Switching Constrained Maxweight Scheduling for Wireless Networks

Soumya Basu

joint with Sanjay Shakkottai

The University of Texas at Austin, USA

Motivation

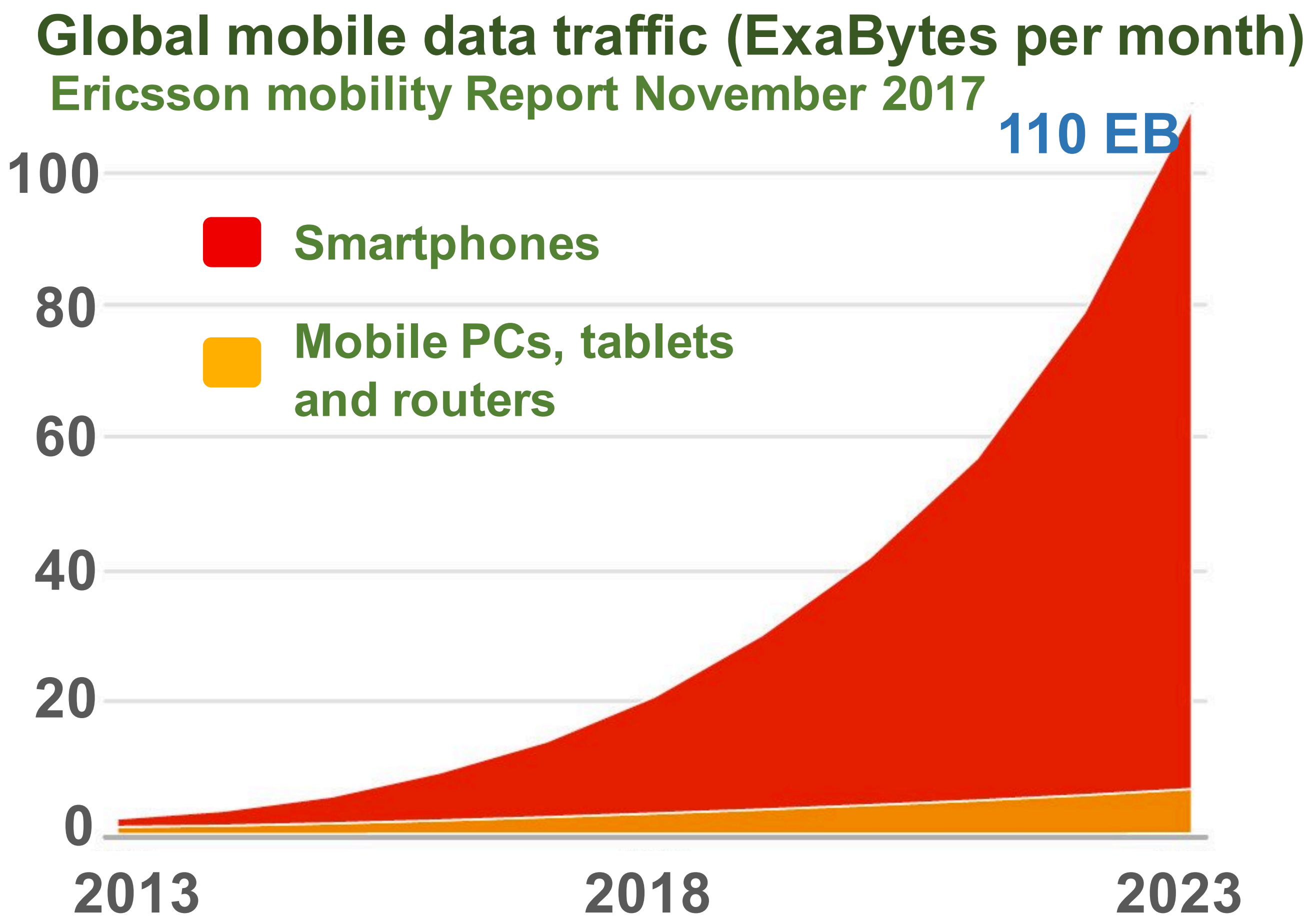

Image credit: Ericsson

Motivation

- **Dense deployment** of base stations (BS) to support peak data traffic
- **Dynamic** activation and de-activation of BS to **optimize energy** usage
- **Fast activation** dynamics is **used** to serve the incoming **data rate**
- **Fast activation** dynamics leads to **large switching overhead**,

  e.g. **hand-offs**, **state exchange** among BSs, and BS **start-up** costs

System Model

- **Downlink** time-slotted system with multiple BSs and Users
- Each BS has a separate queue for each connected user
- Cost of operation + switching: $\boldsymbol{c_1(\#Active\ BS) + c_0(\#BS\ switch)}$

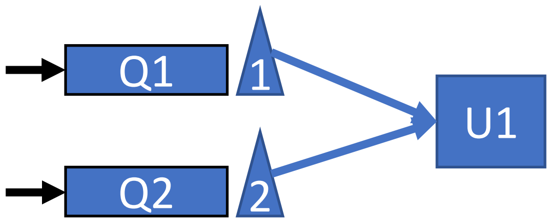

Example: 2 BS, 1 User

System Model

1. I.i.d. Arrival and Channel Realization
2. BS Activation (When to switch? What to switch to?)
3. Channel Observation from Active BS
4. Scheduling and Departure (What to Schedule?)

