

Exploring the Benefits of Carbon-Aware Routing

SAWSAN EL-ZAHR, Department of Engineering Science, University of Oxford, UK

PAUL GUNNING, BT Research & Network Strategy, UK

NOA ZILBERMAN, Department of Engineering Science, University of Oxford, UK

Carbon emissions associated with fixed networks can be significant. However, accounting for these emissions is hard, requires changes to deployed equipment, and has contentious benefits. This work sheds light on the benefits of carbon aware networks, by exploring a set of potential carbon-related metrics and their use to define link-cost in carbon-aware link-state routing algorithms. Using realistic network topologies, traffic patterns and grid carbon intensity, we identify useful metrics and limitations to carbon emissions reduction. Consequently, a new heuristic carbon-aware traffic engineering algorithm, CATE, is proposed. CATE takes advantage of carbon intensity and routers' dynamic power consumption, combined with ports power down, to minimize carbon emissions. Our results show that there is no silver bullet to significant carbon reductions, yet there are promising directions without changes to existing routers' hardware.

1 INTRODUCTION

The fast development and deployment of the Internet has widely focused on reliability, scalability, speed and security. Starting in 2001, many initiatives tackled the power efficiency of Information and Communications Technology (ICT) for wireless networks [12, 48] and then in 2003 for wired networks [40]. In 2015, the Paris agreement set new sustainability goals of achieving 45% less carbon emissions by 2030 and reaching net zero by 2050 [54]. With this trend, and while ICT carbon footprint contributed to 2% of the overall carbon emissions in 2010 [23], ICT companies try to minimize their carbon emissions. Most works addressed data centers, improving power usage effectiveness (PUE) from 2.0 and above to the order of 1.1 for hyperscale data centers [39], as well as improving across all compute aspects: from CPU design to server, software and data center design.

Compared to data centers, fixed wired networks have seen limited improvement [69]. The improvements in this field are limited by the absence of standard power and carbon accounting metrics [19]. Fast technological advancements affect the contribution of different components within a router to the overall power consumption and thus, power metrics vary substantially with time. On the other hand, carbon metrics require visibility into the energy generation mix of the local power grid [28] which is difficult to integrate into the routing stack of deployed network elements [43].

In the past, estimating the carbon emissions of an entity was hard, in the absence of carbon-related information. However, the carbon net zero contention resulted in a plethora of carbon-related information, such as historical information of the carbon intensity of the electrical grid [21] as well as initiatives to forecast the carbon intensity per geographical region [10, 51, 76]. Such information is valuable for carbon-aware traffic routing and scheduling. The main goal of this work is to explore new carbon-related metrics, and quantify their importance in reducing the scope 2¹ carbon emissions associated with wired networks.

Existing metrics for traffic routing algorithms consist of the delay, hop count, bandwidth, Quality of Service (QoS) metrics and security metrics. None of the aforementioned metrics targets the power and carbon optimization problems in networks. Consequently, it is urgent to standardize new cost metrics that reflect the actual carbon emissions associated with the equipment in the network.

¹Scope 2 carbon emissions are indirect emissions due to the production of the energy that an entity consumes.

These carbon metrics cannot be separated from the power metrics because of the proportionality between the two parameters so power metrics are necessary to be investigated as well.

Traffic engineering (TE) algorithms can solve complex routing optimization problems [75]. In this work, we propose a carbon-aware traffic engineering (CATE) algorithm that optimizes for the overall carbon emissions of a network given the knowledge of a traffic matrix. Carbon metrics are incorporated into CATE algorithm [36, 37] to identify the least utilized links associated with the highest carbon intensity. The selected links will be shutdown to further increase carbon savings.

In this paper, we make the following contributions:

- We propose energy and carbon related metrics, and evaluate the carbon-savings associated with integrating these metrics into the link costs.
- We propose a new heuristic TE algorithm, minimizing carbon emissions further by using carbon-related link metrics and turning off unused links.
- We demonstrate in simulation the effect of real traffic patterns on carbon savings, using 2 Wide-Area Network (WAN) topologies and historical carbon-intensity information.

2 BACKGROUND

Intra-Domain, link-state protocols define the routing policy inside one autonomous system (AS). In particular, link state routing is the prevailing intra-domain protocol currently in use inside ASes. It presents a reliable and convergent process primarily based on sharing information about connected neighbors with every other router in the network. Its main steps include: initializing neighborhood relationships, exchanging the Link State Databases (LSDB) and finally, using Dijkstra's Algorithm to compute the optimal paths.

The optimal paths computed can follow a range of objectives, in particular energy optimization and carbon optimization. It is important to highlight the difference between these two problems. For the former, routers are classified based on power efficiency and priority routing is applied that prioritizes paths with the most power efficient routers. However, for the latter, routers are classified based on both the power efficiency and carbon intensity of the energy sources available at the node's location. This adds the geographical dimension to the routing problem and can prioritize a less efficient router fed by greener energy resulting in less carbon emissions.

In broader terms, the power grid consists of a mix of power stations that generate energy based on demand. There is an energy market that decides when to sell and which type of energy to sell [61]. And then, within a national or international network, generation in different areas will have a different carbon footprint based on the type of energy sources used. By weighting the imports and exports of energy between regions and running a power flow analysis, the overall carbon intensity of the energy consumed per region can be calculated. The carbon intensity metric refers to the amount of carbon emitted to produce 1 kWh of electricity whereas the overall carbon emissions are the product of the energy consumed and the average carbon intensity. For example, in the UK, the generation in Scotland is mostly based on renewable energy sources (especially, wind and hydro) resulting in very low carbon intensity. On the other hand, renewable energy sources are not that abundant in London and consequently, gas is one alternative to cope with the demand. This energy partition implies that a network Point-of-Presence (PoP) placed in Scotland can easily access renewable energy, will have a lower carbon intensity, and thus achieve a lower carbon footprint.

Nowadays, we have many sources of carbon reporting [10, 51, 76] that have visibility to the power grid. Figures 1 and 2 show the variation of the carbon intensity on national and regional levels, respectively [10]. The carbon intensity varies per day, per season and per region and noticeable changes can be seen within a few hours. Thus, accounting for the daily or monthly average of carbon




Fig. 1. National Carbon Intensity of the UK for 4 seasons




Fig. 2. Regional Carbon Intensity in the UK for three regions in a typical day in Fall 2022

intensity reduces the accuracy and may result in missed opportunities of low carbon intervals. This variability is the main motivation for this work: to accurately adapt the routing to greener paths.

3 RELATED WORK

Most of previous green-networking research focused on reducing the energy consumption (e.g., [2, 13, 15–17, 22, 46, 73, 78, 79]), while tackling the carbon efficiency of networks received limited attention (e.g., [47, 52, 65, 72, 74, 77]).

Moving from power efficiency to carbon efficiency is a non-trivial problem, as it requires knowledge of the energy sources powering the network equipment. Wang et al. [74] proposed an extension of OSPF that used the energy sources available at the nodes for path costs. It used a linear scale to express path costs based on the energy sources, without considering the actual carbon intensity of sources. In [72], the actual carbon emissions at the node were set as the link weights. This involves the local energy consumption and the overall average carbon intensity at the nodes. Traffic engineering was used in [65] to minimize the overall carbon emissions, but assumed that the renewable energy per node is static over time, which is imprecise. Partitioning energy into renewable and non-renewable energy was assumed in [47, 52, 77]. These works assumed additional renewable energy sources are available at every node and the goal was to optimize the non-renewable energy consumption. Routers were assumed to be able to instantaneously differentiate between the renewable and non-renewable energy portions available. This, however, is not a feature of current routers and requires technological changes.

Inter-domain carbon-aware routing was tackled in [53] and [69]. A new path attribute is added to BGP in [53] based on a minimum path emission (MPE) factor. On the other hand, [69] takes advantage of the SCION architecture in the context of path-aware networks. It derives an estimated carbon footprint for every inter-domain path and then, end-domains choose the paths with the least emissions. This should not be confused with the network-wide carbon optimization problem that removes any overlapping between end-domains.

Telecom companies achieved some improvement in terms of power and carbon efficiency. The main initiatives are placing more efficient equipment as well as adding dedicated renewable energy sources in their own facilities. Telefonica reported 49 MWh/PB and 353,346 tCO₂ in 2022 [70] that is equivalent to 87% and 80.5% reduction from 2015, respectively, while traffic increased by 637.6%. Google reported in 2021 6,757,312 tCO₂ of operational carbon emissions reduced to 0 after compensations by buying carbon credits elsewhere [38].

Moreover, the networking community has a wide discussion about the environmental impact of the Internet through the e-impact list under the Internet Engineering Task Force (IETF). One main goal is to identify the correct metrics that depict the contribution of every network entity to the overall carbon emissions. Finally, the work in [80] dissected the technical challenges of achieving carbon-intelligent routing. Steps go from standardizing metrics and reporting real-time network telemetry to time- and space-shifting data transmission.

3.1 Limitations and Gaps in Previous Works

The progress made in renewable-energy research and deployment, and changes in energy market behavior [42], creates a gap between the seminal research in green routing, and current state-of-the-art works, especially when applied to variation in carbon intensity and access to this carbon intensity information. Previous work either (1) considered unmixed energy sources (only one type of energy is available at every node of the system) [74] or (2) assigned inaccurate weights to different energy sources [47, 52, 77] or (3) used the average value of carbon intensity over a duration and assumed it to be a constant [72] or (4) used old meteorological data to estimate the amount of renewable energy at a certain location and time [77]. A different type of limitations concerns the power model of routers, assuming that the power is constant and independent of the traffic load [47, 52, 74]. Energy consumption of network devices has changed significantly over time, but only old values are still used [52, 77]. In addition, assumptions about switching chassis and line-cards into sleeping mode in marginal time are also used [47, 52], despite practical limitations.

