

Stackelberg Equilibria for Two-Player Network Routing Games on Parallel Networks

David Grimsman, João P. Hespanha and Jason R. Marden

Abstract—We consider a two-player zero-sum network routing game in which a router wants to maximize the amount of legitimate traffic that flows from a given source node to a destination node and an attacker wants to block as much legitimate traffic as possible by flooding the network with malicious traffic. We address scenarios with asymmetric information, in which the router must reveal its policy before the attacker decides how to distribute the malicious traffic among the network links, which is naturally modeled by the notion of Stackelberg equilibria. The paper focuses on parallel networks, and includes three main contributions: we show that computing the optimal attack policy against a given routing policy is an NP-hard problem; we establish conditions under which the Stackelberg equilibria lead to no regret; and we provide a metric that can be used to quantify how uncertainty about the attacker’s capabilities limits the router’s performance.

I. INTRODUCTION

This paper addresses a network routing game between a player that wants to route legitimate traffic from a source node to a destination node and another player that wants to block traffic by flooding the network with malicious traffic. We refer to these players as the *router* and the *attacker*. Motivated by network security problems, we are interested in scenarios of asymmetric information, where the router exposes its policy to the attacker before the attacker needs to select its policy. The problem formulation considered here is motivated by the so-called Crossfire attack in which an attacker persistently degrades network connectivity by targeting a selected set of links within the network, while adjusting to changes in routing policies [1]. The defense against such attacks has been the subject of recent work [2], [3], [4], [5].

The Nash equilibrium is an attractive solution concept for noncooperative games because it leads to very strong notions of equilibria, in that neither player regrets its choice after the outcome of the game is revealed [6]. However, such equilibria often do not exist in problems of asymmetric information. The Stackelberg equilibrium is an alternative solution concept where one player (the leader) must select and reveal its policy before the other player (the follower) makes a decision [7]. This type of equilibrium specifically addresses the information asymmetry that we consider here

and has been applied to domains closely related to the problem considered in this paper, including network routing [8], scheduling [9], and channel allocation for cognitive radios [10], but also has application in supply chain and marketing channels [11] among other fields. The Stackelberg equilibrium is a concept that is also well-suited for security of critical infrastructure systems [12] and has been applied to surveillance problems that include the ARMOR program at the Los Angeles International Airport [13], the IRIS program used by the US Federal Air Marshals [14], power grid security [15], and defending oil reserves [12]. These two types of equilibria have also been studied extensively for various types of security games [16].

This paper includes three main contributions:

- 1) Theorem 1 establishes that computing the attacker’s optimal response to a routing policy is an NP-hard problem.
- 2) Theorem 2 determines conditions on the network under which Stackelberg equilibria lead to no-regret policies (i.e., are also Nash).
- 3) Section IV explores how uncertainty in knowledge about the capabilities of the attacker translates into performance loss for the router. Theorem 3 provides a closed-form expression which quantifies this for a two-link network.

We focus on a network consisting solely of N parallel links that directly connect source and destination. Even within this simple set of networks, the computation of the optimal attack policy turns out to have higher complexity than one might expect. For any fixed routing policy, we show in Section III that the computation of the “optimal” distribution of a fixed budget of attack traffic among the parallel links is an NP-hard problem with respect to the scaling parameter N . From the attacker’s perspective, “optimal” means that the attacker can prevent as much traffic as possible from reaching the destination, by flooding network links so that legitimate traffic in excess the links’ capacity is dropped.

As noted above, Nash equilibria have the desirable feature that they lead to no regret by both players, a feature that is generally not shared by Stackelberg equilibria. It turns out that in the network routing games considered here, Stackelberg equilibria only lead to no-regret (i.e., are also Nash equilibria) in the extreme cases where the attacker controls a very large or a very small amount of traffic. We show this to be true for parallel networks in Section IV. For these two extreme cases, we actually provide explicit formulas for the optimal Stackelberg/Nash routing policies. Not surprisingly

D. Grimsman (davidgrimsman@ucsb.edu), J. R. Marden (jrmarden@ece.ucsb.edu), and J. P. Hespanha (hespanha@ece.ucsb.edu) are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA

This research was supported by NSF Grant #ECCS-1638214, U.S. Office of Naval Research (ONR) grant #N00014-17-1-2060 and ONR MURI grant No. N00014-16-1-2710.

in view of the NP-hardness result, no explicit formulas are provided for intermediate levels of attack traffic.

