

Network Function Virtualization: State-of-the-art and Research Challenges

Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Niels Bouten, Filip De Turck, Raouf Boutaba

Abstract—Network Function Virtualization (NFV) has drawn significant attention from both industry and academia as an important shift in telecommunication service provisioning. By decoupling Network Functions (NFs) from the physical devices on which they run, NFV has the potential to lead to significant reductions in Operating Expenses (OPEX) and Capital Expenses (CAPEX) and facilitate the deployment of new services with increased agility and faster time-to-value. The NFV paradigm is still in its infancy and there is a large spectrum of opportunities for the research community to develop new architectures, systems and applications, and to evaluate alternatives and trade-offs in developing technologies for its successful deployment. In this paper, after discussing NFV and its relationship with complementary fields of Software Defined Networking (SDN) and cloud computing, we survey the state-of-the-art in NFV, and identify promising research directions in this area. We also overview key NFV projects, standardization efforts, early implementations, use cases and commercial products.

Index Terms—Network function virtualization, virtual network functions, future Internet, software defined networking, cloud computing.

I. INTRODUCTION

Service provision within the telecommunications industry has traditionally been based on network operators deploying physical proprietary devices and equipment for each function that is part of a given service. In addition, service components have strict chaining and/or ordering that must be reflected in the network topology and in the localization of service elements. These, coupled with requirements for high quality, stability and stringent protocol adherence, have led to long product cycles, very low service agility and heavy dependence on specialized hardware.

However, the requirements by users for more diverse and new (short-lived) services with high data rates continue to increase. Therefore, Telecommunication Service Providers (TSPs) must correspondingly and continuously purchase, store and operate new physical equipment. This does not only require high and rapidly changing skills for technicians operating and managing this equipment, but also requires dense deployments of network equipment such as base stations. All these lead to high CAPEX and OPEX for TSPs [1], [2].

Moreover, even with these high customer demands, the resulting increase in capital and operational costs cannot be translated in higher subscription fees, since TSPs have learned

that due to the high competition, both among themselves and from services being provided over-the-top on their data channels, increasing prices only leads to customer churn. Therefore, TSPs have been forced to find ways of building more dynamic and service-aware networks with the objective of reducing product cycles, operating & capital expenses and improving service agility.

NFV [3], [4] has been proposed as a way to address these challenges by leveraging virtualization technology to offer a new way to design, deploy and manage networking services. The main idea of NFV is the decoupling of physical network equipment from the functions that run on them. This means that a network function - such as a firewall - can be dispatched to a TSP as an instance of plain software. This allows for the consolidation of many network equipment types onto high volume servers, switches and storage, which could be located in data centers, distributed network nodes and at end user premises. This way, a given service can be decomposed into a set of Virtual Network Functions (VNFs), which could then be implemented in software running on one or more industry standard physical servers. The VNFs may then be relocated and instantiated at different network locations (e.g., aimed at introduction of a service targeting customers in a given geographical location) without necessarily requiring the purchase and installation of new hardware.

NFV promises TSPs with more flexibility to further open up their network capabilities and services to users and other services, and the ability to deploy or support new network services faster and cheaper so as to realize better service agility. To achieve these benefits, NFV paves the way to a number of differences in the way network service provisioning is realized in comparison to current practice. In summary, these differences are as follows [5]:

Decoupling software from hardware. As the network element is no longer a composition of integrated hardware and software entities, the evolution of both are independent of each other. This allows separate development timelines and maintenance for software and hardware.

Flexible network function deployment. The detachment of software from hardware helps reassign and share the infrastructure resources, thus together, hardware and software, can perform different functions at various times. This helps network operators deploy new network services faster over the same physical platform. Therefore, components can be instantiated at any NFV-enabled device in the network and their connections can be set up in a flexible way.

Dynamic scaling. The decoupling of the functionality of the network function into instantiable software components pro-

R. Mijumbi, J. Serrat and J.L. Gorricho are with the Network Engineering Department, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain.

N. Bouten and F. De Turck are with Department of Information Technology, Ghent University - iMinds, B-9050 Ghent, Belgium.

R. Boutaba is with the D.R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada.

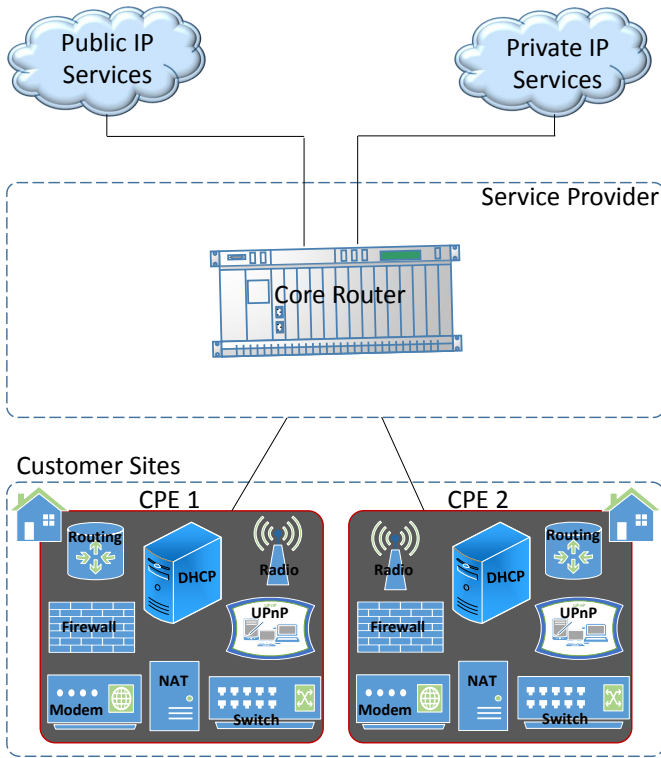


Fig. 1. Traditional CPE Implementations

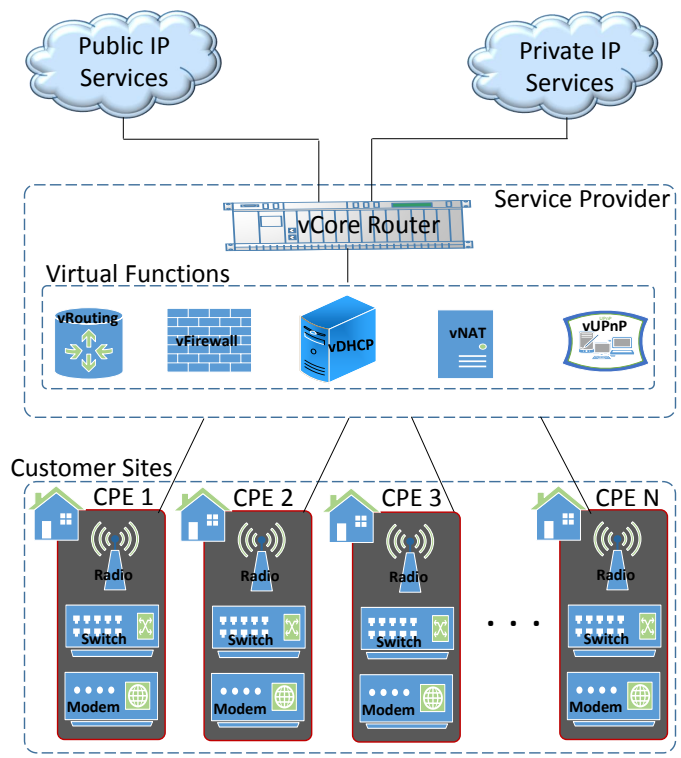


Fig. 2. Possible CPE Implementation with NFV

vides greater flexibility to scale the actual VNF performance in a more dynamic way and with finer granularity, for instance, according to the actual traffic for which the network operator needs to provision capacity.

It is worth remarking that the general concept of decoupling NFs from dedicated hardware does not necessarily require virtualization of resources. This means that TSPs could still purchase or develop software (NFs) and run it on physical machines. The difference is that these NFs would have to be able to run on commodity servers. However, the gains (such as flexibility, dynamic resource scaling, energy efficiency) anticipated from running these functions on virtualized resources are very strong selling points of NFV. Needless to mention, it is also possible to have hybrid scenarios where functions running on virtualized resources co-exist with those running on physical resources. Such hybrid scenarios may be important in the transition towards NFV.

A. History of Network Function Virtualization

The concept and collaborative work on NFV was born in October 2012 when a number of the world’s leading TSPs jointly authored a white paper [4] calling for industrial and research action. In November 2012 seven of these operators (AT&T, BT, Deutsche Telekom, Orange, Telecom Italia, Telefonica and Verizon) selected the European Telecommunications Standards Institute (ETSI)[6] to be the home of the Industry Specification Group for NFV (ETSI ISG NFV)¹.

¹In the rest of this paper, the acronyms ETSI and ETSI ISG NFV are used synonymously.

Now, more than two years later, a large community of experts are working intensely to develop the required standards for NFV as well as sharing their experiences of its development and early implementation. The membership of ETSI has grown to over 245 individual companies including 37 of the world’s major service providers as well as representatives from both telecoms and IT vendors [6]. ETSI has successfully completed Phase 1 of its work with the publication of 11 ETSI Group Specifications [7]. These specifications build on the first release of ETSI documents published in October 2013 and include an infrastructure overview, updated architectural framework, and descriptions of the compute, hypervisor and network domains of the infrastructure. They also cover Management and Orchestration (MANO), security and trust, resilience and service quality metrics.

Since ETSI is not a standards body, its aim is to produce requirements and potential specifications that TSPs and equipment vendors can adapt for their individual environments, and which may be developed by an appropriate standards development organization (SDO). However, since standards bodies such as the 3GPP [8] are in liaison with the ETSI, we can expect these proposals will be generally accepted and enforced as standards. 3GPP’s Telecom Management working group (SA5) is also studying the management of virtualized 3GPP network functions.

B. NFV Examples

The ETSI has proposed a number of use cases for NFV [9]. In this subsection, we will explain how NFV may be applied

to Customer Premises Equipment (CPE), and to an Evolved Packet Core (EPC) network.

1) *Customer Premises Equipment (CPE)*: In Figures 1 and 2, we use an example of a CPE to illustrate the economies of scale that may be achieved by NFV. Fig. 1 shows a typical (current) implementation of a CPE which is made up of the functions: Dynamic Host Configuration Protocol (DHCP), Network Address Translation (NAT), routing, Universal Plug and Play (UPnP), Firewall, Modem, radio and switching. In this example, a single service (the CPE) is made up of eight functions. These functions may have precedence requirements. For example, if the functions are part of a service chain², it may be required to perform firewall functions before NAT. Currently, it is necessary to have these functions in a physical device located at the premises of each of the customers 1 and 2. With such an implementation, if there is a need to make changes to the CPE, say, by adding, removing or updating a function, it may be necessary for a technician from the ISP to individually talk to or go to each of the customers. It may even require a complete change of the device in case of additions. This is not only expensive (operationally) for the ISPs, but also for the customers.

In Figure 2, we show a possible implementation based on NFV in which some of the functions of the CPE are transferred to a shared infrastructure at the ISP, which could also be a data center. This makes the changes described above easier since, for example, updating the DHCP for all customers would only involve changes at the ISP. In the same way, adding another function such as parental controls for all or a subset of customers can be done at once. In addition to saving on operational costs for the ISP, this potentially leads to cheaper CPEs if considered on a large scale.

2) *Evolved Packet Core*: Virtualizing the EPC is another example of NFV that has attracted a lot of attention from industry. The EPC is the core network for Long Term Evolution (LTE) as specified by 3GPP [8]. On the left side of Fig. 3, we show a basic architecture of LTE without NFV. The User Equipment (UE) is connected to the EPC over the LTE access network (E-UTRAN). The evolved NodeB (eNodeB) is the base station for LTE radio. The EPC performs essential functions including subscriber tracking, mobility management and session management. It is made up of four NFs: Serving Gateway (S-GW), Packet Data Network (PDN) Gateway (P-GW), Mobility Management Entity (MME), and Policy and Charging Rules Function (PCRF). It is also connected to external networks, which may include the IP Multimedia Core Network Subsystem (IMS). In the current EPC, all its functions are based on proprietary equipment. Therefore, even minor changes to a given function may require a replacement of the equipment. The same applies to cases when the capacity of the equipment has to be changed.

On the right side of Fig. 3, we show the same architecture in which the EPC is virtualized. In this case, either all functions




Fig. 3. Virtualization of the EPC

in the EPC, or only a few of them are transferred to a shared (cloud) infrastructure. Virtualizing the EPC could potentially lead to better flexibility and dynamic scaling, and hence allow TSPs to respond easily and cheaply to changes in market conditions. For example, as represented by the number of servers allocated to each function in Fig. 3, there might be a need to increase user plane resources without affecting the control plane. In this case, VNFs such as a virtual MME may scale independently according to their specific resource requirements. In the same way, VNFs dealing with the data plane might require a different number of resources than those dealing with signaling only. This flexibility would lead to more efficient utilization of resources. Finally, it also allows for easier software upgrades on the EPC network functions, which would hence allow for faster launch of innovative services.

C. Related Work and Open Questions

While both industry and academia embrace NFV at unprecedented speeds, the development is still at an early stage, with many open questions. As TSPs and vendors look at the details of implementing NFV and accomplishing its foreseen goals, there are concerns about the realization of some of these goals and whether implementation translates to the benefits initially expected. There are important unexplored research challenges such as testing and validation [10], resource management, inter-operability, instantiation, performance of VNFs, etc, that should be addressed. Even areas being explored such as MANO still have open questions especially with regard to support for heterogeneity.

There have been recent efforts to introduce NFV, explain its performance requirements, architecture, use cases and potential approaches to challenges [3]. A discussion of challenges to introducing NFV in mobile networks, with a focus on virtualized evolved packet core is presented in [11], while

²The chain of functions that make up a service for which the connectivity order is important is known as VNF Forwarding Graph (VNFFG) [9]. In addition to sequencing requirements, the links in a VNFFG may split (i.e. from one function, packets could take one of many paths which lead to similar functionality), or may join.

the reliability challenges of NFV infrastructures are examined in [12]. However, all efforts in current literature are narrow in at least one of the following main ways: (1) with regard to scope, they do not consider important aspects of NFV, such as its relationship with SDN and cloud computing, (2) limited review and analysis of standardization activities, and (3) incomplete descriptions of ongoing research and state-of-the-art efforts and research challenges.

This paper examines the state-of-the-art in NFV and identifies key research areas for future exploration. In addition, we explore the relationship between NFV and two closely related fields, SDN [13] and cloud computing [14]. We also describe the different research and industrial initiatives and projects on NFV, as well as early implementation, proof of concepts and product cases. To the best of our knowledge, this paper presents the most comprehensive state-of-the-art survey on NFV to date.

D. Organization

The rest of this paper is organized as follows: Section II presents the NFV architecture that has been proposed by ETSI, and discusses its limitations. We propose a reference business model and identify important design considerations in section III. In section IV, we introduce SDN and cloud computing, describing the relationship between them and NFV, as well as current efforts to implement environments involving all of them. In section V, we survey the major projects on NFV as well as early implementations, use cases and commercial products. Based on a qualitative analysis of the state-of-the-art, section VI identifies key research areas for further exploration, and section VII concludes this paper.

II. NFV ARCHITECTURE

According to ETSI, the NFV Architecture is composed of three key elements: Network Function Virtualization Infrastructure (NFVI), VNFs and NFV MANO [15]. We represent them graphically in Fig. 4. In this section these elements are defined [5], [15], [16].

A. NFV Infrastructure (NFVI)

The NFVI is the combination of both hardware and software resources which make up the environment in which VNFs are deployed. The physical resources include commercial-off-the-shelf (COTS) computing hardware, storage and network (made up of nodes and links) that provide processing, storage and connectivity to VNFs. Virtual resources are abstractions of the computing, storage and network resources. The abstraction is achieved using a virtualization layer (based on a hypervisor), which decouples the virtual resources from the underlying physical resources. In a data center environment, the computing and storage resources may be represented in terms of one or more Virtual Machines (VMs), while virtual networks are made up of virtual links and nodes. A virtual node is a software component with either hosting or routing functionality, for example an operating system encapsulated in a VM. A virtual link is a logical interconnection of two virtual nodes, appearing to them as a direct physical link with dynamically changing properties [17].

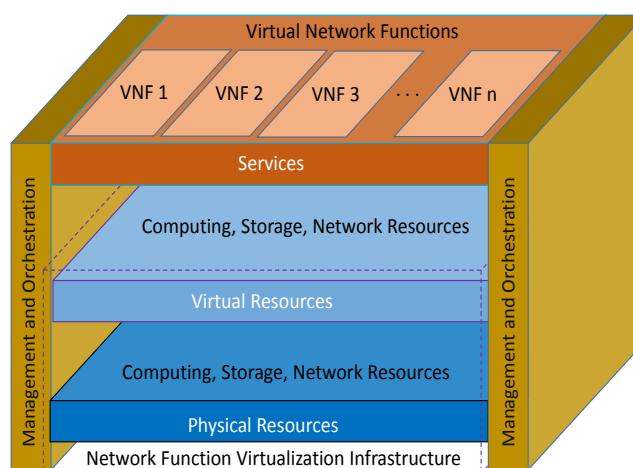


Fig. 4. Network Function Virtualization Architecture

B. Virtual Network Functions and Services

A NF is a functional block within a network infrastructure that has well defined external interfaces and well-defined functional behaviour [15]. Examples of NFs are elements in a home network, e.g. Residential Gateway (RGW); and conventional network functions, e.g. DHCP servers, firewalls, etc. Therefore, a VNF is an implementation of an NF that is deployed on virtual resources such as a VM. A single VNF may be composed of multiple internal components, and hence it could be deployed over multiple VMs, in which case each VM hosts a single component of the VNF [5]. A service is an offering provided by a TSP that is composed of one or more NFs. In the case of NFV, the NFs that make up the service are virtualized and deployed on virtual resources such as a VM. However, in the perspective of the users, the services—whether based on functions running dedicated equipment or on VMs—should have the same performance. The number, type and ordering of VNFs that make it up are determined by the service’s functional and behavioral specification. Therefore, the behaviour of the service is dependent on that of the constituent VNFs.

C. NFV Management and Orchestration (NFV MANO)

According to the ETSI’s MANO framework [18], NFV MANO provides the functionality required for the provisioning of VNFs, and the related operations, such as the configuration of the VNFs and the infrastructure these functions run on. It includes the orchestration and lifecycle management of physical and/or software resources that support the infrastructure virtualization, and the lifecycle management of VNFs. It also includes databases that are used to store the information and data models which define both deployment as well as lifecycle properties of functions, services, and resources. NFV MANO focuses on all virtualization-specific management tasks necessary in the NFV framework. In addition the framework defines interfaces that can be used for communications between the different components of the NFV MANO, as well as coordination with traditional network management systems such as Operations

Support System (OSS) and Business Support Systems (BSS) so as to allow for management of both VNFs as well as functions running on legacy equipment.

