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 popular *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 (locally) 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 (inverse) proportion of attackers caught by the 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. More specifically, we identify graph-theoretic *thresholds* for the number of defenders that determine the possibility of optimizing Defense-Ratio. In this vein, we obtain, through an extensive combinatorial analysis of Nash equilibria, a comprehensive collection of *trade-off* results:

- When the number of defenders is either sufficiently small or sufficiently large, there *are* cases where the Defense-Ratio can be optimized. The corresponding optimization problem is then computationally tractable when the number of defenders is large. The problem becomes $\mathcal{NP}$-complete when the number of defenders is small; the intractability is shown by reduction from a previously unconsidered combinatorial problem in *Fractional Graph Theory.*
- Perhaps paradoxically, there is a middle range of values for the number of defenders where optimizing the Defense-Ratio is *impossible* (in every case).
- It is *always* possible to apply a simple and efficient *replication* technique on the defenders in order to achieve an arbitrarily good approximation to a *Defense-Optimal* Nash equilibrium.

Due to the space constraints, almost all technical proofs are shifted to the Appendix. The Appendix may be consulted at the discretion of the Program Committee.

*This work has been partially supported by the IST Program of the European Union under contract numbers IST-2004-001907 (DELIS) and 15964 (AEOLUS).

†Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus. Currently visiting Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, 33102 Paderborn, Germany. Email: `mavronic@cs.ucy.ac.cy`

‡Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, 33102 Paderborn, Germany. Email: `bm@upb.de`

§Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus. Email: `viki@cs.ucy.ac.cy`

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 *servers* (*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 a challenging task 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 antivirus 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 to adopt 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. To that end, the owner has subcontracted the task to a set of independent, deployable agents. Clearly, each such agent is selfish, trying to optimize the protection he offers in order to be paid more.

We materialize 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 to the number of attackers caught. We assume that there are ν attackers and μ 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 ν, 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].

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 economic investments 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*), 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 μ smaller, *vertex-disjoint* Fractional Perfect Matchings so that the total weight (inherited from the Fractional Perfect Matching) in each partite 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 (cf. 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 μ 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 7.2 and 7.3).
- For the the middle range $\frac{|V|}{2} < \mu < \beta'(G)$ of values of μ, 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 μ, 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 replication 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, since it reduces to recognizing a graph with a Fractional Perfect Matching.)

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 $\mathsf{Vertices}_G(F) = \{v \in V \mid (u, v) \in F \text{ for some } u \in V\}$. 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 Perfect Matching f, denote as $E_f = \{e \in E \mid f(e) > 0\}$. A *Fractional Perfect Matching* is a Fractional Matching f with $\sum_{e|v\in e} f(e) = 1$ for all vertices $v \in V$. We observe a simple property of Fractional Perfect Matchings:

Lemma 2.1 *For a Fractional Perfect Matching f, the graph $G(E_f)$ has no pendant edges.*

Lemma 2.1 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 first prove:

Proposition 2.2 *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.*

To prove Proposition 2.2, we present and analyze the algorithm EliminateEvenCycles:

Algorithm EliminateEvenCycles
INPUT: A graph $G(V, E)$ and a Fractional Matching f for G.
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 $f := f'$.

Proposition 2.3 *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 any odd cycle in the graph $G(E_{f'})$ is a component of it.*

To prove Proposition 2.3, we present and analyze the algorithm IsolateOddCycles:

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 $g : E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \rightarrow \left\{+1, -1, +\frac{1}{2}, -\frac{1}{2}\right\}$ so that

$$g(e) = \begin{cases} +1 \text{ or } -1 \text{ (alternatily, starting with } -1), & \text{if } e = (v_i, v_{i+1}) \text{ with } 0 \leq i \leq l-1 \\ +\frac{1}{2} \text{ or } -\frac{1}{2} \text{ (alternatily, starting with } +\frac{1}{2}), & \text{if } e \in E(\mathcal{C}) \\ +\frac{1}{2} \text{ or } -\frac{1}{2} \text{ (alternatily, starting with } -\mathtt{sgn}(g(v_{l-1}, v_l)), & \text{if } e = (v_i, v_{i+1}) \text{ with } l \leq i \leq r-1 \end{cases}$$

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

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

(2/e) For each edge $e \in E$, set

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

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

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.

3 Framework

Basics. 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 ν ***attackers*** A_i and $\mathcal{N}_\mathsf{D}$ contains μ ***defenders*** D_i.
- The ***strategy set*** S_{A_i} of attacker A_i is V, and the ***strategy set*** S_{D_i} of defender 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 A_i is a function $\mathsf{IP}_{\mathsf{A}_i} : S \rightarrow \{0, 1\}$ 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 A_i chooses vertex v, he receives 0 if it is caught by a defender; otherwise, he receives 1.
  - The ***Individual Profit*** of defender D_j is a function $\mathsf{IP}_{\mathsf{D}_j} : S \rightarrow \mathbb{R}$ with

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

  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 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 D_i on vertex v in the profile $\mathbf{s}$ is given by $\mathsf{Prop}_\mathbf{s}(\mathsf{D}_i, v) = \frac{1}{|\mathsf{defenders}_\mathbf{s}(v)|}$.

$\boxed{\textbf{Pure Nash equilibria.}}$ 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 that G is ***pure*** if there is a pure Nash equilibrium for the strategic game $\Pi_{\nu,\mu}(G)$.

$\boxed{\textbf{Mixed profiles.}}$ 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 A_i chooses vertex v, and $s_{\mathsf{D}_j}(e)$ is the probability that defender D_j chooses edge e.

Fix now a mixed profile $\mathbf{s}$. The ***support*** of player $i \in \mathcal{N}$ in the profile $\mathbf{s}$, 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 $\mathsf{Supports}_{\mathbf{s}}(\mathsf{A}) = \bigcup_{\mathsf{A}_i \in \mathcal{N}_{\mathsf{A}}} \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$; denote as $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D}) = \bigcup_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. A vertex v is ***multidefender*** in the profile $\mathbf{s}$ if $|\{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}} \mid$ there is an edge $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$ 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 $\mathbf{s}$ is ***unidefender*** if every vertex $v \in V$ is unidefender in $\mathbf{s}$.

A mixed profile $\mathbf{s}$ 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 $\mathsf{Hit}(\mathsf{D}_i, v)$ 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, $\mathsf{Hit}(v) = \bigcup_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathsf{Hit}(\mathsf{D}_i, v)$. Hence, by the *Principle of Inclusion-Exclusion,* $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = \sum_{j \in [\mu]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_{\mathsf{D}} | |\mathcal{D}| = j} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}} \left(\mathsf{Hit}(\mathsf{D}_k, v) \right)$. From this expression, we immediately observe:

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

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

Lemma 3.2 *For a profile* $\mathbf{s}$, $\sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \leq 2\,\mu$. *(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 $\mathbf{s}$. Denote as $\mathsf{VP}_{\mathbf{s}}(v)$ the expected number of attackers choosing vertex v (according to $\mathbf{s}$); so, $\mathsf{VP}_{\mathbf{s}}(v) = \sum_{i \in [\nu]} s_{\mathsf{A}_i}(v)$. 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* $\mathbf{s}$, $\mathsf{MinHit}_{\mathbf{s}} \leq \dfrac{2\mu}{|V|}$.

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.

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

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

$$\begin{aligned} \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) &= \sum_{j \in [m]} \frac{1}{j} \sum_{\mathcal{D} \subseteq \mathcal{N}_{\mathsf{D}} \setminus \{\mathsf{D}_i\} | |\mathcal{D}| = j-1} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)) \prod_{\mathsf{D}_k \notin \mathcal{D} \cup \{\mathsf{D}_i\}} (1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v))) \\ &= \sum_{j \in [m]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_{\mathsf{D}} \setminus \{\mathsf{D}_i\} | |\mathcal{D}| = j-1} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)), \end{aligned}$$

  where Lemma A.8 was used for the last equality.

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

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

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

Nash equilibria. 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 Existence Theorem of Nash [13, 14], $\Pi_{\nu,\mu}(G)$ has at least one Nash equilibrium.) Clearly, in a Nash equilibrium $\mathbf{s}$, for each attacker A_i, $\mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i, v)$ is *constant* over all vertices $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$; for each defender D_i, $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i, e)$ is *constant* over all edges $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. It follows that in a Nash equilibrium $\mathbf{s}$, for each attacker A_i, $\mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i) = 1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v))$ for any vertex $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$; for each defender D_i, $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i) = \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{s}}(u) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v)$, for any edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. Hence, for each attacker A_i, $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v))$ is *constant* over all vertices $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$.

Some notation. Set $\mathsf{Edges}_{\mathbf{s}}(v) = \{(u, v) \in E \mid (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D})\}$; so, $\mathsf{Edges}_{\mathbf{s}}(v)$ contains all edges incident to v that are included in the union of supports of the defenders. For a vertex set $U \subseteq V$, set $\mathsf{Edges}_{\mathbf{s}}(U) = \{(u, v) \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{D}) \mid u \in U\}$; so, $\mathsf{Edges}_{\mathbf{s}}(U)$ contains all edges incident to a vertex in U that are included in the union of supports of the defenders. For an edge set $F \subseteq E$, set $\mathsf{Vertices}_{\mathbf{s}}(F) = \{u \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A}) \mid (u, v) \in F \text{ for some } v \in V\}$.

Some special profiles. A profile $\mathbf{s}$ is ***uniform*** if each player uses a *uniform* probability distribution on its support; so, for each attacker A_i, for each vertex $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$, $s_{\mathsf{A}_i}(v) = \frac{1}{|\mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)|}$, and for each defender D_i, for each edge $e \in E$, $s_{\mathsf{D}_i}(e) = \frac{1}{|\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)|}$. A profile $\mathbf{s}$ is ***attacker symmetric*** (resp., ***defender symmetric***) if for all pairs of attackers A_i and A_k (resp., all pairs of defenders D_i and 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 A_i (resp., for each defender D_i), $\mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i) = V$ (resp., $\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) = V$).

