

Towards a Metric for Communication Network Vulnerability to Attacks: A Game Theoretic Approach*

Assane Gueye, Vladimir Marbukh ^{†‡}

February 13, 2012

Abstract

In this paper, we propose a quantification of the vulnerability of a communication network when links are subject to failures due to the actions of a strategic adversary. We model the adversarial nature of the problem as a 2-player game between a network manager who chooses a spanning tree of the network as communication infrastructure and an attacker who is trying to disrupt the communication by attacking a link. We use previously proposed models for the value of a network to derive payoffs of the players and propose the network's expected *loss-in-value* as a metric for vulnerability. In the process, we generalize the notion of *betweenness* centrality: a metric largely used in Graph Theory to measure the relative *importance* of a link within a network. Furthermore, by computing and analyzing the Nash equilibria of the game, we determine the actions of both the attacker and the defender. The analysis reveals the existence of subsets of links that are more critical than the others. We characterize these critical subsets of links and compare them for the different network value models. The comparison shows that critical subsets depend both on the value model and on the connectivity of the network. Knowing the critical parts of a network is crucial for network design and improvement. We describe an efficient algorithm that can be used to compute critical subsets of a graph.

***This writeup is still an ongoing effort. Finished version is to be coming soon. We apologize for any inconvenience.**

[†]Assane Gueye and Vladimir Marbukh are with the National Institute of Standard and Technology, email: {assane.gueye|vladimir.marbukh}@nist.gov.

[‡]The authors would like to thank Prof. Jean C. Walrand for his detailed suggestions and reviews of preliminary drafts of this work.

1 Introduction

“....one cannot manage a problem if one cannot measure it...”

This study is an effort to derive a metric that quantifies the vulnerability of a communication network when the links are subject to failures due to the actions of a strategic attacker. Such metric can serve as guidance when designing new networks in adversarial environments. Also, knowing such value helps identify the most critical/vulnerable links and/or nodes of the network, which is an important step towards improving an existing network. We quantify the vulnerability as the expected *loss-in-value* of a network when links are attacked by an adversary. Naturally, the first question towards such quantification is: “what is the *value* of a communication network?”

The value of a network depends on several parameters including the number of agents who can communicate over it. Indeed, the utility of a network increases as it adds more members. The Internet did have the big impact we know today only after it was made public, allowing more users to access it. Social networks platforms as well as email providers are investing a lot of efforts to attract more users, and the popularity of a radio/TV broadcast show is measured by the number of listeners/viewers who are following it. Hence, it seems that everyone agrees that the more members a network has, the more valuable it is. But, there ends the consensus. There is certainly no unanimity on how much this value increases when new members are added, and there is very little (if not zero) agreement on how important a given node or link is for a network, neither do people concur on how much value a given network has.

Attempts to assess the utility of a communication network as a function of the number of its members include the proposition by David Sarnoff [1] who viewed the value of a network as a linear function of its number of nodes n . Robert Metcalfe [10] has suggested that the value of a network grows as a function of the total number of possible connections (roughly n^2). David Reed ([6], [21], [22]) has proposed an exponential (2^n) model for the utility of a network. For Odlyzko *et. al.* ([17], [2]) a more reasonable approximation of the value of a network as a function of the number of nodes is $n \log(n)$. Finally, the authors of the present paper have considered a power law model where the value of a network is estimated as n^{1+a} , $a \leq 1^*$. The parameter a is a design parameter and needs to be specified. Details of these value models are discussed later in section 2.1.

*The authors thank Prof. Jean C. Walrand for suggesting this law.

Each of these very generic models is suitable for a particular network setting, as we will see later. However, they all have a number of limitations; two of which are particularly of interest to us: *They do not take into account the topology of the network neither do they consider the way in which traffic is being carried over the network.* In this paper, we build upon these models and use them in the process to quantifying the vulnerability of a network. More precisely, we use the models as proof of concept to defining the importance of network links *relative to spanning trees*. With this definition, we are implicitly considering networks where information flows over spanning trees. The topology is also taken into account because the set of spanning trees of the network has a one-to-one correspondence with its topology. We are particularly interested in an adversarial situation where links are the target of an attacker. We use a game theoretic approach to model the strategic interaction between the attacker and the defender[†].

Our focus on spanning trees is not a limitation as the techniques of the paper can be used to study other scenarios where the network manager chooses *some* subset of links (shortest path, Hamiltonian cycle, etc...) and the attacker is targeting more than one link as can be seen in [11, Chap. 4]. However, spanning trees have a number of desirable properties that have made them one central concept in communication networking. A spanning tree of a connected graph G is a minimal set of edges that connect all vertices. It can also be defined as a maximal set of edges of G that contains no cycle. Also, in a spanning tree there is a unique path between any two nodes in the network. As a consequence, communicating through a spanning tree guarantees *connectivity* using minimum amount of resources while preventing undesirable loops in the network. The Spanning-Tree Protocol (STP-802.1D 1998) is the standard link management protocol used in Ethernet networks. [19] and [20] review the details of the protocol.

When communication is carried over a spanning tree, any node can reach any other node. In that sense, a spanning tree can be said to deliver the maximum *value* of the network (indeed this ignores the cost of communication). This value can be determined by using one of the models cited above. Now, assuming that information flow through over a given spanning tree, two scenarios are possible when a link of the network fails. If the link does not belong to the spanning, then its failure does not affect the communication. If, on the other hand, the link belongs to the spanning tree, then some exchanges that originally could be carried become impossible. In fact, the spanning tree is separated into two subtrees, each of them being a connected subnetwork. The nodes belonging to each subtree can reach each other, hence each subnetwork delivers some value. However, the sum of the values delivered by the two subnetworks is expected to be less than the value

[†]Throughout this paper we call the defender a “network manager”. This should be thought of as an automata that implements the game.

of the original network (where all nodes could communicate). We define the *importance* of the link, relative to the spanning tree, to be this *loss-in-value* (LIV): *the difference between the value of the original connected network and the sum of the values of the two connected subnetworks obtained as a consequence of the link being down*. The details of this process is discussed in section 2.2.

Link failures occur because of random events (faults) such as human errors, power outage, and machine crashes. These types of failures are dealt with under the subject of *reliability* and *fault tolerance* [16]. Link downtime also can be the result of the action of a malicious attacker. Such failures are the main focus in this study: a strategic attacker is targeting links in the network in order to disrupt the communication. A network manager would like to avoid this disruption by choosing an appropriate communication infrastructure. We model this scenario as a 2-player game where the network manager is choosing a spanning tree to carry the communication in anticipation of the action of an intelligent attacker who is trying to inflict the most damage. The attacker also plans in anticipation of the defense. We use the links' LIV discussed above to derive payoffs for both players.

Applying game theoretic models to the security problem is a natural process and it has recently attracted a lot of interest (see surveys [23], [15]). In this paper, we setup a game on the graph of a network and consider Nash equilibrium concept. We propose the expected “loss-in-value” of the game for the network manager as a metric for the vulnerability. As we will see later, this value captures how much loss an adversary can inflict to the network manager by attacking links. Furthermore, by analyzing the Nash equilibria of the game, we determine the actions of both the attacker and the defender and identify the set of links that are most critical for the network. We repeat the analysis for each of the network value models cited above and discuss how the set of critical links depends on the network value model. Knowing the critical parts of a network is crucial for network design and for improvement. We discuss an efficient algorithm that can be used to compute critical subsets of a graph.

The quantification of the *importance* of a communication link as discussed above is done with respect to spanning trees. For a given link, we compute a different value for each spanning tree. This does not tell how important the link is for the whole communication process. To compute the relative significance of a link (or node) within a network, graph theorists have been using *centrality* measures. There are a few different notions of centrality ([3], [8], [18]). Freeman [8] has classified them into *degree*, *closeness* and *betweenness* centrality. The betweenness centrality ([26], [18]) of a link (node) is the fraction of time that it is needed in the shortest paths from all possible sources to all possible destinations.

In this definition of betweenness centrality, source and destination pairs are indifferent with respect to which of the available paths is carrying their communications. Also, the *importance* of a link to a given path is quantified in a binary manner with a value of 1 if the link belongs to the path a 0 otherwise. In general, certain paths might be more preferred than others and the importance of a link to a path does not necessarily have a binary weight.

In this paper, we propose a generalization of the notion of betweenness centrality[‡] assuming networks where information flow over spanning trees; we hence use spanning trees in lieu of paths . Our generalization allows both the consideration of none-binary weights of the links as well as preference for spanning tree utilization. The weights of the links, with respect to spanning trees, are taken to be the links' LIV mentioned above, and the preference among the spanning trees is given by the mixed strategy Nash equilibria of the game between the network manager and the attacker. Indeed, the mixed strategy equilibrium for the manager is a natural metric to quantify how much one spanning tree is preferred to another: it is the best response to the actions of an intelligent adversary who is trying to cause the most damage.

The remainder of this paper is organized as follow. The next section 2.1 discusses the different network value models that we briefly introduced above. We use these models to compute the relative importance of the links with respect to spanning trees. This is shown in section 2.2, followed by our generalization of the notion of betweenness centrality in section 2.3. The strategic interaction between the network manager and the attacker is modeled as a 2-player game which is presented in section 3.1. The Nash equilibrium theorem of the game is stated in section 3.2 followed by a discussion and analysis of its implications in section 4. Section 4.1 discusses our choice of metric for the vulnerability of a network. In section 4.2 we compare the critical subsets of a network for the different value models cited above. The algorithm to compute a critical subset is briefly presented in section 4.3. Concluding remarks and future directions are presented in section 5. All our proofs as well as the details of the algorithm to compute a critical subset are presented in the appendix.

2 On the Value of a Communication Networks

The value of a network depends on several parameters including the number of nodes, the number of links, the topology, and the type of communication/information that is carried over the network. Assessing such

[‡]Notice that in this paper, we are only interested in the value of links. Node value is an equally important topic that we will consider in subsequent studies.

value is a subjective topic and, to the knowledge of the authors, there is no systematic quantification of the value of a communication network. In the next section, we discuss some attempts that have been made to measure the utility of a network as a function of its number of nodes.

2.1 Network value models

Sarnoff's Law:

Sarnoff's law [1] states that *the value of a broadcast network is proportional to the number of users*. This law has long been used as a measure of the popularity of television and radio programs. The high advertising cost during prime time shows and other popular events like super bowl game night can be explained by Sarnoff's law. Indeed as more viewers are expected to watch a program, a higher price is charged per second of advertising. Although Sarnoff's law has been widely accepted as a good model for broadcast network, many critics say that it underestimates the value of general communication networks such as the Internet.

Metcalf's Law:

Metcalf's law [7] was first formulated by George Gilder (1993) and attributed to Robert Metcalfe who used it mostly in the context of the Internet. The law states that *the value of a communication network is proportional to the square of the number of node*. Its foundation is the observation that in a general network with n nodes, each node can establish $n - 1$ connections. As a consequence, the total number of undirected connections is equal to $n(n - 1)/2 \sim O(n^2)$. This observation is particularly true in Ethernet networks where everything is "logically" connected to everything else. Metcalfe's law, has long been held up along side with Moore's law as the foundation of Internet growth. Reed Hunt (former FCC chairman 1996) once claimed that Metcalfe's law "gives us the best foundation for understanding the Internet".

Walrand's Law:

Walrand's law generalizes the previous laws by introducing a parameter a . The intuition behind this law is as follows. Imagine a large tree of degree d that is rooted at you. Your direct children in the tree are your friends. The children of these children are the friends of your friends, and so on. Imagine that there are $L \geq 2$ levels. The total number of nodes without the root is $n = d + d^2 + \dots + d^L = d(d^L - 1)/(d - 1)$. If d is large, this number can be roughly approximated by $n \approx d^L$. However, you don't care about the friends of friends of friends... Assume you care only about your friends, i.e., about d people. Then the value of the network to you is $d = n^a$ where $a = 1/L$. If you care about your friends and their friends, then you care about approximately d^2 people, i.e., n^b where $b = 2/L$. If you care about friends at all levels, then

you care about approximately $d^L \approx n$. Repeating the same reasoning for each user (node), the total value of the network is approximately equal to $n * n^a = n^{1+a}$ with $a \leq 1$. The parameter a is a characteristic of the network and needs to be determined. Notice that if all nodes value children at all levels, the total value of the network becomes n^2 which corresponds to the Metcalfe's law ($a = 2$), and if $a = 0$, we get Sarnoff's model.

Reed's Law:

Reed's law also called Group-Forming law was introduced by David Reed ([21],[6], [22]) to quantify the value of networks that support the construction of communicating group. A group forming network resembles a network with smart nodes that, on-demand, form into such configurations. Indeed, the number of *possible groups that can be formed over a network of n nodes is $\sim 2^n$* . Reed's law has been used to explain many new social network phenomenons. Important messages posted on social networking platforms such as Twitter and Facebook have been witnessed to spread exponentially fast.

Odlyzky, Briscoe, and Tilly (OBT)'s Law:

Odlyzky, Briscoe, and Tilly ([2], [17]), have proposed an $n \log(n)$ rule for the valuation of a network of size n . Their law is mostly inspired by Zipf's law that states that *if we order a large collection of items by size or popularity, the second element in the collection will be about half the measure of the first, the third element will be about 1/3 of the first, and the k -th element will measure about $1/k$ of the first*. Setting the measure of the first element (arbitrarily) to 1, the sequence looks like $(1, 1/2, 1/3, \dots, 1/k, \dots, 1/n)$. Now, assuming that each node in the network assigns value to the other nodes according to Zipf law, then the total value of the network to any given node will be proportional to the harmonic sum $1 + 1/2 + 1/3 + \dots + 1/(n-1)$, which approaches $\log(n)$. Summing over the nodes, we get the $n \log(n)$ rule. This growth rate is faster than the linear growth of Sarnoff's law and does not have the overestimating downside that is inherent to Reed and Metcalfe. It also has a diminishing return property which is missing in Walrand's model.

2.2 Assessing *importance* of links via spanning trees

Assuming that a model has been determined for the value of a network, we quantify the importance of a network link with respect to a spanning tree as the *loss-in-value* (LIV) when the link fails while communication is carried over the tree. We will use this to derive players' payoff in the game model we present in section 3.1. As we have stated earlier our choice of spanning tree is not a limitation as our techniques can be utilized to define relative importance of a link with respect other notions (paths, cycles, etc...). However, spanning trees




Figure 1: Tree-Link model for the value of a network link. a) Complete network with link 'e' of interest shown in bold. b) A particular spanning tree 'T' of the graph. c) Disconnected network when link 'e' is removed.

have a number of desirable properties that have made them a central concept in communication networking and the focus of this paper.

The loss-in-value of a link, relative to a given spanning tree, is determined as follow. Since a spanning tree is a connected subgraph of the network (see figure (1).b)), communicating over such subgraph can be seen as delivering the total “value” of the network. We let $f(n)$ denote such value, where n is the number of nodes in the network. If $\eta(T)$ is the cost of maintaining spanning tree T , the net delivered value is $f(n) - \eta(T)$. The function $f(\cdot)$ can be chosen to be one of the models discussed earlier. In general it depends on other aspects of the network. However, in this paper, $f(n)$ is assumed to be only a function of the number of nodes. We let $f(0) = 0$, if the network contains 0 node (i.e is empty).

