Game Theory for Wireless Engineers
Chapter 3, 4

Zhongliang Liang
ECE@Mcmaster Univ

October 8, 2009

Outline

3.1 Definition of A Strategic Form Game - 1

The difference between a traditional optimization and a game is that a player's decisions potentially affects the utility accrued by everyone else in the game.

A game consists of:

- a principal, sets the rules of the game.
- and a finite set of players $\mathbf{N} = \{1, 2, ...N\}$.
- each player $i \in \mathbf{N}$ selects a strategy $s_i \in \mathbf{S_i}$ to maximize his utility u_i.

In this book qwe focus on noncooperative games, where players select strategies without information exchanging.

3.1 Definition of A Strategic Form Game - 2

1. The strategy profile $\mathbf{s}$ is defined as: $\mathbf{s} = (s_i)_{i \in \mathbf{N}} = (s_1, s_2, ...s_N)$.

2. $\mathbf{s_{-i}}$ is the collective strategies of all players except player i.

3. Joint space is defined as the Cartesian Product of the individual strategy spaces: $\mathbf{S} = \times_{i \in \mathbf{N}} \mathbf{S_i}$.

4. Similarly, $\mathbf{S_{-i}} = \times_{j \in \mathbf{N}, j \neq i} \mathbf{S_j}$.

5. Utility function is a mapping from joint space to real number: $u_i(\mathbf{s}) : \mathbf{S} \rightarrow \mathbb{R}$.

Table: Example 3.1

	$s_2 = 0$	$s_2 = 1$
$s_1 = 0$	(0,0)	(1,-1.5)
$s_1 = 1$	(-1.5,1)	(-0.5,-0.5)

3.2 Dominated Strategies and Iterative Deletion of Dominated Strategies - 1

Table: Example 3.2

	$s_2 = L$	$s_2 = R$
$s_1 = L$	(1,1)	(0.5,1.5)
$s_1 = M$	(2,0)	(1,0.5)
$s_1 = R$	(0,3)	(0,2)

Definition 9. A pure strategy s_i is strictly dominated for player i if there exists $s_i' \in \mathbf{S_i}$ such that $u_i(s_i', \mathbf{s_{-i}}) > u_i(s_i, \mathbf{s_{-i}}) \forall \mathbf{s_{-i}} \in \mathbf{S_{-i}}$.

Furthismore, we say that s_i is strictly dominated with respect to $\mathbf{A_{-i}} \subseteq \mathbf{S_{-i}}$ if there exists $s_i' \in \mathbf{S_i}$ such that $u_i(s_i', \mathbf{s_{-i}}) > u_i(s_i, \mathbf{s_{-i}}) \forall \mathbf{s_{-i}} \in \mathbf{A_{-i}}$.

3.2 Dominated Strategies and Iterative Deletion of Dominated Strategies - 2

Using these definitions, we define notation for the undominated strategies with respect to A_{-i} available to player i:

$$\mathbf{D_i}(\mathbf{A_{-i}}) = \{s_i \in \mathbf{S_i} | s_i \text{ is not strictly dominated with respect to } \mathbf{A_{-i}}\}$$

We next define notation for the undominated strategy profiles with respect to a set of strategy profiles $\mathbf{A} \subseteq \mathbf{S}$ (the previous one is for player i, this is the set for all players):

$$\mathbf{D}(\mathbf{A}) = \times_{i \in \mathbf{N}} \mathbf{D_i}(\mathbf{A_{-i}})$$

3.2 Dominated Strategies and Iterative Deletion of Dominated Strategies - 3

The term $\mathbf{D(S)}$ represents the set of all strategy profiles in which no player is playing a dominated strategy.

Likewise,$\mathbf{D^2(S)} = \mathbf{D(D(S))}$ represents the set of all strategy profiles in which no player is playing a strategy that is dominated with respect to the set of undominated strategy profiles D(S).

