

A Case for Ecological Efficiency in Database Server Lifecycles

Thomas Bodner, Martin Boissier, Tilmann Rabl, Ricardo Salazar-Díaz, Florian Schmeller,
Nils Strassenburg, Ilin Tolovski, Marcel Weisgut, Wang Yue
Hasso Plattner Institute, University of Potsdam
Potsdam, Germany
firstname.lastname@hpi.de

ABSTRACT

Like other software systems, database systems benefit from hardware performance improvements. For the longest time, acquiring new hardware resulted in significant software efficiency gains due to exponential improvements of hardware capabilities. Physical limits in hardware manufacturing have brought former niche designs into standard components, such as multiple cores and specialized circuits. Even with these new designs, hardware improvements are decreasing, while software and applications are still becoming increasingly complex and resource demanding. Given the resource consumption of hardware manufacturing, the ideal lifecycle of hardware naturally has to extend from an efficiency aspect.

In this paper, we try to estimate efficiency of lifecycle duration of database hardware. We calculate the reduction in performance improvements of hardware using publicly available performance numbers, as well as our own benchmarks, and relate them to the specified thermal design power to get the power efficiency. Incorporating estimations on hardware and power production carbon intensity, we challenge current wisdom on hardware replacement frequencies and try to establish new rules of thumb on the ideal hardware lifecycles for database deployments. We present opportunities for future research trends.

1 INTRODUCTION

Processor clock frequency increased steadily over multiple decades. Running an application on a newer CPU with higher frequency improved its single-thread performance without any modifications. In recent years, the CPU clock frequency has stagnated. For continuous performance improvement, vendors have increased the number of CPU cores with two major consequences.

First, the performance of applications does not improve anymore by simply running them on a newer processor. Instead, they must use parallel algorithms to leverage the hardware improvements [28], which database systems have been successfully doing [33].

Second, processors are becoming larger and their production uses more energy. At the same time, their per-watt computation power is still increasing and their operation becomes more energy efficient.

This opens the question, whether server lifecycles need to be adapted to these changes. While industry answers such questions mostly based on economic decisions, another aspect that needs to be addressed is ecological efficiency.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution, provided that you attribute the original work to the authors and CIDR 2025. 15th Annual Conference on Innovative Data Systems Research (CIDR '25), January 19-22, Amsterdam, The Netherlands

In this paper, we analyze the performance and efficiency improvements of 15 years of CPU generations. Using publicly available benchmark data and hardware specifications, we model end-to-end hardware carbon footprints and break-even points for hardware replacement from an ecological viewpoint.

Our analysis confirms previous findings that hardware improvements are slowing down and embodied carbon from manufacturing hardware makes up an increasing part of the overall footprint of server usage. Our own benchmarks show that this effect is drastically increased for database workloads, to the extent that from a performance and efficiency improvement aspect hardware upgrades are hardly justified. This opens an opportunity for researchers and industry to extend server lifetime by focusing on long-term hardware support, considering hardware durability in software design, and research in usability of aging hardware. We make the case for conducting new research in the field of aging hardware.

We make the following contributions:

- We analyze CPU performance trends for 15 years of CPU generations.
- We calculate end-to-end carbon footprints of database server deployments and their break-even points when replaced with upgraded hardware.
- We discuss implications of our findings and point out relevant, novel research directions.

The rest of this paper is structured as follows. [Section 2](#) gives an overview of CPU performance trends. In [Section 3](#), we present our model for server life-cycle carbon dioxide (and equivalents) emissions and present model calculations for different server generations and workloads. We discuss our findings and research opportunities in [Section 4](#). We propose new metrics for benchmarking efficiency in [Section 5](#). In [Section 6](#), we present related work before concluding in [Section 7](#).

2 HARDWARE EVOLUTION

Breakthroughs in hardware research have continuously improved performance for many decades. With decreasing transistor size, hardware manufacturers can increase the frequency of a processor while the power stays the same. Recently, hardware design has run into thermal design challenges, which decrease energy efficiency and increase circuit degradation at higher frequencies. As a solution, manufacturers have started increasing core counts on increasingly large chips, as well as specializing circuits and turning off transistors adaptively in dark silicon designs [10]. In this section, we analyze the performance of several CPU generations.

