Economics and dynamics in networking: three case studies

Fernando Paganini
Universidad ORT, Uruguay

Collaborators:
- Enrique Mallada (Univ. ORT, now at Cornell).
- Andrés Ferragut (Univ. ORT and Univ. de la Rep, Uruguay)
- Pablo Belzarena (Universidad de la Republica, Uruguay)

Outline:
1. Congestion control with multipath routing.
2. Controlling fairness through number of TCP connections.
3. Auctions for resource allocation in overlay networks.

1. Congestion control with multipath routing

Optimal WELFARE problem:

max $\sum U_k(x^k)$ subject to $y \le c$.

- Rate y_l
- Capacity c_l.

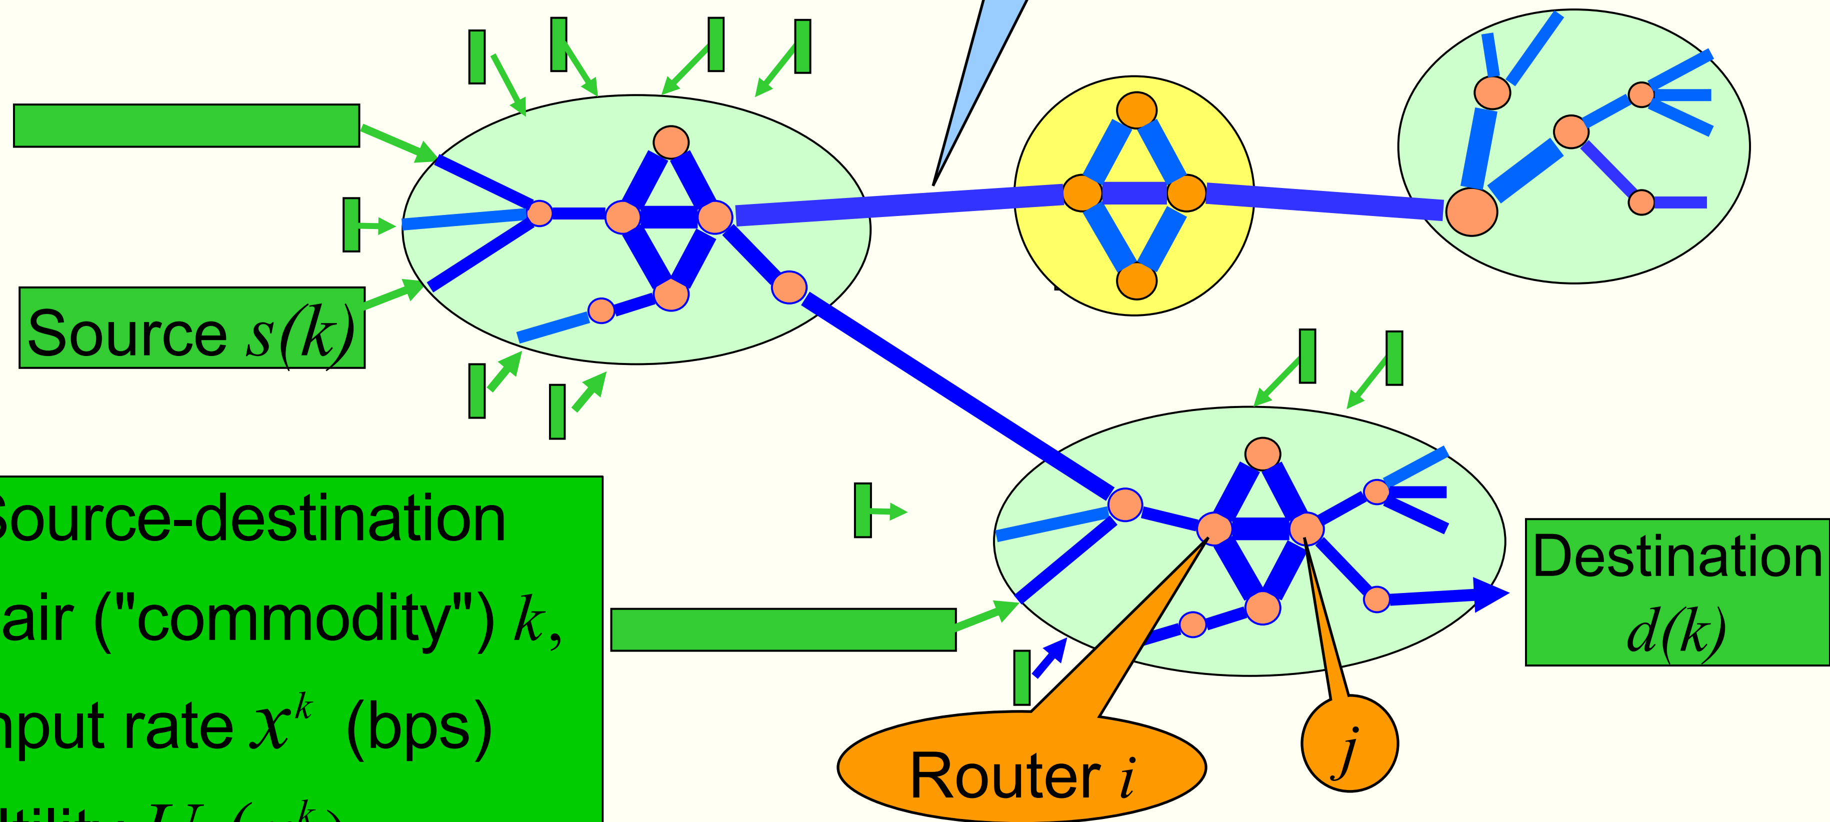

Source-destination pair ("commodity") k, input rate x^k (bps)

Utility $U_k(x^k)$.

$\alpha_{i,j}^d$: fraction of traffic destined to d, sent to node j. $\alpha_i^d \ge 0$, $\sum_j \alpha_{i,j}^d = 1$.

A simple network

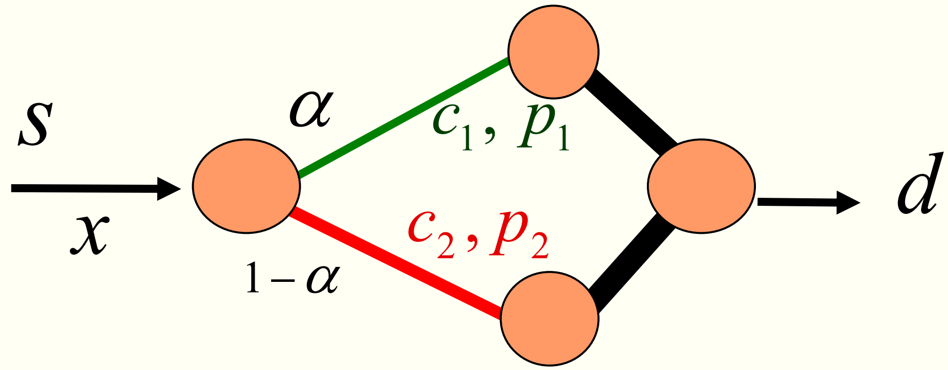

The "customer":

- elastic traffic source, rate follows "demand curve" $x = f(q)$, $f = U'^{-1}$
- $q = \alpha p_1 + (1-\alpha) p_2$, mean price.

The "broker": multipath router, adapts routing fraction α slowly in direction of cheaper prices,

$\dot{\alpha} = \beta(p_2 - p_1)$.

The resources: link capacities. Prices p_1, p_2 (e.g. queueing delays) indicate their scarcity.

Optimal welfare equililbrium:

$$x = c_1 + c_2 = f(q),$$

$$\alpha = \frac{c_1}{c_1 + c_2}, \quad p_1 = p_2 = q,$$

Does the system reach this equilibrium?

We implemented this in the packet simulator ns2:

- Source runs TCP-FAST, responds to delay.
- Router split traffic, adapt split to measured delays.

