Adaptivity in Network Interdiction

Bastián Bahamondes[1], José Correa[1], Jannik Matuschke[2(✉)], and Gianpaolo Oriolo[3]

[1] Department of Industrial Engineering, Universidad de Chile, Santiago, Chile
{bastian.bahamondes,correa}@uchile.cl
[2] TUM School of Management and Department of Mathematics, Technische Universität München, Munich, Germany
jannik.matuschke@tum.de
[3] Department of Civil Engineering and Computer Science, Università di Roma Tor Vergata, Rome, Italy
oriolo@disp.uniroma2.it

Abstract. We study a network security game arising in the interdiction of fare evasion or smuggling. A defender places a security checkpoint in the network according to a chosen probability distribution over the links of the network. An intruder, knowing this distribution, wants to travel from her initial location to a target node. For every traversed link she incurs a cost equal to the transit time of that link. Furthermore, if she encounters the checkpoint, she has to pay a fine.

The intruder may adapt her path online, exploiting additional knowledge gained along the way. We investigate the complexity of computing optimal strategies for intruder and defender. We give a concise encoding of the intruders optimal strategy and present an approximation scheme to compute it. For the defender, we consider two different objectives: (i) maximizing the intruder's cost, for which we give an approximation scheme, and (ii) maximizing the collected fine, which we show to be strongly NP-hard. We also give a paramterized bound on the worst-case ratio of the intruders best adaptive strategy to the best non-adaptive strategy, i.e., when she fixes the complete route at the start.

1 Introduction

Network interdiction problems model the control or halting of an adversary's activity on a network. Typically, this is modelled as the interaction between two adversaries—an *intruder* and a *defender*—in the context of a Stackelberg game. The defender allocates (or removes) scarce resources on the network in order to thwart the objective of the intruder, who—knowing the defender's strategy—reacts by choosing the response strategy optimizing his own objective. Such models are used to great effect in applications such as disease containment [11,13], drug traffic interdiction [17], airport security [16], or fare inspection [5].

In order to mitigate the intruder's advantage of observing the defender's actions first, the defender may opt to employ a randomized strategy. The intruder can only observe the probability distribution of the defender's actions, but she

© Springer International Publishing AG 2017
S. Rass et al. (Eds.): GameSec 2017, LNCS 10575, pp. 40–52, 2017.
DOI: 10.1007/978-3-319-68711-7_3

does not know the exact realization. In this work, we study a variant of a network interdiction problem in which the defender employs such randomization, but the intruder gains additional information about the realization while she is acting, and may use this information to adapt her strategy.

Our game is played on a network. The defender randomly establishes a security checkpoint on one of the arcs. The intruder wants to move from her initial location to a designated target node, preferably without being detected by the defender. Her objective is to minimize her expected cost, which consists of movement costs for traversing arcs and a fine, which has to be paid if she traverses the arc with the checkpoint. Knowing the probability distribution specified by the defender and that only one arc is subjected to inspection, the intruder gains additional information while traveling through the network, observing whether or not the inspected arc was among those she traversed so far. She may use this information in order to decide which arc to take next. This type of path-finding strategy is called *adaptive*, as opposed to a *non-adaptive strategy*, in which she commits to an origin-destination-path at the start and does not deviate from it.

In this paper, we investigate the computational complexity of finding optimal adaptive and non-adaptive strategies for the intruder as well as optimal randomized strategies for the defender, considering two objectives: (i) the *zero-sum* objective of maximizing the intruder's cost and (ii) the *profit maximization* objective of maximizing the expected collected fine. We also provide bounds on the cost ratio between optimal adaptive and non-adaptive strategies and the impact of adaptivity on the defender's objective.

1.1 Related Work

Stackelberg games, and in particular network interdiction models, are widely used in the context of security applications; see the textbook by Tambe for an overview of applications in airport security [16].

A very basic version of a Stackelberg game is the security game studied by Washburn and Wood [17]. In this zero-sum game, an inspector strives to maximize the probability of catching an evader, who chooses a path minimizing that probability. The authors show that optimal strategies for both players can be computed by a network flow approach. The optimization problem of maximizing the defender's profit has been extensively studied in the context of *Stackelberg pricing games* [3,9,14]. Here, the defender sets tolls for a subset of the edges of the network, trying to collect as much tolls as possible from the intruder, who chooses a path minimizing the sum of the travel costs plus the tolls. As opposed to the zero-sum game mentioned above, these pricing games are usually computationally hard to solve.

