



HAL
open science

Assessing VoD pressure on network power consumption

Gaël Guennebaud, Aurélie Bugeau, Antoine Dudouit

► **To cite this version:**

Gaël Guennebaud, Aurélie Bugeau, Antoine Dudouit. Assessing VoD pressure on network power consumption. ICT4S - International Conference on Information and Communications Technology for Sustainability, Jun 2023, Rennes, France. hal-04059523

HAL Id: hal-04059523

<https://hal.science/hal-04059523>

Submitted on 6 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Assessing VoD pressure on network power consumption

Gaël Guennebaud
Inria, Univ. Bordeaux
Talence, France
gael.guenebaud@inria.fr

Aurélié Bugeau
Univ. Bordeaux, CNRS, LaBRI UMR 5800
IUF
Talence, France
aurelie.bugeau@u-bordeaux.fr

Antoine Dudouit
Univ. Bordeaux, LaBRI, UMR 5800, Inria
Talence, France

Abstract—Assessing the energy consumption or carbon footprint of data distribution of video streaming services is usually carried out through energy or carbon intensity figures (in Wh or gCO₂e per GB). In this paper, we first review the reasons why such approaches are likely to lead to misunderstandings and potentially to erroneous conclusions. To overcome those shortcomings, we propose a new methodology whose key idea is to consider a video streaming usage at the whole scale of a territory, and evaluate the impact of this usage on the network infrastructure. At the core of our methodology is a parametric model of a simplified network and Content Delivery Network (CDN) infrastructure, which is automatically scaled according to peak usage needs. This allows us to compare the power consumption of this infrastructure under different scenarios, ranging from a sober baseline to a generalized use of high bitrate videos. Our results show that classical efficiency indicators do not reflect the power consumption increase of more intensive Internet usage, and might even lead to misleading conclusions.

Index Terms—Access network, Internet power consumption, video streaming, peak usage

I. INTRODUCTION

Internet traffic has grown exponentially in the last decade. Cisco [1] projected an increase of IP traffic from 122 exabytes per month in 2017 to 365 in 2022. Studying energy consumption and impacts on climate change of data transmission over the Internet has therefore received much attention over the last decade. Such work usually strives to estimate the overall yearly energy consumption of the Internet (in TWh/year) from which energy intensity estimates of data transmission (in Wh/GB) are extracted. A recent review [2] of these works shows high variability in the results and even inconsistencies between overall energy consumption, energy intensity, and Internet traffic. Those variations are likely explained by differences in Internet modeling, system boundaries, hypotheses, and methodologies.

Limits of energy intensity estimates: Despite those high uncertainties, such energy intensity estimates are frequently used to assess the energy consumption, or environmental impacts, of data transmission of a given Internet service such as, for instance, video-streaming. A commonly raised question is: “what is the energy consumption of transferring one GB of data?”, which exhibits severe limitations.

This study has been financially supported by the French Research Agency through the PostProdLEAP project (ANR-19-CE23-0027-01)

Firstly, we argue such a question is ill posed as it depends on numerous variables that go way beyond distinguishing the core, fixed-access, and mobile-access networks. Some other variable examples include the technological maturity of the considered network (from old energy-intensive equipment, to the newest generation) or the actual route taken by the data: a data intensive service hosted in the US but used in Europe will have a very different impact on the network than another service hosted in the same city than its primary users.

Secondly, unlike what their units convey, those numbers exhibit a poor proportionality with the physical reality. Indeed, one can quickly come to the false conclusion that, e.g., reducing by two the amount of data of a given usage will reduce by two its impacts. This limitation might not be a problem when those numbers are exclusively used as an *attributional* key to allocate the overall shared footprint across the different usages retrospectively. Their use in a *consequential* manner is, however, very frequent and misleading both for a short or long term point of view [3]. On the shorter term because the infrastructure is permanently switched on, and the volume of data passing through it at a given time has very little influence on the power consumption of the equipment (especially for fixed network equipment). On the longer term, one could expect a correlation because if the traffic volume increases, the traffic peak is expected to increase too, yielding to an increase in the infrastructure equipment, and thus an increase of the overall consumption of the infrastructure [4]. Conversely, if the traffic is maintained or decreased, oldest equipments might be replaced by smaller and more efficient ones when renewed. However, this long term correlation is only partial because i) only a subset of network hardware is subject to such correlation with peak demands, ii) the energy efficiency of such equipment improves quickly over time, and iii) two identical volumes of transferred data might have very distinct effects on local traffic peaks (because of different bitrates, different routes, or different *burstiness* [5]).

Thirdly, such intensity numbers (in Wh/GB or gCO₂e/GB) are only *efficiency* indicators hiding the true *absolute* energy consumption or absolute impacts, which are the only numbers that really matter. This observation combined with the aforementioned second point yield a paradox: increasing the total amount of traffic increases load percentage and enables

scaling gains. Both lead to an improvement (i.e, a decrease) of those efficiency indicators, while the absolute impacts increase. In contrast, sobriety behaviors are certain to maintain or decrease absolute impacts even though they might degrade those efficiency indicators.

Fourthly, such indicators tend to put the responsibility of the impacts solely on the consumer side, while we believe environmental impacts are systemic problems that concern manufacturers, content providers and users. In other words, by hiding the global absolute impacts, such indicators tend to lead to individualization at the expense of a collective vision at which different and more effective levers could emerge.

Contributions: To address those shortcomings and convey a more realistic understanding of data transmission, we propose a new methodology whose central idea is to consider a given Internet usage at the whole scale of an appropriately chosen territory, and evaluate the impact of this usage, or variants of this usage, on the IT infrastructure. This is accomplished through a parametric bottom-up network modeling of a simplified network infrastructure. Starting from a minimalistic infrastructure (the *baseline*), our model automatically scales the required hardware according to peak usage scenarios, from which absolute power and energy consumption can be estimated and compared to the baseline or other scenarios. Our simplified model relies on a tree representation of the network infrastructure, allowing us to adjust it to peak access rates in a hierarchical manner.

Having a precise estimate of the overall energy intensity of the Internet is out of the scope of this paper. The version of the parametric model we propose in this paper is rather designed to analyze and compare given use-cases relatively. It is an approximation of the reality with the aim of comparing past and future scenarios for a known usage under the same boundaries. For the sake of simplicity and clarity, we restrict ourselves to the scale of a *territory* with a limited geographical extent. The general principle of our methodology is presented in Section III.

In Section IV, we demonstrate our methodology on video-on-demand (VoD) at the scale of the France territory. This use case implies a high data traffic playing an important role on the scale of current infrastructures. Through our experiments, we propose an evaluation of the impacts that would result if watching only HD video in contrast to higher quality streams. Our model enables analyzing which network equipment is most impacted by each parameter of the scenario. The proposed methodology is therefore a first step towards a better understanding of the consequences of politic, industrial or societal decisions on infrastructure sizing. In particular, our results confirm the aforementioned claims and paradox, and show that using a simple efficiency indicator may lead to misleading decisions. Owing to the lack of data, in this paper we restrict our analysis to the estimation of power consumption, however, our model could easily be extended to account for the manufacturing and other life-cycle phases to estimate other environmental indicators.

