

Game Theory in Wireless Networks: A Tutorial

Mark Felegyhazi, Jean-Pierre Hubaux

EPFL – Switzerland

email: {mark.felegyhazi, jean-pierre.hubaux}@epfl.ch

EPFL Technical report: LCA-REPORT-2006-002,

submitted on Feb. 21, 2006; revised on Jun. 28, 2006

Abstract—The behavior of a given wireless device may affect the communication capabilities of a neighboring device, notably because the radio communication channel is usually shared in wireless networks. In this tutorial, we carefully explain how situations of this kind can be modelled by making use of game theory. By leveraging on four simple running examples, we introduce the most fundamental concepts of non-cooperative game theory. This approach should help students and scholars to quickly master this fascinating analytical tool without having to read the existing lengthy, economics-oriented books. It should also assist them in modelling problems of their own.

1 INTRODUCTION

Game theory [6, 7, 13] is a discipline aimed at modelling situations in which decision makers have to make specific actions that have mutual, possibly conflicting, consequences. It has been used primarily in economics, in order to model competition between companies: for example, should a given company enter a new market, considering that its competitors could make similar (or different) moves? Game theory has also been applied to other areas, including politics and biology.¹

The first textbook in this area was written by von Neumann and Morgenstern, in 1944 [19]. A few years later, John Nash made a number of additional contributions [11, 12], the cornerstone of which is the famous *Nash equilibrium*. Since then, many other researchers have contributed to the field, and in a few decades game theory has become a very active discipline; it is routinely taught in economics curricula. An amazingly large number of game theory textbooks have been produced, but almost all of them consider economics as the premier application area (and all their concrete examples are inspired by that field). Our tutorial is inspired by three basic textbooks and we mention them in the ascending order of complexity. Gibbons [7] provides a very nice, easy-to-read introduction to non-cooperative game theory with many examples using economics. Osborne and Rubinstein [13] introduce the game-theoretic concepts very precisely, although this book is more difficult to read because of the more formal development. This is the only book out of the three that covers cooperative game theory as well. Finally, Fudenberg and Tirole's [6] book covers many advanced topics, in addition to the basic concepts.

Not surprisingly, game theory has also been applied to networking, in most cases to solve routing and resource allocation

problems in a competitive environment. The references are so numerous that we cannot list them due to space constraints. A subset of these papers is included in [1]. Recently, game theory was also applied to wireless communication: the decision makers in the game are devices willing to transmit or receive data (e.g., packets). They have to cope with a limited transmission resource (i.e., the radio spectrum) that imposes a conflict of interests. In an attempt to resolve this conflict, they can make certain moves such as transmitting now or later, changing their transmission channel, or adapting their transmission rate.

There is a significant amount of work in wired and wireless networking that make use of game theory. Oddly enough, there exists no comprehensive tutorial specifically written for wireless networkers². We believe this situation to be unfortunate, and this tutorial has been written with the hope of contributing to fill this void. As game theory is still rarely taught in engineering and computer science curricula, we assume that the reader has no (or very little) background in this field; therefore, we take a basic and intuitive approach. Because in most of the strategic situations in wireless networking the players have to agree on sharing or providing a common resource in a distributed way, our approach focuses on the theory of *non-cooperative games*. Cooperative games require additional signalization or agreements between the decision makers and hence a solution based on them might be more difficult to realize.

In a non-cooperative game, there exist a number of decision makers, called *players*, who have potentially conflicting interests. In the wireless networking context, the players are the *users* controlling their devices. As we assume that the devices are bound to their users, we will refer to devices as players and we use the two terms interchangeably throughout the paper. In compliance with the practice of game theory, we assume that the players are *rational*, which means that they try to maximize their *payoff* or alternatively to minimize their *costs*³. This assumption of rationality is often questionable, given, for example, the altruistic behavior of some animals. Herbert A. Simon was the first one was to question this assumption and introduced the notion of *bounded rationality*

²To the best of our knowledge, there exist only two references: a monograph [10] that provides a synthesis of lectures on the topic, and a survey [1] that focuses mostly on wired networks.

³In game theory one usually uses the concept of payoff maximization, whereas cost minimization comes from the optimal control theory community. As it is more appropriate for this tutorial, we use the payoff maximization objective.

¹The name itself of “game theory” can be slightly misleading, as it could be associated with parlor games such as chess and checkers. Yet, this connection is not completely erroneous, as parlor games do have the notion of players, payoffs, and strategies - concepts that we will introduce shortly.




Figure 1. The network scenario in the Forwarder's Dilemma game.

[18]. But, we believe that in computer networks, most of the interactions can be captured using the concept of rationality, with the appropriate adjustment of the payoff function. In order to maximize their payoff, the players act according to their *strategies*. The strategy of a player can be a single *move* (as we will see in Section 2) or a set of moves during the game (as we present in Section 4).

In this tutorial, we devote particular attention to the selection of the examples so that they match our focus on wireless networks. For the sake of clarity, and similarly to classic examples, we define these examples for two decision makers, hence the corresponding games are two-player games. Note that the application of game theory extends far beyond two-person games. Indeed, in most networking problems, there are several participants.

We take an intuitive top-down approach in the protocol stack to select the examples in wireless networking as follows. Let us first assume that the time is split into *time slots* and each device can make one move in each time slot.

- 1) In the first game called the *Forwarder's Dilemma*⁴, we assume that there exist two devices as players, p_1 and p_2 . Each of them wants to send a packet to his receiver, r_1 and r_2 respectively, in each time slot using the other player as a forwarder. We assume that the communication between a player and his receiver is possible only if the other player forwards the packet. We show the Forwarder's Dilemma scenario in Figure 1. If player p_1 forwards the packet of p_2 , it costs player p_1 a fixed *cost* $0 < c \ll 1$, which represents the energy and computation spent for the forwarding action. By doing so, he enables the communication between p_2 and r_2 , which gives p_2 a *reward* of 1. The payoff is the difference of the reward and the cost. We assume that the game is symmetric and the same reasoning applies to the forwarding move of player p_2 . The dilemma is the following: Each player is tempted to drop the packet he should forward, as this would save some of his resources; but if the other player reasons in the same way, then the packet that the first player wanted to be relayed will be dropped. They could, however, do better by mutually relaying each other's packet. Hence the dilemma.
- 2) In the second example, we present a scenario, in which a sender se wants to send a packet to his receiver r in each time slot. To this end, he needs *both devices* p_1 and p_2 to forward for him. Thus, we call this game the *Joint Packet Forwarding Game*. Similarly to the previous example, there is a forwarding cost $0 < c \ll 1$ if a player forwards



Figure 2. The Joint Packet Forwarding Game.

the packet of the sender. If both players forward, then they each receive a reward of 1 (e.g., from the sender or the receiver). We show this packet forwarding scenario in Figure 2.

- 3) The third example, called *Multiple Access Game*⁵, introduces the problem of medium access. Assume that there are two players p_1 and p_2 who want to send some packets to their receivers r_1 and r_2 using a shared medium. We assume that the players have a packet to send in each time slot and they can decide to transmit it or not. Suppose furthermore that p_1 , p_2 , r_1 and r_2 are in the power range of each other, hence their transmissions mutually interfere. If player p_1 transmits his packet, it incurs a transmission cost of $0 < c \ll 1$, similarly to the previous examples. The packet transmission is successful if p_2 does not transmit (stays quiet) in that given time slot, otherwise there is a collision. If there is no collision, player p_1 gets a reward of 1 from the successful packet transmission.
- 4) In the last example, we assume that player p_1 wants to transmit a packet in each time slot to a receiver r_1 . In this example, we assume that the wireless medium is split into two channels ch_1 and ch_2 according to the Frequency Division Multiple Access (FDMA) principle [14, 16]. The objective of the *malicious* player p_2 is to prevent player p_1 from a successful transmission by transmitting on the same channel in the given time slot. In wireless communication, this is called *jamming*, hence we refer to this game as the *Jamming Game*⁶. Clearly, the objective of p_1 is to succeed in spite of the presence of p_2 . Accordingly, he receives a payoff of 1 if the attacker cannot jam his transmission and he receives a payoff of -1 if the attacker jams his packet. The payoffs for the attacker p_2 are the opposite of those of player p_1 . We assume that p_1 and r_1 are synchronized, which means that r_1 can always receive the packet, unless it is destroyed by the malicious player p_2 . Note that we neglect the transmission cost c , since it applies to each payoff (i.e., the payoffs would be $1 - c$ and $-1 - c$) and does not change the conclusions drawn from this game.

