

DeltaLCA: Comparative Life-Cycle Assessment for Electronics Design

ZHIHAN ZHANG*, FELIX HÄHNLEIN*, and YUXUAN MEI*, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

ZACHARY ENGLHARDT, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

SHWETAK PATEL, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

ADRIANA SCHULZ, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

VIKRAM IYER, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA




Fig. 1. DeltaLCA enables designers to rapidly compare the environmental impact (EI) of two PCB designs. In the traditional Life Cycle Assessment (LCA) pipeline (top row), LCA experts manually generate an inventory of parts and match them to a database to calculate the EI. DeltaLCA (bottom row) first automates inventory generation and estimates partial EI for known components, then performs a user-in-the-loop comparison using domain-specific heuristics to determine if Design A has a greater EI than Design B.

*These authors contributed equally to this research.

Authors' addresses: Zhihan Zhang, zzhihan@cs.washington.edu; Felix Hähnlein, fhahnlei@cs.washington.edu; Yuxuan Mei, ym2552@cs.washington.edu, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; Zachary Enghardt, zacharye@cs.washington.edu, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; Shwetak Patel, shwetak@cs.washington.edu, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; Adriana Schulz, adriana@cs.washington.edu, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; Vikram Iyer, vsiyer@uw.edu, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

2474-9567/2024/1-ART30

<https://doi.org/10.1145/3643561>

Reducing the environmental footprint of electronics and computing devices requires new tools that empower designers to make informed decisions about sustainability during the design process itself. This is not possible with current tools for life cycle assessment (LCA) which require substantial domain expertise and time to evaluate the numerous chips and other components that make up a device. We observe first that informed decision-making does not require absolute metrics and can instead be done by comparing designs. Second, we can use domain-specific heuristics to perform these comparisons. We combine these insights to develop DeltaLCA, an open-source interactive design tool that addresses the dual challenges of automating life cycle inventory generation and data availability by performing comparative analyses of electronics designs. Users can upload standard design files from Electronic Design Automation (EDA) software and the tool will guide them through determining which one has greater carbon footprints. DeltaLCA leverages electronics-specific LCA datasets and heuristics and tries to automatically rank the two designs, prompting users to provide additional information only when necessary. We show through case studies DeltaLCA achieves the same result as evaluating full LCAs, and that it accelerates LCA comparisons from eight expert-hours to a single click for devices with ~30 components, and 15 minutes for more complex devices with ~100 components.

CCS Concepts: • **Human-centered computing** → **Interactive systems and tools**; • **Applied computing** → *Operations research*; • **Computing methodologies** → *Knowledge representation and reasoning*.

Additional Key Words and Phrases: Sustainable Computing, Life Cycle Assessment, Domain-Specific Heuristics, Linear Programming

ACM Reference Format:

Zhihan Zhang, Felix Hähnlein, Yuxuan Mei, Zachary Englhardt, Shwetak Patel, Adriana Schulz, and Vikram Iyer. 2024. DeltaLCA: Comparative Life-Cycle Assessment for Electronics Design. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 1, Article 30 (January 2024), 29 pages. <https://doi.org/10.1145/3643561>

1 INTRODUCTION

For the past 70 years, computing devices and systems have been designed with increasing complexity driven solely by demand for computational power and little foresight to sustainability or disposal. As a result, current estimates of climate warming emissions from the overall information and communication technology (ICT) sector range from 2.1-3.9% [23] of total global emissions and are projected to grow rapidly to 8% over the next decade if left unchecked [5]. Reducing these emissions requires not only switching to carbon-free energy sources, but also addressing the more specific problem of embodied carbon from device manufacturing. Particularly for ubiquitous consumer devices like smartphones and laptops, manufacturing often accounts for 80% or more of life cycle emissions [9, 10, 24].

This presents an opportunity to reduce environmental impact (EI) by optimizing future device designs with embodied carbon in mind. Making sustainable design decisions however requires the ability to iteratively analyze a design and determine its carbon footprint, which is impossible with today's tools. The environmental impacts of a device, such as those quoted above, are typically quantified retrospectively through a manual Life Cycle Assessment (LCA) in which a human expert analyzes the impacts of the product's production, usage, and disposal. This is challenging for computing devices, which are complex systems composed of numerous parts, including a CPU, memory, power management circuitry, and much more. For example, the printed circuit board (PCB) of a consumer device may have 500+ parts [40]. Moreover, mapping this complex inventory to environmental costs requires data about the semiconductor fabrication processes used to make them, which are often proprietary to protect intellectual property. As a result, while tools exist to support users in computing LCA, performing a rigorous assessment requires significant expertise and could take months for complex designs [14, 51], and even experts often lack perfect data. This makes it impractical for designers to compute multiple LCAs for different design variations, thereby making it hard to take sustainability into account during the design process itself.

We propose a novel approach to sustainability-focused design for reducing the manufacturing carbon footprint of computing devices at the PCB design stage. Our approach is anchored in two principal ideas:

Insight 1: Our first key insight is that this design tool should be centered around empowering designers with informed decision-making during both the drafting and iteration stages. This approach recognizes that conducting conventional LCAs for each design alternative is not only costly but also impractical as search spaces become larger. We observe that informed decision-making does not require absolute metrics, but can be done instead with relative comparisons between different designs.

Insight 2: Our second key insight is that even if we do not have perfect information, we can leverage domain-specific knowledge of electronics manufacturing to reason about relative differences between parts or designs.

We combine these insights to develop DeltaLCA, an open-source interactive design tool that addresses the dual challenges of automating life cycle inventory generation and data availability by performing comparative analyses for electronics. Users can input two standard PCB design files from Electronic Design Automation (EDA) software, and the tool will guide the users in determining which one has a greater carbon footprint. DeltaLCA leverages electronics-specific LCA datasets and heuristics and tries to automatically rank the two designs, prompting users to provide additional information only when necessary, through a user interface designed to highlight missing data. This approach dramatically simplifies LCA computation by canceling out common components, automatically estimating the carbon footprint of remaining components where possible, and providing a user-friendly interface to input new information or use heuristics to compare the remaining parts automatically.

Designing a tool that can compare the environmental impact of electronic designs requires addressing two fundamental challenges. The first challenge is that typical LCAs require developing a highly detailed life cycle inventory (LCI) of the raw materials used, their extraction impacts, manufacturing process steps, and energy consumed. This is particularly challenging in electronics where every device is a complex combination of integrated circuits (ICs) and other components that result from complex semiconductor manufacturing pipelines and global distribution networks. Component distributors like DigiKey have millions of unique parts in their catalogs, and there are no comprehensive databases of their environmental impact.

To address this challenge, we develop an automated pipeline to extract relevant LCI data from the outputs of common PCB design tools such as KiCAD and EAGLE. We develop a custom Bill-of-Materials (BOM) generator that goes beyond existing tools focused on generating fabrication files to extract detailed information about part categories, footprint sizes, and specifications relevant to environmental analyses using large online parts databases. We then analyze this data to automatically group classes of devices, such as passive components and classes of ICs. We automatically infer essential information for LCA computation, such as die sizes and process technology nodes, and combine these with open-source measurement data to automate carbon footprint estimation for a large class of parts.

While the solutions above provide significant insights into the composition and environmental footprint of a device, there remain significant gaps in the data. This brings us to the second challenge: data availability is a fundamental problem in LCA. To address this, we revisit our motivation of empowering designers to make sustainable decisions at various levels, including individual components, sub-modules of a partial design, and complete candidate designs. We also acknowledge that a common design goal is to reduce EIs in successive iterations of a product [2]. In these cases, we care more about the *relative improvement* of the carbon footprint than the raw impact numbers themselves. As a result, if two designs have a shared set of components with unknown impacts, computing the delta between the designs cancels out these unknowns.

To address the remaining unique parts that are not canceled out, we can take this idea a step further and reduce relative improvement to a binary decision of whether the impact of Design A is greater than Design B. To do this, we leverage both an estimate of the environmental impact (EI) for some parts, whenever these can be computed, and domain-specific heuristics (DSH) which enable us to perform a comparative LCA between parts without having access to full impact data. The final design comparison can then be abstracted into the problem of matching parts of Design A with those from Design B using this partial comparative information. We formulate




Fig. 2. a) LCA stages and boundaries. b) Contribution of each life cycle stage to total carbon footprint in commercial devices with available data: Dell XPS 15 9500 [18], MacBook Pro 13 [6], Samsung Galaxy S23 [45], Kindle Fire 10 [3], iPad 10th Gen [7], Apple Watch 9 [9], Fitbit Charge 6 [24], Pixel 8 Pro [26], iPhone 15 Pro [10], iPhone 14 Pro [8]. c) Example LCA methodology used by Amazon Devices Sustainability [2].

this matching problem as an integer program, which can be solved efficiently. In cases where a complete match is unattainable, we offer a user interface that displays the unmatched parts, allowing users to contribute additional rules for a conclusive comparison.

We summarize our contributions below:

- DeltaLCA is the first user-in-the-loop design tool that directly empowers electronics designers to make sustainable design decisions by comparing the environmental footprints of two different designs.
- We develop an end-to-end pipeline that directly integrates with common electronics EDA tools to automatically generate a parts inventory and extract the key parameters for LCA computation.
- We develop methods to infer proprietary information such as die sizes and process technology nodes and use them to automatically estimate the carbon footprint of parts for which information is available. Using these, we demonstrate fully automated carbon estimates for two designs with full LCAs and prove our method produces a correct comparison result.
- We create a set of five domain-specific heuristics for comparing parts for which data is unavailable using proxies such as die size and minimum package area to identify which has a higher environmental impact. In case studies, we find that our tool can match an average of 88% of parts using as few as three heuristics.
- We develop a matching algorithm that sorts unknown parts into classes and synthesizes the environmental impact estimates and domain-specific heuristics to determine which design is more sustainable. Our algorithm can process inputs with ~2,500 variables in 0.5 s on a laptop, enabling real-time comparisons. Our case studies incorporating user input for conclusive comparisons when needed accelerated LCA comparisons to 15 min.

