The Price of Defense and Fractional Matchings*

Marios Mavronicolas[1], Vicky Papadopoulou[1], Giuseppe Persiano[2], Anna Philippou[1], and Paul Spirakis[3]

[1] Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus. {mavronic,viki,annap}@ucy.ac.cy

[2] Dipartimento di Informatica ed Applicazioni "Renato M. Capocelli", Università di Salerno, Italy. giuper@dia.unisa.it

[3] Research Academic Computer Technology Institute, Greece & Department of Computer Engineering and Informatics, University of Patras, Greece. spirakis@@ucy.ac.cy

Abstract. Consider a *network* vulnerable to security *attacks* and equipped with *defense* mechanisms. How much is the loss in the provided security guarantees due to the selfish nature of attacks and defenses? The *Price of Defense* was recently introduced in [7] as a *worst-case* measure, over all associated *Nash equilibria*, of this loss. In the particular *strategic game* considered in [7], there are two classes of confronting randomized players on a graph $G(V, E)$: ν *attackers*, each choosing vertices and wishing to minimize the probability of being caught, and a single *defender*, who chooses edges and gains the expected number of attackers it catches.

In this work, we continue the study of the Price of Defense. We obtain the following results:

- The Price of Defense is at least $\frac{|V|}{2}$; this implies that the *Perfect Matching Nash equilibria* considered in [7] are *optimal* with respect to the Price of Defense, so that the lower bound is *tight*.
- We define *Defense-Optimal graphs* as those admitting a Nash equilibrium that attains the (tight) lower bound of $\frac{|V|}{2}$. We obtain:
  - A graph is Defense-Optimal if and only if it has a *Fractional Perfect Matching*. Since graphs with a Fractional Perfect Matching are recognizable in polynomial time, the same holds for Defense-Optimal graphs.
  - We identify a very simple graph that is Defense-Optimal but has no Perfect Matching Nash equilibrium.
- Inspired by the established connection between Nash equilibria and Fractional Perfect Matchings, we transfer a known bivaluedness result about Fractional Matchings to a certain class of Nash equilibria. So, the connection to *Fractional Graph Theory* may be the key to revealing the combinatorial structure of Nash equilibria for our network security game.

* This work was partially supported by the European Union under IST FET Integrated Project 015964 AEOLUS.

1 Introduction

Motivation, Framework and Summary. Consider a complex distributed system such as the Internet with security *attacks* and corresponding *defense* mechanisms. Assume that both attacks and defenses exhibit a selfish behavior, aiming at maximizing the security harm and the security protection, respectively. How much is the loss in security due to this selfish behavior? In a recent work, Mavronicolas *et al.* [7] introduced the *Price of Defense* as a *worst-case* measure for this loss.

More specifically, Mavronicolas *et al.* [7] focused on the concrete case where the distributed system is a network modeled as a graph $G(V, E)$; nodes are vulnerable to infection by ν threats, called *attackers*. Available to the network is a security software (or *firewall* [3]), called the *defender*, cleaning a limited part of the network. This model has been motivated by *Network Edge Security* [6], a new distributed firewall architecture. (For details on motivation, see [7, Section 1.1].) The model was introduced in [8] and further studied in [4, 7, 9].

Each *attacker* (called *vertex player*) targets a node of the network chosen via its own probability distribution on nodes; the *defender* (called *edge player*) chooses a single *edge* via its own probability distribution on edges. A node chosen by an attacker is harmed unless it is incident to the edge protected by the defender. The *Individual Profit* of an attacker is the probability that it escapes; the *Individual Profit* of the defender is the expected number of caught attackers. In a *Nash equilibrium* [12, 13], no single player can unilaterally deviate from its randomized strategy in order to increase its Individual Profit. The *Price of Defense* is the worst case ratio, over all Nash equilibria, of the ratio of ν over the Individual Profit of the defender. For a particular Nash equilibrium, this ratio is called its *Defense Ratio*. The Price of Defense can be cast as the particular case of Price of Anarchy [5] induced by taking Social Cost to be the Individual Profit of the defender.

In this work, we continue the study of the Price of Defense. More specifically, we provide a *tight* lower bound on the Price of Defense, and we determine a characterization of graphs admitting a Nash equilibrium that attains this lower bound. The characterization establishes a connection to *Fractional Graph Theory* [14]; we further investigate this connection to shed some light into the combinatorial structure of Nash equilibria for our graph-theoretic network security game.

Contribution. We obtain the following results:

- We prove that the Price of Defense is at least $\frac{|V|}{2}$ (Theorem 5). This implies that the *Perfect Matching Nash equilibria*, a special class of Nash equilibria considered in [7] and known to have a Defense Ratio equal to $\frac{|V|}{2}$, are *optimal* with respect to the Price of Defense. It also naturally raises the question whether Perfect Matching Nash equilibria are the only such optimal Nash equilibria; more generally, which are the graphs that admit optimal Nash equilibria with respect to the Price of Defense?

