

Toward Carbon-Aware Networking

Noa Zilberman
University of Oxford

Eve M. Schooler
Intel

Uri Cummings
Intel

Rajit Manohar
Yale University

Dawn Nafus
Intel

Robert Soulé
Yale University

Rick Taylor
Ori Industries

Abstract

Data transmission tends to be neglected when considering the carbon efficiency of systems, even though the electricity usage of data networks as a whole is as large, or larger, than that of data centers. Accounting for carbon cost of the movement of data is hard, and is often assumed to be the responsibility of the receiver or an intermediate provider.

To be able to account for the carbon footprint of networks, mutually agreed metrics are required, covering the end-to-end environmental cost of data transmission and up-the-stack network software costs of data processing, rather than merely the independent network devices.

Beyond discussing the considerations for defining these metrics, this paper suggests building upon existing practices, such as network telemetry, programmable network elements and cost-aware routing to enable carbon-intelligent networking, a concept that goes beyond network energy efficiency and considers the impact of energy decarbonization on the routing and scheduling of data transmission.

1 Introduction

Current models predict that if nothing is done to slow climate change, global temperatures may increase by 4 degrees Celsius or more by the year 2100 [8, 22, 26], posing an existential threat to humanity. Although there are many contributing factors, energy production and consumption has one of the most direct impacts on the environment. The Information/Communications Technology (ICT) industry is one of the biggest consumers of electricity. The ICT industry’s estimated consumption of worldwide electricity stands at 2-3% today, and it is predicted to increase to between 8-21% by 2030 [3, 4]. It is therefore critical that the technology industry considers not only how to reduce its electricity consumption, but also the transition to cleaner sources of energy.

However, quantifying, and subsequently reducing, the consumption of electricity is no easy task. Let’s consider a typical web-based mobile application—for convenience, let’s call it

Application X—and ask ourselves: *what is the electricity consumption of Application X?*

There are many contributors. We must account for the electricity usage of the mobile device running the application. We also must account for the infrastructure that carries the application message over the radio link to a cell tower. Then, we add the cell tower, which is typically shared by more than one carrier, and, the fiber connecting the tower to the Internet backbone networks, owned by various Internet service providers. We must include the data center, which runs the application logic in a cloud platform shared by different businesses. The application logic is likely distributed, sending additional messages in the form of remote procedure calls (RPCs) to other services. If we care about sustainability, we need to measure the aggregated impact of each of these components.

There have been several promising initiatives focusing on some of the individual contributors. For example, reducing the power consumption of microprocessors was a major focus of the computer architecture community in the 2000s, leading to both new research as well as tools for modeling power consumption [9]. Reducing data center power consumption has long been a goal for industry, perhaps because there is a direct operational cost. However, there is still a lot to be accounted for when trying to accurately measure the electricity consumed, let alone the carbon footprint of sending a packet from one’s phone to a service running in a data center.

The accounting gap can be attributed to several factors, including a lack of awareness of the problem; a lack of standards for who should be responsible for collecting and attesting to what data; and a lack of tools for collecting data. In the systems research community, much of the work on power consumption has focused on data centers, rather than end-to-end networked systems [23, 31, 42].

On the standards side, when there are multiple parties involved, it is often unclear who should take “ownership”. For example, if there are Netflix applications running on a server at an ISP’s Point of Presence (PoP), then it is expected that the ISP would account for energy consumption [29]. On the tools side, there are a wide range of network monitoring tools

for collecting traditional metrics, such as latency, throughput, and packet loss. However, these metrics are related to performance or security, rather than power consumption and its link to carbon impact.

In his book on software dynamics, Sites [35] argues that the network is one of the four fundamental resources that must be considered when reasoning about end-to-end performance issues. He calls the unattributed communication time “the slop”. Similarly, in this paper, we argue that the network is also one of the fundamental resources that need to be observed when measuring energy consumption. We need to measure the electricity consumption of “the slop”, including all of its components.

Knowing the total electricity usage is not enough. We must also be able to derive the carbon-intensity of that electricity, in order to determine the carbon footprint of network elements, whether hardware or software, when they are pressed into service. Technically, carbon intensity is defined as the amount of carbon by weight emitted per unit of energy consumed. For the purposes of this paper, we use the term informally to convey how “green” is an energy source. Simply put, the carbon intensity is an important factor in the carbon footprint equation; the lower the better. In this paper, we acknowledge that the terms power-efficiency and -consumption, energy-efficiency and -consumption and carbon-efficiency and -emissions, are interlinked but not interchangeable.

