

Quantifying All-to-One Network Topology Robustness Under Budget Constraints*

[Extended Abstract]

Aron Laszka

Laboratory of Cryptography and System Security
Budapest University of Technology and
Economics

laszka@crysys.hu

Assane Gueye

National Institute of Standards and Technology
(NIST)

Gaithersburg, MD, USA

assane.gueye@nist.gov

ABSTRACT

To design robust network topologies that resist strategic attacks, one must first be able to quantify *robustness*. In a recent line of research, the theory of *network blocking games* has been used to derive robustness metrics for topologies. However, these previous works did not consider the budget constraints of the network operator. In this paper, we introduce a budget limit on the operator and study two budget constraint formulations: the maximum and the expected cost constraints. For practical applications, the greatest challenge posed by blocking games is their computational complexity. Therefore, we show that the expected cost constraint formulation leads to games that can be solved efficiently, while the maximum cost constraint leads to NP-hard problems. As an illustrative example, this paper discusses the particular case of All-to-One (e.g., sensor or access) networks.

Keywords

network topology robustness, robustness metrics, game theory, blocking games, computational complexity

1. INTRODUCTION

Designing network topologies that are robust and resilient to attacks has been and continues to be an important and challenging topic in the area of communication networks. One of the main difficulties resides in quantifying the robustness of a network in the presence of an intelligent attacker who might exploit the structure of the network topology to design harmful attacks. Quantifying the robustness or, equivalently, the vulnerability of topologies has been extensively studied (e.g., [6, 11, 2]); however, the simultaneous and strategic decision making of the defender and the adversary, which is key to the security of information systems, has received only little attention.

To capture the strategic nature of the interactions between a defender and an adversary, game theoretic models have been gaining a lot of interest. In a recent line of research ([4],[5],[9],[8],[3]), *network blocking games* (NBGs) have been introduced and applied to the analysis of the robustness of network topologies. An NBG takes as input the communication model (e.g., the All-to-One model for modeling access

and sensor networks) and the topology of a network and models the strategic interactions between an adversary and the network operator as a two-player game. The Nash equilibrium strategies are then used to predict the most likely attacker's actions and the attacker's Nash equilibrium payoff¹ serves as a quantification of the vulnerability (inverse robustness) of the network. As a consequence, being able to efficiently compute a NE is crucial for NBG models.

Zero-sum, two-player games can be cast as linear programs and, hence, can be solved "efficiently" using linear programming tools, provided that the game is of reasonable size. In the case of NBG models, however, the games are generally exponential in size, which makes them challenging to deal with. In the series of NBG papers cited above, new algorithms have been developed that compute a Nash equilibrium *efficiently* for a number of communication models: All-to-One (e.g., access and sensor) networks [8], All-to-All (e.g., Ethernet) networks [4, 9], and Supply-Demand networks [3]. These algorithms are based on the theory of network flows or, for some models, on the minimization of submodular functions.

One implicit assumption of the NBG model is that the defender (i.e., the network operator) can use the network resources (links) at zero cost. However, in reality, network links have positive usage costs (e.g., operation or protection costs) and, in general, these costs are nonuniform. Furthermore, network operators do not have an unlimited budget, which would allow them to use any combination of network resources. In [3], a cost of security as well as a budget constraint have been introduced for the particular case of Supply-Demand (S-D) networks. The budget constraint limits the operator to using only those sets of resources (links) whose associated costs do not exceed her budget.

In the present paper, we extend the budget constraint idea and provide complexity results with regard to the computation of the equilibrium payoff. Recall that the aim of solving these models is to derive a quantification of the robustness of the network in the presence of a strategic adversary, and the equilibrium payoff is used as such a quantification. Thus, being able to efficiently compute the equilibrium payoff is of central importance in these models. As an illustrative example, we discuss the particular case of the All-to-One communication model. The main contributions of this paper can be summarized as follows:

- We generalize the (All-to-One) network blocking game

¹It has been shown that the payoffs are the same in every equilibrium of an NBG; thus, it suffices to find a single equilibrium in order to characterize the robustness of a network.

*U.S. Government work not protected by U.S. copyright.

model by introducing the maximum cost constraint, which is based on a similar constraint previously proposed for Supply-Demand networks, and a novel constraint formulation, called the expected cost constraint.

