

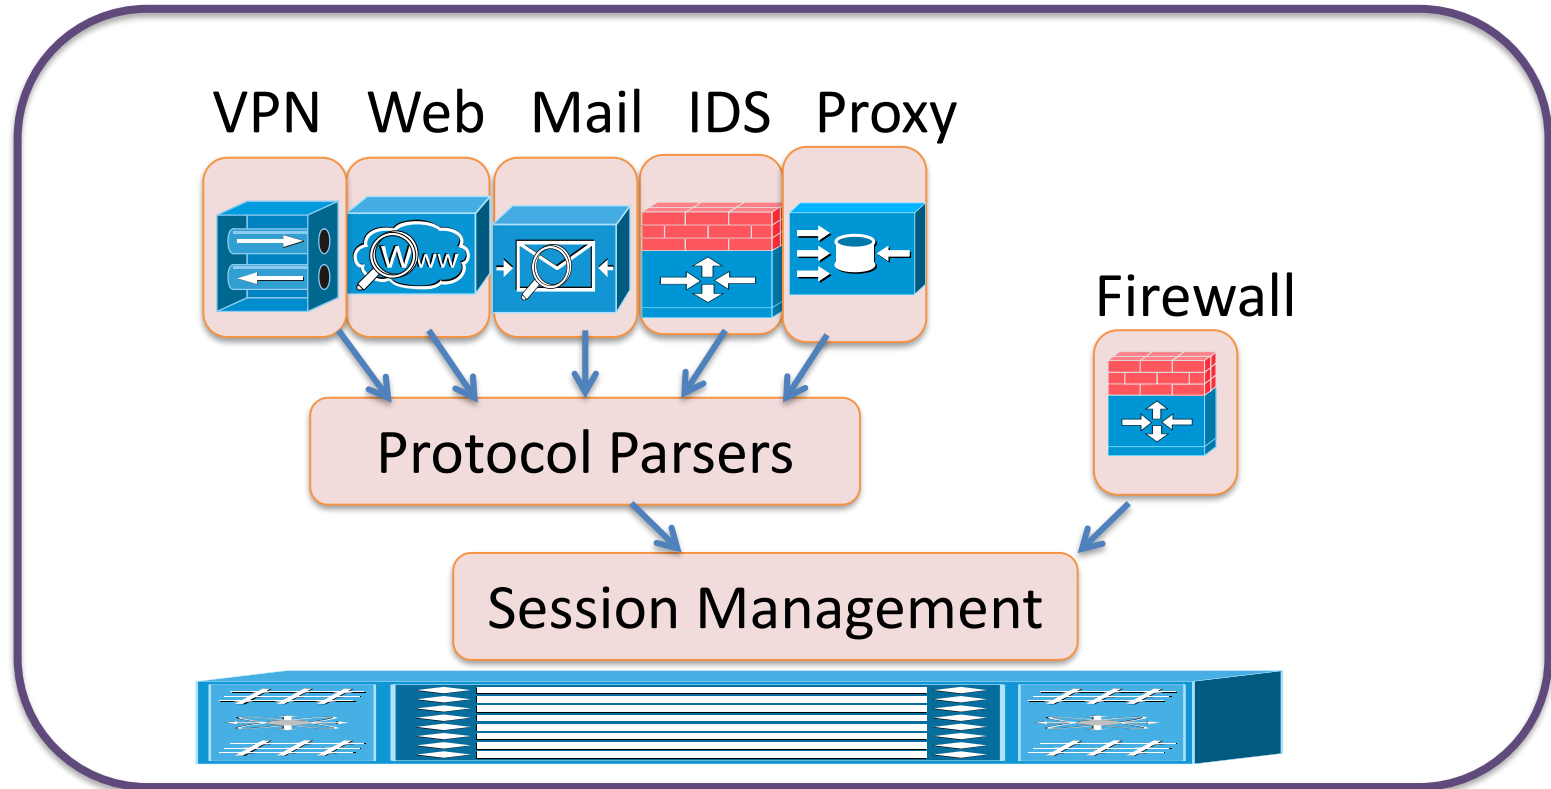
Outline

- Review
 - Midterm
 - SDN and Middleboxes
- SDN Wireless Networks
 - Motivation
 - Data Plane Abstraction: OpenRadio
 - Control Plane Architecture
 - Radio Access Networks: SoftRAN
 - Core Networks: SoftCell

Review of Previous Lecture: Middlebox Basics

- A middlebox is any traffic processing device except for routers and switches.
- Why do we need them?
 - Security
 - Performance
 - Functionality (e.g. echo cancellation, video transcoding)
- Deployments of middlebox functionalities:
 - Embedded in switches and routers (e.g., packet filtering)
 - Specialized devices with hardware support of SSL acceleration, DPI, etc.
 - Virtual vs. Physical Appliances
 - Local (i.e., in-site) vs. Remote (i.e., in-the-cloud) deployments
- They can break end-to-end semantics (e.g., load balancing)

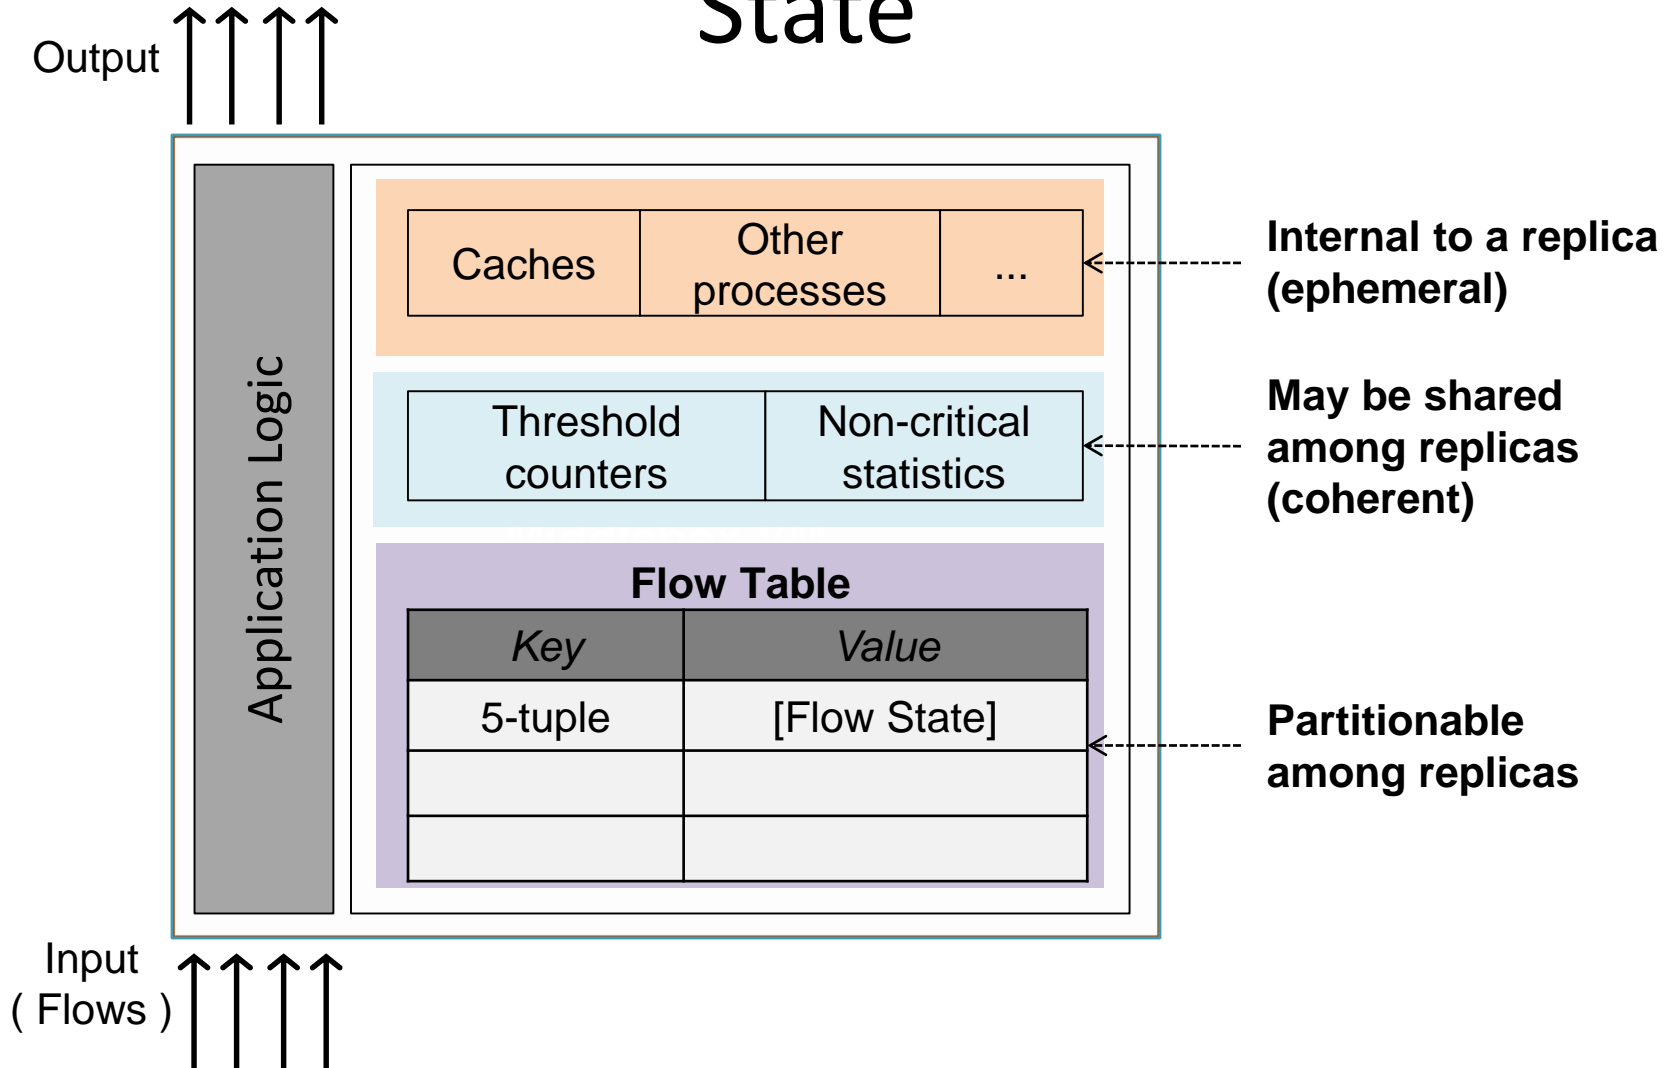
Review of Previous Lecture: Middlebox Consolidation



Contribution of reusable modules: 30 – 80 %

Review of Previous Lecture: Middlebox

State

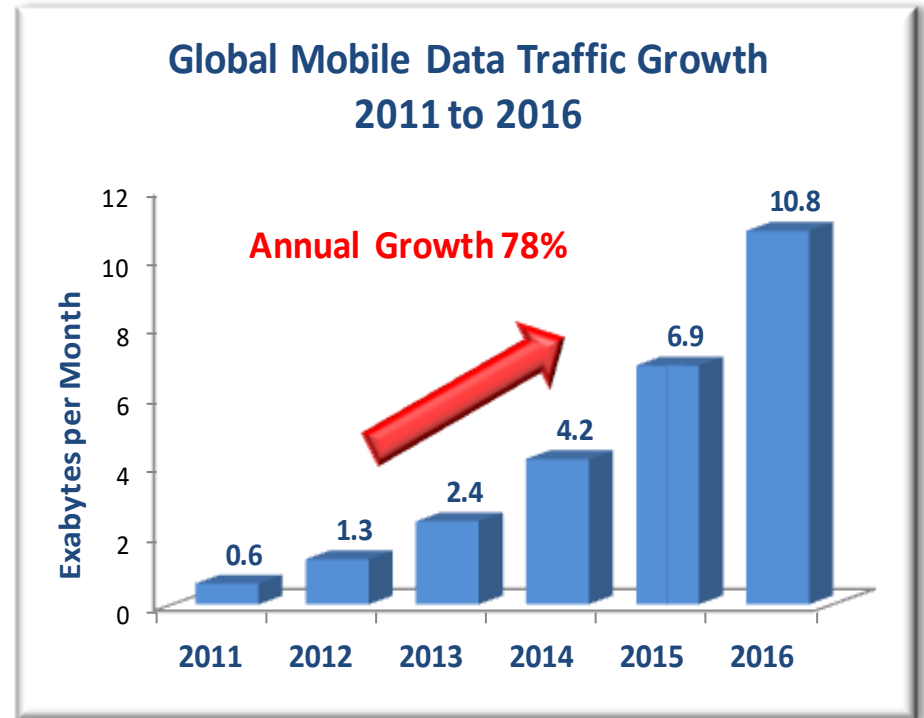


Outline

- Review
 - Midterm
 - SDN and Middleboxes
- SDN Wireless Networks
 - Motivation
 - Data Plane Abstraction: OpenRadio
 - Control Plane Architecture
 - Radio Access Networks: SoftRAN
 - Core Networks: SoftCell

Wireless Data Growth

- AT&T
 - Wireless data growth 20,000% in the past 5 years




Question: How to substantially improve wireless capacity?

Source: CISCO Visual Networking Index (VNI)
Global Mobil Data Traffic Forecast 2011 to 2016

OpenRadio: Access Dataplane

OpenRadio APs built with merchant DSP (digital signal processing) & ARM (Advanced RISC Machine) silicon

- Single platform capable of **LTE, 3G, WiMax, WiFi**
- OpenFlow for Layer 3
- Inexpensive (\$300-500)



Exposes a match/action interface to program how a flow is forwarded, scheduled & encoded

Design goals and Challenges

Programmable wireless dataplane using off-the-shelf components

- At least 40MHz OFDM-complexity performance
 - More than 200 GLOPS computation
 - Strict processing deadlines, eg. 25us ACK in WiFi
- Modularity to provide ease of programmability
 - Only modify affected components, reuse the rest
 - Hide hardware details and stitching of modules

Design principle I

Judiciously scoping flexibility


- Provide just enough flexibility
- Keep blocks coarse
- Higher level of abstraction
- High performance through hardware acceleration
 - Viterbi co-processor
 - FFT co-processor
- Off-the-shelf heterogeneous multicore DSPs
 - TI, CEVA, Freescale etc.

