Carbon-Intelligent Content Scheduling in CDNs

Sawsan El-Zahr
University of Oxford
Oxford, UK

William Nathan
University of Oxford
Oxford, UK

Noa Zilberman
University of Oxford
Oxford, UK

Abstract

Content delivery networks scaling with users' demand can meet quality of service needs by placing caches close to users. However, distributing content to caches is not for free. Not only it has a monetary cost, but also an environmental footprint. This work proposes a carbon-intelligent content delivery algorithm that accounts for cost and traffic demands as well as carbon emissions. Given the spatial and temporal dynamism of carbon intensity, this work aligns data delivery to caches with low carbon intensity periods and low distribution costs. The solution is evaluated on the Netflix Open Connect network, showing operational carbon savings between 13% and 64% while conserving other constraints.

1 Introduction

Content delivery networks (CDNs) are a major driver of Internet traffic. In 2023, streaming accounted for 65% of downstream traffic, with Netflix alone responsible for 13.7% of all Internet traffic [34].

CDNs deploy caches closer to users than origin servers [23]. This achieves several objectives: (1) lower latency and improved user experience [9], (2) reduced traffic on the backbone and load on the origin servers [3], and (3) lower ISP transit costs by serving traffic locally [3].

Content in edge caches should be up-to-date. Content delivery is often not uniform across the globe. Different countries and regions have different popularity figures, making content delivery to edge caches more challenging.

Currently, new content is typically pushed to edge caches during off-peak hours [27]. This option minimizes bandwidth cost in the backbone network [13]. Moreover, during off-peak hours, fill traffic does not compete with other traffic, meaning less congestion and fewer re-transmissions. And finally, this separates fill and stream times for caches, ensuring appliances are not overwhelmed. This presents a reliable and cost-efficient approach. However, sustainability is currently not a critical component in the content delivery schedule.

Streaming and content delivery companies such as Netflix, Disney+ and Akamai are committed to the 2030 carbon reduction goals [2, 16, 26]. Their carbon footprint comprises cloud processing, CDN storage, transmission networks, and end-user devices, with each component contributing differently depending on the company's scope [39].

CDNs have infrastructure spread across the globe. Caches located in different places will consume energy from a different mix of sources, as energy generation varies between geographical regions. Some locations have renewable energy in abundance, while some rely mostly on fossil fuel. It is not enough to know the energy consumption of every appliance, but one should also account for the carbon intensity of the region. The carbon emissions of an appliance are the integration over time of its energy consumption multiplied by instantaneous carbon intensity.

In this paper, a carbon-intelligent scheduling scheme is proposed for content delivery to caches that adds carbon-related information to the existing design principles of the CDN. Carbon-intelligent content delivery scheduling could assist streaming companies in their sustainability goals, providing tangible value as well as social benefit.

In summary, this paper makes the following contributions:

- We propose a new carbon-intelligent scheduling scheme for content delivery to caches.
- We demonstrate the solution on the Netflix Open Connect network and show that it can lead to carbon savings without compromising on existing constraints.
- We suggest a static carbon-efficient content delivery schedule that saves computations' emissions.
- We discuss further considerations in the scheduling of carbon-efficient content distribution.

2 Content Delivery

From a high level perspective, the objective of a CDN is to serve users' demand as fast as possible. Users want up-to-date content without waiting. Therefore, planning ahead for content delivery is a key step to achieve high performance.

CDNs use tiered caches [10, 29]. This can be largely divided into three levels: (1) origin servers where all content is stored and is first available once produced, (2) high-level caches which are almost a copy of origin servers and are spread over a large geographical area. Often, these are co-located with Internet Exchange Points (IXP) for better reachability to other networks [12, 31], (3) and finally low-level caches that have a subset of the content but are placed closer to users. These smaller caches can be embedded inside ISPs and are customized based on local content popularity. Consequently, a content delivery schedule should be from origin servers to high-level caches and then to low-level caches.

Filling caches with content should not interrupt or impact the main purpose of caches, which is serving content. During peak traffic times, caches would be downstreaming content at a high utilization. If filling caches with content coincides with offloading of content, then performance is affected. Hence, off-peak hours are the target time-slots for content filling. In practice, the filling window starts around 2AM [27].

