

A Network Game with Attackers and a Defender

Marios Mavronicolas · Vicky Papadopoulou ·
Anna Philippou · Paul Spirakis

Received: 29 March 2006 / Accepted: 25 September 2006 / Published online: 30 October 2007
© Springer Science+Business Media, LLC 2007

Abstract Consider an information network with *threats* called *attackers*; each attacker uses a probability distribution to choose a *node* of the network to damage. Opponent to the attackers is a protector entity called *defender*; the defender scans and cleans from attacks some part of the network (in particular, a *link*), which it chooses independently using its own probability distribution. Each attacker wishes to maximize the probability of escaping its cleaning by the defender; towards a conflicting objective, the defender aims at maximizing the expected number of attackers it catches.

We model this network security scenario as a non-cooperative *strategic* game on graphs. We are interested in its associated *Nash equilibria*, where no network entity can unilaterally increase its local objective. We obtain the following results:

A preliminary version of this work appeared in the Proceedings of the 16th Annual International Symposium on Algorithms and Computation, X. Deng and D. Du, eds., Lecture Notes in Computer Science, vol. 3827, pp. 288–297, Springer, December 2005. This work has been partially supported by the IST Program of the European Union under contract 001907 (DELIS), and by research funds at University of Cyprus.

M. Mavronicolas · V. Papadopoulou (✉) · A. Philippou
Department of Computer Science, University of Cyprus, Nicosia 1678, Cyprus
e-mail: wiki@ucy.ac.cy

M. Mavronicolas
e-mail: mavronic@ucy.ac.cy

A. Philippou
e-mail: annap@ucy.ac.cy

P. Spirakis
Department of Computer Engineering and Informatics, University of Patras, Rion, Patras 26500,
Greece

P. Spirakis
Research Academic Computer Technology Institute, Rion, Patras 26500, Greece
e-mail: spirakis@cti.gr

- We obtain an algebraic characterization of (*mixed*) Nash equilibria.
- No (non-trivial) instance of the graph-theoretic game has a *pure* Nash equilibrium. This is an immediate consequence of some *covering* properties we prove for the *supports* of the players in all (*mixed*) Nash equilibria.
- We coin a natural subclass of mixed Nash equilibria, which we call *Matching Nash equilibria*, for this graph-theoretic game. Matching Nash equilibria are defined by enriching the necessary covering properties we proved with some additional conditions involving other structural parameters of graphs, such as *Independent Sets*.
 - We derive a characterization of graphs admitting Matching Nash equilibria. All such graphs have an *Expanding Independent Set*. The characterization enables a non-deterministic, polynomial time algorithm to compute a Matching Nash equilibrium for any such graph.
 - Bipartite graphs are shown to satisfy the characterization. So, using a polynomial time algorithm to compute a Maximum Matching for a bipartite graph, we obtain, as our main result, a deterministic, polynomial time algorithm to compute a Matching Nash equilibrium for any instance of the game with a bipartite graph.

Keywords Game theory · Graph theory · Security · Nash equilibria · Attacks and defenses

1 Introduction

1.1 Motivation and Framework

1.1.1 Attacks and Defenses

The huge growth of the Internet has significantly extended the importance of *Network Security* [25]. As it is well known, many widely used Internet systems and components are prone to security risks (see, for example, [6]). Such risks are inevitable since some parts of a large system such as the Internet with many computers and a wide range of software are expected to have weaknesses that expose them to security attacks. Some such risks have led to successful and well-publicized attacks [24]. Typically, an *attack* exploits the discovery of loopholes in the security mechanisms of the Internet; the latter are also known as *defenses*.

Attacks and defenses are currently attracting a lot of interest in major forums of communication research; see, for example, Session 37 (*Attacks and defenses in emerging networks*) in *INFOCOM 2006*. A current challenge is to invent and study appropriate theoretical models of security attacks and defenses for emerging networks like the Internet. In this work, we introduce and analyze such a model for a very simple case of attacks and defenses.

1.1.2 The Network Security Game

We consider a network whose nodes are vulnerable to infection by *threats* (e.g., viruses, worms, Trojan horses or eavesdroppers [9]), called *attackers*. Available to the network is a security software (or *firewall*), called the *defender*. However, due to economic and performance reasons, the defender is only able to clean a limited part of the network. Such reasons stem from financial costs (e.g., the prohibitive cost of purchasing a global security software) or from performance bottlenecks (e.g., the reduced usability of the protected part of the network).

The defender seeks to protect the network as much as possible; on the other hand, each attacker wishes to avoid being caught so as to be able to damage the network. Both the attackers and the defender make individual decisions for their placement in the network while seeking to maximize their contrary objectives. Each attacker targets a *node* of the network chosen via its own probability distribution on nodes; the defender cleans a single *link* chosen via its own probability distribution on links. The node chosen by an attacker is damaged unless it is incident to the link being cleaned by the defender.

Since attacks and defenses over the Internet are self-interested procedures that seek to maximize damage and protection, respectively, it is natural to model this network security scenario as a non-cooperative *strategic game* on graphs with two kinds of *players*: the *vertex players* representing the attackers and the *edge player* representing the defender. The (*Expected*) *Individual Profit* of an attacker is the probability that it is not caught by the defender; the (*Expected*) *Individual Profit* of the defender is the expected number of attackers it catches.

We are interested in the *Nash equilibria* [21, 22] associated with this graph-theoretic game, where no player can unilaterally improve its (*Expected*) *Individual Profit* by switching to another probability distribution. Nash's celebrated result [21, 22] guarantees that the graph-theoretic game has at least one (*mixed*) Nash equilibrium.

1.1.3 Three Motivating Examples

Network Edge Security A first motivating example for our network security game comes from *Network Edge Security* [16], a new distributed firewall architecture designed to counter threats undetected by existing firewall implementations that usually protect individual *servers*. In this new architecture, a firewall is implemented in a distributed way, rather than being locally installed; so, the firewall protects the subnetwork spanned by the nodes that participate in the distributed implementation. The simplest case occurs when the subnetwork is just a single link with its two incident nodes; this starting case offers an initial basis for our theoretical model.

Understanding the mathematical intricacies of attacks and defenses in this simplest theoretical model is a necessary prerequisite for making rigorous progress in the analysis of distributed firewall architectures with more involved topologies.

Corrupted Databases A second motivating example comes from considering a collection of databases, each stored at an individual computer. One way to detect corruption is to store hashed versions (or *fingerprints*) of the database on some other

computers. Assume that whenever one computer stores a fingerprint for another, the second also does the same for the first. This assumption induces an undirected graph on computers in the natural way.

An attacker is an entity that chooses a computer to corrupt its database. The defender is the distributed database administrator whose goal is to maintain data integrity. To do so, it chooses two adjacent computers in the induced graph and compares their two databases to the corresponding fingerprints in order to detect corruption. These checks are performed on a regular basis, and the administrator is rewarded for any corruption it detects in a proportional way.

Intranets and Gateways A third motivating example comes from representing the Internet as a graph whose edges represent *intranets* and vertices represent *gateways* to an intranet; adjacent (trusting) intranets are modeled by collapsing their corresponding gateways down to one. So, a vertex refers to a collection of mutually trusted intranets. The simplest (non-trivial) case arises when each intranet participates in exactly two such trusts, and our network security game models this simplest case.

An attacker chooses a gateway in order to harm the corresponding trust of intranets. A defender chooses an edge in order to monitor and control all communication into and out of the corresponding intranet; so, a defender acts as a firewall for the intranet it chooses.

1.1.4 The Assumption of Multiple Attackers

The assumption of multiple attackers models better the heterogeneity and independence of security attacks to the Internet; there, diverse threats with varying destruction goals act independently of each other and may have different privileges to the availability of (encrypted) mobile code or to the access of large resources.

A feature implicit in our modeling with multiple attackers is that there is no interaction among the attackers: if the same vertex is attacked by multiple attackers, all of them profit unless they are caught. This feature fits well into the context of (nilpotent) mischievous attacks to the same target, which benefit from each other in an indirect way: to achieve the common goal, it suffices that at least one attack succeeds.

We remark that an alternative system with a single attacker choosing multiple edges would fail to capture the essence of multiple, independent security attacks, while it might be suitable for modeling a single orchestrated attack by *cooperating* threats with mutual interests. However, a Nash equilibrium for a system with multiple attackers must satisfy multiple constraints, as opposed to a Nash equilibrium for the alternative single-attacker system. This implies that the two systems may have, in general, different sets of Nash equilibria. In turn, this allows for the possibility that each of multiple attackers employs a different strategy in a Nash equilibrium for the multiple attackers case; this possibility gets closer to reality featuring attacks with ranging scales and policies.

1.2 Contribution

We obtain a multitude of results for several classes of Nash equilibria and for some special class of graphs, namely the bipartite graphs.

To describe our contribution, we will need some game-theoretic terminology, which we review here. (For precise definitions, see Sect. 2.) A *profile* is a tuple of probability distributions, one for each player. The *support* of the edge player is the set of all edges to which it assigns strictly positive probability; the *support* of a vertex player is the set of all vertices to which it assigns strictly positive probability, and the *support* of the vertex players is the union of the supports of all vertex players.

1.2.1 Mixed Nash Equilibria

Characterization

We discover that a mixed Nash equilibrium enjoys some elegant algebraic characterization (Theorem 3.1). The characterization is a precise algebraic formulation of the requirement that no player can unilaterally improve its Expected Individual Profit in a Nash equilibrium. In more detail, the characterization provides a system of equalities and inequalities to be satisfied by the players' probabilities.

Graph-Theoretic Structure

We proceed to study the graph-theoretic structure of mixed Nash equilibria. We discover two interesting *covering* properties of Nash equilibria. In more detail, we prove that in a Nash equilibrium, the support of the edge player is an *Edge Cover* of the graph (Proposition 4.1); the support of the vertex players is a *Vertex Cover* of the graph induced by the support of the edge player (Proposition 4.2). So, these covering properties represent necessary graph-theoretic conditions for Nash equilibria.

