

Dynamic and Application-Aware Provisioning of Chained Virtual Security Network Functions

Roberto Doriguzzi-Corin^α, Sandra Scott-Hayward^β, Domenico Siracusa^α, Marco Savi^α, Elio Salvadori^α

^αCREATE-NET, Fondazione Bruno Kessler - Italy

^βCSIT, Queen's University Belfast - Northern Ireland

Abstract—A promising area of application for Network Function Virtualization is in network security, where chains of Virtual Security Network Functions (VSNFs), i.e., security-specific virtual functions such as firewalls or Intrusion Prevention Systems, can be dynamically created and configured to inspect, filter or monitor the network traffic. However, the traffic handled by VSNFs could be sensitive to specific network requirements, such as minimum bandwidth or maximum end-to-end latency. Therefore, the decision on which VSNFs should apply for a given application, where to place them and how to connect them, should take such requirements into consideration. Otherwise, security services could affect the quality of service experienced by customers.

In this paper we propose PESS (Progressive Embedding of Security Services), a solution to efficiently deploy chains of virtualised security functions based on the security requirements of individual applications and operators' policies, while optimizing resource utilization. We provide the PESS mathematical model and heuristic solution.

Simulation results show that, compared to state-of-the-art application-agnostic VSNF provisioning models, PESS reduces computational resource utilization by up to 50%, in different network scenarios. This result ultimately leads to a higher number of provisioned security services and to up to a 40% reduction in end-to-end latency of application traffic.

Index Terms—Network Function Virtualization, Network Service Chaining, Progressive Embedding, Application-Aware Network Security.

I. INTRODUCTION

Network security implemented by Telecommunication Service Providers (TSPs) has traditionally been based on the deployment of specialized, closed, proprietary Hardware Appliances (HAs). Such HAs are inflexible in terms of functionalities and placement in the network, which means that even slight changes in the security requirements generally necessitate manually intensive and time-consuming re-configuration tasks, the replacement of existing HAs or the deployment of additional HAs.

The Network Function Virtualization (NFV) [1] initiative has been proposed as a possible solution to address the operational challenges and high costs of managing proprietary HAs. The main idea behind NFV is to transform network functions (e.g. firewalls, intrusion detection systems etc.) based on proprietary HAs, into software components (called Virtual Network Functions (VNFs)) that can be deployed and executed in virtual machines on commodity, high-performance servers. By decoupling software from hardware, this approach allows any (security) network function to be deployed in any server

connected to the network through an automated and logically centralized management system.

The centralized management system, called NFV Management and Orchestration (NFV MANO), controls the whole life-cycle of each VNF. In addition, the NFV MANO can dynamically provision complex network services in the form of sequences (often called chains) of VNFs. Indeed, Network Service Chaining (NSC) is a technique for selecting subsets of the network traffic and forcing them to traverse various VNFs in sequence. For example, a firewall followed by an Intrusion Prevention System (IPS), then a Network Address Translation (NAT) service and so on. NSC and NFV enable flexible, dynamic service chain modifications to meet the real time network demands.

A promising area of application for NSC and NFV is in network security, where chains of Virtual Security Network Functions (VSNFs), i.e., security-specific VNFs such as a firewall or an IPS, can be dynamically created and configured to inspect, filter or monitor the network traffic. The flexibility of the NSC and NFV paradigms brings many benefits, among others: (i) highly customizable security services based on the needs of the end-users, (ii) fast reaction to new security threats or variations of known attacks, and (iii) low Operating Expenditure (OpEx) and Capital Expenditure (CapEx) for the operator. On the other hand, compared to specialized HAs, VSNFs may have a significant impact on the performance of the network and on the Quality of Service (QoS) level experienced by the users. The virtualization overhead, the utilization level of the servers and the techniques adopted to implement the VSNFs are the most significant contributors to the QoS degradation.

We argue that, for a wide adoption of NSC and NFV technologies in network security, the provisioning strategies should take into account not only the security requirements, but also specific QoS needs of applications (a list of common classes of applications is provided in Table I with their corresponding security and QoS requirements). Omitting the latter may lead, for instance, to a model that blindly forces all the user traffic to traverse the whole chain of VSNFs. As a result, computationally demanding VSNFs such as IPSs may cause a noticeable performance degradation to latency-sensitive applications (e.g. online games [5]) or bandwidth sensitive applications (e.g. video streaming). On the other hand, beyond the overall resource consumption and QoS requirements, the model must also take into account the specific

TABLE I
SECURITY AND QoS REQUIREMENTS OF APPLICATIONS.

Application class	Description	Related threats	Relevant VSNF ¹	QoS requirements
CCTV systems	Closed Circuit TV for video surveillance accessible remotely	Port scanning, DDoS password cracking	Firewall, DPI, IDS, IPS	Bandwidth: 10Mbps (5 cameras, 720p, 15fps, H.264, medium quality) Latency: 200ms (PTZ ² two-way latency [2])
Email	Electronic mail	Malware, spam, phishing, data exfiltration	DPI, Antispam, IDS, DLP	–
Instant messaging	Real-time text-based Internet chat	Malware, DDoS, phishing (out-of-band)	DPI, Antispam, IDS, IPS	–
Media streaming	Audio/video content accessed over the Internet	Inappropriate content	Parental control	Bandwidth ³ 5Mbps (HD) 25Mbps (UHD)
Remote storage	File transfer over the network	Data exfiltration	VPN, Data Encryption	Bandwidth
Network services (DNS, VoD, file sharing, WWW, SSH)	Server application accessed by remote client applications	DDoS, SQL injection, remote code execution	Firewall, IDS, WAF, HoneyPot	–
Online gaming	Video games played over the Internet	Online game cheating (out-of-band attacks) DDoS (in-band attacks)	DPI, Antispam, IDS, IPS	Latency: 100ms (first-person games [3])
Peer-to-peer	File sharing over peer-to-peer networks	DDoS, malware	DPI, IDS, IPS	–
Video conferencing	Real-time audio/video over the Internet	DDoS	Firewall, IPS	Latency: 150ms [4]
Web browsing	Applications for browsing the World Wide Web	Cross-site scripting, phishing, malware, inappropriate content	DPI, WAF, Parental control	Latency: 400ms [4]

¹Acronyms of VSNFs: Deep Packet Inspection (DPI), Intrusion Detection System (IDS), Intrusion Prevention System (IPS), Data Loss/Leakage Prevention (DLP), Virtual Private Network (VPN), Web Application Firewall (WAF). ²PTZ: Pan, Tilt, Zoom. ³Netflix Internet connection speed recommendations.

security best practices and policies. Omitting such aspects may result in the inappropriate placement of a firewall in the middle of the network, thus allowing unauthorized traffic to reach hosts that should be protected.

In this paper, we present PESS (Progressive Embedding of Security Services), a novel approach to provision security services by composing chains of VSNFs according to the specific QoS needs of user applications and the security policies defined by the TSP. TSP's security policies (given as an input to PESS) include: the kind of VSNFs (e.g., firewall, IPS, etc.) that should be deployed for a specific class of applications, their order (e.g., firewall first, then IPS, etc.), and more (e.g., a parental control should be installed close to the user's premises).

PESS defines an Integer Linear Programming (ILP) formulation and a heuristic algorithm to tackle the provisioning problem in dynamic network scenarios, where the service requests are not known in advance. In contrast, advance knowledge of service requests is assumed by the majority of related works. Although the PESS formulation and implementation presented in this paper focus on security-specific services, the proposed approach is also suitable for more complex scenarios, where heterogeneous network services provided by means of generic VNFs coexist (e.g., security, video broadcasting, content caching, etc.).

To the best of our knowledge, this work is the first attempt to tackle the challenges of progressively provisioning application-aware security services onto operational infrastructures, in which the QoS performance of existing services may be compromised by adding new services.

A preliminary ILP formulation of our application-aware approach has been introduced in [6]. This work extends [6] with the following contributions:

- A taxonomy of the most relevant VSNFs along with a list of popular implementations available in the open-source

community and in today's market.

- A mathematical formulation for the progressive provisioning of security services (the PESS ILP model). With respect to the ILP model presented in [6], this work formulates the estimation of the processing delay based on residual computing resources of the physical nodes, and it formalizes the impact on the end-to-end latency of operational services when allocating computing resources for new requests.
- A heuristic algorithm, called PESS heuristic, to obtain near-optimal solutions of the embedding problem in an acceptable time frame (in the order of a few milliseconds even in large network scenarios).
- An evaluation of the heuristic's performance in terms of quality of the solutions (deviation from optimality) and scalability performed on real-world and randomly generated topologies.
- Comparison of the proposed application-aware approach against the baseline approach, in which security services are provided without taking into account the specific requirements of applications.

Our simulations prove that the PESS solution can deploy more security services over the same infrastructure compared to an application-agnostic approach (the baseline), while still respecting the security policies and best practices defined by the TSP. With PESS, each traffic flow generated by each application can be served by the strict subset of VSNFs that are necessary to ensure its security, which means that no flow is burdened with any unnecessary security function that could affect its smooth execution, as would be the case with an application-agnostic approach. Of course, the capability of the VSNFs to properly react to any security attack depends on the specific implementation of the VSNF itself and is beyond the scope of this work.