The sets $\mathbf{D^3(S)}, \mathbf{D^4(S)}, \ldots$ are similarly defined, and it can be shown that $\mathbf{D^{(k+1)}(S)} \subseteq \mathbf{D^k(S)}$.

The set $\mathbf{D^\infty(S)} = \lim_{\mathbf{k}\to\infty} \mathbf{D^k(S)}$ is well defined and nonempty. It's the set of strategy profiles that survive the iterated deletion of dominated strategies.

Notice by doing this we are deleting the obviously worst strategies, but that doesn't necessarily mean we'll find the best. However, at least we know the best strategies, if exist, are among those who have survived the deletion.

3.2 Dominated Strategies and Iterative Deletion of Dominated Strategies - 4

Unfortunately, for some very well known games $\mathbf{D}^{\infty}(\mathbf{S})$ may be equal to $\mathbf{S}$, yielding no predictive power whatsoever! (The paper-scissorsrock game in the next section is an example of such a game.)

Hence, we need a stronger predictive notion. We therefore move to the broader concept of Nash equilibria, while noting that every Nash equilibrium (in pure strategies) is a member of the set$\mathbf{D}^{\infty}(\mathbf{S})$. First, though, we introduce mixed strategies.

3.3 Mixed Strategies - 1

Thus far, we have been assuming that each player picks a single strategy in his strategy set.

However, an alternative is for player i to randomize over his strategy set, adopting what is called a *mixed strategy*. For instance, a player could decide to choose a strategy with some probability $0 < p < 1$.

We denote a mixed strategy available to player i as σ_i(a random variable) .We denote by $\sigma_i(s_i)$ the probability that σ_i assigns to s_i .

Clearly, $\sum_{s_i \in \mathbf{S_i}} \sigma_i(s_i) = 1$.

The space of player i's mixed strategies is Σ_i (similar as $\mathbf{s_i}$).

3.3 Mixed Strategies - 2

We note that the expected utility of player i under joint mixed strategy σ is given by

$$u_i(\sigma) = \sum_{\mathbf{s} \in \mathbf{S}} \left(\prod_{j=1}^{\mathbf{N}} \sigma_j(s_j) \right) u_i(\mathbf{s})$$

It is convenient to define the support of mixed strategy σ_i as the set of pure strategies to which it assigns positive probability: supp $\sigma_i = \{s_i \in \mathbf{S_i} : \sigma_i(s_i) > 0\}$ (excluded the one with probability 0).

There are numerous games where no pure strategy can be justified (or, no equilibria in pure strategies), and where the logical course of action is to randomize over pure strategies. An example is the well-known paper-scissors-rock game, whose logical strategy is to randomize among the three pure strategies, each with probability $1/3$.

3.4 Nash Equilibrium - 1

The Nash equilibrium is a joint strategy where no player can increase his utility by unilaterally deviating. In pure strategies, that means:

Definition 10. Strategy $\mathbf{s} \in \mathbf{S}$ is a Nash equilibrium if $u_i(\mathbf{s}) \geq u_i(\hat{s}_i, \mathbf{s}_{-i})\ \forall \hat{s}_i \in \mathbf{S_i}, \forall i \in \mathbf{N}$.

3.4 Nash Equilibrium - 2

An alternate interpretation of the definition is that it is a mutual best response from each player to other players strategies.

Let us first define the *best-reply correspondence* for player i as a point-to-set mapping that associates each strategy profile $\mathbf{s} \in \mathbf{S}$ with a subset of $\mathbf{S_i}$:

$$\mathbf{M_i}(\mathbf{s}) = \{\arg\max_{\hat{s}_i \in \mathbf{S_i}} u_i(\hat{s}_i, \mathbf{s_{-i}})\}$$

The best-reply-correspondence for the game is then defined as

$$\mathbf{M}(\mathbf{s}) = \times_{i \in \mathbf{N}} \mathbf{M_i}(\mathbf{s})$$

We can now say that strategy s is a Nash equilibrium if and only if

$s \in \mathbf{M}(\mathbf{s})$. Note that this definition is equivalent to (and, indeed, a corollary of) Definition 10.

