On congestion control, multipath routing, and admission control

Fernando Paganini
Universidad ORT Uruguay

- Part I (with Enrique Mallada). Combined congestion control and node-based multipath routing: new results on stability since CISS'06.

- Part II (with Andrés Ferragut, in CISS'08 paper). Achieving network stability and user fairness through admission control of TCP connections.

Congestion control with multipath routing

[Paganini, CISS'06, ECC07
Mallada-P' Netcoop '07]

Links

- Rate y_l
- Capacity c_l, or cost $\phi_l(\, y_l\,)$.

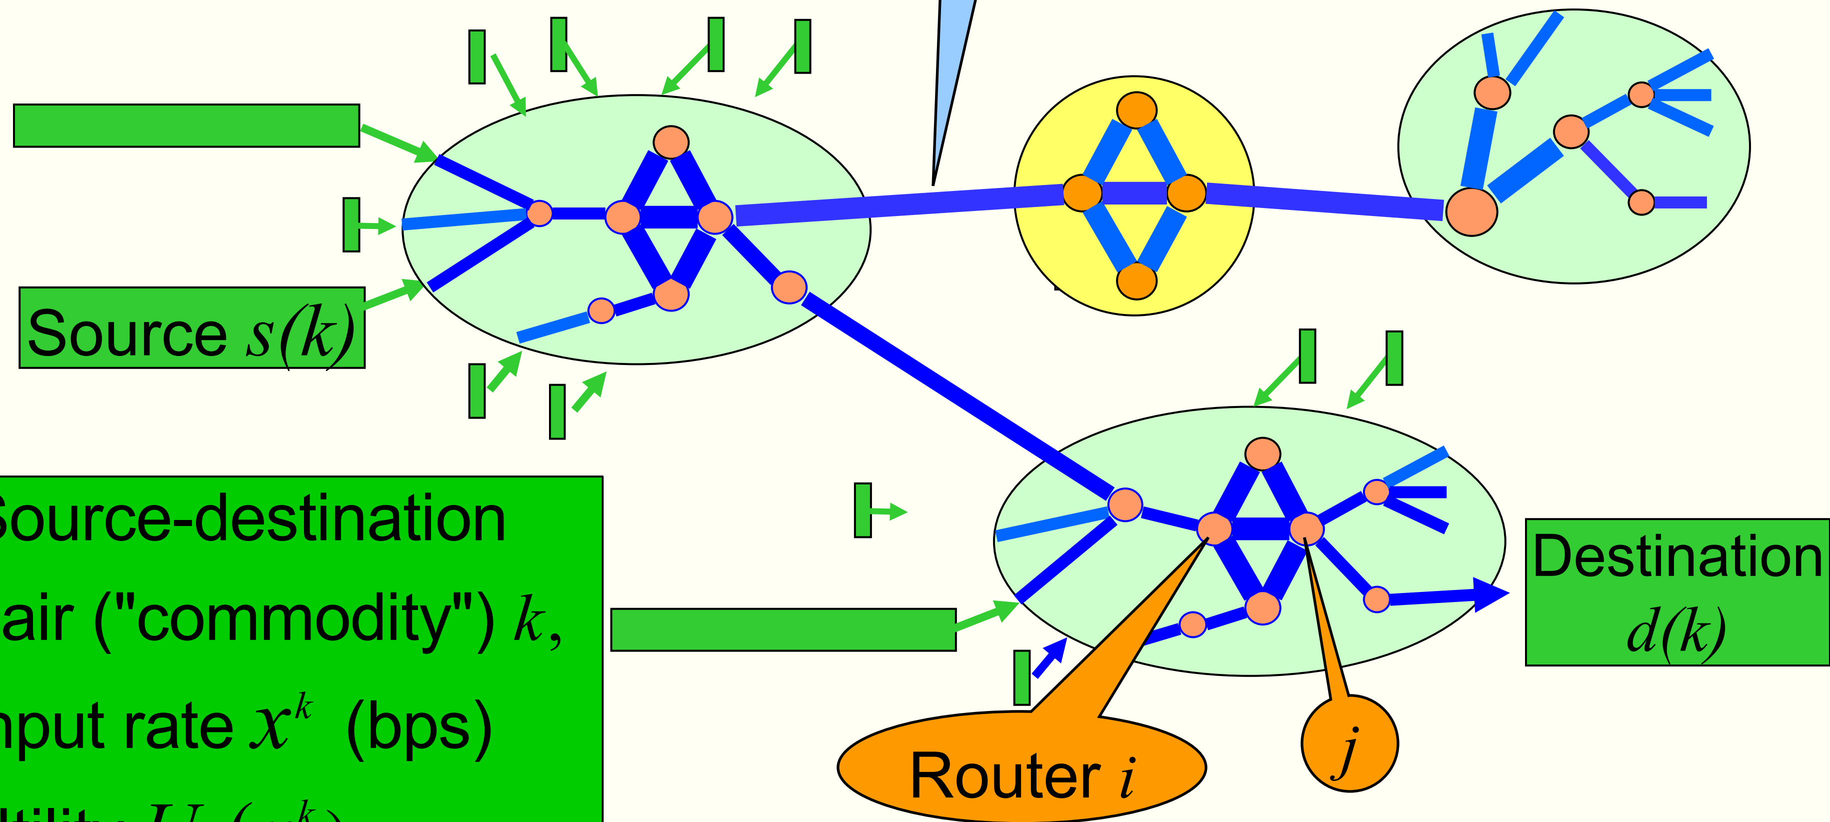

Source-destination pair ("commodity") k,
input rate x^k (bps)
Utility $U_k(x^k)$.

$\alpha_{i,j}^d$: fraction of traffic destined to d, sent to node j. $\alpha_i^d = \{\alpha_{i,j}^d\} \in \Delta_i$ (simplex)

Congestion prices

- Define: link congestion price p_l. (e.g., $p_l = \phi'(y_l)$ or $\frac{dp_l}{dt} = \gamma_l [y_l - c_l]$.)
- q_i^d : average price from node i to destination d

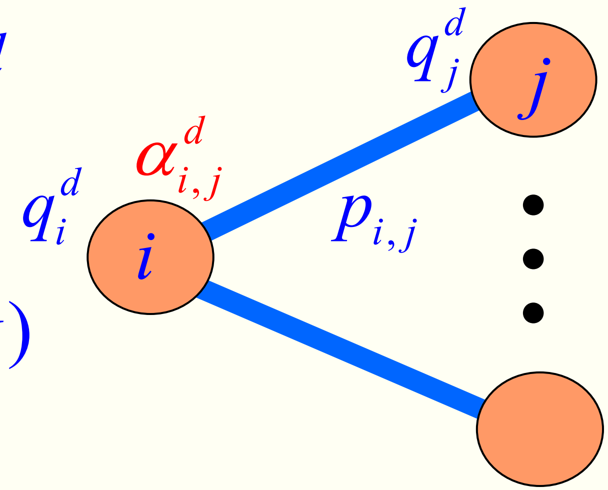

$$q_d^d = 0, \qquad q_i^d = \sum_{j:(i,j)\in L} \alpha_{i,j}^d \left(p_{i,j} + q_j^d \right), \; i \neq d$$

$q^k := q_{s(k)}^{d(k)}$, average price seen by source $s(k)$

Multipath routing control.

Routers control $\alpha_i^d := \{\alpha_{i,j}^d\}_{(i,j)\in L}$ based on seen prices $\pi_i^d := \{p_{i,j} + q_j^d\}_{(i,j)\in L}$

- First choice (essentially from Gallager '77):

follow negative price gradient. $\dot{\alpha}_i^d = \beta_i E_{\alpha_i^d}\left[-\pi_i^d\right]$.

The projection $E_{\alpha_i^d}$ keeps $\alpha_i^d \in \Delta_i$.

