From Component to System: Rethinking Edge Computing Design Through a Carbon-Aware Lens

XUESI CHEN, Cornell Tech, USA
ARIEL GOLDNER, Columbia University, USA
EREN YILDIZ, Georgia Institute of Technology, USA
ILAN MANDEL, Cornell Tech, USA
TINGYU CHENG, Georgia Institute of Technology, USA
JOSIAH HESTER, Georgia Institute of Technology, USA
UDIT GUPTA, Cornell Tech, USA

As edge devices see increasing adoption across a wide range of applications, understanding their environmental impact has become increasingly urgent. Unlike cloud systems, edge deployments consist of tightly integrated microcontrollers, sensors, and energy sources that collectively shape their carbon footprint. In this paper, we present a carbon-aware design framework tailored to embedded edge systems. We analyze the embodied emissions of several off-the-shelf microcontroller boards and peripheral components and examine how deployment context—such as workload type, power source, and usage duration—alters the carbon-optimal configuration. Through empirical case studies comparing battery- and solar-powered scenarios, we find that the lowest-emission choice is often workload- and context-specific, challenging assumptions that energy-efficient or renewable powered systems are always the most sustainable. Our results highlight the need for fine-grained, system-level reasoning when designing for sustainability at the edge and provide actionable insights for researchers and practitioners seeking to reduce the carbon cost of future deployments.

CCS Concepts: • **Hardware → Impact on the environment**; • **Computer systems organization → Embedded hardware**.

Additional Key Words and Phrases: Sustainable computing, embedded system, life cycle analysis, energy

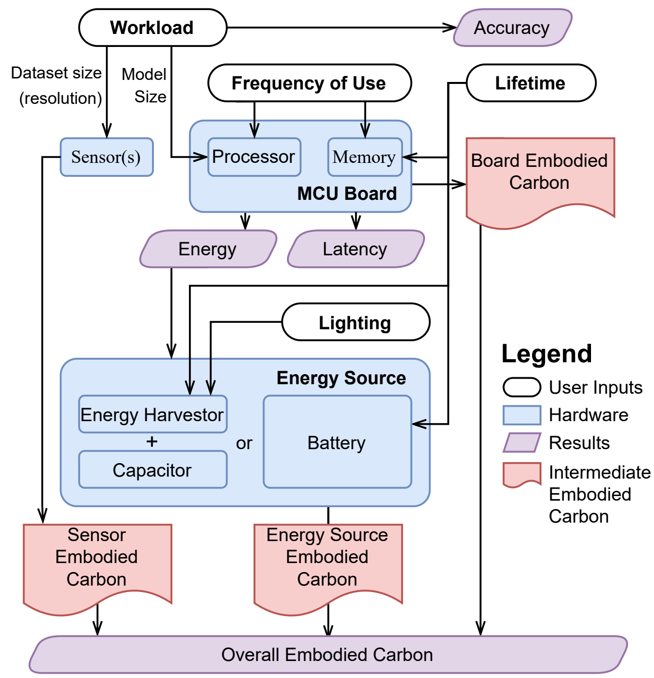

Fig. 1. Carbon-aware edge computing system design. Carbon emissions arise not only from the choice of sensors and microcontrollers but also from how systems are deployed and powered. Energy efficiency alone does not ensure lower carbon emissions, underscoring the need for a holistic and quantified design framework that considers hardware selection, workload characteristics, and provisioning strategy.

1 INTRODUCTION

Fueled by advancements in artificial intelligence and hardware efficiency, edge devices are becoming increasingly ubiquitous, appearing in a wide range of applications from basic air quality monitoring [5] to smartphones and wearables. By the end of 2023, there were 16.6 billion connected IoT devices worldwide, a 15% increase from 2022 [2]. Collectively, edge devices account for approximately one-third of emissions from the Information and Communication Technology (ICT) sector, which itself contributes around 4% of global carbon emissions [10].

Despite this massive and growing deployment, the carbon footprint of edge computing remains overgeneralized and understudied. Much of the existing sustainability work has centered on datacenters and general-purpose compute systems [10, 13, 14, 27], with relatively little attention paid to the embodied and operational carbon emissions of embedded edge devices. Existing life-cycle assessment (LCA) studies often focus narrowly on processors and memory [4, 13, 18, 33], overlooking the tightly coupled nature of edge systems that integrate sensors, energy harvesting modules, and custom power management [16, 26].

In this paper, we present a comprehensive carbon-aware design framework for edge systems to address this gap, shown in Figure 1. We begin by modeling the embodied carbon of a range of off-the-shelf microcontroller boards and peripheral components. Our cradle-to-gate analysis reveals significant variation in embodied emissions—ranging from 0.52 kg CO_2e for the Raspberry Pi Pico to 2.59 kg CO_2e for the Coral Dev Board Micro, with integrated circuits as primary carbon footprint contributors in compute-heavy devices. In contrast, for compute-lean MCUs, components such as

Authors' addresses: Xuesi Chen, xc562@cornell.edu, Cornell Tech, USA; Ariel Goldner, ag4967@columbia.edu, Columbia University, USA; Eren Yildiz, eyildiz8@gatech.edu, Georgia Institute of Technology, USA; Ilan Mandel, im334@cornell.edu, Cornell Tech, USA; Tingyu Cheng, tcheng32@gatech.edu, Georgia Institute of Technology, USA; Josiah Hester, josiah@gatech.edu, Georgia Institute of Technology, USA; Udit Gupta, ugupta@cornell.edu, Cornell Tech, USA.

sensors and passive elements contribute a substantial share of the total carbon footprint.