Discussion: The ETSI-proposed NFV reference architecture specifies initial functional requirements and outlines the required interfaces. However, the ETSI's scope of work is rather limited, excluding aspects such as control and management of legacy equipment [5]. This could make it difficult to specify the operation and MANO of an end-to-end service involving both legacy functions and VNFs. In addition, standards and/or de-facto best practices and reference implementations of the VNFs, infrastructure, MANO and detailed definitions of required interfaces are not yet available.

In particular, it can be seen from current NFV solutions that vendors have differing ideas on what constitutes an NFVI and VNFs, and how both of them can be modeled. There remains a number of open questions such as: (1) which NFs should be deployed in data center nodes, and which ones in operator nodes; (2) which functions should be deployed on dedicated VMs and which ones in containers³; (3) what quantity and types of NFVI resources will be required to run specific functions; and (4) operational requirements of environments that involve both VNFs and those running on legacy equipment. While many of these questions such as inter-operability and interface definition will be addressed in the second Phase of ETSI's work, time is of the essence. Since both vendors and TSPs are already investing significantly in NFV, we could reach a point where it is impossible to reverse the vendor-specific solutions.

III. BUSINESS MODEL AND DESIGN CONSIDERATIONS

Using the architecture represented in Fig. 4, and based on business models for network virtualization [20] and cloud computing [14], we identify five main players in a NFV environment and propose a reference business model that illustrates the possible business relationships between them as shown in Fig. 5. We also discuss important NFV system design considerations.

A. Business Model

1) *Infrastructure Provider (InP)*: InPs deploy and manage physical resources in form of data centers and physical networks. It is on top of these resources that virtual resources may be provisioned and leased through programming interfaces to one or more TSPs. The InPs may also determine how the pool of the available resources are allocated to the TSPs. In NFV, examples of InPs could be public data centers such as those by Amazon, or private servers owned by TSPs. If a given InP is not able to provide resources fully or in part to a given TSPs, negotiations and hence coalitions can be formed with other InPs so as to provision multi-domain VNFs [21].

³In fact, even the fact whether containers may be used to host VNFs and the corresponding ecosystem still needs research [19].

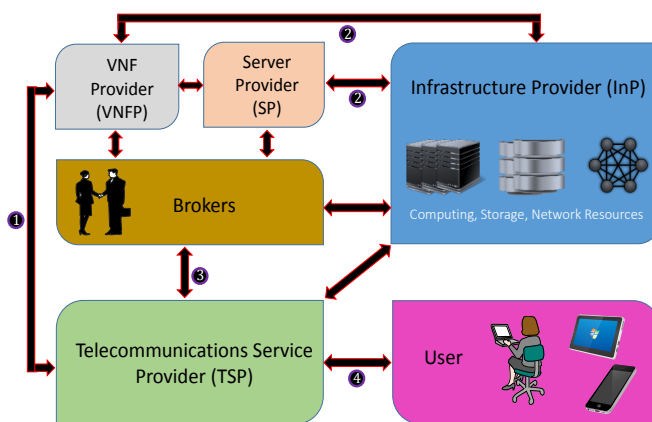


Fig. 5. Proposed NFV Business Model

2) *Telecommunications Service Provider (TSP)*: TSPs⁴ lease resources from one or more InPs, which they use for running VNFs. They also determine the chaining of these functions to create services for end users. In a more general case, TSPs may sub-lease their virtual resources to other TSPs. In such a case, the reselling TSP would take up the role of a InP. In cases where the InP is private or in-house, e.g. provided by TSP network nodes or servers, then the InP and TSP may be one entity.

3) *VNF Providers (VNFPs) and Server Providers (SPs)*: NFV splits the role of traditional network equipment vendors (such as Cisco, Huawei, HP and Alcatel-Lucent) into two: VNFPs and SPs. VNFPs provide software implementations for NFs. These functions may either be provided directly to TSPs (via interface 1), or VNFPs could provide them to InPs (via interface 2), who would then provide both infrastructure as well as VNFs to TSPs. It is also possible that TSPs develop (some of) their own NFs (software). In this case, VNFPs and TSPs would be one entity.

In the same way, SPs provide industry standard servers on which VNFs can be deployed. These servers may be provided to InPs (in case the functions will be run in a cloud), or to TSPs (in case the functions will be run in the network nodes of TSPs). It is worth noting that these entities (VNFPs and SPs) may in fact be one company. The main difference is that the functions they provide are not tied to running on equipment with specialized functionality or made by a specific vendor. In other words, a TSP could purchase VNFs from one entity, and servers from a different one.

4) *Brokers*: In some cases, a TSP may need to purchase functions which make up a single service from multiple VNFPs, and/or to deploy and manage the resulting end-to-end services running on resources from multiple InPs. In this case, it may be necessary to have a brokerage role. The brokers would receive resource and/or functions requirements from TSPs and then discover, negotiate and aggregate resources and functions from multiple InPs, VNFPs and SPs to offer them as

⁴In this paper, we use the term TSP to generally mean all service providers. This includes service providers such as Netflix that deploy services with caches in different locations, as well as the traditional TSPs such as Telefonica and Deutsche Telekom.

a service to the TSP. This role is only included in the model for completeness as it may not be required in all cases of the NFV ecosystem.

5) *End User*: End users are the final consumers of the services provided by TSPs. They are similar to the end users in the existing Internet, except that the existence of multiple services from competing TSPs enables them to choose from a wide range of services. End users may connect to multiple TSPs for different services.

Finally, the arrows in Fig. 5 indicate business relationships or interfaces between the the different entities. For example, VNFPs and/or SPs use interfaces 1 and 2 to negotiate and/or provide VNFs and commodity servers respectively, to TSPs and InPs, while TSPs use interfaces 3 and 4 for their interactions with brokers and users respectively.

B. NFV Design Considerations

As NFV matures, it is important to note that it is not only sufficient to deploy NFs over virtualized infrastructures. Network users are generally not concerned with the complexity (or otherwise) of the underlying network. All users require is for the network to allow them access to the applications they need, when they need them. Therefore, NFV will only be an acceptable solution for TSPs if it meets key considerations identified below.

1) *Network Architecture and Performance*: To be acceptable, NFV architectures should be able to achieve performance similar to that obtained from functions running on dedicated hardware. This requires that all potential bottlenecks at all layers of the stack are evaluated and mitigated. As an example, if VNFs belonging to the same service are placed in different VMs, then there must be a connection between these two VMs, and this connection must provide sustained, aggregated high bandwidth network traffic to the VNFs. To this end, it may be important for the network to be able to take advantage of connections to the network interfaces that are high-bandwidth and low latency due to processor offload techniques such as direct memory access (DMA)[22] for data movement and hardware assist for CRC computation [23], [24].

In addition, some VNFs such as Deep Packet Inspection (DPI) are network and compute intensive, and may require some form of hardware acceleration [25] to be provided by the NFVI to still meet their performance goals [26]. Some recent efforts [27] have studied the implications of utilizing Data Plane Development Kits (DPDKs) for running VNFs and shown that near-native (i.e., similar to non-virtualized) performance for small and large packet processing can be achieved. In addition, Field-Programmable Gate Arrays (FPGAs) have also been shown to enhance performance of VNFs [28], [29]. Finally, VNFs should only be allocated the storage and computation resources they need. Otherwise, NFV deployments may end up requiring more resources, and hence there would be no justification for transiting to NFV.

2) *Security and Resilience*: The dynamic nature of NFV demands that security technologies, policies, processes and practices are embedded in its genetic fabric [30]. In particular, there are two important security risks that should be considered

in NFVI designs: (1) functions or services from different subscribers should be protected/isolated from each other. This helps to ensure that functions are resilient to faults and attacks since a failure or security breach in one function/service would not affect another. (2) the NFVI (physical and virtual resources) should be protected from the delivered subscriber services. One way to secure the NFVI is to deploy internal firewalls within the virtual environment [24]. These would allow for the NFV MANO to access to the VNFs without letting malicious traffic from the customer networks into the NFVI. Finally, to make service deployment resilient, it may be necessary for functions that make up the same service not be hosted by physical resources in the same fault or security domain during deployment.

3) *Reliability and Availability*: Whereas in the IT domain outages lasting seconds are tolerable and a user typically initiates retries, in telecommunications there is an underlying service expectation that outages will be below the recognizable level (i.e. in the order of milliseconds), and service recovery is performed automatically. Furthermore, service impacting outages need to be limited to a certain amount of users (e.g. a certain geography) and network wide outages are not acceptable [31]. These high reliability and availability needs are not only a customer expectation, but often a regulatory requirement, as TSPs are considered to be part of critical national infrastructure, and respective legal obligations for service assurance/business continuity are in place. However, not every function has the same requirements for resiliency: For example, whereas telephony usually has the highest requirements for availability, other services, e.g. Short Messaging Service (SMS), may have lower availability requirements. Thus, multiple availability classes may be defined which should be supported by a NFV framework [31]. Again, functions may be deployed with redundancy to recover from software or hardware failures.

4) *Support for Heterogeneity*: The main selling point of NFV is based on breaking the barriers that result from proprietary hardware-based service provision. It is therefore needless to mention that openness and heterogeneity will be at the core of NFV's success. Vendor-specific NFV solutions with vendor-specific hardware and platform capabilities defeat the original NFV concept and purpose. Therefore, any acceptable NFV platform must be an open, shared environment capable of running applications from different vendors. InPs must be free to make their own hardware selection decisions, change hardware vendors, and deal with heterogeneous hardware. In addition, such platforms should be able to shield VNFs from the specifics of the underlying networking technologies (e.g., optical, wireless, sensor etc.) [32]. Finally, and equally important, platforms should allow for possibilities of an end-to-end service to be created on top of more than one infrastructural domain without restrictions, and without need for technology specific solutions. While virtualization within a single InP reduces cost, inter-provider NFV enables the "productization" of the same internal software functions and results in opportunities for revenue growth [33]. As an example, if a mobile user subscribing to given TSP roams into the coverage of another TSP, the user should not be restricted

to voice, data and simple messaging services. The real power of NFV would be realized if such a user is able to choose a firewall or security service from the current TSP, or use a combination of functions from the host TSP and others from the one for which he has coverage.

5) *Legacy Support*: Backward compatibility will always be an issue of high concern for any new technology. NFV is not an exception. It is even more important for the telecommunications industry, given that even for a given operator that decides to make the transition to NFV, it may take time for this to be complete, let alone the fact that some operators will do this faster than others. Therefore, support for both physical and virtual NFs is important for operators making the transition to NFV as they may need to manage legacy physical assets alongside virtualized functions for some time. This may necessitate having an orchestration strategy that closes the gap between legacy services and NFV. It is important to maintain a migration path toward NFV, while keeping operators' current network investments in place [34]. InPs must be able to function in an environment whereby both virtualized and physical network functions operate on the network simultaneously.

6) *Network Scalability and Automation*: In order to achieve the full benefits of NFV, a scalable and responsive networking solution is necessary. Therefore, while meeting the above design considerations, NFV needs to be acceptably scalable to be able to support millions of subscribers. To give an example, most current NFV proof-of-concepts are based on deploying a VM to host a VNF. Just like a single VM may not be able to meet the requirements of a given function, it is not economical to deploy a VM per NFV, as the resulting VM footprint would be too large, and would lead to scalability problems at the virtualization layer. However, NFV will only scale if all of the functions can be automated. Therefore, automation of processes is of paramount importance to the success of NFV [4]. In addition, the need for dynamic environments requires that VNFs can be deployed and removed on demand and scaled to match changing traffic.

IV. RELATED CONCEPTS

The need for innovativeness, agility and resource sharing is not new. In the past, the communications industry has invented and deployed new technologies to help them offer new and multiple services in a more agile, cost and resource effective way. In this section, we introduce two such concepts that are closely related to NFV; cloud computing and SDN. We also discuss the relationship between NFV and each of them, as well as current attempts to enable all three to work together.

A. Cloud Computing

According to NIST [35] cloud computing is "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction". In a cloud computing environment, the traditional role of service provider is divided

into two: the infrastructure providers who manage cloud platforms and lease resources according to a usage-based pricing model, and service providers, who rent resources from one or many infrastructure providers to serve the end users [14]. The cloud model is composed of five essential characteristics and three service models [35]. We briefly introduce these in the following subsections.

1) *Essential Characteristics of Cloud Computing: On-demand self-service*. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access. Capabilities (e.g. compute resources, storage capacity) are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.

Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand.

Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts).

2) *Cloud Computing Service Models*: The three service models of cloud computing are shown in Fig. 6, and defined below [35].

Software as a Service (SaaS). The user is able to use the providers applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface.

Platform as a Service (PaaS). The user is able to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider.

Infrastructure as a Service (IaaS). The user is able to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications.

3) *Relationship between Cloud Computing and NFV*: In general, NFV is not restricted to functions for services in telecommunications. In fact, many IT applications already run on commodity servers in the cloud [40]. However, since most of the promising use cases for NFV originate from the telecommunications industry, and because the performance and reliability requirements of carrier-grade functions are higher than those of IT applications, the discussions in this paper consider that acceptable NFV performance should be carrier-class. In Fig. 6, we have mapped the cloud service models to part of the NFV architecture. It can be observed that IaaS corresponds to both the physical and virtual resources in

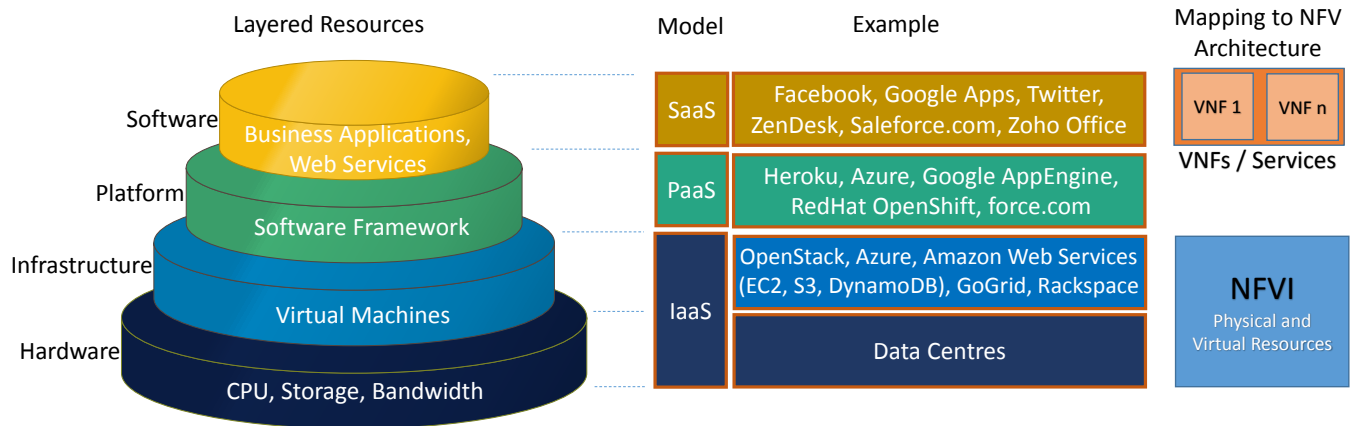


Fig. 6. Cloud Computing Service Models and their Mapping to Part of the NFV Reference Architecture

TABLE I
COMPARISON OF NFV IN TELECOMMUNICATION NETWORKS AND CLOUD COMPUTING

Issue	NFV (Telecom Networks)	Cloud Computing
Approach	Service/Function Abstraction	Computing Abstraction
Formalization	ETSI NFV Industry Standard Group	DMTF Cloud Management Working Group [36]
Latency	Expectations for low latency	Some latency is acceptable
Infrastructure	Heterogeneous transport (Optical, Ethernet, Wireless)	Homogeneous transport (Ethernet)
Protocol	Multiple Control Protocols (e.g OpenFlow [37], SNMP [38])	OpenFlow
Reliability	Strict 5 NINES availability requirements [39]	Less strict reliability requirements [40]
Regulation	Strict Requirements e.g NEBS [41]	Still diverse and changing

the NFVI, while the services and VNFs in NFV are similar to the SaaS service model in cloud computing.

Being the cheapest choice for testing and implementation, most NFV proof of concepts and early implementations have been based on deploying functions on dedicated VMs in the cloud. The flexibility of cloud computing, including rapid deployment of new services, ease of scalability, and reduced duplication, make it the best candidate that offers a chance of achieving the efficiency and expense reduction that are motivating TSPs towards NFV.

However, deploying NFs in the cloud will likely change every aspect of how services and applications are developed and delivered. While work continues to be done with respect to networked clouds and inter-cloud networking [42], [43], telecommunication networks differ from the cloud computing environment in at least three ways: (1) data plane workloads in telecom networks imply high pressure on performance, (2) telecom network topologies place tough demands on the network and the need for global network view for management [44], (3) the telecom industry requires scalability, five-nines availability and reliability. In traditional telecom networks, these features are provided by the site infrastructure. If NFV should be based on cloud computing, these features need to be replicated by the cloud infrastructure in such a way that they can be orchestrated, as orchestrated features can be exposed through appropriate abstractions, as well as being coupled with

advanced support for discoverability and traceability [45]. It is therefore worth stressing that NFV will require more considerations than just transferring carrier class network functions to the cloud. There is need to adapt cloud environments so as to obtain carrier-class behaviour [44]. In Table I, we summarize the relationship between NFV for telecom networks and cloud computing.

4) *Research on Cloud-based NFV:* In order for NFV to perform acceptably in cloud computing environments, the underlying infrastructure needs to provide a certain number of functionalities which range from scheduling to networking and from orchestration to monitoring capacities. While OpenStack has been identified as one of the main components of a cloud-based NFV architectural framework, it currently does not meet some NFV requirements. For example, through a gap analysis in [46], it was noted that, among other gaps, OpenStack neither provides detailed description of network resources including Quality-of-Service (QoS) requirements, nor supports a resource reservation service and consequently it does not provide any interface for resource reservation.