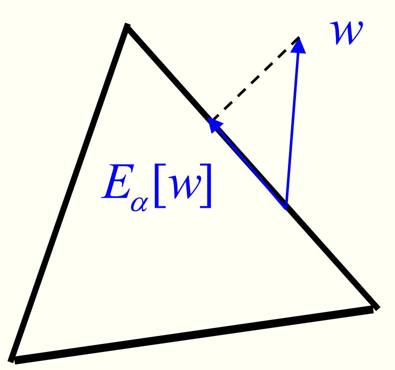

Primal congestion control under gradient control of routing fractions

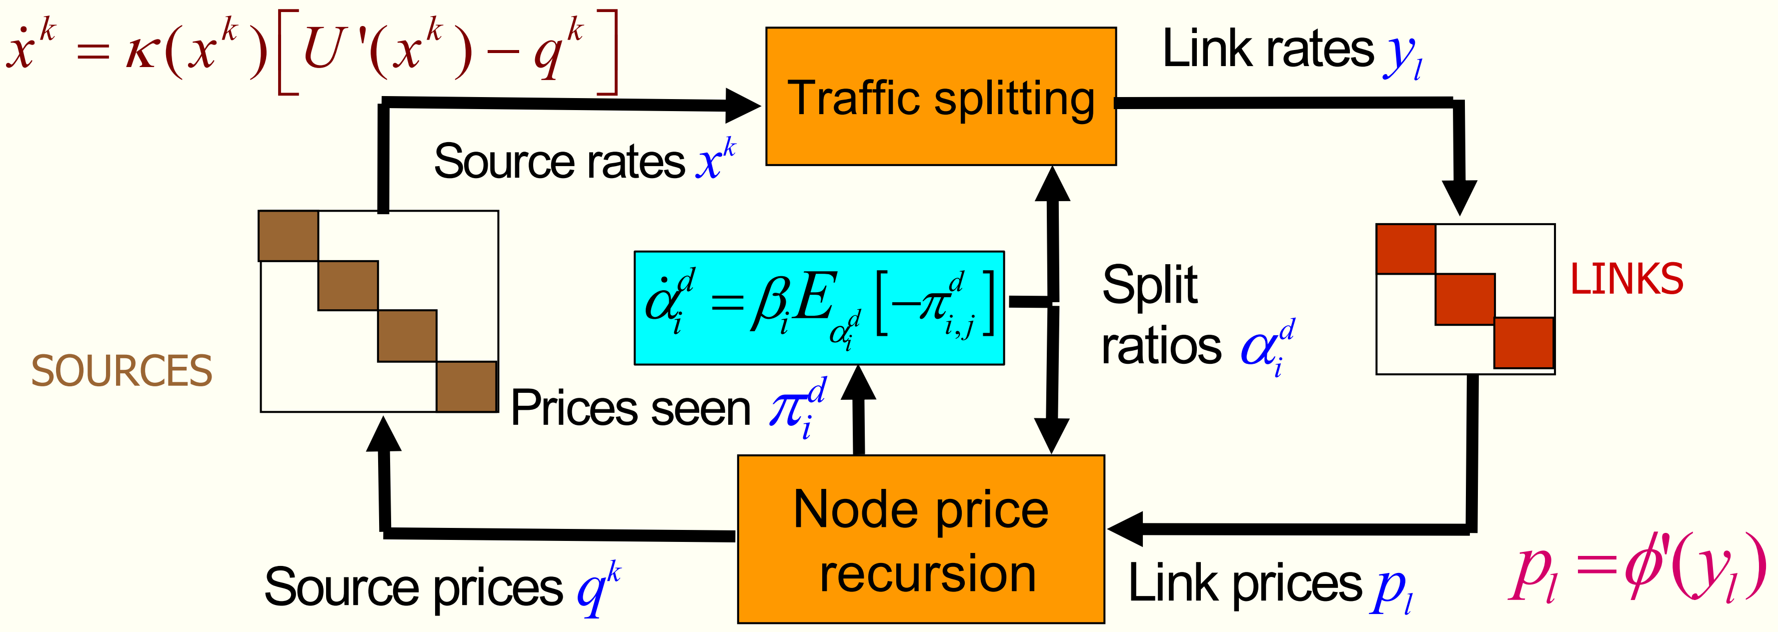

Theorem: with these dynamics the system converges globally to the optimum of the SURPLUS problem,

$$\max S := \sum U^k(x^k) - \sum \phi_l(y_l)$$

[See P' CISS '06], extends Gallager '77.

Also in [Xi & Yeh CISS'06], combined with wireless power control.

Dual congestion control under gradient control of routing fractions

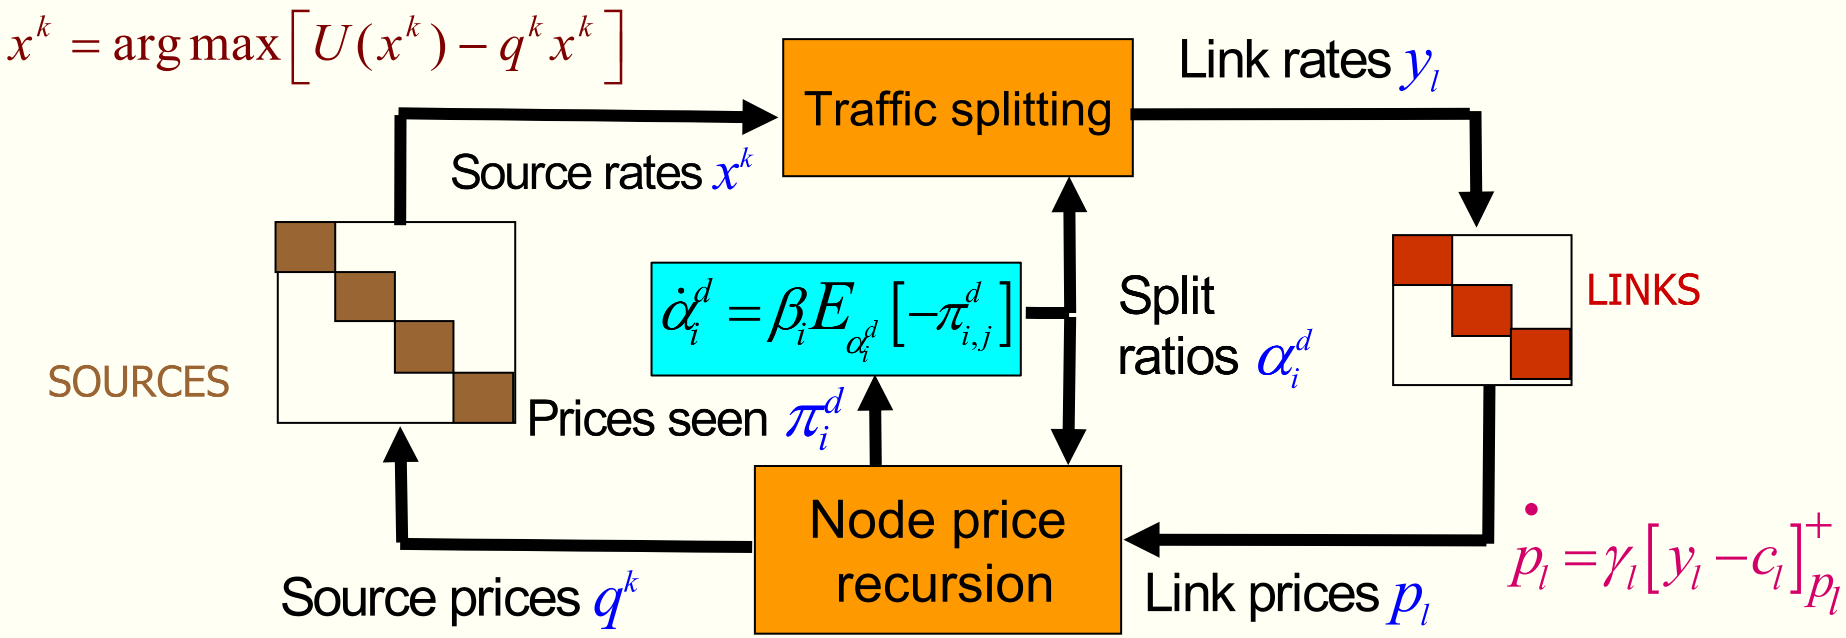

Equilibrium: solution to optimal WELFARE problem,

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

Convergence to equilibrium?

No! Simple examples exhibit harmonic oscillations.

Which model is correct for queuing delay:

Static $p_l = \varphi(y_l)$ or integrator $\dot{p}_l = \frac{1}{c_l}[y_l - c_l]^+_{p_l}$?

Example of instability,

packet simulation in ns2

Single source, two bottlenecks.

