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

Assane Gueye, Vladimir Marbukh [*][†][‡]

Abstract

In this paper, we propose a quantification of the vulnerability of a communication network where 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 related them to notions in Graph Theory. Using numerical examples, we show that critical subsets depend both on the value model and on the connectivity of the network. Finally we discuss how the critical subsets and the vulnerability metric can be used for network improvement.

1 Introduction

"....one cannot manage a problem if one cannot measure it..."

[*] 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.

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 a metric can serve as guidance when designing new networks in adversarial environments. Also, knowing such a 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 *loss-in-value* of a network when links are attacked by an adversary.

To analyze the vulnerability of networks, a common approach has been to use the theory of graphs. Indeed, the interactions and interdependencies between members, parts and subsystems of a network can be represented by (weighted) graphs [14]. Numerous papers and books have been written for this subject and various vulnerability metrics have been proposed depending on the type of network (communication, transport, financial, power grid), and the (sub)graph structure (flow and source-destination pairs–Ford, Fulkerson [15]; shortest path–Freeman [5]; spanning tree–Tutte [23], Nash-Williams [19]; neighborhood and cluster–Watts and Strogatz [24]; etc...) used to analyze the graph.

This paper follows the approaches by Tutte and Nash-Williams and uses spanning trees as the (sub)graph structure to study and analyze graphs. Spanning trees have a number of desirable properties that have made them a 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 bridged Ethernet local area networks. Spanning trees (or the directed graph equivalent of spanning arborescences) also can be used to model access networks and sensor networks where there is a central node that every other node is trying to connect to.

Although the present paper is about spanning trees, the approach presented here can be used with other (sub)graph structures such as shortest path, flows and source-destination pairs, Hamiltonian cycle, etc... In Gueye, Walrand, Anantharam, [8, Chap. 4] an example is provided where the focus is on flows between a set of sources to a set of destinations.

We use spanning trees to quantify the *importance* of links as the relative *loss-in-value* of the network when the links fail. Naturally, the first question towards such quantification is: “what is the *value* of a communication network?" The value of a network depends on many 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 a value is a subjective topic and, to the knowledge of the authors, there is no systematic quantification of the value of a communication network. Attempts have, however, been made to quantify (as a first order approximation) the value of a network as a function of its number of nodes. These attempts include the laws by David Sarnoff [1], Robert Metcalfe [7], David Reed ([21]), Odlyzko *et. al.* [3] and other laws mentioned later in the paper.

We build upon these models and use them in the process to quantify the vulnerability of a network. More precisely, we use the models as a proof of concept for defining the importance of network links *relative to spanning trees* in the follow way: 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. When one link that does not belong to the spanning tree fails, communication can still be carried on between all nodes. If, on the other hand, the link belongs to the spanning tree, then some exchanges that originally could be carried out become impossible. 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 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) due to the failure of the link.

In this process, we propose a generalization of the notion of *betweenness centrality* ([25], [20]), which is a measure of the *importance* a link has within a network.

Link failures, in traditional network *reliability* and *fault tolerance* analysis [18], are assumed to occur because of random events (faults) such as human errors and/or machine failures. As such, probabilistic tools have been used to model failures: the probability distribution often chosen in an *arbitrary* way. In this paper, we consider scenarios where link failures are due to the action of a malicious and *strategic* adversary who can analyze the network structure to decide which links to attack. The attack is conceptual in this study but can be thought of as a physical attack, jamming or launching a denial-of-service (DOS) against a particular link.

We use a 2-player game theoretic model to study the *static* interaction between a defender (network manager) an a malicious attacker and analyze the network vulnerability. Applying game theoretic models to (this) security problem is a natural process and it has recently attracted a lot of interest (see surveys [22], [17]). In the present paper, we use the graph of a network to set up a game between a network manager who would like to choose a spanning tree as communication infrastructure in anticipation of an intelligent attack by a malicious attacker who is trying to inflict the most damage. The adversary also chooses a link to attack in anticipation of the choice of spanning tree. We use the network value models and the links' LiV discussed above to derive payoffs for both players.

By analyzing the Nash equilibria of the game, we determine the actions of both the attacker and the defender. We also identify the set of links that are most critical for the network and use them to define a network vulnerability metric. The critical subsets depend both on the connectivity of the network as well as the value model used to compute the links' LiV.

Knowing the critical parts of a network is crucial for network design and improvement. For the general model, computing a critical subset involves objects that are exponential in number (e.g.; spanning trees). Thus, finding a critical subset is a priori a challenging task. We discuss two particular models where efficient algorithms have been derived to compute critical subsets of a graph. We are still investigating efficient algorithms for the general case.

The reader should be advised that the use of Game Theory in this paper is not meant to capture the actual *active* interaction between a defender who dynamically chooses a spanning tree and an attacker who dynamically tries to disrupt the communication. Game Theory is rather used here as a modeling tool to study network configurations in adversarial environment. Instead of defining/setting failure probabilities as it is done in conventional *reliability* analysis, we consider the worst-case scenario where a strategic adversary studies the network structure to figure out the most damaging attack. The defender also does the same.

The game theoretic model has allowed us to understand the attacker most damaging attack and the defender's best defense strategy; to identify the most critical links and relate them to well defined Graph Theory notions; and to define a vulnerability metric for the network. We believe that the study in this paper can be a good starting point for more *realistic* and certainly more complicated games where the strategies are dynamically chosen; all the game parameters (e.g.; graph topology, attack cost, etc...) are not known to all players; and the actual network protocol, the link capacities as well as the nodes in the network are taken into consideration.

The remainder of this paper is organized as follows. 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 presents an analysis and interpretation of the players' strategies. Section 4.2 discusses our choice of metric for the vulnerability of a network. In Section 4.3 we discuss the critical subsets in two models where closed-form characterization are found for the critical subsets and there exists polynomial time algorithm to compute them. Two numerical examples are used in Section 4.4 to show 1) the dependency of the critical subsets to the value model and to the connectivity of the network, 2) how the vulnerability metric and the notion of critical subset can be used for network improvement. Concluding remarks and future directions are presented in Section 5.

2 On the Value of a Communication Networks

The present paper is mainly concerned with quantifications of the value of a network as functions of its number of nodes. In the next section, we discuss some of such quantifications.

The reader should be advised that the purpose of this paper is neither to compare the different value models presented next, nor to decide which one is more realistic than the others[*] Rather, for a given model, we are interested in quantifying the network vulnerability and deciding which network configuration is more vulnerable, and for a given network configuration what are the links that are most critical.

2.1 Network value models

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 $f(G) = |V|$. This law was mainly designed for radio/TV broadcast networks where the popularity of a program is measured by the number of listeners/viewers. Considering networks where

[*] In fact, we believe that each model makes "some sense" for the purpose it was designed for.

each node can in principle communication with each other node (e.g. Ethernet), Robert Metcalfe [7] has suggested that the value of a network grows as a function of the total number of possible connections ($f(G) = |V|^2$). David Reed ([21]) has proposed an exponential ($f(G) = 2^{|V|}$) model for the utility of a group forming network. This is the total number of possible groups that can be formed in a (connected) network of $|V|$ nodes. For Odlyzko *et. al.* [3] a more reasonable approximation of the value of a network as a function of the number of nodes is $f(G) = |V|log(|V|)$. This law particularly has a diminishing return property which is missing in other models. Considering a friendship network, Walrand has proposed a power law model where the value of a network is estimated as $f(G) = |V|^{1+a},\ a \leq 1$. This law combines and generalizes Sarnoff and Metcalfe's laws. The parameter a is a design parameter and needs to be specified. Details about these value models can be found in [10]. In a very recent paper [16], Aron *et.al.* have considered a *many-to-one* (e.g.; Sensor Network) where there exists a designated node ($\mathbf{S}$) (e.g.; Gateway) to which every other node would like to connect to. In their model, the value of a network is given by $f(G) = |V|\mathbf{1}_{S\in V}$, where $\mathbf{1}_{S\in V}$ indicates whether node $\mathbf{S}$ belongs to the network or not. While, we were finishing editing the present paper, the authors in [16] have indicated that they have studied a model where the value of a the network is equal to the size of its largest connected component (i.e.; the maximum number of nodes that can communicate). In this model, $f(G) = \max_i |V(G_i)|$, where G_i is a connected component of a the graph G and $V(G_i)$ is its node set.

