Theoretical Computer Science 409 (2008) 341–350

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Games to induce specified equilibria ☆

Antoniy Ganchev [a], Lata Narayanan [a,*], Sunil Shende [b]

[a] Department of Computer Science, Concordia University, Montreal, Canada

[b] Department of Computer Science, Rutgers University, Camden, NJ 08102, USA

ARTICLE INFO

Article history:
Received 2 October 2007
Received in revised form 14 April 2008
Accepted 5 June 2008
Communicated by D. Peleg

ABSTRACT

Media access protocols in wireless networks require each contending node to wait for a *backoff time*, chosen randomly from a fixed range, before attempting to transmit on a shared channel. However, nodes acting in their own selfish interest may not follow the protocol. In this paper, a static version of the problem is modeled as a strategic game played by non-cooperating, rational players (the nodes). The objective is to design a game which exhibits a *unique*, *a priori* mixed-strategy Nash equilibrium. In the context of the media access problem, the equilibrium of the game would correspond to nodes choosing backoff times randomly from a given range of values, according to the given distribution. We consider natural variations of the problems concerning the number of actions available to the players and show that it is possible to design such a game when there are at least two players that each have the largest number of possible actions among all players. In contrast, we show that if there are exactly two players with different number of actions available to them, then it becomes impossible to design a strategic game with a unique such Nash equilibrium.

1. Introduction

A number of recent papers have tried to address the problem of selfishness of autonomous agents using the tools of game theory and algorithmic mechanism design. In this paper we are interested in the media access problem in a wireless network. In the IEEE 802.11 protocol, for instance, all nodes wishing to access a common link must follow an algorithm whereby each one chooses a random *backoff value* in a specified range. After waiting for the amount of time indicated by the backoff value, the node attempts to transmit. The node with the smallest backoff value, if it is unique, gains access to the medium. However, if two or more nodes attempt to transmit at the same time, a collision results. In the event of a collision, all colliding nodes double the range from which the backoff value was chosen, and retry. Nodes that did not collide keep the originally chosen backoff value, appropriately decremented, for the next round. Since nodes in a wireless network are autonomous agents, we cannot be sure that they will follow the protocol as specified. In particular, some nodes may try to cheat by always choosing a small backoff value, and getting an unfair share of access to the medium. If two or more nodes cheat simultaneously in this manner, then repeated collisions among cheating nodes may reduce the network throughput to zero, effectively making the network inoperative.

Our problem is just one instance of a general class of problems that are concerned with eliciting compliance to network layer protocols among nodes in the network [6]. Punishment-based approaches work by trying to isolate the misbehaving node [3,10,11]. In contrast, incentive-based or pricing-based approaches attempt to give some incentive to participating

☆ Research supported by NSERC, Canada. A preliminary version appears in the Proceedings of the 2nd International Workshop on Internet & Network Economics (WINE 2006).

* Corresponding author. Tel.: +1 514 848 2424.
E-mail addresses: a_ganche@cs.concordia.ca (A. Ganchev), lata@cs.concordia.ca (L. Narayanan), shende@crab.rutgers.edu (S. Shende).

0304-3975/$ – see front matter
doi:10.1016/j.tcs.2008.06.011

nodes to cooperate with the protocol [2,15]. For the media access problem, Kyasanur and Vaidya propose a modification to the 802.11 protocol, which supposes the presence of some trusted nodes [9]. In the modified protocol, instead of the sender choosing the backoff value, the receiver selects a random backoff value and sends it to the sender.

One approach is to use game theoretic tools [14] including various notions of *equilibria* to study the behavior of selfish agents in networks. The problem of media access in a wireless network has been modeled as a game by Cagalj et al. [5], where the authors show that non-cooperative behavior by more than one cheater can lead to network collapse. The equilibria of a game modeling an Aloha network with selfish users are analyzed by MacKenzie and Wicker [12]. In recent work, Chen et al. [4] describe the modeling of a media access protocol using random access games and initiate the study of their equilibria.

We are primarily concerned with an abstraction of the media-access problem: is it possible to design a strategic game with actions and utility functions that automatically induces protocol-compliant behavior among utility-maximizing selfish players? Since each player may have its own valuation of any given outcome, the utility function will include not only the agent's intrinsic valuation of the outcome, but also an incentive or payment that the mechanism will pay to the player to elicit honest play. Nisan and Ronen [13] introduced the term *algorithmic mechanism design* for their framework of studying algorithms that assume that the participants all act according to their own self-interest. Their model is specific to optimization problems, and much of the work that followed (for example, [7]) has focused on the same class of problems, and the mechanisms designed are the so-called VCG mechanisms, in which truth-telling is shown to be a dominant strategy for every player. In the wireless network setting, Anderegg and Eidenbenz propose a routing protocol for ad hoc networks, called Ad hoc VCG, which implements a VCG mechanism that is guaranteed to find the minimum energy path in the network [1]. As far as we know, there has been no work that uses a mechanism design approach to the wireless media access problem. In this paper, we demonstrate theoretical results that may have a bearing on designing such a mechanism for protocol compliance.

