Go with the Flow: Analyzing the Carbon Footprint of Green Streaming

Jorge Murillo
University of Massachusetts Amherst
Amherst, MA, USA
jrmurillo@umass.edu

Walid A. Hanafy
University of Massachusetts Amherst
Amherst, MA, USA
whanafy@cs.umass.edu

David Irwin
University of Massachusetts Amherst
Amherst, MA, USA
irwin@ecs.umass.edu

Ramesh Sitaraman
University of Massachusetts Amherst
Amherst, MA, USA
ramesh@cs.umass.edu

Prashant Shenoy
University of Massachusetts Amherst
Amherst, MA, USA
shenoy@cs.umass.edu

Abstract

Video streaming has significantly surged over the past two decades and now accounts for 50% of Internet traffic, raising concerns regarding its environmental impact. Prior work on carbon-aware video streaming, often denoted as "green streaming", has generally focused on streaming from low-carbon data center regions. However, these approaches often consider only the source data center or rely on coarse-grained attribution models that fail to capture network emissions incurred during streaming. To understand and account for carbon emissions associated with video streaming, we first highlight the limitations of state-of-the-art carbon attribution methods. We then introduce a fine-grained, end-to-end model that enables clients to attribute the average emissions of a video streaming session across the playback devices, data centers, and the network path. We also demonstrate the utility of our carbon attribution model in analyzing the carbon footprint of video streaming using real-world datasets of network topology and routes, as well as the actual cloud and edge locations of commercial services. Our analysis covers 1322, 637, and 235 cloud and edge data center locations in North America, Europe, and Asia, respectively, as well as grid carbon intensity data from 374 regions worldwide. Our results show that streaming from the greenest data center region without considering the incurred network emissions can result in 52%, 44%, and 43% higher attributed emissions, on average, for clients in North America, Europe, and Asia, respectively. Our results also demonstrate that streaming from the nearest edge site is frequently suboptimal, while streaming from a slightly further edge or cloud region can lower the average emissions for a video stream more than 77% of the time globally.

ACM Reference Format:
Jorge Murillo, Walid A. Hanafy, David Irwin, Ramesh Sitaraman, and Prashant Shenoy. 2026. Go with the Flow: Analyzing the Carbon Footprint of Green Streaming. In *Proceedings of e-Energy '26.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

e-Energy '26, Banff, Canada

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Video streaming over the Internet has grown significantly in popularity over the past two decades. Today, streaming services, such as YouTube, Netflix, and Twitch, are seen as popular sources of daily entertainment, with many users spending up to four hours per day consuming digital media [14]. Recent studies highlight that more than half of the total Internet traffic in many countries originates from video streaming [61], and today's streaming services consume a significant fraction of Internet resources, as well as server resources in cloud and edge data centers where the content resides. The popularity of streaming services, along with their associated data center and network workloads, has raised concerns about the environmental footprint of video streaming. For example, the IEA estimates that one hour of streaming video emits 36g·CO_2eq [37]. Given that large streaming services have hundreds of millions of active users, such services can result in a considerable overall carbon footprint [47].

One approach to reduce the carbon emissions of video streaming is the use of spatial shifting, often denoted as green streaming[1], where users stream the video content from a data center location with low carbon intensity. Although spatial workload shifting is an effective technique for reducing the carbon footprint of data centers, especially for workloads with spatial flexibility [11, 15, 56], such techniques have often focused on compute-intensive tasks, primarily concentrating on data center-side carbon optimizations, and have overlooked the network impact, as these workloads do not utilize significant network bandwidth. Thus, server-side-only methods are often inadequate for optimizing video streaming emissions, which consume significant network bandwidth and server resources. Even recent work on spatial workload shifting in content delivery networks (CDNs) has made similar assumptions, resulting in an underestimate of the carbon footprint of video streams [23, 50]. Another approach is to design green network transfer methods that choose the greenest path between the server and the client [18, 24, 72]. However, since video streaming also consumes significant server resources, choosing the greenest network path for streaming while ignoring the server impact can also be suboptimal. Finally, there have been decades of work on enhancing the energy efficiency of both the network [6, 36, 52, 59, 63] and the data centers [38, 66]. While increasing energy efficiency can also

[1]We use these two terms interchangeably in the paper

reduce the carbon footprint of streaming services, they do not directly consider region-specific differences in grid carbon emissions, which can impact the overall emissions of video streaming.

Therefore, to design effective and practical green streaming approaches, there is an urgent need to analyze and account for the carbon emissions of both servers and the network paths used for video streaming. Such analysis can be done using attributional methods, which use average carbon intensity data of different grids and mean energy usage of servers and network elements to determine the average carbon emissions of each stream. Such methods are useful for carbon reporting purposes and for understanding the average per-stream emissions incurred by a CDN streaming service. Alternatively, the analysis can use consequential methods that estimate the marginal emissions of each stream, an approach that is useful for decision making [8, 13, 19]. Our goal in this paper is to perform an attributional analysis to understand the average per-stream emissions incurred by a CDN when streaming videos from different data center locations and along different network paths.

Prior attributional analysis on quantifying the environmental impact of video streaming [3, 37, 47, 70] have used *coarse-grain* models that consider a fixed energy per stream (kWh/GB) [37, 47, 51, 70] and often assign an average region-agnostic carbon intensity estimate for data centers and networks [3, 47]. However, the energy of the end-to-end network path varies significantly across client and server locations due to variations in the number of network hops and traversed distance. In addition, recent work has already shown that carbon intensity can vary significantly across different spatial regions and times through which the network traffic flows [28, 72]. Hence, the use of prior coarse-grain methods does not reveal the full carbon impact of video streaming. For instance, a streaming session that utilizes a data center with low carbon intensity but traverses a long network path or one that goes through regions with high carbon intensity may have a higher overall carbon footprint than a session that uses a data center with higher carbon intensity but traverses a network path with lower carbon emissions. In Section 3.1, we show how a coarse-grained model does not fully capture such impacts and can under- or over-estimate the carbon emissions of a video streaming session. Thus, it is crucial to develop fine-grain models that capture both data center and network carbon emissions when analyzing the full emissions of video streaming.

Motivated by these limitations, we propose a fine-grained attribution carbon accounting model that couples time- and location-varying grid intensity with streaming device and server energy and path-specific network emissions across edge and cloud data centers. Unlike prior work, which employs coarse-grain methods, our approach aims to develop fine-grain estimates at a per-stream level, considering both the data center location and the entire network path used to deliver the stream. We then utilize this model to conduct a model-driven measurement study that empirically quantifies the attributed end-to-end carbon footprint of video streaming. Importantly, our model and measurement study also consider the carbon impact of both cloud streaming and edge streaming, since the latter is widely used by content distribution networks employed by streaming services. Using this lens, we also expose cases where the *apparent* choices—serving from low-carbon data centers or from the nearest edge data center—or relying on coarse-grained network energy consumption paradoxically have higher attributed carbon emissions to a video streaming session due to the high carbon intensity of edge data centers. We then show how our model can better capture end-to-end attributed carbon emissions in video streaming, offering more accurate estimates and better carbon accounting for video streaming applications.

Our model-driven attributional analysis seeks to answer several key questions: (i) What is the relative contribution of the devices, network, and servers to the overall attributed carbon footprint of a video streaming session, and which one is dominant? (ii) Does streaming from the nearest location lead to the lowest carbon emissions, and how does our fine-grained attribution model affect server selection? (iii) From an emissions standpoint, is streaming from the edge always better than streaming from the cloud? (iv) How do these results vary across different parts of the world, different times of the day, and settings?

In developing our model and measurement study, our paper makes the following contributions:

(1) **Fine-grained model.** We develop a fine-grained attribution model that estimates end-to-end carbon emissions of video streaming by accounting for the playback device, the server, and network energy, data-center location, end-to-end path, and spatiotemporal variations of energy's carbon intensity.
(2) **Real-world Analysis.** To answer the above questions, we conduct a model-driven analysis of the carbon footprint attributed to a streaming session when the video is streamed from different data center locations and along different network paths. Our analysis uses real-world datasets of network topology and routes, actual cloud and edge locations of commercial services, including 1322, 637, and 235 cloud and edge data center locations in North America, Europe, and Asia, as well as grid carbon intensity data from 374 regions worldwide.
(3) **Key Findings:** Our study reveals many interesting insights. (i) We demonstrate that network emissions are a crucial factor for designing green streaming methods, particularly when streaming from distant data centers, highlighting the utility of our fine-grained attribution model. Our results show that streaming from the greenest server location, without considering the network path, increases attributed carbon emissions by 52%, 44%, and 43% for clients in North America, Europe, and Asia, respectively. (ii) We also show that edge streaming is not always a greener choice than cloud streaming, even though it incurs fewer hops to the client, where the relative greenness of edge and cloud streaming depends on the specific data center locations and network paths used. In many cases, a slightly more distant edge or cloud data center may offer lower attributed emissions for a video stream. Our results show that using the nearest edge data center leads to higher emissions in 90%, 82%, and 77% of cases, across North America, Europe, and Asia, respectively. (iii) Finally, we find that the results vary significantly across different parts of the world and different times of the day, highlighting the importance of considering spatio-temporal variations in carbon intensity and the need for adaptive green streaming methods that can dynamically select the greenest option based on the current conditions.

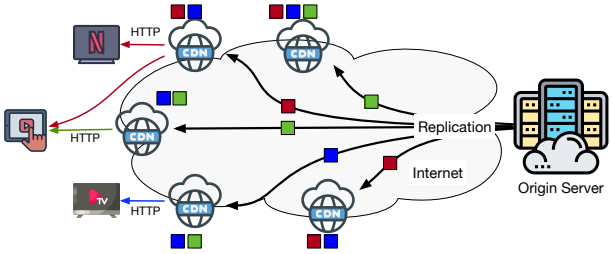

Figure 1: Main components of a video streaming service.

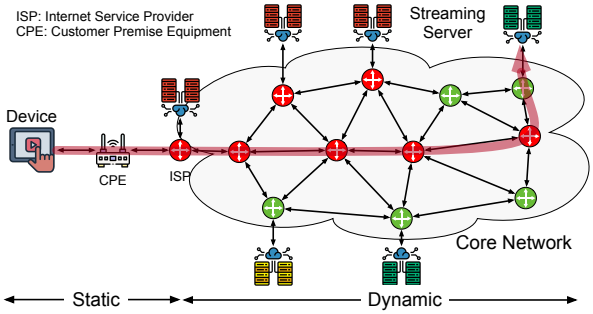

Figure 2: Network path for a green streaming approach, incorporating both static and dynamic network components, highlighting how a green server might be offset by a brown network path.

