Queues under feedback Control

Sergio D Servetto

School of Electrical & Computer Engineering
Cornell University
http://people.ece.cornell.edu/servetto/

Acknowledgement

- Răzvan Cristescu
  (Ecole Polytechnique fédérale de Lausanne)

On the Shannon capacity of a Single-Server Queue

A simple queueing model:

- Packets of size 1 bit
- Arrivals rate λ departures rate μ

Shannon capacity $\frac{2}{e}\log_2(e)\,\mu \approx 1.0615\,\mu$ b/s

[V. Anantharam S. Verdú Bits Through Queues Trans IT 1/1996]

HOW IS THIS POSSIBLE ???

On the Shannon Capacity of a Single-Server Queue

Answer: <u>timing</u> information.

Example: assume service rate has variance 0

- Choose a real number $s >$ at tx.
- Inject one packet at time 0.
- — — — — — s.

$\Rightarrow$

- Packet 1 departs exactly at time μ.
- — 2 — — — — $\mu + s$.

$\Rightarrow$ Rx can recover s exactly

$\Rightarrow$ INFINITE CAPACITY!

On the Shannon Capacity of a Single-Server Queue

A&V's key insight: idle periods of the queue can be used to convey information

Basic tradeoff:

Communicate via packet contents → loaded queue and no timing data

vs

Communicate via idle periods → empty queue and no packet data

So, for large packets, it "expensive" to leave the queue idle

A natural generalization:

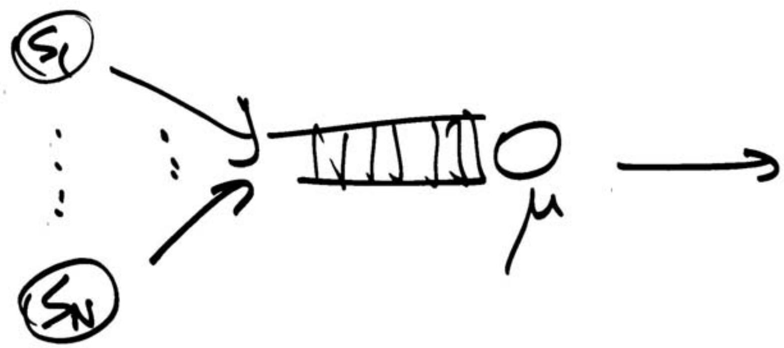

Key questions to address here:

- How do sources gain access to the shared queue?
- What information do they have available to make these decisions?
- Performance bounds?

Issues in the back of our mind

(or, why should we care?)

- From A&V: IT efficiency ⇒ "almost stochastic instability" (at least for queues)
- IT considerations in flow control problems
- Decentralized control
- Analysis of systems in which observations perturb state

Nice if we could set up a "toy" problem in which to say interesting things about these issues.

Talk Outline

- System Model
- The Control Problem
  - Complete Observations
  - Partial Information
  - Numerical Results
- Performance analysis
- Summary / Conclusions / Applications in Large Scale Sensor Networks

System Model

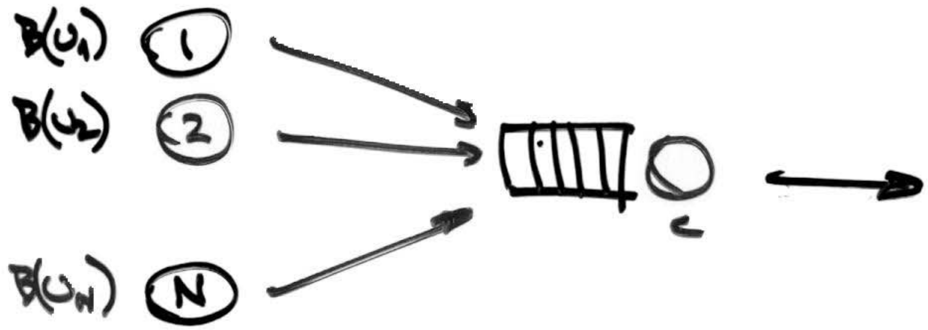

- Discrete time
- ON/OFF sources, # is Markov in time
- Observations 1 −1, 0 (finite buffer)

  No communication among sources
- Controllable intensities, and everyone runs the same controller.

