

A Network Game with Attackers and a Defender: A Survey*

Marios Mavronicolas[†] Vicky Papadopoulou[†] Anna Philippou Paul Spirakis[‡]

Abstract

We survey a research line recently initiated by Mavronicolas *et al.* [14, 15, 16], concerning a *strategic game* on a graph $G(V, E)$ with two confronting classes of randomized players: ν *attackers* who choose vertices and wish to minimize the probability of being caught by the *defender*, who chooses edges and gains the expected number of attackers it catches. So, the defender captures system rationality. In a *Nash equilibrium*, no single player has an incentive to unilaterally deviate from its randomized strategy. The *Price of Defense* is the worst-case ratio, over all Nash equilibria, of the *optimal* gain of the defender (which is ν) over the gain of the defender at a Nash equilibrium. We present a comprehensive collection of trade-offs between the Price of Defense and the computational efficiency of Nash equilibria proved in [14, 15, 16].

- We present 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 proved for the *supports* of the players in all (*mixed*) Nash equilibria.
- We present a reduction of the game to a *Zero-Sum Two-Players Game* that proves that a general Nash equilibrium can be computed via Linear Programming in polynomial time. However, the reduction does not provide any apparent guarantees on the Price of Defense.
- To obtain guarantees on Price of Defense, we present an analysis of several structured Nash equilibria:
 - In a *Matching Nash equilibrium*, the support of the defender is an *Edge Cover* of the graph. Matching Nash equilibria are shown to still be computable in polynomial time, and that they incur a Price of Defense of $\alpha(G)$, the *Independence Number* of G .
 - In a *Perfect Matching Nash equilibrium*, the support of the defender is a *Perfect Matching* of the graph. Perfect Matching Nash equilibria are shown to be computable in polynomial time, and that they incur a Price of Defense of $\frac{|V|}{2}$.
- We consider a generalization of the basic model with an increased power for the defender: it is able to scan a simple *path* of the network instead of a single edge. Deciding existence of a pure Nash equilibrium is shown to be an \mathcal{NP} -complete problem for this model.

*This work was partially supported by the IST Program of the European Union under contract number IST-2004-001907 (DELIS), and by research funds at University of Cyprus.

[†]Department of Computer Science, University of Cyprus, P.O. Box 20537, Nicosia CY-1678, Cyprus. Email: {mavronic,viki,annap}@ucy.ac.cy

[‡]Research Academic Computer Technology Institute (RACTI), Rion, Patras 26500, Greece & Department of Computer Engineering and Informatics, University of Patras, Rion, Patras 26500, Greece. Email: spirakis@cti.gr

1 Introduction

1.1 Motivation and Framework

Here, we overview a recent result line concerning a network game with *attackers* and a *defender*, introduced recently by Mavronicolas *et al.* [14] and further studied in [7, 15, 16]; the game was conceived as an appropriate theoretical model of security *attacks* and *defenses* in emerging networks like the Internet. In this network game, nodes are vulnerable to infection by *threats*, called *attackers*. Available to the network is a security software (or *firewall* [2]), called the *defender*, cleaning some part of the network.

This network game is partially motivated by *Network Edge Security* [13], a new distributed firewall architecture where a firewall is implemented in a *distributed* way and protects the subnetwork spanned by the nodes participating in the distributed implementation. The simplest case where the subnetwork is a single link (with its two incident nodes) offers the initial basis for the theoretical model of Mavronicolas *et al.* [14]. Understanding the mathematical pitfalls 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.

Each attacker (called a *vertex player*) targets a *node* of the network chosen via its own probability distribution on nodes; the defender (called the *edge player*) chooses a single *link* via its own probability distribution on links. A node chosen by an attacker is destroyed unless it crosses the link being cleaned by the defender. The *Individual Profit* of an attacker is the probability that it is not caught; the *Individual Profit* of the defender is the expected number of attackers it catches.

The game with attackers and defenders is partially motivated from the problem of *data integrity* on systems implementing a *distributed database*. A distributed database is realized through a set of nodes that store parts of the database, and cooperate each other in a distributed way for the database management. Given a node holding some database, one way to detect corruption of the data is to store a hashed version (or *fingerprint*) of the data on some other node. Checking the database and the fingerprint you can detect with high probability any corruption. Of course, one can store multiple (possibly different) fingerprints of the database on a number of other nodes. This operation induces a directed graph with links from nodes with some database to nodes holding a fingerprint of that database. If we assume that each fingerprint holder y for the database on computer x has the fingerprint of its own database stored on x , then the graph becomes undirected.

We can view this network scenario, as a game with attackers and defender entities: attackers represent “intelligent” (hence, *rational*) entities of the network that may choose a computer and corrupt its database (and possibly any stored fingerprints), such as viruses, hackers, etc. The defender represents the distributed database administrator who’s goal is to maintain the integrity of the data. The administrator does this by choosing two computers connected by an edge (i.e., two computers that hold each other’s fingerprints) and compares their two databases to the corresponding fingerprints. With high probability this will detect any corruption. These checks are performed on a regular basis, and the administrator is rewarded for any corruptions it detects. Moreover, the more attackers attack a database, the more corrupted the database becomes. So, the defender is not simply awarded a reward for detecting a corrupted database, but rather a reward depending on how corrupted the database was.