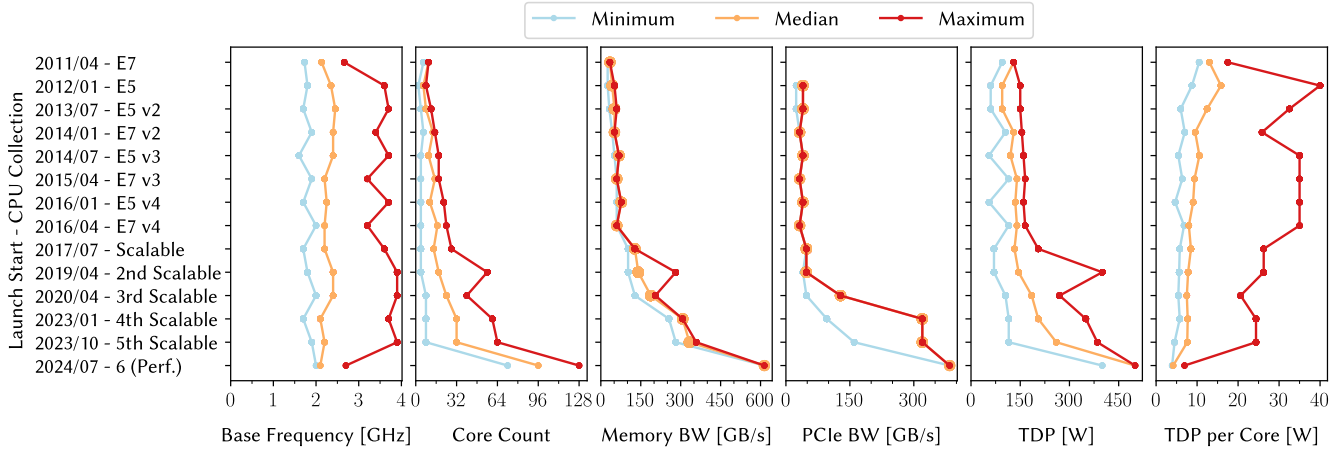


Figure 1: Intel Xeon CPU evolution.

2.1 CPU Trends

We analyze CPU specifications [21] to identify how CPU performance and energy consumption have evolved over the last decade. **Performance Metrics.** We investigate several CPU metrics that impact an application’s performance. The CPU’s frequency determines how many cycles an application has available per second. Increasing the frequency also increases the number of instructions executed per second if the workload is not limited by another resource (e.g., memory). The number of cores multiplies the cycles executed per second. This allows for parallelizing an application so that multiple threads utilize the cores simultaneously, increasing the overall performance. Besides CPU cycles, memory bandwidth impacts performance as it defines how much data an application can read from and write to memory per time unit. The data transfer bandwidth between CPUs and peripheral devices is another aspect influencing database performance. The IT industry is moving towards utilizing specialized hardware [29], which increases the relevance of interconnect technology between a CPU and peripheral devices. PCI Express (PCIe) is a common interconnect for attaching peripheral devices, such as solid-state drives (SSDs) [16], accelerator cards, and smart network interface cards (sNICs). The number of PCIe lanes supported by a CPU influences the number of PCIe devices that can be attached to the CPU. Combined with the supported PCIe generation and its data transfer rate, the number of supported lanes defines the theoretical bandwidth between the CPU and PCIe devices. We include the bandwidth trends in our hardware evolution analysis since memory bandwidth [38] and interconnect bandwidth [35] improvements can alleviate performance bottlenecks in data systems. Advances in the CPU architecture, such as wider vector registers, can lead to single-thread performance improvements, which are orthogonal to frequency increases. The efficiency considerations in this work use CPU benchmark numbers as performance references. We argue that the CPU benchmark numbers reflect architectural CPU advances.

Energy Consumption Metrics. We use the thermal design power (TDP) provided by the hardware vendors as an indicator of the power (in watts) necessary to run a CPU. TDP is defined as the

power needed to maintain the base frequency under a compute-intensive workload [6, 26]. Hardware manufacturers calculate the TDP based on different criteria. Intel calculates the TDP based on the maximum power consumption under a sustained workload necessary to maintain the base frequency [6]. On the other hand, AMD uses the maximum cooling power required to maintain the specification performance under a sustained workload [26]. Since the CPU’s thermal footprint over a sustained period aligns with the CPU’s total power draw, the definitions of both vendors amount to the CPU’s maximal power output. Both definitions do not include the power output during short-running performance bursts exceeding the peak frequency (overclocking).

