# How Many Attackers Can Selfish Defenders Catch?[*]

*Marios Mavronicolas*[†]      *Burkhard Monien*[‡]      *Vicky G. Papadopoulou*[†]

## Abstract

*In a distributed system with attacks and defenses, an economic investment in defense mechanisms aims at increasing the degree of system protection against the attacks. We study such investments in the selfish setting, where both attackers and defenders are self-interested entities. In particular, we assume a reward-sharing scheme among interdependent defenders; each defender wishes to maximize its own fair share of the attackers caught due to him (and possibly due to the involvement of others).*

*Addressed in this work is the fundamental question of determining the* maximum *amount of protection achievable by a number of such defenders against a number of attackers if the system is in a Nash equilibrium. As a measure of system protection, we adapt the Defense-Ratio [12], which describes the expected proportion of attackers caught by defenders. In a Defense-Optimal Nash equilibrium, the Defense-Ratio is optimized. We discover that the answer to this question depends in a quantitatively subtle way on the invested number of defenders. We identify graph-theoretic* thresholds *for the number of defenders that determine the possibility of optimizing a Defense-Ratio. In this vein, we obtain, through an extensive combinatorial analysis of Nash equilibria, a comprehensive collection of* trade-off *results.*

## 1 Introduction

**The Model and its Motivation.** *Safety* and *security* have traditionally been included among the key issues for the design and operation of a distributed system. With the unprecedented advent of the Internet, there is a growing interest among the *Distributed Computing* community in formalizing, designing and analyzing distributed systems prone to security *attacks* and *defenses*. A new dimension is that Internet *hosts* and *clients* are controlled by *selfish* agents whose interest is the local maximization of their own benefits (rather than optimizing global performance). So, it is challenging to consider the *simultaneous* impact of selfish and *malicious* behavior of Internet agents. In this work, a distributed system is modeled as a graph $G = (V, E)$; nodes represent the *hosts* and edges represent the *links*.

An *attacker* (also called *virus*) is a malicious client that targets a host to destroy. Associating attacks with nodes make sense since malicious attacks are often targeted at destroying individual servers. A *defender* is a non-malicious client modeling the *antivirus software* implemented on a link in order to protect its two connected hosts. Associating defenses with edges is motivated by *Network Edge Security* [8]; this is a recently proposed, distributed, *firewall architecture*, where antivirus software, rather than being statically installed and licensed at a host, is implemented by a *distributed algorithm* running on a specific subnetwork. Such distributed implementations are attractive since they offer to the hosts more fault-tolerance and the benefit of sharing the licensing costs. In this work, we focus on the simplest possible case where the subnetwork is just a *single* link; a precise understanding of the mathematical pitfalls of attacks and defenses for this simplest case is a necessary prerequisite to making progress for the general case.

Since malicious attacks are *independent*, each trying to maximize the amount of harm it causes during its lifetime, it is natural to model each attacker as a *strategic player* wishing to maximize the chance of escaping the antivirus software; thus, the strategy of one attacker does not (directly) affect the profit of another. In contrast, one may consider at least three approaches for modeling the defenses: (1) Defenses are not strategic at all; such an assumption would lead to a (*centralized*) optimization problem of computing the best locations for the defenders (given that attackers are strategic). (2) Defenses are strategic, and they *cooperate* to maximize the number of caught viruses. This is modeled by assuming a *single* (strategic) defender, which centrally chooses multiple links and it has been studied in [5]. (3) Defenses are strategic and *non-cooperative*.

We have chosen the third approach. This choice is motivated as follows: (1) In a large network, the defense policies are *independent* and *decentralized*. Hence, it may not be so realistic to assume that a *centralized* entity coordinates all defenses. (2) There are financial incentives offered by hosts to heterogeneous (locally installed) defense mechanisms on the basis of *effectiveness* (i.e., number of sustained attacks); for example, prices for antivirus software may be determined on the basis of *recommendation systems*, which collect data about effectiveness from scrutinized hosts. Such incentives induce a natural *competition* among the defenses. (3) Think of a *network owner*, who is interested in maximizing the protection of the network against attacks; the selfish owner has subcontracted the task to a set of independent, deployable agents, and tries to optimize the protection in order to be paid more.

We justify the assumption that defenses are non-cooperative by considering an intuitive *reward-sharing* scheme among the defenders. When more than one colocated defenders are extinguishing the same attacker(s), each will be rewarded with the *fair share* of the number of attackers caught. Thus, each defender is a strategic player wishing to maximize its fair share of the number of attackers caught. We assume that there are $\nu$ attackers and $\mu$ defenders; they are allowed to use mixed strategies. In a *Nash equilibrium* [13, 14], no *player* can unilaterally increase its (expected) *profit*. Motivated by the *Price of Stability* [1], we study *Defense-Optimal* Nash equilibria, where the ratio of the expected number of attackers extinguished by the defenders, over the optimum $\nu$, called *Defense-Ratio*, is as small as possible. (Contrast this to *worst-case* equilibria and the *Price of Anarchy* [6].) The very special but yet highly non-trivial case of this model with a single defender was already introduced in [12] and further studied in [5, 9, 10, 11].

**The Game.** Fix integers $\nu \geq 1$ and $\mu \geq 1$. Associated with $G$ is a *strategic game* $\Pi_{\nu,\mu}(G)$ on $G$:

- The set of *players* is $\mathcal{N} = \mathcal{N}_{\mathsf{A}} \cup \mathcal{N}_{\mathsf{D}}$, where $\mathcal{N}_{\mathsf{A}}$ contains $\nu$ *attackers* $\mathsf{A}_i$ and $\mathcal{N}_{\mathsf{D}}$ contains $\mu$ *defenders* $\mathsf{D}_i$.

- The *strategy set* $S_{\mathsf{A}_i}$ of attacker $\mathsf{A}_i$ is $V$, and the *strategy set* $S_{\mathsf{D}_i}$ of defender $\mathsf{D}_i$ is $E$. So, the *strategy set* $S$ of the game is $S = (\times_{\mathsf{A}_i \in \mathcal{N}_{\mathsf{A}}} S_{\mathsf{A}_i}) \times (\times_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} S_{\mathsf{D}_i}) = V^{\nu} \times E^{\mu}$.
  A *profile* (or *pure profile*) is a $(\nu + \mu)$-tuple $\mathbf{s} = \langle s_{\mathsf{A}_1}, \ldots, s_{\mathsf{A}_{\nu}}, s_{\mathsf{D}_1}, \ldots, s_{\mathsf{D}_{\mu}} \rangle \in S$.

- – The *Individual Profit* of attacker $\mathsf{A}_i$ is a function $\mathsf{IP}_{\mathsf{A}_i} : S \to \{0, 1\}$ with

$$
\mathsf{IP}_{\mathsf{A}_i}(\mathbf{s}) = \begin{cases} 0, & s_{\mathsf{A}_i} \in \bigcup_{\mathsf{D}_j \in \mathcal{N}_{\mathsf{D}}} \{s_{\mathsf{D}_j}\} \\ 1, & s_{\mathsf{A}_i} \notin \bigcup_{\mathsf{D}_j \in \mathcal{N}_{\mathsf{D}}} \{s_{\mathsf{D}_j}\} \end{cases}
$$

Intuitively, when the attacker $\mathsf{A}_i$ chooses vertex $v$, he receives 0 if it is caught by a defender; otherwise, he receives 1.

– The *Individual Profit* of defender $\mathsf{D}_j$ is a function $\mathsf{IP}_{\mathsf{D}_j} : S \to \mathbb{R}$ with

$$
\begin{aligned}
\mathsf{IP}_{\mathsf{D}_j}(\mathbf{s}) = {} & \frac{1}{|\mathsf{defenders}_{\mathbf{s}}(u)|} \cdot |\{\mathsf{A}_i \mid s_{\mathsf{A}_i} = u\}| \\
& + \frac{1}{|\mathsf{defenders}_{\mathbf{s}}(v)|} \cdot |\{\mathsf{A}_i \mid s_{\mathsf{A}_i} = v\}|,
\end{aligned}
$$

where $(u, v) = s_{\mathsf{D}_j}$ and for each vertex $v \in V$, $\mathsf{defenders}_{\mathbf{s}}(v) = \{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}} \mid v \in s_{\mathsf{D}_i}\}$ Intuitively, the defender $\mathsf{D}_j$ receives the *fair share* of the total number of attackers choosing each of the two end vertices of the edge it chooses.