2.1 Content Delivery Sustainability

Caches consume energy when sending and receiving content. Regardless of the delivery schedule, for the same content, a given cache will consume the same amount of energy. However, to derive the carbon footprint of this operation, the carbon intensity of the energy consumed should be accounted for. Carbon intensity is the measure of grams of CO_2 emitted to produce 1 kWh of energy [44]. This metric varies based on the source of the energy generated, and when consumed from a national grid is commonly a mix of energy from multiple sources. Therefore, the location of the cache that is consuming energy matters to derive the carbon footprint. If local generation of renewable energy is used, this should also be taken into account.

Carbon intensity varies significantly over hours, days and seasons. Figure 1 shows how the national hourly average carbon intensity varies in Italy, France, UK over a week [19]. For the US, it shows how the regional carbon intensity of New York varies over time [19]. Some regions are almost always greener than others, and there are also local minimum carbon time slots specific to every region.

Carbon intensity can be a useful metric when scheduling content delivery. Sending content from an origin server to a cache consumes energy on both sides. However, picking a time-slot that minimizes the carbon intensity on both ends means that the carbon footprint of the operation is reduced.

3 Carbon Intelligent Content Scheduling

In this work the delivery of content considered is static or does not require frequent updates. As a leading example, consider video streaming services. Each week, new shows are released and should be delivered to caches. Currently, off-peak hours are used for this task. A sustainable content delivery scheduling scheme for CDNs would build on the knowledge of both forecasted carbon emissions (combining energy and carbon intensity) of sites and traffic patterns per region. The content delivery schedule would minimize the overall carbon emissions while not violating traffic or cost constraints. Content transfer can be delayed or sent earlier to align with low carbon intensity and meet other constraints.

In [44], carbon-aware routing was defined as combining carbon emissions with standard routing methods, while carbon-intelligent routing was defined as combining carbon

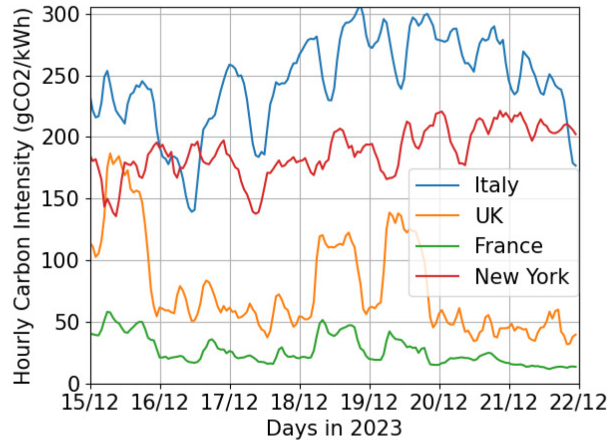

Figure 1: Average values of national (Italy, UK, France) and regional (New York) hourly carbon intensity over a week in Dec 2023.

emissions with special routing methods. Similar definitions are used to address content delivery scheduling: a carbon-aware schedule will use forecasted carbon intensity as an additional metric when setting the delivery schedule, without changes to the delivery path or other practices. On the other hand, a carbon-intelligent approach considers beyond carbon intensity also changes to the routing, such as using intermediate nodes to store content until carbon-intensity decreases. A carbon intelligent scheduler jointly accounts for all caches in a system and aims for an overall minimum carbon footprint, meaning a global minimization problem.

The origin server where all content is first stored, may not be the greenest node in the CDN. Some regions benefit from renewable energy, such as Sweden or Seattle (USA) where hydropower is abundant. CDN sites located in such greener regions, once received some new content, can act as substitute senders to other nodes, especially if the carbon intensity at the origin server's location is high.

Figure 2 illustrates three types of scheduling considered in this paper. The first approach (block dashed line) is a naive approach where content is sent at 2AM from the origin server to each cache separately, in line with [27]. The second approach (purple dashed line) is a one-hop carbon-aware approach where the time-slot of the lowest carbon intensity of both ends is selected. Hence, content is sent from the origin server to Cache 1 at 4AM and to Cache 2 at 6AM. Finally, the third approach (green dashed line) uses carbon intelligent content scheduling; multiple hops are supported, and content can be sent from cache to cache, rather than just from the origin, reducing the overall carbon emissions. In the example of Figure 2, the delivery schedule with the lowest carbon footprint sends content from the origin server to Cache 1 at 4AM and then from Cache 1 to Cache 2 at 6AM.