The remainder of this paper is structured as follows: Section

II reviews and discusses the related work. Section III gives the relevant background information. Section IV provides the motivation behind this work. Section V details the mathematical formulation of the ILP model, while Section VI describes the heuristic algorithm that we implemented to solve the problem. In Section VII, the heuristic algorithm is evaluated on real-world and random topologies. Finally, the conclusions are provided in Section VIII.

II. RELATED WORK

With the recent “softwarisation” of network resources, a plethora of research initiatives has emerged in the last few years to address the problem of the optimal placement of chained VNFs. Most of these tackle the problem by using linear programming techniques and by proposing heuristic algorithms to cope with large scale problems. In this section, we classify and review the most relevant works for our studies.

A. QoS-driven VNF Placement

QoS-driven approaches primarily focus on the QoS requirements of specific services without considering network security aspects. In this regard, the proposed mathematical models include bandwidth and latency constraints (similar to Constraints (10) and (12) presented in Section V) or define objective functions that require minimization of the total bandwidth and latency of created chains.

The ILP model in [7] considers computing and bandwidth constraints to minimize the costs related to (i) VNF deployment, (ii) energy consumption of the servers, and (iii) forwarding traffic. The end-to-end delay requirement is formulated as a penalty in the objective function. However, the computation of the end-to-end delay only considers link propagation delays without including the processing delay introduced at each VNF. In [8], the placement problem is formulated as a Mixed Integer Quadratically Constrained Problem with respect to bandwidth, number of used nodes and latency. The processing delay at each VNF is also not considered in this work. The study in [9] proposes an ILP formulation and a heuristic algorithm for the VNF placement problem focusing on QoS parameters such as end-to-end delay and NSC availability. The ILP model formulation presented in the paper does not discuss how the processing delay introduced by the VNFs is computed. This limitation is reflected in the assumptions made for the evaluation, where the processing delay is considered independent from the VNF type/implementation and from the computing capacity of the physical node where VNFs are placed. In [10], Tajiki et al. present a resource allocation architecture for softwarized networks. The proposed architecture includes two resource allocation modules whose goal is configuring the network while satisfying QoS constraints and optimizing the energy consumption and the number of flow entries in the network. Although the authors tackle the problem of progressively allocating resources for newly arrived flows, neither the ILP formulation nor the heuristic algorithm consider the effects of the resource allocation on servers whose computing capacity is close to the limit. As discussed in Sections V-B and VI, this may lead to a degradation of the

QoS of existing services in terms of higher end-to-end latency.

B. Placement of VNFs/VSNFs

In addition to the research work on QoS-driven VNFs placement, there are a number of works that specifically consider the placement of VSNFs.

The method proposed in [11] is based on light-weight, protocol-specific intrusion detection VNFs. The system dynamically invokes a chain of these IDSs according to the traffic characteristics. The placement of the chains is based on a user-defined or common shortest-path algorithm such as Dijkstra, without consideration of the application QoS requirements or available network/computing resources.

In [12], the authors argue that reactive mechanisms used by cloud providers to deploy VSNFs do not ensure an optimal resource allocation. To address this, the authors propose a novel resource allocation scheme, which estimates the behaviour of the traffic load by monitoring the history of the current VSNFs, and pro-actively provisions new instances of those VSNFs as a countermeasure to any incoming resource pressure. The proposed algorithm does not tackle the problem of VSNF chaining. Instead, it focuses on the optimal placement of new instances of VSNFs, which are part of existing chains. It also assumes infinite network and computing resources.

In [13], Dermici et al. tackle the VSNFs placement problem by proposing an ILP formulation whose objective is the minimization of the energy consumption of servers. This solution does not consider any security nor QoS constraints. The aim of the NSC embedding model presented in [14] is to minimize the end-to-end latency of cross-domain chains of VSNFs. The main limitation of the proposed ILP formulation is that it only considers link propagation delays, while ignoring the processing delay introduced at each VSNF.

C. Security-driven VSNF Placement

Although the literature reviewed in Section II-B addresses the placement of VSNFs, few solutions have been proposed with a focus on the network security requirements of the VSNF placement. In [15], the authors propose a model for the placement of VSNFs that takes into account security deployment constraints. Such constraints are necessary to avoid incorrect deployment of security functions such as placing an IDS on an encrypted channel. The authors propose an ILP formulation of the problem and validate their model by measuring the execution time in four different scenarios and by comparing the model with other heuristics in terms of placement cost. However, the proposed optimization algorithm is always computed for all flows in the network. Therefore, it does not scale well. The authors mitigate the problem by partitioning the network into independent blocks. Nevertheless, the partitioning scheme is limited to fat-tree topologies. Furthermore, the end-to-end latency is not considered among the constraints of the proposed model, which limits its application space. The authors of [16] propose an ILP formulation and a heuristic algorithm for efficiently composing chains of virtual security

TABLE II
TAXONOMY OF SECURITY VNFs.

VNF	Description	Use in security	Implementations
Antispam	Email filtering	Malware detection, Phishing prevention	SpamAssassin, rspamd, ASSP, Juniper vSRX
Antivirus	Email/Web scanning/Endpoint security	Virus/Trojan/Malware detection	ClamAV, ClamWin, Juniper vSRX
DLP	Data Loss/Leakage Prevention	Data exfiltration detection	myDLP, OpenDLP
DPI	Payload analysis	Spam Filtering, Intrusion detection, DDoS detection, Malware detection, Security Analytics	OpenDPI, nDPI, L7-filter, Libprotoident, PACE, NBAR, Cisco ASAv
Honeypot	Traffic redirection and inspection	Spam filtering, Malware detection, SQL database protection, Security Analytics	HoneyD, SpamD, Kippo, Kojoney, Dionaea, Glastopf
IDS	Traffic inspection (header and payload)	Intrusion detection, Malware detection, DDoS detection, Security Analytics	Snort, Bro, Suricata, ACARM-ng, AIDE, OSSEC, Samhain, Cuckoo, Cisco ASAv
IPS	Traffic filtering based on header and payload	Intrusion prevention, DDoS prevention	Snort, Suricata, ACARM-ng, Fail2Ban, Juniper vSRX
NAT ¹	IP address mapping	Intrusion prevention	Netfilter, IPFilter, PF
Packet Filter Firewall	Header-based packet filtering	Intrusion prevention	Netfilter, nftables, NuFW, IPFilter, ipfw, PF, Juniper vSRX, VMWare vShield, Fortigate FW
Parental Control	Media content filtering	Blocking access to inappropriate content	OpenDNS, SquidGuard, DansGuardian, pfsense
VPN Gateway	Site-to-site VPN connection over unsecured networks	Data Tunneling/Encryption	OpenVPN, strongSwan, Juniper vSRX, Cisco ASAv, Fortigate VPN
WAF	HTTP traffic monitoring, filtering, logging	Prevention of SQL injection, cross-site scripting	ModSecurity

¹ NAT is not a security function but inherently provides packet filtering similar to a firewall.

functions. The ILP formulation includes a single security-related constraint to ensure that the security level of each deployed VSNF instance is higher than the security level required by the service request. However, this work does not take into account basic security aspects, such as order and operational mode (stateful/stateless) of the chained VSNFs. Moreover, the proposed formulation does not consider the mutual interference between security services caused by the concurrent access to the (finite) computing resources available in the infrastructure. The latter aspect is particularly relevant in a TSP scenario (see also part II of [17]), where the security services are provisioned in a dynamic manner based on the incoming customers' requests.

III. BACKGROUND

This work is underpinned by two emerging network technologies; Network Function Virtualization (NFV) and Software-Defined Networking (SDN) and their integration to provision network security solutions.

A. Network Function Virtualization

Today's network functions such as firewalling, DPI, IDSs, etc. are provided by specialized proprietary hardware appliances (also called *middleboxes*) strategically deployed in the network. The NFV paradigm separates the network functions from the underlying hardware by moving the functions from specialized devices to off-the-shelf commodity equipment such as industry standard servers or high-performance network devices. Therefore, network services can be decomposed into multiple VNFs running on physical or virtual machines, which could be located in data centers, network nodes or at the end-user premises.

In contrast to middleboxes, the configuration of which requires intensive and time-consuming manual intervention, NFV allows an automated and flexible deployment of network

functions on any NFV-enabled device. The lifecycle management of VNFs and hardware resources is achieved through a centralized software component called the Orchestrator.

B. Software-Defined Networking

SDN is often referred to as a paradigm for network environments where the control plane is physically separated from the data plane and a logically centralized control plane controls several devices. This differs from traditional networks in which nodes are autonomous systems unaware of the overall state of the network. In SDN deployments, the nodes are remotely controlled via standard protocols (e.g. OpenFlow [18]) by a logically centralized intelligent module called the *SDN controller*, which bases routing decisions on a global (domain) view of the network.

The controller is a software component which runs on commodity hardware appliances and provides an open Application Programming Interface (API) to program the network for configuration, monitoring and troubleshooting purposes. Such programmability enables automated and dynamic network configurability and fine-grained control of the traffic based on the values of the packets' header fields (e.g., source/destination IP/MAC addresses, VLAN tags, TCP/UDP ports, etc.).

C. Service Function Chaining

Service Function Chaining (also known as Network Service Chaining) is a technique for selecting and steering data traffic flows through network services. The network services can be traffic management applications such as load balancing, or security applications such as those detailed in Section III-D. Service function chaining combines the capabilities of SDN and NFV to connect a distributed set of VNFs.