The particular game we study in this article arises from a variation of two toll/fare inspection models introduced by Borndörfer et al. [2] and Correa et al. [5], respectively. In these models, the defender, who represents the network operator, decides an inspection probability for each arc, subject to budget limiting the total sum of inspection probabilities. The intruder (toll evading

truck drivers/fare evading passengers) tries to get to her destination minimizing a combined objective of travel time and expected cost for the fine when being discovered. Correa et al. [5] also study an adaptive version of the problem, in which the intruders adapt their behavior as they traverse the network. They propose an efficient algorithm based on a generalized flow decomposition, and give a tight bound on the adaptivity gap of 4/3; see Sect. 5 for details. In both the above models, the event of an inspection occurring on a given arc is independent to that on all other arcs. In contrast, in our model, the checkpoint can only be located on a single arc, leading to a different optimization problem for the intruder.

A different notion of adaptive path-finding was previously studied by Adjiashvili et al. [1] in the so-called *Online Replacement Path* problem. Here, a routing mechanism must send a package between two nodes in a network trying to minimize transit cost. An adversary, knowing the intended route, may make one of the arcs fail. Upon encountering the failed arc, the package may be rerouted to its destination along a different path. Note that in this setting the failing arc is chosen by the adversary *after* the routing has started, whereas in our settings the inspection probabilities are determined *before* the intruder chooses her path. Computationally, adaptive path-finding is related to shortest path problems in which there is a trade-off between two cost functions. The restricted shortest path problem [6,8,10] and the parametric shortest path problem [4,12] are representative examples of such problems.

1.2 Contribution

We study both adaptive and the non-adaptive path-finding strategies for the intruder. After observing that the non-adaptive intruder's problem reduces to the standard shortest path problem, we turn into the adaptive version, which turns out to be much more intricate. We show that an optimal adaptive strategy of the intruder can always be represented by a simple, i.e., cycle-free, path. We then devise fully polynomial time approximation scheme (FPTAS) for computing the a near-optimal adaptive strategy with adjustable precision.

By using an approximate version of the equivalence of separation and optimization [15], we also obtain an FPTAS for maximizing the defender's zero-sum objective. For the profit objective, on the other hand, we show that the defender's optimization problem is strongly NP-hard, ruling out the existence of an FPTAS (unless $P = NP$).

We further study the impact of adaptivity on the intruder's and defender's objective. Extending a result by Correa et al. [5], we show that the intruder's best non-adaptive strategy is within a factor of 4/3 of the optimal adaptive strategy and that this ratio decreases for instances where the intruder does not deviate significantly from her shortest path (which is a natural assumption, e.g., in the context of transit networks). We also mention that our bound on the adaptivity gap for the intruder directly translates to several guarantees for the defender's zero sum game, e.g., bounding his loss in pay-off when he wrongly believes the intruder is non-adaptive.

2 The Model

Before we can describe our model in detail, we establish some notation. Throughout this article, we are given a directed graph $G = (V, E)$ with $n := |V|$ nodes and $m := |E|$ arcs. For two nodes $u, v \in V$ an u-v-walk in G is a sequence of edges $(e_1, \dots, e_k)$ with $e_i = (v_{i-1}, v_i) \in E$ and $v_0 = u$ and $v_k = v$. A u-v-path is a u-v-walk in which no arc or node is repeated, i.e., $e_i \neq e_j$ and $v_i \neq v_j$ for $i \neq j$. For a u-v-path P, we let $V(P)$ be the set of nodes visited by P and for $u', v' \in V(P)$ such that P visits u' before v', we let $P[u', v']$ denote the u'-v'-path contained in P. We denote the set of all u-v-walks in G by $\mathcal{W}_{uv}$ and the set of all u-v-paths in G by $\mathcal{P}_{uv}$.

In our model, the intruder starts at a designated node s and wants to reach a node t (both nodes are also known to the defender). Each arc $e \in E$ is equipped with a cost $c_e \in \mathbb{Z}_+$ that is incurred to the intruder when she traverses e. Furthermore, there is a fine F, which the defender charges to the intruder, if she runs into the defender's security checkpoint. In the first level of our interdiction game, the defender specifies the random distribution of the checkpoint, i.e., he specifies for every arc $e \in E$ the probability π_e of placing the checkpoint at e. In the second level, the intruder takes her way from s to t, having full knowledge of the probability distribution chosen by the defender. We distinguish two variants of the intruder's path-finding strategy:

non-adaptive: At the start, the intruder selects an s-t-path $P \in \mathcal{P}_{st}$ and follows this path to t. For every arc $e \in P$ she pays the transit cost c_e of that arc. In addition, if the security checkpoint is located on one of the arcs of P, she has to pay the fine F.

adaptive: From her current location, the intruder moves along one of the outgoing arcs e to a neighboring vertex, paying the transit cost c_e. She observes whether the security checkpoint is located at the arc she traverses (in which case she additionally has to pay the fine F). Knowing this information, she decides which arc to take next. This procedure continues until she reaches her destination (after a finite number of steps).