4 SCOPE OF THIS WORK

This work focuses on the minimization of carbon emissions of intra-domain backbone networks. The goal is to quantify the potential benefits of carbon-aware routing. This work explores the performance limits that carbon-aware routing can achieve without any joint QoS optimization. We take a cautious approach, examining different metrics, topologies and traffic patterns, and trying to understand the drawbacks of different approaches. We therefore report both positive and negative results. As the timescales considered are in the order of 10's of minutes, this paper does not address orchestration or network stabilization. These are important problems left to future work. Carbon intensity information is gathered from open sources [10, 51] with a granularity of 30minutes-1hour, and we do not aim to predict it. Power models are based on publicly available data, and we explore the sensitivity of the results to this data, acknowledging that there are differences between vendors, platforms and equipment grades. The sensitivity analysis of static/dynamic power ratio, provided in section §8.5, reflects on the difference in power consumption models between routers' families.

Unlike previous works [47, 52, 77], we do not assume that renewable energy sources are available to nodes at all locations. We also seek network-wide carbon-reduction maxima, rather than local optimization. This work studies the status quo of network infrastructure, and tries to quantify the benefits of short-term and long-term approaches to carbon savings using carbon-aware routing.

5 METRICS

Reducing the carbon footprint of a network requires metrics related to energy and carbon emissions. Such metrics go beyond conventional metrics such as link capacity, delay, and utilization. The following sections discuss a set of potential energy-related and carbon-related metrics, with the goal of identifying practical and impactful metrics. The metrics are evaluated in subsequent sections of this paper. Figure 3 provides a simplistic overview of the system. In every point of presence there is a router. The energy-related metrics depend on the type of router and its utilization. The mix of energy sources used in the specific region of the router gives the carbon intensity metric. Finally, the carbon intensity combined with the router's information give the carbon emissions metric.




Fig. 3. Overview diagram of the system

5.1 Energy-Related Metrics

As long as PoPs are not completely provisioned from renewable energy sources, the carbon emissions of network devices are a function of their energy consumption. However, as previously pointed out [19, 80], energy consumption can be reported in different ways. An illustration of a potential power-consumption model of a router is shown in Figure 4. The figure shows a model where a router has a component of power that is static (P_{idle}), meaning independent of traffic load, and a power component that depends on the load ($P_{dynamic}$). While the figure shows a linear curve, we acknowledge that it may not be linear, e.g., due to the addition of line-cards with traffic or non-linear data plane's processing power. However, this assumption is based on our own validation experiments on available platforms as well as previous work [71].

To find the routes that minimize carbon emissions, it is required to identify metrics that can capture variation in energy and emissions with traffic. A particular challenge is the idle power of routers. Even if a router is not part of a green-routing plan, it is often kept working, for redundancy in the event of failures and to minimize recovery time. Therefore, the emissions from unused or underutilized routers need to be accounted for as well. A required feature of a chosen metric is to be time-invariant or slowly changing over time, such that updates can be (relatively) infrequent.

5.1.1 Typical Power: A first metric to consider is the typical power (P_{typ} as in Figure 4) of a router, usually defined as the average power consumption at 50% utilization [1, 3–8, 18, 25, 30, 41, 50, 55, 56]. It is found in datasheets, derived from idle and maximum power, or measured in the lab.

5.1.2 Energy Labeling: Energy labels are commonly used to indicate the energy efficiency of many appliances. While a legal requirement in many areas, there is still no standard for energy labels for routers. In this paper, we define a formula for quantizing the energy consumption based on the typical power of the router and the maximum capacity that it can operate with. In this paper, we propose an energy labeling solution that captures the energy consumption of a router as a function of its capacity.

We define the energy rating metric as the ratio of typical power in watts (W) divided by the maximum packet processing capacity C_{max} in Megapackets per second (Mpps), indicating the increase of power per packet. The typical power and the maximum capacity are indicators of the power consumption level and the size of the router, respectively.

Based on examining a range of routers available in the market [1, 3–8, 18, 25, 30, 41, 50, 55, 56], a 7-star scale can be suggested as shown in Figure 5, ranging from “A” to “G” rating. The value of the power-to-capacity ratios falls within the range [0;1]. The range is then divided into 7 intervals. Every interval is representative of a label starting from label “A” mapped to interval [0;0.1), reaching to label “G” mapped to interval [0.7;1]. To map these labels into the link costs, the scale should fit in the range [1;65535] (65535 is 2^{16} that corresponds to 16 bits used to express a link cost). Thus,




Fig. 4. A power model of a router




Fig. 5. A scale for energy labeling of a router based on ratio of typical power and packet rate

every label will be advertised by a cost proportional to the upper limit of its corresponding interval (i.e. interval $[0;0.1]$ of label “A” is advertised by 10, proportional to 0.1, “B” is advertised by 20...) as shown in Figure 5. The range of the 7 intervals are non-uniform because, based on available routers information, the density of energy label ratios falling in the interval $(0;0.5]$ is higher than the interval of $(0.5;1]$ corresponding to the least efficient routers. Quantizing this range of ratios from continuous to discrete, helps in advertising equal-cost paths and consequently avoiding congestion. To elaborate, to split the flows between equal-cost paths, the cost of these paths should be identical. Given an unquantized ratio of typical power over the packet processing rate as a cost-metric, even a minor difference in the cost results in ECMP ignoring one path that might have a very close cost. With the quantization in place, more paths with equal cost are available for ECMP, which helps reduce congestion.

Several alternatives can be used to the proposed labeling method, such as using capacity in bandwidth rather than packet rate, considering maximum power or dynamic power. We opted for the packet rate metric as different packet sizes at a similar bandwidth result in different power consumption levels for the router. Further power measurements are part of future work.

5.1.3 Incremental Dynamic Power Per Data Rate: This metric captures the change in a router’s energy consumption under load, meaning the slope of the curve in Figure 4 (assumed to be linear). This metric ignores the idle power of a router. It can be defined as the incremental power per processed packet or bit, meaning $P_{max} - P_{idle}$ normalized by C_{max} . With the power expressed in Watts and C_{max} expressed in Megabits per second (Mbps), the resulting metric unit is $[W/Mbps]$.

5.2 Carbon-Related Metrics

5.2.1 Carbon Intensity: Carbon intensity is a metric relating energy consumption to carbon emissions. It is defined as the amount of carbon emitted per unit of energy consumed [80]. Carbon intensity is a property of energy sources, rather of a router using these sources. It is computed as a weighted sum of the individual carbon contribution of every energy source connected to the power grid. For example, if the power grid is fed by 50% renewable energy sources (associated with 0 gCO₂/kWh), 20% coal (associated with 937 gCO₂/kWh) and 30% gas (associated with 394 gCO₂/kWh), then, the resulting carbon intensity of the grid would be $0.5 * 0 + 0.2 * 937 + 0.3 * 394 = 305.6$ gCO₂/kWh. A router does not need to differentiate between renewable and non-renewable energy sources, as the carbon intensity metric captures the mix of both. When a router is powered by a local energy source then the carbon intensity used is that of the local source rather than the grid. In our evaluation (§8) we assume that all routers are powered from the grid.

The evolution of the energy market enables today, with the aid of Machine Learning (ML) algorithms, to forecast the carbon intensity up to 24-48 hours ahead of time [10, 51]. Currently, the available information offers carbon intensity values per country, energy provider, and sometimes

per region. In addition, many energy providers report retroactively the carbon intensity of their grid [31, 59, 66, 68].

5.2.2 Carbon Emissions: The overall carbon emissions of a device are a product of its energy consumption multiplied by carbon intensity, integrated over time. For the purpose of carbon-aware routing, this metric is better weighted over 30 minutes or 1 hour intervals, and used for the next time interval. Importantly, the energy consumption is affected by multiple metrics, including idle power, dynamic power and utilization over the considered time interval.

6 EXPLORING THE INFLUENCE OF METRICS

A network topology can be modeled as a graph $G = (V, E)$ where vertices V and edges E of the graph are the network routers and links, respectively. In this section, an extension of OSPF is considered, where link costs vary based on the metrics described in §5. A similar methodology can be applied to other protocols, such as IS-IS.

For the purpose of exploration, a router $i \forall i = 1, \dots, |V|$ is assumed to have the knowledge of all the aforementioned metrics: its typical power $P_{typ,i}$, energy label $E_{label,i}$, incremental dynamic power per unit of traffic denoted by λ_i , the current carbon intensity denoted by c_i , and its average energy consumption over the previous time interval $e_{i,t}$. Links are bidirectional and link's costs differ between directions. A packet traversing a link $l_{i \rightarrow j}$, from node i to node j , will result in additional dynamic carbon emissions equal to the product $\lambda_j * c_j$ with j being the index of the receiving router. The receiving router is the node that will process that packet.

To study the impact of every metric on the potential energy consumption and carbon emissions reduction, multiple cases of individual metrics and a mix of metrics are considered, as detailed in Table 1. For every case, links are assigned a cost related to the corresponding metrics. The cost of a link in OSPF is a 16-bit number in the range $[1; 64K]$, with 0 forbidden, as it causes instability in shortest path calculation. To map to this range, cost metrics are scaled, as detailed in the last column of Table 1.

All energy-related metrics considered in this paper are static and change only when equipment is replaced, hence, can be done once per changed router by the administrator. Moreover, as suggested in [80], routers can acquire their power consumption information from smart Power Distribution Units (PDUs) connected to them. On the other hand, carbon intensity information is updated periodically from grid forecasting services that have visibility into the grid. This is accessed locally through a carbon intensity API as detailed in [80]. Information about any additional renewable energy sources (solar panels or wind turbines) placed at a certain PoP can be predicted locally. This can be based on local data such as historical data about batteries charging from renewable energy and then, the local carbon intensity value is adjusted accordingly. These features are not supported by current routers available in the market but there are no underlying technical challenges to implement them [80]. This is left for future work.