Inspired by the shown covering properties, we define a *Covering profile* as one that satisfies the two necessary covering conditions for Nash equilibria we proved. It is natural to ask whether a Covering profile is necessarily a Nash equilibrium. We provide a simple counterexample to show that a Covering profile is *not* necessarily a Nash equilibrium (Proposition 4.4). This implies that a Covering profile must be enriched with some additional condition(s) in order to provide a set of sufficient graph-theoretic conditions for Nash equilibria.

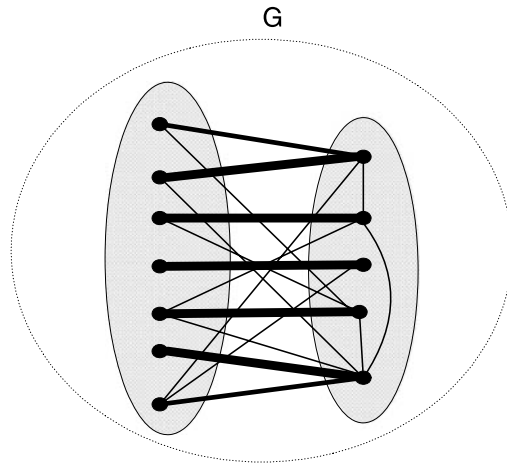
We attempt such enrichment in our definition of an *Independent Covering profile* (Definition 4.2). Loosely speaking, the following two additional conditions are included in the definition of an Independent Covering profile:

- The support of the vertex players is an *Independent Set* of the graph.
- Each vertex in the support of the vertex players is incident to *exactly* one edge from the support of the edge player.

Note that, intuitively, the first condition in the definition of an Independent Covering profile favors a decrease to the expected number of vertex players caught by the edge player. Moreover, intuitively, the second condition favors a decrease to the probability that some vertex player be caught by the edge player. So, by its two additional conditions, an Independent Covering profile is one that, intuitively, favors the vertex players.

In addition, the following two auxiliary conditions are included in the definition of an Independent Covering profile:

Fig. 1 An illustration of a graph with a Matching Nash equilibrium. The *left shaded region* represents the support of the vertex players, which is an Independent Set; the *right shaded region* represents its complement. Note that there may be edges joining vertices in the *right region*. The support of the edge player is indicated by *bold edges*; its contained Matching is indicated by *double bold edges*



- All vertex players have the same support.
- Each player uses a uniform probability distribution on its support.

These two conditions provide some more intuitive, simplifying assumptions that may facilitate the computation of an Independent Covering profile. In particular, the first auxiliary condition provides some kind of symmetry for the vertex players; the second auxiliary condition provides some kind of symmetry for the support of the edge player.

We prove that an Independent Covering profile is a Nash equilibrium (Proposition 4.6). The proof verifies that an Independent Covering profile satisfies the characterization of a Nash equilibrium (Theorem 3.1). So, an Independent Covering profile provides sufficient graph-theoretic conditions for Nash equilibria.

Moreover, we prove that in an Independent Covering profile, the support of the edge player contains a suitable *Matching* that matches each vertex outside the support of the vertex players to some vertex in the support of the vertex players. So, an Independent Covering profile will henceforth be called a *Matching Nash* equilibrium. Figure 1 provides an illustration for a Matching Nash equilibrium.

1.2.2 Pure Nash Equilibria

We prove that the graph-theoretic game has no pure Nash equilibrium unless the graph is trivial (Theorem 4.3). This follows as an immediate consequence of one of the covering properties of Nash equilibria that we establish (Proposition 4.1).

1.2.3 Matching Nash Equilibria

Characterization We provide a very simple and intuitive characterization of graphs admitting a Matching Nash equilibrium (Theorem 5.1). Specifically, we prove that a graph admits a Matching Nash equilibrium if and only if it has an Independent Set whose complementary vertex set is an *Expander* for the graph. Call such an Independent Set an *Expanding Independent Set*. To the best of our knowledge, Expanding Independent Sets have not been considered before in the literature.

Computation The proof of the characterization of graphs admitting a Matching Nash equilibrium is *constructive* (specifically, the sufficiency part of the proof). Hence, it immediately yields a non-deterministic algorithm to compute a Matching Nash equilibrium for any graph that satisfies the characterization (Algorithm MatchingNE in Fig. 4). The only remaining non-deterministic step in the algorithm is the step of choosing an Expanding Independent Set for any arbitrary graph that has one. At present, we do not know if there is a deterministic, polynomial time algorithm to compute an Expanding Independent Set for an *arbitrary* graph (that has one). Hence, we do not know if there is a deterministic, polynomial time algorithm to compute a Matching Nash equilibrium for an *arbitrary* graph that admits one.

Other than this step, the remaining steps of the algorithm MatchingNE run in polynomial time. The dominating such step invokes the computation of a *Maximum Matching* of some intermediate bipartite graph; this step can be done in polynomial time using the (currently most efficient) algorithm of Feder and Motwani [8]. Hence, the non-deterministic algorithm MatchingNE to compute a Matching Nash equilibrium for an arbitrary graph that has an Expanding Independent Set runs in polynomial time.

Non-Triviality

We use the characterization of graphs admitting Matching Nash equilibria (Theorem 5.1) to prove that the class of graphs admitting Matching Nash equilibria is *non-trivial* (Theorem 5.6): there are both graphs admitting and not admitting Matching Nash equilibria (Propositions 5.4 and 5.5). Interestingly, we discover that an *odd cycle* does not admit a Matching Nash equilibrium, while an *even cycle* does. König's Theorem [15] shows that a graph has no odd cycle if and only if it is bipartite; hence, due to our counterexamples, it is natural to consider the class of bipartite graphs as a candidate class of graphs admitting Matching Nash equilibria.

1.2.4 Bipartite Graphs

We prove that a bipartite graph satisfies the characterization of graphs admitting a Matching Nash equilibrium (Proposition 6.1); so, a bipartite graph admits a Matching Nash equilibrium (Corollary 6.2).

The proof of this property of bipartite graphs is *constructive*. Hence, it immediately completes a deterministic algorithm to compute a Matching Nash equilibrium for bipartite graphs by providing an algorithm to compute an Expanding Independent Set of the graph. More important, we observe that an Expanding Independent Set of a bipartite graph can be computed in polynomial time by reduction to computing a *Minimum Vertex Cover* of a bipartite graph; in turn, this can be computed in polynomial time by reduction to computing a *Maximum Matching* of a bipartite graph. Hence, a polynomial time algorithm to compute a Matching Nash equilibrium for a bipartite graph follows (Theorem 6.3), as the main result of our study.

1.3 Significance

To the best of our knowledge, our work is the *first* (with an exception of [2]) to formulate a network security problem as a strategic, graph-theoretic game and study its

associated Nash equilibria. Our formulated game is the *first* game where the network is explicitly modeled as a distinct, non-cooperative player. We believe that our work invites a simultaneously game-theoretic and graph-theoretic analysis of network security problems for which not only threats seek to maximize their caused damage to the network but also the network seeks to protect itself as much as possible.

Our results are specific to the particular graph-theoretic game we have introduced and studied; however, they exhibit a novel interaction of ideas, arguments and techniques from two seemingly diverse fields, namely *Game Theory* and *Graph Theory*. We believe that this interaction is quite promising. In particular, we believe that Matching Nash equilibria (and suitable extensions of them) will find further applications to other graph-theoretic games motivated by network security problems with a similar flavor.

Our work joins the booming area of *Algorithmic Game Theory* [20]. In particular, to the best of our knowledge, our work is the *first* to study algorithmic properties of a non-cooperative strategic game with two *distinct* kinds of opponent players.

1.4 Related Work and Comparison

Our graph-theoretic game can be seen as a graph-theoretic generalization of the classical *Matching Pennies* game [26] where two players choose simultaneously between two strategies; one wants to choose like the other, whereas the second wants to choose differently from the first. The generalization lies in that one of the two players is now a group of players (the attackers), and in that the two opponent players have different strategies; these are the vertices and edges of the graph for the attackers and the defender, respectively. So, the two kinds of players have distinct strategy sets, although the individual strategies (vertices and edges) may intersect. Hence, the inexistence of pure Nash equilibria for our graph-theoretic game does not follow from the inexistence of pure Nash equilibria for Matching Pennies.

In the so called *Interdependent Security* games studied by Kearns and Ortiz [13], a large number of players must make individual decisions related to security. The ultimate safety of each player may depend in a complex way on the actions of the entire population. Our graph-theoretic game may be seen as a particular instance of Interdependent Security games with some kind of limited interdependence: there is interdependence between the actions of the defender and an attacker, while there is no interdependence among the actions of attackers.

Aspnes et al. [2] consider a variation of Interdependent Security games. In particular, they consider an interesting graph-theoretic game with a similar security flavor, modeling containment of the spread of *viruses* on a network with installable *antivirus* software. In the graph-theoretic game of Aspnes et al., the antivirus software may be installed at individual nodes; a virus damages a node if it can reach the node starting at a random initial node and proceeding to it without crossing a node with installed antivirus software. Aspnes et al. [2] prove several algorithmic properties for their graph-theoretic game and establish connections to a certain graph-theoretic problem called *Sum-of-Squares Partition*.

A particular graph-theoretic game with two players, and its connections to the *k-server problem* and *network design*, has been studied by Alon et al. [1].

Bonifaci et al. [4, 5] have independently studied *uniformly mixed Nash equilibria* where each player uses a uniform probability distribution on its support. Our Matching Nash equilibria are uniformly mixed Nash equilibria that satisfy additional (graph-theoretic) requirements.

1.5 Road Map

Section 2 provides the framework for our study. Section 3 provides a characterization of (mixed) Nash equilibria. Section 4 presents several structural conditions that are either necessary or sufficient for Nash equilibria. Matching Nash equilibria are treated in Sect. 5. Section 6 considers Matching Nash equilibria for the special case of bipartite graphs. We conclude, in Sect. 7, with a discussion of our results and some open problems.

