Available online at www.sciencedirect.com
ScienceDirect
Procedia Computer Science 56 (2015) 276 – 283
Procedia Computer Science

The 10th International Conference on Future Networks and Communications (FNC 2015)

Distributed control in virtualized networks

L. Zuccaro[a], F. Cimorelli[a], F. Delli Priscoli[a], C. Gori Giorgi[a], S. Monaco[a], V. Suraci[b]*

[a]*University of Rome "Sapienza", Via Ariosto 25, Rome, 00185, Italy*
[b]*Università degli studi e-Campus, Via Isimbardi 10, Novedrate, 22060, Italy*

Abstract

The increasing number of the Internet connected devices requires novel solutions to control the next generation network resources. The cooperation between the Software Defined Network (SDN) and the Network Function Virtualization (NFV) seems to be a promising technology paradigm. The bottleneck of current SDN/NFV implementations is the use of a centralized controller. In this paper, different scenarios to identify the pro and cons of a distributed control-plane were investigated. We implemented a prototypal framework to benchmark different centralized and distributed approaches. The test results have been critically analyzed and related considerations and recommendations have been reported. The outcome of our research influenced the control plane design of the following European R&D projects: PLATINO, FI-WARE and T-NOVA.

Peer-review under responsibility of the Conference Program Chairs

Keywords: Distributed control, Network Function Virtualization, Software Defined Network, Future Internet, 5G networks.

1. Introduction

In the last decade, the economic crisis moved from the virtual world of the high finance, to the real world of millions of European entrepreneurs and workers. Internet is the backbone of the European economy recovery and it makes event more urgent to develop a faster, more secure and reliable infrastructure. On top of Internet can grow novel research areas and related business opportunities, such as Internet of Things (IoT), Big Data and Future Internet. These technologies transformed the Internet into a virtual world in which anyone and anything can

* Corresponding author. Tel.: +39 06 77274043.
E-mail address: vincenzo.suraci@uniecampus.it

Peer-review under responsibility of the Conference Program Chairs
doi:10.1016/j.procs.2015.07.209

exchange, consume and provide resources, services, applications, data, information or knowledge no matter what device, location, context, situation or communication technology they have[1,2,3]. The heart of the Internet is its network infrastructure: on top of it, all the ICT novelties have flourished. The 5G networks represent the beyond the state of the art Internet infrastructures. 5G networks will face an unprecedented an challenging set of requirements, the management, control and supervision of such a complex, heterogeneous, networked system requires highly scalable approaches and cost effective solutions.

A first attempt to face this challenge has been done by the 5G Public Private Partnership (5G-PPP) and Future Internet Public Private Partnership (FI-PPP) initiatives, supported by other related projects such as PLATINO[4] and FI-CORE[5]. In particular, these projects aim at defining the Future Internet as a platform that provides a rich library of Generic Enablers supporting the easy and seamless development of innovative applications.

To enable the Future Internet paradigm, a seamless access to control and manage the underlying technologies is crucial. The virtualization of the underlying technologies will allow the Future Internet to boost its performance and to have a unique access to the available resources[6]. In this respect, a promising technology[7] is the combined use of Network Function Virtualization (NFV) and Software Defined Network (SDN)[8,9]. In this paper, our experience gained in the most recent European research projects (T-NOVA[10], FI-CORE[5]) envisages that a crucial role to the success of the 5G-PPP and FI-PPP visions will be the adoption of scalable, high-performance NFV control-plane. In section 2, the evolution of the current Internet architecture is shown and the key role played by NFV is highlighted. In section 3, the role of centralized SDN controllers to support NFV is analyzed and a distributed and elastic solution is proposed to cope with the requirements of the next generation Internet, with a particular focus on reducing the latencies introduced by the control-plane. In section 4, a prototypal proof-of-concept has been developed to demonstrate the validity of the proposed approach analyzing the results of our stress tests.