In addition, through measurements some performance degradation has been reported [47]. Some efforts have already been dedicated to study the requirements needed to make the performance of cloud carrier-grade [48], [49], [50]. In particular, OpenANFV [28] proposes an OpenStack-based framework which uses hardware acceleration to enhance the

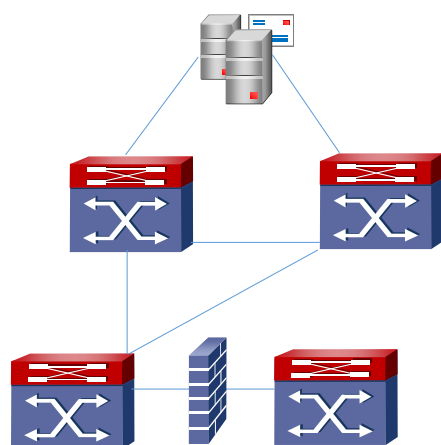


Fig. 7. Distributed Control and Middleboxes (e.g. Firewall, Intrusion Detection, etc.) in Traditional Networks

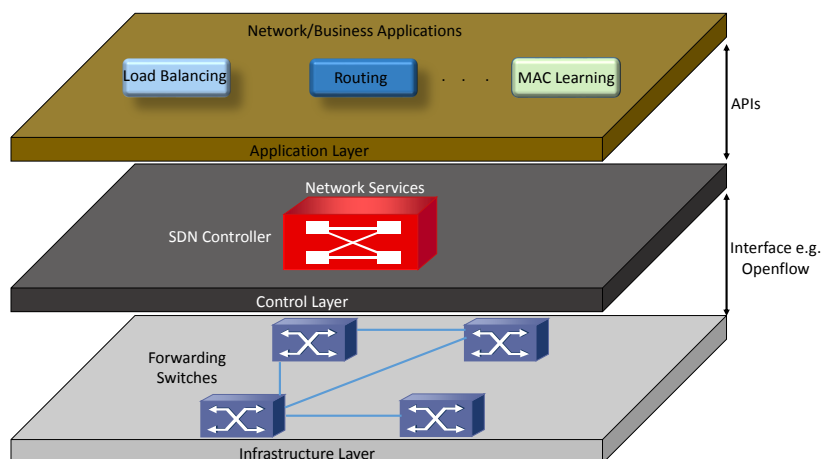


Fig. 8. Logical Layers in a Software Defined Network

performance of VNFs. The author’s efforts are motivated by the observation that for some functions (e.g., DPI, network deduplication (Dedup) and NAT), industry standard servers may not achieve the required levels of performance. Therefore, OpenNFV aims at providing elastic, automated provisioning for hardware acceleration to VNFs in OpenStack. To this end, the tested VNFs (DPI, Dedup and NAT) were allowed access to a predefined set of accelerated behavior and to communicate through a hardware-independent interface with the hypervisor to configure the accelerator. The authors reported performances 20, 8 and 10 times better for DPI, Dedup and NAT respectively.

B. Software Defined Networking (SDN)

SDN [51] is currently attracting significant attention from both academia and industry as an important architecture for the management of large scale complex networks, which may require re-policing or re-configurations from time to time. As shown in Figures 7 and 8, SDN decouples the network control and forwarding functions. This allows network control to become directly programmable via an open interface (e.g., ForCES [52], OpenFlow [53], etc) and the underlying infrastructure to become simple packet forwarding devices (the data plane) that can be programmed.

While the SDN control plane can be implemented as pure software which runs on industry-standard hardware, the forwarding plane requires an SDN agent [54], and may therefore require to be implemented in specialized hardware. However, depending on the performance and capacity needs of the SDN networking element, and depending on whether specialized hardware transport interfaces are required, the forwarding plane may also be implemented on commodity servers [55]. For example, VMware’s NSX platform [56] includes a virtual switch (vSwitch) and controller both of which implement SDN protocols without requiring specialized hardware.

SDN has the potential to dramatically simplify network management and enable innovation and evolution [57]. According to the Open Network Foundation (ONF) [58], SDN addresses the fact that the static architecture of conventional

networks is ill-suited for the dynamic computing and storage needs of today’s data centers, campuses, and carrier environments. The SDN architecture is [59]:

Programmable. SDN makes network control directly programmable since control is decoupled from forwarding functions. This programmability can be used to automate network configuration in such a way that network administrators can run ‘SDN apps’ that help to optimize particular services such as VoIP so as to ensure a high Quality-of-Experience (QoE) for phone calls.

Agile. Abstracting control from forwarding lets administrators dynamically adjust network-wide traffic flow to meet changing needs. This makes the network more agile since logic is now implemented in a software running on commodity hardware, which has shorter release cycles than device firmware.

Centrally managed. Network intelligence is (logically) centralized in software-based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch.

Open standards-based and vendor-neutral. When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.

1) *Relationship between SDN and NFV:* NFV and SDN have a lot in common since they both advocate for a passage towards open software and standard network hardware. Specifically, in the same way that NFV aims at running NFs on industry standard hardware, the SDN control plane can be implemented as pure software running on industry standard hardware. In addition, both NFV and SDN seek to leverage automation and virtualization to achieve their respective goals. In fact, NFV and SDN may be highly complimentary, and hence combining them in one networking solution may lead to greater value. For example, if it is able to run on a VM, an SDN controller may be implemented as part of a service chain. This means that the centralized control and management applications (such as load balancing, monitoring and traffic analysis) used in SDN can be realized, in part, as VNFs, and hence benefit from NFV’s reliability and elasticity features.

TABLE II
COMPARISON OF SOFTWARE DEFINED NETWORKING AND NETWORK FUNCTION VIRTUALIZATION CONCEPTS

Issue	NFV (Telecom Networks)	Software Defined Networking
Approach	Service/Function Abstraction	Networking Abstraction
Formalization	ETSI	ONF
Advantage	Promises to bring flexibility and cost reduction	Promises to bring unified programmable control and open interfaces
Protocol	Multiple control protocols (e.g. SNMP, NETCONF)	OpenFlow is de-facto standard
Applications run	Commodity servers and switches	Commodity servers for control plane and possibility for specialized hardware for data plane
Leaders	Mainly Telecom service providers	Mainly networking software and hardware vendors
Business Initiator	Telecom service providers	Born on the campus, matured in the data center

In the same way, SDN can accelerate NFV deployment by offering a flexible and automated way of chaining functions, provisioning and configuration of network connectivity and bandwidth, automation of operations, security and policy control [60]. It is however worth stressing that most of the advantages expected from both NFV and SDN are promises that have not been proven yet.

However, SDN and NFV are different concepts, aimed at addressing different aspects of a software-driven networking solution. NFV aims at decoupling NFs from specialized hardware elements while SDN focuses on separating the handling of packets and connections from overall network control. As stated by the ONF in the description of the SDN architecture [54], “the NFV concept differs from the virtualization concept as used in the SDN architecture. In the SDN architecture, virtualization is the allocation of abstract resources to particular clients or applications; in NFV, the goal is to abstract NFs away from dedicated hardware, for example to allow them to be hosted on server platforms in cloud data centers”. It can be observed that the highest efforts in promoting and standardizing SDN is in data center and cloud computing areas while telecom carriers are driving similar efforts for NFV. Finally, an important distinction is that while NFV can work on existing networks because it resides on servers and interacts with specific traffic sent to them, SDN requires a new network construct where the data and control planes are separate. We summarize the relationship between SDN and NFV in Table II.

2) *Research on SDN-based NFV*: There is currently a lot of work involving the combination of SDN and NFV to enhance either of them; including: a ForCES-based framework [61], NFV-based monitoring for SDN [62], an abstraction model for both the forwarding model and for the network functions [61]. As these efforts show, the unique demands of NFV will potentially necessitate a massively complex forwarding plane, blending virtual and physical appliances with extensive control and application software, some of it proprietary [63]. There are two major aspects of SDN that may need to be improved in order to meet the requirements of NFV: the Southbound API (mainly OpenFlow), and controller designs. We discuss advances in each of these two aspects below.

a) *Southbound API*: OpenFlow is the de-facto implementation of a southbound API for SDN. However, before we consider NFV support, even in current SDN environments OpenFlow is by no means a mature solution [64]. Since OpenFlow targets L2-L4 flow handling, it has no application-layer protocol support and switch-oriented flow control. Therefore, users have to arrange additional mechanism for upper-layer flow control. Furthermore, executing a lot of flow matching on a single switch (or virtual switch) can cause difficulties in network tracing and overall performance degradation [65].

Therefore, OpenFlow will have to be extended to include layers L5-L7 to be able to support NFV. Basta et al. [66] investigated the current OpenFlow implementation in terms of the basic core operations such as QoS, data classification, tunneling and charging, concluding that there is a need for an enhanced OpenFlow to be able to support some functions in an NFV environment. In an implementation of a virtual EPC function [9], [67] extends OpenFlow 1.2 by defining virtual ports to allow encapsulation and to allow flow routing using the GTP Tunnel Endpoint Identifier (TEID).

Finally, while OpenFlow assumes a logically centralized controller, which ideally can be physically distributed, most current deployments rely on a single controller. This does not scale well and can adversely impact reliability. In addition, network devices in an NFVI require collaboration to be able to provide services, which cannot currently be provided by SDN. There is therefore still a need to improve SDN by considering distributed architectures [68], [69]. It may also be important for TSPs, InPs and ETSI to consider other possible solutions such as NETCONF [70].

b) *Controller Design*: While there are multiple controllers that may be used in an SDN environment, all of them require improvements to be able to support NFV requirements, especially with regard to distributed network management and scalability. OpenNF [71], [65] proposes a control plane that allows packet processing to be redistributed across a collection of NF instances, and provides a communication path between each NF and the controller for configuration and decision making. It uses a combination of events and forwarding updates to address race conditions, bound overhead, and accommodate a variety of NFs. [72] also designed a protocol to implement the

communication between the controllers and the VNFs. Finally, [73] proposes an architecture that considers the control of both SDN and NFV.

OpenDaylight [74] is one of the few SDN control platforms that supports a broader integration of technologies in a single control platform [75]. A collaborative project hosted by the Linux Foundation, OpenDaylight is a community-led and industry-supported open source framework to accelerate adoption, foster new innovation and create a more open and transparent approach to SDN and NFV. The objective of the OpenDaylight initiative is to create a reference framework for programmability and control through an open source SDN and NFV solution. The argument of OpenDaylight is that building upon an open source SDN and NFV controller enables users to reduce operational complexity, extend the life of their existing infrastructure hardware and enable new services and capabilities only available with SDN.

C. Summary: NFV, SDN and Cloud Computing

To summarize the relationship between NFV, SDN, and cloud computing, we use Fig. 9⁵. We observe that each of these fields is an abstraction of different resources: compute for cloud computing, network for SDN, and functions for NFV. The advantages that accrue from each of them are similar; agility, cost reduction, dynamism, automation, resource scaling etc.

The question is not whether NFs will be migrated to the cloud, as this is in fact the general idea of NFV. It is whether the cloud will be a public one like Amazon, or if TSPs will prefer to user private ones distributed across their infrastructure. Either way, work will have to be done to make the cloud carrier-grade in terms of performance, reliability, security, communication between functions, etc.

On the other hand, NFV goals can be achieved using non-SDN mechanisms, and relying on the techniques currently in use in many data centers. However, approaches relying on the separation of the control and data forwarding planes as proposed by SDN can enhance performance, simplify compatibility with existing deployments, and facilitate operation and maintenance procedures. In the same way, NFV is able to support SDN by providing the infrastructure upon which the SDN software can be run. Finally, the modern variant of a data center (the cloud and it's self-service aspect) relies on automated management that may be obtained from SDN and NFV. In particular, aspects such as network as a service, load balancing, firewall, VPN etc. all run in software instantiated via APIs

V. STATE-OF-THE-ART

As the ETSI continues work on NFV, several other standards organizations, academic and industrial research projects and vendors are working in parallel with diverse objectives, and some of them in close collaboration with the ETSI. In this section, we explore these NFV activities.

⁵It is worth remarking that OpenFlow is not the only SDN protocol. In the same way, OpenStack is not the only cloud computing platform. The reason we present only these two in Fig. 9 is that, as already mentioned, they have received more attention in general, and with regard to NFV.

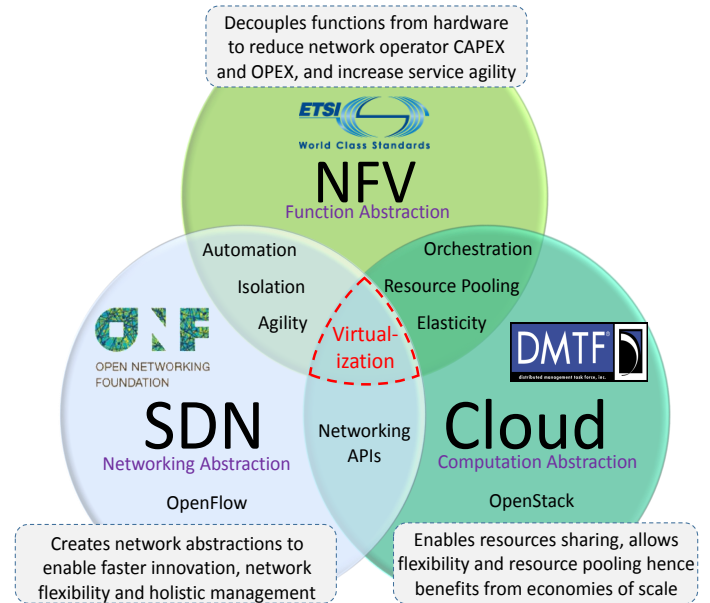


Fig. 9. Relationship between NFV, SDN & Cloud Computing

A. NFV Standardization Activities

1) *IETF Service Function Chaining Working Group*: Functions in a given service have strict chaining and/or ordering requirements that must be considered when decisions to place them in the cloud are made. The Internet Engineering Task Force (IETF) [76] has created the Service Function Chaining Working Group (IETF SFC WG) [77] to work on function chaining. The IETF SFC WG is aimed at producing an architecture for service function chaining that includes the necessary protocols or protocol extensions to convey the service function chain (SFC) and service function path information [78] to nodes that are involved in the implementation of service functions and SFCs, as well as mechanisms for steering traffic through service functions.

2) *IRTF NFV Research Group (NFVRG)*: The Internet Research Task Force (IRTF) has created a research group, NFVRG [79], to promote research on NFV. The group is aimed at organizing meetings and workshops at premier conferences and inviting special issues in well-known publications. The group focuses on research problems associated with NFV-related topics and on bringing a research community together that can jointly address them, concentrating on problems that relate not just to networking but also to computing and storage aspects in such environments.

3) *ATIS NFV Forum*: The ATIS NFV Forum [33] is an industry group created by the Alliance for Telecommunications Industry Solutions (ATIS), a North American telecom standards group. The group is aimed at developing specifications for NFV, focusing on aspects of NFV which include inter-carrier inter-operability and new service descriptions and automated processes. ATIS NFV Forum plans to develop technical requirements, the catalog of needed capabilities and the service chaining necessary for a third party service provider or enterprise to integrate the functions into a business application. This process is expected to result in creation of specifications

that are complementary with existing industry work products and that extend the current environment for inter-provider NFV. The forum also engages open source activities for the implementation of these capabilities in software.

4) *Broadband Forum*: The Broadband Forum (BB Forum) [80] is an industry consortium dedicated to developing broadband network specifications. Members include telecommunications networking and service provider companies, broadband device and equipment vendors, consultants and independent testing labs (ITLs). BB Forum collaborates with the ETSI after agreeing a formal liaison relationship in 2013. The BB Forum is working on how NFV can be used in the implementation of the multi-service broadband network (MSBN). To this end, the forum has many work items in progress, including: migrating to NFV in the context of TR-178 (WT-345), introducing NFV into the MSBN (SD-340), virtual business gateway (WT-328), flexible service chaining (SD-326) [81].

5) *Standardization of Related Paradigms*: In addition to the NFV standardization efforts, other bodies continue to work on standardization of related fields, SDN and cloud computing, which may also play a significant role in the success of NFV. The DMTF defined the Open Virtualization Format (OVF) [82] to address the portability and deployment of physical machines, virtual machines and appliances. OVF enables the packaging and secure distribution of virtual machines or appliances, providing cross-platform portability and simplified deployment across multiple platforms including cloud environments. OVF has adopted by both ANSI as a National Standard and ISO as the first international virtualization and cloud standard. It takes advantage of the DMTF's Common Information Model (CIM) [83], where appropriate, to allow management software to clearly understand and easily map resource properties by using an open standard. OVF and CIM may be used as one option for capturing some or all of the VNF package and/or Virtual Deployment Unit (VDU) descriptor [18], [84]. Although OVF does a great job enabling the provisioning of workloads across various clouds, it is still insufficient for new era cloud applications and runtime management.

In the same way, the ONF is standardizing the OpenFlow protocol and related technologies. ONF defines OpenFlow as the first standard communications interface defined between the control and forwarding layers of an SDN architecture. ONF has more than 123 member companies, including equipment vendors, semiconductor companies, computer companies, software companies, telecom service providers, etc.

In Table III, we summarize all the activities in the standardization of NFV and related technologies. In general, it can be said that there is sufficient involvement of standards bodies in NFV activities. While many of them work in liaison with the ETSI, some of them such as ATIS and 3GPP SA5 have identified and are working on specific aspects of NFV that have not yet been sufficiently developed by the ETSI. What remains to be seen is whether the output in terms of standards will match with the speed at which vendors and TSPs propose NFV solutions.

B. Collaborative NFV Projects

1) *Zoom*: Zero-time Orchestration, Operations and Management (ZOOM) [85] is a TM Forum project aimed at defining an operations environment necessary to enable the delivery and management of VNFs, and identifying new security approaches that will protect NFVI and VNFs. To achieve these objectives, the project regularly conducts a range of hands-on technology demos each of which is developed from what they call a catalyst project. Each catalyst project is sponsored by one or more network operators and equipment and software vendors in a real-world demo. The project currently runs about 9 catalysts with a focus on NFV aspects such as end-to-end automated management, security orchestration, function and service modeling, and using big data technologies and open software principles for workload placement.

2) *Open Platform for NFV (OPNFV)*: OPNFV [86] is an open source project founded and hosted by the Linux Foundation, and composed of TSPs and vendors. It aims to establish a carrier-grade, integrated, open source reference platform to advance the evolution of NFV and to ensure consistency, performance and inter-operability among multiple open source components. The first outcome of the project is referred to as OPNFV Arno [87], and was released in June 2015. The release provides an initial build of the NFVI and Virtual Infrastructure Manager (VIM) components of the ETSI architecture. It is developer-focused, and can therefore be used to explore NFV deployments, develop VNF applications, or to evaluate NFV performance and for use case-based testing. In particular, Arno has capabilities for integration, deployment and testing of components from other projects such as Ceph, KVM, OpenDaylight, OpenStack and Open vSwitch. In addition, end users and developers can deploy their own or third party VNFs on Arno to test its functionality and performance in various traffic scenarios and use cases.