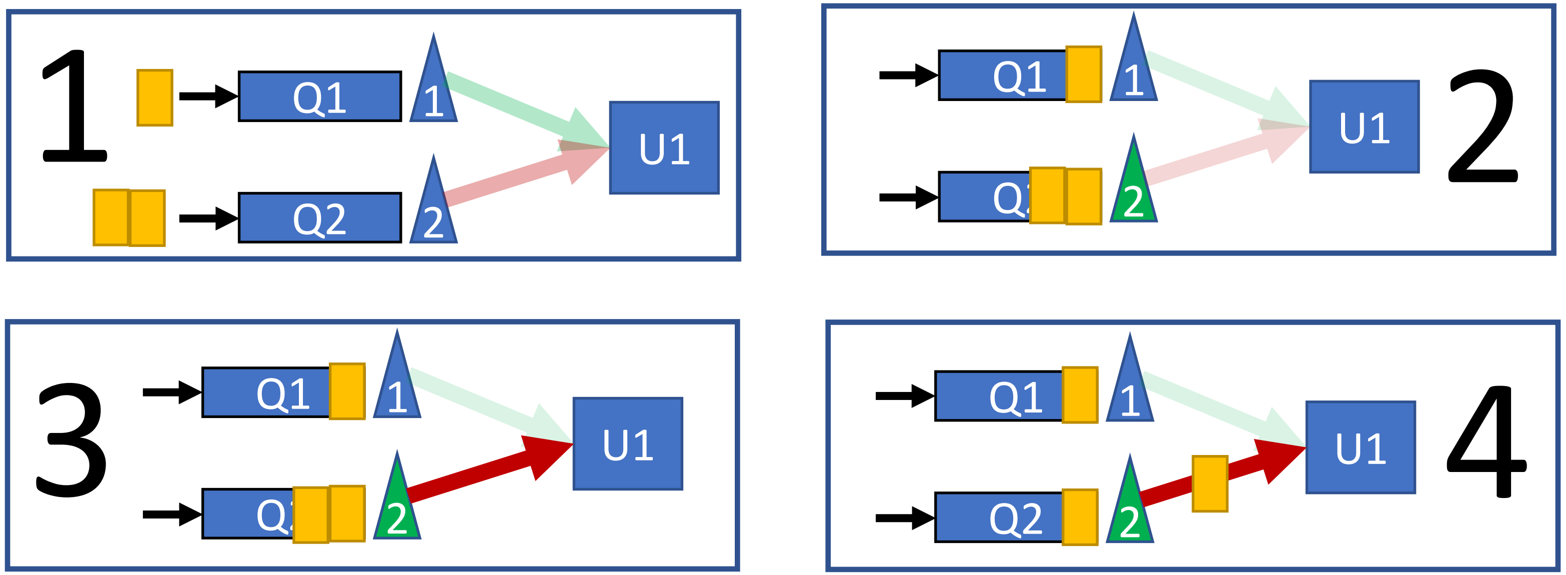

Problem Definition

- Performance Metrics: 'Asymptotic time average of expected …'
  1. Average Cost (Operational + Switching)
  2. Average Queue Length
  3. Queue Length Tail

Minimize	Average Cost
Subject to	Bounded Average Queue Length (Stability)
	Exponential decay in Queue Length Tail
	Over all Causal policies
	(A policy is **causal** iff it only depends on the **history)**

Fixed BS activation

- Capacity region for a fixed BS activation

  Set of arrival rate for which **Stability** is feasible
- What to schedule using active BSs? **[Solved]** [L. Tassiulas et al. ’92]

  Static-split among Active BS-User matching

  Max-weight Active BS-User matching
- We focus on **BS Activation and Switching**

Capacity Region

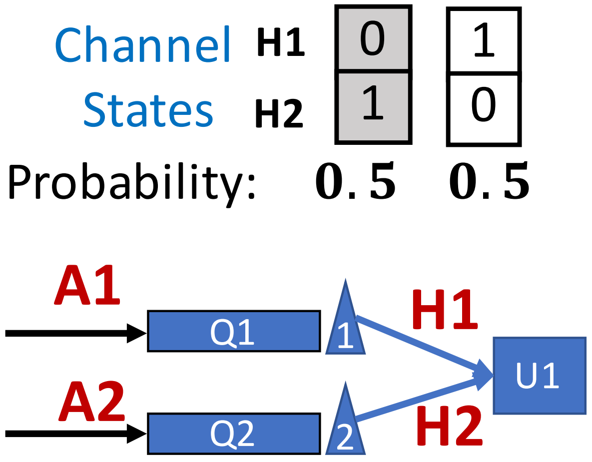

Example: 2 BS 1 User

- Capacity region of different BS activation

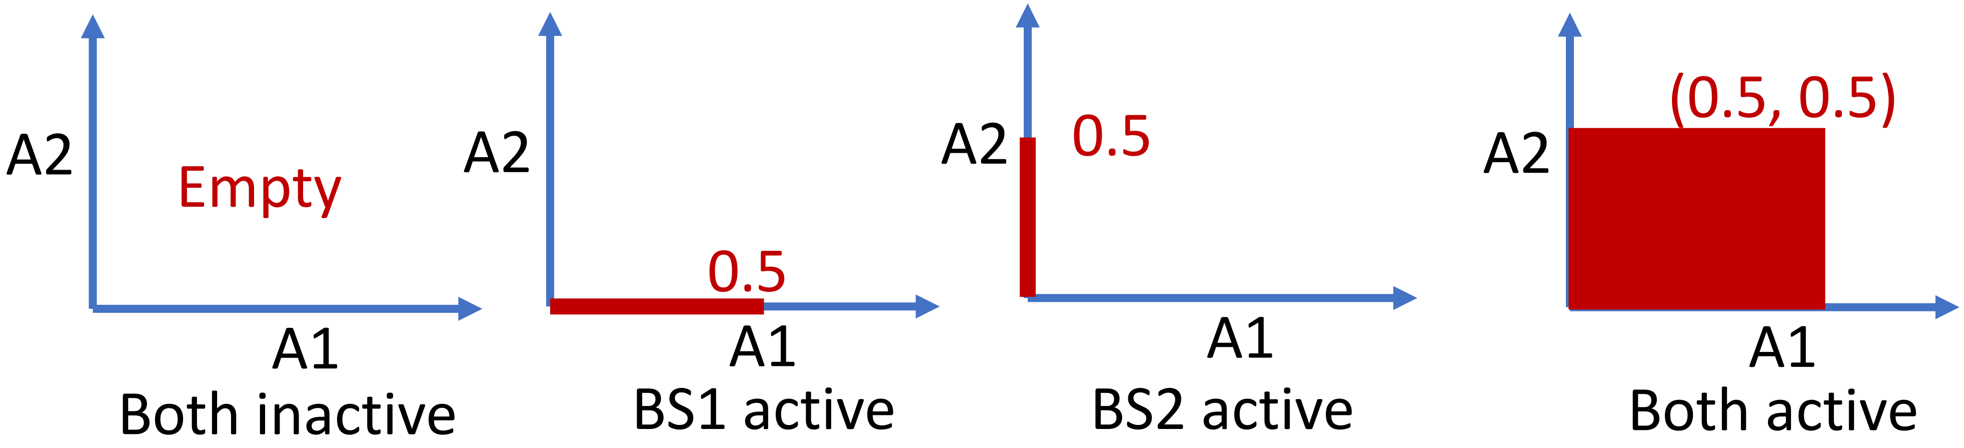

Optimal Activation

- Optimally time share between activations (I.i.d. BS activation)