– To address the last question, we introduce the class of Defense-Optimal graphs: a graph is *Defense-Optimal* if it admits a Nash equilibrium whose Defense Ratio is $\frac{|V|}{2}$. Clearly, the class of graphs admitting a Perfect Matching Nash equilibrium is contained in this class; an efficient characterization for that class is shown in [7, Theorem 6.2] (repeated as Theorem 2 in this paper). (This class is a strict subclass of the class of graphs with a Perfect Matching.) We have obtained the following results:
  - A graph is Defense-Optimal if and only if it has a *Fractional Perfect Matching* (Theorem 8). Our proof is constructive: Given a Fractional Perfect Matching, we construct a Defense-Optimal Nash equilibrium (Theorem 6), and vice-versa (Theorem 7). These dual constructions exhibit an interesting, perhaps unexpected connection between Nash equilibria for our graph-theoretic game and Fractional (Perfect) Matchings [14, Chapter 2] in graphs.

– We observe that the class of graphs admitting Perfect Matching Nash equilibria is strictly contained into the class of Defense-Optimal graphs. Towards this end, we identify the *simplest* Defense-Optimal graph that does not admit a Perfect Matching Nash equilibrium (Theorem 9).

– We further investigate the established equivalence between (Defense-Optimal) Nash equilibria and Fractional Perfect Matchings. Our starting point is a result from *Fractional Graph Theory* [14] stating that for any graph, there is a *Fractional Maximum Matching* with only two distinct (non-zero) values on edges [14, Theorem 2.1.5]. We establish a corresponding fact for *Defender Uniform* Nash equilibria. (These are Nash equilibria where the defender uses a uniform probability distribution on its support.) Specifically, we prove that from a Defender Uniform Nash equilibrium, one can obtain in polynomial time another (Defender Uniform) Nash equilibrium where the expected number of vertex players choosing each vertex may take only two distinct (non-zero) values (Theorem 11).

  We believe that a further investigation of the connection between Nash equilibria for our graph-theoretic game and Fractional Matchings will provide further key insights into the (yet not so well understood) combinatorial structure of these Nash equilibria.

Related Work and Significance. Our work continues the study of the game-theoretic virus model with attackers and a defender introduced by Mavronicolas *et al.* [8] and further studied in [4, 7, 9]. In particular, our work continues the study of the Price of Defense introduced in [7].

A different game-theoretic model of virus attack and propagation has been introduced by Aspnes *et al.* [1] and recently studied by Moscibroda *et al.* [11]. Moscibroda *et al.* [11] introduced the *Price of Malice* to quantify the impact of malicious players on the Price of Anarchy (without malicious players) for the game of Aspnes *et al.* [1]. Note that we do not consider malicious players for our game; we assume that all players are strategic. So, there is no apparent relation between Price of Malice and Price of Defense.

Our work is part of a currently major effort to introduce game-theoretic models in Computer Science in order to obtain insights into the reality of contemporary distributed systems such as the Internet. Work on game-theoretic analysis of complex distributed systems is now featured in major conferences of Distributed Computing.

2 Background, Definitions and Preliminaries

Graph Theory. Throughout, we consider an undirected graph $G = \langle V, E\rangle$ with no isolated vertices. We sometimes treat an edge as the set of its two vertices. For a vertex $v \in V$, denote as $\mathsf{Neigh}_G(v)$ the set of neighboring vertices of v in G; denote $\mathsf{Edges}_G(v)$ the set of edges incident to v. For a vertex set $U \in V$, $\mathsf{Neigh}_G(U) = \{v \in V \backslash U : u \in U \text{ and } (v, u) \in E\}$. For a vertex $v \in V$, denote $d_G(v)$ the *degree* of vertex v in G. For an edge set $F \subseteq E$, denote $G(F)$ the subgraph of G induced by F. For any integer $n \geq 1$, denote as K_n the *clique* graph of size n.

A vertex set $IS \subseteq V$ is an *Independent Set* 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* is a vertex set $VC \subseteq V$ such that for each edge $(u, v) \in E$ either $u \in VC$ or $v \in VC$. An *Edge Cover* 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* is a set $M \subseteq E$ of non-incident edges. A *Maximum Matching* is one that has maximum size. A *Perfect Matching* is a Matching that is also an Edge Cover.