Algorithm	WiFi	LTE	3G	DVB-T
FIR / IIR	√	√	√	√
Correlation	√	√	√	√
Spreading			√	
FFT	√	√		√
Channel Estimation	√	√	√	√
QAM Mapping	√	√	√	√
Interleaving	√	√	√	√
Convolution Coding	√	√	√	√
Turbo Coding		√	√	
Randomization	√	√	√	√
CRC	√	√	√	


Design principle II

Processing-Decision separation


- Logic pulled out to decision plane
- Blocks and actions are branch-free
 - Deterministic execution times
 - Efficient pipelining, algorithmic scheduling
 - Hardware is abstracted out



Regular compilation	OpenRadio scheduling
Instructions	Atomic processing blocks
Heterogeneous functional units	Heterogeneous cores
Known cycle counts	Predictable cycle counts
Argument data dependency	FIFO queue data dependency



Prototype

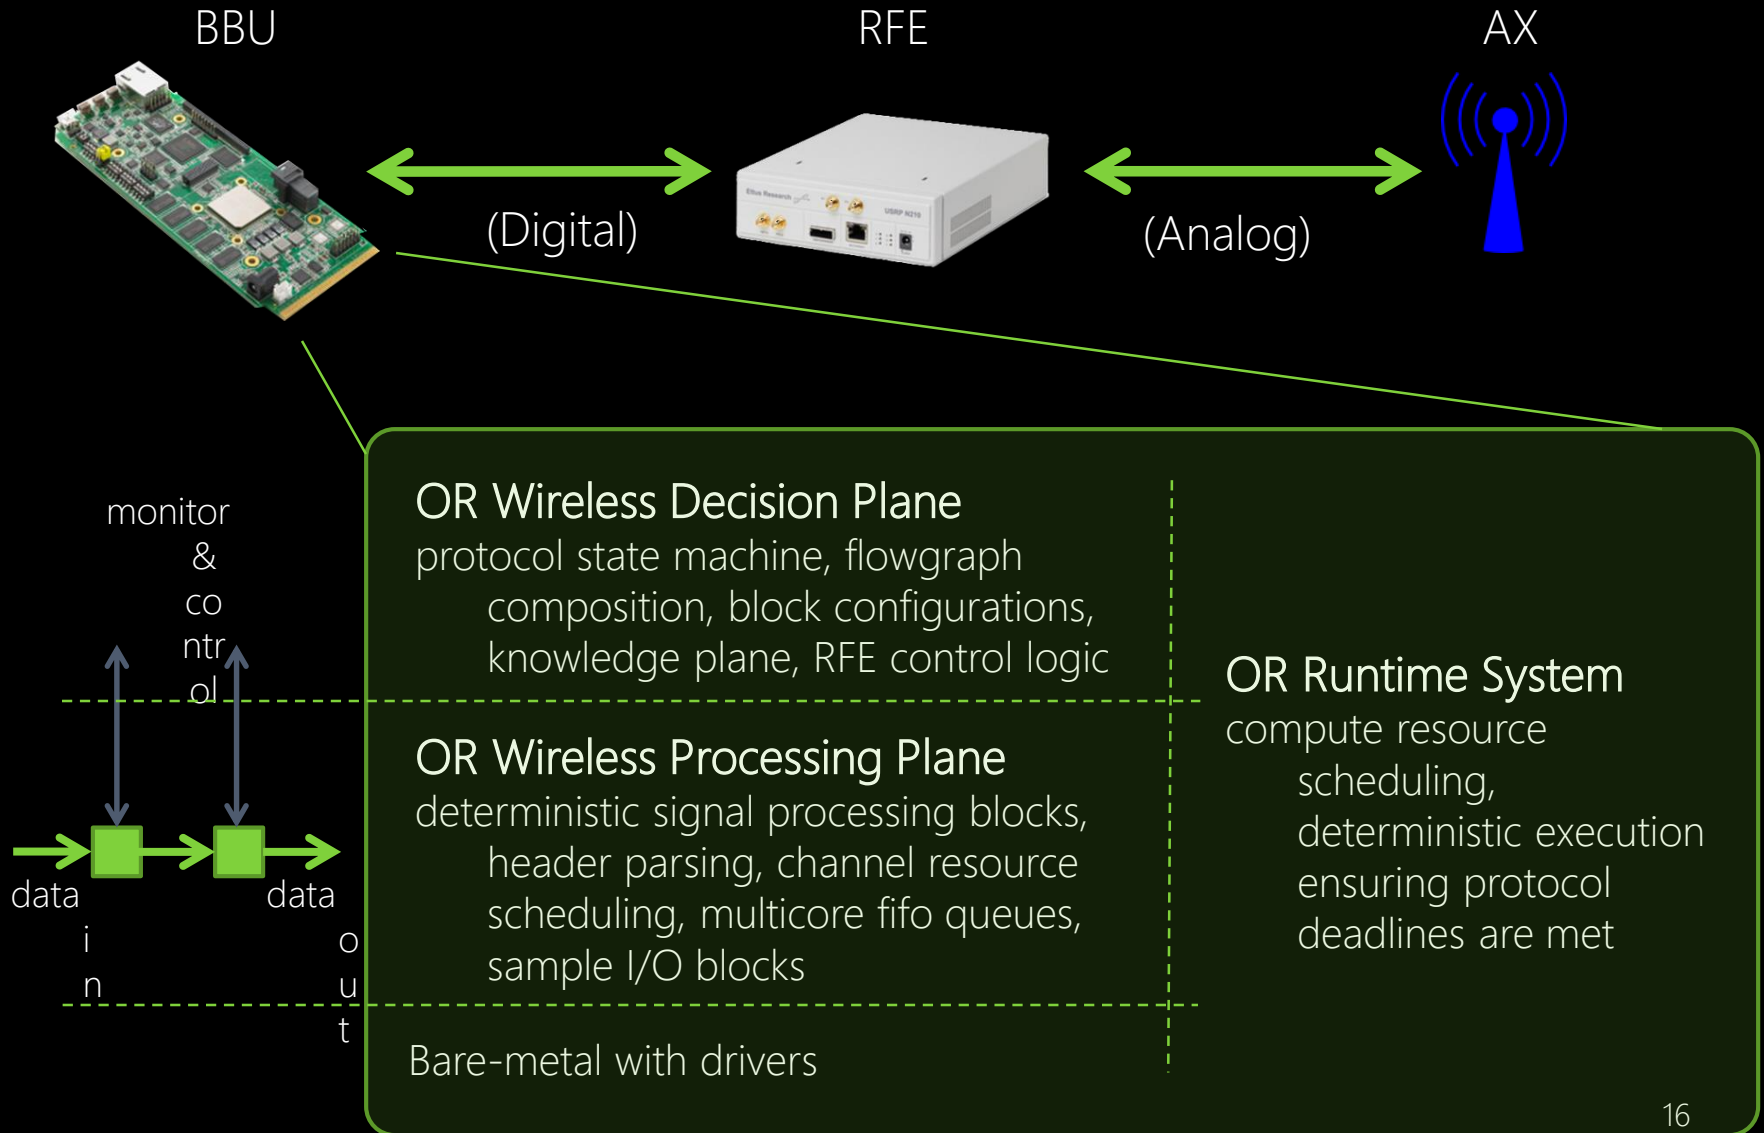


- **COTS TI KeyStone multicore DSP platform** (EVM6618, two chips with 4 cores each at 1.2GHz, configurable hardware accelerators for FFT, Viterbi, Turbo)
- **Prototype can process 40MHz, 108Mbps 802.11g on one chip using 3 of 4 cores**

OpenRadio: Current Status

- OpenRadio APs with full WiFi/LTE software on **TI** C66x DSP silicon
- OpenRadio commodity WiFi APs with a firmware upgrade
- Network OS

Software architecture



Summary

Provides programmatic interfaces to monitor and program wireless networks

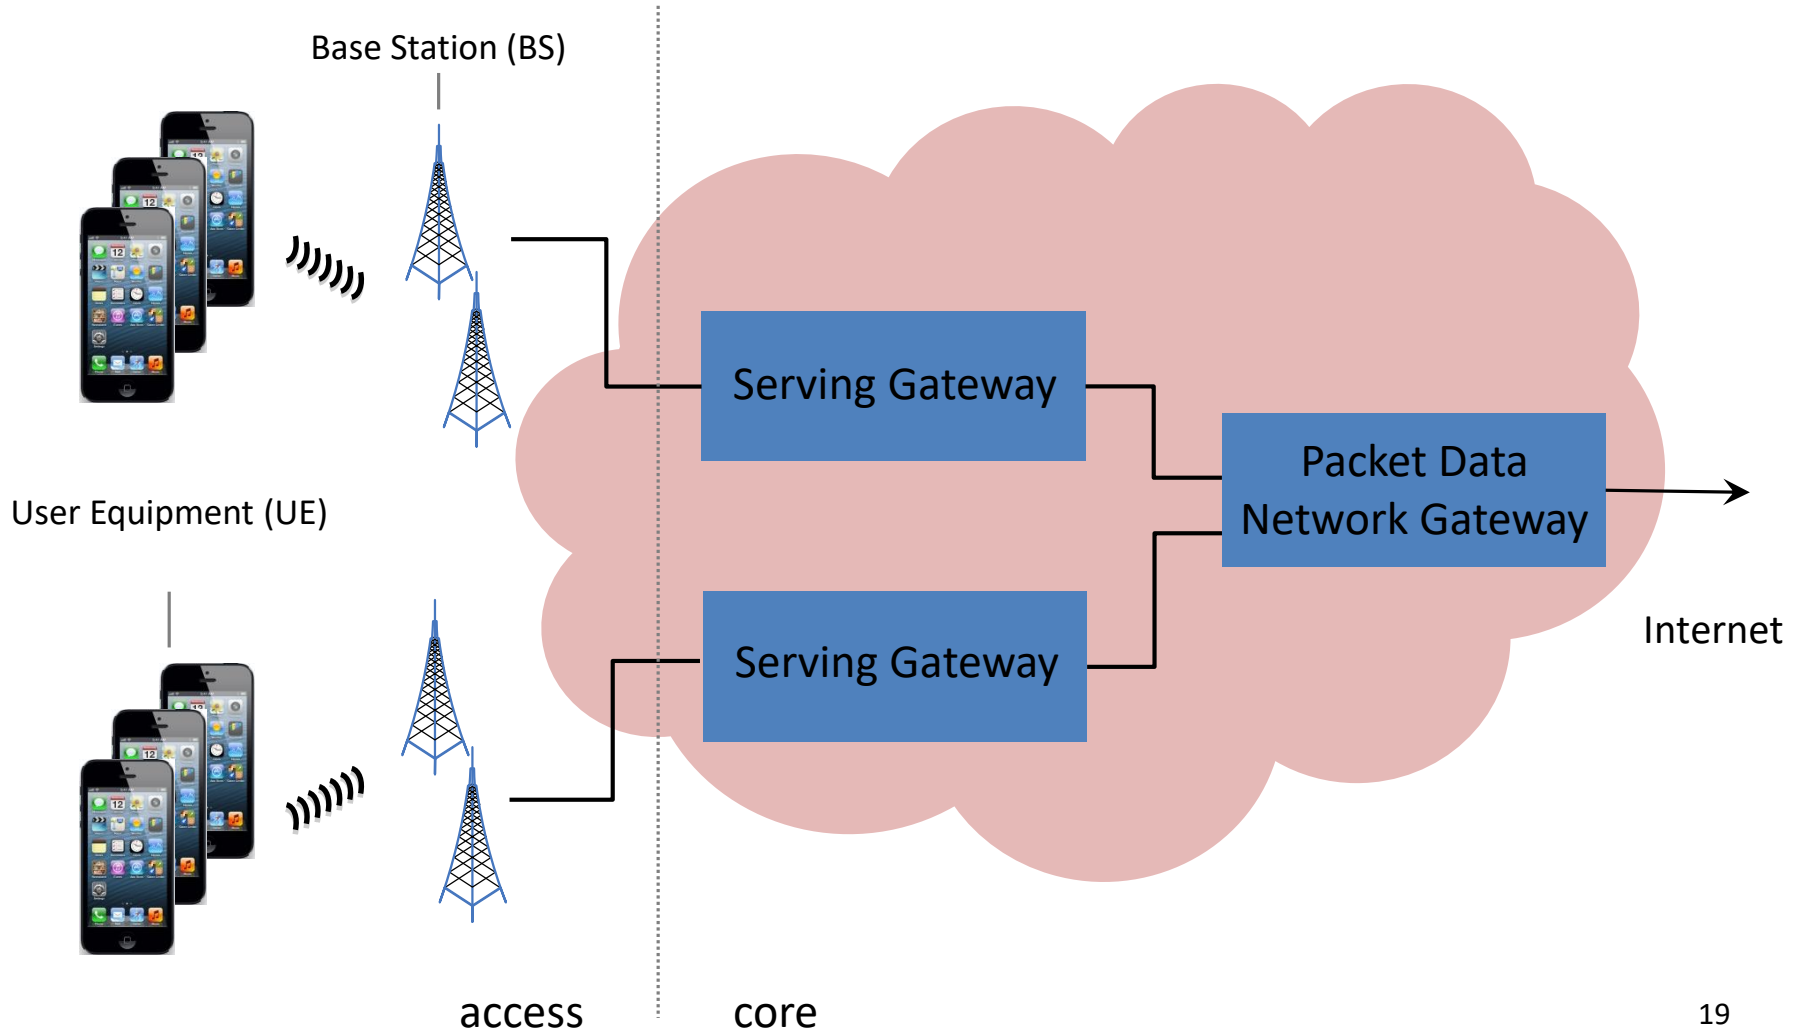
- High performance substrate using merchant silicon

