

The Expected Impact of Smart Devices Virtualization

Raffaele Bolla, Franco Davoli, Chiara Lombardo
CNIT – Research Unit of the University of Genoa
DITEN – University of Genoa
Genoa, Italy

Roberto Bruschi
CNIT – Research Unit of the University of Genoa
Genoa, Italy

Laura Masullo
DITEN – University of Genoa
Genoa, Italy

Abstract—Overcoming the typical ossification of the traditional TCP/IP-based Internet infrastructure will allow opening the way for more flexible and energy efficient paradigms, able to provide a sustainable support to the constantly increasing number of devices and services. To this goal, the INPUT Project will overcome current limitations by introducing computing and storage capabilities to edge network devices in order to allow users/telecom operators to create/manage private clouds “in the network”. In addition, these new capabilities will allow replacing smart devices, such as network-attached storage servers, set-top-boxes, video sensors etc. with their virtual images. Although this virtualization process can clearly bring to a reduction of the emissions, along with the lowering of capital (CAPEX) and operational (OPEX) expenditures, on the other hand it requires more computational capacity at server level, which may erase the savings produced by virtualization in the absence of a thorough management and planning. In this respect, this paper presents a mathematical model that analyzes the impact of different levels of virtualization on the overall energy efficiency by thoroughly outlining how the carbon footprint varies depending on the virtualization level of a device.

Keywords—energy efficiency evaluation, power management, SDN.

I. INTRODUCTION

The new access and service opportunities that have emerged on the market in the last few years, as a consequence of the diffusion of smart objects, like smartphones and tablets, have put a spotlight on the inefficiency of the current Internet infrastructures. In fact, although the Green House Gases (GHG) emissions due to the Information and Communication Technology (ICT) have been a serious concern for several years now ([1], [2], [3], [4]), the need for further computational and storage resources of the above mentioned smart devices has quickly exacerbated the issue, in particular by contributing to the exponential growth of the data centers needed to host cloud services.

Unfortunately, the existing network infrastructures are not suitable for the introduction of environmentally sustainable and highly performing alternatives [1], [5], as they still rely heavily on TCP/IP proprietary solutions, which are hardware-based and do not provide the flexibility needed to introduce new protocols and architectures for increasing scalability and energy efficiency.

In this respect, the INPUT Project [6] aims at overcoming the current Internet limitations by designing a novel infrastructure and paradigm to support personal cloud services in a more scalable and sustainable way. The INPUT technologies will enable next-generation cloud applications to go beyond classical service models (i.e., IaaS, PaaS, and SaaS) by moving cloud services closer to end-users and smart devices, in order both to avoid pointless network infrastructure and datacenter overloading, and to provide lower latency reactivity to services. To this goal, INPUT will exploit the capabilities provided by Network Function Virtualization (NFV) [7] for the design of the architectural solutions deployed in the access and edge network, and Software Defined Networking (SDN) [8] to support the control needed for the introduction of computing and storage capabilities to edge network devices, and consequently for moving cloud services closer to end-users and smart devices.

In addition, INPUT aims at replacing physical Smart Devices (SD), usually placed in users’ homes (e.g., network-attached storage servers, set-top-boxes, video recorders, home automation control units, etc.) or deployed around for monitoring purposes (e.g., sensors), with their “virtual images,” providing them to users “as a Service” (SD as a Service – SDaaS). This aspect is of crucial importance, as in 2011 and only in USA, set-top-boxes were costing Americans 3 billion dollars in electricity charges each year, a figure equivalent to nine power plants and to GHG emissions of 15 MtCO_{2e} per year [9]. Similar figures can also be extended to Europe.

However, the advantages obtained from the complete dematerialization of set-top-boxes and similar end-user devices would be vanishing if the increased computational load on Telco servers caused GHG emissions overcome the savings obtained through virtualization. For this reason, this paper presents a mathematical model to determine under what circumstances replacing physical objects with their virtual images reduces Carbon Footprint and GHG emissions. To do so, the model analyzes the impact of different levels of virtualization on the overall energy efficiency by thoroughly outlining how the carbon footprint varies depending on the virtualization level of a device.