System Dynamics

- $x_k \in \{1 \ldots N\}$: # of sources ON at time k.
- $r_k^{(i)} \in \{-1, 0, 1\}$: i-th source observation at time k.
- $u_k^{(i)} \in (0, 1]$: i-th source intensity at time k.
- A-priori information, assumed known:
  - $P_{sv} = P_r(x_k = s \mid x_{k-1} = v)$, indep of $u_k^{(i)}$ & k.
  - $P_r(x_0)$, initial distribution
- Information available via measurements:
  - $P_r(r \mid x, u)$

System Dynamics – Example

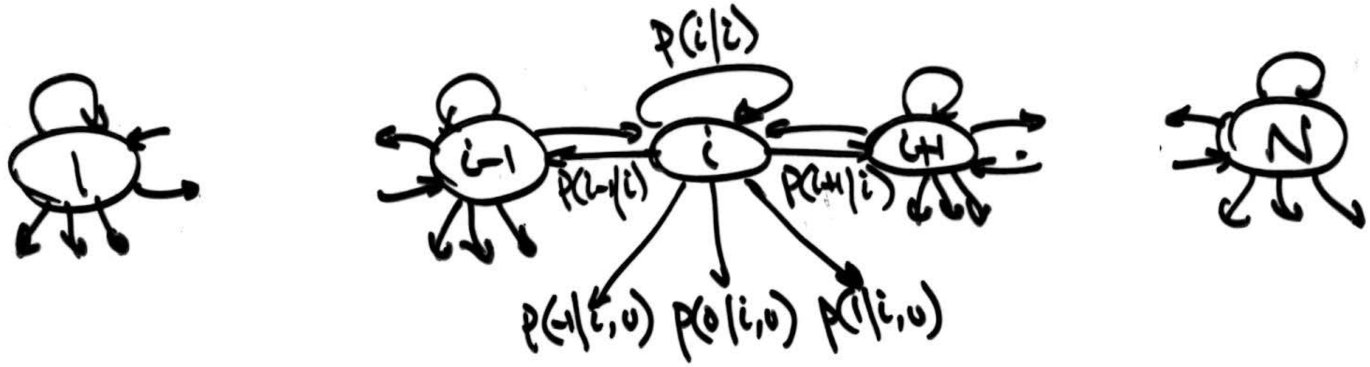

(a birth & death chain model)

Steps:

Chain on x_k → flip a biased coin $B(u_k)$ → get an observation r_k → move to x_{k+1}

Falls in the framework on Controlled MCs...

IV The Control Problem

Formal statement: find a policy g s.t.

$$\boxed{\begin{aligned} &\max_{g} \ \limsup_{K\to\infty} \frac{1}{K}\sum_{k=0}^{K-1} p(r_k = 1 \mid x_k, u_k) \\ &\text{subject to: } p(r_k = -1 \mid x_k, u_k) \le T, \ \forall k \end{aligned}}$$

- $T \in (0,1]$ is a parameter.
- lim sup because we don't know yet that the limit exists (but it does!).

The Control Problem

"Warming up" finite horizon, complete observations

$\Rightarrow$ solution given by DP!

Define $\vec{c}(u) = [c(1\ u)\quad c(N\ u)]^T$, $c(i\ u) = p(1|i\ u)$

Then $V_K = 0$

$$V_k = \sup_{u_T} \{\vec{c}(u) + P(u) V_{k+1}\}$$

$$= \sup_{u_T} \vec{c}(u) + \text{const}$$

(range of u_T depends on constraints sup on each component of $\vec{c}$)

The Control Problem

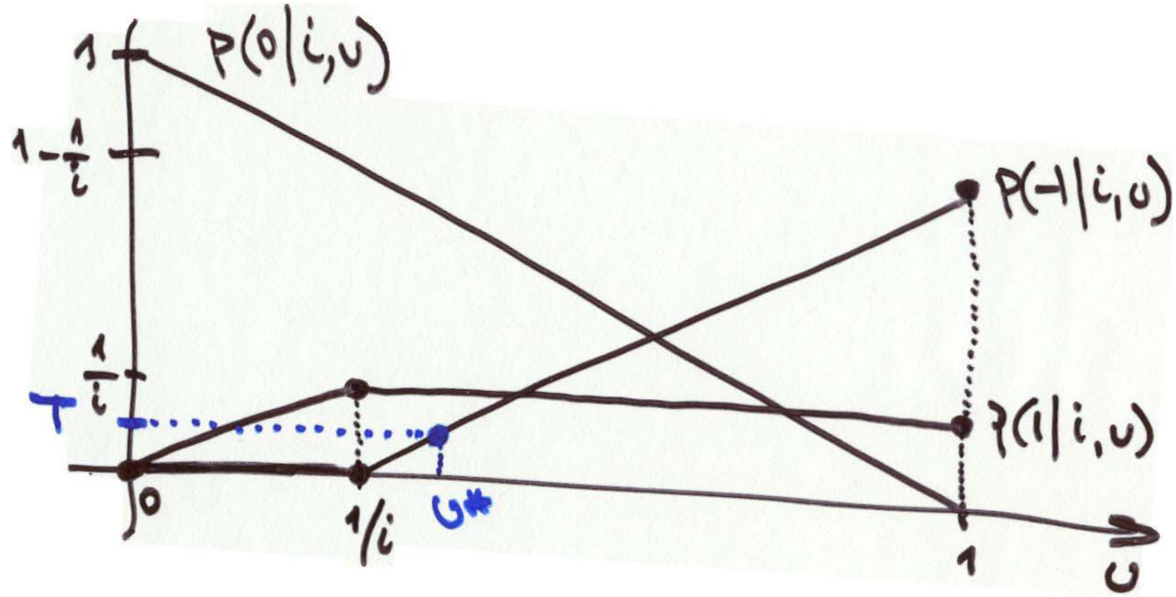

u^* = largest u in $(0,1]$ s.t.

