Spatial-Temporal Embodied Carbon Models for the Embodied Carbon Accounting of Computer Systems

Xiaoyang Zhang
The Hong Kong Polytechnic Univ.
xiaoyang.zhang@connect.polyu.hk

Yijie Yang
The Hong Kong Polytechnic Univ.
yi-jie.yang@connect.polyu.hk

Dan Wang
The Hong Kong Polytechnic Univ.
dan.wang@polyu.edu.hk

ABSTRACT

Embodied carbon is the total amount of carbon released from the processes associated with a product from cradle to gate. In many industry sectors, embodied carbon dominates the overall carbon footprint. *Embodied carbon accounting*, i.e., to estimate the embodied carbon of a product, has become an important research topic.

Existing studies derive the embodied carbon through life cycle analysis (LCA) reports. Current LCA reports only provide the carbon emission of a *product class*, e.g., 28nm CPU, yet a *product instance* can be manufactured from diverse regions and in diverse time periods, e.g., in the winter in Ireland (Intel). It is known that carbon emissions depend on the electricity generation process, which has spatial and temporal dynamics. Therefore, the embodied carbon of a specific product instance can largely differ from its product class. In this paper, we present new Spatial-Temporal Embodied Carbon (STEC) models for embodied carbon accounting. We observe significant differences between current embodied carbon models and STEC models, e.g., for 7nm CPU the difference can be 13.69%. We further examine the impact of STEC models on existing embodied carbon accounting schemes on key computer applications, such as Large Language Model (LLM) inference and LLM training. We observe that using STEC models leads to much greater differences in the embodied accounting of certain applications as compared to others (e.g., 32.26% vs. 6.35%). This is because the hardware requirements of certain applications allow for a wider range of hardware choices, while others have greater restrictions.

CCS CONCEPTS

• **Computer systems organization → Architectures**.

KEYWORDS

Sustainable computing, Carbon accounting, Computer architecture

ACM Reference Format:
Xiaoyang Zhang, Yijie Yang, and Dan Wang. 2024. Spatial-Temporal Embodied Carbon Models for the Embodied Carbon Accounting of Computer Systems. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24), June 04–07, 2024, Singapore, Singapore.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3632775.3661939

E-Energy '24, June 04–07, 2024, Singapore, Singapore

ACM ISBN 979-8-4007-0480-2/24/06
https://doi.org/10.1145/3632775.3661939

1 INTRODUCTION

In recent years, awareness of the importance of sustainability [10, 16, 17, 31] and carbon reduction [21–23, 38] has been growing. *Embodied carbon* is the total amount of carbon released from the processes associated with a product from cradle to gate [15]. In many industry sectors, embodied carbon dominates the overall carbon footprint of a product as compared to its *operational carbon* [4]. For example, the embodied carbon of an iPhone 11 accounts for 79% of its overall carbon footprint [14].

Embodied carbon accounting, i.e., to estimate the embodied carbon of a product, has become an important topic [6, 7, 24, 40]. There are studies on embodied carbon accounting for computer the hardware of processors, memory, and storage [14, 39]. The methodology is to leverage life cycle analysis (LCA) reports [29]. For example, in the Environmental, Social, and Governance (ESG) report of SK hynix, the embodied carbon of memory (LPDDR4) is 48 g/GB [18].

The embodied carbon of a product depends heavily on the *carbon intensity* [26, 41, 43] of the electricity used in the process of manufacturing this product. Specifically, carbon intensity is the amount of carbon emitted when generating a unit of electricity; and different energy sources (e.g., coal or solar) can lead to different carbon emissions when generating a unit of electricity. The carbon intensity of electricity has spatial and temporal dynamics. The spatial dynamics come from the energy policies of the geographic locations. For example, the electricity generated in Taiwan has a higher carbon intensity than that in Ireland, since Taiwan's energy policy is to rely on traditional energy sources due to Taiwan's lack of renewable energy sources. The temporal dynamics come from the environmental dynamics, which affect the amount of renewable energy sources when generating electricity [28, 32]. For example, the electricity generated in Ireland in the winter has less carbon intensity than that in the summer, when wind sources are abundant.

None of the existing studies on embodied carbon accounting has taken spatial and temporal dynamics into consideration. Existing LCA reports on the embodied carbon of a product represent a *product class* (e.g., 28nm CPU) with the same manufacturing process. Yet a *product instance* can be manufactured from diverse regions and in diverse time periods, e.g., in the winter in Ireland (Intel) or in the summer in Taiwan (TSMC).

In this paper, we present new embodied carbon models that can extend existing embodied carbon models to capture the spatial and temporal dynamics in different granularities.[1] Specifically, we observe that (1) energy policies have granularity at the country-level; at the treaty-zone-level of multiple countries with energy treaties, (e.g., the European Union (EU), Association of Southeast Asian Nations (ASEAN), etc.); and at the global-level and (2) environmental dynamics have granularity at a day-level, at a season-level; and at a

[1]We make our codes available: https://github.com/stuabc/STEC

year-level. Embodied carbon models at different granularities will be useful for different applications. For example, manufacturers that regularly purchase computer hardware for their assembly lines may be concerned with embodied carbon at a year-level, while consumers with relatively flexible needs may be more concerned about embodied carbon with finer granularity at a season-level or even a day-level (e.g., individuals with environmentally aware individuals [30]).

Existing models can be considered spatial-temporal embodied carbon (STEC) models of a global-level and a year-level (STEC-GY). In this paper, we also study (1) an STEC model of a country-level and a day-level (STEC-CD), (2) an STEC model of a country-level and a season-level (STEC-CS), and (3) an STEC model of a zone-treaty-level and a year-level (STEC-ZY). We study these models due to the availability of data. Note that one challenge to developing our models is the need to collect public data, which is spread throughout various reports. We have made and effort to collect and *organize* data. We plan to make the data available through an open-source format when this paper goes public.

