Green Inter-Domain Routing in the SCION Internet Architecture

Master Thesis

Author(s):
Tabaeiaghdaei, Seyedali

Publication date:
2021-04-20

Permanent link:
https://doi.org/10.3929/ethz-b-000486929

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Master Thesis

Network Security Group, Department of Computer Science, ETH Zurich

Green Inter-Domain Routing in the SCION Internet Architecture

by Seyedali Tabaeiaghdaei

Spring 2021

ETH student ID: 19-954-346
E-mail address: tabaeias@student.ethz.ch

Supervisors: Dr. Markus Legner
Prof. Dr. Adrian Perrig

Date of submission: April 20, 2021

Abstract

The growth in the number of Internet users and the amount of Internet data in past years has significantly increased the Information and Communication Technology (ICT) sector's energy consumption. It is estimated that the ICT sector is responsible for 8 percent of electricity consumption and 2.7 percent of global CO_2 emissions. Communication networks are responsible for the considerable share of 20 percent of the whole electricity consumption of the ICT sector. Therefore, numerous studies have been trying to reduce the energy consumption or CO_2 emission of communication networks; however, all of them propose approaches that can be applied only within one domain or between two neighboring domains.

In this work, we propose an inter-domain approach based on the SCION next-generation Internet architecture to decrease the CO_2 emission of inter-domain communications by routing packets through more energy-efficient or more carbon-efficient routes. Using simulations, we show that this approach can decrease the CO_2 emission by at least 50 percent for the communication paths between more than half of the domain pairs in a large-scale subset of today's Internet topology. Furthermore, as some end users want to reduce their contribution to the CO_2 emission, they select greener paths. This causes a shift of traffic from more-polluting carrier networks to less-polluting ones. Therefore, to attract more traffic, carrier networks compete with each other to use more renewable energy resources. Using simulations, we show that this competition reduces the CO_2 emission of the whole network by 87 percent.

Acknowledgements

First and foremost, I would like to thank Prof. Dr. Adrian Perrig for giving me the opportunity to work on this fascinating project. I appreciate his guidance during the project.

I would like to express my sincere gratitude to Dr. Markus Legner for his supervision and valuable insights on the thesis. Without his help, this thesis would not have been possible.

I would like to thank Simon Scherrer for his economic insights on the project, and Dr. Huayi Duan for his valuable feedback on this work.

Lastly, my special thank goes to my family, especially my wife, for her unconditional support during my studies.

Contents

1 Introduction

In this section we provide an overview of the thesis. In Section 1.1 we describe what motivates us to perform this study. In Section 1.2 we explain contributions in this work. Finally, we explain the organization of the thesis in Section 1.3.

1.1 Motivation

It is estimated that the communication and information technology (ICT) sector consumes 8 percent of the total electricity generation of the world and is responsible for 2.7 percent (1.43 Gt per year) of total global CO_2 emissions [32, 4, 5, 6]. Studies show that the energy consumption of the ICT sector has been growing rapidly in the past decade and will continue to grow [4, 5, 6]. For example, the ICT sector's contribution to the whole electricity consumption of the world was 3.9 and 4.6 percent in 2007 and 2012, respectively [7]. Researchers predict that the ICT sector can contribute to the total CO_2 emission by up to 23 percent in 2030 [5].

It is estimated that communication networks are responsible for 15–20 percent of the total energy consumption of the ICT sector [18, 4]. Among different parts of the communication networks, core networks consume 30 percent of the total energy consumption of communication networks [40, 25]. Also, it is expected that the core networks become the major power-consuming and therefore CO_2-emitting part of communication networks since the annual increase in global traffic is more than the increase in the energy efficiency of core network devices [15, 25].

Therefore, numerous studies have been conducted to decrease the core networks' carbon footprint either by proposing energy-aware traffic engineering methods or by routing traffic through paths whose energy resources are mostly renewable [37, 22, 41, 34]. However, to the best of our knowledge, all these approaches can be applied only to the *intra-domain* routing scope or, at best, between two neighboring autonomous systems (ASes)[1]. The main problem with such methods is that they cannot find paths with the lowest emission between every two ASes in the Internet. Furthermore, most prior work only focuses on either the energy efficiency of the network or the type of energy resources of network devices. However, the carbon emission of a network depends on both of these

---

[1] In this work we use *domain* and *AS* interchangeably. Also most ASes in this work are *internet service providers* (ISPs)

variables, and focusing on either of them cannot find the optimal carbon-efficient paths.

1.2 Contributions

In this work, we propose a new method for disseminating information about the emission of inter-domain core paths as well as selecting and propagating the paths with the lowest CO_2 emission toward every destination AS in the SCION Internet architecture [35]. SCION is a path-aware next-generation Internet architecture [2] that provides end hosts with multiple paths to every destination in the Internet and provides connectivity using a packet-carried forwarding state, meaning that end hosts put the path they want their packets to take in every packet and ASes forward packets based on the path specified in their header. The path exploration process in SCION provides ASes with a framework that enables them to encode different types of information about the part of an inter-domain path that crosses them into routing messages. Each AS can take this information into account for selecting and propagating paths to its neighbors. Different types of applications running on end hosts can use this feature to select optimal available paths based on their needs.

This property of SCION enables us to encode information about the CO_2 emission of an inter-domain path in PCBs and inform end hosts about the amount of CO_2 emission they are responsible for by selecting each path. To measure the amount of CO_2 that sending a bit of data on a path emits, we need a model for the energy consumption of devices on a path inside the core network. The CO_2 emission of a device is then computed as the multiplication of the average CO_2 emission of its electricity sources and its energy consumption. To the best of our knowledge, all prior studies that provide a model for energy consumption of the Internet or the core network try to estimate the energy intensity of the whole Internet or core network, and they do not propose methods for estimating the energy consumption of a single path. Therefore, we propose a model for the per-bit energy consumption and CO_2 emission of a single path inside the core network. Based on that model, we propose a method for ASes to calculate the CO_2 emission of their intra-domain paths that are parts of inter-domain paths. Furthermore, we propose an algorithm for selecting and propagating the paths with the lowest CO_2 emission toward every destination AS.

We test our method using simulations on a realistic topology and show that this method decreases the CO_2 emission by at least 50 percent for paths between more than half of the AS pairs. Moreover, by modeling the traffic between different ASes in this set of ASes, we show that this model can reduce CO_2 emission of this network by 210 000 ton/year.

Furthermore, as SCION enables end-users to select paths based on their desired properties, we expect that selection of greener paths by some users shifts some traffic from ASes that are not energy efficient or do not use renewable energy sources, which causes competition between ASes to make themselves more energy efficient or increase their use of renewable energy resources. We propose a model for this competition, and, by simulating it, we show that this leads to a state where the total CO_2 emission of the whole network is 87 percent less than the initial state in which all ASes use the same energy resources as they use in BGP. Furthermore, using simulations on a smaller topology, we perform a sensitivity analysis on this model by changing the values of the parameters used in it. The results show that for the majority of parameter sets, the competition leads to a state where all ISPs that carry traffic between other ASes use 100 percent renewable electricity resources and the total CO_2 emission of the whole network is 53 percent less than the initial state.

1.3 Thesis organization

The thesis is organized as follows. Chapter 2 provides a background of the SCION next-generation Internet architecture and the Internet's model which is mainly used for analyzing the energy consumption. In Chapter 3, we propose a model for the CO_2 emission of inter-domain paths per bit of data, then we explain how each AS on an inter-domain path can calculate the CO_2 emission of the part of the path that crosses its network using the model that we proposed. Furthermore, in this section, we explain how the CO_2 emission information will be encoded in routing messages. Then, we propose an algorithm for constructing the set of paths with the lowest CO_2 emission to every destination. In Chapter 4, we explain how we evaluate our proposed method. In Chapter 5, we analyse the results of our evaluations. In Chapter 6, we discuss related work. Finally, in Chapter 7 we conclude our work.

2 Background

In this chapter we provide an overview of SCION next-generation Internet architecture in Section 2.1 and introduce the model which researchers mostly us for analyzing the energy consumption of different parts of the Internet in Section 2.2.

2.1 SCION

SCION is a next-generation Internet architecture that provides security, high availability, isolation, and scalability. SCION provides end hosts with multiple authentic inter-domain paths towards every destination in the Internet. In SCION, every packet carries its own inter-domain forwarding path, selected by the packet's sender. SCION border routers forward packets on the path specified by the end host. In this section, we provide an overview of the SCION architecture, then describe the path-construction, path-registration, and path-lookup procedures in SCION, and finally explain how packets are forwarded in this Internet architecture. In this work, we use the SCION book [35] as the reference for the SCION Internet architecture except where we explicitly state that we use another reference.

2.1.1 SCION architecture overview

To provide isolation and scalability, SCION introduces *isolation domains* (ISDs). An ISD consists of multiple ASes that agree on a set of trust roots defined in a *trust root configuration* (TRC). The TRC defines the roots of trust that are used to validate bindings between names and public keys or addresses. Each ISD is governed by a set of ASes known as core ASes, which are responsible for providing connectivity to other ISDs as well as governing the trust roots.

In SCION, there are three kinds of relations between ASes: core, provider-customer, and peer-peer. Core relation is the relation between core ASes, and provider-customer and peer-peer relations are the relations between two ASes of which at least one is not a core AS.

An address in SCION is a tuple of ⟨*ISD number, AS number, local address*⟩, in which ISD number is the ISD identifier where the end host resides, the AS number is the AS identifier where the end host resides, the local address is the host's identifier inside its AS. The local address is not used for inter-domain routing or forwarding.

Each SCION AS needs to run multiple services to provide SCION connectivity: a beacon service, which discovers paths; a path service, which disseminate path information; a certificate service, which assists with validating path information; and a name service that provides name resolution to SCION addresses.

SCION provides secure inter-domain routing by introducing a control-plane public-key infrastructure (CP-PKI), which distributes public keys used for authentication of SCION paths (and other control-plane messages). Each AS signs the paths it advertises by its private key, and other ASes check the signature using the public key of that AS.

2.1.2 Path construction and registration

The SCION control plane is responsible for generating paths between every two ASes. Each SCION path consists of at least one and at most three path segments. SCION path segments are of three types: up-segment, core-segment, and down-segment. An up-segment specifies the path from a customer AS to a provider AS by crossing only provider-customer links from customer to the provider. A core-segment specifies a path from a core AS to another core AS by crossing only core links, and a down-segment specifies a path from a provider AS to one of its direct or indirect customers by crossing only provider-customer from provider to customer.

The granularity of SCION path segments is inter-AS interfaces that connect neighbor ASes. The inter-AS interfaces of an AS are the interfaces on its border routers that connect its border routers to another AS's border routers. This property gives the opportunity of having multiple paths for the same AS-level path.

Path segments in SCION are constructed in a process called *beaconing* by sending control-plane messages called *path-segment construction beacons* (PCBs). The beaconing process is conducted on two hierarchy levels: (1) among all core ASes in all ISDs known as *core beaconing*, and (2) within each ISD known as *intra-ISD beaconing*.

Core beaconing

In core beaconing, the beacon service of each core AS periodically initiates PCBs for all of its interfaces connecting it to its core neighbor ASes, signs them with its private key from the CP-PKI, and sends them to border routers connected to those interfaces. When a border router receives a PCB from its local beacon service, it sends out the PCB on the egress interface specified by the PCB. Upon reception of a core PCB from a core neighbor AS the border router of a core AS sends the received PCB to the beacon service of its AS. The beacon service first verifies the PCB. Then it decides whether to register the path specified by the PCB to the AS's path service or not. It also decides whether to import the PCB into the beacon store or not. The beacon store is a data structure for storing PCBs that are candidates for propagating to core neighbor ASes. The beacon service of each core AS periodically selects the best-received core PCBs per destination core AS and selects the best egress interfaces to propagate those PCBs further to its core neighbor ASes. In the case of deciding to propagate a PCB further to an interface, the beacon service appends its own AS entry to the path segment of the PCB, computes its signature over all AS entries in the PCB using its private key, and sends it to the border router connected to the egress interface.

Intra-ISD beaconing

Intra-ISD beaconing is similar to core beaconing except that PCB propagation's direction is from providers to customers. In intra-ISD beaconing, only the beacon services of core ASes initiate beaconing, and beacon services of non-core ASes only propagate PCBs to their customers. In intra-ISD beaconing, ASes with peering links with other ASes can embed their peering links in their AS entry that they append to the PCB. If a beacon service decides to register a path specified by a non-core PCB as an up-segment, it creates an up-segment from the PCB's path and registers it into the local path service. Also, if it decides to register the path as a down-segment, it creates a down-segment out of the PCB's path and registers it to the path service of the core AS that initiated the PCB. Hosts in local AS use the up-segment to reach the core AS that initiated the PCB. Hosts in upstream ASes use the down-segment to reach the local AS.

PCB format

Figure 2.1 shows the structure of a PCB for propagating a path segment. Each PCB has one *InfoField* and multiple *ASEntries*. The *InfoField* contains some flags, its timestamp, the initiator ISD number, and the length of the path segment. Each *ASEntry* contains necessary information about each AS hop that will be encoded in the data packets for forwarding, information about peering links in intra-ISD beaconing, the signature of the AS, and extra information about the AS hop. The *HopField* contains information like ingress and egress interfaces, expiration time, and the hop field's *message authentication code* (MAC). The MAC is calculated over the *HopField* by the AS's secret key which is not shared with other ASes and is used by border routers for forwarding packets to check whether this *HopField* was really generated by this AS or not.

SCION PCBs can disseminate additional information about each AS hop on a path segment by adding this information into the *Extensions* field of the corresponding *ASEntry* in the PCB. The *Extensions* field is used to signal extra information about each AS hop. There are multiple types of these extensions defined in SCION. One of them is the *StaticInfoExtension* which is designed for carrying static information about AS hops on a path segment such as latency, maximum bandwidth, etc. This extension type is not defined in the SCION book and we use SCION development documents as the reference [20].

The beacon services in SCION ASes can use the information encoded in PCBs about paths to decide which paths to register in the path services and which paths to propagate further to neighbor ASes. They can optimize path for different criteria on different path quality metrics such as latency or bandwidth. The optimization for different criteria can be performed jointly or in parallel in each beacon service and depends on the local implementation of the beacon service.

2.1.3 Path lookup

Whenever a host h_A inside SCION AS A wants to send packets to a remote host h_B in another SCION AS B, it must know h_B's ISD number, its AS number, and its local address inside B. Therefore, name resolution in SCION results in ⟨*ISD number, AS number, local address*⟩ tuples. Then, h_A asks A's path service for a path towards h_B's ISD and AS. If the path towards B is registered or cached in A's path service, it returns the path to h_A. Otherwise, it asks the path service of a core AS in its own ISD for a path towards B. If B is in the same ISD as A or is a core AS in another ISD, the path towards B has been registered in the core path service. Therefore, the core path service returns a path segment towards B to A's path service. Otherwise, the local ISD's core path service asks the remote core path service of B's ISD for a path towards B. The remote core path service returns a down-segment towards B to the local core path service. The local core path service returns the core-segment to reach the remote core AS as well as the down-segment

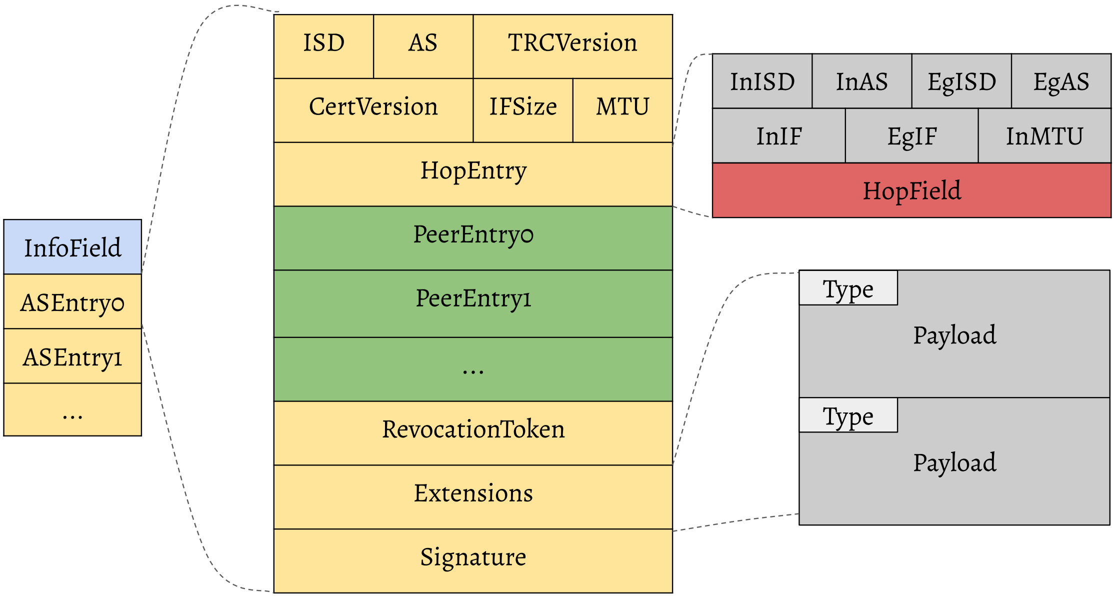

Figure 2.1: Structure of a PCB in SCION. Extensions can disseminate metadata about each AS hop. The figure is adapted from the Fig. 15.13 of the SCION book [35].

to reach *B* from the remote core AS. This process is called *path lookup*. Many path lookup queries can be resolved faster by caching paths to popular destinations in the local host, in the local path service, and in the core path service.