The paper is organized as follows. Section II presents the model characterizing the carbon footprint in the two considered scenarios, while Section III provides an extension of such model

in the presence of power management capabilities. A use case based on the former model can be found in Section IV, and conclusions are drawn in Section V.

II. BREAKDOWN OF THE CARBON FOOTPRINT

The huge number of connected devices people have at home and in offices, often replicated in both places, has a twofold negative impact on the GHG emissions: (i) producing many objects means consuming a large amount of resources, material, energy, which implies the emission of non-negligible quantities of GHG in the atmosphere; (ii) although such objects are actually used for a small portion of time during the day, they are often left switched on all day long, collectively causing the additional emission of big quantities of GHG in the atmosphere.

Among the technologies that could bring to a reduction of the emissions and to a lowering of CAPEX and OPEX definitely appreciable by service providers, virtualization represents a viable answer. In this respect, the INPUT technologies could have a significant potential impact on the ICT carbon footprint: thanks to the replacement of the physical objects with their virtual images, the production of and subsequent energy supply to millions of physical connected smart devices could be reduced.

To this goal, this section will introduce a simple model that describes how the carbon footprint varies depending on the virtualization level of a device (partially, fully or not virtualized). Moreover, our model also takes into account the possibility of optimizing the network infrastructure by means of power management and consolidation criteria, as will be presented in Section III. The model has the goal of understanding how different levels of virtualization could impact on the overall energy efficiency, and whether replacing physical objects with their virtual images could reduce Carbon Footprint and GHG emissions. To do this, we will compare the Business As Usual (BAU) and the INPUT scenarios.

The term of comparison will be the Total Carbon Footprint, defined as:

$$TCF = ECF + OCF \quad (1)$$

where *ECF* (Embodied Carbon Footprint) is the amount of GHG emitted in the atmosphere due to manufacturing, transporting, packing, and disposing of the product the product, and *OCF* (Operating Carbon Footprint) is the amount of GHG due to the utilization of the product. Both are measured in kgCO₂eq units, which represents the contribution of the GHG emissions to the global warming in kg of CO₂ equivalent. The next subsections will provide a characterization of the carbon footprint of the BAU and INPUT scenarios.

A. The Business as Usual Scenario

In the BAU scenario a large amount of computational resources is replicated multiple times and used sparsely, causing a waste of these resources and of the consequent power. In order to evaluate a device total energy consumption, it is necessary to take into account several factors: firstly, the energy consumed to manufacture, transport, pack, and also dispose of the product, which contributes to the *ECF*; then, the average life of the product (the shorter the life, the greater will be its negative

impact on the total energy consumption per year); and, finally, the energy consumption caused by its utilization.

The *TCF* per year of a single device in the BAU scenario can be defined as :

$$tcf_{1y}^{BAU} = \frac{ecf^{(A)} + \tau * ocf_{1y}^{(A)}}{\tau} \quad (2)$$

where τ is the lifetime of the appliance, expressed in years, $ecf^{(A)}$ is the *ECF* as defined above; $ocf_{1y}^{(A)}$ is the *OCF* per year, calculated as in (3):

$$ocf_{1y}^{(A)} = F * 24 * 365 * \varphi \quad (3)$$

Note that from now on, we will denote the *TCF*, *ECF* and *OCF* of a single device by using the lowercase characters, whereas the uppercase format will be used for denoting the carbon footprint quantities of multiple devices. 24 and 365 are the factors used to convert the power, φ , to the annual energy consumption. Hence, $ocf_{1y}^{(A)}$ is obtained multiplying the annual energy consumption by F , which is the conversion factor between energy consumption, in kWh, and kgCO₂eq units. Its value mainly depends on the energy source (coal, gas, oil, hydro, etc.). Values of F used in similar studies are 0.6 kgCO₂eq/kWh [10], 0.4582 kgCO₂eq/kWh [11]. In our example, we have decided to use the value, 0.6 kgCO₂eq/kWh, as suggested by [10].

