Algorithmic Game Theory | Summer 2015, Week 1

Motivation

ETH Zürich | Paul Dütting

Many real-world problems have both a computational and a strategic aspect. Our overarching goal will be to provide the necessary tools and techniques to “solve” these problems. We begin by introducing three paradigmatic problems, and the type of questions that arise. We will return to each of these problems as we progress.

1 Example 1: Routing Traffic

The first example is that of routing traffic in a ntework, such as vehicles in a road network. Routing clearly has a computational aspect. Given the data—such as the number of vehicles, their departure times, start and end points—we could compute routes such that some overall objective is optimized. However, it also has a strategic aspect. Drivers in a road network, for example, are free to choose their routes and it seems reasonable to assume that they would just choose the route that promises the shortest travel time.

1.1 Braess’ Paradox

A famous example that illustrates the sometimes counter-intuitive implications of selfish behavior in road networks is known as Braess’ Paradox. It says that under selfish behavior adding a new road to a network may actually increase travel times.

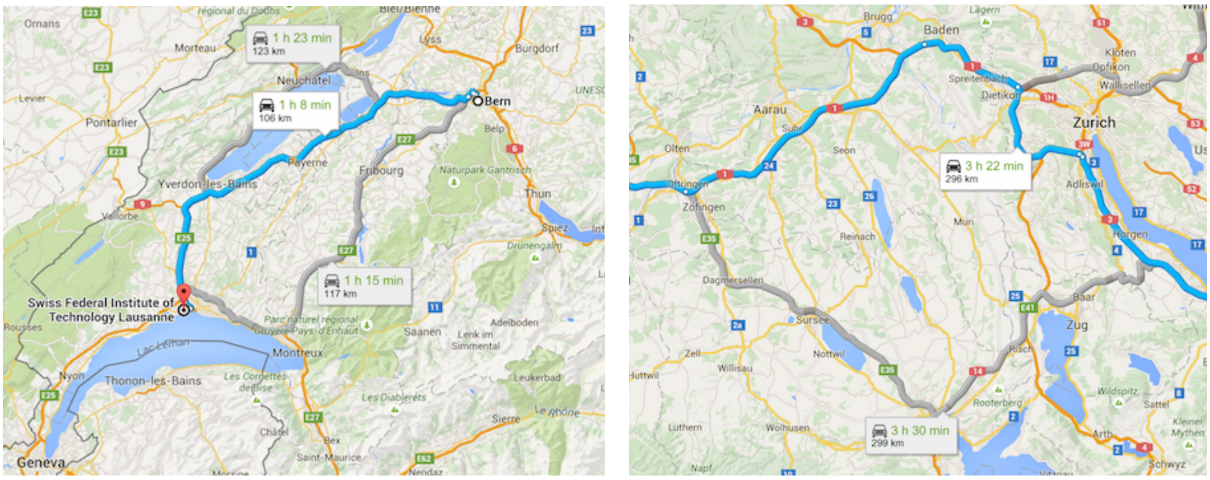

Figure 1: Road networks between Bern and Lausanne and between Bern and Zürich. (Source: `http://maps.google.ch/`, September 16, 2015.)

The paradox is known to occur in situations like the ones depicted in Figure 1. The first shows two alternative routes from Bern to Lausanne and the second shows two alternative routes from Bern to Zürich (in blue and grey, respectively). In both cases the two alternative routes take roughly the same time. An important difference between the two, however, is that in the latter there is a fast cross-link that connects the two routes.

To see the paradox at work, let’s abstract the two examples as depicted in Figure 2. Suppose there is a source node s and a target node t and that there are so many drivers that want to

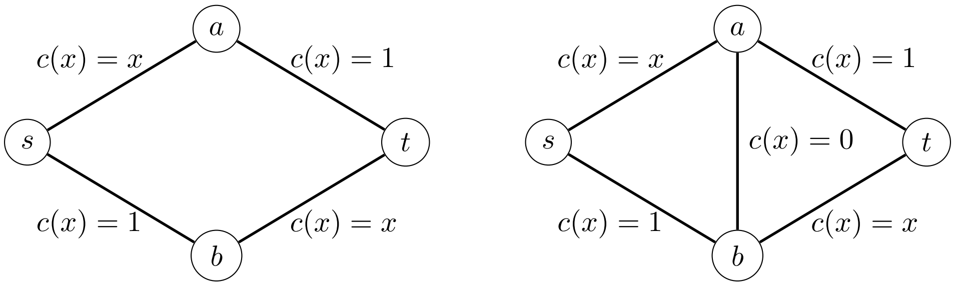

Figure 2: Road network exhibiting Braess' paradox.

travel from s to t so that we can treat the entirety of drivers as one infinitely divisible unit of traffic. Our global objective is to minimize the maximum travel time of any driver, while the drivers themselves want to minimize their individual travel times. Suppose further that there are two possible routes from s to t. One via node a and one via node b. The travel times from s to a and from b to t depend on the fraction of drivers that take this route. Formally, on both routes, the travel time $c(x)$ as a function of the fraction of drivers x taking this route is $c(x) = x$. The travel times on the other two links, from s to b and from a to t, are $c(x) = 1$ independent of how many drivers take this route.

