Bandwidth auctions and their parallels to power

Fernando Paganini.
Universidad ORT, Uruguay.

Outline:

1. Intro: resource allocation and pricing in comm networks.
2. Background on auctions.
3. Auctions for Internet bandwidth.
4. Connections to auctions in the power grid.

Lund, May 2011

1. Intro: Rate allocation in the Internet

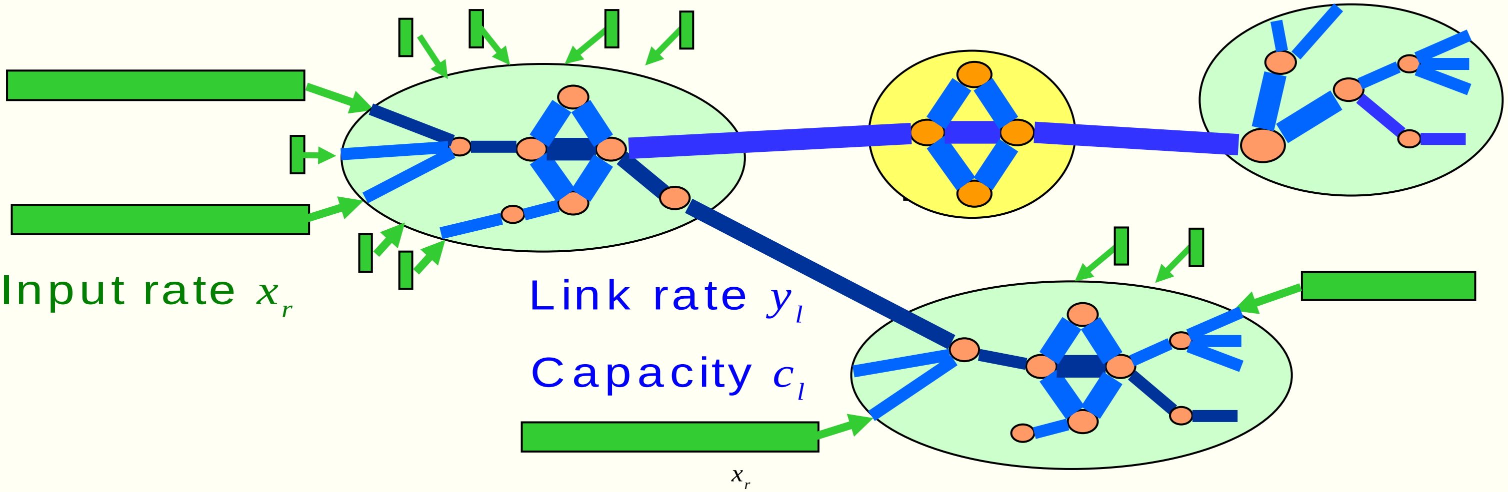

- Network of links, indexed by l, with capacity c_l (e.g. Mbps).
- End-to-end flows, indexed by r, rate x_r
- Link rate aggregation: $y_l = \sum_r R_{lr} x_r$, where $R_{lr} = \begin{cases} 1 & \text{if route } r \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}$

Network Utility Maximization (Kelly ´98)

$$\max_x \sum_r \underbrace{U_r(x_r)}_{\text{SOURCE UTILITY}}, \quad \text{subject to} \quad \underbrace{Rx \le c}_{\text{LINK CAPACITY CONSTRAINTS}}$$

Solution through duality (Low-Lapsley ’99)

Lagrangian $L(x,p) = \sum_r [U_r(x_r) - q_r x_r] + \sum_l p_l c_l.$

- Lagrange multiplier p_l: "congestion price" of link l.
- $q = R^T p \quad (q_i = \sum_l R_{li} p_l:$ total price of route i).