- We show that (in the All-to-One model) the problem of determining the equilibrium payoff is NP-hard under the maximum cost constraint, but can be solved in polynomial time under the expected cost constraint. To the best of our knowledge, these are the first computational complexity results for any budget constrained network blocking game model.

2. NETWORK BLOCKING GAMES

The concept of network blocking game was first introduced in [4]. Here, we discuss it in the context of the All-to-One communication model, which we introduce next.

2.1 All-to-One Communication Model

In an *All-to-One* network [8], the primary goal of the network manager is to enable all nodes to communicate with a designated node r . This models sensor and access networks, where all of the nodes are trying to reach a gateway or data collection node (or, alternatively, a set of nodes, which can be modeled by a super-designated node).

The topology of the network is represented by a connected simple graph $G = (V, E)$ with node set V , link set E , and designated node $r \in V(G)$. To get all nodes connected to r , the network manager chooses a subset of links T that forms a spanning tree. In practice, this spanning tree could be implemented, for example, as the next-hop forwarding table entries for r , which are stored at the individual nodes of the network. We let \mathcal{T} designate the set of all spanning trees.

Let the network be connected using a spanning tree T . If a given link $e \in E$ fails (because of the action of an attacker), some nodes might no longer be able to communicate with r and, thus, can be considered as lost for the network operator. We let $\lambda(T, e)$ designate the number of those nodes that are disconnected from r . Notice that, if $e \notin T$, then $\lambda(T, e) = 0$. In the next section, we use this loss to define the payoffs in our game model.

2.2 Game-Theoretic Measure of Robustness

The game is played on the topology of the network by a defender (the network operator or manager) and a strategic attacker. The operator wants to guarantee that all nodes can communicate with r . For this, she chooses a spanning tree $T \in \mathcal{T}$ (i.e., her strategy space is the set \mathcal{T} of all spanning trees). At the same time, a strategic and malicious adversary is trying to disrupt the communication by removing a link (i.e., her strategy space is the set E of links in the network). We model this attacker-defender “interaction” as a zero-sum, two-player game where the payoffs are defined as follows: when the operator picks spanning tree T and the attacker targets link e , the defender loses $\lambda(T, e)$ (as defined above), which goes as a payoff to the attacker.

We consider mixed strategy Nash equilibria, where the network operator chooses a distribution (denoted by α) over the set \mathcal{T} and the attacker chooses a distribution (denoted by β) over the set E . It is assumed that the operator tries to minimize her *expected loss* (i.e., negative payoff), while the attacker tries to maximize her *expected payoff*. Formally, the operator chooses α to minimize $L(\alpha, \beta)$, while the attacker chooses β to maximize $L(\alpha, \beta)$, where

$$L(\alpha, \beta) = \sum_{T \in \mathcal{T}} \sum_{e \in E} \alpha_T \beta_e \lambda(T, e). \quad (1)$$

In the analysis of the general NBG, it has been shown that the equilibrium expected loss $L(\alpha^*, \beta^*)$ is a property of (i.e., solely determined by) the topology of the network and the communication model (which, in this paper, is considered to be the All-to-One model). A low $L(\alpha^*, \beta^*)$ indicates that operating the network has low expected loss due to attack, that is, the network is robust against attacks. If, on the other hand, $L(\alpha^*, \beta^*)$ is high, then the expected loss is also high, and the network can be considered vulnerable. As such, $L(\alpha^*, \beta^*)$ can be used as a measure of network topology vulnerability (i.e., inverse robustness). It is noteworthy that, for the All-to-One model, this measure was shown to be the inverse of the *persistence* of the graph of the network [8], which had been previously proposed in [1] as a metric for graph robustness in a non-game theoretic framework.

The model above discusses scenarios where the defender can choose from the set of spanning trees at zero cost. For such (unconstrained) models, efficient algorithms have been derived to compute a NE for multiple NBG models (as stated earlier). In the following section, we introduce a budget constraint on the defender and, then, analyze the complexity of computing an NE in Section 4.

