Dynamic Power Allocation and Routing for Time Varying Wireless Networks

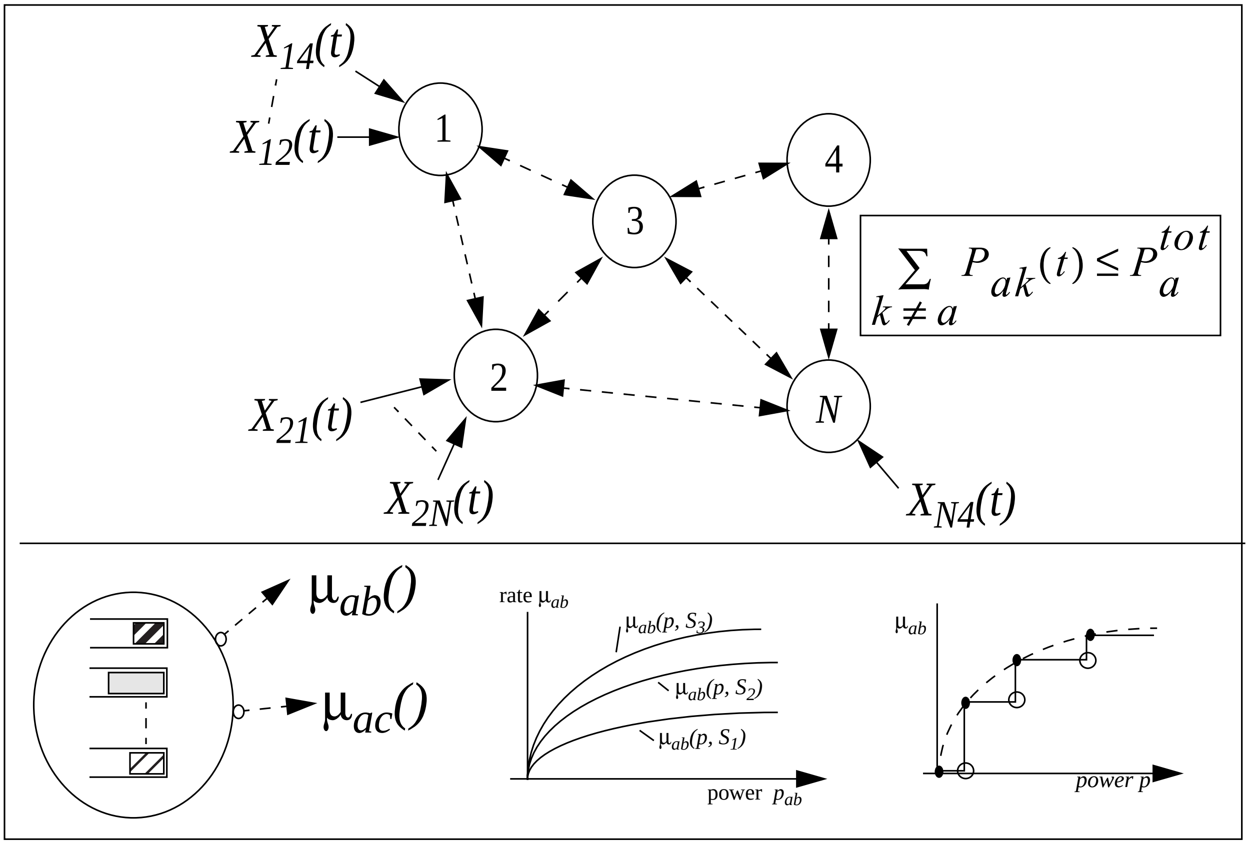

Michael J. Neely
Ph.D. Candidate
MIT -- LIDS
mjneely@mit.edu
http://web.mit.edu/mjneely/www

Network Optimization:

Kelly, Maulloo, Tan -- Journ. of Op Res. 1998

Marbach -- INFOCOM 2002

Xiao, Johansson, Boyd -- Allerton 2001

Network Control:

Elbatt, Ephremides -- INFOCOM 2002

Cruz, Santhanam -- Allerton 2002

Tassiulas and Ephremides -- Trans. Auto. Contr. '92

Network Capacity:

Gupta, Kumar -- Trans. Inf. Th. 2000

Grossglauser and Tse -- TON 2002

General Problem Formulation:
Joint Routing, Power Allocation, and Scheduling

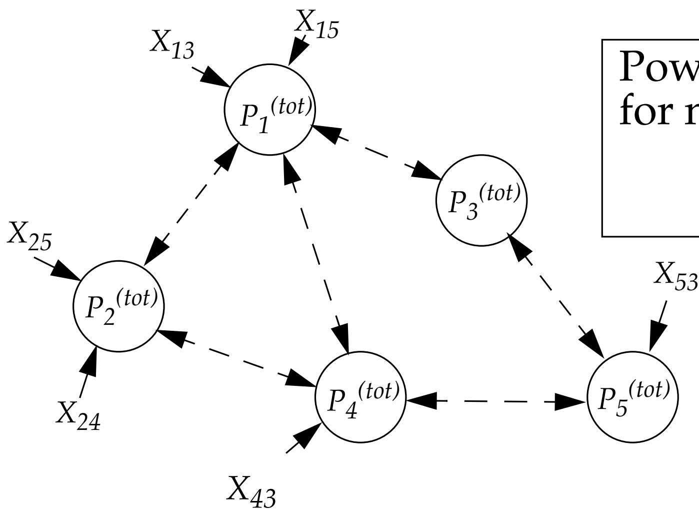

Power Constraint for node *i*:

$$\sum_j P_{ij}(t) \leq P_i^{(tot)}$$

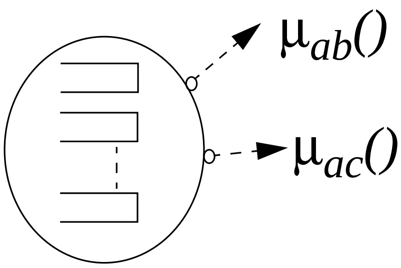

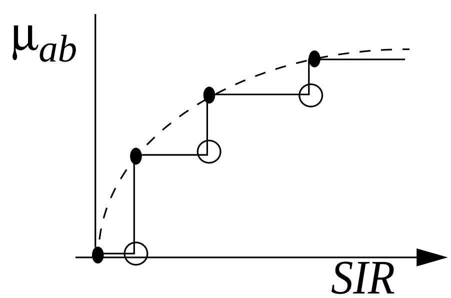

Assumptions:

-Random Traffic (Markov Modulated, bursty, etc.)

-Slotted Time with slot length *T*.

-Time Varying Channel States $\underline{C}(t) = (C_{ij}(t))$

Steady State Channel Probability $\pi_{\underline{C}}$

-Rate Function: $\mu_{ij}(\underline{P}(t), \underline{C}(t))$ (perhaps discontinuous)

-Power Constraint: $\underline{P}(t) \in \Pi$

The rate function μ*()*:

Can model a wired link with fixed capacity:

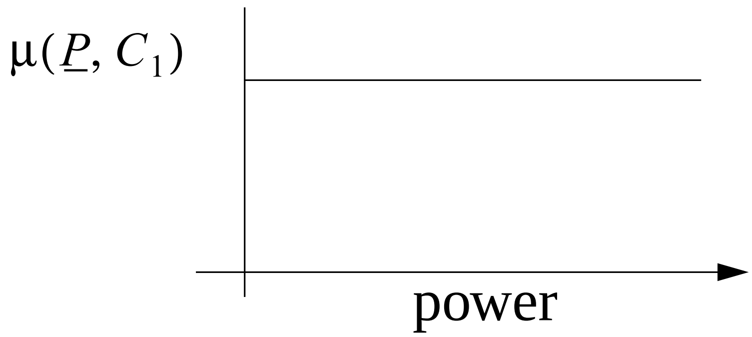

Or a broken link:

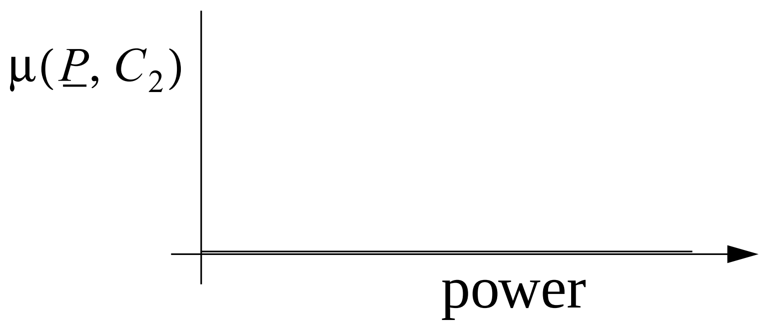

Or Server Allocation:

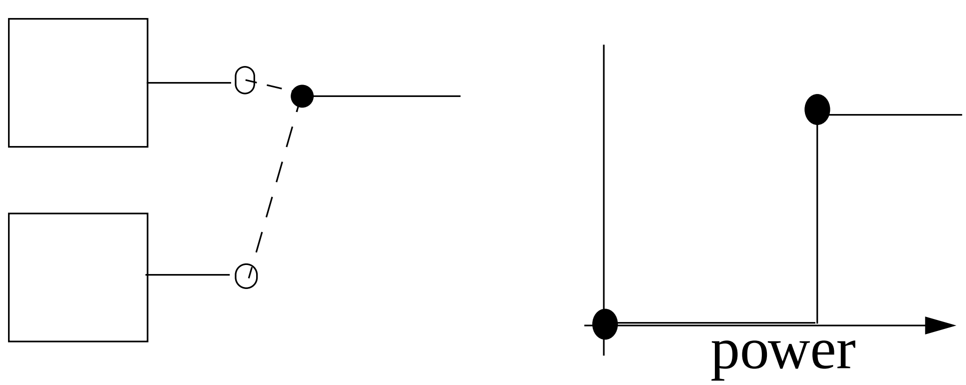

The Dynamic Control Problem:

