

GRASS: Green Ranking of Autonomous Systems

Antonia Affinito, Cristian Hesselman[†], Savvas Kastanakis
 University of Twente, Netherlands; [†]SIDN Labs, Netherlands
 Emails: {a.affinito, c.e.w.hesselman, s.kastanakis}@utwente.nl

Abstract—The rapid growth of the Internet has raised concerns about the carbon emissions linked to data transmission. Autonomous Systems (ASes), which make inter-domain routing decisions, influence how traffic moves through the network and, indirectly, where emissions are generated. Although carbon-aware routing is gaining interest, limited visibility into emission sources remains a key challenge. In this paper, we present GRASS (Green Ranking of Autonomous SystemS), a framework for estimating the CO₂ intensity of ASes based on their geographic distribution and the carbon efficiency of regional electricity grids. GRASS combines geolocation, routing, and carbon data to infer probabilistic location profiles and compute a CO₂ intensity score for each AS. Using GRASS, we analyze emission patterns across AS popularity tiers, and customer cone sizes, and assess the carbon intensity of AS-to-AS links based on associated organizations. Our results show that a small number of structurally central ASes and links account for a large fraction of total Internet-related emissions. We propose a novel approach to assessing the CO₂ intensity of ASes and their interconnections, offering insights that enable targeted interventions for greener network operations and routing policies.

I. INTRODUCTION

The global expansion of Internet connectivity has brought significant economic and societal benefits. However, it has also contributed to a rising environmental footprint. In 2020, the Information and Communication Technology (ICT) sector accounted for approximately 4% of global electricity consumption and about 1.4% of total greenhouse gas (GHG) emissions [1]. Within the broader ICT sector, network infrastructure plays a key role in enabling global Internet connectivity. As networks continue to grow, the processes required for data transmission involve increasingly complex computations, which in turn lead to higher energy consumption. Routing protocols, like BGP, select paths based on various factors, including operator policies and business agreements, but not on carbon footprint. This raises the question: are the most efficient paths also the most sustainable? Recent works have started exploring carbon-aware routing [2]–[5], yet how emissions are distributed across the Internet’s topology remains unclear. In this paper, we introduce GRASS (Green Ranking of Autonomous SystemS), a framework for estimating and comparing the CO₂ intensity of ASes. The core premise of GRASS is that emissions from network operations are location-dependent: infrastructure such as routers, switches, and interconnection points consumes electricity locally, and the carbon intensity of that electricity varies significantly across countries due to differences in how power is generated. Despite this strong geographic dependency, detailed data on the physical footprint of ASes is often lacking. To the best of

our knowledge, this is the first study to address this challenge by integrating multiple complementary data sources to infer the likely geographic footprint of ASes and their interconnections (ASes links). GRASS combines facility-level information (e.g., from PeeringDB) with routing-based geolocation (e.g., from RIPEstat and BGPStream), and links this to country-level carbon intensity data from the Green Web Foundation. We apply GRASS to analyze emission patterns across AS popularity tiers, which indicate how often networks appear in routing paths, customer cone sizes, which represent the set of networks reachable through customer–provider links, and the carbon intensity of AS-to-AS connections. The result includes CO₂ intensity scores for both ASes and their interconnections, capturing how emissions vary across the Internet.

Our contributions are: (i) We introduce the first green ranking of ASes and their interconnections at Internet scale. (ii) We show how CO₂ intensity varies with AS popularity tiers, customer cone sizes, and AS-to-AS links based on associated organizations. (iii) We release an interactive website presenting per-AS and per-link CO₂ intensity scores to support future research on carbon-aware routing: <https://ut-dacs.github.io/GRASS/>. All data and code are publicly available to support reproducibility: <https://github.com/ut-dacs/GRASS>.

In addition to providing a framework that maps CO₂ intensity among ASes, we show that a small set of structurally central ASes and links concentrate a large portion of estimated network emissions, driven by their broad geographic footprint and structural importance in the global routing system. Network operators, policymakers, and researchers can use these insights to design targeted interventions, for example, by prioritizing greener peering arrangements, rerouting traffic through lower-emission paths.