CPU Models. In this section, we analyze Intel server CPUs, which have a worldwide market share of over 60% from 2012 to 2024 [37]¹. Figure 1 shows performance and efficiency metrics for Intel Xeon CPU families launched since 2011. We exclude *Xeon W Processor*, *Xeon E Processor*, *Xeon D Processor*, and *Xeon E3 CPU collections*. These are suitable for workstations, edge computing, and embedded applications rather than large-scale data centers or enterprise-level infrastructure. We exclude *Xeon Max Series CPUs* (HPC-optimized with high-bandwidth memory). The *Xeon 6* family contains models with either efficiency or performance cores. We exclude CPU models with efficiency cores.

Base Frequency. Figure 1 shows the base frequency, number of cores, memory bandwidth, PCIe bandwidth, TDP, and TDP per core of the selected Intel Xeon CPUs. The minimum and maximum frequencies differ significantly since each CPU collection contains a variety of CPU models, from entry-level CPUs for light workloads to high-performance CPUs for compute-intensive workloads, such as real-time analytics. The median base frequency slightly varies between 2.0 GHz and 2.5 GHz. It does not show an increasing or declining trend over time.

Core Count. With a stagnating median frequency over time, CPU performance can still be improved with increasing cores, which is the trend in the last decade (see Figure 1). A steeper increase in

¹In our extended experiments, we observe that other architectures exhibit very similar trends.




Figure 2: Performance and energy efficiency of TPC CPUs.

cores can be observed starting from 2017, rising from 28 cores in 2017 to 64 cores in 2023. The latest Intel Xeon 6 processors provide up to 128 cores. The core count grew by 16× from 2011 to 2024.

Memory Bandwidth. We investigate the theoretical memory bandwidth of the selected CPUs. The CPU specifications provide the number of supported memory channels and the maximum supported double data rate (DDR) memory speed. We calculate the theoretical memory bandwidth per CPU as:

$$\text{Memory BW} = \# \text{ Memory channels} \times \text{Memory speed} \times 8 \text{ B}$$

where 8 bytes are the data payload per transfer. Figure 1 shows the memory bandwidth increase over the different CPU families. The median memory bandwidth has increased slightly from 2011 to 2016, with values between 30 and 80 GB/s. After 2016, the bandwidth has increased significantly until 2024. The Xeon 6 CPUs (2024) support a median bandwidth of 10.3× the bandwidth of the E7 v4 Family (2016). Compared to 2011, the bandwidth increased by 18×.

PCIe Bandwidth. We investigate the theoretical PCIe bandwidth of the selected CPUs. The CPU specifications provide the number of supported PCIe lanes and the PCIe generation. We derive the transfer rate of the PCIe physical layer from the generation. The selected CPUs support either PCIe 3, 4, or 5, which translates to a transfer rate of 8, 16, and 32 GT/s, respectively. We calculate the theoretical PCIe bandwidth per CPU as:

$$\text{PCIe BW} = \# \text{ Supported PCIe lanes} \times \text{Transfer rate} \times 1 \text{ bit}$$

where 1 bit is the data payload per transfer. We omit the bandwidth for the E7 Family CPUs (2011) due to a lack of PCIe generation information in the specifications. Figure 1 shows that the total theoretical PCIe bandwidth barely changed from 2012 to 2019 with values ranging between 24 and 48 GB/s. The median bandwidth has increased steeply starting in 2020, reaching a maximum of 384 GB/s in 2024. The bandwidth increased by 12× until 2024, compared to the median bandwidth of 40 GB/s in 2012.

Increasing TDP. The TDP stayed constant between 2011 and 2017, except for a lower level median TDP for Xeon E5 (2012) and Xeon E5 v2 (2013) CPUs (see Figure 1). Since 2017, it has significantly increased, correlating with the increase in core count. From 2011 up to 2024, the median TDP grew by 5.3×.