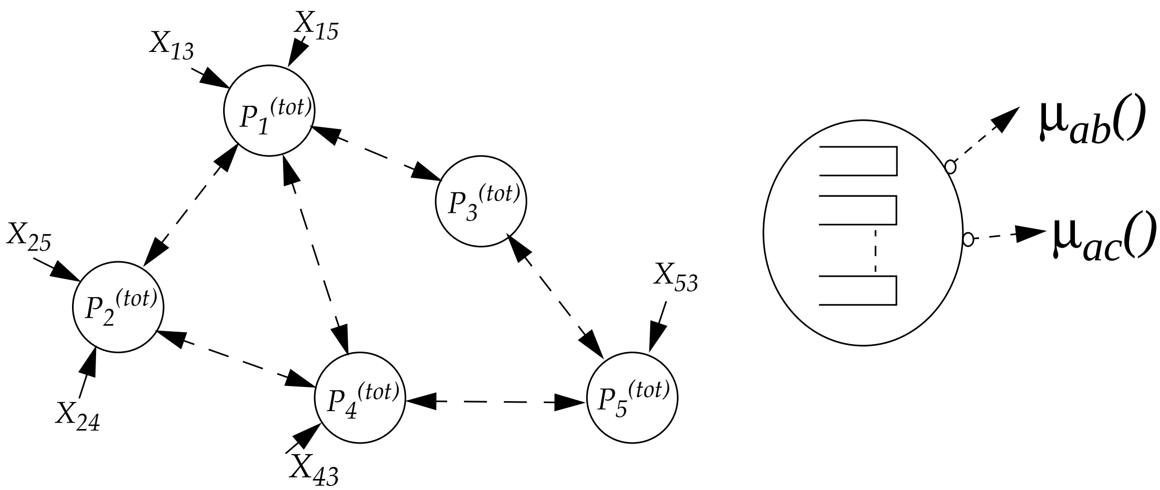

Every Timeslot, Observe:

$\underline{U}(t) = (U_{ij}(t))$ (Unfinished work Matrix)
$\underline{C}(t) = (C_{ij}(t))$ (Channel State Matrix)

Network Controller Decides:
-Power Allocation $\underline{P}(t)$ ----> $\underline{\mu}(\underline{P}(t), \underline{C}(t))$
-Routing directions for next hop

Goal: Achieve Network Capacity with low delay

(Decentralized version: View from a single node)

What is an optimal, capacity achieving strategy?

<u>Example Problem</u>: Data sources X_{1i}, X_{2j}, X_{3k}.
Destinations *i*, *j*, *k*. Two Intermediate stations.

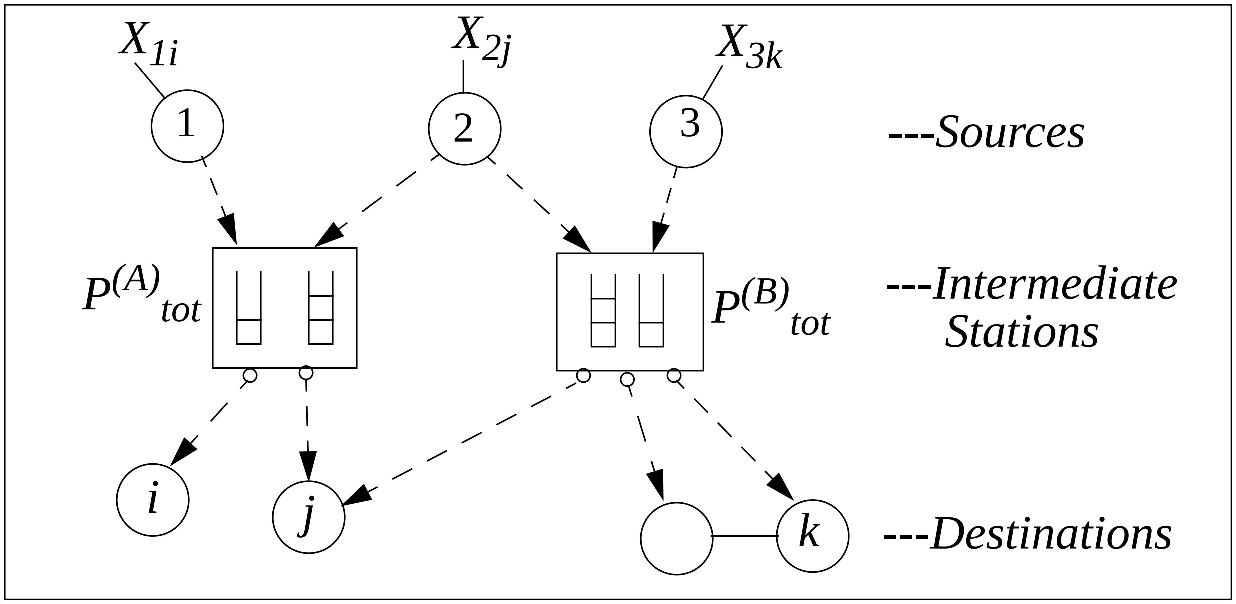

Information known at slot *t*:

Channel states $\vec{C}(t)$ and the queue backlogs $\vec{U}(t)$:

Channel state: $\vec{C}(t) = (C_{Ai}(t), C_{Aj}(t), C_{Bj}(t), C_{Bk}(t), C_{Bz}(t))$

Unfinished Work Backlog: $\vec{U}(t) = (U_{Ai}(t), U_{Aj}(t), U_{Bj}(t), U_{Bk}(t))$

<u>*Routing*</u>: In which station do we put packets from source 2?

<u>*Power Allocation*</u>: For all time, we are constrained so that:

$$p_{Ai}(t) + p_{Aj}(t) \le P^{(A)}_{tot}$$
$$p_{Bj}(t) + p_{Bz}(t) + p_{Bk}(t) \le P^{(B)}_{tot}$$

What is the <u>Capacity region</u> of the system?