II. RELATED WORKS

This section presents prior related works on energy consumption of data transmission. The ICT infrastructure is commonly decomposed into three parts: datacenters, transmission networks and user equipment. Many works have tackled the estimation of energy intensity for each part. We here mainly focus on the network. Note that some earliest works estimating the energy consumption of the Internet included datacenters while others did not, but it is now more common to consider them apart [6]. Early works [7], [8] were bottom-up approaches. They were based on an inventory of all US computing and networking equipment and their average yearly energy consumption to compute the total energy of networking devices in the US. Those overall energy consumption were then normalized by some estimates of the overall traffic volume, yielding to energy intensity indicators. Such indicators have been re-evaluated on a regular basis, and we refer to Aslan *et al.* [9] and Coroama [2] for recent surveys and analysis. In the rest of this section, we rather focus on network models and alternative approaches.

Baliga *et al.* [10] proposed one of the first bottom-up network infrastructure model. It is composed of the access, metro and core networks. The total power is the sum of the power P_i of each equipment multiplied by i) the power usage effectiveness (PUE), ii) a redundancy factor η to ensure functionality in case of failure, and iii) a scaling factor based on peak access rate R_i over individual capacity C_i :

$$P = \sum_i PUE \times \eta \times P_i \times \left[\frac{R_i}{C_i} \right].$$

Baliga's model has inspired many later works [6], [11]–[14]. In particular, Hinton *et al.* [13] presented an extension to assess the energy consumption of optical networks for different services and scenarios. Each network element is associated with an affine power profile distinguishing the static power P_{idle} and the linear component which is assumed to be proportional to the current throughput (in bit/s). They observed that for such fixed network elements, this proportional part is very small with $P_{idle} > 0.9P_{max}$. This observation is also confirmed by Malmodin [15]. With such an affine power model, allocating the proportional part boils down to a simple volume-based allocation. To accommodate for the idle power, they proposed different allocation strategies, and in particular one based on relative throughputs. Malmodin's power model [15] is also based on an affine power-profile. The idle power is first equally spread to each line, and then distributed to the potentially multiple users and devices using this line, which can be rather sketchy to do in practice. In a similar vein, Ullrich *et al.* [16] described an hybrid allocation: duration-based for the customer premises equipment (CPE) and access network equipment, and volume-based for the core network. All those strategies [13], [15], [16], however, ignore the energy consumption during standby time, while the last two assume that the idle power consumption is unrelated to traffic demand, which tends to artificially minimize the network part of the

impacts of a given usage. Our approach overcomes those shortcomings by replacing the allocation issues by a more systemic view, and estimating the ideal idle-power for a given usage or service.

To better understand network energy consumption or GHG emissions, several studies focused on a narrow and specific use case. For instance, Schien *et al.* [17] used traceroute data to estimate the number and type of network devices involved in digital media transmission. Coroama *et al.* [12] considered a 40 Mbps videoconferencing transmission between Switzerland and Japan, and modeled all Internet nodes and links along the way, distributing the energy according to the relative traffic volumes. Ficher *et al.* [18] estimated the carbon footprint of transmitting one gigabyte of data on a specific segment of the RENATER network. Golard *et al.* [19] evaluated and projected the total energy consumption of broadband radio networks at the scale of Belgium. Another use-case that has received attention is the assessment of the carbon footprint of watching one hour of video streaming, as discussed in Section IV.

III. METHODOLOGY

In this section, we present a general overview of our methodology. A concrete instance on the VoD streaming use case is given in Section IV. Let us recall that our goal here is not to estimate the impact of an existing infrastructure, but rather to estimate the impacts of a given service or use-case through its pressure on the dimensioning of a *hypothetical* infrastructure. It is instanced by a bottom-up model that automatically scales the infrastructure to different *scenarios* and hypotheses. From this infrastructure, we can then estimate its *absolute* power consumption and other environmental impacts. To estimate the absolute impacts of a given use-case, one starts to define a minimalist *baseline* scenario from which a baseline infrastructure is generated and evaluated. Then, a second infrastructure is generated and evaluated for the given use-case and the difference between the absolute impacts of these two scenarios is attributed to this use-case. In practice, this approach also permits to compare different hypotheses for the same use-case, hence enabling a better understanding of the consequences of different choices on the infrastructure. Through this exercise, it is important to consider the whole service/use-case at the scale of a large-enough *territory* to be representative of the service/use-case at hand. Those few methodological principles are key to avoid the pitfalls discussed in the introduction, but also to avoid tricky allocation issues of shared or multi-purpose equipment.

Below we present our methodology as four main steps.

Step 1 - Use-case

In order to guide the next steps, we first need to define the service or use-case that we aim to model, evaluate, and analyze. Examples encompass video streaming, videoconferencing, large file downloading, email communication, etc. In addition, one also has to identify the main parameters and variables associated to this use-case, and their range of values that will be explored (e.g., video resolutions, number

of viewers, server localization, file sizes, frequencies, etc.). At this step, one can already define the *baseline scenario* through the choice of the most sober values for those variables (e.g., no streaming, a few emails a day without attachments, etc.).

Step 2 - Boundary

This step covers two aspects. First, which parts of the Internet infrastructure are included: datacenters, core, edge, fixed and/or radio network, fiber and/or copper, customer premises equipment (CPE), etc. Second, which geographical territory: a city, a country, the world? The choice of a territory might be dictated by the purpose of the evaluation, e.g., one might be interested in evaluating a service for a given country. Otherwise, for the sake of simplicity, it might be wise to choose the smallest possible territory that is representative of the scenarios identified in step 1.

Step 3 - Design of the parametric infrastructure model

This step consists in designing the parametric model that will generate infrastructures according to some dimensioning variables. To this end one must start to design a minimalist infrastructure, e.g., every home and datacenters of the considered territory must be connected with the capacity to exchange some bits, the radio network must cover 99% of the territory, every user of our scenarios possesses at least one smartphone, tablet or laptop, etc. This minimalist instance completes the baseline scenario identified at step 1. The model has to be designed to be able to scale-up and to cover the range of use-cases and boundaries defined in the previous steps.

For the network, as a proof of concept, in this work we propose to use a simplified tree representation that goes from the user houses up to the main datacenters hosting the considered service, and passing through nodes representing the different pooling points of the fixed-access, edge, and core network layers. More complex structure representations shall be used in future work. Figure 1 illustrates part of the tree of the infrastructure. Putting aside the end-user devices, the leaves correspond to the home-routers, and the nodes represent converging/splitting points where congestion points might occur. Edges represent fiber links which might either include passive-only equipments, but also active equipments for long-distance hops. Following previous bottom-up network models, its main dimensioning variables are the bandwidth capacities required at the different nodes and links of the tree.




