How to Choose Communication Links in an Adversarial Environment?

Assane Gueye, Jean C. Walrand, and Venkat Anantharam

University of California at Berkeley, EECS Department, Berkeley CA 94720, USA
{agueye,wlr,ananth}@eecs.berkeley.edu

Abstract. Given the topology of a network, characterized by an undirected graph, we consider the following game situation: a network manager is choosing (as communication infrastructure) a spanning tree of the graph, and an attacker is trying to disrupt the communication tree by attacking one link of the network. Attacking a link has a certain cost for the attacker who also has the option of not attacking. We model the interaction between the network manager and the attacker as a bimatrix game and study the set of mixed strategy Nash equilibria. We define the notion of critical subset of links and determine the structure of a particular set of Nash equilibria when the attack cost is nonzero. In each NE of this set, the attacker targets edges in critical subsets and all edges in the same critical subset are attacked with the same probability. For the game of zero cost of attack considered in [8], we characterize the set of all Nash equilibria. Some implications of the results are discussed and a detailed proof of the NE theorem is provided.

Keywords: Network Topology, Connectivity, Graph Vulnerability, Spanning Trees, Minimum Cut Set, Game Theory, Nash Equilibrium, Linear Programming, Blocking pairs of polyhedra.

1 Introduction

In [8], we have studied the *strategic* interaction between a network manager whose goal is to choose a spanning tree of a network as communication infrastructure, and an attacker who tries to disrupt the communication tree by attacking one link in the network. Therein, we assumed that the cost of attack is equal to zero for the attacker and we discussed the notions of vulnerability and critical subset of links. We have also shown that there always exists a Nash equilibrium under which the attacker targets uniformly, at random, links in a critical set.

In the present paper, we generalize our results to the case where the attacker incurs a positive cost by attacking a given link of the graph. We revisit the notions of *vulnerability* and *criticality* of a subset of links and show that the critical subset attack theorem in [8] generalizes to the present case. We determine a particular set of Nash equilibria for the game of positive attack cost, and for the game with zero cost of attack, we characterize the set of all NE. We also provide a unifying proof of the Nash equilibrium theorem that applies both to the game presented in [8], and the one studied here.

This paper is organized as follows. We present the model and the environment of the game in Section 2. The notion of critical subset is discussed in section 3 followed by the

main result of this paper in section 4. We provide a proof of the main theorem in section 6. The proof requires the notion of blocking pair of matrices. Appendix B gives a brief introduction to this notion and presents a lemma that unifies the proofs of the theorems in [8] and in this paper. The implications of the result as well as illustrative examples are presented in section 5. This paper ends with concluding remarks discussed in section 7.

2   Model

The network topology is given by a connected undirected graph $G = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{E}| = m$ links and $|\mathcal{V}| = n$ nodes. Let $\mathcal{T}$ be the set of spanning trees, and let $N = |\mathcal{T}|$. Each edge $e \in \mathcal{E}$ is associated with some cost $\boldsymbol{\mu}(e)$ that an attacker needs to spend to successfully attack that link. Each tree T has the same cost for the network manager that we assume to be equal to 1.

To get all nodes connected in a cycle-free way, the network manager chooses a spanning tree $T \in \mathcal{T}$ of the graph. The attacker simultaneously chooses a link $e \in \mathcal{E}$ to attack. The attacker wins if the attacked link belongs to the chosen spanning tree, otherwise the network wins. More precisely, for a choice pair (T, e) of tree and edge, the attack loss is $L(T, e) = \mathbf{1}_{e \in T}$ for the network, while the net attack reward is equal to $R(T, e) = \mathbf{1}_{e \in T} - \boldsymbol{\mu}(e)$ for the attacker.

The manager picks a spanning tree according to a chosen distribution $\boldsymbol{\alpha}$ on $\mathcal{T}$ to minimize the expected attack loss. Similarly, the attacker chooses a link according to some distribution $\boldsymbol{\beta}$ on $\mathcal{E}$ to maximize the expected attack reward. We assume 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 formulate this interaction as a one-shot 2-player game between the network manager and the attacker. Their respective pure strategy sets are the set $\mathcal{T}$ of spanning trees and the set $\mathcal{E}$ of edges of the graph. We are interested in analyzing mixed strategy Nash equilibria of this game.

Let $\mathcal{A} := \{\boldsymbol{\alpha} \in \Re_+^N \mid \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T = 1\}$ be the set of mixed strategies for the network manager, and $\mathcal{B} := \{\boldsymbol{\beta} \in \Re_+^m \mid \sum_{e \in \mathcal{E}} \boldsymbol{\beta}_e = 1\}$ the set of mixed strategies for the attacker. Define A as the loss matrix for the manager, with $A_{T,e} = \mathbf{1}_{e \in T}$ and B is the reward matrix of the attacker with $B_{T,e} = \mathbf{1}_{e \in T} - \boldsymbol{\mu}(e)$. The average expected loss $L(\alpha, \beta)$ for the manager, and reward $R(\alpha, \beta)$ for the attacker are given by

$$L(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \boldsymbol{\alpha}' A \boldsymbol{\beta} = \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\sum_{e \in T} \boldsymbol{\beta}_e \right), \tag{1}$$

$$R(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \boldsymbol{\alpha}' B \boldsymbol{\beta} = \sum_{e \in \mathcal{E}} \boldsymbol{\beta}_e \left(\sum_{T \ni e} \boldsymbol{\alpha}_T - \boldsymbol{\mu}(e) \right). \tag{2}$$

3   Critical Subsets

In this section we define the notions of vulnerability and critical subset of a graph.

Definition 1. *For any nonempty subset of links $E \subseteq \mathcal{E}$, define*

$$\mathcal{M}(E) := \min_{T \in \mathcal{T}} | T \cap E |, \quad and \quad \theta(E) := \frac{\mathcal{M}(E) - \mu(E)}{|E|}. \tag{3}$$

$\theta(E)$ is called the vulnerability *of E. It is the minimum fraction of edges that E has in common with any tree minus the average cost of attacking E. Here and throughout the paper, we use the notation $\mu(E) = \sum_{e \in E} \mu(e)$.*
A nonempty subset of edges E is said to be critical *if it has maximum vulnerability:*
$\theta(E) = max_{E' \subseteq \mathcal{E}} \{\theta(E')\}$. We let $\mathcal{C}$ denote the set of all critical subsets.
The vulnerability of the graph θ is defined to be equal to the vulnerability of its critical subset(s).