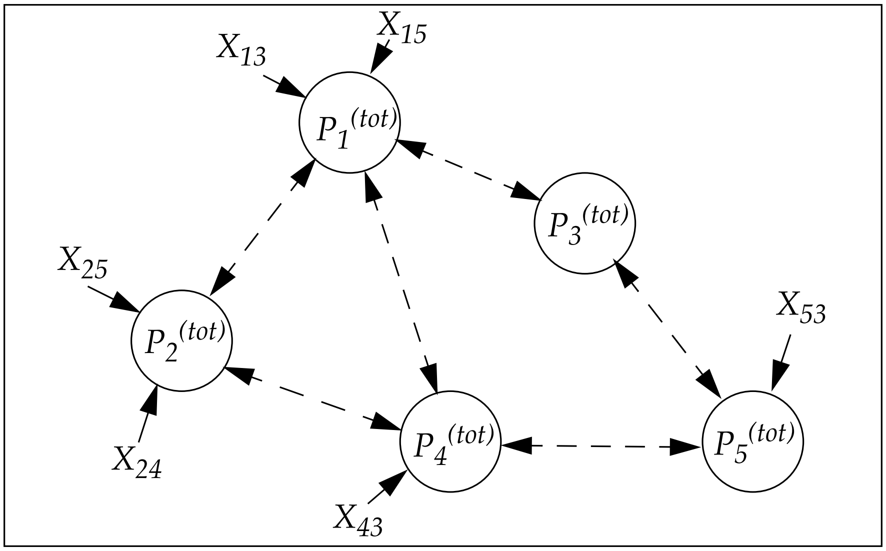

Definition of the Capacity Region Λ:

Let λ_{ij} be the bit rate of stream X_{ij}.

The Capacity region Λ is the set of all rate matrices such that:

-The network is necessarily unstable whenever $(\lambda_{ij}) \notin \Lambda$ (even if future known).

-The network can be stabilized if (λ_{ij}) is *strictly interior* to Λ.

<u>A note on Stability</u>: Consider a queue with input process *X(t)* and processing rate μ*(t)*

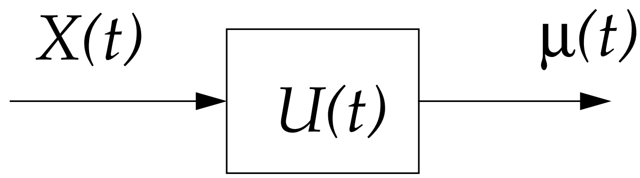

X(t) = amount of bits that arrived in [0, *t*].
μ(t) = instantaneous processing rate.
U(t) = Uninished work in queue at time *t*.

(Need to consider general--potentially non-ergodic case).

Definition: The *overflow function g(M):*

$$g(M) = \limsup_{t \to \infty} \left[\frac{1}{t} \int_0^t 1_{\{U(\tau) > M\}} d\tau \right]$$

Definition: A queueing system is *stable* if $g(M) \to 0$ as $M \to \infty$.

A network is stable if all queues are stable.

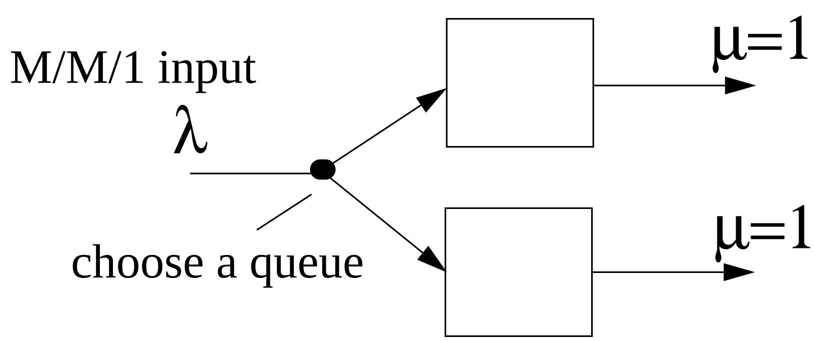

The *lim sup* definition is essential to obtain the correct notion of stability. The above system is stable whenever λ<μ. If *lim inf* is used, it is stable for all λ< ∞.

The simplest possible network:
A single queue with slotted time, timeslot size = T

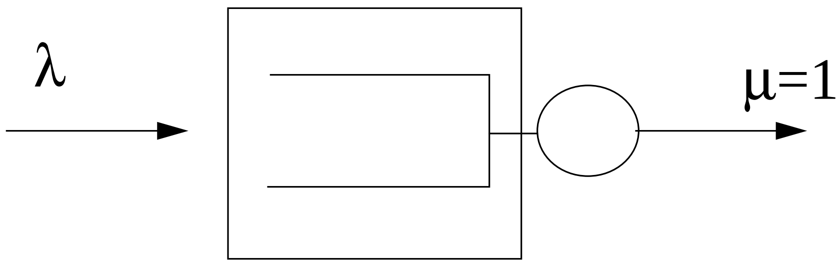

Capacity and Delay:

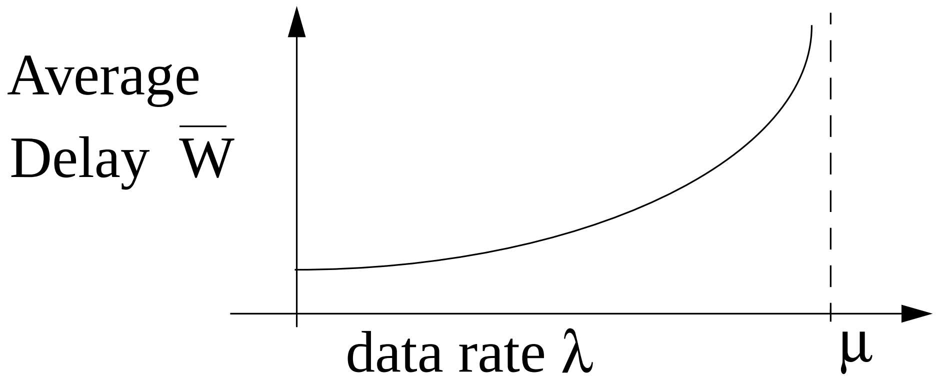

Well Known P-K Formula for M/G/1 Queue:

$$\overline{W} = \frac{\lambda T E(A^2)}{2(\mu - \lambda)}$$

2^{nd} moment of arrivals

$\varepsilon = \mu - \lambda$

New Result for Bursty Data and/or Time Varying $\mu(t)$:

$$\overline{W} \leq \frac{KT\left(A_{max}^2 + \mu_{max}^2\right)}{2(\mu_{av} - \lambda)}$$

