Congestion control: fairness, pricing and stability

Frank Kelly
Statistical Laboratory
University of Cambridge
f.p.kelly@statslab.cam.ac.uk

15th IEEE
Computer Communications Workshop
Captiva Island, Florida
16 October 2000

Acknowledgements

- Statistical Laboratory:
  Richard Gibbens, Ramesh Johari,
  Frank Kelly, Aman Maulloo, David Tan,
  Stephen Turner, Damon Wischik
- Microsoft Research:
  Paul Barham, Peter Key, Koenraad Laevens,
  Laurent Massoulié, Derek McAuley
- Economics:
  Robin Mason
- UCL, Computer Science:
  Jon Crowcroft, Philippe Oechslin
- Pointers to papers available at
  `www.statslab.cam.ac.uk/~frank/pf`
  `research.microsoft.com/network/disgame.asp`

Background

Heterogeneity of users, applications, capacities, routes, propagation delays

Variability over many time-scales, unpredictable evolution of traffic patterns

To an extent, the Internet copes by organizing itself: e.g. adaptive rate control

Where it fails to do so, difficulty with resource allocation or payment is often involved: e.g. interconnection agreements, "quality of service"

Can we understand how global behaviour emerges from packet level processes in the current Internet, and how this may/should evolve?

For example, can ECN allow a *simple* packet network to carry heterogeneous applications?

Important related work

- S.J. Golestani and S. Bhattacharyya:
  'A class of end-to-end congestion control algorithms for the Internet'
- S. H. Low and D. E. Lapsley:
  'Optimization Flow Control'
- S. Kunniyur and R. Srikant:
  'Congestion Control and Pricing for a Heterogeneous Internet'
- J. Mo and J. Walrand,
  R. J. La and V. Anantharam:
  delay based congestion control
- Pointers to papers available at
  Floyd and Mo's page:
  'Global optimization with end-to-end congestion control: pointers to the literature'
  `www.aciri.org/floyd/`

Outline

- Utility framework
- Max-min and proportional fairness
- Primal and dual algorithms
- Random effects and time-lags
- TCP
- Routing, admission control
- Some Internet pricing proposals
- References

The basic model

J — set of *resources*

R — set of *routes*

A route r is a subset $r \subset J$.

Let

$$\begin{aligned} A_{jr} &= 1 \quad \text{if } j \in r \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

For $j \in J$, let C_j be the *capacity* of resource j.

Each route has a *user*: if x_r is the *rate* on route r, then the *utility* to user r is $U_r(x_r)$.

$U_r(\cdot)$ — increasing, concave — *elastic* traffic.

(Assume $U_r(\cdot)$ continuously differentiable on $(0, \infty)$ with $U_r'(0) = \infty, U_r'(\infty) = 0$)

Let $C = (C_j, j \in J)$, $x = (x_r, r \in R)$. Then

$$Ax \leq C.$$

The system problem

$$\begin{aligned} &\text{maximize} && \sum_{r \in R} U_r(x_r) \\ &\text{subject to} && Ax \leq C \\ &\text{over} && x \geq 0. \end{aligned}$$

- Maximize aggregate utility, subject to capacity constraints

The user problem

$$\begin{aligned} &\text{maximize} && U_r\left(\frac{w_r}{\lambda_r}\right) - w_r \\ &\text{over} && w_r \geq 0. \end{aligned}$$

- User r chooses an amount to pay per unit time, w_r, and receives in return a flow $x_r = w_r / \lambda_r$

The network problem

$$\begin{array}{ll} \text{maximize} & \sum_{r\in R} w_r \log x_r \\ \text{subject to} & Ax \le C \\ \text{over} & x \ge 0. \end{array}$$

- *As if* the network maximizes a logarithmic utility function, but with constants $(w_r, r \in R)$ chosen by the users

Three optimization problems

SYSTEM(U, A, C)

$$\begin{array}{ll} \text{maximize} & \sum_{r\in R} U_r(x_r) \\ \text{subject to} & Ax \le C \\ \text{over} & x \ge 0. \end{array}$$

USER$_r(U_r; \lambda_r)$

$$\begin{array}{ll} \text{maximize} & U_r\left(\frac{w_r}{\lambda_r}\right) - w_r \\ \text{over} & w_r \ge 0. \end{array}$$

NETWORK$(A, C; w)$

$$\begin{array}{ll} \text{maximize} & \sum_{r\in R} w_r \log x_r \\ \text{subject to} & Ax \le C \\ \text{over} & x \ge 0. \end{array}$$

Theorem (problem decomposition)

There exist vectors $\lambda = (\lambda_r, r \in R)$, $w = (w_r, r \in R)$ and $x = (x_r, r \in R)$ such that

(i) $w_r = \lambda_r x_r$ for $r \in R$

(ii) w_r solves USER$_r(U_r; \lambda_r)$

(iii) x solves NETWORK$(A, C; w)$.