Outline

- Review
 - Midterm
 - SDN and Middleboxes
- SDN Wireless Networks
 - Motivation
 - Data Plane Abstraction: OpenRadio
 - Control Plane Architecture
 - Radio Access Networks: SoftRAN
 - Core Networks: SoftCell

LTE Radio Access Networks


- Goal: high capacity wide-area wireless network



Coping with Increasing Traffic


- Increasing demand on wireless resources
 - Dense deployments
 - Radio resource management (RRM) decisions made at one base station affect neighboring base stations
 - RRM needs to be coordinated

Radio Resource Management: Interference



- Power used at BS1 affects interference seen at Client 2
- Interference seen at Client 2 affects power required at BS2

Radio Resource Management : Mobility




- Coordination required to decide handovers
- Load at BS1 reduces and load at BS2 increases

LTE-RAN: Current Architecture


- Distributed control plane
 - Control signaling grows with density
 - Inefficient RRM decision making
 - Harder to manage and operate the network
 - Clients need to resynchronize state at every handover
- Works fine with sparse deployments, but problems compound in a dense network

SoftRAN: Big Base Station Abstraction




Radio Resource Allocation

Flows




3D Resource Grid

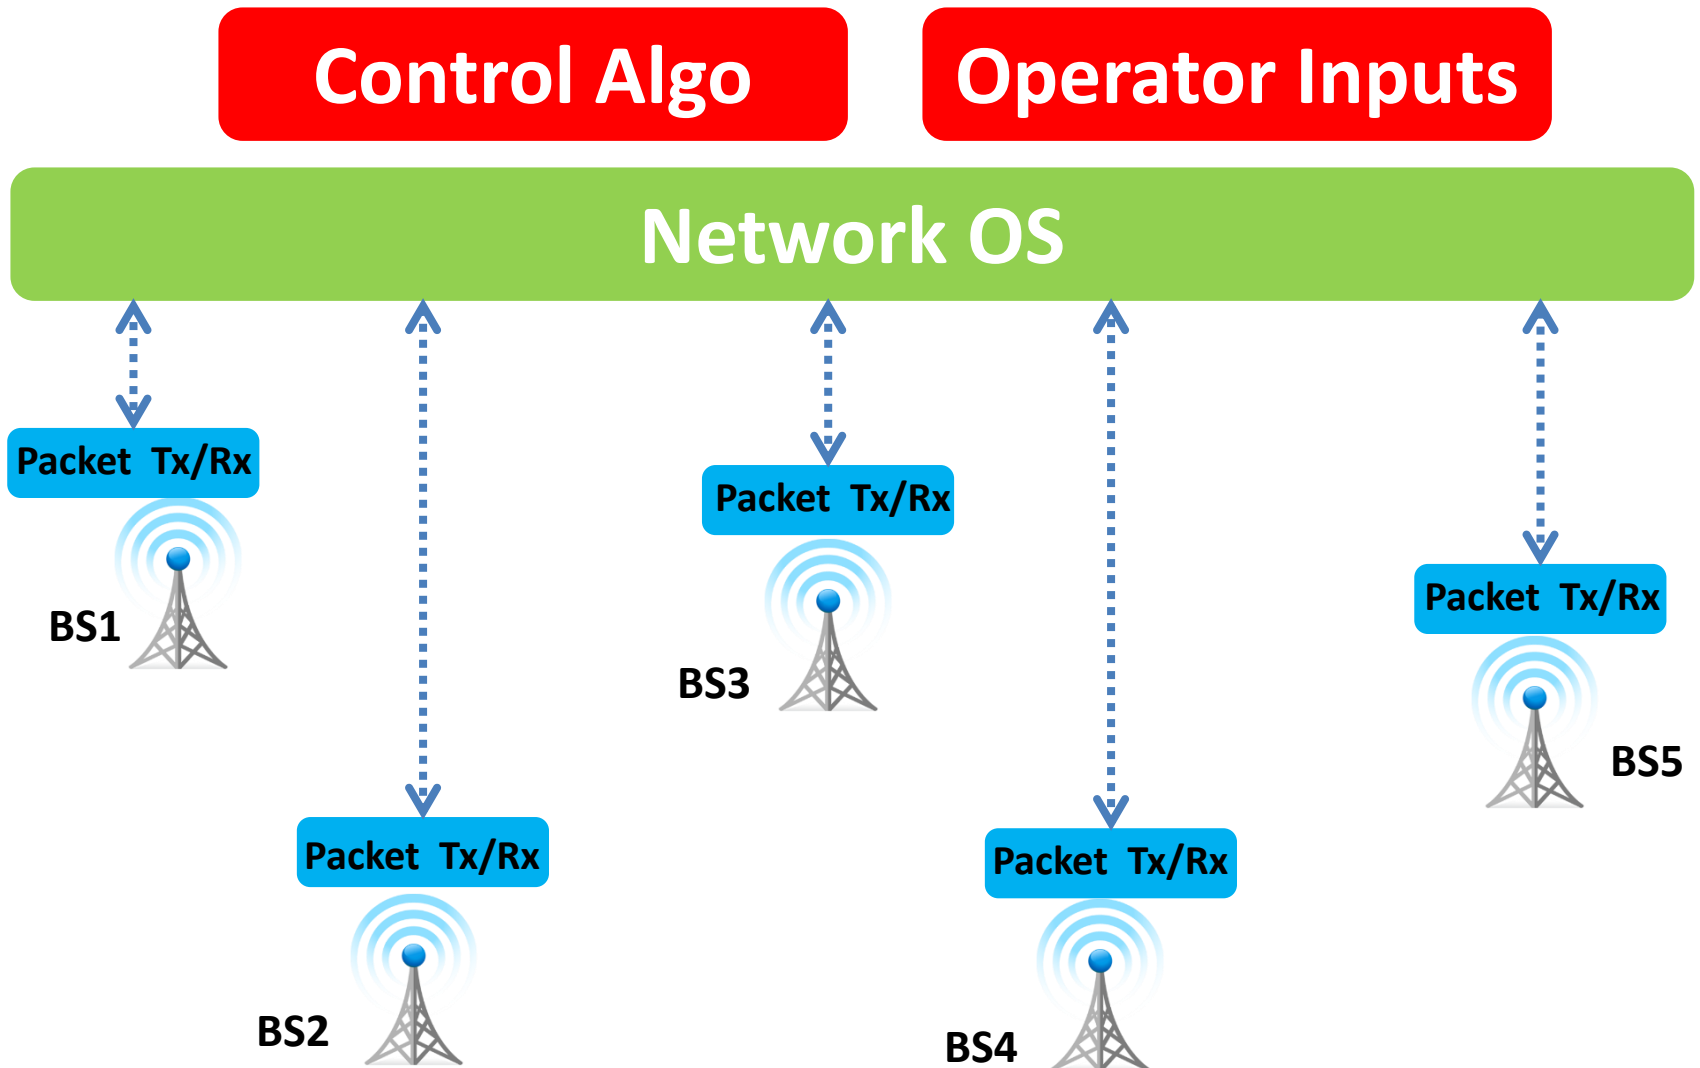


SoftRAN: SDN Approach to RAN

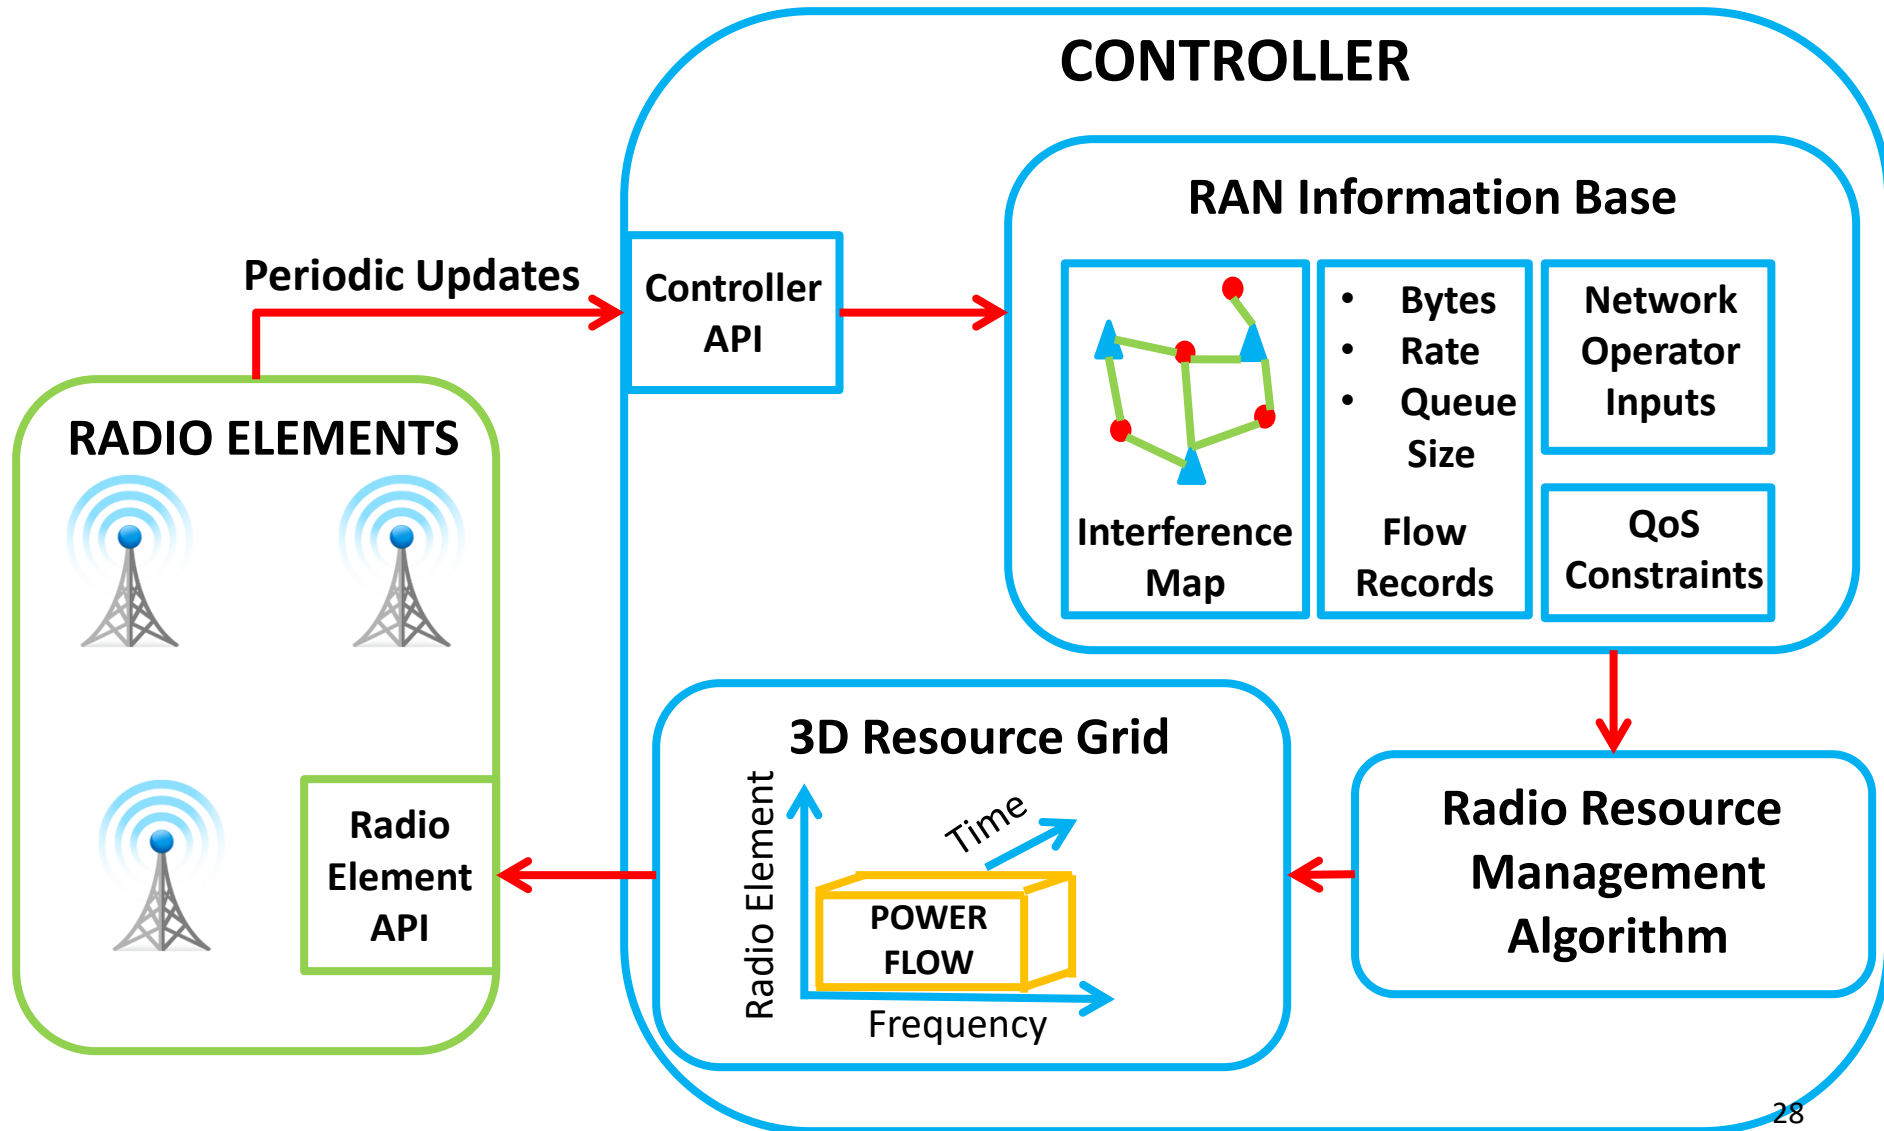
↔ Coordination :
X2 Interface



SoftRAN: SDN Approach to RAN



SoftRAN Architecture Summary



SoftRAN Architecture: Updates

- Radio element -> controller (updates)
 - Flow information (downlink and uplink)
 - Channel states (observed by clients)
- Network operator -> controller (inputs)
 - QoS requirements
 - Flow preferences