To the best of our knowledge, the network game of Mavronicolas *et al.* [14] is the *first* strategic game where the network (*system*) is *explicitly* modeled as a distinct, non-cooperative player (namely, the defender). Unlike previously studied games that evaluated the effect of selfish behavior on system performance using the *Price of Anarchy* [12, 20] (which *implicitly* modeled the system), we pursue this evaluation by defining and using the *Price of Defense* as the worst case ratio of ν over the Individual Profit of the defender.

We are interested in analyzing the Price of Defense for *Nash equilibria* [17, 18], where no single player has an incentive to deviate from its randomized strategy. How does the Price of Defense vary

with Nash equilibria? Are there Nash equilibria that both are computationally tractable and offer good Price of Defense? We present a comprehensive collection of trade-offs between Price of Defense and the computational complexity of Nash equilibria proved in [14, 15, 16].

1.2 Overview

Here, we overview the most important results of [14, 15, 16], providing a multitude of results for general and special classes of Nash equilibria. To describe the results, we need some game-theoretic terminology, which we review here. (For precise definitions, see Section 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 present an elegant algebraic characterization of mixed Nash equilibrium proved in [14] (Theorem 4.1). 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 present two interesting *covering* properties of Nash equilibria showed in [14]. In more detail, in a Nash equilibrium, the support of the edge player to be is an *Edge Cover* of the graph (Proposition 5.1); the support of the vertex players is a *Vertex Cover* of the graph induced by the support of the edge player (Proposition 5.2). So, these covering properties represent necessary graph-theoretic conditions for Nash equilibria.

Inspired by the shown covering properties, the authors of [14], introduced a *Covering profile* as one that satisfies the two necessary covering conditions for Nash equilibria that were proved. It is natural to ask whether a Covering profile is necessarily a Nash equilibrium. We show a simple counterexample to show that a Covering profile is *not* necessarily a Nash equilibrium (Proposition 5.4), first presented in [14]. 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.

Such enrichment is achieved in [14] via the definition of an *Independent Covering profile* (Definition 5.2). Loosely speaking, the following two additional conditions are included in the definition of an Independent Covering profile: (i) The support of the vertex players is an *Independent Set* of the graph. (ii) 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: (i) All vertex players have the same support. (ii) 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.

An Independent Covering profile is proved to be a Nash equilibrium (Proposition 5.6) [14]. The proof verifies that an Independent Covering profile satisfies the characterization of a Nash equilibrium (Theorem 4.1). So, an Independent Covering profile provides sufficient graph-theoretic conditions for Nash equilibria.

Moreover, 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, as shown in [14]. So, an Independent Covering profile will henceforth be called a *Matching Nash* equilibrium.

1.2.2 Pure Nash Equilibria

The graph-theoretic game has no pure Nash equilibrium unless the graph is trivial (Theorem 5.3) [14]. This follows as an immediate consequence of one of the covering properties of a Nash equilibrium (Proposition 5.1).

1.2.3 General Nash Equilibria

A (mixed) Nash equilibrium for the network game is computable in polynomial time (Theorem 6.1) [16]. The proof is by reduction to the case of two players (one attacker and one defender), which is shown to be *constant-sum*. *Constant-sum Two-Players* games are reducible to *Linear Programming* [19], hence solvable in polynomial time [10]. However, the reduction to Linear Programming hides the Price of Defense. This invites the consideration of special classes of Nash equilibria with sufficient structure for the evaluation of the incurred Prices of Defense.

1.2.4 Matching Nash Equilibria

For Matching Nash equilibria we present the following results obtained in [16]:

- We present a new characterization of graphs admitting Matching Nash equilibria (Theorem 7.3) proved in [16]. Such graphs have their *Independence Number* equal to their *Edge Covering Number*. The characterization improves an earlier one from [14]. The characterization benefits from an improved understanding of structural (graph-theoretic) properties of Matching Nash equilibria. In particular, we prove that in a Matching Nash equilibrium, the support of the vertex players is a *Maximum Independent Set* of the graph (Proposition 7.1) and the support of the edge player is a *Minimum Edge Cover* of the graph (Proposition 7.2).
- We present a polynomial time algorithm which translates the characterization to (decide the existence of and) compute a Matching Nash equilibrium (Theorem 7.4). This relies on obtaining a polynomial time algorithm for the (new) graph-theoretic problem of deciding, given a graph G , whether its *Independence Number* $\alpha(G)$ and *Edge Covering Number* $\beta'(G)$ are equal, and yielding, if so, a Maximum Independent Set for the graph (Proposition 3.1). In turn, the graph-theoretic algorithm relies on computing a *Minimum Edge Cover* (via computing a *Maximum Matching*) and a subsequent reduction to 2SAT.
- The Price of Defense for a Matching Nash equilibrium is shown to be $\alpha(G)$ (Proposition 7.5). This relies on its modeling assumption that all vertex players are symmetric and uniform, and on its shown property that the support of the vertex players is a Maximum Independent Set.