Next, every router, knowing its local values of the metrics, can compile them with its links state information and transmit the corresponding Link State Advertisements (LSAs) to neighboring nodes. Finally, with the global view achieved by LSA packets of OSPF, routers compute the greenest paths in parallel based on Dijkstra's algorithm. A router's ability to acquire its own power and carbon information is independent of the size of the network. However, the frequency of updates can affect the stability of the routing tables. Our model assumes standard network operations (e.g., using OSPF), with the change in link costs happening periodically in the order of 15 minutes to 1 hour (significantly higher than RTT) and hence, should not lead to system convergence issues.

Table 1. Metrics Considered

#	Metrics	Label	Cost($I_{i \rightarrow j}$)	Description & Scaling
1	Typical Power	Ptyp	$P_{typ,j}$	Average power under 50% load. It can be in the order of 50 kW for large chassis-based routers.
2	Energy Label	E-label	$E_{label,j}$	The mapping scale is presented in Figure 5, with a range of [10;100].
3	Incremental Dynamic Power per data rate	IncD	$\alpha * \lambda_j$	λ_j - the incremental dynamic power per Mbps (in W/Mbps) of node j . As values are typically in the range [0.00001;0.1] it is scaled up by a factor $\alpha \in [100k; 640K]$.
4	Carbon Intensity	C	$1 + c_j$	c_j - the carbon intensity (in gCO ₂ /kWh) at the location of the node j . Its range is [0;950]. Valid minimum value guaranteed by +1.
5	Combined Carbon Intensity and Typical Power	C + Ptyp	$1 + \beta * c_j * P_{typ,j}$	Estimate of carbon emissions based on typical power. Combines metrics #1 and #4. β - a down-scaling factor. $c_j * P_{typ,j} \in [0; 950 * P_{typ,max}]$. $\beta_{max} = 64K / (950 * P_{typ,max})$. Valid minimum value guaranteed by +1.
6	Combined Carbon Intensity and Energy Label	C + E-label	$1 + c_j * E_{label,j}$	Estimates the <i>scale</i> of carbon emissions based on energy label. $c_j * E_{label,j} \in [0; 9500]$. Valid minimum value guaranteed by +1.
7	Combined Carbon Intensity and Inc Dynamic Power	C + IncD	$1 + \lambda_j * c_j$	Estimates traffic-induced carbon emissions. The product $\lambda_j * c_j \in [0; 95]$. Valid minimum value guaranteed by +1.
8	Carbon Emissions	Emis-CE	$1 + \beta * c_j * (P_{idle,j} + \lambda_j * U_{j,\Delta t})$	Estimates carbon emissions based on full knowledge. $U_{j,\Delta t}$ is the utilization at node j during the previous time interval Δt . $(P_{idle,j} + \lambda_j * U_{j,\Delta t})$ is the actual power consumption at node j , with a maximum value of P_{max} . The overall range is $[0; 950 * P_{max}]$. β - a downscaling factor. $\beta_{max} = 64K / (950 * P_{max})$. Valid minimum value guaranteed by +1.

7 CATE: A HEURISTIC TRAFFIC ENGINEERING ALGORITHM

So far, this paper has assumed passive carbon-aware routing, where routing decisions do not lead to active changes to routers, beyond the routing tables. In this section, we explore the potential for further benefits by reducing the power-consumption of under-utilized routers.

Objective: Changing links' cost while the routers are kept on at all times, results in energy and carbon savings in terms of the dynamic power only. A different approach to enable saving more carbon emissions can be to reduce the idle power of the router as well. This, however, can be applied only to a limited degree. A common and straight-forward approach is to turn-off unused ports [47, 52, 77]. The power consumed by a transceiver and associated PHYs is constant, and therefore by incrementally disabling unnecessary ports, the idle power of a router can be reduced. Figure 6 illustrates the breakdown of a router's idle vs dynamic power components. Further to the

model in Figure 4, the idle power of a router can be divided into two parts: (1) the static power, including among others the idle power of the chassis, line cards (except ports), switching fabrics and the supervisor module, and (2) the power associated with ports.




Fig. 6. A breakdown of the power components of a router: idle vs dynamic power

High-level Description: If a link is shut down, then the ports at both ends are disabled. This results in a modified version G' of the graph topology G . However, shutting down links cannot be done randomly, and should consider the state of the network: the graph should remain connected, and if the network is congested, then running the same traffic on a reduced topology would worsen the situation. Hence, knowledge of a traffic matrix T is necessary to identify low-utilization links.

IP-based Traffic Engineering (TE) is the basis of this approach [75]. It changes the weights of links based on the specified optimization problem while keeping traditional OSPF routing for packet forwarding. This reduces the computation complexity because in case of a link failure, traffic can take alternative shortest paths set by OSPF. It accounts for the traffic matrix information in addition to other metrics considered previously. Other Traffic Engineering approaches include Resource Reservation Protocol Traffic Engineering (RSVP-TE) [9] and Segment Routing [29].

To reduce carbon emissions beyond passive carbon-aware routing, we explore the combination of carbon-aware routing metrics and shutting down links. Based on preliminary exploration (results reported in §8), link costs are selected to be based on combined carbon intensity and incremental dynamic power ($C + IncD$). A heuristic approach is used, determining recursively the link with the least traffic load and highest carbon emissions, and shutting it down subject to three conditions: (1) the graph G' stays connected, (2) the network can still satisfy all traffic demands without exceeding any link capacity and, (3) the improvement in carbon emissions introduced by shutting down this link is positive. Additional conditions, such as redundancy, can be supported too. The QoS metric of path stretch is evaluated in §8, but is not used by the algorithm (as justified later).

This algorithm does not shut down important paths in the topology for three reasons: (1) the link load is considered, meaning that if a link is a critical one in the topology then, this link would not be turned off. (2) If the alternative path to a link is very long it would result in higher carbon emissions and the algorithm will automatically deselect this link. (3) Additional QoS constraints such as the maximum end-to-end delay can be added to the algorithm, thought in this work, we opted to explore only the highest potential level of carbon reductions without QoS constraints.

Algorithm Description: Let T be a given traffic matrix with source-destination (s - d) pairs. Based on the Dijkstra algorithm, paths can be computed and the resulting flow intensity at every node and link are calculated. We denote by y_l and x_n the intensity of flows on a link l (sum of flows in both directions) and on a certain node n , respectively. To simplify the notations, let $Y, X = \text{Dijkstra}(G, T)$ be the function returning the set of flow intensities $Y = y_l \forall l \in E$ and $X = x_n \forall n \in V$, given a graph topology G and a corresponding traffic matrix T . The improvement is assessed based on the reduction in carbon emissions. The carbon emissions saved in this case are a function of carbon intensity, and both the dynamic power and the power of enabled ports. The total carbon emissions within a time interval is denoted by C_{tot} and is expressed as in Equation (1). The static power is not part of the optimization and hence, is not included in Equation (1).

$$C_{tot}(G, X, \lambda, c, \beta) = \sum_{n=1}^V (x_n * \lambda_n * c_n) + \sum_{l=1}^E \beta_l (c_{l_{n_1}} + c_{l_{n_2}}) \quad (1)$$

where β_l is the incremental power consumption of enabling one port and l_{n_1} and l_{n_2} are the nodes at both ends of a link l . Algorithm 1 briefly explains the steps of the routing operation.

Implementation: CATE is a centralized offline TE algorithm, that takes as inputs the power consumption model and carbon intensity values of every router in the network, as well as a traffic matrix. The traffic matrix is usually predicted based on previous traffic traces available for an AS. CATE is run periodically every time the carbon intensity values or the traffic matrix of connectivity are updated. The carbon intensity values vary every 15 minutes – 1 hour while traffic matrix updates frequency varies from order of minutes to hours [64]. The frequency of rerunning CATE can change on a use-case basis, but is on a minutes-scale or more. Information about disabled links is propagated at the start of every interval using an extended version of the LSAs. Adjustments to the format of the OSPF-LSA account for the additional metrics and differentiate between a disabled link and a link failure.

Response to link failure and link overflow: Unexpected link failures and link overflows need an additional online step to re-enable necessary links and recover from packet drops. CATE can be extended to provide the operator with a sorted list of links to be re-enabled as alternatives to failed or congested links. The predictability of carbon intensity ahead of time enables planning the response to topology changes beforehand. Adjusting to unexpected traffic load can be, for example, by setting a utilization threshold at every node, and if the utilization increases beyond this threshold, the algorithm starts to re-enable locally disabled links. Any TE algorithm requires an online version to adjust to the dynamics of the network [75]. This is an additional level of complexity but is essential for proper operation of the system.

Complexity and Optimality: The bulk complexity of the CATE algorithm stems from the use of Dijkstra's algorithm at every iteration. The complexity of Dijkstra's algorithm is $O(|E|\log|V|)$ where $|E|$ and $|V|$ are the number of links and routers in the network, respectively [62]. At every iteration, a link is checked to be removed from the graph and hence the number of iterations can be at most equal to the number of links $|E|$. Therefore, the complexity of this algorithm is $O(|E|^2\log|V|)$.

An optimal solution would consider all the possible permutations of used network links and calculate the possible carbon savings based on it. This however would lead in the worst case to a factorial complexity of the order $O(|E|!(|E|\log|V|))$. CATE is a sub-optimal solution that sorts the links based on their utilization and associated carbon emissions, which reduces the complexity of the algorithm at the expense of reduced accuracy. Moreover, for large networks with a high number of links, it is possible to omit multiple links at the same time (modifying line 9 in Algorithm 1 to pick L links with the highest cost instead of only 1) to further reduce the runtime of the algorithm. An Integer Linear Programming (ILP) based solution is likely to replace the Dijkstra algorithm in a more advanced version of CATE that uses segment-routing, with additional QoS constraints. In our algorithm, to assign the flows to paths, we used Dijkstra whereas MPLS-TE uses ILP, which increases the complexity. Dijkstra's algorithm is sufficient in this case to explore the potential carbon savings of the network.