SoftRAN Architecture: Controller Design

- RAN information base (RIB)
 - Update and maintain global network view
 - Interference map
 - Flow records
- Radio resource management
 - Given global network view: maximize global utility
 - Determine RRM at each radio element

SoftRAN Architecture: Radio Element API

- Controller -> radio element
 - Handovers to be performed
 - RF configuration per resource block
 - Power allocation and flow allocation
 - Relevant information about neighboring radio elements
 - Transmit Power being used

SoftRAN: Backhaul Latency



Refactoring Control Plane

- Controller responsibilities:
 - Decisions influencing global network state
 - Load balancing
 - interference management
- Radio element responsibilities:
 - Decisions based on frequently varying local network state
 - Flow allocation based on channel states

SoftRAN Advantages

- Logically centralized control plane:
 - Global view on interference and load
 - Easier coordination of radio resource management
 - Efficient use of wireless resources
 - Plug-and-play control algorithms
 - Simplified network management
 - Smoother handovers
 - Better user-experience

SoftRAN: Evolving the RAN

- Switching off radio elements based on load
 - Energy savings
- Dynamically splitting the network into Big-BSs
 - Handover radio elements between Big-BSs

Implementation: Modifications

- SoftRAN is incrementally deployable with current infrastructure
 - No modification needed on client-side
 - API definitions at base station
 - Femto API : Standardized interface between scheduler and L1 (<http://www.smallcellforum.org/resources-technical-papers>)
 - Minimal modifications to FemtoAPI required

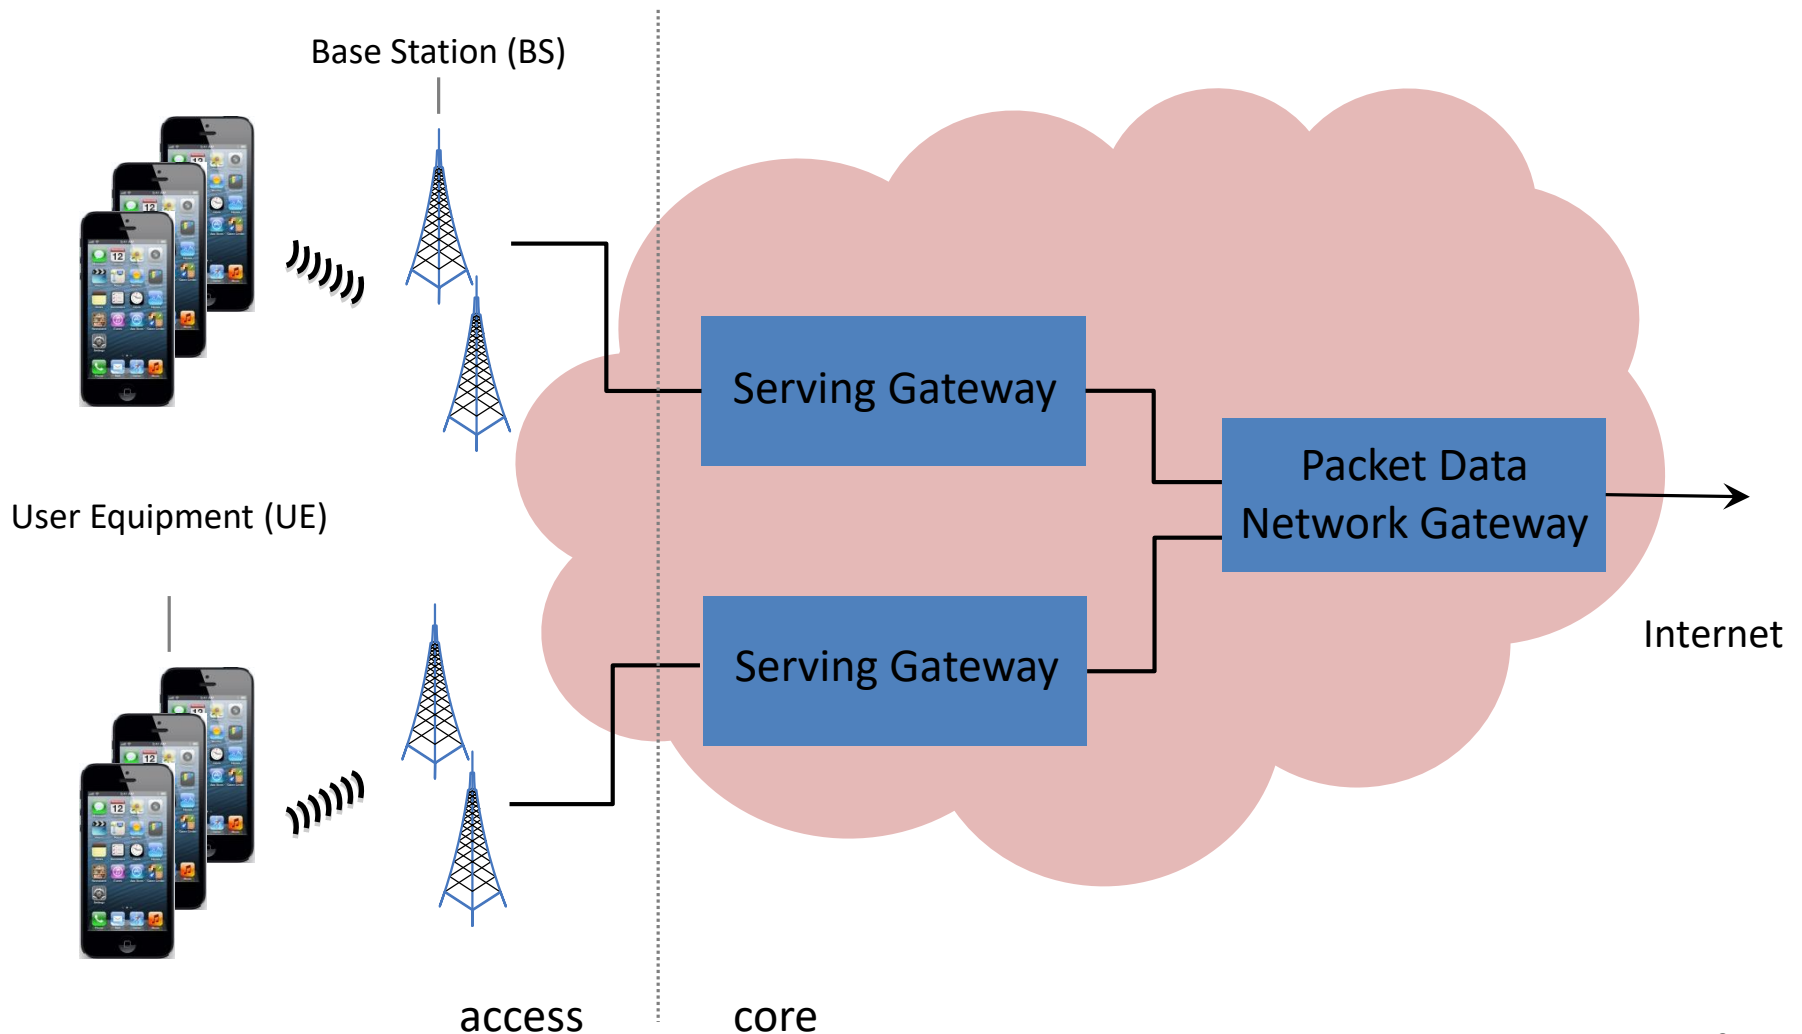
Implementation (Ongoing): Controller

- Floodlight : controller implementation
- Radio resource management algorithm
 - Load balancing
 - Interference management
 - QoS constraints
 - Network operator preferences

Outline

- Review
 - Midterm
 - SDN and Middleboxes
- SDN Wireless Networks
 - Motivation
 - Data Plane Abstraction: OpenRadio
 - Control Plane Architecture
 - Radio Access Networks: SoftRAN
 - Core Networks: SoftCell

Cellular Core Network Architecture



Cellular core networks are not flexible

- Most functionalities are implemented at **Packet Data Network Gateway**
 - Content filtering, application identification, stateful firewall, lawful intercept, ...