In the sequel, we will, by abuse of notation, use $\mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i)$ and $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i)$ for $\mathsf{IP}_{\mathsf{A}_i}(\mathbf{s})$ and $\mathsf{IP}_{\mathsf{D}_i}(\mathbf{s})$, respectively; we do so in order to emphasize reference to the player rather than to $\mathbf{s}$.

Assume that $v \in s_{\mathsf{D}_i}$. Then, the *proportion* $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v)$ of defender $\mathsf{D}_i$ on vertex $v$ in the profile $\mathbf{s}$ is given by $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) = \frac{1}{|\mathsf{defenders}_{\mathbf{s}}(v)|}$.

The profile $\mathbf{s}$ is a *pure Nash equilibrium* [13, 14] if for each player $i \in \mathcal{N}$, it maximizes $\mathsf{IP}_i(\mathbf{s})$ over all profiles $\mathbf{t}$ that differ from $\mathbf{s}$ only with respect to the strategy of player $i$; so, a pure Nash equilibrium is a local maximizer for the Individual Profit of each player. Say that $G$ *admits a pure Nash equilibrium*, or $G$ is *pure*, if there is a pure Nash equilibrium for the strategic game $\Pi_{\nu,\mu}(G)$.

A *mixed strategy* for player $i \in \mathcal{N}$ is a probability distribution over $S_i$; so, a mixed strategy for an attacker (resp., a defender) is a probability distribution over vertices (resp., edges). A *mixed profile* (or *profile* for short) $\mathbf{s} = \langle s_{\mathsf{A}_1}, \ldots, s_{\mathsf{A}_{\nu}}, s_{\mathsf{D}_1}, \ldots, s_{\mathsf{D}_{\mu}} \rangle$ is a collection of mixed strategies, one for each player; $s_{\mathsf{A}_i}(v)$ is the probability that attacker $\mathsf{A}_i$ chooses vertex $v$, and $s_{\mathsf{D}_j}(e)$ is the probability that defender $\mathsf{D}_j$ chooses edge $e$. The mixed profile $\mathbf{s}$ induces also an *Expected Individual Profit* $\mathsf{IP}_i(\mathbf{s})$ for each player $i \in \mathcal{N}$, which is the expectation (according to $\mathbf{s}$) of the Individual Profit of player $i$. A mixed profile $\mathbf{s}$ is a *Nash equilibrium* [13, 14] if for each player $i \in \mathcal{N}$, it maximizes $\mathsf{IP}_i(\mathbf{s})$ over all profiles $\mathbf{t}$ that differ from $\mathbf{s}$ only with respect to the mixed strategy of player $i$; so, a Nash equilibrium is a local maximizer of the Expected Individual Profit of each player. (Note that by the celebrated Theorem of Nash [13, 14], $\Pi_{\nu,\mu}(G)$ has at least one Nash equilibrium.)

The *Defense-Ratio* $\mathsf{DR}_{\mathbf{s}}$ of a Nash equilibrium $\mathbf{s}$ is the ratio of the *optimal* gain $\nu$ of the defenders over their expected gain in $\mathbf{s}$; so, $\mathsf{DR}_{\mathbf{s}} = \frac{\nu}{\sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i)}$. Clearly, it is desirable that a Nash equilibrium $\mathbf{s}$ maximizes the sum $\sum_{\mathsf{D}_i \in \mathcal{N}} \mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i)$, representing the total gain of all defenders; equivalently, $\mathbf{s}$ should minimize $\mathsf{DR}_{\mathbf{s}}$.

**Summary of Results.** We are interested in the possibility of achieving, and the complexity of computing, a Defense-

Optimal Nash equilibrium using a *given* number of defenders. Note that the number of defenders in this theoretical model directly translates into the real cost of purchasing and installing several units of (licensed) antivirus software. So, this question addresses the *cost-effectiveness* of an economic investment in security for a distributed system. Through a comprehensive collection of results, we discover that the answer depends in a quantitatively subtle way on the number of defenders: There are two graph-theoretic *thresholds,* namely $\frac{|V|}{2}$ and $\beta'(G)$ – the size of a *Minimum Edge Cover* (cf. Section 2, second paragraph), which determine this possibility. (Recall that always $\frac{|V|}{2} \leq \beta'(G)$.)

- When either $\mu \leq \frac{|V|}{2}$ or $\mu \geq \beta'(G)$, there are cases with a Defense-Optimal Nash equilibrium.

  - For $\mu \leq \frac{|V|}{2}$, we provide a combinatorial characterization of graphs admitting a Defense-Optimal Nash equilibrium (Theorem 5.3). Roughly speaking, these make a subclass of the class of graphs with a *Fractional Perfect Matching* where it is possible to partition *some* Fractional Perfect Matching into $\mu$ smaller, *vertex-disjoint* Fractional Perfect Matchings so that the total weight (inherited from the Fractional Perfect Matching) in each part is the same $\left(\text{and equal to } \frac{|V|}{2\,\mu}\right)$. We prove that the recognition problem for this subclass, a previously unconsidered, combinatorial problem in *Fractional Graph Theory* [15], is $\mathcal{NP}$-complete (Proposition 5.7). Hence, the decision problem for the existence of a Defense-Optimal Nash equilibrium is $\mathcal{NP}$-complete as well $\left(\text{for } \mu \leq \frac{|V|}{2}\right)$ (Corollary 5.8). A further interesting consequence of the combinatorial characterization $\left(\text{for } \mu \leq \frac{|V|}{2}\right)$ is that if there is a Defense-Optimal Nash equilibrium, then $\mu$ divides $|V|$ (Corollary 5.4).

    On the positive side, we identify a more restricted subclass of graphs (within the class of graphs with a Fractional Perfect Matching), namely those with a *Perfect Matching,* that admit a Defense-Optimal Nash equilibrium in certain, well-characterized and polynomial time recognizable cases (Theorem 5.9).

  - When there are $\mu \geq \beta'(G)$ defenders, we identify two cases where there are Defense-Optimal Nash equilibria with some special structure (namely, the *balanced* Nash equilibria);

these can be computed in polynomial time (Theorems 7.1 and 7.2).

- For the the middle range $\frac{|V|}{2} < \mu < \beta'(G)$ of values of $\mu$, we provide a combinatorial proof that there is *no* graph with a Defense-Optimal Nash equilibrium (Theorem 6.1). This is somehow paradoxical, since with *fewer* defenders $\left(\mu \leq \frac{|V|}{2}\right)$, we already identified cases with a Defense-Optimal Nash equilibrium. Since the value of the Defense-Ratio changes around $\mu = \frac{|V|}{2}$, this paradox may not be wholly surprising.

- For any number of defenders $\mu$, it is always possible to apply a *replication* technique on the defenders in order to transform a Nash equilibrium for the case of one defender into a Nash equilibrium for $\mu > 1$ defenders (Theorem 8.2). Since a Nash equilibrium for the case of one defender can be computed in polynomial time [9], this implies that the same holds for the general case as well. Whenever the original Nash equilibrium (for $\mu = 1$) is Defense-Optimal, the resulting Nash equilibrium (for $\mu > 1$) may get arbitrarily close to (but *never* be) a Defense-Optimal Nash equilibrium. We propose this technique as a compensation for the cases with no Defense-Optimal Nash equilibria.

**Related Work.** We emphasize that the assumption of $\mu > 1$ defenders has required a far more challenging combinatorial and graph-theoretic analysis than for the case of one defender studied in [5, 9, 10, 11, 12]. Hence, we view our work as a **major** generalization of the work in [5, 9, 10, 11, 12] towards the more realistic case of $\mu > 1$ defenders. The notion of Defense-Ratio generalizes a corresponding definition from [9] to the case of $\mu > 1$ defenders. The special case where $\mu = 1$ of Theorem 5.3 was shown in [10]. (Note that this special case allowed for a polynomial time algorithm to decide the existence of and compute a Defense-Optimal Nash equilibrium.)

