Scalable Cost-Efficient Placement and Chaining of Virtual Network Functions

Marcelo Caggiani Luizelli¹, Luciana Salete Buriol¹, Luciano Paschoal Gaspary¹

¹Instituto de Informática – Universidade Federal do Rio Grande do Sul (UFRGS) Caixa Postal 15.064 – 91.501-970 – Porto Alegre – RS – Brazil

Abstract. While Network Function Virtualization (NFV) is increasingly gaining momentum, with promising benefits of flexible service function deployment and reduced operations and management costs, there are several challenges that remain to be properly tackled, so that it can realize its full potential. One of these challenges, which has a significant impact on the NFV production chain, is effectively and (cost) efficiently deploying service functions, while ensuring that service level agreements are satisfied and making wise allocations of network resources. Despite recent research activity in the field, little has been done towards scalable and cost-efficient placement & chaining of virtual network functions (VNFs) – a key feature for the effective success of NFV. In this thesis, we approach VNF placement and chaining as an optimization problem in the context of Inter- and Intra-datacenter. We formalize the Virtual Network Function Placement and Chaining (VNFPC) problem and propose a mathematical model to solve it. Our model has established one of the first baseline comparison in the field of resource management in NFV and has been widely used in the recent literature. We also address scalability of VNFPC problem to solve large instances by proposing a novel fix-and-optimize-based heuristic algorithm for tackling it. Further, we extensively measure the performance limitations of realistic NFV deployments. Based on that, we propose an analytical model that accurately predict incurred operational costs. Then, we develop an optimal Intra-datacenter service chain deployment mechanism based on our cost model. Finally, we tackle the problem of monitoring service chains in NFV-based environments efficiently.

1. Motivation

Network Functions (NF) play an essential role in today's networks, as they support a diverse set of functions ranging from security (*e.g.*, firewalling and intrusion detection) to performance (*e.g.*, caching and proxying) [Martins et al. 2014]. As implemented, middleboxes are difficult to deploy and maintain. This is mainly because cumbersome procedures need to be followed, such as dealing with a variety of custom-made hardware interfaces and manually chaining middleboxes to ensure the desired network behavior. Further, studies show that the number of middleboxes in enterprise networks (as well as in datacenter and ISP networks) is similar to the number of forwarding devices [Sekar et al. 2012]. Thus, the aforementioned difficulties are exacerbated by the complexity imposed by the high number of network functions that a network provider has to cope with, leading to high operational expenditures. Moreover, in addition to costs related to manually deploying and chaining middleboxes, the need for frequent hardware upgrades adds up to substantial capital investments.

Network Function Virtualization (NFV) has been proposed to shift middlebox processing from specialized hardware appliances to software running on commoditized

hardware [Group 2012]. In addition to potentially reducing acquisition and maintenance costs, NFV is expected to allow network providers to make the most of the benefits of virtualization on the management of network functions (*e.g.*, elasticity, performance, flexibility, etc.). In this context, Software-Defined Networking (SDN) can be considered a convenient complementary technology, which, if available, has the potential to make the chaining of the aforementioned network functions much easier. In fact, it is not unreasonable to state that SDN has the potential to revamp the Service Function Chaining (SFC) problem.

In short, the problem consists of making sure network flows go efficiently through end-to-end paths traversing sets of network functions [Mechtri et al. 2017]. In the NFV/SDN realm and considering the flexibility offered by this environment, the problem consists of (sub)optimally defining how many instances of virtual network functions (VNF) are necessary and where to place them in the infrastructure. Furthermore, the problem encompasses the determination of end-to-end paths over which known network flows have to be transmitted so as to pass through the required placed network functions.

2. Problem Statement

There have been significant achievements in NFV, addressing aspects from effective planning and deployment to efficient operation and management. Nevertheless, NFV is a relatively new and yet maturing paradigm, with various research questions open. As mentioned, one of the most challenging aspects is how to efficiently find a proper VNF placement and chaining.