We compare our models with ACT [14], a state-of-the-art embodied carbon model. We observe significant differences in the embodied carbon accounting when the spatial-temporal factors are taken into consideration and when they are not. First, the differences are significant at any granularity (e.g., 13.69% for 7nm CPU at the country-level and season-level). Second, we find greater dynamics on finer granularity. For example, the average maximum difference in STEC-CD is greater than in STEC-ZY (40.54% vs. 19.29%). Last, we observe that the embodied carbon in some countries has seasonal patterns and can be significantly affected by certain types of weather.

There are an increasing number of studies on embodied carbon accounting for key computer applications, such as Large Language Model (LLM) inference (e.g., CarbonMin [7]) and LLM training (e.g., LLMCarbon [13]). We further examine the impact of STEC models on these embodied carbon accounting schemes. Note that, different applications require different kinds of supporting hardware. We observe that STEC models affect applications differently from other models. For example, the embodied carbon of LLM inference accounted for by STEC-CS can be 0.21g to 0.35g, as compared to 0.31g by CarbonMin; a difference of 32.26%. This is because CarbonMin requires ordinary hardware with a wide range of manufacturing choices across locations and time periods. Yet, LLMCarbon has restrictions on their hardware choices (e.g., Samsung 20nm SSD), and the difference is only 6.35%.

The contributions of the paper can be summarized as follows:

- We observe that the embodied carbon accounting for computer systems can be affected by spatial-temporal factors. Existing embodied carbon models have yet to take these factors into consideration; thus, they can have inaccuracies.
- We develop new STEC models by leveraging two existing models: the ACT model for the embodied carbon of computer components and the electricity carbon intensity model for the spatial-temporal carbon intensity in electricity.
- We evaluate STEC at both the hardware level and the application level. Our evaluations show that there can be non-trivial inaccuracies if the spatial-temporal factors are not taken into account.

2 SPATIAL-TEMPORAL EMBODIED CARBON (STEC) MODELS

2.1 A Framework for STEC Models

We can classify embodied carbon models according to different spatial and temporal granularities (see Table 1). Along the spatial dimension, there is granularity at the country-level, treaty-zone-level, and global-level. Along the temporal dimension, there is granularity at the day-level, season-level, and year-level.

The key difficulty is to collect data on renewable energy sources at different granularities. We made an effort to collect data in order to study STEC-CD/CS/ZY; and we omitted STEC-CY due to space limitations. Current data cannot support us in studying STEC-TD/TS/GD/GS. For clarity, we summarize key acronyms in Table 13 in Appendix A.4.

Table 1: Classification of embodied carbon models based on different spatial and temporal granularity

Spatial \ Temporal	Day-level	Season-level	Year-level
Country-level	STEC-CD	STEC-CS	-
Treaty-zone-level	×	×	STEC-ZY
Global-level	×	×	[14][20][39]

2.2 Spatial-Temporal Carbon Intensity

We first present a state-of-the-art embodied carbon model, the Architectural CO_2 Tool (ACT) Model [14]. It is developed for three kinds of fundamental computer hardware: processors (e.g., CPU, GPU), memory (e.g., DRAM), and storage (e.g., SSD, HDD). ACT estimates the embodied carbon of memory and storage by directly applying the embodied carbon from a corporate annual ESG report. For processors, ACT estimates the embodied carbon by modeling the carbon emitted during manufacturing as three components: (1) the carbon released by the gas (e.g., fluorinated compounds) in the manufacturing process, denoted as *GPS* (gas per-size); (2) the carbon released by procuring raw materials, denoted as *MPS* (material per-size); and (3) the carbon released by the electricity consumed during the manufacturing process. This is calculated by the electricity consumption per unit of size, denoted as *EPS*, and by the carbon intensity of a power grid, denoted by *CI*. *CI* is the amount of carbon emitted when generating a unit of electricity and it can be obtained from the reports on a power grid [12, 27]. Both GPS and EPS can be obtained from a related research paper [3], and MPS can be obtained from industrial research reports [5].

The key to developing an STEC model is to study the three carbon emission components in the manufacturing process in the spatial and temporal dimensions. For GPS and MPS, we conjecture that they are less dynamic in the spatial and temporal dimensions since they are less related to renewable energy, although we admit that further investigation should be carried out.

A greater contribution comes from the carbon emissions from the electricity component. For example, for a 7nm CPU, the carbon emission contributions are GPS 11%, MPS 28%, and EPS 61%. Carbon emissions in electricity depend on the energy sources used to generate a unit of electricity (e.g., solar or coal), as well as on the amount of renewable energy available for use (e.g., solar). The former is diverse in terms of locations, while the latter is diverse in terms of weather and seasons. There have been studies on the spatial and temporal carbon emissions in electricity and we adopt the model in [25].

Here, the carbon intensity $CI(s,t)$ is calculated according to a specific time t and location s. Let $\mathcal{E}$ be the set of energy sources. Let $E^k(s,t)$ be the electricity generated by source k at time t and location s. Clearly $E(s,t) = \sum_{k\in\mathcal{E}} E^k(s,t)$. The carbon emitted by each type of energy source in electricity generation differs. Let ef^k be the *carbon emission factor* of an energy source k. We show the carbon emission factors in Table 9 in Appendix A.2. Finally, the carbon intensity of electricity is the ratio of the total amount of carbon emitted against the total amount of electricity generated ($E(s,t)$).

$$CI(s,t) = \frac{\sum_{k\in\mathcal{E}} ef^k \times E^k(s,t)}{E(s,t)} \tag{1}$$

2.3 Spatial-Temporal Embodied Carbon Models

We now present our models. We capture the spatial and temporal dynamics in the carbon intensity $CI(s,t)$, and the embodied carbon can be calculated using $CI(s,t)$ with the electricity consumption. To develop our spatial and temporal models, we need electricity consumption data; which is difficult to find in these annual reports on memory and storage. Therefore, we first reversely compute the electricity consumption from the information on the embodied carbon.

