Pathways to Sustainability: Carbon-Aware Routing for Global AI Data Transfers

Nikolas EJ Schmitz
Institute of Pathology
RWTH Aachen University Hospital
nikschmitz@ukaachen.de

Dayana Savostianova
Institute of Pathology
RWTH Aachen University Hospital
dsavostianov@ukaachen.de

Leon Niggemeier
Institute of Pathology
RWTH Aachen University Hospital
lniggemeier@ukaachen.de

Martin Strauch
Institute of Pathology
RWTH Aachen University Hospital
mastrauch@ukaachen.de

Peter Boor
Institute of Pathology
RWTH Aachen University Hospital
pboor@ukaachen.de

Abstract

AI-driven applications require massive amounts of training data that are not necessarily located close to the computational infrastructure that processes them. The energy consumption of transmitting data, both for training and for routine applications of deployed AI methods, can be substantial, but is often overlooked. However, carbon emissions from data transfers could be reduced through carbon-aware routing that selects lower-emission paths through the network. These paths are selected based on the carbon intensity and the time of day in the countries whose network infrastructure is used along the path, reflecting the country-specific share of green energy. Here, we present a carbon-aware routing based on a weighted graph representation of the global internet infrastructure where time-dependent edge weights capture both the energy consumption and carbon emissions associated with data transmission across submarine cables and terrestrial links. We performed an empirical evaluation of the savings that could be achieved if such a routing was implemented in practice, showing that the carbon emissions of data transfer could be reduced by on average 40.07% if the "greenest" path was chosen over the baseline. Our work raises awareness for the fact that cloud computing causes substantial carbon emissions through the data transfer alone, and that intelligent routing could serve to reduce the carbon footprint of AI in the future.

1 Introduction

The exponential growth in artificial intelligence (AI) applications has created unprecedented demands on global data transmission networks: Modern AI training requires massive datasets that can span terabytes to petabytes, while the data volumes associated with the operational use of deployed models, such as in medicine [11], add up to even larger amounts in the long run. As the world's high-performance AI infrastructure is concentrated in a limited number of locations and countries [5], these data often need to be transferred across vast geographical distances to reach specialized compute clusters. Data centers and telecommunication networks each account for approximately 1-1.5% of the global electricity consumption [7, 8], however the carbon footprint of an individual data transmission,

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2025.

such as from a user to a data center, has so far received limited attention.
Internet data transfer generates substantial carbon emissions through the energy consumption of the network infrastructure. Estimates vary but converge around 0.06-0.2 kWh per gigabyte of data transmitted [1, 2, 9], translating to approximately 15-50 grams of CO_2 and equivalents (CO_2eq) per gigabyte depending on the carbon intensity of electricity grids. With global internet traffic projected to continue growing exponentially [10], and AI contributing increasingly to this demand, there is a growing need for carbon-aware data transmission strategies.

Recent research has demonstrated the feasibility and benefits of carbon-aware routing in computer networks. Sawsan El-Zahr et al. showed that carbon-aware routing algorithms can achieve substantial carbon footprint reductions by considering both the energy consumption characteristics of network equipment and the temporal variability of regional electricity grid carbon intensity [15]. Similarly, carbon-aware global routing in path-aware networks has shown potential for 20% emission reductions through intelligent path selection [13].
The carbon intensity of electricity varies drastically across geographical regions and temporal scales. In 2024, the average carbon intensity ranged from below 20 gCO_2/kWh in regions with a high contribution of renewable energies to over 800 gCO_2/kWh in coal-dependent regions [3]. This variability creates opportunities for intelligent routing algorithms that can leverage cleaner energy sources along different network paths. Furthermore, the dynamic nature of renewable energy generation introduces temporal variations in carbon intensity that can be exploited for time-aware routing strategies.
Our Contribution: We propose a comprehensive framework for analyzing carbon emissions from global data transmission networks, with a particular focus on the transfer of large datasets. Our approach models the internet as a weighted graph where each edge represents a physical cable connection. The graph exhibits dynamic behavior as carbon intensities fluctuate with changing electricity generation mixes. Our framework enables the evaluation of greedy path selection algorithms that can achieve reductions of on average 40.07% in carbon emissions compared to traditional shortest-path routing and a real-world routing baseline. By considering both the spatial distribution of network infrastructure and the temporal variability of regional carbon intensities, our approach provides a foundation for practical carbon-aware routing protocols that can be implemented in real-world internet infrastructure.