2 ELECTRONICS LCA PRIMER

Performing a Life Cycle Assessment (LCA) involves systematically evaluating the EIs associated with stages of a product's life cycle, from raw material extraction through materials processing, manufacturing, distribution, use, repair and maintenance, and disposal or recycling. This analysis of the total EIs enables manufacturers and consumers to make more informed decisions. The process begins with defining the scope and goal of the assessment, which includes identifying the product to be assessed and the boundaries of the study as shown in Figure 2a. Common LCA boundaries include “cradle-to-gate” assessments which account for raw material extraction through product manufacturing, and “cradle-to-grave” which also includes transport to the consumer,

usage throughout the product's lifetime, and end-of-life disposal. The next step is performing an inventory analysis, where data is collected on every input (energy, water, and materials) and output (e.g., emissions and waste) associated with each stage of the product's life cycle. The impact assessment phase interprets this data to understand the EIs, such as global warming potential, ocean acidification, or resource depletion. Finally, the results are interpreted to provide insights and recommendations for reducing the EI.

Figure 2b compiles information from publicly available product environment reports for several recent mobile devices. The results show that for many mobile devices, the cradle-to-gate emissions of manufacturing and raw materials comprise 70-80% of their lifetime carbon footprint and dominate emissions. This is because, unlike devices like servers and desktops, which are both physically larger and require constant power during their use phase, mobile devices must be portable and designed with energy consumption in mind to enable long battery life. For this class of devices in particular, sustainable design is critical for further emission reduction beyond switching to clean energy sources. As a result, in this work, we focus primarily on evaluating carbon footprints in the manufacturing or cradle-to-gate phase.

How is an LCA computed? We summarize an example of LCA Methodology published by Amazon Devices Sustainability as a case study of how this process is typically conducted by a domain expert when evaluating the cradle-to-gate segment [2]. The high-level block diagram is illustrated in Figure 2c. A device's bill of materials ("BOM") is first obtained from an internal Product Lifecycle Management ("PLM") system. Parts are then assigned to categories as either a mechanical ("ME") component or process versus an electronic ("EE") component. Yield losses are accounted for at each level of the BOM from a series of manufacturing and assembly processes based on a multi-level BOM structure. Manufacturing energy is computed by measuring the total energy required for a process and dividing it by the number of devices produced to determine unit energy cost. The emissions for each part E_{part_n} in the BOM are summed to determine the total emissions for manufacturing, $E_{MFG} = \sum_{n=1}^N E_{part_n}$.

To calculate the emissions, an expert must map each part to the most appropriate emissions factor available from databases such as ecoinvent, GaBi, industry sources, or academic literature. These emissions factors are generally industry average estimates unless specific supplier data is available. Data availability, however, is a persistent problem in LCA [13]. An emissions factor for the IC is chosen based on information available about die size, process technology node, and package type. When data is unavailable, the die size is measured for a device's top five critical ICs and modeled with a conservative emissions factor. Similarly, PCB production is based on total board area, number of layers, and/or mass. Components such as ICs, PCBs, capacitors, and resistors are scaled by mass or area. Scaling by mass is performed by multiplying the component mass m by an emissions factor ef and dividing by the reference mass m_{ref} and loss factor L to account for waste in the production process,

$$E_{part,EE} = \frac{m \times ef}{m_{ref} L}.$$

A similar scaling is applied when using area. The lack of information about IC production presents a key source of uncertainty in this method. Experts from the sustainability science team must thoroughly review the resulting data to evaluate data quality issues.

Why is electronics LCA hard? While large companies are able to produce these environmental reports for individual high-value products, it is a challenging and time-intensive process for multiple reasons. The first major challenge in electronics LCA is the sheer number of manufactured parts within a device. For example, a device may be composed of over 500 different components across multiple PCBs [40]. Within a device like a laptop or smartphone, some parts are themselves sub-assemblies, such as an SSD which has multiple ICs on it. These components include various ICs, resistors, capacitors, inductors, and more. Computing a thorough LCA requires evaluating the impacts of each of these individual parts. Second, individual components are manufactured parts that result from a highly complex and resource-intensive fabrication pipeline. For example, IC fabrication is a complex multi-step process involving photolithography, etching and deposition processes, and more. Third, the manufacturing processes may differ significantly even within classes of parts such as ICs or capacitors. This

makes it difficult to use a single generic process model for each. For example, capacitors use a wide variety of different dielectric materials. Different ICs may use different semiconductor materials optimized for high performance (e.g., GaAs for high-frequency RF chips), and certain process technology nodes (often denoted by the minimum feature size, 65 nm, 7 nm, etc.) require fundamentally different technologies with significantly higher environmental costs.

These fundamental challenges make it very difficult to evaluate LCA during the design phase. This is compounded by the fact that there is a significant disconnect between tools created to perform LCAs and EDA tools for electronics designers. Traditional LCA tools such as GaBi do not have interfaces to interact with PCB design software and operate fundamentally differently. For example, although GaBi has a database of electronic parts, because of the complexity and variety of semiconductor fabrication processes, its models are industry averages not specific to a chip and sorted instead by package type. Additionally, PCB design tools such as KiCAD and EAGLE do not have any way to output the information required to compute sustainability metrics such as a chip's die size. This combination makes it a highly manual process requiring significant domain expertise to map a component to one's best guess or approximation in a database. In addition to requiring domain expertise, software such as GaBi is proprietary with licensing fees exceeding \$20,000, making it highly inaccessible to electronics designers who seek to optimize for sustainability.

3 RELATED WORK

Assessing environmental impacts has been explored in multiple domains. We present a survey of related work on relevant topic areas below.

3.1 Life Cycle Inventory

Life cycle assessment (LCA) comprises two phases: Life cycle inventory (LCI) and Environmental impact assessment (EIA). LCI serves as the foundation of LCA [31]. It involves the compilation and quantification of inputs, outputs, and the potential EIs for a given product throughout its life cycle [49]. Therefore, the generation of LCI involves comprehensive data collection about all inputs (e.g., raw materials and energy) and outputs (e.g., emissions and waste) associated with a product's life cycle, which is often the most costly and time-consuming phase [53]. Different approaches for generating LCI have been explored for decades with two main focuses: accuracy and system boundary completeness [34]. Beyond the accuracy of LCI databases, system boundaries are important to determine which processes and supply chain pathways are included in the LCA study. Process-based LCI examines every process involved in the life cycle, often used in the studies of products, such as energy products [28] and electronic products [4, 47]; input-output-based LCI connects the economic IO model to LCA, offering a broader perspective [38]; hybrid approach combines elements of these two, offers greater accuracy and boundary completeness at the cost of greater time and complexity. Despite decades of research, the generation of LCI remains complex and time-consuming and requires a significant amount of manual work.

3.2 Environmental Impact Assessment for ICT

A number of works have attempted to estimate the overall impact of the ICT industry as a whole [5, 23]. Additionally, others have developed tools to model embodied and operational emissions, corresponding to the production and use stages of computing devices. The dominating source of emissions, as discussed by Gupta et al. [29, 30], is shifting from operational activities towards hardware production for many devices.

Computing devices typically consist of numerous and varied components, each class of which includes a diverse range of materials and processes involved in the production which complicates LCA computation [21]. In response to these challenges, there has been a growing body of work focused on developing modeling tools. Much of this work, however, has been concentrated on larger-scale applications, such as data centers [1, 22, 52],

or IC level designs [29]. We seek to complement these works by developing a tool for circuit board-level designs to empower designers in the ubiquitous and mobile computing communities to consider sustainability metrics.

3.3 Comparative Life Cycle Assessment

Prior work on sustainable development has explored the idea of performing comparative LCA as a decision-making tool, enabling stakeholders to evaluate and compare the EIs of various alternatives and successors [39, 54]. Comparative LCA is also crucial for making public comparative claims. Despite its importance, these works have only done this by evaluating two complete LCAs and comparing the resulting outputs. This approach to comparative LCAs presents several challenges. First, the cost and time associated with full LCA studies are often prohibitive, particularly in industries like electronics characterized by rapid product development cycles. Second, the complexity of modern supply chains, especially in sectors like electronics, makes the quantification of all processes a time-consuming and resource-intensive task [2, 14, 51]. Third, the requirement for extensive and often confidential data makes full LCAs daunting.

In response to these challenges, recent work developed various methods to streamline the LCA process, i.e., simplified or screening LCA. For instance, certain analysis, known as cradle-to-gate, gate-to-gate, etc, ignores certain upstream or downstream processes [27]. Others focus on a narrower range of evaluated environmental metrics [12], which streamlines the inventory phase by reducing the inventory parameter considered in the selected categories. Other approaches simplify the analysis by using mass as a coarse-grained indicator of raw material used [35, 42] to perform quicker assessments.

Furthermore, traditional LCA is susceptible to variability in LCA methodologies and software tools between studies. Inconsistencies in results can be attributed to differences in:

- **System boundaries.** System boundaries determine which processes and life cycle stages are included in the assessment. Even one stage difference can lead to significant variations in the results.
- **Database differences.** If two studies use different data sources, for example, choosing thermal power and nuclear power electricity within the same region could cause 164x of difference [20]. Variations across states and even countries can also cause large differences in the calculated LCA results.
- **Methodological differences.** LCA methods can vary in terms of their assumptions. Differences in factors like allocation methods, emissions factors, and modeling assumptions can lead to variations in results.

This further complicates the LCA comparison between different designs. We propose a different approach that enables decision-making through relative comparisons between different stages of designs. Canceling out common components between design iterations dramatically simplifies the number of parts to evaluate, and we can then develop ways to reason about which components have more environmental impact.

3.4 Sustainable Computing in HCI and UbiComp

The field of HCI and UbiComp has recently seen an increased interest in sustainable computing. This trend aligns with the broader societal shift towards environmental consciousness. Previous research has focused on unmaking [16, 48], the development of biodegradable [11, 17, 50] and recyclable materials [55] for substrates, and innovative design tools that facilitate the reuse of electronic components [40]. It is worth noting that *EcoEDA* [40] aims to extend the lifespan of electronic parts, thus mitigating the EI of electronic waste.

