How Many Attackers Can Selfish Defenders Catch?*

Marios Mavronicolas† *Burkhard Monien*‡ *Vicky Papadopoulou Lesta*§

(May 24, 2013)

Abstract

In a distributed system with *attacks* and *defenses,* both *attackers* and *defenders* are self-interested entities. We assume a *reward-sharing* scheme among *interdependent* defenders; each defender wishes to (locally) maximize her own total *fair share* to the attackers extinguished due to her involvement (and possibly due to those of others). What is the *maximum* amount of protection achievable by a number of such defenders against a number of attackers while the system is in a *Nash equilibrium*? As a measure of system protection, we adopt the *Defense-Ratio* [20], which provides the expected (inverse) proportion of attackers caught by the defenders. In a *Defense-Optimal* Nash equilibrium, the Defense-Ratio matches a simple lower bound.

We discover that the existence of Defense-Optimal Nash equilibria depends in a subtle way on how the number of defenders compares to two natural graph-theoretic thresholds we identify. In this vein, we obtain, through a combinatorial analysis of Nash equilibria, a collection of trade-off results:

- When the number of defenders is either sufficiently small or sufficiently large, Defense-Optimal Nash equilibria may exist. The corresponding decision problem is computationally tractable for a large number of defenders; the problem becomes $\mathcal{NP}$-complete for a small number of defenders and the intractability is inherited 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 Defense-Optimal Nash equilibria do not exist.

*A preliminary version of this work appeared in the *CD-ROM Proceedings of the 41st Hawaii International Conference on System Sciences*, Track on Software Technology, Minitrack on Algorithmic Challenges in Emerging Applications of Computing, January 2008. 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. Part of this work was performed while this author was visiting the Faculty of Electrical Engineering, Computer Science and Mathematics, University of Paderborn. Email `mavronic@ucy.ac.cy`

‡Faculty of Electrical Engineering, Computer Science and Mathematics, University of Paderborn, 33102 Paderborn, Germany. Part of this work was performed while this author was visiting the Department of Computer Science, University of Cyprus. Email `bm@upb.de`

§Department of Engineering, Computer Science and European University Cyprus, Nicosia 1516, Cyprus. Part of this work was performed while this author was visiting the Faculty of Electrical Engineering, Computer Science and Mathematics, University of Paderborn. Email `v.papadopoulou@euc.ac.cy`

1 Introduction

1.1 The Model and its Rationale

Safety and *security* are key issues for the design and operation of a distributed system. With the unprecedented advent of the Internet, there is a growing interest to formalize, design and analyze distributed systems prone to *malicious attacks* and (*non-malicious*) *defenses*. A new dimension stems from the fact that Internet servers and clients are controlled by *selfish* agents interested in the local maximization of their own benefits rather than the optimization of global performance [1, 4, 8, 9, 10]. So, it is a challenge to formalize and analyze 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* models a *virus*; it is a malicious client that targets a host to destroy. A *defender* is a non-malicious server modeling the *antivirus software* installed and licensed at a single host in order to protect it. We assume two species with α attackers and δ defenders. So, δ is proportional to the actual cost of purchasing and installing several units of (licensed) antivirus software.

Associating attacks with nodes makes sense since security attacks are often directed to individual hosts such as commercial and public sector entities. Associating defenses with edges is motivated by *Network Edge Security* [16], a distributed *firewall architecture* where antivirus software is implemented by a *distributed algorithm* running on a subnetwork. Such architectures are attractive since they offer increased fault-tolerance and the benefit of sharing licensing costs among hosts. We focus on the simplest case where the subnetwork is a *single* link.

In reality, malicious attackers are *independent*; each attacker tries to maximize on her own the amount of harm it causes (cf. [26, 11, 13]). So, each attacker is modeled as a strategic *player* seeking to escape the antivirus software and receive *payoff* (or *utility*) 1. The strategy of an attacker does not (directly) affect the payoff of another.

On the other hand, defenders are strategic but *non-cooperative*; so, each defender competes to maximize the number of viruses she catches on the basis of an intuitive *reward-sharing* scheme: Whenever more than one colocated defenders are extinguishing the attacker(s) targeting a host, each defender is rewarded with the *fair share* of the number of attackers extinguished. So, each defender is modeled as a strategic player seeking to maximize her payoff of total fair share.

Both selfish species may use *mixed* strategies. In a *Nash equilibrium* [22, 23], no player can unilaterally increase her expected utility. To evaluate Nash equilibria, we employ the *Defense-Ratio*; this is the ratio of the given number α over the expected number of attackers extinguished by the defenders (cf. [17, 18]). Motivated by *best-case* Nash equilibria and the

Price of Stability [2], we introduce *Defense-Optimal* Nash equilibria where the Defense-Ratio attains the value $\max\left\{1, \frac{|V|}{2\delta}\right\}$ (Definition 6.1); we chose this value since it is observed to be a (tight) lower bound on Defense-Ratio (Corollary 6.2). A *Defense-Optimal* graph (for a given δ) is one with a Defense-Optimal Nash equilibrium.

1.2 Contribution

We explore the existence and the computational complexity of a Defense-Optimal Nash equilibrium for a given number of defenders δ. We discover that both depend on two graph-theoretic thresholds for δ: $\frac{|V|}{2}$ and $\beta'(G) \geq \frac{|V|}{2}$, the size of a *Minimum Edge Cover*.

Our chief tool is a combinatorial characterization of Nash equilibria we obtain (Proposition 5.1). For *Pure* Nash equilibria where both species use *pure* strategies, the characterization yields some necessary graph-theoretic conditions for Nash equilibria (Proposition 5.7); it also yields sufficient conditions for Defense-Optimal Nash equilibria (Theorems 6.3 and 6.5). Our end findings amount as follows: When either $\delta \leq \frac{|V|}{2}$ (***few defenders***) or $\delta \geq \beta'(G)$ (***too many defenders***), there *are* cases allowing for a Defense-Optimal Nash equilibrium. When $\frac{|V|}{2} < \delta < \beta'(G)$ (***many defenders***), there are no such cases. We continue with the details.

For the case of few defenders, we provide a combinatorial characterization of Defense-Optimal graphs, which points out an interesting connection to *Fractional* (*Perfect*) *Matchings* [24, Chapter 2]. Roughly speaking, these graphs are a strict subset of graphs with a Fractional Perfect Matching: for a Defense-Optimal graph, and assuming that $\delta \leq \frac{|V|}{2}$, it is possible to partition *some* Fractional Perfect Matching of it into δ smaller, *vertex-disjoint* Fractional Perfect Matchings so that the total weight inherited to each partite equals $\frac{|V|}{2\,\delta}$ (Theorem 7.4). Call such a Fractional Perfect Matching a *δ-Partitionable Fractional Perfect Matching*; this is a previously unconsidered, combinatorial concept in *Fractional Graph Theory* [24].

We prove that the recognition problem for graphs with a δ-Partitionable Fractional Perfect Matching is $\mathcal{NP}$-complete (Corollary 2.12); this intractability result holds for an *arbitrary* value of δ. Hence, so is the decision problem for a Defense-Optimal Nash equilibrium $\left(\text{for } \delta \leq \frac{|V|}{2}\right)$ (Corollary 7.7). To establish the $\mathcal{NP}$-completeness, we develop some techniques of independent interest for the reduction of Fractional (Perfect) Matchings (Section 2.2). A further interesting number-theoretic consequence of the combinatorial characterization for Defense-Optimal graphs $\left(\text{for } \delta \leq \frac{|V|}{2}\right)$ is that δ divides $|V|$ (Corollary 7.5).

On the positive side, we identify another restriction of graphs with a Fractional Perfect Matching that are Defense-Optimal in some well-characterized cases (Theorem 7.8); these are the graphs with a *Perfect Matching*.

For the case of too many defenders, we identify two cases where there are Defense-Optimal Nash equilibria, namely the *vertex-balanced* Nash equilibria, with a special structure enabling their efficient computation (Theorems 9.2 and 9.5). The two corresponding algorithms use the computation of a *Minimum Edge Cover*; the second algorithm requires that 2δ divides α.

For the case of many defenders, we provide a combinatorial proof that there is *no* Defense-Optimal graph for $\frac{|V|}{2} < \delta < \beta'(G)$ (Theorem 8.1). This is paradoxical since with fewer defenders $\left(\delta \leq \frac{|V|}{2}\right)$ there are Defense-Optimal graphs. Since the Defense-Ratio in a Defense-Optimal Nash equilibrium has a transition around $\delta = \frac{|V|}{2}$, this paradox may not be surprising.

1.3 Related Work and Comparison

Our network game with multiple attackers and defenders is a particular instance of *Interdependent Security Games*, studied by Kearns and Orliz [11], where a number of players must make individual decisions related to security which affect the ultimate safety and effectiveness of other players.

The very special but yet highly non-trivial case with a single defender was originally introduced in [20] and further studied in [7, 17, 18, 19]; the case of $\delta > 1$ defenders is requiring a far more challenging combinatorial and graph-theoretic analysis. Hence, we view this work as a major generalization of previous related work. Very recently, a variant of our network game with the roles of attackers and defenders interchanged was studied in [25].

The notion of Defense-Ratio generalizes a corresponding definition from [17, Section 3.4] from $\delta = 1$ to $\delta > 1$ defenders. The special case with $\delta = 1$ of Theorem 7.4 was considered in [18, Corollary 2]; it allowed for a polynomial time algorithm for the decision problem for a Defense-Optimal Nash equilibrium by reduction to the recognition problem for a graph with a Fractional Perfect Matching. In contrast, for an arbitrary $\delta \leq \frac{|V|}{2}$, the decision problem for a Defense-Optimal Nash equilibrium reduces to the recognition problem for graphs with a δ-partitionable Fractional Perfect Matching and becomes $\mathcal{NP}$-complete (Corollary 7.7).

The latter recognition problem simultaneously generalizes a tractable and an intractable recognition problem: those for graphs with a Perfect Matching [5] and graphs whose vertex set can be partitioned into triangles [6, GT11], respectively.

2 Background and Preliminaries from Graph Theory

Throughout the paper, for an integer $n \geq 1$, denote $[n] = \{1, \ldots, n\}$; for a number $x \neq 0$, $\mathtt{sgn}(x)$ denotes the *sign* of x (which is $+1$ or -1).

We shall consider a simple undirected graph $G = \langle V, E \rangle$. The *trivial* graph consists of a single edge. Denote as $\mathsf{d}_G(v)$ the *degree* of vertex v in G. An edge $(u, v) \in E$ is *pendant* if $\mathsf{d}_G(u) = 1$ but $\mathsf{d}_G(v) > 1$. A *path* is a sequence of distinct vertices $v_1, v_2, \cdots v_{\ell+1}$ successively connected by edges. In a cycle $\mathcal{C}$, $v_{\ell+1} = v_1$; $\mathcal{C}$ has length ℓ, and it is *even* (resp., *odd*) if ℓ is even (resp., odd). A *triangle* is a cycle of length three. There are polynomial time algorithms to compute an odd cycle (cf. [12, Proposition 2.27]) or an even cycle [14, 21, 28].

Vertex and edge sets induce subgraphs. For a vertex set $U \subseteq V$, $G(U)$ is the subgraph of G induced by U; denote $\mathsf{Edges}_G(U) = \{(u, v) \in E \mid u, v \in U\}$ so that $G(U) = \langle U, \mathsf{Edges}_G(U)\rangle$. For an edge set $F \subseteq E$, $G(F)$ is the subgraph of G induced by F; denote $\mathsf{Vertices}_G(F) = \bigcup_{(u,v)\in F}\{u, v\}$ so that $G(F) = \langle \mathsf{Vertices}_G(F), F\rangle$. A *component* of G is a maximal connected subgraph. A cycle is *isolated* if it is a component; else, it is *non-isolated*. A component containing a cycle is *cyclic*; else it is *acyclic*. So, an acyclic component has a pendant edge.

A *Vertex Cover* (resp., *Edge Cover*) is a vertex (resp., edge) set $VC \subseteq V$ (resp., $EC \subseteq E$) such that for each edge $(u, v) \in E$ (resp., for each vertex $v \in V$) either $u \in VC$ or $v \in VC$ (resp., there is an edge $(u, v) \in EC$); a *Minimum Vertex Cover* (resp., *Minimum Edge Cover*) is one that has minimum size, denoted as $\beta(G)$ (resp., $\beta'(G)$). Clearly, $\frac{|V|}{2} \leq \beta'(G)$. Denote as $\mathcal{EC}(G)$ the set of all Edge Covers of G.

2.1 Matchings and Fractional (Perfect) Matchings

A *Matching* is a set $M \subseteq E$ of non-incident edges; a *Maximum Matching* is one with maximum size. Computing a Minimum Edge Cover is polynomial time reducible to computing a Maximum Matching—see, e.g., [27, Theorem 3.1.22] or [15], which is done in polynomial time (cf. [5]).

A *Perfect Matching* is a Matching that is also an Edge Cover; so, a Perfect Matching has size $\frac{|V|}{2}$. A *Perfect-Matching* graph G is one that has a Perfect Matching; so, $\beta'(G) = \frac{|V|}{2}$. Any (polynomial time) algorithm to compute a Maximum Matching yields a (polynomial time) algorithm to recognize Perfect-Matching graphs and compute a Perfect Matching.

A *Fractional Matching* is a map $f : E \to [0, 1]$ where for each vertex $v \in V$, $\sum_{e\in E|e\ni v} f(e) \leq 1$. (Matching is the special case where $f(e) \in \{0, 1\}$ for each edge $e \in E$; Perfect Matching is the special case of Matching where $\sum_{e\in E|e\ni v} f(e) = 1$.) The *range* of a Fractional Matching f is the set $\mathsf{Range}(f) = \{f(e) \mid e \in E\}$; so, $\mathsf{Range}(f)$ is a finite set of values from $[0, 1]$. Induced by f is the set $E(f) = \{e \in E \mid f(e) > 0\}$; call $|E(f)|$ the *size* of f. Given two Fractional Matchings f and f', write $f' \subseteq f$ (resp., $f' \subset f$) if $E(f') \subseteq E(f)$ (resp., $E(f') \subset E(f)$).

A *Fractional Perfect Matching* is a Fractional Matching f such that for each vertex $v \in V$, $\sum_{e\in E|e\ni v} f(e) = 1$. So, for each vertex $v \in V$, there is at least one edge $e \in E$ with $e \ni v$ such that $f(e) > 0$, so that $e \in E(f)$ and $E(f)$ is an Edge Cover. A *Fractional Maximum Matching* is a Fractional Matching f that maximizes $\sum_{e \in E} f(e)$ among all Fractional Matchings. A Fractional Perfect Matching is a Fractional Maximum Matching (but not vice versa). 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 edge.*

Proof. Assume, by way of contradiction, that $G(E(f))$ has a pendant edge (u, v) with $\mathsf{d}_{G(E(f))}(u) = 1$ but $\mathsf{d}_{G(E(f))}(v) > 1$. Since f is a Fractional Perfect Matching, $\sum_{e \in E | e \ni v} f(e) = 1$ and $\sum_{e \in E | e \ni v} f(e) = 1$. By assumption on u, the first equality implies that $f((u, v)) = 1$. By assumption on v, the second equality implies that $f((u, v)) < 1$. A contradiction. ■

Lemma 2.1 implies that for a Fractional Perfect Matching f, each component of $G(E(f))$ is either a single edge or a (non-trivial) subgraph of G with no pendant edges; hence, each acyclic component of $G(E(f))$ is a single edge. The proof for [24, Theorem 2.1.5] establishes that a Fractional Maximum Matching with smallest size induces no pendant edge; so, Lemma 2.1 identifies another special case of Fractional Maximum Matching inducing no pendant edge.

Say that two maps $f : E \rightarrow [0, 1]$ and $f' : E \rightarrow [0, 1]$ are *equivalent* if for each vertex $v \in V$, $\sum_{e \in E | e \ni v} f(e) = \sum_{e \in E | e \ni v} f'(e)$. Clearly, a map $f : E \rightarrow [0, 1]$ that is equivalent to a Fractional (Perfect) Matching is also a Fractional (Perfect) Matching.

The class of graphs with a Fractional Perfect Matching is recognizable in polynomial time via a Linear Programming formulation [3]. (The same holds for the corresponding search problem.)

2.2 Reductions of Fractional (Perfect) Matchings

We recall a combinatorial property of a special case of Fractional Maximum Matchings.

