A Proofs of Propositions and Theorems

A.1 Proof of Proposition 1

Proof. To prove this Proposition, it is sufficient to show that $U_d(\mathbf{x}, \mathbf{f}) \geq U_d(\mathbf{x}', \mathbf{f})$ for all $\mathbf{x}' \in \mathcal{X}$. Since $\langle \mathbf{x}, \mathbf{f} \rangle$ forms the Nash (Stackelberg) equilibrium of the restricted *NFIG* with restricted sets $\mathcal{S}'$ and $\mathcal{P}'$, we have:

$$U_d(\mathbf{x}, \mathbf{f}) \geq U_d(S', \mathbf{f}) \quad \forall S' \in \mathcal{S}'$$

Therefore, if $S \in \mathcal{S}'$, we have: $U_d(\mathbf{x}, \mathbf{f}) \geq U_d(S, \mathbf{f})$. For any mixed strategy $\mathbf{x}' \in \mathcal{X}$, we have:

$$U_d(\mathbf{x}', \mathbf{f}) = \sum_{S' \in \mathcal{S}} x_{S'} U_d(S', \mathbf{f}) \leq U_d(S, \mathbf{f})$$

since S is the best response pure strategy against $\mathbf{f}$ and $\sum_{S' \in \mathcal{S}} x_{S'} = 1$.

Thus, $U_d(\mathbf{x}, \mathbf{f}) \geq U_d(\mathbf{x}', \mathbf{f})$ for all $\mathbf{x}' \in \mathcal{X}$. □

A.2 Proof of Theorem 2

Proof. We show the NP-hardness of the Column Generation problem by reducing the set cover problem to it. **The Set-Cover problem:** Given a set U, a collection $\mathcal{Q}$ of subsets of U (that is, $\mathcal{Q} \subseteq 2^U$), and an integer k. The question is whether there is a cover $\mathcal{C} \subseteq \mathcal{Q}$ of size k or less, that is, $\bigcup_{q \in \mathcal{C}} q = U$ and $|\mathcal{C}| \leq k$.

Given an instance $\langle U, \mathcal{Q}, k \rangle$ of set cover problem, we reduce it to a *ColG* instance $\langle G, \mathcal{P}', I, \mathbf{f} \rangle$ as follows: Except the two unique source and sink nodes s, t, for each subset $q \in \mathcal{Q}$, we assign a vertex v_q in G and an inspection station i_q on vertex v_q with inspection probability $\tau_{i_q} = 1$. Now, for each element $u \in U$, we construct a path $p_u \in \mathcal{P}'$ as follows: p_u starts from s, passing through vertices v_q such that $u \in q$, and end at t; The attacker flow $\mathbf{f}$ is defined on paths of $\mathcal{P}'$ such that $f_p = 1$ for all $p \in \mathcal{P}'$. Obviously, this reduction can be done in polynomial time. Next, we show that the set cover problem instance $\langle U, \mathcal{Q}, k \rangle$ has a cover $\mathcal{C}$ of size k or less if and only if the Column Generation problem has a solution S such that all the flow in $\mathbf{f}$ is intervened, i.e., $U_a(S, \mathbf{f}) = 0$.

The "if" direction: If S intervenes all the flow in $\mathbf{f}$, which means each path p_u in $\mathcal{P}'$ passes through some inspection station i_q in S, then the collection $\mathcal{C} = \{q \in \mathcal{Q} : i_q \in S\}$ is a cover of size k or less.

The "only if" direction: If there exists a cover $\mathcal{C}$ of size k or less for the set cover problem, the defender pure strategy $S = \{v_q : q \in \mathcal{C}\}$ intervenes all the paths $p_u \in \mathcal{P}'$ and hence all the flow in $\mathbf{f}$ is blocked. □

A.3 Proof of Theorem 3

Proof. We provide a dynamic programming algorithm for solving the *ColG* subproblem on tree-like graph G with source node s in polynomial time. First, we transform G into a tree G' in following way: for sink node t, supposing its incoming edges are $(v_1, t), ..., (v_m, t)$, we replace t with m new sinks $t_1, ..., t_m$ and replace incoming edges with m new edges $(v_1, t_1), ..., (v_m, t_m)$. Figure 3 shows an example of the tree G' transformed from tree-like graph G. Let s be the root of tree G'. Given a vertex v, let T_v be the subtree rooted at v, C_v be the set of child nodes of v, c_v^i be the i-th child node of v, and T_v^j be the subtree rooted at v with first j branches.