Defense-Ratio. The ***Defense-Ratio*** $\mathsf{DR}_{\mathbf{s}}$ of a Nash equilibrium $\mathbf{s}$ is the ratio of the *optimal* gain ν of the defenders over their expected gain in $\mathbf{s}$; so, $\mathsf{DR}_{\mathbf{s}} = \frac{\nu}{\sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i)}$.

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* $\mathbf{s}$ *is a Nash equilibrium if and only if the following conditions hold:*

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

(2) *For each defender* D_i*, for each edge* $(u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$*,*

$$\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v)\,\mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u)\,\mathsf{VP}_{\mathbf{s}}(u) \quad = \quad \max_{(u',v') \in E} \{\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v')\,\mathsf{VP}_{\mathbf{s}}(v') + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u')\,\mathsf{VP}_{\mathbf{s}}(u')\} \, .$$

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 $e \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$, $\mathsf{VP}_{\mathbf{s}}(e) = \max_{e' \in E} \{\mathsf{VP}_{\mathbf{s}}(e')\}$. We continue to prove:

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

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

Corollary 4.3 *In a Nash equilibrium* $\mathbf{s}$*,* $\mathsf{DR}_{\mathbf{s}} = \frac{1}{\mathsf{MinHit}_{\mathbf{s}}}$*.*

Corollary 4.3 implies that $\mathsf{DR}_{\mathbf{s}} \geq 1$. Furthermore, we observe:

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

We are now ready to provide a significant definition:

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

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*** if there is a Defense-Optimal Nash equilibrium for the strategic game $\Pi_{\nu,\mu}(G)$. We continue to prove:

Proposition 4.5 *In a Nash equilibrium* $\mathbf{s}$*,* $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ *is an Edge Cover.*

Proposition 4.6 *In a Nash equilibrium* $\mathbf{s}$*,* $\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})$ *is a Vertex Cover of* $G(\mathsf{Supports}_{\mathbf{s}}(\mathsf{D}))$*.*
We use Propositions 4.5 and 4.6 to prove:

Proposition 4.7 (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 4.8 *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 $\mathbf{s}$ has $\mathsf{DR}_{\mathbf{s}} = \max\left\{1, \frac{|V|}{2\,\mu}\right\} = \frac{|V|}{2\mu}$. 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* μ***-partitionable*** *if the edge set* E_f *can be partitioned into* μ *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 that for a Fractional Perfect Matching, which can be solved in polynomial time [2]. We observe a preliminary property of μ-partitionable Fractional Perfect Matchings:

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

We now 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* μ*-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,* μ *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]:

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

Note that the restriction to instances for which μ divides $|V|$ is motivated from Lemma 5.2; it is made to restrict to the set of interesting instances. We use Propositions 2.2 and 2.3 to prove:

Proposition 5.5 *Assume that G contains a μ-partitionable Fractional Perfect Matching f. Then, it also has a μ-partitionable Fractional Perfect Matching f' such that $G(E_{f'})$ consists only of either 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 μ-partitionable Fractional Perfect Matching if and only if E can be partitioned into a collection $E_1, \cdots, E_\mu$ of μ 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 μ-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 a triangle graph?

This problem is known to be $\mathcal{NP}$-complete [4, GT11, attribution to (personal communication with) Scheafer]. To prove that μ-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 μ-PARTITIONABLE FRACTIONAL PERFECT MATCHING, we get that μ-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* $\mathsf{Support}_{\mathbf{s}}(\mathsf{D})$ *is a Perfect Matching if and only if $2\,\mu$ divides $|V|$.*

Note that Corollary 5.4 applies to *all* graphs, while Proposition D.9 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 D.9 *unless* μ is odd. (This is because 2 divides $|V|$ and μ divides $|V|$ imply together that 2μ divides $|V|$ exactly when μ 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 $\mathsf{DR}_{\mathbf{s}} = \max \left\{ 1, \frac{|V|}{2\,\mu} \right\} = 1$. We show:

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

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.

Definition 7.1 (Balanced profile) *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 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 observe that a defender-pure balanced profile goes half the way towards a Nash equilibrium:

Lemma 7.1 *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 7.2 *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 7.2, 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 ν and μ 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 7.2 allows for a (Defense-Optimal) *Pure* Nash equilibrium.

Theorem 7.3 *Assume that* $\mu \geq \beta'(G)$ *and* $2\,\mu$ *divides* ν*. Then,* G *admits a Defense-Optimal, Pure Nash equilibrium, which can be computed in polynomial time.*

To prove Theorem 7.3, we present a polynomial time algorithm Defender-Pure&Balanced to compute a Pure Balanced Nash equilibrium:

Algorithm $\mathsf{Pure\&BalancedNE}$
INPUT: A graph $G(V,E)$ and integers μ and ν such that $\beta'(G) \leq \mu$ 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 $s_{\mathsf{D}_i} := (v_{i \bmod \beta'(G)}, u_{i \bmod \beta'(G)})$.

(3) Arbitrarily assign pure strategies to the attackers so that for each vertex $v \in V$, $\mathsf{VP}_{\mathbf{s}}(v) = \mathsf{defenders}_{\mathbf{s}}(v) \cdot \frac{\nu}{2\mu}$.

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* $\mathsf{MinHit}_{\mathbf{t}} = 1 - (1 - \mathsf{MinHit}_{\mathbf{s}})^{\mu}$.

By Corollary 4.3, Proposition 8.1 immediately implies:

Theorem 8.2 (From Single Defender to Symmetric Defenders) *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* $\Pi_{\nu,\mu}(G)$ *with Defense-Ratio* $\mathsf{DR}_{\mathbf{t}} = \frac{1}{1-\left(1-\frac{1}{\mathsf{DR}_{\mathbf{s}}}\right)^{\mu}}$.

It is simple to see that in the setting of Theorem 8.2, $\mathsf{DR}_{\mathbf{t}} \geq \frac{\mathsf{DR}_{\mathbf{s}}}{\mu}$. (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 $\mathsf{DR}_{\mathbf{s}} = \frac{|V|}{2}$.)

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 Computer Science,* pp. 295–304, 2004.

[2] J.-M. Bourjolly and W. R. Pulleyblank, "König-Egerváry Graphs, 2-Bicritical Graphs and Fractional Matchings," *Discrete 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 Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Co., 1979.

[5] 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*, 2006.

[6] E. Koutsoupias and C. H. Papadimitriou, "Worst-Case Equilibria," *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 Exposition,* Vol. 1, pp. 279-286, 2001.

[9] 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.

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

[11] 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.

[12] 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. Extended version accepted to *Algorithmica.*

[13] J. F. Nash, "Equilibrium Points in N-Person Games", *Proceedings 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 Mathematics,* Vol. 54, pp. 286–295, 1951.

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

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

A Proofs from Section 2:

A.1 Lemma 2.1:

Assume, by way of contradiction, that $G(E_f)$ has a pendant edge (u, v) with $d_{G(E_f)}(u) = 1$ and $d_{G(E_f)}(v) > 1$. Since f is a Fractional Perfect Matching, $\sum_{e|u\in e} f(e) = 1$ and $\sum_{e|v\in e} f(e) = 1$. By assumption on u, the first condition implies that $f((u, v)) = 1$. By definition of $G(E_f)$, the second condition implies that $f((u, v)) < 1$. A contradiction.

A.2 Proposition 2.2 (Continued):

We first prove:

Lemma A.1 *For each loop iteration of* EliminateEvenCycles*, upon completion of Step* (3)*, f' is equivalent to f.*

Proof: Take any vertex $v \in V$. Then,

$$\begin{aligned}\sum_{e|v\in e} f'(e) &= \textstyle\sum_{e\in E(\mathcal{C})|v\in e} f'(e) + \sum_{e\notin E(\mathcal{C})|v\in e} f'(e)\\ &= \textstyle\sum_{e\in E(\mathcal{C})|v\in e} f'(e) + \sum_{e\notin E(\mathcal{C})|v\in e} f(e) \quad \text{(by Step 3).}\end{aligned}$$

Note that if there are no edges $e \in E(\mathcal{C})$ such that $v \in e$, then $\sum_{e\in E(\mathcal{C})|v\in e} f'(e) = \sum_{e\in E(\mathcal{C})|v\in e} f(e) = 0$, and we are done. Otherwise, there are two edges e_1, $e_2 \in E(\mathcal{C})$ such that $v \in e_1$ and $v \in e_2$. Hence,

$$\begin{aligned}\sum_{e|v\in e} f'(e) &= f'(e_1) + f'(e_2) + \sum_{e\notin E(\mathcal{C})|v\in e} f(e)\\ &= f(e_1) + g(e_1)\cdot f(e_0) + f(e_2) + g(e_2)\cdot f(e_0) + \sum_{e\notin E(\mathcal{C})|v\in e} f(e) && \text{(by Step (3))}\\ &= f(e_1) + f(e_2) + ((g(e_1) + g(e_2))\cdot f(e_0) + \sum_{e\notin E(\mathcal{C})|v\in e} f(e)\\ &= f(e_1) + f(e_2) + \sum_{e\notin E(\mathcal{C})|v\in e} f(e) && \text{(since } g(e_1) + g(e_2) = 0 \text{ (by Step (2)))}\\ &= \sum_{e|v\in e} f(e),\end{aligned}$$

and we are done. ■

Since f is a Fractional Matching, Lemma A.1 immediately implies:

Corollary A.2 *f' is a Fractional Matching.*

We continue to prove:

Lemma A.3 *For each loop iteration of* EliminateEvenCycles*, upon completion of Step* (3)*, $f' \subset f$ and the even cycle $\mathcal{C}$ is eliminated from $G(E_{f'})$.*

Proof: We proceed by case analysis.

- Consider first an edge $e \notin E(\mathcal{C})$. By Step (3), $f'(e) = f(e)$. So, $e \in E_{f'}$ if and only if $e \in E_f$.