2 Methodology

Data We construct a comprehensive model of the internet's physical infrastructure, representing both terrestrial and submarine cables using geospatial and emission data. Our analysis integrates heterogeneous datasets to enable scenario-driven evaluation of carbon-aware routing algorithms. We source cable topology data from several public resources: *terrestrial* network data was obtained from the ITU Broadband Map [6], with information on the endpoints, physical routes, and lengths of individual cable segments; for *submarine* cable connectivity, we utilize detailed geographic metadata from the Submarine Cables Map project [12], which provides paths and landing sites for intercontinental links. Each cable, whether terrestrial or submarine, is mapped to the specific countries it traverses, enabling precise assignment of region-specific environmental properties. While country information for the terrestrial cables is provided, for the submarine cables we consider the following rule to assign countries responsible for providing the electricity and thus responsible for the carbon emissions: each country with a landing point of the given cable is equally responsible for the electricity provided to each segment of the cable.

In addition to the graph describing the cable infrastructure, we utilize a temporal dataset capturing the hourly carbon intensity of electricity generation for each country worldwide. Carbon intensity values measured in grams of CO_2 per kilowatt-hour (gCO_2/kWh) were obtained from Electricity Maps [3], which compiles real-time and historical CSV records of grid composition and emission rates. This enables our model to reflect both geographic and temporal variability in the emissions associated with data transmission.
Methods Building on these datasets, we represent the global network as a weighted graph $G = (V, E)$, where each vertex denotes a network location (city or landing point of an undersea cable) and each edge encodes a cable segment. The weight of each edge is analytically constructed to approximate the environmental cost of transmitting data over that particular path. Specifically, for a given data volume, the weight assigned to edge e corresponds to the product of the length of the cable, the data volume, a literature-derived metric for energy consumption per gigabyte-kilometer $(\mathrm{kWh}/\mathrm{GB} \cdot \mathrm{km})$ [4], and the carbon intensity of the jurisdiction through which the cable segment passes:

$$w_e = \text{Energy consumption}\left(\frac{\text{kWh}}{\text{GB}\cdot\text{km}}\right) \times \textsf{Cable length}_e(\text{km}) \times$$
$$\times\ \text{Data volume}\,(\text{GB}) \times \text{Carbon Intensity}_{\,c,t}\left(\frac{\text{gCO}_2}{\text{kWh}}\right), \quad (1)$$

where e is the edge or cable section in the cable graph, c is a country in which the edge is located, t is the date and time of the data transmission. The time dependency enables a dynamic reassignment of edge weights for any chosen hour, following e.g. the changes in solar energy production in the course of the day.

To analyze the potential for emission reductions, we implement two path selection strategies for simulated data transfers between arbitrary endpoints: (1) Shortest path: Shortest path routing, which minimizes the sum of the cable lengths, and (2) Lowest carbon intensity (CI) path: Carbon-optimal path selection, which aims to minimize the total emissions incurred for a given data transfer.
Path computation leverages standard algorithms for weighted graphs, such as Dijkstra's algorithm, adapted to handle time-dependent and heterogeneous edge weights as dictated by hourly emission profiles. For cables crossing multiple national borders, segments are treated separately to associate the appropriate carbon intensity for each portion. For each transfer scenario and hourly carbon intensity snapshot, we calculate and compare the estimated energy usage and carbon emissions incurred by both routing strategies. While the model does not capture network traffic, congestion, or operational constraints, its purpose is to estimate the upper bound of emission savings that could be achieved under real-world electricity mix variability.
Baseline path: The baseline is defined as traced network routes between two servers. The My Traceroute (MTR) service from PerfOps was used to observe the intermediate servers (hops) passed by IP packets on their way from the source, a server hosted by PerfOps, to the destination, an Amazon Web Services (AWS) S3 server. Each of these hops was geolocated using the website ipinfo.io to determine its physical location. To ensure a fair comparison, we then applied Dijkstra's algorithm to interpolate the paths between MTR-measured hops, selecting the sequence of links with the lowest possible carbon cost. This construction effectively grants the baseline the advantage of always choosing the most carbon-efficient route available among the observed hops. By defining the baseline in this way, we provide a conservative reference point against which our model can be evaluated.