For the sake of brevity, in the remaining of the paper, we will set:

$$G = F * 24 * 365 \quad (4)$$

B. The INPUT Scenario

The adoption of the INPUT technologies allows removing physical appliances, currently located at users' homes, and providing users with the virtual images of these appliances. The process of virtualization can remove redundancy and reduce consumption at the end-user level, but at the same time it can clearly cause more computational effort to be moved to the Telco Operator and Cloud Service Provider sides.

Thus, in order to evaluate whether it is worth replacing physical objects with their virtual images, it is necessary to estimate the *ECF* of servers, and how their energy consumption changes in the presence of this additional computational effort. In particular, it is important to evaluate how many objects could be hosted on a server: this quantity depends on the functionalities to be virtualized, therefore on the computational and storage requirements, and on the category of server used (i.e., volume, mid-range, high-end). In order to limit the cost and to scale, a large number of virtualized devices need to be integrated on a limited number of servers: understanding which economy of scale could allow decreasing consumption is one of the objectives of this study.

As also envisaged in [12], there could be coexistence between virtualized and non-virtualized devices. The same concept could also apply to single devices, allowing functions to be handled separately and different virtualization levels to be reached step-by-step.

TABLE I. DATA RELEVANT TO A SINGLE STB: EMBODIED CARBON FOOTPRINT, AVERAGE LIFETIME, AVERAGE NUMBER OF STBS IN A HOUSEHOLD, AND ANNUAL ENERGY CONSUMPTION.

Device	ecf^{STB} [kgCO ₂ eq]	τ^{STB} [y]	Ω [#]	W_{1y}^{STB} [kWh]
STB	25.00	3.9	0.52	52.85

In order to characterize how the carbon footprint changes in the presence of different virtualization paradigms, we define the following parameters:

- ξ , ($0 \leq \xi \leq 1$), as the “grade of virtualization” of a smart device, i.e., how much an appliance can be virtualized as a service on a server ($\xi = 1$, object fully virtualized, $\xi < 1$, object partially virtualized);
- γ , as the portion of a server capacity, needed to run the fully virtualized smart object. In other words, γ is the inverse of the number of smart objects of the same type, n , that could run on a server, i.e.:

$$\gamma = \frac{1}{n} \quad (5)$$

With these parameters, the TCF per year for a single device in the INPUT scenario can be written as:

$$tcf_{1y}^{INPUT} = \frac{ecf^{(s)} + \tau^{(s)} * ocf_{1y}^{(s)}}{\tau^{(s)}} \gamma \xi + \frac{ecf^{(A)} + \tau^{(A)} * ocf_{1y}^{(A)}}{\tau^{(A)}} (1 - \xi) \quad (6)$$

where the first term refers to the amount of TFC due to the virtualized part that runs on servers (ECF , τ and OCF relevant to the server), and the second term refers to the non-virtualized physical appliance at the user’s home.

III. THE IMPACT OF POWER SAVING MECHANISMS ON THE CARBON FOOTPRINT

Virtualizing smart objects also adds further advantages. Although objects at users’ homes are most of the time switched on, they are actually used only for a fraction of the day. The architecture and capabilities of the servers located on the operator network facilitate the incorporation of optimization strategies for resource management, for instance to support power management.

In this respect, let us now focus on power management and how dynamic power scaling and stand-by mechanisms impact on the server energy consumption, i.e. on $ocf_{1y}^{(s)}$.

Dynamic power scaling enables the modulation of energy absorption according to the actual workload. Sleeping/standby approaches are used to smartly and selectively drive unused network/device portions to low standby modes and to wake them up only if necessary. Details on this topic can be found in [13].

In Figure 1, a qualitative power profile of a server supporting both stand-by (flat part) and dynamic power scaling (linear part between $\varphi_{min}^{(s)}$ and $\varphi_{max}^{(s)}$) versus workload is represented.