2.1.4 Path header in data packets

A SCION data packet contains its forwarding path in its header. The path consists of at least one and at most three segments. An up-segment can be followed by either a core-segment or a down-segment. A core-segment can be followed by only a down-segment. A down-segment cannot be followed by another segment. Each segment in a path consists of one info field and at least one hop field. Every SCION border router forwards each SCION packet based on the ingress and egress interfaces identified in the packet's current hop field.

2.1.5 Differences between SCION and BGP

In this section we describe some of the most fundamental differences between SCION and the Border Gateway Protocol (BGP)—the current Internet's inter-domain routing protocol that make the inter-domain carbon-aware routing possible in SCION.

SCION can provide fine-grained information about paths

Since the granularity of SCION path segments is inter-domain interfaces, SCION can provide more fine-grained information about an inter-domain path than BGP, in which the granularity of inter-domain paths is AS numbers. Using an example, we illustrate the benefits of having interface-level paths. Assume three ASes *A*, *B*, and *C* where *A* is connected to *B* in two different locations and *C* is also connected to *B* in two different locations, but *A* and *C* are not connected directly to each other. Therefore, there are four different interface-level paths and only one AS-level path between *A* and *C* via *B* in each direction. Since these four interface-level paths connect four different locations in *B*'s network, they have different properties such as different latency, different maximum bandwidths, different energy consumption per bit, etc. Since SCION differentiates these four paths, encoding each path's different metrics information into AS entries of the PCB during the beaconing process helps beacon services and end hosts to select the highest-quality path out of multiple interface-level paths.

SCION provides multiple paths

Multi-path routing in the current Internet's inter-domain routing protocol (BGP) is not scalable because it needs to store multiple paths per entry in each router's forwarding state. Currently, there are more than 800 000 entries in the complete

BGP forwarding table [27]. By storing only two paths per destination, the forwarding table's size exceeds the size of the memory dedicated to the forwarding table in most routers. Routers store forwarding tables in high-speed memory called ternary content addressable memory (TCAM), so they can look up the next hop at the line speed. TCAMs are expensive and have small capacities; therefore, storing multiple paths per destination in a TCAM is impossible. In contrast, SCION packets carry their own forwarding state, and SCION border routers forward packets based on the forwarding path encoded in their header. Therefore, SCION routers do not need to perform look-ups in a (global) forwarding table. Having forwarding-table-free border routers enables multi-path routing in SCION. Furthermore, in SCION, end hosts select the forwarding path among multiple paths depending on their path-selection policy and based on information encoded in the PCB about different metrics. For example, one host needs high-bandwidth paths towards an AS, while the other host in the same AS needs low-latency paths towards the same destination. While SCION can satisfy the requirements of both of these hosts, none of the BGP multi-path extensions can satisfy these requirements because in these extensions, each router locally decides on which path to forward traffic regardless of source requirements.

SCION is capable of constructing all possible paths

BGP routers usually adhere to the Gao-Rexford routing policy [21]. This policy suggests that ASes prefer the paths they learn from their customers to the paths they learn from their peers and prefer the paths they learn from their peers to the paths they learn from their providers. In this policy, ASes advertise paths learned from customers to all their neighbors and advertise paths learned from their providers and peers only to their customers. There are two main reasons for adopting Gao-Rexford policy: (i) it adheres to the commercial relationships between ASes since customers pay providers for the traffic they send to them, but peers usually do not pay each other, and (ii) it guarantees deterministic convergence of global routing state among all BGP speakers. Using this policy, some available paths are not advertised even if they have desirable properties for end domains.

In SCION intra-ISD beaconing, the Gao-Rexford policy is respected due to commercial relations. However, as SCION border routers do not have routing and forwarding table, they are stateless, and there is no need for convergence among border routers. Therefore, the Gao-Rexford policy can be violated in SCION inter-ISD beaconing and all possible paths can be discovered. So, core paths that cannot be advertised in BGP can be constructed in SCION as Scherrer et al. showed in their work [39]. They also showed that violating Gao-Rexford policy can make sense economically. As core ASes are highly inter-connected, this leads to discovering a variety of paths with different properties.

2.2 Overview of the Internet and core networks

In this section we provide an overview of the model that researchers mostly used for analyzing the energy consumption of different parts of the Internet and core networks.

Internet model

The standard method for analyzing the Internet energy consumption is to divide the network into access, metro/edge, and core networks [25, 26, 40]. Figure 2.2 shows this segmentation of the Internet. Access networks connect end hosts to their local exchanges using a wide range of technologies like digital subscriber line (DSL), which uses telephone copper pairs; fiber to the premises (FTTP), which uses optical links; or wireless technologies such as WiMAX. Local exchanges in the same city or different cities are connected through metro and edge networks. Metro and edge networks also connect access networks to the core network. The main devices used in metro and edge networks are edge Ethernet switches, broadband network gateways (BNG), and edge routers. Edge Ethernet switches are connected to many access networks and a few BNGs concentrating traffic from access networks and sending it to BNGs. BNGs provide rate control, authentication, and security services and send traffic to provider edge routers. These routers are connected to the core network.

The core network provides connectivity to any network in the Internet. It consists of few large routers (core routers) interconnected by wavelength-division multiplexed (WDM) fiber links. Core routers within each provider's network are highly meshed while they have a few links to other providers' networks [25]. Every inter-domain path in the Internet only consists of network devices in the core network and edge routers in the edge network with a similar energy consumption model to the core network. Therefore, to model the energy consumption of inter-domain paths, we need to analyze the path's energy consumption from the point a packet enters the core network to the point it leaves the core network.

Core network

The core network is an optical multi-layer network known as IP-over-WDM network [24], which is divided into two layers: (i) the Internet Protocol / Multiprotocol Label Switching (IP/MPLS) layer and (ii) the optical layer [3].

The IP/MPLS layer consists of core routers such as the Cisco CRS series and Juniper T-series and performs layer-3 routing and forwarding. Every core router has three main building blocks [24]:

- *basic node*, which contains chassis, switching fabrics, routing engine, internal cooling equipment, and power supply;
- *line cards*, which contains the forwarding engine which forwards packets using the forwarding table populated by the routing engine; and

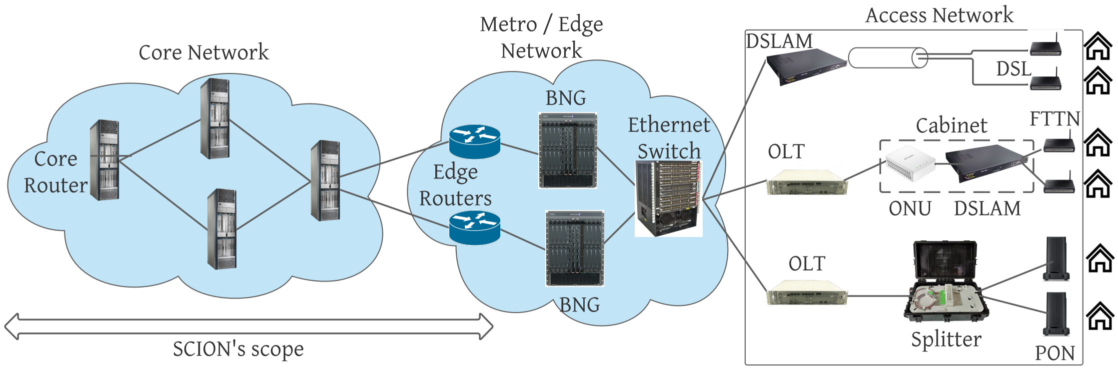

Figure 2.2: Core, metro/edge, and access networks in the Internet. SCION's scope only contains the core and (part of the) edge networks. This figure is adapted from Fig. 1 in Hinton et al. [25].

- *port cards*, which provide the physical connection to the optical layer and provide an interface between the Ethernet layer (layer 2) and the IP/MPLS layer (layer 3).

The WDM layer (also known as optical layer) provides a broadband physical connection between core routers. The optical layer contains the following devices [24]:

- *Transceivers*, which convert an optical signal to an electrical signal and vice versa. They are included in the port cards of the router.
- *Transponders*, which convert different optical wavelengths to each other. They typically convert a gray optical signal (1300 nm) from/to a dense-WDM-band one (1500 nm).
- *Muxponders*, which convert lower-rate signals to higher-rate WDM signals.
- *WDM terminal systems*, which multiplex individual channels into one into the fiber and de-multiplex the signal from fiber to multiple individual channels. They are usually characterized by the number of channels they support.
- *Optical switches (OADMs, OXCs)*, which switch wavelength channels directly in the optical layer without converting the optical signal to electrical and vice versa. Optical add-drop multiplexers (OADM) provide two bidirectional transit fiber ports, while optical cross-connects (OXCs) provide more than two transit fiber ports and can cross-connect wavelength channels.
- *Optical line amplifiers (OLAs)*, which amplify optical signals approximately after every 80 km.

- *Regenerators*, which regenerate optical signals by re-timing, re-shaping and re-transmitting them. Typically, the optical signal needs to be regenerated every 1500 km. This distance can be increased or decreased depending on fiber quality, transponder type, data rate, and modulation.

3 Green routing in SCION

SCION gives us the opportunity to reduce the Internet's CO_2 emissions by providing multiple network paths to destinations, and giving ASes and end hosts the opportunity to select paths based on their needs. In this section, we introduce a SCION-based approach to construct green paths (the ones with the lowest CO_2 emission) to every destination. In Section 3.1, we develop a model for calculating CO_2 emission of inter-domain paths. In Section 3.2, we provide a method for embedding emission information in SCION PCBs. In Section 3.3, we introduce methods to import and construct green path segments and create green paths from green path segments. Finally, in Section 3.4, we discuss the effect of advertising green paths on the profit of ISPs and how the competition between different ISPs to become greener can make the whole Internet greener.

3.1 Modeling CO_2 emission of inter-domain paths

To design a global CO_2-efficient routing and forwarding approach based on SCION, we need a model for CO_2 emission of inter-domain paths in the Internet backbone. An inter-domain path is the sequence of network devices each data packet traverses from the point it exits from the source AS until the point it enters the destination AS. In this section, we introduce a model for estimating the CO_2 emission of the Internet inter-domain paths. The model that we propose, can also be applied to intra-domain communications that cross core paths; however, since intra-domain routing and forwarding is out of the scope of SCION, we focus on inter-domain communications. The amount of CO_2 emitted by sending a bit of data on an inter-domain path depends on (i) the amount of electrical energy that the network devices on the path consume for sending that bit of data and (ii) the electricity source mix of each device on the path that determines the amount of CO_2 emitted per unit of electrical energy consumed by that device. Therefore, we introduce a model for energy consumption of inter-domain paths; then, we model CO_2 emission based on the energy consumption model.

3.1.1 Per-bit energy consumption in core paths

Every packet entering the core network from an edge network to reach another edge network is forwarded by a sequence of core routers until it reaches its exiting point of the core network. Every two consecutive core routers in this sequence are usually connected by logical layer-3 links. The optical layer provides logical layer-3 connectivity between each two neighbor core routers. Therefore, every packet forwarded by a core router to reach the neighboring core router crosses a variable number of optical devices before arriving at the neighboring core router. Every device on this path from the core network's ingress point until its egress point consumes some amount of energy to process or transfer the packet. The amount of energy consumed for sending a packet through the core network is computed as

$$E_{path}(S_{\mathrm{pkt}}) = \rho_{\mathrm{pr}}\left[\sum_{R\in path.routers} \rho_{R,\mathrm{c}} E_R(S_{\mathrm{pkt}}) + \sum_{OD\in path.optdevices} \rho_{OD,\mathrm{c}} E_{OD}(S_{\mathrm{pkt}})\right]. \quad (3.1)$$

In this equation, *path* denotes a path in the core network and is the sequence of network devices in all layers including routers and optical devices that a packet crosses on that path. *path.routers* is the subsequence of routers in the *path*, and *path.optdevices* is the subsequence of optical devices in the *path*. R and OD denote a core router and an optical device, respectively. We consider both routers and optical devices as network devices and use the unified symbol D for them when we do not differentiate between them. The ρ_{pr} is the provisioning factor for traffic protection and is equal to the number of primary and backup paths between the ingress and the egress points of the network. Backup paths are used in the case of the primary path failure, and usually no traffic is sent over them but they still consume energy. The $\rho_{R,\mathrm{c}}$ and $\rho_{OD,\mathrm{c}}$ are the cooling and facilities provisioning factor of routers and optical devices at their points of presence, respectively.[1] S_{pkt} is the packet size (in bit). $E_R(S_{\mathrm{pkt}})$ is the amount of energy a core router consumes to process a packet and breaks down into per-packet forwarding energy consumption ($E^{\mathrm{fwd}}_{R,\mathrm{pkt}}$) and per-bit switching fabric energy consumption ($E^{\mathrm{S\&F}}_{R,\mathrm{bit}}$). $E^{\mathrm{fwd}}_{R,\mathrm{pkt}}$ is the amount of energy consumed by the forwarding engine in the line cards to perform per-packet forwarding and is independent of the packet size. $E^{\mathrm{S\&F}}_{R,\mathrm{bit}}$ is the amount of energy consumed mainly by switching fabrics and buffering and depends on the packet size. Therefore, $E_R(S_{\mathrm{pkt}})$ for a packet of size S_{pkt} can be written as

$$E_R(S_{\mathrm{pkt}}) = S_{\mathrm{pkt}} \times E^{\mathrm{S\&F}}_{R,\mathrm{bit}} + E^{\mathrm{fwd}}_{R,\mathrm{pkt}}. \quad (3.2)$$

$E_{OD}(S_{\mathrm{pkt}})$ is the amount of energy an optical device consumes to process a packet and depends on the size of the packet and the energy consumption of the optical device per bit. So, it can be written as

$$E_{OD}(S_{\mathrm{pkt}}) = S_{\mathrm{pkt}} \times E_{OD,\mathrm{bit}}. \quad (3.3)$$

We denote the per-bit energy consumption of both routers and optical devices as $E_{D,\mathrm{bit}}$.

---

[1] For a network device, the cooling and facilities provisioning factor is equal to the ratio of the power consumption of the building that the network device is located in to the actual power consumption of all network devices in that building. This overhead is usually characterized by *power usage effectiveness* (PUE) [10].