Motivated by the nontrivial dependence of the Stackelberg policy on the total amount of traffic r^a controlled by the attacker, we also study how uncertainty in r^a affects routing performance. Previous work in this area has modeled this type of uncertainty as a distribution over the possible values of r^a , giving rise to routing policies that give an optimal expected value on the cost function [17]. However, in this work, we define a metric for the “value of information” about the power of the attacker that compares the amount of traffic that the attacker could block if the router knew precisely r^a versus the amount of traffic it could block if the router had to select a policy without precise knowledge of r^a . The latter scenario generally leads to an increase in blocked traffic. We show in Section V a closed-form expression for the value of information in two-link networks.

II. MODEL

This paper focuses on a two-player network routing game where the system operator is tasked with deriving a routing policy to maximize the throughput of a given single source / single destination parallel network in the presence of an adversary. The network is comprised of a set of edges E , where each edge $e \in E$ is associated with a given capacity $c_e \geq 0$, and we denote $C(E')$ as the sum of the capacities of all edges in $E' \subseteq E$. The system operator, which we will henceforth refer to as the *router*, is tasked with designing a routing profile $f = \{f_e\}_{e \in E}$ which routes $r \geq 0$ units of traffic across this network. A feasible routing profile satisfies $\sum_{e \in E} f_e = r$ and $0 \leq f_e \leq c_e$ for all edges $e \in E$. We denote the convex set of all admissible routing profiles as $\mathcal{F}(c, r)$ where $c = \{c_e\}_{e \in E}$ denotes the capacities of all edges.

This work considers the existence of an attacker whose goal is to block as much routed traffic as possible by reducing the capacities of the edges in the network through a cross-fire style attack where the attacker can send up to $r_a \geq 0$ units of non-responsive traffic on various edges in the network. An adversarial attack can be characterized by a routing profile $f^a = \{f_e^a\}_{e \in E}$ which satisfies $\sum_{e \in E} f_e^a = r^a$ and $0 \leq f_e^a \leq c_e$ for all edges $e \in E$. We denote the set of all admissible adversarial attack policies as $\mathcal{F}^a(c, r^a)$. We will often refer to r^a as the attack budget of the adversary. Given an admissible routing profile $f \in \mathcal{F}(c, r)$ and an adversarial attack $f^a \in \mathcal{F}^a(c, r^a)$, the amount of legitimate traffic blocked on any edge $e \in E$ is defined as

$$B_e(f, f^a, c) := \max\{f_e + f_e^a - c_e, 0\}, \quad (1)$$

and the total blocked traffic in the system as $B(f, f^a, c) = \sum_{e \in E} B_e(f, f^a, c)$. Since the routing policy is non-responsive, the adversarial choice effectively reduces the capacity on each edge e from c_e to $c_e - f_e^a$. Lastly, we will often omit highlighting the functional dependence on the parameters c , r , and r^a for brevity, e.g., express $\mathcal{F}^a(c, r^a)$ as merely \mathcal{F}^a , when this dependence is clear.

One focus of this paper is to characterize different forms of equilibria in this two-player network routing game. In general, we will assume that a router is required to choose the routing strategy first and the adversary can respond accordingly. The most natural class of equilibria that captures this phenomena is that of Stackelberg equilibria (SE), which consists of any pair of routing profiles (f, f^a) such that

$$f \in \arg \inf_{\bar{f} \in \mathcal{F}} \sup_{\bar{f}^a \in \mathcal{F}^a} B(\bar{f}, \bar{f}^a, c), \quad (2)$$

$$f^a \in \arg \sup_{\bar{f}^a \in \mathcal{F}^a} B(f, \bar{f}^a, c). \quad (3)$$

If f^a satisfies (3), we refer to f^a as a *best response attack* to f . A second class of equilibria that we focus on is Nash equilibria (NE), which focuses on situations where both the router and adversary are required to select their strategy without knowledge of the other’s choice. A NE is defined as any pair of profiles (f, f^a) such that

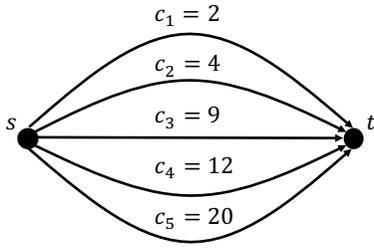
$$f \in \arg \inf_{\bar{f} \in \mathcal{F}} B(\bar{f}, f^a, c), \quad (4)$$

$$f^a \in \arg \sup_{\bar{f}^a \in \mathcal{F}^a} B(f, \bar{f}^a, c). \quad (5)$$

