



HAL
open science

What is the carbon footprint of one hour of video streaming?

Joanna Moulhierac, Guillaume Urvoy-Keller, Marco Dinuzzi, Zhejiayu Ma

► **To cite this version:**

Joanna Moulhierac, Guillaume Urvoy-Keller, Marco Dinuzzi, Zhejiayu Ma. What is the carbon footprint of one hour of video streaming?. Université Côte d'Azur. 2023. hal-04069500v2

HAL Id: hal-04069500

<https://hal.science/hal-04069500v2>

Submitted on 2 Nov 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.

Estimating the Carbon Footprint of Video Streaming

Joanna Moulierac*, Guillaume Urvoy-Keller*, Marco Dinuzzi*, Zhejiayu Ma†

* Université Côte d’Azur, France

† Easybroadcast, France

November 2, 2023

Abstract

Video streaming dominates the Internet traffic. Assessing the carbon footprint of video streaming has received recently a significant attention with a number of models proposed to associate a CO₂ cost to one hour of streaming.

In this work, we compare the modeling assumptions and computation methods used by five recent works to inform the debate. Indeed, initial results can be at odds, with up to one order of magnitude difference in the estimates. Our contributions are: (i) we relate the difference in the results primarily to the perimeter of the study, e.g. including production cost or not, (ii) we question some of the modeling assumptions made using a real deployment of a streaming server in a controlled environment with up to 2000 clients and (iii) we propose a technique to reconcile the models and obtain a CO₂ estimate in between 60 and 140 grams when considering the average worldwide carbon intensity of electricity.

1 Introduction

Video streaming dominates the Internet traffic [3, 15] and calls for a clear assessment of its environmental impact. To this end, a number of models have been developed recently to assess the CO₂ emissions of one hour of video streaming [18, 10, 13, 8].

Our objective in this work is to investigate the differences observed in the results provided by these models. The topic is controversial as exemplified by the SHIFT Project vs IEA dispute [8]: the SHIFT project publicly advertised an estimate of 0.769 kWh for one hour of video streaming while IEA experts obtained 0.078 kWh, an order of magnitude smaller value.

We compare and try to explain the differences of five different models [18, 10, 5, 13, 8]. Our contributions are as follows: (1) We relate the difference in the results to the perimeter of the study, e.g. including grey energy (production cost) or not, and also to the choice of inner

parameters of the models, e.g. the energy efficiency per GB of the delivery network; (2) We analyse and compare the models under study in a number of scenarios; (3) We question some of the modeling assumptions made using a real deployment of a streaming server in a controlled environment with up to 2000 clients; (4) We propose a technique to reconcile the models and obtain a CO₂ estimate in between 60 and 140 grams when considering the average worldwide electricity efficiency.

2 Models Analysis

Borrowing the life cycle assessment (LCA) vocabulary, the functional unit we consider corresponds to streaming one hour of video from a server in the cloud to an end user equipment.

2.1 High level overview

Let us first provide a high level overview of the common modeling assumptions shared by all these models. The classical approach to tackle the problem is to break the streaming distribution scheme into three tiers: (1) Data center, (2) Network, (3) User Equipment. As network and data centers are shared facilities, all models use as unit the amount traffic serviced by network devices and servers. This means that we need estimates for both the total energy consumption and the total traffic conveyed or processed by the device to obtain a per byte estimate. It is clearly a daunting task as data centers operators and ISPs may not publish this level of detail [18]. Note also that a direct consequence for the computation methodology (divide costs by traffic volume) for the networking contribution is that the value will be independent from the path length. On the contrary, the end user equipment being a non shared facility, one traditionally counts its usage with respect to usage duration, hence one hour in our case.

Figure 1 provides a high level overview of the different components. One clearly see two different sets of impacts that can be assessed: the impact due to the usage and the

impact due to the production (manufacturing) of the various devices (often called grey energy). As we will see, some models consider both usage and production, while some consider usage only. When considering production, either the production of all or only of some of the devices is considered.

Some models consider a mobile or fixed network access while others focus on one technology only. Also, depending on the models, either a single type of user equipment is considered or different types (laptop, TV, smartphone).

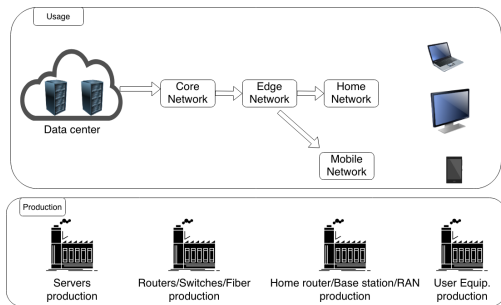


Figure 1: Models components and Emissions Causes

Usage and manufacturing costs need in general to be aligned. Indeed, the usage cost depends on the number of years of utilisation while the production cost is a fixed cost, “payed” before the usage phase. The latter is obtained from the manufacturers environmental sheets [9] which provide CO₂ footprints. As for usage, it can be either in kWh or CO₂ depending on the study. Translating kWh into CO₂ is country dependent, since each country use a different set of technologies (solar, wind, coal, etc) to produce electricity [11].