Equation (3.1) gives the actual amount of energy consumed for sending a packet over a core path; however, network devices consume considerable amount of energy when they are completely idle. It is estimated that idle power consumption of network devices, especially routers, can be responsible for up to 90 percent of the their peak power consumption [50]. So, not including the idle power consumption of network devices can cause considerable error. Therefore, instead of modeling the actual amount of energy consumed for sending a packet on a core path, we model the expected amount of energy consumed by sending a packet on a core path by including idle power consumption of devices on the path and re-write Eq. (3.1) as

$$\overline{E}_{path}(\overline{S}_{\mathrm{pkt}}) = \rho_{\mathrm{pr}}\left[\sum_{R \in path.routers} \rho_{R,\mathrm{c}}\overline{E}_R(\overline{S}_{\mathrm{pkt}}) + \sum_{OD \in path.optdevices} \rho_{OD,\mathrm{c}}\overline{E}_{OD}(\overline{S}_{\mathrm{pkt}})\right] \tag{3.4}$$

where $\overline{S}_{\mathrm{pkt}}$ is the average size of packets (in bit) and $\overline{E}_R(\overline{S}_{\mathrm{pkt}})$ and $\overline{E}_{OD}(\overline{S}_{\mathrm{pkt}})$ are expected amounts of energy consumed for processing a packet on a router and an optical device, respectively. More formally, we define them as

$$\overline{E}_R(\overline{S}_{\mathrm{pkt}}) = \frac{P_R}{l_{R,\mathrm{pkt}}} = \overline{S}_{\mathrm{pkt}} \times \frac{P_R}{l_{R,\mathrm{bit}}}, \tag{3.5}$$

and

$$\overline{E}_{OD}(\overline{S}_{\mathrm{pkt}}) = \frac{P_{OD}}{l_{OD,\mathrm{pkt}}} = \overline{S}_{\mathrm{pkt}} \times \frac{P_{OD}}{l_{OD,\mathrm{bit}}}, \tag{3.6}$$

where P_R and P_{OD} denote the power consumption of a router and an optical device at a given time, and $l_{R,\mathrm{pkt}}$ and $l_{OD,\mathrm{pkt}}$ denote the traffic load on each device in packet/s at the same time. The unit of $\overline{E}_R(\overline{S}_{\mathrm{pkt}})$ and $\overline{E}_{OD}(\overline{S}_{\mathrm{pkt}})$ is J and their values can vary over time.

The total power consumption of a network device is the the sum of its idle power consumption ($P_{D,\mathrm{idle}}$), which is the power consumption of the device when there is no traffic load on it, and its traffic-dependant power consumption. For core routers we write the total power consumption as

$$P_R = P_{R,\mathrm{idle}} + l_{R,\mathrm{bit}} \times E^{\mathrm{S\&F}}_{R,\mathrm{bit}} + l_{R,\mathrm{pkt}} \times E^{\mathrm{fwd}}_{R,\mathrm{pkt}}, \tag{3.7}$$

where $P_{R,\mathrm{idle}}$ is the idle power consumption of the router. For optical devices we write the total power consumption as

$$P_{OD} = P_{OD,\mathrm{idle}} + l_{OD,\mathrm{bit}} \times E_{OD,\mathrm{bit}}, \tag{3.8}$$

where $P_{OD,\mathrm{idle}}$ is the idle power consumption of the optical device.

By dividing both sides of Eq. (3.7) by $l_{R,\text{pkt}}$, and both sides of Eq. (3.8) by $l_{OD,\text{pkt}}$ we reach the following formula for $\overline{E}_R(\overline{S}_{\text{pkt}})$ and $\overline{E}_{OD}(\overline{S}_{\text{pkt}})$:

$$\overline{E}_R(\overline{S}_{\text{pkt}}) = \frac{P_R}{l_{R,\text{pkt}}} = \overline{S}_{\text{pkt}} \times \left(\frac{P_{R,\text{idle}}}{l_{R,\text{bit}}} + E^{\text{S\&F}}_{R,\text{bit}}\right) + E^{\text{fwd}}_{R,\text{pkt}}, \tag{3.9a}$$

$$\overline{E}_{OD}(\overline{S}_{\text{pkt}}) = \frac{P_{OD}}{l_{R,\text{pkt}}} = \overline{S}_{\text{pkt}} \times \left(\frac{P_{OD,\text{idle}}}{l_{OD,\text{bit}}} + E_{OD,\text{bit}}\right). \tag{3.9b}$$

In Eq. (3.9a), we can also use the actual packet size (S_{pkt}) instead of the average packet size to calculate the expected consumed energy for a single packet with a particular size. However, we do not use it in this work since we want to optimize paths for flows with different packet sizes.

The term $\frac{P_{D,\text{idle}}}{l_{D,\text{bit}}}$ in Eq. (3.9a) distributes the idle power consumption of a device over traffic that it receives and its unit is J/bit. We call this term *share of idle energy* (*SIE*) per bit. It determines the amount of idle energy consumption that a single bit of data is responsible for. The *SIE* metric does not only take the energy efficiency of a single device into account but also the energy efficiency of the whole network. The smaller the *SIE* is, the more energy-efficient the network is. For example, in networks using green traffic-engineering mechanisms, they aggregate traffic from underutilized paths into a few paths; therefore, some links are offloaded, and the line cards connected to them are turned off. In such networks, the *SIE*s of both offloaded routers and the routers that receive more traffic decrease. In offloaded routers, some line cards are turned off, and therefore the idle power of the routers decreases, so the numerator of the fraction is decreased. In routers that receive more traffic, the current load, which is the *SIE* fraction's denominator, increases, so the fraction's value decreases.

By substituting Eqs. (3.9a) and (3.9b) into the Eq. (3.4) we reach the following formula for calculating the expected energy consumption for sending a packet on a core path:

$$\overline{E}_{path}(\overline{S}_{\text{pkt}}) = \rho_{\text{pr}} \left[\overline{S}_{\text{pkt}} \times \sum_{D \in path} \rho_{D,\text{c}} \left(E_{D,\text{bit}} + \frac{P_{D,\text{idle}}}{l_{D,\text{bit}}}\right) + \sum_{R \in path.routers} \rho_{R,\text{c}} E^{\text{fwd}}_{R,\text{pkt}}\right]. \tag{3.10}$$

By dividing Eq. (3.10) by the packet's size, we obtain Eq. (3.11) to model the amount of energy consumed per bit of a packet that is forwarded on a particular path in the core network.

$$\overline{E}_{path,\text{bit}}(\overline{S}_{\text{pkt}}) = \rho_{\text{pr}} \left[\sum_{D \in path} \rho_{D,\text{c}} \left(E_{D,\text{bit}} + \frac{P_{D,\text{idle}}}{l_{D,\text{bit}}}\right) + \frac{1}{\overline{S}_{\text{pkt}}} \sum_{R \in path.routers} \rho_{R,\text{c}} E^{\text{fwd}}_{R,\text{pkt}}\right] \tag{3.11}$$

3.1.2 Per-bit CO_2 emission in core paths

The amount of CO_2 emission per one bit of data (*CEPB*) that is sent on each path in the core network is equal to the sum of CO_2 emissions of every device on that path for processing that bit of data. The per-bit CO_2 emission of every network device is equal to the product of the expected per-bit energy consumption of that device ($\overline{E}_{D,\text{bit}}(\overline{S}_{\text{pkt}})$) and the carbon intensity (CI_D) of its electricity source, which is the CO_2 emission per unit of electrical energy it generates. Each network device may have multiple electrical energy sources with different *CI*s. Therefore, the *CI* of the device is obtained by calculating the weighted mean of *CI*s of all its electricity sources in which each source's weight is its share of the total electricity consumption of the device. Furthermore, the *CI* of a device can change over time depending on its energy sources. For example, if a router is powered by solar electricity during the day and by coal during the night, its *CI* drops significantly after sunrise and increases significantly after sunset.

$$CEPB_D(\overline{S}_{\text{pkt}}) = CI_D \times \overline{E}_{D,\text{bit}}(\overline{S}_{\text{pkt}}) \tag{3.12a}$$

$$CI_D = \sum_{s \in sources} W_s \times CI_s \quad \text{where} \quad \sum_{s \in sources} W_s = 1 \tag{3.12b}$$

By combining Eqs. (3.11) and (3.12a), we obtain the following equation for calculating the amount of CO_2 emission per bit of data on every core path:

$$\begin{aligned} CEPB_{path}(\overline{S}_{\text{pkt}}) = \rho_{\text{pr}} \Bigg[& \sum_{D \in path} \rho_{D,\text{c}} CI_D \left(E_{D,\text{bit}} + \frac{P_{D,\text{idle}}}{l_{D,\text{bit}}} \right) \\ & + \frac{1}{\overline{S}_{\text{pkt}}} \sum_{R \in path.routers} \rho_{R,\text{c}} CI_R E^{\text{fwd}}_{R,\text{pkt}} \Bigg]. \end{aligned} \tag{3.13}$$

As CI_D and $l_{D,\text{bit}}$ vary over time, the *CEPB* of the core path also varies over time.

Modeling *CEPB* of inter-domain paths using the *CEPB* of intra-domain paths

Since SCION path-segments are constructed incrementally by adding each AS's *ASEntry* to the PCB, and each AS can only calculate the *CEPB* of its own internal paths, we need to change our model such that it can estimate the *CEPB* of the whole inter-domain path using the *CEPB* of the intra-domain paths of which it consists.

As the core network consists of multiple carrier ASes, every packet that enters the core network may cross multiple carrier ASes along the path towards its destination. Therefore, an inter-domain path's per-bit CO_2 emission equals the sum of per-bit CO_2 emissions of intra-AS paths inside every carrier AS on the path. These intra-AS paths provide connectivity between the border router that a packet enters an AS (ingress router) and the border router that the packet exits the AS (egress router). We use the $[ing, eg]$ notation for the intra-AS path between ingress and egress routers in each AS.

$$CEPB_{path}(\overline{S}_{\text{pkt}}) = \sum_{AS \in AS\ path} CEPB_{AS,[ing,eg]}(\overline{S}_{\text{pkt}}) \tag{3.14}$$

The per-bit CO_2 emission of every intra-AS path between each pair of border routers of an AS is calculated using Eq. (3.15) which is similar to Eq. (3.13) but the entering and exiting points are border routers of each AS.

$$\begin{aligned} CEPB_{AS,\ [ing,\ eg]}(\overline{S}_{\text{pkt}}) = \rho_{\text{pr}} \Bigg[& \sum_{D \in [ing,eg]} \rho_{D,\text{c}} CI_D \left(E_{D,\text{bit}} + \frac{P_{D,\text{idle}}}{l_{D,\text{bit}}} \right) \\ & + \frac{1}{\overline{S}_{\text{pkt}}} \sum_{R \in [ing,eg].routers} \rho_{R,\text{c}} CI_R E^{\text{fwd}}_{R,\text{pkt}} \Bigg] \end{aligned} \tag{3.15}$$

Having a carbon-emission model for inter-domain paths at the granularity of inter-domain interfaces on the path allows us to disseminate this information using SCION PCBs as they construct paths in the same granularity.

3.2 Disseminating *CEPB* of SCION paths

In this section, we explain how we disseminate greenness information of paths in SCION. We also propose a method by which ASes can calculate their internal paths' carbon emission between their border routers based on the model we proposed in the previous section.

3.2.1 Embedding *CEPB* into PCBs

To disseminate information about a path segment's greenness, we introduce a new metadata type for the *StaticInfoExtension*. This metadata type, which we call *greenness metadata*, disseminates the *CEPB* of intra-AS paths between the egress interface of an *ASEntry* and other interfaces of the AS. This metadata is a map where its keys are integers specifying interface numbers and values are doubles equal to the *CEPB* of the intra-AS path between the egress interface of the *ASEntry* and the interface specified by the key. At every propagation or initiation interval, the beacon service of an AS selects a set of interfaces and embeds the *CEPB*s of internal paths from each of them to the egress interface of a PCB. The selection of these interfaces depends on whether the AS is initiating the PCB and whether the PCB contains a core path segment or not. Therefore, there are four different cases:

1. The core beacon service is the initiator of a core PCB. In this case, it does not embed any information about the *CEPB*.

2. The core beacon service is not the initiator of a core PCB and propagates it to its core neighbors. In this case, it only embeds the *CEPB* of the internal path between the ingress and egress interfaces of its own *ASEntry* that appends to the PCB.

3. The core beacon service is initiating an intra-ISD PCB. In this case, it embeds the *CEPB*s of the internal paths between the egress interface of the PCB and all other interfaces that connect the core AS to its neighbors except the ones connecting it the AS to which the PCB is propagating to. This information is needed to construct the greenest combination of up- and core-segments, the greenest combination of core- and down-segments, the greenest path containing a shortcut, and the greenest path containing a peering link. When an up-segment and a down-segment intersect at an AS, a shortcut is formed.

4. The non-core beacon service is propagating an intra-ISD PCB to its AS's customers. In this case, it embeds the *CEPB* of internal paths connecting the egress interface of the PCB to the ingress interface of the PCB, interfaces connecting the AS to its peers, and interfaces connecting the AS to its other customers. This information is needed to construct the the greenest path containing a shortcut at the non-core AS, and the greenest path containing a peering link.

3.2.2 Greenness measurement

To calculate the *CEPB* of an intra-AS path between two inter-domain interfaces the beacon service needs the values of CI_D, $E_{D,\mathrm{bit}}$, $P_{D,\mathrm{idle}}$, $l_{D,\mathrm{bit}}$ of all devices on the intra-AS path, and $E^{\mathrm{fwd}}_{R,\mathrm{pkt}}$ of all routers on the intra-AS path. Also, it needs the values of ρ_{pr} for all paths between border routers and $\rho_{D,\mathrm{c}}$ for all devices on these paths. Therefore, either the beacon service should calculate these variables or another service should calculate them and provide the beacon service with them. Also, as CI_D and $l_{D,\mathrm{bit}}$ vary over time, the beacon service needs to know the up-to-date value of these variables and re-calculate the *CEPB* of the path frequently and if it has changed significantly compared to the previous measurements, the beacon service encode updated value of *CEPB* in the PCB at every propagation or initiation interval. The frequency of re-calculating the *CEPB* depends on energy resources types and the traffic load at any time and each AS should adjust it based on its own energy resources types and the traffic behavior it observes. For example, if an AS uses solar power during the day, and coal power during the night, it must re-calculate its *CEPB* at least every 12 hours. Beacon services can pre-calculate these variables for each date and time, weather condition, and load matrix and use the pre-calculated values instead of calculating them frequently.[2] We propose methods for calculating variables used in Eq. (3.15) by each AS assuming that the AS administrators know

- the model and the location of all devices inside their AS;
- their electrical energy source mix at all of their presence points based on their electricity contracts;
- the sequence of routers on the intra-AS path between every two border routers (using their intra-domain routing protocol such as OSPF);
- the model and location of all optical devices between every consecutive routers on the path;
- the PUE of all their points of presence; and
- the provisioning factor for traffic protection for the path between each pair of their border routers.

The CI_D variable is calculated using Eq. (3.12b) and by having the location information of the device and the electricity contract information at the device's location. Since the mix of electricity sources can vary over time, it is necessary to re-calculate CI_D in the case of a significant change in the mix such as after sunrise or sunset. The $P_{D,\mathrm{idle}}$ variable is either provided by the vendor of the device, or can be measured by plugging-in the device and measuring its power consumption while there is no traffic load on it before installing it in the network. This is a one-time measurement. The $E_{D,\mathrm{bit}}$ can be approximated as a function of the idle power consumption of the device ($P_{D,\mathrm{idle}}$), its maximum power consumption ($P_{D,\max}$) and the maximum data rate it can process ($C_{D,\mathrm{bit}}$). The values of $P_{D,\max}$ and $C_{D,\mathrm{bit}}$ are provided by the vendor. Using these variables, the $E_{D,\mathrm{bit}}$ can be obtained by the following one-time calculation for every device:

$$E_{D,\mathrm{bit}} = \frac{P_{D,\max} - P_{D,\mathrm{idle}}}{C_{D,\mathrm{bit}}}. \tag{3.16}$$

The $E^{\mathrm{fwd}}_{R,\mathrm{pkt}}$ is the average of the per-packet energy consumption of all line cards ($E_{LC,\mathrm{pkt}}$) installed in a router. To forward a packet, each line card on a router looks up the forwarding table. As this operation is performed once per packet, the maximum number of packets a line card can process per second is equal to the maximum number of lookup operations per second it can perform, which is provided by its vendor. Also, we assume that the idle power consumption of a line card ($P_{LC,\mathrm{idle}}$), its maximum power consumption ($P_{LC,\max}$), and the maximum

---

[2] It is also possible that ASes include multiple *CEPB* values in PCBs that are calculated based on different values of the model's variables or for different time intervals in the future until the expiration of the PCB—similar to a system previously proposed for satellite networks [23, 29].

number of packets it can process per second ($C_{LC,\mathrm{pkt}}$) are provided by its vendor. Therefore, we calculate a line card's $E_{LC,\mathrm{pkt}}$ using the following equation:

$$E_{LC,\mathrm{pkt}} = \frac{P_{LC,\max} - P_{LC,\mathrm{idle}}}{C_{LC,\mathrm{pkt}}}. \tag{3.17}$$

Measuring the value of $l_{D,\mathrm{bit}}$ is a real-time measurement, and the AS administrator should monitor the traffic load on every device and provide it to the beacon service. As PCB propagation is performed periodically, we compute the average traffic load on the device between the last propagation interval and the current one.

