A Perspective on Carbon-aware Networking

Eve M. Schooler (Intel), Rick Taylor (Ori Industries), Noa Zilberman (Oxford), Robert Soulé (Yale), Dawn Nafus (Intel), Rajit Manohar (Yale), Uri Cummings (Intel)

October 31, 2022

Introduction

Given the inherently distributed nature of modern applications and the system infrastructure that supports them, everything relies on the often hidden capabilities of the networks that connect them - and everything incurs environmental costs due to those networks. As a consequence, it is important to understand how to measure and/or model the environmental impact of communication networks.

Despite the attention that hyperscaler data centers (DCs) continue to receive regarding their electricity usage, networks have been found to have as large or larger an impact, by as much as 1.5x the electricity consumption [8]. With ever-increasing projections for the volume of data expected to be created [17], the need for increased data transmission capacity, and the desire to perform data analytics and compute earlier rather than later in the data lifecycle, the carbon footprint for the information communication technology (ICT) industry, and the footprint of the networks that underpin it, must be expected to grow at a similar rate.

To combat the growth in environmental impact, one emerging solution in large cloud DCs is to employ *carbon-aware computing*, i.e., to time- or space-shift compute workloads to maximize the usage of low *carbon-intensity* energy [11]. Technically, carbon intensity is defined as the amount of carbon by weight emitted per unit of energy consumed. We use the term in this position paper 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. Hyperscalers are able to employ carbon-awareness by physically co-locating their resources with hydro-electric or other on-grid energy storage systems, but those solutions are largely only applicable to the core of the cloud ecosystem and are proprietary. By under- standing the availability and origin of energy at locations within the hyperscaler's infrastructure , suitable workloads can be postponed or migrated to minimize the carbon footprint of hosting the services.

While carbon-aware computing aims to reduce the carbon footprint of ICT's computational load, it has the beneficial side effect that it solves a long-standing problem that renewable energy generation is *variable*, meaning it is only generated when the wind is blowing or the sun is shining, and thus may or may not align with energy demands of Internet consumers. As a result, carbon-aware computing techniques are notable in that they can be used to shift workloads to consume *excess* renewable energy, which would otherwise go unused and cause instability in the electric grid. In this regard, carbon-aware computing can be considered a means of load balancing the grid or serving as a virtual battery [1]. This is a particularly acute problem in parts of the globe where renewables might account for a sizable and growing percentage of energy generation (such as CA or Germany) [2].

With hyperscaler DCs operationalizing carbon-aware computing [5] [9], a natural question to raise is how to bring carbon-awareness not only to orchestrated computing workloads, but to all of ICT [10]: across not only computing, but also networking and storage-oriented systems; from hardware to software components, inclusive of programmable and virtualized elements; vertically across the software stack resident on a single platform and horizontally across end-to-end transactional systems. For the purposes of this position paper, building on thoughts captured in [13][19], we focus on carbon-aware networking as a key ingredient in the arsenal of mechanisms to rein in the carbon footprint of ICT and likely to result in substantial savings [14].

Facets of Carbon-aware Networking

There are several facets to carbon-aware networking worthy of discussion in the IAB workshop on Environmental Impact of Internet Applications and Services that raise open research questions:

Carbon-aware routing. Is it possible to teach routing protocols to select more carbon-efficient paths over others? These would be paths that result in a lower carbon footprint, where carbon footprint for each $\texttt{node}_i$ encountered along the path is calculated as `energy`$_i$ `x carbon-intensity`$_i$. With better insights into the (excess) availability and carbon intensity of the energy, traffic could be routed along network paths that not only minimize carbon footprint, but also help to stabilize renewables integration. If this excess energy has a reduced price, this path selection should yield operating cost benefits. This isn't to say that the minimization of carbon-intensity is the only objective function; it is likely that path selection will be a joint optimization problem, which also considers minimization of traditional attributes such as latency, packet loss, and jitter.

Time-variant routing. There is a movement afoot to make routing topology changes schedulable (versus event-driven), for example when a renewable energy source becomes (un)available. In other words, a path might only become available when its carbon consumption is below an established maximum threshold or is close to zero. Thus the routing topology might change when a wind or solar farm has come on line (or gone off line), making a part of the network infrastructure available (or unavailable) that wasn't previously. Effectively, this makes the routing infrastructure more closely coupled with requirements for sustainable operation. By using TVR intelligently, we could allow bits of the platform and broader system to go to sleep, without losing context, something that currently can be highly disruptive as such a loss is considered a "failure" by the route maintenance subsystem. The key question: can TVR adequately solve the disruption problem? If the routing protocol peering system has sufficient warning to be able to declare a priori that it will disconnect (or reconnect) in some number of minutes, then the system can react more sensibly as well as more carbon efficiently.

