



HAL
open science

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

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

► To cite this version:

Joanna Moulierac, Guillaume Urvoy-Keller, Marco Dinuzzi, Zhejiayu Ma. What is the carbon footprint of one hour of video streaming?. Université Côte d'Azur. 2023. hal-04069500v1

HAL Id: hal-04069500

<https://hal.science/hal-04069500v1>

Submitted on 14 Apr 2023 (v1), last revised 2 Nov 2023 (v2)

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.

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

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

Université Côte d'Azur, France
Easybroadcast, France

Abstract

Video streaming traffic dominates the Internet traffic. In such a context, 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. The topic is even becoming sensitive, as highlighted by the recent debate between the SHIFT project in France and the International Energy Agency (IEA) experts.

Our objective in this work is to compare the modeling assumptions and computation methods used by a number of recent works to inform the debate. Indeed, initial results can be at odds, with up one order of magnitude difference in the results. We focus on five different models. Our contributions are: (i) we relate the difference in the results primarily to the perimeter of the study, e.g. including grey energy (production cost) or not, (ii) we demonstrate that some of the assumptions, while necessary, are difficult to estimate due to a lack of public data and (iii) 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 (iv) we propose a technique to reconcile the models and obtain a CO₂ estimate in between 60 and 400 grams when considering the average worldwide electricity efficiency.

1 Introduction

Video streaming traffic now dominates the Internet traffic [3, 16]. Combined with an apparently unbounded growth of Internet traffic [8], video streaming becomes a significant source of CO₂ emissions in the Internet.



Figure 1: Yann LeCun tweet: “ Medieval obscurantism at the EcoInfo group of the CNRS:

"We will not be able to control the energy consumption and the environmental impacts of mobile networks without imposing a form of limitation in use."

What?"

The tweet by Yann LeCun in fig. 1, a senior researcher at Meta, stresses the importance of this debate. His use of the term "obscurantisme médiéval" (medieval obscurantism) is a blatant illustration of how some members of our society, particularly (but not exclusively) those working in the IT sector, consider that, due to its many benefits, the digital sector should not be imposed any bound. Yann LeCun states, in the remaining of his tweet feed, that “40g of CO₂ per hour of streaming is probably a gross overestimate. Recent estimates are in the order of 1 to 2g per

hour of streaming. And the figure is constantly decreasing thanks to the progress of technology”.

To be able to inform the debate and check such a claim, we need to quantify the actual impact of video streaming. To this end, a number of models have been developed recently to assess the impact of one hour of video streaming [19, 11, 14, 9].

Our objective in this work is to investigate the differences observed in the results provided by these models. Indeed, every group of researchers devised its own methodology for calculating the impact of the streaming industry. They further all face a problem of lack of data. For instance, to understand the variability in network energy intensity, more accurate and consistent data on network energy and carbon intensity would be useful. However, it is acknowledged that network operators may not wish to publish this level of detail [19]. The consequence is that results vary significantly from one model to the other. A clear example of this debate is given by the SHIFT Project vs IEA case[9]: The SHIFT Project used their so called ‘1-byte model’ developed in 2019 and obtained an estimate of 0.769 kWh for one hour of video streaming while IEA argued that their colleagues greatly over-estimated the result which was, instead, according to their methodology, 0.078 kWh[9].

In this work, we consider, compare and try to reconcile five models that we term LIMITS [11], CARBONTRUST [14], RENATER[5], IEA [9] and SHIFT [14]. Our contributions are as follows:

- 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;
- We analyse and compare the models under study in a number of scenarios;
- We demonstrate that some of the assumptions, while necessary, are difficult to estimate due to a lack of public data;
- 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;
- We propose a technique to reconcile the models and obtain a CO₂ estimate in between 60 and 400 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 an end user equipment.**

The focus will be on global warming (CO₂ emissions), which is one out of the eight impact factors typically considered in LCA. When focusing on utilisation (hence excluding production costs) we will produce results in kWh.