Now, assume that communication is carried over spanning tree T and a particular link e is removed from the network. If $e \in T$ (assuming a spanning tree to be a set of links), then T is partitioned into 2 subtrees; each subtree T_i , $i \in \{1, 2\}$ represents a connected components with n_i nodes, where $n_1 + n_2 = n$ (see figure (1).c)). The net value of the resulting disconnected network is $f(n_1) + f(n_2) - \eta(T)$, where $f(n_i)$ is the value of the connected component i .

When link e is removed, some exchanges that could be carried on the original network become impossible. As of such, it is reasonable to assume that $f(\cdot)$ is such that $f(n) \geq f(n_1) + f(n_2)$. Notice that this is the case for all the network value models cited above. Hence, removing link $e \in T$ has reduced the value of the network from $f(n) - \eta(T)$ to $f(n_1) + f(n_2) - \eta(T)$. We define the LIV of link e , relative to spanning tree T , to be equal to $f(n) - (f(n_1) + f(n_2))$.

If the link does not belong to the spanning tree, then removing it will leave the network connected and the value of network is still $f(n) - \eta(T)$. As a consequence, if link $e \notin T$, its LIV, relative to spanning tree T , is

equal to 0. In conclusion, the value of link e , relative to spanning tree T is defined as

$$\bar{\lambda}(T, e) = f(n) - (f(n_1) + f(n_2)), \quad (1)$$

with the understanding that if $e \notin T$, $n_1 = n$ and $n_2 = 0$, giving $\bar{\lambda}(T, e) = 0$.

Writing this expression for all spanning trees and all links of the network, we build the tree-link LIV matrix $\bar{\Lambda}$ defined by $\bar{\Lambda}[T, e] = \bar{\lambda}(T, e)$.

Finally, since $f(n)$ is assumed to be the same for all T , we can define the normalized LIV of a link e relative to a spanning tree T as

$$\lambda(T, e) = 1 - \frac{f(n_1) + f(n_2)}{f(n)}. \quad (2)$$

The normalized tree-link LIV matrix Λ is defined accordingly.

Remark 1 *With the definition in (1), the LIV of a link relative to any spanning tree is always equal to zero under Sarnoff's law (i.e. $\lambda(T, e) = 0$, $\forall e$ and T). This is the effect of the linear rule (indeed $n = n_1 + n_2$) and the fact that the broadcast network has a tree structure (every edge belongs to the tree). If there were links between receivers, then attacking such links would not disrupt the communication and the value of the network would still be n . We generalize this idea to arbitrary network topology by using the model introduced in [13]), which we denote GWA model. It is a simple model that gives the same normalized LIV of 1 if the link e belongs to the spanning tree and 0 otherwise (i.e. $\lambda(T, e) = \mathbf{1}_{e \in T}$). The model basically assumes that whenever a link on the spanning tree is removed (hence disconnecting it), the network loses its entire value. This is the case in distributed applications where each single node has to receive the sent information in order for an operation to be carried (e.g. consensus).*

Table (1) shows the LIV of links for the different models presented above. It is assumed that removing link e divide spanning tree T into two subtrees with respectively n_1 and n_2 nodes ($n_1 + n_2 = n$)

2.3 A generalization of the betweenness centrality measure

The quantification we have described above for the significance of a link is relative to spanning trees: *there is a different value for each different tree*. In general, one would like to get a sense of the importance of a link for the overall communication process. Betweenness centrality is a measure that have long been used for that purpose. Next, we propose a quantification of the importance a link within a network that generalizes

Table 1: Normalized LIV of link e relative to spanning tree T for the different laws. Removing link e from spanning tree T divide the network into two subnetwork with respectively n_1 and n_2 nodes ($n_1 + n_2 = n$).

Model	Normalized LIV
GWA	$1_{e \in T}$
Metcalf	$1 - \frac{n_1^2 + n_2^2}{n^2}$
Reed	$1 - 2^{-n_1} - 2^{-n_2}$
BOT	$1 - \frac{n_1 \log(n_1) + n_2 \log(n_2)}{n \log(n)}$
Walrand	$1 - \frac{n_1^{1+a} + n_2^{1+a}}{n^{1+a}}$

the notion of betweenness. We start by recalling the betweenness centrality measure as it was defined by Freeman [8].

For link e , and nodes i and j , let $g_{i,j}$ be the number of shortest paths between i and j and let $g_{ij}(e)$ the numbers of those paths that contain e . The partial betweenness centrality of e with respect to i and j is defined as $\vartheta_{ij}(e) = \frac{g_{ij}(e)}{g_{ij}}$ and the betweenness centrality of e is defined as $\vartheta(e) = \sum_{i < j} \vartheta_{ij}(e)$.

Freeman has given a probabilistic interpretation of the partial betweenness. If we assume that the two points i and j are indifferent with respect to which of several alternative geodesics carries their communications, the probability of using any one is $\frac{1}{g_{ij}}$. The partial betweenness is hence equal to the expected number of shortest paths between i and j that use e . Also, notice that in this definition, the *importance* of link e with respect to a path p_{ij} is equal to 1 if $e \in p_{ij}$ and 0 if not. $g_{ij}(e)$ is the total *value* of e for a communication between i and j .

Using this interpretation, we can easily generalize the betweenness centrality to quantify the importance of a link as

$$\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \sum_{i < j} \alpha_{p_{ij}} \lambda(p_{ij}, e), \quad (3)$$

where the parameter $\boldsymbol{\lambda}$ is a vector which entries are the relative importance $\lambda(p_{ij}, e)$ of link e to path p_{ij} , and $\boldsymbol{\alpha}$ is a vector having as entries the probabilities (preference) $\alpha_{p_{ij}}$ of using path p_{ij} . In the standard definition of betweenness centrality, parameter $\lambda(p_{ij}, e) = g_{ij}(e)$ and $\alpha_{p_{ij}} = \frac{1}{g_{ij}}$. In general, these parameters can be determined by considering relevant aspects of the communication network. For example, in this paper, the parameters $\boldsymbol{\lambda}$ are chosen to be equal to the LIV of the links relative to spanning trees, as was defined in the previous section. Again, here we focus on networks where information flows through spanning trees. With this, the definition in (3) becomes

$$\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \sum_T \alpha_T \lambda(T, e), \quad (4)$$

where the summation is now over spanning trees.

The preference parameter $\alpha(T)$ can be driven by many factors. The cost of utilizing the links, the overall communication delay and the vulnerability of links are, among others, reasons why, in a communication network, some spanning trees might be more preferred than others. In this paper, where the focus is on adversarial environments, the choice of spanning tree is dictated by the vulnerability factor. More precisely, the preference $\alpha(T)$ are chosen to be the mixed strategy Nash equilibrium in a game between a network manager who chooses a spanning tree as communication infrastructure and an adversary who tries to cut off the communication by attacking one link. Details of the game are presented next.

3 Game Theoretic Approach

3.1 Game model

The game is over the links of the network with a topology given by a connected undirected graph $G = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{E}| = m$ links and $|\mathcal{V}| = n$ nodes. The set of spanning trees is denoted $\mathcal{T}$; we let $N = |\mathcal{T}|$.

To get all nodes connected in a cycle-free way, the network manager chooses a spanning tree $T \in \mathcal{T}$ of the graph. Running the communication on spanning tree T requires a maintenance cost of $\eta(T)$ to the network manager. The attacker simultaneously selects an edge $e \in \mathcal{E}$ to attack. Each edge $e \in \mathcal{E}$ is associated with some cost $\mu(e)$ that an attacker needs to spend to launch a successful attack on e . It also has a *loss-in-value* For the network manager, the value $\lambda(T, e)$ of link e relative to spanning tree T is given by (2). This is how much the network manager loses when he chooses tree T and link e happens to be attacked. This loss goes to the attacker. More precisely, for a choice pair (T, e) of tree and edge, the net loss is $\eta(T) + \lambda(T, e)$ for the network, while the net attack reward is equal to $\lambda(T, e) - \mu(e)$ for the attacker. It is also assumed that the attacker has the option of not attacking, which results in a zero net reward for the attacker and a zero loss for the manager. We let $e_\emptyset$ denote that option.

The pure strategy sets are the set $\mathcal{T}$ of spanning trees for the manager and the set $\mathcal{E}$ of edges for the attacker. We are mainly interested in analyzing mixed strategy Nash equilibria of the game. We let $\{\alpha \in \mathbb{R}_+^N \mid \sum_{T \in \mathcal{T}} \alpha_T = 1\}$ be the set of mixed strategies for the network manager, and $\{\beta \in \mathbb{R}_+^m \mid \sum_{e \in \mathcal{E}} \beta_e = 1\}$ the set of mixed strategies for the attacker. The defender is choosing α to minimize the expected net

communication cost $L(\alpha, \beta)$ while the attacker is choosing β to maximize the expected net reward $R(\alpha, \beta)$.

$$L(\alpha, \beta) = \sum_{T \in \mathcal{T}} \alpha_T \left(\eta(T) + \sum_{e \in T} \beta_e \lambda(T, e) \right), \quad (5)$$

$$R(\alpha, \beta) = \sum_{e \in \mathcal{E}} \beta_e \left(\sum_{T \ni e} \alpha_T \lambda(T, e) - \mu(e) \right). \quad (6)$$

In the sequel, we have focused on the case where $\eta(T) = \eta$ is constant; hence not relevant to the optimization in (24), which now becomes the minimization of $\sum_{T \in \mathcal{T}} \alpha_T \sum_{e \in T} \beta_e \lambda(T, e)$. As a consequence, we ignore $\eta(T)$ for the rest of this paper. The general case of $\eta(T)$ will be considered in subsequent studies.

3.2 Nash equilibrium theorem

To state the NE theorem of the game, we need to make a certain number of definitions.

For each subset of edges $E \subseteq \mathcal{E}$, we let Λ_E be the matrix Λ where columns corresponding to links not in E are set to zero. Matrix Λ is defined in section 2.2 and its entries are given in (2). Now, assume that the attacker's equilibrium strategy β has support E . If there is a spanning T such that $E \cap T = \emptyset$, then the defender's best response is to choose T ; and the attack would give zero payoff to the attacker. Hence, if an attacker were to focus her attack on a set E , she would like that each spanning tree contains at least some link of E . In the context our game, E should be such that for all spanning tree T , there exists $e \in E$ such that $\lambda(T, e) > 0$. Thus, if a mixed strategy equilibrium of the attacker has the form $\beta = \mathbf{y}/(\mathbf{1}'\mathbf{y})$ with support E , one can "correctly" guess that $\mathbf{y}$ verifies $\Lambda_E \mathbf{y} \geq \mathbf{1}$ (since each entry of the vector $\Lambda_E \mathbf{y} > 0$ we can always divide by the smallest entry). Let $\mathcal{B}_E$ be the collection of all such $\mathbf{y}$, $\mathcal{B}_E = \{\mathbf{y} \in \mathbb{R}_+^m \mid \Lambda_E \mathbf{y} \geq \mathbf{1}\}$. $\mathcal{B}_E$ is a polyhedron and can be characterized by its set of vertex points.

For such strategy $\beta (= \mathbf{y}/(\mathbf{1}'\mathbf{y}))$ of the attacker, the defender's best response is to choose an α that put positive weight only on rows of Λ_E (i.e. spanning trees T) that achieve the minimum

$$\min_T \left(\sum_{e \in T \cap E} \beta_e \lambda(T, e) \right) = \min_T \left(\sum_{e \in T \cap E} \frac{y_e}{\mathbf{1}'\mathbf{y}} \lambda(T, e) \right) = \frac{\min_T (\sum_{e \in T \cap E} y_e \lambda(T, e))}{\mathbf{1}'\mathbf{y}}. \quad (7)$$

Now, if $\mathbf{y}$ is a vertex of the polyhedron $\mathcal{B}_E$, one can show that there exists a spanning tree T_o such that the minimum in the numerator is $\sum_{e \in T_o \cap E} y_e \lambda(T_o, e) = 1$. We discuss the details of this assertion in the appendix. This is the minimum value possible because of the constraints $\Lambda_E \mathbf{y} \geq \mathbf{1}$. With these choices of

strategies by both players, the attacker's expected payoff can be written as:

$$\bar{R} = \frac{1}{\mathbf{1}'\mathbf{y}} - \sum_{e \in E} \frac{\mathbf{y}_e}{\mathbf{1}'\mathbf{y}} \mu(e), \quad (8)$$

where again E is the support of $\mathbf{y}$. The attacker would like this payoff to be maximized. As a consequence, she wants to minimize the quantity $\mathbf{1}'\mathbf{y}$ among all $\mathbf{y}$ in $\mathcal{B}_E$ (this is obvious—at least when the attack cost $\mu = \mathbf{0}$). This leads us to the following definition.

Definition 1 *For any subset of links $E \subseteq \mathcal{E}$, the function $\kappa(E)$ is defined as*

$$\begin{aligned} \kappa : 2^{\mathcal{E}} &\longrightarrow \mathbb{R}_+ \\ E &\longmapsto \kappa(E) = \min\{\mathbf{1}'\mathbf{y}, \mathbf{y} \in \mathcal{B}\}. \end{aligned} \quad (9)$$

$\kappa(E)$ is the value of a linear program (LP) that might be infeasible (e.g. when a row of Λ_E is all zeros). However, its dual is always feasible (see Appendix E), and when the dual LP is bounded, the primal is necessarily feasible [4]. Let $\mathbf{y}_E$ be a solution of the primal program whenever the dual LP is bounded. If this dual is unbounded for some subset E , we let $\mathbf{y}_E = K\mathbf{1}_m$, for an arbitrary large constant K , where $m = |\mathcal{E}|$, and $\mathbf{1}_m$ is the all-ones vector of length m . With such “fix”, $\kappa(E) = m * K$ when the dual LP is unbounded. Hence, we can define the following two quantities that are induced by E .

Definition 2 *The induced probability distribution is defined as $\beta_E = \mathbf{y}_E / \kappa(E)$.*

And the induced expected net reward for the attacker is denoted by $\theta(E)$, with

$$\theta(E) := \frac{1}{\kappa(E)} - \sum_{e \in E} \beta_E(e) \mu(e). \quad (10)$$

Finally, we call a subset of vertices E critical if its induced expected rewards is maximum, i.e. if

$$\theta^* := \theta(E) = \max_{\tilde{E}} \left(\theta(\tilde{E}) \right), \quad (11)$$

and we let $\mathcal{C}$ be the set of all critical subsets. We call θ^ the vulnerability of the network.*

Remark 2 *When proving the Nash equilibrium theorem below (see Appendix E), we will argue that a critical subset of edges E must be a minimal disconnecting set. For such subsets of links the dual of the LP defining*

$\kappa(E)$ is always bounded (hence the primal is always feasible). As a consequence, for a critical subset E , we always have that $0 < \kappa(E) < \infty$ and the corresponding $\mathbf{y}_E$ (hence β_E) is always well-defined.