A *Fractional Matching* is a function $f : E \rightarrow [0, 1]$ such that for each vertex $v \in V$, $\sum_{e \in \mathsf{Edges}(v)} f(e) \leq 1$. (If $f(e) \in \{0, 1\}$ for each edge $e \in E$, then f is just a Matching, or precisely, the indicator function of a Matching.) The *Fractional Matching Number* $\alpha'_F(G)$ of a graph G is the supremum of $\sum_{e \in E} f(e)$ over all Fractional Matchings f. A *Fractional Maximum Matching* is one that attains the Maximum Matching Number. It is a basic fact that $\alpha'_F(G) \leq \frac{|V|}{2}$ (see, for example, [14, Lemma 2.1.2]). A *Fractional Perfect Matching* is a Fractional Matching f with $\sum_{e \in \mathsf{Edges}(v)} f(e) = 1$ for all vertices $v \in V$. Hence, for a Fractional Perfect Matching f, $\sum_{e \in E} f(e)$ achieves the upper bound on $\alpha'_F(G)$, so that $\sum_{e \in E} f(e) = \frac{|V|}{2}$.

Note that the Fractional Matching Number of a graph can be computed in polynomial time by formulating (and solving) the Fractional Matching Number problem as a polynomial size (in fact, $|V| \cdot |E|$ size) Linear Program. (See, also, [2] for an efficient combinatorial algorithm.) Since a graph $G = (V, E)$ has a Fractional Perfect Matching if and only if its Fractional Matching Number is equal to $\frac{|V|}{2}$, it follows that the class of graphs with a Fractional Perfect Matching is recognizable in polynomial time.

Game Theory. We consider a *strategic game* $\Pi(G) = \langle \mathcal{N}, \{S_i\}_{i \in \mathcal{N}}, \{\mathsf{IP}\}_{i \in \mathcal{N}} \rangle$:

– The set of *players* is $\mathcal{N} = \mathcal{N}_{vp} \cup \mathcal{N}_{ep}$, where $\mathcal{N}_{vp}$ has ν *vertex* players vp_i, called *attackers*, $1 \leq i \leq \nu$ and $\mathcal{N}_{ep}$ has *edge* player ep, called *defender*.
– The *strategy set* S_i of vertex player vp_i is V, and the *strategy set* S_{ep} of the edge player ep is E. So, the *strategy set* $\mathcal{S}$ of the game is $\mathcal{S} = \left(\times_{i \in \mathcal{N}_{vp}} S_i \right) \times S_{ep} = V^\nu \times E$.
– Fix any *profile* $\mathbf{s} = \langle s_1, \ldots, s_\nu, s_{ep} \rangle \in \mathcal{S}$, also called a *pure profile*.
  - The *Individual Profit* of vertex player vp_i is a function $\mathsf{IP}_\mathbf{s}(i) : \mathcal{S} \rightarrow \{0,1\}$ such that $\mathsf{IP}_\mathbf{s}(i) = \begin{cases} 0, & s_i \in s_{ep} \\ 1, & s_i \notin s_{ep} \end{cases}$; 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 $\mathsf{IP}_\mathbf{s}(ep) : \mathcal{S} \rightarrow \mathbb{N}$ such that $\mathsf{IP}_\mathbf{s}(ep) = |\{i : s_i \in s_{ep}\}|$; intuitively, the edge player ep receives the number of vertex players it catches.

A *mixed strategy* for player $i \in \mathcal{N}$ is a probability distribution over S_i. A (*mixed*) *profile* $\mathbf{s} = \langle s_1, \ldots, s_\nu, s_{ep} \rangle$ is a collection of mixed strategies, one for each player; $s_i(v)$ is the probability that 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 player $i \in \mathcal{N}$ in the mixed profile $\mathbf{s}$, denoted $\mathsf{Support}_\mathbf{s}(i)$, is the set of pure strategies in its strategy set to which i assigns a strictly positive probability in $\mathbf{s}$. Denote $\mathsf{Support}_\mathbf{s}(vp) = \bigcup_{i \in \mathcal{N}_{vp}} \mathsf{Support}_\mathbf{s}(i)$. Set $\mathsf{Edges}_\mathbf{s}(v) = \{(u,v) \in E : (u,v) \in \mathsf{Support}_\mathbf{s}(ep)\}$. So, $\mathsf{Edges}_\mathbf{s}(v)$ contains all edges incident to v that are included in the support of the edge player. For an edge $e = (u,v) \in E$, set $\mathsf{Vertices}_\mathbf{s}(e) = \{w \in \{u,v\} : w \in \mathsf{Support}_\mathbf{s}(vp)\}$.

