From control to networks, and into microeconomics?

Fernando Paganini
UCLA Electrical Engineering

Acknowledgements: Steven Low, John Doyle
Zhikui Wang, Sarma Gunturi,
Ahmad Fattahi.

Optimal control: an authoritarian view

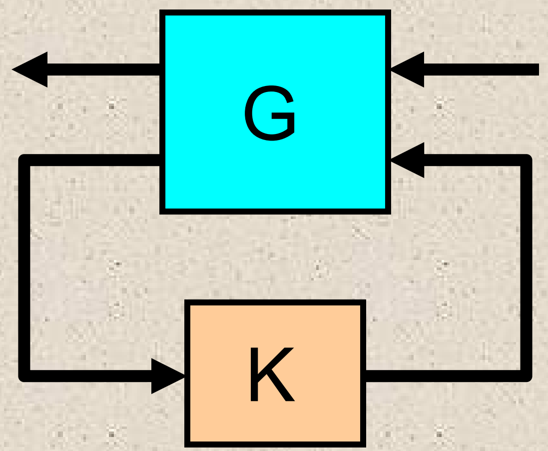

Theory of control design applies mainly to problems with:

- Single performance objective.
- Single decision maker (centralized information).

Complexity in analysis or design comes from the description of the plant G: nonlinearity, or the amount of structure in an uncertainty description $G \in \mathbf{G}$.

- Q: Is this picture relevant to large-scale systems?
- A: Rarely. As dimension grows, we run into
  - Multiple objectives
  - Decentralized information
- These requirements lead to complexity, even for a simple G.

Control with no central authority: the Internet.