II. RELATED WORK

Environmental sustainability in networking has gained increasing attention as global data traffic and energy consumption continue to grow. In response to these challenges, the IETF GREEN Working Group [6] has recently initiated efforts to develop standards for energy-efficient networking, focusing on terminology, use cases, and data models to support energy-aware operations. In addition to this, Tabaeiaghdaei et al. [3] introduced CIRo, a carbon-aware inter-domain routing system that leverages path-aware architectures such as SCION [7]. By forecasting and sharing carbon intensity data for inter-domain paths, CIRo allows endpoints to select lower-emission routes without altering core routing protocols, demonstrating

the feasibility of edge-level carbon-aware decisions. Similarly, Sutherland and Phillips proposed Low-Carb BGP [2], a carbon-aware extension to BGP that introduces green metrics into routing decisions via an SDN-based control layer. Their design avoids changes to BGP itself, instead offering a lightweight, deployable mechanism for reducing the carbon footprint of inter-domain routing while incentivizing operators to improve energy performance.

El-Zahr et al. developed CATE [4], a traffic engineering algorithm that incorporates real-time carbon intensity data and router power models to minimize emissions within backbone networks. Their work demonstrates the challenges of integrating carbon metrics into existing routing logic and highlights the trade-offs between precision, feasibility, and carbon savings.

Ruiz-Rivera et al. [8] examined the unintended consequences of energy-saving techniques, such as shutting down idle links, on inter-domain routing protocols. Their findings show that such actions can trigger hot-potato routing changes and degrade BGP stability, pointing to the need for routing-aware energy strategies that do not compromise network performance. Zilberman et al. [5] outlined a vision for carbon-intelligent networking that goes beyond energy efficiency to include full-stack, end-to-end carbon accounting. They emphasize the need for standardized carbon metrics, improved visibility into energy sources, and coordinated actions across devices, applications, and routing protocols to support meaningful decarbonization of network operations.

In contrast to protocol-level solutions, our work takes a complementary, data-driven approach by estimating AS-level carbon footprints based on geographic presence and regional electricity carbon intensity. This offers a macroscopic view of network CO₂ intensity values to support sustainability analysis and inform greener infrastructure and routing strategies.

III. METHODOLOGY

To estimate the carbon intensity of ASes based on their geographic presence and the carbon efficiency of local electricity grids, we have developed a multi-step methodology.

A. Data Collection

To capture the geographic distribution of AS infrastructure, we adopted a methodology inspired by a previous study [9] and employed two complementary datasets: PeeringDB [10] and RIPEstat [11]. We subsequently integrated the inferred geolocation profiles with country-level carbon intensity metrics from the Green Web Foundation (GWF) [12] and inter-domain routing data from BGPStream [13].

1) *PeeringDB*: We used a snapshot of the PeeringDB database from May 1, 2025, to map the geographic presence of ASes. PeeringDB [10] is a community-maintained database that records the interconnection infrastructure of ASes, including the facilities (colocation centers or PoPs) and Internet Exchange Points (IXPs) where they are present. This data serves as a useful proxy for identifying where an AS maintains physical or peering infrastructure. We consider each instance

of an AS appearing in a country, whether through a facility or an IXP, as a distinct indicator of physical infrastructure presence. We aggregate these signals to create a per-AS list of observed countries, and then normalize these lists into percentage distributions. The result is a probabilistic geographic footprint for each AS, where each country is associated with a share of the AS's observed infrastructure presence. This facility- and IXP-level view reflects physical interconnections rather than logical routing, helping identify edge and metro deployments often missed by IP-based data. We assume that every recorded presence indicates meaningful infrastructure and not merely a placeholder or administrative registration.

2) *RIPEstat*: While infrastructure-level presence shows where ASes operate, it can be incomplete, especially for smaller networks or those not thoroughly documented in PeeringDB. To address this, we used a complementary geolocation method based on IP prefix advertisements from RIPEstat [11], which integrates MaxMind's GeoLite database [14]. For each AS, we retrieved all announced IPv4 and IPv6 prefixes and mapped them to countries using MaxMind's heuristics, based on IP allocation records, geocoding, and user localization data.