Notice that while the last two models implicitly assume that the graph is (already) disconnected, the other models are defined for connected graph. If the graph is (still) connected, the last two models are equivalent to Sarnoff's model. In this paper, we are interested in scenarios where the graph is "first" connected, hence has a value of $f(G)$, and "later" gets disconnected (to become G') because of the action of a malicious attacker. We compare the value $f(G)$ to $f(G')$. Details of this is discussed in the next section.

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.

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 $f(G)$ denote such a value. In general this value depends on other aspects of the network (e.g.; topology, type of traffic, number of links). However, in this paper, $f(\cdot)$ is assumed to be only a function of the number of nodes and can be chosen to be one of the models discussed earlier. Also, there is in general a cost associated to the maintenance of each spanning tree. In this paper, this cost is intentionally ignored and left for future studies. Finally, we assume that $f(\emptyset) = 0$; 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 $T_i,\ i \in \{1, 2\}$ induces a connected component which we denote $G_i^{(T,e)}$ (where we add the superscript (T, e) to stress the fact that the subgraphs $G_i^{(T,e)}$ are obtained by removing link e from spanning tree T). We let $V(G_i^{(T,e)})$ denote the node set of $G_i^{(T,e)}$, with $n_i = |V(G_i^{(T,e)})|$ (see Figure (1).c)). The value of the resulting disconnected network $G^{(T,e)} = G_1^{(T,e)} \cup G_2^{(T,e)}$ is assumed to be $f(G^{(T,e)}) = f(G_1^{(T,e)}) + f(G_2^{(T,e)})$. As an example, for the Metcalfe model discussed above, $f(G^{(T,e)}) = |V(G_1^{(T,e)})|^2 + |V(G_2^{(T,e)})|^2$. For Aron *et. al.* model presented in [16], $f(G^{(T,e)}) = |V(G_1^{(T,e)})|\mathbf{1}_{S \in V(G_1^{(T,e)})} + |V(G_2^{(T,e)})|\mathbf{1}_{S \in V(G_1^{(T,e)})}$.

When link e is removed, some exchanges that could be carried on the original network become impossible. As such, it is reasonable to assume that $f(\cdot)$ verifies $f(G) \geq f(G_1^{(T,e)}) + f(G_2^{(T,e)})$, i.e.; the function $f(\cdot)$ is *superadditive*. 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 $f(G)$ to $f(G_1^{(T,e)}) + f(G_2^{(T,e)})$. We define the loss-in-value (LiV) of link e, relative to spanning tree T, to be equal to $f(G) - \Big(f(G_1^{(T,e)}) + f(G_2^{(T,e)})\Big)$.

If the link does not belong to the spanning tree, then removing it leaves the network connected and the value of network is still $f(G)$. As a consequence, if link $e \notin T$, its LiV, relative to spanning tree T, is equal to 0. In conclusion, the normalized LiV of link e, relative to spanning tree T is defined as

$$\boldsymbol{\lambda}(T, e) = 1 - \frac{f(G_1^{(T,e)}) + f(G_2^{(T,e)})}{f(G)}. \tag{1}$$

with the understanding that if $e \notin T$, $G_1^{(T,e)} = G$ and $G_2^{(T,e)} = \emptyset$, giving $\boldsymbol{\lambda}(T, e) = 0$.

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

Remark 1 *With the definition in (1), the normalized LiV of a link relative to any spanning tree is always* *equal to zero under Sarnoff's model (i.e $\boldsymbol{\lambda}(T,e) = 0$, $\forall e$ and T). This makes the discussion below irrelevant for that model. Consequently, we drop Sarnoff's model for the remainder of this paper. Instead, we discuss a close model introduced in [12]), 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.* $\boldsymbol{\lambda}(T,e) = 1 - \mathbf{1}_{e \notin T}$*). 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 network. Betweenness centrality is a measure that has long been used for that purpose. Next, we propose a quantification of the importance of a that generalizes the notion of betweenness. We start by recalling the betweenness centrality measure as it was defined by Freeman [5].

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 $\boldsymbol{\alpha}(P_{ij}) = \frac{1}{g_{ij}}$. 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 it is not. $\boldsymbol{\lambda}(p_{ij}, e) := g_{ij}(e)$ is the total *value* of e for a communication between i and j.

We use this interpretation to generalize the betweenness centrality to quantify the importance of a link. Since we are interested in spanning trees, we write the definition with respect to spanning trees (instead of paths).

$$\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \sum_{T} \boldsymbol{\alpha}(T) \boldsymbol{\lambda}(T, e), \tag{2}$$

where the summation is now over spanning trees.

Our generalization allows both the consideration of non-binary weights of the links as well as a *preference* for spanning tree utilization. Here, the weights $\boldsymbol{\lambda}(T,e)$ of the links, with respect to spanning trees, are taken to be the links' LIV mentioned above. In general some spanning trees are preferable to others. There are many reasons for such a preference among spanning trees. They include, among others, the cost of utilizing the links, the overall communication delay and the vulnerability of links and the vulnerability to attack. In this paper, where the focus is on adversarial environments, the preference of spanning tree $\boldsymbol{\alpha}(T)$ is derived from the mixed strategy Nash equilibria 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 $N = |\mathcal{T}|$. These assumptions (connected and undirected) can be easily relaxed by considering spanning forests and spanning arborescence instead of spanning trees.

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 $\boldsymbol{\mu}(e)$ that an attacker needs to spend to launch a successful attack on e. For the network manager, the LiV $\boldsymbol{\lambda}(e,T)$ of link e relative to spanning tree T is given by (1). 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 $\boldsymbol{\eta}(T) + \boldsymbol{\lambda}(T,e)$ for the network, while the net attack reward is equal to $\boldsymbol{\lambda}(T,e) - \boldsymbol{\mu}(e)$ 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

$\{\boldsymbol{\alpha} \in \Re_+^N \mid \sum_{T\in\mathcal{T}} \boldsymbol{\alpha}(T) = 1\}$ be the set of mixed strategies for the network manager, and $\{\boldsymbol{\beta} \in \Re_+^m \mid \sum_{e\in\mathcal{E}} \boldsymbol{\beta}(e) = 1\}$ the set of mixed strategies for the attacker. The defender is choosing $\boldsymbol{\alpha}$ to minimize the expected net communication cost $L(\alpha, \beta)$ while the attacker is choosing $\boldsymbol{\beta}$ to maximize the expected net reward $R(\alpha, \beta)$.

$$
\begin{aligned}
L(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \sum_{T\in\mathcal{T}} \boldsymbol{\alpha}(T) \left(\boldsymbol{\eta}(T) + \sum_{e\in T} \boldsymbol{\beta}(e) \boldsymbol{\lambda}(T,e) \right), && (3)\\
R(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \sum_{e\in\mathcal{E}} \boldsymbol{\beta}(e) \left(\sum_{T\ni e} \boldsymbol{\alpha}(T) \boldsymbol{\lambda}(T,e) - \boldsymbol{\mu}(e) \right). && (4)
\end{aligned}
$$

In the sequel, we have focused on the case where $\boldsymbol{\eta}(T) = \eta$ is constant; hence not relevant to the optimization in (3), which now becomes the minimization of $\sum_{T\in\mathcal{T}} \boldsymbol{\alpha}(T) \sum_{e\in T} \boldsymbol{\beta}(e)\boldsymbol{\lambda}(T,e)$. As a consequence, we ignore $\eta(T)$ for the rest of this paper. The general case of $\boldsymbol{\eta}(T)$ will be considered in subsequent studies.