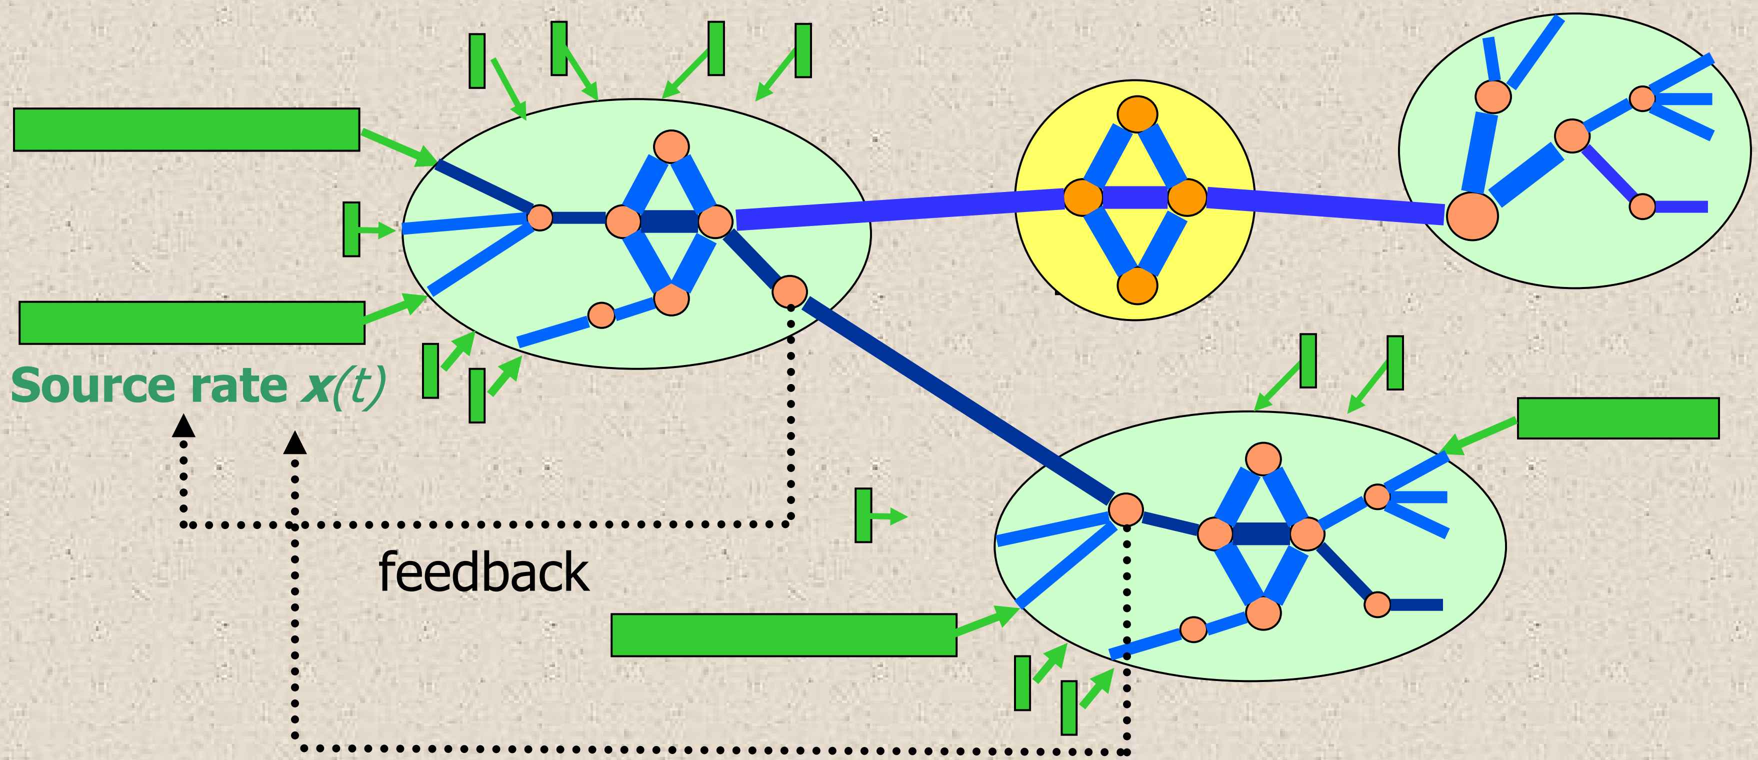

- Multiple decision makers, decentralized information. Nobody knows the system, let alone others' signals.
- Multiple, and not very explicit objectives.
- Anarchy? Not quite, some standards, IETF, etc. But limited quantitative understanding of its behavior.

Bringing some order: Kelly's formulation.

S source-destination pairs (long flows) share L links.

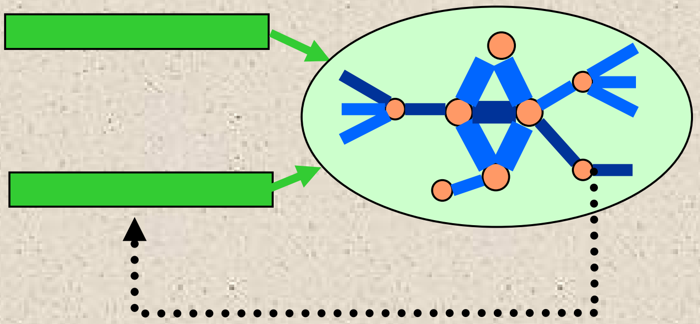

Routing matrix R (fixed for now)

$$R_{li} = \begin{cases} 1 & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}$$

- Source i has rate x_i, receives utility $U_i(x_i)$.
- Link l has rate $y_l = \sum R_{li} x_i$:, and capacity c_l.

$$\boxed{y = Rx}$$

Social welfare optimization: $\max_x \sum_i U_i(x_i)$ subject to $Rx \le c$.

A market way to decentralize the optimization:

- Link l generates congestion measure or price p_l
- Source i sees price $q_i = \sum R_{li} p_l$, optimizes own profit: $\max_{x_i} U_i(x_i) - q_i x_i$

Issue: generating prices that align individual and social welfare; they should be the Lagrange multipliers of the dual of the convex problem.

Finding them dynamically: $\dot{p}_l = \gamma \left[y_l - c_l \right]^+_{p_l}$ (Low-Lapsley '99)

Beyond quasi-statics → Control viewpoint.