- Consider now edge $e_0 \in E(\mathcal{C})$; clearly, $e_o \in E_f$. Then,

$$\begin{array}{rcl} f'(e_0) = & f(e_0) + g(e_0) \cdot f(e_0) & \text{(by Step (3))} \\ = & f(e_0) - f(e_0) & \text{(by Step (2))} \\ = & 0. & \end{array}$$

  It follows that $e \notin E_{f'}$. Hence, cycle $\mathcal{C}$ is eliminated from $G(E_{f'})$.

- Consider finally an edge $e \in E(\mathcal{C})\backslash\{e_0\}$; clearly, $e_o \in E_f$. Then,

$$\begin{array}{rcl} f'(e) = & f(e) + g(e) \cdot f(e_0) & \text{(by Step (3))} \\ \geq & f(e) - f(e_0) & \text{(by Step (2))} \\ \geq & 0 & \text{(by Step (1))}. \end{array}$$

  It follows that if $e \in E_{f'}$ then $e \in E_f$.

It follows that $f' \subset f$, as needed. ■

Lemmas A.1 and A.3 together imply that the output f' of algorithm EliminateEvenCycles is a Fractional Matching which is equivalent to and strictly contained in f, and which contains no even cycles.

Recall that an even cycle for an undirected graph can be found in polynomial time [7]. Lemma A.3 implies that there at most $|E|$ loop iterations; each loop execution takes $O(|E|)$ time. It follows that algorithm EliminateEvenCycles runs in polynomial time.

A.3 Proposition 2.3 (Continued):

Since f is a Fractional Perfect Matching, Lemma 2.1 implies that $G(E_f)$ has no pendant edges. Thus, we immediately obtain that Step (2/a) is necessarily successful:

Lemma A.4 *There is a DFS path $v_l, \cdots, v_r$ with $v_r = v_l$ for some l, $1 \leq l < r - 1$.*

Note that since $G(E_f)$ has no even cycle, it follows that the cycle $v_l, \cdots, v_r$ is odd. We continue to prove a preliminary property of the algorithm IsolateOddCycles:

Lemma A.5 *The DFS path $v_1, v_2, \cdots, v_r$ is disjoint from $\mathcal{C}$.*

Proof: By way of contradiction, assume otherwise. Then, there is some vertex v_k, $1 \leq k \leq r$, such that $v_k \in V(\mathcal{C})$. Since $\mathcal{C}$ has odd length, the vertices v_0 and v_k partition $\mathcal{C}$ into two paths $\mathcal{C}_1$ and $\mathcal{C}_2$ of odd and even length, respectively. Consider the concatenations of the path $v_1, \cdots, v_k$ with $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively; each of them is a cycle in $G(E_f)$. Clearly, one of these cycles has even length. A contradiction. ■

Lemma A.5 implies that the function g is well-defined. We prove:

Lemma A.6 *For each loop iteration of* IsolateOddCycles*, upon completion of* Step (2/e)*, f' is equivalent to f.*

Proof: Take any vertex $v \in V$. Then,

$$\begin{array}{rcll} \sum_{e \mid v \in e} f'(e) & = & \sum_{e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \mid v \in e} f'(e) + \sum_{e \notin E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \mid v \in e} f'(e) & \\ & = & \sum_{e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \mid v \in e} f'(e) + \sum_{e \notin E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \leq i \leq r-1\} \mid v \in e} f(e) & \text{(by Step (2/e))}. \end{array}$$

Note that if there are no edges $e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \le i \le r-1\}$ such that $v \in e$, then $\sum_{e \in E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} f'(e) = \sum_{e \in E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} f(e) = 0$, and we are done. Otherwise,

$$\begin{aligned}
\sum_{e|v\in e} f'(e) &= \sum_{e \in E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} f'(e) + \sum_{e \notin E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} f(e) \\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} (f(e) + g(e) \cdot \lambda) + \sum_{e \notin E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} f(e) \\
&= \sum_{e \,|\, v \in e} f(e) \qquad (\text{since } \textstyle\sum_{e \in E(\mathcal{C}) \cup \{(v_i,v_{i+1}) \,|\, 0 \le i \le r-1\} \,|\, v \in e} g(e) = 0 \text{ (by Step (2/b))}),
\end{aligned}$$

and we are done. ■

We continue to prove:

Lemma A.7 *For each loop iteration of* IsolateOddCycles*, upon completion of* Step (2/e)*, $f' \subset f$, and the odd cycle $\mathcal{C}$ (which is not a component in $G(E_f)$) is eliminated from $G(E_{f'})$.*

Proof: We proceed by case analysis.

- Consider first $e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \le i \le r-1\}$. By Step (2/e), $f'(e) = f(e)$. So, $e \in E_{f'}$ if and only if $e \in E_f$.
- Consider now edge $e' \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \le i \le r-1\}$ that realizes $\min\{\min_{0 \le i \le l-1} f((v_i, v_{i+1})),\ 2\min_{e \in E(\mathcal{C})} f(e),\ 2\min_{l \le i \le r-1} f((v_i, v_{i+1}))\}$. Clearly, $e' \in E_f$. Then,

$$\begin{aligned}
f'(e') &= f(e') + g(e') \cdot f(e') && \text{(by Step (2/d))} \\
&= f(e') - f(e') && \text{(by Step (2/c))} \\
&= 0.
\end{aligned}$$

  It follows that $e' \notin E_{f'}$.
- Consider finally an edge $e \in E(\mathcal{C}) \cup \{(v_i, v_{i+1}) \mid 0 \le i \le r-1\} \backslash \{e'\}$; clearly, $e \in E_f$. Then,

$$\begin{aligned}
f'(e) &= f(e) + g(e) \cdot f(e') \qquad \text{(by Step (2/e))} \\
&\ge f(e) + \begin{cases} -1 \cdot \min_{0 \le i \le l-1} f((v_i, v_{i+1})), & \text{if } e = (v_i, v_{i+1}) \text{ with } 0 \le i \le l-1 \\ -\frac{1}{2} \cdot 2\min_{e \in E(\mathcal{C})} f(e), & \text{if } e \in E(\mathcal{C}) \\ -\frac{1}{2} \cdot 2\min_{l \le i \le r-1} f((v_i, v_{i+1})), & \text{if } e = (v_i, v_{i+1}) \text{ with } l \le i \le r-1 \end{cases} \\
&\qquad \text{(by Steps (2/b), (2/c) and (2/d))} \\
&\ge 0.
\end{aligned}$$

  It follows that if $e \in E_{f'}$ then $e \in E_f$.

It follows that $f' \subset f$. It remains to show that the odd cycle $\mathcal{C}$ (which is not a component in $G(E_f)$) is eliminated from $G(E_{f'})$. Recall that $e' \notin E_{f'}$. There are three possible cases to consider:

- Assume first that $e' \in E(\mathcal{C})$. Then, clearly, cycle $\mathcal{C}$ is eliminated from $G(E_{f'})$.
- Assume first that $e' \in \{(v_i, v_{i+1}) \mid 0 \le i \le l-1\}$. Since f' is equivalent to f (by Lemma A.6) and f is a Fractional Perfect Matching, it follows that f' is also a Fractional Perfect Matching. Lemma 2.1 implies that f' has no pendant edges. Since $e' \notin E_{f'}$, this implies that either $(v_0, v_1) \notin E_{f'}$ or (in case (v_0, v_1) is an isolated edge in $G(E_{f'})$) both edges in $E(\mathcal{C})$ containing v_0 are eliminated from $E_{f'}$.

- Assume finally that $e' \in \{(v_i, v_{i+1}) \mid l \leq i \leq r-1\}$. Since f' is equivalent to f (by Lemma A.6) and f is a Fractional Perfect Matching, it follows that f' is also a Fractional Perfect Matching. Lemma 2.1 implies that f' has no pendant edges. Since $E_{f'}$ has no pendant edges and $e \notin E_{f'}$, it follows that there is some edge $e'' \in \{v_i, v_{i+1} \mid 0 \leq i \leq l-1\}$ such that $e'' \notin E_{f'}$, and we are reduced in the previous case.

The proof is now complete. ■

Lemmas A.6 and A.7 imply together that the output f' of algorithm IsolateOddCycles is a Fractional Perfect Matching which is equivalent to and strictly contained in f, and which contains no odd cycle that is not a component.

Recall that an odd cycle for an undirected graph can be found in polynomial time [7]. Lemma A.7 implies that there at most $|E|$ loop iterations; each loop execution takes $O(|E|)$ time. It follows that algorithm IsolateOddCycles runs in polynomial time.

A.4 A Combinatorial Lemma.

Here, we will prove a combinatorial lemma (of independent interest) that will be useful later. For a *probability* x, with $0 \leq x \leq 1$, we make two *probability literals,* the *positive* literal x and the *negative* literal $\bar{x} = 1 - x$. A *probability product* is a product of probability literals $x_1 \cdots x_n$ for any $n \geq 1$. For any integer $j \leq n$, the probability product $x_1 \cdots x_n$ is j*-positive* if exactly j of its probability literals are positive. For each $j \leq n$, denote as $\mathsf{Pos}_j(x_1, \ldots, x_n)$ the set of all j-positive probability products that use literals made from the probabilities $x_1, \ldots, x_n$.

Lemma A.8 *For each integer $n \geq 2$,*

$$\sum_{j\in[n]} \frac{1}{j} \sum_{x_2\ldots x_n \in \mathsf{Pos}_{j-1}(x_2,\ldots,x_n)} x_2 \cdots x_n \quad = \quad \sum_{j\in[n]} (-1)^{j-1} \cdot \frac{1}{j} \cdot \sum_{x_2\ldots x_j \in \mathsf{Pos}_{j-1}(x_2,\ldots,x_n)} x_2 \cdots x_j \, .$$

Proof: It suffices to establish that each term $x_2 \cdots x_j \in \mathsf{Pos}_{j-1}(x_2, \ldots, x_n)$ in the right-hand side (RHS) appears in the left-hand side (LHS) with the same coefficient.

- Consider first the *constant* term $x_2 \cdots x_j|_{j=1} = 1$ in the RHS. Its coefficient is $(-1)^{1-1} \cdot \frac{1}{1} = 1$.