The intruder's objective is to minimize her expected cost. For a set of arcs S, we use $c(S) := \sum_{e \in S} c_e$ to denote the sum of the transit times and $\pi(S) := \sum_{e \in S} \pi_e$ to denote the probability that the security checkpoint is located within the set of arcs S (note that we can sum up these probabilities since there is a single checkpoint, so these are disjoint events). Therefore, in the non-adaptive case, the expected cost of following a path P is

$$f_{\mathrm{N},\pi}(P) := c(P) + \pi(P)F = \sum_{e \in P} (c_e + \pi_e F)\,.$$

We denote the optimization problem of finding an optimal non-adaptive strategy for the intruder by

$$\min_{P \in \mathcal{P}_{st}} f_{\mathrm{N},\pi}(P). \tag{$\textsc{Int}_{\mathrm{N}}$}$$

Thus, it is straightforward to note that an optimal non-adaptive strategy for the intruder is to follow a shortest path with respect to arc weights $c_e + \pi_e F$. Such a path can be computed efficiently, e.g., using Dijkstra's Algorithm. Therefore we conclude the following result.

Proposition 1. $\textsc{Int}_{\text{N}}$ *reduces to the Shortest Path Problem and can be solved in polynomial time.*

The optimal adaptive strategy is less obvious. In principle, the intruder's choice of where to go next from her current location can depend on the set of arcs she has visited so far and the information whether the security checkpoint is located at one of these arcs. Let us consider any such adaptive strategy. Note that, because the intruder has to reach t after a finite number of steps, for each fixed realization of the checkpoint location, the strategy determines an s-t-walk. We distinguish two cases.

First, assume the intruder encounters the checkpoint in every realization. Then for the given strategy, she pays the fine with probability 1. Obviously, the non-adaptive strategy of simply following the shortest path with respect to c has at most the same cost than the considered strategy.

Now assume that there is a realization in which the intruder reaches t without being inspected. Let $W = (e_1, \dots, e_k)$ be the walk she takes in this realization, with $e_i = (v_{i-1}, v_i)$, $v_0 = s$, and $v_k = t$. Observe that W is the same for all realizations where the intruder is not inspected, as her decisions are based only on whether or not she encountered the checkpoint so far. We now define a new adaptive strategy, in which the intruder follows W starting at s until she either reaches t or encounters the security checkpoint at some arc e_i of W. In the latter case, after traversing e_i she simply follows a shortest path with respect to c from her current location v_i to t. It is easy to check that the cost of the new strategy is at most the cost of the strategy considered originally.

We have thus shown that for every adaptive strategy there is a strategy of at most the same cost which is completely defined by an s-t-walk W that the intruder follows while not being inspected. Note that W can contain cycles and arcs can appear multiple times along W. Define $\tilde{\pi}_i := \pi_{e_i}$ if $e_i \neq e_j$ for all $j < i$, i.e., the ith position is the first appearance of the arc e_i on W, and $\tilde{\pi}_i := 0$ otherwise, i.e., if arc e_i occurred on W before the ith position. Furthermore, let $\text{SP}_c(v, w) := \min_{P \in \mathcal{P}_{vw}} c(P)$ be the length of a shortest path w.r.t. c from v to w. Then the intruder's expected cost for following W can be expressed as follows:

$$f_{\text{A},\pi}(W) := \sum_{i=1}^{k} \tilde{\pi}_i \left(\sum_{j=1}^{i} c_{e_j} + F + \text{SP}_c(v_i, t) \right) + \left(1 - \sum_{i=1}^{k} \tilde{\pi}_i \right) \sum_{i=1}^{k} c_{e_i}$$

Here, each summand of the first sum corresponds to the event that the checkpoint is encountered at arc e_i (which can only happen if it is the first occurrence of this arc along the walk). In this case, the intruder traverses the walk W until e_i, pays the fine, and then follows the shortest path from v_i to t. The second sum

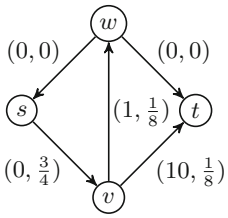

Fig. 1. Example network for the intruder's best response problem. Labels (c_e, π_e) at the arcs denote transit times and inspection probabilities. A possible adaptive strategy for the intruder is to follow s-t-walk s-v-w-s-v-t and deviating to a shortest path when encountering the security checkpoint. For a fine $F = 7$, the expected cost of this strategy is 9.25, whereas following the underlying simple path s-v-t deviating to a shortest path after inspection has a higher expected cost of 9.375.

represents the event that none of the arcs in W contains the checkpoint, in which case the intruder simply traverses W from start to end.