2 Framework

Throughout, we consider an undirected graph $G = G(V, E)$ with no isolated vertices; G is *non-trivial* whenever it has more than one edges, otherwise it is *trivial*. We will sometimes model an edge as the set of its two end vertices; we also say that a vertex is *incident* to an edge (or that the edge is incident to the vertex) if the vertex is one of the two end vertices of the edge. Sections 2.1 and 2.2 provide a summary of tools from Game Theory and Graph Theory, respectively, we shall employ. For more background on Game Theory and Graph Theory, we refer the reader to the authoritative textbooks [23] and [27], respectively.

2.1 Game Theory

2.1.1 The Strategic Game $\Pi(G)$

Associated with G is a *strategic game* $\Pi(G) = \langle \mathcal{N}, \{S_i\}_{i \in \mathcal{N}}, \{IP\}_{i \in \mathcal{N}} \rangle$ on G defined as follows:

- The set of *players* is $\mathcal{N} = \mathcal{N}_{vp} \cup \mathcal{N}_{ep}$, where:
 - \mathcal{N}_{vp} is a finite set of v *vertex* players vp_i , called *attackers*, $1 \leq i \leq v$;
 - \mathcal{N}_{ep} is a singleton set of an *edge* player ep , called *defender*.
- The strategy sets of the players are as follows:
 - The *strategy set* S_i of vertex player vp_i is V .
 - The *strategy set* S_{ep} of the edge player ep is E .

So, the *strategy set* \mathcal{S} of the game is $\mathcal{S} = \left(\prod_{i \in \mathcal{N}_{vp}} S_i \right) \times S_{ep} = V^v \times E$.
- Fix an arbitrary *strategy profile* $\mathbf{s} = \langle s_1, \dots, s_v, s_{ep} \rangle \in \mathcal{S}$, also called a *profile*.
 - The *Individual Profit* of vertex player vp_i is a function $IP_{\mathbf{s}}(i) : \mathcal{S} \rightarrow \{0, 1\}$ such that $IP_{\mathbf{s}}(i) = \begin{cases} 0, & s_i \in s_{ep} \\ 1, & s_i \notin s_{ep} \end{cases}$; intuitively, the vertex player vp_i receives 1 if it is not caught by the edge player, and 0 otherwise.
 - The *Individual Profit* of the edge player ep is a function $IP_{\mathbf{s}}(ep) : \mathcal{S} \rightarrow \mathbb{N}$ such that $IP_{\mathbf{s}}(ep) = |\{i : s_i \in s_{ep}\}|$; intuitively, the edge player ep receives the number of vertex players it catches.

2.1.2 Pure Strategies and Pure Nash Equilibria

The profile \mathbf{s} is a *pure Nash equilibrium* [21, 22] if for each player $i \in \mathcal{N}$, it maximizes $IP_{\mathbf{s}}(i)$ over all profiles \mathbf{t} that differ from \mathbf{s} only with respect to the strategy of player i . Intuitively, in a pure Nash equilibrium, no vertex player (resp., the edge player) can (resp., cannot) improve its Individual Profit by switching to a different vertex (resp., edge). In other words, a pure Nash equilibrium is a *local maximizer* for the Individual Profit of each player. Say that G admits a pure Nash equilibrium if there is a pure Nash equilibrium for the strategic game $\Pi(G)$.

2.1.3 Mixed Strategies

A *mixed strategy* for player $i \in \mathcal{N}$ is a probability distribution over its strategy set S_i ; thus, a mixed strategy for a vertex player (resp., the edge player) is a probability distribution over vertices (resp., over edges) of G . A *mixed profile* $\mathbf{s} = \langle s_1, \dots, s_v, s_{ep} \rangle$, or *profile* for short, is a collection of mixed strategies, one for each player; so, $s_i(v)$ is the probability that the vertex player vp_i chooses vertex v and $s_{ep}(e)$ is the probability that the edge player ep chooses edge e .

The *support* of a player $i \in \mathcal{N}$ in the mixed profile \mathbf{s} , denoted $\text{Support}_{\mathbf{s}}(i)$, is the set of pure strategies in its strategy set to which i assigns strictly positive probability in \mathbf{s} . Denote $\text{Support}_{\mathbf{s}}(vp) = \bigcup_{i \in \mathcal{N}_{vp}} \text{Support}_{\mathbf{s}}(i)$; so, $\text{Support}_{\mathbf{s}}(vp)$ contains all pure strategies (that is, vertices) to which some vertex player assigns a strictly positive probability in \mathbf{s} ; $\text{Support}_{\mathbf{s}}(vp)$ will be called the *support* of the vertex players. Denote $\text{Edges}_{\mathbf{s}}(v) = \{(u, v) \in E : (u, v) \in \text{Support}_{\mathbf{s}}(ep)\}$. So, $\text{Edges}_{\mathbf{s}}(v)$ contains all edges incident to v that are included in the support of the edge player.

We shall often deal with profiles with some special structure. A mixed profile is *uniform* if each player uses a uniform probability distribution on its support. Consider a uniform profile \mathbf{s} . Then, for each vertex player $vp_i \in \mathcal{N}_{vp}$, for each vertex $v \in \text{Support}_{\mathbf{s}}(i)$, $s_i(v) = \frac{1}{|\text{Support}_{\mathbf{s}}(i)|}$; for the edge player ep , for each $e \in \text{Support}_{\mathbf{s}}(ep)$, $s_{ep}(e) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$. A profile \mathbf{s} is *vp-symmetric* if for all vertex players $vp_i, vp_k \in \mathcal{N}_{vp}$, $\text{Support}_{\mathbf{s}}(i) = \text{Support}_{\mathbf{s}}(k)$. Clearly, a uniform, vp-symmetric profile is completely determined by the support of the vertex players and the support of the edge player.

2.1.4 Probabilities and Expectations

We now determine some probabilities and expectations according to the profile \mathbf{s} that will be of interest. For a vertex $v \in V$, denote $\text{Hit}(v)$ the event that the edge player ep chooses an edge that contains the vertex v . Denote as $P_{\mathbf{s}}(\text{Hit}(v))$ the probability (according to \mathbf{s}) of the event $\text{Hit}(v)$ occurring. Clearly,

$$P_{\mathbf{s}}(\text{Hit}(v)) = \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e).$$

For a vertex $v \in V$, denote as $VP_s(v)$ the expected number of vertex players choosing vertex v according to s ; so,

$$VP_s(v) = \sum_{i \in \mathcal{N}_{vp}} s_i(v).$$

Clearly, for a vertex $v \notin \text{Support}_s(vp)$, $VP_s(v) = 0$. Also, in a symmetric, vp-uniform profile s , for a vertex $v \in \text{Support}_s(vp)$, $VP_s(v) = \sum_{i \in \mathcal{N}_{vp}} s_i(v) = \frac{v}{|\text{Support}_s(vp)|}$. For each edge $e = (u, v) \in E$, denote as $VP_s(e)$ the expected number of vertex players choosing either the vertex u or the vertex v according to s ; so,

$$\begin{aligned} VP_s(e) &= VP_s(u) + VP_s(v) \\ &= \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)). \end{aligned}$$

2.1.5 Expected Individual Profit and Conditional Expected Individual Profits

A mixed profile s induces an *Expected Individual Profit* $IP_s(i)$ for each player $i \in \mathcal{N}$, which is the expectation according to s of the Individual Profit of player i .

Induced by the mixed profile s is also the *Conditional Expected Individual Profit* $IP_s(i, v)$ of vertex player $vp_i \in \mathcal{N}_{vp}$ on vertex $v \in V$, which is the conditional expectation according to s of the Individual Profit of player vp_i had it chosen vertex v . So,

$$\begin{aligned} IP_s(i, v) &= 1 - P_s(\text{Hit}(v)) \\ &= 1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e). \end{aligned}$$

Clearly, for the vertex player $vp_i \in \mathcal{N}_{vp}$,

$$\begin{aligned} IP_s(i) &= \sum_{v \in V} s_i(v) \cdot IP_s(i, v) \\ &= \sum_{v \in V} s_i(v) \cdot \left(1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e) \right). \end{aligned}$$

Finally, induced by the mixed profile s is the *Conditional Expected Individual Profit* $IP_s(ep, e)$ of the edge player ep on edge $e = (u, v) \in E$, which is the conditional expectation according to s of the Individual Profit of player ep had it chosen edge e . So,

$$\begin{aligned} IP_s(ep, e) &= VP_s(e) \\ &= \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)). \end{aligned}$$

Clearly, for the edge player ep ,

$$\begin{aligned} IP_s(ep) &= \sum_{e \in E} s_{ep}(e) \cdot IP_s(ep, e) \\ &= \sum_{e=(u,v) \in E} s_{ep}(e) \cdot \left(\sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)) \right). \end{aligned}$$

2.1.6 Mixed Nash Equilibria

The mixed profile \mathbf{s} is a *mixed Nash equilibrium* [21, 22] if for each player $i \in \mathcal{N}$, it maximizes $IP_s(i)$ over all mixed profiles \mathbf{t} that differ from \mathbf{s} only with respect to the mixed strategy of player i . In other words, a Nash equilibrium \mathbf{s} is a *local maximizer* for the Expected Individual Profit of each player. By Nash’s celebrated result [21, 22], there is at least one mixed Nash equilibrium for the strategic game $\Pi(G)$; so, every graph G admits a mixed Nash equilibrium.

The particular definition of Expected Individual Profits implies that a Nash equilibrium has two significant properties:

- First, for each vertex player $vp_i \in \mathcal{N}_{vp}$ and vertex $v \in V$ such that $s_i(v) > 0$, all Conditional Expected Individual Profits $IP_s(i, v)$ are the same and no less than any Conditional Expected Individual Profit $IP_s(i, v')$ with $s_i(v') = 0$. It follows that for each vertex player vp_i , for any vertex $v \in \text{Support}_s(i)$,

$$IP_s(i) = 1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e).$$

Thus, the Expected Individual Profit of a vertex player in a Nash equilibrium is determined by any vertex in its support and the mixed strategy of the edge player.