1.2.5 Perfect Matching Nash Equilibria

A *Perfect Matching Nash equilibrium* is a Matching Nash equilibrium where, additionally, the support of the edge player is a *Perfect Matching* of the graph.

- We provide a characterization of graphs admitting a Perfect Matching Nash equilibria (Theorem 8.2) presented in [16]. Such graphs have a Perfect Matching and their Independence Number equals $\frac{|V|}{2}$ (V is the vertex set). The characterization benefits from a structural (graph-theoretic) property of Perfect Matching Nash equilibria we prove, namely that the support of the vertex players has size $\frac{|V|}{2}$ (Proposition 8.1).
- The characterization can be translated into a polynomial time algorithm to (decide the existence and) compute a Perfect Matching Nash equilibrium (Theorem 8.3) [16]. This relies on obtaining a polynomial time algorithm for the (new) graph-theoretic problem of deciding, given a graph G with a Perfect Matching, whether its Independence Number equals $\frac{|V|}{2}$, and yielding, if so, a Maximum Independent Set for the graph (Proposition 3.2). In turn, the graph-theoretic algorithm relies on computing a Perfect Matching and a subsequent reduction to 2SAT.
- The Price of Defense for a Perfect Matching Nash equilibrium is $\frac{|V|}{2}$ (Theorem 8.4) [16]. This relies on its modeling assumption that all vertex players are symmetric and uniform, and on its shown property that the support of the vertex players has size $\frac{|V|}{2}$.

The relation between the Prices of Defense for Perfect Matching and Matching Nash equilibria is precisely the relation between $\frac{|V|}{2}$ and $\alpha(G)$ for the graph G . For graphs that have both Matching and Perfect Matching Nash equilibria, Theorem 8.2 implies that $\alpha(G) = \frac{|V|}{2}$ and the two Prices of Defense coincide (as also do the two classes of equilibria). Consider a graph that has a Matching Nash equilibrium but not a Perfect Matching Nash equilibrium. By the characterization of Matching Nash equilibria in [14, Theorem 3], $\alpha(G) \geq \frac{|V|}{2}$ (else, there could not be enough vertices inside an Independent Set to which vertices outside have to be matched). Thus, the Price of Defense of a Perfect Matching Nash equilibrium may not exceed that of a Matching Nash equilibrium.

1.2.6 Generalizations

Finally, we consider two generalization of the basic model with an increased power for the defender:

- In the first variation, introduced in [15], the defender is able to scan a simple *path* of the network instead of a single edge. The problem of deciding the existence of pure Nash equilibria in this model is shown to be \mathcal{NP} -complete (Theorem 9.2) [15]. This result opposes interestingly with the corresponding non-existence result of the our basic model, proved before and indicates some fascinating dimensions of the yet unexplored research area considered here.
- In the second variation, introduced in [7], the defender is able to scan and protect a set of k links of the network. For this model, we discuss a generalized class of Nash equilibria, called *k-matching* Nash equilibria and introduced in [7].

2 Background, Definitions and Preliminaries

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.

2.1 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. A *simple path*, P , is a path of G with no repeated vertices, i.e. $P = \{v_1, \dots, v_i, \dots, v_k\}$, where $1 \leq i \leq k \leq n$, $v_i \in V$, $(v_i, v_{i+1}) \in E(G)$ and each $v_i \in V$ appears at most once in P . Denote $\mathcal{P}(G)$ the set of all possible paths in G . For a vertex set $U \subseteq V$, denote $\text{Neigh}_G(U) = \{u \notin U : (u, v) \in E \text{ for some vertex } v \in U\}$. 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* [3, 11] 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, [1, Theorem 10-2-1, p. 180]). For the class of bipartite graphs, the currently most efficient algorithm to compute a Maximum Matching for a bipartite graph is due to Feder and Motwani [5] and runs in time $O\left(\sqrt{|V|} \cdot |E| \cdot \log_{|V|} \frac{|V|^2}{|E|}\right)$.

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 $U' \subseteq U$, $|U'| \leq |\text{Neigh}_G(U') \cap (V \setminus U)|$. 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 .

2.2 Game Theory

2.2.1 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 ν *vertex* players vp_i , called *attackers*, $1 \leq i \leq \nu$;
 - \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 $\mathcal{S} = \left(\prod_{i \in \mathcal{N}_{vp}} S_i \right) \times S_{ep} = V^\nu \times E$.
- Fix an arbitrary *strategy profile* $\mathbf{s} = \langle s_1, \dots, s_\nu, s_{ep} \rangle \in \mathcal{S}$, 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.