Time-scales determined by the plant dynamics: network delays.

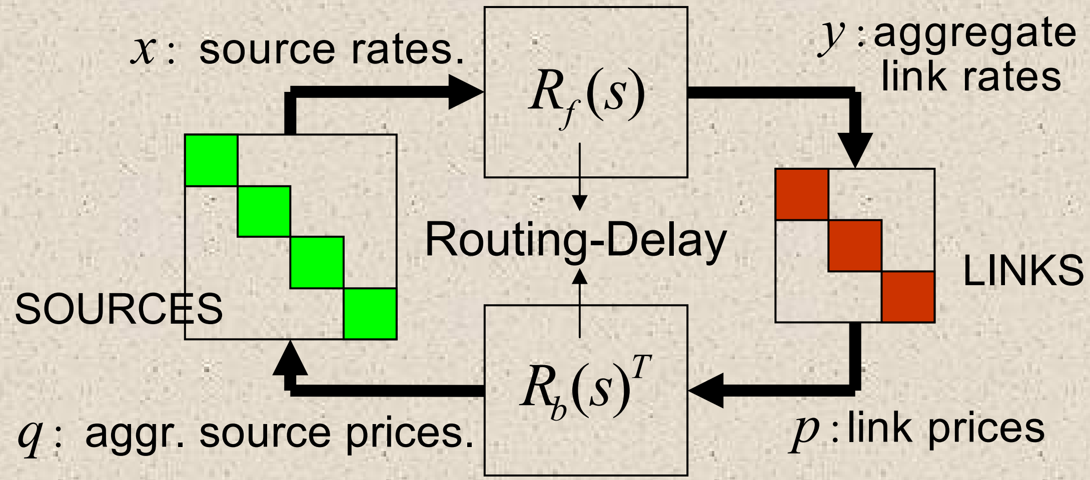

$$q_i(t) = \sum_l R_{li} p_l(t - \tau_{i,l}^b)$$

$$y_l(t) = \sum_i R_{li} x_i(t - \tau_{i,l}^f)$$

Delayed model in Laplace transform (for fixed delays):

$$\left[R_f(s)\right]_{l,i} = \begin{cases} e^{-\tau_{i,l}^f s} & \text{if } i \text{ uses } l \\ 0 & \text{otherwise} \end{cases}$$

$$\text{RTT}: \tau_i = \tau_{i,l}^f + \tau_{i,l}^b$$

- Challenges for design of source and link laws:
  - Highly decentralized information.
  - Unknown network: topology, parameters.
  - Accommodating both equilibrium point and dynamic behavior.
- Solutions feature:
  - Handcrafted designs, use equilibrium information to bound loop gains.
  - Mathematical proofs of local stability in any network, using multivariable control.
  - Time-scale separation between desired equilibrium and dynamic behaviors.

Example: A "primal-dual" congestion control

[joint work with Z. Wang, S. Low and J. Doyle, Infocom '03]

$$\dot{x}_i = -\frac{\alpha_i}{M_i \tau_i} x_i \dot{q}_i + \beta_i \frac{(U_i^{'}(x_i) - q_i) x_i}{\tau_i}$$