A profile $\mathbf{s}$ is *Fully Mixed* [10] if for each vertex player vp_i, $\mathsf{Support}_\mathbf{s}(i) = V$ and $\mathsf{Support}_\mathbf{s}(ep) = E$; so, the support of each player is its strategy set. A profile $\mathbf{s}$ is *Uniform* if each player uses a uniform probability distribution on its support; that is, for every vertex player $vp_i \in \mathcal{N}_{vp}$ and $v \in \mathsf{Support}_\mathbf{s}(i)$, $s_i(v) = \frac{1}{|\mathsf{Support}_\mathbf{s}(i)|}$, and, for the edge player ep, for each $e \in \mathsf{Support}_\mathbf{s}(ep)$, $s_{ep}(e) = \frac{1}{|\mathsf{Support}_\mathbf{s}(ep)|}$. A profile $\mathbf{s}$ is *Attacker Symmetric* [7] if for all vertex players vp_i, $vp_k \in \mathcal{N}_{vp}$, $s_i(v) = s_k(v)$, for each $v \in V$. An *Attacker Symmetric and Uniform* profile is an Attacker Symmetric profile where each attacker uses a uniform probability distribution on the common support; an Attacker Symmetric, Uniform and Fully Mixed profile is an Attacker Symmetric and Uniform profile where the common support is V. A profile is *Defender Uniform* [7] if the edge player uses a uniform probability distribution on its support.

For a vertex $v \in V$, the probability that the edge player ep chooses an edge that contains the vertex v is denoted by $P_\mathbf{s}(\mathsf{Hit}(v))$. So, $P_\mathbf{s}(\mathsf{Hit}(v)) = \sum_{e \in \mathsf{Edges}_\mathbf{s}(v)} s_{ep}(e)$. For a vertex $v \in V$, denote as $\mathsf{VP}_\mathbf{s}(v)$ the expected number of vertex players choosing vertex v according to $\mathbf{s}$; so, $\mathsf{VP}_\mathbf{s}(v) = \sum_{i \in \mathcal{N}_{vp}} s_i(v)$. Further, in an Attacker Symmetric and Uniform profile $\mathbf{s}$, for a vertex $v \in$

$\mathsf{Support}_{\mathbf{s}}(vp)$, $\mathsf{VP}_{\mathbf{s}}(v) = \sum_{i \in \mathcal{N}_{vp}} s_i(v) = \frac{\nu}{|\mathsf{Support}_s(vp)|}$. For each edge $e = (u, v) \in E$, $\mathsf{VP}_{\mathbf{s}}(e)$ is the expected number of vertex players choosing either the vertex u or the vertex v; so, $\mathsf{VP}_{\mathbf{s}}(e) = \mathsf{VP}_{\mathbf{s}}(u) + \mathsf{VP}_{\mathbf{s}}(v)$. We provide a preliminary observation which will be useful later.

Lemma 1. *In a profile* $\mathbf{s}$, $\sum_{v \in V} P_{\mathbf{s}}(\mathsf{Hit}(v)) = 2$.

A mixed profile $\mathbf{s}$ induces an *Expected Individual Profit* $\mathsf{IP}_{\mathbf{s}}(i)$ for each player $i \in \mathcal{N}$, which is the expectation according to $\mathbf{s}$ of the Individual Profit of player i. One may easily show that for the edge player ep, $\mathsf{IP}_{\mathbf{s}}(ep) = \sum_{i \in \mathcal{N}_{vp}} \left(\sum_{v \in V} s_i(v) \cdot (P_{\mathbf{s}}(\mathsf{Hit}(v)) \right)$; alternatively, $\mathsf{IP}_{\mathbf{s}}(ep) = \sum_{v \in V} \mathsf{VP}_{\mathbf{s}}(v) \cdot P_{\mathbf{s}}(\mathsf{Hit}(v))$.

The mixed profile $\mathbf{s}$ is a *(mixed) Nash equilibrium* [12, 13] if, for each player $i \in \mathcal{N}$, it maximizes $\mathsf{IP}_{\mathbf{s}}(i)$ over all mixed profiles that differ from $\mathbf{s}$ only with respect to the mixed strategy of player i. By Nash's result [12, 13], there is at least one Nash equilibrium. We use a characterization of them from [8]:

Theorem 1 ([8]). *A profile* $\mathbf{s}$ *is a Nash equilibrium if and only if* (1) *for each vertex* $v \in \mathsf{Support}_{\mathbf{s}}(vp)$, $P_{\mathbf{s}}(\mathsf{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v'))$, *and* (2) *for each edge* $e \in \mathsf{Support}_{\mathbf{s}}(ep)$, $\mathsf{VP}_{\mathbf{s}}(e) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')$.

Call $\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v'))$ the *Minimum Hitting Probability* associated with $\mathbf{s}$.

