Potential Functions in Strategic Games ⋆

Paul G. Spirakis[1,2] and Panagiota N. Panagopoulou[2]

[1] Computer Engineering and Informatics Department, University of Patras
[2] Computer Technology Institute & Press "Diophantus"
spirakis@cti.gr, panagopp@cti.gr

Abstract. We investigate here several categories of strategic games and antagonistic situations that are known to admit potential functions, and are thus guaranteed to either possess pure Nash equilibria or to stabilize in some form of equilibrium in cases of stochastic potentials. Our goal is to indicate the generality of this method and to address its limits.

1 Introduction

A *strategic game* is a model of interactive decision making, helping us in analyzing situations in which two or more individuals, called *players*, make decisions (or choose *actions*) that will influence one another's welfare. The most important solution concept of a strategic game is the well-known *Nash equilibrium* [14], which captures a steady state of the game, in the sense that no player has an incentive to change her action if all the other players preserve theirs. The classical theorem of Nash [14] proves that every finite game has a *randomized* Nash equilibrium, i.e., there exists a combination of *mixed strategies*, one for each player, such that no player can increase her expected payoff by unilaterally deviating. A mixed strategy for a player is actually a probability distribution over the set of her available actions. However, Nash's proof of existence is non-constructive, and the problem of computing a Nash equilibrium has been identified among the "inefficient proofs of existence" [15], and it was eventually shown to be complete for the complexity class PPAD [2].

The apparent difficulty of computing a randomized Nash equilibrium, together with the fact that the concept of non-deterministic strategies has received much criticism in game theory, raises the natural question of what kind of games possess a *pure* Nash equilibrium, i.e., a Nash equilibrium where each player chooses deterministically one of her available actions. In this paper we discuss strategic games that are guaranteed to have pure Nash equilibria via the existence of *potential functions*, which enable the application of optimization theory to the study of pure Nash equilibria. More precisely, a potential function for a game is a real-valued function, defined on the set of possible action

⋆ This research has been co-financed by the European Union (European Social Fund – ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) – Research Funding Program: Thales. Investing in knowledge society through the European Social Fund.

combinations (or *outcomes*) of the game, such that the pure equilibria of the game are precisely the local optima of the potential function. If a game admits a potential function, there are nice consequences for the existence and tractability of pure Nash equilibria. In particular, if the game is finite, i.e., the player set and actions sets are finite, then the potential function achieves a global optimum, which is also a local optimum, and hence the game has at least one equilibrium in pure strategies. Going one step further, in any such game, *the Nash dynamics converges*. This means that the directed graph with outcomes as nodes and payoff-improving defections by individual players as edges has no cycles, implying that it has a sink corresponding to a pure Nash equilibrium, and that if we start with an arbitrary action combination and let one player at a time perform an improvement step, i.e., change her action to increase her payoff, then such a sink will be eventually reached.

We survey here several categories of strategic games that are known to admit potential functions (or some kind of generalization of a potential). Our goal is to give some insight on both the characterization of games guaranteed to possess pure Nash equilibria and the complexity of reaching such an equilibrium via a sequence of selfish payoff-improving steps. We note that potential games are not the only class of games known to possess pure Nash equilibria; another well-studied one is the class of *games of strategic complementarities*, introduced in [21]. Roughly speaking, a game has strategic complementarities if there is an order on the set of the players' pure strategies such that an increase in one player's strategy makes the other players want to increase their strategies as well. Based on a fixpoint theorem due to Tarski [20], games of strategic complementarities are known to possess at least one equilibrium in pure strategies, and furthermore, the set of pure Nash equilibria has a certain order structure. Uno [22] revealed an important relation between potential games and games with strategic complementarities: any finite game with strategic complementarities admits a *nested* pseudo-potential (which are extensions of potentials). A couple of natural questions are raised here: (1) How exactly are these two classes of games related, i.e., does the existence of (even a generalized notion of) a potential function is equivalent to the existence of strategic complementarities? (2) Are these two notions necessary in order to characterize the existence of pure Nash equilibria, or do there exist classes of games neither admitting any kind of potential nor having complementarities that are nevertheless guaranteed to have pure equilibria? The answer to these questions could give us more insight in order to better understand the nature of games that have pure Nash equilibria.

2 Games and potential functions

Strategic games and Nash equilibria. A *game* refers to any situation in which two or more decision-makers interact. We focus on *finite games in strategic form*: A **finite strategic form game** is any Γ of the form $\Gamma = \langle N, (C_i)_{i\in N}, (u_i)_{i\in N}\rangle$ where N is a finite nonempty set, and, for each $i \in N$, C_i is a finite nonempty set and u_i is a function from $C = \times_{j\in N} C_j$ into the set of real numbers $\mathbb{R}$. In the above definition, N is the set of players in the game Γ. For each player $i \in N$, C_i is the set of *actions* available to player i. When the strategic form game Γ is played, each player i must choose one of the actions in the set C_i. For each combination of actions, or *pure strategy profile* $c = (c_j)_{j \in N} \in C$ (specifying one action for each player), the number $u_i(c)$ represents the *payoff* that player i would get in this game if c were the combination of actions implemented by the players. When we study a strategic form game, we assume that all the players choose their actions simultaneously.

A *pure Nash equilibrium* of a strategic game is a combination of actions, one for each player, from which no player has an incentive to unilaterally deviate.