In the figure, the following parameters can be identified:

- $\varphi_{st-by}^{(s)}$, the stand-by power, when the device is sleeping in standby mode;

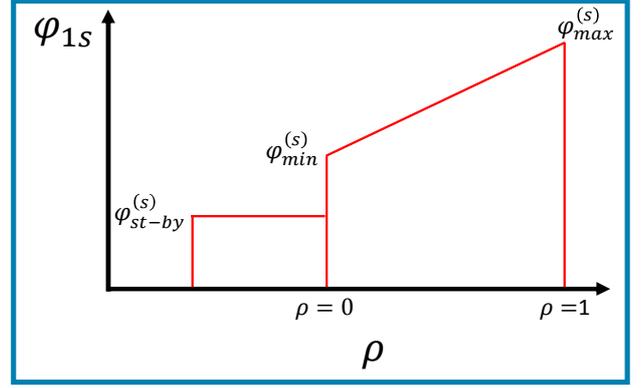


Figure 1. Power profile of a server due to power management techniques: namely, dynamic power scaling and stand-by.

- $\varphi_{min}^{(s)}$, the minimum power, when the device is active, but operating at the minimum speed;
- $\varphi_{max}^{(s)}$, the maximum power, when the device is fully active and operating at the maximum speed;
- ρ , the fraction of time in which the device is active at the maximum power.

First, let us analyze only the energy consumption related to dynamic adaptation. For the sake of simplicity but without losing generality, we assumed that power consumption values between the $\varphi_{min}^{(s)}$ and $\varphi_{max}^{(s)}$ vary linearly with respect to ρ ; therefore, the power profile of an active server can be written as:

$$\varphi_{1s} = \rho \varphi_{max}^{(s)} + (1 - \rho) \varphi_{min}^{(s)} \quad (7)$$

If N_s is the total number of servers present in the scenario, we can define $M < N_s$, as the number of servers staying active after applying consolidation policies and $N_s - M$ as the number of servers in stand-by. The total workload previously handled by servers now in stand-by is:

$$(N_s - M) \cdot \rho \quad (8)$$

Notice that each server is supposed to have normalized load ρ . This workload has to be redistributed on the M active servers. Therefore, each active server has an extra load of

$$\varepsilon = \frac{(N_s - M) \rho}{M} \quad (9)$$

By adding the ε term to each ρ in Equation (7), the power of each active server is modified in:

$$\varphi_{1s}^{PM} = \rho \frac{N_s}{M} \varphi_{max}^{(s)} + \left(1 - \rho \frac{N_s}{M}\right) \varphi_{min}^{(s)} \quad (10)$$

and the total power, due to both active and stand-by servers, is

$$\varphi_{TOT}^{PM} = M \varphi_{1s}^{PM} + (N_s - M) \varphi_{st-by}^{(s)} \quad (11)$$

By substituting (10) in (11), the OCF per year for all the N_s servers can be calculated, according to (3), as:

$$OCF_{1y}^{INPUT} = G \left[\rho N_s \varphi_{max}^{(s)} + (M - \rho N_s) \varphi_{min}^{(s)} + (N_s - M) \varphi_{st-by}^{(s)} \right] \quad (12)$$

Therefore, from (6), the TCF_{1y}^{INPUT} per year due to fully virtualized smart objects ($\xi = 1$) can be written as:

$$TCF_{1y}^{INPUT} = \frac{N_s * ecf^{(s)} + \tau^{(s)} * G [\rho N_s \varphi_{max}^{(s)} + (M - \rho N_s) \varphi_{min}^{(s)} + (N_s - M) \varphi_{st-by}^{(s)}]}{\tau^{(s)}} \quad (13)$$

This global equation takes into account all the above-mentioned factors; namely, virtualization, dynamic adaptation and stand-by/consolidation techniques. It is worth noting that, in the absence of consolidation policies, all the N_s servers are always kept on and active, which results in setting $N_s = M$ in (13). In the same way, removing dynamic adaptation means setting, for all the N_s servers, the power at the maximum value, without any modulation, i.e. $\varphi_{max}^{(s)} = \varphi_{min}^{(s)} = \varphi^{(s)}$. In this scenario, all servers are up, operating at the maximum speed.

