

Elsevier Editorial System(tm) for Discrete Applied Mathematics
Manuscript Draft

Manuscript Number:

Title: How Many Attackers Can Selfish Defenders Catch?

Article Type: Contribution

Keywords: Game Theory, Graph Theory, Security, Nash equilibria, Attacks and defenses

Corresponding Author: Assistant Professor Vicky G Papadopoulou, Ph.D.

Corresponding Author's Institution: European University Cyprus

First Author: Vicky G Papadopoulou, Assistant Professor

Order of Authors: Vicky G Papadopoulou, Assistant Professor; Marios Marios Mavronicolas, Professor; Burkhard Monien, Professor; Vicky G Papadopoulou, Ph.D.

Abstract:

In a distributed system with $\{it attacks\}$ and $\{it defenses,\}$ both $\{it attackers\}$ and $\{it defenders\}$ are self-interested entities. We assume a $\{it reward-sharing\}$ scheme among $\{it interdependent\}$ defenders; each defender wishes to (locally) maximize her own total $\{it fair share\}$ to the attackers extinguished due to her involvement (and possibly due to those of others). What is the $\{lem maximum\}$ amount of protection achievable by a number of such defenders against a number of attackers while the system is in a $\{it Nash equilibrium\}$? As a measure of system protection, we adopt the $\{it Defense-Ratio\}$ ~\cite{MPPS05a}, which

provides the expected (inverse) proportion of attackers caught by the defenders. In a Defense-Optimal Nash equilibrium, the Defense-Ratio is optimized.

We discover that the possibility of optimizing the Defense-Ratio (in a Nash equilibrium) 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:

`\begin{itemize}`

`\item` When the number of defenders is either sufficiently small or sufficiently large, there are cases where the Defense-Ratio can be optimized. The optimization 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.

`\item` Perhaps paradoxically, there is a middle range of values for the number of defenders where optimizing the Defense-Ratio is never possible.

`\end{itemize}`

How Many Attackers Can Selfish Defenders Catch?*

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

(NOVEMBER 21, 2008)

*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 Computer Science, Electrical Engineering and Mathematics, University of Paderborn. Email mavronic@cs.ucy.ac.cy

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

1
2
3 **Abstract**
4

5 In a distributed system with *attacks* and *defenses*, both *attackers* and *defenders* are self-
6 interested entities. We assume a *reward-sharing* scheme among *interdependent* defenders;
7 each defender wishes to (locally) maximize her own total *fair share* to the attackers extin-
8 guished due to her involvement (and possibly due to those of others). What is the *maximum*
9 amount of protection achievable by a number of such defenders against a number of attack-
10 ers while the system is in a *Nash equilibrium*? As a measure of system protection, we adopt
11 the *Defense-Ratio* [24], which provides the expected (inverse) proportion of attackers caught
12 by the defenders. In a *Defense-Optimal* Nash equilibrium, the Defense-Ratio is optimized.
13

14 We discover that the possibility of optimizing the Defense-Ratio (in a Nash equilibrium)
15 depends in a subtle way on how the number of defenders compares to two natural graph-
16 theoretic thresholds we identify. In this vein, we obtain, through a combinatorial analysis
17 of Nash equilibria, a collection of trade-off results:
18

- 19 • When the number of defenders is either sufficiently small or sufficiently large, there
20 are cases where the Defense-Ratio can be optimized. The optimization problem is
21 computationally tractable for a large number of defenders; the problem becomes \mathcal{NP} -
22 complete for a small number of defenders and the intractability is inherited from a
23 previously unconsidered combinatorial problem in *Fractional Graph Theory*.
24
- 25 • Perhaps paradoxically, there is a middle range of values for the number of defenders
26 where optimizing the Defense-Ratio is never possible.
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 Introduction

1.1 The Model and its Rationale

Safety and *security* are key issues for the design and operation of a distributed system; see, e.g., [1] or [6, Chapter 7]. Indeed, 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 whose interest is the local maximization of their own benefits rather than the optimization of global performance [2, 5, 11, 12, 13]. So, it is a challenging task to formalize and analyze the *simultaneous* impact of selfish and malicious behavior of Internet agents (cf. [17]).

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* represents a *virus*; it is a malicious client that targets a host to destroy. A *defender* is a non-malicious server representing the *antivirus software* implemented on a subnetwork in order to protect all hosts thereby connected. Here is the rationale and motivation for these modeling choices:

- Associating attacks with nodes makes sense since computer 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* [20]; this is a recently proposed, distributed *firewall architecture* where antivirus software, rather than being statically installed and licensed at a host, is implemented by a *distributed algorithm* running on a subnetwork. Such distributed implementations are attractive since they offer increased fault-tolerance and the benefit of sharing the licensing costs to the hosts.

We focus here on the simplest possible case where the subnetwork is just a *single* link; a precise understanding of the mathematical pitfalls of attacks and defenses for this simplest case is a necessary prerequisite to mastering the general case.

In reality, malicious attackers are *independent*; each (financially motivated) attacker tries to maximize on her own the amount of harm it causes during her lifetime (cf. [30]). Hence, it is natural to model each attacker as a strategic *player* seeking to maximize the chance of escaping the antivirus software; so, the strategy of one attacker does not (directly) affect the payoff of another. In contrast, there are at least three approaches to modeling the defenses:

- Defenses are *not* strategic; this approach would imply the (*centralized*) optimization problem of computing locations for the defenders that maximize the system protection given that attackers are strategic.

- Defenses are strategic and they *cooperate* to maximize the number of trapped viruses. This is modeled by assuming a *single* (strategic) defender, which centrally chooses multiple links. This approach has been pursued in [10].
- Defenses are strategic but *non-cooperative*; so, each defender still tries to maximize the number of trapped viruses she catches, while competing with the other defenders.

We have chosen to adopt the third approach. Our choice of approach is motivated as follows:

- In a large network, the defense policies are *independent* and *decentralized*. Hence, it may be not so realistic to assume that a *centralized* (even selfish) entity coordinates all defenses.
- There are financial incentives offered by hosts to defense mechanisms on the basis of the number of sustained attacks; consider, for example, the following scenarios:
 - Prices for antivirus software are determined through *recommendation systems*, which collect data from networks where scrutinized hosts were witnessed. Such price incentives induce a competition among defenders, resulting to non-cooperation.
 - Think of a *network owner* interested in maximizing the network protection. Towards that end, the owner has subcontracted the protection task to a set of independent, deployable agents. Clearly, each such agent tries to optimize the protection she offers in order to increase her reward; again, this manifests non-cooperation.

We materialize the assumption that defenses are independent and non-cooperative 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 will be rewarded with the *fair share* of the number of attackers extinguished. So, each defender is modeled as a strategic player seeking to maximize her *total* fair share to the number of extinguished attackers.

We assume two selfish species with α attackers and δ defenders; both species may use mixed strategies. Note that δ is proportional to the real cost of purchasing and installing several units of (licensed) antivirus software. The very special but yet highly non-trivial case with a single defender was originally introduced in [24] and further studied in [10, 21, 22, 23]. In a *Nash equilibrium* [26, 27], no player can unilaterally increase her (expected) *utility*.

To evaluate Nash equilibria, we employ the *Defense-Ratio*; this is the ratio of the optimum number α over the expected number of attackers extinguished by the defenders (cf. [21, 22]). Motivated by *best-case* Nash equilibria and the *Price of Stability* [3], 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 choose this value since we observe that it is a (tight) lower bound on Defense-Ratio

(Corollary 6.2). (Contrast Defense-Optimal Nash equilibria and the smallest possible Defense-Ratio to *worst-case* Nash equilibria and the *Price of Anarchy* from the seminal work of Koutsoupias and Papadimitriou [16].) A *Defense-Optimal* graph (for a given δ) is one that admits a Defense-Optimal Nash equilibrium.

1.2 Contribution

We are interested in the possibility of achieving, and the complexity of computing, a Defense-Optimal Nash equilibrium for a *given* number of defenders δ . We discover that this possibility and the associated complexity depend on δ in a quantitatively subtle way: They are determined by two graph-theoretic thresholds for δ , namely $\frac{|V|}{2}$ and $\beta'(G)$ (the size of a *Minimum Edge Cover*). (Recall that $\frac{|V|}{2} \leq \beta'(G)$.)

Our chief tool is a combinatorial characterization of the associated Nash equilibria we obtain (Proposition 5.1). For Pure Nash equilibria where both species use pure strategies, this characterization yields some interesting necessary graph-theoretic conditions for Nash equilibria (Proposition 5.9). Furthermore, this characterization yields some sufficient conditions for Defense-Optimal Nash equilibria (Theorems 6.3 and 6.5).

Our end findings are as follows:

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

- The case of *few* defenders (with $\delta \leq \frac{|V|}{2}$): We provide a combinatorial characterization of Defense-Optimal graphs (Theorem 7.4), which points out an interesting connection to *Fractional (Perfect) Matchings* [29, Chapter 2]. Roughly speaking, these graphs make a strict subset of the class 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 from the original Fractional Perfect Matching) in each partite is equal to $\frac{|V|}{2\delta}$ (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* [29].

We prove that the recognition problem for the class of graphs with a δ -Partitionable Fractional Perfect Matching is \mathcal{NP} -complete (Corollary 2.19); this intractability result holds for an *arbitrary* value of δ . Hence, so is the decision problem for a Defense-Optimal Nash equilibrium (for $\delta \leq \frac{|V|}{2}$) (Corollary 7.7). To establish the

\mathcal{NP} -completeness of the recognition problem, we develop some techniques for the reduction of Fractional (Perfect) Matchings (Section 2.3); these may be of independent interest.

We note that the recognition problem for the class of graphs with a δ -Partitionable Fractional Perfect Matching simultaneously generalizes a tractable and an intractable recognition problem: the first one concerns the class of graphs with a Perfect Matching [7], while the second concerns the class of graphs whose vertex set can be partitioned into triangles [9, GT11].

A further interesting number-theoretic consequence of the combinatorial characterization we have derived for Defense-Optimal graphs (for $\delta \leq \frac{|V|}{2}$) is that δ divides $|V|$ in a Defense-Optimal graph (Corollary 7.5).

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

- The case of *too many* defenders (with $\delta \geq \beta'(G)$): We identify two cases where there are Defense-Optimal Nash equilibria with a special structure, namely the *vertex-balanced* Nash equilibria (Definition 9.1); their structure enables their polynomial time computation (Theorems 9.2 and 9.5). The two corresponding algorithms rely on the efficient computation of *Minimum Edge Cover*; the second algorithm requires some relation between δ and α (namely, that 2δ divides α).
- The case of *many* defenders (with $\frac{|V|}{2} < \delta < \beta'(G)$): We provide a combinatorial proof that there is *no* Defense-Optimal graph for $\frac{|V|}{2} < \delta < \beta'(G)$ (Theorem 8.1). This is somehow paradoxical since with fewer defenders ($\delta \leq \frac{|V|}{2}$), we already identified cases with a Defense-Optimal Nash equilibrium. However, since the Defense-Ratio in a Defense-Optimal Nash equilibrium has a transition around the value $\delta = \frac{|V|}{2}$, this paradox may not be wholly surprising.

Our techniques have identified several new classes of graphs for any arbitrary pair of values of δ and α , such as graphs with δ -Partitionable Fractional Perfect Matching, Defense-Optimal graphs and *Pure graphs* (which admit Pure Nash equilibria); each such class was defined to support the existence of some Nash equilibria with a particular structure (for example, Defense-Optimal Nash equilibrium or Pure Nash equilibrium). Our results have revealed a fine structure among these classes, which is summarized in Figure 1.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 1: 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.

1.3 Related Work and Comparison

We emphasize that the assumption of $\delta > 1$ defenders has required a far more challenging combinatorial and graph-theoretic analysis than those used for the case of a single defender in [10, 21, 22, 23, 24]. Hence, we view this work as a major generalization of the work in [10, 21, 22, 23, 24] towards the more realistic case of $\delta > 1$ defenders.

The notion of Defense-Ratio generalizes a corresponding definition from [21, Section 3.4] to the case of $\delta > 1$ defenders. The special case where $\delta = 1$ of Theorem 7.4 was considered in [22, Corollary 2]; this case allowed for a polynomial time algorithm to decide the existence of (and compute) a Defense-Optimal Nash equilibrium by reduction to the recognition problem for a graph with a Fractional Perfect Matching. In contrast, the decision problem for a Defense-Optimal Nash equilibrium (for an arbitrary $\delta \leq \frac{|V|}{2}$) is \mathcal{NP} -complete (Corollary 7.7).

Schechter and Smith [28] considered the complementary question of determining the minimum number of defenders to catch a *single* attacker in a related model of economic threats.

1.4 Road Map

The rest of this paper is organized as follows. Section 2 collects together some background and preliminaries from Graph Theory. A preliminary combinatorial lemma is formulated and proved in Section 3. Section 4 presents the game-theoretic framework. The combinatorial structure of the associated Nash equilibria is treated in Section 5. Section 6 considers Defense-Optimal Nash equilibria. Sections 7, 8 and 9 treat the cases of few, many and too many defenders, respectively. We conclude, in Section 10, with a discussion of the results and some open problems.

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

2 Background and Preliminaries from Graph Theory

Some basic definitions are articulated in Section 2.1. Fractional Matchings are recalled in Section 2.2. Some reduction techniques for Fractional (Perfect) Matchings are developed in Section 2.3. Section 2.4 treats δ -Partitionable Fractional Perfect Matchings.

2.1 Basics

We consider a simple undirected graph $G = \langle V, E \rangle$ (with no isolated vertices). The *trivial* graph consists of a single edge. We will sometimes model an edge as the set of its two vertices.

Denote as $d_G(v)$ the *degree* of vertex v in G . An edge $(u, v) \in E$ is *pendant* if $d_G(u) = 1$ but $d_G(v) > 1$. A *path* is a sequence of vertices v_1, v_2, \dots, v_{n+1} from V such that for each index $k \in [n]$, $(v_k, v_{k+1}) \in E$; in a *cycle* \mathcal{C} , $v_{n+1} = v_1$. The cycle \mathcal{C} has length n , and \mathcal{C} is *even* (resp., *odd*) if n is even (resp., odd). A *triangle* is a cycle of length three. We shall sometimes treat \mathcal{C} as a set of vertices; $E(\mathcal{C})$ denotes the edge set induced by \mathcal{C} in the natural way.

Vertex sets and edge sets induce subgraphs in the natural way. For a vertex set $U \subseteq V$, denote as $G(U)$ the subgraph of G induced by U ; denote $\text{Edges}_G(U) = \{(u, v) \in E \mid u, v \in U\}$. For an edge set $F \subseteq E$, denote as $G(F)$ the subgraph of G induced by F ; denote $\text{Vertices}_G(F) = \bigcup_{(u,v) \in F} \{u, v\}$. (We shall sometimes omit the index G when it is clear from context.) A *component* is a maximal connected subgraph. A cycle is *isolated* (as a subgraph) if it is a component; else, it is *non-isolated*. A component is *cyclic* if it contains a cycle; else, it is *acyclic*.

For an undirected graph, both an odd and an even cycle are computable in polynomial time. A linear time algorithm to compute an odd cycle is based on incorporating *breadth-first search* into the constructive proof for the characterization of bipartite graphs due to König [14] (cf. [15, Proposition 2.27]). Polynomial time algorithms to compute an even cycle have appeared in [18, 25, 32].

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

A *Matching* is a set $M \subseteq E$ of non-incident edges; a *Maximum Matching* is one that has maximum size. The first polynomial time algorithm to compute a Maximum Matching is due to Edmonds [7]. It is known that computing a Minimum Edge Cover is polynomial time reducible to computing a Maximum Matching—see, e.g., [31, Theorem 3.1.22] or [19].

