Stable Scheduling Policies for Broadcast Channels

A. Eryılmaz, R. Srikant (University of Illinois)
and J.R. Perkins (Boston University)

Outline

- System Model & Objective
- A set of Scheduling Policies
  - $\rightarrow$ Deterministic Analysis
  - $\rightarrow$ Stochastic Analysis
  - $\rightarrow$ Simulations
- A different set of Scheduling Policies
  - $\rightarrow$ Results
  - $\rightarrow$ Simulations
- Conclusions

System Model

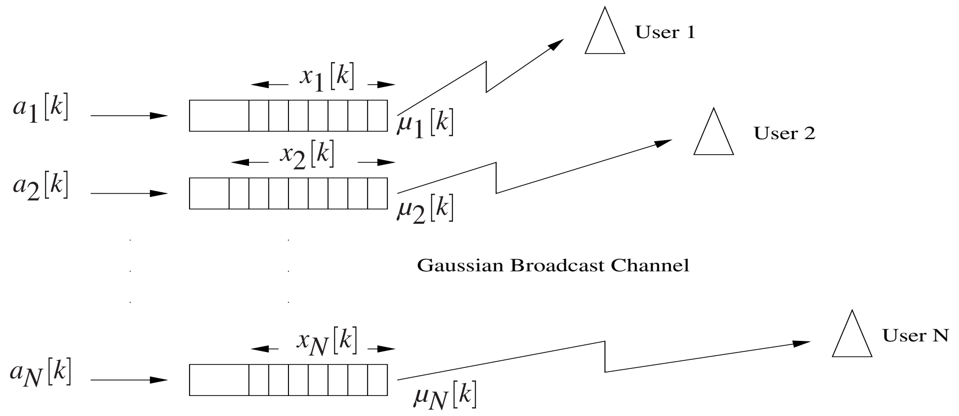

- N separate queues of infinite length, each maintained for a different user.
- Time is slotted.
- $x_i[k]$: length of queue i at the beginning of slot k.
- $a_i[k]$: number of bits that arrive at queue i at slot k.
- $\mu_i[k]$: number of bits of service offered to queue i in slot k.
- $x_i[k+1] = (x_i[k] + a_i[k] - \mu_i[k])^+$, where $(y)^+ = y$ if $y > 0$ and zero otherwise.

Goal

- Without knowing the arrival rate, the goal is to

  FIND A STABILIZING POWER ALLOCATION STRATEGY.

- Such a policy is called *throughput optimal* in the sense that, it achieves any mean throughput that is achievable under stability.
- Throughput optimal scheduling policies are offered and analyzed

  $\longrightarrow$ first using a deterministic model,

  $\longrightarrow$ then in the stochastic setting.

Assumptions

- The achievable rate region, $\mathcal{S}$,
  - $\rightarrow$ is convex and bounded.
  - $\rightarrow$ is known by the transmitter.
  - $\rightarrow$ Broadcast capacity region is a specific example.
- The arrival processes, $a_i[k]$,
  - $\rightarrow$ are independent for each queue and at each time slot.
  - $\rightarrow$ $\vec{\lambda} := (\lambda_1, \cdots, \lambda_N)$, where $\lambda_i := E[a_i[1]]$, lies within $\mathcal{S}$.
  - $\rightarrow$ $E[a_i^2[1]] = A_i < \infty$.
  - $\rightarrow$

$$\lim_{M \to \infty} E\left[\frac{1}{M}\left|\sum_{n=0}^{M-1} a_i[n] - \lambda_i\right|\right] = 0, \qquad \forall i.$$