Due to page constraints, most proofs have been omitted; they may be found in the full version of this paper available at `http://www.cs.ucy.ac.cy/~mavronic/`.

## 2 Background and Preliminaries

**Graph Theory.** For an integer $n \geq 1$, denote $[n] = \{1, \ldots, n\}$. Throughout, we consider a simple undirected graph $G = \langle V, E \rangle$ (with no isolated vertices). We will sometimes model an edge as the set of its two end vertices. For a vertex set $U \subseteq V$, denote as $G(U)$ the subgraph of $G$ induced by $U$. For an edge set $F \subseteq E$, denote as $G(F)$ the subgraph of $G$ induced by $F$; denote as

$\text{Vertices}_G(F) = \{v \in V \mid (u,v) \in F \text{ for some } u \in V\}$. A *component* of $G$ is a maximal connected subgraph of it. Denote as $d_G(u)$ the *degree* of vertex $u$ in $G$. An edge $(u,v) \in E$ is *pendant* if $d_G(u) = 1$ but $d_G(v) > 1$.

A *Vertex Cover* is a vertex set $VC \subseteq V$ such that for each edge $(u,v) \in E$ either $u \in VC$ or $v \in VC$; a *Minimum Vertex Cover* is one that has minimum size (denoted as $\beta(G)$). An *Edge Cover* is an edge set $EC \subseteq E$ such that for each vertex $v \in V$, there is an edge $(u,v) \in EC$; a *Minimum Edge Cover* is one that has minimum size (denoted as $\beta'(G)$). Denote as $\mathcal{EC}(G)$ the set of all Edge Covers of $G$.

A *Matching* is a set $M \subseteq E$ of non-incident edges; a *Maximum Matching* is one that has maximum size. The first polynomial time algorithm to compute a Maximum Matching appears in [3]. It is known that computing a Minimum Edge Cover reduces to computing a Maximum Matching. (See, e.g., [16, Theorem 3.1.22].) A *Perfect Matching* is a Matching that is also an Edge Cover; so, a Perfect Matching has size $\frac{|V|}{2}$. A *Fractional Matching* is a function $f : E \to [0,1]$ such that for each vertex $v \in V$, $\sum_{e|v \in e} f(e) \leq 1$. (Matching is the special case where $f(e) \in \{0,1\}$ for each edge $e \in E$.) For a Fractional Matching $f$, denote as $E_f = \{e \in E \mid f(e) > 0\}$. A Fractional Maximum Matching is a Fractional Matching $f$ such that $\sum_{e \in E} f(e)$ is maximum over all Fractional Matchings. A Fractional Edge Minimal Matching is a Fractional Matching $f$ such that $|E_f|$ is minimum over all Fractional Matchings. A *Fractional Perfect Matching* is a Fractional Matching $f$ with $\sum_{e|v \in e} f(e) = 1$ for all vertices $v \in V$. It is known that the class of graphs with a Fractional Perfect Matching is recognizable in polynomial time. (See [2] for an efficient combinatorial algorithm.) The same holds for the corresponding search problem.

We prove that for a Fractional Perfect Matching $f$, the graph $G(E_f)$ has no pendant edges. This implies that each component of the graph $G(E_f)$ is either a single edge or a subgraph without pendant edges. Given two Fractional Matchings $f$ and $f'$, write that $f' \subseteq f$ ($f' \subset f$) if $E_{f'} \subseteq E_f$ ($E_{f'} \subset E_f$). Say that two Fractional Matchings $f$ and $f'$ are *equivalent* if for each vertex $v \in V$, $\sum_{e|v \in e} f'(e) = \sum_{e|v \in e} f(e)$. We present two reduction techniques for the simplification of Fractional (Perfect) Matchings. We prove:

**Proposition 2.1** *There is a polynomial time algorithm to transform a Fractional Matching $f$ for a graph $G$ into an equivalent Fractional Matching $f' \subseteq f$ for $G$ such that $G(E_{f'})$ has no even cycle.*

**Proposition 2.2** *Consider a Fractional Perfect Matching $f$ for a graph $G$ such that $G(E_f)$ has no even cycle. Then, there is a polynomial time algorithm to transform $f$ into an equivalent Fractional Perfect Matching $f' \subseteq f$ such that any odd cycle in the graph $G(E_{f'})$ is a component of $f'$.*

To prove Propositions 2.1 and 2.2, we present and analyze two algorithms:

---

**Algorithm EliminateEvenCycles**

INPUT: A graph $G(V,E)$ with a Fractional Matching $f$.
OUTPUT: An equivalent Fractional Matching $f' \subseteq f$ for $G$ such that $G(E_{f'})$ has no even cycle.

**While** $G(E_f)$ has an even cycle $\mathcal{C}$ **do**:

(1) Choose an edge $e_0 \in E(\mathcal{C})$ such that
$$f(e_0) = \min_{e \in E(\mathcal{C})} f(e).$$

(2) Define a function $g : E(\mathcal{C}) \to \{-1, 0, +1\}$ with $g(e) = +1$ or $-1$ alternately, starting with $g(e_0) = -1$.

(3) For each edge $e \in E$, set
$$f'(e) := \begin{cases} f(e) + g(e) \cdot f(e_0), & \text{if } e \in E(\mathcal{C}) \\ f(e), & \text{if } e \notin E(\mathcal{C}) \end{cases}$$

(4) Set $f := f'$.

---

**Algorithm IsolateOddCycles**

INPUT: A graph $G(V,E)$ and a Fractional Perfect Matching $f$ for $G$ such that $G(E_f)$ has no even cycles.
OUTPUT: An equivalent Fractional Perfect Matching $f' \subseteq f$ for $G$ such that any odd cycle in $G(E_{f'})$ is a component.

**While** $G(E_f)$ has an odd cycle $\mathcal{C}$ that is not a component **do**:

(1) Take any vertex $v_0 \in V(\mathcal{C})$ with $d_{G(E_f)}(v_0) \geq 3$ and an edge $e_0 = (v_0, v_1) \in E_f$ with $v_1 \notin V(\mathcal{C})$.

(2) **While** $E(\mathcal{C}) \cup \{e_0\} \subseteq E_f$ **do**:

  (2/a) Find a DFS path $v_1, v_2, \cdots, v_r$ with $v_r = v_l$ for some $l$, $1 \leq l < r - 1$.

  (2/b) Define a function $g : E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r - 1\} \to \{+1, -1, +\frac{1}{2}, -\frac{1}{2}\}$ so that
$$g(e) = \begin{cases} +1 \text{ or } -1 \text{ (alternatily, starting with } -1), \\ \qquad \text{if } e = (v_i, v_{i+1}) \text{ with } 0 \leq i \leq l-1 \\ +\frac{1}{2} \text{ or } -\frac{1}{2} \text{ (alternatily, starting with } +\frac{1}{2}), \\ \qquad \text{if } e \in E(\mathcal{C}) \\ +\frac{1}{2} \text{ or } -\frac{1}{2} \text{ (alternatily, starting with} \\ -\text{sgn}(g(v_{l-1}, v_l)), \\ \qquad \text{if } e = (v_i, v_{i+1}) \text{ with } l \leq i \leq r-1 \end{cases}$$

  (2/c) Find $e'$ that realizes $\min\{\min_{0 \leq i \leq l-1} f((v_i, v_{i+1}))$ $2\min_{e \in E(\mathcal{C})} f(e), 2\min_{l \leq i \leq r-1} f((v_i, v_{i+1}))\}$.

  (2/d) If $g(e') > 0$ then set $g := -g$.

  (2/e) For each edge $e \in E$, set
$$f'(e) := \begin{cases} f(e) + g(e) \cdot \min\{\min_{0 \leq i \leq l-1} f((v_i, v_{i+1})), \\ 2\min_{e \in E(\mathcal{C})} f(e), 2\min_{l \leq i \leq r-1} f((v_i, v_{i+1}))\}, \\ \qquad \text{if } e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \\ f(e), \qquad \text{otherwise} \end{cases}$$

  (2/f) Set $f := f'$.