  In the LHS, the only constant term is the constant term in the sum $\sum_{x_2\cdots x_n \in \mathsf{Pos}_{j-1}(x_2,\ldots,x_n)} x_2 \cdots x_n\Big|_{j=1} = \overline{x}_2 \cdots \overline{x}_n$. Clearly, this constant term is 1 and its coefficient is $\frac{1}{1} = 1$, as needed.

- Consider now any term $x_2 \cdots x_j \in \mathsf{Pos}_{j-1}(x_2, \ldots, x_n)$ from the sum $\sum_{x_2\cdots x_j \in \mathsf{Pos}_{j-1}(x_2,\ldots,x_n)} x_2 \cdots x_j$ in the RHS with $j \geq 2$. Note that all such terms (in the RHS) have the same coefficient, which is $(-1)^{j-1} \cdot \dfrac{1}{j}$ We calculate the coefficient of this term in the LHS.

  Clearly, a k-positive term with $k \geq j$ in the LHS cannot include $x_2 \cdots x_j$ in its expansion. So, we only need to consider contributions from k-positive terms with $0 \leq k \leq j-1$ (in the LHS) to the coefficient of the term $x_2 \cdots x_j$.

  - Note that there are $\binom{j-1}{k}$ ways to choose k positive literals (and therefore $j-1-k$ negative literals) out of the $(j-1)$ literals $x_2, \ldots, x_j$ in order to form a k-positive term that expands to $x_2 \cdots x_j$ multiplied with a coefficient. (The literals $x_{j+1}, \ldots, x_n$ all have to be negative since they do not appear in the product $x_2 \cdots x_j$.)
  - The sign of the resulting k-positive term is $(-1)^{(j-1)-k}$, since each of the $(j-1)-k$ negative literals in it contributes one minus to the sign. (The negative literals $x_{j+1}, \ldots, x_n$ do not contribute to the sign.).
  - The (absolute value of the) coefficient of the resulting k-positive term is $\frac{1}{k+1}$ (since a $(j-1)$-positive term in the LHS is multiplied by $\dfrac{1}{j}$).

So, in total, the coefficient of $x_2 \cdots x_j$ in the LHS is

$$
\begin{aligned}
\sum_{0 \leq k \leq j-1} \binom{j-1}{k} (-1)^{(j-1)-k} \frac{1}{k+1} &= \sum_{0 \leq k' \leq j-1} \binom{j-1}{k'} (-1)^{k'} \frac{1}{j-k'} \\
&= \sum_{0 \leq k \leq j-1} \binom{j-1}{k} (-1)^{k} \frac{1}{j-k} \\
&= \frac{1}{j} \sum_{0 \leq k \leq j-1} \binom{j}{k} (-1)^{k} \\
&= \frac{1}{j} \left(\sum_{0 \leq k \leq j} \binom{j}{k} (-1)^{k} - \binom{j}{j} (-1)^{j} \right) \\
&= \frac{1}{j} \left(0 + (-1)^{j-1}\right) \\
&= \frac{1}{j} (-1)^{j-1} \,,
\end{aligned}
$$

and the claim follows.

■

B Proofs from Section 3:

B.1 Lemma 3.2:

Assume first that there is no multidefender vertex. So, for each vertex v, there is at most one defender $\mathsf{D}(v)$ such that $v \in e$ for some edge $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}(v))$. Then,

$$
\begin{aligned}
& \sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \\
= & \quad \textstyle\sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{D}(v), v) \\
= & \quad \textstyle\sum_{v \in V} \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} s_{\mathsf{D}(v)}(e) \\
= & \quad \textstyle\sum_{e \in E} \sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} (2\, s_{\mathsf{D}_i}(e)) \qquad \text{(by changing the order of summation)} \\
= & \quad 2 \textstyle\sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \sum_{e \in E} s_{\mathsf{D}}(e) \\
= & \quad 2\, \mu \,.
\end{aligned}
$$

Assume now that there is a multidefender vertex $v \in V$. Then,

$$
\begin{aligned}
\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) &= \sum_{j \in [\mu]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_{\mathsf{D}} || \mathcal{D}| = j} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}} \left(\mathsf{Hit}_{\mathbf{s}}(\mathsf{D}_k, v)\right) \\
&< \sum_{v \in V} \sum_{\mathsf{D}_k \in \mathcal{N}_{\mathsf{D}}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)) \\
&= \sum_{v \in V} \sum_{\mathsf{D}_k \in \mathcal{N}_{\mathsf{D}}} \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} s_{\mathsf{D}_k}(e) \\
&= \sum_{e \in E} \left(2 \sum_{\mathsf{D}_k \in \mathcal{N}_{\mathsf{D}}} s_{\mathsf{D}_k}(e) \right) \\
&= 2 \sum_{\mathsf{D}_k \in \mathcal{N}_{\mathsf{D}}} \sum_{e \in E} s_{\mathsf{D}_k}(e) \\
&= 2\, \mu \,.
\end{aligned}
$$

B.2 Lemma 3.3:

Assume, by way of contradiction, that $\mathsf{MinHits_s} > \frac{2\mu}{|V|}$. Then, $\sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \geq |V| \cdot \mathsf{MinHits_s} > 2\mu$, a contradiction to Lemma 3.2.

B.3 Lemma 3.4:

Clearly,

$$\begin{aligned}
& \sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) \cdot \mathsf{Prop_s}(\mathsf{D}_i, v) \\
= & \sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) \cdot \left(\sum_{j \in [\mu]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} \setminus \{\mathsf{D}_i\} | |\mathcal{D}| = j-1} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)) \right) \\
= & \sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \sum_{j \in [\mu]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} \setminus \{\mathsf{D}_i\} | \mathsf{D}_i \in \mathcal{U} \wedge |\mathcal{D}| = j-1} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)) \\
= & \sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \sum_{j \in [\mu]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} \setminus \{\mathsf{D}_i\} | |\mathcal{D}| = j} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_k, v)) \,.
\end{aligned}$$

C Proofs from Section 4:

C.1 Proposition 4.1:

Assume first that $\mathbf{s}$ is a Nash equilibrium. To show Condition (1), consider any vertex $v \in \mathsf{Supports_s}(\mathsf{A})$; so, $v \in \mathsf{Support_s}(\mathsf{A}_i)$ for some attacker A_i. Since $\mathbf{s}$ is a Nash equilibrium, $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v'))$ is *constant* over all vertices $v' \in \mathsf{Support_s}(\mathsf{A}_i)$. So, consider any vertex $u \notin \mathsf{Support_s}(\mathsf{A}_i)$. We prove that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u)) \geq \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v))$.

> Assume, by way of contradiction, that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u)) < \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v))$. Construct from $\mathbf{s}$ the mixed profile $\mathbf{t}$ by only changing s_{A_i} to t_{A_i} so that $u \in \mathsf{Support_s}(\mathsf{A}_i)$. Then,
>
> $$\begin{aligned}
> & \mathsf{IP}_t(\mathsf{A}_i) \\
> = & \ 1 - \mathbb{P}_{\mathbf{t}}(\mathsf{Hit}(u)) \quad \text{(since } u \in \mathsf{Support_t}(\mathsf{A}_i)\text{)} \\
> = & \ 1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u)) \quad \text{(by construction)} \\
> > & \ 1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \quad \text{(by assumption)} \\
> = & \ \mathsf{IP_s}(\mathsf{A}_i) \,,
> \end{aligned}$$
>
> which contradicts the fact that $\mathbf{s}$ is a Nash equilibrium.

It follows that for each vertex $v \in \mathsf{Support_s}(\mathsf{A}_i)$, $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \leq \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u))$ for all other vertices $u \in V$.

Consider now any vertex $u \notin \mathsf{Support_s}(\mathsf{A}_i)$ such that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \leq \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u))$. Since $\mathbf{s}$ is a local maximizer of the Expected Individual Profit of each other attacker A_k, it follows that there is no attacker A_k such that $u \in \mathsf{Support_s}(\mathsf{A}_k)$. This implies that $u \notin \mathsf{Supports_s}(\mathsf{A})$. It follows that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = \mathsf{MinHits_s}$, as needed.

For Condition (2), fix any defender D_i and consider any edge $e = (u, v) \in \mathsf{Support_s}(\mathsf{D}_i)$. Recall that

$$\mathsf{IP_s}(\mathsf{D}_i) \quad = \quad \mathsf{Prop_s}(\mathsf{D}_i, v) \cdot \mathsf{VP_s}(v) + \mathsf{Prop_s}(\mathsf{D}_i, u) \cdot \mathsf{VP_s}(u) \,,$$

so that in a Nash equilibrium, the quantity $\mathsf{Prop_s}(\mathsf{D}_i, v') \cdot \mathsf{VP_s}(v') + \mathsf{Prop_s}(\mathsf{D}_i, u') \cdot \mathsf{VP_s}(u')$ is *constant* over all edges $(u', v') \in \mathsf{Support_s}(\mathsf{D}_i)$. So, consider any edge $(u', v') \notin \mathsf{Support_s}(\mathsf{D}_i)$. Assume, by way of contradiction, that

$$\mathsf{Prop_s}(\mathsf{D}_i, v') \cdot \mathsf{VP_s}(v') + \mathsf{Prop_s}(\mathsf{D}_i, u') \cdot \mathsf{VP_s}(u') \quad > \quad \mathsf{Prop_s}(\mathsf{D}_i, v) \cdot \mathsf{VP_s}(v) + \mathsf{Prop_s}(\mathsf{D}_i, u) \cdot \mathsf{VP_s}(u) \,.$$