$$P(-1|i,u) = T$$

(Intuitively very pleasing)

The Control Problem

"The real thing": partial information

⟹ solution given by DP on info states!

Structure of an optimal controller:

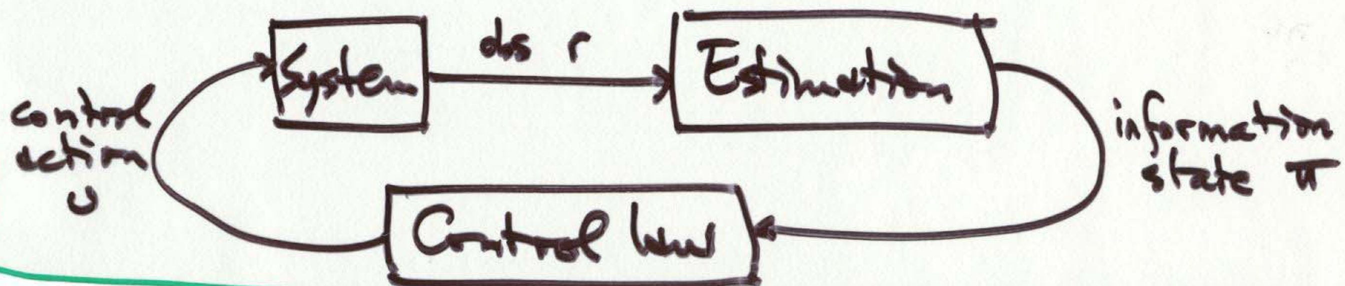

[P.R. Kumar, P. Varaiya Stochastic Systems: Estimation, Identification & Adaptive Control Prentice-Hall, 1986]

The Control Problem

What is an information state?

- π_k function of $r_0 \quad r_{k-1}\ u_0 \cdots u_{k-1}$
- π_{k+1} can be obtained from π_k, r_k, u_k

"Typically" $\pi_k = P_r(x_k \mid r_0 \cdots r_{k-1}\ u_0 \cdots u_{k-1})$

Intuitively:

- π_k is everything we can infer about the state given everything we have seen.
- π_k can be updated recursively: natural ref. ~~for~~ something we call "state"

The Control Problem

New utility function: $\underline{c(\pi, u) = p(1 \mid \pi, u)}$

$$\boxed{c(\pi_k, u_k)} = p\Big(r_k = 1 \;\Big|\; \underbrace{p_r(x_k \mid r_0^{k-1} u_0^{k-1})}_{\pi_k}, u_k\Big)$$

$$= \sum_{x_k} p(r_k = 1 \mid x_k, \pi_k, u_k)\, p(x_k \mid \pi_k u_k)$$

$$= \sum_{x_k} p(r_k = 1 \mid x_k, u_k)\, p(x_k \mid \pi_k u_k)$$

$$= \sum_{x_k} p(r_k = 1 \mid x_k, u_k)\, p(x_k \mid r_0^{k-1} u_0^{k-1})$$

$$= \boxed{E_{\pi_k}\left(c(x_k, u_k)\right)}$$

(also intuitively very pleasing!)

The Control Problem

New DP equations:

- $V_K(\pi) = \odot$
- $V_k(\pi) = \sup_{u \in U_\pi} E_\pi \left\{ c(x_k, u) + V_{k+1}\big(F(r_k, u, \pi)\big) \right\}$

$$= \sup_{u \in U_\pi} \left\{ \sum_{i=1}^{N} c(i,u)\,\pi(i) + \sum_{r=-1}^{1} V_{k+1}\big(F(r,u,\pi)\big)\, p(r \mid \pi, u) \right\}$$

where

- F : state transition function (updates of π).
- U_π : $\{ u \mid E_\pi[p(-1 \mid x, u)] \le T \}$

The Control Problem

So, an optimal controller chooses

$$U_k^* = g(\pi_k) = \underbrace{\arg \sup_{u \in U_\pi} E_\pi(\cdots)}_{V_k(\pi)}$$