Fig. 1. Abstract tree representation of the infrastructure. In this example, tree goes from a main datacenter to many devices through nodes and links. A node can, for instance, represent an IXP that includes a CDN or routers. The detailed infrastructure for our use-case is presented in Figure 3.




Fig. 2. General OTT VoD infrastructure. Steps in black are the one included in our boundary. Figure modified from [20].

At each node and link l , the equipment is scaled up to the minimal quantity enabling:

- i) a connection to every subscribers (households)
- ii) a bandwidth capacity equal or greater to R_l (in bit/s) which has to be adjusted with respect to the peak access rate estimated at the node or link l for the given scenario (see Step-4). Throughout this paper, peak rates are assumed to refer to averaged traffic rates over a few seconds.

This quantity is then multiplied by the redundancy factor η . The global power consumption is estimated as a sum over all equipments and facilities of the infrastructure. To this end, each equipment must be associated with a power consumption profile, i.e., a static (or idle) power (in W) and a power intensity factor (in W/Gbps) for the dynamic power consumption part which is assumed to be proportional to the actual traffic. Because we scale the infrastructure to peak needs, we are able to understand the physical reality and bottlenecks behind power consumption.

Step 4 - Scenario evaluation and peak demand modeling

Finally, to evaluate one of our scenarios, we need to translate it to the dimensioning variables exposed by the model defined in the previous step. In our case, this mainly requires to estimate the capacities R_i from peak access rates estimated at every node and link of the tree. This step usually embeds a growing margin factor (α in the following) allowing to anticipate exceptional traffic peaks, and for future growth provisioning.

Our scenarios are constructed based on global average statistics of different kinds. The first kind are expressed in term of percentage of “active” users (e.g., percentage of active subscribers for the baseline, or the percentage of simultaneous VoD watchers). The second kind are numbers such as the average number of inhabitants per house. Whereas using such statistic means would be sufficient when considering a large pool of inhabitants, they cannot be used to reflect worst case scenario near the leaves of the tree where some equipments are shared by few dozens to a few hundreds of inhabitants only. For a better accuracy, we propose to consider their respective distributions, say d^n , for a given sub number of inhabitants n . To neglect the most unlikely occurrences only, we use the smallest quantity q such that the probability of having an occurrence x greater than q is extremely low, i.e., $d^n(x > q) < \epsilon$. We used $\epsilon = 10^{-9}$.

For the statistics of the first kind, defined by a percentage s of “active” inhabitants, the distribution d_s^n corresponds to a hypergeometric distribution. Assuming that n is small compared to the total number of inhabitants, it is well approximated by the simpler binomial distribution, and we define the function $q_s(n)$ as the smallest $q \leq n$ such that $d_s^n(x > q) < \epsilon$, which corresponds to the $1 - \epsilon$ percentile, which is itself computed through the inverse of the cumulative distribution function (CDF) of d_s^n . Note that for very large n , we have $q_s(n) \approx s \times n$, but $q_s(n)$ can be significantly greater than $s \times n$ otherwise (e.g., $q_{3\%}(64)/64 = 21\%$).

For our statistics of the second kind, such as the number of inhabitants per home, we first have to define the respective discrete distribution $d_i(x)$, and consider the sum of n random variables having this discrete distribution. Again, we assume that n is small compared to the total number of inhabitants. The resulting distribution $d_i^n(x)$ is thus obtained by the convolution of d_i with itself n times. As before, we then define $q_i(n)$ as the $1 - \epsilon$ percentile of $d_i^n(x)$. Since there is no closed-form formula for d_i^n , computing $q_i(n)$ for many different values of n can be very tedious. Instead, we found that $q_i(n)$ can be very well approximated by a function of the form:

$$q_i(n) \approx \max(a_i n + b_i n^{c_i}, n \bar{d}_i) \quad (1)$$

where \bar{d}_i is the mean of $d_i(x)$. The three coefficients a_i , b_i , c_i are found numerically to interpolate three points taken at $n \in \{16, 128, 1024\}$. Values of d_i for our use-case presented in next section can be found in Table I.

TABLE I
DISTRIBUTION OF INHABITANTS PER HOUSE IN FRANCE IN 2019 [21].

x	1	2	3	4	5	≥ 6
$d_i(x)$	0.369	0.326	0.135	0.113	0.041	0.016

IV. CASE STUDY: VoD STREAMING

To illustrate our methodology, we consider VoD streaming. This use-case has already been addressed by many studies [16], [20], [22]–[27]. All of them, however, focus on estimating the electricity intensity of one hour of video or yearly electricity usage of a video service, while allocating the network part based on volume of data (Wh/GB), or a mix of volume and duration. As motivated in the introduction, those *a posteriori* attributional allocation strategies can hardly be used to predict the impacts of intensifying VoD streaming, nor to understand the real pressure of VoD streaming on the network infrastructure. Applying our methodology to this use-case will thus reveal new insights.

This section follows the four steps presented in the previous section. For the sake of clarity, the steps 3 and 4 are detailed

per kind of equipment. Owing to the lack of power profile data, and since we consider a fixed network, we will focus on the estimation of the static power only, leaving the dynamic part for discussions in Section V-A.

A. Design overview

1) *Step 1 - Use-cases definition:* Video streaming in general covers a large range of different types of services, each having their own infrastructure design with different usage patterns. In this study, we limit ourselves to over-the-top (OTT) streaming from a unique service provider for the whole territory. We further assume a bounded catalog controlled by the service provider. We thus exclude other types of video streaming such as Youtube (unbounded catalog), live streaming, IPTV, and advertising videos.

Figure 2 depicts the main components of an OTT VoD service. After the content creation stage, the service provider manages its catalog with many different encoding settings and redundant storage in its main datacenters. In general, videos are not delivered directly from the main datacenters but from closer servers belonging to content delivery networks (CDNs). CDNs are storage servers usually located at Internet service provider (ISP) or Internet exchange points (IXP). CDNs reduce data traffic in the core network (and in particular in submarine and longhaul optical fibers) and improve user experience with faster loads. Those CDN servers are partly updated every days during low traffic periods. The videos are then delivered to the customers through the core and access networks.

The main variables are the video quality R_v (ranging from 0 to 30 Mbps), and the percentage s_v of inhabitants accessing the service simultaneously during the peak period. Other variables include: the size of the catalog and the percentage of video content accessed through the CDN.