We continue to introduce the class of Perfect Matching Nash equilibria from [7]. A *Covering profile* is a profile $\mathbf{s}$ such that (1) $\mathsf{Support}_{\mathbf{s}}(ep)$ is an Edge Cover of G and (2) $\mathsf{Support}_{\mathbf{s}}(vp)$ is a Vertex Cover of the graph $G(\mathsf{Support}_{\mathbf{s}}(ep))$. It is shown in [8] that a Nash equilibrium $\mathbf{s}$ is a Covering profile, but not vice versa. An *Independent Covering profile* [8] is an Attacker Symmetric and Uniform Covering profile $\mathbf{s}$ such that (1) $\mathsf{Support}_{\mathbf{s}}(vp)$ is an Independent Set of G and (2) each vertex in $\mathsf{Support}_{\mathbf{s}}(vp)$ is incident to exactly one edge in $\mathsf{Support}_{\mathbf{s}}(ep)$. In the same work, it was proved that an Independent Covering profile is a Nash equilibrium, called a *Matching Nash equilibrium* [8]. A *Perfect Matching Nash equilibrium* is a Matching Nash equilibrium such that the support of the edge player is a Perfect Matching of G. Call a graph *Perfect-Matching* if it admits a Perfect Matching Nash equilibrium. (This should not be confused with the strictly larger class of graphs with a Perfect Matching.) A characterization of Perfect-Matching graphs is provided in [7]:

Theorem 2 ([7]). *A graph G is Perfect-Matching if and only if G has a Perfect Matching and* $\alpha(G) = \frac{|V|}{2}$.

A *Defender Uniform Nash equilibrium* is a Defender Uniform profile that is a Nash equilibrium. Call a graph *Defender-Uniform* if it admits a Defender Uniform Nash equilibrium. We use a characterization from [7]:

Theorem 3 ([7]). *A graph G is Defender-Uniform if and only if there are non-empty sets* $V' \subseteq V$, *partitioned as* $V' = V'_i \cup V'_r$, *and* $E' \subseteq E$, *and an integer* $r \geq 1$ *such that:*

(1/a) *For each* $v \in V'$, $d_{G(E')}(v) = r$.

(1/b) *For each* $v \in V \backslash V'$, $d_{G(E')}(v) \geq r$.
(2/a) *For each* $v \in V_i'$, *for each* $u \in \mathsf{Neigh}_G(v)$, *it holds that* $u \notin V'$.
(2/b) *The graph* $\langle V_r', \mathsf{Edges}_G(V_r') \cap E' \rangle$ *is an* r*-regular graph.*
(2/c) *The graph* $\langle V_i' \cup (V \backslash V'), \mathsf{Edges}_G(V_i' \cup (V \backslash V')) \cap E' \rangle$ *is a* $(V_i', V \backslash V')$*-bipartite graph.*

An inspection of the proof of Theorem 3 in [7] implies a partial but more specific version of Theorem 3 that suffices for our purposes.

Theorem 4. *Consider a Defender Uniform Nash equilibrium* $\mathbf{s}$*. Then, for the choices*

- $V' = \mathsf{Support}_{\mathbf{s}}(vp)$, *with* (i) $V_i' := \{v \in V' \mid \mathsf{VP}_{\mathbf{s}}(v) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')\}$ *and* (ii) $V_r' := V' \backslash V_i'$;
- $E' = \mathsf{Support}_{\mathbf{s}}(ep)$;
- $r = d_{G(\mathsf{Support}_{\mathbf{s}}(ep))}(v)$ *for any vertex* $v \in \mathsf{Support}_{\mathbf{s}}(vp)$,

the graph $\langle V_r', \mathsf{Edges}_G(V_r') \cap E' \rangle$ *is an* r*-regular graph.*

We prove a useful property of Defender Uniform Nash equilibria:

Lemma 2. *Consider a Defender Uniform Nash equilibrium* $\mathbf{s}$ *and the induced subgraph* $\langle V_r', \mathsf{Edges}_G \ (V_r') \cap E' \rangle$*, where* $V_r' = V \backslash \{v \in V' \mid \mathsf{VP}_{\mathbf{s}}(v) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')\}$ *and* $E' = \mathsf{Support}_{\mathbf{s}}(ep)$*. Then, over all vertices* v *in each connected component of the subgraph, the variable* $\mathsf{VP}_{\mathbf{s}}(v)$ *takes on at most two distinct (non-zero) values, which occur an equal number of times.*

For a Nash equilibrium $\mathbf{s}$, the ratio $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)}$ is called the *Defense Ratio* of $\mathbf{s}$. The *Price of Defense* [7], denoted PoD_G, is the *worst-case* Defense Ratio of $\mathbf{s}$, over all Nash equilibria $\mathbf{s}$. It is known that the Defense Ratio of every Perfect Matching Nash equilibrium is $\frac{|V|}{2}$ [7, Theorem 6.4]. Hence, restricted to Perfect Matching Nash equilibria, the Price of Defense is $\frac{|V|}{2}$.

3 A Lower Bound on the Price of Defense

We first use Theorem 1 to evaluate the Defense Ratio of a Nash equilibrium:

Proposition 1. *For a Nash equilibrium* $\mathbf{s}$, $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)} = \frac{1}{\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v'))}$.

Using Lemma 1 we show:

Proposition 2. *Assume a Nash equilibrium* $\mathbf{s}$*. Then,* $\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v')) \leq \frac{2}{|V|}$.

Theorem 5. *The Price of Defense is at least* $\frac{|V|}{2}$.

