



HAL
open science

Efficient Management of Virtualized Information-Centric Networks

Maxim Claeys, Steven Latré, Filip De Turck

► **To cite this version:**

Maxim Claeys, Steven Latré, Filip De Turck. Efficient Management of Virtualized Information-Centric Networks. 8th IFIP International Conference on Autonomous Infrastructure, Management and Security (AIMS), Jun 2014, Brno, Czech Republic. pp.42-46, 10.1007/978-3-662-43862-6_4. hal-01401285

HAL Id: hal-01401285

<https://inria.hal.science/hal-01401285>

Submitted on 23 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Efficient Management of Virtualized Information-Centric Networks

Maxim Claeys¹, Steven Latré², and Filip De Turck¹

¹ Department of Information Technology, Ghent University - iMinds
{maxim.claeys, filip.deturck}@intec.ugent.be

² Department of Mathematics and Computer Science, University of Antwerp - iMinds
steven.latre@uantwerpen.be

Abstract. The Internet has rapidly evolved from a network, connecting a couple of dozens of computers, to a network containing billions of devices. Furthermore, the current Internet is mostly used to deliver complex services with increasingly stringent Quality of Service (QoS) requirements. However, the underlying network model has remained the same, making the Internet not well suited to optimally support the current user trends and services. Currently, a lot of effort is being made in the area of network virtualization and Information-Centric Networking (ICN) to support the evolution towards the QoS constraint distribution of large amounts of information. Even though both directions offer a lot of opportunities, multiple important challenges have to be faced when managing the placement of content inside the network and guaranteeing delivery efficiency. These challenges are further increased when a combination of both trends is considered. This paper gives an overview of these challenges and how this PhD will deal with the mutual influences of network virtualization and ICN in an efficient way.

1 Introduction

Since its emergence in the 70's, the Internet has known a phenomenal growth. However, the fundamental principles have remained the same. In recent years, Internet usage has shifted to a variety of complex services that require the distribution of large amounts of information. These services include, amongst others, the delivery of Video on Demand (VoD) and popular social media like Twitter and Facebook. The delivery of these complex services is in strong contrast to the traditional usage of the Internet. It is the information itself, and no longer the communicating hosts, that plays a central role in the Internet. The source of the information is of minor importance to the end-user, as long as the delivery meets increasingly stringent Quality of Service (QoS) requirements.

These observations reveal multiple strong management challenges for the future Internet. In this area, two novel research directions can be identified. The development of Information-Centric Networking (ICN) architectures and algorithms [1], where the information is decoupled from its location at the network level, addresses the increasing importance of content distribution. This

allows content to be duplicated everywhere in the network, enabling in-network caching [2]. Addressing individual data elements is in strong contrast to the current host-based IP addressing scheme.

On the other hand, a lot of research effort is being made in the area of network virtualization [3]. Just as the virtualization of computing resources enabled a lot of novel applications like cloud computing in data centers, network virtualization may offer similar perspectives in communication networks [4]. Network virtualization allows the differentiation of distinct services over a shared physical network. Furthermore, decoupling a service from a fixed physical location allows services to be dynamically configured and reconfigured. As these services are fully isolated, they can be managed individually. In this way, individual QoS requirements can be imposed.

Although both research areas offer a lot of new possibilities, multiple important challenges have to be dealt with. A first challenge can be found in ICN-based routing. Next to the fact that the number of data objects is a level of magnitude bigger than the number of end systems, multiple duplicates of data objects, spread across the network, further complicates the routing problem. Second, in-network caching can strongly benefit from efficient caching strategies. Advanced algorithms, considering pro-active content placement next to current reactive caching strategies, can strongly increase the delivery performance in the network. A third challenge can be found in finding the optimal resource allocation scheme in the virtualized network [5]. This challenge increases even further when reallocation in case of changing requirements or network conditions is considered.

2 Virtualized ICN Networks

This PhD research covers the design of an architecture and algorithms for the efficient management of virtualized information-centric networks. More specific, we focus on the elastic allocation of virtual network resources and on pro-active content placement. The remainder of this section will focus on the different algorithmic challenges of virtualized information-centric networks that will be addressed in this PhD.

2.1 Static resource-allocation and content placement

In this research phase, algorithms will be developed to calculate the optimal network configuration that meets the specified QoS requirements while minimizing the related costs. The resulting configuration will define both the content placement and the allocation of virtual network slices. This process is illustrated in Fig. 1. The system takes into account both the predicted requests and related QoS requirements for the service (A) and the physical network topology (B). Based on this input, the algorithm defines the optimal content placement in order to provide a cost-efficient service delivery and allocates the virtual network slices. The resulting solution satisfies both the service requirements and the physical resource capacities.

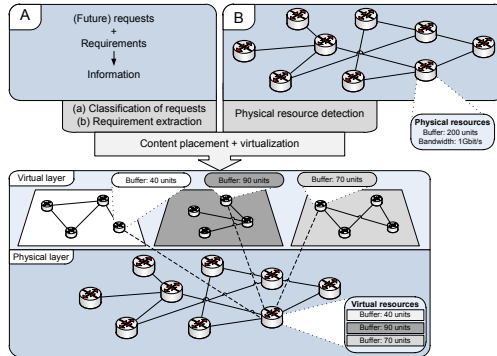


Fig. 1. Schematic overview of the static resource-allocation and content placement process.