Propositions 2.1 and 2.2 are reminiscent to (but different than) some results about Fractional Matchings from [15, Theorem 2.1.5]. Specifically, it is shown in [15, Theorem 2.1.5] that for every Fractional Maximum Matching $f$ that is also Minimal (*i*) $G(E_f)$ has no even cycle, and (*ii*) every odd cycle in $G(E_f)$ is a component. Properties (*i*) and (*ii*) imply that the corresponding, naturally defined decision problems about Fractional Maximum and Minimal Matchings are both trivial. In contrast, the corresponding decision problems about Fractional Perfect Matchings are non-trivial since there are Fractional Perfect Matchings $f$ such that $G(E_f)$ has an even cycle, and there are ones such that $G(E_f)$ has an odd cycle that is not a component. (Note that the classes of Fractional Perfect Matchings and Fractional Maximum and Minimal Matchings are *incomparable*.) Propositions 2.1 and 2.2 establish that *search problems* corresponding to these decision problems are solvable in polynomial time. We note that the proofs for Propositions 2.1 and 2.2 have been inspired by those for [15, Theorem 2.1.5].

## 3 Framework

Fix now a mixed profile **s**. The ***support*** of player $i \in \mathcal{N}$ in the profile **s**, denoted as $\mathsf{Support}_\mathbf{s}(i)$, is the set of pure strategies in $S_i$ to which $i$ assigns strictly positive probability. Denote as $\mathsf{Supports}_\mathbf{s}(A) = \bigcup_{A_i \in \mathcal{N}_A} \mathsf{Support}_\mathbf{s}(A_i)$; denote as $\mathsf{Supports}_\mathbf{s}(D) = \bigcup_{D_i \in \mathcal{N}_D} \mathsf{Support}_\mathbf{s}(D_i)$. A vertex $v$ is ***multidefender*** in the profile **s** if $|\{D_i \in \mathcal{N}_D \mid$ there is an edge $e \in \mathsf{Support}_\mathbf{s}(D_i)$ such that $v \in e\}| \geq 2$; that is, a multidefender vertex is "hit" by more than one defenders. Else, the vertex $v$ is ***unidefender***. A profile **s** is ***unidefender*** if every vertex $v \in V$ is unidefender in **s**.

A mixed profile **s** induces a probability measure $\mathbb{P}_\mathbf{s}$ in the natural way. Fix a vertex $v \in V$ and an edge $e \in E$. For a defender $D_i$, denote as $\mathsf{Hit}(D_i, v)$ the event that defender $D_i$ chooses an edge incident to vertex $v$. Denote as $\mathsf{Hit}(v)$ the event that some defender chooses an edge incident to vertex $v$. Clearly, $\mathsf{Hit}(v) = \bigcup_{D_i \in \mathcal{N}_D} \mathsf{Hit}(D_i, v)$. Hence, by the *Principle of Inclusion-Exclusion,* $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \sum_{j \in [\mu]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_D \mid |\mathcal{D}|=j} \prod_{D_k \in \mathcal{D}} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_k, v))$. From this expression, we immediately observe:

**Lemma 3.1** *Assume that vertex $v$ is multidefender in* **s**. *Then,* $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) < \sum_{D_i \in \mathcal{N}_D} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_i, v))$.

A vertex $v \in V$ is ***maxhit*** in the profile **s** if $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 1$; a defender $D_i \in \mathcal{N}_D$ is a ***maxhitter*** in **s** if there is an edge $e \in \mathsf{Support}_\mathbf{s}(D_i)$ such that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_i, v)) = 1$ for some vertex $v \in e$. We prove:

**Lemma 3.2** *For a profile* **s**, $\sum_{v \in V} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) \leq 2\mu$. *(and* $< 2\mu$ *if there is a multidefender vertex).*

Denote as $\mathsf{MinHit}_\mathbf{s} = \min_{v \in V} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))$, the ***Minimum Hitting Probability*** associated with **s**. Denote as $\mathsf{VP}_\mathbf{s}(v)$ the expected number of attackers choosing vertex $v$ (according to **s**); so, $\mathsf{VP}_\mathbf{s}(v) = \sum_{i \in [\nu]} s_{A_i}(v)$. For an edge $e = (u, v) \in E$, $\mathsf{VP}_\mathbf{s}(e) = \mathsf{VP}_\mathbf{s}(u) + \mathsf{VP}_\mathbf{s}(v)$. We observe:

**Lemma 3.3** *For a profile* **s**, $\mathsf{MinHit}_\mathbf{s} \leq \frac{2\mu}{|V|}$.

- Induced by **s** is also the ***Conditional Expected Individual Profit*** $\mathsf{IP}_\mathbf{s}(A_i, v)$ of attacker $A_i \in \mathcal{N}_A$ on vertex $v$, which is the conditional expectation (according to **s**) of the Individual Profit of attacker $A_i$ had he chosen vertex $v$. So, $\mathsf{IP}_\mathbf{s}(A_i, v) = 1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))$. Then, the Expected Individual Profit $\mathsf{IP}_\mathbf{s}(A_i)$ is $\mathsf{IP}_\mathbf{s}(A_i) = \sum_{v \in V} s_{A_i}(v) \cdot \mathsf{IP}_\mathbf{s}(A_i, v) = \sum_{v \in V} s_{A_i}(v) \cdot (1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)))$.

- The ***Conditional Expected Proportion*** $\mathsf{Prop}_\mathbf{s}(D_i, v)$ of defender $D_i \in \mathcal{N}_D$ on vertex $v$ is his conditional expected proportion on vertex $v$ had he chosen an edge incident to vertex $v$:

$$\mathsf{Prop}_\mathbf{s}(D_i, v)$$
$$= \sum_{j \in [m]} \frac{1}{j} \sum_{\mathcal{D} \subseteq \mathcal{N}_D \setminus \{D_i\} \mid |\mathcal{D}|=j-1} \prod_{D_k \in \mathcal{D}} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_k, v))$$
$$\prod_{D_k \notin \mathcal{D} \cup \{D_i\}} (1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_k, v)))$$
$$= \sum_{j \in [m]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_D \setminus \{D_i\} \mid |\mathcal{D}|=j-1}$$
$$\prod_{D_k \in \mathcal{D}} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_k, v)).$$

- The ***Conditional Expected Individual Profit*** $\mathsf{IP}_\mathbf{s}(D_i, v)$ of defender $D_i$ on edge $e = (u, v) \in E$ is the conditional expectation (according to **s**) of the Individual Profit of defender $D_i$ had he chosen edge $e$. So, $\mathsf{IP}_\mathbf{s}(D_i, e) = \mathsf{Prop}_\mathbf{s}(D_i, u) \cdot \mathsf{VP}_\mathbf{s}(u) + \mathsf{Prop}_\mathbf{s}(D_i, v) \cdot \mathsf{VP}_\mathbf{s}(v)$. Then, the Expected Individual Profit $\mathsf{IP}_\mathbf{s}(D_i)$ of defender $D_i$ takes a particularly simple form:

$$\mathsf{IP}_\mathbf{s}(D_i) = \sum_{v \in V} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_i, v)) \cdot \mathsf{Prop}_\mathbf{s}(D_i, v) \cdot \mathsf{VP}_\mathbf{s}(v).$$

**Lemma 3.4** *Fix a mixed profile* **s**. *Then, for any* $v \in V$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \sum_{D_i \in \mathcal{N}_D} \mathbb{P}_\mathbf{s}(\mathsf{Hit}(D_i, v)) \cdot \mathsf{Prop}_\mathbf{s}(D_i, v)$.