When focusing on the global footprint of one hour of streaming, we use CO₂ emissions, while we use kWh when considering usage phase. In this paper, we use either the french electricity factor¹ of 108 g of CO₂/kWh similarly to [5] or a World (2019) electricity factor of 340g of CO₂ per kWh, extracted from the IEA web site².

We consider five models: LIMITS [10], CARBONTRUST [18], RENATER[5], IEA [8] and SHIFT [13]. We named them based on the name of the organization (Shift, Carbontrust and IEA), the conference where it was published (LIMITS) or the network where the measurements were made (RENATER). A few remarks need to be done before delving into the details of the models:

- The LIMITS model has been published in a peer-reviewed paper. The RENATER model has been extracted from a research report (in French) ; it is

¹The energy factor indicates how much CO₂ grams are emitted to produce one kWh of electricity.

²<https://www.iea.org/reports/global-energy-review-2019/co2-emissions>

however based, for the network model on a peer-reviewed paper [6].

- The SHIFT and CARBONTRUST models are described in white papers. The SHIFT model is based on the so-called 1-byte model available as an excel spreadsheet [14].
- The RENATER model is not a video streaming model per se but a video conferencing model. The model is complex as it accounts for service details, such as the number of virtual machines deployed or the maintenance cost. We propose here a simplified version in line with the three tier approach and considering a physical (non shared) server. A key point of interest for this paper is that the model is calibrated with actual measurements made in the Renater³ network and data center since the application, the network, the data center and the clients are all operated by Renater.

We do not have access to the detailed equations of the IEA models but an online simulator is available on a Web page of the IEA web site [8]. The procedure to extract the data used in our result section will be detailed after the presentation of the test scenarios, at the end of section 3.1

2.2 Analysis

We follow the three-tier approach (data center/network/user equipment) to present the details of the different models. Table 1 presents the uniformed notations used for all the models (except IEA).

2.2.1 Data centers Contribution

LIMITS The energy consumption is computed at the scale of the data center [10]. The authors consider three typical sizes of data centers in terms of number of servers and power consumption and divide this consumption by the number of active streams. The corresponding values vary between 50 kW and more than 2500 kW, based on [17] and 10k to 50k streams. Note that it is implicitly assumed that the data center is dedicated to the streaming service. The resulting data center energy per stream is computed as follows:

$$DC_{energy} = \frac{D \cdot t}{\# Streams}$$

where:

1. D is the power demand in kilowatts (kW) based on the data center size,

³Renater is the French academic ISP.

Global Notation	
D	power demand (in kW) based on the data center size
t	Duration of the video being streamed (in hours)
R	Bit rate of the stream as a function of quality level
Q	Traffic during the streaming session: $Q = R \cdot t$
Data Center contribution	
# Streams	Number of simultaneous streams served by the data center.
ID_C	Energy intensity of the data center (in W). Estimated as 1.3W.
YC_{DC}	Annual energy consumption in kWh of the data center.
Y_{Traf}	Traffic serviced by the datacenter in GB in 1 year
PUE	Power Usage Effectiveness. Ratio of the electricity used to power the IT devices (servers and network) to the total energy consumed by the data center that includes notably air conditioning.
EF	Energy factor, the number of grams of CO ₂ emitted to produce one kWh
EI_{DC}	Energy used for the data center in order to process 1 byte of data (in kWh).
$Prod_{DC}$	Production cost in kg of CO ₂ to be accounted during the streaming session.
Lt_{EQUIP}	Life Time of equipment
$TotD_{VOD}$	Cumulative video streaming over one year
Network contribution	
I_{ISP}	Energy intensity of the access network and customer premises (in kWh/GB)
I_{FN}	Energy intensity of fixed network transmission (in kWh/GB)
I_{MN}	Energy intensity of mobile network transmission (in kWh/GB)
I_{HR}	Energy intensity the home router (in kWh/GB)
EI_{NET}	Energy used for the network in order to process 1 byte of data (in kWh)
$Prod_{NET}$	Production cost (kg of CO ₂) of server
I	Energy intensity: sum of edge, metro and long haul energy intensities
User contribution	
W_S	average power consumption of screens in kW
W_P	average power consumption of peripherals in kW
EI_{UD}	power consumption of the device for 1 min of video streaming (in kWh)
Lt_{ROUTER}	Life Time of Router
$Prod_{DEVICE}$	Production cost of user device
$Prod_{ROUTER}$	Production cost of router

Table 1: Notation used throughout the paper.

- t is the duration of the video being streamed in hours (1 hour in our case).
- # Streams is the number of simultaneous streams being served by the data center. It is implicitly assumed that the number of simultaneous streams does not vary during the considered time period (1 hour).

panies running some large data center, we can consider the large data center case. In the latter case, the power consumption varies from $\frac{2500kW}{50k} = 50W/stream$ to $\frac{2500kW}{10k} = 250W/stream$. Note that it is assumed here that the data center consumption does not vary with the load.