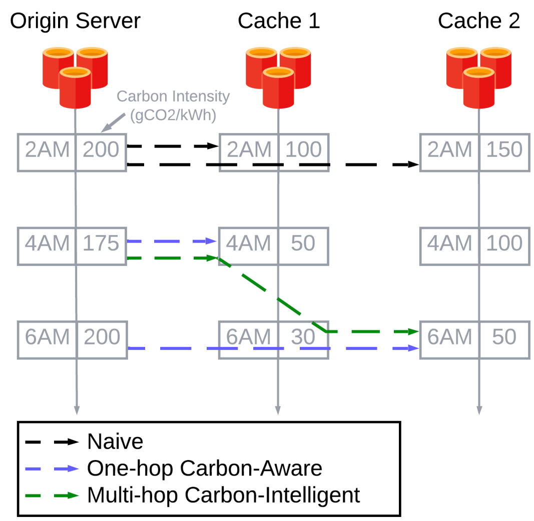

Figure 2: Illustration of the three scheduling options. The carbon intensity (in grey) varies between locations. The naive approach sends the content at 2AM (off-peak). The one-hop carbon-aware approach picks the lowest carbon intensity time-slots for every source-destination pair. The multi-hop carbon intelligent approach builds a schedule that sends first to Cache 1 and later from Cache 1 to Cache 2.

4 Carbon Emissions Model

Assume a given network, where all equipment is already deployed, and consumes idle power. Routers draw nearly constant power, both due to router properties [22] and because the cache uses at most one port per router– equivalent to less than 3% of the router's dynamic power for a small 32 port router. Therefore, the scheduling of filling caches is dominated by the dynamic power of *caches*.

The carbon model associated with the content filling can be expressed as: $\sum_{t=1}^{T}\sum_{i=1}^{N} e_{i,t} * c_{i,t}$, where N is the number of caches and T is the number of time slots (of duration 1 hour), $e_{i,t}$ is the dynamic energy of cache i due to reading and writing content during the time interval t, and $c_{i,t}$ is the carbon intensity at the location of cache i. The cache appliances are assumed to be fully power proportional. While in practice some of the power consumption will be static, drawn regardless of activity, this element is a constant that only offsets the results. As caches are CPU-based, they have better power proportionality than routers [21].

From a holistic perspective of the carbon footprint, the delivery schedule uses existing infrastructure, which impacts only operational emissions while leaving embodied emissions unchanged. Cache placement is an orthogonal research question, where embodied carbon should be considered.

5 Evaluation

In the following sections, several questions are addressed:

- Can carbon-related metrics be used for content delivery without violating other constraints?
- Is there a need for a dynamic content delivery plan based on carbon-intensity?
- What are the environmental benefits of carbon-aware and carbon-intelligent content delivery scheduling?

To answer these questions, the focus in this paper is on one case-study: distributing content across Netflix's content distribution network, Open Connect.

5.1 Scope

This work primarily asks if it is beneficial to include carbon-related metrics when scheduling content delivery. It is based solely on publicly available information, meaning that results should be considered relative to the used baseline, rather than as absolute numbers. There are some known gaps in the evaluation, such as the lack of full content delivery list or energy consumed by the transmission network, including repeaters. These are further addressed in §6.

5.2 Datasets

This work relies on publicly available data. Netflix Open Connect is chosen due to its transparency [29] and its significant share of Internet traffic (13.7%) [34], making carbon reductions impactful. Collected data includes content popularity, network traffic patterns and carbon intensity.

Open Connect Network: The detailed locations of Internet exchange (IX) points are used [31]. Netflix also provides the type of appliances available and their power ratings [30].

Top 10 Content: A list of the weekly top 10 movies and series per country is available from Netflix [33]. This information is collected from July 2021 to June 2024, but only 2023 data is used. Our analysis shows that about 50% of films and 25% of TV shows are new entrants to the Top 10 each week. The dataset includes 6885 unique titles, which is a significant portion of the 15,994 titles in the Netflix library [32], and the ones most likely to be in the cache. It includes both global and local titles. All IX nodes are assumed to share the same content, that is all IX points have a full library of content. Moreover, content needs to be added only once, and evicting content does not consume energy (overwriting disk space).

The average content distributed weekly is 3.53 TB. The breakdown of the size of content is estimated separately for movies and TV shows. This size depends on the length of content and the compression level used. Netflix stores each show in multiple compression levels to be streamed according to the user's available bandwidth. Netflix uses at least basic, standard, high, full and ultra-high definition options for content viewing [28]. The estimated size per hour sums up to a total of 12GB for all 5 compression levels [28]. The length of content is estimated as 1 hour per episode in TV series (of 10 episodes on average [7]) and 2 hours per movie [20]. Hence, the estimated size of content is 24GB per movie and 120GB per TV series.

