Chapter 19

Equilibria and algorithms

Economic and game-theoretic reasoning —specifically, how agents respond to economic incentives as well as to each other's actions– has become increasingly important in algorithm design. Examples: (a) Protocols for networking have to allow for sharing of network resources among users, companies etc., who may be mutually cooperating or competing. (b) Algorithm design at Google, Facebook, Netflix etc.—what ads to show, which things to recommend to users, etc.—not only has to be done using objective functions related to economics, but also with an eye to how users and customers *change* their behavior in response to the algorithms and to each other.

Algorithm design mindful of economic incentives and strategic behavior is studied in a new field called *Algorithmic Game Theory.* (See the book by Nisan et al., or many excellent lecture notes on the web.)

Last lecture we encountered zero sum games, a simple setting. Today we consider more general games.

19.1 Nonzero sum games and Nash equilibria

Recall that a 2-player game is *zero sum* if the amount won by one player is the same as the amount lost by the other. Today we relax this. Thus if player 1 has n possible actions and player 2 has m, then specifying the game requires two a $n \times m$ matrices A, B such that when they play actions i, j respectively then the first player wins A_{ij} and the second wins B_{ij}. (For zero sum games, $A_{ij} = -B_{ij}$.)

A Nash equilibrium is defined similarly to the equilibrium we discussed for zero sum games: a pair of strategies, one for each player, such that each is the optimal response to the other. In other words, if they both announce their strategies, neither has an incentive to deviate from his/her announced strategy. The equilibrium is *pure* if the strategy consists of deterministically playing a single action.

Example 44 (Prisoners' Dilemma) This is a classic example that people in myriad disciplines have discussed for over six decades. Two people suspected of having committed a crime have been picked up by the police. In line with usual practice, they have been placed in separate cells and offered the standard deal: help with the investigation, and you'll be treated with leniency. How should each prisoner respond: Cooperate (i.e., stick to the story he and his accomplice decided upon in advance), or Defect (rat on his accomplice and get a reduced term)?

Let's describe their incentives as a 2×2 matrix, where the first entry describes payoff for the player whose actions determine the row. If they both cooperate, the police can't

	Cooperate	Defect
Cooperate	3, 3	0, 4
Defect	4, 0	1, 1

prove much and they get off with fairly light sentences after which they can enjoy their loot (payoff of 3). If one defects and the other cooperates, then the defector goes scot free and has a high payoff of 4 whereas the other one has a payoff of 0 (long prison term, plus anger at his accomplice).

The only pure Nash equilibrium is (Defect, Defect), with both receiving payoff 1. In every other scenario, the player who's cooperating can improve his payoff by switching to Defect. This is much worse for both of them than if they play (Cooperate, Cooperate), which is also the social optimum —where the sum of their payoffs is highest at 6—is to cooperate. Thus in particular the social optimum solution is not a Nash equilibrium. ((OK, we are talking about criminals here so maybe social optimum is (Defect, Defect) after all. But read on.)

One can imagine other games with similar payoff structure. For instance, two companies in a small town deciding whether to be polluters or to go green. Going green requires investment of money and effort. If one does it and the other doesn't, then the one who is doing it has incentive to also become a polluter. Or, consider two people sharing an office. Being organized and neat takes effort, and if both do it, then the office is neat and both are fairly happy. If one is a slob and the other is neat, then the neat person has an incentive to become a slob (saves a lot of effort, and the end result is not much worse).

Such games are actually ubiquitous if you think about it, and it is a miracle that humans (and animals) cooperate as much as they do. Social scientists have long pondered how to cope with this paradox. For instance, how can one change the game definition (e.g. a wise governing body changes the payoff structure via fines or incentives) so that cooperating with each other —the socially optimal solution—becomes a Nash equilibrium? The game can also be studied via the *repeated game* interpretation, whereby people realize that they participate in repeated games through their lives, and playing nice may well be a Nash equilibrium in that setting. As you can imagine, many books have been written. □

EXAMPLE 45 (CHICKEN) This dangerous game was supposedly popular among bored teenagers in American towns in the 1950s (as per some classic movies). Two kids would drive their cars at high speed towards each other on a collision course. The one who swerved away first to avoid a collision was the "chicken."How should we assign payoffs in this game? Each player has two possible actions, *Chicken* or *Dare*. If both play Dare, they wreck their cars and risk injury or death. Lets call this a payoff of 0 to each. If both go Chicken, they both live and have not lost face, so let's call it a payoff of 5 for each. But if one goes Chicken and the other goes Dare, then the one who went Dare looks like the tough one (and presumably attracts more dates), whereas the Chicken is better of being alive than dead but lives in shame. So we get the payoff table:

	Chicken	Dare
Chicken	5, 5	1, 6
Dare	6, 1	0, 0

This has two pure Nash equilibria: (Dare, Chicken) and (Dare, Chicken). We may think of this as representing two types of behavior: the reckless type may play Dare and the careful type may play Chicken.

