A robust optimization approach for placing virtual network functions to counter cyberattacks in a 5G network

Céline Gicquel[1], Sonia Vanier[2], Alexandros Papadimitriou[3]

[1] Laboratoire Interdisciplinaire des Sciences du Numérique, Université Paris Saclay, France
celine.gicquel@universite-paris-saclay.fr

[2] Laboratoire d'Informatique de l'Ecole Polytechnique, Université Paris 1, France
vanier@lix.polytechnique.fr

[3] Orange Labs Products & Services, France

Mots-clés : *Cybersecurity, Network function virtualization, Robust optimization*

1 Introduction

A distributed denial of service (DDoS) is a type of cyberattack in which multiple compromised computer systems attack a target, such as a server or a website, and cause a denial of service for its legitimate users. DDoS attacks are among the top threats to network operators and internet service providers (ISPs) as they can be very damaging for the organization they target.

Network Function Virtualization (NFV) is a recent network architecture concept in which network functions are implemented as software and deployed as virtual machines running on general purpose commodity hardware. NFV offers new possibilities to counter DDoS attacks. In particular, its flexibility and reactivity allows to postpone the DDoS defense deployment after the attack is detected and to adjust the deployed defense mechanisms to the type and scale of the DDoS attack. We take here the perspective of an ISP such as Orange aiming at providing a NFV-based DDoS mitigation service to its customers in a 5G network. Among the key features of 5G networks is network slicing. Network slicing is an architecture in which the physical network infrastructure managed by an ISP is partitioned into multiple virtual independent networks termed slices. Each slice is an isolated end-to-end network which is lent by the ISP to a single customer. On each slice of the network, the routing of the flow will be managed by its customer which will rely on its own proprietary routing algorithms. This significantly enhances the difficulty for the ISP of providing a DDoS mitigation service as it will not control the exact routing of the malicious flow that needs to be stopped.

We thus present a robust optimization (RO) model to optimally design an NFV-based DDoS mitigation infrastructure in the context of 5G network slicing. This model explicitly takes into account the fact that the ISP is not aware of the exact routing of the attack flow and considers the malicious flow routing as an uncertain input parameter.

2 Problem modeling

The network topology is modeled by a digraph $G = (\mathcal{N}, \mathcal{L})$ in which $\mathcal{N}$, the set of nodes, represents specific equipment in the network and $\mathcal{L}$, the set of arcs, corresponds to the links that can be used to route the traffic. The routing of the traffic in the network is limited by the bandwidth b_l of each link l.

The illegitimate traffic corresponding to the on-going DDoS attack is represented as a set $\mathcal{A}$ of attacks : attack $a \in \mathcal{A}$ corresponds to an illegitimate traffic of F^a Mbps between a source $s^a \in \mathcal{N}$ and the target $t \in \mathcal{N}$ of the DDoS attack. The exact routing in the network of the malicious flow to be stopped is not known by the ISP at the time when it has to decide about the NFV-based DDoS mitigation infrastructure. Let $\mathcal{P}^a$ be the set of all potential paths between s^a and t for attack a. The amount of malicious flow of attack $a \in \mathcal{A}$ on path $p \in \mathcal{P}^a$, denoted by $\tilde{f}_p^a$, is thus subject to uncertainty. However, we know that the total amount of malicious flow routed on the paths belonging to $\mathcal{P}^a$ may not be greater than F^a, the amount of illegitimate traffic of attack a, and that it should comply with the limited bandwidth of each link. This means that the uncertain malicious flow routing, $\tilde{f}$, belongs to the uncertainty set $\mathcal{U}$ defined by : $\mathcal{U} = \{\tilde{f} \geq 0 | \sum_{p \in \mathcal{P}^a} \tilde{f}_p^a \leq F^a, \forall a \in \mathcal{A}; \sum_{a \in \mathcal{A}} \sum_{p \in \mathcal{P}^a \text{ s.t. } l \in \mathcal{L}_p^a} \tilde{f}_p^a \leq b_l, \forall l \in \mathcal{L}\}$.

Virtual Network Functions (VNFs) are used to stop the illegitimate traffic before it reaches its target. A VNF instantiated on a node $n \in \mathcal{N}$ of the network will thus stop the malicious flow going through n. However, it has a limited capacity which corresponds to the maximum amount of malicious flow it can block. A VNF of type $v \in \mathcal{V}$ is characterized by its filtering capacity ϕ^v, its cost K^v and its computing resources consumption. The set of computing resources (CPU, memory, etc.) is denoted by $\mathcal{R} = \{1, ..., R\}$. Let k^{rv} be the amount of computing resource r required by the instantiation of one VNF of type v and Cap_n^r the amount of computing resource r available at node n.

The optimization problem consists in identifying the location and number of VNFs to be placed in the network so as to stop all the malicious flow before it reaches its target, and this whatever its routing through the network, while minimizing the cost of the instantiated VNFs. Let x_n^v be the decision variables representing the number of VNFs of type v placed at node n. The problem can be formulated as follows :

$$Z^* = \min \sum_{v \in \mathcal{V}} \sum_{n \in \mathcal{N}} K^v x_n^v \tag{1}$$

$$\sum_{v \in \mathcal{V}} k^{rv} x_n^v \leq Cap_n^r \qquad \forall n \in \mathcal{N}, \forall r \in \mathcal{R} \tag{2}$$

$$\sum_{n \in \mathcal{N}(\tilde{f})} \sum_{v \in \mathcal{V}} \phi^v x_n^v \geq \sum_{a \in \mathcal{A}} \sum_{p \in \mathcal{P}^a} \tilde{f}_p^a \qquad \forall \tilde{f} \in \mathcal{U} \tag{3}$$

$$x_t^v = 0 \qquad \forall v \in \mathcal{V} \tag{4}$$

$$x_n^v \text{ integer} \qquad \forall n \in \mathcal{N}, \forall v \in \mathcal{V} \tag{5}$$

The objective (1) is to minimize the total costs of the deployed VNFs. Constraints (2) ensure that the VNFs installed at each node n do not consume more than the available computing capacity for each computing resource. Constraints (3) impose that, for each possible routing $\tilde{f}$, the total filtering capacity installed on the nodes traversed by a strictly positive amount of malicious flow in the routing $\tilde{f}$, i.e on the nodes belonging to $\mathcal{N}(\tilde{f})$, is larger than the total malicious flow actually routed through the network in $\tilde{f}$. Constraints (4) forbid any filtering at the targeted node.

3 Solution approach and computational results

We propose an adversarial approach based on the decomposition of the initial problem into a master problem and an adversarial sub-problem. The master problem is a restricted version of the original RO problem in which only a finite number of extreme points $\mathcal{U}_R \subset \mathcal{U}$ of the uncertainty set are used to express the robust constraints. This problem is a deterministic optimization problem with a finite number of constraints and is thus computationally tractable. Given the solution provided by the master problem, the adversarial sub-problem seeks to find an extreme point of $\mathcal{U}$ for which this solution is infeasible. If no such extreme point can be found, the current solution of the decision maker problem is optimal for the initial RO problem. Otherwise, we add the newly found extreme point to $\mathcal{U}_R$ and reiterate the process. The finite convergence of this algorithm is ensured by the fact that $\mathcal{U}$ has a finite number of extreme points. Our computational results show that the proposed algorithm is able to provides optimal solutions to medium-size randomly generated instances within short computation times.