Traffic Patterns: Content delivery is usually done during off-peak hours. To check off-peak hours, Cloudflare Radar [11] is used, which monitors traffic levels per country. Off-peak hours vary slightly between regions, however, the local time range starting at midnight and ending at around 6-7 AM can be perceived as a minimum off-peak time-interval with traffic levels below 60% across the globe.

Carbon Intensity: The carbon intensity data is collected on an hourly basis from Electricity Maps [19] for the year 2023. The locations of IX nodes in Open Connect are mapped to the geographical zones provided by Electricity Maps. While historic data is used, in real time, there are forecasts of the carbon intensity per region ahead of time [6] which enables the suggested scheduling scheme.

5.3 Implementation

To find the carbon-aware and carbon-intelligent content delivery schedule, the problem is formulated as an integer linear programming (ILP) problem. The scheduling problem is represented by a graph of all IX nodes where all nodes should receive their new content demand. There are constraints on content delivery times, i.e. off-peak hours. The objective is to reduce the overall carbon emissions. The model is formulated in Python and solved using the Gurobi optimizer [1]. The code used in this paper is available at [18].

5.4 Results

We compare the three scheduling approaches previously described: naive scheduling case (i.e. sending at 2AM), carbon-aware (one-hop) scheduling and carbon-intelligent (multi-hop) scheduling cases. The new schedules abide by the off-peak time constraints 100% of the time while the naive scheduling has 20% chance of violating the off-peak hours because of the difference in time-zones.

Figure 3 shows the weekly carbon emissions for content delivery in the US and Europe calculated over a year for a 7-days-ahead schedule. Both approaches that use carbon intensity reduce carbon emissions compared with the naive case. The one-hop carbon-aware schedule saves on average 13% of carbon emissions in the US and 17% in Europe, compared with the multi-hop carbon-intelligent schedule that saves 53% in the US and 64% in Europe. These are significant savings, showing that collectively optimizing a content delivery schedule benefits from leveraging the greenest nodes.

Figure 4 shows the effect on carbon savings of extending the duration of a schedule. As new shows need to be available

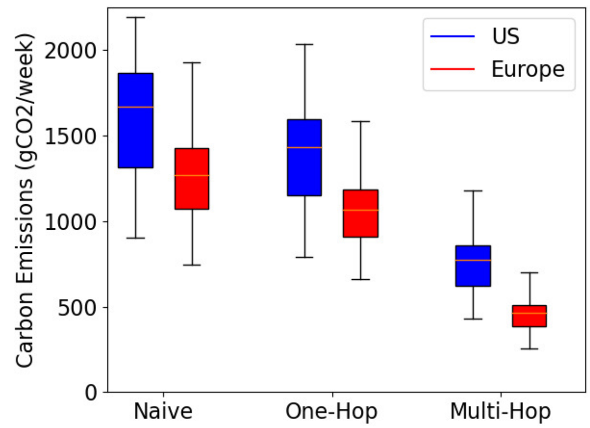

Figure 3: Carbon emissions for content delivery in the US and Europe, comparing the naive case (sending 2AM) with carbon-aware (one-hop) and carbon-intelligent (multi-hop) scheduling

in a cache by a specific date, this approach allows to deploy them up to a week before the required date. In this manner, there are more opportunities to benefit from lower carbon time-slots and as the figure shows savings can increase by 18% for a sample week. The savings curve shows the biggest difference is changing the schedule from a single day to a span of 3 days, and then the change in savings is less significant. This indicates that it is advantageous to plan at least 3 days ahead to get the highest carbon savings possible.

A geographical illustration of a weekly delivery schedule in the US is presented in Figure 5. The color code represents the day of the week for sending content. In the naive schedule, all nodes receive content at 2AM on Monday. With the one-hop carbon-aware schedule, there is variability in sending time, based on local carbon intensity. The multi-hop carbon-intelligent schedule is partly the same as the carbon-aware schedule, while others have a delayed or an earlier time slot. The time slot changes only if a better path combination of low carbon slots is found.