Scheduling Policy

Policy 1: Let $\{K_1, \cdots, K_N\}$ be a set of positive constants and $\alpha > 0$. Given the queue length vector $\vec{x}[k] := (x_1[k], \cdots, x_N[k])' \neq \vec{0}$, and $\mathcal{S}$, choose the service rate vector $\vec{\mu}[k] = (\mu_1[k], \cdots, \mu_N[k])'$ such that it is the solution to the following nonlinear program:

$$\max_{C \geq 0} \quad C$$

subject to

$$\mu_i[k] = CK_i(x_i[k])^{\alpha}, \qquad \forall i,$$

and $\vec{\mu}[k] \in \mathcal{S}$. If $\vec{x}[k] = \vec{0}$, then $\vec{\mu}[k]$ can be chosen arbitrarily. $\diamond$

Scheduling Policy [Cont'd]

Remark: An equivalent definition would be to choose $\vec{\mu}[k]$ such that

$$\frac{K_1(x_1[k])^{\alpha}}{\mu_1[k]} = \frac{K_2(x_2[k])^{\alpha}}{\mu_2[k]} = \cdots = \frac{K_N(x_N[k])^{\alpha}}{\mu_N[k]} \tag{1}$$

and $\vec{\mu}[k]$ lies on the boundary of $\mathcal{S}$. ◇

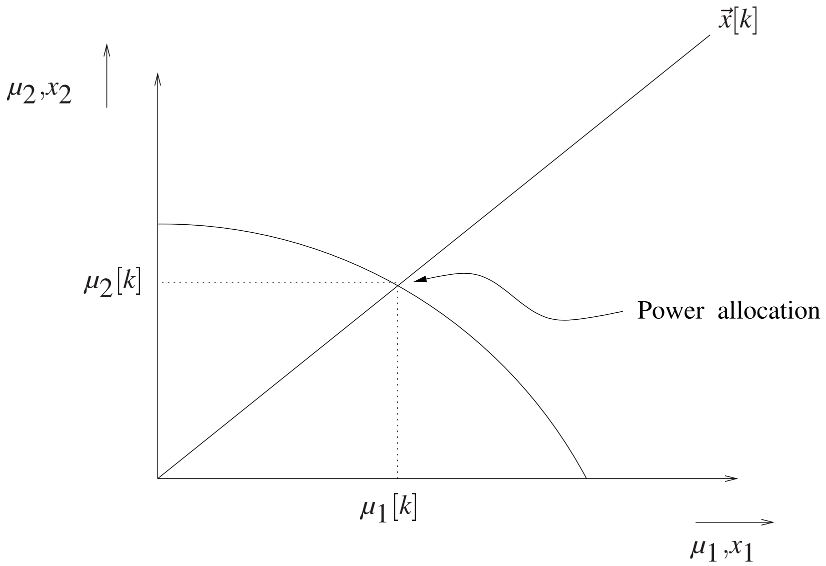

Figure 1: Power allocation when $K_1 = K_2 = \alpha = 1$.

Deterministic, Discrete-time model

- Arrivals to each queue is constant at their mean, λ_i, $\forall i$.
- The queue length evolution is governed by

$$x_i[k+1] = (x_i[k] + \lambda_i - \bar{\mu}_i[k])^+, \qquad \forall i.$$

- Facts about the policy:

  $\rightarrow$ There exists $\{\tilde{\mu}_i\}_{i=1}^N$ on the boundary of $\mathcal{S}$ satisfying

$$\frac{K_1(\tilde{\mu}_1 - \lambda_1)^\alpha}{\tilde{\mu}_1} = \frac{K_2(\tilde{\mu}_2 - \lambda_2)^\alpha}{\tilde{\mu}_2} = \cdots = \frac{K_N(\tilde{\mu}_N - \lambda_N)^\alpha}{\tilde{\mu}_N},$$

  such that $\tilde{\mu}_i > \lambda_i$.

  $\rightarrow$ If, at any time $k = k_0$, $\vec{x}[k] \neq \vec{0}$ and

$$\frac{x_1[k]}{\tilde{\mu}_1 - \lambda_1} = \frac{x_2[k]}{\tilde{\mu}_2 - \lambda_2} = \cdots = \frac{x_N[k]}{\tilde{\mu}_N - \lambda_N}, \qquad (attracting\ line)$$

  then $\bar{\mu}_i[k] = \tilde{\mu}_i\ \forall i$. Further, $\forall k > k_0$, until $\vec{x}[k] = \vec{0}$, the equalities continue to hold and $\bar{\mu}_i[k] = \tilde{\mu}_i\ \forall i$.