D. A Taxonomy of security VNFs

As introduced in Section III-A, a VNF is a software implementation of a network function which is deployed on a

virtual resource such as a Virtual Machine. Table II provides a list of the most common security functions. Traditionally, the majority of these functions would have been implemented on dedicated hardware (middleboxes) to process the network traffic along the data path. Today, these functions are deployed as VNFs. Table II includes a short description of each VNF and some of the publicly available open-source implementations or commercial products.

IV. MOTIVATION

We motivate our work by describing two use case scenarios, namely web browsing and online gaming, where the TSP exploits the NSC and NFV technologies to provide security services tailored to specific users' application requirements.

Web browsing. Parental control is applied to Web traffic to block unwanted media and social-media content, while an IDS might be used to intercept malicious software (malware). Stateful VSNFs provide security functionality by tracking the state of network connections (e.g. Layer 4 firewall, NAT). In this case, the same VSNF instance must be traversed by all traffic flows of a network conversation in order to maintain the correct state of the connection. More flexible provisioning schemes can be adopted for stateless VSNFs, where multiple instances of the same VSNF might be deployed on different servers for load balancing. This example also illustrates the security best-practice that unwanted traffic should be blocked as soon as it enters the network by placing firewalls and IDS/IPS close to the border of the TSP domain. Another generally accepted practice, is to place firewalls before IDS/IPS (from the point of view of incoming traffic). Firewalls are generally designed to drop unauthorized traffic very quickly, thus reducing the burden on IDS/IPS, which are more computationally expensive.

Online gaming. An IDS might also be used to detect possible threats due to the misuse of chat tools integrated within the gaming software (e.g., phishing [19], social engineering [20], etc.). As the communication between the client and the server relies on timely delivery of packets, IDS operations are not executed on the *in-game* traffic. In this case the security is enforced by a faster VSNF such as a Firewall, which checks the packet headers without any deep-payload analysis. It should be noted that web traffic and chat conversations are often encrypted by TLS/SSL cryptographic protocols. Although encryption preserves the confidentiality of the traffic, it also prevents IDS-based VSNFs such as Parental Control and IDS from inspecting the packets, thus allowing an attacker to obfuscate malicious data in encrypted payloads. However, the TSP could overcome this limitation either using a Transparent Proxy VSNF or by exploiting recent advances in network security [21], [22].

These are just two examples of how the security service can be tailored to the user's application requirements by appropriate selection and placement of VSNFs. A list of common classes of applications supported in the TSP use-case is provided in Table I with their corresponding security

and QoS requirements and relevant VSNFs. One of them, the remotely accessible CCTV system, will be used in the rest of this paper as a running example to illustrate various aspects of our work.

According to the motivations provided above, we can summarize the rationale behind the PESS approach as follows: (i) a user's application should never under-perform because of VSNF operations and (ii) the VSNF placement must obey the TSP's security best-practices in terms of application security requirements, position in the network, operational mode (stateless or stateful VSNF), and order with respect to the direction of the traffic. In the next section, we present the PESS mathematical model for the placement of VSNF chains based on these criteria.

V. PESS VSNF PLACEMENT MODEL

The PESS model (Fig. 1) is a mathematical model to progressively embed service requests, formed by one or multiple VSNF chains, onto a physical network substrate by considering the available resources and realistic constraints.




Fig. 1. PESS placement model workflow.

PESS takes as input a model of the physical network including the current status of computing and network resources of servers and links, a security service request and the TSP's security policies (expressed in the form of constraints for PESS). The output of PESS is the mapping of the VSNFs onto the physical network (position of the VSNFs and one or more paths between them) and an updated model of the physical network taking into account the resources used to provision the service. The updated model is used as input for the next request.

Physical network model. We represent the physical network as a weighted graph $\mathcal{G} = (N, E)$, i.e. a graph where weights are assigned to nodes and edges.

Without loss of generality and to simplify the model, we assume that every node $i \in N$ is a NFVI-POP (Network Function Virtualization Infrastructure Point of Presence) [23] consisting of a set of servers and a local network composed of routers and switches. Each node i is characterized by the total computing resources of the servers $\gamma_i \in \mathbb{N}^+$ expressed in CPU cycles/sec. A link $(k, l) \in E$ is a wired connection between two nodes k and $l \in N$. It is characterized by its capacity $\beta_{k,l} \in \mathbb{N}^+$ and its propagation delay $\lambda_{k,l} \in \mathbb{N}^+$. Both are expressed as positive integer numbers representing bandwidth (bits/sec) and latency (sec).

Regions in a physical network are defined as subsets of nodes sharing some high-level features. Examples of regions are: (i) a set of nodes in the TSP network providing the same

cloud service (e.g. multimedia caching, data storage, etc.), or (ii) the set of egress nodes that connect the TSP network to the Internet (called *border region* in the rest of this paper).

Security service request. We model a security service request as a set of independent weighted directed graphs:

$$\mathcal{G}_s = \{(U^c, U_{pairs}^c) : c \in C_s\}$$

where C_s is the set of unidirectional chains composing the service request. Each graph includes nodes and arcs. Nodes $U^c = A^c \cup V^c$ comprise user and remote applications (A^c , the endpoints of chain c) as well as a subset of all VSNFs (V^c). Each arc in U_{pairs}^c delineates the order of traversing the VSNFs $\in V^c$ between endpoints in A^c .

Each chain $c \in C_s$ is characterized by its requirements in terms of minimum bandwidth β^c and maximum latency λ^c . Each endpoint in A^c is characterized by an identifier, which specifies where the endpoint must be placed in the physical network. The user application is characterized by the identifier of the physical node to which the user is attached (called *ep1* in the rest of the paper). A remote application is characterized by the identifier of a region in the physical network (called *EP2*). For instance, the border region if the endpoint represents a remote gaming server located outside the physical network. In this paper, *ep1* and *EP2* are referred to as *physical endpoints* of the service request \mathcal{G}_s .

A VSNF $u \in V^c$ is characterized by its requirements in terms of CPU units γ_u expressed in CPU cycles/bit. u is also characterized by the latency $\lambda_{i,u}^c$ it introduces in the dataplane to process a packet of chain c on node i . As formalized in Eq. (13), this latency is a function of the residual computing capacity of the node i where u is placed, the computing requirements γ_u of the VSNF, the average packet size σ^c of chain c and the traffic load of the chain (whose upper bound is β^c). Finally, a VSNF is characterized by its operational mode (either stateless or stateful) and by the identifier of a region in the physical network where it must be placed, if required by the TSP security policies.

Illustrative example. An example of a security service request for a CCTV system (see Table I) is represented in Fig. 2. The request in the example is composed of three chains ($c1$, $c2$, and $c3$), each one identified by the type of traffic and its direction.




Fig. 2. Example of security service request for the CCTV system.

Chain $c1$ is applied to the live video stream captured by the cameras and accessible over the Internet. The chain

comprises a L3 firewall to ensure that the stream is only transmitted to authorized endpoints. As specified in Table I, the most relevant requirement in this case is the bandwidth (β^{c1}) which depends on the frame rate, frame size and video codec of the CCTV system. In this case, a deep inspection of the video stream packets (e.g., with an IPS) would not provide any additional protection but would possibly reduce the frame rate of the video streaming, thus compromising the detection of anomalous events. On the other hand, the bi-directional control/management traffic is inspected by the IPS and the firewall included in chains $c2$ and $c3$. Such VSNFs protect the CCTV system from attacks such as Mirai [24] perpetrated through bots maliciously installed on Internet-connected devices, while the latency requirements λ^{c2} and λ^{c3} guarantee the responsiveness of the remote control of the CCTV cameras (pan, tilt, zoom, etc.).

A. ILP formulation

Definitions. Let us first define two binary variables:

- $x_{i,u}^c = 1$ iff node $u \in U^c$ is mapped to $i \in N$.
- $y_{k,l,i,j,u,v}^c = 1$ iff physical link $(k,l) \in E$ belongs to the path between nodes i and j to which $u, v \in U^c$ are mapped.

The residual capacity of a link, $\beta'_{k,l}$, is defined as the total amount of bandwidth available on link $(k,l) \in E$:

$$\beta'_{k,l} = \beta_{k,l} - \sum_{\substack{c \in C, i,j \in N \\ (u,v) \in U_{pairs}^c}} \beta^c \cdot y_{k,l,i,j,u,v}^c \quad (1)$$

thus, it is the nominal capacity of link (k,l) minus the bandwidth required by the chains $c \in C$ already mapped on that link.

Similarly, the residual capacity of a node is defined as its nominal CPU capacity minus the computing resources used by the VSNFs v instantiated on the node:

$$\gamma'_i = \gamma_i - \sum_{\{c \in C, u \in V^c\}} \gamma_u^c \cdot x_{i,u}^c \quad (2)$$

Problem formulation. Given a physical network \mathcal{G} , for each security service request \mathcal{G}_s , find a suitable mapping of all its unidirectional chains on the physical network, which minimizes the physical resources of \mathcal{G} expended to map \mathcal{G}_s , also known as the *embedding cost*.