$$\dot{p}_l = \frac{1}{c_{0l}} \left[y_l - c_{0l} \right]^+_{p_l}$$

α_i, β_i parameters,

τ_i: round-trip time

M_i : bound on # of bottlenecks

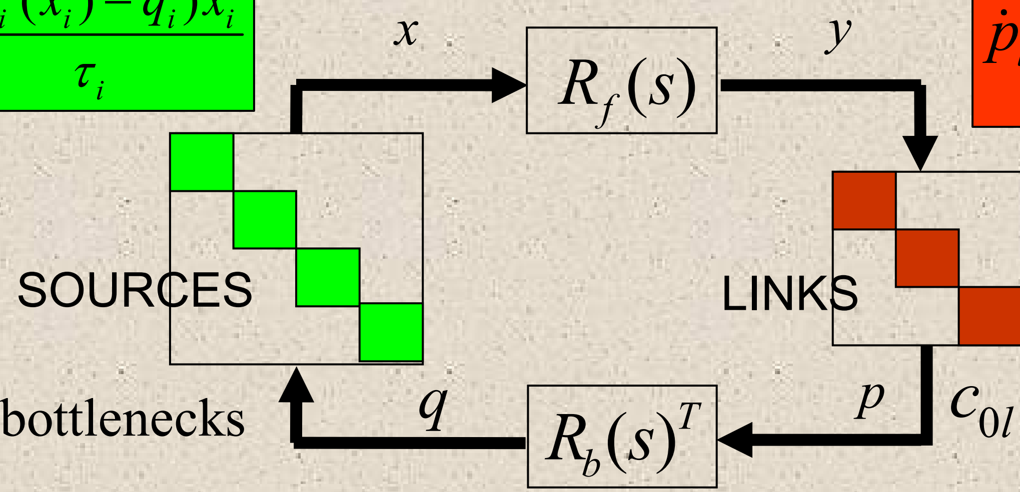

c_{0l}: virtual capacity, slightly below c_l

- Design strategy:
  - Start from local dynamics, compensate for RTT at the source, and use network equilibrium to control the overall gain in a distributed manner.
  - Incorporate equilibrium features as a slower dynamics
- Given a bound on RTT, one can provably satisfy:
  - Almost full network utilization, empty equilibrium queues.
  - Resource allocation according to utility functions.
  - Local dynamic stability for any network.
  - Global stability: for simple networks, conditions similar to local ones.

Still, we have a single objective.

We have avoided the inherently multi-objective nature of the problem by imposing operation in the maximizer of $\sum_i U_i(x_i)$. Justification?

More precisely, consider the "utility possibility set" $\mathbf{U}$ achievable by sources, given the capacity constraints $Rx \le c$.
It is clearly desirable to be Pareto optimal, but which one to pick?

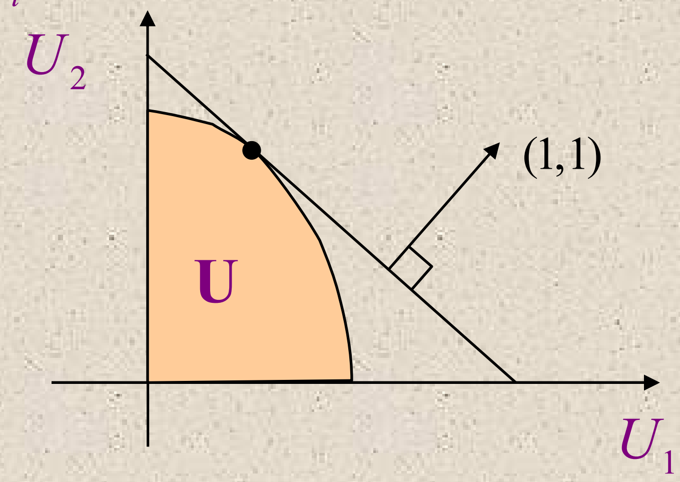

An argument from microeconomics:

if utilities are expressed in units of a tradeable "numeraire" good (money), then this is the only market equilibrium : from any other point, a side payment between users makes it advantageous to switch to this point.

This, however, does not appear too realistic for the networking context, especially at fast time-scales.

The viewpoint of welfare economics.

A policy maker can arbitrate between multiple utilities by maximizing a social welfare function $W(u_1,\ldots,u_S)$, increasing and concave. Examples:

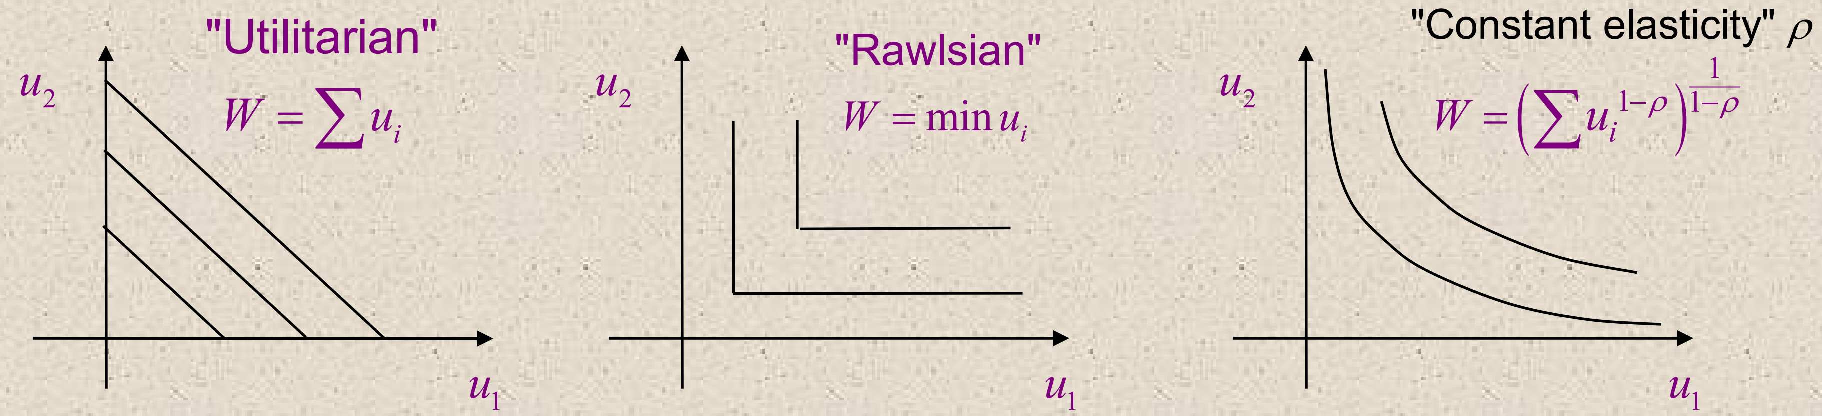

Used before — "Max-min" fairness. — Intermediate family

In the limit as $\rho \to 1$, the constant elasticity welfare functions gives $W=\sum \log(u_1)$. The resulting Pareto optimal point gives the so-called "Nash Bargaining Solution", which satisfies the following axioms:

(i) Symmetry, (ii) Invariance to change of scale in utilities

(iii) Independence of irrelevant alternatives: reducing $\mathbf{U}$ by non-optimal points does not change choice.

Two interpretations for "proportionally fair" rate allocation.

$$\max_x \sum_i \log(x_i) \quad \text{subject to} \quad Rx \le c.$$

1. (Kelly): Utility is $\log(x_i)$, use utilitarian welfare function.
2. (Yaiche-Mazumdar): Utility is x_i, use Nash bargaining solution.

Asymmetric generalization: $\max_x \sum_i w_i \log(x_i)$

1. (Kelly) Utility is $w_i \log(x_i)$, w_i is "willingness to pay".
2. Utility is x_i, use social welfare function $\sum_i w_i \log(u_i)$, Nash with asymmetric bargaining power.

The supply side of this economy

A dual problem would be to fix a demand, and allocate network resources so as to minimize a certain cost.

- At long time-scales, this could include bandwidth provisioning.
- At short time scales, the main degree of freedom is *routing.*