The vector x then also solves SYSTEM(U, A, C).

- Thus the system problem may be solved by solving simultaneously the network and user problems

Sketch of proof

Lagrangian for system:

$$L_{\text{system}}(x, z; \mu) = \sum_{r\in R} U_r(x_r) + \mu^T(C - Ax - z)$$

$$\frac{\partial}{\partial x_r} L_{\text{system}}(x, z; \mu) = U_r'(x_r) - \sum_{j\in r} \mu_j$$

Lagrangian for network:

$$L_{\text{network}}(x, z; \mu) = \sum_{r\in R} w_r \log x_r + \mu^T(C - Ax - z)$$

$$\frac{\partial}{\partial x_r} L_{\text{network}}(x, z; \mu) = \frac{w_r}{x_r} - \sum_{j\in r} \mu_j$$

User objective: $U_r\left(\frac{w_r}{\lambda_r}\right) - w_r$

$$\frac{\partial}{\partial w_r}\left[U_r\left(\frac{w_r}{\lambda_r}\right) - w_r\right] = \frac{1}{\lambda_r} U_r'\left(\frac{w_r}{\lambda_r}\right) - 1.$$

If user r chooses w_r so that $U_r'\left(\frac{w_r}{\lambda_r}\right) = \lambda_r$, then earlier derivatives become identical, under the identification

$$x_r = \frac{w_r}{\lambda_r}, \quad \lambda_r = \sum_{j \in r} \mu_j.$$

- μ_j - *shadow price* per unit of flow through resource j
- λ_r - shadow price along route r

$$\max_{x,z \geq 0} L_{\text{network}}(x, z; \mu) = \sum_{r \in R} w_r \log \frac{w_r}{\sum_{j \in r} \mu_j} + \mu^T C,$$

and so the dual to NETWORK$(A, C; w)$ becomes:

DUAL$(A, C; w)$

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{r \in R} w_r \log\Big(\sum_{j \in r} \mu_j\Big) - \sum_{j \in J} \mu_j C_j \\ \text{over} & \mu \geq 0. \end{array}$$

Max-min fairness

$x = (x_r, r \in R)$ is *feasible* if $x \geq 0$ and $Ax \leq C$.

$x = (x_r, r \in R)$ is *max-min fair* if it is feasible, and if, for any other feasible vector y,

$$\exists r : y_r > x_r \Rightarrow \exists s : y_s < x_s \leq x_r.$$

- The compactness and convexity of feasible region imply that such a vector exists and is unique
- As a concept, max-min fairness has been much discussed by political philosophers (major work: J. Rawls, *A Theory of Justice*, 1971)

Proportional fairness

$x = (x_r, r \in R)$ is *proportionally fair* if it is feasible, and if, for any other feasible vector y, the aggregate of proportional changes is zero or negative:

$$\sum_{r \in R} \frac{y_r - x_r}{x_r} \leq 0.$$

Let $w = (w_r, r \in R)$ be a vector of weights, or charges.

$x = (x_r, r \in R)$ is *weighted proportionally fair* if it is feasible, and if, for any other feasible vector y,

$$\sum_{r \in R} w_r \frac{y_r - x_r}{x_r} \leq 0.$$

Proposition

A vector x solves NETWORK$(A, C; w)$ if and only if it is weighted proportionally fair.

Sketch of proof

Consider a perturbation from x to y, where $y_r = x_r + \delta x_r, r \in R$. This increases the objective function $\sum_{r \in R} w_r \log x_r$ by an amount

$$\sum_{r \in R} w_r \frac{\delta x_r}{x_r} + o(\delta x).$$

Thus result follows, from convexity of the feasible region and strict concavity of the objective function.