What is the global optimum in this situation? Intuitively, because of symmetry, we want to split the traffic in two halves and route half of the traffic via a and the other half via b. In fact, this solution, which results in travel times of $1 + c(1/2) = 3/2$ hours for each driver indeed minimizes the maximum travel time of any driver. Moreover, this globally optimal solution is also stable or at equilibrium as we will say. No individual driver can lower his travel time by switching from one route to the other.

Now suppose we add a new link between a and b with zero travel time. What does this change? With this new link the zig-zag connection from s to t via a and b looks more attractive. It promises a travel time of $1/2 + 1/2 = 1$ hour. So any rational driver would switch. But what if all of them switch? The individual travel times will go up to $1 + 1 = 2$ hours!

Did our reasoning fail? Is it unreasonable to assume that all drivers switch and have to live with longer travel times? It turns out that even if all drivers switch and take the zig-zag connection, no individual driver can achieve a shorter travel time by switching to an alternate route. So this solution is again stable or at equilibrium. In fact, in this modified network, choosing the zig-zag route is what game-theorists call a dominant strategy. The zig-zag route promises the shortest travel time, no matter what the other drivers do. Moreover, all drivers taking the zig-zag route is the only equilibrium in this network.

We conclude that adding a new road to a network can actually increase travel times when drivers choose routes selfishly. Another corollary of this example is that selfish behavior can lead to suboptimal solutions.

1.2 Outlook

The above examples are instances of what is called a congestion game in the literature. Congestion games will be our running example for the first few lectures. We will use them to introduce basic notions from game theory and to address questions such as: Do natural networks, in which entities behave selfishly, have stable states? How are these stable states reached? Can we compute them efficiently? Does natural game playing dynamic lead to them? How bad can stable states be compared to the global optimum?

2 Example 2: Sponsored Search Auctions

Ever wondered with what search engines, such as Google, earn so much money? Right, it is the ads—or sponsored search results—that you see alongside the organic search results on the results page of a web search. If you were working for a search engine, how would you sell the ads? It turns out that this is as much an engineering as an economic problem, and that things can (and did) go horribly wrong initially.

2.1 Sawtooth Bidding

Starting in 1994, internet ads were primarily sold via contracts. Advertisers could buy a certain number of so-called ad impressions. This changed with a company called Overture in 1997. The basic idea behind Overture's system was to let advertisers bid on keywords and to then run an auction to determine which ads are shown and in which order, and to only charge advertisers when someone clicks on their ad.