Hence, the solution of the problem is represented by a set of $x_{i,u}^c$ and $y_{k,l,i,j,u,v}^c$ such that the cumulative usage of physical resources for all the chains in \mathcal{G}_s is minimized:

$$\min \sum_{\substack{c \in C_s, i,j \in N, \\ (k,l) \in E, (u,v) \in U_{pairs}^c}} b_{k,l} \cdot \beta^c \cdot y_{k,l,i,j,u,v}^c + \alpha \sum_{c \in C_s, i \in N, u \in V^c} c_i \cdot \gamma_u^c \cdot x_{i,u}^c \quad (3)$$

Here, α is a factor that can be used to tune the relative weight of the cost components (we have used $\alpha = 1$ for the experiments described in Section VII).

$b_{k,l}$ and c_i are the costs for allocating bandwidth and CPU:

$$b_{k,l} = \frac{1}{\beta'_{k,l} + \delta} \quad c_i = \frac{1}{\gamma'_i + \delta}$$

They penalize nodes and links with less residual capacity with the aim to increase the chances of accommodating more security service requests on the given physical network. $\delta \rightarrow 0$ is a small positive constant used to avoid dividing by zero in computing the value of the function.

B. Constraints

Routing Constraint (4) ensures that each node $u \in U^c$ is mapped to exactly one physical node $i \in N$. With Constraint (5), a physical link (k, l) can belong to a path between two nodes i and j for an arc $(u, v) \in U_{pairs}^c$ of chain $c \in C_s$ only if u and v are mapped to these nodes. Constraint (6) ensures that the path created for arc (u, v) starts at exactly one edge extending from node i to where VSNF (or start/endpoint) u is mapped. Similarly, (7) ensures the correctness and the uniqueness of the final edges in the path. Constraints (5-7) can be easily linearized with standard techniques such as the ones presented in [25]. Constraint (8) is the classical *flow conservation constraint*. That is, an outbound flow equals an inbound flow for each intermediate node l (intermediate nodes cannot consume the flow). Together with Constraint (8), Constraint (9) prevents multiple incoming/outgoing links carrying traffic for a specific flow in the intermediate node l , i.e., we only consider unsplittable flows.

$$\sum_{\{i \in N\}} x_{i,u}^c = 1 \quad \forall c \in C_s, \forall u \in U^c \quad (4)$$

$$y_{k,l,i,j,u,v}^c \leq x_{i,u}^c \cdot x_{j,v}^c \quad (5)$$

$$\forall c \in C_s, \forall i, j \in N, \forall (u, v) \in U_{pairs}^c, \forall (k, l) \in E$$

$$\sum_{\{(i,k) \in E, j \in N\}} y_{i,k,i,j,u,v}^c \cdot x_{i,u}^c \cdot x_{j,v}^c = 1 \quad (6)$$

$$\forall c \in C_s, \forall (u, v) \in U_{pairs}^c$$

$$\sum_{\{(k,j) \in E, i \in N\}} y_{i,k,i,j,u,v}^c \cdot x_{i,u}^c \cdot x_{j,v}^c = 1 \quad (7)$$

$$\forall c \in C_s, \forall (u, v) \in U_{pairs}^c$$

$$\sum_{\substack{k \in N \\ (k,l) \in E}} y_{k,l,i,j,u,v}^c = \sum_{\substack{m \in N \\ (l,m) \in E}} y_{i,m,i,j,u,v}^c \quad (8)$$

$$\forall c \in C_s, \forall i, j \in N, \forall l \in N, l \neq i, l \neq j, \forall (u, v) \in U_{pairs}^c$$

$$\sum_{\substack{k \in N \\ (k,l) \in E}} y_{k,l,i,j,u,v}^c \leq 1 \quad (9)$$

$$\forall c \in C_s, \forall i, j \in N, \forall l \in N, l \neq i, l \neq j, \forall (u, v) \in U_{pairs}^c$$

Resource Constraints (10-11) ensure that the resources consumed by a security service do not exceed the available bandwidth and computing capacities.

$$\sum_{\substack{c \in C_s, i, j \in N \\ (u,v) \in U_{pairs}^c}} y_{k,l,i,j,u,v}^c \cdot \beta^c \leq \beta'_{k,l} \quad \forall (k, l) \in E \quad (10)$$

$$\sum_{\{c \in C_s, u \in V^c\}} x_{i,u}^c \cdot \gamma_u^c \leq \gamma'_i \quad \forall i \in N \quad (11)$$

QoS Constraint (12) verifies that the requirements in terms of maximum end-to-end latency are met. It takes into consideration the propagation delay of physical links, the processing

TABLE III
GLOSSARY OF SYMBOLS.

Sets	
N	Set of physical nodes
E	Set of physical links
C	Set of all unidirectional chains already embedded in the network
C_s	Set of all unidirectional chains in the service request \mathcal{G}_s
U^c	Set of virtual nodes in the chain c
U_{pairs}^c	Set of unidirectional arcs in the chain c
A^c	Set of endpoints of the chain c . $A^c \subset U^c$
V^c	Set of VSNFs in the chain c . $V^c \subset U^c$
R_u	Region of N where VSNF u must be placed (<i>region</i> constraint)
M	Region of N where no VSNFs can be placed (<i>veto</i> constraint)
$ep1, EP2$	Physical endpoints of a service request. $ep1 \in N, EP2 \subset N$
Parameters	
γ_i	Nominal computing resources of node i (CPU cycles/sec)
γ'_i	Residual computing resources of node i (CPU cycles/sec)
γ_u	CPU cycles required by u to process one bit of a network packet (CPU cycles/bit)
γ_u^c	Computing resources required by node u of chain c (CPU cycles/sec). $\gamma_u^c = \gamma_u \cdot \beta^c$
$\beta_{k,l}$	Nominal capacity of link (k, l) (bits/sec)
$\beta'_{k,l}$	Residual capacity of link (k, l) (bits/sec)
β^c	Minimum bandwidth required by chain c (bits/sec)
$\lambda_{k,l}$	Propagation delay: the time spent by a packet to traverse the link (k, l) (secs)
$\lambda_{k,l,i,j}^c$	Queuing delay: time spent by a packet of chain c to traverse the network devices (routers and switches) in the local networks of adjacent nodes k and l (secs). $\lambda_{k,l,i,j}^c > 0$ iff $k = i$ or $l = j$.
$\lambda_{i,u}^c$	Processing delay: the time spent by a packet to traverse VSNF u of chain c placed on node i (secs)
λ^c	Maximum latency tolerated by chain c (secs)
π^c	Estimated latency between the TSP network and the remote endpoint of chain c (secs). $\pi^c = 0$ if the endpoint belongs to the TSP network.
σ^c	Average packet size of chain c (bits).
$b_{k,l}$	Cost for allocating a unit of bandwidth on link (k, l)
c_i	Cost for allocating a unit of CPU on node i
Decision variables	
$x_{i,u}^c$	Binary variable such that $x_{i,u}^c = 1$ iff node $u \in U^c$ is mapped to $i \in N$
$y_{k,l,i,j,u,v}^c$	Binary variable such that $y_{k,l,i,j,u,v}^c = 1$ iff physical link $(k, l) \in E$ belongs to the path between nodes i and j to which $u, v \in U^c$ are mapped

delay of VSNFs and the queuing delay through network devices. Note that the minimum bandwidth requirement is verified against the bandwidth resource Constraint (10).

$$\pi^c + \sum_{i \in N, u \in V^c} x_{i,u}^c \cdot \lambda_{i,u}^c + \sum_{\substack{i, j \in N, (k, l) \in E \\ (u, v) \in U_{pairs}^c}} y_{k,l,i,j,u,v}^c \cdot (\lambda_{k,l} + \lambda_{k,l,i,j}^c) \leq \lambda^c \quad \forall c \in C_s \quad (12)$$

π^c is an estimation of the propagation delay between the TSP network and the remote endpoint of chain c , in case the endpoint is outside the TSP network. We assume that this value is independent from the TSP's network egress node. Clearly π^c is 0 for those chains whose remote endpoint is part of the TSP network (e.g., a cloud data center managed by the TSP).

The processing delay $\lambda_{i,u}^c$ is the time spent by a packet to traverse VSNF u on physical node i . It contributes to the overall end-to-end delay of chain c only if VSNF u is placed on node i (i.e., $x_{i,u}^c = 1$). $\lambda_{i,u}^c$ includes the time taken by the VSNF to process the packet and the overhead of the virtualization technology (VMware, KVM, QEMU virtual machines, Docker containers, etc.). For simplicity, we do not model the delays due to the CPU scheduler operations implemented on the physical node [26]. Based on the observations in [27], [28], [29], $\lambda_{i,u}^c$ is modeled as a convex function of the traffic load of

the chain, and its value is computed by considering the impact of other VSNFs co-located on the same physical node.

$$\lambda_{i,u}^c = \frac{\gamma_u \cdot \sigma^c}{(\gamma'_i - \gamma_u \cdot \beta^c) + \delta} = \frac{\gamma_u \cdot \sigma^c}{(\gamma'_i - \gamma_u^c) + \delta} \quad (13)$$

In Eq. (13), $\gamma_u \cdot \sigma^c$ is the average amount of CPU cycles used by VSNF u to process a packet of chain c (virtualization overhead included). The latency overhead caused by co-located VSNFs depends on the amount of computing resources of the node they use or, equivalently, on the residual computing resources of the node γ'_i . $\gamma_u \cdot \beta^c = \gamma_u^c$ is the amount of CPU cycles/sec used by VSNF u on node i , which depends on the traffic load of the chain. δ is a small positive constant used to avoid dividing by zero in the case that u consumes all the residual computing resources of node i .