Example:
Channel States:

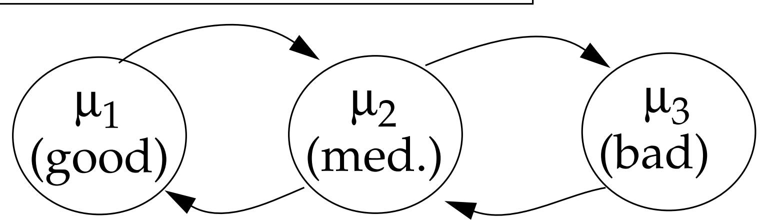

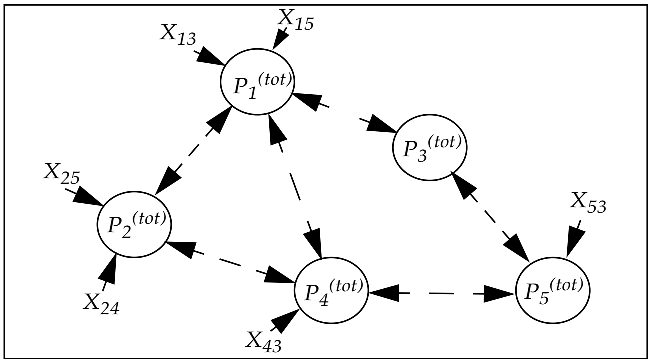

<u>Theorem 1</u>: Wireless Network Capacity Region Λ is the set of all (λ_{ij}) rates s.t. there are flows $f_{ab}^{(c)}$ with:

$$f_{ab}^{(c)} \geq 0 \qquad \text{(non-negativity)}$$

$$\lambda_{ic} + \sum_a f_{ai}^{(c)} = \sum_b f_{ib}^{(c)} \quad \forall i \neq c \qquad \text{(flow conservation)}$$

$$\sum_a \lambda_{ac} = \sum_a f_{ac}^{(c)} \quad \forall c \qquad \text{(sink the data)}$$

$$\sum_c f_{ab}^{(c)} \in \Gamma \qquad \text{(link capacity constraint)}$$

where <u>graph family Γ</u> is the set of all feasible 1-hop link rates achievable by some power alloc. strategy:

$$\Gamma = \sum_{\underline{C}} \pi_{\underline{C}} Conv\{\mu(\underline{P}, \underline{C}) \mid \underline{P} \in \Pi\}$$

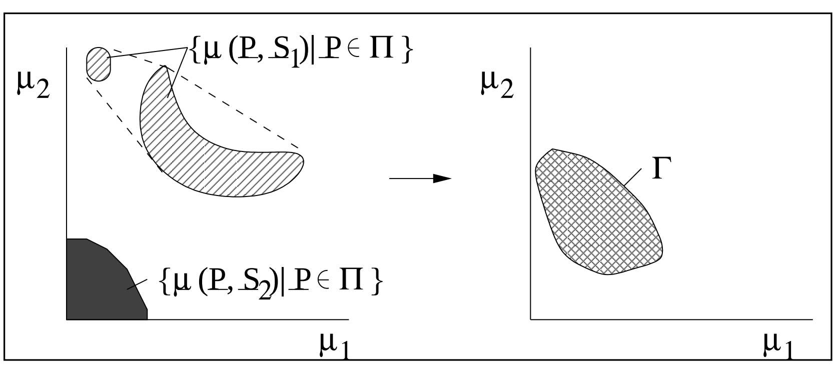

Note: Capacity region depends only on steady state channel probability distribution $\pi_{\underline{C}}$.

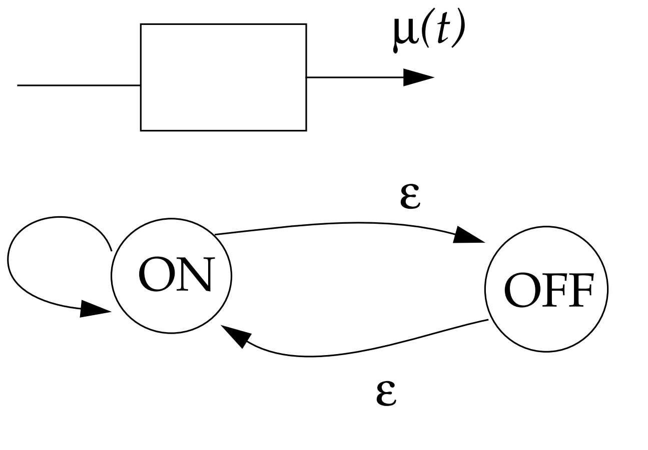

$$\mu(t) = \begin{cases} \text{ON prob } 1/2 \\ \text{OFF prob } 1/2 \end{cases}$$

Thus, any channel state evolution with the same steady state probabilities yields the same capacity region.