As the increase in the $l_{D,\mathrm{bit}}$ increases the power efficiency of a device, it makes the path containing it looks greener and therefore more customers select the path which is already under traffic load. This can lead to congestion. To address this problem, we suggest that ASes use the real value $l_{D,\mathrm{bit}}$ for every device whenever it is below a certain threshold and they use this threshold as $l_{D,\mathrm{bit}}$ when the real load on the device is above the threshold.

The last variable in the Eq. (3.15) is the average size of a packet $\overline{S}_{\mathrm{pkt}}$. For that we suggest that ASes use the average of the packet size they see to the destination of a PCB that is coming from the neighbor AS they want to propagate the PCB to.

3.3 Constructing green paths

Embedding each hop's greenness information in PCBs gives all ASes that receive the PCB the opportunity to import and construct greener path segments. In this section, we propose algorithms for selecting and creating green path segments using the greenness information embedded in the PCB. Since AS beacon services can be configured to find different sets of paths with different properties, a beacon service in an AS can run the green import policy and beaconing algorithm in parallel with other beaconing algorithms such as the baseline beaconing algorithm.

3.3.1 Importing green path segments

Importing path segments in SCION consists of two different but connected procedures: (i) inserting the PCB to the beacon store for further propagation to neighbor ASes and (ii) registering the PCB's path segment into local and/or remote path services. Each of these data structures has an upper bound on the number of path segments per destination AS. Whenever a new PCB originated at a destination AS is received and there is no more space in the beacon store for storing more PCBs initiated by that destination AS, an algorithm must decide which path segments should be thrown away. A new PCB is a PCB that its path segment contains a sequence of ASes and interfaces that differs from those in the beacon store. A new path segment is the path segment of a new PCB. Reaching the upper bound of the number of path segments per destination happens in highly-connected topologies like the SCION inter-ISD network, where many ASes are directly connected to each other in multiple locations.

To select the greenest set of path segments per destination, we propose a simple policy called *green import policy* that imports every path segment initiated by a destination AS to the beacon store and the path service until the number of path segments initiated by that the destination AS reaches its upper bound. In that case, whenever a new PCB initiated by that destination AS arrives at the beacon service, the algorithm compares the *CEPB* of the new path segment with the maximum *CEPB* of all path segments stored in the beacon store. If the new path segment's *CEPB* is less than the maximum *CEPB* of all path segments in the beacon store, it removes the old PCB associated with the maximum *CEPB* from the beacon store and the path service and inserts the new path segment to the beacon store and registers it to the path service. Otherwise, it discards the new path segment. This policy selects the greenest set of path segments to reach a destination from a source AS. However, since this policy discards path segments based on their own greenness not the greenness of their combination with egress interfaces, it might cause creating sub-optimal greenness for the path segments that are propagated to neighboring ASes. This happens when the combination of a discarded path segment with an egress interface is greener than all combinations of existing path segments with all egress interfaces.

3.3.2 Constructing green path segments

At every propagation interval, the beacon service of an AS selects a set of ⟨received PCB, egress interface⟩ tuples per destination AS and per neighbor AS to construct new PCBs by appending its own *ASEntry* which includes the selected egress interface to the selected received PCB. Then it sends the constructed PCB to the neighbor AS connected to the selected egress interface. In core beaconing, it is usually impossible to select all combinations of received PCBs and egress interfaces since the number of all received PCBs, and the number of interfaces is large due to the topology's high connectivity. Therefore, a *beaconing algorithm* is needed to select ⟨received PCB, egress interface⟩ tuples such that the resulting path segments meet specific quality criteria. We propose a *green beaconing algorithm* that is tuned to construct path segments with the least possible *CEPB* between every two ASes; this algorithm is shown in Algorithm 1.

At every propagation interval, the *green beaconing algorithm* finds the set of least polluting path segments per destination AS and per neighbor AS. For every destination AS and every eligible neighbor AS (core neighbor in core beaconing, customer neighbor in intra-ISD beaconing), it repeats the following process. It first reads the set of all PCBs initiated by the destination AS as well as the interfaces they are received from (ingress interface) from the beacon store. It also retrieves all interfaces that connect the local AS to the neighbor AS and considers them as the set of egress interfaces for the neighbor. Then, it computes the *CEPB* of all prospective new PCBs that can be constructed by combining all current PCBs

Algorithm 1: green beaconing algorithm

```
Result: green combinations of PCB and egress interfaces
selectedCombinations ← [];
for d ∈ set of destination ASes of all PCBs do
    for n ∈ eligible neighbor ASes do
        CEPBToComb ← new sorted_multimap ();
        for PCB ∈ set of PCBs initiated by d do
            for eg_iface ∈ interfaces connected to n do
                comb ← ⟨PCB, eg_iface⟩;
                $CEPB_{PCB} \leftarrow \sum_{ASEntry} CEPB_{ASEntry}$;
                $CEPB_{[ing,\, eg]} \leftarrow$ Eq. (3.15);
                $CEPB_{comb} \leftarrow CEPB_{PCB} + CEPB_{[ing,\, eg]}$;
                CEPBToComb.insert($CEPB_{comb}$, comb);
        cnt ← 0;
        for (CEPB, comb) ∈ CEPBToComb do
            selectedCombinations.append(comb);
            cnt ← cnt + 1;
            if cnt >= MAX_PCBs_TO_SEND then
                break;
```

and all egress interfaces. For each ⟨current PCB, egress interface⟩ combination, the *CEPB* of the resulting PCB is equal to the *CEPB* of the current PCB plus the *CEPB* of the intra-AS path between the ingress interface and the egress interface of the local AS hop. The *CEPB* of the current PCB is equal to the sum of all *greenness metadata*s associated with intra-AS paths from ingress interfaces to egress interfaces of all *ASEntries* in the PCB. The *CEPB* of the intra-AS path of the local AS hop is computed using methods proposed in Section 3.2.2. Finally, the algorithm sorts these combinations based on their *CEPB* and selects the set of n combinations that have the least *CEPB*s.

We call the version of SCION which uses both the *green import policy* and the *green beaconing algorithm* to construct green path segments the *SCION with the green import policy and beaconing algorithm* (*SCI-GIB*) and use the abbreviation in the rest of the text.

3.3.3 Creating green paths from green path segments

Once the green core path segments are constructed between each two core ASes and green up- and down-segments are constructed between each non-core AS and its core providers in the ISD core, and also the *CEPB* information for every combination of path segments is disseminated, green paths can be constructed between every two ASes in the Internet. When end hosts ask their local path services for a green path to a destination ⟨ISD, AS⟩, the local path service asks the core path service of the core AS that can be reached via the greenest up-segment to return a green path to the destination ⟨ISD, AS⟩ pair. If the destination AS is inside the local ISD, it returns the greenest down-segment to the local path service. Otherwise, it asks the path service of the destination ISD's core AS, which can be reached via the greenest core-segment, to return a green down-segment to the destination AS. After recursively returning the path segments to the host, the host combines these green path segments and creates a green path. If the destination AS can be reached via peering links as well, the local path service compares the greenness of the paths that can be constructed using peering links with the greenest path that can be constructed by combining up, core, and down path segments. It then returns the greenest one to the end host.

3.4 Virtuous feedback

Providing end users with the green routes to every destination and giving them the authority to select their path to send their traffic, the environmentally conscientious users configure their devices to select the greener paths even if they are more expensive or have lower quality than other paths. Therefore, some amount of traffic shifts to greener paths, and ASes on polluting paths lose some portion of traffic they used to transit. As the business model of ISPs has low profit margins and high static and low variable costs, a small decrease in the amount of traffic they transit drops their profit significantly, or even makes their business unprofitable [48]. This incentivizes ASes that lose traffic because of the traffic shift to greener ASes to make their electricity resources greener or make their devices and traffic engineering mechanisms more energy-efficient so they can win back their lost traffic. This creates a virtuous cycle between different ISPs to attract traffic by making themselves more energy efficient and greener.

4 Evaluation method

In this chapter, we describe our method to evaluate the proposed *SCI-GIB*. In Section 4.1, we introduce data sets we have used in our evaluations. In Section 4.2, we explain our method to estimate the *CEPB* of both SCION and BGP paths. In Section 4.3, we explain how we estimate the global inter-domain traffic matrix, which we then use in Section 4.4 to estimate the absolute amount of reduction in CO_2 emission caused by the deployment of the *SCI-GIB*. In Section 4.5, we propose a model for the virtuous feedback cycle between ISPs to make their energy resources greener and explain our method to estimate the CO_2 reduction caused by competition to attract traffic.

4.1 Data sets

In this section we introduce data sets we used in our evaluations as well as the abbreviation for each of them that we use in the following:

- AS-Rel-Geo: *CAIDA AS relationship with geographic annotations data set* [12], which has the relationships between the largest 12 000 ASes along with the locations where they are interconnected.
- AS-Rel: *CAIDA AS relationhips and customer cones data set* [11], which contains the relationships between all neighbor ASes in the Internet as well as the customer cone of all ASes.
- ITDK: *CAIDA Internet Topology Data Kit data set* [13], which has the locations of 82 million routers as well as their degree and the ASes they belong to.
- Pfx2AS: *Routeviews prefix to AS mappings data set* [14], which maps all prefixes in the global BGP routing table to the AS that has announced them.

4.2 Estimating *CEPB*

In this section, we describe our method to estimate the *CEPB* of BGP paths and SCION core path segments constructed in two different setups:

1. The beacon store of each AS only embeds the *CEPB* of its AS hop into PCBs but its import policy, and path-selection and propagation algorithms are all CO_2-agnostic. In this configuration, only end hosts take *CEPB* information

into account to select the greenest path segments. We call this version of SCION the *SCION baseline with CEPB encoded in PCBs* (*SCI-BCE*) and use the abbreviation in the rest of the text.

2. The *SCI-GIB* in which the beacon store uses the green import policy and beaconing algorithm described in Chapter 3 to import, select, and propagate the greenest path segments.

As SCION is so far deployed only in a small number of ASes, we use simulations to analyze the effect of deploying *SCI-GIB* on the *CEPB* of the discovered paths. Also, to make a fair comparison to BGP, we simulate BGP to find inter-domain paths in the current Internet and compute their *CEPB*.

In our evaluations we only consider core path segments due to the following reasons:

1. Core ASes are highly interconnected and most of neighboring core ASes are interconnected in multiple locations. Therefore, there are lots of potential paths between every two core ASes that have different properties that can be used for different application types with different needs. Since the SCION core beaconing does not have to comply with the Gao-Rexford policy, all these diverse paths can be discovered and used. Due to the high diversity of paths among core ASes it is more likely that paths with very low CO_2 emission can be constructed. In contrast, in intra-ISD routing the path diversity is usually much smaller, so ASes and end hosts are not provided with many paths from which they can select the greener ones.

2. Core ASes carry a large amount of traffic between different ISDs. Therefore, reducing the *CEPB* of core-segments causes notable saving in the absolute CO_2 emission.

3. Since non-core ASes in the same ISD are usually located in the same region, their electricity resources are also similar. Therefore, using *SCI-GIB* in intra-ISD paths cannot make large difference. This is mainly true for our simulations where we do not have access to the exact electricity resource mix of ASes, so we assume that they use the same mix as their country. In reality, two ASes in the same region can have completely different electricity resources.

4.2.1 Constructing a large-scale AS topology

In our simulations we need a large-scale AS topology so that we can evaluate how the *SCI-GIB* behaves globally. We construct a topology consisting of the most-interconnected 2000 Tier-1 and Tier-2 ASes. The size of the topology is reasonable as it is expected that most ASes in a SCION ISD are non-core ASes and only ASes with many connections to other ASes participate in the ISD core.

Therefore, if we assume that all ASes in each country constitute one ISD and each ISD have 10 core ASes we would have around 2000 core ASes. We extract this topology from the AS-Rel-Geo data set. We first remove all ASes that are neither a Tier-1 nor a Tier-2 AS. Then among all remaining ASes—which are only Tier-1 and Tier-2 ASes—we incrementally remove ASes with the fewest interconnections with other ASes until 2000 ASes remains in the topology.

4.2.2 Estimating *CEPB* of intra-domain paths

To estimate the *CEPB* of every intra-domain path between ingress and egress interfaces of each AS hop, we need to adjust Eq. (3.15) so that it can be applicable to simulations. The first adjustment is to use the maximum power consumption of each device ($P_{D,\text{max}}$) instead of $E_{D,\text{bit}}$, $E^{\text{fwd}}_{R,\text{pkt}}$, and $P_{D,\text{idle}}$ as we do not have accurate measurements for these variables. The second one is to use $C_{D,\text{bit}}$ instead of $l_{D,\text{bit}}$ since we cannot measure the real-time traffic load on each device on intra-domain paths as the simulator only simulates the control plane and not the data plane. Moreover, we replace the CI_D of a device on the path with the CI_{AS} of the AS in which it is located. With all these modifications to Eq. (3.15) we obtain a simpler version in Eq. (4.1). In other words, we model the carbon emission of every AS hop on a path as the product of average carbon intensity of energy resources of the AS and the sum of per bit energy consumption of all devices on the intra-domain path between its ingress and egress interfaces:

$$CEPB_{AS,[ing,eg]} = \rho_{\text{pr}} CI_{AS} \sum_{D \in [ing,eg]} \rho_{D,\text{c}} \frac{P_{D,\text{max}}}{C_{D,\text{bit}}} \tag{4.1}$$

Finding the average carbon intensity of AS energy resources

We calculate the CI_{AS} of every AS by computing the weighted mean of the *CI*s of its energy resources at its points of presence (PoPs). We infer an AS's PoPs from the ITDK data set. We calculate each PoP's *CI* using Eq. (3.12a) assuming that the electrical resource mix of the PoP is the same as the electrical resource mix of the country where it is located in. We consider the degree of the router located at each PoP as the PoP's weight as the greater the router's degree is, the more energy the router and all other network devices connected to it consume.

We take countries energy resource mixes from the International Energy Agency (IEA) [28] and the *CI* of each electricity resource from the report written by Edenhofer at the International Panel on Climate Change (IPCC) [19]. Table 4.1 shows the 50^{th} percentile of carbon intensity of different electrical energy resources.

Finding the sum of per bit energy consumption of devices on the path

The summation term in Eq. (4.1) depends on the number of IP-layer hops (routers) between ingress and egress interfaces (including border routers themselves), the

Table 4.1: 50^{th} percentile CO_2 emission of different energy resources. From Edenhofer's report at the International Panel on Climate Change (IPCC) [19].

Energy resource	CO_2 intensity (g/kWh)
Coal	1001
Natural gas	469
Biomass	230
Solar	46
Geothermal	45
Nuclear	16
Wind	12
Hydroelectric	4

distance between routers, and the energy efficiency of devices on the path in all layers (IP and WDM).

We compute the intra-domain IP-layer path (router path) between every two inter-domain interfaces of each AS by applying the Dijkstra algorithm on the router topology of the AS given by the ITDK data set. Since we use the AS-Rel-Geo data set for finding the locations of inter-AS connections, in some cases there is no router in the ITDK data set at the same location and in the same AS as the geolocation data set suggests for an interconnection between two ASes. In such cases, we compute the distance between the inter-domain link location of an AS in geolocation data set and all routers of the same AS in the ITDK data set and pick the nearest one and consider it as the border router suggested by the geolocation data set. We also assume that the location in the geolocation data set is correct.

The number of WDM-layer hops depends on the number of IP-layer hops (routers) and the distance between every two WDM switches. Based on the model that Heddeghem et al. propose [24], WDM switches, transpoders, muxponders, and terminals are present at the same location as the IP routers are. WDM line amplifiers and regenerators are needed every 80 km and 1500 km between two consecutive WDM switches. We compute the distance between two consecutive routers on an intra-domain path using their locations in the ITDK topology and the great-circle distance formula, assuming the link between them is on the great circle connecting the two routers' locations. Table 4.2 shows the energy efficiencies of typical devices in IP-layer and WDM-layer devices which are proposed by Heddeghem et al. [24] and we use in our simulations. Based on the same study, the typical values of $\rho_{D,c}$ and ρ_{pr} are both 2.