2.2.2 Pure Strategies and Pure Nash Equilibria

The profile \mathbf{s} is a *pure Nash equilibrium* (abbreviated as *pure NE*) [17, 18] 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 $\Pi(G)$.

2.2.3 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* $\mathbf{s} = \langle s_1, \dots, s_\nu, s_{ep} \rangle$, 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 $v \in V$, $s_i(v) = \frac{1}{|\text{Support}_{\mathbf{s}}(i)|}$; for the edge player ep , for each $e \in E$, $s_{ep}(e) = \frac{1}{|\text{Support}_{\mathbf{s}}(ep)|}$. 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.

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

2.2.5 Expected Individual Profit and Conditional Expected Individual Profits

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

Induced by the mixed profile \mathbf{s} is also the *Conditional Expected Individual Profit* $\text{IP}_{\mathbf{s}}(i, v)$ of vertex player $vp_i \in \mathcal{N}_{vp}$ on vertex $v \in V$, which is the conditional expectation according to \mathbf{s} of the Individual Profit of player vp_i had it chosen vertex v . So, $\text{IP}_{\mathbf{s}}(i, v) = 1 - P_{\mathbf{s}}(\text{Hit}(v)) = 1 - \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e)$.

Clearly, for the vertex player $vp_i \in \mathcal{N}_{vp}$, $\text{IP}_{\mathbf{s}}(i) = \sum_{v \in V} s_i(v) \cdot \text{IP}_{\mathbf{s}}(i, v) = \sum_{v \in V} s_i(v) \cdot (1 - \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e))$.

Finally, induced by the mixed profile \mathbf{s} is the *Conditional Expected Individual Profit* $\text{IP}_{\mathbf{s}}(ep, e)$ of the edge player ep on edge $e = (u, v) \in E$, which is the conditional expectation according to \mathbf{s} of the Individual Profit of player ep had it chosen edge e . So, $\text{IP}_{\mathbf{s}}(ep, e) = \text{VP}_{\mathbf{s}}(e) = \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v))$.

Clearly, for the edge player ep , $\text{IP}_{\mathbf{s}}(ep) = \sum_{e \in E} s_{ep}(e) \cdot \text{IP}_{\mathbf{s}}(ep, e) = \sum_{e=(u,v) \in E} s_{ep}(e) \cdot (\sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)))$.

2.2.6 Mixed Nash Equilibria

The mixed profile \mathbf{s} is a *mixed Nash equilibrium* (abbreviated as mixed *NE*) [17, 18] if for each player $i \in \mathcal{N}$, it maximizes $\text{IP}_{\mathbf{s}}(i)$ 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 [17, 18], 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 $\text{IP}_{\mathbf{s}}(i, v)$ are the same and no less than any Conditional Expected Individual Profit $\text{IP}_{\mathbf{s}}(i, v')$ with $s_i(v') = 0$. It follows that for each vertex player vp_i , for any vertex $v \in \text{Support}_{\mathbf{s}}(i)$, $\text{IP}_{\mathbf{s}}(i) = 1 - \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e)$. 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 $\text{IP}_{\mathbf{s}}(ep, e)$ are the same and no less than any Conditional Expected Individual Profit $\text{IP}_{\mathbf{s}}(ep, e')$ with $s_{ep}(e') = 0$. It follows that for the edge player ep , for any edge $(u, v) \in \text{Support}_{\mathbf{s}}(ep)$, $\text{IP}_{\mathbf{s}}(ep) = \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v))$. 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.

A simple but crucial fact about mixed Nash equilibria is proved in [14]:

Lemma 2.1 [14] *Fix a mixed Nash Equilibrium \mathbf{s} . Then, for any pair of vertex players vp_i and vp_k , $IP_{\mathbf{s}}(i) = IP_{\mathbf{s}}(k)$.*

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 $IP_{\mathbf{s}}(ep) > 0$ for a Nash equilibrium \mathbf{s} .

We study algorithmic problems of existence and computation of various classes of Nash equilibria for the considered game.

CLASS NE EXISTENCE INSTANCE: A graph $G(V, E)$. QUESTION: Does $\Pi(G)$ admit a CLASS Nash equilibrium?	FIND CLASS NE INSTANCE: A graph $G(V, E)$. OUTPUT: A CLASS Nash equilibrium of G .
---	--

Variable CLASS takes values GENERAL, MATCHING, PERFECT MATCHING, it determines the classes of general, Matching, Perfect Matching, respectively. We note that for all values of CLASS, membership of a profile in CLASS can be verified in polynomial time. Since a Nash equilibrium can be verified in polynomial time (by Theorem 4.1), it follows that CLASS NE EXISTENCE $\in \mathcal{NP}$.

The Price of Defense is the worst-case ratio, over all Nash equilibria \mathbf{s} , of $\frac{\nu}{IP_{\mathbf{s}}(ep)}$.