Clearly, in a Nash equilibrium **s**, for each attacker $A_i$, $\mathsf{IP}_\mathbf{s}(A_i, v)$ is *constant* over all vertices $v \in \mathsf{Support}_\mathbf{s}(A_i)$; for each defender $D_i$, $\mathsf{IP}_\mathbf{s}(D_i, e)$ is *constant* over all edges $e \in \mathsf{Support}_\mathbf{s}(D_i)$. It follows that in a Nash equilibrium **s**, for each attacker $A_i$, $\mathsf{IP}_\mathbf{s}(A_i) = 1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))$ for any vertex $v \in \mathsf{Support}_\mathbf{s}(A_i)$; for each defender $D_i$, $\mathsf{IP}_\mathbf{s}(D_i) = \mathsf{Prop}_\mathbf{s}(D_i, u) \cdot \mathsf{VP}_\mathbf{s}(u) + \mathsf{Prop}_\mathbf{s}(D_i, v) \cdot \mathsf{VP}_\mathbf{s}(v)$, for any edge $e = (u, v) \in \mathsf{Support}_\mathbf{s}(D_i)$. Hence, for each attacker $A_i$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))$ is *constant* over all vertices $v \in \mathsf{Support}_\mathbf{s}(A_i)$.

**Some notation.** Set $\mathsf{Edges_s}(v) = \{(u, v) \in E \mid (u, v) \in \mathsf{Support_s}(\mathsf{D})\}$. For an edge set $F \subseteq E$, set $\mathsf{Vertices_s}(F) = \{u \in \mathsf{Supports_s}(\mathsf{A}) \mid (u, v) \in F$ for some $v \in V\}$.

**Some special profiles.** A profile **s** is *uniform* if each player uses a *uniform* probability distribution on its support; so, for each attacker $\mathsf{A}_i$, for each vertex $v \in \mathsf{Support_s}(\mathsf{A}_i)$, $s_{\mathsf{A}_i}(v) = \frac{1}{|\mathsf{Support_s}(\mathsf{A}_i)|}$, and for each defender $\mathsf{D}_i$, for each edge $e \in E$, $s_{\mathsf{D}_i}(e) = \frac{1}{|\mathsf{Support_s}(\mathsf{D}_i)|}$. A profile **s** is *attacker symmetric* (resp., *defender symmetric*) if for all pairs of attackers $\mathsf{A}_i$ and $\mathsf{A}_k$ (resp., all pairs of defenders $\mathsf{D}_i$ and $\mathsf{D}_k$) for all vertices $v \in V$, (resp., all edges $e \in E$) $s_{\mathsf{A}_i}(v) = s_{\mathsf{A}_k}(v)$ (resp., $s_{\mathsf{D}_i}(v) = s_{\mathsf{D}_k}(v)$). A profile is *attacker symmetric uniform* (resp., *defender symmetric uniform*) if it is attacker symmetric (resp., defender symmetric) and each attacker (resp., defender) uses a uniform probability distribution on his support. A profile is *attacker fully mixed* (resp., *defender fully mixed*) if for each attacker $\mathsf{A}_i$ (resp., for each defender $\mathsf{D}_i$), $\mathsf{Support_s}(\mathsf{A}_i) = V$ (resp., $\mathsf{Support_s}(\mathsf{D}_i) = V$).

# 4 The Structure of Nash Equilibria

We provide an extensive combinatorial analysis of Nash equilibria. We first prove:

**Proposition 4.1 (Characterization of Nash Equilibria)**
*A profile **s** is a Nash equilibrium if and only if the following conditions hold:*

(1) *For each vertex $v \in \mathsf{Supports_s}(\mathsf{A})$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \mathsf{MinHit_s}$.*

(2) *For each defender $\mathsf{D}_i$, for each edge $(u, v) \in \mathsf{Support_s}(\mathsf{D}_i)$,*

$$\mathsf{Prop_s}(\mathsf{D}_i, v)\,\mathsf{VP_s}(v) + \mathsf{Prop_s}(\mathsf{D}_i, u)\,\mathsf{VP_s}(u) =$$
$$\max_{(u', v') \in E} \{\mathsf{Prop_s}(\mathsf{D}_i, v')\,\mathsf{VP_s}(v') + \mathsf{Prop_s}(\mathsf{D}_i, u')\,\mathsf{VP_s}(u')\}$$

We remark that Proposition 4.1 generalizes a corresponding characterization of Nash equilibria for $\Pi_{\nu,1}(G)$ shown in [12], where Condition (2) had the simpler counterpart: (2') For each edge $e \in \mathsf{Supports_s}(\mathsf{D})$, $\mathsf{VP_s}(e) = \max_{e' \in E} \{\mathsf{VP_s}(e')\}$. We continue to prove:

**Proposition 4.2** *In a Nash equilibrium **s**, $\sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \mathsf{IP_s}(\mathsf{D}_i) = \nu \cdot \mathsf{MinHit_s}$.*

By the definition of Defense-Ratio, Proposition 4.2 immediately implies:

**Corollary 4.3** *In a Nash equilibrium **s**, $\mathsf{DR_s} = \frac{1}{\mathsf{MinHit_s}}$.*

Corollary 4.3 implies that $\mathsf{DR_s} \geq 1$. Furthermore, we observe:

**Lemma 4.4** *For a Nash equilibrium **s**, $\mathsf{DR_s} \geq \frac{|V|}{2\mu}$.*

We are now ready to provide a significant definition:

**Definition 4.1** *A Nash equilibrium **s** is **Defense-Optimal** if $\mathsf{DR_s} = \max\left\{1, \frac{|V|}{2\mu}\right\}$.*

We will later construct Defense-Optimal Nash equilibria; so, $\max\left\{1, \frac{|V|}{2\mu}\right\}$ is a *tight* lower bound on Defense-Ratio, and this will justify our definition of Defense-Optimal Nash equilibria. Say that $G$ admits a ***Defense-Optimal Nash equilibirum*** or that $G$ is ***Defense-Optimal*** (with respect to the particular parameter $\mu$) if there is a Defense-Optimal Nash equilibrium for the strategic game $\Pi_{\nu,\mu}(G)$. This leads to the formulation of a natural decision problem:

---

DEFENSE OPTIMAL GRAPH
INSTANCE: A graph $G = \langle V, E \rangle$ and an integer $\mu$.
QUESTION: Is $G$ Defense-Optimal (with respect to $\mu$)?

---

We continue to prove:

**Proposition 4.5** *In a Nash equilibrium **s**, $\mathsf{Supports_s}(\mathsf{D})$ is an Edge Cover, and $\mathsf{Supports_s}(\mathsf{A})$ is a Vertex Cover of $G(\mathsf{Supports_s}(\mathsf{D}))$.*

We use Propositions 4.5 and 4.5 to prove:

**Proposition 4.6 (Necessary Condition for Pure Nash Equilibria)** *Assume that $G$ is pure. Then, $\mu \geq \beta'(G)$ and $\nu \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$.*

We finally prove:

**Proposition 4.7** *A Defender Pure Nash equilibrium is Defense-Optimal.*

# 5 Few Defenders

We consider the case of *few* defenders where $\mu \leq \frac{|V|}{2}$. There, a Defense-Optimal Nash equilibrium **s** has $\mathsf{DR_s} = \max\left\{1, \frac{|V|}{2\mu}\right\} = \frac{|V|}{2\mu}$. We start with a structural property of Defense-Optimal Nash equilibria:

**Proposition 5.1** *Assume that $\mu \leq \frac{|V|}{2}$. Then, a Defense-Optimal Nash equilibrium is unidefender.*

**Characterization of Defense-Optimal Graphs.** We continue with a new graph-theoretic definition.

**Definition 5.1** *Fix an integer $\mu \geq 1$. A Fractional Perfect Matching $f : E \to \mathbb{R}$ is $\mu$-**partitionable** if the edge set $E_f$ can be partitioned into $\mu$ non-empty, vertex-disjoint subsets $E_1, \cdots, E_\mu$ so that for each subset $E_i$, $\sum_{e \in E_i} f(e) = \frac{|V|}{2\mu}$.*

Note that for $\mu = 1$, the existence problem for a 1-partitionable Fractional Perfect Matching is trivially the one for a Fractional Perfect Matching, which can be solved in polynomial time [2]. We observe a preliminary property of $\mu$-partitionable Fractional Perfect Matchings:

**Lemma 5.2** *Assume that $G$ has a $\mu$-paritionable Fractional Perfect Matching. Then, $\mu$ divides $|V|$.*

We prove a characterization of Defense-Optimal graphs:

**Theorem 5.3** *Assume that $\mu \leq \frac{|V|}{2}$. Then, a graph $G$ is Defense-Optimal if and only if $G$ has a $\mu$-partitionable Fractional Perfect Matching.*

Theorem 5.3 immediately implies:

**Corollary 5.4** *For $\mu \leq \frac{|V|}{2}$, assume that $G$ is Defense-Optimal. Then, $\mu$ divides $|V|$.*

**Complexity of Defense-Optimal Graphs.** By Theorem 5.3, the complexity of recognizing Defense-Optimal graphs is that of the following, previously unconsidered combinatorial problem from Fractional Graph Theory [15]:

---

$\mu$-PARTITIONABLE FRACTIONAL PERFECT MATCHING
INSTANCE: A graph $G = \langle V, E \rangle$ and an integer $\mu$ such that $\mu$ divides $|V|$.
QUESTION: Is there a $\mu$-partitionable Fractional Perfect Matching for $G$?