Table 4.2: Energy efficiencies of typical devices in IP and WDM layers [24].

Device type	Energy consumption (W/Gbps = J/Gb)
Core router	10
WDM switch (OXC)	0.05
Trans/Mux -ponder	1.5
Amplifier	0.03
Regenerator	3

4.2.3 Estimating *CEPB* of SCION paths

To simulate SCION beaconing process on large-scale topologies we have developed a discrete-time simulator[1] in the ns-3 network-simulation framework [1]. This simulator provides the basic beaconing functionality of SCION, and enables researchers to implement their own import policy and beaconing algorithm. In the simulator, the beacon service of each AS calculates the $CEPB_{AS,[ing,eg]}$ of its AS hop using the simplified model we proposed in Section 4.2.2 and encode this information into the PCBs. It finds the *CEPB* of a path by summing the $CEPB_{AS,[ing,eg]}$ of all AS hops on the path using Eq. (3.14).

We simulate the two beaconing algorithms mentioned in the beginning of Section 4.2.2 in different configurations by changing the maximum number of PCBs to store per destination AS in the beacon store and the maximum number of PCBs per destination to propagate to every neighbor. So, we can analyze the effect of the *SCI-GIB* on the greenness of the final path segments.

4.2.4 Estimating *CEPB* of BGP paths

To calculate the *CEPB* of the current Internet paths, we simulate BGP using the SimBGP simulator [49] on the same topology we used in the SCION simulations. We assume that all border routers of the same AS have pairwise iBGP sessions. Each BGP router adheres to the Gao-Rexford policy [21] for importing and exporting routes from/to its iBGP and eBGP peers. In the case that multiple paths have the same priority using the Gao-Rexford policy, the router selects the route with the shortest AS path. If a router receives multiple routes containing the preferred AS path from its BGP peers, it selects the route whose next hop is the geographically nearest next hop.

Using this policy all border routers of an AS select the same AS path for all destination prefixes located in the same destination AS. However, as source ASes can forward their locally-sourced packets to all their border routers that are connected to the first AS hop along the preferred AS path towards the destination

---

[1] The code base of the simulator is available in a private repository on GitLab [46]. Access will be provided upon request.

prefix, the number of paths from the source AS to the destination AS would be equal to the number of such border routers. We calculate the *CEPB* of such paths by accumulating $CEPB_{AS,[ing,eg]}$ of all AS hops on each path.

4.3 Traffic matrix synthesis

To estimate the absolute amount of CO_2 emission by SCION and BGP paths and also modeling the virtuous feedback cycle between ISPs, we require an estimate of the amount of traffic between different ASes in our topology. To the best of our knowledge there is no global inter-domain traffic matrix data set. Hence, we need to synthesize the traffic matrix based on a model for inter-domain traffic matrices.

Inter-domain traffic matrix model

We use a modified version of the model that Mikians et al. propose for global inter-domain traffic matrices [33]. The amount of traffic from AS_i to AS_j is the sum of the amount of traffic that AS_i's users send to AS_j to request for contents in AS_j and the amount of traffic that servers in AS_i send to users in AS_j in response to their request. Therefore, the amount of traffic between each AS pair depends on the number of their users and the relative popularity of each AS for the other's users. Also, the traffic in the two directions of a connection is not symmetric; instead, in their model the amount of traffic in one direction is proportional to the amount of traffic in the opposite direction. Therefore, the amount of traffic between AS_i and AS_j is estimated using the following equation:

$$T_{i,j} = \sum_{A \in applications} m_A \Big(S_i p_i^A(j) + d_A S_j p_j^A(i) \Big) \tag{4.2}$$

Here, S_i and S_j denote the number of users of ASes i and j; $p_i^A(j)$ denotes the relative popularity of AS_j for users of AS_i with respect to application A; the (a)symmetry of traffic in the two directions for every application is denoted by d_A; and m_A denotes the contribution of application A to the whole amount of traffic. In this work we consider HTTP(S) data and media-streaming applications as well as popular video-streaming services which together amount to around 60 percent of the total traffic of the Internet [38].

HTTP(S) traffic matrix generation

For HTTP(S) applications the $\log_{10} d_{http}$ falls into the range of $(0.4, 1.5)$ according to Mikians et al., so we pick $\log_{10} d_{http} = 1$. For each AS we find the number of its customers (S_i) by finding the number of IPv4 addresses it announces using the Pfx2AS data set. However, for large CDNs, video-streaming services and Tier-1 carriers we assume they do not have any users.

For each AS_i we find the popularity of HTTP(S) content in other ASes subjective to its users ($p_i^{http}(j)$) similar to the method Mikians et al. proposed. We first generate a vector of n random numbers from a Zipf distribution with a slope of 1.2 where n is the number of other ASes in the network. Then we sort the vector's elements in the descending order and normalize them. We assign the largest popularity numbers in the vector to the most globally-popular ASes in the same order. Then we assign the remaining popularity numbers to other ASes randomly. To find the most popular ASes, we first find the list of top 1 million domain names provided by Pochat et al. [36]. Then we find the IP addresses of these domains using DNS queries, and their ASes using the Pfx2AS data set.

To sort ASes based on their popularity, we first sort popular domains from the least popular ones to the most popular ones and assign them a popularity index equal to their ranking in this sorted list. The most popular domain gets the highest index, and the least popular one receives the lowest index. The popularity index of an AS would be the sum of the popularity indices of the domains it hosts. We sort ASes based on their popularity indices from the most popular one to the least popular one.

The HTTP(S) traffic matrix that we obtain from the this procedure does not contain absolute traffic amounts between ASes but the relative ones. To find the absolute traffic between two ASes, we first find the relative contribution of each matrix element to the whole HTTP(S) traffic by dividing it by the sum of all elements in the matrix. Then we multiply each element's relative contribution by the real global amount of HTTP(S) traffic per month, which amounts to 28 percent of total global traffic according to the Sandvine Global Internet Phenomena Report 2018 [38]. Since the total worldwide traffic is predicted to reach 293 Exabytes per month in 2022 [44], the HTTP(S) traffic per month would reach 82 Exabytes.

Popular video streaming services traffic

For popular video streaming services we consider only the top three video streaming ASes, namely Netflix, YouTube, and Amazon Prime Video, which are responsible for 15, 11.4, and 3.7 percent of total Internet traffic, respectively, according to the Sandvine Global Internet Phenomena Report 2018 [38]. So the absolute their traffic would be 43.95, 33.4 ,and 10.8 Exabytes per month, respectively. Based on the model we introduced above, to achieve the amount of video traffic each of these ASes sends to any other AS, we distribute their total amount of traffic they send between all other ASes proportional to their number of users. The result would be three arrays specifying the amount of traffic per month from these three ASes to all other ASes. For Netflix and YouTube, we substitute their HTTP(S) traffic arrays in the HTTP(S) traffic matrix with their video traffic arrays, but for Amazon we add its HTTP(S) and video traffic arrays since we could not find a specific AS for Amazon Prime Video and Amazon has lots of HTTP(S) traffic along with their streaming services as well.

Constructing core traffic matrix

The resulting traffic matrix is a full traffic matrix between all ASes in the Internet. However, our topology is a subset of this topology containing only 2000 core ASes. Therefore, we need to compute the traffic matrix of core ASes from the full traffic matrix we have computed. We make the following assumptions and simplifications to obtain a traffic matrix between core ASes:

- The amount of traffic from one core AS to another core AS would be the sum of traffic from all ASes in the customer cone of the first core AS to all ASes in the customer cone of the second core AS
- If an AS is in the customer cone of multiple ASes, its traffic is distributed evenly between its providers.
- We discard the traffic from a core AS to itself.
- No traffic between the ASes with different core providers is sent over peering links between them.

We use the AS-Rel data set for finding the customer cone of each core AS.[2]

4.4 Estimating CO_2 emission for SCION and BGP

Once we have the *CEPB* of SCION and BGP paths and the traffic between every two ASes in our topology, we compute the amount of CO_2 emission for each path by multiplying the *CEPB* of the path with the amount of traffic on that path. Since multiple paths may exist between each two ASes, we assume that the traffic between them is evenly distributed over all the paths between them.

4.5 Modeling the virtuous feedback

In this section we propose a model for the competition between ASes to attract more traffic by making their energy resources greener. We first assume that we have a full deployment of SCION. The initial state would be the baseline SCION in which both beacon services and end domains select the n shortest paths to any other ASes and do not take the *CEPB* of paths into account for path selection. We simulate SCION core beaconing with this configuration in our simulator and find the n shortest paths from every core AS to any other core ASes. Then we compute the amount of core traffic that each core AS transits by summing the amount of traffic that is being sent over all selected core path segments that cross that

---

[2] The code base of the traffic generator as well as all pre- and post-simulation computations are available in a private repository on GitLab [45]. Access will be provided upon request.

particular AS. We assume that the amount of traffic between every source and destination core AS is distributed evenly over all selected paths between them.

Then we simulate SCION core beaconing with beacon services that run both baseline beaconing algorithm and the *SCI-GIB* the same time. Therefore, all core ASes receive at least n shortest paths and n greenest paths to any other core AS. Then we compute the amount of core traffic each core AS transits. We assume that a portion of customers—which we call *green customers* (gc)—would prefer the greenest paths to the shortest paths. Thus, the ASes on the paths discovered by baseline SCION would receive $1 - gc$ of the traffic they used to receive in baseline SCION and ASes on the green paths would receive an additional share gc of traffic between two core ASes. After calculating the amount of transit traffic of each core AS, we compare their transit traffic with baseline SCION and assume that if their transit traffic falls below a *transit loss threshold* (lt) they make a portion of their total electricity demand green-sourced. We call this portion the *willingness to become greener* (wg).

If after the first round of simulating *SCI-GIB* some ASes decide to make their electricity resources greener, we increase their green energy resources by wg of their total energy resources and simulate the beaconing process again. We repeat the same procedure until either there is no more AS that wants to make its energy resources greener or we have reached the upper bound for the number of repetitions.

To compute the CO_2 reduction in the final state of the simulation compared to the initial state, we compute the weighted mean of CI_R for all routers of core ASes in the topology using *ITDK* data set in both the initial and the final states. We consider the number of links connected to each router as its weight. In the initial state, we consider the energy mix of each country where the router is located as its CI_R and for the final state we consider the final value of CI_{AS} of its AS as its CI_R. The amount of CO_2 reduction is the relative reduction in the average CI before and after the virtuous feedback simulation.

5 Results

In this section, we present and analyze the results of our experiments to evaluate the concept of green routing in SCION. In Section 5.1, we present the result that shows the amount of CO_2 that can be saved per unit of traffic. In Section 5.2, we investigate the effect of green routing on the other quality metrics of paths. In Section 5.3, we estimate the absolute amount of CO_2 emission that can be reduced by deploying the proposed green routing approach. Finally, in Section 5.4, we analyze the impact of the virtuous feedback cycle between ISPs to make their network greener on the total CO_2 emission of the network.

5.1 Paths emission per unit of traffic

In this section, we report the amount of CO_2 per gigabit of transferred data that the green routing in SCION can save.

Absolute CO_2 emission per gigabit

Figure 5.1 shows the cumulative distribution of all core AS pairs (near 4 million pairs) as the function of the per-gigabit CO_2 emission of paths between them in different routing protocols and configurations. We compare such CDFs for the greenest paths and all paths discovered by BGP, the greenest path, the 5-greenest, and the n-greenest paths (n = number of BGP paths) discovered by *SCI-GIB* and by *SCI-BCE*. In both SCION versions, beacon stores keep a maximum of 120 paths per destination AS and send 10 paths per destination AS to every neighbor at every propagation interval.

The number of BGP paths is equal to the number of interfaces between the source AS and the next AS in the BGP-selected AS-path towards the destination. Figure 5.2 shows the cumulative distribution of core AS pairs as the function of the number of BGP paths between each pair. It shows that 75 percent of all pairs have only one path and less than 5 percent of AS pairs have more than 5 BGP paths. Since such interfaces are located in different places, different BGP paths can have different CO_2 emissions. Figure 5.1 suggests that for almost all AS pairs, the greenest path emission and the n-greenest paths average emission in *SCI-GIB* are both less than 0.05 g/Gb while these amounts for *SCI-BCE* and BGP are more than 0.05 g/GB for around 2 percent and 20 percent of all AS pairs, respectively. Moreover, the greenest path emission and all paths average emission in BGP are both more than 0.064 g/Gb, 0.075 g/Gb, and 0.1 g/Gb for 10, 5, and 2

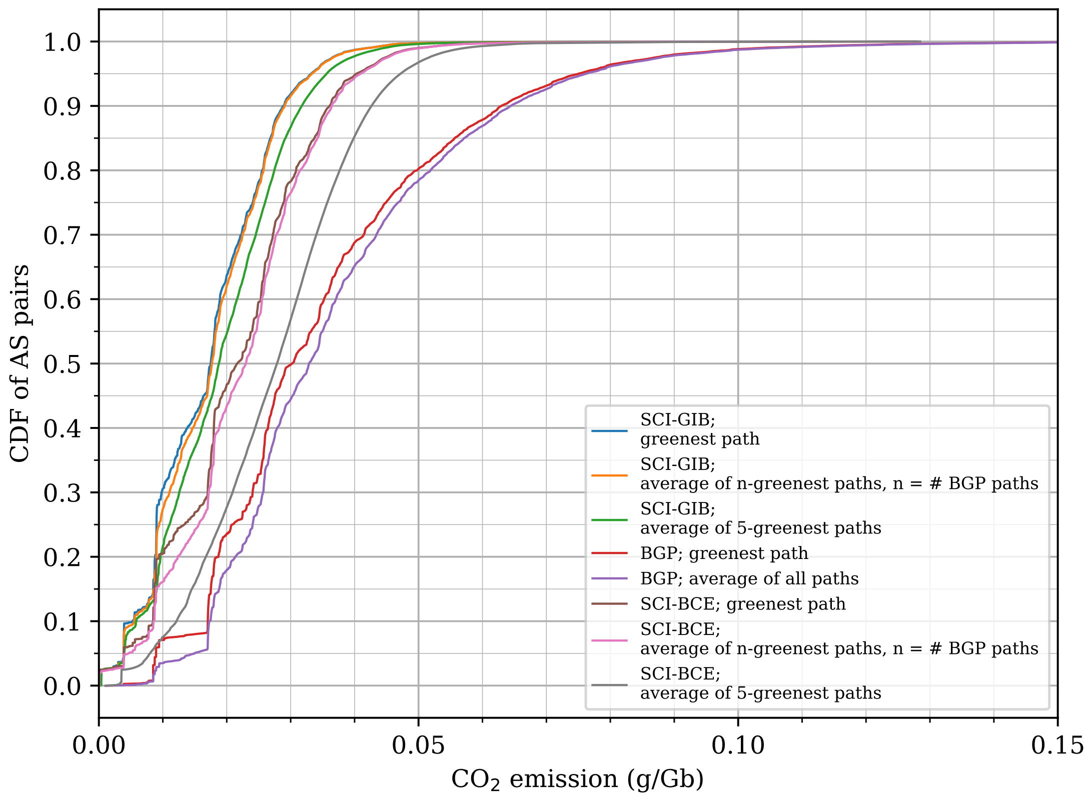

Figure 5.1: The cumulative distribution of core AS pairs as the function of per-Gb CO_2 emission of paths between them in *SCI-BCE*, *SCI-GIB*, and BGP.

percent of AS pairs, respectively. The medians of the greenest path emission and the n-greenest paths average emission in *SCI-GIB* are both 0.017 g/Gb while the medians of the greenest path emission in *SCI-BCE* and BGP are 0.022 g/Gb, and 0.03 g/Gb, respectively. The medians of the n-greenest paths average emission in *SCI-BCE* and BGP are 0.023 g/Gb, 0.032 g/Gb, respectively. The greenest path emission and the n-greenest paths average emission are almost the same in both beaconing algorithms in SCION for almost all pairs of ASes as shown by Fig. 5.1. The reason is that the number of BGP paths is less than 5 for most AS pairs. This figure also suggests that the 5-greenest paths average emission differs from the greenest path emission and the n-greenest paths average emission for both beaconing algorithms. However, for every percentile, this difference in *SCI-GIB* is about 3 times less than the one in *SCI-BCE*. This means that the 5-greenest paths average emission in *SCI-GIB* is much closer to the greenest path emission than *SCI-BCE*. Also, for every percentile, the difference between the 5-greenest paths average emissions of the two different beaconing algorithms is 2 times more than the difference between their greenest path emissions. From both of these observations, we conclude that *SCI-GIB* is better at finding the greenest set of paths. Although the greenest path emission in *SCI-BCE* is close to the greenest path emission in *SCI-GIB*, *SCI-BCE* cannot find a set of paths with low emission

