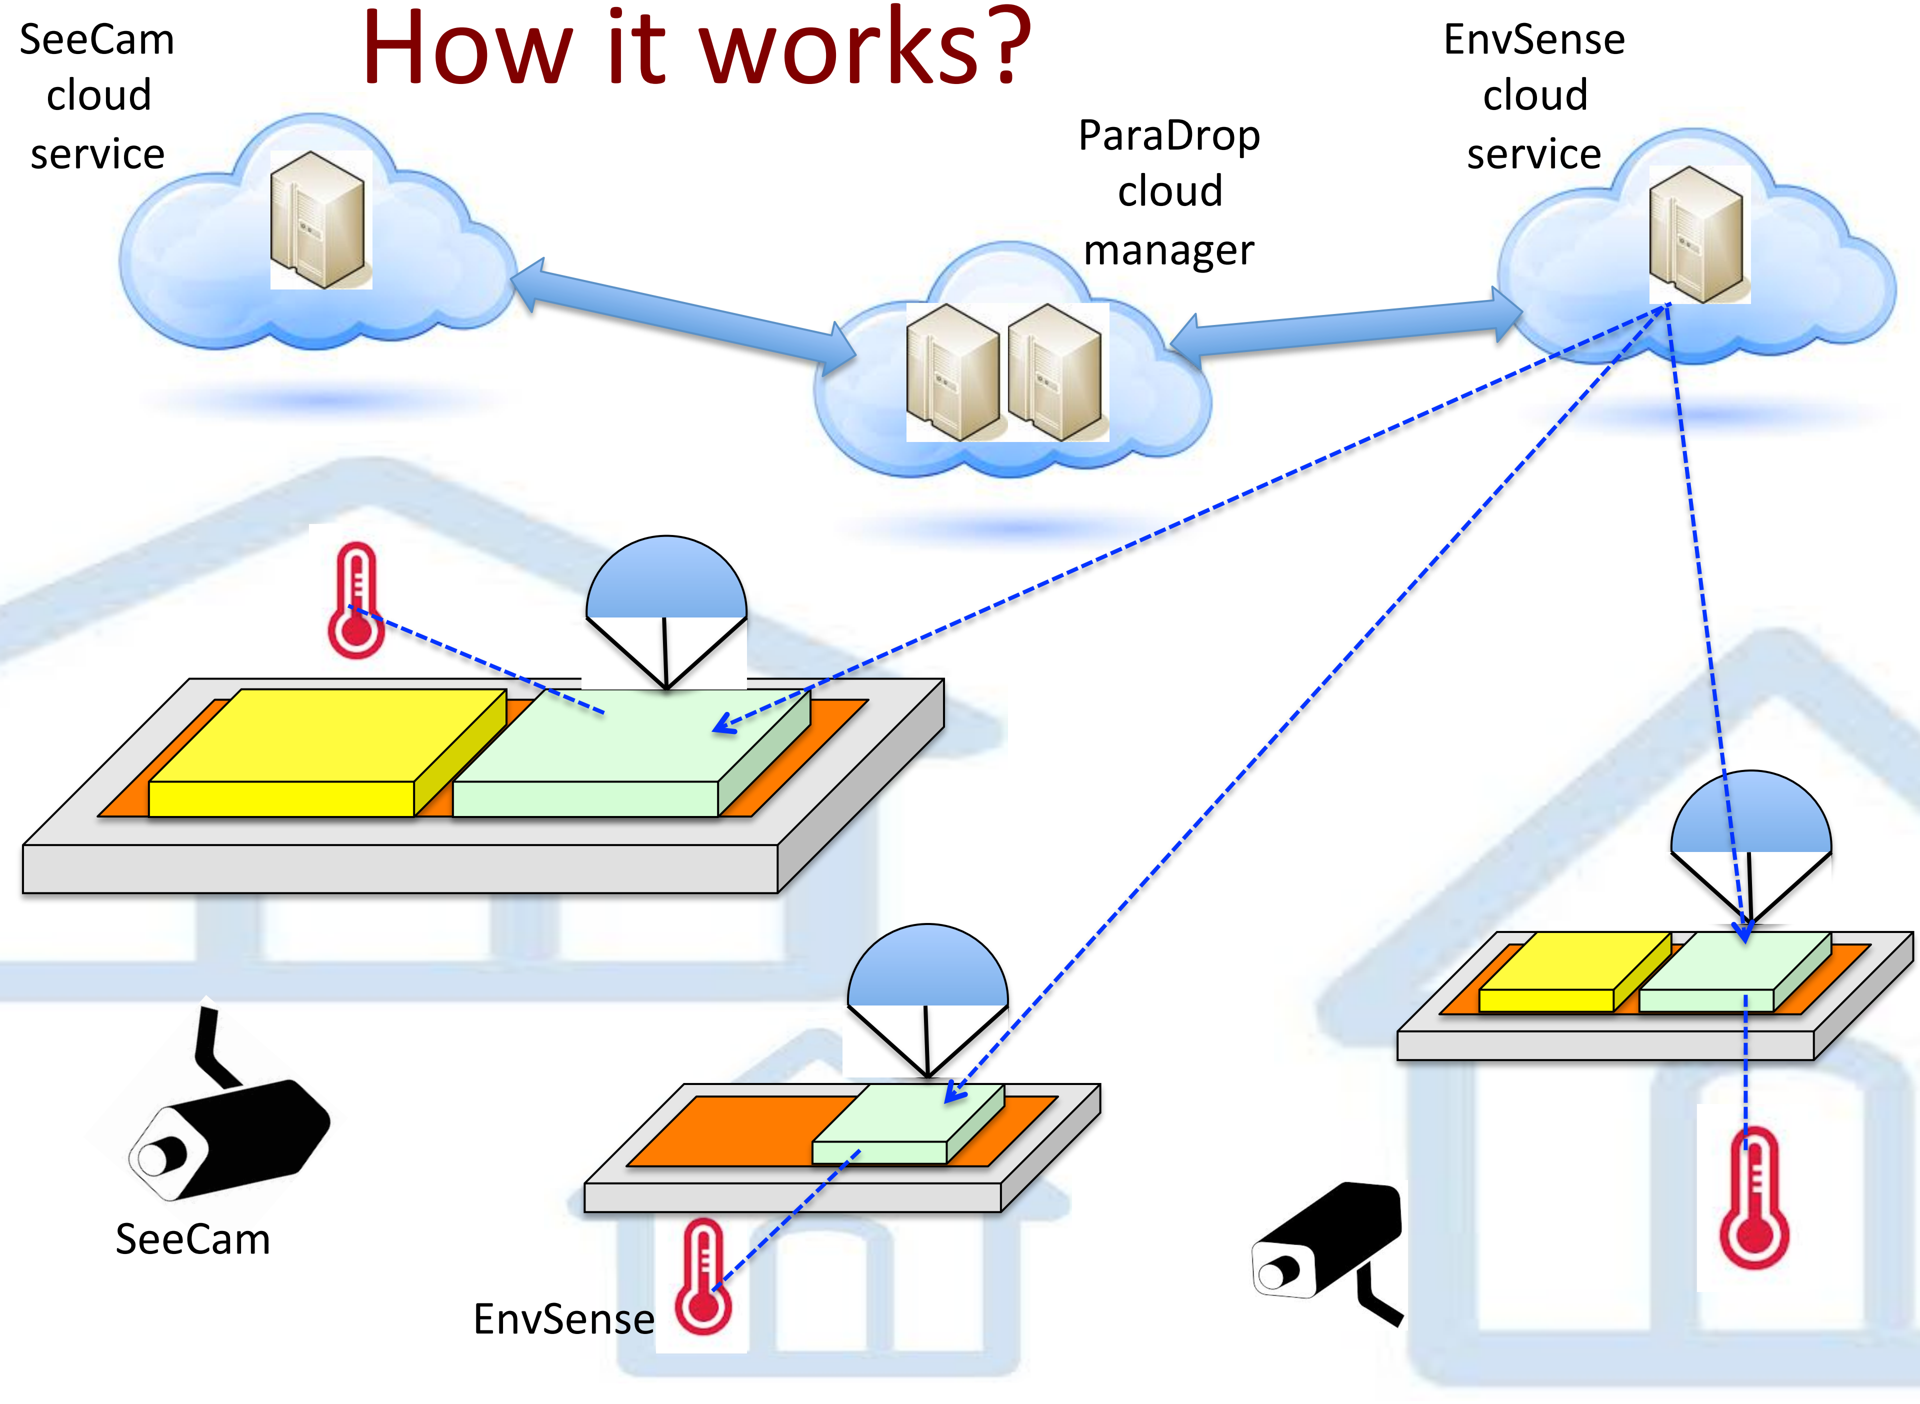
SeeCam cloud service
ParaDrop cloud manager
EnvSense cloud service
SeeCam
EnvSense

Example: Transcoding

- Transcode video to adapt to wireless channel conditions

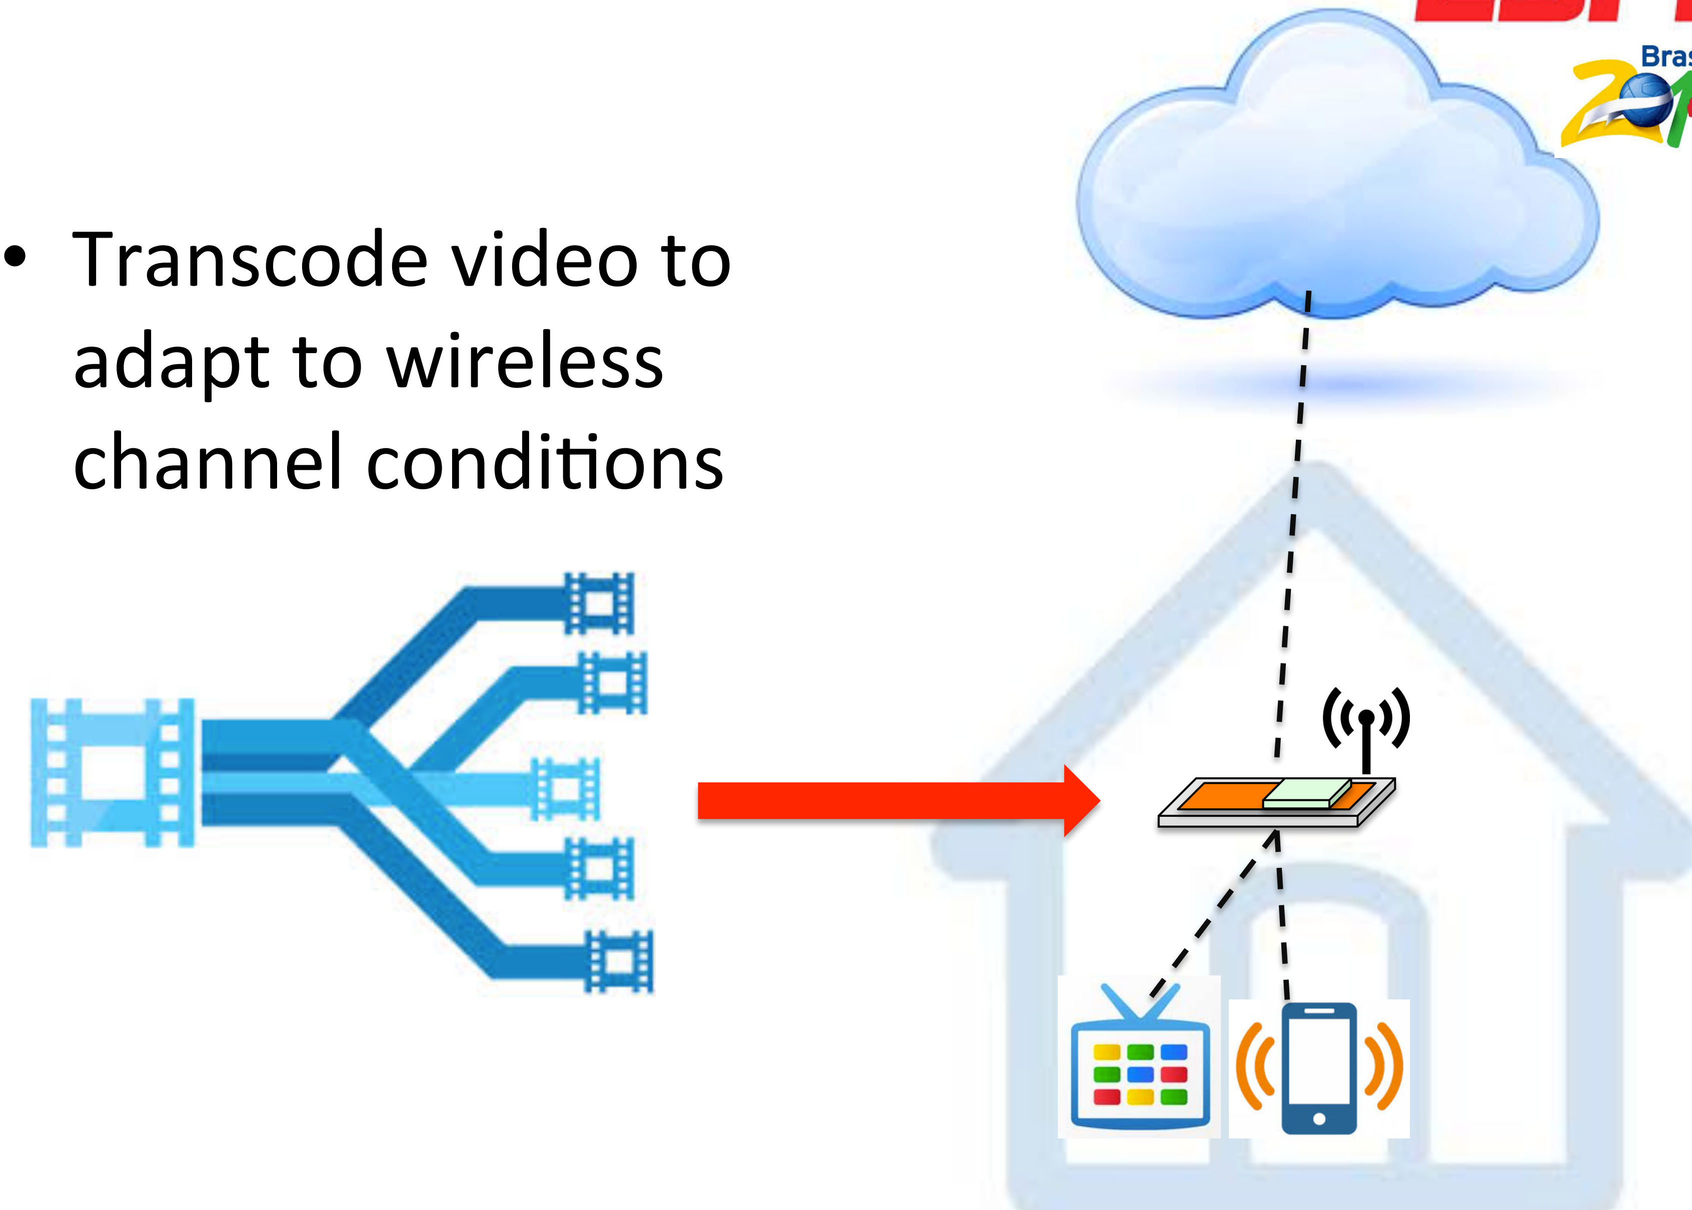

Example: Caching

- Cache movies in router from head of instant queue

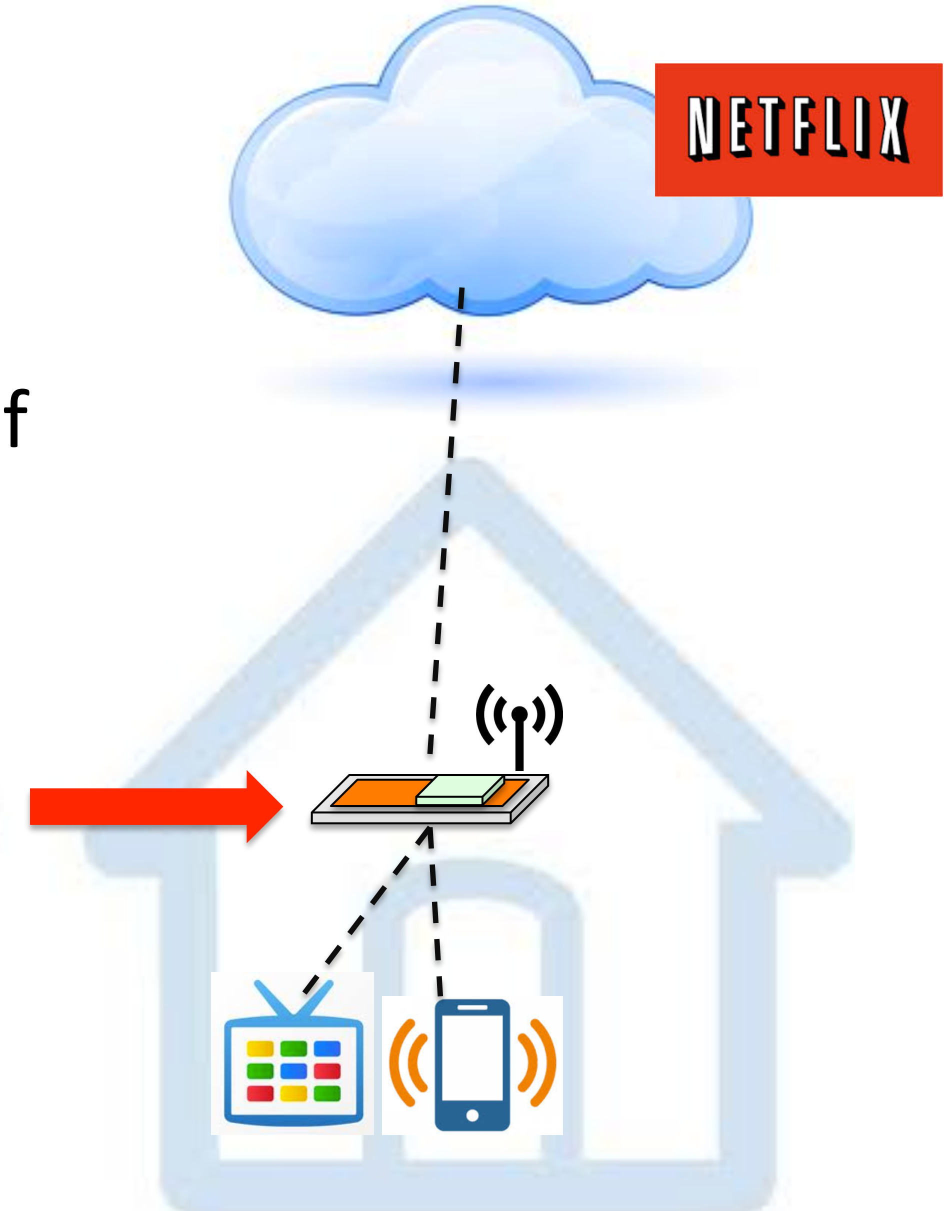

Additional capabilities

- A wireless context API
  - What else is happening in the wireless environment

Additional capabilities

- A wireless context API
  - What else is happening in the wireless environment
  - Where are devices located?
    - Which room?

Additional capabilities

- A wireless context API
  - What else is happening in the wireless environment
  - Where are devices located?
    - Which room?
  - Which devices are co-located?

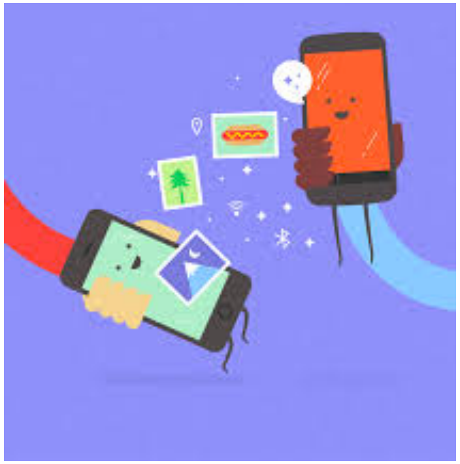

Talk outline

- Edge computing in the extreme and apps
  - ParaDrop design
    - Virtualization, RF management
  - **Application 1: Home environment management**
  - Application 2: Fuel efficient driving
  - Application 3: Drive monitoring

Application 1: Energy management

- Fine-grained home environment control

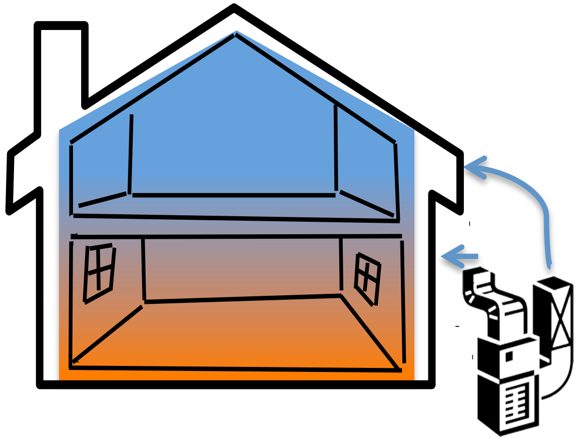

Fine-grained environment control

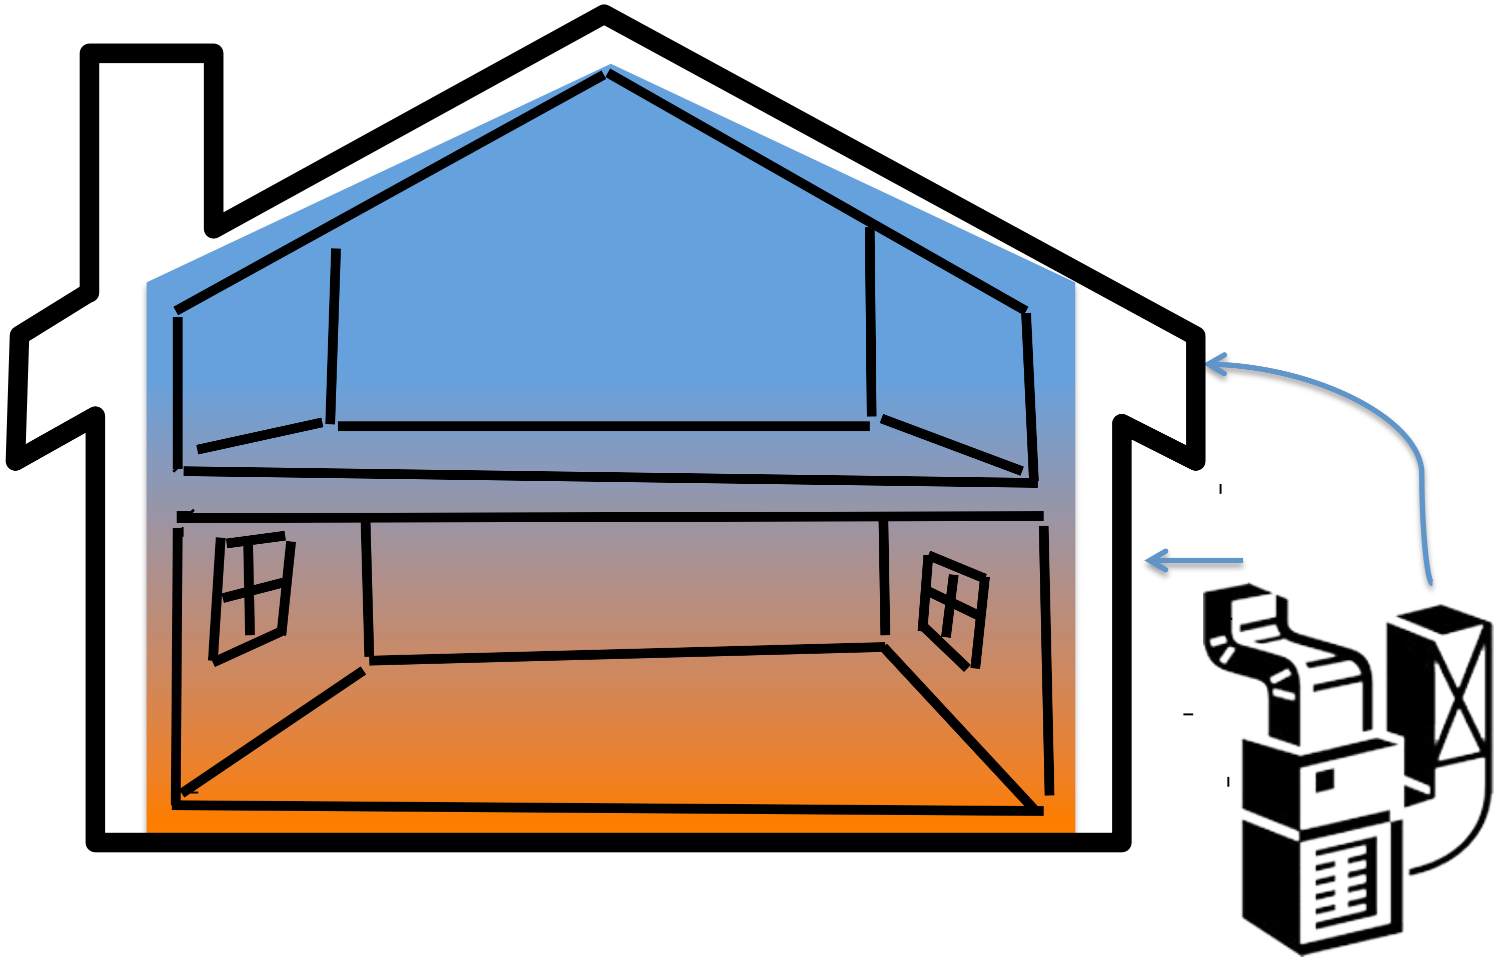

One solution: Crank it up

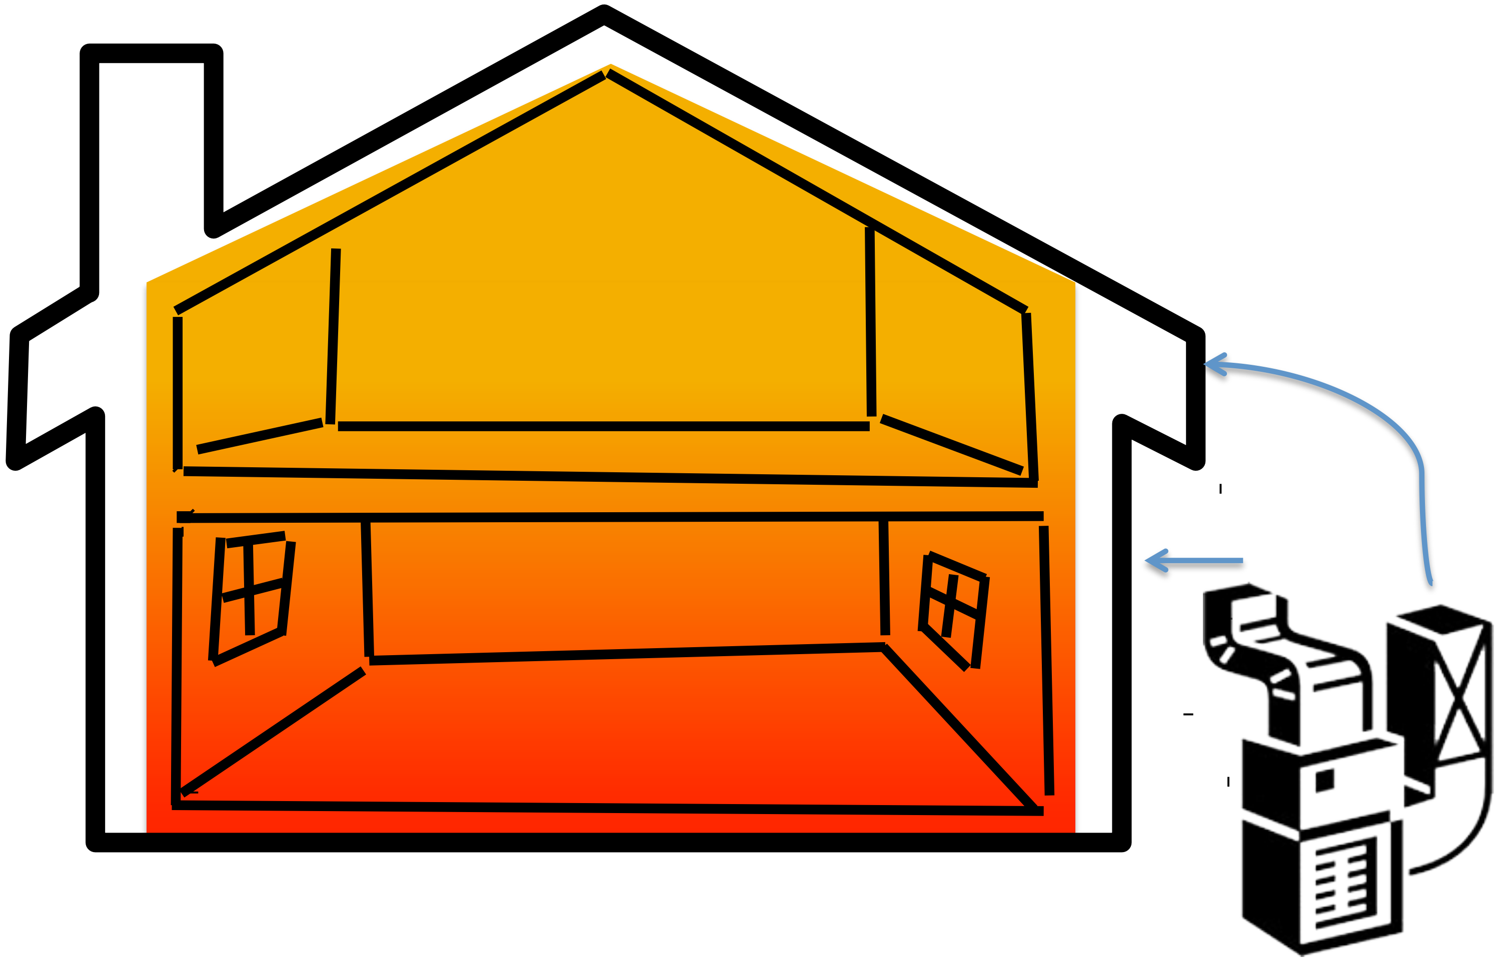

Set Point

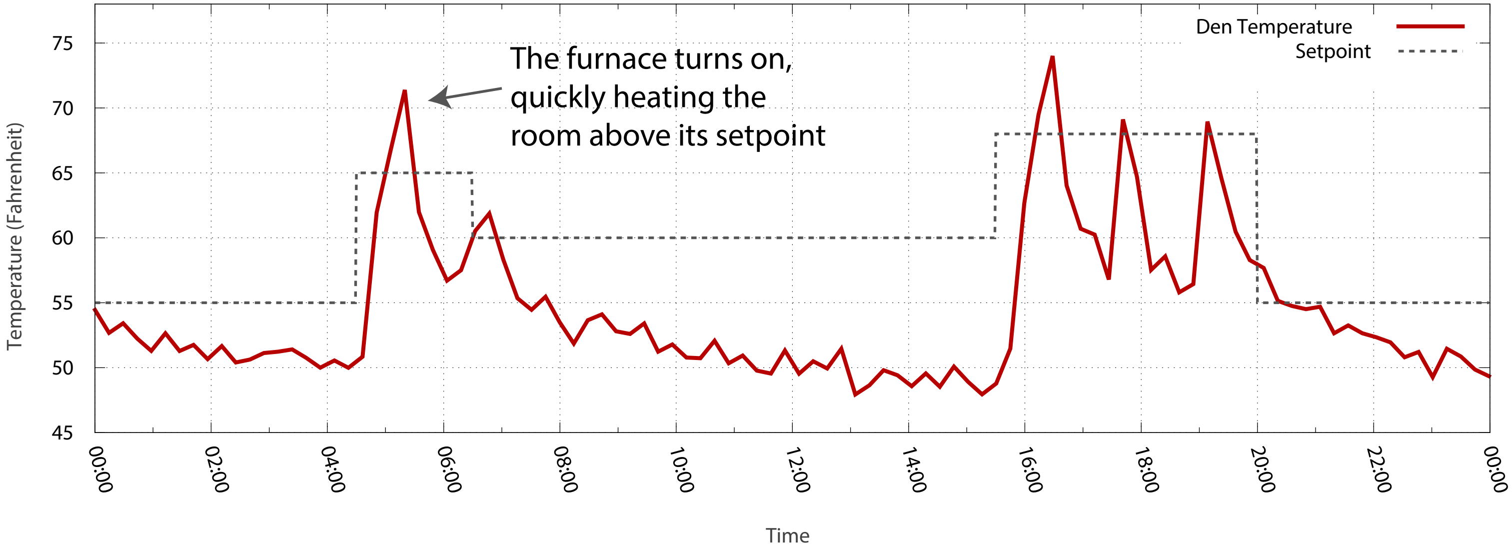

Den Temperature
Setpoint
The furnace turns on,
quickly heating the
room above its setpoint
Temperature (Fahrenheit)
75
70
65
60
55
50
45
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
00:00
Time