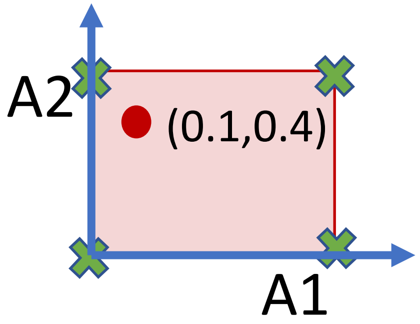

Rate	Activation Probability				Energy Cost	Switching Cost
	None	BS1	BS2	Both		
●	0.2	0	0.6	0.2	1	**0.64**

High Switching Cost

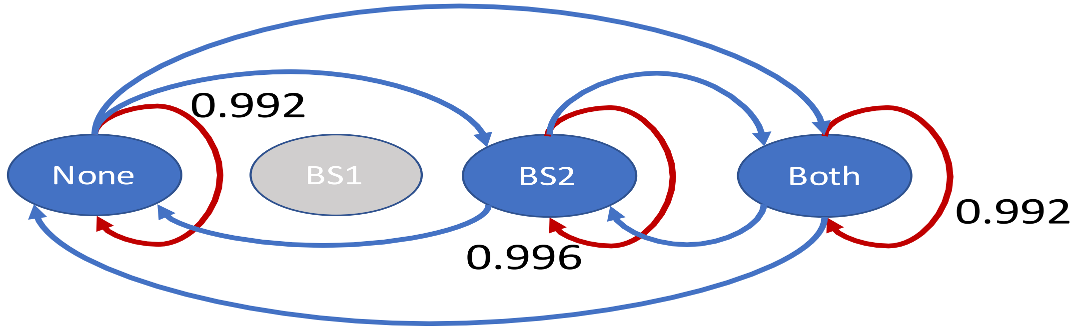

Activation Markov Chain

Energy Cost: 1 Switching Cost: 0.0064

- Switching cost driven to arbitrarily low values
- Slow Markov Activation + Max-weight Scheduling [S. Krishnasamy et al. '17]

Scope for Improvement

- Static activation is not adaptive to the queue lengths

  BS1 activated w.p. 0.2, even if Q1 Large and Q2 Small !

- Linear decay in queue length tail under Slow Markov Activation

$$\mathbb{P}(\boldsymbol{Sum\ of\ Queue\ lengths} \geq \boldsymbol{x}) \leq \frac{\boldsymbol{c}}{\boldsymbol{x}}$$

- Queue Dependent BS activation with constrained switching

- Prior works without constrained switching

  A Gopalan et al. ’07, MJ Neely et al. ’08, MJ Neely et al. ’12

Queue Dependent BS Activation

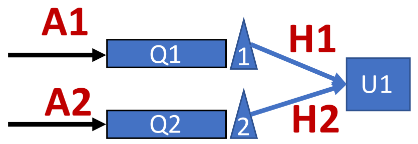

Example: 2 BS 1 User

- Prioritize service for large queues greedily
  Drift(Channel, BS activation) := Sum of (Queue Length × Departure)
- Penalize BS activation for operational cost. **Penalty** := # Active BS
- Constraint Switching between activation
- **Algorithm I (LASS-Static):**

  W.p. ϵ_{sw} Activate

  ArgMax **Expected** (Drift – **V** × Penalty)

  O/w: Stick to Previous BS activation

  V: Penalty Scale

  ϵ_{sw} : Switching Rate

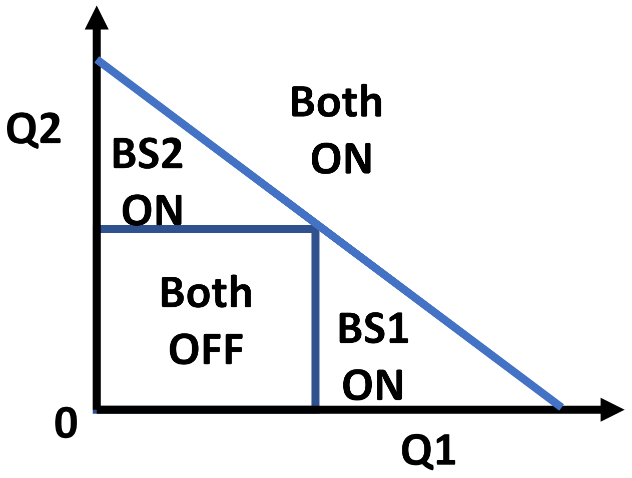

BS Activation Function

Drawback of Static Switching

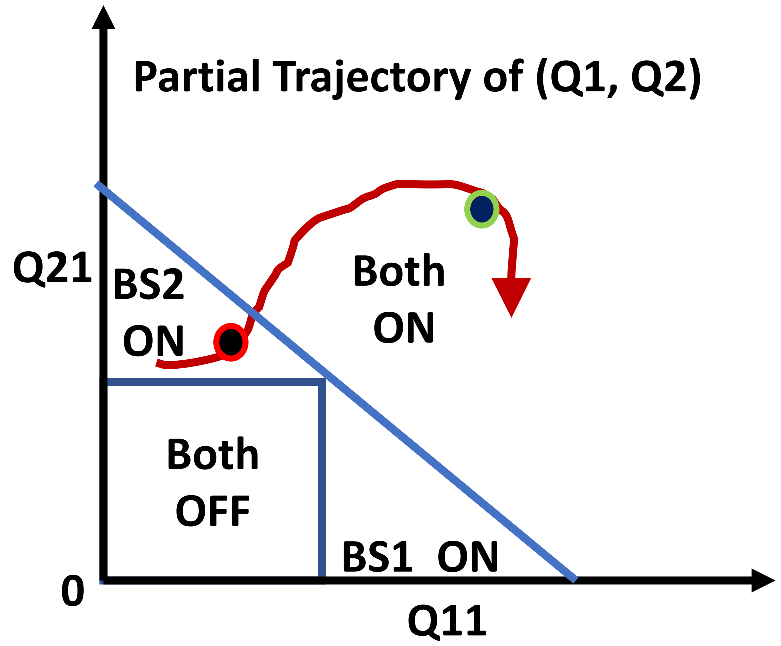

- We start with BS2 ON
- In points and switching of BSs is allowed
- In **Previous BS Subset = Best BS Subset**
  No switching even though it is allowed
  Switching resource/opportunity is wasted
- In **Previous BS Subset $\neq$ Best BS Subset**
  We switch to Both ON from BS2 ON

Dynamic Switching

Algorithm II (LASS-Dynamic): Only differs in 'When to Switch?'