As video streaming is dominated by a few large com-

Datacenter		
	Usage	Manufacturing
LIMITS	$\frac{D \cdot t}{\# \text{ Streams}}$	
CARBONTRUST	$ID_C \cdot t$	
RENATER	$YC_{DC} \cdot \frac{Q}{Y_{Traf}} \cdot EF \cdot PUE$	$Prod_{DC} \cdot \frac{t}{Lt_{EQUIP}}$
SHIFT	$EI_{DC} \cdot Q$	

Table 2: DC contribution. Time corresponds to 1 hour and Total traffic to the amount of bytes streamed during this one hour session, which depends on the quality level.

CARBONTRUST In the CARBONTRUST model the energy consumption at the data center is estimated as follows:

$$DC_{energy} = ID_C \cdot t$$

where ID_C is the energy intensity of the data centre, estimated as 1.3W.

This value is significantly smaller than in the LIMITS case. It could correspond, in the LIMITS approach, to a small DC consuming 50 kW operating at the maximum load of 50k streams in parallel (1W/stream).

RENATER The RENATER model considers both usage and manufacturing cost. The usage cost is computed as follows:

$$DC_{energy} = YC_{DC} \cdot \frac{Q}{Y_{Traf}} \cdot PUE \cdot EF$$

where:

1. YC_{DC} is the annual energy consumption in kWh of the data center.
2. Q indicates the traffic generated by the user ($Q = R \cdot t$) and Y_{Traf} is the traffic serviced by the datacenter in 1 year.
3. PUE stands for Power Usage Effectiveness. It is the ratio of the electricity used to power the IT devices (servers and network) to the total energy consumed by the data center that includes notably air conditioning. RENATER considers it to be 1.5, i.e. a 50% overhead for non IT operations in the data center.
4. EF is the the energy factor, i.e. the number of grams of CO_2 emitted to produce one kWh, since the RENATER paper produces CO_2 emissions only. We will set aside this factor when comparing the usage contribution of the different models as they use different EF, e.g. France for RENATER and USA for LIMITS.

As compared to the other models, the RENATER model also considers the manufacturing cost of the data center computed as :

$$DC_{man} = Prod_{DC} \cdot \frac{t}{Lt_{EQUIP}}$$

where:

1. $Prod_{DC}$ is the total manufacturing cost (provided by the manufacturer) in kg of CO_2 .
2. Lt_{EQUIP} is the life time of the equipment.

SHIFT The Shift project estimates the datacenter energy consumption as follows:

$$DC_{energy} = EI_{DC} \cdot Q$$

where:

1. EI_{DC} is the energy used in order to process 1 byte. It is estimated to be $7.2 \cdot e^{-11}$ kWh/byte [14].
2. Q is the video size in bytes.

Table 2 summarizes the equations of the different models for what relates to the data center contribution. We separated usage and production costs. Focusing only on the usage part, one can observe a variety of approaches. The CARBONTRUST and SHIFT models seem to make implicit assumptions on the load of the data center, while the LIMITS and RENATER models try to make them explicit. The RENATER model has a clear advantage that the authors have access to the real statistics of the data center. However, this can also be interpreted as very specific figures. In contrast, the LIMITS paper is considering typical large scale data centers.

2.2.2 Network Contribution

Let us now focus on the network energy part for every model.

Network		
	Usage	Manufacturing
LIMITS	$I \cdot Q + I_{ISP} \cdot t$	$\text{Prod}_{\text{NET}} \cdot \frac{t}{\text{Lt}_{\text{EQUIP}}}$
CARBONTRUST	$(I_{\text{FN}} + I_{\text{HR}}) \cdot Q$ or $(I_{\text{FN}} + I_{\text{MN}}) \cdot Q$	
RENATER	$(\text{Router} + \text{NOC} + \text{Fiber}) \cdot Q \cdot EF$	
SHIFT	$EI_{\text{NET}} \cdot Q$	

Table 3: Network Contribution

LIMITS In the Limits paper, the total network cost is expressed as:

$$\text{Net}_{\text{energy}} = \text{CoreNet}_{\text{energy}} + \text{ISP}_{\text{energy}}$$

The $\text{CoreNet}_{\text{energy}}$ accounts for the energy intensity of data traffic through the edge, metro and long-haul networks, which includes undersea cable links. For a given stream, it is computed as:

$$\text{CoreNet}_{\text{energy}} = I \cdot Q$$

where:

1. $Q = R \cdot t$ and R is the number of GB/s per second that can be transmitted (depends on the stream quality level)
2. I (0.0523 kWh/GB) is the energy intensity and it is considered as the sum of 0.0043 for edge, 0.0200 for metro, 0.0280 for long haul, based on [16].

The $\text{ISP}_{\text{energy}}$ is computed as:

$$\text{ISP}_{\text{energy}} = I_{\text{ISP}} \cdot t$$

where: $I_{\text{ISP}} = 52 \cdot e^{-3}$ kW is the energy intensity of the access network and customer premise equipment (excluding user device).

CARBONTRUST This model accounts for the energy consumption over both fixed and mobile networks. The energy consumption from fixed network transmission is computed as:

$$\text{FixedNet}_{\text{energy}} = I_{\text{FN}} \cdot Q$$

where:

1. I_{FN} is the energy intensity of fixed network transmission (0.0065 kWh/GB) [18]
2. R is the data transmission rate (depends on the quality of streaming perceived by user).