The sum $\lambda_{k,l} + \lambda_{k,l,i,j}^c$ in Eq. (12) is the total time spent by a packet travelling between two adjacent nodes k and l . It includes the propagation delay $\lambda_{k,l}$, proportional to the distance between k and l , and the queuing delay $\lambda_{k,l,i,j}^c$, proportional to the number of network devices (switches and routers) the packet traverses within the local networks of k and l . The queuing delay is influenced by the buffer size of network devices' ports and by the traffic load [30]. For the sake of simplicity, we assume that the buffers are correctly dimensioned, i.e. no dropped packets due to buffer overflow. In addition, we estimate the queuing delay $\lambda_{k,l,i,j}^c$ as a traffic-load independent value; a function of the maximum queue capacity of the ports and of the VSNFs placement (hence a function of indices k, l, i and j). Specifically, $\lambda_{k,l,i,j}^c > 0$ if at least one VSNF is mapped either on k ($k = i$), or on l ($l = j$), meaning that a packet of chain c must traverse the local network of either k , or l (or both) to reach the VSNFs running on the nodes' servers. Otherwise, the local networks of k and l are by-passed by the traffic of c , resulting in $\lambda_{k,l,i,j}^c = 0$.

Constraint (14) ensures that the current security service \mathcal{G}_s does not compromise the end-to-end latency of chains $\hat{c} \in C$ in operational security services (also called *operational chains* in the rest of the paper).

$$\pi^{\hat{c}} + \sum_{i \in N, \hat{u} \in V^{\hat{c}}} \bar{x}_{i,\hat{u}}^{\hat{c}} \cdot \lambda_{i,\hat{u}}^{\hat{c}} + \sum_{\substack{i,j \in N, (k,l) \in E \\ (\hat{u}, \hat{v}) \in U_{pairs}^{\hat{c}}}} \bar{y}_{k,l,i,j,\hat{u},\hat{v}}^{\hat{c}} \cdot (\lambda_{k,l} + \lambda_{k,l,i,j}^{\hat{c}}) \leq \lambda^{\hat{c}} \quad \forall \hat{c} \in C \quad (14)$$

In Eq. (14), $\bar{x}^{\hat{c}}$ and $\bar{y}^{\hat{c}}$ are the values of decision variables x and y computed for the placement of chain \hat{c} . $\lambda_{i,\hat{u}}^{\hat{c}}$ is the updated value of the processing delay introduced to the traffic of chain \hat{c} by VSNF \hat{u} when running on node i .

$$\lambda_{i,\hat{u}}^{\hat{c}} = \frac{\gamma_{\hat{u}} \cdot \sigma^{\hat{c}}}{(\gamma'_i - \sum_{\{c \in C_s, u \in V^c\}} x_{i,u}^c \cdot \gamma_u^c) + \delta} \quad (15)$$

In Eq. (15), the value of $\lambda_{i,\hat{u}}^{\hat{c}}$ is updated by considering the computing resources consumed on node i by VSNFs of the security service request \mathcal{G}_s . Approximation of Eq. (15) can be achieved by using piecewise linearization techniques and Special-Ordered Set (SOS) variables and constraints available in most commercial solvers (e.g., [31]).

Security constraints ensure that the TSP's security policies

are applied. Specifically, Constraint (16) forces a subset C'_s of the chains in the request to share the same VSNF instance in case of stateful flow processing.

$$x_{u,i}^{c_1} = x_{u,i}^{c_2} \quad \forall c_1, c_2 \in C'_s \subset C_s, i \in N, u \in V^c \quad (16)$$

Constraint (17) forces the algorithm to place the VSNF $u \in V^c$ in a specific region of the network defined as a subset of nodes $R_u \subset N$.

$$\sum_{\{i \in R_u\}} x_{i,u}^c = 1 \quad R_u \subset N, R_u \neq \emptyset, u \in V^c \quad (17)$$

We use Constraint (17) to enforce the security close to the user by placing VSNFs on *ep1* ($R_u = \{ep1\}$), or to protect a portion of the TSP's network, such as the border region or a distributed data center ($R_u = EP2$) from potentially malicious user traffic. Furthermore, Constraint (17) can be used to place a VSNF on a physical node with special hardware characteristics (e.g., hardware acceleration for encryption). Similarly, the *veto* Constraint (18) can be used to prevent the placement of any VSNFs on a pre-defined subset of nodes $M \subset N$. A TSP may choose to do this to protect specific nodes (called *veto nodes*) that host sensitive data or critical functions from user traffic.

$$\sum_{\{i \in M, u \in V^c\}} x_{i,u}^c = 0 \quad \forall c \in C_s, M \subset N, M \neq \emptyset \quad (18)$$

Finally, for each chain $c \in C_s$, the correct order of VSNFs in V^c is ensured by Constraints (4-9), plus Constraint (17) applied to user and remote applications $u \in A^c$ with $R_u = \{ep1\}$ and $R_u = EP2$ respectively. Note that, the order can be specified per application (chain), as different applications may require the same VSNFs but in different order.

These four security constraints enable fulfillment of the security policies/practices defined by the TSP e.g. the order in which the VSNFs are executed, the position of the VSNFs in the network, and the operational mode of VSNFs (either stateful or stateless).

VI. THE PESS HEURISTIC ALGORITHM

The embedding problem presented in Section V has been solved using a commercial solver. However, given the complexity of the ILP model, the solver is unable to produce solutions in an acceptable time frame, as required for dynamic scenarios such as the ones under study. For this reason, we have also implemented a heuristic algorithm to find near optimal solutions in much shorter time.

The logic behind the PESS heuristic is based on assuring that Constraints (4-18) are applied in an efficient manner. In particular, the security constraint (16) ensures that a stateful VSNF specified in two or more chains in the same service request \mathcal{G}_s is placed on the same node. However, as different chains might share more than one stateful VSNF (possibly in a different order), the correct placement of a multi-chain security service request may become a computationally expensive operation. For this reason, given a path between *ep1* and one of the nodes $ep2 \in EP2$, the heuristic places all the VSNFs specified in \mathcal{G}_s on a maximum of three nodes of the path with the following strategy: (i) place each region-specific VSNF

$u \in V^c$ ($R_u \neq \emptyset$) either on $ep1$ or on $ep2 \in EP2$ depending on R_u (i.e., either $R_u = \{ep1\}$ or $R_u = EP2$), (ii) place all the other VSNFs in \mathcal{G}_s on the node with the highest residual capacity in the path to minimize the embedding cost (Eq. 3).

The solution is obtained by selecting the candidate path between $ep1$ and $EP2$ where the embedding of all the chains in \mathcal{G}_s fulfills the constraints described in Section V at the lowest cost, as computed with the objective function (Eq. 3).

Algorithm 1 the PESS algorithm.

Input: Physical network substrate (\mathcal{G}), security service request (\mathcal{G}_s), set of active chains in the network (C)

Output: The mapping of the security service onto the physical substrate (*solution*). *None* if no feasible mappings are found.

```

1: procedure PESS( $\mathcal{G}, \mathcal{G}_s, C$ )
2:    $\beta \leftarrow \sum_{c \in C_s} \beta^c$  ▷ total required bandwidth
3:    $\gamma \leftarrow \sum_{c \in C_s, u \in V^c} \gamma_u^c$  ▷ total required CPU
4:    $P = \{p_{[ep1, ep]} \mid ep \in EP2\} \leftarrow \text{DIJKSTRA}(ep1, EP2, \beta)$ 
5:   if  $P = \emptyset$  then
6:     return None
7:   end if
8:    $S = \{s_{[ep1, ep]} \mid ep \in EP2\} \leftarrow \text{EMBED}(P, \beta, \gamma)$ 
9:    $N_S \leftarrow \{i \in S \mid s \in S\}$  ▷ physical nodes in the initial solutions
10:   $E \leftarrow \{i \in N \mid i \notin N_S \cup M, \gamma_i' > \gamma_j' \forall j \in N_S\}$ 
11:   $\bar{s}_{[ep1, ep2]} \leftarrow \underset{s \in S}{\text{argmin}} \text{cost}(s)$  ▷ best initial solution
12:   $P_1 \leftarrow \text{DIJKSTRA}(ep1, E, \bar{\beta})$ 
13:   $P_2 \leftarrow \text{DIJKSTRA}(ep2, E, \bar{\beta})$ 
14:   $S \leftarrow S \cup \text{EMBED}(P_1 \cup P_2, \bar{\beta}, \bar{\gamma})$  ▷ expanded solution set
15:  solution  $\leftarrow$  None
16:   $S \leftarrow \text{SORTEDDECREASINGCOST}(S)$ 
17:  for all  $cs \in S$  do
18:    if  $\text{LATENCYOPCHAINS}(\mathcal{G}, C, cs)$  is True then
19:      solution  $\leftarrow cs$ 
20:    break
21:  end if
22: end for
23: if solution is None then
24:   return None
25: end if
26:  $\text{UPDATERESOURCES}(\text{solution}, \mathcal{G})$ 
27:  $\text{STORESOLUTION}(\mathcal{G}, C, \text{solution})$ 
28: return solution
29: end procedure

```

Initial solution. The embedding process starts at line 4 in Algorithm 1 with a greedy approach based on the Dijkstra's algorithm. At this stage, we compute the shortest path tree between the two endpoints $ep1$ and $EP2$ using the residual bandwidth as link weight computed as $b_{k,l} \cdot \beta^c$ in Eq. (3). The Dijkstra algorithm stops when all the nodes $ep2 \in EP2$ are marked as *visited*, i.e. before building the whole tree of paths. For each path between $ep1$ and $EP2$, the algorithm places the VSNFs in the chains according to the aforementioned strategy, the order of the VSNFs as specified in the service request, the latency Constraint (12), and the security Constraints (16-18) (line 8). The output of this first step is a set of candidate solutions S with different embedding costs. S is passed as input to the next step.