3.2 Nash equilibrium theorem

The Nash equilibrium theorem was established in Gueye *et. al.* [8, Chap. 4] using the theory in Blocking Pairs of Polyhedra. For completeness purposes, we briefly recall it here. A quick tutorial on the theory of Blocking Pairs of Polyhedra can be found in that reference. In this paper, we show that the vertices of the blocker of the payoff matrix correspond to *Minimal Disconnecting Subsets (MDS)* (see appendix A) of links. As such, we reduce the discussion of the Nash equilibrium theorem below to subsets of links of the graph.

Let Λ be the $N \times m$ nonnegative tree-link payoff matrix whose entries are defined in (1). The polyhedron P_Λ associated with Λ is defined as the vector sum of the convex hull of its rows $(\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_N)$ and the nonnegative orthant:

$$P_\Lambda = \text{conv.hull}\,(\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_N) + \mathbb{R}_+^m. \qquad (5)$$

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

$$bl(P_\Lambda) = \left\{ \mathbf{y} \in \mathbb{R}_+^m : \ \Lambda\mathbf{y} \geq \mathbf{1}_{\mathcal{T}} \right\}. \qquad (6)$$

We assume that the graph of the network contains more than 1 spanning tree (otherwise the game is trivial). With this assumption and the definition in (1), it is easy to verify that matrix Λ does not contain an all-zeros row. As a consequence, both P_Λ and its blocker $bl(P_\Lambda)$ are well defined and non-trivial.

Now, let $\boldsymbol{\omega}$ be a vertex (i.e.; an extreme point) of $bl(P_\Lambda)$. We have shown in the appendix that the support of $\boldsymbol{\omega}$ is a minimal disconnecting set (MDS); call it $E_{\boldsymbol{\omega}}$. To ease notation, we drop the subscript (unless needed) and assume, in the sequel, that E is the support of $\boldsymbol{\omega}$ and the pair $(\boldsymbol{\omega}, E)$ is always given together. Let $\mathcal{W}$ the set of all MDS that correspond to vertices $\boldsymbol{\omega} \in bl(P_\Lambda)$. We write $\boldsymbol{\omega} = (\boldsymbol{\omega}(e), e \in \mathcal{E})$, also we define $\kappa(E) := \sum_{e\in\mathcal{E}} \boldsymbol{\omega}(e)$. Note that $\boldsymbol{\omega}(e) \geq 0$ for all $e \in \mathcal{E}$ (with strict inequality only if $e \in E$) and $\boldsymbol{\omega}(\mathcal{E}) > 0^{\dagger}$; so that $\boldsymbol{\beta} = (\frac{\boldsymbol{\omega}(e)}{\kappa(E)}, e \in \mathcal{E})$ is a probability distribution on $\mathcal{E}$. We call it the probability distribution associated to $(\boldsymbol{\omega}, E)$.

Also for each $E \in \mathcal{W}$, consider the quantity

$$\frac{\min_{T\in\mathcal{T}} \left(\sum_{e\in E} \boldsymbol{\omega}(e)\boldsymbol{\lambda}(T,e)\right)}{\kappa(E)}; \tag{7}$$

This quantity is the minimum loss seen by the defender if the attacker were to choose a target according to the distribution $(\frac{\boldsymbol{\omega}(e)}{\kappa(E)}, e \in \mathcal{E})$ associated to $(\boldsymbol{\omega}, E)$. A closer look at the expression above shows that the numerator is equal to 1. In fact, we already know that if $\boldsymbol{\omega}$ belongs to the blocker $bl(P_\Lambda)$, $\sum_{e\in\mathcal{E}} \boldsymbol{\omega}(e)\boldsymbol{\lambda}(T,e) \geq 1$ for all $T \in \mathcal{T}$. Now, if $\boldsymbol{\omega}$ is a vertex of the blocker $bl(P_\Lambda)$, one can easily show that, there must exist some T_o such that $\sum_{e\in E} \boldsymbol{\omega}(e)\boldsymbol{\lambda}(T_o e) = 1$. If this is not the case, then we can divide $\boldsymbol{\omega}$ by this minimum sum and obtain $\tilde{\boldsymbol{\omega}}$ which belongs to $bl(P_\Lambda)$ and is strictly dominated by $\boldsymbol{\omega}$. This contradicts the assumption that $\boldsymbol{\omega}$ is a vertex. Thus we can drop the minimization and rewrite the quantity above as $\frac{1}{\kappa(E)}$.

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

$$\theta(E) := \frac{1}{\kappa(E)}\left(1 - \sum_{e\in E} \boldsymbol{\omega}(e)\boldsymbol{\mu}(e)\right). \tag{8}$$

$\theta(E)$ is the expected attack reward associated $(\boldsymbol{\omega}, E)$ if the attacker were to choose a resource to attack according to the distribution $\boldsymbol{\beta} = (\frac{\boldsymbol{\omega}(e)}{\kappa(E)}, e \in \mathcal{E})$.

We call the subset $E \subseteq \mathcal{W}$ a *critical* subset if

$$\theta(E) = \max_{F\in\mathcal{W}} \theta(F)\ . \tag{9}$$

We define $\theta^* := \max_{F\in\mathcal{W}} \theta(F)$ and we let $\mathcal{C}$ denote the set of all critical subsets.

[†] 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 (P_Λ is not empty).

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. $\boldsymbol{\beta}(e_\emptyset) = 1$) is always an optimal strategy for the attacker. In this case, the equilibrium strategy $(\boldsymbol{\alpha}(T),\ T \in \mathcal{T})$ for the defender is such that*

$$\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \sum_{T \in \mathcal{T}} \boldsymbol{\alpha}(T)\boldsymbol{\lambda}(T, e) \leq \boldsymbol{\mu}(e), \quad \forall e \in \mathcal{E}. \tag{10}$$

   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 $(\boldsymbol{\beta}(e), e \in \mathcal{E})$ defined by*

$$\boldsymbol{\beta}(e) = \sum_{E \in \mathcal{E}} \gamma_E \boldsymbol{\beta}_E(e) \tag{11}$$

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

$$\begin{cases} \vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) - \boldsymbol{\mu}(e) = \theta^* & \text{for all } \ e \in \mathcal{E} \ \text{ such that } \ \boldsymbol{\beta}(e) > 0. \\ \vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) - \boldsymbol{\mu}(e) \leq \theta^* & \text{for all } \ e \in \mathcal{E}. \end{cases} \tag{12}$$

   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

$$r(\gamma) := \sum_{E \in \mathcal{C}} \frac{\gamma_E}{\kappa(E)}. \tag{13}$$

3. *If $\boldsymbol{\mu} = 0$, then every Nash equilibrium pair of strategies $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ has the form given in (11) and (12).*