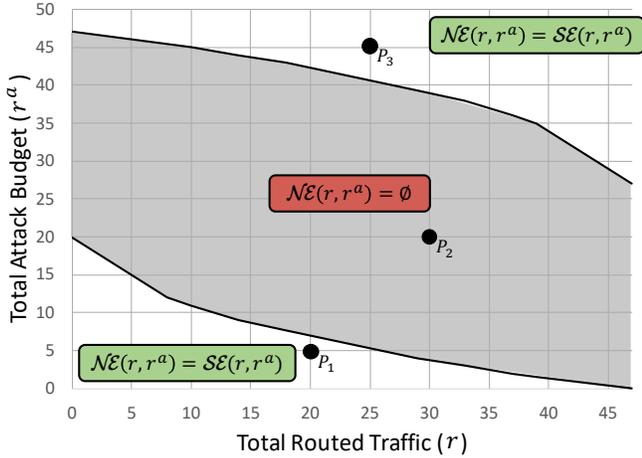
We refer to $\mathcal{SE}(r, r^a)$ as the set of all SE for values r, r^a , and likewise $\mathcal{NE}(r, r^a)$ for NE. Note that given the definitions above, $\mathcal{NE}(r, r^a) \subseteq \mathcal{SE}(r, r^a)$. In the event where $\mathcal{NE}(r, r^a) = \mathcal{SE}(r, r^a)$, this implies that the router is not strategically disadvantaged by having to reveal its choice before the adversary selects its policy. However, while a SE will always exist, the same does not hold true for NE. Furthermore, this paper will address how knowledge of the exact value of r^a impacts the existence and efficacy of such equilibria.

Example 1: We begin with the following example highlighting the complexity of computing NE and SE in such a routing game. To that end, consider the example shown in Figure 1a with $r = 25$ and $r^a = 20$ and denote the edge set as $E = \{1, 2, \dots, 5\}$ and edge capacities as $c = \{2, 4, 9, 12, 20\}$. Given a routing profile $f = \{1, 1, 5, 10, 8\}$ and an attack profile $f^a = \{2, 4, 4, 4, 6\}$, it follows from (1) that the traffic blocked on each edge is 1, 1, 0, 2 and 0, respectively. Note that these strategy profiles (f, f^a) neither capture a NE or SE as there are numerous adversarial strategies that could increase the total blocked traffic given the routing profile f , e.g., $\bar{f}^a = \{0, 0, 8, 12, 0\}$.

The plot in Figure 1b highlights the distinction between NE and SE for the considered routing problem for all pairs (r, r^a) satisfying $47 \geq r, r^a \geq 0$. For instance, when $r = 20$ and $r^a = 5$ (see point P_1 in Figure 1b), any SE is also a NE. One such routing profile is $f = \{0, 0, 0, 5, 15\}$, as this does not allow the attacker to block any traffic. When $r = 25$ and $r^a = 45$ (point P_3), we see a similar phenomenon, where the attacker has much more power. In fact, observe that the routing profile $f = \{2, 4, 6.\bar{3}, 6.\bar{3}, 6.\bar{3}\}$ and attack profile $f^a = \{0, 4, 9, 12, 20\}$ constitute both a SE and NE. The router is able to design a policy such that the attacker can only block $\sum_{e \in E} c_e - r^a$ traffic, the best the router can achieve given r^a . Thus the router has no incentive to deviate,



(a) An example network. Suppose that $f = \{1, 1, 5, 10, 8\}$ and $f^a = \{2, 4, 4, 4, 6\}$. Then, for instance on edge 4, since the capacity is 12, 2 units of traffic are blocked. In total we see that $B_1 = 1$, $B_2 = 1$, $B_3 = 0$, $B_4 = 2$, and $B_5 = 0$, which results in $B(f, f^a, c) = 4$.



(b) This figure showcases one of the contributions of this paper: a characterization of where no NE exist (gray region) and when all SE are also NE (white regions) for the network in (a). For any values (r, r^a) , one of those two properties must hold. See Example 1 and Theorem 2 for more details.

Fig. 1: An example network showcasing the model and regions of (r, r^a) where NE exist.

and clearly the attacker cannot. Lastly, when $r = 30$ and $r^a = 20$ (point P_2) we begin to notice a discrepancy between NE and SE in the sense that given any profiles (f, f^a) , if (5) is satisfied then (4) is not satisfied. For example, consider the profiles $f = \{1.4, 4, 6.4, 6.4, 11.8\}$ and $f^a = \{0, 0, 0, 0, 20\}$ and note that f^a satisfies (5). If the attacker implements this policy, then (f, f^a) is not a NE, since the router would benefit unilaterally by moving some traffic from edge 5 to another unblocked edge. The forthcoming Theorem 2 provides the characterization shown in Figure 1b.