Airflow issues

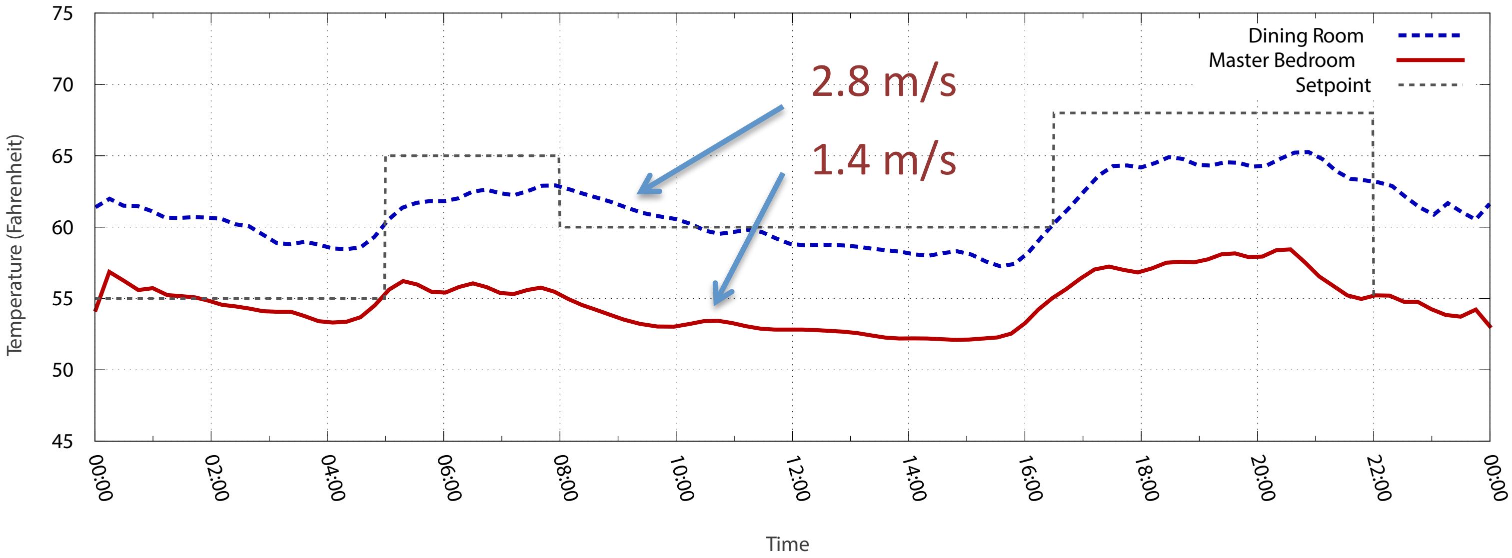

Our Solution

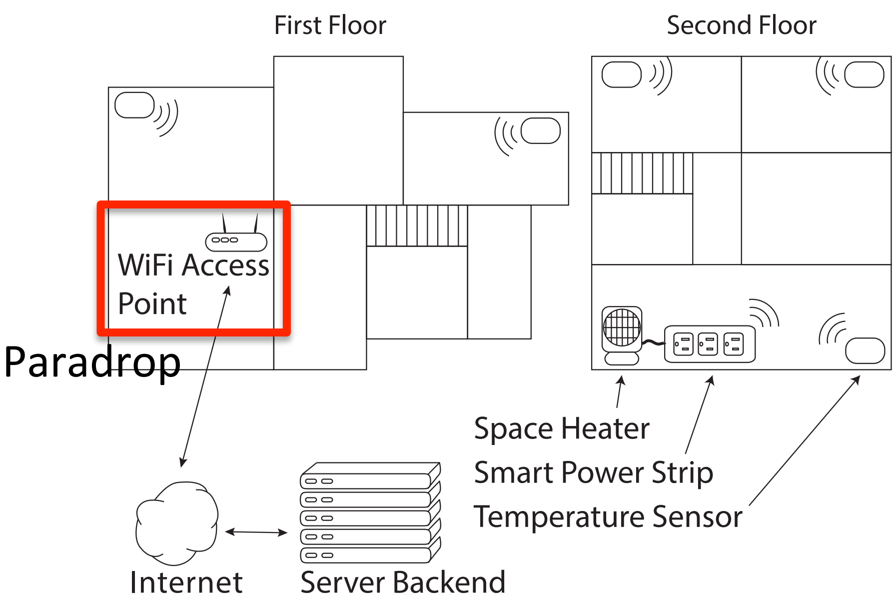

- Network-connected temperature sensors
- Register booster fans
- Space heaters
- Network-connected smart power strips
- Paradrop-based software control

Register Booster Fans

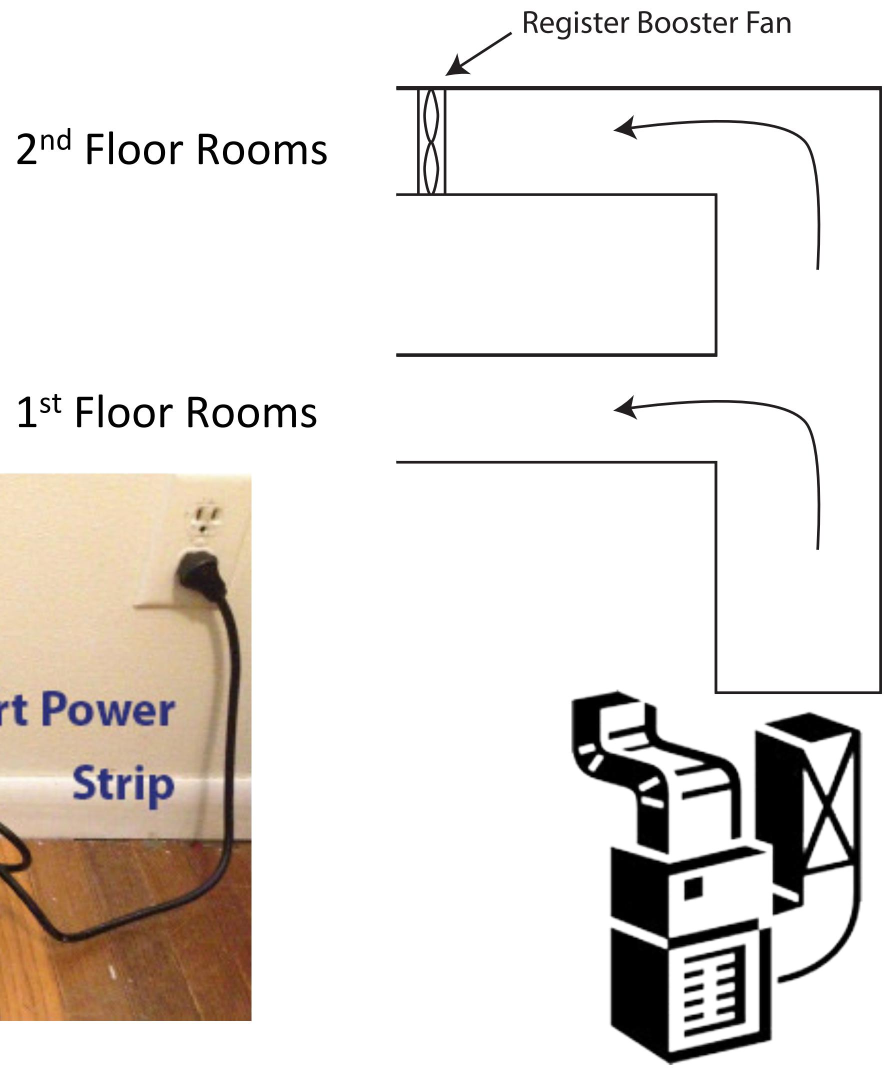

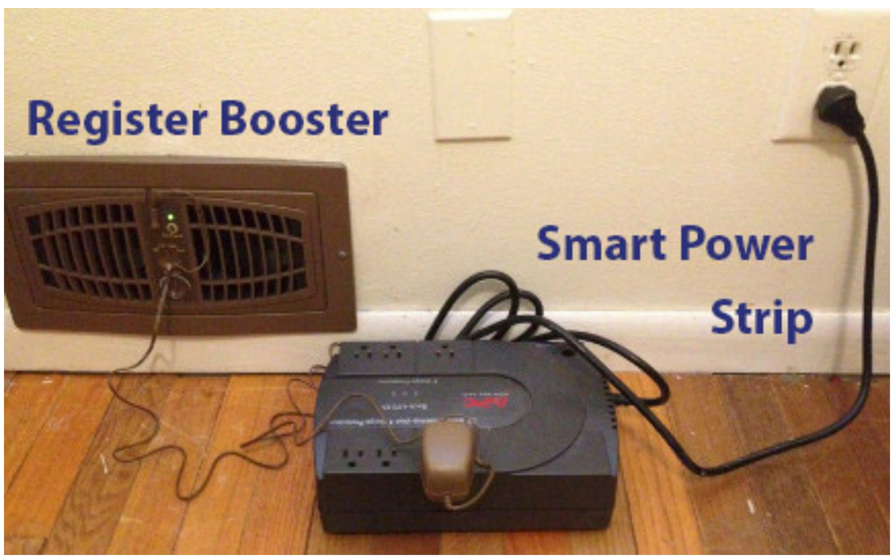

Register Booster Fans

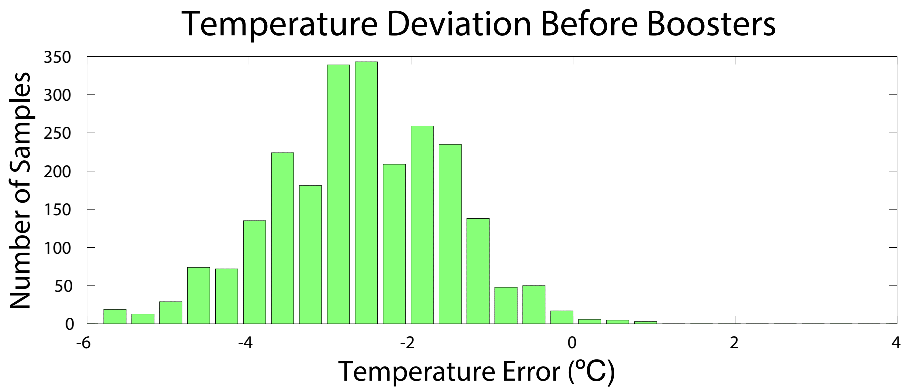

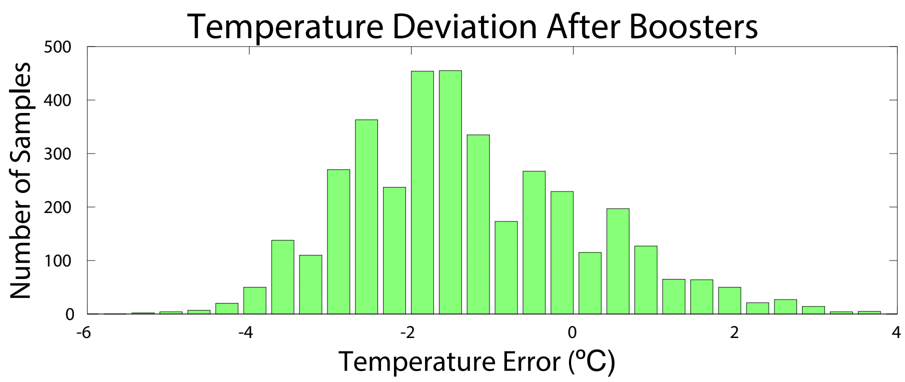

- Reduce temperature error without using much extra electricity

Space Heaters

Before our system

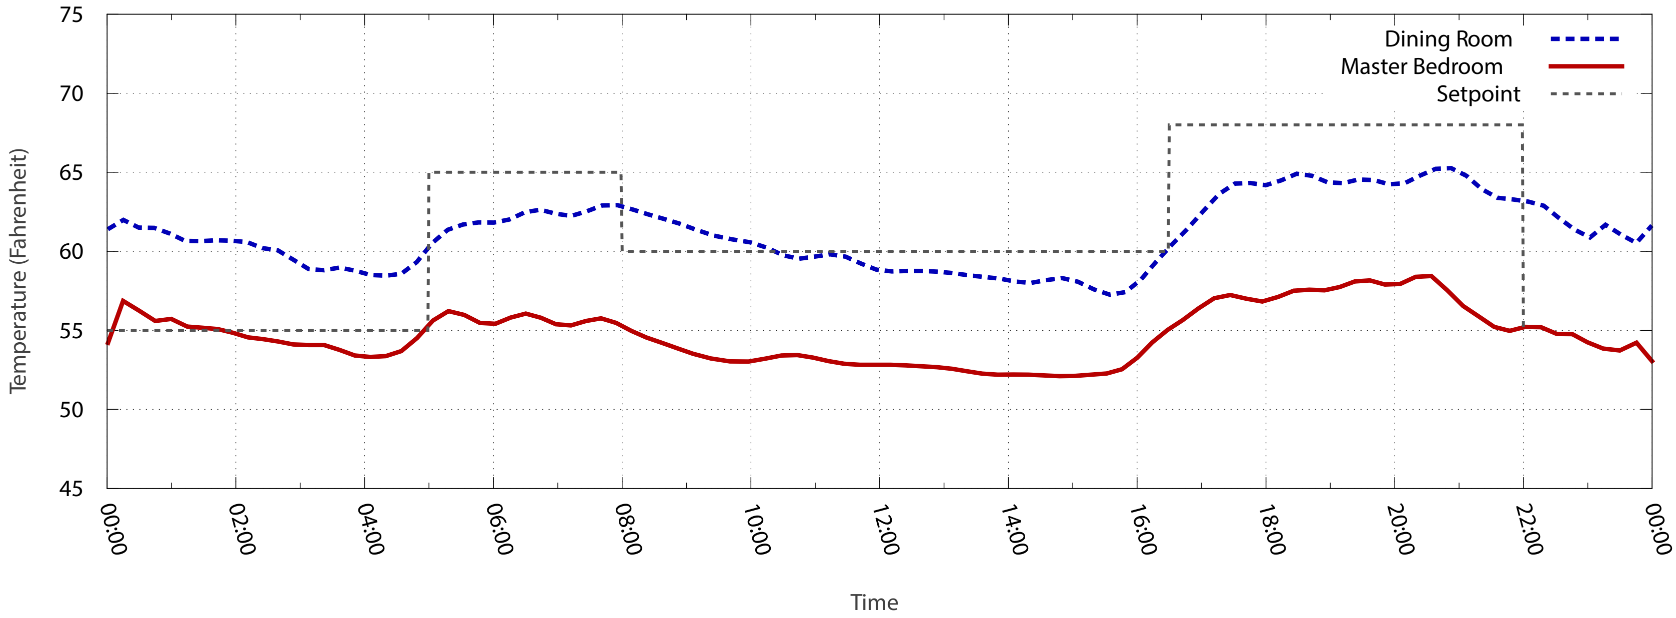

And after

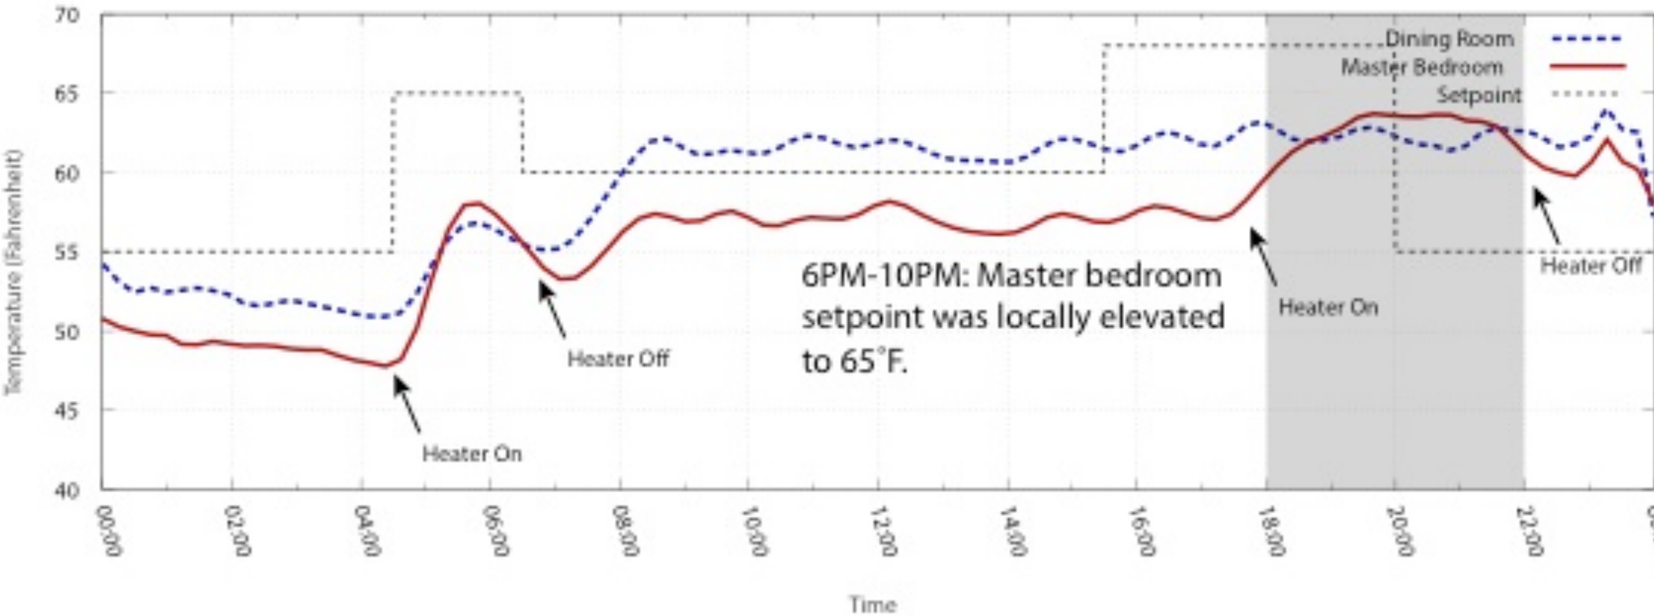
Dining Room
Master Bedroom
Setpoint
Temperature (Fahrenheit)
70
65
60
55
50
45
40
Time
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
00:00
Heater On
Heater Off
6PM-10PM: Master bedroom setpoint was locally elevated to 65˚F.
Heater On
Heater Off

Energy Savings

	Heating °C Days	Natural Gas Usage	Electric Heating
Furnace Only	804	107 Therms	0 kWh
Local Control	659	87 Therms	54 kWh
Improvement	18%	20 Therms = 586 kWh (**18.7 %**)	(54 kWh)

"The master bedroom is amazingly warm! I'm very comfortable!"

Talk outline

- Edge computing in the extreme and apps
  - ParaDrop design
    - Virtualization, RF management
  - Application 1: Home environment management
  - **Application 2: Fuel efficient driving**
  - Application 3: Drive monitoring

EcoDrive Overview

→ Automate drive actions to be fuel efficient

20% ~ 30% fuel savings

More possible if sacrificing travel time

Illustrative Example

- Conservative driver : Cruise at 15 km/h (less efficient speed)
- Aggressive driver: Accelerate to 40km/h in 1 second and cruise to the end (less efficient acceleration)
- EcoDrive: Calculate the fuel consumptions of

EcoDrive Architecture

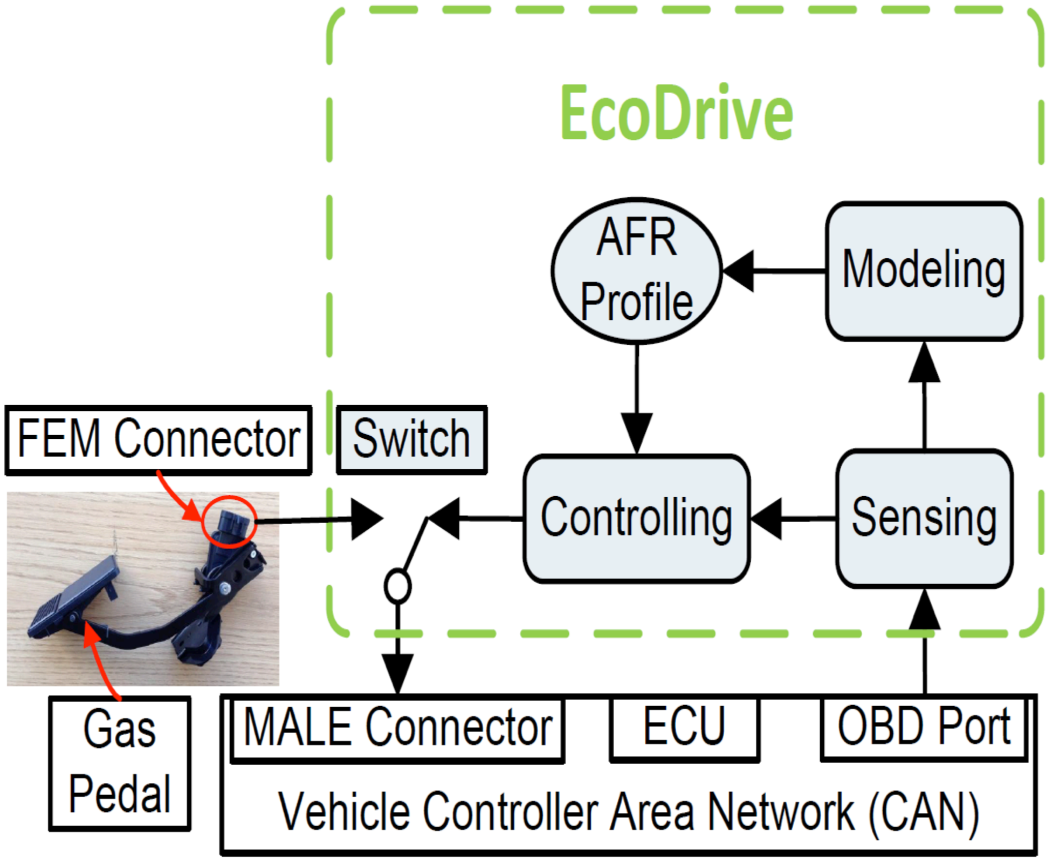

- ➔ Sensing OBD parameters
- ➔ Modeling vehicle forces
- ➔ Controlling air/fuel injection rate

OBD Parameters and Power

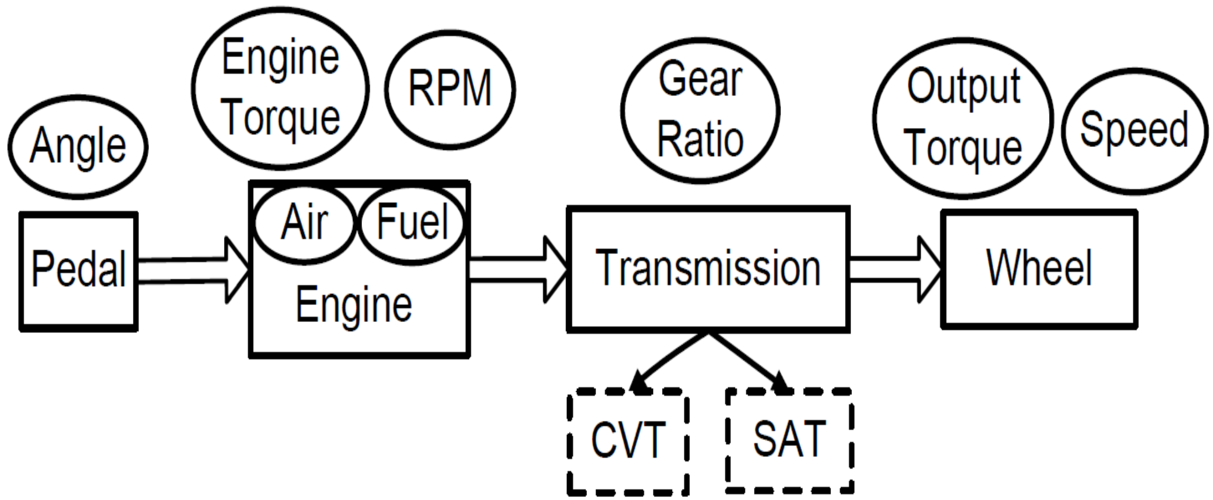

1	Gas Pedal Angel	Angle of the gas pedal, controlled by driver
2	Air/Fuel Flow Rate	Air/Fuel injection rate, controlled by gas pedal angle
3	Engine RPM	Engine rotation speed, transit power to wheel through transmission
4	Vehicular Speed	Speed of the vehicle

Model Vehicle Forces

➔Engine Propulsion

- ◆ Function of air/fuel rate and gear ratio (estimated by vehicular speed and engine RPM)

➔Drivetrain loss and wind resistance

- ◆ Function of propulsion when driving in constant speed

➔Grade resistance

- ◆ Function of altitude changes (extracted from online elevation dataset)

Build AFR Profile (A Lookup Table)

AFR(v, a) : the air/fuel rate when accelerates at a (m/s/s) under speed v(km/h)

Air/Fuel Rate	0.0 m/s/s	0.1 m/s/s	0.2 m/s/s	… …
1 km/h				
2 km/h		AFR(2, 0.1)		
3 km/h				
4 km/h				

Edge Controlled Gas Pedal

➔Gas Pedal (drive-by-wire)

- ◆ Human driver press the gas pedal
- ◆ The position sensor senses gas pedal position
- ◆ The gas pedal sends the position value to the Electronic Control Unit (ECU)
- ◆ ECU controls the volumes of air/fuel injected into the engine

➔EcoDrive Controller (Emulate gas pedal)

- ◆ It calculates the gas pedal position value
- ◆ It sends the value to the ECU through an Arduino

Dynamic Programming based Driving Strategy

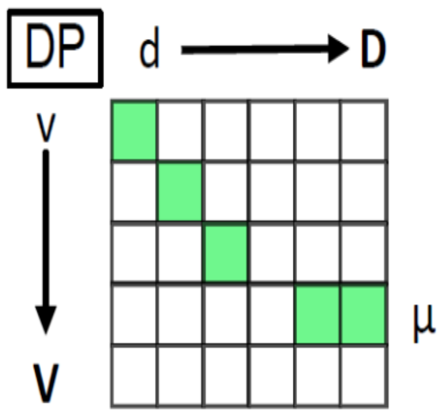

- D: road segment length
- V: speed limit
- S(v, d): minimum fuel cost at distance d with speed v

➔ Case 1: The car cruises to state S(v + 1, d + 1)

- ◆ S(v + 1, d + 1) = S(v + 1, d) + AFR(v + 1, 0) * __**time**__

➔ Case 2: The car accelerates to state S(v + 1

Implementation

- ➔ Hardware
  - ◆ OBD Scanner with ELM327 USB interface
  - ◆ Arduino board converts digital gas pedal position to voltage signals
- ➔ Software
  - ◆ One thread writes commands to OBD interface through serial communication
  - ◆ One thread reads OBD parameters and write gas pedal position to Arduino board

In-vehicle Setup

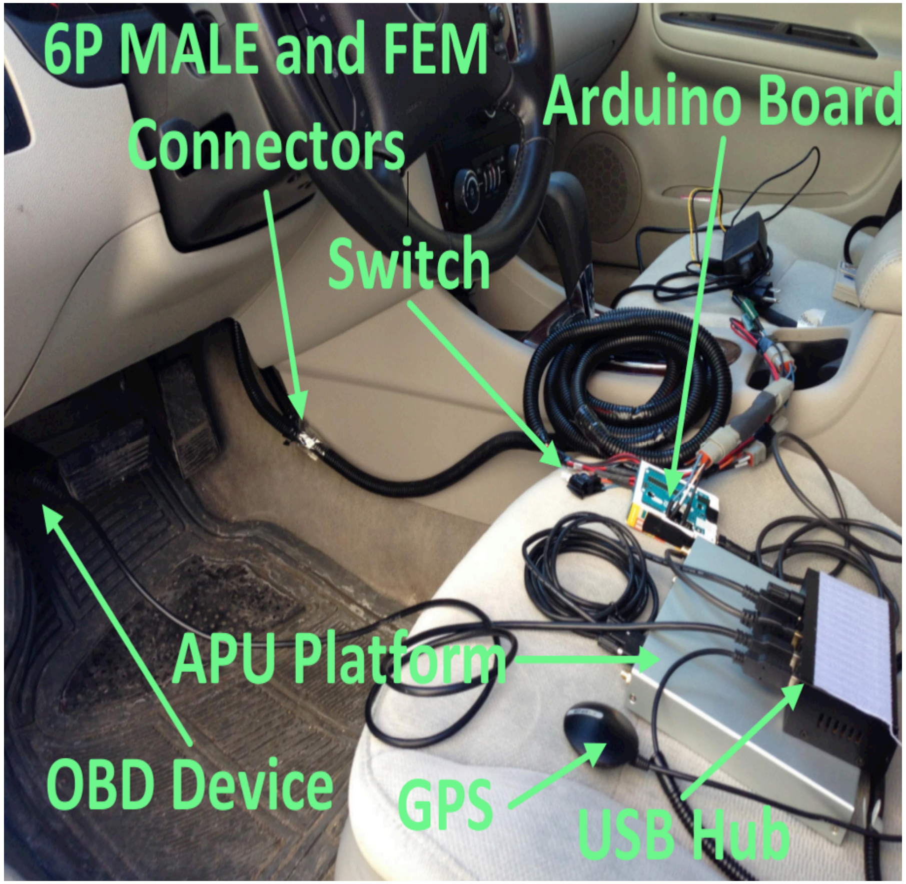

Evaluation

- ➔Real road test [over 100 miles]
  - ◆ Urban: Road segments with various lengths (50-1000m)
  - ◆ Highway: Various highway segments (2km each)
- ➔Comparison [Kilometer per Liter (KPL) and Travel Time]
  - ◆ Theoretical value
  - ◆ Human drivers
  - ◆ Cruise control