Deterministic, Discrete-time model [Cont'd]

Define the Lyapunov function

$$V(\vec{x}) = \max_i \frac{x_i}{\tilde{\mu}_i - \lambda_i}.$$

Lemma: Under our scheduling policy, for any given M, $\exists R_M < \infty$, such that for all $\vec{x}[k]$ satisfying

$$\sum_{i=1}^{N} \frac{x_i[k]}{\tilde{\mu}_i - \lambda_i} > R_M,$$

we have

$$\Delta V^{(M)}(\vec{x}[k]) := V(\vec{x}[k+M]) - V(\vec{x}[k]) \leq -M\delta$$

with some $\delta > 0$.

Stochastic Analysis

- Arrivals are no longer assumed to be constant.

Theorem 1 *The system introduced is stable in the mean under the assumptions stated earlier, i.e.,*

$$\limsup_{K\to\infty} \frac{1}{K} \sum_{n=0}^{K-1} E\left[x_i[n]\right] < \infty \qquad \forall i.$$

Stochastic Analysis[Cont'd]

Sketch of the Proof: To prove Theorem 1, the following Lyapunov function is utilized

$$W(\vec{x}) := (V(\vec{x}))^2.$$

Then the M-step mean drift of W is examined.

$$\begin{aligned}\Delta W^{(M)}(\vec{x}) &:= E\left[W(\vec{x}[k+M]) - W(\vec{x}[k]) \quad \middle| \quad \vec{x}[k] = \vec{x}\right] \\ &= E\big[(V(\vec{x}[k+M]) - V(\vec{x}[k])) \\ &\qquad \cdot (V(\vec{x}[k+M]) + V(\vec{x}[k])) \quad \big| \quad \vec{x}[k] = \vec{x}\big]\end{aligned}$$

By choosing M and R_M sufficiently large, it is shown that the mean drift of the stochastic model cannot deviate from the drift of the deterministic model by much, hence the result follows.

$\rightarrow$ *Even when the channel state is allowed to fluctuate, under reasonable conditions (such as ergodicity and information on channel statistics), Theorem 1 still holds.*

Stochastic Analysis[Cont'd]

Theorem 2 *If the system is further restricted to take only discrete queue length values, then it is still stable in the mean under the assumptions stated earlier, i.e.,*

$$\limsup_{K\to\infty} \frac{1}{K} \sum_{n=0}^{K-1} E\left[x_i[n]\right] < \infty \qquad \forall i.$$

→ If the arrival processes is such that the queuing system forms an aperiodic Markov chain with a single communicating class, then Theorem 2 implies that the chain is *positive recurrent*.

Simulations

- A two queue system is simulated.
- The mean arrival rates, (5,10), lie inside and very close to the boundary of the capacity region.
- Deterministic model simulations

  $\rightarrow$ Arrival process to the first queue has rate $\lambda_1 = 5$. The arrival process to the second queue has rate $\lambda_2 = 10$.
- Stochastic model simulations

  $\rightarrow$ Arrival process to the first queue is Poisson with mean, $\lambda_1 = 5$. The arrival process to the second queue is uniform with mean, $\lambda_2 = 10$.