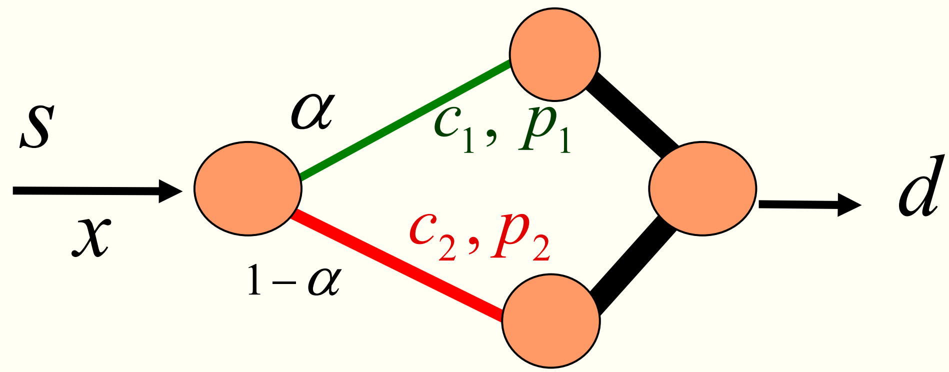

Sources rate reaches desired equilibrium.

but....

Queues and routing splits oscillate!

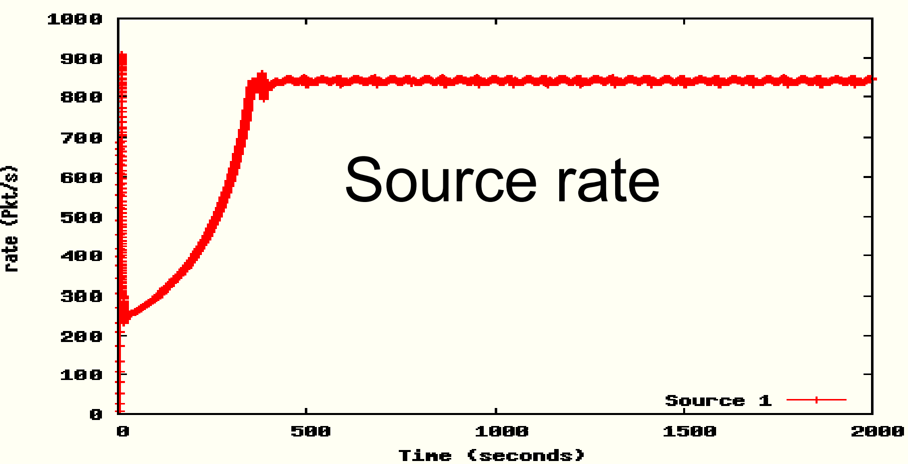

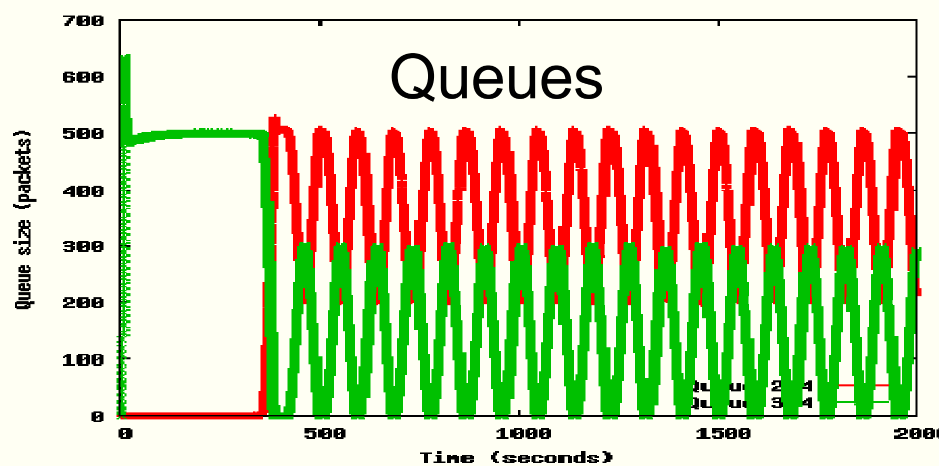

Can we explain this with flow models?

A commonly model for delay is the "latency function" $p_l = \varphi(y_l)$, where y_l is the link rate.

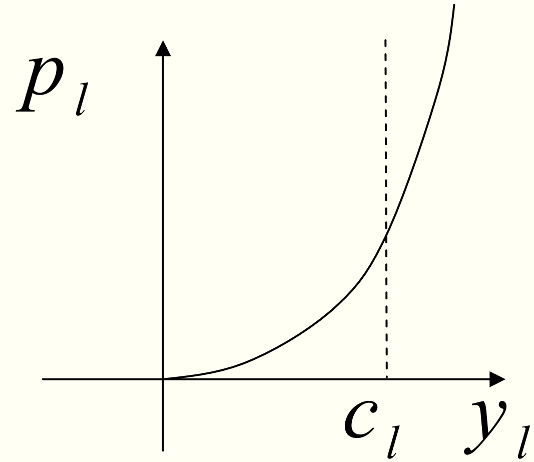

This model implies global convergence to equilbrium in α, p, x, using a Lyapunov argument.

Something's wrong....

The latency model, from queueing theory in *steady state*, not relevant to dynamic studies.

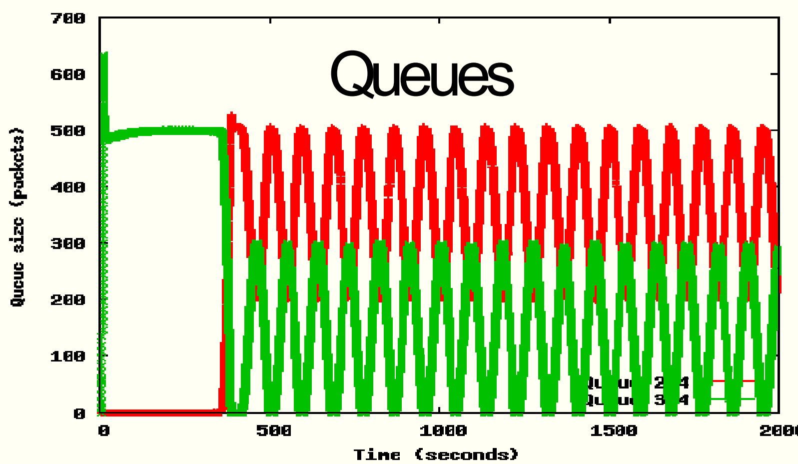

Another fluid model for queueing delay $\dot{p}_l = \frac{1}{c_l}\left[y_l - c_l\right]^+_{p_l}$

We can prove $x \to c_1 + c_2$.

Equilibrium:

$$x^* = c_1 + c_2, \quad \alpha^* = \frac{c_1}{x^*},$$

$$p_1^* = p_2^* = q^* = U'(x^*).$$

Take $x \equiv c_1 + c_2$, around equilibrium

we have the linear dynamics

$$\begin{bmatrix} \delta\dot{\alpha} \\ \delta\dot{p}_1 \\ \delta\dot{p}_2 \end{bmatrix} = \begin{bmatrix} 0 & -\beta & \beta \\ \gamma_1 x^* & 0 & 0 \\ -\gamma_2 x^* & 0 & 0 \end{bmatrix} \begin{bmatrix} \delta\dot{\alpha} \\ \delta p_1 \\ \delta p_2 \end{bmatrix} \Rightarrow$$ UNSTABLE (imaginary modes)

Mass-spring like system, frequency matches packet simulations.

Conclusion:

- Route oscillations are not trivial to avoid.
- A naive equilibrium viewpoint (most of econ) misses all this.
- Beware of simplistic models for delay!