OBD Data Collection

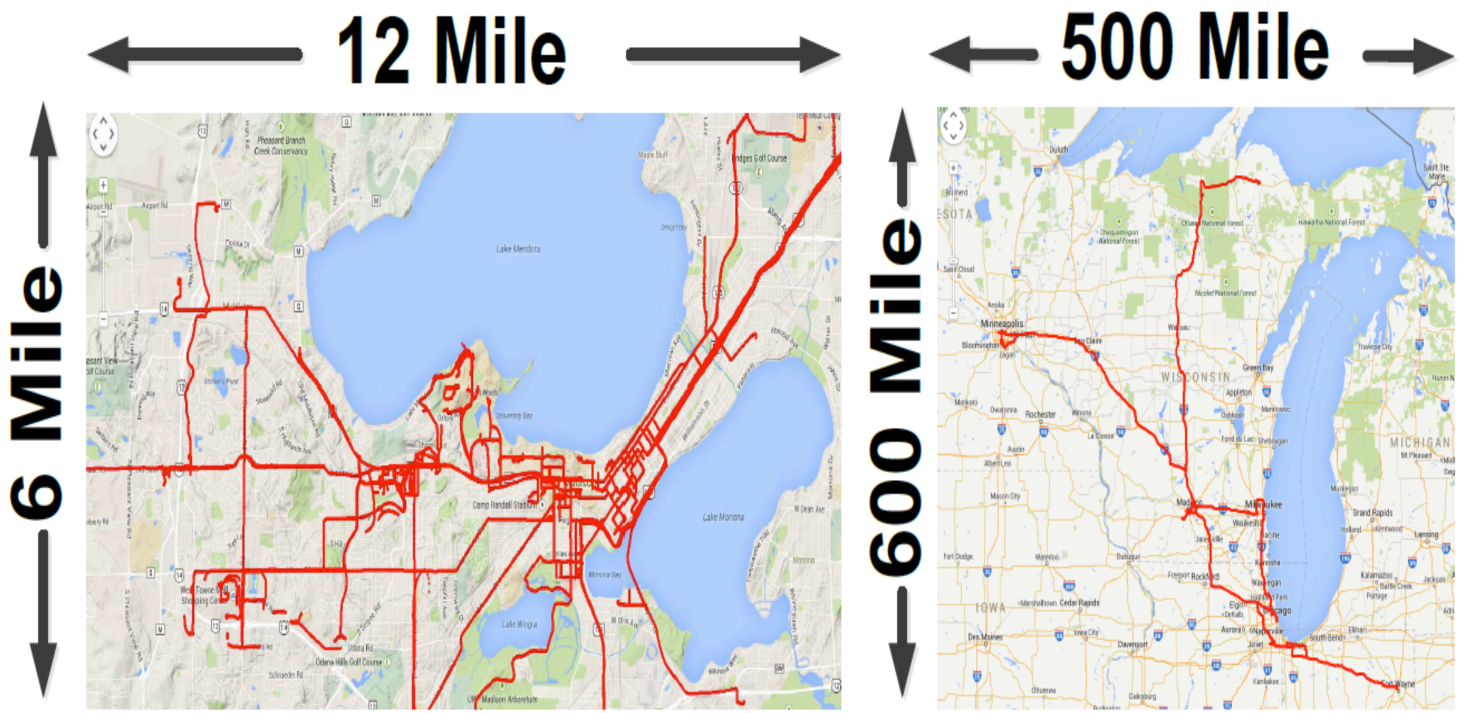

- →Urban: Madison and Chicago, 5000+ miles
- →Highway: Local highways and cross-

OBD Data Collection

No.	Car Model	Urban	Highway
1	Chevrolet Impala 2011	1051	852
2	Nissan Rogue 2011	1198	1063
3	Subaru Forester 2011	651	757
4	Buick LaCrosse 2006	599	649
5	Volkswagen Tiguan 2014	600	347
6	Honda Accord 2013	173	840
7	Toyota Camry 2011	35	338
8	Volkswagen Touareg 2014	21	156
9	Nissan Altima 2014	193	271
10	Nissan Rogue 2011	105	0
11	Subaru Legacy 2015	119	30
12	Mazda CX5 2014	202	89

Case Study: Cruise Control

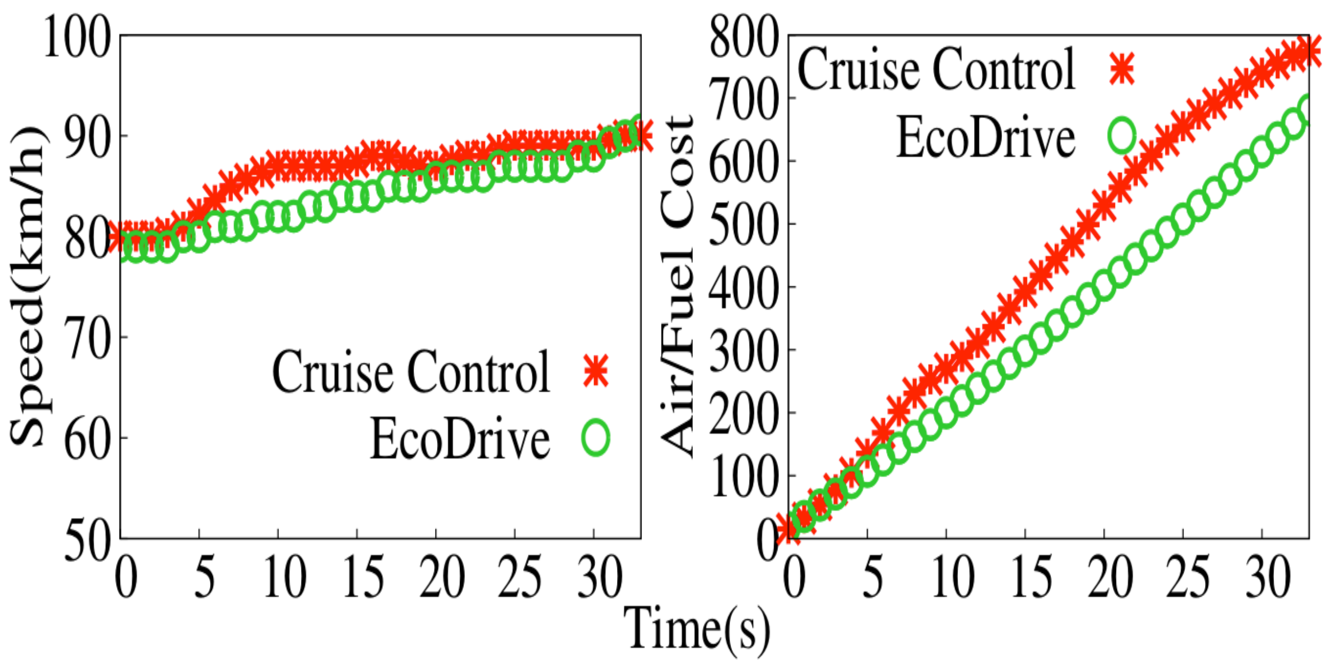

- Cruise control: accelerate aggressively on upslope or human manipulation
- EcoDrive: gradually change air/fuel injection

Travel Time vs. Fuel Efficiency

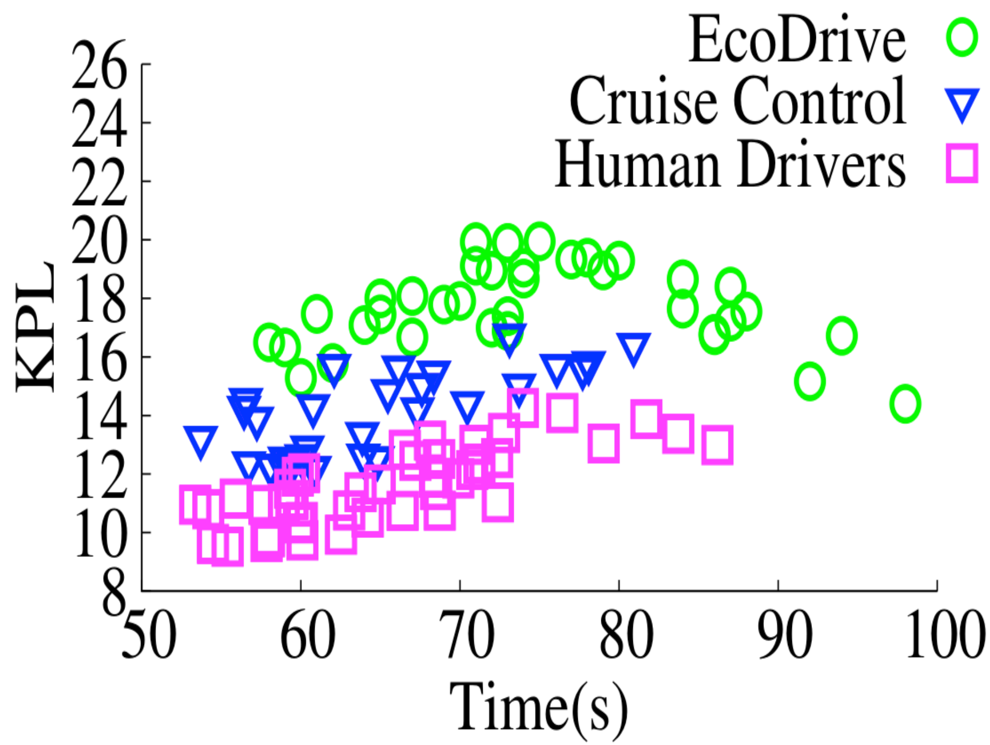

- EcoDrive achieves higher KPL than human drivers in similar travel time

Talk outline

- Edge computing in the extreme and apps
  - ParaDrop design
    - Virtualization, RF management
  - Application 1: Home environment management
  - Application 2: Fuel efficient driving
  - **Application 3: Drive monitoring**

Understanding Driving Behavior

Safe Driving Apps

DriveWell
Cambridge Mobile
Telematics

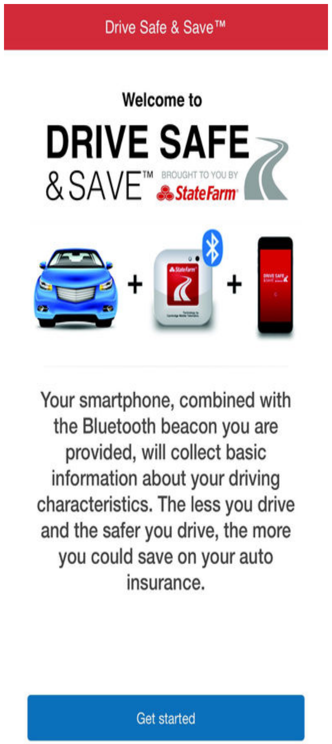

DriveSa
fe
StateFar

Sensing Vehicle Dynamics

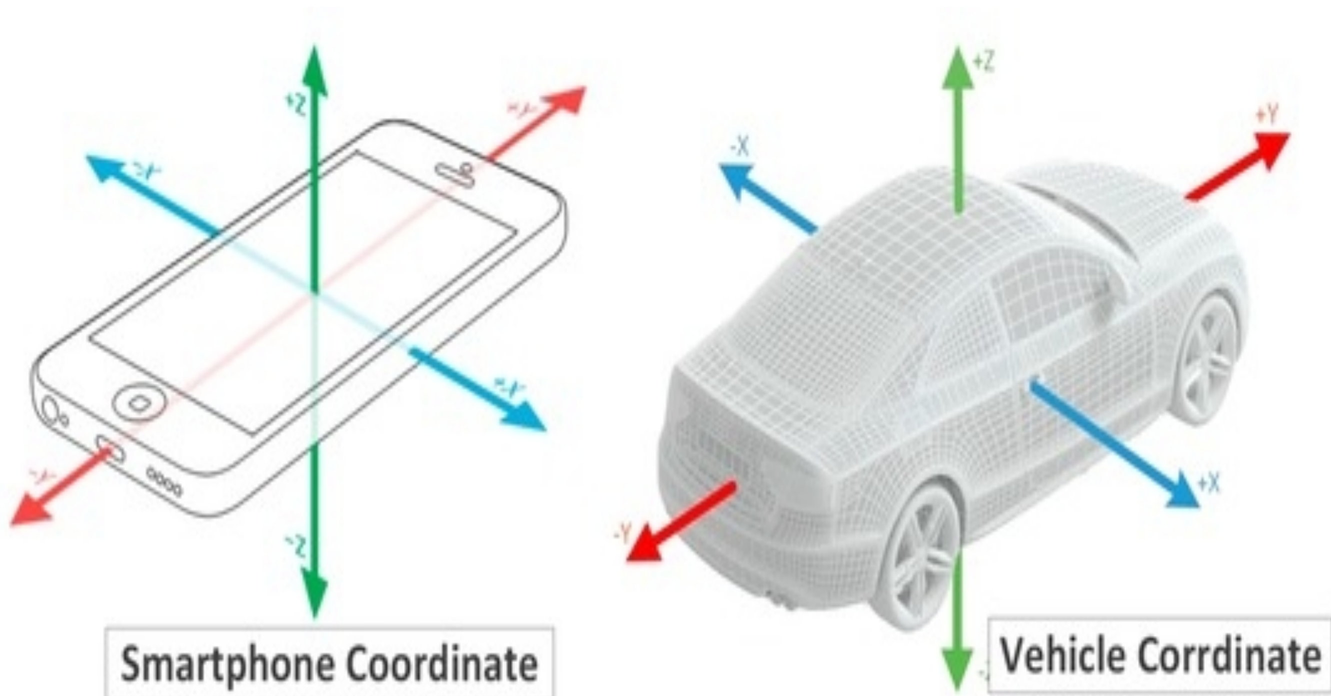

Hard Brake Detection

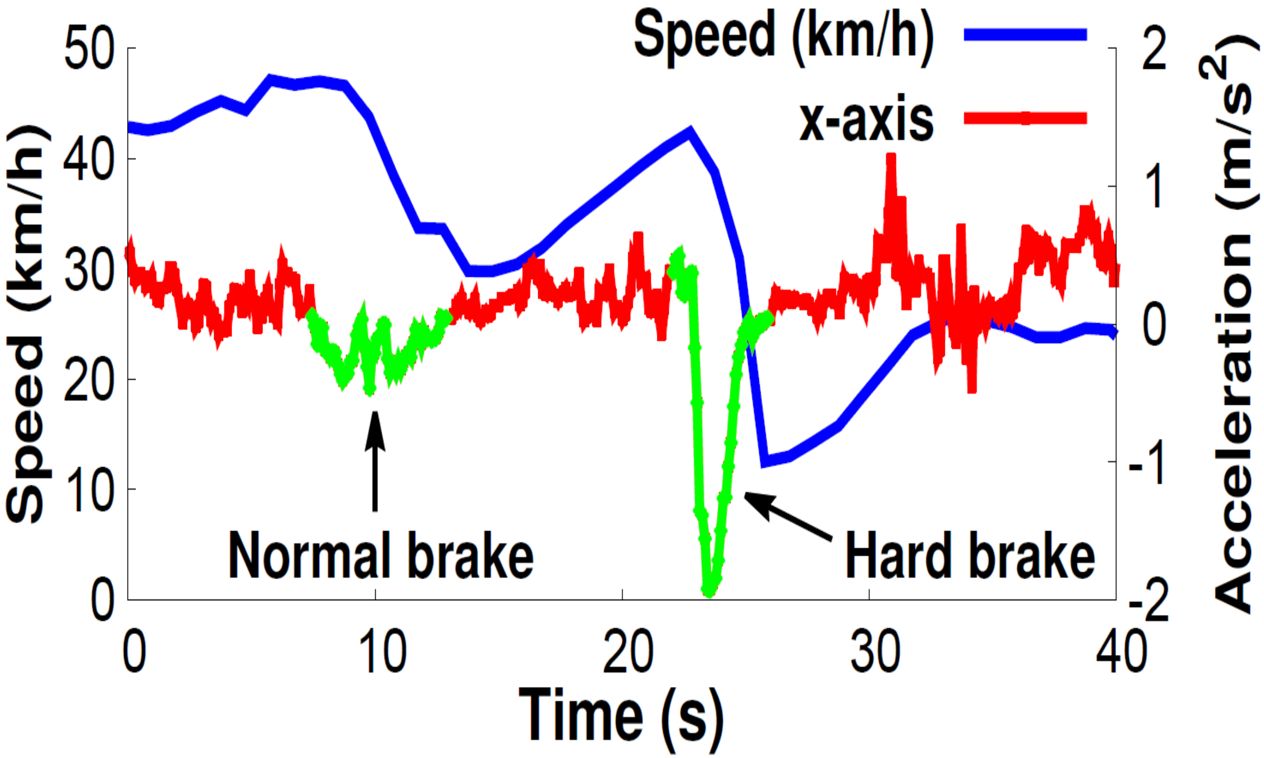

Turn & Lane Change Detection

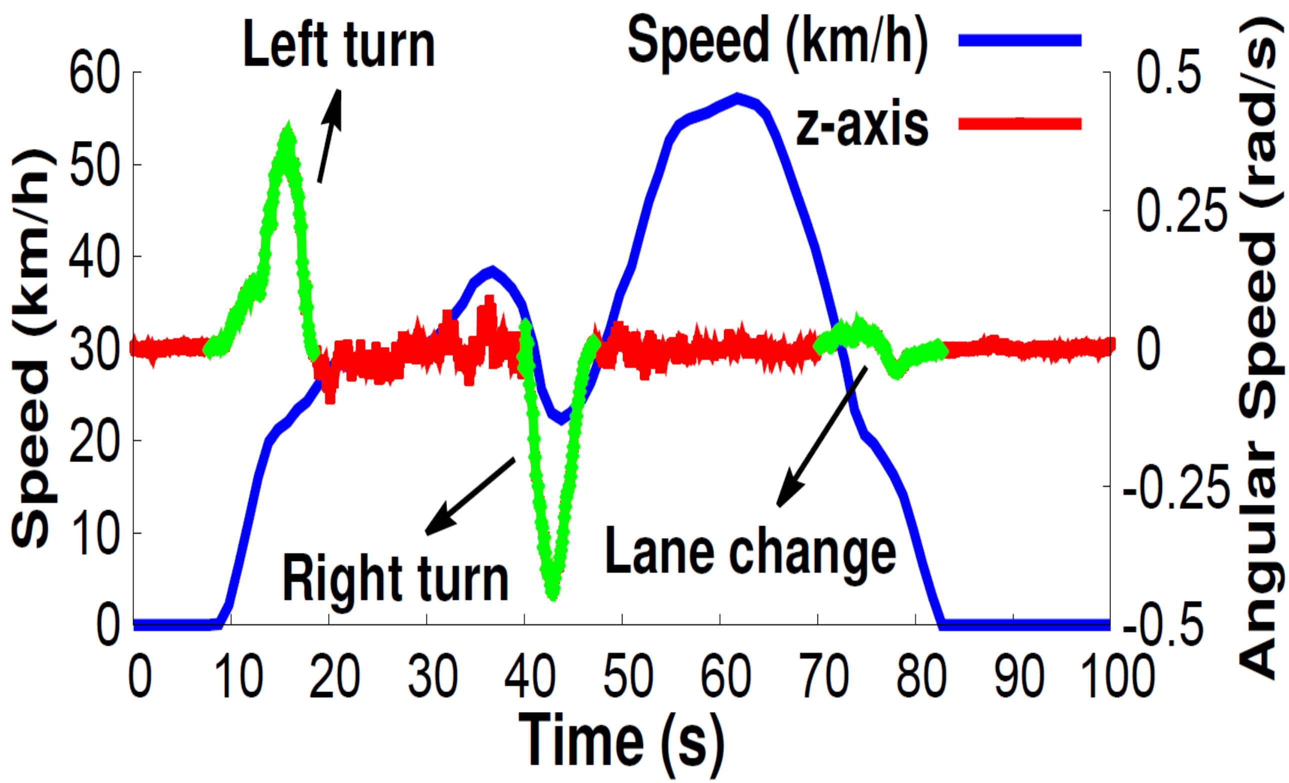

Evaluating driving behaviors

- IMU sensors can provide accurate analytics on ***what*** motions happened during a drive
  - Hard brakes, sudden lane changes, etc.
- Data from IMU sensors does not answer ***why*** the driver acted in that manner
  - Driver distraction, surrounding traffic, etc.

Finding Missing Information

Drive Analytics using Audio Visual Sensors

1. **Motion Sensors** — Detecting driving events through motion sensors, e.g. accelerometer, gyroscope, magnetic field sensor, GPS
2. **Visual-Audio Sensors** — Monitoring driver behaviors and surrounding traffic through audio-visual sensors. e.g. camera, microphone
3. **Evaluate Driving behaviors** — Augment driving behavior analysis by combining motion and visual-audio sensors

DriveAQ Architecture

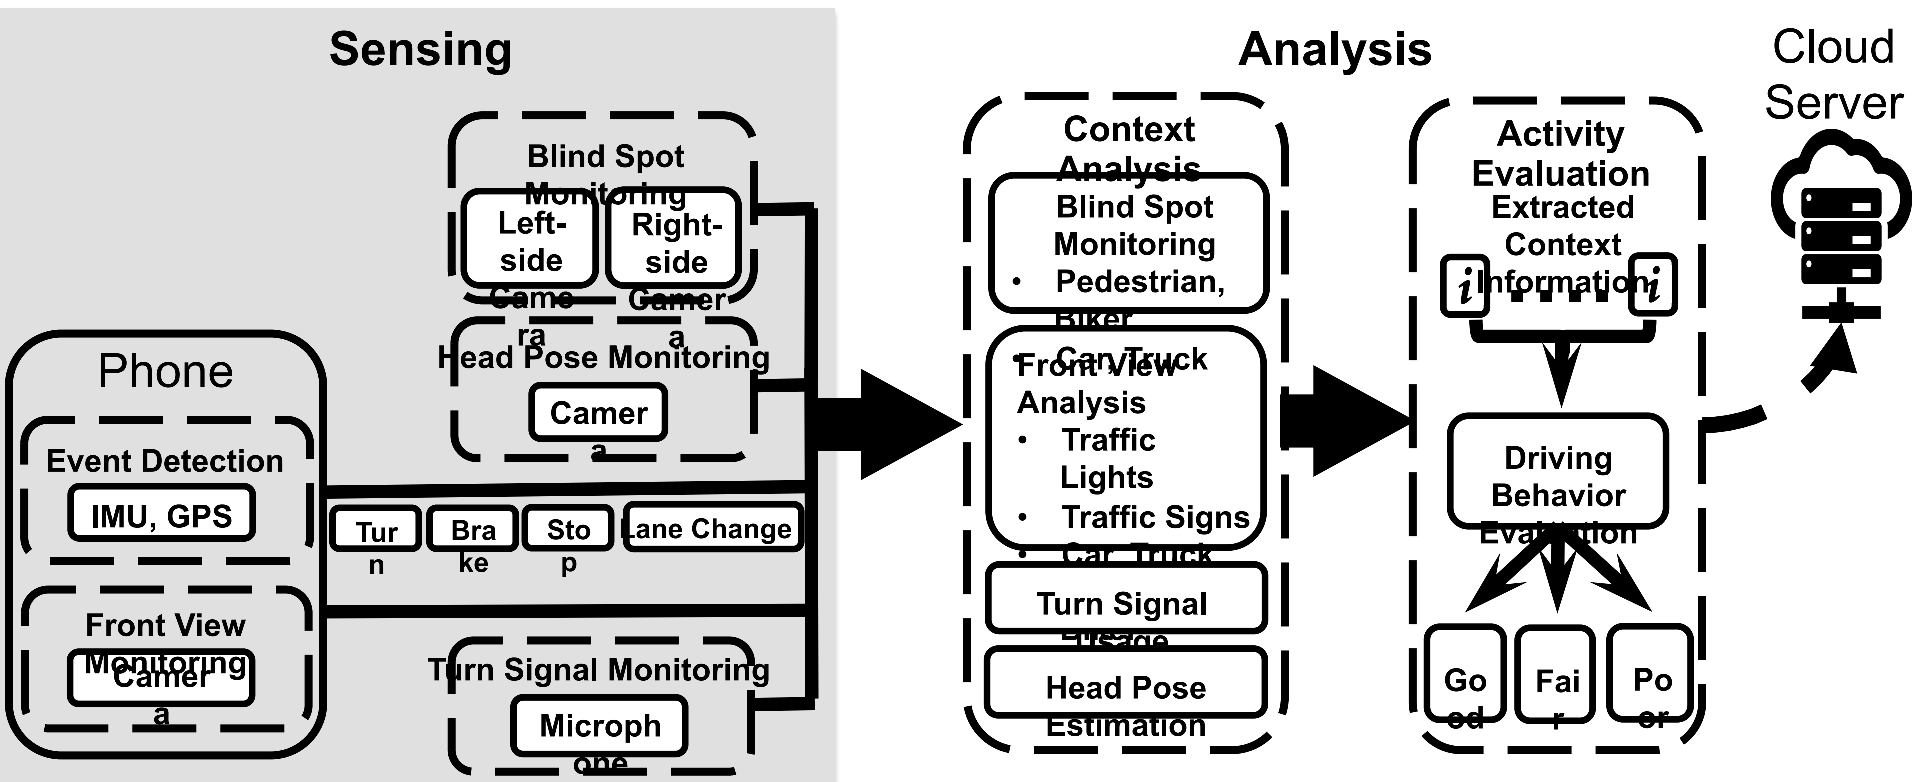

Why Processing Data at Edge?

- Large volume of data
- Many vehicles need to stay on the road for long and contiguous periods of time
- Opportunities to offload the data to remote data centers are infrequent and inconvenient.