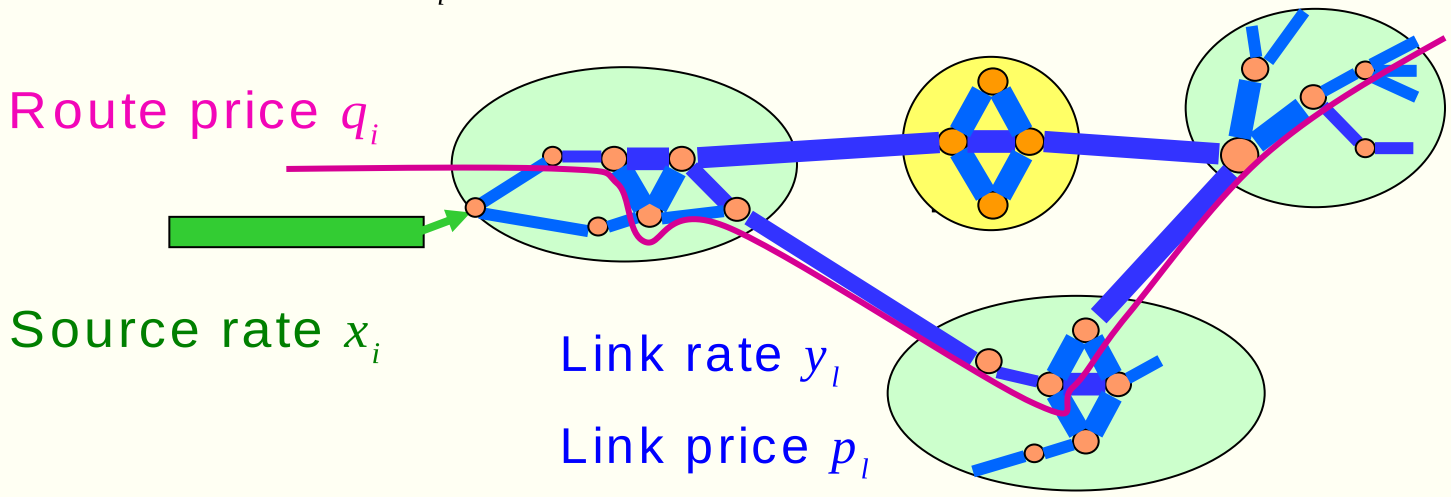

Dual algorithm

- Sources solve $x_r := \operatorname{argmax}_{x_r} [U_r(x_r) - q_r x_r],$
- Links update prices as $\alpha_l := [\alpha_l + \gamma_l (y_l - c_l)]^+,$ prices sent back.

Uses for this “virtual economy” of bandwidth

- Interpretation of the equilibrium and dynamic properties of current TCP congestion control protocols.
- Guides in design of new protocols with:
  - Better dynamic properties (convergence, etc.)
  - Use other utility functions, achieving other notions of fairness.
- Extension to cross-layer optimization, including other layers of the protocol stack, for both wired and wireless networks:
  - Routing
  - Medium access control (Scheduling, random access).
  - Physical layer control (power control, modulation,…)
- However, the “real” economy of bandwidth doesn’t work this way. Why?
  - Bandwidth has been abundant, not crucial to optimally allocate it.
  - This control is faster than the human time-scale.

1. Auctions

- Popular trading mechanism
  - Fast, reliable and transparent way of setting market price.
  - Various mechanisms exist, single and multi-unit auctions.

- Open auctions for sale of a single unit
  - English: ascending bids, open-outcry, terminates when one bidder is left.
  - Dutch: descending bids, open-outcry, terminates when one bidder shouts “mine”.

- Closed, sealed-bid auctions:
  - First price: highest bidder wins, pays his/her bid.
  - Second price (Vickrey): highest bidder wins, pays 2nd bid.

Vickrey Auctions and truth revelation

Bidder of $5 bid wins the auction,
but pays $3.7 for the item.

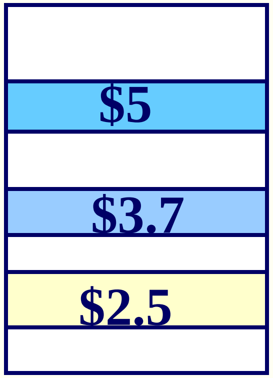

- In a second-price auction, it is rational for participants to bid their true valuation for the item:
  - They gain no reduction in payment by bidding below their valuation.
  - They might lose the auction if they do so.

Multiple unit auctions

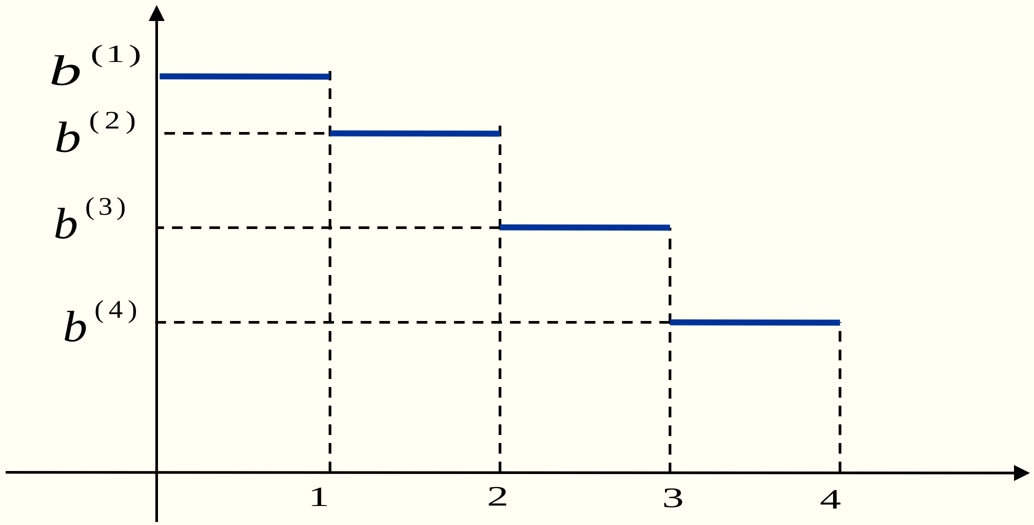

Example: sell 3 units, choose 3 highest bidders. Alternatives:

- Charge bid amount: gives incentives for bidding below valuation.
- Vickrey-Clarke-Groves (VCG) principle: charge users the loss of valuation imposed to others by their presence.

In this case: presence of 2nd bidder changes others' total valuation

from $b^{(1)} + b^{(3)} + b^{(4)}$ to $b^{(1)} + b^{(3)}$ $\Rightarrow$ Should charge $b^{(4)}$.

- It is shown VCG makes it rational to reveal the true valuations.