III. PROBLEM HARDNESS

In this section, we show that finding the best response attack policy in (3) is NP-Hard. We formally define it as follows:

Problem 1: Given a parallel network with edges E , corresponding capacities c , a routing policy f , and attack power r^a , find f^a which satisfies (3), i.e., a best response attack policy.

Note that an instance of the problem can be defined by (E, c, f, r^a) , and we show how the complexity of the

problem scales with the number of edges in the parallel network.

Theorem 1: Problem 1 is NP-Hard on the scaling variable $|E|$.

The theorem is proved by reducing the 0-1 Knapsack Problem (KP), a known NP-Hard problem, to Problem 1. We do this by showing that if all f_e are “sufficiently small”, then any best response attack must either block all traffic on an edge or block none of it. Thus finding the best response attack is simply finding the set of edges to fully block, corresponding to the discrete nature of the items in the 0-1 KP. This implies any method for solving these instances of Problem 1 will also solve the 0-1 KP.

The following lemma defines “sufficiently small” in this context:

Lemma 1: Consider an instance of Problem 1 (E, c, f, r^a) , where

$$f_e < \min_{E' \subseteq E: r^a - C(E') > 0} r^a - C(E'), \quad (6)$$

for some $e \in E$. Then $B_e(f, f^a) \in \{0, f_e\}$ for any f^a which is a solution to Problem 1.

Proof: We prove the contrapositive statement. Let e be such that $B_e(f, f^a) \notin \{0, f_e\}$. Define $E^{\text{block}} := \{e' \in E : f_{e'} + f_{e'}^a > c_{e'}\}$, and observe by definition that $e \in E^{\text{block}}$. Then it must be true that $f_e > r^a - C(E^{\text{block}} \setminus \{e\}) > 0$, otherwise the attacker could block more routed traffic by redistributing as much attack traffic as possible from e to the other edges in E^{block} . Therefore, (6) must be false. ■

Given this, we proceed with the proof of Theorem 1. The 0-1 KP can be defined as follows: assume we have n items, where each item e has a cost w_e and a value v_e . Given a total cost constraint W , find the combination of items with maximum total which does not exceed W . More formally stated, determine

$$\begin{aligned} & \underset{x}{\text{maximize}} && \sum_e v_e x_e \\ & \text{subject to} && \sum_e x_e w_e \leq W, \quad x_e \in \{0, 1\}, \end{aligned} \quad (7)$$

where $x := [x_e]$. This problem is known to be NP-Hard in the number of items [18].

Mapping a 0-1 KP to Problem 1 can be done with the following method: let every item be mapped to an edge in a parallel network, $r^a = W$, $c_e = w_e$, and $f_e = \varepsilon v_e$, where $\varepsilon > 0$ satisfies

$$\varepsilon v_e < \min_{E' \subseteq E: r^a - C(E') > 0} r^a - C(E'), \quad (8)$$

for all $e \in E$. By Lemma 1, we know that any solution to this subset of instances of Problem 1 has the property that every edge will either have all routed traffic blocked or none. Therefore, the problem can be reformulated as

$$\begin{aligned} & \underset{x}{\text{maximize}} && \sum_e f_e x_e \\ & \text{subject to} && \sum_e x_e c_e \leq r^a, \quad x_e \in \{0, 1\}. \end{aligned} \quad (9)$$

This problem yields an equivalent solution to that in (7), since the constraints are the same, and each objective function is a scaled version of the other. Thus solving this instance of Problem 1 will also solve 0-1 KP and shows that Problem 1 is NP-Hard. ■

IV. EQUILIBRIA

In this section, we present results that describe precisely the relationship between SE and NE in our model. We then give a two examples: one which illustrates this relationship and one which showcases why generalizing these results to even slightly more complex networks is nontrivial.

Theorem 2: Consider a parallel network with capacities c , routing demand r , and adversarial routing power r^a . The set of Nash Equilibria $\mathcal{NE}(r, r^a)$ is nonempty and $\mathcal{NE}(r, r^a) = \mathcal{SE}(r, r^a)$ if and only if one of the following is satisfied:¹

$$r^a \leq \max_{E' \subseteq E} \frac{C(E') - r}{|E'|} \quad (10)$$

$$r^a \geq C(E) - \max_{E' \subseteq E} \frac{r - C(E \setminus E')}{|E'|}. \quad (11)$$

The proof for Theorem 2 can be found in [19]. An implication of this proof is that the following routing profiles are of some importance:

$$f_e^{\text{lo}} := \max \left\{ c_e - \max_{E' \subseteq E} \frac{C(E') - r}{|E'|}, 0 \right\} \quad (12)$$

$$f_e^{\text{hi}} := \min \left\{ c_e, \max_{E' \subseteq E} \frac{r - C(E \setminus E')}{|E'|} \right\}, \quad (13)$$

namely, that when (10) holds, then (12) is a SE (and NE) routing profile; and that when (11) holds, (13) is a SE (and NE) routing profile.