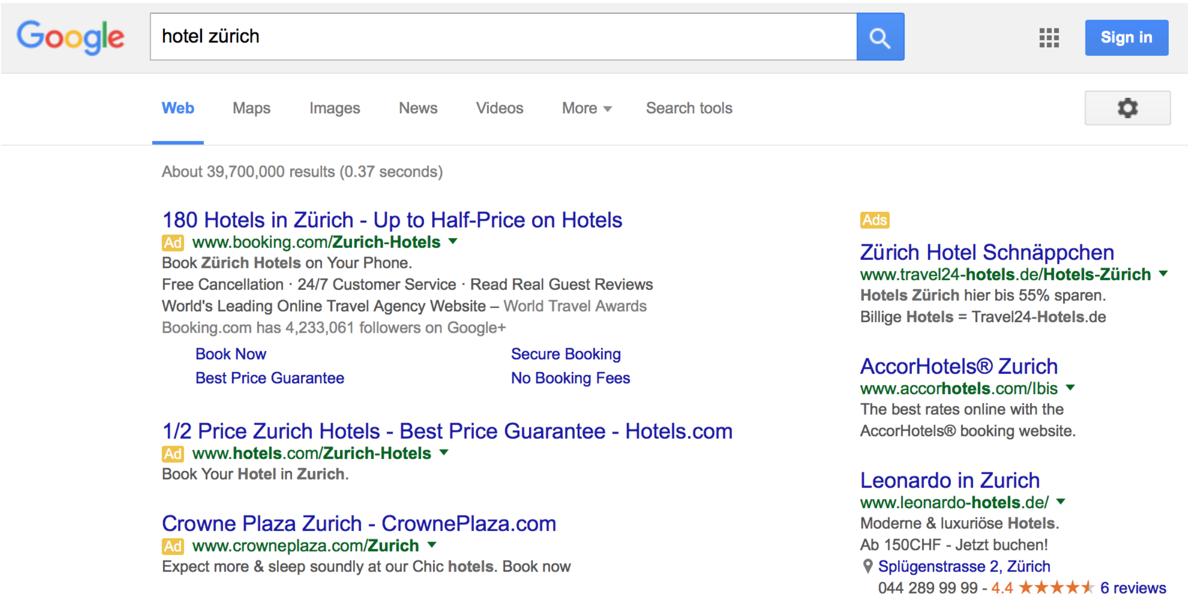

Figure 3: Google's search results for keyword "hotel zürich". (Source: `http://www.google.ch/`, September 16, 2015.)

A simple model is as follows: There are n advertisers N. Advertiser $i \in N$ has a value of v_i per click. This value is not known to the search engine. There are k slots in which an ad can be shown. The ads are arranged in a vertical list with slot 1 at the top and slot k at the bottom. Slots towards the top receive, more clicks than slots toward the bottom. Let $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$ denote the number of clicks received by slots 1 through k.

The original auction design was to ask each advertiser $i \in N$ for a bid b_i. To sort the advertisers by non-increasing bid, so that $b_1 \geq b_2 \geq \cdots \geq b_n$, and to then (a) assign advertiser $i \in \{1, \ldots, k\}$ to slot i and (b) charge each advertiser $i \in \{1, \ldots, k\}$ his bid b_i whenever someone clicks on his ad. So overall, for bids $b = (b_1, \ldots, b_n)$, the utility of advertiser $i \leq k$ is $u_i(b) = \alpha_i \cdot (v_i - b_i)$ and for advertiser $i > k$ it is $u_i(b) = 0$.

Is this a good system from an engineering perspective? It definitely has the advantage that it is automated, and it also has the advantage that it lets advertisers target a specific keyword and ad slot. However, game-theoretic reasoning quickly exhibits a potential flaw. The flaw is

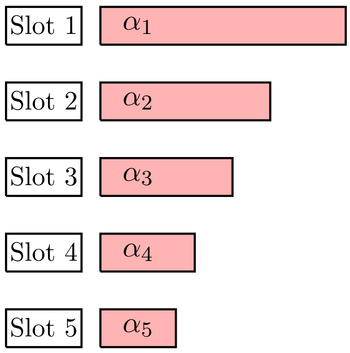

Figure 4: A basic model for sponsored search auctions.

that this auction format may not have an equilibrium and so advertisers, especially those using automated bidding tools, will keep updating their bid!

For concreteness suppose there are two slots and three advertisers. Let $\alpha_1 = 200$, $\alpha_2 = 100$ and let $v_1 = 20$, $v_2 = 14$, $v_3 = 6$. Suppose advertiser 2 bids 6.01, to guarantee that he gets a slot. Then advertiser 1 will bid 6.02, i.e., just above advertiser 2, to get slot 1. But then advertiser 2 will find it more beneficial to raise his bid to 6.03 so as to outbid advertiser 1. Advertiser 1 in turn will raise his bid to 6.04, and so on. But notice that as soon as this process would require advertiser 2 to bid $b_2 > 10$, he will find it more beneficial to fall back to his initial bid $b_2 = 6.01$ and rather take the second slot. However, as a result, advertiser 1 can lower his bid to $b_1 = 6.02$ and still win the first slot, and the process repeats.