Proposition 2.2 ([24], Theorem 2.1.5) *Consider a Fractional Maximum Matching f with smallest size. Then, $G(E(f))$ has only single edges and odd cycles.*

Implicit in the constructive proof for [24, Theorem 2.1.5], which assumes a Fractional Maximum Matching with smallest size, are two polynomial time reductions for a Fractional Maximum Matching with smallest size; for self-containment, they are formally stated here, extended to Fractional Perfect Matchings. The first reduction (Algorithm EliminateEvenCycles in Figure 1) will eliminate all induced even cycles from an *arbitrary* Fractional Matching without assuming that the Fractional Matching is Maximum and has smallest size. The second reduction (Algorithm IsolateOddCycles in Figure 2) will eliminate all induced non-isolated odd cycles from a Fractional Perfect Matching with no induced even cycle; this complements the corresponding

Algorithm EliminateEvenCycles
<u>INPUT</u>: A graph G and a Fractional Matching f for G.
<u>OUTPUT</u>: An equivalent Fractional Matching $f'' \subseteq f$ with no even cycle.

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

While $G(E(f))$ contains 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 $g{:}E(\mathcal{C}) \rightarrow \{-1, +1\}$ with $g(e) = -1$ or $+1$ (alternately, starting with $g(e_0) = -1$).

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

$$f''(e) \quad := \quad \begin{cases} f(e) + g(e) \cdot f(e_0), & \text{if } e \in E(\mathcal{C}) \\ f(e), & \text{if } e \notin E(\mathcal{C}) \end{cases} .$$

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

Figure 1: The algorithm EliminateEvenCycles consisting of a single loop; upon termination, there will be no even cycle for the output f''. Step (1) chooses an edge e_0 on the cycle $\mathcal{C}$ on which f is minimized; Step (2) assigns a sign to each edge e on $\mathcal{C}$, alternating for consecutive edges of the cycle. The new values for f'' are assigned in Step (3); note that $f''(e_0) = 0$. Step (4) prepares the input f for the next loop iteration.

reduction in the proof of [24, Theorem 2.1.5] which considers a Fractional Maximum Matching with smallest size and with no induced even cycle. Respectively, we prove:

Proposition 2.3 *Consider a Fractional Matching f. Then, there is a polynomial time algorithm to transform f into an equivalent Fractional Matching $f'' \subseteq f$ with no even cycle.*

Proposition 2.4 *Consider a Fractional Perfect Matching f with no even cycle. Then, there is a polynomial time algorithm to transform f into an equivalent Fractional Perfect Matching $f' \subseteq f$ with no non-isolated odd cycle.*

The proofs of Propositions 2.3 and 2.4 are deferred to the Appendix. We now prove:

Proposition 2.5 *Consider a Fractional Perfect Matching f. Then, there is a polynomial time algorithm to transform f into an equivalent Fractional Perfect Matching $f' \subseteq f$ with only single edges and odd cycles.*

We consider the algorithm EliminateEven&IsolateOddCycles in Figure 3: the sequential cascade of the algorithms EliminateEvenCycles and IsolateOddCycles from Figures 1 and 2, respectively.

Algorithm IsolateOddCycles

INPUT: A graph G and a Fractional Perfect Matching f for G with no even cycle.

OUTPUT: An equivalent Fractional Perfect Matching $f' \subseteq f$ with no non-isolated odd cycle.

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

While $G(E(f))$ contains a non-isolated odd cycle $\mathcal{C}$ **do**

(1) Choose a vertex $v_0 \in \mathcal{C}$ with $\mathsf{d}_{G(E(f))}(v_0) \geq 3$ and an edge $(v_0, v_1) \in E(f)$ with $v_1 \notin \mathcal{C}$.

(2) **While** $E(f)$ includes all edges from $E(\mathcal{C}) \cup \{(v_0, v_1)\}$ **do**

(2/a) Choose a path $v_1, v_2, \cdots, v_r$ with $v_r = v_l$ for some $l \in 0 \cup [r-2]$.

(2/b) Define a function $g : E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} \rightarrow \left\{+1, -1, +\frac{1}{2}, -\frac{1}{2}\right\}$ with

$$g(e) = \begin{cases} +\frac{1}{2} \text{ or } -\frac{1}{2}, & \text{if } e \in E(\mathcal{C}) \text{ (alternately, starting with } +\frac{1}{2} \text{ for an edge incident to } v_0) \\ +1 \text{ or } -1, & \text{if } e = (v_k, v_{k+1}) \text{ for } 0 \leq k \leq l-1 \text{ with } l > 0 \text{ (alternately, starting with } -1) \\ +\frac{1}{2} \text{ or } -\frac{1}{2}, & \text{if } e = (v_k, v_{k+1}) \text{ for } l \leq k \leq r-1 \text{ (alternately, starting with a sign opposite} \\ & \text{to the sign of the last value assigned by } g) \end{cases}.$$

(2/c) Choose an edge $e_0 \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$ such that

$$\frac{f(e_0)}{|g(e_0)|} = \min \left\{ \min_{e \in E(\mathcal{C})} \frac{f(e)}{|g(e)|}, \min_{\substack{l>0 \\ 0 \leq k \leq l-1}} \frac{f((v_k, v_{k+1}))}{|g((v_k, v_{k+1}))|}, \min_{l \leq k \leq r-1} \frac{f((v_k, v_{k+1}))}{|g((v_k, v_{k+1}))|} \right\};$$

(2/d) If $g(e_0) > 0$, then set $g := -g$.

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

$$f'(e) := \begin{cases} f(e) + g(e) \cdot \frac{f(e_0)}{|g(e_0)|}, & \text{if } e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} \\ f(e), & \text{otherwise} \end{cases}.$$

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

Figure 2: The algorithm IsolateOddCycles consisting of an outer loop, which includes an inner loop (Step (2)); upon termination, there will be no non-isolated odd cycle for f'. For Step (1), note that a vertex $v_0 \in \mathcal{C}$ with $\mathsf{d}_{G(E(f))}(v_0) \geq 3$ exists since $\mathcal{C}$ is non-isolated; v_0 has two incident edges from $\mathcal{C}$ and at least one incident edge (v_0, v_1) outside $\mathcal{C}$. The precondition for the inner loop is the inclusion of all edges from $E(\mathcal{C}) \cup \{(v_0, v_1)\}$ in $E(f)$; note that $\mathcal{C}$ remains a (non-isolated) cycle (and the inner loop is executed) as long as no such edge has been eliminated from f (by Step (2/e)). For Step (2/a), note that a path $v_1, \cdots, v_r$ with $v_r = v_l$ for some $l \in 0 \cup [r-2]$ exists since $G(E(f))$ has no pendant edges (by Lemma 2.1); this path together with $\mathcal{C}$ make a *bicycle graph.* For Step (2/b), it will follow from Lemma B.1 that for any vertex v_k with $0 < k \leq r$, it holds that $v_k \notin \mathcal{C} \backslash \{v_0\}$. So, Step (2/b) assigns a signed coefficient to each edge $e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$. Step (2/c) chooses an edge e_0 on either $\mathcal{C}$ or path $v_1, v_2, \cdots, v_r$ that minimizes the ratio of values of f and $|g|$. Step (2/d) adjusts g so that $g(e_0) < 0$. The new values for f' are assigned in Step (2/e); note that $f'(e_0) = 0$ (by Step (2/d)). Step (2/f) prepares the input (f) for the next (inner) loop iteration.

Algorithm EliminateEven&IsolateOddCycles
<u>INPUT</u>: A graph G and a Fractional Perfect Matching f for G.
<u>OUTPUT</u>: A Fractional Perfect Matching f' with no even cycle and no non-isolated odd cycle.

(1) Apply EliminateEvenCycles on f to obtain f''.

(2) Apply IsolateOddCycles on f'' to obtain f'.

Figure 3: The algorithm EliminateEven&IsolateOddCycles, cascading the algorithms EliminateEvenCycles and IsolateOddCycles from Figures 1 and 2, respectively.

Proof. By Proposition 2.3, $f'' \subseteq f$ is a Fractional Perfect Matching with no even cycle, which is equivalent to f. By Proposition 2.4, $f' \subseteq f''$ is a Fractional Perfect Matching with no non-isolated odd cycle, which is equivalent to f''. It follows that (*i*) f' is equivalent to f and $f' \subseteq f$, and (*ii*) f' has no even cycle and no non-isolated odd cycle. Since f is a Fractional Perfect Matching, Condition (*i*) implies that f' is a Fractional Perfect Matching; by Lemma 2.1, this implies that each acyclic component of $G(E(f'))$ is a single edge. By Condition (*ii*), it follows that each cyclic component of $G(E(f'))$ is an isolated odd cycle. ■

2.3 δ-Partitionable Fractional Perfect Matchings

2.3.1 Definition and Preliminaries

Fix an integer $\delta \geq 1$. A Fractional Perfect Matching $f : E \to \mathbb{R}$ is ***δ-Partitionable*** if the edge set $E(f)$ can be partitioned into δ (non-empty,) vertex-disjoint partites $E_1, \cdots, E_\delta$ so that for each partite E_j, $\sum_{e \in E_j} f(e) = \frac{|V|}{2\,\delta}$. So, a 1-Partitionable Fractional Perfect Matching is a Fractional Perfect Matching. Consider now a δ-Partitionable Fractional Perfect Matching f; note that for a partite E_j and any vertex $v \in V(E_j)$, $\sum_{e \in E | e \ni v} f(e) = 1$ (since f is a Fractional Perfect Matching). Since the partites are vertex-disjoint, this implies that for any vertex $v \in V(E_j)$, $\sum_{e \in E_j | e \ni v} f(e) = 1$. We now prove a necessary condition on partites:

Proposition 2.6 *Consider a δ-Partitionable Fractional Perfect Matching f, and fix a partite E_j. Then,* $|V(E_j)| = \frac{|V|}{\delta}$.

Proof. Clearly, $\sum_{e \in E_j} f(e) = \frac{1}{2} \sum_{v \in V(E_j)} \left(\sum_{e \in E_j | e \ni v} f(e) \right) = \frac{1}{2} \sum_{v \in V(E_j)} 1 = \frac{|V(E_j)|}{2}$. Hence, $|V(E_j)| = 2 \cdot \sum_{e \in E_j} f(e) = 2 \cdot \frac{|V|}{2\,\delta} = \frac{|V|}{\delta}$, as needed. ■

Proposition 2.6 immediately implies:

Corollary 2.7 *If G has a δ-Partitionable Fractional Perfect Matching, then δ divides $|V|$, so that $\delta \leq \frac{|V|}{2}$.*

We observe that the equality in the condition $\delta \leq \frac{|V|}{2}$ in Corollary 2.7 is not necessary:

Proposition 2.8 *There is a graph G and an integer δ such that G has a δ-Partitionable Fractional Perfect Matching while $\delta < \frac{|V|}{2}$.*

Proof. Consider the cycle graph $\mathcal{C}_3$ with $\delta = 1$. The map $f : E(\mathcal{C}_3) \rightarrow [0, 1]$ with $f(e) = \frac{1}{2}$ for each edge $e \in E(\mathcal{C}_3)$ is an 1-Partitionable Fractional Perfect Matching. ∎

We finally identify an equivalence relation on the set of Fractional Perfect Matchings:

Proposition 2.9 *Consider a δ-Partitionable Fractional Perfect Matching f and an equivalent Fractional Perfect Matching $f' \subseteq f$. Then, f' is δ-Partitionable.*

Proof. Consider the δ (non-empty,) vertex-disjoint partites $E_1, \cdots, E_\delta$. Define edge sets $E'_1, \cdots, E'_\delta$ so that for each $j \in [\delta]$, $E'_j = \{e \in E_j \mid f'(e) > 0\}$. Since $f' \subseteq f$, it follows that for each $j \in [\delta]$, $E'_j \subseteq E_j$; hence, the collection $E'_1, \cdots, E'_\delta$ partitions $E(f')$. Also, by construction, the edge sets $E'_1, \cdots, E'_\delta$ are vertex-disjoint; so call them partites.

Fix a partite E'_j. Then, $\sum_{e \in E'_j | e \ni v} f'(e) = \sum_{e \in E_j | e \ni v} f'(e)$ (since $f(e) = 0$ for each edge $e \in E_j \backslash E'_j$). Since f' is equivalent to f, for any vertex $v \in V$, $\sum_{e \in E | v \in e} f'(e) = \sum_{e \in E | v \in e} f(e)$. Since $E'_1, \cdots, E'_\delta$ are vertex-disjoint, $\sum_{e \in E'_j | e \ni v} f'(e) = \sum_{e \in E_j | e \ni v} f(e)$. Since f is δ-Partitionable, it follows that $\sum_{e \in E'_j | e \ni v} f'(e) = \frac{|V|}{2\,\delta}$. So, f' is δ-Partitionable. ∎

2.3.2 Characterization

We show:

Proposition 2.10 *A graph G has a δ-Partitionable Fractional Perfect Matching if and only if E contains a collection of δ (non-empty) vertex-disjoint edge sets $E_1, \cdots, E_\delta$ such that* (C1) *$\bigcup_{j \in [\delta]} E_j$ is an Edge Cover, and* (C2) *for each edge set E_j with $j \in [\delta]$,* (C2/a) *E_j consists of single edges and odd cycles, and* (C2/b) *$|V(E_j)| = \frac{|V|}{\delta}$.*

Proof. Assume first that G has a δ-Partitionable Fractional Perfect Matching f. By Proposition 2.5, there is an equivalent Fractional Perfect Matching $f' \subseteq f$ with only single edges and odd cycles. Since f is δ-Partitionable and $f' \subseteq f$, Proposition 2.9 implies that f' is δ-Partitionable. So, the edge set $E(f')$ can be partitioned into δ (non-empty), vertex-disjoint partites $E_1, \cdots, E_\delta$ so that for each partite E_j with $j \in [\delta]$, $\sum_{e \in E_j} f(e) = \frac{|V|}{2\,\delta}$.

Consider now the (vertex-disjoint) edge sets $E_1, \cdots, E_\delta$. Since f' is a Fractional Perfect Matching, $E(f')$ is an Edge Cover; so, $\bigcup_{j \in [\delta]} E_j$ is an Edge Cover and (C1) follows. Consider now any edge set E_j with $j \in [\delta]$. Since f' consists of single edges and odd cycles, (C2/a) follows; since f' is a Fractional Perfect Matching, (C2/b) follows from Proposition 2.6.

Assume now that E satisfies the condition of the claim. To prove that G has a δ-Partitionable Fractional Perfect Matching, define the function $f : E \to [0,1]$ with

$$f(e) \quad = \quad \begin{cases} 1, & \text{if } e \in E_j \text{ with } j \in [\delta] \text{ and } E_j \text{ is a single edge} \\ \frac{1}{2}, & \text{if } e \in E_j \text{ with } j \in [\delta] \text{ and } E_j \text{ is an odd cycle} \\ 0, & \text{if } e \in E \backslash \bigcup_{j \in [\delta]} E_j \end{cases} .$$

Consider any vertex $v \in V$. By Condition (C1), $v \in V(E_j)$ for some set E_j. By Condition (C2/a), in either case for E_j (a single edge or an odd cycle), by construction, $\sum_{e \in E | e \ni v} f(e) = 1$. Thus, f is a Fractional Perfect Matching.

Fix now a partite E_j. Since f is Perfect, $\sum_{e \in E_j} f(e) = \frac{1}{2} \sum_{v \in V(E_j)} 1 = \frac{1}{2} \cdot |V(E_j)| = \frac{|V|}{2\,\delta}$, by Condition (C2/b). Hence, f is δ-Partitionable. ■

We observe an interesting special case of Proposition 2.10:

Proposition 2.11 *A graph G has a $\frac{|V|}{2}$-Partitionable Fractional Perfect Matching if and only if G is Perfect-Matching.*

Proof. Assume first that G has a $\frac{|V|}{2}$-Partitionable Fractional Perfect Matching f. By Proposition 2.6, f can be partitioned into $\frac{|V|}{2}$ vertex-disjoint partites E_j with $|V(E_j)| = \frac{|V|}{|V|/2} = 2$. So, each partite E_j is a single edge and f is a Perfect Matching.

Assume that G has a Perfect Matching M. Consider the *indicator function* $f : E \to \{0,1\}$ for M, where $f(e) = 1$ if and only if $e \in M$; so, f is a Fractional Perfect Matching and it remains to show that f is $\frac{|V|}{2}$-Partitionable. For each edge $e_j \in M$, define the partite $E_j := \{e_j\}$. Since M is a Perfect Matching, the partites are vertex-disjoint. So, for each partite E_j, $\sum_{e \in E_j} f(e) = 1 = \frac{|V|}{2\,\frac{|V|}{2}}$, and this completes the proof. ■

2.3.3 Complexity

We define a natural decision problem about δ-Partitionable Fractional Perfect Matchings:

δ-PARTITIONABLE FPM
INSTANCE: A graph $G = \langle V, E\rangle$ and an integer δ which divides $|V|$.
QUESTION: Is there a δ-Partitionable Fractional Perfect Matching for G?

The restriction to instances for which δ divides $|V|$ is inherited from Corollary 2.7 to exclude the non-interesting instances.