QoS Constraints: Additional constraints can be added to CATE, such as maintaining a certain level of redundancy, setting a maximum packet delay or a maximum link load, and many others. The further constraints can be defined by amending line 14 in Algorithm 1 to check for them. Advanced QoS constraints, such as load balancing and congestion control, require further modifications to the CATE algorithm by for example, using ILP techniques instead of Dijkstra's algorithm to compute

Algorithm 1: CATE Algorithm

Given: $G = (V, E), T, \lambda, c, \beta$; /* The weights in G are based on the values of IncD λ and carbon intensity c as in setup 7 in Table 1 */

Result: $G' = (V, E')/E' \subset E$

```

1   $Y, X = \text{Dijkstra}(G, T)$ 
2   $C_1 = C_{tot}(G, X, \lambda, c, \beta)$ 
3   $G' = G, I_{set} = 0$ ;          /* Set the initial values of  $G'$  and the improvement  $I_{set}$  */
4  while true do
5       $G_{test} = G'$ ;          /* Define a graph to test on it */
6       $\Psi = \emptyset$ ;          /* Define a set of necessary links for  $G'$  to stay connected */
7       $G_u = \text{GetUndGraph}(G_{test}, Y, \lambda, c)$ ; /* Define the undirected graph  $G_u$  */
8      while true do
9           $l = \max\{G_u\}/l \notin \Psi$ ; /* Pick the link with maximum cost */
10         if  $l = \emptyset$  then
11             return  $G'$ ; /* No more links can be shut down */
12         end
13          $G_{test} = G_{test} - \{l\}$ ; /* Remove  $l$  from  $G_{test}$  */
14         if  $G_{test}$  not connected then
15              $\Psi.add\{l\}$ ;  $G_{test} = G_{test} + \{l\}$ ;
16         else
17             break;
18         end
19     end
20      $Y, X = \text{Dijkstra}(G_{test}, T)$ ; /* Recompute the flow intensities */
21     if  $\exists l \in E / y_l > C_l$  then
22         return  $G'$ ; /* link capacity exceeded */
23     end
24      $C_2 = C_{tot}(G_{test}, X, \lambda, c, \beta)$ ;
25      $I = C_1 - C_2$ ; /* Check the improvement */
26     if  $(I > I_{set})$  then
27          $G' = G_{test}$ ;  $I_{set} = I$ ;
28     else
29         return  $G'$ ; /* Improvement started to diminish */
30     end
31 end
32 function  $\text{GetUndGraph}(G, Y, \lambda, c)$  :
33      $G = (V, E)$ ;  $G' = G$ ;
34     for  $l \in E'$  do
35         if  $y_l = 0$  then
36              $G'(l) = MAX$ ; /* MAX is the maximum integer value */
37         else
38              $G'(l) = (\lambda_{l_1} * c_{l_1} + \lambda_{l_2} * c_{l_2})/y_l$ ; /* the higher  $y_l$  the least the cost */
39         end
40     end
41     return  $G' = (V, E')$ 
42 end

```

the link and node loads in lines 1 and 20 of Algorithm 1. This paper explores the performance limits that carbon-aware routing can achieve without any joint QoS optimization.

8 EVALUATION

The evaluation of carbon-aware routing benefits is done in simulation, using ns-3 [60]. The following section describes the evaluated network topologies and routers setup. It then focuses on three aspects: the effect of using different metrics, the benefits of shutting down links, and the effect of using different routers.

This work focuses on general routing by Internet Service Providers (ISPs). Carbon-aware routing can be applied also for other real-world applications. For example, it can be applied to content distribution networks (CDNs), and to cloud computing as part of carbon-intelligent computing [14, 24, 32, 49, 67]. We leave such applications to future work. Carbon-aware routing cannot be used within data centers, as all paths share the same energy source and in turn the same carbon intensity and are therefore distinguishable only based on their energy consumption.

8.1 Evaluated Topologies

Two topologies are used in the evaluation: one national and one international. British Telecom (BT) [63] is a national ISP centered in the UK; GEANT is an international research and education network that provides connectivity spanning Europe. These two Wide Area Networks (WAN) have a contrasting hierarchy and different levels of redundancy. The main reason to choose these topologies is the availability of both nodes location and carbon intensity information.

8.1.1 BT-UK. The BT network spans across the UK, offering both residential and business services. As of 2022, BT connects 45% of UK households and 1.2 million enterprise customers [11]. Figure 7 describes the BT core network hierarchy, based on [63]. It consists of 1008 nodes and 3111 bidirectional optical links with up to 100 Gbps bandwidth. The nodes are hierarchically divided into: 20 core nodes, 63 metro nodes and 925 tier 1 nodes [63]. Access nodes are managed by Openreach (a subsidiary of BT) and omitted. Every PoP carries also local traffic by serving its connected macro-cells in addition to routing backbone traffic. The arrows in Figure 7 represent the local traffic received by every node.




Fig. 7. BT Network Topology




Fig. 8. Variation of BT traffic load with the time of the day [45]

The traffic intensity for the BT network varies across a 24-hours period [45]. Figure 8 shows a peak at around 8:00 PM while traffic is relatively flat during the working hours of the day. Similar traffic patterns are reported by Cloudflare Radar [20]. Most traffic peaks that are in excess of the average diurnal variation are anticipated, e.g., new gaming product releases, live premiership football fixtures, and new series from streaming service providers. Based on that, two traffic patterns can be identified that explain the variation in traffic intensity during each 24-hours period:

- **Daytime Traffic:** Traffic from business customers during working hours [9AM - 5PM] that is mostly symmetric (any-to-any) [45].
- **Evening Traffic:** Residential customer traffic dominates the evening hours, with a peak between 7PM and 8PM [45]. The traffic is predominantly downstream of content (90%), taken from third-party content caches co-located within metro-nodes.

8.1.2 GEANT-Europe. GEANT network is a main research network based in Europe, interconnecting multiple national research and education network (NREN) partners. It differs hierarchically from BT and consists of 46 nodes based in 39 countries with 68 bidirectional optical links [35]. The network is in constant development, but links are assumed to have 100 Gbps capacity. According to [33], GEANT carries daily 7 PB of data and traffic can be assumed to be symmetric (all-to-all).

8.2 Simulation Setup

8.2.1 Equipment. GEANT implements its packet layer using Juniper MX routers [34], and BT uses Nokia 77xx routers [57]. However, detailed power consumption information about these routers is not available in the public-domain. Instead, we use information that is in the public domain of equivalent equipment from Arista Networks. This information includes both typical and maximum power usage. Hence, in our simulations we select: the Arista 7280CR3-96 [5], which is approximately equivalent to the equipment deployed in BT's core and metro nodes; and the Arista 7060CX-32S [4], which is approximately equivalent to the equipment used in both BT's tier-1 nodes and GEANT's PoPs. Our choice of routers is matched to the maximum number of links per node, allowing for the evaluation of both: the incremental dynamic power and the effect of router design (The static-dynamic trade-off is evaluated in §8.5.)

The power for enabling a singular 100G transceiver is estimated to be within the range 3.5 - 5.5W [44, 58]. We select a value of 4.5W to use in our simulations, however, our sensitivity analysis showed similar improvements for port power values that spanned between 3.5 and 5.5W.

8.2.2 Carbon Intensity Data. Historic carbon intensity data is used in the evaluation. Regional UK historic carbon intensity information is available from [10], and two typical days are chosen, one in the summer (02/08/2022) and one in the winter (01/02/2022). The level of carbon savings was closely the same for all the days of simulation, including fall and spring days. Only winter and summer are shown for conciseness. Carbon intensity values in Europe are taken from [51]. As only limited information is available from it, the simulation for GEANT topology uses only one typical winter day (02/12/2022).




Fig. 9. BT: Energy and Carbon Savings with respect to OSPF for Day-Traffic (winter)




Fig. 10. BT: Energy and Carbon Savings with respect to OSPF for Evening-Traffic (winter)

8.3 Effect of Metrics

The first step in the evaluation is to explore the potential carbon and energy savings using each of the metrics discussed in §6 and §7. The overall carbon and energy savings over the BT topology

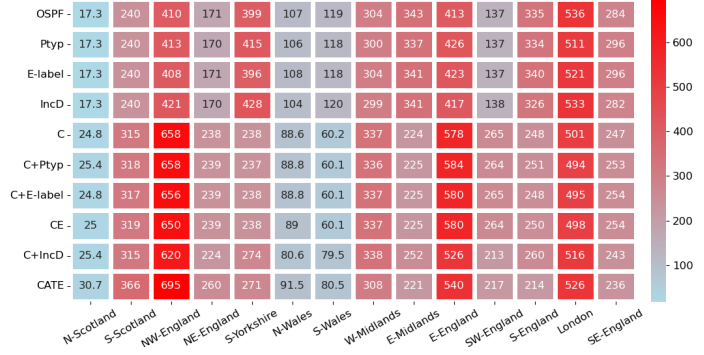



Fig. 11. BT: CDF of Overall Delay for Day-Traffic (winter)

Fig. 12. BT: Heat Map of Intensity of Flows per Region (Gbps) for Day-Traffic (winter)






Fig. 13. GEANT: Carbon Savings per interval (winter) with respect to OSPF

Fig. 14. GEANT: CDF of Delay (winter)

Fig. 15. BT: Improvement (w.r.t OSPF) when varying the Idle/Dynamic ratios for day-traffic




Fig. 16. GEANT: Heat map of intensity of flows routed per country in Gbps

are presented in Figure 9 (day traffic) and Figure 10 (evening traffic). Figure 11 shows a CDF of the overall packet delay corresponding to Figure 9. The labels in the figures match Table 1.