(altho the exact dynamics can significantly effect delay)

Corollary: The capacity region is preserved if we consider channel states $\underline{C}$ which are chosen iid each timeslot.

Capacity Achieving Strategy:

A generalization of the Tassiulas Backpressure strategy [Tassiulas,Ephremides 92]

Define: $U_i^{(c)}$ = Unfinished Work in i destined for c

<u>Joint Routing and Power Allocation</u>: Every timeslot, and for each link (i,j), find the commodity $d(i,j)$ that has the largest *differential backlog* $U_i^{d(i,j)}(t) - U_j^{d(i,j)}(t)$. Route this commodity from i to j, using the power allocation $\underline{P}(t)$ determined by:

maximize: $$\sum_{i,j} \mu(\underline{P}, \underline{C}(t))\left[U_i^{d(i,j)}(t) - U_j^{d(i,j)}(t)\right]$$

subject to: $$\sum_j P_{ij} \le P_i^{tot} \quad \text{for all } i.$$

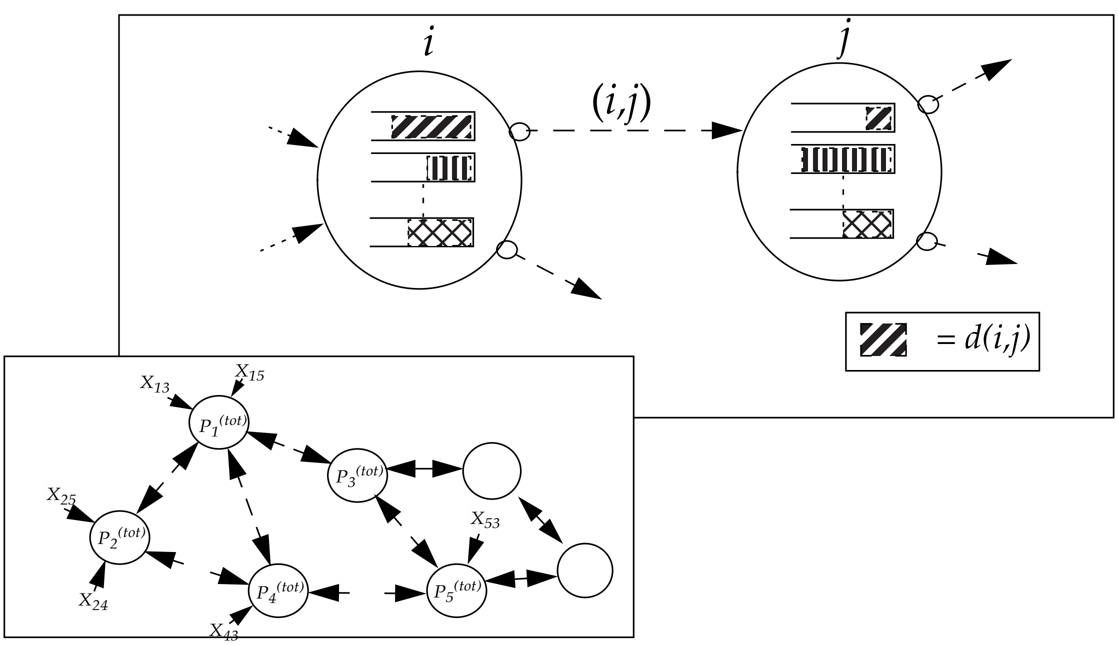

<u>Theorem</u>: The Differential Backlog Policy stabilizes the system whenever possible, <u>without requiring knowledge</u> of the arrival processes or channel state processes, and ensures the following delay guarantee:

Average Delay:

$(\lambda_{ij} + \varepsilon) \in \Lambda$

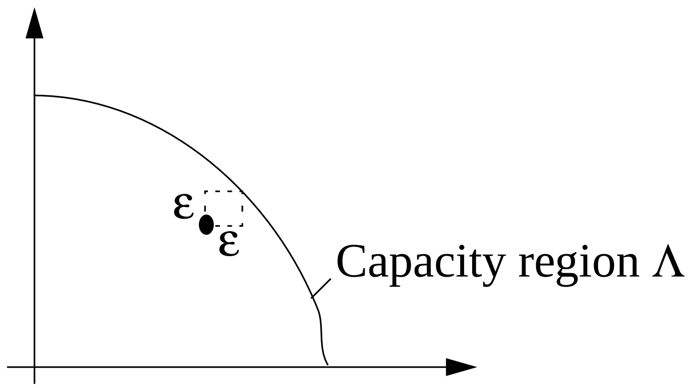

ε can be viewed as the "distance" to the boundary of the capacity region Ω.

$$\text{Average Delay in Network} \leq \frac{KT \sum_{i,j} \left(E\left[A_{ij}^2\right] + E\left[\mu_{ij}^2\right] \right)}{2\lambda_{tot}\varepsilon}$$

□

Note Fundamental Similarity to M/G/1 queue:

$$\text{Average Delay}_{(M/G/1)} = \frac{\lambda T E(A^2)}{2(\mu - \lambda)}$$

Can prove the result using the theory of Lyapunov Drift.

When viewed from above...