4. *For all $\boldsymbol{\mu} \geq 0$, then in every Nash equilibrium the attacker payoff is equal to θ^*.*

Proof: Items 1-3 of the theorem have been proved in [8, Chap. 4] for a more general game. The proof of Item 4 is tedious and not necessary for the understanding of the discussion made here. We skip it here and present a version in our online report [9]. ■

4 Discussion and Analysis

The Nash equilibrium theorem gives a characterization of the attacker's maximum net attack gain θ^* in any Nash equilibrium; and depending on the sign of this net gain, there are two decisional situations. If the maximum gain is negative ($\theta^* < 0$), the attacker will not launch an attack and the defender randomly chooses a spanning tree according to a distribution $\boldsymbol{\alpha}$ that satisfies (10). This randomization is necessary for the equilibrium to hold. When the gain is nonnegative ($\theta^* \geq 0$), the equilibrium strategy for the attacker is to always launch an attack that focuses only on edges belonging to critical subsets. Her randomized strategy is a convex combination of the probability distributions induced by the critical subsets. In this case, the defender chooses a spanning tree according to a probability distribution that satisfies (12). When there is no cost associated with launching an attack ($\boldsymbol{\mu} = 0$), the attacker's maximum net gain is always positive and she will always launch an attack. In this case, the theorem gives a closed form characterization (Item 3) of all Nash equilibrium pairs $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ of the game: $\boldsymbol{\beta}$ has the form in (11), while $\boldsymbol{\alpha}$ has to satisfy (12).

4.1 Attacker and Defender Equilibrium Strategies

The attacker's equilibrium strategy is completely defined by the structure of the graph, the network value model and the cost to attack ($\boldsymbol{\mu}$–whenever it is not zero). In fact, the attacker's strategy is derived from a vertex of the blocker polyhedron ($bl(P_\Lambda)$), which is solely dependent on the graph structure and the network value model. *This indicates that a sophisticated attacker would analyze the topology of the graph to decide which links to attack.* This contrasts with conventional reliability models where the failure probability of a link is chosen without any consideration of the structure of the graph.

Also, the *attacker's strategy focuses only on links that are critical* (see equation (11)). If there is no cost associated to launching an attack (i.e.; $\boldsymbol{\mu} = \mathbf{0}$–which corresponds to the most powerful adversary), the attacker targets only links that have maximum expected loss-in-value (LiV–$\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha})$). This can be seen in inequalities in (12), which, when $\boldsymbol{\mu} = \mathbf{0}$, become $\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha}) \leq \theta^*, \text{ for all } e \in \mathcal{E}$; with equality whenever link e is targeted with positive probability ($\boldsymbol{\beta}(e) > 0$). *These links, which are attacked with positive probability, correspond in this case to the most important ones for the defender.* For general cost of attack $\boldsymbol{\mu} \geq 0$, the attacker's strategy is such that *each attacked link gives the maximum possible net reward achievable in any Nash equilibrium* (see equation (12) and Item 4 of the theorem). When attacking become expensive and the maximum net attack reward is negative ($\theta^* < 0$), the attacker does not launch one.

The defender's equilibrium strategy $\boldsymbol{\alpha}$ can be interpreted as the *best way to choose a spanning tree in the presence of a strategic adversary.* In fact, as a best response to the attacker's strategy, $\boldsymbol{\alpha}$ minimizes the total expected LiV. Each entry $\boldsymbol{\alpha}(T)$ of the distribution vector is an indication about the *robustness* of spanning tree T–whenever $\boldsymbol{\alpha}(T) = 0$ choosing spanning tree T implies high expected loss due to an attack.

The defender's choice of $\boldsymbol{\alpha}$ also quantifies how much *importance* a given link has within the network. Here, a link with 'high importance' is one that has a high average LiV. When there is no attack cost associated to an attack, the probability distribution $\boldsymbol{\alpha}$ is such that the *most important links corresponds to the most critical ones.* When attacking requires a relatively substantial effort the maximum expected net attack reward can be negative $\theta^* < 0$. In this case the defender chooses the distribution $\boldsymbol{\alpha}$ such that the attacker has no incentive to attack. *Such a choice can be seen as a deterrence tactic for the defender.*

4.2 Vulnerability Metric: Critical vs Important Links

The vulnerability metric (θ^*) proposed in this paper *reflects both the importance of network links as well as the willingness of an attacker to attack them.* This is a desirable feature for a vulnerability metric because no rational adversary will launch an attack if the expected net attack reward is less than zero. On the other hand, links with high importance and low cost of attack are very attractive to adversaries.

For the analysis of the vulnerability metric θ^* one needs to make the interesting "distinction" between the *importance* of a link and its *criticality.* The importance that a link has within the network is quantified by its average LiV $\vartheta(e, \boldsymbol{\lambda}, \boldsymbol{\alpha})$. It is specified by the defender's choice of $\boldsymbol{\alpha}$ and the network value model $f(G)$. The *criticality* indicates the net gain an attacker receives by attacking the link. It depends not only on the *importance* of a link, but also on the cost of attacking the link. *The vulnerability metric θ^* corresponds to the* criticality *of the most critical links.* When there is no cost associated to launching an attack ($\boldsymbol{\mu} = \mathbf{0}$), *criticality* and *importance* coincide: $\theta^* = \vartheta(e^*, \boldsymbol{\lambda}, \boldsymbol{\alpha})$, where link e^* is a critical link.

Notice that, similar to the attacker's equilibrium strategy, the vulnerability metric θ^* is completely defined by the structure of the graph, the network value model and the cost to launch attacks $\boldsymbol{\mu}$. This is also the case for critical subsets (θ^* is equal to the criticality of the most critical subsets). Example 1 below is an illustration of this dependency while Example 2 shows how the vulnerability metric can be used in network improvement.

4.3 Critical Subsets

We have shown (by the NE theorem and appendix A) that a critical subset of links is necessarily a *minimal disconnecting set* (MDS) of the graph. In some particular cases of network value model, there exists a closed-form characterization of critical subsets. Next, we discuss GWA [11] and Aron's [16] models. We assume that the attack cost $\boldsymbol{\mu} = 0$ (which corresponds to the most powerful attacker).

Recall that for the GWA model introduced in [11] we have $\boldsymbol{\lambda}_{T,e} = 1 - \mathbf{1}_{e \notin T}$. In this case, the vertices of the blocker polyhedron have been characterized [6] and $\kappa(E) = \frac{|E|}{Q(G \backslash E)-1}$, for any subset of links E such that $G\backslash E$ is disconnected; $Q(G\backslash E)$ is the number of connected components of the graph G when edges in E are removed. In this case, the vulnerability θ^* is defined as

$$\theta^* = \max_{E \subseteq \mathcal{E}} \left(\frac{Q(G\backslash E) - 1}{|E|} \right). \tag{14}$$

This vulnerability metric (actually its inverse) has been previously proposed by Gusfield [13] in a related but slightly different (non-game theoretic) context. In [13] $\sigma(G) = \min_{E \subseteq \mathcal{E}} \left(\frac{|E|}{Q(G\backslash E)-1} \right) = \frac{1}{\theta^*}$ was used as a measure of the *invulnerability* of the graph, i.e. the smaller this is the more vulnerable the graph is.

Inspired by Cunningham [4] an algorithm is presented in [11] that computes θ^* and a critical subset in polynomial time. The algorithm is essentially based on the Max-Flow/Min-Cut algorithm.