In the above discussion, we assumed that the intruder may walk along cycles and even traverse arcs multiple times. Although all transit costs are non-negative, such detours cycles could—in principle—help the intruder, because along the way she gains additional information. In fact, Fig. 1 depicts an example of an s-t-walk containing a cycle where the intruder's expected cost increases when omitting the cycle. However, one can show that there always exists an optimal adaptive solution without a cycle, i.e., defined by an s-t-path.

Lemma 1. *Let P be a shortest s-t-path w.r.t. c. Then $f_{\mathrm{A},\pi}(P) \leq \mathrm{SP}_c(s,t) + F$.*

Lemma 2. *There is an s-t-path P such that $f_{\mathrm{A},\pi}(P) \leq f_{\mathrm{A},\pi}(W)$ for all s-t-walks W.*

The problem of finding an optimal adaptive strategy thus reduces to finding an s-t-path minimizing $f_{\mathrm{A},\pi}$. We denote this optimization problem by

$$\min_{P \in \mathcal{P}_{st}} f_{\mathrm{A},\pi}(P). \tag{$\textsc{Int}_{\mathrm{A}}$}$$

3 Approximating the Intruder's Optimal Strategy

A *fully polynomial time approximation scheme* (FPTAS) for a minimization problem is an algorithm that takes as input an instance of the problem as well as a precision parameter $\varepsilon > 0$, and computes in polynomial time in the size of the input and $1/\varepsilon$ a solution to that instance with cost at most $(1+\varepsilon)\mathrm{OPT}$, where OPT denotes the cost of the optimal solution.

In this section, we design such an FPTAS for $\textsc{Int}_{\mathrm{A}}$. The algorithm is based on a label propagating approach, where each label at node v represents an s-v-path, that is extended by propagating the label along the outgoing edges of v. In order

to keep the number of distinct labels small and achieve polynomial running time, we discard an s-v-path when we find another s-v-path with similar objective function value but higher inspection probability (intuitively higher inspection probability at equal objective value means that any completion of the new path to an s-t-path will be cheaper than the corresponding completion of the former path). An additional challenge, that arises when propagating the labels in the graph, is to ensure that the constructed paths are cycle free. To deal with this issue we argue that there is a way to avoid cycles without overlooking potentially good paths.

The Algorithm. Given $\varepsilon > 0$, let $\alpha := 1 + \frac{\varepsilon}{2n}$. From Lemma 1, we know that the cost of an optimal strategy is in the interval $[0, \mathrm{SP}_c(s,t) + F]$. We divide this interval geometrically by powers of α. Let $K := \lceil \log_\alpha (\mathrm{SP}_c(s,t) + F) \rceil$ and define $I_0 := [0,1)$ as well as $I_k := \left[\alpha^{k-1}, \alpha^k\right)$ for $k \in \{1, \ldots, K\}$. At every node v we maintain an array $L_v^0, \ldots, L_v^K$, where L_v^i is either empty or contains a label (f,q,P) such that P is an s-v-path with $f = f_{\mathrm{A},\pi}(P) \in I_k$ and $q = \pi(P)$.

Initially, only the label $L_s^0 = (0,0,\emptyset)$ is present. In each iteration, the algorithm propagates all labels at each vertex v along all outgoing arcs (v,w). When propagating label (f,q,P) at node v along arc $e = (v,w)$, we get a label (f',q',P') at node w with $f' = f + (1-q)c_e + \pi_e(F + \mathrm{SP}_c(w,t))$, $q' = q + \pi_e$, and $P' = P \cup \{e\}$. In order to avoid cycles, the propagation of (f,q,P) along $e = (v,w)$ only takes place if $w \notin V(P)$. Moreover, if the propagation of a label along an arc gives rise to two different labels (f',q',P') and (f'',q'',P'') for a node such that $f', f'' \in I_k$ for some k, we discard the label with the lower inspection probability (breaking ties arbitrarily). The full description is given in Algorithm 1.