Definition 1. *A* ***pure Nash equilibrium*** *of game* $\Gamma = \langle N, (C_i)_{i \in N}, (u_i)_{i \in N} \rangle$ *is a pure strategy profile* $c = (c_j)_{j \in N} \in C$ *such that, for all players* $i \in N$*,*

$$u_i(c) \geq u_i(c_i', c_{-i}) \quad \forall c_i' \in C_i \ .$$

Not all games are guaranteed to possess a pure Nash equilibrium; however, if we extend the strategy set of each player to include any probability distribution on her set of actions, then a (*mixed strategy*) Nash equilibrium is guaranteed to exist [14].

Potential games. Potential games, defined in [13], are games with the property that the incentive of all players to unilaterally deviate from a pure strategy profile can be expressed in one global function, the potential function.

Fix an arbitrary game in strategic form $\Gamma = \langle N, (C_i)_{i \in N}, (u_i)_{i \in N} \rangle$ and some vector $\mathbf{b} = (b_1, \ldots, b_{|N|}) \in \mathbb{R}_{>0}^{|N|}$. A function $P : C \to \mathbb{R}$ is called

- An *ordinal potential* for Γ if, $\forall c \in C$, $\forall i \in N$, $\forall a \in C_i$,

$$P(a, c_{-i}) - P(c) > 0 \Longleftrightarrow u_i(a, c_{-i}) - u_i(c) > 0 \ . \tag{1}$$

- A $\mathbf{b}$-*potential* for Γ if, $\forall c \in C$, $\forall i \in N$, $\forall a \in C_i$,

$$P(a, c_{-i}) - P(c) = b_i \cdot (u_i(a, c_{-i}) - u_i(c)) \ . \tag{2}$$

- An *exact potential* for Γ if it is a $\mathbf{b}$-potential for Γ where $b_i = 1$ for all $i \in N$.

It is straightforward to see that the existence of an ordinal, exact, or $\mathbf{b}$-potential function P for a finite game Γ guarantees the existence of at least one pure Nash equilibrium in Γ: each local optimum of P corresponds to a pure Nash equilibrium of Γ and vice versa. Thus the problem of finding pure Nash equilibria of a potential game Γ is equivalent to finding local optima for the optimization problem with state space the pure strategy space C of the game and objective the potential function of the game.

Furthermore, the existence of a potential function P for a game $\Gamma = \langle N, (C_i)_{i \in N}, (u_i)_{i \in N} \rangle$ implies a straightforward algorithm for constructing a pure Nash equilibrium of Γ: The algorithm starts from an arbitrary strategy profile $c \in C$ and, at each step, one single player performs a *selfish step*, i.e., switches to a pure strategy that strictly improves her payoff. Since the payoff of the player increases, P increases as well. When no move is possible, i.e., when a pure strategy profile $\hat{c}$ is reached from which no player has an incentive to unilaterally deviate, then $\hat{c}$ is a pure Nash equilibrium and a local optimum of P. This procedure however does not imply that the computation of a pure Nash equilibrium can be done in polynomial time, since the improvements in the potential can be very small and too many.

3 Congestion games and selfish routing

Congestion games, introduced in [19], are games in which each player chooses a particular subset of resources out of a family of allowable subsets for her (her strategy set), constructed from a basic set of primary resources for all the players. The *delay* associated with each primary resource is a non-decreasing function of the number of players who choose it, and the total delay received by each player is the sum of the delays associated with the primary resources she chooses. In [19] it was shown that any game in this class possesses at least one Nash equilibrium in pure strategies, and this result follows from the existence of a potential function. Later, Monderer and Shapley [13] showed that every finite potential game is *isomorphic* to a congestion game.

3.1 Congestion Games

A *congestion model* $\langle N, E, (\Pi_i)_{i\in N}, (d_e)_{e\in E}\rangle$ is defined as follows. N denotes the set of players $\{1, \ldots, n\}$. E denotes a finite set of resources. For $i \in N$ let Π_i be the set of strategies of player i, where each $\varpi_i \in \Pi_i$ is a nonempty subset of resources. For $e \in E$ let $d_e : \{1, \ldots, n\} \to \mathbb{R}$ denote the delay function, where $d_e(k)$ denotes the cost (e.g. delay) to each user of resource e, if there are exactly k players using e. The *congestion game* associated with this congestion model is the game in strategic form $\langle N, (\Pi_i)_{i\in N}, (u_i)_{i\in N}\rangle$, where the payoff functions u_i are defined as follows: Let $\Pi \equiv \times_{i\in N}\Pi_i$. For all $\varpi = (\varpi_1, \ldots, \varpi_n) \in \Pi$ and for every $e \in E$ let $\sigma_e(\varpi)$ be the number of users of resource e according to the configuration ϖ: $\sigma_e(\varpi) = |\{i \in N : e \in \varpi_i\}|$. Define $u_i : \Pi \to \mathbb{R}$ by $u_i(\varpi) = -\sum_{e\in\varpi_i} d_e(\sigma_e(\varpi))$. In a *network congestion game* the families of subsets Π_i are represented implicitly as paths in a network. We are given a directed network $G = (V, E)$ with the edges playing the role of resources, a pair of nodes $(s_i, t_i) \in V \times V$ for each player i and the delay function d_e for each $e \in E$. The strategy set of player i is the set of all paths from s_i to t_i. If all origin-destination pairs (s_i, t_i) of the players coincide with a unique pair (s, t) we have a *single-commodity network congestion game* and then all users share the same strategy set, hence the game is symmetric.