We are now ready to state the Nash equilibrium theorem of the game. We claim that:

Theorem 1 *For the game defined above, the following always hold.*

1. If $\theta^* \leq 0$, then “No Attack” (i.e. $\beta(e_\emptyset) = 1$) is always an optimal strategy for the attacker. In this case, the equilibrium strategy $(\alpha_T, T \in \mathcal{T})$ for the defender is such that

$$\vartheta(e, \lambda, \alpha) = \sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) \leq \mu(e), \quad \forall e \in \mathcal{E}. \quad (12)$$

The corresponding payoff is 0 for both players.

2. If $\theta^* \geq 0$, then for every probability distribution $(\gamma_E, E \in \mathcal{C})$ on the set of critical subsets, the attacker’s strategy $(\beta(e), e \in \mathcal{E})$ defined by

$$\beta(e) = \sum_{E \in \mathcal{E}} \gamma_E \beta_E(e) \quad (13)$$

is in Nash equilibrium with any strategy $(\alpha_T, T \in \mathcal{T})$ of the defender that satisfies the following properties:

$$\begin{cases} \vartheta(e, \lambda, \alpha) - \mu(e) = \theta^* & \text{for all } e \in \mathcal{E} \text{ such that } \beta(e) > 0. \\ \vartheta(e, \lambda, \alpha) - \mu(e) \leq \theta^* & \text{for all } e \in \mathcal{E}. \end{cases} \quad (14)$$

Furthermore, there exists at least one such strategy α .

The corresponding payoffs are θ^* for the attacker, and $r(\gamma)$ for the defender, where

$$r(\gamma) := \sum_{E \in \mathcal{C}} \frac{\gamma_E}{\kappa(E)}. \quad (15)$$

3. If $\mu = 0$, then every Nash equilibrium pair of strategies for the game is of this type.

The proof of the theorem is presented in the Appendix. Its implications are discussed next.

4 Discussion and Analysis

The NE theorem has three parts. If the quantity θ^* is negative then the attacker has no incentive to attack. For such choice to hold in an equilibrium, the defender has to choose his strategy α as given in (12). Such

α always exists. When $\theta^* \geq 0$ there exists an equilibrium under which the attacker launches an attack that focuses only on edges of critical subsets. The attack strategies (probability of attack of the links) are given by convex combinations of the induced distributions of critical subsets. The corresponding defender's strategies are given by (14). When there is no attack cost, the attacker always launches an attack ($\theta^* > 0$) and the theorem states that all Nash equilibria of the game have the structure in 14.

4.1 Vulnerability metric and importance of links

For simplicity, let's first assume that there is no attack cost i.e $\mu = \mathbf{0}$. In this case, $\theta(E) = \frac{1}{\kappa(E)}$ and $\theta^* > 0$. Also, a subset of link E is critical if and only if $\kappa(E)$ is minimal. Since in this case the game is zero-sum, the defender's expected loss is also $\theta^* = (\min_E \kappa(E))$. θ^* depends only on the graph and the network value model ($f(n)$). It measures the worst case loss/risk that the network manager is expecting in the presence of any (strategic) attacker. Notice that in our setting, a powerful attacker is one who does not have a cost of attack $\mu = \mathbf{0}$. When θ^* is high, the potential loss in connectivity is high. When it is low, an attacker a very little to gain in attacking, hence the risk low. As of such, we propose θ^* as a measure of the vulnerability of the graph.

This vulnerability metric also corresponds to a quantification of the *importance* of the most critical links. This is captured by the inequalities in (14), which, when $\mu = \mathbf{0}$, become

$$\vartheta(e, \lambda, \alpha) \leq \theta^* \text{ for all } e \in \mathcal{E}, \quad (16)$$

with equality whenever link e is targeted with positive probability ($\beta(e) > 0$) at equilibrium. From (13) we see that $\beta(e) > 0$ only if edge e belongs to a critical subset, and hence is critical. Thus, the attacker focuses its attack only on critical links, which inflict the maximum loss to the defender.

For the defender, since the game is zero-sum, the Nash equilibrium strategy corresponds to the *min-max* strategy. In other words, his choice of α minimizes the maximum expected loss. Hence, the defender's equilibrium strategy α can be interpreted as the best way (in the min-max sense) to choose a spanning tree in the presence of a strategic adversary. Using this interpretation with our generalization of betweenness centrality in (4), we get a way to quantify the *importance* of the links to the overall communication process. The inequalities in (16) above say that the links that are the most important to the defender (i.e. with maximum $\vartheta(e, \lambda, \alpha)$) are the ones that are targeted by the attacker (the most critical). This unifies the

positive view of *importance* of links when it comes to participation to the communication process to the *negative* view of *criticality* when it comes to being the target of a strategic adversary. This is not surprising because since the attacker’s goal is to cause the maximum damage to the network, it makes sense that she targets the most important links.

When the cost of attack is not zero ($\mu \neq \mathbf{0}$), our vulnerability metric θ^* takes it into account. For instance, if the attacker has spent too much effort to successfully launch an attack, to the point where (the expected net reward) θ^* is negative, the theorem tells that, unsurprisingly, the attacker will choose to not launch an attack. To “force” the attacker to hold to such choice (i.e to maintain the equilibrium), the defender has to randomly pick a spanning tree according to (12). With this choice, the relative value of any link is less than the amount of effort needed to attack it. When μ is known, such strategy can be seen as a deterrence tactic for the defender.

If the vulnerability θ^* is greater than zero, then there exists an attack strategy that only targets critical links. To counter such attack, the defender has to draw a spanning tree according to the distribution α in (14). For such choice of a tree, the relative importance of any critical link, offset by the cost of attacking the link is equal to the θ^* . For any other link, this difference is less than θ^* . In this case, the criticality of a link is determined not only by how much importance it has for the network, but also how much it would take to the adversary to successfully attack it. Hence, when $\mu \geq \mathbf{0}$, θ^* is a measure of the willingness of an attacker to launch an attack. It includes the loss-in-value for the defender as well as the cost of attack for the attacker.

Observe that when $\mu \geq \mathbf{0}$ the theorem does not say anything about the existence of other Nash equilibria. It is our conjecture (verified in all experiments) that even if there were other equilibria, θ^* is still the maximum payoff that the attacker could ever receive. Hence, it measures the worst case scenario for the defender.

4.2 Critical subsets and network value models

In this section we discuss how the critical subsets depend on the model used for the value of the network. Figure 2 shows an example of network with the critical subsets for the different value models discussed earlier. The example shows a core network with a set of bridges connecting it to peripheral nodes. A bridge is a single link the removal of which disconnects the network. In all figures, the critical subset of links is shown in dashed. In this discussion we mainly assume that the attack cost μ is equal to zero.




Figure 2: Example of critical subsets for different value models. a) GWA model b) OBT, Walrand, and Metcalfe's models. c) Reed's model.

Figure 2.a) shows the critical subset corresponding to the GWA link cost model introduced in [12] for which $\lambda(T, e) = \mathbf{1}_{e \in T}$. With this model, the defender loses everything (i.e. 1) whenever the attacked link belongs to the chosen spanning tree. Since a bridge is contained in any spanning tree, attacking a bridge gives the maximum outcome to the attacker. As a consequence, the critical subsets correspond to the set of bridges as can be observed in the figure. In fact, for the GWA value model, we have shown in [12] that $\kappa(E) = \frac{|E|}{\mathcal{M}(E)}$, where $\mathcal{M}(E) = \min_T (|T \cap E|)$. Notice that if E is a disconnecting set (i.e. removing the edges in E divides the graph into 2 or more connected components), $\mathcal{M}(E) \geq 1$. Now, if e is a bridge, $|T \cap \{e\}| = 1$ for all spanning trees T , implying that $\mathcal{M}(\{e\}) = 1$ and $\theta(\{e\}) = \kappa(\{e\}) = 1$, which is the maximum possible value of θ^* . As a consequence, each bridge is a critical subset and any convex combination over the bridges yields an optimal attack.

Figure 2.b) depicts the critical subsets with the Metcalfe, OBT, and Walrand ($a = 0.6$) models. For all these models (as well as for Reed's model), the function $f(x) - (f(x_1) + f(x_2))$, where $x_1 + x_2 = x$, is maximized when $x_1 = x_2 = x/2$. This suggests that attacks that target links that evenly divide (most) spanning trees are optimal. This *conjecture "seems"* to be confirmed by the examples shown in the figure. The most critical links are the *innermost* or *core* links of the network for all three models. The attack distributions are slightly different for the 3 models. The distribution on links (1, 2, 3, 4, 5) is equal to (0.1875, 0.2500, 0.1875, 0.1875, 0.1875) for Metcalfe, (0.1911, 0.2356, 0.1911, 0.1911, 0.1911) for OBT, and (0.1884, 0.2465, 0.1884, 0.1884, 0.1884) for Walrand ($a = 0.6$). Notice that for all models, the middle link (2) is attacked with a higher probability.

Although Reed's (exponential) model also has the same property discussed in the previous paragraph, the critical subset with Reed is different, as can be seen in figure 2.c). While Metcalfe, OBT, and Walrand models lead to the core network being critical, with Reed's model, the critical links are the access to the




Figure 3: Comparison of the loss functions $1 - \frac{f(n_1) + f(n-n_1)}{f(n)}$ when a link belonging to the chosen spanning tree is cut, dividing it into 2 subtrees of n_1 and $n - n_1$ nodes. For GWA, since $\lambda_{Te} = \mathbf{1}_{e \in T}$, the loss is always 1. The models GWA, Reed, and Walrand (for large values of a), overlap in a wide region of values of n_1 .

core. Each of the links is attacked with the same probability. This might be a little surprising because it contradicts the conjecture that innermost links tend to be more critical. However, observing the attack's reward function $\left(1 - \frac{f(n_1) + f(n-n_1)}{f(n)}\right)$ shown in figure 3, Reed's model coincides with the GWA model in a wide range of n_1 . This means that any link that separates (most of the spanning) into subtrees of n_1 and $n - n_1$ nodes gives the maximum reward to the attacker, for most values of n_1 . Also, notice that since the core network is "well connected", the defender has a many options for choosing a spanning tree. This means that in the core, the attacker has less chances of disrupting the communication. Links accessing the core, on the other hand, deliver high gain and better chances of disrupting the communication. Hence, the best strategy for the attacker is, in this case, to target access to the core. Notice that Metcalfe, OBT, and Walrand ($a \leq 1$) models do not have this optimal *tradeoff choice*.

By choosing the parameter a to be sufficiently large in the Walrand model, we have (experimentally) observed that the critical subset moves from being the core, to corresponding to the one in GWA model (the bridges) for very large values of a . In fact, with all network topologies we have considered (more than 50), we could always choose the parameter of the Walrand so that the critical subset matches the one in GWA model. This implies that as the model loss function $\left(1 - \frac{f(n_1) + f(n-n_1)}{f(n)}\right)$ gets closer to the GWA function $\mathbf{1}_{e \in T}$, the critical subset moves away from the inner links to the outer links.

These observations indicate that the critical subsets of a graph depend on the value model used to setup the game. The value model is however not the only factor that characterizes the critical subset(s) of a graph.




Figure 4: Example of critical subsets for different value models. a) GWA model b) OBT, Walrand, and Metcalfe's models. c) Reed's model

Figure 4 shows the same network as in the previous example with one additional (core) link. With this additional link, the connectivity of the core network is enhanced. The critical subset does not change for the GWA models. However, for all other 4 models, the critical subset is now the access to the core. This suggests that *connectivity* is another factor that characterizes the critical subset(s).

As was (experimentally) observed on the previous example, in this case also, when the parameter a of Walrand's model is chosen sufficiently large, the critical subsets become the same as the GWA critical subsets.

4.3 Computing a critical subset: an algorithm

Computing a critical subset requires two things: an oracle that returns $\kappa(\cdot)$ and an algorithm that computes the maximal value of $\theta(E)$ over the set of subsets of links.

The oracle can be implemented by any efficient algorithm that solves a linear program [4]. The simplex method [25] can be used for that purpose. The worst case complexity of the simplex is exponential [14], however, for most practical cases, it runs very efficiently.

To compute θ^* and the corresponding maximizer E (critical subset), we use the fact that the function $\kappa(E)$ is a submodular function [24]. When the cost of attack is equal to zero, maximizing $\theta(E)$ is the same as minimizing $\kappa(E)$. The minimization of a submodular function has been largely studied [24] and polynomial algorithms have been proposed for that purpose. Using one of these algorithms with efficient methods to solve linear program, one can build an efficient algorithm to compute critical subsets of a network. We discuss the details of this in the appendix.

5 Conclusion and Future Work

In this study, we quantify the vulnerability of a communication network where links are subject to failures due to the actions of a strategic attacker. Such a metric can serve as guidance when designing new communication networks and determining it is an important step towards improving existing networks.

We build upon previously proposed models for the value of a network, to quantify the *importance* of a link, relative to a spanning tree, as the *loss-in-value* when communication is carried over the tree and the link is failed by a strategic attacker. We use these values to setup a 2-player game where the defender (network manager) chooses a spanning tree of the network as communication infrastructure and the attacker tries to disrupt the communication by attacking one link. We propose the equilibrium's expected *loss-in-value* as a metric for the vulnerability of the network. We analyze the set of Nash equilibria of the game and discuss its implications. The analysis shows the existence of subsets of links that are more critical than the others. We characterize these sets of critical subsets and, using examples, we show that such critical subsets depend on the network value model as well as the connectivity of the graph. The nature of this dependency is an interesting question that we are planning to investigate in future studies. We also discuss an efficient algorithm that can be used to compute critical subsets. Finally, we propose a generalization of the notion of betweenness centrality that allows different weights for the links as well as preference among the graph structures that carry the communication (e.g. spanning trees for this paper).

Several future directions are being considered as followup to this paper. First, in here, we have discussed the critical subsets using illustrative examples. To get a better intuition about the relationship between the value function and the critical subset of the network, a more rigorous analysis of the game value function ($\kappa(E)$) is needed. With such analysis we will be able to integrate and understand more realistic (and potentially more complicated) network value models. Also, in this paper, we use spanning trees to define the relative importance of links. This implicitly considers only networks in which things flow through spanning trees. However, our result is general and can be used to study games on other type of networks. One interesting extension is the situation where the network manager chooses $p \geq 1$ spanning trees (example $p = 2$ is the situation where the manager chooses a communication tree and a backup one), and the attacker has a budget to attack $k \geq 1$ links. Also, we have assumed, in this paper, that the cost of communicating over any spanning tree is the same. In the future, we will study versions of the problem where some spanning trees might be more costly than others. Finally, this study has focused on the failure of links in a network.

Nodes also are subject failures: whether random or strategic. A more thorough study should consider both links and nodes.

