

Optimization Models for Flexible and Adaptive SDN Network Virtualization Layers

Johannes Zerwas, Andreas Blenk, Wolfgang Kellerer
Chair of Communication Networks, Technical University of Munich, Germany
{johannes.zerwas, andreas.blenk, wolfgang.kellerer}@tum.de

ABSTRACT

Network hypervisors provide the network virtualization layer for Software Defined Networking (SDN). They enable virtual network (VN) tenants to bring their SDN controllers to program their logical networks individually according to their demands. In order to make use of the high flexibility of virtual SDN networks and to provide high performance, the deployment of the virtualization layer needs to adapt to changing VN demands. This paper initializes the study of the optimization of dynamic SDN network virtualization layers. Based on the definition of reconfiguration events, we formalized mixed integer programs to analyze the multi-objective problem of adapting virtualization layers. Our initial simulation results demonstrate Pareto frontiers of conflicting objectives, namely control plane latency and hypervisor and control path reconfigurations.

CCS Concepts

•**Networks** → *Network performance modeling; Logical / virtual topologies; Network dynamics; Programmable networks*; •**General and reference** → Performance;

Keywords

Software defined networking; network virtualization; reconfiguration; multi-objective optimization

1. INTRODUCTION

SDN network hypervisors provide the ability to control virtual networks (VNs) via SDN-based controllers. Hypervisors connect tenant SDN controllers with their logical VNs. Combining SDN and network virtualization provides the benefits of both worlds, i.e., flexible network resource sharing due to virtualization [25, 19, 3] and programmability and adaptability due to open interfaces provided by SDN [18, 23, 16]. Accordingly, many network hypervisor solutions have already been proposed [4, 2, 22, 13, 21, 1, 9, 8]. While these works only focus on the implementation aspect of network hypervisors, we draw attention to the occurring optimization problems when virtualizing SDN networks with focus on VN dynamics, e.g., due to changing user demands.

As it has been stated in the controller placement problem (CPP) [12], the location of the control logic and the number of controllers can severely impact the performance of SDN networks. For instance, long flow setup times can directly affect the perceived user quality, e.g., in case of web page requests where every millisecond counts [24]. Thus, as network hypervisors provide the control logic for virtualization, their placement inside the network needs to be optimized [5, 6]. In order to pay attention to the dynamic nature of VNs [14], e.g., due to changing user demands [17, 10, 7], we focus on the dynamic hypervisor placement problem (DHPP).

To always gain high efficiency, e.g., low control plane latency for VNs, hypervisor locations need to be adapted in case of changing VNs at runtime. However, such operation leads to reconfigurations of the existing hypervisor placement, e.g., migration of hypervisor entities or changing the routing of control paths between hypervisors and tenant SDN controllers. Generally, network reconfigurations can lead to significant side effects, such as high latency or even network outages [20, 11]. In this paper, we model the different reconfiguration types of virtual SDN networks. Using mixed integer programming models for varying latency performance measures [5, 6], we conducted an initial simulation study that focus on trade-offs between latency and reconfiguration measures. Our initial results gain first insights into trade-offs among conflicting optimization goals, i.e., latency and reconfigurations.

2. RECONFIGURATION MODELS

In virtual SDN environments, multiple tenant SDN controllers (SDN-C), hypervisor entities (HVs) and physical nodes hosting virtual switches (Vs) exist. The SDN-Cs are connected via the hypervisor instances to their logical switches (Vs). A virtual SDN control path (VCP) is identified via a tenant SDN controller (SDN-C), a logical SDN node (V), and a hypervisor entity (HV), which connects both. Inside a network, a hypervisor entity (HV) is placed at a specific location (LOC). The solution of a static hypervisor placement, i.e., one that does not consider reconfigurations, provides the location of HVs and the routing of VCPs, i.e., which tenant controller is connected via which entity to its logical switch,




Figure 1: Reconfigurations for a virtual control path (VCP) in a virtualized SDN environment. Tenant controller (C), illustrated via boxes, connects via hypervisor (H), illustrated as hexagons, to virtual switch (V).

for an objective like average control plane latency.

In case of adapting an HV placement, three different VCP reconfiguration events can be classified. Fig. 1 shows an initial HV placement (a) and the three reconfigurations (b-d) for one controller (C) and its logical switch (V) (both illustrated via squares), and two HV instances 1 and 2 (ill. via hexagons). The substrate consists of four nodes (circles) and four physical links (dashed lines). In the initial placement (Fig. 1(a)), the VCP (solid line) between V and C is routed via HV 1. A hypervisor migration, i.e., VCP location change (LOC), is shown in Fig. 1(b) where HV 1 is migrated to the right node, whereas HV 2 is shutdown, e.g., due to energy savings. Here, the VCP is still routed via HV 1, however the routing has been adapted. Fig 1(c) shows an HV change for the VCP. Here, the HV entity changed whereas the VCP routing stays. Finally, the HV entity and the routing (LOC) are changed in (Fig. 1(d)). For considering the reconfigurations while solving the HPP, two metrics provide the number of reconfigurations of VCPs between the current and the new HV placement. $R_{VCP,LOC}$ determines the number of VCP location changes (Fig.1(a)). $R_{VCP,HV}$ determines VCPs that migrated from one HV entity to another.