Construct from $\mathbf{s}$ the mixed profile $\mathbf{t}$ by only changing s_{D_i} to t_{D_i} so that $e' \in \mathsf{Support}_{\mathbf{t}}(\mathsf{D}_i)$. Then,

$$\begin{array}{lll} & \mathsf{IP}_{\mathbf{t}}(\mathsf{D}_i) & \\ = & \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v') \cdot \mathsf{VP}_{\mathbf{t}}(v') + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u') \cdot \mathsf{VP}_{\mathbf{t}}(u') & (\text{since } e' \in \mathsf{Support}_{\mathbf{t}}(\mathsf{D}_i)) \\ = & \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') \cdot \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') \cdot \mathsf{VP}_{\mathbf{s}}(u') & (\text{by construction of } \mathbf{t} \text{ from } \mathbf{s}) \\ > & \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{s}}(u) & (\text{by assumption}) \\ = & \mathsf{IP}_{\mathbf{t}}(\mathsf{D}_i) & (\text{since } e \in \mathsf{Support}_{\mathbf{t}}(\mathsf{D}_i))\,, \end{array}$$

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

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

- Consider first any attacker A_i. Then, for any vertex $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A}_i)$, for any other vertex $u \in V$,

$$\begin{array}{lll} & \mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i, v) & \\ = & 1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) & (\text{since } v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)) \\ \geq & 1 - \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(u))(\text{by Condition } (1)) & \\ = & \mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i, u)\,. & \end{array}$$

- Consider now any defender D_i. Then, for any edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$, for any edge $e' = (u', v') \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$,

$$\begin{array}{ll} & \mathsf{IP}_{\mathbf{s}}(\mathsf{A}_i, v) \\ > & \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{s}}(u) \\ \geq & \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') \cdot \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') \cdot \mathsf{VP}_{\mathbf{s}}(u') \quad (\text{by Condition } (2))\,. \end{array}$$

It follows that $\mathbf{s}$ is a Nash equilibrium.

C.2 Proposition 4.2:

Clearly,

$$\begin{array}{lll} & \sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i) & \\ = & \sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) \cdot \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v) & \\ = & \sum_{v \in V} \left(\sum_{\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) \cdot \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \right) \cdot \mathsf{VP}_{\mathbf{s}}(v) & (\text{changing the order of summation}) \\ = & \sum_{v \in V} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \cdot \mathsf{VP}_{\mathbf{s}}(v) & (\text{by Lemma 3.4}) \\ = & \sum_{v \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A})} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \cdot \mathsf{VP}_{\mathbf{s}}(v) & \\ = & \sum_{v \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A})} \mathsf{MinHit}_{\mathbf{s}} \cdot \mathsf{VP}_{\mathbf{s}}(v) & \\ = & \sum_{v \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A})} \mathsf{VP}_{\mathbf{s}}(v) \cdot \mathsf{MinHit}_{\mathbf{s}} & \\ = & \nu \cdot \mathsf{MinHit}_{\mathbf{s}}\,, & \end{array}$$

as needed.

C.3 Proposition 4.5:

Assume, by way of contradiction, that $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ is *not* an Edge Cover. Then, there is a vertex $v \in V$ such that $v \notin \mathsf{Vertices}(\mathsf{Supports}_{\mathbf{s}}(\mathsf{D}))$. So, $\mathsf{Edges}_{\mathbf{s}}(v) = \emptyset$ and $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = 0$.

Since $\mathbf{s}$ is a local maximizer for the Expected Individual Profit of each attacker A_i, which is at most 1, it follows that attacker A_i chooses some such v with probability 1. It follows that for each edge $e = (u, v) \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$, $\mathsf{VP}_{\mathbf{s}}(e) = 0$. So, for any defender D_i,

$$\begin{aligned}\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i) &= \sum_{e=(u,v)\in\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)} s_{\mathsf{D}_i}(e)\cdot(\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v)\cdot\mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u)\cdot\mathsf{VP}_{\mathbf{s}}(u))\\ &= 0\,.\end{aligned}$$

Since $\mathbf{s}$ is a Nash equilibrium, $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i) > 0$. A contradiction.

C.4 Proposition 4.6:

Assume, by way of contradiction, that $\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})$ is *not* a Vertex Cover of the graph $G(\mathsf{Supports}_{\mathbf{s}}(\mathsf{D}))$. Then, there is some edge $e \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A})$ such that $e \in \mathsf{Edges}_{\mathbf{s}}(\mathsf{Supports}_{\mathbf{s}}(\mathsf{A}))$. So, $\mathsf{Vertices}_{\mathbf{s}}(e) = 0$ and $\mathsf{VP}_{\mathbf{s}}(e) = 0$. Since $\mathbf{s}$ is a local maximizer for the Expected Individual Profit of each defender D_i, it follows that $s_{\mathsf{D}}(e) = 0$. So, $e \notin \mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$. A contradiction.

C.5 Proposition 4.7:

Assume first that $\mu < \beta'(G)$. Consider any pure profile $\mathbf{s}$. Clearly, $|\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})| \leq \mu < \beta'(G)$. Hence, $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ is not an Edge Cover. Hence, Proposition 4.5 implies that $\mathbf{s}$ is not a Nash equilibrium.

Assume now that $\nu < \min_{EC\in\mathcal{EC}(G)} \beta(G(EC))$. Consider any pure profile $\mathbf{s}$. Thus, $|\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})| \leq \nu < \min_{EC\in\mathcal{EC}(G)} \beta(G(EC))$. By Lemma 4.5, $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ is an Edge Cover. It follows that $|\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})| < \beta(G(\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})))$. Hence, $\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})$ is not a Vertex Cover of $G(\mathsf{Supports}_{\mathbf{s}}(\mathsf{D}))$. Proposition 4.6 implies that $\mathbf{s}$ is not a Nash equilibrium.

C.6 Proposition 4.8:

Fix an arbitrary Defender-Pure Nash equilibrium $\mathbf{s}$. Then,

$$\begin{aligned}&\mathsf{DR}_{\mathbf{s}}\\ =\;& \frac{\nu}{\sum_{\mathsf{D}_i\in\mathcal{N}_{\mathsf{D}}}\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i)}\\ =\;& \frac{\nu}{\sum_{\mathsf{D}_i\in\mathcal{N}_{\mathsf{D}}}(\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v_i)\cdot\mathsf{VP}_{\mathbf{s}}(v_i) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u_i)\cdot\mathsf{VP}_{\mathbf{s}}(u_i))} && \text{(by Condition (2) in Proposition 4.1)}\\ =\;& \frac{\nu}{\sum_{\mathsf{D}_i\in\mathcal{N}_{\mathsf{D}}}\left(\frac{\mathsf{VP}_{\mathbf{s}}(v_i)}{\mathsf{defenders}_{\mathbf{s}}(v_i)} + \frac{\mathsf{VP}_{\mathbf{s}}(u_i)}{\mathsf{defenders}_{\mathbf{s}}(u_i)}\right)} && \text{(since } \mathbf{s} \text{ is Defender-Pure)}\\ =\;& \frac{\nu}{\sum_{v\in\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})}\mathsf{defenders}_{\mathbf{s}}(v)\cdot\frac{\mathsf{VP}_{\mathbf{s}}(v)}{\mathsf{defenders}_{\mathbf{s}}(v)}}\\ =\;& \frac{\nu}{\sum_{v\in\mathsf{Supports}_{\mathbf{s}}(\mathsf{A})}\mathsf{VP}_{\mathbf{s}}(v)}\\ =\;& 1\,,\end{aligned}$$

which implies that $\mathbf{s}$ is Defense-Optimal.

D Proofs from Section 5:

D.1 Proposition 5.1:

Assume, by way of contradiction, that there is a Defense-Optimal Nash equilibrium $\mathbf{s}$ for which there is a multidefender vertex. Since $\mathbf{s}$ is Defense-Optimal, Corollary 4.3 implies that $\mathsf{MinHit}_{\mathbf{s}} = \frac{2\mu}{|V|}$. So, $\sum_{v\in V}\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) \geq \sum_{v\in V}\mathsf{MinHit}_{\mathbf{s}} = 2\mu$. Since there is a multidefender vertex, Lemma 3.2 implies that $\sum_{v\in V}\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) < 2\mu$. A contradiction.

D.2 Lemma 5.2:

Consider a μ-Fractional Perfect Matching f. Then, E_f can be partitioned into μ 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 each set E_i is itself a Fractional Perfect Matching; thus, $\sum_{e \in E_i} f(e) = \frac{|V_i|}{2}$. It follows that for each $i \in [\mu]$, $|V_i| = \frac{|V|}{\mu}$. Since $|V_i|$ is an integer, it follows that μ divides $|V|$.

D.3 Theorem 5.3:

Theorem 5.3 will follow from Propositions D.1 and D.4.

Proposition D.1 *Assume that G has a μ-partitionable Fractional Perfect Matching. Then, G is Defense-Optimal.*

Proof: Consider a μ-partitionable Fractional Perfect Matching f and the corresponding edge sets $E_i, \cdots, E_\mu$. Construct a profile $\mathbf{s}$ as follows:

- For each index $i \in [\mu]$, $\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) = E_i$; for each edge $e \in E_i$, $s_{\mathsf{D}_i}(e) = \frac{2\mu}{|V|} \cdot f(e)$.
- $\mathbf{s}$ is Attacker Symmetric Uniform and Attacker Fully Mixed. So, for each attacker $\mathsf{A}_i \in \mathcal{N}_{\mathsf{A}}$, for each vertex $v \in V$, $s_{\mathsf{A}_i}(v) = \frac{1}{|V|}$; thus, for each vertex $v \in V$, $\mathsf{VP}_{\mathbf{s}}(v) = \frac{\nu}{|V|}$.

To show that $\mathbf{s}$ is a profile, we prove:

Claim D.2 *For each defender $\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}$, s_{D_i} is a probability distribution, so that $\mathbf{s}$ is a profile.*

Proof: Clearly,

$$\begin{array}{rll} & \displaystyle\sum_{e \in E} s_{\mathsf{D}_i}(e) & \\ = & \sum_{e \in E_i} s_{\mathsf{D}_i}(e) & (\text{since } \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) = E_i) \\ = & \sum_{e \in E_i} \frac{2\mu}{|V|} \cdot f(e) & (\text{by construction}) \\ = & \frac{2\mu}{|V|} \sum_{e \in E_i} f(e) & \\ = & 1, & (\text{since } f \text{ is a } \mu\text{-partitionable Fractional Perfect Matching}) \end{array}$$

which implies that $\mathbf{s}$ is a profile. ■

We continue to prove that $\mathbf{s}$ is a Nash equilibrium. We will establish Conditions (1) and (2) in the characterization of Nash equilibria (Proposition 4.1).

