

# The Dark Side of Network Functions Virtualization: A Perspective on the Technological Sustainability

R. Bolla<sup>\*†</sup>, R. Bruschi<sup>\*</sup>, F. Davoli<sup>\*\*</sup>, C. Lombardo<sup>\*\*</sup>, J. F. Pajo<sup>\*\*</sup>, O. R. Sanchez<sup>†\*</sup>

<sup>†</sup>DITEN - University of Genoa - Genoa, Italy

<sup>\*</sup>CNIT - University of Genoa Research Unit - Genoa, Italy

**Abstract**— The Network Functions Virtualization (NFV) paradigm is undoubtedly a key technological advancement in the Information and Communication Technology (ICT) community, especially for the upcoming 5G network design. While most of its promise is quite straightforward, the implied reduction of the power consumption/carbon footprint is still debatable, and not in line with the energy efficiency perspective forecasted by the ETSI NFV working group (WG). In this paper, we provide an estimate of the possible future requirements of this upcoming technology when deployed according to the virtual Evolved Packet Core (vEPC) use case specified by the ETSI NFV WG. Our estimation is based on real performance levels, certified by independent third-party laboratories, and datasheet values provided by existing commercial products for both the legacy and NFV network architectures, under different deployment scenarios. Obtained results show that a massive deployment of the current NFV technologies in the EPC may lead to a minimum increase of 106% in the carbon footprint/energy consumption with respect to the Business As Usual (BAU) network solutions. Moreover, these values tend to increase at a very high pace when the most suitable software/hardware combination is not applied, or when packet processing latency is taken into account.

**Keywords**—NFV; vEPC; SGW; Carbon Footprint; Energy Efficiency

## I. INTRODUCTION

As the underlying physical infrastructure of today's networks is becoming inadequate for future demands, it is a common viewpoint among technological and business experts that state-of-the-art Information and Communication Technologies (ICT) – primarily cloud computing, Internet of Things (IoT) and the latest generation mobile networks – are pushing for a new technological and business revolution that will have a significant impact on our everyday way of being. The commodification of new (virtual/ cloud / personal/ social) services, as well as the commoditization of hardware and software resources (e.g., storage resources), will significantly increase the ICT resource demand and will lead to a sustainability problem. Indeed, the energy requirements and the carbon footprint have already reached the warning levels with the current technologies and business models. Already in 2012, close to 4.7% of the electrical energy worldwide was consumed by ICT, releasing roughly 1.7% of the total greenhouse gas emissions into the atmosphere [1] [2].

The revolutionary concept of Network Functions Virtualization (NFV) [3] is seen as the key solution by network

operators to detach network functions (e.g., routers/ firewalls, etc.) from the physical infrastructure and virtualize them into a more open industry standard of high volume servers. While this paradigm appears to be a valid way of addressing most of the above-mentioned challenges, the reduction of power consumption and related carbon footprint is still debatable. Although energy efficiency and cost reductions are supposed to be the pillars of NFV, there is a lack of clear figures supporting this claim. During the recent and somehow unstoppable rise of this technology, the most frequent tacit question of many experts from the academic and industry communities regarded the real efficiency and the performance levels that general-purpose hardware and IT virtualization means can provide. Understanding whether the (much lower) deployment cost of these systems will be supported by a sufficiently high efficiency level – capable of assuring operating expenses comparable with the ones of legacy network technologies – is still an open point on the NFV viability.

In an attempt to seek clarification on this point, this paper explores the virtualization of the Evolved Packet Core (vEPC), which is one of the use cases defined by the ETSI NFV working group (WG), in order to provide a breakdown of the power and performance behavior that could be expected with the deployment of NFV technologies in real edge networks. The reported estimates are based on real data provided by hardware manufacturers and, where made, assumptions on NFV deployment and performance tend to the optimistic. However, results show that, in the best case, an NFV deployment may lead to approximately double the carbon footprint and, consequently, the energy requirement of network infrastructures, with respect to the Business As Usual (BAU) scenario.

The remainder of this paper is organized as follows: Section II discusses the architectural steps that will bring to the evolution of the Point of Presence (PoP) towards NFV. Section III presents the mathematical model used for the estimation of the annual carbon emission. Section IV describes the reference scenarios that will be used for the test results, which can be found in Section V. Finally, the concluding statements are given in Section VI.

## II. EVOLUTION OF THE POINT OF PRESENCE TOWARDS NFV

Among the nine use cases specified in [4], the virtualization of the mobile core is one of the hottest topics for network providers. Its main motivation is to reduce the Total Cost of Ownership (TCO), by allocating different network functions on

general-purpose hardware. Flexible allocation of resources on demand enables a more efficient operation.

As defined by 3GPP [5], the latest development of the mobile core in a Long Term Evolution (LTE) network is the Evolved Packet Core (EPC), which unifies voice and data traffic on IP, removing the traditional circuit-switching for voice. A basic EPC consists of a Serving Gateway (SGW), Packet Data Network Gateway (PGW), Home Subscriber Server (HSS), and Mobility Management Entity (MME).