In [16], Aron *et. al* consider a Many-to-One network model where there exists a designated node S to which all other nodes want to communicate with. Indeed, connecting all nodes to S using the minimum number of links results to a spanning tree that is rooted at S. With this model, the links' LiV are defined as $\boldsymbol{\lambda}_{T,e} = 1 - \frac{n_1 \mathbf{1}_{S \in V_1} + n_2 \mathbf{1}_{S \in V_2}}{n}$, where it is assumed that removing link e from spanning tree T divides the nodes of the network into two groups V_1 and V_2 of sizes $V_1 = n_1$ and $V_2 = n2$, $n_1 + n_2 = n$. In this case, $\kappa(E) = \frac{|E|}{\rho(E)}$, for any subset of links E such that $G\backslash E$ contains at least one node that is disconnected from S; $\rho(E)$ is the number of nodes that are disconnected from S. The vulnerability is given by

$$\theta^* = \frac{1}{n} \max_{E \subseteq \mathcal{E}} \left(\frac{\rho(E)}{|E|} \right) = \frac{1}{n} \frac{1}{\pi(G)}, \tag{15}$$

where $\pi(G)$ is defined (actually renamed) as the *persistence* of the graph G. This metric was also previously proposed [4] as a measure for the *strength* of a directed graph.

The authors of [16] have derived a polynomial time algorithm to compute θ^* and a critical subset in this case. As in the previous case, the algorithm proposed here is also based on the Max-Flow/Min-Cut algorithm.

While we were finishing editing this paper, the same authors of [16] have indicated (in a preprint which we received a copy of) that they have studied a model where the value of a (disconnected) network is equal to the size of its biggest connected component. With this model, they have also derived a polynomial algorithm to compute θ^* and a critical subset.

As seen earlier, efficient algorithms was derived to compute critical subsets in the particular cases discussed above. For the general case, a naive approach to compute a critical subset would be to consider each MDS (there is an exponential number of them), and for each one, solve a linear program which involves the spanning trees of the graph (there is also an exponential number of them) to compute the function $\kappa(\cdot)$. This is intractable for any practical example of network. We are investigating the general case to find out whether an efficient algorithm can be derived; and if not, for which network value models an efficient algorithm exists.

4.4 Numerical Examples

Example 1 **Critical Subsets, Value Model and Topology**

The vulnerability of a network (hence its critical subsets of links) depend on the model used for the value of the network and on its topology. The present example gives an illustration.

Figure 2(a) shows an example of a 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 line. In this discussion we mainly assume that the attack cost $\boldsymbol{\mu}$ is equal to zero.

Figure 2(a).a shows the critical subset corresponding to the GWA t model for which $\boldsymbol{\lambda}_{T,e} = 1 - \mathbf{1}_{e \notin T}$. 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 gain to the attacker. As a consequence, the critical subsets correspond to the set of bridges as can be observed in the figure.

For the other models, first notice that the LiV function $g(x) = f(n) - (f(x) + f(n-x))$ is increasing in the range $0 \leq x \leq n/2$ ($g(\cdot)$ is symmetric around $n/2$). This indicates that the maximum loss-in-value is achieved when a spanning tree is divided into subtrees of comparable sizes. Hence, in a network, one can expect *innermost* links to be more critical (because they belong to many spanning trees and divide each one in comparable size subtrees). This *intuition seems* to be confirmed by the examples shown in Figure 2(a).b which depicts the critical subsets with the Metcalfe, OBT, and Walrand ($a = 0.6$) models. For all three models the most critical links are the *innermost* or *core* links of the network. Reed's model, however, deviate from this observation as can be seen in Figure 2(a).c. With this model, the critical links are the links giving access to the core network. One explanation of this deviation can be found by having a closer look at the LiV function $g(x)$ whose curve is plotted in Figure 2(c). We can see that Reed's model coincides with the GWA model for a wide range of x. This indicates that links that are critical with GWA (i.e.; bridges in this case) tend to remain critical with Reed's model. But, because of the property of $g(\cdot)$ mentioned above, inner links are preferable to *outer* links in Reed's model. That is exactly what we observe in Reed's model where inner bridges constitute the critical links. These observations suggest that the critical subsets of a graph depend on the value model used to setup the game.

Connectivity is another factor that characterizes the critical subset(s). Figure 2(b) demonstrates this point in an example of network that has the same topology 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.

Example 2 **Network Improvement**

Identifying the network vulnerability and the critical subsets (the weakest points of the network) of a network is necessary for network management, risk assessment, and prioritization for link maintenance and repair. In this example, we discuss network improvement using the vulnerability metric and the notion of critical subset discussed in this paper. The graphs in Figure (3) are an illustration.

Assume that the network topology is given by the graph in Figure (3(a))–denoted 'N0'; the node S represents the gateway in Aron's model. In all other models (Walrand ($a = 0.6$), Metcalfe, BOT, Reed, and GWA), there is no designated node and all nodes are assumed to play a similar role. Table (2) (column 2) shows the vulnerability of the network for the models cited above. Figure (3(b)) shows the critical links in Aron's model while Figure (3(c)) shows the critical links for all the other models (they have the same critical subset in this particular case–but not the same vulnerability, as can be seen in Table (2) (column 2)).

Now, suppose that the network manager has an extra link to add to this network and would like to know the optimal way to do so. Figures (3(d))-(3(f)) show 3 different ways to add the link resulting to 3 different networks (respectively denoted 'N1', 'N2', 'N3'). Table (2) (columns 3-5) shows the vulnerability of those networks for the different network value models discussed in this paper. Figures (3(g))-(3(o)) show the corresponding critical subsets. Figures (3(g))-(3(i)) show the critical subsets for Aron's model; Figures (3(j))-(3(l)) those for the GWA model; and finally, Figures (3(m))-(3(o)) show the critical for Walrand, Reed, Metcalfe, and OBT models. Column 6 of Table (2) shows the comparison of the vulnerability for the 3 configurations for the different models.

The comparison shows that for all models, configuration 'N3' corresponds to the less vulnerable network. Notice that, in general, among all spanning trees, a *star* spanning tree is the most robust against the type of attack described in this paper (this is not true when attacks target nodes instead of links). This explains why configuration 'N3' (which is a *star-like* network–node S is connected to all but 1 node) is preferable to 'N1' (except for GWA model where they have the same vulnerability). Surprisingly, 'N2' is also *star-like* (in the same sense as 'N3'), however, in all models, 'N3' is strictly less vulnerable to 'N2'. Furthermore, for Reed and GWA models, network 'N1' has a smaller vulnerability than 'N2'. One reason *might be* that 'N3' and 'N1' contain more independent cliques than 'N2' (cliques that do not share a link). Indeed, each clique contains $|C|$ different *star* spanning trees, where $|C|$ is its size (i.e.; number of nodes). This variety of choices can explain why configuration 'N2' is less robust (recall here that the network manager chooses a spanning tree as communication infrastructure–hence the more independent spanning trees, the better it is). Notice that, as in the previous discussion, GWA and Reed models deviate from the other models. A better understanding of these observations requires a thorough analysis of the different model functions, which is the subject of future studies.

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 taken down 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 rest. 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 show, via numerical examples, how the vulnerability metric and the notion of critical subset can be used for network improvement. We also discuss models where efficient algorithms that can be used to compute critical subsets were derived. 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 a followup to this paper. First, in the present publication, 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 types 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 then 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.

A Minimum Disconnecting Subset (MDS)

We start by the following definition.

Definition 1 *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 $G\backslash E$ (the graph obtained by removing, from G, the edges in E) decreases its number of connected components by 1.*

Figure (4) 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 in not a minimum cut.