We then conduct a detailed case study comparing three microcontrollers (Arduino Nano 33 BLE Sense, Coral Dev Board Micro, and Raspberry Pi Pico) under realistic workloads and deployment conditions. We evaluate two energy supplying scenarios: solar energy harvesting and AA battery power. Our results show that energy efficiency alone does not guarantee carbon optimality. For example, while the Coral Dev Board Micro is the most energy-efficient for a keyword spotting workload, its higher embodied carbon makes it suboptimal for short-term battery-powered deployments, where the Arduino Nano yields lower lifetime emissions. Conversely, under long-term use or solar provisioning, energy-efficient devices can amortize their embodied cost and become the more sustainable choice.

These findings challenge the conventional wisdom that systems powered by renewable sources are always more sustainable, or that energy-efficient compute inherently results in lower carbon emissions. We demonstrate that the carbon-optimal edge configuration is highly sensitive to workload, deployment time, and power source. As such, carbon-aware edge system design demands holistic analysis and context-specific provisioning. Our work serves as a blueprint for researchers, designers, and manufacturers seeking to align embedded computing with sustainability goals.

2 BACKGROUND AND RELATED WORK

Carbon emissions of computing can be categorized into two types: operational and embodied. Operational carbon refers to the carbon emission that occurred during hardware use: the combination of energy consumption and carbon intensity of the energy source that powers the device. Embodied carbon refers to the carbon emission that occurred during the chip manufacturing process [13, 14]

While several prior works [13, 14, 18, 27, 33] have LCA of computing devices, most focus on modeling compute and memory chips. However, edge computing has its distinct paradigm, characterized by tightly integrated sensors, communication modules, and energy-constrained power sources and necessitates its own sustainability modeling and carbon-aware design strategies.

Carbon Modeling for Edge Computing Despite the increasing deployment of edge devices, sustainability work specific to IoT and intermittently powered systems remains limited. A notable exception is the bottom-up LCA framework proposed in [26], which categorizes emissions from IoT device components (processors, sensors, actuators, etc.) and estimates cradle-to-gate emissions for consumer devices such as the Google Home Mini and Apple Watch. However, their approach generalizes hardware specifications and relies heavily on size- and quantity-based estimations, often overlooking critical technology-level variations that significantly influence carbon emissions, especially in compute, memory, and sensing categories.

Some prior work has studied the LCA of specific electronic components, such as CMOS image sensors [1], processors [4, 13, 18], memory [33, 37], and capacitors [28]. Fairphone [7], on the other hand, discloses detailed LCA data for an entire cellphone product, serving as a valuable reference for edge computing LCAs. Since calculating carbon emissions for all electronic components is not always straightforward, DeltaLCA [39] adopts a comparative approach, avoiding absolute metrics by analyzing trade-offs between different electronic designs.

Last but not least, prior modeling efforts have largely focused on the cradle-to-gate embodied carbon emissions of IoT devices, often neglecting how application workloads and usage patterns (e.g., device lifetime, duty cycle) can influence onboard resource provisioning—and ultimately, the device's embodied emissions. Our work addresses this gap by integrating embodied carbon modeling with system design. We consider user requirements early in the design phase and provision onboard components accordingly, enabling more carbon-aware edge system design.

Sustainability for Edge Computing While sustainable edge computing is still an emerging area, several studies have explored the carbon trade-offs involved in intermittent computing and the sensing pipeline from edge to cloud. CO_2CoDe[25] introduces a carbon-aware hardware–software co-design framework for intermittent computing, demonstrating how choices in energy storage can influence scheduling decisions. However, this work is limited to a fixed PCB size and a specific processor with a predefined carbon footprint, and the evaluation only accounts for the carbon emissions of the capacitors used. Desai et al. [16] has examined the carbon impact of the full sensing pipeline—from the sensing device to the cloud—analyzing how different energy harvesting approaches and data storage methods affect the overall carbon footprint of ML-enabled devices. Despite these efforts, conclusions in prior studies are typically drawn from highly-specific workloads, hardware platforms, and sensing use cases. In contrast, our work aims to provide generalizable design guidance for commonly available off-the-shelf boards and representative workloads.

3 CRADLE-TO-GATE EMBODIED CARBON FOR OFF-THE-SHELF EMBEDDED SYSTEM COMPONENTS

To understand the carbon impact of edge devices and enable carbon-aware design, we first quantify the emissions of common edge computing components, including processors, sensors, and their power supplies (e.g., batteries, capacitors, and energy harvesters, etc.). This section presents a cradle-to-gate analysis of off-the-shelf embedded components, encompassing raw material extraction, manufacturing, and, where available, transportation.

We begin by detailing how to estimate emissions for high confidence components—such as SoCs, memory, image sensors, resistors, and solar panels—based on prior literature [1, 13, 29, 37]. These estimates are specific to manufacturing processes and technologies but often lack transportation data. For the remaining components, we use ecoinvent [11], which includes transportation but relies on broader categorizations.