Packet Data
Network Gateway

- This is not flexible

Combine functionality from different vendors



Easy to add new functionality

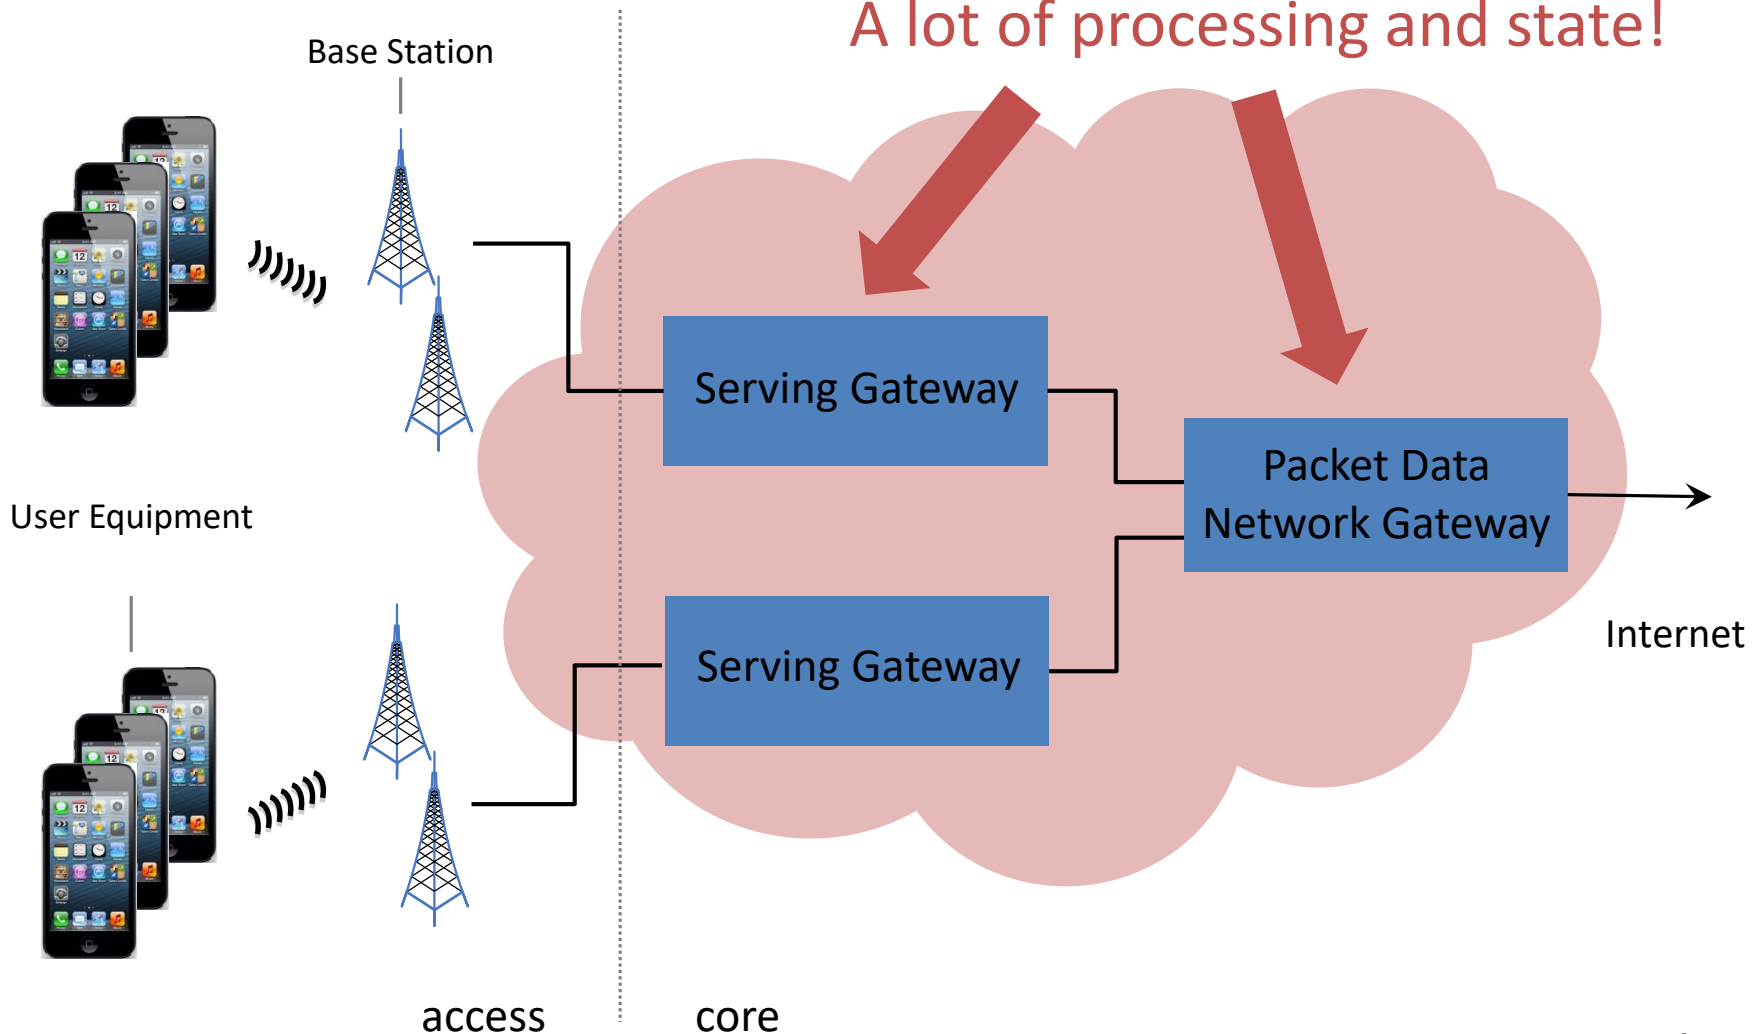


Only expand capacity for bottlenecked functionality

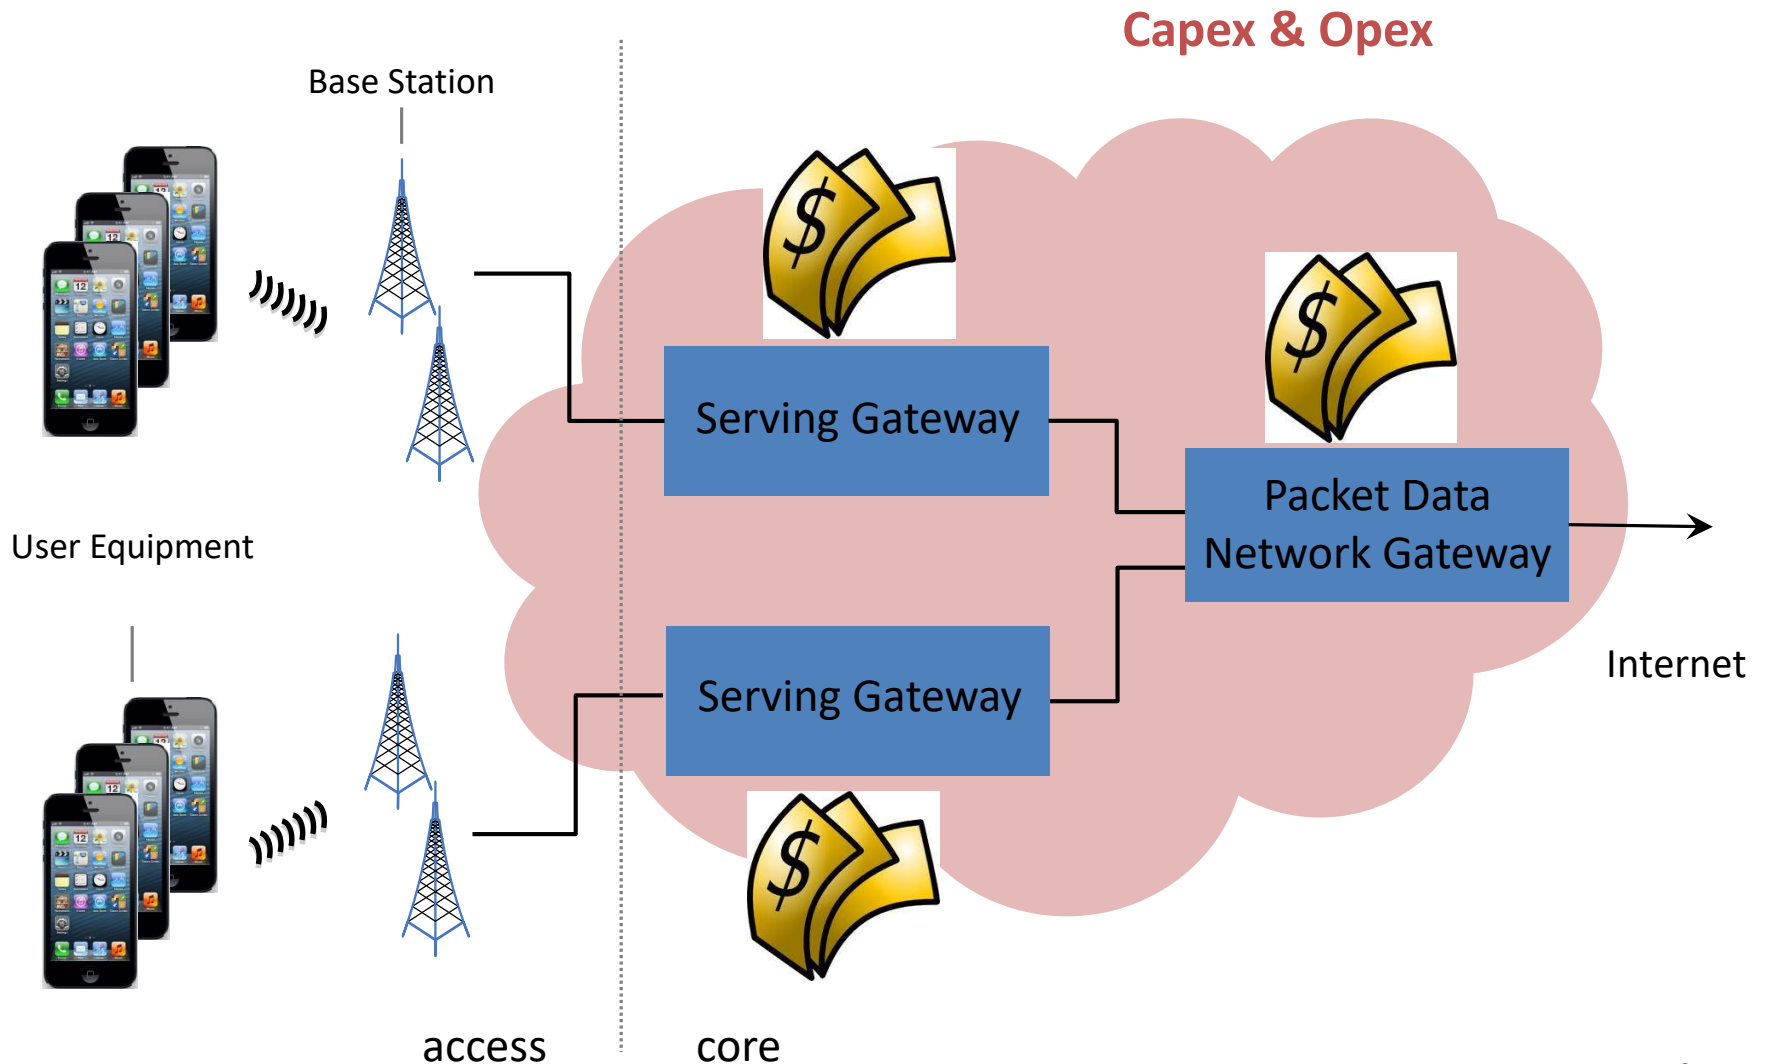


Cellular core networks are not **scalable**

A lot of processing and state!



Cellular core networks are not **cost-effective**



Can we make cellular core networks like data center networks?



- ✓ *Flexible*
- ✓ *Scalable*
- ✓ *Cost-Effective*

Can we make cellular core networks like data center networks?

Yes! With **SoftCell!**


- ✓ *Flexible*
- ✓ *Scalable*
- ✓ *Cost-Effective*

SoftCell Overview

No change

Commodity hardware
+ **SoftCell** software

No change



SoftCell: Taking Control of Cellular Core Networks

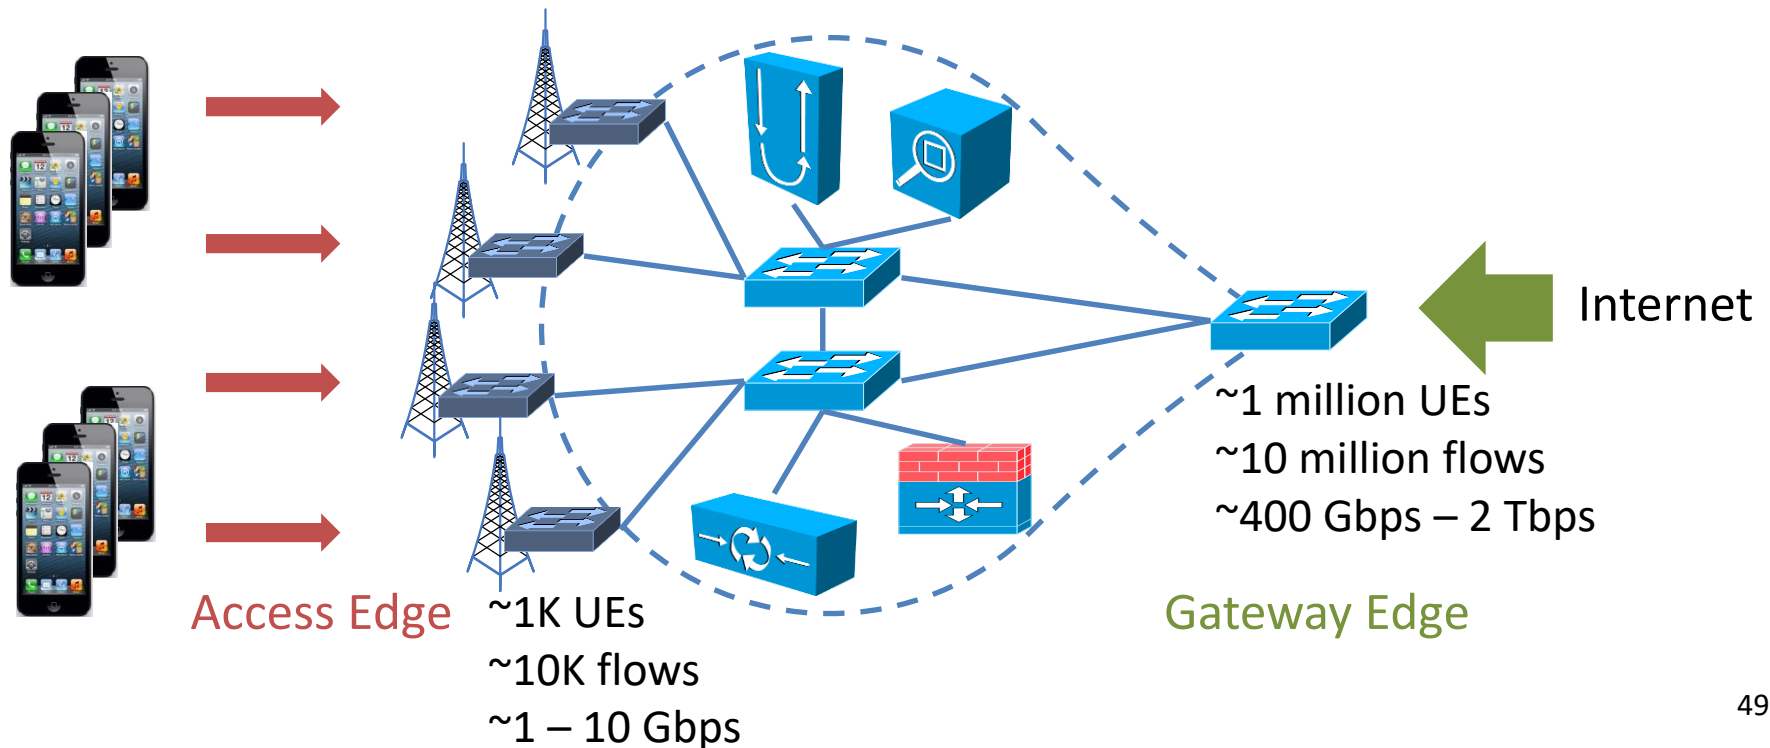
- Characteristics of Cellular Core Networks
- Scalable Data Plane
 - Asymmetric Edge: Packet Classification
 - Core: Multi-Dimensional Aggregation
- Scalable Control Plane
 - Hierarchical Controller

Characteristics of Cellular Core Networks

1. “North south” traffic pattern: in cellular core networks, most traffic is from/to the Internet
 - In data centers, 76% traffic is intra data center traffic.
[Cisco Global Cloud Index]
 - cellular core networks have asymmetric edge. The access edge has lower bandwidth than the gateway edge

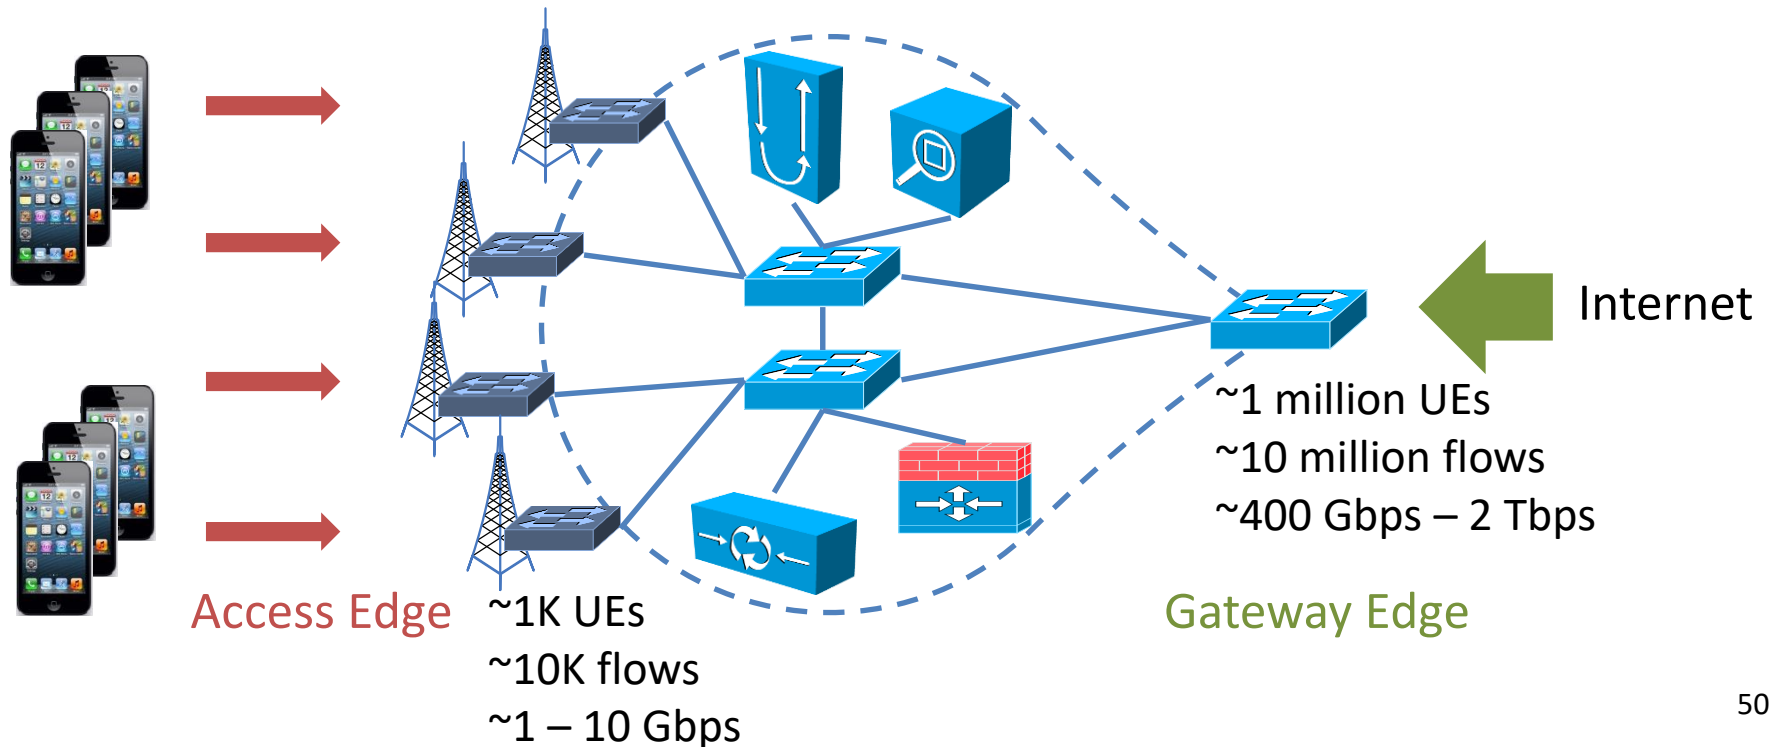
Characteristics of Cellular Core Networks

1. “North south” traffic pattern
2. **Asymmetric edge**: low-bandwidth access edge vs. high-bandwidth gateway edge



Characteristics of Cellular Core Networks

1. “North south” traffic pattern
2. Asymmetric edge
3. Traffic initiated from low-bandwidth access edge



Characteristics of Cellular Core Networks

1. “North south” traffic pattern
2. Asymmetric edge
3. Traffic initiates from low-bandwidth access edge

Goal: Scalable support of fine-grained policies in such a network



with diverse needs!