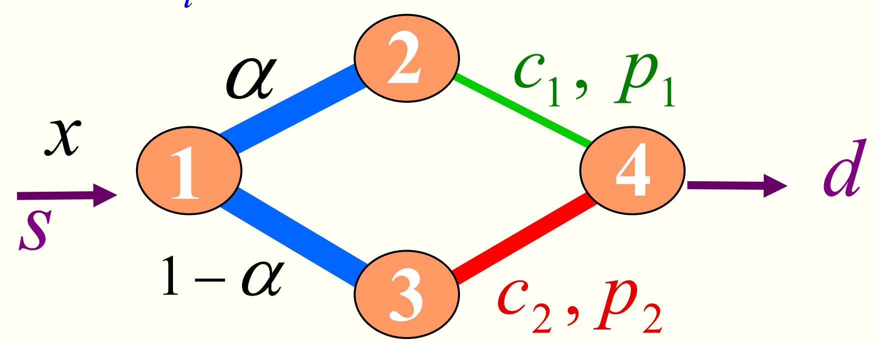

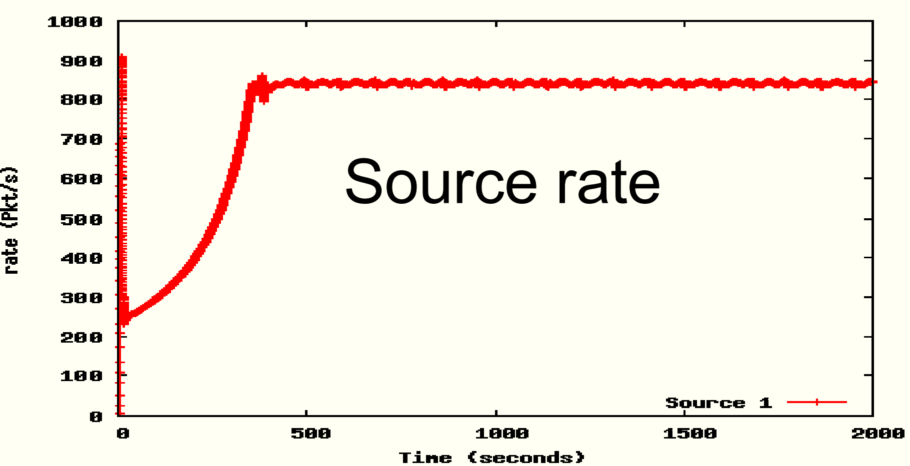

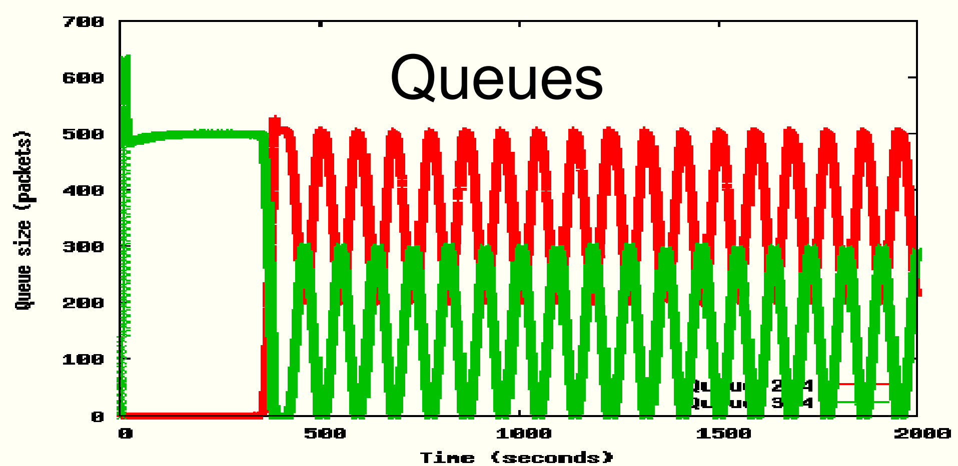

- Dual model of queue is more appropriate.
- Indeed, it predicts correctly oscillation period.

Solving the problem

Adapt α_i^d based on anticipated (rather than current) price $\boxed{\pi_i^{d*} = \pi_i^d + \nu_i \dot{\pi}_i^d}$

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

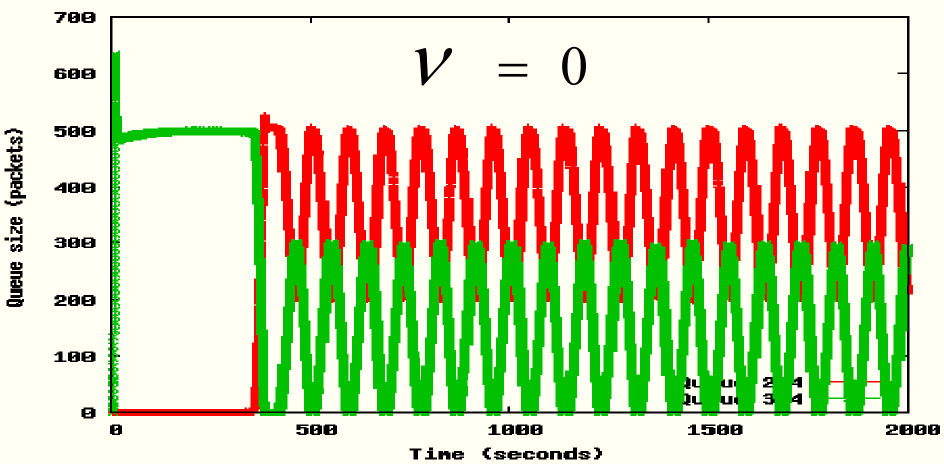

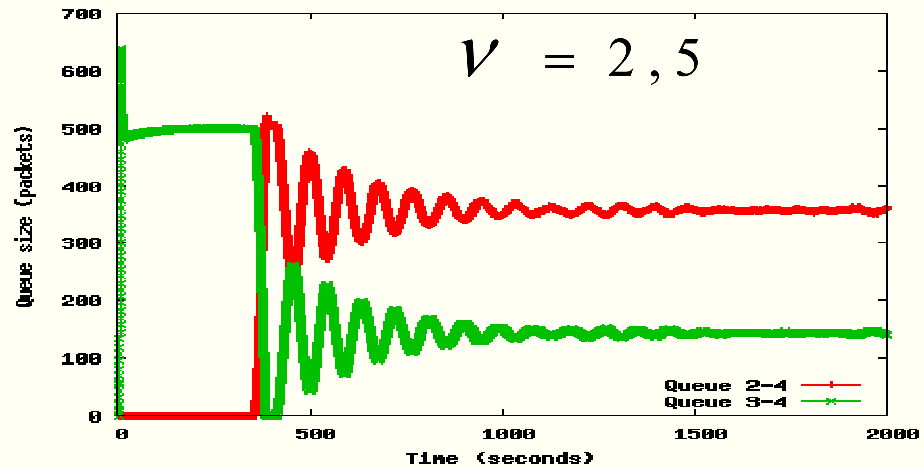

Theorems [P'-Mallada, submitted to ToN, CDC'08]

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

- locally asymptotically stable in an arbitrary network
- globally asymptotically stable in a network of parallel links.

Packet implementation: variants of TCP-FAST and RIP.

On congestion control, multipath routing, and admission control

Fernando Paganini
Universidad ORT Uruguay

- Part I (with Enrique Mallada). Combined congestion control and node-based multipath routing: new results on stability since CISS'06.

- Part II (with Andrés Ferragut, in CISS'08 paper). Achieving network stability and user fairness through admission control of TCP connections.

Stability and user-level fairness

User rate ρ_k

Utility $U_k(\rho_k)$

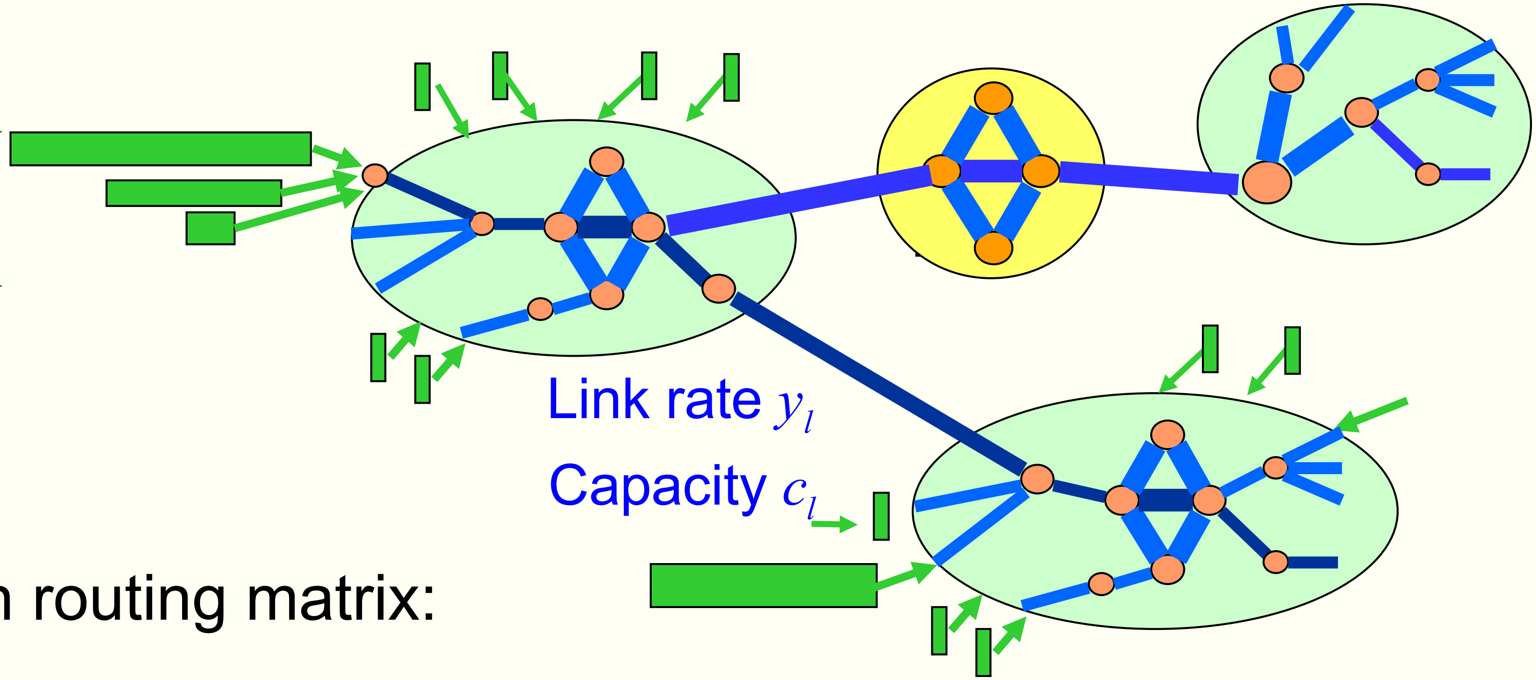

Single path routing matrix:

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

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

Contrast with flow-level fairness of TCP

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

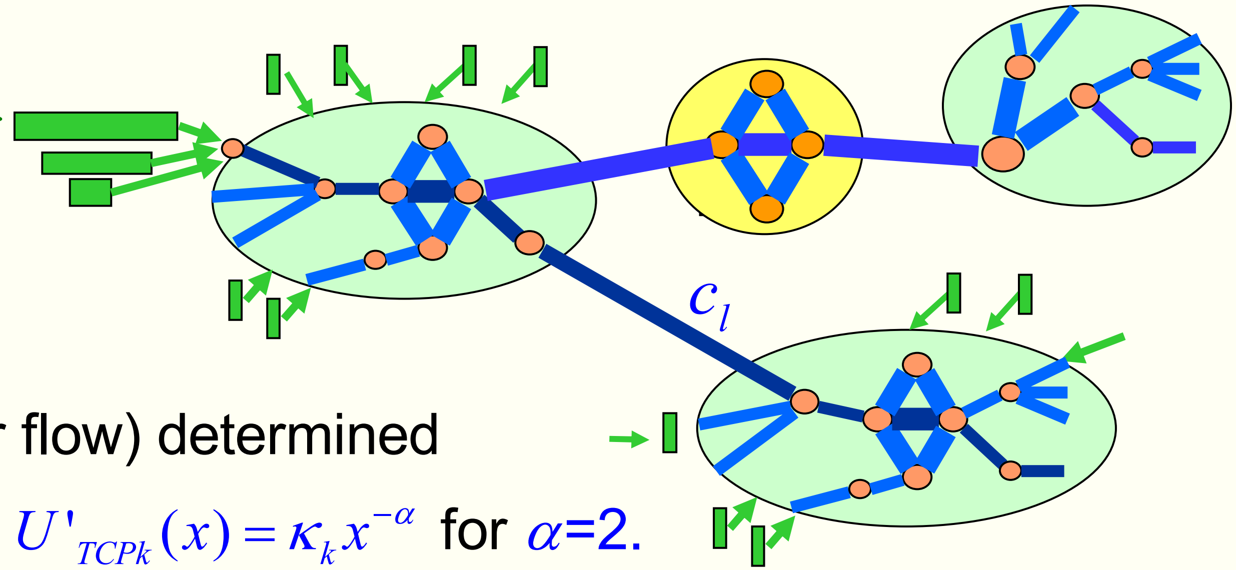

TCP utility U_{TCPk} (per flow) determined by the protocol, e.g., $U'_{TCPk}(x) = \kappa_k x^{-\alpha}$ for α=2.

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

- 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?
- If we could control n_k, can we induce with TCP the SYSTEM problem allocation?

(similar to "user problem" setting weight parameter in Kelly-Maulloo-Tan '98)

On stochastic stability of a network served by TCP
[deVeciana, Lee, Konstantopoulos '99, Bonald-Massoulié '01]

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

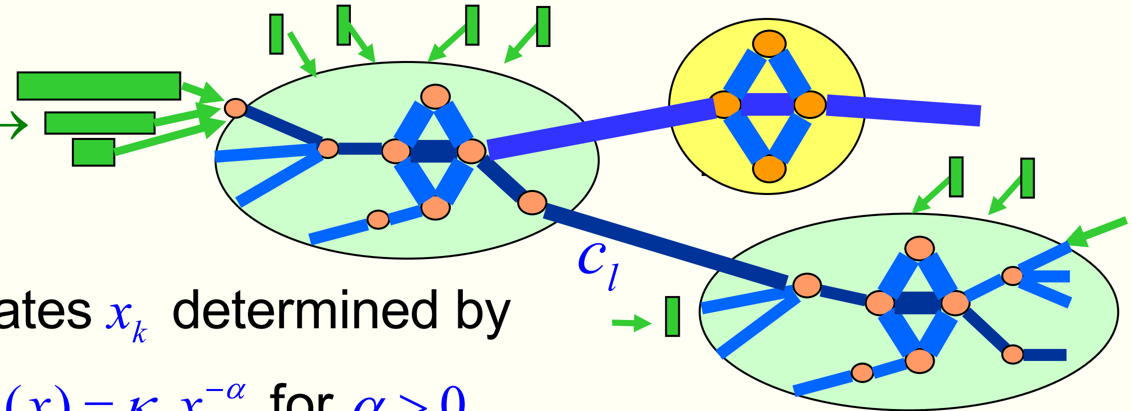

For each fixed $\{n_k\}$, service rates x_k determined by TCP congestion control $U'_{TCPk}(x) = \kappa_k x^{-\alpha}$ for $\alpha > 0$.

Result: n_k 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 stability, and resource allocation 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:

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

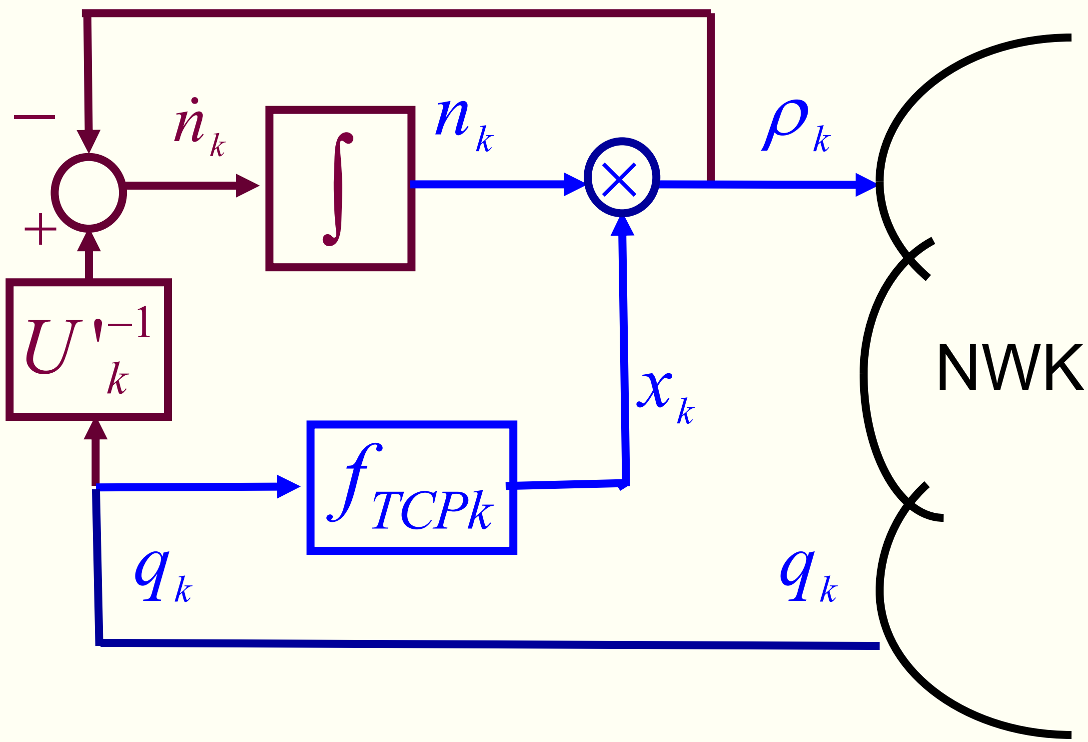

Other recent work on controlling no. of flows: Chen - Zakhor '06 (for TCP over wireless).

Analysis using dual TCP congestion control,

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

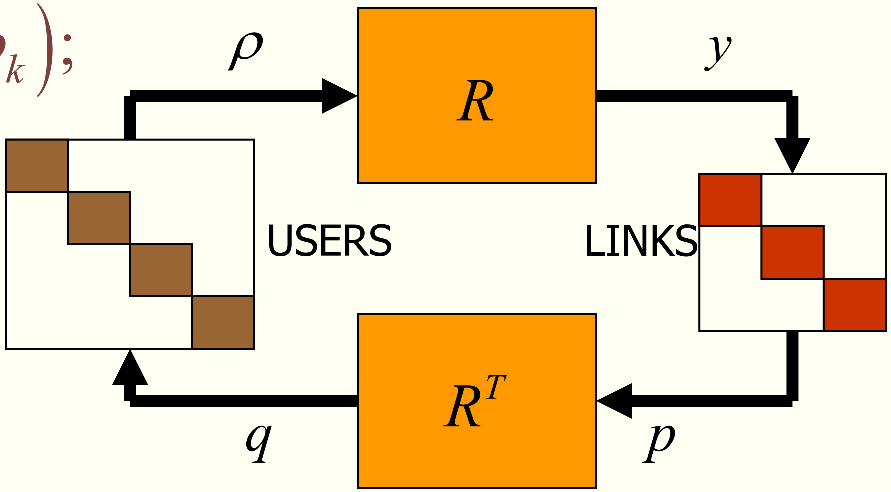

$$\dot{p}_l = \gamma_l\left[y_l - c_l\right]^+_{p_l}$$

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 router.
- User utility $U_k(\rho_k)$ describes the SLA: admit new connection $\Leftrightarrow U'^{-1}_k(q_k) > \rho_k$

Stochastic model. Poisson(λ_k) arrivals, exp(μ_k) workloads.