This process results in a set of country-level associations for each AS, where each country is assigned a coverage percentage, representing the fraction of an AS's total advertised prefixes believed to originate in that country. These percentages serve as probabilistic indicators of presence, reflecting the AS logical footprint in address space utilization, similar to the facility-based model.

The output is a per-AS distribution across countries, normalized so the country-level values sum to 100%. This provides a network-centric view of geolocation, capturing not where an AS has infrastructure, but where its address space is advertised and likely routed from. This method complements the PeeringDB-based approach in several ways. First, it can detect presence in countries where the AS is active but not listed in facility databases. It also captures edge deployments and content delivery strategies that rely on localized IPs, even when the infrastructure is outsourced. Finally, it introduces a routing-oriented view of presence, reflecting how ASes engage with global Internet topology.

Together, the prefix-based and facility-based methods form a more comprehensive model of AS geography, enabling better carbon impact estimates by capturing both physical infrastructure and logical network behavior.

3) *Green Web Foundation*: To estimate the environmental impact of AS operations, we used country-level carbon intensity data from the GWF [12], which has received attention in recent years [15], [16]. The GWF reports average CO₂ emissions per kilowatt-hour (gCO₂/kWh), calculated using official energy statistics and national generation mixes. We used the average carbon intensity dataset, which assigns each country a single value based on its typical energy sources (e.g., coal, gas, hydro, solar). We downloaded the most recent update available as of May 2025. We matched each country in our geolocation dataset, whether sourced from PeeringDB or RIPEstat, to its corresponding CO₂ intensity value. We

excluded countries not covered by the GWF dataset to avoid introducing uncertainty. By aligning geographic presence with CO₂ intensity factors, we translated each AS's international footprint into a carbon-relevant signal.

4) *BGPStream*: To complement geolocation data and assess the real-time routing behavior, we also incorporated data from BGPStream [13] for May 1, 2025. BGPStream offers access to real-time and historical BGP routing updates, allowing for a dynamic understanding of how ASes announce their prefixes and how routes propagate across the global Internet. This data is critical for inferring the logical paths that traffic might take by measuring link frequency. It can help validate or refine the probabilistic geographic footprints derived from PeeringDB and RIPEstat, offering insights into the broader topological context relevant to carbon intensity.

B. Mapping CO₂ Intensity to ASes

After establishing country-level presence profiles for each AS and retrieving corresponding national carbon intensity data, we computed a composite carbon footprint score for each AS. This score reflects the average CO₂ intensity of the electricity grids in the countries where the AS operates.

1) *Weighted CO₂ Estimation*: We computed each AS's carbon intensity as a weighted average of national CO₂ intensities, where the weights correspond to the AS's normalized presence across countries. Formally, for a given AS:

$$\text{CO}_2(\text{AS}) = \sum_{c \in C} w_c \cdot I_c$$

where:

- C is the set of countries in which the AS is active,
- w_c is the normalized weight of the AS's presence in country c ,
- I_c is the average carbon intensity of electricity in country c , measured in grams of CO₂ per kilowatt-hour (gCO₂/kWh), as provided by the Green Web Foundation.

This assumes an AS's carbon impact is proportional to its geographic footprint and the carbon intensity of the grids it relies on.

2) *Handling Incomplete Data*: When an AS was present in countries without carbon intensity data, we excluded those countries and renormalized the remaining weights to maintain consistency. This prevented missing data from skewing results or causing underestimation.

This process yields a single CO₂ intensity value per AS, reflecting the environmental efficiency of its infrastructure locations. Lower values indicate presence in cleaner-energy countries, while higher values suggest reliance on more carbon-intensive grids.