Note that the socially optimal solution —both players play chicken, which maximises their total payoff—is not a Nash equilibrium.

Many games do not have any pure Nash equilibrium. Nash's great insight during his grad school years in Princeton was to consider what happens if we allow players to play a *mixed* strategy, which is a probability distribution over actions. An equilibrium now is a pair of mixed strategies x, y such that each strategy is the optimum response (in terms of maximising expected payoff) to the other.

Theorem 32 (Nash 1950)
For every pair of payoff matrices A, B there exists a mixed equilibrium.

(In fact, Wilson's theorem from 1971 says that for random matrices A, B, the number of equilbria is odd with high probability.)

Unfortunately, Nash's proof doesn't yield an efficient algorithm for computing an equilibrium: when the number of possible actions is n, computation may require $\exp(n)$ time. Recent work has shown that this may be inherent: computing Nash equilibria is PPAD-complete (Chen and Deng'06).

The Chicken game has a mixed equilibrium: play each of Chicken and Dare with probability 1/2. This has expected payoff $\frac{1}{4}(5+1+6+0) = 3$ for each, and a simple calculation shows that neither can improve his payoff against the other by changing to a different strategy.

19.2 Multiplayer games and Bandwidth Sharing

One can define multiplayer games and equilibria analogously to single player games. One can also define games where each player's set of moves comes from a continuous set like the interval $[0, 1]$. Now we do this in a simple setting: multiple users sharing a single link of fixed bandwidth, say 1 unit. They have different utilities for internet speed, and different budgets. Hence the owner of the link can try to allocate bandwidth using a game-theoretic view, which we study using a game introduced by Frank Kelly.

The Setting: There are n users. If user i gets x units of bandwidth by paying w dollars, his/her utility is $U_i(x) - w$, where the *utility* function U_i is nonnegative, increasing, *concave*[1]

[1] Concavity implies that the going from 0 units to 1 brings more happiness than going from 1 to 2, which in turn brings more happiness than going from 2 to 3. For twice differentiable functions, concavity means the second derivative is negative.

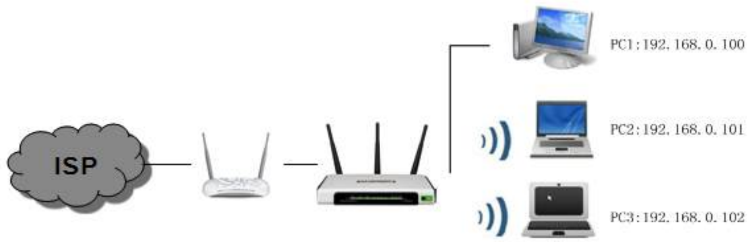

Figure 19.1: Sharing a fixed bandwidth link among many users

and twice-differentiable. If a unit of bandwidth is priced at p, this utility describes the amount of bandwidth desired by a utility-maximizing user: the ith user demands x_i that maximises $U_i(x_i) - px_i$. This maximum can be computed by calculus.

Social Optimum: First, let us consider how an infinitely wise and all-knowing *social planner* would go about solving this problem, assuming all users pay the same price per unit of bandwidth. The planner would find a socially optimal per unit price p^* and bandwidth allocations $x_1, x_2, \ldots, x_n \geq 0$ satisfying $\sum_i x_i = 1$ so as to maximise

$$\sum_i (U_i(x_i) - p^* x_i),$$

where $U_i(x_i) - p^* x_i$ is the utility of user i less the money he/she paid.

But this means that x_i is the allocation that maximises $U_i(x_i) - p^* x_i$, which is the unique x_i that satisfies $U_i'(x_i) = p^*$. (This x_i is unique since U_i' is a decreasing function, as follows from the fact that U_i" is negative.) Now let's see that the planner can find x_i's such that they sum to 1. This follows from realising that $\sum_i x_i$ is a decreasing continuous function of p^*, taking the value $+\infty$ if $p^* = 0$ and 0 if $p^* = +\infty$. So by the mean value theorem there exists a unique p^* such that $\sum_i x_i = 1$. This is the social optimum.

But the social optimum involves an all-knowing planner. Can we have a market mechanism that finds a solution almost as good?

Kelly's game: Each user i submits a bid where he/she offers to pay a sum of w_i. The link owner then allocates $w_i / \sum_j w_j$ portion of the bandwidth to user i. Thus the entire bandwidth is used up and the effective price for the entire bandwidth is $\sum_j w_j$.