In this work, the focus will be on the SGW, since it is the main function that will undergo deep changes in the evolution of the mobile core from the traditional infrastructure to the NFV adoption. In fact, it is expected that most of the functionalities originally available in the legacy device as an SGW will be deployed in the NFV scenario as Virtual Network Functions (VNFs), which will be composed of a number of Virtual Machines (VMs) running on general purpose servers. Fig. 1 shows how the SGW function in the PoP changes from the traditional dedicated hardware (represented in the box) to the infrastructure supporting NFV. This new paradigm will remove the components originally performing those operations from the legacy device, in which only the physical components and capabilities needed to interconnect the PoP to and from the Radio Access Network (RAN) and the backbone network will remain, using long haul interfaces such as GPON (Gigabit Passive Optical Network) and WDM (Wavelength Division Multiplexing).

As can be seen in Fig. 1, the removal of the SGW functionalities and their deployment as VNFs require the presence of servers to host the related VMs. The number of such servers, their performance level and the availability of power management capabilities are important factors to determine the overall energy requirements of the deployment, which can be higher than that of the BAU case due to the resulting virtualization overhead, as pointed out in [6].

In addition, it is worth noting that the servers must be interconnected to the SGW and among themselves for the proper communication and operation of the VMs composing the VNFs. The knowledge of the topologies used inside the datacenters, of which the fat tree topology is a common choice, allows the determination of the number of switches needed for the proper interconnections. The energy consumption of a typical switch has grown in the last several years, due to the improved capacity.

In the following section, these aspects will be taken as input to estimate the energy requirement of the NFV network architectures and compare it against the legacy ones. Since the comparison of SGWs in the BAU scenario and its equivalent server architecture in the NFV case is the focus, for simplicity, other factors such as cooling, lighting, maintenance of devices, etc., will not be considered here, assuming they are constant for both cases (and will be studied in future work). The estimation will be performed by defining a mathematical model and by quantifying the energy consumed in both scenarios, according to the real performance levels of existing commercial products.

### III. CARBON FOOTPRINT ESTIMATION

This section aims to present how the annual carbon emission of the mobile core network has been estimated. The carbon

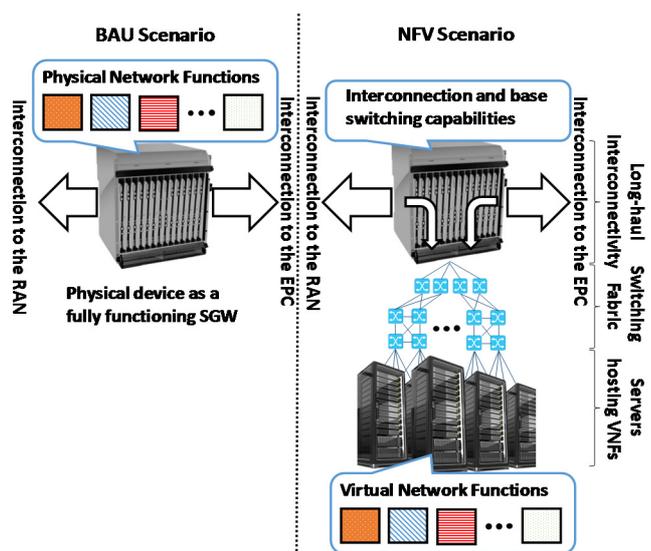


Fig. 1. Evolution of the SGW in the presence of NFV.

emissions are computed through the Total Carbon Footprint (TCF). The latter is the sum of the Embodied Carbon Footprint (ECF), which is the volume of Green House Gases (GHG) in the atmosphere due to the acquisition of the physical device (e.g., fabrication, packaging, etc.) and the Operating Carbon Footprint (OCF), which is the amount of GHG due to its operational functions, or simply, to the operating power consumption [7]. However, by following [7], it can be estimated that more than 90% of the annual TCF of a single mid-range server (e.g., rack server) is due to its OCF. Therefore, it is reasonable to focus on estimating the OCF for both BAU and NFV architectures. The general equation of the OCF for a year can be computed by using the equation from [7]:

$$OCF^{1y} = c \cdot 365 \cdot \varphi^{day} \quad (1)$$

where  $c$  is the conversion factor of power in kWh into kg of  $CO_2$ , with the value 0.37 kg/kWh for Western Europe [8], and  $\varphi^{day}$  is the total power consumption of the network for a day. The latter can be computed as

$$\varphi^{day} = \sum_{h=1}^{24} (N_h \cdot \varphi_s) + 24 \cdot (N_{sw} \cdot \varphi_{sw} + \delta \cdot N_{gw} \cdot \varphi_{gw}) \quad (2)$$

The first term is composed of the dynamic number of active NFV servers in each hour of the day,  $N_h$ , multiplied by  $\varphi_s$ , the hourly power consumption of a server.  $N_{sw}$  and  $\varphi_{sw}$  are the corresponding values related to the switches, which are mainly needed for interconnecting the NFV servers, while  $N_{gw}$  and  $\varphi_{gw}$  are those due to the SGWs described in Section II, with  $\delta$  representing the remaining fraction of consumption of SGWs, as the functionalities are removed.