3. INITIAL OPTIMIZATION RESULTS

To evaluate the interaction of the latency and reconfiguration objectives, our existing HPP framework [5, 6] is extended to maintain initial HV placements (HP) and address the reconfiguration objectives. We investigate the dynamic hypervisor placement problem for the case where an additional vSDN needs to be added to the network, i.e., for an increasing number of hosted vSDNs. In particular, we would like to analyze the impact of adding vSDNs on the average control plane latency and the amount of reconfigurations, which are needed to adapt the HP towards an optimal latency placement.

Multi-stage optimization is considered as approach to perform multi-objective optimization for three objectives, namely minimum average HP control plane latency (L_{avg}), minimum LOC changes, and minimum HV entity changes. In Stage 1, the average HP control

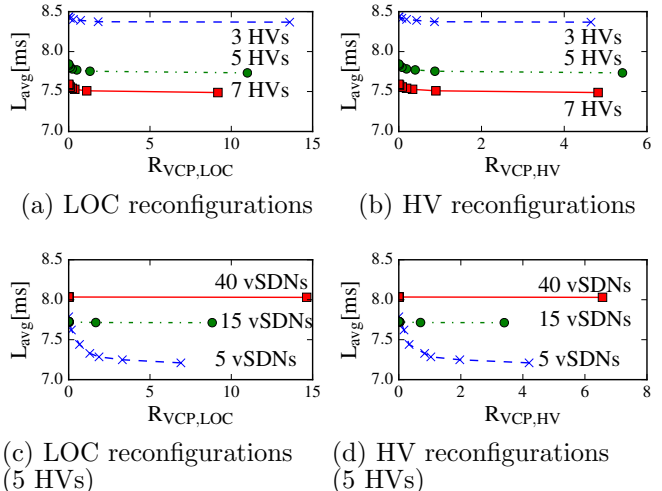


Figure 2: Figures showing trade off between average control plane latency L_{avg} and number of reconfigurations.

plane latency (L_{avg} in milliseconds [ms]) is optimized. The LOCs are minimized in Stage 2. In Stage 3, the HV entity changes are minimized. The latency value obtained in Stage 1 is used to upper bound the latency of Stage 2 and 3. A latency relaxation factor ρ in Stage 1 and 2 is introduced to determine the Pareto frontier.

The ATT MPLS network from the Topology Zoo [15] is used as physical network. The size and embedding of the virtual SDN networks (vSDNs) is randomly distributed on the physical network. The number of initially embedded vSDNs varies between 5 and 40. For each number of embedded vSDNs, 30 different sets are generated to gain statistical confidence. We analyze the impact of the number of HVs (3, 5, 7) and the impact of the number of vSDNs (5, 15, 40). Fig. 2(a-b) show L_{avg} against the number of reconfigurations for 3, 5, 7 HVs and 5–40 vSDNs. The marker represent the mean values over number of vSDNs and runs for ρ varying between 0% and 10%. We observe that a small latency relaxation of 1% already decreases both $R_{VCP,LOC}$ and $R_{VCP,HV}$ significantly. For larger ρ , the reduction of the reconfiguration metrics diminishes. A $\rho = 10\%$ results in avoiding reconfigurations on average.

Fig. 2(c)-(d) illustrate the impact of initially embedded vSDNs on the reconfigurations for 5 HVs. For 5 vSDNs, avoiding reconfigurations requires a larger latency increase than for setups that contain more vSDNs. A latency increase from 7.17 ms to 7.67 ms (7%) is required to reduce $R_{VCP,LOC}$ to 0. In contrast, for 40 vSDNs only a latency increase of about 0.017ms avoids all reconfigurations. Considering $R_{VCP,ID}$, similar observations are made. Thus, the number of vSDNs should be taken into account while performing the DHPP.

4. CONCLUSIONS

SDN network hypervisors enable tenants to bring their

own controller. In this paper, we initialized the study of the dynamic hypervisor placement problem. The dynamic problem introduces conflicting objectives, namely latency and reconfigurations. Based on mixed integer programming formulations, we conducted a study of the multi-objective optimization problem. We conclude that a latency relaxation of 7% already leads to no reconfigurations at all. Generally, our models can serve as baselines when designing heuristic solutions. For future work, we plan to extend our study with different objectives and analyze it for larger substrate networks.

Acknowledgment

This work has been performed in part in the framework of the CELTIC EUREKA project SENDATE-PLANETS (Project ID C2015/3-1) and is partly funded by the German BMBF (Project ID 16KIS0473), and in part in the framework of the EU project FlexNets funded by the European Research Council under the European Unions Horizon 2020 research and innovation program (grant agreement No 647158 - FlexNets). The authors alone are responsible for the content of the paper.