The energy consumption from mobile network transmission is represented similarly as:

$$\text{MobileNet}_{\text{energy}} = I_{\text{MN}} \cdot Q$$

where I_{MN} is the energy intensity of mobile network transmission (0.1 kWh/GB) [18]

The last network element is the subscriber premise equipment, i.e. the home router, and only applies to video streaming over fixed network. The energy consumption of the home router is expressed as:

$$\text{Router}_{\text{energy}} = I_{\text{HR}} \cdot Q$$

where I_{HR} is the energy intensity the home router (0.025kWh/GB) [18]. Therefore if the user uses a fixed network we will have:

$$\text{Net}_{\text{energy}} = (I_{\text{FN}} + I_{\text{HR}}) \cdot D \cdot R$$

Instead if the user downloads through a mobile network we will have:

$$\text{Net}_{\text{energy}} = (I_{\text{FN}} + I_{\text{MN}}) \cdot D \cdot R$$

RENATER This model, as in the data center case, accounts for both usage and production. This leads to the following equation for the network⁴:

$$\text{Net}_{\text{energy}} = (\text{Router} + \text{NOC} + \text{Fiber}) \cdot Q \cdot EF +$$

$$\text{Prod}_{\text{NET}} \cdot \frac{t}{\text{Lt}_{\text{EQUIP}}}$$

where:

1. Router, NOC (Network Operating Center) and Fiber costs are utilisations in kWh/GB
2. Prod_{NET} is in kg of CO_2/GB .

⁴The Renater equations unit is CO_2 . One needs to divide by EF to get the results in kWh

SHIFT The approach is similar to the CARBON-TRUST model:

$$Network_{energy} = EI_{NET} \cdot Q$$

where EI_{NET} is the energy used for the network in order to process 1 byte of data. It is estimated as $4.29 \cdot e^{-10}$ kWh/byte for a fixed access and $8.84 \cdot e^{-10}$ kWh/byte for mobile.

Table 3 summarizes the equations of the different models for what relates to the network contribution. A few key points are noteworthy. The RENATER model is the only one that takes into account manufacturing, as well as the fiber and NOC operation costs. However, the RENATER and LIMITS models consider a fixed network access only and will thus underestimate the network in a mobile access scenario.

2.2.3 User Equipment

The last source of emission to be analysed is the user equipment. Depending on the study, the perimeter might encompass only the user device (LIMITS and CARBON-TRUST) or also the network equipment at home (RENATER).

LIMITS The user device energy during the usage phase is modeled as:

$$Device_{energy} = D_{USER} \cdot t$$

where D_{USER} is the power demand in kilowatts (kW) of a given user device. They use respectively 3.7W for a smartphone, 24.4W for a laptop⁵ and 43 W for a smartTV.

The LIMITS models also accounts for the manufacturing phase. Assuming that the manufacturing phase accounts typically for 90% of the total CO₂ emissions during the life time of the equipment:

$$Prod_{DEVICE} = 9 \cdot Device_{energy}$$

This assumption might seem a bit too high, as typical values lay in the range [60%, 80%] [9]. It has however the clear merit to work around the problem of having a different EF (emission factor) at the production phase (that depends on the possibly many countries where the plants producing the parts are located) and at usage phase (in general, a single country).

⁵Table 3 in [10] states 7.4W which must be a typo as it is far too small. The reference for the table is [25], which leads to 24.4W, a clearly better fit.

CARBONTRUST In this study, the authors account both for the user devices' screens (e.g. TVs, laptops, smartphones) and peripherals (e.g. set-top boxes and gaming consoles).

The energy consumption of user devices is estimated as:

$$Device_{energy} = (W_S + W_P) \cdot t$$

where:

1. W_S is the average power consumption of screens.
2. W_P is the average power consumption of peripherals.

$W_S = 1W$ for a Smartphone, 22W for a Laptop and 100W for a TV. As for W_P , it is equal to 89W for a gaming console and 18 W for a set top box. In section 3, we will take $W_P = 0$, i.e. not consider usage of a gaming console nor of a set top box.

RENATER The RENATER model considers as user device both the home router and the personal equipment of users. They account each time for the manufacturing and usage energies. The equations for the personal equipment are:

$$Device_{energy} = \frac{YC_{USER}}{TotD_{VOD}} \cdot Q \cdot EF$$

$$Prod_{DEVICE} = \frac{ManufImpact \cdot t}{Lt_{EQUIP}}$$

where the manufacturing impact is the total CO₂ impact reported by the manufacturer. As for the router part, it is:

$$Router_{energy} = \frac{YC_{ROUTER}}{Y_{Traf}} \cdot Q \cdot EF$$

$$Prod_{ROUTER} = \frac{ManufImpact \cdot t}{Lt_{ROUTER}}$$

SHIFT The equation used to compute the energy consumption of a user device is the following in the SHIFT model:

$$Device_{energy} = EI_{UD} \cdot t$$

where EI_{UD} is the amount of kWh consumed by the device for 1 min of video streaming: $1.1 \cdot e^{-4}$ kWh/min for smartphones and $3.2 \cdot 10^{-4}$ kWh/min for laptops. These values correspond to power demands of 6.6W and 19.2W respectively.