C. Mapping CO₂ Intensity to AS Links

Following the methodology of mapping CO₂ intensity to ASes, we estimated CO₂ emissions associated with AS-level links, enabling more detailed analysis of end-to-end path footprints. We compute a weighted average of national CO₂ intensities when the two ASes are present in one or more of

the same countries, weighting each country by their combined presence. If they have no countries in common but are each present in exactly one country, we use a simple average. The resulting intensity is then scaled by how often the link appears in observed paths. Formally, for a link between AS₁ and AS₂:

$$\text{CO}_2(\text{Link}) = \left(\frac{\left(\sum_{c \in C} (w_c^{(1)} + w_c^{(2)}) \cdot I_c \right)}{\sum_{c \in C} (w_c^{(1)} + w_c^{(2)})} \right) \cdot N$$

where:

- C is the set of countries where both AS₁ and AS₂ are present,
- $w_c^{(1)}$ and $w_c^{(2)}$ are the normalized presence weights of AS₁ and AS₂ in country c ,
- I_c is the carbon intensity of electricity in country c , expressed in grams of CO₂ per kilowatt-hour (gCO₂/kWh),
- N is the number of times the link AS₁–AS₂ appears in observed potential AS paths (i.e., usage factor).

We included only links where the two ASes shared at least one country of presence, or where both ASes were active in exactly one country (even if different). This approach ensures that emissions are attributed to plausible points of interconnection. We assigned each link a usage factor N , representing how often the AS pair appeared consecutively in BGP paths, reflecting its relevance in real-world routing.

To compute the usage factor N , we parsed BGP RIB snapshots from the RIPE RIS project dated May 1, 2025.¹ Our pipeline retained only AS paths that met the following criteria: the prefix was not listed in the fullbogons files, AS path prepending was removed, and paths with loops were discarded. We then matched prefixes against the RouteViews Prefix-to-AS mappings dataset², which provides daily IPv4 and IPv6 associations between IP prefixes and origin ASes. This ensured that only paths with valid, known origin ASes were considered.

IV. ANALYSIS

In this section, we examine how carbon intensity varies across ASes and their links, estimating CO₂ emissions based on geographic presence. We relate these values to popularity, customer cone size, and interconnections to reveal structural patterns in Internet sustainability.

A. CO₂ Intensity by AS Popularity

To investigate whether the popularity of an AS correlates with its carbon footprint, we combined our CO₂ intensity estimates with AS popularity data. To measure AS popularity, we followed the approach of Naab et al. [17], who aggregate domain-based top lists (e.g., Tranco, Umbrella) to network prefixes and ASes. We resolve each domain to one or more IP addresses, which we then map to prefixes and AS numbers.

¹<https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/>

²<https://www.caida.org/catalog/datasets/routeviews-prefix2as/>




Fig. 1. Distribution of CO₂ intensity across AS popularity groups. While the median remains stable, broader groups include more low-intensity ASes, increasing the spread.




Fig. 2. Log-log plot of total CO₂ intensity values aggregated over each AS's customer cone versus cone size. Each point represents an individual AS from the CAIDA May 2025 dataset.

We apply a Zipf distribution to assign popularity weights based on domain ranks, and aggregate these weights at the AS level. This approach gives us a weighted list of ASes that reflects their prominence based on the popularity of hosted domains. By correlating these metrics, we investigate to what extent more prominent ASes are associated with higher environmental impact. Towards that end, we group ASes by popularity into four tiers: *Top 100*, *Top 1,000*, *Top 10,000*, and the full set of ASes. For each group, we plot the distribution of CO₂ intensity using violin plots, as shown in Figure 1. The results show that median CO₂ values are consistent across popularity groups, suggesting that highly popular ASes do not necessarily emit more carbon than less popular ones. As group size increases, the distribution becomes broader and more variable, with a higher concentration of low-emitting ASes.