2 Background

In this section, we provide background on video streaming, our energy and carbon accounting methodology, and networks, and approaches to designing green streaming services.

2.1 Video Streaming

A video streaming service, e.g., Netflix, delivers a rich set of video content to its users with the ability to serve millions or billions of video files and may serve tens of millions of users. Let us consider how an individual video file is streamed from a content/source data center to a user's device. Figure 1, depicts the various components of modern video services. Typically, the content server is located in a cloud data center, often referred to as the content origin; however, many streaming service providers rely on content delivery networks (CDNs) to cache content at edge services and stream it from locations closer to users. We will refer to streaming content from a cloud data center as "cloud streaming", and to streaming from an edge data center as "edge streaming." Moreover, video content is often replicated across multiple cloud and edge data centers, allowing a particular file to be streamed from any of the edge or cloud replicas. It is then up to the streaming service or the underlying CDN to select a specific data center or CDN location. For example, a streaming service may choose a region-specific data center to service a particular client (e.g., European clients may be serviced by CDNs/Cloud data centers in Europe [21]). Similarly, a CDN, which replicates content across a number of edge data centers, may choose the closest edge data center to serve a client, or the edge data center that minimizes the bandwidth cost [1].

Once a data center has been chosen to serve the client, a video streaming protocol is used to deliver content from the server to the client device. Modern video streaming services, also known as adaptive streaming [73], often rely on an HTTP server that hosts the video content. The streaming server then runs an adaptable streaming protocol, such as HLS [55], and MPEG-DASH [68], which allows for streaming service providers to divide videos into small *chunks* (e.g., a few seconds) with different versions (frame rate, quality, and sizes), collectively denoted as a bit rate ladder [2, 22, 73]. At streaming time, the client (or the client device) selects one of the versions, based on device characteristics (screen size), network conditions (4G or WiFi), or user preference [2, 22, 73]. In addition, to adapt to variable network conditions, video streaming applications employ adaptive bit rate (ABR) optimizations [2, 3, 22, 73]. ABR approaches, such as SODA [12] and BOLA [69], automatically select the *right* segment to navigate the trade-offs between video quality and sizes and avoid video freezes. In addition, approaches such as GreenABR[74] and E-WISH[44] enhance the selection process by incorporating the energy consumption of both the device and the network.

2.2 Carbon Accounting

The Green House Gas (GHG) protocol specifies two types of methods for estimating emissions: *attributional* and *consequential* [8]. Attributional methods focus on fair sharing of system-wide emissions and tend to use average grid carbon intensity values and average energy usage across the system when estimating the emissions of a particular entity (or stream). Such methods are useful for accounting and reporting purposes and for system-wide analysis. In contrast, consequential methods focus on actions that reduce system-wide emissions from an action or a decision. Recently, the networking community has also highlighted the differences between the operational and consequential methods when analyzing the energy (and carbon) usage of network workloads [13, 17, 19] and challenges in estimating the emissions for either accounting methodology.

Regardless of which method is used, the emissions themselves come in two forms: operational and embodied. Operational emissions are the emissions resulting from consuming energy during the operations of a system or direct emissions; these are also known as Scope 1 and 2 emission in the GHG protocol. Embodied emissions (also known as scope 3 emissions) are the emissions resulting from the manufacture of a product (e.g., manufacture of server and router hardware), its transportation, and setup [29, 39]. Recent systems research has developed carbon optimization methods such as temporal and spatial workload shifting to optimize the operational emissions of data centers [23, 31, 50, 71], while research in the architecture community has focused on optimizing embodied emissions of computing hardware. Interestingly, the operational emissions of one entity can become the embodied emissions of another entity that uses the former. For example, the operational emissions of the servers and GPUs used to train a machine learning model become the embodied emissions of the model when it is deployed for inference. Similarly, the operational emissions from the server and network elements used by a CDN to stream a video will

become the embodied emissions of the end-user's video application that consumes it.

In this work, we use attributional analysis to estimate the operational emissions from streaming a video from different data center locations and along different network paths. Our goal is to highlight the differences in operational emissions across data center and client locations and across different paths. The key challenge in estimating the energy consumption for a video streaming session, at the flow level, is that the energy usage of the server and network routers depends on their utilization levels. However, full knowledge of utilization levels of servers and network elements and the network state such as throughput [19] is challenging to obtain and most Internet traces do not capture this information. Fortunately, servers and network elements such as routers and switches are not energy proportional and, as a result, changes in traffic and workload level have less of an impact on their power consumption [18, 19, 25, 35, 51, 67]. Since our focus is on attributional, rather than consequential, analysis, this allows us to rely on static or average power consumption values for servers and network devices when conducting our model-driven analysis.

2.3 Energy and Carbon Attribution

Regardless of the streaming protocol, each video stream must go through: (1) a server at a particular edge or cloud data center location that streams the content, (2) a network path over the internet, comprising multiple network hops, over which this content traverses, and (3) a client device at the user's location that receives and renders the content. Given these components, for each video stream, we consider the energy consumption at the server for sending the video, as well as the energy consumption along the network path for transmitting this video. Once the energy consumption of the video stream is estimated, it can be translated into a carbon footprint for that streaming server by considering the carbon intensity (CI) of the electricity supplied across the server and various network components. In general, the carbon intensity of electricity, measured in g·CO_2eq/kWh, represents the amount of emitted CO_2 per unit of energy consumed. The CI is computed as a weighted average of the electricity supply mix for a particular grid and time period. More importantly, as highlighted in earlier research [71], energy's carbon intensity varies across locations (i.e., electricity grids) and over time, with the availability of renewable energy sources (e.g., solar or wind) yielding lower CI (see Figure 5). The carbon footprint of a video stream is simply the sum of the product of the energy consumption at each element (server, network) and the CI of that location at the time of streaming.

2.4 Green Streaming and Networking

Recent efforts to quantify the carbon footprint of video streaming services make several assumptions to simplify their analysis. For example, several studies make coarse-grain assumptions and use a single static value for the network energy cost of streaming [37, 47, 51, 70]. Moreover, these studies do not capture the path-specific nature or assume a single average value of carbon intensity for the electricity supply [16, 37]. However, as shown in Figure 2, a selected server can lead to different network paths with varying numbers of hops or carbon intensity. While these assumptions are viable when performing large-scale analysis, they do not accurately capture the actual emissions attributed to a streaming session.

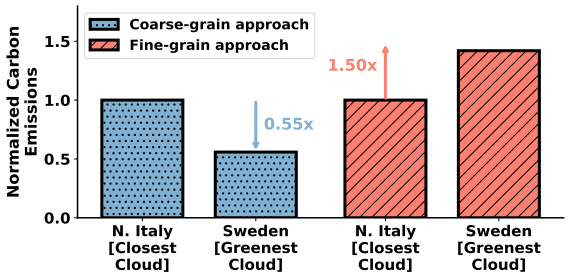

Figure 3: Streaming carbon emissions from the closest and greenest cloud data center in Europe for a client in Perugia, Italy, using coarse-grain and fine-grain accounting models.

To optimize the carbon emissions attributed to a video streaming session, recent work has leveraged the spatial variations of carbon intensity across locations. In the case of video streaming, earlier approaches have relied on the availability of content across data centers and have developed methods that select a replica in regions with low carbon intensity [46, 50]. However, as shown in Figure 2, a green streaming server might be accessed via a brown network path, which may overtake these gains. Conversely, researchers have also exploited the time-varying nature of carbon intensity and designed approaches that schedule data transfer tasks [24] during periods of low carbon intensity. Moreover, authors have proposed carbon-aware routing and traffic engineering approaches [18, 54, 72], enabling the selection of network paths that traverse locations with low carbon intensity. Nonetheless, these approaches either do not consider the network [46, 50], or do not consider the server emission [24], or rely on coarse-grained estimates[47].

As a result, prior approaches do not fully capture one or more of the following: (1) detailed server energy cost, (2) full network route and energy cost, (3) carbon cost for a streaming server edge or cloud location and time, (4) carbon cost of routes based on the specific location and time along the path. Next, we present our end-to-end model that addresses these requirements when attributing carbon emissions to a video streaming session.

3 Methodology

This section introduces a fine-grained model for estimating the carbon emissions attributed to a video streaming session, attributing emissions across the device, network, and server components for content served from edge or cloud data centers. First, we motivate the need for a fine-grained model to attribute carbon emissions from video streaming across different sources accurately. Then, we present an overview of our proposed model and its key elements. Next, we outline our proposed attribution approach, which resolves emissions along physical network routes across time, geography, and network segments. Lastly, we conclude with a list of carbon-aware ("green") streaming methodologies that we use in our analysis.

3.1 Motivation

To estimate the carbon emissions of video streaming, prior work often employs coarse-grained models–either a fixed energy-per-stream (kWh/GB) across the Internet or high-level averages per network type [37, 47, 51, 63, 70]. While useful for rough estimates, these abstractions cannot accurately estimate carbon emissions attributed to a video streaming session as they implicitly assume that network energy depends only on data volume, without considering the end-to-end path. In reality, however, physical paths differ across clients and serving locations, due to changing (i) Device counts (L2/L3 hops, Internet Exchange Points (IXPs)), (ii) Distance (and hence amplifiers/regenerators), and (iii) Carbon intensity of the regions traversed. Consequently, transmitting the same video stream from a different region can significantly increase or decrease network energy consumption and carbon emissions attributed to it. For instance, an edge data center in the same metro may traverse a handful of L2/L3 hops. In contrast, a remote cloud region 1,000–2,000 km away will traverse multiple autonomous systems (ASs), i.e., many L2/L3 hops and numerous transponders and inline amplifiers. Thus, even a "green" location can be offset by an energy- or carbon-intensive transit. A path-aware model is therefore required to accurately attribute carbon emissions across different locations and paths.