3.4 Nash Equilibrium - 3*

Let us generalize the previous discussion by taking into consideration mixed strategies. We begin by revisiting the concept of best reply:

Definition 11. The best reply correspondence in pure strategies for player $i \in \mathbf{N}$ is a correspondence $\mathbf{r_i} : \Sigma \rightrightarrows \mathbf{S_i}$ defined as $\mathbf{r_i}(\sigma) = \{\arg\max_{s_i \in \mathbf{S_i}} u_i(s_i, \sigma_{-\mathbf{i}})\}$.

This definition describes a player's pure strategy best response(s) to opponents' mixed strategies. However, it is possible that some of a players best responses would themselves be mixed strategies, leading us to define:

Definition 12. The best reply correspondence in mixed strategies for player $i \in \mathbf{N}$ is a correspondence $\mathbf{r_i} : \Sigma \rightrightarrows \mathbf{\Sigma_i}$ defined as $\mathbf{mr_i}(\sigma) = \{\arg\max_{\sigma_i \in \mathbf{\Sigma_i}} u_i(\sigma_i, \sigma_{-\mathbf{i}})\}$.

We can now expand on our previous definition of Nash equilibrium to allow for mixed strategy equilibria:

Definition 13. A mixed strategy profile $\sigma \in \Sigma$ is a Nash equilibrium if $u_i(\sigma) \geq u_i(s_i, \sigma_{-\mathbf{i}})\ \forall i \in \mathbf{N}, \forall s_i \in \mathbf{S_i}$.

3.4 Nash Equilibrium - 5

The Nash equilibrium is a consistent prediction of the outcome of the game in the sense that if all players predict that a Nash equilibrium will occur, then no player has an incentive to choose a different strategy.

Furthismore, if players start from a strategy profile that is a Nash equilibrium, no player will deviate, and the system will be in equilibrium.

But what happens if players start from a nonequilibrium strategy profile? Will it converge to equilibrium? And what will happen if there are multiple Nash Equilibria?

In addition, a Nash equilibrium, even if it is not vulnerable to unilateral deviation by a single player, may be vulnerable to deviations by a coalition of players.

3.5 Existence of Nash Equilibria - 1*

As fixed point theorems are key to establishing the existence of a Nash equilibrium, we start by reviewing the concept of a fixed point:

Definition 14. Consider a function with identical domain and range: $f : \mathbf{X} \rightarrow \mathbf{X}$. We say that fixed point of function f if $f(x) = x$.

This definition can be generalized to apply to point-to-set functions (i.e., correspondences):

Definition 15. Consider a correspondence that maps each point $\mathbf{x} \in \mathbf{X}$ to a set $\phi(x) \subset X$. Denote this correspondence by $\phi : \mathbf{X} \rightrightarrows \mathbf{X}$. We say x is a fixed point of ϕ if $x \in \phi(x)$.

Then, by the definition of the Nash equilibrium as a mutual best response, we see that any fixed point of **mr** is a Nash equilibrium. Even more, any Nash equilibrium is a fixed point of **mr**.

3.5 Existence of Nash Equilibria - 2*

Definition 16. A correspondence ϕ from a subset $\mathbf{T}$ of Euclidean space to a compact subset $\mathbf{V}$ of Euclidean space is upper hemicontinuous at point $x \in \mathbf{T}$ if $x_r \to x, y_r \to y$, where $y_r \in \phi(x_r) \forall r$, implies $y \in \phi(x)$. The correspondence is upper hemicontinuous if it is upper hemicontinuous at every $x \in \mathbf{T}$.

Theorem 10 (Kakutani)*. Let $\mathbf{X} \subset \mathbb{R}^m$ be compact and convex. Let the correspondence $\phi : X \rightrightarrows X$ be upper hemicontinuous with nonempty convex values. Then ϕ has a fixed point.

