

# Carbon-Aware Global Routing in Path-Aware Networks

Seyedali Tabaeiaghdaei  
ETH Zürich

Jonghoon Kwon  
ETH Zürich

Simon Scherrer  
ETH Zürich

Adrian Perrig  
ETH Zürich

## ABSTRACT

The growing energy consumption of Information and Communication Technology (ICT) has raised concerns about its environmental impact. However, the carbon footprint of data transmission over the Internet has so far received relatively modest attention. This carbon footprint can be reduced by sending traffic over carbon-efficient inter-domain paths. However, challenges in estimating and disseminating carbon intensity of inter-domain paths have prevented carbon-aware path selection from becoming a reality.

In this paper, we take advantage of path-aware network architectures to overcome these challenges. In particular, we design CIRo, a system for forecasting the carbon intensity of inter-domain paths and disseminating them across the Internet. We implement a proof of concept for CIRo on the codebase of the SCION path-aware Internet architecture and test it on the SCIONLab global research testbed. Further, through large-scale simulations, we demonstrate the potential of CIRo for reducing the carbon footprint of endpoints and end domains: With CIRo, half of domain pairs can reduce the carbon intensity of their inter-domain traffic by at least 47%, and 87% of end domains can reduce their carbon footprint of Internet use by at least 50%.

## CCS CONCEPTS

• **Applied computing** → **Forecasting**; **Multi-criterion optimization and decision-making**; • **Networks** → **Network measurement**; **Network simulations**; **Network performance modeling**; **Data path algorithms**; **Control path algorithms**; Topology analysis and generation; • **Hardware** → **Renewable energy**.

## KEYWORDS

Green Networking, Internet Carbon-Emission Modeling and Measurement, Inter-Domain Routing, Carbon-Aware Routing, SCION

### ACM Reference Format:

Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon-Aware Global Routing in Path-Aware Networks. In *The 14th ACM International Conference on Future Energy Systems (e-Energy '23)*, June 20–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3575813.3595192>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*e-Energy '23, June 20–23, 2023, Orlando, FL, USA*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0032-3/23/06...\$15.00

<https://doi.org/10.1145/3575813.3595192>


## 1 INTRODUCTION

Growing concerns regarding climate change encourage companies to measure and reduce their carbon footprint, i.e., the amount of carbon emission that can be attributed to them. This also applies to their use of Information and Communication Technology (ICT), as ICT has a notable contribution of 2.7% to global CO<sub>2</sub> emissions [39], which is expected to grow significantly – approximately four times – until 2030 [3]. Hence, reducing the carbon footprint of ICT use is becoming increasingly relevant for enterprises, manifesting in carbon-neutrality statements of major technology corporations.

While these efforts are laudable and impactful, promising opportunities for further carbon-footprint reduction exist. Indeed, previous research has identified a range of such opportunities. However, most of these proposals apply to local aspects: intra-domain networking (i.e., within a single domain), data-center optimizations, or neighbor-domain cooperation (cf. §8). In contrast, inter-domain networking (i.e., among multiple domains), which accounts for around 13% of total ICT energy consumption, has so far received less attention. An exception is the work by Zilberman et al. [70], who identify carbon-aware networking as a high-potential research area and sketch the concept of carbon-intelligent routing, i.e., to leverage differences in network paths' carbon intensity (i.e., carbon emission per unit of data transmitted) to reduce the carbon footprint of communications.

Previous research on green inter-domain networking applies carbon efficiency to the optimization metric of the Border Gateway Protocol (BGP) [42]. Unfortunately, this direction faces several challenges. **Inefficient Green Route:** A strict carbon-optimal path can result in a highly inefficient end-to-end path in terms of monetary cost, latency, bandwidth, loss, or jitter (cf. §7.1). Depending on the application requirements, an optimization subject to all these constraints needs to be made, requiring path selection within a fine-grained metric space. **Ossification:** Carbon-optimal paths can thus only be offered as additional options, not as replacements for the conventional BGP route. When using BGP to provide carbon-efficient alternative paths, routers would thus require multiple forwarding tables, and packets would need to indicate the desired optimization criteria. Updating BGP and router hardware represents a challenge – as we have experienced in securing the BGP protocol through BGPSEC [37], which has been an effort for over two decades. **Volatility of Carbon Emissions:** Renewable energy sources (e.g., solar and wind) can fluctuate greatly within short time spans. The dynamic nature of carbon data used in routing metrics would introduce a large number of re-routing events, likely pushing BGP beyond its scalability limits.

This paper aims to overcome these challenges to enable carbon-aware inter-domain routing. Concretely, we pose two main research questions in this paper:



**Figure 1: The concept of carbon-aware inter-domain routing in a PAN architecture.**

- (1) How can we design a viable system that enables carbon-aware inter-domain routing?
- (2) How beneficial for endpoints can this system be in terms of the carbon footprint of their Internet usage?

To answer the first question, we identify a promising opportunity in Path-Aware Networking (PAN). PAN architectures provide endpoints with path information and allow them to choose paths based on their individual criteria. Thus, endpoints can optimize paths for multiple criteria, thus addressing the inefficient green route problem. Furthermore, some PANs employ the packet-carried forwarding state where packets convey the forwarding path information in the header field, making routers stateless and addressing the ossification problem.

One way to achieve carbon-aware routing in PANs is to provide endpoints with paths' carbon information and let them perform fine-grained path optimization, weighing carbon intensity against other performance criteria [14]. Fig. 1 illustrates carbon-aware inter-domain routing in a PAN architecture. To that end, a PAN architecture needs to be extended to provide carbon information in a practical way. More precisely, path carbon intensity needs to be estimated and disseminated in a way that (1) provides up-to-date information, (2) is scalable despite the volatility of carbon intensity. We discuss these requirements in more detail in §3.1, and (3) does not require autonomous systems (AS<sup>1</sup>) to disclose their topology and electricity providers.

In this paper, we design CIRo (Carbon-aware Inter-domain Routing), a system that forecasts and disseminates carbon intensity of inter-domain paths in a PAN architecture. It satisfies all the aforementioned requirements by *forecasting* day-ahead carbon intensity of inter-domain paths in a *distributed* manner and disseminating the data via the path distribution mechanism of PAN (cf. §3.2 and §5). To compute these forecasts, CIRo uses our new model for the carbon intensity of inter-domain paths (cf. §4).

To prove that scalable green routing can be deployed and used in the short term, we build CIRo based on the commercially deployed SCION PAN [13, 34] architecture, implement a proof of concept on the open-source SCION codebase [51], and operate it on the SCIONLab [35] testbed. We anticipate that a concrete open-source system provides an important baseline for future research to extend and improve upon.

To answer the second question of the paper, i.e., the carbon-footprint benefits for endpoints and end domains, we carry out

an investigation by simulating CIRo on a large-scale realistic Internet topology. This topology reflects environmentally relevant characteristics of today's Internet and allows for the simulation of carbon-aware path selection. Our simulations show that CIRo can reduce the carbon intensity of inter-domain traffic by at least 47% between half of the domain pairs. Further we show that 87% of end domains can benefit from at least 50% reduction in the footprint of their Internet usage, i.e., the amount of carbon emission that can be attributed to them for their Internet use.

Importantly, the reduction in the carbon footprint of end domains only means that less carbon emission can be attributed to individual end domains, and it does not necessarily translate to an actual reduction in carbon emission. The high-carbon footprint paths that are avoided still cause carbon emissions due to their idle power consumption. However, the shift of traffic of environmentally conscientious endpoints from polluting paths to greener ones is expected to eventually result in actual reductions in carbon emissions through competition among ISPs that is enabled by CIRo: to attract more traffic and increase revenue, ISPs enhance their carbon efficiency to entice environmentally conscientious endpoints to send traffic through their networks, which will, in turn, encourage other ISPs to do the same. This phenomenon is studied in detail in a companion publication [49].

**Contributions.** This paper makes the following contributions:

- Propose a carbon-intensity model for inter-domain paths;
- design CIRo, a system to forecast the carbon intensity of inter-domain paths and disseminate it to endpoints;
- implement a proof of concept for CIRo on SCION's open-source codebase and operate it on the SCIONLab testbed;
- investigate the impact of carbon-aware inter-domain path selection on the carbon footprint of end domains by implementing CIRo in the ns-3-based SCION simulator and simulating it on a large-scale topology.

## 2 BACKGROUND

As CIRo is designed and implemented based on the SCION PAN architecture, we provide a brief overview of SCION. Furthermore, to compute the carbon-intensity forecast of network paths, we employ power-grid carbon-intensity forecasting tools, which we briefly present.

### 2.1 PAN Architectures and SCION

PAN architectures provide endpoints with information about different network paths, and let them select paths according to provided information. Thus, PAN architectures introduce a promising opportunity for carbon-aware path selection. SCION is a deployed PAN architecture that can help concretize the principles of carbon-aware inter-domain networking throughout this paper. We next present the SCION features that are directly relevant for the remainder of the paper; a comprehensive description of the architecture can be found in dedicated works [13, 34].

**Architecture.** SCION groups ASes in *Isolation Domains* (ISDs). In each ISD, *core ASes* provide connectivity to other ISDs. Other ASes in an ISD are direct or indirect customers of core ASes.

**Data Plane.** SCION uses packet-carried forwarding state to forward

<sup>1</sup>In this paper, we use AS and domain interchangeably.

inter-domain traffic. This means that endpoints encode AS-level inter-domain paths into packet headers, based on which border router of every AS forwards packets. SCION paths are specified at the granularity of ingress and egress interfaces representing links between neighboring ASes, providing endpoints with fine-grained control over paths. Border routers do not need to store inter-domain forwarding tables to make forwarding decisions, enabling scalable multi-path forwarding and allowing path optimization across multiple criteria, including carbon intensity.

**Control Plane.** SCION constructs and disseminates inter-domain paths using a hierarchical control plane with two levels: (1) among all core ASes, and (2) within every ISD along provider-customer links. Each level of the hierarchy is responsible for constructing path segments and disseminating them to endpoints, which combine these segments into complete end-to-end forwarding paths.