In a *weighted congestion model* we allow the users to have different demands, and thus affect the resource delay functions in a different way, depending on their own weights. The *weighted congestion game* associated with a weighted congestion model is the game in strategic form $\langle (w_i)_{i\in N}, (\Pi_i)_{i\in N}, (u_i)_{i\in N}\rangle$, where the payoff functions u_i are defined as follows. For any configuration $\varpi \in \Pi$ and for all $e \in E$, let $\Lambda_e(\varpi) = \{i \in N : e \in \varpi_i\}$ be the set of players using resource e according to ϖ. The cost $\lambda^i(\varpi)$ of user i for adopting strategy $\varpi_i \in \Pi_i$ in a given configuration ϖ is equal to the cumulative delay $\lambda_{\varpi_i}(\varpi)$ on the resources that belong to ϖ_i: $\lambda^i(\varpi) = \lambda_{\varpi_i}(\varpi) = \sum_{e\in\varpi_i} d_e(\theta_e(\varpi))$, where, for all $e \in E$, $\theta_e(\varpi) \equiv \sum_{i\in\Lambda_e(\varpi)} w_i$ is the load on resource e with respect to the configuration ϖ. The payoff function for player i is then $u_i(\varpi) = -\lambda^i(\varpi)$. A configuration $\varpi \in \Pi$ is a pure Nash equilibrium if and only if, for all $i \in N$, $\lambda_{\varpi_i}(\varpi) \leq \lambda_{\pi_i}(\pi_i, \varpi_{-i}) \quad \forall \pi_i \in \Pi_i$, where (π_i, ϖ_{-i}) is the same configuration as ϖ except for user i that has now been assigned to path π_i. In a *weighted network congestion game* the strategy sets Π_i are represented implicitly as $s_i - t_i$ paths in a directed network $G = (V, E)$.

Since the payoff functions u_i of a congestion game can be implicitly computed by the resource delay functions d_e, in the following we will denote a general (weighted or unweighted) congestion game by $\langle N, E, (\Pi_i)_{i\in N}, (w_i)_{i\in N}, (d_e)_{e\in E}\rangle$.

The following theorem [19], [13] proves the strong connection of unweighted congestion games with the exact potential games.

Theorem 1 ([19], [13]). *Every (unweighted) congestion game is an exact potential game.*

Proof. Fix an arbitrary (unweighted) congestion game $\Gamma = \langle N, E, (\Pi_i)_{i\in N}, (d_e)_{e\in E}\rangle$. For any pure strategy profile $\varpi \in \Pi$, the function

$$\Phi(\varpi) = \sum_{e\in\cup_{i\in N}\varpi_i} \sum_{k=1}^{\sigma_e(\varpi)} d_e(k) \tag{3}$$

(introduced in [19]) is an exact potential function for Γ. □

The converse of Theorem 1 does not hold in general, however in [13] it was proven that every (finite) exact potential game Γ is *isomorphic* to an unweighted congestion game. In [9] it was shown that this does not hold for the case of weighted congestion games. In particular, it was shown that there exist weighted single commodity network congestion games with resource delays being either linear or 2-wise linear functions of the loads, for which pure Nash equilibria cannot exist. Furthermore, there exist weighted single-commodity network congestion games which admit no exact potential function, even when the resource delays are identical to their loads. On the other hand, it was shown that any weighted (multi-commodity) network congestion game with linear resource delays admits a weighted potential function, yielding the existence of a pure Nash equilibrium which can be constructed in pseudo-polynomial time:

Theorem 2 ([9]). *For any weighted multi-commodity network congestion game $\langle N, E, (\Pi_i)_{i\in N}, (w_i)_{i\in N}, (d_e)_{e\in E}\rangle$ with linear resource delays, i.e., $d_e(x) = a_e x + b_e, e \in E, a_e, b_e \geq 0$, function $\Phi : \Pi \to \mathbb{R}$ defined for the configuration $\varpi \in \Pi$ as*

$$\Phi(\varpi) = \sum_{e\in E} d_e(\theta_e(\varpi))\theta_e(\varpi) + \sum_{i=1}^{n} \sum_{e\in\varpi_i} d_e(w_i)w_i$$

is a **b***-potential for* $b_i = 1/2w_i, i \in N$.

In [16], it was shown that weighted (multi-commodity) network congestion games with resource delays *exponential* to their loads also admit a weighted potential function:

Theorem 3 ([16]). *For any weighted multi-commodity network congestion game* $\langle N, E, (\Pi_i)_{i \in N}, (w_i)_{i \in N}, (d_e)_{e \in E} \rangle$ *with exponential resource delays, i.e.,* $d_e(x) = \exp(x)$*, function* $\Phi : \Pi \to \mathbb{R}$ *defined for the configuration* $\varpi \in \Pi$ *as*

$$\Phi(\varpi) = \sum_{e \in E} \exp(\theta_e(\varpi))$$

is a **b***-potential for* $b_i = \frac{\exp(w_i)}{\exp(w_i)-1}, i \in N$.