IV. THE SET-TOP-BOX EXAMPLE

In order to evaluate the impact of the INPUT technologies on the ICT carbon footprint, and to understand how replacing physical objects with their virtual counterparts could change the GHG emission in the atmosphere, in the following we will analyze a practical example of smart device virtualization. Specifically, we have selected an object that we estimate could be fully virtualized, namely the Set-Top-Box (STB). In particular, we have estimated the Carbon Footprint produced by STBs in both scenarios, BAU and INPUT, and we also present results for two partial scenarios: the first one, where only virtualization is considered, but power management is not available; and the second one, where both virtualization and dynamic adaptation are considered, but no consolidation policies are applied.

To do so, we will use real data, collected by analyzing and comparing several studies and surveys conducted on the energy consumption of home equipment (for example from [1], [10], [14], [15]) regarding STBs, and studies related to ICT consumption that focus on servers ([1], [16]). The STB example refers to the European scenario: also in this case, numbers relevant to the European households are collected from other studies [2].

In the following, we will first calculate the TCF for the BAU scenario, according to the equations in Section II.A; then, we will repeat the same evaluation for the INPUT scenario, as stated above, also including intermediate scenarios that do not take into account Dynamic Adaptation and Stand-by/Consolidation policies, according to the equations in Section III. Lastly, we will compare the different results.

A. STBs Footprint in the BAU Scenario

Values relevant to STBs found in the literature considerably differ from one another, mainly due to the typology of the considered STB (simple, complex, with or without High Definition feature, etc.) [2], [14], [17].

In our example, to evaluate the TCF_{1y}^{BAU} according to (2), we use data relevant to Free-to-air STB, as collected in [14] and reported in Table I. The ECF value has been taken from [10].

In Table I, τ^{STB} is the average lifetime of a STB, the time it “lives” in a house before being substituted, Ω is the average number of appliances per household, and W_{1y}^{STB} is the average annual consumption per appliance. The value can be easily worked out, by considering a power of almost 10.2 W in active

mode (for a time of 4 hours a day), and 5.2 W in stand-by mode (for the remaining 20 hours a day).

The total number of STBs is:

$$N_{STB} = N_{HH} \Omega \quad (14)$$

where N_{HH} , which amounts for 165 million, is the estimated number of European (EU-15) households in 2010 [2].

The total amount of the TCF per year in Europe due to all the estimated STBs is:

$$TCF_{1y}^{BAU} = N_{STB} \frac{ecf^{STB} + \tau^{STB} * ocf_{1y}^{STB}}{\tau^{STB}} \quad (15)$$

where ocf_{1y}^{STB} is calculated, according to (3), as:

$$ocf_{1y}^{STB} = 0.6 * W_{1y}^{STB} \quad (16)$$

The result is:

$$TCF_{1y}^{BAU} = 3.27 \text{ MtCO}_2\text{eq} \quad (17)$$

which corresponds to 38.12 kgCO₂eq for each appliance per year.

B. STBs Footprint in the INPUT Scenario

In order to calculate TCF_{1y}^{INPUT} , numbers related to servers are required. With this goal, a short digression about servers is necessary.

Servers can be categorized in 3 groups, according to [1]: (i) volume servers, <\$25,000 per unit; (ii) mid-range servers, between \$25,000 and \$500,000 per unit; (iii) high-end servers, >\$500,000 per unit. For our calculation, we believe the correct class of servers to be used is the mid-range.