---

The restriction to instances for which $\mu$ divides $|V|$ is motivated from Lemma 5.2 to restrict to the set of interesting instances. We use Propositions 2.1 and 2.2 to prove:

**Proposition 5.5** *Assume that $G$ has a $\mu$-partitionable Fractional Perfect Matching $f$. Then, it also has a $\mu$-partitionable Fractional Perfect Matching $f'$ such that $G(E_{f'})$ consists only of single edges and odd cycles. Furthermore, $f'$ can be computed from $f$ in polynomial time.*

We are now ready to prove:

**Proposition 5.6** *A graph $G$ has a $\mu$-partitionable Fractional Perfect Matching if and only if $E$ can be partitioned into a collection $E_1, \cdots, E_\mu$ of $\mu$ vertex-disjoint subsets and corresponding vertex sets $V_1, \cdots, V_\mu$, so that $\bigcup_{i \in [\mu]} V_i = V$, each $E_i$ is a collection of single edges and odd cycles, and $|V_i| = \frac{|V|}{\mu}$, where $i \in [\mu]$.*

We shall show an interesting relation of the problem of deciding the existence of a $\mu$-partitionable Fractional Perfect Matching to a well known graph-theoretic problem:

---

PARTITION INTO TRIANGLES
INSTANCE: A graph $G = \langle V, E \rangle$ with $|V| = 3q$ for some integer $q$.
QUESTION: Can the vertices of $G$ be partitioned into $q$ disjoint sets $V_1, \cdots, V_q$, each containing exactly three vertices, such that the subgraph of $G$ induced by each $V_i$ is $K_3$?

---

This problem is $\mathcal{NP}$-complete [4, GT11, attribution to (personal communication with) Scheafer]. To prove that $\mu$-PARTITIONABLE FRACTIONAL PERFECT MATCHING is $\mathcal{NP}$-complete, we consider a special case of it:

---

SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING
INSTANCE: A graph $G = \langle V, E \rangle$ with $|V| = 3q$ for some integer $q$.
QUESTION: Is there a $\frac{|V|}{3}$-partitionable Fractional Perfect Matching for $G$?

---

**Proposition 5.7** SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING $\equiv$ PARTITION INTO TRIANGLES

Proposition 5.7 gives that SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING is $\mathcal{NP}$-complete. Since SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING is a special case of $\mu$-PARTITIONABLE FRACTIONAL PERFECT MATCHING, we get that $\mu$-PARTITIONABLE FRACTIONAL PERFECT MATCHING is $\mathcal{NP}$-complete as well. Hence, Theorem 5.3 implies:

**Corollary 5.8** *Assume that $\mu \leq \frac{|V|}{2}$. Then, the recognition problem for Defense-Optimal graphs is $\mathcal{NP}$-complete.*

**Graphs with Perfect Matchings.** We now restrict to graphs with Perfect Matchings. We show:

**Theorem 5.9** *Consider a graph $G$ with a Perfect Matching and an integer $\mu \leq \frac{|V|}{2}$. Then, $G$ admits a Defense-Optimal Nash equilibrium $\mathbf{s}$ where $\mathsf{Support_s}(\mathsf{D})$ is a Perfect Matching if and only if $2\mu$ divides $|V|$.*

Theorem 5.9 will follow from Propositions 5.10 and 5.11.

**Proposition 5.10** *Assume a graph $G$ with a Perfect Matching and an integer $\mu \leq \frac{|V|}{2}$ such that $2\mu$ divides $|V|$. Then, $G$ admits a Defense-Optimal Nash equilibrium $\mathbf{s}$ where $\mathsf{Support_s}(\mathsf{D})$ is a Perfect Matching.*

**Proof:** Consider a Perfect Matching $M$. Construct a profile $\mathbf{s}$ as follows:

- $\mathbf{s}$ is Attacker Symmetric Uniform and Attacker Fully Mixed. So, for each attacker $\mathsf{A}_i \in \mathcal{N}_\mathsf{A}$, for each vertex $v \in V$, $s_{\mathsf{A}_i}(v) = \frac{1}{|V|}$ and $\mathsf{VP_s}(v) = \frac{\nu}{|V|}$.

- Partition $M$ into $\mu$ subsets, each with $\frac{|V|}{2\mu}$ edges; each defender uses a uniform probability distribution over each one subset. Thus, $\mathsf{Support_s}(\mathsf{D}) = M$ and each edge is unidefender in $\mathbf{s}$.

We now establish Conditions (1) and (2) in the characterization of Nash equilibria (Proposition 4.1).

- For Condition (1), fix any vertex $v \in V$. Since $M$ is a (Perfect) Matching, there is a single edge $e \in \mathsf{Edges_s}(v)$. Since $e$ is unidefender (say by defender $\mathsf{D}_i$), it follows that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v)) = \frac{2\mu}{|V|}$. Now, Condition (1) follows trivially.

- For Condition (2), consider any defender $\mathsf{D}_i \in \mathcal{N}_\mathsf{D}$. Fix an edge $e = (u, v) \in \mathsf{Support_s}(\mathsf{D}_i)$. Since each edge $e$ is unidefender, it follows that $\mathsf{Prop_s}(\mathsf{D}_i, v) = \mathsf{Prop_s}(\mathsf{D}_i, u) = 1$. It follows that $\mathsf{Prop_s}(\mathsf{D}_i, v) \cdot \mathsf{VP_s}(v) + \mathsf{Prop_s}(\mathsf{D}_i, u) \cdot \mathsf{VP_s}(u) = \frac{2\nu}{|V|}$. On the other hand, fix any edge $e' = (v', u') \notin \mathsf{Support_s}(\mathsf{D}_i)$. Since $M$ is an Edge Cover, it follows that for the vertex $v'$ (resp., vertex $u'$), there is a defender $\mathsf{D}_j$ such that $v' \in \mathsf{Vertices}(\mathsf{Support_s}(\mathsf{D}_j))$ (resp., $u' \in \mathsf{Vertices}(\mathsf{Support_s}(\mathsf{D}_j))$). (Note that $\mathsf{D}_j \neq \mathsf{D}_i$ since $M$ is a Matching.) It follows that $\mathsf{Prop_s}(\mathsf{D}_i, v') < 1$ (resp., $\mathsf{Prop_s}(\mathsf{D}_i, u') < 1$). Thus, $\mathsf{Prop_s}(\mathsf{D}_i, v') \cdot \mathsf{VP_s}(v') + \mathsf{Prop_s}(\mathsf{D}_i, u') \cdot \mathsf{VP_s}(u') < \mathsf{VP_s}(v') + \mathsf{VP_s}(u') = \frac{2\nu}{|V|}$. Now, Condition (2) follows.

