Accepted Manuscript

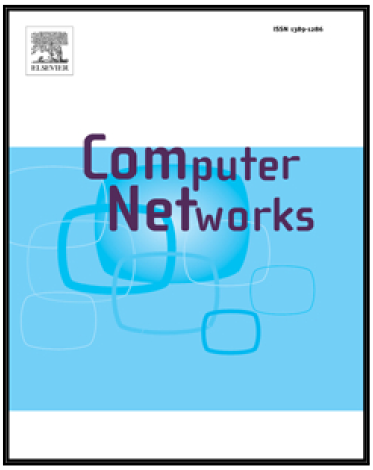

Towards Full Virtualization of SDN Infrastructure

Junfeng Li, Dan Li, Yirong Yu, Yukai Huang, Jing Zhu, Jinkun Geng

PII: S1389-1286(18)30385-2
DOI: 10.1016/j.comnet.2018.06.014
Reference: COMPNW 6524

To appear in: *Computer Networks*

Received date: 4 January 2018
Revised date: 16 May 2018
Accepted date: 13 June 2018

Towards Full Virtualization of SDN Infrastructure

Junfeng Li, Dan Li, Yirong Yu, Yukai Huang, Jing Zhu and Jinkun Geng

Abstract—With the emerging multi-tenant virtual networks in the cloud data center and the network slicing in wide-area network, *full virtualization* of the hardware network infrastructure is required. In this work, we propose SVirt, a full virtualization architecture of the software defined networking (SDN) infrastructure, which can flexibly support multiple virtual networks (tenants) to simultaneously run SDN on top of the same physical network. In order to simultaneously run multiple virtual SDN switches with different processing logics on the same physical switch, we redesign the physical switch's processing pipeline with a "*late-binding key extractor (LBX)*". In the control plane, SVirt enables "*many-to-one*" and "*one-to-many*" mapping when allocating the physical resource for a virtual network, which embraces arbitrary topology and ternary content-addressable memory (TCAM) resource demanded by a virtual network. In the data plane, SVirt explicitly carries the forwarding context information in the packets, overcoming the "*context-loss problem*" in a virtual network. We develop a NetFPGA prototype of SVirt switch. Event-driven simulations and prototype-based experiments demonstrate that, compared with traditional approaches, SVirt significantly enhances the cloud's capability to accept various virtual SDN requests with resource utilization ratio of more than 90% and improves the network's throughput by around 236%. Besides, SVirt can hold more flows than traditional approaches, thus reducing two-order more table misses.

Index Terms—Software defined networking, Network Virtualization, Virtual network embedding, Switch.

1 Introduction

Network virtualization draws more and more attention in recent years. In public cloud, tenants outsource their infrastructure to the cloud data center, and each tenant has the specific requirement for his virtual network in the cloud [2]. In the wide-area network, people are thinking on network slicing, so as to support multiple logical (virtualized) network infrastructure, such as phone network and IoT network, in the future 5G core network [3], [4]. Software defined networking (SDN) significantly enhances the programmability of physical networks [5]. By P4 language [6] and programmable switches [16], network operators can arbitrarily define the matching fields and actions to process packets in the switches. However, a desired SDN virtualization (slicing) architecture requires the routers/switches to simultaneously support multiple different processing logics in the same pipeline stage, which recent works cannot well address. Previous SDN virtualization solutions [10], [11], [12], [13] seek to achieve flexible topology and isolated flow space for a virtual network, but lack consideration of the constraint to a virtual network caused by the processing logic and resources of individual physical switches. Network function virtualization (NFV) is another important enabling technology for network virtualization, but existing software network platforms cannot compete against the hardware routers/switches when handling complex packet processing logic [14].

Full virtualization is a common item in NFV and it has several definitions with some trivial differences. In our paper, "full virtualization" means that the experience of a tenant obtains from the virtual network is the same as if he builds a physical network infrastructure by his own. In order to provide *full virtualization* service of the *hardware* SDN infrastructure to the tenants, we abstract the physical network as a single and large *network resource pool*. By enabling highly flexible mapping from a virtual network to the physical network, we mask the details of the physical network, such as the network topology, the number of switches, the packet processing pipeline and ternary content-addressable memory (TCAM) space in an individual switch, from tenants. A tenant gets his virtual network service based on whatever network topology and routing/security policy he defines on virtual switches. If we can achieve this goal of full virtualization, tenants can get the same experience as building physical networks by themselves, which should be ultimate goal of network virtualization.

It is non-trivial to realize full virtualization of the hardware network infrastructure. First, the *packet processing pipeline* of today's switches is fixed, either by hard-coded switching chips or by reconfigurable switching chips [8], [16], [27]. But in network full virtualization, we need to map multiple virtual switches with different, flexibly-defined processing pipelines into the same physical switch. Hence, the physical switch should be redesigned, so as to let each pipeline stage of the switch be able to simultaneously process multiple different logics. Second, considering the mismatch between the resource (such as the number of switches, the length of processing pipeline and the TCAM space) required by the virtual network and that of the physical network, it is necessary to enable "*many-to-one*" mapping and "*one-to-many*" mapping. When this occurs,

- *Junfeng Li, Dan Li, Yirong Yu, Yukai Huang, Jing Zhu and Jinkun Geng are with Department of Computer Science and Technology, Tsinghua University, Beijing 10084, China.*

In this paper, we add the following work based on a previous version [1] published at IEEE ICNP '15: 1) We make the motivation and position of this work more clearly; 2) We formulate the network embedding problem and prove it is NP-Hard; 3) We describe the network embedding algorithm in detail; 4) We add more simulations to evaluate the solution.

No conflict of interest exists in the submission of this manuscript. None of the funding sources played a role in the design, collection, analysis or interpretation of the data, in the writing of the report, or in the decision to submit the manuscript for publication.