2. Evolution of the Internet architecture

The traditional approaches to model the Internet architecture (e.g. layered or hierarchical) are all somehow limited, since they intrinsically tend to organize the Internet infrastructure into rigid schemes. In our vision, we want to abstract from traditional architectural models, trying to be more agnostic as possible. From a functional point of view, the aforementioned Future Internet platform is in charge of meeting two different entities, the *actors* and the *resources*, by means of dedicated *applications*. An actor represents the entity whom requirement fulfilment is the main goal of the Future Internet. For example, an actor could be an individual user, a content prosumer, an app developer, a network operator, a service provider or a cloud owner. A resource represents any entity that can be exploited to satisfy the actors' needs. Examples of resources include, but are not limited to, services, contents, terminals, devices, functionalities, storage, computation, connectivity or networking capabilities. An application is any means used by the actors to exploit the available resources with the aim of fulfilling their requirements. Social networking, context-aware services, on-line games, interactive multimedia services, cloud storage and processing, collaborative services or automation services are examples of applications. The Internet evolution will promote those solutions where applications transparently, efficiently and flexibly exploit the available resources while satisfying the expectations of the involved actors.

Nowadays Internet architecture has several limitations that slow down this evolution process. A first limitation is due to the heterogeneity of the access and transport networks, developed according to different standard, scopes, services and customers led to a coexistence of technology-dependent, proprietary solutions. In this respect, the requirement of virtualizing the resources, to ease their management, requires the introduction of a convergent-layer between the resources and the applications. A valid solution to implement a convergence-layer is to use a virtualization framework. In the specific case of network and cloud resources, the join use of SDN and NFV represents the best candidate to implement the virtualization layer.

Another limitation is the traditional layering architecture of the current Internet. It forces to maintain independent one another those algorithms and procedures that operate at different layers. In addition, even in the framework of the same layer, algorithms and procedures dealing with different tasks are often designed independently one another. The short-term advantage of a layered approach is to simplify the overall design of a growing telecommunication network and to reduce the processing capabilities, since the overall network control problem is decoupled in a number of simpler sub-problems. Nevertheless, the long-term limitation of this approach derives from the fact that algorithms and procedures are poorly coordinated, impairing the efficiency of the overall network control. Solving this issue requires more cooperation between the involved control-planes to apply for a distributed control of the underlying telecommunication network (TLC and IT resources)[11,12,13].

Another important limitation is related to the fact that most of the algorithms and procedures embedded in the telecommunication networks are open-loop. They are setup and configured on the base of off-line, reasonable estimations of network variables, rather than on real-time measurements of such variables. This static approach does not fit well with the dynamicity of a 5G network. This claims for an evolution towards closed-loop algorithms and procedures, which are able to react properly on the base of appropriate real-time feedback information.

In order to cope with these limitations, the Internet architecture developed in FI-CORE[5] and PLATINO[4] research projects, evolved to clearly define how actors can exploit TLC and IT resources through proper applications. In particular the Internet architectural concept, shown in Fig. 1, is based on the following functionalities:

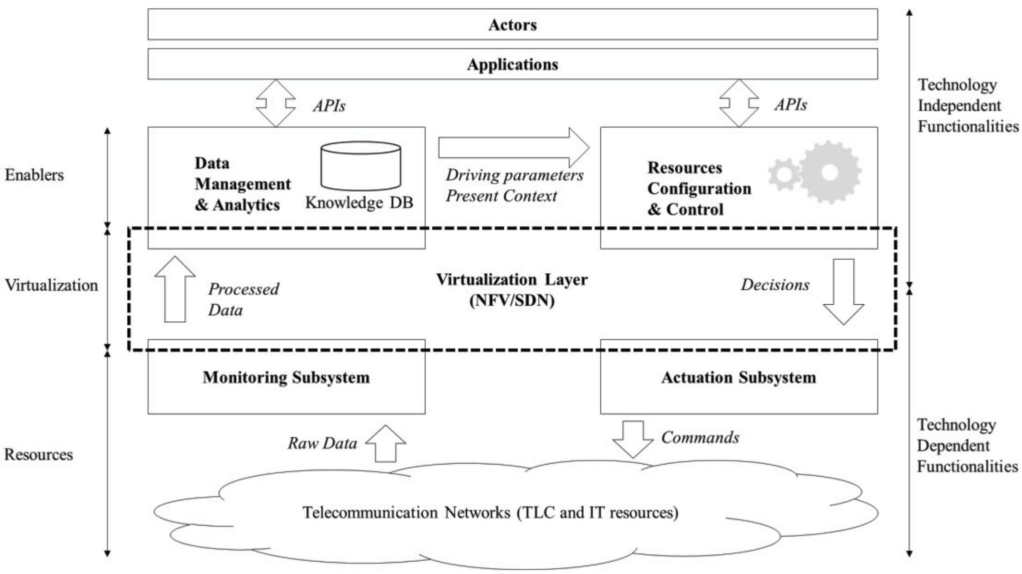

Fig. 1. Internet architectural concept

- A *Monitoring Subsystem* in charge of gathering and pre-processing heterogeneous, technology-dependent, multi-layer and multi-network measurements and information coming from the TLC and IT resources (*Raw Data*), with the aim of producing *Processed Data*, which can be useful for analysis and control purposes.
- A *Virtualization Layer* in charge of assuring the decoupling between the technology dependent functionalities offered by the underlying resources and the technology independent functionalities offered to the overlying applications. The virtualization layer applies the convergence among heterogeneous technologies in different domains: information, communications, energy, health, gaming, office automation, transport, agri-food, etc. Concrete examples of virtualization layer instances are, among others: cloud storage, virtual servers, network function virtualization. The virtualization layer hosts technology-independent, interoperating *Enablers* which offers appropriate *APIs*[7] to the overlying applications. Such Enablers can be grouped into two sets: Data Analytics and Resource Configuration & Control.
- An *Actuation Subsystem* in charge of putting in place, on the underlying resources, the Decisions taken by the virtualization layer enablers. Such task is performed by means of proper technology-dependent Commands.
- The *Data Management & Analytics* Enablers are in charge of: (i) the collection and the preliminary filtering of the heterogeneous inter-layer/inter-network data produced by the Monitoring Subsystem; (ii) the abstraction (e.g. by using appropriate ontological description languages) of such heterogeneous data in order to obtain

homogeneous metadata; (iii) the appropriate aggregation and semantic enrichment of these metadata. The above-mentioned operations are dynamically performed and eventually produce the so-called *Present Context*. The Present Context continuously feeds a *Knowledge Database* which hence stores both the present and the past aggregated and semantically enriched metadata; (iv) inferring (for instance by adopting machine learning and pattern recognition techniques) from the data stored in the Knowledge Database and from the interaction with the applications and the actors through proper APIs, the so-called *Driving Parameters*. These parameters drive the Resource Configuration & Control Enablers.
- The *Resources Configuration and Control* Enablers which, exploiting the virtualization layer functionalities and on the basis of the Present Context and of the Driving Parameters, have to dynamically produce control and configuration Decisions aiming at the satisfaction of the personalized requirements and needs of the Actors while using given Applications.

The proposed Internet architectural concept can be easily mapped into different domains. In this paper we focus on the control of 5G network resources. In this context, the technology dependent network resources are virtualized through the use of a SDN paradigm. SDN offers a feasible solution to virtualize the basic, per-flow, monitor and control network functionalities. To offer to the overlying applications more complex network functionalities for the data management & analytics and for the resource configuration & control (e.g. routing[14], load balancing[15], admission control[16], bandwidth on demand[17,18], resource allocation[19,20], quality of experience[21], security[22], etc.), a NFV approach is proposed. The conjunct use of SDN and NFV allows to overcome all the aforementioned current Internet limitations. In our work, we studied the impact on the latency of 5G networks when adopting a centralized or a distributed implementation of the SDN/NFV control-plane.

3. Elastic SDN Control Plane

The heterogeneity of technologies, algorithms, procedures and hardware constitute one of the major limitations of the existing network infrastructure. As highlighted in section 2, the virtualization layer represents a feasible solution to seamlessly manage network and IT resources in a technology independent manner[6]. In this context, NFV enhances the design of future network infrastructures by translating network hardware appliances in software solutions[9], dynamically deployed on one or more VMs, connected and chained to create network services.
The SDN paradigm looks highly complementary to NFV having the potential to provide a scalable, elastic and on-demand network infrastructure[8]. Additionally, the centralized view of the network allows the SDN network controller to make optimal forwarding decisions. However, the controller may be subject to overload or failure issues and increasing the computational and memory capacity may not be enough. These issues have an impact on the control plane reactiveness, and consequently degrade the overall network latency. It became more evident when the network size grows, thus a way to overcome these limitations is needed to make SDN/NFV a pillar technology of 5G networks.