3) *OpenMANO*: OpenMANO [88] is an open source project led by Telefonica, which is aimed at implementing ETSI's NFV MANO framework. Specifically, it attempts to address aspects related to performance and portability by applying Enhanced Platform Awareness (EPA) [89] principles. The OpenMANO architecture is made up of three main components: openmano, openvim and a graphical user interface (GUI). OpenMANO has a northbound interface (openmano API), based on REST, where MANO services are offered including the creation and deletion of VNF templates, VNF instances, network service templates and network service instances. Openvim is a lightweight, NFV-specific virtual infrastructure manager implementation directly interfacing with the compute and storage nodes in the NFVI, and with an openflow controller in order to create the infrastructural network topology. It offers a REST-based northbound interface (openvim API) where enhanced cloud services are offered including the lifecycle management of images, flavors, instances and networks. The REST interface of openvim is an extended version of the OpenStack API to accommodate EPA.

4) *Mobile Cloud Networking (MCN)*: MCN [90] is a consortium consisting of network operators, cloud providers, vendors, university and research institutes, as well as SMEs. The objective is to cloudify all components of a mobile network

TABLE III
SUMMARY OF NETWORK FUNCTION VIRTUALIZATION STANDARDIZATION EFFORTS

	Description	Focus Area	Description of NFV-Related Work
ETSI	Industry-led ETSI Standards Group	NFV	NFV architectural framework, infrastructure description, MANO, security and trust, resilience and service quality metrics.
3GPP SA5	3GPP's Telecom Management working group	Mobile Broadband	Working in liaison with the ETSI. Studying the management of virtualized 3GPP network functions.
IETF SFC WG	IETF Working Group	NFV	To propose a new approach to service delivery and operation, an architecture for service function chaining, management and security implications.
IRTF NFVRG	IRTF Research Group	NFV	Organizing NFV-related research activities in both academia and industry through workshops, research group meetings etc. at premier conferences.
ATIS NFV Forum	Industry-led Standards Group	NFV	Developing specifications for NFV, focusing on inter-carrier interoperability.
ONF	Industry-led consortium for standardization of OpenFlow	SDN	Standardizing the OpenFlow protocol and related technologies. Defines OpenFlow as the first standard communications interface defined between the control and forwarding layers of an SDN architecture.
DMTF OVF	Industry-led consortium	Cloud	DMTF's OVF and the CIM may be used as one option for capturing some or all of the VNF package and/or VDU [18] Descriptor.
BB Forum	Industry-led consortium that develops broadband network specifications	NFV in Broadband Networks	Collaborating with the ETSI to achieve a consistent approach and common architecture for the infrastructure needed to support VNFs.

operation such as: the access - Radio Access Network (RAN); the core - EPC; the services - IP Multimedia Subsystem (IMS), Content Delivery Networks (CDN) and Digital Signage (DSS); the Operational Support Systems (OSS) and the Business Support Systems (BSS).

5) *UNIFY*: UNIFY [91] is aimed at researching, developing and evaluating the means to orchestrate, verify and observe end-to-end service delivery from home and enterprise networks through aggregation and core networks to data centers. To this end, the project plans to develop an automated, dynamic service creation platform, leveraging a fine-granular service chaining architecture. They will also create a service abstraction model and a service creation language to enable dynamic and automatic placement of networking, computing and storage components across the infrastructure. Finally, they will develop a global orchestrator with optimization algorithms to ensure optimal placement of elementary service components across the infrastructure.

6) *T-NOVA*: T-NOVA [92] aims at promoting the NFV concept, by proposing an enabling framework, allowing operators not only to deploy VNFs for their own needs, but also to offer them to their customers, as value added services. For this purpose, T-NOVA leverages SDN and cloud management architectures to design and implement a management/orchestration platform for the automated provision, configuration, monitoring and optimization of Network Functions-as-a-Service (NFaaS) over virtualized network/IT infrastructures.

7) *CONTENT*: CONTENT [93] is an EU funded project aimed at offering a network architecture and overall infrastructure solution to facilitate the deployment of conventional cloud computing as well as mobile cloud computing. The main objectives of the project include:

(1) proposing a cross-domain and technology virtualization solution allowing the creation and operation of infrastructure slices including subsets of the network and computational physical resources, and (2) supporting dynamic end-to-end service provisioning across the network segments, offering variable QoS guarantees, throughout the integrated network.

Summary: To summarize, in Table IV we present all the projects giving their main objective, their focus with respect to NFV and related areas, and entities leading or funding them. All these projects are guided by the proposals coming out of the standardization described earlier, in particular ETSI, 3GPP and DMTF. It is interesting to observe that all the three industrial projects (ZOOM, OPNFV and OpenMANO) surveyed are focused on MANO. This underlines the importance of MANO in NFV. MANO is a critical aspect towards ensuring the correct operation of the NFVI as well as the VNFs. Just like the decoupled functions, NFV demands a shift from network management models that are device-driven to those that are aware of the orchestration needs of networks which do not only contain legacy equipment, but also VNFs. The enhanced models should have improved operations, administration, maintenance and provisioning focused on the creation and lifecycle management of both physical and virtualized functions. For NFV to be successful, all probable MANO challenges should be addressed at the current initial specification, definition and design phase, rather than later when real large scale deployments commence.

C. NFV Implementations

In order to demonstrate the possibility to implement the ideas proposed by NFV, and to determine performance char-

TABLE IV
SUMMARY OF NETWORK FUNCTION VIRTUALIZATION PROJECTS

	Project Type	Leader and/or Funding	Focus Areas	Main Objective
ZOOM	Association of SPs	TM Forum	NFV	Enable more rapid deployment of services by automating the provisioning process and modernizing OSS/BSS models.
OPNFV	Collaborative Project	Linux Foundation	NFV	Build an open source reference platform to advance the evolution of NFV.
OpenMANO	Vendor Project	Telefonica	SDN, NFV	Implementation of ETSI's MANO framework.
MCN	Research Project	European Union	SDN, NFV	Cloudify all components of a mobile network operation.
UNIFY	Research Project	European Union	NFV	Develop an automated, dynamic service creation platform, leveraging fine-granular service chaining.
T-NOVA	Research Project	European Union	SDN, NFV	Design and implement a MANO platform for NFV.
CONTENT	Research Project	European Union	Mobile Networks, Cloud	Providing a technology platform interconnecting geographically distributed computational resources that can support a variety of Cloud and mobile Cloud services.
OpenStack	Working Group	OpenStack Foundation	Cloud, NFV	Identify requirements needed to deploy, manage, and run telecom services on top of OpenStack.
OpenDaylight	Collaborative Project	Linux Foundation	SDN, NFV	Develop an open platform for SDN and NFV.

acteristics, a number of use cases for NFV, mostly based on those defined by ETSI [9], have already been implemented. These have mainly been based on implementing single virtual functions such as routing [94], Broadband Remote Access Server [95], policy server [96], deep packet inspection [97], EPC [98], [73], RAN [99], [100], [101], [102], monitoring [62], CPE [103], [104], [105], [106], [107], [108], GPRS [109] and access control [110], in cloud environments. All these originate from the research community. Perhaps not surprisingly, the biggest implementations have arisen from equipment vendors. In the remainder of this section, we introduce some key NFV implementations and products from industry.

1) *HP OpenNFV*: The HP OpenNFV [111] is a platform, based on HP's NFV Reference Architecture, upon which services and networks can be dynamically built. The HP NFV Reference Architecture is aligned towards providing solutions to each of the functional blocks defined in the ETSI architecture, as a starting point. The NFVI and VNFs parts of the architecture mainly include HP servers and virtualization products, while MANO is based on three solutions; NFV Director, NFV Manager, and Helion OpenStack. The NFV Director is an orchestrator that automatically manages the end-to-end service, by managing its constituent VNFs. It also performs global resource management, allocating resources from an appropriate pool based on global resource management policies. VNF managers are responsible for the VNFs lifecycle actions, for example, by deciding to scale in or out. It also includes a Helion OpenStack cloud platform for running VNFs.

2) *Huawei NFV Open Lab*: The Huawei NFV Open Lab [112] is aimed at providing an environment to ensure that NFV solutions and carrier grade infrastructure are compatible with

emerging NFV standards and with the OPNFV [86]. The lab is dedicated to being open and collaborative, expanding joint service innovations with partners, and developing the open eco-system of NFV to aggregate values and help customers achieve business success. They also plan to collaborate with the open source community to innovate on NFV technologies to provide use cases for multi-vendors inter-operability around NFVI, and VNF-based services.

3) *Intel Open Network Platform (Intel ONP)*: Intel ONP [113] is an ecosystem made up of several initiatives to advance open solutions for NFV and SDN. The initiatives are focused on Intel product development (such as the Intel ONP Server), participation in open source development and standardization activities and collaborations with industry for proof of concepts and trials.

The main result of the ONP so far is the Intel ONP Server. This is a reference architecture that integrates open-source and hardware ingredients optimized for SDN/NFV. It is aimed at enabling manageability by exposing health, state, and resource availability, for optimal workload placement and configuration. Its software stack consists of released open-source software based on the work done in community projects, including contributions provided by Intel. Some of the key open-source software ingredients forming the Intel ONP Server software stack are OpenStack, OpenDaylight, DPDK, Open vSwitch, and Linux KVM.

4) *CloudNFV*: CloudNFV [114] is an NFV, SDN and cloud computing platform resulting from cooperation between six companies (6WIND, CIMI Corporation, Dell, EnterpriseWeb, Overture Networks, and Qosmos). CloudNFV proposed their own NFV architecture [114] which is made up of 3 main elements: active virtualization, NFV orchestrator, and NFV Manager. Active virtualization is a data model which rep-

resents all aspects of services, functions and resources. The VNF orchestrator has policy rules, which, combined with service orders and the status of available resources, determines the location of the functions that make up the service as well as connections between them. The VNF Manager uses a data/resource model structured according to TMF rules and the concept of “derived operations” is used to manage VNFs. Derived operations are used to integrate the status of available resources with the resource commitments for functions of a given NFV service. The main difference between the ETSI NFV MANO and CloudNFV is that unlike the former, the latter considers both management and orchestration as applications that can run off a unified data model.

5) *Alcatel-Lucent CloudBand*: Alcatel-Lucent’s CloudBand [115] is a two-level platform implementing NFV. First, it includes nodes that provide resources like VMs and storage, and then, the CloudBand Management System which is the functional heart of the process. It operates as a work distributor that makes hosting and connection decisions based on policy, acting through cloud management APIs. Virtual functions are deployed using *recipes* that define *packages* of deployable components and instructions for their connection. The recipes can be used to set policies and determine how specific components are instantiated and then connected. The platform uses the Nuage SDN technology [116] and its related links to create an agile connection framework for the collection of nodes and functions, and to facilitate traffic management.

Alcatel-Lucent recently teamed with RedHat [117] such that the latter could fill the gaps required to use CloudBand and OpenStack to promote the inclusion of more NFV requirements in the OpenStack upstream and hence build a solution that is optimized for telco NFV environments. Within this collaboration, the CloudBand node uses the RedHat Enterprise Linux OpenStack platform as the VIM.

6) *Broadcom Open NFV*: The Broadcom Open NFV platform [118] is aimed at accelerating creation of NFV applications across multiple system on chip (SoC) [119] processors, and to allow system vendors to be able to migrate virtual functions between platforms based on various vendor solutions. Broadcom’s platform supports open API standards such as Linaro’s Open Data Plane (ODP)[120] to access acceleration components for scaling critical functionality and reducing time-to-market. The ETSI has recently accepted a VNF state migration and inter-operability proof of concept in which Broadcom is demonstrating an implementation of an EPC and migrating the virtual function state from operating on one instruction set architecture (ISA) to a different ISA.

7) *Cisco Open Network Strategy*: Cisco’s Open Network Strategy (OPN) [121] includes an Evolved Services Platform (ESP) and an Evolved Programmable Network (EPN). The ESP and EPN include a service orchestrator, a VNF manager, and a SDN controller, all of which are aimed at providing implementations for some of the functional blocks of ETSI’s MANO framework. The service orchestrator is responsible for providing the overall lifecycle management at the network service level. The VNF manager provides scalable, automated VNF lifecycle management, including the creation, provisioning, and monitoring of both Cisco and third-party VNFs. The

VNF manager is also responsible for the scale-up and scale-down of the VNFs based on dynamic and fluctuating service demands. It uses cloud-computing resource managers such as OpenStack and VMware at the VIM layer to configure and provision compute and storage resources across multi-vendor data center networks. Finally, the SDN Controller is responsible for connecting the virtualized services (a VNF or a set of chained VNFs) to the service provider VPNs, the Internet, or both. It is designed around open standards and APIs and uses a holistic systems-based approach to manage multi-vendor and multi-tenant data centers, and a common policy-based operating model to reduce costs.

8) *F5 Software Defined Application Services*: F5 Software Defined Application Services (F5 SDAS) [122], [123], [124] provides Layer 4-7 capabilities to supplement existing Layer 2-3 network and compute initiatives such as SDN. It enables service injection, consumption, automation, and orchestration across a unified operating framework of pooled resources. It is comprised of three key components: (1) The application service platform supports programmability of both control and data paths. It is extensible and enables new service creation. (2) The application services fabric provides core services such as scalability, service isolation, multi-tenancy, and integration with the network, and (3) Application services, which are the heart of F5 SDAS, are a rich catalog of services across the application delivery spectrum.

9) *ClearWater*: ClearWater [125] is an open source implementation of an IMS built using web development methods to provide voice, video and messaging services to users. It leans heavily on established design patterns for building and deploying scalable web applications, adapting these design patterns to fit the constraints of SIP and IMS. In particular, all components scale out horizontally using simple, stateless load-balancing. In addition, long-lived state is not stored on cluster nodes, avoiding the need for complex data replication schemes. Instead, long-lived state is stored in back-end service nodes using cloud-optimized storage technologies such as Cassandra. Finally, interfaces between the front-end SIP components and the back-end services use RESTful web services APIs. Interfaces between the various components use connection pooling with statistical recycling of connections to ensure load is spread evenly as nodes are added and removed from each layer.

Metaswitch [126] contributed the initial code base for the ClearWater project to software developers and systems integrators, and continues to drive the evolution of the code base.

10) *Overture Virtual Service Edge (vSE)*: Overture vSE [24] is an open carrier Ethernet platform for hosting VNFs at the service edge. It allows TSPs to instantly deploy on-demand VNFs at the customer premise. It combines carrier Ethernet access with the benefits of virtualization, openness and software-defined services. The result is a single platform for both services and network access, which allows for VNFs to be turned up, down, expanded and removed dynamically so that compute and storage resources are used only when needed. Additionally, it supports multiple wireline and wireless connections to the WAN, allowing access to all end

TABLE V
SUMMARY OF STATE-OF-THE-ART NFV IMPLEMENTATIONS

	Functionality	Platform	Driving Standards
HP OpenNFV	Open standards-based NFV reference architecture, labs as a sandbox in which carriers and equipment vendors can test vEPC.	OpenStack	ETSI
NFV Open Lab	Supports the development of NFV infrastructure, platforms and services.	OpenStack, OpenDaylight	ETSI
Intel ONP	Provides developers with a validated template for quickly developing and showcasing next-generation, cloud-aware network solutions.	OpenStack, OpenDaylight	3GPP or TMF
CloudNFV	Provides a platform for virtual network service creation, deployment, and management.	OpenStack	TMF and ETSI
Alcatel CloudBand	Can be used for standard IT needs as well as for CSPs who are moving mobile networks into the cloud.	Red Hat Linux OpenStack Platform	ETSI
BroadBand NFV	Migrate virtual functions between platforms based on various vendor solutions.		ETSI
Cisco ONS	Automated service delivery, improved network and data center use, fast deployment of personalized offerings.	OpenStack, OpenDaylight	ETSI
F5 SDAS	Extensible, context-aware, multi-tenant system for service provisioning	OpenStack, BIG-IP, BIG-IQ [122]	IETF, 3GPP, GSMA, ETSI, ONF
ClearWater	SIP-based call control for voice and video communications and for SIP-based messaging applications.	Apache Cassandra, Memcached	3GPP IMS, ETSI TS
Overture vSE	Host multiple VNFs in one box, Accelerate service creation, activation and assurance, Decrease inventory and management costs, Optimize service flexibility, Eliminate trucks rolls	Linux Overture Ensemble OSA [24], OpenStack	

customer locations.

The platform implements an Ethernet access as a VNF, and is based on a virtualization platform comprising a Linux KVM/QEMU hypervisor, an optimized virtual switch, and includes supports for OpenStack integration with another product - the Ensemble Service Orchestrator.

Summary: In Table V we summarize the different state-of-art implementations stating their functionality, the standards bodies they closely follow and platforms on which they run. It is worth remarking that although NFV is gaining momentum, it is still an emerging technology and solutions based on final specifications, and widespread deployments for end-users may take a few years to appear. As the survey above shows, many organizations are investing in and are willing to test NFV-based solutions. In addition, it can be observed from these early implementations and platforms, that two aspects re-appear in a big number of them: (1) the high focus on open source, and (2) the ability of current SDN and cloud technologies to support NFV.

VI. RESEARCH CHALLENGES

Even with all the anticipated benefits, and despite the immense speed at which it is being accepted by both academia and industry, NFV is still in early stages. There still remain important aspects that should be investigated and standard practices which should be established. This section discusses crucial research directions that will be invaluable as NFV matures.

A. Management and Orchestration

The deployment of NFV will greatly challenge current management systems and will require significant changes to the way networks are deployed, operated and managed. Such changes are required, not just to provide network and service solutions as before, but also to exploit the dynamism and flexibility made possible by NFV [127], [128]. It will likely lead to scenarios where functions that provide a service to a given customer are scattered across different server pools. The challenge then will be to have an acceptable level of orchestration to make sure that on a per service (or user) level, all the required functions are instantiated in a coherent and on-demand basis, and to ensure that the solution remains manageable [129].

ETSI is working on a MANO framework [18] required for the provisioning of VNFs, and the related operations, such as the configuration of the VNFs and the infrastructure these functions run on. In a related effort, Cloud4NFV [130], [131] has proposed an end-to-end management platform for VNFs, which is based on the ETSI architectural specification. Clayman et al. [132] describe an architecture based on an orchestrator that ensures the automatic placement of the virtual nodes and the allocation of network services on them supported by a monitoring system that collects and reports on the behaviour of the resources. NetFATE [133] proposes an orchestration approach for virtualized functions, taking into account the service chains needed by traffic flows and the desired QoE. In addition, other MANO frameworks and architectures have been proposed in [134], [135], [136], [137], [138], [139], [140].