3 Results

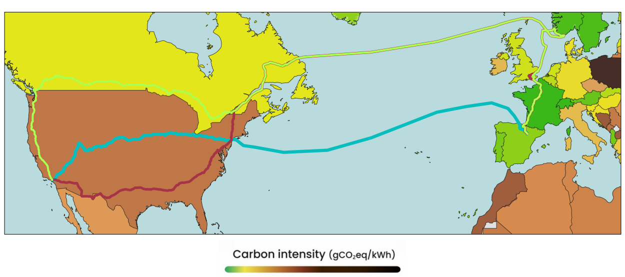

(a) Los Angeles,USA to Zaragoza,Spain with carbon costs 38.16, 51.67 (+35.38%), 22.41 (**-41.29%**)

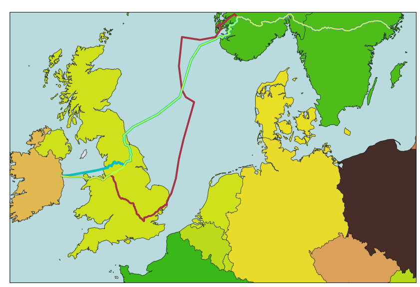

(b) Dublin, Ireland to Stockholm, Sweden with carbon costs 2.93, 2.55 (-13.14%), 2.07 (**-29.54%**)

Figure 1: Baseline, Shortest and Lowest CI (ours) paths are shown for two exemplary routes. The carbon costs of the paths (gCO_2/GB and relative change compared to Baseline) are provided in the caption in the respective color. Each country is colored according to its carbon emissions at the time of routing, from highest in dark brown to lowest in green (color legend under under left figure).

We computed 670 routes between international locations and compared the carbon-optimal Lowest CI path (ours) to the Shortest path and the Baseline path. Figure 1 visualizes exemplary routes and their carbon emissions. In Figure 2, we show the carbon emissions incurred by the carbon-optimal path compared to the baseline. Matching pairs (routes between the same locations) are connected by lines. For the majority of routes (green lines), we observed a pronounced decrease in carbon emissions with respect to the baseline. Only for few routes (yellow), in particular those that already had low emissions for the baseline, no decrease could be observed. Similarly, we observed reduced

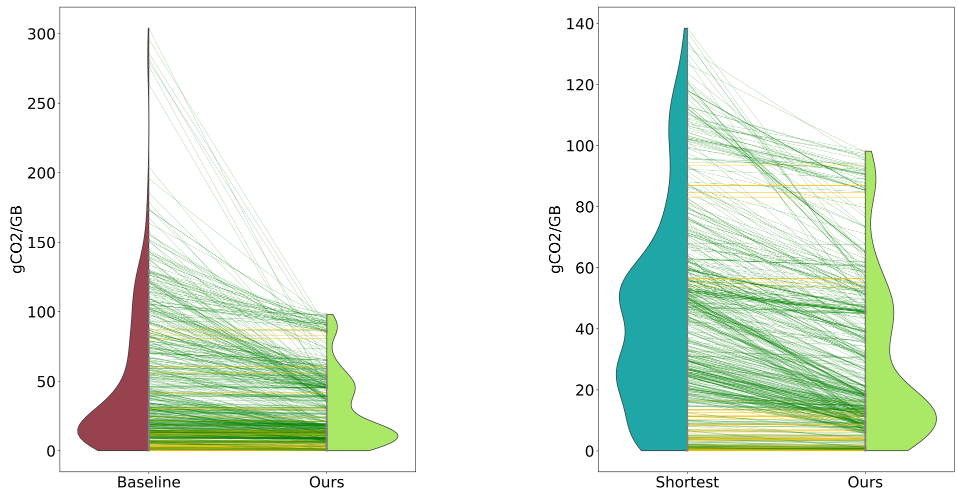

Figure 2: Carbon emissions in gCO_2/GB for Baseline vs. Lowest CI (ours) (N=670) and Shortest vs. Lowest CI (ours) (N=648: shortest paths may go through regions without CI information).
Mean (standard deviation) emissions for Baseline: 46.89 (49.05), Lowest CI: 28.1 (25.34) and Shortest: 43.98 (32.68). Corresponding pairs of routes are connected by lines. Green lines: reduction in carbon emissions; yellow lines: equal carbon emissions.