The delivery schedule in Europe using the multi-hop carbon-intelligent schedule is presented in Figure 6. The color indicates the number of hops travelled by the content to reach the edge caches. Interestingly, when multiple hops are allowed, content delivery is from an origin server in the east

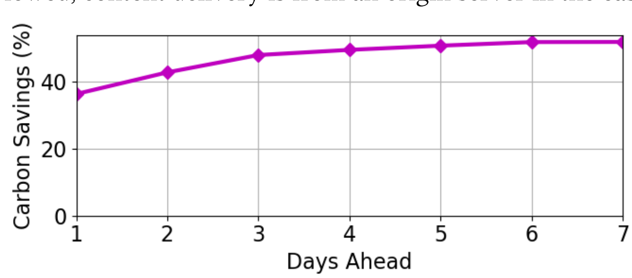

Figure 4: Improvement in Carbon Savings when planning for more days ahead for the multi-hop carbon-intelligent approach in the US for one sample week

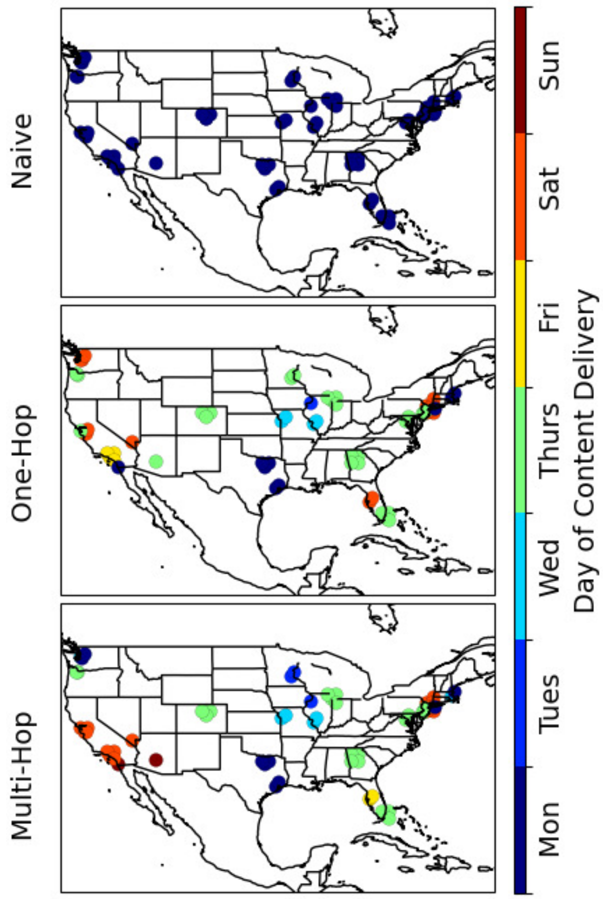

Figure 5: Geographical representation of content delivery in the US over a week, comparing the naive case, one-hop carbon-aware and multi-hop carbon-intelligent approaches

coast to Sweden, from there to France and then to other nodes. Sweden and France are the "greenest" in Europe and the transmission to other nodes is split between the two.

6 Discussion

Traffic Constraints. The off-peak hours constraint for content delivery limits the available scheduling time-slots. In some regions, e.g., those that rely more on solar energy, carbon intensity may be lower outside the range of hours used. Finding a common time between the origin server and some regions is difficult, especially with large differences in time zones. However, when an intermediate green node is used, new joint low carbon slots are available. As peak traffic hours are not simultaneous across the globe, there are more opportunities for the scheduler.

Scheduling Window. The time between the content is first available and the time to release it to users may not be long. If time allows, extending the scheduling window allows the scheduler to explore more low-carbon intervals. In one sample week in the US, carbon savings ranged from 34% for planning 1 day ahead to 52% for planning 6-7 days ahead.

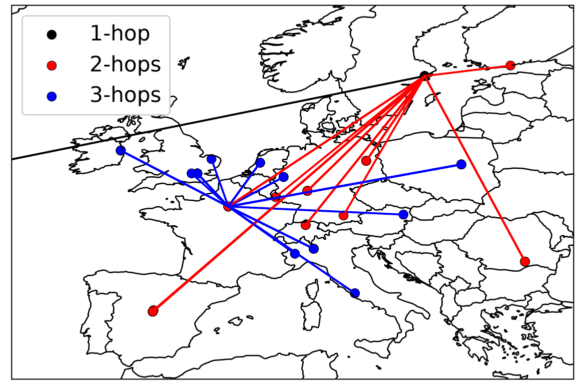

Figure 6: Geographical representation of content delivery in Europe for one week using the multi-hop carbon-intelligent approach

Caches Location. In this paper, the evaluation used only the advertised Netflix caches at IX points. These locations do not span all countries, e.g., Norway or Portugal. However, embedded caches inside local ISPs are an important part of the overall CDN [4] and are even more geographically spread. These caches also require content scheduling, but are not included in this work due to lack of public data. However, the suggested algorithm is applicable also to in-ISP caches.