Proposition 2.11 identifies a tractable special case of δ-PARTITIONABLE FPM (namely, $\frac{|V|}{2}$-PARTITIONABLE FPM). We shall show that for an arbitrary δ, δ-PARTITIONABLE FPM is $\mathcal{NP}$-complete. Towards that end, we observe the coincidence of another special case of the problem to a well known, intractable graph-theoretic problem:

PARTITION INTO TRIANGLES
INSTANCE: A graph $G = \langle V, E\rangle$ with $|V| = 3\delta$ for some integer δ.
QUESTION: Can V be partitioned into δ disjoint vertex sets $V_1, \cdots, V_\delta$, each containing exactly three vertices, such that for each V_j, $E(V_j)$ is a triangle?

The restriction to instances for which $|V| = 3\delta$ is made to exclude the non-interesting instances. This problem is known to be $\mathcal{NP}$-complete [6, GT11, attribution to (personal communication with) Schaefer]. To prove that δ-PARTITIONABLE FPM is $\mathcal{NP}$-complete for an arbitrary δ, it suffices to consider the special case with $\delta = \frac{|V|}{3}$:

$\frac{|V|}{3}$-PARTITIONABLE FPM
INSTANCE: A graph $G = \langle V, E\rangle$ with $|V| = 3\delta$ for some integer δ.
QUESTION: Is there a $\frac{|V|}{3}$-Partitionable Fractional Perfect Matching for G?

The restriction to instances with $|V| = 3\delta$ expresses the assumption that $\delta = \frac{|V|}{3}$. To prove that this special case is intractable, we prove that it coincides with PARTITION INTO TRIANGLES: it incurs an identical set of positive instances. We prove:

Proposition 2.12 $\frac{|V|}{3}$-PARTITIONABLE FPM $\equiv$ PARTITION INTO TRIANGLES

Proof. Consider a graph $G = \langle V, E\rangle$ with $|V| = 3\delta$ for some integer δ. Assume first that G is a positive instance for $\frac{|V|}{3}$-PARTITIONABLE FPM. By Proposition 2.10, E contains a collection of $\frac{|V|}{3}$ (non-empty,) vertex-disjoint edge sets $E_1, \cdots, E_{\frac{|V|}{3}}$ such that (C1) $\bigcup_{j\in\left[\frac{|V|}{3}\right]} E_j$ is an Edge Cover, and (C2) each edge set E_j consists of single edges and odd cycles with $|V(E_j)| = 3$.

It follows that each edge set E_j is a triangle. This implies that G is a positive instance for PARTITION INTO TRIANGLES (with vertex sets $V(E_1), \cdots, V\left(E_{\frac{|V|}{3}}\right)$).

Assume now that G is a positive instance for PARTITION INTO TRIANGLES. Consider the corresponding partition of V into $\delta = \frac{|V|}{3}$ disjoint vertex sets $V_1, \cdots, V_{\frac{|V|}{3}}$. This partition induces a corresponding partition of E into a collection of $\frac{|V|}{3}$ vertex-disjoint partites $E_1, \cdots, E_{\frac{|V|}{3}}$, where each partite E_j is a single triangle. Proposition 2.10 implies that G has a $\frac{|V|}{3}$ -Partitionable Fractional Perfect Matching. Hence, G is a positive instance for $\frac{|V|}{3}$-PARTITIONABLE FPM, and we are done. ■

By Proposition 2.12, it follows that $\frac{|V|}{3}$-PARTITIONABLE FPM is $\mathcal{NP}$-complete. Since $\frac{|V|}{3}$-PARTITIONABLE FPM is a special case of δ-PARTITIONABLE FPM, this implies:

Corollary 2.13 δ-PARTITIONABLE FPM *is $\mathcal{NP}$-complete.*

3 A Combinatorial Lemma

For a *probability* x, we define two *probability literals,* or *literals* for short: the *positive* x and the *negative* literal $\bar{x} = 1 - x$. A *probability product*, or *product* for short, is a product of probability literals $x_1 \cdots x_n$ for any $n \geq 1$; we adopt the convention that $x_{\ell_1} \cdots x_{\ell_2} = 1$ whenever $\ell_2 < \ell_1$. A *constant* probability product is the trivial one with no literals that equals to 1.

The probability product $x_1 \cdots x_n$ is *positive* if all its probability literals are positive. More generally, for any integer $\ell \leq n$, the probability product $x_1 \cdots x_n$ is *ℓ-positive* if exactly ℓ of its literals are positive; so, an n-positive probability product $x_1 \ldots x_n$ is a positive probability product. For each $\ell \in [n]$, denote as $\mathsf{Pos}_\ell(x_1, \ldots, x_n)$ the collection of all ℓ-positive probability products with literals defined from the probabilities $x_1, \ldots, x_n$.

The *expansion* of a probability product is obtained when substituting each negative literal $\overline{x}$ with $1 - x$. So, the expansion is an algebraic sum of positive probability products where the coefficient of each included positive product is $+1$ or -1. Note that the expansion of a k-positive product $x_1 \cdots x_n \in \mathsf{Pos}_k(x_1, \ldots, x_n)$ with $k > \ell$ cannot include $x_1 \cdots x_\ell$. We prove:

Lemma 3.1 *For each integer $n \geq 2$,*

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

Note that the right-hand side (RHS) is the weighted sum of sums of $(\ell-1)$-positive probability products over the literals $x_2, \cdots, x_\ell$, with $\ell \in [n]$, with weights $\frac{1}{\ell}$ and signs alternating with ℓ: the sign is $+1$ for odd ℓ and -1 for even ℓ. In contrast, the left-hand side (LHS) is the weighted sum of sums of $(\ell-1)$-positive probability products over the literals $x_2, \cdots, x_n$, with $\ell \in [n]$, with positive weight: each summed $(\ell-1)$-positive product in the LHS is multiplied by $\frac{1}{\ell}$. So, the RHS contains only positive probability products, while the LHS contains arbitrary, not necessarily positive, products.

Proof. It suffices to establish that for each $\ell \in [n]$, each (positive) probability product $x_2 \cdots x_\ell \in \mathsf{Pos}_{\ell-1}(x_2, \ldots, x_n)$ from the RHS appears in the sum of expansions of probability product in the LHS with the same coefficient. We proceed by case analysis on ℓ.

- Assume first that $\ell = 1$, and fix any product $x_2 \cdots x_\ell \in \mathsf{Pos}_{\ell-1}(x_2, \ldots, x_n)$ with $\ell = 1$ in the RHS. By convention, there is only one such product and it is 1. The coefficient of this product in the RHS is $(-1)^{1-1} \cdot \frac{1}{1} = 1$.

  In the LHS, the only constant term is the constant term in the sum

$$\left. \sum_{x_2 \cdots x_n \in \mathsf{Pos}_{\ell-1}(x_2,\ldots,x_n)} x_2 \cdots x_n \right|_{\ell=1} \;=\; \overline{x}_2 \cdots \overline{x}_n .$$

  Clearly, this constant term is 1 and its coefficient is $\frac{1}{1} = 1$. The claim follows for $\ell = 1$.

- Assume now that $\ell \geq 2$, and fix any product $x_2 \cdots x_\ell \in \mathsf{Pos}_{\ell-1}(x_2, \ldots, x_n)$ from the sum $\sum_{x_2 \cdots x_\ell \in \mathsf{Pos}_{\ell-1}(x_2,\ldots,x_n)} x_2 \cdots x_\ell$ with $\ell \geq 2$ in the RHS. Clearly, all products in $\mathsf{Pos}_{\ell-1}(x_2, \ldots, x_n)$ (in the RHS) have the same coefficient, which is $(-1)^{\ell-1} \cdot \frac{1}{\ell}$. We calculate the coefficient of this particular product in the LHS.

  We only need to consider contributions from the expansions of k-positive products with $0 \leq k \leq \ell - 1$ (in the LHS) to the coefficient of the product $x_2 \cdots x_\ell$. Consider such a k-positive product $x_2 \ldots x_n$.

  - All literals $x_{\ell+1}, \ldots, x_n$ have to be negative in the product $x_2 \ldots x_n$ since they do not appear in the product $x_2 \ldots x_\ell$.
  - There are $\binom{\ell-1}{k}$ ways to choose k positive literals out of the $(\ell-1)$ literals $x_2, \ldots, x_\ell$ in the product $x_2 \ldots x_n$ to form a k-positive product including $x_2 \cdots x_\ell$ in its expansion.
  - The sign of the resulting k-positive product is $(-1)^{(\ell-1)-k}$, since each of the remaining $(\ell - 1) - k$ negative literals contributes one minus sign. (The negative literals $x_{\ell+1}, \ldots, x_n$ do not contribute to the sign.).
  - The absolute value of the coefficient of the resulting k-positive product is $\frac{1}{k+1}$.

So, the coefficient of $x_2 \cdots x_\ell$ in the sum of expansions in the LHS is the sum

$$\begin{aligned}
\sum_{0 \le k \le \ell-1} \binom{\ell-1}{k} (-1)^{(\ell-1)-k} \frac{1}{k+1} &= \sum_{0 \le k \le \ell-1} \binom{\ell-1}{k} (-1)^{k} \frac{1}{\ell-k} \\
&= \frac{1}{\ell} \sum_{0 \le k \le \ell-1} \binom{\ell}{k} (-1)^{k} \\
&= \frac{1}{\ell} (-1)^{\ell-1},
\end{aligned}$$

as needed. ■

4 Game-Theoretic Framework

4.1 The Strategic Game $\mathsf{AD}_{\alpha,\delta}(G)$

Fix integers $\alpha \geq 1$ and $\delta \geq 1$. Associated with a graph G is a ***game*** $\mathsf{AD}_{\alpha,\delta}(G)$:

- The set of ***players*** is $\mathcal{A} \cup \mathcal{D}$; there are α ***attackers*** $\mathsf{a}_i \in \mathcal{A}$ and δ ***defenders*** $\mathsf{d}_j \in \mathcal{D}$.
- The ***strategy set*** S_a of an attacker a is V; the ***strategy set*** S_d of a defender d is E.

  A ***(pure) profile*** is an $(\alpha+\delta)$-tuple $\mathbf{s} = \langle s_{\mathsf{a}_1}, \ldots, s_{\mathsf{a}_\alpha}, s_{\mathsf{d}_1}, \ldots, s_{\mathsf{d}_\delta} \rangle$; the profile $\mathbf{s}_{-\mathsf{b}} \diamond t_\mathsf{b}$ is obtained from $\mathbf{s}$ and a strategy t_b for player b by substituting t_b for s_b in $\mathbf{s}$.

  For a vertex $v \in V$, $\mathsf{A}_\mathbf{s}(v) = \{\mathsf{a} \in \mathcal{A} \mid s_\mathsf{a} = v\}$ and $\mathsf{D}_\mathbf{s}(v) = \{\mathsf{d} \in \mathcal{D} \mid s_\mathsf{d} \ni v\}$. For a vertex $v \in s_\mathsf{d}$, the ***proportion*** $\mathsf{Prop}_\mathbf{s}(\mathsf{d}, v)$ of defender d on v in $\mathbf{s}$ is $\mathsf{Prop}_\mathbf{s}(\mathsf{d}, v) = \dfrac{1}{|\mathsf{D}_\mathbf{s}(v)|}$.

  - The ***Utility*** of attacker a is a function $\mathsf{U}_\mathsf{a} : S \to \{0,1\}$ with

    $$\mathsf{U}_\mathsf{a}(\mathbf{s}) = \begin{cases} 0, & \text{if } s_\mathsf{a} \in s_\mathsf{d} \text{ for some defender } \mathsf{d} \in \mathcal{D} \\ 1, & \text{if } s_\mathsf{a} \notin s_\mathsf{d} \text{ for every defender } \mathsf{d} \in \mathcal{D} \end{cases}.$$

    So, an attacker receives 0 if it is caught by a defender; else she receives 1.
  - The ***Utility*** of defender d is a function $\mathsf{U}_\mathsf{d} : S \to \mathbb{Q}$ with

    $$\mathsf{U}_\mathsf{d}(\mathbf{s}) = \frac{|\mathsf{A}_\mathbf{s}(u)|}{|\mathsf{D}_\mathbf{s}(u)|} + \frac{|\mathsf{A}_\mathbf{s}(v)|}{|\mathsf{D}_\mathbf{s}(v)|},$$

    where $s_\mathsf{d} = (u, v)$. So, the defender d receives the ***fair share*** of the total number of attackers choosing each of the two vertices of the edge it chooses.

4.2 Pure Nash Equilibria, Mixed Profiles and Supports

The profile s is a ***Pure Nash equilibrium*** [22, 23] if for each player $\mathsf{b} \in \mathcal{A} \cup \mathcal{D}$, for each strategy $t_{\mathsf{b}} \in S_{\mathsf{b}}$, $\mathsf{U}_{\mathsf{b}}(\mathbf{s}) \geq \mathsf{U}_{\mathsf{b}}(\mathbf{s}_{-\mathsf{b}} \diamond t_{\mathsf{b}})$; so, a Pure Nash equilibrium is a local maximizer for the Utility of each player. Say that G ***admits a Pure Nash equilibrium***, or G is ***Pure***, if there is a Pure Nash equilibrium for the strategic game $\mathsf{AD}_{\alpha,\delta}(G)$.

A ***mixed strategy*** for a player is a probability distribution over her strategy set. A **(*mixed*) *profile*** $\boldsymbol{\sigma} = \langle \sigma_{\mathsf{a}_1}, \ldots, \sigma_{\mathsf{a}_\alpha}, \sigma_{\mathsf{d}_1}, \ldots, \sigma_{\mathsf{d}_\delta} \rangle$ is a collection of mixed strategies, one for each player; $\boldsymbol{\sigma}$ induces naturally a probability measure $\mathbb{P}_{\boldsymbol{\sigma}}$ on profiles. The ***support*** $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{b})$ of player b in the mixed profile $\boldsymbol{\sigma}$ is the set of pure strategies in S_{b} to which $\boldsymbol{\sigma}$ assigns strictly positive probability. Denote $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A}) = \bigcup_{\mathsf{a} \in \mathcal{A}} \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$ and $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}) = \bigcup_{\mathsf{d} \in \mathcal{D}} \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. So, in a profile $\mathbf{s}$, $|\mathsf{Supports}_{\mathbf{s}}(\mathcal{A})| \leq \alpha$ and $|\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})| \leq \delta$.

4.2.1 Expectations about Attackers

For each vertex $v \in V$, denote as $|\mathsf{A}|_{\boldsymbol{\sigma}}(v)$ the expected number of attackers choosing vertex v in $\boldsymbol{\sigma}$; so, $|\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \sum_{\mathsf{a} \in \mathcal{A}} \sigma_{\mathsf{a}}(v)$. For an edge $(u,v) \in E$, denote $|\mathsf{A}|_{\boldsymbol{\sigma}}((u,v)) = |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + |\mathsf{A}|_{\boldsymbol{\sigma}}(v)$. By interchanging the order of summation, we obtain:

Observation 4.1 *For a mixed profile* $\boldsymbol{\sigma}$, $\sum_{v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} |\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \alpha$.

4.2.2 Hitting Events and Vertices

Fix a vertex $v \in V$. For a defender d, denote as $\mathsf{Hit}(\mathsf{d}, v)$ the event that defender d chooses an edge incident to vertex v; clearly, for the mixed profile $\boldsymbol{\sigma}$,

$$\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) = \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}) | e \ni v} \sigma_{\mathsf{d}}(e).$$

Denote as $\mathsf{Hit}(v)$ the event that *some* defender chooses an edge incident to vertex v; so, $\mathsf{Hit}(v) = \bigcup_{\mathsf{d} \in \mathcal{D}} \mathsf{Hit}(\mathsf{d}, v)$. Finally, denote as $\mathsf{D}_{\boldsymbol{\sigma}}(v)$ the set of defenders "hitting" vertex v; so,

$$\mathsf{D}_{\boldsymbol{\sigma}}(v) = \Big\{ \mathsf{d} \in \mathcal{D} \mid \text{ there is an edge } e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}) \text{ such that } e \ni v \Big\};$$

A vertex $v \in V$ is ***multidefender*** in the profile $\boldsymbol{\sigma}$ if $|\mathsf{D}_{\boldsymbol{\sigma}}(v)| \geq 2$; that is, a multidefender vertex is "hit" by more than one defenders. A vertex $v \in V$ is ***unidefender*** in $\boldsymbol{\sigma}$ if $|\mathsf{D}_{\boldsymbol{\sigma}}(v)| \leq 1$; v is ***monodefender*** in $\boldsymbol{\sigma}$ if $|\mathsf{D}_{\boldsymbol{\sigma}}(v)| = 1$. So, for a unidefender (resp., monodefender) vertex v, there is at most (resp., exactly) one defender d with an edge $e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$ such that $e \ni v$; if there is such a defender, denote it as $\mathsf{d}_{\boldsymbol{\sigma}}(v)$, else, set $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_{\boldsymbol{\sigma}}(v), v)) := 0$.