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 is one that has a Perfect Matching; note that in a Perfect-Matching graph, $\beta'(G) = \frac{|V|}{2}$. Since a Perfect Matching is a Maximum Matching, any polynomial time algorithm to compute a Maximum Matching yields a polynomial time algorithm to recognize Perfect-Matching graphs and compute a Perfect Matching.

2.2 Fractional (Perfect) Matchings

A *Fractional Matching* is a function $f : E \rightarrow [0, 1]$ where for each vertex $v \in V$, $\sum_{e \in E \mid v \in e} f(e) \leq 1$. (Matching is the special case where $f(e) \in \{0, 1\}$ for each edge $e \in E$.) For a Fractional

Matching f , induced is the set $E(f) = \{e \in E \mid f(e) > 0\}$; $|E(f)|$ is the *size* of f . The *range* of a Fractional Matching f is the set $\text{Range}(f) = \{f(e) \mid e \in E\}$; so, $\text{Range}(f) \subseteq [0, 1]$.

Given two Fractional Matchings f and f' , write $f' \subseteq f$ (resp., $f' \subset f$) if $E(f') \subseteq E(f)$ (resp., $E(f') \subset E(f)$). Say that two functions $f : E \rightarrow [0, 1]$ and $f' : E \rightarrow [0, 1]$ are *equivalent* if for each vertex $v \in V$, $\sum_{e \in E \mid v \in e} f(e) = \sum_{e \in E \mid v \in e} f'(e)$. Clearly, a function $f : E \rightarrow [0, 1]$ that is equivalent to a Fractional Matching is also a Fractional Matching.

A *Fractional Perfect Matching* is a Fractional Matching f such that for each vertex $v \in V$, $\sum_{e \in E \mid v \in e} f(e) = 1$. (Perfect Matching is the special case where $f(e) \in \{0, 1\}$ for each edge $e \in E$. Note that in this special case, $E(f)$ is a Perfect Matching; for an arbitrary Fractional Perfect Matching, $E(f)$ need not be a Perfect Matching.) In this case, for each vertex $v \in V$, there is at least one edge $e \in E$ with $v \in e$ such that $f(e) > 0$, so that $e \in E(f)$; hence, for a Fractional Perfect Matching f , $E(f)$ is an Edge Cover. Note that a function $f : E \rightarrow [0, 1]$ which is equivalent to a Fractional Perfect Matching is also a Fractional Perfect Matching.

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 $d_{G(E(f))}(u) = 1$ and $d_{G(E(f))}(v) > 1$. Since f is a Fractional Perfect Matching, $\sum_{e \in E \mid u \in e} f(e) = 1$ and $\sum_{e \in E \mid v \in e} f(e) = 1$. 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; in particular, each acyclic component of $G(E(f))$ is a single edge. The proof for [29, Theorem 2.1.5] establishes as a by-product that a Fractional Maximum Matching f with smallest size has no pendant edge; so, Lemma 2.1 provides a complementary property for the special case of Fractional Perfect Matchings.

The class of graphs with a Fractional Perfect Matching is recognizable in polynomial time via a Linear Programming formulation. (See [4] for an efficient combinatorial algorithm.) The same holds for the corresponding search problem.

2.3 Reductions of Fractional (Perfect) Matchings

Our starting point is a combinatorial property of a special case of a Fractional Maximum Matching; this property is reported in [29, Theorem 2.1.5].

1
2
3 **Proposition 2.2** Consider a Fractional Maximum Matching f with smallest size. Then, f has
4 only single edges and odd cycles.
5
6

7
8 Proposition 2.2, outlaws, in particular, the induction of even cycles and non-isolated odd cycles
9 in a Fractional Maximum Matching with smallest size. In the spirit of Proposition 2.2, we
10 shall present two new reduction techniques for a Fractional (Perfect) Matching. The first
11 reduction will eliminate all induced even cycles from an *arbitrary* Fractional Matching. The
12 second reduction is applicable only to Fractional Perfect Matchings; it will eliminate all induced
13 non-isolated odd cycles when run on a Fractional Perfect Matching with no induced even cycles.
14
15

16
17 The corresponding elimination algorithms (EliminateEvenCycles and IsolateOddCycles in Fig-
18 ures 2 and 4, respectively) are inspired from the corresponding inexistence proof for Proposition
19 2.2. In more detail, that proof assumes the existence of an even or a non-isolated odd cycle and
20 derives a contradiction by relying on the property that the Fractional Matching is a Maximum
21 one of smallest size; the contradiction is derived by eliminating edges to get a Fractional (Max-
22 imum) Matching with less size. In contrast, our elimination algorithms compute in polynomial
23 time an even or a non-isolated odd cycle (as long as there are such), respectively; they keep
24 eliminating edges (as long as possible) till there are no more even or non-isolated odd cycles,
25 respectively.
26
27
28
29

30 31 32 **2.3.1 Elimination of Even Cycles**

33
34 We prove:
35
36

37 **Proposition 2.3** Consider a Fractional Matching f . Then, there is a polynomial time algo-
38 rithm to transform f into an equivalent Fractional Matching $f' \subseteq f$ with no even cycle.
39
40

41 To prove the claim, we present the algorithm EliminateEvenCycles in Figure 2.
42

43 **Proof.** We start with a first invariant of the algorithm EliminateEvenCycles:
44
45

46 **Lemma 2.4** For each loop iteration of EliminateEvenCycles, upon completion of Step (3), f' is
47 a Fractional Matching equivalent to f .
48
49

50
51 Note that the input Fractional Matching f is already modified in the first (if any) loop iteration
52 of EliminateEvenCycles (in Step (4)), while the statement of Lemma 2.4 refers to the input
53 Fractional Matching f . The proof of Lemma 2.4 will use the current Fractional Matching f ;
54 reference to the input f will be restored in an inductive way upon completing the proof.
55
56
57
58

Algorithm EliminateEvenCycles**INPUT:** A graph $G = \langle V, E \rangle$ and a Fractional Matching f for G .**OUTPUT:** An equivalent Fractional Matching $f' \subseteq f$ with no even cycle.**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

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

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

Figure 2: The algorithm `EliminateEvenCycles`, which consists of a single loop. The precondition for the loop is the existence of an even cycle \mathcal{C} ; so, upon termination, there will be no even cycle for the output f' . (Note that if there are no loop iterations, then $f' = f$.) Step (1) chooses an edge e_0 on the cycle \mathcal{C} on which f is minimized, while Step (2) assigns a sign to each edge e on \mathcal{C} . (Since \mathcal{C} is an even cycle, alternating signs are possible.) 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. An example execution of the algorithm `EliminateEvenCycles` is illustrated in Figure 3.

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

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

If there is no edge $e \in E(\mathcal{C})$ such that $v \in e$, then $\sum_{e \in E(\mathcal{C}) | v \in e} f'(e) = \sum_{e \in E(\mathcal{C}) | v \in e} f(e) = 0$, and we are done. 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

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 3: An example execution of the algorithm `EliminateEvenCycles` on a graph with a Fractional Perfect Matching f ; the execution terminates after two loop iterations. For each loop iteration, edges in $E(f)$ are drawn thick; edges on the cycle \mathcal{C} are drawn clouded. A number next to each (thick) edge $e \in E(f)$ indicates the value $f(e)$; the sign of $g(e)$ is also indicated for each (clouded) edge e on the cycle \mathcal{C} .

Step (3),

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

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

We continue with a second invariant of the algorithm `EliminateEvenCycles`:

Lemma 2.5 *For each loop iteration of `EliminateEvenCycles`, upon completion of Step (3), (i) $f' \subset f$ and (ii) the even cycle \mathcal{C} is eliminated from $G(E(f'))$.*

Similarly to Lemma 2.4, the statement of Lemma 2.5 (Condition (i)) refers to the input Fractional Matching f . The proof of Lemma 2.5 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 `EliminateEvenCycles`, upon completion of Step (3). Consider any edge $e \in E$. We proceed by case analysis.

- Assume that $e \notin E(\mathcal{C})$. Then, Step (3), implies that $e \in E(f')$ if and only if $e \in E(f)$.
- 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 2.4 and Lemma 2.5 (Condition (i)) together imply that the output f' of algorithm `EliminateEvenCycles`, which contains no even cycle due to the loop precondition, is a Fractional

1
2
3 Matching which is equivalent to and contained in f . (By Lemma 2.5 (Condition (ii)), contain-
4 ment is strict exactly when there is at least one loop iteration.)
5

6 Lemma 2.5 (Condition (i) or (ii)) implies that at least one edge is eliminated from f in each
7 loop iteration and no edge is added. Hence, there are at most $|E|$ loop iterations. Note that
8 each loop iteration takes $O(|E|)$ time. Since an even cycle is computable in polynomial time,
9 it follows that the algorithm `EliminateEvenCycles` is polynomial time, and we are done. ■
10
11
12

13 14 **2.3.2 Elimination of Non-Isolated Odd Cycles**

15
16 We prove:
17

18
19 **Proposition 2.6** *Consider a Fractional Perfect Matching f with no even cycle. Then, there*
20 *is a polynomial time algorithm to transform f into an equivalent Fractional Perfect Matching*
21 *$f' \subseteq f$ with no non-isolated odd cycle.*
22
23

24
25 To prove the claim, we present the algorithm `IsolateOddCycles` in Figure 4.
26

27 **Proof.** Since $G(E(f))$ has no even cycle, the cycle $v_l, \dots, v_r = v_l$ determined in Step (2/a) is
28 odd. We now prove a preliminary property of the algorithm `IsolateOddCycles`:
29
30

31
32 **Lemma 2.7** *The path v_1, v_2, \dots, v_r is disjoint from $\mathcal{C} \setminus \{v_0\}$.*
33
34

35 **Proof.** By way of contradiction, assume that there is a vertex v_k with $k \in [r]$ such that
36 $v_k \in \mathcal{C} \setminus \{v_0\}$. Since \mathcal{C} has odd length, the vertices v_0 and v_k partition \mathcal{C} into two paths \mathcal{C}_1 and
37 \mathcal{C}_2 of odd and even length, respectively. Consider the two concatenations of the path v_1, \dots, v_k
38 with \mathcal{C}_1 and \mathcal{C}_2 , respectively; each of them is a cycle in $G(E(f))$ and one of them has even
39 length. A contradiction. ■
40
41

42 We start with a first invariant of the algorithm `IsolateOddCycles`.
43
44

45
46 **Lemma 2.8** *For each inner loop iteration in an outer loop iteration of `IsolateOddCycles`, upon*
47 *completion of Step (2/e), f' is a Fractional Perfect Matching equivalent to f .*
48
49

50 Note that the input Fractional Perfect Matching f is already modified in the first inner loop
51 iteration in the first outer loop iteration of `IsolateOddCycles` (in Step (2/f)). Reminiscent of
52 Lemma 2.4, the statement of Lemma 2.8 refers to the input Fractional Perfect Matching f .
53 The proof of Lemma 2.8 will use the current Fractional Perfect Matching f ; reference to the
54 input f will be restored in an inductive way upon completing the proof.
55
56
57
58

Algorithm IsolateOddCycles**INPUT:** A graph $G = \langle V, E \rangle$ 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.**While** $G(E(f))$ contains a non-isolated odd cycle \mathcal{C} **do**(1) Choose a vertex $v_0 \in \mathcal{C}$ with $d_{G(E(f))}(v_0) \geq 3$ 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, \dots, 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 $e_0 \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}$ that realizes the quantity

$$f_0 := \min \left\{ \min_{e \in E(\mathcal{C})} \frac{f(e)}{|g(e)|}, \min_{\substack{l > 0 \\ 0 \leq k \leq l-1}} \frac{f((v_k, v_{k+1}))}{|g((v_k, v_{k+1}))|}, \min_{l \leq k \leq r-1} \frac{f((v_k, v_{k+1}))}{|g((v_k, v_{k+1}))|} \right\};$$

(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 4: The algorithm `IsolateOddCycles`, which consists of an outer loop; the outer loop includes an inner loop (Step (2)). The precondition for the outer loop is the existence of a non-isolated odd cycle; so, upon termination, there will be no non-isolated odd cycle for f' . (Note that if there are no (outer) loop iterations, then $f' = f$.) For Step (1), note that a vertex $v_0 \in \mathcal{C}$ with $d_{G(E(f))}(v_0) \geq 3$ 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 continues) 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, \dots, 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), note that Lemma 2.7 implies that for any vertex v_k with $0 < k \leq r$, it holds that $v_k \notin \mathcal{C} \setminus \{v_0\}$. 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 the cycle \mathcal{C} or the outgoing path v_1, v_2, \dots, v_r that minimizes a certain quantity f_0 determined from f and g ; so, $f_0 = \frac{f(e_0)}{|g(e_0)|}$. 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. An example execution of the algorithm `IsolateOddCycles` is illustrated in Figure 5.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 5: An example execution of the algorithm `IsolateOddCycles` on the graph with a Fractional Perfect Matching f (with no induced even cycle) from Figure 3(c). The execution terminates after two outer loop iterations; the first outer loop iteration incurs one inner loop iteration, while the second outer loop iteration incurs two inner loop iterations. For each (inner or outer) loop iteration, edges in $E(f)$ are drawn thick; edges on the cycle \mathcal{C} are drawn clouded. A number next to each (thick) edge $e \in E(f)$ indicates $f(e)$; the sign of $g(e)$ (for each edge e on the cycle \mathcal{C}), the vertex v_0 and the edge e_0 are also indicated for each iteration.

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 2.9 $\text{Range}(f') \subseteq [0, 1]$.

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 ,

$$\begin{aligned} f'(e) &= f(e) - |g(e)| \cdot \frac{f(e_0)}{|g(e_0)|} \\ &\geq 0, \end{aligned}$$

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, \dots, v_l (with $l > 0$) or both on the cycle $v_l, \dots, v_r = v_l$.

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

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

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, \dots, v_l (with $l > 0$) or on the cycle $v_l, \dots, v_r = v_l$. Since both the cycle \mathcal{C} and the cycle $v_l, \dots, v_r = v_l$ 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 $v_l, \dots, v_r = v_l$. Hence, edge e lies on the path v_0, \dots, 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,

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

as needed. The proof is now complete. ■

We continue with the second technical claim:

Claim 2.10 f' is equivalent to f .

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

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

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

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

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

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

Hence, by Step (2/e),

$$\begin{aligned}
\sum_{e \in E | v \in e} f'(e) &= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} | v \in e} f'(e) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}) | v \in e} f(e) \\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} | v \in e} (f(e) + g(e) \cdot f_0) + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}) | v \in e} f(e) \\
&= \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} | v \in e} f(e) + f_0 \cdot \sum_{e \in E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\} | v \in e} g(e) \\
&\quad + \sum_{e \in E \setminus (E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq k \leq r-1\}) | v \in e} f(e) \\
&= \sum_{e \in E | v \in e} f(e),
\end{aligned}$$

1
2
3 which implies that f' is equivalent to f . By Step (2/f), it follows inductively that f' is
4 equivalent to the input Fractional Perfect Matching f . ■

5
6 Since f is a Fractional Perfect Matching, Claims 2.9 and 2.10 imply together that f' is a
7 Fractional Perfect Matching, and the claim follows. ■

8
9 We continue with a second invariant of the algorithm `IsolateOddCycles`:

10
11 **Lemma 2.11** *For each outer loop iteration of `IsolateOddCycles`, (a) for each inner loop itera-*
12 *tion, upon completion of Step (2/e), (i) $f' \subset f$, and (ii) some edge from $E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid 0 \leq$*
13 *$k \leq r - 1\}$ is eliminated from $E(f')$, and (b) for the last inner loop iteration, upon completion*
14 *of Step (2/e), the non-isolated odd cycle \mathcal{C} is eliminated from $G(E(f'))$.*

15
16 Similarly to Lemma 2.8, the statement of Lemma 2.11 (Condition (a/i)) refers to the input
17 Fractional Perfect Matching f . The proof of Lemma 2.11 will use the current Fractional Perfect
18 Matching f ; reference to the input f will be restored in an inductive way upon completing the
19 proof.