Solution of the network problem

- Strategy: design algorithms, based on the primal or dual problems for the network, to implement proportional fairness
- Several algorithms possible - try to mimic design choices made in existing standards (e.g. Jacobson's TCP, or explicit rate algorithms for ABR)

A primal algorithm

$$\begin{aligned} \frac{d}{dt}x_r(t) &= \kappa_r\Big(w_r - x_r(t)\sum_{j\in r}\mu_j(t)\Big) \\ \mu_j(t) &= p_j\Big(\sum_{s:j\in s}x_s(t)\Big) \end{aligned}$$

where $p_j(\cdot)$ is a non-negative, continuous, increasing function.

Interpretation:

- resource j generates feedback signals at rate $\mu_j(t)$
- signals reach each user r whose route passes through resource j
- multiplicative decrease in flow x_r at rate proportional to stream of feedback signals received
- linear increase in flow x_r at rate proportional to w_r

Theorem (global stability)

$$\mathcal{U}(x) = \sum_{r\in R} w_r \log x_r - \sum_{j\in J}\int_0^{\sum_{s:j\in s}x_s} p_j(y)dy$$

is a Lyapunov function for the primal algorithm. The unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.

Proof

$$\frac{\partial}{\partial x_r}\mathcal{U}(x) = \frac{w_r}{x_r} - \sum_{j\in r} p_j\Big(\sum_{s:j\in s}x_s\Big),$$

and so

$$\frac{d}{dt}\mathcal{U}\big(x(t)\big) = \sum_{r\in R}\frac{\partial U}{\partial x_r}\frac{d}{dt}x_r(t)$$

$$= \sum_{r\in R}\frac{\kappa_r}{x_r(t)}\left(w_r - x_r(t)\sum_{j\in r}p_j\Big(\sum_{s:j\in s}x_s(t)\Big)\right)^2$$

$\geq 0.$

Note that

$$C_j\big(\sum_{s:j\in s}x_s\big) = \int_0^{\sum_{s:j\in s}x_s} p_j(y)dy$$

is a cost function penalizing proximity to the capacity constraint for resource j.

If

$$\begin{aligned} p_j(y) &= \infty \qquad y > C_j \\ &= 0 \qquad y \leq C_j \end{aligned}$$

then maximization of the Lyapunov function $\mathcal{U}(x)$ becomes the primal problem NETWORK$(A, C; w)$.

But we shall see later that, in the presence of propagation delays, large values of the derivatives $p_j'(\cdot)$ compromise stability.

Rate of convergence

Let x be unique vector maximizing $\mathcal{U}(x)$, let $\mu_j = p_j(\sum_{s:j\in s} x_s)$ and suppose $p_j(\cdot)$ is differentiable at this point, with derivative p_j'. Let

$$x_r(t) = x_r + x_r^{1/2} y_r(t).$$

Then, linearizing about x,

$$\begin{aligned}\frac{d}{dt}y(t) &= -\kappa(WX^{-1} + X^{1/2}A^T P' A X^{1/2})y(t)\\ &= -\kappa\Gamma^T\Phi\Gamma y(t)\end{aligned}$$

where

$$\begin{aligned}X &= \mathrm{diag}(x_r, r\in R)\\ W &= \mathrm{diag}(w_r, r\in R)\\ P' &= \mathrm{diag}(p_j', j\in J)\\ \Phi &= \mathrm{diag}(\phi_r, r\in R) \qquad \text{– the matrix}\end{aligned}$$

of (positive, real) eigenvalues.

Random effects

$$dx_r(t) =$$
$$\kappa_r\left(w_r dt - \sum_{j\in r}\varepsilon dN_{jr}\left(\varepsilon^{-1}\int_0^t x_r(\tau)\mu_j(\tau)d\tau\right)\right)$$

where $(N_{jr}(\tau), \tau\geq 0), j\in J, r\in R$, are independent unit rate Poisson processes.

Interpretation:

- feedback signals from resource j to user r arise at rate $\varepsilon^{-1}x_r(t)\mu_j(t)$
- user r reacts to a feedback signal by reducing $x_r(t)$ by an amount $\kappa_r\varepsilon$

General Brownian model

$$dy(t) = -\kappa\big(\Gamma^T\Phi\Gamma y(t)dt + FdB(t)\big)$$

$$F - |R|\times|I| \quad \text{matrix}$$

$B(t) = \big(B_i(t), i\in I\big)$ – independent standard Brownian motions.