Path segments are constructed through the *beaconing* process, a collaborative hop-by-hop process among *beacon services* of participating ASes that involves exchanging special routing messages called Path Construction Beacons (PCBs). The two hierarchy levels of beaconing are: (1) *core beaconing* among all core ASes in all ISDs, constructing core-path segments, and (2) *intra-ISD beaconing* within each ISD along provider-customer links, constructing up- and down-path segments. During the beaconing process, each AS encodes information about its AS hop in an ASEntry in the PCB. An ASEntry can contain a StaticInfoExtension field [20], providing path metadata.

Beacon services extract path segments from PCBs and register them to *path services*, responsible for resolving endpoints' queries for path segments. To construct an end-to-end forwarding path, an endpoint retrieves path segments from the path service and selects a desirable combination of up, core, and down segments based on their metadata.

## 2.2 Forecasting Power-Grid Carbon Intensity

The carbon intensity of network paths depends on the Carbon Intensity of the Electricity (*CIE*, i.e., CO<sub>2</sub> mass emitted per unit of generated electrical energy in g/kWh) that network devices consume. Therefore, to forecast the carbon intensity of network paths, we need *CIE* forecasts, which is an emerging research area. Commercial tools such as ElectricityMaps [18], WattTime [67], and research projects such as DACF [40] are example systems that provide short-term *CIE* forecasting. These tools provide day-ahead hourly forecasts of *CIE*, i.e., a sliding window of forecast values with one-hour update frequency.

## 3 DESIGN PRINCIPLES

This section describes the design requirements and challenges of a system that estimates and disseminates carbon-intensity of inter-domain paths in a path-aware Internet and provides an overview of CIRo, which overcomes these challenges.

### 3.1 Requirements and Challenges

We identify two requirements for estimating and disseminating network paths' carbon intensity across the Internet.

**R1: Reliable Carbon Information.** The carbon intensity of network paths must be quantified in an accurate and reliable man-

ner. This characterization is complex because the carbon intensity of a network path varies over time and depends on the location of all devices on the path, and their electricity mix, which in turn depends on the daytime and the weather conditions.

**R2: Scalable Dissemination.** Carbon-intensity information must be disseminated in a scalable way. This is challenging as paths' carbon intensity fluctuates and needs to be disseminated frequently. However, frequent dissemination of information about numerous Internet paths can overwhelm participating ASes in terms of communication and computation costs.

**R3: Not Requiring ASes to Disclose their Topology.** The dissemination of carbon intensity information of ASes must not require them to disclose their internal topology or electricity providers. However, accurate computation of an inter-domain path's carbon intensity requires fine-grained information about all devices on the path across multiple ASes.


### 3.2 System Overview

We design CIRo such that it satisfies requirements R1–R3 above. We satisfy the reliability requirement (R1) by proposing a model for the carbon intensity of inter-domain paths and using forecasting. The scalability requirement (R2) is met by providing forecasts instead of real-time information to avoid frequent re-dissemination upon changes. Lastly, CIRo does not require ASes to disclose their topology or their electricity providers (R3) through a distributed design with instances in all participating ASes: each AS deploys a local instance of CIRo that forecasts and disseminates the carbon intensity of intra-domain paths without disclosing detailed information about the AS's topology. The carbon intensity of an *inter-domain* path is then calculated by accumulating the carbon intensities of all its constituent *intra-domain* segments, each traversing one AS (cf. §4).

Each CIRo instance consists of two modules: (1) the forecasting module, and (2) the information dissemination module. The interaction of these modules is illustrated in Fig. 2.

The **forecasting module** computes day-ahead hourly forecasts (i.e., forecast for 24 hours) for the carbon intensity of intra-domain paths that connect interfaces of the AS. These forecasts are then inserted into the forecast database. The main building block of the forecasting module is the model we propose for the Carbon Intensity of Data Transmission over inter-domain paths (*CIDT*, i.e., CO<sub>2</sub> mass emitted per unit of traffic being forwarded on a path in g/bit, cf. §4). This model takes into account the Energy Intensity of Data Transmission for network devices (*EIDT*, i.e., their energy consumption per bit of forwarded traffic in kWh/bit) and the carbon intensity of their used electricity (*CIE*, cf. §2.2).

The **information dissemination module** disseminates the carbon intensity forecasts across the Internet. We design this module as an extension to the control-plane infrastructure of a PAN architecture. We extend the functionality of the control plane to encode carbon-intensity forecasts within routing messages. As a result, endpoints receive this information along with paths without requiring additional queries, reducing standardization efforts as well as communication and computation overheads.



**Figure 2: Overview of CIRo.** A participating AS runs a CIRo instance. An instance consists of a forecasting module and an information dissemination module. The forecasting module computes path carbon intensity forecasts using a model we introduce, the topology of the AS, and power-grid carbon-intensity (CI) forecasts, and inserts its forecasts into the forecast database. The dissemination module disseminates forecasts to other ASes using routing messages. It also provides endpoints within an AS with these forecasts.

## 4 MODELING CARBON INTENSITY OF INTER-DOMAIN PATHS

In this section, we develop a model enabling CIRo to compute and disseminate carbon-intensity forecasts of inter-domain paths in a distributed manner that does not require ASes to disclose their internal topology.

### 4.1 Carbon Intensity of Data Transmission

Data transmission over a network path may cause CO<sub>2</sub> emission through two levels of indirection: forwarding traffic causes (electrical) energy consumption on network devices, and the electricity used by each device is generated from energy resources in a process that may emit CO<sub>2</sub>, depending on the energy resources.

We define **Carbon Intensity of Data Transmission** over a network path, i.e., *CIDT* in g/bit, as the CO<sub>2</sub> emission that can be attributed to a unit of data for being transmitted over that path. Since energy is an additive quantity, the CO<sub>2</sub> emission of energy consumption and the CO<sub>2</sub> emission of data transmission are additive.

### 4.2 Inter-Domain CIDT

Using CIDT's additive property, we write the CIDT of an inter-domain path (denoted by  $\{AS^{src}, \dots, AS^{dst}\}$ ) as the sum of CIDTs of consecutive AS hops it consists of:

$$CIDT_{\{AS^{src}, \dots, AS^{dst}\}} = \sum_{AS_{[ing, eg]}^i \in \{AS^{src}, \dots, AS^{dst}\}} CIDT_{AS_{[ing, eg]}^i}, \quad (1)$$

where  $AS_{[ing, eg]}$  denotes an AS hop on the inter-domain path.  $[ing, eg]$  denotes the ingress and egress interface pair from/to which the inter-domain path enters/exits the AS hop.

The additive property is essential for distributed functionality of CIRo: each CIRo instance only needs to compute the CIDT of paths within one AS without disclosing topology information.

Eq. (1) does not explicitly include inter-domain *links* connecting the border routers of neighboring ASes. That is because such links are either (1) direct links connecting two routers in the same data center, or (2) peerings via IXPs, IXP federations, IXP port resellers, or layer-2 links to an IXP. In the first case, the preceding AS on a path includes the CIDT of the link in its hop's CIDT. In the second case, previous research suggests IXPs and resellers constitute ASes in a PAN context [34]. Thus, they compute their CIDT as normal ASes in the same way we propose in this section.

### 4.3 Per-Hop CIDT

An AS may forward traffic on multiple internal paths, e.g., using equal-cost multi-path (ECMP). Therefore, we define the CIDT of  $AS_{[ing, eg]}$  as the mean CIDT of all active *intra-domain* paths from its ingress interface to its egress interface. Thus,

$$CIDT_{AS_{[ing, eg]}} = \frac{1}{|\mathcal{P}_{[ing, eg]}|} \sum_{P_{[ing, eg]} \in \mathcal{P}_{[ing, eg]}} CIDT_{P_{[ing, eg]}} \quad (2)$$

where  $P_{[ing, eg]}$  is an *intra-domain* path connecting *ing* and *eg* interfaces within the AS, and  $\mathcal{P}_{[ing, eg]}$  is the set of all such paths. In case of weighted ECMP (WECMP), Eq. (2) is substituted with a weighted mean over all paths.

### 4.4 CIDT of a Single Intra-Domain Path

The CIDT over a network path within an AS is composed of two components: (1) the marginal CIDT, which is the result of the actual amount of energy consumed to forward a bit of data over a path, and (2) amortized CIDT, which is the result of energy consumed to keep network paths operational and is independent of traffic volume over the path. Therefore,

$$CIDT_{P_{[ing, eg]}} = CIDT_{P_{[ing, eg]}}^{marginal} + CIDT_{P_{[ing, eg]}}^{amortized} \quad (3)$$

Because of the additive property, we can write

$$CIDT_{P_{[ing, eg]}} = \sum_{D \in P_{[ing, eg]}} CIDT_D^{marginal} + CIDT_D^{amortized} \quad (4)$$

where  $D$  denotes a network device on  $P_{[ing, eg]}$ . Since inter-domain communication mostly takes place over backbone networks, we consider typical backbone-network devices in this work, e.g., IP/MPLS core routers and optical wavelength-division multiplexed (WDM) devices [1, 25].

**4.4.1 Marginal CIDT.** For each device  $D$ ,  $CIDT_D^{marginal}$  is the result of (1) the marginal energy the device consumes to forward one bit of data (or marginal **Energy Intensity of Data Transmission**,  $EIDT_D^{marginal}$  in kWh/bit), and (2) the marginal per-bit energy consumption of the device's external cooling and facilities, which is proportional to  $EIDT_D^{marginal}$  [25]. Thus, the total marginal energy is also proportional to  $EIDT_D^{marginal}$  by a constant factor ( $\eta_{D,c}$ ) determined by the power-usage efficiency (PUE) [7, 25].

To compute  $CIDT_D^{marginal}$ , we multiply the total marginal energy by the *CIE* (in g/kWh cf. §2.2) in the location of that device ( $CIE_D$ ). Therefore,

$$CIDT_D^{marginal} = \eta_{D,c} EIDT_D^{marginal} CIE_D \quad (5)$$

$EIDT_D^{marginal}$  can be computed using the following formula [27], irrespective of device type:

$$EIDT_D^{marginal} = \frac{1}{3600} \frac{P_{D,max} - P_{D,idle}}{C_{D,max}} \quad (6)$$

where  $P_{D,max}$  is the maximum power consumption of the device (in kW),  $P_{D,idle}$  is its idle power consumption (in kW), and  $C_{D,max}$  is its maximum capacity (in bps). Note that 1 kWh/bit = 3600 kW/bps.