3 Some Problems from Graph Theory

For the positive results, we will consider two graph-theoretic problems:

MAXIMUM INDEPENDENT SET EQUAL MINIMUM EDGE COVER

INSTANCE: A graph $G(V, E)$.

OUTPUT: A Maximum Independent Set of G of size $\beta'(G)$ if $\alpha(G) = \beta'(G)$, or No if $\alpha(G) < \beta'(G)$.

MAXIMUM INDEPENDENT SET EQUAL HALF ORDER

INSTANCE: A graph $G(V, E)$.

OUTPUT: A Maximum Independent Set of G of size $\lfloor \frac{|V|}{2} \rfloor$ if $\alpha(G) = \lfloor \frac{|V|}{2} \rfloor$, or No if such does not exist.

For these two new problems, the authors of [16] use reductions to 2SAT (solvable in polynomial time [4]) to prove:

Proposition 3.1 [16] MAXIMUM INDEPENDENT SET EQUAL MINIMUM EDGE COVER $\in \mathcal{P}$

Sketch of Proof. Compute a Minimum Edge Cover EC of G . Recall that EC consists of vertex-disjoint star graphs. Use EC to construct a 2SAT instance ϕ with variable set V as follows:

- (1) For each edge $(u, v) \in E$, add the clause $(\bar{u} \vee \bar{v})$ to ϕ .
- (2) For each single-edge star graph $(u, v) \in EC$, add the clause $(u \vee v)$ to ϕ .
- (3) For each multiple-edge star graph of EC with center vertex u , add the clause $(\bar{u} \vee \bar{u})$ to ϕ .

We prove that G has an Independent Set of size $|EC|$ (hence, $\alpha(G) = \beta'(G)$) if and only if ϕ is satisfiable; when ϕ is satisfiable, the set $\{u \mid \chi(u) = 1\}$ is such a Maximum Independent Set. ■

Similar to Proposition 3.1, in [16] it is proved that,

Proposition 3.2 [16] MAXIMUM INDEPENDENT SET EQUAL HALF ORDER $\in \mathcal{P}$, when restricted to the class of graphs having a Perfect Matching.

4 A Characterization of Nash Equilibria

Theorem 4.1 (Characterization of Nash Equilibria) [14] *A profile \mathbf{s} is a Nash equilibrium if and only if the following two conditions hold:*

- (1) *For any vertex $v \in \text{Support}_{\mathbf{s}}(vp)$, $P_{\mathbf{s}}(\text{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s}}(\text{Hit}(v'))$.*
- (2) *For any edge $e \in \text{Support}_{\mathbf{s}}(ep)$, $\text{VP}_{\mathbf{s}}(e) = \max_{e' \in E} \text{VP}_{\mathbf{s}}(e')$.*

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

For any vertex $v' \in \text{Support}_{\mathbf{s}}(i)$, it holds similarly that $\text{IP}_{\mathbf{s}}(i) = 1 - \sum_{e \in \text{Edges}_{\mathbf{s}}(v')} s_{ep}(e)$. It follows that $\sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e) = \sum_{e \in \text{Edges}_{\mathbf{s}}(v')} s_{ep}(e)$ or $P_{\mathbf{s}}(\text{Hit}(v)) = P_{\mathbf{s}}(\text{Hit}(v'))$. So, consider any vertex $v' \notin \text{Support}_{\mathbf{s}}(i)$. Assume, by way of contradiction, that $P_{\mathbf{s}}(\text{Hit}(v')) < P_{\mathbf{s}}(\text{Hit}(v))$, or equivalently that $\sum_{e \in \text{Edges}_{\mathbf{s}}(v')} s_{ep}(e) < \sum_{e \in \text{Edges}_{\mathbf{s}}(v)} s_{ep}(e)$. Construct from \mathbf{s} the mixed profile \mathbf{s}' by only changing s_i to s'_i so that $v' \in \text{Support}_{\mathbf{s}'}(i)$. Then,

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

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

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

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

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 $v \in \text{Support}_{\mathbf{s}}(i)$,

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

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 $e \in \text{Support}_{\mathbf{s}}(ep)$,

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

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. ■

5 Structure of Nash Equilibria

In this section, we present several graph-theoretic properties of (mixed) Nash equilibria proved in [15]. Necessary and sufficient graph-theoretic conditions are presented in Sections 5.1 and 5.2, respectively.

5.1 Necessary Conditions

In this section, we present necessary, graph-theoretic conditions for Nash equilibria.