We first note that the media access problem, in its full generality, corresponds naturally to realizing a *dynamic* game; the nodes can (and do) modify their actions in response to the outcome of previous rounds. In this paper, we study a simpler abstraction: games that try to model *a single round* of the media access problem. To wit, the k nodes (the players) may be competing for access to the shared wireless medium using a backoff protocol where the jth node *should* choose a backoff value *uniformly at random* from a range given $[1, n_j]$ (the contention window). Our goal is to design the corresponding game, *i.e.* to specify utility functions for the players that would induce a *unique* mixed-strategy Nash equilibrium which corresponds *exactly* to each player faithfully following the protocol, *viz.* choosing a backoff value uniformly at random. We stress here that it is not difficult to construct a game with a mixed-strategy Nash equilibrium that corresponds to the uniform distribution (or indeed, any other distribution); the challenge lies in ensuring that this equilibrium is *unique* and not just one among many possible equilibria, so that rational play automatically leads to protocol compliance.

Hence, we start with the presumption that we are given an *a priori* distribution profile α^* that is desired. We wish to *design* a strategic game that realizes *exactly* this distribution as its unique mixed-strategy Nash equilibrium. We show that this is possible when there exist at least two players with the largest number of actions; we arrive at this result by generalizing the construction of a game achieving a desired equilibrium distribution among players with exactly the same number of actions. On the negative side, we prove that if we only have two players with different numbers of actions, then it is impossible to construct a game that realizes a given, full-support profile as its unique Nash equilibrium.

2. Preliminaries

For any fixed positive integer m, let $[1, m]$ denote the set of integers $\{1, 2, \dots, m\}$. We shall be concerned with two such sets that arise in our strategic games:

- A finite ordered set of $k \geq 2$ players, $P = [1, k]$.
- A finite set of $n_j \geq 1$ possible actions (or strategies), $A_j = [1, n_j]$, for each player $j \in P$.

We use the terminology **profile** for an ordered tuple that is typically indexed by an index set such as P. Following standard game-theoretic notation, an **outcome** of the game is represented by an *action profile*, $s = (i_j)_{j\in P}$, with the interpretation that every player $j \in P$ performs the corresponding action $i_j \in A_j$ in the outcome. The space of all possible outcomes is denoted by S.

A **utility function** is a function $u : A_1 \times A_2 \times \cdots \times A_k \to \mathbb{R}$ that associates the real value $u(s)$ with the action profile s. Every player has its own utility function. Collectively, the *utility function profile*, $(u_j)_{j\in P}$, is interpreted as follows: for any action profile s, the corresponding utility profile for the players is simply $(u_j(s))_{j\in P}$. We assume rational players that play independently (without any collusion) but seek to maximize their respective utilities, *i.e.* a player will always prefer an action $a \in A_j$ over some other action $b \in A_j$ if the corresponding utility is strictly higher.

We briefly review the notion of Nash equilibria for our games; for details, the reader is referred to Osborne and Rubinstein [14]. Given an action profile s, we denote by s_{-j} the partial profile containing the actions in s for all players except player j. For any player j and any partial profile s_{-j}, the set

$$B_j(s_{-j}) = \{a \in A_j : u_j(s_{-j}, a) \geq u_j(s_{-j}, b) \text{ for all } b \in A_j\}$$

is the set-valued **best-response function** for player j. A **pure-strategy Nash equilibrium** is a specific action profile $s^* = (i_j^*)_{j \in P}$ such that all the individual actions in s^* are *simultaneously* best responses for the players with respect to the corresponding partial action profiles. More formally, it is the case that

$$i_j^* \in B_j(s^*_{-j}).$$

Thus, in a pure-strategy Nash equilibrium, players cannot gain any utility by *unilaterally* changing their actions.

Not all finite games necessarily have pure-strategy Nash equilibria. However, they must have **mixed-strategy Nash equilibria** which generalize pure-strategy Nash equilibria. For a player $j \in P$, a **mixed strategy**, α_j, is a *discrete probability distribution* over its action set A_j. One interpretation of a mixed strategy is that the player j chooses any action $a \in A_j$ independently with the corresponding probability $\alpha_j(a)$. In particular, note that every outcome (*i.e.* action profile) of the game corresponds to a degenerate mixed strategy in which every player assigns probability 1 to its action in the outcome.

Collectively, the mixed strategies of all the players constitute a *distribution profile*, $\alpha = (\alpha_j)_{j \in P}$, with the following interpretation of utilities:

- Player j's utility for any particular pure strategy $a \in A_j$ is the expected value of the utility function u_j conditioned on the event that player j chooses strategy a. Let $U_j(\alpha, a)$ denote this expected value.
- Player j's overall utility for distribution profile α is given by

$$u_j(\alpha) = \sum_{a \in A_j} \alpha_j(a) U_j(\alpha, a). \tag{1}$$

We say an action $a \in A_j$ is in the *support* of a mixed strategy α_j for player j if $\alpha_j(a) > 0$. Furthermore, we say that $\alpha = (\alpha_j)_{j \in P}$ is a *full-support* distribution profile if for every player $j \in P$, the support of the distribution α_j is the entire action set A_j, *i.e.* $\alpha_j(a) > 0$ for all $j \in P$, $a \in A_j$. In this paper, we will be primarily interested in full-support distribution profiles.

A *mixed-strategy Nash equilibrium* is a special distribution profile, $\alpha^* = (\alpha_j^*)_{j \in P}$, with the property that a player *cannot* increase its (expected) utility by unilaterally changing its own distribution in the profile. In other words, every player's mixed strategy at equilibrium is a best response to the mixed strategies of the other players. We have assumed a finite set of actions for each player. This allows us to use a very useful alternative characterization of a mixed-strategy Nash equilibrium α^*:

> For every player j, *every action* in the support of the distribution α_j^* has the *same expected utility*:
>
> $$U_j(\alpha^*, a) = U_j(\alpha^*, b)$$
>
> for any two actions $a, b \in A_j$.