From the previous discussion, the following lemma is straightforward:

Lemma 3. *If Algorithm 1 creates a label (f,q,P) in a node $v \in V$, then P is a (s,v)-path with $f_{A,\pi}(P) = f$ and $\pi(P) = q$.*

Now let P^* be an s-t-path minimizing $f_{\mathrm{A},\pi}(P^*)$. Let $(e_1, \ldots, e_k)$ be the arcs of P^*, with $e_i = (v_{i-1}, v_i)$, $v_0 = s$ and $v_k = t$. Define $f_i^* := f_{\mathrm{A},\pi}(P^*[s,v_i])$ and $q_i^* := \pi(P^*[s,v_i])$. For $x \in \mathbb{R}$, let $(x)_+$ denote the positive part of x, i.e. $(x)_+ := \max\{x, 0\}$. We call an iteration of the outer for loop of Algorithm 1 a *round*. The following lemma can be proved by induction on the rounds of the algorithm, using a sequence of careful estimates on the cost of paths and subpaths.

Lemma 4. *After round i of Algorithm 1, there is a label (f_i, q_i, P_i) at node v_i with $f_i \leq \alpha^i f_i^* - (q_i^* - q_i)_+ \cdot c(P^*[v_i,t])$.*

Lemma 4 in particular implies that, at the end of round n, the algorithm has found an s-t-path P with $f_{\mathrm{A},\pi}(P) \leq \alpha^n f_{\mathrm{A},\pi}(P^*)$. Note that $\alpha^n = (1 + \frac{\varepsilon}{2n})^n \leq (1+\varepsilon)$ for all $\varepsilon < 1$. It is also easy to verify that the algorithm runs in time polynomial in $1/\varepsilon$ and the input size.

Theorem 1. *Algorithm 1 is an FPTAS for* $\textsc{Int}_\mathrm{A}$.

Algorithm 1. FPTAS for $\textsc{Int}_{\text{A}}$

Compute $\text{SP}_c(v,t)$ for all $v \in V$.
Let $\alpha \leftarrow 1 + \frac{\varepsilon}{2n}$ and $K \leftarrow \lceil \log_\alpha (\text{SP}_c(s,t) + F) \rceil$
Let $L_s^0 \leftarrow (0,0,\emptyset)$ and $L_v^k \leftarrow \emptyset$ for all $(v,k) \in V \times \{0,\dots,K\} \setminus \{(s,0)\}$
for $i = 1,\dots,(n-1)$ **do**
  for all $e = (v,w) \in E$ **and** $k = 0,\dots,K$ **do**
    if $L_v^k \neq \emptyset$ **then**
      $\textsc{push}(L_v^k, e)$
Let $(f^*, q^*, P^*) \in \operatorname{argmin}\left\{f \,:\, (f,q,P) \in L_t^k \text{ for some } k\right\}$
Return P^*

procedure $\textsc{push}(L = (f,q,P), e = (v,w))$
  if $w \notin V(P)$ **then**
    Let $f' \leftarrow f + (1-q)c_e + \pi_e\,(\text{SP}_c(w,t) + F)$
    Let $q' \leftarrow q + \pi_e$
    Let $P' \leftarrow P \cup \{e\}$
    Let $k \leftarrow \min\left\{\ell \in \mathbb{Z}_+ \,:\, f' < \alpha^\ell\right\}$
    if $L_w^k = \emptyset$ **then**
      $L_w^k \leftarrow (f', q', P')$
    else
      Let $(f'', q'', P'') \leftarrow L_w^k$
      if $q' > q''$ **then**
        $L_w^k \leftarrow (f', q', P')$

4 Complexity of the Defender's Problem

We study the defender's optimization problem for deciding the inspection probabilities on every edge of the network, for both the adaptive and non-adaptive intruder. We analyze two different objectives: maximizing the minimum expected intruder's cost and collecting the highest possible fine from inspections.

4.1 The Zero-Sum Objective

We first consider the defender's problem of maximizing the intruder's expected cost. This problem can be stated as

$$\max_{\substack{\sum_{e\in E}\pi_e = 1 \\ \pi \geq 0}} \min_{P \in \mathcal{P}_{st}} f_{X,\pi}(P), \tag{$\textsc{Def}_X^{\text{cost}}$}$$

where $X \in \{\text{A}, \text{N}\}$, depending on whether the intruder is adaptive or non-adaptive. Note that for a fixed path $P \in \mathcal{P}_{st}$, the function $f_{X,\pi}(P)$ is affine

linear in π, both for $X = \mathrm{A}$ and $X = \mathrm{N}$. Therefore, we can reformulate the defender's problem as a linear program:

$$\begin{aligned} \max_{\lambda\in\mathbb{R},\pi\in\mathbb{R}^E} \quad & \lambda & & & \\ \text{s.t.} \quad & \lambda & \leq & f_{X,\pi}(P) & \forall\, P \in \mathcal{P}_{st} \\ & \sum_{e\in E} \pi_e & = & 1 & \\ & \pi_e & \geq & 0 & \forall\, e \in E. \end{aligned} \tag{$\mathrm{LP}_X^{\mathrm{cost}}$}$$