2.1 High level overview

Before introducing the details of each of the models we consider, let us first provide a high level overview of the common modeling approaches shared by all these models, in terms of functions and components. The classical approach 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 to obtain both estimates and data centers and ISPs are apparently reluctant to provide these

figures. 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 2 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).

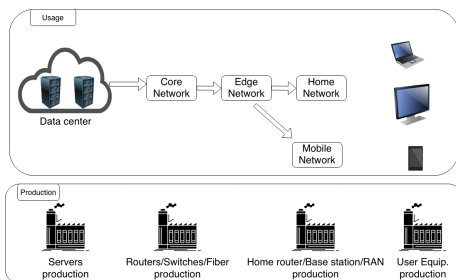


Figure 2: Models components and Emissions Causes

Usage and manufacturing costs need to be aligned. Indeed, the production cost is a fixed cost, “payed” at the beginning of the usage phase. As for the usage cost, it is linearly dependent on the number of years of utilisation. A key assumption behind each model relates thus to the life time of the equipment. Given that manufacturing costs are high in general, the choice of the life time is crucial. Another point to consider is the unit. Production cost are obtained from

manufacturers environmental shifts [10] which provide total CO₂ figures. As for usage, it can be either in kWh or CO₂ depending on the study. Translating kWh into CO₂ for the usage part is dependent on the country considered since the energy factor¹ highly depends on the production plants producing electricity, see for instance [12].

When focusing on the global footprint of one hour of streaming, we use CO₂ emissions, while we use kWh when considering usage phase. Unless stated otherwise, we use the french electricity factor of 108 g of CO₂/kWh similarly to [5].

We consider five models: LIMITS [11], CARBONTRUST [14], RENATER[5], IEA [9] and SHIFT [14]. 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 modeling 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 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 [15].
- 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

¹The energy factor indicates how much CO₂ grams are emitted to produce on kWh.

with actual measurements made in the Renater network and data center since the application (called Rendez-vous), 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 [9]. 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 [11]. 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 [18] 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:

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

where:

1. Pw is the power demand in kilowatts (kW) based on the data center size,
2. D is the duration of the video being streamed in hours (1 hour in our case).

3. $\# Streams$ 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).

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

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

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

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:

$$DC_{energy} = AnnualC_{DC} \cdot \frac{Q}{AnnualTraf} \cdot PUE \cdot EF$$

where:

1. $AnnualC_{DC}$ is the annual energy consumption in kWh of the data center.
2. Q indicates the traffic generated by the user and Annual Traffic 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 CO₂ emitted to produce one kWh, since the RENATER paper produces CO₂ 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 :

$$DC_{man} = \text{Prod}_{\text{DC}} \cdot \frac{D}{\text{Life}_{\text{EQUIP}}}$$

where:

1. Prod_{DC} is the total manufacturing cost (provided by the manufacturer) in kg of CO₂ .
2. $\text{Life}_{\text{EQUIP}}$ is the life time of the equipment. Note that the original paper presents a much more complex formula as they account for the lifetime of the VM hosting the service. We assume here that a physical server is dedicated to the streaming service.

SHIFT The Shift project estimates the data-center energy consumption as follows:

$$DC_{energy} = \text{EI}_{\text{DC}} \cdot Q$$

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 [15].
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.

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

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

The CoreNet_{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:

$$\text{CoreNet}_{energy} = I \cdot D \cdot R$$

where:

1. 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 [17].

The ISP_{energy} is computed as:

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

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:

$$FixedNet_{energy} = I_{FN} \cdot D \cdot R$$

where:

1. I_{FN} is the energy intensity of fixed network transmission (0.0065 kWh/GB) [19]
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:

$$MobileNet_{energy} = I_{MN} \cdot D \cdot R$$

where I_{MN} is the energy intensity of mobile network transmission (0.1 kWh/GB) [19]