This characterization follows from the linearity of $u_j(\alpha^*)$ in α_j^* in Eq. (1) above: if $U_j(\alpha^*, a) = U_j(\alpha^*, b)$, it would be unilaterally profitable for player j to *shift* all the (positive) probability from action b to action a thus contradicting the status of α^* as being a Nash equilibrium. We do not provide a complete proof of this characterization; the interested reader is referred to Lemma 33.2 in [14] for details.

3. Designing games with identical player strategies

We first consider the situation where players have the same set of actions, $A = [1, n]$. We will use the index set $P = [1, k]$ for the players.[1] Suppose that we have an *a priori* known full-support distribution profile $\alpha^* = (\alpha_j^*)_{j \in P}$. We are interested in the following general question: is it possible to design a strategic game with a *unique* Nash equilibrium that is given by the profile α^*? In what follows, we will construct such a game.

For ease of description, we will treat both the index sets P and A as being *circularly ordered*, *i.e.* with player $k + 1$ being interpreted as player 1, with action $n + 1$ being interpreted as action 1, and so on. We design our game with the following property:

> The utility function for player $j \in P$ depends only on its own actions and those of its predecessor, player $j - 1$.

This property allows us to present the game using an abbreviated version of the usual strategic form of presentation. Kearns et al. study a graphical representation of games where each player is represented by a node and a utility matrix, and the utility to a player i is affected only by the actions of player i and its neighbors in the underlying undirected graph [8]. The game presented in this section is a graphical game but for a *directed* simple cycle on k nodes/players numbered from 1 through k: for every node/player i, there is a directed edge from i to node/player $(i + 1)$ with $k + 1$ interpreted as 1.

The utility function u_j, for each player $j \in P$, can be represented concisely by a two-dimensional matrix M_j with n rows and n columns. The interpretation of this matrix is as follows: $M_j(a, b)$ is the value of the utility function u_j when applied to *every action profile* s in which player $j - 1$ performs action b and player j performs action a. Thus, one thinks of the rows

[1] In principle, players need not have exactly the same set of actions; what matters is that they have the *same number* of possible actions.

Actions	Players 1	2	...	k
a				$\odot$
$a+1$	0			

a

Action	Players 1	2	...	j	$j+1$	...	k
a				$\odot$	0		

b

Key: The symbol $\odot$ indicates an α-deficient action.

Fig. 1. Implications for r values in Lemma 1.

of M_j as being indexed by the pure strategies of player j and the columns of M_j as being indexed by the pure strategies of player $j-1$.

For ease of description, the entries in matrix M_j can be specified over two steps. Let I_n be the identity matrix with n rows and columns. Consider the following matrix, V_n, obtained from I_n by shifting down (circularly) the rows of I_n:

$$V_n := \begin{bmatrix} 0 & \dots & 0 & 1 \\ 1 & \dots & 0 & 0 \\ & \dots & & \\ 0 & \dots & 1 & 0 \end{bmatrix}. \tag{2}$$

For every player $j \in P$, an intermediate matrix $\hat{M}_j$ is defined as follows:

$$\hat{M}_j = \begin{cases} V_n & \text{if } j = 1 \\ I_n & \text{otherwise.} \end{cases} \tag{3}$$

Now, let $\alpha^* = (\alpha_j^*)_{j\in P}$ be the desired unique mixed-strategy equilibrium profile. Then, the utility matrix M_j for player $j \in P$ is defined as follows. Recall that the columns of M_j correspond to actions of the previous player $j-1$ in the circular ordering of P. For any pair of actions $a, b \in A$, we have

$$M_j(a, b) = \hat{M}_j(a, b)/\alpha_{j-1}^*(b). \tag{4}$$

In other words, we obtain M_j from $\hat{M}_j$ by *scaling* each entry by the reciprocal of player $(j-1)$'s column probability for that column. Note that the matrices are well defined since we have assumed that α^* is a full-support distribution profile and therefore, the probability in the denominator of (4)'s right-hand side is always non-zero. We will now establish that the game defined above has a *unique* Nash equilibrium where player j chooses *exactly* the corresponding desired mixed strategy α_j^*.

Consider any mixed strategy, $\alpha = (\alpha_j)_{j\in P}$, for the game. Under this mixed strategy, player j's expected payoff for an action a is easily shown via (4) to be

$$U_\alpha^j(a) = \begin{cases} \alpha_k(a-1)/\alpha_k^*(a-1) & \text{if } j = 1 \\ \alpha_{j-1}(a)/\alpha_{j-1}^*(a) & \text{otherwise.} \end{cases} \tag{5}$$

Specifically, for the case when $\alpha = \alpha^*$, the right-hand side is identically equal to 1 for *all actions of all the players*, thus establishing that every player has equal payoffs for all its pure actions under distribution profile α^*. Thus, α^* is indeed a mixed-strategy Nash equilibrium for the game. It remains to show that this equilibrium is unique. We start with a useful definition specific to our game.

Definition 1. Let $\alpha = (\alpha_j)_{j\in P}$ be a mixed-strategy profile that *differs* from α^*. For a given action a and player j, let $r_j(a) = \alpha_j(a)/\alpha_j^*(a)$. We say that action a is *α-deficient* for player j if $r_j(a) < 1$.