Most existing research in sustainable computing has been directed toward the end-of-life stage of electronics and primarily focused on new materials. While novel materials are an important part of holistic sustainability solutions, the EI of a PCB is also significantly influenced by decisions made during its design stage, including the choice of ICs, size, and the integration level of components. Kaebemick et al. [35] emphasize the importance of incorporating environmental considerations into the early design stages. Our work aims to fill the gap in tools to




Fig. 3. Our comparison algorithm uses parts either for pairwise comparisons (grey nodes) via heuristics (edges) or for carbon footprint comparison (green nodes). (1) The initial comparison results are first delivered to the user. (2) Drawing from the user's domain knowledge, the user adds comparison rules using any parts available to solve the unmatched parts (orange nodes). (3) The results can be updated, taking into account the user-defined rules.

evaluate EI and empower the computing research community to integrate sustainability metrics into the PCB design stage.

4 SYSTEM OVERVIEW

This section presents an overview of DeltaLCA, a system to assess the relative EI between two input designs $\mathcal{A}$ and $\mathcal{B}$. DeltaLCA can be used at various stages in the design process where $\mathcal{A}$ and $\mathcal{B}$ could be individual components, partial designs, or complete designs. We note that the specific tools and UI below focus on the use case where a designer inputs complete PCB design files from EDA software to enable seamless integration with typical workflows. The differences between the traditional way of comparing two designs and DeltaLCA are illustrated in Figure 1.

Traditional LCA workflow. As described in Section 2, a traditional LCA is performed in two phases, the inventory phase and the assessment phase (green and orange boxes), respectively, in Figure 1. Both phases are entirely performed by LCA experts. In the inventory phase, the expert lists all relevant parts of the design file to establish a Complete Parts List. Then, for each part in that list, the expert needs to fill in EI data from an LCA database, such as ecoinvent [53]. Parts, especially recent electronic parts, are often not readily available in these databases, and the expert often needs to search for the best approximation using their domain knowledge, making it impossible to automate this process. In the assessment phase, the expert imports additional processes and establishes a flow diagram in an LCA software, such as GaBi. Creating a flow diagram requires an understanding of both the LCA-specific software and knowledge about the secondary processes and resources needed to manufacture parts. This entire process is done for both input designs, resulting in two EI results. Finally, once the expert has these numbers, they can compare the two designs.

DeltaLCA workflow. In DeltaLCA, we completely automate the inventory phase. We then replace the two assessment phases with a single *comparative LCA phase*, which requires minimal user feedback.

For the inventory phase, we make use of both insights from Section 1 to automatically establish a Complete Parts List and a Partial LCI. Focusing on creating a domain-specific tool, in our case a tool for electronic designs, we can leverage our domain knowledge to build an LCA-targeted parser of design files (Insight 2). For each part,

we automatically parse all the information available in public datasets and auxiliary sources. The result of this step is a Partial LCI (see Section 5).

A Partial LCI contains three kinds of parts:

- (1) Parts which are identical and which are present in both $\mathcal{A}$ and $\mathcal{B}$. We can cancel these parts out before the comparison algorithm.
- (2) Unique parts for which we have sufficient LCA information to compute their EIs.
- (3) Unique parts for which we have only partial information.

Parts from type (1) are removed during pre-processing, whereas parts from type (2) are usually encountered in a traditional LCA. To compare parts with partial information (3), we use expert-designed *heuristics*, which allow us to determine that part a_i from $\mathcal{A}$ will have greater EI than part b_j from $\mathcal{B}$ without ever computing the actual EI of either part. Intuitively, these heuristics can be thought of as conservative rules, e.g. a heuristic may state that if two ICs are manufactured with the same processes but one is significantly larger than the other, the larger one will have greater EI. This means that the key algorithmic challenge for reasoning about parts with partial information is to perform a *pairwise matching* between two sets of parts using these heuristics. Importantly, this matching algorithm must also optimally leverage the information that enables computing full EIs for some of the parts. We tackle this challenge with a comparison algorithm which is explained in Section 6.

User-in-the-loop. The output of the comparison algorithm is two design subsets $\mathcal{A}_\delta$ and $\mathcal{B}_\delta$ (nodes except for the orange ones in Figure 3) for which we can show that the EI of the former is bigger than the EI of the latter. For the remaining parts (orange nodes), i.e. $\mathcal{A} \setminus \mathcal{A}_\delta$ and $\mathcal{B} \setminus \mathcal{B}_\delta$, our algorithm either has insufficient information about the parts, or design $\mathcal{A}$ is simply not more environmentally impactful than design $\mathcal{B}$. In both cases, the user is presented with the comparison result and the remaining parts. To refine the assessment, they can then provide more information to the system by completing part information that the inventory phase has been unable to retrieve or by providing additional, user-defined rules about the relative impact between parts. Once more information has been provided, either about the parts themselves or about the relationship between parts, the comparison result can be updated. Our interactive feedback loop is illustrated in Figure 3: (1) the two designs are automatically compared, and the results are delivered to the user; (2) the user can provide additional information to refine the comparison; (3) the comparison result gets updated, taking into account the new information.

In summary, DeltaLCA automates the inventory phase and the assessment phase as much as possible by leveraging domain-specific knowledge and by canceling out uncertainties through part comparisons. As opposed to the traditional LCA comparison, with DeltaLCA the user is only needed at the end for the actual comparison and not throughout the entire process.

5 AUTOMATED LIFE CYCLE INVENTORY

The first step in our DeltaLCA pipeline is generating an LCI from PCB design files. Traditional methods for LCI require time and expertise for two key reasons. First, the component models available in LCA software databases do not have a one-to-one mapping to the parts on a PCB which prevents automation. Second, data for many parts is unavailable.

To achieve the desired level of accuracy and automation, we implemented a four-phase pipeline: 1) parsing the raw design files locally. This step aims to extract basic information about the components and board configuration within the PCB design relevant to LCA. 2) interfacing with external data sources to enrich our resulting part lists with up-to-date attributes of each component. 3) inferring core specifications, typically kept confidential by manufacturers, with all available information. These specifications, such as the die area of an IC, are crucial for assessing the EIs of ICs. 4) estimating the carbon footprints of standard components with complete information. We describe each phase in detail below.

5.1 Create Parts List

Generating parts list from raw PCB design files and accurately categorizing components, while simultaneously extracting detailed information such as size and series, is challenging. We note that computing LCA requires information beyond a typical BOM, which is designed for purchasing parts and fabrication and does not contain information about device size or specifications. Generating an accurate parts list is inherently complex due to several key factors: 1) even when PCB designs are created using the same design software, such as EAGLE, the components library can significantly differ between projects, resulting in diverse and often customized component symbols and footprints. This inherent variability makes establishing a uniform and robust parsing algorithm challenging. 2) The physical packaging of electronic components can differ significantly, making it challenging to employ a one-size-fits-all algorithm for component categorization. Moreover, classifying between similar components, such as different passive components (e.g., resistors, capacitors, inductors) versus IC packages, necessitates a high degree of domain-specific knowledge. 3) Component information is scattered throughout the design files, requiring sophisticated data extraction pipelines to gather the relevant data. 4) Detailed electrical specifications of components, such as power consumption, are often not included in the PCB design files. This lack of comprehensive data necessitates the integration of external databases. We focus on three primary EI sources in our implementation: ICs, PCB substrates, and other electronic components.

5.1.1 Parse Raw PCB Design File. Our automation pipeline is implemented in Python, starting with extracting pertinent information from the raw PCB layout files. We choose to start with layout files to streamline the LCA process for users, as PCB layout represents the most fundamental stage in electronics design. We note that in modern PCB EDAs, the layout file is synchronized with the schematic. Furthermore, the layout file contains additional information on board dimensions, which is necessary for assessing the EI of the substrates. The parsing method involves several critical steps. First, we calculate the board's dimensions by identifying and processing the coordinates of the outermost wires in the PCB plane. Then, we assess the number of layers by enumerating the activated routing layers between the *Top* (which contains the copper on the top of the board) and *Bottom* layer markers, indicative of a multilayer PCB structure. The core functionality of parsing begins with systematically iterating through each element in the PCB design, and the goal is to sort all electronic components into four main categories: passive components (e.g., resistor, capacitor, inductor), active components (e.g., diode, transistor), ICs (e.g., microcontroller, sensor), and misc (e.g. connectors).

Drawing on research in electronic components guidelines and standard topologies, we propose the following heuristics for the initial categorization:




Fig. 4. Our automated LCI pipeline. We start by parsing the PCB design files from common design software into a parts list, then inferring core specifications like die sizes leveraging online resources, and finally generating the partial inventory with partial EIs based on publicly available data.

- **Element name prefix.** We can effectively categorize SMD resistors, capacitors, and inductors into their respective classes by recognizing the prefixes ‘R’, ‘C’, and ‘L’, respectively. These are followed by numeric size identifiers such as ‘0402’ or ‘0603’, which signify the length and width of the package. We employ a set of regular expressions that scan for characteristic prefixes in the designators of components in the raw PCB layout.
- **Footprint topology.** The physical layout and pin configuration of components also provide significant clues. Specific topology and pad positions are often unique to certain component types. For instance, discrete surface mount transistors are commonly associated with three pins where the three pads in the PCB layout may form a triangular shape; operational amplifiers (op-amps) are often available in 8-pin DIP or SOP packages and have symmetrical pin layouts; microcontrollers are characterized by a higher pin count, typically having 16 pins or more.
- **Component correlation.** We analyze the network of connections to further aid in their identification. For example, microprocessors and microcontrollers, which usually serve as the main processing unit on a PCB, can be differentiated by examining the highest density of connections to their footprints; IC sensors can be determined through the connections of bus lines I2C (SCL and SDA) or SPI (SCLK, MOSI, and MISO) traces to the microcontroller. This allows us to distinguish between various ICs and other components that cannot be classified solely based on their element information.

We also determine the core package size of each IC by analyzing the *tDocu* or *tPlace* layer, i.e., the silkscreen. We calculate the wires’ extremities that define the components’ positioning in the PCB layout. We exclude components such as fiducial holes and screw terminals from the parsing process to maintain a focus on electronic components. All parsed information is stored with a set of attributes, including component name, package type, package area, and a quantity which is incremented for each duplicate component.