Declining TDP per Core. Figure 1 also shows the CPU’s overall TDP per core. We calculate it as:

$$\text{TDP per core} = \frac{\text{TDP}}{\text{Core Count}}$$

While the CPU’s overall TDP indicates an increasing trend, the TDP per core is declining. There are significant TDP per core differences between concrete CPU models. For example, Xeon E5 CPUs show the most significant difference of more than 30 watts per core.

2.2 Efficiency of CPUs for Database Workloads

After identifying the trends for a large set of Intel Xeon CPUs, we investigate the changes in energy efficiency for selected CPUs used for database workloads.

SPEC CPU Integer Performance Data. *SPEC CPU2006* [17] and *SPEC CPU2017* [5] are CPU performance benchmark suites by the Standard Performance Evaluation Corporation (SPEC). SPEC CPU contains *SPECspeed* and *SPECrate* benchmark suites. *SPECspeed* focuses on single-threaded workloads, using tasks like data compression and text processing for evaluating general-purpose CPU performance. *SPECrate* measures multi-thread performance for multi-core systems, simulating environments like databases and web servers. *SPECspeed* and *SPECrate* exist in two versions: one with integer benchmarks and one with floating-point benchmarks [36]. In this work, we use integer benchmark results² as CPU performance metrics. We use *SPECspeed* results for single-thread performance and *SPECrate* for multi-thread performance. The older SPEC CPU 2006 benchmark suite was replaced by the SPEC CPU 2017 benchmark suite. We combine the 2006 and 2017 integer result datasets for *SPECspeed* and *SPECrate*. The 2006 performance values are higher than the 2017 values, requiring a conversion factor for combining the datasets. We quantify the performance difference of both datasets to ensure comparability. For the same CPU models, number of sockets, and memory configuration, the performance differs by 8.7× for *SPECspeed* and 9.4× for *SPECrate*, on average, with a standard deviation of 0.4 for both. This aligns with factor nine as empirically determined by Rupp.³ We multiply the performance

²SPEC benchmark results: <https://www.spec.org/results.html>

³Rupp’s Microprocessor Trend Data: <https://github.com/karlrupp/microprocessor-trend-data/blob/f2121ab2b83a156b114fe34e62e1edeca98e8e38/newdata.txt>

Table 1: CPU specifications of the evaluation servers.

Processor	Frequency [GHz]	# Cores	Cache Size			
			L1d [KiB]	L1i [KiB]	L2 [MiB]	L3 [MiB]
E7-4880 v2	2.5	15	32	32	0.25	37.5
E7-4850 v4	2.1	16	32	32	0.25	40
8180	2.5	28	32	32	1	38.5
8259CL	2.5	24	32	32	1	35.75
8352Y	2.2	32	48	32	1.25	48
8480CL	2.0	56	48	32	2	105

value of the 2017 datasets by a factor of nine. We refer to a combined set as SPECspeed and SPECrate, respectively. The SPEC datasets contain numbers for systems with one or more CPUs. When referring to either SPECspeed or SPECrate performance, we report the performance per CPU, i.e., the performance value divided by the number of CPUs.

CPU Selection. The Transaction Processing Performance Council (TPC) publishes result data for the TPC benchmarks. We select CPUs of systems used for reported TPC-C and TPC-H V2 benchmark results.⁴ We use TPC-H results measured with a dataset size of 1000 GB. We argue that selecting CPUs based on systems used for TPC database benchmarks results in a set of CPUs employed by the industry for database workloads. We focus on systems with Intel Xeon CPUs, which were used for more than 55% of the reported TPC results. Given the resulting set of CPUs C , we filter the SPECspeed and SPECrate datasets. The datasets only contain entries for CPUs that are present in C . Figure 2 shows the relative change in single-thread and multi-thread performance, as well as energy efficiency (i.e., multi-thread performance per TDP) for the TPC CPUs. For each metric, we use the mean value of the 2006 CPUs as the reference value for calculating the relative change.

Single- and Multi-Thread Performance. Figure 2 shows the change in single-thread (SPECspeed) and multi-thread (SPECrate) performance, as well as the efficiency of the TPC CPUs (SPECrate per TDP) over their launch years. Despite the stable CPU base frequency, single-thread performance has increased over time. We attribute the improvements to advances in the CPU architecture (e.g., wider vector registers). However, the increase in single-thread performance is marginal compared to the multi-thread performance, which provides a better insight into a CPU’s overall performance. The increase in multi-thread performance over time is significant, matching the increasing trend in the number of cores (see Figure 1). **Efficiency.** While the moving average of energy efficiency (i.e., multi-threaded performance per TDP) increases over time, the growth is significantly higher until 2016 and started flattening since then.