Energy and Carbon Savings: As discussed in §5, the metrics can be divided into two groups: those exclusively using energy information and those using regional carbon intensity knowledge (in addition to any other metric). As Figures 9 to 11 show, the second group of metrics achieves higher carbon emissions reduction, but at the expense of increased packet delay. However, this additional delay is small as the overall path stretch is small (6.6% increase on average). In particular, combining carbon intensity and incremental dynamic power (C+IncD) achieves the highest carbon savings of 3.71% during day time. Carbon intensity (C) metric alone is the second best, with 2.39% carbon savings during the day. The carbon emissions (CE) metric that accounts for the utilization of nodes, does not provide an advantage over the carbon intensity metric alone. For energy-based metrics,

using incremental dynamic power metric provides the highest energy saving of 2%, suggesting that C+IncD is indeed a good choice for a combined metric.

Change in Intensity of Flows per Region: Figure 12 presents the variation in the intensity of flows per region (in Gbps) using different cost metrics. It shows the shift of the flows to regions with lower carbon intensity. While it would have been desirable to see a considerable reduction of flows in London, which has a relatively high carbon intensity, this does not happen. The reason is that London has a high density of nodes, and its central location within the network necessitates other flows to pass through it.

Day-time vs Evening-time Traffic: Compared with day-time traffic, Figure 10 shows a limited improvement in carbon and energy during evening-time. This is mainly for two reasons: (1) During the evening downstream traffic is dominant, users downstream from the closest cache, which may be only 1 or 2 hops away. Hence, flows have a limited number of short paths to choose from (1.06 hops for downstream traffic, 2.5 hops on average for all traffic). The motivation to use carbon-aware routing is the variability of carbon emissions between paths, therefore given the limited choice, no matter the metric used, the carbon savings improvement is small. (2) The peak volume of traffic in the evening uses a high portion of the network capacity and most links are highly utilized and cannot be disabled, leading to the limited carbon savings in the evening.

Results for GEANT Topology: In the GEANT topology, all nodes are assumed to use the same routers and therefore have the same power parameters. Hence, only carbon-related metrics are applied and compared for this topology. The carbon savings over time and the corresponding CDF of the packet delay are presented in Figures 13 and 14, respectively. The carbon intensity metric shows the highest levels of savings level, but is almost matched by the carbon emissions metric (which uses the traffic matrix of the *previous* period), which has lower energy consumption. The packet delay CDFs of all the metric are similar. A heat map of the intensity of flows passing through routers in different countries is shown in Figure 16. This figure shows that nodes in Germany (DE), France (FR) and Italy (IT) are the most central nodes in the topology. France and Italy have an increased flow intensity because of the abundant renewable energy in these countries. However, Germany presents a bottleneck in GEANT because flows often have no choice but to route through Germany. Further experimentation shows that replacing the equipment in these three countries by twice more power-efficient devices can lead to 14.53% carbon improvement without disabling any links. We conclude that carbon-aware routing cannot completely deviate the traffic from high carbon intensity nodes, but can identify bottlenecks in a topology. By placing additional renewable energy sources at bottleneck nodes, carbon savings can be increased dramatically (§8.5). This matches a similar insight in [72].

In summary, the carbon intensity metric is the key metric to reducing the overall carbon emissions. Combined with energy metrics, specifically incremental dynamic power per unit of traffic, it can achieve even higher savings. However, these savings are only with respect to the dynamic part of the router power, whereas the carbon emissions due to the idle power are constant for all setups.

Further simulation results under different scenarios show similar patterns. Appendix A presents the values of the carbon emissions, energy consumption, delay, hop count and maximum link load.

8.4 CATE: Effect of Shutting Down Links

Passive carbon-aware routing, as shown above, can save a limited percentage of carbon emissions. Therefore, we explore the active approach using CATE. The portion of carbon emissions associated with ports is not negligible. Applied to the BT topology, CATE recommends disabling up to 32% of the links during day-time and up to 12.8% for evening-traffic, with 0% of links disabled at peak times. On the other hand, GEANT topology does not have a high level of redundancy and thus only up to 8% of its links can be shut down. Consequently, carbon improvements reach 9.68% for GEANT

and 17.96% in BT (day-traffic). The delay introduced by CATE is small, and the maximum link-load does not increase significantly after shutting down links. The maximum link-load is defined as the maximum instantaneous load transmitted on any link per direction. It indicates the potential level of congestion per metric.

While previous green-routing work focused on any-to-any traffic, the downstreaming-pattern in BT's evening-traffic was overlooked². The short paths needed for this traffic limit potential carbon savings using alternate paths. Given the limited ability to shut down links, a potential way to improve this scenario is move content even closer to the user while time-shifting loads to intervals with more renewable energy available.

8.5 Effect of Changing Equipment

Architecture and design decisions result in different ratios of idle/dynamic power of routers. Chassis-based routers, especially carrier-grade, are characterized by a large ratio of idle power for the chassis, contributed by management cards, fabric cards, redundant power supplies and fans, etc. The Arista routers used in the previous evaluation have a ratio of idle-to-dynamic power of 1:7, meaning that most of the power consumption is dynamic. However, large chassis and older routers can have up to 75%-80% idle power. To explore the effect of router design on potential carbon savings, we vary the idle/dynamic ratio, with the dynamic power as in §8.3 and the idle power increasing. The results, shown in Figure 15, indicate that the benefits of carbon-aware routing diminish as the ratio of idle/dynamic power increases. Similarly, improvements introduced by CATE decrease as well. Our results indicate that the best approach to reduce carbon emissions is invest in replacing equipment with ones of low idle power and higher dynamic power ratio.

8.6 Comparison with State of the Art

Already in 2012, Wang et al. [74] estimated that rerouting traffic without shutting down elements can introduce up to 16% savings for the links' carbon emissions. However, they assumed a single unmixed (weighted) energy source per node. Similarly, estimations in [72] were based on static carbon intensity, achieving 2% to 23% (average 3.7%) per-path carbon reduction for national research and education networks (NRENs). Actual carbon intensity values were not available in [72]. GEANT is a comparable network to the work in [72], with average 7.48% carbon improvement and a maximum of 8.36% using passive carbon aware routing. The higher average carbon saving is potentially as our work dynamically accounts for the carbon intensity. BT presents lower savings with passive carbon aware routing because the network has more nodes, leading to a higher idle power contribution.

Adding renewable energy sources was suggested in [47, 52, 65, 77]. Given guaranteed constant renewable energy, carbon savings were up to 50% in [65]. Shutting down entire nodes achieved up to 37% and 60% reduction in non-renewable energy consumption in [47] and [52], respectively. Finally, by only shutting down unnecessary links and given additional renewable energy source deployed for every router, up to 20% reduction in non-renewable energy was estimated in [77]. In this paper, we take a more practical and immediate approach, considering immediate potential carbon savings, using the current power grid. CATE introduces considerable improvement (11.93-17.96%) in carbon savings without changes to energy sources nor equipment. Using fine-grain carbon intensity provides higher accuracy than zero-one weight assignment for all renewable and all non-renewable energy sources.

The highest levels of savings were shown by shutting down entire nodes in [47, 52]. However, similar to servers, shutting down is undesirable given current technologies, as it affects reliability.

²There is significant prior research on green video streaming, but not in the same caching-routing context

Moreover, for the BT topology there is always local traffic, so no nodes could be shut down, though potentially some equipment within a PoP could be. This highlights inaccuracies in assumptions made in previous works.

The addition of renewable energy sources at specific sites increases the level of savings greatly [47, 52, 65, 77]. This is not considered in our work. However, as shown in §8, carbon-aware routing can help identify the nodes with least carbon emissions and enable traffic to benefit from such nodes.

In summary, this paper looked at practical current considerations, different to assumptions in previous works. It estimated similar carbon savings while accounting for more fine-grained carbon intensity data, technical operating considerations of routers, without assumptions on additional renewable energy deployments.

9 DISCUSSION AND RECOMMENDATIONS

Best metrics for carbon-aware routing. Carbon intensity and incremental dynamic power are the best metrics for carbon-reduction. Carbon intensity significantly varies during the day, and there is no single set of paths that optimizes carbon. The dynamic adaption of the routing to the carbon intensity metric can achieve some improvement with respect to the dynamic power of routers. Using incremental dynamic power per unit of traffic reflects the most on the actual dynamics of the network, therefore the combination of the two is most effective.

High idle power limits carbon savings. The carbon emissions associated with the idle power of routers are high, and may dominate over the dynamic part. Given a high idle/dynamic ratio and specifically for large networks, the carbon emissions resulting from processing incoming packets are negligible. Efforts should be directed towards reducing the idle power of routers with greener design techniques.

Routing bottlenecks limit carbon savings. Topologies have some routing bottlenecks. The majority of traffic flows will traverse the central nodes of the network independent of the routing priority. The geographical location of these nodes is often associated with low availability of renewable energy (e.g., London in UK, Germany in Europe). Efforts should focus on placing more efficient equipment and additional renewable energy sources specifically at the bottleneck nodes, as this will have an impact on the most flows. Moreover, in the long term, using structured finance for energy swap can transform high-carbon energy in central locations into low-carbon energy, and is a promising direction towards net zero carbon goals.

Using energy labels. Understanding the power efficiency of a router is a hard task, therefore energy labels can simplify customers' purchasing decisions, and are easy to use as part of carbon-aware routing. The definition of energy labeling in this paper accounts for the typical power and the maximum processing rate of a router. With this metric, an efficient router should minimize both the idle and dynamic power for high processing rates. A router with high idle power or with high dynamic power results in a less efficient rating. However, energy labels do not capture the effect of idle/dynamic ratio, thereby skewing carbon-aware routing decisions. It may therefore be useful to consider idle (or zero-port) and dynamic power as separate labels. It should be noted that while companies should buy more efficient routers, this needs to be balanced against the scope 1 and scope 3 emissions of newly manufactured and retired equipment, correspondingly.

Carbon optimization is application-specific. Application-wise, the length of the path taken by any flow will determine the amount of carbon that can be saved. Long paths mean a high variability of regions with different carbon intensity from which to select. However, flows associated with streaming applications have very short paths and cannot benefit from this approach. Only replacing or turning off some network equipment can slightly save on carbon in such case. Time-shifting these flows until more renewable energy is available might help, yet an application may have latency deadlines. Thus, application-specific carbon optimization should be considered [43].