Expanded solution set. The algorithm now evaluates whether high-capacity nodes not included in the initial solution set S can be used to build new solutions with lower embedding cost. Hence, given the initial set of solutions S , the algorithm identifies the physical nodes in the network with these two

properties (set E defined at line 10): (i) not included in the initial set of solutions S nor *veto* nodes, and (ii) higher computing capacity with respect to the nodes included in S . The algorithm then computes the shortest path tree twice, once from $ep1$ towards E and once from $ep2 \in EP2$ towards E (lines 12-13), where $\{ep1, ep2\}$ are the physical endpoints of the solution in S with the lowest embedding cost (line 11).

The resulting subpaths are joined to form a new set of paths between $ep1$ and $ep2$. Afterwards, the algorithm performs the placement of the VSNFs on each of the new paths with the strategy described earlier in this section. The feasible solutions are added to the initial set S (line 14).

The set of candidate solutions is sorted in descending value of embedding cost (line 16). The first one that satisfies Constraint (14) is the accepted solution (lines 18-19). Finally, the algorithm updates the values of γ_i' and β_i' by removing the resources consumed with the accepted solution and stores the mapping of its chains in the set C that records all the active chains in the network (lines 26-27).

Latency of operational chains. Given a candidate solution $cs \in S$, function LATENCYOPCHAINS is invoked to verify whether embedding cs compromises the end-to-end latency of operational chains (line 18 in Algorithm 1). Instead of verifying the inequality in Eq. (14) for each operational chain, LATENCYOPCHAINS implements a heuristic approach, which reduces the time complexity of this operation from $O(n)$, with n the number of operational chains, to $O(1)$.

Each time a chain $c \in C$ becomes operational, the algorithm computes $\langle \gamma \rangle^c$, a threshold value obtained from Eq. (12) and (13) as follows:

$$\langle \gamma \rangle^c = \frac{\sum_{i \in N, u \in V^c} \bar{x}_{i,u}^c \cdot \gamma_u \cdot \sigma^c}{\lambda^c - \pi^c - \sum_{\substack{i,j \in N, (k,l) \in E \\ (u,v) \in U_{pairs}}} \bar{y}_{k,l,i,j,u,v}^c \cdot (\lambda_{k,l} + \lambda_{k,l,i,j}^c)} - \delta \quad (19)$$

In Eq. (19), \bar{x} and \bar{y} are the values of decision variables x and y used to embed c . $\langle \gamma \rangle^c$ estimates the minimum average residual computing capacity necessary to satisfy the inequality in Eq. (12). Therefore, the algorithm records and monitors those operational chains with the highest values of $\langle \gamma \rangle^c$ to establish whether a candidate solution is feasible or not, as inequality in Eq. (12) is violated earlier for such chains than for the others.

The algorithm stores one operational chain per physical node in a data structure, i.e. the chain with the highest value of $\langle \gamma \rangle^c$ with at least one VSNF mapped on that node. Hence, given the physical nodes mapped in the candidate solution cs , the algorithm computes Eq. (14) only for the operational chains in the data structure linked to such nodes by using the values of variables x and y of solution cs . If the inequality is not satisfied for one of those chains, cs is rejected.

As the maximum number of physical nodes used to provision a security service is three ($ep1$ and $EP2$ to fulfill the region constraint and the node with the highest residual capacity in the path), the worst-case time complexity of this process is $O(1)$, thus constant in the number of operational chains and with respect to the size of the network. There-

fore, the overall time complexity of the PESS heuristic is $O(|E| + |N| \log(|N|))$, i.e., the worst-case time complexity of the Dijkstra's algorithm.

VII. EVALUATION

We first assess the PESS heuristic by comparing its solutions against the optimal embeddings as computed by a commercial solver (Gurobi [32]). We then prove the benefits of the proposed application-aware approach against the baseline (the application-agnostic approach adopted, for instance, in [15]), in which security services are provided without taking into account the specific requirements of applications. We finally analyze the scalability of PESS by measuring the average embedding time on different network sizes.

A. Test configuration

The PESS heuristic has been implemented as a single-threaded Python program, while the ILP model formalized in Section V has been implemented with the Gurobi Python API version 7.5 [33]. All experiments are performed on a server-class computer equipped with 2 Intel Xeon Silver 4110 CPUs (16 cores each running at 2.1 GHz) and 64GB of RAM.

1) *Topology*: The simulations are performed on synthetic topologies randomly generated based on the Barabási-Albert model [34]. We generate topologies of different sizes and densities to evaluate the performance of the PESS heuristic in a variety of generic network scenarios.

We also validate PESS with two realistic network models. One is the Stanford University backbone [35], a medium-scale campus network consisting of 46 links, 14 operational zone routers, 10 Ethernet switches, and 2 border routers connecting the University to the Internet. We assume one NFVI-POP for each network device and 10Gbps links. The second model is the Italian education and research network (consortium GARR [36]). The GARR network covers the entire Italian national territory ([37]), comprising 83 links and 46 nodes. Along with the actual view of the physical topology, [37] provides the specification of the egress nodes, i.e. the nodes that connect the GARR network to the Internet and that compose the *border* region in our evaluation (nodes *FII*, *MI2*, *PD2*, *RM2* and *TO1*, as indicated in [37]). As we have no information related to data center distribution in the GARR network, we have assumed one NFVI-POP for each node. In addition, we set the nominal capacity of the links $\beta_{k,l}$ using the values specified in [38].

Given the relatively small size of the Stanford network, we assume no propagation delay between its nodes, i.e. $\lambda_{k,l} = 0 \forall k, l$. For the other networks, random and GARR, we compute the propagation delay of each link with the following formula:

$$\lambda_{k,l} = \frac{d_{k,l} \cdot r_index}{C}$$

where $r_index = 1.5$ is an approximation of the refractive index of optical fibers, $C \simeq 3 \cdot 10^8$ m/s is the speed of light in the vacuum and $d_{k,l}$ is the distance between two nodes k and l . In the case of random networks, $d_{k,l}$ is a random positive value ranging from 10 to 100 Km, while for the GARR

TABLE IV
CPU REQUIREMENTS FOR SOME VSNF IMPLEMENTATIONS.

VSNF	Virtualization	γ_u (cycles/bit) ¹
Snort IDS/IPS	VirtualBox	9.5 [39]
Suricata IDS/IPS	VirtualBox	8.2 [39]
OpenVPN with AES-NI	KVM/QEMU	31 [40]
strongSwan with AES-NI	KVM/QEMU	16 [40]
Fortigate-VM NGFW	FortiOS	9 [41]
Fortigate-VM SSL VPN	FortiOS	13.6 [41]
Fortigate-VM IPSec VPN	FortiOS	14.5 [41]
Fortigate-VM Threat protection	FortiOS	11.3 [41]
Cisco ASAv Stateful IDS	VMware ESX/ESXi	4.2 [42]
Cisco ASAv AES VPN	VMware ESX/ESXi	6.9 [42]
Juniper vSRX FW	VMware VMXNET3	2.3 [43]
Juniper vSRX IPS	VMware VMXNET3	2.4 [43]
Juniper vSRX AppMonitor	VMware VMXNET3	1.5 [43]

¹ $\gamma_u = \text{CPU clock} \cdot (\text{CPU usage}) / \text{Throughput}$. CPU usage is set to 1 (i.e. 100%) when the value is not specified.

network $d_{k,l}$ is computed by approximating the coordinates of the nodes based on the information available on the web site.

As introduced in Section V-B, we estimate the worst-case queuing delay $\lambda_{k,l,i,j}^c$ as a traffic-load independent value using the queue capacity of switch ports reported in [30] (80 μ s for 10Gbps ports with a 100 KB buffer). Specifically, we assume a three-tier local network at each node of GARR and random topologies, resulting in a maximum of $12 \times 80 \mu$ s queuing delay introduced at each node. This reflects the maximum queuing delay experienced by each packet crossing a node to be processed by one or more VSNFs mapped on the node, which involves traversing three network devices (hence, six 10Gbps ports) to reach the servers where the VSNFs are running, and traversing three network devices before leaving the node (six further 10Gbps ports).

For the campus scenario, implemented using the Stanford University topology, we instead assume only one network device per node; the device specified in the network topology. Hence, the maximum queuing latency for a packet crossing a Stanford node is $4 \times 80 \mu$ s.

For each node of the three evaluation scenarios we assume one server with computing capacity of 32x2.1 GHz (a 32-core CPU running at 2.1 GHz).

2) *Security service requests*: As introduced in Section V, a security service request is configured by the TSP to provision security for user applications (see the CCTV example in Section V). For evaluation purposes, we automatically generate requests composed of a random number of chains, ranging between 1 and 5. Each chain comprises a random subset of VSNFs from the list presented in Table IV, with a maximum of 3 VSNFs per chain (i.e., up to 15 VSNFs per user application). Based on the use case scenarios illustrated in Sections IV and V (web browsing, online gaming, CCTV system), we believe these are reasonable values.