A mixed profile $\boldsymbol{\sigma}$ is ***unidefender*** (resp., ***monodefender***) if every vertex $v \in V$ is unidefender (resp., monodefender) in $\boldsymbol{\sigma}$; else $\boldsymbol{\sigma}$ is ***multidefender***. So, for a unidefender (resp., monodefender) profile $\boldsymbol{\sigma}$, for each edge $e \in E$, there is at most (resp., exactly) one defender d such that $\sigma_{\mathsf{d}}(e) > 0$; if there is such a defender d, denote it as $\mathsf{d}_{\boldsymbol{\sigma}}(e)$, else set $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{d}_{\boldsymbol{\sigma}}(e), e) := 0$.

4.2.3 Hitting Probabilities

Since the events $\mathsf{Hit}(\mathsf{d}_j, v)$ and $\mathsf{Hit}(\mathsf{d}_{j'}, v)$ with $j \neq j'$ are independent and not mutually exclusive (for a fixed vertex v), we immediately obtain a strengthening of the Union Bound:

Observation 4.2 *For each vertex $v \in V$,*

$$\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \quad \begin{cases} < & \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)), \quad \textit{if } v \textit{ is multidefender in } \boldsymbol{\sigma} \\ = & \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)), \quad \textit{if } v \textit{ is unidefender in } \boldsymbol{\sigma} \end{cases}.$$

By the *Principle of Inclusion-Exclusion,* we immediately observe:

Lemma 4.1 *For a vertex $v \in V$,*

$$\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \;=\; \sum_{l \in [\delta]} (-1)^{l-1} \sum_{\mathcal{D}' \subseteq \mathcal{D} | |\mathcal{D}'| = l} \prod_{\mathsf{d} \in \mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}\left(\mathsf{Hit}(\mathsf{d}, v)\right).$$

We continue to prove:

Lemma 4.2 *For a mixed profile $\boldsymbol{\sigma}$,*

$$\sum_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \begin{cases} < \; 2\,\delta, & \textit{if } \boldsymbol{\sigma} \textit{ is multidefender} \\ = \; 2\,\delta, & \textit{if } \boldsymbol{\sigma} \textit{ is unidefender} \end{cases}.$$

Proof. Clearly,

$$\begin{aligned} \sum_{v \in V} \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) &= \sum_{v \in V} \sum_{\mathsf{d} \in \mathcal{D}} \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}) | e \ni v} \sigma_{\mathsf{d}}(e) \\ &= 2 \sum_{e \in E} \sum_{\mathsf{d} \in \mathcal{D}} \sigma_{\mathsf{d}}(e) \\ &= 2 \sum_{\mathsf{d} \in \mathcal{D}} \sum_{e \in E} \sigma_{\mathsf{d}}(e) \\ &= 2\,\delta\,. \end{aligned}$$

Hence, by Observation 4.2,

$$\sum_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \quad \begin{cases} < \sum_{v \in V} \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)), & \text{if } \boldsymbol{\sigma} \text{ is multidefender} \\ = \sum_{v \in V} \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)), & \text{if } \boldsymbol{\sigma} \text{ is unidefender} \end{cases}$$

$$\begin{cases} < 2\,\delta, & \text{if } \boldsymbol{\sigma} \text{ is multidefender} \\ = 2\,\delta, & \text{if } \boldsymbol{\sigma} \text{ is unidefender} \end{cases},$$

as needed. ■

4.2.4 Minimum Hitting Probability, Maxhit Vertices and Maxhitters

Denote as $\mathsf{MinHit}_{\boldsymbol{\sigma}} = \min_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$ the ***Minimum Hitting Probability*** associated with the mixed profile $\boldsymbol{\sigma}$. A simple proof by contradiction (appealing to Lemma 4.2) yields:

Lemma 4.3 *For a mixed profile* $\boldsymbol{\sigma}$, $\mathsf{MinHit}_{\boldsymbol{\sigma}} \leq \frac{2\delta}{|V|}$.

A vertex $v \in V$ is ***maxhit*** in the mixed profile $\boldsymbol{\sigma}$ if $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))=1$; a defender $\mathsf{d} \in \mathcal{D}$ is a ***maxhitter*** in $\boldsymbol{\sigma}$ if there is a vertex $v \in \mathsf{Vertices}(\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}))$ with $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v))=1$. We observe:

Lemma 4.4 *Consider a maxhit vertex* v *in a mixed profile* $\boldsymbol{\sigma}$*. Then, there is a (maxhitter) defender* d *(in* $\boldsymbol{\sigma}$*) with* $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) = 1$.

Proof. Assume, by way of contradiction, that for each defender $\mathsf{d} \in \mathcal{D}$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) < 1$. Since the set $\{\mathsf{Hit}(\mathsf{d}, v) \mid \mathsf{d} \in \mathcal{D}\}$ is a family of independent, non-full events, this implies that $\mathsf{Hit}(v) = \bigcup_{\mathsf{d} \in \mathcal{D}} \mathsf{Hit}(\mathsf{d}, v)$ is non-full. So, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) < 1$ and v is *not* maxhit. A contradiction. ■

4.3 Expected Utilities

The mixed profile $\boldsymbol{\sigma}$ induces an ***Expected Utility*** $\mathsf{U}_{\mathsf{b}}(\boldsymbol{\sigma})$ for each player b: the expectation of her Utility. We shall derive formulas for Expected Utilities. We first define and derive formulas for the Conditional Expected Proportions induced by the defenders; we then use it to derive an expression for the Conditional Expected Utility for each attacker. The Expected Utility of each attacker or defender is a weighted sum of Conditional Expected Utilities.

4.3.1 Conditional Expected Proportion

Induced by $\boldsymbol{\sigma}$ is the ***Conditional Expected Proportion*** $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v)$ of defender $\mathsf{d} \in \mathcal{D}$ on vertex v: the expectation (induced by $\boldsymbol{\sigma}$) of the proportion of defender d on vertex v had she chosen an edge incident to vertex v. Clearly,

$$\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) = \sum_{\ell \in [\delta]} \frac{1}{\ell} \sum_{\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\} | |\mathcal{D}'| = \ell - 1} \prod_{\mathsf{d}_k \in \mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)) \prod_{\mathsf{d}_k \notin \mathcal{D}' \cup \{\mathsf{d}\}} (1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)))$$

Lemma 3.1 implies now an alternative expression for Conditional Expected Proportion.

Lemma 4.5 *For each pair of a defender* $\mathsf{d} \in \mathcal{D}$ *and a vertex* $v \in V$*,*

$$\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) = \sum_{\ell \in [\delta]} \frac{1}{\ell} (-1)^{\ell - 1} \sum_{\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\} | |\mathcal{D}'| = \ell - 1} \prod_{\mathsf{d}_k \in \mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)) .$$

4.3.2 Attackers

Induced by $\boldsymbol{\sigma}$ is the ***Conditional Expected Utility*** $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond v)$ of attacker a on vertex v: the conditional expectation of the Utility of a had she chosen vertex v. Clearly, $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond v) = 1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$. By the Law of Conditional Alternatives, we immediately obtain:

Lemma 4.6 *Fix a mixed profile* $\boldsymbol{\sigma}$*. Then, the Expected Utility* $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma})$ *of an attacker* $\mathsf{a} \in \mathcal{A}$ *is*

$$\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}) = \sum_{v \in V} \sigma_{\mathsf{a}}(v) \cdot (1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))) .$$

We continue with a preliminary observation:

Lemma 4.7 *Fix a mixed profile* $\boldsymbol{\sigma}$*. Then, for each vertex* $v \in V$*,*

$$\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) .$$

Proof. By Lemma 4.5,

$$\begin{aligned}
& \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \\
= & \sum_{\mathsf{d} \in \mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \left(\sum_{\ell \in [\delta]} \frac{1}{\ell} (-1)^{\ell - 1} \sum_{\substack{\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\} \\ |\mathcal{D}'| = \ell - 1}} \prod_{\mathsf{d}_k \in \mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)) \right) \\
= & \sum_{\mathsf{d} \in \mathcal{D}} \sum_{\ell \in [\delta]} \frac{1}{\ell} (-1)^{\ell - 1} \sum_{\substack{\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\} \\ |\mathcal{D}'| = \ell - 1}} \prod_{\mathsf{d}_k \in \mathcal{D}' \cup \{\mathsf{d}\}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)) \\
= & \sum_{\ell \in [\delta]} (-1)^{\ell - 1} \cdot \frac{1}{\ell} \sum_{\mathsf{d} \in \mathcal{D}} \sum_{\substack{\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\} \\ |\mathcal{D}'| = \ell - 1}} \prod_{\mathsf{d}_k \in \mathcal{D}' \cup \{\mathsf{d}\}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)).
\end{aligned}$$

For each integer $\ell \in [\delta]$, for each set $\mathcal{D}'' \subseteq \mathcal{D}$ with $|\mathcal{D}''| = \ell$, there are ℓ pairs of a defender $\mathsf{d} \in \mathcal{D}$ such that $\mathsf{d} \in \mathcal{D}''$ and a set $\mathcal{D}' \subseteq \mathcal{D}''$ such that $\mathcal{D}' \subseteq \mathcal{D} \setminus \{\mathsf{d}\}$ and $|\mathcal{D}'| = \ell - 1$. So,

$$\sum_{\substack{\mathcal{D}''\subseteq\mathcal{D}\\|\mathcal{D}''|=\ell}} \prod_{\mathsf{d}_k\in\mathcal{D}''} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)) = \frac{1}{\ell} \sum_{\mathsf{d}\in\mathcal{D}} \sum_{\substack{\mathcal{D}'\subseteq\mathcal{D}\setminus\{\mathsf{d}\}\\|\mathcal{D}'|=\ell-1}} \prod_{\mathsf{d}_k\in\mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v)).$$

Hence, by Lemma 4.1,

$$\begin{aligned}\sum_{\mathsf{d}\in\mathcal{D}} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) &= \sum_{\ell\in[\delta]} (-1)^{\ell-1} \sum_{\substack{\mathcal{D}'\subseteq\mathcal{D}\\|\mathcal{D}'|=\ell}} \prod_{\mathsf{d}_k\in\mathcal{D}'} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_k, v))\\ &= \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)),\end{aligned}$$

as needed. ■

4.3.3 Defenders

Induced by $\boldsymbol{\sigma}$ is the ***Conditional Expected Utility*** $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v))$ of defender d on edge $(u,v) \in E$: the conditional expectation (induced by $\boldsymbol{\sigma}$) of the Utility of d had she chosen (u,v). So,

$$\mathsf{U}_{\mathsf{d}}((\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v))) = \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)\,.$$

We prove:

Lemma 4.8 *Fix a mixed profile $\boldsymbol{\sigma}$. Then, the Expected Utility of a defender $\mathsf{d} \in \mathcal{D}$ is*

$$\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = \sum_{v\in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)\,.$$

Proof. By the Law of Conditional Alternatives,

$$\begin{aligned}\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) &= \sum_{(u,v)\in E} \sigma_{\mathsf{d}}((u, v)) \cdot \mathsf{U}_{\mathsf{d}}((\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)))\\ &= \sum_{(u,v)\in E} \sigma_{\mathsf{d}}((u, v)) \cdot \Big(\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)\Big)\\ &= \sum_{v\in V} \left(\sum_{e|v\in e} \sigma_{\mathsf{d}}(e)\right) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)\\ &= \sum_{v\in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)\,,\end{aligned}$$

as needed. ■

4.4 Nash Equilibria

A mixed profile $\boldsymbol{\sigma}$ is a ***Nash equilibrium*** [22, 23] if for each player b, for each mixed strategy τ_{b} of player b, $\mathsf{U}_{\mathsf{b}}(\boldsymbol{\sigma}) \geq \mathsf{U}_{\mathsf{b}}(\boldsymbol{\sigma}_{-\mathsf{b}} \diamond \tau_{\mathsf{b}})$; so, a Nash equilibrium is a local maximizer of the Expected Utility of each player. A (necessary and) sufficient condition for a Nash equilibrium $\boldsymbol{\sigma}$ is that for each player b, for each pure strategy t_{b} of player b, $\mathsf{U}_{\mathsf{b}}(\boldsymbol{\sigma}) \geq \mathsf{U}_{\mathsf{b}}(\boldsymbol{\sigma}_{-\mathsf{b}} \diamond t_{\mathsf{b}})$. By Nash's Theorem [22, 23], $\mathsf{AD}_{\alpha,\delta}(G)$ has at least one Nash equilibrium.

Clearly, in a Nash equilibrium $\boldsymbol{\sigma}$, for each attacker a, $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond v)$ is *constant* over all vertices $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$; for each defender d, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond e)$ is *constant* over all edges $e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Hence, in a Nash equilibrium $\boldsymbol{\sigma}$, for each attacker a, $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}) = 1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$ for any vertex $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$. So, for each attacker $\mathsf{a} \in \mathcal{A}$, the quantity $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$ is *constant* over all vertices $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$. In the same way, for each defender d,

$$\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) \quad = \quad \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)$$

for any edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. So, for each defender d, the quantity $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)$ is *constant* over all edges $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Note that in a Nash equilibrium $\boldsymbol{\sigma}$, for each defender $\mathsf{d} \in \mathcal{D}$, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) > 0$; in contrast, it is possible that $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}) = 0$ for some attacker $\mathsf{a} \in \mathcal{A}$. (For an example where this occurs, see the proof of Theorem 9.2.)

4.5 Some Special Profiles

A mixed profile $\boldsymbol{\sigma}$ is ***uniform*** if each player uses a *uniform* distribution on its support; so, for each attacker a (resp., defender d) for each vertex $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$ (resp., for each edge $e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$), $\sigma_{\mathsf{a}}(v) = \frac{1}{|\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})|}$ $\left(\text{resp., } \sigma_{\mathsf{d}}(e) = \frac{1}{|\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})|}\right)$.

A mixed profile $\boldsymbol{\sigma}$ is ***attacker-symmetric*** (resp., ***defender-symmetric***) if for all pairs of attackers a_i and a_k (resp., all pairs of defenders d_j and d_l), for all vertices $v \in V$ (resp., all edges $e \in E$) $\sigma_{\mathsf{a}_i}(v) = \sigma_{\mathsf{a}_k}(v)$ (resp., $\sigma_{\mathsf{d}_j}(e) = \sigma_{\mathsf{d}_l}(e)$). A mixed profile is ***attacker-uniform*** (resp., ***defender-uniform***) if each attacker (resp., defender) uses a uniform probability distribution on her support. Now, ***Attacker-Symmetric*** (resp., ***Defender-Symmetric***) ***Nash equilibria*** and ***Attacker-Uniform*** (resp., ***Defender-Uniform***) ***Nash equilibria*** are defined in the natural way. A ***Symmetric Nash equilibrium*** is both Attacker-Symmetric and Defender-Symmetric. A ***Uniform Nash equilibrium*** is both Attacker-Uniform and Defender-Uniform.

A mixed profile $\boldsymbol{\sigma}$ is ***attacker-fullymixed*** (resp., ***defender-fullymixed***) if for each attacker a (resp., defender d), $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a}) = V$ (resp., $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}) = E$). Now, ***Attacker Fullymixed*** (resp., ***Defender-Fullymixed***) ***Nash equilibria*** are defined in the natural way. A ***Fullymixed Nash equilibrium*** is both Attacker-Fullymixed and Defender-Fullymixed.

A mixed profile $\boldsymbol{\sigma}$ is ***defender-pure*** if each defender chooses a single edge with probability 1. Now ***Defender-Pure Nash equilibria*** are defined in the natural way. Say that G is ***Defender-Pure***, if there is a Defender-Pure Nash equilibrium for the strategic game $\mathsf{AD}_{\alpha,\delta}(G)$.

Fix a Perfect-Matching graph. A profile is ***perfect-matching*** if $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ is a Perfect Matching. Now ***Perfect-Matching Nash equilibria*** are defined in the natural way.

4.6 Notation

Fix a mixed profile $\boldsymbol{\sigma}$. For a vertex $v \in V$, set $\mathsf{Edges}_{\boldsymbol{\sigma}}(v) := \{e \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}) \mid e \ni v\}$. For a vertex set $U \subseteq V$, set $\mathsf{Edges}_{\boldsymbol{\sigma}}(U) := \bigcup_{v \in U} \mathsf{Edges}_{\boldsymbol{\sigma}}(v)$. For an edge $e \in E$, set $\mathsf{Vertices}_{\boldsymbol{\sigma}}(e) := \{v \in e \mid v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})\}$; so, $|\mathsf{Vertices}_{\boldsymbol{\sigma}}(e)| \leq 2$. For an edge set $F \subseteq E$, set $\mathsf{Vertices}_{\boldsymbol{\sigma}}(F) := \bigcup_{e \in F} \mathsf{Vertices}_{\boldsymbol{\sigma}}(e)$.