We deliberately chose these examples to represent a wide range of problems over different protocol layers (as shown in Figure 3). There are indeed fundamental differences between these games as follows. The Forwarder's Dilemma is a symmetric *nonzero-sum game*, because the players can mutually increase their payoffs by cooperating (i.e., from zero to $1 - c$). The conflict of interest is that they have to provide the packet

⁴We have chosen this name as a tribute to the famous Prisoner's Dilemma game in the classic literature [2, 7, 6, 13].

⁵In the classic game theory textbooks, this type of game is referred to as the "Hawk-Dove" game, or sometimes the "Chicken" game.

⁶In the classic game theory literature, this game corresponds to the game of "Matching Pennies."




Figure 3. The classification of the examples according to protocol layers.

forwarding service for each other. Similarly, the players have to establish the packet forwarding service in the Joint Packet Forwarding Game, but they are not in a symmetric situation anymore. The Multiple Access Game is also a nonzero-sum game, but the players have to share a common resource, the wireless medium, instead of providing it. Finally, the Jamming Game is a *zero-sum game* because the reward of one player represents the loss of the other player, meaning that $\sum_{i \in \mathcal{N}} (\text{reward}_i - \text{cost}_i) = 0$. These properties lead to different games and hence to different strategic analyses, as we will demonstrate in the next section.

2 STATIC GAMES

In this section, we assume that there exists only one time slot, which means that the players have only one move as a strategy. In the game-theoretic terms this is called a *static game*. We will demonstrate how game theory can be used to analyze the games introduced before and to identify the possible outcomes of the strategic interactions of the players.

2.1 Static Games in Strategic Form

We define a game $G = (P, S, U)$ in *strategic form* (or normal form) by the following three elements. P is the set of *players*. Note that in our paper we have two players, $p_1, p_2 \in P$, but we present each definition such that it holds for any number of players. For convenience, we will designate by subscript $-i$ all the players belonging to P except i himself. These players are often designated as being the *opponents* of i . In our games, player i has one opponent referred to as j . S_i corresponds to the *pure strategy space* of player i . This means that the strategy assigns zero probability to all moves, except one (i.e., it clearly determines the move to make). We will see in Section 2.4, that the players can also use *mixed strategies*, meaning that they choose different moves with different probabilities. We designate the joint set of the strategy spaces of all players as follows $S = S_1 \times \dots \times S_{|\mathcal{N}|}$. We will represent the pure strategy space of the opponents of player i by $S_{-i} = S \setminus S_i$. The set of chosen strategies constitutes a *strategy profile* $s = \{s_1, s_2\}$. In this tutorial, we have the same strategy space for both players, thus $S_1 = S_2$. Note that our examples have two players and thus we refer to the strategy profile of the opponents as $s_{-i} = s_j \in S$. The *utility*⁷ or payoff

⁷Note that there is a subtle difference between the concepts of utility and payoff. Payoff is usually used for the outcome of a stage, i.e. stage payoff, whereas utility denotes the objective of a player in the whole game. Hence, the two notions are the same for static games, but different for dynamic games.

$u_i(s)$ expresses the benefit of player i given the strategy profile s . In our examples, we have $U = \{u_1(s), u_2(s)\}$. Note that the objectives (i.e., utility functions) might be different for the two players, as for example in the Jamming Game.

At this point of the discussion, it is very important to explicitly state that we consider the game to be with *complete information*.

Definition 1 A game with complete information is a game in which each player has full knowledge of all aspects of the game.

In particular, complete information means that the players know each element in the game definition: (i) who the other players are, (ii) what their possible strategies are and (iii) what payoff will result for each player for any combination of moves. The concept of complete information should not be confused with the concept of *perfect information*, another concept we present in detail in Section 3.3.

Let us first study the Forwarder's Dilemma in a static game. As mentioned before, in the static game we have only one time slot. The players can decide to forward (F) the packet of the other player or to drop it (D); this decision represents the strategy of the player. As mentioned earlier, this is a nonzero-sum game, because by helping each other to forward, they can achieve an outcome that is better for both players than mutual dropping.

Matrices provide a convenient representation of strategic-form games with two players. We can represent the Forwarder's Dilemma game as shown in Table I. In this table, p_1 is the row player and p_2 is the column player. Each cell of the matrix corresponds to a possible combination of the strategies of the players and contains a pair of values representing the payoffs of players p_1 and p_2 , respectively.

$p_1 \setminus p_2$	F	D
F	(1-c, 1-c)	(-c, 1)
D	(1, -c)	(0, 0)

Table I. The Forwarder's Dilemma game in strategic form, where p_1 is the row player and p_2 is the column player. Each of the players has two strategies: to forward (F) or to drop (D) the packet of the other player. In each cell, the first value is the payoff of player p_1 , whereas the second is the payoff of player p_2 .

2.2 Iterated Dominance

Once the game is expressed in strategic form, it is usually interesting to *solve* it. Solving a game means predicting the strategy of each player, considering the information the game offers and assuming that the players are rational. There are several possible ways to solve a game; the simplest one consists in relying on *strict dominance*.

Definition 2 Strategy s'_i of player i is said to be strictly dominated by his strategy s_i if,

$$u_i(s'_i, s_{-i}) < u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i} \quad (1)$$

Coming back to the example of Table I, we solve the game by *iterated strict dominance* (i.e., by iteratively eliminating

strictly dominated strategies). If we consider the situation from the point of view of player p_1 , then it appears that for him the F strategy is strictly dominated by the D strategy. This means that we can eliminate the first row of the matrix, since a rational player p_1 will never choose this strategy. A similar reasoning, now from the point of view of player p_2 , leads to the elimination of the first column of the matrix. As a result, the solution of the game is (D, D) and the payoff is $(0, 0)$. This can seem quite paradoxical, as the pair (F, F) would have led to a better payoff for each of the players. It is the lack of trust between the players that leads to this suboptimal solution.⁸

The technique of iterated strict dominance cannot be used to solve every game. Let us now study the Joint Packet Forwarding Game. The two devices have to decide whether to forward the packet simultaneously, before the source actually sends it⁹. Table II shows the strategic form.

$p_1 \backslash p_2$	F	D
F	$(1-c, 1-c)$	$(-c, 0)$
D	$(0, 0)$	$(0, 0)$

Table II. The Joint Packet Forwarding Game in strategic form. The players have two strategies: to forward (F) or to drop (D) the packet sent by the sender. Both players p_1 and p_2 get a reward, but only if each of them forwards the packet.

In the Joint Packet Forwarding Game, none of the strategies of any player strictly dominates the other. If player p_1 drops the packet, then the move of player p_2 is indifferent and thus we cannot eliminate his strategy D based on strict dominance. To overcome the requirements defined by strict dominance, we define the concept of *weak dominance*.

Definition 3 Strategy s'_i of player i is said to be weakly dominated by his strategy s_i if,

$$u_i(s'_i, s_{-i}) \leq u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i} \quad (2)$$

with strict inequality for at least one $s_{-i} \in S_{-i}$.

Using the concept of weak dominance, one can notice that the strategy D of player p_2 is weakly dominated by the strategy F . One can perform an elimination based on *iterated weak dominance*, which results in the strategy profile (F, F) . Note, however, that the solution of the iterated strict dominance technique is unique, whereas the solution of the iterated weak dominance technique might depend on the sequence of eliminating weakly dominated strategies, as explained at the end of Section 2.3.

It is also important to emphasize that the iterated elimination techniques are very useful, even if they do not result in a single strategy profile. These techniques can be used to reduce the size of the strategy space (i.e., the size of the strategic-form matrix) and thus to ease the solution process.