Table II reports data related to this group. $\tau^{(s)}$ is the average lifetime; we have decided to use 5 years for servers, evaluating data reported, in the literature ([18]-[21]), $ecf^{(s)}$ is the Embodied Carbon Footprint for a server, taken from [16], $\varphi^{(s)}$ is the instantaneous power of a server in active mode. To be noted that $\varphi^{(s)}$, as in [1], is calculated as

$$\varphi^{(s)} = \varphi * (1 + SPC + CPC) * PUE \quad (18)$$

where φ is the power per server (607 W), SPC is the percentage of storage power consumption (24% of total server power consumption) and CPC is the percentage of communication power consumption (15% of total server power consumption); moreover, the total server power consumption is increased by a factor of 1.83 (PUE , Power Usage Effectiveness), which comprises cooling, power provisioning and power backup systems.

STBs are supposed to be fully virtualized, ($\xi = 1$); therefore, the second term equals zero and, from this point forward, (6) for TCF_{1y}^{INPUT} will be reduced to:

TABLE II. DATA RELEVANT TO A SINGLE MID-RANGE SERVER: AVERAGE LIFETIME, EMBODIED CARBON FOOTPRINT AND POWER IN ACTIVE MODE.

	$\tau^{(s)}$ [y]	$ecf^{(s)}$ [kgCO ₂ eq]	$\varphi^{(s)}$ [kW]
Mid-range server	5	360	1.54

$$tcf_{1y}^{INPUT} = \frac{ecf^{(s)} + \tau^{(s)} * ocf_{1y}^{(s)}}{\tau^{(s)}} \gamma \quad (19)$$

To determine the TCF_{1y}^{INPUT} value considering the data relevant to the European panorama, we have to determine the number of servers needed to virtualize the number of STBs. We can calculate N_s as:

$$N_s = N_{STB} \gamma = N_{HH} \Omega \gamma \quad (20)$$

where N_{STB} is calculated according to (14) and γ , the portion of a server needed to virtualize a STB, is defined according to (5).

According to Figure 1, the power profile for a server implementing Power Management techniques can be expressed with the values reported in Table III.

The value for $\varphi_{max}^{(s)}$ is calculated from (18); values for $\varphi_{min}^{(s)}$ and $\varphi_{st-by}^{(s)}$ have been experimentally measured on servers belonging to the mid-range category.

Therefore, from (13), the TCF_{1y}^{INPUT} can be calculated as:

$$TCF_{1y}^{INPUT} = N_s \frac{ecf^{(s)} + \tau^{(s)} * G * [\rho \varphi_{max}^{(s)} + \frac{M}{N_s} \varphi_{min}^{(s)} - \rho \varphi_{min}^{(s)} + (\frac{N_s - M}{N_s}) \varphi_{st-by}^{(s)}]}{\tau^{(s)}} \quad (21)$$

As stated in Section II.B, when stand-by/consolidation policies are not implemented, it means that all servers are on. The TCF for this scenario can be obtained by setting $N_s = M$: in this way, we have TCF_{1y}^{DA} , where virtualization and Dynamic Adaptation are considered, but not the consolidation effect:

$$TCF_{1y}^{DA} = N_s \frac{ecf^{(s)} + \tau^{(s)} * G * [\rho \varphi_{max}^{(s)} + (1 - \rho) \varphi_{min}^{(s)}]}{\tau^{(s)}} \quad (22)$$

Lastly, TCF_{1y}^V , due only to virtualization, without any Power Management technique, can be obtained by setting $\varphi_{max}^{(s)} = \varphi_{min}^{(s)} = \varphi^{(s)}$ in (23):

$$TCF_{1y}^V = N_s \frac{ecf^{(s)} + \tau^{(s)} * G * \varphi^{(s)}}{\tau^{(s)}} \quad (23)$$

For our calculation, we fix $\rho = \frac{4}{24}$, presuming STBs are used only 4 hours a day [2]. Moreover, let us assume that, for reliability reasons, the consolidation policy always guarantees a ratio between the number of active servers, M , and total number of server N_s equal to:

$$\frac{M}{N_s} = 2\rho \quad (24)$$

In order to compare the Total Carbon Footprint of the INPUT scenarios with the BAU one, we define saving as:

$$S = \frac{TCF_{1y}^{xxx} - TCF_{1y}^{BAU}}{TCF_{1y}^{BAU}} \quad (25)$$

TABLE III. POWER VALUES FOR A SERVER WITH POWER MANAGEMENT CAPABILITIES.