Active sessions served with rate x_k obtained from the network.

Continuous time Markov chain with state $n = \{n_k\}$.

Transition rates: $q_{n,n+e_k} = \lambda_k 1\{U'^{-1}_k(\hat{q}_k(n)) > \hat{\rho}_k(n)\}$; $q_{n,n-e_k} = \mu_k \hat{\rho}_k(n)$

Single bottleneck, two users, utility $U_k(\rho_k)=K\log(\rho_k)$.

State space and transition rates:

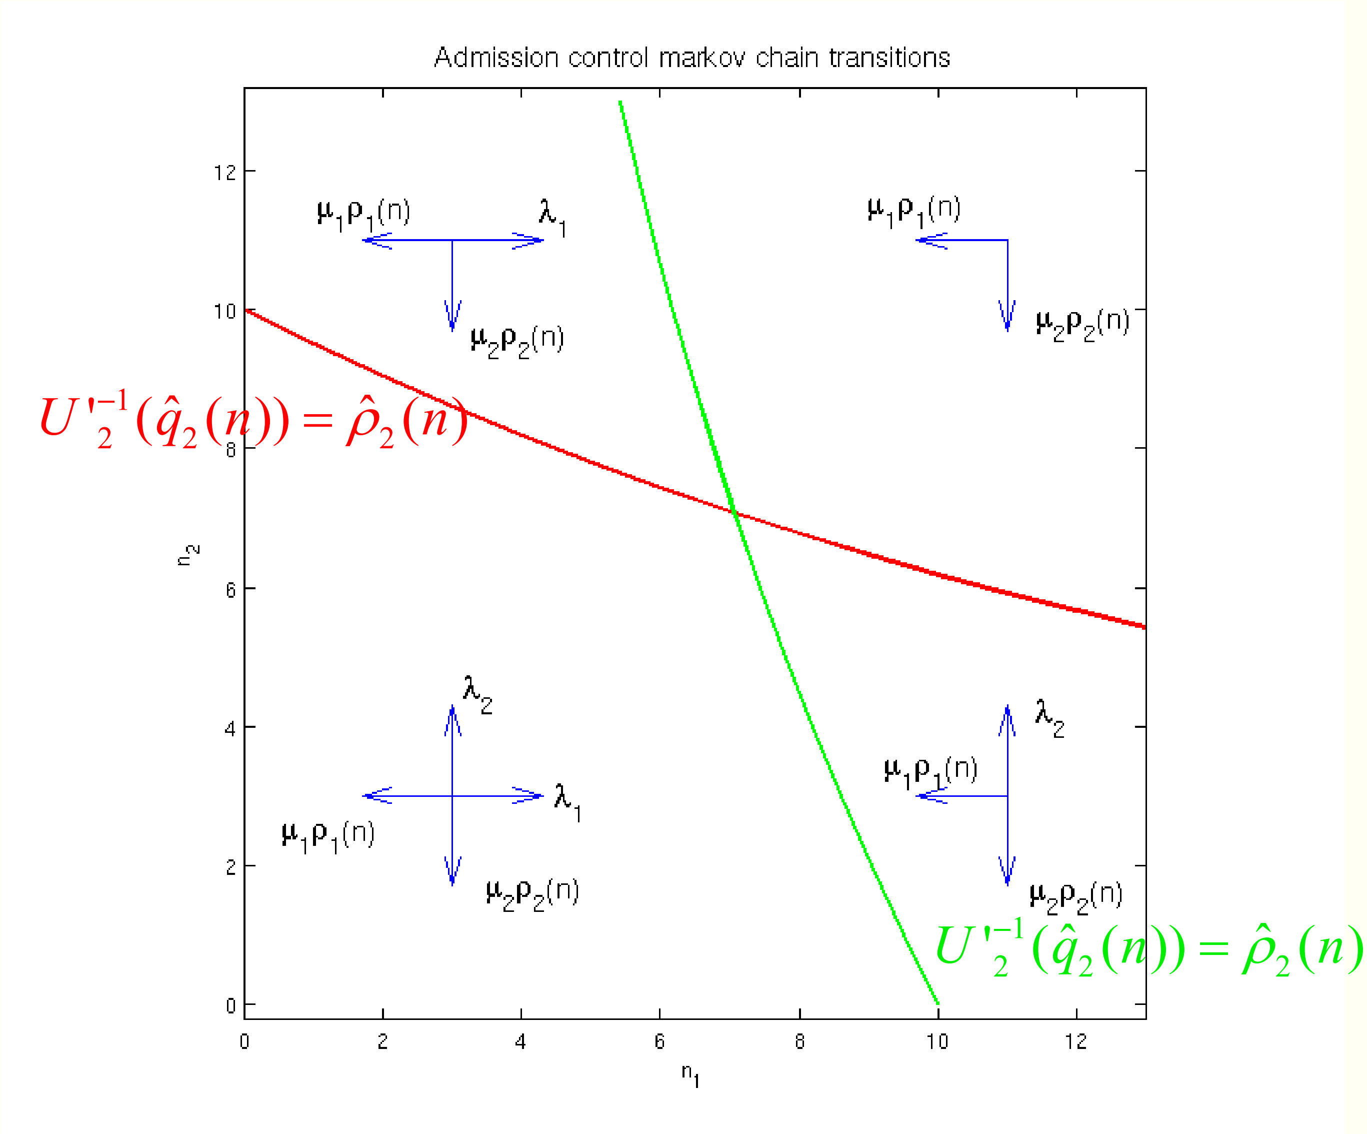

Irreducible set around $n=0$ is bounded

$\Rightarrow$ Stability,

Independently of λ_k, μ_k

Fairness? Simulations: JAVA-based tool with random arrivals and workload, simulated dual congestion control.

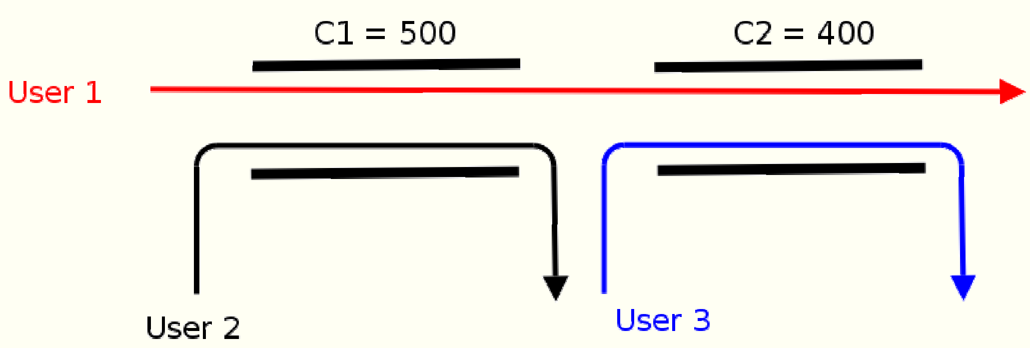

TCP utility $U'_{TCPk}(x) = \kappa_k x^{-2}$.

User utility $U'_k(x) = \kappa' x^{-5}$.

emulates max-min fairness.