We first develop basic models for the memory, storage, and processor ($EC_M(s,t), EC_S(s,t), EC_P(s,t)$). The embodied carbon is estimated according to the per-unit size of the respective computer hardware; specifically, the per-unit die size (in cm2) for processors and the per-unit storage size (in Gigabytes/GB) for memory and storage. We then develop STEC-CD, STEC-CS, and STEC-ZY, where we differentiate between different granularities on s (e.g., country-level, treaty-zone-level of multiple countries) and t (e.g., day, season).

2.3.1 Basic models. The three models are illustrated as follows:

Memory: The spatial-temporal embodied carbon of the memory ($EC_M(s,t)$) contains two components: (1) the carbon released by the electricity consumed during the process of manufacturing the memory. This is calculated by multiplying the electricity consumption per unit of size (EPS) by the carbon intensity of a power grid ($CI(s,t)$); and (2) the carbon released independent of electricity, such as raw materials, distribution, and packaging, denoted as α_M. The calculation of α_M involves yearly EC_M, EPS, and BD (bit density), along with the annual carbon intensity CI, all of which can be obtained from the reports. This relationship is depicted in Eq. 2.

$$\alpha_M = EC_M - CI \times EPS \div BD \tag{2}$$

As such, $EC_M(s,t)$ is calculated using Eq. 3.

$$EC_M(s,t) = CI(s,t) \times EPS \div BD + \alpha_M \tag{3}$$

Storage: The spatial-temporal embodied carbon of the storage ($EC_S(s,t)$) contains two components: (1) the carbon emissions from the process of manufacturing the storage, which is determined by multiplying the electricity consumed during the manufacturing process (EPG) by the annual carbon intensity CI; and (2) the carbon released independent of electricity such as raw materials, distribution, and packaging, denoted as α_S. α_S can be found in the industry reports. EPG can be calculated through the annual EC_S in the reports and the annual carbon intensity CI, which is shown in Eq. 4.

$$EPG = (EC_S - \alpha_S)/CI \tag{4}$$

As such, $EC_S(s,t)$ can be calculated using Eq. 5.

$$EC_S(s,t) = CI(s,t) \times EPG + \alpha_S \tag{5}$$

Processor: The spatial-temporal embodied carbon of processor ($EM_P(s,t)$) contains three components as shown in Eq. 6: (1) the carbon released by the electricity consumed during manufacturing, which is calculated by multiplying (EPS) by $CI(s,t)$; (2) the carbon released by raw materials (MPS); and (3) the carbon released by gas (GPS). All of the parameters are shown in Appendix A.1.

$$EC_P(s,t) = GPS + MPS + CI(s,t) \times EPS \tag{6}$$

2.3.2 STEC-CD, STEC-CS, and STEC-ZY. In the STEC-CD model, $t \in \{day_i, \forall i \in [1,365]\}$, $s \in \{country\}$. $\{country\}$ is the set of all of the CPU production places, e.g., China, Korea, USA, etc.

In the STEC-CS model, $t \in \{spring, summer, fall, winter\}$, $s \in \{country\}$. $\{country\}$ is the set of all of the CPU production places.

In the STEC-ZY model, $t \in \{year\}$, $s \in \{zone\}$. $\{year\}$ is the set of production years. $\{zone\}$ is the set of the treaty-zone-level of multiple countries with energy treaties, e.g., ASEAN, EU, etc. Intrinsically, ACT can be considered an STEC-GY model.

2.4 Evaluation

In this section, we compare STEC models with STEC-GY (the ACT model), and we compare them in terms of three types of hardware: CPU, Memory, and Storage.

Table 7 (a) shows the data sources for the carbon emissions of electricity; and (b) shows the data sources for the carbon emissions of other components. From the temporal perspective, we have yearly-level, monthly-level, and daily-level. From the spatial perspective, we have two zones (EU and ASEAN) and countries. Note that we study the regions with semiconductor industry. We select the top regions according to their scales, i.e., the annual semiconductor production. This is the total production of the leading semiconductor companies in the region; and according to the rank in [42], the top companies are Samsung, Intel, TSMC, SK Hynix, Micron, Seagate. The result is {Taiwan, the USA, China, South Korea, Ireland, Italy}. ElectricityMaps lists data for Taiwan, the USA, South Korea, Ireland, and Italy. Data for China were drawn from EMBIR. The data granularity in ElectricityMaps is daily and that in EMBIR is monthly. We evaluate our STEC-CD model using five regions {Taiwan, USA, South Korea, Ireland, and Italy} and STEC-CS using six regions {Taiwan, USA, China, South Korea, Ireland, Italy}.

Table 2: The comparison between STEC-CD and STEC-GY on CPU (7nm), memory (10nm DDR4), SSD [36], and HDD [35].