Lemma 1. *Suppose that the profile $\alpha = (\alpha_j)_{j\in P}$ is a mixed-strategy Nash equilibrium for the game. Then, the following implications hold:*

(A) *If action a is α-deficient for player k, then $\alpha_1(a+1) = r_1(1) = 0$.*
(B) *For $1 \le j < k$, if action a is α-deficient for player j, then $\alpha_{j+1}(a) = r_{j+1}(a) = 0$.*

Proof. Let α be a Nash equilibrium for the game. To prove the first implication (Lemma 1(A) above), assume that the hypothesis holds. Both α_k and α_k^* being probability distributions, $\sum_{b\in A} \alpha_k(b) = 1 = \sum_{b\in A} \alpha_k^*(b)$, and from this it follows that, since $r_k(a) < 1$, there must be another action $b \neq a$ for which $r_k(b) > 1$.

Applying (5) above with $j = 1$, we conclude that under the mixed strategy α, player 1 will have a strictly larger payoff for playing the pure strategy $b+1$ as compared to playing the pure strategy $a+1$. Consequently, action $a+1$ cannot be in the support of the equilibrium strategy α_1 for player 1, and hence, $\alpha_1(a+1) = 0$ (which, in turn, implies that $r_1(a+1) = 0$). An almost identical argument works for the second part of the lemma except that we use (5) for the case when $1 < j \le k$. Fig. 1 is a schematic illustration of these implied dependencies among relevant r values for the players. □

Theorem 1. *The game outlined above is a k-player, n-strategy game that has the unique mixed-strategy Nash equilibrium given by the full-support distribution profile $(\alpha_j^*)_{j\in P}$.*

Proof. We have already shown that α^* is a Nash equilibrium for the game. To establish uniqueness, we will show that assuming a different equilibrium profile, $\alpha \neq \alpha^*$, yields a contradiction. Fig. 1 (based on Lemma 1) provides the intuition for this: it shows that if an action a is α-deficient for a player j then so is action a for player $j+1$, except in the case when $j = k$, and then it is action $a+1$ that is α-deficient for player 1.

More formally, let $\alpha \neq \alpha^*$. Then there must be some player j for whom there is an action a that is α-deficient. If player j is someone other than player k (*i.e.* where $1 \leq j < k$), then Lemma 1(B) implies that $\alpha_{j+1}(a) = 0$. In fact, we can apply Lemma 1(B) to obtain that $r_j(a) < 1$ implies $a_{j+1}(a) = 0$ and therefore $r_{j+1}(a) = 0 < 1$. By the same reasoning, and by repeatedly applying Lemma 1(B), we obtain that $\alpha_i(a) = 0, \forall i \in \{j+2, k\}$. Thus we deduce that player k must have an α-deficient action if α differs from α^*.

Without loss of generality, let a be an α-deficient action for player k such that action $a+1$ is *not* α-deficient for player k, *i.e.* with $r_k(a+1) \geq 1$. Now, Lemma 1(A) applies, and we get $\alpha_1(a+1) = 0$. Thus, action $(a+1)$ is α-deficient for player 1, and by applying Lemma 1(B) in succession $(k-1)$ times, we conclude that $\alpha_k(a+1) = 0$ (see Fig. 1). This contradicts our earlier assertion that action $a+1$ is not α-deficient for player k. Thus, contrary to assumption, α cannot differ from α^*; the game exhibits the unique Nash equilibrium profile α^*. □

We note that Theorem 1 can be easily specialized to the case that is relevant for the single round version of the medium access control problem, with all backoff values in the range $[1, n]$ and the desired distribution being the discrete uniform distribution for all players.

4. Designing games with non-identical player strategies

In this section, we consider games where the players do not have the same number of strategies available to them, but we still wish to achieve a target distribution profile that is a unique Nash equilibrium with full-support component distributions. A brief synopsis of the results in this section follows.

Using standard ideas from linear programming, we show in Section 4.1 that it is impossible to design such a game when there are only two players. With three or more players, however, it is indeed possible – under fairly non-restrictive conditions – to realize a game that has an *a priori* given distribution as its unique Nash equilibrium with full-support component distributions. In Section 4.2, we first show how such a game can be designed if the players can be partitioned into groups where each group contains at least two players having the same number of actions. More generally, we show that it suffices to have at least two players that have the maximum number of actions among all players; under this condition, a given profile is realizable as the unique Nash equilibrium of an appropriately designed game.

4.1. Games with two players

We assume, without loss of generality, that player 1 has the set of strategies $A_1 = [1, m]$ and player 2 has the set of strategies $A_2 = [1, n]$, with $m > n$. As before, we are given an *a priori* full-support distribution profile (α_1^*, α_2^*).

Theorem 2. *Given players* 1, 2, *and the full-support distribution profile* (α_1^*, α_2^*) *as described above, there is no two-player strategic game which can realize the given profile as its unique mixed-strategy Nash equilibrium.*

Proof. Suppose, to the contrary, that such a game can be realized with utility functions represented by the matrices M_1 and M_2 for players 1 and 2 respectively. We will assume, as before, that the rows are indexed by the corresponding player's actions, *i.e.* that $M_1(a, b)$ (respectively, $M_2(b, a)$) is the utility for player 1 (respectively, for player 2) when player 1's action is $a \in [1, m]$ and player 2's action is $b \in [1, n]$. Recall that $m > n$.