Is there a cost for truth revelation in loss of revenue for the seller?

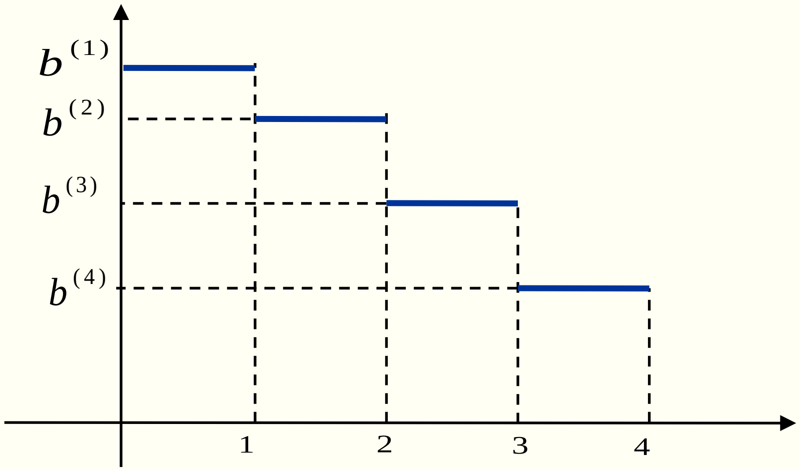

Revenue Equivalence Theorem (Vickrey, Myerson, Riley-Samuelson)

Assume:

- Risk neutral buyers, valuations drawn from a known distribution.
- Mechanism assigns objects to bidders with highest valuation.

$\Rightarrow$ At Bayesian Nash equilibrium, all auction mechanisms yield the same expected revenue for the auctioneer.

However, equivalence does not hold if buyers are risk averse, do not know the distribution, or do not have unbounded rationality.

- In such cases a first-price auction may give higher revenue than VCG.

Procurement auctions

Auctioneer buys one or more items from lowest offers (asks).

- e.g., buy 2 items from offers 1, 2.
- VCG: pay $a^{(3)}$ to both. .

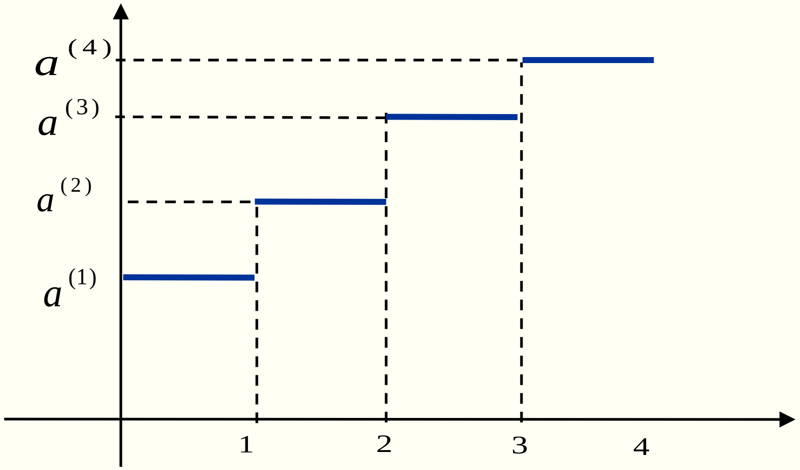

Double or two-sided auctions

Bids for buying and selling

- Sell $\max\left\{k : b^{(k)} \geq a^{(k)}\right\}$. Here, 2 items.
- Equilibrium price: crossing point.

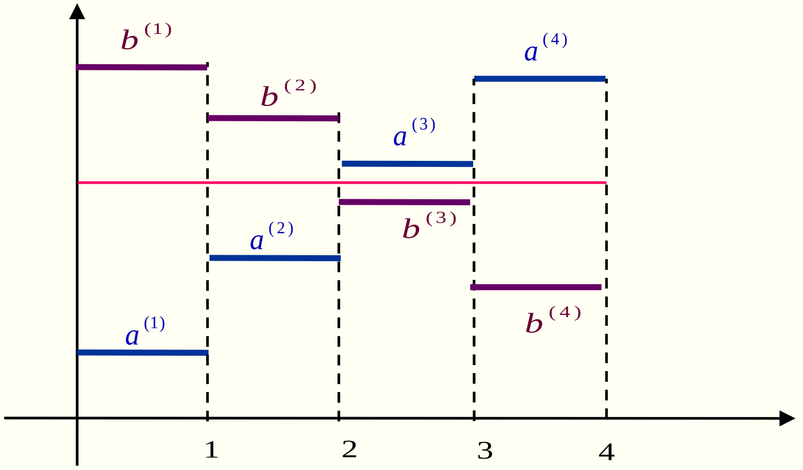

2. Internet Bandwidth Auctions

Based on 2011 paper in *Computer Networks.* Joint work with:
- Pablo Belzarena (Universidad de la República, Uruguay)
- Andrés Ferragut (Universidad ORT, Uruguay)

- Scenario: a network periodically auctions capacity.
- Users submit bids for amounts of end-to-end bandwidth.
- A distributed algorithm must optimally assign capacity.
- Motivation: overlay for premium services over the Internet.