Refer again to the network in Figure 1. At point P_1 , $r = 20$ and $r^a = 5$. Here we calculate $\max_{E' \subseteq E} (C(E') - r)/|E'| = 7$, which means that r^a satisfies (10). Thus the router can use the policy $f^{\text{lo}} = \{0, 0, 2, 5, 13\}$ to ensure that the attacker cannot block any traffic. By Theorem 2, this also implies that (f, f^a) is both a SE and a NE for any $f^a \in \mathcal{F}^a(c, r)$. At point P_3 , $r = 25$ and $r^a = 45$. Here we calculate $C(E) - \max_{E' \subseteq E} (r - C(E \setminus E'))/|E'| = 40.\bar{6}$, which means that r^a satisfies (11). Thus the router can use the policy $f^{\text{hi}} = \{2, 4, 6.\bar{3}, 6.\bar{3}, 6.\bar{3}\}$, and from Theorem 2, (f, f^a) is a NE and SE for any $f^a \in \mathcal{F}^a(c, r)$. At point P_2 , $r = 30$ and $r^a = 20$. We calculate that $\max_{E' \subseteq E} (C(E') - r)/|E'| = 3.75$ and $C(E) - \max_{E' \subseteq E} (r - C(E \setminus E'))/|E'| = 59$, therefore r^a does not satisfy (10) or (11). By Theorem 2, we know that no NE can exist at this point.

Example 2: Consider now the example in Figure 2, a graph where two parallel networks are connected in series. We present this as a simple example to showcase the complexities that arise when studying the SE of non-parallel

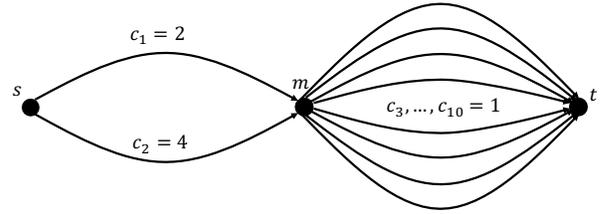


Fig. 2: Two parallel networks in series. We use this example to illustrate the complexities for finding SE and NE in more general networks than just parallel. For instance, one cannot simply decompose the optimal attack problem into either attacking the set of edges between s and m , and attacking the edges between m and t . Even if we limited our scope to such attacks, which set of edges to attack depends on the value of r^a , not merely on f and c . See Example 2 for more details.

networks. For more complex networks, one might think that finding a best response attack could be limited to attacking a minimal cut-set in the network. However, even in this very simple example, we show that this is not the case, and in fact, a best response attack will often incorporate edges of multiple cut-sets in the network. Thus investigating parallel networks in this paper gives a natural simplification of the problem in order to address the questions of interest.

In Figure 2, denote E_{sm} as the cut-set of edges between s and m and E_{mt} as the cut-set of edges between m and t . Observe that regardless of the attacker's capability, there always exists a SE route where all edges in E_{mt} have the same amount of traffic routed on them. We assume in the following cases that the router always uses such a policy, and therefore, we need only focus on the routing strategy across E_{sm} .

Let $r = 2$ and $r^a = 5$. If the attacker restricts its attacks to a single cut-set E_{sm} or E_{mt} , then the router can choose its policy accordingly, for instance $f_e = 1$ for $e \in E_{sm}$, and $f_e = 0.25$ for $e \in E_{mt}$. Note that across each cut-set, this route satisfies (2). Attacking only E_{sm} , the attacker can block 1 unit of traffic, but attacking only E_{mt} , the attacker can block 1.25 units of traffic. This may seem unintuitive, since the total capacity of E_{sm} is less than that of E_{mt} . Furthermore, the best response for the attacker is to block some traffic on E_{sm} and some on E_{mt} . For instance, the attacker could block the 1 unit of traffic on edge 1, and then block all traffic on 3 of the edges in E_{mt} . Assuming that the router evenly distributes the remaining 1 unit of routed traffic that arrives at node m , this attack would block 1.375 units of traffic. Therefore, solving for a SE must include all attacks across multiple cut-sets.