Let Λ be the $N \times m$ nonnegative tree-link payoff matrix whose entries are defined in (1). The polyhedron P_Λ associated with Λ is defined as the vector sum of the convex hull of its rows $(\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_N)$ and the nonnegative orthant:

$$P_\Lambda = \text{conv.hull}\,(\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_N) + \mathbb{R}_+^m. \tag{16}$$

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

$$bl(P_\Lambda) = \left\{\mathbf{y} \in \mathbb{R}_+^m : \ \Lambda\mathbf{y} \geq \mathbf{1}_{\mathcal{T}}\right\}. \tag{17}$$

Lemma 1 *Let $\boldsymbol{\omega} \in \mathrm{bl}(P_\Lambda)$ be an extreme point (vertex) of $\mathrm{bl}(P_\Lambda)$ and let $S(\boldsymbol{\omega})$ be its support (i.e.; the set of indices with positive entry). Then, $S(\boldsymbol{\omega})$ corresponds to a MDS of the graph.*

Proof: The proof works as follow. First, notice that if a subset E is not a DS, there exists a spanning tree T such that $T \cap E = \emptyset$. Hence, for any vector $\mathbf{y}$ with support E, $\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T,e)\mathbf{y}_e = \sum_{e\in E} \boldsymbol{\lambda}(T,e)\mathbf{y}_e = 0$.

As a consequence $\mathbf{y}$ cannot belong to $bl(P_\Lambda)$. So, the support of any vector in $bl(P_\Lambda)$ is a DS. Now assume that $\mathbf{y}$ is a vertex and its support $S(\mathbf{y})$ is not an MDS. Then $S(\mathbf{y})$ contains a an edge $e_o = (u, v)$ such that u and v belongs to the same connected component (say V_o) of $G \backslash S(\mathbf{y})$. Furthermore, V_o has a path p_{uv} that does not contain e_o. We have argued in Section (3.2) that since $\mathbf{y}$ is a vertex there exists a spanning tree T_o such that $\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T_o, e)\mathbf{y}_e = 1$. Now, if $\boldsymbol{\lambda}(T, e_o) = 0$ for all T such that $\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T, e)\mathbf{y}_e = 1$, we can build $\tilde{\mathbf{y}} = \mathbf{y} - \epsilon \mathbf{1}_{e_o}$, where $\mathbf{1}_{e_o}$ is the vector with all zero entries except entry e_o which is equal to 1. For ϵ chosen small enough $\tilde{\mathbf{y}}$ belongs to $bl(P_\Lambda)$ and is strictly dominated by $\mathbf{y}$; which contradicts the fact that $\mathbf{y}$ is a vertex.
If $\boldsymbol{\lambda}(T_o, e_o) > 0$ for some T_o verifying $\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T_o, e)\mathbf{y}_e = 1$, then let $T_1 = T_o - e_o + e_1$, where e_1 belongs to the path $p_{u,v}$. Notice that $\mathbf{y}_{e_1} = 0$ because e_1 does not belong to the support $S(\mathbf{y})$. We have that

$$\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T_1, e)\mathbf{y}_e = \sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T_o, e)\mathbf{y}_e - \boldsymbol{\lambda}(T_o, e_o)\mathbf{y}_{e_o} + \boldsymbol{\lambda}(T_o, e_1)\mathbf{y}_{e_1} = 1 - \boldsymbol{\lambda}(T_o, e_o)\mathbf{y}_{e_o} + \boldsymbol{\lambda}(T_o, e_1)\mathbf{y}_{e_1}.$$

But, $\boldsymbol{\lambda}(T_o, e_1)\mathbf{y}_{e_1} = 0$ and $\boldsymbol{\lambda}(T_o, e_o)\mathbf{y}_{e_o}$, hence, $\sum_{e\in\mathcal{E}} \boldsymbol{\lambda}(T_1, e)\mathbf{y}_e < 1$, which means that $\mathbf{y}$ does not belong to $bl(P_\Lambda)$; a contradiction. Hence, we get the lemma. ■

B Proof of Item 4 of Theorem 1

Note: The proof below is inspired by a pre-print by Lemonia Dritsoula, Patrick Loiseau, and John Musacchio.

For recall, for $\boldsymbol{\omega}$ a vertex of the blocker polyhedron (17) we define $\kappa(\boldsymbol{\omega}) = \sum_{e\in\mathcal{E}} \boldsymbol{\omega}(e)$ and

$$\theta(\boldsymbol{\omega}) := \frac{1}{\kappa(\boldsymbol{\omega})} \left(1 - \sum_{e\in\mathcal{E}} \boldsymbol{\omega}(e)\boldsymbol{\mu}(e)\right). \tag{18}$$

and

$$\theta^* = \max_{\boldsymbol{\omega}} \theta(\boldsymbol{\omega}) \ . \tag{19}$$

We would like to show that if $\boldsymbol{\beta}$ is a Nash equilibrium strategy for the attacker, than the associated attacker's payoff is equal to θ^*. For this, we will first argue that a Nash equilibrium strategy $\boldsymbol{\beta}$

To understand this, first recall from Item 2 of the theorem that there exists a Nash equilibrium strategy of the attacker that is derived from the vertices $\boldsymbol{\omega}$ of the blocker polyhedron. The attacker's payoff associated to this strategy is θ^*. Now, consider any attacker's strategy $\boldsymbol{\beta}$. Let $\theta(\boldsymbol{\beta})$ be the payoff associated to $\boldsymbol{\beta}$

when the defender plays her best response. Since the defender wants to minimize the expected loss given in (3),

$$\theta(\boldsymbol{\beta}) = \min[\Lambda\boldsymbol{\beta}] - \boldsymbol{\mu}'\boldsymbol{\beta}, \tag{20}$$

where the minimum above is taken on the entries of the vector $\Lambda\boldsymbol{\beta}$.

Now, let us recall the notion of *best response polyhedron* introduced in [2]. The best response polyhedron for the defender is defined as the set of her mixed strategies with the *upper envelope* of expected payoffs (and any larger payoff) ν to the attacker.

$$\left\{(\mathbf{x}, \nu) \in \mathbb{R}^N \times \mathbb{R} \mid \mathbf{x} \geq 0,\ \mathbf{1}'\mathbf{x} = 1,\ (\Lambda - U)'\,\mathbf{x} \leq \nu\mathbf{1}_m\right\}, \tag{21}$$

where the matrix U has the same dimensions as Λ and has each row equal to $\boldsymbol{\mu}'$. Notice that, since $\mathbf{x}$ is assumed to be a probability distribution, $U'\mathbf{x} = \boldsymbol{\mu}$. Thus, the equation above can be written as

$$\left\{(\mathbf{x}, \nu) \in \mathbb{R}^N \times \mathbb{R} \mid \mathbf{x} \geq 0,\ \mathbf{1}'\mathbf{x} = 1,\ \Lambda'\mathbf{x} - \boldsymbol{\mu} \leq \theta(\boldsymbol{\beta})\mathbf{1}_m\right\}, \tag{22}$$

where in the above equation, we have replaced the attacker's expected payoff ν by its value $\theta(\boldsymbol{\beta})$.