<u>A Dual Formulation</u>: Consider just testing if a rate matrix (λ_{ij}) is inside of the capacity region:

Maximize 1 Subject to

$$f_{ab}^{(c)} \geq 0 \qquad \text{(non-negativity)}$$

$$\lambda_{ic} + \sum_a f_{ai}^{(c)} = \sum_b f_{ib}^{(c)} \quad \forall i \neq c \qquad \text{(flow conservation)}$$

$$\sum_a \lambda_{ac} = \sum_a f_{ac}^{(c)} \quad \forall c \qquad \text{(sink the data)}$$

$$\sum_c f_{ab}^{(c)} \in \Gamma \qquad \text{(link capacity constraint)}$$

Dual:

Maximize: $fun(\mu_{ij}(p_{ij}), U_{ij})$ (to find a subgradient)

Update : $U_i^{(c)}(t+T) = \left(U_i^{(c)}(t) - T\sum_b \mu_{ib}^{(c^*)}\right)^+ + T\sum_a \mu_{ai}^{(c^*)} + T\lambda_{ic}$

A relationship between static method for computing a multi-commodity flow and a dynamic backpressure scheme that achieves capacity...thru the unifying framework of convex duality.

Application To Ad-Hoc Networks

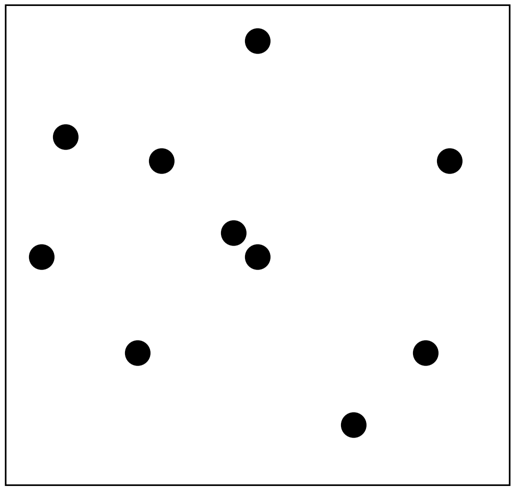

Capacity:

Static Networks: Gupta, Kumar -- $O(1/\sqrt{N})$

Mobile Networks: Grossglauser, Tse -- $O(1)$

$$\sum_{i,j} \mu(\underline{P}, \underline{C}(t))\left[U_i^{d(i,j)}(t) - U_j^{d(i,j)}(t)\right]$$

The DRPC algorithm achieves this capacity with the optimal coefficient in mobile or non-mobile case.

Can we achieve full capacity in a distributed way?
Conjecture: No (unless channels independent)

Can we achieve asymptotic capacity?
Answer: Yes, with distributed approximations...

Implementation for Mobile Ad-Hoc Networks:

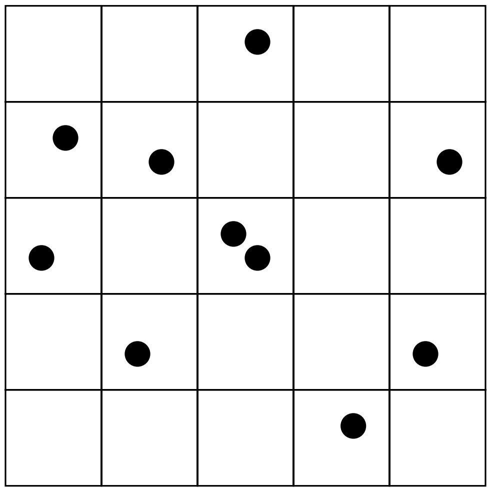

Discretize Location Space of Network to a simple 5 x 5 grid

10 Users randomly moving (prob. 1/2 they stay in same cell, prob. 1/2 they move to an adjacent cell).

Attenuation Model (i.e., a $1/r^4$ loss characteristic)

$$SIR_{ab}(\underline{P}, \underline{C}) = \frac{\text{Attenuated Signal at } b}{N_b + \text{Atten. Interference at } b}$$

modulation	bits/symbol	power/symbol
2 PAM	1	$0.25\Delta^2$
4 QAM	2	$0.5\Delta^2$
16 QAM	4	$1.25\Delta^2$
64 QAM	6	$5.25\Delta^2$

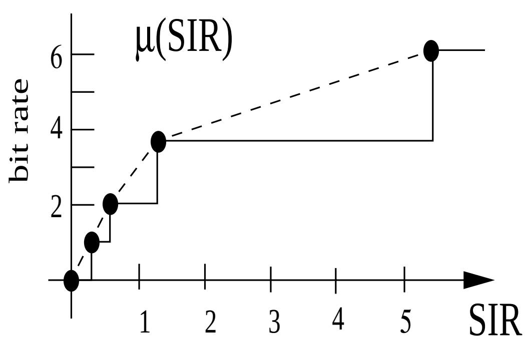

Distributed Implementation:
Use side channel to exchange backlog info with neighbors, and learn local link attenuations α_{ij}. Then:

1. At beginning of each timeslot, each node randomly decides to transmit (at full power P_{tot}) or remain idle,with prob 1/2. A control signal of power γP_{tot} is transmitted.

2. Define Ω as the set of all transmitting nodes. Each node b measures its total interference

$$\gamma I_b = \sum_{i \in \Omega} \alpha_{ib} \gamma P_{tot}$$

and sends this scalar quantity to all neighbors.