- Second, for each edge $e \in E$ such that $s_{ep}(e) > 0$, all Conditional Expected Individual Profits $IP_s(ep, e)$ are the same and no less than any Conditional Expected Individual Profit $IP_s(ep, e')$ with $s_{ep}(e') = 0$. It follows that for the edge player ep , for any edge $(u, v) \in \text{Support}_s(ep)$,

$$IP_s(ep) = \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)).$$

Thus, the Expected Individual Profit of the edge player in a Nash equilibrium is determined by any edge in its support and the mixed strategies of the vertex players.

We now prove a simple but crucial fact about mixed Nash equilibria:

Lemma 2.1 *Fix a mixed Nash Equilibrium \mathbf{s} . Then, for any pair of vertex players vp_i and vp_k , $IP_s(i) = IP_s(k)$.*

Proof Assume, by way of contradiction, that there are vertex players vp_i and vp_k such that $IP_s(i) \neq IP_s(k)$. Assume, without loss of generality, that $IP_s(i) < IP_s(k)$. Recall that $IP_s(k) = 1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e)$ for any vertex $v \in \text{Support}_s(k)$. Construct

from \mathbf{s} a mixed profile \mathbf{s}' by changing s_i to s_k . For any vertex $v \in \text{Support}_{\mathbf{s}'}(i)$, our construction implies that $v \in \text{Support}_{\mathbf{s}}(k)$. So,

$$\begin{aligned} \text{IP}_{\mathbf{s}'}(i) &= 1 - \sum_{e \in \text{Edges}_{\mathbf{s}'}(v)} s'_{ep}(e) \quad (\text{since } v' \in \text{Support}_{\mathbf{s}'}(i)) \\ &= 1 - \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e) \quad (\text{since } s'_{ep} = s_{ep}) \\ &= \text{IP}_{\mathbf{s}}(k) \quad (\text{since } v \in \text{Support}_{\mathbf{s}}(k)) \\ &> \text{IP}_{\mathbf{s}}(i) \quad (\text{by assumption}), \end{aligned}$$

which contradicts the fact that \mathbf{s} is a Nash equilibrium. □

Note that for each vertex player vp_i , there is some vertex v such that $s_i(v) > 0$; since a Nash equilibrium \mathbf{s} maximizes the Individual Profit of the edge player ep , it follows that $\text{IP}_{\mathbf{s}}(ep) > 0$ for a Nash equilibrium \mathbf{s} .

2.2 Graph Theory

For a vertex set $U \subseteq V$, denote $\text{Neigh}_G(U) = \{u \notin U : (u, v) \in E \text{ for some vertex } v \in U\}$; denote $G(U) = (V(U), E(U))$ the subgraph of G induced by the vertices in U . (So, $V(U) = U$ and $E(U) = \{(u, v) : u \in U, v \in U, \text{ and } (u, v) \in E\}$.) For the edge set $F \subseteq E$, denote $\text{Vertices}(F) = \{v \in V : (u, v) \in F \text{ for some vertex } u \in V\}$. For edge set $F \subseteq E$, denote $G(F) = (V(F), E(F))$ the subgraph of G induced by the edges in F . (So, $E(F) = F$ and $V(F) = \{u \in V : (u, v) \in F \text{ for some vertex } v \in V\}$.) Given any vertex set $U \subseteq V$, the graph $G \setminus U$ is obtained by removing from G all vertices of U and their incident edges.

The graph G is *bipartite* if $V = V_1 \cup V_2$ for some disjoint vertex sets $V_1, V_2 \subseteq V$ so that for each edge $(u, v) \in E$, $u \in V_1$ and $v \in V_2$. Call (V_1, V_2) a *bipartition* of the bipartite graph G .

A vertex set $IS \subseteq V$ is an *Independent Set* of the graph G if for all pairs of vertices $u, v \in IS$, $(u, v) \notin E$. A *Maximum Independent Set* is one that has maximum size; denote $\alpha(G)$ the size of a Maximum Independent Set of G .

A *Vertex Cover* of G is a vertex set $VC \subseteq V$ such that for each edge $(u, v) \in E$ either $u \in VC$ or $v \in VC$. A *Minimum Vertex Cover* is one that has a minimum size; denote $\beta(G)$ the size of a Minimum Vertex Cover of G . It is immediate to see that for any graph G , $\alpha(G) + \beta(G) = |V|$. An *Edge Cover* of G is an edge set $EC \subseteq E$ such that for every vertex $v \in V$, there is an edge $(v, u) \in EC$.

A *Matching* of G is a set $M \subseteq E$ of non-incident edges. For an edge $(u, v) \in M$, say that the Matching M *matches* vertex u to vertex v . A *Maximum Matching* is one that has maximum size; denote $\nu(G)$ the size of a Maximum Matching of G . The classical *König-Egerváry Minimax Theorem* [7, 14] shows that for a bipartite graph G , $\beta(G) = \nu(G)$. Implicit in the proof is a polynomial time algorithm to compute a Minimum Vertex Cover of a bipartite graph though computing a Maximum Matching of the graph (see, for example, [3, Theorem 10-2-1, p. 180]). For the class of bipartite graphs, the currently most efficient algorithm to compute a Maximum Matching is due to Feder and Motwani [8] and runs in time $O(\sqrt{|V|} \cdot |E| \cdot \log_{|V|} \frac{|V|^2}{|E|})$.

Fix now a vertex set $U \subseteq V$. The graph G is a U -Expander graph (and the set U is an *Expander* for G) if for each set $U' \subseteq U$, $|U'| \leq |\text{Neigh}_G(U') \cap (V \setminus U)|$. *Hall's Theorem* [12] establishes a necessary and sufficient condition for one of the bipartitions of a bipartite graph to be an Expander for the graph.

Theorem 2.2 (Hall's Theorem) *Consider a bipartite graph $G = (V, E)$ with bipartition (V_1, V_2) . Then, V_1 is an Expander for G if and only if G has a Matching that matches each vertex in V_1 to a vertex in V_2 .*

Hall's Theorem implies a more general result, known as the *Marriage Theorem*, which is due to Frobenius [10] and holds for all graphs.

Theorem 2.3 (Marriage Theorem) *Consider a graph $G = (V, E)$ and a vertex set $U \subseteq V$. Then, U is an Expander for G if and only if G has a Matching that matches each vertex of U to a vertex of $V \setminus U$.*

An *Expanding Independent Set* of the graph G is an Independent Set IS of G such that the complementary vertex set $V \setminus IS$ is an Expander for G . The following is an immediate consequence of the definition of an Expanding Independent Set and Theorem 2.3:

Corollary 2.4 *Consider a graph $G = (V, E)$ with an Expanding Independent Set IS . Then, $|V \setminus IS| \leq |IS|$.*

3 A Characterization of Nash Equilibria

We prove:

Theorem 3.1 (Characterization of Nash Equilibria) *A profile s is a Nash equilibrium if and only if the following two conditions hold:*

- (1) For any vertex $v \in \text{Support}_s(vp)$, $P_s(\text{Hit}(v)) = \min_{v' \in V} P_s(\text{Hit}(v'))$.
- (2) For any edge $e \in \text{Support}_s(ep)$, $VP_s(e) = \max_{e' \in E} VP_s(e')$.

Proof Assume first that s is a mixed Nash Equilibrium. To show (1), consider any vertex $v \in \text{Support}_s(vp)$; so, $v \in \text{Support}_s(i)$ for some vertex player vp_i . Recall that, $IP_s(i) = 1 - \sum_{e \in \text{Edges}_s(v)} sep(e)$.

For any vertex $v' \in \text{Support}_s(i)$, it holds similarly that $IP_s(i) = 1 - \sum_{e \in \text{Edges}_s(v')} sep(e)$. It follows that $\sum_{e \in \text{Edges}_s(v)} sep(e) = \sum_{e \in \text{Edges}_s(v')} sep(e)$ or $P_s(\text{Hit}(v)) = P_s(\text{Hit}(v'))$. So, consider any vertex $v' \notin \text{Support}_s(i)$. Assume, by way of contradiction, that $P_s(\text{Hit}(v')) < P_s(\text{Hit}(v))$, or equivalently that $\sum_{e \in \text{Edges}_s(v')} sep(e) < \sum_{e \in \text{Edges}_s(v)} sep(e)$. Construct from s the mixed profile s' by only changing s_i to

s'_i so that $v' \in \text{Support}_{s'}(i)$. Then,

$$\begin{aligned} \text{IP}_{s'}(i) &= 1 - \sum_{e \in \text{Edges}_{s'}(v')} s'_{ep}(e) \quad (\text{since } v' \in \text{Support}_{s'}(i)) \\ &= 1 - \sum_{e \in \text{Edges}_s(v')} s_{ep}(e) \quad (\text{since } s'_{ep} = s_{ep}) \\ &> 1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e) \quad (\text{by assumption}) \\ &= \text{IP}_s(i) \quad (\text{since } v \in \text{Support}_s(i)) \end{aligned}$$

which contradicts the fact that \mathbf{s} is a Nash equilibrium.

To show (2), consider any edge $e \in \text{Support}_s(ep)$. Recall that, $\text{IP}_s(ep) = \text{VP}_s(e)$. For any edge $e' \in \text{Support}_s(ep)$, it similarly holds that $\text{IP}_s(ep) = \text{VP}_s(e')$. It follows that, $\text{VP}_s(e) = \text{VP}_s(e')$. So, consider any edge $e' \notin \text{Support}_s(ep)$. Assume, by way of contradiction, that $\text{VP}_s(e') > \text{VP}_s(e)$. Construct from \mathbf{s} the mixed profile \mathbf{s}' by only changing s_{ep} to s'_{ep} so that $e' \in \text{Support}_{s'}(ep)$. Then,

$$\begin{aligned} \text{IP}_{s'}(ep) &= \text{VP}_{s'}(e') \quad (\text{since } e' \in \text{Support}_{s'}(ep)) \\ &= \text{VP}_s(e') \quad (\text{since } s'_i = s_i \text{ for all vertex players } vp_i \in \mathcal{N}_{vp}) \\ &> \text{VP}_s(e) \quad (\text{by assumption}) \\ &= \text{IP}_s(ep) \quad (\text{since } e \in \text{Support}_s(ep)) \end{aligned}$$

which contradicts the fact that \mathbf{s} is a Nash equilibrium.