Hardware	Ave. Difference (%)	Max. Difference (%)
CPU	11.29	42.05
SSD	13.22	49.24
HDD	8.52	31.75
Memory	10.50	39.12
Average	10.88	40.54

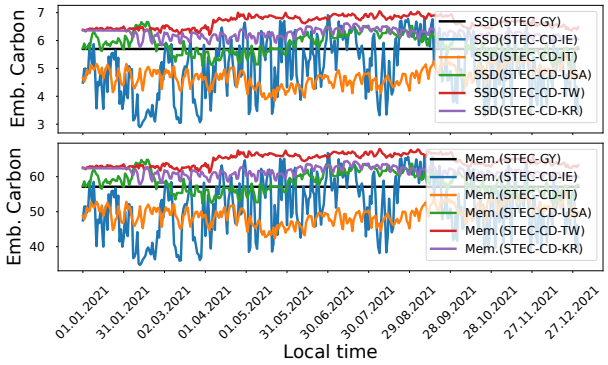

Figure 1: The comparison between STEC-CD and STEC-GY.

2.4.1 Evaluation of STEC-CD. Table 2 presents a comparison between STEC-CD and STEC-GY (baseline). We find the following. (1) The difference in performance between STEC-CD and STEC-GY is significant. Specifically, the average difference is 10.88%, and the average maximum difference is 40.54%. (2) The SSD [36] has a greater difference than the HDD [35] (13.22% vs. 8.52%). The reason for this is that SSD is manufactured using a more advanced process than HDD, resulting in the consumption of more electricity and leading to more dynamics. Therefore, advanced processes will amplify the difference in embodied carbon. (3) The proportion of variable renewable energy sources (VRE) in electricity can affect the dynamics of the embodied carbon. For example, there is more VRE in the grids of Ireland (31.45%) than in the grids of Taiwan (6.44%). This brings more dynamics for the embodied carbon of the memory manufactured in Ireland (with a variance of 63.63) than in Taiwan (with a variance of 2.74), as shown in Fig. 1.

2.4.2 Evaluation of STEC-CS. Table 3 compares STEC-CS and STEC-GY (baseline) in the six major IC manufacturing regions. Fig. 2 visually shows the embodied carbon (g/cm^2) output by the STEC-CD and STEC-GY in the CPU in Ireland and Italy from 2019 to 2022. We find the following. (1) The difference is significant (the average difference: 13.33%, the average maximum difference: 27.50%). For each type of hardware, the average difference is 13.69% (CPU), 16.51% (SSD), 10.36% (HDD), and 12.74% (Memory). (2) The embodied carbon has seasonal patterns in some regions. As Fig. 2 shows, the embodied carbon is higher in summer and lower in winter in Ireland, while, the opposite trend is true in Italy. The pattern depends on the climate and on the composition of energy sources in the grids of a region. Specifically, wind power dominates in the variable renewable energy sources in Ireland, while, solar power dominates in the variable renewable energy sources in Italy. In addition, wind is abundant in the winter in Ireland, while solar power is abundant in the summer in Italy. These together lead to different patterns in the two countries.

Table 3: The comparison between STEC-CS and STEC-GY on CPU (7nm), memorgy (10nm DDR4), SSD [36], and HDD [35].

Hardware	Ave. Difference (%)	Max. Difference (%)
CPU	13.69	27.32
SSD	16.51	34.46
HDD	10.36	21.63
Memory	12.74	26.60
Average	13.33	27.50

Table 4: The comparison between STEC-ZY and STEC-GY on CPU (7nm), memory (10nm DDR4), SSD [36], and HDD [35].

Hardware	STEC-ZY		STEC-GY	Ave. difference (%)	Max. difference (%)
	Emb.carbon in EU	Emb.carbon in ASEAN	Emb.carbon		
CPU	1266.88	1848.10	1557.49	18.65	19.99
SSD	4.77	7.41	6.09	21.64	23.19
HDD	4.36	5.81	5.08	14.28	15.3
Memory	70.72	49.72	60.22	17.42	18.67
Average	/	/	/	18.00	19.29

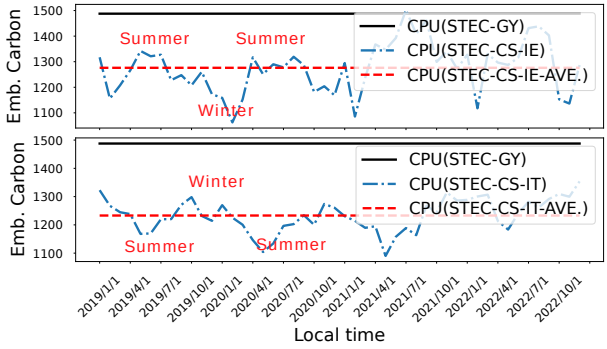

Figure 2: The comparison between STEC-CS and STEC-GY.

2.4.3 Evaluation of STEC-ZY. Table 4 presents a comparison between STEC-ZY and STEC-GY (baseline) in the two treaty zones (ASEAN and EU). Overall, the average difference is 18.00%, and the average maximum difference is 19.29%. We find that (1) the STEC-CD/CS/ZY models all have significant differences with the STEC-GY (baseline) model. (2) Finer granularity can bring more dynamics. For example, the average maximum difference between STEC-CD and STEC-GY is greater than that between STEC-CS and STEC-GY, STEC-ZY and STEC-GY (40.54% compared to 27.50%, 40.54% compared to 19.29%).

2.4.4 A Case on Storm Malik Effect. Fig. 3 visually presents dynamics in the embodied carbon of the CPU and the energy sources in grids in Ireland during the period of Storm Malik. For example, the embodied carbon reaches the maximum value (1276.87 g/cm^2) at 2

a.m. on the day before Storm Malik arrives. When Storm Malik arrives, the embodied carbon decreases rapidly, reaching a minimum ($894.01\ g/cm^2$) at 5 p.m. The reason for this decline is that wind power increases rapidly under the influence of the storm. We can find that the embodied carbon has more variance in hourly time scale affected by weather.