**4.4.2 Amortized CIDT.** Equivalently to the marginal carbon intensity, the amortized carbon intensity  $CIDT_D^{amortized}$  of a device  $D$  is determined by the amortized idle energy consumption  $EIDT_{D,idle}^{amortized}$ , scaled by the PUE factor  $\eta_{D,c}$ . In addition, the model takes into account the set  $\mathcal{R}(D)$  of *redundant* devices associated with device  $D$ , i.e., devices enabled when the primary device  $D$  becomes unavailable. We assume that these redundant devices do not forward any traffic. As a result, we can write:

$$CIDT_D^{amortized} = \eta_{D,c} EIDT_{D,idle}^{amortized} CIE_D + \sum_{D' \in \mathcal{R}(D)} \eta_{D',c} EIDT_{D',idle}^{amortized} CIE_{D'} \quad (7)$$

Since redundant devices can be located in a different location from the primary device, they could have different  $CIE$ s.

To amortize the idle energy consumption of a device over a bit of data crossing the device, we hold that bit of data accountable for the device's total *idle* energy consumption during its processing time. The capacity of the device (in bps) determines the processing time: On average, processing of a bit on a device takes time  $\frac{1}{C_{D,max}}$  (in bit/bps = s). By multiplying the idle power consumption of the device by this duration, we arrive at the idle energy consumption of the device during this interval:

$$EIDT_{D,idle}^{amortized} = \frac{1}{3600} \frac{P_{D,idle}}{C_{D,max}} \quad (8)$$

In case of redundant devices, we substitute their  $P_{D',idle}$  in Eq. (8). However, since the *primary* device forwards the traffic, the processing time is still determined by the maximum capacity of the *primary* device. Thus,

$$EIDT_{D',idle}^{amortized} = \frac{1}{3600} \frac{P_{D',idle}}{C_{D,max}} \quad (9)$$

In §7.3, we discuss another method for amortizing idle energy consumption.

## 4.5 Final Form of Per-Hop CIDT

We combine Eqs. (2)–(9) and derive the following formula for the CIDT of a path within an AS between an interface pair:

$$CIDT_{AS[ing, eg]} = \frac{1}{3600 \cdot |\mathcal{P}[ing, eg]|} \sum_{P[ing, eg] \in \mathcal{P}[ing, eg]} \sum_{D \in P[ing, eg]} \left( \eta_{D,c} \frac{P_{D,max}}{C_{D,max}} CIE_D + \sum_{D' \in \mathcal{R}(D)} \eta_{D',c} \frac{P_{D',idle}}{C_{D,max}} CIE_{D'} \right) \quad (10)$$

In Eq. (10), the only time-variant variable is  $CIE$ , which depends on the availability of renewable electricity, which in turn depends on the weather conditions and the time of day. As a result,  $CIDT_{AS[ing, eg]}$  is also time-variant.

## 5 CIRO: DESIGN DETAILS

This section describes the design details of the forecasting and information dissemination modules, for which we implement a proof of concept on the SCION codebase [51], available on GitHub [58], and run it on the SCIONLab [35] testbed.

### 5.1 Forecasting Module

The forecasting module computes day-ahead hourly CIDT forecasts between interfaces of each AS. This module is distributed across participating ASes, where each instance is administered by one AS. Thus, it does not disclose internal information about the AS to any external entity.

**5.1.1 Required Inputs.** Each instance of the forecasting module performs its predictions using Eq. (10), based on the following AS information.

#### Topology Information:

- (1) *AS Interfaces*: Set of AS interfaces and the corresponding border routers.
- (2) *Intra-Domain Paths*: Device-level intra-domain paths connecting border routers of the AS, computed by intra-domain routing protocols and listed in routers' intra-domain routing information base (RIB).

#### Device Information:

- (1) *Device Locations*: Locations of all network devices (both primary and redundant) of the AS, including routers and optical devices. If an AS leases optical fibers from other companies, the location and redundancy provisioning of optical devices might not be available. In that case, the module uses models [25] of IP over WDM networks to approximate locations. Moreover, the same model suggests one redundant device identical to and in the same place as the primary device.
- (2) *Device Specifications*: Maximum power consumption ( $P_{max}$ ), maximum capacity ( $C_{max}$ ), and idle power consumption ( $P_{idle}$ ) of all network devices in the AS network, available in device specifications. If this information is unavailable for device  $D$  (primary or redundant), the module substitutes  $\frac{P_{D,max}}{C_{D,max}}$  in Eq. (10) with typical values for the energy intensity of network devices [25] (cf. Table 2 in Appendix C).

#### Electricity Information:

- (1) *Power-Usage Efficiency (PUE)*: Power-usage efficiency at all AS's points of presence (PoPs), determining  $\eta_{D,c}$ . If unknown, the module uses a typical value of 2 [25].
- (2) *Carbon Intensity of Electricity (CIE)*: Day-ahead hourly forecast of  $CIE$  at all PoPs from electricity carbon-intensity forecasting sources (cf. §2.2).

**5.1.2 Forecasting Procedure.** A periodic procedure retrieves updated inputs and re-computes day-ahead hourly CIDT forecasts for all interface pairs. The periodic re-computation is necessary as (1)  $CIE$  forecasts are updated with a period of one hour, and (2) paths between interfaces could change even within these 1-hour intervals. Thus, we consider two re-computation periods:  $T_{CIE}$ , which is set to 1 hour to retrieve updated  $CIE$  forecasts at the beginning of every hour (\*\*:00), and  $T_{path}$  to collect updated intra-domain RIBs from routers.  $T_{path}$ 's value depends on the frequency of path changes in



one AS and is set by each AS independently.  $T_{path}$  is only relevant if it is less than  $T_{CIE}$ .

At each  $T_{CIE}$  (i.e., beginning of every hour), the module first retrieves all inputs mentioned in §5.1.1. Then, it computes the day-ahead hourly CIDT forecast for every pair of interfaces in both directions (i.e., both could be ingress or egress interfaces) using Eq. (10). The result for each interface pair in each direction is a vector of 24 CIDT values associated with the next 24 hours. Finally, the module inserts the results into the *CIDT database* shared with the information dissemination module.

At each  $T_{path}$ , the module checks whether intra-domain paths between every interface pair have changed compared to the last  $T_{path}$ . If so, the module computes CIDT forecasts for the new path and updates the database accordingly.

The database stores two records for each interface pair, one per direction. Each record has 26 columns; the first two columns store the ingress and the egress interface, respectively, and the following 24 columns store CIDT forecasts for the next 24 hours associated with that pair in the direction from the ingress interface to the egress interface.

## 5.2 Information Dissemination Module

This module uses the routing infrastructure (control plane) of a PAN architecture to disseminate CIDT forecasts. It disseminates forecasts across ASes along with path construction by the means of a routing-message extension we introduce. Endpoints can then receive this information when they request paths from the path service of their ASes.

The benefits of building this module based on already existing infrastructure are two-fold: (1) low message complexity and communication overhead as it does not send additional messages dedicated to CIDT, and (2) low design and implementation effort by re-using existing protocols.

In Appendix B, we develop a concrete prototype of the design we propose in this section by tuning it to work based on the control plane of the deployed SCION PAN architecture. Using SCION's control plane further improves this module in terms of (1) scalability to large topologies due to SCION's hierarchical routing architecture [34], and (2) providing up-to-date carbon information due to the periodic routing in SCION. However, this hierarchical routing only gives CIDT forecasts for path segments, not the entire path. We address this challenge in Appendix B, which is also applicable to other hierarchical routing architectures.

We now introduce CIDTExtension, an extension to routing messages to disseminate CIDT forecasts. Later, we explain how ASes use this extension to convey the CIDT forecasts of their paths during the routing process.

**5.2.1 CIDTExtension.** Fig. 3 demonstrates the structure of this extension within a routing message. Every AS can include one CIDTExtension in its AS Hop to disseminate its CIDT forecasts.

A CIDTExtension provides CIDT forecasts for both the forward path from ingress to egress interface ( $CIDT_{AS[ing, eg]}$ ) and the backward path from egress to ingress interface ( $CIDT_{AS[eg, ing]}$ ). It is a sequence of 48 contiguous unsigned 8-bit integer numbers, each specifying one CIDT value in mg/Gbit. We choose this unit based on our observation of CIDT range in Internet paths (cf. §6.2). The




Figure 3: A routing message containing CIDTExtension.

first 24 values of the vector represent  $CIDT_{AS[ing, eg]}$  forecasts for the 24-hour interval starting from the greatest hour before the routing message's TimeStamp. The following 24 values provide similar information for the reverse path ( $CIDT_{AS[eg, ing]}$ ). This design ensures the time-alignment of forecasts added by different ASes at different times.

For hierarchical routing architectures like SCION, where path segments are combined, CIDTExtension is a map from relative interface identifiers to CIDT forecast vectors, and the structure is different in each routing hierarchy level (cf. Appendix B).

**5.2.2 Timing of CIDT Dissemination.** Since CIDT forecasts have a 24-hour time window, they need to be disseminated regularly to preserve the temporal continuity of CIDT information.

This periodic dissemination, however, does *not* require equal frequency or synchronization among ASes as all CIDTExtensions in one routing message are time-aligned.

Therefore, every AS locally chooses its CIDT dissemination period, with a lower bound of 1 hour and an upper bound of 2.6 hours. The lower bound is derived from the update period of CIDT forecasts by the forecasting module, i.e., 1 hour, making it unnecessary to disseminate CIDT forecasts more frequently. The upper bound is required to guarantee the temporal continuity of CIDT forecasts. It is derived for the worst case scenario, where all ASes on a path segment synchronously disseminate routing messages with the same frequency. In that case, the temporal continuity is satisfied when the last AS on the path segment receives routing messages with CIDT forecasts that are valid for at least one period. Therefore, the upper bound is  $\frac{24}{n-1}$ , where  $n$  is the number of AS hops on the path segment. Assuming the longest path segment is 10 AS hops long (twice as long as the average AS-hop length in today's Internet [28]), the upper bound period is 2.6.

**5.2.3 Process of CIDT Dissemination.** When disseminating CIDT forecasts, an AS first selects a set of received routing messages per origin AS (i.e., initiator of the routing message). In this selection, the newest routing messages and the ones with the highest ratio of AS Hops with CIDTExtension are prioritized to maximize the temporal and spatial continuity of CIDT information.