References

- [1] USN Admiral James Stavridis. Channeling David Sarnoff, Sept 2006. <http://www.aco.nato.int/saceur/channeling-david-sarnoff.aspx>.
- [2] Andrew Odlyzko Bob Briscoe and Benjamin Tilly. Metcalfe’s Law is Wrong. *IEEE Spectrum*, pages 26–31, July 2006.
- [3] Stephen P. Borgatti and Martin G. Everett. A Graph-Theoretic Perspective on Centrality. *Social Networks*, 28(4):466 – 484, 2006.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, March 2004.
- [5] S Chopra. On the Spanning Tree Polyhedron. *Operations Research Letters*, 8(1):25 – 29, 1989.
- [6] Marketing Conversation. Reeds Law States that Social Networks Scale Exponentially, August 2007. <http://marketingconversation.com/2007/08/28/reeds-law/>.
- [7] Marketing Conversation. A Short discussion on Metcalfe’s Law for Social Networks, May 2008. <http://marketingconversation.com/2007/08/28/reeds-law/>.
- [8] L. Freeman. Centrality in Social Networks Conceptual Clarification. *Social Networks*, 1(3):215–239, 1979.
- [9] D R Fulkerson. Blocking and Anti-Blocking Pairs of Polyhedra. *Math. Programming*, (1):168–194, 1971.
- [10] George Gilder. Metcale’s Law And Legacy, Nov 1995. <http://www.seas.upenn.edu/~gaj1/metgg.html>.
- [11] Assane Gueye. *A Game Theoretical Approach to Communication Security*. PhD dissertation, University of California, Berkeley, Electrical Engineering and Computer Sciences, March 2011.
- [12] Assane Gueye, Jean C. Walrand, and Venkat Anantharam. Design of Network Topology in an Adversarial Environment. In *GameSec 2010, Conference on Decision and Game Theory for Security*, pages 1–20. Springer-Verlag Berlin Heidelberg 2010, November 2010.
- [13] Assane Gueye, Jean C. Walrand, and Venkat Anantharam. How to Choose Communication Links in an Adversarial Environment. In *2nd International ICST Conference on Game Theory for Networks*, Shanghai, China, April 2011.
- [14] Victor Klee and George J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities*, volume III, pages 159–175. Academic Press, New York, 1972.
- [15] Mohammadhossein Manshaei, Quanyan Zhu, Tansu Alpcan, Tamer Basar, and Jean-Pierre Hubaux. Game Theory Meets Network Security and Privacy. Technical report, EPFL, Lausanne, 2010.
- [16] Deepankar Medhi. *Network Reliability and Fault-Tolerance*. John Wiley & Sons, Inc., 2007.
- [17] Andrew Odlyzko and Benjamin Tilly. A refutation of Metcalfe’s Law and a better estimate for the value of networks and network interconnections.

- [18] Tore Opsahl, Filip Agneessens, and John Skvoretz. Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths. *Social Networks*, 32(3):245 – 251, 2010.
- [19] Cisco Press. Spanning Tree Protocol: Introduction, Aug 2006. http://www.cisco.com/en/US/tech/tk389/tk621/tsd_technology_support_protocol_home.html.
- [20] Cisco Press. Understanding and Configuring Spanning Tree Protocol (STP) on Catalyst Switches, Aug 2006. http://www.cisco.com/en/US/tech/tk389/tk621/technologies_configuration_example09186a008009467c.shtml.
- [21] David P. Reed. That Sneaky Exponential: Beyond Metcalfe’s Law to the Power of Community Building, Spring 1999. <http://www.reed.com/dpr/locus/gfn/reedslaw.html>.
- [22] David P. Reed. Weapon of Math Destruction, Feb 2003. <http://www.immagic.com/eLibrary/ARCHIVES/GENERAL/GENREF/C030200D.pdf>.
- [23] Sankardas Roy, Charles Ellis, Sajjan Shiva, Dipankar Dasgupta, Vivek Shandilya, and Qishi Wu. A Survey of Game Theory as Applied to Network Security. *Hawaii International Conference on System Sciences*, 0:1–10, 2010.
- [24] Alexander Toshev. Submodular Function Minimization. Technical report, University of Pennsylvania, Philadelphia, 2010.
- [25] Robert J. Vanderbei. *Linear Programming: Foundations and Extensions*. Springer, second edition, 2001.
- [26] Douglas R. White and Stephen P. Borgatti. Betweenness centrality measures for directed graphs. *Social Networks*, 16(4):335 – 346, 1994.
- [27] Laurence A. Wolsey and George L. Nemhauser. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1 edition, November 1999.

A Game Model and NE Theorem

In this appendix section, we discuss a generalized model of the game described in section 3.1 and provide a proof of the Nash equilibrium theorem. The proof require the notion of blocking pairs of polyhedra that we revise next.

A.1 Blocking Pair of Matrices

The discussion in this section is mostly based on [27, pp. 99-101] and [9].

Let Λ be a $N \times m$ nonnegative matrix with non-zero rows. The polyhedron P_Λ associated with Λ is defined as the vector sum of the convex hull of its rows $(\lambda_1, \dots, \lambda_N)$ and the nonnegative orthant:

$$P_\Lambda = \text{conv.hull}(\lambda_1, \dots, \lambda_N) + \mathbb{R}_+^m. \quad (17)$$

A row λ_i of Λ is said to be *inessential* if it dominates a convex combination of other rows of Λ , otherwise we say that λ_i is *essential*. If all the rows of Λ are essential, we say that Λ is *proper*. The set of essential rows corresponds to the set of *extreme* points of the polyhedron P_Λ . Since inessential rows are not relevant for the definition of P_Λ we will drop them and assume that Λ is proper.




Figure 5: Example of polyhedron P_Λ defined by a nonnegative proper matrix Λ and its corresponding blocker $bl(P_\Lambda)$. The extreme points of the blocker define the nonnegative proper matrix Ω .

The example in Figure (5) shows the associated polyhedron P_Λ to the nonnegative matrix Λ given below:

$$\Lambda = \begin{pmatrix} 1/3 & 3 \\ 2/3 & 2 \\ 2 & 1 \\ 4 & 1/2 \end{pmatrix}, \quad \Omega = \begin{pmatrix} 0 & 2 \\ 1/6 & 2/3 \\ 3/10 & 2/5 \\ 2/3 & 10/36 \\ 4 & 0 \end{pmatrix} \quad (18)$$

Given Λ and its associated polyhedron, we define the blocker of the polyhedron P_Λ as follow.

Definition 3 The blocker $bl(P_\Lambda)$ of P_Λ is the polyhedron given as:

$$bl(P_\Lambda) = \{\mathbf{y} \in \mathbb{R}_+^m : \mathbf{y}'\mathbf{x} \geq 1, \forall \mathbf{x} \in P_\Lambda\}, \quad (19)$$

where $\mathbf{y}'\mathbf{x}$ is the inner product of $\mathbf{y}$ and $\mathbf{x}$. Recall that in this thesis we use the prime sign ($'$) for vector and matrix transpose.

We are interested in characterizing the polyhedron P_Λ and its blocker $bl(P_\Lambda)$. For that, we use the following theorem by Fulkerson [9]. It is based on the fact that there is a one-to-one correspondence between the essential rows of Λ and the extreme points of P_Λ .

Theorem 2 (Fulkerson, 1971) Let the N -by- m matrix Λ be proper with rows $\lambda_1, \dots, \lambda_N$, and let the polyhedron P_Λ be defined as in ((17)). Let $\omega_1, \dots, \omega_K$ be the extreme points of $bl(P_\Lambda)$, and let Ω be the matrix having those points as rows. Then,

1. The blocker $bl(P_\Lambda)$ of P_Λ is given by $bl(P_\Lambda) = \{\mathbf{x} \in \mathbb{R}_+^m : \Lambda\mathbf{x} \geq \mathbf{1}_N\}$.
2. Ω is proper, and the polyhedron P_Λ can be described as $P_\Lambda = \{\mathbf{x} \in \mathbb{R}_+^m : \Omega\mathbf{x} \geq \mathbf{1}_K\}$.
3. The blocker of the blocker $bl(P_\Lambda)$ verifies $bl(bl(P_\Lambda)) = P_\Lambda$.

Λ and Ω are said to form a blocking pair of matrices.

Equations ((18)) show a blocking pair of matrices Λ and Ω , and the corresponding polyhedra are shown in Figure (5).

Blocking pairs of matrices play an important role in the combinatorial problem of *maximum packing* (see Fulkerson[9]). In this thesis, we use the theory of blocking pairs to provide an easy argument for the existence of a probability distribution that satisfies a certain number of constraints.

For instance, consider the following linear program:

$$\begin{aligned} & \text{Maximize } \mathbf{1}'_{\mathcal{T}} \mathbf{x} \\ & \text{subject to } \Lambda' \mathbf{x} \leq \mathbf{w}, \text{ and } \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{20}$$

where the constraints Λ form a nonnegative matrix, and $\mathbf{w}$ is a given nonnegative vector.

We are interested to knowing whether the value of the program is greater than 1 or not. If this is the case, one can easily derive a probability distribution by normalizing a feasible solution of the program. Indeed, since the normalizing factor is greater than 1, the constraints will still be satisfied. The following lemma gives an answer to our question.

Lemma 1 *The value of the LP in ((20)) is greater than 1 if and only if $\mathbf{w}$ belongs to the polyhedron P_{Λ} defined by Λ .*

Proof: The proof of the lemma is as follow.

First, notice that strong duality holds for this LP. In fact, Slater's condition [4] is satisfied for any nonnegative and nonzero $\mathbf{w}$. The dual of the LP is given as:

$$\begin{aligned} & \text{Minimize } \mathbf{w}' \mathbf{y} \\ & \text{subject to } \Lambda \mathbf{y} \geq \mathbf{1}_{\mathcal{T}}, \text{ and } \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{21}$$

The constraints of the dual program ((21)) define the blocker

$$bl(P_{\Lambda}) = \{\mathbf{y} \in \mathbb{R}_+^m : \Lambda \mathbf{y} \geq \mathbf{1}_{\mathcal{T}}\}$$

of the polyhedron P_{Λ} . Now, if $\mathbf{w}$ belongs to P_{Λ} , then for all $\mathbf{y} \in bl(P_{\Lambda})$, we have that $\mathbf{w}' \mathbf{y} \geq 1$.

Conversely, if $\mathbf{w}' \mathbf{y} \geq 1$ for all $\mathbf{y} \in bl(P_{\Lambda})$, then $\mathbf{w}$ must be in the blocker of $bl(P_{\Lambda})$, which by Fulkerson's theorem (2), is P_{Λ} . This implies that the value of the dual program is greater than 1. Combined with the strong duality property, we get that the value of the primal program is at least 1. ■

A.2 Generalized Game Model

This section presents the blocking game model and introduces some notations that we will need to characterize the Nash equilibria of the game.

We consider that there is a nonempty, finite set S and two nonempty collections $\mathcal{T} = \{T_1, \dots, T_N\}$ and $\mathcal{E} = \{e_1, \dots, e_m\}$ of nonempty subsets of S . We call $\mathcal{E}$ the collection of resources. The defender selects a subset $T \in \mathcal{T}$ to perform a *mission critical* task. Each subset $T \in \mathcal{T}$ needs some set of resources $e_{T_1}, e_{T_2}, \dots, e_{T_p} \in \mathcal{E}$ in order to fulfill the task. To disrupt the mission, an attacker targets one resource $e \in \mathcal{E}$ to attack. Each resource $e \in \mathcal{E}$ has a cost of attack $\mu(e)$ that is the amount of effort that the attacker needs to spend to successfully launch the attack. The attacker also has the option of not attacking ("No Attack"); which we materialize by the choice of $e_{\emptyset}$. This choice results to zero loss for the defender and zero reward for the attacker.

Whenever the defender chooses subset T and resource e is attacked, he loses some value $\lambda(T, e)$. This loss goes to the attacker. It is conceivable that $\lambda(T, e) = 0$ if subset T does not need resource e . Hence, when the pair (T, e) is selected, the defender's loss is $\lambda(T, e)$ and the attacker's net reward is equal to $\lambda(T, e) - \mu(e)$.

This scenario can be modeled as a two-player matrix game where the players (the defender and the attacker) choose their pure strategies in the nonempty and finite sets $\mathcal{T}$ and $\mathcal{E} \cup \{e_\emptyset\}$, respectively (with $|\mathcal{E}| = m$ and $|\mathcal{T}| = N$).

The defender and attacker's respective payoff matrices are given by

$$\tilde{\Lambda} = [\Lambda|0], \text{ and } \tilde{\Pi} = [\Pi|0], \quad (22)$$

where Λ and Π are N -by- m matrices; $(\Lambda, [\Lambda](T, e) = \lambda(T, e))$ is a nonnegative matrix with no zero rows[§], and such that there is no column of Λ that dominates all other columns[¶]. Π is given by,

$$\Pi = \Lambda - \begin{pmatrix} \mu(1) & \mu(2) & \dots & \mu(m) \\ \mu(1) & \mu(2) & \dots & \mu(m) \\ \vdots & \vdots & \ddots & \vdots \\ \mu(1) & \mu(2) & \dots & \mu(m) \end{pmatrix}, \quad (23)$$

The “last” all-zero column in the definition of $\tilde{\Lambda}$ and $\tilde{\Gamma}$ captures the zero-loss for the defender ($\lambda_T(e_\emptyset) = 0$) and zero-reward for the attacker ($\mu(e_\emptyset) = 0$) when this latter chooses the “No Attack” strategy ($e_\emptyset$). For notational simplicity, we will only mention Λ and Γ .

We consider mixed strategies of this game where the defender chooses a distribution $(\alpha_T, T \in \mathcal{T})$ on $\mathcal{T}$ and the attacker chooses a distribution $(\beta(e), e \in \mathcal{E} \cup \{e_\emptyset\})$ on $\mathcal{E} \cup \{e_\emptyset\}$. The goal of the defender is to minimize the expected loss

$$L(\alpha, \beta) = \alpha' \tilde{\Lambda} [\beta; \beta(e_\emptyset)] = \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right), \quad (24)$$

while the attacker is trying to maximize the expected reward

$$R(\alpha, \beta) = \alpha' \tilde{\Pi} [\beta; \beta(e_\emptyset)] = \sum_{e \in \mathcal{E}} \beta(e) \left(\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) - \mu(e) \right). \quad (25)$$

In our notation, $[\beta; \beta(e_\emptyset)]$ is the column vector obtained by appending the additional row (entry) $\beta(e_\emptyset)$ to the column vector β .

Let P_Λ be the polyhedron associated with Λ given in ((22)), and let $bl(P_\Lambda)$ denote its blocker. From the discussion in the previous section and from Fulkerson's theorem, the blocker $bl(P_\Lambda) \subseteq \mathbb{R}_+^m$ is the polyhedron associated with the nonnegative matrix Ω whose rows are the vertices of $bl(P_\Lambda)$.