Since (α_1^*, α_2^*) is a mixed-strategy Nash equilibrium for the game, and since both the component distributions have full support, it follows that the expected utility of any distinct pair of pure strategies for player 1 (using matrix M_1) and for player 2 (using matrix M_2) are equal. More to the point, the distribution α_1^* satisfies the system of equations

$$\sum_{j=1}^{m} q_j[M_2(1, j) - M_2(i, j)] = 0 \quad \text{for } 2 \leq i \leq n$$

$$\sum_{j=1}^{m} q_j = 1 \tag{6}$$

and the distribution α_2^* satisfies the system of equations

$$\sum_{i=1}^{n} p_i[M_1(1, i) - M_1(j, i)] = 0 \quad \text{for } 2 \leq j \leq m$$

$$\sum_{i=1}^{n} p_i = 1. \tag{7}$$

Since the system of equations (6) has more variables than equations, it must either have no solutions or an infinite number of solutions. We know that it has at least one solution, *viz.* α_1^*. Hence, it must have infinitely many solutions and by convexity, at least one of these solutions must be a distribution with full support that differs from α_1^*. Let β be such a distribution. Then, it follows that *any* strategy by player 2 would be a best response to player 1 adopting the strategy β. Hence, α_2^* is a best response to β. Similarly, since α_2^* is a full-support distribution that satisfies the system of equations (7), we also conclude that any strategy, including strategy β, would be a best response by player 1 to player 2 adopting strategy α_2^*.

Hence, the game has at least two mixed-strategy Nash equilibria given by the profiles (α_1^*, α_2^*) and (β, α_2^*), a contradiction. □

4.2. Games with three or more players

We note that Theorem 2 implies that in order to be able to apply Theorem 1 to the case of two players, the target distribution profile must have full support for both players. From now on, we will assume that when a given distribution profile is to be realized, we limit the set of actions for any player to be exactly those actions that are in the support of the component that corresponds to the player's distribution. A simple corollary of Theorem 1 gives us the following result.

Theorem 3. *Consider a set of players P that can be partitioned into subsets $P_1, P_2, \ldots, P_m$, where $|P_j| > 1$ for all $1 \le j \le m$. If the set of strategies for all players in subset P_j is $[1, n_j]$, then for any given full-support distribution profile α^*, there exists a game whose unique Nash equilibrium is the profile α^*.*

Proof. We note that the case when the number of groups m equals 1 is the case handled by Theorem 1. More generally, let $|P_1| = k_1 > 1$. Then we create utility matrices for the players in set P_1 according to the k_1-player, n_1-strategy game in Theorem 1. Similarly we create utility matrices for the players in each subset P_j. Thus the utility for any player in P_j depends only on a single other player within the same subset P_j; the game is essentially partitioned into disjoint games, one for each subset of players, P_j. The topology of the graph corresponding to this game is a set of m disjoint cycles, each cycle corresponding to a set of players P_j. It is straightforward to see that the game represented by these utility matrices realizes the target profile as its unique Nash equilibrium. □

The conditions in the above theorem can be relaxed. We show that so long as there are at least *two* players with the (same) *largest* number of strategies, we can create a game corresponding to any *a priori* given, full-support distribution profile α^*. Our basic idea is to create utility matrices so that player j's utilities depend on the actions chosen by player $j-1$ and player k.

Consider the players arranged in non-decreasing order of the number of actions available to them. Let n_j be the number of actions available to player j. Then $n_{k-1} = n_k$; the last two players have the same number of actions. For the remaining players, we will make the simplifying assumption that for all $j \in [1, k-2]$, $n_j < n_{j+1}$. It will be obvious from the construction how this assumption can be relaxed.[2] The game itself is specified over two stages starting with unscaled utility matrices which will subsequently be scaled appropriately in a second stage. For convenience, we will assume that $n_0 = 0$ henceforth.

In the first stage, we will represent the utilities for j as simple unscaled matrices with player j's actions represented by the rows. Recall that for any $n \ge 1$, I_n is the identity matrix with n rows and n columns. The utility matrices are described below:

- Player 1's utilities only depend on player k's actions; the utility matrix (unscaled) is

  $$\hat{M}_1 = \left[V_{n_1} \mid 0 \right]$$

  where V_{n_1} is the identity matrix with its rows shifted down once (circularly). The 0 sub-matrix corresponds to actions $n_1 + 1, \ldots, n_k$ of player k.
- Player k's utilities only depend on player $(k-1)$'s actions; the utility matrix (unscaled) is

  $$\hat{M}_k = I_{n_{k-1}}.$$

  Note that $n_{k-1} = n_k$.
- For every other player j (hence, $2 \le j \le (k-1)$), the utilities depend both on player k as well as the previous player $(j-1)$. We represent the utilities as separate $(n_j \times n_{j-1})$ matrices for each of the actions $1, \ldots, n_k$ of player k. The matrices are divided into three groups; each matrix has an *upper sub-matrix* consisting of the first n_{j-1} rows and a *lower sub-matrix* consisting of the remaining $(n_j - n_{j-1})$ rows:
  - For player k's action $a \in [1, n_{j-2}] \cup [n_j + 1, n_k]$, the matrix is

    $$\hat{M}_j^a = \left[\begin{array}{cccc} & I_{n_{j-1}} & & \\ \hline 0 & 0 & \ldots & 1 \\ & & \ldots & \\ 0 & 0 & \ldots & 1 \end{array} \right].$$

    Every row in the lower sub-matrix is identical, and equals $[0, 0, \ldots, 1]$, the last row of the identity sub-matrix.

[2] Alternatively, if there are two or more players that have the same number of strategies, we can deal with this set of players in isolation using Theorem 3.