In our baseline scenario we consider a traffic expected to represent a “minimalist and sober” use of the Internet as communication and information sharing means, excluding all traffic-intensive usages such as video streams, large file transfers, heavy web pages, etc. (both R_v and s_v are set to 0). We further ignore all the B2B traffic. At peak hours, this baseline traffic is modeled as a global percentage of active customers $s_b = 2\%$, and a per customer speed rate $R_b = 10$ Mbps. Those numbers yields an average peak rate of 200 kbps per customers, which matches average peak rates for copper lines in 2013 in France. Those numbers are thus rather high for a “minimalist and sober” use of Internet, but they can be considered as conservative. We further assume that the baseline and VoD traffic peaks are fully correlated, which is also a rather conservative choice.

2) *Step 2 - VoD streaming boundaries:* For the sake of simplicity, we chose a narrow boundary in this study. Parts in dark gray in Figure 2 are those that we include, while we ignore steps in light gray (content creation, encoding, customer management, end-user devices, ...). To be consistent with our tree structure, we consider a unique CDN that includes both servers and dedicated routers.

Territory: The territory is chosen to be representative of the geographical scale typically covered by a unique CDN. This makes Metropolitan France, which hosts the ICT4S 2023 conference, a rather good candidate to scale our network scenarios. We thus consider $65e6$ inhabitants for about $\#home = 30e6$ households, each having an internet connection and being a customer of the VoD service. We assume that $2/3$ of our baseline traffic stays in this territory. The CDN is naturally located in an IXP in Paris.

Main datacenters: The main datacenter servers and routers of the VoD service provider are expected to be shared by many countries. We thus chose to ignore their own power consumption. However, in order to account for the Internet traffic load required to update our CDN, as well as to deliver videos that are not cached in the CDN servers, we still consider one main datacenter located in North America at about 900km from the Atlantic submarine cable landing point. We also assume that the $1/3$ of the baseline traffic coming from international sources goes through this same route.

Network: We include both the core and edge network active equipment, but ignore passive ones. For the access network, since the VoD service we consider is mostly used at home, we consider only a fixed-access network that we assume to be fully implemented through the GPON (Gigabit Passive Optical Network) FTTH (fiber-to-the-home) architecture.

CPE: Since our study focuses on the network, on the customer side, we consider only the ONU needed for the GPON fiber architecture, but ignore all other devices such as home-routers, set-top-boxes, TVs, laptops, etc.

Summary of the infrastructure: Figure 3 summarizes the infrastructure for our use-case. Our infrastructure and the number of equipments, is scaled according to peak usage that depends on the scenario. In the next subsection, we detail each node and how we compute the quantity and power of all equipments. Some general parameters are considered. The *PUE* indicator for the network is set to 1.8 [9] for all experiments, while the redundancy factor in case of failure is $\eta = 2$. Our design is largely inspired by the model of Baliga et al. [10], some power consumption coming from this paper. When it is the case, we update energy intensity (W/Gbps) of equipments considering the energy efficiency gains through years. The formula proposed by the authors: $I = \frac{P}{C} = \frac{P_0}{C_0} \times (1 - \gamma)^t$ considers the t years between 2020 (our reference year) and 2008. P_0 is power and C_0 capacity from 2008 and $\gamma = 0.1$.

B. GPON based FTTH access network

Step 3 - Access network design: To model FTTH, we have chosen the GPON architecture as our outermost link (Figure 4). Its main component is the Optical Line Terminal (OLT) hosted in central offices (hubs) of the operator. It is primarily composed of GPON cards with up to $\#gpon/card = 16$ GPON ports per card, and up to 16 cards per OLT. Each GPON port has a maximal capacity of 2.5Gbps shared by $\#sub/gpon \leq 128$ subscribers through a unique output fiber. This fiber is then split in a tree structure through passive




Fig. 3. Network modeling of our video on demand use case. We restrict our boundary to fiber to the home network, a unique datacenter and one CDN.




Fig. 4. Representation of the fiber to the home structure that connects core network to access network with GE and GPON ports.

optical elements. This tree ends with one Optical Network Unit (ONU) in each household. This active equipment is required to deal with the temporal multiplexing of the underlying GPON protocol. In our model, the OLT is connected to the backbone through 10 Gigabit Ethernet ports (10GE ports). We have at least one 10GE port per OLT and thus at most $256 \times \#sub/gpon$ subscribers per 10GE port. The number $\#hub$ of central offices is fixed and determined to cover the whole territory. An opendata list of central offices in France in 2020¹ reports about 1500 hubs dedicated to OLT hosting, and 4300 other ones tagged as “copper”. In practice, many of them also host OLTs, and the presence of four main operators also tends to increase the number of hubs compared to a single operator scenario. All in all, we took $\#hub = 3000$ (in practice this parameter has a very small influence on the overall results). For the sake of simplicity, we also assume that the subscribers are equally spread among the hubs. Static power consumption and capacity of those equipments are given in Table II.

From this design, the total number of GPON ports is:

$$\#gpon = \left[\frac{\#home}{\#sub/gpon} + \#hub \times \#gpon/card \right]. \quad (2)$$

The right-hand side term accounts for the fact that GPON cards are, in practice, not completely filled. We therefore add the equivalent of one GPON card with $\#gpon/card = 16$ ports per hub. The number of 10GE ports is scaled with respect to the peak-rate demand R_{olt}^* at one OLT:

$$\#ge = \eta \#olt \left[\frac{R_{olt}^*}{C_{GE}} \right]. \quad (3)$$

The number of OLTs is obtained as:

$$\#olt = \max \left(\#hub, \left[\frac{\#gpon}{\#gpon/card \times \#card/olt} \right] \right), \quad (4)$$

with $\#card/olt$ the maximal number of GPON cards per OLT ($\#card/olt = 16$). The total power of the access network is the cumulative power of all GPON and GE cards:

$$P_{access} = PUE \times (\#gpon \times P_{gpon} + \#ge \times P_{ge}). \quad (5)$$

Step 4 - Access network peak demand: In order to evaluate the above equations, we have to compute the actual maximal number of subscribers per GPON tree ($\#sub/gpon$) such that the peak demand traffic of a given scenario does not exceed its bandwidth capacity, as well as the demand rate R_{olt}^* . To this end, we will exploit the statistical tools presented in

TABLE II
STATIC POWER P AND CAPACITY C OF SOME NETWORK EQUIPMENT.
NUMBERS ARE ADAPTED FROM [10] AND [28].