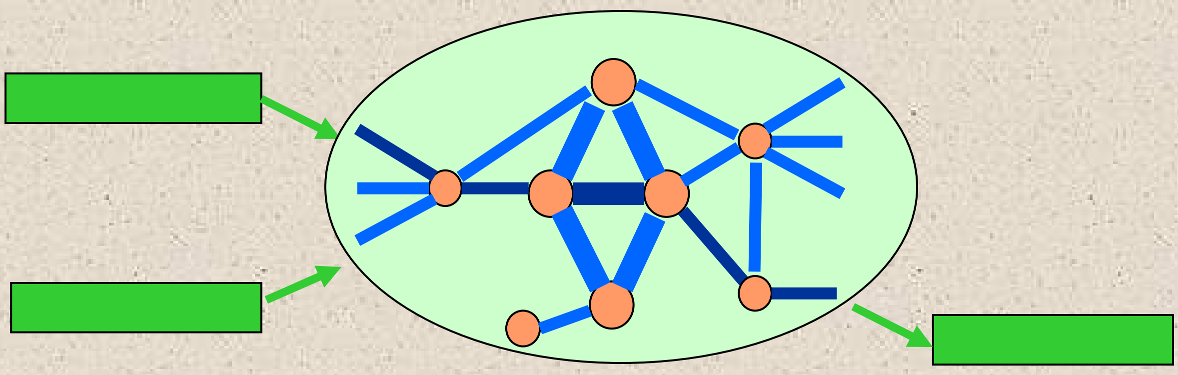

Traffic engineering problem: given a traffic matrix of flows between source-destination nodes, minimize $\sum \phi_l(y_l)$ subject to meeting demands. Cost $\phi_l(y_l)$ could represent, e.g. a delay at the link.

For single-path (IP) routing, the problem is hard.

Allowing multiple routes per source: multicommodity flow problem: let z_l^i be the rate of source i going through link l, solve

$$\min \sum_l \phi_l(y_l) \text{ subject to}$$

$$y_l = \sum_i z_l^i,$$

$$\sum_{l \text{ in}} z_l^i = \sum_{l \text{ out}} z_l^i \text{ at interior node,}$$

$$\sum_{l \text{ in}} z_l^i = x_i \text{ at source node.}$$

- If $\phi_l(y_l)$ are convex, optimization can be solved by a central planner.
- Using duality, we can obtain weights for which the routes are which give routes via shortest path.
- Decentralized versions involve sources, or edge routers keeping track of routes.

Combining demand and supply?

A "partial equilibrium" formulation from microeconomics would be to maximize the "Marshallian surplus" $\sum U_i(x_i) - \sum \phi_l(y_l)$ subject to "market-clearing" (supply=demand).

For a simple divisible good, solution is "competitive equilibrium" with price $p = U_i'(x_i) = \phi_l'(y_l)$.

- Applying this to networks:
  - Combining Traffic Engineering and congestion control; dynamics, time-scales (see J.Wang et al '03)
  - Justification? Note that "network cost" is really QoS for sources.
  - If demand is an aggregate of elastic flows with random duration, what is a good model?

Other results from the economic viewpoint.

- Demand problem in which sources are not price-takers. i.e., they anticipate their influence on prices. Problem becomes a game, overall utility deteriorates by < 25% (Johari-Tsitsiklis ’04, Sanghavi-Hajek-’04).
- Traffic engineering problem with “selfish routing”. Again, this is a game, and there are results on deterioration of the overall cost (Roughgarden & Tardos ‘02).
- The network as a profit-seeking entity: monopoly and oligopoly theory, (Hajek & Gopal ’04).
- Bidding formulations to differentiated services: Vickrey-Clarke-Groves method (e.g., Shu & Varaiya ‘03).
- Open questions:
  - Decentralization, and simplicity of implementation?
  - Dynamics.

Another layer: Wireless power control

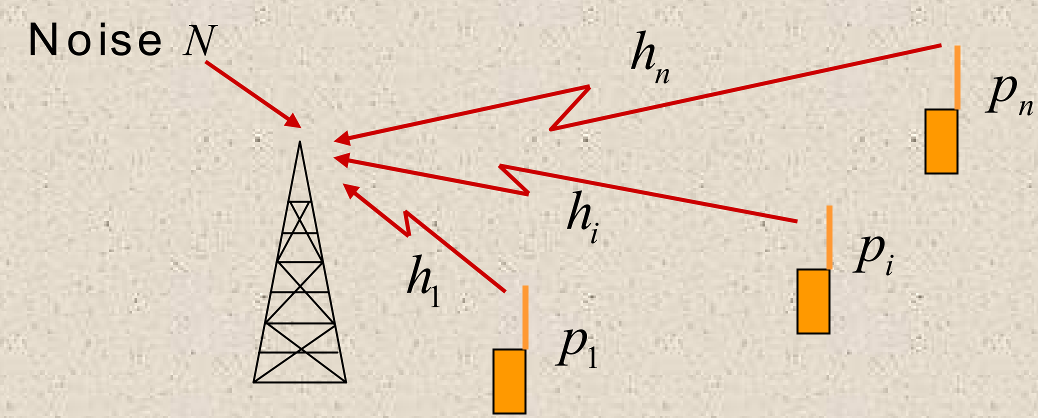

CDMA: SIR for mobile i :