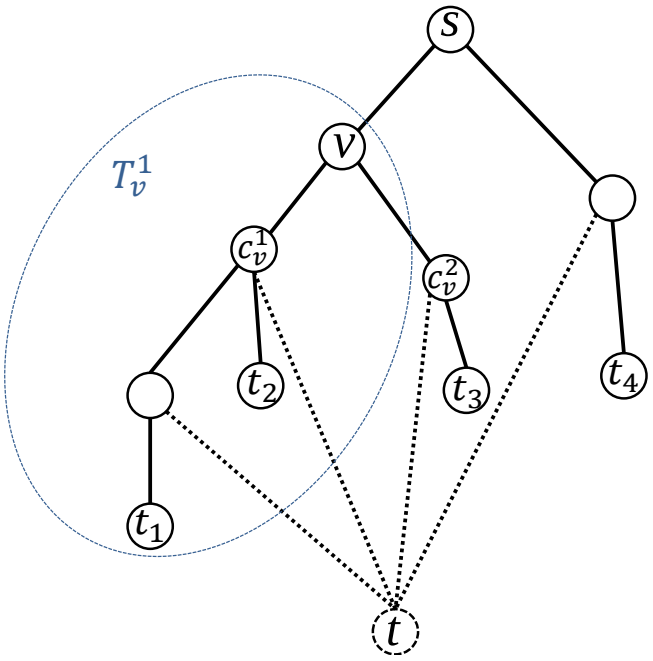

Figure 3: Dynamic Programming for Tree-Like Graph.

Let $S_v(r, j)$ denote the optimal allocation of r resources on T_v^j, and $S_v(r) = S_v(r, |C_v|)$ be the optimal allocation of r resources on T_v. A key property here is that since all flows passing through T_v^j come from the unique path from s to v (G' is a tree), the resource allocation outside T_v^j will decrease f_p with the same proportion for all paths p from s to $v \in T_v^j$. Thus, the optimal resource allocation $S_v(r, j)$ is independent with allocation of resources on the complementary part $G' - T_v^j$.

The above property leads to the following recursion:

$$S_v(r, j+1) = \arg\min_{r' < r} \min \begin{cases} U(S_v(r - r', z)) + U(S_{v'}(r')) \\ U(S_v(r - r'), z) + k_v^{v'} U(S_{v'}(r' - 1)) \end{cases} \tag{10}$$

where $U(S_v(r, j))$ measures the amount of flow passing through T_v^j with resource allocation $S_v(r, j)$ on T_v^j and no resource put outside T_v^j, $v' = c_v^{j+1}$, i.e., the $j+1$-th child node of v, and $k_v^{v'}$ be the inspection probability of the checkpoint on edge (v, v'). The idea of this recursion is that allocating r resources on T_v^{j+1} can be done in two possible ways: i) putting $r - r'$ resources on T_v^j and the rest on $T_{v'}$; ii) putting $r - r'$ resources on T_v^j, $r' - 1$ resources on $T_{v'}$, and one resource on (v, v'). The best way minimizing the amount of flow passing through T_v^j gives $S_v(r, j+1)$. Here, we assume that checkpoints are all on edges, which is reasonable because in the tree graph, each node v has a unique parent node $\tilde{v}$. Thus, we can always insert a new node v' at the middle of edge $(\tilde{v}, v)$ and move the checkpoint on v to the new edge (v', v), while keeping the game unchanged.

With the above recursion, we can compute $S_v(r, j)$ for every nodes following direction from sink nodes to s with r from 1 to the budget k and j from 1 to $|C_v|$. Thus, the algo-

rithm computes the optimal solution $S^* = S_s(k)$ for *ColG* in $O(k|V|^2)$ runtime.