3. BUDGET CONSTRAINT

In the basic NBG, the operator is only interested in minimizing her expected loss due to attack, without taking her operating costs into account. In practice, however, network operators have to take economic goals and constraints into consideration when deciding their strategies. These economic decisions are affected by the topology of the network as links (and, hence, spanning trees) can have varying (unit) costs of usage.

3.1 Unit Usage / Protection Cost

In [3], a (per unit) usage cost model was introduced and discussed for the particular case of the S-D communication model. Here, we extend this cost model to the All-to-One communication model. First, recall that, in the All-to-One model, $\lambda(T, e)$ is the number of nodes that communicate with r through e . Thus, $\lambda(T, e)$ is proportional to the traffic on e and, hence, to the usage of link e . We assume that each link is associated with some unit usage cost $w(e)$, so that using link e has a *net* per link cost of $w(e)\lambda(T, e)$ to the operator. With this definition, the total cost of a spanning tree T is

$$w(T) := \sum_{e \in E} \lambda(T, e) w(e). \quad (2)$$

If the spanning tree T is chosen according to some distribution α , then we define the expected usage cost as

$$w(\alpha) := \sum_{T \in \mathcal{T}} \alpha_T w(T) = \sum_{e \in E} w(e) \sum_{T \in \mathcal{T}} \alpha_T \lambda(T, e). \quad (3)$$

We assume that, to operate the network, the operator has a fixed *budget* $b \in \mathbb{R}_{\geq 0}$ to spend. Therefore, her objective is to minimize her expected loss (see Equation 1) by choosing a strategy such that the constraint posed by the given budget is satisfied. The budget constraint can be formulated in multiple ways. In the following sections, we introduce and

study two straightforward formulations, the maximum and the expected (or average) cost budget constraints.

3.2 Maximum Cost Budget Constraint

In our first budget constraint formulation, which is referred to as *maximum cost constraint* (MCC), we require that, for a given budget b , the operator can use a spanning tree T only if its total cost (in Equation 2) is less than or equal to b . Formally, the pure-strategy set of the operator is restricted to

$$\mathcal{T}^{(b)} = \{T \in \mathcal{T} \mid w(T) \leq b\} . \quad (4)$$

The maximum cost constraint is best-suited for budget limits that are determined by the amount of preallocated resources available. In this case, the cost of a link can be the amount of resources needed (e.g., energy consumption) to operate the link and the budget limit is the amount of resources available (e.g., amount of power available).

3.3 Expected Cost Budget Constraint

The maximum cost constraint misses to capture certain situations. For instance, when the amount of allocated resources can be modified during operation, e.g., resources can be leased, the budget limit applies to the average or, equivalently, the expected cost of a strategy during continuous operation. Thus, in our second budget constraint formulation, which we will refer to as the *expected cost constraint* (ECC), we only require the expected (or average) cost of the operator not to exceed the budget limit.

Under the *expected cost constraint* with a budget limit b , the operator can employ a mixed strategy only if its expected cost (in Equation 3) is less than or equal to b . Formally, the set of mixed strategies available to the operator is

$$\mathcal{A}^{(b)} = \left\{ \alpha \in \mathbb{R}^{|\mathcal{T}|} \mid w(\alpha) \leq b \right\} . \quad (5)$$

Note that the above formulation generalizes the classic notion of mixed strategies in game-theory, where the set of mixed strategies is always the set of *all* distributions over the set of pure strategies. Here, a mixed strategy is chosen from a predefined subset of distributions.

3.4 Constrained Game

Having defined the set of available strategies (pure for MCC and mixed for ECC), we can now setup the constrained game (in a similar way as presented in Subsection 2.2). We are interested in mixed strategy Nash equilibria where the operator picks a distribution α over $\mathcal{T}^{(b)}$ (for MCC) or from the set $\mathcal{A}^{(b)}$ (for ECC), while the attacker chooses a distribution β over the set of links. The Nash equilibrium payoff is denoted $L^{(b)}(\alpha^*, \beta^*)$ for a game with budget limit b .