- For condition (1), fix any vertex $v \in V$. Since E_f is an Edge Cover and f is a μ-partitionable Fractional Perfect Matching, there is exactly one edge set $E_i \subseteq E_f$ such that $v \in \mathsf{Vertices}(E_i)$. Since $\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) = E_i$, the definition of a μ-partitionable Fractional Perfect Matchings implies that vertex v is unidefender in the profile $\mathbf{s}$. Then, we prove:

  Claim D.3 $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = \frac{2\mu}{|V|}$

Proof:

$$
\begin{array}{lll}
& \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) & \\
= & \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) & \\
= & \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) | v \in e} s_{\mathsf{D}_i}(e) & \\
= & \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) | v \in e} \frac{2\mu}{|V|} \cdot f(e) & \text{(by construction of } \mathbf{s}\text{)} \\
= & \frac{2\mu}{|V|} \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i) | v \in e} f(e) & \\
= & \frac{2\mu}{|V|} \sum_{e \in \mathsf{Edges}_{\mathbf{s}}(v)} f(e) & \text{(since } v \text{ is multidefender in } \mathbf{s}\text{)} \\
= & \frac{2\mu}{|V|} & \text{(since } f \text{ is a Fractional Perfect Matching)}
\end{array}
$$

■

It now follows that Condition (1) holds trivially.

- For Condition (2), consider any defender $\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}$. Fix an edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. Since each vertex is unidefender in $\mathbf{s}$, it follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) = \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) = 1$. It follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{s}}(u) = \mathsf{VP}_{\mathbf{s}}(v) + \mathsf{VP}_{\mathbf{s}}(u) = \frac{2\nu}{|V|}$. On the other hand, fix any edge $e' = (v', u') \notin \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. Since E_f is an Edge Cover, it follows that for the vertex v' (resp., vertex u'), there is a defender D_j such that $v' \in \mathsf{Vertices}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_j))$ (resp., $u' \in \mathsf{Vertices}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_j))$).

  Assume first that $\mathsf{D}_j = \mathsf{D}_i$; since v' is unidefender, it follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') = 1$. (resp., $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') = 1$). Assume now that $\mathsf{D}_j \neq \mathsf{D}_i$; since v' is unidefender, $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') < 1$. (resp., $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') < 1$). So, in all cases, $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') \leq 1$ and $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') \leq 1$. Thus, $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') \cdot \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') \cdot \mathsf{VP}_{\mathbf{s}}(u') \leq \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{VP}_{\mathbf{s}}(u') = \frac{2\nu}{|V|}$. 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, $\mathsf{MinHit}_{\mathbf{s}} = \frac{2\mu}{|V|}$. By Corollary 4.3, it follows that $\mathsf{DR}_{\mathbf{s}} = \frac{|V|}{2\mu}$.

■

Proposition D.4 *Assume that G is Defense-Optimal. Then, G has a μ-partitionable Fractional Perfect Matching.*

Proof: Consider a Defense-Optimal Nash equilibrium $\mathbf{s}$; so, $\mathsf{DR}_{\mathbf{s}} = \frac{|V|}{2\mu}$. Corollary 4.3 implies that $\mathsf{MinHit}_{\mathbf{s}} = \frac{2\mu}{|V|}$. Hence, Lemma 3.2 implies that for each vertex $v \in V$, $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = \frac{2\mu}{|V|}$. Fix a vertex $v \in V$. Since $\mathbf{s}$ is Defense-Optimal, Proposition 5.1 implies that there is at most one defender D_k such that $v \in \mathsf{Vertices}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. By Proposition 4.5, $\mathsf{Support}_{\mathbf{s}}(\mathsf{D})$ is an Edge Cover, so that there is at least one defender D_k such that $v \in \mathsf{Vertices}_{\mathbf{s}}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. It follows that there is exactly one defender D_k such that $v \in \mathsf{Vertices}_{\mathbf{s}}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. This implies that for each edge $e \in E$, there is at most one defender D_k such that $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k)$. Thus, for all pairs of defenders $\mathsf{D}_j, \mathsf{D}_k$, $\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_j) \cap \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k) = \emptyset$. We define now a function $f : E \to \mathbb{R}$; we will then prove that f is a μ-partitionable Fractional Perfect Matching.

- For any edgee $\in E$, if there is a defender D_i such that $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$, then set $f(e) := \frac{|V|}{2\mu} \cdot s_{\mathsf{D}_i}(e)$; else, set $f(e) := 0$. Thus, $E_f = \mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$.

Claim D.5 $\sum_{e \in E} f(e) = \frac{|V|}{2}$.

Proof: Clearly,

$$
\begin{aligned}
& \sum_{e \in E} f(e) \\
= \quad & \textstyle\sum_{\mathsf{D}_i \in \mathcal{N}_\mathsf{D}} \sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)} \frac{|V|}{2\mu} \cdot s_{\mathsf{D}_i} && \text{(by construction of } f\text{)} \\
= \quad & \textstyle\frac{|V|}{2\mu} \sum_{e \in \mathcal{N}_\mathsf{D}} \sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)} \cdot s_{\mathsf{D}_i} \\
= \quad & \frac{|V|}{2\mu} \cdot \mu && \text{(since } \mathbf{s} \text{ is a profile)} \\
= \quad & \frac{|V|}{2}.
\end{aligned}
$$

■

One the other hand, fix any vertex $v \in V$. Recall that there is exactly one defender D_i such that $v \in \mathsf{Vertices}_\mathbf{s}(\mathsf{Support}_\mathbf{s}(\mathsf{D}_i))$. This implies that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = \mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v))$. We prove:

Claim D.6 $\sum_{e \in \mathsf{Edges}_\mathbf{s}(\mathsf{v})} f(e) = 1$.

Proof: Clearly,

$$
\begin{aligned}
& \sum_{e \in \mathsf{Edges}_\mathbf{s}(\mathsf{v})} f(e) \\
= \quad & \textstyle\sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i) | v \in e} f(e) \\
= \quad & \textstyle\sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i) | v \in e} \frac{|V|}{2\mu} \cdot s_{\mathsf{D}_i}(e) && \text{(by construction)} \\
= \quad & \textstyle\frac{|V|}{2\mu} \sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i) | v \in e} s_{\mathsf{D}_i}(e) \\
= \quad & \frac{|V|}{2\mu} \cdot \mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}_i, v)) \\
= \quad & \frac{|V|}{2\mu} \cdot \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) \\
= \quad & \frac{|V|}{2\mu} \cdot \frac{2\mu}{|V|} \\
= \quad & 1.
\end{aligned}
$$

■

It follows that f is a Fractional Perfect Matching. To prove that f is μ-partitionable, define edge sets $E_i := \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)$ for $i \in [\mu]$. Clearly, $\bigcup_{i \in [\mu]} E_i = \bigcup_{i \in [\mu]} \mathsf{Support}_\mathbf{s}(\mathsf{D}_i) = \mathsf{Support}_\mathbf{s}(\mathsf{D}) = E_f$. Since for all pairs of defenders D_j and D_k, $\mathsf{Support}_\mathbf{s}(\mathsf{D}_j) \cap \mathsf{Support}_\mathbf{s}(\mathsf{D}_k) = \emptyset$, it follows that the sets E_i, $i \in [\mu]$, partition the set E_f. In addition, for each edge set E_i, with $i \in [\mu]$, we prove:

Claim D.7 $\sum_{e \in E_i} f(e) = \frac{|V|}{2\mu}$

Proof: Clearly,

$$
\begin{aligned}
& \sum_{e \in E_i} f(e) \\
= \quad & \textstyle\sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)} f(e) && \text{(by construction of the edge sets)} \\
= \quad & \textstyle\sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)} \frac{|V|}{2\mu} \cdot s_{\mathsf{D}_i}(e) && \text{(by construction of } f\text{)} \\
= \quad & \textstyle\frac{|V|}{2\mu} \sum_{e \in \mathsf{Support}_\mathbf{s}(\mathsf{D}_i)} s_{\mathsf{D}_i}(e) \\
= \quad & \frac{|V|}{2\mu}, && \text{(since } \mathbf{s} \text{ is a profile)}
\end{aligned}
$$

■

Claim D.7 implies that f is a μ-partitionable Perfect Matching. The proof is now complete.

■

D.4 Corollary 5.4:

Since G is Defense-Optimal, Theorem 5.3 implies that G has a μ-partitionable Fractional Perfect Matching f. Lemma 5.2 implies now the claim.

D.5 Proposition 5.5:

Apply the polynomial time algorithm EliminateEvenCycles (from Proposition 2.2), to compute from f an equivalent Fractional (Perfect) Matching $f'' \subseteq f$ such that $G(E_{f''})$ has no even cycle. Since f is μ-partitionable and f'' is equivalent to f, it follows that f'' is also μ-partitionable.

Then, apply the polynomial time algorithm IsolateOddCycles (from Proposition 2.3), to compute from f'' an equivalent Fractional Perfect Matching f' such that any odd cycle in $G(E_{f'})$ is a component of it. Since f'' is μ-partitionable and f' is equivalent to it, it follows that f' is also μ-partitionable. Since f' is perfect, Lemma 2.1 implies that $G(E_{f'})$ has no pendant edges. Hence, it follows that $G(E_{f'})$ consists only of single edges and odd cycles.