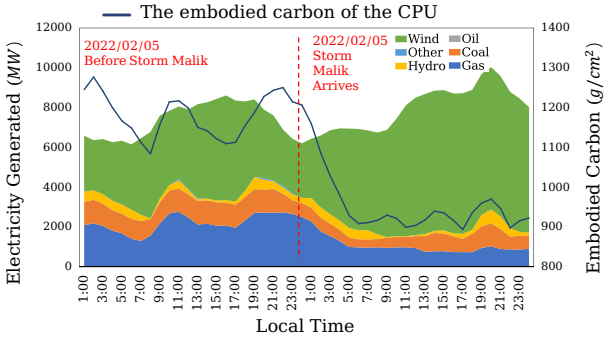

Figure 3: A case on storm Malik effect in Ireland.

3 STEC ACCOUNTING FOR APPLICATIONS

Embodied Carbon Accounting for Computer Applications: Nowadays, there are studies on embodied carbon accounting and optimization on computer applications, such as CarbonMin for LLM inference [7], LLMCarbon for LLM training [13], Carbond for operating systems [33], and others. The embodied carbon of these applications is accounted for by the hardware that they require, which differs across applications to meet their performance requirements. For example, LLMCarbon performs carbon accounting on LLM training and requires more high-end memory and storage (e.g., 64×256GB Mircron 18nm, 64×32TB Samsung 20nm), whereas CarbonMin performs LLM inference and requires ordinary memory and storage (e.g., 900GB memory, 6000GB SSD).

Existing studies have estimated the carbon of specific hardware from ESG reports, which have been selected an ad hoc manner. Due to space limitations, in this paper we examine the embodied carbon accounting of two applications, LLM inference (CarbonMin) and LLM training (LLMCarbon) under the STEC-CS model, i.e., we recalculate CarbonMin and LLMCarbon using the STEC-CS model. Details of the calculation are given in Appendix A.3.

Results: The main criterion for CarbonMin is embodied carbon (gram) per 1000 inference requests (specifically, ChatGPT-3 inference). The main criterion of LLMCarbon is the embodied carbon (ton) of the training of a large language model (specifically, the XLM model [9] with 570M parameters).

The hardware of CarbonMin and LLMCarbon is listed in Table 5 (a) and Table 6 (a). The embodied carbon is listed in Table 5 (b) and Table 6 (b), respectively. We develop algorithms to find the location and time period with the minimum amount of embodied carbon to manufacture the hardware, as well as with the 20th percentile, median, 80 percentile, and maximum amount (see Table 5 (b) and Table 6 (b)). We observe that STEC-CS leads to great differences for the LLM inference. Specifically, the embodied carbon of the ChatGPT-3 inference accounted for by CarbonMin is the 0.31g/1K inference.

Table 5: The comparison between CarbonMin and STEC.

Table (a)	Hardware	Number/Size	Technology Nodes	Die Size
	CPU	1	7nm	74 mm2
	GPU	8	7nm	826 mm2
	Memory	900GB	/	/
	Storage	6000GB	/	/

Table (b)	Embodied carbon (g)/1k inference requests				
	Min.	20th percentile	Median	80th percentile	Max.
CarbonMin			0.31		
STEC-CS	0.21	0.25	0.28	0.33	0.35
Dif. (%)	32.26%	19.35%	9.68%	6.45%	12.90%

Table 6: The comparison between LLMCarbon and STEC.

Table (a)	Hardware	Description	Unit	Number
	GPU	V100 TSMC 12nm	815 mm2	512
	CPU	TSMC 16nm	147 mm2	64
	SSD	Samsung 20nm	32TB	64
	DRAM	Micron 18nm	256GB	64
	Other	/	/	64

Table (b)	Embodied carbon (t)				
	Min.	20th percentile	Median	80th percentile	Max.
LLMCarbon			0.63		
STEC-CS	0.61	0.62	0.64	0.66	0.67
Dif. (%)	3.17%	1.59%	1.59%	4.76%	6.35%

Under STEC-CS, the embodied carbon can be the 0.21g/1K inference (or the 0.25g/1K inference under the 20th percentile), a difference of 32.26% (or 19.35%). As a comparison, STEC-CS leads to small differences for LLM training, e.g., the differences between LLMCarbon and STEC-CS is 3.17% (or 1.59% under the 20th percentile). This is because the LLM inference requires ordinary hardware with great spatial-temporal choices in manufacturing, whereas the hardware for LLM training, to the best of our knowledge, is more restricted.

4 CONCLUSION AND DISCUSSIONS

This paper observes that for the same product class of computer hardware, e.g., 28nm CPU, the carbon emissions of its product instances (e.g., made in the summer in Taiwan or in the winter in Ireland) can differ. We present new spatial-temporal embodied carbon (STEC) models by extending the embodied carbon model in the spatial-temporal dimension. We evaluated STEC with a state-of-the-art embodied carbon accounting model, and also applied STEC for embodied carbon accounting for two computer applications. The results showed that significant differences can exist; in other words, neglecting the spatial and temporal factors may lead to non-trivial inaccuracies in embodied carbon accounting.

We would comment that STEC also has inaccuracies. Inaccuracy comes from using averaged/aggregated data in various aspects. There is an industry value-chain, and our calculation of the carbon emissions in the upstream industry value-chain uses industry averages. As one example, when we estimate the embodied carbon of processors, we use the value of $500 g/cm^2$ for MPS (material per size). This is the carbon released from upstream raw materials, and this is an average value which has inaccuracies.

ACKNOWLEDGMENTS

Dan Wang's work is supported by RGC GRF 15209220, 15200321, 15201322, RGC-CRF C5018-20G, and ITC ITF-ITS/056/22MX.