2 The Energy Consumption of Networks

Today, it is hard to quantify the energy consumption of networks that comprise what we think of as the Internet infrastructure. There have been a few studies that analyze networks within the broader ICT footprint [20,24,25], but they are somewhat limited due to the scarcity of data available. Moreover, there was a long delay from when that data was collected and when those results were published, meaning that existing studies have stale information that fail to capture ongoing changes in networking hardware, protocols, scale, traffic density, etc. Given that the number of Internet users is expected to grow by 29.4% between 2018—2023, and mobile IoT connections 3.6-fold over the same period [12], this lag in reporting is much less acceptable.

The data that we do have suggests that networks are a dominant component in the carbon footprint of digital infrastructure. Some studies even suggest that networks have $1.5\times$ the electricity usage of data centers [24], although their scope is unclear. Barroso [7] indicates that within data centers, networks already account for 10%-20% of the energy.

While cloud service providers such as Amazon consumed 24TWh of electricity in 2020 [2], and Google consumed 15.4TWh of energy in 2021 [16], Internet service providers (ISP) consume energy on the same scale. For example, BT consumed 3.3TWh [11], Vodafone consumed 5.8TWh [40] and Telefonica consumed 6.1TWh [36]. AT&T consumed

17.1TWh [5] across all its businesses, which may not be indicative of the digital infrastructure alone. Given the large number of ISPs worldwide, the aggregate energy consumed by communications is clearly substantial.

ISPs are already significantly invested in improving their networks’ energy consumption. For example, Telefonica reports its energy consumption per petabyte (PB) of data, and set a goal to reduce this consumption by 90% between 2015 and 2025, already achieving 86% reduction by 2021. However, Telefonica’s traffic cost is still 54MWh/PB, and it has seen $6.7\times$ growth in traffic between 2015 and 2021. It is conservative to assume that network demand will only increase, with next generation networks roll-out, the growth in use of edge computing, the adoption of IoT, and the sheer volume of data being created by users and for users [39].

3 Carbon Reporting Metrics: Technical Challenges

Most of the sustainability reports previously quoted [2, 11, 16, 36, 40] report their energy consumption and carbon efficiency as a bulk number, the aggregation of all consumption contributions. This approach, however, is ill fitted where networking is considered.

First, the carbon efficiency of a network service or a networked-application operation can cross multiple administrative domains, and may be mixed with other services sharing the network. Second, similar to other computing hardware such as CPUs, the power consumption may be sensitive to traffic load, and different devices, based on different technologies (e.g., ASIC, SoC, FPGA) and manufacturing processes, will have a different utilization-to-power profiles. Some devices will be more power-proportional (e.g., ASIC), while others (e.g., FPGA) will have an almost constant consumption for a given program [37].

On top of the above, the carbon footprint of a network platform is composed of multiple components: a network chip (e.g., switch ASIC), transceivers, fans, control and management devices, power supply and control, and others. A vendor may choose to present a device as environmentally friendly, but neglect additional carbon sources such as the effect of transceivers (e.g., $64\times 400G$ vs $256\times 100G$) or the need for co-processing devices. Additionally, the work done in the network stack at the end-point also contributes to the overall carbon footprint.

To properly account for the carbon efficiency of networking, we argue for an end-to-end approach, as is commonly recommended for applications in security, fault tolerance, and reliable delivery [32]. Specifically: (1) Devices should be able to report their real-time or near real-time electricity consumption, (2) Devices should be able to report the carbon-intensity or quality of consumed electricity, (3) There should be a mechanism to tie items 1 and 2 back to an application, e.g., what

percentage of electricity does an individual application consume, (4) Applications and services should react in (near) real-time to carbon-related information collected from the network.

With these four items, we could create end-to-end networking solutions optimized for carbon efficiency with maximum coverage. Below, we discuss the technical challenges related to the above.

Reporting real-time electricity consumption. Today, network equipment manufacturers tend to report the maximum power consumption of a platform, for power and cooling purposes, but this may not be a suitable metric; the difference between average and maximum may be large (or small), and may not represent the actual platform carbon emissions.