The examples shown in Figure 1 illustrate the definitions presented above. The network in Figure 1(a) has a vector of attack cost $\mu = [0.5, 0.5, 0.5, 2, 0.5, 0.5, 0.5]$. It contains a bridge that has a relatively high cost of attack ($\mu(4) = 2$). As a consequence it is not critical. There are two critical subsets $E_1 = \{1, 2, 3\}$ and $E_2 = \{5, 6, 7\}$ shown respectively by the dashed and dash-dotted lines. This example illustrates the impact of the attack cost. When a link is too costly to attack, it becomes less critical.

Figures 1(b) and 1(c) show the same network topology with different costs of attack. In the first one, the attack costs are $\mu = [5, 3, 3, 5, 5, 4, 3, 3, 5, 5, 4, 5, 5, 3]/14$. For these values of the costs of attack, the minimum cutset of the graph (links 6 and 8) is critical. If the attack costs are equal to $\mu = [2, 5, 1, 2, 1, 1, 6, 5, 3, 7, 1, 4, 3, 6]/21$ (second case), the minimum cutset is no longer critical. It has vulnerability $\theta(6, 8) = \frac{1 - (4+3)/14}{2} = 1/4$. One critical subset of the graph is given by the set $E = \{1, 2, 3, 4, 5, 6, 7, 8, \}$. Its vulnerability is $\theta(E) = 0.3631$.


(a) Graph with bridge   (b) Network where mini-   (c) Minimum cutset is not
                         mum cutset is critical   critical

Fig. 1. Examples illustrating definition 1. The vector of attack costs are $\mu = [0.5, 0.5, 0.5, 2, 0.5, 0.5, 0.5]$ for Figure 1(a), $\mu = [5, 3, 3, 5, 5, 4, 3, 3, 5, 5, 4, 5, 5, 3]/14$ for Figure 1(b), and $\mu = [2, 5, 1, 2, 1, 1, 6, 5, 3, 7, 1, 4, 3, 6]/21$ for Figure 1(c). Each set of dashed (or dash-dotted) lines is a critical subset.

4   Critical Subset Attack Theorem

Now, we give the critical subset attack theorem for the game defined above, a proof of which is provided in section 6.

Theorem 1 (Critical Subset Attack Theorem). *For the game defined in section 2 with attack costs μ, the following always holds.*

1. If $\theta = \max_{E \subseteq \mathcal{E}} \left(\frac{\mathcal{M}(E) - \boldsymbol{\mu}(E)}{|E|} \right) \leq 0$, then there is a NE under which the attacker will opt to not launch an attack. The equilibrium strategy $(\boldsymbol{\alpha}_T, \ T \in \mathcal{T})$ for the defender is such that

$$\boldsymbol{\alpha}(e) := \sum_{T \ni e} \boldsymbol{\alpha}_T \leq \boldsymbol{\mu}(e), \quad \forall e \in \mathcal{E}. \tag{4}$$

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 $\mathcal{C}$ of critical subsets, the attacker's strategy $(\boldsymbol{\beta}(e), e \in \mathcal{E})$ defined by

$$\boldsymbol{\beta}(e) = \sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|}, \tag{5}$$

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

$$\begin{cases} \boldsymbol{\alpha}(e) - \boldsymbol{\mu}(e) = \theta \ \text{for all} \ e \in \mathcal{E} \ \text{such that} \ \boldsymbol{\beta}(e) > 0. \\ \boldsymbol{\alpha}(e) - \boldsymbol{\mu}(e) \leq \theta \ \text{for all} \ e \in \mathcal{E}. \end{cases} \tag{6}$$

Furthermore, there exists at least one such strategy $\boldsymbol{\alpha}$.

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

$$r(\gamma) := \sum_{E \in \mathcal{C}} \gamma_E \frac{\mathcal{M}(E)}{|E|}. \tag{7}$$

3. If $\boldsymbol{\mu} = 0$, then every Nash equilibrium pair of strategies for the game is of the type in (2.) above.

5   Analyzing the NE Theorem

We discuss the NE theorem by considering a game on the graph shown in Figure 2. Table 1 shows the parameters and results of the game. The first column shows different values of the attack costs $\boldsymbol{\mu}$ and the second column shows the corresponding critical subset(s). The third column displays the vulnerability of the graph. For each vector of attack costs, we compute the Nash equilibria of the game. The next two columns of the table show the Nash equilibrium strategies, respectively $\boldsymbol{\alpha}$ for the network manager, and $\boldsymbol{\beta}$ for the attacker. The equilibrium payoffs are displayed in the last column. In all equilibria, we have chosen the distribution γ_E to only focus on a particular critical subset (the ones shown on the table). Note that we have not shown all Nash equilibria.


Fig. 2. Example of graph and its spanning trees. The left figure is the original graph with the 5 edges labeled with their number. The right figures are the 8 spanning trees of the graph also labeled with their numbers.

Table 1. Game with positive attack cost played for different values of the cost of attack μ

Attack Cost μ	Critical Set E_c	Vulnerability θ	Nash Equilibria α	β	Payoffs $(\lambda_{min}, \mu_{max})$
$[3\ 3\ 5\ 3\ 4]/5$	(1,4)	-0.1	$[0\ 0\ 2\ 0\ 0\ 1\ 0\ 2]/5$	No Attack	$(0,0)$
			$[0\ 0\ 2\ 1\ 0\ 0\ 0\ 2]/5$	No Attack	$(0,0)$
$[5\ 2\ 3\ 2\ 3]/5$	(1-5)&(2-4)	0	$[0\ 2\ 1\ 2\ 0\ 0\ 0\ 0]/5$	$[1\ 1\ 1\ 1\ 1]/5$	$(0.6,0)$
			$[0\ 2\ 1\ 2\ 0\ 0\ 0\ 0]/5$	$[0\ 1\ 1\ 1\ 1]/4$	$(0.5,0)$
			$[0\ 2\ 1\ 2\ 0\ 0\ 0\ 0]/5$	No Attack	$(0,0)$
$[5\ 4\ 2\ 4\ 2]/8$	(3,5)	0.25	$[0\ 1\ 0\ 0\ 1\ 2\ 0\ 0]/4$	$[0\ 0\ 1\ 0\ 1]/2$	$(0.5,0.25)$
			$[0\ 1\ 0\ 1\ 0\ 1\ 1\ 0]/4$	$[0\ 0\ 1\ 0\ 1]/2$	$(0.5,0.25)$
$[4\ 3\ 2\ 4\ 3]/8$	(1-5)	0.2	$[0\ 11\ 1\ 5\ 11\ 12\ 0\ 0]/40$	$[1\ 1\ 1\ 1\ 1]/5$	$(0.6,0.2)$
			$[0\ 11\ 1\ 16\ 0\ 1\ 11\ 0]/40$	$[1\ 1\ 1\ 1\ 1]/5$	$(0.6,0.2)$
			$[1\ 10\ 1\ 16\ 0\ 0\ 12\ 0]/40$	$[1\ 1\ 1\ 1\ 1]/5$	$(0.6,0.2)$