Carbon-aware transport. Popping up the stack, carbon-aware transport protocols must exist that can react to signaling from lower-layer carbon-aware routing protocols. These are the transport mechanisms that would perform the deliberate time- and/or space-shifting of a data transmission to minimize the carbon cost of delivering the payload. They likely also require novel congestion control mechanisms. These transports are meant to support time-elastic usages, i.e., use cases not unduly disrupted by stretching or compressing the time window within which the transmission is bounded. This is admittedly inspired by Delay Tolerant Networking (DTN), and this potentially identifies a new use for DTN, driven not by the intermittency of deep space connectivity but rather by the inter- mittency of available renewables or other forms of clean energy.

Carbon-aware traffic engineering. In the deterministic networking (DetNet) realm [3], it is not uncommon to encounter a latency budget within which a transmission has been asked to fit. Here, we propose to extend network determinism to include carbon-intensity and demand that a flow stay within a carbon consumption budget. Normally, to meet an end-to-end service level agreement (SLA), a path-pinning mechanism could pro-actively perform a hop-by-hop accumulation of a route's aggregate estimated latency, and assess its ability to stay below the requested envelope. To guarantee the SLA typically requires that all nodes along the path reserve buffers to avoid any congestion and thus minimize latency. For carbon-aware traffic engineering, instead the reservations are for clean energy resources (predicted to be generated in real-time or stored in battery), to deliver a minimum carbon-intensity path.

Carbon-aware telemetry. To support carbon-awareness more broadly, 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 payload. 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 piece 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 [6] 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.

Carbon-aware applications. Because applications are the most in the know about their own requirements and contextual details, they mustn't be shoehorned by the lower layers (or by other points in the architectural stack) into using functionality that might best be performed at the application layer. In fact there should be many entry points (e.g., APIs) in the software stack that provide insight into carbon-awareness, which applications or their users can select amongst.

Technical Challenges

To properly account for the carbon efficiency of networking requires an end-to-end approach. 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 discover and collate the energy usage and carbon efficiency of network paths, (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 technical challenges to realizing these aims.

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. 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. However, unlike some forms of in-network telemetry, the gathering electricity consumption metrics of a path cannot be measured by the telemetry mechanism itself, but must be gathered from data available along the path, introducing problems of authenticity and accuracy.

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 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 [4] [12][15][18]. While these are already being used by cloud service providers, network operators have yet to embrace this knowledge operationally. 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.

Discovering and collating energy usage and carbon efficiency of network paths. Fundamentally the problem is that energy consumption of the network infrastructure, and in turn its carbon footprint, is not something that routing protocols have full insight into at present. Nor do the transport or application layer protocols. This is a function of the prevalence of middleboxes that live along a route, and that would be largely hidden or omitted from router-level computations. Essentially, we do not have the tools as yet to measure inside the network what we need. In the past, tools such as traceroute provided per host telemetry data and reachability status, as well as cumulative pathway features (latency and hop count). With the proliferation of NATs, VPNs, multi-path, and middleboxes, traceroute is no longer a reliable forensic tool: a particular IP path may be overlaid on top of other topologies and lower link-layers, and therefore a simple hop-by-hop accumulation of data may miss out on relevant information hidden by layers of abstraction in the network. What would it take to create a carbon-aware traceroute, one that accurately assesses the carbon efficiency of a path between a given source and destination? If we could not derive a path's precise measurements, what about being able to guarantee at least an estimate of the order of magnitude of environmental impact? Moreover, how are measurements impacted by the supporting cast of network hardware equipment at layer 2, in addition to those at layer 3? How does billing and network management issues (leasing and ownership) further complicate carbon-awareness, as these too are hidden costs?

Reacting in (near) real-time to carbon-related information collected from the network. 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, 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, tie it to an application, or even packet level, using an in-network telemetry solution [7]. To turn the information into a working, useful solution, it is assumed that the 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. While we outline an intra-AS strategy, to scale up, solutions likely will need to extend and to simplify inter-AS techniques that leverage trustworthy peering relationships (such as [14]), sharing carbon-efficiency metrics between ASes.