Proposition 5.1 [14] *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 $e = (u, v) \in \text{Support}_{\mathbf{s}}(ep)$, $\text{VP}(e) = \sum_{i \in \mathcal{N}_{vp}} (s_i(u) + s_i(v)) = 0$, so that $\text{IP}_{\mathbf{s}}(ep) = \sum_{e \in \text{Support}_{\mathbf{s}}(ep)} s_{ep}(e) \cdot \text{VP}_{\mathbf{s}}(e) = 0$. Since \mathbf{s} is a Nash equilibrium, $\text{IP}_{\mathbf{s}}(ep) > 0$. A contradiction. ■

Using similar arguments, it can be shown that,

Proposition 5.2 [14] *For a Nash Equilibrium \mathbf{s} , $\text{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of $G(\text{Support}_{\mathbf{s}}(ep))$.*

We remark that the necessary conditions in Propositions 5.1 and 5.2 express *covering* properties of Nash equilibria. From Proposition 5.1 it can be proved:

Theorem 5.3 [14] *The graph G admits no pure Nash equilibrium unless it is trivial.*

Inspired by the necessary graph-theoretic conditions in Propositions 5.1 and 5.2 is the definition of a Covering profile that follows.

Definition 5.1 *A Covering profile is a profile \mathbf{s} such that $\text{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G and $\text{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of the graph $G(\text{Support}_{\mathbf{s}}(ep))$.*

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

Proposition 5.4 [14] *A Covering profile is not necessarily a Nash equilibrium.*

5.2 Sufficient Conditions

In this section, we present sufficient, graph-theoretic conditions for Nash equilibria presented in [14]. 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 5.2 *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 $\text{Support}_{\mathbf{s}}(ep)$.

An preliminary property of Independent Covering profiles is the following.

Lemma 5.5 [14] *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 $\text{Support}_{\mathbf{s}}(vp)$.*

The above analysis is used in [14] to show:

Proposition 5.6 [14] *An Independent Covering profile is a Nash equilibrium.*

A necessary condition for Independent Covering profiles is proved in [14].

Proposition 5.7 [14] *For an Independent Covering profile \mathbf{s} , there is a Matching $M \subseteq \text{Support}_{\mathbf{s}}(ep)$ that matches each vertex in $V \setminus \text{Support}_{\mathbf{s}}(vp)$ to some vertex in $\text{Support}_{\mathbf{s}}(vp)$.*

An immediate consequence of Proposition 5.7 follows:

Corollary 5.8 [14] *For an Independent Covering profile \mathbf{s} , $|\text{Support}_{\mathbf{s}}(vp)| \leq |V \setminus \text{Support}_{\mathbf{s}}(vp)|$.*

Propositions 5.6 and 5.7 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*.

6 General Nash Equilibria

Denote as $\hat{\Pi}(G)$ the special case of $\Pi(G)$ with $\nu = 1$; so, $\hat{\Pi}(G)$ is a Two-Players game. Consider a Nash equilibrium $\hat{\mathbf{s}}$ of $\hat{\Pi}(G)$. Construct from $\hat{\mathbf{s}}$ a vp-symmetric profile \mathbf{s} for $\Pi(G)$ where for each vertex player vp_i , for each vertex $v \in V$, $s_i(v) = \hat{s}_{vp}(v)$, where vp denotes the (single) vertex player of $\hat{\Pi}(G)$; for the edge player ep , for each edge $e \in E$, $s_{ep}(e) = \hat{s}_{ep}(e)$. We prove that \mathbf{s} satisfies the characterization of Nash equilibria in Theorem 4.1 (assuming that $\hat{\mathbf{s}}$ does); so, \mathbf{s} is a Nash equilibrium for $\Pi(G)$. Hence, a Nash equilibrium \mathbf{s} for $\Pi(G)$ can be computed from a Nash equilibrium $\hat{\mathbf{s}}$ for $\hat{\Pi}(G)$ in polynomial time.

It can be proved that the two players game $\hat{\Pi}(G)$ is a *constant-sum* (two players) game: for each profile $\hat{\mathbf{s}}$, $\text{IP}_{\hat{\mathbf{s}}}(vp) + \text{IP}_{\hat{\mathbf{s}}}(ep)$ is a constant (independent of $\hat{\mathbf{s}}$). Clearly,

$$\begin{aligned}
 \text{IP}_{\hat{\mathbf{s}}}(vp) + \text{IP}_{\hat{\mathbf{s}}}(ep) &= \sum_{v \in V} \hat{s}_{vp}(v) \cdot \left(1 - \sum_{e \in \text{Edges}_{\hat{\mathbf{s}}}(v)} \hat{s}_{ep}(e) \right) + \sum_{(u,v)=e \in E} \hat{s}_{ep}(e) \cdot (\hat{s}_{vp}(u) + \hat{s}_{vp}(v)) \\
 &= \sum_{v \in V} \hat{s}_{vp}(v) - \sum_{v \in V} \hat{s}_{vp}(v) \left(\sum_{e \in \text{Edges}_{\hat{\mathbf{s}}}(v)} \hat{s}_{ep}(e) \right) + \sum_{(u,v)=e \in E} \hat{s}_{ep}(e) \cdot (\hat{s}_{vp}(u) + \hat{s}_{vp}(v)) \\
 &= 1 - \sum_{(u,v)=e \in E} \hat{s}_{ep}(e) \cdot (\hat{s}_{vp}(u) + \hat{s}_{vp}(v)) + \sum_{(u,v)=e \in E} \hat{s}_{ep}(e) \cdot (\hat{s}_{vp}(u) + \hat{s}_{vp}(v)) = 1
 \end{aligned}$$

Since a Nash equilibrium of a constant-sum, two-players game can be computed in polynomial time via reduction to Linear Programming [19] (which can be solved in polynomial time [10]), it is obtain:

Theorem 6.1 [16] FIND GENERAL NE $\in \mathcal{P}$

7 Matching Nash Equilibria

We first present some graph-theoretic properties of Matching NE proved in [16].