To estimate the carbon emissions of a video streaming session, we show the difference between two carbon emission attribution models for the network: (1) *Coarse-grain model:* We consider the model presented in [47] of 0.0523 kWh/GB, and assign the carbon intensity of the network as the average between the source and destination [27]; (2) *Fine-grain model:* our proposed model that computes the network route, detailed in Section 3.3. Figure 3 illustrates the drawbacks of using a coarse-grained model. This figure shows the normalized carbon emissions when streaming to a client in Perugia, Italy. We consider the closest data center, Azure Italy North, and compare it to streaming from a low-carbon location, i.e., Azure Sweden Central, as noted in [71]. As shown, when considering the coarse-grain model, streaming from Sweden results in 55% *lower* attributed carbon emissions when compared to Italy. However, when considering the fine-grain model, which accounts for the physical network and the carbon intensity along the entire path, the attributed carbon emissions are 50% *higher* than Italy. *This discrepancy highlights the importance of considering the physical route as well as the data center emissions in correctly attributing carbon emissions and understanding the true impact of locations and paths.*

3.2 Carbon Attribution Model Overview

To address these concerns, we present a fine-grained streaming carbon emissions attribution model that captures (i) the playback device, (ii) the streaming source (edge or cloud), and (iii) the physical network route between origin and destination, combined with time- and location-varying grid carbon intensity. This structure lets us evaluate how changing the serving region alters both the equipment mix and the carbon intensity of the electricity supplied to each element. The carbon emissions of video depend on the device, the content bitrate/codec (video size), and the streaming origin, which dictates the network path. The network path itself, as shown in Figure 2, comprises static elements that are always traversed (e.g., the CPE device and the Internet gateway), and dynamic components that depend on the streaming source (e.g., L3 Hops). We define the carbon emissions $\mathbf{C}^{Q}_{stream}$ attributed to a video streaming per streaming quality Q, as follows:

$$\mathbf{C}^{Q}_{stream} = C_{\mathcal{D}} + C_{\mathcal{N}} + C_{\mathcal{DC}} \tag{1}$$

where $C_{\mathcal{D}}, C_N, C_{\mathcal{DC}}$ are the carbon emissions of the user device, network, and data center.

Device Emissions: The carbon emissions of the user device $C_{\mathcal{D}}$ is computed as:

$$C_{\mathcal{D}} = \int_0^{T_{stream}} P^{Q}_{\mathcal{D}}(t) \cdot I_{\mathcal{D}}(t)\, dt \tag{2}$$

where $P^{Q}_{\mathcal{D}}(t)$ is the power consumption in kW of the streaming device, given a video quality Q, and T_{stream} is the stream length, and $I_{\mathcal{D}}(t)$ is the time varying carbon intensity of the device location.

Network Emissions: The carbon emissions of the network C_N is defined as:

$$C_{\mathcal{N}} = \int_0^{T_{stream}} \sum_{R_i \in \mathcal{R}} \delta_i \sum_{r \in R_i} P^{Q}_{r}(t) \cdot I_r(t)\, dt \tag{3}$$

where $\mathcal{R}$ is the set of network paths between a client and a server, and δ_i is the share of the traffic going through path R_i. R_i is a network path containing a set network hops that constitute the physical network route, $P^{Q}_{r}(t) : r \in R_i$ is the typical power consumption of a network hop per stream with a video quality Q, including routers, switches, and PUE overheads, as we will detail later, and $I_r(t) : r \in R$ is the time varying carbon intensity for the network hop r. We note that our analysis uses a fixed value for energy consumption (e.g., typical power consumption [18] or maximum power consumption), as network devices are not power-proportional [18, 19, 35, 67] and network device utilization is often not known in practice [19, 72].

Data Center Emissions: The carbon emissions of the edge or cloud data center $C_{\mathcal{DC}}$ is:

$$C_{\mathcal{DC}} = \int_0^{T_{stream}} P^{Q}_{\mathcal{DC}}(t) \cdot I_{\mathcal{DC}}(t)\, dt \tag{4}$$

where $P^{Q}_{\mathcal{DC}}$ is the power consumption per streaming server for a streaming quality Q, including streaming stack, and PUE, and $I_{\mathcal{DC}}$ is the carbon intensity of the data center.

Our analysis uses the above model for accurately attributing carbon emissions to a video stream. The key challenge, however, is to estimate accurately the physical network paths and the locations of intermediate infrastructure. To address this, we utilize state-of-the-art network topology and routes datasets to infer end-to-end routes and component locations, as detailed in the next section.

3.3 Estimating the Network Route

To compute the physical network route, we need to estimate the number and location of the network hops accurately. To do so, we utilize CAIDA's Internet Technology Data Kit (ITDK) [9], which is collected based on traceroute measurements and Border Gateway Protocol (BGP) routing tables. The dataset represents the *logical*

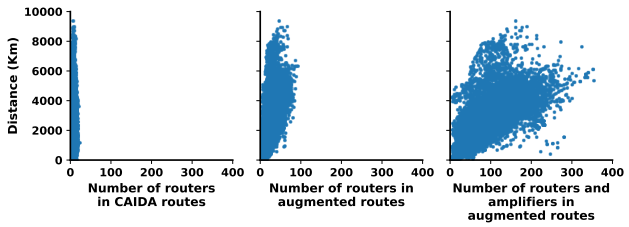

Figure 4: Number of network devices per route for a) original CAIDA dataset, b) augmented CAIDA routes and c) augmented CAIDA routes with amplifiers

structure of the internet, including router-level network topologies and router-to-autonomous system (AS) assignment. In addition, to infer the physical locations of the network route, we rely on CAIDA's geolocations [45] and iGDB geolocation datasets [5]. While this CAIDA dataset is widely used by the networking community to infer paths, we note that it does not include the existence of private networks, which are not visible in public routing tables, resulting in inferred paths that differ from actual ones [28].

We draw inspiration from recent efforts to infer the physical network path [5, 28] for our own path computations. A key difference though is that [28] estimates network paths with the lowest energy consumption, while we assume that the CDN uses standard Internet paths based on shortest path routing, which we use to infer our paths. We implement the following steps to infer the physical network route between the client and a video streaming source while considering connections across multiple clients, edge and cloud data centers, and traffic replication.

(1) **Compute the logical network path.** Given a source device location and a destination server, we utilize the CAIDA data set [9] to compute a set of network routes $\mathcal{R} : R_i \leftarrow \text{find\_routes}(\mathcal{D}, \mathcal{DC})$. To capture multi-path routing, we compute the top-k shortest paths using Dijkstra's shortest-path algorithm [30], where nodes are routers and edges represent logical links from the CAIDA dataset.
(2) **Infer the physical path.** For each network hop $r \in R_i$, we infer a latitude/longitude. Then, following the insight from iGDB [5] that long-haul links tend to follow rights-of-way (e.g., highways and railways), we transform the logical hop sequence into a piecewise physical route. Concretely, for each consecutive pair of geo-located hops $(r, r') : r, r' \in R_i$, we map these endpoints to nearby rights-of-way and compute a physical route.
(3) **Path augmentation.** We enrich the physical path with infrastructure implied by span lengths and nearby sites. Specifically, we (i) attach nearby PoPs, Internet exchanges, or metro rings when the path traverses major hubs, and (ii) insert optical repeaters/regenerators along long-haul spans based on distance thresholds typical for the medium.
(4) **ISO mapping.** Finally, we map routers, PoPs, and physical segments to electricity-grid regions to support carbon analysis. Each segment is assigned to an ISO/RTO (or equivalent) polygon. Using region labels, we later join with time-varying carbon intensity data from ElectricityMap [20] to attribute per-segment emissions and aggregate along the end-to-end route.

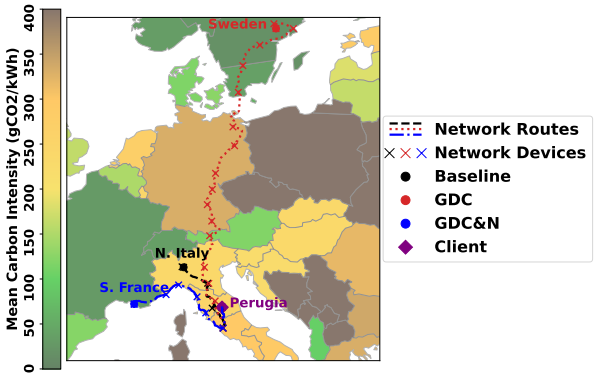

Figure 5: Evaluating Data Center selection for a client in Perugia, where the streaming from Sweden [71] yields 43% higher attributed emissions. In contrast, streaming from South France yields 38% lower emissions.

To assess the benefits of our inference pipeline, Figure 4 illustrates the relation between the haversine distance and the number of networks devices across 50000 client-server pairs. As shown, using raw CAIDA ITDK, which only includes routers that cross autonomous systems (AS), aside from distance, misrepresents the number of network devices, including routers and amplifiers, and shows no correlation between distance and the number of network devices. However, when considering the augmentation steps and including possible network hops along the path, the number of network devices becomes correlated with the distance, yielding a more plausible representation. Lastly, when incorporating the needed network amplifiers to ensure high performance and minimize the effects of noise and signal attenuation, the number of devices increases, while preserving the correlation.

3.4 Impact of Data Center Location.

To show the impact of data center location on the attributed emissions of a video streaming session, we consider three data center choices that span the spectrum from the closest data center to fully carbon–aware. Each policy operates over a candidate set of servers (data centers or CDN origins). For simplicity, we assume that the content is already replicated at these locations.

Closest Data Center (Baseline). This network–centric baseline selects the geographically (or RTT) closest server. It captures today's common practice in which path length and user experience dominate the decision rule. We use it to quantify the incremental gains of carbon-aware decision making.

Greenest Data Center (`GDC`). Motivated by spatial shifting approaches in the literature [71], this policy chooses to stream from the greenest data center location while ignoring network emissions.

Greenest Data Center & Network (`GDCN`). This policy minimizes the total attributed carbon emissions by accounting for both the data center and the network. This end-to-end view captures cases where a "green" data center is reachable only via carbon-intensive routes, or conversely, where a very clean backbone connects a modestly green server.

This set of policies allows us to isolate where the carbon benefits arise (server selection versus network route) and how much the physical path affects the scheduling decision. Figure 5 highlights the difference between these three policies and details why the greenest cloud approach may increase attributed carbon emissions, as underlined in Figure 3. As shown, when streaming from the greenest data center, using the GDC policy, the end-to-end attributed carbon emissions are higher by 43%. The reason for this is that the video stream not only traverses a significantly longer network path (8 more network hops than the closest data center) but also passes through a region with high carbon intensity. In contrast, under the greenest data center and network policy (GDCN), the video is streamed from Marseille, passing through 7 network hops and regions with low carbon intensity, resulting in 38% lower end-to-end carbon emissions.

4 Evaluation

In this section, we answer the research questions in Section 1, where we estimate the carbon emissions for different choices of data center locations and path using our fine-grained carbon attribution model. We begin by outlining our experimental setup, which includes the real-world datasets and the evaluation metrics. Next, we analyze the attributed carbon emissions when streaming from different cloud locations. We then generalize these findings to settings that involve both edge and cloud data centers. Finally, we conduct a sensitivity analysis to examine how different assumptions influence our results.