We summarize the equations for the user side contribution in table 4. The LIMITS and RENATER models likely lead too much higher values for the device energy consumption since they consider also the production cost.

After this detailed review of the models, we exemplify and compare them in the next section for different scenarios in terms of network set-up, end user device and the quality level of the multimedia flow.

User Device		
	Usage	Manufacturing
LIMITS	$D \cdot t$	$9 \cdot D \cdot t$
CARBONTRUST	$(W_S + W_P) \cdot t$	
RENATER	$DeviceEnergy + RouterEnergy$	$Prod_{DEVICE} + Prod_{ROUTER}$
SHIFT	$EI_{UD} \cdot t$	

Table 4: User Equipment contribution

3 Results

3.1 Scenarios

We use the following parameters: (i) User device in (TV, Laptop, Smartphone), (ii) Video Resolution in (SD, HD, 4k) and (iii) Network Type in (Fixed or Mobile) to build scenarios, e.g. (TV, 4k, Fixed).

Similarly to [18], the bit rates at each resolution are set to: SD = 2.22 Mbps, HD = 6.67 Mbps and 4k = 15.56 Mbps. These values are the (maximum) ones advertised by Netflix at high quality on a "web browser" [12].

For the LIMITS model, we considered a large data center servicing 10k flows.

3.2 IEA Model

Now that we introduced the scenarios, we can explain how we proceeded to extract the data from the IEA web site for the IEA model. We ran all the possible scenarios with the online simulator [8]. The values, presented in Table 5, are in grams of CO₂. We converted them back in kWh by considering the World (2019) case and an energy factor of 340g of CO₂ per kWh, extracted from the IEA web site⁶. We can observe from Table 5 that the user device contribution is time based and not volume based, as all the other models. In contrast, the network and data center components are volume dependent, as all the models, except the LIMITS one that accounts for the global load of the data center.

3.3 Overall comparison

Let us first compute the overall energy consumption for all models for two typical scenarios: (TV, 4k, Fixed) and (Laptop, SD, Mobile). We assume for both scenarios a large and lightly loaded data center for the LIMITS model. In Figures 2 and 3, we report the breakdown for the three tiers (data center, network and user device), as well as the total values. We included the manufacturing costs when available. We observe that:

⁶<https://www.iea.org/reports/global-energy-review-2019/co2-emissions>

User Device			
	Smartphone	Laptop	TV
SD	0.56	10.19	55.56
HD	0.56	10.19	55.56
4k	0	10.19	55.56
Network			
	Smartphone	Laptop	TV
SD	1.66	8.37	8.37
HD	5.28	8.49	8.49
4k	0	8.69	8.69
Data Center			
	Smartphone	Laptop	TV
SD	0.64	0.64	0.64
HD	2.75	2.75	2.75
4k	0	6.42	6.42

Table 5: IEA model: CO₂ emissions in grams

1. IEA always returns the smallest values, among all models, while the SHIFT model features the largest ones.
2. When comparing the TV and Smartphone case, we see clearly the impact of the size of the screen on the user device cost: it is negligible for the smartphone case, but becomes significant for the TV case. Also, a higher screen requiring a higher resolution, this results in a higher load in the network.
3. The SHIFT model features the highest network cost among all models, which explains most of its cost.
4. Both the LIMITS and SHIFT models, on the other hand, have a high data center cost. In the LIMITS case, it is exacerbated by the light load we chose.

In the following sections, we will focus on the contribution of each of the three tiers of the model. The grey energy will be treated separately.

3.4 User Device

We present in Figure 4 the energy consumption (production is excluded) of the user equipment for the different




Figure 2: Total Emissions (TV, 4k, Fixed)

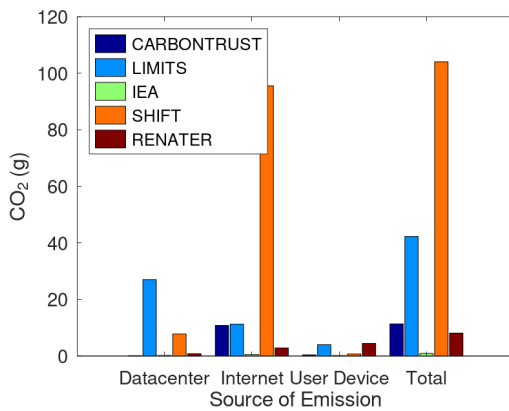


Figure 3: Total Emissions (Smartphone, SD, Mobile)




Figure 4: Comparison of the user devices (in kWh)

models. We chose a scenario with HD resolution and Fixed Network. The results are qualitatively similar for the other cases. We express the results in kWh and not grams of CO₂, since we focus on usage.

Figure 4 shows a clear agreement among the models for the smartphone and laptop cases. For the TV case, there is more variability (the SHIFT model does not implement the TV scenario, hence its zero value). These results can be explained as follows. The utilisation contribution of the devices is computed with a similar equation of the form $Constant_{Device} \cdot Time$ for all types of devices and all models. The difference between the models stems from the choice of the $Constant_{Device}$ value. For the smartphone case, it ranges between $1e-03$ kWh for the IEA model to $7.4e-03$ kWh for the RENATER case. As the constants are relatively similar, the final user device contributions are similar also. For the TV case, the CARBONTRUST model assumes a power demand of 100W while the LIMITS one picked 43W. The IEA web page assumes a large TV of 55 inches, which might explain its large value.