3.2 Concurrent congestion games and coalitions

In this section we study the effect of concurrent greedy moves of players in atomic congestion games where n selfish agents (players) wish to select a resource each (out of m resources) so that their selfish delay there is not much. This problem of "maintaining" global progress while allowing concurrent play is examined and answered in [7] with the use of potential functions. Two orthogonal settings are examined: (i) A game where the players decide their moves without global information, each acting "freely" by sampling resources randomly and locally deciding to migrate (if the new resource is better) via a random experiment. Here, the resources can have quite arbitrary latency that is load dependent. (ii) An "organised" setting where the players are pre-partitioned into selfish groups (coalitions) and where each coalition does an improving coalitional move.

Concurrent congestion games. A setting where selfish players perform best improvement moves in a sequential fashion and eventually reach a Nash equilibrium is not appealing to modern networking, where simple decentralized distributed protocols better reflect the essence of the networks liberal nature. In fact, it is unrealistic to assume that global coordination between the players can be enforced and that the players are capable of monitoring the configuration of the entire network. Furthermore, even if a player can grasp the whole picture, it is computationally demanding to decide her best move.

In [7], the advantages and the limitations of such a distributed protocol for congestion games on parallel edges under very general assumptions on the latency functions are investigated. A restricted model of distributed computation that allows a limited amount of global knowledge is adopted. In each round, every player can only select a resource uniformly at random and check its current latency. Migration decisions are made concurrently on the basis only of the current latency of the resource to which a player is assigned and the current latency of the resource to which the player is about to move. Migration decisions take advantage of local coordination between the players currently assigned to the same resource, in the sense that at most one player is allowed to depart from each resource. The only global information available to the players is an upper bound α on the slope of the latency functions.

In this setting, an (ε, α)-approximate equilibrium ((ε, α)-EQ), which is dictated by the very limited information available to the players, is a state where at most εm resources have latency either considerably larger or considerably smaller than the current average latency. This definition relaxes the notion of exact pure Nash equilibria and introduces a meaningful notion of approximate (bicriteria) equilibria for the myopic model of migrations described above. In particular, an (ε, α)-EQ guarantees that unless a player uses an overloaded resource (i.e., a resource with latency considerably larger than the average latency), the probability that she finds (by uniform sampling) a resource to migrate and significantly improve her latency is at most ε. Furthermore, it is unlikely that any (ε, α)-EQ reached by the protocol assigns a large number of players to overloaded resources (even though this possibility is allowed by the definition of an (ε, α)-EQ).

A simple oblivious protocol for this restricted model of distributed computation is presented in [7]. According to this myopic protocol, in parallel each player selects a resource uniformly at random in each round and checks whether she can significantly decrease her latency by moving to the chosen resource. If this is the case, the player becomes a potential migrant. The protocol uses a simple local probabilistic rule that selects at most one (this is a local decision between players on the same resource) potential migrant to defect from each resource.

If the number of players is $\Theta(m)$ and they start from a random initial allocation, the protocol reaches an (ε, α)-EQ in $O(\log(E[\Phi(0)]/\Phi_{\min}))$ time, where $E[\Phi(0)]$ is Rosental's expected potential value as the game starts and $\Phi_{\min}$ is the corresponding value at a Nash equilibrium. The proof of convergence given in [7] is technically involved and is omitted here.

Congestion games with coalitions. In many practical situations, the competition for resources takes place among coalitions of players instead of individuals. In such settings, it is important to know how the competition among coalitions affects the rate of convergence to an (approximate) pure Nash equilibrium. A *congestion game with coalitions* is a natural model for investigating the effects of non-cooperative resource allocation among static coalitions. In congestion games with coalitions, the coalitions are static and the selfish cost of each coalition is the total delay of its players.

In this setting, [7] present an upper bound on the rate of (sequential) convergence to approximate Nash equilibrium in single-commodity linear congestion games with static coalitions. The restriction to linear latencies is necessary because this is the only class of latency functions for which congestion games with static coalitions is known to admit a potential function and a pure Nash equilibrium. Sequences of ε-moves are considered, i.e., selfish deviations that improve the coalitions' total delay by a factor greater than ε. Combining the approach of [3] with the potential function of [8], it is shown that if the coalition with the largest improvement in its total delay moves in every round, an approximate Nash equilibrium is reached in a small number of steps.

More precisely, for any initial configuration s_0, every sequence of largest improvement ε-moves reaches an approximate Nash equilibrium in at most $\frac{kr(r+1)}{\varepsilon(1-\varepsilon)} \log \Phi(s_0)$ steps, where k is the number of coalitions, $r = \lceil \max_{j\in[k]}\{n_j\} / \min_{j\in[k]}\{n_j\} \rceil$ denotes the ratio between the size of the largest coalition and the size of the smallest coalition, and $\Phi(s_0)$ is the initial potential. This bound holds even for coalitions of different size, in which case the game is not symmetric. This bound implies that in network congestion games, where a coalition's best response can be computed in polynomial time, an approximate Nash equilibrium can be computed in polynomial time. Moreover, in the special case that the number of coalitions is constant and the coalitions are almost equisized, i.e. when, $k = \Theta(1)$ and $r = \Theta(1)$, the number of ε-moves to reach an approximate Nash equilibrium is logarithmic in the potential of the initial state.