To evaluate the embodied carbon of off-the-shelf microcontrollers, we extract the Bill of Materials (BoM) from public disclosures and schematics. Each BoM is categorized by component type, and the total carbon footprint is computed accordingly. The five off-the-shelf microcontrollers and microprocessors we evaluate in this paper are Coral Dev Board Micro, Raspberry Pi Pico, MAX780000 FTHR, Raspberry Pi model 3 A+ and Arduino Nano 33 BLE Sense.

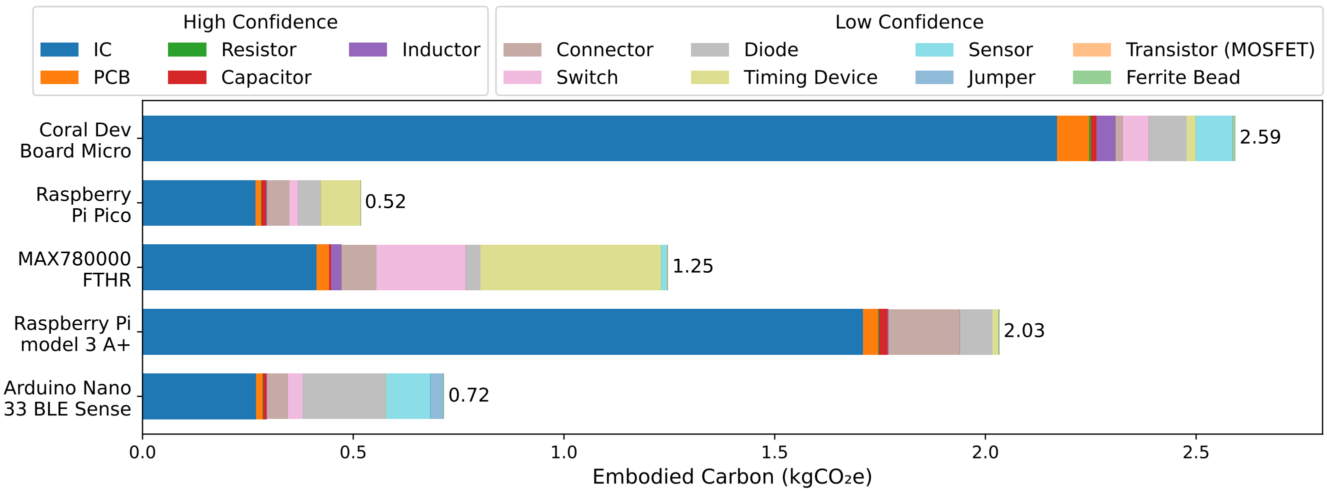

Fig. 2. Cradle-to-gate embodied carbon emissions of common off-the-shelf microcontrollers breaking into categories. Components estimated using well-documented literature and process-specific data sources are marked as high confidence and shown with hatching, while components estimated using emission factors from ecoinvent [11] are marked as low confidence.

Image Sensor	**Resolution**	**Embodied Carbon (g)**
HM5040	2592 × 1944	88.4
HM2140	1920 × 1080	76.0
HM1091	1280 × 720	20.4
OV2640	1632 × 1232	81.0
OV7670	656 × 488	42.8

Table 1. Embodied carbon of common off-the-shelf CMOS image sensors

3.1 High-confidence Components

Integrated Circuit For all Integrated Circuit (IC) embodied emissions calculations, we adopt the methodology described in ACT [13]. The embodied carbon of a given System-On-Chip (SoC) is estimated as the product of the carbon per area and the area of the IC. Carbon per area of the given IC is dependent on the technology node of the chip and fabrication location. We relied on ACT [13] and Life-cycle Assessment of Semiconductors [4] to generate the embodied carbon emission per area for each technology node for each corresponding manufacturer. To refer the die area for each chip, we rely on the package-to-die ratio published in DeltaLCA [39] and Lall, et al [19]. For memory chips, we estimate embodied carbon based on memory size and type, using data disclosed in prior works [13, 33, 37].

CMOS Image Sensors The embodied carbon of a CMOS image sensor depends on its pixel size, resolution, manufacturing location, and technology node. We estimate silicon area using Equation 1, which multiplies pixel dimensions by resolution and adds a 0.3 mm margin for connectivity [1]. Since datasheets often omit technology node details, we approximate it using a 20:1 pixel-to-node ratio [35]. While some advanced sensors use dual-layer transistor pixel technology [24], most IoT-class sensors follow a conventional stacked architecture, making them compatible with our area-based estimation model. We present some of the common off-the-shelf image sensors in Table 1.

$$\hat{S} = (Res_X \cdot Pix_W + 0.3\,\text{mm}) \cdot (Res_Y \cdot Pix_H + 0.3\,\text{mm}) \qquad (1)$$

Capacitors, Resistors, Inductors We use a combination of weight-based and package-based analyses to estimate the environmental impact of capacitors, resistors, and inductors. For package-based assessments, we rely on the component-level data disclosed in DeltaLCA [39]. When capacitor package data is unavailable, we apply a weight-based approach using emission factors specific to different capacitor types: multi-layer ceramic capacitors (MLCC), tantalum electrolytic capacitors (TEC) [28], and aluminum electrolytic capacitors [38]. For resistors and inductors not covered by DeltaLCA [39]'s package list, we use the generic emission factors for resistors and inductors provided by ecoinvent [11].