5.1.2 Call External API. For each IC, we initiate a keyword search via Digikey’s Product Information API, with the part name as the search query. Integrating real-time API calls to comprehensive online databases of electronic component distributors like Digikey enriches the parts list with additional up-to-date technical attributes, such as operational frequency and number of GPIO, that are not typically included within the layout design files themselves. This also enables cross-validation for the information extracted from the user files to resolve custom component libraries.

In the case of empty returns, we attempt modified fuzzy searches by stripping off the suffix to accommodate different naming conventions and improve the search completion rate. Upon successfully receiving a response, DeltaLCA parses the detailed attributes of the top-matched search result returned. We extract the key attributes of ICs, such as supplier device package, memory size, and number of pins from the product’s attributes. These attributes will be used to infer the specifications of the die, which will be discussed in the section below. Additionally, a fuzzy string matching algorithm is applied to accurately map and retrieve product attributes that may be represented by various terminologies across different data entries.

In [Table 1](#), we compare our method against a range of existing BOM generation tools. The baselines include the built-in BOM generators within commonly used PCB design software such as EAGLE and KiCAD, as well as popular open-source tools like KiCost [36] which can extract component data from several distributors’ web servers and InteractiveHtmlBom [33] which has robust local file parsing capabilities. Our parser is the only one that can extract all of the available part information needed to compute LCA. Moreover, we conducted cross-validation comparing the BOM generated by our method against these baselines. We found that our method consistently performs as well as or outperforms the best baseline in each category. However, we note that in cases where components have null information or users have substituted components with identical footprints but failed to update the component attributes, all methods, ours included, fail to parse these components accurately.

BOM Generator	Device Name	Category	Package	Footprint Size	Detailed Specs	Quantity
EAGLE	✓	✓	✓	✗	✗	✓
KiCAD	✓	✗	✓	✗	✗	✓
InteractiveBom	✓	✗	✓	✓	✗	✓
KiCost	✓	✗	✓	✗	✓	✓
Ours	✓	✓	✓	✓	✓	✓

Table 1. Feature comparison of our method against existing BOM generation tools.

5.2 Generate Partial Inventory

Having established a robust pipeline for local design file extraction with external API integration, we next focus on inferring die specifications in ICs and then generating a partial inventory.

Integrated Circuits, such as microcontrollers, represent the most substantial portion of the environmental footprint in the device manufacturing stage and are more significant than passive components like resistors and capacitors or the PCB substrate itself [11, 55]. This predominance stems from the high energy cost for semiconductor fabrication plants and raw material cost during the complicated fabrication processes [29].

The die size and process technology node of IC are two main indicators of the environmental footprint from the manufacturing stage because 1) die size determines the amount of raw material used; 2) the process technology node determines the complexity of fabrication: the lithography equipment, patterning process, and associated energy consumed during fabrication. More advanced (smaller) technology nodes and larger die sizes also typically lead to lower yields [41], further increasing the fraction of the fabrication energy and resources used per chip. Consequently, inferring the die size and process technology node of ICs within the desired error range is crucial for estimating the environmental footprint of a PCB design.




Fig. 5. (a) Performance by package type of our die size estimation in Mean Absolute Percentage Error (MAPE) ($N \geq 3$). QFN: Quad Flat No-lead; QFP: Quad Flat Package; DIP: Dual In-line Package; BGA: Ball Grid Array; SSOP: Shrink Small Outline Package; SO: Small Outline. (b) Trends of process technology nodes across 15 popular MCU series. Each data point represents an MCU series, labeled with its process node in nm, plotted against its maximum operational frequency (MHz) and active current consumption normalized by frequency (mA/MHz).

5.2.1 Infer Die Area. In practice, calculating the precise die size of an IC often requires access to the original design files from semiconductor packaging suppliers. However, these types of proprietary information are typically not readily disclosed by design houses, posing a significant barrier. Alternative methods such as decapsulation or three-dimensional imaging [32, 37] can precisely determine die size; however these methods require both specialized equipment and access to the physical part making them challenging to scale.

We aim to develop a generalized alternative method for approximating die size with readily available information: the dimension of the IC package and package type. The dimensions of the package and its structural characteristics determine the maximum silicon area it can house. For example, QFP packages usually have a traditional lead frame with leads extending from all four sides of the package, requiring space for the metal contacts; QFNs do not have leads protruding from the package, and instead, they have an exposed metal pad on the bottom surface; WLCSPs integrate leads within the chip fabrication process, making these a close approximation of the bare silicon die size.

Using the GaBi Electronics Extension dataset, we compile a table of die area coefficients, which represent the percentage of each package type occupied by actual silicon. This enables us to rapidly estimate the die area directly from the package size and type with sufficient accuracy.

Using a dictionary of specific package types and their corresponding die size coefficients, we scale each package size by its coefficient to obtain an estimated die area. The package is assumed to be square for packages with only one dimension available, and the area is calculated accordingly. A fuzzy string matching algorithm is applied to determine the best match for the package type within our coefficient dictionary, accounting for any variations in naming conventions and missing package info in the dictionary.

To evaluate our method, we used the Mean Absolute Percentage Error (MAPE), a normalized error metric, as our primary measure. This choice is relevant because the die sizes differ across different MCU series, resulting in varying scales of ground truth values. As shown in Figure 5a, our method demonstrates MAPE of 50.20%, 76.22%, 7.93%, 69.65%, 42.81%, 69.83% for QFN, QFP, DIP, BGA, SSOP, and SO packages, respectively. These values indicate the relative deviation of our estimated die areas from the actual values. Notably, more complex packages like QFP and BGA exhibit higher indeterminacy due to their complex internal structures.

5.2.2 Infer Process Technology Node. To infer the process technology node used in an MCU, we establish a correlation between the maximum operational frequency and the active current consumption. As the process node size decreases, which implies smaller transistor dimensions, the gate capacitance is reduced. This reduction is proportional to the transistor's switching time due to the diminishing resistive-capacitive (RC) delay, and lower power consumption by minimizing the power dissipation through the channel [46]. Consequently, smaller process nodes enable higher operational frequencies while concurrently decreasing active current.

Such dual observations act as a heuristic to infer the process technology node. For instance, an MCU operating at a higher frequency with a relatively lower active current than another MCU with similar functionality indicates the use of a more advanced process node. Figure 5b shows the process technology node for 15 commonly used MCU series with publicly available data in relation to their operational frequency and corresponding active current consumption normalized by frequency. We use this data to construct a lightweight classification model to estimate the process node by fitting the known specifications of an MCU with established trends in semiconductors. This method leverages empirical data from semiconductor technology advancements to provide a non-invasive estimation technique that does not rely on direct semiconductor foundry data.

5.3 Estimate Partial Environmental Impact

To facilitate a standardized comparison of the environmental footprint in PCB design, we first estimate the cradle-to-gate carbon footprints of standard fixed components such as resistors, capacitors, inductors, and substrates,

that are commonly present in every PCB, in CO2 equivalent. We aim to streamline the process, particularly for users who are new to LCA.

Despite the variability in the materials used for common passive components (e.g., capacitors can be made from multilayer ceramic, tantalum, or electrolytic dielectric layers), prior work has shown their EIs are relatively similar [47]. Furthermore, compared to major footprint sources like PCB substrates and ICs, the environmental footprint of these passive components is small. Consequently, assigning a single LCA value per component does not result in significant errors.

In our assessment, we standardize the LCA values for these components based on their package size and the results are shown in Table 2. The baseline values are derived from previous research [47] with a functional unit of 1 kg for each component type (e.g., 1 kg MLCC capacitors equates to 97 kg of CO2-eq). For resistors, the environmental footprint ranges from 0.01 grams of CO2 equivalent per piece for 0201 package to 0.6 grams for 0805. Capacitors follow a similar trend, where the EI increases with package size, starting at 0.036 grams of CO2 equivalent for the 0201 size, and rising to 1.067 grams for 0805. Inductors range from 0.022 grams of CO2 equivalent for 0201 to 1.358 grams for 0805. For the PCB substrate material, we defaulted to 1.6 mm thick FR-4, the most commonly used substrate, which has an EI of 0.006125 grams of CO2 equivalent per square millimeter per 1mm-thick layer [39, 44, 55]. The number of substrate layers is equal to the routing layer minus 1. This assumption is due to the absence of substrate material information in PCB design files because users typically determine the material, thickness, and other substrate specifications at the last stage, i.e., manufacturing quote. However, it is noteworthy that users could select or customize different substrate materials according to their specific requirements. For ICs with complete information including die area and process node, we also estimate their carbon footprints using the aggregate carbon footprints per area from Gupta et al.'s work [29].

Package (Imperial)	Resistor	Capacitor	Inductor
0201	0.010	0.036	0.022
0402	0.040	0.146	0.078
0603	0.120	0.611	0.330
0805	0.600	1.067	1.358

Table 2. Standardized environmental impacts in gram CO2-eq for common passive electronic components (resistors, capacitors, and inductors).

6 COMPARATIVE IMPACT ASSESSMENT

6.1 Overview

The comparison algorithm takes the partial LCIs from the inventory phase for two designs $\mathcal{A}$ and $\mathcal{B}$. Without loss of generality, we assume that we want to show that the environmental impact, e.g., carbon footprint, of $\mathcal{A}$ is *bigger* than $\mathcal{B}$. In general, our algorithm will be able to show this for two subsets $\mathcal{A}_\delta$ and $\mathcal{B}_\delta$. Our primary goal is to design an algorithm where $\mathcal{B}_\delta$ is maximized, thereby enhancing the likelihood of demonstrating that $\mathcal{B}$ has the lowest environmental impact with minimal or no additional user input (i.e. when $\mathcal{B}_\delta \equiv \mathcal{B}$).