Note that the number of constraints in the above LP can be exponential in the size of the network, as it contains one constraint for every path. A standard way to solve such non-compact LPs is to devise a *separation routine*: A famous result by Grötschel, Lovasz, and Schrijver [7] shows that in order to solve a linear program with the ellipsoid method, it is sufficient to determine for a given setting of the variables, whether it is a feasible solution, and if not, find a violated inequality.

Indeed checking whether a given solution (π, λ) is feasible for $\mathrm{LP}_X^{\mathrm{cost}}$ boils down to determining whether there is a path P with $f_{X,\pi}(P) < \lambda$. For this, it is sufficient to determine the intruder's optimal path. As discussed in Sect. 2, this can be done efficiently for the non-adaptive setting. We thus obtain the following theorem.

Theorem 2. $\textsc{Def}_{\mathrm{N}}^{\mathrm{cost}}$ *can be solved in polynomial time.*

For the adaptive intruder problem, we do not know an exact polynomial time algorithm. However, we can use the FPTAS presented in Sect. 3 as an *approximate separation routine*. This enables us to employ an approximation version of the equivalence of separation and optimization [15], obtaining an FPTAS for $\textsc{Def}_{\mathrm{A}}^{\mathrm{cost}}$.

Theorem 3. *There is an FPTAS for* $\textsc{Def}_{\mathrm{A}}^{\mathrm{cost}}$.

4.2 The Profit Maximization Objective

Next we address the problem of maximizing the expected fine collected by the defender through inspections, that is

$$\begin{aligned} \max \ & \sum_{e\in P} \pi_e F & & \\ \text{s.t.} \ & \sum_{e\in E} \pi_e = 1, \quad \pi \geq 0 & & \\ & P \in \operatorname{argmin}\{f_{X,\pi}(P') \,:\, P' \in P \in \mathcal{P}_{st}\}, & & \end{aligned} \tag{$\textsc{Def}_X^{\mathrm{fine}}$}$$

where again $X \in \{\mathrm{A}, \mathrm{N}\}$ specifies whether the intruder employs an adaptive or non-adaptive path-finding strategy, respectively.

This problem shares many features with the Stackelberg network pricing problem, which is defined as follows: in the first stage, the defender sets tolls on a given subset of "tollable" edges. In the second stage the intruder chooses a path between two fixed nodes minimizing the sum of travel times plus the tolls of the traversed arcs. The defender's objective is to maximize the collected revenue from the tolls. Roch et al. [14] showed that this problem is NP-hard.

We show that also $\textsc{Def}_{\mathrm{N}}^{\mathrm{fine}}$ is NP-hard, even when all arc costs are in $\{0, 1, 2\}$. Such a hardness for instances with small input numbers is referred to as *strong* NP-hardness. Our reduction resembles that of Roch et al., but we have to introduce some modifications to accommodate for non-tollable arcs, which exist in the Stackelberg network pricing problem but not in $\textsc{Def}_{\mathrm{N}}^{\mathrm{fine}}$.

Theorem 4. $\textsc{Def}_{\mathrm{N}}^{\mathrm{fine}}$ *is strongly NP-hard.*

Although we do not provide a hardness result for $\textsc{Def}_{\mathrm{A}}^{\mathrm{fine}}$, we expect it to be NP-hard as well, as the adaptive intruder's first stage problem becomes as least as hard than it is in the $\textsc{Def}_{\mathrm{N}}^{\mathrm{fine}}$ setting.

5 The Impact of Adaptivity

5.1 Adaptivity Gap for the Follower

Let $\mathrm{OPT}_{\mathrm{A}}$ and $\mathrm{OPT}_{\mathrm{N}}$ the optimal values for $\textsc{Int}_{\mathrm{A}}$ and $\textsc{Int}_{\mathrm{N}}$ respectively. Correa et al. [5] showed that for their model (in which inspections are independent events) the ratio of the best non-adaptive strategy to the best adaptive strategy is bounded by 4/3. Indeed, their proof does not use the fact that arc inspections are independent events and thus translates to our setting.

Theorem 5 (Correa et al. [5]). $\mathrm{OPT}_{\mathrm{N}} \leq \frac{4}{3}\mathrm{OPT}_{\mathrm{A}}$.

