Computer Networks 68 (2014) 1–4

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Editorial

Cloud networking and communications

Cloud computing is having an important impact on communication networks, both fixed and mobile, stimulating intensive research and standardization of new network architectures, protocols and resource management mechanisms. In the literature, significant attention has been devoted to system aspects of Cloud computing. More recently, however, the focus is shifting towards Cloud networking and communications with evolutionary and revolutionary propositions. The goal is to offer stronger interworking and interoperability between system and network elements.

With the proliferation of Cloud offerings, new networking and communication challenges have also emerged. For instance, Data-centers are becoming containers of virtual provider networks, whose embedding can be optimized as a function of cost, customer demands, energy, dynamic scaling, Quality-of-Service (QoS) and Quality-of-Experience (QoE). Users access to Clouds can be subjected to provider-level filtering and shaping, functions implementable in novel network middle-boxes based on Cloud Service Level Agreements (C-SLA). Cloud computing has an increasing impact on mobile access networks as well. In particular, network offloading protocols and advanced service migration and caching techniques are being leveraged to offload cellular provider networks and improve user QoE. Efficient resource management in data center and Cloud networks is an open research challenge that needs to be addressed in order to provide bandwidth guarantees and performance isolation. This is particularly important with the increasing reliance on bandwidth-demanding Virtual Machine (VM) migrations for resource consolidation and energy management within data centers and across geographically distributed data centers. In general, Cloud applications and their requirements are evolving so fast that new problems are faced by the telecommunication and Cloud providers every day asking for novel networking and communication architectures, protocols and resource management mechanisms.

This special issue presents recent research in the area of communications and networking in the Cloud. Over sixty papers were submitted for this special issue. After extensive review and discussion, twelve papers were selected for publication. The selected papers address four critical topics that play central role in Cloud communications and networking: Cloud network architectures and protocols, Cloud resource allocation and management, and Cloud application deployment and provisioning.

Three papers in the special issue address Cloud network architectures and protocols.

In "Cloud-integrated WOBAN: an offloading-enabled architecture for service-oriented access networks", Reaz et al. [1] propose the design of CIW; an access network that integrates a Cloud with a wireless-optical broadband access network (WOBAN). CIW is designed to create an infrastructure platform to provide different Cloud services from within the access network while improving resource utilization by offloading traffic from wireless links, and providing higher scalability by reducing bottlenecks from the gateways of the WOBAN. The authors propose a Mixed-Integer Linear Program (MILP) formulation of the optimal Cloud components placement problem in the WOBAN and solve it in a realistic case-scenario. They also propose a novel energy-saving routing mechanism, called Green Routing for CIW (GRC) that allows Cloud-Integrated WOBANs to self-manage the activation of network components in order to minimize the overall energy consumption. GRC performs load-balanced anycast routing across active devices. The performance evaluation of GRC shows that the routing mechanism achieves significant energy savings with low average packet delay.

In "Deadline and Incast aware TCP for Cloud data center networks", Hwang et al. [2] present DIATCP, a transport protocol for Cloud data center applications that is both deadline-aware and incast-avoidable. As opposed to prior works where deadline awareness is achieved through host-based or network-based approaches, DIATCP is deployed at an aggregator server which directly controls the peer servers' sending rate to avoid incast congestion and, more importantly, to meet the application deadline. This is under the key observation that the aggregator knows the bottleneck link status as well as its workers' information under the Partition/Aggregate traffic pattern.

http://dx.doi.org/10.1016/j.comnet.2014.05.012
1389-1286/

Through detailed ns-3 simulations and real testbed experiments, DIATCP is shown to significantly outperform state-of-the-art protocols in the Cloud data center environment.

In "A TRILL-based multi-tenant data center network", Amamaou et al. [3] present a multi-tenant network architecture for virtualized data center networks based on the TRILL protocol. The architecture leverages bridging and routing into a layer-2 network ensuring scalability, efficiency, fault-tolerance and simplified management. The architecture also integrates the VNT (Virtual Network over TRILL) solution, enabling the creation of large numbers of logical networks within a virtualized data center. Experimental evaluations of the system show that VNT is able to support several hundred switches, thousands of physical nodes, and hundreds of thousands of tenants with only low overhead.