It is also known that $bl(P_\Lambda)$ is the vector sum of the convex hull of rows of Ω with the positive orthant $\mathbb{R}_+^m$, and that its blocking polyhedron is P_Λ (see [9] and [27, pp. 99-101]). Also, Theorem (2) gives that

$$P_\Lambda = \{\mathbf{x} \in \mathbb{R}_+^m, \text{ s.t. } \Omega \mathbf{x} \geq \mathbf{1}_\mathcal{T}\} \quad (26)$$

Now, for ω row of Ω ^{||}, which we denote as $\omega \in \Omega$, we write $\omega = (\omega(e), e \in \mathcal{E})$, and let $\omega(\mathcal{E}) := \sum_{e \in \mathcal{E}} \omega(e)$. Note that $\omega(e) \geq 0$ for all $e \in \mathcal{E}$ and $\omega(\mathcal{E}) > 0^{**}$; so that $(\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$ is a probability distribution on $\mathcal{E}$. We call it the probability distribution induced by ω .

[§]A row with all zeros would lead to a trivial game because it means that there exists a subset $T \in \mathcal{T}$ that will always result to zero loss. The defender would then always select such strategy and the game ends.

[¶]This would also lead to a trivial game if the attack cost is not too high. In fact, it means that there is a resource that will always give higher attack gain, independently of the subset chosen by the defender. The attacker would always target such a resource.

^{||}Notice that this is an abuse of language because ω is a column vector.

^{**}This is because the blocker $bl(P_\Lambda)$ is not empty (Λ is not a one-rowed zero matrix), and does not contain the all-zero vector (the origin) (this could not give an inner product with rows of Λ that is greater than 1).

Also define, for each $\omega \in \Omega$, consider the quantity

$$\frac{\min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \omega(e) \lambda_T(e) \right)}{\omega(\mathcal{E})}; \quad (27)$$

This quantity is the minimum loss seen by the defender if the attacker were to choose a target according to the distribution $(\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$ induced by ω . A closer look at the expression above shows that the numerator is equal to 1. In fact, we already know that if ω belongs to the blocker of P_Λ , $\sum_{e \in \mathcal{E}} \omega(e) \lambda_T(e) \geq 1$ for all $T \in \mathcal{T}$. Now, if ω is a vertex of the blocker of P_Λ , one can easily show that

Lemma 2 *There must exist a some T_o such that $\sum_{e \in \mathcal{E}} \omega(e) \lambda_{T_o}(e) = 1$.*

Proof: Assume that the assertion is not true and $\sum_{e \in \mathcal{E}} \omega(e) \lambda_T(e) > 1$ for all T . Let $\tau = \min_T (\sum_{e \in \mathcal{E}} \omega(e) \lambda_T(e))$ and let $\omega_o = \frac{\omega}{\tau}$. Then, $\omega_o \in \Omega$ because $\sum_{e \in \mathcal{E}} \omega_o(e) \lambda_T(e) \geq 1$ for all T . But, ω strictly dominates ω_o which contradicts the fact that ω is a vertex of the blocker of P_Λ . ■ Thus, we can drop the minimization and define the following function on the vertices ω of the blocker $bl(P_\Lambda)$

$$\kappa : \Omega \longrightarrow \mathbb{R} \quad (28)$$

$$\omega \longmapsto \kappa(\omega) := \omega(\mathcal{E}) = \sum_{e \in \mathcal{E}} \omega(e). \quad (29)$$

Finally, let us define $\theta(\omega)$ as

$$\theta(\omega) := \frac{1}{\kappa(\omega)} \left(1 - \sum_{e \in \mathcal{E}} \omega(e) \mu(e) \right). \quad (30)$$

The expression of $\theta(\omega)$ is composed with two terms. If the attacker were to choose a resource to attack according to the distribution $\omega = (\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$, then the first term would have been the loss seen by the defender, which, as we have assumed, goes to the attacker. The second term is the average cost of attack corresponding to the distribution induced by ω . Hence, $\theta(\omega)$ can be seen as the expected attack reward associated ω .

We call the vertex ω of the blocking polyhedron a *critical* vertex if

$$\theta(\omega) = \max_{\tilde{\omega} \in \Omega} \theta(\tilde{\omega}). \quad (31)$$

A critical vertex is one whose induced probability distribution gives a maximum rewards to the attacker (considering the defender's response).

We define $\theta := \max_{\tilde{\omega} \in \Omega} \theta(\tilde{\omega})$ to be the maximum achievable value in the preceding expression, and we let Ω_{max} denote the matrix having as rows the critical vertices of $bl(P_\Lambda)$.

We use θ as the *vulnerability* for the defender's task.

A.3 Nash Equilibrium Theorem

This section presents the main results of the two-player matrix game defined in section (A.2). We claim that:

Theorem 3 *For the game defined above, the following always hold.*

1. If $\theta \leq 0$, then “No Attack” (i.e. $\beta(e_\emptyset) = 1$) is always an optimal strategy for the attacker. In this case,

the equilibrium strategy $(\alpha_T, T \in \mathcal{T})$ for the defender is such that

$$\bar{\lambda}\alpha(e) := \sum_{T \in \mathcal{T}} \alpha_T \lambda_T(e) \leq \mu(e), \quad \forall e \in \mathcal{E}. \quad (32)$$

The corresponding payoff is 0 for both players.

2. If $\theta \geq 0$, then for every probability distribution $(\gamma_{\omega}, \omega \in \Omega_{max})$, the attacker's strategy $(\beta(e), e \in \mathcal{E})$ defined by

$$\beta(e) = \sum_{\omega \in \Omega_{max}} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E})} \quad (33)$$

is in Nash equilibrium with any strategy $(\alpha_T, T \in \mathcal{T})$ of the defender that satisfies the following properties:

$$\begin{cases} \bar{\lambda}\alpha(e) - \mu(e) = \theta & \text{for all } e \in \mathcal{E} \text{ such that } \beta(e) > 0. \\ \bar{\lambda}\alpha(e) - \mu(e) \leq \theta & \text{for all } e \in \mathcal{E}. \end{cases} \quad (34)$$

Further, there exists at least one such strategy α .

The corresponding payoffs are θ for the attacker, and $r(\gamma)$ for the defender, where

$$r(\gamma) := \sum_{\omega \in \Omega_{max}} \frac{\gamma_{\omega}}{\kappa(\omega)}. \quad (35)$$

3. If $\mu = 0$, then every Nash equilibrium pair of strategies for the game is of this type.

Note: $\theta = 0$ is a particular case where both cases (1) and (2) can occur. In all cases, the maximum achievable attack reward is equal to 0. There exist equilibria where the attacker decides to not attack. In those cases, the defender has to choose α according to ((32)). There might also exist equilibria where the attack launches an attack but gets a expected reward of 0. In such equilibrium, α has to satisfy ((34)). In our experiments we did not find any equilibrium where the attacker would mix between $e_{\emptyset}$ and some resources in $\mathcal{E}$.

Proof: The proof of the theorem is presented in Appendix D ■




Figure 6: Examples of disconnecting sets (DS) and minimal disconnecting sets (MDS). The chosen subset is shown in dashed line. The leftmost example is not a DS. The next example is a DS but is not a MDS because the red link connect two nodes belonging to the same connected component. The third example is a MDS. It is also a minimum edge-cut of the graph (the one with the minimum size). The rightmost example is a MDS but not a minimum edge-cut.

B Network Value Function $\kappa(\cdot)$

B.1 Minimum Disconnecting Subset (MDS)

We start by the following definition.

Definition 4 Let $\mathcal{E}$ be the set of edges of the graph G , and let $E \subseteq \mathcal{E}$ be a subset of edges.

1. E is said to be a disconnecting set (DS) of G , if removing the edges on E disconnects the graph (and results into 2 or more connected components).
2. E is said to be a minimal disconnecting subset (MDS) if E is a disconnecting set such that, for every edge $e \in E$, adding e to $G \setminus E$ (the graph obtained by removing, from G , the edges in E) decreases its number of connected components by 1.
3. We denote by $\hat{E}$ the MDS of maximum cardinality that is included in E .

Figure (6) shows examples of DS's and MDS's. The first example (leftmost) is not a DS. The second example is a DS but is not a MDS. The third example is a MDS. It is also the minimum cut of the graph. However, in general MDS's are not minimum cuts. The last example shows a MDS that is not a minimum cut.

The notion of MDS is equivalent to the notion of *feasible partition* presented by Chopra [5]. Next, we recall the definition of a feasible partition.

Definition 5 A feasible partition Π of the vertices of the graph G ($V = V_1 \cup V_2 \cup \dots \cup V_{|\Pi|}$, $i \neq j \Rightarrow V_i \cap V_j = \emptyset$) is such that every partition member V_i is a connected component (or equivalently that two vertices u and v are in the same partition member V_i if and only if there exists a path p_{uv} from u to v that uses only edges of the subgraph $(V_i, V_i(\mathcal{E}))$).

In [11], we have shown the following lemma.

Lemma 3 1. For every feasible partition Π of the vertices of the graph G , the set $\mathcal{E}(\Pi)$ of edges that go between vertices in distinct elements of Π is a MDS.
 2. Every MDS E is the set of edges going across the elements of some feasible partition (that we will denote Π_E).

In [5], Chopra has shown that the spanning tree polyhedron can be characterized by assigning weights to edges going across members of feasible partitions (this characterization was previously given by Fulkerson,

Chopra provided a proof in the cited paper). Here, we show that the blocker of the payoff matrix defined above (via spanning trees and value functions) can be characterized by solving some linear program on the MDS's.

First, we give two important properties of MDS's.

Property 1 *If $E \subseteq F \in DS(\mathcal{E})$, then $\widehat{E} \subseteq \widehat{F}$.*

Proof: If $e \in \widehat{E}$, then e must be a bridge in the graph $\{G \setminus E\} \cup \{e\}$. Hence, e has to be a bridge in $\{G \setminus F\} \cup \{e\}$, which implies that $e \in \widehat{F}$.

Better Proof: If $e \in \widehat{E}$, then it must connect two different members of the partition Π_E . But since $E \subseteq F$, the partition Π_F further partitions the members of Π_E and hence e connects two different members of Π_F , implying that $e \in \widehat{F}$. ■

Property 2 *Let e be an edge in $\mathcal{E} \setminus E$. Then, if e is not a bridge on the subgraph $G \setminus E$, we have that $\widehat{E} = \widehat{E + e}$.*

Recall that a bridge is an edge $e \in \mathcal{E}$ the removal of which increases the number of connected components of the graph.

Proof: Let Π_E be the partition obtained by removing the edges of E from the graph and considering that u and v belongs to the same class if there exists a path p_{uv} from u to v . If $e = (u, v)$, the vertices u and v belong to the same class of the partition Π_E . Call it V_{uv} . Furthermore, if e is not a bridge in $G \setminus E$, then there exists another path p_{uv} that does not contain the vertex $e = (u, v)$ and that only uses edges that belong to that same class V_{uv} . Thus, removing e from $G \setminus E$ does not disconnected the subgraph $(V_{uv}, V_{uv}(\mathcal{E}))$ and hence does not change the partition. Also, the same edges that go across the classes of partition Π_E are those going across the classes of partition Π_{E+e} . Thus, $\widehat{E} = \widehat{E + e}$. ■

Next we define a function that is related to the polyhedron P_Λ , where for any nonnegative matrix Λ , P_Λ is defined as follow

$$P_\Lambda = \text{conv.hull}(\lambda_1, \dots, \lambda_N) + \mathbb{R}_+^m. \quad (36)$$

In the definition above, λ_i , $i = 1 \dots N$ are the rows of the matrix Λ .

Definition 6 *Let $\mathcal{E} = \{1, 2, \dots, n\}$ be the set of indices of the column of the tree-link payoff matrix Λ defined above, and let $DS(\mathcal{E}) \subseteq 2^\mathcal{E}$ be the set of all disconnecting sets of the graph. We define the function $\kappa(\cdot)$ on $DS(\mathcal{E})$ as follow:*

$$\kappa : DS(\mathcal{E}) \longrightarrow \mathbb{R} \quad (37)$$

$$E \longmapsto \kappa(E) = \max\{\mathbf{1}^T \mathbf{x}, \Lambda_{\widehat{E}}^T \mathbf{x} \leq \mathbf{1}\}, \quad (38)$$

where $\widehat{E}$ is the MDS of maximum size that is contained in E and Λ_F is the matrix that contains the columns indexed by $F \subseteq \mathcal{E}$ and zeros in any other column.

This later definition (compared to the one using the dual LP) generalizes to all subsets of $\mathcal{E}$ and does not require to focus on MDS's (I have to show that the maximum is attained at a MDS (actually $\widehat{E}$)). With this definition, if E is not a disconnecting set, then the maximum is equal to $+\infty$. We set an upper bound K to all numbers (can be set as large as possible but is fixed ones it is chosen) and define $\kappa(\emptyset) = K$.

As an example, if Λ is the spanning tree-link incidence matrix of a connecting graph, for any minimal disconnecting set E , $\kappa(E) = \frac{|E|}{Q(G \setminus E) - 1}$, where $Q(G \setminus E)$ is the number of connected components of the graph $G \setminus E$. This can be verified by Fulkerson's theorem or the theorem stating that the solution of certain integer

program are integral (fractional) (give some references). Related to this (and in anticipation to the characterization of the blocker), the vertices of the blocker of spanning tree polyhedron correspond to the essential (i.e. that do not dominate a convex combination of the others) vectors of the family defined as follow:
for any MDS E , construct vector $\mathbf{y}_E$ by assigning weight $\frac{1}{Q(G \setminus E) - 1}$ to each link in E and zero to any other link.

The function $\kappa(\cdot)$ has a certain number of properties that we list below. In the sequel we will apply, whenever it is required, the appropriate extension on non-disconnecting sets.

Property 3 For any $E \subseteq \mathcal{E}$, $\kappa(E) \geq 0$.

This property trivially follows from the definition.

Property 4 $\kappa(\cdot)$ is none-increasing on the set of disconnecting sets $DS(\mathcal{E})$; i.e. if $E \subseteq F \in DS(\mathcal{E})$, then $\kappa(E) \geq \kappa(F)$.

Proof: The reason is that if a nonnegative vector $\mathbf{y}$ satisfies the constraints (44) for $\widehat{E} \subseteq \widehat{F}$ (by property (1)), it must satisfy them for $\widehat{F}$ as well. Indeed, since both Λ and $\mathbf{y}$ are nonnegative, adding the columns indexed by $\widehat{F} \setminus \widehat{E}$ to $\Lambda_{\widehat{E}}$ can only increase the values in the LHS's of the constraints, hence making the constraints even looser. ■

Property 5 $\kappa(\cdot)$ is submodular, i.e. it satisfies the following

$$\kappa(E) - \kappa(E + e_1) \leq \kappa(E + e_2) - \kappa(E + e_1 + e_2), \quad \forall E \subseteq \mathcal{E}, \quad e_1, e_2 \in \mathcal{E} \setminus E, \quad (39)$$

where in the above and throughout this paper, we use the notation $E + e$ as a shortcut for $E \cup \{e\}$.

In other words, $\kappa(\cdot)$ has increasing returns on the set of edges of the graph. Again, here we have assumed the extended definition on $2^{\mathcal{E}}$.