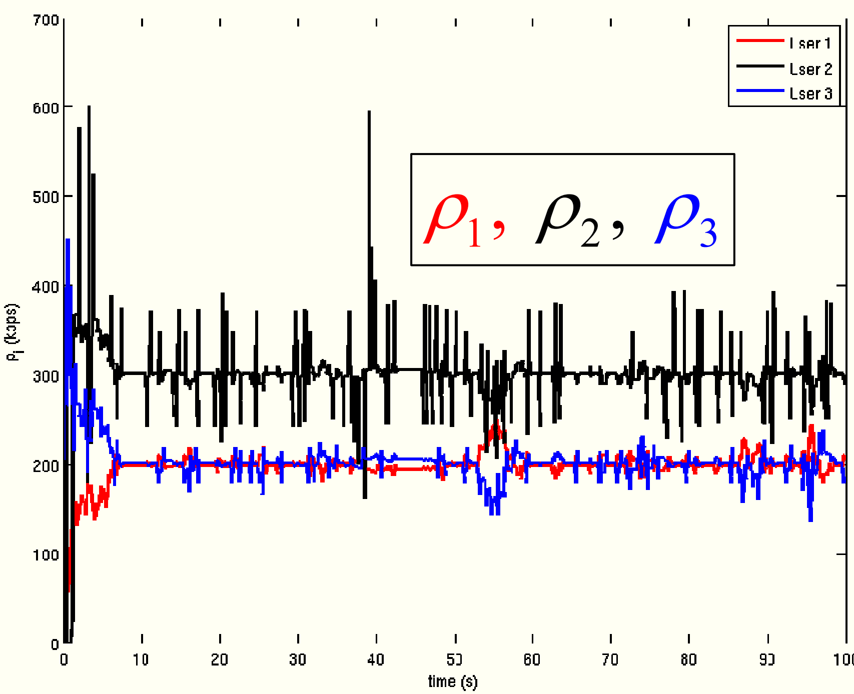

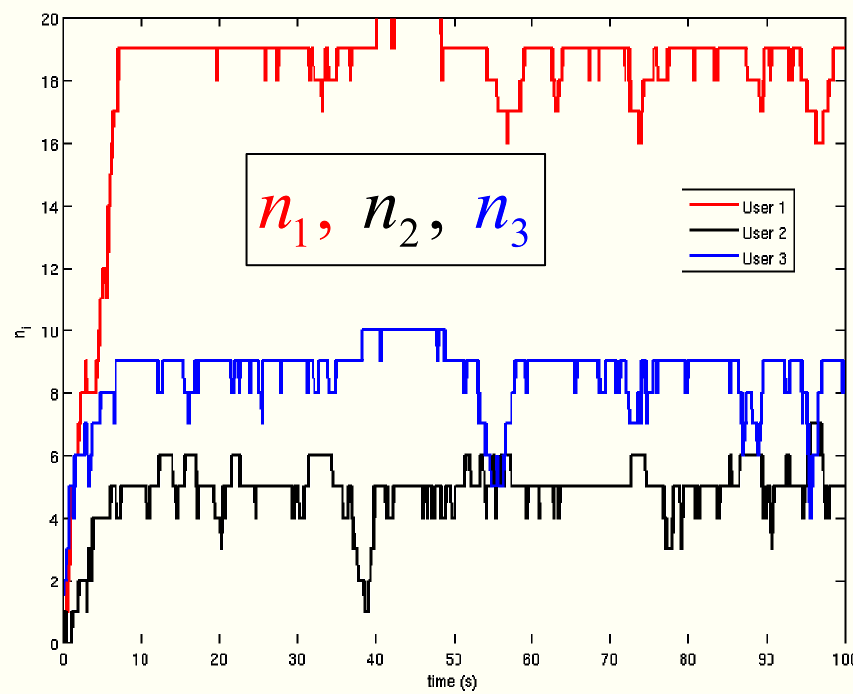

Fluid modeling of admission control.

Simulation results show admission control achieves optimal allocation, provided the offered loads $\frac{\lambda_k}{\mu_k}$ are larger than the equilibrium fairshare ρ^*_k.

We seek analytical proof, and also understanding of the non-greedy case.

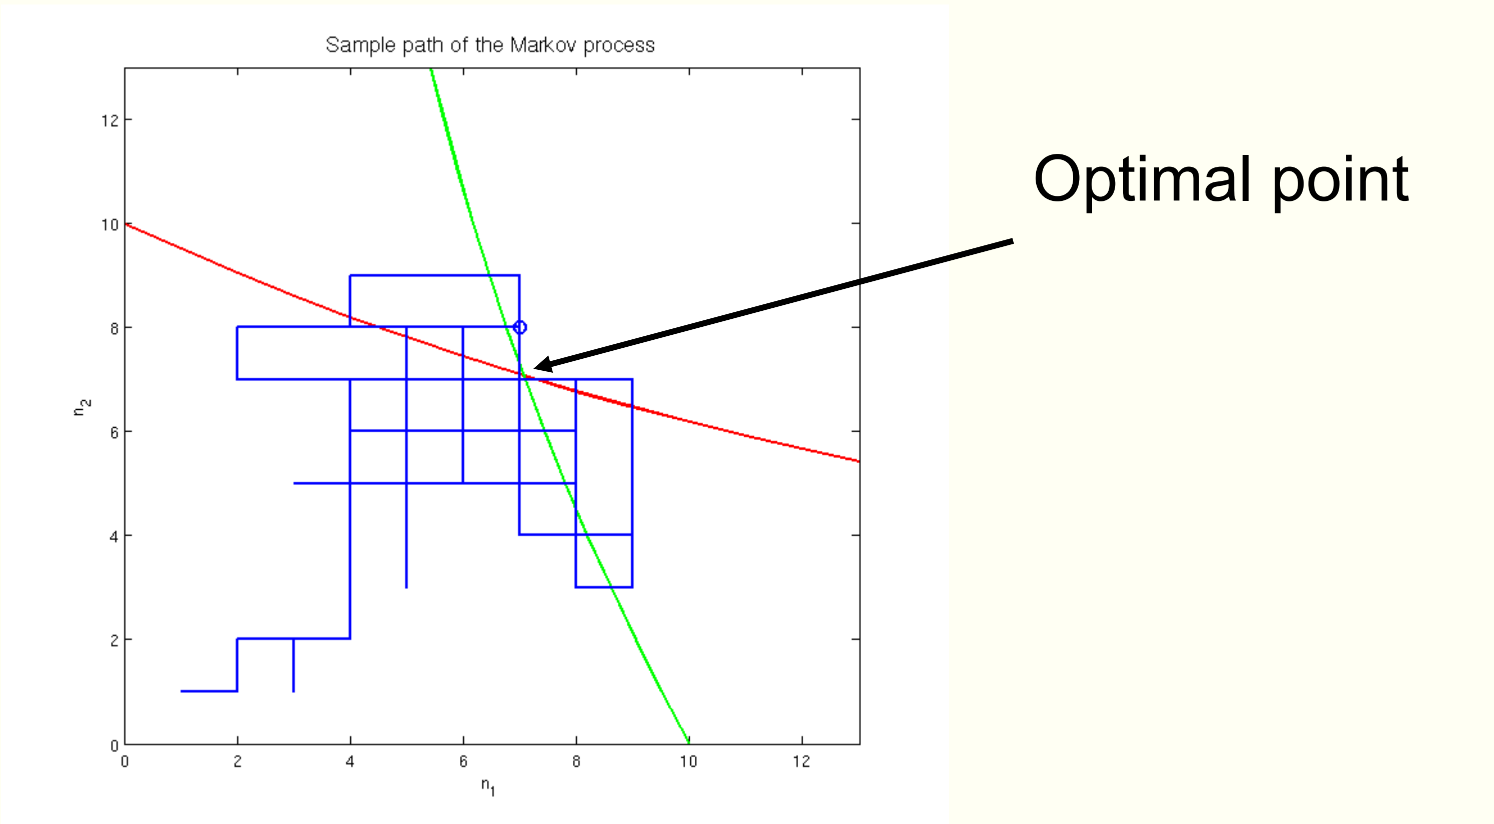

Fluid modeling of admission control.

Try a large network asymptotic, scaling capacity and user demand curve.

$c^{(L)} = c_0 L,$ $U_k^{(L)}(\rho) = U_k^{(1)}\left(\rho / L\rho_0\right).$ Rescaled simulation plot $\frac{n^{(L)}{}_k}{L}$

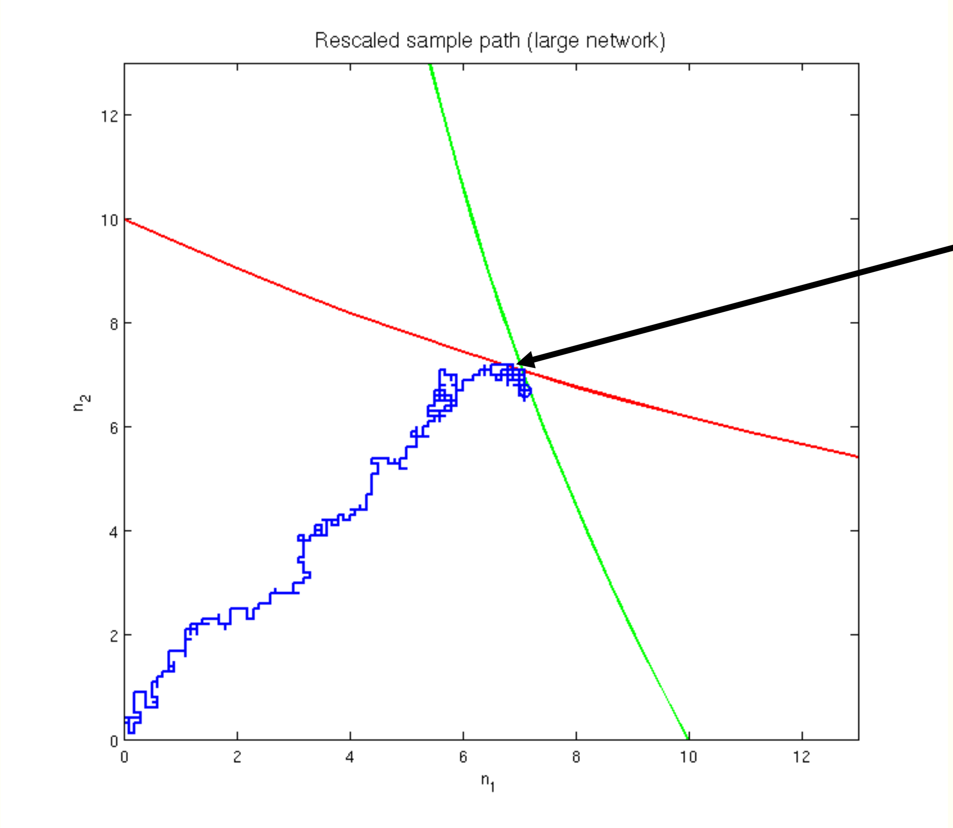

Optimal point

Fluid limit.