B. CO₂ Intensity by Customer Cone Size

To examine how carbon intensity scales per the structural reach of an AS, we analyze the relationship between total CO₂ output and customer cone size. The customer cone of an AS, as defined in prior work [18], consists of all ASes that can be reached through provider-to-customer links, recursively. This metric serves as a proxy for the AS's influence and its potential role in traffic transit. Using CAIDA's AS relationship dataset (May 2025)³, we compute per-AS customer cones and

aggregate the estimated CO₂ intensity of all ASes within each cone, including the root. The result is a metric of *total cone CO₂*, a structural measure of an AS's indirect carbon footprint. Figure 2 shows the relationship between customer cone size and total cone CO₂ on a log-log scale. The distribution reveals a clear superlinear trend, consistent with a power-law relationship. A linear regression in log-log space (red line) yields a slope greater than one, indicating that CO₂ increases faster than cone size. This indicates that as the customer cone grows, total intensity increases at a faster-than-linear rate.

This behavior likely reflects topological centralization: large ASes serve many customers and accumulate higher indirect emissions due to their central role in the transit hierarchy. Moreover, variance among small and medium cone sizes suggests heterogeneity in AS operational profiles, possibly driven by differences in energy mix, infrastructure density, or functional role. These results indicate that a small set of structurally central ASes account for a large share of the Internet's carbon footprint. As such, targeted mitigation strategies, such as prioritizing greener energy sources, integrating carbon intensity into route selection policies, or applying carbon-aware traffic engineering, could therefore deliver significant global impact.

C. CO₂ Intensity per AS Link

To understand how carbon intensity varies across the Internet infrastructure, we analyze CO₂ intensity at the level of individual AS-to-AS links, as detailed in Section III. This metric estimates the carbon emissions associated with data exchanged between two ASes and is normalized by dividing each link's CO₂ intensity by the total across all links, enabling direct comparison. Figure 3 shows the cumulative share of CO₂ emissions as a function of the fraction of AS-to-AS links, ranked from most to least impactful. The distribution of CO₂ emissions is highly concentrated: the top 10% of AS-to-AS links are responsible for about 93.65% of total emissions, as shown in the CDF. This highlights a small but critical subset of links that dominate Internet-related emissions. To further investigate this concentration, Figure 4 shows the CO₂ intensity between the 10 organizations whose AS-level




Fig. 3. Cumulative distribution of CO₂ emissions across AS-to-AS links, ordered from most to least impactful. The distribution is highly skewed: the top 10% of links account for approximately 93.65% of total emissions.

³<https://publicdata.caida.org/datasets/as-relationships/>




Fig. 4. CO₂ Intensity of Traffic Between Top 10 Network Organizations.

interconnections generate the highest total emissions. Each cell aggregates the emissions of all interconnections between ASes affiliated with the corresponding source and destination organizations. Darker shades indicate higher intensity, while a ‘x’ marks the absence of observed interconnection. The heatmap reveals that most organization pairs maintain relatively low-impact links, while a few exhibit high CO₂ intensity. Notably, interconnections involving Saudi Telecom Company JSC, NTT America, and Level 3 Parent appear among the most carbon-intensive. These links may reflect transcontinental routes, high-throughput exchanges, or routing through carbon-heavy infrastructure. This aligns with the CDF analysis because emissions are not evenly distributed but concentrated in a small number of organizational relationships. This concentration suggests that targeted interventions on key interconnections could yield substantial reductions in the Internet’s overall carbon footprint.

V. CONCLUSION

We introduced GRASS (Green Ranking of Autonomous SystemS), a novel framework that estimates the carbon intensity of ASes based on their geographic distribution and the carbon efficiency of regional electricity grids. GRASS integrates facility records and IP geolocation, and maps these locations to national carbon intensity data to assign carbon scores to ASes. This allows us to assign carbon scores that reflect the carbon footprint of both AS operations and their interconnections. Our results highlight the unequal distribution of estimated network emissions, with a small number of ASes and interconnections responsible for a significant share of the Internet’s carbon footprint. This concentration suggests that targeted interventions, such as adopting greener practices among major ASes, could have a considerable impact. The effect could be further amplified through regulatory support and

economic incentives that encourage environmentally conscious routing decisions. Moreover, it is essential to avoid unintended routing consequences, such as congestion or instability from overusing low-emission paths, that could undermine sustainability goals. We also provide a public website with estimated CO₂ emissions of ASes and their potential interconnections, making all values openly accessible.