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:

$$Router_{energy} = I_{HR} \cdot D \cdot R$$

where I_{HR} is the energy intensity the home router (0.025kWh/GB) [19]. Therefore if the user uses a fixed network we will have:

$$Net_{energy} = (I_{FN} + I_{HR}) \cdot D \cdot R$$

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

$$Net_{energy} = (I_{FN} + I_{MN}) \cdot D \cdot R$$

RENATER This model, as in the data center case, accounts for both usage and production. This leads to the following equation for the network²:

$$Net_{energy} = (Router + NOC + Fiber) \cdot Q \cdot EF +$$

$$Prod_{NET} \cdot \frac{D}{Life_{EQUIP}}$$

where:

²The Renater equations unit is CO₂. One needs to divide by EF to get the results in kWh

1. Router, NOC (Network Operating Center) and Fiber costs are utilisations in kWh/GB
2. $Prod_{NET}$ is in kg of CO₂/GB.

SHIFT The approach is similar to the CARBONTRUST 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 CARBONTRUST) or also the network equipment at home (RENATER).

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

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

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 laptop³ 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

³Table 3 in [11] 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.

total CO₂ emissions during the life time of the equipment:

$$\text{Prod}_{\text{DEVICE}} = 9 \cdot \text{Device}_{\text{energy}}$$

This assumption might seem a bit too high, as typical values lay in the range [60%, 80%] [10]. 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).

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:

$$\text{Device}_{\text{energy}} = (W_{\text{S}} + W_{\text{P}}) \cdot D$$

where:

1. W_{S} is the average power consumption of screens.
2. W_{P} is the average power consumption of peripherals.

$W_{\text{S}} = 1\text{W}$ 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 18W for a set top box. In section 3, we will take $W_{\text{P}} = 0$, 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:

$$\text{Device}_{\text{energy}} = \frac{\text{Annual}C_{\text{USER}}}{\text{Tot}D_{\text{VOD}}} \cdot Q \cdot EF$$

$$\text{Prod}_{\text{DEVICE}} = \frac{\text{ManufImpact} \cdot D}{\text{Life}_{\text{EQUIP}}}$$

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

$$\text{Router}_{\text{energy}} = \frac{\text{Annual}C_{\text{ROUTER}}}{\text{AnnualTraf}} \cdot Q \cdot EF$$

$$\text{Prod}_{\text{ROUTER}} = \frac{\text{ManufImpact} \cdot D}{\text{Life}_{\text{ROUTER}}}$$

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

$$\text{Device}_{\text{energy}} = \text{EI}_{\text{UD}} \cdot D$$

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 setup, end user device and the quality level of the multimedia flow.

3 Results

3.1 Scenarios

We compare in this section the RENATER, LIMITS, CARBONTRUST SHIFT and IEA. Values for the latter are obtained thanks to the web interface provided by the author – see details below. We decided to focus on the following scenarios:

- User device (TV, Laptop, Smartphone).

- Video Resolution (SD, HD, 4k)
- Network Type (Fixed or Mobile)

Similarly to [19], the bit rates at each resolution are set to:

- SD = 2.22 Mbps
- HD = 6.67 Mbps
- 4k = 15.56 Mbps

These values are the (maximum) ones advertised by Netflix at high quality on a "web browser"[13].

We also account for other factors such as the dimension of the data center (Big, Medium, Small) and the number of active flows from 10k to 50k, following the values considered in the LIMITS paper.

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 a number of scenarios:

- Smartphone, Laptop, TV
- Mobile or Wifi (high) defined as " case [that] assumes that the streaming video device is the only connection"
- SD, HD and 4k, the latter being able only for Laptop and TV, which makes sense.

The values extracted 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 site⁴.

In table 5, one can observe 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.

⁴<https://www.iea.org/reports/global-energy-review-2019/co2-emissions>

Eventually note that the IEA web site also presents a comparison for some scenarios with the SHIFT model. We checked the SHIFT values presented and established that the so-called 'Shift project (corrected)' is the original 1byte model available in [15], that we presented in the section 2.2. We thus use this model for the results presented hereafter and not the IEA data.

3.2 Overall comparison

Let us first consider a typical scenario and compute the overall energy consumption for all models. We consider two emblematic scenarios:

- Scenario "TV", with a TV receiving a 4k stream at home;
- Scenario "Smartphone", with a smartphone receiving a SD stream with a mobile connection.