Comparing different ISPs. The metrics provided are not enough to compare different ISPs. Based on our evaluation of the BT and GEANT topologies, we cannot claim that one is greener than the other. Each ISP has a different hierarchy and looking at the core networks only cannot give a complete view. Many aspects differentiate between ISPs such as the number of nodes, coverage, routers' type, type of customers and the services provided. The metrics discussed in this paper aim to explore the level of carbon savings that can be achieved without giving deterministic or comparable results on carbon efficiency of a network.

Limitations This work has several limitations. First, only public information is used in this work. We acknowledge that the routers used in this work are not the same routers used in the actual network, nor carrier grade, yet the required power figures required for the analysis are not publicly available. Second, our evaluation considers only inter-PoP traffic and no intra-PoP traffic, relevant for content streaming. However, our simulation considers a "best case" scenario while we claim limited improvement, and our sensitivity analysis compensates for lack of information and supports the results' trend. Considering these issues provides a balanced picture of benefits, limitations and considerations for carbon aware routing.

10 CONCLUSION AND FUTURE WORK

This work explored the potential benefits of carbon-aware routing, using different metrics, traffic patterns and real-world carbon intensity information. Our results show that routing is one way to reduce carbon emissions but is insufficient by itself. Re-design of network equipment by vendors to reduce idle power is necessary. While this work has introduced some spatial shifts in traffic routing, a next step is to tackle the temporal dimension, shifting traffic to intervals of time when renewable energy is more abundant.

The metrics considered in this paper will help guide the current discussion in the networking community about reducing the environmental impact of the Internet. The inclusion of the carbon and power metrics into routing decisions is one important step towards achieving a greener Internet.

ACKNOWLEDGMENTS

This research was funded in part by the Margaret Thatcher Scholarship Trust and the Qatar Fund For Development. We thank Eve Schooler for her constructive feedback. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript (AAM) version arising from this submission.

REFERENCES

- [1] Accton. 2012. ES5652BT1 Top-of-Rack Switch. <https://people.ucsc.edu/~warner/Bufs/Accton5652.pdf> [Online, accessed June 27, 2023].
- [2] Bernardetta Addis, Antonio Capone, Giuliana Carello, Luca G Gianoli, and Brunilde Sansò. 2013. Energy Management through Optimized Routing and Device Powering for Greener Communication Networks. *IEEE/ACM Transactions on Networking* 22, 1 (2013), 313–325.
- [3] Arista. 2015. Arista - 7060X4 Series 100/200/400G Data Center Switches. <https://people.ucsc.edu/~warner/Bufs/7060X4-Datasheet.pdf> [Online, accessed June 27, 2023].
- [4] Arista. 2016. Arista - 7060X and 7260X Series 10/25/40/50/100G Data Center Switches. https://www.arista.com/assets/data/pdf/Datasheets/7060X_7260X_DS.pdf [Online, accessed June 27, 2023].
- [5] Arista. 2023. 7280R3 Series Data Center Switch Router - Data Sheet - Arista Networks. <https://www.arista.com/assets/data/pdf/Datasheets/7280R3-Data-Sheet.pdf> [Online, accessed June 27, 2023].
- [6] Arista. 2023. Arista - 7020R Series Data Center Switches. https://people.ucsc.edu/~warner/Bufs/7020R-48_Datasheet.pdf [Online, accessed June 27, 2023].
- [7] Arista. 2023. Arista - 7050 Series 10/40G Data Center Switches. https://people.ucsc.edu/~warner/Bufs/7050S_Datasheet.pdf [Online, accessed June 27, 2023].

- [8] Arista. 2023. Arista - 7050X4 Series 100/200/400G Data Center Switches. <https://people.ucsc.edu/~warner/BuFs/7050X4-Datasheet.pdf> [Online, accessed June 27, 2023].
- [9] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow. 2001. RFC 3209: RSVP-TE: Extensions to RSVP for LSP Tunnels. <https://datatracker.ietf.org/doc/html/rfc3209>
- [10] Alasdair Bruce, Lyndon Ruff, James Kelloway, Fraser MacMillan, and Alex Rogers. 2020. Carbon Intensity Methodology. <https://www.carbonintensity.org.uk/> [Online, accessed June 27, 2023].
- [11] BT. 2022. BT Group plc Annual Report 2022. <https://www.bt.com/bt-plc/assets/documents/investors/financial-reporting-and-news/annual-reports/2022/2022-bt-annual-report.pdf> [Online, accessed June 27, 2023].
- [12] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. 2001. Span: An Energy-Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*. 85–96.
- [13] Luca Chiaraviglio, Antonio Cianfrani, Marco Listanti, Luigi Mignano, and Marco Polverini. 2015. Implementing Energy-Aware Algorithms in Backbone Networks: a Transient Analysis. In *2015 IEEE International Conference on Communications (ICC)*. IEEE, 142–148.
- [14] Luca Chiaraviglio and Ibrahim Matta. 2010. GreenCoop: Cooperative Green Routing with Energy-Efficient Servers. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. 191–194.
- [15] Antonio Cianfrani, Vincenzo Eramo, Marco Listanti, Marco Marazza, and Enrico Vittorini. 2010. An Energy Saving Routing Algorithm for a Green OSPF protocol. In *2010 INFOCOM IEEE Conference on Computer Communications Workshops*. IEEE, 1–5.
- [16] Antonio Cianfrani, Vincenzo Eramo, Marco Listanti, and Marco Polverini. 2011. An OSPF Enhancement for Energy Saving in IP Networks. In *2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 325–330.
- [17] Antonio Cianfrani, Vincenzo Eramo, Marco Listanti, Marco Polverini, and Athanasios V Vasilakos. 2012. An OSPF-Integrated Routing Strategy for QoS-Aware Energy Saving in IP Backbone Networks. *IEEE Transactions on Network and Service Management* 9, 3 (2012), 254–267.
- [18] Cisco. 2016. Cisco Nexus 3048 Switch. <https://people.ucsc.edu/~warner/BuFs/Cisco-3048.pdf> [Online, accessed June 27, 2023].
- [19] Alexander Clemm, Lijun Dong, Greg Mirsky, Laurent Ciavaglia, Jeff Tantsura, Marie-Paule Odi, Eve Schooler, and A Rezaki. 2023. *Green Networking Metrics*. Internet-Draft draft-cx-opsawg-green-metrics-00. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-cx-opsawg-green-metrics/> Work in Progress.
- [20] CloudFlare. [n. d.]. Traffic - Cloudflare Radar. <https://radar.cloudflare.com/traffic/gb> [Online, accessed June 27, 2023].
- [21] European Commission. 2022. EDGAR - Emissions Database for Global Atmospheric Research. https://edgar.jrc.ec.europa.eu/report_2022 [Online, accessed June 27, 2023].
- [22] 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.
- [23] Mauro De Sanctis, Ernestina Cianca, and Viraj Joshi. 2011. Energy Efficient Wireless Networks towards Green Communications. *Wireless Personal Communications* 59 (2011), 537–552.
- [24] Nassim Dehouche. 2015. A Bi-criteria Algorithm for Low-Carbon and QoS-Aware Routing in Cloud Computing Infrastructures. In *2015 International Conference on Cloud Computing (ICCC)*. IEEE, 1–5.
- [25] Dell. 2015. Dell Force10 S-Series. https://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell_Force10_S4810_Spec_sheet.pdf [Online, accessed June 27, 2023].
- [26] S. El-Zahr and N. Zilberman. 2023. CATE. <https://github.com/ox-computing/CATE>
- [27] S. El-Zahr and N. Zilberman. 2023. CATE. <https://doi.org/10.5281/zenodo.10035068>
- [28] National Grid ESO. [n. d.]. Live System Information. <https://data.nationalgrideso.com/system/live-system-information> [Online, accessed June 27, 2023].
- [29] C. Filsfil, Previdi Ed., Ed. S., L. Ginsberg, B. Decraene, S. Litkowski, and R. Shakir. 2018. RFC 8402: Segment Routing Architecture. <https://datatracker.ietf.org/doc/rfc8402/>
- [30] Fujitsu. 2015. Datasheet Brocade VDX 6940 Switch. https://www.fujitsu.com/sk/Images/ds_brocade_vdx_6940_ap.pdf [Online, accessed June 27, 2023].
- [31] British Gas. 2022. Where does British Gas energy come from? <https://www.britishgas.co.uk/energy/guides/energy-sources.html> [Online, accessed June 27, 2023].
- [32] Mirko Gattulli, Massimo Tornatore, Riccardo Fiandra, and Achille Pattavina. 2012. Low-Carbon Routing Algorithms for Cloud Computing Services in IP-over-WDM Networks. In *2012 IEEE International Conference on Communications (ICC)*. IEEE, 2999–3003.
- [33] GEANT. [n. d.]. Géant Network. <https://network.geant.org/> [Online, accessed June 27, 2023].
- [34] GEANT. 2020. Technology - Géant Network. <https://network.geant.org/technology/> [Online, accessed June 27, 2023].