Reporting the electricity consumption of network devices is an engineering project more than a technology challenge. The ability to monitor the power consumption of different components of a platform exists — but requires vendors to add support for it on their platforms. In particular, power consumption information needs to be fed back to the platform itself, and continuously so.

The absence of hardware support within the platform does not mean that we need to wait for new devices to come to the market. It is possible to leverage proxy data that will indicate usage (e.g., in switches, there is a correspondence between throughput and power consumption).

To make use of the real-time information, there is a need for an end-to-end reporting mechanism. Just like in-network telemetry is used to analyze end-to-end network performance, the reporting mechanism will utilize the network itself for the purpose of collecting statistics, combined with measurements of the software stack on the endpoint. This mechanism can be viewed as “the network telemetry of carbon efficiency.”

Reporting electricity carbon intensity. Electricity consumption is not an indication of carbon emissions, as the carbon intensity of the energy source must be factored in. Therefore, it is required that a device not only reports electricity consumption, but also the carbon intensity of the electricity it consumed. When comparing network elements, a coarse grain distinction could be made between elements consuming electricity from renewable energy sources versus fossil fuel, while more fine grain distinctions might include the embodied carbon or energy losses.

Several organizations have created APIs to deliver near real-time measurements for carbon-intensity [13] [41] [34] [38]. While these are already being used by cloud service providers to make decisions about where best to place workloads resulting in the least carbon footprint, network operators have yet to embrace this knowledge operationally (discussed further in Section 5.1).

Although these APIs exist, the availability of carbon intensity data is not without its challenges. While many regions

globally are making carbon intensity data available publicly, coverage is incomplete. Additionally, the frequency of the data updates varies considerably across regions. Some operators report relatively static average values over large regions, whereas others report at finer granularities such as minutes or hours, depending on the nature of the energy supply (solar, wind, battery). Traditionally in networking, high-resolution information is preferred, for example, to detect micro-bursts of traffic. This isn’t necessarily a requirement for reporting carbon intensity, a measurement that comes from the electrical grid. As more renewables are integrated into the grid, there will be a proliferation of smaller regions reporting carbon-intensity measurements, for a growing number of non-utility-owned distributed energy resources. Thus there is a need for finer grain spatial data, beyond the coarse-grain zone boundaries currently defined by the independent system operators that coordinate, control and monitor the operation of the electrical grid. The measurement data must also be verifiable, especially if it is being used to prove regulatory compliance to emissions thresholds or reductions, and a need for these independent resources to communicate with the broader electric grid infrastructure.

Smart grids have been taking advantage of Advanced Metering Infrastructure (AMI) for years now, using bi-directional power line communication, to collect information from smart meters [18]. The same technology can be used by the energy supplier to provide information about the source and the quality of the electricity in-band with the supply. Both the energy provider, and the network operator, would need to add new support: the energy provider by sending electricity-quality information (periodically or through dedicated API), and the network operator by processing and propagating this information as part of in-network telemetry (INT) [21] updates.

Tying electricity consumption and carbon intensity to applications. Given the reporting of real-time electricity consumption, and the carbon intensity of the electricity, the two can be exposed to applications. On the end host side, carbon emissions of all devices could become an IO device on the system, exposing the data through a system level API (e.g., `/dev/carbon`). Existing telemetry infrastructure (e.g., Intel’s DeepInsight) could read this information, and tie it to an application, or even packet level, using an in-network telemetry solution [21].

Section 5 discusses some challenges and proposals to turn the information into a working, useful solution. In particular, at the first stage it is assumed that a solution will be limited to a single administrative system, where the operator has full knowledge and control of the deployed network platforms. As many applications will be running in parallel, exposing application-specific insights from telemetry information remains a challenge.

4 Carbon Reporting Metrics: Policy

A network platform has both an embedded carbon footprint, from its manufacturing, and an operational carbon footprint, from its usage. In this paper it is assumed that the operational footprint eclipses the embedded footprint, based on the expected lifetime of network elements and their “size” (higher performance devices that consume a large amount of electricity) [20]. However, this may not be a valid assumption as more consumer devices, take on networking services.

We expect network device and equipment vendors will want to be able to leverage the energy ratings of their devices for comparative purposes, attracting customers seeking sustainable designs; thus these metrics need to be adapted for a carbon-aware networking context.