Conclusion and Call to Action

Networking needs to become carbon-aware, like any other part of the ICT digital infrastructure. Its disaggregated and distributed nature makes the accounting harder, but possible. Carbon-awareness in networking promises to fundamentally change the way networks operate. It will play across multiple layers of the network stack: from the physical and routing layers to the transport and application layers. Having visibility in each of these layers into the environmental consequences of data transmission means that better decisions can be made: applications and services that consume less carbon resources will be preferred, and routes will be taken where the carbon footprint is minimized. To achieve concrete progress, however, standard environmental metrics need to be defined; although we have focused on carbon, others have provided additional candidate metrics [20] for consideration. We call on the community to join the effort to define these metrics.

Designing carbon-intelligent routing solutions is the next big challenge in networking. The proposed facets described here are early building blocks for a more holistic carbon-aware solution. New algorithms are needed, as is collaboration with the electrical grid community to standardize carbon-intensity data and to build an understanding of how to balance energy consumption with decarbonization. We must attend to the dynamic nature of changes in efficiency, and provide accountability and interpretability. Together, we can make networking truly green.

References

[1] AGARWAL, A., SUN, J., NOGHABI, S., IYENGAR, S., BADAM, A., CHANDRA, R., SESHAN, S., AND KALYANARAMAN, S. , Redesigning data centers for renewable energy, HotNets 2021, 45–52.

[2] CALIFORNIA INDEPENDENT SYSTEM OPERATOR, Managing Oversupply. Oneline, 2022. http://www.caiso.com/informed/Pages/ManagingOversupply.aspx [Accessed: October 2022]

[3] FINN, N., THUBERT, P., VARGA, B., FARKAS, J., Deterministic Networking Architecture, RFC 8655, October 2019, https://www.rfc-editor.org/rfc/rfc8655

[4] ELECTRICITY MAP. Electricity Map API. Online, 2022. https://static.electricitymaps.org/api/docs/index.htm. [Accessed: October 2022]

[5] KONINGSTEIN, R., We now do more computing where there's cleaner energy. Online, 2021. https://blog.google/outreach-initiatives/sustainability/carbon-aware-computing-location/ [Accessed: October 2022]

[6] GOOGLE CARBON-FREE ENERGY BLOG. Operating on 24/7 Carbon-Free Energy by 2030. Online, 2022. https://sustainability.google/progress/energy/ [Accessed: October 2022]

[7] KIM, C., SIVARAMAN, A., KATTA, N., BAS, A., DIXIT, A., ANDWOBKER, L. J. In-band network telemetry via programmable dataplanes. In *ACM SIGCOMM* (2015), vol. 15.

[8] LORINCZ, J., CAPONE, A., AND WU, J. Greener, energy-efficient and sustainable networks: State-of-the-art and new trends. *Sensors 19*, 22 (2019).

[9] JAMES, A., SCHIEN, D., A Low Carbon Kubernetes Scheduler, *CEUR Workshop*, vol. 2382 (Jan 2019)

[10] NAFUS, D., SCHOOLER, E. M., AND BURCH, K. A., Carbon responsive computing: Changing the nexus between energy and computing. *Energies 14,* 21 (2021).

[11] RADOVANOVIC, A., KONINGSTEIN, R., SCHNEIDER, I., CHEN, B., DUARTE, A., ROY, B., XIAO, D., HARIDASAN, M., HUNG, P., CARE, N., ET AL. Carbon-aware computing for datacenters. *arXiv preprint arXiv:2106.11750* (2021).

[12] SINGULARITY. Singularity API. Online, 2022. https://singluarity-docs.stoplight.io [Accessed June 2022].

[13] SCHOOLER, E.M., TAYLOR, R., Non-Traditional Network Metrics, IAB 2021 Workshop on Measuring Network Quality for End-Users (Sept 2021).

[14] SCION ARCHITECTURE. Online, 2022. https://scion-architecture.net/ [Accessed: Oct 2022]

[15] Footprints on the path: how routing data could reduce the internet's carbon toll, World Economic Forum, Digital Communications (Mar 2021). https://www.weforum.org/agenda/2021/03/internet-carbon-emissions-data-path-scion [Accessed: October 2022]

[16] UK NATIONAL GRID ESO. National Grid ESO Carbon Intensity API. Online, 2022. https://carbonintensity.org.uk/ [Accessed: October 2022]

[17] VAN DER MEULEN, R. 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-leaders [Accessed: October 2022]

[18] WATTTIME. WattTime API. Online, 2022. https://www.watttime.org/api-documentation/ [Accessed: October 2022].

[19] ZILBERMAN, N., SCHOOLER, E.M., CUMMINGS, U., MANOHAR, R., NAFUS, D., SOULE, R., TAYLOR, R., Toward Carbon-Aware Networking, *HotCarbon 22* (June 2022).

[20] 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).