3.5 Network

Figure 2 and Figure 3 highlighted that the models return significantly different estimates for the network part. The IEA model reports the smallest values among all models while the SHIFT model returns significantly higher estimates. The reported values depend on the bit rate of the stream, hence the differences between SD and 4k cases. The RENATER model factors a manufacturing cost in its network footprint estimation, which accounts for about 40% of the network cost.

3.6 Data center

A few strategies have been considered to model the data center cost:

- The LIMITS model is a top-down model, as it considers the total energy consumption at the data center scale.
- The SHIFT and CARBONTRUST models are bottom-up models that focus on the streaming server consumption only. However, the SHIFT model considers that the consumption is a function of the amount of bytes while the CARBONTRUST consider time duration. The latter means that the energy consumption is independent of the actual load.
- The RENATER model relies on the quantity of data rather than the usage duration. The explicit use of the PUE further aims at accounting for the global data center contribution. In addition, the RENATER model accounts for the server production.
- We eventually note that the IEA model features the smallest estimate among among all models.

We will further discuss these different modeling assumptions in Section 4.2 using measurements with a real streaming server.

3.7 Grey energy

Two of the considered models take into account the production cost. First, the LIMITS model that accounts for it at the user device level. Second, the RENATER model that assigns a production cost to all tiers. The approach for the LIMITS model is straightforward: one takes the estimated utilisation cost in CO₂ and multiply it by a factor of 9. This has the clear advantage of aligning easily the production and usage costs. One could argue that this figure might be an overestimation as the observation that the production cost make 90% is valid for the full lifetime of the device while, here, we focus on a usage period where the power demand might be high.

To illustrate this point, we performed the following experiment with a iPhone 13 max. When watching a Netflix series in Wifi, it consumes 1.2-1.3W. This ramps up to 2-2.5W with a 5G access. Strangely enough, the base consumption might be higher, e.g. the locked screen consumes around 1.7W. Hence we observe that the streaming on a smartphone is not necessarily resource intensive. The LIMITS approach thus appears justified.

Let us now focus on the RENATER model. We report in Figure 5 both the utilisation and production cost of each tier. We consider the same (TV, 4k, Fixed) scenario as in section 3.3.

We can observe that the manufacturing costs for the data center and networking are high as compared to their usage cost. This is astonishing, as one could have expected that the long lifetime of the equipment would wipe out the production cost [4]. As for the user device, we are far from the 1:9 ratio advocated by the LIMITS model. There are two reasons for this observation. First, the lifetime of the considered TV was 10 years, the upper bound of the values reported in [10]. Second, we assume that the TV streams for 3h/day. The ratio 1:9 between production and usage holds at the scale of the day but if we consider a specific streaming hour, the usage cost is 24/3 higher than the average usage cost. Hence, production and usage cost are now on par.

4 Reconciling the models

In this section, we present our preliminary efforts to reconcile the models, by selecting the best approach at each tier. This is relatively easy for the user device and network tiers. It is however more complex for the data center. This is why we performed a measurement campaign over the




Figure 5: RENATER: Usage versus Manufacturing

Grid’5000 testbed with a streaming server servicing up to 2000 clients.

4.1 Towards a unified model

For the user device part, we picked the same power demand for all models (except IAE) for a specific device, e.g. 70W for the TV. For the production part of the LIMITS model, we used a 1:1 ratio and not a 1:9 ratio, see section 3.7.

Concerning the networking part, the CARBONTRUST and LIMITS model feature similar results while the IEA model is way smaller and the SHIFT model way higher. As advocated by the CARBONTRUST authors, a key reference in the domain is [2]. This paper focuses on the usage energy of the Internet and reconciles the previous observed estimates by proposing a simple empirical law that states that the energy consumption of the Internet “has decreased by half approximately every 2 years since 2000”. Following their model, one expects to reach 0.01 kWh/GB in 2020. This is in line with the measurements that fuel the RENATER model. Indeed, table 11 (p35) of [5] reports an network emission factor for the FAN equal to $1.7e^{-3}$ CO₂/GB, which corresponds to 0.016 kWh/GB. We used this value also for the SHIFT model.

Note that the above results are valid for the fixed part of the network and not for the mobile access network. The last tier is the data center. This is where reconciling the models is the most difficult. In a first attempt to address this issue, we propose a measurement campaign performed in Grid’5000.