Assume now that \mathbf{s} is a mixed profile that satisfies conditions (1) and (2). We will prove that \mathbf{s} is a (mixed) Nash equilibrium.

- Consider first any vertex player vp_i . Then, for any vertex $v \in \text{Support}_s(i)$,

$$\begin{aligned} \text{IP}_s(i) &= 1 - \sum_{e \in \text{Edges}_s(v)} s_{ep}(e) \quad (\text{since } v \in \text{Support}_s(i)) \\ &\geq 1 - \sum_{e \in \text{Edges}_s(v')} s_{ep}(e) \quad (\text{by condition (1)}), \end{aligned}$$

for any vertex $v' \in V$. So, the vertex player vp_i cannot increase its Expected Individual Profit according to \mathbf{s} by changing its mixed strategy s_i so that its support would include vertex v' . Since its Expected individual Profit only depends on its support (and not on its probability distribution), it follows that the vertex player vp_i cannot increase its Expected Individual Profit by changing its mixed strategy.

- Consider now the edge player ep . Then, for any edge $e \in \text{Support}_s(ep)$,

$$\begin{aligned} \text{IP}_s(ep) &= \text{VP}_s(e) \quad (\text{since } e \in \text{Support}_s(ep)) \\ &\geq \text{VP}_s(e') \quad (\text{by condition (2)}), \end{aligned}$$

for any edge $e' \in E$. So, player ep cannot increase its Expected Individual Profit according to \mathbf{s} by changing its mixed strategy s_{ep} so that its support would include

edge e' . Since its Expected Individual Profit only depends on its support (and not on its probability distribution), it follows that the edge player ep cannot increase its Expected Individual Profit by changing its mixed strategy.

Hence, it follows that \mathbf{s} is a Nash equilibrium. \square

4 Structure of Nash Equilibria

In this section, we prove several graph-theoretic properties of (mixed) Nash equilibria. Necessary and sufficient graph-theoretic conditions are presented in Sects. 4.1 and 4.2, respectively.

4.1 Necessary Conditions

In this section, we present our necessary, graph-theoretic conditions for Nash equilibria. We first prove:

Proposition 4.1 *For a Nash Equilibrium \mathbf{s} , $\text{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G .*

Proof Assume, by way of contradiction, that $\text{Support}_{\mathbf{s}}(ep)$ is not an Edge Cover of G . Consider any vertex $v \in V$ such that $v \notin \text{Vertices}(\text{Support}_{\mathbf{s}}(ep))$. Thus, $\text{Edges}_{\mathbf{s}}(v) = \emptyset$ and $P_{\mathbf{s}}(\text{Hit}(v)) = 0$.

Since \mathbf{s} is a local maximizer for the Expected Individual Profit of each player $vp_i \in \mathcal{N}_{vp}$, which is at most 1, it follows that vertex player vp_i chooses some such v with probability 1 while $s_i(u) = 0$ for each vertex $u \in \text{Vertices}(\text{Support}_{\mathbf{s}}(ep))$. It follows that for each edge $e = (u, v) \in \text{Support}_{\mathbf{s}}(ep)$,

$$\begin{aligned} \text{VP}(e) &= \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)) \\ &= 0, \end{aligned}$$

so that


$$\begin{aligned} \text{IP}_{\mathbf{s}}(ep) &= \sum_{e \in \text{Support}_{\mathbf{s}}(ep)} s_{ep}(e) \cdot \text{VP}_{\mathbf{s}}(e) \\ &= 0. \end{aligned}$$

Since \mathbf{s} is a Nash equilibrium, $\text{IP}_{\mathbf{s}}(ep) > 0$. A contradiction. \square

We continue to prove:

Proposition 4.2 *For a Nash Equilibrium \mathbf{s} , $\text{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of the graph $G(\text{Support}_{\mathbf{s}}(ep))$.*

Fig. 2 The graph G used in the proof of Proposition 4.4



Proof Assume, by way of contradiction, that $\text{Support}_s(vp)$ is *not* a Vertex Cover of the graph induced by $\text{Support}_s(ep)$. Consider any edge $e \in \text{Support}_s(ep)$ such that $e \notin \text{Edges}(\text{Support}_s(vp))$. Thus, $\text{Vertices}_s(e) = \emptyset$ and $\text{VP}_s(e) = 0$. Since s is a local maximizer for the Expected Individual Profit of the edge player ep , it follows that $s_{ep}(e) = 0$. So, $e \notin \text{Support}_s(ep)$. A contradiction. \square

We remark that the necessary conditions in Propositions 4.1 and 4.2 express *covering* properties of Nash equilibria. We now use Proposition 4.1 to prove:

Theorem 4.3 *The graph G admits no pure Nash equilibrium unless it is trivial.*

Proof Note first that the trivial graph admits the (trivial) pure Nash equilibrium where the edge player chooses the single edge and each vertex player chooses any of the two vertices. So consider a non-trivial graph G . Assume, by way of contradiction, that G admits a pure Nash equilibrium s . Clearly, the support of the edge player is a single edge. Proposition 4.1 implies that the graph G has an Edge Cover consisting of a single edge. It follows that G is trivial. A contradiction. \square

Inspired by the necessary graph-theoretic conditions in Propositions 4.1 and 4.2 is the definition of a Covering profile that follows.

Definition 4.1 A *Covering profile* is a profile s such that $\text{Support}_s(ep)$ is an Edge Cover of G and $\text{Support}_s(vp)$ is a Vertex Cover of the graph $G(\text{Support}_s(ep))$.

It is now natural to ask whether a Covering Profile is necessarily a Nash equilibrium. We provide a negative answer to this question:

Proposition 4.4 *A Covering profile is not necessarily a Nash equilibrium.*

Proof Consider the graph G is Fig. 2 and a profile s with $\text{Support}_s(vp) = \{v_1, v_4\}$ and $\text{Support}_s(ep) = \{(v_1, v_2), (v_3, v_4), (v_4, v_5)\}$. Clearly, $\text{Support}_s(ep)$ is an Edge Cover of G and $\text{Support}_s(vp)$ is a Vertex Cover of the graph $G(\text{Support}_s(ep))$. We will prove that s is *not* a Nash equilibrium. Assume, by way of contradiction, that s is a Nash equilibrium. Theorem 3.1 implies that $P_s(\text{Hit}(v_3)) \geq P_s(\text{Hit}(v_4))$, or $s_{ep}((v_3, v_4)) \geq s_{ep}((v_3, v_4)) + s_{ep}((v_4, v_5))$. It follows that $s_{ep}((v_4, v_5)) = 0$, so that $(v_4, v_5) \notin \text{Support}_s(vp)$. A contradiction. \square

4.2 Sufficient Conditions

In this section, we present sufficient, graph-theoretic conditions for Nash equilibria. In particular, we will enrich the definition of a Covering profile with additional conditions; we will then prove that the enriched set of conditions is a set of sufficient conditions for Nash equilibria. We start with the definition of an *Independent Covering* profile.

Definition 4.2 An *Independent Covering profile* is a uniform, vp-symmetric Covering profile \mathbf{s} satisfying the additional conditions:

- (1) $\text{Support}_{\mathbf{s}}(vp)$ is an Independent Set of G .
- (2) Each vertex in $\text{Support}_{\mathbf{s}}(vp)$ is incident to exactly one edge in $\text{Support}_{\mathbf{s}}(ep)$.

We first prove a preliminary property of Independent Covering profiles.

Lemma 4.5 Consider an Independent Covering profile \mathbf{s} . Then, for each edge $e = (u, v) \in \text{Support}_{\mathbf{s}}(ep)$, exactly one of u and v is in $\text{Support}_{\mathbf{s}}(vp)$.

Proof Since $\text{Support}_{\mathbf{s}}(vp)$ is an Independent Set of G (by additional Condition (1)) either $u \in \text{Support}_{\mathbf{s}}(vp)$ or $v \in \text{Support}_{\mathbf{s}}(vp)$ but not both. Since $\text{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of the graph $G(\text{Support}_{\mathbf{s}}(ep))$, it follows that either $u \in \text{Support}_{\mathbf{s}}(vp)$ or $v \in \text{Support}_{\mathbf{s}}(vp)$ or both. It follows that exactly one of u and v is in $\text{Support}_{\mathbf{s}}(vp)$, as needed. \square

We are now ready to prove:

Proposition 4.6 An Independent Covering profile is a Nash equilibrium.

Proof We shall employ the characterization of Nash equilibria in Theorem 3.1 consisting of Conditions (1) and (2). We first prove Condition (1). Consider any vertex $v \in \text{Support}_{\mathbf{s}}(vp)$.

- Consider first a vertex $v' \in \text{Support}_{\mathbf{s}}(vp)$. By additional Condition (2), v' is incident to exactly one edge $e \in \text{Support}_{\mathbf{s}}(ep)$. So, $P_{\mathbf{s}}(\text{Hit}(v')) = s_{ep}(e)$. Since \mathbf{s} is uniform, $s_{ep}(e) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$. It follows that $P_{\mathbf{s}}(\text{Hit}(v')) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$. In particular, $P_{\mathbf{s}}(\text{Hit}(v)) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$.
- Consider now a vertex $v' \notin \text{Support}_{\mathbf{s}}(vp)$. Since $\text{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G (by Condition (1) in the definition of a Covering Profile), there is some edge $e \in \text{Support}_{\mathbf{s}}(ep)$ such that $v' \in e$. Then, clearly, $P_{\mathbf{s}}(\text{Hit}(v')) \geq s_{ep}(e) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$ (since \mathbf{s} is uniform).

It follows that $P_{\mathbf{s}}(\text{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s}}(\text{Hit}(v'))$, which proves (1).

We now prove Condition (2). Consider any edge $e = (u, v) \in \text{Support}_{\mathbf{s}}(ep)$.