The defender's best response is a solution of the following linear program:

$$\begin{aligned} &\text{Minimize}_{\boldsymbol{\alpha}} \boldsymbol{\alpha}'\Lambda\boldsymbol{\beta} \\ &\text{subject to } \Lambda'\boldsymbol{\alpha} - \mu \leq \theta(\boldsymbol{\beta})\mathbf{1}_m \\ &\qquad\qquad \boldsymbol{\alpha} \geq \mathbf{0},\ \mathbf{1}'\boldsymbol{\alpha} = 1 \end{aligned} \tag{23}$$

The dual of this LP can be written as

$$\begin{aligned} &\text{Maximize }_{\mathbf{y},z} \left(-\mathbf{1}_m\theta(\boldsymbol{\beta}) - \mu\right)'\mathbf{y} + z \\ &\text{subject to } z\mathbf{1}_{\mathcal{T}} - \Lambda\mathbf{y} \leq \Lambda\boldsymbol{\beta} \\ &\qquad\qquad \mathbf{y} \geq \mathbf{0},\ z \geq 0 \end{aligned} \tag{24}$$

Looking at the constraints of $z\mathbf{1}_{\mathcal{T}} \leq \Lambda(\boldsymbol{\beta} + \mathbf{y})$ 24, we see that the optimal value for z should be chosen as

$z = min[\Lambda(\boldsymbol{\beta} + \mathbf{y})]$. Replacing this back to the dual LP, we get

$$\begin{aligned} \text{Maximize }_{\mathbf{y}} & -\mathbf{1}'_m \mathbf{y}\theta(\boldsymbol{\beta}) - \mu'\mathbf{y} + min[\Lambda(\boldsymbol{\beta} + \mathbf{y})] \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{25}$$

Next, we show that

Lemma 2 *If $\boldsymbol{\beta}$ is a Nash equilibrium strategy, than the LP defined in (23) is feasible.*

This must be true because if $\boldsymbol{\beta}$ is a NE strategy, than there exists a corresponding defender's strategy which is a best response to $\boldsymbol{\beta}$. Thus, the best response polyhedron is not empty.

Now, we have all the ingredients to prove Item 4 of the theorem, which we re-state here (in slightly different words) as a lemma for the reader's interest.

Lemma 3 *If $\boldsymbol{\beta}$ is a NE strategy for the attacker, then $\theta(\boldsymbol{\beta})$ is maximum and is equal to θ^*.*

Notice that, since we know (by Item 2) that there exists at least one NE strategy derived from the vertices of the blocker polyhedron and it gives a payoff of θ^*, if the lemma is true, then $\theta(\boldsymbol{\beta}) = \theta^*$.

Proof: Suppose that the attacker's strategy is such that $\theta(\boldsymbol{\beta}) < \theta(\xi)$ where $\xi = \arg\max_{\tilde{\boldsymbol{\beta}}} \theta(\tilde{\boldsymbol{\beta}})$, is a NE strategy ($\mathbf{1}'_m\xi = 1$). Let $\mathbf{y} = \tau\xi$ for $\tau \gg 1$ chosen large enough. Then, (25) becomes

$$\text{Maximize }_{\tau} - \tau\mathbf{1}'_m\xi\theta(\boldsymbol{\beta}) - \tau\mu'\xi + \min[\Lambda(\boldsymbol{\beta} + \tau\xi)] \tag{26}$$

Notice that $\Lambda\boldsymbol{\beta} > \mathbf{0}$. In fact, if this vector has a zero entry, then $\min[\Lambda\boldsymbol{\beta}] = 0$ and $\theta(\boldsymbol{\beta}) < 0$ which means that $\boldsymbol{\beta}$ cannot be a NE strategy. Similarly $\Lambda\xi > \mathbf{0}$ Hence, for large $\tau \gg 1$, only the entries of $\Lambda(\tau\xi)$ are relevant to find the minimum in the last term of (26). In other terms,

$$min[\Lambda(\boldsymbol{\beta} + \tau\xi)] = \min[\Lambda(\tau\xi)] + C^{st} = \tau\min[\Lambda\xi] + C^{st}, \tag{27}$$

where the constant above is the value of $\Lambda\boldsymbol{\beta}$ at the entry with minimum value in $\Lambda\xi$. The objective function in (26) can be re-written as

$$
\begin{aligned}
-\tau\mathbf{1}_m'\xi\theta(\boldsymbol{\beta}) - \tau\mu'\xi + \tau\min[\Lambda\xi] + C^{st} &= -\tau\mathbf{1}_m'\xi\theta(\boldsymbol{\beta}) - \tau\mu'\xi + \tau\min[\Lambda\xi] + C^{st} \\
&= -\tau\theta(\boldsymbol{\beta}) + \tau\left(\tau\min[\Lambda\xi] - \mu'\xi\right) + C^{st} \qquad (28) \\
&= \tau\left(\theta(\xi) - \theta(\boldsymbol{\beta})\right) + C^{st}, \qquad (29)
\end{aligned}
$$

where in (28) we use the fact that $\mathbf{1}_m'\xi = 1$ and in (29), we use the definition in (20).

Now, since $\theta(\xi) > \theta(\boldsymbol{\beta})$, (29) can be made as large as possible. This means that the dual (24) is unbounded, which implies that the primal (23) is infeasible. But this contradicts Lemma 2 (with the assumption that $\boldsymbol{\beta}$ is a NE strategy). Thus, $\theta(\boldsymbol{\beta})$ must be maximal. ■

This results also explains why in Item 1 of the theorem, the attacker never launches an attack when $\theta^* < 0$. Since any NE will give θ^* and not attacking always guarantees 0, the attacker is better off not attacking when $\theta^* < 0$.

References

[1] USN Admiral James Stavridis. Channeling David Sarnoff, Sept 2006. http://www.aco.nato.int/saceur/channeling-david-sarnoff.aspx.

[2] David Avis, Gabriel Rosenberg, Rahul Savani, and Bernhard von Stengel. Enumeration of Nash Equilibria for Two-Player Games. *Economic Theory*, 42:9–37, 2010.

[3] Andrew Odlyzko Bob Briscoe and Benjamin Tilly. Metcalfe's Law is Wrong. *IEEE Spectrum*, pages 26–31, July 2006.

[4] William H. Cunningham. Optimal Attack and Reinforcement of a Network. *J. ACM*, 32(3):549–561, 1985.

[5] L. Freeman. Centrality in Social Networks Conceptual Clarification. *Social Networks*, 1(3):215–239, 1979.

[6] D R Fulkerson. Blocking and Anti-Blocking Pairs of Polyhedra. *Math. Programming*, (1):168–194, 1971.

[7] George Gilder. Metcale's Law And Legacy, Nov 1995. http://www.seas.upenn.edu/~gaj1/metgg.html.

[8] Assane Gueye. *A Game Theoretical Approach to Communication Security*. PhD dissertation, University of California, Berkeley, Electrical Engineering and Computer Sciences, March 2011.

[9] Assane Gueye, Vladimir Marbukh, and Jean C. Walrand. Towards a Quantification of Communication Network Vulnerability to Attacks: A Game Theoretic Approach. Technical report, National Institute of Standards and Technology, December 2011. http://www.nist.gov/itl/math/cctg/assane.cfm.

[10] Assane Gueye, Vladimir Marbukh, and Jean C. Walrand. Towards a Quantification of Communication Network Vulnerability to Attacks: A Game Theoretic Approach. In *3rd International ICST Conference on Game Theory for Networks*, Vancouver, Canada, May 2012.

[11] 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.

[12] 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.

[13] D Gusfield. Connectivity and Edge-Disjoint Spanning Trees. *Information Processing Letters*, (16):87–89, 1983.

[14] Erik Jenelius, Tom Petersen, and Lars-Gran Mattsson. Importance and exposure in road network vulnerability analysis. *Transportation Research Part A: Policy and Practice*, 40(7):537–560, August 2006.