Stationary solution?

$$y(t) = -\kappa\int_{-\infty}^t e^{-\kappa(t-\tau)\Gamma^T\Phi\Gamma}FdB(\tau).$$

But this is just a linear transformation of the Gaussian process $(B(\tau), \tau < t)$. Thus

$$y(t)\sim N(0,\Sigma)$$

where

$$\begin{aligned}\Sigma &= \mathbb{E}[y(t)y(t)^T]\\ &= \kappa^2\int_{-\infty}^0 e^{\kappa\tau\Gamma^T\Phi\Gamma}FF^Te^{\kappa\tau\Gamma^T\Phi\Gamma}d\tau\\ &= \kappa\Gamma^T\left[\int_{-\infty}^0 e^{\tau\Phi}\Gamma FF^T\Gamma^Te^{\tau\Phi}d\tau\right]\Gamma.\end{aligned}$$

Define $[\Gamma F;\Phi]$ by

$$\begin{aligned}[\Gamma F;\Phi]_{rs} &= \left[\int_{-\infty}^0 e^{\tau\Phi}\Gamma FF^T\Gamma^Te^{\tau\Phi}d\tau\right]_{rs}\\ &= \frac{[\Gamma FF^T\Gamma^T]_{rs}}{\phi_r+\phi_s}.\end{aligned}$$

Then

$$\Sigma = \kappa\Gamma^T[\Gamma F;\Phi]\Gamma.$$

A dual algorithm

$$\frac{d}{dt}\,\mu_j(t) = \kappa_j\Big(p_j\big(\sum_{r:j\in r} x_r(t)\big) - \mu_j(t)\Big)$$

$$x_r(t) = \frac{w_r}{\sum_{k\in r}\mu_k(t)}$$

Interpretation:

- user r has rate x_r fixed by prices μ_j
- rate of change of prices $\mu_j(t)$ proportional to measured error of current prices

Theorem (global stability)

$$\mathcal{V}(\mu) = \sum_{r\in R} w_r \log\Big(\sum_{j\in r}\mu_j\Big) - \sum_{j\in J}\int_0^{\mu_j} p_j^{-1}(\eta)d\eta$$

is a Lyapunov function for the dual algorithm. The unique value μ maximizing $\mathcal{V}(\mu)$ is a stable point of the system, to which all trajectories converge.

User adaptation

Suppose user r monitors its rate $x_r(t)$, and chooses $w_r(t)$ to track optimum to $USER_r(U_r;\lambda_r(t))$, where $\lambda_r(t) = w_r(t)/x_r(t)$.

Let $D_r(\lambda)$ be the resulting *demand* of user r when sensing a charge per unit flow of λ.

Suppose also

$$p_j(y) = \frac{d}{dy}\,C_j(y).$$

Then both the primal algorithm and the dual algorithm maximize

$$\mathcal{U}(x) = \sum_{r\in R} U_r(x_r) - \sum_{j\in J} C_j\Big(\sum_{s:j\in s} x_s\Big).$$

This function is a Lyapunov function for the primal algorithm, while its conjugate

$$\mathcal{V}(\mu) = \sum_{r\in R}\int^{\sum_{j\in r}\mu_j} D_r(\zeta)d\zeta - \sum_{j\in J}\int_0^{\mu_j} p_j^{-1}(\eta)d\eta$$

is a Lyapunov function for the dual algorithm.

MulTCP (congestion avoidance)

Jacobson's TCP maintains a window of transmitted but not yet acknowledged packets; the rate x and the window size `cwnd` satisfy the approximate relation $\texttt{cwnd} = xT$.

- each positive acknowledgement increases the window size `cwnd` by $1/\texttt{cwnd}$
- each congestion indication halves the window size.

Crowcroft and Oechslin's MulTCP: users are allowed to set a parameter m, which

- multiplies by m the rate of additive increase
- makes $1 - 1/2m$ the multiplicative decrease factor

in Jacobson's algorithm. MulTCP behaves in many respects as a collection of m single TCP connections.