□

A.4 Proof of Proposition 4

Proof. To prove the convexity of $U_a(S, \mathbf{f})$, it is enough to prove that $\Phi(S, p)$ is convex. To do so, we provide an equivalent representation of $\Phi(S, p)$ defined in the following equation:

$$\Phi(S, p) = \prod_{i \in I} (1 - \tau_i)^{\alpha_{pi} S_i},$$

where $\alpha_{pi} = 1$ if path p passes through inspection station i and $\alpha_{pi} = 0$ otherwise. The Hessian matrix of $\Phi(S, p)$ is denoted by $\mathbf{H}$, such that:

$$H_{ij} = \frac{\partial^2 \mathbf{H}}{\partial S_i \partial S_j} = \alpha_{pi}\alpha_{pj} ln(1 - \tau_i) ln(1 - \tau_j) \Phi(S, p)$$

It is easy to verify that $\mathbf{H}$ can be represented as:

$$\mathbf{H} = \mathbf{y}\mathbf{y}^{\mathbf{T}} \Phi(S, p)$$

where $y_i = \alpha_{pi} ln(1 - \tau_i)$. Therefore, $\mathbf{H}$ is positive semidefinite, and $U_a(S, \mathbf{f})$ is convex function over $S \in [0, 1]^{|I|}$. □

A.5 Proof of Theorem 5

Proof. The defender utilities are always negative, therefore an approximation guarantee in terms of utility is meaningless. To facilitate analysis, we first define a non-negative function h for each defender allocation S such that $|S| \leq k$:

$$h_{\mathbf{f}}(S) = U_d(S, \mathbf{f}) - U_d(\emptyset, \mathbf{f}) = \sum_{p \in \mathcal{P}} (1 - \Phi(S, p)) f_p$$

Note h denotes the marginal benefit of allocation S over the non-defending defender strategy.

Given the adversarial flow $\mathbf{f}$, we first prove that $h_{\mathbf{f}}(S)$ is a monotone submodular function, i.e., i) $h_{\mathbf{f}}(S_1) \leq h_{\mathbf{f}}(S_2)$, and ii) $h_{\mathbf{f}}(S_1 \cup \{i\}) - h_{\mathbf{f}}(S_1) \geq h_{\mathbf{f}}(S_2 \cup \{i\}) - h_{\mathbf{f}}(S_2)$, for each $S_1 \subseteq S_2 \subseteq I, i \in I \setminus S_2$.

Monotone Remember $\Phi(S, p)$ represents the proportion of adversary flow on path p not interdicted by the operated inspection stations given S. Given $S_1 \subseteq S_2$, we have $\Phi(S_1, p) \geq \Phi(S_2, p)$ for all $p \in \mathcal{P}$. Thus, $h_{\mathbf{f}}(S_1) \leq h_{\mathbf{f}}(S_2)$.

Submodulariry For each $S \subseteq I$ and $i \in I \setminus S$, we have:

$$h_{\mathbf{f}}(S \cup \{i\}) - h_{\mathbf{f}}(S) = \sum_{p \in \mathcal{P}} (\Phi(S, p) - \Phi(S \cup \{i\}, p)) f_p$$

According to Eq.(2):

$$\Phi(S, p) = \prod_{i \in I : i \in p} (1 - \tau_i)^{S_i}$$

We have:

$$h_{\mathbf{f}}(S \cup \{i\}) - h_{\mathbf{f}}(S) = \sum_{p \in \mathcal{P} : i \in p} \tau_i \cdot \Phi(S, p) f_p$$

Since $\Phi(S_1, p) \geq \Phi(S_2, p)$ for each $S \subseteq I$ and $i \in I \setminus S$ and $p \in \mathcal{P}$, we get: $h_{\mathbf{f}}(S_1 \cup \{i\}) - h_{\mathbf{f}}(S_1) \geq h_{\mathbf{f}}(S_2 \cup \{i\}) - h_{\mathbf{f}}(S_2)$.