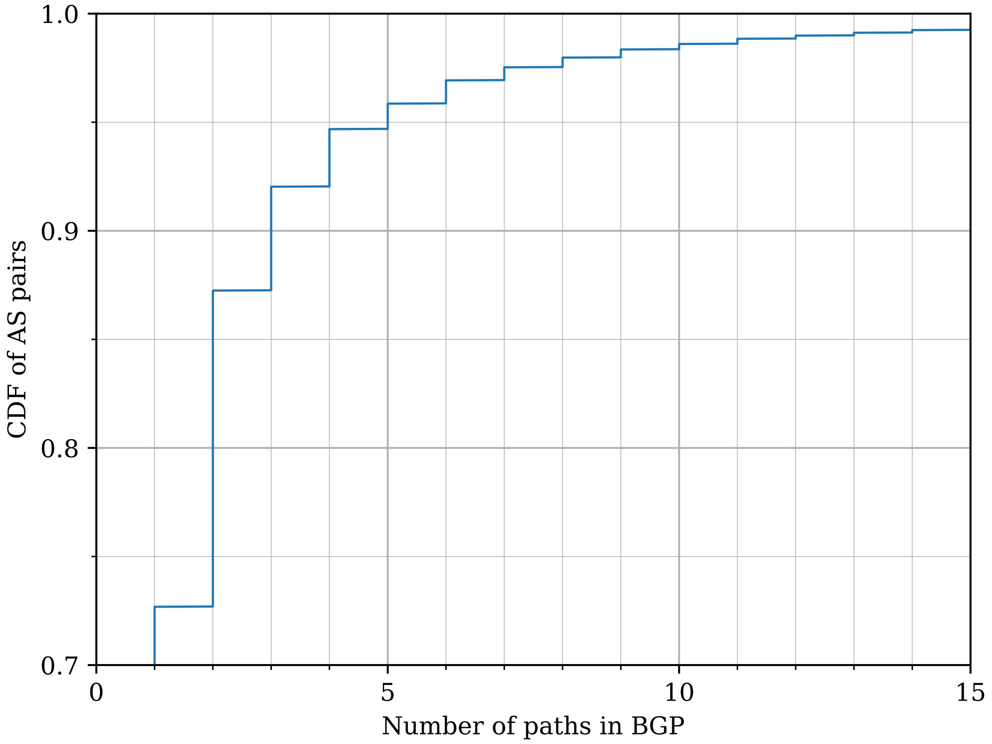

Figure 5.2: The cumulative distribution of core AS pairs as the function of the number of paths per AS path discovered by BGP. The number of BGP paths from a source AS to a destination AS is equal to the number of interfaces by which the source AS is connected to the next AS in the AS path towards the destination.

and the average emission of paths it constructs is much higher than the average emission of paths that *SCI-GIB* constructs.

Per-AS-pair differential reduction in CO_2 emissions

Figure 5.3 shows the cumulative distribution of core AS pairs as the function of the difference between the *SCI-GIB* and *SCI-BCE* greenest and n-greenest paths emissions with the greenest and all BGP paths emissions. A negative difference means that we have less CO_2 emission than BGP. As this figure suggests, the greenest path in *SCI-GIB* reduces CO_2 emission for 80 percent of AS pairs compared to the greenest BGP path, while the greenest path in *SCI-BCE* shows such reduction for 68 percent of AS pairs. Furthermore, the n-greenest paths in *SCI-GIB* reduce emission for 85 percent of AS pairs compared to all BGP paths. On the contrary, the n-greenest paths in *SCI-BCE* shows such reduction for 75 percent of AS pairs. Moreover, the greenest and n-greenest paths in *SCI-BCE* as well as the greenest path in *SCI-GIB* increase CO_2 emission for around 3 percent of AS pairs compared to BGP. However, the n-greenest paths in *SCI-GIB* increase emission only for a few AS pairs. The reason for the increase in emission compared

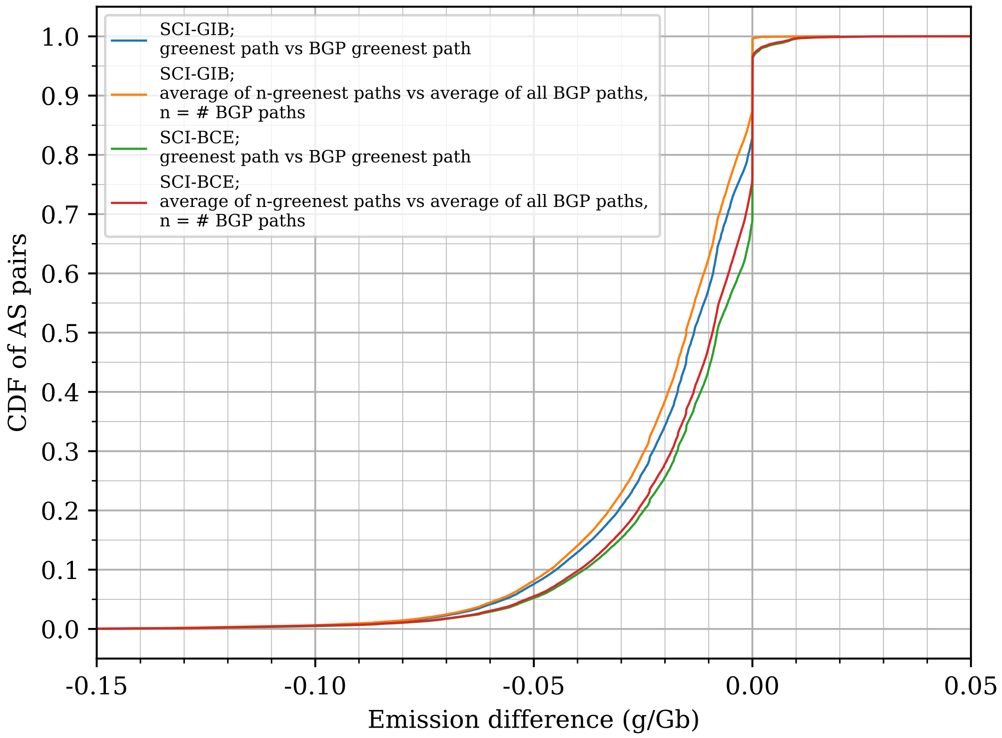

Figure 5.3: The cumulative distribution of core AS pairs as the function of the difference between the emission of SCION paths and the emission of BGP paths per Gb of traffic.

to BGP in *SCI-GIB* is that this algorithm is a greedy algorithm that does not find the greenest paths for all AS pairs since the green import policy discards some paths whose combination with some egress interfaces could make greener paths than the ones constructed by the combination of paths selected by the import policy and egress interfaces.

Based on this figure, in all configurations, the emission reduction is less than 0.05 g/Gb for more than 90 percent of all AS pairs. Also, the medians of emission reductions by the greenest paths in *SCI-GIB* and *SCI-BCE* are 0.013 g/Gb and 0.008 g/Gb, respectively. For the average of the n-greenest path, the median emission reductions are 0.015 g/Gb and 0.009 g/Gb for *SCI-GIB* and *SCI-BCE*, respectively.

Per-AS-pair relative reduction in CO_2 emissions

Figure 5.4 shows the cumulative distribution of core AS pairs as the function of the greenest and the n-greenest paths emissions in *SCI-GIB* and *SCI-BCE* relative to the greenest and all BGP paths emissions. Based on this figure, the greenest path and the n-greenest paths in *SCI-GIB* reduce the CO_2 emission by at least 45 and 48 percent for more than half of the AS pairs, respectively. In *SCI-BCE*, these reductions are at least 25 and 30 percent for more than half of the AS pairs,

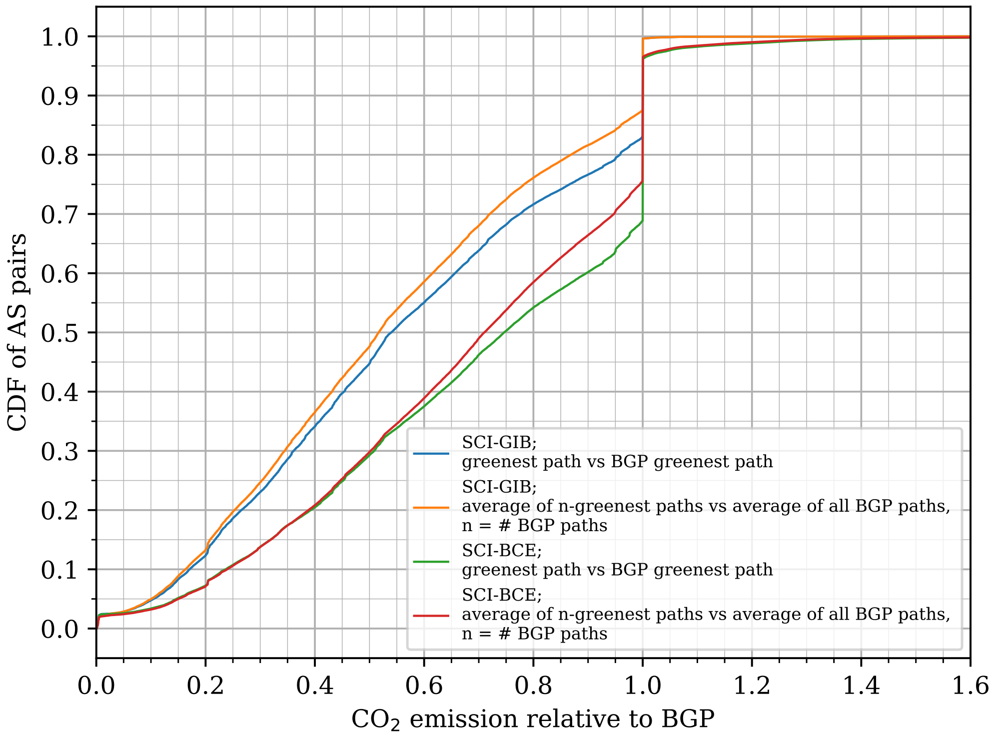

Figure 5.4: The cumulative distribution of core AS pairs as the function of the ratio of the emission of SCION paths to the emission of BGP paths that connect each pair.

respectively. Despite the fact that the differential emission reduction of *SCI-GIB* and *SCI-BCE* are close especially for the greenest path, Fig. 5.4 suggests that the relative emission reduction of *SCI-GIB* is significantly more than the relative emission reduction of *SCI-BCE*. This means that *SCI-GIB* performs significantly better than *SCI-BCE* in finding the greenest paths.

Comparison between *SCI-GIB* and *SCI-BCE*

Figure 5.5 shows the cumulative distribution of core AS pairs as the function of the greenest, the 5-greenest, and all paths emissions in *SCI-GIB* relative to the emissions of the *SCI-BCE*. As we do not consider BGP, we ignore the n-greenest paths. Also, the total number of paths is significantly more than BGP. For all AS pairs, we have at least 60 paths, and for most of them, we have 120 paths which are all included in our comparisons for all paths.

For around 60 percent of AS pairs, the greenest *SCI-GIB* path emits less than the greenest *SCI-BCE* path, and for almost all AS pairs, the greenest *SCI-GIB* path does not emit more than the greenest *SCI-BCE* path. For around 97 percent of all AS pairs, the 5-greenest *SCI-GIB* paths emit less than the 5-greenest *SCI-BCE* paths, and for more than half of AS pairs, the 5-greenest *SCI-GIB* paths reduce the 5-greenest *SCI-BCE* paths emission by at least 30 percent. For almost all

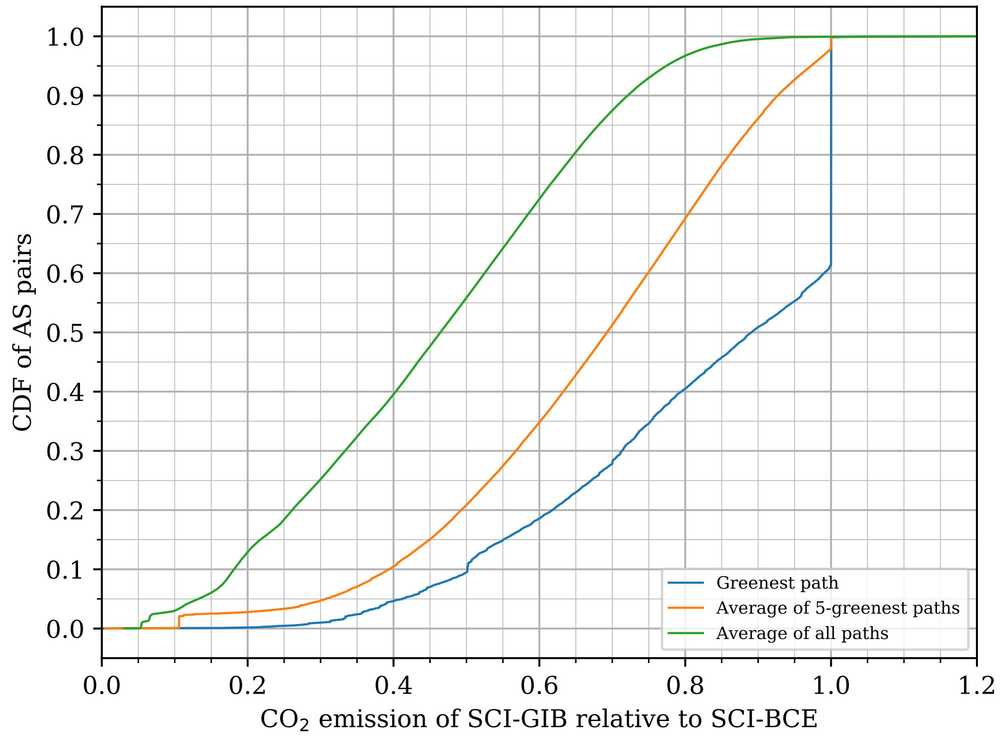

Figure 5.5: The cumulative distribution of core AS pairs as the function of *SCI-GIB* paths emissions relative to *SCI-BCE* paths emissions.

AS pairs, the set of all *SCI-GIB* paths emit less than the set of all *SCI-BCE* path. Moreover, for more than half of all AS pairs, the set of all *SCI-GIB* paths reduce the emission of the set of all *SCI-BCE* paths by at least 54 percent. All these observations suggest that although *SCI-GIB* is more efficient than *SCI-BCE* in finding the greenest possible path, it is significantly more efficient in finding the greenest set of paths.

5.2 Impact on path quality

To analyze the effect of finding the paths with the lowest amount of emission on other path-quality metrics, we compare the latency of the paths discovered by *SCI-GIB* and *SCI-BCE* with the latency of BGP paths. As we do not have the actual bandwidths of inter-domain links, we cannot compare the path bandwidths. To calculate each path latency, we add the latency between each AS hop's ingress and egress interfaces on the path. To estimate each AS hop latency, we multiply the great circle distance between the two interfaces of the AS hop by the speed of light in the fiber. This is a lower bound for the actual latency, which differs, on average, by a factor of $sim1.5$ [43].