In the BAU case, we assume that no power saving algorithms are implemented and no additional switches are required; thus, only the term  $24 \cdot N_{gw} \cdot \varphi_{gw}$  remains, which corresponds to the consumption of the fully functioning physical SGWs ( $\delta=1$ ). In the NFV case, all three terms in (2) would appear, but the role of the SGW is now reduced to perform the interconnection and base switching capabilities between the RAN, the EPC, and the datacenter implementing the VNFs (see Fig. 1). This scaling of functionalities of the SGW is taken into account by the coefficient  $\delta < 1$ . In order to estimate a reasonable value for  $\delta$ ,

the situation can be compared to that of scaling the power consumption of routers vs. switches; according to the trends that were derived in [9] for the edge network (and to the power consumption breakdown among device components operated in [10], among others), it is possible to conclude that the power consumed by the “downsized” SGW would be approximately 50% of that in the BAU case ( $\delta=0.5$ ).

One of the goals of NFV is traffic consolidation [3]; resources can be aggregated and servers not in use can be turned off or on in real-time. This possibility allows to save power, by dynamically activating the minimum number of servers to support the current traffic. Taking this into account,  $N_h$  can be calculated as:

$$N_h = \left\lceil \frac{a_h \gamma A}{\lambda_s} \right\rceil \quad (3)$$

where  $a_h \in [0,1]$  is the normalized hourly factor of the peak traffic, i.e.,  $a_h = 1$  corresponds to the hour  $h$  of the day that has reached the peak traffic (this will be reintroduced in the next section),  $\gamma$  is a factor used to convert  $A$ , the average traffic intensity obtained from [8], into peak traffic, and  $\lambda_s$  is the maximum throughput of a server; the final figure is rounded up to the nearest integer. According to [11], the peak traffic is currently 3 times the average traffic and this factor will grow up to 3.3 in 2019. Dimensioning exactly according to the peak ( $\gamma=3$ ) is unrealistic, since network operators usually over-dimension their system according to double of the peak traffic ( $\gamma=6$ ). Hence,  $\gamma$  serves as an over-dimensioning factor in such a case, and the total number of deployed servers can be determined by

$$N_s = \left\lceil \frac{\gamma A}{\lambda_s} \right\rceil \quad (4)$$

$N_{gw}$  is similarly computed as in (4), since the SGWs are assumed not to implement power saving schemes, but by using the throughput of the SGW,  $\lambda_{gw}$ , instead of  $\lambda_s$ .

$N_{sw}$  can be computed by using the  $k$ -ary fat-tree topology [12], which consists of 3 switching layers (i.e., core, aggregation, and access layers). The topology consists of  $k$ -racks, and for every rack we need  $k/2$   $k$ -port switches for the aggregation and access layers, and an additional  $(k/2)^2$  core switches. The  $k$ -port switches support  $k^3/4$  servers. According to [12],

$$N_{sw} = k + \left\lceil \frac{k^2}{2} \right\rceil \quad (5)$$

with  $k$  computed as:

$$k = \left\lceil 2 + \sqrt{N_s} \right\rceil \quad (6)$$

giving the number of switches required to interconnect the deployed servers.

Since NFV leads to a radical change in criteria to be used for assuring network and services reliance, we decided to explicitly take this aspect into account. In detail, like other technologies based on IT virtualization means, NFV allows sensibly reducing the density of redundant hardware with respect to BAU networking solutions (where most of the devices are deployed in 1+1 hot standby protection). In order to achieve similar resiliency levels, NFV deployments can rely on  $N+M$  hardware protection, with  $M \ll N$ , since virtual machines on a server

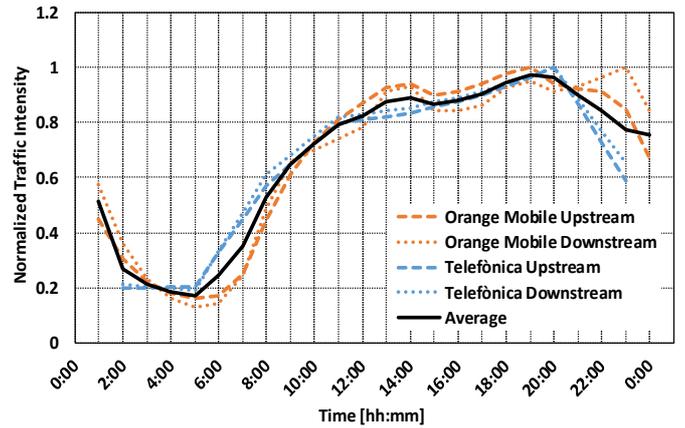


Fig. 2. Normalized hourly mobile Internet traffic [13].

experiencing a fault can be recovered on any other server in the datacenter.