4.1 Experimental Setup

4.1.1 ***Real-world Traces.*** Our study relies on multiple real-world data sets, including:

Network Topology and Locations. As highlighted in Section 3, it utilizes the CAIDA ITDK [9] and iGDB [5] datasets to infer and augment the network route and physical path for different client and data center pairs.

Cloud Data Center Traces. Our analysis considers 28, 37, and 40 hyper-scale data centers across North America, Europe, and Asia, spanning AWS, Microsoft Azure, and Google Cloud.

Edge Data Center Traces. In addition, to emulate real-world edge deployment, we utilize Akamai CDN traces, which include the 1294, 600, and 195 locations of edge data centers across North America, Europe, and Asia, identifiable by their coordinates.

Power Measurements. Table 1 lists our power measurement data and sources for network devices. However, our model is generalizable and can use more recent measurements [19, 35]. For the server, we use the recently estimated power consumption per stream server of 1 W [70] and account for the streaming device; we consider the power consumption of a phone and use measurements from [7]. Nonetheless, in Section 4.4, we illustrate how network energy efficiency and device characteristics may affect our results.

Carbon Intensity Traces. Lastly, we use carbon intensity data from ElectricityMaps [20] for 2024. The traces provide hourly average carbon intensity information, measured in grams of carbon dioxide equivalent per kilowatt-hour ($g{\cdot}CO_2eq/kWh$), for 374 zones worldwide.

Table 1: Power measurement of different devices.

Device	Power Consumption
Streaming Server	1 W
Router	2 W/Gbps [75]
WDM switch (OXC)	0.05 W/Gbps
Transponders	0.09 W/Gbps
Amplifier	0.03 W/Gbps
Regenerator	3 W/Gbps
Network hop*	2.28 W/Gbps
PUE	1.5
Phone	~0.5 W [7]

*We assume that each hop consists of a router, two switches, and two transponders.

4.1.2 ***Experimental Scenarios.*** We analyze the differences in attributed emissions when considering the baseline (closest data center), Greenest Data Center (GDC), and Greenest Data Center & Network (GDCN). We consider a constant streaming quality Q where clients stream a 1-hour 4K or UHD video, approximately ~7GB or 15.9Mbps [53]. In our scenarios, we consider 375, 484, and 540 clients across North America, Europe, and Asia, distributed among major cities within these continents. Our evaluation considers two main video streaming scenarios: *1) Cloud Streaming*, where clients stream to cloud data centers, and *2) Edge Streaming*, where clients stream from an edge data center.

4.1.3 ***Metrics.*** Our evaluation relies on the following metrics to assess the impact of server selection and the underlying network.

Carbon Emissions ($g{\cdot}CO_2eq$). The carbon emissions attributed to a video streaming session between a server at location i and a client device at location j.

Network hops (#). The number of network hops when streaming a video from a server at location i and a client device at location j.

4.2 Cloud Streaming

We begin by evaluating and illustrating the attributed carbon emissions across different data center locations for cloud streaming, and by highlighting the emissions of the data center, network, and devices. In this case, we report the attributed emissions based on the three selection criteria detailed in Section 3.4.

4.2.1 ***Green Streaming in Action.*** Figure 6 illustrates the carbon emissions of video streaming across three locations in North America, Europe, and Asia. In these scenarios, we consider the closest data centers, the top 5 green locations, and the top 5 locations with the lowest end-to-end emissions. The figures highlight that the GDC policy, which does not consider the network, often leads to higher carbon emissions. For instance, as shown in Figure 6a, serving a client in Louisiana from a data center in Quebec — the greenest region in North America — has 34% higher attributed carbon emissions relative to the baseline. The reason is that while using the greenest cloud location reduces the server-side attributed emissions by 90.33%, the additional network carbon footprint offsets those gains. In contrast, when considering end-to-end carbon emissions, the GDCN policy yields 21% lower attributed emissions for a video streaming session, even though the server's carbon emissions are higher by 148% compared to Quebec.

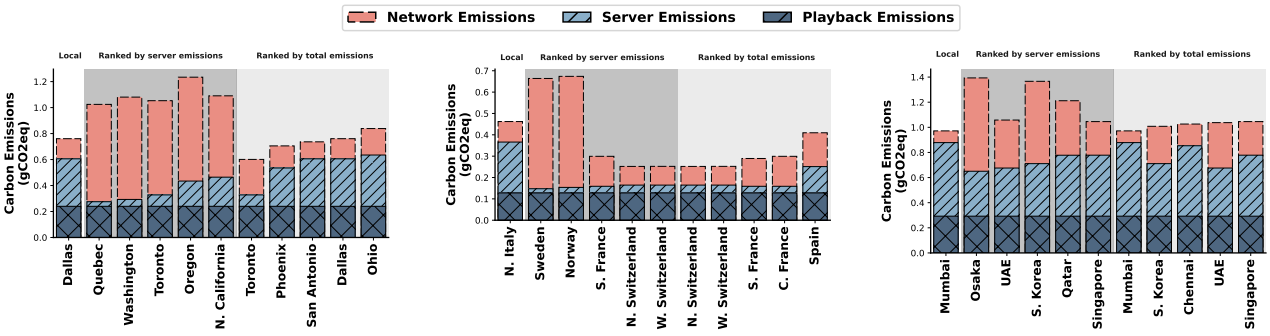

(a) Monroe, Louisiana (b) Perugia, Italy (c) Greater Noida, India

Figure 6: Carbon Emissions of streaming video when sourcing the video from the closest cloud datacenter, the top 5 greenest cloud datacenters (GDC), and the top 5 greenest cloud datacenters when considering the network (GDCN). The emissions from playback devices stay constant regardless of the streaming source.

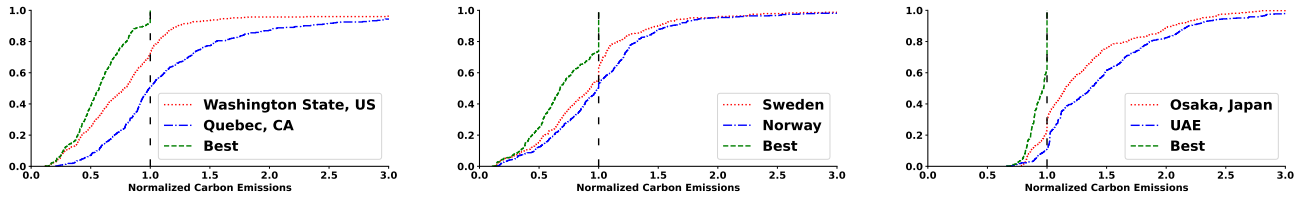

(a) North America (b) Europe (c) Asia

Figure 7: Normalized carbon emissions of video streaming across different destinations when considering the greenest data centers across locations using the GDC policy, and when considering the best selection using the GDCN policy.

Figure 6b shows another example, in Perugia, Italy, where the same issue occurs. Here, due to the availability of nearby regions with low carbon emissions, the graph shows that the GDCN approach incurs 38% lower attributed emissions when streaming from North Switzerland over streaming from the closest location. Lastly, Figure 6c highlights a slightly different situation in Greater Noida, India, where green cloud data centers are located far away, making streaming from the nearest data center a more viable option.

4.2.2 ***Aggregated Results.*** To further understand how the network paths impact stream emissions, we compute the attributed carbon emissions when considering the 375, 484, and 540 across North America, Europe, and Asia when streaming from two of the greenest data centers per continent: Quebec, CA, and Washington, US for North America, Sweden and Norway for Europe, and Osaka, Japan, and UAE in Asia. Figure 7 shows a CDF of the normalized carbon emissions considering the greenest locations (i.e., using the GDC policy) as well as the greenest (best) choice using the GDCN policy. As shown in all cases, streaming from the greenest locations can incur higher attribued emissions in 52%, 44%, and 43% of the cases across North America, Europe, and Asia. However, as shown, by considering the greenest streaming server and network path for each client, i.e., the GDCN policy, incurs 23%, 32%, and 12% lower attributed emissions across North America, Europe, and Asia. Lastly, the figure shows that there are still many cases where streaming from the nearest data center is optimal, as in 8%, 26%, and 38% of cases across North America, Europe, and Asia, respectively.

4.2.3 ***Impact of Temporal and Seasonality Variations.*** In addition to the spatial variations in carbon intensity, carbon intensity exhibits temporal variations that impact the benefits of spatial shifting [15, 71]. Figure 8 highlights situations where the best selection, considering both data center and network emissions (i.e., GDCN), changes over time and across seasons. For instance, Figure 8a illustrates the total carbon emissions attributed to streaming sessions to Monroe, Louisiana, USA, between April 1st and April 7th, 2024. We consider the selection among the five greenest locations (see Figure 6a). As shown, carbon emissions vary significantly between regions across time, with up to 2.11×, making it often inadequate to select a single location. For instance, while Toronto is the best streaming location for this client on average, in many cases, streaming from Phoenix or San Antonio results in lower carbon emissions. Figure 8b highlights how the optimal selection varies among seasons and months. For instance, between February and May, the carbon intensity of Phoenix and San Antonio, along with the carbon intensity of their respective network paths, is low, while in March, streaming from Phoenix and San Antonio yields the lowest emissions in 36.4% and 14.1% of the time, respectively. Lastly, Figure 8c highlights the net change in attributed emissions across different selections, showing that a naive fixed selection can yield

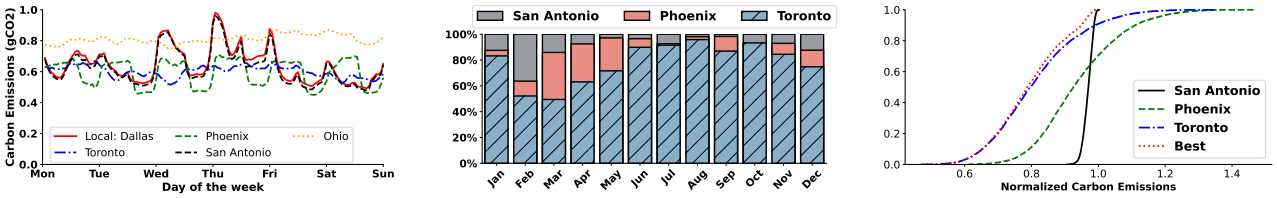

(a) Carbon Emissions (b) Optimal Cloud Selection (c) CDF

Figure 8: Impact of temporal and seasonal variations on carbon emissions in Monroe, Louisiana, USA. (a) Carbon Emissions across the first week of April, 2024. (b) Optimal selection ratios across months. (c) Normalized carbon emissions across locations and selections.