Then, for each selected routing message and each egress interface (EgIntf), the AS constructs the corresponding CIDTExtension, adds it to the AS Hop, extends the routing message with the AS Hop and disseminates it on the egress interface to the next-hop AS.

To construct the CIDTExtension for each routing message, the dissemination module first queries the *CIDT database* for CIDT forecasts between IngIntf and the EgIntf in both directions. However, the retrieved results may not be time-aligned with the TimeStamp of the routing message. This is because the database always pro-

vides CIDT forecasts for the next 24 hours from the current time, but the CIDTExtension has to provide CIDT forecasts for the next 24 hours after the routing message's TimeStamp, which is in the past. To solve this problem, it shifts the vectors retrieved from the database to the right by the number of hours past from the TimeStamp's hour.

## 6 CIRO'S IMPACT ON CARBON FOOTPRINT

This section answers the second research question of the paper: To what extent could carbon-aware inter-domain path selection reduce the carbon footprint of the Internet usage of endpoints and end domains (end ASes, i.e., source or destination of the traffic)? This potential is critical to the attractiveness of CIRO, and in turn to the prospect, span, and pace of CIRO's deployment. We describe our method in §6.1, and we present our results in §6.2.

### 6.1 Methodology

Endpoints could benefit from carbon-aware path selection by sending their traffic over paths with lower CIDT (compared to the BGP path in the traditional Internet). This benefit can be found by computing the CIDT difference between CIRO and BGP paths and multiplying this difference by the traffic volume between domain pairs.

Since SCIONLab is an overlay network with a modest number of virtual ASes, it does not resemble real Internet paths. Therefore, we cannot rely on our proof-of-concept implementation on SCIONLab to approximate real CIDT of paths in the Internet. Thus, we develop a method based on simulations and large-scale realistic datasets. Our method consists of the following steps:

- (1) Approximation of the Internet topology (§6.1.1),
- (2) reconstruction of intra-domain paths (§6.1.2),
- (3) carbon-intensity estimation for intra-domain paths (§6.1.3),
- (4) reconstruction of BGP paths for comparison (§6.1.4),
- (5) discovery of CIRO paths (§6.1.5), and
- (6) synthesis of a traffic matrix (§6.1.6).

**6.1.1 Creating an Inter-Domain Topology.** To find realistic inter-domain paths, we use the inter-domain topology suggested by Krähenbühl et al. [34], which contains a total of 2000 ASes. These are the highest-degree Tier-1 and Tier-2 ASes of the CAIDA AS-Rel-Geo dataset [9]. We extract this topology from the AS-Rel-Geo dataset by iteratively removing the lowest-degree ASes. The benefits of CIRO are most likely realized within these top-tier ASes, as the main differences in the CIDT of paths stem from path diversity provided by massively interconnected ASes in the Internet core.

**6.1.2 Finding Intra-Domain Paths.** To compute CIDT of inter-domain paths, we need the CIDT of intra-domain segments they consist of. In turn, the CIDT of intra-domain segments relies on finding intra-domain paths. Since intra-domain paths are not publicly available, we compute them by applying Dijkstra's algorithm [16] on the internal topology of each AS.

First, we locate border routers associated with AS interfaces, where the interfaces are available from the AS-Rel-Geo dataset [9]: For each interface between AS  $A_1$  and  $A_2$ , we identify all routers of AS  $A_1$  that are also connected to AS  $A_2$ , and select the router closest to the interface. This information is available via the CAIDA ITDK dataset [10]. If no routers satisfy these conditions, we add

a router at the interface location. Given these border routers, we compute the shortest router-level path between every pair of border routers by running the Dijkstra algorithm [16] on the intra-domain topology of each AS, given by the ITDK dataset. If we cannot find such a path, we assume that the number of routers is an increasing integer-valued step function of the distance between interfaces with the range of  $\{1, 2, 3, 4, 5\}$  and steps at 1 km, 20 km, 100 km, and 1000 km distance.

Once router-level paths are established, we add optical devices between consecutive routers. As each packet crosses every WDM switch, transponder, and muxponder once before entering a router and once after leaving it, we assume two of these devices per router on the path [25]. The number of amplifiers and regenerators, however, depends on the distance between consecutive routers. We assume one amplifier at each 80 km interval, and one regenerator at each 1500 km interval [25].

**6.1.3 Estimating CIDT of Intra-Domain Paths.** The CIDT of intra-domain paths is necessary to compute the CIDT of inter-domain paths. Thus, we compute  $CIDT_{AS[ing, eg]}$  for all ingress and egress interface pairs of all ASes in the topology using Eq. (10) and computed paths in §6.1.2.


However, not all the required information by Eq. (10) is available. First, we do not know using which simultaneous internal paths ASes achieve multi-path routing. Second, we do not know device specifications, the number and location of redundant devices per primary device, and the PUE at PoPs of each AS. Third, hourly CIE forecasts are not available for all device locations.

To address the first issue, we assume single-path intra-domain routing in simulations, as opposed to the real deployment of CIRO. To solve the second problem, we use the model Heddeghem et al. [25] propose. They suggest typical values for the energy intensity of various types of backbone network devices (cf. Table 2 in Appendix C). They also propose the typical value of 2 for PUE at each PoP. Further, for each primary device, they consider one identical redundant device. We further assume that both devices are at the same location. To overcome the third challenge, we use each country's *annual average CIE* as the  $CIE_D$  of all devices located in that country. We compute the CIE of countries using their mix of electricity-production technology [29], and the carbon intensity of electricity-production technologies [17]. To find the location of routers, we use the CAIDA ITDK dataset [10].

**6.1.4 Finding BGP Paths and their CIDT.** As not all ASes publish their routing tables, we simulate BGP using the SimBGP simulator [62] to find the BGP-selected paths. Then, we compute their CIDT as the sum of CIDTs of all their intra-domain segments.

**6.1.5 Finding Green Inter-Domain Paths.** To compute green inter-domain paths in the mentioned topology, we implement CIRO in the ns-3-based [60] SCION simulator [34, 56]. Furthermore, we change the beaconing (routing) algorithm such that all ASes optimize the CIDT of paths, i.e., they advertise paths with the lowest CIDT. However, as this algorithm is heuristic, it cannot guarantee optimality.

We use simulations for practicality as setting up a large-scale testbed is not feasible in terms of cost and effort. Further, Krähenbühl et al. [34] have shown that the SCION simulator is highly



**Figure 4: Distribution of path CIDT across AS pairs, distinguished by path types: (1) the greenest path discovered by CIRO, (2) the greenest BGP path among the  $k$  BGP paths in RIBs of the source AS's border routers,  $k$  being the number of border routers, (3) the average of available BGP paths, and (4) the average of the  $k$ -greenest paths discovered by CIRO.**

predictive of actual testbed results in terms of discovered paths, so simulation does not affect the reliability of results.

**6.1.6 Synthesizing a Traffic Matrix.** We synthesize a traffic matrix between all ASes using the model proposed by Mikians et al. [41]. We only consider HTTP(S) data and media, and popular video streaming services as they contribute to 60% of the Internet's traffic [47]. Details are available in Appendix E.

## 6.2 Results

This section presents the result of our study on the effect of carbon-aware inter-domain path selection on the carbon footprint of end domains. We first compare the CIDT of greenest available paths with BGP paths. Then, we present the absolute and relative reductions in CIDT of communication between each pair of ASes. Finally, we demonstrate the reduction in the carbon footprint of end ASes.

Note that since the granularity of the traffic matrix is end ASes, not endpoints, we can only study the impact on CIDT (i.e., the intensity of carbon footprint) of end domains.

**6.2.1 Path CIDT Comparison.** Fig. 4 shows the CIDT distribution of different types of paths available between AS pairs. The results indicate that CIRO-paths almost halve CIDT for almost all percentiles of AS pairs (and their endpoints). It also suggests that CIRO can find *multiple* paths with significantly lower CIDTs than BGP paths.

**6.2.2 Absolute CIDT Reduction.** Fig. 5a demonstrates the distribution of CIDT difference between the greenest CIRO path and the greenest BGP path, across all AS pairs. It suggests that 83% of AS pairs (and their endpoints) can benefit from CIDT reduction by using CIRO's greenest path. This reduction is at least 0.015 g/Gbit for half of AS pairs. Only a negligible portion of AS pairs experience increased CIDT as the used heuristic algorithm does not guarantee optimality.

**Table 1: Carbon-footprint reduction for the outbound *inter-domain traffic* (HTTP(S) and video streaming) of the top-8 beneficiaries of carbon-aware path selection. This results are with the assumption that content providers do not have caches in local ISPs, and all their traffic to other ASes traverses inter-domain paths.**

Source AS (ASN)	Carbon footprint reduction (t/yr)	Relative carbon footprint reduction for outbound inter-domain traffic (%)	Traffic (EB/yr)
Netflix (2906)	47,640	68	423
YouTube (36040)	43,032	55	326
Amazon (16509)	30,240	65	230
Cloudflare (13335)	28,728	79	147
Google (15169)	7536	69	57
Fastly (54113)	5520	57	55
Microsoft (8075)	4284	74	29
Incapsula (19551)	3876	76	22

**6.2.3 Relative CIDT.** Fig. 5b illustrates the distribution of CIDT of the greenest CIRO path relative to the greenest BGP path, across all AS pairs. According to this figure, half of AS pairs (and their endpoints) can use paths that are at least 47% greener.

### 6.2.4 Reduction in Carbon Footprint of End Domains.

Fig. 5c depicts the distribution of ASes' relative reduction in annual carbon footprint of their outbound inter-domain traffic enabled by CIRO. According to this figure, this reduction is at least 50% for more than 85% of ASes.

Table 1 demonstrates absolute and relative annual carbon-footprint reductions of outbound inter-domain traffic through carbon-aware path selection for its top-8 beneficiaries. It also suggests that the most popular ASes, i.e., the largest CDNs, can benefit from the largest *absolute* carbon-footprint reduction.

## 6.3 Limitations

Our results provide an *upper* bound on the carbon-footprint reduction for endpoints and end-domains due to two assumptions:

**No local caching.** In our inter-domain traffic-matrix synthesis, we assume all the traffic between any pair of ASes traverses inter-domain paths, and thus, is not cached by local ISPs.