Proposition 7.1 [16] In a Matching NE s , $\text{Support}_s(vp)$ is a Maximum Independent Set of G .

Proposition 7.2 [16] In a Matching NE s , $\text{Support}_s(ep)$ is a Minimum Edge Cover of G .

Using these properties it can be proved that,

Theorem 7.3 [16] The graph G admits a Matching NE if and only if $\alpha(G) = \beta'(G)$.

The constructive parts of the sufficiency proofs of Proposition 3.1 and Theorem 7.3 yield together a polynomial time algorithm **MatchingNE** to compute a Matching NE, if one exists:

Algorithm MatchingNE

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

OUTPUT: The supports in a Matching NE s for G , or NO if such does not exist.

1. Compute a Minimum Edge Cover EC of G .
2. Construct an instance ϕ of 2SAT as follows:
 - For each edge $(u, v) \in E$, add the clause $(\bar{u} \vee \bar{v})$ to ϕ .
 - For each single-edge star graph $(u, v) \in EC$, add the clause $(u \vee v)$ to ϕ .
 - For each multiple-edge star graph of EC with center vertex u , add the clause $(\bar{u} \vee \bar{u})$ to ϕ .
3. Compute a satisfying assignment χ of ϕ , or output NO if such does not exist.
4. Set $IS = \{u \mid \chi(u) = 1\}$.
5. Set $\text{Support}_s(ep) := EC$ and $\text{Support}_s(vp) := IS$.

Theorem 7.4 [16] Algorithm MatchingNE solves FIND MATCHING NE in time $O\left(\sqrt{|V|}|E| \cdot \log_{|V|} \frac{|V|^2}{|E|}\right)$.

Proposition 7.5 [16] In a Matching NE, the Price of Defense is $\alpha(G)$.

8 Perfect Matching Nash Equilibria

A *Perfect Matching NE* is a Matching NE s such that $\text{Support}_s(ep)$ is a Perfect Matching of G . We first present a graph-theoretic property of Perfect Matching NE proved in [16].

Proposition 8.1 [16] For a Perfect Matching NE s , $|\text{Support}_s(vp)| = \frac{|V|}{2}$.

Using this property, a characterization of Perfect Matching NE is proved in [16]:

Theorem 8.2 [16] A graph G admits a Perfect Matching NE if and only if G has a Perfect Matching and $\alpha(G) = \frac{|V|}{2}$.

Algorithm PerfectMatchingNEINPUT: A graph $G(V, E)$.OUTPUT: The supports in a Perfect Matching NE s for G , or NO if such does not exist.

1. Compute a Perfect Matching M of G , or output NO if such does not exist.
2. Construct an instance ϕ of 2SAT as follows:
 - For each edge $(u, v) \in E$, add the clause $(\bar{u} \vee \bar{v})$ to ϕ .
 - For each edge $(u, v) \in M$, add the clause $(u \vee v)$ to ϕ .
3. Compute a satisfying assignment χ of ϕ , or output NO if such does not exist.
4. Set $IS = \{u \mid \chi(u) = 1\}$.
5. Set $\text{Support}_s(ep) := M$ and $\text{Support}_s(vp) := IS$.

The constructive parts of the sufficiency proof of Proposition 3.2 and Theorem 8.2 yield together a polynomial time algorithm **PerfectMatchingNE** to compute a Perfect Matching NE, if one exists.

Theorem 8.3 [16] *Algorithm PerfectMatchingNE solves FIND PERFECT MATCHING NE in time $O\left(\sqrt{|V|}|E| \cdot \log_{|V|} \frac{|V|^2}{|E|}\right)$.*

Observe that a Perfect Matching NE is a Matching NE for which, by Theorem 8.2, $\alpha(G) = \frac{|V|}{2}$. Hence, Proposition 7.5 implies:

Theorem 8.4 [16] *In a Perfect Matching Nash equilibrium, the Price of Defense is $\frac{|V|}{2}$.*

9 Generalizations

The Path Model

Here, we consider a generalization of the basic model, where the defender has increased power: it is able to clean, each time a simple *path* of the network instead of a single edge. We denote the advanced defender as the *path player* and denote it as *pp*. The resulting model, called as *Path model*, is denoted as $\Pi_P(G)$ is identical to $\Pi(G)$, with the only difference that now $|\mathcal{S}_{pp}| = |\mathcal{P}(G)|$. Thus, the path player have an exponential number of pure strategies. We provide the following characterization of pure Nash Equilibria in the Path model.