This problem is particularly challenging for many reasons. First, NFV is an inherently distributed network design based on small cloud nodes spread over the network infrastructure. Therefore, depending on how VNFs are positioned and chained in the infrastructure, end-to-end latencies may become intolerable. This problem is aggravated by the fact that processing times tend to be higher, due to the use of virtualization, and may vary, depending on the type of network function and the hardware configuration of the device hosting it. Second, even when deploying VNFs on a single (small) data center, network services might face performance penalties (and limitations) on critical network metrics (such as throughput, latency, and jitter) depending on how network functions are chained and deployed onto physical servers. Third, resource allocation must be performed in a cost-effective manner, preventing over- or under-provisioning of resources. Therefore, placing network functions and programming network flows in a cost-effective manner while ensuring the required network service performance (e.g., maximum tolerable end-to-end delays) represent an essential step towards enabling the use of NFV in production environments. In the following paragraphs, we provide an overview of the virtual network function placement and chaining (VNFPC) problems tackled in this thesis.

Inter-datacenter Virtual Network Function Placement and Chaining Problem. We begin with a general description of the Inter-datacenter placement and chaining problem with an example, illustrated in Figure 1. It involves the deployment of two Service Function Chaining requests (SFCs, also referred to as "composite services") onto a backbone network. For the first one, incoming flows must be passed through (an instance of) virtual network function (VNF) 1 (e.g., a firewall), and then VNF 2 (e.g., a load balancer). The second specifies that incoming flows must also be passed through VNF 1, and then VNF 3 (e.g., a proxy). Both SFCs are sketched in Figures 1(a) and 1(b), respectively.




Figure 1. Examples of SFCs, and partial view of the network backbone (focusing on the set of N-PoPs available for placing VNFs) considered in our scenario.

The VNF instances required for each composite service must be placed onto Network Points of Presence (N-PoPs). N-PoPs are infrastructures (*e.g.*, servers, clusters or even datacenters) spread in the network and on top of which (virtual) network functions can be provisioned. Without loss of generality, we assume the existence of one N-PoP associated with every major forwarding device comprising a backbone network (see circles in Figure 1(c)). Each N-PoP has a certain amount of resources (*e.g.*, computing power) available. Likewise, each SFC has its own requirements. For example, functions composing an SFC (*e.g.*, caching) are expected to sustain a given load, and thus are associated with a computing power requirement. Also, traffic between functions can be expected to reach some peak throughput, which must be handled by the physical path connecting the N-PoPs hosting those functions.

Given the context above, the first problem we approach is finding a proper placement of VNFs onto distributed N-PoPs and chaining of placed functions, so that overall network resource commitment is minimized. The placement and chaining must ensure that each of the SFC requirements, as well as network constraints, are met. In our illustrating example, a possible deployment is shown in Figure 1(c). The endpoints, represented as filled circles, denote flows originated/destinated from/to devices/networks attached to core forwarding devices. Observe that both composite services share a same instance of VNF 1, placed on N-PoP (2), therefore minimizing resource allocation as desired. Although relatively simple for small instances, the complexity of solving a VNF placement and chaining problem is NP-complete as proven in our paper [Luizelli et al. 2017a].

Intra-datacenter Virtual Network Function Placement and Chaining Problem. The second problem we approach deals specifically with the placement of VNFs onto servers in datacenters (N-PoPs). A solution for the above problem (*i.e.*, the Interdatacenter VNFPC) involves the placement and chaining of SFC requests into multiple (distributed) locations. Particularly, this is the case when SFCs have stringent network requirements (*e.g.*, very low end-to-end delays). As a resulting of the planning, SFCs are broken down into sub-chains (*i.e.*, subgraphs) individually placed and chained into specific datacenters. Then, these partial SFC requests are deployed on top of available commodity servers.

In the context of Intra-datacenter VNFPC, identifying deployment mechanisms that minimize the provisioning cost of service chains has recently received significant attention from both academia and industry¹. However, to the best of our knowledge, ex-

¹To further references on recent NFV efforts, please refer to the thesis.