**Device energy consumption.** For the energy consumption of network devices, we make an effort to use data sources that are both peer-reviewed and comprehensive, i.e., report energy intensity for a wide range of network devices, and not just a small sample of devices. Hence, we rely on the most recent article that satisfies both conditions [25], with the caveat that it does not reflect the growing energy efficiency of modern hardware.

## 6.4 We Report Reduction in Carbon Footprint, Not Emission

As mentioned in §1, the *carbon footprint* of Internet usage of an endpoint (or an end domain) is the amount of carbon emission that






Figure 5: Greenness comparison between carbon-aware path selection, and BGP.

can be attributed to the Internet data transmission of the endpoint.

Our system provides end domains with the opportunity to reduce their carbon footprint, i.e., emission they are responsible for, by sending their traffic on greener paths, which is beneficial for carbon accounting of end domains. However, these reductions in carbon footprint *do not directly translate into global carbon-emission reduction*: Even after shifting traffic to greener paths, the network devices on more polluting paths consume energy, and cause carbon emissions. Since idle energy consumption dominates energy consumption in network devices, the traffic shift to greener paths does not significantly reduce global carbon emission – the reduction in global carbon emissions, would be affected through competition dynamics among ISPs to counteract the traffic shifts [49].

## 7 DISCUSSION

This section discusses several important aspects of carbon-aware inter-domain routing and outlines future research directions related to carbon-aware routing.

### 7.1 Effects on Transmission Quality

**Latency.** We evaluate the impact of optimizing CIDT of paths on their propagation delay as an indicator of latency. We compute the propagation delay of an inter-domain path as the sum of the propagation delays of all AS hops, which we approximate using the great circle delay between ASes' border routers available in the CAIDA dataset. This is a lower bound for the actual latency, which on average underestimates by a factor of  $\sim 1.5$  [53].

Fig. 6 illustrates the distribution of relative propagation delay compared to BGP for paths discovered by CIRo, across AS pairs. According to this figure, optimizing paths only for CIDT does not increase the delay for 75% of AS pairs, and even reduces it for 35% of AS pairs. However, around 5% of AS pairs can suffer from significant inflated delay by a factor of more than 2 due to *inefficient green routes* resulting from optimizing paths exclusively for CIDT. Nevertheless, endpoint path selection in PANs provides an opportunity for multi-criteria path optimization to address this issue. If endpoints specify the AS-level path in every packet and border routers can thus be stateless (as in SCION [13, 34]), this multi-criteria optimization can also be scalable. Developing PAN-based multi-criteria path-




Figure 6: AS-pair distribution regarding the relative propagation delay compared to BGP for (1) the greenest path discovered by CIRo, and (2) the average of the  $k$ -greenest paths discovered by CIRo,  $k$  being the number of the source AS's border routers.

optimization algorithms is an interesting task for future work.

**Congestion.** By enabling endpoints to select network paths based on carbon intensity, many endpoints might select green paths, raising concern regarding congestion on such paths. However, due to the vast diversity of paths in the Internet, every AS pair is connected by numerous low-emission paths, with almost the same carbon intensity as the greenest possible paths. Fig. 4 confirms that multiple green paths with carbon intensities as low as the greenest path exist between all AS pairs. Thus, endpoints have multiple choices to reduce the carbon footprint of their inter-domain communication, lowering the probability of congestion on one single green path.

**Traffic Oscillation.** Oscillation arises in load-adaptive path selection, where delayed information about path load entices endpoints to switch paths. This shift causes a high load on the newly selected path and subsequent abandonment of that path [19, 48]. Therefore, such oscillation arises if the path attribute used as a selection crite-

tion changes in response to the selection of the path. In contrast to path load, path carbon intensity is generally no such responsive criterion and therefore is unlikely to cause oscillation. However, we note that path selection based on carbon intensity may be combined with load-adaptive path selection to avoid congestion on green paths. The load-adaptive part of this multi-criteria strategy then needs to follow the guidelines for oscillation-free path selection developed in previous research [19, 32]. We expect congestion risk on green paths to decrease continuously as competitive pressure erodes greenness differences between paths.

## 7.2 Operational Overhead and Carbon Footprint

**Overhead of CIRo.** An important concern is the energy consumption and thus the carbon footprint of running CIRo itself resulting from its computation and communication overhead. In the following paragraphs, we provide a detailed analysis of CIRo overhead, showing that a single mid-size server is more than sufficient to run a whole CIRo instance for a large AS with a global operation. Thus, CIRo introduces negligible overhead.

For each instance of CIRo, the overhead of the forecasting module includes *hourly* querying CIE forecasts, collecting AS internal topology and routes, computing CIDT for all interface pairs using Eq. (10), and inserting them into the forecast database. The overhead of the dissemination module includes querying the forecast database, and encoding forecast information into routing messages, increasing the message size.

The overhead of the forecasting module only depends on the size of the AS it is located in. For a large AS with global operation with  $O(10^4)$  interfaces, the per-period (hourly) forecasting overhead translates to  $O(10^3)$  CIE queries (on the order of number of geographical zones in ElectricityMaps), writing  $O(10^8)$  rows in a database table, and  $O(10^{10})$  addition and multiplication operations. An efficient implementation using modern databases can carry out this procedure in  $O(10^2)$  seconds on a single machine with a few CPU cores with  $O(10^1)$  GB of memory. We do not consider collecting internal topology and routes as an overhead of CIRo, since ASes already continuously monitor and collect information about their network, irrespective of CIRo's presence.

The dissemination module's database queries per period (at least once in 2.6 hours. cf. §5.2) are upper bounded by the square of the number of interfaces in an AS; thus, a large AS with  $O(10^4)$  interfaces requires  $O(10^8)$  database queries.

The communication overhead of the dissemination module depends on the number of interfaces of the AS, the number of ASes in the inter-domain topology, and the number of paths per origin AS that the AS advertises to its neighbors. Our SCIONLab experiments with 100 ASes suggest a maximum of 125 kB overhead per period per interface. By extrapolating to  $O(10^4)$  ASes in the network, we arrive at a maximum of 12.5 MB overhead per period per interface. Each inter-domain link needs to transfer this additional data in a period, which is a negligible effort considering the period length and the typical capacity of inter-domain links.

**Overhead of PAN Architectures.** Since the stateless routers in a PAN architecture are significantly more energy efficient than routers in today's Internet [57], using a PAN architecture can further reduce the carbon footprint of the Internet.

## 7.3 Amortizing Idle Carbon Emission

In this work, we take each bit of data responsible for the idle energy consumption of fully utilized devices by amortizing their idle power consumption over their capacity, irrespective of traffic load.

Another approach is to amortize idle power based on the traffic load on a device. In this way, a bit of data is responsible for more carbon emission if the device is under-utilized and vice versa. This can cause a positive feedback on greener paths: with more traffic they attract, because of their initial greenness, their carbon intensity reduces, making them more attractive. However, this method complicates carbon intensity prediction because (1) it requires predicting traffic on a device level over the next 24 hours, and (2) it creates hard to predict dynamics through the positive feedback (i.e., paths becoming greener with more traffic).

## 7.4 Carbon Transparency

By providing endpoints with predictions of paths' carbon intensity, CIRo takes an essential step towards carbon transparency on the public Internet and helps end domains with their carbon footprint accounting. However, to achieve full carbon transparency and accurate accounting, real-time information about paths' carbon intensity should be distributed globally. This requires a complementary system to work in parallel with CIRo. Designing such a system has its own specific challenges, tackling which is out of the scope of the present work.

## 7.5 Trustworthy Carbon Information

Green path selection by endpoints could tempt ASes to claim false carbon-intensity information to attract more traffic. Hence, carbon trust roots are needed to certify carbon-intensity information claimed by ISPs. The main challenges for these trust roots are: to estimate the carbon intensity of paths within ISPs, to ensure that ISPs do not deviate from their certified carbon intensity after obtaining a certificate, and to periodically re-certify carbon intensities.

## 7.6 Deployment and Ossification

The deployment of CIRo in a PAN architecture with stateless routers and a control plane separate from the data plane is relatively easy, as it introduces no change to the data plane and only modest, backward-compatible extensions of the control plane. Thus, no hardware or protocol update is required. While the deployment of a PAN Internet itself requires significant effort, the SCION Internet has in fact been experiencing growing deployment [34] in recent years as a result of its incremental deployability.

## 7.7 Economic Effects of Green Routing

From an economic perspective, it is central to predict the impact of carbon routing on the business performance of ISPs and on the competitive environment in the Internet, considering that ISPs might reduce their carbon footprint if they can thereby attract traffic. These competitive dynamics may result in a virtuous cycle, where ISPs compete for traffic shares by offering paths with ever-decreasing carbon intensity. In fact, this effect is confirmed by recent work based on game theory and economic modeling [49].

## 7.8 Green Data Centers and CDNs

A path-aware Internet architecture can also enable endpoints to select the data center providing a service or content [38]. Therefore, providing carbon-intensity information of data centers in combination with information about communication paths enables endpoints to monitor and optimize the *total* carbon footprint of their requests, providing CO<sub>2</sub> optimization for computation-intensive and communication-intensive applications.

## 8 RELATED WORK

Much work has studied how to reduce the carbon footprint of the ICT sector by either improving the energy efficiency of end devices and networks [2, 21, 30, 31, 42, 50, 52, 54, 66, 69], or increasing the utilization of available renewable energy resources [22, 45, 46, 61].

**Green data centers.** A large body of research proposes ecological improvements in the area of data centers [23]: energy-efficient software [31], capacity planning for better resource utilization [54], improved energy management [21], better virtualization technologies [66], power-saving cooling systems [30], greenness-oriented load balancing across data centers [38], or simply using more renewable energy [45]. In a recent work, Radovanović et al. [44] propose a method to minimize the carbon footprint of computing among globally distributed data centers by temporally delaying flexible workloads. They achieve this goal by predicting the day-ahead carbon intensity and computation demands.

**Carbon-/energy-aware networking.** Research conducted on green routing and traffic engineering falls into two main categories: they either make the network more energy-efficient or route packets through paths whose energy resources are green.