We also assumed for both scenarios a large data center that is lightly loaded for the LIMITS model.

In fig. 3 and fig. 4, we report, for the two scenarios, the breakdown for the three tiers, namely data center, network and user device, as well as the total (sum) values. We included the manufacturing costs in this figure and used the energy emission factor of France (108 g CO₂/kWh as in [5]) for the utilisation factors.

We can make the following observations for the two scenarios:

- IEA always returns the smallest values, among all models, while the SHIFT model features the largest ones.
- 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 mobile case, but becomes significant for the wired case. Also, a higher screen requiring a higher resolution, this results in a higher load in the network.
- The SHIFT model features the highest network cost among all models, which explains most of its cost.

- 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.




Figure 3: Total Emissions of the represented scenario for scenario (TV,4k,Fixed)




Figure 4: Total Emissions of the represented scenario for scenario (Smartphone,SD,Mobile)

3.3 User Device

We present in fig. 5 the energy consumption (production is excluded) of the device for the different models and different types of devices. We




Figure 5: Comparison of the energy consumption all considered user device

chose a scenario with HD resolution and Fixed Network as the results are qualitatively similar for the other cases. We express the results in kWh and not grams of CO₂, since it is a direct electricity consumption.

Figure 5 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 for the model assumes a large TV of 55 inches, which might explain its large value.

3.4 Network

Figure 3 and fig. 4 highlight 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, which correspond here to 2.22 and 15.56 Mbps respectively.

The RENATER model factors a manufacturing cost in its network footprint estimation. It accounts for about 40% of the network cost.

3.5 Data center

As can be observed in table 2, 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 more bottom-up models that focus on the streaming server consumption only. However, the SHIFT model considers that the consumption is a function of a amount of byte while the CARBONTRUST consider the 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. 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.

3.6 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 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. One could argue that this figure might be an 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 fig. 6 both the utilisation and production cost of each tier. We focus on the same (TV, 4k, Fixed) scenario as in section 3.2. We directly observe that the obtained manufacturing contributions are especially significant for the networking and user device levels. While the latter is in line with the environmental reports provided by the user device manufacturer, the former is more astonishing, as one could have expected that the long lifetime of the equipment would wipe out the production effect. Also the production cost at the data center while small, is noticeable. It is again against what is often argued in the literature, e.g. [4].

4 Reconciling the models

In this section, we present our preliminary ideas to reconcile the models, by selecting the best approach for each tier. As we will see, this approach 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 Grid'5000 testbed with a streaming server servicing up to 2000 clients.




Figure 6: Usage versus Manufacturing - RENATER model

4.1 Towards a unified model

Reconciling the models is relatively easy for the utilisation part of the user device impact. Indeed, it simply consists in picking the same power demand for all models (except IEA) for a specific device, e.g. 70W for the TV. For the production part, the approach of the LIMITS model offers a simple yet easy to compute solution.

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 in the literature 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 a network emission factor for the FAN equal to $1.7e^{-3}$ CO₂/GB, which corresponds to 0.016 kWh/GB.

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 2000 clients. Clients arrive in batches of 200 clients as can be seen in fig. 7a. All clients 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 7b 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). 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 a most 80 to 100 W with 2000 clients, as can be seen in fig. 7d, where each boxplot is obtained with all the one-second long samples in the corresponding phase (constant number of clients) of the experiment. This means that the server consumption is around 0.05-0.3 W/client.