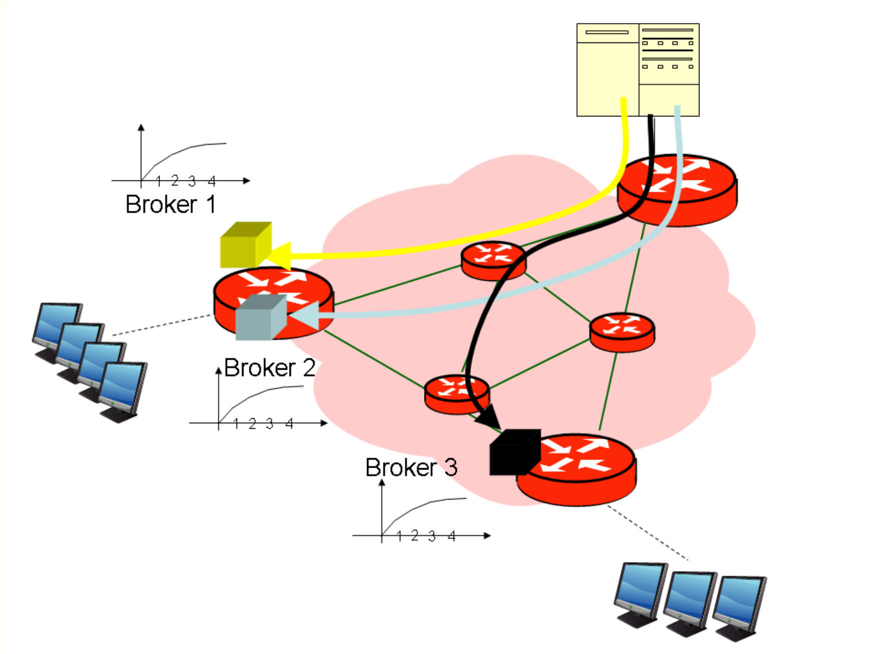

- Related work: Lazar – Semret '00, Shu – Varaiya '03, Reichl – Wrzaczek'05, Maillé – Tuffin '06, Courcoubetis et al. '07.

Three issues and solution features

1. Auction allocation/payment mechanism:
   - Optimize the value of accepted bids.
   - Charge 1st price, VCG would have high complexity (Maillé-Tuffin ´07).
   - Revenue equivalence argument.

2. Distributed auction over a general network topology.
   - Bids submitted to "bandwidth brokers" distributed across the network.
   - Bidders need not know network topology, capacity, etc.
   - Brokers run a distributed algorithm to allocate the auction.

3. Inter-temporal constraints.
   - Auctions are held periodically, for currently available capacity.
   - Service may be longer than the auction period, and reservations are in place: a connection, once assigned, cannot be displaced by future bids.
   - So the seller optimize over the risk of future bids: selling capacity now with a low bid can cause the rejection of a better bid in the future.

Notation: auction for a single service

- N bids $b^{(1)} \geq b^{(2)} \geq b^{(3)} \geq \cdots \geq b^{(N)}$ for σ units of bandwidth.

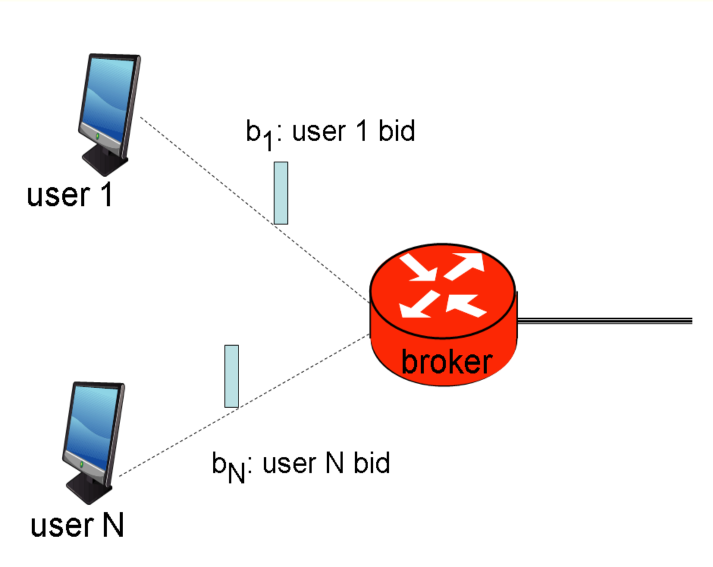

- Revenue from allocating bandwidth $a = N\sigma$: (first-price auction)

$$U_b(a) = \sum_{i=1}^{a/\sigma} b^{(i)}.$$

Interpolate to piecewise linear, concave function.