– For player k's action $a \in [n_{j-2}+1, n_{j-1}]$, the matrix is

$$\hat{M}_j^a = \left[\begin{array}{cccc} \multicolumn{4}{c}{I_{n_{j-1}}} \\ \hline x_j & x_j & \dots & 1+x_j \\ & & \dots & \\ x_j & x_j & \dots & 1+x_j \end{array}\right].$$

Every row in the lower sub-matrix is identical, and equals $[x_j, x_j, \dots, 1+x_j]$ for a value $x_j > 0$ to be determined later.

– Let $y_j > 0$ be a value to be determined. For action $(n_{j-1}+i)$ of player k, with $i \in [1, n_j - n_{j-1}]$, the corresponding matrix is

$$\hat{M}_j^{n_{j-1}+i} = \left[\begin{array}{c} I_{n_{j-1}} \\ \hline C_{i,j} \end{array}\right]$$

where $C_{i,j}$, the lower sub-matrix, has, as its ith row, the row vector
$[-(1+y_j), -(1+y_j), \dots, -(1+y_j), -y_j]$,
and whose remaining rows are all identically equal to the row vector
$[-y_j, -y_j, \dots, -y_j, 1-y_j]$.

As before, we obtain the actual utility matrices by scaling the above matrix entries by the reciprocals of the α^*-probabilities of the actions of the relevant players that influence any given entry. Thus, for instance, the scaled matrix M_j^a for any player $j \in [2, k-1]$ is obtained from the corresponding unscaled matrix $\hat{M}_j^a$ by multiplying each entry in column b by

$$\frac{1}{\alpha_{j-1}^*(b)\alpha_k^*(a)},$$

and so on. The scaled utility matrices define our game. The topology of the graph corresponding to this game is a directed cycle on k nodes, with additional edges from node k to every node in $[2, k-2]$.

Suppose that the players use a distribution profile α for their mixed strategies. As usual, let $r_j(a) = \alpha_j(a)/\alpha_j^*(a)$ be the ratio of the actual probability to the desired one for action a by player j. Some more notation comes in handy when describing various expected payoffs. Let $B_j = [n_{j-1}+1, n_j]$ denote the jth *block* of actions (recall that $n_0 = 0$). For any player $m \geq j$, we let

$$R_m^j = \sum_{a \in B_j} r_m(a)$$

denote the sum of the r_m values over the B_j-actions. Also, for any player j, the sum of its ratios over *all* its actions is given by

$$R_j = \sum_{a \in [1, n_j]} r_j(a).$$

Let us calculate the expected payoffs for player j when the players use mixed-strategy profile α. Then, the payoffs are as follows.

1. For player 1, the expected payoff for action $a \in [1, n_1]$ is

$$U_\alpha^1(a) = \begin{cases} r_k(a-1) & \text{when } 1 < a \leq n_1 \\ r_k(n_1) & \text{when } a = 1. \end{cases} \tag{8}$$

2. For player k, the expected payoff for action $a \in [1, n_k]$ is

$$U_\alpha^k(a) = r_{k-1}(a). \tag{9}$$

3. For every other player $j \in [2, k-1]$, it follows from the construction above that for an action $a \in [1, n_{j-1}]$, the identity sub-matrix (the upper sub-matrix) determines the expected payoff, which is simply

$$U_\alpha^j(a) = r_{j-1}(a)R_k \tag{10}$$

since the sub-matrix is repeated for all actions of player k. Next, consider the payoff for action $n_{j-1}+i$ of player j; this action corresponds to the ith row of each of the lower sub-matrices in the three groups. The respective contributions from the three groups of matrices appear below on separate lines in the expression for the expected payoff:

$$\begin{aligned} U_\alpha^j(n_{j-1}+i) \ &= r_{j-1}(n_{j-1})[R_k - R_k^{j-1} - R_k^j] + R_k^{j-1}[x_j R_{j-1} + r_{j-1}(n_{j-1})] \\ &\quad - y_j R_{j-1} R_k^j - R_{j-1} r_k(n_{j-1}+i) + R_k^j r_{j-1}(n_{j-1}). \end{aligned} \tag{11}$$

We now proceed to show that for appropriate values of x_j and y_j in the matrices above, the profile α^* is the unique Nash equilibrium for this game. Note that if the players use the mixed-strategy profile α^*, then every ratio $r_j(a)$ equals 1 and hence, by construction, the expected payoffs (under strategy α^*) for players 1 and k also equal 1 regardless of the action played.

For player $j \in [2, k-1]$, the payoff for any action in the upper sub-matrices is equal to $R_k = n_k$ from Eq. (10) since all ratios are 1. Observing Eq. (11), we can see that the contributions from all three groups of lower sub-matrices are independent of the action $n_{j-1} + i$ when the ratios are equal to 1; the net payoff, on simplification, can be seen to be

$$U^j_{\alpha^*}(n_{j-1} + i) = n_k + x_j n_{j-1}(n_{j-1} - n_{j-2}) - y_j n_{j-1}(n_j - n_{j-1}) - n_{j-1}.$$

Hence, equating the payoff for the upper sub-matrix actions to the right-hand side of the above equation gives us the following necessary and sufficient relationship between x_j and y_j *for α^* to be a Nash equilibrium*:

$$x_j(n_{j-1} - n_{j-2}) = y_j(n_j - n_{j-1}) + 1. \tag{12}$$

We can always choose $x_j, y_j > 0$ to ensure that Eq. (12) holds, *e.g.* when $x_j = 1$ and $y_j = (n_{j-1} - n_{j-2} - 1)/(n_j - n_{j-1})$.