Given these complexities with even very simple non-parallel networks, the characterizations of SE and NE in Theorem 2 only apply to parallel networks. While this class of networks is sufficiently rich to ask the questions and showcase the phenomena that are relevant to this work, future work can ask similar questions in a broader setting.

¹While finding the maxima in (10) and (11) may appear to be computationally intractable given the number of edges in the network, it is true that the maximizing E' for both (10) and (11) is of the form $\{1, 2, \dots, k\}$, where the edges are ordered starting with highest capacity to the lowest. Therefore, finding either maxima is equivalent to finding the best value of k , which can be completed in linear time.

V. THE VALUE OF INFORMATION

In this section, we present preliminary results about the value to the router of knowing information about the attack power r^a . In order to do this, we introduce some notation. We define

$$B^*(f, r^a) := \max_{f^a \in \mathcal{F}^a(r^a)} B(f, f^a), \quad (14)$$

in other words, $B^*(f, r^a)$ measures how much traffic is blocked in the attacker's best response to f , given r^a . We also define

$$B^{\text{SE}}(r, r^a) := B(f, f^a), \quad (15)$$

where $(f, f^a) \in \mathcal{F}(r) \times \mathcal{F}^a(r^a)$ is a SE. Recall that for the pair (r, r^a) the same amount of traffic will be blocked by any SE (f, f^a) .

As an example of both these functions, consider the plot in Figure 3 for a three-link parallel network where $c = \{2, 3, 5\}$ and $r = 5$. For the fixed route $f = \{0.5, 2, 3, 5\}$, the gray line represents how $B^*(f, r^a)$ changes as a function of r^a . Likewise, the orange line showcases $B^{\text{SE}}(r, r^a)$ as a function of r^a . Observe that $B^*(f, r^a) \geq B^{\text{SE}}(r, r^a)$ for all values of r^a .

A. Limited information

We limit the router's knowledge of r^a by stating that the router only knows that r^a is in some interval $\pi^a = [\underline{\pi}^a, \bar{\pi}^a]$. In light of this uncertainty, if the router chooses policy f , then we can define the risk of f on interval π^a as

$$R(f, \pi^a) := \max_{r^a \in \pi^a} (B^*(f, r^a) - B^{\text{SE}}(r, r^a)). \quad (16)$$

Intuitively, the value $B^*(f, r^a) - B^{\text{SE}}(r, r^a)$ represents how much more traffic the attacker is able to block because the router chose policy f instead of a SE policy for that value of r^a . Thus the risk $R(f, \pi^a)$ is the maximum such value across all $r^a \in \pi^a$. In other words, this measurement of risk shows, in the worst case, the advantage that the attacker gains by the router not knowing the true value of r^a .

As an example, consider again the plot in Figure 3. If we assume that the router has no knowledge of r^a (i.e., $\pi^a = [0, 10]$), then the risk associated with the route $f = \{0.5, 2, 3, 5\}$ is the maximum difference between the gray and orange lines, which is achieved at $r^a = 8$. Therefore, in this case we see that $R(f, \pi^a) = 1.5$.

It turns out that the maximization in (16) can be restricted to a finite set of points in π^a .

Lemma 2: For a parallel network,

$$R(f, \pi^a) = \max_{r^a \in (\alpha \cap \pi^a) \cup \{\underline{\pi}^a, \bar{\pi}^a\}} B^*(f, r^a) - B^{\text{SE}}(r, r^a), \quad (17)$$

where α is the finite set $\{r^a : \exists E' \subseteq E \text{ where } r^a = C(E')\}$, which has at most $2^{|E|}$ elements.

The full proof is given in [19], however here we provide some intuition: consider the plot in Figure 3. The orange line, $B^{\text{SE}}(r, r^a)$, is piecewise linear, with no line slope being greater than 1. The grey line, $B^*(f, r^a)$, is also a piecewise linear function, with lines slopes either 0 or 1. The value