⁸Unfortunately, we can find many examples of this situation in the history of mankind, such as the arms race between countries.

⁹In Section 3, we will show that the game-theoretic model and its solution changes if we consider a sequential move of the players (i.e., if player p_2 knows the move of player p_1 at the moment he makes a move).

2.3 Nash Equilibrium

In general, the majority of the games cannot be solved by the iterated dominance techniques. As an example, let us consider the Multiple Access Game introduced at the beginning. Each of the players has two possible strategies: either transmit (T) or not transmit (and thus to stay quiet) (Q). As the channel is shared, a simultaneous transmission of both players leads to a collision. The game is represented in strategic form in Table III.

$p_1 \backslash p_2$	Q	T
Q	$(0, 0)$	$(0, 1-c)$
T	$(1-c, 0)$	$(-c, -c)$

Table III. The Multiple Access Game in strategic form. The two moves for each player are: transmit (T) or be quiet (Q).

It can immediately be seen that no strategy is dominated in this game. To solve the game, let us introduce the concept of *best response*. If player p_1 transmits, then the best response of player p_2 is to be quiet. Conversely, if player p_2 is quiet, then p_1 is better off transmitting a packet. We can write $b_i(s_{-i})$, the best response of player i to an opponents' strategy vector s_{-i} as follows.

Definition 4 The best response $b_i(s_{-i})$ of player i to the profile of strategies s_{-i} is a strategy s_i such that:

$$b_i(s_{-i}) = \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}) \quad (3)$$

One can see, that if two strategies are mutual best responses to each other, then no player would have a reason to deviate from the given strategy profile. In the Multiple Access Game, two strategy profiles exist with the above property: (Q, T) and (T, Q) . To identify such strategy profiles in general, Nash introduced the concept of Nash equilibrium in his seminal paper [11]. We can formally define the concept of *Nash equilibrium (NE)* as follows.

Definition 5 The pure strategy profile s^* constitutes a Nash equilibrium if, for each player i ,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \forall s_i \in S_i \quad (4)$$

This means that in a Nash equilibrium, none of the users can unilaterally change his strategy to increase his utility. Alternatively, a Nash equilibrium is a strategy profile comprised of mutual best responses of the players.

A Nash equilibrium is *strict* [8] if we have:

$$u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*), \forall s_i \in S_i \quad (5)$$

It is easy to check that (D, D) is a Nash equilibrium in the Forwarder's Dilemma game represented in Table I. This corresponds to the solution obtained by iterated strict dominance. This result is true in general: Any solution derived by iterated strict dominance is a Nash equilibrium. The proof of this statement is presented in [6]. In the Multiple Access Game, however, the iterated dominance techniques do not help us derive the solutions. Fortunately, using the concept of Nash equilibrium, we can identify the two pure-strategy Nash

equilibria: (Q, T) and (T, Q) . Note that the best response $b_i(s_{-i})$ is not necessarily unique. For example in the Joint Packet Forwarding Game presented in Table II, player p_2 has two best responses (D or F) to the move D of player p_1 . Multiple best responses are the reason that the solutions of the iterated weak dominance technique in a given game might depend on the order of elimination.

2.4 Mixed Strategies

In the examples so far, we have considered only pure strategies, meaning that the players clearly decide on one behavior or another. But in general, a player can decide to play each of these pure strategies with some probabilities; in game-theoretic terms such a behavior is called a *mixed strategy*.

Definition 6 The mixed strategy $\sigma_i(s_i)$, or shortly σ_i , of player i is a probability distribution over his pure strategies $s_i \in S_i$.

Accordingly, we will denote the mixed strategy space of player i by Σ_i , where $\sigma_i \in \Sigma_i$. Hence, the notion of profile, which we defined earlier for pure strategies, is now characterized by the probability distribution assigned by each player to his pure strategies: $\sigma = \sigma_1, \dots, \sigma_{|P|}$, where $|P|$ is the cardinality of P . As in the case of pure strategies, we denote the strategy profile of the opponents by σ_{-i} . For a finite strategy space, i.e. for so called *finite games*¹⁰ [6] for each player, player i 's utility to profile σ is then given by:

$$u_i(\sigma) = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}) \quad (6)$$

Each of the concepts that we have considered so far for pure strategies can be also defined for mixed strategies. As there is no significant difference in these definitions, we refrain from repeating them for mixed strategies.

Let us first study the Multiple Access Game. We call x the probability with which player p_1 decides to transmit, and y the equivalent probability for p_2 (this means that p_1 and p_2 stay quiet with probability $1 - x$ and $1 - y$, respectively).

The payoff of player p_1 is:

$$u_1 = x(1 - y)(1 - c) - xyc = x(1 - c - y) \quad (7)$$

Likewise, we have:

$$u_2 = y(1 - c - x) \quad (8)$$

As usual, the players want to maximize their utilities. Let us first derive the best response of p_2 for each strategy of p_1 . In (8), if $x < 1 - c$, then $(1 - c - x)$ is positive, and u_2 is maximized by setting y to the highest possible value, namely $y = 1$. Conversely, if $x > 1 - c$, u_2 is maximized by setting $y = 0$ (these two cases will bring us back to the two pure-strategy Nash equilibria that we have already identified). More interesting is the last case, namely $x = 1 - c$, because here u_2 does not depend on y anymore (and is always equal to 0); hence, *any* strategy of p_2 (meaning any value of y) is a

¹⁰The general formula for infinite strategy space is slightly more complicated. The reader can find it in [6] or [13].




Figure 4. Best response functions in the Multiple Access Game. The best response function of player p_1 (x as a function of y) is represented by the dashed line; that of player p_2 (y as a function of x) is represented by the solid one. The two dots at the edges indicate the two pure-strategy Nash equilibria and the one in the middle shows the mixed-strategy Nash equilibrium.

best response. The game being symmetric, reversing the roles of the two players leads of course to the same result. This means that $(x = 1 - c, y = 1 - c)$ is a mixed-strategy Nash equilibrium for the Multiple Access Game.

We can graphically represent the best responses of the two players as shown in Figure 4. In the graphical representation, we refer to the set of best response values as the *best response function*¹¹. Relying on the concept of mutual best responses, one can identify the Nash equilibria as the crossing points of these best response functions.

Note that the number of Nash equilibria varies from game to game. There are games with no pure-strategy Nash equilibrium, such as the Jamming Game. We show the strategic form of this game in Table IV.

p_1 (sender) \ p_2 (jammer)	ch_1	ch_2
ch_1	(-1,1)	(1,-1)
ch_2	(1,-1)	(-1,1)

Table IV. The Jamming Game in strategic form.

The reader can easily verify that the Jamming Game cannot be solved by iterated strict dominance. Moreover, this game does not even admit a pure-strategy Nash equilibrium. In fact, there exists only a mixed-strategy Nash equilibrium in this game that dictates each player to play a uniformly random distribution strategy (i.e., select one of the channels with probability 0.5).

The importance of mixed strategies is further reinforced by the following theorem of Nash [11, 12]. This theorem is a crucial existence result in game theory. The proof uses the Brouwer-Kakutani fixed-point theorem and is provided in [7].

Theorem 1 (Nash, 1950) Every finite strategic-form game has a mixed-strategy Nash equilibrium.

2.5 Equilibrium Selection

As we have seen so far, the first step in solving a game is to investigate the *existence* of Nash equilibria. Theorem 1

¹¹Let us emphasize that, according to the classic definition of a function in calculus, the set of best response values does not correspond to a function, because there might be several best responses to a given opponent strategy profile.

states that in a broad class of games there always exists at least a mixed-strategy Nash equilibrium. However, in some cases, such as in the Jamming Game, there exists no pure-strategy Nash equilibrium. Once we have verified that a Nash equilibrium exists, we have to determine if it is a *unique* equilibrium point. If there is a unique Nash equilibrium, then we have to study its *efficiency*. Efficiency can also be used to select the most appropriate solutions from several Nash equilibria. *Equilibrium selection* means that the users have identified the desired Nash equilibrium profiles, but they also have to coordinate which one to choose. For example, in the Multiple Access Game both players are aware that there exist three Nash equilibria with different payoffs, but each of them tries to be “the winner” by deciding to transmit (in the expectation that the other player will be quiet). Hence, their actions result in a profile that is not a Nash equilibrium. The topic of equilibrium selection is one of the hot research fields in game theory [5, 15].