Some numerical simulations

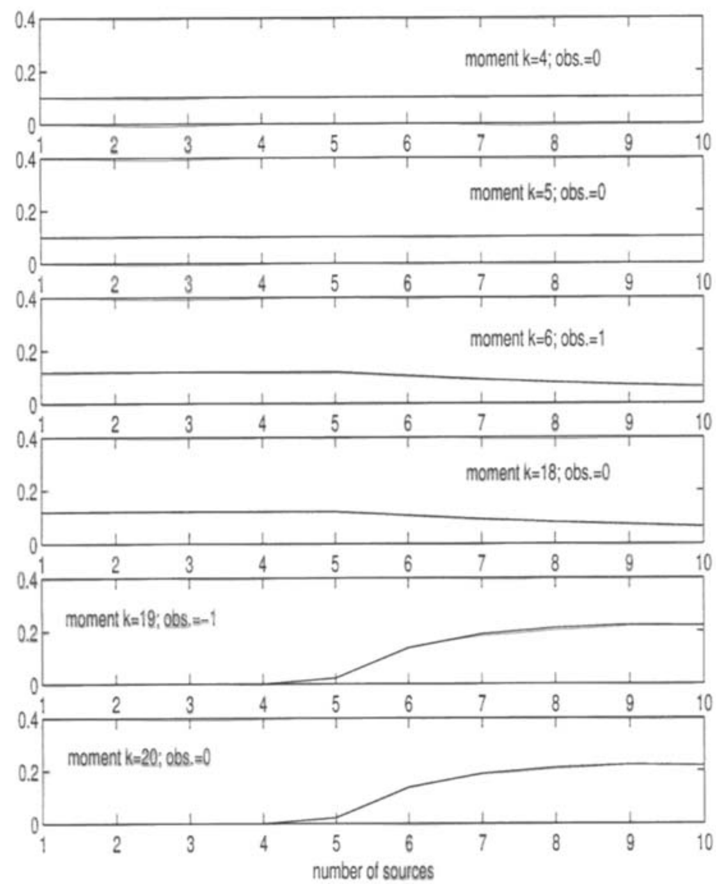

Figure 1: Probability over the number of active sources for a fixed source, at different times, for parameters: $T = 0.04$, $p = 0.001$, $N = 10$ (p is the probability that the chain switches to one more or one less active source, in a birth-and-death model). Observe the shifts in probability mass: positive ack at time 6, negative ack at time 19.

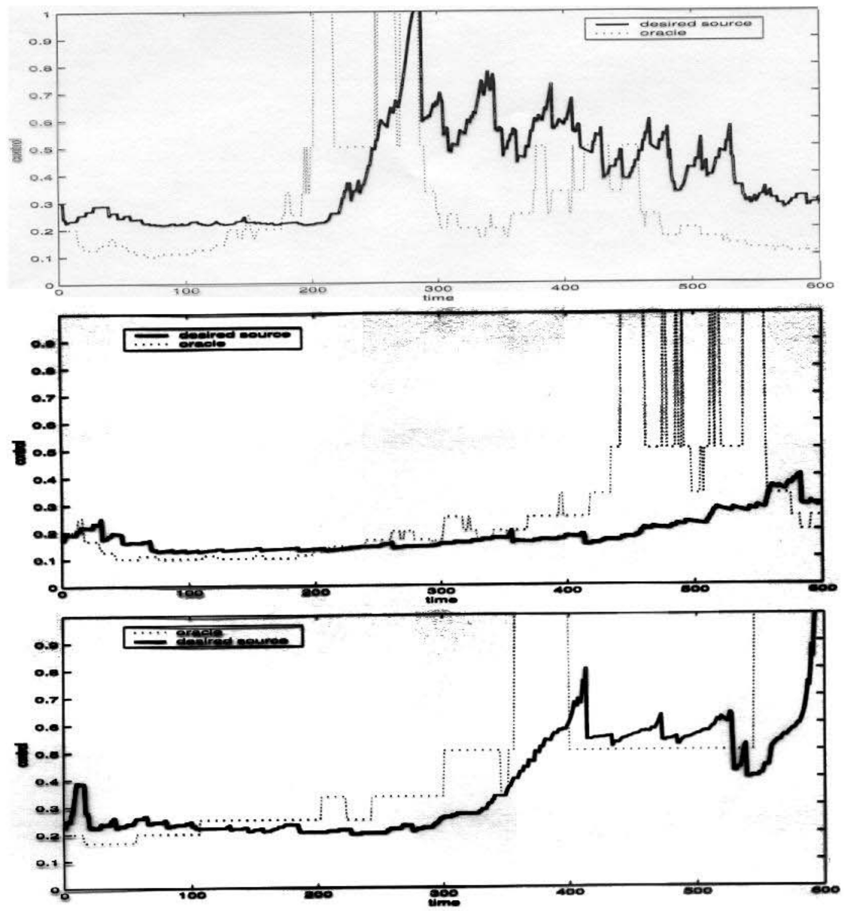

Figure 2: Control sequences for a fixed source and two different thresholds, a symmetric birth-and-death chain with transition probability p and $N = 10$ sources. Up: $T = 0.1, p = 0.1$, middle: $T = 0.02, p = 0.1$, down: $T = 0.05, p = 0.02$. "Oracle" refers to an ideal controller which can actually observe the hidden state.