2.3 CPU Improvements for Data Management

The SPECrate benchmark suite is highly relevant for integer performance of server systems. However, it does not reflect typical operations found in database contexts. We evaluate different Intel CPUs with launch dates ranging from 2014 (Ivy Bridge) to 2023 (Sapphire Rapids), which have been either available to us in our lab or

can be rented as bare metal instances on Amazon EC2 (m5n.metal with Intel Xeon 8259CL). Table 1 shows the CPU specifications of the evaluation servers. We evaluate two tasks representative of data management: (i) the analytical TPC-H benchmark and (ii) parallel sorting. We run TPC-H with a scale factor of 10 and 25 query streams on Hyrise [8] to evaluate analytical database workloads. Hyrise is a relational research database system that keeps data memory-resident. For sorting, we use the parallel implementation of `std::sort` in `libstdc++ 3.4.32` with GCC 13. Sorting is a typical task that is hard to fully parallelize and frequently used in most data processing systems. We generate a vector of four billion random integer values (`uint32_t`, 16 GB), and measure the time to sort the entire vector. We compare both results with the SPECrate 2017 scores (see Section 2.2). We assume that compiler-driven optimization incorporates advances in CPU architecture. For TPC-H and sorting, we use all logical threads of a single CPU.

The performance results are shown in Figure 3a. Over the span of nine years, the multi-threading performance of the SPECrate scores improves by a factor of 7.26 \times . For complex database workloads, such as the analytical TPC-H benchmark, the difference is lower with a factor of 4.44 \times . While analytical workloads with concurrent query streams scale to large core counts, performance is often not CPU-bound but limited by the main memory bandwidth. For the single process sorting 32 GB in parallel, the factor is 3.03 \times , being significantly lower than the improvement of the SPECrate score. Improvements in sorting performance reflect recent increases in memory bandwidth.

To assess efficiency improvements, we also measure the power consumption using `perf` and the *Running Average Power Limit* (RAPL) interface. Previous research showed that RAPL-based measurements are sufficiently accurate for recent Intel CPU generations [23]. The efficiency results are shown in Figure 3b. We make two observations. First, for all three benchmarks, efficiency improves over time. However, for SPECrate and sorting, efficiency improves slower than performance. While the SPECrate scores improved by 7.26 \times , efficiency only improved by 2.7 \times . An exception is the most recent Sapphire Rapids CPU (Intel Xeon 8480CL), whose efficiency is 11.13 \times better than the Ivy Bridge CPU. We used the same concurrent query load for all systems with core counts ranging from 15 to 56⁵. For the Sapphire Rapids, this load does not fully utilize all logical threads allowing the CPU to reduce its power consumption. Second, while CPUs usually get faster with every generation, efficiency improvements are less steady. For the parallel sorting experiment, the 2017 Skylake CPU (8180) is less efficient than the more than three years older Ivy Bridge CPU. The following Cascade Lake and Ice Lake CPU generations improve significantly over the Skylake CPU before the Sapphire Rapids CPU shows a slight degradation in terms of efficiency again.

3 HARDWARE LIFECYCLE EFFICIENCY

An organization may consider upgrading its database server hardware, looking for an improvement in performance and energy efficiency. Usually, the latest CPU generations outperform their predecessors and are more energy efficient. The new generation can

⁴TPC benchmark results: https://www.tpc.org/tpch/results/tpch_results5.asp?version=3

⁵Running 50 concurrent query streams on the Sapphire Rapids server reduces the TPC-H runs per kJoule to 0.43.