In this regard, the concept of a "distributed, but logically centralized" controller has been investigated to develop an instance of the virtualization layer applying for a distributed control of the network elements. The proposed SDN/NFV Control Plane (CP) has been developed firstly within the PLATINO project and it has been enhanced in the T-NOVA project[10]. Our solution is based on the virtualization of the network controller through multiple controller instances organized in *cluster*, while keeping the benefits of centralized network control by means of a distributed data store. The key concept is to deploy each instance of SDN controller on dedicated virtual machines, favoring the distribution of the network control workload across the cluster. In this way, the controller virtualization may help in overcoming scalability and centralization issues, which affect the SDN controller performances in large data center hosting NFV applications.
The high-level architecture of the SDN Control Plane, designed to support deployments in a large-scale scenario, introduces the following functional components:

- *Distributed Data Store:* it is responsible for consistently maintaining a global view (topology and the state of the network) across the control plane instances belonging to the cluster. Northbound applications/internal CP components can take advantage of the global network view in making forwarding and policy decisions

- *Northbound Request Handler:* it is mainly responsible for spreading the northbound requests among the available controller instances, it is essential to make the network control plane accessible through the northbound API as a unique single instance.
- *CP Coordinator:* it supervises the operation in the cluster. Specifically it has to dynamically configure the controller-to-switch connections; decide whether to add or remove a controller instance to the cluster depending on the network needs. This role is played by one of the instances available in the cluster, by means of a procedure of leader election.
- *CP Agent:* it collects information about the resource utilization (CPU load, memory usage, control messages arrival rate, etc.) at each CP instance and enforces the switch-to-controller instance connection rules used by each switch to identify the controller instance/s to which the southbound requests must be forwarded.

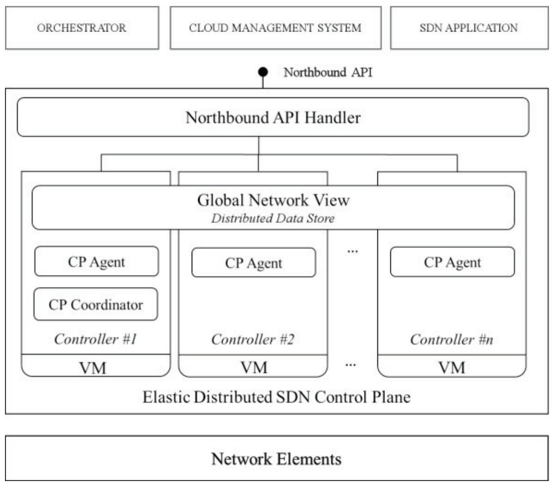

Fig. 2. SDN Control Plane Architecture

4. Experimental results

In order to evaluate the benefits of having multiple controller instances, a basic test environment was setup using Mininet[23] and OpenDaylight[24], the most popular network emulator and controller in the SDN context.

The testing environment in Fig. 3 is composed by a SDN controller cluster of 2 instances and 26 virtual switches created using OpenvSwitch[25]. Every switch Si is connected to the Si+1 switch and to the host Hi that simulates the data traffic generator. The control traffic is generated by the switch when there is no flow entry for an incoming data packet that the host wants to send. Hence, the switch encapsulates the data packet in a control packet (packet-in) and sends it to its controllers. Then, exactly one controller should send a message containing the flow entry that the switch will install in its table. Let flow-mod be the name of these messages.

We performed the test using the following physical devices: n.1 Quad core Intel Q8300 machine with 4 GB RAM and a gigabit Ethernet card hosting the Mininet network emulator and the tcpdump packet sniffer; n. 2 Quad core AMD Athlon X4 750K with 8 GB RAM and a gigabit Ethernet card, both running the OpenDayight controller; n. 1 gigabit switch. The performance is measured in the machine with Mininet as all packet-in and flow-mod messages pass through its network interface (NIC in figure 3).