Simulations[Cont'd]

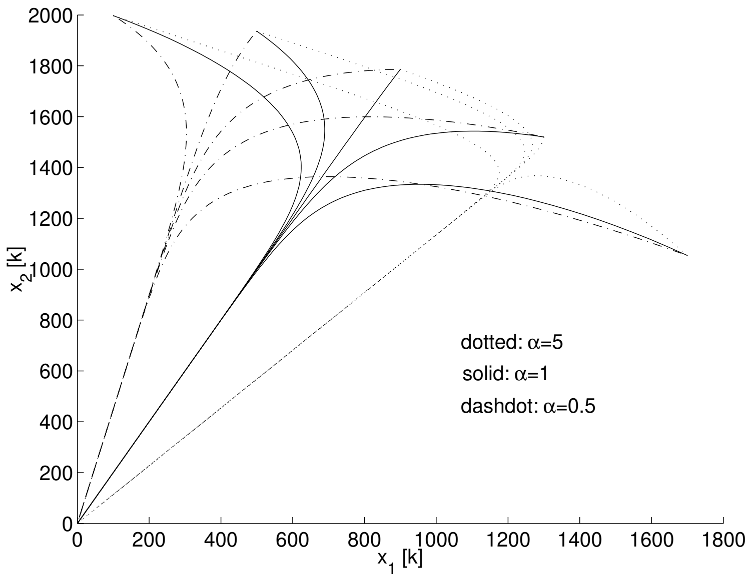

Figure 2: Evolution in the state space, $K_1 = K_2$.

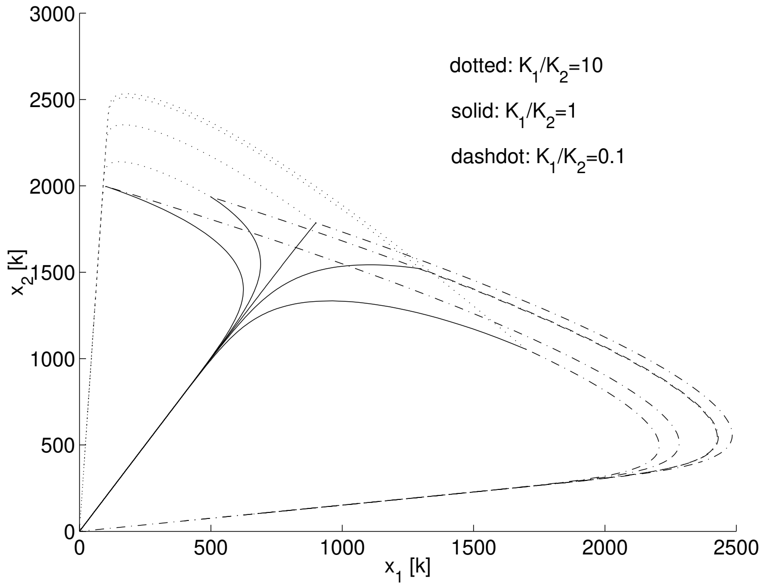

Figure 3: Effect of K_i with $\alpha = 1$.