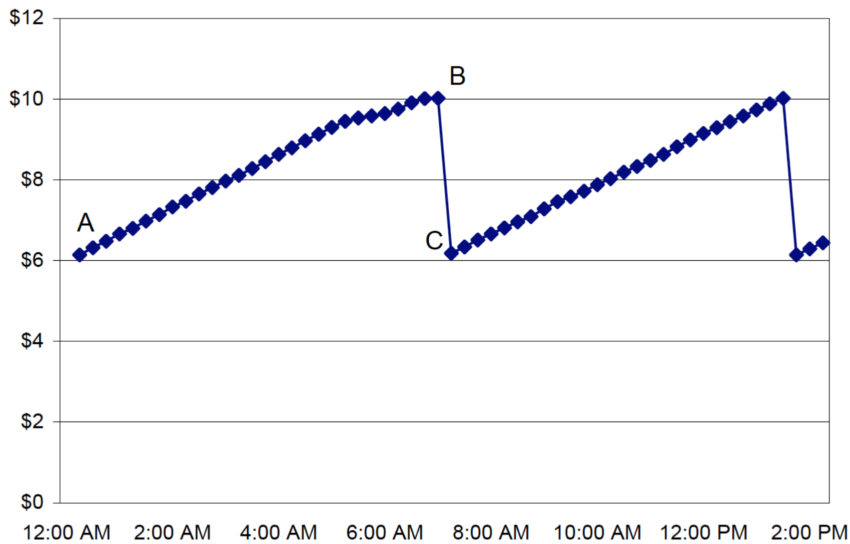

Figure 5: Characteristic bidding patterns as observed in early ad auctions. (Source: B. Edelman and M. Ostrovsky. Strategic Bidding Behavior in Sponsored Search Auctions. Decision Support Systems, 43:192–198, 2007.)

While this obviously is an undesirable situation from the engineering perspective (just imagine the amount of traffic needed for the bid updates), it also has undesirable economic effects such as: (a) It will typically not be the case that the bidders with the highest values win the highest slots. (b) The revenue achieved by such a system will typically be lower than the revenue that one could achieve with a different auction format.

2.2 Outlook

The above example nicely illustrates that the "rules of the game matter". We will use auctions for a single item and their extension to ad auctions as an introduction to a field called Mechanism Design, which is concerned with designing the rules of a game so that the equilibria of the induced game satisfy the designer's goals. In particular, we will see how to "fix" the above problem by changing the rules of the auction.

3 Example 3: Spectrum Auctions

As a third and final example consider spectrum auctions such as the upcoming FCC Incentive Auction, which is scheduled for early 2016. The goal of this auction, which was initiated by the U.S. Federal Communication Commission (FCC) in 2012, is to free some of the electromagnetic spectrum, by buying back licences from TV broadcasters, and to then sell the freed spectrum to mobile internet providers.

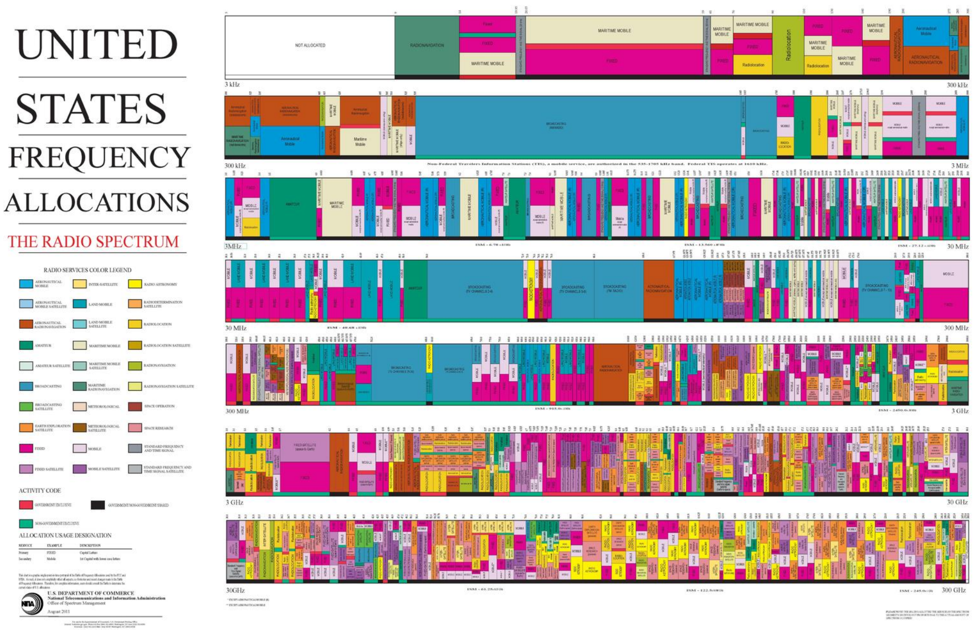

Figure 6: Chart showing the allocations of the electromagnetic spectrum in the United States. (Source: U.S. Department of Commerce, August 2011.)

3.1 Computation vs. Incentives

What makes spectrum auctions special is that they are what one calls a combinatorial auction: There is a collection of items, and bidders want to acquire subsets of items. These auctions are challenging for at least two reasons: (a) Typically, the underlying optimization problem is computationally intractable. (b) They are often used in highly strategic settings and individual bidders or groups of bidders may try to game the system.