PCB FR-4 is a common PCB substrate widely used in edge devices. We estimate the carbon emissions of PCBs by multiplying their area by an emission factor of 0.006125 grams of CO_2 equivalent per mm^2 per 1mm-thick layer [9, 22, 39]. The number of PCB layers is determined based on the component schematics.

3.2 Low-Confidence Components Based on ecoinvent Data

While conducting the LCA for microcontrollers, we observed that not all components are thoroughly documented or studied using up-to-date methodologies and data. Several components lack detailed academic assessments of their emissions. These include: connectors, switches, diodes, timing devices, jumpers, single surface mount transistors, thermistors, ferrite beads, and microphones.

Emissions factors for these components are available from private databases used in industrial LCAs. In this study, we accessed Sustainable Minds [31], which leverages the EcoInvent Version 2.0 database [11]. While some components had exact category matches in the database, others required approximation through broader or similar classifications. We estimated the emissions of each component by multiplying its weight by the corresponding emissions factor.

Boards	CPU	ML accelerator	RAM	Flash	IDE	Cost ($)
Coral Dev Board Micro	NXP i.MX RT1176 MCU (ARM Cortex-M7 @ 800 MHz, ARM Cortex-M4 @ 400 MHz)	Google Edge TPU coprocessor: 4 TOPS (int8);	64MB	128MiB	FreeRTOS	79.99
Arduino Nano 33 BLE Sense	nRF52840 (ARM Cortex-M4 with FPU @ 64MHz)	N/A	256KB	1MB	Arduino	38.7
Raspberry Pi Pico	Dual-core Arm Cortex-M0 @ 133 MHz	N/A	264KB	2MB	Pico SDK	4

Table 2. List of development boards used in the experiment with processor and memory information

Workload	Dataset	Model	Size
KWS-S [34]	Speech Command [36]	CNN	117KB
KWS-L [3]	Speech Command [36]	DS-CNN	384KB
PD [34]	VWW Dataset [21]	MobileNet	294KB

Table 3. Inference workloads used in our experiments, including person detection (PD) and keyword spotting models—small (KWS-S) and large (KWS-L).

3.3 LCA Results and Discussion

The embodied carbon of the five off-the-shelf microcontrollers and microprocessors we examined ranges from 0.53 kg CO_2e to 2.59 kg CO_2e. A detailed breakdown of total embodied carbon emissions is shown in Figure 2, where shaded regions indicate components estimated with high confidence. Integrated circuits are the dominant contributors to embodied carbon across most devices. However, in compute-lean microcontrollers—such as the Raspberry Pi Pico, MAX78000 FTHR, and Arduino Nano 33 BLE Sense—the integrated circuit accounts for less than 50% of total embodied emissions. This is partly due to the substantial contribution of timing devices and diodes, which may be overestimated due to the limitations of the ecoinvent [11] database. Specifically, timing devices are categorized as active components, sharing the same emission factor as integrated circuits and displays, which likely inflates their estimated carbon impact.

While some microcontrollers, such as the Arduino Nano 33 BLE Sense, include onboard sensors, others, like the Raspberry Pi Pico, require external sensor modules. As a result, the embodied carbon of devices like the Raspberry Pi Pico can increase depending on the sensor type and deployment context. We account for this variability in our case study in section 4.

4 CARBON-AWARE IOT DESIGN CASE STUDY

In this section, we examine the decision-making process for selecting between off-the-shelf microcontroller (MCU) boards, listed in Table 2 for always-on inference with sustainability in mind. We consider two scenarios: one where the edge system is powered by solar panels paired with capacitors, and another where it is powered by AA batteries.

Experimental Setup: We use the DS-CNN keyword spotting model from the MLPerf Tiny v1.1 benchmark [3], trained on the Speech Commands dataset with 90% accuracy, as well as the micro speech and person detection model provided by Tensorflow Lite Micro [6], as summarized in Table 3. Runtime and per-inference energy are measured using Nordic Semiconductor Power Profiler kit II [23] and X-NUCLEO-LPM01A [30]. Results are shown in Figure 3.

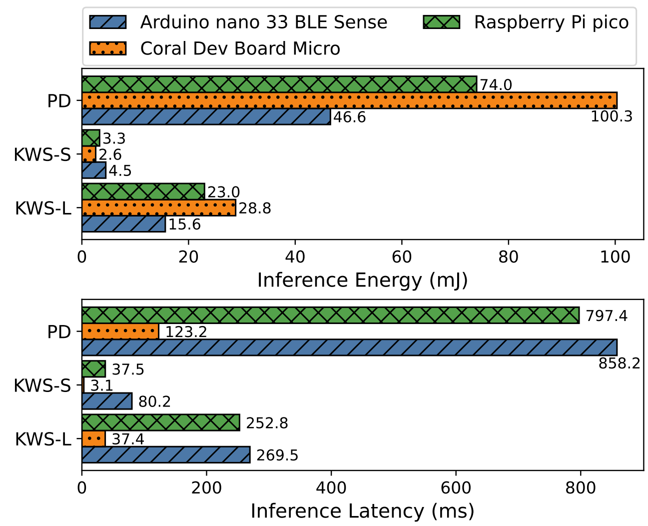

Fig. 3. Per-inference energy and latency trade-offs for three MCUs running three edge ML workloads. The Coral Dev Board Micro achieves significantly lower inference latency, with varying energy trade-offs depending on the workload.