Having established that α^* is a Nash equilibrium profile for our game for appropriate choices of x_j and y_j values as above, we now turn our attention to showing that the equilibrium is unique. We approach this in a spirit similar to that in Section 3, *i.e.* we first show that it suffices to consider distributions that differ in probability from the given distribution α^* for some actions of player k, and then show that this will lead to a contradiction.

Assume that our game has a Nash equilibrium α that differs from α^*. Recall that, by definition, the action a for any player j is α-deficient if and only if the ratio $r_j(a) = \alpha_j(a)/\alpha^*_j(a)$ is less than 1.

Lemma 2. *Suppose that the profile $\alpha = (\alpha_j)_{j \in P}$ is a mixed-strategy Nash equilibrium for the game. For $1 \le j < k$, if action a is α-deficient for player j, then $\alpha_{j+1}(a) = 0$.*

Proof. The proof is very similar to that of Lemma 1(B). If action a is α-deficient, then there must be another action b that is not α-deficient for player j. Hence, the expected utility to player $j + 1$ for action b will dominate the utility for action a via Eq. (10) above. Hence, if α is an equilibrium profile, it will assign zero probability to action a for player $j + 1$. □

Now, repeated applications of Lemma 2 allow us to conclude that a purported Nash equilibrium α that differs from α^* must witness some α-deficient action a for player k. In turn, this implies that player k has at least one action b for which $r_k(b) > 1$. Let A_d be the proper subset of $[1, n_k]$ that contains all the α-deficient actions for player k. We denote by $\bar{A_d}$ the complement of A_d in the set $[1, n_k]$. Recall that $B_j = [n_{j-1} + 1,\ n_j]$ is the jth block of actions (with $n_0 = 0$ tacitly). We have two exclusive (and exhaustive) cases to consider:

1. Within some block B_j of player k's actions, there is at least one action from each of the sets A_d and $\bar{A_d}$. Without loss of generality, let a and $a + 1$ be a pair of consecutive actions (circularly within block B_j) such that $a \in A_d$ and $a + 1 \in \bar{A_d}$. Then

$$r_k(a) < 1 \le r_k(a + 1)$$

by definition. Let us examine the difference in payoffs between actions $a + 1$ and a for player j. From Eq. (11), we observe for all actions within block B_j, only the contributions from the third group of utility matrices for player j differ. In fact, for consecutive actions like a and $a + 1$, it is only the matrices corresponding to actions a and $a + 1$ for player k that differ in their contribution. Upon simplification, we see that

$$\begin{aligned} U^j_\alpha(a) - U^j_\alpha(a + 1) &= R_{j-1}[r_k(a + 1) - r_k(a)] \\ &> 0 \end{aligned}$$

since R_{j-1} is positive and $r_k(a + 1) > r_k(a)$ by assumption. Consequently, if α is a Nash equilibrium profile as claimed, it will yield $\alpha_j(a + 1) = 0$. By repeatedly applying Lemma 2, we obtain the fact that $\alpha_k(a + 1) = 0$, which contradicts the assumed non-deficiency of action $a + 1$ for player k.

2. Every block of actions for player k either contains only actions from A_d or only actions from $\bar{A_d}$. Then there must be consecutive blocks B_{j-1} and B_j (in circular order of the blocks) such that block B_{j-1} contains only A_d-actions and block B_j contains only $\bar{A_d}$-actions or vice versa. Without loss of generality, consider the first possibility. Then by assumption, action n_{j-1}, the last action in block B_{j-1}, is deficient while action $n_{j-1} + 1$, the first action in block B_j, is not deficient. We now compare the expected payoffs for player j for these two actions. Observing Eq. (11), we see that while the first group of matrices does not contribute to any difference in payoffs, the latter two groups do provide non-zero contributions. Since block B_{j-1} is entirely deficient for player k, it must be the case that

$$\begin{aligned} R^{j-1}_k &= \sum_{b \in B_{j-1}} r_k(b) \\ &< \mid B_{j-1} \mid \\ &= (n_{j-1} - n_{j-2}). \end{aligned}$$

Similarly, $R^j_k \ge (n_j - n_{j-1})$ and $r_k(n_{j-1} + 1) \ge 1$ since block B_j is assumed to be wholly non-deficient.

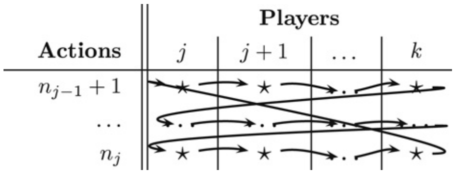

Fig. 2. Dependencies for r values in various blocks.

Putting these inequalities together, we can simplify and show that for action $a = n_{j-1}$

$$\begin{aligned} U_\alpha^j(a) - U_\alpha^j(a+1) &= R_{j-1}[y_j R_k^j + r_k(n_{j-1}+1) - x_j R_k^{j-1}] \\ &> R_{j-1}[y_j(n_j - n_{j-1}) + 1 - x_j(n_{j-1} - n_{j-2})] \\ &= 0 \end{aligned}$$

where the last step above follows directly from Eq. (12) that holds for the chosen values of x_j and y_j. Again, we conclude first that $\alpha_j(a+1) = \alpha_j(n_{j-1}+1) = 0$ and hence, that repeated applications of Lemma 2 imply $\alpha_k(n_{j-1}+1) = 0$. This violates our assumption that action $n_{j-1}+1$ is non-deficient for player k.