Hardware

Nvidia Jetson TX2

USB Camera

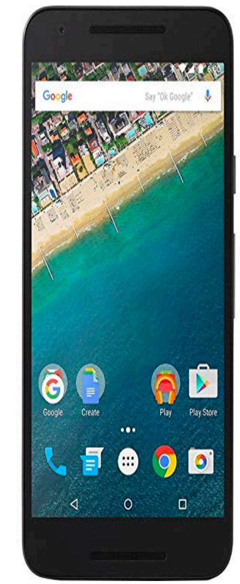

Pho ne

Object Detection

- Face detection and head pose estimation
- Objects in blind spots
  - Vehicles and pedestrians
- Objects in front view
  - Traffic signs, vehicles, bikers, and pedestrians

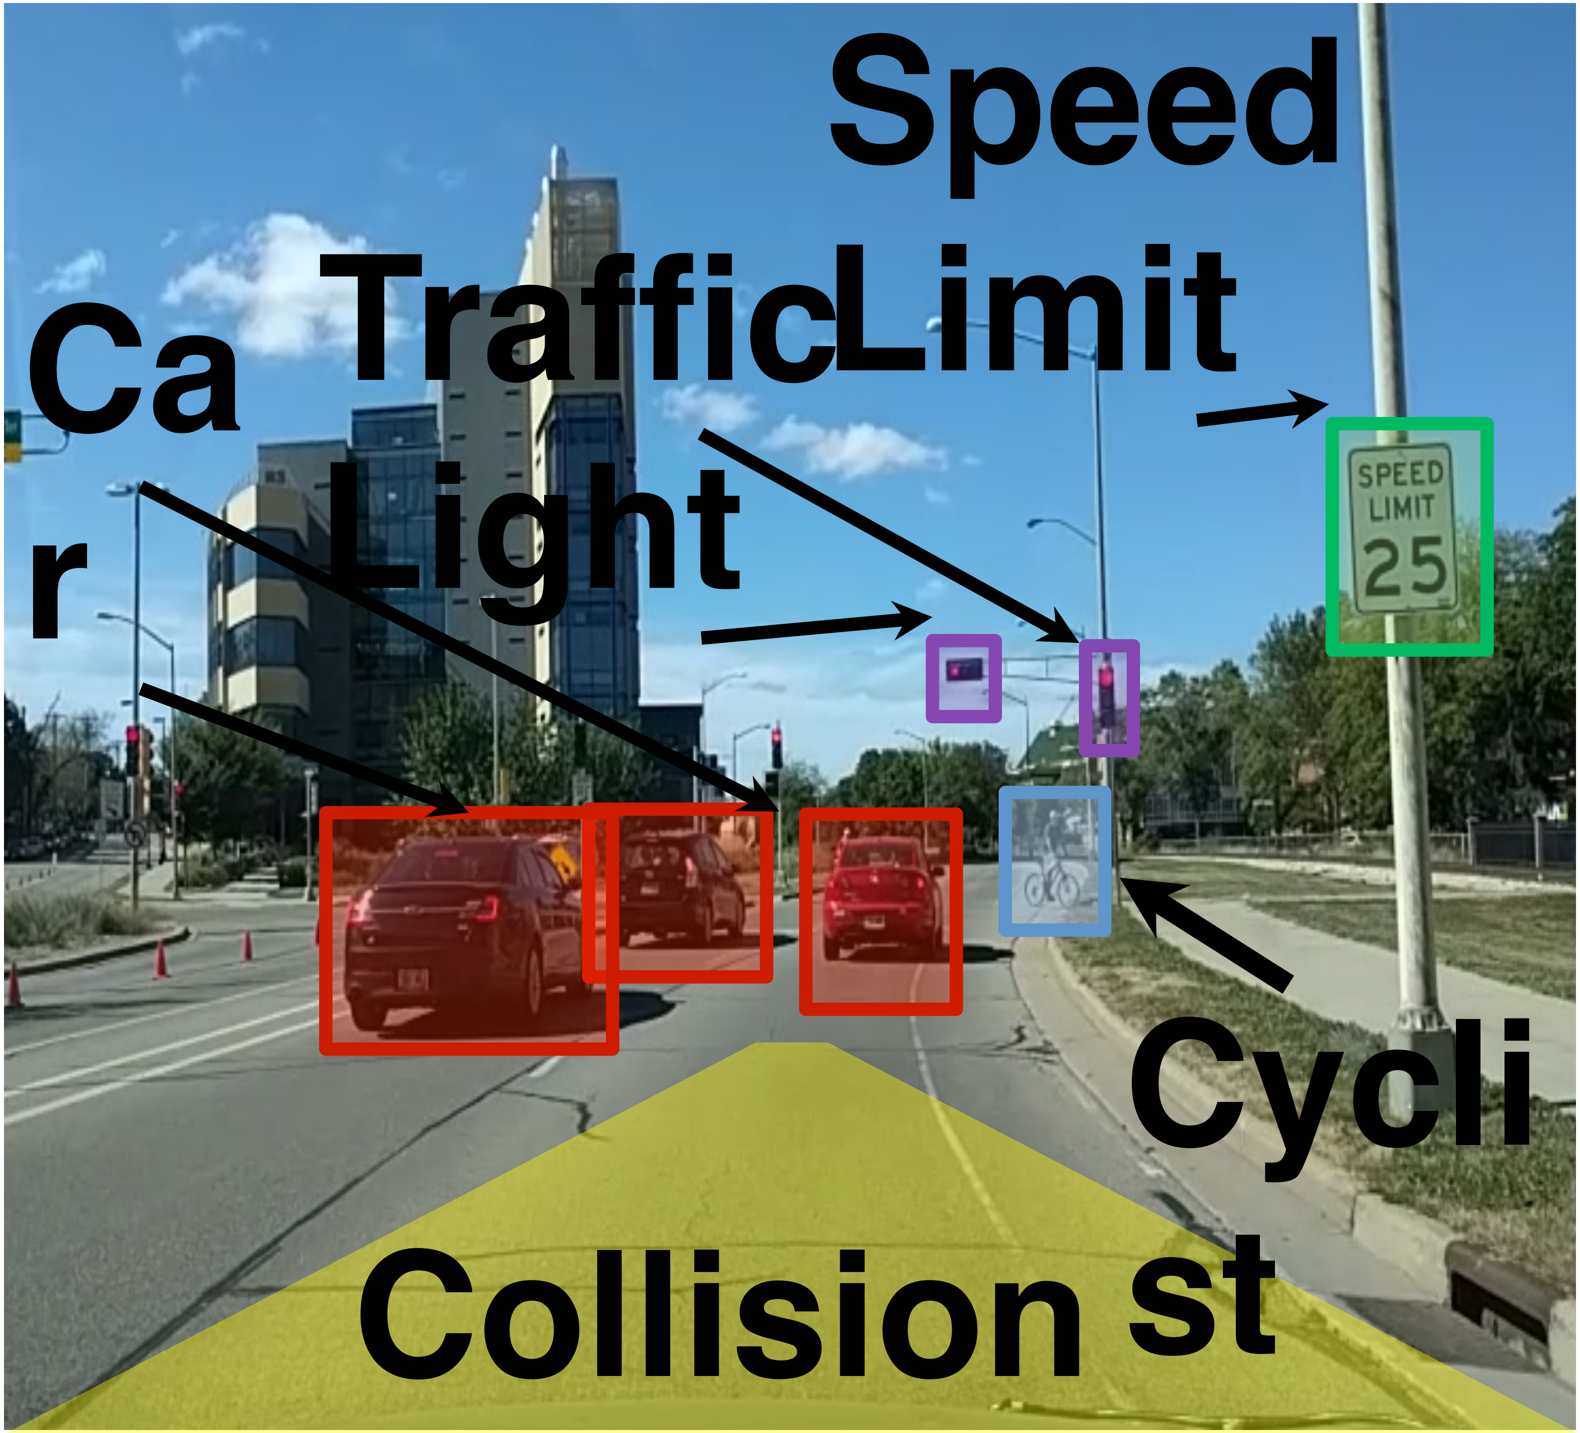
Speed
Limit
Traffic
Light
Ca
r
Cycli
st
Collision
SPEED
LIMIT
25


Pedestri
an
Stop
Sign
Ca
r
STOP
Collision

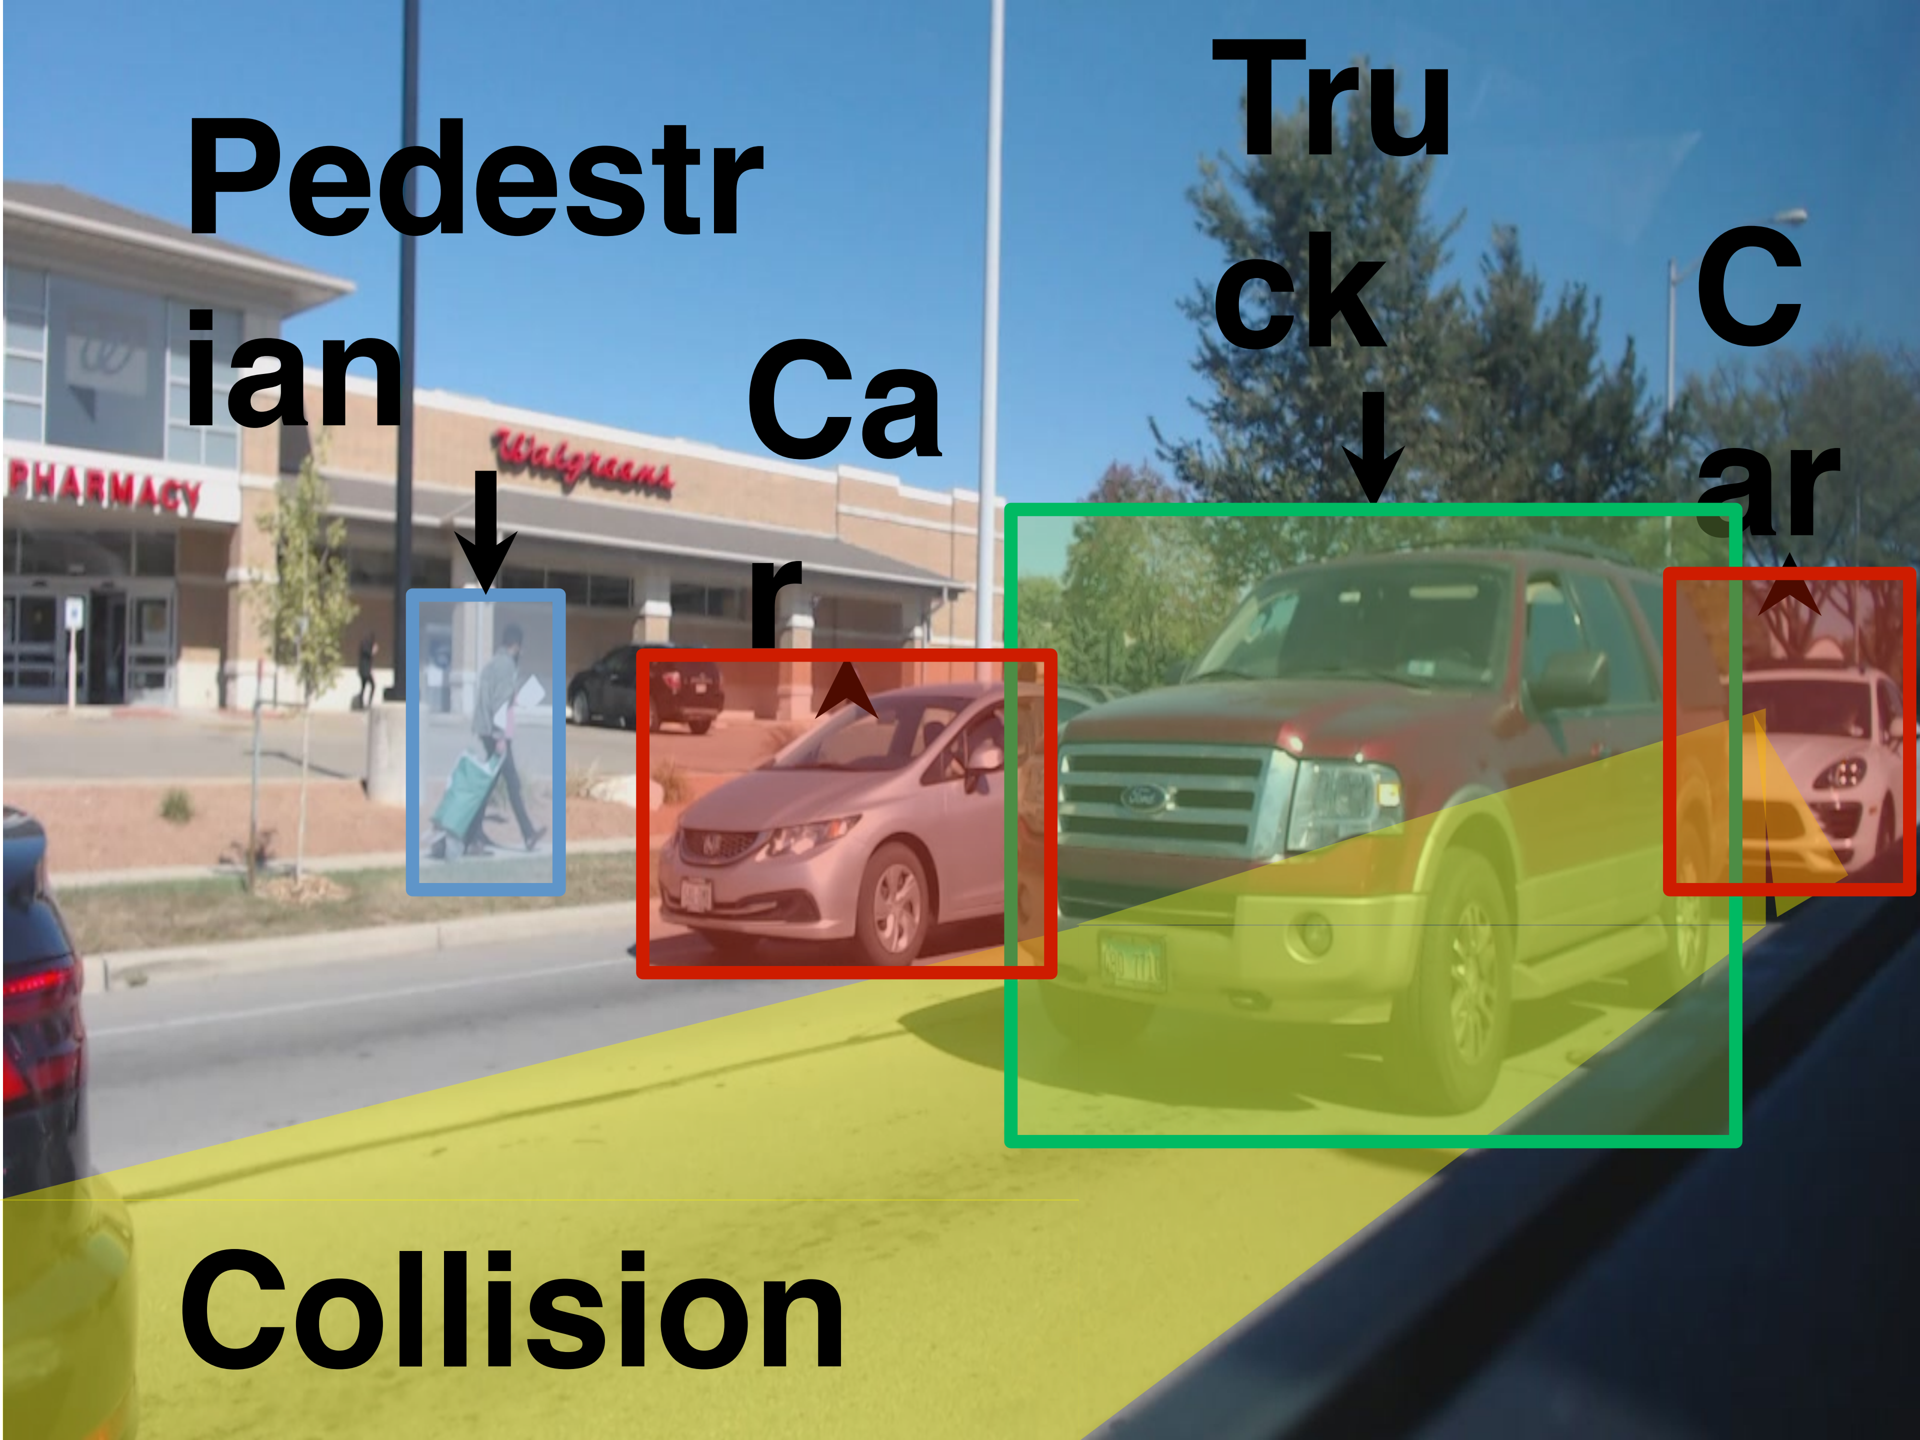
Pedestr
ian
Ca
r
Tru
ck
C
ar
Collision
PHARMACY
Walgreens

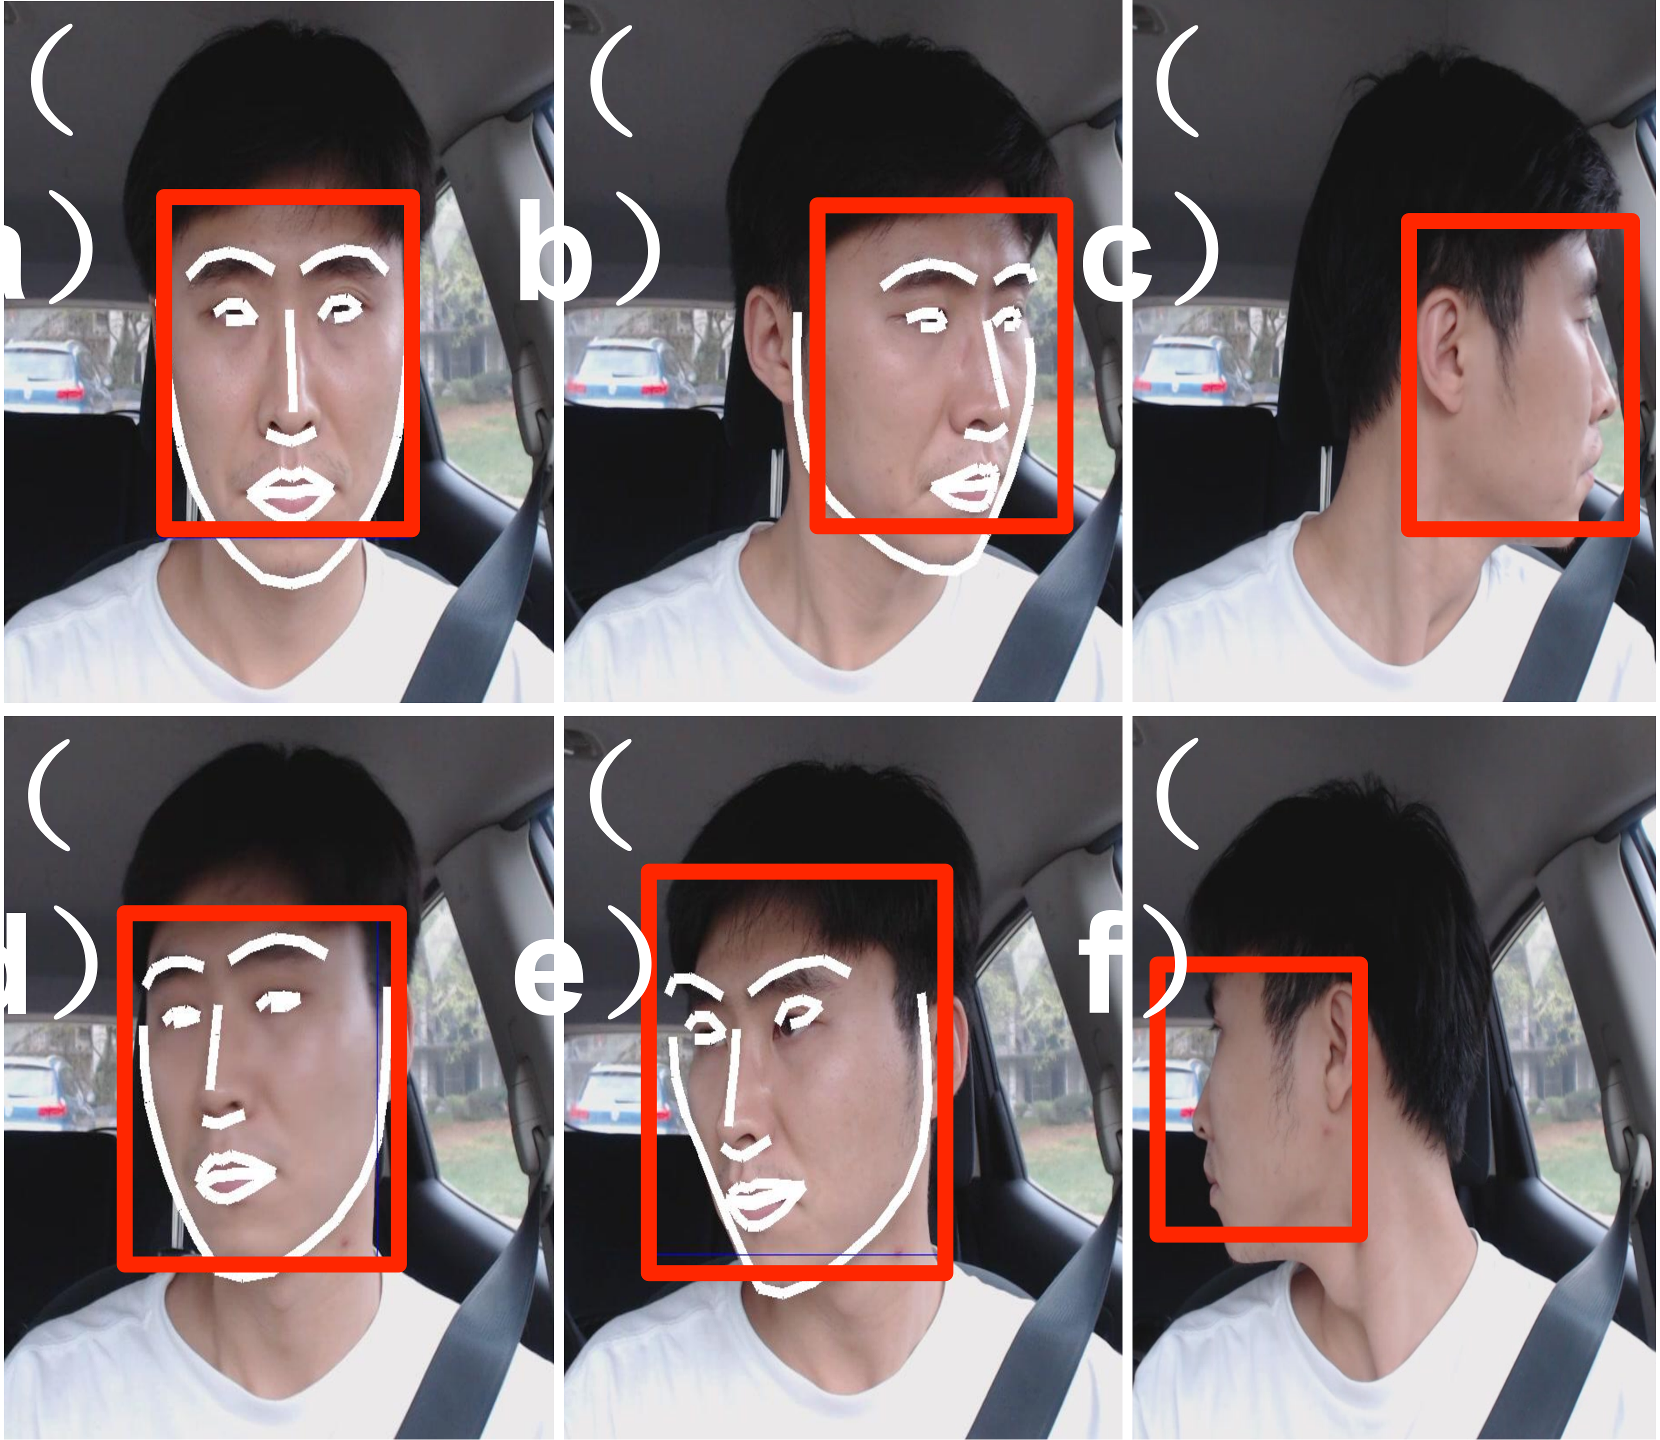
(b)
(c)
(e)
(f)

Driving Activity Evaluation

- Develop deep understanding of each activity by combining information acquired from multiple sensors
  - Motion sensors, cameras, and microphone
- Use a decision tree for driving behavior evaluation
  - Turn, lane change, and hard brake
- Take different factors into consideration for different activities
  - E.g., five factors for lane change, they are front and rear vehicle distance, turn signal usage, head pose and maneuver duration

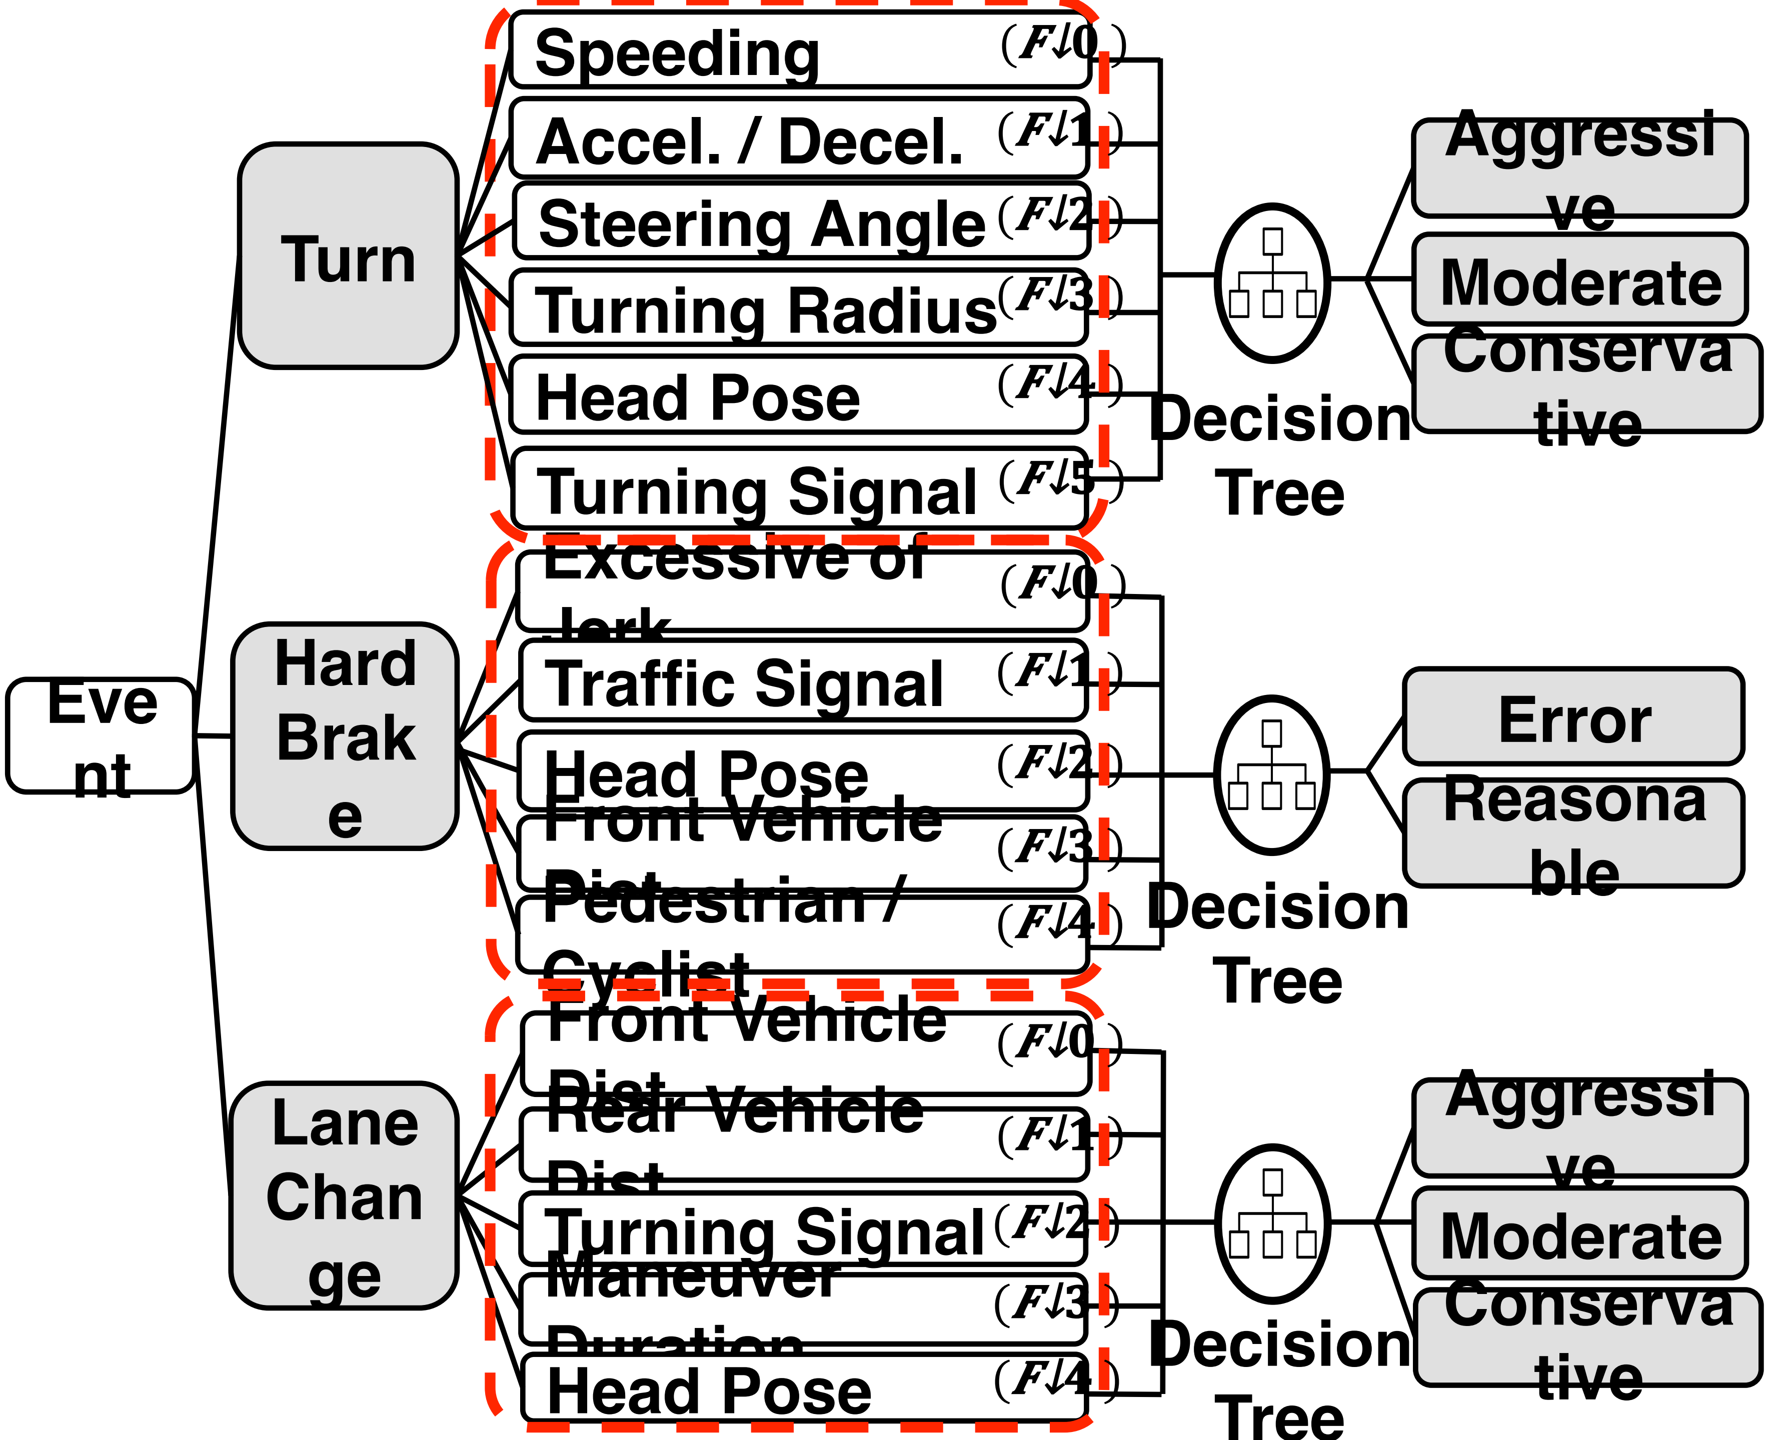
Event
Turn
Speeding
Accel. / Decel.
Steering Angle
Turning Radius
Head Pose
Turning Signal
Decision Tree
Aggressive
Moderate
Conservative
Hard Brake
Excessive of Jerk
Traffic Signal
Head Pose
Front Vehicle Dist.
Pedestrian / Cyclist
Decision Tree
Error
Reasonable
Lane Change
Front Vehicle Dist.
Rear Vehicle Dist.
Turning Signal
Maneuver Duration
Head Pose
Decision Tree
Aggressive
Moderate
Conservative

Demo

- A phone and three cameras are placed in the vehicle
- Turn, lane change and brake are monitored in real time