20
21 **Proof.** Consider any outer loop iteration. For Condition (a), consider any inner loop iteration
22 within this outer loop iteration, upon completion of Step (2/e). Consider any edge $e \in E$.

- 23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
- 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 $e \in E(f)$.

38
39 The case analysis implies that $f' \subseteq f$. Since $f'(e_0) = 0$ while $f(e_0) > 0$, this implies that
40 $f' \subset f$. By Step (2/f), Condition (a/i) follows inductively.

41
42 Since $f'(e_0) = 0$, e_0 is eliminated from $E(f')$, so that some edge from $E(\mathcal{C}) \cup \{(v_k, v_{k+1}) \mid$
43 $0 \leq k \leq r - 1\}$ is eliminated from $E(f')$, and Condition (a/ii) follows.

44
45 To prove Condition (b), note that Condition (a/i) implies that there is a last inner loop
46 iteration (and the outer loop terminates). So, consider the last inner loop iteration. The
47 precondition for the inner loop implies that some edge from $E(\mathcal{C}) \cup \{(v_0, v_1)\}$ has been eliminated
48 from $E(f')$. Hence, the non-isolated odd cycle \mathcal{C} is eliminated from $G(E(f'))$, and Condition
49 (b) follows. ■

50
51
52 Lemma 2.8 and Lemma 2.11 (Condition (a/i)) together imply that the output f' of the
53 algorithm `IsolateOddCycles`, which contains no non-isolated odd cycle due to the outer loop
54 precondition, is a Fractional Perfect Matching which is equivalent to and contained in f . (By
55
56
57

1
2
3 Lemma 2.11 (Condition (b)), containment is strict exactly when there is at least one outer loop
4 iteration.)
5

6 Lemma 2.11 implies that at least one edge is eliminated from f in each inner loop iteration
7 and no edge is added. Hence, there are at most $|E|$ inner loop iterations in all outer loop
8 iterations. Note that each iteration of the inner loop takes $O(|E|)$ time. Since an odd cycle
9 is computable in polynomial time, it follows that the algorithm `IsolateOddCycles` is polynomial
10 time, and we are done. ■
11
12
13

14 15 **2.3.3 Recap**

16
17 We are now ready to prove:
18
19
20

21 **Proposition 2.12** *Consider a Fractional Perfect Matching f . Then, there is a polynomial*
22 *time algorithm to transform f into an equivalent Fractional Perfect Matching $f' \subseteq f$ with only*
23 *single edges and odd cycles.*
24
25

26
27 To prove the claim, we present the algorithm `EliminateEven&IsolateOddCycles` in Figure 6. The
28 algorithm is the sequential cascade of the algorithms `EliminateEvenCycles` and `IsolateOddCycles`
29 from Figures 2 and 4, respectively.
30
31

32 **Algorithm** `EliminateEven&IsolateOddCycles`

33 INPUT: A Fractional Perfect Matching f for a graph G .

34 OUTPUT: A Fractional Perfect Matching f' with no even cycle and no non-isolated odd cycle.

- 35
36 (1) Apply `EliminateEvenCycles` on f to obtain f'' .
37
38 (2) Apply `IsolateOddCycles` on f'' to obtain f' .
39
40

41 Figure 6: The algorithm `EliminateEven&IsolateOddCycles`, incorporating the algorithms `Elimi-`
42 `nateEvenCycles` and `IsolateOddCycles` from Figures 2 and 4, respectively.
43
44

45 **Proof.** By Proposition 2.3, $f'' \subseteq f$ is a Fractional Perfect Matching with no even cycle,
46 which is equivalent to f . By Proposition 2.6, $f' \subseteq f''$ is a Fractional Perfect Matching with no
47 non-isolated odd cycle, which is equivalent to f'' . It follows that (1) f' is equivalent to f and
48 $f' \subseteq f$, and (2) f' has no even cycle and no non-isolated odd cycle. Since f is a Fractional
49 Perfect Matching, Condition (1) implies that f' is a Fractional Perfect Matching; by Lemma
50 2.1, this implies that each acyclic component of $G(E(f'))$ is a single edge. By Condition (2),
51 it follows that each cyclic component of $G(E(f'))$ is an (isolated) odd cycle. It follows that f'
52 consists of single edges and odd cycles, and the proof is complete. ■
53
54
55
56
57
58

2.4 δ -Partitionable Fractional Perfect Matchings

2.4.1 Definition and Preliminaries

We introduce a special class of Fractional Perfect Matchings:

Definition 2.1 Fix an integer $\delta \geq 1$. A Fractional Perfect Matching $f : E \rightarrow \mathbb{R}$ is δ -**Partitionable** if the edge set $E(f)$ can be partitioned into δ (non-empty,) vertex-disjoint partites E_1, \dots, E_δ so that for each partite E_j with $j \in [\delta]$, $\sum_{e \in E_j} f(e) = \frac{|V|}{2\delta}$.

Note that a 1-Partitionable Fractional Perfect Matching is a Fractional Perfect Matching. Hence, the decision problem for a 1-Partitionable Fractional Perfect Matching is solved in polynomial time. Note also that the restriction of a δ -Partitionable Fractional Perfect Matching to each partite E_j with $j \in [\delta]$ is a Fractional Perfect Matching; so, for each partite E_j with $j \in [\delta]$, for each vertex $v \in V(E_j)$, $\sum_{e \in E_j | v \in e} f(e) = 1$. Since the partites are vertex-disjoint, this implies that for each partite E_j with $j \in [\delta]$, for each vertex $v \in V(E_j)$, $\sum_{e \in E_j | v \in e} f(e) = 1$. We now prove a necessary condition for a δ -Partitionable Fractional Perfect Matching:

Proposition 2.13 Consider a δ -Partitionable Fractional Perfect Matching f . Then, for each partite E_j with $j \in [\delta]$, $|V(E_j)| = \frac{|V|}{\delta}$.

Proof. Fix a partite E_j with $j \in [\delta]$. Then,

$$\begin{aligned} \sum_{e \in E_j} f(e) &= \frac{1}{2} \sum_{v \in V(E_j)} \left(\sum_{e \in E_j | v \in e} f(e) \right) \\ &= \frac{1}{2} \sum_{v \in V(E_j)} 1 \\ &= \frac{|V(E_j)|}{2}. \end{aligned}$$

Since f is δ -Partitionable, it follows that

$$\begin{aligned} |V(E_j)| &= 2 \cdot \sum_{e \in E_j} f(e) \\ &= 2 \cdot \frac{|V|}{2\delta} \\ &= \frac{|V|}{\delta}, \end{aligned}$$

as needed. ■

Proposition 2.13 immediately implies:

Corollary 2.14 *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 necessary condition $\delta \leq \frac{|V|}{2}$ in Corollary 2.14 is not always necessary:

Proposition 2.15 *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 and fix $\delta = 1$. Clearly, $\beta'(\mathcal{C}_3) = 2$ so that $\delta < \frac{|V|}{2}$. Consider the function $f : E(\mathcal{C}_3) \rightarrow [0, 1]$ with $f(e) = \frac{1}{2}$ for each edge $e \in E(\mathcal{C}_3)$. Clearly, f is a 1-Partitionable Fractional Perfect Matching, and the claim follows. ■

We finally prove that the equivalence relation on Fractional Perfect Matchings preserves δ -Partitionability under a certain containment assumption:

Proposition 2.16 *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, \dots, E_δ . Define edge sets E'_1, \dots, E'_δ 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 $j \in [\delta]$, $E'_j \subseteq E_j$. This implies that the collection E'_1, \dots, E'_δ partitions $E(f')$. Since the partites E_1, \dots, E_δ are vertex-disjoint, this also implies that the edge sets E'_1, \dots, E'_δ are vertex-disjoint; so call them partites. Fix any partite E'_j with $j \in [\delta]$. Then,

$$\begin{aligned}
& \sum_{e \in E'_j} f'(e) \\
&= \sum_{e \in E_j} f'(e) \quad (\text{since } f'(e) = 0 \text{ for each } e \in E_j \setminus E'_j) \\
&= \frac{1}{2} \cdot \sum_{v \in V(E_j)} \sum_{e \in E_j \mid v \in e} f'(e) \\
&= \frac{1}{2} \cdot \sum_{v \in V(E_j)} \sum_{e \in E \mid v \in e} f'(e) \\
&= \frac{1}{2} \cdot \sum_{v \in V(E_j)} \sum_{e \in E \mid v \in e} f(e) \quad (\text{since } f \text{ and } f' \text{ are equivalent}) \\
&= \frac{1}{2} \cdot \sum_{v \in V(E_j)} 1 \quad (\text{since } f \text{ is Perfect}) \\
&= \frac{1}{2} \cdot |V(E_j)| \\
&= \frac{|V|}{2\delta} \quad (\text{by Proposition 2.13}).
\end{aligned}$$

Hence, f' is δ -Partitionable, as needed. ■

2.4.2 Characterization

We show:

Proposition 2.17 *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, \dots, E_δ such that (1) $\bigcup_{j \in [\delta]} E_j$ is an Edge Cover, and (2) for each edge set E_j with $j \in [\delta]$, (i) E_j consists of single edges and odd cycles, and (ii) $|V(E_j)| = \frac{|V|}{\delta}$.*

Note that the edge sets E_1, \dots, E_δ need not form a partition of E ; in contrast, by Condition (1), the induced vertex sets $V(E_1), \dots, V(E_\delta)$ are required to form a partition of V .

Proof. Assume first that G has a δ -Partitionable Fractional Perfect Matching f . By Proposition 2.12, 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 2.16 implies that f' is δ -Partitionable. So, the edge set $E(f')$ can be partitioned into δ (non-empty), vertex-disjoint partites E_1, \dots, E_δ 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, \dots, E_δ . Since f' is a Fractional Perfect Matching, $E(f')$ is an Edge Cover; so, $\bigcup_{j \in [\delta]} E_j$ is an Edge Cover. Consider now any edge set E_j with $j \in [\delta]$. Since f' consists of single edges and odd cycles, Condition (2/i) follows; since f' is a Fractional Perfect Matching, Condition (2/ii) follows from Proposition 2.13.

Assume now that E contains a collection of δ (non-empty,) vertex-disjoint edge sets E_1, \dots, E_δ such that (1) $\bigcup_{j \in [\delta]} E_j$ is an Edge Cover, and (2) for each edge set E_j with $j \in [\delta]$, (i) E_j is a collection of single edges and odd cycles, and (ii) $|V(E_j)| = \frac{|V|}{\delta}$. We shall prove that G has a δ -Partitionable Fractional Perfect Matching f . The proof is constructive. Define the function $f : E \rightarrow [0, 1]$ with

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

To prove that f is a Fractional Perfect Matching, consider any vertex $v \in V$. Since $\bigcup_{j \in [\delta]} E_j$ is an Edge Cover, this implies that $v \in V(E_j)$ for some $j \in [\delta]$. There are two cases:

- Assume that E_j is a single edge e_j . Then, by construction, $\sum_{e \in E|v \in e} f(e) = f(e_j) = 1$.
- Assume that E_j is an (isolated) odd cycle, so that $v = e_j \cap e'_j$ for a pair of consecutive edges e_j, e'_j on the cycle. Then, by construction, $\sum_{e \in E|v \in e} f(e) = f(e_j) + f(e'_j) = \frac{1}{2} + \frac{1}{2} = 1$.

The case analysis implies that f is a Fractional Perfect Matching. To prove that f is δ -Partitionable, consider the partites E_1, \dots, E_δ . Fix any partite E_j with $j \in [\delta]$. Then,

$$\begin{aligned}
& \sum_{e \in E_j} f(e) \\
&= \frac{1}{2} \sum_{v \in V(E_j)} \sum_{e \in E_j | v \in e} f(e) \\
&= \frac{1}{2} \sum_{v \in V(E_j)} 1 \quad (\text{since } f \text{ is Perfect}) \\
&= \frac{1}{2} |V(E_j)| \\
&= \frac{|V|}{2\delta} \quad (\text{by Condition (2/ii)}).
\end{aligned}$$

It follows that f is δ -Partitionable, and the proof is now complete. \blacksquare

We observe an interesting special case of Proposition 2.17:

Proposition 2.18 *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. Proposition 2.17 implies that E contains a collection of $\frac{|V|}{2}$ (non-empty,) vertex-disjoint edge sets $E_1, \dots, E_{\lfloor \frac{|V|}{2} \rfloor}$ such that (1) $\bigcup_{j \in [\lfloor \frac{|V|}{2} \rfloor]} E_j$ is an Edge Cover, and (2) for each edge set E_j with $j \in [\lfloor \frac{|V|}{2} \rfloor]$, (i) E_j consists of single edges and odd cycles, and (ii) $|V(E_j)| = 2$. By Conditions (2/i) and (2/ii), it follows that each edge set E_j with $j \in [\delta]$ is a single edge. Hence, the collection of the (vertex-disjoint) edge sets is a Matching. By Condition (1), this implies that the collection of the edge sets is a Perfect Matching.

Assume that G is Perfect-Matching with a Perfect Matching M . Consider the *indicator function* $f : E \rightarrow \{0, 1\}$ 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 $e_j \in M$ with $j \in [\frac{|V|}{2}]$, define the partite $E_j := \{e_j\}$. Since M is a Perfect Matching, the partites are vertex-disjoint. So, for each E_j with $j \in [\delta]$, $\sum_{e \in E_j} f(e) = 1 = \frac{|V|}{2 \frac{|V|}{2}}$, and this completes the proof. \blacksquare

2.4.3 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 ?

Note that the restriction to instances for which δ divides $|V|$ is inherited from Corollary 2.14 in order to exclude the non-interesting instances.

Proposition 2.18 identifies a tractable special case of δ -PARTITIONABLE FPM (namely, $\frac{|V|}{2}$ -PARTITIONABLE FPM). We shall use Proposition 2.17 to show that in the general case where δ is arbitrary, δ -PARTITIONABLE FPM is \mathcal{NP} -complete. To do so, we shall observe an interesting relation of some other (intractable) special case of the problem to a well known 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, \dots, V_δ , each containing exactly three vertices, such that for each $j \in [\delta]$, $E(V_j)$ is a triangle?

This problem is known to be \mathcal{NP} -complete [9, GT11, attribution to (personal communication with) Schaefer]. (This restriction to graphs $G = \langle V, E \rangle$ with $|V| = 3\delta$ is made in order to exclude the non-interesting instances.) To prove that δ -PARTITIONABLE FPM is \mathcal{NP} -complete (for an arbitrary δ), we consider the special case of it 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 graphs $G = \langle V, E \rangle$ with $|V| = 3\delta$ is necessary since $\delta = \frac{|V|}{3}$ is an integer.) 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 2.19 $\frac{|V|}{3}$ -PARTITIONABLE FPM = 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 2.17, E contains a collection of $\frac{|V|}{3}$ (non-empty,) vertex-disjoint edge sets $E_1, \dots, E_{\frac{|V|}{3}}$ such that (1) $\bigcup_{j \in [\frac{|V|}{3}]} E_j$ is an Edge Cover, and (2) 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 with $j \in [\frac{|V|}{3}]$ is a triangle. This implies that G is a positive instance for PARTITION INTO TRIANGLES (with vertex sets $V(E_1), \dots, V(E_{\frac{|V|}{3}})$).