Therefore, $h_{\mathbf{f}}(S)$ is a non-negative monotone submodular function. According to [Nemhauser and Wolsey, 1978], the greedy algorithm for submodular function maximization can achieve a $1 - \frac{1}{e}$ approximation ratio. Hence, $U_d(S', \mathbf{f}) - U_d(\emptyset, \mathbf{f}) \geq (1 - \frac{1}{e})(U_d(S^*, \mathbf{f}) - U_d(\emptyset, \mathbf{f}))$. □

A.6 Proof of Theorem 6

Proof. Let S' be the greedy solution of Algorithm 2 against attacker's network flow $\mathbf{f}'$, and S^* be the defender best response pure strategy against $\mathbf{f}'$, i.e.,

$$S^* = \arg\max_{S \in \mathcal{S}} U_d(S, \mathbf{f}') \tag{11}$$

According to Theorem 5, we have:

$$U_d(S', \mathbf{f}') - U_d(\emptyset, \mathbf{f}') \geq (1 - \frac{1}{e})(U_d(S^*, \mathbf{f}') - U_d(\emptyset, \mathbf{f}'))$$

Since $\mathbf{x}'$ is the defender best response strategy against $\mathbf{f}'$ in the restricted *NFIG* defined on restricted defender pure strategy space $\mathcal{S}'$, and $S' \in \mathcal{S}'$ according to the convergence criteria of *CCG* algorithm, we have:

$$\begin{aligned} &U_d(\mathbf{x}', \mathbf{f}') - U_d(\emptyset, \mathbf{f}') \\ &\geq U_d(S', \mathbf{f}') - U_d(\emptyset, \mathbf{f}') \\ &\geq (1 - \frac{1}{e})(U_d(S^*, \mathbf{f}') - U_d(\emptyset, \mathbf{f}')) \end{aligned} \tag{12}$$

According to Eq.(11), we have:

$$U_d(\mathbf{x}^*, \mathbf{f}') = \sum_{S \in \mathcal{S}} x^*_S U_d(S, \mathbf{f}') \leq U_d(S^*, \mathbf{f}')$$

Since $\mathbf{f}^*$ is the attacker best response flow against $\mathbf{x}^*$ in the original *NFIG*, we have:

$$U_d(\mathbf{x}^*, \mathbf{f}^*) \leq U_d(\mathbf{x}^*, \mathbf{f}') \leq U_d(S^*, \mathbf{f}') \tag{13}$$

Finally, substitute Eq.(12) with Eq.(13):

$$\begin{aligned} &U_d(\mathbf{x}', \mathbf{f}') - U_d(\emptyset, \mathbf{f}') \\ &\geq (1 - \frac{1}{e})(U_d(\mathbf{x}^*, \mathbf{f}^*) - U_d(\emptyset, \mathbf{f}')) \end{aligned}$$

□

A.7 Proof of Proposition 7

Proof. Supposing the optimal solution of $CoreLP(\mathcal{S}', \mathcal{P}')$ is $(\mathbf{x}, \mathbf{u})$, then the current equilibrium flow $\mathbf{f}$ over $\mathcal{P}'$, which is the dual solution with respect to inequality (6c) on paths in $\mathcal{P}'$, satisfies

$$U_a(\mathbf{x}, \mathbf{f}) = \sum_{e \in E} c_e u_e$$

based on the dual theory [Bertsimas and Tsitsiklis, 1997]. If the ConG generates an s-t path p with maximal reduced cost such that $R(p) \leq 0$, we have

$$R(p') \leq R(p) \leq 0 \quad \forall p' \in \mathcal{P}$$

In other words, $(\mathbf{x}, \mathbf{u})$ is also the optimal solution of $CoreLP(\mathcal{S}', \mathcal{P})$ as it satisfies inequality (6c) with respect to all $p \in \mathcal{P}$. Notice that $CoreLP(\mathcal{S}', \mathcal{P})$ computes the attacker

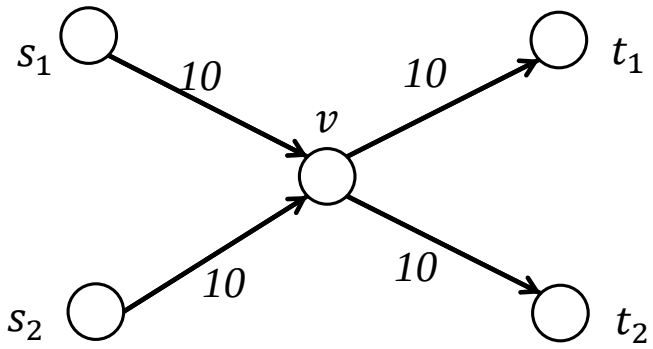

Figure 4: Example of Compact Representation.