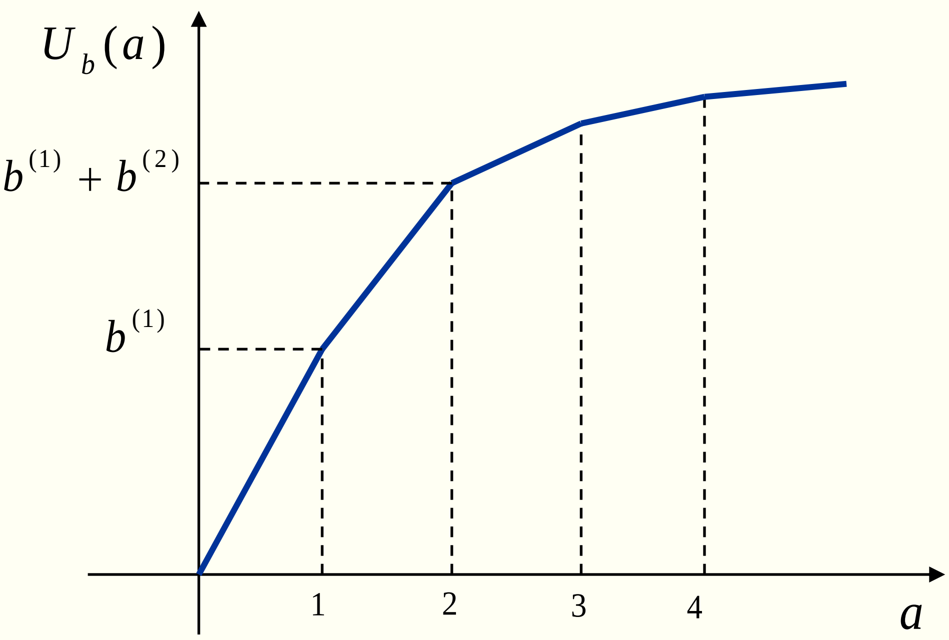

Auction over a network

Let r represent a service, characterized by a reserved bandwidth σ_r between 2 endpoints.

Single-path case: service has a fixed route, defined by a routing matrix R:

$R_{lr} = 1$ iff route r uses link l.

Multipath generalizations are available.

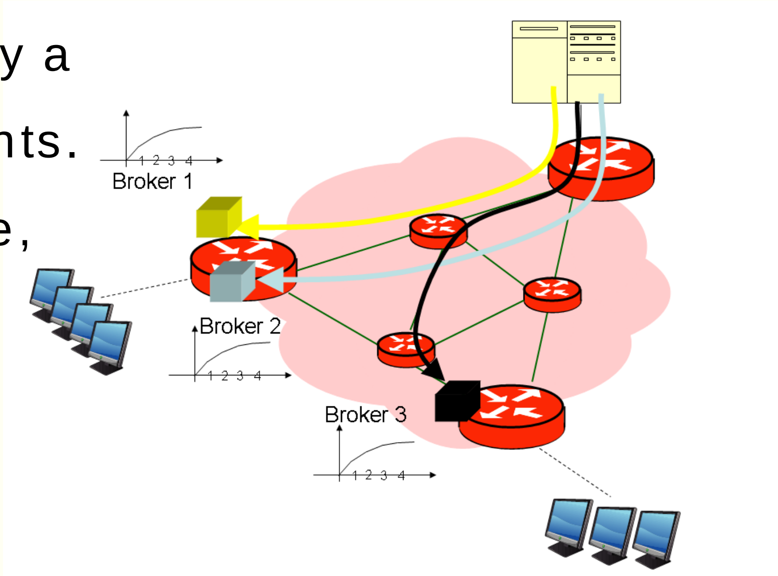

Broker collects bids for this service: $b_r^{(1)} \geq b_r^{(2)} \geq b_r^{(3)} \geq \cdots \geq b_r^{(N_r)}$.

If we admit bids for a total rate a_r, the total revenue is $U_{b_r}(a_r) = \sum_{i=1}^{a_r/\sigma_r} b_r^{(i)}$.

Network optimal revenue auction

$$\max \sum_r U_{b^r}(a^r)$$

subject to $\sum_r R_{lr} a_r \leq c_l, \qquad \frac{a_r}{\sigma_r} \in Z$

Integer program.

Relaxation is a concave utility maximization as introduced in Kelly '98.

Distributed allocation algorithm

Lagrangian $L(a,\alpha) = \sum_r U_{b^r}(a^r) + \sum_l \alpha_l [c_l - \sum_s R_{lr} a_r] + \sum_l \alpha_l c_l$

Dual algorithm (from Low-Lapsley '99)

- Brokers solve $a_r := \mathrm{argmax}_{a_r} [U_{b_r}(a_r) - q_r a_r]$, with $q_r = \sum_l R_{rl} \alpha_l$ route price. Amounts to selecting bids better than q_r.
- Communicate a_r to links, who update prices as $\alpha_l := [\alpha_l + \gamma_l (y_l - c_l)]^+$, with $y_l = \sum_s R_{lr} a_r$ link bandwidth demand. Prices sent back to brokers.
- Can be implemented in the control plane, variant of RSVP protocol.

Difficulty:

- $U_{b_r}(a_r)$ not strictly concave, algorithm will ``chatter" around optimum.
- Solution: proximal optimization with extra variable d_r

$$\max \sum_r U_{b_r}(a_r) - \sum_s \frac{\kappa_r}{2}(a_r - d_r)^2, \quad \text{subject to} \quad \sum_r R_{lr} a_r \le c_l,$$

(see Lin-Shroff '06, also useful for multipath case).

Periodic auctions for one link

- Service of bandwidth $\sigma = 1$, single link of capacity C.
- Collect bids for time of length T, allocate a^k bandwidth units at time kT
- Allocated users have a reservation for service duration, assume $\exp(\mu)$.
- $p = e^{-\mu T}$: probability that a connection is still active at the next auction.