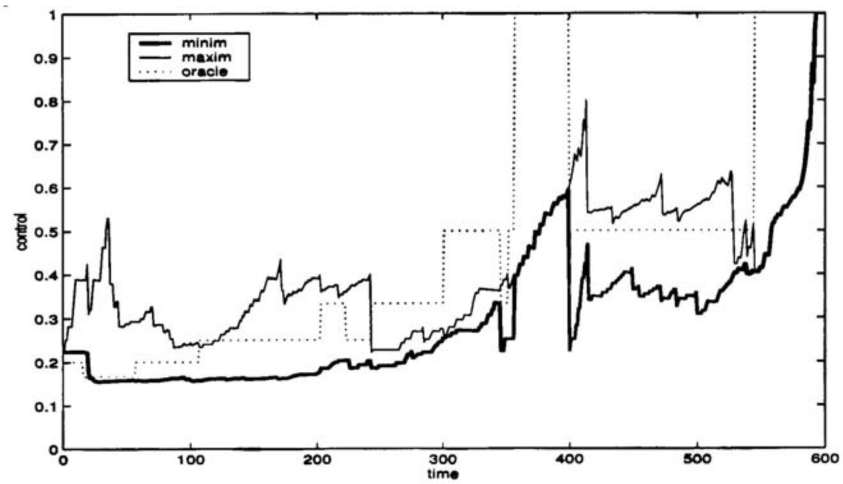

Figure 3: Minimum and maximum values of the control set over time. Note: the sources transmitting at minimum/maximum value are not necessarily the same at different time moments.

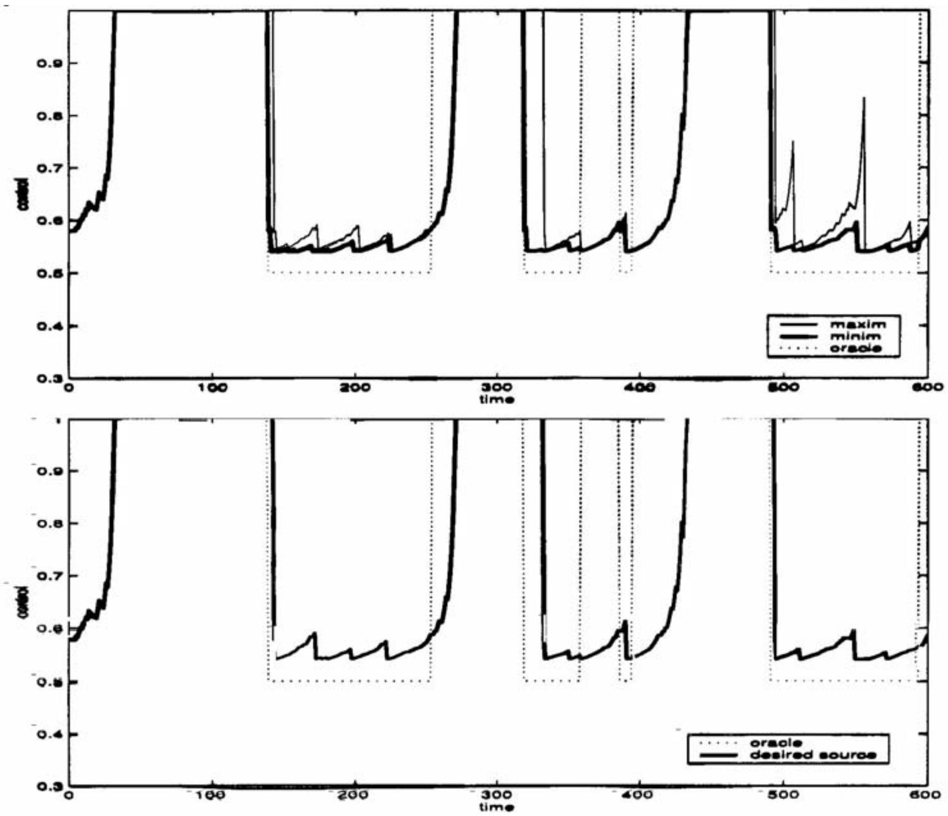

Figure 4: Control sequences for $p = 0.02$, $N = 2$ sources, and threshold $T = 0.04$. Up: minimum/maximum values, down: values for a fixed source.