The CPU requirements for the VSNFs are presented in Table IV. It should be noted that the values of γ_u (cycle/bit) reported in Table IV are estimated based on the results of experiments reported in scientific papers or product datasheets and obtained under optimal conditions, with only one VSNF running at a time. The impact on the network traffic caused by concurrent

TABLE V
COMPARISON BETWEEN THE PESS HEURISTIC AND THE PESS
ILP-BASED ALGORITHM ON THREE NETWORK SCENARIOS.

Network model	Heuristic embedding cost overhead ¹	Average Time (sec)	
		Heuristic	Solver
Random	0.06%	0.002	150
Stanford	0.07%	0.003	700
GARR	0.5%	0.003	1500

¹ Average overhead with respect to the solver embedding cost.

VSNFs running on the same node are estimated with Eq. (13) and (15). These values of γ_u have been used to perform the evaluation tests described in the remainder of this section, with the aim of enabling interested readers to replicate the experiments in similar conditions. However, we also obtained comparable results using random values.

B. Comparison between solver and heuristic

Methodology. In this experiment, we compare the PESS ILP-based algorithm implemented with the solver and the PESS heuristic on the Stanford and GARR network models, and on Barabási-Albert random topologies with 20 nodes and 36 links.

The security service requests are generated using a Poisson process with exponential distribution of inter-arrival and holding times. Once a service expires, the resources allocated to it are released.

We start by simulating the processing of 10^5 service requests using the PESS heuristic. Once a stable network utilization (load) is reached, we save the subsequent service requests along with the network state and the heuristic solution. In a second stage, we run the solver to compute the optimal solution for each of the requests saved in the previous stage and we compare the results with the recorded heuristic solutions. This process is repeated with values of network load ranging between 1000 and 20000 Erlang.

Metrics. (i) Heuristic embedding cost overhead over optimal solutions and (ii) embedding time.

Discussion. As explained in Section VI, the PESS heuristic places all the chains of a service request on a single path to efficiently guarantee that the QoS Constraint (12) and the region Constraint (17) are respected. Once the path is found, the heuristic places the VSNFs of all the chains on a maximum of three nodes in the chosen path: the one with the highest residual computing capacity and the ones specified with the region constraint (if any). Such implementation choices reduce the solution space in case of requests with multiple chains and VSNFs. On the other hand, Constraints (12) and (17) also narrow down the solution space for the solver, often resulting in single-path optimal solutions. As a result, we measure a marginal embedding cost overhead of the heuristic solutions with respect to the optimal solutions on all three evaluation scenarios (see Table V).

It is worth analyzing the reason behind nearly one order of magnitude difference between the GARR topology and the other two network scenarios. When the initial solution is

computed, the heuristic algorithm selects the endpoint $ep2 \in EP2$ to further explore the solution space, thus excluding the other endpoints in $EP2$ (line 11 in Algorithm 1). This strategy improves the scalability of the heuristic in case of large endpoint sets $EP2$, at the cost of slightly reducing the quality of the solutions.

In this regard, on the GARR network the border region is used as endpoint $EP2$ for 80% of the requests, to simulate a real-world TSP network where most of the traffic is directed towards the Internet. Hence, good solutions involving four of the five nodes in the border are not considered during the second stage of the heuristic, possibly leading to less accurate solutions. Conversely, a border region of only two nodes is defined in the Stanford topology (the two border routers), while no special regions at all are configured for the random networks (thus, always $|EP2| = 1$), resulting in more precise embeddings.

As reported in Table V, the embedding time measured for the heuristic is 3 ms, on average, with the Stanford and GARR topologies, and below 3 ms, on average, with the random topologies. In contrast, the solver takes between 150 and 1500 s, on average, to find the optimal solutions on the three network scenarios. Please note that, the results related to the GARR network are limited to service requests with less than 10 VSNFs. Due to the size of the GARR topology (46 nodes and 83 links), above this threshold the solver runs out of memory and it is terminated by the operating system.


C. PESS vs application-agnostic provisioning

Methodology. We start two experiments in parallel using two identical copies of the same physical network graph. At each iteration, we generate a service request with application-specific QoS and security requirements. In *Experiment 1*, the security service is provisioned on one copy of the network with the PESS heuristic. In *Experiment 2*, the service is provisioned on the second copy of the network by simulating the standard approach (adopted, for instance, in [15] and used in this test as baseline), where two application-agnostic chains of VSNFs (one for each direction of the traffic) are applied to the user traffic to fulfill all the security requirements regardless of the specific needs of the applications. At the end of each iteration, the two copies of the network are updated according to the resources consumed by the respective provisioning approach.


As in the previous experiment, security service requests are generated using a Poisson process with exponential distribution of inter-arrival and holding times. We run 10^5 iterations, starting to collect statistics after the first $8 \cdot 10^4$ requests (once a stable network load is reached). The two parallel experiments are repeated with different network load values.

Metrics. Blocking probability, consumption of computing resources, end-to-end latency of the chains and number of active services in the network.


Discussion. Fig. 3 compares the performance of the PESS application-aware service provisioning algorithm (*PESS* in the figure) and the baseline approach (*Base*) on random networks.




(a) Consumed CPU resources.



(b) Blocking probability.




(c) Average load.




(d) Processing delay ratio (Base/PESS).

Fig. 3. Comparison between the baseline (Base) and the PESS approaches on random networks (20 nodes and 36 links).




(a) Consumed CPU resources.




(b) Blocking probability.

Fig. 4. Comparison between the baseline (Base) and the PESS approaches on the GARR network.



(a) Consumed CPU resources.



(b) Blocking probability.

Fig. 5. Comparison between the baseline (Base) and the PESS approaches on the Stanford backbone network.

The experimental results are plotted as functions of the network load, which is expressed in terms of the average number of security service requests in the network (Erlang).

The efficient usage of the computing resources reported in Fig. 3(a) is a major benefit of the application-aware provisioning mechanism proposed in this work. In particular, PESS avoids inefficiencies, such as a high bandwidth video stream being processed by a high resource demanding IPS (see the CCTV example in Section V), ultimately leading to a lower

blocking probability and to a higher number of active services in the network, as shown in Fig. 3(b) and 3(c) respectively.

The benefits of PESS in terms of reduced end-to-end latency are reported in Fig. 3(d). The plot illustrates the ratio between the average end-to-end latency of the chains in *Experiment 2* (Baseline), and the average end-to-end latency of the chains in *Experiment 1* (PESS). At low loads, when the nodes in the networks of both experiments are only partially busy, the value of this ratio is between 1.1 and 1.4. In other words, under





Fig. 6. Heuristic execution time as a function of the number of physical nodes. The plot on the left reports the results with $m = 1, 3, 5$ in the Barabási-Albert model when generating random graphs, while the one on the right side shows the measurements with fixed $m = 5$ at different sizes of endpoint $EP2$.

typical operational conditions, the average end-to-end latency of chains provisioned with our approach is 10-40% lower than the baseline. Moreover, when the nodes in *Experiment 2* are heavily loaded, the processing delay introduced by busy nodes becomes very high, as modeled with Eq. (12). This phenomenon produces high ratios, represented by the spike in the plot, which gradually decrease at high loads when the nodes in the network of *Experiment 1* also become fully loaded.

Fig. 4 reports the results of the simulations performed with the GARR network. In this case, we are particularly interested in observing the behavior of our approach in the presence of a critical region (from the security viewpoint) such as the border of the network. In order to analyse this, we empirically configure the random generator of service requests to generate 80% of requests directed towards the Internet (i.e., crossing the border of the network). In Fig. 4(b), it can be noted that both PESS and the baseline have similar blocking probability at low loads (below 6000). This is a consequence of the bandwidth usage on links towards the border region, which is almost always identical for *Experiment 1* and *Experiment 2*. The two curves start diverging at load 6000, i.e. when the border region runs out of computing resources with the baseline approach (as shown with dashed curves in Fig. 4(a)). The probability curves in Fig. 4(b) begin to re-converge at load 12000, when the border region with PESS also becomes full. Solid curves in Fig. 4(a) indicate that, between loads 1000 and 6000, when the blocking probability of the two experiments is comparable, PESS requires around 50% less computing resources than the baseline to provision the security services.

The results obtained with the Stanford network model are presented in Fig. 5. In contrast to the GARR network, where busy links in sparsely connected areas cause rejected requests at low loads, in these experiments we see a non-zero blocking probability only when the border region of the Stanford network runs out of computing resources, i.e. at loads 4000 and 6000 for the baseline and PESS, respectively.

Similar to the random networks scenario, we can observe a higher number of active services and a lower end-to-end latency with PESS in both GARR and Stanford networks. The plots are omitted due to space constraints.

D. Scalability evaluation

Methodology. We evaluate the scalability of the PESS heuristic on Barabási-Albert random topologies of between 10 and 1000 nodes. For each of these topologies, we simulate the processing of 1000 service requests and report the average execution time.

Metrics. Average execution time.