	Power (W)	Capacity
ONU	$P_{ONU} = 2.5$	-
GPON port	$P_{GPON} = 15$	$C_{GPON} = 2.5$ Gbps
10GE port	$P_{GE} = 30$	$C_{GE} = 10$ Gbps
Ethernet switch module	$P_{eth} = 60$	$C_{eth} = 40$ Gbps
BNG module	$P_{BNG} = 75$	$C_{BNG} = 40$ Gbps
Edge router module	$P_{erm} = 120$	$C_{erm} = 40$ Gbps
Core router module	$P_{crm} = 1400$	$C_{crm} = 560$ Gbps
Flash server	$P_{flash} = 320$	$C_{flash} = 190$ Gbps
Storage server	$P_{sto} = 400$	-

¹<https://www.data.gouv.fr/fr/datasets/localisations-des-noeuds-de-raccordement-abonnes-na-et-optiques-nr-dans-le-pays-breton/>

Section III-Step-4. Since the GPON technology cannot support more than 128 subscribers per GPON port, $\#sub/gpon$ is the largest integer $n \leq 128$ such that:

$$\alpha_t R(n) < C_{GPON}, \quad (6)$$

where $R(n)$ is the peak demand for a pool of n subscribers:

$$R(n) = R_v \times \frac{q_{s_v} \circ q_i(n)}{S} + R_b \times q_{s_b}(n). \quad (7)$$

We recall that α_t is a growing margin factor (Section III-Step-3), s_b is the average percentage of active subscribers with baseline usage, s_v is the average percentage of VoD viewers among the population at peak hours, and S is the average number of viewers sharing the same device/video flux ($S = 1.5$ [20]). We used $\alpha_t = 1.5$. In practice, we compute this solution using a binary search while accelerating the evaluations of $q_{s_v} \circ q_i(n)$ and $q_{s_b}(n)$ by approximating both of them by functions of the same form as equation (1) with coefficients computed using the same three points interpolation strategy. The logic for the 10GE uplink ports is slightly different. Indeed, once the number of subscribers per GPON tree is known, the number of subscribers per OLT is fixed, and we thus avoid the need for the binary search by directly setting the peak-rate demand:

$$R_{olt}^* = \alpha_t R(\lceil \#sub/\#olt \rceil). \quad (8)$$

C. National edge and core network

This subsection focuses on the part of the network tree connecting the OLTs to the main IXP located in Paris. The link from the IXP to North America will be addressed in the next subsection.

Step 3 - Edge/core network design: To connect the entire country to the main IXP, we use a tree topology that interconnects children to core router nodes with a star topology. We consider that each core router node has 8 child nodes, and 3 core levels L_0 to L_2 , with L_0 corresponding to the main IXP in Paris. A core router is made of a variable number of modules. Given a capacity requirement of R Gbps, the actual number of modules is given by:

$$\#CN(R) = \left\lceil \frac{R}{C_{crm}} \right\rceil,$$

with an electrical power of:

$$P_{CN}(R) = PUE \times \eta \times \#CN(R) \times P_{crm}.$$

Each of the 64 core nodes of level 2 are then linked to 8 edge nodes (level 3). An edge node is composed of a modular edge router, a modular broadband network gateway (BNG), and a modular Ethernet switch connected to the OLTs. The actual quantity of modules composing one edge node and its respective power is computed just as $\#CN(R)$ and P_{CN} but using the capacity and power features given in Table II:

$$P_{EN}(R) = PUE \times \eta \times \sum_{k \in \{eth, BNG, erm\}} \left\lceil \frac{R}{C_k} \right\rceil \times P_k.$$

Core and edge routers are connected through a wavelength division multiplexing (WDM) transport system, composed of two terminal multiplexers (one at each extremity, with a power of 4.6 W per channel), amplifiers every 100 km (with a power of 3.5 W per channel), and a capacity of 40Gbps per channel. For a capacity R and distance $dist$, the electrical power of such a link is thus:

$$P_{wdm}(R, dist) = \eta \left\lceil \frac{R}{40} \right\rceil \left(PUE \times 9.2 + \left\lceil \frac{dist}{100} - 1 \right\rceil \times 3.5 \right).$$

Let R_l^* be the required capacity at a node level l , and $dist_l$ be the average distance separating a pair of nodes at level $l-1$ and l . In our design, $dist_1 = dist_2 = 300$ km, and $dist_3 = 100$. The overall power P_{nat} of national core and edge nodes and links are finally:

$$P_{nat} = 2^3 \times P_{EN}(R_3^*) + \sum_{l \in \{0,2\}} 2^l (P_{CN}(R_l^*) + P_{wdm}(R_{l+1}^*, dist_{l+1})).$$

Step 4 - Edge/core network peak demand modeling: The capacity R_l^* required at each level l is estimated from the peak demand of our scenarios in a similar fashion than for the uplink ports of the access network. Observing that the number of subscribers related to one node is equal to $\#home/2^l$, we directly have: $R_l^* = \alpha_t R(\#home/2^l)$.

D. International longhaul link

Step 3 - International longhaul design: The main IXP is connected to the main datacenter in North America through a longhaul WDM transport system made of two terrestrial parts of 600 and 900 km respectively, and one WDM submarine section of 8000km. We further assume that about 7 core nodes are crossed over along this route. Let R_u^* be the required capacity for this line. The electrical power for the terrestrial part is thus:

$$P_{intT} = 7 \times P_{CN}(R_u^*) + \sum_{dist \in \{600, 900\}} \left(P_{wdm}(R_u^*, dist) \right).$$

Submarine WDM systems are slightly more complex to model. It is composed of two terminal multiplexers (35W per channel), with repeaters (0.2W per channel) placed to amplify the signal every 50 kilometers and a capacity of 40 Gbps per channel. The repeaters are powered by electrical suppliers located on the coast with 80% energy efficiency through cables having a resistance yielding power loss of about 0.004 W/km [10]. This sums up to 142 W per channel for our 8000 km cable, yielding for the final electrical intensity of the undersea connection:

$$P_{intU} = PUE \times \eta \times \alpha_u \times \left\lceil \frac{R_u^*}{40} \right\rceil \times 142. \quad (9)$$

Submarine cables are often designed with a larger margin factor than terrestrial links, hence the dedicated scale factor $\alpha_u = 2$ [10].

Step 4 - International longhaul peak demand modeling: The peak demand rate R_u^* between the main datacenter and our main IXP/CDN is the maximum between the peak rate R_{fill}^* to fill CDN servers and the peak traffic rate corresponding to the baseline and VoD scenarios. For the later, we assume that only 1/3 of the baseline traffic goes through this international link, and that %CDN percent of the VoD traffic is handled by the CDN (we used %CDN = 80%). We end up with:

$$R_u^* = \max(R_{fill}^*, (1 - \%CDN)R_v^*) + \frac{1}{3}R_b^*. \quad (10)$$

This equation assumes that the hours of CDN filling might overlap with the baseline peak, which is unlikely but conservative. The quantities R_b^* and R_v^* are the global peak rates for all customers for the baseline and VoD usages respectively:

$$R_b^* = s_b \times \#home \times R_b; R_v^* = s_v \times \bar{d}_i \times \#home \times R_v.$$

E. CDN

Step 3 - CDN design: Our CDN includes different kind of servers. Following Netflix's CDN design [28], we consider a CDN made of a few *storage* servers having a large storage capacity to hold a large percentage of the catalog content, and many *flash* servers having a limited storage capacity, but a very high throughput. It is completed with dedicated edge routers, yielding an overall electrical power modeled as:

$$P_{CDN} = PUE \times \left(\left[\frac{R_{CDN}^*}{C_{flash}} \right] P_{flash} + \#sto P_{sto} + \eta \left[\frac{R_{CDN}^*}{C_{erm}} \right] P_{erm} \right). \quad (11)$$

where R_{CDN}^* is the required throughput capacity. The average power and throughput of the flash and storage servers are given in Table II.