Proof: To show this, we use the correspondence between MDS's and subset of edges going across members of vertex partitions.

First assume that e_1 is not a bridge of the graph $G \setminus \{E + e_2\}$ or in other terms removing e_1 , does not increase the number of connected components of the graph $G \setminus \{E + e_2\}$. Then, from Lemma (2), we have that $\widehat{E + e_2} = \widehat{E + e_1 + e_2}$. As a consequence, $\kappa(E + e_2) = \kappa(E + e_1 + e_2)$. Thus, since $\kappa(\cdot)$ is none-increasing, we only need to show that $\kappa(E) = \kappa(E + e_1)$. But, this is indeed the case because if e_1 is not a bridge in $G \setminus \{E + e_2\}$, it cannot be a bridge in $G \setminus E$ (we remove less edges in $G \setminus E$ than in $G \setminus \{E + e_2\}$). Thus, $\widehat{E} = \widehat{E + e_1}$ implying that $\kappa(E) = \kappa(E + e_1)$.

Now assume that e_1 is a bridge in $G \setminus (E + e_2)$ (notice that we always have that $\kappa(E + e_2) \geq \kappa(E + e_2 + e_1)$). In this case one needs to consider three cases: e_2 connects two vertices inside the same connected component of the graph $G \setminus \{E + e_1 + e_2\}$ (figure 7.a)), e_1 and e_2 go across the same connected components of $G \setminus \{E + e_1 + e_2\}$ (figure 7.b)) (they must be the only edges across those two connected components), or e_1 and e_2 go across different connected components of $G \setminus \{E + e_1 + e_2\}$ (figures 7.c) and 7.d)),

- If e_2 connects two vertices inside the same connected component of the graph $G \setminus \{E + e_1 + e_2\}$, then e_2 is neither a bridge of the graph $G \setminus E$ nor one of $G \setminus \{E + e_1\}$. These implies that $\widehat{E} = \widehat{E + e_2}$ and $\widehat{E + e_1} = \widehat{E + e_1 + e_2}$. As a consequence, $\kappa(E) = \kappa(E + e_2)$ and $\kappa(E + e_1) = \kappa(E + e_1 + e_2)$. Thus the RHS of (39) is equal to its LHS.




Figure 7: Illustrative proof of the submodularity: possible scenarios when e_1 is a bridge in $G \setminus (E + e_2)$.




Figure 8: Illustrative proof of the submodularity: possible scenarios when e_1 is a bridge in $G \setminus (E + e_2)$.

- If e_1 and e_2 go across the same connected components of $G \setminus \{E + e_1, e_2\}$ say V_1 and V_2 (they must be the only edges across those two connected components), then $\widehat{E} = \widehat{E} + e_1$ because $G \setminus E$ is “almost” the same as $G \setminus \{E + e_1 + e_2\}$ except that in $G \setminus E$ the classes V_1 and V_2 are connected by e_1 and e_2 . Thus, removing e_1 from $G \setminus E$ will let V_1 and V_2 connected (by e_2). Hence the partitions Π_E and Π_{E+e_1} have the same classes and thus $\widehat{E} = \widehat{E} + e_1$. This implies that $\kappa(E) = \kappa(E + e_1)$, and $\kappa(E) - \kappa(E + e_1) = 0 \leq \kappa(E + e_2) - \kappa(E + e_1 + e_2)$.
- If e_1 and e_2 go across different connected components of $G \setminus \{E + e_1, e_2\}$, then $\widehat{E} \neq \widehat{E} + e_1$, $\widehat{E} \neq \widehat{E} + e_2$, and $\widehat{E} + e_1 = \widehat{E} + e_1 + e_2$. To show that the LHS of (39) is smaller than its RHS, we consider the changes in the value of the LP in the definition of $\kappa(\cdot)$ (38) when we go from E to $E + e_1$ to $E + e_1 + e_2$ and from E to $E + e_2$ to $E + e_1 + e_2$. Notice that since the definition of $\kappa(\cdot)$ is a maximization, going from E to $E + e_1$ introduces a new constraint and hence decreases the value of the LP from $\kappa(E)$ to $\kappa(E + e_1)$. Let the difference be $d_1 = \kappa(E) - \kappa(E + e_1)$. Similarly, we define $d_2 = \kappa(E) - \kappa(E + e_2)$ as the difference of going from E to $E + e_2$, $d_{1 \rightarrow 12} = \kappa(E + e_1) - \kappa(E + e_1 + e_2)$ as the difference of going from $E + e_1$ to $E + e_1 + e_2$, and finally $d_{2 \rightarrow 21} = \kappa(E + e_2) - \kappa(E + e_1 + e_2)$ as the difference of going from $E + e_2$ to $E + e_1 + e_2$. Figure 8 gives an illustration of this process of going from one LP solution to another. Notice that the value of the final LP p_{12}^* can be written as

$$\kappa(E + e_1 + e_2) = \kappa(E) - (d_1 + d_{1 \rightarrow 12}) = \kappa(E) - (d_2 + d_{2 \rightarrow 21}).$$

Also, we have that

$$\kappa(E + e_1) = \kappa(E) - d_1, \quad \text{and} \quad \kappa(E + e_2) = \kappa(E) - d_2$$

Replacing these definition in the inequality (39) gives us

$$\begin{aligned} \kappa(E) - \kappa(E + e_1) &\leq \kappa(E + e_2) - \kappa(E + e_1 + e_2) \\ &\Leftrightarrow \end{aligned} \quad (40)$$

$$\begin{aligned} \kappa(E) - (\kappa(E) - d_1) &\leq (\kappa(E) - d_2) - (\kappa(E) - (d_2 + d_{2 \rightarrow 21})) \\ &\Leftrightarrow \end{aligned} \quad (41)$$

$$d_1 \leq d_{2 \rightarrow 21} \quad (42)$$

Thus, we only need to show that $d_1 \leq d_{2 \rightarrow 21}$. To show this, we consider the dual of the LP presented in the definition of the function $\kappa(\cdot)$ (see (38)).

$$\kappa : DS(\mathcal{E}) \longrightarrow \mathbb{R} \quad (43)$$

$$E \longmapsto \kappa(E) = \min\{\mathbf{1}^T \mathbf{y}, \Lambda_E \mathbf{y} \geq \mathbf{1}\}, \quad (44)$$

Now, notice that going from E to $E + e_1$ introduce a new variable $\mathbf{y}_{e_1}$ to the optimization (44). This adds one degree of freedom and hence the solution of the LP defined with E can be improved to get that of the LP defined by $E + e_1$. The net improvement be computed by using the dual simplex algorithm. The algorithm starts from the current solution (with $\mathbf{y}_{e_1} = 0$) (notice that the current solution is feasible for the new LP) and makes a series of pivots to find the new optimal solution. Since the dual is a minimization, this new optimal solution is a feasible vector (for the new LP) that maximizes the improvement d_1 obtained by the introduction of the new variable $\mathbf{y}_{e_1}$ (in LP terms, we say that $\mathbf{y}_{e_1}$ has become a basic variable). Again, the constraints to the maximization of the improvement is to maintain feasibility ($\widehat{A_{E+e_1}} \mathbf{y} \leq \mathbf{1}$). Making the same reasoning for going from $E + e_2$ to $E + e_1 + e_2$, we have that $d_{2 \rightarrow 12}$ is the maximum improvement to the solution of the LP defined by $E + e_2$ (with $\mathbf{y}_{e_1} = 0$) with the constraints that ($\widehat{A_{E+e_1+e_2}} \mathbf{y} \leq \mathbf{1}$). Now, notice that any solution of the LP defined by E is a solution to the LP defined by $E + e_2$ (with $\mathbf{y}_{e_2} = 0$) and any solution of the LP defined by $E + e_1$ is a solution to the LP defined by $E + e_1 + e_2$. Hence the solution set of the maximization to obtain d_1 is a subset of the solution set of the maximization to obtain $d_{2 \rightarrow 12}$. But, $\max\{f(\mathbf{x}), \mathbf{x} \in S\}$ is an increasing function of the solution set S . As a consequence, $d_1 \leq d_{2 \rightarrow 12}$, which gives the inequality in (42). ■

Next, we discuss one nice property of the normalized tree-link (payoff) matrix introduced in sections 2.2.

Property 6 *If the graph $G = (\mathcal{V}, \mathcal{E})$ is simple (does not contains parallel or loop links of) and 2-vertex connected, then each column $\Lambda[:, e]$ of the payoff matrix contains all possible values of $\lambda(T, e)$.*

In other terms, the matrix Λ is *balanced* in the sense that there is not an entry that largely dominates (or is dominated by) the other entries. Other way to state the lemma is that for each edge $e \in \mathcal{E}$ and for each partition (N_1, N_2) of N (i.e. $N_1 + N_2 = N$), there exists a spanning tree T containing e such that removing e disconnects T into 2 subtrees; one containing N_1 nodes and the other N_2 nodes. We denote that property by $T - \{e\} = T_{N_1} \cup T_{N_2}$.

If the graph contains a bridge 9.a, then the column corresponding to the bridge contains only one value. If the graph contains a cut vertex (as in figure 9.b), then there might exist edges that are not “balanced”. In figure 9.b the dashed edge cannot separate any spanning tree into 2 subtrees with respectively 2 and 5 nodes. Figure 9.c is a simple graph that is not 2-vertex connected but on which the property holds.

Proof: We show the lemma by induction on the size of the graph (n, m) (number of vertices, number of edges).




Figure 9: Examples of simple graphs that are not 2-vertex connected.

The property is trivial for the simple and 2-vertex connected graph of size $(2, 1)$ (the only not one value of $\lambda(T, e)$ is 0). Also, it can easily be checked for the simple and 2-vertex connected graph of size $(3, 3)$. Again, the only possible value of $\lambda(T, e)$ is $1 - \frac{f(1)+f(2)}{f(3)}$ and each column of Λ has 2 entries with this value and one entry with a 0.

Now suppose that the property is satisfied for all graphs of size small than the size (n, m) of some arbitrary graph G . Then let's show that it must be satisfied for G .

First recall that if graph is 2-vertex connected, every edge is either *contractable* or *deletable* (one or the other), i.e. if the edge is contracted (else/or deleted), then the resulting graph (after removing parallel and loop links) is another 2-vertex connected graph.

Now consider an arbitrary edge $e(u, v)$ of the graph G . We would like to show that for any pair (n_1, n_2) such that $n_1 + n_2 = n$, there exist a spanning tree T containing f such that removing f separates T into 2 subtrees one with n_1 nodes and the other with n_2 nodes (i.e. $T - \{e\} = T_{n_1} \cup T_{n_2}$). Let f be a vertex incident to e . If f is deletable, then we are done. We delete f and apply the induction hypothesis to the new graph $G' = G - \{f\}$ of size $(n, m - 1)$. Any spanning tree on that graph is a spanning tree in G . If f is not deletable, then it must be contractable. For each pair (n_1, n_2) , we will show the property by exhibiting a spanning tree that e divides into subtrees of respectively n_1 and n_2 nodes, i.e. $T - \{e\} = T_{n_1} \cup T_{n_2}$. We do so by contracting e (and removing all loops and parallel links). The contraction leads to a smaller graph G' of size $(n - 1, m')$, $m' < m$, for which the induction hypothesis applies. Hence, there exists a spanning tree T' of G' containing e such that $T' - \{e\} = T'_1 \cup T'_{n-2}$. Without loss of generality we let $T'_1 = \{u\}$. Two scenarios are possible.

- If f is incident to v (see figure 10.a) and 10.ab), then uncontracting f transform T'_{n-2} to another subtree $T'_{n-1} = T'_{n-2} + \{f\}$ that contains $n - 1$ nodes. Furthermore, $T = T' + \{f\}$ is a spanning tree of the original graph G and $T - \{e\} = T' + \{f\} - \{e\} = T'_1 \cup T'_{n-1}$.
- If f is incident to u (see figure 10.c)), then $f = (u, u')$ for some node u' that must have a neighbor v' among the nodes of T'_{n-2} (because the graph is 2-vertex connected). Let $f' = (u', v')$. Then, $T = T' + \{f'\}$ is a spanning tree of the origin graph G and $T - \{e\} = T' + \{f'\} - \{e\} = T'_1 \cup (T'_{n-2} + \{f'\})$.

For any other pair $(n_1, n_2) > (1, 1)$ we can find a spanning tree T of G such $T - \{e\} = T_{n_1} \cup T_{n_2}$ as follow. By the induction hypothesis, there exists a spanning tree T' of G' such that $T' - \{e\} = T'_{n_1-1} \cup T'_{n_2}$ (see figures 11.a) and 11.b)). If f is incident to a node in T'_{n_1-1} otherwise we can find T' such that $T' - \{e\} = T'_{n_1} \cup T'_{n_2-1}$. Uncontracting f gives a spanning tree $T = T' + \{f\}$ such that $T - \{e\} = T' + \{f\} - \{e\} = (T'_{n_1-1} + \{f\}) \cup T'_{n_2}$ (if f is incident to a node in T'_{n_1-1} otherwise $T - \{e\} = T' + \{f\} - \{e\} = T'_{n_1} \cup (T'_{n_2-1} + \{f\})$). This ends the proof of the lemma. ■

B.2 Blocker of the payoff matrix

For any non negative vector $\mathbf{v} \in \mathbb{R}_+^m$, we let $\chi_{\mathbf{v}}$ be the entries of $\mathbf{v}$ that have positive value.

Conjecture

Vertices of the blocker of the payoff matrix are vectors ω such that χ_{ω} is a MDS. Each such vertex is




Figure 10: Illustrative figures for the proof of Lemma 6: building a spanning tree T such that $T - \{e\} = T_1 \cup T_{n-1}$.




Figure 11: Illustrative figures for the proof of Lemma 6: building a spanning tree T such that $T - \{e\} = T_{n_1} \cup T_{n_2}$.

obtained by solving an LP of the form

$$\begin{aligned}
& \text{Minimize } \omega^T \mathbf{y} \\
& \text{subject to } \Lambda_{\hat{E}} \mathbf{y} \geq \mathbf{1} \\
& \mathbf{y} \geq \mathbf{0},
\end{aligned} \tag{45}$$

where the weights ω depend on the value models.

C Algorithm to compute a critical subset

In this section we present the algorithm to compute a critical subset of a critical subset of a graph.

D Proof of the NE theorem

In this section we provide a proof of the Nash equilibrium theorem presented in section ??.

D.1 “No Attack” Option

We start by showing that if $\theta \leq 0$ “No Attack” and α verifying (32) are best responses to each other.

D.1.1 Best Responses

First, notice that if the attacker chooses “No Attack”, then any α will result to the minimum loss of zero for the defender (in particular the one given in the theorem). Now, assume that α satisfies (32). Then, “No Attack” is a *dominant* strategy for the attacker. In fact, the expected attack reward is

$$R(\alpha, \beta) = \sum_{e \in E} \beta(e) (\bar{\lambda}\alpha(e) - \mu(e)), \quad (46)$$

which is less than zero if $\bar{\lambda}\alpha(e) - \mu(e) \leq 0$, $\forall e \in \mathcal{E}$, and for any β . On the other hand, zero reward can always be achieved by playing $e_\emptyset$. As a consequence, not attacking is a best response to α satisfying (32).