In a visionary paper, Gupta and Singh point out the energy utilization of Internet routers, and suggest energy savings by placing devices in sleep mode [24]. Later, Zhang et al. [69] propose a heuristic to reduce energy by turning off line cards and rerouting traffic to underutilized links. Vasic et al. [64] perform an energy optimization on real-world networks by inactivating unnecessary network elements. Vasic and Kostic [63] propose Energy-Aware Traffic Engineering (EATe), to allocate traffic in an energy-optimal manner within an ISP. Andrews et al. [4] study network optimization from a theoretical perspective, with energy minimization as an objective. Chabarek et al. [12] study the power consumption of core and edge routers, optimizing the intra-domain energy consumption.

Among the studies also taking energy-efficient routing into account, Van der Veldt et al. [61] propose a path provisioning method to reduce the CO<sub>2</sub> emission of communications in a national research and education network (NREN). Aksanli et al. [2] design green energy-aware routing policies for wide area traffic between data centers. Ricciardi et al. [46] use integer linear programming for routing and wavelength assignment considering a network node's carbon intensity. Gattulli et al. [22] propose a CO<sub>2</sub> and energy-aware routing mechanism for intra-domain routing. They find two paths for each source and destination pair; one with routers whose energy resources have the lowest emission, and one with the lowest energy consumption. Then, they compare these paths and select the one with the lower emission.

**Green inter-domain communication.** To the best of our knowl-

edge, only two previous studies propose methods to reduce the carbon footprint of inter-domain communications. Shi et al. [52] extend the aforementioned traffic aggregation method [69] by taking into account the traffic between border routers of an AS. Thus, it does not tackle global carbon-aware routing. Nafarieh et al. [42] propose extensions for OSPF-TE and BGP to propagate information about the emission of links on each path to routers inside and outside an AS. Using this information, each router selects the path with the minimum path emission to all other routers in its own AS, and each border router shares this information with the neighboring AS. During BGP route propagation, the emission of paths is accumulated in update messages, and used by multi-homed ASes to select the provider of the greenest route. However, this BGP-based approach provides limited transparency and control over the carbon footprint of inter-domain communications compared to CIRo.

**Internet Energy models.** Since we propose a model for the carbon-intensity of Internet paths, we provide an overview of the related work in the area of modeling the energy consumption of the Internet in Appendix A as energy consumption and carbon intensity can be considered as correlated topics.

## 9 CONCLUSION

This work presents an important step towards realizing carbon-aware global routing by designing and implementing CIRo, a practical system to forecast and disseminate carbon-intensity information of inter-domain paths. Our large-scale simulation results suggest that CIRo enables endpoints to substantially reduce their carbon footprint of Internet usage.

Importantly, we note that CIRo may also support the renewable-energy transition of entire energy systems in two main ways. First, renewable-energy sources such as solar and wind have distinct production peaks, which make it challenging to optimally utilize and monetize the electricity produced during these peaks. In this regard, CIRo is valuable because it allows to shift energy consumption to locations where production peaks occur and abundant green energy is currently available. Second, CIRo creates a global competitive environment in which ISPs worldwide can attract traffic by using green energy. This new competitive pressure is also felt in countries where the renewable-energy transition has so far been hindered by inert and monopolistic electricity markets. In these countries, ISPs thus have an incentive to convince local electricity producers to expand the green-energy supply, which may accelerate the green-energy transition especially in developing and emerging countries. In our further research, we are eager to investigate the role carbon-aware routing can play in the global quest for sustainable energy.

## 10 ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their insightful feedback. We would also like to thank Markus Legner, Cedric Meury, and Tobias Schmid for their insightful discussions and valuable ideas and comments.

## REFERENCES

- [1] Arsalan Ahmad, Andrea Bianco, Edoardo Bonetto, Luca Chiaraviglio, and Filip Idzikowski. 2013. Energy-Aware Design of Multilayer Core Networks. *Optical*

- Communications and Networking* 5 (Oct 2013), A127–A143. <https://doi.org/10.1364/JOCN.5.00A127>
- [2] Baris Aksanli, Tajana Simunic Rosing, and Inder Monga. 2012. Benefits of green energy and proportionality in high speed wide area networks connecting data centers. In *Proceedings of the Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 175–180. <https://doi.org/10.1109/DAT.2012.6176458>
  - [3] Anders S. G. Andrae and Tomas Edler. 2015. On Global Electricity Usage of Communication Technology: Trends to 2030. *Challenges* 6, 1 (2015), 117–157. <https://doi.org/10.3390/challe6010117>
  - [4] M. Andrews, A. Anta, L. Zhang, and W. Zhao. 2010. Routing for energy minimization in the speed scaling model. In *Proceedings of the International Conference on Computer Communications (INFOCOM)*. IEEE.
  - [5] J. Baliga, R. Ayre, K. Hinton, W. Sorin, and R. Tucker. 2009. Energy Consumption in Optical IP Networks. *Lightwave Technology* 27, 13 (2009), 2391–2403. <https://doi.org/10.1109/JLT.2008.2010142>
  - [6] Jayant Baliga, Kerry Hinton, and Rodney Tucker. 2007. Energy Consumption of the Internet. In *Proceedings of the Joint International Conference on the Optical Internet and the Australian Conference on Optical Fibre (COIN-ACOFT)*. IEEE. <https://doi.org/10.1109/COINACOFT.2007.4519173>
  - [7] C. Belady and Andy Rawson. 2008. Green Grid Data Center Power Efficiency Metrics: PUE and DCIE.
  - [8] Center for Applied Internet Data Analysis (CAIDA). [n. d.]. AS-Relationships. <https://www.caida.org/data/as-relationships/> (last visited Apr 19, 2021), archived at <https://perma.cc/RYN2-BSNB>. ([n. d.])
  - [9] Center for Applied Internet Data Analysis (CAIDA). [n. d.]. AS Relationships – with Geographic Annotations. <https://www.caida.org/data/as-relationships-geo/> (last visited Apr 19, 2021), archived at <https://perma.cc/S4XX-Y2EK>. ([n. d.])
  - [10] Center for Applied Internet Data Analysis (CAIDA). [n. d.]. Macroscopic Internet Topology Data Kit (ITDK). <https://www.caida.org/data/internet-topology-data-kit/> (last visited Apr 19, 2021), archived at <https://perma.cc/CU8X-7GRU>. ([n. d.])
  - [11] Center for Applied Internet Data Analysis (CAIDA). [n. d.]. Routeviews Prefix to AS Mappings Dataset (pfx2as) for IPv4 and IPv6. <https://www.caida.org/data/routing/routeviews-prefix2as.xml> (last visited Apr 19, 2021), archived at <https://perma.cc/RB6L-PEVA>. ([n. d.])
  - [12] Joseph Chabarek, Joel Sommers, Paul Barford, Cristian Eitan, David Tsang, and Steve Wright. 2008. Power Awareness in Network Design and Routing. In *Proceedings of the International Conference on Computer Communications (INFOCOM)*.
  - [13] Laurent Chuat, Markus Legner, David Basin, David Hausheer, Samuel Hitz, Peter Müller, and Adrian Perrig. 2022. *The Complete Guide to SCION. From Design Principles to Formal Verification*. Springer International Publishing AG. <https://link.springer.com/book/10.1007/978-3-031-05288-0>
  - [14] A. Clemm, L. Dong, G. Mirsky, L. Ciavaglia, J. Tantsura, and M-P. Odini. 2022. Green Networking Metrics. <https://datatracker.ietf.org/doc/draft-cx-green-metrics/>. (July 2022).
  - [15] Vlad C. Coroama, Lorenz M. Hilty, Ernst Heiri, and Frank M. Horn. 2013. The Direct Energy Demand of Internet Data Flows. *Industrial Ecology* 17, 5 (2013), 680–688. <https://doi.org/10.1111/jiec.12048>
  - [16] Edsger W Dijkstra. 1959. A note on two problems in connexion with graphs. *Numerische mathematik* 1, 1 (1959), 269–271.
  - [17] Ottmar Edenhofer. 2012. *Renewable Energy Sources and Climate Change Mitigation: Special Report of the Intergovernmental Panel on Climate Change*. Cambridge university Press.
  - [18] Electricity Maps. 2022. Electricity Maps. <https://www.electricitymaps.com> (last visited Nov 24, 2022), archived at <https://perma.cc/S4ND-PRAZ>. (2022).
  - [19] Simon Fischer and Berthold Vöcking. 2009. Adaptive routing with stale information. *Theoretical Computer Science* 410, 36 (2009), 3357–3371.
  - [20] Matthias Frei. 2021. Path Metadata in Beacons. <https://github.com/scionproto/scion/blob/master/doc/beacon-metadata.rst> (last visited Apr 19, 2021), archived at <https://perma.cc/9NXU-KDNR>. (Feb 2021).
  - [21] Saurabh Kumar Garg and Rajkumar Buyya. 2012. Green cloud computing and environmental sustainability. In *Harnessing Green IT: Principles and Practices*. 315–340.
  - [22] M. Gattulli, M. Tornatore, R. Fiandra, and A. Pattavina. 2012. Low-Carbon Routing Algorithms for Cloud Computing Services in IP-over-WDM Networks. In *IEEE International Conference on Communications (ICC)*. 2999–3003. <https://doi.org/10.1109/ICC.2012.6364347>
  - [23] Sukhpal Singh Gill and Rajkumar Buyya. 2018. A Taxonomy and Future Directions for Sustainable Cloud Computing: 360 Degree View. *Comput. Surveys* 51, 5, Article 104 (Dec 2018), 33 pages.
  - [24] M. Gupta and S. Singh. 2003. Greening of the Internet. In *Proceedings of the SIGCOMM Conference (SIGCOMM)*. Association for Computing Machinery.
  - [25] Ward Heddeghem, Filip Idzikowski, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester. 2012. Power Consumption Modeling in Optical Multilayer Networks. *Photonic Network Communications* 24 (Oct 2012). <https://doi.org/10.1007/s11107-011-0370-7>
  - [26] Kerry Hinton, Jayant Baliga, Michael Feng, Robert Ayre, and Rodney Tucker. 2011. Power Consumption and Energy Efficiency in the Internet. *Network, IEEE* 25 (May 2011), 6 – 12. <https://doi.org/10.1109/MNET.2011.5730522>
  - [27] Kerry Hinton, Fatemeh Jalali, and Ashrar Matin. 2015. Energy Consumption Modeling of Optical Networks. *Photonic Netw. Commun.* 30, 1 (Aug 2015), 4–16. <https://doi.org/10.1007/s11107-015-0491-5>
  - [28] Geoff Huston. 2021. BGP in 2020 – The BGP Table. <https://blog.apnic.net/2021/01/05/bgp-in-2020-the-bgp-table/> (last visited Apr 19, 2021), archived at <https://perma.cc/9U97-98UN>. (Jan 2021).
  - [29] International Energy Agency (IEA). 2018. Data & Statistics. <https://www.iea.org/data-and-statistics?country=WORLD&fuel=Electricity%20and%20heat&indicator=ElecGenByFuel> (last visited Apr 19, 2021), archived at <https://perma.cc/72QW-26D9>. (2018).
  - [30] Weixiang Jiang, Ziyang Jia, Sirui Feng, Fangming Liu, and Hai Jin. 2019. Fine-Grained Warm Water Cooling for Improving Datacenter Economy. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*. Association for Computing Machinery.
  - [31] Chao Jin, Bronis R. de Supinski, David Abramson, Heidi Poxon, Luiz DeRose, Minh Ngoc Dinh, Mark Endrei, and Elizabeth R. Jessup. 2017. A survey on software methods to improve the energy efficiency of parallel computing. *Int. J. High Perform. Comput. Appl.* 31, 6 (2017), 517–549.
  - [32] Frank Kelly and Thomas Voice. 2005. Stability of end-to-end algorithms for joint routing and rate control. *ACM SIGCOMM Computer Communication Review (CCR)* (2005).
  - [33] Jonathan Koomey, Huimin Chong, Woonsien Loh, Bruce Nordman, and Michele Blazek. 2004. Network Electricity Use Associated with Wireless Personal Digital Assistants. *Infrastructure Systems* 10, 3 (2004), 131–137. [https://doi.org/10.1061/\(ASCE\)1076-0342\(2004\)10:3\(130\)](https://doi.org/10.1061/(ASCE)1076-0342(2004)10:3(130))
  - [34] Cyrill Krähenbühl, Seyedali Tabaeiaighdaei, Christelle Gloor, Jonghoon Kwon, Adrian Perrig, David Hausheer, and Dominik Roos. 2021. Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure. In *Proceedings of ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. Association for Computing Machinery.
  - [35] Jonghoon Kwon, Juan A. Garcia-Pardo, Markus Legner, François Wirz, Matthias Frei, David Hausheer, and Adrian Perrig. 2020. SCIONLAB: A Next-Generation Internet Testbed. In *Proceedings of the International Conference on Network Protocols (ICNP)*. IEEE, 1–12. <https://doi.org/10.1109/ICNP49622.2020.9259355>
  - [36] Steven Lanzisera, Bruce Nordman, and Richard E. Brown. 2011. Data Network Equipment Energy Use and Savings Potential in Buildings. *Energy Efficiency* 5, 2 (2011), 149–162. <https://doi.org/10.1007/s12053-011-9136-4>
  - [37] Matt Lepinski and Kotikalapudi Sriram. 2017. BGPsec Protocol Specification. RFC 8205. (Sept. 2017). <https://doi.org/10.17487/RFC8205>
  - [38] Zhenhua Liu, Minghong Lin, Adam Wiernan, Steven H Low, and Lachlan LH Andrew. 2011. Greening geographical load balancing. *ACM SIGMETRICS Performance Evaluation Review* 39, 1 (2011), 193–204.
  - [39] Josip Lorincz, Antonio Capone, and Jinsong Wu. 2019. Greener, Energy-Efficient and Sustainable Networks: State-Of-The-Art and New Trends. *Sensors* 19 (Nov 2019), 4864. <https://doi.org/10.3390/s19224864>
  - [40] Diptyarop Maji, Ramesh K. Sitaraman, and Prashant Shenoy. 2022. DACF: Day-Ahead Carbon Intensity Forecasting of Power Grids Using Machine Learning. In *Proceedings of the International Conference on Future Energy Systems (e-Energy) (e-Energy)*. Association for Computing Machinery, New York, NY, USA, 188–192. <https://doi.org/10.1145/3538637.3538849>
  - [41] J. Mikians, N. Laoutaris, A. Dhamdhere, and P. Barlet-Ros. 2013. ITMgen – A First-Principles Approach to Generating Synthetic Interdomain Traffic Matrices. In *2013 IEEE International Conference on Communications (ICC)*. 2507–2512. <https://doi.org/10.1109/ICC.2013.6654910>
  - [42] 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. <https://doi.org/10.1016/j.procs.2013.09.006>
  - [43] Victor Pochat, Tom Van Goethem, Samaneh Tajalizadehkhooob, Maciej Korczynski, and Wouter Joosen. 2022. Tranco. <https://tranco-list.eu> (last visited Jan 11, 2022), archived at <https://perma.cc/PK5Q-Z2W5>. (Jan 2022).
  - [44] Ana Radovanovic, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, Saurav Talukdar, Eric Mullen, Kendal Smith, Mariellen Cottman, and Walfredo Cirne. 2022. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems* (2022), 1–1. <https://doi.org/10.1109/TPWRS.2022.3173250>
  - [45] Syed Raza, Isam Janajreh, and Chaouki Ghenai. 2014. Sustainability index approach as a selection criteria for energy storage system of an intermittent renewable energy source. *Applied Energy* 136 (Dec 2014), 909–920.
  - [46] S. Ricciardi, D. Careglio, F. Palmieri, U. Fiore, G. Santos-Boada, and J. Sole-Pareta. 2011. Energy-Aware RWA for WDM Networks with Dual Power Sources. In *IEEE International Conference on Communications (ICC)*. 1–6. <https://doi.org/10.1109/icc.2011.5962432>
  - [47] Sandvine. 2018. The Global Internet Phenomena Report. <https://www.sandvine.com/hubs/downloads/phenomena/2018-phenomena-report.pdf> (last visited Apr 19, 2021), archived at <https://perma.cc/SY2E-C7N8>. (Oct 2018).
  - [48] Simon Scherrer, Markus Legner, Adrian Perrig, and Stefan Schmid. 2020. Incentivizing stable path selection in future internet architectures. *Performance*