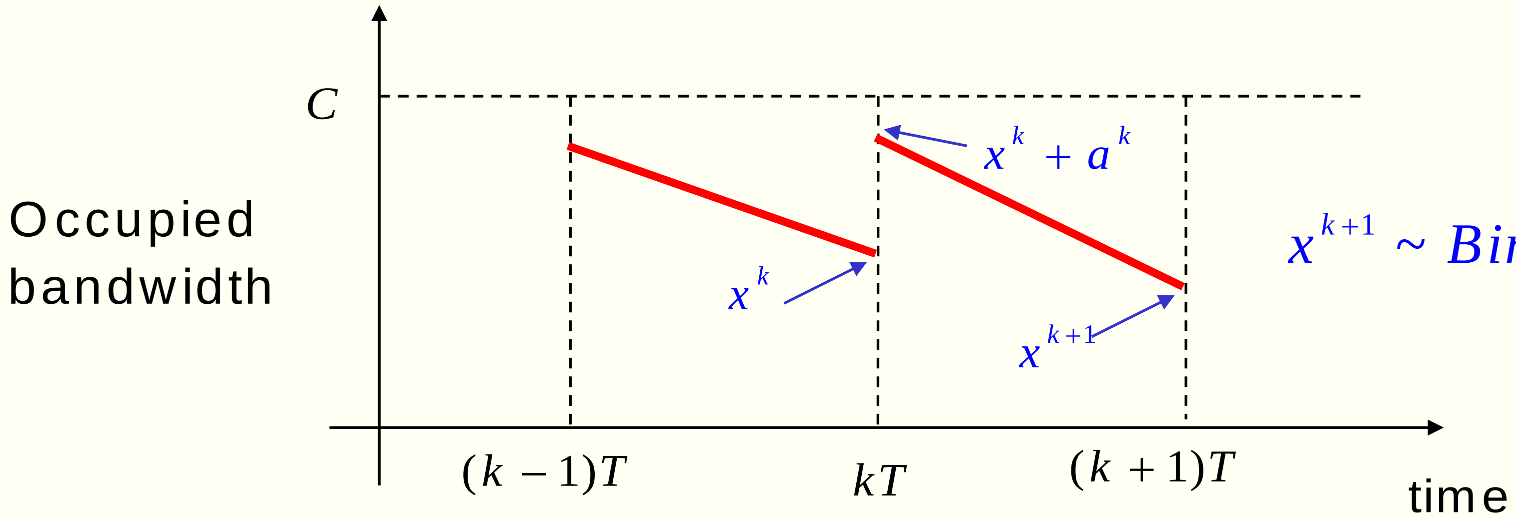

$$x^{k+1} \sim Bin(x^k + a^k, p).$$

Myopic policy: $a^k = C - x^k$, sell all currently available capacity.

May miss higher bids in the future.

What is the optimal policy?

Optimal revenue problem $\max \lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} E\left[U_b(a^k)\right]$

- Expectation with respect to bids (assume known distribution, otherwise can be estimated) and the departure process.
- This is a Markov Decision Process (MDP).
- Solution is a policy $a_k = a(s_k)$, where the state is $s_k = (x_k, b_k)$.
- $a(s)$ can be found numerically, large compupational cost.

Receding horizon approximation:

$$a^0 = \arg\max_{a \le C - x^0} [U_b(a) + E_{x^1}\bar{U}(C - x^1)]$$

Here $\bar{U}(a) = E\left[U_b(a)\right] = \sum_{i=1}^{a} E[b^{(i)}]$

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

- Reduces to intersection of decreasing marginal utilities (bids) with acceptance thresholds, w_k: cost of missed future opportunities.