As mentioned in Section 4, at this stage, partial LCIs contain parts with complete LCA information for which we can perform direct carbon footprint comparison and parts with only partial information which will be compared using pairwise heuristics. But how do we combine heuristic-based comparisons with carbon footprint-based comparisons? One straightforward approach would be first to perform an LCA comparison for parts with complete information, so parts that have carbon footprint information, and then, in a second stage, perform a pairwise

comparison based on heuristics. However, this could easily lead to selection problems. Suppose $\mathcal{A}$ has only parts with complete information and $\mathcal{B}$ has parts with partial information. The first step would consume all parts of $\mathcal{A}$, even if not all of the per-part carbon footprints are necessary to weigh up against the carbon footprint of $\mathcal{B}$. Then, $\mathcal{A}$ would not have any parts left for the pairwise comparison. This example shows that the selection of parts used for the carbon footprint comparison and the pairwise, heuristic-based comparison are interdependent. We formulate this interdependent selection problem as an integer program, as described in Section 6.3. Finally, we discuss our implementation and evaluate the performance in Section 6.4.

6.2 Heuristics

We define a *heuristic* H as a map from a subset of parts of design $\mathcal{A}$ and a subset of parts of design $\mathcal{B}$ to a decision about which subset has a higher carbon footprint, or if neither of them subsumes the other. The heuristic decision is based on certain part attributes, and we only assign heuristics between parts with the necessary attributes available. Heuristics are based on domain knowledge of the application domain and example heuristics to compare PCB components are listed below. We note that these can be extended and modified by the user based on available information:

- (1) **Package size for the same type of passive components.** For passive components, such as capacitors and resistors, those of identical dimensions consume comparable amounts of raw materials, making their environmental impact comparable, irrespective of variations in material composition.
- (2) **Chips with the same core and architecture.** Chips that share a common standard core and architecture (e.g., ARM Cortex-M), are considered equivalent in our analysis if they differ only in packaging type. This is underpinned by two factors: the environmental impact of the package is relatively overshadowed by the silicon part, and a shared core implies identical manufacturing processes and material usage, thereby equating their environmental footprints.
- (3) **Equivalent process.** In cases where components are analogous in dimensions and weight, we assume that their transportation-related environmental costs are comparable. Furthermore, for components of the same category, the electricity cost in the assembly process, such as soldering, is consistent.
- (4) **Chip die size.** The impact of chip die sizes is twofold. Larger dies not only suggest increased raw material consumption but also tend to result in lower production yields[15, 41], both contributing to a larger environmental impact.
- (5) **Process technology node.** More advanced (smaller) process technology node requires more steps and complexity in fabrication. This includes the need for advanced lithography equipment and multiple patterning processes, culminating in elevated energy consumption.
- (6) **Diode size.** size is an effective indicator of raw material consumption across various diode types, such as photodiodes, LEDs, and Schottky diodes. Given their similar semiconductor structures, the size of the diode correlates directly with the environmental costs associated with raw materials.

We use the first three heuristics to identify parts in $\mathcal{A}$ and $\mathcal{B}$ which are "identical", i.e., which parts we assume to amount to the same carbon footprint, given the partial information available to use. These parts are being removed before the comparison algorithm.

Note that we might encounter a situation where two heuristics evaluate to conflicting results. Two heuristic edges h_i and h_j conflict with each other if both consider the same subset of parts from $\mathcal{A}$ and the same subset of parts from $\mathcal{B}$ and if h_i and h_j make different decisions about which subset of parts has a higher carbon footprint. For example, if between a chip a_i and a chip b_j , a_i has a bigger die size (heuristic 4), but b_j has a smaller process node (heuristic 5), conflicting directed edges arise: heuristic 4 would conclude that $a_i > b_j$ while heuristic 5 would decide $b_j > a_i$. We have no means of reasoning about them without further information. We therefore remove all conflicting heuristic edges between parts before our comparison algorithm.




Fig. 6. Input to this example problem is the bipartite graph, including the heuristic edges (solid lines represent selected heuristics and unselected are dashed) and the carbon footprint weights, see Table 3 for detailed values and variable assignments. The example constraints illustrate the three constraints from our integer program formulation.

Carbon footprints $\mathcal{A}$	Carbon footprint variables $\mathcal{A}$	Carbon footprints $\mathcal{B}$	Carbon footprint variables $\mathcal{B}$	Part variables $\mathcal{B}$	Heuristic variables
$w_{a_1} = N/A$	$c_{a_1} = 0$	$w_{b_1} = N/A$	$c_{b_1} = 0$	$b_1 = 0$	$h_{22} = 1$
$w_{a_2} = N/A$	$c_{a_2} = 0$	$w_{b_2} = N/A$	$c_{b_2} = 0$	$b_2 = 1$	$h_{23} = 0$
$w_{a_3} = 10$	$c_{a_3} = 0$	$w_{b_3} = 10$	$c_{b_3} = 1$	$b_3 = 1$	$h_{34} = 1$
$w_{a_4} = 10$	$c_{a_4} = 1$	$w_{b_4} = 10$	$c_{b_4} = 0$	$b_4 = 1$	$h_{44} = 0$
$w_{a_5} = 10$	$c_{a_5} = 1$				

Table 3. Carbon footprint values and variable assignments for the example problem from Figure 6. Parts for which we do not have complete LCA information have no available carbon footprint. The binary variables are assigned such that the objective function is maximized.

6.3 Comparison Algorithm

In the following, we explain the general integer program formulation for the comparison of two designs. Please refer to Figure 6 for a toy example.

Objective function. The goal of the selection problem is to select parts of $\mathcal{A}$ and parts of $\mathcal{B}$ such that each part of $\mathcal{B}$ is shown to be of lesser carbon footprint than some part of $\mathcal{A}$, either through direct carbon footprint comparison or by selecting pairwise heuristics. More specifically, we formulate this selection problem as a binary integer program where we associate to each part b_i a binary variable $b_i \in \{0, 1\}$. b_i is equal to 1 if the part b_i has been included in the LCA comparison. The objective function is shown in Eq.1. We are maximizing the number of compared parts of $\mathcal{B}$ since we want to do a comparison for as many parts as possible of that design.

$$\max \sum_i b_i \quad (1)$$

Constraints. b_i should only be put to 1 if it has been taken into account by a pairwise heuristic comparison or by means of its carbon footprint, which we translate via the constraint in Eq.2.

$$(C1) : b_i \leq \left(\sum_{h \in \mathcal{H}_{b_i}} h \right) + c_{b_i} \quad (2)$$

where $\mathcal{H}_{b_i}$ is the set of all binary heuristic variables for heuristics which show that the part $\mathcal{B}_i$ has been subsumed by some part of $\mathcal{A}$ and c_{b_i} is a binary variable which is equal to 1 if $\mathcal{B}_i$ has been used in the carbon footprint

computation. In the toy example in Figure 6, b_2 has been successfully put to 1 because the heuristic h_{22} has been selected and b_2 's carbon footprint has not been used.

As mentioned earlier, each part $\mathcal{A}_i$ of $\mathcal{A}$ can only be used in one heuristic. But if an $\mathcal{A}_i$ has already been used in the carbon footprint computation, it cannot be used for pairwise heuristic comparison. This is enforced by the constraint in Eq.3.

$$(C2) : \sum_{h \in \mathcal{H}_{a_i}} h \leq 1 - c_{a_i} \quad (3)$$

where $\mathcal{H}_{a_i}$ is the set of all binary heuristic variables for heuristics which use part $\mathcal{A}_i$ and c_{a_i} is a binary variable which is equal to 1 if $\mathcal{A}_i$ has been used in the carbon footprint computation. Intuitively, if c_{a_i} is equal to 1, then the right-hand side of the constraint is equal to 0 and no heuristic for $\mathcal{A}_i$ can be selected. If the carbon footprint of $\mathcal{A}_i$ has not been used for comparison, then c_{a_i} is equal to 0, the right-hand side is equal to 1 and at most one heuristic for $\mathcal{A}_i$ can be selected. The effect of this constraint can be observed in Figure 6, where h_{44} could not have been selected because c_{a_4} has already been set to 1.

Finally, we have to include the carbon footprint comparison. For each part with complete information, we assign a weight w_i which is its actual carbon footprint. To show that the carbon footprint of $\mathcal{A}$ is bigger than the carbon footprint of $\mathcal{B}$, we must ensure the weighted sum of the parts selected for carbon footprint comparison of the former is bigger than the weighted sum of the latter, which is expressed via the constraint in Eq.4.

$$(C3) : \sum w_{a_i} c_{a_i} \geq \sum w_{b_i} c_{b_i} \quad (4)$$

In the toy example in Fig.6, the left-hand side of this constraint sums up to 20, which is greater than the right-hand side which selects only $w_{b_3} = 10$.

Summary. The objective of the integer program is to maximize the sum of b_i (Eq.1). However, to assign b_i to 1, either a corresponding heuristic variable h or its carbon footprint variable c_{b_i} have to be assigned to 1 (Eq.2). Whereas a heuristic can only be selected if the counterpart in $\mathcal{A}$ has not already been used by another heuristic (Eq.3), the carbon footprint can only be used if we have enough carbon footprint in $\mathcal{A}$ to counterbalance the use of $\mathcal{B}_i$ (Eq.4). And we ensure that resources in $\mathcal{A}$ are only used once via Eq.3.

Note that theoretically, a part $\mathcal{B}_i$ could be shown to be subsumed by $\mathcal{A}$ via multiple heuristic comparisons and/or carbon footprint comparison at the same time, i.e., we do not enforce comparison uniqueness as for $\mathcal{A}_i$ in Eq.3. However, it is not advantageous for the solver to put both kinds of variables for $\mathcal{B}_i$ to 1, since both would have to be weighed up against by parts in $\mathcal{A}$.

6.4 Implementation

Integer Program. We implemented the integer program using the Python API of OR-tools[25] and for common comparisons, the runtime is a fraction of a second. Although integer programming is NP-complete and combinatorial in nature, in practice, solutions to a problem of a reasonable size can be found efficiently. Since for each part of $\mathcal{A}$, we introduce heuristics with all compatible parts of $\mathcal{B}$, the number of heuristic variables scales quadratically w.r.t. the problem size, i.e., the number of parts. This means that the runtime also scales approximately quadratically w.r.t. to the number of parts. We evaluated the runtime of our comparison algorithm by making pairwise comparisons among eight PCB boards (the three complex designs from the case study (Section 7.1.2) and five smaller designs). As shown in Figure 7, our comparison algorithm's average runtime for a problem of a size similar to the complex design comparisons (i.e., 40 ~ 50 parts in comparison) is around 0.5 seconds and thus enables our UI to run at interactive speeds.

