

CS364A: Algorithmic Game Theory

Lecture #1: Introduction and Examples*

Tim Roughgarden[†]

September 23, 2013

1 Mechanism Design: The Science of Rule-Making

This course is roughly organized into three parts, each with its own overarching goal. Here is the first.

Course Goal 1 Understand how to design systems with strategic participants that have good performance guarantees.

We begin with a cautionary tale. In 2012, the Olympics were held in London. One of the biggest scandals of the event concerned, of all sports, women's badminton. The scandal did not involve any failed drug tests, but rather a failed tournament design that did not carefully consider *incentives*.

The tournament design that was used is familiar from World Cup soccer. There are four groups (A,B,C,D) of four teams each. The tournament has two phases. In the first "round-robin" phase, each team plays the other three teams in its group, and does not play teams in other groups. The top two teams from each group advance to the second phase, the bottom two teams from each group are eliminated. In the second phase, the remaining eight teams play a standard "knockout" tournament (as in tennis, for example): there are four quarterfinals (with the losers eliminated), then two semifinals (with the losers playing an extra match to decide the bronze medal), and then the final (the winner gets the gold, the loser the silver).

The incentives of participants and of the Olympics committee (and fans) are not necessarily aligned in such a tournament. What does a team want? To get as good a medal as possible, of course. What does the Olympics committee want? They didn't seem to think carefully about this question, but in hindsight it's clear that they want every team to try their best to win every match. Why, you ask, would a team ever want to lose a match?

*©2013, Tim Roughgarden.

[†]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

Indeed, in the second "knockout" phase of the tournament, where losing leads to instant elimination, it's obvious that winning is better than losing.

To understand the incentive issues, we need to explain how the eight winners from the round-robin phase are paired up in the quarterfinals. The team with the best record from group A plays the second-best team from group C in the first quarterfinal; similarly with the best team from group C and the second-best time from group A in the third quarterfinal; and similarly with the top two teams from groups B and D in the second and fourth quarterfinals. The dominoes started to fall when, on the last day of round-robin competition, there was a shocking upset: the Danish team of Pedersen and Juhl (PJ) beat the Chinese team of Qing and Wunlei (QW), and as a result PJ won group D with QW coming in second (so both teams advanced to the knockout round).

The first controversial match involved another team from China, Xiaoli and Yang (XY), and the South Korean team of Kyung-eun and Ha-na (KH). Both teams had a 2-0 record in group A play. Thus, both were headed for the knockout stage, with the winner and loser of this match the top and second-best team from the group, respectively. Here was the issue: the group-A winner would meet the fearsome QW team (which came in only second in group D) in the semifinals of the knockout tournament (where a loss means a bronze medal at best), while the the second-best team in group-A would not have to face QW until the final (where a silver medal is guaranteed). Both the XY and KH teams found the difference between these two scenarios significant enough to try to *deliberately lose the match*.¹ This unappealing spectacle led to scandal, derision, and, ultimately, the disqualification of the XY and KH teams, as well as two group C teams (from Indonesia and another team from South Korea) that used the same strategy, for the same reason.²

The point is that, in systems with strategic participants, *the rules matter*. Poorly designed systems suffer from unexpected and undesirable results. The burden lies on the system designer to anticipate strategic behavior, not on the participants to behave against their own interests. Can we blame the badminton players for acting to maximize their medal placement? To quote Hartline and Kleinberg [5]:

"The next time we bemoan people exploiting loopholes to subvert the intent of the rule makers, instead of asking 'What's wrong with these people?'" let's instead ask, 'What's wrong with the rules?' and then adopt a scientifically principled approach to fixing them."

There is a well-developed science of rule-making, the field of *mechanism design*. The goal in this field is to design rules so that strategic behavior by participants leads to a desirable outcome. Killer applications of mechanism design, which we will discuss in detail later, include Internet search auctions, wireless spectrum auctions, matching medical residents to hospitals, matching children to schools, and kidney exchange markets.

¹That the teams feared the Chinese team QW far more than the Danish team PJ seems justified in hindsight: PJ were knocked out in the quarterfinals, while QW won the gold medal.

²If you're having trouble imagining what a badminton match looks like when both teams are trying to lose, I encourage you to track down the video on YouTube.




Figure 1: Braess's Paradox. The addition of an intuitively helpful edge can adversely affect all of the traffic.

We'll study some of the basics of the traditional economic approach to mechanism design, along with several complementary contributions from computer science that focus on computational efficiency, approximate optimality, and robust guarantees.