5 The Structure of Nash Equilibria

5.1 Combinatorial Characterization

We show:

Proposition 5.1 (Characterization of Nash Equilibria) *A mixed profile $\boldsymbol{\sigma}$ is a Nash equilibrium if and only if the following conditions hold:*

(C1) *For each vertex $v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \mathsf{MinHit}_{\boldsymbol{\sigma}}$.*

(C2) *For each defender $\mathsf{d} \in \mathcal{D}$, for each edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$,*

$$\begin{aligned} & \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) \\ = \quad & \max_{(u',v') \in E} \left\{ \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v') \right\}. \end{aligned}$$

Proof. Assume first that $\boldsymbol{\sigma}$ is a Nash equilibrium. To establish (C1), consider any vertex $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$ for some attacker a, and note that $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v'))$ is constant over all vertices $v' \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$. Hence, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$ is constant over all vertices $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathcal{A})$.

Consider now another vertex $u \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathcal{A})$. Note that $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u) = 1 - \mathbb{P}_{\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u}(\mathsf{Hit}(u))$. Then, since $\boldsymbol{\sigma}$ is a Nash equilibrium, $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u) \leq \mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond v)$. It follows that $\mathbb{P}_{\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u}(\mathsf{Hit}(u)) \geq \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$. Note that, by construction, $\mathbb{P}_{\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u}(\mathsf{Hit}(u)) = \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(u))$. Thus, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(u)) \geq \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v))$ and (C1) follows.

We now prove (C2). Fix a defender d. Note that, by construction, for any edge $(u', v') \in E$, $|\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u',v')}(u') = |\mathsf{A}|_{\boldsymbol{\sigma}}(u')$. Consider an edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Since $\boldsymbol{\sigma}$ is a Nash equilibrium, the quantity $\mathsf{Prop}_{\boldsymbol{\sigma}}(\mathsf{d}, v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) + \mathsf{Prop}_{\boldsymbol{\sigma}}(\mathsf{d}, u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u)$ is constant over all edges $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. So, consider any edge $(u', v') \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$.

Since $\boldsymbol{\sigma}$ is a Nash equilibrium, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) \leq \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v))$. Thus, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')}(v') \leq \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)}(u)$ $+\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)}(v)$. Since $|\mathsf{A}|_{\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')}(u') = |\mathsf{A}|_{\boldsymbol{\sigma}}(u')$, it follows that $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v') \leq \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u, v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v)$, as required for (C2).

Assume now that the mixed profile $\boldsymbol{\sigma}$ satisfies (C1) and (C2). Consider first an attacker $\mathsf{a} \in \mathcal{A}$. For any pair of vertices $u \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$ and $v \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a})$, by (C1), $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}) = 1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \geq 1 - \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(u))$, so that $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}) \geq \mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond u)$. Consider now a defender $\mathsf{d} \in \mathcal{D}$. For any pair of edges $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$ and $(u', v') \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$, by (C2), $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) \geq \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v')$, so that $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) \geq \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v'))$. Hence, $\boldsymbol{\sigma}$ is a Nash equilibrium. ■

We remark that Proposition 5.1 generalizes a corresponding characterization of Nash equilibria for $\mathsf{AD}_{\alpha,1}(G)$ of [20, Theorem 3.1], differing only in Condition (C2) which is simpler there.

5.2 Necessary Conditions

We establish necessary conditions for Nash equilibria, following from Proposition 5.1. We first prove a simple expression for the *total* Expected Utility of the defenders:

Proposition 5.2 *In a Nash equilibrium $\boldsymbol{\sigma}$, $\sum_{\mathsf{d} \in \mathcal{D}} \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = \alpha \cdot \mathsf{MinHit}_{\boldsymbol{\sigma}}$.*

Proof. Clearly,

$$
\begin{aligned}
&\sum_{\mathsf{d} \in \mathcal{D}} \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) \\
&= \textstyle\sum_{\mathsf{d} \in \mathcal{D}} \sum_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) \cdot \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) && \text{(by Lemma 4.8)} \\
&= \textstyle\sum_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) && \text{(by Lemma 4.7)} \\
&= \textstyle\sum_{v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) \\
&= \textstyle\sum_{v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} \mathsf{MinHit}_{\boldsymbol{\sigma}} \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) && \text{(by Proposition 5.1 (Condition (2)))} \\
&= \mathsf{MinHit}_{\boldsymbol{\sigma}} \cdot \textstyle\sum_{v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} |\mathsf{A}|_{\boldsymbol{\sigma}}(v) \\
&= \alpha \cdot \mathsf{MinHit}_{\boldsymbol{\sigma}} && \text{(by Observation 4.1),}
\end{aligned}
$$

as needed. ■

We continue to show:

Proposition 5.3 *In a Nash equilibrium $\boldsymbol{\sigma}$, $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ is an Edge Cover.*

Proof. Assume, by way of contradiction, that $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ is *not* an Edge Cover. Then, there is a vertex $v \in V$ such that $v \notin \mathsf{Vertices}(\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}))$. So, $\mathsf{Edges}_{\boldsymbol{\sigma}}(v) = \emptyset$ and $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 0$. Fix an attacker $\mathsf{a} \in \mathcal{A}$. Since $\boldsymbol{\sigma}$ is a local maximizer for the Expected Utility of a, which is at most 1, it follows that $\sigma_{\mathsf{a}}(v) = 1$. Hence, for each $(u', v') \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$, $|\mathsf{A}|_{\boldsymbol{\sigma}}((u', v')) = 0$, since both $u' \neq u$ and $v' \neq v$ (by the choice of vertex v). So, $|\mathsf{A}|_{\boldsymbol{\sigma}}(u') = |\mathsf{A}|_{\boldsymbol{\sigma}}(v') = 0$. Thus, for any defender $\mathsf{d} \in \mathcal{D}$, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = 0$. Since $\boldsymbol{\sigma}$ is a Nash equilibrium, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) > 0$. A contradiction. ■

Proposition 5.3 immediately implies:

Corollary 5.4 *A unidefender Nash equilibrium is monodefender.*

We finally show:

Proposition 5.5 *In a Nash equilibrium $\boldsymbol{\sigma}$, $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$ is a Vertex Cover of the graph $G(\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}))$.*

Proof. Assume, by way of contradiction, that $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathcal{A})$ is *not* a Vertex Cover of the graph $G(\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}))$. Then, there is an edge $(u, v) \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ such that both $u \notin \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$ and $v \notin \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$. So, $|\mathsf{A}|_{\boldsymbol{\sigma}}((u, v)) = 0$. Assume that $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$ for some defender $\mathsf{d} \in \mathcal{D}$. Since $\boldsymbol{\sigma}$ is a local maximizer for the Expected Utility of defender d, it follows that $\sigma_{\mathsf{d}}((u, v)) = 0$. So, $(u, v) \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. A contradiction. ■

5.3 Pure Nash Equilibria

We observe that for the special case of Pure Nash equilibria, Proposition 5.1 simplifies to:

Proposition 5.6 (Characterization of Pure Nash Equilibria) *A profile $\mathbf{s}$ is a Pure Nash equilibrium if and only if the following conditions hold:*

(C1) $\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})$ *is an Edge Cover.*

(C2) *For each attacker $\mathsf{d} \in \mathcal{D}$, for each edge $(u, v) \in \mathsf{Support}_{\mathbf{s}}(\mathsf{d})$,*

$$\frac{|\mathsf{A}_{\mathbf{s}}(u)|}{|\mathsf{D}_{\mathbf{s}}(u)|} + \frac{|\mathsf{A}_{\mathbf{s}}(v)|}{|\mathsf{D}_{\mathbf{s}}(v)|} \quad = \quad \max_{(u',v') \in E} \left\{ \frac{|\mathsf{A}_{\mathbf{s}}(u')|}{|\mathsf{D}_{\mathbf{s}_{-j}}(u')| + 1} + \frac{|\mathsf{A}_{\mathbf{s}}(v')|}{|\mathsf{D}_{\mathbf{s}_{-j}}(v')| + 1} \right\} .$$

We shall now use Propositions 5.3 and 5.5 to show:

Proposition 5.7 (Necessary Conditions for Pure Nash Equilibria) *Assume that G is Pure. Then,* (C1) $\delta \geq \beta'(G)$ *and* (C2) $\alpha \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$.

Proof. Consider a Pure Nash equilibrium $\mathbf{s}$. For (C1), assume, by way of contradiction, that $\delta < \beta'(G)$. Then, $|\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})| < \beta'(G)$. Hence, $\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})$ is *not* an Edge Cover. A contradiction to Proposition 5.3.

For (C2), assume, by way of contradiction, that $\alpha < \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. Since $\mathbf{s}$ is a pure profile, it follows that $|\mathsf{Supports}_{\mathbf{s}}(\mathcal{A})| < \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. By Proposition 5.3, $\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})$ is an Edge Cover; so, $\beta(G(\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}))) \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. It follows that $|\mathsf{Supports}_{\mathbf{s}}(\mathcal{A})| < \beta(G(\mathsf{Supports}_{\mathbf{s}}(\mathcal{D})))$. Thus, $\mathsf{Supports}_{\mathbf{s}}(\mathcal{A})$ is not a Vertex Cover of the graph $G(\mathsf{Supports}_{\mathbf{s}}(\mathcal{D}))$. A contradiction to Proposition 5.5. ■

We remark that (C1) (resp., (C2)) in Proposition 5.7 is necessary for Defender-Pure (resp., Attacker-Pure) Nash equilibria. We report a counterexample to the converse of Proposition 5.7:

Proposition 5.8 *There is a graph G and integers α and δ such that* (C1) $\delta \geq \beta'(G)$ *and* (C2) $\alpha \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$ *while G is not Pure.*

6 Defense-Optimal Nash Equilibria

6.1 Definitions

The ***Defense-Ratio*** $\mathsf{DR}_{\boldsymbol{\sigma}}$ of a Nash equilibrium $\boldsymbol{\sigma}$ is the ratio of the *optimal* total Utility α of the defenders over their total Expected Utility in $\boldsymbol{\sigma}$; so, $\mathsf{DR}_{\boldsymbol{\sigma}} = \frac{\alpha}{\sum_{\mathsf{d} \in \mathcal{D}} \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma})}$. By the definition of the Defense-Ratio, Proposition 5.2 immediately implies:

Corollary 6.1 *For a Nash equilibrium $\boldsymbol{\sigma}$,* $\mathsf{DR}_{\boldsymbol{\sigma}} = \frac{1}{\mathsf{MinHit}_{\boldsymbol{\sigma}}}$.

Clearly, $\mathsf{DR}_{\boldsymbol{\sigma}} \geq 1$. Furthermore, Lemma 4.3 implies a second lower bound on the Defense-Ratio:

Corollary 6.2 *For a Nash equilibrium $\boldsymbol{\sigma}$,* $\mathsf{DR}_{\boldsymbol{\sigma}} \geq \frac{|V|}{2\,\delta}$.

Our next major definition encompasses these two lower bounds on the Defense-Ratio.

Definition 6.1 *A Nash equilibrium $\boldsymbol{\sigma}$ is* ***Defense-Optimal*** *if* $\mathsf{DR}_{\boldsymbol{\sigma}} = \max\left\{1, \frac{|V|}{2\,\delta}\right\}$.

The justification for the definition of a Defense-Optimal Nash equilibrium will come later, when constructing Defense-Optimal Nash equilibria in two cases (Proposition 7.9 and Theorems 9.2 and 9.5); these constructions will establish that $\max\left\{1, \frac{|V|}{2\delta}\right\}$ is a *tight* lower bound on Defense-Ratio. Say that G is ***Defense-Optimal*** if G admits a Defense-Optimal Nash equilibrium.

6.2 Sufficient Conditions

We show:

Theorem 6.3 *Assume that G has a δ-Partitionable Fractional Perfect Matching. Then, G is Defense-Optimal.*

Proof. Consider a δ-Partitionable Fractional Perfect Matching f and the corresponding (non-empty) partites $E_1, \cdots, E_\delta$. Recall that $E(f)$ is an Edge Cover. Construct $\boldsymbol{\sigma}$ as follows:

- For each attacker $\mathsf{a} \in \mathcal{A}$, for each vertex $v \in V$, set $\sigma_{\mathsf{a}}(v) := \frac{1}{|V|}$; so, $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{a}) = V$. So, for each vertex $v \in V$, $|\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \frac{\alpha}{|V|}$.
- For each defender $\mathsf{d}_j \in \mathcal{D}$, for each edge $e \in E$, set $\sigma_{\mathsf{d}_j}(e) := \frac{2\delta}{|V|} \cdot f(e)$ if $e \in E_j$, and 0 otherwise; so, $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) = E_j$ and all values of σ_{d_j} are non-negative.

So, $\boldsymbol{\sigma}$ is attacker-symmetric, attacker-uniform, attacker-fullymixed and defender-symmetric; moreover, $\boldsymbol{\sigma}$ is monodefender. Furthermore, for each vertex $v \in V$, $\mathsf{Edges}_{\boldsymbol{\sigma}}(v) = \{e \in E(f) \mid v \in e\}$. We now prove that for each defender d_j, σ_{d_j} is a distribution (on E): by construction and the assumption that f is δ-Partitionable, $\sum_{e \in E} \sigma_{\mathsf{d}_j}(e) = 1$.

We now prove that $\boldsymbol{\sigma}$ is a Nash equilibrium. We shall verify $(\mathsf{C1})$ and $(\mathsf{C2})$ in Proposition 5.1. For $(\mathsf{C1})$, fix a vertex $v \in V$. Since $E(f)$ is an Edge Cover, there is a partite $E_j \subseteq E(f)$ such that $v \in \mathsf{Vertices}(E_j)$. Since the partites $E_1, \cdots, E_\delta$ are vertex-disjoint and $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) = E_j$, it follows that vertex v is monodefender in $\boldsymbol{\sigma}$ with $\mathsf{d}_{\boldsymbol{\sigma}}(v) = \mathsf{d}_j$. We prove:

Claim 6.4 $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \frac{2\delta}{|V|}$.

Proof. Clearly,

$$
\begin{array}{lll}
 & \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) & \\
= & \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_j, v)) & \text{(since } v \text{ is monodefender in } \boldsymbol{\sigma}) \\
= & \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) \mid v \in e} \sigma_{\mathsf{d}_j}(e) & \\
= & \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) \mid v \in e} \frac{2\delta}{|V|} \cdot f(e) & \text{(by construction of } \boldsymbol{\sigma}) \\
= & \frac{2\delta}{|V|} \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) \mid v \in e} f(e) & \\
= & \frac{2\delta}{|V|} \sum_{e \in \mathsf{Edges}_{\boldsymbol{\sigma}}(v)} f(e) & \text{(since } v \text{ is monodefender in } \boldsymbol{\sigma}) \\
= & \frac{2\delta}{|V|} \sum_{e \in E(f) \mid e \ni v} f(e) & \\
= & \frac{2\delta}{|V|} \cdot \sum_{e \in E \mid e \ni v} f(e) & \text{(since } f(e) = 0 \text{ for } e \notin E(f)) \\
= & \frac{2\delta}{|V|} \cdot 1 & \text{(since } f \text{ is a Fractional Perfect Matching),}
\end{array}
$$

as needed. ■

By Claim 6.4, (C1) holds trivially. For (C2), fix a defender d and an edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Since $\boldsymbol{\sigma}$ is monodefender, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) = \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) = 1$. Hence, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \frac{2\alpha}{|V|}$.

Fix now an edge $(u', v') \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Since $E(f)$ is an Edge Cover, there are edges $e_{u'}$ and $e_{v'} \in E_f$ such that $u' \in e_{u'}$ and $v' \in e_{v'}$. By the construction of $\boldsymbol{\sigma}$, this implies that there are defenders $\mathsf{d}_{u'}$ and $\mathsf{d}_{v'}$ such that $e_{u'} \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_{u'})$ and $e_{v'} \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_{v'})$. There are two cases for $\mathsf{d}_{u'}$ (resp., $\mathsf{d}_{v'}$): either $\mathsf{d}_{u'} = \mathsf{d}$ or $\mathsf{d}_{u'} \neq \mathsf{d}$ (resp., $\mathsf{d}_{v'} = \mathsf{d}$ or $\mathsf{d}_{v'} \neq \mathsf{d}$).