Virtual Objects: Switch Queue and **Switch Counter**

- W.p. ϵ_{sw} **remove one packet** from Switch Queue [Black dots in the plot]
- Current BS activation is not optimal: **Increment** Switch Counter [2 to 3]

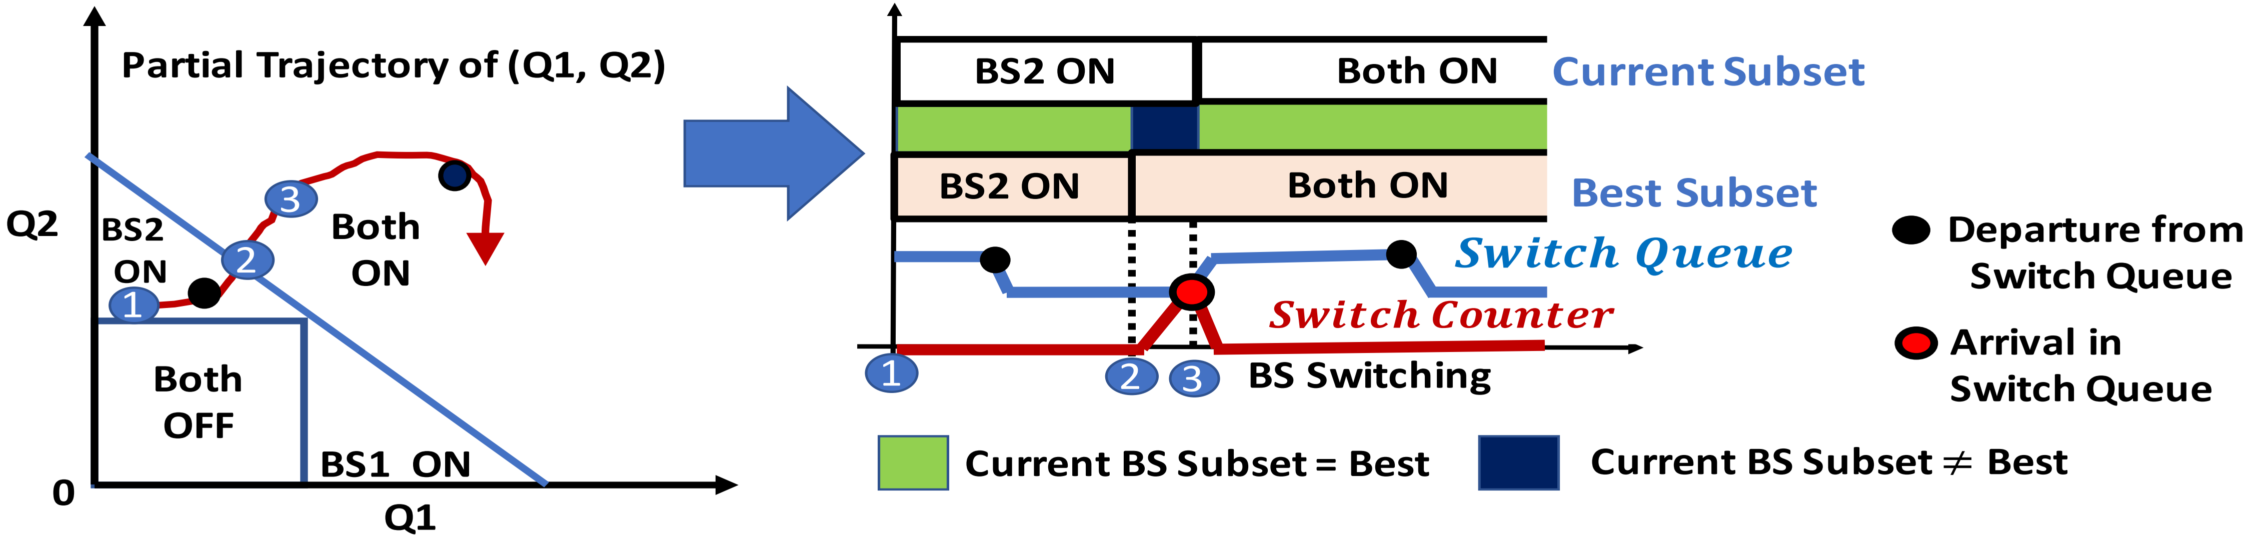

Dynamic Switching

- BS switching when **Switch Counter** $\geq$ **Switch Queue** [3: Red Dot]
- BS switching: 1) **Reset** Switch Counter
  2) **Add packet** to Switch Queue