– The first game considers a case where $\mu = [3\ 3\ 5\ 3\ 4]/5$. Here, edge 3 has a particularly high cost (equal to the cost of a tree). In this case, the vulnerability of the graph ($\theta = -0.1$) is negative and the attacker does not make any gain by attacking. Her best strategy is to "*not attack*" and the network manager chooses a tree according to a distribution α that satisfies (4). There exist many such distributions α; two of which are shown in the table. Since there is no attack, each player gets a payoff of zero.

   This game models scenarios where attacking requires an investment from the attacker that is larger than the maximum possible expected reward. As a rational decision maker, the attack will opt to not attack. The network manager needs to randomize his choice of trees to deter the attacker from attacking. In fact, if the network were to pick a fixed tree, then the attacker could get a positive reward by attacking the cheapest link (of cost 3/5) of that tree. In other word, the randomization is necessary for the NE to hold.

– In the next game (second row of the table), the cost of attack is $\mu = [5\ 2\ 3\ 2\ 3]/5$. In this case, the maximum attack reward is exactly equal to zero, and it can be achieved by several attack strategies as can be seen in the table (column 5). Although the attacker cannot gain by launching an attack, the damage she can cause to the network varies depending on the attack she launches.

   This game illustrates the importance of knowing the *type/nature* of an opponent. For example, if the attacker is a competitor who also wants to maximize the loss to the network, then, she will likely attack a link at random with the same probability (which gives a loss of 0.6). However, if the attacker is just interested in her own payoff, then she will probably not launch an attack.

– From these two examples and the first part of the theorem, one can infer that if the network manager is able to influence the attack costs μ, for example making the links harder to attack by investing on security (physical protection, Firewalls, Intrusion Prevention Systems (IPS) to avoid Denial of Service (DoS), etc...), then he can deter the attacker from attacking. This can be done by investing on the links to the point that $\mathcal{M}(E) \leq \mu(E)$ for all subsets of edges $E \subseteq \mathcal{E}$. One can compute the optimal investment by solving an *optimal reinforcement* like problem. The network

reinforcement problem of [4] is related to minimizing the price of increasing the cost of attack of individual edges in order to achieve a target vulnerability (here 0) for the graph. For details see [4]. If the cost of attack can be estimated by some means, this can be a very good candidate for preventive security.

- The last two games are examples where the maximum attack reward is strictly positive. In the first one, the attacker only targets the links that are less costly which turn out to be the minimum cutset of the graph (seen by the attacker). In the second example, the minimum cut seen by the attacker corresponds to links 3 and 5. However, the attack's reward is maximized by targeting all the links with the same probability as it is shown in the table.

- If all links of the graph have the same cost $\boldsymbol{\mu}(e) = \mu$, then the vulnerability of a subset E (defined in equation 3) becomes $\theta(E) = \frac{\mathcal{M}(E) - \mu(E)}{|E|} = \frac{\mathcal{M}(E)}{|E|} - \mu$, and a critical subset is one that maximizes the ratio $\frac{\mathcal{M}(E)}{|E|}$. This definition of criticality corresponds to the one given in [8] where the cost of attack was assumed to be zero. Theorem 1 implies that if the cost of attack μ is larger than $\frac{\mathcal{M}(E)}{|E|}$ for all E, the attacker will not attack. In fact, the net gain of attacking will be negative. If, in the other hand $\mu > \max_{E \subseteq \mathcal{E}} \left(\frac{\mathcal{M}(E)}{|E|} \right)$, then the second part of Theorem 1 corresponds to the critical subset attack theorem in [8] with $\gamma_E = 1$ for some critical subset E. The attacker can take any convex combination of uniform attack on the links in a critical subset, and the manager will choose trees according to (6).

- If $\gamma_{E_c} = 1$ for a some critical subset E_c, we have that the corresponding attack is to target uniformly links in E_c. The defense strategy should verify $\sum_{T \ni e} \boldsymbol{\alpha}_T - \boldsymbol{\mu}(e) \leq \frac{\mathcal{M}(E_c)}{|E_c|}$ for all $e \in \mathcal{E}$, and equality holds for each $e \in E_c$. Also, by the Nash equilibria conditions it must be that for any spanning tree T

$$\sum_{e \in T} \boldsymbol{\beta}(e) = \sum_{e \in T} \frac{1_{e \in E_c}}{|E_c|} = \frac{|E_c \cap T|}{|E_c|} \geq \frac{\mathcal{M}(E_c)}{|E_c|}. \tag{8}$$

The minimum value in the equation above is achieved at each T for which $\boldsymbol{\alpha}_T > 0$.

Since the defender's payoff is equal to $\frac{\mathcal{M}(E_c)}{|E_c|}$, we have that $\mathcal{M}(E_c) = \min_T (|E \cap T|) = |E_c \cap T|$ for each T for which $\boldsymbol{\alpha}_T > 0$. In other words, the defender will select only spanning trees that cross the critical subset in the minimum number of links. Furthermore, the net reward ($\sum_{T \ni e} \boldsymbol{\alpha}_T - \boldsymbol{\mu}(e)$) is the same at each link e of the critical subset E_c. This quantity is equal to θ, the vulnerability of the subset E_c. For any other link, this quantity should be less than θ.

- We have seen in [8] that if the cost of attack is zero, the attacker targets edges on a given critical subset with the same probability. The theorem of this paper tells that this still holds even with positive cost of attack. The attack operates by taking convex combination of uniform strategies on critical subsets.

This uniformity of attack on critical subsets comes from the geometry of the blocker $\mathcal{P}_A^b$ of the spanning tree polyhedron $\mathcal{P}_A$ induced by the defender's payoff matrix A (which is the spanning tree incidence matrix – see appendix B). The attack is no longer uniform if the payoff matrix changes. To see this, assume for example that the defender incurs a certain operation cost η_T by choosing spanning tree T. In

this case his payoff matrix is given as $A_{T,e} = 1_{e \in T} + \eta_T$. We consider the simplest none-trivial topology of two nodes connected by two parallel edges e_1 and e_2. With this topology, edge e_1 corresponds to tree T_1. Similarly for link e_2 and tree T_2. We further simplify by assuming that the attack cost $\boldsymbol{\mu} = 0$.