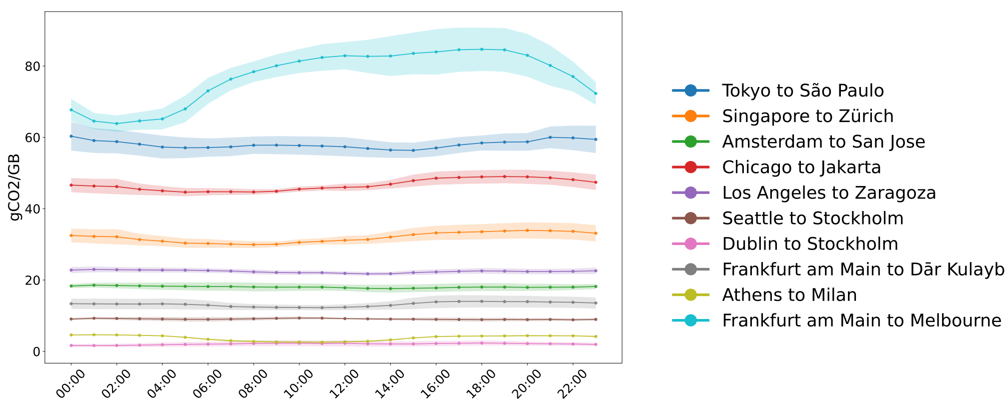

Figure 3: For each hour (measurements on 6th of January, April, June, and October to capture seasonal changes), mean and SEM of Lowest CI paths connecting selected city pairs.

carbon emissions for the majority of routes when compared to the shortest path (Figure 2).
While Figure 2 reports carbon emissions averaged over the hours of the day, Figure 3 visualizes the carbon emissions of selected routes over time: Hourly fluctuations of renewable energy production can lead to changes in carbon intensity of the paths by the hour. It is possible to combine weather data with historical carbon emission data to accurately predict the carbon intensity of the region that uses a high percentage of renewable energy sources to provide electricity [14], thus enabling an efficient time-dependent carbon-aware data transmission.

4 Conclusion

By incorporating the temporal and spatial heterogeneity of electricity grid carbon intensity into the routing process, we demonstrate that significant reductions in emissions can be achieved compared to the baseline, a realistic approximation of data transmission via internet cables. Our graph-based framework demonstrates how carbon-aware routing can reduce the carbon emissions of data transfer by on average 40.07% (over the baseline: Figure 2) through aligning network traffic with the availability of cleaner energy sources. As cloud data centers and AI hardware are concentrated in a limited number of countries [5], AI computations cause substantial data transmission traffic in the form of huge training datasets or through daily operational use. Reducing the carbon emissions of AI thus requires carbon-optimized computations as well as carbon-optimized routing.

References

[1] Joshua Aslan et al. "Electricity Intensity of Internet Data Transmission: Untangling the Estimates". In: *Journal of Industrial Ecology* 22.4 (Aug. 2018), pp. 785–798. ISSN: 1088-1980, 1530-9290. DOI: 10.1111/jiec.12630. (Visited on 04/11/2025).

[2] V.C. Coroama et al. "The Direct Energy Demand of Internet Data Flows". In: *Journal of Industrial Ecology* 17.5 (2013), pp. 680–688. ISSN: 1088-1980. DOI: 10.1111/jiec.12048.

[3] *Electricity Maps*. https://app.electricitymaps.com/map/.

[4] Gaël Guennebaud, Aurélie Bugeau, and Antoine Dudouit. "Assessing VoD Pressure on Network Power Consumption". In: *ICT4S - International Conference on Information and Communications Technology for Sustainability*. Rennes, France, June 2023. (Visited on 11/05/2024).

[5] Zoe Hawkins, Vili Lehdonvirta, and Boxi Wu. "AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators". In: *SSRN* (2025). URL: http://dx.doi.org/10.2139/ssrn.5312977.

[6] *International Telecommunication Union (ITU) interactive maps*. https://bbmaps.itu.int/bbmaps/.