Both the allocation of virtual networks, taking into account physical resource capacities, and the content placement problem are known to be NP-hard [5, 6]. Given the mutual influence between both problems, a combined approach is required. Because of the algorithmic complexity of the individual problems, an optimal solution can not be obtained in a scalable way. Therefore, we will have to reside to heuristic approaches. During this PhD, multiple algorithms will be developed, both using meta-heuristics and custom heuristics, tailored to the specific problem. For this purpose, multiple existing heuristics, both for the resource allocation problem (e.g. [7, 8]) and the content placement problem (e.g. [9, 10]), will be investigated in depth.

Even though we stated in the previous paragraph that an optimal solution can not be guaranteed in a scalable way, it will be very useful to model the problem as an Integer Linear Programming (ILP) problem. Using this formulation, an optimal solution can be calculated for problems of limited scale. These results can then be used as a benchmark to evaluate the developed heuristics.

2.2 Elastic resource-allocation and content placement

Given the high algorithmic complexity of the algorithms introduced in Section 2.1, the allocation of virtual network slices and the content placement can only be executed periodically in an off-line way on a relative large timescale. During the period in between two executions however, strong fluctuations in the request pattern can occur. Furthermore, requests for additional services can cause the optimal network configuration and content placement to change. For this reason, it is crucial to monitor the conditions and optimize the configuration on a short timescale if needed. In this way, virtual networks can be scaled in an elastic way by dynamically allocating more or less resources and moving content inside the network, based on the perceived requirements, analogous to virtual machines in cloud computing. It is important to note that these resources are not limited to traditional hardware resources like CPU and bandwidth capacity,

but can additionally include higher level resources, e.g. the availability of specific protocols.

The elastic approach is illustrated in Fig. 2. It is important to note that, for reasons of clarity, the requirements are visualized in only one dimension. In reality however, the problem is more complex. As illustrated in this example, the allocated resources for the virtual network slice scale when the requirements change significantly. Furthermore, this reallocation could introduce a better content placement solution. In this way, the elastic network will be able to meet the QoS requirements at any time, while minimizing the related costs.

Requirements are very likely to fluctuate lightly during the monitoring period. However, small fluctuations do not necessarily cause the current configuration to be unsatisfactory. When the perceived requirements stay within a specific range, reconfigurations are not required. To detect consistent anomalous behaviour, anomaly detection techniques can be applied. Anomaly detection identifies patterns that do not correspond to the perceived normal behaviour. When such anomalies are detected, the resource allocation and content placement has to be reconfigured. A first step in this area has been taken with the implementation of a heuristic, dynamically adjusting the resource allocation of virtual networks using reinforcement learning [11]. In this way, the heuristic can adapt its behaviour based on the performance of previous decisions. Similar to the static case described in Section 2.1, this problem will be modelled as an ILP formulation to serve as a benchmark for the developed algorithms.




Fig. 2. Simplified schematic overview of the elastic scaling of virtual network slices, based on perceived requirements.

3 Conclusion

For the Internet to efficiently deal with the delivery of current complex services, this paper proposed a combined approach to the problem of resource allocation in network virtualization and the information-centric content placement problem. A combined approach is crucial given the mutual influence between both problems. As a first step, heuristic approaches are proposed to define the optimal network configuration and content placement in a static way on a large time-scale. On a shorter time-scale, self-learning heuristics are proposed to adapt the network configuration and content placement to perceived fluctuations in the requests and requirements patterns. Using this approach, the elastic network is able to meet the QoS requirements in a cost-efficient way at any time.

Acknowledgements. Maxim Claeys is funded by grant of the Agency of Innovation by Science and Technology in Flanders (IWT). This work was partly funded by Flamingo, a Network of Excellence project (318488) supported by the European Commission under its Seventh Framework Programme.

References

1. Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutscher, D., Ohlman, B.: A survey of information-centric networking. *IEEE Communications Magazine* **50** (2012) 26–36
2. Psaras, I., Chai, W.K., Pavlou, G.: Probabilistic in-network caching for information-centric networks. In: *ICN Workshop on Information-centric Networking*. (2012)
3. Chowdhury, N.M.K., Boutaba, R.: A survey of network virtualization. *Computer Networks* **54** (2010) 862–876
4. Khan, A., Zugenmaier, A., Jurca, D., Kellerer, W.: Network virtualization: a hypervisor for the internet? *IEEE Communications Magazine* **50** (2012) 136–143
5. Haider, A., Potter, R., Nakao, A.: Challenges in resource allocation in network virtualization. In: *20th ITC Specialist Seminar*. (2009)
6. Sourlas, V., Flegkas, P., Paschos, G.S., Katsaros, D., Tassiulas, L.: Storage planning and replica assignment in content-centric publish/subscribe networks. *Computer Networks* **55** (2011) 4021–4032
7. Balasubramaniam, S., Mineraud, J., Perry, P., Jennings, B., Murphy, L., Donnelly, W., Botvich, D.: Coordinating allocation of resources for multiple virtual iptv providers to maximize revenue. *IEEE Transactions on Broadcasting* **57** (2011) 826–839
8. Leon-Garcia, A., Mason, L.: Virtual network resource management for next-generation networks. *Communications Magazine, IEEE* **41** (2003) 102–109
9. Qiu, L., Padmanabhan, V.N., Voelker, G.M.: On the placement of web server replicas. In: *IEEE InfoComm*. (2001)
10. Kangasharju, J., Roberts, J., Ross, K.W.: Object replication strategies in content distribution networks. *Comput. Commun.* **25** (2002) 376–383
11. Mijumbi, R., Serrat, J., Gorricho, J.L., Claeys, M., De Turck, F., Latré, S.: Design and evaluation of learning algorithms for dynamic resource management in virtual networks. In: *IFIP/IEEE Network Operations and Management Symposium (NOMS) 2014, Krakow, Poland (May 2014)*