Proof. Consider any Nash equilibrium $\mathbf{s}$. By Proposition 1, we get that $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)} = \frac{1}{\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v'))}$. By Proposition 2, this implies that $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)} \geq \frac{|V|}{2}$. Since $\mathsf{PoD}_G \geq \frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)}$, the claim follows. □

A Nash equilibrium $\mathbf{s}$ is *Defense-Optimal* if its Defense Ratio $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)}$ equals to $\frac{|V|}{2}$. A graph G is *Defense-Optimal* if it admits a Defense-Optimal Nash equilibrium. Proposition 1 immediately implies:

Corollary 1. *Consider a Defense-Optimal Nash equilibrium* $\mathbf{s}$. *Then,* $\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v')) = \frac{2}{|V|}$.

Together with Proposition 2, Corollary 1 implies that Defense-Optimal Nash equilibria maximize the Minimum Hitting Probability.

4 Defense-Optimal Graphs

We provide a characterization of Defense-Optimal graphs. We first prove:

Theorem 6. *Assume that G has a Fractional Perfect Matching. Then, G is Defense-Optimal.*

Sketch of Proof. Consider a Fractional Perfect Matching $f : E \to [0, 1]$. Define an Attacker Symmetric, Uniform and Fully Mixed profile $\mathbf{s}$ as follows:

– For each edge $e \in E$, $s_{ep}(e) = \frac{2}{|V|} \cdot f(e)$.

It can be easily shown that s_{ep} is a probability distribution for the edge player. We first prove that $\mathbf{s}$ is a Nash equilibrium. It suffices to prove Conditions (1) and (2) in the characterization of Nash equilibria (Theorem 1).

– For Condition (1), consider any vertex $v \in V$. Clearly,

$$\begin{aligned} & P_{\mathbf{s}}(\mathsf{Hit}(v)) \\ = & \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} s_{ep}(e) \\ = & \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} \frac{2}{|V|} \cdot f(e) \\ = & \frac{2}{|V|} \cdot \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} f(e) \\ = & \frac{2}{|V|} \quad \text{(since } f \text{ is a Fractional Perfect Matching).} \end{aligned}$$

  Thus, in particular, for any vertex $v \in \mathsf{Support}_{\mathbf{s}}(vp)$, $P_{\mathbf{s}}(\mathsf{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v'))$ and Condition (1) holds.

– For Condition (2), consider any edge $e = (u, v) \in E$. Clearly,

$$\begin{aligned} & \mathsf{VP}_{\mathbf{s}}(e) \\ = & \mathsf{VP}_{\mathbf{s}}(u) + \mathsf{VP}_{\mathbf{s}}(v) \\ = & \frac{\nu}{|V|} + \frac{\nu}{|V|} \quad \text{(since } \mathbf{s} \text{ is Attacker Symmetric, Uniform and Fully Mixed)} \\ = & \frac{2\nu}{|V|}. \end{aligned}$$

  Thus, in particular, for any edge $e \in \mathsf{Support}_{\mathbf{s}}(ep)$, $\mathsf{VP}_{\mathbf{s}}(e) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')$ and Condition (2) holds.

It follows that $\mathbf{s}$ is a Nash equilibrium. We finally prove that $\mathbf{s}$ is Defense-Optimal. Clearly, for any edge $e \in \mathsf{Support}_{\mathbf{s}}(ep)$, $\mathsf{IP}_{\mathbf{s}}(ep) = \mathsf{VP}_{\mathbf{s}}(e)$, so that $\frac{\nu}{\mathsf{IP}_{\mathbf{s}}(ep)} = \frac{\nu}{\mathsf{VP}_{\mathbf{s}}(e)} = \frac{|V|}{2}$, so that $\mathbf{s}$ is Defense-Optimal. The claim follows. □

Theorem 7. *Assume that G is Defense-Optimal. Then, G has a Fractional Perfect Matching.*

Sketch of Proof. Consider a Defense-Optimal Nash equilibrium $\mathbf{s}$ for G. By Proposition 1, $\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v')) = \frac{2}{|V|}$. By Lemma 1, $\sum_{v \in V} P_{\mathbf{s}}(\mathsf{Hit}(v)) = 2$. It follows that for each vertex $v \in V$, $P_{\mathbf{s}}(\mathsf{Hit}(v)) = \frac{2}{|V|}$. Define a function $f : E \rightarrow [0,1]$ as follows:

- For each edge $e = (u,v) \in E$, $f(e) = \frac{s_{ep}(e)}{P_{\mathbf{s}}(\mathsf{Hit}(v))}$.