$\varphi_{max}^{(s)}$ [kW]	$\varphi_{min}^{(s)}$ [kW]	$\varphi_{st-by}^{(s)}$ [kW]
1.540	0.300	0.012

where TCF_{1y}^{xxx} can assume, each time, the value calculated in (21), (22), or (23) for each of the three INPUT scenarios, respectively.

C. Carbon Footprint Savings through the INPUT Adoption

Figure 2 reports the percentage of carbon footprint saving obtained with the introduction of the INPUT technologies (all the 3 scenarios) with respect to the BAU scenario. The INPUT scenarios are calculated by varying the number of STBs, here indicated as n_{STB} , from 25 to 800.

In Figure 2, the blue line represents the savings obtained by considering only the presence of Virtualization (V), the red line savings obtained by adopting both Virtualization and Dynamic Adaptation (DA), and finally the green line also includes Consolidation policies (C).

The figure shows that the virtualization of STBs provides positive savings only for a number of virtualized devices per server above $n_{STB} > 225$. Below this number, the savings obtained through virtualization are not sufficient to cover for the increased server consumption due to the added computational load. As the value of n_{STB} increases from 225, savings quickly rise for all the three INPUT scenarios. The Virtualization scenario reaches a saving of 40% for $n_{STB}=325$, and then continues gaining at a slower pace up to 75% for $n_{STB}=775$. The further adoption of Dynamic Adaptation visibly increases the savings with respect to the single Virtualization: 40% saving is obtained with only $n_{STB}=225$, and total savings reach up to 90%. The scenario that includes also Consolidation presents visible differences from the DA one for lower values of n_{STB} . This behavior is due to the stronger impact that consolidation policies can have on less loaded servers, which allows to power off a higher number of sub-components, while such impact cannot be as critical on a fully loaded server.

V. CONCLUSIONS

The INPUT technologies could have a significant potential impact on the ICT carbon footprint. In fact, thanks to the replacement of the physical objects with their virtual images, the INPUT paradigms may allow reducing the production of and subsequent energy supply to millions of physical connected smart devices, e.g., Set-Top-Box, Network Attached Storage.

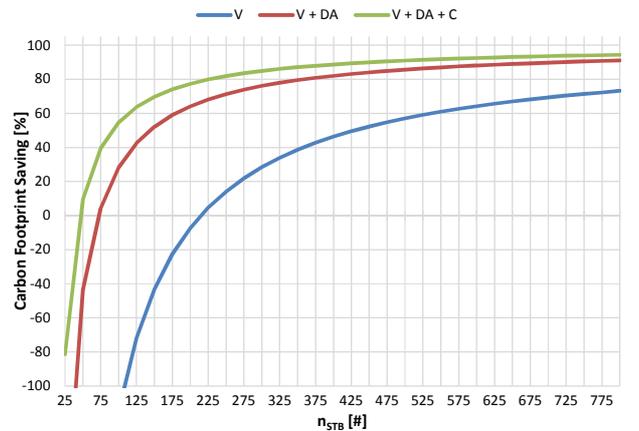


Figure 2. Saving, in terms of TCF, of the INPUT scenario respect the BAU one.

However, as removing physical objects from homes implies adding servers at the network edges, the operational expenditure of the network itself is increased.

Thus, in order to evaluate if virtualizing is actually worthwhile, both in terms of energy efficiency and economically, this paper has proposed a simple model to describe how the carbon footprint varies depending on the virtualization level of a device taking also into account the availability of power management and consolidation criteria. Results have outlined the potential savings that can be achieved in the presence of different power management capabilities and server load conditions. In the presence of adaptive rate and stand-by capabilities, virtualization can provide up to 90% of savings compared to the carbon footprint obtained 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 (Call H2020-ICT-2014-1, Grant no. 644672).