2.6 Pareto-optimality

So far, we have seen how to identify Nash equilibria. We have also seen that there might be several Nash equilibria, as in the Joint Packet Forwarding Game. One method for identifying the desired equilibrium point in a game is to compare strategy profiles using the concept of Pareto-optimality. To introduce this concept, let us first define Pareto-superiority.

Definition 7 *The strategy profile s is Pareto-superior to the strategy profile s' if for any player $i \in N$:*

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s'_{-i}) \quad (9)$$

with strict inequality for at least one player.

In other words, the strategy profile s is *Pareto-superior* to the strategy profile s' , if the utility of a player i can be increased by changing from s' to s without decreasing the utility of other players. The strategy profile s' is defined as *Pareto-inferior* to the strategy profile s . Note that the players might need to change their strategies simultaneously to reach the Pareto-superior strategy profile s .

Based on the concept of Pareto-superiority, we can identify the most efficient strategy profile or profiles.

Definition 8 *The strategy profile s^{po} is Pareto-optimal if there exists no other strategy profile that is Pareto-superior to s^{po} .*

In a Pareto-optimal strategy profile, one cannot increase the utility of player i without decreasing the utility of at least one other player. Using the concept of Pareto-optimality, we can eliminate the Nash equilibria that can be improved by changing to a more efficient (i.e. Pareto-superior) strategy profile. Note that we cannot define s^{po} as the strategy profile that is Pareto-superior to all other strategy profiles, because a game can have several Pareto-optimal strategy profiles. It is important to stress that a Pareto-optimal strategy profile is not necessarily a Nash equilibrium.

We can now use the concept of Pareto-optimality to study the efficiency of pure-strategy Nash equilibria in our running examples.

- In the Forwarder’s Dilemma game, the Nash equilibrium (D, D) is not Pareto-optimal. The strategy profile (F, F) is Pareto-optimal, but not a Nash equilibrium.
- In the Joint Packet Forwarding game, both strategy profiles (F, F) and (D, D) are Nash equilibria, but only (F, F) is Pareto-optimal.
- In the Multiple Access Game, both pure strategy profiles (T, Q) and (Q, T) are Nash equilibria and Pareto-optimal.
- In the Jamming game, there exists no pure-strategy Nash equilibrium, and all pure strategy profiles are Pareto-optimal.

We have seen that the Multiple Access Game has three Nash equilibria. We can notice that the mixed-strategy Nash equilibrium $\sigma = (p = 1 - c, q = 1 - c)$ results in the expected payoffs $(0, 0)$. Hence, this mixed-strategy Nash equilibrium is Pareto-inferior to the two pure-strategy Nash equilibria. In fact, it can be shown in general that there exist no mixed strategy profile that is Pareto-superior to all pure strategy profiles, because any mixed strategy of a player i is a linear combination of his pure-strategies with positive coefficients that sum up to one.

3 DYNAMIC GAMES

In the strategic-form representation it is usually assumed that the players make their moves *simultaneously* without knowing what the other players do. This might be a reasonable assumption in some problems, for example in the Multiple Access Game. In most of the games, however, the players might have a *sequential* interaction, meaning that the move of one player is conditioned by the move of the other player (i.e., the second mover knows the move of the first mover before making his decision). These games are called *dynamic games* [3] and we can represent them in an *extensive form*. We refer to a *game with perfect information*, if the players have a perfect knowledge of all previous moves in the game at any moment they have to make a new move.

3.1 Extensive Form with Perfect Information

In the *extensive form*, the game is represented as a tree, where the root of the tree is the start of the game and shown with an empty circle. We refer to one level of the tree as a *stage*. The nodes of a tree, denoted by a filled circle, show the possible unfolding of the game, meaning that they represent the sequence relation of the moves of the players. This sequence of moves defines a *path* on the tree and is referred to as the *history h* of the game. It is generally assumed that a single player can move when the game is at a given node.¹² This player is represented as a label on the node. Note that this is a tree, thus each node is a complete description of the path preceding it (i.e., each node has a unique history). The moves that lead to a given node are represented on each branch of the tree. Each terminal node (i.e., leaf) of the tree defines a

¹²Osborne and Rubinstein [13] define a game where a set of players can move in one node. Also, there exist specific examples in [6], in which different players move in the same stage. For the clarity of presentation, we do not discuss these specific examples in this tutorial.

potential end of the game called *outcome* and it is assigned the corresponding payoffs. In addition, we consider *finite-horizon games*, which means that there exist a finite number of stages.

Note that the extensive form is a more convenient representation, but basically every extensive form can be transformed to a strategic form and vice versa. However, extensive-form games can be used to describe *sequential* interactions more easily than strategic-form games. In extensive form, the *strategy* of player i assigns a move $m_i(h)$ to every non-terminal node in the game tree with the history h . For simplicity, we use pure strategies in this section. The definition of Nash equilibrium is basically the same as the one provided in Definition 5.

To illustrate these concepts, let us consider the *Sequential Multiple Access Game*. This is a modified version of the Multiple Access Game supposing that the two transmitters p_1 and p_2 are not perfectly synchronized, which means that p_1 always moves first (i.e., transmits or not) and p_2 observes the move of p_1 before making his own move.¹³ We show this extensive form game with perfect information in Figure 5. In this game, the strategy of player p_1 is to transmit (T) or to be quiet (Q). But the strategy of player p_2 has to define a move given the previous move for player p_1 . Thus, the possible strategies of p_2 are TT , TQ , QT and QQ , where for example TQ means that player p_2 transmits if p_1 transmits and he remains quiet if p_1 is quiet. Thus, we can identify the pure-strategy Nash equilibria in the Sequential Multiple Access Game. It turns out that there exist three pure-strategy Nash equilibria: (T, QT) , (T, QQ) and (Q, TT) .




Figure 5. The Sequential Multiple Access Game in extensive form.

Kuhn formulated a relevant existence theorem about Nash equilibria in finite extensive-form games in [9]. The intuition of the proof is provided in [6].

Theorem 2 (Kuhn, 1953) *Every finite extensive-form game of perfect information has a pure-strategy Nash equilibrium.*

The proof relies on the concept of backward induction, which we introduce in the following.

3.2 Backward Induction and Stackelberg Equilibrium

We have seen that there exist three Nash equilibria in the Sequential Multiple Access Game. For example, if player p_2

¹³In fact, this is called the *carrier sense* and it is the basic technique in the CSMA/CA protocols [14, 16].

plays the strategy TT , then the best response of player p_1 is to play Q . Let us notice, however, that the claim of player p_2 to play TT is an *incredible (or empty) threat*. Indeed, TT is not the best strategy of player p_2 if player p_1 chooses T in the first round.

We can eliminate equilibria based on such incredible threats using the technique of *backward induction*. Let us first solve the Sequential Multiple Access Game presented in Figure 5 with the backward induction method as shown in Figure 6.




Figure 6. The backward induction solution of the Sequential Multiple Access Game in extensive form.

The Sequential Multiple Access Game is a finite game with complete information. Hence, player p_2 knows that he is the player that has the last move. For each possible history, he predicts his best move. For example, if the history is $h = T$ in the game, then player p_2 concludes that the move Q results in the best payoff for him in the last stage. Similarly, player p_2 defines T as his best move following the move Q of player p_1 . In Figure 6, similarly to the examples in [7], we represent these best choices with thick solid lines in the last game row. Given all the best moves of player p_2 in the last stage, player p_1 calculates his best moves as well. In fact, each reasoning step reduces the extensive form game by one stage. Following this backward reasoning, we arrive at the beginning of the game (the root of the extensive-form tree). The continuous thick line from the root to one of the leaves in the tree gives us the backward induction solution. In the Sequential Multiple Access Game, we can identify the backward induction solution as $h = \{T, Q\}$. Backward induction can be applied to any finite game of perfect information. This technique assumes that the players can reliably forecast the behavior of other players and that they believe that the other can do the same. Note, however, that this argument might be less compelling for longer extensive-form games due to the complexity of prediction.