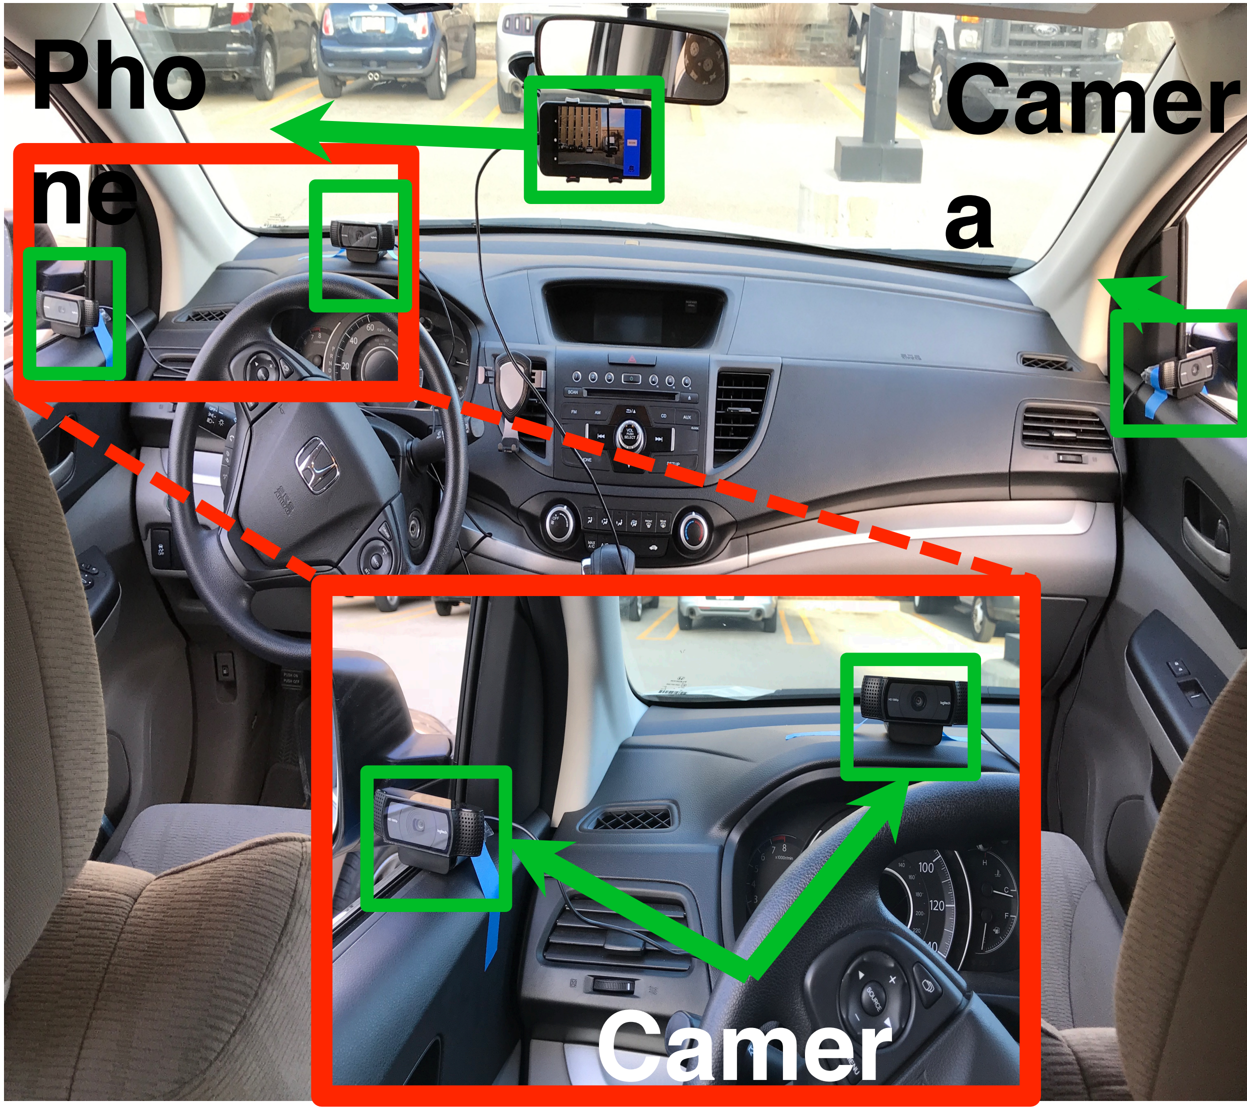

Right Turn

Steps for making a right turn

1	Check Right Wing Mirror	Using camera to track driver's head poses
2	Check Right Blind Spot	Using camera to track driver's head poses
3	Turn on Turn Signal	Using microphone to track turn signal usage
4	Make a Right Turn	Using gyroscope to identify right turn

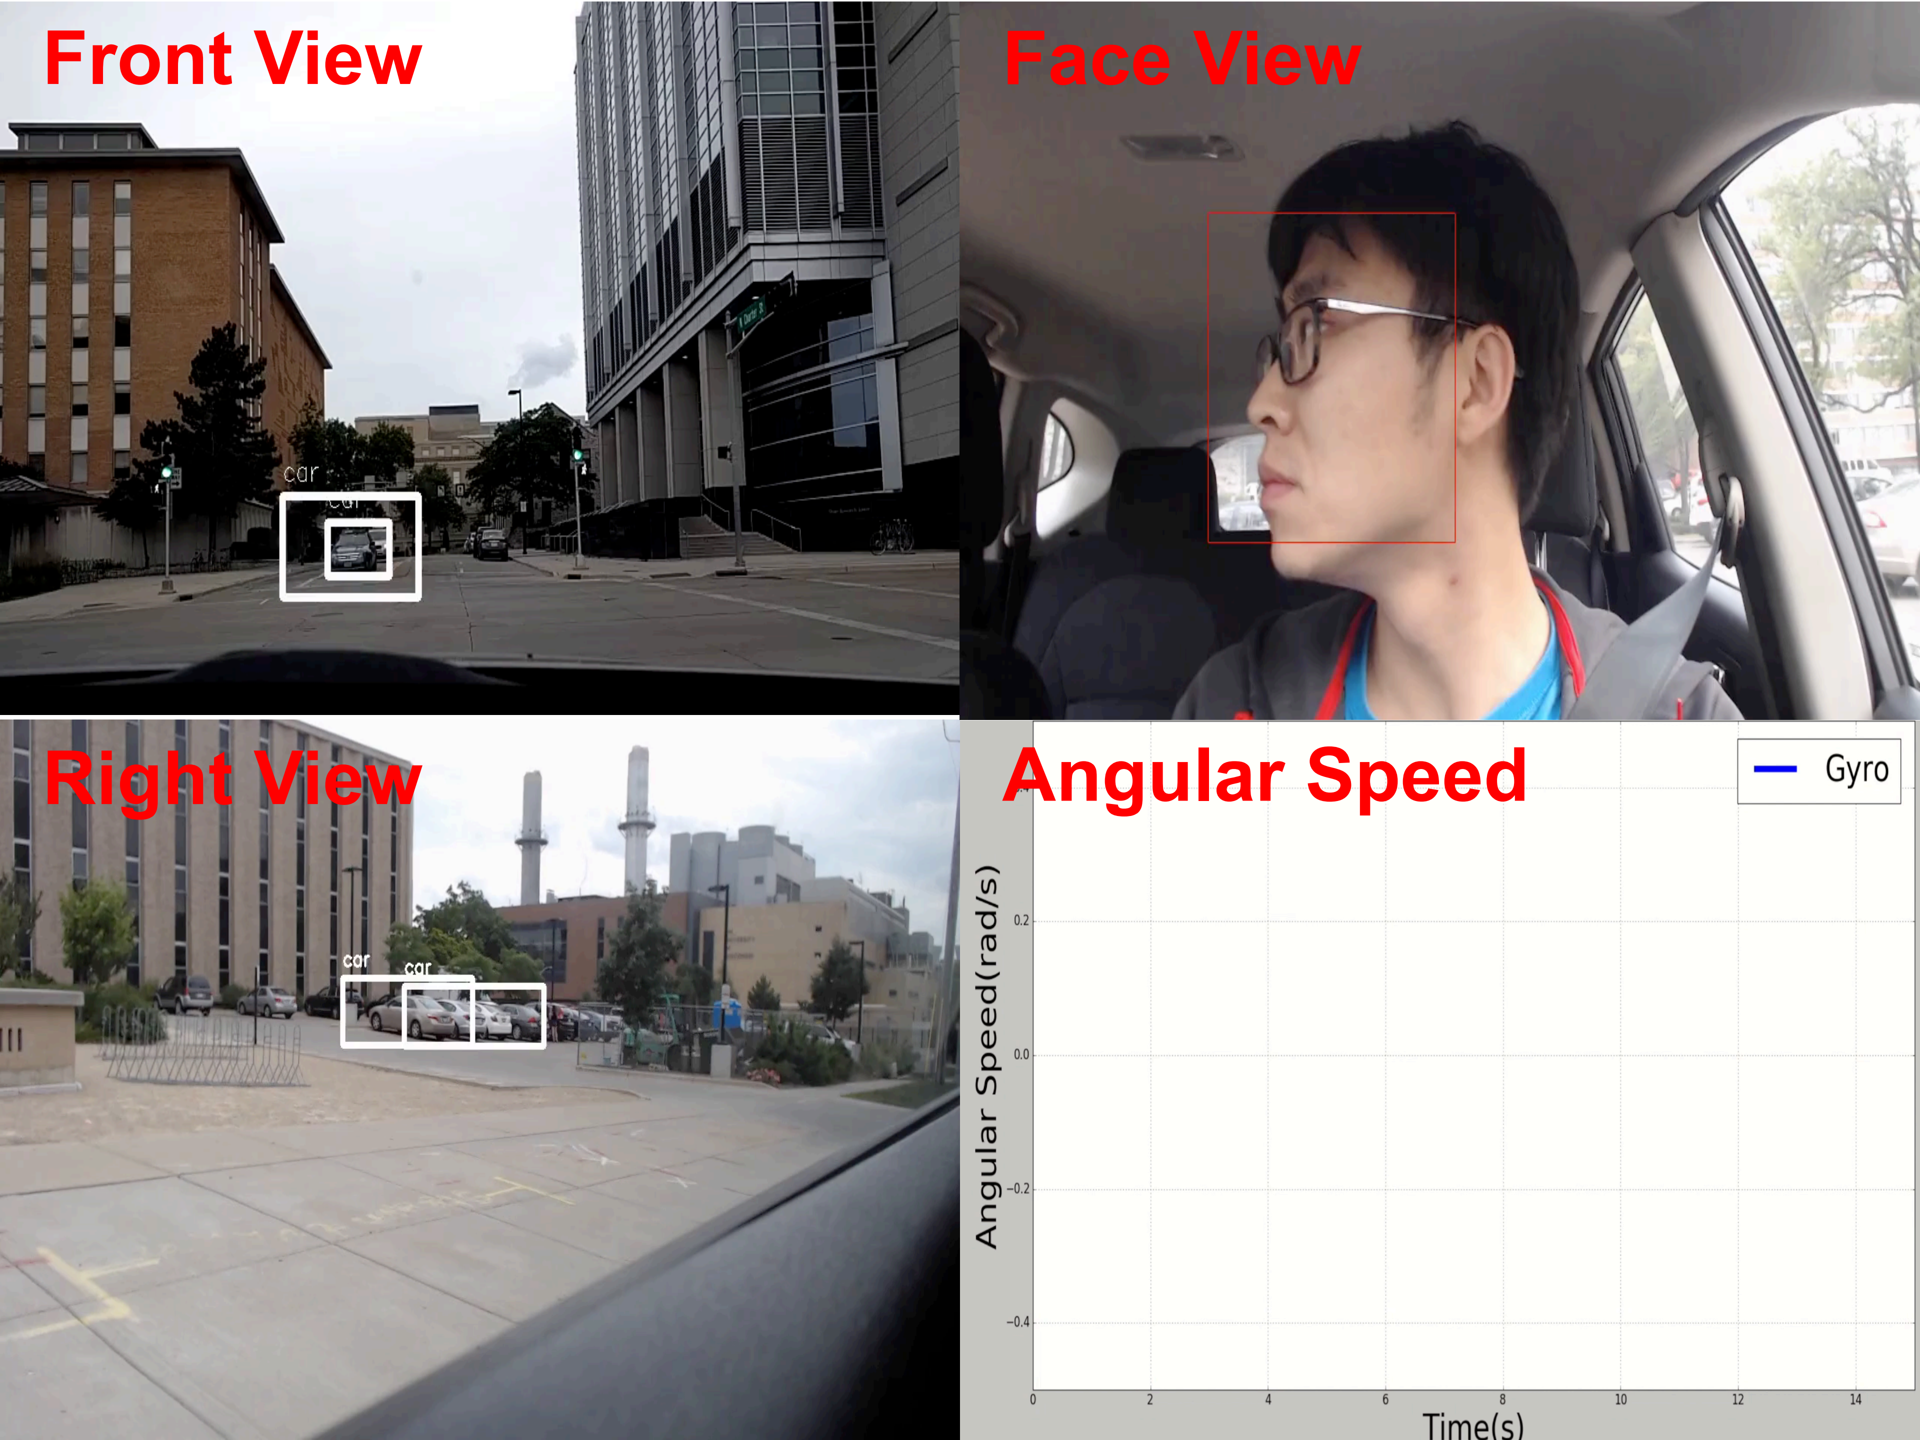
Front View
car
car
Face View
Right View
car
car
Angular Speed
Gyro
Angular Speed(rad/s)
0.2
0.0
−0.2
−0.4
0
2
4
6
8
10
12
14
Time(s)

Many more apps …

- Kids router
- Framily router

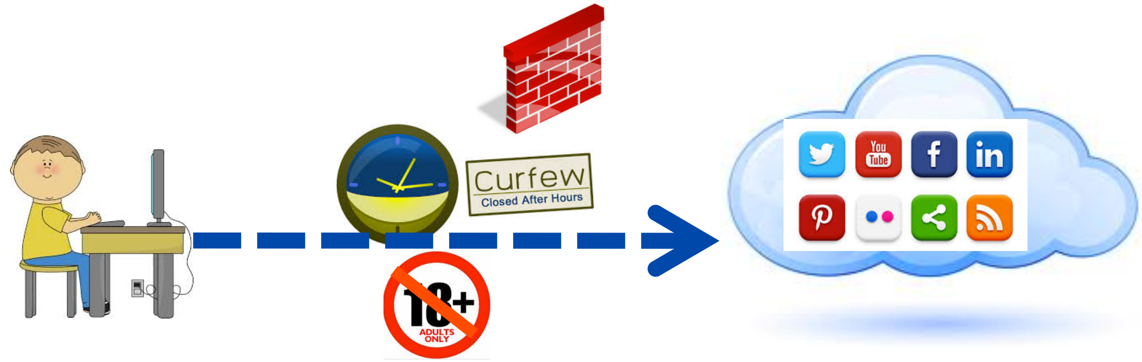

- TOR in your router
- Monitor your ISP (a la BisMark)
- As a wireless and networking educational aid

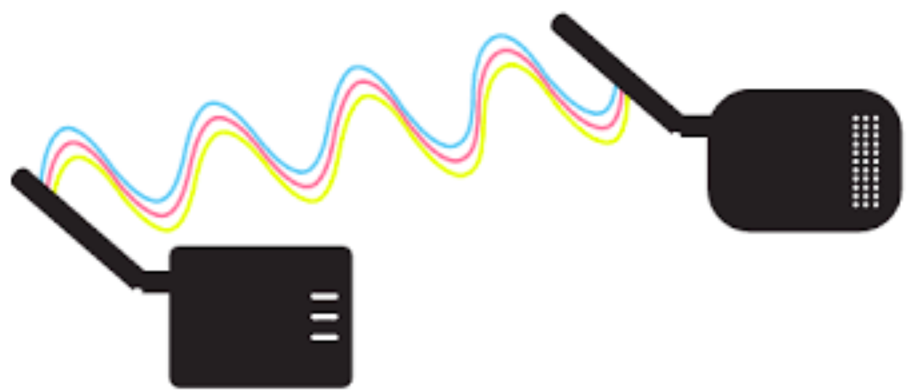

What could you do?

- Your idea here

Water quality management

- Water hardness

- Water softeners --- Protect the expensive water heater from calcification

Water quality management

Water hardness

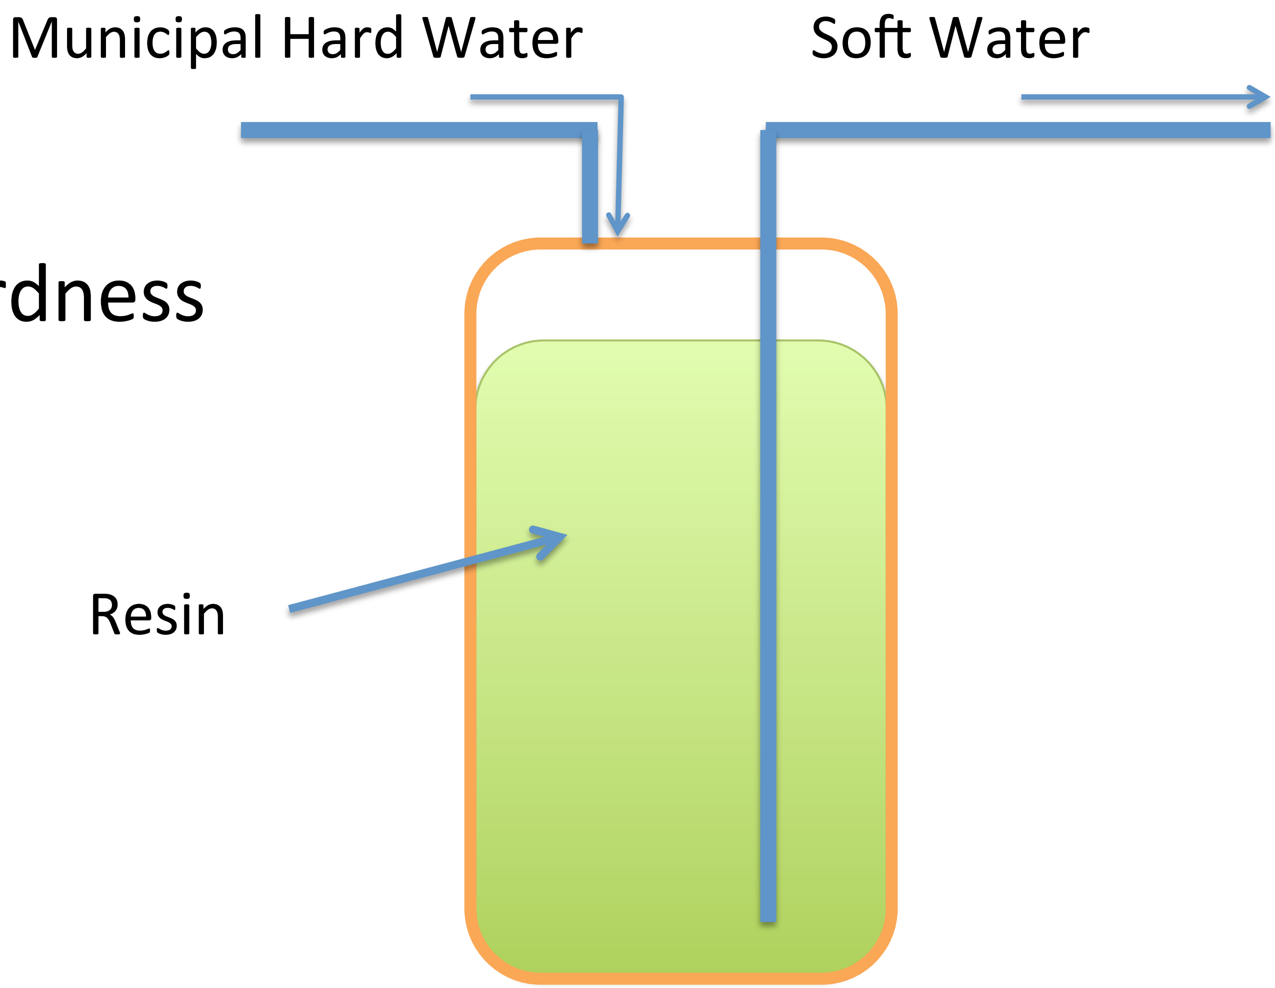

Water hardness in the US

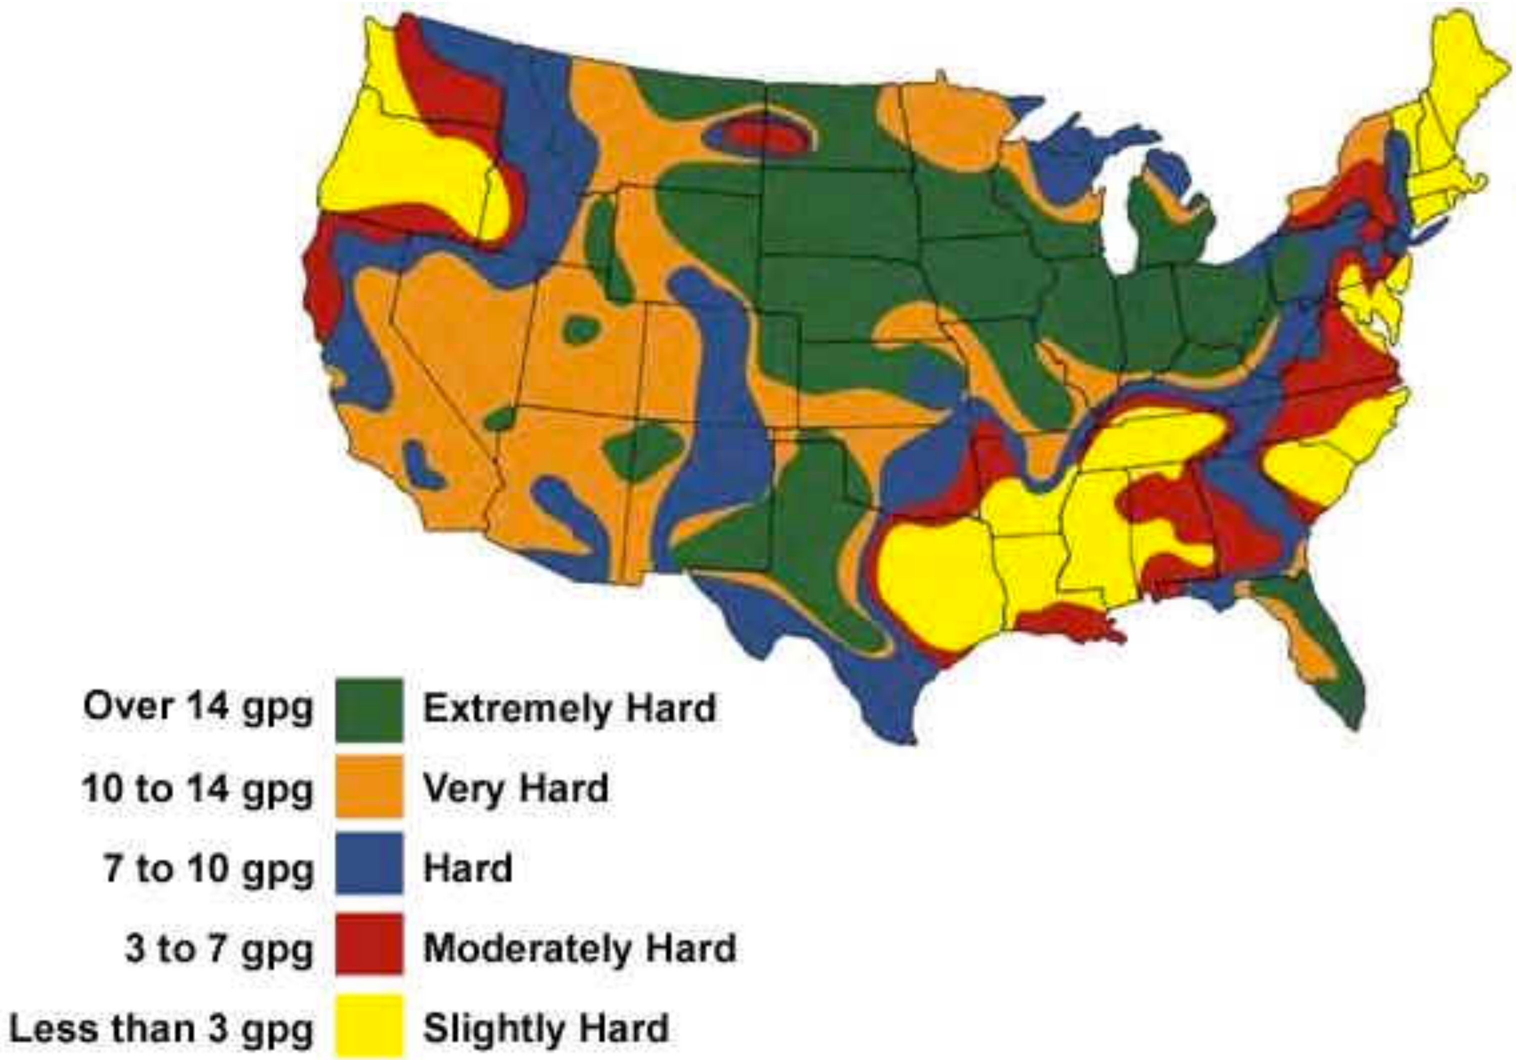

How Water Softeners Work
SOFTENING CYCLE

Water Softener

Hard water enters

Many gallons later

Capacity to soften water now ending

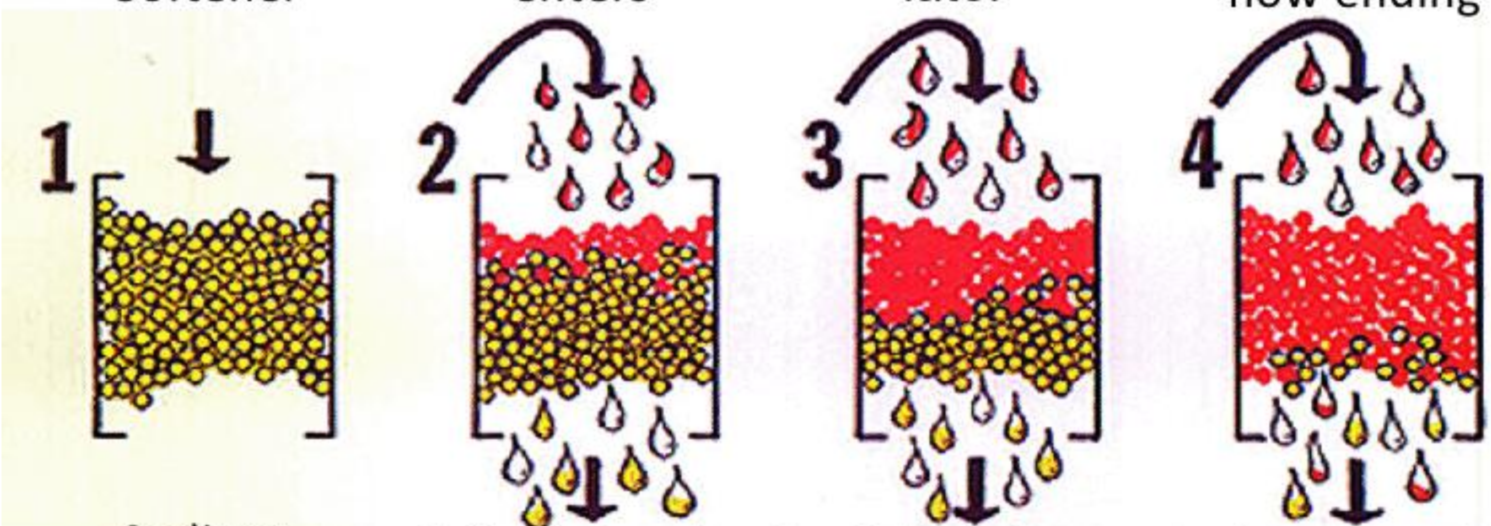

Sodium charged ion exchange resin

Softened water containing sodium ions

Partially exhausted, but still delivering 100% soft water

Regeneration with salt brine now required

Lime in a
Water Heater

But here's the thing

- We don't actually know if the water coming out is soft (unless you test)
- So, to be safe, people typically regenerate 10-30% more frequently than they should
- Regenerating too frequently wastes salt