Thus, in order to handle these aspects, the proposed estimation model has been completed by adding a parameter  $M$  to the values of  $N_h$  derived as in (3).  $M$  has been dimensioned to have a spare server for every rack. Regarding the switching fabric, the redundancy is already provided by the fat-tree topology, while for the physical device providing the long-haul interconnectivity, it is assumed to be deployed according to the 1+1 protection scheme.

#### IV. REFERENCE SCENARIOS

This section describes the input data and the physical composition of the servers needed for the EPC for both the BAU and NFV scenarios that will be used in the tests to be presented in Section V.

For the daily traffic dynamics, we have used the hourly traffic of both the fixed and mobile networks reported in [13] for two major European operators – Telefónica and Orange France – Internet Service Providers (ISPs). Fig. 2 shows the dynamic mobile network traffic from [13]. These data were averaged from the 28th of November 2013 to the 4th of December 2013 and from the 23rd to the 29th of March 2013, respectively, for the Orange and Telefónica ISPs. The mobile Internet upstream and downstream were studied separately for both ISPs. Fig. 2 shows that the peak traffic for the Orange network is at noon and night-time, in correspondence with the break interval and after work. This trend, though still evident, is more smoothed in the Telefónica curve, where the data are measured with a 3-hour resolution. The traffic model used for this study is the average of the four curves, which corresponds to the hourly factor  $a_h$  in (3).

The number of servers needed to handle the data throughput in the SGW function has been estimated for both the BAU and the NFV case. BAU architectures use dedicated hardware, as described in Section IV.A, while the NFV case is composed of general-purpose servers virtualizing the network functions that compose the SGW, as shown in Section IV.B. The 2016 network model and the network size of Western Europe have been chosen in [8] to determine the monthly average network load. Considering that the mobile core is in the edge, the

corresponding average load  $A=20.1138$  Exabytes per month was used. It is worth noting that, in the description of the two scenarios and in the results reported in Section V, signaling in the EPC has been neglected. According to [14], operators can expect that for every million smartphone users 31,000 signaling transactions are generated per second during peak hours in an LTE network. However, it is estimated in [15] that these account for less than 10% of the network throughput. For simplicity, the effect of signaling in the mobile core is neglected for both cases and only the estimation for handling the data throughput is considered.

### A. BAU Architecture

The Nokia 7750 family of Service Routers (SRs) is one of the packet processing engines that are currently being used by network providers [16] for SGWs/PGWs. Three SRs were selected for this study and their specifications are reported in Table I. The reported maximum consumption values are in the worst-case scenario (i.e., minimum size Ethernet frames of 64 bytes, plus 20 bytes of Preamble, Delimiter and Inter-frame Gap).

TABLE I. NOKIA SERVICE ROUTER DATASHEET VALUES

Model name	Max throughput [Mpps]	Max power consumption [W]
7750 SR7	2976	3750
7750 SR12	5952	6480
7750 SR12e	10713	12000

The maximum throughput and power consumption values in Table I have been used as  $\lambda_s$  and  $\varphi_s$ , respectively. It can be noted that SR12 is the most efficient among the three SRs in handling the data traffic in terms of throughput per power consumption.

### B. NFV Scenario

The number of needed NFV servers highly depends on the throughput that each can handle, which in turn depends on its configuration (i.e., processor type, NICs, etc.). In a study conducted by an independent third party, the European Advanced Networking Test Center Action Group (EANTC) [17], the virtual Service Router (vSR) throughput for the data path of the worst-case scenario described above reached 57.5 million packets per second (Mpps). In this set-up, the vSR consists of two compute nodes, one each for the Data Path and the Control Path VMs.

To characterize an NFV server with a vSR focusing only on the data plane, we consider a hosting server with one VM inside solely for the data path; so, we remove the control path VM considered in [17], and use this resource slot for another processor for the data path. A vSR server for the NFV case must at least consist of a motherboard, processor, and Network Interface Card (NIC), as shown in Table II. Regarding the

TABLE II. COMPOSITION OF THE NFV SERVER

Unit	Model	Quantity	Power consumption [W/Unit]
Processor	Intel Processor a-j	2	80-140
Motherboard	Intel Server Board	1	145
NIC	Intel 10GbE Card	8	6.8