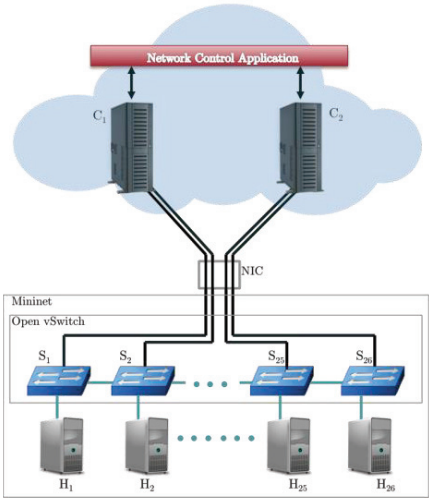

Fig. 3. The testing environment

In Fig. 4 the testing scenarios are presented. In the configuration *(a)* all the switches are connected to one instance of controller.

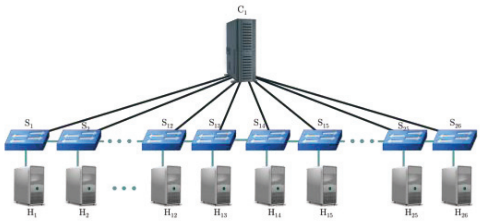

a. Single connection (One Controller)

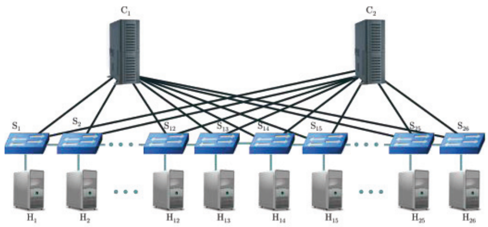

b. All connections (Two controllers)

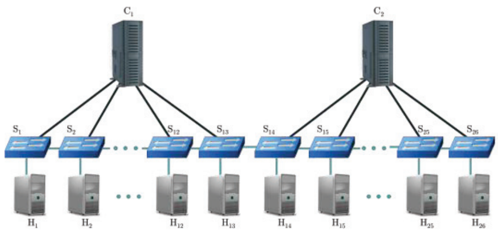

c. Selective connections (Two controllers)

Fig. 4. Test scenarios

The scenario *(b)* is interesting to see what happens when every switch sends the *packet-in* control message to more than one controllers. One controller should answer to the request while the other one should ignore it. However, the ignoring decision takes some computational resources and l see the its impact. In *(c)* there are two controllers and the first 13 switches are connected to the first one while the remaining ones are connected to the second controller. Figure 5 depicts the 95-percentile response time *(a)* and the throughput *(b)* of the 3 cases described above.

One can notice that the difference between the *All* and the *Selective* connection case becomes relevant when the number of packet-in/s is greater than 40,000. At 50,000 this difference becomes more significant (about 8 seconds).

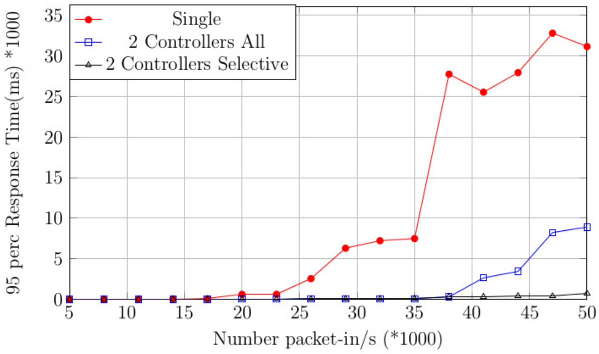

a. *Single connection (One Controller)*

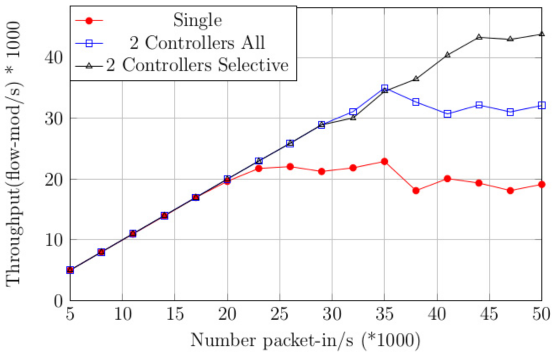

b. *All connections (Two controllers)*

Fig. 5. 95-percentile response time (*a*) and throughput (*b*) comparison between the *Single* (red), *All* (blue) and *Selective* (black) test configurations

Similar results are reported in Figure 5b. When the network load is high the throughput of the Selective test is the highest one, being 66% greater than the one in the All test scenario and 100% greater than the one in the *Single* test scenario. These plots clearly highlight the benefits of a control plane with multiple controllers. It is easy to observe that relevant performance boost is obtained only under certain level of load. Hence it is important to find the load levels that should trigger an increment or a decrease of the number of controllers. Moreover, shrinking the cluster when the load drops under some decided value may reduce the operational cost.