III Performance Analysis

Using the controller from II, we have "decoupled the problem:

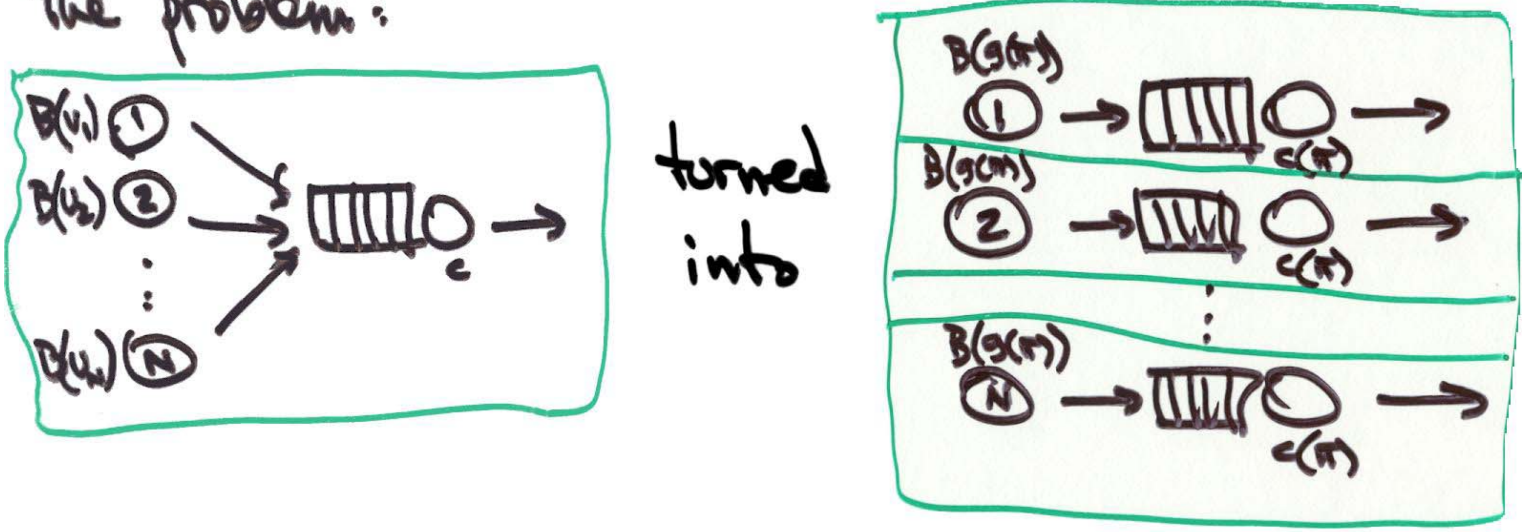

Q: Average throughput? Loss rates?