Letting $(\alpha, 1-\alpha)$ and $(\beta, 1-\beta)$ respectively be the defender's and the attacker's strategy, the expected attack loss for the defender can be written as $L(\alpha, \beta) = \alpha(2\beta + \eta_1 - \eta_2 - 1) + 1 - \beta + \eta_2$. The attacker's expected reward is $R(\alpha, \beta) = \beta(2\alpha - 1) + 1 - \alpha$. By analyzing these payoff functions, we see that the NE is given as follow. If $\eta_1 \geq 1 + \eta_2$, then $\alpha = 0$ and $\beta = 0$. If $\eta_2 \geq 1 + \eta_1$, then $\alpha = 1$ and $\beta = 1$. If $0 < |\eta_1 - \eta_2| < 1$, then $\alpha = 1/2$ and $\beta = \frac{\eta_2 - \eta_1 + 1}{2}$. Hence, the attacker's mixed strategy equilibrium is not in general uniform. We get the uniform distribution only if $\eta_1 = \eta_2$.

This shows the importance of the geometry of the problem (namely the polyhedron induced by the payoff matrix and its blocker) for the determination of the NE structure. The authors have found that [7] for the quasi-zero-sum game defined in section 2 with arbitrary nonnegative payoff matrix A, and attack cost $\boldsymbol{\mu} \geq 0$, the attacker's NE strategies are obtained by normalizing *critical vertices* of the blocker polyhedron $\mathcal{P}_A^b$. In the case of the spanning tree game, the blocker is such that the normalized vertices correspond to uniform distributions. For a general payoff matrix, normalized vertices can give arbitrary distribution.

- The Nash equilibria characterization provided in this paper (and which the authors have studied in a more general setting [7], [6]) can be considered as an application of the result in [1] to the particular case of *quasi zero-sum game*. Although Avis *et al.* were not interested in characterizing Nash equilibria (which would be very laborious for an arbitrary two-player matrix game) and did not explicitly consider the notion of blockers, all the ingredients we have used in our NE characterization can be derived from their results. Our use of the combinatorial notion of blocker was the key to our success in characterizing the mixed strategy Nash equilibria of the game. To our knowledge, such notion was not used before in the context of computing Nash equilibria.

6   Proof of the Critical Subset Attack Theorem

In this section we provide a proof of the Nash equilibrium theorem presented in section 4. In the first part of the proof, we argue that the strategies given in the theorem for $\theta \leq 0$ and $\theta \geq 0$ are best responses to each other. The second part shows the existence of a distribution $\boldsymbol{\alpha}$ that satisfies (4) if $\theta \leq 0$ and (6) if $\theta \geq 0$. The last part of the proof shows that when $\boldsymbol{\mu} = 0$, all Nash equilibria have the form given in part (2) of the theorem. The proof requires the notion blocking pair of polyhedra that we define in the appendix section B.

6.1   Best Responses

First, notice that if the attacker chooses to not attack, then any $\boldsymbol{\alpha}$ will result to the minimum loss of zero for the defender (in particular the one given in the theorem).

Also, if $\boldsymbol{\alpha}$ is such that $\boldsymbol{\alpha}(e) - \boldsymbol{\mu}(e) \leq 0$, $\forall\, e \in \mathcal{E}$, one can easily see from (2) that not attacking is a *dominant* strategy for the attacker. Thus, if $\theta \leq 0$, the strategies given in the theorem are best responses to each other.

Next, we show that if $\theta \geq 0$ the strategies given in (5) and (6) are best response to each other. We start by showing that:

Lemma 1. *If* $\theta \geq 0$, *then not attacking is a* dominated *strategy for the attacker. The domination is strict if* $\theta > 0$.

The Lemma implies that if $\theta \geq 0$ the attacker can always at least do as well as than not attacking (and better strictly better if $\theta > 0$).

Proof Sketch: The proof follows from the fact that if $\theta \geq 0$, then, the attacker can always get a nonnegative attack reward by uniformly targeting the edges of a critical subset E. Indeed, there always exists at least one critical subset. The reward of such attack is lower bounded by $\frac{\mathcal{M}(E) - \boldsymbol{\mu}(E)}{|E|}$, which is greater than zero under the assumption that $\theta \geq 0$. The bound is strict if $\theta > 0$.

Now, suppose that the defender plays a strategy $\boldsymbol{\alpha}$ that satisfies (6). Then, any distribution $\boldsymbol{\beta}$ of the form $\boldsymbol{\beta}(e) = \sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|}$ for some distribution $\gamma = (\gamma_E,\ E \in \mathcal{C})$, achieves a reward of θ. This is the maximum possible reward that the attacker can get. To see this, observe that for any $\boldsymbol{\beta}$,

$$R(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e)\, (\boldsymbol{\alpha}(e) - \boldsymbol{\mu}(e)) \leq \sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e)\theta = \theta. \tag{9}$$

The upper bound of θ is achieved by any $\tilde{\boldsymbol{\beta}} = (\frac{1_{e \in E}}{|E|}, e \in \mathcal{E})$ uniform on a critical subset $E \in \mathcal{C}$. In fact, replacing such $\tilde{\boldsymbol{\beta}}$ in (2) and reordering the terms, we get

$$R(\boldsymbol{\alpha}, \tilde{\boldsymbol{\beta}}) = \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} 1_{e \in T} - \sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} \boldsymbol{\mu}(e) \right) \tag{10}$$

$$= \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\frac{|E \cap T|}{|E|} - \sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} \boldsymbol{\mu}(e) \right) \tag{11}$$

$$\geq \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \frac{\mathcal{M}(E)}{|E|} - \frac{\boldsymbol{\mu}(E)}{|E|} = \frac{\mathcal{M}(E)}{|E|} - \frac{\boldsymbol{\mu}(E)}{|E|} = \theta, \tag{12}$$

where in the last step we use the fact that E is critical.

As a consequence, any distribution of the form $(\frac{1_{e \in E}}{|E|}, e \in \mathcal{E})$ for $E \in \mathcal{C}$ critical is a best response and any convex combination of those distributions is also a best response.