Future work may include comparing our BGP-based estimates with traffic data from IXPs, modeling cross-country AS links, and using more fine-grained carbon intensity data.

ACKNOWLEDGEMENTS

This research received funding from the Dutch Research Council (NWO) under the projects UPIN and CATRIN.

REFERENCES

- [1] J. Malmodin, N. Lövhagen, P. Bergmark, and D. Lundén, “Ict sector electricity consumption and greenhouse gas emissions – 2020 outcome,” [Online; accessed 2025-06-26]. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4424264
- [2] E. Sutherland and I. Phillips, “Low-carb bgp: A carbon-aware inter-domain routing extension to bgp,” in *Proceedings of the 2024 Applied Networking Research Workshop*, 2024, pp. 53–59.
- [3] S. Tabaeiaghdaei, S. Scherrer, J. Kwon, and A. Perrig, “Carbon-aware global routing in path-aware networks,” in *Proceedings of the 14th ACM International Conference on Future Energy Systems*, 2023, pp. 144–158.
- [4] S. El-Zahr, P. Gunning, and N. Zilberman, “Exploring the benefits of carbon-aware routing,” *Proceedings of the ACM on Networking*, vol. 1, no. CoNEXT3, pp. 1–24, 2023.
- [5] N. Zilberman, E. M. Schooler, U. Cummings, R. Manohar, D. Nafus, R. Soulé, and R. Taylor, “Toward carbon-aware networking,” *ACM SIGENERGY Energy Informatics Review*, vol. 3, no. 3, pp. 15–20, 2023.
- [6] Internet Engineering Task Force (IETF), “Green WG Documents,” <https://datatracker.ietf.org/group/green/documents/>, 2025, accessed on September 21, 2025.
- [7] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen, “Scion: Scalability, control, and isolation on next-generation networks,” in *2011 IEEE Symposium on Security and Privacy*. IEEE, 2011, pp. 212–227.
- [8] A. Ruiz-Rivera, K.-W. Chin, and S. Soh, “A novel framework to mitigate the negative impacts of green techniques on bgp,” *Journal of Network and Computer Applications*, vol. 48, pp. 22–34, 2015.
- [9] S. Kastanakis, V. Giotsas, I. Livadariu, and N. Suri, “Investigating location-aware advertisements in anycast ip networks,” in *Proceedings of the 2024 Applied Networking Research Workshop*, 2024, pp. 15–22.
- [10] “PeeringDB,” <https://www.peeringdb.com>, 2025.
- [11] R. NCC, “RIPEstat data api,” <https://stat.ripe.net/>.
- [12] T. G. W. Foundation, “The green web foundation,” <https://www.thegreenwebfoundation.org/>.
- [13] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, “Bgp-stream: a software framework for live and historical bgp data analysis,” in *Proceedings of the 2016 Internet Measurement Conference*, 2016, pp. 429–444.
- [14] RIPE NCC, “RIPEstat MaxMind GeoLite API,” 2025, accessed: March 20, 2025. [Online]. Available: <https://stat.ripe.net/docs/02.data-api/maxmind-geo-lite.html>
- [15] M.-A. Karyotakis and N. Antonopoulos, “Web communication: A content analysis of green hosting companies,” *Sustainability*, vol. 13, no. 2, p. 495, 2021.
- [16] “Nlnet; the open green web,” [accessed 2025-06-29]. [Online]. Available: <https://nlnet.nl/project/GreenWebSearch/>
- [17] J. Naab, P. Sattler, J. Jelten, O. Gasser, and G. Carle, “Prefix top lists: Gaining insights with prefixes from domain-based top lists on dns deployment,” in *Proceedings of the Internet Measurement Conference*, 2019, pp. 351–357.
- [18] M. Luckie, B. Huffaker, A. Dhamdhare, V. Giotsas, and K. Claffy, “As relationships, customer cones, and validation,” in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 243–256.