Solving the problem

Adapt α based on anticipated (rather than current) price $\boxed{\pi_l = p_l + \nu \dot{p}_l}$

In control terms, add derivative action. Same equilibrium. Simulations:

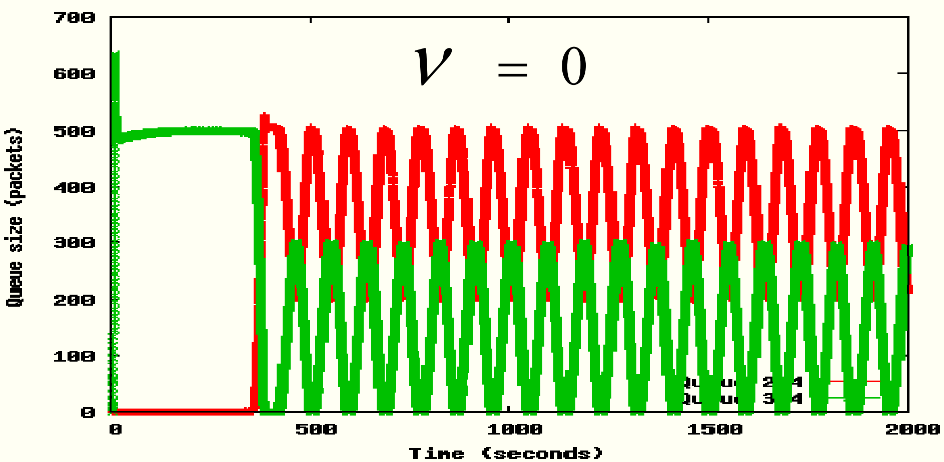

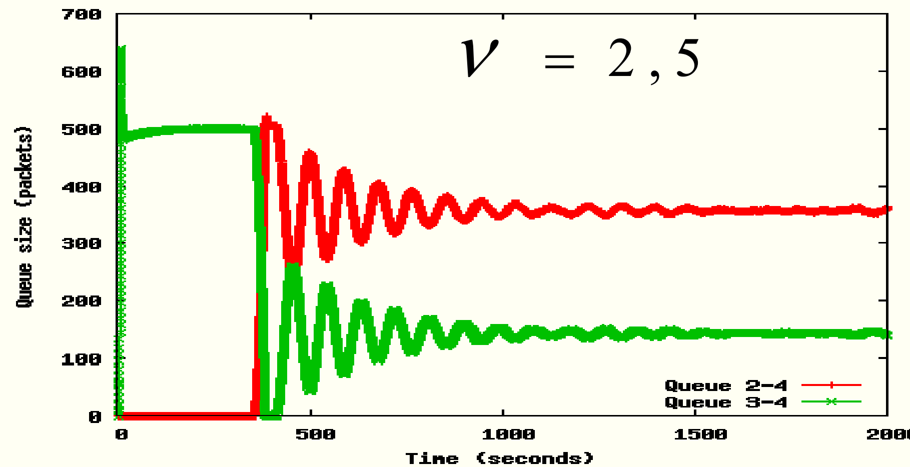

Theorems [P'-Mallada, to appear in IEEE ToN]

the equilibrium point (optimum max $\sum U^k(x^k)$)

is locally asymptotically stable in an arbitrary network

Packet implementation: variants of TCP-FAST and RIP.

2. Stability and user-level fairness

[Ferragut-P', CISS '08, Stochastic Networks '08]

User rate ρ_k

Utility $U_k(\rho_k)$

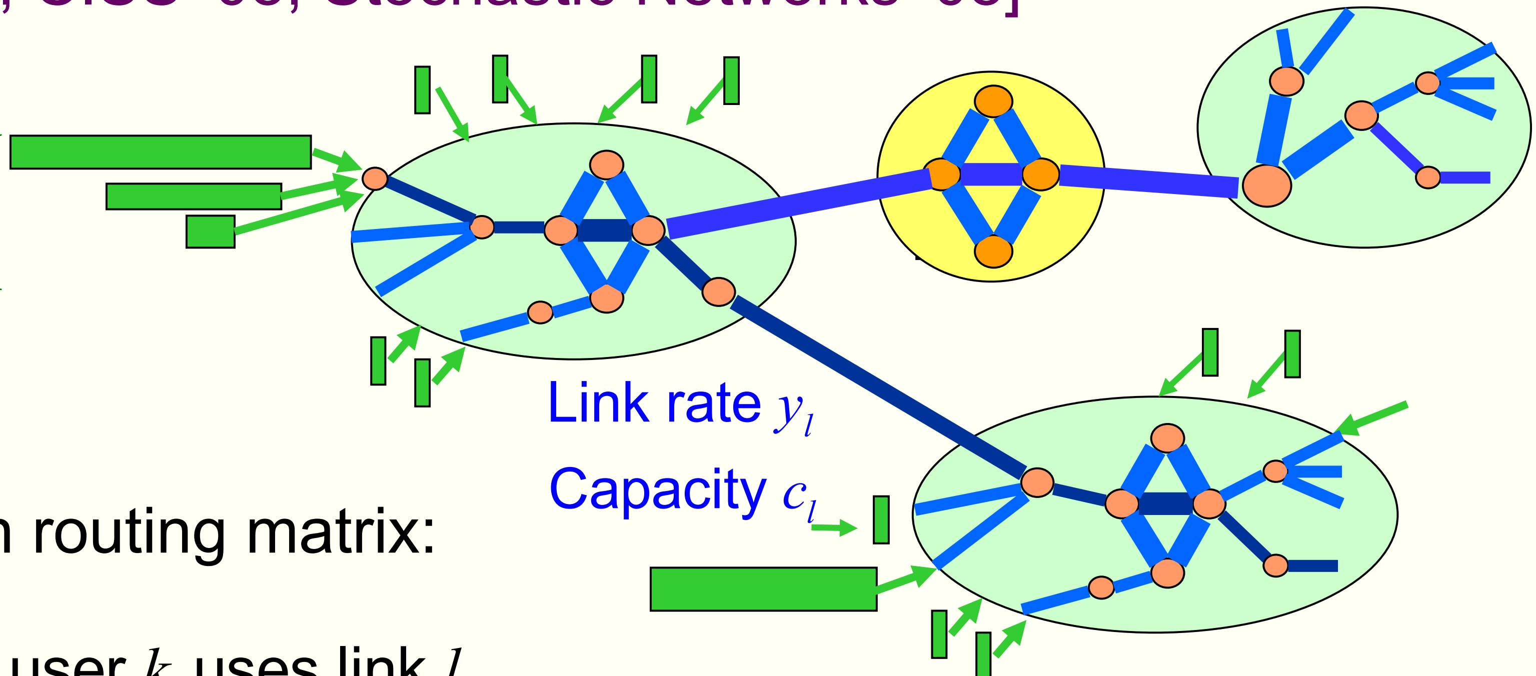

Single path routing matrix:

$$R_{lk} = \begin{cases} 1 & \text{if user } k \text{ uses link } l \\ 0 & \text{otherwise} \end{cases} \qquad y = R\rho$$

$$\boxed{KELLY's\ SYSTEM\ PROBLEM} : \max_{\rho} \sum_k \underbrace{U_k(\rho_k)}_{\text{USER UTILITY FUNCTION}}, \quad \text{subject to} \quad \underbrace{R\rho \le c}_{\text{LINK CAPACITY CONSTRAINTS}}$$

Contrast with flow-level fairness of TCP

Rate x_k per flow $\rightarrow$
Utility $U_{TCPk}(x_k)$
n_k flows per user.