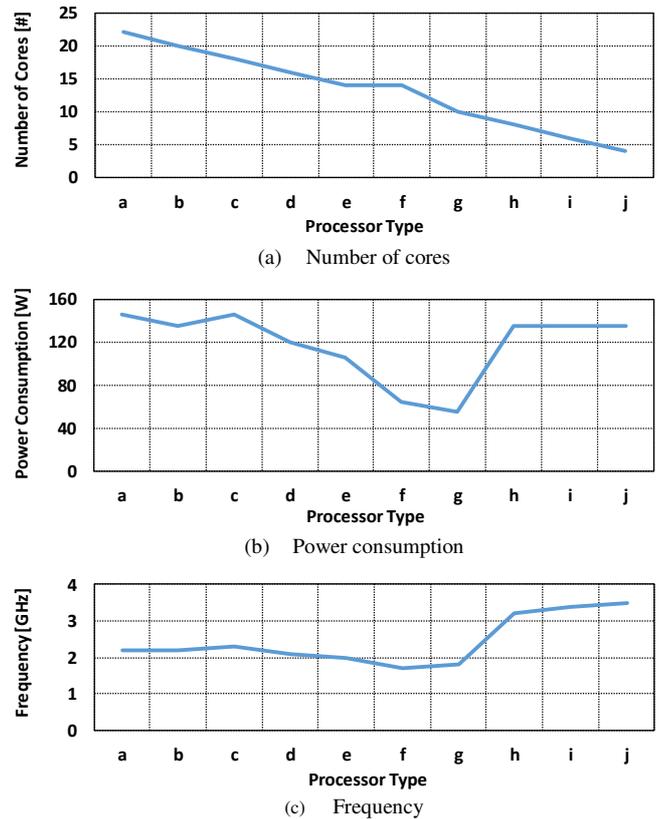


Fig. 3. Selected processors types.

processors, we consider the Intel Xeon Processor E5-2600 v4 Product Family [18] for our comparison. In this family, ten processors are available, which have been labeled with the letters a-j in Table II and Fig. 3. The first five processors (a-e) are optimized according to the number of cores, the next two (f-g) are optimized according to the power consumption, and the last three (h-j) are optimized according to frequency [18], as shown in Fig. 3.

Each of these processors would yield a different NFV server in terms of throughput and power consumption. The power consumption for a server configuration can be computed based on the values in Table II. For the throughput estimation inside a server, the value in [17] was used as the source value and scaled through Amdahl's law [19], to obtain a non-linear increase in throughput. The law states that the speed-up (or performance) achievable through parallelization as the number of resources (e.g., CPUs/cores) is increased can be computed as:

$$s_c(n) = \frac{1}{(1-p) + \frac{p}{n}} \quad (7)$$

where  $n$  is the number of cores needed to attain the speed-up and  $p$  is the percentage of parallelization that a system achieves; thus,  $1-p$  is the percentage of the workload that will remain serial.

The throughput of an NFV server is then obtained through:

$$\lambda_s = 2 \cdot \lambda_c \cdot \frac{f_a}{f_c} s_c(n) \quad (8)$$

where  $\lambda_c$  and  $f_c$  are the source throughput and processor frequency, respectively, while  $f_a$  and  $s_c(n)$  are the selected processor frequency and the speed-up that will be obtained

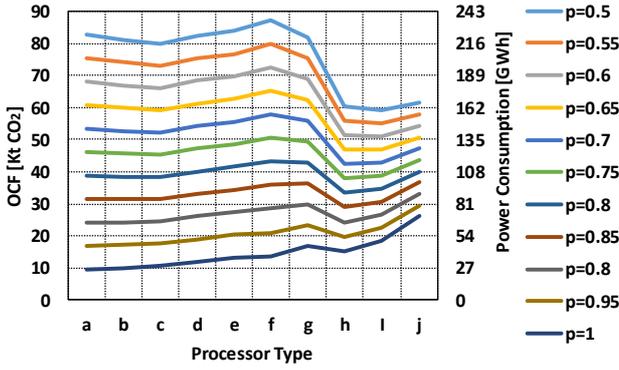


Fig. 4. OCF and power consumption per year obtained on different processors by varying  $p$ .

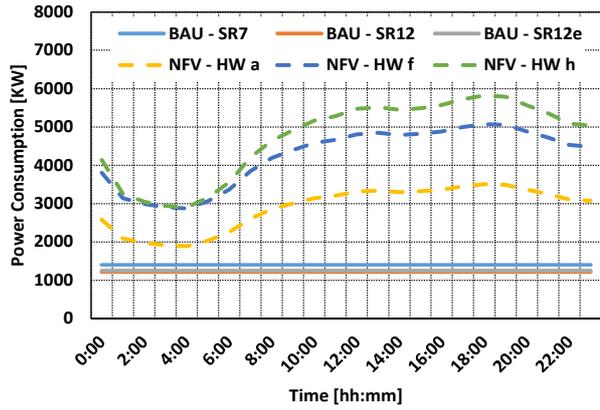


Fig. 5. Hourly server consumption in a day for the BAU and NFV scenarios.

according to the number of cores. Unless explicitly indicated, the software is assumed to be completely parallelizable (i.e.,  $p=1$ ). It was also assumed that doubling the number of processors in an NFV server also doubles the throughput, whence the value 2. This assumption is the best-case scenario, since it is assumed that the overall performance of the server scales linearly with the number of processors.

## V. PERFORMANCE RESULTS

This section compares the results of the annual OCF and energy consumption for the BAU and NFV cases, along with the effects of varying the model parameters  $p$ ,  $\delta$ , and  $\gamma$  (which will be introduced/reintroduced shortly). The effects of the processor type, switch configuration and latency requirement for the NFV case are also discussed.