[7] G. Kamiya and P. Bertoldi. *Energy consumption in data centres and broadband communication networks in the EU*. Publications Office of the European Union, 2024. URL: https://data.europa.eu/doi/10.2760/706491.

[8] B Knowles. *ACM TechBrief: Computing and Climate Change*. ACM, 2021. ISBN: 978–1-4503-9836-7.

[9] Jens Malmodin et al. "Life Cycle Assessment of ICT". In: *Journal of Industrial Ecology* 18.6 (2014), pp. 829–845. DOI: https://doi.org/10.1111/jiec.12145. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jiec.12145.

[10] Janine Morley, Kelly Widdicks, and Mike Hazas. "Digitalisation, Energy and Data Demand: The Impact of Internet Traffic on Overall and Peak Electricity Consumption". In: *Energy Research & Social Science* 38 (Apr. 2018), pp. 128–137. ISSN: 2214-6296. DOI: 10.1016/j.erss.2018.01.018. (Visited on 02/14/2025).

[11] Alireza Vafaei Sadr et al. "Operational greenhouse-gas emissions of deep learning in digital pathology: a modelling study". In: *The Lancet Digital Health* 6.1 (2024), e58–e69.

[12] *Submarine Cable Map project*. https://www.submarinecablemap.com/.

[13] Seyedali Tabaeiaghdaei et al. "Carbon-Aware Global Routing in Path-Aware Networks". In: *Proceedings of the 14th ACM International Conference on Future Energy Systems*. e-Energy '23. New York, NY, USA: Association for Computing Machinery, June 16, 2023, pp. 144–158. ISBN: 979-8-4007-0032-3. DOI: 10.1145/3575813.3595192. URL: https://dl.acm.org/doi/10.1145/3575813.3595192 (visited on 08/06/2025).

[14] Leyi Yan et al. "EnsembleCI: Ensemble Learning for Carbon Intensity Forecasting". In: *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems (e-Energy)*. 2025. DOI: 10.1145/3679240.3734630.

[15] Sawsan El-Zahr, Paul Gunning, and Noa Zilberman. "Exploring the Benefits of Carbon-Aware Routing". In: *Proceedings of the ACM on Networking* 1 (CoNEXT3 Nov. 27, 2023), pp. 1–24. ISSN: 2834-5509. DOI: 10.1145/3629165. URL: https://dl.acm.org/doi/10.1145/3629165 (visited on 08/07/2025).

A Further Visualizations and Analysis

Additional exemplary routes: Even the lowest carbon cost for the baseline in 24 hours is greater than the carbon cost of the optimized Lowest CI: Figure 4. We provide examples for the lowest carbon intensity paths in Figure 5, for baseline paths in Figure 6 and for shortest paths in Figure 7 for the same city pairs as in Figure 3. Details on the carbon costs are available in Table 1.