- Assume first that $\mathsf{d}_{u'} = \mathsf{d}$ (resp., $\mathsf{d}_{v'} = \mathsf{d}$); since u' is monodefender, it follows that $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') = 1$ (resp., $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') = 1$).
- Assume now that $\mathsf{d}_{u'} \neq \mathsf{d}$ (resp., $\mathsf{d}_{v'} \neq \mathsf{d}$); since v' is monodefender, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') < 1$ (resp., $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') < 1$).

So, in all cases, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \leq 1$ and $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') \leq 1$. Thus, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v') \leq |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + |\mathsf{A}|_{\boldsymbol{\sigma}}(v') = \frac{2\alpha}{|V|}$. Now, (C2) follows. Hence, by Proposition 5.1, $\boldsymbol{\sigma}$ is a Nash equilibrium. By Claim 6.4 and (C1) in Proposition 5.1, it follows that $\mathsf{MinHit}_{\boldsymbol{\sigma}} = \frac{2\delta}{|V|}$. By Corollary 6.1, it follows that $\mathsf{DR}_{\boldsymbol{\sigma}} = \frac{|V|}{2\delta}$. Since G has a δ-Partitionable Fractional Perfect Matching, Corollary 2.7 implies that $\delta \leq \frac{|V|}{2}$, so that $\max\left\{1, \frac{|V|}{2\,\delta}\right\} = \frac{|V|}{2\,\delta}$. This implies that $\mathsf{DR}_{\boldsymbol{\sigma}} = \max\left\{1, \frac{|V|}{2\,\delta}\right\}$ and $\boldsymbol{\sigma}$ is Defense-Optimal, as needed. ■

We continue with another sufficient condition:

Theorem 6.5 *Assume that G is Defender-Pure. Then, G is Defense-Optimal.*

Proof. Fix a Defender-Pure Nash equilibrium $\boldsymbol{\sigma}$. Since $\mathsf{MinHit}_{\boldsymbol{\sigma}} = 1$, Corollary 6.1 implies that $\mathsf{DR}_{\boldsymbol{\sigma}} = 1$. For a defender d, let $s_{\mathsf{d}} = (u_{\mathsf{d}}, v_{\mathsf{d}}) \in E$. Since $\boldsymbol{\sigma}$ is a Nash equilibrium,

$\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u_{\mathsf{d}}, v_{\mathsf{d}}))$. So, since $\boldsymbol{\sigma}$ is Defender-Pure,

$$
\begin{aligned}
\mathsf{DR}_{\boldsymbol{\sigma}} &= \frac{\alpha}{\sum_{\mathsf{d}\in\mathcal{D}} \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma})} \\
&= \frac{\alpha}{\sum_{\mathsf{d}\in\mathcal{D}} \mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u_{\mathsf{d}}, v_{\mathsf{d}}))} \\
&= \frac{\alpha}{\sum_{\mathsf{d}\in\mathcal{D}} \left(\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u_{\mathsf{d}}) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u_{\mathsf{d}}) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v_{\mathsf{d}}) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v_{\mathsf{d}})\right)} \\
&= \frac{\alpha}{\sum_{\mathsf{d}\in\mathcal{D}} \left(\frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(u_{\mathsf{d}})}{|\mathsf{D}_{\boldsymbol{\sigma}}(u_{\mathsf{d}})|} + \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v_{\mathsf{d}})}{|\mathsf{D}_{\boldsymbol{\sigma}}(v_{\mathsf{d}})|} \right)} \\
&= \frac{\alpha}{\sum_{v\in V} \sum_{\mathsf{d}\in\mathcal{D} | v \in \mathsf{Vertices}_{\boldsymbol{\sigma}}(\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}))} \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|}} \\
&= \frac{\alpha}{\sum_{v\in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} |\mathsf{D}_{\boldsymbol{\sigma}}(v)| \cdot \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|}} \\
&= \frac{\alpha}{\sum_{v\in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})} |\mathsf{A}|_{\boldsymbol{\sigma}}(v)} \\
&= 1\,.
\end{aligned}
$$

By Corollary 6.1, it follows that $\mathsf{MinHit}_{\boldsymbol{\sigma}} = 1$. Hence, Lemma 4.3 implies that $\delta \geq \frac{|V|}{2}$. Since $\mathsf{DR}_{\boldsymbol{\sigma}} = 1$, it follows that $\mathsf{DR}_{\boldsymbol{\sigma}} = \max\left\{1, \frac{|V|}{2\,\delta}\right\}$, so that G is Defense-Optimal. ■

We compare the sufficient conditions for a Defense-Optimal graph from Theorems 6.3 and 6.5:

Proposition 6.6 *There is a graph G and an integer δ such that G has a δ-Partitionable Fractional Perfect Matching while G is not Defender-Pure.*

Proof. Consider a graph $G = (V, E)$ consisting of two triangles $\mathcal{T}_1$ and $\mathcal{T}_2$ and a single edge $e' = (u, v)$ connecting them with $u \in V(\mathcal{T}_1)$ and $v \in V(\mathcal{T}_2)$. Consider the function $f : E \rightarrow [0, 1]$ with $f(e) = \frac{1}{2}$ for each edge $e \in E \setminus \{e'\}$ and $f(e') = 0$. Clearly, f is a 2-Partitionable Fractional Perfect Matching with $E_1 = \mathcal{T}_1$ and $E_2 = \mathcal{T}_2$. Since $\delta = 2$ and $\beta'(G) = 3$, it follows by Proposition 5.7 (($\mathsf{C1}$)), that G is not Defender-Pure. ■

7 Few Defenders

In this case, $\delta \leq \frac{|V|}{2}$; so, a Defense-Optimal Nash equilibrium $\boldsymbol{\sigma}$ has $\mathsf{DR}_{\boldsymbol{\sigma}} = \max\left\{1, \frac{|V|}{2\,\delta}\right\} = \frac{|V|}{2\delta}$, so that, by Corollary 6.1, $\mathsf{MinHit}_{\boldsymbol{\sigma}} = \frac{2\delta}{|V|}$, and $\sum_{v\in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) \geq 2\delta$. By Lemma 4.2, it follows that $\sum_{v\in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 2\delta$, so that $\boldsymbol{\sigma}$ is unidefender. By Corollary 5.4, $\boldsymbol{\sigma}$ is monodefender.

7.1 Necessary Conditions

We show a necessary condition for Defense-Optimal graphs:

Proposition 7.1 *Assume that $\delta \leq \frac{|V|}{2}$. Then, a Defense-Optimal graph has a δ-Partitionable Fractional Perfect Matching.*

Proof. Consider a Defense-Optimal Nash equilibrium $\boldsymbol{\sigma}$. Recall that $\boldsymbol{\sigma}$ is monodefender. Since $\mathsf{MinHit}_{\boldsymbol{\sigma}}(v) = \frac{2\delta}{|V|}$ and $\sum_{v \in V} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 2\delta$, it follows that for each vertex $v \in V$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \frac{2\delta}{|V|}$. We now define a function $f : E \to \mathbb{R}$; we will then prove that f is a δ-Partitionable Fractional Perfect Matching. For each edge $e \in E$, set

$$f(e) \quad := \quad \begin{cases} \frac{|V|}{2\delta} \cdot \sigma_{\mathsf{d}_{\boldsymbol{\sigma}}(e)}(e), & \text{if } e \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}) \\ 0, & \text{otherwise} \end{cases} .$$

By construction, $E(f) = \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$. So, for each vertex $v \in V$, $\{e \in E(f) \mid e \ni v\} = \mathsf{Edges}_{\boldsymbol{\sigma}}(v)$. Since $\boldsymbol{\sigma}$ is monodefender, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_{\boldsymbol{\sigma}}(v), v))$. We use this fact and the definition of f to prove:

Claim 7.2 *For each vertex $v \in V$, $\sum_{e \in \mathsf{Edges}_{\boldsymbol{\sigma}}(v)} f(e) = 1$.*

Since $\mathsf{Edges}_{\boldsymbol{\sigma}}(v) = \{e \in E(f) \mid v \in e\}$, Claim 7.2 implies that f is a Fractional Perfect Matching. To prove that f is δ-Partitionable, define the (non-empty) sets $E_1, \cdots, E_\delta$ where for each $j \in [\delta]$, $E_j := \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j)$. Clearly, $\bigcup_{j \in [\delta]} E_j = \bigcup_{j \in [\delta]} \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) = \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}) = E(f)$. Since $\boldsymbol{\sigma}$ is monodefender, the sets $E_1, \cdots, E_\delta$ partition $E(f)$. We observe:

Claim 7.3 *For each index $j \in [\delta]$, $\sum_{e \in E_j} f(e) = \frac{|V|}{2\delta}$.*

Proof. By the constructions of f and $E_1, \cdots, E_\delta$, $\sum_{e \in E_j} f(e) = \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j)} f(e) = \frac{|V|}{2\delta} \sum_{e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j)} \sigma_{\mathsf{d}_j}(e) = \frac{|V|}{2\delta}$. ■

Claim 7.3 implies that f is δ-Partitionable, and we are done. ■

Proposition 7.1 establishes that the sufficient condition for a Defense-Optimal graph from Theorem 6.3 is also necessary when $\delta \leq \frac{|V|}{2}$.

7.2 Characterization and Complexity of Defense-Optimal Graphs

We now state a combinatorial characterization of Defense-Optimal graphs $\left(\text{for } \delta \leq \frac{|V|}{2}\right)$. Sufficiency and necessity follow from Theorem 6.3 and Proposition 7.1, respectively.

Theorem 7.4 *Assume that $\delta \leq \frac{|V|}{2}$. Then, G is Defense-Optimal if and only if G has a δ-Partitionable Fractional Perfect Matching.*

A first implication is an immediate consequence of Theorem 7.4 and Corollary 2.7.

Corollary 7.5 *Assume that $\delta \leq \frac{|V|}{2}$ and G is Defense-Optimal. Then, δ divides $|V|$.*

The second implication is an immediate consequence of Theorem 7.4 and Proposition 2.11.

Corollary 7.6 *Assume that $\delta = \frac{|V|}{2}$. Then, G is Defense-Optimal if and only if it is Perfect-Matching.*

Corollary 7.6 yields that for $\delta = \frac{|V|}{2}$, the recognition problem for Defense-Optimal graphs is tractable. For the third implication, Theorem 7.4 and Proposition 2.12 immediately imply:

Corollary 7.7 *Assume that $\delta \leq \frac{|V|}{2}$. Then, the recognition problem for Defense-Optimal graphs is $\mathcal{NP}$-complete.*

7.3 Perfect-Matching Graphs

We show:

Theorem 7.8 *Assume that $\delta \leq \frac{|V|}{2}$ for a Perfect-Matching graph G. Then, G admits a Defense-Optimal, Perfect-Matching Nash equilibrium if and only if $2\,\delta$ divides $|V|$.*

Proof. The claim will follow from Propositions 7.9 and 7.10.

Proposition 7.9 *Assume that $\delta \leq \frac{|V|}{2}$ for a Perfect-Matching graph G, where $2\,\delta$ divides $|V|$. Then, G admits a Defense-Optimal, Perfect-Matching Nash equilibrium.*

Proof. Consider a Perfect Matching M. Construct a mixed profile $\boldsymbol{\sigma}$ as follows:

- For an attacker $\mathsf{a} \in \mathcal{A}$ and each vertex $v \in V$, set $\sigma_{\mathsf{a}}(v) := \frac{1}{|V|}$. So, $\boldsymbol{\sigma}$ is attacker-symmetric, attacker-uniform and attacker-fully mixed, and for each vertex $v \in V$, $|\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \frac{\alpha}{|V|}$.

- Partition M into δ sets $M_1, \cdots, M_\delta$, each with $\frac{|V|}{2\delta}$ edges; each defender d_j uses a uniform distribution over M_j. So, for each edge $e \in M_j$, set $\boldsymbol{\sigma}_{\mathsf{d}_{\boldsymbol{\sigma}}(e)}(e) := \frac{2\delta}{|V|}$. Thus, $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_j) = M_j$ for each defender d_j, so that $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}) = M$. Clearly, each vertex $v \in V$ is monodefender in $\boldsymbol{\sigma}$ with $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}_{\boldsymbol{\sigma}}(v), v)) = \frac{2\delta}{|V|}$.

We shall verify (C1) and (C2) in the characterization of Nash equilibria (Proposition 5.1). For (C1), fix a vertex $v \in V$. Since $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \frac{2\delta}{|V|}$, (C1) follows trivially. For (C2), consider a defender $\mathsf{d} \in \mathcal{D}$.

- Fix an edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Since each edge is monodefender in $\boldsymbol{\sigma}$, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) = \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) = 1$. Hence, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u) + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \frac{2\alpha}{|V|}$.

- Fix now an edge $(u', v') \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Since M is an Edge Cover, there are edges $e_{u'}, e_{v'} \in M$ such that $u' \in e_{u'}$ and $v' \in e_{v'}$. By the construction of $\boldsymbol{\sigma}$, this implies that there are defenders $\mathsf{d}_{u'}$ and $\mathsf{d}_{v'}$ such that $e_{u'} \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_{u'})$ and $e_{v'} \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_{v'})$. Since each vertex is monodefender in $\boldsymbol{\sigma}$, it follows that $\mathsf{d} \neq \mathsf{d}_{u'}$ and $\mathsf{d} \neq \mathsf{d}_{v'}$. Hence, $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \leq \frac{1}{2}$ and $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v) \leq \frac{1}{2}$, so that $\mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond u') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(u') + \mathsf{Prop}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond v') \cdot |\mathsf{A}|_{\boldsymbol{\sigma}}(v') \leq \frac{1}{2} \cdot (|\mathsf{A}|_{\boldsymbol{\sigma}}(u') + |\mathsf{A}|_{\boldsymbol{\sigma}}(v')) = \frac{\alpha}{|V|}$.

Now, (C2) follows. Hence, by Proposition 5.1, $\boldsymbol{\sigma}$ is a Nash equilibrium.

Now recall that for each vertex $v \in V$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = \frac{2\delta}{|V|}$. Hence, $\mathsf{MinHit}_{\boldsymbol{\sigma}} = \frac{2\delta}{|V|}$. By Corollary 6.1, $\mathsf{DR}_{\boldsymbol{\sigma}} = \frac{|V|}{2\delta} = \max\left\{1, \frac{|V|}{2\,\delta}\right\}$ $\left(\text{since } \delta \leq \frac{|V|}{2}\right)$. Hence, $\boldsymbol{\sigma}$ is Defense-Optimal. ■

We continue to prove:

Proposition 7.10 *Assume that $\delta \leq \frac{|V|}{2}$ for a Perfect-Matching graph G, which admits a Defense-Optimal, Perfect-Matching Nash equilibrium. Then, $2\,\delta$ divides $|V|$.*

Proof. Consider such a Nash equilibrium $\boldsymbol{\sigma}$, and recall that $\mathsf{MinHit}_{\boldsymbol{\sigma}} = \frac{2\delta}{|V|}$. Consider an edge $(u, v) \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$; so, $e \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$ for some defender $\mathsf{d} \in \mathcal{D}$. Proposition 5.5 implies that $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$ is a Vertex Cover of the graph $G(\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D}))$. Hence, either $u \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$ or $v \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$. Assume, without loss of generality, that $u \in \mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A})$. Since $\boldsymbol{\sigma}$ is monodefender, there is a single defender d_k such that $u \in \mathsf{Vertices}(\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d}_k))$. Hence, d_k is identified with d. Since $\boldsymbol{\sigma}$ is Perfect-Matching, $\mathsf{Support}_{\boldsymbol{\sigma}}(\mathcal{D})$ is a Perfect Matching; this implies that $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = s_{\mathsf{d}}(e)$. We prove:

Claim 7.11 $|\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})| = \frac{|V|}{2\delta}$

Proof. Clearly, $\sum_{e\in\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})} \sigma_{\mathsf{d}}(e) = \sum_{e\in\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})} \mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = |\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})| \cdot \frac{2\delta}{|V|}$. Since $\boldsymbol{\sigma}$ is a mixed profile, $\sum_{e\in\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})} \sigma_{\mathsf{d}}(e) = 1$. Hence, $|\mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})| = \frac{|V|}{2\delta}$, as needed. ■

Claim 7.11 immediately implies that 2δ divides $|V|$, as needed. ■

The claim follows now from Propositions 7.9 and 7.10. ■

While Corollary 7.5 applies to *all* graphs, Proposition 7.10 applies to Perfect-Matching graphs. The application of Corollary 7.5 to Perfect-Matching graphs does *not* imply Proposition 7.10 *unless* δ is odd. (In detail, 2 divides $|V|$ and δ divides $|V|$ imply together that 2δ divides $|V|$ only if δ is odd.) So, Proposition 7.10 strictly strengthens Corollary 7.5 to the case of even δ.