3.3 Social ignorance in congestion games

Most of the work on congestion games focuses on the full information setting, where each player knows the precise weights and the actual strategies of all players, and her strategy selection takes all this information into account. In many typical applications of congestion games however, the players have incomplete information not only about the weights and the strategies, but also about the mere existence of (some of) the players with whom they compete for resources. In fact, in many applications, it is both natural and convenient to assume that there is a *social context* associated with the game, which essentially determines the information available to the players. In particular, one may assume that each player has complete information about the set of players in her *social neighborhood*, and limited (if any) information about the remaining players.

In this section we investigate how such social-context-related information considerations affect the inefficiency of pure Nash equilibria and the convergence rate to approximate pure Nash equilibria. To come up with a manageable setting that allows for some concrete answers, we make the simplifying assumption that each player has complete information about the players in her social neighborhood, and no information whatsoever about the remaining players. Therefore, since each player is not aware of the players outside her social neighborhood, her individual cost and her strategy selection are not affected by them. In fact, this is the model of *graphical congestion games*, introduced by Bilò et al. [1]. The new ingredient in the definition of graphical congestion games is the social graph, which represents the players social context. The social graph is defined on the set of players and contains an edge between each pair of players that know each other. The basic idea (and assumption) behind graphical congestion games is that the individual presumed cost of each player only depends on the players in her social neighborhood, and thus her strategy selection is only affected by them.

The social graph $G(V, E)$ is defined on the set of players $V = N$ and contains an edge $\{i, j\} \in E$ between each pair of players i, j that know each other. We consider graphical games with weighted players and simple undirected social graphs.

Given a graphical congestion game with a social graph $G(V,E)$, a configuration s and a resource e, let $V_e(s) = \{i \in V : e \in s_i\}$ be the set of players using e in s, let $G_e(s)(V_e(s), E_e(s))$ be the social subgraph of G induced by $V_e(s)$, and let $n_e(s) = |V_e(s)|$ and $m_e(s) = |E_e(s)|$. For each player i (not necessarily belonging to $V_e(s)$), let $\Gamma_e^i(s) = \{j \in V_e(s) : \{i,j\} \in E\}$ be i's social neighborhood among the players using e in s. In any configuration s, a player i is aware of a *presumed congestion* $s_e^i = w_i + \sum_{j \in \Gamma_e^i(s)} w_j$ on each resource e, and of her *presumed cost* $p_i(s) = \sum_{e \in s_i} w_i(a_e s_e^i + b_e)$. We note that the presumed cost coincides with the actual cost if the social graph is complete. For graphical congestion games, a configuration s is a pure Nash equilibrium if no player can improve her *presumed cost* by unilaterally changing her strategy.

In [1] it is shown that graphical linear congestion games with unweighted players are potential games:

Theorem 4. *Every graphical linear congestion game defined over an undirected social graph is an exact potential game, and thus always converges to a Nash equilibrium.*

Proof. The potential function establishing the result is

$$\Phi(s) = \sum_{e \in E} \left[a_e(m_e(s) + n_e(s)) + b_e n_e(s)\right] \quad .$$

□

In [6] it is shown that graphical linear congestion games with weighted players also admit a potential function:

Theorem 5. *Every graphical linear congestion game with weighted players admits a potential function, and thus a pure Nash equilibrium.*

Proof. The potential function establishing the result is

$$\Phi(s) = \sum_{e \in R} \left[a_e \left(\sum_{i \in V_e(s)} w_i^2 + \sum_{\{i,j\} \in E_e(s)} w_i w_j \right) + b_e \sum_{i \in V_e(s)} w_i \right] \quad .$$

□

4 Potential functions in population dynamics: Generalized Moran process

In this section we consider the Moran process, as generalized by Lieberman et al. [12]. A population resides on the vertices of a finite, connected, undirected graph and, at each time step, an individual is chosen at random with probability proportional to its assigned “fitness” value. It reproduces, placing a copy of itself on a neighboring vertex chosen uniformly at random, replacing the individual that was there. The initial population consists of a single mutant of fitness $r > 0$ placed uniformly at random, with every other vertex occupied by an individual of fitness 1. The main quantities of interest are the probabilities that the descendants of the initial mutant come to occupy the whole graph (fixation) and that they die out (extinction); almost surely, these are the only possibilities. In general, exact computation of these quantities by standard Markov chain techniques requires solving a system of linear equations of size exponential in the order of the graph so is not feasible. In [4] a potential function is used to show that, with high probability, the number of steps needed to reach fixation or extinction is bounded by a polynomial in the number of vertices in the graph. This bound allows to construct fully polynomial randomized approximation schemes (FPRAS) for the probability of fixation (when $r \geq 1$) and of extinction (for all $r > 0$).

In the following, we consider only finite, connected, undirected graphs $G = (V, E)$ of order $n = |V|$, and r denotes the fitness of the initially introduced mutant in the graph. Given a set $X \subseteq V$, we denote by $W(X) = r|X| + |V \setminus X|$ the *total fitness* of the population when exactly the vertices of X are occupied by mutants.

We first show that the Moran process on a connected graph G of order n is expected to reach absorption in a polynomial number of steps. To do this, we use the potential function given by $\Phi(X) = \sum_{x \in X} \frac{1}{\deg(x)}$ for any state $X \subseteq V$. Note that $1 < \Phi(V) \leq n$ and that, if $(X_i)_{i \geq 0}$ is a Moran process on G then $\Phi(X_0) = \frac{1}{\deg(x)} \leq 1$ for some vertex $x \in V$ (the initial mutant). The following lemma shows that the potential strictly increases in expectation when $r > 1$ and strictly decreases in expectation when $r < 1$.