Dynamic behavior of π_k

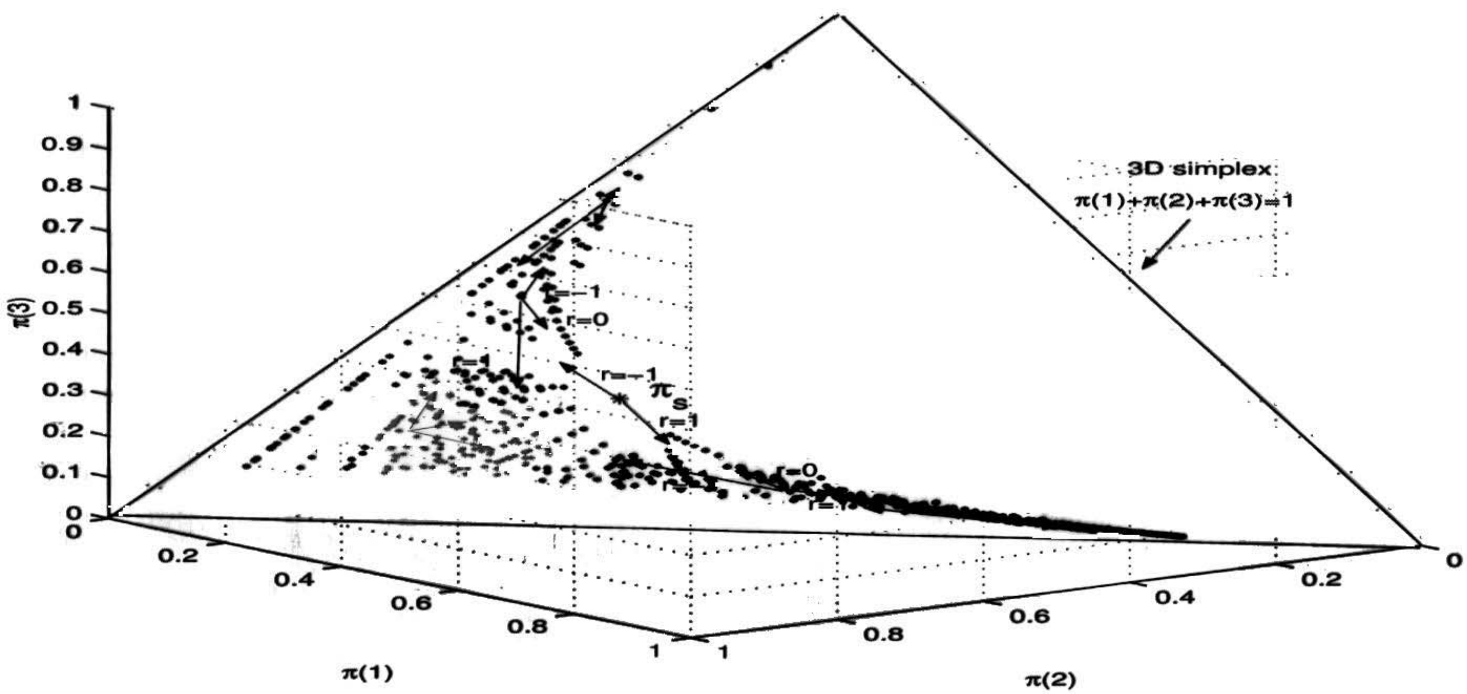

Performance Analysis

- Average throughput:

$$\lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} p(1 \mid x_k, g(\pi_k)) \overset{?}{=} \int_{\{x,\pi\}} p(1 \mid x, g(\pi))\, d\nu(x,\pi)$$

(note: this is $J(g)$!)

- Loss rate:

$$\lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} p(-1 \mid x_k, g(\pi_k)) \overset{?}{=} \int_{\{x,\pi\}} p(-1 \mid x, g(\pi))\, d\nu(x,\pi)$$

(strong suspicion: $= T$)

Q: Does ν exist???

Performance Analysis

Answer is yes but we will have to sweat it!

- T. Kaijser. A Limit Theorem for Partially Observed Markov Chains. Annals of Probability, 1975
- A J. Goldsmith, P Varaiya Capacity, Mutual Information & Coding for Finite State Markov Channels Trans IT, 1996.
- V Sharma S.K. Singh Entropy & Channel Capacity in the Regenerative Setup w/ App to Markov Channels. ISIT 2001

Feedback breaks one hypothesis or another in all cases

Performance Analysis

Elements of the proof. (limit of "simple measures")

1) π_k is Markov
2) For all "small enough" discretizations of the simplex Π $\exists$ an obs out of the cell with positive probability
3) π_s is reachable from anywhere
4) The discretized chain is positive recurrent
5) Limit measure exists.

Performance Analysis

Step 1: π_k is Markov

- π_k is everything we can infer about x_k
- Coupling due to control is only a <u>memoryless</u> function of π.

Note only true for the optimal controller.

Performance Analysis

Step 2: $\exists C > 0$, s.t. for any π, there is an observation r for which

$$\| \pi - F(\pi, g(\pi), r) \| \geq \varepsilon , \quad \forall\, 0 < \varepsilon \leq C$$

The only π for which $\pi \simeq F(\pi, g(\pi)\; r)$ is π_s. But $r = 1$ or $r = -1$ we saw here the effect of shifting mass around (away from π_s!)

Hence, C exists

Performance Analysis

Step 3: π_s is reachable from anywhere

$$\pi_{k+1} = C \; \pi_k \; D(u_k r_k) \; P,$$

with C : normalizing constant.
P : transition matrix of underlying chain

$$D(u, r) = \begin{pmatrix} p(r|1,u) & & 0 \\ & \ddots & \\ 0 & & p(r|N,u) \end{pmatrix}$$

So, $D(u, 0) = \begin{pmatrix} 1-u & & 0 \\ & \ddots & \\ 0 & & 1-u \end{pmatrix}$, constant

$$\pi_{k+1} = \pi_k P$$

Performance Analysis

Step 4: $q_\varepsilon(\pi)$ is positive recurrent on a non-empty set.

- Chain starts in $q_\varepsilon(\pi_S)$, wanders around, eventually comes back
  - $\Rightarrow$ $q_\varepsilon(\pi_S)$ is a recurrent state
- Any cell reachable from $q_\varepsilon(\pi_S)$ is also recurrent.

Note: we really don't know if every possible cell is recurrent...