Now assume that $\boldsymbol{\beta}$ is given as in (5) for some distribution $(\gamma_E,\ E \in \mathcal{C})$. Then, the distribution $(\boldsymbol{\alpha}_T,\ T \in \mathcal{T})$ in (6) achieves a loss of $r(\gamma) = \sum_{E \in \mathcal{C}} \gamma_E \frac{\mathcal{M}(E)}{|E|}$. This is the minimum possible loss. To see this, use this expression for $\boldsymbol{\beta}$ to rewrite the expected loss (1) (for any $\boldsymbol{\alpha}$) as

$$L(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\sum_{E \in \mathcal{C}} \gamma_E \left(\sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} 1_{e \in T} \right) \right) \tag{13}$$

$$\geq \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\sum_{E \in \mathcal{C}} \gamma_E \frac{\mathcal{M}(E)}{|E|} \right) \tag{14}$$

$$= \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T r(\gamma) = r(\gamma). \tag{15}$$

To get equation (13) from (1), we have reversed the order of the summations over $\mathcal{E}$ and over $\mathcal{C}$.

The lower bound $r(\gamma)$ can be achieved by choosing $\boldsymbol{\alpha}$ such that $\sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T 1_{e \in T} = \theta + \boldsymbol{\mu}(e)$ for each $e \in \mathcal{E}$ such that $\boldsymbol{\beta}(e) > 0$ (the existence of such $\boldsymbol{\alpha}$ is shown in the second part of the theorem). This can be seen by using $\sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T 1_{e \in T} = \theta + \boldsymbol{\mu}(e)$ and $\boldsymbol{\beta}(e) = \sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|}$ in (1) to get

$$L(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e) \left(\sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T 1_{e \in T} \right) = \sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e) \left(\theta + \boldsymbol{\mu}(e) \right) \tag{16}$$

$$= \theta + \sum_{e \in \mathcal{E}} \left(\sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|} \boldsymbol{\mu}(e) \right) \tag{17}$$

$$= \theta + \sum_{E \in \mathcal{C}} \gamma_E \left(\sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} \boldsymbol{\mu}(e) \right) \tag{18}$$

$$= \theta + \sum_{E \in \mathcal{C}} \gamma_E \left(\frac{\boldsymbol{\mu}(E)}{|E|} \right) \tag{19}$$

$$= \theta + \sum_{E \in \mathcal{C}} \gamma_E \left(\frac{\mathcal{M}(E)}{|E|} - \theta \right) \tag{20}$$

$$= \sum_{E \in \mathcal{C}} \gamma_E \frac{\mathcal{M}(E)}{|E|} = r(\gamma) \tag{21}$$

This implies that the distribution $(\boldsymbol{\alpha}_T, T \in \mathcal{T})$ in (6) is a best response to $\boldsymbol{\beta}$ given in (5).

6.2 Existence of the Equilibrium Distribution $\boldsymbol{\alpha}$

In the previous section we have shown that the strategies given in the theorem are best responses to each other. The distribution in (5) exists by definition. However, a priori, one does not know if there exists a probability distribution that satisfies (4) if $\theta \leq 0$. Similarly, if $\theta \geq 0$, one needs to show the existence of a distribution that verifies the conditions in (6). Using the results discussed in appendix B, we show the existence of such distributions. More concretely, we will show that:

- if $\theta \leq 0$, there exists $\boldsymbol{\alpha}$ verifying, $\boldsymbol{\alpha} \geq 0$, $\mathbf{1}'_{\mathcal{T}} \boldsymbol{\alpha} = 1$, and $A' \boldsymbol{\alpha} \leq \boldsymbol{\mu}$,
- if $\theta \geq 0$, there exists $\boldsymbol{\alpha}$ verifying, $\boldsymbol{\alpha} \geq 0$, $\mathbf{1}'_{\mathcal{T}} \boldsymbol{\alpha} = 1$, and $A' \boldsymbol{\alpha} \leq \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$, with equality in the constraints for each e such that $\boldsymbol{\beta}(e) > 0$.

Recall that 'A' is the tree-link incidence matrix $A_{T,e} = 1_{e \in T}$. Also, the spanning tree polyhedron $\mathcal{P}_A$ is characterized by (see appendix B, [8], and [3])

$$\mathcal{P}_A = \{\mathbf{x} \in R_+^m \mid \mathbf{x}(\mathcal{E}(P)) \geq |P| - 1, \text{ for all feasible partitions } P = \{V_1, V_2, \ldots, V_{|P|}\}\}. \tag{22}$$

P is said to be a *feasible* partition of the nodes $\mathcal{V}$ of G if each V_i induces a connected subgraph $G(V_i)$ of G. We let $\mathcal{E}(P)$ denote the set of edges going from one member of the partition to another, and $G_{\bar{\mathcal{E}}(P)}$ be the graph obtained by removing from G the edges going across P. The number of connected components of $G_{\bar{\mathcal{E}}(P)}$ is denoted $Q(G_{\bar{\mathcal{E}}(P)})$ and is equal the size of the partition P. We have also shown in [8] that $\mathcal{M}(E) = Q(G_{\bar{E}}) - 1$ for all $E \subseteq \mathcal{E}$.

Now, we claim that,

Lemma 2.   – If $\theta \leq 0$, then $\boldsymbol{\mu} \in \mathcal{P}_A$.
  – If $\theta \geq 0$, then $(\theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}) \in \mathcal{P}_A$.

Using the first part of this lemma, and Lemma 3 of Appendix B, we conclude that if $\theta \leq 0$, the value of the following LP is greater than 1.

$$\text{Maximize } \mathbf{1}_T' \mathbf{x}, \quad \text{subject to } A' \mathbf{x} \leq \boldsymbol{\mu}, \text{ and } \mathbf{x} \geq \mathbf{0}. \tag{23}$$

Using this, we can construct a distribution $\boldsymbol{\alpha}$ satisfying (4) by normalizing any solution of this LP.

Similarly, if $\theta \geq 0$, we can construct a distribution $\boldsymbol{\alpha}$ that satisfies $A' \boldsymbol{\alpha} \leq \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$. This gives an $\boldsymbol{\alpha}$ for which we still need to show that equality holds whenever $\boldsymbol{\beta}(e) > 0$, where $\boldsymbol{\beta}$ is a distribution of the form (5). For that, we make the following additional claims.