A conclusion we can draw from this experiment is that for a popular video where the content is continuously present in the RAM of the server, the power consumption per client is small, for the simple reason that the video streaming workload is not CPU intensive. This justifies a modeling approach where the energy consumption of the server is based on the session duration and not the number of clients in parallel. This does not rule out the LIMITS model, which points to a blind spot of studies in the domain

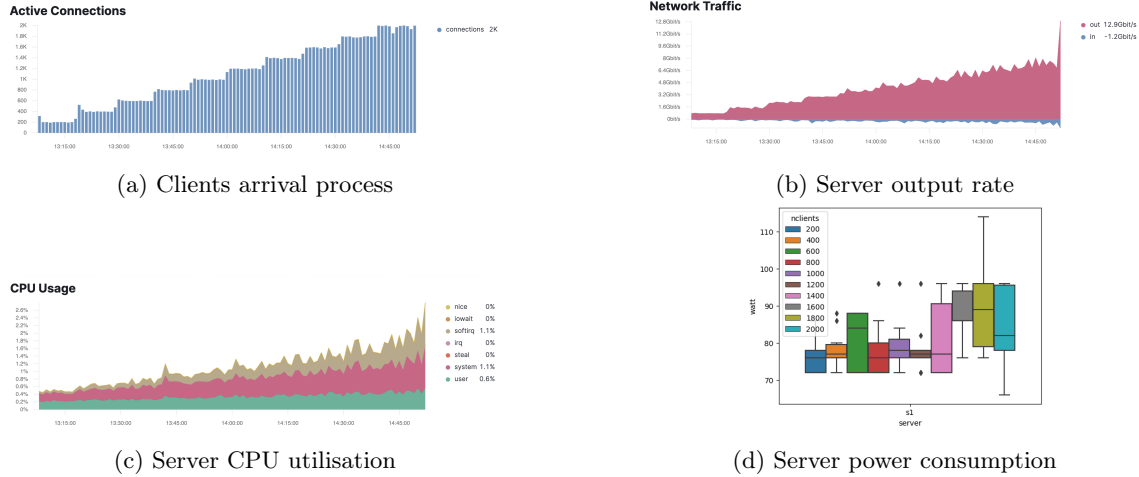


Figure 7: Grid 5000 experiment with one streaming server

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 scenarios should be explored, e.g. a server that needs to populate its cache with new videos regularly or quality changes. We leave it for future work.

4.3 Adjusting the values

In a first attempt to align the models and obtain a range of CO₂ emissions associated to one hour of streaming, we proceeded as follows:

- For the user device, we picked for the utilisation contributions, the smallest power demand among all the models and used it as input for each model.
- For the SHIFT model, after discussions with the main author of the initial model, we updated the values for the network and data center efficiencies based on [1].
- We considered a high load in terms of number of flows for the LIMITS model as this leads to smaller per flow consumption, in line with our measurement cam-

paign and the data center intensity of CARBONTRUST (1.3W).

We present in fig. 8 a (TV,4k,Fixed) scenario, which has the advantage of featuring 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 [20,120] grams of CO₂. The largest estimate comes from the LIMITS paper, not because of the way the data center is modeled, but mostly because of the production cost for the TV. As discussed previously, the LIMITS approach makes sense and can't be ruled out. As for the IEA model, it apparently underestimates the networking cost.

As a final remark, please note that the results of fig. 8 have been obtained with the energy factor of France, which is 3 times smaller than the average worldwide energy factor. The resulting value will thus be in the range of 60 to 400 grams of CO₂ if we consider an average world scenario.

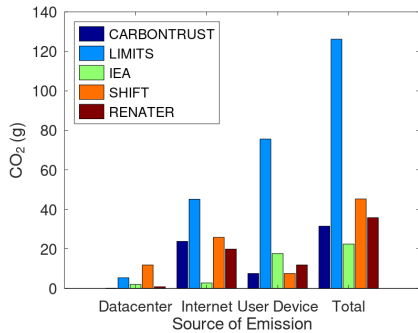


Figure 8: "Aligned" models

5 Discussion

A number of models have been proposed to estimate the carbon footprint of a typical hour of streaming. Our contribution in this work has been two consider a number of these models and dissect them 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) the choice of the choice made to account for the manufacturing cost or not. It is especially important in the network and user device tiers. There exist various approaches at the data center that are difficult to reconcile. However, as the data center contribution seems less important than the two other levels, the resulting discrepancies do not affect too much the global trend.

A key difficulty of the exercise we made is the complexity to obtain good estimates for the constant 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 the Renater network.