Using the same interpretation as in Subsection 2.2, the NE payoff $L^{(b)}(\alpha^*, \beta^*)$ can be used to quantify the vulnerability (i.e., inverse robustness) of the network at budget limit b , which is the minimum vulnerability that the defender can achieve when her budget is b . By varying b , one can draw the Pareto frontier between the region of achievable vulnerability/budget points and the region of unachievable ones, as was done in [3] for the particular case of S-D networks with the maximum cost constraint. In the next two sections, we discuss the complexity of computing $L^{(b)}(\alpha^*, \beta^*)$.

4. COMPUTATIONAL COMPLEXITY

Recall that the game considered in this paper (and generally in NBGs) is exponential in size (the number of spanning trees grows exponentially in the number of nodes). As a consequence, usual solution techniques (such as linear programs) become impractical. However, it has been shown that an equilibrium can be computed efficiently for a number of communication models ([4], [9], [8], and [3]) in the unconstrained game. The next theorem gives the computational complexity for the constrained game.

THEOREM 1. *Computing the NE payoff for the All-to-One communication model is*

- *NP-hard under a maximum cost budget constraint and*
- *can be solved in polynomial time under an expected cost budget constraint.*

In the following two subsections, we give an outline of the proof. The complete proof can be found in [7].

4.1 Proof: Maximum Cost Budget Constraint

We show NP-hardness by reducing a well-known NP-hard problem, the *Partition Problem* (PP) [10], to the problem of deciding whether the equilibrium payoff in a given network is greater than a certain value. We refer to this problem as the *Equilibrium Problem* (EP).

First, we construct an instance of EP (i.e., a network topology, a budget limit, and an equilibrium payoff value) from an instance of PP (i.e., a multiset of positive integers $\{x_1, \dots, x_n\}$ that has to be partitioned into two subsets of equal size) in polynomial time as follows:

- Let the topology be the following: The designated node r is connected with zero cost edges to $2n$ nodes, denoted by $1_a, 1_b, \dots, n_a$ and n_b , in the form of a large star rooted at r . Additionally, there are n “outer” nodes, denoted by $1, \dots, n$. Node i is connected to nodes i_a and i_b with edges having costs of x_i and 0.
- Let the budget be $b = \frac{1}{2} \sum_{i=1}^n x_i$.
- Finally, let the equilibrium payoff value be $\frac{3}{2}$.

Second, we show that the equilibrium payoff in the above network is greater than $\frac{3}{2}$ iff the PP does not have a solution.

4.2 Proof: Expected Cost Budget Constraint

In order to overcome the computational complexity caused by the exponential size of the operator’s pure strategy space, we use a network flow based characterizations of the operator’s mixed strategy space, which was introduced in [8].

Let L_{opt} denote the value of the following LP:

$$\text{Maximize } -L \quad (6)$$

$$\text{subject to } \sum_{(u,v) \in E} f(u,v)w(u,v) \leq b \quad (7)$$

$$\forall (u,v) \in E : f(u,v) \leq L \quad (8)$$

$$\forall v \in V \setminus \{r\} : \sum_{(u,v) \in E} f(u,v) - \sum_{(v,w) \in E} f(v,w) \leq -1 , \quad (9)$$

where $L \in \mathbb{R}_{\geq 0}$ and $\forall (u,v) \in E : f(u,v) \in \mathbb{R}_{\geq 0}$.

First, we show that an operator strategy that achieves at most L_{opt} loss (regardless of the adversary’s strategy) can be computed in polynomial time. Second, we show that an adversarial strategy that achieves at least L_{opt} payoff (regardless of the operator’s strategy) can be computed in polynomial time. Finally, by combining these results, we have that the equilibrium payoff is L_{opt} .

5. APPLICATION EXAMPLE

In this section, we present a practical example, and show how the choice of budget constraint can impact an operator’s investment decisions. The example compares the two budget constraint formulations to each other and to the unconstrained game. Note that, because of the small size of the example, the computational complexity of the maximum cost constraint is not an issue.

Assume that we have to operate a small sensor network, whose topology is shown in Figure 1. The unit cost of each link is assumed to be 1. To run the network, we are given a budget of $b = 12$ units. Notice that this budget is large enough so that, for the original topology, the robustness is the same in the constrained and the unconstrained games (as can be seen in the first row of Table 1).




Figure 1: The topology of the example network. The designated node is represented by the shaded disk. Original links are represented by solid lines.