1
2
3 Assume now that G is a positive instance for PARTITION INTO TRIANGLES. Consider
4 the corresponding partition of V into $\delta = \frac{|V|}{3}$ disjoint vertex sets $V_1, \dots, V_{\frac{|V|}{3}}$. This par-
5
6
7 tion induces a corresponding partition of E into a collection of $\frac{|V|}{3}$ vertex-disjoint partites
8
9 $E_1, \dots, E_{\frac{|V|}{3}}$, where each partite E_j is a single triangle. Proposition 2.17 implies that G has
10
11 a $\frac{|V|}{3}$ -Partitionable Fractional Perfect Matching. Hence, G is a positive instance for $\frac{|V|}{3}$ -
12
13 PARTITIONABLE FPM, and we are done. ■

14
15 By Proposition 2.19, it follows that $\frac{|V|}{3}$ -PARTITIONABLE FPM is \mathcal{NP} -complete. Since $\frac{|V|}{3}$ -
16
17 PARTITIONABLE FPM is a special case of δ -PARTITIONABLE FPM, this implies:
18

19 **Corollary 2.20** δ -PARTITIONABLE FPM is \mathcal{NP} -complete.
20

21 22 23 3 A Combinatorial Lemma

24
25
26 In this section, we prove a combinatorial lemma that will be useful later.
27

28 For a *probability* x , we define two *probability literals*, or *literals* for short: the *positive* literal
29
30 x and the *negative* literal $\bar{x} = 1 - x$. A *probability product*, or *product* for short, is a product
31
32 of probability literals $x_1 \cdots x_n$ for any $n \geq 1$; we adopt the convention that $x_{\ell_1} \cdots x_{\ell_2} = 1$
33
34 whenever $\ell_2 < \ell_1$. A *constant* probability product is the trivial one which equals to 1 and has
35
36 no literals. The *expansion* of a probability product is obtained when substituting each negative
37
38 literal \bar{x} with $1 - x$. So, an expansion contains positive literals and no negative literals.

39 The probability product $x_1 \cdots x_n$ is *positive* if all its probability literals are positive. More
40
41 generally, for any integer $\ell \leq n$, the probability product $x_1 \cdots x_n$ is ℓ -*positive* if exactly ℓ of its
42
43 probability literals are positive; so, an n -positive probability product is a positive probability
44
45 product. For each $\ell \in [n]$, denote as $\text{Pos}_\ell(x_1, \dots, x_n)$ the collection of all ℓ -positive probability
46
47 products with literals defined from the probabilities x_1, \dots, x_n . We prove a combinatorial
48
49 identity for sums of probability products:

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

$$51 \sum_{\ell \in [n]} \frac{1}{\ell} \cdot \sum_{x_2 \dots x_n \in \text{Pos}_{\ell-1}(x_2, \dots, x_n)} x_2 \cdots x_n = \sum_{\ell \in [n]} (-1)^{\ell-1} \cdot \frac{1}{\ell} \cdot \sum_{x_2 \dots x_n \in \text{Pos}_{\ell-1}(x_2, \dots, x_n)} x_2 \cdots x_n .$$

52
53 Note that the right-hand side (RHS) is a weighted sum of positive probability products, with
54
55 weights of alternating signs. In contrast, the left-hand side (LHS) is a weighted sum of arbitrary
56
57

(not necessarily positive) probability products, with positive weights; an $(\ell-1)$ -positive product in the LHS is multiplied by $\frac{1}{\ell}$.

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

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

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

$$\sum_{x_2 \cdots x_n \in \text{Pos}_{\ell-1}(x_2, \dots, x_n)} x_2 \cdots x_n \Big|_{\ell=1} = \bar{x}_2 \cdots \bar{x}_n.$$

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 $x_2 \cdots x_\ell \in \text{Pos}_{\ell-1}(x_2, \dots, x_n)$ from the sum $\sum_{x_2 \cdots x_\ell \in \text{Pos}_{\ell-1}(x_2, \dots, x_n)} x_2 \cdots x_\ell$ in the RHS. Note that all products in $\text{Pos}_{\ell-1}(x_2, \dots, x_n)$ (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 expansion of the LHS.

Clearly, a k -positive product with $k \geq \ell$ in the LHS cannot include $x_2 \cdots x_\ell$ in its expansion. So, 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 $x_2 \cdots x_\ell$ in the expansion of the LHS.

- Note that there are $\binom{\ell-1}{k}$ ways to choose k positive literals (or $\ell - 1 - k$ negative literals) out of the $(\ell - 1)$ literals x_2, \dots, x_ℓ in order to form a k -positive product that includes $x_2 \cdots x_\ell$ (multiplied with a coefficient) in its expansion. (All literals $x_{\ell+1}, \dots, x_n$ have to be negative since they do not appear in the product $x_2 \cdots x_\ell$.)
- The sign of the resulting k -positive product is $(-1)^{(\ell-1)-k}$, since each of the $(\ell-1)-k$ negative literals in it contributes one minus sign. (The negative literals $x_{\ell+1}, \dots, x_n$ 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 expansion of the LHS is

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

and the claim follows for $\ell \geq 2$.

The proof is now complete. ■

4 Game-Theoretic Framework

Section 4.1 introduces the strategic game $AD_{\alpha,\delta}(G)$. The associated pure Nash equilibria are defined in Section 4.2. Section 4.3 considers mixed profiles; their associated Expected Utilities are treated in Section 4.4. (Mixed) Nash equilibria are introduced in Section 4.5. Some special profiles and corresponding special classes of Nash equilibria are treated in Section 4.6. Some notation is articulated in Section 4.7.

4.1 The Strategic Game $AD_{\alpha,\delta}(G)$

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

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- The set of *players* is $\mathcal{A} \cup \mathcal{D}$; \mathcal{A} contains α *attackers* a_i with $i \in [\alpha]$, and \mathcal{D} contains δ *defenders* d_j with $j \in [\delta]$.

- The *strategy set* S_a of each attacker a is V ; the *strategy set* S_d of each defender d is E . So, the *strategy space* S is $S = (\times_{a \in \mathcal{A}} S_a) \times (\times_{d \in \mathcal{D}} S_d) = V^\alpha \times E^\delta$.

A *profile* (or *pure profile*) is an $(\alpha + \delta)$ -tuple $\mathbf{s} = \langle s_{a_1}, \dots, s_{a_\alpha}, s_{d_1}, \dots, s_{d_\delta} \rangle \in S$. The profile $\mathbf{s}_{-b} \diamond t_b$ is obtained from the profile \mathbf{s} and a strategy t_b for player $b \in \mathcal{A} \cup \mathcal{D}$ by substituting t_b for s_b in the profile \mathbf{s} .

For each vertex $v \in V$, $A_{\mathbf{s}}(v) = \{a \in \mathcal{A} \mid s_a = v\}$ and $D_{\mathbf{s}}(v) = \{d \in \mathcal{D} \mid v \in s_d\}$. Assume that $v \in s_d$. Then, the *proportion* $\text{Prop}_{\mathbf{s}}(d, v)$ of defender d on vertex v in the profile \mathbf{s} is given by $\text{Prop}_{\mathbf{s}}(d, v) = \frac{1}{|D_{\mathbf{s}}(v)|}$.

- The *Utility* of attacker a is a function $U_a : S \rightarrow \{0, 1\}$ with

$$U_a(\mathbf{s}) = \begin{cases} 0, & \text{if } s_a \in s_d \text{ for some defender } d \in \mathcal{D} \\ 1, & \text{if } s_a \notin s_d \text{ for every defender } d \in \mathcal{D} \end{cases}.$$

Intuitively, when the attacker a chooses vertex v , she receives 0 if it is caught by a defender; otherwise, she receives 1.

- The *Utility* of defender d is a function $U_d : S \rightarrow \mathbb{Q}$ with

$$U_d(\mathbf{s}) = \frac{|A_{\mathbf{s}}(u)|}{|D_{\mathbf{s}}(u)|} + \frac{|A_{\mathbf{s}}(v)|}{|D_{\mathbf{s}}(v)|},$$

where $s_d = (u, v)$. Intuitively, 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

The profile \mathbf{s} is a *Pure Nash equilibrium* [26, 27] if for each player $b \in \mathcal{A} \cup \mathcal{D}$, for each strategy $t_b \in S_b$, $U_b(\mathbf{s}) \geq U_b(\mathbf{s}_{-b} \diamond t_b)$; 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 $\text{AD}_{\alpha, \delta}(G)$.

4.3 Mixed Profiles

A *mixed strategy* for a player is a probability distribution over her strategy set; so, a mixed strategy for an attacker (resp., a defender) is a probability distribution over vertices (resp., edges). A *mixed profile* (or *profile* for short) $\sigma = \langle \sigma_{a_1}, \dots, \sigma_{a_\alpha}, \sigma_{d_1}, \dots, \sigma_{d_\delta} \rangle$ is a collection

of mixed strategies, one for each player; $\sigma_a(v)$ is the probability that attacker a chooses vertex v , and $\sigma_d(e)$ is the probability that defender d chooses edge e .

4.3.1 Supports

Fix now a mixed profile σ . The **support** of player $b \in \mathcal{A} \cup \mathcal{D}$ in the profile σ , denoted as $\text{Support}_\sigma(b)$, is the set of pure strategies in S_b to which b assigns strictly positive probability. Denote $\text{Supports}_\sigma(\mathcal{A}) = \bigcup_{a \in \mathcal{A}} \text{Support}_\sigma(a)$ and $\text{Supports}_\sigma(\mathcal{D}) = \bigcup_{d \in \mathcal{D}} \text{Support}_\sigma(d)$. A mixed profile σ induces a probability measure \mathbb{P}_σ (on pure profiles) in the natural way. Note that in a pure profile s , $\text{Supports}_s(\mathcal{A}) \leq \alpha$ and $\text{Supports}_s(\mathcal{D}) \leq \delta$.

4.3.2 Expectations about Attackers

For each vertex $v \in V$, denote as $|A|_\sigma(v)$ the expected number of attackers choosing vertex v in σ ; so,

$$|A|_\sigma(v) = \sum_{a \in \mathcal{A}} \sigma_a(v).$$

Clearly, $|A|_\sigma(v) > 0$ if and only if $v \in \text{Supports}_\sigma(\mathcal{A})$. For an edge $(u, v) \in E$, denote

$$|A|_\sigma((u, v)) = |A|_\sigma(u) + |A|_\sigma(v).$$

We observe:

Observation 4.1 *For a mixed profile σ ,*

$$\sum_{v \in \text{Supports}_\sigma(\mathcal{A})} |A|_\sigma(v) = \alpha.$$

Proof. Clearly,

$$\begin{aligned} \sum_{v \in \text{Supports}_\sigma(\mathcal{A})} |A|_\sigma(v) &= \sum_{v \in \text{Supports}_\sigma(\mathcal{A})} \sum_{a \in \mathcal{A}} \sigma_a(v) \\ &= \sum_{a \in \mathcal{A}} \sum_{v \in \text{Supports}_\sigma(\mathcal{A})} \sigma_a(v) \\ &= \sum_{a \in \mathcal{A}} 1 \\ &= \alpha, \end{aligned}$$

as needed. ■

4.3.3 Hitting Events and Vertices

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

$$\mathbb{P}_\sigma(\text{Hit}(d, v)) = \sum_{e \in \text{Support}_\sigma(d) | v \in e} \sigma_d(e).$$

Denote as $\text{Hit}(v)$ the event that some defender chooses an edge incident to vertex v . Clearly,

$$\text{Hit}(v) = \bigcup_{d \in \mathcal{D}} \text{Hit}(d, v).$$

Finally, denote as $D_\sigma(v)$ the set

$$D_\sigma(v) = \left\{ d \in \mathcal{D} \mid \text{there is an edge } e \in \text{Support}_\sigma(d) \text{ such that } v \in e \right\};$$

so, $D_\sigma(v)$ is the set of defenders “hitting” vertex v .

A vertex $v \in V$ is **multidefender** in the profile σ if $|D_\sigma(v)| \geq 2$; that is, a multidefender vertex is “hit” by more than one defenders. A vertex $v \in V$ is **undefender** in σ if $|D_\sigma(v)| \leq 1$; v is **monodefender** in σ if $|D_\sigma(v)| = 1$. So, for each undefender (resp., monodefender) vertex v , there is at most (resp., exactly) one defender d with an edge $e \in \text{Support}_\sigma(d)$ such that $v \in e$; if there is such a defender, denote it as $d_\sigma(v)$, else, set, by convention, $\mathbb{P}_\sigma(\text{Hit}(d_\sigma(v), v)) = 0$.

A profile σ is **undefender** (resp., **monodefender**) if every vertex $v \in V$ is undefender (resp., monodefender) in σ ; else the profile σ is **multidefender**. Note that for a undefender (resp., monodefender) profile σ , for each edge $e \in E$, there is at most (resp., exactly) one defender d such that $\sigma_d(e) > 0$; if there is such a defender d , denote it as $d_\sigma(e)$, else set, by convention, $\mathbb{P}_\sigma(d_\sigma(e), e) = 0$.

4.3.4 Hitting Probabilities

Since the events $\text{Hit}(d_j, v)$ and $\text{Hit}(d_{j'}, v)$ 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 $v \in V$,*

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

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

1
2
3 **Lemma 4.1** For a vertex $v \in V$,

$$4 \quad \mathbb{P}_\sigma(\text{Hit}(v)) = \sum_{l \in [\delta]} (-1)^{l-1} \sum_{\mathcal{D}' \subseteq \mathcal{D} \mid |\mathcal{D}'|=l} \prod_{d \in \mathcal{D}'} \mathbb{P}_\sigma(\text{Hit}(d, v)).$$

5
6
7
8 We continue to prove:

9
10
11 **Lemma 4.2** For a mixed profile σ ,

$$12 \quad \sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) \begin{cases} < 2\delta, & \text{if } \sigma \text{ is multidefender} \\ = 2\delta, & \text{if } \sigma \text{ is undefender} \end{cases}.$$

13
14
15
16
17 **Proof.** Clearly,

$$18 \quad \begin{aligned} 19 \quad \sum_{v \in V} \sum_{d \in \mathcal{D}} \mathbb{P}_\sigma(\text{Hit}(d, v)) &= \sum_{v \in V} \sum_{d \in \mathcal{D}} \sum_{e \in \text{Support}_\sigma(d) \mid v \in e} \sigma_d(e) \\ 20 &= 2 \sum_{e \in E} \sum_{d \in \mathcal{D}} \sigma_d(e) \\ 21 &= 2 \sum_{d \in \mathcal{D}} \sum_{e \in E} \sigma_d(e) \\ 22 &= 2\delta. \end{aligned}$$

23
24
25
26
27
28 Hence, by Observation 4.2,

$$29 \quad \sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) \begin{cases} < \sum_{v \in V} \sum_{d \in \mathcal{D}} \mathbb{P}_\sigma(\text{Hit}(d, v)), & \text{if } \sigma \text{ is multidefender} \\ = \sum_{v \in V} \sum_{d \in \mathcal{D}} \mathbb{P}_\sigma(\text{Hit}(d, v)), & \text{if } \sigma \text{ is undefender} \\ < 2\delta, & \text{if } \sigma \text{ is multidefender} \\ = 2\delta, & \text{if } \sigma \text{ is undefender} \end{cases},$$

30
31
32 as needed. ■

33 4.3.5 Minimum Hitting Probability, Maxhit Vertices and Maxhitters

34
35
36 Denote as

$$37 \quad \text{MinHit}_\sigma = \min_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v))$$

38
39
40 the *Minimum Hitting Probability* associated with the mixed profile σ . We observe:

41
42
43 **Lemma 4.3** For a mixed profile σ ,

$$44 \quad \text{MinHit}_\sigma \leq \frac{2\delta}{|V|}.$$

Proof. Assume, by way of contradiction, that $\text{MinHit}_\sigma > \frac{2\delta}{|V|}$. Then,

$$\begin{aligned} \sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) &\geq |V| \cdot \text{MinHit}_\sigma \\ &> 2\delta, \end{aligned}$$

a contradiction to Lemma 4.2. ■

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

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

Proof. Assume, by way of contradiction, that for each defender $d \in \mathcal{D}$, $\mathbb{P}_\sigma(\text{Hit}(d, v)) < 1$. Since the set $\{\text{Hit}(d, v) \mid d \in \mathcal{D}\}$ is a family of independent events with none of them being certain, this implies that the event $\text{Hit}(v) = \bigcup_{d \in \mathcal{D}} \text{Hit}(d, v)$ is not certain. So, $\mathbb{P}_\sigma(\text{Hit}(v)) < 1$. A contradiction. ■

4.4 Expected Utilities

The mixed profile σ induces an *Expected Utility* $U_b(\sigma)$ for each player $b \in \mathcal{A} \cup \mathcal{D}$, which is the expectation (according to σ) of the Utility of player b . We shall derive some formulas for Expected Utilities. To do so, we first define and derive formulas for some auxiliary quantities. In more detail, we define the Conditional Expected Proportion associated with the defenders; we then use it to derive an expression for the Conditional Expected Utility for each attacker. The Expected Utility of each attacker is then derived as a weighted sum of Conditional Expected Utilities. Similarly, the Expected Utility of each defender is derived as a weighted sum of Conditional Expected Utilities defined for the defenders in the natural way.