TABLE VI
SUMMARY OF CHANGES IN ENERGY CONSUMPTION FROM VIRTUALIZING NETWORK FUNCTIONS

	Traffic (EXABYTES/MONTH)	Total Efficiency (MBITS/J)	Total Power (MWATTS)	Power Savings (MWATTS)	Cummulative Savings (2013 - 2018) GJ
Baseline Network	1,153.05	0.0328510	116,203	0.0	
Virtual EPC	1,153.05	0.0422222	92,159.8	24,044.1	5.0×10^9
Virtual CPE	1,205.11	0.0352130	113.500	2,703.63	5.5×10^9
Virtual RAN	1,227.88	0.0463708	89,599.5	26,604.4	7.5×10^9
Virtual Video CDN	810.22	0.0346562	80,029.3	36,174.6	7.5×10^9
Virtual Broadband Network Gateway	1,169.69	0.0333016	116.260	-76.794	-1.7×10^7
Virtual Provider Edge	1,151.91	0.0328255	116,180	22.9517	3.8×10^6

However, there are still some open issues. Current approaches are focused on NFV management, without considering the management challenges in SDN [141]. While traditional management approaches must be improved to accommodate each one of them, the demands for management are even higher in environments including both. In such cases, we no longer just need to create dynamic traffic flows, but the switching points (locations of functions) are also changing dynamically. Therefore, a complete management solution should combine requirements from both SDN and NFV.

In addition, support for inter-operability is a key requirement for NFV. However, looking at the ETSI MANO framework, most effort has been on defining intra-operator interfaces, without clear guidelines on inter-operability. This is why, while current vendor products are “based on the ETSI MANO framework”, most of them use custom models and/or representation for functions and services. Furthermore, the need for dynamism in function means that functions will likely be moved from one VM to another. This underscores the importance of a higher focus on possibilities of an availability monitoring mechanism as part of the end-to-end management solution. Finally, while the ETSI-proposed NFV MANO framework considers the management and orchestration requirements of both virtualized and non-virtualized functions via interfaces to traditional network management functions OSS/BSS, the relationship between them is yet to be fully defined [142]

B. Energy Efficiency

Since energy bills represent more than 10% of TSPs’ OPEX [143], reduced energy consumption is one of the strong selling points of NFV. The argument is that with the flexibility and ability to scale resource allocations up and down, as traffic demands ebb and flow, TSPs could potentially reduce the number of physical devices operating at any point, and hence reduce their energy bills. Yet, NFV will likely make data centers an integral part of telecommunication networks. According to an analysis in the SMARTer 2020 report from GeSI [144], the cloud, if it were a country, would rank 6th in the world in terms of its energy demand, and yet this demand is expected to increase by 63% by 2020 [145]. While some progress on energy efficient cloud computing has been

made, the fast growing energy needs of data centers continue to receive a lot of attention [146], [147]. Therefore, there is an urgent need to study whether NFV will meet its energy savings expectations, or whether—like the NFs—the energy consumption will just be transferred to the cloud.

China Mobile recently published [148] their experiences in deploying a Cloud Radio Access Network (C-RAN). One of the tests was performed on their 2G and 3G networks, where it was observed that by centralizing the RAN, power consumption could be reduced by 41% due to shared air-conditioning. In addition, Shehab et al. [149] analyzed the technical potential for energy savings associated with shifting U.S. business software to the cloud. The results suggested a substantial potential for energy savings. In fact, the authors noted that if all U.S. business users shifted their email, productivity software, and CRM software to the cloud, the primary energy footprint of these software applications could be reduced by as much as 87%.

In order to determine the possible effect of energy consumption on the evolution to VNFs, Bell Labs has recently extended its G.W.A.T.T. tool [143]. The tool is able to show the effect of virtualizing different network functions based on forecasts for traffic growth. G.W.A.T.T. divides the network into six domains (Home & Enterprise, Access & Aggregation, Metro, Edge, Core and Service Core & Data Centers). Each network domain can be edited to select different network models and technologies and hence analyze its energy impact. Based on the tool’s default settings and using EPC network models for 2015, the tool shows that total network energy efficiency is 0.0422222 MBITS/J, total energy consumption is 92,159.8 MWATTS, and that the energy savings resulting from virtualizing the EPC would be 24,044.1 MWATTS. For the same use case, the tool showed that the total energy savings over a five year period (using 2013 as baseline) would be 5.0×10^9 GJ, and that the energy efficiency of the core network 1.86393 MBITS/J. The results for some other NFV use cases, including those for the baseline network⁶ are summarized in Table VI. However, while the tool is an important step in attaching numbers to the energy savings expected from NFV, it can still be improved. In particular, it does not yet have a

⁶A baseline network is one where all functions are run in physical equipment, using the tool’s default technologies and settings.

detailed technical documentation. For example, Cisco's visual networking index [150] forecasts that annual global IP traffic will reach 1000 exabytes in 2016. Based on this, the (monthly, 2015) traffic values in Table VI seem to be too high, yet it is currently not possible to know how these values are derived.

Therefore, we expect that the energy efficiency of cloud based NFs will continue to receive attention. NFV will put InPs under even more pressure to manage energy consumption [138] not to only to cut down energy expenses, but also to meet regulatory and environmental standards. Topics with regard to energy efficient hardware which could allow reductions in CPU speeds and partially turning off some hardware components, more energy-aware function placement, scheduling and chaining algorithms, will be important. An example could be to track the cheapest prices for energy costs and adapt the network topology and/or operating parameters to minimize the cost of running the network [60]. However, all these should be carefully considered to ensure that there is a balance in the trade-off between energy efficiency and function performance or service level agreements.

C. NFV Performance

The concept of NFV is to run NFs on industry standard servers. This means that server providers should produce equipment without knowledge of the characteristics of functions that could run on them in future. In the same way, VNF providers should ensure that the functions will be able to run on commodity server. This raises the question of whether functions run on industry standard servers would achieve a performance comparable to those running on specialized hardware, and whether these functions would be portable between the servers [60]. Finding answers to these questions has been another focus of the ETSI, and resulted into a "Performance & Portability Best Practises" specification [151]. The specification gives performance test results on NFV use cases such as DPI, C-RAN, BRAS, etc. The results proved that if "best practices were followed" it was not only possible to achieve high performance (upto 80 Gbps for a server) in a fully virtualized environment, but that the performance was predictable, consistent and in vendor-agnostic manner, leveraging features commonly available in current state-of-the-art servers [60].

In a related effort, results from China Mobile's C-RAN deployment [148] indicated that the Common Public Radio Interface (CPRI) [152] over a wavelength-division multiplexing (WDM) front-haul transport solution gives ideal performance, with no impact on radio performance. The tests also verified the feasibility of using a general purpose platform (GPP) and the NFV implementation. In particular, a GPP based C-RAN prototype with the ability to support as many as 90 TD-LTE carriers, 15 FDD-LTE carriers and 72 GSM carriers was developed. The prototype demonstrated a similar level of performance to the traditional DSP/FPGA based systems.

However, performance at high speeds is an issue even in non-virtualized NFs [29], [153]. Therefore, techniques such as hardware acceleration will also be important for NFV. In fact, hardware acceleration has been shown to improve the

performance of some VNFs. Ge et. al [28] determine that for some functions (e.g. DPI, Dedup and NAT), industry standard servers may not achieve the required levels of performance. From the authors' tests, a virtualized Dedup could only achieve 267 Mbps throughput in each core at most. It was also proved by Yamazaki et. al [154] who reported achieving a better performance and energy efficiency by deploying a virtualized DPI on Application Specific Instruction-set Processor (ASIP) rather than commodity servers.

Therefore, there are some high performance NFs that may be difficult to virtualize without degradation in performance. While hardware acceleration may be used for such functions, such specialization is against the concept of NFV which aims at high flexibility. There should be defined ways of managing the trade-off between performance and flexibility. It will also be appropriate to have phased migrations to NFV where those functions that have acceptable performance are virtualized first and allowed to run alongside unvirtualized or physical ones.

D. Resource Allocation

To achieve the economies of scale expected from NFV, physical resources should be used efficiently. It has been shown that default deployment of some current use cases may result in sub-optimal resource allocation and consumption [10].

This calls for efficient algorithms to determine on to which physical resources (servers) network functions are *placed*, and be able to move functions from one server to another for such objectives as load balancing, energy saving, recovery from failures, etc. The task of placing functions is closely related to virtual network embedding [155] and virtual data center embedding [156] and may therefore be formulated as an optimization problem, with a particular objective. Such an approach has been followed by [157], [158], [159], [160], [161].

For example, Basta et. al [157] investigated the influence of virtualizing the S-GW and P-GW functions on the transport network load and data-plane delay. For these two functions, the authors showed differences in performance (of upto 8 times) when the functions were either fully virtualized and when their data and control planes were separated. The authors proposed a model for placing the functions in a way that minimizes the network load overheads introduced by the SDN control plane interactions. In addition to placement, Mehraghdam [160] proposes a model for formalizing the chaining of NFs. To this end, for each service deployment request, their approach constructs a VNFFG which is then mapped to the physical resources, considering that the network resources are limited and that functions have specific requirements. The mapping is formulated as a Mixed Integer Quadratically Constrained Program (MIQCP). The authors concluded that in order to obtain efficient use of resources, the placement of functions should be different according to the desired placement objective (i.e. remaining data rate, latency, number of used network nodes). Finally, Moens et. al [158] formulate the placement problem as an Integer Linear Program (ILP) with an objective of allocating a service chain onto the physical network minimizing the number of servers used.

However, when formulated as an optimization problem, function placement and chaining would reduce to a binary integer program, which is NP-Hard [162], and hence intractable for big instances of the problem. This calls for heuristics such as those proposed in [163], [164], [165], [132]. For example, Xia et. al [163] formulate the placement and chaining problem as binary integer programming (BIP), and propose a greedy heuristic to improve computational efficiency. The proposed greedy algorithm first sorts VNFs according their resource demand, and thereafter, VNFs with the highest resource demands are given priority for placement and chaining.

In addition, NFV systems should allow for one or a group of VNFs to be migrated to disparate physical servers. The physical servers may be in different InP domains, and hence use different tunneling addresses or be managed by different protocols. This does not only call for efficient algorithms to determine where the functions can be moved, but will also require comprehensive management of function and server states, as well as maintain communications. ViRUS [166] allows the runtime system to switch between blocks of code that perform equivalent functionality at different QoS levels when the system is under stress, while [167] presents a model that can be used to derive some performance indicators, such as the whole service downtime and the total migration time, so as to make function migration decisions.

Finally, to ensure scalable NFV implementations, functions should only be allocated the resources they need. Contrary to most current proof of concept implementations, it is not feasible to deploy a VM per subscriber or per function as the resulting VM footprint would be too high. This is because each VM is like a computer running its own operating system, and is meant to be isolated from other VMs and hence independent on a network level. This approach could become wasteful of resources for two reasons: (1) some of the functions such as DHCP in a CPE are so light that they would not justify a dedicated operating system on the scale of multiple functions per user, (2) some functions do not need to be strictly isolated from each other. Therefore, depending on the requirements of a given function, containers could be a more efficient way to use resources. Linux containers [168] are an alternative to dedicated VMs in which a Docker [169] may be used to achieve the automated resource isolation and namespacing which allows for partitioning of memory, network, processes etc. The use of containers avoids the overhead of starting and maintaining virtual machines since they do not require a complete duplication of an operating system. Using containers could lead to up to a 30% savings in server costs to support the same number of virtual hosting points [170].

Moreover, even if given functions must utilize the same resources in a VM's operating system, it is possible to use scheduling techniques to allow the functions to share the resources. To this end, the proposals in [171], [172], [173] formulate the problem as a Resource Constrained Project Scheduling Problem (RCPSP) [174] and solve it using a job shop scheduling approach [175]. Specifically, Mijumbi et. al [171] formulate an online VNF mapping and scheduling problem and propose a set of greedy algorithms and a Tabu Search (TS) [176] heuristic for solving it. The greedy algorithms

perform the mapping and scheduling of VNFs based on a greedy criterion such as available buffer capacity for the node or the processing time of a given VNF on the possible nodes, while the TS algorithm starts by creating an initial solution randomly, which is iteratively improved by searching for better solutions in its neighborhood.

In addition, existing scheduling tools such as Google's Borg [177] and Apache Mesos [178] may be considered for scheduling of VNFs. Borg uses task-packing, over-commitment, and machine sharing with process-level performance isolation to run multiple jobs, from many applications, across a number of clusters. Users of the Borg system submit jobs consisting of one or several tasks that are run from the same executable. The scheduler in Borg monitors queues and schedules jobs considering the resources available on individual machines. The jobs may have requirements such as CPU and OS. However, unlike the functions in NFV, the tasks in Borg are run directly on hardware not in a virtualized environment. In addition, while Borg may have the scalability (cells usually contain 10K servers) that would be required in an NFV environment, it would have to be improved to meet carrier class requirements. For example, unlike the functions that make up a service in NFV, the tasks considered in Borg do not have ordering requirements. Finally, a task start up latency of 25s, and the 4 nines (99.99%) availability that Borg is able to give may need to be enhanced for NFV.

Therefore, it can be observed that there are still many open areas with regard to how physical resources are shared among the VNFs. First of all, the results in each of the above areas may still be improved. In particular, the efficiency and applicability of containers needs to be studied more, just like there is need to study and propose more efficient function scheduling algorithms. In addition, given the dynamic requirements of NFV, there is need for resource allocation proposals that are able to find solutions online, consider multi-domain and distributed VNFs [179], [180], network survivability [181], dynamic resource management [182] etc.

E. Security, Privacy and Trust

Despite the enormous potential of cloud computing, consumer uncertainty and concern regarding issues of privacy, security and trust remain a major barrier to the switch to cloud models [183]. Therefore, cloud privacy issues will be among the key concerns for TSPs if they have to move to public clouds. Because the functions to be virtualized represent subscriber services, personally identifiable information may be transferred to the cloud. This will present unique challenges especially as the functions will be distributed, making it hard to know where this data is and who has access to it. In the case where the functions are deployed in third party clouds, users and Telecom service providers would not have access to the physical security system of data centers. Even if the service providers do specify their privacy and security requirements, it may still be hard to ensure that they are fully respected.

Emphasizing its importance, ETSI constituted a security expert group to focus on this concern. The group started by identifying potential security vulnerabilities of NFV and

TABLE VII
POTENTIAL SECURITY THREATS IN NFV [184]

Security Threat
Topology Validation & Enforcement
Availability of Management Support Infrastructure
Secured Boot
Secure Crash
Performance Isolation
User/Tenant Authentication, Authorization and Accounting
Authenticated Time Service
Private Keys within Cloned Images
Back-Doors via Virtualized Test & Monitoring Functions
Multi-Administrator Isolation

establishing whether they are new problems, or just existing problems in different guises [184]. The evaluation confirmed that indeed NFV creates new security concerns as shown in Table VII. After identifying the possible threats, the group proposed some solutions. In particular, they have provided a security and trust guidance that is unique to NFV development, architecture and operation [30]. However, this does not consist of prescriptive requirements or specific implementation details.

However, it was noted that while solutions for these threats are available, there are currently no processes to take advantage of these solutions and, once in place, they will add procedural complexity [60], [184]. Moreover, for some of the threats (such as topology validation, network performance isolation and multi-administrator isolation), the group determined that solutions are not yet available [60]. As NFV gets deployed and more important functions virtualized, we can expect it to attract even more security and privacy threats. More than ever, there will be threats based on data interception (whether lawful or otherwise). Therefore, security, privacy and trust are other important research directions in NFV.

F. Modeling of Resources, Functions and Services

NFV's potential is based on its ability to deliver high levels of automation and flexibility. However, the resources and functions in NFV will be provided by different entities. Therefore, the availability of well understood, open and standardized descriptors for these multi-vendor resources, functions and services will be key to large-scale NFV deployments. Models should consider both initial deployment as well as lifecycle management - reconfiguration. As part of the MANO specification [18], the ETSI provided a possible set of models that may be useful in NFV. These include OVF, TOSCA, YANG and SID. OVF was introduced in section V-A5. In what follows, we introduce the other three models.

1) *Topology & Orchestration Standard for Cloud Application (TOSCA)*: TOSCA [185] is an OASIS standard language to describe a topology of cloud based web services, their components, relationships, and the processes that manage them. It describes what is needed to be preserved across

service deployments in different environments to enable interoperable deployment of cloud services and their management when the applications are ported over alternative cloud environments [18]. TOSCA may be used for VNF definition, node monitoring and active policies like healing and scaling.

2) *NETCONF/YANG*: NETCONF [186] is a protocol defined by the IETF to “install, manipulate, and delete the configuration of network devices”. NETCONF operations are realized on top of a Remote Procedure Call (RPC) [187] layer using an XML encoding and provide a basic set of operations to edit and query configuration on a network device. NETCONF is based on the YANG data modeling language. YANG is used to model both configuration and state data of network elements. Furthermore, YANG can be used to define the format of event notifications emitted by network elements and it allows data modelers to define the signature of remote procedure calls that can be invoked on network elements via the NETCONF protocol.

3) *Information Framework (SID)*: SID [188] is a component of TM Forum's Framework aimed at providing an information model and common vocabulary for all the information shared among things of interest (entities) to an enterprise such as customer, location and network element, and relationships (associations) between these entities, such as a network element is situated at a location. Entities are further characterized by facts (attributes) that describes them and their behavior (operations) that describe how the entities work. SID was originally based on Unified Modeling Language (UML) [189], but was extended to include XML Schema Definition (XSD) representations.

Discussion: Table VIII summarizes the information and data modeling possibilities for NFV. All the models defined above have relatively wide adoption, and may therefore be considered for modeling of resources and functions in NFV. For example, to enable simple and scalable gradual deployment of VNFs and other NFV concepts, VNFs need to co-exist with traditional non NFV-based NFs. To provide an integration with existing OSS/BSS systems, end-to-end network services that include VNFs or VNF Forwarding Graphs may be able to be mapped to the SID service model [18].