Fig. 4 shows the effects of different values of  $p$ , the percentage of parallelization, on the annual OCF of an NFV system deployed using ten processor types presented in Section IV.B. NFV heavily decreases its consumption as  $p$  increases, because the number of servers can be reduced thanks to the increased parallelization. However, the type of processor providing the lowest consumption changes for different values of  $p$ . For  $0 < p < 0.85$ , the most efficient processors belong to the frequency-optimized type, which in theory is the best if the program is less likely to be parallelized. From  $p=0.85$  onwards, the core-optimized configuration becomes more efficient, while

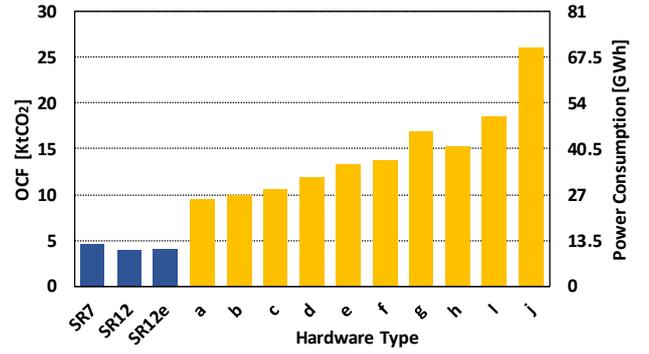


Fig. 6. OCF and power consumption per year in the BAU and NFV case obtained for  $p = 1$ .

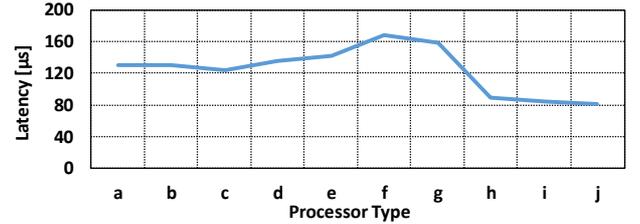


Fig. 7. Estimated frame-forwarding latency of the NFV servers.

frequency-optimized processors become the worst, beaten even by the power-optimized processors.

The hourly consumption in a day for both the NFV and BAU cases are compared in Fig. 5, which reports data related to the consumption of BAU packet processing engines and of the best performing vSRs (for each optimization type), for  $p=1$ .

The possibility of applying power saving policies on the general-purpose servers enables NFV to consume power proportionally to the traffic level. This means saving power in the hours from 02:00 to 07:00, and having very high consumptions during the peak hours, from 12:00 to 23:00. For the BAU scenario without power saving methods, the maximum power is constantly consumed, even in the non-peak hours. However, NFV consumptions are higher than in the BAU case for all server types, and even in the non-peak hours.

A comparison of the resulting annual OCF and power consumption for both cases is shown in Fig. 6. For all processor types, the OCF of the NFV case is at least 106% higher than that of the BAU SRs. It is worth noting that this figure shows the best case for NFV, i.e.  $p=1$  in Fig. 4, which may even not be true in the real-world scenario, as programs retain at least some tasks which are not parallelizable. For example, selecting a more realistic  $p=0.95$  from Fig. 4, the best performing processor for the NFV case (i.e.,  $a$ ) gives consumptions that are almost triple with respect to the BAU case.

The core-optimized configuration is the most efficient one as  $p$  moves closer to 1; however, this configuration is constrained by the latency requirements, as shown in Fig. 7. This estimation is obtained by scaling a fraction of 129  $\mu\text{s}$ , the reported average latency value obtained from [17]. Since the reported minimum latency reached a value of 13  $\mu\text{s}$ , we assumed a portion of it (specifically, around 5  $\mu\text{s}$ ) to represent the fixed part of the

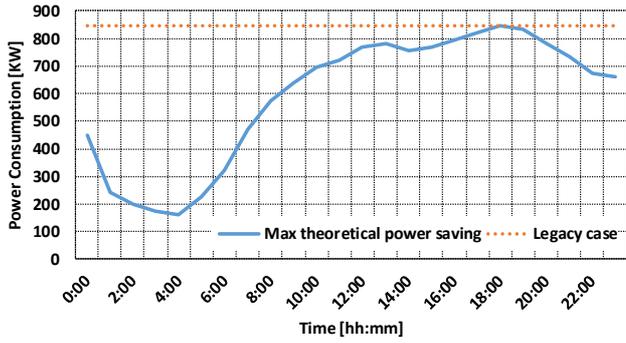


Fig. 8. NfV switch consumption in the power saving case vs legacy case.