3. Using knowledge of the interference, attenuation, and queue backlogs of neighbors, user a transmits with full power to user b who maximizes the function

$$W_{ab} \log\left(1 + \frac{\alpha_{ab} P_{tot}}{N_b + I_b - \alpha_{ab} P_{tot}}\right)$$

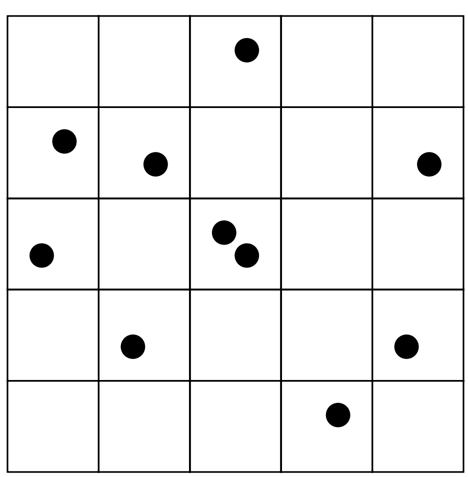

$$\text{maximize:} \quad \sum_{i,j} \mu(\underline{P}, \underline{C}(t))\left[U_i^{d(i,j)}(t) - U_j^{d(i,j)}(t)\right]$$

$$\text{subject to:} \quad \sum_j P_{ij} \le P_i^{tot} \quad \text{for all } i.$$

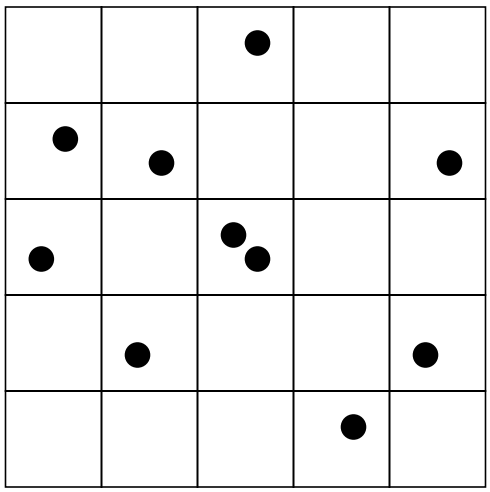

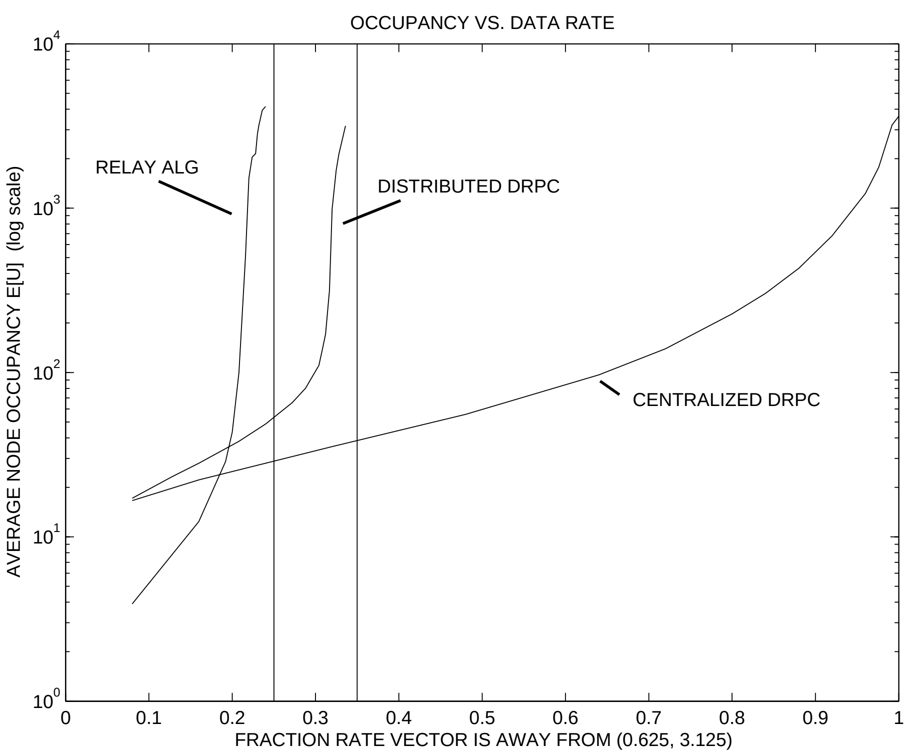

Concluding Summary:

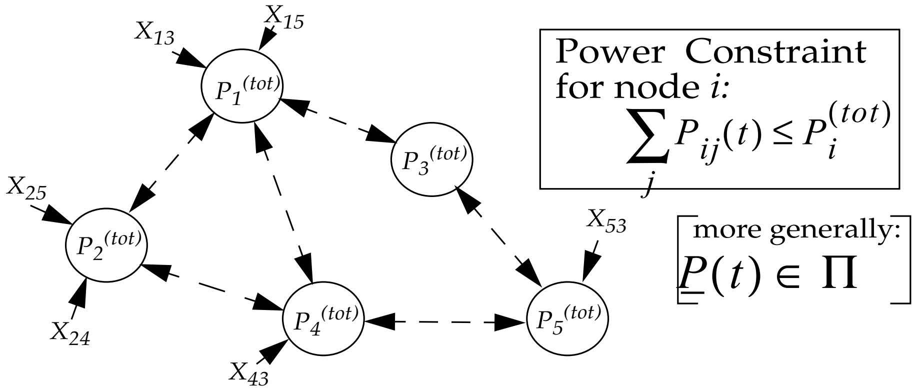

General Power Allocation Formulation for a Wireless Network

Learned: There are some principles that can be applied to all networks, but every type of network has distinct structure which must be understood for development of control algorithms.

Dynamic Power Allocation Algorithm:

-Issues of implementation complexity

-How much control information is needed?