4.2 Measurement campaign

The Grid’5000 tested allows to book dedicated physical servers and track their energy consumption [7]. We set up an experiment with a single streaming server with up to

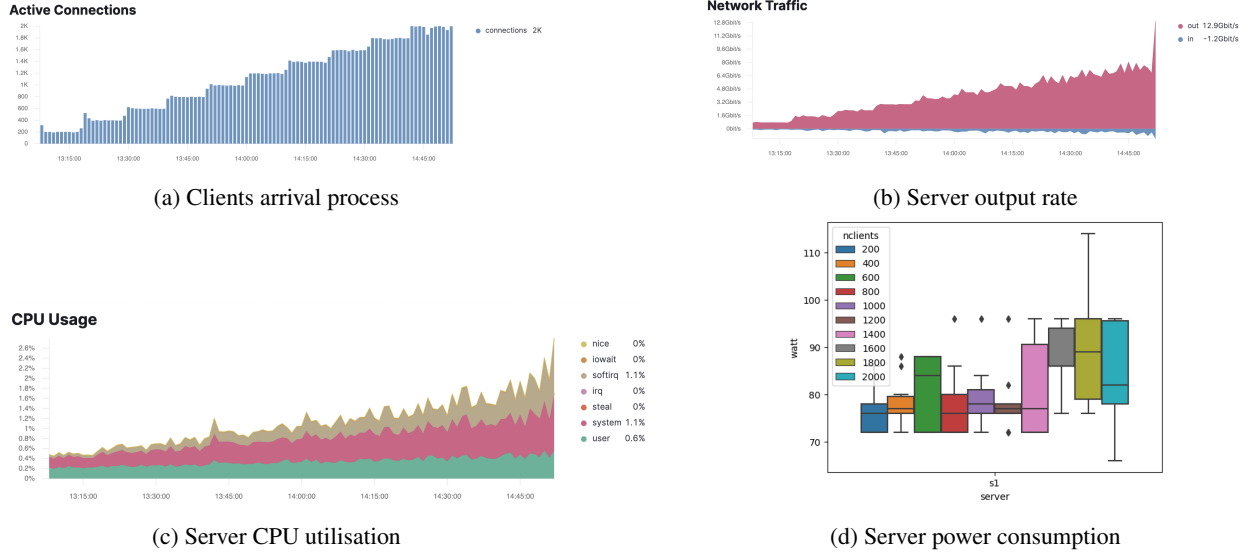


Figure 6: Grid 5000 experiment with one streaming server

2000 clients. Clients arrive in batches of 200 clients as can be seen in Figure 6a. and request the same video at the same quality level. There is thus a single session and the video chunks are kept in RAM. The streaming server is in a dedicated physical server while clients are deployed in containers in a different set of physical servers.

The experiment being made in a controlled environment, all runs are very similar and we present a single representative trajectory of the system. Figure 6b represents the output rate of the server. We reach, at the maximum, a bit rate of about 7 Gb/s (the peak at the end is a measurement artifact). While the bit rate of the server reaches 7Gb/s with 2000 clients, a key observation is that the CPU consumption remains very low as the cache server process involves mostly transfers from the RAM to the network interface card, resulting in a number of software interruptions (softirq). The maximum observed CPU usage is about 3%. As a result, the power consumption of the server raises from 75 W (with 200 clients) to at most 80 to 100 W with 2000 clients, as can be seen in Figure 6d, where each boxplot is obtained with all the one-second long samples in the corresponding phase (constant number of clients) of the experiment. Thus, the server consumption is in the range 0.05-0.3 W/client.

For a popular video where the content is continuously present in the RAM of the server, the power consumption per client is thus small. This however does not rule out the LIMITS model, which points to a blind spot of studies in the domain by considering that if a data center is devoted to a single service, one should consider also idle servers, hence the total consumption of data center. Clearly, more efforts should be done to understand real deployment, like

the one of Netflix⁷ that mixes data centers and their open-connect CDN. In particular, all the hardware, including storage and CDN servers should be accounted.

4.3 Adjusting the values

In addition to the changes discussed in section 4.1, we proceeded as follows to align the models:

- ★ For the SHIFT model, after discussions with the main author of the initial model, we updated the values for the data center efficiencies based on [1].
- ★ We considered a high load of flows for the LIMITS model as this leads to smaller per flow consumption, in line with our measurement campaign and the data center intensity of CARBONTRUST (1.3W/stream).

We present in Figure 7 the (TV,4k,Fixed) scenario, which features a fixed access network that the RENATER and SHIFT models take into account. One observes that some discrepancies remain but the results are (as expected) more similar. A clear tendency is that the network and the user device represent the majority of the cost. The final values are in the range [20,45] grams of CO₂. The smallest estimate comes from the IEA model, which apparently underestimates the networking cost.

As a final remark, note that the results of Figure 7 have been obtained with the energy factor of France. If we consider an average world scenario, the resulting value will be in the range of 60 to 140 grams of CO₂.

⁷<https://openconnect.netflix.com/en/appliances/>

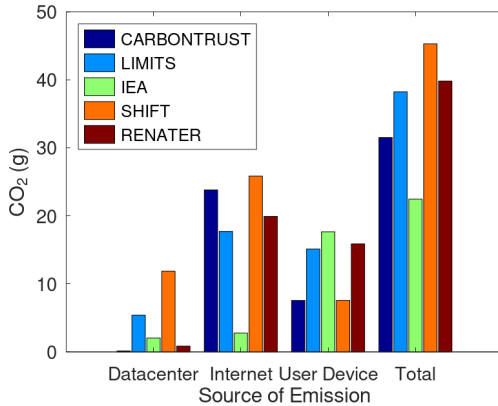


Figure 7: "Aligned" models for (TV,4k,Fixed)

5 Discussion

A number of models have been proposed to estimate the carbon footprint of one hour of streaming. Our contribution in this work has been to dissect a number of these models to understand the roots of the observed discrepancies.