8 Many Defenders

In this case, $\frac{|V|}{2} < \delta < \beta'(G)$; so, a Defense-Optimal Nash equilibrium $\boldsymbol{\sigma}$ has Defense-Ratio $\mathsf{DR}_{\boldsymbol{\sigma}} = \max\left\{1, \frac{|V|}{2\,\delta}\right\} = 1$. By Corollary 6.1, this implies that $\mathsf{MinHit}_{\boldsymbol{\sigma}} = 1$. It follows that for a vertex $v \in V$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 1$ and the number of maxhit vertices in $\boldsymbol{\sigma}$ is $|V|$. We show:

Theorem 8.1 *Assume that* $\frac{|V|}{2} < \delta < \beta'(G)$*. Then,* G *is not Defense-Optimal.*

Proof. Towards a contradiction, consider a Defense-Optimal Nash equilibrium $\boldsymbol{\sigma}$. Consider any (maxhit) vertex $v \in V$. By Lemma 4.4, there is a maxhitter $\mathsf{d} \in \mathcal{D}$ in $\boldsymbol{\sigma}$ with $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(\mathsf{d}, v)) = 1$. Use $\boldsymbol{\sigma}$ to construct a defender-pure profile $\boldsymbol{\tau}$ as follows: Fix a defender $\mathsf{d} \in \mathcal{D}$. If d is maxhitter in $\boldsymbol{\sigma}$, then $\boldsymbol{\tau}_{\mathsf{d}}$ is any edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$ such that d is maxhitter in $\boldsymbol{\sigma}$ for the vertex $v \in V$; else, $\boldsymbol{\tau}_{\mathsf{d}}$ is any arbitrary edge $(u, v) \in \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$.

By construction of $\boldsymbol{\tau}$, it holds that: (*i*) $|\mathsf{Supports}_{\boldsymbol{\tau}}(\mathcal{D})| \leq \delta$. (*ii*) Each maxhit vertex in $\boldsymbol{\sigma}$ is maxhit in $\boldsymbol{\tau}$; so, the number of maxhit vertices in $\boldsymbol{\tau}$ is $|V|$. Since $\delta < \beta'(G)$, (*i*) implies that $|\mathsf{Supports}_{\boldsymbol{\tau}}(\mathcal{D})| < \beta'(G)$. Hence, $\mathsf{Supports}_{\boldsymbol{\tau}}(\mathcal{D})$ is not an Edge Cover. So, there is a vertex $v \in V$ with $\mathbb{P}_{\boldsymbol{\tau}}(\mathsf{Hit}(v)) = 0$, and the number of maxhit vertices in $\boldsymbol{\tau}$ is less than $|V|$. A contradiction.■

9 Too Many Defenders

In this case, $\frac{|V|}{2\,\delta} \leq \frac{|V|}{2\beta'(G)} \leq 1$; so, a Defense-Optimal Nash equilibrium $\boldsymbol{\sigma}$ has $\mathsf{DR}_{\boldsymbol{\sigma}} = 1$. By Corollary 6.1, $\mathsf{MinHit}_{\boldsymbol{\sigma}} = 1$, so that for each vertex $v \in V$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 1$.

9.1 (Defender-Pure and Pure,) Vertex-Balanced Profiles

A mixed profile $\boldsymbol{\sigma}$ is ***vertex-balanced*** if there is a constant $c > 0$ such that for each vertex $v \in V$, $\frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|} = c$. It follows trivially that: (C1) $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ is an Edge Cover, and this matches the necessary condition for an arbitrary Nash equilibrium from Proposition 5.3, and (C2) $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{A}) = V$.

We shall consider defender-pure, vertex-balanced profiles and pure, vertex-balanced profiles. We prove an interesting property of defender-pure, vertex-balanced profiles.

Proposition 9.1 *A defender-pure, vertex-balanced profile is a local maximizer for the Expected Utility of each defender.*

Proof. Consider such a profile $\boldsymbol{\sigma}$ and a defender $\mathsf{d} \in \mathcal{D}$ with $\sigma_{\mathsf{d}} = (u, v)$. Clearly, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}) = \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(u)}{|\mathsf{D}_{\boldsymbol{\sigma}}(u)|} + \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|} = 2c$. Fix now an edge $(u', v') \notin \mathsf{Support}_{\boldsymbol{\sigma}}(\mathsf{d})$. Clearly, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{\sigma}_{-\mathsf{d}} \diamond (u', v')) = \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(u')}{|\mathsf{D}_{\boldsymbol{\sigma}}(u')| + 1} + \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v')}{|\mathsf{D}_{\boldsymbol{\sigma}}(v')| + 1} < \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(u')}{|\mathsf{D}_{\boldsymbol{\sigma}}(u')|} + \frac{|\mathsf{A}|_{\boldsymbol{\sigma}}(v')}{|\mathsf{D}_{\boldsymbol{\sigma}}(v')|} = 2c$, and the claim follows. ■

By Proposition 9.1, a defender-pure, vertex-balanced profile, which is also a local maximizer for the Expected Utility of each attacker, is a Nash equilibrium. We shall present polynomial time algorithms for Defender-Pure (and Pure), Vertex-Balanced Nash equilibria which are Defense-Optimal for the case where $\delta \geq \beta'(G)$; the second algorithm requires an additional assumption.

9.2 Defense-Optimal, Defender-Pure, Vertex-Balanced Nash Equilibria

We show:

Theorem 9.2 *Assume that $\delta \geq \beta'(G)$. Then, G admits a Defender-Pure, Vertex-Balanced Nash equilibrium, which is computable in polynomial time.*

To prove the claim, we present the algorithm DefenderPure&VertexBalancedNE in Figure 4.

Proof. By construction (Steps (1) and (2)) and the assumption $\delta \geq \beta'(G)$, $\mathsf{Supports}_{\boldsymbol{\sigma}}(\mathcal{D})$ is a Minimum Edge Cover. Since $\boldsymbol{\sigma}$ is defender-pure, this implies that for each vertex $v \in V$, $\mathbb{P}_{\boldsymbol{\sigma}}(\mathsf{Hit}(v)) = 1$; hence, for each attacker $\mathsf{a} \in \mathcal{A}$, $\mathsf{U}_{\mathsf{a}}(\boldsymbol{\sigma}_{-\mathsf{a}} \diamond v) = 0$. So, $\boldsymbol{\sigma}$ is a local maximizer for the Expected Utility of each attacker. By Proposition 9.1, $\boldsymbol{\sigma}$ is a Nash equilibrium. ■

By Theorems 6.5 and 9.2, it immediately follows:

Corollary 9.3 *Assume that $\delta \geq \beta'(G)$. Then, G is Defense Optimal.*

Algorithm DefenderPure&VertexBalancedNE
INPUT: A graph $G = \langle V, E \rangle$ such that $\delta \geq \beta'(G)$.
OUTPUT: A Defense-Optimal, Defender-Pure Vertex-Balanced Nash equilibrium $\boldsymbol{\sigma}$.

(1) Choose a Minimum Edge Cover EC.

(2) Assign each defender to a distinct edge from EC in a round-robin fashion.

(3) Determine a solution $\left\{ \mathsf{A}(v) \mid v \in V \right\}$ to the following linear system:

  (a) For each vertex $v \in V$, $\dfrac{\mathsf{A}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|}$ is constant.

  (b) $\sum_{v \in V} \mathsf{A}(v) = \alpha$.

(4) Assign a mixed strategy $\boldsymbol{\sigma}$ to each attacker in an arbitrary way so that for each vertex $v \in V$, $|\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \mathsf{A}(v)$.

Figure 4: The algorithm DefenderPure&VertexBalancedNE. By Step (2) and since $\delta \geq \beta'(G)$, $\boldsymbol{\sigma}$ is defender-pure. Step (3) provides for $\boldsymbol{\sigma}$ to be vertex-balanced; towards this end, it provides for the ratio $\dfrac{\mathsf{A}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|}$ to be constant over all vertices $v \in V$. Thus, by Step (4), it follows that $\boldsymbol{\sigma}$ is vertex-balanced. Since a Minimum Edge Cover is computable in polynomial time, the algorithm DefenderPure&VertexBalancedNE is polynomial time.

By Theorem 5.7 (Condition (i)) and Theorem 9.2, it finally follows:

Corollary 9.4 *G is Defender-Pure if and only if $\delta \geq \beta'(G)$.*

Since a Minimum Edge Cover is computable in polynomial time, Corollary 9.4 implies that Defender-Pure graphs are recognizable in polynomial time (for an arbitrary value of δ). For the case where $2\,\delta$ divides α, we strengthen Proposition 9.2 as follows.

Theorem 9.5 *Assume that $\delta \geq \beta'(G)$ and $2\,\delta$ divides α. Then, G admits a Defense-Optimal, Pure, Vertex-Balanced Nash equilibrium, which is computable in polynomial time.*

To prove the claim, we consider a modification of algorithm DefenderPure&VertexBalancedNE which differentiates from the original algorithm only in the following steps: In Step (3/a), the constant is set to $\frac{\alpha}{2\,\delta}$. In Step (4), a pure strategy is assigned for each attacker instead of a mixed one, with the same properties as in the original algorithm (see below). Note that by Step (3/a), $\dfrac{\mathsf{A}(v)}{|\mathsf{D}_{\boldsymbol{\sigma}}(v)|} = \dfrac{\alpha}{2\delta}$ for each vertex $v \in V$; hence, $\mathsf{A}(v)$ is integer. So, at Step (4) pure strategies *can* be assigned to the attackers that induce the integer value $|\mathsf{A}|_{\boldsymbol{\sigma}}(v) = \mathsf{A}(v)$ for each vertex $v \in V$. The rest of the proof is identical to that of Theorem 9.2.

10 Epilogue

We proposed and analyzed a new combinatorial model for a distributed system like the Internet with selfish, malicious attacks and selfish, non-malicious, interdependent defenses. Through an extensive combinatorial analysis of Nash equilibria for this model, we derived a comprehensive collection of (in some cases surprising) trade-off results between the number of defenders and the best possible Defense-Ratio of associated Nash equilibria.

We have identified several new classes of graphs for an arbitrary pair of values of δ and α, such as graphs with δ-Partitionable Fractional Perfect Matching, Defense-Optimal and Pure graphs; each such class guarantees the existence of a Nash equilibrium with a particular property. Our results reveal a fine structure among these classes, which is summarized in Figure 5.

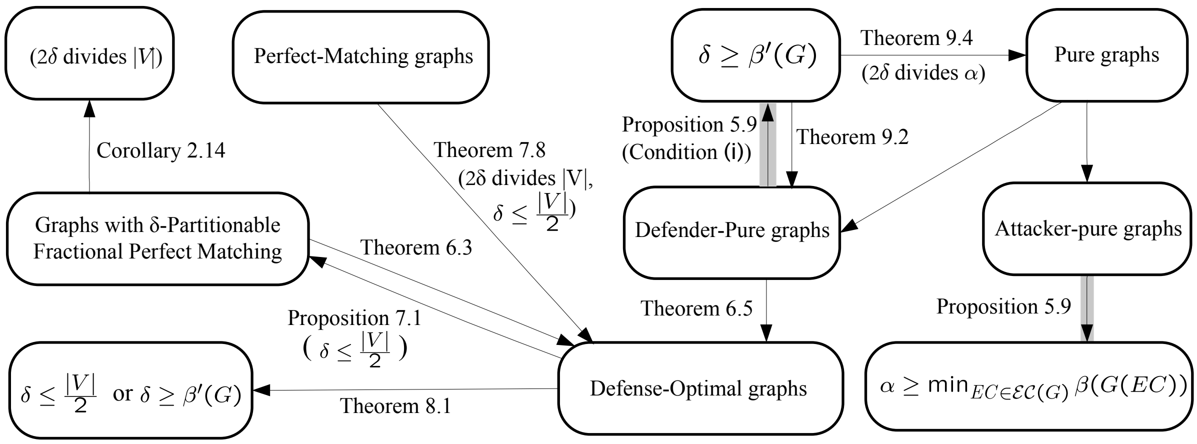

Figure 5: Some inclusion relationships among the graph classes associated with Nash equilibria we have introduced. A directed edge from class $\mathcal{C}_1$ to class $\mathcal{C}_2$ indicates that $\mathcal{C}_1 \subseteq \mathcal{C}_2$; a condition on the edge indicates the condition under which the inclusion holds. Clouded directed edges indicate inclusions that have been demonstrated to be non-strict.

Our work leaves numerous open problems relating to (i) the *worst-case* Nash equilibria, (ii) the investigation of alternative reward-sharing schemes for the defenders and (iii) the complexity of computing and verifying (Defense-Optimal) Nash equilibria in this model.

Acknowledgments. We thank Martin Gairing, Loizos Michael, Florian Schoppmann, Karsten Tiemann and the anonymous *HICSS 2008* and *Discrete Applied Mathematics* reviewers for many helpful comments and suggestions on earlier versions of this paper.

References

[1] R. Anderson, "Why Information Security is Hard—An Economic Perspective," *Proceedings of the 17th Annual Computer Security Applications Conference*, pp. 358–365, December 2001.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler and T. Roughgarden, "The Price of Stability for Network Design with Fair Cost Allocation," *SIAM Journal on Computing*, Vol. 38, No. 4, pp. 1602–1623, November 2008.

[3] 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.

[4] N. Christin, J. Grossklags and J. Chuang,"Near Rationality and Competitive Equilibria in Networked Systems," *Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 213–219, August 2004.

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

[6] M. R. Garey and D. S. Johnson, *Computers and Intractability—A Guide to the Theory of $\mathcal{NP}$-Completeness*, W. H. Freeman and Co., 1979.

[7] M. Gelastou, M. Mavronicolas, V. Papadopoulou, A. Philippou and P. G. Spirakis, "The Power of the Defender," *CD-ROM Proceedings of the 2nd International Workshop on Incentive-Based Computing*, in conjunction with the *26th IEEE International Conference on Distributed Computing Systems*, July 2006.

[8] L. Gordon and M. Loeb, "The Economics of Information Security Investment," *ACM Transactions on Information and System Security*, Vol. 5, No. 4, pp. 438–457, November 2002.

[9] J. Grossklags, N. Christin and J. Chuang, "Predicted and Observed User Behavior in the Weakest-Link Security Game," *Proceedings of the USENIX Workshop on Usability, Psychology, and Security*, pp. 1–6, April 2008.

[10] J. Grossklags, N. Christin and J. Chuang, "Secure or Insure? A Game-Theoretic Analysis of Information Security Games," *Proceedings of the 17th International World Wide Web Conference*, pp. 209–218, April 2008.

[11] M. Kearns and L. Ortiz, "Algorithms for Interdependent Security Games", *Advances in Neural Information Processing Systems*, Vol. 16, S. Thrun, L. Saul and B. Schölkopf eds., The MIT Press, 2004.

[12] B. Korte and J. Vygen, *Combinatorial Optimization—Theory and Algorithms*, Springer, 2000.

[13] H. Kunreuther and G. Heal, "Interdependent Security", *Journal of Risk and Uncertainty*, Special Issue on Terrorist Risks, Vol. 26, No. 2/3, pp. 231–249, March 2003.

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

[15] E. L. Lawler, *Combinatorial Optimization—Networks and Matroids*, Holt, Rinehart and Winston, 1976.

[16] 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, June 2001.

[17] 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, Lecture Notes in Computer Science, Springer-Verlag, August/September 2006.

[18] 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, Lecture Notes in Computer Science, Springer-Verlag, December 2006.

[19] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, "A Graph-Theoretic Network Security Game," *International Journal of Autonomous and Adaptive Communications Systems*, Vol. 1, No. 4, pp. pp. 390–410, November 2008.

[20] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, "A Network Game with Attackers and a Defender," *Algorithmica,* Vol. 51, No. 3, pp. 315–341, July 2008.

[21] B. Monien, "The Complexity of Determining a Shortest Cycle of Even Length," *Computing*, Vol. 31, No. 4, pp. 355–369, December 1983.

[22] 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.

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

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

[25] A. Termimi Ab Ghani and K. Tanaka, “Network Security Games With and Without Synchroneity”, *Proceedings of the 2nd International Conference on Decision and Game Theory for Security*, pp. 87–103, Vol. 7037, Lecture Notes in Computer Science, Springer-Verlag, November 2011.

[26] N. Weaver and V. Paxson, “A Worst-Case Worm,” *Proceedings of the 3rd Annual Workshop on Economics of Information Security,* May 2004. Available online at `http://www.icir.org/vern/papers/worst-case-worm.WEIS04.pdf`

[27] D. B. West, *Introduction to Graph Theory,* Prentice Hall, Second Edition, 2001.

[28] R. Yuster and U. Zwick, “Finding Even Cycles Even Faster,” *SIAM Journal on Discrete Mathematics*, Vol. 10, No. 2, pp. 209-222, May 1997.