Figure 2. Deployment strategies for a given a set of service chains Φ and a set of k servers.

isting studies have neglected the actual operational cost of NFV deployments. Therefore typical proposed models (*e.g.*, those implemented in NFV orchestrators) might either lead to infeasible solutions (*e.g.*, in terms of CPU requirements) or suffer high penalties on the expected performance of network services.

In an attempt to address this gap, we first focused on evaluating and modeling the virtual switching cost in a real NFV-based infrastructure [Luizelli et al. 2017b, Luizelli et al. 2018]. In this environment, virtual switching is an essential building block that enables flexible communication between VNFs. However, its operation comes with an extra cost in terms of computing resources that are allocated specifically to software switching in order to steer the traffic through running services (in addition to computing resources required by VNFs). This cost depends primarily on the way VNFs are internally chained, packet processing requirements, and accelerating technologies.

Figure 2 illustrates possible Intra-datacenter deployment solutions for a set of service chains on three identical physical servers. For simplicity, we assume that all traffic steering is done by servers' internal virtual switching, all VNFs require the same amount of processing power to perform their tasks, and that each service chain is associated with a known amount of traffic it has to process. Figures 2(a) and 2(b) illustrate two widely applied VNF deployment strategies in OpenStack Cloud Orchestrator. In Figure 2(a), we depict a deployment strategy where all VNFs of a single SFC are deployed on the same server (referred to as "gather"). In contrast, Figure 2(b) illustrates a deployment strategy where each VNF of a same SFC is deployed onto different servers (referred to as "distribute"). Observe that servers in Figure 2 that are hosting the same number of VNFs are under the same CPU processing requirements. However, determining the amount of processing resources required for switching inside each server is far from being straightforward even when considering simple deployment strategies such the ones illustrated in Figure 2(a) and 2(b). Therefore, understanding these operational overheads in real NFV deployments is of paramount importance for three main reasons: (i) to ensure performance requirements of deployed network services in light of performance limitations (e.g., maximum throughput of a given SFC); (ii) to design efficient Intra-datacenter placement strategies and accurately estimate the incurred cost for any arbitrary deployments (e.g., the one illustrated in Figure 2(c)); and (iii) to reduce the operational cost (in particular, CPU consumption of software switching) of NFV providers. Intuitively, our goal is to develop a placement function that decomposes each service chain into arbitrary decomposition of the service chain, while minimizing the operational cost of the network traffic switching (such as in Figure 2(c)).

3. Related Work

To the best of our knowledge, network function placement and chaining has not been investigated before the inception of NFV, *e.g.*, for planning and deployment of physical middleboxes. One of the most similar problems in the networking literature is Virtual Network Embedding (VNE).

With respect to placement and chaining of network functions, Barkai *et al.* [Barkai et al. 2013] have taken a first step toward modeling this problem. They propose mechanisms to program network flows to an SDN substrate taking into account the position of virtual network functions. Moens *et al.* [Moens and De Turck 2014] were the first to address VNF placement, by formalizing it as an optimization problem. Lewin-Eytan *et al.* [Lewin-Eytan et al. 2015] follow a similar direction, and use an optimization model along with approximation algorithms to solve the problem. Mehraghdam et al. [Mehraghdam et al. 2014], Luizelli et al. [Luizelli et al. 2015] and Bari et al. [Bari et al. 2015] introduce joint optimization problems for placement and chaining of VNFs. However, as they do not consider insights on which VNFs to bring up/down, their deployment solution may end up increasing operational expenditures more than the expected footprint savings.