Performance Analysis

Step 5. Limit exists: $\mu(g_\varepsilon(\pi)) \xrightarrow[\varepsilon \to 0]{} \mu$

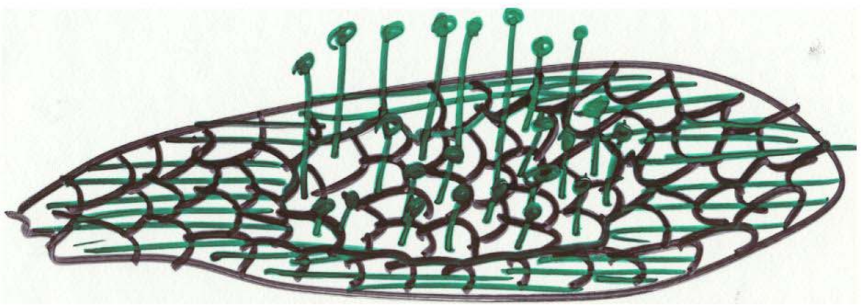

With step 4, we assign probabilities to any compact set contained in the simplex Π.

$\Longrightarrow$ a theorem of Meyn & Tweedie does the trick.

[S.P. Meyn, R.L. Tweedie MCs & Stochastic Stability Springer '93]

Summary & Conclusions

Interesting:

- Observations affect the state: send packet → communicate → learn something.
- Controlling arrivals introduces memory.

BUT, something that looks

is in fact

!

Summary & Conclusions

What next?

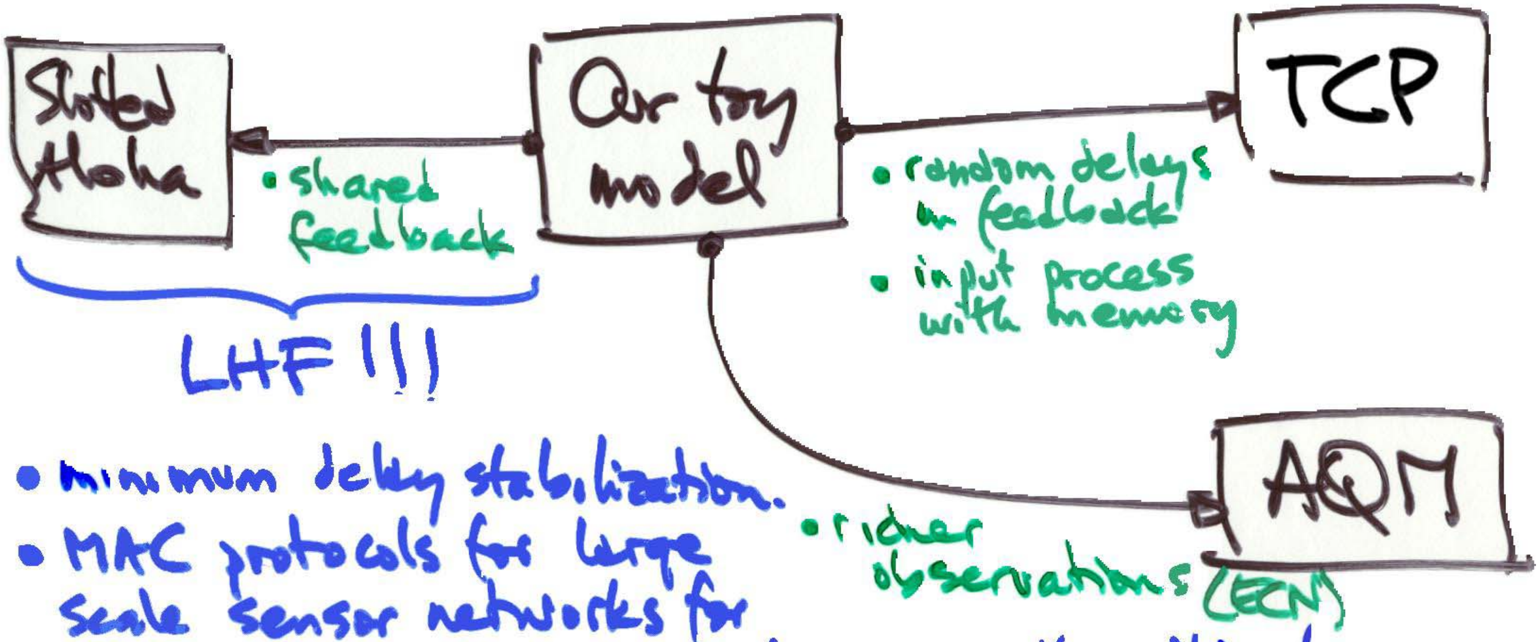

• minimum delay stabilization.

• MAC protocols for large scale sensor networks for which we can <u>prove</u> things!

• the ultimate performance bounds against which to measure heuristics

Main Corollary..

http://people.ece.cornell.edu/servetto/publications/