$$\gamma_i = \frac{W}{B} \cdot \frac{h_i p_i}{N + \sum_{j \neq i} h_i p_j}$$

Assume:

- Elastic traffic sources, utility is proportional to rate x_i.

  For simplicity, assume the rate depends on SIR through the Gaussian Shannon capacity formula. So pick $U_i = k_i B \log(1 + \gamma_i)$.

Non-convex utility possibility set.

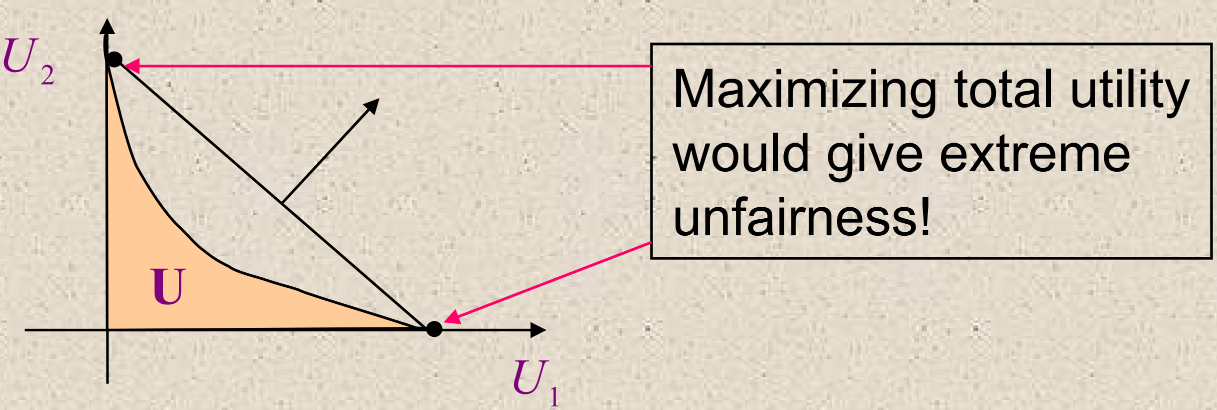

Other ways of arbitration.

Extending Nash bargaining

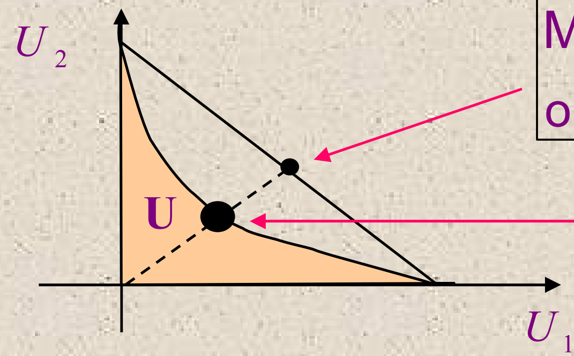

Maximizer of $\sum \log(u_1)$ over the convex hull.

Choose Pareto optimum which is proportional.

Extended Kazai-Smorodnisky solution

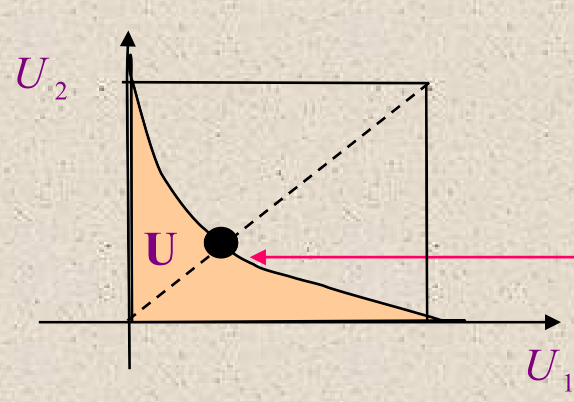

Choose Pareto optimum where each source gets a fixed proportion of their maximum utility.