MulTCP (differential equations)

$$\frac{d}{dt}\,x_r(t) = \frac{m_r}{T_r^2} - \left(\frac{m_r}{T_r^2} + \frac{x_r(t)^2}{2m_r}\right)\sum_{j\in r}\mu_j(t)$$

Interpretation:

- for MulTCP, expected change in the congestion window `cwnd` per update step is about
$$\frac{m}{\texttt{cwnd}}\,(1-p) - \frac{\texttt{cwnd}}{2m}\,p$$
where p is the probability of congestion indication at the update step.
- time between update steps is about $T/\texttt{cwnd}$
- change in rate x per unit time is about
$$\frac{\left(\frac{m}{\texttt{cwnd}}\,(1-p) - \frac{\texttt{cwnd}}{2m}\,p\right)/T}{T/\texttt{cwnd}} = \frac{m}{T^2} - \left(\frac{m}{T^2} + \frac{x^2}{2m}\right)\,p.$$

Theorem (global stability)

$$\mathcal{U}(x) = \sum_r \frac{\sqrt{2}m_r}{T_r} \arctan\left(\frac{x_r T_r}{\sqrt{2}m_r}\right) - \sum_{j\in J} \int_0^{\sum_{s:j\in s} x_s} p_j(y)dy$$

is a Lyapunov function for the MulTCP differential equations. The unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.

Proof

$$\frac{\partial}{\partial x_r}\mathcal{U}(x) = \frac{m_r}{T_r^2}\left(\frac{m_r}{T_r^2} + \frac{x_r^2}{2m_r}\right)^{-1} - \sum_{j\in r} p_j\Big(\sum_{s:j\in s} x_s\Big)$$

and so

$$\begin{aligned} \frac{d}{dt}\mathcal{U}(x(t)) &= \sum_{r\in R} \frac{\partial \mathcal{U}}{\partial x_r}\cdot\frac{d}{dt}x_r(t) \\ &= \sum_{r\in R}\left(\frac{m_r}{T_r^2} + \frac{x_r(t)^2}{2m_r}\right)^{-1}\left(\frac{d}{dt}x_r(t)\right)^2 \\ &\geq 0. \end{aligned}$$

Current TCP algorithm?

Rate allocated to route r is approximately

$$x_r = \frac{c}{T_r\left(\sum_{j\in r} p_j\right)^{1/2}}.$$

This rate is optimal for a system where

$$U_r(x_r) = constant - \frac{c^2}{T_r^2 x_r}$$

(rapidly diminishing returns as x_r increases, and an inverse square law with distance) and

$$C_j(y) = \int_0^y p_j(\eta)d\eta$$

(should be $yp_j(y)$, if $p_j(\cdot)$ drop probability).

The resource share for route r is proportional to

$$\frac{1}{T_r}\Big(\sum_{j\in r} p_j\Big)^{1/2}$$

(relative to proportional fairness, distance is over-penalized and use of congested resources is under-penalized).

Time-lags and stability

$$\frac{d}{dt}x(t) = \kappa\Big(w - x(t-T)p(x(t-T))\Big)$$

where T is the round-trip time.

The unique equilibrium point of this system does not depend on the round-trip time T, but its stability does.

By linearizing this delay equation we find that its equilibrium point is locally stable and non-oscillatory if

$$\kappa T\,(p + xp') < e^{-1}$$

For the dual algorithm

$$\frac{d}{dt}\mu(t) = \kappa\Big(p\Big(\frac{w}{\mu(t-T)}\Big) - \mu(t)\Big)$$

equivalent criterion is

$$\kappa T \exp(1+\kappa T)\frac{xp'}{p} < 1.$$

Note importance of the product κT, and of the derivative p'.

Primal or dual?

How should the gain parameters $\kappa_r, r\in R$, for the primal, or $\kappa_j, j\in J$, for the dual, be chosen? Equivalently, how should the time-constants (κ^{-1}) in EWMA estimation schemes be chosen?

The primal averages at end-systems, while the dual averages at resources.

Advantages of averaging at end-systems:

- easier to control $\kappa_r T_r$ (e.g. self-clocking window algorithms can do this implicitly); harder to control $\kappa_j T_r$ uniformly over r;
- price (like traffic) has fluctuations on all time-scales; users with short RTTs can (and should) exploit fast fluctuations
- choice of κ_r affects the *variance* of x_r - different users may have different preferences.

Currently TCP does averaging at end-systems. RED attempts some averaging at resources, but time-constants hard to choose!

Routing

S — set of *users*

A user s is a subset $s \subset R$. Let

$$\begin{aligned} H_{sr} &= 1 \quad \text{if } r \in s \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

Let $s(r)$ be the (unique) user served by route r.