Lemma 1. *Let $(X_i)_{i \geq 0}$ be a Moran process on a graph $G = (V, E)$ and let $\emptyset \subset S \subset V$. If $r \geq 1$, then $E\left[\Phi(X_{i+1}) - \Phi(X_i) | X_i = S\right] > \left(1 - \frac{1}{r}\right) \cdot \frac{1}{n^3}$. Otherwise, $E\left[\Phi(X_{i+1}) - \Phi(X_i) | X_i = S\right] < \frac{r-1}{n^3}$.*

Proof. Write $W(S) = n + (r-1)|S|$ for the total fitness of the population. For $\emptyset \subset S \subset V$, and any value of r, we have

$$E\left[\Phi(X_{i+1}) - \Phi(X_i) | X_i = S\right] = \frac{r-1}{W(S)} \sum_{xy \in E, x \in S, y \in \overline{S}} \frac{1}{\deg(x)\deg(y)}.$$

The sum is minimized by noting that there must be at least one edge between S and $\overline{S}$ and that its endpoints have degree at most $(n-1) < n$. The greatest weight configuration is the one with all mutants if $r \geq 1$ and the one with no mutants if $r < 1$. Therefore, if $r \geq 1$, we have $E\left[\Phi(X_{i+1}) - \Phi(X_i) | X_i = S\right] > \frac{r-1}{rn} \cdot \frac{1}{n^2} = \left(1 - \frac{1}{r}\right) \frac{1}{n^3}$ and, if $r < 1$, $E\left[\Phi(X_{i+1}) - \Phi(X_i) | X_i = S\right] < (r-1)\frac{1}{n^3}$. □

Martingale techniques are used to bound the expected absorption time. It is well known how to bound the expected absorption time using a potential function that decreases in expectation until absorption. This has been made explicit by

Hajek [10] and [4] uses the following formulation based on that of He and Yao [11].

Theorem 6. *Let $(Y_i)_{i\geq 0}$ be a Markov chain with state space Ω, where Y_0 is chosen from some set $I \subseteq \Omega$. If there are constants $k_1, k_2 > 0$ and a non-negative function $\psi : \Omega \to \mathbb{R}$ such that*

- *$\Psi(S) = 0$ for some $S \in \Omega$,*
- *$\Psi(S) \leq k_1$ for all $S \in I$ and*
- *$E[\Psi(Y_i) - \Psi(Y_{i+1})|Y_i = S] \geq k_2$ for all $i \geq 0$ and all S with $\Psi(S) > 0$,*

then $E[\tau] \leq k_1/k_2$, where $\tau = \min\{i : \Psi(Y_i) = 0\}$.

The above theorem is useful in bounding the absorption time of the Moran process:

Theorem 7. *Let $G = (V, E)$ be a graph of order n. For $r < 1$, the absorption time τ of the Moran process on G satisfies $E[\tau] \leq \frac{1}{1-r} n^3$.*

Proof. Let $(Y_i)_{i\geq 0}$ be the process on G that behaves identically to the Moran process except that, if the mutants reach fixation, we introduce a new non-mutant on a vertex chosen uniformly at random. That is, from the state V, we move to $V - x$, where x is chosen u.a.r., instead of staying in V. Writing $\tau' = \min\{i : Y_i = \emptyset\}$ for the absorption time of this new process, it is clear that $E[\tau] \leq E[\tau']$. The function Φ meets the criteria for Ψ in the statement of Theorem 6 with $k_1 = 1$ and $k_2 = (1 - r)n^{-3}$. The first two conditions of the theorem are obviously satisfied. For $S \subset V$, the third condition is satisfied by Lemma 1 and we have

$$E[\Phi(Y_i) - \Phi(Y_{i+1})|Y_i = V] = \frac{1}{n} \sum_{x\in V} \frac{1}{\deg(x)} > \frac{1}{n} > k_2 \ .$$

Therefore, $E[\tau] \leq E[\tau'] \leq \frac{1}{1-r} n^3$. □

The following corollary is immediate from Markov's inequality.

Corollary 1. *The Moran process on G with fitness $r < 1$ reaches absorption within t steps with probability at least $1 - \varepsilon$, for any $\varepsilon \in (0, 1)$ and any $t \geq \frac{1}{1-r} n^3/\varepsilon$.*

For $r > 1$, the proof needs slight adjustment because, in this case, Φ increases in expectation, and we obtain:

Corollary 2. *The Moran process on G with fitness $r > 1$ reaches absorption within t steps with probability at least $1 - \varepsilon$, for any $\varepsilon \in (0, 1)$ and any $t \geq \frac{r}{r-1} n^3 \Phi(G)/\varepsilon$.*

We refer to [4] for detailed proofs.

5 Bounding the chromatic number of graphs

One of the central optimization problems in Graph Theory and Computer Science is the problem of *vertex coloring* of graphs: Given a graph $G = (V, E)$ with n vertices, assign a color to each vertex of G so that no pair of adjacent vertices gets the same color (i.e., so that the coloring produced is *proper*) and so that the total number of distinct colors used is minimized. The global optimum of vertex coloring is the *chromatic number* $\chi(G)$, defined as the minimum number of colors required to properly color the vertices of graph G.