Figure 5.6 shows the cumulative distribution of core AS pairs as the function of the difference between the *SCI-GIB* and *SCI-BCE* greenest and n-greenest paths latencies with the greenest and all BGP paths latencies. This figure suggests that

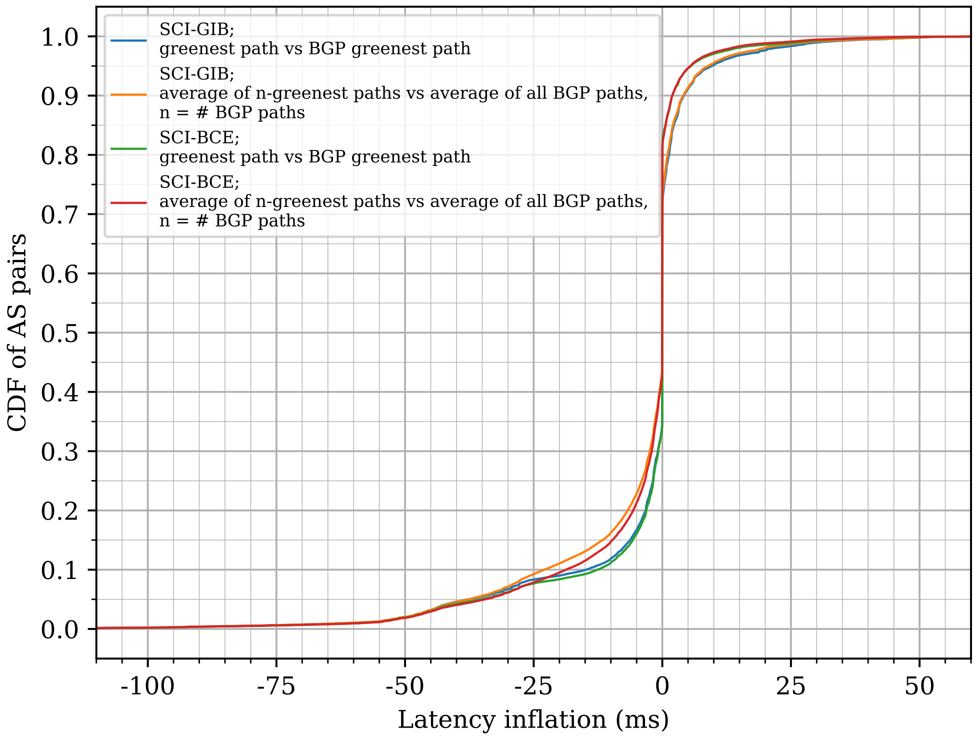

Figure 5.6: The cumulative distribution of core AS pairs as the function of latency inflation caused by selecting greener paths instead of BGP paths.

for almost 75 percent of all AS pairs, neither the greenest path nor the n-greenest paths increase the latency compared to BGP paths. Also, for only 5 percent of AS pairs, they introduce inflation of more than 10 ms. Furthermore, the *SCI-GIB* and *SCI-BCE* greenest and n-greenest paths reduce the latency for around 40 and 33 percent of all AS pairs, respectively. These observations suggest that using green paths instead of BGP paths does not affect communication quality so much and can even provide end domains with better paths in some cases. However, a latency-optimizing beaconing algorithm would find paths with lower latency.

5.3 Estimated reduction in CO_2 emission per year

By multiplying the estimated traffic between each AS pair in one year by the difference between the average per-Gb CO_2 emission of n-greenest paths in *SCI-GIB* and average per-Gb CO_2 emission of all BGP paths, we estimate the reduction in CO_2 emission that we can obtain by the deployment of *SCI-GIB*in one year. Then, for each AS we accumulate the possible reductions for the traffic it sends to any other AS. Table 5.1 shows the largest per-source-AS CO_2 reductions as well as the whole CO_2 reduction in one year by the deployment of *SCI-GIB*. It also shows the amount of core traffic that our model estimates that each of these ASes sends to other core ASes and its ratio to the total global traffic. The table suggests

Table 5.1: Modeled core traffic that popular core ASes send to other core ASes and reduction in the amount of CO_2 emission for delivering traffic from them to other core ASes in one year by using green inter-core-AS paths discovered by *SCI-GIB* instead of BGP paths. The sum for all source ASes is also written in the last row.

AS (ASN)	Core traffic (EB/year)	Contribution of core traffic to global traffic (%)	CO_2 reduction (ton/year)
Netflix (2906)	423	12	47 640
YouTube (36040)	326	9	43 032
Amazon (16509)	230	6	30 240
Cloudflare (13335)	147	4	28 728
Google (15169)	57	1.6	7536
Fastly (54113)	55	1.6	5520
Microsoft (8075)	29	0.8	4284
Akamai (20940)	10	0.3	1728
All core ASes	2004	56	210 564

that the ASes hosting the most popular services, can reduce their carbon footprint more than other ASes as they send the huge amounts of traffic to other ASes.

Although Fig. 5.4 shows significant reduction in per-Gb CO_2 emission for traffic between each AS pair, we do not see significant amount of global CO_2 reduction in Table 5.1. There are two main reasons. First, there is little traffic between many AS pairs, so the multiplication of per-Gb CO_2 reduction for those AS pairs by the amount of traffic between them is small. Furthermore, even for AS pairs with lots of traffic, sending traffic over greener paths does not make more polluting paths emit less CO_2 because the network devices on those paths are still running, and the idle power consumption of network devices contributes to the majority of power consumption of network devices.

5.4 The effect of the virtuous feedback cycle

In this section, we analyze the impact of the green routing on the competition between ASes to attract the traffic of customers who decide to use green paths instead of the shortest paths and how this competition can make the whole network emit less CO_2.

5.4.1 Virtuous feedback on a large topology

We simulate the virtuous feedback cycle on the same topology as we used in all previous experiments which consists of 2000 core ASes. We set parameters *gc*,

Table 5.2: The relative CO_2 emission reduction after each round of the virtuous feedback cycle compared to the initial state for $gc = 0.2$, $lt = 0.9$, and $wg = 0.25$ on the topology of 2000 core ASes. After 6 rounds, the increase in the relative reduction is less than 0.001.

Round	Relative CO_2 emission reduction compared to initial state
0	0
1	0.05
2	0.3
3	0.59
4	0.869
5	0.870
6	0.871

lt, and *wg* to 0.2, 0.9, and 0.25, respectively and simulate the virtuous feedback cycle for a maximum of 10 rounds. Table 5.2 shows the relative amount of CO_2 emission reduction in this topology after each round of simulations compared to the initial state (where all ASes use the same energy resources as they use in BGP) until the difference between the relative reduction in two consecutive rounds are less than 0.001. Although our simulation suggests that the virtuous feedback cycle does not converge until the 10^{th} round, the change in the relative CO_2 emission reduction is negligible after 6 rounds. After 6 rounds of simulation, the CO_2 emission of the whole network has decreased by 87 percent compared to the initial state.

5.4.2 Sensitivity analysis on a small topology

We perform a sensitivity analysis on the virtuous feedback cycle to determine how the three parameters *gc*, *lt*, and *wg* we introduced in Section 4.5 affect the competition between ASes to become greener. Therefore, we run the virtuous feedback simulations for all combinations of $gc \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}$, $lt \in \{0.95, 0.9, 0.85, 0.8\}$, and $wg \in \{0.25, 0.5, 0.75, 1.0\}$. Also, we set an upper bound of 10 for the number of beaconing simulation repetitions. As the simulations on a topology of 2000 ASes take a considerable amount of time, we shrink topology to 500 by picking the most highly-connected 500 ASes from the larger topology.

Figure 5.7 shows the number of simulation rounds it takes for core ASes to converge to a state in which they do not change their energy resources any more. In configurations that the simulation converges in zero rounds, no AS changes its energy sources since its transit traffic does not fall below the threshold specified. In all other configurations, the simulation converges to the same state where all core carrier ASes (ASes that carry traffic between other core ASes) make their

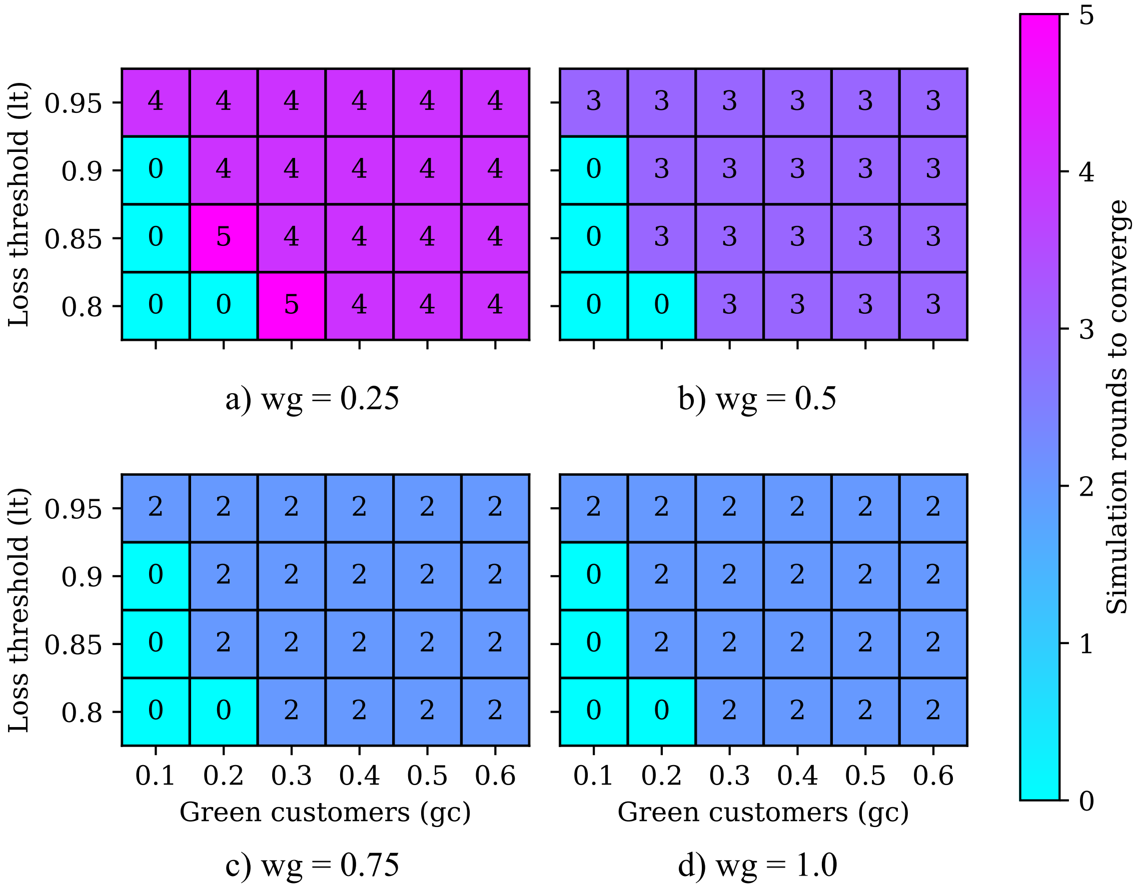

Figure 5.7: The number of rounds it takes for virtuous feedback simulations to converge in different configuration for the topology consisting of 500 core ASes.

energy sources completely green. This means that in such configurations, when some ASes make their energy source green to win back their lost traffic, more polluting ASes' traffic falls below the specified threshold, and they need to make their energy source green even if they have done it previously. This continues until all carrier ASes become 100 percent green and no traffic migration happens afterward.

In the final state, the CO_2 emission of this topology of 500 core ASes is reduced by 53 percent compared to the initial state. The reason for achieving only 53 percent of reduction and not 100 percent is that not all core ASes carry traffic between other core ASes. Therefore, not all of them change their energy source types. However, since in the actual network, all these core ASes also carry traffic from their customer cones to other core ASes and vice versa, all of them are affected by the virtuous feedback since customer ASes can select through which core provider to send their traffic based on their greenness. Also, the topology we used here for the sensitivity analysis is a subset of the larger core topology. As we showed in Section 5.4.1, we see a reduction of 87 percent in the CO_2 emission in the larger core topology. In that topology, more ASes carry traffic between other ASes which results in more ASes participating in the competition and more CO_2 reduction. Therefore, we expect that the larger the topology is, the more CO_2 reduction we get through the virtuous feedback cycle.

Based on our observation, we can conclude that although the emission reduction by the green routing without considering the effect of the virtuous feedback between ASes is not significant, as we discussed in Section 5.3, it provides customers with paths that decreases the per-Gb emission significantly. When customers use these paths to reduce their contribution to CO_2 emission, they shift their traffic from polluting ASes to greener ASes causing polluting ASes to lose traffic. The competition between ASes to attract more traffic by making their energy sources greener make the whole network emit at least 53 percent less than before deploying the green routing approach, as our simulations suggest.

6 Related work

As the current work has two different but related contributions, we investigate related work separately for each of these contributions. Therefore, we classify the related work into two categories: work that proposes models for the Internet's energy intensity, especially the core network, which will be discussed in Section 6.1, and (ii) work that suggests methods for routing packets in energy-efficient or carbon-efficient ways, which will be discussed in Section 6.2.

6.1 Modeling the Internet's energy consumption

Numerous studies in the literature propose different methods for analyzing or modeling the energy intensity of the Internet. These works fall into three categories: top-down approaches, bottom-up approaches, and model-based approaches [16].

6.1.1 Top-down approaches

In top-down approaches, researchers divide the total electricity consumption of the Internet or a part of it by the Internet traffic of that part within a period of time. The result would be the average energy consumption of the Internet per transferred data. These methods usually lead to overestimating the energy intensity of the Internet as they do not take the idle energy consumption of network devices which is responsible for the majority of the energy consumption of network devices. Studies conducted by Koomey et al. [30], Taylor et al. [47], Weber et al. [51], Lanzisera et al. [31], and Andrea et al. [5] fall into this category.

6.1.2 Bottom-up approaches

In bottom-up analyses, researchers generalize the energy intensity values they have obtained in some case studies through direct measurement or observation. For example, Coroama et al. [17] presented a pure bottom-up assessment of the energy intensity of a video conferencing transmission between Switzerland and Japan. In this study, they knew the exact path between end domains and the type of network devices on the path, so this work's estimation is accurate. However, the generalization to the whole Internet may lead to considerable error.

6.1.3 Model-based approaches

In model-based analyses, researchers model parts of the Internet based on network design principles and then apply vendor-provided energy consumption of each device in that part of the Internet to their model to find the total energy consumption of that part. Baliga et al. [9] proposed one of the earliest studies in this area. In this work, they propose a model for the energy consumption of the core routers and WDM links as the function of the number of Internet subscribers and their access rate. In their later work [8] they propose models for the energy consumption of core, metro/edge, and access networks. Hinton et al. [50] propose a model for power consumption of high-capacity switches and routers in metro and edge networks taking both idle and dynamic power into account and verified their model with direct measurements. In another work [25], they predict the future trends in the power consumption of different parts of the Internet with and without considering the improvement in devices' energy efficiencies.

However, to the best of our knowledge, none of them has tried to model the energy consumption and the CO_2 emission of a path that a single packet takes from its source to its destination. That is because there has not been an application for such a model.

6.2 Green networking protocols

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.

Zhang et al. [52] propose a heuristic method to maximize the energy-saving by turning off as many line cards as possible and rerouting traffic to other underutilized links such that all traffic constraints are met. Their method is designed only for intra-domain traffic engineering. Shi et al. [42] extend this algorithm such that it also takes into account the traffic between an AS's border routers along with intra-domain traffic. However, it cannot be considered a global inter-domain method since each domain independently deploys the method without cooperation with other domains. Ricciardi et al. [37] propose a method for routing and wavelength assignment formulation by taking every network node's carbon intensity into account. They solve this formulation using integer linear programming. Gattulli et al. [22] propose a CO_2 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 check whether the greener path emits less CO_2 or the one which is more energy-efficient and select the one with the lower emission. Schöndienst et al. [41] propose a heuristic algorithm for grooming traffic to shift the optical-electrical-optical conversions that are among the most power-hungry operations in IP backbone to nodes whose energy resources are green.

Nafarieh et al. [34] proposed extensions for both OSPF-TE and BGP to propagate information about the emission of links on each path to routers inside and outside an AS. In this approach, internal routers encode their own electricity resource information in the control plane messages they send to their neighbors. They also select the paths with the minimum path emission (MPE) toward all other routers inside their AS and advertise them to their neighboring routers. Each AS's border routers store the MPEs toward any other router in their AS, including other border routers, in a matrix. They exchange this matrix with their neighboring border routers of other ASes using a TLV extension to BGP traffic engineering. Therefore, each border router stores two MPE matrices, one for its local AS and one for the neighbor AS. During propagation of a route along the BGP path, its $\mathrm{MPE}_{\mathrm{total}}$ is calculated by accumulating all MPEs along the path. Multi-homed end domains can use the $\mathrm{MPE}_{\mathrm{total}}$ of each path to select the greenest path to every destination. This approach, however, does not affect the path selection of BGP and only propagates the MPE information of all AS hops along the BGP-selected route. Only multi-homed end domains can select to which the first-hop provider they want to forward their traffic based on the greenness of the paths these providers provide.