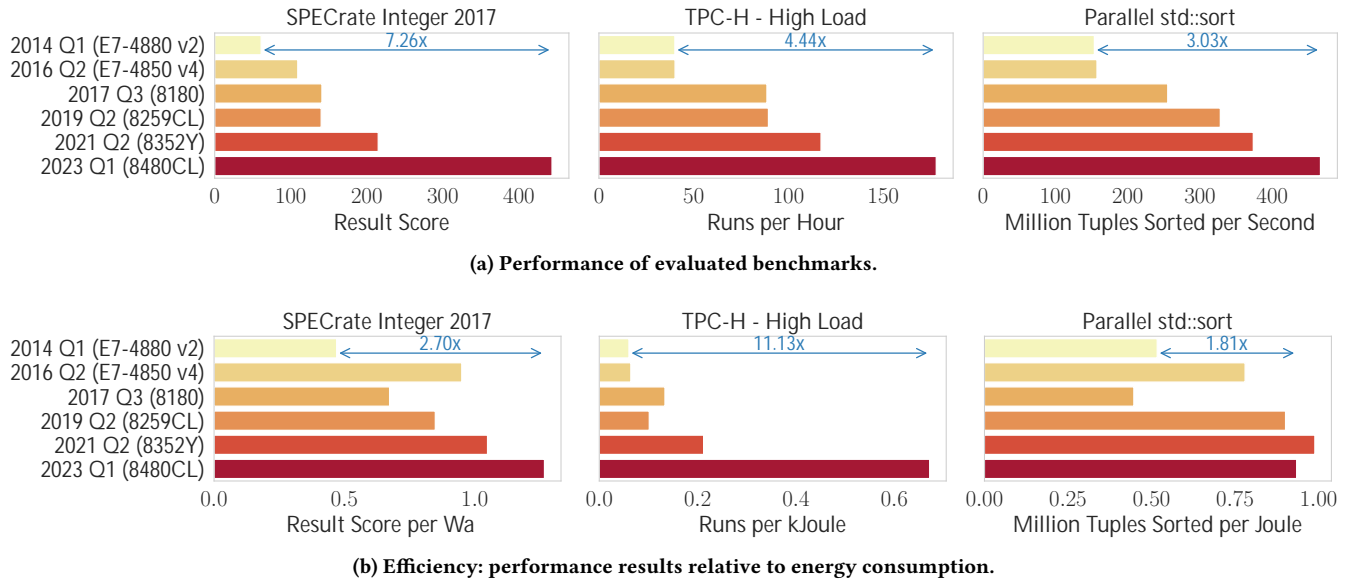


Figure 3: Performance and efficiency results for different generations of Intel Xeon server CPUs and three different benchmarks. The blue arrow denotes the improvement factor comparing the oldest and newest CPU generations.




Figure 4: Lifecycle of a database server. The dimmed stages are not considered in our carbon footprint model.

handle more work per time unit and requires less energy per work unit compared to the previous generation. However, the carbon footprint from manufacturing the new hardware increases as the hardware capability increases [15]. This increase is primarily due to carbon emissions from mining raw materials, manufacturing, assembly, packaging, and transportation. Facility infrastructure construction and chip manufacturing generate capital expenditure (capex)-related carbon emissions, while hardware use and energy consumption generate operational expenditure (opex)-related carbon emissions [15]. Recent research has shown that capex-related emissions are 23x higher than opex-related emissions for data center operators and cloud providers [15]. Thus, evaluating the environmental efficiency of upgrading a server’s hardware should ideally include a "cradle-to-grave" analysis of the hardware, including its capex and opex-related carbon emissions.

Typically, decision-makers use cost-benefit, lifecycle, and break-even analyses to assess a project’s economic viability. In this work, we focus on the environmental break-even analysis and provide a framework to determine the ecological efficiency of replacing a database server. Intuitively, the project will reach the break-even point when the carbon output from the operational use equals

that from hardware manufacturing. We model the amortization of hardware manufacturing by operating hardware over time.

3.1 Manufacturing and Operational Footprint

The lifecycle of a database server includes mining raw materials, transportation, manufacturing, operation, and recycling, as depicted in Figure 4. Each stage contributes to the server’s overall carbon footprint. In the procurement phase, the footprint is mostly impacted by the mining of materials and manufacturing processes, while in the operation phase, energy consumption has the most significant impact. As a result, our analysis primarily focuses on these two stages without transportation and cooling.