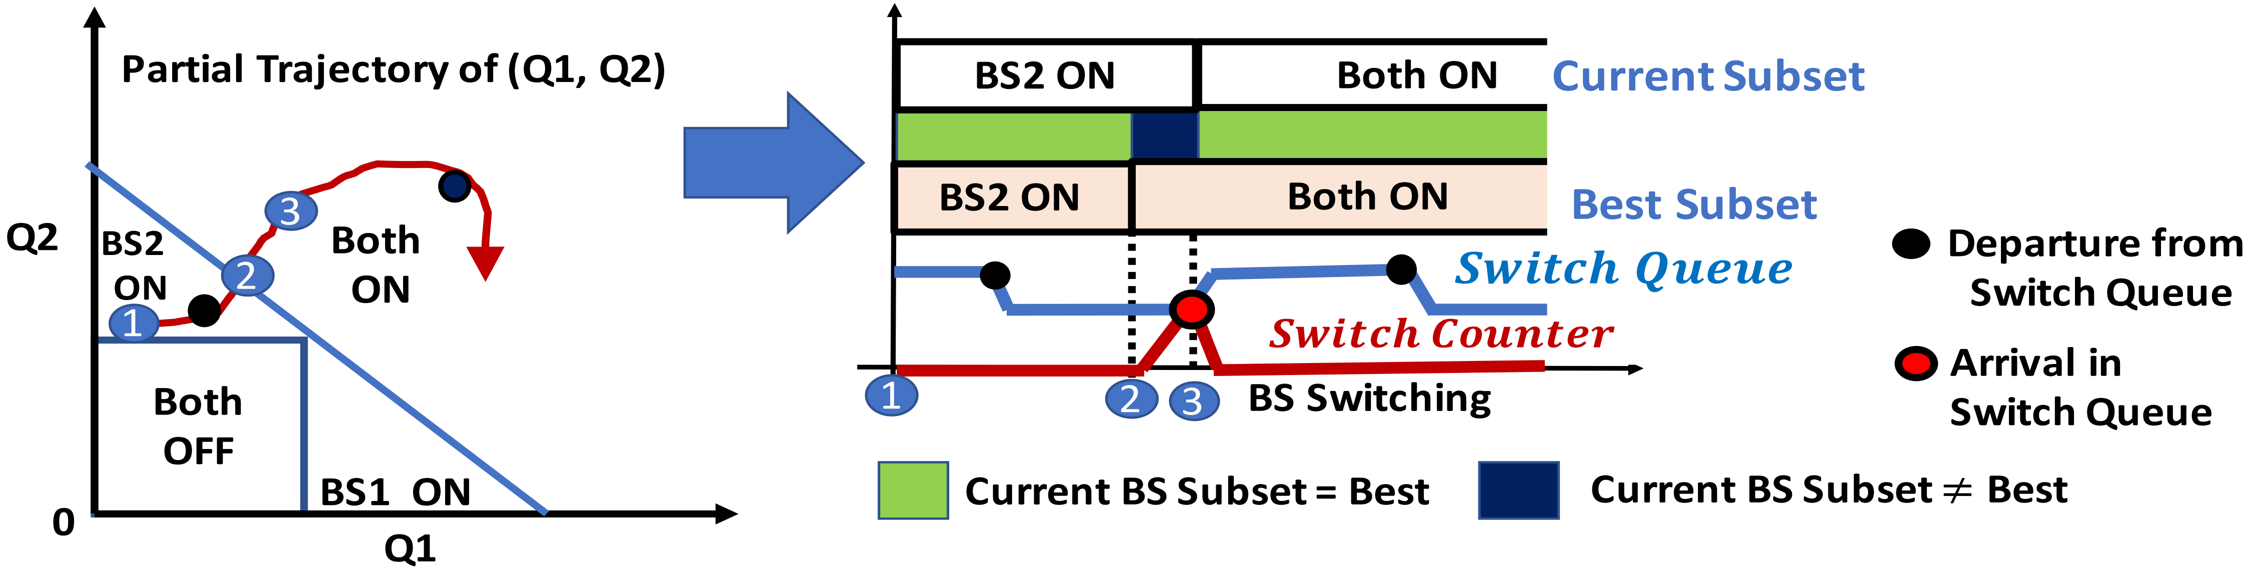

Main Results

Parameters

- # BS: **N** # Users: **M** Capacity gap: $\epsilon_g > 0$
- Queue Length at time **t**: $Q(t)$
- **Optimal cost without switching cost**: C^*_{avg}
- Switching rate: ϵ_{sw}, Penalty scale: V **(Tuning Knobs)**

Time Average bounds

Both LASS-Static and LASS-Dynamic

$$Q_{avg} \leq O\left(\frac{C^*_{avg}}{\epsilon_g} + V + \frac{NM}{\epsilon_g \epsilon_{sw}}\right) \qquad C_{avg} \leq C^*_{avg} + O\left(\epsilon_{sw} + \frac{NM}{V \epsilon_{sw}}\right)$$

$$V \uparrow, \epsilon_{sw} \downarrow \Rightarrow Q_{avg} \uparrow, (C_{avg} - C^*_{avg}) \downarrow$$

Main Results

Tails Bounds

For large enough x and all time t

- For **LASS Static**: $\mathbb{P}(|\boldsymbol{Q(t)}| \geq \boldsymbol{x}) \leq \mathbf{exp}\left(-\boldsymbol{\Theta}\left(\boldsymbol{\epsilon_{sw}\epsilon_g}\right)\boldsymbol{x}\right) + \boldsymbol{O}\left(\frac{\mathbf{log}(\boldsymbol{t})}{\boldsymbol{t}}\right)$
- For **LASS Dynamic**: $\mathbb{P}(|\boldsymbol{Q(t)}| \geq \boldsymbol{x}) \leq \mathbf{exp}\left(-\boldsymbol{\Theta}\left(\boldsymbol{\epsilon_g}\right)\boldsymbol{x}\right) + \boldsymbol{O}\left(\frac{\mathbf{log}(\boldsymbol{t})}{\boldsymbol{t}}\right)$

Decay rate of LASS Static Depends on $\boldsymbol{\epsilon_{sw}}$

Simulation Results

- Three algorithms for **8 Users and 3 BSs** simulated until convergence
  - **DP**: Drift + Penalty (Baseline with NO Switching Cost)
  - **LSG**: LASS Static
  - **LD:** LASS Dynamic

Simulation Results

- **First Plot:** $\mathbf{Q}_{avg}$ of **DP < LD < LSG** (V = 100, load = 0.9)
- **Second plot:** C_{avg} of **DP > LD ≈ LSG** (V = 100, ϵ_{sw} = 0.1)