To ensure a fair comparison, we normalize the number of inferences executed per day across all devices. For example, for MLPerf tiny's keyword spotting (KWS) workloads, we use the Arduino Nano running at a 100% duty cycle as the baseline. Due to its faster runtime, the Coral Dev Board Micro operates at a 14.7% duty cycle, while the Raspberry Pi Pico runs at 93.9% for the same model. Based on the measured energy consumption and runtime per inference, we size the required capacitors and energy harvesters for each board accordingly.

Energy Harvester–Powered Case: We estimate the required area for solar panels based on real irradiance power traces (in $\mu W/cm^2$) sourced from the EnHANTs dataset [12]. We assume a panel efficiency of 16.65% [20], reflecting the performance of typical polycrystalline solar panels. For each device, we calculate the average power needed over the course of a day and divide this value by the average daily irradiance (adjusted for efficiency) to determine the minimum panel area required for sustained operation. To convert the solar panel area to its embodied carbon, we apply an emission factor of 0.0227 kg CO_2e per cm^2, based on values reported in a systematic review of life cycle assessments of crystalline silicon photovoltaic electricity generation [17].

In our evaluation, we utilized the following real-world setups from the EnHANTs dataset [12]:

(1) *High-light setup (mobile-indoor-outdoor)*: A pedestrian walked around the university campus, which included both indoor and outdoor environments, while carrying a sensor.
(2) *Low-light setup (mobile-car)*: A car-based road trip where the sensor was attached to the dashboard.

We pair microcontrollers with different numbers of $470\mu F$ MLCC capacitors based on the energy needed to do a single inference. A voltage regulator is added to the Coral Dev Board Micro setup to boost the supply voltage to 5V as required.

The breakdown of carbon emissions of the energy harvesters and microcontroller boards is shown in Figure 4. The carbon emission of capacitors and voltage regulators is negligible in comparison to energy harvesters and microcontrollers.

While energy efficiency and runtime are often used as primary metrics for selecting microcontrollers, our analysis shows that these do not always align with carbon-optimal design. As illustrated in Figure 4, the carbon-optimal edge system configuration depends not only on the hardware and energy source, but also on the specific workload and deployment environment. For example, the Arduino Nano 33 BLE Sense consistently yields the lowest embodied carbon for both the kws MLPerf tiny and person detection workloads across high- and low-light conditions due to its low board-level emissions. However, for the kws tflm workload, the Raspberry Pi Pico is optimal under high-light conditions, while the Coral Dev Board Micro becomes preferable in low-light settings because its energy efficiency offsets the higher board carbon through a smaller solar panel footprint. These findings highlight that even with the same workload and set of devices, the lowest-carbon solution can vary based on deployment context. Therefore, carbon-aware edge system design must account for workload characteristics, environmental conditions, and energy provisioning strategy—not just energy efficiency alone.

Battery-powered Case: We construct the battery-powered scenario by modeling always-on inference systems powered by AA batteries. In our calculation, each battery provides 11,250 J of energy [8] and carries an embodied carbon footprint of 0.107 kg CO_2e [15]. Based on the measured per-inference energy consumption of each device and its adjusted duty cycle, we estimate the number of batteries required to sustain operation over time and calculate the resulting lifetime carbon emission.

As shown in Figure 5, the slope of each line corresponds to the device's energy efficiency. The Coral Dev Board Micro is the most energy-efficient among the three devices when running kws tflm; although it starts with relatively high embodied carbon, its operational efficiency results in the lowest lifetime emissions over time—surpassing the Arduino Nano 33 BLE Sense after day 97 and the Raspberry Pi Pico after day 270. Notably, the relative energy efficiency of the same set of devices can shift depending on the workload. For instance, while the Arduino Nano 33 BLE Sense is the least efficient for kws tflm, it becomes the most energy-efficient option when running kws MLPerf tiny and pd tflm. These results underscore that carbon-optimal hardware choices are not fixed but depend on the interplay between workload characteristics and energy consumption.

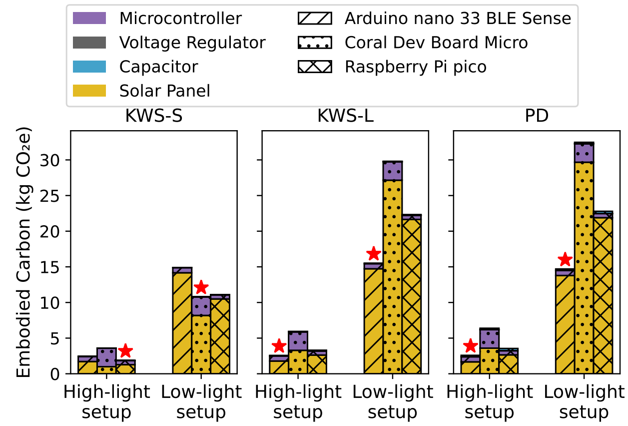

Fig. 4. Carbon emissions of three always-on inference workloads deployed on three MCU boards under varying light conditions. The carbon-optimal solution varies by workload and environment, highlighting the need to jointly consider hardware choice, deployment strategy, and operating conditions in carbon-aware edge system design.