Following the classical breakdown approach into data center, network and user device, we highlighted that the differences stem mostly from (i) the energy efficiency value of the network components (core network and fixed or mobile access) and (ii) accounting for the manufacturing cost or not. Various approaches exist at the data center that are difficult to reconcile. One could argue that the data center contribution appears less important than the two other levels. However, a full answer to this question requires to better understand typical deployments, by accounting for storage server, distribution servers and cache servers. The result is likely to depend on the video broadcaster architecture.

A key difficulty of this comparison is the complexity to obtain good estimates for the constants of the models. There is a clear lack of transparency and data sharing from streaming provider. From this perspective, the RENATER model is interesting as it is based on measurements performed in its own network.

As future work, we would like to further explore the reconciling approach initiated in this paper, and test other scenarios on Grid 5000, e.g. several sessions in parallel and the cache populating phase impact.

We would like also explore the benefits of a typical hybrid distribution scheme where the load on the streaming server is alleviated as clients exchange video chunks with one another. One could argue at first sight that this won't be beneficial as the server consumption is not that high. However, we can expect to trade long range data transfers

with shorter one. The measurements in the Renater network [6] highlighted that the network cost depends on the length of the path. It might thus be interesting, especially for fixed clients.

References

- [1] Anders SG Andrae. New perspectives on internet electricity use in 2030. *Engineering and Applied Science Letters*, 2020.
- [2] Joshua Aslan, Kieren Mayers, Jonathan G Koomey, and Chris France. Electricity intensity of internet data transmission: Untangling the estimates. *Journal of industrial ecology*, 2018.
- [3] Thomas Barnett, Shruti Jain, Usha Andra, and Taru Khurana. Cisco visual networking index (vni) complete forecast update, 2017–2022. *Americas/EMEAR Cisco Knowledge Network (CKN) Presentation*, 2018.
- [4] Lotfi Belkhir and Ahmed Elmeligi. Assessing ict global emissions footprint: Trends to 2040 & recommendations. *Journal of cleaner production*, 177, 2018.
- [5] Francoise Berthoud and Marion Ficher. Évaluation de l’empreinte carbone d’une visioconference entre deux utilisateurs du service rendez-vous. Research report, CNRS - EcoInfo, 2022.
- [6] Marion Ficher, Françoise Berthoud, Anne-Laure Ligozat, Patrick Sigonneau, Maxime Wisslé, and Badis Tebbani. Assessing the carbon footprint of the data transmission on a backbone network. In *24th Conference ICIN*. IEEE, 2021.
- [7] Grid5000. Energy consumption monitoring tutorial. https://www.grid5000.fr/w/Energy_consumption_monitoring_tutorial, 2023.
- [8] George Kamiya. The carbon footprint of streaming video fact checking the headlines. <https://www.iea.org/commentaries/the-carbon-footprint-of-streaming-video-fact-checking>, 2022.
- [9] P-VP Clément Louis-Philippe, Quentin ES Jacquemotte, and Lorenz M Hilty. Sources of variation in life cycle assessments of smartphones and tablet computers. *Environmental Impact Assessment Review*, 2020.

- [10] Stephen Makonin, Laura U. Marks, Radek Przepiński, Alejandro Rodriguez-Silva, and Ramy El-Mallah. Calculating the carbon footprint of streaming media: Beyond the myth of efficiency. *Eighth Workshop on Computing within Limits 2022*, 2022.
- [11] Electricity Map. Electricity map. <https://www.electricitymaps.com/>.
- [12] Netflix. How to control how much data netflix uses. <https://help.netflix.com/en/node/87>, 2023.
- [13] Shift Project. Climate crisis: The unsustainable use of online video. <https://theshiftproject.org/en/article/unsustainable-use-online-video/>, 2019.
- [14] The Shift Project. 1-byte model. <https://theshiftproject.org/wp-content/uploads/2018/10/Lean-ICT-Materials-1byte-Model-2018.xlsx>, 2018.
- [15] Sandvine. 2022 global internet phenomena report. https://www.sandvine.com/hubfs/Sandvine_Redesign_2019/Downloads/2022/Phenomena\%20Reports/GIPR\%202022/Sandvine\%20GIPR\%20January\%202022.pdf, 2022.
- [16] Daniel Schien, Vlad C Coroama, Lorenz M Hilty, and Chris Preist. The energy intensity of the internet: edge and core networks. In *ICT Innovations for Sustainability*. Springer, 2015.
- [17] Arman Shehabi, Sarah J Smith, Eric Masanet, and Jonathan Koomey. Data center growth in the united states: decoupling the demand for services from electricity use. *Environmental Research Letters*, 2018.
- [18] Andie Stephens, Chloe Tremlett-Williams, Liam Fitzpatrick, Luca Acerini, Matt Anderson, and Noor Crabbendam. Carbon impact of videostreaming. <https://www.carbontrust.com/our-work-and-impact/guides-reports-and-tools/carbon-impact-of-video-streaming>, 2021.