Most of the effort has been focused on engineering ways of running network functions on top of commodity hardware and, more recently, optimizing aspects in NFV-based environments. Regarding optimization, many studies have been proposed in the context of VNF placement and chaining problems. In spite of their potentialities, the investigations referred above do not properly scale to large network settings and several SFCs submitted in parallel. Our proposal in this thesis enhances the state-of-the-art in that it outperforms the solutions of [Lewin-Eytan et al. 2015, Luizelli et al. 2015], coming up with feasible, high quality solutions for larger scenarios in a timely fashion. Additionally, none of these studies provides optimized operational costs to NFV deployments. Therefore, most of these solutions are not practical in NFV deployments as they might lead to either infeasible or low-performance solutions. Therefore, since most of these works focus on arbitrary cost functions (*e.g.*, reducing the amount of deployed VNFs)— the operational cost of maintaining deployed network services is still an open research field. In this regard, we propose a general deployment mechanism specifically tailored to provide performance-oriented deployment.

4. Contributions of this thesis

This thesis contributes directly to the fields of Computer Networks and Optimization regarding specifically the efficient management of NFV environments. We group the main contributions of this thesis into the five areas below:

1. Inter-Datacenter Placement & Chaining of VNF [Luizelli et al. 2015].

- We formalize the VNFPC problem by means of an Integer Linear Programming model.
- $\bullet~$ We prove the NP-completeness nature of the VNFPC problem.
- We proposed a heuristic procedure that dynamically and efficiently guides the search for solutions performed by commercial solvers.

2. Limited Scalability of State-of-the-Art Solutions [Luizelli et al. 2017a].

- We address scalability of VNFPC by proposing a novel fix-and-optimize-based heuristic algorithm.
- We demonstrate that our proposed method scales to large NFV infrastructures and that our proposed method come up with solution on average 5x better than state-of-the-art solutions.

3. Operational Costs [Luizelli et al. 2017b].

- We provide in-depth insights of SFC deployment strategies in a real NFV infrastructure by measuring overheads and incurred costs.
- We develop an analytical model to properly estimate the cost of software switching for different placement strategies in NFV infrastructures.

4. Intra-Datacenter Placement & Chaining of VNF [Luizelli et al. 2018].

- We generalize our previous analytical model in order to properly estimate any arbitrary SFC deployment placement strategy.
- We develop an online algorithm that significantly reduces the operational costs when compared to commonly employed algorithms in NFV orchestrators (up to a factor of 4).

5. Efficient Network Traffic Monitoring of SFCs.

- We formalize the Distributed Network Monitoring (DNM) problem and propose and Integer Linear Programming model.
- We conducted a set experiments in an NFV-based environment in order to assess the cost
 of sampling network traffic to monitoring systems.

5. Final Remarks

We published the results of this thesis in the most prestigious conferences and journals in the area of Computer Networks including the well-known conference IEEE INFOCOM (h-index \geq 80). The list of currently published papers follows:

- Luizelli et al. Piecing Together the NFV Provisioning Puzzle: Efficient Placement and Chaining of Virtual Network Functions. IFIP/IEEE International Symposium on Integrated Network Management (IM), 2015. Qualis A2.
 - (Best Student Paper Award)
- Luizelli et al. A Fix-and-Optimize Approach for Efficient and Large Scale Virtual Network Function Placement and Chaining. Elsevier Computer Communications, 2016. Qualis A2. Impact factor: 3.33.
- 3. Luizelli et al. **The Actual Cost of Software Switching for NFV Chaining.** IFIP/IEEE International Symposium on Integrated Network Management 2017 (IM), 2017. Qualis A2.
- Luizelli et al. Optimizing NFV Chain Deployment Through Minimizing the Cost of Virtual Switching. IEEE International Conference on Computer Communications (INFOCOM), 2018. Qualis A1.

In addition to the aforementioned main outcomes of this thesis, we further authored/coauthored nine other studies on correlated optimization problems in the context of Computer Networks – including papers in ACM SIGCOMM and IEEE INFOCOM. These publications are listed next.