- Evaluation 144 (2020), 102137.
- [49] Simon Scherrer, Seyedali Tabaeiaghdaei, and Adrian Perrig. 2023. Quality Competition Among Internet Service Providers in a Path-Aware Internet. *arXiv preprint arXiv:2305.06811* (2023).
- [50] Thilo Schöndienst and Vinod Vokkarane. 2014. Renewable Energy-Aware Grooming in Optical Networks. *Photonic Network Communications* 28 (Aug 2014), 71–81. <https://doi.org/10.1007/s11107-014-0436-4>
- [51] SCION. 2023. Scalability, Control and Isolation On next-generation Networks (SCION). <https://github.com/scionproto/scion>. (Jan. 2023).
- [52] J. Shi and B. Zhang. 2014. Making Inter-Domain Routing Power-Aware?. In *Proceedings of the International Conference on Computing, Networking and Communications (ICNC)*. 932–938. <https://doi.org/10.1109/ICNC.2014.6785462>
- [53] Ankit Singla, Balakrishnan Chandrasekaran, P. Brighten Godfrey, and Bruce Maggs. 2014. The Internet at the Speed of Light. In *Proceedings of the Workshop on Hot Topics in Networks (HotNets)*. Association for Computing Machinery. <https://doi.org/10.1145/2670518.2673876>
- [54] Erica Sousa, Fernando Lins, Eduardo Tavares, Paulo Cunha, and Paulo Maciel. 2015. A Modeling Approach for Cloud Infrastructure Planning Considering Dependability and Cost Requirements. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, 4 (2015), 549–558.
- [55] Statista. 2019. Consumer Internet Data Traffic Worldwide by Application Category from 2016 to 2022 (in Eb per Month). <https://www.statista.com/statistics/454951/mobile-data-traffic-worldwide-by-application-category/> (last visited Apr 19, 2021), archived at <https://perma.cc/XLH7-J8W3>. (Feb 2019).
- [56] Seyedali Tabaeiaghdaei and Christelle Gloor. 2022. ns-3-scion. <https://gitlab.com/SeyedaliTabaeiaghdaei/ns-3-scion>. (2022).
- [57] Seyedali Tabaeiaghdaei and Adrian Perrig. 2022. Data-Plane Energy Efficiency of a Next-Generation Internet Architecture. In *Proceedings of the Symposium on Computers and Communications (ISCC)*. IEEE, 1–8. <https://doi.org/10.1109/ISCC55528.2022.9912464>
- [58] Tabaeiaghdaei, Seyedali. 2023. CIRo. <https://github.com/SeyedaliTabaei/CIRo>. (Jan. 2023).
- [59] Cody Taylor and Jonathan Koomey. 2008. Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising. (Jan 2008).
- [60] University of Washington NS-3 Consortium. 2022. ns-3 Network Simulator. <https://www.nsnam.org/> (last visited Jan 18, 2022), archived at <https://perma.cc/4S4N-8VGV>. (2022).
- [61] Karel van der Veldt, Cees de Laat, Inder Monga, Jon Dugan, and Paola Grosso. 2014. Carbon-aware path provisioning for NRENs. In *Proceedings of the International Green Computing Conference (IGCC)*. IEEE. <https://doi.org/10.1109/IGCC.2014.7039147>
- [62] Laurent Vanbever. 2014. dragon\_simulator. [https://github.com/network-aggregation/dragon\\_simulator](https://github.com/network-aggregation/dragon_simulator) (last visited Apr 19, 2021), archived at <https://perma.cc/4QCZ-7JWU>. (Jun 2014).
- [63] N. Vasic and D. Kotic. 2010. Energy-Aware Traffic Engineering. In *Proceedings of the International Conference on Energy-Efficient Computing and Networking (e-Energy)*. Association for Computing Machinery.
- [64] Nedeljko Vasic, Dejan Novakovic, Satyam Shekhar, Prateek Bhurat, Marco Canini, and Dejan Kotic. 2011. Identifying and Using Energy-Critical Paths. In *Proceedings of the International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. Association for Computing Machinery.
- [65] A. Vishwanath, K. Hinton, R. W. A. Ayre, and R. S. Tucker. 2014. Modeling Energy Consumption in High-Capacity Routers and Switches. *IEEE Journal on Selected Areas in Communications* 32, 8 (2014), 1524–1532. <https://doi.org/10.1109/JSAC.2014.2335312>
- [66] Xiaoying Wang, Guojing Zhang, Mengqin Yang, and Lei Zhang. 2017. Green-Aware Virtual Machine Migration Strategy in Sustainable Cloud Computing Environments. In *Cloud Computing-Architecture and Applications*.
- [67] WattTime. 2022. WattTime. <https://www.watttime.org> (last visited Nov 24, 2022), archived at <https://perma.cc/B63V-8D9L>. (2022).
- [68] Christopher L. Weber, Jonathan G. Koomey, and H. Scott Matthews. 2010. The Energy and Climate Change Implications of Different Music Delivery Methods. *Industrial Ecology* 14, 5 (2010), 754–769. <https://doi.org/10.1111/j.1530-9290.2010.00269.x>
- [69] M. Zhang, C. Yi, B. Liu, and B. Zhang. 2010. GreenTE: Power-Aware Traffic Engineering. In *The 18th IEEE International Conference on Network Protocols*. 21–30. <https://doi.org/10.1109/ICNP.2010.5762751>
- [70] Noa Zilberman, Eve M. Schooler, Uri Cummings, Rajit Manohar, Dawn Nafus, Robert Soule, and Rick Taylor. 2022. Toward Carbon-Aware Networking. In *Workshop on Sustainable Computer Systems Design and Implementation*.
- [71] George Kingsley Zipf. 1949. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Addison-Wesley. 573 pages.