100K lbs salt pass through Madison sewers each day

A WatEr Softener Online OptiMization Engine

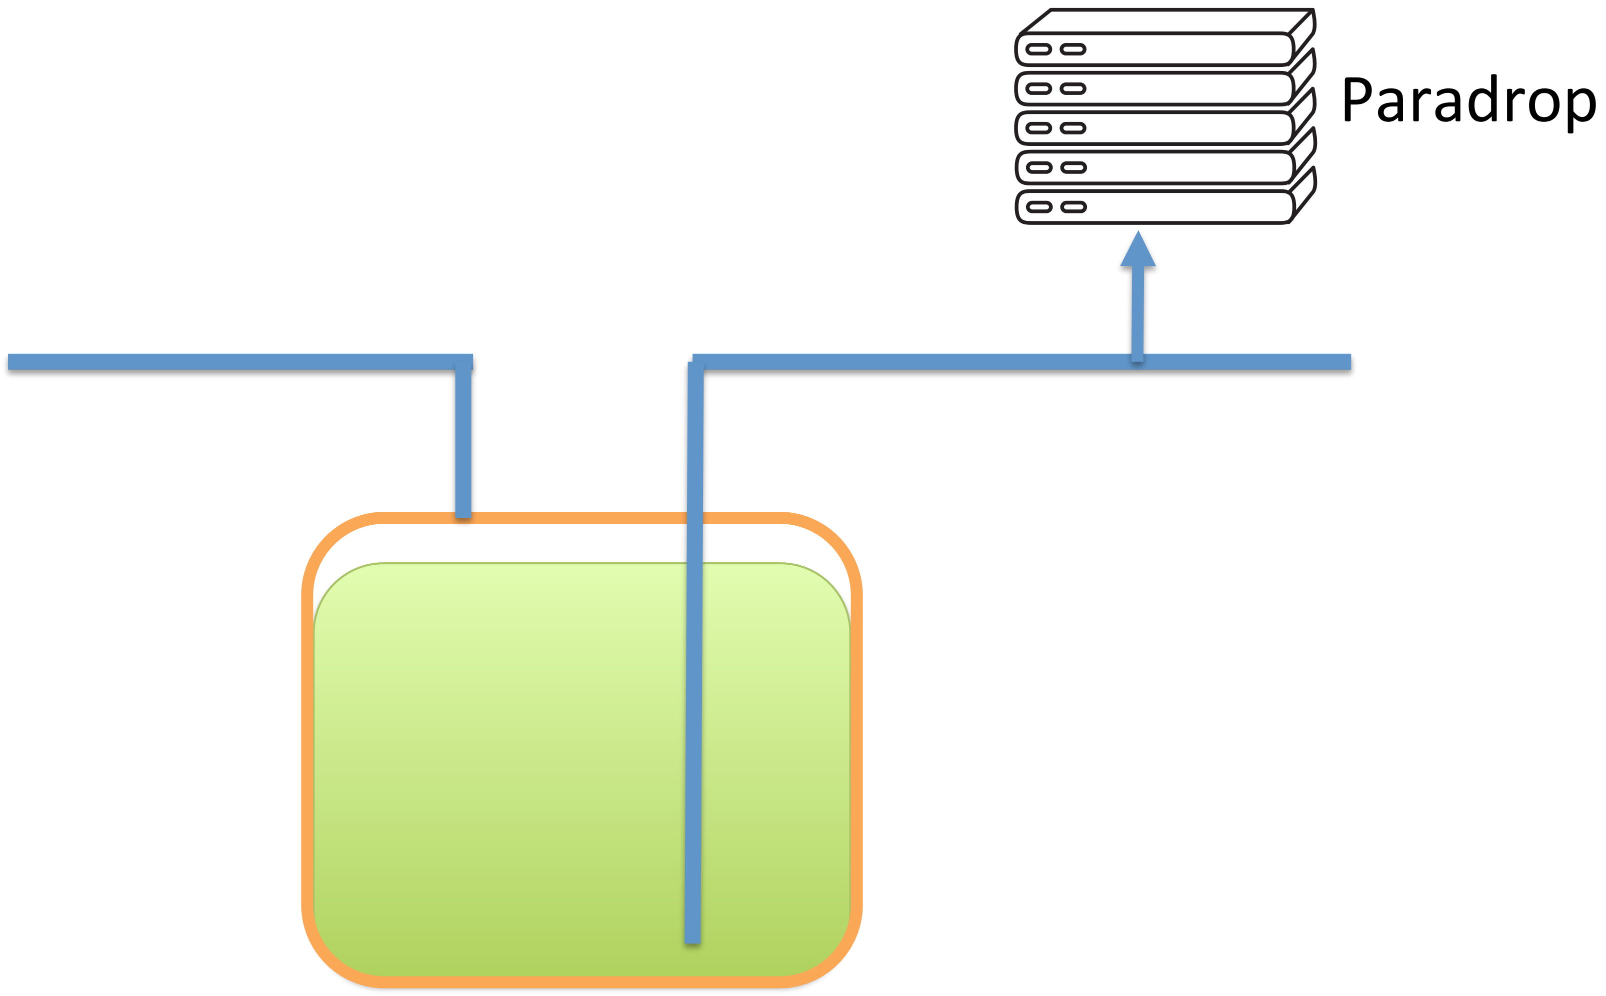

AWESOME

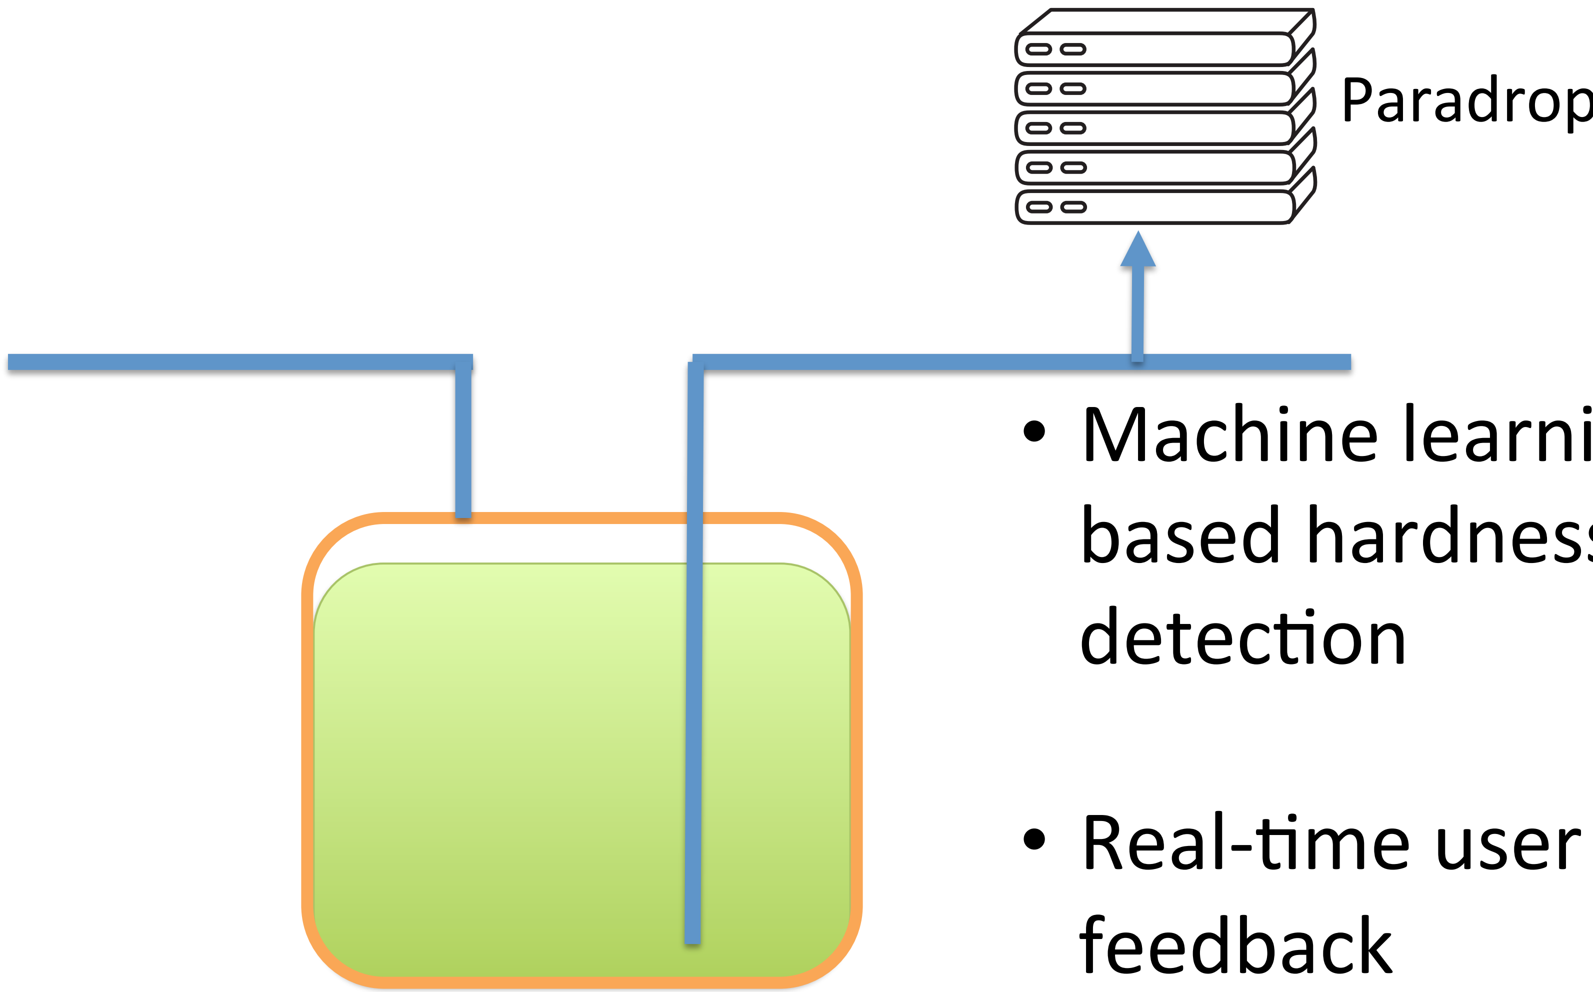

- Machine learning based hardness detection
- Real-time user feedback

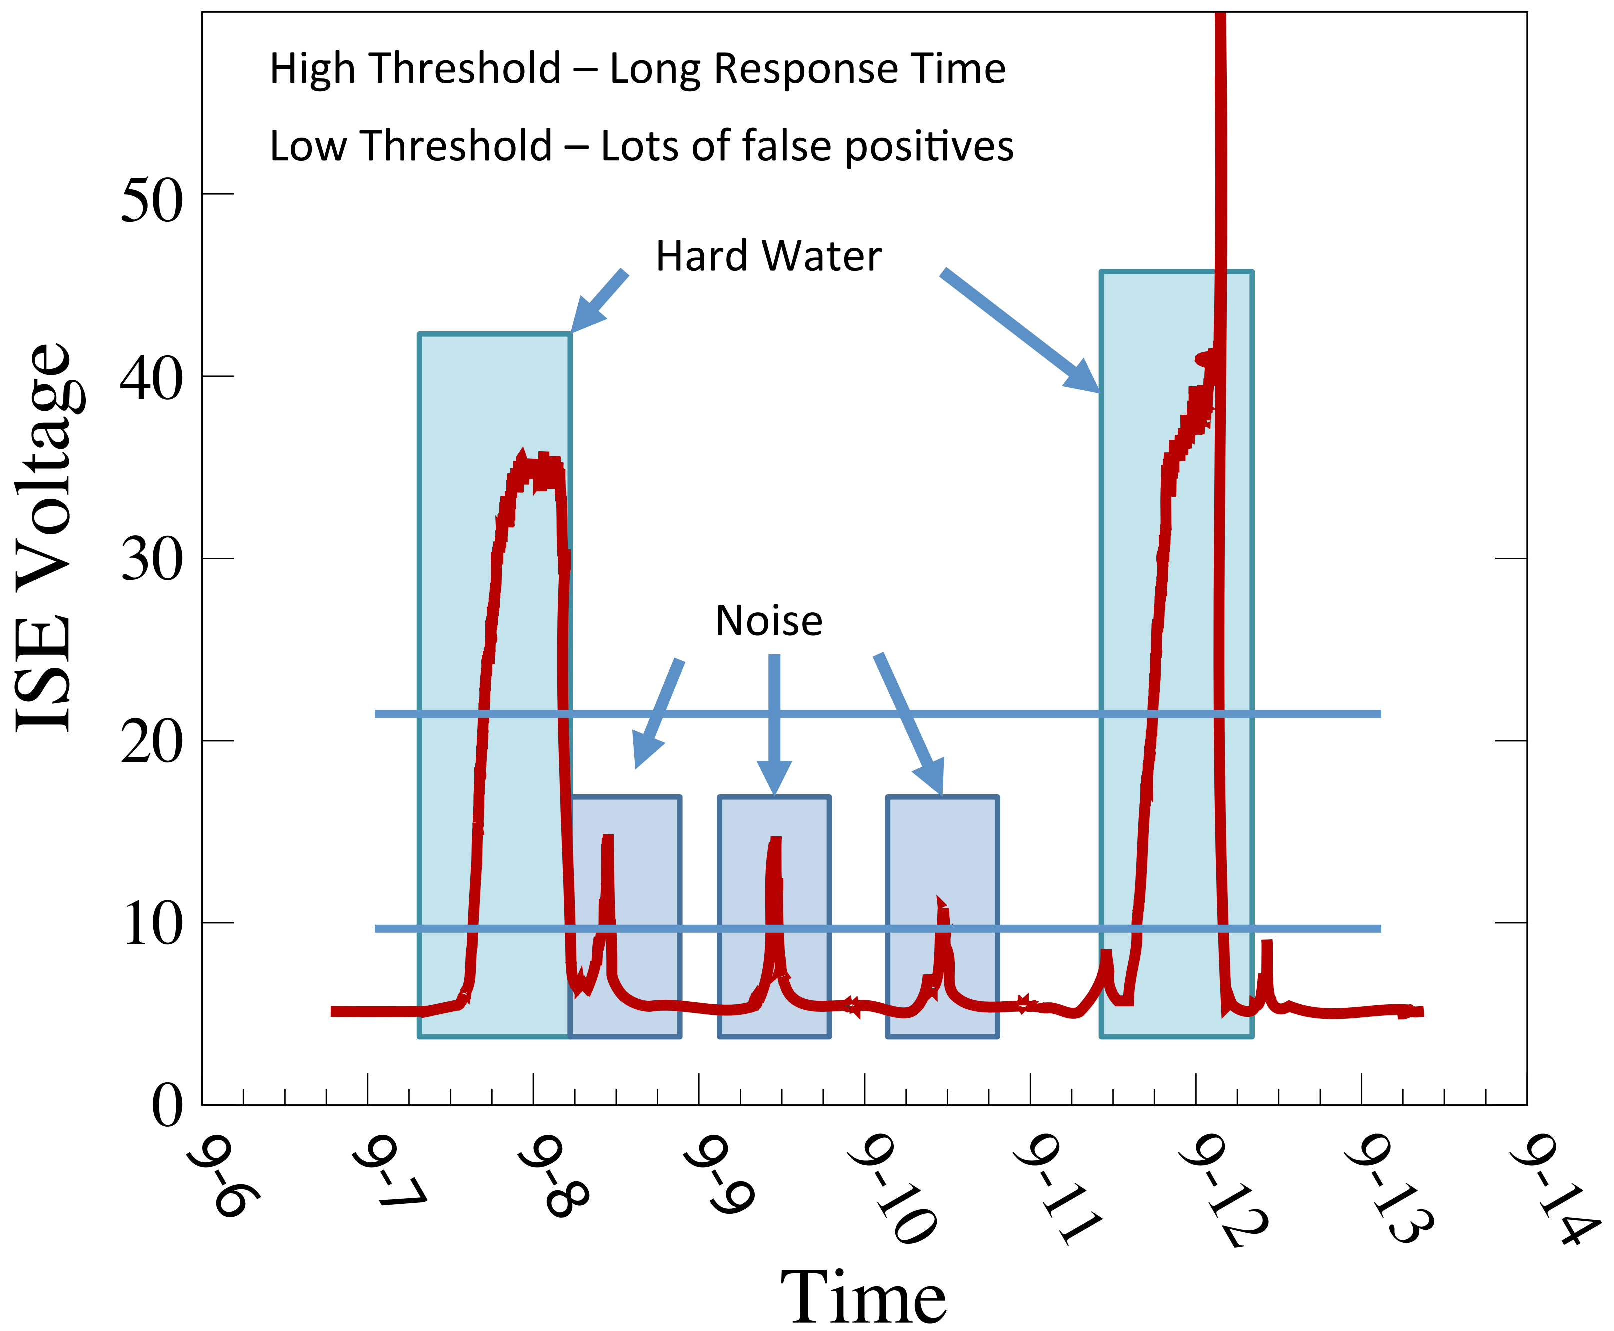

High Threshold – Long Response Time
Low Threshold – Lots of false positives
Hard Water
Noise
ISE Voltage
0
10
20
30
40
50
9-6
9-7
9-8
9-9
9-10
9-11
9-12
9-13
9-14
Time

Learning algorithm

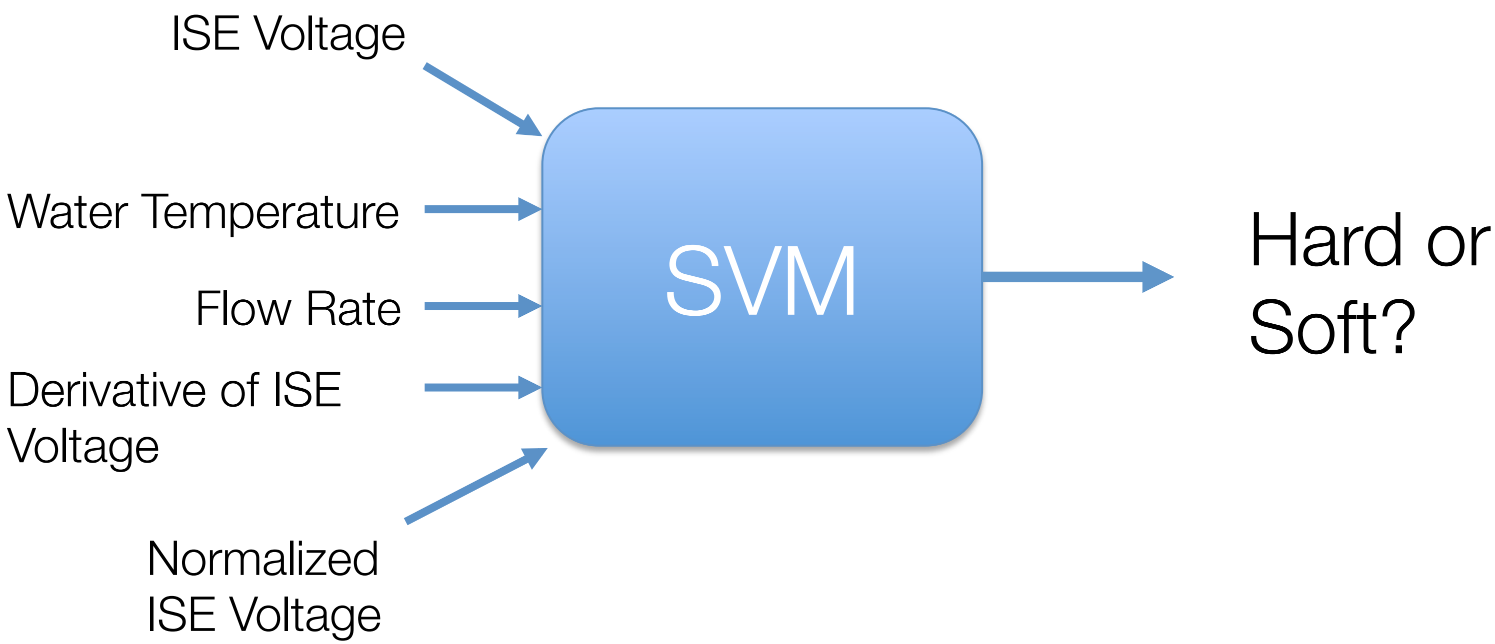

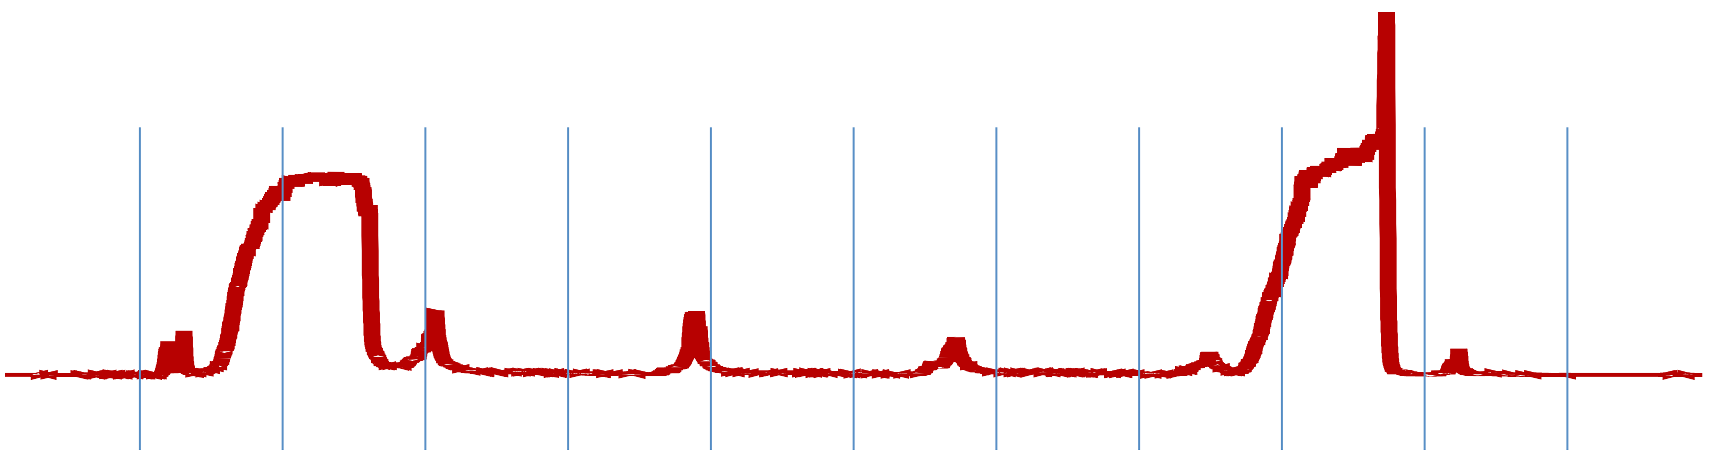

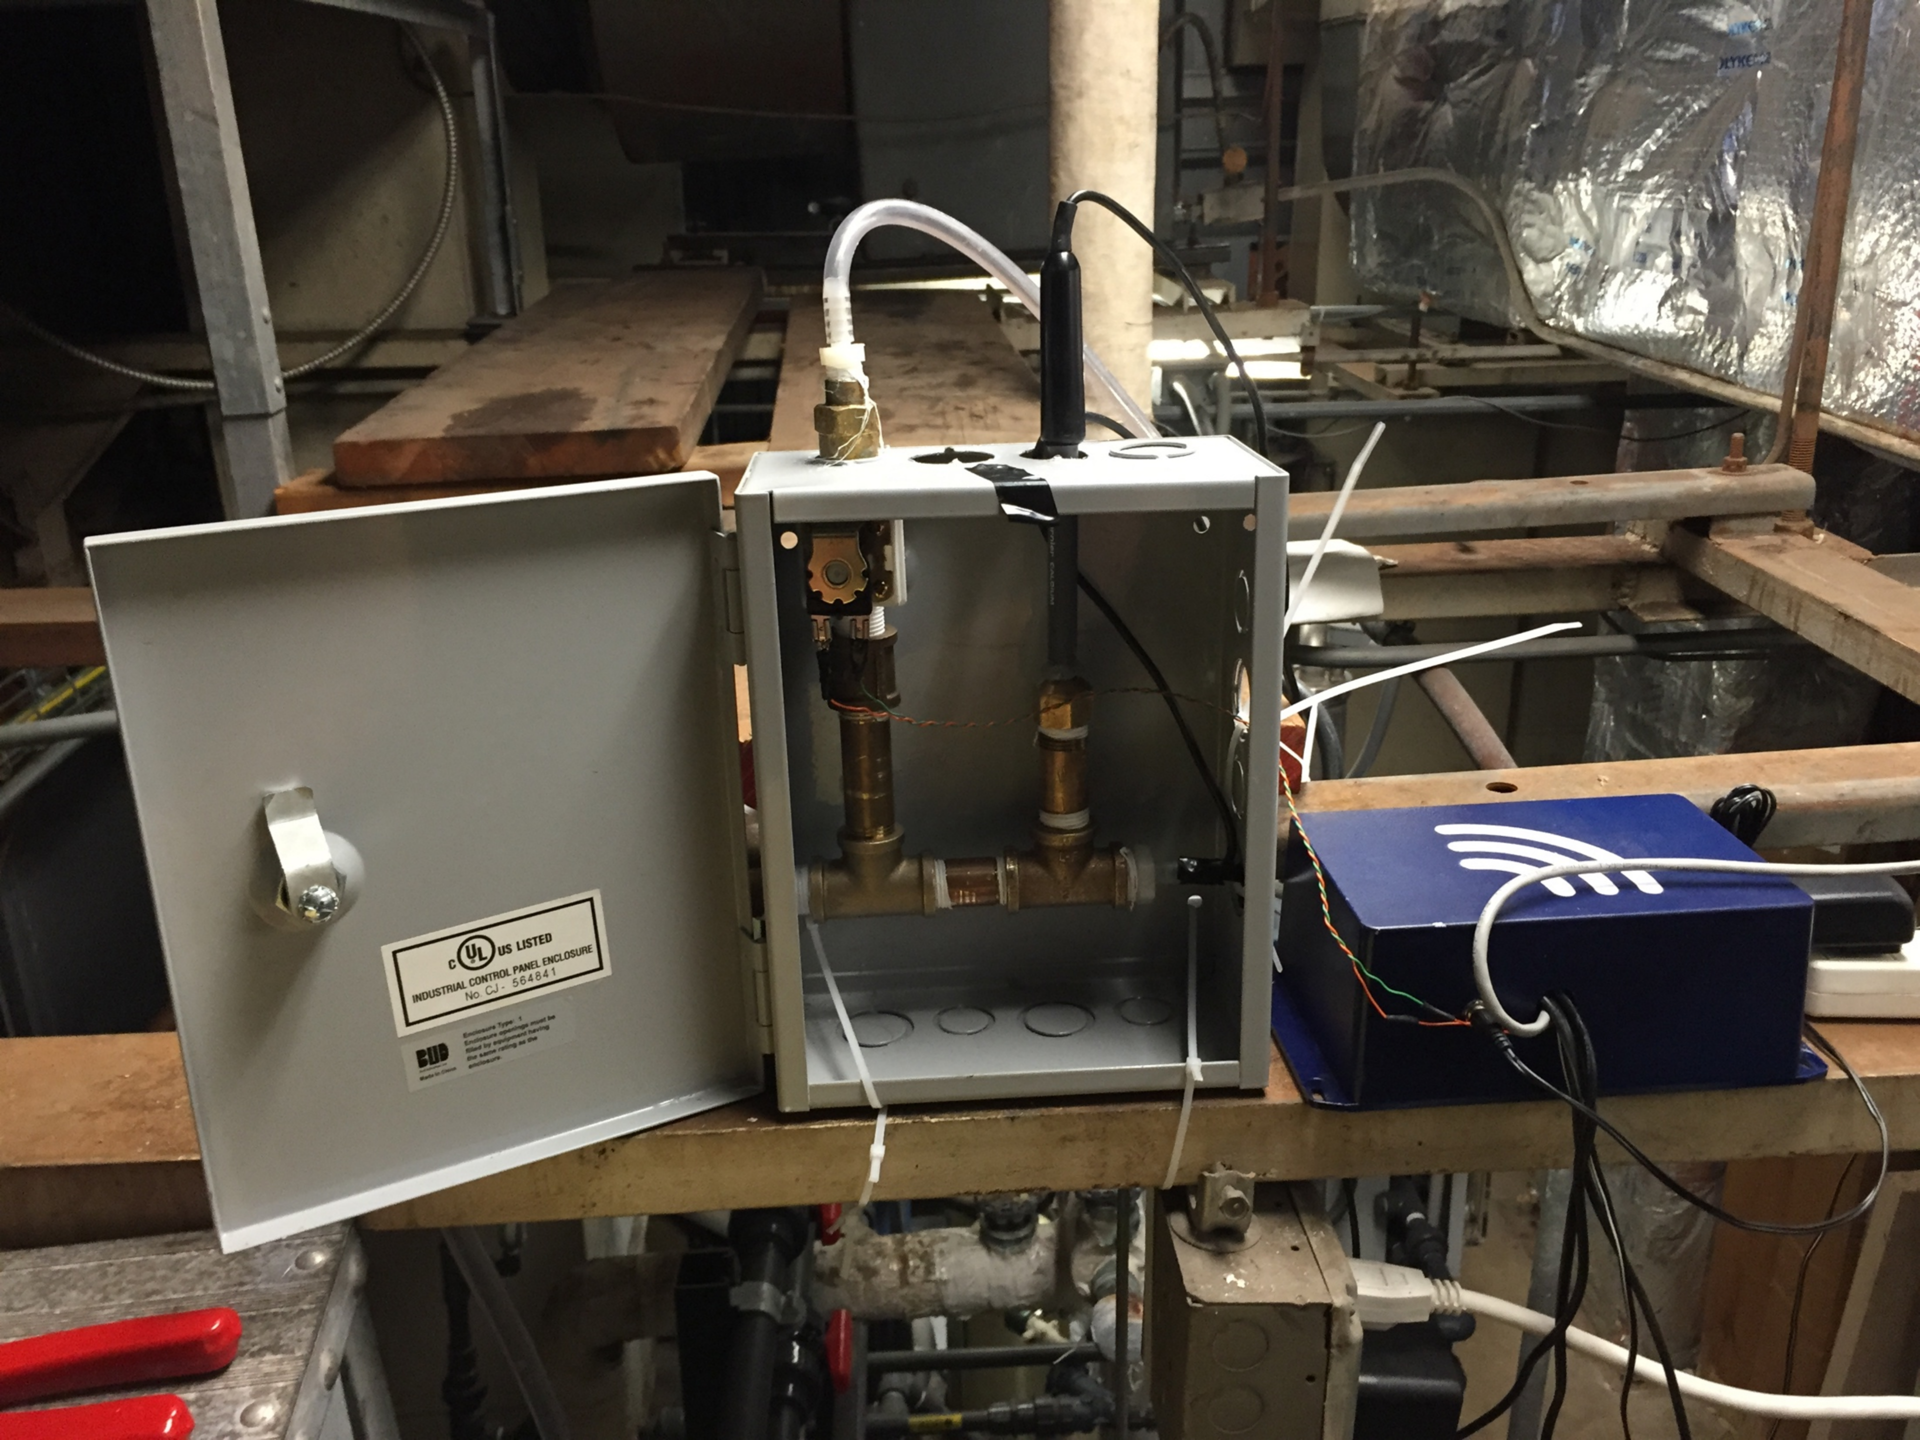
c UL US LISTED
INDUSTRIAL CONTROL PANEL ENCLOSURE
No. CJ- 564841
Enclosure Type: 1
Enclosure openings must be
filled by equipment having
the same rating as the
enclosure.
BUD
Made in China