However, these models were not initially developed with explicit considerations for some of the more specific requirements expected by NFV deployments and can therefore only be used as starting points and should continue to evolve for this purpose. For example, portability of data models and support for federated services have been identified [190], [191] as outstanding improvements for TOSCA. TOSCA also needs improvement to support run-time management of services. With regard to NETCONF/YANG, there is need to improve them to be able to cope with situations when multiple administrators (multi-domain environment) are present [192]. A lot of work is ongoing to extend some of the models for NFV. For example, SID has been extended using the ZOOM information model [193] to define four concepts (VirtualResource, NetworkFunction, NetworkService, and Graph) aimed at modeling NFV-based systems. In addition, The TOSCA TC recently formed a workgroup focused on creating a “TOSCA Simple Profile

TABLE VIII
SUMMARY OF CHOICE OF INFORMATION AND DATA MODELS FOR NFV

	OVF	TOSCA	NETCONF/YANG	SID
Organization	DTMF	OASIS	IETF	TMF
Objective	Describe the packaging and distribution of software to be run in one or more VMs	Standardize interaction between cloud platforms and to provide cross-platform compatibility for applications and services.	Install, manipulate, and delete the configuration of network devices	Identify the business entities that play role in the business processes of a telecommunications service provider.
Roots	Server Virtualization	IT applications	Configuration of network services, devices	Identification and modelling of TSP business processes
Data Model	CIM	YAML / XML	YANG	UML
Applicability to NFV	Capturing some or all of the VNF package and/or VDU descriptor	Function template modelling for deployment	Runtime configuration of VNFs	May be used on the interface between OSS/BSS and NFV MANO and also on the interface between OSS/BSS and EM.
Encoding	XSD		XML	XSD
Language		Declarative, Imperative	Procedural (Yang)	
NFV Project Using Model		ClearWater, Multi-vendor PoCs [199], ExperiaSphere [200]	ONF	ZOOM
Research Challenges	Support for runtime management, possibly more stringent requirements of VNFs	Portability of data models, support for federated services, runtime management	Capability to support easier modelling/deployment template designs, protocol independence	As SID is fundamentally an information model with a defined data model, it still lacks protocol and implementation details

for NFV.”

As the models continue to improve, it may be important to have solutions that combine them so as to avoid some of their disadvantages. For example, A TOSCA template can install a virtual router, but it cannot subsequently create/modify/delete configuration on demand on the same router during run-time. Therefore, fulfilling VNF requirements requires more than TOSCA. In the same way, YANG is designed for writing machine readable schema, and is hence difficult to use for design of templates for initial service deployment. In this case, TOSCA may be combined with NETCONF/YANG where the file-based templates in TOSCA may be used for deploying VNFs on cloud infrastructure, while NETCONF can be used to provide a runtime API both for configuring VNFs after they have been installed, bringing VNFs to a state of operational readiness, and while they are running in the cloud, fulfilling the service requirements of a particular customer [142].

G. Research Directions in Selected NFV Use Cases

1) *The Internet of Things*: Like NFV, the Internet of Things (IoT) [194] paradigm has recently drawn a lot of industrial attention. The IoT is a network of physical objects or “things” into which sensors with unique identifiers are embedded. Such sensors may collect and transfer various kinds of (big) data over a network without requiring human-to-human or human-to-computer interaction. Inevitably, by networking zillions of devices, the IoT will lead to networks of unbelievable scale and complexity with tremendous implications on network management. It will lead to security, scalability and resource management challenges in networks that should simultaneously transport, process and act on this data in real time.

NFV has been proposed as a key enabler of the IoT [195], [196]. The idea in [195] is to limit the functionality embedded

in deployed sensors, and provide virtualized functions such as security, intelligence, computation and storage to the devices. These would take advantage of the scalable distribution capabilities of NFV as well as the configuration flexibility of SDN. On the other hand, Omnes et. al [196] propose multi-layered IoT architecture involving SDN and NFV, and illustrate how the proposed architecture is able to cope with some of the challenges in IoT.

However, there are serious questions on the management of big amounts of IoT-generated data with better network efficiency. It is therefore critical to study efficient ways of transporting (big) data over such softwarized networks, and whether current cloud data management applications such as Hadoop and Cassandra would be able to support the real time requirements in such environments.

2) *Information-Centric Networking*: Motivated by the fact that the Internet is increasingly used for information dissemination rather than for pair-wise communication between end hosts, Information-Centric Networking (ICN) [197] has emerged as a promising candidate for the architecture of the Future Internet. ICN addresses named data rather than named hosts. This way, content distribution is implemented directly into the network fabric rather than relying on the complicated mapping, availability, and security mechanisms currently used to map content to a single location.

The separation between information processing and forwarding in ICN is related to both the decoupling of functions from devices in NFV, and to the decoupling of control from data plane in SDN. While the relationship between NFV, SDN and cloud computing has already received some attention, that between NFV and ICN has not. Yet, ICN may be used in NFV to determine the best position to place network functions. For example, Arumaithurai et al. [198] propose a *function-centric service chaining* (FCSC) approach

TABLE IX
SUMMARY OF STATE-OF-THE-ART AND RESEARCH CHALLENGES

Challenge	Description	Reference	Contribution / Objective	Research Opportunities
Management and Orchestration	ETSI MANO Framework	[18]	Specifies a management and orchestration framework for NFV	Traffic and function monitoring, inter-operability and interfacing, programmability and Intelligence, distributed management, combined management of cloud, SDN and NFV, autonomic (self) management technologies in NFV (e.g., processing of alarms)
	Vendor Products	[111], [114], [115], [121]	Vendor specific products for different components or specifications of the ETSI MANO framework	
	Projects	[85], [86], [88], [91], [92], [130], [131]	Implementation and/or proposals based on the ETSI MANO framework	
	Research Papers	[132], [133], [134], [135], [136], [137], [138], [139], [140]	Managements and orchestration frameworks and architectures	
NFV Performance	ETSI Performance & Portability Best Practices	[151]	Defines the "best practices" that need to followed to obtain acceptable performance in NFV. Also gives performance test results on on NFV use cases such as DPI, C-RAN, BRAS, etc	More studies on the applicability of hardware acceleration to some NFs, and on the resulting trade-off between performance and flexibility
	Practical Measurements	[148]	experiences in deploying a C-RAN on a 2G and 3G network	
	Hardware Acceleration	[26], [28], [29], [154]	Various proposals for applying hardware acceleration to enhance the performance of some VNFs such as DPI, dedup and NAT	
Energy Efficiency	Practical Measurements	[148], [149]	Measurements on the effect of transferring network and user functions to the cloud	Still limited number of real world deployments to give actual vales, energy efficient hardware, energy-aware function placement chaining, consideration of inter-data center communications
	Simulation	[143]	Vendor tool that simulates possible energy saving resulting from NFV	
Resource Allocation	Placement	[157], [158], [159], [160], [161]	Deciding the optimal placement of functions in the operator's network or the cloud, following specific functions requirements and resource constraints	Use of containers, function scheduling, multi-domain function placement and chaining, survivability of VNFs in case of network failures, dynamic resource allocation (scaling up and down)
	Migration	[166], [167]	Allow for one or a group of VNFs to be migrated to disparate physical servers	
	Scheduling	[171], [172], [173]	Allow multiple VNFs to be hosted in a single VM and schedule their efficient utilization of resources	
Security, Privacy, Trust	ETSI Security Problem Statement	[184]	Defines the security, trust and privacy threats in NFV	Topology validation, network performance isolation, multi-administrator isolation, data interception
	ETSI Security Guidance	[30]	Provides guidance on how security, privacy and trust may be achieved in NFV.	
NFV Use Cases	IoT	[195], [196]	Architecture combining NFV and IoT, and an application scenario involving a virtualized sensor function	Monitoring and metering of carrier-scale virtualized networks. Application of big data approaches, ICN-based placement of VNFs, proof of concepts and implementations involving chains of VNFs
	ICN	[198]	exploits ICN to provide flexibility and dynamism in placing VNFs	
	ETSI Use Cases	[62], [73], [94], [95], [96], [97], [98], [100], [101], [102], [103], [104], [105], [106], [107], [108], [109], [110]	Implementation, demonstrations and proofs of concepts based on the ETSI use cases	

which exploits ICN to provide flexibility and dynamism in placing VNFs.

Summary: In Table X, we summarize the state-of-art in each of the identified research challenges, as well as specific open questions in each one of them. We have noted that despite the significant and rapidly increasing activity on NFV, there are still major gaps especially with regard to standardization that may slow down NFV deployment and undermine the possibilities to fulfill its anticipated business case. While the ETSI-defined reference architecture covers most of the aspects needed to operationalize NFV, current specifications are still too general to envelope all the essential pillars of required evolution such as inter-operability, legacy support, and management of both legacy and NFV-based systems.

For example, currently, different vendors depend on different languages to model resources and functions in NFV. TOSCA has been used in modeling services for a multi-vendor E2E proof of concept [199], for ClearWater, and ExperiaSphere [200]. On the other hand, the descriptors in the HP NFV Director are not based on TOSCA, and ONF has chosen YANG as the modeling language. Similar examples can be given for vendor implementations of the NFVI, VNFs and MANO. This could result into inter-operability challenges where vendor-specific Command Line Interfaces (CLIs) require manual configuration or expensive integration by service providers themselves or systems integrators with their own proprietary tools and equipment-specific adapters. Therefore, though there are many options for modeling of functions and resources, the techniques remain generally in their infancy.

With regard to performance, most current PoCs are based on a rather limited list of use cases proposed by the ETSI. While these PoCs are important to prove technical principles unique to NFV, they do not give a complete view of performance and benefits for a wide range of end-to-end services. Finally, research on possible enablers of NFV such as ICN, and on the application areas such as IoT are still largely unexplored.

VII. CONCLUSION

Due to user demands for real-time, on-demand, online, inexpensive, short-lived services, TSPs have been forced to look for new ways of delivering these services in ways that are agile, and with OPEX and CAPEX savings. NFV has emerged as a possible approach to make network equipment more open, and hence allow TSPs to become more flexible, faster at service innovations and reduce operation & maintenance (O&M) costs. It is clear that NFV, together with the closely related and complementary fields of SDN and cloud computing may be big parts of future telecommunication service provision.

In this paper, we introduced NFV, described its architecture as defined by ETSI, proposed a reference business model, and explored important design considerations. We then compared NFV with closely related fields, SDN and cloud computing, discussing current research for combining them. We have also presented major specification and standardization efforts, research projects, commercial products and early NFV proof

of concept implementations. Finally, we discussed the key research areas that will be pivotal to the success of NFV as well as to its application to ICN and IoT, and summarized the findings of the survey. We believe that before these areas are explored, TSPs who deploy NFV may end up being reliant on vendor-proprietary solutions to solve these gaps, which would be against the original objective of NFV.

We have noted that many current NFV solutions, especially from the industry, have been mainly about pooling vendor specific resources hosted in a cloud rather than real support for flexibility, inter-operability, integrated management, orchestration and service automation all of which are core requirements for NFV. It is expected that such implementations will continue to increase before NFV gets completely standardized. As NFV moves from labs and PoCs to trials and commercial deployments, vendors are investing significant resources to develop these NFV solutions. It is therefore urgent for specification and standardization bodies to complete specifications before it becomes too late for the standards to change or influence what has already been deployed.

ACKNOWLEDGMENT

The authors are indebted to the Editor-in-Chief for coordinating the review process, and to the anonymous reviewers for their insightful comments and suggestions. This work was partly funded by FLAMINGO, a Network of Excellence project (318488) supported by the European Commission under its Seventh Framework Programme, and project TEC2012-38574-C02-02 from Ministerio de Economía y Competitividad.

REFERENCES

- [1] J. Wu, Z. Zhang, Y. Hong, and Y. Wen, "Cloud radio access network (C-RAN): a primer," *Network, IEEE*, vol. 29, no. 1, pp. 35–41, Jan 2015.
- [2] China Mobile Research Institute, "C-RAN: The Road Towards Green RAN. White Paper. Version 2.5." October 2011.
- [3] B. Han, V. Gopalakrishnan, L. Ji, and S. Lee, "Network function virtualization: Challenges and opportunities for innovations," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 90–97, Feb 2015.
- [4] R. Guerzoni, "Network Functions Virtualisation: An Introduction, Benefits, Enablers, Challenges and Call for Action. Introductory white paper," in *SDN and OpenFlow World Congress*, June 2012.
- [5] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV 002 V1.2.1: Network Functions Virtualisation (NFV); Architectural Framework," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01_02_01_60/gs_NFV002v010201p.pdf, December 2014.
- [6] ETSI, "European Telecommunications Standards Institute, Industry Specification Groups (ISG) - NFV," <http://www.etsi.org/technologies-clusters/technologies/nfv>, 2015, Accessed: June 03, 2015.
- [7] ETSI Industry Specification Group (ISG) NFV, "ETSI Group Specifications on Network Function Virtualization. 1st Phase Documents," <http://docbox.etsi.org/ISG/NFV/Open/Published/>, January 2015.
- [8] "The 3rd Generation Partnership Project (3GPP)," <http://www.3gpp.org/about-3gpp/about-3gpp>, 2015, Accessed: February, 10 2015.
- [9] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV 001 V1.1.1: Network Function Virtualization. Use Cases," www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf, October 2013.
- [10] P. Veitch, M. J. McGrath, and V. Bayon, "An instrumentation and analytics framework for optimal and robust NFV deployment," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 126–133, Feb 2015.
- [11] H. Hawilo, A. Shami, M. Mirahmadi, and R. Asal, "NFV: state of the art, challenges, and implementation in next generation mobile networks (vEPC)," *Network, IEEE*, vol. 28, no. 6, pp. 18–26, Nov 2014.

- [12] D. Cotroneo, L. De Simone, A. Iannillo, A. Lanzaro, R. Natella, J. Fan, and W. Ping, "Network Function Virtualization: Challenges and Directions for Reliability Assurance," in *Software Reliability Engineering Workshops (ISSREW), 2014 IEEE International Symposium on*, Nov 2014, pp. 37–42.
- [13] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-Defined Networking: A Comprehensive Survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, Jan 2015.
- [14] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: state-of-the-art and research challenges," *J. Internet Services and Applications*, vol. 1, no. 1, pp. 7–18, 2010.
- [15] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV 003 V1.2.1: Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/003/01.02.01_60/gs_NFV003v010201p.pdf, December 2014.
- [16] Q. P and N. T, "Service Function Chaining Problem Statement. Internet-Draft draft-ietf-sfc-problem-statement-10," Active Internet-Draft, IETF Secretariat, Tech. Rep., August 2014.
- [17] R. Mijumbi, "Self-managed Resources in Network Virtualization Environments," Ph.D. dissertation, Technical University of Catalunya, Barcelona, Spain, November 2014.
- [18] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV-MAN 001 V1.1.1: Network Functions Virtualisation (NFV); Management and Orchestration," http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf, December 2014.
- [19] —, "ETSI GS NFV-INF 004 V1.1.1: Network Functions Virtualisation (NFV); Infrastructure; Hypervisor Domain," http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/004/01.01.01_60/gs_NFV-INF004v010101p.pdf, January 2015.
- [20] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862 – 876, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128609003387>
- [21] F. Samuel, M. Chowdhury, and R. Boutaba, "PolyViNE: policy-based virtual network embedding across multiple domains," *Journal of Internet Services and Applications*, vol. 4, no. 1, p. 6, 2013.
- [22] Harvey, A. F. et. al, "DMA Fundamentals on Various PC Platforms. Application Note 011, National Instruments," <http://cires.colorado.edu/jimenez-group/QAMSRsources/Docs/DMAFundamentals.pdf>, April 1991.
- [23] W. Peterson and D. Brown, "Cyclic Codes for Error Detection," *Proceedings of the IRE*, vol. 49, no. 1, pp. 228–235, Jan 1961.
- [24] "Overture 65vSE: Open Platform for Virtualization at the Service Edge," <http://www.overturenetworks.com/products/overture-65vse>, 2015, Accessed: February, 06 2015.
- [25] S. Byma, J. Steffan, H. Bannazadeh, A. L. Garcia, and P. Chow, "Fpgas in the cloud: Booting virtualized hardware accelerators with openstack," in *Field-Programmable Custom Computing Machines (FCCM), 2014 IEEE 22nd Annual International Symposium on*, May 2014, pp. 109–116.
- [26] Z. Bronstein, E. Roch, J. Xia, and A. Molkho, "Uniform handling and abstraction of nfv hardware accelerators," *Network, IEEE*, vol. 29, no. 3, pp. 22–29, May 2015.
- [27] J. DiGiglio and D. Ricci, "High Performance, Open Standard Virtualization with NFV and SDN. Joint Technical White Paper by Intel Corporation and Wind River," http://www.windriver.com/whitepapers/ovp/ovp_whitepaper.pdf, 2015, Accessed: June, 05 2015.
- [28] X. Ge, Y. Liu, D. H. Du, L. Zhang, H. Guan, J. Chen, Y. Zhao, and X. Hu, "OpenANFV: Accelerating Network Function Virtualization with a Consolidated Framework in Openstack," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 353–354.
- [29] L. Nobach and D. Hausheer, "Open, elastic provisioning of hardware acceleration in nfv environments," in *Networked Systems (NetSys), 2015 International Conference and Workshops on*, March 2015, pp. 1–5.
- [30] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV-SEC 003 V1.1.1: Network Functions Virtualisation (NFV); NFV Security; Security and Trust Guidance," http://www.etsi.org/deliver/etsi_gs/NFV-SEC/001_099/003/01.01.01_60/gs_NFV-SEC003v010101p.pdf, December 2014.
- [31] ETSI ISG NFV, "ETSI GS NFV-REL 001 V1.1.1: Network Functions Virtualisation (NFV); Resiliency Requirements," http://www.etsi.org/deliver/etsi_gs/NFV-REL/001_099/001/01.01.01_60/gs_NFV-REL001v010101p.pdf, January 2015.
- [32] Andreas Lemke, Alcatel Lucent, "Why service providers need an NFV platform: Strategic White Paper," January 2015.
- [33] ATIS NFV, "Alliance for Telecommunications Industry Solutions, Network Functions Virtualization Forum," <http://www.atis.org/NFV/index.asp>, January 2015.
- [34] Allot, "Evolution of Network Service Enablement Utilizing Network Functions Virtualization (NFV). Technical White Paper," Allot Communications, Tech. Rep., July 2013.
- [35] M. Peter and G. Timothy, "The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technology," <http://www.nist.gov/itl/cloud/>, National Institute of Standards and Technology(NIST) Special Publication 800-145, Tech. Rep., September 2011.
- [36] Distributed Management Task Force, "Cloud Management Working Group (CMWG)," <http://www.dmtf.org/standards/cmwg>, February 2015.
- [37] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1355734.1355746>
- [38] D. R. Mauro and K. J. Schmidt, *Essential SNMP, Second Edition*. O'Reilly Media, Inc., 2005.
- [39] G. Wedge and L. Barbara, "Carrier-Grade: Five Nines, the Myth and the Reality," *Pipeline Publications*, vol. 3, no. 1, June 2006.
- [40] K. V. Vishwanath and N. Nagappan, "Characterizing Cloud Computing Hardware Reliability," in *Proceedings of the 1st ACM Symposium on Cloud Computing*, ser. SoCC '10. New York, NY, USA: ACM, 2010, pp. 193–204. [Online]. Available: <http://doi.acm.org/10.1145/1807128.1807161>
- [41] "Network Equipment-Building System (NEBS)," http://en.wikipedia.org/wiki/Network_Equipment-Building_System, 2015, Accessed: February, 07 2015.
- [42] M. Scharf, T. Voith, W. Roome, B. Gaglianella, M. Steiner, V. Hilt, and V. Gurbani, "Monitoring and abstraction for networked clouds," pp. 80–85, Oct 2012.
- [43] A. Mandal, Y. Xin, I. Baldine, P. Ruth, C. Heerman, J. Chase, V. Orlikowski, and A. Yumerefendi, "Provisioning and Evaluating Multi-domain Networked Clouds for Hadoop-based Applications," in *Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on*, Nov 2011, pp. 690–697.
- [44] D. Lopez, "Network functions virtualization: Beyond carrier-grade clouds," in *Optical Fiber Communications Conference and Exhibition (OFC), 2014*, March 2014, pp. 1–18.
- [45] Henrik, Basilier and Marian, Darula and Joe, Wilke, "Virtualizing network services the telecom cloud. Technical White Paper," Ericsson, Tech. Rep., March 2014.
- [46] ETSI Industry Specification Group (ISG) NFV, "ETSI, OpenStack Liason Statement: NFV Requirements," [https://wiki.openstack.org/w/images/c/c7/NFV\(14\)000154r2_NFV_LS_to_OpenStack.pdf](https://wiki.openstack.org/w/images/c/c7/NFV(14)000154r2_NFV_LS_to_OpenStack.pdf), September 2014, NFV(14)000154r2.
- [47] F. Callegati, W. Cerroni, C. Contoli, and G. Santandrea, "Performance of Network Virtualization in cloud computing infrastructures: The OpenStack case," in *Cloud Networking (CloudNet), 2014 IEEE 3rd International Conference on*, Oct 2014, pp. 132–137.
- [48] R. Glitho, "Cloudifying the 3GPP IP Multimedia Subsystem: Why and How?" in *New Technologies, Mobility and Security (NTMS), 2014 6th International Conference on*, March 2014, pp. 1–5.
- [49] A. Tolonen, "Dynamic Virtualized Network Functions on an OpenStack Cloud," Master's thesis, Aalto University, Espoo, Finland, 9 2014.
- [50] H. Jamjoon, D. Williams, and U. Sharma, "Don't Call Them Middleboxes, Call Them Middlepipes," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 19–24.
- [51] F. Hu, Q. Hao, and K. Bao, "A Survey on Software-Defined Network and OpenFlow: From Concept to Implementation," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 4, pp. 2181–2206, Fourthquarter 2014.
- [52] A. Doria, J. H. Salim, R. Haas, W. Wang, L. Dong, and R. Gopal, "Forwarding and Control Element Separation (ForCES)," <http://www.ietf.org/rfc/rfc5810.txt>, March 2010, Protocol Specification. Internet Engineering Task Force.
- [53] A. Lara, A. Kolasani, and B. Ramamurthy, "Network Innovation using OpenFlow: A Survey," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 1, pp. 493–512, First 2014.
- [54] Open Networking Foundation, "SDN Architecture. Issue 1. ONF TR 502," <https://www.opennetworking.org/images/stories/downloads/>