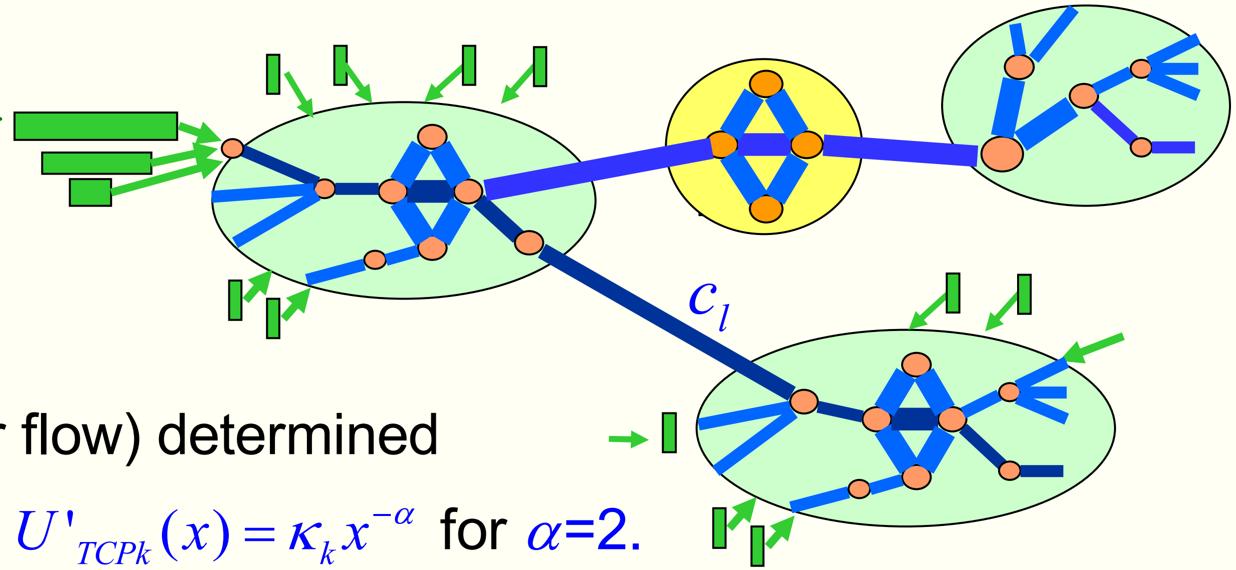

TCP utility U_{TCPk} (per flow) determined

by the protocol, e.g., $U'_{TCPk}(x) = \kappa_k x^{-\alpha}$ for α=2.

$$\boxed{\begin{array}{c} TCP: NETWORK \\ PROBLEM \end{array}} : \max_x \sum n_k \underbrace{U_{TCPk}(x_k)}_{\text{TCP UTILITY FUNCTION}} \quad \text{subject to} \sum_k R_{lk} \underbrace{n_k x_k}_{\rho_k} \le c_l$$

- Without control of number of connections, fairness per flow is moot (Briscoe'07).
- Incentives to employ many TCP flows (e.g., p2p) . Tragedy of the commons?

On stochastic stability of a network served by TCP

[deVeciana et al '99, Bonald-Massoulié '01]

User: Poisson (λ_k) arrivals, $\exp(\mu_k)$ workloads. $\rightarrow$

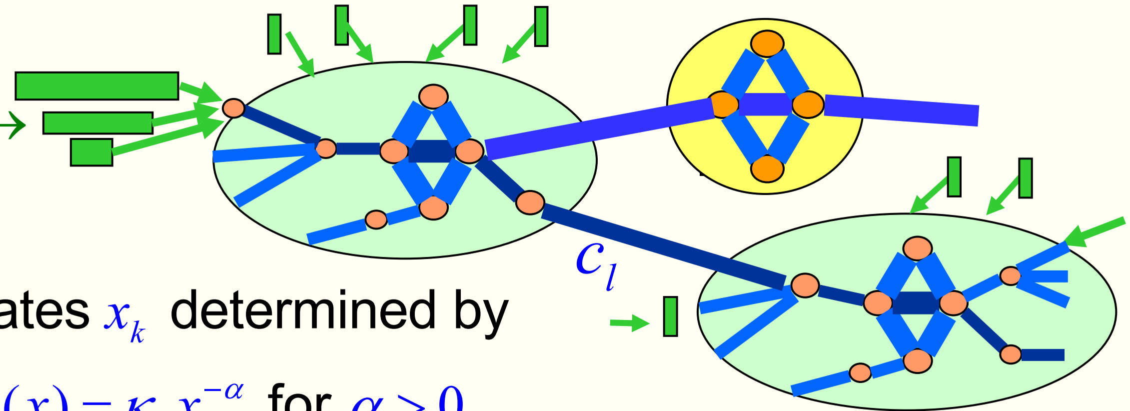

For each fixed $\{n_k\}$, service rates x_k determined by TCP congestion control $U'_{TCPk}(x) = \kappa_k x^{-\alpha}$ for $\alpha > 0.$

Result: Markov chain $\{n_k\}$ stable if and only if $\sum_k R_{lk} \frac{\lambda_k}{\mu_k} < c_l \quad \forall l.$

Remark: congestion control ensures neither stability nor fairness.

- Both depend solely on users' "open loop" demands $\frac{\lambda_k}{\mu_k}.$
- Fairness choice per flow (e.g., value of α) has minimal impact.

  A heavy user will compensate a low TCP rate by increasing n_k, until ρ_k serves demand, if feasible. If not n_k's grow without bounds.

Closing the loop on n_k for user-level fairness

Assume that for fixed n_k, the flow rate x_k is determined by TCP:

$x_k = f_{TCPk}(q_k)$ where q_k is the congestion price seen by the source,

and $f_{TCPk} = (U'_{TCPk})^{-1}$, TCP demand curve. The user rate is $\rho_k = n_k x_k$.

Objective: control n_k so that the system converges to an equilibrium where

$\rho_k = n_k x_k$ solves $\max_\rho \sum_k U_k(\rho_k)$, s.t. $R\rho \le c$, with utilities defined by users.

Control law for continuous n_k:

$$\dot{n}_k = \beta\left(U'^{-1}_k(q_k) - \rho_k\right).$$

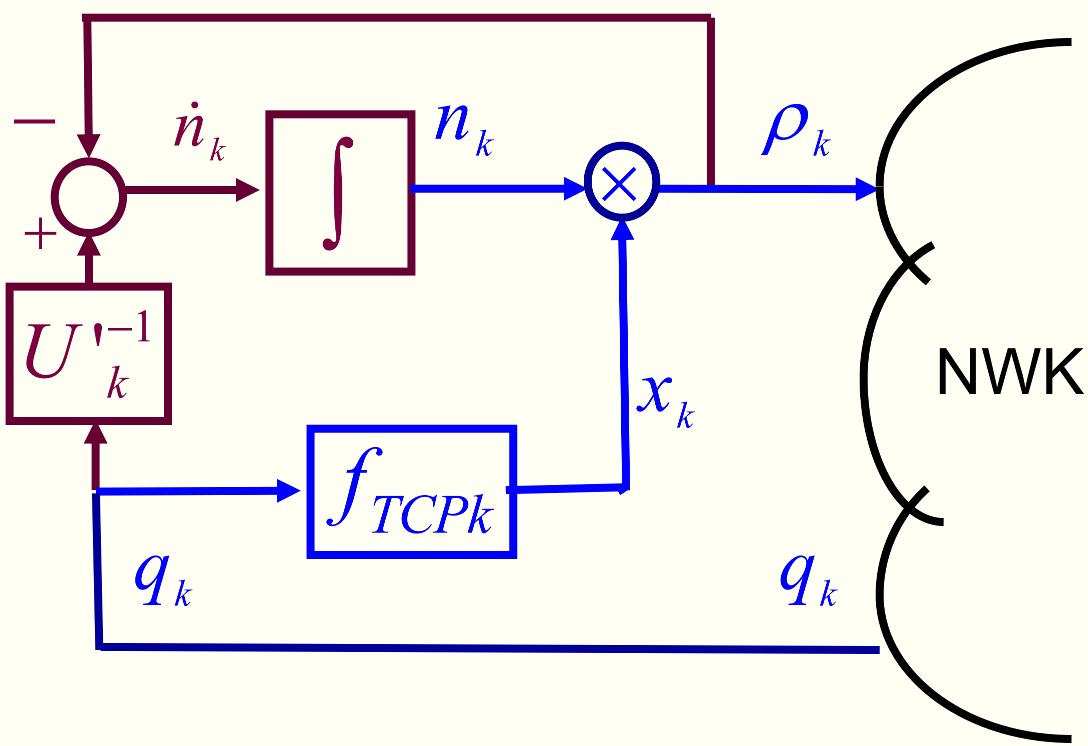

Analysis using dual TCP congestion control,

$$\dot{n}_k = \beta\left(U'^{-1}_k(q_k) - \rho_k\right);$$

$$\rho_k = n_k x_k;$$

$$x_k = f_{TCPk}(q_k).$$

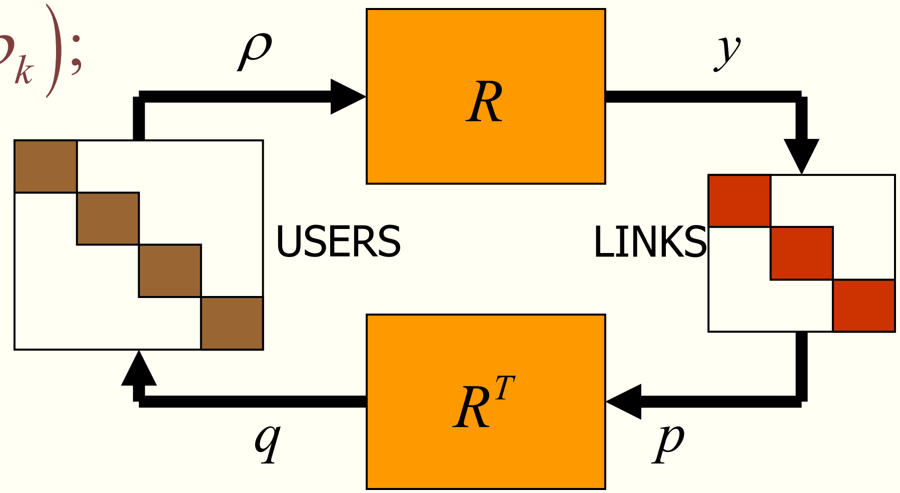

$$\dot{p}_l = \gamma_l \left[y_l - c_l\right]^+_{p_l}$$

Theorem 1 (arbitrary network).

The equilibrium satisfies $\max_\rho \sum_k U_k(\rho_k)$, subject to $R\rho \le c$, and is locally asymptotically stable. Proof: passivity argument (as in Wen-Arcak '03).

Theorem 2 (single bottleneck).

Assume time-scale separation: for fixed $n = \{n_k\}$, let $\hat{q}_k(n)$, $\hat{x}_k(n)$ be the equilibrium values from dual congestion control, and $\hat{\rho}_k(n) = n_k \hat{x}_k(n)$. Then the "slow" dynamics $\dot{n}_k = \beta\left(U'^{-1}_k(\hat{q}_k(n)) - \hat{\rho}_k(n)\right)$ are globally convergent to a point n^* where the corresponding $\hat{\rho}_k(n^*)$ are at the optimum welfare point.

From fluid control to admission control.

In practice, n_k is discrete (number of TCP connections). Furthermore:

- Real-time control at sources' (application layer) is impractical, incentives?
- Killing an ongoing TCP connection to reduce n_k is undesirable.

More practical alternative:

- Control increase of n_k (admit new connections), rely on natural termination.
- Admission control carried out by edge ro
- User utility $U_k(\rho_k)$ describes the SLA:

admit new connection $\Leftrightarrow U'^{-1}_k(q_k) > \rho_k$

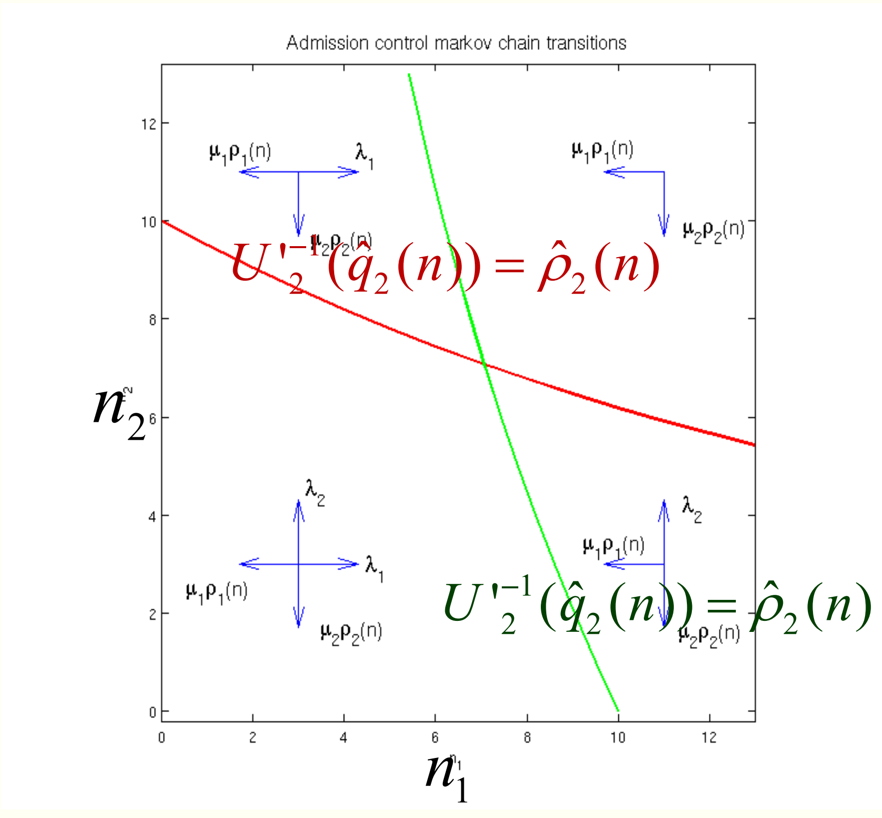

Markov chain model.

Poisson(λ_k) arrivals,

exp(μ_k) workloads.

Active sessions served with rate x_k obtained from TCP.

Theory and simulations

Fluid limit model,

$$\dot{n}_k = \lambda_k \mathbf{1}_{\{U'^{-1}_k(\hat{q}_k(n)) > \hat{\rho}_k(n)\}} - \mu_k \hat{\rho}_k(n)$$

For $\dfrac{\lambda_k}{\mu_k} > \rho^*_k$ (optimal fairshare)

simulations show convergence to optimal point,

we can prove it in simple networks.

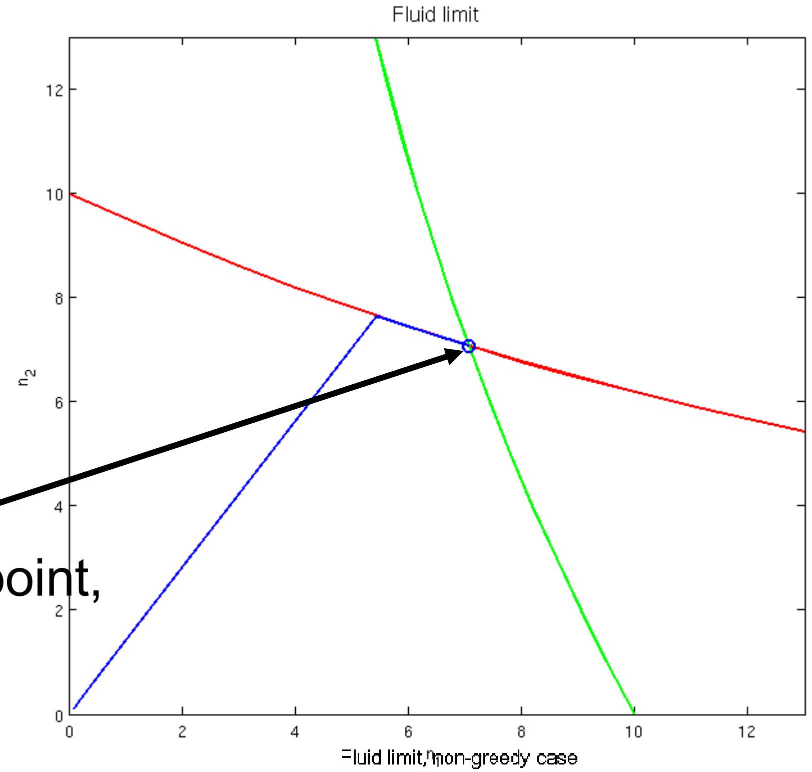

If a certain user has $\dfrac{\lambda_k}{\mu_k} < \rho^*_k$, convergence

to solution of $\max_\rho \sum_k \bar{U}_k(\rho_k)$, s.t. $R\rho \le c$,

where $\bar{U}_k(\bullet)$ corresponds to a demand curve

saturated at rate $\rho_k = \dfrac{\lambda_k}{\mu_k}$.

Non-greedy user is protected.