4.4.1 Conditional Expected Proportion

Induced by σ is the *Conditional Expected Proportion* $\text{Prop}_d(\sigma_{-d} \diamond v)$ of defender $d \in \mathcal{D}$ on vertex v , which is the expectation (induced by σ) of the proportion of defender d on vertex v had she chosen an edge incident to vertex v . Clearly,

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

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

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

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

4.4.2 Attackers

Induced by σ is the **Conditional Expected Utility** $U_a(\sigma_{-a} \diamond v)$ of attacker $a \in \mathcal{A}$ on vertex v , which is the conditional expectation (induced by σ) of the Utility of attacker a had she chosen vertex v . Clearly,

$$U_a(\sigma_{-a} \diamond v) = 1 - \mathbb{P}_\sigma(\text{Hit}(v)).$$

By the Law of Conditional Alternatives, we immediately obtain:

Lemma 4.6 Fix a mixed profile σ . Then, the Expected Utility $U_a(\sigma)$ of an attacker $a \in \mathcal{A}$ is

$$U_a(\sigma) = \sum_{v \in V} \sigma_a(v) \cdot (1 - \mathbb{P}_\sigma(\text{Hit}(v))).$$

We continue with a preliminary observation:

Lemma 4.7 Fix a mixed profile σ . Then, for each vertex $v \in V$,

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

Proof. By Lemma 4.5,

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

Note that 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 $d \in \mathcal{D}$ such that $d \in \mathcal{D}''$ and a set $\mathcal{D}' \subseteq \mathcal{D}''$ such that $\mathcal{D}' \subseteq \mathcal{D} \setminus \{d\}$ and $|\mathcal{D}'| = \ell - 1$. Hence,

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

It follows that

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

as needed. ■

4.4.3 Defenders

Induced by σ is also the *Conditional Expected Utility* $U_d(\sigma_{-d} \diamond (u, v))$ of defender d on edge $(u, v) \in E$, which is the conditional expectation (induced by σ) of the Utility of defender d had she chosen edge (u, v) . Clearly,

$$U_d((\sigma_{-d} \diamond (u, v))) = \text{Prop}_d(\sigma_{-d} \diamond u) \cdot |A|_\sigma(u) + \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v).$$

We prove:

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

$$U_d(\sigma) = \sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(d, v)) \cdot \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v).$$

Proof. By the Law of Conditional Alternatives,

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

as needed. ■

4.5 Nash Equilibria

A mixed profile σ is a *Nash equilibrium* [26, 27] if for each player $b \in \mathcal{A} \cup \mathcal{D}$, for each mixed strategy τ_b of player b , $U_b(\sigma) \geq U_b(\sigma_{-b} \diamond \tau_b)$; so, a Nash equilibrium is a local maximizer of the Expected Utility of each player. A (necessary and) sufficient condition for a Nash equilibrium σ is that for each player $b \in \mathcal{A} \cup \mathcal{D}$, for each pure strategy t_b of player b , $U_b(\sigma) \geq U_b(\sigma_{-b} \diamond t_b)$. By the celebrated Theorem of Nash [26, 27], $AD_{\alpha,\delta}(G)$ has at least one Nash equilibrium. Say that G *admits* a Nash equilibrium with a particular property if the game $AD_{\alpha,\delta}$ has a Nash equilibrium with this particular property.

Clearly, in a Nash equilibrium σ , for each attacker $a \in \mathcal{A}$, $U_a(\sigma_{-a} \diamond v)$ is *constant* over all vertices $v \in \text{Support}_\sigma(a)$; for each defender $d \in \mathcal{D}$, $U_d(\sigma_{-d} \diamond e)$ is *constant* over all edges $e \in \text{Support}_\sigma(d)$. It follows that in a Nash equilibrium σ , for each attacker $a \in \mathcal{A}$,

$$U_a(\sigma) = 1 - \mathbb{P}_\sigma(\text{Hit}(v)),$$

for any vertex $v \in \text{Support}_\sigma(a)$. So, for each attacker $a \in \mathcal{A}$, the quantity $\mathbb{P}_\sigma(\text{Hit}(v))$ is constant over all vertices $v \in \text{Support}_\sigma(a)$. In the same way, for each defender $d \in \mathcal{D}$,

$$U_d(\sigma) = \text{Prop}_d(\sigma_{-d} \diamond u) \cdot |A|_\sigma(u) + \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v),$$

for any edge $(u, v) \in \text{Support}_\sigma(d)$. So, for each defender $d \in \mathcal{D}$, the quantity $\text{Prop}_d(\sigma_{-d} \diamond u) \cdot |A|_\sigma(u) + \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v)$ is constant over all edges $(u, v) \in \text{Support}_\sigma(d)$. Note that in a Nash equilibrium σ , for each defender $d \in \mathcal{D}$, $U_d(\sigma) > 0$; in contrast, it is possible that $U_a(\sigma) = 0$ for some attacker $a \in \mathcal{A}$. (See, for an example, the proof of Theorem 9.2.)

4.6 Some Special Profiles

A profile σ is *uniform* if each player uses a *uniform* probability distribution on its support; so, for each attacker $a \in \mathcal{A}$, for each vertex $v \in \text{Support}_\sigma(a)$, $\sigma_a(v) = \frac{1}{|\text{Support}_\sigma(a)|}$, and for each defender $d \in \mathcal{D}$, for each edge $e \in \text{Support}_\sigma(d)$, $\sigma_d(e) = \frac{1}{|\text{Support}_\sigma(d)|}$.

A profile σ 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_k), for all vertices $v \in V$, (resp., all edges $e \in E$) $\sigma_{a_i}(v) = \sigma_{a_k}(v)$ (resp., $\sigma_{d_j}(e) = \sigma_{d_k}(e)$). A profile is *attacker-uniform* (resp., *defender-uniform*) if each attacker (resp., defender) uses a uniform probability distribution on his 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.

1
2
3 A profile σ is *attacker-fullymixed* (resp., *defender-fullymixed*) if for each attacker a
4 (resp., for each defender d), $\text{Support}_\sigma(a) = V$ (resp., $\text{Support}_\sigma(d) = E$). Now, **Attacker-**
5 **Fullymixed** (resp., **Defender-Fullymixed**) **Nash equilibria** are defined in the natural way.
6 A **Fullymixed Nash equilibrium** is both Attacker-Fullymixed and Defender-Fullymixed.
7
8

9 A profile σ is *defender-pure* if each defender chooses a single strategy with probability 1 in
10 σ . Now **Defender-Pure Nash equilibria** are defined in the natural way. Say that G admits
11 a **Defender-Pure Nash equilibrium**, or G is **Defender-Pure**, if there is a Defender-Pure
12 Nash equilibrium for the strategic game $\text{AD}_{\alpha,\delta}(G)$.
13
14

15 Fix now a Perfect-Matching graph. Say that a profile is *perfect-matching* if $\text{Supports}_\sigma(\mathcal{D})$
16 is a Perfect Matching. Now, **Perfect-Matching Nash equilibria** are defined in the natural
17 way.
18
19

20 21 22 4.7 Notation

23
24 Fix a mixed profile σ . For a vertex $v \in V$, set

$$25 \quad \text{Edges}_\sigma(v) = \{e \in \text{Supports}_\sigma(\mathcal{D}) \mid v \in e\};$$

26
27 so, $\text{Edges}_\sigma(v)$ consists of all edges incident to v that are included in the union of supports of
28 the defenders. For a vertex set $U \subseteq V$, set

$$29 \quad \text{Edges}_\sigma(U) = \bigcup_{v \in U} \text{Edges}_\sigma(v);$$

30
31 so, $\text{Edges}_\sigma(U)$ consists of all edges incident to a vertex in U that are included in the union of
32 supports of the defenders.
33
34

35 For an edge $e \in E$, set

$$36 \quad \text{Vertices}_\sigma(e) = \{v \in e \mid v \in \text{Supports}_\sigma(\mathcal{A})\};$$

37
38 so, $|\text{Vertices}_\sigma(e)| \leq 2$. For an edge set $F \subseteq E$, set

$$39 \quad \text{Vertices}_\sigma(F) = \bigcup_{e \in F} \text{Vertices}_\sigma(e);$$

40
41 so, $\text{Vertices}_\sigma(F)$ consists of all vertices incident to an edge in U that are included in the union
42 of supports of the attackers.
43
44
45
46
47
48
49
50
51
52

5 The Structure of Nash Equilibria

We provide an analysis of the combinatorial structure of the Nash equilibria associated with the strategic game $AD_{\alpha,\delta}(G)$. Section 5.1 presents a combinatorial characterization of Nash equilibria. Some necessary conditions for Nash equilibria are derived in Section 5.2. Section 5.3 treats the special case of Pure Nash equilibria.

5.1 Combinatorial Characterization

We show:

Proposition 5.1 (Characterization of Nash Equilibria) *A profile σ is a Nash equilibrium if and only if the following conditions hold:*

(1) For each vertex $v \in \text{Supports}_\sigma(\mathcal{A})$, $\mathbb{P}_\sigma(\text{Hit}(v)) = \text{MinHit}_\sigma$.

(2) For each defender $d \in \mathcal{D}$, for each edge $(u, v) \in \text{Support}_\sigma(d)$,

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

Proof. Assume first that σ is a Nash equilibrium. To establish Condition (1), consider any vertex $v \in \text{Supports}_\sigma(\mathcal{A})$; so, $v \in \text{Support}_\sigma(a)$ for some attacker a . Since σ is a Nash equilibrium, $\mathbb{P}_\sigma(\text{Hit}(v'))$ is constant over all vertices $v' \in \text{Support}_\sigma(a)$. We prove:

Lemma 5.2 *Fix any vertex $u \notin \text{Support}_\sigma(a)$. Then,*

$$\mathbb{P}_\sigma(\text{Hit}(u)) \geq \mathbb{P}_\sigma(\text{Hit}(v)).$$

Proof. Assume, by way of contradiction, that $\mathbb{P}_\sigma(\text{Hit}(u)) < \mathbb{P}_\sigma(\text{Hit}(v))$. Define $\tau = \sigma_{-a} \diamond \tau_a$, where τ_a is any mixed strategy of attacker a such that $u \in \text{Support}_\tau(a)$. So, by construction, $\mathbb{P}_\tau(\text{Hit}(u)) = \mathbb{P}_\sigma(\text{Hit}(u))$. Then,

$$\begin{aligned} & U_a(\sigma_{-a} \diamond \tau_a) \\ = & 1 - \mathbb{P}_\tau(\text{Hit}(u)) \quad (\text{since } u \in \text{Support}_\tau(a)) \\ = & 1 - \mathbb{P}_\sigma(\text{Hit}(u)) \\ > & 1 - \mathbb{P}_\sigma(\text{Hit}(v)) \quad (\text{by assumption}) \\ = & U_a(\sigma) \quad (\text{since } v \in \text{Support}_\sigma(a)), \end{aligned}$$

a contradiction. ■

1
2
3 We are now ready to prove Condition (1). Consider any vertex $u \notin \text{Support}_\sigma(\mathbf{a})$ such
4 that $u \in \text{Support}_\sigma(\mathbf{a}_k)$ for some attacker \mathbf{a}_k . (If such a vertex does not exist, then we are
5 done). By Lemma 5.2, $\mathbb{P}_\sigma(\text{Hit}(v)) \leq \mathbb{P}_\sigma(\text{Hit}(u))$. Assume, by way of contradiction, that
6 $\mathbb{P}_\sigma(\text{Hit}(v)) < \mathbb{P}_\sigma(\text{Hit}(u))$. Since σ is a local maximizer of the Expected Utility of attacker \mathbf{a}_k ,
7 and $U_{\mathbf{a}_k}(\sigma) = 1 - \mathbb{P}_\sigma(\text{Hit}(u))$. Thus, $u \notin \text{Support}_\sigma(\mathbf{a}_k)$. A contradiction.
8
9

10 For Condition (2), fix a defender \mathbf{d} and consider an edge $(u, v) \in \text{Support}_\sigma(\mathbf{d})$. Since σ is a
11 Nash equilibrium, the quantity $\text{Prop}_\sigma(\mathbf{d}, v') \cdot |A|_\sigma(v') + \text{Prop}_\sigma(\mathbf{d}, u') \cdot |A|_\sigma(u')$ is constant over
12 all edges $(u', v') \in \text{Support}_\sigma(\mathbf{d})$. So, consider any edge $(u', v') \notin \text{Support}_\sigma(\mathbf{d})$. Assume, by way
13 of contradiction, that
14

$$\begin{aligned} & \text{Prop}_\sigma(\mathbf{d}, u') \cdot |A|_\sigma(u') + \text{Prop}_\sigma(\mathbf{d}, v') \cdot |A|_\sigma(v') \\ & > \text{Prop}_\sigma(\mathbf{d}, u) \cdot |A|_\sigma(u) + \text{Prop}_\sigma(\mathbf{d}, v) \cdot |A|_\sigma(v). \end{aligned}$$

15 Denote $\tau = \sigma_{-\mathbf{d}} \diamond \tau_{\mathbf{d}}$, where $\tau_{\mathbf{d}}$ is any mixed strategy of defender \mathbf{d} such that $(u', v') \in$
16 $\text{Support}_\tau(\mathbf{d})$. So, by construction, $|A|_\tau(u') = |A|_\sigma(u')$ and $|A|_\tau(v') = |A|_\sigma(v')$. Then,
17
18

$$\begin{aligned} & U_{\mathbf{d}}(\sigma_{-\mathbf{d}} \diamond (u', v')) \\ & = \text{Prop}_\sigma(\mathbf{d}, u') \cdot |A|_\tau(u') + \text{Prop}_\sigma(\mathbf{d}, v') \cdot |A|_\tau(v') \quad (\text{since } e' \in \text{Support}_\tau(\mathbf{d})) \\ & = \text{Prop}_\sigma(\mathbf{d}, u') \cdot |A|_\sigma(u') + \text{Prop}_\sigma(\mathbf{d}, v') \cdot |A|_\sigma(v') \\ & > \text{Prop}_\sigma(\mathbf{d}, u) \cdot |A|_\sigma(u) + \text{Prop}_\sigma(\mathbf{d}, v) \cdot |A|_\sigma(v) \quad (\text{by assumption}) \\ & = U_{\mathbf{d}}(\sigma) \quad (\text{since } (u, v) \in \text{Support}_\sigma(\mathbf{d})), \end{aligned}$$