- sdn-resources/technical-reports/TR_SDN_ARCH_1.0_06062014.pdf, June 2014.
- [55] M. Taylor, "A Guide to NFV and SDN. White Paper by Metaswitch Networks," http://www.metaswitch.com/sites/default/files/Metaswitch_WhitePaper_NFVSDN_final_rs.pdf, December 2014.
- [56] VMWARE, "The VMware NSX Network Virtualization Platform. Technical White Paper," <https://www.vmware.com/files/pdf/products/nsx/VMware-NSX-Network-Virtualization-Platform-WP.pdf>, 2015, Accessed: June, 05 2015.
- [57] B. Nunes, M. Mendonca, X.-N. Nguyen, K. Obraczka, and T. Turletti, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 3, pp. 1617–1634, Third 2014.
- [58] "Open Networking Foundation (ONF)," <https://www.opennetworking.org/>, 2015, Accessed: February, 03 2015.
- [59] M. Jarschel, T. Zinner, T. Hossfeld, P. Tran-Gia, and W. Keller, "Interfaces, attributes, and use cases: A compass for SDN," *IEEE Communications Magazine*, June 2014.
- [60] C. Cui et. al, "Network Functions Virtualisation: Network Operator Perspectives on Industry Progress. White Paper No. 3, Issue 1," in *SDN and OpenFlow World Congress, Dusseldorf-Germany*, October 2014.
- [61] E. Haleplidis, J. Hadi Salim, S. Denazis, and O. Koufopavlou, "Towards a Network Abstraction Model for SDN," *Journal of Network and Systems Management*, pp. 1–19, 2014.
- [62] T. Choi, S. Kang, S. Yoon, S. Yang, S. Song, and H. Park, "SuVMF: Software-defined Unified Virtual Monitoring Function for SDN-based Large-scale Networks," in *Proceedings of The Ninth International Conference on Future Internet Technologies*, ser. CFI '14. New York, NY, USA: ACM, 2014, pp. 4:1–4:6.
- [63] Z. Michael, A. David, C. Marc, D. Nabil, K. Christos, M. Jeff, M. Serge, M. Dave, R. Evelyne, and S. Meral, "OpenFlow-enabled SDN and Network Functions Virtualization. ONF Solution Brief," Open Networking Foundation, Tech. Rep., February 2014.
- [64] W. Xia, Y. Wen, C. Foh, D. Niyato, and H. Xie, "A Survey on Software-Defined Networking," *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–1, 2014.
- [65] R. Kawashima, "vNFC: A Virtual Networking Function Container for SDN-Enabled Virtual Networks," in *Network Cloud Computing and Applications (NCCA), 2012 Second Symposium on*, Dec 2012, pp. 124–129.
- [66] A. Basta, W. Kellerer, M. Hoffmann, K. Hoffmann, and E.-D. Schmidt, "A Virtual SDN-Enabled LTE EPC Architecture: A Case Study for S-/P-Gateways Functions," in *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, Nov 2013, pp. 1–7.
- [67] J. Kempf, B. Johansson, S. Pettersson, H. Luning, and T. Nilsson, "Moving the mobile Evolved Packet Core to the cloud," in *Wireless and Mobile Computing, Networking and Communications (WiMob), 2012 IEEE 8th International Conference on*, Oct 2012, pp. 784–791.
- [68] A. Tootoonchian and Y. Ganjali, "HyperFlow: A Distributed Control Plane for OpenFlow," in *Proceedings of the 2010 Internet Network Management Conference on Research on Enterprise Networking*, ser. INM/WREN'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 3–3. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1863133.1863136>
- [69] V. Yazici, M. Sunay, and A. Ercan, "Architecture for a distributed openflow controller," in *Signal Processing and Communications Applications Conference (SIU), 2012 20th*, April 2012, pp. 1–4.
- [70] J. Yu and I. Al Ajarmeh, "An Empirical Study of the NETCONF Protocol," in *Networking and Services (ICNS), 2010 Sixth International Conference on*, March 2010, pp. 253–258.
- [71] A. Gember-Jacobson, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and A. Akella, "OpenNF: Enabling Innovation in Network Function Control," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 163–174.
- [72] J. Batalle, J. Ferrer Riera, E. Escalona, and J. Garcia-Espin, "On the Implementation of NFV over an OpenFlow Infrastructure: Routing Function Virtualization," in *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, Nov 2013, pp. 1–6.
- [73] M. R. Sama, L. M. Contreras, J. Kaipallimalil, I. Akiyoshi, H. Qian, and H. Ni, "Software-defined control of the virtualized mobile packet core," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 107–115, Feb 2015.
- [74] "OpenDayLight," <http://www.opendaylight.org/project>, 2015, Accessed: February, 06 2015.
- [75] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-Defined Networking: A Comprehensive Survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, Jan 2015.
- [76] "The Internet Engineering Task Force (IETF)," <https://www.ietf.org/>, 2015, Accessed: February, 06 2015.
- [77] "The Internet Engineering Task Force (IETF) Service Function Chaining (SFC) Working Group (WG)," <https://datatracker.ietf.org/wg/sfc/charter/>, 2015, Accessed: February, 06 2015.
- [78] "The Internet Engineering Task Force (IETF) Service Function Chaining (SFC) Working Group (WG). Documents," <https://datatracker.ietf.org/wg/sfc/documents/>, 2015, Accessed: February, 06 2015.
- [79] "Internet Research Task Force, Network Function Virtualization Research Group (NFVRG)," <https://irtf.org/nfvrg>, 2015, Accessed: February, 03 2015.
- [80] "The Broadband Forum," <https://www.broadband-forum.org/>, 2015, Accessed: February, 10 2015.
- [81] "The Broadband Forum, Technical Work in Progress," <https://www.broadband-forum.org/technical/technicalwip.php>, 2015, Accessed: February, 28 2015.
- [82] Distributed Management Task Force (DMTF), "Open Virtualization Format (OVS)," <http://www.dmtf.org/standards/ovf>, June 2015, Accessed: June 14, 2015.
- [83] DMTF, "Common Information Model (CIM)," <http://www.dmtf.org/standards/cim>, June 2015, Accessed: June 14, 2015.
- [84] Huawei Technologies, "Huawei Observation to NFV. White Paper," February 2015.
- [85] TMF, "Zero-time Orchestration, Operations and Management (ZOOM)," TeleManagement Forum, Tech. Rep., August 2014.
- [86] "Open Platform for NFV (OPNFV)," <https://www.opnfv.org/about>, 2015, Accessed: January, 26 2015.
- [87] OPNFV, "OPNFV's Arno Release," <https://www.opnfv.org/arno>, June 2015.
- [88] D. R. Lopez, "OpenMANO: The Dataplane Ready Open Source NFV MANO Stack," in *IETF Meeting Proceedings, Dallas, Texas, USA*, March 2015.
- [89] Intel, "OpenStack Enhanced Platform Awareness. White Paper," https://01.org/sites/default/files/page/openstack-epa_wp_fin.pdf, March 2015.
- [90] "MCN: The Mobile Cloud Networking EU Project," <http://www.mobile-cloud-networking.eu/site/>, 2015, Accessed: February, 12 2015.
- [91] A. Csaszar, W. John, M. Kind, C. Meirosu, G. Pongracz, D. Staessens, A. Takacs, and F.-J. Westphal, "Unifying Cloud and Carrier Network: EU FP7 Project UNIFY," in *Utility and Cloud Computing (UCC), 2013 IEEE/ACM 6th International Conference on*, Dec 2013, pp. 452–457.
- [92] "T-NOVA: Network Functions-as-a-Service (NFaaS) over Virtualized Infrastructures," <http://www.t-nova.eu/>, 2015, Accessed: January, 26 2015.
- [93] "Convergence of Wireless Optical Network and IT Resources IN Support of Cloud Services (CONTENT) EU Project," <http://content-fp7.eu/theproject.html>, January 2015, accessed on February 09, 2015.
- [94] B. O. Josep, "Experimentation on Virtualized Routing Function Migration Using OpenFlow," Master's thesis, Universitat Politècnica de Catalunya, Barcelona, Spain, 9 2014.
- [95] H. Masutani, Y. Nakajima, T. Kinoshita, T. Hibi, H. Takahashi, K. Obana, K. Shimano, and M. Fukui, "Requirements and design of flexible NFV network infrastructure node leveraging SDN/OpenFlow," in *Optical Network Design and Modeling, 2014 International Conference on*, May 2014, pp. 258–263.
- [96] T. C. Bálan, "Network Policy Function Virtualization via SDN and Packet Processing," *Review of the Air Force Academy*, vol. 27, no. 3, pp. 73–78, 2014.
- [97] A. Bremier-Barr, Y. Harchol, D. Hay, and Y. Koral, "Deep packet inspection as a service," in *Proceedings of the 10th ACM International Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '14. New York, NY, USA: ACM, 2014, pp. 271–282.
- [98] S. Gebert, D. Hock, T. Zinner, P. Tran-Gia, M. Hoffmann, M. Jarschel, E.-D. Schmidt, R.-P. Braun, C. Banse, and A. Köpsel, "Demonstrating the Optimal Placement of Virtualized Cellular Network Functions in Case of Large Crowd Events," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 359–360.
- [99] M. Peng, Y. Li, J. Jiang, J. Li, and C. Wang, "Heterogeneous cloud radio access networks: a new perspective for enhancing spectral and energy efficiencies," *Wireless Communications, IEEE*, vol. 21, no. 6, pp. 126–135, December 2014.

- [100] R. Wang, H. Hu, and X. Yang, "Potentials and Challenges of C-RAN Supporting Multi-RATs Toward 5G Mobile Networks," *Access, IEEE*, vol. 2, pp. 1187–1195, 2014.
- [101] C.-L. I, J. Huang, R. Duan, C. Cui, J. Jiang, and L. Li, "Recent Progress on C-RAN Centralization and Cloudification," *Access, IEEE*, vol. 2, pp. 1030–1039, 2014.
- [102] D. Sabella, P. Rost, Y. Sheng, E. Pateromichelakis, U. Salim, P. Guitton-Ouhamou, M. di Girolamo, and G. Giuliani, "RAN as a service: Challenges of designing a flexible RAN architecture in a cloud-based heterogeneous mobile network," in *Future Network and Mobile Summit (FutureNetworkSummit)*, 2013, July 2013, pp. 1–8.
- [103] R. Vilalta, R. Muoz, R. Casellas, R. Martinez, V. Lopez, and D. Lopez, "Transport PCE network function virtualization," in *Optical Communication (ECOC)*, 2014 European Conference on, Sept 2014, pp. 1–3.
- [104] R. Vilalta, R. Munoz, A. Mayoral, R. Casellas, R. Martinez, V. Lopez, and D. Lopez, "Transport Network Function Virtualization," *Lightwave Technology, Journal of*, vol. PP, no. 99, pp. 1–1, 2015.
- [105] "Arinet vCPE," <https://www.sdxcentral.com/products/virtual-cpe-vcpe-framework/>, 2015, Accessed: January, 26 2015.
- [106] "Anuta Networks vCPE," <http://www.anutanetworks.com/road-to-virtualizing-cpe/>, 2015, Accessed: January, 26 2015.
- [107] "Calsoftlabs vCPE," <http://sdn.calsoftlabs.com/network-function-virtualization/virtual-cpe.html>, 2015, Accessed: January, 26 2015.
- [108] S. Aleksic and I. Miladinovic, "Network virtualization: Paving the way to carrier clouds," in *Telecommunications Network Strategy and Planning Symposium (Networks)*, 2014 16th International, Sept 2014, pp. 1–6.
- [109] M. Nagy and I. Kotuliak, "Utilizing OpenFlow, SDN and NFV in GPRS Core Network," in *Testbeds and Research Infrastructure: Development of Networks and Communities*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, V. C. Leung, M. Chen, J. Wan, and Y. Zhang, Eds. Springer International Publishing, 2014, vol. 137, pp. 184–193.
- [110] E. Jacob, J. Matias, A. Mendiola, V. Fuentes, J. Garay, and C. Pinedo, "Deploying a virtual network function over a software defined network infrastructure: experiences deploying an access control VNF in the University of Basque Countrys OpenFlow enabled facility," 2014.
- [111] "HP OpenNFV Reference Architecture," <http://www8.hp.com/us/en/cloud/nfv-overview.html?>, 2015, Accessed: February, 05 2015.
- [112] "Huawei NFV Open Lab," <http://pr.huawei.com/en/news/>, January 2015, Accessed: February, 09 2015.
- [113] "Intel Open Network Platform," <http://www.intel.com/ONP>, 2015, Accessed: February, 03 2015.
- [114] "CloudNFV," <http://www.cloudnfv.com/>, 2015, Accessed: January, 26 2015.
- [115] "Alcatel-Lucent's ClouBand," <http://www.alcatel-lucent.com/solutions/cloudband>, 2015, Accessed: January, 26 2015.
- [116] Nuage Networks, "Nuage Networks Virtualized Services Platform," <http://www.nuagenetworks.net/>, February 2015, An Alcatel-Lucent venture.
- [117] Alcatel-Lucent, "CloudBand with OpenStack as NFV Platform. Strategic White Paper, NFV Insights Series," www.alcatel-lucent.com/, Alcatel Lucent and RedHat, Tech. Rep., August 2014.
- [118] "Broadcom Open NFV," <http://www.broadcom.com/press/release.php?id=s827048>, 2015, Accessed: January, 26 2015.
- [119] A. Khan, "Recent developments in high-performance system-on-chip IC design," in *Integrated Circuit Design and Technology, 2004. ICICDT '04. International Conference on*, 2004, pp. 151–158.
- [120] ODP, "OpenDataPlane project," <http://www.opendataplane.org/>, February 2015.
- [121] Cisco, "NFV Management and Orchestration: Enabling Rapid Service Innovation in the Era of Virtualization," <http://www.cisco.com/>, Cisco, Tech. Rep., June 2015.
- [122] F5, "NFV: Beyond Virtualization. Technical White Paper," F5 Networks, Inc., Tech. Rep., February 2014.
- [123] F. Yue, "Network Functions Virtualization - Everything Old Is New Again. Technical White Paper," F5 Networks, Inc., Tech. Rep., February 2014.
- [124] L. MacVittie, "Software Defined Application Services. Technical White Paper," F5 Networks, Inc., Tech. Rep., February 2014.
- [125] "Clearwater," <http://www.metaswitch.com/clearwater>, 2015, Accessed: January, 26 2015.
- [126] "Metaswitch," <http://www.metaswitch.com/nfv>, 2015, Accessed: January, 26 2015.
- [127] J. Keeney, S. v. d. Meer, and L. Fallon, "Towards real-time management of virtualized telecommunication networks," in *Network and Service Management (CNSM)*, 2014 10th International Conference on, Nov 2014, pp. 388–393.
- [128] L. Bondan, C. R. P. d. Santos, and L. Z. Granville, "Management requirements for ClickOS-based Network Function Virtualization," in *Network and Service Management (CNSM)*, 2014 10th International Conference on, Nov 2014, pp. 447–450.
- [129] Z. Bronstein and E. Shraga, "NFV virtualisation of the home environment," in *Consumer Communications and Networking Conference (CCNC)*, 2014 IEEE 11th, Jan 2014, pp. 899–904.
- [130] J. Soares, M. Dias, J. Carapinha, B. Parreira, and S. Sargento, "Cloud4NFV: A platform for Virtual Network Functions," in *Cloud Networking (CloudNet)*, 2014 IEEE 3rd International Conference on, Oct 2014, pp. 288–293.
- [131] J. Soares, C. Goncalves, B. Parreira, P. Tavares, J. Carapinha, J. P. Barraca, R. L. Aguiar, and S. Sargento, "Toward a telco cloud environment for service functions," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 98–106, Feb 2015.
- [132] S. Clayman, E. Maini, A. Galis, A. Manzalini, and N. Mazzocca, "The dynamic placement of virtual network functions," in *Network Operations and Management Symposium (NOMS)*, 2014 IEEE, May 2014, pp. 1–9.
- [133] V. Riccobene, A. Lombardo, A. Manzalini, and G. Schembra, "Network Functions At The Edge (NetFATE): Design and Implementation Issues," 2014.
- [134] E. Maini and A. Manzalini, "Management and Orchestration of Virtualized Network Functions," in *Monitoring and Securing Virtualized Networks and Services*, ser. Lecture Notes in Computer Science, A. Sperotto, G. Doyen, S. Latr, M. Charalambides, and B. Stiller, Eds. Springer Berlin Heidelberg, 2014, vol. 8508, pp. 52–56.
- [135] W. Shen, M. Yoshida, T. Kawabata, K. Minato, and W. Imajuku, "vConductor: An NFV management solution for realizing end-to-end virtual network services," in *Network Operations and Management Symposium (APNOMS)*, 2014 16th Asia-Pacific, Sept 2014, pp. 1–6.
- [136] W. Shen, M. Yoshida, K. Minato, and W. Imajuku, "vConductor: An enabler for achieving virtual network integration as a service," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 116–124, Feb 2015.
- [137] P. Donadio, G. Fioccola, R. Canonico, and G. Ventre, "A PCE-based architecture for the management of virtualized infrastructures," in *Cloud Networking (CloudNet)*, 2014 IEEE 3rd International Conference on, Oct 2014, pp. 223–228.
- [138] R. Bolla, C. Lombardo, R. Bruschi, and S. Mangialardi, "DROPv2: energy efficiency through network function virtualization," *Network, IEEE*, vol. 28, no. 2, pp. 26–32, March 2014.
- [139] W. Shen, M. Yoshida, T. Kawabata, K. Minato, and W. Imajuku, "vconductor: An nfv management solution for realizing end-to-end virtual network services," in *Network Operations and Management Symposium (APNOMS)*, 2014 16th Asia-Pacific, Sept 2014, pp. 1–6.
- [140] K. Giotis, Y. Kryftis, and V. Maglaris, "Policy-based orchestration of nfv services in software-defined networks," in *Network Softwarization (NetSoft)*, 2015 1st IEEE Conference on, April 2015, pp. 1–5.
- [141] J. Wickboldt, W. De Jesus, P. Isolani, C. Both, J. Rochol, and L. Granville, "Software-defined networking: management requirements and challenges," *Communications Magazine, IEEE*, vol. 53, no. 1, pp. 278–285, January 2015.
- [142] C. Chappell, "Deploying Virtual Network Functions: The Complementary Roles of TOSCA and NETCONF/YANG," February 2015, Technical White Paper. Heavy Reading, Cisco, Alcatel-Lucent.
- [143] Bell Labs, Alcatel Lucent, "G.W.A.T.T. (Global What if Analyzer of NeTwork Energy ConsumpTion). Bell Labs application able to measure the impact of technologies like SDN & NFV on network energy consumption. White Paper." <http://gwatt.net/intro/1>, 2015.
- [144] "Global e-Sustainability Initiative (GeSI) SMARTer2020," <http://gesi.org/SMARTer2020>, 2015, Accessed: June, 08 2015.
- [145] GREENPEACE, "Clicking Clean: How Companies are Creating the Green Internet," www.greenpeace.org, April 2014.
- [146] The Natural Resources Defense Council (NRDC), "Data Center Efficiency Assessment. Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers. Issue Paper." August 2014.
- [147] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Y. Zomaya, "A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems." *Advances in Computers*, vol. 82, pp. 47–111, 2011. [Online]. Available: <http://dblp.uni-trier.de/db/journals/ac/ac82.html#BeloglazovBLZ11>