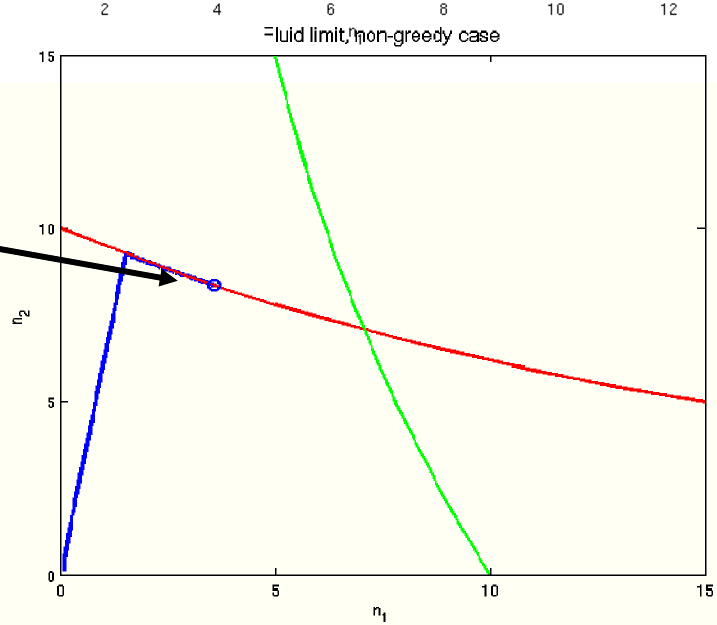

Multipath extensions are available.

3. Distributed Auctions for Resource Allocation in Overlay Networks

[Belzarena-Ferragut-P', Net-Coop '08]

- Scenario: an operator sells premium services (e.g. video-on-demand) over a network.
- QoS is guaranteed by reserving bandwidth in fixed amounts, and controlling access.
- One related proposal Service Overlay Newtork (Zuan et al.): here the service is offered by leasing bandwidth from several ISPs, installing distributed content servers.
- Objective: selling the service to maximize revenue.

Distributed Auction

- End users submit bids for certain end-to-end service to a broker within their domain.
- Periodically, an auction is held to decide which users are assigned bandwidth.
- Broker decisions are “spatially” coupled.

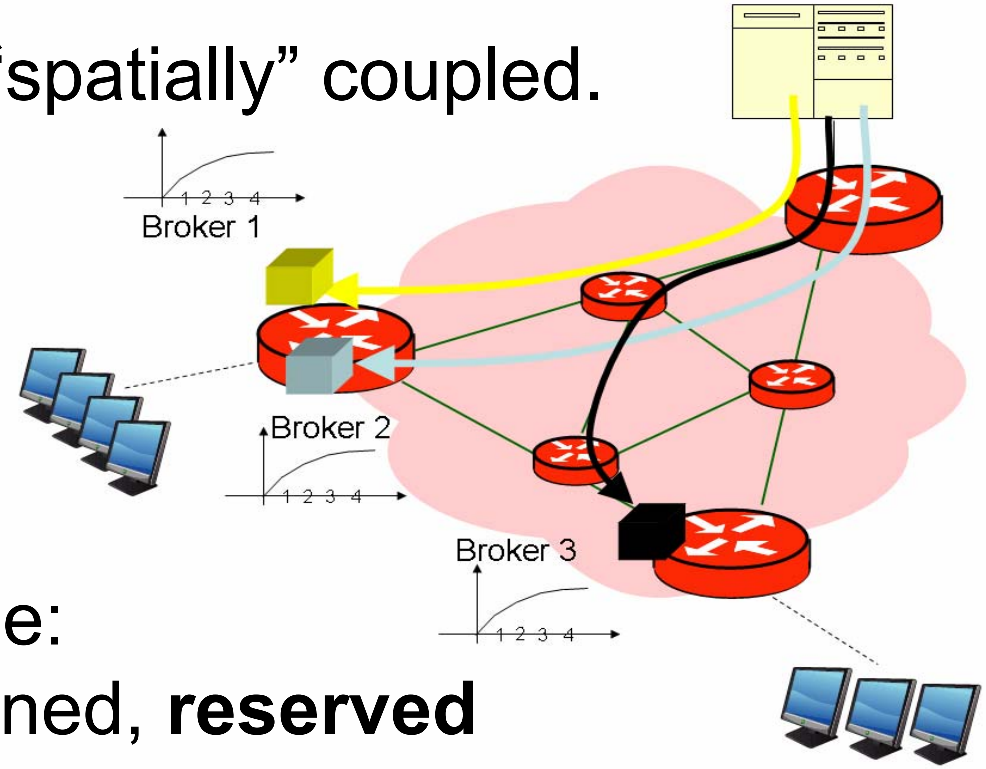

- Also coupled over time:
  once a circuit is assigned, **reserved** for entire connection (e.g. movie).

Auctions for a single link

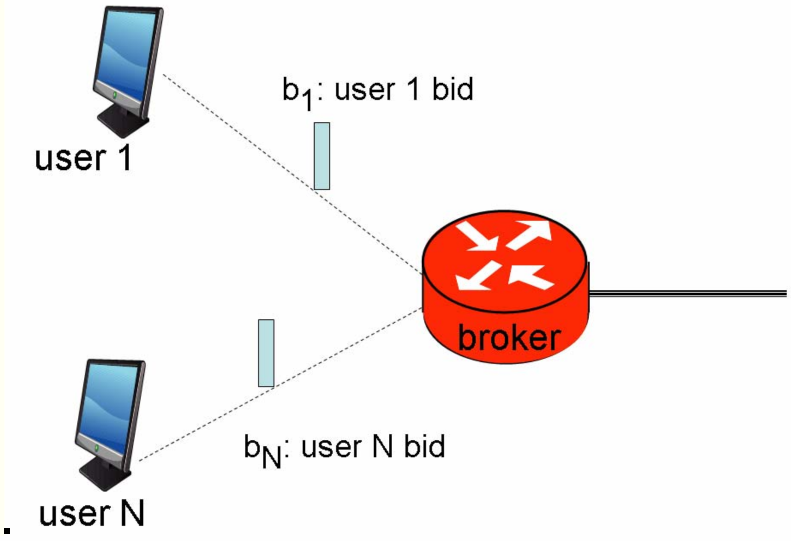

Assume:

- At $k-$th auction period: N bids $b_k^{(1)} \geq b_k^{(2)} \geq b_k^{(3)} \geq \cdots \geq b_k^{(N)}$ for a unit of bandwidth. Link capacity C.
- Bids have known probability distribution.

  - Revenue from accepting a_k bids

(first-price auction) $U_b(a_k) = \sum_{i=1}^{a_k} b_k^{(i)}$.

Interpolate to piecewise linear, concave function.

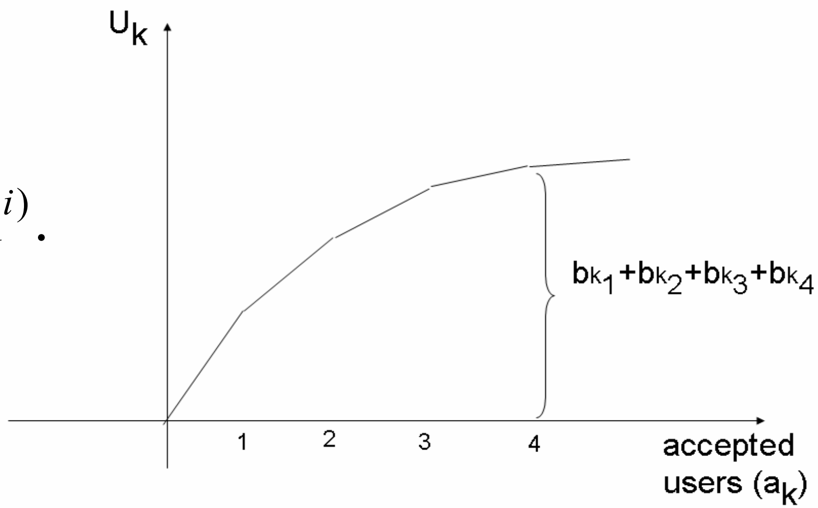

Operator revenue in auction k

- Expected revenue from a_k bids: $\bar{U}(a_k) = E\left[U_b(a_k)\right] = \sum_{i=1}^{a_k} E[b_k^{(i)}]$