19 a contradiction.
20
21
22
23
24
25
26
27
28
29
30
31
32

33 Assume now that the mixed profile σ satisfies Conditions (1) and (2). We will prove that
34 σ is a Nash equilibrium.
35
36
37

- 38 • Consider first an attacker $\mathbf{a} \in \mathcal{A}$. Then, for any vertex $u \notin \text{Support}_\sigma(\mathbf{a})$,

$$\begin{aligned} & U_{\mathbf{a}}(\sigma) \\ & = 1 - \mathbb{P}_\sigma(\text{Hit}(v)) \quad (\text{where } v \in \text{Support}_\sigma(\mathbf{a})) \\ & \geq 1 - \mathbb{P}_\sigma(\text{Hit}(u)) \quad (\text{by Condition (1)}) \\ & = U_{\mathbf{a}}(\sigma_{-\mathbf{a}} \diamond u). \end{aligned}$$

- 39 • Consider now a defender $\mathbf{d} \in \mathcal{D}$. Then, for any edge $(u', v') \notin \text{Support}_\sigma(\mathbf{d})$,

$$\begin{aligned} & U_{\mathbf{d}}(\sigma) \\ & = \text{Prop}_\sigma(\mathbf{d}, u) \cdot |A|_\sigma(u) + \text{Prop}_\sigma(\mathbf{d}, v) \cdot |A|_\sigma(v) \quad (\text{where } (u, v) \in \text{Support}_\sigma(\mathbf{d})) \\ & \geq \text{Prop}_\sigma(\mathbf{d}, u') \cdot |A|_\sigma(u') + \text{Prop}_\sigma(\mathbf{d}, v') \cdot |A|_\sigma(v') \quad (\text{by Condition (2)}). \end{aligned}$$

It follows that σ is a Nash equilibrium. The proof is now complete. \blacksquare

We remark that Proposition 5.1 generalizes a corresponding characterization of Nash equilibria for $\text{AD}_{\alpha,1}(G)$ shown in [24, Theorem 3.1], where Condition (2) had the simpler counterpart (2'):

Proposition 5.3 *A profile σ is a Nash equilibrium if and only if the following conditions hold:*

- (1) For each vertex $v \in \text{Supports}_{\sigma}(\mathcal{A})$, $\mathbb{P}_{\sigma}(\text{Hit}(v)) = \text{MinHit}_{\sigma}$.
- (2') For each edge $e \in \text{Supports}_{\sigma}(\mathcal{D})$, $|A|_{\sigma}(e) = \max_{e' \in E} \{|A|_{\sigma}(e')\}$.

5.2 Necessary Conditions

We now establish necessary conditions for Nash equilibria, which will follow from their characterization (Proposition 5.1). We first prove a very simple expression for the *total* Expected Utility of the defenders:

Proposition 5.4 *In a Nash equilibrium σ ,*

$$\sum_{d \in \mathcal{D}} U_d(\sigma) = \alpha \cdot \text{MinHit}_{\sigma}.$$

Proof. Clearly,

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

as needed. \blacksquare

We continue to show:

Proposition 5.5 *In a Nash equilibrium σ , $\text{Supports}_{\sigma}(\mathcal{D})$ is an Edge Cover.*

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

Fix an attacker $\mathbf{a} \in \mathcal{A}$. Since σ is a local maximizer for the Expected Utility of \mathbf{a} , which is at most 1, it follows that $\sigma_{\mathbf{a}}(v) = 1$. Hence, for each $(u', v') \in \text{Supports}_\sigma(\mathcal{D})$, $|\mathbf{A}|_\sigma((u', v')) = 0$, since both $u' \neq u$ and $v' \neq v$ (by the choice of vertex v). So, $|\mathbf{A}|_\sigma(u') = |\mathbf{A}|_\sigma(v') = 0$. This implies that for any defender $\mathbf{d} \in \mathcal{D}$,

$$\begin{aligned} U_{\mathbf{d}}(\sigma) &= \sum_{(u,v) \in \text{Support}_\sigma(\mathbf{d})} \sigma_{\mathbf{d}}((u,v)) \cdot (\text{Prop}_{\mathbf{d}}(\sigma_{-\mathbf{d}} \diamond u) \cdot |\mathbf{A}|_\sigma(u) + \text{Prop}_{\mathbf{d}}(\sigma_{-\mathbf{d}} \diamond v) \cdot |\mathbf{A}|_\sigma(v)) \\ &= 0. \end{aligned}$$

Since σ is a Nash equilibrium, $U_{\mathbf{d}}(\sigma) > 0$. A contradiction. \blacksquare

Proposition 5.5 immediately implies:

Corollary 5.6 *A undefender Nash equilibrium is monodefender.*

We finally show:

Proposition 5.7 *In a Nash equilibrium σ , $\text{Supports}_\sigma(\mathcal{A})$ is a Vertex Cover of the graph $G(\text{Supports}_\sigma(\mathcal{D}))$.*

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

5.3 Pure Nash Equilibria

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

Proposition 5.8 (Characterization of Pure Nash Equilibria) *A pure profile \mathbf{s} is a Pure Nash equilibrium if and only if the following conditions hold:*

- (1) $\text{Supports}_\sigma(\mathcal{D})$ is an Edge Cover.
- (2) For each attacker $\mathbf{d} \in \mathcal{D}$, for each edge $(u, v) \in \text{Support}_\sigma(\mathbf{d})$,

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

We now use Propositions 5.5 and 5.7 to show:

Proposition 5.9 (Necessary Conditions for Pure Nash Equilibria) *Assume that G is Pure. Then, (i) $\delta \geq \beta'(G)$ and (ii) $\alpha \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$.*

Proof. By contradiction. Consider a Pure Nash equilibrium \mathbf{s} . For Condition (i), assume that $\delta < \beta'(G)$. Since $|\text{Supports}_{\mathbf{s}}(\mathcal{D})| \leq \delta$, it follows that $|\text{Supports}_{\mathbf{s}}(\mathcal{D})| < \beta'(G)$. Hence, $\text{Supports}_{\mathbf{s}}(\mathcal{D})$ is not an Edge Cover. A contradiction to Proposition 5.5.

For Condition (ii), assume that $\alpha < \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. Since $|\text{Supports}_{\mathbf{s}}(\mathcal{A})| \leq \alpha$, it follows that $|\text{Supports}_{\mathbf{s}}(\mathcal{A})| < \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. By Proposition 5.5, $\text{Supports}_{\mathbf{s}}(\mathcal{D})$ is an Edge Cover; so, $\beta(G(\text{Supports}_{\mathbf{s}}(\mathcal{D}))) \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$. It follows that $|\text{Supports}_{\mathbf{s}}(\mathcal{A})| < \beta(G(\text{Supports}_{\mathbf{s}}(\mathcal{D})))$. Thus, $\text{Supports}_{\mathbf{s}}(\mathcal{A})$ is not a Vertex Cover of the graph $G(\text{Supports}_{\mathbf{s}}(\mathcal{D}))$. A contradiction to Proposition 5.7. \blacksquare

We remark that Condition (i) (resp., Condition (ii)) in Proposition 5.9 is necessary for Defender-Pure (resp., Attacker-Pure) Nash equilibria. We finally provide a counterexample to the converse of Proposition 5.9:

Proposition 5.10 *There is a graph G and integers α and δ such that (i) $\delta \geq \beta'(G)$ and (ii) $\alpha \geq \min_{EC \in \mathcal{EC}(G)} \beta(G(EC))$ while G is not Pure.*

Proof. Consider the graph $G = (V, E)$ in Figure 7, and fix $\alpha = 2$ and $\delta = 6$. Clearly, $\beta'(G) = 6$ and $\min_{EC \in \mathcal{EC}(G)} \beta(G(EC)) = 2$. So, Conditions (i) and (ii) from the claim hold. Towards a contradiction, assume that G is Pure; consider a Pure Nash equilibrium \mathbf{s} .

- By Proposition 5.5, $\text{Support}_{\mathbf{s}}(\mathcal{D})$ is an Edge Cover. By the construction of G , this implies that $\text{Support}_{\mathbf{s}}(\mathcal{D}) = \{(v_2, v_3), (v_4, v_5), (v_4, v_6), (v_4, v_7), (v_4, v_8), (v_1, v)\}$, where $v \in \{v_2, v_4\}$. Since $\delta = 6$, it follows that for each edge $e \in \text{Support}_{\mathbf{s}}(\mathcal{D})$, there is a unique defender d such that $s_d = e$.
- By Proposition 5.7, $\text{Support}_{\mathbf{s}}(\mathcal{A})$ is a Vertex Cover of the graph $G(\text{Support}_{\mathbf{s}}(\mathcal{D}))$. By the construction of G , this implies that $\text{Support}_{\mathbf{s}}(\mathcal{A}) = \{v_2, v_4\}$. (Note that $\{v_2, v_4\}$ is the *unique* Vertex Cover of the graph $G(\text{Support}_{\mathbf{s}}(\mathcal{D}))$ with size at most 2.) Since $\alpha = 2$, it follows that $\text{Support}_{\mathbf{s}}(a_1) = v_2$ and $\text{Support}_{\mathbf{s}}(a_2) = v_4$.

Consider now the (unique) defender $d \in \mathcal{D}$ such that $s_d = (v_4, v_5)$. Clearly, $U_d(\mathbf{s}) = \frac{1}{4}$, but

$$\begin{aligned} U_d(\mathbf{s}_{-d} \diamond (v_2, v_3)) &= \begin{cases} \frac{1}{3}, & \text{if } v = v_2 \\ \frac{1}{2}, & \text{if } v = v_4 \end{cases} \\ &\geq \frac{1}{3}. \end{aligned}$$

So, $U_d(\mathbf{s}_{-d} \diamond (v_2, v_3)) > U_d(\mathbf{s})$. A contradiction to the fact that \mathbf{s} is a Nash equilibrium. \blacksquare




Figure 7: The graph G used in the proof of Proposition 6.6. Edges in $\text{Support}_s(\mathcal{D})$ are drawn thick; vertices in $\text{Support}_s(\mathcal{A})$ are squared.

6 Defense-Optimal Nash Equilibria

Section 6.1 introduces Defense-Optimal Nash equilibria and Defense-Optimal graphs. Some sufficient conditions for Defense-Optimal graphs are presented in Section 6.2.

6.1 Definitions

The *Defense-Ratio* DR_σ of a Nash equilibrium σ is the ratio of the *optimal* total Utility α of the defenders over their total Expected Utility in σ ; so,

$$\text{DR}_\sigma = \frac{\alpha}{\sum_{d \in \mathcal{D}} U_d(\sigma)}.$$

By the definition of Defense-Ratio, Proposition 5.4 immediately implies:

Corollary 6.1 *For a Nash equilibrium σ ,*

$$\text{DR}_\sigma = \frac{1}{\text{MinHit}_\sigma}.$$

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

Corollary 6.2 *For a Nash equilibrium σ ,*

$$\text{DR}_\sigma \geq \frac{|V|}{2\delta}.$$

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

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

The justification for the definition of a Defense-Optimal Nash equilibrium will come later, when we construct Defense-Optimal Nash equilibria in two particular 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 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, \dots, E_δ . Recall that $E(f)$ is an Edge Cover. Construct σ as follows:

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

Clearly, σ is attacker-symmetric, attacker-uniform, attacker-fullymixed and defender-symmetric; moreover, σ is monodefender. Furthermore, for each vertex $v \in V$, $\text{Edges}_\sigma(v) = \{e \in E(f) \mid v \in e\}$. To prove that σ is a (mixed) profile, we prove that for each defender $\mathbf{d}_j \in \mathcal{D}$, $\sigma_{\mathbf{d}_j}$ is a probability distribution (on E). Clearly,

$$\begin{aligned}
 & \sum_{e \in E} \sigma_{\mathbf{d}_j}(e) \\
 &= \sum_{e \in E_j} \sigma_{\mathbf{d}_j}(e) \quad (\text{since } \text{Support}_\sigma(\mathbf{d}_j) = E_j) \\
 &= \sum_{e \in E_j} \frac{2\delta}{|V|} \cdot f(e) \quad (\text{by construction}) \\
 &= \frac{2\delta}{|V|} \sum_{e \in E_j} f(e) \\
 &= 1 \quad (\text{since } f \text{ is } \delta\text{-Partitionable});
 \end{aligned}$$

so, σ_{d_j} is a probability distribution, which establishes that σ is a profile.

We continue to prove that σ is a Nash equilibrium. We shall verify Conditions (1) and (2) in the characterization of Nash equilibria (Proposition 5.1).

For Condition (1), fix a vertex $v \in V$. Since $E(f)$ is an Edge Cover, there is a partite $E_j \subseteq E(f)$ such that $v \in \text{Vertices}(E_j)$. Since the partites E_1, \dots, E_δ are vertex-disjoint and $\text{Support}_\sigma(d_j) = E_j$, it follows that vertex v is monodefender in σ with $d_\sigma(v) = d_j$. We prove:

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

Proof. Clearly,

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

as needed. ■

By Claim 6.4, Condition (1) holds trivially.

For Condition (2), consider a defender $d \in \mathcal{D}$. Fix an edge $(u, v) \in \text{Support}_\sigma(d)$. Since σ is monodefender, $\text{Prop}_d(\sigma_{-d} \diamond u) = \text{Prop}_d(\sigma_{-d} \diamond v) = 1$. Hence,

$$\begin{aligned}
\text{Prop}_d(\sigma_{-d} \diamond u) \cdot |A|_\sigma(u) + \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v) &= |A|_\sigma(u) + |A|_\sigma(v) \\
&= \frac{2\alpha}{|V|}.
\end{aligned}$$

Fix now an edge $(u', v') \notin \text{Support}_\sigma(d)$. 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 σ , this implies that there are defenders $d_{u'}$ and $d_{v'}$ such that $e_{u'} \in \text{Support}_\sigma(d_{u'})$ and $e_{v'} \in \text{Support}_\sigma(d_{v'})$.

There are two cases for $d_{u'}$ (resp., $d_{v'}$): either $d_{u'} = d$ or $d_{u'} \neq d$ (resp., $d_{v'} = d$ or $d_{v'} \neq d$). We shall treat each of them separately.

- Assume first that $d_{u'} = d$ (resp., $d_{v'} = d$); since u' is monodefender, it follows that $\text{Prop}_d(\sigma_{-d} \diamond u') = 1$ (resp., $\text{Prop}_d(\sigma_{-d} \diamond v') = 1$).
- Assume now that $d_{u'} \neq d$ (resp., $d_{v'} \neq d$); since v' is monodefender, $\text{Prop}_d(\sigma_{-d} \diamond u') < 1$ (resp., $\text{Prop}_d(\sigma_{-d} \diamond v') < 1$).