Note that the technique of backward induction is analogous to the technique of iterated strict dominance in strategic-form games. It is an elimination method for reducing the game. Furthermore, the backward induction procedure is a technique to identify *Stackelberg equilibria* in the extensive-form game. Let us call the first mover the *leader* and the second mover the *follower*.¹⁴ Then, we can define a Stackelberg equilibrium as follows.

¹⁴Note that in the general description of the Stackelberg game, there might be several followers, but there is always a single leader.

Definition 9 The strategy profile s is a Stackelberg equilibrium, with player p_1 as the leader and player p_2 as the follower, if player p_1 maximizes his payoff subject to the constraint that player p_2 chooses according to his best response function.

Let us now derive the Stackelberg equilibrium in the Sequential Multiple Access Game by considering how the leader p_1 argues. If p_1 chooses T , then the best response for p_2 is to play QQ or QT , which results in the payoff of $1 - c$ for p_1 . However, if p_1 chooses Q , then the best response of p_2 is TQ or TT , which results in the payoff of zero for leader p_1 . Hence, p_1 will choose T and (T, QT) or (T, QQ) are the Stackelberg equilibria in the Sequential Multiple Access Game. We can immediately establish the connection between this reasoning and the backward induction procedure.

We have seen in the above example that the leader can exploit his advantage if the two players have conflicting goals. In this game, the leader can enforce the equilibrium beneficial to himself.

Let us now briefly discuss the extensive form of the other three wireless networking examples with sequential moves. In the extensive-form version of the Forwarder's Dilemma, the conclusions do not change. Both players will drop each others' packets. In the extensive form of the Joint Packet Forwarding Game, if player p_1 chooses D , then the move of player p_2 is irrelevant. Hence by induction, we deduce that the Stackelberg equilibrium is (F, F) . Finally, in the Jamming Game, let us assume that p_1 is the leader and the jammer p_2 is the follower. In this case, the jammer can easily observe the move of p_1 and jam. Hence, being the leader does not necessarily result in an advantage. We leave the derivation of these claims to the reader as an exercise.

3.3 Imperfect Information and Subgame Perfect Equilibria

In this section, we will extend the notions of history and information. As we have seen, in the game with perfect information, the players always know the moves of all the other players when they have to make their moves. However, in the examples with simultaneous moves (e.g., the static games in Section 2), the players have an *imperfect information* about the unfolding of the game. To define perfect information more precisely, let us first introduce the notion of *information set* $h(n)$, i.e. the amount of information the players have at the moment they choose their moves in a given node n . The information set $h(n)$ is a partition of the nodes in the game tree. The intuition of the information set is that a player moving in n is uncertain if he is in node n or in some other node $n' \in h(n)$. We can now formally define the concept of *perfect information*¹⁵.

¹⁵Note that two well-established textbooks on game theory, [6] and [13], have different definitions of perfect information. We use the interpretation of [6], which we believe is more intuitive. The authors of [13] define, in Chapter 6 of their book, a game with simultaneous moves also as a game with perfect information, where the players are substituted with a set of players, who make their moves. Indeed, there seems to be no consensus in the research community either.

Definition 10 The players have a perfect information in the game if every information set is a singleton (meaning that each player always knows the previous moves of all players when he has to make his move).

It is not a coincidence that we use the same notation for the information set as for the history. In fact, the concept of information set is a generalized version of the concept of history.

To illustrate these concepts, let us first consider the extensive form of the original Multiple Access Game shown in Figure 7. Recall that this is a game with imperfect information. The dashed line represents the information set of player p_2 at the time he has to make his move. The set of nodes in the game tree circumvented by the dashed line means that player p_2 does not know whether player p_1 is going to transmit or not at the time he makes his own move, i.e. that they make simultaneous moves.




Figure 7. The original Multiple Access Game in extensive form. It is a game with imperfect information.

The *strategy* of player i assigns a move $m_i(h(n))$ to every non-terminal node n in the game tree with the information set $h(n)$. Again, we deliberately restrict the strategy space of the players to pure strategies, but the reasoning holds for mixed strategies as well [6, 13]. The possible strategies of each player in the Multiple Access Game are to transmit (T) or be quiet (Q). As we have seen before, both (T, Q) and (Q, T) are pure-strategy Nash equilibria. Note that in this game, player p_2 cannot condition his move on the move of player p_1 .

As we have seen in Section 3.2, backward induction and the concept of Stackelberg equilibrium can be used to eliminate incredible threats. Unfortunately, the elimination technique based on backward induction cannot always be used. To illustrate this, let us construct the game called *Multiple Access Game with Retransmissions* and solve it in the pure strategy space. In this game, the players play the Sequential Multiple Access Game, and they play the Multiple Access Game if there is a collision (i.e., they both try to transmit). We show the extensive form in Figure 8.

Note that the players have many more strategies than before. Player p_1 has four strategies, because there exist two information sets, where he has to move; and he has two possible moves at each of these information sets. For example, the strategy $s_1 = TQ$ means that player p_1 transmits at the beginning and he does not in the second Multiple Access Game. Similarly, player p_2 has $2^3 = 8$ strategies, but they are less trivial to identify. For example, each move in the strategy $s_2 = QTT$ means the following: (i) the first move




Figure 8. The Multiple Access Game with Retransmissions in extensive form. It is also a game with imperfect information.

means that player p_2 stays quiet if player p_1 transmitted, or (ii) p_2 transmits if p_1 was quiet and (iii) p_2 transmits in the last stage if they both transmitted in the first two stages. This example highlights an important point: *The strategy defines the moves for a player for every information set in the game, even for those information sets that are not reached if the strategy is played.* The common interpretation of this property is that the players may not be able to perfectly observe the moves of each other and thus the game may evolve along a path that was not expected. Alternatively, the players may have incomplete information, meaning that they have certain beliefs about the payoffs of other players and hence, they may try to solve the game on this basis. These beliefs may not be precise and so the unfolding of the game may be different from the predicted unfolding. Game theory covers these concepts in the notion of Bayesian games [6], but we do not present this topic in our tutorial due to space constraints.

It is easy to see that the Multiple Access Game with Retransmissions cannot be analyzed using backward induction, because the Multiple Access Game in the second stage is of imperfect information. To overcome this problem, Selten suggested the concept called *subgame perfection* in [17, 8]. In Figure 8, the Multiple Access Game in the second stage is a *proper subgame* of the Multiple Access Game with Retransmissions. Let us now give the formal definition of a proper subgame.

Definition 11 *The game G' is a proper subgame of an extensive-form game G if it consists of a single node in the extensive-form tree and all of its successors down to the leaves. Formally, if a node $n \in G'$ and $n' \in h(n)$, then $n' \in G'$. The information sets and payoffs of the subgame G' are inherited from the original game G : this means that n and n' are in the same information set in G' if they are in the same information set in G ; and the payoff function of G' is the restriction of the original payoff function to G' .*

Now let us formally define the concept of subgame perfection. This definition reduces to backward induction in finite games with perfect information.

Definition 12 *The strategy profile s is a subgame-perfect equilibrium of a finite extensive-form game G if it is a Nash equilibrium of any proper subgame G' of the original game G .*

One can check the existence of subgame perfect equilibria by applying the *one-deviation property*.

Definition 13 *The one-deviation property requires that there must not exist any information set, in which a player i can gain by deviating from his subgame perfect equilibrium strategy and conforming to it in other information sets.*

A reader somewhat familiar with dynamic programming may wonder about the analogy between the optimization in game theory and in dynamic programming [4]. Indeed, the one-deviation property corresponds to the *principle of optimality* in dynamic programming, which is based on backward induction. Hence, strategy profile s is a *subgame-perfect equilibrium* of a finite extensive-form game G if the one-deviation property holds.