Hence, by Proposition 4.1, $\mathbf{s}$ is a Nash equilibrium. To prove that $\mathbf{s}$ is Defense-Optimal, recall that for each vertex $v \in V$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \frac{2\mu}{|V|}$. Hence, $\mathsf{MinHit_s} = \frac{2\mu}{|V|}$. By Corollary 4.3, it follows that $\mathsf{DR_s} = \frac{|V|}{2\mu}$. ∎

**Proposition 5.11** *Assume a graph $G$ with a Perfect Matching and an integer $\mu \leq \frac{|V|}{2}$ such that $G$ admits a Defense-Optimal Nash equilibrium $\mathbf{s}$ where $\mathsf{Support_s}(\mathsf{D})$ is a Perfect Matching. Then, $2\mu$ divides $|V|$.*

**Proof:** Consider such a Nash equilibrium $\mathbf{s}$. Since $\mathbf{s}$ is Defense-Optimal, Corollary 4.3 implies that $\mathsf{MinHit_s} = \frac{2\mu}{|V|}$. Consider any edge $e = (u, v) \in \mathsf{Support_s}(\mathsf{D})$; so, $e \in \mathsf{Support_s}(\mathsf{D}_i)$ for some defender $\mathsf{D}_i \in \mathcal{N}_\mathsf{D}$. Since $\mathsf{Support_s}(\mathsf{A})$ is a vertex Cover of the graph $G(\mathsf{Support_s}(\mathsf{D}))$, it follows that $v \in \mathsf{Support_s}(\mathsf{A})$ or $u \in \mathsf{Support_s}(\mathsf{A})$ (or both). Since $\mathsf{Support_s}(\mathsf{D})$ is a Perfect Matching (and therefore an Edge Cover), there is at least one defender $\mathsf{D}_k$ such that $v \in \mathsf{Vertices_s}(\mathsf{Support_s}(\mathsf{D}_k))$. Since $\mathbf{s}$ is Defense-Optimal, Proposition 5.1 implies that there is at most one defender $\mathsf{D}_k$ such that $v \in \mathsf{Vertices_s}(\mathsf{Support_s}(\mathsf{D}_k))$. It follows that there is exactly one defender $\mathsf{D}_k$ such that $v \in \mathsf{Vertices_s}(\mathsf{Support_s}(\mathsf{D}_k))$. So, clearly, $\mathsf{D}_k$ is $\mathsf{D}_i$. Since $\mathsf{Support_s}(\mathsf{D})$ is a Perfect Matching, this implies that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = s_{\mathsf{D}_i}(e)$. We prove that $|\mathsf{Support_s}(\mathsf{D}_i)| = \frac{|V|}{2\mu}$. Since $|\mathsf{Support_s}(\mathsf{D}_i)|$ is an integer, this implies that $2\mu$ divides $|V|$. ∎

Note that Corollary 5.4 applies to *all* graphs, while Proposition 5.11 applies only to graphs with a Perfect Matching. However, the restriction of Corollary 5.4 to graphs with a Perfect Matching does *not* imply Proposition 5.11 *unless* $\mu$ is odd. (This is because 2 divides $|V|$ and $\mu$ divides $|V|$ imply together that $2\mu$ divides $|V|$ exactly when $\mu$ is odd.)

# 6  Many Defenders

We now consider the case of *many* defenders, where $\frac{|V|}{2} < \mu < \beta'(G)$. Note that in this case, a Defense-Optimal Nash equilibrium has Defense-Ratio $\mathsf{DR_s} = \max\left\{1, \frac{|V|}{2\mu}\right\} = 1$.

**Theorem 6.1 (Non-existence of Defense-Optimal)**
*Assume that $\frac{|V|}{2} < \mu < \beta'(G)$. Then, $G$ admits no Defense-Optimal Nash equilibrium.*

**Proof:** Assume, by way of contradiction, that $G$ admits a Defense-Optimal Nash equilibrium $\mathbf{s}$. So, $\mathsf{DR}_\mathbf{s} = 1$. Corollary 4.3 implies that $\mathsf{MinHit}_\mathbf{s} = 1$. It follows that for each vertex $v \in V$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 1$; so, all vertices are maxhit. So, fix a (maxhit) vertex $v \in V$. The expression

$$\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \sum_{j \in [\mu]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} || \mathcal{D}| = j} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_\mathbf{s}(\mathsf{Hit}_\mathbf{s}(\mathsf{D}_k, v))$$

implies that there is at least one maxhitter $\mathsf{D}_i \in \mathcal{N}_\mathsf{D}$ (a defender $\mathsf{D}_i$ such that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v)) = 1$). There are two cases for each maxhitter $\mathsf{D}_i$: *(i)* $\mathsf{D}_i$ uses a pure strategy $(u, v)$, so that there are two vertices $u, v \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)$ such that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v)) = \mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, u)) = 1$, or *(ii)* $\mathsf{D}_i$ uses a mixed strategy, in which case there is a *single* vertex $v \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)$ such that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v)) = \mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, u)) = 1$.

Use $\mathbf{s}$ to construct a defender-pure profile $\mathbf{t}$ as follows: The pure strategy of each (multihitter or not) defender $\mathsf{D}_k$ is some edge from $\mathsf{Support}_\mathbf{s}(\mathsf{D}_k)$. Note that, by construction of $\mathbf{t}$, (1) $\mathsf{Support}_\mathbf{t}(\mathsf{D}) \leq \mu$, and (2) the number of maxhit vertices in $\mathbf{s}$ is at most the number of maxhit vertices in $\mathbf{t}$. Since $\mu < \beta'(G)$, (1) implies that $|\mathsf{Supports}_\mathbf{t}(\mathsf{D})| < \beta'(G)$. So, there is some vertex $v \in V$ such that $\mathbb{P}_\mathbf{t}(\mathsf{Hit}(v)) = 0$. It follows that the number of maxhit vertices in $\mathbf{t}$ is at most $|V| - 1$. By (2), it follows that the number of maxhit vertices in $\mathbf{s}$ is at most $|V| - 1$. A contradiction. ∎

## 7 Too Many Defenders

We finally turn to the case of *too many* defenders, where $\mu \geq \beta'(G)$. Note that, in this case, a Defense-Optimal Nash equilibrium $\mathbf{s}$ has Defense-Ratio $\mathsf{DR}_\mathbf{s} = \max\left\{1, \frac{|V|}{2\mu}\right\} \leq \max\left\{1, \frac{|V|}{2\beta'(G)}\right\} = 1$ (since $\frac{|V|}{2} \leq \beta'(G)$ for every graph $G$). For the analysis, we will use a special class of profiles that we introduce. A profile $\mathbf{s}$ is *balanced* if there is a constant $c > 0$ such that for each pair of a defender $\mathsf{D}_i \in \mathcal{N}_\mathsf{D}$ and a vertex $v \in V$, $\mathsf{Prop}_\mathbf{s}(\mathsf{D}_i, v) \cdot \mathsf{VP}_\mathbf{s}(v) = c$.

Clearly, in a balanced profile, *(i)* for each defender $\mathsf{D}_i$ and each vertex $v \in V$, $\mathsf{Prop}_\mathbf{s}(\mathsf{D}_i, v) > 0$; and *(ii)* for each vertex $v \in V$, $\mathsf{VP}_\mathbf{s}(v) > 0$. From *(i)*, it follows that the support of each defender is an Edge Cover; note that this (necessary) condition is *stronger* than the necessary condition in Proposition 4.5. From *(ii)*, it follows that the supports of attackers is $V$; note that this (necessary) condition is *weaker* than the condition in the definition of an attacker fully mixed profile. Note also that by definition, a balanced profile satisfies Condition (2) in the characterization of Nash equilibria (Proposition 4.1). We have been unable to construct

*mixed* balanced profiles for the general case. So, we focused on the special case of pure strategies. A *defender-pure balanced* profile is a defender-pure profile $\mathbf{s}$ such that there is a constant $c > 0$ such that for each vertex $v \in V$, $\frac{\mathsf{VP}_\mathbf{s}(v)}{\mathsf{defenders}_\mathbf{s}(v)} = c$. We prove that that a defender-pure balanced profile is a local maximizer for the Individual Profit of each defender.

We will present polynomial time algorithms to compute Defender-Pure Balanced Nash equilibria in two cases. Both algorithms will rely on a polynomial time algorithm for computing a Minimum Edge Cover.