Theorem 2. *Let* $\mathbf{x}^*$ *be the solution of the following LP:*

$$\text{Maximize } \mathbf{1}_T' \mathbf{x}$$
$$\text{subject to } A' \mathbf{x} \leq \mathbf{b}, \ \mathbf{x} \geq \mathbf{0}. \tag{24}$$

where $\mathbf{b} = \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$. *Then,*
  a) $\mathbf{1}_T' \mathbf{x}^* \leq 1$;
  b) $A' \mathbf{x}^*(e) = \mathbf{b}(e), \forall e \in \mathcal{E}$ *for which* $\boldsymbol{\beta}(e) > 0$, *where* $\boldsymbol{\beta}$ *is given in (5).*

Notice that, from Lemma 2 we have that the value of the linear program is greater than 1. This, combined with part a) of the theorem, imply that the value of the LP is exactly 1. Part b) of the theorem gives the equality conditions that we needed. As a consequence, $\mathbf{x}^*$ satisfies (6) and implies the existence of the NE distribution $\boldsymbol{\alpha}$ when $\theta \geq 0$.

6.2.1   Proof of Lemma 2

  – By definition of θ, we have $\theta \leq 0 \Leftrightarrow \boldsymbol{\mu}(E) \geq \mathcal{M}(E)$ for all $E \subseteq \mathcal{E}$. [8, Lemma 1] gives $\mathcal{M}(E) = Q(G_{\bar{E}}) - 1$, where $Q(G_{\bar{E}})$ is the number of connected components of the graph G when all edges in E are removed. Thus, $\theta \leq 0 \Leftrightarrow \boldsymbol{\mu}(E) \geq Q(G_{\bar{E}}) - 1$ for all $E \subseteq \mathcal{E}$.

Now, let P be a *feasible* partition of the nodes $\mathcal{V}$ of G. Using the above observations, we can conclude that

$$\theta \leq 0 \Leftrightarrow \boldsymbol{\mu}\left(\mathcal{E}(P)\right) \geq Q(G_{\mathcal{E}(P)}) - 1 = |P| - 1 \tag{25}$$

Since the partition P is feasible, $\boldsymbol{\mu}\left(\mathcal{E}(P)\right) \geq |P| - 1$ implies that $\boldsymbol{\mu} \in \mathcal{P}_A$, which ends the proof of the first part of the lemma.

– To prove that the vector $\mathbf{b} = \theta \mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu} \geq \mathbf{0}$ belongs to the polyhedron $\mathcal{P}_A$ whenever $\theta \geq 0$, we argue that

$$\mathbf{b}\left(\mathcal{E}(P)\right) \geq |P| - 1, \quad \text{for all feasible partitions } P. \tag{26}$$

Recall, from the above that for all feasible partitions P

$$\mathcal{M}(\mathcal{E}(P)) = |P| - 1. \tag{27}$$

Now, assume that $\mathbf{b}$ does not verify (26)– i.e $\mathbf{b}\left(\mathcal{E}(P)\right) < |P| - 1$, for some feasible partition P. Then one must have,

$$|P| - 1 > \sum_{e \in \mathcal{E}(P)} \mathbf{b}_e = \theta\left(\mathcal{E}(P)\right) \sum_{e \in \mathcal{E}(P)} 1 + \sum_{e \in \mathcal{E}(P)} \mu(e) \tag{28}$$

$$= \theta\left(\mathcal{E}(P)\right)|\mathcal{E}(P)| + \boldsymbol{\mu}\left(\mathcal{E}(P)\right) \tag{29}$$

$$= \mathcal{M}\left(\mathcal{E}(P)\right) - \boldsymbol{\mu}\left(\mathcal{E}(P)\right) + \boldsymbol{\mu}\left(\mathcal{E}(P)\right) = \mathcal{M}\left(\mathcal{E}(P)\right) \tag{30}$$

which contradicts (27). Thus, $\mathbf{b}\left(\mathcal{E}(P)\right) \geq |P| - 1$ for all feasible P, or equivalently $\mathbf{b} \in \mathcal{P}_A$.

6.2.2 Proof of Theorem 2

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

$$\boldsymbol{\beta}' A' \mathbf{x} = \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e) 1_{e \in T} \right) \tag{31}$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{e \in \mathcal{E}} \left(\sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|} \right) 1_{e \in T} \right) \tag{32}$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{E \in \mathcal{C}} \gamma_E \left(\sum_{e \in \mathcal{E}} \frac{1_{e \in E}}{|E|} 1_{e \in T} \right) \right) \tag{33}$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{E \in \mathcal{C}} \gamma_E \left(\frac{|E \cap T|}{|E|} \right) \right) \tag{34}$$

$$\geq \sum_{T \in \mathcal{T}} \mathbf{x}_T \left(\sum_{E \in \mathcal{C}} \gamma_E \frac{\mathcal{M}(E)}{|E|} \right) \tag{35}$$

$$= \sum_{T \in \mathcal{T}} \mathbf{x}_T r(\gamma) \tag{36}$$

$$= r(\gamma) \mathbf{1}_{\mathcal{T}}' \mathbf{x} \tag{37}$$

On the other hand, from the constraints $A'\mathbf{x} \leq \mathbf{b} = \theta\mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}$ and using the same arguments as in (16)-(21), we have that

$$\boldsymbol{\beta}'\Lambda'\mathbf{x} \leq \boldsymbol{\beta}'\left(\theta\mathbf{1}_{\mathcal{E}} + \boldsymbol{\mu}\right) = \theta + \boldsymbol{\beta}'\boldsymbol{\mu} = r(\gamma). \tag{38}$$

Combining (37) and (38), it follows that,

$$r(\gamma)\mathbf{1}_{\mathcal{T}}'\mathbf{x} \leq \boldsymbol{\beta}'\Lambda'\mathbf{x} \leq r(\gamma). \tag{39}$$

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) Notice from the above and from the conclusion of Lemma 2 that for $\theta \geq 0$ the value of the LP defined in Theorem 2 is exactly equal to 1. Thus,

$$\boldsymbol{\beta}'A'\mathbf{x}^* = r(\gamma)\mathbf{1}_{\mathcal{T}}'\mathbf{x}^* = r(\gamma). \tag{40}$$

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

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

$$\boldsymbol{\beta}'A'\mathbf{x}^* = \sum_{e\in\mathcal{E}}\boldsymbol{\beta}(e)A'\mathbf{x}^*(e) \tag{41}$$

$$< \sum_{e\in\mathcal{E}}\boldsymbol{\beta}(e)(\theta + \boldsymbol{\mu}(e)) \tag{42}$$

$$= \theta + \sum_{e\in\mathcal{E}}\boldsymbol{\beta}(e)\boldsymbol{\mu}(e) \tag{43}$$

$$= r(\gamma), \tag{44}$$

where the last equality is obtained by using the same arguments as in (16)-(21). This contradicts observation (40). As a consequence, $A'\mathbf{x}^*(e) = \theta + \boldsymbol{\mu}(e)$ for all $e \in \mathcal{E}$ with $\boldsymbol{\beta}(e) > 0$.