Subgame perfection provides a method for solving the Multiple Access Game with Retransmissions. We can simply replace the Multiple Access Game subgame (the second one with simultaneous moves) with one of his pure-strategy Nash equilibria. Hence, we can obtain one of the game trees presented in Figure 9. Solving the reduced games with backward induction, we can derive the following solutions. In the game shown in Figure 9a, we have the subgame perfect equilibrium (QQ, TTT) . In Figure 9b we obtain the subgame perfect equilibria $(TT, Q*Q)$, where $*$ means any move from $\{T, Q\}$.

Because any game is a proper subgame of itself, a subgame-perfect equilibrium is necessarily a Nash equilibrium, but there might be Nash equilibria in G that are not subgame-perfect. In fact, the concept of Nash equilibrium does not require that the one-deviation property holds. We leave it to the reader as an exercise to verify that there are more Nash equilibria than subgame-perfect equilibria in the Multiple Access Game with Retransmissions.

The concept of subgame perfection has often been criticized with arguments based on equilibrium selection (recall the issue from Section 2.5). Many researchers point out that the players might not be able to determine how to play if several Nash equilibria exist in a given subgame. As an example, they might both play T in the Multiple Access Game with Retransmissions in the second subgame as well. This disagreement can result in an outcome that is not an equilibrium according to the definitions considered so far.

4 REPEATED GAMES

So far, we have assumed that the players interact only once and we modelled this interaction in a static game in strategic form in Section 2 and partially in Section 3. Furthermore, we have seen the Multiple Access Game with Retransmissions, which was a first example to illustrate repeated games, although the number of stages was quite limited. As we have seen in Section 3, the extensive form provides a more convenient representation for sequential interactions. In this




Figure 9. Application of subgame perfection to the Multiple Access Game with Retransmissions. In a) the proper subgame is replaced by one of the Nash equilibria of that game, namely (Q, T) . Solution b) represents the case, where the subgame is replaced by the other Nash equilibrium (T, Q) . The thick lines show the result of the backward induction procedure on the reduced game trees.

section, we assume that the players interact several times and hence we model their interaction using a *repeated game*. The analysis of repeated games in extensive form is basically the same as presented in Section 3, hence we focus on the strategic form in this section. To be more precise, we consider repeated games with *observable actions* and *perfect recall*: this means that each player knows all the moves of others, and that each player knows his own previous moves at each stage in the repeated game.

4.1 Basic Concepts

In repeated games, the players interact several times. Each interaction is called a *stage*. Note that the concept of stage is similar to the one in extensive form, but here we assume that the players make their moves simultaneously in each stage. The set of players is defined similarly to the static game presented in Section 2.1.

As a running example, let us consider the *Repeated Forwarder's Dilemma*, which consists of the repetition of the Forwarder's Dilemma stage game. In such a repeated game, all past moves are common knowledge at each stage t . The set of the past moves at stage t is commonly referred to as the history $h(t)$ of the game. We call it a history (and not an information set), because it is uniquely defined at the beginning of each stage. Let us denote the move of player i in stage t by $m_i(t)$. We can formally write the history $h(t)$ as follows:

$$h(t) = \{(m_1(t), \dots, m_{|N|}(t)), \dots, (m_1(0), \dots, m_{|N|}(0))\} \quad (10)$$

For example, at the beginning of the third stage of the Repeated Forwarder's Dilemma, if both players always cooperate, the history is $h(2) = \{(F, F), (F, F)\}$.

The *strategy* s_i defines a move for player i in the next stage $t + 1$ for each history $h(t)$ of the game¹⁶.

$$m_i(t + 1) = s_i(h(t)) \quad (11)$$

Note that the initial history $h(0)$ is an empty set. The strategy s_i of player i must define a move $m_i(0)$ for the initially empty history, which is called the *initial move*. For a moment, suppose that the Repeated Forwarder's Dilemma has two

stages. Then one example strategy of each player is $FFFFF$, where the entries of the strategy define the forwarding behavior for the following cases: (i) in the first stage, i.e. as an initial move, (ii) if the history was $h(1) = \{(F, F)\}$, (iii) if the history was $h(1) = \{(F, D)\}$, etc. As we can notice, the strategy space grows very quickly as the number of stages increases: In the two-stage Repeated Forwarder's Dilemma, we have $|S| = 2^5 = 32$ strategies for each player. Hence in repeated games, it is typically infeasible to make an exhaustive search for the best strategy and hence for Nash equilibria.

The utility in the repeated game might change as well. In repeated games, the users typically want to maximize their utility for the whole duration T of the game. Hence, they maximize:

$$u_i = \sum_{t=0}^T u_i(t, s) \quad (12)$$

In some cases, the objective of the players in the repeated game can be to maximize their payoffs only for the next stage (i.e., as if they played a static game). We refer to these games as *myopic games* as the players are short-sighted optimizers. If the players maximize their total utility during the game, we call it a *long-sighted game*.

Recall that we refer to a finite-horizon game if the number of stages T is finite. Otherwise, we refer to an infinite-horizon game. We will see in Section 4.3 that we can also model finite-horizon games with an unpredictable end as an infinite-horizon game with specific conditions.

4.2 Nash Equilibria in Finite-horizon Games

Let us first solve the finite Repeated Forwarder's Dilemma using the concept of Nash equilibrium. Assume that the players are long-sighted and want to maximize their total utility (the outcome of the game). As we have seen, it is computationally infeasible to calculate the Nash equilibria based on strategies that are mutual best responses to each other as the number of stages increases. Nevertheless, we can apply the concept of backward induction we have learned in Section 3.2. Because the game is of complete information, the players know the end of it. Now, in the last stage game, they both conclude that their dominant strategy is to drop the opponent's packet (i.e., to play D). Given this argument, their

¹⁶Recall that in the static game, the strategy was a single move.

best strategy is to play D in the penultimate stage. Following the same argument, this technique of backward induction dictates that the players should choose a strategy that plays D in every stage. Note that many strategies exist with this property.

In repeated games in general, it is computationally infeasible to consider all possible strategies for every possible history, because the strategy space increases exponentially with the length of the game. Hence, one usually restricts the strategy space to a reasonable subset. One widely-used family of strategies is the *strategies of history-1*. These strategies take only the moves of the opponents in the *previous stage* into account (meaning that they are “forgetful” strategies, because they “forget” the past behavior of the opponent). In the games we have considered thus far, we have two players and hence the history-1 strategy of player i in the repeated game can be expressed by the *initial move* $m_i(0)$ and the following strategy function:

$$m_i(t+1) = s_i(m_j(t)) \quad (13)$$

Accordingly, we can define the strategies in the Repeated Forwarder’s Dilemma as detailed in Table V. Note that these strategies might enable a feasible analysis in general, i.e., if there exists a large number of stages.

We can observe that in the case of some strategies, such as All-D or All-C, the players do not condition their next move on the previous move of the opponents. One refers to these strategies as *non-reactive strategies*. Analogously, the strategies that take the opponents’ behavior into account are called *reactive strategies* (for example, TFT or STFT).

Let us now analyze the Repeated Forwarder’s Dilemma assuming that the players use the history-1 strategies. We can conclude the same result as with the previous analysis.

Theorem 3 *In the Repeated Forwarder’s Dilemma, the strategy profile (All-D, All-D) is a Nash equilibrium.*

Although not proven formally, the justification of the above theorem is provided in [2].

4.3 Infinite-horizon Games with Discounting

In the game theory literature, infinite-horizon games with discounting are used to model a finite-horizon game in which the players are not aware of the duration of the game. Clearly, this is often the case in strategic interactions, in particular in networking operations. In order to model the unpredictable end of the game, one decreases the value of future stage payoffs. This technique is called *discounting*. In such a game, the players maximize their *discounted total utility*:

$$u_i = \sum_{t=0}^{\infty} u_i(t, s) \cdot \delta^t \quad (14)$$

where δ denotes the *discounting factor*. The *discounting factor* δ determines the decrease of the value for future payoffs, where $0 < \delta < 1$ (although in general, we can assume that δ

is close to one). The discounted total utility expressed in (14) is often normalized, and thus we call it the *normalized utility*:

$$u_i = (1 - \delta) \cdot \sum_{t=0}^{\infty} u_i(t, s) \delta^t \quad (15)$$

The role of the factor $1 - \delta$ is to let the stage payoff of the repeated game be expressed in the same unit as the static (stage) game. Indeed, with this definition, if the stage payoff $u_i(t, s) = 1$ for all $t = 0, 1, \dots$, then the normalized utility is equal to 1, because $\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$.