Successive auctions

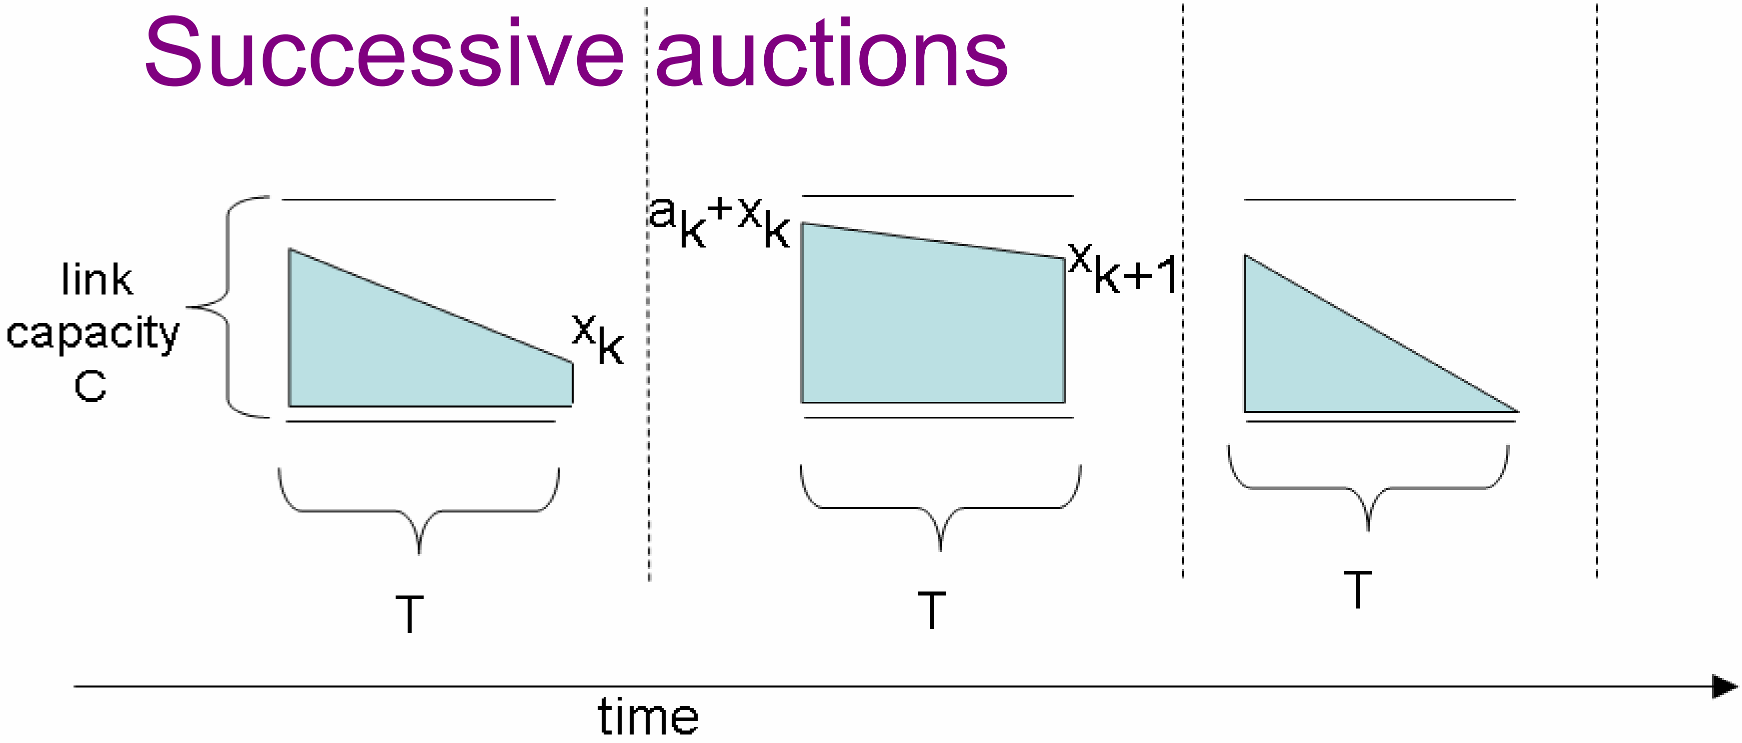

p: probability that a connection remains active at the end of the period of time T

x_k: number of active connections at kT-

a_k: number of accepted connections in auction k

x_k+a_k: active connections in kT+

Myopic policy: $a_k = C - x_k$, sells all available capacity.

May miss better bids in the future.

Optimal revenue problem

Maximize $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{a_k} E\left[U_b(a_k)\right]$

- Expectation is w.r. t the bids, and the departure process.
- This is a Markov Decision Process (MDP).
- Solution is a policy $a_k = a(s_k)$, where the state $s_k = (x_k, b_k)$
- $a(s)$ can be found numerically, large compuational cost.

Receding horizon approximation:

$$a_0 = \arg\max_{a \le C - x_k} [U_b(a) + E_{x_1} \bar{U}(C - x_1)]$$

- Optimize current revenue+ expected revenue of next auction, assuming all remaining capacity will be sold off at that time.
- Take auction a_0, and repeat recursively.

Receding horizon policy: $a_0 = \arg\max_{a \le C - x_k} [U_b(a) + E_{x_1} \bar{U}(C - x_1)]$

- For the second term, $x_1 \sim Bin(x_0 + a, p)$.
- Some calculations reduce it to $-W(x_0 + a)$, where $W()$ piecewise linear, increasing and convex (cost of missed future opportunities).
- Maximum: intersection of decreasing marginal utilities (bids) with increasing marginal costs (acceptance thresholds, w_k).

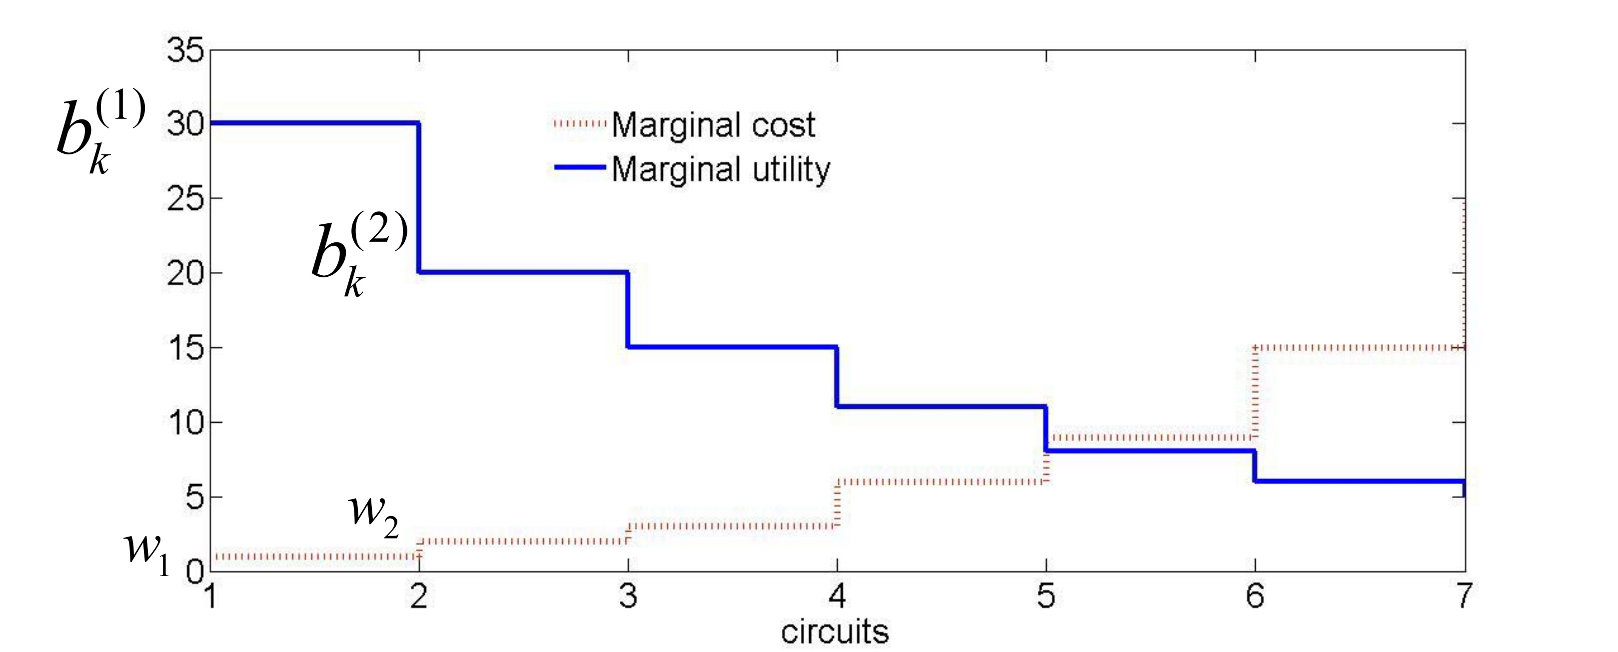

Comparisons (C=1)