5. REFERENCES

- [1] A. Al-Shabibi, M. De Leenheer, M. Gerola, A. Koshibe, W. Snow, and G. Parulkar. OpenVirteX: a network hypervisor. In *Proc. of the Open Networking Summit*, 2014.
- [2] A. Al-Shabibi, M. D. Leenheer, and B. Snow. OpenVirteX: Make your virtual SDNs programmable. In *Proc. of the ACM Workshop HotSDN*, pages 25–30, 2014.
- [3] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. *Computer*, 38(4):34–41, 2005.
- [4] A. Blenk, A. Basta, M. Reisslein, and W. Kellerer. Survey on network virtualization hypervisors for Software Defined Networking. *IEEE Commun. Surv. & Tut.*, 18(1):655–685, Jan. 2016.
- [5] A. Blenk, A. Basta, J. Zerwas, and W. Kellerer. Pairing SDN with network virtualization: The network hypervisor placement problem. In *Proc. of the IEEE NFV-SDN*, pages 198–204, San Francisco, CA, Dec. 2015.
- [6] A. Blenk, A. Basta, J. Zerwas, M. Reisslein, and W. Kellerer. Control Plane Latency with SDN Network Hypervisors: The Cost of Virtualization. *IEEE TNSM*, 4537(c), 2016.
- [7] A. Blenk and W. Kellerer. Traffic pattern based virtual network embedding. In *Proc. of the ACM CoNEXT Student Workshop*, pages 23–26, New York, NY, 2013. ACM Press.
- [8] Z. Bozakov and P. Papadimitriou. AutoSlice: automated and scalable slicing for software-defined networks. In *Proc. of the ACM CoNEXT Student Workshop*, pages 3–4, Nice, France, 2012.
- [9] Z. Bozakov and P. Papadimitriou. Towards a scalable software-defined network virtualization platform. In *Proc. of the IEEE/IFIP NOMS 2014*, pages 1–8, 2014.
- [10] J. Eрман and K. Ramakrishnan. Understanding the super-sized traffic of the super bowl. *Proc. of the ACM SIGCOMM Internet measurement conference - IMC '13*, pages 353–360, 2013.
- [11] P. Gill, N. Jain, and N. Nagappan. Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications. In *Proc. of the ACM SIGCOMM 2011 conference*, 2011.
- [12] B. Heller, R. Sherwood, and N. McKeown. The controller placement problem. *ACM SIGCOMM CCR*, 42(4):473, 2012.
- [13] X. Jin, J. Gossels, J. Rexford, and D. Walker. CoVisor: A compositional hypervisor for software-defined networks. In *Proc. of the USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, pages 87–101, 2015.
- [14] E. Keller, D. Arora, S. Ghorbani, M. Caesar, and J. Rexford. Live Migration of an Entire Network (and its Hosts). *Princeton University Computer Science Technical Report*, pages 109–114, 2012.
- [15] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The internet topology zoo. *IEEE JSAC*, 29:1765–1775, 2011.
- [16] D. Kreutz, F. M. V. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig. Software-defined networking: A comprehensive survey. *Proc. of the IEEE*, 103(1):14–76, 2015.
- [17] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On dominant characteristics of residential broadband internet traffic. In *Proc. of the 9th ACM SIGCOMM Internet measurement conference - IMC '09*, pages 90–102, New York, New York, USA, 2009.
- [18] N. McKeown, et al. OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM CCR*, 38(2):69–74, 2008.
- [19] N. Nikaiein, E. Schiller, R. Favraud, K. Katsalis, D. Stavropoulos, I. Alyafawi, Z. Zhao, T. Braun, and T. Korakis. Network Store: Exploring Slicing in Future 5G Networks. In *Proc. of MobiArch '15*, pages 8–13, 2015.
- [20] A. Pathak, M. Zhang, Y. C. Hu, R. Mahajan, and D. Maltz. Latency inflation with MPLS-based traffic engineering. In *Proc. of the ACM SIGCOMM Internet measurement conference - IMC '11*, pages 463–472, New York, New York, USA, 2011.
- [21] J. Rexford and D. Walker. Incremental update for a compositional SDN hypervisor categories and

- subject descriptors. In *Proc. of the ACM Workshop HotSDN*, pages 187–192, 2014.
- [22] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. Mckeown, and G. Parulkar. FlowVisor: A network virtualization layer. Technical report, OpenFlow Consortium, 2009.
- [23] V. Sivaraman, T. Moors, H. Habibi Gharakheili, D. Ong, J. Matthews, and C. Russell. Virtualizing the access network via open APIs. In *Proc. of the ACM CoNEXT '13*, pages 31–42, New York, New York, USA, 2013.
- [24] X. S. Wang. Speeding up Web Page Loads with Shandian. In *Proc. of the USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, pages 109–122, 2016.
- [25] K. Webb, a.C. Snoeren, and K. Yocum. Topology switching for data center networks. *Proc. of the 11th USENIX conference on Hot topics in management of internet, cloud, and enterprise networks and services*, pages 14–14, 2011.