In many real-life scenarios, it is reasonable to assume that the ratio of the length of the path chosen by the intruder to the shortest path (w.r.t. c) is not too large. E.g., most passengers in transit systems would pay a ticket rather than choosing a path with twice the transit time just in order to avoid inspection. We extend the proof by Correa et al. [5] to give a parameterized bound that takes this ratio into account and gives stronger guarantees for realistic values; also see Fig. 2.

Theorem 6. *If* $\mathrm{SP}_c(s,t) > 0$*, then* $\mathrm{OPT}_{\mathrm{N}} \leq \frac{\Delta^2}{2(1-\Delta)^{3/2}+3\Delta-2}\mathrm{OPT}_{\mathrm{A}}$*, where* $\Delta := \mathrm{SP}_c(s,t)/c(P^*)$ *and* P^* *is an optimal solution to* $\textsc{Int}_{\mathrm{A}}$.

Proof. We first observe that $\mathrm{OPT}_{\mathrm{N}} \leq \min\{\mathrm{SP}_c(s,t) + F,\, c(P^*) + \pi(P^*)F\}$ as both following the shortest path or following P^* are feasible non-adaptive strategies. On the other hand, observe that $\mathrm{OPT}_{\mathrm{A}} = f_{\mathrm{A},\pi}(P^*) \geq (1 - \pi(P^*))c(P^*) + \pi(P^*)(\mathrm{SP}_c(s,t) + F)$, as the total amount of transit cost will always be at least

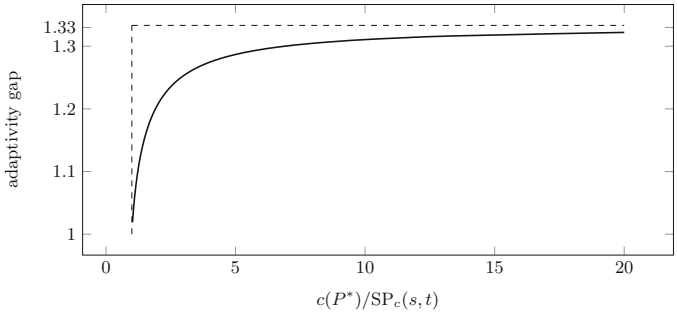

Fig. 2. Upper bound on the adaptivity gap $\mathrm{OPT_N}/\mathrm{OPT_A}$ given in Theorem 6 parameterized by $\Delta^{-1} = c(P^*)/\mathrm{SP}_c(s,t)$, where P^* is an optimal solution to $\textsc{Int}_\mathrm{A}$.

as much as the length of a shortest s-t-path. Defining $S := \mathrm{SP}_c(s,t)$, $C := c(P^*)$, and $Q := \pi(P^*)$, we obtain

$$\frac{\mathrm{OPT_A}}{\mathrm{OPT_N}} \geq \frac{(1-Q)C + Q(S+F)}{\min\{S+F,\, C+QF\}} = \frac{(1-Q)C + Q(\Delta C+F)}{\min\{\Delta C+F,\, C+QF\}}.$$

In order to prove the bound, we fix Δ and treat C, F, Q as variables of an optimization problem subject to $Q \in [0,1]$ and $F, C \geq 0$.

$$\frac{\mathrm{OPT_A}}{\mathrm{OPT_N}} \geq \min_{F,C\geq 0, Q\in[0,1]} \frac{(1-Q)C + Q(\Delta C+F)}{\min\{\Delta C+F,\, C+QF\}}.$$

It is easy to see that in an optimal solution, the minimum in the denominator is attained by both terms, i.e., $\Delta C + F = C + QF$. Substituting $F = \frac{1-\Delta}{1-Q}C$ we get

$$\frac{\mathrm{OPT_A}}{\mathrm{OPT_N}} \geq \min_{C\geq 0, Q\in[0,1]} \frac{(1-Q)C}{\left(1+Q\frac{1-\Delta}{1-Q}\right)C} + Q = \min_{Q\in[0,1]} \frac{(1-Q)^2}{1+\Delta Q} + Q.$$

By computing the derivative of the righthand side term, we observe that the minimum is attained at $Q = \frac{1-\sqrt{1-\Delta}}{\Delta}$, which gives the desired bound. □

5.2 Defender Gaps

We consider three gaps concerning the defender in the context of the zero-sum objective. Let π_A and π_N be the inspection probabilities that maximize the intruder's costs against an adaptive and non-adaptive intruder respectively, and let $f_X(\pi_Y) := \min_{P\in\mathcal{P}_{st}} f_{X,\pi_Y}(P)$ denote the defender's pay-off, where $X, Y \in \{\mathrm{A}, \mathrm{N}\}$.