Clearly, for each edge $e = (u,v) \in E$, $P_{\mathbf{s}}(\mathsf{Hit}(v)) \geq s_{ep}(e)$, so that $f(e) \leq 1$. Moreover, for each vertex $v \in V$,

$$\begin{aligned}\sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} f(e) &= \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} \frac{s_{ep}(e)}{P_{\mathbf{s}}(\mathsf{Hit}(v))} \\ &= \frac{1}{P_{\mathbf{s}}(\mathsf{Hit}(v))} \cdot \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} s_{ep}(e) \\ &= 1.\end{aligned}$$

Hence, f is a Fractional Perfect Matching, as needed. □

Theorems 6 and 7 together imply:

Theorem 8 (Characterization of Defense-Optimal Graphs). *A graph is Defense-Optimal if and only if it has a Fractional Perfect Matching.*

Since the class of graphs with a Fractional Perfect Matching is recognizable in polynomial time, Theorem 8 immediately implies:

Corollary 2. *Defense-Optimal graphs are recognizable in polynomial time.*

By Theorem 2, the class of Perfect-Matching graphs is (strictly) contained in the class of graphs with a Perfect Matching. Since a Perfect Matching is a special case of a Fractional Perfect Matching, it follows that the class of Perfect-Matching graphs is (strictly) contained in the class of graphs with a Fractional Perfect Matching. Hence, Theorem 8 implies that the class of Perfect-Matching graphs is (strictly) contained in the class of Defense-Optimal graphs. We provide a particular example to demonstrate the strict inclusion.

Theorem 9. *K_3 is a Defense-Optimal graph but not a Perfect-Matching graph.*

5 Bivalued Nash Equilibria

Our starting point is a bivaluedness result about Fractional Maximum Matchings, which appears in [14, Theorem 2.1.5].

Theorem 10. *For any graph G, there is a Fractional Maximum Matching f such that for each edge $e \in E$, $f(e) \in \left\{0, \frac{1}{2}, 1\right\}$.*

We prove a game-theoretic analog of Theorem 10 with Nash equilibria (of Defender-Uniform graphs) in place of Fractional Maximum Matchings.

Theorem 11. *For a Defender-Uniform graph G, there is a Defender Uniform Nash equilibrium $\mathbf{s}$ such that for each $v \in \mathsf{Support}_{\mathbf{s}}(vp)$, $\frac{\mathsf{VP}_{\mathbf{s}}(v)}{\max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')} \in \left\{\frac{1}{2}, 1\right\}$.*

Sketch of Proof. Transform a Defender Uniform Nash equilibrium $\mathbf{s}'$ for G into an Attacker Symmetric (and still Defender Uniform) profile $\mathbf{s}$:

1. $\mathbf{s}'_{ep} := \mathbf{s}_{ep}$.
2. For each player $vp_i \in \mathcal{NP}_{vp}$, for each vertex $v \in V$:

$$s_i(v) := \begin{cases} \frac{\max_{e' \in E} \mathsf{VP}_{\mathbf{s}'}(e')}{\nu}, & \text{if } \mathsf{VP}_{\mathbf{s}'}(v) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}'}(e') \\ \frac{\max_{e' \in E} \mathsf{VP}_{\mathbf{s}'}(e')}{2\nu}, & \text{if } 0 < \mathsf{VP}_{\mathbf{s}'}(v) < \max_{e' \in E} \mathsf{VP}_{\mathbf{s}'}(e') \\ 0, & \text{if } \mathsf{VP}_{\mathbf{s}'}(v) = 0 \end{cases}$$

Note that, by construction, $\mathsf{Support}_{\mathbf{s}}(ep) = \mathsf{Support}_{\mathbf{s}'}(ep)$ and $\mathsf{Support}_{\mathbf{s}}(i) = \mathsf{Support}_{\mathbf{s}'}(vp)$. We prove:

Lemma 3. *For each edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(ep)$, $\mathsf{VP}_{\mathbf{s}}(e) = \max_{e' \in E} \mathsf{VP}_{\mathbf{s}}(e')$.*

Lemma 4. $\sum_{v \in V} \mathsf{VP}_{\mathbf{s}}(v) = \nu$.

Sketch of Proof. By Theorem 4, the graph $G(E')$ is partitioned into two subgraphs: *(i)* the r-regular graph $\langle V'_r,\ \mathsf{Edges}_G(V'_r) \cap E' \rangle$, and *(ii)* the graph $\langle V'_i \cup (V \backslash V'), \mathsf{Edges}_G(V'_i \cup (V \backslash V')) \cap E' \rangle$. We will separately calculate the sums $\sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}}(v)$ and $\sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}}(v)$.