A similar argument works as well for the symmetric case where every B_j-action is deficient and every B_{j-1}-action is not; we will not repeat the details. An illustration of ripple-effect of an α-deficient action causing further deficiencies is given in Fig. 2; the two parts of the figure respectively show the dependencies for the action block B_j (with $1 \le j \le k-2$) and action block B_k.

We have thus established the following result:

Theorem 4. *Suppose we are given k players with action sets* $[1, n_j]$ *for every player j such that* $n_1 < n_2 < \cdots < n_{k-1} = n_k$. *Then given any full-support distribution profile* α^* *for this collection of players and actions, there is a game whose unique Nash equilibrium is* α^*.

It is easy to see that the construction described in this section can be easily extended to the case when $n_1 \le n_2 \le \cdots \le n_{k-1} = n_k$. For instance, a game can be constructed such that, for all players j with $n_j = n_{j-1}$, the utilities depend only on the actions of players j and $j-1$ as in the proof of Theorem 1, and for other players, the utilities are defined as in the proof of Theorem 4. Alternatively, as mentioned earlier, if there are two or more players that have the same number of strategies, we can deal with this set of players in isolation using the techniques of Theorem 3. Using either method, we can derive a game for the case where possibly multiple players have action set $[1, n_j]$ for any j; the game will have an *a priori* distribution as the only possible Nash equilibrium.

5. Concluding remarks

In this paper, we have shown that, under certain conditions, it is possible to reverse-engineer a game that has as its unique Nash equilibrium a certain given distribution profile of the players' actions. Our motivation was the static media-access problem in wireless networks, where the given profile corresponds to the random choices expected from the nodes (the players) participating in a single round of a typical backoff protocol. In other words, protocol compliance in that round can be achieved if the players play the designed game.

Of course, what is more desirable is a specific mechanism for *enforcing* protocol compliance. Any such mechanism will need to take into account domain-specific details such as private types/valuations of the players, incentive payments, and cost recovery by the network authority. Also, in practice, wireless media access more closely resembles a dynamic game; in the event of a collision, all colliding nodes double their values of contention window, and retry. Nodes that did not collide keep the originally chosen backoff value, appropriately decremented, for the next round. As a result of this, in any given round, nodes not only have different contention window values, but the backoff value is related to the history of previous rounds. In future work, we expect to address these issues in a mechanism design framework.

Acknowledgments

We wish to thank the anonymous referees for their detailed and insightful comments that have been incorporated into the paper.

References

[1] L. Anderegg, S. Eidenbenz, Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents, in: Proceedings of the 9th Annual International Conference on Mobile Computing and Networking, MobiCom '03, September 2003, pp. 245–259.

[2] L. Buttyan, J.-P. Hubaux, Nuglets: A virtual currency to stimulate cooperation in self-organized mobile Ad Hoc networks, Technical Report DSC/2001, Swiss Federal Institue of Technology, 2001.

[3] S. Buchegger, J.Y. Le Boudec, Performance analysis of the CONFIDANT protocol, in: Proceedings of MOBIHOC, June 2002, pp. 226–236.
[4] L. Chen, T. Cui, S.H. Low, J.C. Doyle, A game-theoretic model for medium access control, in: Proceedings of WICON, 2007.
[5] M. Cagalj, S. Ganeriwal, I. Aad, J.-P. Hubaux, On selfish behavior in CSMA/CA networks, in: Proceedings of IEEE INFOCOM 2005, vol. 4, March 2005, pp. 2513–2524.
[6] M. Conti, E. Gregori, G. Maselli, Cooperation issues in mobile ad hoc networks, in: Proceedings of the 24th International Conference on Distributed Computing Systems Workshops - W7: EC, ICDCSW'04, March 2004, pp. 803–808.
[7] J. Feigenbaum, C. Papadimitriou, R. Sami, S. Shenker, A BGP-based mechanism for lowest-cost routing, in: Proceedings of the 21st Symposium on Distributed Computing, 2002, pp. 173–182.
[8] M. Kearns, M. Littman, S. Singh, Graphical models for game theory, in: Proceedings of UAI 2001, 2001.
[9] P. Kyasanur, N. Vaidya, Selfish MAC layer misbehavior in wireless networks, IEEE Transactions on Mobile Computing 4 (5) (2005) 502–516.
[10] S. Marti, T. Guili, K. Lai, M. Baker, Mitigating routing misbehavior in mobile ad hoc networks, in: Proceedings of ACM MOBICOM, 2000, pp. 255–265.
[11] P. Michiardi, R. Molva, CORE: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks, in: Proceedings of Communications and Multimedia Security, September 2002, pp. 107–121.
[12] A.B. MacKenzie, S.B. Wicker, Selfish users in Aloha: A game theoretic approach, in: Proceedings of IEEE Vehicular Technology Conference, October 2001, pp. 1354–1357.
[13] N. Nissan, A. Ronen, Algorithmic mechanism design, in: Proceedings of the thirty-first annual ACM symposium on Theory of computing, STOC '99, May 1999, pp. 129–140.
[14] M. Osborne, A. Rubinstein, A Course in Game Theory, The MIT Press, 1996.
[15] S. Zhong, J. Chen, Y.R. Yang, Sprite: A simple, cheat-proof, credit-based system for mobile ad hoc networks, in: Proceedings of IEEE INFOCOM 2003, vol. 3, March/April 2003, pp. 1987–1997.