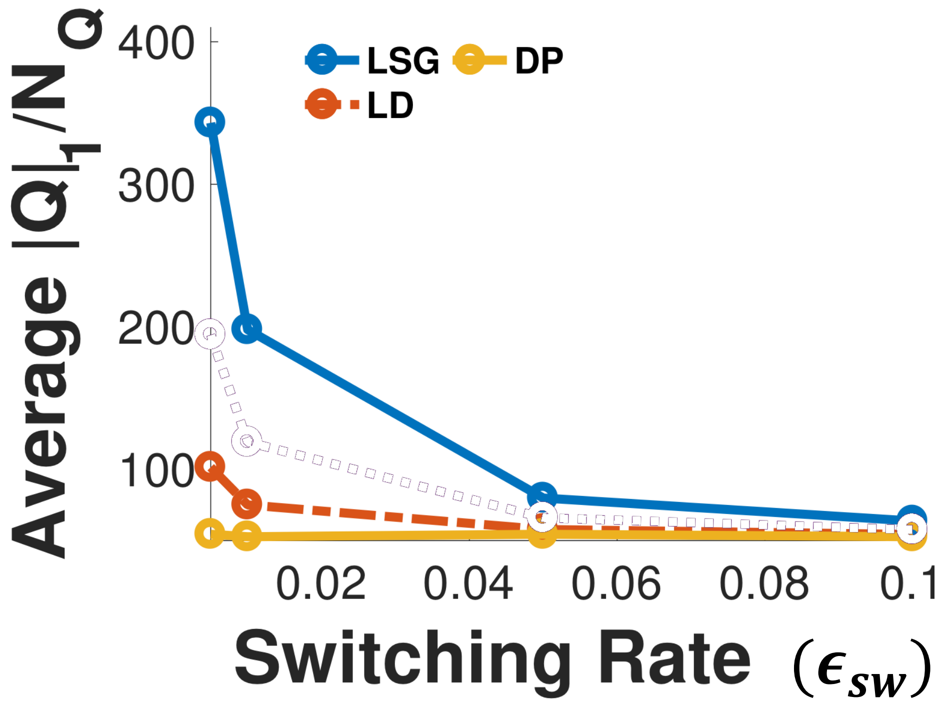

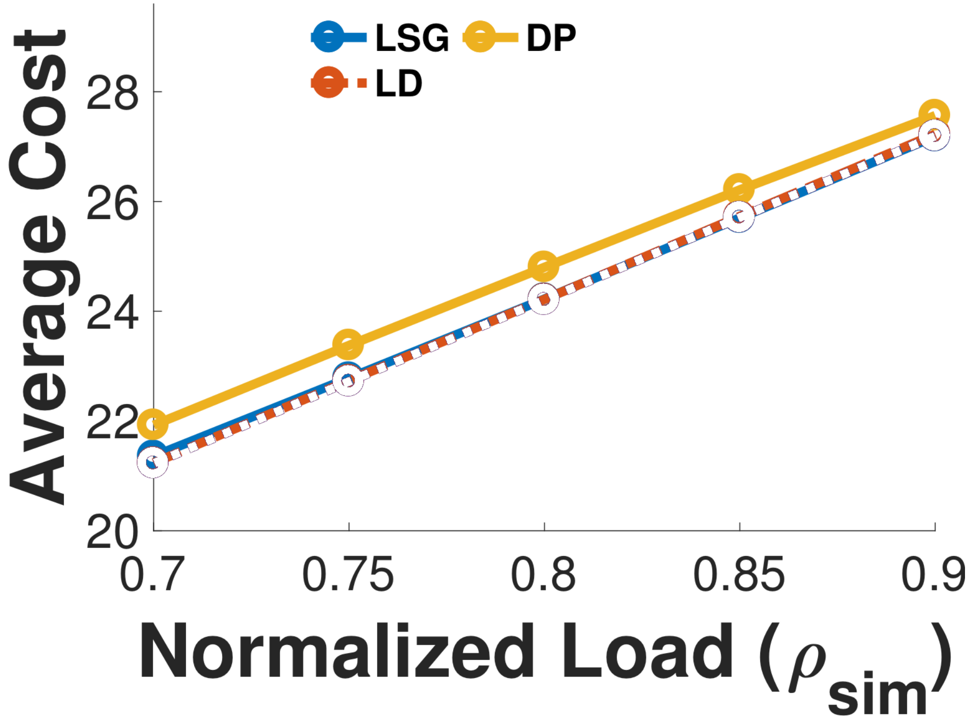

Simulation Results

- Separation of queue length tail distribution
  - **DP < LD << LSG** (V = 100, load = 0.9)
  - Differences are more pronounced for smaller ϵ_{sw}

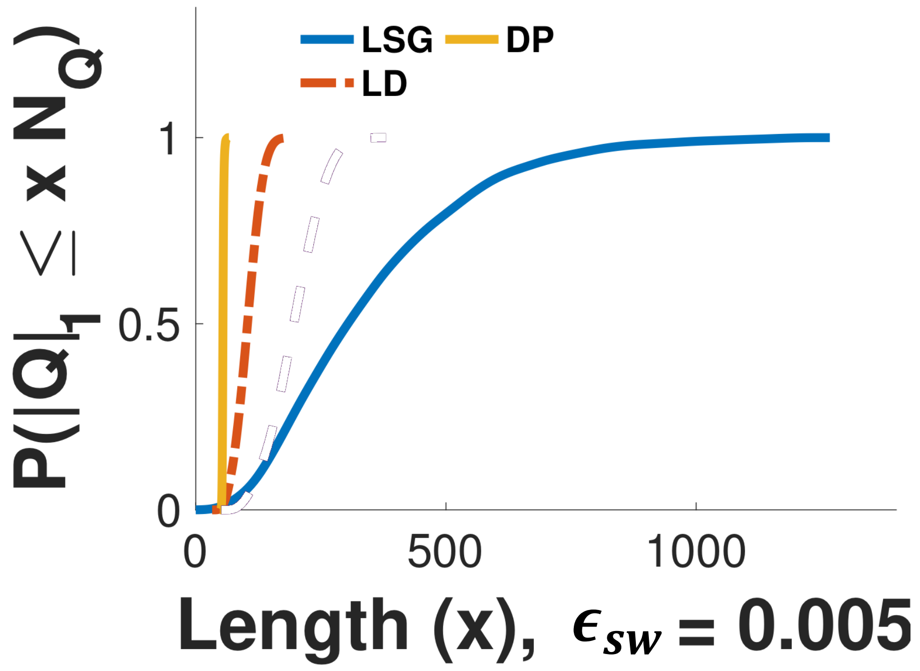

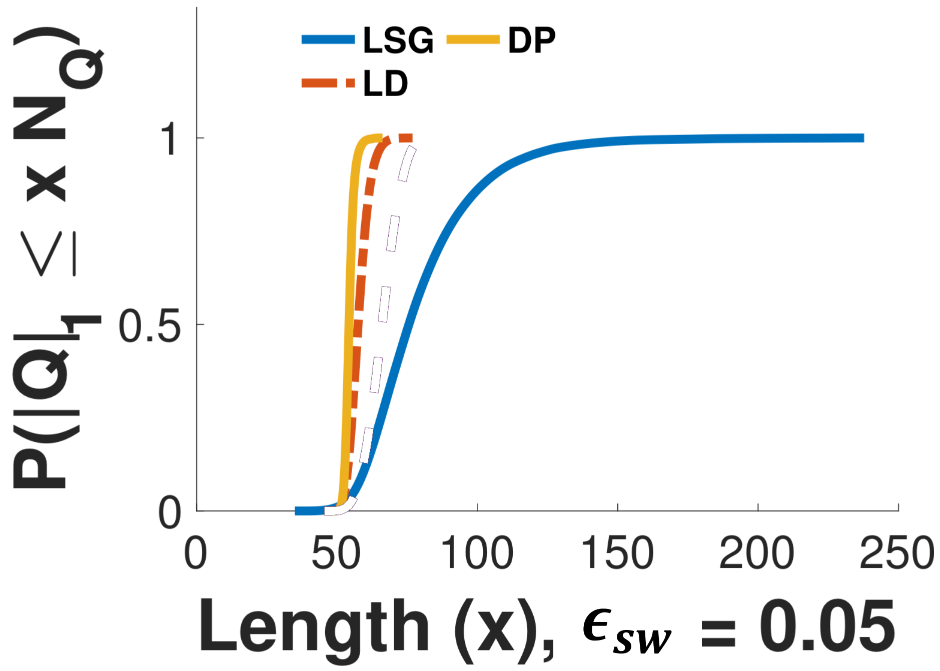

Thanks!