Adaptivity Gap (η_A): This measures the defender's pay-off loss when the intruder is adaptive, as opposed to when she is non-adaptive.

Pay-off Gap (η_P): When the intruder is adaptive, this gap measures the deviation of the defender's pay-off from his own estimation if he wrongly assumes she is non-adaptive.

Approximation Gap (η_{App}): This is the approximation factor achieved by the defender against an adaptive intruder when playing the optimal strategy for non-adaptive intruders π_N as an approximation for π_A.

$$\eta_A = \frac{f_N(\pi_N)}{f_A(\pi_A)}, \quad \eta_P = \frac{f_N(\pi_N)}{f_A(\pi_N)}, \quad \eta_{App} = \frac{f_A(\pi_A)}{f_A(\pi_N)}.$$

As a straightforward consequence of Theorem 5, all of these gaps are upper bounded by 4/3.

6 Conclusion

In this paper, we investigated different variants of a Stackelberg network game in which the follower can gain and exploit information about the realization of the leader's random strategy while traversing the network. In the present work, we confined ourselves to the model in which a single arc is subjected to inspections. Future work will focus on the natural generalization in which several checkpoints are placed simultaneously and possibly in a correlated fashion, getting closer to real-world security scenarios.

Acknowledgements. This work was supported by the Alexander von Humboldt Foundation with funds of the German Federal Ministry of Education and Research (BMBF), by the Millennium Nucleus Information and Coordination in Networks Grant ICM/FIC RC130003, and by a CONICYT grant (CONICYT-PCHA/MagísterNacional/2014 - 22141563).

References

1. Adjiashvili, D., Oriolo, G., Senatore, M.: The online replacement path problem. In: Bodlaender, H.L., Italiano, G.F. (eds.) ESA 2013. LNCS, vol. 8125, pp. 1–12. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40450-4_1
2. Borndörfer, R., Omont, B., Sagnol, G., Swarat, E.: A Stackelberg game to optimize the distribution of controls in transportation networks. In: Krishnamurthy, V., Zhao, Q., Huang, M., Wen, Y. (eds.) GameNets 2012. LNICSSITE, vol. 105, pp. 224–235. Springer, Heidelberg (2012). doi:10.1007/978-3-642-35582-0_17
3. Brotcorne, L., Labbé, M., Marcotte, P., Savard, G.: A bilevel model for toll optimization on a multicommodity transportation network. Transp. Sci. **35**(4), 345–358 (2001)
4. Carstensen, P.J.: The complexity of some problems in parametric linear and combinatorial programming (1983)
5. Correa, J.R., Harks, T., Kreuzen, V.J.C., Matuschke, J.: Fare evasion in transit networks. Oper. Res. **65**(1), 165–183 (2017)
6. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, New York (2002)

7. Grötschel, M., Lovász, L., Schrijver, A.: Geometric Algorithms and Combinatorial Optimization. Algorithms and Combinatorics, vol. 2. Springer, Heidelberg (1988)
8. Hassin, R.: Approximation schemes for the restricted shortest path problem. Math. Oper. Res. **17**(1), 36–42 (1992)
9. Joret, G.: Stackelberg network pricing is hard to approximate. Networks **57**(2), 117–120 (2011)
10. Lorenz, D.H., Raz, D.: A simple efficient approximation scheme for the restricted shortest path problem. Oper. Res. Lett. **28**(5), 213–219 (2001)
11. Manfredi, P., Posta, P.D., d'Onofrio, A., Salinelli, E., Centrone, F., Meo, C., Poletti, P.: Optimal vaccination choice, vaccination games, and rational exemption: an appraisal. Vaccine **28**(1), 98–109 (2009)
12. Nikolova, E.V.: Strategic algorithms. Ph.D. thesis, Massachusetts Institute of Technology (2009)
13. Panda, S., Vorobeychik, Y.: Stackelberg games for vaccine design. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, pp. 1391–1399. International Foundation for Autonomous Agents and Multiagent Systems (2015)
14. Roch, S., Savard, G., Marcotte, P.: An approximation algorithm for Stackelberg network pricing. Networks **46**(1), 57–67 (2005)
15. Schulz, A.S., Uhan, N.A.: Approximating the least core value and least core of cooperative games with supermodular costs. Discrete Optim. **10**(2), 163–180 (2013)
16. Tambe, M.: Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. Cambridge University Press, Cambridge (2012)
17. Washburn, A., Wood, K.: Two-person zero-sum games for network interdiction. Oper. Res. **43**(2), 243–251 (1995)