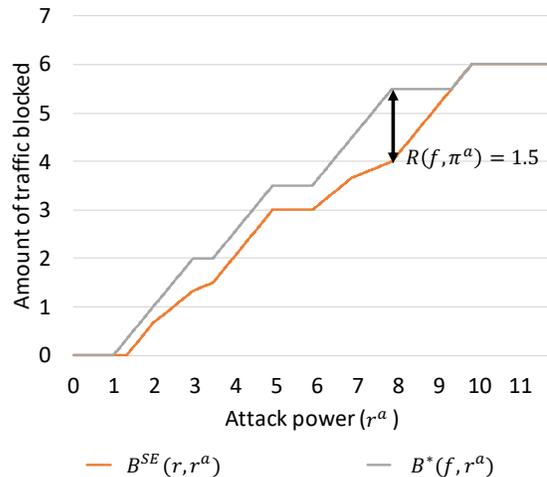


Fig. 3: A plot showing the amount of traffic blocked by an optimal attack for a SE routing policy (orange) versus the traffic blocked when the router selects specific routing policy f (regardless of the value of r^a). Here $c = \{2, 3, 5\}$, and $r = 5$. The fixed policy represented by the gray line is $f = \{0.5, 2, 3, 5\}$. We show the values of the risk $R(f, \pi^a)$ for $\pi^a = [0, 10]$.

of the risk $R(f, \pi^a)$ is incurred at $r^a = C(\{2, 3\}) = 8$, where the attacker's best response against f is to fully block edges 2 and 3. Because the two lines are piecewise linear, the largest distance must take place at one of the points of discontinuity for the gray line inside the interval π^a .

Finally, we define the *value of information* to the router for an interval π^a as the minimum amount of risk that can be incurred for any routing policy. More formally stated,

$$V(\pi^a) := \min_{f \in \mathcal{F}} R(f, \pi^a) \quad (18)$$

$$= \min_{f \in \mathcal{F}} \max_{r^a \in \pi^a} (B^*(f, r^a) - B^{\text{SE}}(r, r^a)) \quad (19)$$

We also denote the routing policy which minimizes (18) by f^π . This value of information is meant to reflect how valuable (i.e., how much less traffic would be blocked) if the router knew the exact value of r^a . For instance, if $V(\pi^a) = 0$, then there exists a route which satisfies (2) for any value of $r^a \in \pi^a$, thus the router does not need to know the exact value. However, when $V(\pi^a)$ is high, knowing r^a would allow the router to ensure that less traffic is blocked. Figure 4 shows $V(\pi^a)$ and f^π for a two-link network.

B. The Value of Information in Two-Link Networks

Lemma 2 provides a numerical procedure to compute the risk for a routing policy f against an attack power interval π^a for general parallel networks. For two-link networks, this means there exists a closed-form solution for $R(f, \pi^a)$ and subsequently $V(\pi^a)$.

Theorem 3: Consider a two-link parallel network, where $c_1 \leq c_2$. Suppose that the router only knows that $r^a \in \pi^a =$

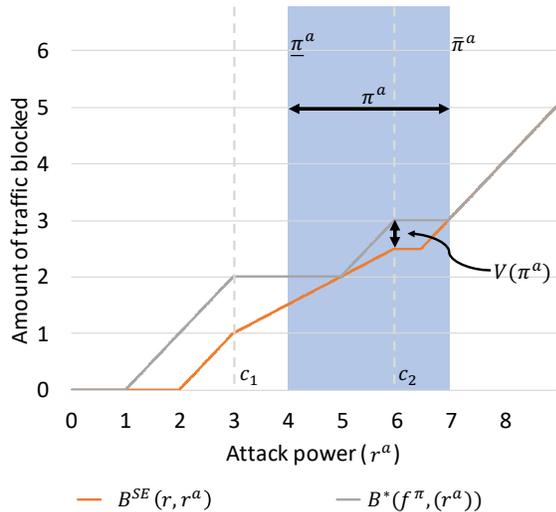


Fig. 4: A plot with an example two-link parallel network that shows a graphical interpretation for $V(\pi^a)$. The edge capacities in the network are $\{3, 6\}$, and $r = 5$. The orange line represents how much traffic is blocked at the SE for each value of r^a , and the gray line is how much is blocked by a best response attack against $f^\pi = \{2, 3\}$ for each value of r^a . The interval $\pi^a = [4, 7]$ is the blue shaded region. The value of information $V(\pi^a)$ is the maximum difference between the two lines within the blue region.

$[\underline{\pi}^a, \bar{\pi}^a]$. Then the value of information is

$$V(\pi^a) = \begin{cases} 0, & \text{if } \pi^a \cap [c_1, c_2] = \emptyset \\ \frac{1}{4}(\min\{\bar{\pi}^a, c_2\} - \max\{\underline{\pi}^a, c_1\}), & \text{otherwise.} \end{cases} \quad (20)$$

The theorem proof is given in [19], but here we give an example to provide some intuition. Consider the plot in Figure 4. In this network, $c = \{3, 6\}$, and $r = 5$. If the router knows the exact value of r^a , it can choose a SE routing policy, which will make the difference between the lines 0 at that value of r^a . If we assume that the router only knows that $r^a \in \pi^a = [4, 7]$, then it must choose a policy to mitigate the risk associated with that loss of information. In this scenario, the router's best option is to use $f^\pi = \{2, 3\}$ (gray line), which minimizes the maximum difference between the two lines on π^a . The value of the router knowing r^a is then this minimum maximum difference, i.e., $V(\pi^a) = 0.5$.