The Nash Equilibrium. What n-tuple of strategies $w_1, w_2, \ldots, w_n$ is a Nash equilibrium? Note that this n-tuple implies a per unit price p of $\sum_j w_j$, and for each i his received amount is optimal at this price if w_i is the solution to $\max_w U_i(\frac{w}{w + \sum_{j \neq i} w_j}) - w$, which requires (by chain rule of differentiation, and using the shorthands $p = w + \sum_{j \neq i} w_j$ and $x_i = w/p$):

$$\begin{aligned} U_i'(x_i)(\frac{1}{p} - \frac{w}{p^2}) &= 1 \\ \Rightarrow U_i'(x_i)(1 - x_i) &= p. \end{aligned}$$

This implicitly defines x_i in terms of p. Furthermore, the left hand side is easily checked to be a decreasing function of x_i. (Specifically, its derivative is $(1-x_i)U_i"(x_i) - U'(x_i)$, whose first term is negative by concavity and the second because $U_i'(x_i) \geq 0$ by our assumption that U_i is an increasing function.) Thus $\sum_i x_i$ is a decreasing function of p. When $p = +\infty$, the x_i's that maximise utility are all 0, whereas for $p = 0$ the x_i's are all 1, which violates the constraint $\sum_i x_i = 1$. By the mean value theorem, there must exceed a choice of p between 0 and $+\infty$ where $\sum_i x_i = 1$, and the corresponding values of w_i's then constitute a Nash equilibrium.

How does this equilibrium compare to the social optimum? The social optimum satisfies $U_i'(x_i) = p^*$ whereas the Nash equilibrium price p_N corresponds to solving $U_i'(x_i)(1-x_i) = p_N$. If the number of users is large (and the utility functions not "too different"so that the x_i's are not too different) then each x_i is small and $1 - x_i \approx 1$. Thus the Nash equilibrium price is close to but not the same as the socially optimal choice.

Price of Anarchy

One of the notions highlighted by algorithmic game theory is *price of anarchy,* which is the ratio between the cost of the Nash equilibrium and the social optimum. The idea behind this name is that Nash equilibrium is what would be achieved in a free market, whereas social optimum is what could be achieved by a planner who knows everybody's utilities. One identifies a family of games, such as bandwidth sharing, and looks at the *maximum* of this ratio over all choices of the players' utilities. The price of anarchy for the bandwidth sharing game happens to be 4/3. Please see the chapter on inefficiency of equilibria in the AGT book.

19.3 Correlated equilibria

In HW 3 you were asked to simulate two strategies that repeatedly play Rock-Paper-Scissors while minimizing regret. The Payoffs were as follows:

	Rock	Paper	Scissor
Rock	0,0	0, 1	1, 0
Paper	1, 0	0, 0	0, 1
Scissor	0, 1	1, 0	0, 0

Possibly you originally guessed that they would converge to playing Rock, Paper, Scissor randomly. However, this is not regret minimizing since it leads to payoff 0 every third round in the expectation. What you probably saw in your simulation was that the players converged to a *correlated* strategy that guarantees one of them a payoff every other round. Thus they learnt to game the system together and maximise their profits.

This is a subcase of a more general phenomenon, whereby playing low-regret strategies in general leads to a different type of equilibrium, called *correlated equilibrium.*

Example 46 In the game of Chicken, the following is a correlated equilibrium: each of the three pairs of moves other than (Dare, Dare) with probability 1/3. This is a correlated strategy: there is a global random string (or higher agency) that tells the players what to do. Neither player knows what the other has chosen.

Suppose we think of the game being played between two cars approaching a traffic intersection from two directions. Then the correlated equilibrium of the previous paragraph has a nice interpretation: a traffic light! Actually, it is what a traffic light would look like if there were no traffic police to enforce the laws. The traffic light would be programmed to repeatedly pick one of three states with equal probability: (Red, Red), (Green, Red), and (Red, Green). (By contrast, real-life lights cycle between (Red, Green), and (Green, Red); where we are ignoring Yellow for now.) If a motorist arriving at the intersection sees Green, he knows that the other motorist sees Red and so can go through without hesitation. If he sees Red on the other hand, he only knows that there is equal chance that the other motorist sees Red or Green. So acting rationally he will come to a stop since otherwise he has probability 1/2 of getting into an accident. Note that this means that when the light is (Red, Red) then the traffic would be sitting at a halt in both directions.

The previous example illustrates the notion of correlated equilibrium, and we won't define it more precisely. The main point is that it can be arrived at using a simple algorithm, namely, multiplicative weights (this statement also has caveats; see the relevant chapter in the AGT book). Unfortunately, correlated equilibria are also not guaranteed to maximise social welfare.

Bibliography

1. Algorithmic Game Theory. Nisan, Roughgarden, Tardos, Vazirani (eds.), Cambridge University Press 2007.

2. The mathematics of traffic in networks. Frank Kelly. In *Princeton Companion to Mathematics* (T. Gowers, Ed.). PU Press 2008.

3. Settling the Complexity of 2-Player Nash Equilibrium. X. Chen and X. Deng. IEEE FOCS 2006.