[15] D Fulkerson L Ford. Flows in Networks. pages 453–460. Princeton Univ. Press, 1962.

[16] Aron Laszka, David Szeszlér, and Levente Buttyán. Game-theoretic Robustness of Many-to-one Networks. In *3rd International ICST Conference on Game Theory for Networks*, Vancouver, Canada, May 2012.

[17] Mohammadhossein Manshaei, Quanyan Zhu, Tansu Alpcan, Tamer Basar, and Jean-Pierre Hubaux. Game Theory Meets Network Security and Privacy. Technical report, EPFL, Lausanne, 2010.

[18] Deepankar Medhi. *Network Reliability and Fault-Tolerance.* John Wiley & Sons, Inc., 2007.

[19] J A Nash-Williams. Edge-Disjoint Spanning Trees of Finite Graphs. *Journal London Math. Soc*, (36):445–450, 1961.

[20] Tore Opsahl, Filip Agneessens, and John Skvoretz. Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths. *Social Networks*, 32(3):245 – 251, 2010.

[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] 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.

[23] W T Tutte. On the Problem of Decomposing a Graph into N Connected Factors. *Journal of the London Mathematical Society*, (36):221–230, 1961.

[24] D. J. Watts and S. H. Strogatz. Collective dynamics of'small-world'networks. *Nature*, 393(6684):409–10, 1998.

[25] Douglas R. White and Stephen P. Borgatti. Betweenness centrality measures for directed graphs. *Social Networks*, 16(4):335 – 346, 1994.

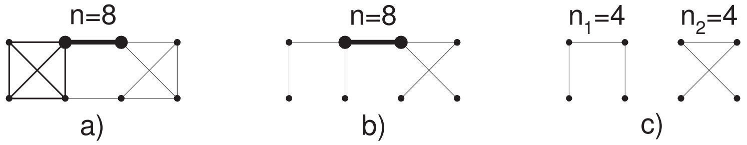

Figure 1: Tree-Link model for the value of a network link. a) Complete network with link 'e' of interest shown in bold. b) A particular spanning tree 'T' of the graph. c) Disconnected network when link 'e' is removed.

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	$1_{e\in T}$
Metcalfe	$1 - \frac{n_1^2+n_2^2}{n^2}$
Reed	$1 - 2^{-n_1} - 2^{-n_2}$
OBT	$1 - \frac{n_1\log(n_1)+n_2\log(n_2)}{n\log(n)}$
Walrand	$1 - \frac{n_1^{1+a}+n_2^{1+a}}{n^{1+a}}$
Aron	$1 - \frac{n_1\mathbf{1}_{S\in V_1}+n_2\mathbf{1}_{S\in V_2}}{n}$

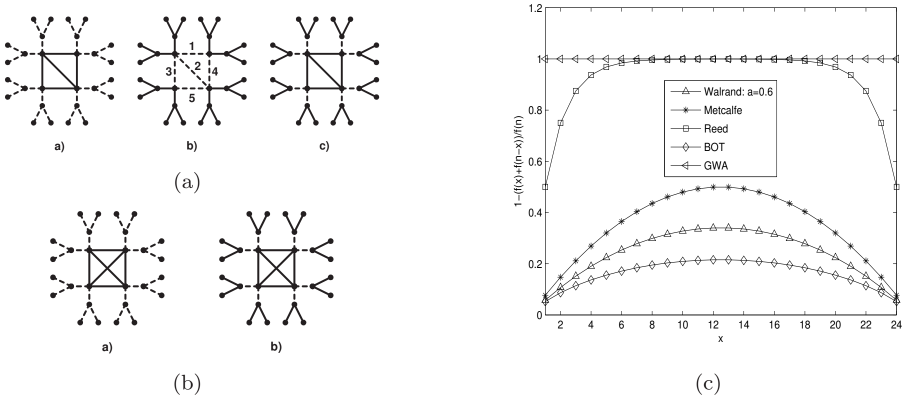

(a)

(b)

(c)

Figure 2: Critical subsets. 1-Figure 2(a): dependency on network value models: a) GWA b) OBT, Walrand, and Metcalfe c) Reed. 2-Figure 2(b): dependency on connectivity: a) GWA b) OBT, Walrand, and Metcalfe c) Reed. 3-Figure 2(c):Comparison of the loss functions $1 - \frac{f(x)+f(n-x)}{f(n)}$.

Table 2: Network vulnerability.

Model	N0	N1	N2	N3	vulnerability
Aron	3/(2x7)	3/(2x7)	6/(5x7)	6/(5x7)	$N3 = N2 < N1$
Walrand	0.1853	0.1669	0.1461	0.1321	$N3 < N2 < N1$
Metcalfe	0.2682	0.2449	0.2099	0.1909	$N3 < N2 < N1$
OBT	0.2049	0.1755	0.1659	0.1494	$N3 < N2 < N1$
Reed	0.5752	0.4689	0.4775	0.4268	$N3 < N1 < N2$
GWA	3/4	3/5	2/3	3/5	$N3 = N1 < N2$

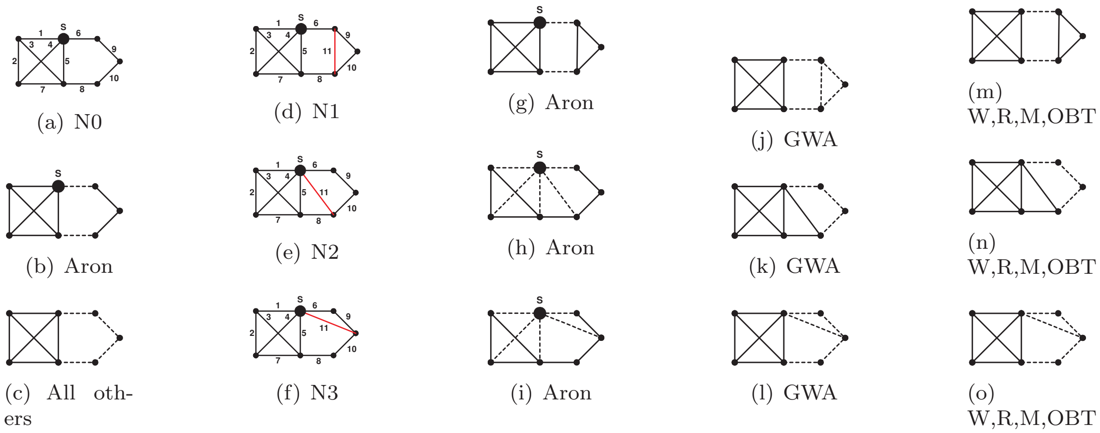

(a) N0

(b) Aron

(c) All others

(d) N1

(e) N2

(f) N3

(g) Aron

(h) Aron

(i) Aron

(j) GWA

(k) GWA

(l) GWA

(m) W,R,M,OBT

(n) W,R,M,OBT

(o) W,R,M,OBT

Figure 3: Network improvement example.

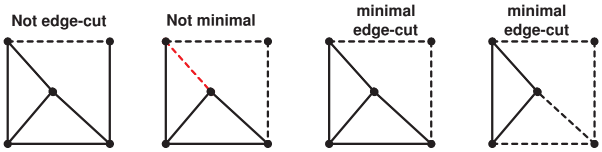

Figure 4: Examples of disconnecting sets (DS) and minimal disconnecting sets (MDS). The chosen subset is shown in dashed line. The leftmost example is not a DS. The next example is a DS but is not a MDS because the red link connect two nodes belonging to the same connected component. The third example 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.