best response flow $\mathbf{f}^*$ against defender mixed strategy $\mathbf{x}$ over pure strategy set $\mathcal{S}'$ also satisfying

$$U_a(\mathbf{x}, \mathbf{f}^*) = \sum_{e \in E} c_e u_e$$

according to the dual theory. Thus, $U_a(\mathbf{x}, \mathbf{f}) = U_a(\mathbf{x}, \mathbf{f}^*)$, and the current equilibrium flow $\mathbf{f}$ is an attacker best response pure strategy against $\mathbf{x}$. □

B Infeasibility of Compact Representation for the Attacker Strategy

We demonstrate the infeasibility of compact representation for the attacker's flow with a simple example shown in Figure 4 with 2 sources and two sinks. With compact representation, the flow $\mathbf{f}$ is defined over edges E such that f_e denotes the amount of flow passing through edge e. In this example, $f_e = 10$ for all edges. Now, suppose the defender allocates two resources on edges (s_1, v) and (v, t_1) with inspection probabilities 1. In this case, we cannot compute the exact amount the flow interdicted by the defender. It can be 10 units passing through $s_1 - v - t_2$ and 10 units passing through $s_2 - v - t_1$, and the attacker utility is 0 in this case; It may also be 10 units passing through $s_1 - v - t_1$ and 10 units passing through $s_2 - v - t_2$, and the attacker utility is 10 now. The reason for this kind of infeasibility comes from the fact that each compact representation can be realized by numerous possible flows defined over paths, and they can have different utilities against a given resource allocation. Thus, the network flow in our *NFIG* is defined over paths.

C Southern Border Patrol Networks

We conduct experiments on two sectors in Southern Border Patrol [GAO, 2005]: i) San Diego sector with 29 nodes and degree of 2.89, two source nodes (San Ysidro and Tecate POEs), one sink node (Los Angeles), and 12 checkpoints (Figure 5); ii) Laredo sector with 20 nodes and degree of 3, one source node (Laredo POE), one sink node (San Antonio), and 6 checkpoints (Figure 6). The roads can be classified into four levels: Interstate, U.S. Hwy, State Hwy, and Hwy. The capacity of a road is randomly drawn from: [0, 35] (Interstate), [0, 30] (U.S. Hwy), [0, 25] (State Hwy), and [0, 20] (Hwy). $k = 6$ for San Diego sector and $k = 3$ for Laredo sector.

Figure 5: San Diego Sector (circles are tactical checkpoints).

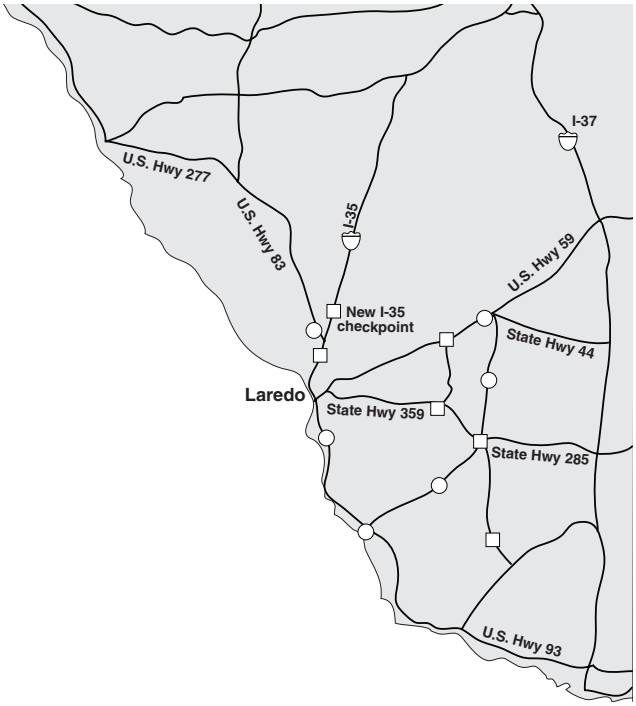

Figure 6: Laredo Sector (circles are tactical checkpoints).