More results

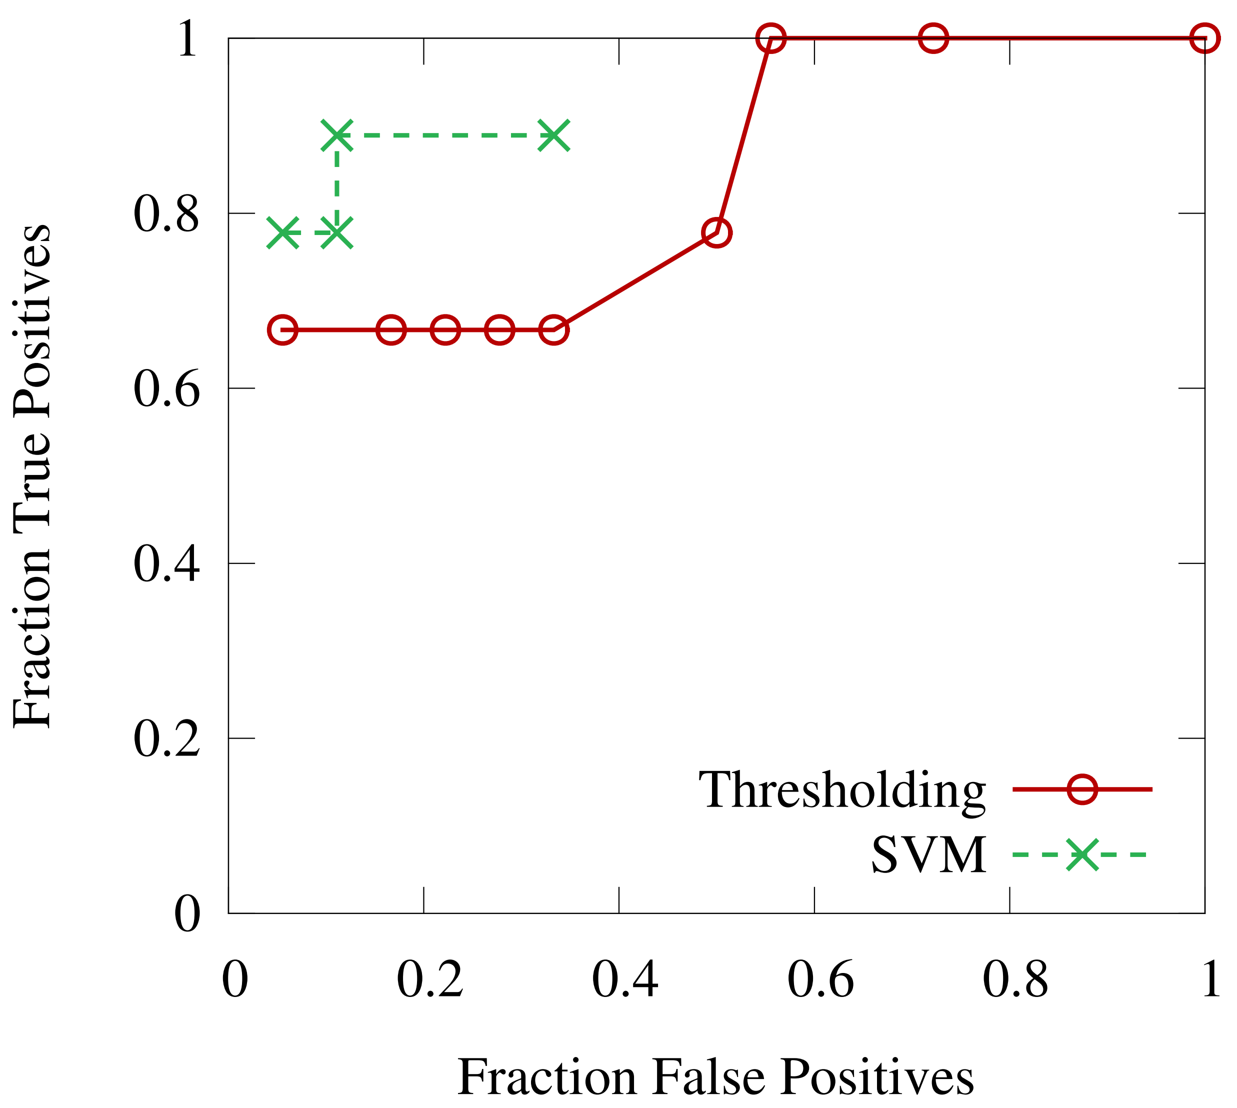

AWESOME

Multiple installs on UW Campus

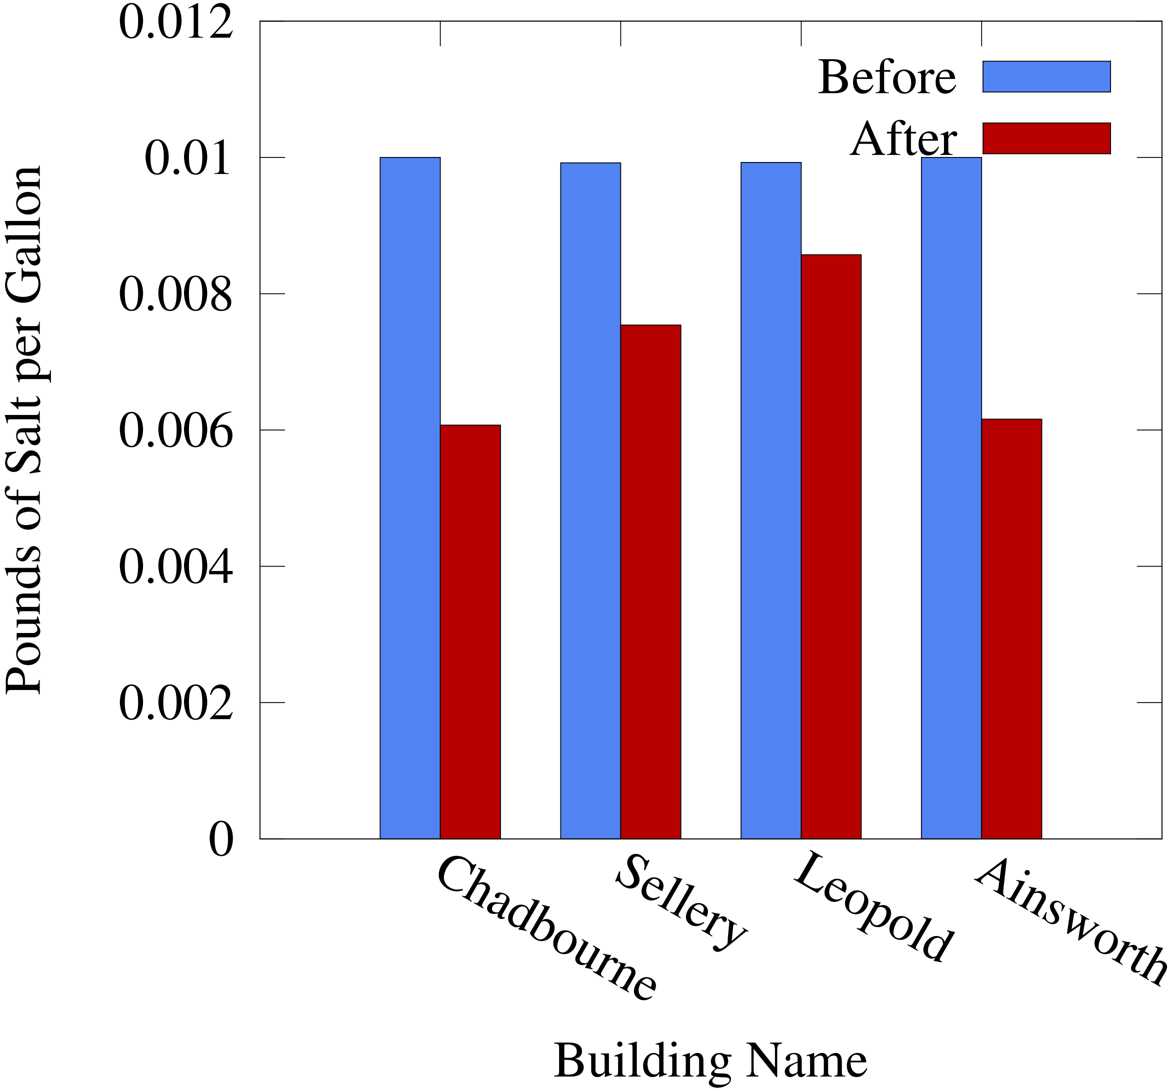

Average savings 27%

Money saved $1500-$3000 per year

RF management in ParaDrop

"Is my wireless network operating well enough?"

A “blood test” for the network

(Comment made by Pravin Bhagwat: CTO, Airtight Networks)

Basic question

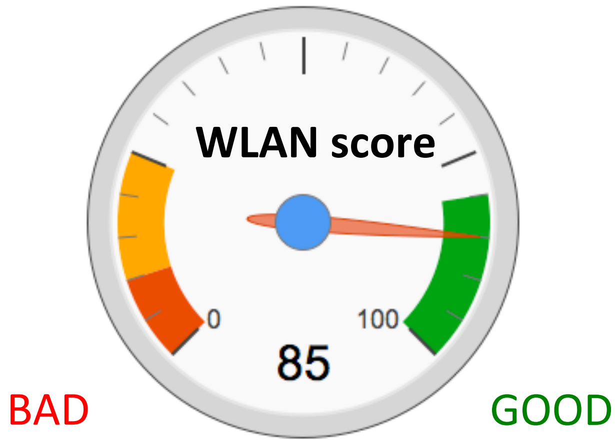

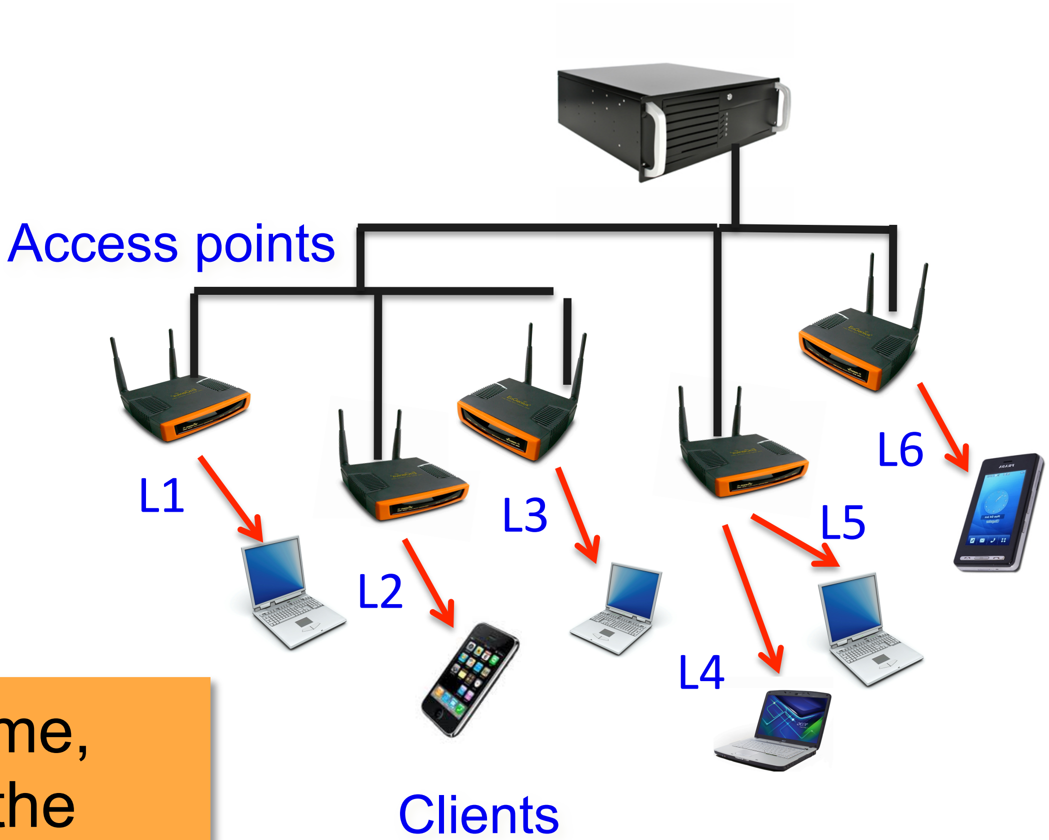

Can we provide a real-time, overall health score for the entire network?

Basic question

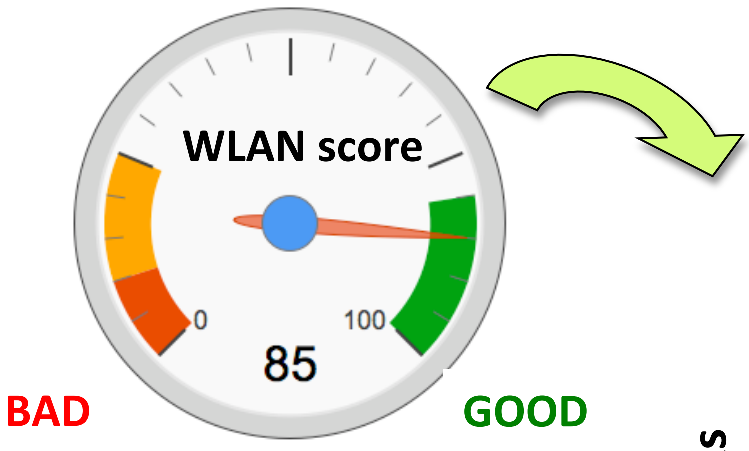

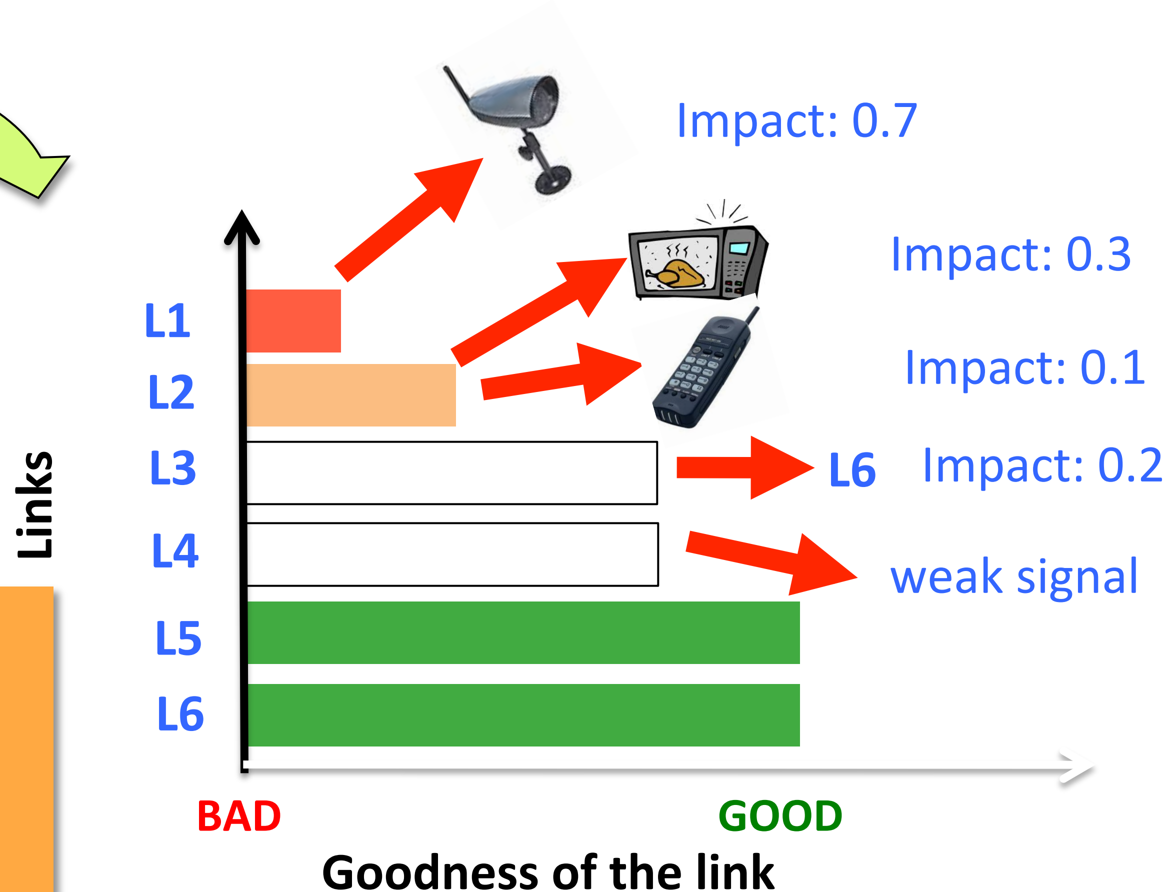

Identify why certain links are under-performing

Many solutions

Real-time Interference estimators

WiFi to WiFi Interference

- **COLLIE** (Collision Inferencing Engine)
  [Rayanchu et. al., Infocom 2007]

- **PIE** (Passive Interference Estimator)
  [Shrivastava et. al., NSDI 2011]

Non-WiFi to WiFi Interference

(use WiFi-only hardware!)

- **Airshark** [Rayanchu et. al., IMC 2011]
  1. detect non-WiFi devices

- **WiFiNet**
  1. quantify interference impact
  2. pin point device location

  [Rayanchu et. al., NSDI 2012]

Detecting non-WiFi activity

- Using software-only mechanisms
  - Benefits in interference debugging

Spectrum at a university cafe

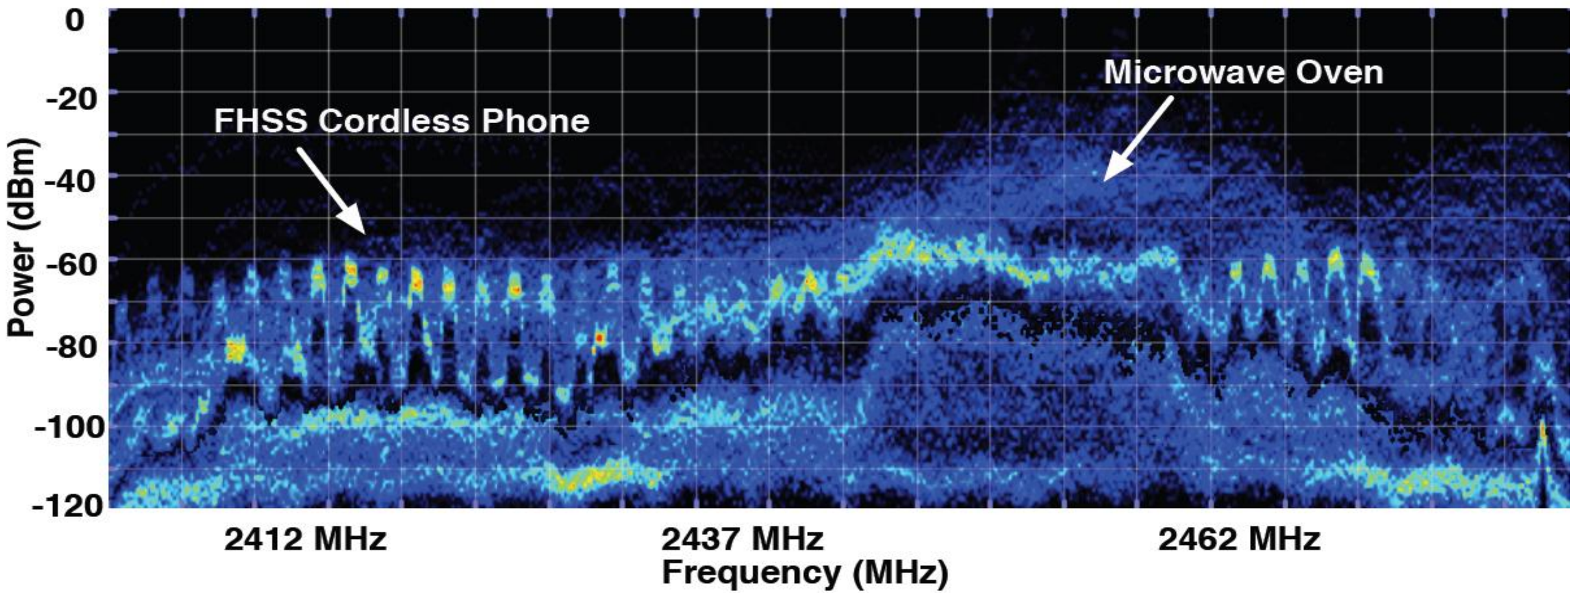

High powered non-WiFi devices share the spectrum with WiFi devices

Use a coarse-grained WiFi lens

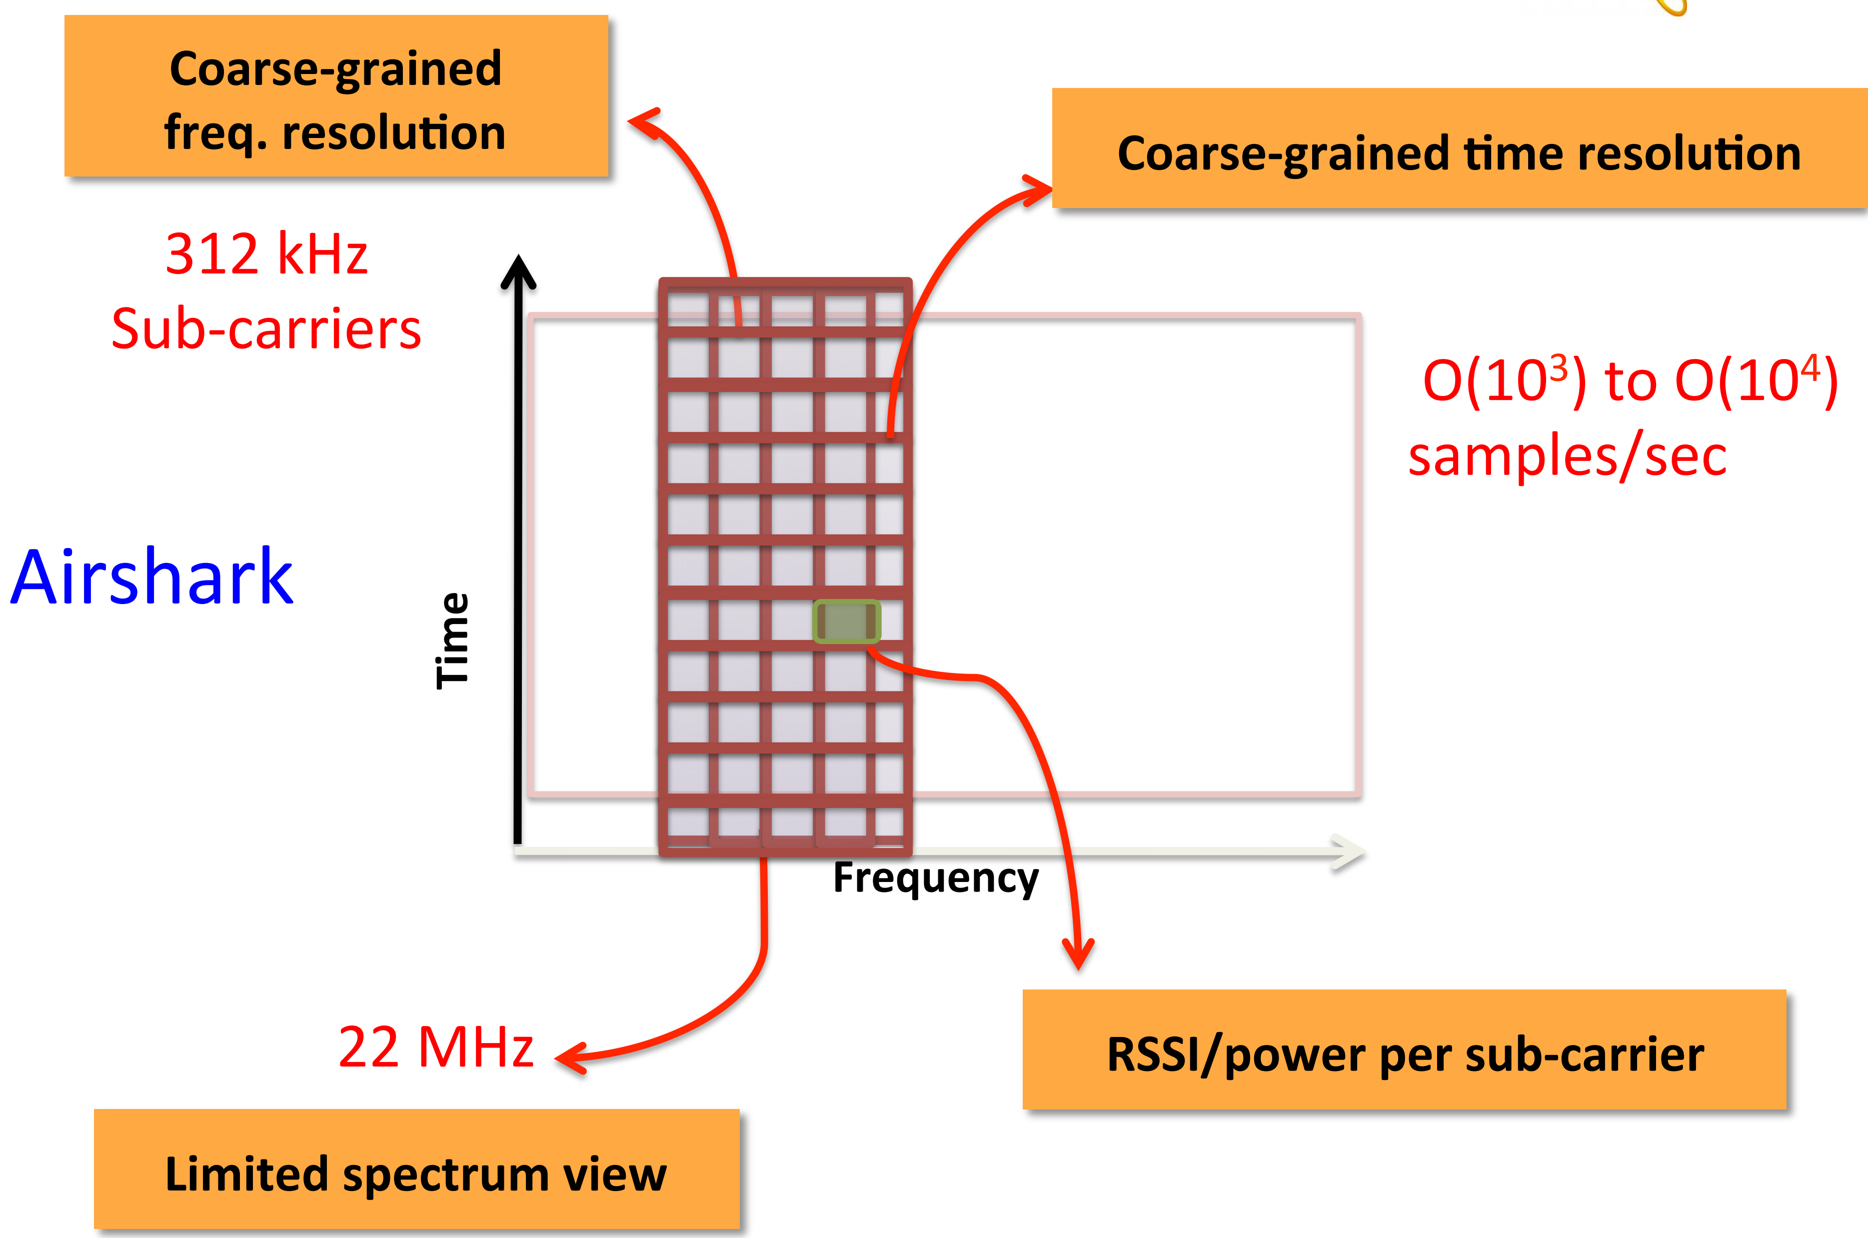

Airshark: *how it works*

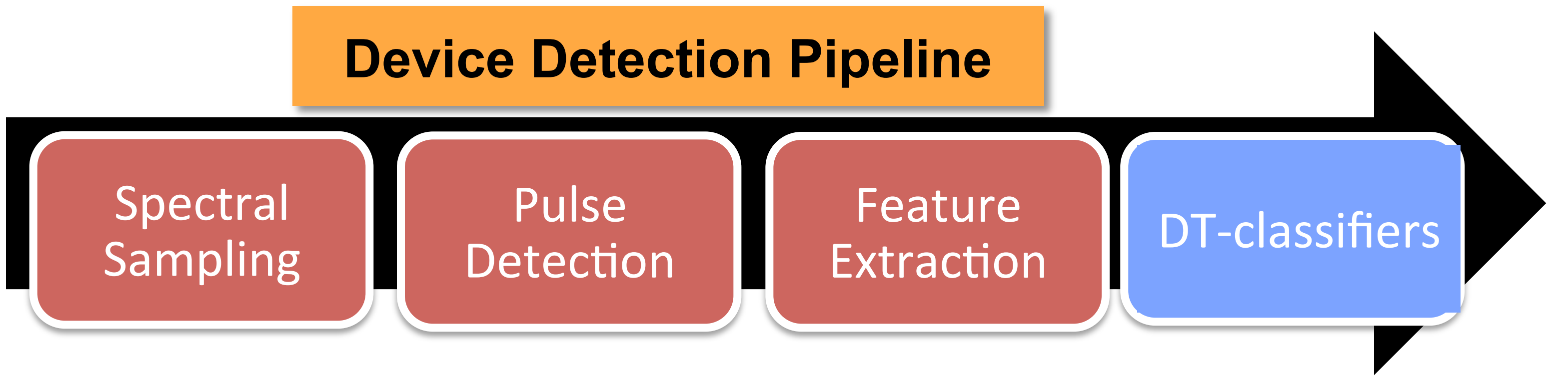

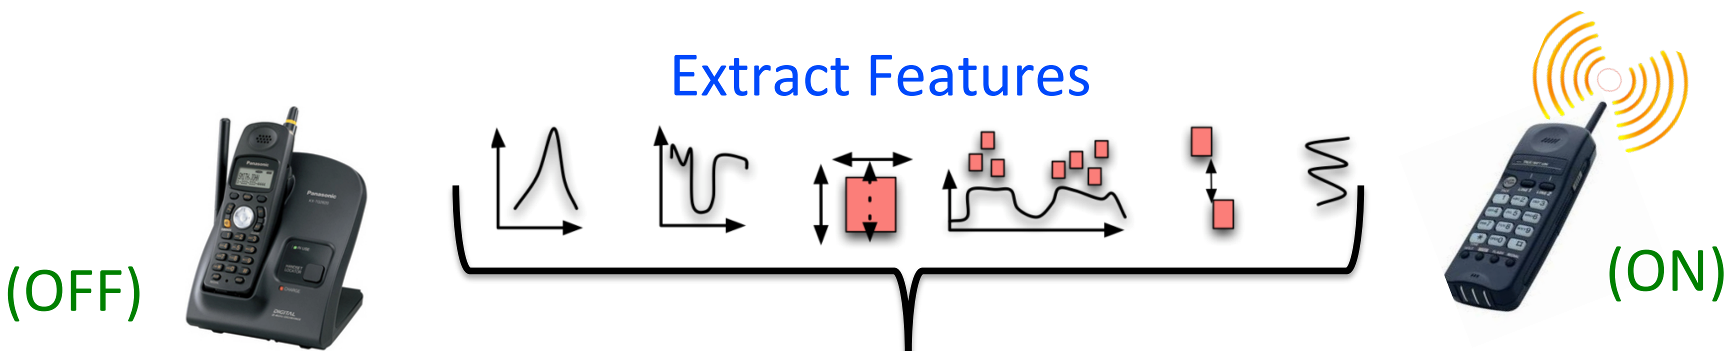

Training / Classification

Decision Tree Classifier

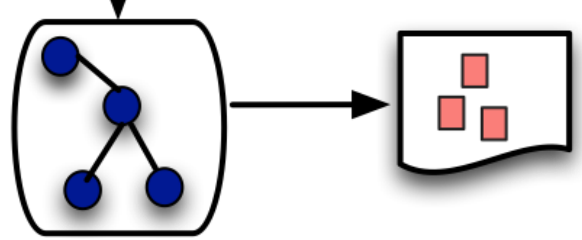

Detection Accuracy

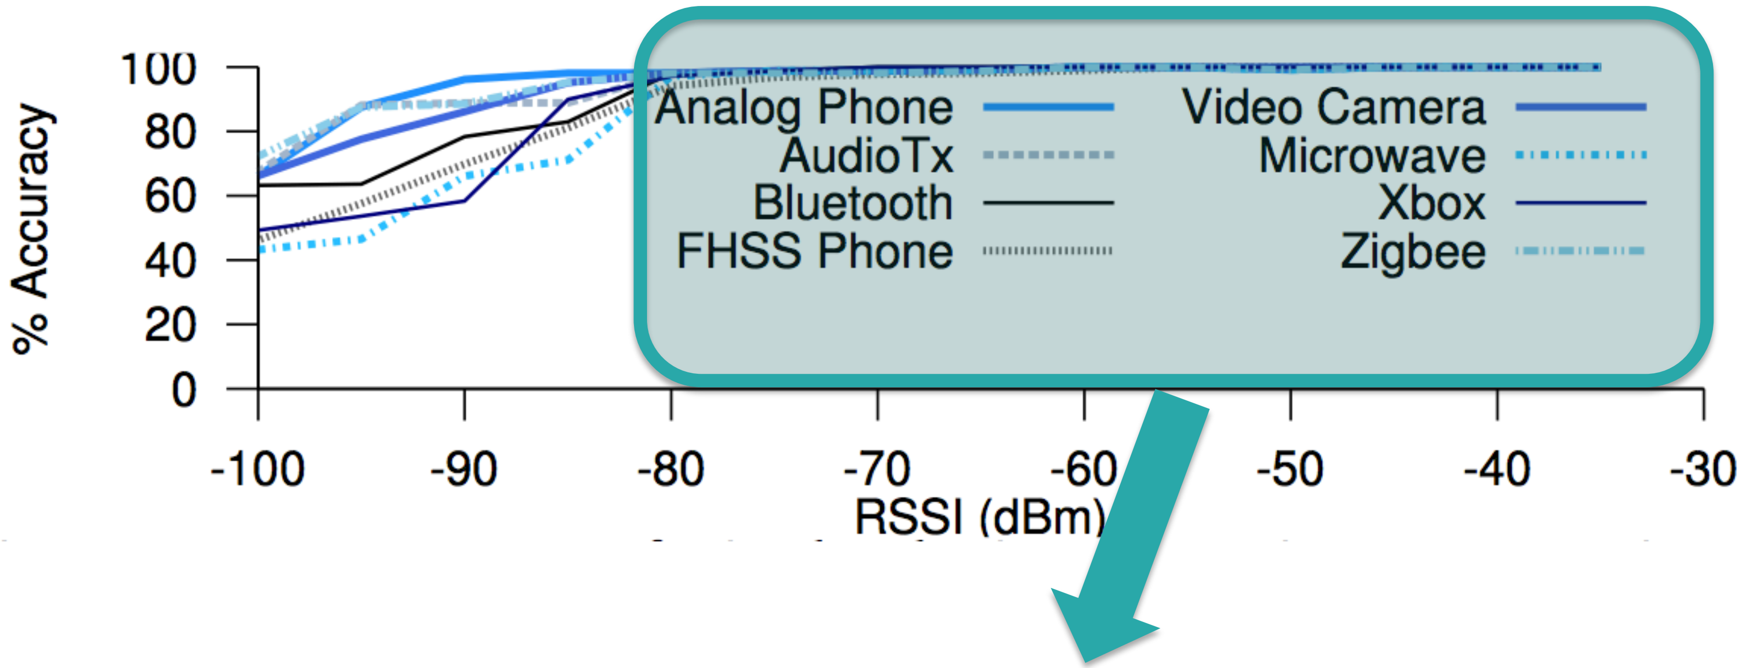

> 98% accuracy at signal strengths >= -80 dBm

WiFiNet: Interference estimation concept

- Quantify the "impact" of each device on each WiFi link
  - Identify the transmission overlaps between WiFi frames and non-WiFi pulses
  - Correlate frame losses and transmission overlaps