Fine-grained and sophisticated policies

“ I want

traffic to customer with parental control to
go through a firewall then a content filter

and

balance the load among
all content filters and firewalls in the network!”

Decouple the problem

“ I want

traffic to customer with parental control to go through a firewall then a content filter

and

balance the load among all content filters and firewalls in the network!”

Decouple the problem

Service Policy: meet customer demand

traffic to customer with parental control to go through a firewall then a content filter

Traffic Management Policy: meet operational goal

balance the load among all content filters and firewalls in the network!”

Decouple the problem

Service Policy: meet customer demand

subscriber attributes + application type
→ an ordered list of middleboxes



Normal Customer



Parental Control

Content Filter <-> Firewall



Normal Customer

Firewall



Government Customer

IPS <-> Firewall



“Gold Plan” Customer




Web Traffic

Web Accelerator <-> Customized Firewall

Decouple the problem

Traffic Management Policy: meet operational goal

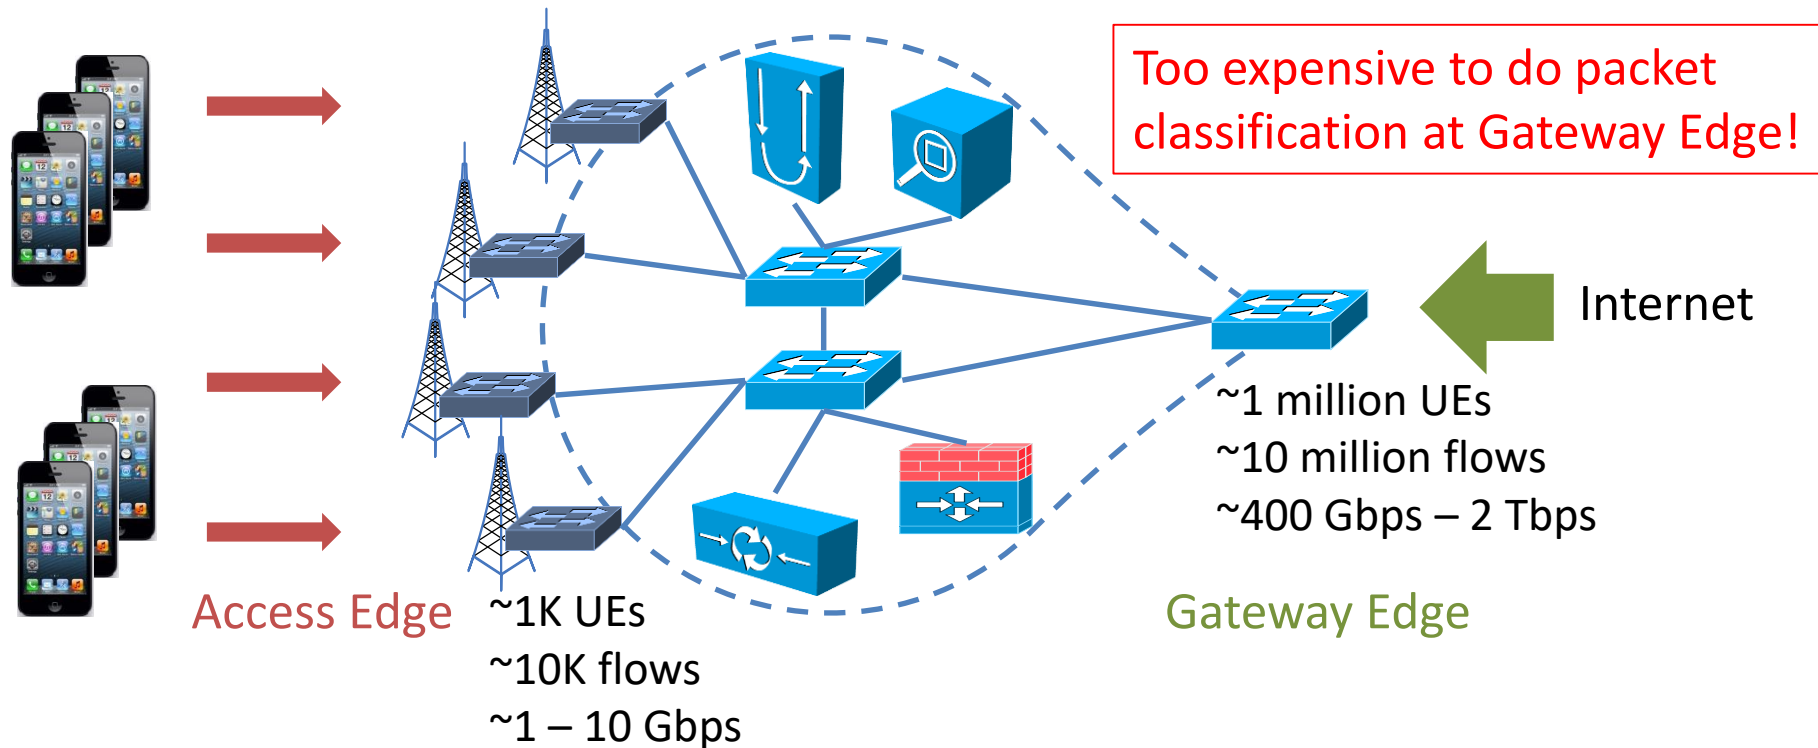
Specify how to allocate network resources, e.g. load balance among multiple middlebox instances



Challenge: Scalability

- **Packet Classification:** decide which **service policy** to be applied to a flow
 - How to classify **millions of flows**?
- **Traffic Steering:** generate switch rules to implement paths given by **traffic management policy**
 - How to implement **million of paths**?

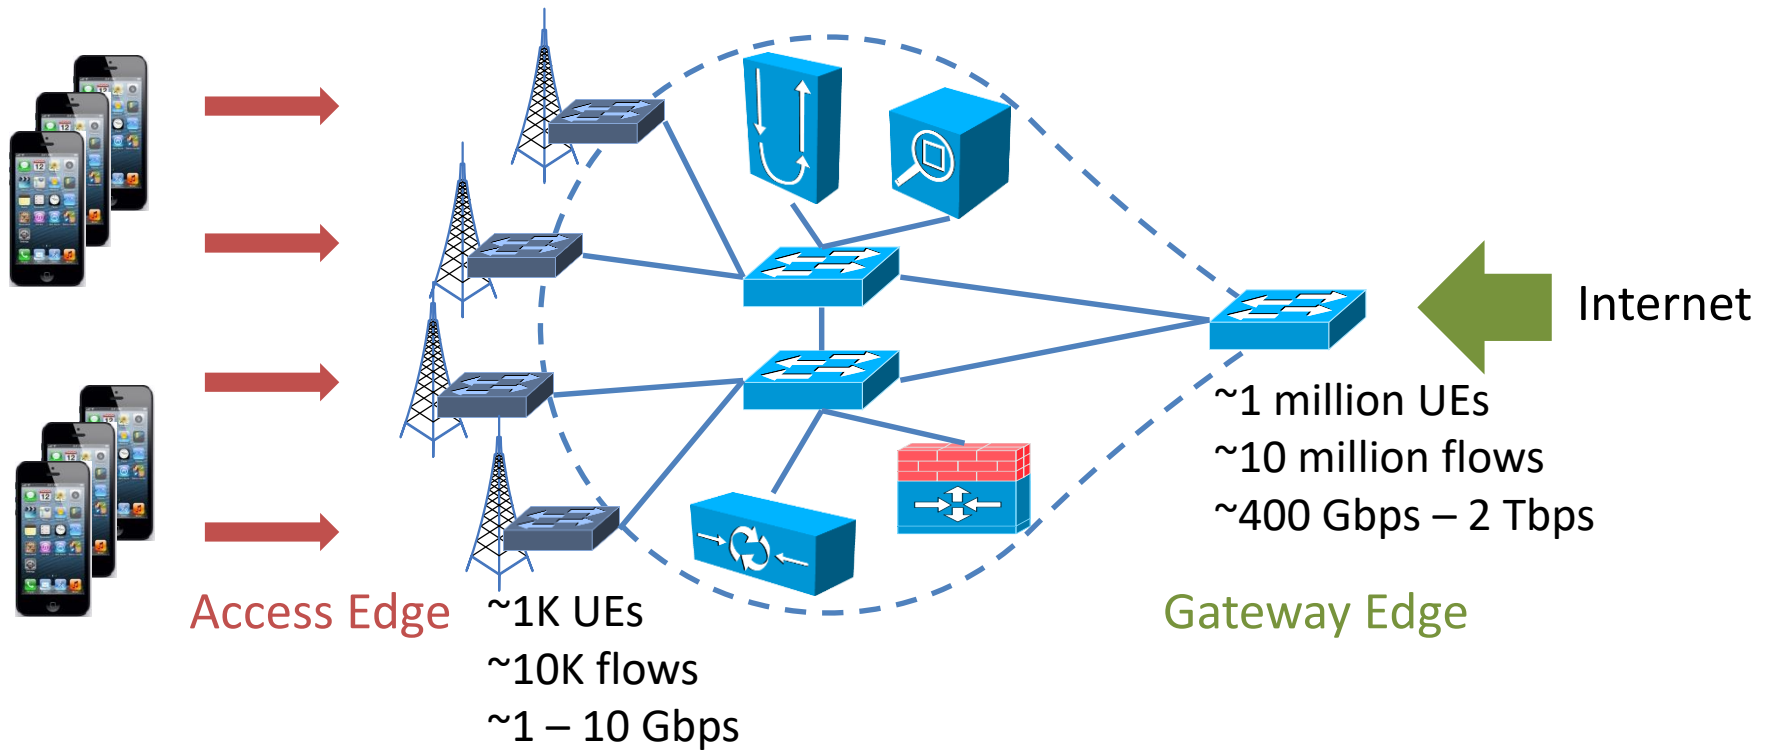
“North south” Traffic Pattern



- Low traffic volume
- Small number of active flows

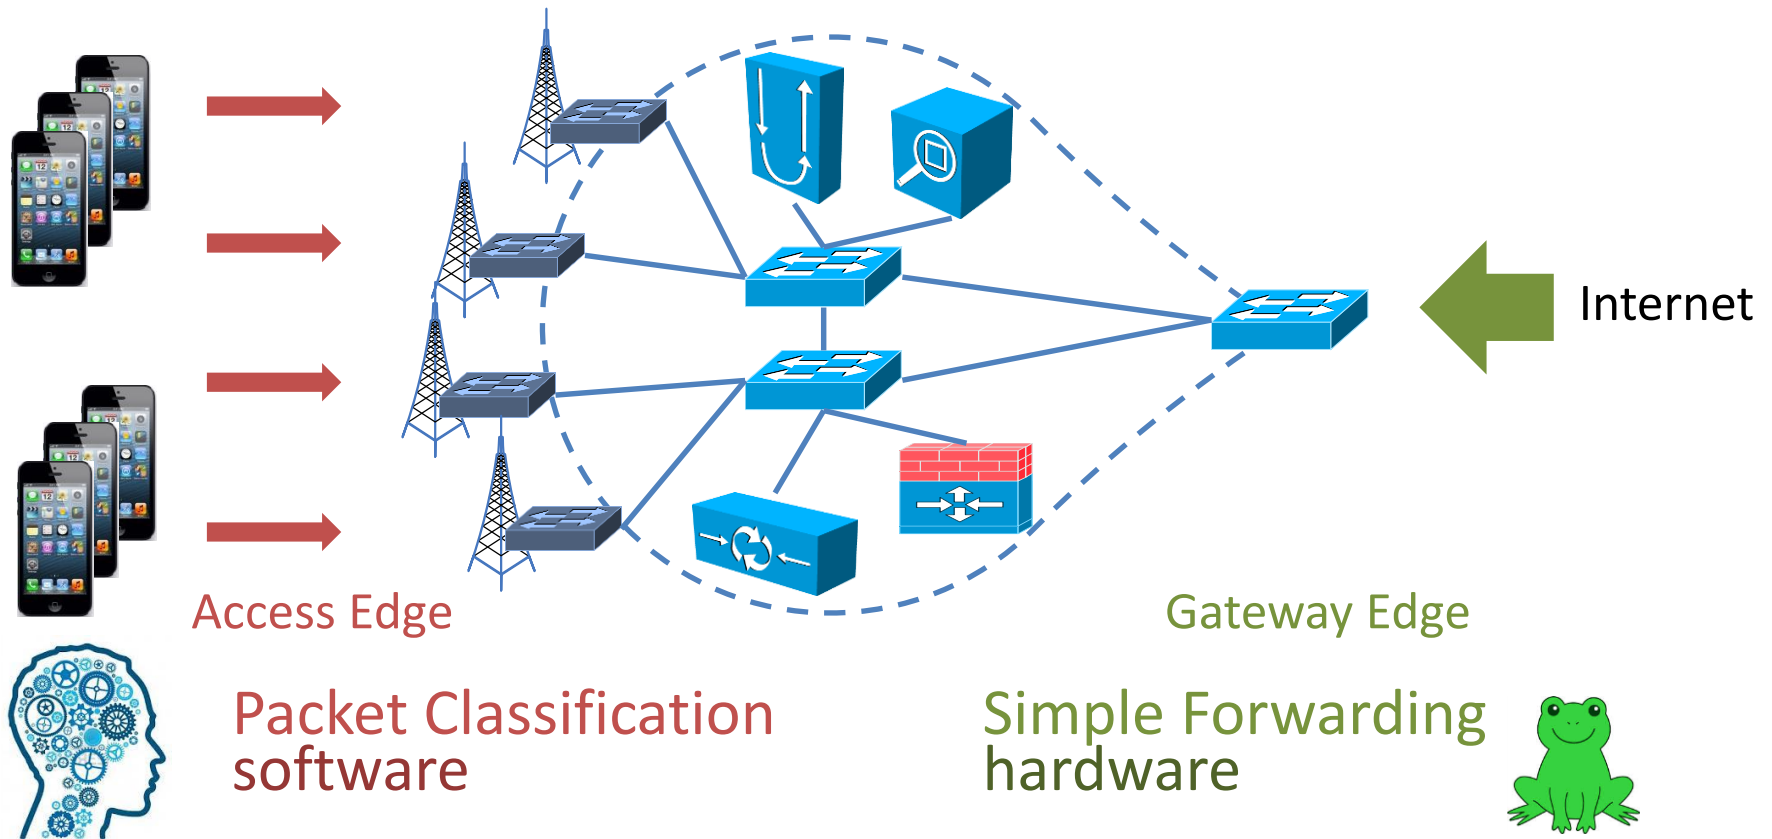
- High traffic volume
- Huge number of active flows