The problem of coloring a graph using the minimum number of colors is NP-hard, and the chromatic number cannot be approximated to within $\Omega(n^{1-\epsilon})$ for any constant $\epsilon > 0$, unless $\mathsf{NP} \subseteq \mathsf{co\text{-}RP}$ [5]. Despite these negative approximation results, several upper bounds on the chromatic number have been proven in the literature; these bounds are related to various graph-theoretic parameters. In particular, given a graph $G = (V, E)$, let n and m denote the number of vertices and number of edges of G. Let $\Delta(G)$ denote the maximum degree of a vertex in G, and let $\Delta_2(G)$ be the maximum degree that a vertex $v \in V$ can have subject to the condition that v is adjacent to at least one vertex of degree no less than the degree of v (note that $\Delta_2(G) \leq \Delta(G)$). Denote by $\omega(G)$ the *clique number* of G, i.e., the maximum size of a clique in G, and by $\alpha(G)$ the *independence number* of G, i.e., the maximum size of an independent set in G. Then, it is known that

$$\chi(G) \leq \min\left\{\Delta_2(G) + 1,\ \frac{n + \omega(G)}{2},\ n - \alpha(G) + 1,\ \frac{1 + \sqrt{1 + 8m}}{2}\right\} . \tag{4}$$

A different, potential-based method for proving all the above bounds on the chromatic number in a unified manner is given in [17]. The method is *constructive*, in the sense that it actually computes *in polynomial time* a proper coloring using a number of colors that satisfies these bounds. In particular, the vertices of a graph $G = (V, E)$ are viewed as players in a strategic game. Given a finite, simple, undirected graph $G = (V, E)$ with $|V| = n$ vertices, we define the *graph coloring game* $\Gamma(G)$ as the game in strategic form where the set of players is the set of vertices V, and the action set of each vertex is a set of n colors $X = \{x_1, \ldots, x_n\}$. A *configuration* or *pure strategy profile* $\mathbf{c} = (c_v)_{v \in V} \in X^n$ is a vector representing a combination of actions, one for each vertex. That is, c_v is the color chosen by vertex v. For a configuration $\mathbf{c} \in X^n$ and a color $x \in X$, we denote by $n_x(\mathbf{c})$ the number of vertices that are colored x in $\mathbf{c}$, i.e., $n_x(\mathbf{c}) = |\{v \in V : c_v = x\}|$. The *payoff* that vertex $v \in V$ receives in the configuration $\mathbf{c} \in X^n$ is

$$\lambda_v(\mathbf{c}) = \begin{cases} 0 & \text{if } \exists u \in N(v) : c_u = c_v \\ n_{c_v}(\mathbf{c}) & \text{else} \end{cases} .$$

A configuration $\mathbf{c} \in X^n$ of the graph coloring game $\Gamma(G)$ is a pure Nash equilibrium, or an *equilibrium coloring*, if, for all vertices $v \in V$ and for all colors $x \in X$, $\lambda_v(x, \mathbf{c}_{-v}) \leq \lambda_v(\mathbf{c})$. A vertex $v \in V$ is *unsatisfied* in the configuration

$\mathbf{c} \in X^n$ if there exists a color $x \neq c_v$ such that $\lambda_v(x, \mathbf{c}_{-v}) > \lambda_v(\mathbf{c})$; else we say that v is *satisfied*. For an unsatisfied vertex $v \in V$ in the configuration $\mathbf{c}$, we say that v performs a *selfish step* if v unilaterally deviates to some color $x \neq c_v$ such that $\lambda_v(x, \mathbf{c}_{-v}) > \lambda_v(\mathbf{c})$.

For any graph coloring game $\Gamma(G)$, define the function $\Phi : P \to \mathbb{R}$, where $P \subseteq X^n$ is the set of all configurations that correspond to proper colorings of the vertices of G, as $\Phi(\mathbf{c}) = \frac{1}{2}\sum_{x\in X} n_x^2(\mathbf{c})$, for all proper colorings $\mathbf{c}$.

Theorem 8. *Φ is an exact potential function for $\Gamma(G)$.*

Proof. Fix a *proper* coloring $\mathbf{c}$. Assume that vertex $v \in V$ can improve its payoff by deviating and selecting color $x \neq c_v$. This implies that the number of vertices colored c_v in $\mathbf{c}$ is at most the number of vertices colored x in $\mathbf{c}$, i.e., $n_{c_v}(\mathbf{c}) \leq n_x(\mathbf{c})$. If v indeed deviates to x, then the resulting configuration $\mathbf{c}' = (x, \mathbf{c}_{-v})$ is again a proper coloring (vertex v can only decrease its payoff by choosing a color that is already used by one of its neighbors, and v is the only vertex that changes its color). The improvement on v's payoff will be $\lambda_v(\mathbf{c}') - \lambda_v(\mathbf{c}) = n_x(\mathbf{c}') - n_{c_v}(\mathbf{c}) = n_x(\mathbf{c}) + 1 - n_{c_v}(\mathbf{c})$. Moreover, $\Phi(\mathbf{c}') - \Phi(\mathbf{c}) = \frac{1}{2}\left(n_x^2(\mathbf{c}') + n_{c_v}^2(\mathbf{c}') - n_x^2(\mathbf{c}) - n_{c_v}^2(\mathbf{c})\right) = \lambda_v(\mathbf{c}') - \lambda_v(\mathbf{c})$. □