To uncover the real cost of networking and adopt sustainable solutions, the following steps should be taken:

- **Use standard metrics.** Manufacturers, service providers and users should all use agreed upon metrics. These metrics should be networking-specific and standardized through organizations such as the ITU, IEEE, and IETF. Two green routing metrics proposed in [19] are carbon intensity and environmental waste.
- **Provide carbon efficiency under different loads.** The carbon efficiency of devices changes under load, and maximum figures may misrepresent power-proportional devices. To this end, carbon efficiency metrics should be reported under different loads. As link-utilization and packet-rate may yield different results, the agreed metrics should consider using the functionality available on devices for reporting purposes (See §5).
- **Provide measured results.** In platforms built from multiple components, the carbon efficiency rating should rely on measurements of the platform as a whole. Using maximum or datasheet-reported value of each individual component is easy to do, but will misrepresent (often for the worse) the real carbon efficiency of a platform. For devices, accounting for mandatory overheads (e.g., power, cooling, transceivers) is needed to truly represent their operational carbon footprint.
- **Using standard evaluation environments.** Using standard evaluation setups is common in electronic devices, but software systems suffer from a reproducibility crisis. To accurately report network platforms carbon efficiency, the evaluation environments and tests must be standard and reproducible, covering both hardware and software aspects. In particular, where network devices are concerned, the internal configuration of devices (e.g., stages, acceleration engines) should be known, as otherwise vendors may turn off functionality to achieve better energy rating.
- **Avoid double counting.** Reporting the end-to-end power efficiency of a system or of networks’ use should avoid

double counting of intermediate elements (e.g., switches) by multiple elements of the system.

- **Trustworthy networking.** The AI community has developed and adopted mechanisms for trustworthy AI, ranging from algorithms and software to systems and hardware [6, 10]. Many mechanisms, such as auditing and interpretability, the latter of which addresses the needs of audiences with diverse expertise, should be adopted by the networking community when addressing carbon-efficiency.
- **Real-Time Observability.** All the previous recommendations indicate that there is a need for network devices to extract, derive, account for, and report their actual energy consumption. Providing real time information will allow the construction of a closer-to-reality picture of the cross-layer effects of networking on the environment.

5 Towards Carbon-Aware Networking

The presence of carbon-efficiency metrics, and the use of INT, enables the development of new carbon-aware routing. Building upon past cost-aware routing algorithms [14, 27, 43], carbon-aware routing will try to use the most carbon-efficient end-to-end route, based on information from network elements along the way.

As suggested in Section 3, a simple approach would limit the routing to a single autonomous system (AS), where the administrator has full control over the network. Using existing knowledge of the devices within the network, or by collecting the energy rating of devices using a variant of INT [21], routes can be energy-rated. The optimization function could then seek the least (best) end-to-end energy-rated path.

The above approach, however, does not account for the length of a route. Would N hops through A-rated switches be better than $N/2$ hops through B-rated devices? Therefore, route calculations must be *weighted*, with a different weight for every energy rating level.

This second approach is not ideal either, as it is likely to lead to congestion and does not account for differences in energy efficiency under different loads. A better approach would consider the current load on different devices in the network and the carbon efficiency gradient of a device. This becomes a multi-route optimization problem. If the problem is limited to a single administrative entity, which also updates the various routing tables, this is feasible.

The discussion so far assumed that identical switches have identical carbon footprints, no matter their location. This is, however, inaccurate. Switches powered by renewable energy will be “greener” than those powered from fossil-fuel sources, and should be preferred. While the trade-off between decarbonization and power efficiency of devices is open to debate (i.e., should a switch powered by solar energy be preferred over one powered by fossil fuel, even if power consumption

is $\times 100$ higher?), this question is beyond the scope of this paper. Nonetheless, we advocate to enable visibility into these measurements for others to decide.

Below, we make the distinction between carbon-aware versus carbon-intelligent, i.e., knowing the carbon emissions and minimizing them while still applying standard routing practices, versus knowing the carbon emissions and taking different approaches to routing and scheduling of data-transfer, such as delaying transmission to align with the availability of (excess) renewables.