The special issue features five papers that address Cloud resource allocation and management.

In "Resource management in IaaS cloud platforms made flexible through programmability", Wickboldt et al. [4] envision that Infrastructure as a Service (IaaS) Cloud administrators could benefit from customizations in resource management strategies to achieve environment specific objectives or to enable application oriented resource allocation. From this perspective, the authors introduce a new concept of Cloud management platform where resource management is made flexible by the addition of programmability to the core of the platform, with a simplified object-oriented API. A proof of concept prototype is described and an evaluation of three resource management programs on an emulated network using Linux virtualization containers and Open vSwitch running the OpenFlow protocol is presented. Results show the feasibility of the presented solution and how optimization programs can achieve different objectives defined by the administrator.

In "Provisioning high-availability datacenter networks for full bandwidth communication", Ni et al. [5] study the capacity allocation problem in datacenter networks where Valiant Load Balancing (VLB) is used to handle highly variable traffic. The authors target full bandwidth communication among all servers, for all valid traffic matrices, and under k arbitrary link failures, and focus on two typical datacenter topologies, VL2 and fat-tree. They derive and compare the minimum edge, respectively core, link capacity as well as the minimum total link capacity required on each topology and characterize the trend in capacity increase with k and with the total number of supported servers. They show that given the same server scale, fat-tree requires less total capacity than VL2 for small k, while for large k, there exists a turning point beyond which VL2 becomes more capacity-efficient than fat-tree. This finding is important for datacenter providers to project their capital expenditures in datacenter design, upgrade, and expansion.

In "Improving the performance of load balancing in software-defined networks through load variance-based synchronization", Guo et al. [6] propose Load Variance-based Synchronization (LVS), a new type of controller state synchronization scheme to improve the load-balancing performance in multi-controller multi-domain SDN networks. As opposed to Periodic Synchronization-based schemes, LVS conducts effective state synchronizations among controllers only when the load of a specific server or domain exceeds a certain threshold, which significantly reduces the synchronization overhead of controllers. Simulation results show that LVS achieves loop-free forwarding and good load-balancing performance with much less synchronization overhead, compared to existing synchronization schemes.

In "VDC-analyst: design and verification of virtual desktop cloud resource allocations", Calyam et al. [7] present VDCAnalyst, a tool that captures critical quality metrics such as Net Utility and Service Response Time, which can be used to quantify the readiness of Virtual Desktop Cloud platforms. VDCAnalyst is designed for Cloud service providers, researchers and educators to design and verify resource allocation schemes with simulations or emulations. The tool operates in two modes: "Run Simulation" and "Run Experiment", respectively. The Run Simulation mode allows users to test and visualize resource provisioning and placement schemes on a simulation framework. Run Experiment mode allows testing on a real Software-Defined Network testbed using emulated virtual desktop application traffic to create a realistic environment. Results from using VDCAnalyst show that a significant increase in perceived Quality of Experience could be achieved by using a combination of techniques integrated in the tool.

In "Multiple bulk data transfers scheduling among datacenters", Wang et al. [8] aim at reducing the network congestion due to bulk data transfer. The Multiple Bulk Data Transfers Scheduling (MBDTS) problem is modeled as a Linear Program (LP) problem after applying an Elastic Time-Expanded Network technique to represent the time-varying network status as a static one with a reasonable expansion cost. An algorithm that solves the LP iteratively is used to obtain the optimal solution to the MBDTS problem. Simulations conducted on a real network topology show that the algorithm can significantly reduce network congestion and balance the entire network traffic with practical computational costs.

The following four papers address Cloud application deployment and provisioning.

In "Optimal application allocation on multiple public clouds", Woo and Mirkovic [9] investigate the benefits of deploying components of a distributed application in a multi-Cloud environment. The authors propose a resource allocation algorithm that minimizes the overall Cloud operation cost, while satisfying required service-level agreements (SLAs). A simulation study that integrates real Cloud performance and cost data shows that, in spite of the additional delays for inter-Cloud communication and the additional costs for inter-Cloud data transfer, multi-Cloud allocation still outperforms single-Cloud allocations in a number of realistic scenarios.