- Consider first an edge $e' = (u', v') \in \text{Support}_{\mathbf{s}}(ep)$. By Lemma 4.5, exactly one of u' and v' is in $\text{Support}_{\mathbf{s}}(ep)$. So, assume, without loss of generality, that $u' \in \text{Support}_{\mathbf{s}}(vp)$ and $v' \notin \text{Support}_{\mathbf{s}}(vp)$. Then, $\text{VP}_{\mathbf{s}}(e') = \text{VP}_{\mathbf{s}}(u')$. Since \mathbf{s} is

vp-symmetric and uniform, $VP_s(u') = \frac{v}{|\text{Support}_s(vp)|}$. So, $VP_s(e') = \frac{v}{|\text{Support}_s(vp)|}$. In particular, $VP_s(e) = \frac{v}{|\text{Support}_s(vp)|}$.

- Consider now an edge $e' = (u', v') \notin \text{Support}_s(ep)$. Since $\text{Support}_s(vp)$ is an Independent Set (by additional Condition (2)), either $u' \in \text{Support}_s(vp)$ or $v' \in \text{Support}_s(vp)$ but not both. So, assume, without loss of generality, that $v' \notin \text{Support}_s(vp)$. Then, $VP_s(e') \leq VP_s(u')$. If $u' \notin \text{Support}_s(vp)$, then $VP_s(u') = 0$; else, since s is vp-symmetric and uniform, $VP_s(u') = \frac{v}{|\text{Support}_s(vp)|}$. It follows that $VP_s(e') \leq \frac{v}{|\text{Support}_s(vp)|}$.

It follows that $VP_s(e) = \max_{e' \in E} VP_s(e')$, which proves (2).

By Theorem 3.1, it follows that s is a Nash equilibrium. □

We continue to prove a necessary condition for Independent Covering profiles.

Proposition 4.7 *For an Independent Covering profile s , there is a Matching $M \subseteq \text{Support}_s(ep)$ that matches each vertex in $V \setminus \text{Support}_s(vp)$ to some vertex in $\text{Support}_s(vp)$.*

Proof Consider any vertex $v \in V \setminus \text{Support}_s(vp)$. Since $\text{Support}_s(ep)$ is an Edge Cover of G , there is an edge $(u, v) \in \text{Support}_s(ep)$ for some $u \in V$. Lemma 4.5 implies that $u \in \text{Support}_s(vp)$. By additional Condition (2) in the definition of an Independent Covering profile, vertex u is incident to exactly one edge in $\text{Support}_s(ep)$. So, clearly, the set $\{(u, v) \in \text{Support}_s(ep) : v \in \text{Support}_s(vp)\}$ is a Matching contained in $\text{Support}_s(ep)$ that maps each vertex in $V \setminus \text{Support}_s(vp)$ to a vertex in $\text{Support}_s(vp)$, as needed □

An immediate consequence of Proposition 4.7 follows:

Corollary 4.8 *For an Independent Covering profile s , $|\text{Support}_s(vp)| \geq |V \setminus \text{Support}_s(vp)|$.*

Propositions 4.6 and 4.7 together imply that an Independent Covering profile is a Nash equilibrium which induces a suitable Matching contained in the support of the edge player. So, in the rest of this paper, an Independent Covering profile will be called a *Matching Nash equilibrium*.

5 Matching Nash Equilibria

A characterization of graphs admitting Matching Nash equilibria is proved in Sect. 5.1; in Sect. 5.2, the characterization is turned into a non-deterministic, polynomial time algorithm to compute a Matching Nash equilibrium. Section 5.3 uses the characterization to establish the non-triviality of the class of graphs admitting Matching Nash equilibria.

5.1 Characterization

We prove:

Theorem 5.1 *A graph G admits a Matching Nash equilibrium if and only if G has an Expanding Independent Set.*

Proof Assume first that G has an Expanding Independent Set IS . The proof that G admits a Matching Nash equilibrium is constructive. For emphasis, the constructive parts of the proof will be boxed.

Since $V \setminus IS$ is an Expander for G , Theorem 2.3 implies that G has a Matching M that matches each vertex in $V \setminus IS$ to a vertex in $V \setminus (V \setminus IS) = IS$. Use the Matching M to partition the Expanding Independent Set IS into vertex sets IS_1 and IS_2 as follows:

- IS_1 consists of vertices $v \in IS$ for which there is an edge $(v, u) \in M$ for some vertex $u \in V \setminus IS$.
- IS_2 consists of the remaining vertices of IS ; that is, $IS_2 = IS \setminus IS_1$.

Intuitively, IS_1 consists of vertices in IS to which M maps some vertex in $V \setminus IS$.

Since IS is an Independent Set and G has no isolated vertices, it follows that for each vertex $v \in IS_2$, there is a vertex $u \in V \setminus IS$ such that $(v, u) \in E$. So, use IS_2 to construct an edge set M_1 as follows:

- For each vertex $v \in IS_2$, M_1 contains exactly one edge (v, u) , where $u \in V \setminus IS$.

Note that by the construction of the edge set M_1 , each edge in M_1 is incident to exactly one vertex in IS_2 .

Use now the sets IS , M and M_1 to construct a uniform, vp-symmetric profile \mathbf{s} as follows:

- $\text{Support}_{\mathbf{s}}(vp) := IS$ and $\text{Support}_{\mathbf{s}}(ep) := M \cup M_1$.

We now prove that the profile \mathbf{s} is a Matching Nash equilibrium. By Proposition 4.6, it suffices to prove that \mathbf{s} is an Independent Covering profile.

- By construction, \mathbf{s} is uniform and vp-symmetric.
- To prove that \mathbf{s} is a Covering profile, we proceed as follows:
 - We first prove that $\text{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G . Consider any vertex $v \in V$. There are two cases to consider:
 - * Assume first that $v \in V \setminus IS$. Recall that v is matched by M to some vertex in IS . So, v is incident to an edge in $M_1 \subseteq \text{Support}_{\mathbf{s}}(ep)$.
 - * Assume now that $v \in IS$. Either $v \in IS_1$ or $v \in IS_2$. If $v \in IS_1$, then by the definition of IS_1 , v is incident to some edge in M . If $v \in IS_2$, then by the definition of M_1 , v is incident to some edge in M_1 .
 So, in all cases, v is incident to some edge in $M \cup M_1 = \text{Support}_{\mathbf{s}}(ep)$. So, $\text{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G .
 - We now prove that $\text{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of $G(\text{Support}_{\mathbf{s}}(ep))$. Consider any edge $e \in \text{Support}_{\mathbf{s}}(ep)$. There are two cases to consider:

- * Assume first that $e \in M$. Recall that M maps each vertex in $V \setminus IS$ to a vertex in IS . Thus, e is incident to a vertex in $IS = \text{Support}_s(vp)$.
- * Assume now that $e \in M_1$. Recall that each edge in M_1 corresponds to some vertex in $IS_2 \subseteq IS = \text{Support}_s(vp)$. Hence, e is incident to some vertex in $\text{Support}_s(vp)$.

So, in all cases, e is incident to some vertex in $\text{Support}_s(vp)$. Hence, $\text{Support}_s(vp)$ is a Vertex Cover of $G(\text{Support}_s(ep))$.

Our proof that s is a Covering profile is now complete. We proceed to prove the additional conditions in the definition of an Independent Covering profile.

- To prove additional Condition (1), recall that, by construction, $\text{Support}_s(vp) = IS$. Since IS is an Independent Set of G , the claim follows.
- We finally prove additional Condition (2). Consider any vertex $v \in \text{Support}_s(vp) = IS$. There are two cases to consider
 - Assume first that $v \in IS_1$. By definition of IS_1 , there is an edge $(v, u) \in M$ (for some vertex $u \in V \setminus IS$). Since M is a Matching, there is exactly one such edge.
 - Assume now that $v \in IS_2$. By definition of IS_2 , there is an edge $(v, u) \in M_1$ (for some vertex $u \in V \setminus IS$).

So, in all cases, v is incident to exactly one edge in $M \cup M_1 = \text{Support}_s(ep)$; thus, additional Condition (2) follows.

Hence, s is an Independent Covering profile. Proposition 4.6 implies that s is a Nash equilibrium.

Assume now that G admits a Matching Nash equilibrium (that is, an Independent Covering profile). By additional Condition (2), in the definition of an Independent Covering profile, $\text{Support}_s(vp)$ is an Independent Set of G . By Proposition 4.7, there is a Matching $M \subseteq \text{Support}_s(ep)$ that maps each vertex in $V \setminus \text{Support}_s(vp)$ to some vertex in $\text{Support}_s(vp)$. Theorem 2.3 implies that $V \setminus \text{Support}_s(vp)$ is an Expander for G . It follows that $\text{Support}_s(vp)$ is an Expanding Independent Set, as needed. □

5.2 Computation

The sufficiency part of the proof of Theorem 5.1 immediately yields a non-deterministic algorithm MatchingNE to compute a Matching Nash equilibrium for a graph satisfying the characterization. Figure 3 presents the algorithm in suitable pseudocode.

The non-deterministic steps (1) and (2) in the algorithm MatchingNE correspond to choosing an Expanding Independent Set for the graph G and a corresponding Matching M , respectively. (Recall that the first exists by assumption, while the second exists by Theorem 2.3.)

Clearly, to turn this non-deterministic algorithm into a polynomial-time, deterministic algorithm, it is necessary to replace (whenever possible) the two non-deterministic steps (1) and (2) with corresponding deterministic, polynomial time algorithms. We are able to achieve such replacement for the non-deterministic step (2).

To do so, define the graph $G(IS)$ *shrunk* by the Expanding Independent Set IS as the subgraph of G that does not include any edges (u, v) with both $v \in V \setminus IS$ and $v \in V \setminus IS$.

Algorithm MatchingNE

INPUT: A graph $G(V, E)$ assumed to have an Expanding Independent Set.

OUTPUT: The support of the vertex players and the support of the edge player of a Matching Nash equilibrium for G .