“North south” Traffic Pattern



Opportunity: Traffic initiated from the access edge!

Asymmetric Edge: Packet Classification



- **Encode** classification results in packet header

- Classification results are **implicitly piggybacked** in header

Challenge: Scalability

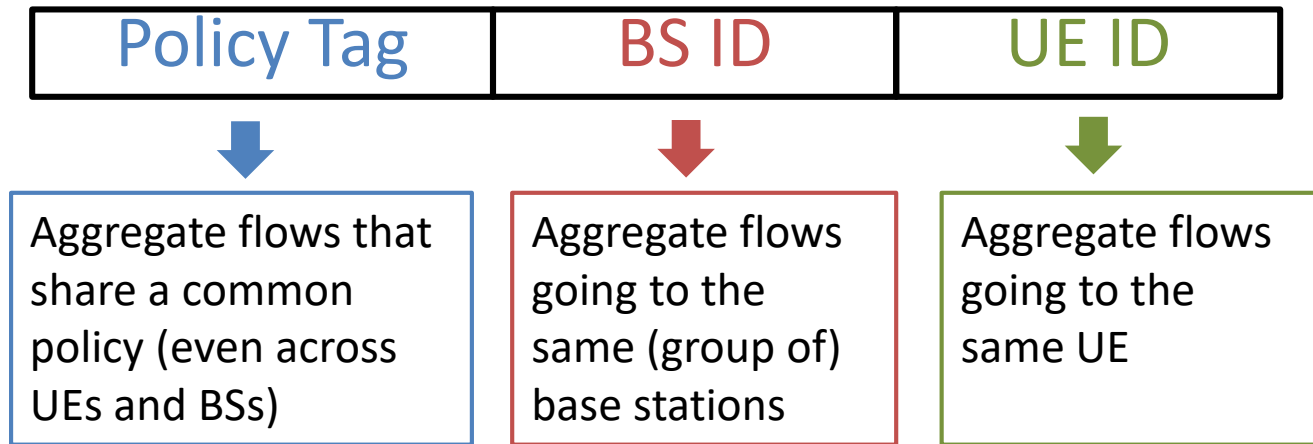
- **Packet Classification:** decide which service policy to be applied to a flow
 - How to classify millions of flows?
- **Traffic Steering:** generate switch rules to implement paths given by traffic management policy
 - How to implement million of paths?

Traffic Steering

- Steering traffic to go through different sequences of middlebox instances
 - Difficult to configure with traditional layer-2 or layer-3 routing
 - [PLayer'08] use packet classifiers, large flow table
- What about use a tag to encode a path?
 - Aggregate traffic of the same path
 - Suppose 1000 service policy clauses, 1000 base stations
 - May result in 1 million paths, need 1 million tags
- Limited switch flow tables: ~1K – 4K TCAM, ~16K – 64K L2/Eth
- Solution: multi-dimensional aggregation


Multi-Dimensional Aggregation

- Use **multi-dimensional** tags rather than flat tags

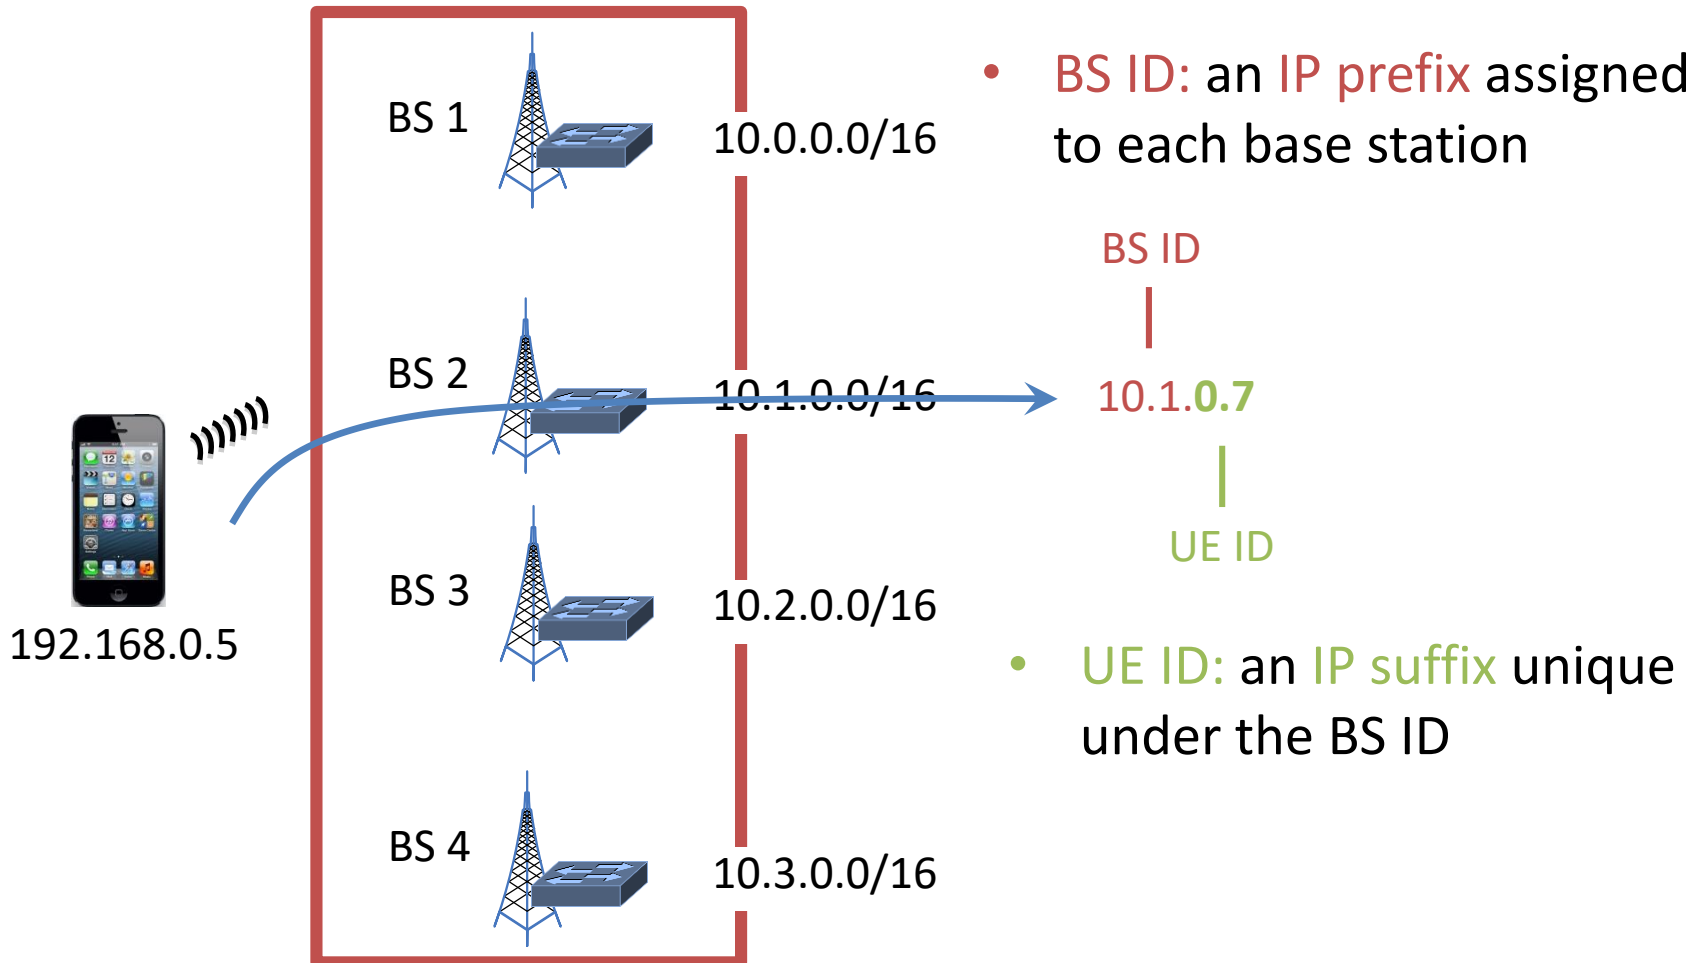


- Exploit **locality** in the network
- **Selectively** match on one or multiple dimensions
 - Supported by TCAM in today's switches