REFERENCES

- [1] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, P. Demeester, "Trends in Worldwide ICT Electricity Consumption from 2007 to 2012," *Comp. Commun.*, vol. 50, pp. 64-76, Sept. 2014.
- [2] European Commission DG INFSO, Final Report – Impacts on ICT on Energy Efficiency, 2008.
- [3] V. Coroama, L. M. Hilty, "Energy Consumed vs. Energy Saved by ICT - A Closer Look", In: Wohlgemuth, V., Page, B., Voigt, K. (Eds.): *Environmental Informatics and Industrial Environmental Protection: Concepts, Methods and Tools*, 23rd International Conference on Informatics for Environmental Protection, pp 353-361, 2009.
- [4] Global e-Sustainability Initiative (GeSI), "SMARTer2020: The Role of ICT in Driving a Sustainable Future", Report, URL: <http://gesi.org/SMARTer2020>.
- [5] A. Manzalini, "Future edge ICT networks," *IEEE ComSoc MMTC E-Lett.*, vol. 7, no. 7, pp. 1-4, Sep. 2012.
- [6] The In-Network Programmability for next-generation personal cloud service support (INPUT) Project, URL: <http://www.input-project.eu/>.
- [7] M. Chiosi, et al., "Network Functions Virtualization: An Introduction, Benefits, Enablers, Challenges & Call For Action", ETSI White Paper, Oct. 2012, URL: http://portal.etsi.org/NFV/NFV_White_Paper.pdf.
- [8] Open Network Foundation (ONF), "OpenFlow Switch Specification," Version 1.4.0, URL: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf>.
- [9] U.S. Department of Energy, "U.S. Energy Department, Pay-Television Industry and Energy Efficiency Groups Announce Set-Top Box Energy Conservation Agreement," Dec. 2013, URL: energy.gov/articles/us-energy-department-pay-television-industry-and-energy-efficiency-groups-announce-set-top.
- [10] J. Malmodin, D. Lundén, A. Moberg, G. Andersson, Nilsson, "Life Cycle Assessment of ICT - Carbon Footprint and Operational Electricity Use from the Operator, National, and Subscriber Perspective in Sweden," *Journal of Industrial Ecology*, vol. 18, no. 6, pp. 829-845, Dec. 2014.
- [11] European Commission DG INFSO, Final Report – Impacts on ICT on Energy Efficiency, 2008.
- [12] ETSI GS NFV, "Network Function Virtualisation (NFV); Use Cases," GS NFV 001 V1.1.1.
- [13] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, D. Suino, C. Vassilakis, A. Zafeiropoulos, "Cutting the Energy Bills of Internet Service Providers and Telecoms through Power Management: An Impact Analysis," *Comp. Netw.*, vol. 56, no. 10, pp. 2320-2342, Jul. 2012.
- [14] http://www.energyrating.gov.au/wp-content/uploads/Energy_Rating_Documents/Library/Standby_Power/Standby_Power/E3-2010-Intrusive-Survey-FINAL-Report.pdf.
- [15] A. de Almeida, P. Fonseca, B. Schlomann, N. Feilberg, "Characterization of the Household Electricity Consumption in the EU, Potential Energy savings and Specific Policy Recommendations," *Energy and Buildings*, vol. 43, no.8, pp. 1884-1894, Aug. 2011.
- [16] P. Teehan, M. Kandlikar, "Comparing Embodied Greenhouse Gas Emissions of Modern Computing and Electronics Products," *Environ. Sci. Technol.* Vol. 47, no. 9, pp. 3997-4003, 2013.
- [17] G. Hardy et al., "Better Viewing, Lower Energy Bills, and Less Pollution: Improving the Efficiency of Television Set-Top Boxes," *Natural Resource Defense Council*, 2011.
- [18] <http://www.intel.com/support/motherboards/server/sb/CS-034303.htm>
- [19] http://www.dell.com/content/topics/global.aspx/power/en/ps3q02_shetty?c=us
- [20] <https://fusetg.com/life-span-server/>
- [21] <http://ipoint-tech.com/determining-your-server-life-expectancy/>