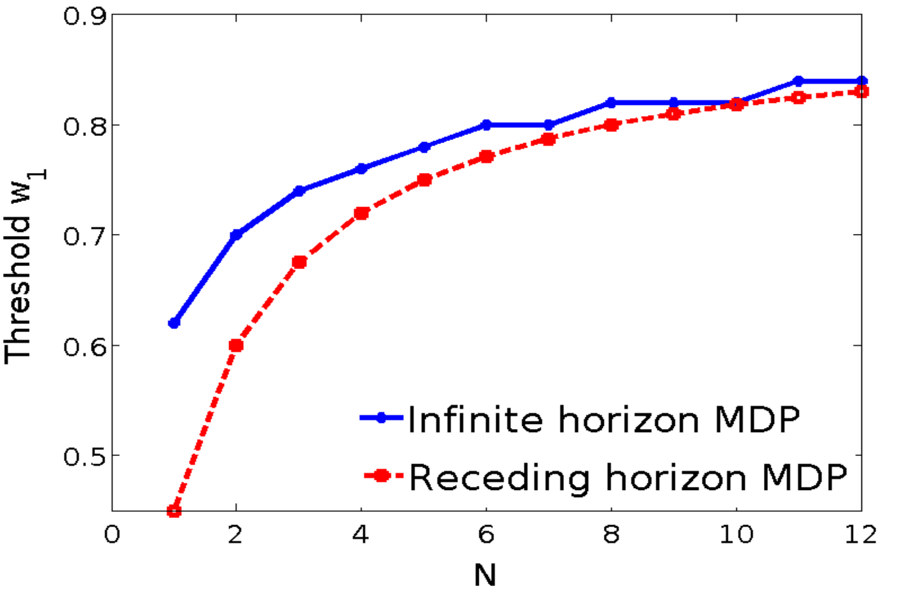

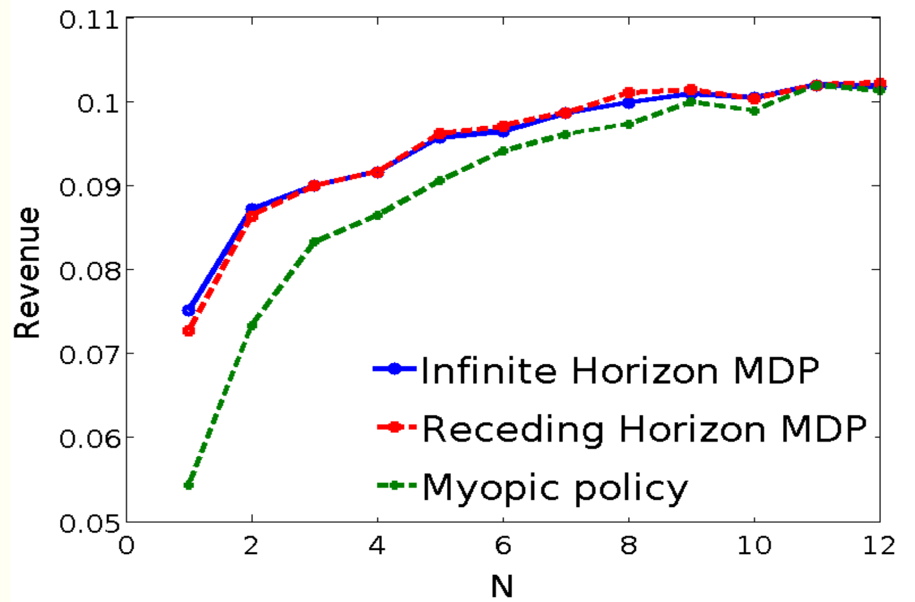

A fluid aproximation: replace $E_{x_1}\bar{U}(C-x_1)$ by $\bar{U}(C-E[x_1])$

The problem reduces to the convex program

Maximize $\underbrace{U_b(a)}_{\text{CURRENT REVENUE}} + \underbrace{\bar{U}(z)}_{\text{EXPECTED REVENUE IN NEXT STEP}}$

subject to $\underbrace{x_0 + a \le C,}_{\text{CURRENT CAPACITY CONSTRAINT}} \quad \underbrace{p(x_0 + a) + z \le C.}_{\text{EXPECTED FUTURE CAPACITY CONSTRAINT}}$

The network case

Assign a broker to each r, route/service in the network.

a^r : current allocation, broker r

z^r: expected next-step allocation.

R: routing matrix, as before.

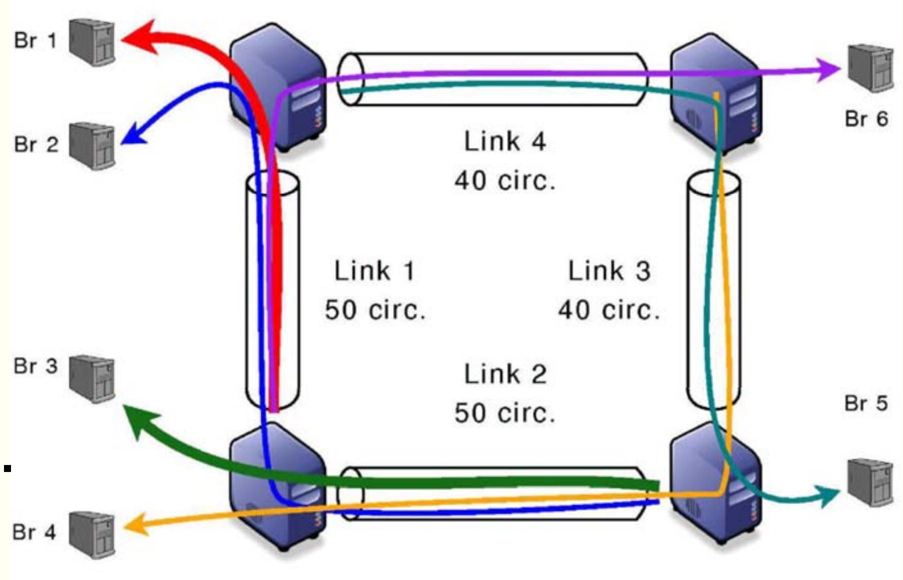

Network allocation problem: Maximize $\sum_r U_{b^r}(a^r) + \bar{U}(z^r)$

subject to $R(x_0 + a) \le C, \quad RP(x_0 + a) + z \le C$

- Similar to network utility maximization in congestion control
- Additional (one-step ahead) variables and constraints. Requires additional price variables.
- Distributed solution via message passing in the control plane (modification to RSVP).

Conclusions

We studied three problems in cross layer control and optimization:

1. Congestion control with multipath routing.
2. Controlling fairness through number of TCP connections.
3. Auctions for resource allocation in overlay networks.

Common features:

- Economic (utility based) models.
- Dynamics play a non-trivial role.
- Distributed solutions.

Progression from “virtual” to real economics (utilities as protocol representations, versus real monetary utilities).

Grand challenge for the future: an integrated view of network control and economics.