This ends the proof of the theorem and establishes the existence of an $\boldsymbol{\alpha}$ satisfying (6) for any $\boldsymbol{\beta}$ defined as in (5).

6.3   Enumerating All Nash Equilibria

In this section, we consider the zero-sum game where $\boldsymbol{\mu} = 0$ and show that all Nash equilibria of the game have the form given in Theorem 1 equations (5) and (6).

In this case, since there is no cost of attack, $\theta > 0$. We claim that for any strategy pair $(\boldsymbol{\alpha}_T, T \in \mathcal{T})$ and $(\boldsymbol{\beta}(e), e \in \mathcal{E})$ that are in Nash equilibrium, it must be the case that $(\boldsymbol{\beta}(e), e \in \mathcal{E})$ is given by

$$\boldsymbol{\beta}(e) = \sum_{E\in\mathcal{C}}\gamma_E\frac{\mathbf{1}_{e\in E}}{|E|}, \tag{45}$$

for some probability distribution $(\gamma_E, E \in \mathcal{C})$ on the set of critical subsets.

As a consequence of this, we will conclude that $\boldsymbol{\alpha}$ must be in the form given in the Nash equilibrium theorem.

Because of space limitations, we describe the main points of the proof in appendix A and for the full proof, we refer the interested reader to [6] and [7].

7   Conclusion and Future Work

This paper studies a generalization of the topology design game defined in [8], where a network manager is choosing a spanning tree of a graph as communication infrastructure, and an attacker is trying to disrupt the communication tree by attacking one link of the graph. Assuming that the attacker incurs a positive cost by attacking any given link of the network, we revisit the notions of vulnerability and criticality of a subset of links.

We have determined the values of the attack costs for which a rational attacker will opt to not launch an attack. When the attacker decides to attack, we have shown that there always exists a NE under which she attacks randomly, with the same probability, links in a given critical subset. The randomization can also be done across critical subsets. The network manager chooses only spanning trees that cross the critical set in the minimum number of edges, and such that the sum of the probabilities of all trees going through any link in the critical set minus the cost of attacking that link, is the same.

For the game of zero costs of attack studied in [8], we have characterized the set of all Nash equilibria. The NE strategies are such that the attacker will always target links in critical subsets and attacks all links in the same critical subset with the same probability.

We have shown, by a simple example, that the uniformity of the attack on each critical subset is a consequence of the geometry of the problem. Mainly, the vertices of the blocker of the spanning tree polyhedron are such that if normalized, they results to uniform distribution. This is not always the case. For instance, if the defender incurs different cost of choosing different spanning trees, the attack strategies are no longer uniform on critical subsets.

The proof concepts presented in this paper have been generalized to identify Nash equilibria for a class of *quasi zero-sum* games. For details of the general study, we refer the interested readers to [7] and [6].

References

1. Avis, D., Rosenberg, G., Savani, R., von Stengel, B.: Enumeration of Nash Equilibria for Two-Player Games. Economic Theory 42, 9–37 (2010)
2. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press (March 2004)
3. Chopra, S.: On the Spanning Tree Polyhedron. Operations Research Letters 8(1), 25–29 (1989)
4. Cunningham, W.H.: Optimal Attack and Reinforcement of a Network. J. ACM 32(3), 549–561 (1985)
5. Fulkerson, D.R.: Blocking and Anti-Blocking Pairs of Polyhedra. Math. Programming (1), 168–194 (1971)

6. Gueye, A.: A Game Theoretical Approach to Communication Security. PhD dissertation, University of California, Berkeley, Electrical Engineering and Computer Sciences (March 2011)
7. Gueye, A., Walrand, J.C., Anantharam, V.: Bloking Games. Technical report, University of California, Berkeley (December 2010),
   `http://www.eecs.berkeley.edu/~agueye/index.html`
8. Gueye, A., Walrand, J.C., Anantharam, V.: Design of Network Topology in an Adversarial Environment. In: Alpcan, T., Buttyán, L., Baras, J.S. (eds.) GameSec 2010. LNCS, vol. 6442, pp. 1–20. Springer, Heidelberg (2010)
9. Wolsey, L.A., Nemhauser, G.L.: Integer and Combinatorial Optimization, 1st edn. Wiley-Interscience (November 1999)

A   Proof Sketch of the NE Enumeration Claim

Theorem 1, tells that if the attack costs $\boldsymbol{\mu} = 0$, then all Nash equilibrium pairs $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ of the game have the form given in (6) for $\boldsymbol{\alpha}$ and in (5) for $\boldsymbol{\beta}$. To show this, we claim that for any strategy pair $(\boldsymbol{\alpha}_T, T \in \mathcal{T})$ and $(\boldsymbol{\beta}(e), e \in \mathcal{E})$ that are in Nash equilibrium, it must be the case that $\boldsymbol{\beta}$ is given by

$$\boldsymbol{\beta}(e) = \sum_{E \in \mathcal{C}} \gamma_E \frac{1_{e \in E}}{|E|} \, , \tag{46}$$

for some probability distribution $(\gamma_E, E \in \mathcal{C})$ on the set of critical subsets.

We prove this claim by scaling any mixed strategy $\boldsymbol{\beta}$ (seen as a vector in $R_+^{|\mathcal{E}|}$) with a proper constant such that it belongs to the blocker of the spanning tree polyhedron. The proof is based on the following ideas.

– Since the spanning tree polyhedron $\mathcal{P}_A$ and its blocker $\mathcal{P}_A^b$ are given in terms of feasible partitions, we establish a correspondence between feasible partitions and critical subsets of the graphs. Basically, we show that every critical subset is the set of edges going across the elements of some feasible partition. We define the notion of critical partitions Π_C (corresponding to critical subsets) and show the following equivalent claim:

$$\boldsymbol{\beta}(e) = \sum_{P \in \Pi_C} \gamma_P \frac{1_{e \in \mathcal{E}(P)}}{|\mathcal{E}(P)|} \, , \tag{47}$$

where γ is now viewed as a distribution on the critical partitions.
– Because the game is zero-sum, we know that all NE $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ have payoff $\theta > 0$ which is given as:

$$\theta = \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}_T \left(\sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e) 1_{e \in T} \right) > 0. \tag{48}$$

We argue that $\sum_{e \in \mathcal{E}} \boldsymbol{\beta}(e) 1_{e \in T} > 0$ for all T, and can be scaled by a constant κ so that $\sum_{e \in \mathcal{E}} \kappa \boldsymbol{\beta}(e) 1_{e \in T} \geq 1$. This means that the vector $\kappa \boldsymbol{\beta}$ belongs to the blocker of the spanning tree polyhedron (see Theorem 3 of the appendix). Recall that [8] the vertices of this blocker are vectors of the form $(\frac{1_{e \in \mathcal{E}(P)}}{|P|-1}, \ e \in \mathcal{E})$, for some feasible partitions P.