## A RELATED WORK ON ESTIMATING THE ENERGY CONSUMPTION OF THE INTERNET

Numerous studies estimate the energy consumption and energy intensity of the Internet, where the methodologies can be categorized in bottom-up, top-down, and model-based approaches.

**Bottom-up approaches.** Bottom-up approaches generalize the energy-intensity values obtained through direct measurement of selected network devices. The study conducted by Coroama et al. falls into this category [15]. In that study, they estimate the energy intensity of the communication between two conference locations in Japan and Switzerland, where the communication path was determined in advance. This work is the only work we are aware of that has estimated the energy intensity of a network *path*, and thus has a relation to our proposed estimation of the *carbon intensity* of inter-domain paths.

**Top-down approaches.** In top-down approaches, researchers divide the total energy consumption of a network by the amount of traffic transited by the network over a particular time period. Studies conducted by Koomey et al. [33], Taylor et al. [59], Weber et al. [68], Lanzisera et al. [36], and Andrae et al. [3] are examples of top-down approaches.

**Model-based approaches.** Model-based approaches rely on modelling parts of the Internet based on network-design principles, and on energy-consumption information of device vendors in order to find the total energy consumption of specific Internet parts. Baliga et al. [5, 6], Vishwanath et al. [65], and Hinton et al. [26] propose model-based approaches to estimate the energy consumption of the Internet.

## B SCION-SPECIFIC DETAILS OF THE INFORMATION DISSEMINATION MODULE

Fig. 7 demonstrates the structure of this extension within a PCB. Every AS can include one CIDTExtension in its ASEntry to disseminate its CIDT forecasts. A CIDTExtension is a map from relevant interface identifiers (IntfId) of the AS to vectors of CIDT forecasts. Having more than one relevant interface is useful for composition of SCION path segments.

**SCION-Specific CIDTExtension.** An IntfId, as a key of this map, represents the intra-domain segment from the IntfId to the egress interface (EgIntf) of the ASEntry.

The CIDT forecast vector associated with an IntfId provides CIDT forecasts for the forward path from IntfId to EgIntf ( $CIDT_{AS[ing, eg]}$ ) and for the backward path from EgIntf to IntfId ( $CIDT_{AS[eg, ing]}$ ).

A CIDT forecast vector is a sequence of 48 contiguous unsigned 8-bit integer numbers, each specifying one CIDT value in mg/Gbit. We choose this unit based on our observation of the CIDT range in Internet paths (cf. §6.2). The first 24 values of the vector represent  $CIDT_{AS[ing, eg]}$  forecasts for the 24-hour interval starting from the full hour time point before the PCB's TimeStamp. The 24 following values provide similar information for  $CIDT_{AS[eg, ing]}$ . This design ensures time-alignment of forecasts added by different ASes at different times. **Database Query by Dissemination Module.** For






Figure 7: A SCION PCB containing CIDTExtension.

each PCB to disseminate, the dissemination module queries CIDT forecasts between the EgIntf and all IntfIds in relevant interfaces.

**CIDT Dissemination in Core Beaconsing.** The forecasting module uses core beaconsing to disseminate CIDT forecasts of core-path segments. For each core PCB that is being disseminated by a core AS, the beacon service constructs a CIDTExtension containing the CIDT forecasts only for the intra-domain segment between the ingress and egress interfaces of the PCB.

Including forecasts for only one interface is essential to the scalability of forecasting module among core ASes. This is because core ASes are large ASes with numerous interfaces, so including information about all interfaces of every hop on the path would introduce significant overhead to each PCB. This per-PCB overhead in conjunction with the high message complexity of core beaconsing [34] would hinder the scalability of the forecasting module. This encoding, however, does not impact the expressiveness of CIDT forecasts since each core PCB only specifies one core-path segment and core-path segments do not combine.

**CIDT Dissemination in Intra-ISD Beaconsing.** The forecasting module uses intra-ISD beaconsing to disseminate CIDT forecasts of up- and down-path segments. Therefore, beacon services include the CIDT forecasts between ingress and egress interfaces of disseminated PCBs in CIDTExtensions. However, since an end-to-end path could be a combination of path segments and peering links, it is essential to provide CIDT forecast of their combination as well to arrive at the end-to-end CIDT forecast. To cover the combination with core-path segments, a core AS originating a non-core PCB includes the CIDT forecasts between the EgIntf of the PCB and all interfaces connecting to all other core ASes. To cover the combination of up- and down-path segments, any AS disseminating a non-core PCB includes the CIDT forecasts between the EgIntf of the PCB and all interfaces connecting to all other customer ASes whose IntfId is smaller than EgIntf. This way, the required information is either encoded in the up segment or in the down segment, halving the overhead. To cover the combination of up- or down-path segments with peering links, any AS disseminating a non-core PCB includes the CIDT forecasts between the EgIntf of the PCB and all interfaces connecting to all other peering ASes.

**Obtaining the CIDT of Full Inter-Domain Paths.** Once an endpoint retrieves path segments to construct an inter-domain path it computes the CIDT of the full path for the current hour using CIDTExtensions provided by all segments. Note that different path segments have different origination time, thus, their CIDT forecast vectors are not necessarily aligned. Therefore, endpoints compute the index of current time for each path segment independently.

Then, they accumulate the CIDT forecast of all AS hops associated with the current hour to arrive at the CIDT of the end-to-end path.

## C TYPICAL ENERGY INTENSITY OF NETWORK DEVICES

Table 2: Energy intensities of typical devices in IP and WDM layers [25].

Device type	Energy intensity ( $I_E$ in W/Gbps = J/Gbit)
Core router	10
WDM switch (OXC)	0.05
Trans/Mux -ponder	1.5
Amplifier	0.03
Regenerator	3

## D CARBON INTENSITY OF ENERGY RESOURCES

Table 3: 50<sup>th</sup> percentile CO<sub>2</sub> emission of different energy resources. From Edenhofer's IPCC report [17].

Energy resource	Carbon intensity (gCO <sub>2</sub> /kWh)
Coal	1001
Natural gas	469
Biomass	230
Solar	46
Geothermal	45
Nuclear	16
Wind	12
Hydroelectric	4

## E TRAFFIC MATRIX SYNTHESIS

We compute the global inter-domain traffic matrix for all ASes in the CAIDA AS-Rel dataset [8] using the model provided by Mikians et al. [41]. In this model, the amount of traffic from AS<sub>*i*</sub> to AS<sub>*j*</sub> is the sum of traffic sent from all users and servers in AS<sub>*i*</sub> to all servers and users in AS<sub>*j*</sub> for all application types:

$$T_{i,j} = \sum_{A \in \text{applications}} m_A(S_i p_i^A(j) + d_A S_j p_j^A(i)), \quad (11)$$

where  $p_i^A(j)$  is the relative popularity of servers in AS<sub>*j*</sub> for users in AS<sub>*i*</sub> with respect to application A,  $S_i$  represents the size of an AS in the number of its IP addresses,  $d_A$  denotes the asymmetry factor for application A, which is the ratio of the response flow size to the request flow size for application A, and  $m_A$  scales the computed relative traffic for application A to its absolute real traffic volume.

We set  $\log_{10} d_{http} = 1$  which is in the range of (0.4, 1.5) for HTTP(S) data and media [41]. Furthermore, we assume the size of each AS to be its number of IPv4 addresses from the CAIDA Pfx2AS dataset [11], except Tier-1 ASes with no users and large CDNs with

no users requesting a service. Moreover,  $p_i^{http}(j)$  follows the Zipf distribution [71] with a slope of 1.2 [41]. For each source AS, we thus generate a vector of relative popularities from the Zipf distribution. Then, we assign the largest popularity values to the most popular destination ASes, and the remaining popularity values randomly to other ASes. We define the most popular ASes as the ASes whose hosted websites have the largest accumulated inverse ranks in the Tranco1M dataset [43]. We find the website to AS mappings using DNS queries, and the Pfx2AS dataset. Once the matrix is computed, we scale it to the global HTTP(S) traffic, which is estimated to be 82 Exabytes per month [47, 55].

In addition to HTTP(s) traffic, we also include video traffic in the traffic matrix. For this video traffic, we consider Netflix, YouTube,

and Amazon Prime Video, which are responsible for 15, 11.4, and 3.7 percent of total Internet traffic, respectively [47]. We construct the video traffic matrix by assuming that the amount of traffic any other AS receives from these services is proportional to its number of users in Pfx2AS dataset, and these services receive negligible traffic from other ASes. Finally, we add the video traffic matrix to the HTTP(S) traffic matrix.

To construct the traffic matrix for the core topology of 2000 ASes used in our simulation, we assume that the amount of traffic between two core ASes is the sum of all traffic in the global traffic matrix between their customer cone ASes in the AS-Rel dataset. If an AS is in the customer cones of multiple core ASes, its inbound and outbound traffic is evenly divided across its core AS providers.