5. Conclusions

In this paper, the future Internet trends from the European research projects perspective was studied. Once identified current Internet limitations, an architectural solution exploiting SDN/NFV was proposed. Starting from the current SDN/NFV centralized solution, an evolution based on a distributed and elastic control plane that can easily scale was described. We developed a proof of concept demonstrator to perform stress tests. After these tests, the *Selective* connection proved to be better than any other configuration. These findings open the way to future research streams. In particular, the authors intend to perform a deep study of pro/cons between the proposed configurations to select the best controlling strategy on the base of varying contextual situations, such as the network topology, the network congestion or the available virtualization resources. As a matter of fact a significant attention is on the application of controlled cloud and virtualization functions to the Smartgrid sector[26], for massive deployment of control strategies involving grid equipment and distributed energy resources[27,28].

Acknowledgments

The research leading to these results has received funding from the Italian project PLATINO, under Grant Agreement n° PON01_01007, the EU project FI-WARE/FI-CORE, under Grant Agreement n° 285248, the EU project T-NOVA, under Grant Agreement n° 619520. The authors wish to thank Eng. Daniel Gheorghita for setting up the testbed to run the presented experiments.

References

1. M. Castrucci, F. Delli Priscoli, A. Pietrabissa, V. Suraci (2011). A Cognitive Future Internet Architecture. In: Domingue, J.; Galis, A.; Gavras, A.; Zahariadis, T.; Lambert, D.; Cleary, F.; Daras, P.; Krco, S.; Müller, H.; Li, M.-S.; Schaffers, H.; Lotz, V.; Alvarez, F.; Stiller, B.; Karnouskos, S.; Avessta, S.; Nilsson, M. (Eds.). The Future Internet. p. 91-102, SPRINGER, doi: 10.1007/978-3-642-20898-0_7.
2. M. Castrucci, M. Cecchi, F. Delli Priscoli, L. Fogliati, P. Garino, V. Suraci (2011). Key Concepts for the Future Internet Architecture. In: Paul Cunningham and Miriam Cunningham (Eds). Future Network and MobileSummit 2011 Conference Proceedings. Warsaw, Poland, 15 - 17 June 2011, IIMC International Information Management Corporation, ISBN: 9781905824236.