Energy Efficiency vs Carbon Efficiency. In this study, one type of appliance is used and it is assumed that the same content will be distributed to all nodes. Hence, the energy consumed by the three approaches is the same. The carbon savings are therefore only due to carbon intensity variability. Energy consumption will vary if different appliances are used or if the nodes store different content. This can be supported by our algorithm.

US vs Europe. The US and Europe have similar cache counts, but US emissions are consistently higher. Moreover, carbon savings peak at 53% in the US vs. 64% in Europe. Renewable energy is more available in Europe and better low-carbon time-slots are selected by the carbon-intelligent algorithm.

Static Delivery Schedule. The results show that the greenest nodes in a CDN are frequently chosen to deliver content to other caches. Hence, a static delivery schedule might be used instead of recomputing the paths. Caches can be ranked based on two metrics: (1) average carbon intensity and (2) strategic location based on the density of connectivity to other nodes. The top greenest countries/regions with Netflix IX sites in order are: Seattle, Sweden, Brazil, France, Costa Rica, Uruguay, Finland, Switzerland and New Zealand. However, some locations like New Zealand are far from remaining nodes, while others like Uruguay may lack direct connectivity to other locations [37]. Beyond choosing the best intermediate nodes, off-peak delivery times need to be chosen. An estimate of these times can be derived from a few runs of the carbon-intelligent algorithm. A static schedule built for US and Europe saves on average 39% and 47% carbon emissions, respectively. While this is significant, it is 13% and 16% less savings, respectively, compared to the dynamic carbon-intelligent approach.

The Carbon Cost of Building a Schedule. As a delivery plan needs to be calculated anyway, adding carbon intensity as a metric and implementing our scheduling algorithms adds only a small overhead. The carbon emissions of calculating a multi-hop carbon-intelligent schedule were about $6gCO_2$, and less than $1gCO_2$ for the carbon-aware schedule.

The Carbon Emissions of the Network Content delivery involves not only caches, but also the network infrastructure connecting them (e.g., routers, repeaters). These components add to the overall carbon footprint. However, with the complexity of inter-domain routing and the lack of public information, it is currently hard to account for their emissions [44]. Moreover, the power variation of routers is negligible as the static power of routers is dominant [22, 41].

Limitations. This work has some limitations. First, only public data is used for the evaluation. The estimation of the weekly demand matrix is based on the weekly Top 10 list published by Netflix, and does not cover all new content. Off-peak time intervals are assumed between 12 AM and 7 AM (inclusive) for all countries, based on our analysis, but we acknowledge that this range may vary between countries and ISPs based on business agreements. As previously mentioned, the network emissions are not accounted for. In addition, the granularity of the carbon intensity data is per energy provider in the US and on a country-level for other countries. More granular data will increase accuracy. Still, the assumptions used in this paper indicate potential carbon savings through carbon-intelligent content scheduling.

Applications of the algorithm Carbon-intelligent scheduling can be used by applications that require prepositioning of content in cache. Beyond video on demand, this includes for example planned software updates and game releases.

Does It Really Matter? A common concern regarding carbon-reduction techniques in networks is that their effect is negligible compared with other contributors, such as the carbon footprint of creating content, e.g., producing a movie. While this is true, it does not excuse our community from developing environmentally-friendly solutions [44]. Furthermore, one of the reasons the environmental footprint of the Internet isn't bigger is because of the continuing efforts to reduce it [22]. The approach suggested in this paper enables carbon reductions with no infrastructure changes, and negligible implementation costs, making it an appealing optimization.

7 Related Work

The greening of video streaming is an active research area, including energy-conscious content steering [36] and efficient encoding techniques such as AI-based content aware encoding in [35]. More general aspects include energy-efficient practices for data centers such as smart cooling approaches [42] and energy-efficient workload management [5]. A CDN can opt for a greener streaming through green routing approaches such as energy-aware routing [14, 25, 40] and carbon-aware routing [17, 24, 43]. The Dagstuhl Seminar report [8] highlights the need for greener CDN operations for video streaming, including time-shifting content caching and energy-aware traffic pacing.

Power proportionality is important for energy savings. The memory in a server is estimated to consume 30% of its power [15]. The energy model of a memory can be estimated by the sum of static power and the dynamic power proportional to the read and write throughput values [15]. The choice of the storage device type such as Solid State Drives (SSD) and Hard Disk Drives (HDD), gives trade-offs on energy and performance as well [38]. For the Netflix use-case, both HDD-based and SSD-based appliances are used [30].