- (1) Choose an Expanding Independent Set IS of G .
- (2) Choose a Matching M of G that matches each vertex in $V \setminus IS$ to a vertex in $V \setminus (V \setminus IS) = IS$.
- (3) Partition IS into vertex sets IS_1 and IS_2 as follows:
 - IS_1 consists of vertices $v \in IS$ for which there is an edge $(v, u) \in M$ for some vertex $u \in V \setminus IS$.
 - IS_2 consists of the remaining vertices of IS ; that is, $IS_2 = IS \setminus IS_1$.
- (4) Determine an edge set $M_1 \subseteq E$ such that for each vertex $v \in IS_2$, M_1 contains exactly one edge (v, u) , where $u \in V \setminus IS$.
- (5) $\text{Support}_s(vp) := IS$ and $\text{Support}_s(ep) := M \cup M_1$.

Fig. 3 The non-deterministic, polynomial time algorithm MatchingNE

- By assumption, IS is an Independent Set of the graph G . The construction of the graph $G(IS)$ implies that IS is also an Independent Set of the graph $G(IS)$.
- By construction, $V \setminus IS$ is an Independent Set of the graph $G(IS)$.

It follows that $(IS, V \setminus IS)$ is a bipartition of the graph $G(IS)$, so $G(IS)$ is bipartite. Consider a Maximum Matching M of the bipartite graph $G(IS)$. Since IS is an Expanding Independent Set of the graph G , Corollary 2.4 implies that $|V \setminus IS| \leq |IS|$. It follows that $|M| \leq |V \setminus IS|$.

We will prove that, in fact, $|M| = |V \setminus IS|$. Since $V \setminus IS$ is an Expander for G , Theorem 2.3 implies that G has a Matching M' that matches each vertex in $V \setminus IS$ to a vertex in IS ; so, $|M'| \geq |V \setminus IS|$. Clearly, M' induces a Matching M'' for the graph $G(IS)$ that matches each vertex in $V \setminus IS$ to a vertex in IS ; clearly, $|M''| = |V \setminus IS|$. Since $|M| \geq |M''|$, $|M| \leq |V \setminus IS|$ and $|M''| = |V \setminus IS|$, it follows that $|M| = |V \setminus IS|$.

Since $|M| = |V \setminus IS|$, it follows that M matches each vertex in $V \setminus IS$ to a vertex in IS . Since $G(IS)$ is a subgraph of G , the Matching M of $G(IS)$ is also a Matching of G . Hence, we obtain:

Lemma 5.2 *A Maximum Matching M of $G(IS)$ is a Matching of G that matches each vertex in $V \setminus IS$ to a vertex in IS .*

Lemma 5.2 implies that the non-deterministic step (2) in the non-deterministic algorithm MatchingNE can be replaced by a deterministic, polynomial time algorithm to compute a Maximum Matching of the bipartite graph $G(IS)$. The modified algorithm MatchingNE appears in Fig. 4.

Clearly, the dominating step in the modified, still non-deterministic algorithm MatchingNE is step (2): computing a Maximum Matching for a bipartite graph. Hence, it follows:

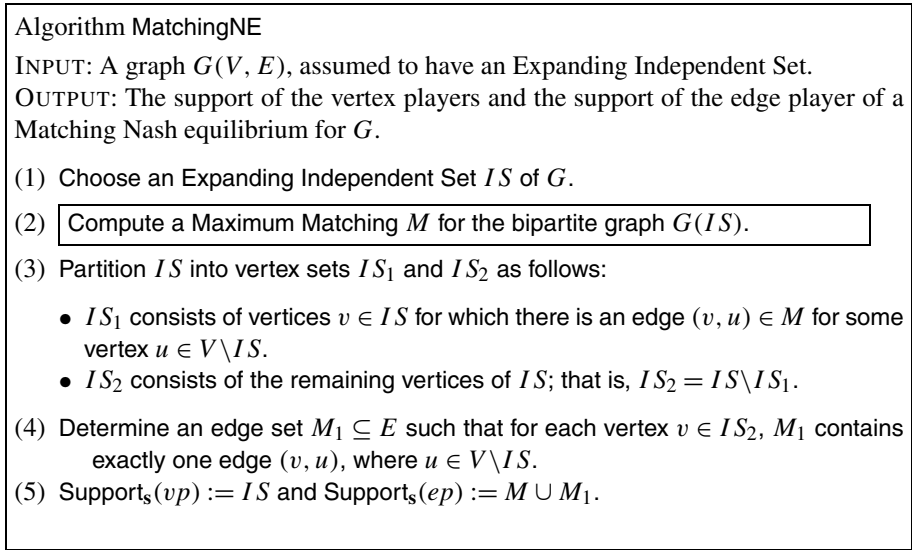


Fig. 4 The modified non-deterministic, polynomial time algorithm MatchingNE

Theorem 5.3 *The modified non-deterministic algorithm MatchingNE computes a Matching Nash equilibrium in time $O(\sqrt{|V|} \cdot |E| \cdot \log_{|V|} \frac{|V|^2}{|E|})$.*

5.3 Non-Triviality

We now prove that the class of graphs admitting a Matching Nash equilibrium is *non-trivial* by presenting a graph that does not admit a Matching Nash equilibrium and a second graph that admits one. Interestingly, the first graph is the *odd cycle* (a cycle with an odd number of vertices); the second graph is the *even cycle* (a cycle with an even number of vertices). We prove:

Proposition 5.4 *An odd cycle admits no Matching Nash equilibrium.*

Proof Consider an odd cycle C_{2n+1} with an odd number of vertices $2n + 1$, where $n \geq 1$. Clearly, the size of a Maximum Independent Set of C_{2n+1} is $\alpha(C_{2n+1}) = n$. Assume, by way of contradiction, that C_{2n+1} has a Matching Nash equilibrium s . Theorem 5.1 implies that C_{2n+1} has an Expanding Independent Set IS . Corollary 2.4 implies that $|IS| \geq n + 1$. However, $|IS| \leq \alpha(C_{2n+1}) = n$. A contradiction. \square

Finally, we prove:

Proposition 5.5 *An even cycle admits a Matching Nash equilibrium.*

Proof Consider an even cycle $C_{2n} = (V(C_{2n}), E(C_{2n}))$ with an even number of vertices $2n$, where $n \geq 1$. Take any Maximum Independent Set IS of C_{2n} ; clearly,

$|IS| = n$. (There are only two Maximum Independent Sets for \mathcal{C}_{2n} : the set of vertices with even index and the set of vertices with odd index.)

Note that there is an one-to-one mapping from $V(\mathcal{C}_{2n}) \setminus IS$ to IS , which maps a vertex to its left neighbor. This induces an one-to-one mapping from sets of vertices $V' \subseteq V(\mathcal{C}_{2n}) \setminus IS$ to sets of neighboring vertices $V'' \subseteq IS$, with $|V''| = |V'|$. This implies that $V(\mathcal{C}_{2n}) \setminus IS$ is an expander for \mathcal{C}_{2n} , so that IS is an Expanding Independent Set of \mathcal{C}_{2n} . Hence, Theorem 5.1 implies that \mathcal{C}_{2n} has a Matching Nash equilibrium, as needed. \square

Propositions 5.4 and 5.5 together imply:

Theorem 5.6 *The class of graphs admitting Matching Nash equilibria is non-trivial.*

Theorem 5.6 implies that deciding whether a given graph admits a Matching Nash equilibrium is a *non-trivial* decision problem.

6 Bipartite Graphs

In this section, we further investigate Matching Nash equilibria for the case of bipartite graphs. In Sect. 6.1, we establish that every bipartite graph admits a Matching Nash equilibrium; the proof is turned, in Sect. 6.1, into a deterministic, polynomial time algorithm to compute a Matching Nash equilibrium for a given bipartite graph.

6.1 Existence

Proposition 6.1 *A bipartite graph has an Expanding Independent Set.*

Proof Consider a bipartite graph $G(V, E)$. Our proof is constructive. For emphasis, the constructive parts will be boxed.

- Choose a Minimum Vertex Cover VC of the graph G .

We will use VC to construct an Expanding Independent Set IS of G .

Consider the bipartition (V_1, V_2) of the bipartite graph G . Consider the induced subgraphs $G_1 = G((VC \cap V_1) \cup (V_2 \setminus VC))$ and $G_2 = G((VC \cap V_2) \cup (V_1 \setminus VC))$. That is, the vertices of G_1 are either vertices of V_1 that also belong to the Vertex Cover VC or vertices of V_2 that do not belong to the Vertex Cover VC ; the vertices of G_2 are either vertices of V_2 that also belong to the Vertex Cover VC or vertices of V_1 that do not belong to the Vertex Cover VC . Since $(VC \cap V_1) \cup (VC \cap V_2)$ is a Vertex Cover of the graph G , there is no edge between $V_1 \setminus VC$ and $V_2 \setminus VC$.

We claim that for each set $U' \subseteq VC \cap V_1$, $|\text{Neigh}_{G_1}(U')| \geq |U'|$. Assume otherwise; that is, assume that $|\text{Neigh}_{G_1}(U')| < |U'|$. Then, the set $(VC \setminus U') \cup \text{Neigh}_{G_1}(U')$ is a Vertex Cover of G which is smaller than VC , a contradiction. So, the claim follows. This implies that the vertex set $VC \cap V_1$ is an Expander for the graph G_1 . Hence, Theorem 2.2 implies that G_1 has a Matching M_1 that matches each vertex in $VC \cap V_1$ to a vertex in $V_2 \setminus VC$. By identical reasoning, we obtain that G_2 has a Matching M_2 that matches each vertex in $VC \cap V_2$ to a vertex in $V_1 \setminus VC$.

- Choose a Matching M_1 of the graph $G_1 = G((VC \cap V_1) \cup (V_2 \setminus VC))$ that matches each vertex in $VC \cap V_1$ to a vertex in $V_2 \setminus VC$.
- Choose a Matching M_2 of the graph $G_2 = G((VC \cap V_2) \cup (V_1 \setminus VC))$ that matches each vertex in $VC \cap V_2$ to a vertex in $V_1 \setminus VC$.
- $M := M_1 \cap M_2$.