Step 4 - CDN peak demand modeling: Following Netflix documentation [28], we set the number of storage servers to $\#sto = 40$, and kept it fixed for all our VoD scenarios even though one could slightly adjust it with respect to the maximal video bitrate. The throughput capacity R_{CDN}^* is estimated from the global peak demand R_v^* , and the percentage %CDN of content effectively provided by the CDN:

$$R_{CDN}^* = \alpha_t \times \%CDN \times R_v^*.$$

Finally, we found that the bitrate R_{fill}^* needed to update the $\#sto$ CDN storage servers on a daily basis to be quite negligible. According to Netflix [28], $\sim 1.8\%$ of the content of CDN storage servers are updated every day. With 320 TB of capacity each, and assuming they are updated during a period of 8h, this yields $R_{fill} = 72$ Gbps, which is ~ 30 times lower than our estimated baseline peak for the international link.

V. RESULTS ON DIFFERENT SCENARIOS

In this section, we compare our baseline with other scenarios. We first analyze the influence of video quality and discuss proportionality (Section V-A), and then evaluate an hybrid OTT+DTT scenario (Section V-B). We implemented our model as a static web-application allowing the user to modify

all the parameters and hypotheses of our model, including the capacity and static power of each network element. This application will be made available online upon acceptance.

A. Effect of video quality

With our consequential methodology, we wish to understand what could be or have been the consequences of some restrictions over infrastructure dimension and energy consumption. In this section, we propose to compare the influence of several video qualities, namely HD (1280×720 at 3 Mbps), FHD (1920×1080 at 5 Mbps), UHD (3840×2160 at 16 Mbps), and a fourth UHD++ scenario with 4K resolution, high-dynamic-range (HDR) color depth, and 60 frame-per-second (at 27 Mbps). In all our scenarios, we assume a baseline peak rate $R_b = 10$ Mbps for a percentage $s_b = 2\%$ of subscriptions. All our VoD scenarios assume a peak percentage $s_v = 20\%$ of viewers through OTT. We recall that contrary to s_b , this percentage is relative to the whole population.

Moreover, in order to understand the effect of different usage patterns, we also sketched a fictitious download scenario (“DL”) of very large files such as OS updates and AAA video games. The latters become larger and larger with the top 26 ranging from 64GB to 200GB at the end of 2021 [29], while generating heavy loads on the network at every release or patch update of the most famous titles. Our VoD architecture can easily be adapted to such a use-case by setting: $R_v = 200$ Mbps, $S_v = 3\%$, %CDN = 95%, and replacing the VoD term of equation (6) by a simpler subscriber-based term: $R_v q_{s_v}(n)$. For this term, we also relaxed the confidence parameter to $\epsilon = 10^{-7}$ (Section III-step-4), hence accepting that a few users will very likely experience slightly degraded download rates.

Applying the model detailed in the previous section to our scenarios, and integrating static powers over one year, leads to annual power consumptions presented in Table III. This table also reports the percentage of energy consumption increase relative to the baseline scenario.

A first observation is that the overall energy consumption is not directly proportional to bitrates. This is especially true when including the ONUs that are plugged 24/7 in every home. The increase of the access network is negligible for HD/FHD bitrates, and it remains limited even for the FHD scenario. This is because the baseline 1:128 configuration of the GPON trees is enough to handle such bitrates. Only the 10GE uplinks has to be upscaled. The UHD++ scenario, however, yields a much higher pressure on the GPON trees that have to be upscaled to a 1:75 configuration. For all scenarios,

TABLE III
ANNUAL POWER CONSUMPTION IN GWh FOR EACH NETWORK PART.

Scenario	ONU	Access	National Core+Edge	Int. longhaul	CDN	Total
Baseline	667	69	9	2.8	0	748
HD	667	71	17	10	2.2	767 (+3%)
FHD	667	72	25	15	2.5	782 (+5%)
UHD	667	84	70	43	11	875 (+17%)
UHD++	667	136	114	70	18	1005 (+34%)
DL	667	271	140	2.8	32	1113 (+49%)




Fig. 5. Energy consumption difference between our VoD/DL scenarios and the baseline. The respective video bitrates R_v are shown as dots.

we observed that the submarine cable counts for about 33% of the whole international longhaul connection.

Figure 5 shows the energy consumption difference between our VoD/DL scenarios and the baseline. One can observe that for the VoD scenarios, for which only the streaming bitrate changes, this difference is much more proportional to the VoD bitrate than absolute consumption, though not perfect. However, when including a different use-case as our large-file download scenario, this apparent correlation breaks.

Owing to the lack of precise power consumption profiles for each of the considered equipments, the previous energy estimates cover static power consumption only, ignoring the dynamic part. In order to get a rough estimate of what could be the effect of accounting for the dynamic power consumption, we used the average dynamic intensity factor estimated by Malmodin [15] for a fixed line (i.e., ~ 0.1 Wh/B). Considering an average of 2GB (resp. 25GB) of data per subscriber per month for the baseline (resp. DL) use-case, and an average of 3.2h of OTT video per day per subscriber, we obtained the total yearly traffic volume and updated yearly energy consumptions reported in Table IV. As expected, for fixed lines the dynamic power consumption part is quite negligible. Comparing the UHD++ and DL scenarios, we see that both the absolute (1048 vs 1114 GWh) and incremental (300 vs 366 GWh) energy consumptions are even less correlated to volume (433 vs 10 EB) than bitrates (27 vs 200 Mbps). This table also reports the energy intensity indicators obtained as the ratio of those two numbers in Wh/GB. This clearly shows that increasing the video bitrate enables decreasing the relative data transmission consumption when expressed in Wh/GB, confirming that such an efficiency indicator does not reflect the power consumption increase of more intensive Internet usage.

B. Evaluation of DTT caching

The LoCaT report [24] estimates GHG emissions associated with serving TV content across different platforms, including digital terrestrial television (DTT), IPTV and OTT. Their functional unit was delivering one hour of video to a TV set. In one of their scenario, the authors studied the opportunity of viewing VoD content through DTT broadcasting. This requires one additional home-caching device to every home, which would enable receiving content via the DTT signal, cataloging, and storing content locally. When the user requests a given video, the device would first search for the content locally and would only stream it through OTT if absent. In a sense,

TABLE IV
EFFICIENCY ESTIMATION BASED ON YEARLY VOLUME AND ENERGY CONSUMPTION.