To the best of our knowledge, no previous work addressed green scheduling for the pre-caching phase in CDNs.

8 Conclusion

This work presented the concept of carbon-intelligent content scheduling in CDNs, using a case study of Netflix based on open-source data. The evaluation shows that it is a win-win solution: traffic constraints are conserved while carbon emissions are reduced. This approach does not trade-off on performance, but adds the carbon intensity metric into the content delivery scheduling, without infrastructure changes or additional costs.

Acknowledgements

This work was partly funded by NSF CBET-UKRI EPSRC TECAN (EP/X040828/1) and by the John Fell Oxford University Press Research Fund. We thank Eve Schooler and Te-Yuan Huang for their constructive feedback.

References

[1] 2024. Gurobi Optimization. https://www.gurobi.com/solutions/gurobi-optimizer/ accessed 20 Feb 2025.

[2] Akamai. 2023. 2023 ESG Impact Report. https://www.akamai.com/site/en/documents/akamai/2024/akamai-2023-esg-impact-report.pdf accessed 20 Feb 2025.

[3] Akamai. 2024. What Is a CDN? https://www.akamai.com/glossary/what-is-a-cdn accessed 20 Feb 2025.

[4] Timm Böttger, Felix Cuadrado, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2018. Open connect everywhere: A glimpse at the internet ecosystem through the lens of the Netflix CDN. *ACM SIGCOMM Computer Communication Review* 48, 1 (2018), 28–34.

[5] Rickard Brännvall, Tina Stark, Jonas Gustafsson, Mats Eriksson, and Jon Summers. 2023. Cost Optimization for the Edge-Cloud Continuum by Energy-Aware Workload Placement. In *e-Energy*. 79–84.

[6] Alasdair Bruce, Lyndon Ruffa, James Kelloway, Fraser, and Alex Rogers. 2024. Carbon Intensity API. https://carbonintensity.org.uk/ accessed 20 Feb 2025.

[7] Travis Clark. 2019. Netflix loves 10-episode TV seasons and reportedly doesn't see the value in longer original shows. https://www.businessinsider.com/netflix-favors-tv-shows-with-10-episode-seasons-report-2019-3 accessed 20 Feb 2025.

[8] 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.

[9] Cloudflare. 2024. CDN Benefits. https://www.cloudflare.com/learning/cdn/performance/ accessed 20 Feb 2025.

[10] Cloudflare. 2024. Cloudflare CDN Reference Architecture. https://developers.cloudflare.com/reference-architecture/architectures/cdn/#tiered-cache accessed 20 Feb 2025.

[11] Cloudflare. 2024. Cloudflare Radar. https://radar.cloudflare.com/ accessed 20 Feb 2025.

[12] Cloudflare. 2024. What is a content delivery network (CDN)? | How do CDNs work? https://www.cloudflare.com/learning/cdn/what-is-a-cdn/ accessed 20 Feb 2025.

[13] Michael Costello and Ellen Livengood. 2016. Netflix and Fill. https://netflixtechblog.com/netflix-and-fill-c43a32b490c0 accessed 20 Feb 2025.

[14] Maurizio D'Arienzo and Simon Pietro Romano. 2016. GOSPF: An energy efficient implementation of the OSPF routing protocol. *Journal of Network and Computer Applications* 75 (2016), 110–127.

[15] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. 2016. Data Center Energy Consumption Modeling: A Survey. *IEEE Communications Surveys & Tutorials* (2016). doi:10.1109/COMST.2015.2481183

[16] Disney. 2023. 2023 Sustainability & Social Impact Report. https://impact.disney.com/app/uploads/2024/03/2023-SSI-Report-4.pdf accessed 20 Feb 2025.

[17] 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

[18] Sawsan El Zahr, William Nathan, and Noa Zilberman. 2025. CICS. https://github.com/ox-computing/CICS/ accessed 16 June 2025.

[19] Electricity Maps. 2024. Electricity Maps. https://www.electricitymaps.com/get-our-data accessed 20 Feb 2025.

[20] Przemyslaw Jarzabek. 2018. Are New Movies Longer than They Were 10, 20, 50 Year Ago? *Towards Data Science* (2018).

[21] Chaoqiang Jin, Xuelian Bai, Chao Yang, Wangxin Mao, and Xin Xu. 2020. A review of power consumption models of servers in data centers. *applied energy* 265 (2020), 114806.