So, in all cases, $\text{Prop}_d(\sigma_{-d} \diamond u') \leq 1$ and $\text{Prop}_d(\sigma_{-d} \diamond v') \leq 1$. Thus,

$$\begin{aligned} \text{Prop}_d(\sigma_{-d} \diamond u') \cdot |A|_{\sigma}(u') + \text{Prop}_d(\sigma_{-d} \diamond v') \cdot |A|_{\sigma}(v') &\leq |A|_{\sigma}(u') + |A|_{\sigma}(v') \\ &= \frac{2\alpha}{|V|}. \end{aligned}$$

Now, Condition (2) follows.

Hence, by Proposition 5.1, σ is a Nash equilibrium. By Claim 6.4 and Condition (1) of Proposition 5.1, it follows that $\text{MinHit}_{\sigma} = \frac{2\delta}{|V|}$. By Corollary 6.1, it follows that $\text{DR}_{\sigma} = \frac{|V|}{2\delta}$. Since G has a δ -Partitionable Fractional Perfect Matching, Corollary 2.14 implies that $\delta \leq \frac{|V|}{2}$, so that $\max\left\{1, \frac{|V|}{2\delta}\right\} = \frac{|V|}{2\delta}$. This implies that $\text{DR}_{\sigma} = \max\left\{1, \frac{|V|}{2\delta}\right\}$. Hence, σ is Defense-Optimal, as needed. \blacksquare

We continue with another sufficient condition:

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

Proof. Fix an arbitrary Defender-Pure Nash equilibrium σ ; so $\text{DR}_{\sigma} = 1$. For each defender $d \in \mathcal{D}$, denote $s_d = (u_d, v_d) \in E$. Since σ is a Nash equilibrium, $U_d(\sigma) = U_d(\sigma_{-d} \diamond (u_d, v_d))$. So,

$$\begin{aligned} \text{DR}_{\sigma} &= \\ &= \frac{\alpha}{\sum_{d \in \mathcal{D}} U_d(\sigma)} \\ &= \frac{\alpha}{\sum_{d \in \mathcal{D}} U_d(\sigma_{-d} \diamond (u_d, v_d))} \\ &= \frac{\alpha}{\sum_{d \in \mathcal{D}} (\text{Prop}_d(\sigma_{-d} \diamond u_d) \cdot |A|_{\sigma}(u_d) + \text{Prop}_d(\sigma_{-d} \diamond v_d) \cdot |A|_{\sigma}(v_d))} \\ &= \frac{\alpha}{\sum_{d \in \mathcal{D}} \left(\frac{|A|_{\sigma}(u_d)}{|D_{\sigma}(u_d)|} + \frac{|A|_{\sigma}(v_d)}{|D_{\sigma}(v_d)|} \right)} \quad (\text{since } \sigma \text{ is Defender-Pure}) \\ &= \frac{\alpha}{\sum_{v \in V} \sum_{d \in \mathcal{D} | v \in \text{Vertices}_{\sigma}(\text{Support}_{\sigma}(d))} \frac{|A|_{\sigma}(v)}{|D_{\sigma}(v)|}} \\ &= \frac{\alpha}{\sum_{v \in \text{Supports}_{\sigma}(\mathcal{A})} |D_{\sigma}(v)| \cdot \frac{|A|_{\sigma}(v)}{|D_{\sigma}(v)|}} \\ &= \frac{\alpha}{\sum_{v \in \text{Supports}_{\sigma}(\mathcal{A})} |A|_{\sigma}(v)} \\ &= 1. \end{aligned}$$




Figure 8: The graph G used in the proof of Proposition 6.6. Edges in $E(f)$ for the 2-Partitionable Fractional Perfect Matching f are drawn thick.

By Corollary 6.1, it follows that $\text{MinHit}_\sigma = 1$. Hence, Lemma 4.3 implies that $\delta \geq \frac{|V|}{2}$. Since $\text{DR}_\sigma = 1$, it follows that $\text{DR}_\sigma = \max \left\{ 1, \frac{|V|}{2\delta} \right\}$, and Condition (i) follows. ■

We finally 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 the graph $G = (V, E)$ in Figure 8, and fix $\delta = 2$. Consider the function $f : E \rightarrow [0, 1]$ with $f(e) = \frac{1}{2}$ for each edge $e \in E \setminus \{(v_3, v_4)\}$ and $f(e) = 0$ for $e = (v_3, v_4)$. Clearly, f is a 2-Partitionable Fractional Perfect Matching with $E_1 = \{(v_1, v_2), (v_2, v_3), (v_1, v_3)\}$ and $E_2 = \{(v_4, v_5), (v_5, v_6), (v_4, v_6)\}$. Since $\delta = 2$ and $\beta'(G) = 3$, it follows by Proposition 5.9 (Condition (i)) that G is not Defender-Pure. ■

7 Few Defenders

We consider the case of *few defenders* where $\delta \leq \frac{|V|}{2}$; there, a Defense-Optimal Nash equilibrium σ has Defense-Ratio $\text{DR}_\sigma = \max \left\{ 1, \frac{|V|}{2\delta} \right\} = \frac{|V|}{2\delta}$, so that by Corollary 6.1, $\text{MinHit}_\sigma = \frac{2\delta}{|V|}$. This implies that $\sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) \geq 2\delta$. By Lemma 4.2, it follows that $\sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) = 2\delta$, so that σ is undefender. By Corollary 5.6, σ is monodefender.

Section 7.1 provides some necessary conditions for Defense-Optimal Nash equilibria and Defense-Optimal graphs. A combinatorial characterization of Defense-Optimal graphs is presented in Section 7.2, with an implication on the associated complexity. Section 7.3 considers the special case of Perfect-Matching graphs.

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 σ . Recall that σ is monodefender. Since $\text{MinHit}_\sigma(v) = \frac{2\delta}{|V|}$ and $\sum_{v \in V} \mathbb{P}_\sigma(\text{Hit}(v)) = 2\delta$, it follows that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = \frac{2\delta}{|V|}$.

We now define a function $f : E \rightarrow \mathbb{R}$; we will then prove that f is a δ -Partitionable Fractional Perfect Matching. For each edge $e \in E$, set

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

By construction, $E(f) = \text{Supports}_\sigma(\mathcal{D})$; so, for each vertex $v \in V$, $\{e \in E(f) \mid v \in e\} = \text{Edges}_\sigma(v)$. Since σ is monodefender, it follows that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = \mathbb{P}_\sigma(\text{Hit}(\mathbf{d}_\sigma(v), v))$. We prove:

Claim 7.2 *For each vertex $v \in V$,*

$$\sum_{e \in \text{Edges}_\sigma(v)} f(e) = 1.$$

Proof. By the construction of f ,

$$\begin{aligned} \sum_{e \in \text{Edges}_\sigma(v)} f(e) &= \sum_{e \in \text{Support}_\sigma(\mathbf{d}_\sigma(v)) \mid v \in e} f(e) \\ &= \sum_{e \in \text{Support}_\sigma(\mathbf{d}_\sigma(v)) \mid v \in e} \frac{|V|}{2\delta} \cdot \sigma_{\mathbf{d}_\sigma(v)}(e) \\ &= \frac{|V|}{2\delta} \sum_{e \in \text{Support}_\sigma(\mathbf{d}_\sigma(v)) \mid v \in e} \sigma_{\mathbf{d}_\sigma(v)}(e) \\ &= \frac{|V|}{2\delta} \cdot \mathbb{P}_\sigma(\text{Hit}(\mathbf{d}_\sigma(v), v)) \\ &= \frac{|V|}{2\delta} \cdot \mathbb{P}_\sigma(\text{Hit}(v)) \\ &= \frac{|V|}{2\delta} \cdot \frac{2\delta}{|V|} \\ &= 1, \end{aligned}$$

as needed. ■

Since $\text{Edges}_\sigma(v) = \{e \in E(f) \mid v \in e\}$, 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, \dots, E_δ where for each $j \in [\delta]$, $E_j := \text{Support}_\sigma(d_j)$. Clearly,

$$\begin{aligned} \bigcup_{j \in [\delta]} E_j &= \bigcup_{j \in [\delta]} \text{Support}_\sigma(d_j) \\ &= \text{Supports}_\sigma(\mathcal{D}) \\ &= E(f). \end{aligned}$$

Since σ is monodefender, it follows that for all pairs of distinct defenders d_k and d_l , $\text{Support}_\sigma(d_k) \cap \text{Support}_\sigma(d_l) = \emptyset$. Hence, it follows that the sets E_1, \dots, E_δ , partition the set $E(f)$; so, we shall call them partites. 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 construction of f and the partites E_1, \dots, E_δ ,

$$\begin{aligned} \sum_{e \in E_j} f(e) &= \sum_{e \in \text{Support}_\sigma(d_j)} f(e) \\ &= \sum_{e \in \text{Support}_\sigma(d_j)} \frac{|V|}{2\delta} \cdot \sigma_{d_j}(e) \\ &= \frac{|V|}{2\delta} \sum_{e \in \text{Support}_\sigma(d_j)} \sigma_{d_j}(e) \\ &= \frac{|V|}{2\delta}, \end{aligned}$$

as needed. ■

Claim 7.3 implies that f is δ -Partitionable, and the claim follows. ■

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 (for $\delta \leq \frac{|V|}{2}$); 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.

We observe three implications of Theorem 7.4. The first one is an immediate consequence of Theorem 7.4 and Corollary 2.14.

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 2.18.

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 identifies a particular value of δ (namely, $\delta = \frac{|V|}{2}$) for which the recognition problem for Defense-Optimal graphs is tractable. For the third implication, Theorem 7.4 implies that the complexity of recognizing Defense-Optimal graphs is that of δ -PARTITIONABLE FPM. Hence, Proposition 2.19 immediately implies:

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 2δ 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 2δ divides $|V|$. Then, G admits a Defense-Optimal, Perfect-Matching Nash equilibrium.

Proof. Consider a Perfect Matching M . Construct a profile σ as follows:

- For each attacker $\mathbf{a} \in \mathcal{A}$ and for each vertex $v \in V$, set

$$\sigma_{\mathbf{a}}(v) := \frac{1}{|V|}.$$

So, σ is attacker-symmetric, attacker-uniform and attacker-fullymixed. Clearly, for each vertex $v \in V$, $|\mathbf{A}|_{\sigma}(v) = \frac{\alpha}{|V|}$.

- 1
2
3
4 • Partition M into δ sets M_1, \dots, M_δ , each with $\frac{|V|}{2\delta}$ edges; each defender d_j with $j \in [\delta]$
5 uses a uniform probability distribution over the set M_j . So, for each edge $e \in M_j$, set
6

$$7 \quad \sigma_{d_\sigma(e)}(e) := \frac{2\delta}{|V|}.$$

8
9
10 Thus, $\text{Support}_\sigma(d_j) = M_j$ for each $j \in [\delta]$, so that $\text{Supports}_\sigma(\mathcal{D}) = M$. Clearly, each
11 vertex $v \in V$ is monodefender in σ with $\mathbb{P}_\sigma(\text{Hit}(v)) = \mathbb{P}_\sigma(\text{Hit}(d_\sigma(v), v)) = \frac{2\delta}{|V|}$.
12
13

14 We shall verify Conditions (1) and (2) in the characterization of Nash equilibria (Proposi-
15 tion 5.1). For Condition (1), fix a vertex $v \in V$. Since $\mathbb{P}_\sigma(\text{Hit}(v)) = \frac{2\delta}{|V|}$, Condition (1) follows
16 trivially. For Condition (2), consider any defender $d \in \mathcal{D}$.
17
18

- 19
20 • Fix an edge $(u, v) \in \text{Support}_\sigma(d)$. Since each edge is monodefender in σ , it follows that
21 $\text{Prop}_d(\sigma_{-d} \diamond u) = \text{Prop}_d(\sigma_{-d} \diamond v) = 1$. Hence,
22

$$23 \quad \text{Prop}_d(\sigma_{-d} \diamond u) \cdot |A|_\sigma(u) + \text{Prop}_d(\sigma_{-d} \diamond v) \cdot |A|_\sigma(v) = \frac{2\alpha}{|V|}.$$

- 24
25
26
27 • Fix now an edge $(u', v') \notin \text{Support}_\sigma(d)$. Since M is an Edge Cover, there are edges
28 $e_{u'}, e_{v'} \in M$ such that $u' \in e_{u'}$ and $v' \in e_{v'}$. By the construction of σ , this implies that
29 there are defenders $d_{u'}$ and $d_{v'}$ such that $e_{u'} \in \text{Support}_\sigma(d_{u'})$ and $e_{v'} \in \text{Support}_\sigma(d_{v'})$.
30 Since each vertex is monodefender in σ , it follows that $d \neq d_{u'}$ and $d \neq d_{v'}$. Hence,
31 $\text{Prop}_d(\sigma_{-d} \diamond u') \leq \frac{1}{2}$ and $\text{Prop}_d(\sigma_{-d} \diamond v') \leq \frac{1}{2}$, so that
32
33

$$34 \quad \begin{aligned} 35 \quad \text{Prop}_d(\sigma_{-d} \diamond u') \cdot |A|_\sigma(u') + \text{Prop}_d(\sigma_{-d} \diamond v') \cdot |A|_\sigma(v') &\leq \frac{1}{2} \cdot (|A|_\sigma(u') + |A|_\sigma(v')) \\ 36 &= \frac{\alpha}{|V|}. \end{aligned}$$

37
38
39
40 Now, Condition (2) follows. Hence, by Proposition 5.1, σ is a Nash equilibrium.
41

42 To prove that σ is Defense-Optimal, recall that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = \frac{2\delta}{|V|}$.
43 Hence, $\text{MinHit}_\sigma = \frac{2\delta}{|V|}$. By Corollary 6.1, it follows that $\text{DR}_\sigma = \frac{|V|}{2\delta}$. Since $\delta \leq \frac{|V|}{2}$, it follows
44 that $\text{DR}_\sigma = \max \left\{ 1, \frac{|V|}{2\delta} \right\}$. Hence, σ is Defense-Optimal. ■
45
46
47
48
49

50 We continue to prove:

51
52 **Proposition 7.10** *Assume that $\delta \leq \frac{|V|}{2}$ for a Perfect-Matching graph G , which admits a*
53 *Defense-Optimal, Perfect-Matching Nash equilibrium. Then, 2δ divides $|V|$.*
54
55
56
57
58

Proof. Consider such a Nash equilibrium σ , and recall that $\text{MinHit}_\sigma = \frac{2\delta}{|V|}$. Consider an edge $(u, v) \in \text{Supports}_\sigma(\mathcal{D})$; so, $e \in \text{Support}_\sigma(\mathbf{d})$ for some defender $\mathbf{d} \in \mathcal{D}$. Proposition 5.7 implies that $\text{Supports}_\sigma(\mathcal{A})$ is a Vertex Cover of the graph $G(\text{Supports}_\sigma(\mathcal{D}))$. Hence, either $u \in \text{Supports}_\sigma(\mathcal{A})$ or $v \in \text{Supports}_\sigma(\mathcal{A})$ (or both). Assume without loss of generality, that $u \in \text{Supports}_\sigma(\mathcal{A})$. Since σ is monodefender, there is a single defender \mathbf{d}_k such that $u \in \text{Vertices}(\text{Support}_\sigma(\mathbf{d}))$. Hence, \mathbf{d}_k is identified with \mathbf{d} . Since σ is Perfect-Matching, $\text{Support}_\sigma(\mathbf{d})$ is a Perfect Matching; this implies that $\mathbb{P}_\sigma(\text{Hit}(v)) = s_{\mathbf{d}}(e)$. We prove:

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

Proof. Clearly,

$$\begin{aligned} \sum_{e \in \text{Support}_\sigma(\mathbf{d})} \sigma_{\mathbf{d}}(e) &= \sum_{e \in \text{Support}_\sigma(\mathbf{d})} \mathbb{P}_\sigma(\text{Hit}(v)) \\ &= |\text{Support}_\sigma(\mathbf{d})| \cdot \frac{2\delta}{|V|} \end{aligned}$$

Since σ is a profile, $\sum_{e \in \text{Support}_\sigma(\mathbf{d})} \sigma_{\mathbf{d}}(e) = 1$. Hence, $|\text{Support}_\sigma(\mathbf{d})| = \frac{|V|}{2\delta}$, as needed. \blacksquare

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

The claim follows now from Propositions 7.9 and 7.10. \blacksquare

Note that while Corollary 7.5 applies to *all* graphs, Proposition 7.10 applies only to Perfect-Matching graphs. However, the restriction of Corollary 7.5 to Perfect-Matching graphs does *not* imply Proposition 7.10 *unless* δ is odd. (This is because 2 divides $|V|$ and δ divides $|V|$ imply together that 2δ divides $|V|$ exactly when δ is odd.) Hence, Proposition 7.10 strictly strengthens Corollary 7.5 for the case where δ is even.