As future work, we would like to further explore the reconciling approach initiated at the

end of the paper, possibly incorporating more models. Obtaining more information on the production cost for the network and data center will be a key point. Following the approach for the RENATER model [5], we would like to check also if additional components (storage server, maintenance, etc) can be taken into account. At some point, we expect to reach a trade-off between the level of details and the availability of operational data.

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*, 3(2):19–31, 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*, 22(4):785–798, 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*, pages 1–30, 2018.
- [4] Lotfi Belkhir and Ahmed Elmeligi. Assessing ict global emissions footprint: Trends to 2040 & recommendations. *Journal of cleaner production*, 177, 2018.

- [5] Françoise Berthoud and Marion Ficher. Évaluation de l’empreinte carbone d’une visioconférence entre deux utilisateurs du service rendez-vous. Research report, CNRS - EcoInfo, March 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 on Innovation in Clouds, Internet and Networks and Workshops, ICIN 2021, Paris, France, March 1-4, 2021*, pages 105–109. IEEE, 2021.
- [7] Grid5000. Energy consumption monitoring tutorial. https://www.grid5000.fr/w/Energy_consumption_monitoring_tutorial, 2023.
- [8] IAE. Digitalisation and energy. <https://www.iea.org/reports/digitalisation-and-energy>, 2017.
- [9] 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-the-headlines>, 2022.
- [10] 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*, 84:106416, 2020.
- [11] Stephen Makonin, Laura U. Marks, Radek Przedpejski, Alejandro Rodriguez-Silva, and Ramy ElMallah. Calculating the carbon footprint of streaming media: Beyond the myth of efficiency. *Eighth Workshop on Computing within Limits 2022*, 2022.
- [12] Electricity Map. Electricity map. <https://www.electricitymaps.com/>.
- [13] Netflix. How to control how much data netflix uses. <https://help.netflix.com/en/node/87>, 2023.
- [14] Shift Project. Climate crisis: The unsustainable use of online video. <https://theshiftproject.org/en/article/unsustainable-use-online-video/>, 2019.
- [15] The Shift Project. 1-byte model. <https://theshiftproject.org/wp-content/uploads/2018/10/Lean-ICT-Materials-1byte-Model-2018.xlsx>, 2018.
- [16] 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.
- [17] 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*, pages 157–170. Springer, 2015.
- [18] 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*, 13(12):124030, 2018.
- [19] 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.

Global Notation	
P_w	power demand (in kW) based on the data center size
D	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: $Q = R \cdot D$
Data Center contribution	
# Streams	Number of simultaneous streams served by the data center.
ID_C	Energy intensity of the data centre (in W). Estimated as 1.3W.
Annual C_{DC}	Annual energy consumption in kWh of the data center.
Annual T_{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.
Life $_{EQUIP}$	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)
Life $_{ROUTER}$	Life Time of Router
Prod $_{DEVICE}$	Production cost of user device
Prod $_{ROUTER}$	Production cost of router

Table 1: Notation used throughout the paper.

Datacenter		
	Usage	Manufacturing
LIMITS	$\frac{P_w \cdot D}{\#Streams}$	
CARBONTRUST	$ID_C \cdot D$	
RENATER	$AnnualC_{DC} \cdot \frac{Q}{Annual_{Traf}} \cdot EF \cdot PUE$	$Prod_{DC} \cdot \frac{D}{Life_{EQUIP}}$
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.

Network		
	Usage	Manufacturing
LIMITS	$I \cdot D \cdot R + I_{ISP} \cdot D$	
CARBONTRUST	$(I_{FN} + I_{HR}) \cdot D \cdot R$ or $(I_{FN} + I_{MN}) \cdot D \cdot R$	
RENATER	$(Router + NOC + Fiber) \cdot Q \cdot EF$	$Prod_{NET} \cdot \frac{D}{Life_{EQUIP}}$
SHIFT	$EI_{NET} \cdot Q$	

Table 3: Network Contribution

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

Table 4: User Equipment contribution

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