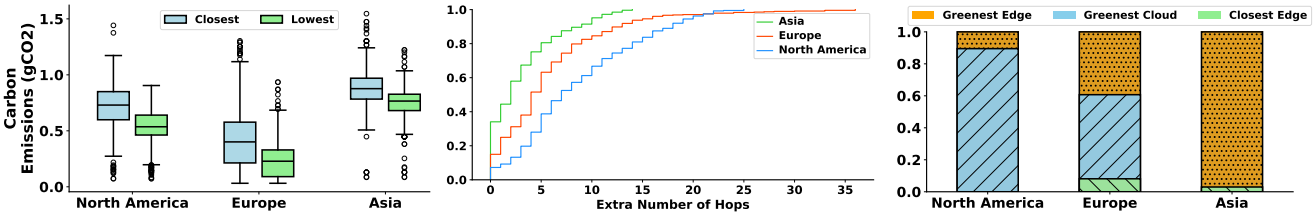

(a) Carbon Emissions (b) Extra Number of Network Hops (c) Fraction

Figure 9: Carbon Emissions and optimal selection fraction per continent when considering the carbon cost of the network.

higher emissions in ~10% of the time and can increase attributed emissions by up to 30%.

Key Takeaways. *Streaming from the greenest data center location without considering the network path emission can yield higher attributed emissions then when both the location and path are considered. Temporal and seasonal variations in carbon intensity have a significant impact on emissions, often rendering static selection ineffective.*

4.3 Edge Streaming

Edge computing brings computation closer to users (i.e., reducing network overhead). However, the constrained nature of edge resources and the increase in the number of more capable data centers make cloud streaming a reasonable alternative [4, 65]. In this section, we answer the question: *Is edge streaming a greener choice than cloud streaming?*

4.3.1 ***Edge vs Cloud Streaming.*** Figure 10 highlights the carbon emissions when considering edge and cloud resources. We consider streaming from the closest and greenest data center and network (`GDCN`) selection across the cloud and edge data center locations. The figure shows that in many cases, the closest cloud represents a reasonable choice. For instance, in North America (see Figure 10a), when considering the closest selection a video streaming session is attributed average emissions of 0.70 and 0.76 g·CO_2eq, respectively, and the greenest end-to-end emissions (i.e., using the `GDCN` policy) have average emissions of 0.54 and 0.58 g·CO_2eq, across edge and cloud data centers, respectively. Similarly, in Europe (see Figure 10b), the closest selection yield averages of carbon emissions of 0.41 and 0.58 g·CO_2eq, respectively, and the greenest end-to-end emissions (i.e., using the `GDCN` policy) have carbon emissions equal to 0.25 and 0.26 g·CO_2eq, across edge and cloud data centers, respectively. Nonetheless, in situations where cloud data centers are further away, attributed emissions are much higher, which is visible in Europe (see Figure 10b) and Asia (see Figure 10c), where the average emissions, when considering the closest edge or cloud data center, is higher by 34% and 48% across Europe and Asia, respectively.

4.3.2 ***Edge and Cloud Streaming.*** In practice, however, streaming spans edge and cloud resources. Figure 9 depicts the attributed carbon emissions under different scenarios, when considering the `GDCN` policy, while allowing selection to be distributed between cloud and edge resources. Figure 9a shows how `GDCN` can yield lower carbon emissions—by an average of 24%, 39%, and 14% across North America, Europe, and Asia. More importantly, the graph highlights how spatial variations in carbon intensity between North America and Europe can lead to lower carbon emissions, due to the presence of extremely low-carbon regions such as Sweden and France in Europe, as well as California and Ontario in North America. In contrast, in Asia, where carbon intensity is similar across regions, the effect of such policies is limited. Figure 9b shows how the `GDCN` policy increases the number of network hops, a proxy for Quality of Experience (QoE), across North America, Europe, and Asia. As shown, in 7%, 15%, 34% of the scenarios, there is no increase in the number of network hops, with an average increase of 8.56%, 5.63%, 2.97% across North America, Europe, and Asia.

To understand the role of edge and cloud data centers, Figure 9c shows how the selection differs among the three selection approaches. As shown, in North America, due to the high penetration of cloud data centers, 90% of clients benefit more from cloud streaming to lower their carbon emissions, while in only 10% of the cases, the greenest edge is a more viable option. In contrast, in Europe, due to the higher density of edge data centers, 8% and

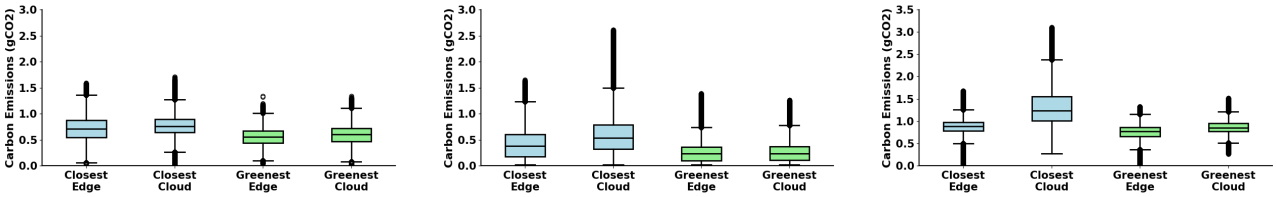

(a) North America (b) Europe (c) Asia

Figure 10: Carbon emissions when considering the closest and greenest end-to-end selection across edge and cloud data centers.

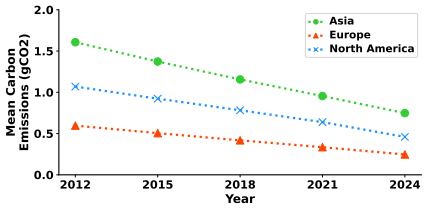

Figure 11: Effect of network energy efficiency on carbon emissions.

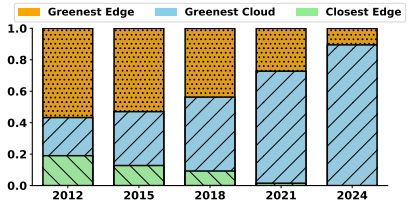

Figure 12: Effect of network energy efficiency on selection in North America.

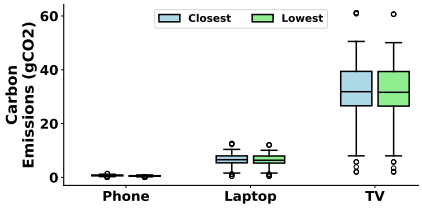

Figure 13: Carbon Emissions across devices in North America.

39% of clients prefer to stream from the closest and greenest edge resources, while 53% of clients stream from a cloud resource. Lastly, in Asia, all clients stream from edge resources.[2]

Key Takeaways. *Carbon emissions of cloud streaming are comparable with edge streaming in North America (resp. Europe), with averages of 0.70 and 0.76 (resp. 0.54 and 0.58) g·CO_2eq. Cloud streaming is a greener option for 89% of clients in North America and 52% of clients in Europe.*

4.4 Sensitivity Analysis

In this section, we evaluate the effect of different settings on the emissions of a video streaming session and the optimal selection.

4.4.1 ***Effect of Network Energy Efficiency.*** Earlier research has highlighted that the energy efficiency of network components has significantly increased over the years [6, 28], often following Koomey's law [38]. Thus, as network energy efficiency changes over time, the importance of the network in the optimization objectives may vary. To understand the effect of energy efficiency of different network infrastructures, we consider a scenario where the energy efficiency of the network varies between values reported in [76] and recent measurements as highlighted in Section 4.1.1, and interpolate values between them. To focus on network energy efficiency, we do not consider changes in server and device energy efficiency, aside from their replacement cycles [26] or changes in electricity carbon intensity across years.

Figure 11 shows the carbon emissions when considering different energy efficiencies across years using the `GDCN` policy. As shown in the figure, there is a decreasing trend across all regions. For instance, between 2024 and 2012, carbon emissions are lower by 57%, 58%, and 53% across North America, Europe, and Asia, respectively.

[2]We highlight that our cloud data set may not represent popular cloud providers in Asia.

While increases in network energy efficiency significantly affect end-to-end emissions, they also reduce their relative contribution to attributed emissions, allowing clients to stream from farther and greener locations. Figure 12 highlights the changes in the placement across the edge and cloud resources for clients in North America. As shown, with a lower network energy efficiency (higher energy intensity), edge computing becomes a more preferred choice. In fact, when considering the energy efficiency from [76] (i.e., 2012), streaming from the edge constitutes 75% of the choices. Notably, as shown, the energy efficiency of the network has improved over time (its energy intensity has decreased), making the cloud a more viable alternative. Lastly, we note that variations in the computing energy efficiency can also change the relative weights of the data center and network. However, we omitted such an evaluation for brevity.

4.4.2 ***Effect of Streaming Device.*** In an earlier experiment, we considered streaming from a phone [7]. Figure 13 shows the effect and contribution of the streaming device to the end-to-end emissions. We consider three device, including a phone [7], a laptop [33], and a TV [34] with average power consumption of 0.5W, 13.5W, 70W, respectively. For simplicity, we assumed that all devices stream the same video quality. As shown, while the absolute difference between the original and the optimized selections are similar at 0.17 g·CO_2eq, as the device energy consumption increases, the impact of the network optimization shrinks. For instance, when considering phones, data center and network optimizations yield 25% lower attributed emissions, while the same optimizations yield only 1% lower emissions of the streaming session for a TV.

4.4.3 ***Effect of Multipath Transport.*** While our earlier experiments assumed a single end-to-end path between the sender and

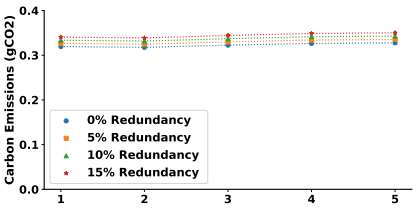

Figure 14: Effect of multi-path communication when considering top-k routes and redundancy on attributed emissions.

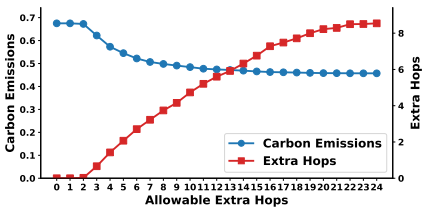

Figure 15: Effect of network energy overheads on carbon in North America.