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

For $\frac{\lambda_k}{\mu_k} > \rho^*_k$ (optimal fairshare) fluid simulations converge to optimal point.

We can prove it in simple cases.

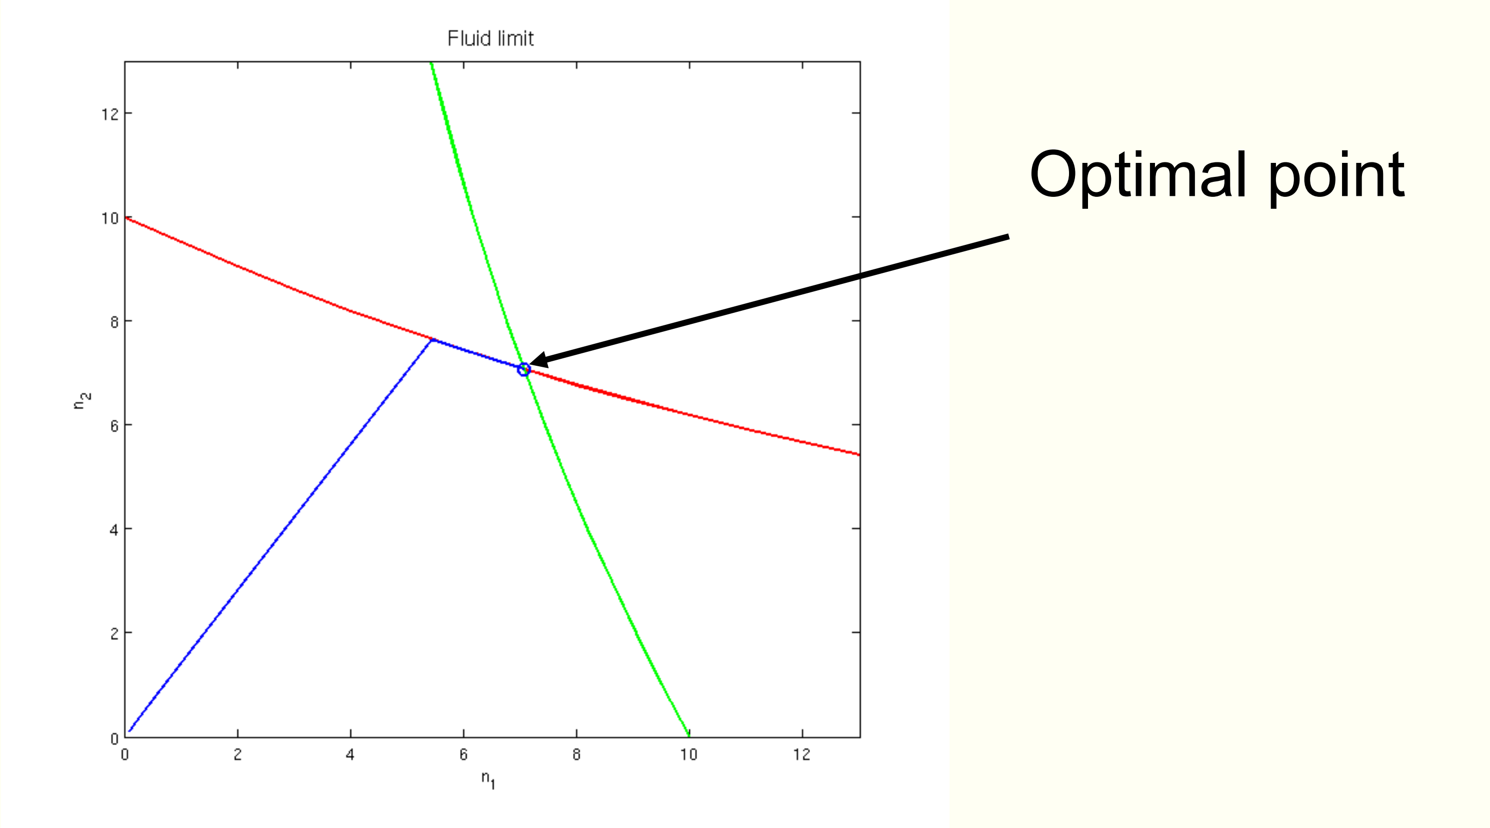

Fluid limit for the case of non-greedy users:

$\dot{n}_k = \lambda_k \mathbf{1}_{\{U'^{-1}_k(\hat{q}_k(n)) > \hat{\rho}_k(n)\}} - \mu_k \hat{\rho}_k(n);$ assume a certain user has $\frac{\lambda_k}{\mu_k} < \rho^*_k$

Example: $c = 1,$ user 1 with $\frac{\lambda_1}{\mu_1} = 0.3,$ greedy user 2.

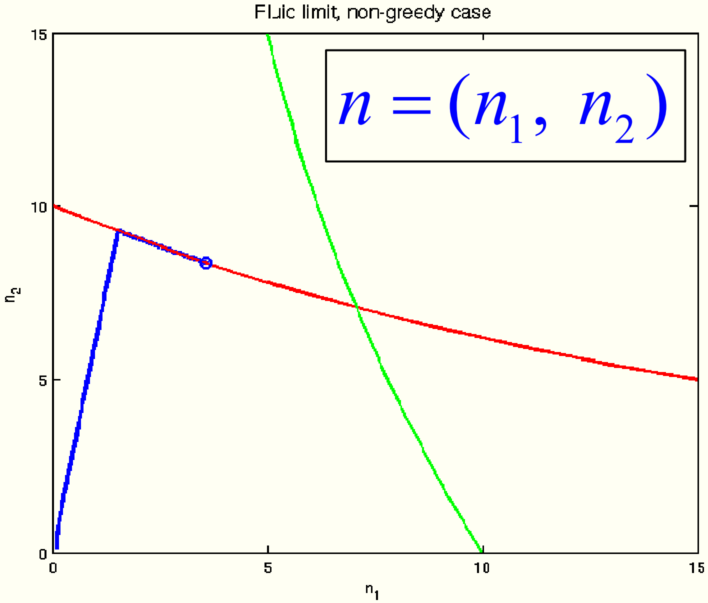

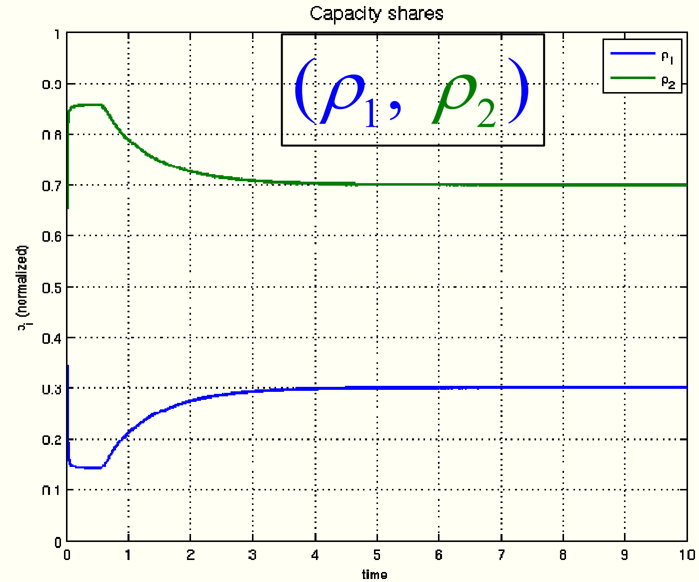

Conjecture (verified in simulations so far): converges to solution of

$\max_{\rho} \sum_k \bar{U}_k(\rho_k),$ s.t. $R\rho \le c,$ where $\bar{U}_k(\bullet)$ corresponds to a demand curve

saturated at rate $\rho_k = \frac{\lambda_k}{\mu_k}$. Non-greedy user is protected.

Back to multipath, work in progress.

Suppose: edge router can choose in which path to route a new flow.

Given: prices q_k^r of the various candidate routes, a natural policy is:

- Admit new connection $\Leftrightarrow \min_r q_k^r < U'_k(\rho_k)$, where $\rho_k = \sum_r \rho_k^r$
- If admitted, select cheapest path.