A server’s carbon footprint (S_{CF}) consists of emissions from the manufacturing of its various hardware components (E_{CF}) and carbon emissions from its operation (O_{CF}) until the end of its life (EOL):

$$S_{CF} = E_{CF} + \sum_{t=0}^{EOL} O_{CF,t}$$

The ecological break-even point (t_{be}) of replacing a server occurs when the carbon footprint of operating the current server (c) over

time equals the embodied carbon of the new server (n) and its projected operational carbon footprint:

$$\int_{t=0}^{\infty} O_{CF_t}^c = E_{CF}^n + \int_{t=0}^{\infty} O_{CF_t}^n$$

The carbon footprint of manufacturing hardware includes procuring raw materials, gas emissions from the manufacturing process, and transporting hardware to the operational site.

Embodied carbon data from manufacturing is typically not publicly available or incomplete. We model the embodied emissions from manufacturing (E_{CF}) and from operating the server over time (O_{CF}) following the framework proposed by Gupta et al. [14]. We consider the carbon footprint (CF) of manufacturing the server's main components: CPU (ECF_{CPU}), DRAM (ECF_{DRAM}), and SSD (ECF_{SSD}) to model the embodied carbon from manufacturing (E_{CF}):

$$E_{CF} = ECF_{CPU} + ECF_{DRAM} + ECF_{SSD}$$

In this model, the embodied carbon footprint from manufacturing CPUs, DRAM, and SSDs is calculated as follows:

$$ECF_{CPU} = \frac{(CI_{fab} \times EPA + GPA + MPA) \times Area}{Yield}$$

$$ECF_{DRAM} = E_{DRAM} \times Capacity$$

$$ECF_{SSD} = E_{SSD} \times Capacity$$

We model the operational carbon footprint as the sum of the carbon emissions from operating the server's main components, CPUs, DRAM, and SSDs:

$$O_{CF} = OCF_{CPU} + OCF_{DRAM} + OCF_{SSD}$$

We model the yearly operational carbon emissions from the different server components as follows:

$$OCF_{CPU} = kWh_{max/y} \times NPC \times GCI$$

$$OCF_{DRAM} = kWh_{/y} \times Capacity \text{ Factor} \times GCI$$

$$OCF_{SSD} = kWh_{/y} \times GCI$$

We estimate the yearly power draw ($kWh_{max/y}$) for the CPU using the CPU's thermal design power (TDP), which is the power consumption under maximum theoretical load [6]. We consider the grid carbon intensity (GCI) to convert power draw to carbon emissions.

We model the normalized power consumption (NPC) as a linear function of utilization, drawing inspiration from the observations reported by Barroso and Hölze regarding the relationship between power, energy efficiency, and utilization [3]:

$$NPC = P_0 + P_{slope} \times Utilization$$

In line with previous research on energy-efficient servers, we define *utilization* as a performance measure – for instance, the ratio of queries per second to maximum queries per second [3, 13, 30]. Based on the findings of Barroso and Hölze, who monitored thousands of Google servers over six months, we assume that servers are rarely idle or functioning at their maximum theoretical capacity and typically operate between 10% and 50% of their maximum utilization levels [3].

For DRAM, we align the power value with values reported by Lee et al. [27]. As for the SSD, we consider the manufacturer's power specifications.

Table 2: Model parameters

Parameter	Description	Selected Value	Range	Ref.
MPA	Procure material	0.5	0.5 kg CO ₂ /cm ²	[14]
EPA	Fabrication energy	2.15	0.8-3.5 kWh/cm ²	[14]
CI _{fab}	Fab CO ₂ intensity	0.365	0.03-0.7 kg CO ₂ /kWh	[14]
GPA	GHG from fabrication	0.3	0.1-0.5 kg CO ₂ /cm ²	[14]
Yield	Fab yield	0.875	0-1	[14]
E _{DRAM}	DRAM emb. CO ₂	0.3	0-0.6 kg CO ₂ /GB	[14]
E _{SSD}	SSD emb. CO ₂	0.015	0-0.03 kg CO ₂ /GB	[14]
E _{HDD}	HDD emb. CO ₂	0.06	0-0.12 kg CO ₂ /GB	[14]
W _{DRAM}	DRAM power	0.1	0.1 Watts/GB	[27]
W _{SSD}	SSD power	3	2-3 Watts	[7]
P ₀	Min. power (% of TDP)	50	0-100	[3]
P _{slope}	Power slope	0.5	0-1	[3]

Table 2 outlines the parameters used in the model, including the chosen values and a reference for further explanation of these selections.