D.1.2 Existence of the Equilibrium Distribution

Now, we need to prove the existence of a distribution α that satisfies (32) whenever $\theta \leq 0$. To summarize, we are looking for α verifying:

$$\alpha : \begin{cases} \alpha \geq 0 \\ \mathbf{1}'_{\mathcal{T}} \alpha = 1 \\ \lambda' \alpha \leq \mu, \end{cases} \quad (47)$$

We first show the following lemma.

Lemma 4

$$\theta \leq 0 \Rightarrow \mu \in P_{\lambda}. \quad (48)$$

Proof: For this we show that if $\theta \leq 0$, then $\mu' \omega \geq 1$, for any vertex ω of $bl(P_{\lambda})$. As a result of this, μ belongs to the blocker of $bl(P_{\lambda})$ which is P_{λ} .

If $\theta \leq 0$, then we have that, for all $\omega \in \Omega$,

$$\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) \geq \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right). \quad (49)$$

Or, equivalently

$$\mu' \omega \geq \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \omega(e) \lambda(T, e) \right), \text{ for all } \omega \in \Omega. \quad (50)$$

Now, since $\omega \in bl(P_{\lambda})$ we have that

$$\sum_{e \in \mathcal{E}} \omega(e) \lambda(T, e) \geq 1, \quad (51)$$

which implies that $\mu' \omega \geq 1$, $\forall \omega \in \Omega$, or equivalently $\mu \in P_{\lambda}$.

Using Lemmas 1 and 4, we conclude that the value of the following LP is greater than 1.

$$\begin{aligned} & \text{Maximize } \mathbf{1}'_{\mathcal{T}} \mathbf{x} \\ & \text{subject to } \lambda' \mathbf{x} \leq \mu, \text{ and } \mathbf{x} \geq \mathbf{0} \end{aligned} \quad (52)$$

Construct α satisfying (32) by normalizing any solution of this LP. ■

D.2 The “Always Attack” option

As in the previous section, we will first argue that the strategies given in the theorem are best responses to each other, then we show the existence of a distribution $(\alpha_T, T \in \mathcal{T})$ that satisfies (34). We start by the following lemma.

Lemma 5 *If $\theta > 0$, then “No Attack” is a strictly dominated strategy for the attacker.*

Note that if $\theta \geq 0$, all the steps of the proof still hold. However, “No Attack” will only be weakly dominated, and as seen in the previous case, there will exist equilibrium for which the attacker will opt to not launch an attack. In other words, if $\theta = 0$ there exist equilibria for which $\beta_{e_\emptyset} = 1$, as well as equilibria for which $\beta_{e_\emptyset} = 0$.

Proof: Suppose that $\omega \in \Omega_{max}$ is a critical vertex of $bl(P_\lambda)$ and let $\beta = (\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$. We will show that the attacker can achieve positive reward by playing β (independently of α). To see this, first notice that since $\theta(\omega) = \theta > 0$ for all $\omega \in \Omega_{max}$, we have that

$$\lambda(\omega) = \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) > \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) = \frac{\mu' \omega}{\omega(\mathcal{E})} \quad (53)$$

Playing the strategy β , the attacker’s expected reward against any defense strategy α is given by

$$\begin{aligned} R(\alpha, \beta) &= \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \left(\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) - \mu(e) \right) \\ &= \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) - \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) \right) \end{aligned} \quad (54)$$

$$\geq \sum_{T \in \mathcal{T}} \alpha_T \left(\lambda(\omega) - \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) \right) \quad (55)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T \left(\lambda(\omega) - \frac{\mu' \omega}{\omega(\mathcal{E})} \right) \quad (56)$$

$$= \lambda(\omega) - \frac{\mu' \omega}{\omega(\mathcal{E})} \quad (57)$$

$$> 0, \quad (58)$$

where in (55) we use the definition of $\lambda(\omega)$; (58) is implied by (53). ■

As a consequence of this lemma, we conclude that if $\theta > 0$, then the attacker will never play the “No Attack” (i.e. $e_\emptyset$) strategy.

D.2.1 Best Responses

Given the set of critical vertices Ω_{max} and α satisfying (34), any distribution β of the form $\beta(e) = \sum_{\omega \in \Omega_{max}} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E})}$ for some distribution $\gamma = (\gamma_\omega, \omega \in \Omega_{max})$, achieves a reward of θ . This is the maximum possible reward that the attacker can get. To see this, observe that for any β ,

$$R(\alpha, \beta) = \sum_{e \in \mathcal{E}} \beta(e) (\bar{\lambda} \alpha(e) - \mu(e)) \leq \sum_{e \in \mathcal{E}} \beta(e) \theta \leq \theta. \quad (59)$$

The upper bound of θ is achieved by any $\tilde{\beta} = (\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$ and $\omega \in \Omega_{max}$, because for any such $\tilde{\beta}$,

$$R(\alpha, \tilde{\beta}) = \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \left(\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) - \mu(e) \right) \quad (60)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) - \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) \right) \quad (61)$$

$$\geq \sum_{T \in \mathcal{T}} \alpha_T \left(\lambda(\omega) - \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \mu(e) \right) \quad (62)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T \theta \quad (63)$$

$$= \theta, \quad (64)$$

where (62) uses the definition of $\lambda(\omega)$, and in (63) we use the fact that $\omega \in \Omega_{max}$.

As a consequence, any distribution of the form $(\frac{\omega(e)}{\omega(\mathcal{E})}, e \in \mathcal{E})$ for $\omega \in \Omega_{max}$ is a best response and any convex combination of those distributions is also a best response.

Now assume that β is given as in (33) for some distribution $(\gamma\omega, \omega \in \Omega_{max})$. Then, the distribution $(\alpha_T, T \in \mathcal{T})$ in (34) achieves a loss of $r(\gamma) = \sum_{\omega \in \Omega_{max}} \gamma\omega \lambda(\omega)$. This is the minimum possible loss. To see this, note that, for any α , the expected loss seen by the defender is given by

$$L(\alpha, \beta) = \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right) \quad (65)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega_{max}} \gamma\omega \frac{\omega(e)}{\omega(\mathcal{E})} \right) \lambda(T, e) \right) \quad (66)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{\omega \in \Omega_{max}} \gamma\omega \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (67)$$

$$\geq \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{\omega \in \Omega_{max}} \gamma\omega \lambda(\omega) \right) \quad (68)$$

$$= \sum_{T \in \mathcal{T}} \alpha_T r(\gamma) \quad (69)$$

$$= r(\gamma). \quad (70)$$

The lower bound $r(\gamma)$ can be achieved by choosing α such that $\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) = \theta + \mu(e)$ for each $e \in \mathcal{E}$ such that $\beta(e) > 0$ (the existence of such α is shown in the second part of the theorem). This can be seen by rewriting $L(\alpha, \beta)$ as

$$L(\alpha, \beta) = \sum_{e \in \mathcal{E}} \beta(e) \left(\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) \right) \quad (71)$$

$$= \sum_{e \in \mathcal{E}} \beta(e) (\theta + \mu(e)) \quad (72)$$

$$= \theta + \beta' \mu = r(\gamma). \quad (73)$$

The last equality above is justified by

$$\theta + \beta' \mu = \theta + \sum_{\omega \in \Omega_{max}} \gamma \omega \left(\frac{1}{\omega(\mathcal{E})} \omega' \mu \right) - r(\gamma) + r(\gamma) \quad (74)$$

$$= \theta + \sum_{\omega \in \Omega_{max}} \gamma \omega \left(\frac{1}{\omega(\mathcal{E})} \omega' \mu - \lambda(\omega) \right) + r(\gamma) \quad (75)$$

$$= \theta + \sum_{\omega \in \Omega_{max}} \gamma \omega (-\theta) + r(\gamma) \quad (76)$$

$$= \theta - \theta + r(\gamma) \quad (77)$$

$$= r(\gamma), \quad (78)$$

where in (74) and (75) we use the definitions of β and $r(\gamma)$ respectively and write the summation over $\mathcal{E}$ as a product of a row vector and a column vector. In 75 we have also used the definition of $\theta(\omega)$ for a critical vertex $\omega \in \Omega_{max}$.

From this analysis, we see that for the set Ω_{max} of critical vertices, the distributions given in the theorem are best responses to each other, as a consequence, they form a Nash equilibria under the assumption that α exists. Such existence is shown in the next section.

D.2.2 Existence of the Equilibrium Distribution

We claim that for any $(\beta(e), e \in \mathcal{E})$ of the form in the statement (33) of Theorem 3, we can find an associated $(\alpha_T, T \in \mathcal{T})$ of the form (34).

Theorem 4 *Assume that $\theta \geq 0$ and let Ω_{max} be the set of critical vertices. Let $\mathbf{x}^*$ be the solution of the following LP:*

$$\begin{aligned} & \text{Maximize } \mathbf{1}'_{\mathcal{T}} \mathbf{x} \\ & \text{subject to } A' \mathbf{x} \leq \mathbf{b}, \quad \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (79)$$

where $\mathbf{b} = \theta(E) \mathbf{1}_{\mathcal{E}} + \mu$. Then,

$$a) \mathbf{1}'_{\mathcal{T}} \mathbf{x}^* \leq 1;$$

$$b) \mathbf{1}'_{\mathcal{T}} \mathbf{x}^* \geq 1;$$

$$c) A' \mathbf{x}^*(e) = \mathbf{b}(e), \forall e \in \mathcal{E} \text{ for which } \beta(e) > 0.$$

As a consequence, $\mathbf{x}^*$ satisfies (34) and implies the existence of α .

Proof: a) To prove that $\mathbf{1}'_{\mathcal{T}} \mathbf{x}^* \leq 1$, we first observe that

$$\beta' \lambda' \mathbf{x} = \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right) \quad (80)$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega_{max}} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E})} \right) \lambda(T, e) \right) \quad (81)$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{\omega \in \Omega_{max}} \gamma \omega \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (82)$$

$$\geq \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{\omega \in \Omega_{max}} \gamma \omega \lambda(\omega) \right) \quad (83)$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T r(\gamma) \quad (84)$$

$$= r(\gamma) \mathbf{1}'_{\mathcal{T}} \mathbf{x} \quad (85)$$

On the other hand, from the constraints $\lambda' \mathbf{x} \leq \mathbf{b} = \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$ and from (78), we have that

$$\beta' \lambda' \mathbf{x} \leq \beta' (\theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}) = \theta + \beta' \boldsymbol{\mu} = r(\gamma) \quad (86)$$

Combining (85) and (78) it follows that,

$$r(\gamma) \mathbf{1}'_{\mathcal{T}} \mathbf{x} \leq \beta' \lambda' \mathbf{x} \leq r(\gamma) \quad (87)$$

Thus $\mathbf{1}'_{\mathcal{T}} \mathbf{x} \leq 1$ for all feasible $\mathbf{x}$, i.e. the value of the program is at most 1.

b) To prove that $\mathbf{1}'_{\mathcal{T}} \mathbf{x}^* \geq 1$, we use Lemma 1 above to claim that it suffices to verify that the vector $\mathbf{b}$ belongs to the polyhedron P_{λ} . For that, we will show that $\mathbf{b}' \omega \geq 1$ for all $\omega \in \Omega$. In fact*,

$$\omega' \mathbf{b} = \theta \omega' \mathbf{1}_{\mathcal{E}} + \omega' \boldsymbol{\mu} \quad (88)$$

$$= \omega(\mathcal{E}) \left(\theta + \frac{1}{\omega(\mathcal{E})} \omega' \boldsymbol{\mu} \right) \quad (89)$$

$$\geq \omega(\mathcal{E}) \left(\lambda(\omega) - \frac{1}{\omega(\mathcal{E})} \omega' \boldsymbol{\mu} + \frac{1}{\omega(\mathcal{E})} \omega' \boldsymbol{\mu} \right) \quad (90)$$

$$= \omega(\mathcal{E}) \lambda(\omega) \quad (91)$$

$$= \omega(\mathcal{E}) \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \quad (92)$$

$$= \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \omega(e) \lambda(T, e) \right) \quad (93)$$

$$\geq 1 \quad (94)$$

where (94) follows from the fact that $\omega \in bl(P_{\lambda})$. Indeed, if $\omega \in bl(P_{\lambda})$, then by definition of the blocker polyhedron, $\sum_{e \in \mathcal{E}} \omega(e) \lambda(T, e) \geq 1$ for all $T \in \mathcal{T}$. Thus, we have that $\Omega \mathbf{b} \geq 1$, which implies that $\mathbf{b}$ belongs to the blocker of $bl(P_{\lambda})$ which is equal to P_{λ} .

*Again we use vector product for summation.

Now, using Lemma 1, we conclude that the value of the LP is greater than 1.

This, together with the previous part a) imply that the value of the LP is equal to 1; hence, any solution $\mathbf{x}^*$ is a probability distribution on $\mathcal{T}$.

c) We first observe from (85) and (78) that

$$\beta' \lambda' \mathbf{x}^* = r(\gamma) \mathbf{1}_{\mathcal{T}} \mathbf{x}^* = r(\gamma). \quad (95)$$

Also, $\lambda' \mathbf{x}^* \leq \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$ by the constraints of the primal LP above.

Now, assume that $\lambda' \mathbf{x}^*(e) < \theta + \mu(e)$ for some $e \in \mathcal{E}$ with $\beta(e) > 0$. Then,

$$\beta' \lambda' \mathbf{x}^* = \sum_{e \in \mathcal{E}} \beta(e) \lambda' \mathbf{x}^*(e) \quad (96)$$

$$< \sum_{e \in \mathcal{E}} \beta(e) (\theta + \mu(e)) \quad (97)$$

$$= \theta + \sum_{e \in \mathcal{E}} \beta(e) \mu(e) \quad (98)$$

$$= r(\gamma), \quad (99)$$

where the last equality is obtained by using the same arguments as in (74)-(78). This contradicts observation (95). As a consequence, $\lambda' \mathbf{x}^*(e) = \theta + \mu(e)$ for all $e \in \mathcal{E}$ with $\beta(e) > 0$.

This ends the proof of the theorem and establishes the existence of an α satisfying (34) for any β defined as in (33). $\blacksquare$

D.3 Enumerating all Nash Equilibria

In this section, we consider the zero-sum game where $\boldsymbol{\mu} = 0$. In this case, since there is no cost of attack $\theta > 0$. We will show that for any strategy pair $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ that are in Nash equilibrium, it must be the case that $(\beta(e), e \in \mathcal{E})$ is given by

$$\beta(e) = \sum_{\omega \in \Omega_{max}} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E})}, \quad (100)$$

for some probability distribution $(\gamma_{\omega}, \omega \in \Omega_{max})$.

As a consequence of this, we will conclude that α must be in the form given in the Nash equilibrium theorem.

First, notice that since $\boldsymbol{\mu} = 0$, we have that $\lambda(\omega) = \theta(\omega)$ (see the definitions in (30) and (27)).