Lemma 3*. Σ is compact and convex.

Lemma 4*. Let $\mathbf{mr}(\sigma) = \times_{i \in \mathbf{N}} \mathbf{mr_i}(\sigma)$. The correspondence $\mathbf{mr}(\sigma)$ is nonempty and convex.

Lemma 5*. The correspondence mr is upper hemicontinuous.

The three lemmas above, combined with Kakutanis fixed point theorem, establish the following theorem:

Theorem 11 (Nash). Every finite game in strategic form has a Nash equilibrium in either mixed or pure strategies.

Theorem 12*. Consider a strategic-form game with strategy spaces Si that are nonempty compact convex subsets of an Euclidean space. If the payoff functions u_i are continuous in s and quasi-concave in s_i , there exists a Nash equilibrium of the game in pure strategies.

Recall that, a *strategic-form game* is the kind of game that can be represented by a matrix which shows the players, strategies, and payoffs (see the example to the right). More generally it can be represented by any function that associates a payoff for each player with every possible combination of actions.

3.6 Applications - 1

Pricing of Network Resources.

In networks offering different levels of QoS, both network performance and user satisfaction will be directly influenced by the users choices as to what level of service to request. Since each users choice of service may be influenced by both the pricing policy and other users behavior, the problem can naturally be treated under a game-theoretic framework, in which the operating point of the network is predicted by the Nash equilibrium.

For instance, in priority-based networks, a strategy may be the priority level a user requests for his traffic; in networks that support delay or data rate guarantees, a strategy may be the minimum bandwidth to which a user requests guaranteed access. The tradeoff, of course, is that the higher the level of service requested the higher the price to be paid by the user. The network service provider architects the Nash equilibria by setting the rules of the game: the pricing structure and the dimensioning of network resources.

Flow Control.

Flow control, each user determines the traffic load he will offer to the network in order to satisfy some performance objective, is another network mechanism that has been modeled using game theory.

One of the earliest such models was developed in 1. In that model, a finite number of users share a network of queues. Each users strategy is the rate at which he offers traffic to the network at each available service class, constrained by a fixed maximum rate and maximum number of outstanding packets in the network. The performance objective is to select an admissible flow control strategy that maximizes average throughput subject to an upper bound on average delay. The authors were able to determine the existence of an equilibrium for such a system.
[1]

[1] Y. A. Korilis and A. A. Lazar, "On the existence of equilibria in noncooperative optimal flow control", J. ACM, 1995.

Chapter 4 Repeated and Markov Games

This chapter considers the concepts of repeated games and Markov games and takes a brief dip into the waters of extensive form games.

In a repeated game formulation, players participate in repeated interactions within a potentially infinite time horizon. Players must, therefore, consider the effects that their chosen strategy in any round of the game will have on opponents strategies in subsequent rounds. Each player tries to maximize his expected payoff over multiple rounds.

Before we go any further, let us discuss the extensive form representation of a game,

4.1.1 Extensive Form Representation-1

A game in extensive form is represented as a tree, where each node of the tree represents a decision point for one of the players, and the branches coming out of that node represent possible actions available to that player.

An information set is a collection of decision nodes that are under the control of the same player and which the player is unable to distinguish from one another. In other words, if the player reaches any of those nodes, he will not know which node he has reached.

Note that any game in strategic form can also be represented in extensive form (and vice versa). So, the extensive form representation does not necessarily imply that players actions are taken sequentially.

4.1.1 Extensive Form Representation - 2

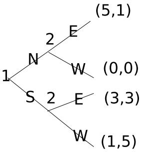

	EE	EW	WE	WW
N	5,1	5,1	0,0	0,0
S	3,3	1,5	3,3	1,5

Let us now construct an example of a **sequential game** where players actions are taken sequentially: when player 2 selects his actions, he is aware of what action player 1 has taken. Player 1 chooses between two directions (North and South) and, once he made his choice, player 2 then chooses between two directions (East and West). This is shown as extensive form in the figure above.