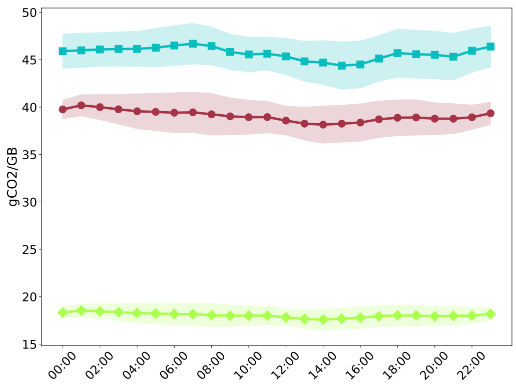

(a) Amsterdam to San Jose

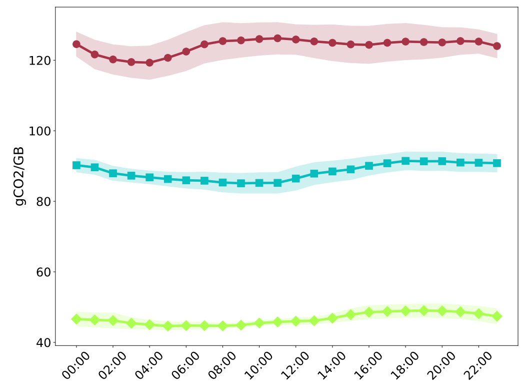

(b) Chicago to Jakarta

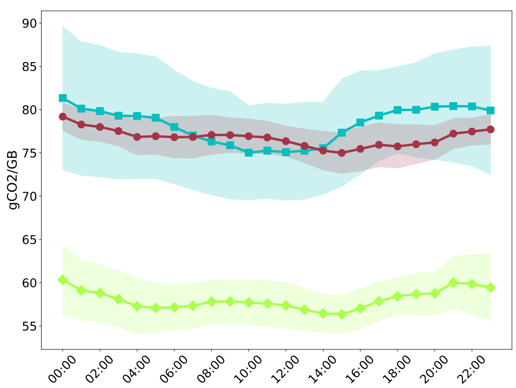

(c) Tokyo to São Paulo

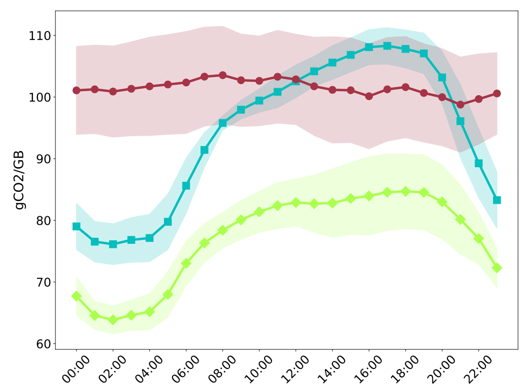

(d) Frankfurt am Main to Melbourne

Figure 4: For each hour (measures on 6th of January, April, June, and October to capture seasonal changes), mean and SEM of carbon emissions for exemplary routes are shown for the following scenarios:Baseline (•) path, Shortest (□) path, and Lowest CI (◇) path.

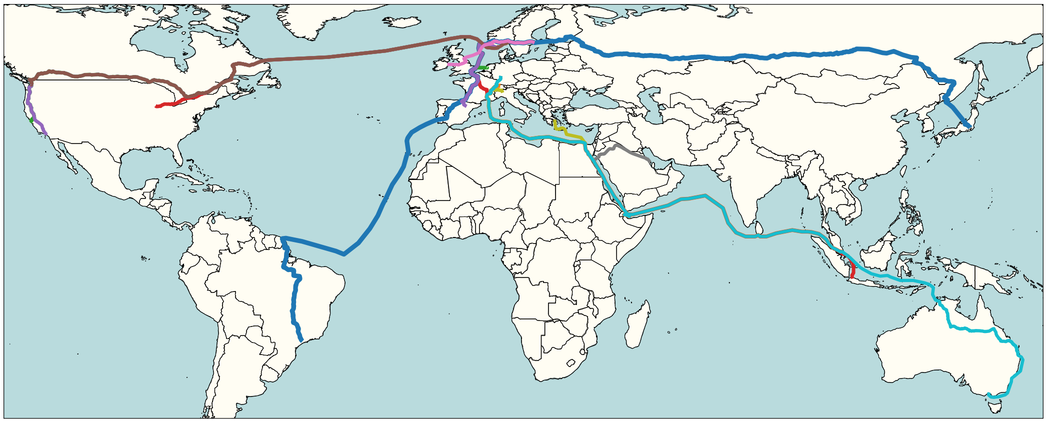

Figure 5: Lowest CI paths for the 10 city pairs in Table 1.