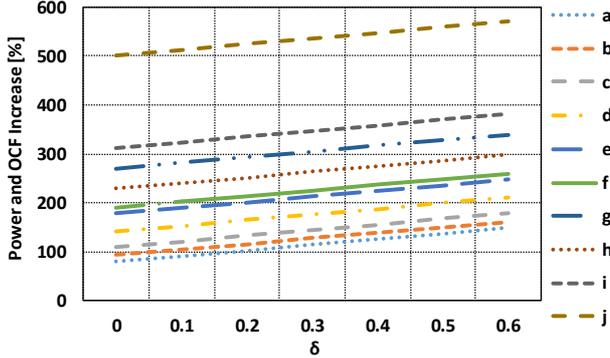


Fig. 9. OCF and power consumption increase of NfV at increasing  $\delta$  with respect to the BAU case.

latency, due to the time spent by packets in the hardware without being processed. The remaining amount of the average latency ( $124\mu\text{s}$ ) is scaled proportionally to the ratio  $f_c/f_a$ . If the latency is strictly imposed, the frequency-optimized configuration must be selected at the expense of heavy consumption.

The switches interconnecting the NfV servers are always turned on in this estimation, even though some of the NfV servers will be turned off at some point in time. To further reduce the consumption of the NfV case, we suppose that switches also adopt power saving schemes that allow for the evaluation of the theoretical minimum consumption of the switches (i.e., maximum power saving), as shown in Fig. 8. This estimation has been performed by computing the least number of switches needed depending on the hourly number of active NfV servers through the  $k$ -ary fat-tree topology. This enhancement would reduce the total consumption of the NfV case by at most 8%, which is still small, as it only improves the OCF of the NfV case to at least 92% (from 106%) more than the BAU cases. In addition, the possible deployment of consolidation schemes might generate VM migrations that would have a heavy impact on the overall costs, which may not be justified by obtaining such small savings.

As NfV does not entirely virtualize all the functions in an SGW, as seen from Fig. 1, it still uses a fraction of the legacy device physical components, defined as  $\delta$  in Fig. 9. The figure shows the impact of this additional SGW functionality (implemented on a SR7 from the BAU case) on the power consumption in an NfV deployment. Even if the SGWs were entirely removed ( $\delta=0$ ), the consumption would still be at least

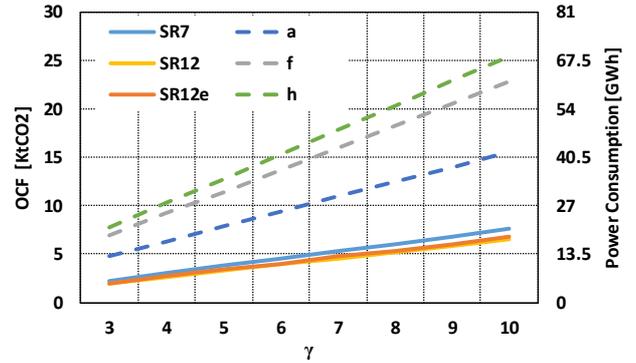


Fig. 10. OCF and power consumption values at increasing  $\gamma$ .

double. It should be noted that the default value,  $\delta=0.5$ , was used for all other results, as defined in Section III.

The results shown in this section have been computed at  $\gamma=6$ , which is the usual over-dimensioning factor applied by the operators, in order to ideally cover the double of the peak traffic. Additionally, in order to take into account the increase of traffic forecasted for the near future, Fig. 10 shows the effect of increasing  $\gamma$ , still considering  $p=1$ .

The increase for both the BAU and NfV cases are still linear, but that of the NfV case grows faster. One of the reasons is that the switches do not increase linearly with the number of NfV servers. In the real-world scenario, a complexity factor must also be taken into account as the number of servers increases.

## VI. CONCLUSIONS

In this paper, we proposed an estimation model for analyzing the environmental sustainability and energy requirements of the upcoming NfV technologies against BAU solutions. The model has been applied to the vEPC use case using real performance indexes of NfV and BAU products available on the market, which have been publicly provided by independent third-party entities or declared in the product datasheets. For the sake of simplicity, we considered only the SGW device since, given its peculiarities that are mostly related to the data plane, it represents the network function that can hardly maintain suitable efficiency and performance levels when virtualized. Estimations have been dimensioned according to the 2016 traffic loads for Western Europe Telecom providers estimated by the same vendor of NfV and BAU products here considered. The NfV advantages related to service elasticity and redundancy have been explicitly considered according to different levels of over-dimensioning of network processing capacity.

Contrary to the common perception that NfV will be intrinsically green, the obtained results (even though limited to a specific, but popular, processor category) underline how, in its current implementation, this technology may lead to at least double of the environmental impact and energy requirements of the EPC. Moreover, our estimate also outlines how the cost of virtualization may rapidly increase up to five times or more, if the best combination of hardware and software is not applied, or if the packet processing latency is considered as an objective. All these figures suggest that, if not considered as a key objective in the mainstream design of NfV evolution, the energy

requirements may become a predominant factor for the TCO of NFV infrastructures, to hinder the foundations of commercial viability of this upcoming technological paradigm.

Therefore, in our opinion, the energy consumption aspect ought to be explicitly taken into account in the design of NFV architectural paradigms. Energy-aware solutions and mechanisms must be intrinsically part of the core specification of both the VNF internal structuring and operations and their control and orchestration planes, in order to allow the overall system to achieve suitable trade-offs between power and performance (within the desired upper bounds) that would consent a level of energy efficiency at least comparable to what could be reached in the BAU scenario.