3. Jing Ren ; Kejie Lu ; Sheng Wang ; Xiong Wang ; Shizhong Xu ; Lemin Li ; Shucheng Liu, "VICN: a versatile deployment framework for information-centric networks". Network, IEEE Volume: 28, Issue: 3. DOI: 10.1109/MNET.2014.6843229. 2014, Page(s): 26 - 34.
4. "Platform for Innovative Services in Future Internet", Italian Ministry of University and Research (MIUR) PLATINO project, (Grant Agreement n° PON01_01007), http://www.progettoplatino.it/.
5. FI-CORE http://www.fi-ware.org/
6. S. Jyotiprakash, S. Mohapatra, and R. Lath. "Virtualization: A survey on concepts, taxonomy and associated security issues." Computer and Network Technology (ICCNT), 2010 Second International Conference on. IEEE, 2010.
7. A. Palo, L. Zuccaro, A. Simeoni, V. Suraci, L. Musto, P. Garino. "A Common Open Interface to Programmatically Control and Supervise Open Networks in the Future Internet". Future Network & Mobile Summit 2013, July 2013
8. J. Raj, and S. Paul. "Network virtualization and software defined networking for cloud computing: a survey." Communications Magazine, IEEE51.11 (2013): 24-31.
9. ETSI ISG NFV group. GS NFV 002. Network Functions Virtualisation (NFV); Architectural Framework. Available online at http://www.etsi.org/technologies-clusters/technologies/nfv.
10. T-NOVA www.t-nova.eu
11. T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker. "Onix: A distributed control platform for large-scale production networks". In OSDI, 2010
12. Advait Abhay Dixit, Fang Hao, Sarit Mukherjee, T. V. Lakshman, and Ramana Rao Kompella. Towards an elastic distributed SDN controller. In HotSDN, 2013.
13. P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, and G. Parulkar, "ONOS: towards an open, distributed SDN OS," in Proceedings of the third workshop on Hot topics in software defined networking, Chicago, Illinois, USA, 2014.
14. C. Bruni, F. Delli Priscoli, G. Koch, A. Pietrabissa, L. Pimpinella, "Network decomposition and multi-path routing optimal control", Transactions on Emerging Telecommunications Technologies, Vol. 24, Issue: 2, March 2013, pp. 154-165, doi: 10.1002/ett.2536
15. Oddi G., Pietrabissa A., Delli Priscoli F., Suraci V., "A decentralized load balancing algorithm for heterogeneous wireless access networks". WTC 2014; World Telecommunications Congress 2014; Proceedings of, 1-6
16. A. Pietrabissa, "A Reinforcement Learning Approach to Call Admission and Call Dropping Control in Links with Variable Capacity", European Journal of Control, Vol. 17, Issue 1, 2011, pp. 89-103, ISSN 0974-3580, DOI: 10.3166/EJC.17.89-103
17. F. Delli Priscoli, A. Pietrabissa, "Load-Adaptive Bandwidth-on-Demand Protocol for Satellite Networks", Proc. of the Conference on Decision and Control (CDC), Las Vegas (U.S.A.), December 2002, pp. 4066-4071.
18. F. Delli Priscoli, A. Pietrabissa, "Control-theoretic Bandwidth-on-Demand Protocol for Satellite Networks", Proc. of Conference on Control Applications (CCA), Glasgow (UK), September 2002, pp. 530-535.
19. Oddi G., Panfili M., Pietrabissa A., Zuccaro, L., Suraci V., "A resource allocation algorithm of multi-cloud resources based on Markov Decision Process" Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, Volume 1, 2013, Article number 6753787, Pages 130-135, DOI: 10.1109/CloudCom.2013.24
20. F. Delli Priscoli, A. Pietrabissa, "Resource Management for ATM-Based Geostationary Satellite Networks with On-board Processing", Computer Networks - Special Issue on Broadband Satellite Systems: A Networking Perspective, Vol. 39, Issue 1, pp. 43-60, May 15, 2002, DOI: 10.1016/S1389-1286(01)00297-3
21. F. Delli Priscoli, A. Pietrabissa, V. Suraci, M. Iannone (2012). Modelling Quality of Experience in Future Internet Networks. In: Proceedings of Future Network & Mobile Summit. p. 1-9, ISBN: 9781467303200, Berlin, 4 - 6 July 2012.
22. P. Capodieci, R. Setola, F. Delli Priscoli, M. Castrucci, V. Suraci et alii, "Improving Resilience of Interdependent Critical Infrasructures via on-line Alerting System", Complexity in Engineering Conf. (Compeng 2010), Rome, December 2010, pp. 88-90.
23. Lantz, Bob, Brandon Heller, and Nick McKeown. "A network in a laptop: rapid prototyping for software-defined networks." Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks. ACM, 2010.
24. OpenDaylight, "OpenDaylight: A Linux Foundation Collaborative Project," 2013. [Online]. Available: http://www.opendaylight.org
25. Pfaff, Ben, et al. "Extending Networking into the Virtualization Layer." Hotnets. 2009.
26. Lanna A., Liberati F., Zuccaro L., Di Giorgio A., "Electric Vehicles Charging Control based on Future Internet Generic Enablers", IEEE International Electric Vehicle Conference IEVC14. IEEE, Firenze, I, December 17 - 19, pp. 1-5, 2014. DOI: 10.1109/IEVC.2014.7056079.
27. Mercurio A., Di Giorgio A., Cioci P., "Open Source Implementation of Monitoring and Controlling Services for EMS/SCADA Systems by Means of Web Services", IEEE Transactions on Power Delivery, Vol. 24, Issue 3, July 2009, pp.1148-1153. DOI: 10.1109/TPWRD.2008.2008461.
28. Di Giorgio A., Liberati F., Canale S., "IEC 61851 compliant electric vehicle charging control in Smartgrids", 21st Mediterranean Conference on Control and Automation MED12, pp. 1329-1335, MED13. IEEE, Chania, GR, June 25 - 28, 2013. DOI: 10.1109/MED.2013.6608892.