Selfish Behavior in Networks

FDIS 2005

Alex Fabrikant

Computer Science Division,

UC Berkeley

Backdrop

- The Internet: hundreds of millions of users, tens of thousands of autonomous systems
- By design, centralized control of the system is difficult:
  - Technical constraints: too much resources required to control the entire network's state
  - Political constraints: ISPs (and even countries) have their own preferences for how to interact with the network
- Hence, selfish decisions by the participants are both a desirable and a necessary consideration

Goals

- The grand motivation behind this line of research is twofold:

  1) Understand how to **Model** existing behavior: how has the current Internet been shaped by selfish interests, and what would it have been like if it weren't?

  2) Learn to **design** protocols under which the participants' selfish behavior leads to good properties of the resulting network.

- Conveniently enough, humans have been selfish since before ARPANET, so these problems aren't completely new

Game Theory

- ...is the study of **games**: systems with selfish "**players**" who each choose a strategy, and get a "**payoff**" which is a function of the **game state** (the combination of everyone's strategy)
- Prisoner's dilemma (classical example):

"Column player"

"Row player"

	C	D
C	1, 1	6, 0
D	0, 6	5, 5

Nash Equilibria

- (Pure) Nash equilibrium: a game state such that no player has motivation to change his strategy (only (D,D) below)
  - Simplest notion of a “stable” position
  - Not necessarily “good” for the players

“Column player” / “Row player”

		C	D
“Row player”	C	1,1	6,0
	D	0,6	5,5

- Pure Nash equilibria are not guaranteed to exist:

	A	B
A	1 , -1	-1 , 1
B	-1 , 1	1 , -1

- (Nash, 1950) But, if we allow each player a distribution over strategies ("**mixed strategy**"), the existence of a **mixed Nash Equilibrium** (no player can improve expectation) is guaranteed (here, pick randomly)

"Social Cost"

- To compare selfish behavior versus centralized control, we establish a **social cost** function, which, given a game state, evaluates how good the overall scenario is for all the players combined.
  - For prisoner's dilemma, the sum of the two payoffs is a reasonable candidate. Note that the optimum is much better than the only Nash Equilibrium!

	C	D
C	1 , 1	6 , 0
D	0 , 6	5 , 5

Two interesting metrics

- In the tradition of computer science, we want to analyze bounds on performance ratios
- Price of Anarchy (Koutsoupias & Papadimitriou '99):

$$\frac{max_{NE} S(NE)}{S(OPT)}$$

- "How bad is the worst equilibrium that the game can settle into if the players are left to their own devices?"

Two interesting metrics (cont)

- Price of Stability (Anshelevich, et al., '03, '04):

$$\frac{min_{NE}\, S(NE)}{S(OPT)}$$

- Assuming a central authority is there to *advise* players about a good equilibirum, how much better can you do than enforcing the global optimum?

Example problems: Network design games

- (Anshelevich, et al., '04) In a graph of nodes and possible links, each player has some pairs of nodes he wants to connect with paths
- Each possible edge has a cost, which is split evenly among all users using it; the user's cost is the sum of the costs of his edges
- Social cost = total price of edges bought
- For N users, the price of stability is O(log N)

Example problems: Network design games (2)

- (FLMPS'03) Players are nodes allowed to buy edges to any other nodes, with the goal of minimizing a weighted sum of:
  - average path lengths to all others in the network
  - cost of edges bought (the cost of an edge is a constant, α)
- Price of anarchy turns out to be somewhere between 5 and $\sqrt{\alpha}$ (bounds may have been recently tightened by Albers, et al)
- The social optimum is a star at low α, and a clique at high α

Example problems: Selfish Routing

- (Roughgarden '01-...) A network, with users wanting to route traffic; edge latency is a function of amount of traffic flowing through (also some old work dating to the '50s from road traffic engineering)
- Price of anarchy turns arbitrarily bad for most latency function families (even polynomials)
- On the other hand, if latency functions are all linear, price of anarchy is at most 4/3rds

Example problems: Selfish Routing (2)

- Braess's Paradox: adding an extra edge may worsen delays:

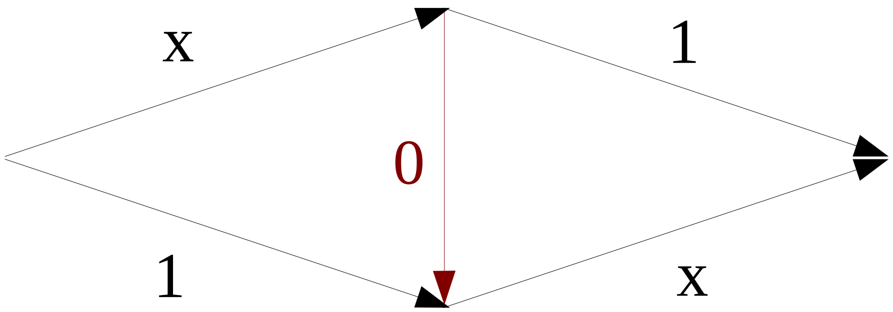

- (Lin, et al, '03) – Computationally intractable to detect whether this is a problem

Can we even find them?

- Computational applications of much of game theory is contingent on being able to compute, or at least approximate equilibria
- (Papadimitriou'94) Finding mixed NE is in an exotic class, PPA (not NP since solution always exists); no polytime alg known
- (von Stengel, et al'04) Best-known local search heuristic is guaranteed to have exponential worst-case
- (Conitzer&Sandholm'03) A plethora of slight tweaks to "find mixed NE" is NP-complete even in the simplest cases

Can we even find them? (2)

- For pure Nash Equilibrium, explicit payoff table is huge, but checking each entry is enough
- Succinct payoff functions? (F, Papadimitriou, Talwar '04) For congestion games (a broad class of games with guaranteed pure NE), finding one is complete for PLS (guarantees various nastiness, like worst-case exponential paths for local search)
- For many special types of games (e.g. symmetric network congestion), poly time.

There's plenty more...

- Dynamic behavior and convergence in games (recent work by Even-Dar and Mansour)
- (Karp, etc.) TCP/IP congestion control as a competitive policy
- Designing auctions for bandwidth (not unrelated to auctions for selling the junk in your attic)
- The big complexity questions are all still open (aren't they always...)

Conclusions

- There's a rich trove of problems involving modeling selfish behavior, which promises to give us a better understanding of what's happening on the net right now, and how to design better mechanisms for the future
- There's a complementary trove of problems about the tractability (and approximability) of the models that result
- We also hope to bring the latter results back to economists and have them reconsider their own models