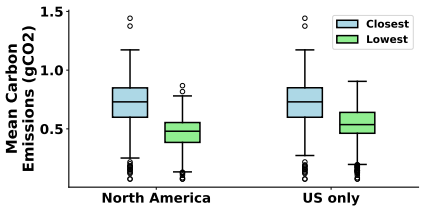

Figure 16: Difference in carbon emissions in North America when data is restricted to the US vs. all of North America.

receiver, real-world deployments increasingly rely on multipath transport protocols such as Multipath TCP (MPTCP) [57], which can distribute and even replicate traffic across these paths to improve throughput, resilience to failures, and tail latency under congestion. Figure 14 shows the effect of multi-path communication for a client in Perugia. The x-axis represents top-k routes, ranging from 1 (single path) to 5, and traffic is assigned relative to the number of hops. As shown, the end-to-end carbon emissions increase with the number of routes and replication factor by up to 9%.

4.4.4 ***Navigating the trade-off.*** As highlighted earlier, research on carbon-aware spatial shifting often involves many trade-offs [32]. One key trade-off in our case is network overhead, where streaming from farther locations may increase network hops and latency. Figure 15 highlights the carbon emissions under different numbers of extra hops across clients in North America. As shown, as we increase the number of allowable hops for streaming, carbon emissions decrease by up to 32% with an average of 8.54 extra hops. More importantly, increasing the number of hops by 6 lowers emissions by 0.15 g·CO_2eq on average.

4.4.5 ***Effect of Regulatory Constraints.*** In earlier examples, we assumed that clients can stream anywhere within their continent. In practice, however, regulatory constraints, privacy laws, and regulations may prohibit applications and data from moving outside a specific geographic jurisdiction [21]. Figure 16 shows an example, where clients in the US are allowed to stream from any edge or cloud data center in North America, or only from data centers in the US. As shown, the average carbon emissions are 24% higher when only considering locations within the US. This is because, in North America, many clients opted to stream from Quebec and Ontario, which are known to have very low carbon intensity [71].

Key Takeaways. *As the relative weights of the network and the streaming devices vary, our fine-grained model becomes increasingly critical for attributing the emissions to a video streaming session.*

5 Discussion

In this paper, we propose a fine-grained methodology to estimate carbon emissions for video streaming sessions and highlight the drawbacks of coarse-grained attribution methods under different destination selection strategies. Next, we highlight the generalization benefits of our approach as well as its limitations.

Generalization. Although our proposed model and evaluation focused on reporting carbon emissions in video streaming, our findings are generally applicable to a wide range of scenarios. For instance, our study highlights critical drawbacks of obvious decisions, such as streaming from the nearest edge or the greenest cloud. Applications such as video analytics [80] and data ingestion [49] can significantly benefit from such insights in carbon reporting. More importantly, our approach highlights the significance of cloud datacenters, which contrasts with the common reliance on edge data centers [4, 65].

Limitations. Our analysis presents a method for estimating carbon emissions from video streaming sessions under different scenarios. However, our analysis has the following limitations: First, our analysis does not account for capacity constraints in edge and cloud data centers or the network. However, in practice, such capacity constraints exist and may affect the client's ability to stream from different locations, as well as the carbon emissions. Analyzing the effects of streaming from different sources on link utilization patterns, addressing how different streaming strategies might skew network load distribution, and their impact on system-wide energy consumption and carbon emissions are left for future work. Second, our analysis focuses on the carbon emissions associated with video streaming and on how they differ when streaming from different locations. However, a key aspect of edge data centers is the delivery bandwidth cost. Such a view highlights the three-way trade-off between cost, energy, and emissions, and presents an important future research direction. Third, our assumption considers fixed power consumption per device. However, while network routes and energy consumption may vary [51, 72, 76], networks do not expose such information. Lastly, in this paper, we focus on carbon emissions *attributed* to streaming sessions; however, this does not necessarily translate to an actual reduction in carbon emissions, as highlighted in earlier research [18, 36, 51, 59, 72].

6 Related Work

Earlier work has focused on optimizing data centers and the energy efficiency and carbon footprint of internet services.

Green Computing. Green computing solutions have long focused on optimizing energy efficiency, the integration of renewable energy in computing services, and data centers [40–43, 48, 62, 78]. For instance, [48] explored power management approaches to save energy while adhering to Service Level Availability (SLA) constraints. In contrast, [42] has explored how renewable-aware load shifting can reduce carbon emissions and operational costs.

However, more recent research has highlighted the diversity in energy's carbon intensity across locations and times [15, 23, 27, 50, 71, 78]. For example, in [15], the authors utilize spatial load shifting

to decrease the carbon emissions of AI workloads while respecting SLA constraints. Similarly, [50] considers large-scale CDN services and implements a carbon-aware spatial shifting approach to reduce carbon emissions. However, these approaches often focus on compute-intensive workloads or ignore the network overheads. In contrast, we focus on video streaming services, a computationally and bandwidth-intensive task, and propose approaches to estimate and optimize carbon emissions for a video streaming session.

Green Networking. Green networking solutions have long focused on analyzing and optimizing the energy consumption of networks [6, 10, 18, 19, 25, 35, 36, 51, 52, 59, 63, 64, 66, 76, 77]. For instance, [6] proposes an estimate for energy intensity and presents some projections based on increases in energy efficiency of hardware, [28] proposes a model to estimate the carbon emissions of data transfer. In contrast, [19] implements a measurement study and proposes a model for the power consumption of a network device. However, in contrast to these approaches, we consider end-to-end emissions attributed to a video streaming across the device, the network, and edge and cloud data centers, and do not consider any changes to the routing protocol or knowledge of network utilization.

Earlier research has also proposed an approach that optimizes the carbon efficiency of network activity [18, 24, 28, 54, 58, 72, 81]. For instance, authors have proposed carbon-aware routing approaches [18, 54, 72], which complement our proposed approach. In addition, researchers proposed carbon-aware scheduling approaches that temporally shift data transfer activities by considering the carbon intensity of the network [24, 28, 58]. In contrast to these approaches, we focus on video streaming services, a task that cannot be temporally shifted.

Green Streaming. Earlier work on green streaming has focused on optimizing the energy efficiency [2, 3, 44, 60, 74, 79]. For instance, [44, 74] propose ABR approaches that minimize the energy consumption of video streaming. In contrast to these approaches, we focus on optimizing the carbon emissions of video streaming, which complements approaches that optimize the energy efficiency of video streaming. In addition, many papers have focused on analyzing the carbon emissions of streaming [6, 16, 37, 47, 51, 70]. For instance, [47] conducts a global analysis of streaming, considers the device, network, and edge, and highlights how the number of users affects the energy/stream. In contrast to these approaches, we rely on a fine-grained model that is essential for reporting emissions across different streaming sources.

7 Conclusion

Video streaming is a significant contributor to the internet's carbon emissions. To understand and account for the carbon emissions of video streaming, we introduced a fine-grained, end-to-end model that estimates the carbon emissions attributed to a video streaming session from cloud and edge data centers. Our model-driven analysis showed that ignoring network emissions can yield higher attributed emissions for 52%, 44%, and 43% of the clients in North America, Europe, and Asia. Our results also demonstrate that the nearest edge site is frequently suboptimal, in more than 77% of the cases globally. Building on this model, our future work will focus on designing CDN-scale streaming algorithms that optimize carbon emissions from video streaming, while accounting for capacity constraints, green routing approaches, and bandwidth costs to balance the trade-off among cost, energy, and carbon.

Acknowledgments

We thank the anonymous e-Energy reviewers for their valuable insight and feedback. This research is supported by National Science Foundation (NSF) grants 2213636, 2105494, 2211302, 2211888, 2325956, 23091241, the U.S. Department of Energy Award DE-EE0010143, US Army contract W911NF-17-2-0196, and support from VMware. This work used Amazon Web Services through the CloudBank, which is supported by NSF grant 19250001.

References

[1] Micah Adler, Ramesh K. Sitaraman, and Harish Venkataramani. 2011. Algorithms for optimizing the bandwidth cost of content delivery. *Computer Networks* 55, 18 (2011), 4007–4020. doi:10.1016/j.comnet.2011.07.015 Internet-based Content Delivery.

[2] Samira Afzal, Narges Mehran, Zoha Azimi Ourimi, Farzad Tashtarian, Hadi Amirpour, R.-C. Prodan, and Christian Timmerer. 2024. A Survey on Energy Consumption and Environmental Impact of Video Streaming. *ArXiv* abs/2401.09854 (2024). https://api.semanticscholar.org/CorpusID:267035135

[3] Samira Afzal, Radu Prodan, and Christian Timmerer. 2024. Green Video Streaming: Challenges and Opportunities. *SIGMultimedia Rec.* 15, 1, Article 3 (Oct. 2024), 1 pages. doi:10.1145/3699843.3699846

[4] Ahmed Ali-Eldin, Bin Wang, and Prashant Shenoy. 2021. The Hidden Cost of the Edge: A Performance Comparison of Edge and Cloud Latencies. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (St. Louis, Missouri) *(SC '21)*. Article 23, 12 pages. doi:10.1145/3458817.3476142

[5] Scott Anderson, Loqman Salamatian, Zachary S. Bischof, Alberto Dainotti, and Paul Barford. 2022. iGDB: connecting the physical and logical layers of the internet. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) *(IMC '22)*. Association for Computing Machinery, New York, NY, USA, 433–448. doi:10.1145/3517745.3561443

[6] Joshua Aslan, Kieren Mayers, Jonathan G. Koomey, and Chris France. 2018. Electricity Intensity of Internet Data Transmission: Untangling the Estimates. *Journal of Industrial Ecology* 22, 4 (2018), 785–798. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/jiec.12630 doi:10.1111/jiec.12630

[7] Youssef Badra and Razvan Stanica. 2025. Experimental Analysis of Energy Consumption in Video Streaming Services. In *2025 IEEE 26th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. 11–20. doi:10.1109/WoWMoM65615.2025.00016

[8] Matthew Brander. 2022. The most important GHG accounting concept you may not have heard of: the attributional-consequential distinction. *Carbon Management* 13, 1 (2022), 337–339. doi:10.1080/17583004.2022.2088402

[9] CAIDA (Center for Applied Internet Data Analysis). 2024. CAIDA Macroscopic Internet Topology Data Kit (ITDK), Release 2024-08. https://www.caida.org/catalog/datasets/internet-topology-data-kit/release-2024-08/.