REFERENCES

[1] AWS. 2023. Amazon EC2. https://aws.amazon.com/ec2/instance-types/?nc1=h_ls.

[2] Microsoft Azure. 2023. ND A100 v4-series. https://learn.microsoft.com/en-us/azure/virtual-machines/nda100-v4-series/.

[3] M Garcia Bardon, P Wuytens, L-Å Ragnarsson, G Mirabelli, D Jang, G Willems, A Mallik, A Spessot, J Ryckaert, and B Parvais. 2020. DTCO including sustainability: Power-performance-area-cost-environmental score (PPACE) analysis for logic technologies. In *2020 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 41–4.

[4] Noman Bashir, David Irwin, and Prashant Shenoy. 2023. On the Promise and Pitfalls of Optimizing Embodied Carbon. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. 1–6.

[5] Sarah B Boyd. 2011. *Life-cycle assessment of semiconductors*. Springer Science & Business Media.

[6] Andrew A Chien. 2023. GenAI: Giga ,TeraWatt-Hours,andGigaTonsofCO2. *Commun. ACM* 66, 8 (2023), 5–5.

[7] Andrew A Chien, Liuzixuan Lin, Hai Nguyen, Varsha Rao, Tristan Sharma, and Rajini Wijayawardana. 2023. Reducing the Carbon Impact of Generative AI Inference (today and in 2035). In *The 2nd Workshop on Sustainable Computer Systems (HotCarbon'23)*. 9.

[8] Jeongdong Choe. 2017. XPoint Memory Comparison Process Architecture. https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2017/20170808_FR12_Choe.pdf.

[9] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 8440–8451.

[10] Yang Deng, Rui Liang, Dan Wang, Ao Li, and Fu Xiao. 2023. Decomposition-based Data Augmentation for Time-series Building Load Data. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 51–60.

[11] EMBIR. 2023. Electricity Sector. https://ember-climate.org/countries-and-regions/.

[12] ENTSOE. 2023. transparency platform. https://transparency.entsoe.eu/dashboard/show.

[13] Ahmad Faiz, Sotaro Kaneda, Ruhan Wang, Rita Osi, Parteek Sharma, Fan Chen, and Lei Jiang. 2023. LLMCarbon: Modeling the end-to-end Carbon Footprint of Large Language Models. *arXiv preprint arXiv:2309.14393* (2023).

[14] 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.

[15] Geoffrey Hammond, Craig Jones, Ed Fiona Lowrie, and Peter Tse. 2011. Embodied carbon. *The inventory of carbon and energy (ICE). Version (2.0)* (2011).

[16] Fang He, Xiaoyang Zhang, and Dan Wang. 2022. Cement-α: an ontology-based data access system for building analytics with multiple data sources. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*. 436–437.

[17] Heidi Hottenroth, C Sutardhio, A Weidlich, I Tietze, S Simon, W Hauser, T Naegler, L Becker, J Buchgeister, T Junne, et al. 2022. Beyond climate change. Multi-attribute decision making for a sustainability assessment of energy system transformation pathways. *Renewable and Sustainable Energy Reviews* 156 (2022), 111996.

[18] SK hynix. 2021. Sustainability Report. https://www.skhynix.com/sustainability/UI-FR-SA1601/.

[19] World in Data. 2023. Carbon intensity of electricity. https://ourworldindata.org/grapher/carbon-intensity-electricity?tab=chart.

[20] Sven Köhler, Benedict Herzog, Henriette Hofmeier, Manuel Vögele, Lukas Wenzel, Andreas Polze, and Timo Hönig. 2023. Carbon-Aware Memory Placement. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. 1–7.

[21] Liuzixuan Lin and Andrew A Chien. 2023. Adapting Datacenter Capacity for Greener Datacenters and Grid. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 200–213.

[22] Liuzixuan Lin and Andrew A Chien. 2023. Reducing Datacenter Operational Carbon Emissions Effectively by Cooperating with the Grid. *arXiv preprint arXiv:2301.03148* (2023).

[23] Liuzixuan Lin, Victor M Zavala, and Andrew A Chien. 2021. Evaluating coupling models for cloud datacenters and power grids. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 171–184.

[24] Diptyaroop Maji, Noman Bashir, David Irwin, Prashant Shenoy, and Ramesh K Sitaraman. 2023. Untangling Carbon-free Energy Attribution and Carbon Intensity Estimation for Carbon-aware Computing. *arXiv preprint arXiv:2308.06680* (2023).

[25] Diptyaroop Maji, Prashant Shenoy, and Ramesh K Sitaraman. 2022. CarbonCast: multi-day forecasting of grid carbon intensity. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 198–207.

[26] Diptyaroop Maji, Ramesh K Sitaraman, and Prashant Shenoy. 2022. DACF: day-ahead carbon intensity forecasting of power grids using machine learning. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*. 188–192.

[27] Electricity Maps. 2019. electricity data. https://www.seagate.com/esg/planet/product-sustainability/.

[28] Stefan Meisenbacher, Benedikt Heidrich, Tim Martin, Ralf Mikut, and Veit Hagenmeyer. 2023. AutoPV: Automated photovoltaic forecasts with limited information using an ensemble of pre-trained models. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 386–414.

[29] Karim Ali Ibrahim Menoufi et al. 2011. Life cycle analysis and life cyle impact assessment methodologies: a state of the art. (2011).

[30] Divya Pandey, Madhoolika Agrawal, and Jai Shanker Pandey. 2011. Carbon footprint: current methods of estimation. *Environmental monitoring and assessment* 178 (2011), 135–160.