Location-Based Hierarchical IP Address

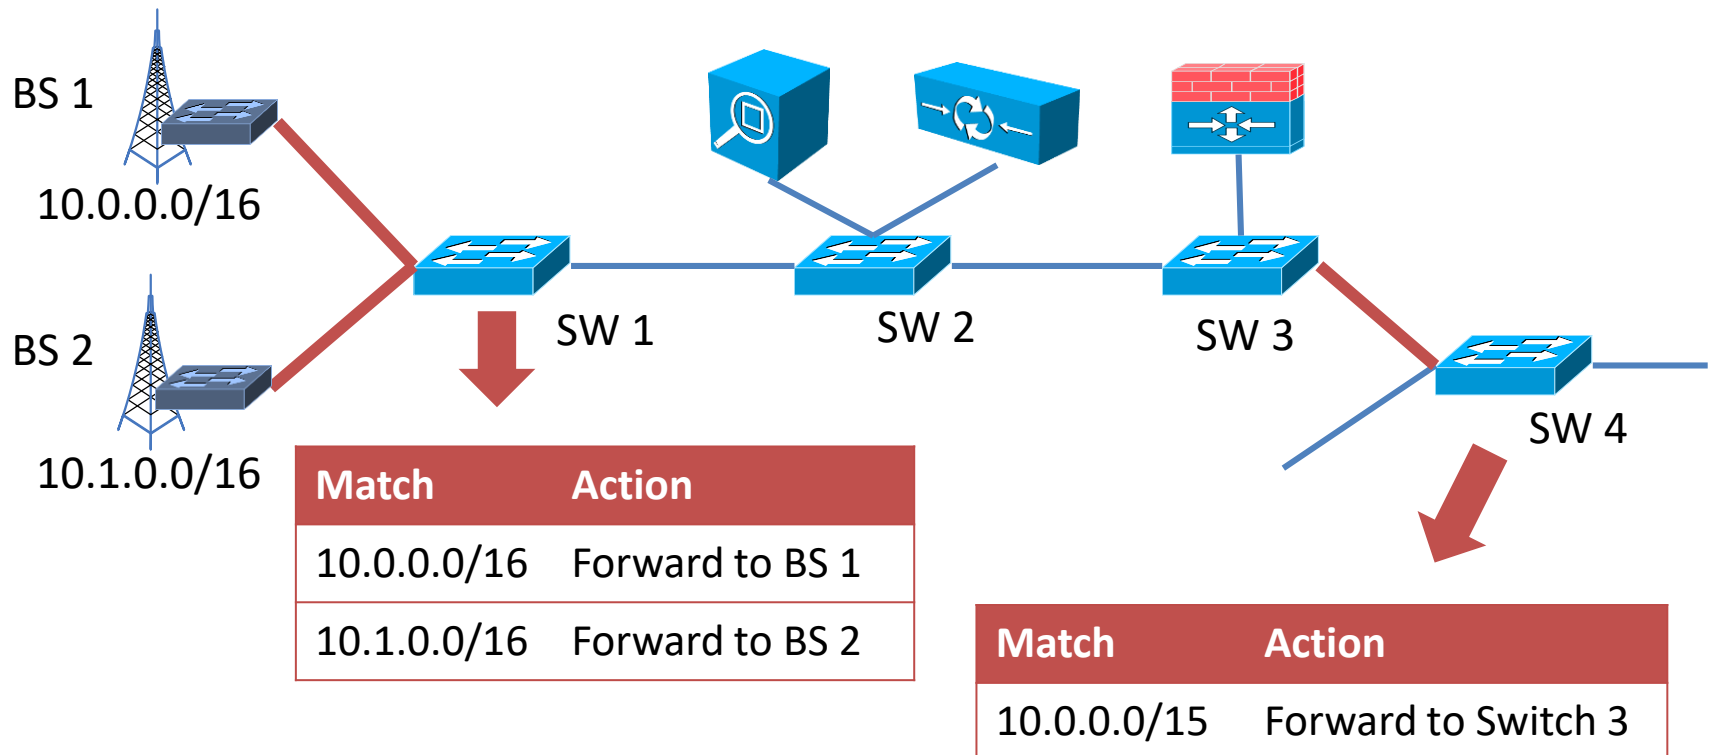


Location-Based Hierarchical IP Address




Route to different BSs with BS ID

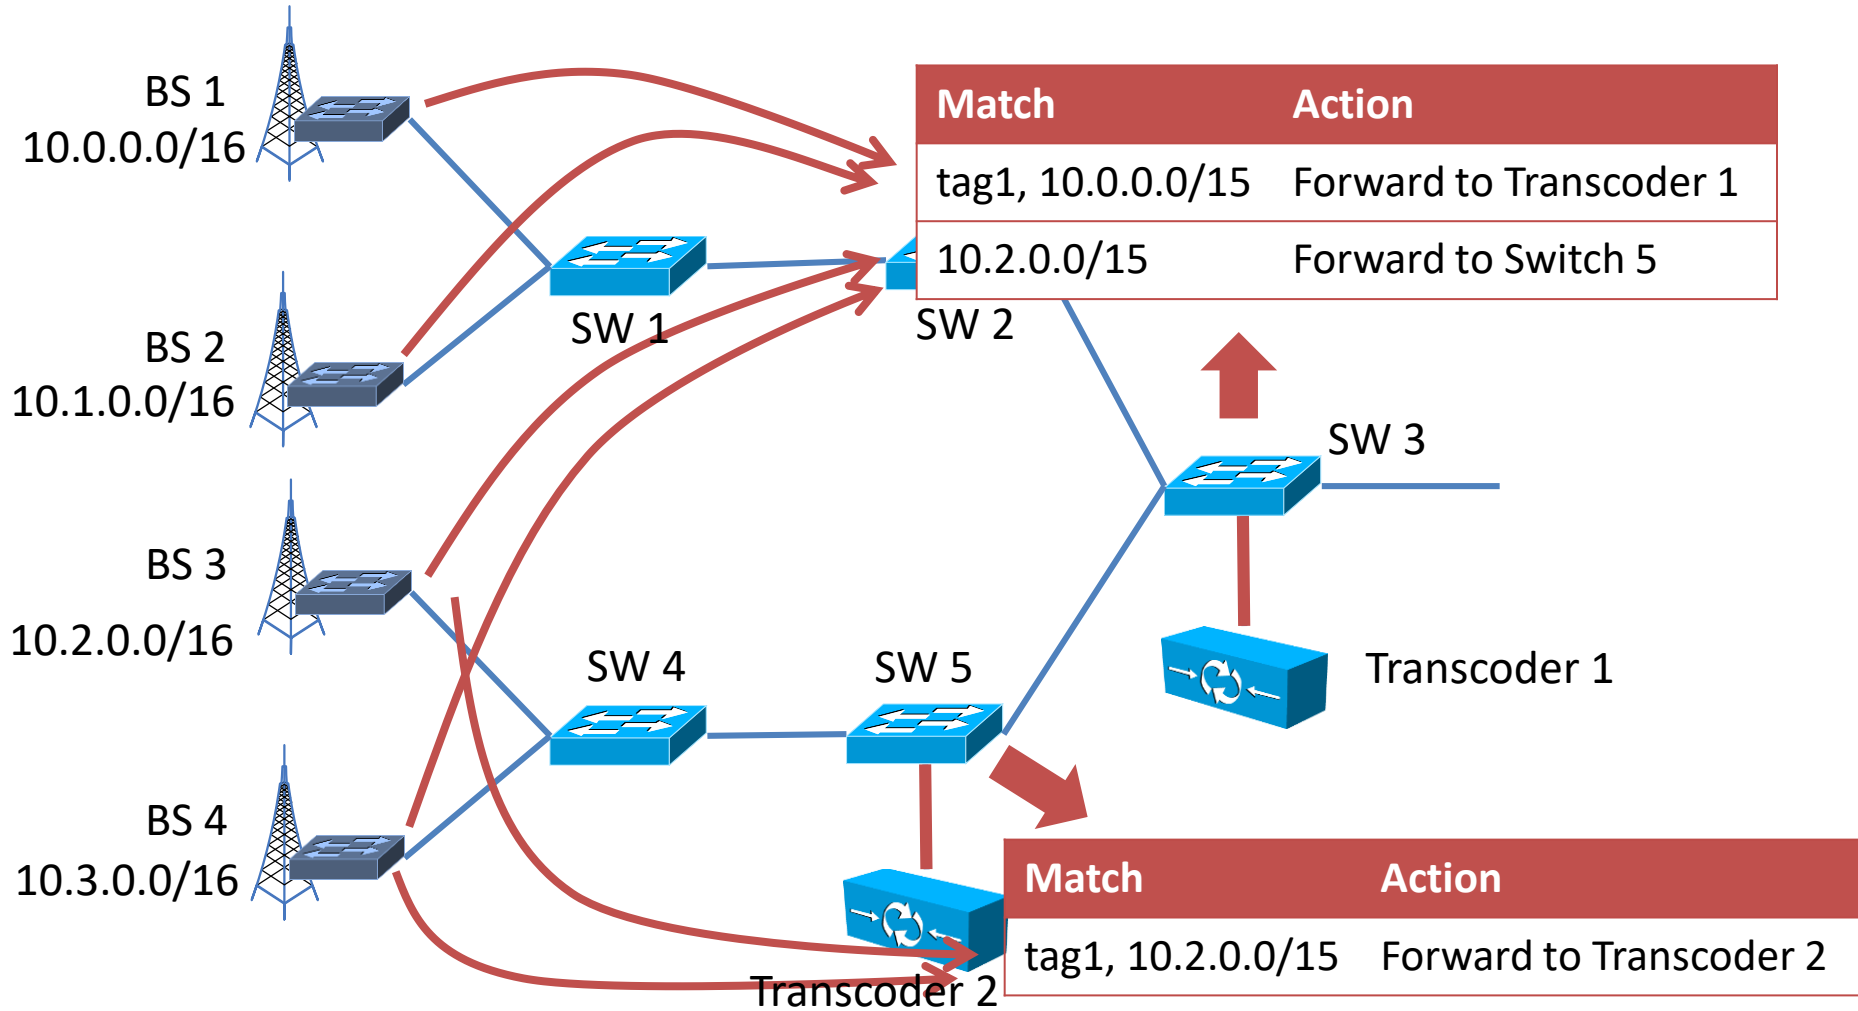
- Forward to base station with **prefix matching**
- Can **aggregate** nearby BS IDs



MB load balancing with policy tag and BS ID



MB load balancing with policy tag and BS ID




Policy Consistency

- **UE Mobility:** frequent, unplanned
- **Policy consistency:**
 - Ongoing flows traverse **the same sequence of middlebox instances**, even in the presence of UE mobility
 - Crucial for **stateful** middleboxes, e.g., stateful firewall


Policy Consistency

- An ongoing flow traverses stateful Firewall 1 before handoff
 - Use **10.0.0.7** (old IP under BS1), go via the old path
- New Flow can go via stateful Firewall 2
 - Use **10.1.0.11** (new IP under BS2), go via the new path



Multi-Dimensional Identifier Encoding

- Encode multi-dimensional identifiers to source IP and source port



- Return traffic from the Internet:
 - Identifiers are **implicitly piggybacked** in destination IP and destination port
- Commodity chipsets (e.g., Broadcom) can wildcard on these bits

Scalable Data Plane Summary


Packet classification based on service policy

Encoding results to packet headers

Traffic steering based on traffic management policy

Selectively multi-dimensional aggregation

Simple forwarding based on multi-dimensional tags



SoftCell: Taking Control of Cellular Core Networks

- Characteristics of Cellular Core Networks
- Scalable Data Plane
 - Asymmetric Edge: Packet Classification
 - Core: Multi-Dimensional Aggregation
- Scalable Control Plane
 - Hierarchical Controller

Control Plane Load

Packet classification


Handle every flow

Frequent switch update

Multi-dimensional aggregation


Handle every policy path

Infrequent switch update



Hierarchical Controller

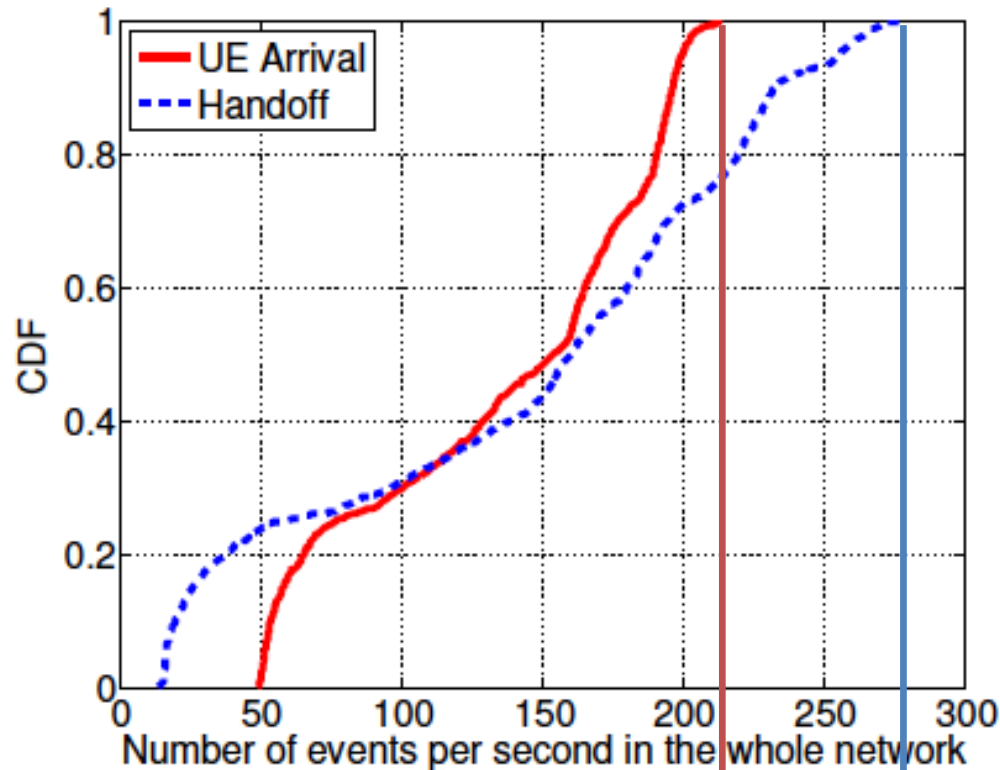
- Local agent (LA) at each base station
- **Offload** packet classification to local agents



Evaluation


- Control Plane: LTE workload characteristics
 - Dataset: 1 week traces from a large LTE network
 - ~1500 base stations, ~1 million UEs
 - Measure:
 - Network wide (Controller load): # of UE arrivals/sec, # of handoffs/sec
 - Per Base station (Local agent load): # of active UEs, # of bearer arrivals/sec
 - Compare with micro benchmark
- Data Plane: large-scale simulations

Network Wide (Controller Load)




99.999th percentile 214 UE arrivals/s 280 handoffs/s

Per Base Station (Local Agent Load)



99.999th percentile

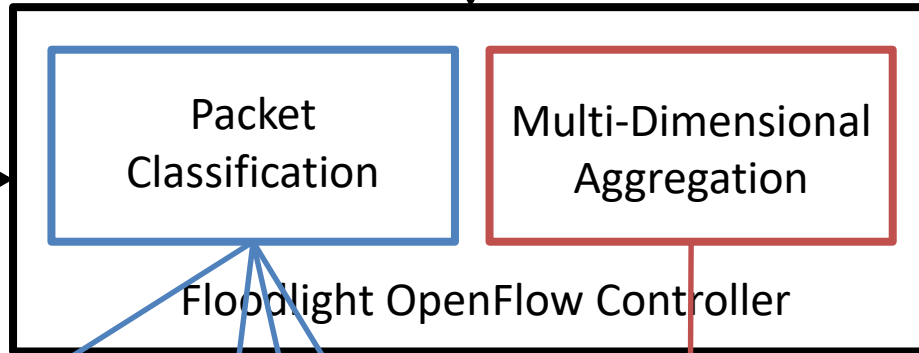
514 active UEs



34 bearer arrivals/s

Micro Benchmark

Service Policy



Subscriber Attributes



Topology

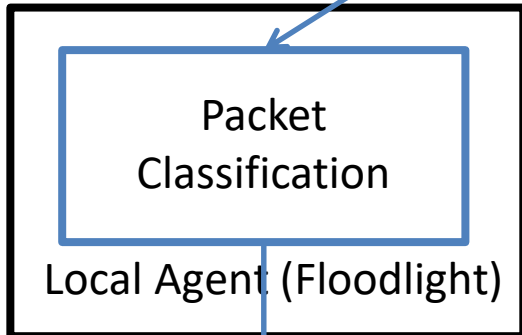


Floodlight OpenFlow Controller

Packet Classifiers

Switch Rules For Path Implementation

For topology with ~1 K BSs and ~1 K service policy clauses, ~10 ms to calculate one path. Can pre-compute.



Emulate 1000 local agents: 2.2 million requests/sec

All packet-in go to controller: 1.8 K requests/sec


All packet-in processed locally: 505.8 K requests/sec

Switch Rules For Header Rewriting

Evaluation

- Control Plane: LTE workload characteristics
- Data Plane: large-scale simulation
 - Synthesized topology [Ceragon'10]: 128 switches, 1280 base stations
 - 8 middlebox types, 10 replicas each type
 - 1000-8000 service policy clauses, traversing 4-8 MBs
 - Measure: switch flow table size (# of rules)

Flow table size vs. # of service policy clauses



Questions?