Theorem 9.1 [15] *For any graph G , $\Pi_P(G)$ has a pure NE if and only if G contains a Hamiltonian path.*

This characterization immediately implies:

Corollary 9.2 [15] *The problem of deciding whether G admits a pure NE for a $\Pi_P(G)$ is \mathcal{NP} -complete.*

The Tuple Model

Another generalization of the basic model $\Pi(G)$, called the *Tuple model*, is introduced and studied in [7]. Here, the defender is able to scan and protect a set of k links of the network. There, the existence problem for pure Nash equilibria is shown to be solvable in polynomial time. Also, a generalized class of Nash equilibria, called *k-matching* Nash equilibria, is introduced. For this class of Nash equilibria, the authors in [7] provide a polynomial-time transformation of any Matching Nash equilibrium of instance of the Edge model into a *k-Matching* Nash equilibrium on a corresponding instance of the Tuple model and vice versa.

References

- [1] A. S. Asratian, T. M. J. Denley and R. Häggkvist, *Bipartite Graphs and Their Applications*, Cambridge Tracts in Mathematics 131, Cambridge University Press, 1998.
- [2] E. R. Cheswick and S. M. Bellovin, *Firewalls and Internet Security*, Addison-Wesley, 1994.
- [3] E. Egerváry, “On Combinatorial Properties of Matrices” (in Hungarian with German summary), *Matematikai és Fizikai Lapok*, Vol. 38, pp. 16-28, 1931.
- [4] S. Even, A. Itai and A. Shamir, “On the Complexity of Timetable and Multicommodity Flow Problems”, *SIAM Journal on Computing*, Vol. 5, No. 4, pp. 691–703, 1976.
- [5] T. Feder and R. Motwani. “Clique Compressions, Graph Partitions and Speeding-up Algorithms”, *Journal of Computer and System Sciences*, Vol. 51, No. 2, pp. 261-272, October 1995.
- [6] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman & Co., 1979.
- [7] M. Gelastou, M. Mavronicolas, V. Papadopoulou, A. Philippou and P. Spirakis, “The Power of the Defender”, *Proceedings of the 2nd International Workshop on Incentive-Based Computing*, in conjunction with the *26th IEEE International Conference on Distributed Computing*, July 2006, to appear.
- [8] A. V. Goldberg and A. V. Karzanov, “Maximum Skew-Symmetric Flows”, *Proceedings of the 3rd Annual European Symposium on Algorithms*, pp. 155–170, Vol. 979, LNCS, 1995.
- [9] P. Hall, “On Representation of Subsets”, *Journal of the London Mathematical Society*, Second Series, Vol. 10, pp. 26-30, 1935.
- [10] L. G. Khachiyan, “A Polynomial Algorithm in Linear Programming”, *Soviet Mathematics Doklady*, Vol. 20, No. 1, pp. 191–194, 1979.
- [11] D. König, “Graphen und Matrizen”, *Matematikai és Fizikai Lapok*, Vol. 38, pp. 116–119, 1931.
- [12] E. Koutsoupias and C. H. Papadimitriou, “Worst-case Equilibria”, *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Vol. 1563, LNCS, 1999.
- [13] T. Markham and C. Payne, “Security at the Network Edge: A Distributed Firewall Architecture”, *Proceedings of the 2nd DARPA Information Survivability Conference and Exposition*, Vol. 1, pp. 279-286, 2001.
- [14] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis. “A Network Game with Attacker and Protector Entities”, *Proceedings of the 16th Annual International Symposium on Algorithms and Computation*, pp. 288–297, Vol. 3827, LNCS, 2005.
Full version: “A Network Game with Attackers and a Defender”, Invited submission to *Algorithmica*, Special Issue with selected papers from *ISAAC 2005*. Electronic version available at <http://www2.cs.ucy.ac.cy/~mavronic/publications.htm>
- [15] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, “A Graph-Theoretic Network Security Game”, *Proceedings of the First International Workshop on Internet and Network Economics*, pp. 969–978, Vol. 3828, LNCS, 2005.
- [16] M. Mavronicolas, L. Michael, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, “The Price of Defense”, Submitted for publication to the *31st International Symposium on Mathematical Foundations of Computer Science 2006*.
- [17] J. F. Nash, “Equilibrium Points in N-Person Games”, *Proceedings of National Academy of Sciences of the United States of America*, pp. 48–49, Vol. 36, 1950.
- [18] J. F. Nash, “Non-Cooperative Games”, *Annals of Mathematics*, Vol. 54, No. 2, pp. 286–295, 1951.
- [19] J. von Neumann, “Zur Theorie der Gesellschaftsspiele”, *Mathematische Annalen*, Vol. 100, pp. 295–320, 1928.
- [20] C. H. Papadimitriou, “Algorithms, Games and the Internet”, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 749–753, 2001.