[10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. 2008. Power Awareness in Network Design and Routing. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*. 457–465. doi:10.1109/INFOCOM.2008.93

[11] Mohak Chadha, Thandayuthapani Subramanian, Eishi Arima, Michael Gerndt, Martin Schulz, and Osama Abboud. 2023. GreenCourier: Carbon-Aware Scheduling for Serverless Functions. In *Proceedings of the 9th International Workshop on Serverless Computing* (Bologna, Italy) *(WoSC '23)*. Association for Computing Machinery, New York, NY, USA, 18–23. doi:10.1145/3631295.3631396

[12] Tianyu Chen, Yiheng Lin, Nicolas Christianson, Zahaib Akhtar, Sharath Dharmaji, Mohammad Hajiesmaili, Adam Wierman, and Ramesh K. Sitaraman. 2024. SODA: An Adaptive Bitrate Controller for Consistent High-Quality Video Streaming. In *Proceedings of the ACM SIGCOMM 2024 Conference* (Sydney, NSW, Australia) *(ACM SIGCOMM '24)*. Association for Computing Machinery, New York, NY, USA, 613–644. doi:10.1145/3651890.3672260

[13] Alexander Clemm, Dirk Kutscher, Michael Welzl, Cedric Westphal, Noa Zilberman, and Simone Ferlin-Reiter. 2025. Greening Networking: Toward a Net Zero Internet (Dagstuhl Seminar 24402). *Dagstuhl Reports* 14, 9 (2025), 167–192. doi:10.4230/DagRep.14.9.167

[14] Ethan Cramer-Flood. 2024. US Time Spent With Media Forecast 2024. https://www.emarketer.com/content/us-time-spent-with-media-forecast-2024. eMarketer, March 21, 2024.

[15] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCario, and Will Buchanan. 2022. Measuring the Carbon Intensity of AI in Cloud Instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. 1877–1894. doi:10.1145/3531146.3533234

[16] Hugues Efoui-Hess and The Shift Project. 2019. *Climate crisis: The unsustainable use of online video*. Technical Report. The Shift Project. https://theshiftproject.org/en/publications/unsustainable-use-online-video/ Accessed: 2025-09-17.

[17] David Ekchajzer. 2024. *Counting Environmental Impacts Using Attributional and Consequential Approaches*. Hubblo. https://hubblo.org/blog/attributional-vs-consequential/ Blog post.

[18] Sawsan El-Zahr, Paul Gunning, and Noa Zilberman. 2023. Exploring the Benefits of Carbon-Aware Routing. *Proc. ACM Netw.* 1, CoNEXT3, Article 20 (Nov. 2023), 24 pages. doi:10.1145/3629165

[19] Sawsan El Zahr and Noa Zilberman. 2025. From Measurement to Emissions: Assessing the Carbon Footprint of Traffic Flows. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* (2025).

[20] Electricity Maps. 2024. Electricity Maps. https://www.electricitymap.org/map.

[21] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. https://data.europa.eu/eli/reg/2016/679/oj.

[22] Reza Farahani, Zoha Azimi, Christian Timmerer, and Radu Prodan. 2024. Towards AI-Assisted Sustainable Adaptive Video Streaming Systems: Tutorial and Survey. arXiv:2406.02302 [cs.MM] https://arxiv.org/abs/2406.02302

[23] Peter Xiang Gao, Andrew R. Curtis, Bernard Wong, and Srinivasan Keshav. 2012. It's Not Easy Being Green. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication* (Helsinki, Finland) *(SIGCOMM '12)*. 211–222. doi:10.1145/2342356.2342398

[24] Jacob Goldverg, Hasibul Jamil, Elvis Rodrigues, and Tevfik Kosar. 2025. Towards Carbon-Aware Data Transfers. *IEEE Internet Computing* 01 (2025), 1–7. doi:10.1109/MIC.2025.3565187

[25] Eduard Goma, Marco Canini, Alberto Lopez Toledo, Nikolaos Laoutaris, Dejan Kostić, Pablo Rodriguez, Rade Stanojević, and Pablo Yagüe Valentin. 2011. Insomnia in the access: or how to curb access network related energy consumption. In *Proceedings of the ACM SIGCOMM 2011 Conference* (Toronto, Ontario, Canada) *(SIGCOMM '11)*. 338–349. doi:10.1145/2018436.2018475

[26] Albert Greenberg, James Hamilton, David A. Maltz, and Parveen Patel. 2009. The cost of a cloud: research problems in data center networks. *SIGCOMM Comput. Commun. Rev.* 39, 1 (Dec. 2009), 68–73. doi:10.1145/1496091.1496103

[27] Viktor Urban Gsteiger, Pin Hong (Daniel) Long, Yiran (Jerry) Sun, Parshan Javanrood, and Mohammad Shahrad. 2024. Caribou: Fine-Grained Geospatial Shifting of Serverless Applications for Sustainability. In *Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles* (Austin, TX, USA) *(SOSP '24)*. 403–420. doi:10.1145/3694715.3695954

[28] Yibo Guo, Amanda Tomlinson, Runlong Su, and George Porter. 2025. The Effect of the Network in Cutting Carbon for Geo-shifted Workloads. arXiv:2504.14022 [cs.NI] https://arxiv.org/abs/2504.14022

[29] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: Designing Sustainable Computer Systems With An Architectural Carbon Modeling Tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture* (New York, New York) *(ISCA '22)*. Association for Computing Machinery, New York, NY, USA, 784–799. doi:10.1145/3470496.3527408

[30] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*. 11–15.

[31] Walid A. Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3, Article 57 (December 2023), 28 pages. doi:10.1145/3626788

[32] Walid A. Hanafy, Qianlin Liang, Noman Bashir, Abel Souza, David Irwin, and Prashant Shenoy. 2024. Going Green for Less Green: Optimizing the Cost of Reducing Cloud Carbon Emissions. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS'24)*. 479—-496. doi:10.1145/3620666.3651374

[33] Christian Herglotz, Werner Robitza, Matthias Kränzler, Andre Kaup, and Alexander Raake. 2022. Modeling of Energy Consumption and Streaming Video QoE using a Crowdsourcing Dataset. In *2022 14th International Conference on Quality of Multimedia Experience (QoMEX)*. 1–6. doi:10.1109/QoMEX55416.2022.9900886

[34] Jackery. 2025. How Many Watts Does a TV Use: 24, 32, 50, 55, 65 Inch TV and More [With a Data Table]. https://www.jackery.com/blogs/knowledge/how-many-watts-does-a-tv-use Updated November 21, 2025.

[35] Romain Jacob, Lukas Röllin, Jackie Lim, Jonathan Chung, Maurice Béhanzin, Weiran Wang, Andreas Hunziker, Theodor Moroianu, Seyedali Tabaeiaghdaei, Adrian Perrig, and Laurent Vanbever. 2025. Fantastic Joules and Where to Find Them. Modeling and Optimizing Router Energy Demand. In *Proceedings of the 2025 ACM Internet Measurement Conference* (USA) *(IMC '25)*. 48–62. doi:10.1145/3730567.3732920

[36] Romain Jacob and Laurent Vanbever. 2023. The Internet of Tomorrow Must Sleep More and Grow Old. *SIGENERGY Energy Inform. Rev.* 3, 3 (Oct. 2023), 27–32. doi:10.1145/3630614.3630620

[37] George Kamiya. 2020. The carbon footprint of streaming video: fact-checking the headlines. *International Energy Agency* (11 December 2020). https://www.iea.org/commentaries/the-carbon-footprint-of-streaming-video-fact-checking-the-headlines

[38] Jonathan Koomey and Samuel Naffziger. 2015. Moore's Law might be slowing down, but not energy efficiency. *IEEE spectrum* 52, 4 (2015), 35.

[39] Baolin Li, Rohan Basu Roy, Daniel Wang, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. 2023. Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Denver, CO, USA) *(SC '23)*. Article 19, 15 pages. doi:10.1145/3581784.3607035

[40] Julia Lindberg, Yasmine Abdennadher, Jiaqi Chen, Bernard C. Lesieutre, and Line Roald. 2021. A Guide to Reducing Carbon Emissions through Data Center Geographical Load Shifting. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems* (Virtual Event, Italy) *(e-Energy '21)*. Association for Computing Machinery, New York, NY, USA, 430–436. doi:10.1145/3447555.3466582

[41] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. 2012. Renewable and cooling aware workload management for sustainable data centers. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems* (London, England, UK) *(SIGMETRICS '12)*. 175–186. doi:10.1145/2254756.2254779

[42] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H. Low, and Lachlan L.H. Andrew. 2011. Geographical Load Balancing with Renewables. *SIGMETRICS Perform. Eval. Rev.* 39, 3 (dec 2011), 62–66. doi:10.1145/2160803.2160862

[43] Zhenhua Liu, Adam Wierman, Yuan Chen, Benjamin Razon, and Niangjun Chen. 2013. Data center demand response: Avoiding the coincident peak via workload shifting and local generation. *Performance Evaluation* 70, 10 (2013), 770–791. doi:10.1016/j.peva.2013.08.014 Proceedings of IFIP Performance 2013 Conference.

[44] Daniele Lorenzi, Minh Nguyen, Farzad Tashtarian, and Christian Timmerer. 2024. E-WISH: An Energy-aware ABR Algorithm For Green HTTP Adaptive Video Streaming. In *Proceedings of the 3rd Mile-High Video Conference* (Denver, CO, USA) *(MHV '24)*. Association for Computing Machinery, New York, NY, USA, 28–33. doi:10.1145/3638036.3640802

[45] Matthew Luckie, Bradley Huffaker, Alexander Marder, Zachary Bischof, Marianne Fletcher, and K Claffy. 2021. Learning to extract geographic information from internet router hostnames. In *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies* (Virtual Event, Germany) *(CoNEXT '21)*. Association for Computing Machinery, New York, NY, USA, 440–453. doi:10.1145/3485983.3494869

[46] Diptyaroop Maji, Ben Pfaff, Vipin P R, Rajagopal Sreenivasan, Victor Firoiu, Sreeram Iyer, Colleen Josephson, Zhelong Pan, and Ramesh K Sitaraman. 2023. Bringing Carbon Awareness to Multi-cloud Application Delivery. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems* (Boston, MA, USA) *(HotCarbon '23)*. Article 6, 6 pages. doi:10.1145/3604930.3605711

[47] Stephen Makonin, Laura U Marks, Radek Przedpełski, Alejandro Rodriguez-Silva, and Ramy ElMallah. 2022. Calculating the carbon footprint of streaming media: Beyond the myth of efficiency. In *Eighth workshop on computing within limits 2022*. LIMITS.