Looking beyond the specific examples above, we also observe that more complex designs with high part counts consist primarily of passive components such as resistors and capacitors that are simplified by DeltaLCA. For




Fig. 7. (a) The size of the integer program illustrated by the number of variables and the number of constraints for each comparison in the case study (Section 7.1.2). (b) Average runtime (on a MacBook Pro with M1 chip and 8GB RAM) in seconds over the problem size (represented by the total number of parts involved in the comparison). Here “parts” refers to the ICs in the design and two special parts generated by the automated life-cycle inventory process: passive components (such as resistors, capacitors, and inductors) and board substrates are precomputed as two single parts.

example, a sample laptop motherboard design has 443 parts, of which 353 (80%) are passives [43]. Similarly, the main electronics board for a Roomba contains a total of 633 parts, but 79% are resistors or capacitors [40]. Additionally, many devices like smartphones and tablets consist of multiple smaller PCB modules [19]. This enables rapid analysis of subsystems and the ability to evaluate a complex overall assembly in parts.




Fig. 8. System Screenshot of DeltaLCA user interface. (1) The user uploads design files using the buttons in the sidebar. (2) The user is presented with matched/unmatched components. (3) The user specifies custom comparison rules. (4) The user checks all added rules. (5) Updates the comparison results with user-defined rules taken into account.

User Interface for Comparing Designs. The UI for DeltaLCA is developed with Python's tkinter framework (shown in Figure 8) and bridges the automated LCI with the comparison algorithm. After the user uploads two design files using the buttons in the sidebar, the system initiates an inventory generation process; the inventory may be partially complete due to the incomplete component information from the design file. Then the initial comparison runs automatically and the user is presented with two tables: one displaying the matched components from designs A and B respectively, and another below listing any unmatched components. The user can analyze individual components in each design via a scrollable frame on the top right. All unmatched components are highlighted in yellow and located at the top of the list to emphasize their selection priority, in contrast to matched components shown in gray. The user can create custom comparison rules by selecting components and specifying their quantities via checkboxes and dropdown menus before the component names, and setting the comparison direction through the comparator option menu. Custom user-defined rules are added to a separate table using the 'Add' button, while the 'Del' button allows for the removal of specific rules. The user can then click the 'Update' button to re-run the comparison algorithm taking into account their custom rules. The comparison process concludes once all components, from the design deemed to have a lower EI, are successfully matched.

7 EVALUATION

Having demonstrated the technical efficacy of our DeltaLCA automated inventory and comparison algorithm and integrating DeltaLCA with a frontend user interface, we evaluated the performance of our end-to-end system with users in the loop. Our study aimed to answer the following questions:

- RQ1: Can DeltaLCA accurately calculate the environmental impact (EI) of simple PCB designs with complete information without extra intervention from users?
- RQ2: To what extent can DeltaLCA compare and analyze complex PCB designs with incomplete component information, and what is the role of the user in this process?
- RQ3: How do electronics designers perceive the importance of EI in PCB design and what are their attitudes towards incorporating tools like DeltaLCA in their workflow?

Below, we detail the methodology and findings of our two-pronged evaluation approach, encompassing both quantitative analysis and qualitative user feedback.

7.1 Case Study

Our case study focused on two key scenarios: simple electronic devices, e.g., those with fewer than 50 components, to assess the tool's direct EI calculation capabilities, and more complex PCBs with missing component information to evaluate the tool's comparison efficiency and the degree of user intervention required.

7.1.1 Complete Inventory with Full LCA (RQ1). We employed DeltaLCA to analyze the EI of two real-world designs obtained from the authors of [11]: a custom biodegradable mouse and a Dell optical mouse. To benchmark the accuracy of DeltaLCA's results, we compared them with the emissions quantified through a traditional Full LCA using GaBi. DeltaLCA successfully computed the total global warming emissions in carbon dioxide equivalent (CO₂-eq) for each design, and determined that BioMouse possesses a lower EI than the Dell mouse (results are shown in Figure 9). The carbon emissions estimated by DeltaLCA are in accordance with GaBi full LCA: the discrepancies stood only at 17.38% for the BioMouse and -46.73% for the Dell mouse. In the context of LCA research, such variances are considered within acceptable bounds. We investigate the difference and find that this results from different environmental impact coefficients for components such as PCBs and capacitors between our analyses. We note that our coefficients were determined from publicly available research papers while the LCAs from [11] were from the Gabi Electronics extension database.

Prior studies, such as those by Smith et al. [47] & Zhang et al. [54] and Liu et al. [39] & Ozkan et al. [44] acknowledge that LCA results can exhibit up to a three-fold difference even for same products because LCA is

susceptible to inconsistencies in methodologies between two studies as we described in Section 3.3 in *Related Work*.

The expert user who performed the LCA from Arroyos et al. [11] also noted that modeling the Dell Mouse with standard commercial processes required 2-3 hours while the Biodegradable mouse required 5 hours due to the complexity of modeling a new process flow for the custom water-soluble circuit board. This duration only accounts for the time to create the GaBi model and excludes the time required for new component inventory generation. In stark contrast, DeltaLCA achieves a correct comparison result and accelerates the process to a single click.

7.1.2 Complex Design with Partial Inventory (RQ2). In the realm of electronics design, assessing more complex devices is crucial, as they typically incorporate a diverse range of components, and the components often have incomplete information. To illustrate the benefit of our comparative assessment tool, we did a case study for comparing the carbon footprint of three Arduino dev boards: (1) Arduino Leonardo R3, (2) Arduino MKR FOX 1200, and (3) Arduino UNO Wifi. We choose these devices for several reasons: 1) Complexity: Unlike simpler devices like the mouse above, these development boards are akin to small computers, featuring intricate circuitry and multiple-layer design. 2) Relevance in UbiComp: these platforms are widely used for prototyping and educational purposes.

Our study involved a cross-comparison of these three devices. One expert with five years of electronics design experience, but no LCA experience, used our tool to establish a relative EI ranking of the three boards, namely $(3) \geq (1) \geq (2)$. For reference, we also collected some basic statistics about the three boards in Table A4. If we only deduce the relative ordering based on the information in this table, it is not obvious what the relative order should be. In fact, the expert user initially tried to prove $(2) \geq (1)$ with our UI but in the end could not figure out a way to do so and realized that it should be in the opposite direction. This demonstrates that our tool could also serve as a verifier to help users verify their speculation.

As shown in Figure 10, for complex devices like these, containing around 100 components and 20~30 ICs, our tool's comparative algorithm can automatically match most components (on average 88% overall, and 94% for the target, design B in this case) without any user intervention.

Figure A12 shows that typically fewer than five user-defined rules (specifically, 4, 2, and 1 for the three comparisons) are needed to complete a comparison, with an average time spent of less than 15 minutes per comparison. This is a significant reduction in time compared to traditional LCAs for simpler designs (mouse examples in the previous section). Additionally, our tool does not require LCA-specific expertise, empowering




Fig. 9. Comparing global warming emissions in carbon dioxide equivalent between DeltaLCA and traditional full LCA using GaBi for (a) a BioMouse and (b) a Dell optical mouse. Note that the y-axis is on a common logarithm scale.




Fig. 10. The matching coverage of our comparison algorithm without any user-defined rules for the three cases. Our method can automatically find a valid matching that covers most of the parts in both designs.

electronics designers to perform evaluations directly. Such efficiency demonstrates the substantial time savings DeltaLCA offers designers.

In our study on complex designs with partial inventories, we highlight the strengths of our comparison algorithm which can utilize both carbon footprint data from automated LCI when available, and also apply expert-defined heuristics when only partial information exists. Moreover, users have the flexibility to create rules for any components across the designs, not just the unmatched ones. This feature is crucial because the matching derived from the integer program may not always be globally optimal and users might have insights on specific components that aren't automatically extracted from PCB design files and online data sources.

7.2 User Evaluation & Survey (RQ3)

To deepen our understanding of the tool's utility and effectiveness in real-world scenarios, we conducted user evaluations to understand how designers perceive sustainable design and our tool. The study is composed of two parts. First, we provided a brief background introduction about EI of electronics and guided the participants through the user flow of DeltaLCA (elucidated in Figure 8). We also provide participants the opportunity to personally experiment with DeltaLCA using the three provided PCB layout files. Second, participants filled out a




Fig. 11. Left: Distribution of users who have considered the environmental impact when designing PCB. Middle: Distribution of users who are inclined to incorporate DeltaLCA into the future design workflow. Right: Distribution of users who think it is important to plug DeltaLCA into existing EDA software, on a scale of 1 to 5 where 1 represents the least importance and 5 represents the most importance.

survey. Please refer to [Table A5](#) for the complete list of our survey questions. Note that this study aims to explore designers' awareness of environmental considerations in their workflows rather than performing comprehensive usability testing of DeltaLCA. Furthermore, we sought to gauge the likelihood of DeltaLCA's integration into designers' workflows.

For our study, we recruited 6 participants: 2 aged 18-24 and 4 aged 25-34; 5 research assistants at university and 1 hardware architect in industry; 2 women and 4 men. We present the summarized outcomes of our user survey below ([Figure 11](#)).

7.2.1 PCB Environmental Impact Awareness. The first aspect of our survey investigated whether PCB designers have previously considered EI during their design process. Our results indicate a relatively low level of environmental consideration: only 33.33% reported having considered EI when designing PCBs, while the majority 66.67% had not. One of the participants who has considered EI mentioned: *"but more on the side of rare metals, not carbon footprint."*

We further probed the participants using heuristic methods by asking about their strategies for minimizing the EI of PCBs and ranking the influential EI factors in electronics design. Responses varied, participants mentioned *"reducing PCB size and number of passive components and ICs"*, *"using chips with smaller die areas"*, and *"design a PCB to be modular and reusable so that I can use it for multiple purposes"*. It is worth noting that one participant highlighted supply chain considerations: *"I would minimize the need for shipping and logistics, for example by sourcing all components from the same place."* The various design optimization approaches suggested by the participants, while touching on several major EI factors of electronic design, are not exhaustive. This underscores the need for automated tools like DeltaLCA to guide designers in making environmentally conscious decisions.