5.1 Carbon-Intelligent Routing

Carbon-intelligent¹ computing [30] has been developed by cloud providers to account for carbon intensity when allocating compute capacity; workloads are time- or space-shifted to maximize the usage of renewable energy, to reduce the ICT footprint. In addition, carbon-intelligent computing can behave like a virtual battery [1] to help the electrical grid to consume excess renewable energy that would otherwise be wasted. Similarly, carbon-intelligent routing has an important role to play to allow the transfer of data across the network in a manner that is most carbon-efficient, employing similar carbon-intelligent techniques.

Inspiration can be taken from deterministic networking technologies such as DetNet [15] that aim to deliver certain guaranteed network behaviors, such as a worst case end-to-end latency. The idea here would be to expand the metrics that define network determinism, to request that data transmission stay below a fixed carbon-intensity threshold or within an overall carbon footprint budget along a route. Delay Tolerant Networks (DTN) can therefore be an ideal technology to enable carbon-intelligent data transfer, supporting the time-shifting or deferral of data transmissions to align with the production and storage of clean energy.

Content Distribution Networks make an even more compelling case, due to the dominance of streaming content in the network [33]. Currently, content providers distribute content to caches during off-peak hours, to reduce network congestion during peak hours. Carbon-intelligent routing will optimize content distribution also by the availability of renewable energy and the optimization of routing for carbon-efficiency.

Carbon-intelligent routing will lead to geographically localized routing decisions. If currently cached video content is distributed from the US to Italy and to Norway around the same time, due to similarity in time zones and peak/off-peak hours, the strategy might be different when energy sources are considered: Norway’s main renewable energy is hydropower that is continuously available, while Italy’s is solar that is only intermittently.

¹The terms carbon-intelligent [30] and carbon-responsive [28] are often used interchangeably.

5.2 Carbon-Intelligent Network Telemetry

To support carbon-intelligent routing, information is required about the properties of the network devices along the routing path. In-network telemetry enables collecting this information, adding reports from every device along a route to a telemetry packet’s header.

Section 5 described some of the challenges and pathways to carbon-aware and carbon-intelligent routing. The ability to develop algorithms for either depends on the information available through telemetry. Information such as the energy rating of a device is easy to collect, as this is static information. The use of renewable energy is another type of information that can be collected, either as a binary indicator (yes/no), a carbon-intensity (percentage of energy mix that is carbon-based) or as a bitmap or heatmap [17] indicating the times when renewable energy is available.

Collecting and stewarding more comprehensive information is an open research question. For example, information about the current power consumption of a network device is typically not available from within the chip. Moreover, this information is insufficient, as additional overheads need to be taken into account (fans, power supplies, transceivers, and more). It is not a major technological challenge to build a CPU routine that collects this information and periodically loads it to a programmable switch’s register, however it will be a challenge to widely deploy such solutions.

6 Conclusion and a Call to Action

Carbon-awareness and intelligence in networking will play across multiple layers: from the physical layer and the routing layer to the application layer, and including the network stack. Having visibility in each of these layers into the consequences of data transfer means that better decisions can be made: applications and services that consume less carbon resources will be preferred, and routes will be taken where carbon footprint is minimized. The confluence of cross-layer efforts is mandatory to achieve a tangible effect.

Networking needs to be carbon-efficient, like any other part of digital infrastructure. Its disaggregated nature makes its accounting harder, but possible. The proposed solutions described here are only early stepping stones. To achieve real progress, standard metrics need to be supported, and reported, by network devices. We call the community to join the effort to define these metrics and to work with the electrical grid community to standardize carbon-intensity data.

Designing carbon-intelligent routing solutions is the next big challenge in networking, encompassing technical, organizational, and cross-disciplinary aspects. New algorithms are needed, and understanding how to balance energy consumption with decarbonization, attending to the dynamic nature of changes in efficiency, and providing accountability and interpretability. Together, we can make networking truly green.