**Defender-Pure Balanced Nash Equilibria.** We show:

**Theorem 7.1** *Assume that $\mu \geq \beta'(G)$. Then, $G$ admits a Defense-Optimal, Defender-Pure Nash equilibrium, which can be computed in polynomial time.*

To prove Theorem 7.1, we present a polynomial time algorithm Defender-Pure&BalancedNE to compute a Defender-Pure Balanced Nash equilibrium:

---

**Algorithm** Defender-Pure&BalancedNE

INPUT: A graph $G(V, E)$ and a pair of integers $\nu$ and $\mu$ such that $\beta'(G) \leq \mu$.

OUTPUT: A Defender-Pure Balanced Nash equilibrium $\mathbf{s}$.

(1) Compute a Minimum Edge Cover $EC = \{(v_i, u_i) \mid i \in [\beta'(G)]\}$.

(2) For each $i \in [\mu]$, set $s_{\mathsf{D}_i} := (v_{i \bmod \beta'(G)}, u_{i \bmod \beta'(G)})$.

(3) Compute a solution $\{\mathsf{VP}(v_i) \mid i \in [|V|]\}$ to the following linear system:

   (a) For each $i \in [|V|]$, $\frac{\mathsf{VP}(v_i)}{\mathsf{defenders}_\mathbf{s}(v_i)} = \frac{\mathsf{VP}(v_1)}{\mathsf{defenders}_\mathbf{s}(v_1)}$; (b) $\sum_{i \in [|V|]} \mathsf{VP}_\mathbf{s}(v_i) = \nu$.

(4) Arbitrarily, assign probability distributions to the attackers so that for each $v_i \in V$, $\mathsf{VP}_\mathbf{s}(v_i) = \mathsf{VP}(v_i)$.

---

**Pure Balanced Nash Equilibria.** We now prove that adding a further constraint to those in Theorem 7.1 allows for a (Defense-Optimal) *Pure* Nash equilibrium.

**Theorem 7.2** *Assume that $\mu \geq \beta'(G)$ and $2\mu$ divides $\nu$. Then, $G$ admits a Defense-Optimal, Pure Nash equilibrium, which can be computed in polynomial time.*

To prove Theorem 7.2, we present a polynomial time algorithm Defender-Pure&Balanced to compute a Pure Balanced Nash equilibrium:

```
Algorithm Pure&BalancedNE
INPUT:  A graph $G(V, E)$ and integers $\mu$ and $\nu$ such that
$\beta'(G) \le \mu$ and $\frac{\nu}{2} \equiv 0 \pmod{\mu}$.
OUTPUT: A Pure Balanced Nash equilibrium $\mathbf{s}$.

 (1) Compute a Minimum Edge Cover $EC = \{(v_i, u_i) \mid i \in
     [\beta'(G)]\}$.

 (2) For each  $i \in [\mu]$,  set  $s_{D_i} := (v_{i \bmod \beta'(G)},$
     $u_{i \bmod \beta'(G)})$.

 (3) Arbitrarily assign pure strategies to the attackers so that
     for each vertex $v \in V$, $\mathsf{VP_s}(v) = \mathsf{defenders_s}(v) \cdot \frac{\nu}{2\mu}$.
```

## 8 Replicated Defenders

We use an involved combinatorial analysis to prove:

**Proposition 8.1** *Consider an arbitrary Nash equilibrium* $\mathbf{s}$
*for the game* $\Pi_{\nu,1}(G)$. *Then, there is a Defender Symmetric
Nash equilibrium* $\mathbf{t}$ *for the game* $\Pi_{\nu,\mu}$ *with* $\mathsf{MinHit_t} = 1 -
(1 - \mathsf{MinHit_s})^{\mu}$.

By Corollary 4.3, Proposition 8.1 immediately implies:

**Theorem 8.2 (From Single to Symmetric Defenders)**
*Consider an arbitrary Nash equilibrium* $\mathbf{s}$ *for the game*
$\Pi_{\nu,1}(G)$. *Then, there is a Defender Symmetric Nash equi-
librium* $\mathbf{t}$ *for the game* $\Pi_{\nu,\mu}(G)$ *with Defense-Ratio* $\mathsf{DR_t} =
\frac{1}{1 - \left(1 - \frac{1}{\mathsf{DR_s}}\right)^{\mu}}$.

It is simple to see that in the setting of Theorem 8.2, $\mathsf{DR_t} \ge
\frac{\mathsf{DR_s}}{\mu}$. (This should be expected since otherwise the lower
bound in Lemma 4.4 could be violated by choosing $\mathbf{s}$ to be
a Perfect Matching Nash equilibrium [9] for $\Pi_{\nu,1}(G)$ with
$\mathsf{DR_s} = \frac{|V|}{2}$.)

## References

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos,
    T. Wexler and T. Roughgarden, "The Price of Stability for
    Network Design with Fair Cost Allocation," *Proceedings of
    the 45th Annual IEEE Symposium on Foundations of Com-
    puter Science,* pp. 295–304, 2004.

[2] J.-M. Bourjolly and W. R. Pulleyblank, "König-Egerváry
    Graphs, 2-Bicritical Graphs and Fractional Matchings," *Dis-
    crete Applied Mathematics,* Vol. 24, No. 1, pp. 63–82, 1989.

[3] J. Edmonds, *Paths, Trees and Flowers,* Canadian Journal of
    Mathematics, Vol. 17, pp. 449–467 1965.

[4] M. R. Garey and D. S. Johnson, *Computers and Intractabil-
    ity: A Guide to the Theory of $\mathcal{NP}$-Completeness*, W. H. Free-
    man and Co., 1979.

[5] M. Gelastou, M. Mavronicolas, V. Papadopoulou, A. Philip-
    pou and P. Spirakis, "The Power of the Defender", *CD-ROM
    Proceedings of the 2nd International Workshop on Incentive-
    Based Computing*, 2006.

[6] E. Koutsoupias and C. H. Papadimitriou, "Worst-Case Equi-
    libria," *Proceedings of the 16th International Symposium
    on Theoretical Aspects of Computer Science,* pp. 404–413,
    Vol. 1563, LNCS, 1999.

[7] A. S. LaPaugh and C. H. Papadimitriou, "The Even-Path
    Problem for Graphs and Digraphs," *Networks*, Vol. 14 No.
    4, pp. 507–513, 1984.

[8] T. Markham and C. Payne, "Security at the Network Edge:
    A Distributed Firewall Architecture," *Proceedings of the 2nd
    DARPA Information Survivability Conference and Exposi-
    tion,* Vol. 1, pp. 279-286, 2001.

[9] M. Mavronicolas, L. Michael, V. Papadopoulou, A. Philip-
    pou and P. Spirakis, "The Price of Defense," *Proceedings of
    the 31st International Symposium on Mathematical Founda-
    tions of Computer Science,* pp. 717–728, Vol. 4162, LNCS,
    2006.

[10] M. Mavronicolas, V. Papadopoulou, G. Persiano, A. Philip-
    pou and P. Spirakis, "The Price of Defense and Fractional
    Matchings," *Proceedings of the 8th International Confer-
    ence on Distributed Computing and Networking,* pp. 115–
    126, Vol. 4308, LNCS, 2006.

[11] M. Mavronicolas, V. Papadopoulou, A. Philippou and P. Spi-
    rakis, "A Graph-Theoretic Network Security Game," *Pro-
    ceedings of the 1st International Workshop on Internet and
    Network Economics,* pp. 969–978, Vol. 3828, LNCS, 2005.

[12] M. Mavronicolas, V. Papadopoulou, A. Philippou and P. Spi-
    rakis, "A Network Game with Attacker and Protector En-
    tities," *Proceedings of the 16th Annual International Sym-
    posium on Algorithms and Computation,* pp. 288–297,
    Vol. 3827, LNCS, 2005.

[13] J. F. Nash, "Equilibrium Points in N-Person Games", *Pro-
    ceedings of the National Academy of Sciences of the United
    States of America,* pp. 48–49, Vol. 36, 1950.

[14] J. F. Nash, "Non-Cooperative Games," *Annals of Mathemat-
    ics,* Vol. 54, pp. 286–295, 1951.

[15] E. R. Scheinerman and D. H. Ullman, *Fractional Graph The-
    ory,* John Wiley & Sons, 1997.

[16] D. B. West, *Introduction to Graph Theory,* Prentice Hall, sec-
    ond edition, 2001.