Since all the previous studies focus on reducing the CO_2 emission of the intra-domain paths and our SCION-based method focuses on reducing the CO_2 emission of inter-domain paths, they are orthogonal to each other and can be used at the same time. Moreover, as our model takes both energy-efficiency of networks and their electricity resources into account, it incentivizes ASes to use more energy-efficient internal routing methods to decrease the CO_2 emission of their paths and attract more traffic.

7 Conclusion

In this section, we first summarize what we have carried out in this thesis. Then, we discuss possible future work that can build on these results.

7.1 Summary

In this work, we proposed a method to reduce the CO_2 emission of the Internet's inter-domain paths using the SCION path-aware next-generation Internet architecture. First, we proposed a model for estimating the energy consumption and CO_2 emission of an inter-domain path.

In contrast to other existing models that model the whole core network energy consumption or part of it, we model the energy consumption and CO_2 emission of *every single core path* per bit of data from the source AS to the destination AS. Moreover, we model the CO_2 emission of a path as the sum of CO_2 emissions of intra-domain sections it consists of, so that each AS on the path can accurately calculate the CO_2 emission of its own section. Furthermore, we model the contribution of each bit of data to the idle power consumption of a path in our model. Using such a model, we explained how every AS on a path could measure variables used in our model and then calculate the per-bit CO_2 emission of the section of path that traverses its network.

We then show that how ASes can encode the per-bit CO_2 emission of these sections into PCBs. We also introduced the green import policy and beaconing algorithm, a greedy distributed algorithm for selecting and propagating the greenest set of paths toward every destination. Since SCION is a path-aware Internet architecture, end hosts can select these explored green paths and decrease their contribution to the CO_2 emission of the network.

We evaluated our proposed method using simulations on a subset of the real Internet inter-domain topology. As we could not accurately measure the variables used in our model, we simplified the model, and we use existing reference values for the variables in the simplified model. However, the complete model can still be used by ASes in real deployments since they can accurately measure the values of the variables used in the model. Using our simulations, we showed that for communications between more than half of AS pairs, we can decrease the CO_2 emission by at least 50 percent. Furthermore, by modeling the traffic between ASes in our topology and having the amount of reduction in CO_2 emission per gigabit of data between every two ASes, we estimated the amount of CO_2 that can be saved by each source AS every year. The sum of this amount for all ASes is about 210 000 ton/year.

We introduced the green virtuous feedback phenomenon in which ISPs compete with each other for using more renewable electricity sources to attract traffic from environmentally-conscious customers. Then we proposed a model for this phenomenon and simulated it based on the traffic matrix that we had generated before. Our results showed that on a topology consisting of 2000 ASes and for one set of the model's parameters, this phenomenon leads to a state where the total CO_2 emission of the whole network is 87 percent less than the initial state. To find the effect of change in the model's parameters, we performed a sensitivity analysis by changing the parameters of the model and repeating the simulations. This time we used a topology consisting of 500 ASes. The result of our analysis showed that for most of the configurations, the virtuous feedback cycle converges to a state in which all ASes that carry traffic between other ASes use only renewable energy resources and the whole network's CO_2 emission is 53 percent less than the initial state.

7.2 Future work

There are multiple directions for continuing the present work.

7.2.1 More accurate analysis of virtuous feedback cycle

One direction is to make the model we proposed for the virtuous feedback cycle between ASes more accurate. Currently, it abstracts from many details and models the extra price of renewable energy indirectly by the ISPs' willingness to become greener (wg), the final consumer price and their willingness to pay for green connectivity indirectly by the portion of green customers (gc), and the profit of ISPs by the loss threshold (lt). Although this model shows how using green routing in SCION leads to a greener Internet, it does not provide the exact behavior of each ISP. Therefore, a more detailed model based on the ISPs' profit model is needed to model the virtuous feedback cycle in more detail. In such a model, the costs of making electricity resources greener and the consumer costs of greener paths should be applied directly. A study can be carried out for taking these considerations into account for modeling the virtuous feedback between ISPs.

7.2.2 Green routing in a partial deployment of SCION

Another direction is to analyze the green routing in SCION while not all ASes deploy SCION. In the present work, we assumed that all ASes deploy SCION. However, only few ASes currently deploy SCION. Therefore, a study should be carried out to analyze the CO_2 emission reduction in the current partial deployment of SCION and determine the minimum number of SCION ASes and their points of presence to form a green SCION backbone which can reduce CO_2 emission considerably.

To send packets using the SCION backbone, customers send their packets over the IP network to an ingress point of the SCION backbone. The SCION backbone routes packets on paths inside the SCION network to an egress point. At an egress point, packets are sent to the destination over the IP network. To provide customers with paths with the lowest CO_2 emission, the SCION backbone should select the egress point such that the combination of the SCION path with the IP path from the egress point to the destination has the lowest CO_2 emission. However, disseminating CO_2 emission information using BGP is not as easy as SCION. Therefore, a study can be conducted to propose methods for finding the greenest combination of SCION paths and egress points in a SCION backbone network. Furthermore, since selecting green egress points affects the traffic pattern outside the SCION backbone, a study can be conducted to analyze the effect of selecting the egress point on the profit of non-SCION ASes.

Another relevant direction for future work is to analyze the competition between the SCION backbone and non-SCION ASes. Assuming a green SCION backbone in which SCION ASes mostly use green energy resources and deploy the green import policy and routing algorithm to propagate paths, we expect a virtuous feedback cycle happens between SCION and non-SCION ASes. Since SCION ASes provide customers with low-emission paths, some customers decide to send their traffic through the SCION backbone. This makes some non-SCION carriers lose their traffic even if they use green energy resources since BGP does not provide a mechanism to share their energy resource information with their customers. Therefore, these providers become incentivized to join the SCION backbone and also make their energy resources greener. This process makes the SCION network grow faster. Therefore, a study should analyze this virtuous feedback between SCION and non-SCION ASes and its effect on the SCION network growth.

7.2.3 More sophisticated beaconing algorithms

The beaconing and import policy algorithm can be the subject of numerous future projects. In this work, we tried only to optimize the CO_2 emission of paths. However, end domains may need paths that are optimized for other quality criteria. Therefore, a new import policy and beaconing algorithms should be proposed to find different sets of paths based on the needs of the neighbor ASes while taking the CO_2 emission into account. Furthermore, the overhead of sending PCBs so frequently can be overwhelming, especially in inter-ISD beaconing. On the other hand, as the energy source mix of Internet paths can change frequently, beacon services should send PCBs frequently. A study can be carried out to find a solution for this issue.

Bibliography

[1] ns-3 Network Simulator. `https://www.nsnam.org/` (last visited Apr 19, 2021), archived at `https://perma.cc/4S4N-8VGV`.

[2] Path Aware Networking RG (panrg) Internet Drafts. `https://www.potaroo.net/ietf/html/ids-wg-panrg.html` (last visited Apr 19, 2021), archived at `https://perma.cc/2RME-WVST`.

[3] A. Ahmad, A. Bianco, E. Bonetto, L. Chiaraviglio, and F. Idzikowski. Energy-Aware Design of Multilayer Core Networks. *Optical Communications and Networking, IEEE/OSA Journal of*, 5:A127–A143, Oct 2013.

[4] A. Andrae. Projecting the Chiaroscuro of the Electricity Use of Communication and Computing from 2018 to 2030. Feb 2019.

[5] A. Andrae and T. Edler. On Global Electricity Usage of Communication Technology: Trends to 2030. *Challenges*, 6(1):117–157, 2015.

[6] A. Andrae, L. Hu, L. Liu, J. Spear, and K. Rubel. Delivering Tangible Carbon Emission and Cost Reduction through the ICT Supply Chain. *International Journal of Green Technology*, 3:1–10, Nov 2017.

[7] J. Aslan, K. Mayers, J. G. Koomey, and C. France. Electricity Intensity of Internet Data Transmission: Untangling the Estimates. *Journal of Industrial Ecology*, 22(4):785–798, 2018.

[8] J. Baliga, R. Ayre, K. Hinton, W. Sorin, and R. Tucker. Energy Consumption in Optical IP Networks. *Journal of Lightwave Technology*, 27(13):2391–2403, 2009.

[9] J. Baliga, K. Hinton, and R. Tucker. Energy Consumption of the Internet. Jul 2007.

[10] C. Belady and A. Rawson. Green Grid Date Center Power Efficiency Metrics: PUE and DCIE. 2008.

[11] Center for Applied Internet Data Analysis (CAIDA). AS-Relationships. `https://www.caida.org/data/as-relationships/` (last visited Apr 19, 2021), archived at `https://perma.cc/RYN2-BSNB`.

[12] Center for Applied Internet Data Analysis (CAIDA). 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`.

[13] Center for Applied Internet Data Analysis (CAIDA). 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`.

[14] Center for Applied Internet Data Analysis (CAIDA). Routeviews Prefix to AS Mappings Data Set (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`.

[15] A. Coiro, M. Listanti, T. Squarcia, A. Valenti, and F. Matera. Energy-Minimised Virtual Topology Design in IP Over WDM Backbone Networks. *IET Optoelectronics*, 6:165–172(7), Aug 2012.

[16] V. C. Coroama and L. M. Hilty. Assessing Internet Energy Intensity: A Review of Methods and Results. *Environmental Impact Assessment Review*, 45:63–68, 2014.

[17] V. C. Coroama, L. M. Hilty, E. Heiri, and F. M. Horn. The Direct Energy Demand of Internet Data Flows. *Journal of Industrial Ecology*, 17(5):680–688, 2013.

[18] D. Costenaro and A. Duer. The Megawatts behind Your Megabytes: Going from Data-Center to Desktop. 2012.

[19] O. Edenhofer. *Renewable Energy Sources and Climate Change Mitigation: Special Report of the Intergovernmental Panel on Climate Change*. Cambridge university Press, 2012.

[20] M. Frei. 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] L. Gao and J. Rexford. Stable Internet Routing Without Global Coordination. In *IEEE/ACM Transactions on Networking (TON)*, volume 9, pages 307–317, Jun 2000.

[22] M. Gattulli, M. Tornatore, R. Fiandra, and A. Pattavina. Low-Carbon Routing Algorithms for Cloud Computing Services in IP-over-WDM Networks. In *IEEE International Conference on Communications (ICC)*, pages 2999–3003, 2012.

[23] G. Giuliari, T. Klenze, M. Legner, D. Basin, A. Perrig, and A. Singla. Internet Backbones in Space. *ACM SIGCOMM Computer Comunication Review (CCR)*, 50(1), Jan 2020.

[24] W. Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. Power Consumption Modeling in Optical Multilayer Networks. *Photonic Network Communications*, 24, Oct 2012.

[25] K. Hinton, J. Baliga, M. Feng, R. Ayre, and R. Tucker. Power Consumption and Energy Efficiency in the Internet. *Network, IEEE*, 25:6 – 12, May 2011.

[26] K. Hinton, F. Jalali, and A. Matin. Energy Consumption Modeling of Optical Networks. *Photonic Netw. Commun.*, 30(1):4–16, Aug 2015.

[27] G. Huston. 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.

[28] International Energy Agency (IEA). 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`.

[29] T. Klenze, G. Giuliari, C. Pappas, A. Perrig, and D. Basin. Networking in Heaven as on Earth. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*, Nov 2018.

[30] J. Koomey, H. Chong, W. Loh, B. Nordman, and M. Blazek. Network Electricity Use Associated with Wireless Personal Digital Assistants. *Journal of Infrastructure Systems*, 10(3):131–137, 2004.

[31] S. Lanzisera, B. Nordman, and R. E. Brown. Data Network Equipment Energy Use and Savings Potential in Buildings. *Energy Efficiency*, 5(2):149–162, 2011.

[32] J. Lorincz, A. Capone, and J. Wu. Greener, Energy-Efficient and Sustainable Networks: State-Of-The-Art and New Trends. *Sensors*, 19:4864, Nov 2019.

[33] J. Mikians, N. Laoutaris, A. Dhamdhere, and P. Barlet-Ros. ITMgen — A First-Principles Approach to Generating Synthetic Interdomain Traffic Matrices. In *2013 IEEE International Conference on Communications (ICC)*, pages 2507–2512, 2013.

[34] A. Nafarieh, Y. Fazili, and W. Robertson. Dynamic Inter-domain Negotiation for Green Algorithms in Optical Networks. *Procedia Computer Science*, 21:25 – 32, 2013. The 4th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN-2013) and the 3rd International Conference on Current and Future Trends of Information and Communication Technologies in Healthcare (ICTH).

[35] A. Perrig, L. Chuat, R. M. Reischuk, and P. Szalachowski. *SCION: A Secure Internet Architecture*. Springer International Publishing, 2017.

[36] V. Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczynski, and W. Joosen. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. Jan 2019.

[37] S. Ricciardi, D. Careglio, F. Palmieri, U. Fiore, G. Santos-Boada, and J. Sole-Pareta. Energy-Aware RWA for WDM Networks with Dual Power Sources. In *IEEE International Conference on Communications (ICC)*, pages 1–6, 2011.

[38] Sandvine. The Global Internet Phenomena Report. `https://www.sandvine.com/hubfs/downloads/phenomena/2018-phenomena-report.pdf` (last visited Apr 19, 2021), archived at `https://perma.cc/SY2E-C7N8`, Oct 2018.

[39] S. Scherrer, M. Legner, A. Perrig, and S. Schmid. Enabling Novel Interconnection Agreements with Path-Aware Networking Architectures. In *Proceedings of the 51th International Conference on Dependable Systems and Networks (DSN)*, Jun 2021.

[40] D. Schien, V. C. Coroama, L. M. Hilty, and C. Preist. The Energy Intensity of the Internet: Edge and Core Networks. In L. M. Hilty and B. Aebischer, editors, *ICT Innovations for Sustainability*, volume 310 of *Advances in Intelligent Systems and Computing*, pages 157–170. Springer, 2015.

[41] T. Schöndienst and V. Vokkarane. Renewable Energy-Aware Grooming in Optical Networks. *Photonic Network Communications*, 28:71–81, Aug 2014.

[42] J. Shi and B. Zhang. Making Inter-Domain Routing Power-Aware? In *International Conference on Computing, Networking and Communications (ICNC)*, pages 932–938, 2014.

[43] A. Singla, B. Chandrasekaran, P. B. Godfrey, and B. Maggs. The Internet at the Speed of Light. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*. Association for Computing Machinery, 2014.

[44] Statista. 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.

[45] S. Tabaeiaghdaei. green_routing. `https://gitlab.inf.ethz.ch/tabaeias/green_routing`, 2021.

[46] S. Tabaeiaghdaei and C. Gloor. ns-3_beaconing_simulator. `https://gitlab.inf.ethz.ch/tabaeias/ns-3_beaconing_simulator/-/tree/green_routing`, 2021.

[47] C. Taylor and J. Koomey. Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising. Jan 2008.

[48] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. Vazirani. How Many Tiers? Pricing in the Internet Transit Market. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM'11*, volume 41, pages 194–205, Aug 2011.

[49] L. Vanbever. dragon_simulator. `https://github.com/network-aggregation/dragon_simulator` (last visited Apr 19, 2021), archived at `https://perma.cc/4QCZ-7JWU`, Jun 2014.

[50] A. Vishwanath, K. Hinton, R. W. A. Ayre, and R. S. Tucker. Modeling Energy Consumption in High-Capacity Routers and Switches. *IEEE Journal on Selected Areas in Communications*, 32(8):1524–1532, 2014.

[51] C. L. Weber, J. G. Koomey, and H. S. Matthews. The Energy and Climate Change Implications of Different Music Delivery Methods. *Journal of Industrial Ecology*, 14(5):754–769, 2010.

[52] M. Zhang, C. Yi, B. Liu, and B. Zhang. GreenTE: Power-Aware Traffic Engineering. In *The 18th IEEE International Conference on Network Protocols*, pages 21–30, 2010.

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Green Inter-Domain Routing in the SCION Internet Architecture

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):	**First name(s):**
Tabaeiaghdaei	Seyedali

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the 'Citation etiquette' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date	**Signature(s)**
Zürich, 20.04.2021	Seyedali Tabaeiaghdaei

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.