[31] Mario Ragwitz, Anke Weidlich, and Dirk Biermann. 2023. Scenarios for a climate-neutral Germany. Technology transformation, consumption reduction and carbon management. Analysis. (2023).

[32] Hartmut Schmeck, Antonello Monti, and Veit Hagenmeyer. 2022. Energy informatics: key elements for tomorrow's energy system. *Commun. ACM* 65, 4 (2022), 58–63.

[33] Andreas Schmidt, Gregory Stock, Robin Ohs, Luis Gerhorst, Benedict Herzog, and Timo Hönig. 2023. carbond: An Operating-System Daemon for Carbon Awareness. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. 1–6.

[34] Seagate. 2019. ESG. https://www.seagate.com/esg/planet/product-sustainability/.

[35] Seagate. 2019. Exos 7E8. *https://www.seagate.com/gb/en/esg/planet/product-sustainability/exos-7e8-sustainability-report/* (2019).

[36] Seagate. 2019. Nytro 3530. *https://www.seagate.com/gb/en/esg/planet/product-sustainability/nytro-3530-sustainability-report/* (2019).

[37] Seagate. 2020. Product Life Cycle Assessments. *https://www.seagate.com/gb/en/esg/planet/product-sustainability/* (2020).

[38] Thanathorn Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. Spatiotemporal Carbon-aware Scheduling in the Cloud: Limits and Benefits. In *Companion Proceedings of the 14th ACM International Conference on Future Energy Systems*.

[39] Swamit Tannu and Prashant J Nair. 2022. The dirty secret of SSDs: Embodied carbon. *arXiv preprint arXiv:2207.10793* (2022).

[40] Philip Ulrich, Tobias Naegler, Lisa Becker, Ulrike Lehr, Sonja Simon, Claudia Sutardhio, and Anke Weidlich. 2022. Comparison of macroeconomic developments in ten scenarios of energy system transformation in Germany: National and regional results. *Energy, Sustainability and Society* 12, 1 (2022), 1–19.

[41] Jan Frederick Unnewehr, Anke Weidlich, Leonhard Gfüllner, and Mirko Schäfer. 2022. Open-data based carbon emission intensity signals for electricity generation in European countries–top down vs. bottom up approach. *Cleaner Energy Systems* 3 (2022), 100018.

[42] Wiki. 2019. Semiconductor industry. https://en.wikipedia.org/wiki/Semiconductor_industry.

[43] Xiaoyang Zhang and Dan Wang. 2023. A GNN-based Day Ahead Carbon Intensity Forecasting Model for Cross-Border Power Grids. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 361–373.

A APPENDIX

To enable further investigation into the spatial-temporal aspects of embodied carbon, we provide detailed sources of data for the STEC models and describe the configurable parameters within the proposed STEC models. Finally, we present detailed calculations of the STEC models of applications.

A.1 Data sources for STEC Models

Table 7 presents the data sources for STEC models. All of the data are public. Table 7(a) shows the sources for the data on electricity data, including the carbon intensity (yearly and monthly) and energy sources (hourly). Table 7(b) shows the sources for the hardware-related parameters.

Table 7: The Data Sources for the STEC models

(a) Data Sources for the Carbon Emissions of Electricity Energy

	EU	ASEAN	China	Taiwan	South Korea	USA	Ireland	Italy	Source
Yearly [2019-2022]	✓	✓	✓	✓	✓	✓	✓	✓	Our World in Data [19]
Monthly [2019-2022]	-	-	✓	✓	✓	✓	✓	✓	EMBIR [11]
Daily [2021]	-	-	-	✓	✓	✓	✓	✓	ENTSOE [12] ElectricityMaps [27]

(b) Data Sources for Other Types of Carbon Emission

Parameter	Description	Unit	Source
EPS	Electricity consumed per die Size	kWh/cm2	Research paper [3]
GPS	Carbon emission from Gas used per die Size	g/cm2	Research paper [3]
MPS	Carbon emission from Material used per die Size	g/cm2	Industrial research reports [5]
BD	Bit density for memory	GB/cm2	Industrial research reports [8]
EPG	Electricity consumed per GB	kWh/GB	Industrial environmental reports of manufacturers [18, 37]

Table 8: Carbon intensity of electricity in six major IC production regions in 2022.

Countriy/Region	Carbon intensity of electricity (g/kWh) (Date Sources: Our World in Data [19])
Taiwan	561
China	531
South Korea	436
United States	367
Italy	372
Ireland	346

Table 9: Carbon emission factors (g/kWh) for different energy sources

Energy sources	Direct Emission Factors (Data Source: Research Paper [25])
Oil	406
Coal	760
Natural gas	370
Nuclear	0
Wind	0
Solar	0
Hydro	0
Geothermal	0
Biomass	0
Other	575

A.2 Embodied Carbon Parameters

Table 8 presents the annual average carbon intensity of electricity in six major IC production regions in 2022. Due to space limitations, we do not present monthly and daily data here. We make all the organized data that we have collected available in our public codes. Table 9 summarizes the direct carbon emissions factors for each energy source. Tables 10, 11, and 12 summarize the related embodied carbon data for processors, memory, and storage, respectively.

Table 10: Embodied carbon parameters in semiconductor manufacturing

Process Node	EPS (kWh/cm2) (Data Source: Research Paper [3])	GPS (g/cm2) (Data Source: Research Paper [3])	MPS (g/cm2) (Data Source: Environmental Report [5])
28	0.9	100	500
20	1.2	110	500
14	1.2	125	500
10	1.475	150	500
7	1.52	200	500
7-EUV	2.15	200	500
7-EUV-DP	2.15	200	500
5nm	2.75	225	500
3nm	2.75	275	500

Table 11: Embodied carbon, bit density, and carbon emissions from electricity consumption of memory [8, 14].