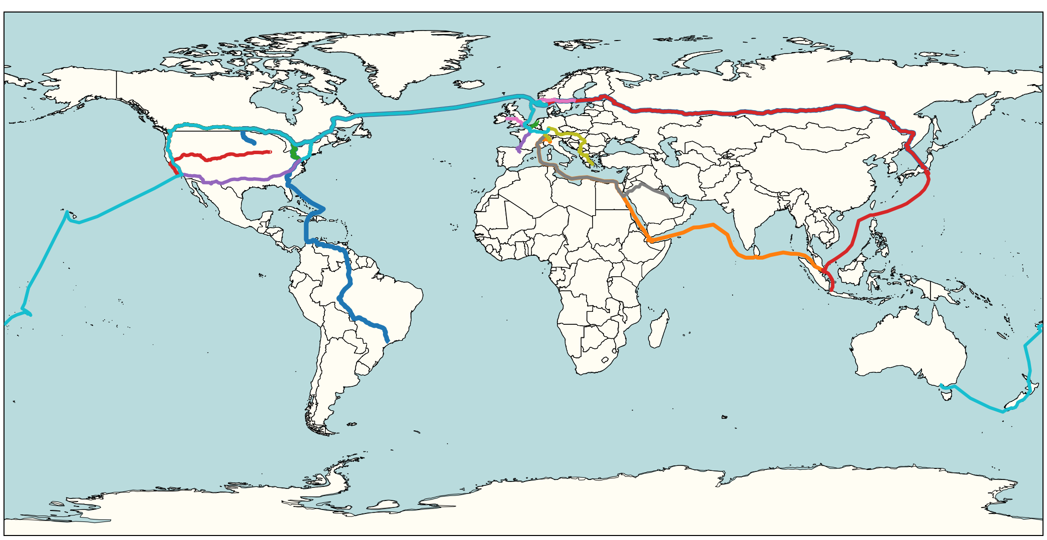

Figure 6: Baseline paths for the 10 city pairs in Table 1.

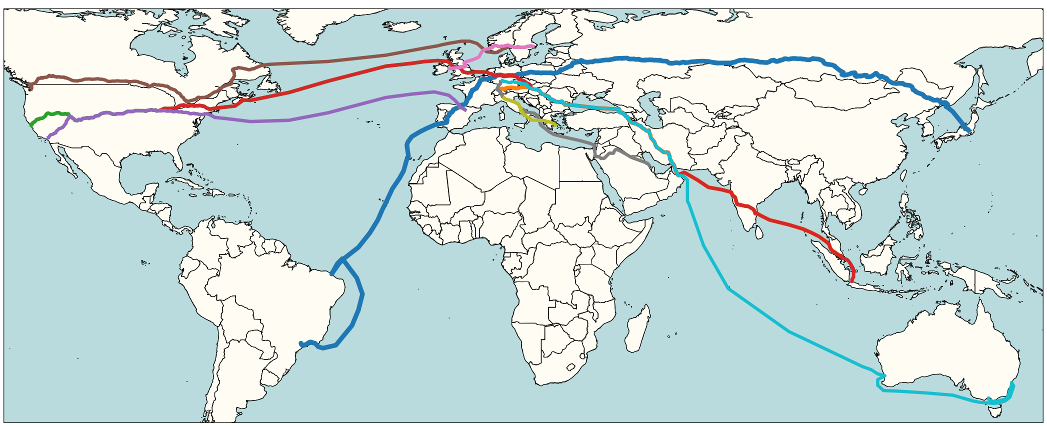

Figure 7: Shortest paths for the 10 city pairs in Table 1.

Table 1: Comparison of carbon emissions of the paths between city pairs for the three different methods. The change of the carbon cost compared to the baseline method in percent is shown in the columns "Δ vs Baseline". The data is averaged across four days in four different months (January, April, June, October) and over the 24 hours of each day.

From City	To City	Baseline		Shortest			Lowest CI (ours)		
		gCO2/GB	std (±)	gCO2/GB	Δ vs Baseline (%)	std (±)	gCO2/GB	Δ vs Baseline (%)	std (±)
Tokyo	São Paulo	76.77	4.08	78.26	1.95	12.72	58.05	**-24.39**	5.57
Singapore	Zürich	32.77	4.19	55.65	69.80	2.51	31.95	**-2.51**	3.31
Amsterdam	San Jose	39.07	3.44	45.69	16.92	4.26	18.06	**-53.79**	1.99
Chicago	Jakarta	124.02	9.18	88.35	-28.76	5.05	46.66	**-62.37**	2.97
Los Angeles	Zaragoza	38.16	1.58	51.67	35.38	2.77	22.41	**-41.29**	1.39
Seattle	Stockholm	15.28	1.81	9.09	-40.51	0.88	9.07	**-40.67**	0.88
Dublin	Stockholm	2.93	2.09	2.55	-13.14	1.61	2.07	**-29.54**	1.32
Frankfurt am Main	Dār Kulayb	13.20	2.66	19.71	49.33	0.91	13.20	**0.00**	2.66
Athens	Milan	8.40	0.97	4.37	-47.95	0.48	3.77	**-55.14**	0.36
Frankfurt am Main	Melbourne	101.49	15.60	94.11	-7.27	6.68	76.99	**-24.14**	8.38