We consider first the sum $\sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}}(v)$ and show that $\sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}}(v) = \sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}'}(v)$. We next consider the sum $\sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}}(v)$ and show that $\sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}}(v) = \sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}'}(v)$. Thus,

$$\begin{aligned} &\sum_{v \in V} \mathsf{VP}_{\mathbf{s}}(v) \\ &= \sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}}(v) + \sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}}(v) \\ &= \sum_{v \in V'_r} \mathsf{VP}_{\mathbf{s}'}(v) + \sum_{v \in V'_i \cup (V \backslash V')} \mathsf{VP}_{\mathbf{s}'}(v) \\ &= \sum_{v \in V} \mathsf{VP}_{\mathbf{s}'}(v) \\ &= \nu \qquad \text{(since } \mathbf{s}' \text{ is a profile).} \end{aligned}$$

□

Lemma 5. $\mathbf{s}$ *is a profile.*

It remains to prove that $\mathbf{s}$ is a Nash equilibrium. We prove that $\mathbf{s}$ satisfies conditions (1) and (2) in the characterization of Nash equilibria (Theorem 1).

- By the construction of $\mathbf{s}$, $s_{ep} = s'_{ep}$. This implies that for each vertex $v \in V$, $P_{\mathbf{s}}(\mathsf{Hit}(v)) = P_{\mathbf{s'}}(\mathsf{Hit}(v))$. Hence, in particular, $\min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v')) = \min_{v' \in V} P_{\mathbf{s'}}(\mathsf{Hit}(v'))$.
  Consider any vertex $v \in \mathsf{Support}_{\mathbf{s}}(vp)$. Since $\mathsf{Support}_{\mathbf{s}}(vp) = \mathsf{Support}_{\mathbf{s'}}(vp)$, $v \in \mathsf{Support}_{\mathbf{s'}}(vp)$. Hence, by Condition (2) in the characterization of Nash equilibria (Theorem 1), $P_{\mathbf{s'}}(\mathsf{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s'}}(\mathsf{Hit}(v'))$. Hence,
  $$\begin{aligned} P_{\mathbf{s}}(\mathsf{Hit}(v)) &= P_{\mathbf{s'}}(\mathsf{Hit}(v)) \\ &= \min_{v' \in V} P_{\mathbf{s'}}(\mathsf{Hit}(v')) \\ &= \min_{v' \in V} P_{\mathbf{s}}(\mathsf{Hit}(v')), \end{aligned}$$
  which proves Condition (1).
- Condition (2) is established in Lemma 3.

The proof is now complete. □

6 Epilogue

In this work, we continued the study of a network security game with attackers and a defender, introduced in [8]. We focused on the Price of Defense, introduced in [7] as a worst-case measure of security loss. We proved an optimal lower bound on the Price of Defense, and we provided an efficient characterization of graphs attaining the optimal lower bound. The characterization revealed a rich connection to Fractional Graph Theory, which we explored to show an interesting combinatorial (bivaluednsess) property of Nash equilibria.

Understanding the combinatorial structure of Nash equilibria for our network security game (and, more generally, for strategic games modeling security attacks and defenses) will provide key insights into the design of defense mechanisms. Quantifying the Price of Defense for other, more realistic variants of the network game remains a thrilling challenge. It will be interesting to see if Fractional Graph Theory will still be handy in this endeavor.

Extending Theorem 11 to the class of all graphs, or proving that such an extension is *not* possible, remains an interesting open problem.

References

1. J. Aspnes, K. Chang and A. Yampolskiy, "Inoculation Strategies for Victims of Viruses and the Sum-of-Squares Problem", *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 43–52, 2005.

2. J.-M. Bourjolly and W. R. Pulleyblank, "König-Egerváry Graphs, 2-Bicritical Graphs and Fractional Matchings", *Discrete Applied Mathematics*, Vol. 24, pp. 63–82, 1989.
3. E. R. Cheswick and S. M. Bellovin, *Firewalls and Internet Security*, Addison-Wesley, 1994.
4. M. Gelastou, M. Mavronicolas, V. Papadopoulou, A. Philippou and P. Spirakis, "The Power of the Defender", *CD-ROM Proceedings of the 2nd International Workshop on Incentive-Based Computing*, July 2006.
5. 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.
6. 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.
7. M. Mavronicolas, L. Michael, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, "The Price of Defense", *Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science*, pp. 717–728. Vol. 4162, LNCS, 2006.
8. 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.
9. M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, "A Graph-Theoretic Network Security Game", *Proceedings of the 1st International Workshop on Internet and Network Economics*, pp. 969–978, Vol. 3828, LNCS, 2005.
10. M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing", *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, 2001.
11. T. Moscibroda, S. Schmid and R. Wattenhofer, "When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game", *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing*, pp. 35–44, 2006.
12. J. F. Nash, "Equilibrium Points in N-Person Games", *Proceedings of National Acanemy of Sciences of the United States of America*, pp. 48–49, Vol. 36, 1950.
13. J. F. Nash, "Non-Cooperative Games", *Annals of Mathematics*, Vol. 54, No. 2, pp. 286–295, 1951.
14. E. R. Scheinerman and D. H. Ullman, *Fractional Graph Theory*, Wiley-Interscience Series in Discrete Mathematics and Optimization, 1997.