References

- [1] Anup Agarwal, Jinghan Sun, Shadi Noghbi, Srinivasan Iyengar, Anirudh Badam, Ranveer Chandra, Srinivasan Seshan, and Shivkumar Kalyanaram. Redesigning data centers for renewable energy. In *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*, HotNets '21, page 45–52, 2021.
- [2] Amazon. Amazon sustainability 2020 report: Further and faster, together. Online, June 2021. <https://sustainability.aboutamazon.com/amazon-sustainability-2020-report.pdf>[Accessed May 2022].
- [3] A. S. Andrae and T. Edler. On global electricity usage of communication technology: trends to 2030. *Challenges*, 6(1):117–157, 2015.
- [4] Andrae, A. S. New perspectives on internet electricity use in 2030. *Engineering and Applied Science Letters*, 3(2):19–31, 2020.
- [5] AT&T. Key Performance Indicators (KPIs). Online. <https://about.att.com/ecms/dam/csr/2021/KPI/2021-KPI-Tables.pdf>[Accessed May 2022].
- [6] Shahar Avin, Haydn Belfield, Miles Brundage, Gretchen Krueger, Jasmine Wang, Adrian Weller, Markus Anderljung, Igor Krawczuk, David Krueger, Jonathan Lebensold, Tegan Maharaj, and Noa Zilberman. Filling gaps in trustworthy development of AI. *Science*, 374(6573):1327–1329, 2021.
- [7] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan. The datacenter as a computer: Designing warehouse-scale machines. *Synthesis Lectures on Computer Architecture*, 13(3):i–189, 2018.
- [8] Sophie Boehm, Katie Lebling, Kelly Levin, Hanna Fekete, Joel Jaeger, Richard Waite, Anna Nilsson, Joe Thwaites, Ryan Wilson, Andreas Geiges, et al. State of climate action 2021: Systems transformations required to limit global warming to 1.5c. *Washington, DC: World Resources Institute*, 2021.
- [9] David M Brooks, Pradip Bose, Stanley E Schuster, Hans Jacobson, Prabhakar N Kudva, Alper Buyuktosunoglu, John Wellman, Victor Zyuban, Manish Gupta, and Peter W Cook. Power-aware microarchitecture: Design and modeling challenges for next-generation microprocessors. *IEEE Micro*, 20(6):26–44, 2000.
- [10] Miles Brundage, Shahar Avin, Jasmine Wang, Haydn Belfield, Gretchen Krueger, Gillian Hadfield, Heidy Khlaaf, Jingying Yang, Helen Toner, Ruth Fong, et al. Toward trustworthy AI development: mechanisms for supporting verifiable claims. *arXiv preprint arXiv:2004.07213*, 2020.
- [11] "BT Group Plc". Tackling climate change and environmental challenges. Online. <https://www.bt.com/bt-plc/assets/documents/digital-impact-and-sustainability/our-report/report-archive/2021/tackling-climate-change-and-environmental-challenges.pdf>[Accessed May 2022].
- [12] Cisco. Cisco annual internet report (2018–2023) white paper. Online, 2020. <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>.
- [13] Electricity Map. Electricity Map API. Online, 2022. <https://static.electricitymap.org/api/docs/index.html>[Accessed June 2022].
- [14] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. A bgp-based mechanism for lowest-cost routing. *Distributed Computing*, 18(1):61–72, 2005.
- [15] Norman Finn, Pascal Thubert, Balazs Varga, and Janos Farkas. Deterministic networking architecture. *RFC8655*, 2019. <https://datatracker.ietf.org/doc/rfc8655/>.
- [16] Google. Google environmental report 2021. Online. <https://www.gstatic.com/gumdrop/sustainability/google-2021-environmental-report.pdf>[Accessed May 2022].
- [17] Google Carbon-free Energy blog. Operating on 24/7 Carbon-Free Energy by 2030. Online, 2022. <https://sustainability.google/progress/energy/>[Accessed May 2022].
- [18] David G Hart. Using ami to realize the smart grid. In *2008 IEEE Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century*, volume 10, 2008.
- [19] Hossain, Md. Mohaimenul, Georges, Jean-Phillippe, Rondeau, Eric, Divoux, Thierry. Energy, carbon and renewable energy: Candidate metrics for green-aware routing? *Sensors*, 19(13), Feb 2019.
- [20] ITU-T Recommendation L.147. Online. <https://www.itu.int/rec/T-REC-L.1470-202001-I/en>[Accessed May 2022].
- [21] Changhoon Kim, Anirudh Sivaraman, Naga Katta, Antonin Bas, Advait Dixit, and Lawrence J Wobker. In-band network telemetry via programmable dataplanes. In *ACM SIGCOMM*, volume 15, 2015.
- [22] Katie Lebling, Ge Mengpin, Kelly Levin, Richard Waite, Johannes Friedrich, Cynthia Elliott, Christina Chan, Katherine Ross, Fred Stolle, and Nancy Harris. State of climate action 2021: Assessing progress toward 2030 and 2050. *Washington, DC: World Resources Institute*, 2022.
- [23] Yang Li, Charles R. Lefurgy, Karthick Rajamani, Malcolm S. Allen-Ware, Guillermo J. Silva, Daniel D. Heimsoth, Saugata Ghose, and Onur Mutlu. A scalable priority-aware approach to managing data center server power. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 701–714, 2019.
- [24] Josip Lorincz, Antonio Capone, and Jinsong Wu. Greener, energy-efficient and sustainable networks: State-of-the-art and new trends. *Sensors*, 19(22), 2019.
- [25] Jens Malmodin and Dag Lundén. The energy and carbon footprint of the global ict and e&m sectors 2010–2015. *Sustainability*, 10(9), 2018.
- [26] V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelečki, R. Yu, and B. Zhou, editors. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 33–144. 2021.
- [27] Dritan Nace and Michal Pióro. Max-min fairness and its applications to routing and load-balancing in communication networks: A tutorial. *IEEE Communications Surveys & Tutorials*, 10(4):5–17, 2008.
- [28] Dawn Nafus, Eve M. Schooler, and Karly Ann Burch. Carbon-responsive computing: Changing the nexus between energy and computing. *Energies*, 14(21), 2021.
- [29] Netflix environmental social governance report 2021. Online. https://s22.q4cdn.com/959853165/files/doc_downloads/2022/03/30/2022_US_EN_Netflix_EnvironmentalSocialGovernanceReport-2021-FINAL.pdf[Accessed May 2022].
- [30] Ana Radovanovic, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, et al. Carbon-aware computing for datacenters. *arXiv preprint arXiv:2106.11750*, 2021.
- [31] Varun Sakalkar, Vasileios Kontorinis, David Landhuis, Shaohong Li, Darren De Ronde, Thomas Blooming, Anand Ramesh, James Kennedy, Christopher Malone, Jimmy Clidas, and Parthasarathy Ranganathan. Data center power oversubscription with a medium voltage power plane and priority-aware capping. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, page 497–511, 2020.
- [32] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, nov 1984.