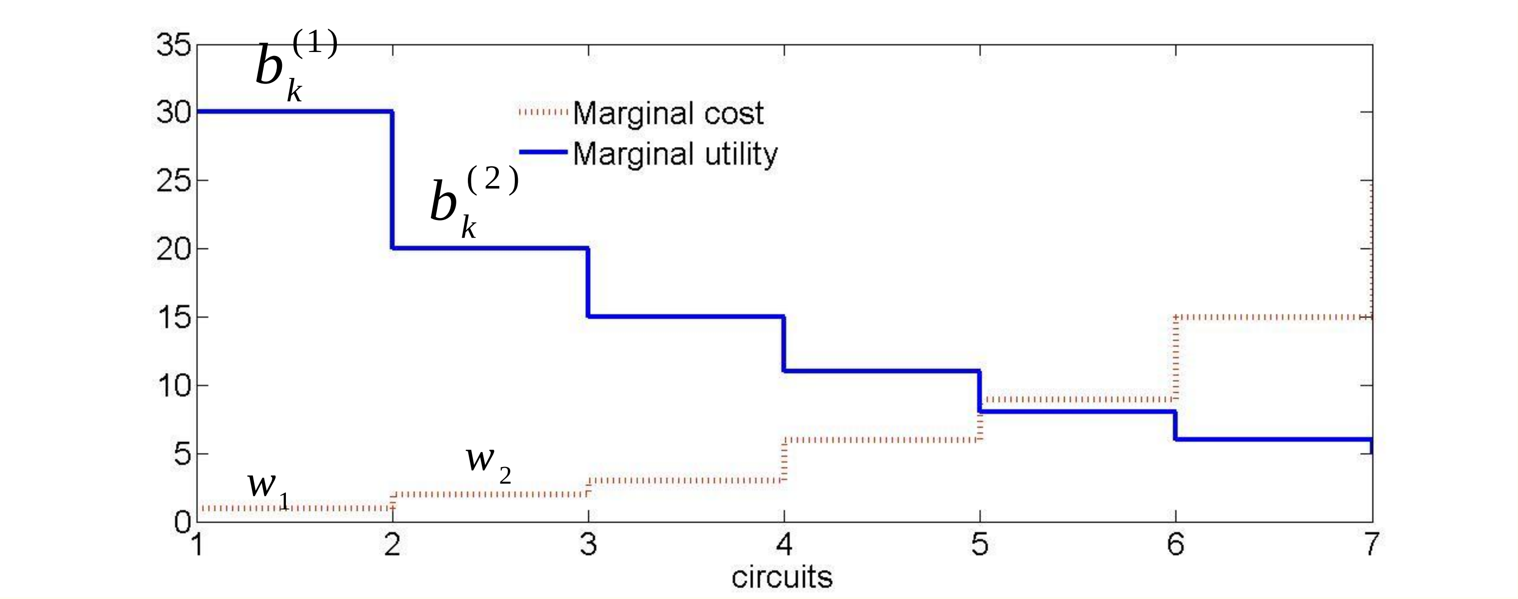

- In simple examples, one-step ahead policy approximates well the optimal revenue.

Fluid aproximation to receding horizon policy:

$$\max \underbrace{U_b(a)}_{\text{CURRENT REVENUE}} + \underbrace{\bar{U}(z)}_{\text{EXPECTED NEXT STEP REVENUE}}, \text{ subject to } \underbrace{x_0 + a \le C}_{\text{CURRENT CAPACITY CONSTRAINT}}, \underbrace{p(x_0 + a) + z \le C}_{\text{EXPECTED FUTURE CAPACITY CONSTRAINT}}.$$

The network case

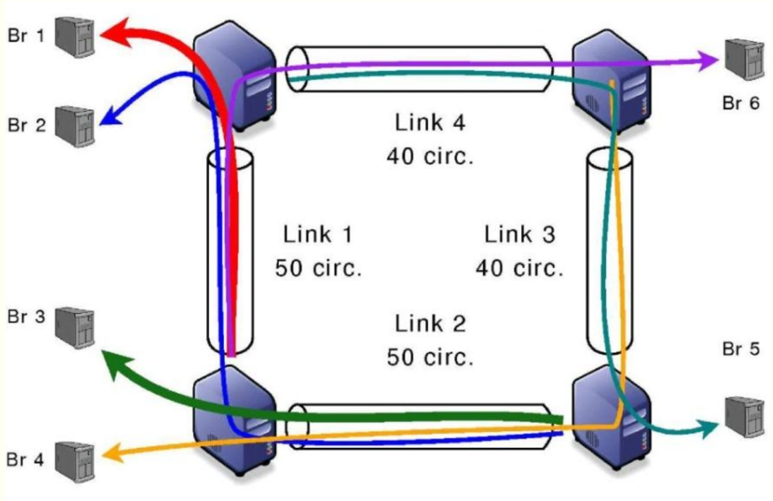

a^r : current allocation, broker r

z^r: expected next-step allocation.

$$\max \sum_r U_{b^r}(a^r) + \bar{U}(z^r)$$

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

- Network utility maximization, with additional (step ahead) variables.
- Distributed implementation: dual algorithm, proximal method.
- Converges to fluid approximation, requires roundoff.
- Extends to multipath routing.

3. Auctions in the power grid

- In deregulated wholesale electricity markets, auctions are carried out by the Independent System Operator (ISO) to buy power for, e.g. one day ahead demand.

- Procurement auction of the type mentioned before, except that power is divisible and may be offered in different amounts.