A convex approximation (Johansson et al.)

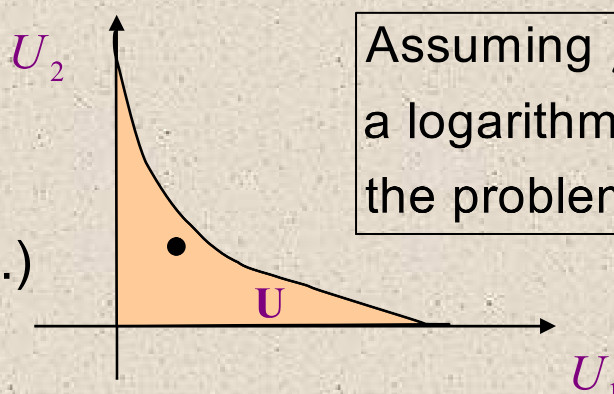

Assuming $\gamma_i >> 1$ so that $U_i \approx \log(\gamma_i)$, a logarithmic change of variables convexifies the problem. Optimize over this set.

Decentralized information? A non-cooperative version.

Goodman-Mandayam '00
Alpcan et al '02,
Gunturi-Paganini '03

- Each user plays the best response assuming others' powers are fixed. If we simply use the utility $U_i(p_i, p_{-i}) = \log(1+\gamma_i)$, then every user will blast at maximum power.

- Idea: add a cost term proportional to power. Users then maximize $U_i(p_i, p_{-i}) = \log(1+\gamma_i) - a_i p_i$ over $p_i \in [0, p_{\max}]$.
  This is convex, the solution is solution $p^+ = F(p) = \text{sat}[Mp + b]$, affine map saturated to $[0, p_{\max}]$.

Properties:

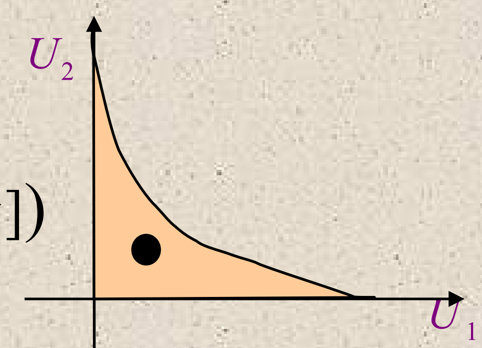

- Nash equilibrium exists, unique if the number of users $n < \frac{B}{W}$. In that case, iteration $p[k+1] = F(p[k])$ converges to it.
- If $a_i = kh_i$, (so cost is proportional to received power from source i) then equilibrium will equalize SIRs.
- Nash equilibrium need not be Pareto optimal.

Economic theory itself

- Most of microeconomic theory applies to characterizations, properties of equilibria. In regard to resource allocation, we have much to draw from this field.
- On the other hand, less is known about dynamics outside equilibrium. One classical model is the "Tatonnement" dynamics for price,

$$\frac{dp}{dt} = \gamma \left[\text{agg.demand} - \text{supply} \right]$$

Exactly what was used before in link laws.

- However economists question its descriptive value.

- Nevertheless, for "designed economies" as in networks, dynamics can have very precise meaning. Control insight could be valuable.
- Another example: dynamics of fictitious play in games with learning (Shamma and Arslan '03) .

Conclusions

- When migrating control theory to large scale networks, we immediately face challenges of decentralized info, and often multiple objectives.
- Feedback design becomes more difficult to automate, so we use handcrafted designs and mathematical analysis (“design-for provability”).
- Elements of economic theory come in naturally in network resource allocation problems: we’re still on a learning curve.
- It is likely, still, that a control viewpoint can go beyond what is known about dynamics, especially for “designed markets” where all models are by definition valid.