In our last example, player 1's strategies can be simply stated as North or South. Player 2s strategies are more complex. One possible strategy is "move East if player 1 moves North but move West if player 1 moves South": this is denoted as strategy EW. While there are two possible strategies for player 1, there are four strategies available to player 2. What strategies are likely outcomes of this game?

4.1.2 Equilibria in Repeated Games - 1*

In addition to the basic notion of Nash equilibria, which can be applied to repeated games virtually unmodified, this section introduces the "subgame perfect equilibrium" – a stronger equilibrium concept that is extremely important in the literature on extensive form games.

Take a node x in an extensive form game(EFG). Let $F(x)$ be the set of nodes and branches that follow x, including x. A subgame is a subset of the entire game such that the following properties hold:

1. the subgame is rooted in a node x, which is the only node of that information set;
2. the subgame contains all nodes $y \in F(x)$; and
3. if a node in a particular information set is contained in the subgame, then all nodes in that information set are also contained.

A proper subgame of an EFG is a subgame whose root is not the root of this EFG. Now, we are ready to define subgame perfect equilibrium:

Definition 17. A subgame perfect equilibrium $\hat{\sigma}$ of an EFG is a Nash equilibrium of this EFG that is also a Nash equilibrium for every proper subgame of this EFG .

4.1.2 Equilibria in Repeated Games - 2*

Again, let us go back to the example. There are two proper subgames for this game: the subgames are rooted in each of the nodes belonging to player 2. There are three Nash equilibria, but the only subgame perfect equilibrium is (N,EW).

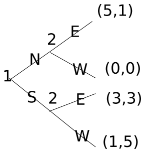

	EE	EW	WE	WW
N	5,1	5,1	0,0	0,0
S	3,3	1,5	3,3	1,5

Let us now introduce some notation for the modeling of repeated games in strategic form.

4.1.3 Repeated Games in Strategic Form -1

Let us now express a repeated game in strategic form. As before, $\mathbf{N}$ denotes the set of players. And $\mathbf{A_i}$ the set of actions available to player i in each round of the game. We use the word *action* to emphasize that these refer to decisions the player makes in a given round; this is to differentiate from *strategy* , which refers to rules that map every possible information state the player can be in into an action.

Similarly, $\mathbf{A} = \times_{i \in \mathbf{N}} \mathbf{A_i}$ and the action profile $\mathbf{a} = (a_i)_{i \in \mathbf{N}}$. Since a player can also randomize his actions, it is convenient to define mixed action α_i as a randomization of actions a_i . Define the payoff at a stage of a game g_i, and $g(\mathbf{a}) = (g_1(\mathbf{a}), ...g_N(\mathbf{a}))$. This game is referred to as the *stage game*. We index the actions that players adopt in each round k, as $\mathbf{a^k} = (a_1^k, a_2^k, ..., a_N^k)$.

Players strive to maximize their expected payoff over multiple rounds of the game. The resulting payoff is often expressed as a sum of single-round payoffs, discounted by a value $0 \leq \delta < 1$. In this manner, players place more weight on the payoff in the current round than on future payoffs. The average discounted payoff can be expressed as:

$$u_i = (1 - \delta) \sum_{k=0}^{\infty} (\delta)^k g_i(\mathbf{a^k})$$

4.1.4 Node Cooperation: A Repeated Game Example - 1

Consider a repeated game played K times, where K is a geometrically distributed random variable with parameter $0 < p < 1$. We can write $p_k = \text{Prob}[K = k] = p(1 - p)k$, $k = 0, 1, 2, \ . \ . \ .$ and therefore a $E[K] = (1 - p)/p$. Note that as $p \rightarrow 1$ the probability that there will be a next round for the game approaches 0.