- [35] GEANT. 2023. Géant Topology. <https://resources.geant.org/wp-content/uploads/2023/01/GN4-3N-Topology-Map-Jan2023.pdf> [Online, accessed June 27, 2023].
- [36] S. Giacalone, D. Ward, J. Drake, A. Atlas, and S. Previdi. 2015. OSPF Traffic Engineering (TE) Metric Extensions. <https://datatracker.ietf.org/doc/rfc7471/>
- [37] L. Ginsberg, Previdi Ed., Ed. S., S. Giacalone, D. Ward, J. Drake, and Q. Wu. 2019. IS-IS Traffic Engineering (TE) Metric Extensions. <https://datatracker.ietf.org/doc/rfc8570/>
- [38] Google. 2022. Google Environmental Report 2022. <https://www.gstatic.com/gumdrop/sustainability/google-2022-environmental-report.pdf> [Online, accessed June 27, 2023].
- [39] Google. 2023. Google Data Centers. <https://www.google.co.uk/about/datacenters/efficiency/> [Online, accessed June 27, 2023].
- [40] Maruti Gupta and Suresh Singh. 2003. Greening of the Internet. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (Karlsruhe, Germany) (SIGCOMM '03). Association for Computing Machinery, New York, NY, USA, 19–26. <https://doi.org/10.1145/863955.863959>
- [41] HP. 2014. HP 5900 Switch Series. <https://people.ucsc.edu/~warner/BuFs/c04111469.pdf> [Online, accessed June 27, 2023].
- [42] IEA. [n. d.]. IEA - Net Zero Emissions. <https://www.iea.org/topics/net-zero-emissions> [Online, accessed June 27, 2023].
- [43] IETF. 2023. Time-Variant Routing (tvr). <https://datatracker.ietf.org/wg/tvr/documents/> [Online, accessed June 27, 2023].
- [44] II-VI. 2022. II-VI Incorporated and ADVA Unveil Industry's First 100ZR Pluggable Coherent Transceiver for the Optical Network Edge. <https://ii-vi.com/news/ii-vi-incorporated-and-adva-unveil-industrys-first-100zr-pluggable-coherent-transceiver-for-the-optical-network-edge/> [Online, accessed October 10, 2023].
- [45] Mark Jackson. 2022. Broadband ISP BT Sees UK Network Traffic Peak at 28Tbps. <https://www.ispreview.co.uk/index.php/2022/07/broadband-isp-bt-sees-uk-network-traffic-peak-at-28tbps.html> [Online, accessed June 27, 2023].
- [46] Romain Jacob and Laurent Vanbever. 2022. The Internet of tomorrow must sleep more and grow old. In *1st Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*.
- [47] Seng-Kyoun Jo, Lin Wang, Max Mülhäuser, Young-Min Kim, and Jussi Kangasharju. 2017. Green Routing using Renewable Energy for IP Networks. In *2017 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*. IEEE, 1–6.
- [48] Christine E Jones, Krishna M Sivalingam, Prathima Agrawal, and Jyh Cheng Chen. 2001. A Survey of Energy Efficient Network Protocols for Wireless Networks. *Wireless Networks* 7 (2001), 343–358.
- [49] Abdallah Khreishah, Haythem Bany Salameh, Issa Khalil, and Ammar Gharaibeh. 2016. Renewable Energy-Aware Joint Caching and Routing for Green Communication Networks. *IEEE Systems Journal* 12, 1 (2016), 768–777.
- [50] Lenovo. 2015. Lenovo RackSwitch G8272. <https://people.ucsc.edu/~warner/BuFs/rackswitch-g8272-datasheet.pdf> [Online, accessed June 27, 2023].
- [51] Electricity Maps. [n. d.]. Live 24/7 co2 emissions of electricity consumption. <https://app.electricitymaps.com/map> [Online, accessed June 27, 2023].
- [52] 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.
- [53] Alireza Nafarieh, Yashar Fazili, and William Robertson. 2013. Dynamic Inter-Domain Negotiation for Green Algorithms in Optical Networks. *Procedia Computer Science* 21 (2013), 25–32.
- [54] United Nations. 2022. Climate Action. <https://www.un.org/en/climatechange/net-zero-coalition#:~:text=To%20keep%20global%20warming%20to,reach%20net%20zero%20by%202050.> [Online, accessed June 27, 2023].
- [55] Edge-Core Networks. 2015. AS5600-52X 10GbE Data Center Switch Bare-Metal Hardware. <https://people.ucsc.edu/~warner/BuFs/AS5600-52X.pdf> [Online, accessed June 27, 2023].
- [56] Juniper Networks. 2023. QFX3500 Switch. <https://people.ucsc.edu/~warner/BuFs/qfx3500-datasheet.pdf> [Online, accessed June 27, 2023].
- [57] Nokia. 2022. Nokia and BT Further Collaborate on Highly Scalable, Power Efficient IP Networks. <https://www.nokia.com/about-us/news/releases/2022/06/30/nokia-and-bt-further-collaborate-on-highly-scalable-power-efficient-ip-networks/> [Online, accessed June 27, 2023].
- [58] Nokia. 2023. Get Ready for the 800GE Reality. https://www.netnod.se/sites/default/files/2023-03/Nr.3_Jonas%20Vermeulen.pdf [Online, accessed June 27, 2023].
- [59] nPower. 2022. nPower - Our Fuel Mix. <https://npowerbusinessolutions.com/company/fuel-mix#> [Online, accessed June 27, 2023].
- [60] ns 3. [n. d.]. ns-3 | a Discrete-Event Network Simulator. <https://www.nsnam.org/> [Online, accessed June 27, 2023].

- [61] Ofgem. [n. d.]. Wholesale Market Indicators. <https://www.ofgem.gov.uk/energy-data-and-research/data-portal/wholesale-market-indicators> [Online, accessed June 27, 2023].
- [62] Nathaniel Akpofure Ojekudo and Nsikan Paul Akpan. 2017. Anapplication of Dijkstra’s Algorithm to shortest route problem. *IOSR Journal of Mathematics (IOSR-JM)* 13, 3 (2017), 14.
- [63] R. Pant, T. Russell, C. Zorn, E. Oughton, and J.W. Hall. 2020. *Resilience Study Research for NIC – Systems Analysis of Interdependent Network Vulnerabilities*. Technical Report. Environmental Change Institute, Oxford University, UK.
- [64] Yarin Perry, Felipe Vieira Frujeri, Chaim Hoch, Srikanth Kandula, Ishai Menache, Michael Schapira, and Aviv Tamar. 2023. DOTE: Rethinking (Predictive) WAN Traffic Engineering. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1557–1581.
- [65] Sergio Ricciardi, Davide Careglio, Francesco Palmieri, Ugo Fiore, Germán Santos-Boada, and Josep Solé-Pareta. 2011. Energy-Aware RWA for WDM Networks with Dual Power Sources. In *2011 IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
- [66] ScottishPower. 2022. Scottish Power - Where Does our Electricity Come from? <https://www.scottishpower.co.uk/about-us/performance/fuel-mix> [Online, accessed June 27, 2023].
- [67] Deepshikha Singh and Gitanjali Chandwani. 2015. Carbon-Aware Routing in Software Defined Inter Data Center Network. In *2015 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE, 1–6.
- [68] SSE. 2021. SSE - Sustainability Statement. <https://www.sse.com/media/u2jk4dgh/sse-hy-sustainability-2021-final.pdf> [Online, accessed June 27, 2023].
- [69] Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon-Intelligent Global Routing in Path-Aware Networks. (2023).
- [70] Telefonica. 2022. Management and Sustainability (ESG) Report 2022. <https://www.telefonica.com/en/shareholders-investors/financial-reports/integrated-annual-report/> [Online, accessed June 27, 2023].
- [71] Yuta Tokusashi, Huynh Tu Dang, Fernando Pedone, Robert Soulé, and Noa Zilberman. 2019. The case for in-network computing on demand. In *Proceedings of the Fourteenth EuroSys Conference 2019*. 1–16.
- [72] Karel van der Veldt, Cees de Laat, Inder Monga, Jon Dugan, and Paola Grosso. 2014. Carbon-Aware Path Provisioning for NRENs. In *International Green Computing Conference*. IEEE, 1–7.
- [73] 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*. 1–12.
- [74] Jiayuan Wang, Sergio Ricciardi, Anna V Manolova, Sarah Ruepp, Davide Careglio, and Lars Dittmann. 2012. OSPF-TE Extensions for Green Routing in Optical Networks. In *2012 17th Opto-Electronics and Communications Conference*. IEEE, 411–412.
- [75] Ning Wang, Kin Hon Ho, George Pavlou, and Michael Howarth. 2008. An Overview of Routing Optimization for Internet Traffic Engineering. *IEEE Communications Surveys & Tutorials* 10, 1 (2008), 36–56.
- [76] WattTime. [n. d.]. The Power to Choose Clean Energy. <https://www.watttime.org/> [Online, accessed October 8, 2023].
- [77] 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.
- [78] Yuan Yang, Mingwei Xu, Dan Wang, and Suogang Li. 2015. A Hop-by-Hop Routing Mechanism for Green Internet. *IEEE Transactions on Parallel and Distributed Systems* 27, 1 (2015), 2–16.
- [79] Mingui Zhang, Cheng Yi, Bin Liu, and Beichuan Zhang. 2010. GreenTE: Power-Aware Traffic Engineering. In *The 18th IEEE International Conference on Network Protocols*. IEEE, 21–30.
- [80] Noa Zilberman, Eve M Schooler, Uri Cummings, Rajit Manohar, Dawn Nafus, Robert Soulé, and Rick Taylor. 2022. Toward Carbon-Aware Networking. In *1st Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*.

A FURTHER SIMULATION RESULTS

A comparative summary of the simulation results for BT topology on a typical day in winter under both day-traffic and evening traffic are presented in Tables 2 and 3, respectively. The units of the parameters in the tables are kgCO₂ for all carbon values, kWh for all energy values, and millisecond for delay values. The overall carbon and energy improvement and the maximum link load are percentage values. The results for BT on a day in summer are presented in Tables 4 and 5. Similarly, results for GEANT topology are presented in Table 6.

Table 2. BT Results for Day-Traffic (Winter). Carbon D/P/S are the carbon emissions due to the dynamic/ports/static power portions, respectively. Same for Energy D/P/S notations. C Imp and E Imp are the overall improvement (%) in carbon and energy, respectively.