Primal algorithm:

$$\frac{d}{dt} y_r(t) = \kappa\Big(w_{s(r)} - \Big(\sum_{a \in s(r)} y_a(t)\Big) \sum_{j \in r} \mu_j(t)\Big)$$

If

$$\mathcal{U}(y) = \sum_{s \in S} w_s \log\Big(\sum_{r \in s} y_r\Big) - \sum_{j \in J} \int_0^{\sum_{r:j \in r} y_r} p_j(y)dy.$$

then for the above dynamical system

$$\frac{d}{dt}\mathcal{U}\big(y(t)\big) > 0$$

except at optimum.

Distributed admission control

Routes $s \in S$ carry non-adaptive traffic. When a call for route s arrives a set of m_s probe packets are transmitted along route s, and the call is rejected and lost if any of these probe packets is marked. Otherwise the call is accepted.

ν_s — offered load on route s.

$y_s(t)$ — aggregate flow along route s.

$$\frac{d}{dt} y_s(t) = \nu_s \left(1 - m_s \sum_{j \in s} \mu_j(t)\right) - y_s(t)$$

$$\mu_j(t) = p_j \left(\sum_{r:j \in r} x_r(t) + \sum_{s:j \in s} y_s(t)\right)$$

A Lyapunov function is:

$$\mathcal{U}(x, y) = \sum_{r \in R} w_r \log x_r + \sum_{s \in S} \frac{1}{m_s}\left(y_s - \frac{y_s^2}{2\nu_s}\right) - \sum_{j \in J} \int_0^{\sum_{r:j \in r} x_r + \sum_{s:j \in s} y_s} p_j(z)dz$$

Some Internet pricing proposals

- Smart Market
  (MacKie-Mason and Varian, 1994)
  —a (dynamic) price per packet, calculated by an auction mechanism at a resource
- Paris Metro Pricing
  (Odlyzko, 1997)
  —two classes of service with different (fixed) prices per packet
- Proportionally Fair Pricing

Proportionally Fair Pricing

- mark packets as overload approached
- charge a receiver a fixed small amount for each marked packet received
- allow intelligent end-nodes to determine their demands on the network

Attempts to combine

- flexibility of smart market
- simplicity of PMP
- evolution from existing standards and proposals (TCP, RED and ECN)

Evolution

- compatible with TCP and RED (with ECN)
- allows marking and flow control strategies to evolve
- able to support arbitrarily differentiated services, defined by users
- no need for large buffers or multiple priorities within network, or for connection acceptance control at edge

Example: multilayer multicast

- at points of replication, an ECN mark is inherited by a single (randomly chosen) daughter packet for a top rate layer
- receivers probe layers

References

- Basic model:
  Rate control in communication networks: shadow prices, proportional fairness and stability
  FK, Aman Maulloo and David Tan
  Journal of the Operational Research Society 49 (1998) 237-252.
- Sample path shadow prices:
  Resource pricing and the evolution of congestion control Richard Gibbens and FK
  Automatica 35 (1999), 1969-1985.
  How to mark fairly Damon Wischik
  ISQE - Workshop on Internet Service Quality Economics, MIT (1999).
- MulTCP:
  Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP
  Jon Crowcroft and Philippe Oechslin
  ACM Computer Communications Review 28 (1998), 53-67.
- Pricing:
  Service differentiation: congestion pricing, brokers and bandwidth futures
  Peter Key. NOSSDAV99

- Distributed network game:
  Differential QoS and pricing in networks: where flow-control meets game theory
  Peter Key and Derek McAuley
  IEE Proceedings Software 146 (1999), 39-43.
  User policies in a network implementing congestion pricing
  Peter Key and Laurent Massoulié - ISQE
  The use of games to assess user strategies for differential Quality of Service in the Internet
  Richard Gibbens and Peter Key - ISQE
- Non-adaptive streams:
  Distributed admission control
  FK, Peter Key and Stan Zachary
  IEEE Journal on Selected Areas in Communications, 18 (2000)
  Models for a self-managed Internet
  FK, Phil. Trans. Royal Society A358 (2000) 2335-2348.
- Pointers to papers available at
  `www.statslab.cam.ac.uk/~frank/pf`
  `research.microsoft.com/network/disgame.asp`