In "Quality of experience in cloud services", Casas and Schatz [10] study the Quality of Experience (QoE) with several Cloud-based services, ranging from services with low requirements in terms of latency and interactivity (e.g., Cloud storage systems), multimedia On-Demand services (e.g., YouTube video streaming), communication and telepresence (e.g., Lync Online videoconferencing) to highly interactive services (e.g., Virtual Cloud Desktop). The services are evaluated according to an experimental assess-

ment methodology developed by the authors based on subjective lab and field tests for quality assessment, and the analysis of interplays and existing mappings between user, application, and network Quality of Service. This study is meant to provide developers and access network service providers a "ground truth basis" for developing Cloud services with QoE consideration and for a better dimensioning of the access network infrastructure, particularly with regard to mobile access technologies.

In "Application delivery in multi-cloud environments using software defined networking", Paul et al. [11] present OpenADN a novel service-centric network design that will allow ASPs to fully leverage multi-Cloud environments for deploying and delivering their applications over a shared, service-centric, wide-area network infrastructure provided by third-party providers including Internet Service Providers (ISPs), Cloud Service Providers (CSPs) and Content Delivery Networks (CDNs). The OpenADN design leverages the Software Defined Networking (SDN) framework to implement and manage the deployment of OpenADN-aware devices.

In "Internet factories: creating application-specific networks on-demand", Strijkers et al. [12] introduce the concept of Internet Factory: a software architecture that enables the on-demand creation and management of application-specific overlay networks using infrastructure-as-a-service Clouds. The authors describe the Internet Factory architecture and report on a proof of concept with three concrete examples including the creation of an IPv6 network over a number of Cloud locations around the world.

We express our thanks to the authors who submitted papers and to the reviewers for their thoughtful comments. It has been a pleasure to put together an issue on such a timely topic. We are grateful to the editor in chief, Harry Rudin, for giving us the opportunity to put together this special issue and for his support throughout the process.

References

[1] Abu Sayeem Reaz, Vishwanath Ramamurthy, Massimo Tornatore, Biswanath Mukherjee, Cloud-integrated WOBAN: an offloading-enabled architecture for service-oriented access networks, Comput. Netw. 68 (2014) 5–19.
[2] Jaehyun Hwang, Joon Yoo, Nakjung Choi, Deadline and Incast Aware TCP for Cloud data center networks, Comput. Netw. 68 (2014) 20–34.
[3] Ahmed Amamou, Kamel Haddadou, Guy Pujolle, A TRILL-based multi-tenant data center network, Comput. Netw. 68 (2014) 35–53.
[4] Juliano A. Wickboldt, Rafael P. Esteves, Márcio B. Carvalho, Lisandro Z. Granville, Resource management in IaaS cloud platforms made flexible through programmability, Comput. Netw. 68 (2014) 54–70.
[5] Wenda Ni, Changcheng Huang, Wu Jing, Provisioning high-availability datacenter networks for full bandwidth communication, Comput. Netw. 68 (2014) 71–94.
[6] Zehua Guo, Mu Su, Yang Xu, Zhemin Duan, Luo Wang, Shufeng Hui, H. Jonathan Chao, Improving the performance of load balancing in software-defined networks through load variance-based synchronization, Comput. Netw. 68 (2014) 95–109.
[7] Prasad Calyam, Sudharsan Rajagopalan, Sripriya Seetharam, Arunprasath Selvadhurai, Khaled Salah, Rajiv Ramnath, VDC-analyst: design and verification of virtual desktop cloud resource allocations, Comput. Netw. 68 (2014) 110–122.
[8] Yiwen Wang, Su Sen, Alex X. Liu, Zhongbao Zhang, Multiple bulk data transfers scheduling among datacenters, Comput. Netw. 68 (2014) 123–137.
[9] Simon S. Woo, Jelena Mirkovic, Optimal application allocation on multiple public clouds, Comput. Netw. 68 (2014) 138–148.
[10] Pedro. Casas, Raimund. Schatz, Quality of experience in cloud services, Comput. Netw. 68 (2014) 149–165.
[11] Subharthi Paul, Raj Jain, Mohammed Samaka, Jianli Pan, Application delivery in multi-cloud environments using software defined networking, Comput. Netw. 68 (2014) 166–186.
[12] Rudolf Strijkers, Marc X. Makkes, Cees de Laat, Robert Meijer, Internet factories: creating application-specific networks on-demand, Comput. Netw. 68 (2014) 187–198.