7.2.2 Future Use of DeltaLCA. The survey results revealed a strong inclination among participants towards incorporating DeltaLCA into their future design processes, while only one participant indicated uncertainty ('Maybe'). Significantly, there was no outright rejection of DeltaLCA, indicating a general openness to its adoption within the design community.

Importance of EDA integration. Even so, we noticed how crucial it is for DeltaLCA to plug in existing EDA software, as this could impact its usability and convenience. When asked about the importance of DeltaLCA's integration with existing EDAs, the majority of respondents (83.33%) rated this aspect as highly crucial (scores of 4 and 5). This finding aligns with one participant's note, *"going out of the way to evaluate this kind of an issue is too hard, but if it's built-in, I can easily check with one button."*

Desired features. When asked about additional desired functionalities in a tool designed for assessing the environmental impact of PCBs, more than half of the participants mentioned *"automatic [design/component] optimization [strategies/guidelines/suggestions] for reducing environmental impact."*

8 DISCUSSION & CONCLUSION

In this paper, we introduce an innovative method for creating design tools that facilitate environmentally-focused decision-making. This approach is grounded in the key understanding that, although Life Cycle Assessment (LCA) calculations are complex, comparative analysis can enable effective decision-making even with incomplete data. Our end-to-end tool, specifically developed for electronics design, demonstrates that leveraging domain-specific knowledge and datasets can significantly automate the process and require minimal user input. We evaluated our tool's effectiveness through case studies and a user survey.

This work presents a new approach to evaluating the environmental impact of electronics and designing for sustainability that opens multiple directions for future research. First, there is considerable potential to scale the system to more complex devices and with different inputs. For example, DeltaLCA currently requires the full design files to compare devices. Future work could explore ways to leverage our heuristic matching approach to

perform comparisons with higher-level information such as the subcomponents within a device (e.g., display, camera module, battery). However, scaling our tool to accommodate this greater complexity of components will require additional engineering efforts. Future research could improve the Automated LCI to be compatible with a broader range of electronic components beyond passive components and ICs such as display sub-assemblies like cameras and radio modules. Potential data sources include publicly available specifications of the products with their size and power consumption as well as data from public product environmental reports of similar devices that could be used to develop heuristic relationships.

Second, improvements to Automated LCI could be achieved by employing novel search algorithms, with advancements in Large Language Models (LLMs) offering possibilities to explore broader databases or to identify freely available web data. This approach could increase automation by augmenting our heuristic matching algorithm that seeks to minimize unmatched components and reduce the number of additional user rules required for making final decisions. Currently, the number of matches is used as a proxy for this goal, but future work may reveal more effective methods, possibly through deeper domain expertise or advanced data analysis.

Third, our approach could be extended to perform comparative assessments at different stages of the design pipeline. Although our current research focuses on comparing two complete designs, the concept of early-stage or abstract comparisons, such as a list of components or even a written paragraph describing the fabrication process of a device, is a promising area for future work. For example, comparing the environmental trade-offs between high-level system architectures could provide valuable information at early design stages. Additionally, zero-knowledge comparisons where design files are not accessible, could provide greater insights into Scope 3 emissions such as components from suppliers that do not provide environmental impact data. Such endeavors would likely require not only direct heuristic comparison but also probabilistic reasoning to incorporate uncertainty in the data and constraints. Furthermore, there is a wealth of opportunities to further expand our insights on comparative approaches to LCA. Future research could adapt this methodology to other domains, such as architecture, material synthesis, and garment production, among others.

In summary, this work introduces a novel approach to sustainable design with multiple avenues for further exploration. We hope that our research will both immediately enable designers in the UbiComp community to develop more sustainable design practices and encourage further investigation into LCA-focused design tools.

ACKNOWLEDGMENTS

We thank Chu Li and Shumin Yu for feedback on the figures; Richard Lin, Jeongwoo Han and Udit Gupta for helpful discussions; and Justin Shen for help with MCU data collection. This research is supported by Amazon Research Awards, a Google Research Scholar, NSF IIS 2212049, NSF IIS 2327136, NSF CCF 2319181, and Meta. Zhihan Zhang is supported by the University of Washington CEI Fellowship.

REFERENCES

- [1] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 2023. Carbon Explorer: A Holistic Framework for Designing Carbon Aware Datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS 2023)*. Association for Computing Machinery, New York, NY, USA, 118–132. <https://doi.org/10.1145/3575693.3575754>
- [2] Amazon. 2023. Amazon Devices Product Carbon Footprint Methodology. <https://sustainability.aboutamazon.com/devices-carbon-methodology.pdf>
- [3] Amazon. 2023. Product Sustainability Fact Sheet Fire HD 10 Tablet 32GB 13th Gen. https://sustainability.aboutamazon.com/devices_fact_sheet_fire_HD_10_lite.pdf
- [4] Anders S.G. Andrae. 2016. Life-Cycle Assessment of Consumer Electronics: A review of methodological approaches. *IEEE Consumer Electronics Magazine* 5, 1 (Jan. 2016), 51–60. <https://doi.org/10.1109/MCE.2015.2484639> Conference Name: IEEE Consumer Electronics Magazine.
- [5] Anders SG Andrae and Tomas Edler. 2015. On global electricity usage of communication technology: trends to 2030. *Challenges* 6, 1 (2015), 117–157.

- [6] Apple. 2020. Product Environmental Report 13-inch MacBook Pro. https://www.apple.com/environment/pdf/products/notebooks/13-inch_MacBookPro_PER_Nov2020.pdf
- [7] Apple. 2022. Product Environmental Report iPad (10th generation). https://www.apple.com/environment/pdf/products/ipad/iPad_PER_Oct2022.pdf
- [8] Apple. 2022. Product Environmental Report: iPhone 14 Pro. https://www.apple.com/environment/pdf/products/iphone/iPhone_14_Pro_PER_Sept2022.pdf
- [9] Apple. 2023. Product Environmental Report Apple Watch Series 9. https://www.apple.com/environment/pdf/products/watch/Apple_Watch_Series_9_PER_Sept2023.pdf
- [10] Apple. 2023. Product Environmental Report: iPhone 15 Pro and iPhone 15 Pro Max. https://www.apple.com/environment/pdf/products/iphone/iPhone_15_Pro_and_iPhone_15_Pro_Max_Sept2023.pdf
- [11] Vicente Arroyos, Maria L K Viitaniemi, Nicholas Keehn, Vaidehi Oruganti, Winston Saunders, Karin Strauss, Vikram Iyer, and Bichlien H Nguyen. 2022. A Tale of Two Mice: Sustainable Electronics Design and Prototyping. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3491101.3519823>
- [12] Alba Bala, Marco Rauegi, Gabriela Benveniste, Cristina Gazulla, and Pere Fullana-i Palmer. 2010. Simplified tools for global warming potential evaluation: when 'good enough' is best. *The International Journal of Life Cycle Assessment* 15, 5 (June 2010), 489–498. <https://doi.org/10.1007/s11367-010-0153-x>
- [13] Tereza Bicalho, Ildo Sauer, Alexandre Rambaud, and Yulia Altukhova. 2017. LCA data quality: A management science perspective. *Journal of Cleaner Production* 156 (2017), 888–898.
- [14] Ole Broberg and Per Christensen. 1999. LCA experiences in Danish industry: results of a survey. *The International Journal of Life Cycle Assessment* 4 (1999), 257–262.
- [15] David V. Campbell. 2010. Yield modeling of 3D integrated wafer scale assemblies. In *2010 Proceedings 60th Electronic Components and Technology Conference (ECTC)*. IEEE, Las Vegas, NV, USA, 1935–1938. <https://doi.org/10.1109/ECTC.2010.5490688>
- [16] Tingyu Cheng, Taylor Tabb, Jung Wook Park, Eric M Gallo, Aditi Maheshwari, Gregory D. Abowd, Hyunjooh Oh, and Andreea Danielescu. 2023. Functional Destruction: Utilizing Sustainable Materials' Physical Transiency for Electronics Applications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3544548.3580811>
- [17] Tingyu Cheng, Zhihan Zhang, Bingrui Zong, Yuhui Zhao, Zekun Chang, Yejun Kim, Clement Zheng, Gregory D. Abowd, and Hyunjooh Oh. 2023. SwellSense: Creating 2.5D interactions with micro-capsule paper. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3544548.3581125>
- [18] Dell. 2020. Dell XPS 15 9500. <https://www.delltechnologies.com/asset/en-us/products/laptops-and-2-in-1s/technical-support/xps-15-9500-pcf-datasheet.pdf>
- [19] P Dempsey. 2017. The Teardown: iPad 2017 tablet computer. *Engineering & Technology* 12, 5 (2017), 80–81.
- [20] Ning Ding, Jingru Liu, Jianxin Yang, and Dong Yang. 2017. Comparative life cycle assessment of regional electricity supplies in China. *Resources, Conservation and Recycling* 119 (April 2017), 47–59. <https://doi.org/10.1016/j.resconrec.2016.07.010>
- [21] Ilija Djekic, Milica Pojić, Alberto Tonda, Predrag Putnik, Danijela Bursać Kovačević, Anet Režek-Jambrak, and Igor Tomasevic. 2019. Scientific Challenges in Performing Life-Cycle Assessment in the Food Supply Chain. *Foods* 8, 8 (Aug. 2019), 301. <https://doi.org/10.3390/foods8080301> Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.
- [22] Mariam Elgamal, Doug Carmean, Elnaz Ansari, Okay Zed, Ramesh Peri, Srilatha Manne, Udit Gupta, Gu-Yeon Wei, David Brooks, Gage Hills, and Carole-Jean Wu. 2023. Carbon-Efficient Design Optimization for Computing Systems. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems (HotCarbon '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3604930.3605712>
- [23] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. 2021. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns* 2, 9 (Sept. 2021), 100340. <https://doi.org/10.1016/j.patter.2021.100340>
- [24] Google. 2023. Fitbit Charge 6 Product environmental report. <https://www.gstatic.com/gumdrop/sustainability/fitbit-charge-6-product-environmental-report.pdf>
- [25] Google. 2023. OR-tools. <https://developers.google.com/optimization>
- [26] Google. 2023. Pixel 8 Pro Product environmental report. <https://www.gstatic.com/gumdrop/sustainability/pixel-8-pro-product-environmental-report.pdf>
- [27] Katja Tasala Gradin and Anna Björklund. 2021. The common understanding of simplification approaches in published LCA studies—a review and mapping. *The International Journal of Life Cycle Assessment* 26, 1 (Jan. 2021), 50–63. <https://doi.org/10.1007/s11367-020-01843-4>
- [28] Begoña Guezuraga, Rudolf Zauner, and Werner Pölz. 2012. Life cycle assessment of two different 2 MW class wind turbines. *Renewable Energy* 37, 1 (Jan. 2012), 37–44. <https://doi.org/10.1016/j.renene.2011.05.008>