Metrics	Carbon D	Carbon P	Carbon S	Energy D	Energy P	Energy S	C Imp	E Imp	Delay (avg/p99)	Hops (avg/p99)	Max linkload
OSPF	65.45	27.79	54.75	524.53	223.98	438.19	0	0	1.6 / 6.76	3.53 / 5.0	12.28
Ptyp	64.48	27.79	54.75	513.49	223.98	438.19	0.66	0.93	1.63 / 7.02	3.54 / 5.0	12.42
Elabel	66.57	27.79	54.75	532.01	223.98	438.19	-0.75	-0.63	1.59 / 6.75	3.54 / 5.0	12.2
IncD	62.47	27.79	54.75	500.78	223.98	438.19	2.01	2	1.68 / 7.18	3.56 / 5.0	13.55
C	61.92	27.79	54.75	539.58	223.98	438.19	2.39	-1.27	1.93 / 9.22	3.78 / 6.0	66.77
C+Ptyp	61.96	27.79	54.75	535.7	223.98	438.19	2.36	-0.94	1.94 / 9.28	3.79 / 6.0	72.43
C+Elabel	62.46	27.79	54.75	543.67	223.98	438.19	2.02	-1.61	1.94 / 9.28	3.79 / 6.0	68.03
CE	62.76	27.79	54.75	548.5	223.98	438.19	1.82	-2.02	1.93 / 9.31	3.79 / 6.0	66.13
C+IncD	59.96	27.79	54.75	510.36	223.98	438.19	3.71	1.19	1.97 / 9.2	3.72 / 6.0	60.13
CATE	56.58	15.22	54.75	518.93	152.01	438.19	14.49	6.54	1.97 / 8.99	3.72 / 6.0	55.93

Table 3. BT Results for Evening-Traffic (Winter)

Metrics	Carbon D	Carbon P	Carbon S	Energy D	Energy P	Energy S	C Imp	E Imp	Delay (avg/p99)	Hops (avg/p99)	Max linkload
OSPF	284.88	25.17	49.63	2233.47	195.98	383.42	0	0	0.97 / 5.83	2.44 / 5.0	78.62
Ptyp	283.11	25.17	49.63	2221.3	195.98	383.42	0.49	0.43	0.98 / 6.01	2.44 / 5.0	78.64
Elabel	285.72	25.17	49.63	2239.25	195.98	383.42	-0.23	-0.21	0.96 / 5.81	2.44 / 5.0	78.98
IncD	281.56	25.17	49.63	2210.51	195.98	383.42	0.92	0.82	1.01 / 6.16	2.45 / 5.0	82.22
C	282.28	25.17	49.63	2243.27	195.98	383.42	0.72	-0.35	1.13 / 7.14	2.55 / 6.0	89.61
C+Ptyp	282.81	25.17	49.63	2245.29	195.98	383.42	0.57	-0.42	1.13 / 7.19	2.56 / 6.0	85.42
C+Elabel	283.72	25.17	49.63	2254.19	195.98	383.42	0.32	-0.74	1.13 / 7.1	2.56 / 6.0	85.44
CE	282.6	25.17	49.63	2242.85	195.98	383.42	0.63	-0.33	1.15 / 7.47	2.56 / 6.0	99.99
C+IncD	279.86	25.17	49.63	2215.21	195.98	383.42	1.39	0.65	1.14 / 7.26	2.51 / 5.0	79.46
CATE	279.52	24.02	49.63	2212.83	190.09	383.42	1.81	0.94	1.15 / 7.39	2.52 / 5.0	99.89

Table 4. BT Results for Day-Traffic (Summer).

Metrics	Carbon D	Carbon P	Carbon S	Energy D	Energy P	Energy S	C Imp	E Imp	Delay (avg/p99)	Hops (avg/p99)	Max linkload
OSPF	63.38	26.92	52.85	524.53	223.98	438.19	0	0	1.6 / 6.76	3.53 / 5.0	12.28
Ptyp	62.64	26.92	52.85	513.49	223.98	438.19	0.52	0.93	1.63 / 7.02	3.54 / 5.0	12.42
Elabel	64.49	26.92	52.85	532.01	223.98	438.19	-0.77	-0.63	1.59 / 6.75	3.54 / 5.0	12.2
IncD	60.75	26.92	52.85	500.78	223.98	438.19	1.84	2	1.68 / 7.18	3.56 / 5.0	13.55
C	61.24	26.92	52.85	533.14	223.98	438.19	1.5	-0.73	1.92 / 9.29	3.72 / 6.0	44.39
C+Ptyp	61.82	26.92	52.85	533.72	223.98	438.19	1.09	-0.77	1.92 / 9.31	3.74 / 6.0	47.84
C+Elabel	62.04	26.92	52.85	539.52	223.98	438.19	0.94	-1.26	1.91 / 9.31	3.73 / 6.0	43.16
CE	62.75	26.92	52.85	547.37	223.98	438.19	0.44	-1.92	1.91 / 9.35	3.75 / 6.0	44.02
C+IncD	59.26	26.92	52.85	506.48	223.98	438.19	2.88	1.52	1.96 / 9.24	3.68 / 6.0	42.93
CATE	56.32	15.33	52.85	507.67	152.01	438.19	13.03	7.49	1.94 / 8.89	3.7 / 6.0	50.1

Table 5. BT Results for Evening-Traffic (Summer).

Metrics	Carbon D	Carbon P	Carbon S	Energy D	Energy P	Energy S	C Imp	E Imp	Delay (avg/p99)	Hops (avg/p99)	Max linkload
OSPF	299.97	26.33	51.79	2233.47	195.98	383.42	0	0	0.97 / 5.83	2.44 / 5.0	78.62
Ptyp	298.23	26.33	51.79	2221.3	195.98	383.42	0.46	0.43	0.98 / 6.01	2.44 / 5.0	78.64
Elabel	300.85	26.33	51.79	2239.25	195.98	383.42	-0.23	-0.21	0.96 / 5.81	2.44 / 5.0	78.98
IncD	296.67	26.33	51.79	2210.51	195.98	383.42	0.87	0.82	1.01 / 6.16	2.45 / 5.0	82.22
C	297.3	26.33	51.79	2236.5	195.98	383.42	0.7	-0.11	1.11 / 6.92	2.52 / 5.0	80.33
C+Ptyp	297.91	26.33	51.79	2238.35	195.98	383.42	0.54	-0.17	1.1 / 6.89	2.53 / 6.0	80.82
C+Elabel	298.14	26.33	51.79	2242.48	195.98	383.42	0.48	-0.32	1.1 / 6.84	2.52 / 5.0	79.94
CE	297.66	26.33	51.79	2235.55	195.98	383.42	0.61	-0.07	1.13 / 7.42	2.53 / 6.0	100.05
C+IncD	295.22	26.33	51.79	2213.78	195.98	383.42	1.25	0.7	1.12 / 7.12	2.49 / 5.0	80.43
CATE	292.05	23.56	51.79	2205.92	180.66	383.42	2.83	1.52	1.25 / 7.74	2.69 / 5.0	100.02

Table 6. GEANT Results (Winter)

Metrics	Carbon D	Carbon P	Carbon S	Energy D	Energy P	Energy S	C Imp	E Imp	Delay (avg/p99)	Hops (avg/p99)	Max linkload
OSPF	9.98	1.44	4.52	37.2	5.51	16.98	0	0	44.67 / 154.26	3.86 / 9.0	41.89
C	8.86	1.44	4.52	36.66	5.51	16.98	6.99	0.89	42.73 / 148.85	4.0 / 10.0	73.76
CE	8.79	1.44	4.52	37.06	5.51	16.98	7.48	0.22	43.45 / 148.61	4.07 / 10.0	69.97
CATE	8.77	1.28	4.52	37.05	5.02	16.98	8.62	1.06	43.75 / 145.04	3.98 / 10.0	71.99

B ARTIFACT DESCRIPTION

The code used in this paper is available at [26, 27]. The provided artifact allows reproducing the results in Figures 9-16.

The simulation is run on an ASUS ESC4000A-E10 server with AMD EPYC 7302P 16-Core Processor with 256 GB memory. The operating system used is Ubuntu 20.04 and the software used are Python 3.8.10 and ns-3 version 3.36.1.

B.1 Data sets

Two network topologies are considered: BT and GEANT.

GEANT Topology: The files associated with GEANT topology are included in the repository under *GEANTFiles/* folder. It contains information about the topology connectivity map in INET format “Topology.txt” [35], the hourly carbon intensity during a day in winter in “winter.txt”, the normalized rate in “RateMean.txt” and the simulation parameters in “parameters.txt”.

BT Topology: The topology for BT network was obtained from [63], and users should contact its authors to request the topology. We include under *BTFiles/* folder in the repository all files used except for the connectivity map.

B.2 Installation

The following steps are needed to install the simulator and import the datasets and scripts:

- (1) Install the ns-3 simulator and configure it with optimized mode.
- (2) Clone the project repository under the directory /ns-allinone-3.36.1/ns-3.36.1/scratch.

B.3 Experiment workflow

The simulation in ns-3 builds the network based on the input topology (“Topology.txt” under *TOPOLOGYFiles/* folder) and then generates the traffic based on the input rate and pattern. The carbon intensity files are loaded (for example from “winter.txt”) and the energy metrics are loaded from the “parameters.txt” file. These values are used as metrics to change the cost of links when computing the routing tables in ns-3. Different metrics and combination of metrics are considered. The mapping of these metrics to the scenarios in the code is explained in the README file. Next, the simulation starts and the code monitors the traffic and the carbon emitted per interval of time

based on the carbon intensity and energy values provided. The simulation program outputs the detailed results of the simulation including the total carbon emissions and energy consumption and the detailed values per interval of time. Additional scripts are used to format the results and generate plots of the carbon emissions, energy consumption, CDF of delay and the flow intensity per region.

The code for CATE (additionally disables unnecessary links) is an extension of the aforementioned ns-3 program. It is included in the repository under *FullSimulation/LinkShutDown/* directory. It further reads the set of links to be shutdown and executes this operation at the start of every interval of time. The set of links to be disabled is generated by the *HeuristicTest.py* script placed in the *FullSimulation/Heuristic/* folder.