Discussion. In the first experiment (reported in the leftmost plot of Fig. 6) we used $|EP2| = 1$ for all the service requests and we varied the attachment parameter m , which determines the number of edges to attach from a new node to existing nodes when generating the random network. This influences the execution time of the shortest path algorithm. For instance, $m = 1$ produces tree-like topologies with $|E| = |N| - 1$. The general rule for computing the number of edges in Barabási-Albert networks is $|E| = m \cdot |N| - m^2$. As illustrated in Fig. 6, even for very large networks with 1000 nodes and 4975 edges ($m = 5$ in the figure), on average, the PESS heuristic can provision a security service in around 200 ms.

In the second experiment, we used a fixed value of $m = 5$ (the worst case in the first experiment) and we varied the size of endpoint $EP2$, as the number of nodes in $EP2$ determines how long PESS takes to compute the initial solution. In the rightmost plot in Fig. 6, the black solid curve is the reference measurement from the first experiment. As shown by the dashed curves in the plot, the average execution time increases linearly with the size of endpoint $|EP2|$, up to around 250 ms in the worst case with $|N| = 1000$, $|E| = 4975$ and $|EP2| = 500$.

As introduced in Section I and formulated in Section V, the VSNFs placement model and heuristic proposed in this work target NFV-enabled systems where security services are dynamically provisioned and updated based on users' applications and their security and QoS requirements. Such systems require efficient provisioning strategies to minimize the exposure of such applications to cyber attacks. With respect to these objectives, the experimental results from the PESS scalability evaluation are encouraging and clearly indicate the potential for practical implementation of the proposed application-aware approach in real-world scenarios.

VIII. CONCLUSIONS

In this paper, we have tackled the problem of the progressive provisioning of security services by means of VSNFs. The proposed approach, called PESS, takes into account security

and QoS requirements of user applications, while ensuring that computing and network resources are accurately utilised. We have discussed the rationale behind our design decisions and presented an ILP formulation and a heuristic algorithm that solve the placement problem. Although we have focused our work on security services, the PESS approach is applicable to more complex scenarios, where heterogeneous network services provided by means of generic VNFs coexist.

The evaluation results demonstrate the benefits of PESS for both users and operators, with savings in resource utilization and in end-to-end latency. We have also shown that the heuristic implementation of the proposed application-aware approach produces near-optimal solutions and scales well in large and dense networks, indicating the potential of PESS in real-world scenarios.

REFERENCES

- [1] R. Mijumbi et al., "Network Function Virtualization: State-of-the-Art and Research Challenges," *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 236–262, 2016.
- [2] R. Hill et al., "Measuring Latency for Video Surveillance Systems," in *Proc. of DICTA*, 2009.
- [3] M. Claypool et al., "Latency and Player Actions in Online Games," *Comm. of the ACM*, vol. 49, no. 11, Nov. 2006.
- [4] Y. Chen et al., "QoS Requirements of Network Applications on the Internet," *Inf. Knowl. Syst. Manag.*, vol. 4, no. 1, pp. 55–76, Jan. 2004.
- [5] M. Claypool et al., "Latency Can Kill: Precision and Deadline in Online Games," in *Proc. of ACM MMSys*, 2010.
- [6] R. Doriguzzi-Corin et al., "Application-Centric Provisioning of Virtual Security Network Functions," in *Proc. of the Third IEEE International Workshop on Security in NFV-SDN (SN-2017)*, 2017.
- [7] F. Bari et al., "Orchestrating Virtualized Network Functions," *IEEE TNSM*, vol. 13, no. 4, pp. 725–739, 2016.
- [8] S. Mehraghdam et al., "Specifying and Placing Chains of Virtual Network Functions," in *Proc. of IEEE CloudNet*, 2014.
- [9] P. Vizaretta et al., "QoS-driven Function Placement Reducing Expenditures in NFV Deployments," in *Proc. of ICC*, 2017.
- [10] M. M. Tajiki et al., "Joint Energy Efficient and QoS-aware Path Allocation and VNF Placement for Service Function Chaining," *IEEE TNSM*, 2018.
- [11] Y. Park et al., "Dynamic Defense Provision via Network Functions Virtualization," in *Proc. of ACM SDN-NFV Sec.*, 2017.
- [12] T. V. Phan et al., "Optimizing resource allocation for elastic security VNFs in the SDNFV-enabled cloud computing," in *Proc. of ICOIN*, 2017.
- [13] S. Demirci et al., "Optimal Placement of Virtual Security Functions to Minimize Energy Consumption," in *Proc. of the IEEE ISNCC*, 2018.
- [14] Q. Xu et al., "Low Latency Security Function Chain Embedding Across Multiple Domains," *IEEE Access*, 2018.
- [15] A. Shamel Sendiet et al., "Efficient Provisioning of Security Service Function Chaining Using Network Security Defense Patterns," *IEEE Transactions on Services Computing*, 2017.
- [16] Y. Liu et al., "A Dynamic Composition Mechanism of Security Service Chaining Oriented to SDN/NFV-Enabled Networks," *IEEE Access*, vol. 6, pp. 53 918–53 929, 2018.
- [17] D. Huang et al., *Software-Defined Networking and Security: From Theory to Practice*. CRC Press, 2018.
- [18] N. McKeown et al., "OpenFlow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 2, pp. 69–74, April 2008.
- [22] S. Canard et al., "BlindIDS: Market-Compliant and Privacy-Friendly Intrusion Detection System over Encrypted Traffic," in *Proc. of ACM ASIA CCS*, 2017.
- [19] S. Gianvecchio et al., "Humans and Bots in Internet Chat: Measurement, Analysis, and Automated Classification," *IEEE/ACM Trans. Netw.*, vol. 19, no. 5, pp. 1557–1571, Oct. 2011.
- [20] J. Yan et al., "A Systematic Classification of Cheating in Online Games," in *Proc. of ACM SIGCOMM NetGames*, 2005.
- [21] J. Sherry et al., "BlindBox: Deep Packet Inspection over Encrypted Traffic," in *Proc. of ACM SIGCOMM*, 2015.
- [23] ETSI, "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV," https://www.etsi.org/deliver/etsi_gs/NFV/001_099/003/01.02.01_60/gs_NFV003v010201p.pdf, [Accessed: 31-May-2019].
- [24] M. Antonakakis et al., "Understanding the Mirai Botnet," in *USENIX Security Symposium*, 2017.
- [25] H.D. Sherali et al., "Reformulation–linearization techniques for discrete optimization problems," in *Handbook of Combinatorial Optimization*, P. M. Pardalos, D.-Z. Du, and R. L. Graham, Eds. Springer US, 2013, pp. 2849–2896.
- [26] M. Savi et al., "Impact of Processing-Resource Sharing on the Placement of Chained Virtual Network Functions," *IEEE Transactions on Cloud Computing*, 2019.
- [27] Meihui Gao et al., "Optimal orchestration of virtual network functions," *Computer Networks*, vol. 142, pp. 108 – 127, 2018.
- [28] D.B. Oljira et al., "A Model for QoS-Aware VNF Placement and Provisioning," in *Proc. of IEEE NFV-SDN*, Nov 2017.
- [29] F. Ben Jemaa et al., "QoS-Aware VNF Placement Optimization in Edge-Central Carrier Cloud Architecture," in *Proc. of IEEE GLOBECOM*, 2016.
- [30] K. Jang et al., "Silo: Predictable Message Latency in the Cloud," in *Proc. of the 2015 ACM Conference on Special Interest Group on Data Communication*, 2015.
- [31] Gurobi Optimization, "Constraints," <http://www.gurobi.com/documentation/7.5/refman/constraints.html>, [Accessed: 31-May-2019].
- [32] "Gurobi Optimizer," <http://www.gurobi.com>, [Accessed: 31-May-2019].
- [33] Gurobi Optimization, "Python API," http://www.gurobi.com/documentation/7.5/refman/py_python_api_overview.html, [Accessed: 31-May-2019].
- [34] A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999. [Online]. Available: <http://science.sciencemag.org/content/286/5439/509>
- [35] P. Kazemian et al., "Header Space Analysis: Static Checking for Networks," in *Proc. of the 9th USENIX Conference on Networked Systems Design and Implementation*, 2012.
- [36] GARR, "Considering the Next Generation of GARR Network," <https://www.garr.it/it/documenti/3474-garr-white-paper-maggio-2017>, 2017, [Accessed: 31-May-2019].
- [37] GARR, "GARR Network Map," https://www.garr.it/images/2018-01_mappaGARR_Net.png, 2018, [Accessed: 31-May-2019].
- [38] GARR, "GARR Backbone," <https://www.garr.it/en/infrastructures/network-infrastructure/backbones>, [Accessed: 31-May-2019].
- [39] S.A.R. Shah et al., "Performance comparison of intrusion detection systems and application of machine learning to Snort system," *Future Generation Computer Systems*, vol. 80, no. Supplement C, pp. 157 – 170, 2018.
- [40] D. Lacković et al., "Performance analysis of virtualized VPN endpoints," in *Proc. of 40th MIPRO*, May 2017.
- [41] Fortinet, "Fortigate Virtual Appliances," https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiGate_VM.pdf, 2018, [Accessed: 31-May-2019].
- [42] Cisco Systems, "Adaptive Security Virtual Appliance (ASAv)," <https://www.cisco.com/c/en/us/products/collateral/security/adaptive-security-virtual-appliance-asav/datasheet-c78-733399.pdf>, 2018, [Accessed: 31-May-2019].
- [43] Juniper Networks, "vSRX Virtual Firewall," <https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000489-en.pdf>, 2018, [Accessed: 31-May-2019].