Now, assume that we are given the opportunity of adding an additional link to make the network topology more robust. To simplify the example, three options are compared to each other: link e_1 , whose unit cost is 4, link e_2 , whose unit cost is 2, and link e_3 , whose unit cost is 1.5.

Table 1: Topology Vulnerability for Different Options and Constraints

	Unconstrained	Maximum Cost Const.	Expected Cost Const.
Original	2.5	2.5	2.5
e_1 added	1.67	2	2
e_2 added	1.75	2	1.86
e_3 added	1.75	1.9	1.9

Table 1 shows the value of the equilibrium payoff, which is a metric of vulnerability (i.e., inverse robustness), for each option in the unconstrained, the maximum cost constrained, and the expected cost constrained models. The optimal option (i.e., the lowest value) for each constraint is marked in bold. The table shows that the optimal option is different for each formulation: In the unconstrained game, the long-range high-cost link is optimal. Under the maximum cost constraint, which is the most restrictive, the cheapest link has to be chosen. Finally, under the expected cost constraint, the medium-cost link is the best as it is a good trade-off between cost and decrease in vulnerability.

6. CONCLUSION & FUTURE WORK

In this paper, we have generalized the All-to-One NBG game by introducing budget constraints on the operator. As the greatest challenge to computing the equilibrium (i.e., the vulnerability of the topology) in practice is the exponential size of the payoff matrix, we have focused our work on computational complexity: we have shown that the maximum

cost formulation leads to NP-hard problems and proposed an efficient algorithm for the expected cost formulation.

Proving that the maximum cost formulation leads to NP-hard problems was a very important first step. Since we now know that no polynomial-time algorithm can solve the game under the MCC, an interesting future work is finding polynomial-time approximation algorithms or efficient heuristics. Another interesting future direction is the study of the cost-security tradeoff problem, where the operator has to maximize security and minimize budget at the same time.

Acknowledgement

This paper has been supported by HSN Lab, Budapest University of Technology and Economics², and by the NIST / UMD American Recovery and Reinvest Act (NIST-ARRA).

7. REFERENCES

- [1] W. Cunningham. Optimal attack and reinforcement of a network. *Journal of the ACM*, 32(3):549–561, 1985.
- [2] T. Grubestic, T. Matisziw, A. Murray, and D. Snediker. Comparative approaches for assessing network vulnerability. *International Regional Science Review*, 31(1):88–112, 2008.
- [3] A. Gueye and V. Marbukh. A game-theoretic framework for network security vulnerability assessment and mitigation. In *Proc. of 3rd Conf. on Decision and Game Theory for Security*. Springer, 2012.
- [4] A. Gueye, J. C. Walrand, and V. Anantharam. Design of network topology in an adversarial environment. In *Proc. of 1st Conf. on Decision and Game Theory for Security*. Springer, 2010.
- [5] A. Gueye, J. C. Walrand, and V. Anantharam. A network topology design game: How to choose communication links in an adversarial environment? In *Proc. of 2nd Int. Conf. on Game Theory for Networks*. Springer, 2011.
- [6] P. Holme, B. Kim, C. Yoon, and S. Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.
- [7] A. Laszka and A. Gueye. Quantifying network topology robustness under budget constraints [Work in Progress]. http://www.crysys.hu/~laszka/papers/quantifying_robustness.pdf.
- [8] A. Laszka, D. Szeszlér, and L. Buttyán. Game-theoretic robustness of many-to-one networks. In *Proc. of 3rd Int. Conf. on Game Theory for Networks*. Springer, 2012.
- [9] A. Laszka, D. Szeszlér, and L. Buttyán. Linear loss function for the network blocking game: An efficient model for measuring network robustness and link criticality. In *Proc. of 3rd Conf. on Decision and Game Theory for Security*. Springer, 2012.
- [10] S. Mertens. The easiest hard problem: Number partitioning. *Computational Complexity and Statistical Physics*, 125(2):125–139, 2006.
- [11] C. Schneider, A. Moreira, J. Andrade Jr, S. Havlin, and H. Herrmann. Mitigation of malicious attacks on networks. *Proceedings of the National Academy of Sciences*, 108(10):3838–3841, 2011.

²<http://www.hsnlab.hu>