We consider homogeneous action spaces for all players $\mathbf{A_i} = 0, 1$, where 1 represents a node decides to share its files, while 0 represents refraining from sharing.

we consider a user's payoff in round k to be the sum of two components:$g_i(\mathbf{a^k}) = \alpha_i(\mathbf{a^k}) + \beta_i(\mathbf{a^k})$. The 1st term represents the benefit accrued by a player from his opponents sharing their resources. On the other hand, there are costs to sharing ones own resources, and those costs are represented by the latter term and $\beta_i(0) = 0$.

Consider a grim-trigger strategy adopted by each node: cooperate as long as all other nodes share their resources; defect if any of the others have deviated in the previous round.

Let us consider any round of the game. If a player cooperates, the payoff he should expect from that point forward is

$$[\alpha_i(N-1)+\beta_i(1)][1+\sum_{k=0}^{\infty} kp(1-p)^k] = \frac{\alpha_i(N-1)+\beta_i(1)}{p}$$

If, on the other hand, a player deviates, his expected payoff from that round on is simply $\alpha_i(N-1)$. So, the grim trigger strategy is a Nash equilibrium if the following inequality holds for all players i:

$$\alpha_i(N-1) > -\frac{\beta_i(1)}{1-p}$$

(Why?)

4.1.5 The "Folk Theorems" - 1*

Before introducing the folk theorems, we need to go through a series of definitions. We start with the concept of feasible payoffs.

Definition 18. The stage game payoff vector $v = (v_1, v_2, ..., v_N)$ is feasible if it is an element of the convex hull $\mathbf{V}$ of pure strategy payoffs for the game:

$$\mathbf{V} = convexhull\{\mathbf{u}|\forall \mathbf{a} \in \mathbf{A} such that g(\mathbf{a}) = \mathbf{u}\}.$$

Definition 19. The min-max payoff for player i is defined as

$$\underline{v_i} = min_{-\alpha_i \in \Delta(\mathbf{A_{-i}})} max_{\alpha_i \in \Delta(\mathbf{A_i})} g_i(\alpha_i, \alpha_{-\mathbf{i}})$$

Definition 20. The set of feasible strictly individually rational payoffs is

$$\{v \in \mathbf{V} | v_i > \underline{v_i} \ \forall i \in \mathbf{N}\}$$

4.1.5 The "Folk Theorems" - 2*

We can now state one of the folk theorems:

Theorem 13. For every feasible strictly individually rational payoff vector v, $\exists \underline{\delta} < 1$ such that $\forall \delta \in (\delta, 1)$ there is a Nash equilibrium of the game $\Gamma^r(\delta)$ with payoffs v.

Theorem 14. Let there be an equilibrium of the stage game that yields payoffs $e = (e_i)_{i \in N}$. Then for every $v \in V$ with $v_i > e_i$ for all players i, $\exists \delta$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a subgame-perfect equilibrium of $\Gamma^r(\delta)$ with payoffs v.

4.2 Markov Games: Generalizing The Repeated Game Idea

There is a very natural relationship between the notion of a repeated game and that of a Markov game. This section introduces Markov game setting, Markov strategies, and Markov equilibrium.

Let us formalize the definition of the game. The game is characterized by state variables $m \in \mathbf{M}$.We must also define a transition probability $q(m^{k+1}|m^k, \mathbf{a^k})$, denoting the probability that the state at the next round is m^{k+1} conditional on being in state m^k during round k and on the playing of action profile $\mathbf{a^k}$.

The history at stage k is $\mathbf{h^k} = (m^0, \mathbf{a^0}, m^1, \mathbf{a^1}, ..., \mathbf{a^{k-1}}, m^k)$. At stage k, each player is aware of the history $\mathbf{h^k}$ before deciding on his action for that stage. Markov strategies are often denoted $\delta_i : \mathbf{M} \rightarrow \Delta(\mathbf{A_i})$, even though this is a slight abuse of notation.