[48] Vimal Mathew, Ramesh K. Sitaraman, and Prashant Shenoy. 2012. Energy-aware Load Balancing in Content Delivery Networks. In *Proceedings IEEE INFOCOM*. 954–962. doi:10.1109/INFCOM.2012.6195846

[49] John Meehan, Cansu Aslantas, Stan Zdonik, Nesime Tatbul, and Jiang Du. 2017. Data Ingestion for the Connected World. In *Cidr*, Vol. 17. 8–11.

[50] Jorge Murillo, Walid A. Hanafy, David Irwin, Ramesh Sitaraman, and Prashant Shenoy. 2024. CDN-Shifter: Leveraging Spatial Workload Shifting to Decarbonize Content Delivery Networks. In *Proceedings of the 2024 ACM Symposium on Cloud Computing* (Redmond, WA, USA) *(SoCC '24)*. Association for Computing Machinery, New York, NY, USA, 505–521. doi:10.1145/3698038.3698516

[51] David Mytton, Dag Lundén, and Jens Malmodin. 2024. Network energy use not directly proportional to data volume: The power model approach for more reliable network energy consumption calculations. *Journal of Industrial Ecology* 28, 4 (2024), 966–980. doi:10.1111/jiec.13512

[52] Sergiu Nedevschi, Lucian Popa, Gianluca Iannaccone, Sylvia Ratnasamy, and David Wetherall. 2008. Reducing network energy consumption via sleeping and rate-adaptation. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation* (San Francisco, California) *(NSDI'08)*. 323–336.

[53] Netflix. [n. d.]. How to control how much data Netflix uses? https://help.netflix.com/en/node/87 Accessed: 2025-06-03.

[54] Daniel Otten, Alexander Brundiers, Timmy Schüller, and Nils Aschenbruck. 2023. Green Segment Routing for Improved Sustainability of Backbone Networks. In *2023 IEEE 48th Conference on Local Computer Networks (LCN)*. 1–9. doi:10.1109/LCN58197.2023.10223317

[55] Roger Pantos and William May. 2017. HTTP Live Streaming. RFC 8216. doi:10.17487/RFC8216

[56] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. 2009. Cutting The Electric Bill for Internet-Scale Systems. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication* (Barcelona, Spain) *(SIGCOMM '09)*. 123–134. doi:10.1145/1592568.1592584

[57] Costin Raiciu, Christoph Paasch, Sebastien Barre, Alan Ford, Michio Honda, Fabien Duchene, Olivier Bonaventure, and Mark Handley. 2012. How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. USENIX Association, San Jose, CA, 399–412. https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/raiciu

[58] Elvis Rodrigues, Jacob Goldverg, and Tevfik Kosar. 2025. Carbon-Aware Temporal Data Transfer Scheduling Across Cloud Datacenters. arXiv:2506.04117 [cs.DC] https://arxiv.org/abs/2506.04117

[59] Lukas Röllin, ROMAIN JACOB, and Laurent Vanbever. 2025. A Sleep Study for ISP Networks: Evaluating Link Sleeping on Real World Data. *SIGENERGY Energy Inform. Rev.* 4, 5 (April 2025), 51–58. doi:10.1145/3727200.3727209

[60] Mohammadhassan Safavi, Saeed Bastani, Zhi Zhang, Martti Forsell, Olli Mämmelä, and Björn Landfeldt. 2016. A study on energy used to deliver H.264/AVC and H.265/HEVC video content. In *2016 IEEE 21st International Workshop on Computer Aided Modelling and Design of Communication Links and Networks (CAMAD)*. 170–176. doi:10.1109/CAMAD.2016.7790353

[61] Sandvine. 2024. *Global Internet Phenomena Report 2024*. Technical Report. Sandvine. https://www.applogicnetworks.com/hubfs/Sandvine_Redesign_2019/Downloads/2024/GIPR/GIPR%202024.pdf Accessed: 2025-06-03.

[62] Nishanth Sastry and Jon Crowcroft. 2010. SpinThrift: saving energy in viral workloads. In *Proceedings of the First ACM SIGCOMM Workshop on Green Networking* (New Delhi, India) *(Green Networking '10)*. 69–76. doi:10.1145/1851290.1851305

[63] Daniel Schien, Vlad C. Coroama, Lorenz M. Hilty, and Chris Preist. 2015. The Energy Intensity of the Internet: Edge and Core Networks. In *ICT Innovations for Sustainability*. Springer International Publishing, Cham, 157–170.

[64] Anders SG Andrae. 2020. New perspectives on internet electricity use in 2030. *Engineering and Applied Science Letter* 3, 2 (2020), 19–31.

[65] Pragya Sharma, Hang Qiu, and Mani Srivastava. 2025. Cloud Is Closer Than It Appears: Revisiting the Tradeoffs of Distributed Real-Time Inference. In *Proceedings of the IEEE ICCCN*.

[66] Arman Shehabi, Sarah Josephine Smith, Dale A. Sartor, Richard E. Brown, Magnus Herrlin, Jonathan G. Koomey, Eric R. Masanet, Nathanial Horner, Ines Lima Azevedo, and William Linter. 2016. *United States Data Center Energy Usage Report*. Technical Report LBNL-1005775. Lawrence Berkeley National Lab (LBL). https://www.osti.gov/biblio/1372902

[67] V. Sivaraman, P. Reviriego, Z. Zhao, A. Sánchez-Macián, A. Vishwanath, J.A. Maestro, and C. Russell. 2014. An experimental power profile of Energy Efficient Ethernet switches. *Computer Communications* 50 (2014), 110–118. doi:10.1016/j.comcom.2014.02.019 Green Networking.

[68] Iraj Sodagar. 2011. The MPEG-DASH Standard for Multimedia Streaming Over the Internet. *IEEE MultiMedia* 18, 4 (2011), 62–67. doi:10.1109/MMUL.2011.71

[69] Kevin Spiteri, Rahul Urgaonkar, and Ramesh K. Sitaraman. 2020. BOLA: Near-Optimal Bitrate Adaptation for Online Videos. *IEEE/ACM Transactions on Networking* 28, 4 (2020), 1698–1711. doi:10.1109/TNET.2020.2996964

[70] Andie Stephens, Chloe Tremlett-Williams, Liam Fitzpatrick, Luca Acerini, Matt Anderson, and Noor Crabbendam. 2021. *Carbon Impact of Video Streaming*. Technical Report. The Carbon Trust. https://www.carbontrust.com/our-work-and-impact/guides-reports-and-tools/carbon-impact-of-video-streaming Accessed: 2025-06-03.

[71] Thanathorn Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2024. On the Limitations of Carbon-Aware Temporal and Spatial Workload Shifting in the Cloud. In *Proceedings of the Nineteenth European Conference on Computer Systems* (Athens, Greece) *(EuroSys '24)*. 924–941. doi:10.1145/3627703.3650079

[72] Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon-Aware Global Routing in Path-Aware Networks. In *Proceedings of the 14th ACM International Conference on Future Energy Systems* (Orlando, FL, USA) *(e-Energy '23)*. Association for Computing Machinery, New York, NY, USA, 144–158. doi:10.1145/3575813.3595192

[73] Christian Timmerer, Hadi Amirpour, Farzad Tashtarian, Samira Afzal, Amr Rizk, Michael Zink, and Hermann Hellwagner. 2025. HTTP Adaptive Streaming: A Review on Current Advances and Future Challenges. *ACM Trans. Multimedia Comput. Commun. Appl.* 21, 7, Article 198 (July 2025), 27 pages. doi:10.1145/3736306

[74] Bekir Oguzhan Turkkan, Ting Dai, Adithya Raman, Tevfik Kosar, Changyou Chen, Muhammed Fatih Bulut, Jaroslaw Zola, and Daby Sow. 2022. GreenABR: energy-aware adaptive bitrate streaming with deep reinforcement learning. In *Proceedings of the 13th ACM Multimedia Systems Conference* (Athlone, Ireland) *(MMSys '22)*. Association for Computing Machinery, New York, NY, USA, 150–163. doi:10.1145/3524273.3528188

[75] Valiant Communications Ltd. 2019. VCL-MX-5040-R MPLS Router Data Sheet. https://www.valiantcom.com/routers/data-sheets/5040-mpls-router.pdf. Revision 9, April 17 2019.

[76] Ward Van Heddeghem, Filip Idzikowski, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester. 2012. Power consumption modeling in optical multilayer networks. *Photonic Network Communications* 24, 2 (2012), 86–102. doi:10.1007/s11107-011-0370-7

[77] Nedeljko Vasić, Prateek Bhurat, Dejan Novaković, Marco Canini, Satyam Shekhar, and Dejan Kostić. 2011. Identifying and using energy-critical paths. In *Proceedings of the Seventh COnference on Emerging Networking EXperiments and Technologies* (Tokyo, Japan) *(CoNEXT '11)*. Association for Computing Machinery, New York, NY, USA, Article 18, 12 pages. doi:10.1145/2079296.2079314

[78] Li Wu, Walid A. Hanafy, Abel Souza, Khai Nguyen, Jan Harkes, David Irwin, Mahadev Satyanarayanan, and Prashant Shenoy. 2025. CarbonEdge: Leveraging Mesoscale Spatial Carbon-Intensity Variations for Low Carbon Edge Computing. doi:10.1145/3731545.3731576

[79] Chaoqun Yue, Subhabrata Sen, Bing Wang, Yanyuan Qin, and Feng Qian. 2020. Energy considerations for ABR video streaming to smartphones: measurements, models and insights. In *Proceedings of the 11th ACM Multimedia Systems Conference* (Istanbul, Turkey) *(MMSys '20)*. Association for Computing Machinery, New York, NY, USA, 153–165. doi:10.1145/3339825.3391867

[80] Haoyu Zhang, Ganesh Ananthanarayanan, Peter Bodik, Matthai Philipose, Paramvir Bahl, and Michael J. Freedman. 2017. Live Video Analytics at Scale with Approximation and Delay-Tolerance. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA, 377–392. https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/zhang

[81] Noa Zilberman, Eve M. Schooler, Uri Cummings, Rajit Manohar, Dawn Nafus, Robert Soulé, and Rick Taylor. 2023. Toward Carbon-Aware Networking. *SIGENERGY Energy Inform. Rev.* 3, 3 (Oct. 2023), 15–20. doi:10.1145/3630614.3630618