We have seen that the Nash equilibrium in the finite Repeated Forwarder’s Dilemma was a non-cooperative one. Yet, this rather negative conclusion should not affect our morale: in most networking problems, it is reasonable to assume that the number of iterations (e.g., of packet transmissions) is very large and a priori unknown to the players. Therefore, as discussed above, games are usually assumed to have an infinite number of repetitions. And, as we will see, infinitely repeated games can lead to more cooperative behavior.

Consider the history-1 strategies All-C and All-D for the players in the Repeated Forwarder’s Dilemma. Thanks to the normalization in (15), the corresponding normalized utilities are exactly those presented in Table I. A conclusion similar to the one we drew in Section 4.2 can be directly derived at this time. The strategy profile (All-D, All-D) is a Nash equilibrium: If the opponent always defects, the best response is All-D. A sketch of proof is provided (for the Prisoner’s Dilemma) in [6].

To show other Nash equilibria, let us first define the *Trigger* strategy. If a player i plays Trigger, then he forwards in the first stage and continues to forward as long as the other player j does not drop. As soon as the opponent j drops his packet, player i drops all packets for the rest of the game. Note that Trigger is not a history-1 strategy. The Trigger strategy applies the general technique of *punishments*.

If no player drops a packet, the payoffs corresponds to (F, F) in Table I, meaning that it is equal to $1 - c$ for each player. If a player i plays $m_i(t) = D$ at stage t , his payoff will be higher at this stage (because he will not have to face the cost of forwarding), but it will be zero for all the subsequent stages, as player j will then always drop. The normalized utility of player i will be equal to:

$$(1-\delta) [(1 + \delta + \dots + \delta^{t-1})(1 - c) + \delta^t \cdot 1] = 1 - c + \delta^t(c - \delta) \quad (16)$$

As $c < \delta$ (remember that, in general, c is very close to zero, whereas δ is very close to one), the last term is negative and the payoff is therefore smaller than $1 - c$. In other words, even a single defection leads to a payoff that is smaller than the one provided by All-C. Hence, a player is better off always forwarding in this infinite-horizon game, in spite of the fact that, as we have seen, the stage game only has (D, D) as an equilibrium point. It can be easily proven that (Trigger, Trigger) is a Nash equilibrium and that it is also Pareto-optimal (the intuition for the latter is the following: there is no way for a player to go above his normalized payoff of $1 - c$ without hurting his opponent’s payoff). Note that by similar arguments,

initial move, $m_1(0)$	$m_1(t) m_2(t) = F$	$m_1(t) m_2(t) = D$	strategy function s_1	name of the strategy
D	D	D	$m_1(t+1) = D$	Always Defect (All-D)
D	F	D	$m_1(t+1) = m_2(t)$	Suspicious Tit-For-Tat (STFT)
D	D	F	$m_1(t+1) = \overline{m_2(t)}$	Suspicious Anti Tit-For-Tat (SATFT)
D	F	F	$m_1(t+1) = F$	Suspicious Always Cooperate (S-All-C)
F	D	D	$m_1(t+1) = D$	Nice Always Defect (Nice-All-D)
F	F	D	$m_1(t+1) = \overline{m_2(t)}$	Tit-For-Tat (TFT)
F	D	F	$m_1(t+1) = m_2(t)$	Anti Tit-For-Tat (ATFT)
F	F	F	$m_1(t+1) = F$	Always Cooperate (All-C)

Table V. History-1 strategies of player 1 in the Repeated Forwarder's Dilemma. The entries in the first three columns represent: the initial move of player p_1 , a move of player p_1 to a previous move $m_2(t) = F$ of player p_2 , and the move of p_1 as a response to $m_2(t) = D$. The bar represents the alternative move (e.g., $\overline{F} = D$). As an example, let us highlight the TFT strategy, which begins the game with forwarding (i.e., cooperation) and then copies the behavior of the opponent in the previous stage.

one can show that (TFT, TFT) is also a Pareto-optimal Nash equilibrium, because it results in the payoff $1 - c$ for each of the players.

It is important to mention that the players cannot predict the end of the game and hence they cannot exploit this information. As mentioned in [6], reducing the information or the strategic options (i.e., decreasing his own payoff) of a player might lead to a better outcome in the game. This uncertainty is the real reason the cooperative equilibrium appears in the repeated version of the Forwarder's Dilemma game.

4.4 The Folk Theorem

We will now explore further the mutual influence of the players' strategies on their payoffs. We will start by defining the notion of *minmax value* (sometimes called the *reservation utility*). The minmax value is the lowest stage payoff that the opponents of player i can force him to obtain with punishments, provided that i plays the best response against them. More formally, it is defined as follows:

$$\underline{u}_i = \min_{s_{-i}} \left[\max_{s_i} u_i(s_i, s_{-i}) \right] \quad (17)$$

This is the lowest stage payoff that the opponents can enforce on player i . Let us denote by $s_{min} = \{s_{i,min}, s_{-i,min}\}$ the strategy profile for which the minimum is reached in (17). We call the $s_{-i,min}$ the *minmax profile* against player i within the stage game.

It is easy to see that player p_i can obtain at least his minmax value \underline{u}_i in any stage and hence we call *feasible payoffs* the payoffs higher than the minmax payoff. In the Repeated Forwarder's Dilemma, the feasible payoffs for any player p_1 are higher than 0. Indeed, by playing $s_1 = \text{All-D}$, he is assured to obtain at least that value, no matter what the strategy of p_2 can be. Similar argument applies to player p_2 . Let us graphically represent the feasible payoffs in Figure 10. We highlight the convex hull of payoffs that are strictly non-negative for both players as the set of feasible payoffs.

The notion of minmax that we have just defined refers to the stage game, but it has a very interesting application in the repeated game, as the following theorem shows.

Theorem 4 *Player i 's normalized payoff is at least equal to \underline{u}_i in any equilibrium of the infinitely repeated game, regardless of the level of the discount factor.*




Figure 10. The feasible payoffs in the Repeated Forwarder's Dilemma.

The intuition can be obtained again from the Repeated Forwarder's Dilemma: a player playing All-D will obtain a (normalized) payoff of at least 0. The theorem is proven in [6].

We are now in a position to introduce a fundamental result, which is of high relevance to our framework: the Folk Theorem¹⁷.

Theorem 5 (Folk Theorem) *For every feasible payoff vector $u = \{u_i\}_i$ with $u_i > \underline{u}_i$, there exists a discounting factor $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a Nash equilibrium with payoffs u .*

The intuition is that if the game is long enough (meaning that δ is sufficiently close to 1), the gain obtained by a player by deviating once is outweighed by the loss in every subsequent period, when loss is due to the punishment (minmax) strategy of the other players.

We have seen the application of this theorem in the infinite Repeated Forwarder's Dilemma. A player is deterred from deviating, because the short term gain obtained by the deviation (1 instead of $1 - c$) is outweighed by the risk of being minmaxed (for example using the Trigger strategy) by the other player (provided that $c < \delta$).

¹⁷This denomination of "folk" stems from the fact that this theorem was part of the oral tradition of game theorists, before it was formalized. Strictly speaking, we present the folk theorem for the discounting criterion. There exist different versions of the folk theorem, each of them is proved by different authors (as they are listed in [13] at end of Chapter 8).

5 DISCUSSION

One of the criticisms of game theory, as applied to the modelling of human decisions, is that human beings are, in practice, rarely fully rational. Therefore, modelling the decision process by means of a few equations and parameters is questionable. In wireless networks, the users do not interact with each other on such a fine-grained basis as forwarding one packet or access the channel once. Typically, they (or the device manufacturer, or the network operator, if any) program their devices to follow a protocol (i.e., a strategy) and it is reasonable to assume that they rarely reprogram their devices. Hence, such a device can be modelled as a rational decision maker. Yet, there are several reasons the application of game theory to wireless networks can be criticized. We detail them here, as they are usually never mentioned (for understandable reasons...) in research papers.