Step I: Arrival and Channel Realization

- Arrival and Channel process
  - I.i.d. across time slots and possibly correlated in a time slot

Arrival States

A1	1	0	0
A2	0	1	0
Probability:	**0.3**	**0.3**	**0.4**

Channel States

H1	0	1
H2	1	0
Probability:	**0.5**	**0.5**

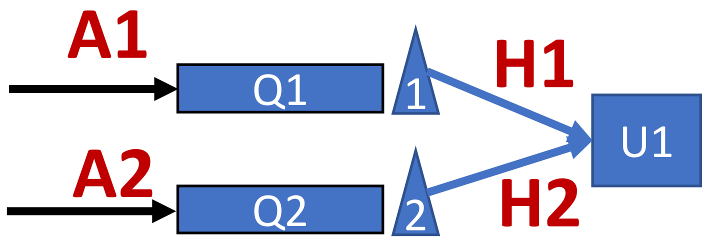

Example: 2 BS 1 User

Step II: Base Station Activation

- Activate a subset of BSs
- Cost of operation + switching at time $\boldsymbol{t}$

$$C(t) = c_1(\#Active\ BS) + c_0(\#BS\ switch)$$

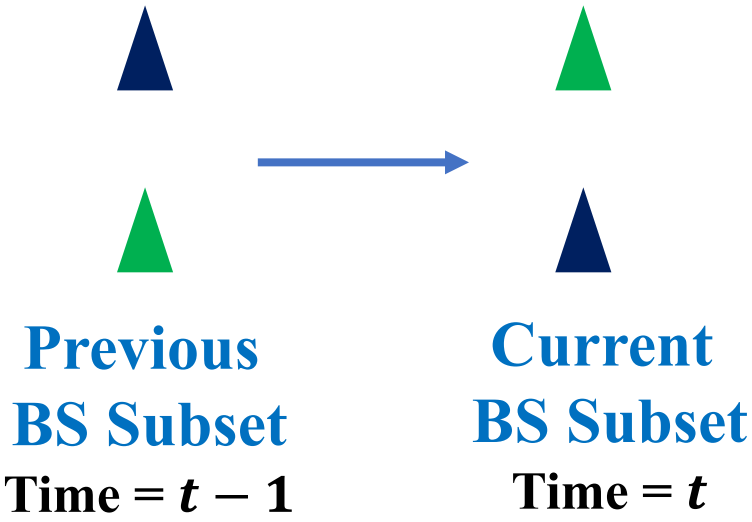

Green: Active
Blue: Inactive

$$C(t) = c_1 + 2c_0$$

Step III: Channel Observation

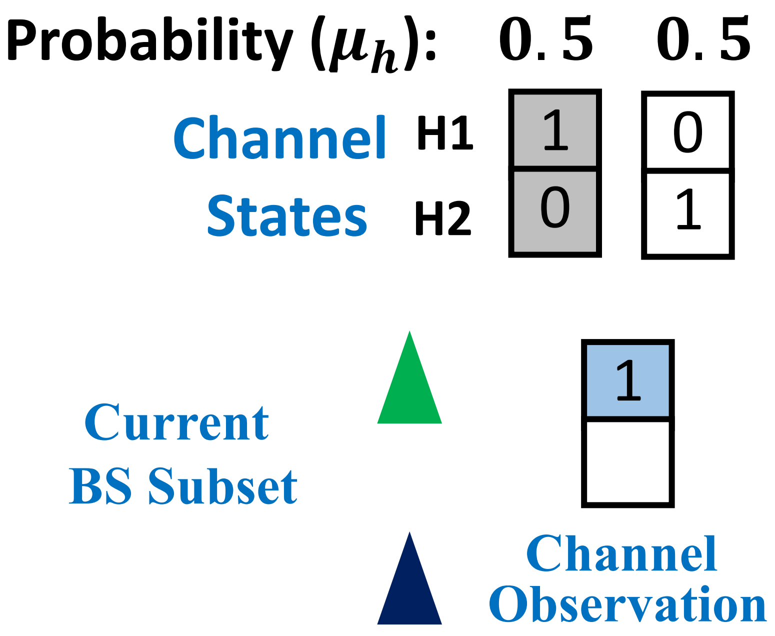

- **Observe channel** after activation
- **Why?** Probing channel requires energy

Step IV: Scheduling

'Active BS'- User matching

- Each user can connect to at most one BS
- Each BS can connect to at most one user

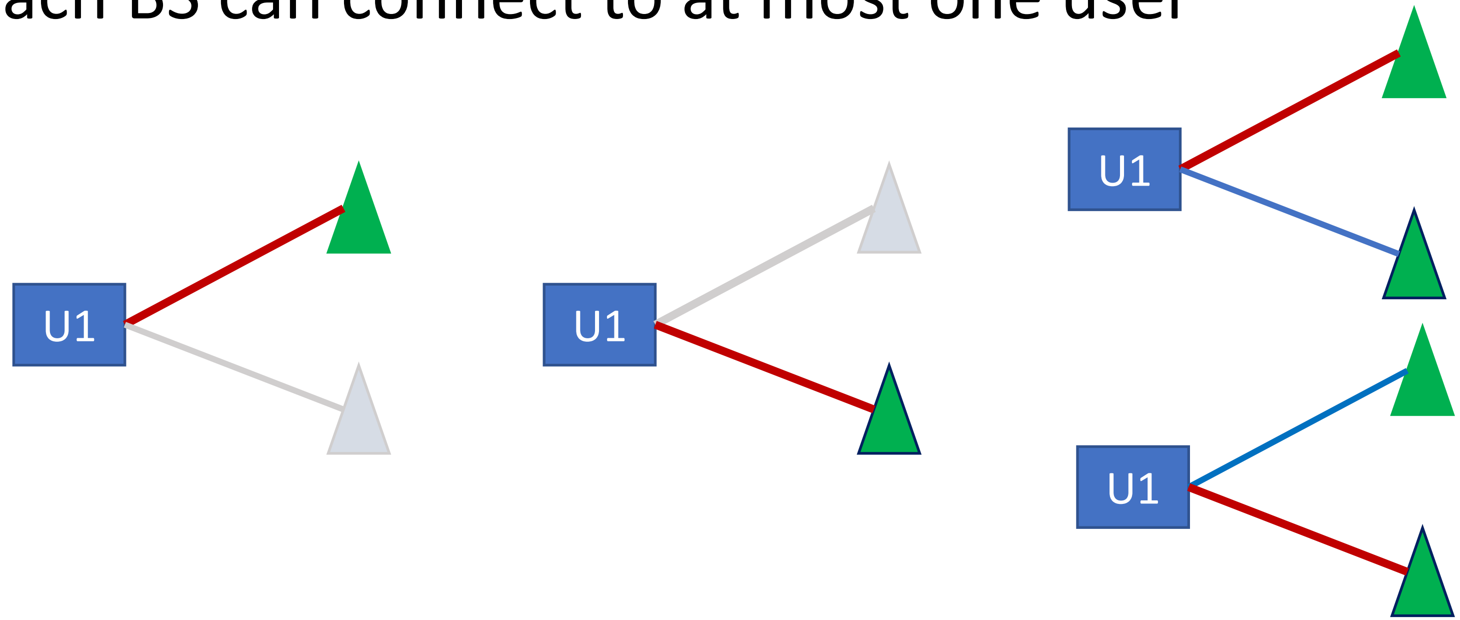

Switching Constrained Max-weight Scheduling

Objective

- **Minimize cost subject to exponential decay**

Minimize C_{avg}^{ϕ}

Subject to:

ϕ is a causal policy

Exponential Decay: $\exists\ c > 0, \forall\ t$, large x

$$\mathbb{P}^{\phi}(|Q(t)|_1 \geq x) \leq exp(-cx)$$

System is stable : $Q_{avg}^{\phi} < \infty$

Exponential Decay implies Stability

Switching Constrained Max-weight Scheduling

How to Schedule?

Edge Weight (BS n, User m):
If BS **n** is active: $\boldsymbol{Q_{nm}(t)H_{nm}(t)}$
Otherwise: **0**

Schedule the Max Weight Matching

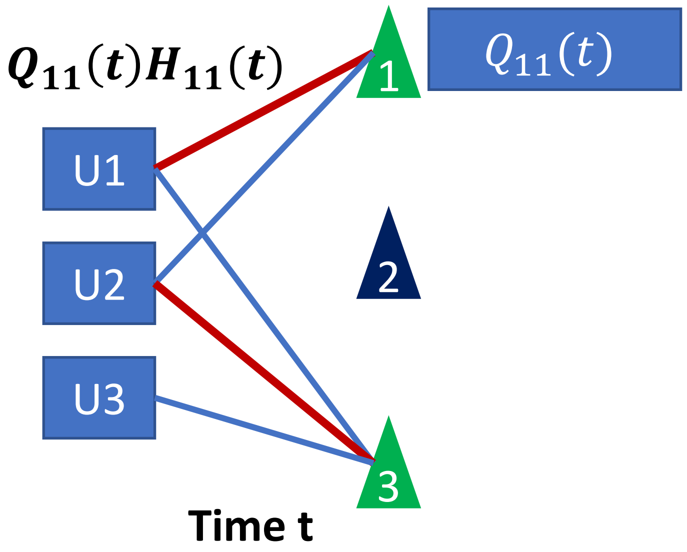

Switching Constrained Max-weight Scheduling
What to Switch to?

Drift+Penalty Method

Best BS Subset maximizes (Expected Weight($\boldsymbol{J}$) – **V|$\boldsymbol{J}$|)**

Expected Weight($\boldsymbol{J}$) = $\sum_{h} \hat{\mu}_h MW_h(\boldsymbol{J})$

$\boldsymbol{J}$: A BS subset $\hat{\mu}_h$: Channel Probability Estimates

$MW_h(\boldsymbol{J})$: Value of Max-weight matching

1) BS Subset $\boldsymbol{J}$ is active
2) Channel state **h** occurs

Switching Constrained Max-weight Scheduling

What to Switch to?

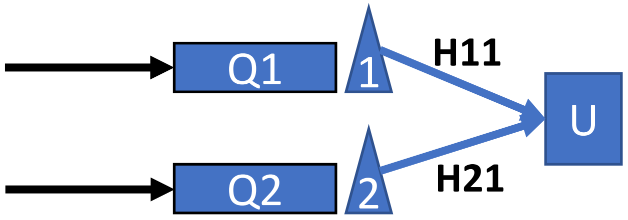

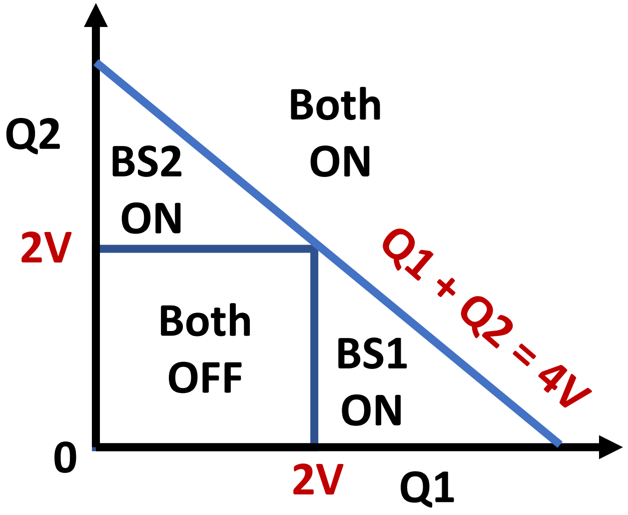

Best BS Subset Vs Queue Lengths

Channel States (CS)		CS 1	CS 2
	H11	1	0
	H21	0	1
Probability:		**0.5**	**0.5**

BS Subset	MaxWeight CS 1	MaxWeight CS 2	Penalty	Total
Both OFF	0	0	0	0
BS1 ON	Q1	0	V	0.5Q1 - V
BS2 ON	0	Q2	V	0.5Q2 - V
Both ON	Q1	Q2	2V	0.5(Q1+Q2)- 2V

Switching Constrained Max-weight Scheduling

What to Switch to?

Argmax over BS Subsets

Channel Estimates

Queue Length

$$J^*(t) = argmax_j \boxed{\sum_h \hat{\mu}_h \, max_{m(j)} \langle Q(t), d(h,m) \rangle} - V c_1 |j|$$

Max over 'Active BS'-User Matching.

Penalty: V x Operational Cost

Drift: Expected Sum of (Queue Length X Departure)

$J^*(t)$: Best BS Subset. The one to switch to.

Switching Constrained Max-weight Scheduling

Drift Equation

$$\mathrm{E}[\mathrm{Drift(t)}|\Psi(t)] \\ \leq -\epsilon_g |Q(t)|_1 + C + \boxed{c'|Q(t)|_1 |\mu_h - \widehat{\mu_h}(t)|_\infty} + \boxed{c''T(t)}$$

Learning Error

Switching Constraint

Use Lyapunov Function: $|Q(t)|_1^2 + \mathrm{T(t)}$