Next, we use the zero-sum structure of the game to observe that it has a well-defined value, which, by the second part of Theorem 3, is equal to,

$$\theta = \max_{\omega \in \Omega} (\lambda(\omega)) = \max_{\omega \in \Omega} \left(\min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right). \quad (101)$$

Thus, we must have that, for any Nash equilibrium pair (α, β) ,

$$\theta = \sum_{T \in \mathcal{T}} \sum_{e \in \mathcal{E}} \alpha_T \beta(e) \lambda(T, e) = \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right) > 0. \quad (102)$$

From the above equation we can argue that there exists a scaling factor $\tilde{\kappa} > 0$ such that $(\tilde{\kappa} \beta(e), e \in \mathcal{E})$

belongs to the blocker $bl(P_{\lambda})$, or equivalently,

$$\sum_{e \in \mathcal{E}} \tilde{\kappa} \beta(e) \lambda(T, e) \geq 1, \quad \text{for all } T \in \mathcal{T}. \quad (103)$$

In fact, by letting α_{max} be the maximum α_T , and defining $\tilde{\kappa} = \frac{N\alpha_{max}}{\theta}$, then $(\tilde{\kappa}\beta(e), e \in \mathcal{E})$ verifies (103). We let κ denote the smallest such scaling that works among all scalings $\tilde{\kappa} > 0$.

Also, observe that since κ is the smallest nonnegative scaling of $(\beta(e), e \in \mathcal{E})$ such that $(\kappa\beta(e), e \in \mathcal{E})$ belongs to $bl(P_{\lambda})$, there must exist some $T_o \in \mathcal{T}$ for which $\sum_{e \in \mathcal{E}} \kappa\beta(e) \lambda_{T_o, e} = 1$.

Indeed this is the case because since $\kappa\beta \in bl(P_{\lambda})$, we have that $\sum_{e \in \mathcal{E}} \kappa\beta(e) \lambda(T, e) \geq 1$ for all $T \in \mathcal{T}$. If this inequality were strict for all $T \in \mathcal{T}$, then considering the strategy T_{min} that minimizes the sum $\sum_{e \in \mathcal{E}} \kappa\beta(e) \lambda(T, e)$ over all T , we can construct

$$\tilde{\kappa} = \frac{\kappa}{\sum_{e \in \mathcal{E}} \beta(e) \lambda_{T_{min}, e}}. \quad (104)$$

$\tilde{\kappa}$ verifies $\tilde{\kappa} < \kappa$ and $\tilde{\kappa}\beta \in bl(P_{\lambda})$. This contradicts the assumption that κ was the smallest such $\tilde{\kappa}$.

Now, we claim that:

Lemma 6 *If $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ form a NE of the game, then we can write*

$$\kappa\beta(e) = \sum_{\omega \in \Omega} \gamma\omega \frac{\omega(e)}{\omega(\mathcal{E})\lambda(\omega)}, \quad (105)$$

for some probability distribution $(\gamma\omega, \omega \in \Omega)$.

We delay the proof of this lemma for later.

Using this expression of $\kappa\beta$, we can write the value of the game θ as:

$$\theta = \frac{1}{\kappa} \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \kappa\beta(e) \lambda(T, e) \right) \quad (106)$$

$$= \frac{1}{\kappa} \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega} \gamma\omega \frac{\omega(e)}{\omega(\mathcal{E})\lambda(\omega)} \right) \lambda(T, e) \right) \quad (107)$$

$$= \frac{1}{\kappa} \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{\omega \in \Omega} \gamma\omega \left(\frac{1}{\lambda(\omega)} \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (108)$$

$$\geq \frac{1}{\kappa} \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{\omega \in \Omega} \gamma\omega \right) \quad (109)$$

$$= \frac{1}{\kappa}, \quad (110)$$

where in (109) we use the fact that $\lambda(\omega) \leq \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e)$ for all $T \in \mathcal{T}$.

Now, since (α, β) is a Nash equilibrium pair, the expression $\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e)$ is minimal for each $T \in \mathcal{T}$ for

which $\alpha_T > 0$. Furthermore, this minimum value is equal to θ . From this, we get that,

$$\theta = \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right) \quad (111)$$

$$= \frac{1}{\kappa} \min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \kappa \beta(e) \lambda(T, e) \right) \quad (112)$$

$$\leq \frac{1}{\kappa}, \quad (113)$$

where in (113) we use the fact that the minimum in (112) is less than $\sum_{e \in \mathcal{E}} \kappa \beta(e) \lambda_{T_o, e}$, which, by definition, is equal to 1. Thus, $\theta = \frac{1}{\kappa}$.

This, combined with (105) that we sum over $e \in \mathcal{E}$, imply:

$$\frac{1}{\theta} = \kappa = \sum_{e \in \mathcal{E}} \kappa \beta(e) \quad (114)$$

$$= \sum_{e \in \mathcal{E}} \sum_{\omega \in \Omega} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)} \quad (115)$$

$$= \sum_{\omega \in \Omega} \gamma_{\omega} \frac{1}{\omega(\mathcal{E}) \lambda(\omega)} \sum_{e \in \mathcal{E}} \omega(e) \quad (116)$$

$$= \sum_{\omega \in \Omega} \gamma_{\omega} \frac{1}{\lambda(\omega)}. \quad (117)$$

Now, recalling (101) that $\theta = \max_{\omega \in \Omega} (\lambda(\omega))$, we conclude that γ_{ω} can be nonzero only for $\omega \in \Omega$ that satisfies $\lambda(\omega) = \max_{\tilde{\omega} \in \Omega} (\lambda(\tilde{\omega}))$. In other terms, $\gamma_{\omega} > 0$ only for $\omega \in \Omega_{max}$.

Hence, we can write

$$\beta(e) = \sum_{\omega \in \Omega_{max}} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E})}. \quad (118)$$

The last thing that remains to be shown to complete the proof of the theorem is that if $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ are in Nash equilibrium and $(\beta(e), e \in \mathcal{E})$ is of the form (33) in the statement Theorem 3, then $(\alpha_T, T \in \mathcal{T})$ must also be of the form in the statement (34) of the theorem. We have already shown that for $(\beta(e), e \in \mathcal{E})$ of the form (33), we must have for *every* strategy $(\tilde{\alpha}_T, T \in \mathcal{T})$

$$\sum_{T \in \mathcal{T}} \tilde{\alpha}_T \sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) = \sum_{T \in \mathcal{T}} \tilde{\alpha}_T \left(\sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega_{max}} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (119)$$

$$= \sum_{T \in \mathcal{T}} \tilde{\alpha}_T \left(\sum_{\omega \in \Omega_{max}} \gamma_{\omega} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (120)$$

$$\geq \sum_{T \in \mathcal{T}} \tilde{\alpha}_T \left(\sum_{\omega \in \Omega_{max}} \gamma_{\omega} \theta \right) \quad (121)$$

$$= \theta, \quad (122)$$

where (121) follows from (101). The minimum value of θ can be achieved by choosing α such that

$\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) = \theta$ whenever $\beta(e) > 0$. To see that, rewrite the summation as

$$\sum_{T \in \mathcal{T}} \tilde{\alpha}_T \sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) = \sum_{e \in \mathcal{E}} \beta(e) \left(\sum_{T \in \mathcal{T}} \tilde{\alpha}_T \lambda(T, e) \right), \quad (123)$$

and observe the claim.

The existence of such α has been shown in the previous section. It also has been shown that such α is a best response to $(\beta(e), e \in \mathcal{E})$. For β to be a best response to α (hence, the (α, β) pair to be in Nash equilibrium), α must also satisfy

$$\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) \leq \theta \text{ for all } e \in \mathcal{E}. \quad (124)$$

Suppose, on the contrary, that this is not the case (i.e. there is some $e \in \mathcal{E}$ for which $\sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e) > \theta$). Then the attacker will prefer to switch to playing strategy e with probability 1 and receive higher reward. This violates the assumption that $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ are in Nash equilibrium. Thus, α satisfies (124). This completes the proof of the theorem, provided that the claim in Lemma 6 can be justified. We give a proof of the lemma in the next section.

D.3.1 Proof of Lemma 6

The claim is that if $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ are in Nash equilibrium, and if $\kappa > 0$ denotes the smallest $\tilde{\kappa} > 0$ for which $(\tilde{\kappa}\beta(e), e \in \mathcal{E}) \in bl(P_\lambda)$, then we must have

$$\kappa\beta(e) = \sum_{\omega \in \Omega} \gamma_\omega \frac{\omega(e)}{\omega(\mathcal{E})\lambda(\omega)}, \quad (125)$$

for some probability distribution $(\gamma_\omega, \omega \in \Omega)$.

Indeed, this needs a proof, because a priori we only know that we can write

$$\kappa\beta(e) = \sum_{\omega \in \Omega} \gamma_\omega \frac{\omega(e)}{\omega(\mathcal{E})\lambda(\omega)} + v(e), \quad (126)$$

for some probability distribution $(\gamma_\omega, \omega \in \Omega)$ and some $(v(e), e \in \mathcal{E})$ such that $v(e) \geq 0$ for all $e \in \mathcal{E}$.

We now provide the proof that works as follow. We consider the expression of this form for $(\kappa\beta(e), e \in \mathcal{E})$ for which $v(\mathcal{E}) := \sum_{e \in \mathcal{E}} v(e)$ is as small as possible. We will assume that $v(\mathcal{E}) > 0$ for this expression, and arrive at a contradiction.

Note that we must have $v(\mathcal{E}) < \kappa$ for this expression, i.e. there has to be a nontrivial “convex hull part”.

First observe that for all $T \in \mathcal{T}$, we have

$$\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) = \frac{1}{\kappa} \sum_{e \in \mathcal{E}} \kappa \beta(e) \lambda(T, e) \quad (127)$$

$$= \frac{1}{\kappa} \sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)} + v(e) \right) \lambda(T, e) \quad (128)$$

$$= \frac{1}{\kappa} \sum_{\omega \in \Omega} \gamma \omega \left(\frac{1}{\lambda(\omega)} \sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) + \frac{1}{\kappa} \sum_{e \in \mathcal{E}} v(e) \lambda(T, e) \quad (129)$$

$$\geq \frac{1}{\kappa} \sum_{\omega \in \Omega} \gamma \omega + \frac{1}{\kappa} v \lambda(T) \quad (130)$$

$$= \frac{1}{\kappa} + \frac{v \lambda(T)}{\kappa}, \quad (131)$$

where in (130) we use the definition of $\lambda(\omega)$, and define $v \lambda(T) := \sum_{e \in \mathcal{E}} v(e) \lambda(T, e)$. As a consequence, we have that

$$\min_{T \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) \right) \geq \frac{1}{\kappa} + \min_{T \in \mathcal{T}} \left(\frac{v \lambda(T)}{\kappa} \right) \quad (132)$$

Next, observe that since $(\alpha_T, T \in \mathcal{T})$ and $(\beta(e), e \in \mathcal{E})$ are in Nash equilibrium, it must be the case that $\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e)$ is the same for every $T \in \mathcal{T}$ such that $\alpha_T > 0$. Also, by the same reasoning as in (113), we have that for all such T

$$\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e) = \min_{\tilde{T} \in \mathcal{T}} \left(\sum_{e \in \mathcal{E}} \beta(e) \lambda_{\tilde{T}, e} \right) \leq \frac{1}{\kappa}. \quad (133)$$

This, combined with (124), implies that $\min_{T \in \mathcal{T}} (\sum_{e \in \mathcal{E}} \beta(e) \lambda(T, e)) = \frac{1}{\kappa}$ and that $v \lambda(T) = 0$ for all $T \in \mathcal{T}$ such that $\alpha_T > 0$.

Now, let

$$\tilde{\beta}(e) := \frac{1}{\kappa - v(\mathcal{E})} \sum_{\omega \in \Omega} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)}. \quad (134)$$

This quantity satisfies $\sum(e) \tilde{\beta}(e) = 1$ and $\tilde{\beta}(e) \geq 0$ for all $e \in \mathcal{E}$. Thus, $(\tilde{\beta}(e), e \in \mathcal{E})$ is a probability distribution on $\mathcal{E}$, and can be used as a strategy by the attacker. This can be verified by summing both sides of (126) over $e \in \mathcal{E}$ to get,

$$\kappa = \sum_{e \in \mathcal{E}} \kappa \beta(e) = \sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)} + v(e) \right) \quad (135)$$

$$= \sum_{e \in \mathcal{E}} \left(\sum_{\omega \in \Omega} \gamma \omega \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)} \right) + v(\mathcal{E}) \quad (136)$$

$$= (\kappa - v(\mathcal{E})) \sum_{e \in \mathcal{E}} \tilde{\beta}(e) + v(\mathcal{E}). \quad (137)$$

This last equation implies that

$$\sum_{e \in \mathcal{E}} \tilde{\beta}(e) = \frac{\kappa - v(\mathcal{E})}{\kappa - v(\mathcal{E})} = 1. \quad (138)$$

For this strategy $(\tilde{\beta})$, in response to $(\alpha_T, T \in \mathcal{T})$, the attack reward is at least $\frac{1}{\kappa - v(\mathcal{E})}$. In fact,

$$\sum_{T \in \mathcal{T}} \alpha_T \sum_{e \in \mathcal{E}} \tilde{\beta}(e) \lambda(T, e) = \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{e \in \mathcal{E}} \left(\frac{1}{\kappa - v(\mathcal{E})} \sum_{\omega \in \Omega} \gamma_{\omega} \frac{\omega(e)}{\omega(\mathcal{E}) \lambda(\omega)} \right) \lambda(T, e) \right) \quad (139)$$

$$= \frac{1}{\kappa - v(\mathcal{E})} \sum_{T \in \mathcal{T}} \alpha_T \left(\sum_{\omega \in \Omega} \gamma_{\omega} \frac{1}{\lambda(\omega)} \left(\sum_{e \in \mathcal{E}} \frac{\omega(e)}{\omega(\mathcal{E})} \lambda(T, e) \right) \right) \quad (140)$$

$$\geq \frac{1}{\kappa - v(\mathcal{E})} \sum_{T \in \mathcal{T}} \alpha_T \sum_{\omega \in \Omega} \gamma_{\omega} \quad (141)$$

$$= \frac{1}{\kappa - v(\mathcal{E})}. \quad (142)$$

However, because of the fact that $v(T) = 0$ for all $T \in \mathcal{T}$ for which $\alpha_T > 0$, the benefit obtained by the attacker by playing the NE strategy $(\beta(e), e \in \mathcal{E})$ is only $\frac{1}{\kappa}$, which is strictly smaller than $\frac{1}{\kappa - v(\mathcal{E})}$ under the standing assumption that $v(\mathcal{E}) > 0$. As a consequence, if $v(\mathcal{E}) > 0$, the attacker can be better off by changing her strategy to $\tilde{\beta}$. But this contradicts the assumption that (α, β) form a NE. Thus, $v(\mathcal{E}) = 0$ implying that $v(e) = 0$ for all $e \in \mathcal{E}$ as we wanted to show.