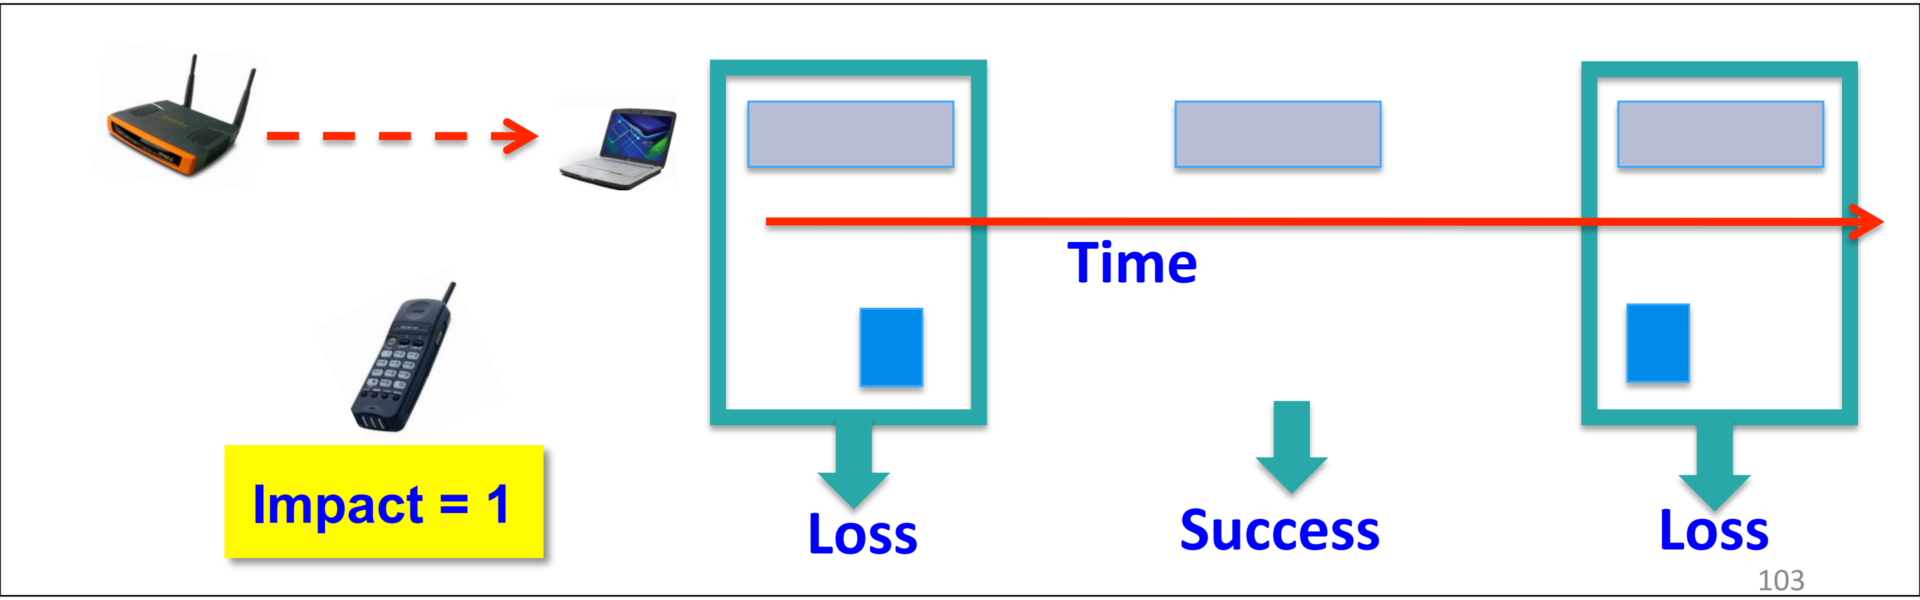

Cloud-based RF Management

COAP: Coordination framework for Open Access Points

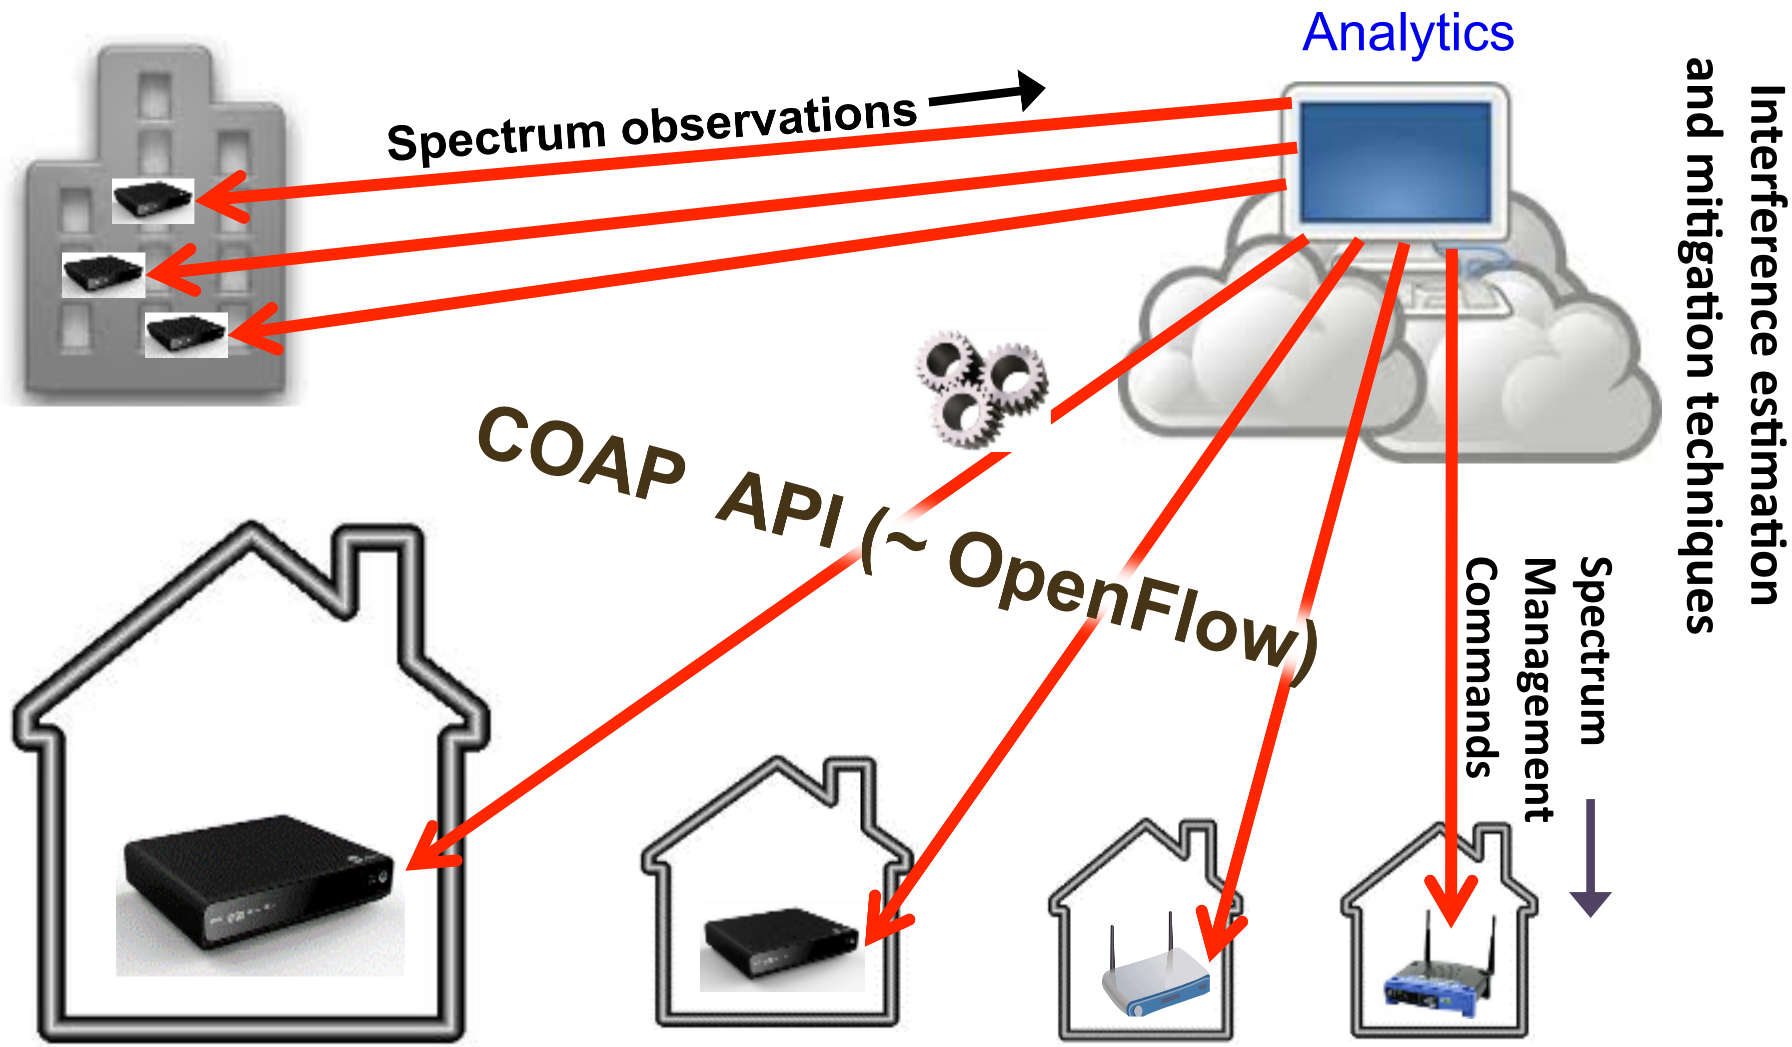

Field Trials

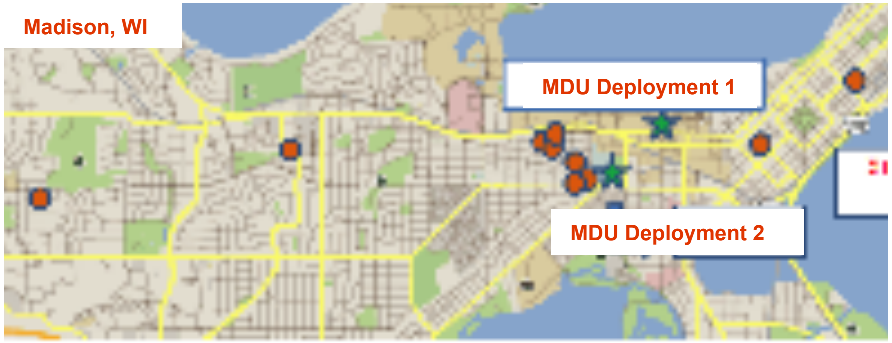

- OpenWRT based APs
  - ALIX 2d2 platform: (500 MHz AMD Geocode CPU, 256 DDR RAM, flash storage)
- 30+ APs deployed in homes & apartment complexes for 3+ years
- Cloud controller hosted in off-the-shelf Linux server

Understanding poor performance

Indicators				Bldg 1		Bldg 2	
A↑	S↓	L↑	R↓	V Poor	Poor	V Poor	Poor
Y	X	X	X	0	18.4	0	1
X	X	Y	X	24.2	**49.5**	25.2	**78.1**
Y	X	Y	X	**61.8**	26.7	2.1	1.4
X	Y	Y	X	2.3	1.1	20	15.8
X	Y	Y	Y	9.4	0	**51.6**	3.4
Others				2.3	4.3	1.1	1.3

Contention experience

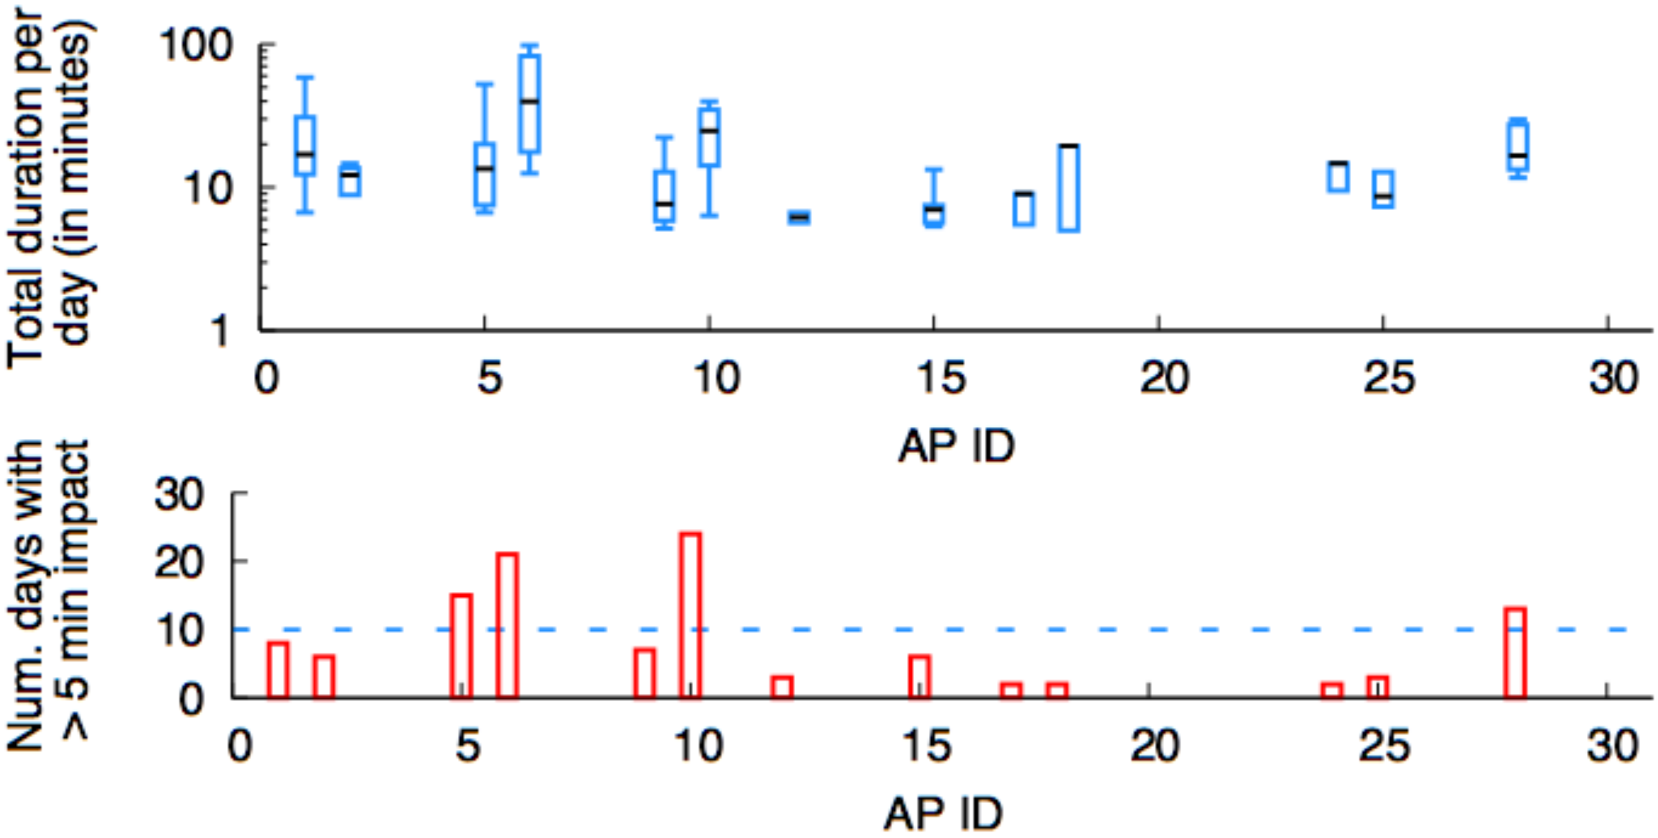

(In a 40 day period)

Microware cuts throughput

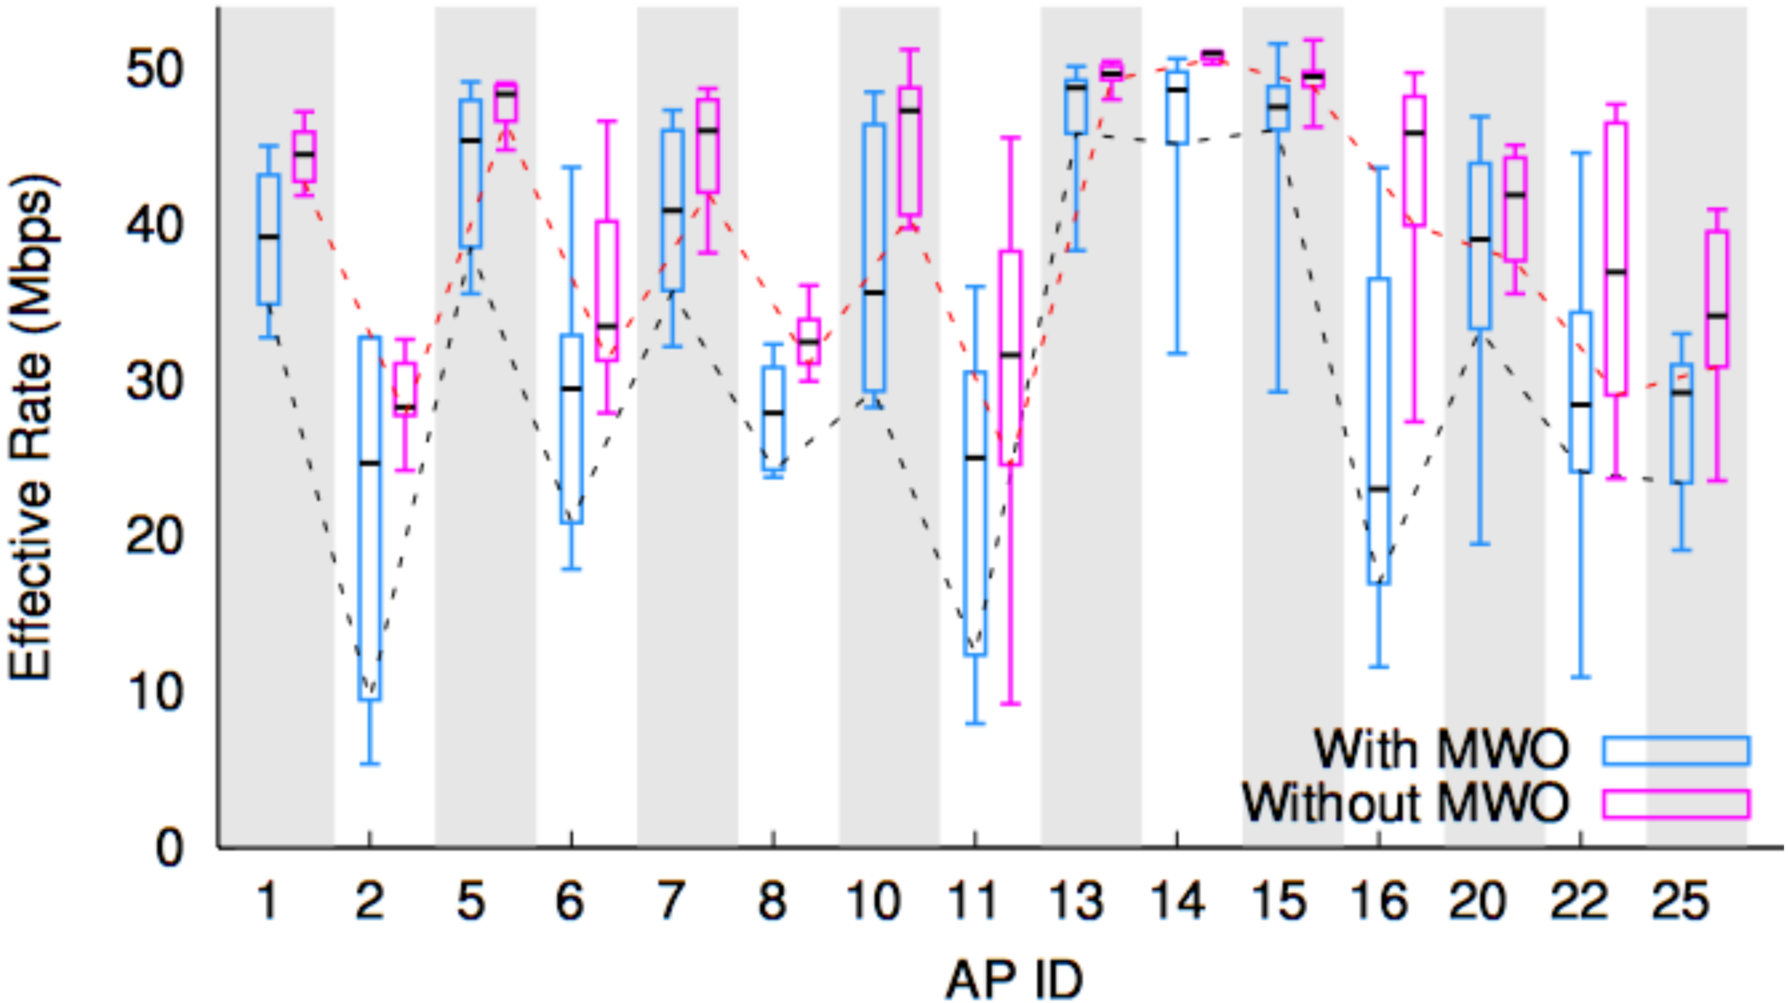

Solution advantage

- Compared solution in one of our apartment buildings

- Approach
  - Day 1: COAP completely disabled
  - Day 2: COAP managed
  - Alternated for nearly two weeks

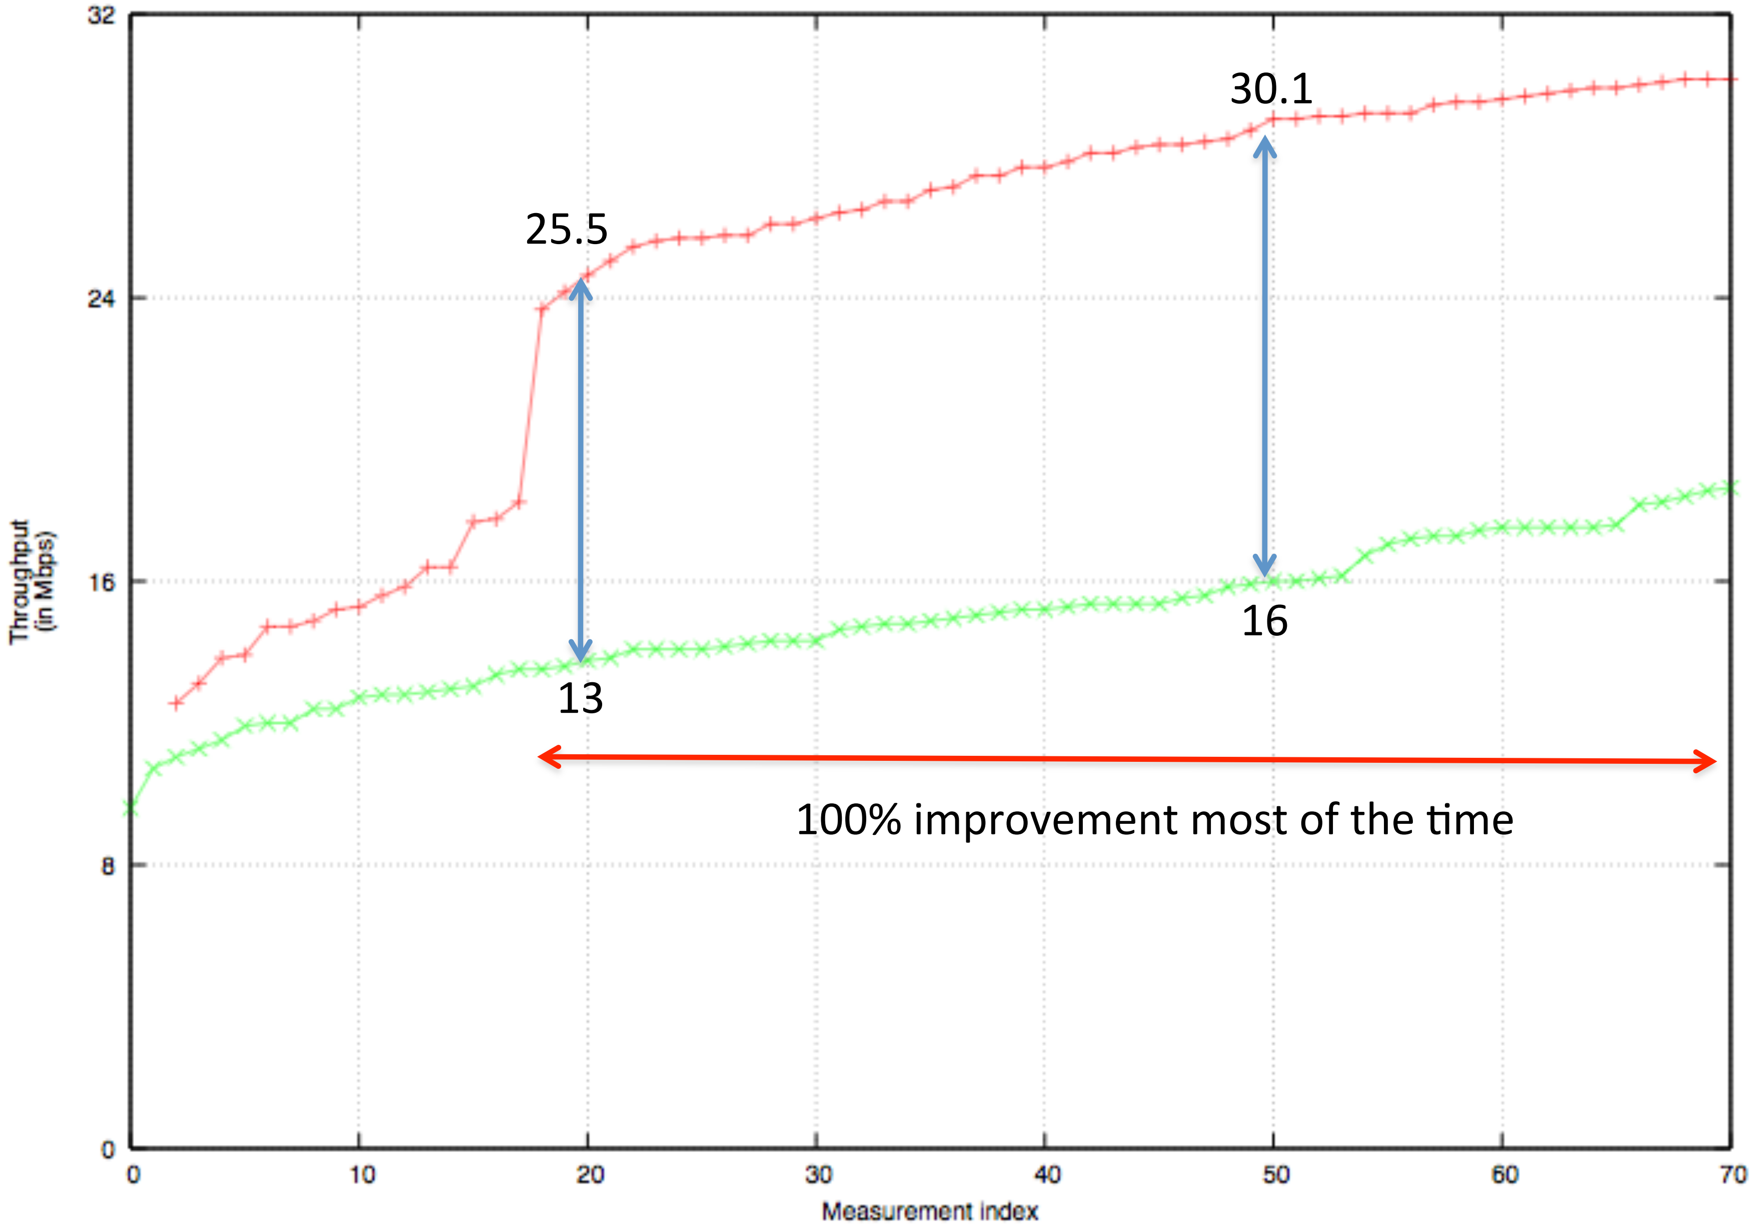

Data is sorted