DIPARTIMENTO DI INGEGNERIA
CORSO DI DOTTORATO IN INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE -
PHD COURSE IN INDUSTRIAL AND INFORMATION ENGINEERING -
34TH CYCLE

Title of the research activity:	**Network function virtualization (NFV) through programmable networking techniques**
State of the Art:	The ongoing refactoring of traditional network function in their virtualized software-based counterparts holds the promise to dramatically enhance service delivery and deployment agility. Indeed, business cycles shrink and, with network function virtualization (NFV) technologies and relevant management and orchestration tools fostered by this proposal, network stakeholders will be able to move quicker than ever, change offerings, promptly add new services when their customers face a need, and get better insight, consistency, troubleshooting and visibility into the network status. Still, NFV's flexibility comes along with an obvious toll: the gap between the speed attainable in software opposed to dedicated HW devices is still very significant, and is not going to decrease in the future. Thus, forthcoming approaches will foster the migration from traditional virtual machine(VM)-based solution to more efficient and scalable virtualization technologies. VM-based virtual network fucntion (VNF) solutions, while traditionally viewed as heavy in terms of provisioning time and runtime overhead, are considered more secure than container-based approaches, that have been shown to be easier to attack. However, leveraging the better VNF agility, especially for on-demand services, that containers offer since they run the VNF directly in the host environment, it can be still interesting to investigate container-based VNF implementation through higher levels of abstractions. At the same time, other solutions involving Unikernels, i.e., tiny virtual machines with a minimal operating system that have shorter booting times than VMs or containers, will be considered to develop a lightweight NFV framework. In addition to the above, software-based, technological advances, new research avenues are enabled by recent results on data plane programmability, which are now sufficiently mature for being tested in the field. Indeed, a disruptive innovation is being currently happening in terms of models, platforms and systems for supporting programmable data planes. We can see this novel direction as a way to overcome both the limitations of current software defined networking (SDN)/OpenFlow-based frameworks. This will bring programmability down to the data plane by extending the elementary "match/action" abstraction provided by OpenFlow. In addition, this will allow overcome the performance limitations of NFV by including hardware acceleration devices tailored to the specific needs of network processing and computing. Worldwide initiatives in this area are getting consensus. The most interesting initiative are the P4 proposal for a data plane programming language, or the recently considered OpenFlow extensions to permit stateful flow processing at line rate.
Short description and objectives of the research activity:	In this research area, two interesting research topics are (i) more efficient and scalable virtualization technologies, and (ii) acceleration and offloading of (performance critical) virtualized components into dedicated network processing devices, specifically smart NICs and programmable switches. The idea is of this research activity is to go beyond the VNFs implementation with state of the art acceleration techniques like DPDK or SR-IOV. The aim is designing packet manipulation processors that are not only able to provide complex

	packet modifications at line rate, but also programmable by means of high level Domain Specific Languages, like P4, hence dramatically improving VNFs deployment time and agility. More specifically, the activity will focus on leveraging programmable smart NICs to offload the performance-critical parts of VNFs, i.e. those (such as traffic control algorithms or packet-level manipulation/processing) which are expected to run at multi-gigabit line rates and which can hardly run on commodity ARM/x86 CPUs, since they require the development of highly specialized processing architectures. It is also requested an effort to identify technologies and techniques which permit to exploit HW-based performance acceleration of VNF without the need to directly implement them in hardware, hence improving deployment time and agility. In order to simplify the access to this high performance networking framework, it is fundamental the use of abstraction layers, which allow masking such technological differences between VNF deployments and their mapping towards different service requirements. This can be done in an innovative way by leveraging the novel concept of intent-based networking to manage virtualized network slices. The main novelty of this paradigm is the idea of a network administrator defining a desired state of the network and having automated network orchestration software implement those policies. This concept can be pushed towards customers, allowing them to express in a simple way "what" their virtualized service slice should provide, leaving to the orchestration engine the task to decide "how". Clearly, in order to implement this paradigm in a successful way, a number of innovations are required inside the network, in particular in the orchestration, interworking, and monitoring functions. With respect to traditional approaches, intent-based networking guarantees a higher degree of automation and abstraction, while it simplifies the management of heterogeneous resources (i.e. NFV implementation options) when dealing with end-to-end slices. This also guarantees a simpler interface towards vertical industries, which is a basic requirement to attract them towards virtualized services, such as those offered by forthcoming 5G networks.
Bibliography:	[1] A. Rostami et al., "Orchestration of RAN and Transport Networks for 5G: An SDN Approach", IEEE Communications Magazine, April 2017. [2] P. Neves et al., "Future mode of operations for 5G – The SELFNET approach enabled by SDN/NFV", Computer Standards & Interfaces 54 (2017) 229–246. [3] 5G PPP Project Slicenet, https://slicenet.eu/. [4] Bosshart, McKeown, et al. "P4: Programming protocol-independent packet processors", SIGCOMM CCR 2014 [5] H. Farhady et al., "Data Plane Programmability in SDN", IEEE ICNP 2014. [6] S. Pontarelli et al., "Smashing SDN "built-in" actions: Programmable data plane packet manipulation in hardware", IEEE NetSoft 2017.
Scientific coordinator (s)	Prof. Gianluca Reali, Dr. Mauro Femminella
Contact (s)	Email:gianluca.reali@unipg.it , skype: gianluca.reali