- 1. Costa Filho et al. **Scalable QoE-aware Path Selection in SDN-based Mobile Networks.** IEEE International Conference on Computer Communications (INFOCOM), 2018. Qualis A1.
- 2. Basat et al. **Constant Time Updates in Hierarchical Heavy Hitters**. ACM Special Interest Group on Data Communication (SIGCOMM), 2017. Qualis A1.
- 3. Enziger et al. Constant Time Weighted Frequency Estimation for Virtual Network Functionalities. IEEE International Conference on Computer Communications and Networks (ICCCN), 2017. Qualis A2.
- Miotto et al. NFV-PEAR: Posicionamento e Encadeamento Adaptativo de Funções Virtuais de Rede. Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC), 2017. Qualis B2.
- 5. Luizelli et al. How physical network topologies affect virtual network embedding quality: A characterization study based on ISP and datacenter networks. Journal of Network and Computer Applications, 2016. Qualis A2. Impact factor: 3.5.
- Nogueira et al. QoS Aware Schedulers for Multi-users on OFDMA Downlink: Optimal and Heuristic. IEEE Latin-American Conference on Communications (LATINCOM), 2016. Qualis B2.
- Luizelli et al. HIPER: Heuristic-based Infrastructure Expansion through Partition Reconnection for Efficient Virtual Network Embedding. IFIP/IEEE/ACM International Conference on Network and Service Management (CNSM), 2014. Qualis A2.

- 8. Muller et al. Survivor: an Enhanced Controller Placement Strategy for Improving SDN Survivability. IEEE Global Communications Conference (GLOBECOM), 2014. Qualis A1.
- Luizelli et al. Reconectando Partições de Infraestruturas Físicas: Rumo a uma Estratégia de Expansão para o Mapeamento Eficiente de Redes Virtuais. Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC), 2014. Qualis B2.

Impact. The impact of this thesis can be measured by the in-depth theoretical and experimental studies developed in the field of Computer Networks. In this thesis, we tackled the VNFPC problem in the context of Inter- and Intra-datacenter. As a major contribution to the state-of-the-art, we first formalized the Inter-datacenter VNFPC and proposed an optimization model to solve it [Luizelli et al. 2015] (awarded best student paper). Our mathematical model has established one of the first baseline comparison in the field of resource management in NFV and has been widely used in the recent related literature (≥ 115 citations). Further, we prove computational complexity results for VNFPC problems and devised a novel Fix-and-Optimize based approach for efficiently solving the problem for realistic and large instances [Luizelli et al. 2017a]. Our proposed approach outperformed state-of-the-art solution by a factor up to 5.

Beyond the theoretical studies developed in this thesis, we backed up our mathematical models by an extensive and realistic set of experiments in an NFV environment (thanks to Nokia Bell Labs infrastructure) [Luizelli et al. 2017b]. These experiments helped us to better comprehend the effects of empirical environments on the performance of deployed network services and to fine-tuning our models. From these empirical experiments, we generalized our CPU cost models for arbitrary NFV setups. Then, we designed the Operational Cost Minimization (OCM) algorithm for the efficient and cost-oriented Intra-datacenter placement of service chains [Luizelli et al. 2018]. OCM mechanism heavily relies on the usage of our previously defined cost models to guide the search for optimal solutions and can outperform commonly used algorithms in NFV orchestrators by a factor of 4. Finally, after diving specifically into the provisioning of NFV services in both Inter- and Intra-datacenter contexts, we tackled the problem of monitoring network traffic of service chains – another essential building block for the proper operation of NFV-based network infrastructures. For that, we formalized the Distributed Network Monitoring (DNM) problem and proposed a mathematical optimization model to effectively coordinate the monitoring of service chains in the context of NFV/SDN deployments. It is worth mentioning that this last work has not yet been published.

This research was partially developed in collaboration with the prestigious research laboratory Nokia Bell Labs and Technion University (both in Israel) during the one-year CNPq granted scholarschip. One of the main outcomes of this thesis, the proposed OCM algorithm, is currently being adapted and evaluated by NOKIA in CloudBand NFV Orchestrator. The development of this thesis allowed to promote a strong relationship with renowned scientists in top-ranked research labs and Universities. It is worth mentioning that this thesis was concluded in 3.5 years so the author of this thesis could assume a position at Federal University of Pampa (Unipampa). Last, as NFV is a recent network paradigm and yet in a maturing phase, there are many open research problems that were out of the scope of this thesis. In this context, many recent papers have been exploring research directions pointed out by this thesis.