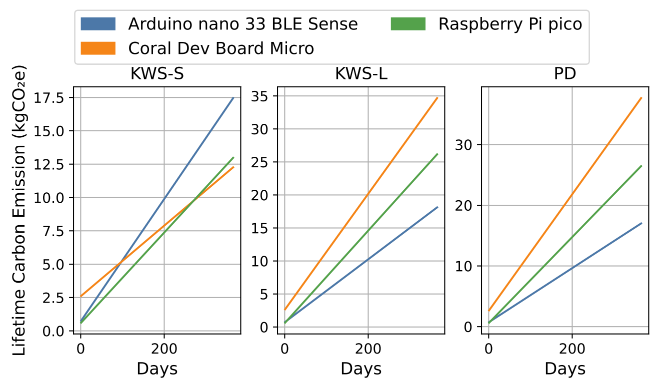

Fig. 5. Carbon emissions of devices running always-on inference with AA batteries over time. For kws tflm, the Coral Dev Board Micro starts with higher embodied carbon, its superior energy efficiency eventually results in the lowest lifetime emissions, highlighting that the most carbon-efficient device can vary depending on the workload and deployment time.

Deployment Time and Settings: Combining the results from both the battery-powered and solar-harvested scenarios, we emphasize that the deployment strategy leading to the lowest carbon emissions is highly use-case dependent. For instance, if a user intends to deploy a person detection model inside a car (a low-light environment) for three months as a temporary security solution, powering the system with batteries results in lower overall carbon emissions than using solar panels. This finding may seem counterintuitive, as renewable energy is often assumed to be the more sustainable choice by default.

These insights highlight the importance of moving beyond one-size-fits-all assumptions. We advocate for a carbon-aware edge system design approach that accounts for deployment context, operational duration, and power provisioning to enable more sustainable decisions.

5 CONCLUSION

This work presents a holistic framework for carbon-aware edge system design, grounded in detailed embodied carbon modeling and deployment-sensitive analysis. We demonstrate that sustainability in embedded systems cannot be evaluated solely based on energy efficiency or compute footprint. Through a case study of off-the-shelf microcontrollers in both solar- and battery-powered scenarios, we show that carbon-optimal decisions are highly dependent on workload characteristics, power source, and deployment duration. Notably, the most energy-efficient or lowest-emission device in isolation may not be the most carbon-optimal choice in specific short-term or constrained-use deployments. These findings highlight the importance of system-level, context-aware design strategies that account for both embodied and operational emissions.

6 LIMITATIONS AND FUTURE WORKS

While our analysis focuses on a subset of representative microcontrollers and ML inference workloads, extending this framework to a wider range of edge devices, sensing modalities, and workloads beyond machine learning would enable broader applicability. In particular, incorporating event-driven workloads and heterogeneous sensor integrations could uncover additional carbon trade-offs across the embedded system design space.

We also aim to refine embodied carbon estimates for components currently modeled with low-confidence data, such as timing devices and diodes, by collaborating with industry partners to obtain more accurate BoM details and leveraging LCA expertise in academic labs. In addition, extending our framework to support end-of-life modeling and circularity scenarios, such as Junkyard Computing [32], could enable more holistic and forward-looking strategies for carbon-aware edge computing system design.

7 ACKOWNLEDGEMENT

We thank Jenna Trost for discussions on the LCA aspects of this work. We also thank the anonymous reviewers, Zhanqiu Hu, and Zev Pogrebin for their valuable feedback. This research was supported by the National Science Foundation under award numbers CCF-2324860 and CNS-2335795.

REFERENCES

[1] 2024. *Streamlined Models of CMOS Image Sensors Carbon Impacts.*

[2] IoT Analytics. 2024. *State of IoT Summer 2024.* Technical Report. IoT Analytics.

[3] Colby Banbury, Vijay Janapa Reddi, Peter Torelli, Jeremy Holleman, Nat Jeffries, Csaba Kiraly, Pietro Montino, David Kanter, Sebastian Ahmed, Danilo Pau, Urmish Thakker, Antonio Torrini, Peter Warden, Jay Cordaro, Giuseppe Di Guglielmo, Javier Duarte, Stephen Gibellini, Videet Parekh, Honson Tran, Nhan Tran, Niu Wenxu, and Xu Xuesong. 2021. MLPerf Tiny Benchmark. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks* (2021). Git tag v1.1 (released Jun 27, 2023). Code: https://github.com/mlcommons/tiny/tree/v1.1.

[4] Sarah B. Boyd. 2009. *Life-cycle Assessment of Semiconductors.* Ph. D. Dissertation. UNIVERSITY OF CALIFORNIA, BERKELEY.