Guest Editors
Raouf Boutaba
University of Waterloo, Canada
E-mail address: rboutaba@uwaterloo.ca

Noura Limam
University of Waterloo, Canada
E-mail address: n2limam@uwaterloo.ca

Stefano Secci
University Pierre and Marie Curie, France
E-mail address: stefano.secci@upmc.fr

Tarik Taleb
NEC Europe, Germany
E-mail address: tarik.taleb@neclab.eu

Raouf Boutaba received the M.Sc. and Ph.D. degrees in computer science from the University Pierre & Marie Curie, Paris, in 1990 and 1994, respectively. He is currently a professor of computer science at the University of Waterloo. His research interests include resource and service management in networks and distributed systems. He is the founding editor in chief of the IEEE Transactions on Network and Service Management (2007–2010) and on the editorial boards of other journals. He has received several best paper awards and other recognitions such as the Premiers Research Excellence Award, the IEEE Hal Sobol, the Fred W. Ellersick, the Joe LociCero, the Dan Stokesbury, and the Salah Aidarous Awards. He is a fellow of the IEEE and the Engineering Institute of Canada.

Noura Limam received the M.Sc. and Ph.D. degrees in computer science from the University Pierre & Marie Curie, Paris VI, in 2002 and 2007, respectively. She is currently a research assistant professor of computer science at the University of Waterloo. Formerly, she was a research fellow at POSTECH, South Korea and researcher at Ucopia Communications Inc., France. She received the IEEE Communications Society Fred W. Ellersick Prize in 2008. Her contributions are in the area of service management and service engineering. Her current research interests are in distributed sensing in large settings and social networking.

Stefano Secci is an Associate Professor at the University Pierre and Marie Curie (UPMC – Paris VI), since 2010. He received the M.Sc. degree in communications engineering from Politecnico di Milano, in 2005, and a dual Ph.D. degree in computer science and networks from the same school and Telecom ParisTech, in 2009. He also worked as a Research Fellow at NTNU, George Mason University, Ecole Polytechnique de Montréal, and Politecnico di Milano, and as a Network Engineer with Fastweb Italia. Dr. Secci is an associate editor for IEEE Communications Surveys and Tutorials since 2012, for Journal of Network and Systems Management (Springer) since 2013. He is Vice-Chair of the Internet Technical Committee (ITC), joint between the IEEE Communication Society and the Internet Society (ISOC), since 2013. His works mostly cover network optimization, protocol design, Internet routing and traffic engineering. His current research interests are about Internet resiliency and Cloud networking.

Tarik Taleb is an IEEE ComSoc Distinguished Lecturer. He is currently working as Senior Researcher and 3GPP Standardization Expert at NEC Europe. Prior to his current position, he worked as assistant professor at Tohoku University. He received his B.E. degree in Information Engineering with distinction, MSc and Ph.D degrees in Information Sciences from Tohoku Univ., in 2001, 2003, and 2005, respectively. His research interests lie in the field of architectural enhancements to mobile core networks, mobile cloud networking, mobile multimedia streaming, and inter-vehicular communications. Dr. Taleb has been also directly engaged in the development and standardization of the Evolved Packet. Dr. Taleb is/was on the editorial board of many IEEE journals and magazines including the IEEE Wireless Communications Magazine and IEEE Transactions on Vehicular Technology. He is serving as vice-chair of the Wireless Communications Technical Committee. Dr. Taleb is the recipient of many awards including the 2009 IEEE ComSoc Asia-Pacific Best Young Researcher award and the 2008 TELECOM System Technology Award. Some of Dr. Taleb's research work has been also awarded best paper awards at prestigious conferences.