2 The Price of Anarchy: When Is Selfish Behavior Near-Optimal?

Sometimes you don't have the luxury of designing the rules of a game from scratch, and want instead to understand a game that occurs "in the wild" — the Internet or a road network, for example.

Course Goal 2 When is selfish behavior essentially benign?

2.1 Braess's Paradox

For a motivating example, consider *Braess's Paradox* (Figure 1) [1]. There is a suburb s , a train station t , and a fixed number of drivers who wish to commute from s to t . For the moment, assume two non-interfering routes from s to t , each comprising one long wide road (with travel time one hour, no matter how much traffic uses it) and one short narrow road (with travel time in hours equal to the fraction of traffic using it) as shown in Figure 1(a). The combined travel time in hours of the two edges on one of these routes is $1 + x$, where x is the fraction of the traffic that uses the route. The routes are therefore identical, and traffic should split evenly between them. In this case, all drivers arrive at their destination 90 minutes after their departure from s .

Now, suppose we install a teleportation device allowing drivers to travel instantly from v to w . The new network is shown in Figure 1(b), with the teleporter represented by edge

(v, w) with constant cost $c(x) = 0$, independent of the road congestion. How will the drivers react?

We cannot expect the previous traffic pattern to persist in the new network. The travel time along the new route $s \rightarrow v \rightarrow w \rightarrow t$ is never worse than that along the two original paths, and it is strictly less whenever some traffic fails to use it. We therefore expect all drivers to deviate to the new route. Because of the ensuing heavy congestion on the edges (s, v) and (w, t) , all of these drivers now experience two hours of travel time when driving from s to t . Braess's Paradox thus shows that the intuitively helpful action of adding a new zero-cost link can negatively impact *all* of the traffic!

Braess's Paradox shows that selfish routing does not minimize the commute time of the drivers — in the network with the teleportation device, an altruistic dictator could dictate routes to traffic in improve everyone's commute time by 25%. We define the *price of anarchy (POA)* to be the ratio between the system performance with strategic players and the best-possible system performance — for the network in Figure 1(b), the ratio between 2 and $\frac{3}{2}$ (i.e., $\frac{4}{3}$).

The POA was first defined and studied by computer scientists. Every economist and game theorist knows that equilibria are generally inefficient, but until the 21st century there had been almost no attempts to quantify such inefficiency in different application domains.

In our study of the POA, the overarching goal is to identify application domains and conditions under which the POA is guaranteed to be close to 1, and thus selfish behavior leads to a near-optimal outcome. Killer applications include network routing, scheduling, resource allocation, and simple auction designs. For example, modest overprovision of network capacity guarantees that the POA of selfish routing is close to 1 [7].

2.2 Strings and Springs

As a final aside, we note that selfish routing is also relevant in systems that have no explicit notion of traffic whatsoever. Cohen and Horowitz [3] gave the following analogue of Braess's Paradox in a mechanical network of strings and springs.

In the device pictured in Figure 2, one end of a spring is attached to a fixed support, and the other end to a string. A second identical spring is hung from the free end of the string and carries a heavy weight. Finally, strings are connected, with some slack, from the support to the upper end of the second spring and from the lower end of the first spring to the weight. Assuming that the springs are ideally elastic, the stretched length of a spring is a linear function of the force applied to it. We can therefore view the network of strings and springs as a traffic network, where force corresponds to traffic and physical distance corresponds to cost.

With a suitable choice of string and spring lengths and spring constants, the equilibrium position of this mechanical network is described by Figure 2(a). Perhaps unbelievably, severing the taut string causes the weight to *rise*, as shown in Figure 2(b)! An explanation for this curiosity is as follows. Initially, the two springs are connected in series, and each bears the full weight and is stretched out to great length. After cutting the taut string, the two springs are only connected in parallel. Each spring then carries only half of the weight, and




Figure 2: Strings and springs. Severing a taut string lifts a heavy weight.

accordingly is stretched to only half of its previous length. The rise in the weight is the same as the improvement in the selfish outcome obtained by deleting the zero-cost edge of Figure 1(b) to obtain the network of Figure 1(a).

This construction is not merely theoretical; on YouTube you can find several physical demonstrations of Braess’s Paradox that were performed (for extra credit) by past students of CS364A.

3 Complexity of Equilibria: How Do Strategic Players Learn?

Some games are easy to play. For example, in the second network of Braess’s Paradox (Figure 1(b)), using the teleporter is a no-brainer for every individual — it is the best route, no matter what other drivers do. In many other games, like the first-price auctions mentioned in the next lecture, it’s much harder to figure out how to play.

Course Goal 3 How do strategic players reach an equilibrium? (Or do they?)