Since $VC \cap V_1$ is disjoint from $V_1 \setminus VC$ and $VC \cap V_2$ is disjoint from $V_2 \setminus VC$, it follows that M is a Matching of G . By the properties of the Matchings M_1 and M_2 , it follows that the Matching M matches each vertex in VC to a vertex in $V \setminus VC$. So, set:

- $IS := V \setminus VC$.

Clearly, $IS = V \setminus VC$ is an Expanding Independent Set of G , as needed. □

Hence, Proposition 6.1 and Theorem 5.1 together imply:

Corollary 6.2 *A bipartite graph admits a Matching Nash equilibrium.*

6.2 Computation

Proposition 6.1 implies that choosing an Expanding Independent Set of a bipartite graph reduces to computing a Minimum Vertex Cover of a bipartite graph. Recall that computing a Minimum Vertex Cover of a bipartite graph reduces to computing a Maximum Matching of the bipartite graph. Hence, for the case of bipartite graphs, the non-deterministic step (1) in the modified non-deterministic algorithm MatchingNE can be substituted by a deterministic, polynomial time algorithm to compute a Minimum Vertex Cover VC of a bipartite graph (via reduction to computing a Maximum Matching), followed by setting IS to include all vertices in the complementary vertex set of VC . This results in the deterministic, polynomial time algorithm BipartiteMatchingNE appearing in Fig. 5. Hence, as our main result, we obtain:

Theorem 6.3 *For a bipartite graph $G = (V, E)$, a Matching Nash Equilibrium can be computed in time $O(\sqrt{|V|} \cdot |E| \cdot \log_{|V|} \frac{|V|^2}{|E|})$.*

7 Epilogue

In this work, we introduced a network game with *attackers* and a *defender* as an abstraction of security attacks and defenses for emerging networks like the Internet.

We focused on the Nash equilibria associated with our network game and proved several structural (graph-theoretic) properties of them. In particular, we obtained necessary graph-theoretic conditions for Nash equilibria and sufficient graph-theoretic conditions for them as well. Unfortunately, the set of sufficient conditions does not coincide with the set of necessary ones: it includes two additional conditions which our sufficiency proofs required. It is a challenging open problem to determine a *single*

Algorithm BipartiteMatchingNE

INPUT: A bipartite graph $G(V, E)$.

OUTPUT: The support of the vertex players and the support of the edge player of a Matching Nash equilibrium for G .

- | |
|--|
| (1/a) Compute a Minimum Vertex Cover VC of the graph G through computing a Maximum Matching of G . (1/b) $IS := V \setminus VC$. |
| (2) Compute a Maximum Matching M for the bipartite graph $G(IS)$. |
- (3) Partition IS into vertex sets IS_1 and IS_2 as follows:
- IS_1 consists of vertices $v \in IS$ for which there is an edge $(v, u) \in M$ for some vertex $u \in V \setminus IS$.
 - IS_2 consists of the remaining vertices of IS ; that is, $IS_2 = IS \setminus IS_1$.
- (4) Determine an edge set $M_1 \subseteq E$ such that for each vertex $v \in IS_2$, M_1 contains exactly one edge (v, u) , where $u \in V \setminus IS$.
- (5) $\text{Support}_s(vp) := IS$ and $\text{Support}_s(ep) := M \cup M_1$.

Fig. 5 The deterministic, polynomial time algorithm BipartiteMatchingNE

set of graph-theoretic conditions (hence, a graph-theoretic characterization) that are both necessary and sufficient for a Nash equilibrium.

We used the set of sufficient graph-theoretic conditions we obtained to define the class of Matching Nash equilibria. Since this set extends our necessary graph-theoretic conditions for Nash equilibria, the new class of Matching Nash equilibria is quite natural.

We also obtained a characterization of graphs admitting Matching Nash equilibria: such graphs have an Expanding Independent Set. We do not know yet the complexity of deciding whether a graph has an Expanding Independent Set. So, the complexity of deciding existence of Matching Nash equilibria remains open for the class of *general* graphs. However, for the class of bipartite graphs, we proved that a Matching Nash equilibrium always exists and provided an efficient graph-theoretic algorithm to compute one. (In the subsequent work [17], we identified additional classes of graphs permitting efficient computation of Matching Nash equilibria.)

Our work has identified some rich links between Game Theory and some central concepts in Graph Theory such as Independent Sets, Covers, Expanders and Matchings. Besides their pure scientific interest, we believe that such links will be instrumental to settling the complexity of computing Nash equilibria for our network game and many possible variants of it. Indeed, what happens, for example, if we have many defenders and one attacker? What if the defender protects some subgraph (rather than a single edge)? Some preliminary steps to such and other extensions already appear in [11, 17–19].

References

1. Alon, N., Karp, R.M., Peleg, D., West, D.: A graph-theoretic game and its application to the k -server problem. *SIAM J. Comput.* **24**(1), 78–100 (1995)
2. Aspnes, J., Chang, K., Yampolskiy, A.: Inoculation strategies for victims of viruses and the sum-of-squares problem. In: *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 43–52 (January 2005)
3. Asratian, A.S., Denley, T.M.J., Häggkvist, R.: *Bipartite Graphs and Their Applications*. Cambridge Tracts in Mathematics, vol. 131. Cambridge University Press, Cambridge (1998)
4. Bonifaci, V., Di Iorio, U., Laura, L.: On the complexity of uniformly mixed Nash equilibria and related regular subgraph problems. In: Liskiewicz, M. (ed.) *Proceedings of the 15th International Symposium on Fundamentals of Computation Theory*. Lecture Notes in Computer Science, vol. 3623, pp. 197–208. Springer, New York (2005)
5. Bonifaci, V., Di Iorio, U., Laura, L.: New results on the complexity of uniformly mixed Nash equilibria. In: Deng, X. (ed.) *Proceedings of the First International Workshop on Internet and Network Economics*. Lecture Notes in Computer Science, vol. 3828, pp. 1023–1032. Springer, New York (2005)
6. Cheswick, E.R., Bellovin, S.M.: *Firewalls and Internet Security*. Addison-Wesley, Reading (1994)
7. Egerváry, E.: On combinatorial properties of matrices. *Mat. Fiz. Lapok* **38**, 16–28 (1931) (in Hungarian with German summary)
8. Feder, T., Motwani, R.: Clique compressions, graph partitions and speeding-up algorithms. *J. Comput. Syst. Sci.* **51**(2), 261–272 (1995)
9. Franklin, M., Alto, P., Galil, Z., Yung, M.: Eavesdropping games: a graph-theoretic approach to privacy in distributed systems. *J. ACM* **47**(2), 225–243 (2000)
10. Frobenius, G.: Über Zerlegbare Determinanten. In: *Sitzungsber der Königlich*, vol. XVIII, pp. 274–277. Preussischen Akademie der Wissenschaften (1917)
11. Gelastou, M., Mavronicolas, M., Papadopoulou, V., Philippou, A., Spirakis, P.: The power of the defender. In: *CD-ROM Proceedings of the 2nd International Workshop on Incentive-Based Computing*, in conjunction with the 26th IEEE International Conference on Distributed Computing, July 2006
12. Hall, P.: On representation of subsets. *J. Lond. Math. Soc. Ser. 2* **10**, 26–30 (1935)
13. Kearns, M., Ortiz, L.: Algorithms for interdependent security games. In: Thrun, S., Saul, L. (eds.) *Advances in Neural Information Processing Systems*, vol. 16. MIT Press, Cambridge (2004)
14. König, D.: Graphen und Matrizen. *Mat. Fiz. Lapok* **38**, 116–119 (1931)
15. König, D.: *Theorie der Endlichen und Unendlichen Graphen*. Akademische Verlagsgesellschaft (1936)
16. Markham, T., Payne, C.: Security at the network edge: a distributed firewall architecture. In: *Proceedings of the 2nd DARPA Information Survivability Conference and Exposition*, vol. 1, pp. 279–286 (June 2001)
17. Mavronicolas, M., Papadopoulou, V., Philippou, A., Spirakis, P.: A graph-theoretic network security game. In: Deng, X. (ed.) *Proceedings of the 1st International Workshop on Internet and Network Economics*. Lecture Notes in Computer Science, vol. 3828, pp. 969–978. Springer, New York (2005)
18. Mavronicolas, M., Michael, L., Papadopoulou, V., Philippou, A., Spirakis, P.: The price of defense. In: Kráľovič, R. (ed.) *Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science*. Lecture Notes in Computer Science, vol. 4162, pp. 717–728. Springer, New York (2006)
19. Mavronicolas, M., Papadopoulou, V., Persiano, G., Philippou, A., Spirakis, P.: The price of defense and fractional matchings. In: Chaudhuri, R.S., Das, S.R., Paul, H.S., Tirthapura, S. (eds.) *Proceedings of the 8th International Conference on Distributed Computing and Networking*. Lecture Notes in Computer Science, vol. 4308, pp. 115–126. Springer, New York (2006)
20. Mavronicolas, M., Spirakis, P.: *Algorithmic Game Theory*. Springer, New York (2007, to appear)
21. Nash, J.F.: Equilibrium points in N -person games. *Proc. Natl. Acad. Sci. USA* **36**, 48–49 (1950)
22. Nash, J.F.: Non-cooperative games. *Ann. Math.* **54**(2), 286–295 (1951)
23. Osborne, M., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
24. Spafford, E.H.: The Internet worm: crisis and aftermath. *Commun. ACM* **32**(6), 678–687 (1989)
25. Stallings, W.: *Cryptography and Network Security: Principles and Practice*, 3rd edn. Prentice-Hall, Englewood Cliffs (2003)
26. von Neumann, J.: Zur Theorie der Gesellschaftsspiele. *Math. Ann.* **100**, 295–320 (1928)
27. West, D.B.: *Introduction to Graph Theory*, 2nd edn. Prentice-Hall, Englewood Cliffs (2001)