D.6 Proposition 5.6:

Assume that G has a μ-partitionable Fractional Perfect Matching f. Proposition 5.5 implies the claim. Conversely, assume that E contains a collection $E_1, \cdots E_\mu$ of vertex-disjoint subsets so that $\bigcup_{i\in[\mu]} V_i = V$, each E_i is a collection of single edges and odd cycles, and $|V_i| = \frac{|V|}{\mu}$, for each $i \in [\mu]$. We show that G admits a μ-partitionable Fractional Perfect Matching f: for each edge $e \in E_i$, $i \in [\mu]$, set $f(e) := 1$ if E_i is a single edge; otherwise, set $f(e) := \frac{1}{2}$ (if E_i is an odd cycle). For any other edge $e \in E \backslash \bigcup_{i\in[\mu]} E_i$, set $f(e) := 0$. We show that f is a μ-partitionable Fractional Perfect Matching of G.

Since E_i are vertex-disjoint and $\bigcup_{i\in[\mu]} V_i = V$, it follows by construction that for each vertex $v \in V$, $\sum_{e|v\in e} f(e) = 1$, so that f is a Fractional Perfect Matching of G. Moreover, observe that by construction, $E_f = \bigcup_{i\in[\mu]} E_i$. Since $|V_i| = \frac{|V|}{\mu}$ for each $i \in [\mu]$, it follows that f is also a μ-partitionable Fractional Perfect Matching of G. The proof is now complete.

D.7 Proposition 5.7:

We employ the (trivially polynomial time) identity reduction. Assume first that SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING is positive for G. Proposition 5.6 implies that E contains a collection $E_1, \cdots, E_{\frac{|V|}{3}}$ of vertex-disjoint subsets and corresponding vertex sets $V_1, \cdots, V_{\frac{|V|}{3}}$, so that $\bigcup_{i\in\left[\frac{|V|}{3}\right]} V_i = V$, each E_i is a collection of single edges and odd cycles, and $|V_i| = \frac{|V|}{\frac{|V|}{3}} = 3$ for $i \in \left[\frac{|V|}{3}\right]$. It follows that each E_i is an odd cycle, so that PARTITION INTO TRIANGLES is positive for G.

Assume now that PARTITION INTO TRIANGLES is positive for G. Thus, V can 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 a triangle graph. This partition induces a corresponding partition of E into a collection $E_1, \cdots E_{\frac{|V|}{3}}$ of $\frac{|V|}{3}$ vertex-disjoint subsets, where each E_i is a single triangle. Proposition 5.6 implies that G has a $\frac{|V|}{3}$-partitionable Fractional Perfect Matching, so that SPECIAL PARTITIONABLE FRACTIONAL PERFECT MATCHING is positive for G.

D.8 Theorem 5.9:

Theorem 5.9 will follow from Propositions D.8 and D.9.

Proposition D.8 *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* $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ *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 $\mathsf{VP}_{\mathbf{s}}(v) = \frac{\nu}{|V|}$.
- Partition M into μ subsets, each with $\frac{|V|}{2\mu}$ edges; each defender uses a uniform probability distribution over each one subset. Thus, $\mathsf{Support}_{\mathbf{s}}(\mathsf{D}) = M$ 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 $e \in \mathsf{Edges}_{\mathbf{s}}(v)$. Since e is unidefender (say by defender 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 $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. Since each edge e is unidefender, it follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) = \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) = 1$. It follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{s}}(v) + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{s}}(u) = \frac{2\nu}{|V|}$. On the other hand, fix any edge $e' = (v', u') \notin \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$. Since M is an Edge Cover, it follows that for the vertex v' (resp., vertex u'), there is a defender D_j such that $v' \in \mathsf{Vertices}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_j))$ (resp., $u' \in \mathsf{Vertices}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_j))$). (Note that $\mathsf{D}_j \neq \mathsf{D}_i$ since M is a Matching.) It follows that $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') < 1$ (resp., $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') < 1$). Thus, $\mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, v') \cdot \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{Prop}_{\mathbf{s}}(\mathsf{D}_i, u') \cdot \mathsf{VP}_{\mathbf{s}}(u') < \mathsf{VP}_{\mathbf{s}}(v') + \mathsf{VP}_{\mathbf{s}}(u') = \frac{2\nu}{|V|}$. 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, $\mathsf{MinHit}_{\mathbf{s}} = \frac{2\mu}{|V|}$. By Corollary 4.3, it follows that $\mathsf{DR}_{\mathbf{s}} = \frac{|V|}{2\mu}$.

■

Proposition D.9 *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* $\mathsf{Supports}_{\mathbf{s}}(\mathsf{D})$ *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 $\mathsf{MinHit}_{\mathbf{s}} = \frac{2\mu}{|V|}$. Consider any edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D})$; so, $e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$ for some defender $\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}$. Since $\mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ is a vertex Cover of the graph $G(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}))$, it follows that $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ or $u \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ (or both). Since $\mathsf{Support}_{\mathbf{s}}(\mathsf{D})$ is a Perfect Matching (and therefore an Edge Cover), there is at least one defender D_k such that $v \in \mathsf{Vertices}_{\mathbf{s}}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. Since $\mathbf{s}$ is Defense-Optimal, Proposition 5.1 implies that there is at most one defender D_k such that $v \in \mathsf{Vertices}_{\mathbf{s}}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. It follows that there is exactly one defender D_k such that $v \in \mathsf{Vertices}_{\mathbf{s}}(\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_k))$. So, clearly, D_k is D_i. Since $\mathsf{Support}_{\mathbf{s}}(\mathsf{D})$ is a Perfect Matching, this implies that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = s_{\mathsf{D}_i}(e)$. We prove:

Claim D.10 $|\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)| = \frac{|V|}{2\mu}$.

Proof: Clearly,

$$\begin{array}{rll} & \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)} s_{\mathsf{D}_i}(e) & \\ = & \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)} \mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) & \\ = & \sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)} \mathsf{MinHit}_{\mathbf{s}} & \text{(since } v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})\text{)} \\ = & |\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)| \cdot \mathsf{MinHit}_{\mathbf{s}} & \\ = & |\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)| \cdot \frac{1}{\mathsf{DR}_{\mathbf{s}}} & \text{(by Corollary 4.3)} \\ = & |\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)| \cdot \frac{2\mu}{|V|} & \text{(since } \mathbf{s} \text{ is Defense-Optimal)} \end{array}$$

Note however, that $\sum_{e \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)} s_{\mathsf{D}_i}(e) = 1$. It follows that $|\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)| = \frac{|V|}{2\mu}$. ■

Since $|\mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)|$ is an integer, Claim D.10 implies that $2\,\mu$ divides $|V|$.

■

E Proofs from Section 6:

E.1 Theorem 6.1:

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

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

implies that there is at least one maxhitter $\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}$ (a defender D_i such that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}_i, v)) = 1$).

There are two cases for each maxhitter D_i: *(i)* 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)* 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 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 $|\mathsf{Supports}_{\mathbf{t}}(\mathsf{D})| < \beta'(G)$. This implies that 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.

F Proofs from Section 7:

F.1 Lemma 7.1:

Fix a defender-pure balanced profile $\mathbf{s}$ and an arbitrary defender $\mathsf{D}_i \in \mathcal{N}_{\mathsf{D}}$. For an edge $e = (u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$, $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i, e) = \frac{\mathsf{VP}_{\mathbf{s}}(v)}{\mathsf{defenders}_{\mathbf{s}}(v)} + \frac{\mathsf{VP}_{\mathbf{s}}(u)}{\mathsf{defenders}_{\mathbf{s}}(u)} = 2c$ (since $\mathbf{s}$ is defender-pure balanced). On the other hand, for an edge $e' = (u', v') \notin \mathsf{Support}_{\mathbf{s}}(\mathsf{D}_i)$, $\mathsf{IP}_{\mathbf{s}}(\mathsf{D}_i, e') = \frac{\mathsf{VP}_{\mathbf{s}}(v')}{\mathsf{defenders}_{\mathbf{s}}(v') + 1} + \frac{\mathsf{VP}_{\mathbf{s}}(u')}{\mathsf{defenders}_{\mathbf{s}}(u') + 1} < 2c$. Hence, $\mathbf{s}$ is a local maximizer for the Individual Profit of defender D_i.

F.2 Theorem 7.2 (Continued):

Note that by construction (Step (2)), $\mathbf{s}$ is defender-pure balanced. Hence, by Lemma 7.1 $\mathbf{s}$ is a local maximizer for the Individual Profit of each defender. So, consider any attacker $\mathsf{A}_i \in \mathcal{N}_\mathsf{A}$. Note that by construction (Steps (1) and (2)) and the assumption that $\mu \geq \beta'(G)$, it follows that $\mathsf{Supports}_\mathbf{s}(\mathsf{D})$ is a Minimum Edge Cover. Since $\mathbf{s}$ is defender-pure, this implies that for each vertex $v \in V$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 1$. Hence, for each vertex $v \in V$, $\mathsf{IP}_\mathbf{s}(\mathsf{A}_i, v) = 1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 0$. So, $\mathbf{s}$ is (vacuously) a local maximizer for the Expected Individual Profit of each attacker. It follows that $\mathbf{s}$ is a Nash equilibrium. Since $\mathbf{s}$ is a Defender-Pure Nash equilibrium, Proposition 4.8 implies that $\mathbf{s}$ is also Defense-Optimal, and we are done.

F.3 Theorem 7.3 (Continued):

Note that by construction (Step (2)), $\mathbf{s}$ is defender-pure, so that $\sum_{v \in V} \mathsf{defenders}_\mathbf{s}(v) = 2\,\mu$. Since $\frac{\nu}{2} \equiv 0 \pmod{\mu}$, it follows that for each vertex $v \in V$, $\mathsf{defenders}_\mathbf{s}(v) \cdot \frac{\nu}{2\,\mu}$ is an integer. Observe that $\sum_{v \in V} \mathsf{defenders}_\mathbf{s}(v) \cdot \frac{\nu}{2\mu} = \frac{\nu}{2\mu} \sum_{v \in V} \mathsf{defenders}_\mathbf{s}(v) = \frac{\nu}{2\mu} \cdot 2\mu = \nu$. Hence, Step (3) yields an attacker pure profile. It follows that $\mathbf{s}$ is a pure profile. Note that by construction (Step (3)), for each vertex $v \in V$, $\frac{\mathsf{VP}_\mathbf{s}(v)}{\mathsf{defenders}_\mathbf{s}(v)} = \frac{\nu}{2\mu}$. Hence, $\mathbf{s}$ is a defender-pure balanced profile. Lemma 7.1 implies that $\mathbf{s}$ is a local maximizer for the Expected Individual Profit for each defender.