Technology	Embodied Carbon (g/GB)	Bit Density (G/mm2)	Carbon Emissions from Electricity Consumption (g/GB)
30nm LPDDR3	230	0.06	67.50
20nm LPDDR3	184	0.11	51.43
10nm DDR4	65	0.19	35.74
LPDDR4	48	0.17	39.04

A.3 The STEC Model of Applications

The embodied carbon of applications (EC_{app}) run on computing systems can be calculated using $EC_{app} = EC_{sys} \times TS \times RS$, where EC_{sys} is the total embodied carbon of the system that the application run on, TS is the time-share (the fraction of the total lifetime of the system used by the application), and RS is the resource-share (the fraction of the total available hardware resources reserved by the application) [33]. The STEC models of applications differ from the previous model in that EC_{sys} is not simply an average value accounted by ACT, but rather a range determined by the hardware range (HR) with spatial-temporal dynamics that make up the system. The suppliers of computing systems usually provide the basic HR, such as Microsoft [2] and AWS [1]. Therefore, EC_{sys} can be calculated as follows:

Table 12: The summary of the embodied carbon of storage from product reports [34]

Types	Technology	Embodied Carbon (g/GB)	Manufacturing Energy Carbon (g/GB)	Other Carbon (g/GB)
Enterprise SSD	Nytro 3530	6.27	4.25	2.02
	Nytro 1551	3.91	1.53	2.38
	Nytro 3331	5.48	0.92	4.56
	Nytro 3332	2.42	0.78	1.64
Consumer SSD	BarraCuda 120 SSD	26.28	23.85	2.43
Enterprise HDD	EXOS X20	0.88	0.36	0.52
	EXOS X18	0.88	0.39	0.49
	Exos 2X14	1.28	0.51	0.78
	Exos 7E8	5.28	2.34	2.94
	Exos 5E8	2.54	1.14	1.40
	Exos 10E2400	10.75	6.94	3.81
	EXOS 15E900	21.62	10.65	10.97
	Exos X16	1.46	0.77	0.69
	Exos X12	1.32	0.53	0.79
Consumer HDD	BarraCuda 3.5	9.40	4.84	4.56
	BarraCuda	4.25	2.08	2.17
	BarraCuda Pro	2.62	1.22	1.40
	FireCuda	5.16	3.81	1.35
	IronWolf	5.28	2.22	3.06
	IronWolf Pro	3.80	1.33	2.47
	Skyhawk 3 TB	9.85	2.17	7.68
	Skyhawk Surveillance HDD	4.37	1.54	2.83
	Skyhawk 6 TB	4.18	1.09	3.09
	Video 3.5 HDD	8.20	3.22	4.98
	Video 3.5 HDD (Pipeline HDD)	9.54	3.23	6.31
External HDD	ULTRA TOUCH	5.54	3.40	2.13
	Rugged Mini	4.22	2.98	1.25

$$EC_{sys} = STEC(HR_{sys}) = \sum_{i \in HR_{sys}} EC_i(s_i, t_i) + N_{sys} \times P \quad (7)$$

where $EC_i(s_i, t_i)$ is the embodied carbon of hardware i with production location s_i and time t_i, P is the packaging footprint (0.15kg [14]), and N_{sys} is the number of the hardware.

A.3.1 The Carbon Footprint of LLM Inference. CarbonMin [7] is a model for estimating and reducing the carbon footprint of LLMs inference. CarbonMin estimates the carbon footprint of ChatGPT inference on Azure ND A100 v4-series instances [2]. We reestimate the embodied carbon based on the data in CarbonMin. Specifically, we assume that the overall lifetime of the hardware is 3 years, and that the *RS* is 1. According to the experimental result of CarbonMin, the average computation per inference request is 2.07 GPU-sec/request. Thus, the TS of one request is 2.07 sec/3 years. Then, the embodied carbon of one ChatGPT-3 inference can be obtained by $EC_{infer} = EC_{sys} \times TS \times RS = EC_{sys} \times 2.07 sec/3 years$. Table 5 (a) presents the *HR* of the Azure ND A100 v4-series used in CarbonMin. The *HR* of CPU and GPU are restricted to TSMC in Taiwan, while the *HR* of storage and memory contains six major IC-producing regions (see Table 8).

A.3.2 The Carbon Footprint of LLM Training. We reestimate the embodied carbon of training the LLM [9] using STEC models based on the data in LLMCarbon. Specifically, the training duration of the LLM is 20.4 days, the lifetime of hardware is 5 years, and *RS* is 1. Consequently, the embodied carbon can be obtained by $EC_{train} = EC_{sys} \times TS \times RS = EC_{sys} \times 20.4 days/5 years$. Table 6 (a) presents the *HR* in LLMCarbon, where the *HR* of each component is restricted to a specific type. Therefore, the dynamics of *HR* mainly come from production time.

Table 13: The summary of acronyms.

Acronyms	Descriptions
CPU	Central Processing Unit
GPU	Graphics processing unit
SSD	Solid-state drive
HDD	Hard disk drive
DRAM	Dynamic random access memory
STEC	Spatial-Temporal Embodied Carbon
STEC-CD	Spatial-Temporal Embodied Carbon in Country and Day level
STEC-CS	Spatial-Temporal Embodied Carbon in Country and Season level
STEC-GY	Spatial-Temporal Embodied Carbon in Gobal and Year level
STEC-ZY	Spatial-Temporal Embodied Carbon in Zone and Year level
GPS	Gas per size
EPS	Energy per size
MPS	Material per size
CI	Carbon intensity

A.4 The summary of acronyms

Here, we summarize the acronyms in Table 13.