- [148] C.-L. I, J. Huang, R. Duan, C. Cui, J. Jiang, and L. Li, "Recent Progress on C-RAN Centralization and Cloudification," *Access, IEEE*, vol. 2, pp. 1030–1039, September 2014.
- [149] E. Masanet, A. Shehabi, L. Ramakrishnan, J. Liang, X. Ma, B. Walker, V. Hendrix, and P. Mantha, "The Energy Efficiency Potential of Cloud-Based Software: A US Case Study. Lawrence Berkely National Laboratory, Berkeley California." 2013.
- [150] Cisco, "The Zettabyte Era: Trends and Analysis," Tech. Rep., May 2015.
- [151] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV-PER 001 V1.1.1: Network Functions Virtualisation (NFV); NFV Performance & Portability Best Practises," http://www.etsi.org/deliver/etsi_gs/NFV-PER/001_099/001/01.01.01_60/gs_nfv-per001v010101p.pdf, June 2014.
- [152] Ericsson AB et. al, "Common Public Radio Interface (CPRI) Specification V6.0," http://www.cpri.info/downloads/CPRI_v_6_0_2013-08-30.pdf, August 2013.
- [153] Napatech, "Time to Rethink SDN and NFV Performance," Tech. Rep., December 2014.
- [154] K. Yamazaki, T. Osaka, S. Yasuda, S. Ohteru, and A. Miyazaki, "Accelerating sdn/nfv with transparent offloading architecture," in *Presented as part of the Open Networking Summit 2014 (ONS 2014)*. Santa Clara, CA: USENIX, 2014. [Online]. Available: <https://www.usenix.org/conference/ons2014/technical-sessions/presentation/yamazaki>
- [155] A. Fischer, J. Botero, M. Till Beck, H. de Meer, and X. Hesselbach, "Virtual Network Embedding: A Survey," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 4, pp. 1888–1906, Fourth 2013.
- [156] M. Rabbani, R. Pereira Esteves, M. Podlesny, G. Simon, L. Zambenedetti Granville, and R. Boutaba, "On tackling virtual data center embedding problem," in *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, May 2013, pp. 177–184.
- [157] A. Basta, W. Kellerer, M. Hoffmann, H. J. Morper, and K. Hoffmann, "Applying NFV and SDN to LTE Mobile Core Gateways, the Functions Placement Problem," in *Proceedings of the 4th Workshop on All Things Cellular: Operations, Applications, & #38; Challenges*, ser. AllThingsCellular '14. New York, NY, USA: ACM, 2014, pp. 33–38.
- [158] H. Moens and F. D. Turck, "Vnf-p: A model for efficient placement of virtualized network functions," in *Network and Service Management (CNSM), 2014 10th International Conference on*, Nov 2014, pp. 418–423.
- [159] M. Bagaa, T. Taleb, and A. Ksentini, "Service-aware network function placement for efficient traffic handling in carrier cloud," in *Wireless Communications and Networking Conference (WCNC), 2014 IEEE*, April 2014, pp. 2402–2407.
- [160] S. Mehroghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *Cloud Networking (CloudNet), 2014 IEEE 3rd International Conference on*, Oct 2014, pp. 7–13.
- [161] M. Bouet, J. Leguay, and V. Conan, "Cost-based placement of vdp functions in nfv infrastructures," in *Network Softwarization (NetSoft), 2015 1st IEEE Conference on*, April 2015, pp. 1–9.
- [162] A. Schrijver, *Theory of Linear and Integer Programming*. New York, NY, USA: John Wiley & Sons, Inc., 1986.
- [163] M. Xia, M. shirazipour, Y. Zhang, H. Green, and A. Takacs, "Network Function Placement for NFV Chaining in Packet/Optical Datacenters," *Lightwave Technology, Journal of*, vol. PP, no. 99, pp. 1–1, 2015.
- [164] M. Yoshida, W. Shen, T. Kawabata, K. Minato, and W. Imajuku, "MORSA: A multi-objective resource scheduling algorithm for NFV infrastructure," in *Network Operations and Management Symposium (APNOMS), 2014 16th Asia-Pacific*, Sept 2014, pp. 1–6.
- [165] O. Hyeonseok, Y. Daeun, C. Yoon-Ho, and K. Namgi, "Design of an Efficient Method for Identifying Virtual Machines Compatible with Service Chain in a Virtual Network Environment," *International Journal of Multimedia and Ubiquitous Engineering*, vol. 11, no. 9, pp. 197–208, 2014.
- [166] L. F. Wanner and M. B. Srivastava, "ViRUS: Virtual Function Replacement Under Stress," in *6th Workshop on Power-Aware Computing and Systems, HotPower '14, Broomfield, CO, USA, October 5, 2014.*, 2014.
- [167] W. Ceroni and F. Callegati, "Live migration of virtual network functions in cloud-based edge networks," in *Communications (ICC), 2014 IEEE International Conference on*, June 2014, pp. 2963–2968.
- [168] LXC, "Linux Containers," <https://linuxcontainers.org/>, January 2015.
- [169] S. Hykes, "Docker," <https://www.docker.com/>, June 2015.
- [170] T. Nolle, "Is NFV and Cloud Computing Missing the Docker Boat?" <http://blog.cimicorp.com/?p=1911>, October 2014.
- [171] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, and S. Davy, "Design and Evaluation of Algorithms for Mapping and Scheduling of Virtual Network Functions," in *IEEE Conference on Network Softwarization (NetSoft)*. University College London, April 2015.
- [172] J. Ferrer Riera, X. Hesselbach, E. Escalona, J. Garcia-Espin, and E. Grasa, "On the complex scheduling formulation of virtual network functions over optical networks," in *Transparent Optical Networks (ICTON), 2014 16th International Conference on*, July 2014, pp. 1–5.
- [173] J. Ferrer Riera, E. Escalona, J. Batalle, E. Grasa, and J. Garcia-Espin, "Virtual network function scheduling: Concept and challenges," in *Smart Communications in Network Technologies (SaCoNeT), 2014 International Conference on*, June 2014, pp. 1–5.
- [174] P. Brucker, A. Drexler, R. Mhring, K. Neumann, and E. Pesch, "Resource-constrained project scheduling: Notation, classification, models, and methods," *European Journal of Operational Research*, vol. 112, no. 1, pp. 3 – 41, 1999.
- [175] J. Baewicz, W. Domschke, and E. Pesch, "The job shop scheduling problem: Conventional and new solution techniques," *European Journal of Operational Research*, vol. 93, no. 1, pp. 1 – 33, 1996.
- [176] F. Glover and M. Laguna, *Tabu Search*. Norwell, MA, USA: Kluwer Academic Publishers, 1997.
- [177] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at google with borg," in *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, 2015, p. 18. [Online]. Available: <http://doi.acm.org/10.1145/2741948.2741964>
- [178] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. Katz, S. Shenker, and I. Stoica, "Mesos: A platform for fine-grained resource sharing in the data center," in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'11. Berkeley, CA, USA: USENIX Association, 2011, pp. 295–308. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1972457.1972488>
- [179] A. Ahmed and P. Panagiotis, "MIDAS: Middlebox Discovery and Selection for On-Path Flow Processing," in *IEEE COMSNETS*, 2015, p. 8.
- [180] R. Rosa, M. Silva Santos, and C. Esteve Rothenberg, "Md2-nfv: The case for multi-domain distributed network functions virtualization," in *Networked Systems (NetSys), 2015 International Conference and Workshops on*, March 2015, pp. 1–5.
- [181] M. G. Rabbani, F. Z. Mohamed, and R. Boutaba, "On Achieving High Survivability in Virtualized Data Centers," *IEICE Transactions*, vol. 97-B, no. 1, pp. 10–18, 2014.
- [182] R. Mijumbi, J.-L. Gorricho, J. Serrat, M. Claeys, F. De Turck, and S. Latre, "Design and evaluation of learning algorithms for dynamic resource management in virtual networks," in *Network Operations and Management Symposium (NOMS), 2014 IEEE*, May 2014, pp. 1–9.
- [183] S. Pearson and G. Yee, *Privacy and Security for Cloud Computing*. Springer, Series on Computer Communications and Networks, 2013.
- [184] ETSI Industry Specification Group (ISG) NFV, "ETSI GS NFV-SEC 001 V1.1.1: Network Functions Virtualisation (NFV); NFV Security; Problem Statement," http://www.etsi.org/deliver/etsi_gs/NFV-SEC/001_099/001/01.01.01_60/gs_NFV-SEC001v010101p.pdf, October 2014.
- [185] "Topology and Orchestration Specification for Cloud Applications Version 1.0," <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf>, November 2013, oASIS Standard.
- [186] J. Schonwalder, M. Bjorklund, and P. Shafer, "Network configuration management using netconf and yang," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 166–173, Sept 2010.
- [187] A. D. Birrell and B. J. Nelson, "Implementing remote procedure calls," *ACM Trans. Comput. Syst.*, vol. 2, no. 1, pp. 39–59, Feb. 1984. [Online]. Available: <http://doi.acm.org/10.1145/2080.357392>
- [188] J. P. Reilly, "Implementing the TM Forum Information Framework (SID). A Practitioner's Guide. Version 1.0," <http://inform.tmforum.org/wp-content/uploads/2014/05/Implementing-the-SID-v1dot0b-Chapters-1-through-3.pdf>, September 2011.
- [189] J. Rumbaugh, I. Jacobson, and G. Booch, *Unified Modeling Language Reference Manual, The (2Nd Edition)*. Pearson Higher Education, 2004.
- [190] G. Katsaros, M. Menzel, A. Lenk, J. Rake-Revelant, R. Skipp, and J. Eberhardt, "Cloud application portability with toasca, chef and openstack," in *Cloud Engineering (IC2E), 2014 IEEE International Conference on*, March 2014, pp. 295–302.

- [191] T. Binz, U. Breitenbücher, F. Haupt, O. Kopp, F. Leymann, A. Nowak, and S. Wagner, "Opentosca—a runtime for toasca-based cloud applications," in *Service-Oriented Computing*. Springer, 2013, pp. 692–695.
- [192] Y. Lee and J. Lee, "Optimizing the operation layer algorithm of netconf protocol," in *Multimedia and Ubiquitous Engineering*, ser. Lecture Notes in Electrical Engineering, J. J. H. Park, S.-C. Chen, J.-M. Gil, and N. Y. Yen, Eds. Springer Berlin Heidelberg, 2014, vol. 308, pp. 249–258. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-54900-7_36
- [193] TMF Forum, "TM Forum Information Framework Enhancements to Support ZOOM. TR224, Release 15.0.0," <https://www.tmfforum.org/resources/suite/gb922-information-framework-sid-r15-0-0/>, Tech. Rep., May 2015.
- [194] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions," *Future Gener. Comput. Syst.*, vol. 29, no. 7, pp. 1645–1660, Sep. 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.future.2013.01.010>
- [195] F. Van den Abeele, J. Hoebeke, G. Teklemariam, I. Moerman, and P. Demeester, "Sensor Function Virtualization to Support Distributed Intelligence in the Internet of Things," *Wireless Personal Communications*, vol. 81, no. 4, pp. 1415–1436, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s11277-015-2481-4>
- [196] N. Omnes, M. Bouillon, G. Fromentoux, and O. Le Grand, "A programmable and virtualized network infrastructure for the internet of things: How can nfv sdn help for facing the upcoming challenges," in *Intelligence in Next Generation Networks (ICIN), 2015 18th International Conference on*, Feb 2015, pp. 64–69.
- [197] G. Xylomenos, C. Ververidis, V. Siris, N. Fiotou, C. Tsilopoulos, X. Vasilakos, K. Katsaros, and G. Polyzos, "A Survey of Information-Centric Networking Research," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 2, pp. 1024–1049, Second 2014.
- [198] M. Arumathurai, J. Chen, E. Monticelli, X. Fu, and K. K. Ramakrishnan, "Exploiting ICN for Flexible Management of Software-defined Networks," in *Proceedings of the 1st International Conference on Information-centric Networking*, ser. INC '14. New York, NY, USA: ACM, 2014, pp. 107–116.
- [199] Intel, Brocade, Cyan, Red Hat, and Telefonica, "End to End Network Function Virtualization Architecture Instantiation," <http://pressoffice.telefonica.com/documentos/EndtoEndNFVArchitectureFinal.pdf>, February 2015, Joint White Paper.
- [200] T. Nolle, "ExperiaSphere: Take The First Step to Open Orchestration," <http://blog.experiasphere.com/>, June 2014.



Juan-Luis Gorricho received a telecommunication engineering degree in 1993, and a Ph.D. degree in 1998, both from the UPC. He is currently an associate professor at the UPC. His recent research interests are in applying artificial intelligence to ubiquitous computing and network management; with special interest on using smartphones to achieve the recognition of user activities and locations; and applying linear programming and reinforcement learning to resource management in virtualized networks and functions.



Niels Bouten obtained a masters degree in computer science from Ghent University, Belgium, in June 2011. In August 2011, he joined the Department of Information Technology at Ghent University, where he is active as a Ph.D. student. His main research interests are the application of autonomic network management approaches in multimedia delivery. The focus of this research is mainly on the end-to-end Quality of Experience optimization, ranging from the design of a single autonomic control loop to the federated management of these distributed loops.



Filip De Turck is a professor at the Department of Information Technology of Ghent University and iMinds in Belgium, where he leads the network and service management research group. His main research interests include scalable software architectures for network and service management, design and performance evaluation of novel QoE-aware multimedia delivery systems. He served as TPC chair of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2012) and the IFIP/IEEE Integrated Network Management Symposium (IM 2013). He is associate editor of the *Journal on Network and System Management*, the *International Journal of Network Management* and *IEEE Transactions on Network and Service Management*.



Rashid Mijumbi obtained a degree in electrical engineering from Makerere University, Uganda in 2009, and a PhD in telecommunications engineering from the Universitat Politècnica de Catalunya (UPC), Spain in 2014. He is currently a Postdoctoral Researcher in the Network Engineering Department at the UPC. His research interests are in autonomic management of networks and services. Current focus is on management of resources for virtualized networks and functions, cloud computing and software defined networks.



Joan Serrat received a degree of telecommunication engineering in 1977 and a PhD in the same field in 1983, both from the Universitat Politècnica de Catalunya (UPC). Currently, he is a full professor at UPC where he has been involved in several collaborative projects with different European research groups, both through bilateral agreements or through participation in European funded projects. His topics of interest are in the field of autonomic networking and service and network management. Currently, he is the contact point of the TM Forum at UPC.



Raouf Boutaba received the MSc and PhD degrees in computer science from the Université de Pierre et Marie Curie, Paris, France, in 1990 and 1994, respectively. He is currently a full professor of computer science at the University of Waterloo, Waterloo, ON, Canada, and a distinguished visiting professor at the Pohang University of Science and Technology (POSTECH), Korea. His research interests include network, resource and service management in wired and wireless networks. He has received several best paper awards and other recognitions such as the Premier's Research Excellence Award, the IEEE Hal Sobol Award in 2007, the Fred W. Ellersick Prize in 2008, the Joe LociCero and the Dan Stokesbury awards in 2009, and the Salah Aidarous Award in 2012. He is a fellow of the IEEE and the Engineering Institute of Canada.