Informally, an equilibrium is a “steady state” of a system where each participant, assuming everything else stays the same, want to remain as-is. Hopefully, you didn’t learn the definition of a Nash equilibrium from the movie *A Beautiful Mind*.

In most games, the best action to play depends on what the other players are doing. Rock-Paper-Scissors, rendered below in “bimatrix” form, is a canonical example.

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

One player chooses a row and the other a column. The numbers in the corresponding matrix entry are the payoffs for the row and column player, respectively.

There is certainly no “deterministic equilibrium” in the Rock-Paper-Scissors game: whatever the current state, at least one player would benefit from a unilateral move.

A crucial idea, developed largely by von Neumann, is to allow *randomized* (a.k.a. *mixed*) strategies. (In a game like Rock-Paper-Scissors, from your perspective, your opponent is effectively randomizing.) If both players randomize uniformly in Rock-Paper-Scissors, then neither player can increase their expected payoff via a unilateral deviation — indeed, every unilateral deviation yields zero expected payoff to the deviator. A pair of probability distributions with this property is a (*mixed-strategy*) *Nash equilibrium*.

Remarkably, with randomization, *every* game has at least one Nash equilibrium.

Nash’s Theorem (’51): Every bimatrix game has a Nash equilibrium.

Nash’s theorem holds more generally in games with any finite number of players.

Another piece of good news, that we’ll cover in the course: if a bimatrix game is *zero-sum* — meaning that the payoff pair in each entry sums to zero, like in Rock-Paper-Scissors — then a Nash equilibrium can be computed in polynomial time. This can be done via linear programming or, if a small amount of error can be tolerated, via simple iterative learning algorithms. These algorithmic results give credence to the Nash equilibrium concept as a good prediction of behavior in zero-sum games.

Far more recently (mid-last decade), computer science contributed an important negative result: under suitable complexity assumptions, there is *no* polynomial-time for computing a Nash equilibrium in general (non-zero-sum) games [2, 4]. While the problem is not *NP*-hard (unless $NP = coNP$), it is something called “PPAD-hard”, which we will explain and interpret in due course.

This hardness result is interesting for at least two different reasons. On a technical level, it shows that computing Nash equilibria is a fundamental computational problem of “intermediate” difficulty (like factoring and graph isomorphism) — unlikely to be in *P* or *NP*-complete. On a conceptual level, many interpretations of an equilibrium concept involve someone — the participants or a designer — determining an equilibrium. For example, the idea that markets implicitly compute a solution to a significant computational problem goes back at least to Adam Smith’s notion of the invisible hand [8]. If all parties are boundedly rational, then an equilibrium can be interpreted as a credible prediction only if it can be computed with reasonable effort. Rigorous intractability results thus cast doubt on the predictive power of equilibrium concepts (a critique that dates back at least to Rabin [6]). While intractability is certainly not the first stone thrown at the Nash equilibrium concept, it

is one that theoretical computer science is ideally situated to make precise. This perspective also provides novel motivation for our study of “easier” equilibrium concepts, like correlated equilibria and coarse correlated equilibria.

4 What Computer Science Brings to the Table

There is, of course, a very rich literature on mechanism design and equilibria in economics. In this course, we’ll see how computer scientists have complemented this literature in a number of novel ways. The traditional economics approach to the topics we study tends to focus on “Bayesian” (i.e., average-case) analyses; emphasizes exact solutions and characterizations; and usually ignores computational issues. Computer science has offered a focus on and a language to discuss computational complexity; popularized the widespread use of approximation bounds to reason about models where exact solutions are unrealistic or unknowable; and encouraged more robust performance guarantees.

5 Target Audience

These notes assume a background in undergraduate theoretical computer science — basic algorithms and NP-completeness. They do not assume any background in game theory or economics. (Conversely, this course is no substitute for a traditional game theory or microeconomics economics.) The level is meant to be accessible to a Masters or 1st-year PhD student with an affinity for theory.

References

- [1] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.
- [2] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- [3] J. E. Cohen and P. Horowitz. Paradoxical behavior of mechanical and electrical networks. *Nature*, 352(8):699–701, 1991.
- [4] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [5] J. D. Hartline and R. D. Kleinberg. Badminton and the science of rule making. *The Huffington Post*, August 13 2012. http://www.huffingtonpost.com/jason-hartline/badminton-and-the-science-of-rule-making_b.1773988.html.

- [6] M. O. Rabin. Effective computability of winning strategies. In M. Drescher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory Games*, volume 3, pages 147–157. Princeton, 1957.
- [7] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [8] A. Smith. *An Inquiry Into the Nature and Causes of the Wealth of Nations*. Methuen, 1776.