- Utility function and cost

The first issue is the notion of utility function: How important is it for a given user that a given packet is properly sent or received? This very much depends on the situation: the packet can be a crucial message, or could just convey a tiny portion of a figure appearing in a game. Likewise, the sensitivity to delay can also vary dramatically from situation to situation.

Similarly, the definition of cost might be a complex issue as well. In our examples (and in the state of the art in the application of game theory to wireless networks), the cost represents the energy consumption of the devices. In some cases, however, a device can be power-plugged, thus its “cost” could be neglected. Likewise, a device whose battery is almost depleted probably has a different evaluation of cost than when his battery is full. Furthermore, cost can include other considerations than energy, such as the previously mentioned delay or the consumed bandwidth. Expressing this diversity properly in the game models is still an open research issue.

- Pricing and mechanism design

Mechanism design is concerned with the question of how to lead the players to a desirable equilibrium by changing (designing) some parameters of the game. In particular, pricing is considered to be a good technique for regulating the usage of a scarce resource by adjusting the costs of the players. Many network researchers have contributed to this field. These contributions provide a better understanding of specific networking mechanisms. Yet it is not clear today, even for wired networks how relevant these contributions are going to be in practice. Usually the pricing schemes used in reality by operators are very coarse-grained, because operators tend to charge based on investment and personnel costs and on the pricing strategy of their competitors, and not on the instantaneous congestion of the network. If a part of the network is frequently congested, they will increase the capacity (deploy more base stations, more optical fibers, more switches) rather than throttle the user consumption by pricing.

Hence, the only area where pricing has practical relevance

is probably for service provisioning *among* operators (e.g., renting transmission capacity); but very little has been published so far on this topic.

- Infinite-horizon games

As mentioned, games in networking are usually assumed to be of infinite horizon, in order to capture the idea that a given player does not know when the interaction with another player will stop. This is, however, not perfectly true. For example, a given player could “know” that he is about to be turned off and moved away (e.g., its owner is about to finish a given session for which the player has been attached at a given access point). Yet we believe this not to be a real problem: indeed, the required “knowledge” is clearly related to the application layer, whereas the games we are considering involve networking mechanisms (and thus are typically related to the MAC and network layers).

- Discounting factor

As we have seen, in the case of infinitely repeated games, it is common practice to make use of the discounting factor. This notion comes from the application of game theory to economics: a given capital at time t_0 has “more value” than the same amount at a later time t_1 because, between t_0 and t_1 , this capital can generate some (hopefully positive) interest. At first sight, transposing this notion into the realm of networking makes sense: a user wants to send (or to receive) information as soon as he expresses the wish to do so.

But this may be a very rough approximation, and the comment we made about the utility function can be applied here as well: The willingness to wait before transmitting a packet heavily depends on the current situation of the user and on the content of the packet. In addition, in some applications such as audio or video streaming, the network can *forecast* how the demand will evolve.

A more satisfactory interpretation of the discounting factor in our framework is related to the uncertainty that there will be a subsequent iteration of the stage game, for example, connectivity to an access point can be lost. With this interpretation in mind, the discounting factor represents the probability that the current round is not the last one.

It is important to emphasize that the average discounted payoff is not the only way to express the payoff in an infinitely repeated game. Osborne and Rubinstein [13] discuss other techniques, such as “Limit of Means” and “Overtaking”. But, none of them captures the notion of users’ impatience, and hence we believe that they are therefore less appropriate for our purpose.

- Reputation

In some cases, a player can include the reputation of another player in order to anticipate his moves. For example, a player observed to be non-cooperative frequently in the past is likely to continue to be so in the future. If the game models individual packet transmissions, this attitude would correspond to the suspicion that another player has been programmed in a highly “selfish” way.

These issues go beyond the scope of this tutorial. For a discussion of these aspects, the reader is referred to [6], Chapter 9.

- Cooperative vs. non-cooperative players

In this tutorial, we assume that each player is a selfish individual, who is engaged in a non-cooperative game with other players. We do not cover the concept of cooperative games, where the players might have an agreement on how to play the game. Cooperative games include the issues of bargaining and coalition formation. These topics are very interesting and some of our problems could be modelled using these concepts. Due to space limitation, the reader interested in cooperative games is referred to [13].

- Information

In this paper, we study games with *complete information*. This means that each player knows the identity of other players, their strategy functions and the resulting payoffs or outcomes. In addition, we consider games with *observable actions* and *perfect recall*. In wireless networking, these assumptions might not hold: For example, due to the unexpected changes of the radio channel, a given player may erroneously reach the conclusion that another player is behaving selfishly. This can trigger a punishment (assuming there is one), leading to the risk of further retaliation, and so on. This means that, for any design of a self-enforcement protocol, special care must be devoted to the assessment of the amount and accuracy of the information that each player can obtain. The application of games with incomplete and imperfect information is an emerging field in wireless networking, with very few papers published so far.

6 CONCLUSION

In this tutorial, we have demonstrated how the non-cooperative game theory can be applied to wireless networking. Using four simple examples, we have shown how to capture wireless networking problems in a corresponding game, and we have analyzed them to predict the behavior of players. We have deliberately focused on the basic notions of non-cooperative game theory and have studied games with complete information. We have modelled devices as players, but there can be problems where the players are other participants, e.g. network operators. In addition, there exists another branch of game theory, namely cooperative game theory. We believe that, due to the distributed nature of wireless networking, this is a less appealing modelling tool in this context. Furthermore, we did not discuss the advanced topic of games with incomplete information, which are definitely a very compelling part of game theory. The purpose of this tutorial is rather to guide the interested readers familiar with computer science through the basics of non-cooperative game theory and to help them to bootstrap their own studies using this fascinating tool.

7 ACKNOWLEDGEMENTS

We would like to thank Tansu Alpcan, Mario Čagalj, Daniel Figueiredo, Jun Luo, Stefan Mangold, Peter Marbach and Panos Papadimitratos for their comments on this paper.

The work presented in this paper was partially supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

REFERENCES

- [1] E. Altman, T. Boulogne, R. E. Azouzi, T. Jimenez, and L. Wynter, "A survey on networking games in telecommunications," *Computers and Operations Research*, vol. 33, no. 2, pp. 286–311, Feb. 2006.
- [2] R. Axelrod, *The Evolution of Cooperation*. New York: Basic Books, 1984.
- [3] T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. London/New York: Academic Press, Jan. 1995.
- [4] R. E. Bellman, *Dynamic Programming*. Princeton, New Jersey: Princeton University Press, 1957.
- [5] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT Press, 1998.
- [6] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.
- [7] R. Gibbons, *A Primer in Game Theory*. Prentice Hall, 1992.
- [8] J. C. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
- [9] H. W. Kuhn, "Extensive games and the problem of information," *Contributions to the Theory of Games II (H. Kuhn and A. Tucker, Eds.)*, Princeton University Press, pp. 193–216, 1953.
- [10] A. B. MacKenzie, L. Dasilva, and W. Tranter, *Game Theory for Wireless Engineers*. Morgan and Claypool Publishers, 2006.
- [11] J. Nash, "Equilibrium points in n -person games," *Proceedings of the National Academy of Sciences*, vol. 36, pp. 48–49, 1950.
- [12] ———, "Non-cooperative games," *The Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951.
- [13] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA: The MIT Press, 1994.
- [14] T. S. Rappaport, *Wireless Communications: Principles and Practice (2nd Edition)*. Prentice Hall, 2002.
- [15] L. Samuelson, *Evolutionary Games and Equilibrium Selection*. MIT Press, 1997.
- [16] M. Schwartz, *Mobile Wireless Communications*. Cambridge Univ. Press, 2005.
- [17] R. Selten, "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit," *Zeitschrift für die gesamte Staatswissenschaft*, vol. 12, pp. 301–324, 1965.
- [18] H. A. Simon, *The Sciences of the Artificial*. MIT Press, 1969.
- [19] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton Univ. Press, Princeton, N.J., 1944.