- [33] Sandvine. Global internet phenomena report. Online, 2022. <https://www.sandvine.com/global-internet-phenomena-report-2022>.
- [34] Singularity. Singularity API. Online, 2022. <https://singularity-docs.stopligh.io>[Accessed June 2022].
- [35] Richard L. Sites. *Understanding Software Dynamics*. Addison-Wesley Professional, 2021.
- [36] Telefónica, S. A. 2021 consolidated management report. Online. <https://www.telefonica.com/en/wp-content/uploads/sites/5/2022/03/consolidated-management-report-2021.pdf>[Accessed May 2022].
- [37] Yuta Tokusashi, Huynh Tu Dang, Fernando Pedone, Robert Soulé, and Noa Zilberman. The case for in-network computing on demand. In *Proceedings of the Fourteenth EuroSys Conference 2019*, pages 1–16, 2019.
- [38] UK National Grid ESO. National Grid ESO Carbon Intensity API. Online, 2022. <https://carbonintensity.org.uk/>[Accessed May 2022].
- [39] Rob van der Meulen. What edge computing means for infrastructure and operations leaders. *Smarter with Gartner*, 2018. <https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leader>[accessed May 2022].
- [40] Vodafone Group Plc. Annual report 2021. Online. <https://investors.vodafone.com/sites/vodafone-ir/files/2021-05/vodafone-annual-report-2021.pdf>[Accessed May 2022].
- [41] WattTime. WattTime API. Online, 2022. <https://www.watttime.org/api-documentation/>[Accessed June 2022].
- [42] Qiang Wu, Qingyuan Deng, Lakshmi Ganesh, Chang-Hong Hsu, Yun Jin, Sanjeev Kumar, Bin Li, Justin Meza, and Yee Jiun Song. Dynamo: Facebook’s data center-wide power management system. In *Proceedings of the 43rd International Symposium on Computer Architecture, ISCA ’16*, page 469–480, 2016.
- [43] Bernard Yaged Jr. Minimum cost routing for static network models. *Networks*, 1(2):139–172, 1971.