A Markov perfect equilibrium is a profile of Markov strategies, which yields a Nash equilibrium in every proper subgame. It can be shown that Markov perfect equilibria are guaranteed to exist when the stochastic game has a finite number of states and actions.

Markov chains are used to model a number of communications and networking phenomena, such as channel conditions, slot occupancy in random channel access schemes, queue state in switches, etc. It is natural that Markov games would find particular applications in this field. In particular, the Markov game results imply that if we can summarize the state of a network or a communications link with one or more "state variables", then we lose nothing by considering only strategies that consider only these state variables without regard to past history or other information.

Power Control in Cellular Networks

The problem of power control in a cellular system has often been modeled as a game. Let us take a CDMA system to illustrate what makes game theory appealing to the treatment of power control.

In a CDMA system, we can model users utilities as an increasing function of signal to interference and noise ratio (SINR) and a decreasing function of power.

This would be a local optimization problem, with each user determining his own optimal response to the tradeoff, if increasing power did not have any effect on others.

The elements of a game are all here: clear tradeoffs that can be expressed in a utility function (what the exact utility function should be is another issue) and clear interdependencies among users decisions.

Let p_i be the power level selected by user i and γ_i the resulting SINR, itself a function of the action profile: $\gamma_i = f(\mathbf{p})$. We can express the utility function as a function of these two factors. One possible utility function is from:

$$u_i(\mathbf{p}) = u_i(p_i, \gamma_i) = \frac{R}{p_i}(1 - 2BER(\gamma_i))^L$$

where R is the rate at which the user transmits, $BER(\gamma_i)$ is the bit error rate achieved given an SINR of γ_i, and L is the packet size in bits.

It has been shown that this problem modeled as a one-stage game has a unique Nash equilibrium. Like the resource sharing game described earlier, this Nash equilibrium happens to be inefficient.

Now consider the same game played repeatedly. It is now possible to devise strategies wherein a user is punished by others if he selects a selfish action (such as increasing power to a level that significantly impairs others SINR). If the users objective is to maximize his utility over many stages of the game, and if the game has an infinite time horizon, the threat of retaliation leads to a Pareto efficient equilibrium.

4.3 Application - 3 (Pareto Efficiency)

Informally, Pareto efficient situations are those in which any change to make any person better off is impossible without making someone else worse off.

Given a set of alternative allocations of, say, goods or income for a set of individuals, a change from one allocation to another that can make at least one individual better off without making any other individual worse off is called a Pareto improvement. An allocation is defined as Pareto efficient or Pareto optimal when no further Pareto improvements can be made.

*Maskin, "Nash Equilibrium and Welfare Optimality", Review of Economic Studies, 1999, 66, 23-38.

Medium Access Control

In this section, we describe a repeated game of perfect information applied to characterizing the performance of slotted Aloha in the presence of selfish users.

In this model, users compete for access to a common wireless channel. During each slot, the actions available to each user are to transmit or to wait. The channel is characterized by a matrix $\mathbf{R} = [\rho_{nk}]$, with ρ_{nk} defined as the probability that k frames are received during a slot where there have been n transmission attempts. The expected number of successfully received frames in a transmission of size n is therefore $r_n = \sum_{k=0}^{n} k\rho_{nk}$. If we further assume that all users who transmit in a given slot have equal probability of success, then the probability that a given users transmission is successful is given by $\frac{r_n}{n}$.

This is a repeated game, as for each slot, each user has the option to attempt transmission or to wait. The cost of a transmission is assumed to be $c \in (0, 1)$, with a successful transmission accruing utility of $1 - c$, an unsuccessful transmission utility $-c$ and the decision to wait giving a utility of 0. Payoffs are subject to a per-slot discount factor of $0 \leq \delta < 1$.

This characterization of medium access control provides insight into the price of distributed decisions (vs. those controlled by a base station or access point) as well as the impact of different channel models on aggregate throughput expected for random access.

End of Chapter 3 and 4, thank you!