[22] Itzik Kiselevsky. 2023. Evolution of switches power consumption. https://eng.ox.ac.uk/media/11vdkdtb/itzikk_evolution-of-switches-power-consumption.pdf accessed 20 Feb 2025.

[23] Balachander Krishnamurthy, Craig Wills, and Yin Zhang. 2001. On the use and performance of content distribution networks. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*. 169–182.

[24] Julien Mineraud, Liang Wang, Sasitharan Balasubramaniam, and Jussi Kangasharju. 2016. Hybrid renewable energy routing for ISP networks. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*. IEEE, 1–9.

[25] Sergiu Nedevschi, Lucian Popa, Gianluca Iannaccone, Sylvia Ratnasamy, and David Wetherall. 2008. Reducing Network Energy Consumption via Sleeping and Rate-Adaptation.. In *NSDI*, Vol. 8. 323–336.

[26] Netflix. 2022. Environmental Social Governance Report 2022. https://s22.q4cdn.com/959853165/files/doc_downloads/2023/06/29/Netflix_2022-ESG-Report-FINAL.pdf accessed 20 Feb 2025.

[27] Netflix. 2024. Fill Patterns. https://openconnect.zendesk.com/hc/en-us/articles/360035618071-Fill-patterns accessed 20 Feb 2025.

[28] Netflix. 2024. How to control how much data Netflix uses. https://help.netflix.com/en/node/87 accessed 20 Feb 2025.

[29] Netflix. 2024. Open Connect. https://openconnect.netflix.com/en/ accessed 20 Feb 2025.

[30] Netflix. 2024. Open Connect Appliances. https://openconnect.netflix.com/en/appliances/ accessed 20 Feb 2025.

[31] Netflix. 2024. Peering with Open Connect. https://openconnect.netflix.com/en/peering/ accessed 20 Feb 2025.

[32] Netflix. 2024. What We Watched the Second Half of 2023. https://about.netflix.com/en/news/what-we-watched-the-second-half-of-2023 accessed 20 Feb 2025.

[33] Netlfix. 2024. Top 10 Movies on Netflix. https://www.netflix.com/tudum/top10/ accessed 20 Feb 2025.

[34] Sandvine. 2023. The Global Internet Phenomena Report January 2023. https://www.sandvine.com/hubfs/Sandvine_Redesign_2019/Downloads/2023/reports/Sandvine%20GIPR%202023.pdf accessed 20 Feb 2025.

[35] Robert Seeliger, Christoph Müller, and Stefan Arbanowski. 2022. Green streaming through utilization of AI-based content aware encoding. In *IEEE IoTaIS*. 43–49. doi:10.1109/IoTaIS56727.2022.9975919

[36] Daniel Silhavy, Will Law, Stefan Pham, Ali C Begen, Alex Giladi, and Alex Balk. 2023. Dynamic CDN switching-dash-if content steering in dash. js. In *Proceedings of the 2nd Mile-High Video Conference*.

[37] TeleGeography. 2024. Submarine Cable Map. https://www.submarinecablemap.com/ accessed 20 Feb 2025.

[38] Erica Tomes and Nihat Altiparmak. 2017. A Comparative Study of HDD and SSD RAIDs' Impact on Server Energy Consumption. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. 625–626. doi:10.1109/CLUSTER.2017.103

[39] Carbon Trust. 2021. Carbon impact of video streaming. https://ctprodstorageaccountp.blob.core.windows.net/prod-drupal-files/documents/resource/public/Carbon-impact-of-video-streaming.pdf accessed 20 Feb 2025.

[40] Nedeljko Vasić and Dejan Kostić. 2010. Energy-aware traffic engineering. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. 169–178.

[41] Arun Vishwanath, Kerry Hinton, Robert WA Ayre, and Rodney S Tucker. 2014. Modeling energy consumption in high-capacity routers and switches. *IEEE Journal on Selected Areas in Communications* 32, 8 (2014), 1524–1532.

[42] Ruihang Wang, Zhiwei Cao, Xin Zhou, Yonggang Wen, and Rui Tan. 2024. Green Data Center Cooling Control via Physics-guided Safe Reinforcement Learning. *ACM Trans. Cyber-Phys. Syst.* 8, 2, Article 19 (may 2024), 26 pages. doi:10.1145/3582577

[43] Yuan Yang, Dan Wang, Dawei Pan, and Mingwei Xu. 2016. Wind blows, traffic flows: Green internet routing under renewable energy. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*. IEEE, 1–9.

[44] 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