- [29] 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 (ISCA '22)*. Association for Computing Machinery, New York, NY, USA, 784–799. <https://doi.org/10.1145/3470496.3527408>
- [30] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2020. Chasing Carbon: The Elusive Environmental Footprint of Computing. <https://doi.org/10.48550/arXiv.2011.02839> arXiv:2011.02839 [cs].
- [31] Stefanie Hellweg and Llorenç Milà i Canals. 2014. Emerging approaches, challenges and opportunities in life cycle assessment. *Science* 344, 6188 (June 2014), 1109–1113. <https://doi.org/10.1126/science.1248361> Publisher: American Association for the Advancement of Science.
- [32] Mirko Holler, Manuel Guizar-Sicairos, Esther H. R. Tsai, Roberto Dinapoli, Elisabeth Müller, Oliver Bunk, Jörg Raabe, and Gabriel Aeppli. 2017. High-resolution non-destructive three-dimensional imaging of integrated circuits. *Nature* 543, 7645 (March 2017), 402–406. <https://doi.org/10.1038/nature21698> Number: 7645 Publisher: Nature Publishing Group.
- [33] InteractiveHtmlBom. 2023. Interactive HTML BOM generation plugin for KiCad, EasyEDA, Eagle, Fusion360 and Allegro PCB designer. <https://github.com/openscopeproject/InteractiveHtmlBom>
- [34] Samantha Islam, S. G. Ponnambalam, and Hon Loong Lam. 2016. Review on life cycle inventory: methods, examples and applications. *Journal of Cleaner Production* 136 (Nov. 2016), 266–278. <https://doi.org/10.1016/j.jclepro.2016.05.144>
- [35] H. Kaebemick, M. Sun, and S. Kara. 2003. Simplified Lifecycle Assessment for the Early Design Stages of Industrial Products. *CIRP Annals* 52, 1 (Jan. 2003), 25–28. [https://doi.org/10.1016/S0007-8506\(07\)60522-8](https://doi.org/10.1016/S0007-8506(07)60522-8)
- [36] KiCost. 2022. KiCost: Build cost spreadsheet for a KiCad project. <https://github.com/hildogjr/KiCost>
- [37] Tsuneo Kurita, Nagayoshi Kasashima, Hiroki Yamakiri, Nobuhito Ichihashi, Nobutoshi Kobayashi, Kiwamu Ashida, and Shinya Sasaki. 2011. Development of the new IC decapsulation technology. *Optics and Lasers in Engineering* 49, 9-10 (2011), 1216–1223.
- [38] Wassily Leontief. 1970. Environmental Repercussions and the Economic Structure: An Input-Output Approach. *The Review of Economics and Statistics* 52, 3 (1970), 262–271. <https://doi.org/10.2307/1926294> Publisher: The MIT Press.
- [39] 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 & Environmental Science* 7, 11 (Oct. 2014), 3674–3682. <https://doi.org/10.1039/C4EE01995D> Publisher: The Royal Society of Chemistry.
- [40] Jasmine Lu, Beza Desta, K. D. Wu, Romain Nith, Joyce E Passananti, and Pedro Lopes. 2023. ecoEDA: Recycling E-waste During Electronics Design. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3586183.3606745>
- [41] Samuel Naffziger, Noah Beck, Thomas Burd, Kevin Lepak, Gabriel H. Loh, Mahesh Subramony, and Sean White. 2021. Pioneering chiplet technology and design for the AMD EPYC™ and Ryzen™ processor families. In *Proceedings of the 48th Annual International Symposium on Computer Architecture (ISCA '21)*. IEEE Press, Virtual Event, Spain, 57–70. <https://doi.org/10.1109/ISCA52012.2021.00014>
- [42] N. F. Nissen, H. Griesse, A. Middendorf, J. Müller, H. Pötter, and H. Reichl. 1997. Comparison of simplified environmental assessments versus full life cycle assessment (LCA) for the electronics designer. In *Life Cycle Networks: Proceedings of the 4th CIRP International Seminar on Life Cycle Engineering 26–27 June 1997, Berlin, Germany*, F.-L. Krause and G. Seliger (Eds.). Springer US, Boston, MA, 301–312. https://doi.org/10.1007/978-1-4615-6381-5_25
- [43] Olimex. 2020. Olimex DIY-Laptop: TERES-PCB1-A64-MAIN PCB Rev.C. <https://github.com/OLIMEX/DIY-LAPTOP/tree/master>
- [44] Elif Ozkan, Nilay Elginöz, and Fatos Germirli Babuna. 2018. Life cycle assessment of a printed circuit board manufacturing plant in Turkey. *Environmental Science and Pollution Research* 25, 27 (Sept. 2018), 26801–26808. <https://doi.org/10.1007/s11356-017-0280-z>
- [45] Samsung. 2023. Product Environmental Report Galaxy S23. https://www.samsung.com/global/sustainability/media/pdf/Galaxy_S23_Environmental_Report_EN_2302.pdf
- [46] Ghavam G. Shahidi. 2019. Chip Power Scaling in Recent CMOS Technology Nodes. *IEEE Access* 7 (2019), 851–856. <https://doi.org/10.1109/ACCESS.2018.2885895> Conference Name: IEEE Access.
- [47] 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 (June 2018), 496–513. <https://doi.org/10.1016/j.apenergy.2018.03.067>
- [48] Katherine W Song and Eric Paulos. 2021. Unmaking: Enabling and Celebrating the Creative Material of Failure, Destruction, Decay, and Deformation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3411764.3445529>
- [49] Sangwon Suh and Gjaltp Huppes. 2005. Methods for Life Cycle Inventory of a product. *Journal of Cleaner Production* 13, 7 (June 2005), 687–697. <https://doi.org/10.1016/j.jclepro.2003.04.001>
- [50] Eldy S. Lazaro Vasquez and Katia Vega. 2019. Myco-accessories: sustainable wearables with biodegradable materials. In *Proceedings of the 2019 ACM International Symposium on Wearable Computers (ISWC '19)*. Association for Computing Machinery, New York, NY, USA, 306–311. <https://doi.org/10.1145/3341163.3346938>

- [51] VentureWell. 2023. Measuring Sustainability. https://venturewell.org/tools_for_design/measuring-sustainability/life-cycle-assessment-content/
- [52] Jaylen Wang, Udit Gupta, and Akshitha Sriraman. 2023. Peeling Back the Carbon Curtain: Carbon Optimization Challenges in Cloud Computing. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems (HotCarbon '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3604930.3605718>
- [53] Gregor Wernet, Christian Bauer, Bernhard Steubing, Jürgen Reinhard, Emilia Moreno-Ruiz, and Bo Weidema. 2016. The ecoinvent database version 3 (part I): overview and methodology. *The International Journal of Life Cycle Assessment* 21, 9 (Sept. 2016), 1218–1230. <https://doi.org/10.1007/s11367-016-1087-8>
- [54] 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 (Nov. 2022), 134044. <https://doi.org/10.1016/j.jclepro.2022.134044>
- [55] Zhihan Zhang, Agni K. Biswal, Ankush Nandi, Kali Frost, Jake A. Smith, Bichlien H. Nguyen, Shwetak Patel, Aniruddh Vashisth, and Vikram Iyer. 2023. Recyclable vitrimer-based printed circuit board for circular electronics. <https://doi.org/10.48550/arXiv.2308.12496> arXiv:2308.12496 [physics].

A APPENDIX

A.1 Case Study Summary

We summarized the basic attributes of the three Arduino dev boards and user interaction with our tool for comparing the EIs of these devices in the case study.

	Board	Board Area (mm²)	#Layers	#Components	#ICs from Automated LCI
(1)	Leonardo R3	4180.6368	2	97	20
(2)	MKR FOX 1200	1696.76	4	92	20
(3)	UNO WiFi	3659.124	2	117	26

Table 4. Basic attributes of the three Arduino dev boards.




(a) Arduino Leonardo R3 $\geq$ Arduino MKR FOX 1200(b) Arduino UNO WiFi $\geq$ Arduino MKR FOX 1200(c) Arduino UNO WiFi $\geq$ Arduino Leonardo R3

Fig. 12. Summary of the user interactions with the UI for comparing the three boards. In the subfigures, the red boxes represent unmatched components, and the gray boxes represent matched components. The blue arrow indicates a user-defined rule between components in design A and design B. Here the user is proving $A > B$ for all three cases, so they stop when all components in B are matched and zero or more components in A are unmatched (indicated with the blue box in the last step).

A.2 User Evaluation & Survey

#	Question
1	Have you ever considered the environmental impact when designing a Printed Circuit Board (PCB)?
2	In your opinion, what strategies or approaches would you employ to design a PCB to minimize its carbon footprint?
3	Please indicate your perception of the factors that significantly influence the carbon footprint of a PCB by ranking the following factors from most important to least important: a. Die area of the Integrated Circuit (IC). b. Process technology node of IC. c. Package size of IC (overall size). d. Board size. e. Number of components. f. Number of PCB layers.
4	Would you consider incorporating this tool into your PCB design workflow?
5	On a scale of 1 to 5, where 1 represents the least importance and 5 represents the most importance, how crucial is it for this tool to integrate with existing Electronic Design Automation (EDA) software (like EAGLE or KiCAD)?
6	What additional features or functionalities would you desire in a tool designed for assessing the environmental impact of PCBs?

Table 5. List of the survey questions.