Scenario	Yearly volume (EB)	Energy (GWh)	Wh/GB
Baseline	0.73	748	1024
HD	49	772	15.8
FHD	81	790	9.7
UHD	257	901	3.5
UHD++	433	1048	2.4
DL	10	1114	111

this home-caching device plays the role of a personal CDN filled through DTT. The authors forecast a reduction of about 25% of OTT downloads, and devices exhibiting a stand-by power of 0.5W, and on power of 10W during 3.5h per day. This yields an average power of 1.9W.

We compare our previous FHD and UHD scenarios, with adding one home-caching device per home, and a very optimistic reduction of both the peak percentage of OTT users (s_v) and total hours of OTT videos by 50% for the FHD scenario, and 25% for the UHD one. This difference could be explained by the fact that FHD videos being significantly smaller, many more videos could be cached in the same device. Yearly energy consumption results are summarized in Table V. In both scenarios, the additional power consumption induced by the home-caching devices cannot counter-balance the minor absolute gains. This result contradicts the conclusions of the LoCaT study which is based on misleading Wh/GB energy intensity indicators. Indeed, as already observed and as confirmed again in this table, reducing by 25% or 50% the OTT video traffic (both in volume and peak bitrate), yields to rather small absolute reduction of the global energy consumption. Surprisingly, a 25% reduction on the bandwidth intensive UHD scenario leads to an energy reduction of 5.4%, whereas a higher usage reduction of 50% on the FHD scenario leads to a smaller relative energy reduction of 3%. Yet another example of how relative indicators can be counter-intuitive.

Finally, these different results suggest that an ideal strategy for such a small caching device would be to insert it directly within the OLTs with a direct forward connection to the GPON cards. This way the gains of this cache on the rest of the network elements (10GE uplinks, edge, core, CDN, etc.) would be preserved, but the additional power consumption of a single caching device would be shared by thousands of users. Using a conservative number of 8000 subscribers per OLT, and a constant power of 30W per caching device would yield a negligible additional energy consumption of about 0.7 GWh. Our model also reveals that filling this cache through DTT or over the network during low traffic periods will not make any practical difference regarding the energy consumption of the rest of the network. Note that this option is purely speculative

TABLE V
OTT VERSUS {DTT+ HOME CACHING DEVICES} SCENARIOS

Scenario	OTT		OTT + DTT caching	
	s_v	GWh	s_v	GWh (network+home-cache)
FHD	20%	790	10%	1271 (768+502)
UHD	20%	901	15%	1362 (860+502)

since it is unclear whether this is even technically possible, and what would be the actual power of such a caching card.

VI. DISCUSSIONS

In this section, we discuss some choices made, limits and future work for our model.

a) Conservative hypotheses: We emphasize that our observations on the non-proportionality (even for the relative plot of Figure 5) and on the sobriety vs efficiency conflict would be exacerbated through less conservative choices such as:

- Home-routers: we considered a very efficient ONU (2.5 W) but adding the associated home-router energy consumption would significantly flatten the absolute consumption variations.
- Peak correlation: decorrelating the baseline and VoD streaming peaks would allow the baseline infrastructure to better absorb part of the VoD traffic.
- Homogeneous tree: a more realistic non-homogeneous distribution of the population would result in some nodes of the baseline scenario to be overdimensioned, hence leading to better handling of the additional VoD traffic.
- Homogeneous technology: all our scenarios are based on the same building elements with a rather low granular capacity. Basing the baseline infrastructure on bleeding-edge technologies (e.g., WDM at 100 Gbps per channel, 10 Gbps GPON) would again yield an overdimensioned baseline network, whereas allowing the most traffic-intensive scenarios to use different technologies would flatten the variations. This last option should, however, be accompanied with an increase of embodied impacts because of anticipated renewal.

We also assumed that the ONUs are always on. Enabling ONU/home-routers to be switched off when unneeded would result in a significant reduction of the absolute energy consumption. It is thus interesting to question how much a given usage is preventing such equipments to be switched off. In this regard, VoD streaming has a rather low impact with an average of 3.2h a day in our scenarios. For the baseline usage, a reasonable assumption would be to assume that they are switched off a few hours over nights and when the households are empty. An obvious worst-case usage is, however, smart-home equipment that require a permanent connection.

b) Complexity of the real world: Just like previous bottom-up model, ours can only offer a simplified vision of the reality which is much more complex. For instance, network equipments exhibit a huge variability both in terms of capacity and efficiency. Equipments evolve with time with increased capacity and efficiency. With many actors deploying network equipment, peak bitrate demand is not the only driver for increase of the infrastructure, but economical competition and geopolitical strategies also play an important role leading to overdimensioning. As future work, it would thus be interesting to integrate all those aspects in such a model.

c) Restricted perimeters: In this work, we have not included user devices, datacenters, nor content creation and encoding. On the datacenter side, maximal video resolution

and quality is expected to have a significant effect on the computing (encoding) and storage resources. Those parameters are also expected to play an indirect but important role in accelerating the renewal of end-user equipments, for instance for larger 4K, HDR-enabled TVs, hence increasing the overall electricity consumption, but also manufacturing impacts.

d) Carbon footprint: Converting energy consumption to carbon footprint requires knowing emission factors. These factors are country-dependent but also time-dependent as the energy mix depends on the time or season. Moreover averaged emission factors, even if made temporally varying, are not necessarily correlated to the *consequential* effect of adding or removing a large body of electricity demand. Therefore, in this study, we have omitted this step on purpose.

e) Life cycle assessment: Most studies are limited to the use phase of equipments, omitting other phases of their life cycle (material production, manufacturing, transport, installation, maintenance, end-of-life). The main reasons are the high level of uncertainty of estimating the emissions of these other phases. Some works do attempt to include embodied impacts [16], [24] from an average ratio method that computes the scale of embodied emissions compared to the use phase. In this paper, we have focused on the use phase only, but we acknowledge the strong importance of including the other phases, as well as accounting for other environmental impacts (e.g., water footprint, human toxicity, abiotic resource depletion, ...). Properly allocating embodied emissions of shared equipment is as tricky as allocating static power consumption. In this regards, since our global approach bypasses the need for arbitrary allocation, we argue that the methodology proposed in this paper is well suited to be extended to estimate embodied impacts. Our model shall therefore be extended to enumerate all passive equipment that we have neglected so far (cables, shelters, buildings, racks, etc.), as well as installation and maintenance operations.