Therefore, if any vertex v performs a selfish step (i.e., changes its color so that its payoff is increased), then the value of Φ is increased as much as the payoff of v is increased. Now, the payoff of v is increased by at least 1. So after any selfish step the value of Φ increases by at least 1. Now observe that, for all proper colorings $\mathbf{c} \in P$ and for all colors $x \in X$, $n_x(\mathbf{c}) \leq \alpha(G)$. Therefore

$$\Phi(\mathbf{c}) = \frac{1}{2}\sum_{x\in X} n_x^2(\mathbf{c}) \leq \frac{1}{2}\sum_{x\in X}(n_x(\mathbf{c}) \cdot \alpha(G)) = \frac{1}{2}\alpha(G)\sum_{x\in X} n_x(\mathbf{c}) = \frac{n \cdot \alpha(G)}{2} \ .$$

Moreover, the minimum value of Φ is $\frac{1}{2}n$. Therefore, if we allow any unsatisfied vertex (but only one each time) to perform a selfish step, then after at most $\frac{n\cdot\alpha(G)-n}{2}$ steps there will be no vertex that can improve its payoff. This implies that a pure Nash equilibrium will have been reached. Of course, we have to start from an initial configuration that is a proper coloring so as to ensure that the procedure will terminate in $O(n \cdot \alpha(G))$ selfish steps; this can be found easily since there is always the trivial proper coloring that assigns a different color to each vertex of G.

In [17] it was also shown that *any* equilibrium coloring of $\Gamma(G)$ uses a number of colors that satisfies all the bounds given in 4. This, combined to the existence of the potential function, implies that, for any graph G, a proper coloring that uses at most $k \leq \min\left\{\Delta_2(G)+1,\ \frac{n+\omega(G)}{2},\ \frac{1+\sqrt{1+8m}}{2},\ n-\alpha(G)+1\right\}$ colors can be computed in polynomial time.

References

1. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Graphical Congestion Games. Proceedings of the 4th International Workshop on Internet and Network Economics (WINE 2008), LNCS vol. 5385, pp. 70-81, 2008.

2. Chen, X., and Xiaotie, D.: Settling the complexity of two-player Nash equilibrium. Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 261–272, 2006.
3. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms (SODA 2007), pp. 169–178, 2007.
4. Díaz, J., Goldberg, L. A., Mertzios, G. B., Richerby, D., Serna, M. J., Spirakis, P. G.: Approximating fixation probabilities in the generalized Moran process. Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012), pp. 954–960, 2012.
5. Feige, U., Kilian, J.: Zero knowledge and the chromatic number, J. Comput. System Sci. **57**(2), pp. 187–199, 1998.
6. Fotakis, D., Gkatzelis, V., Kaporis, A. C., Spirakis, P. G.: The Impact of Social Ignorance on Weighted Congestion Games. Proceedings of the 5th International Workshop on Internet and Network Economics (WINE 2009), pp. 316–327.
7. Fotakis, D., Kaporis, A. C., Spirakis, P. G.: Atomic congestion games: fast, myopic and concurrent. Theory Comput. Syst. **47**(1), pp. 38–59, 2010.
8. Fotakis, D., Kontogiannis, S., Spirakis, P. G.: Atomic congestion games among coalitions. ACM Transactions on Algorithms (TALG) **4**(4), article No. 52, 2008.
9. Fotakis, D., Kontogiannis, S., Spirakis, P. G.: Selfish unsplittable flows. Theoretical Computer Science **348**(2), pp. 226–239, 2005.
10. Hajek, B.: Hitting-time and occupation-time bounds implied by drift analysis with applications. Advances in Applied Probability **14**(3), pp.502–525, 1982.
11. He, J., Yao, X.: Drift analysis and average time complexity of evolutionary algorithms. Artificial Intelligence **127**, pp. 57–85, 2001.
12. E. Lieberman, C. Hauert, and M. A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433:312–316, 2005.
13. Monderer, D., Shapley, L.: Potential games. Games and Economic Behavior **14**, pp. 124–143, 1996.
14. Nash, J. F.: Non-cooperative games. Annals of Mathematics **54**(2), pp. 286–295, 1951.
15. Papadimitriou, C. H.: On the complexity of the parity argument and other inefficient proofs of existence. Journal of Computer and System Sciences **48**(3), pp. 498-532, 1994.
16. Panagopoulou, P. N, Spirakis, P. G.: Algorithms for pure Nash equilibria in weighted congestion games. ACM Journal of Experimental Algorithmics **11**, 2006.
17. Panagopoulou, P. N, Spirakis, P. G.: A game-theoretic approach for efficient graph coloring. Algorithms and Computation, 19th International Symposium (ISAAC 2008), pp. 183–195, 2008.
18. Panagopoulou, P. N, Spirakis, P. G.: Playing a game to bound the chromatic number. The American Mathematical Monthly **119**(9), pp. 771–778, 2012.
19. Rosenthal, R. W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory **2**, pp. 65–67, 1973.
20. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. Pacific Journals of Mathematics **5**, pp. 285–308, 1955.
21. Topkis, D.: Equilibrium points in nonzero-sum n-person submodular games. SIAM Journal of Control and Optimization **17**, pp. 773–787, 1979.
22. Uno, H.: Strategic complementarities and nested potential games. Journal of Mathematical Economics **47**(6), pp. 728-732, 2011.