As a concrete example for the computational intractability consider a so-called single-minded combinatorial auction. There is a set of n bidders N and a set of m items M. Each bidder $i \in N$ is interested in a subset of the items $S_i \subseteq M$. Bidder $i \in N$'s value v_i for a set of items T is $v_i(T) = v_i$ if $T \supseteq S$ and it is $v_i(T) = 0$ otherwise. We say that an allocation $X = (X_1, \ldots, X_n)$, which assigns the set of items $X_i \subseteq M$ to bidder $i \in N$, is feasible if $X_1 \cup \cdots \cup X_i \subseteq M$ and $X_i \cap X_j = \emptyset$ for all $i \neq j$ such that $i, j \in N$.

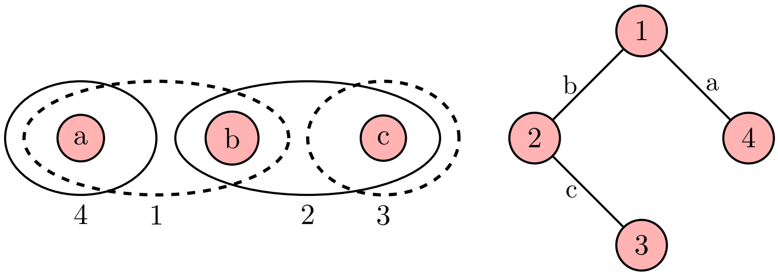

Figure 7: Single-minded CA instance (on the left) and corresponding Maximum Independent Set instance (on the right). There are 4 bidders $N = \{1, 2, 3, 4\}$ and three items $M = \{a, b, c\}$. The "blobs" on the left indicate which items the respective bidders are interested in; an allocation is feasible if no two blobs intersect.

It is not difficult to see that finding the allocation that maximizes total value $\sum_{i \in N} v_i(X_i)$ is $\mathcal{NP}$-hard. We will reduce from Maximum Independent Set: We are given a graph $G = (V, E)$ and we need to find a set of vertices $W \subseteq V$ of maximal cardinality such that for no two vertices $w_1, w_2 \in W$, $(w_1, w_2) \in E$. We identify each vertex $v \in V$ with a bidder and each edge $e \in E$ with an item. A bidder is interested in the items that correspond to edges that are incident to his vertex, and has a value of one for them. Then an allocation with maximal total value corresponds to a maximum independent set.

That bidders in combinatorial auctions in general, and in spectrum auctions in particular, have gamed the system in the past is well-documented. One example of collusive bidding behavior are so-called code bids. A code bid makes use of the fact that in spectrum auctions bids on all licences, except for a few exemptions, are six or more digits. So bidders can use the less significant digits to signal other bidders. For example, they can encode their companies acronyms such as GTE by ending their bids in "483", which spells GTE on the telephone keypad. They can then use this bid to indicate which license they are interested in and/or use it to tag a retaliation bid on a license a direct competitor has showed interest in.

Fixing a specific combinatorial setting, what is the best possible approximation that we can get? What is the best possible approximation that we can get if we also want that it is in each bidder's interest to report his value truthfully? And, last but not least, can we also ensure that no group of bidders has an incentive to misreport their values? How much do we lose with this additional restriction?

3.2 Outlook

We will study the tradeoff between optimization and incentives not only for single-minded combinatorial auctions, but also for other optimization problems. On the way we will answer the above questions, and develop techniques for designing truthful and non-truthful mechanisms with (near-)optimal performance.

Recommended Literature

- D. C. Parkes and Sven Seuken. Economics and Computation, Chapter 1. Unpublished manuscript, September 2015. (General introduction)

- B. Edelman, M. Ostrovsky, M. Schwarz. Selling Billions of Dollars worth of Keywords: The Generalized Second-Price Auction. American Economic Review, 97(1): 242-259, 2007. (Accessible article on sponsored search auctions)
- D. Lehmann, L. O'Callaghan, Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. Journal of the ACM, 49(5): 577–602, 2002. (Accessible article on combinatorial auctions)
- G. Kolata. What if they Closed 42d Street and Nobody Noticed? New York Times, December 25, 1990. Available from: `http://www.nytimes.com/1990/12/25/health/what-if-they-closed-42d-street-and-nobody-noticed.html`. (News article reporting on Braess' Paradox occuring in NYC)
- B. Edelman and M. Ostrovsky. Strategic Bidder Behavior in Sponsored Search Auctions. Decision Support Systems, 43:192-198, 2007. (Article about sawtooth bidding)
- P. Cramton and J. Schwartz. Collusive Bidding in the FCC Spectrum Auctions. Economic Analysis & Policy, 1(1):Article 11, 2002. (Article about code bids)