VI. CONCLUSION

In this paper, we studied a particular set of network routing games, wherein the attacker has full knowledge of the router policy before choosing its own policy. We showed that choosing such a best response attack policy is an NP-Hard problem over the class of parallel networks. We showed that in such networks, a SE policy is also a NE policy when the attack either does not have enough attack power to affect anything, or where the attacker can block nearly everything. We concluded with a study on two-link networks and how the router's uncertainty of the attack power can affect how much traffic is blocked. We also gave a method for designing

routing policies to be as robust as possible against such uncertainty.

Future work will focus on expanding this value of information study first to parallel networks, and then to the set of all networks. Another path is to understand, when the routing policy is not centralized, but distributed, how each router can be incentivized to use local information to determine the proper routing policy.

REFERENCES

- [1] M. S. Kang, S. B. Lee, and V. D. Gligor, "The crossfire attack," in *2013 IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 127–141.
- [2] D. Gkounis, V. Kotronis, and X. Dimitropoulos, "Towards defeating the crossfire attack using SDN," *arXiv preprint arXiv:1412.2013*, 2014.
- [3] A. Aydeger, N. Saputro, K. Akkaya, and M. Rahman, "Mitigating crossfire attacks using SDN-based moving target defense," in *2016 IEEE 41st Conference on Local Computer Networks (LCN)*. IEEE, 2016, pp. 627–630.
- [4] D. Gkounis, V. Kotronis, C. Liaskos, and X. Dimitropoulos, "On the interplay of link-flooding attacks and traffic engineering," *ACM SIGCOMM Computer Communication Review*, vol. 46, no. 2, pp. 5–11, 2016.
- [5] A. Raj, T. Truong-Huu, P. M. Mohan, and M. Gurusamy, "Crossfire attack detection using deep learning in software defined ITS networks," *arXiv preprint arXiv:1812.03639*, 2018.
- [6] J. Nash, "Non-cooperative games," *Annals of mathematics*, pp. 286–295, 1951.
- [7] H. Von Stackelberg, *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- [8] Y. A. Korilis, A. A. Lazar, and A. Orda, "Achieving network optima using Stackelberg routing strategies," *IEEE/ACM transactions on networking*, vol. 5, no. 1, pp. 161–173, 1997.
- [9] T. Roughgarden, "Stackelberg scheduling strategies," *SIAM journal on computing*, vol. 33, no. 2, pp. 332–350, 2004.
- [10] M. Bloem, T. Alpcan, and T. Başar, "A Stackelberg game for power control and channel allocation in cognitive radio networks," in *Proceedings of VALUETOOLS*. ICST (Institute for Computer Sciences, Social-Informatics and ...), 2007, p. 4.
- [11] X. He, A. Prasad, S. P. Sethi, and G. J. Gutierrez, "A survey of Stackelberg differential game models in supply and marketing channels," *Journal of Systems Science and Systems Engineering*, vol. 16, no. 4, pp. 385–413, 2007.
- [12] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, "Defending critical infrastructure," *Interfaces*, vol. 36, no. 6, pp. 530–544, 2006.
- [13] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, "Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport," in *Proceedings of the 7th IFAAMAS: industrial track*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 125–132.
- [14] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez, "Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service," *Interfaces*, vol. 40, no. 4, pp. 267–290, 2010.
- [15] G. G. Brown, W. M. Carlyle, J. Salmeron, and K. Wood, "Analyzing the vulnerability of critical infrastructure to attack and planning defenses," in *Emerging Theory, Methods, and Applications*. INFORMS, 2005, pp. 102–123.
- [16] D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness," *Journal of Artificial Intelligence Research*, vol. 41, pp. 297–327, 2011.
- [17] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus, "Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games," in *Proceedings of the 7th IFAAMAS-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 895–902.
- [18] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack problems*. Springer, Berlin, 2004.
- [19] D. Grimsman, J. P. Hespanha, and J. R. Marden, "Stackelberg equilibria for two-player network routing games on parallel networks." [Online]. Available: <http://arxiv.org/abs/2003.05882>