We argue that by making the proper choice of κ the vector $\kappa\boldsymbol{\beta}$ can be written as

$$\kappa\boldsymbol{\beta}(e) = \sum_{P \in \Pi_C} \gamma_P \frac{1_{e \in \mathcal{E}(P)}}{|P| - 1}, \qquad (49)$$

and show that the proper κ must be in the form $\kappa = \frac{|P|-1}{|\mathcal{E}(P)|}$ for some critical partition $P \in \Pi_C$. But all critical partitions (subsets) have the same ratio $\frac{|P|-1}{|\mathcal{E}(P)|}$. By dividing the equation by κ we get

$$\boldsymbol{\beta}(e) = \sum_{P \in \Pi_C} \gamma_P \frac{1_{e \in \mathcal{E}(P)}}{|\mathcal{E}(P)|}. \qquad (50)$$

Using the correspondence between critical partitions and critical subsets, we get the claim in (46).

B Blocking Pair of Matrices

The discussion in this appendix section is mostly based on [9, pp. 99-101].
Let A be a $r \times m$ nonnegative matrix. The polyhedron $\mathcal{P}_A$ associated with A is defined as the vector sum of the convex hull of its rows $(\mathbf{a}_1, \ldots, \mathbf{a}_r)$ and the nonnegative orthant:

$$\mathcal{P}_A = \text{conv.hull}\,(\mathbf{a}_1, \ldots, \mathbf{a}_r) + \mathbb{R}_+^m. \qquad (51)$$

A row $\mathbf{a}_i$ of A is said to be *inessential* if it dominates a convex combination of other rows of A. If all the rows of A are essential, we say that A is *proper*. In this discussion we will assume that A is proper. For example, if A is the tree-link incidence matrix of the spanning trees of a graph, then A is a proper matrix and $\mathcal{P}_A$ defines the spanning polyhedron of the graph.
Next we define the blocker of the polyhedron $\mathcal{P}_A$.

Definition 2. *The* blocker $\mathcal{P}_A^b$ *of* $\mathcal{P}_A$ *is defined as:*

$$\mathcal{P}_A^b = \left\{ \mathbf{x} \in \mathbb{R}_+^m : \mathbf{x} \cdot \mathbf{y} \geq 1, \ \forall\, \mathbf{y} \in \mathcal{P}_A \right\} \qquad (52)$$

We are interested in characterizing the polyhedron $\mathcal{P}_A$ and its blocker $\mathcal{P}_A^b$. This is given by the following theorem by Fulkerson [5]. It is based on the fact that there is a one-to-one correspondence between the rows of A and the extreme points of $\mathcal{P}_A$.

Theorem 3. *Let the r-by-m matrix A be proper with rows $\mathbf{a}_1, \ldots, \mathbf{a}_r$, and let the polyhedron $\mathcal{P}_A$ be defined as in (51). Let $\mathbf{b}_1, \ldots, \mathbf{b}_s$ be the extreme points of $\mathcal{P}_A^b$, and let B be the matrix having those points as rows. Then,*

  i. *The blocker $\mathcal{P}_A^b$ of $\mathcal{P}_A$ is given by $\mathcal{P}_A^b = \left\{ \mathbf{x} \in \mathbb{R}_+^m : A\mathbf{x} \geq 1 \right\}$.*
  ii. *B is proper, and the polyhedron $\mathcal{P}_A$ can be described as $\mathcal{P}_A = \left\{ \mathbf{x} \in \mathbb{R}_+^m : B\mathbf{x} \geq 1 \right\}$.*
  iii. *The blocker of the blocker $\mathcal{P}_A^b$ verifies $\left(\mathcal{P}_A^b\right)^b = \mathcal{P}_A$.*

A and B are said to form a blocking pair of matrices.

Blocking pairs of matrices play an important role in the combinatorial problem of *maximum packing* (see Fulkerson[5]). In this paper, we use the theory of blocking pair to provide an easy argument for the existence of a probability distribution that satisfies a certain number of constraints.

Consider the following linear program:

$$\text{Maximize } \mathbf{1}'\mathbf{x}$$
$$\text{subject to } A'\mathbf{x} \leq \mathbf{w}, \text{ and } \mathbf{x} \geq \mathbf{0}, \tag{53}$$

where the constraint A' is a nonnegative matrix.

We are interested to knowing whether the value of the program is greater than 1 or not. The following lemma gives an answer to that question.

Lemma 3. *The value of the LP in (53) is greater than 1 if and only if* $\mathbf{w}$ *belongs to the polyhedron* $\mathcal{P}_A$ *defined by* A.

Proof. The proof of the lemma is as follow.

First notice that strong duality holds for this LP. In fact, Slater's condition [2] is satisfied for any nonnegative $\mathbf{w}$. The dual of the LP is given as:

$$\text{Minimize } \mathbf{w}'\mathbf{y}$$
$$\text{subject to } A\mathbf{y} \geq \mathbf{1}, \text{ and } \mathbf{y} \geq \mathbf{0}. \tag{54}$$

The constraints of the dual program (54) define the blocker $\mathcal{P}_A^b = \left\{ \mathbf{y} \in \mathbb{R}_+^m : A\mathbf{y} \geq 1 \right\}$ of the polyhedron $\mathcal{P}_A$.

Now, if $\mathbf{w}$ belongs to $\mathcal{P}_A$, then for all $\mathbf{y} \in \mathcal{P}_A^b$, we have that $\mathbf{w}'\mathbf{y} \geq 1$.

Conversely, if $\mathbf{w}'\mathbf{y} \geq 1$ for all $\mathbf{y} \in \mathcal{P}_A^b$, then $\mathbf{w}$ must be in the blocker of $\mathcal{P}_A^b$ which is $\mathcal{P}_A$.