A Proof of Proposition 2.3

We start with two invariants of the algorithm $\mathsf{EliminateEvenCycles}$:

Lemma A.1 *For each loop iteration of* $\mathsf{EliminateEvenCycles}$*, upon completion of Step* (3)*,* f'' *is a Fractional Matching equivalent to* f*.*

Note that the input Fractional Matching f is already modified in the first (if any) loop iteration of $\mathsf{EliminateEvenCycles}$ (in Step (4)), while the statements of Lemmas A.1 and A.2 refer to the input Fractional Matching f. Their proof will use the current Fractional Matching f; reference to the input f will be restored in an inductive way upon completing the proof.

Proof. Fix any loop iteration of $\mathsf{EliminateEvenCycles}$, upon completion of Step (3). Consider any vertex $v \in V$. Then, by Step (3),

$$\begin{aligned}\sum_{e \in E \mid v \in e} f''(e) &= \sum_{e \in E(\mathcal{C}) \mid v \in e} f''(e) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f''(e) \\ &= \sum_{e \in E(\mathcal{C}) \mid v \in e} f''(e) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f(e)\end{aligned}$$

If there is no edge $e \in E(\mathcal{C})$ such that $v \in e$, then $\sum_{e \in E(\mathcal{C}) \mid v \in e} f''(e) = \sum_{e \in E(\mathcal{C}) \mid v \in e} f(e) = 0$, and equivalence follows. So, assume otherwise. Since $\mathcal{C}$ is a cycle, there are (exactly) two edges e_1, $e_2 \in E(\mathcal{C})$ such that $v \in e_1$ and $v \in e_2$. Note that by Step (2), $g(e_1) + g(e_2) = 0$. Hence, by Step (3),

$$\begin{aligned}\sum_{e \in E \mid v \in e} f''(e) &= f''(e_1) + f''(e_2) + \sum_{eE \setminus E(\mathcal{C}) \mid v \in e} f(e) \\ &= f(e_1) + g(e_1) \cdot f(e_0) + f(e_2) + g(e_2) \cdot f(e_0) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f(e) \\ &= f(e_1) + f(e_2) + (g(e_1) + g(e_2)) \cdot f(e_0) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f(e) \\ &= f(e_1) + f(e_2) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f(e) \\ &= \sum_{e \in E(\mathcal{C}) \mid v \in e} f(e) + \sum_{e \in E \setminus E(\mathcal{C}) \mid v \in e} f(e) \\ &= \sum_{e \in E \mid v \in e} f(e),\end{aligned}$$

so that f'' is equivalent to f. By Step (4), it follows inductively that f'' is equivalent to f. Since f is a Fractional Matching, f'' is a Fractional Matching, and the claim follows. ■

We continue with a second invariant of the algorithm $\mathsf{EliminateEvenCycles}$:

Lemma A.2 *For each loop iteration of* $\mathsf{EliminateEvenCycles}$*, upon completion of Step* (3)*,* (i) $f'' \subset f$ *and* (ii) *the even cycle* $\mathcal{C}$ *is eliminated from* $G(E(f''))$*.*

Proof. Fix any loop iteration of EliminateEvenCycles, upon completion of Step (3). Consider any edge $e \in E$. We proceed by case analysis: (i) Assume that $e \notin E(\mathcal{C})$. Then, Step (3), implies that $e \in E(f'')$ if and only if $e \in E(f)$. (ii) Assume that $e \in E(\mathcal{C})$. Then, $e \in E(f)$; so, it holds vacuously that if $e \in E(f'')$, then $e \in E(f)$. The case analysis implies that $f'' \subseteq f$. Since $f''(e_0) = 0$ while $f(e_0) > 0$, this implies that $f'' \subset f$. By Step (4), Condition (i) follows now inductively. Since $f''(e_0) = 0$, edge e_0 is eliminated from $G(E(f''))$, so that the even cycle $\mathcal{C}$ is eliminated from $G(E(f''))$ and Condition (ii) follows. ■

Lemma A.1 and Lemma A.2 (Condition (i)) together imply that the output f'' of algorithm EliminateEvenCycles, which contains no even cycle due to the loop precondition, is a Fractional Matching which is equivalent to and contained in f. (By Lemma A.2 (Condition (ii)), containment is strict exactly when there is at least one loop iteration.)

Lemma A.2 (Condition (i) or (ii)) implies that at least one edge is eliminated from f in each loop iteration and no edge is added. Hence, there are at most $|E|$ loop iterations. Note that each loop iteration takes $O(|E|)$ time. Since an even cycle is computable in polynomial time, the algorithm EliminateEvenCycles is polynomial time, and we are done.

B Proof of Proposition 2.4

Since $G(E(f))$ has no even cycle, the cycle $v_l, \cdots, v_r = v_l$ determined in Step (2/a) is odd. We now prove a preliminary property of the algorithm IsolateOddCycles:

Lemma B.1 *The path $v_1, v_2, \cdots, v_r$ is disjoint from $\mathcal{C} \backslash \{v_0\}$.*

Proof. Assume by way of contradiction, that there is a vertex v_k with $k \in [r]$ such that $v_k \in \mathcal{C} \backslash \{v_0\}$. Since $\mathcal{C}$ is odd, 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 two concatenations of the path $v_1, \cdots, v_k$ with $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively; each is a cycle in $G(E(f))$ and one is even. A contradiction. ■

We start with a first invariant of the algorithm IsolateOddCycles.

Lemma B.2 *For each inner loop iteration in an outer loop iteration of* IsolateOddCycles*, upon completion of* Step (2/e)*, f' is a Fractional Perfect Matching equivalent to f.*

Note that the input Fractional Perfect Matching f is already modified in the first inner loop iteration in the first outer loop iteration of IsolateOddCycles (in Step (2/f)). Reminiscent of Lemma A.1, the statement of Lemma B.2 refers to the input Fractional Perfect Matching f.

The proof of Lemma B.2 will use the current Fractional Perfect Matching f; reference to the input f will be restored in an inductive way upon completing the proof.

Proof. The proof consists of two technical claims. The first claim determines the range of f'. Fix any inner loop iteration in an outer loop iteration of IsolateOddCycles, upon completion of Step (2/e). We prove:

Claim B.3 Range$(f') \subseteq [0, 1]$.

Proof. By Step (2/e), it suffices to consider inductively an edge e from $E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$. By Step (2/f), it follows inductively that f is a Fractional Perfect Matching.

We first prove that $f'(e) \geq 0$. By Step (2/e), it suffices to consider the case where $g(e) < 0$, so that $g(e) = -|g(e)|$. Then, by Step (2/e) and the choice of the edge e_0, $f'(e) = f(e) - |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} \geq 0$, as needed.

We now prove that $f'(e) \leq 1$. By Step (2/e), it suffices to consider the case where $g(e) > 0$, so that $g(e) = |g(e)|$. We proceed by case analysis on whether there is an edge e' adjacent to e such that e and e' are either both on the cycle $\mathcal{C}$ or both on the path $v_0, \cdots, v_l$ (with $l > 0$) or both on the cycle $v_l, \cdots, v_r = v_l$.

Assume first that there is such an edge e'; clearly, $|g(e')| = |g(e)|$. Then, by Step (2/e),

$$
\begin{array}{lcl}
f'(e) & & \\
= & f(e) + |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} & \\
\leq & 1 - f(e') + \frac{g(e)}{|g(e_0)|} \cdot f(e_0) & \text{(since } f \text{ is a Fractional Matching)} \\
\leq & 1 - |g(e')| \cdot \frac{f(e_0)}{|g(e_0)|} + g(e) \cdot \frac{f(e_0)}{|g(e_0)|} & \text{(by the choice of the edge } e_0\text{)} \\
= & 1 - |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} + g(e) \cdot \frac{f(e_0)}{|g(e_0)|} & \\
= & 1, &
\end{array}
$$

as needed. Assume now that there is no edge e' adjacent to e such that e and e' are either both on the cycle $\mathcal{C}$ or both on the path $v_0, \cdots, v_l$ (with $l > 0$) or on the cycle $v_l, \cdots, v_r = v_l$. Since both the cycle $\mathcal{C}$ and the cycle $v_l, \cdots, v_r = v_l$ are odd, each of them includes at least three edges. It follows that edge e lies neither on the cycle $\mathcal{C}$ nor on the cycle $v_l, \cdots, v_r = v_l$. Hence, edge e lies on the path $v_0, \cdots, v_l$ (with $l > 0$). Since there is no edge e' adjacent to e on this path, it follows that $l = 1$, so that $e = (v_0, v_1)$. So, consider the edges e_1 and e_2 on the cycle $\mathcal{C}$ that are adjacent to e. By the choice of g, it follows that $|g(e_1)| + |g(e_2)| = |g(e)|$. Hence,

$$
\begin{aligned}
&f'(e)\\
= \quad & f(e) + |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} && \text{(by Step (2/e))}\\
\leq \quad & 1 - f(e_1) - f(e_2) + |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} && \text{(since } f \text{ is a Fractional Matching)}\\
\leq \quad & 1 - |g(e_1)| \cdot \frac{f(e_0)}{|g(e_0)|} - |g(e_2)| \cdot \frac{f(e_0)}{|g(e_0)|} + |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} && \text{(by definition of } e_0)\\
= \quad & 1 - \Big(|g(e_1)| + |g(e_2)| - |g(e)| \Big) \cdot \frac{f(e_0)}{|g(e_0)|}\\
= \quad & 1,
\end{aligned}
$$

as needed. The proof is now complete. ■

We continue with the second technical claim:

Claim B.4 *f' is equivalent to f.*

Proof. Consider any vertex $v \in V$. Then, by Step (2/e)),

$$
\begin{aligned}
\sum_{e \in E | v \in e} f'(e) &= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f'(e) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\}) \,|\, v \in e} f'(e)\\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f'(e) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\}) \,|\, v \in e} f(e).
\end{aligned}
$$

If there is no edge $e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$ such that $v \in e$, then

$$
\begin{aligned}
\sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f'(e) &= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f(e)\\
&= 0,
\end{aligned}
$$

and we are done. So, assume otherwise. Note that by Step (2/b),

$$
\sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} g(e) = 0.
$$

Hence, by Step (2/e),

$$
\begin{aligned}
\sum_{e \in E | v \in e} f'(e) &= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f'(e) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\}) \,|\, v \in e} f(e)\\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} \left(f(e) + g(e) \cdot \frac{f(e_0)}{|g(e_0)|} \right) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\}) \,|\, v \in e} f(e)\\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} f(e) + \frac{f(e_0)}{|g(e_0)|} \cdot \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\} \,|\, v \in e} g(e)\\
&\quad + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \,|\, 0 \leq k \leq r-1\}) \,|\, v \in e} f(e)\\
&= \sum_{e \in E \,|\, v \in e} f(e),
\end{aligned}
$$

which implies that f' is equivalent to f. By Step (2/f), it follows inductively that f' is equivalent to the input Fractional Perfect Matching f. ■

Since f is a Fractional Perfect Matching, Claims B.3 and B.4 imply together that f' is a Fractional Perfect Matching, and the claim follows. ■

We continue with a second invariant of the algorithm IsolateOddCycles:

Lemma B.5 *For each outer loop iteration of* IsolateOddCycles*,* (a) *for each inner loop iteration, upon completion of Step* (2/e)*,* (i) $f' \subset f$*, and* (ii) *some edge from* $E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$ *is eliminated from* $E(f')$*, and* (b) *for the last inner loop iteration, upon completion of Step* (2/e)*, the non-isolated odd cycle* $\mathcal{C}$ *is eliminated from* $G(E(f'))$*.*

Similarly to Lemma B.2, the statement of Lemma B.5 (Condition (a/i)) refers to the input Fractional Perfect Matching f. The proof of Lemma B.5 will use the current Fractional Perfect Matching f; reference to f will be restored in an inductive way upon completing the proof.

Proof. Consider any outer loop iteration. For Condition (a), consider any inner loop iteration within this outer loop iteration, upon completion of Step (2/e). Consider any edge $e \in E$.

- Assume that $e \notin E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$. Then, Step (2/e) implies that $e \in E(f')$ if and only if $e \in E(f)$.
- Assume that $e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$. Then, $e \in E(f)$; so, it holds vacuously that if $e \in E(f')$ then $e \in E(f)$.

The case analysis implies that $f' \subseteq f$. Since $f'(e_0) = 0$ while $f(e_0) > 0$, this implies that $f' \subset f$. By Step (2/f), Condition (a/i) follows inductively.

Since $f'(e_0) = 0$, e_0 is eliminated from $E(f')$, so that some edge from $E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$ is eliminated from $E(f')$, and Condition (a/ii) follows.

To prove Condition (b), note that Condition (a/i) implies that there is a last inner loop iteration (and the outer loop terminates), which we consider. The precondition for the inner loop implies that some edge from $E(\mathcal{C}) \cup \{(v_0, v_1)\}$ has been eliminated from $E(f')$. Hence, the non-isolated odd cycle $\mathcal{C}$ is eliminated from $G(E(f'))$, and Condition (b) follows. ■

Lemma B.2 and Lemma B.5 (Condition (a/i)) together imply that the output f' of the algorithm IsolateOddCycles, which contains no non-isolated odd cycle, is a Fractional Perfect Matching which is equivalent to and contained in f. (By Lemma B.5 (Condition (b)), containment is strict exactly when there is at least one outer loop iteration.)

Lemma B.5 implies that at least one edge is eliminated from f in each inner loop iteration and no edge is added. Hence, there are at most $|E|$ inner loop iterations in all outer loop

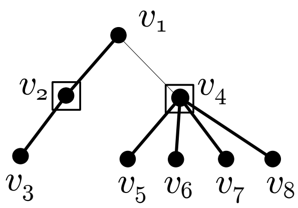

Figure 6: The graph G used in the proof of Proposition 6.6. Edges in $\mathsf{Support}_{\boldsymbol{s}}(\mathcal{D})$ are drawn thick; vertices in $\mathsf{Support}_{\boldsymbol{s}}(\mathcal{A})$ are squared.

iterations. Note that each iteration of the inner loop takes $O(|E|)$ time. Since an odd cycle is computable in polynomial time, the algorithm $\mathsf{IsolateOddCycles}$ is polynomial time.

C Proof of Proposition 5.8

Consider the graph $G = (V, E)$ in Figure 6, and fix $\alpha = 2$ and $\delta = 6$. Clearly, $\beta'(G) = 6$ and $\min_{EC \in \mathcal{EC}(G)} = 2$. So, Conditions (i) and (ii) from the claim hold. Towards a contradiction, assume that G is Pure; consider a Pure Nash equilibirum $\boldsymbol{s}$.

- By Proposition 5.3, $\mathsf{Support}_{\boldsymbol{s}}(\mathcal{D})$ is an Edge Cover. By the construction of G, this implies that $\mathsf{Supports}_{\boldsymbol{s}}(\mathcal{D}) = \{(v_2, v_3), (v_4, v_5), (v_4, v_6), (v_4, v_7), (v_4, v_8), (v_1, v)\}$, where $v \in \{v_2, v_4\}$. Since $\delta = 6$, it follows that for each edge $e \in \mathsf{Supports}_{\boldsymbol{s}}(\mathcal{D})$, there is a unique defender d such that $s_{\mathsf{d}} = e$.

- By Proposition 5.5, $\mathsf{Supports}_{\boldsymbol{s}}(\mathcal{A})$ is a Vertex Cover of the graph $G(\mathsf{Supports}_{\boldsymbol{s}}(\mathcal{D}))$. By the construction of G, this implies that $\mathsf{Supports}_{\boldsymbol{s}}(\mathcal{A}) = \{v_2, v_4\}$. (Note that $\{v_2, v_4\}$ is the *unique* Vertex Cover of the graph $G(\mathsf{Supports}_{\boldsymbol{s}}(\mathcal{D}))$ with size at most 2.) Since $\alpha = 2$, it follows that $\mathsf{Support}_{\boldsymbol{s}}(\mathsf{a}_1) = v_2$ and $\mathsf{Support}_{\boldsymbol{s}}(\mathsf{a}_2) = v_4$.

Consider now the (unique) defender $\mathsf{d} \in \mathcal{D}$ such that $s_{\mathsf{d}} = (v_4, v_5)$. Clearly, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{s}) = \frac{1}{4}$, but $\mathsf{U}_{\mathsf{d}}(\boldsymbol{s}_{-\mathsf{d}} \diamond (v_2, v_3))$ equals either $\frac{1}{3}$ or $\frac{1}{2}$, depending whether $v = v_2$ or $v = v_4$, respectively. So, $\mathsf{U}_{\mathsf{d}}(\boldsymbol{s}_{-\mathsf{d}} \diamond (v_2, v_3)) > \mathsf{U}_{\mathsf{d}}(\boldsymbol{s})$. A contradiction to the fact that $\boldsymbol{s}$ is a Nash equilibrium.