VII. CONCLUSION

In this paper we have presented a novel methodology to assess the relationship between a given usage and network power consumption. Looking at the global energy consumption rather than attempting to arbitrarily allocate power between the different usages allowed us to avoid the classical pitfalls. Our results confirmed that classical efficiency indicators do not reflect the power consumption increase of more intensive Internet usage, and might even lead to misleading conclusions. The bottom-up parametric network model we presented has the notable property of translating global average statistics to local smaller pools of inhabitants. This theoretical network model is, however, necessarily imperfect and this paper discussed many future work opportunities such as variation of the density of population, variability of equipments, broader boundary, adding a broadband radio access network, and modeling other use-cases. Another interesting future work would be to investigate how to extend our methodology to properly account for multiple use-cases whose peak demands are expected not to overlap.

REFERENCES

- [1] Cisco, "Cisco visual networking index: Forecast and trends, 2017–2022," Tech. Rep., 2020.
- [2] V. Coroamă, *Investigating the Inconsistencies among Energy and Energy Intensity Estimates of the Internet, Metrics and Harmonising Values*. Swiss Federal Office of Energy SFOE, 2021.
- [3] D. Schien, P. J. Shabajee, H. B. Akyol, L. Benson, and A. Katsenou, "Help, i shrunk my savings! assessing the carbon reduction potential for video streaming from short-term coding changes," in *2023 IEEE International Conference on Image Processing*, 2023.
- [4] D. Schien, P. J. Shabajee, and C. W. Preist, "Rethinking allocation in high-baseload systems: A demand-proportional network electricity intensity metric," in *IETF Internet Architecture Board workshop on Environmental Impact of Internet Applications and Systems*, 2022.
- [5] R. Guerin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 7, pp. 968–981, 1991.
- [6] V. C. Coroamă, D. Schien, C. Preist, and L. M. Hilty, "The energy intensity of the internet: Home and access networks," in *ICT Innovations for Sustainability*, L. M. Hilty and B. Aebischer, Eds. Cham: Springer International Publishing, 2015, pp. 137–155.
- [7] J. Koomey, H. Chong, W. Loh, B. Nordman, and M. Blazek, "Network electricity use associated with wireless personal digital assistants," *Journal of infrastructure systems*, vol. 10, no. 3, pp. 131–137, 2004.
- [8] M. Gupta and S. Singh, "Greening of the internet," in *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '03. New York, NY, USA: Association for Computing Machinery, 2003, p. 19–26.
- [9] J. Aslan, K. Mayers, J. G. Koomey, and C. France, "Electricity intensity of internet data transmission: Untangling the estimates," *Journal of Industrial Ecology*, vol. 22, no. 4, pp. 785–798, 2018.
- [10] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical ip networks," *Journal of Lightwave Technology*, no. 13, 2009.
- [11] D. Schien, V. C. Coroama, L. M. Hilty, and C. Preist, "The energy intensity of the internet: Edge and core networks," in *ICT Innovations for Sustainability*, L. M. Hilty and B. Aebischer, Eds. Cham: Springer International Publishing, 2015, pp. 157–170.
- [12] V. C. Coroamă, L. M. Hilty, E. Heiri, and F. M. Horn, "The direct energy demand of internet data flows," *Journal of Industrial Ecology*, vol. 17, no. 5, pp. 680–688, 2013.
- [13] K. Hinton, F. Jalali, and A. Matin, "Energy consumption modelling of optical networks," *Photonic Netw. Commun.*, vol. 30, no. 1, p. 4–16, aug 2015.
- [14] A. Wu and P. Ryan, "Intelligent efficiency for data centres and wide area networks," International Energy Agency (IEA), Tech. Rep., 2019.
- [15] J. Malmudin, "The power consumption of mobile and fixed network data services - the case of streaming video and downloading large files," *Electronics Goes Green 2020+*, 2020.
- [16] N. Ullrich, F. M. Piontek, C. Herrmann, A. Saraev, and T. Viere, "Estimating the resource intensity of the internet: A meta-model to account for cloud-based services in lca," *Procedia CIRP*, vol. 105, pp. 80–85, 2022, the 29th CIRP Conference on Life Cycle Engineering, April 4 – 6, 2022, Leuven, Belgium.
- [17] D. Schien, C. Preist, M. Yearworth, and P. Shabajee, "Impact of location on the energy footprint of digital media," in *2012 IEEE International Symposium on Sustainable Systems and Technology (ISSST)*, 2012, pp. 1–6.
- [18] M. Ficher, F. Berthoud, A.-L. Ligozat, P. Sigonneau, B. Tebbani, and M. Wisslé, "Rapport : évaluation de l’empreinte carbone de la transmission d’un gigaoctet de données sur le réseau renater," Tech. Rep., 2021. [Online]. Available: <https://ecoinfo.cnrs.fr/wp-content/uploads/2020/12/Rapport-revise-1Go-VF02-2021.pdf>
- [19] L. Golard, J. Louveaux, and D. Bol, "Evaluation and projection of 4g and 5g ran energy footprints: the case of belgium for 2020–2025," *Annals of Telecommunications*, 11 2022.
- [20] the Carbon Trust and DIMPACT, "Carbon impact of video streaming," 2021. [Online]. Available: <https://www.carbontrust.com/resources/carbon-impact-of-video-streaming>
- [21] "Insee," <https://www.insee.fr/fr/statistiques/2381486>.
- [22] M. Efoui-Hess, *Climat : l’insoutenable usage de la vidéo en ligne - Un cas pratique pour la sobriété numérique*. The Shift Project, 2019. [Online]. Available: <https://theshiftproject.org/article/climat-insoutenable-usage-video/>
- [23] G. Kamiya, "The carbon footprint of streaming video: fact-checking the headlines," International Energy Agency (IEA), Tech. Rep., 2020. [Online]. Available: <https://www.iea.org/commentaries/the-carbon-footprint-of-streaming-video-fact-checking-the-headlines>
- [24] W. Pickett, A. Fiocco, G. Roberts, B. Horn, C. Preist, D. Schien, and P. Shabajee, *Quantitative study of the GHG emissions of delivering TV content*. The Locat Project, 2021. [Online]. Available: <https://thelocatproject.org/>
- [25] L. U. Marks, S. Makonin, A. Rodriguez-Silva, and R. Przedpełski. (2021) Tackling the carbon footprint of streaming media. [Online]. Available: <https://www.sfu.ca/sca/projects---activities/streaming-carbon-footprint/>
- [26] R. Obringer, B. Rachunok, D. Maia-Silva, M. Arbabzadeh, R. Nateghi, and K. Madani, "The overlooked environmental footprint of increasing internet use," *Resources, Conservation and Recycling*, vol. 167, p. 105389, 2021.
- [27] C. Preist, D. Schien, and P. Shabajee, "Evaluating sustainable interaction design of digital services: The case of youtube," in *CHI conference on human factors in computing systems*, 2019, pp. 1–12.
- [28] "Open connect website," <https://openconnect.netflix.com/en/>.
- [29] J. Ridley, "The biggest games by install size, real mighty storage hogs," PC-Gamer, <https://www.pcgamer.com/biggest-game-install-sizes/>, 2021.