Example (symmetric users):

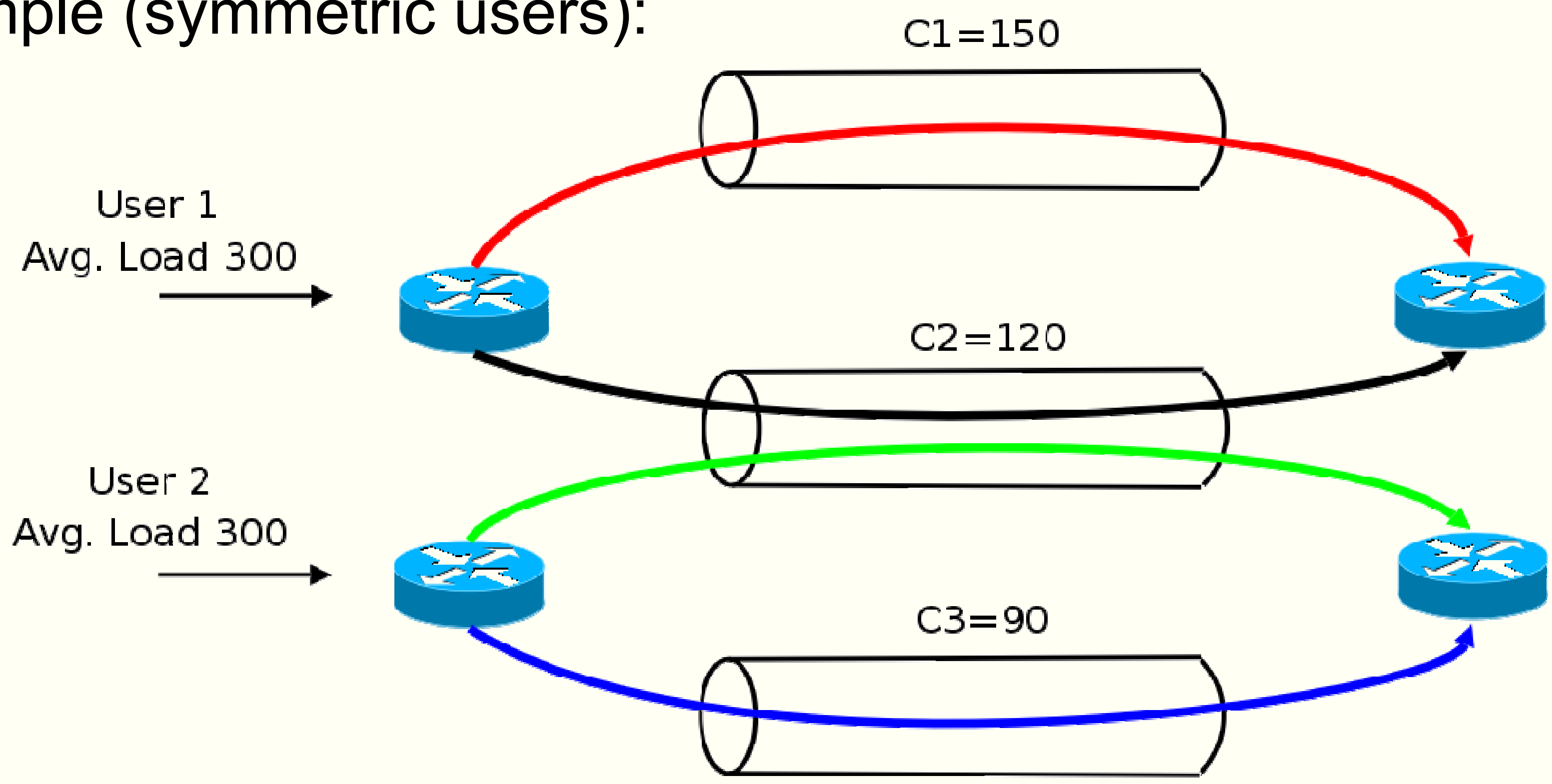

Multipath admission control Simulations

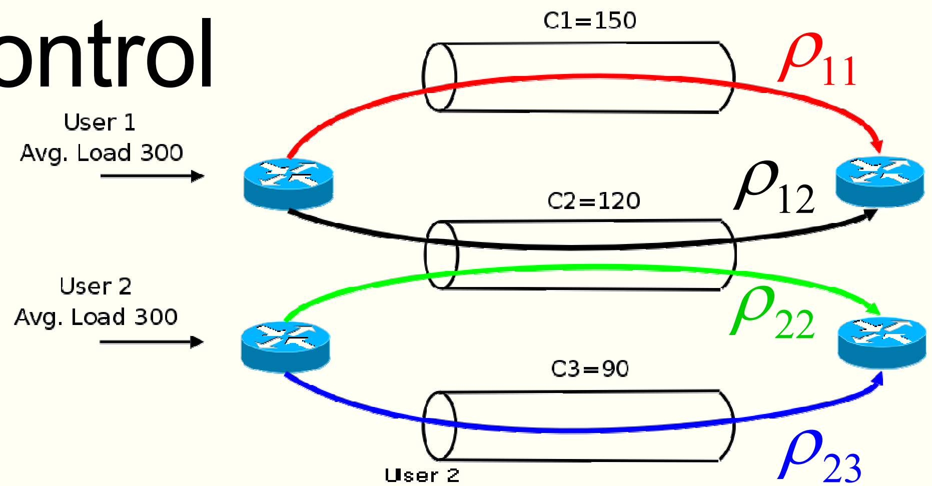

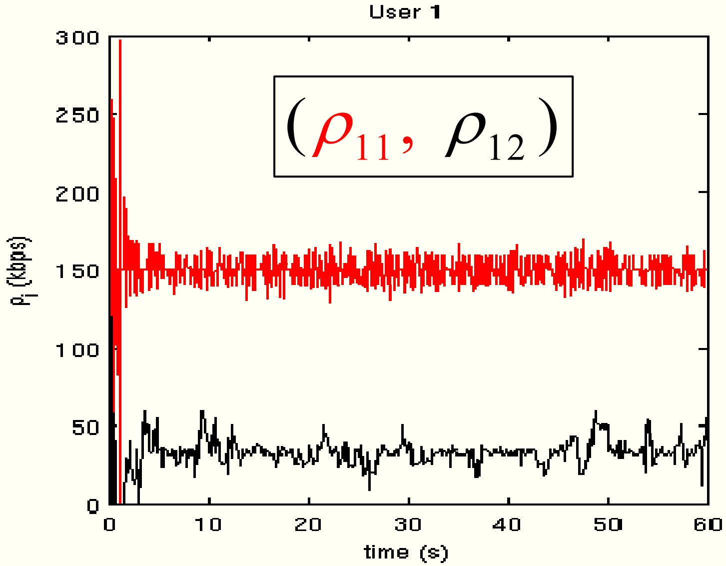

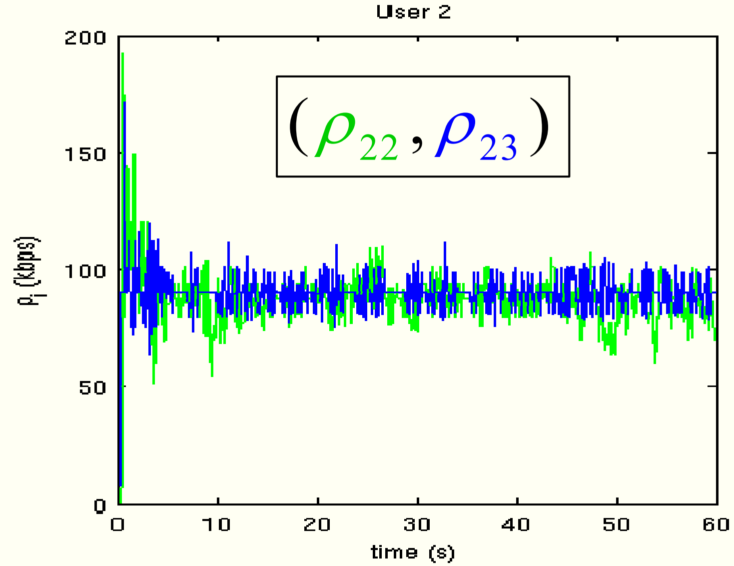

Conjecture: converges to optimal multipath welfare allocation,

$\max_\rho \sum_k U_k(\rho_k)$, subject to $\rho_k = \sum_r \rho_k^r$, $\quad \sum_k \sum_{r \ni l} \rho_k^r \leq c_l$.

Conclusions and Future Work

- We studied two cross-layer resource allocation problems:
  I. Congestion control and multipath routing.
  II. Congestion control and admission control
- Objective: welfare optimization. We designed decentralized control laws based on prices, achieve these equilibria.
- Dynamic analysis: local stability proofs in arbitrary networks, global results for simpler cases.
- Beware on simplistic models for delay!
- Simulation studies confirm and generalize the above theory.
- Future work:
  - Part I: Global proof in arbitrary networks, loss-based implementation.
  - Part II: Stability proofs for fluid limit model, ns2 implementation.
  - Combination: multipath admission control.