References

Bari, M. F., Chowdhury, S. R., Ahmed, R., and Boutaba, R. (2015). On orchestrating virtual network functions. In *Proceedings...*, CNSM '15, pages 50–56, Washington, DC, USA. Inter-

- national Conference on Network and Service Management, IEEE Computer Society.
- Barkai, S., Katz, R., Farinacci, D., and Meyer, D. (2013). Software defined flow-mapping for scaling virtualized network functions. In *Proceedings...*, ACM SIGCOMM HotSDN'13, pages 149–150, New York, NY, USA. ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking, ACM.
- Group, N. F. I. S. (2012). Network function virtualisation (nfv): An introduction, benefits, enablers, challenges and call for action. Available at: https://portal.etsi.org/nfv/. Accessed on: Jan. 20, 2015.
- Lewin-Eytan, L., Naor, J., Cohen, R., and Raz, D. (2015). Near optimal placement of virtual network functions. In *Proceedings...*, INFOCOM'15, pages 1346–1354, New York, NY, USA. IEEE INFOCOM, IEEE.
- Luizelli, M. C., Bays, L. R., Buriol, L. S., Barcellos, M. P., and Gaspary, L. P. (2015). Piecing together the nfv provisioning puzzle: Efficient placement and chaining of virtual network functions. In *Proceedings...*, IM'15, pages 98–106, New York, NY, USA. IFIP/IEEE International Symposium on Integrated Network Management, IEEE.
- Luizelli, M. C., da Costa Cordeiro, W. L., Buriol, L. S., and Gaspary, L. P. (2017a). A fix-and-optimize approach for efficient and large scale virtual network function placement and chaining. *Computer Communications*, 102:67 77.
- Luizelli, M. C., Raz, D., Saar, Y., and Yallouz, J. (2017b). The actual cost of software switching for nfv chaining. In *Proceedings...*, IM'17, pages 335–343, New York, NY, USA. IFIP/IEEE International Symposium on Integrated Network Management, IEEE.
- Luizelli, M. C., Saar, Y., and Raz, D. (2018). Optimizing nfv chain deployment through minimizing the cost of virtual switching. In *Proceedings of IEEE International Conference on Computer Communications*, INFOCOM '18, pages 1–9, Piscataway, NJ, USA. IEEE INFOCOM, IEEE Press.
- Martins, J., Ahmed, M., Raiciu, C., Olteanu, V., Honda, M., Bifulco, R., and Huici, F. (2014). Clickos and the art of network function virtualization. In *Proceedings...*, NSDI'14, Berkeley, CA, USA. USENIX Conference on Networked Systems Design and Implementation, USENIX Association.
- Mechtri, M., Ghribi, C., Soualah, O., and Zeghlache, D. (2017). Nfv orchestration framework addressing sfc challenges. *IEEE Communications Magazine*, 55(6):16–23.
- Mehraghdam, S., Keller, M., and Karl, H. (2014). Specifying and placing chains of virtual network functions. In *Proceedings...*, CloudNet'14, pages 7–13, Piscataway, NJ, USA. IEEE International Conference on Cloud Networking, IEEE Press.
- Moens, H. and De Turck, F. (2014). Vnf-p: A model for efficient placement of virtualized network functions. In *Proceedings...*, CNSM'14 MiniConf, pages 418–423, Piscataway, NJ, USA. International Conference on Network and Service Management, IEEE Press.
- Sekar, V., Egi, N., Ratnasamy, S., Reiter, M. K., and Shi, G. (2012). Design and implementation of a consolidated middlebox architecture. In *Proceedings...*, NSDI'12, page 24, Berkeley, CA, USA. USENIX Conference on Networked Systems Design and Implementation, USENIX Association.