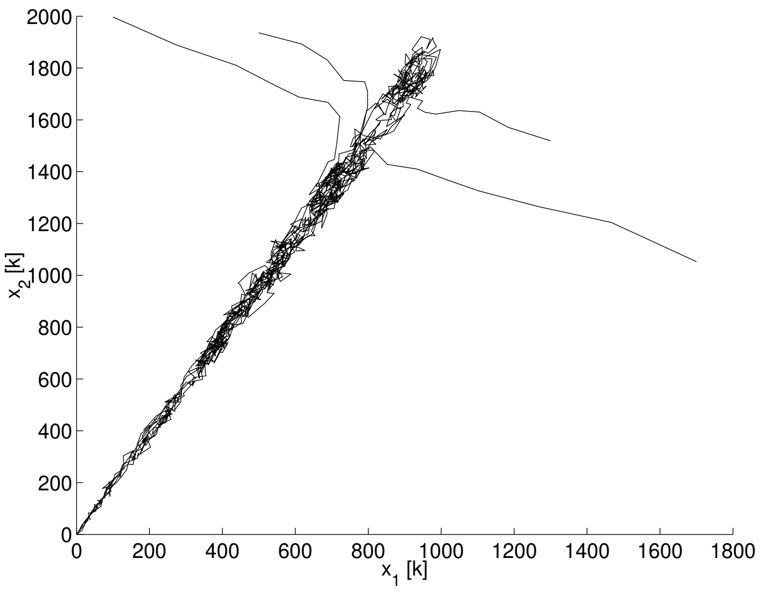

Figure 4: Evolution of the queue lengths for $\alpha = K_1 = K_2 = 1$

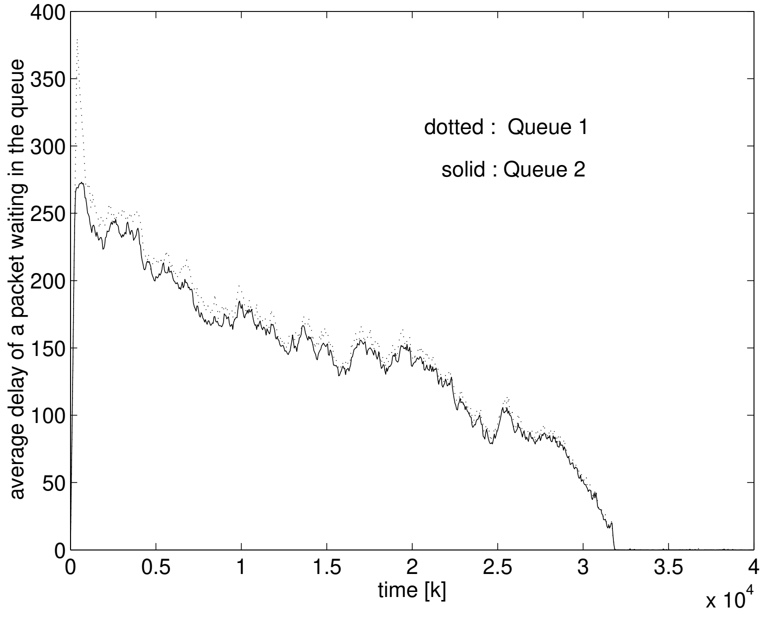

Figure 5: Average delay experienced for $\alpha = K_1 = K_2 = 1$

- Difference between the average delays is kept small throughout the operation.
- Priorities can be assigned through the weighting factor, while maintaining stability.

$\boxed{\text{The policy has a fair average delay property in addition to its throughput optimality}}$

A different set of Scheduling Policies

- The channel state is allowed to vary from one slot to the next, among J possibilities.
- $\mathcal{S}_{s[k]}$ identifies the set of allowable rates at time k, with $s[k] \in \{1, \cdots, J\}$.

Policy 2: Let $f_i(x), i = 1, \cdots, N$ be positive functions for $x \geq 0$. Given the queue length vector $\vec{x}[k] := (x_1[k], \cdots, x_N[k])' \neq \vec{0}$, and the channel state $s[k] \in \{1, \cdots, J\}$ at the beginning of time slot k, choose the service rate vector $\vec{\mu}[k] = (\mu_1[k], \cdots, \mu_N[k])'$ such that it is the solution to the following maximization:

$$\max_{\vec{\mu}[k] \in \mathcal{S}_{s[k]}} \quad \sum_{i=1}^{N} f_i(x_i[k])\mu_i[k]$$

$\diamond$

Additional assumptions

- For $i \in \{1, \cdots, N\}$, $f_i : [0, \infty) \to [0, \infty)$

  $\to$ is an increasing, continuous function..

  $\to$ $f_i(x) \geq 0$ with equality only if $x = 0$.

  $\to$ $\lim_{x \to \infty} f_i(x) = \infty$, for all i.

- The channel state process has a stationary distribution with π_j denoting the probability of being in state j.

  $\to$

$$\lim_{M \to \infty} E\left[\left|\pi_j - \frac{1}{M}\sum_{n=0}^{M-1} I_{s[n]=j}\right|\right] = 0, \qquad \forall j \in \{1, 2, \ldots, J\}.$$

Results

→ The following Lyapunov function is utilized

$$V(\vec{x}) := \sum_{i=1}^{N} g_i(x_i),$$

where $g_i'(x) := f_i(x)$.