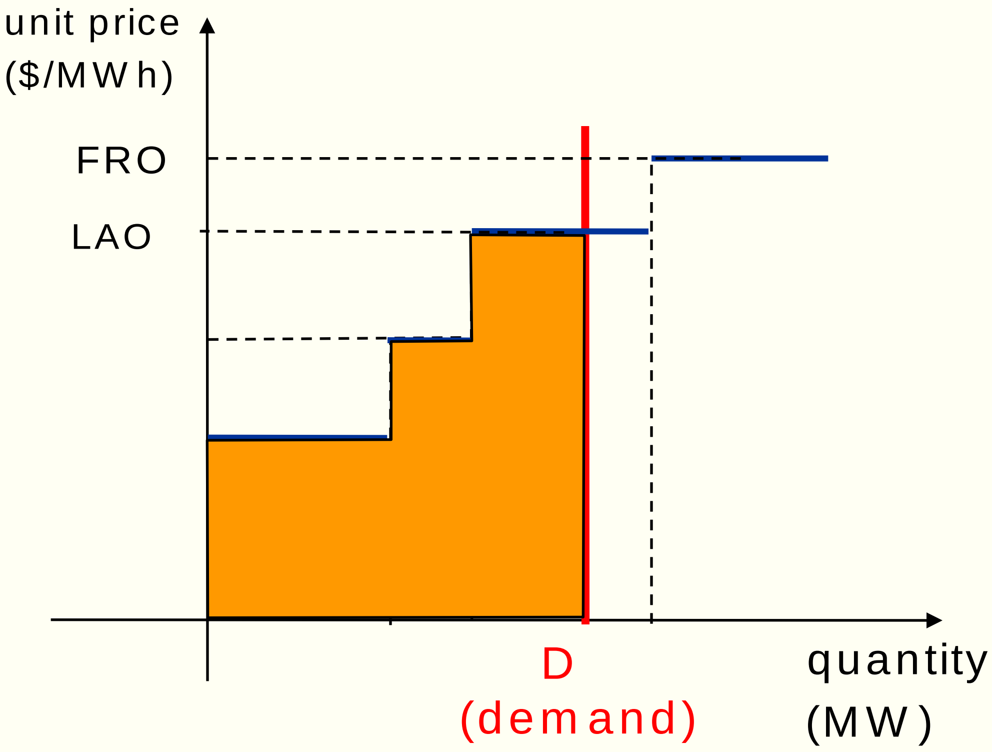

- Typically, a common market price is set. Possible choices:

- Last Accepted Offer,
- First Rejected Offer
- Something in between.

- Two-sided (supply-demand) auctions also possible.

  Ref: Zimmerman '2010.

Recall: three issues in bandwidth auctions

1. Auction allocation/payment mechanism
   - Optimize the value of accepted bids.
   - Charge 1st price, VCG would have high complexity (Maillé-Tuffin ´07).
   - Revenue equivalence argument.

2. Distributed auction over a general network topology.
   - Bids submitted to “bandwidth brokers” distributed across the network.
   - Bidders need not know network topology, capacity, etc.
   - Brokers run a distributed algorithm to allocate the auction.

3. Inter-temporal constraints.
   - Auctions are held periodically, for currently available capacity.
   - Service may be longer than the auction period, and reservations are in place: a connection, once assigned, cannot be displaced by future bids.
   - So the seller optimize over the risk of future bids: selling capacity now with a low bid can cause the rejection of a better bid in the future.

Issue 1 revisited: pricing mechanism

- In the Internet problem, we maximized the value of the bids admitted to the network.
- Analog for a procurement auction: minimizing the cost of acquired power. Good for customers.
- However, first-price (pay offers at their declared value) is not favored, a common market clearing price (MCP) is paid. Not quite VCG, but closer in spirit to 2nd price schemes. Is social welfare of *sellers* the objective?
- Difference with bandwidth case: players (generator firms) are sophisticated, can afford to game the system.
- Still, gaming is possible by combining two offers, e.g. hockey-stick bidding:

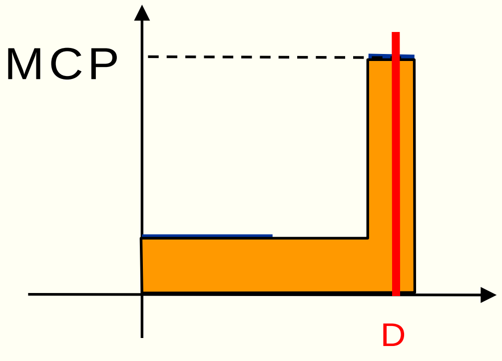

Issue 2 revisited: network topology

- Previous discussion applies to trading at a single location.
- If a transmission network is present:
  - Offers associated with specific network buses.
  - ISO minimizes cost subject to meeting demand, and capacity constraints (the power being dispatchable). Prices become node-dependent (Locational Marginal Prices, LMPs).
  - Integer constraints present?
  - Underlying power flow more complex than convex constraints for the Internet case. Some papers use DC power flow. Recent developments (Lavaei-Low) on convexified OPF may be relevant here.).
- Distributed solutions?
  - Do not appear so relevant in the ISO case.
  - Perhaps for auctions involving mutliple ISOs?

Issue 3 revisited: inter-temporal constraints

- In the power problem, coupling over time can appear from startup constraints: according to technology, some generator plants cannot easily be turned on and off.

In Yan et al. (IEEE Power '08) an additive startup cost considered.

 "Bid cost minimization", $\sum_i \underbrace{a_i}_{\text{offer in \$/MWh}} \cdot \underbrace{q_i}_{\text{quantity}} + \underbrace{S_i}_{\text{startup cost}}$ need not select cheapest a_i's.

If this is done, inconsistent to pay at the MCP (mkt clearing marginal price).

 "Payment cost minimization" $\sum_i MCP \cdot q_i + S_i$ can reduce cost paid out.

But, seems an artifice of the payment choice.

 Startup costs do not represent correctly the multiple period case.

 Block bids for multiple intervals, increase auction complexity.

 Perhaps a receding horizon policy can be useful for this purpose.

Conclusion

- Bandwidth allocation has been the focus of substantial academic research over the last decade, as a test case for distributed optimization and economic theory.
- In particular, we showed work on bandwidth auctions that attempts to put these models to practical use.
- However, the practical Internet rarely follows sophisticated pricing schemes: too much abundance of bandwidth?
- Given the “real” scarcity of energy, perhaps the power grid is a more adequate setting for exploiting these mathematical tools.