So, consider any attacker $\mathsf{A}_i \in \mathcal{N}_\mathsf{A}$. Note that by construction (Steps (1) and (2)) and the assumption that $\mu \geq \beta'(G)$, it follows that $\mathsf{Supports}_\mathbf{s}(\mathsf{D})$ is a Minimum Edge Cover. Since $\mathbf{s}$ is defender-pure, this implies that for each vertex $v \in V$, $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 1$. Hence, for each vertex $v \in V$, $\mathsf{IP}_\mathbf{s}(\mathsf{A}_i, v) = 1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v)) = 0$. So, $\mathbf{s}$ is (vacuously) a local maximizer for the Expected Individual Profit of each attacker. It follows that $\mathbf{s}$ is a Nash equilibrium. Since $\mathbf{s}$ is a Defender-Pure Nash equilibrium, Proposition 4.8 implies that $\mathbf{s}$ is also Defense-Optimal, and we are done.

G Proofs from Section 8:

G.1 Proposition 8.1:

Construct from $\mathbf{s}$ a Defender-Symmetric profile $\mathbf{t}$ for the game $\Pi_{\nu,\mu}$ as follows:

- For each attacker A_i, $t_{\mathsf{A}_i} = s_{\mathsf{A}_i}$. (So, $\mathsf{Supports}_\mathbf{t} = \mathsf{Supports}_\mathbf{s}$.)
- For each defender D_i, $t_{\mathsf{D}_i} = s_\mathsf{D}$. (So, $\mathsf{Supports}_\mathbf{t}(\mathsf{D}) = \mathsf{Supports}_\mathbf{s}(\mathsf{D})$.)

We now show that $\mathbf{t}$ is a Nash equilibrium by proving that it saqtisfies Conditions (1) and (2) in the characterization of Nash equilibria (Proposition 4.1).

- For Condition (1), fix any vertex $v \in V$. Clearly,

$$
\begin{array}{lll}
 & \mathbb{P}_\mathbf{t}(\mathsf{Hit}(v)) & \\
= & \sum_{j \in [m]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} | |\mathcal{D}| = j} \prod_{\mathsf{D}_k \in \mathcal{D}} \mathbb{P}_\mathbf{t}(\mathsf{Hit}(\mathsf{D}_k, v)) & \\
= & \sum_{j \in [m]} (-1)^{j-1} \sum_{\mathcal{D} \subseteq \mathcal{N}_\mathsf{D} | |\mathcal{D}| = j} \left(\mathbb{P}_\mathbf{t}(\mathsf{Hit}(\mathsf{D}_1, v))\right)^j & \text{(since defenders are symmetric)} \\
= & \sum_{j \in [m]} (-1)^{j-1} \binom{\mu}{j} \left(\mathbb{P}_\mathbf{s}(\mathsf{Hit}(\mathsf{D}, v))\right)^j & \text{(by construction)} \\
= & 1 - \left(1 - \mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))\right)^\mu \, . &
\end{array}
$$

 Since $\mathbf{s}$ is a Nash equilibrium, it follows by Condition (1) in the characterization of Nash equilibria (Proposition 4.1) that $\mathbb{P}_\mathbf{s}(\mathsf{Hit}(v))$ is minimized on $\mathsf{Supports}_\mathbf{s}(\mathsf{A})$. Since $\mathsf{Supports}_\mathbf{t}(\mathsf{A}) = \mathsf{Supports}_\mathbf{s}(\mathsf{A})$, the derived formula implies that $\mathbb{P}_\mathbf{t}(\mathsf{Hit}(v))$ is minimized on $\mathsf{Supports}_\mathbf{t}(\mathsf{A})$. So, t satisfies Condition (1) in the characterization of Nash equilibria (Proposition 4.1).

- For Condition (2), fix any defender $\mathsf{D}_i \in \mathcal{N}_\mathsf{D}$. Then, for each vertex $v \in V$,

$$\begin{aligned}
&\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \\
= &\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D}\subseteq\mathcal{N}_\mathsf{D}\setminus\mathsf{D}_i \mid |\mathcal{D}|=j-1} \prod_{\mathsf{D}_k\in\mathcal{D}} \mathbb{P}_{\mathbf{t}}(\mathsf{Hit}(\mathsf{D}_k, v)) \\
= &\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \sum_{\mathcal{D}\subseteq\mathcal{N}_\mathsf{D}\setminus\mathsf{D}_i \mid |\mathcal{D}|=j-1} (\mathbb{P}_{\mathbf{t}}(\mathsf{Hit}(\mathsf{D}_1, v)))^{j-1} && \text{(since defenders are symmetric)} \\
= &\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(\mathsf{D}, v)))^{j-1} && \text{(by construction)} \\
= &\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)))^{j-1} && \text{(since } \mathcal{N}_\mathsf{D} = \{\mathsf{D}\}\text{)}\,.
\end{aligned}$$

Note that for any vertex $v \in \mathsf{Supports}_{\mathbf{s}}(\mathsf{A}) = \mathsf{Supports}_{\mathbf{t}}(\mathsf{A})$, Condition (1) in the characterization of Nash equilibria (Proposition 4.1) implies that $\mathbb{P}_{\mathbf{s}}(\mathsf{Hit}(v)) = \mathsf{MinHit}_{\mathbf{s}}$. Since $\mathsf{Supports}_{\mathbf{s}}(\mathsf{A}) = \mathsf{Supports}_{\mathbf{t}}(\mathsf{A})$, it follows that for each vertex $t \in \mathsf{Supports}_{\mathbf{t}}(\mathsf{A})$.

$$\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \;=\; \sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{t}})^{j-1}\,.$$

Note also that for a vertex $v \notin \mathsf{Supports}_{\mathbf{t}}(\mathsf{A}) = \mathsf{Supports}_{\mathbf{s}}(\mathsf{A})$, $\mathsf{VP}_{\mathbf{t}}(v) = 0$. It follows that for any vertex $v \in V$,

$$\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{t}}(v) \;=\; \begin{cases} \sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{s}})^{j-1} & \text{if } v \in \mathsf{Supports}_{\mathbf{t}}(\mathsf{A}) \\ 0 & \text{otherwise} \end{cases}$$

Consider first any edge $e = (u, v) \in \mathsf{Supports}_{\mathbf{t}}(\mathsf{D})$. Since $\mathsf{Supports}_{\mathbf{t}}(\mathsf{D}) = \mathsf{Support}_{\mathbf{s}}(\mathsf{D})$, it follows that $e \in \mathsf{Support}_{\mathbf{t}}(\mathsf{D})$. Since s is a Nash equilibrium, Lemma 4.6 implies that either $u \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ or $u \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ (or both). There are two cases to consider:

1. Both $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ and $u \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$. Then,

$$\begin{aligned}
&\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{t}}(v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{t}}(u) \\
= &\left(\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{s}})^{j-1}\right) (\mathsf{VP}_{\mathbf{t}}(v) + \mathsf{VP}_{\mathbf{t}}(u)) \\
= &\left(\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{s}})^{j-1}\right) (\mathsf{VP}_{\mathbf{s}}(v) + \mathsf{VP}_{\mathbf{s}}(u)) && \text{(by construction of } \mathbf{t}\text{)} \\
= &\left(\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{s}})^{j-1}\right) \mathsf{VP}_{\mathbf{s}}(e) \\
= &\left(\textstyle\sum_{j\in[m]} \frac{1}{j} (-1)^{j-1} \binom{\mu-1}{j-1} (\mathsf{MinHit}_{\mathbf{s}})^{j-1}\right) \max_{e'\in E} \mathsf{VP}_{\mathbf{s}}(e') && \text{(by Proposition 4.1 (1))}
\end{aligned}$$

2. $v \in \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$ but $u \notin \mathsf{Support}_{\mathbf{s}}(\mathsf{A})$. Then,

$$\begin{aligned}
&\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{t}}(v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{t}}(u) \\
= &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot (\mathsf{VP}_{\mathbf{t}}(v) + 0) \\
= &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot (\mathsf{VP}_{\mathbf{s}}(v) + 0) && \text{(by construction of } \mathbf{t}\text{)} \\
= &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot \max_{e'\in E} \mathsf{VP}_{\mathbf{s}}(e') && \text{(by Proposition 4.1 (1))}
\end{aligned}$$

Consider now any edge $e = (u, v) \notin \mathsf{Support}_{\mathbf{t}}(\mathsf{D})$. Since $\mathsf{Support}_{\mathbf{t}}(\mathsf{D}) = \mathsf{Support}_{\mathbf{s}}(\mathsf{D})$, it follows that $e \notin v\mathsf{Support}_{\mathbf{s}}(\mathsf{D})$. Hence,

$$\begin{aligned}
&\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) \cdot \mathsf{VP}_{\mathbf{t}}(v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u) \cdot \mathsf{VP}_{\mathbf{t}}(u) \\
\leq &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot (\mathsf{VP}_{\mathbf{t}}(v) + \mathsf{VP}_{\mathbf{t}}(u)) \\
= &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot (\mathsf{VP}_{\mathbf{s}}(v) + \mathsf{VP}_{\mathbf{s}}(u)) && \text{(by construction of } \mathbf{t}\text{)} \\
= &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot \mathsf{VP}_{\mathbf{s}}(e) \\
\leq &(\mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, v) + \mathsf{Prop}_{\mathbf{t}}(\mathsf{D}_i, u)) \cdot \max_{e'\in E} \mathsf{VP}_{\mathbf{s}}(e') && \text{(by Proposition 4.1 (1))}
\end{aligned}$$

Condition (2) follows. Hence, $\mathbf{t}$ is a Nash euilibrium.