[5] Madeleine I. G. Daepp, Alex Cabral, Vaishnavi Ranganathan, Vikram Iyer, Scott Counts, Paul Johns, Asta Roseway, Charlie Catlett, Gavin Jancke, Darren Gehring, Chuck Needham, Curtis von Veh, Tracy Tran, Lex Story, Gabriele D'Amone, and Bichlien H Nguyen. 2022. Eclipse: An End-to-End Platform for Low-Cost, Hyperlocal Environmental Sensing in Cities. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN).* 28–40. https://doi.org/10.1109/IPSN54338.2022.00010

[6] Robert David, Jared Duke, Advait Jain, Vijay Janapa Reddi, Nat Jeffries, Jian Li, Nick Kreeger, Ian Nappier, Meghna Natraj, Shlomi Regev, Rocky Rhodes, Tiezhen Wang, and Pete Warden. 2021. TensorFlow Lite Micro: Embedded Machine Learning on TinyML Systems. arXiv:2010.08678 [cs.LG] https://arxiv.org/abs/2010.08678

[7] Lara Eguren David Sánchez, Sarah-Jane Baur. 2024. *Life Cycle Assessment of the Fairphone 5.* Technical Report. Berlin: Fraunhofer IZM.

[8] Energizer Holdings, Inc. 2022. Energizer E91 AA Alkaline Battery Technical Datasheet. https://data.energizer.com/pdfs/e91.pdf. Accessed: 2025-05-19.

[9] Csaba Farkas, László Galá, András Csiszár, Vincent Grennerat, Pierre-Olivier Jeannin, Pascal Xavier, Dániel Rigler, Olivér Krammer, Zbyněk Plachý, Karel Dušek, Róbert Kovács, Anna Éva Fehér, and Attila Géczy. 2024. Sustainable printed circuit board substrates based on flame-retarded PLA/flax composites to reduce environmental load of electronics: Quality, reliability, degradation and application tests. *Sustainable Materials and Technologies* 40 (2024). https://doi.org/10.1016/j.susmat.2024.e00902

[10] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon Blair, and Adrian Friday. 2021. The climate impact of ICT: A review of estimates, trends and regulations. arXiv:2102.02622 [physics.soc-ph] https://arxiv.org/abs/2102.02622

[11] R. Frischknecht, N. Jungbluth, H.-J. Althaus, R. Hischier, G. Doka, R. Dones, T. Heck, S. Hellweg, G. Wernet, T. Nemecek, and et al. 2007. Overview and methodology. Data v2.0 (2007). Ecoinvent report No. 1. (Dec 2007).

[12] Maria Gorlatova, Michael Zapas, Enlin Xu, Matthias Bahlke, Ioannis (John) Kymissis, and Gil Zussman. 2022. CRAWDAD columbia/enhants. https://doi.org/10.15783/C77P47

[13] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: Designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture.* 784–799.

[14] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2021. Chasing Carbon: The Elusive Environmental Footprint of Computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA).* 854–867. https://doi.org/10.1109/HPCA51647.2021.00076

[15] Ramsey Hamade, Raghid Al Ayache, Makram Bou Ghanem, Sleiman El Masri, and Ali Ammouri. 2020. Life Cycle Analysis of AA Alkaline Batteries. *Procedia Manufacturing* 43 (2020), 415–422. https://doi.org/10.1016/j.promfg.2020.02.193 Sustainable Manufacturing - Hand in Hand to Sustainability on Globe: Proceedings of the 17th Global Conference on Sustainable Manufacturing.

[16] Nathan Beckmann Brandon Lucia Harah Desai, Sara McAllister. 2024. Towards Understanding the Carbon Impact in End-to-end Sensing Pipelines. In *HotEthics.*

[17] David D. Hsu, Patrick O'Donoughue, Vasilis Fthenakis, Garvin A. Heath, Hyung Chul Kim, Pamala Sawyer, Jun-Ki Choi, and Damon E. Turney. 2012. Life Cycle Greenhouse Gas Emissions of Crystalline Silicon Photovoltaic Electricity Generation. *Journal of Industrial Ecology* 16, s1 (2012), S122–S135. https://doi.org/10.1111/j.1530-9290.2011.00439.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1530-9290.2011.00439.x

[18] Shixin Ji, Zhuoping Yang, Xingzhen Chen, Stephen Cahoon, Jingtong Hu, Yiyu Shi, Alex K Jones, and Peipei Zhou. 2024. SCARIF: Towards Carbon Modeling of Cloud Servers with Accelerators. In *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI).* IEEE, 496–501.

[19] Pradeep Lall, Naveen Singh, Jeffrey Suhling, Marcia Strickland, and Jim Blanche. 2005. Thermo-mechanical reliability tradeoffs for deployment of area array packages in Harsh environments. *Components and Packaging Technologies, IEEE Transactions on* 28 (10 2005), 457 – 466. https://doi.org/10.1109/TCAPT.2005.854162

[20] Opeyeolu Timothy Laseinde and Moyahabo Dominic Ramere. 2021. Efficiency Improvement in polycrystalline solar panel using thermal control water spraying cooling. *Procedia Computer Science* 180 (2021), 239–248. https://doi.org/10.1016/j.procs.2021.01.161 Proceedings of the 2nd International Conference on Industry 4.0 and Smart Manufacturing (ISM 2020).