Theorem 3 *For the system introduced, if $\vec{\lambda}$ lies within the average achievable rate region $\mathcal{S}$, then the following is true*

$$\limsup_{K\to\infty} \frac{1}{K} \sum_{n=0}^{K-1} E\left[||\vec{f}(\vec{x}[k])||_2 \right] < \infty,$$

where $||\vec{f}(\vec{x}[k])||_2 := \left(\sum_{i=1}^{N} f_i^2(x_i) \right)^{1/2}$.

Simulations[Cont'd]

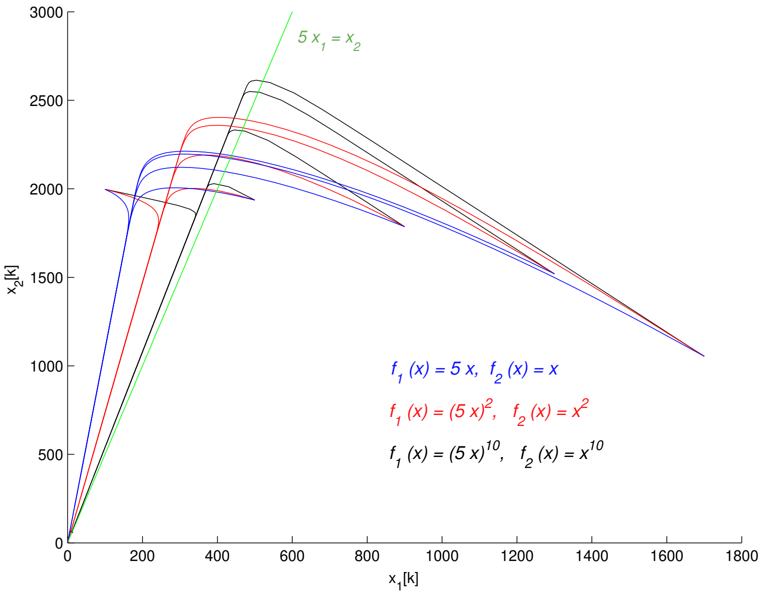

Figure 6: Evolution in the state space

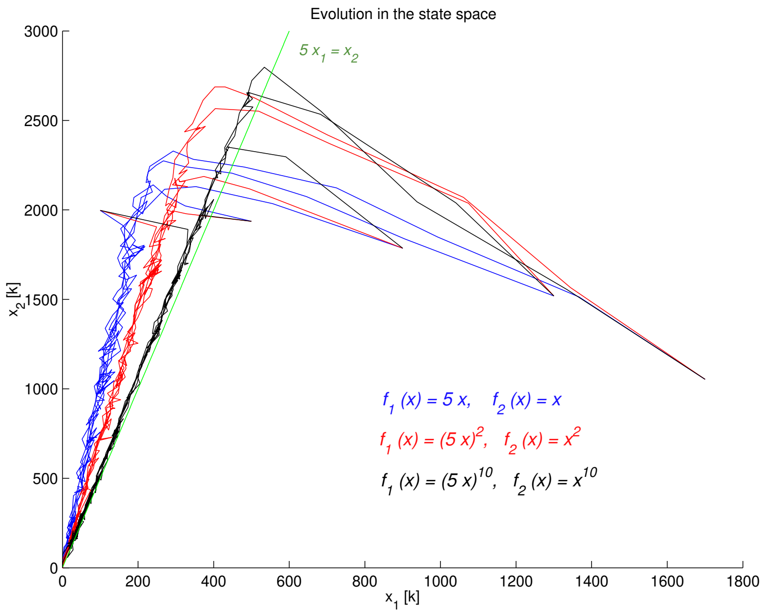

Figure 7: Stochastic model

Conclusions

- Sets of stabilizing scheduling policies are suggested and proved for transmitter with N parallel queues, using a Lyapunov type of argument.
- It is shown that such policies can be applied for a fading broadcast channel scenario.
- Results can be extended to enable imperfect information on queue lengths or channel states.
- Fairness and congestion control?

Thank you!