#### ACKNOWLEDGMENT

This work has been supported by the INPUT (In-Network Programmability for next-generation personal cloUd service support) project, funded by the European Commission under the Horizon 2020 Programme (Grant no. 644672).

#### REFERENCES

- [1] E. Gelenbe, Y. Caseau, "The Impact of Information Technology on Energy Consumption and Carbon Emissions," *Ubiquity*, vol. 2015, pp. 1-15, Jun. 2015. [Online]. Available: <http://doi.acm.org/10.1145/2755977>
- [2] The EINS Consortium. Overview of ICT Energy Consumption (D8.1). Report FP7-2888021. European Network of Excellence in Internet Science. February 2013.
- [3] M. Chiosi *et al.*, "Network Functions Virtualization: An Introduction, Benefits, Enablers, Challenges & Call For Action", ETSI White Paper, Oct. 2012. [Online]. Available: [http://portal.etsi.org/NFV/NFV\\_White\\_Paper.pdf](http://portal.etsi.org/NFV/NFV_White_Paper.pdf).
- [4] ETSI, "NFV Use Cases." [Online]. Available: [http://www.etsi.org/deliver/etsi\\_gs/NFV/001\\_099/001/01.01.01\\_60/gs\\_NFV001v010101p.pdf](http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf)
- [5] 3GPP, "The Evolved Packet Core." [Online]. Available: <http://www.3gpp.org/technologies/keywords-acronyms/100-the-evolved-packet-core>.
- [6] E. Hernandez-Valencia, S. Izzo, B. Polonsky, "How will NFV/SDN Transform Service Provider OpEx?," *IEEE Network*, vol. 29, no. 3, pp. 60-67, May-June 2015.
- [7] R. Bolla, R. Bruschi, F. Davoli, C. Lombardo, L. Masullo, "The Expected Impact of Smart Devices Virtualization," 1st Internat. Workshop on Sustainability, Implementation and Resilience of Energy-Aware Networks (SIREN 2016), Kauai, Hawaii, USA, Feb. 2016.
- [8] Alcatel-Lucent, "Global "What if" Analyzer of NeTwork Energy Consumption (G.W.A.T.T.)." [Online]. Available: <http://gwatt.net/intro/1>.
- [9] H. Mellah, B. Sansò, "Routers vs Switches, How Much More Power Do They Really Consume? A Datasheet Analysis," *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Lucca, Italy, Jun. 2011.
- [10] The ECONET Project, "End-User Requirements, Technology Specification and Benchmarking Methodologies," Deliverable 2.1. [Online]. Available: <https://www.econet-project.eu/Repository/Document/291>.
- [11] Cisco Systems, Inc, "The Zettabyte Era: Trends and Analysis," Tech. Rep., Jun. 2014.
- [12] M. Al-Fares, A. Loukissas, A.Vahdat, "A Scalable, Commodity Data Center Network Architecture," *Proc. ACM SIGCOMM 2008 Conf. on Data Communication*, Seattle, WA, USA, Aug. 2008, pp. 63-74.
- [13] Feknous, T. Houdoin, B. Le Guyader, J. De Biasio, A. Gravey, J. A. T. Gijon, "Internet Traffic Analysis: A Case Study of two Major European Operators," *IEEE Symp. on Computers and Communication (ISCC)*, Madeira, Portugal, Jun. 2014, pp. 1-7.
- [14] "Liquid Net: Signalling is growing 50% faster than data traffic", Nokia Siemens Marketing Material, 2012.
- [15] H. Huang, P. Li, S. Guo, W. Zhuang, "Software-Defined Wireless Mesh Networks: Architecture and Traffic Orchestration", *IEEE Networks*, vol. 29, no. 4, pp. 24-30, Jul. 2015.
- [16] Nokia 7750 Service Router. [Online]. Available: <http://networks.nokia.com/portfolio/products/7750-service-router>.
- [17] C. Rossenhövel *et al.*, "Validating Nokia's IP Routing & Mobile Gateway VNFs", *LightReading*, 2016. [Online]. Available: [http://www.lightreading.com/nfv/validating-nokias-ip-routing-and-mobile-gateway-vnfs/d/d-id/720902?mc=RSS\\_LR\\_EDT](http://www.lightreading.com/nfv/validating-nokias-ip-routing-and-mobile-gateway-vnfs/d/d-id/720902?mc=RSS_LR_EDT)
- [18] Intel, "Introducing the Intel Xeon Processor E5-2600 v4 Product Family." [Online]. Available: <https://newsroom.intel.com/wp-content/uploads/sites/11/2016/04/intel-xeon-processor-e5-2600-v4-fact-sheet-x.pdf>
- [19] D. Woo, H.-S. Lee, "Extending Amdahl's Law for Energy-Efficient Computing in the Many-Core Era", *IEEE Comput.*, vol. 41, no. 12, pp. 24-31, 2008.