[21] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. Microsoft COCO: Common Objects in Context. arXiv:1405.0312 [cs.CV] https://arxiv.org/abs/1405.0312

[22] Jingping Liu, Cheng Yang, Haoyi Wu, Ziyin Lin, Zhexu Zhang, Ronghe Wang, Baohua Li, Feiyu Kang, Lei Shi, and Ching Ping Wong. 2014. Future paper based printed circuit boards for green electronics: fabrication and life cycle assessment. *Energy Environ. Sci.* 7 (2014), 3674–3682. Issue 11. https://doi.org/10.1039/C4EE01995D

[23] Nordic Semiconductor. 2020. Power Profiler Kit II (PPK2). https://www.nordicsemi.com/Products/Development-hardware/Power-Profiler-Kit-2.

[24] Yusuke Oike. 2022. Evolution of Image Sensor Architectures With Stacked Device Technologies. *IEEE Transactions on Electron Devices* 69, 6 (2022), 2757–2765. https://doi.org/10.1109/TED.2021.3097983

[25] Peter Wägemann Phillip Raffeck, Sven Posner. 2024. CO2CoDe: Towards Carbon-Aware Hardware/Software Co-Design for Intermittently-Powered Embedded Systems. In *HotCarbon*.

[26] Thibault Pirson and David Bol. 2021. Assessing the embodied carbon footprint of IoT edge devices with a bottom-up life-cycle approach. *Journal of Cleaner Production* 322 (2021), 128966. https://doi.org/10.1016/j.jclepro.2021.128966

[27] Ian Schneider, Hui Xu, Stephan Benecke, David Patterson, Keguo Huang, Parthasarathy Ranganathan, and Cooper Elsworth. 2025. Life-Cycle Emissions of AI Hardware: A Cradle-To-Grave Approach and Generational Trends. arXiv:2502.01671 [cs.AR] https://arxiv.org/abs/2502.01671

[28] Lucy Smith, Taofeeq Ibn-Mohammed, S.C. Lenny Koh, and Ian M. Reaney. 2018. Life cycle assessment and environmental profile evaluations of high volumetric efficiency capacitors. *Applied Energy* 220 (2018), 496–513. https://doi.org/10.1016/j.apenergy.2018.03.067

[29] Lucy Smith, Taofeeq Ibn-Mohammed, S.C. Lenny Koh, and Ian M. Reaney. 2018. Life cycle assessment and environmental profile evaluations of high volumetric efficiency capacitors. *Applied Energy* 220 (2018), 496–513. https://doi.org/10.1016/j.apenergy.2018.03.067

[30] STMicroelectronics. 2018. X-NUCLEO-LPM01A: STM32 Nucleo Expansion Board for Power Consumption Measurement. https://www.st.com/en/evaluation-tools/x-nucleo-lpm01a.html.

[31] Sustainable Minds, LLC. 2025. Sustainable Minds. https://www.sustainableminds.com/ Accessed: 2025-05-19.

[32] Jennifer Switzer, Gabriel Marcano, Ryan Kastner, and Pat Pannuto. 2022. Junkyard Computing: Repurposing Discarded Smartphones to Minimize Carbon. arXiv:2110.06870 [cs.DC] https://arxiv.org/abs/2110.06870

[33] Swamit Tannu and Prashant J. Nair. 2023. The Dirty Secret of SSDs: Embodied Carbon. *ACM SIGEnergy Energy Informatics Review* 3, 3 (Oct. 2023), 4–9. https://doi.org/10.1145/3630614.3630616

[34] TensorFlow contributors. 2020. TensorFlow Lite Micro Benchmarks. https://github.com/tensorflow/tensorflow/tree/master/tensorflow/lite/micro/benchmarks. Last accessed June 2025; benchmarks for TensorFlow Lite on microcontrollers.

[35] Albert J.P. Theuwissen. 2008. CMOS image sensors: State-of-the-art. *Solid-State Electronics* 52, 9 (2008), 1401–1406. https://doi.org/10.1016/j.sse.2008.04.012 Papers Selected from the 37th European Solid-State Device Research Conference - ESSDERC'07.

[36] Pete Warden. 2018. Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition. *arXiv preprint arXiv:1804.03209* (apr 2018). https://arxiv.org/abs/1804.03209 Speech Commands Dataset.

[37] Olivier Weppe, Thibaut Marty, Sébastien Toussaint, Nicolas Brusselmans, Jean-Christophe Prévotet, Jean-Pierre Raskin, and Maxime Pelcat. 2025. Embodied Carbon Footprint of 3D NAND Memories. In *Inter-national Conference on Computing Frontiers (CF'25)*. Cagliari (Sardinia), France. https://doi.org/10.1145/3706594.3727962

[38] Cheng Zhang, Yu Zheng, Junfeng Jing, Yun Liu, and Haihong Huang. 2022. A comparative LCA study on aluminum electrolytic capacitors: From liquid-state electrolyte, solid-state polymer to their hybrid. *Journal of Cleaner Production* 375 (2022), 134044. https://doi.org/10.1016/j.jclepro.2022.134044

[39] Zhihan Zhang, Felix Hähnlein, Yuxuan Mei, Zachary Englhardt, Shwetak Patel, Adriana Schulz, and Vikram Iyer. 2024. DeltaLCA: Comparative Life-Cycle Assessment for Electronics Design. 8, 1, Article 29 (March 2024), 29 pages. https://doi.org/10.1145/3643561