8 Many Defenders

We now consider the case of *many defenders*, where $\frac{|V|}{2} < \delta < \beta'(G)$. In this case, a Defense-Optimal Nash equilibrium σ has Defense-Ratio $\text{DR}_\sigma = \max\left\{1, \frac{|V|}{2\delta}\right\} = 1$. By Corollary 6.1, this implies that $\text{MinHit}_\sigma = 1$. It follows that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = 1$, so that the number of maxhit vertices in σ 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 σ . Consider any (maxhit) vertex $v \in V$. By Lemma 4.4, there is a maxhitter $d \in \mathcal{D}$ in σ with $\mathbb{P}_\sigma(\text{Hit}(d, v)) = 1$. Use σ to construct a defender-pure profile τ as follows:

- Fix a defender $d \in \mathcal{D}$. If d is maxhitter in σ , then τ_d is any edge $(u, v) \in \text{Support}_\sigma(d)$ such that d is maxhitter in σ for the vertex $v \in V$; else, τ_d is any arbitrary edge $(u, v) \in \text{Support}_\sigma(d)$.

By construction of τ , the following conditions hold:

- (1) $|\text{Supports}_\tau(\mathcal{D})| \leq \delta$.
- (2) Each maxhit vertex in σ remains a maxhit vertex in τ ; so, the number of maxhit vertices in τ is $|V|$.

Since $\delta < \beta'(G)$, Condition (1) implies that $|\text{Supports}_\tau(\mathcal{D})| < \beta'(G)$. Hence, $\text{Supports}_\tau(\mathcal{D})$ is not an Edge Cover. So, there is some vertex $v \in V$ with $\mathbb{P}_\tau(\text{Hit}(v)) = 0$. It follows that the number of maxhit vertices in τ is less than $|V|$. A contradiction. ■

9 Too Many Defenders

We finally turn to the case of *too many defenders* where $\delta \geq \beta'(G)$. In this case, $\frac{|V|}{2\delta} \leq \frac{|V|}{2\beta'(G)} \leq 1$; so, a Defense-Optimal Nash equilibrium σ has Defense-Ratio $\text{DR}_\sigma = 1$. By Corollary 6.1, this implies that $\text{MinHit}_\sigma = 1$; so, that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = 1$.

Section 9.1 introduces vertex-balanced profiles. These profiles give rise to Defender-Pure, Vertex-Balanced Nash equilibria and Pure, Vertex-Balanced Nash equilibria, which will be treated in Sections 9.2 and 9.3, respectively.

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

We start with a significant definition:

Definition 9.1 *A mixed profile σ is vertex-balanced if there is a constant $c > 0$ such that for each vertex $v \in V$,*

$$\frac{|A|_\sigma(v)}{|D|_\sigma(v)} = c.$$

The following properties follow trivially for a vertex-balanced profile σ :

- 1 The set $\text{Supports}_\sigma(\mathcal{D})$ is an Edge Cover. This matches the necessary condition for an arbitrary Nash equilibrium from Proposition 5.5.
- 2 The set $\text{Supports}_\sigma(\mathcal{A})$ is V . Note that this property is strictly weaker than the condition defining an attacker-fullymixed profile σ , which requires that for each attacker $a \in \mathcal{A}$, $\text{Support}_\sigma(a) = V$.

We shall consider defender-pure, vertex-balanced profiles and pure, vertex-balanced profiles. We prove a nice 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 σ and a defender $d \in \mathcal{D}$ with $\sigma_d = (u, v)$. Clearly,

$$\begin{aligned} U_d(\sigma) &= \frac{|A|_\sigma(u)}{|D_\sigma(u)|} + \frac{|A|_\sigma(v)}{|D_\sigma(v)|} \\ &= 2c. \end{aligned}$$

Fix now an edge $(u', v') \notin \text{Support}_\sigma(d)$. Clearly,

$$\begin{aligned} U_d(\sigma_{-d} \diamond (u', v')) &= \frac{|A|_\sigma(u')}{|D_\sigma(u')| + 1} + \frac{|A|_\sigma(v')}{|D_\sigma(v')| + 1} \\ &< \frac{|A|_\sigma(u')}{|D_\sigma(u')|} + \frac{|A|_\sigma(v')}{|D_\sigma(v')|} \\ &= 2c, \end{aligned}$$

and the claim follows. ■

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

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

We show:

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 σ .

- (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 $\{A(v) \mid v \in V\}$ to the following linear system:

- (a) For each vertex $v \in V$, $\frac{A(v)}{|D_\sigma(v)|}$ is constant.
 - (b) $\sum_{v \in V} A(v) = \alpha$.
- (4) Assign a mixed strategy σ to each attacker in an arbitrary way so that for each vertex $v \in V$, $|A|_\sigma(v) = A(v)$.

Figure 9: The algorithm DefenderPure&VertexBalancedNE. By Step (2), σ is defender-pure; note that the assignment exhausts all edges from the Minimum Edge Cover EC due to the assumption that $\delta \geq \beta'(G)$. Step (3) provisions for σ to be vertex-balanced; towards this end, it provides for the ratio $\frac{A(v)}{|D_\sigma(v)|}$ to be constant over all vertices $v \in V$. Finally, Step (4) provides mixed strategies to the attackers that induce $|A|_\sigma(v) = A(v)$ for each vertex $v \in V$; by Step (3/a) this implies that σ is vertex-balanced. Since a Minimum Edge Cover is computable in polynomial time, the algorithm DefenderPure&VertexBalancedNE is polynomial time.

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 9.

Proof. By construction (Steps (1) and (2)) and the assumption that $\delta \geq \beta'(G)$, it follows that $\text{Supports}_\sigma(\mathcal{D})$ is a Minimum Edge Cover. Since σ is defender-pure, this implies that for each vertex $v \in V$, $\mathbb{P}_\sigma(\text{Hit}(v)) = 1$; hence, for each attacker $a \in \mathcal{A}$,

$$\begin{aligned} U_a(\sigma_{-a} \diamond v) &= 1 - \mathbb{P}_\sigma(\text{Hit}(v)) \\ &= 0. \end{aligned}$$

This implies that σ is (vacuously) a local maximizer for the Expected Utility of each attacker. By Proposition 9.1, it follows that σ is a Nash equilibrium. \blacksquare

By Theorem 6.5 and Theorem 9.2, it immediately follows:

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

By Theorem 5.9 (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 the class of Defender-Pure graphs is recognizable in polynomial time (for an arbitrary value of δ).

9.3 Defense-Optimal, Pure, Vertex-Balanced Nash Equilibria

We show:

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

To prove the claim, we present the algorithm Pure&VertexBalanced in Figure 10. The proof of Theorem 9.5 is identical to the proof of Theorem 9.2.

Algorithm Pure&VertexBalancedNE

INPUT: A graph $G = \langle V, E \rangle$ such that $\delta \geq \beta'(G)$ with 2δ divides α .

OUTPUT: A Defense-Optimal, Pure, Vertex-Balanced Nash equilibrium \mathbf{s} .

- (1) Choose a Minimum Edge Cover EC .
- (2) Assign each defender to a distinct edge from EC in a round-robin fashion.
- (3) For each vertex $v \in V$, set $A(v) := |D_{\sigma}(v)| \cdot \frac{\alpha}{2\delta}$.
- (4) Assign each attacker to a vertex from V in an arbitrary way so that for each vertex $v \in V$, $|A|_{\sigma}(v) = A(v)$.

Figure 10: The algorithm Pure&VertexBalancedNE. The algorithm Pure&VertexBalancedNE differs from the the algorithm DefenderPure&VertexBalancedNE only in Steps (3) and (4). The additional assumption that 2δ divides α suffices for Steps (3) and (4) to construct an attacker-pure profile. By Step (2), σ is defender-pure. Step (3) provisions for σ to be vertex-balanced; towards this end, it sets the ratio $\frac{A(v)}{|D_{\sigma}(v)|}$ to the fixed (integer) value $\frac{\alpha}{2\delta}$ for each vertex $v \in V$. Hence, for each vertex $v \in V$, $A(v)$ is integer. Finally, Step (4) assigns pure strategies to the attackers that induce the (integer) value $|A|_{\sigma}(v) = A(v)$ for each vertex $v \in V$; hence, σ is vertex-balanced by construction. Since a Minimum Edge Cover is computable in polynomial time, the algorithm Pure&VertexBalancedNE is polynomial time.

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.

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

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

References

- [1] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, Wiley, Second Edition, 2001.
- [2] R. Anderson, “Why Information Security is Hard—An Economic Perspective,” *Proceedings of the 17th Annual Computer Security Applications Conference*, pp. 358–365, December 2001.
- [3] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler and T. Roughgarden, “The Price of Stability for Network Design with Fair Cost Allocation,” *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 295–304, October 2004.
- [4] 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.
- [5] 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.
- [6] G. Coulouris, J. Dollimore and T. Kindberg, *Distributed Systems—Concepts and Design*, Addison-Wesley, Fourth Edition, 2005.
- [7] J. Edmonds, “Paths, Trees and Flowers,” *Canadian Journal of Mathematics*, Vol. 17, pp. 449–467, 1965.
- [8] J. Edmonds, “Maximum Matching and a Polyhedron with 0,1-Vertices,” *Journal of Research at the National Bureau of Standards*, Section B69, pp. 125–130, 1965.
- [9] M. R. Garey and D. S. Johnson, *Computers and Intractability—A Guide to the Theory of NP-Completeness*, W. H. Freeman and Co., 1979.
- [10] 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.
- [11] 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.

- 1
2
3 [12] J. Grossklags, N. Christin and J. Chuang, “Predicted and Observed User Behavior in
4 the Weakest-Link Security Game,” *Proceedings of the USENIX Workshop on Usability,*
5 *Psychology, and Security*, pp. 1–6, April 2008.
6
7
8 [13] J. Grossklags, N. Christin and J. Chuang, “Secure or Insure? A Game-Theoretic Analysis
9 of Information Security Games,” *Proceedings of the 17th International World Wide Web*
10 *Conference*, pp. 209–218, April 2008.
11
12
13 [14] D. König, “Über Graphen und Ihre Anwendung auf Determinantentheorie und Mengen-
14 lehre,” *Mathematische Annalen*, Vol. 77, pp. 453–465, December 1916.
15
16 [15] B. Korte and J. Vygen, *Combinatorial Optimization—Theory and Algorithms*, Springer,
17 2000.
18
19
20 [16] E. Koutsoupias and C. H. Papadimitriou, “Worst-Case Equilibria,” *Proceedings of the*
21 *16th International Symposium on Theoretical Aspects of Computer Science*, pp. 404–413,
22 Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, March 1999.
23
24
25 [17] C. Landwehr, “Formal Models for Computer Security,” *ACM Computing Surveys*, Vol. 13,
26 No. 3, pp. 247–278, September 1981.
27
28
29 [18] A. S. LaPaugh and C. H. Papadimitriou, “The Even-Path Problem for Graphs and Di-
30 graphs,” *Networks*, Vol. 14 No. 4, pp. 507–513, Winter 1984.
31
32
33 [19] E. L. Lawler, *Combinatorial Optimization—Networks and Matroids*, Holt, Rinehart and
34 Winston, 1976.
35
36 [20] T. Markham and C. Payne, “Security at the Network Edge: A Distributed Firewall Ar-
37 chitecture,” *Proceedings of the 2nd DARPA Information Survivability Conference and Ex-*
38 *position*, Vol. 1, pp. 279–286, June 2001.
39
40
41 [21] M. Mavronicolas, L. Michael, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, “The
42 Price of Defense,” *Proceedings of the 31st International Symposium on Mathematical Foun-*
43 *ditions of Computer Science*, pp. 717–728, Vol. 4162, Lecture Notes in Computer Science,
44 Springer-Verlag, August/September 2006.
45
46
47 [22] M. Mavronicolas, V. G. Papadopoulou, G. Persiano, A. Philippou and P. G. Spirakis,
48 “The Price of Defense and Fractional Matchings,” *Proceedings of the 8th International*
49 *Conference on Distributed Computing and Networking*, pp. 115–126, Vol. 4308, Lecture
50 Notes in Computer Science, Springer-Verlag, December 2006.
51
52
53 [23] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, “A Graph-
54 theoretic Network Security Game,” *Proceedings of the 1st International Workshop on*
55
56
57
58
59
60
61
62
63
64
65

1
2
3 *Internet and Network Economics*, pp. 969–978, Vol. 3828, Lecture Notes in Computer Sci-
4 ence, Springer-Verlag, December 2005. Extended version accepted to *International Journal*
5 *of Autonomous and Adaptive Communications Systems*, Special Issue on Algorithmic Game
6 Theory, R. Kannan and C. Busch eds., June 2008.
7
8

- 9
10 [24] M. Mavronicolas, V. G. Papadopoulou, A. Philippou and P. G. Spirakis, “A Network Game
11 with Attackers and a Defender,” *Algorithmica*, Vol. 51, No. 3, pp. 315–341, July 2008.
12
13 [25] B. Monien, “The Complexity of Determining a Shortest Cycle of Even Length,” *Computing*,
14 Vol. 31, No. 4, pp. 355–369, December 1983.
15
16 [26] J. F. Nash, “Equilibrium Points in N-Person Games,” *Proceedings of the National Academy*
17 *of Sciences of the United States of America*, pp. 48–49, Vol. 36, 1950.
18
19 [27] J. F. Nash, “Non-Cooperative Games,” *Annals of Mathematics*, Vol. 54, pp. 286–295, 1951.
20
21 [28] S. Schechter and M. Smith, “How Much Security Is Enough to Stop a Thief?: The Eco-
22 nomics of Outsider Theft via Computer Systems and Networks,” *Proceedings of the 7th*
23 *International Financial Cryptography Conference*, pp. 122–137, Vol. 2742, Lecture Notes
24 in Computer Science, Springer-Verlag, January 2003.
25
26 [29] E. R. Scheinerman and D. H. Ullman, *Fractional Graph Theory*, John Wiley & Sons, 1997.
27
28 [30] N. Weaver and V. Paxson, “A Worst-Case Worm,” *Proceedings of the 3rd An-*
29 *annual Workshop on Economics of Information Security*, May 2004. Available online at
30 <http://www.icir.org/vern/papers/worst-case-worm.WEIS04.pdf>
31
32 [31] D. B. West, *Introduction to Graph Theory*, Prentice Hall, Second Edition, 2001.
33
34 [32] R. Yuster and U. Zwick, “Finding Even Cycles Even Faster,” *SIAM Journal on Discrete*
35 *Mathematics*, Vol. 10, No. 2, pp. 209–222, May 1997.
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65