Limitations and Assumptions. We make the following simplifying assumptions to make the estimation of CO₂ footprints tractable.

- (1) When estimating the CO₂ for a server's production, we only consider the most relevant components like CPU, DRAM, and SSD.
- (2) When estimating the CO₂ footprint for operating a server, we do not consider the energy consumption and the resulting emitted CO₂ for cooling because these numbers are highly reliant on the overall data center setup.
- (3) We simplify the assumption that for a fixed workload, a newer server with an x % better performance consumes x % less energy, positively affecting its power draw and, thus, its CO₂ footprint during operation.
- (4) For the scenarios discussed in the following sections, we adopt the simplifying assumption that a *utilization rate* of x % reflects a server's average *utilization* over time—for example, over a year. This means that, on average, the server operated at x % of its total peak performance throughout the year.

In the following, we explore scenarios involving replacing hardware, considering different energy mixes and server utilization, and analyzing various systems and workloads.

3.2 Scenarios and Workloads

In the following, we give examples of environmental break-even analysis for two scenarios. We answer the question of when upgrading the current hardware setup to a new one minimizes the environmental impact of manufacturing and operating the server. We use the model presented above to estimate the CO₂ emissions from both the current and new systems, in addition to the manufacturing footprint of the new system. To validate our model estimations, we refer to the HPE Power Advisor [18] for comparison and assessment. Additionally, we provide an overview that compares our findings with previous related work focusing on non-database workloads, highlighting the main differences between our model results and earlier studies.

Scenario 1: SPECrate performance in Germany. In the first scenario, we compare two possible setups based on their SPECrate performance. As our current system, we assume a server with an Intel Xeon Platinum 8352Y CPU from Q2 2021, 8× 64 GB of DDR4 DRAM, and two 1.6 TB NVMe SSDs. We analyze after what amount of time it is reasonable – from a CO₂ footprint perspective

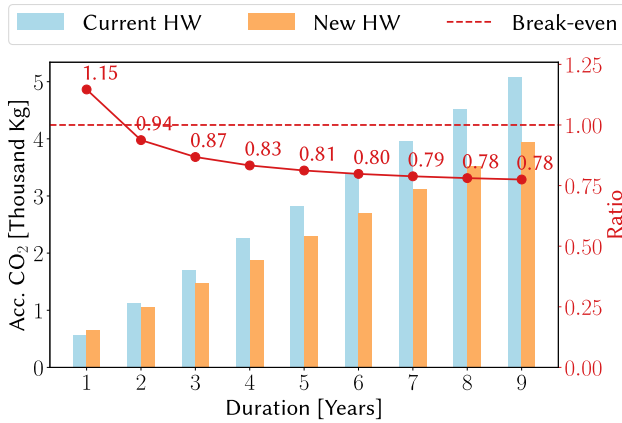


Figure 5: Projected CO₂ accumulated emissions of current and new hardware based on SPECrate numbers for 30% utilization and a German energy mix.




Figure 6: Projected CO₂ accumulated emissions of current and new hardware for a sorting workload, 60% utilization and a Swedish energy mix.

– to replace the current system with a two years newer system having an Intel Xeon Platinum 8480C CPU from Q1 2023 and the same DRAM and SSD capacity. For our comparison, we assume to (1) run both setups in Germany with a GCI of 344g of CO₂ per kWh [9], (2) a workload that leads to a system utilization of 30%, and (3) the SPECrate numbers as a representative indicator of how well the CPUs perform on our given task. We use the model presented in Section 3.1 to estimate the CO₂ emissions emitted for manufacturing and running both systems. Figure 5 shows that replacing the current hardware with the next generation reduces the accumulated CO₂ after roughly 1.7 years. Thus, if we expect to run the new system for less than 1.7 years an upgrade is not advisable. The reason for the short replacement cycle is that, given our assumptions, the more recent hardware is much more efficient such that the CO₂ savings during operation rapidly compensate for the additional CO₂ footprint for its production.

Scenario 2: Sorting workload in Sweden. In a second scenario, we assume the current and new systems to be nine years apart. Both are equipped with 8× 64GB of DDR4 DRAM and two 1.6TB

NVMe SSDs. The CPU of the current system is an Intel Xeon E7-4880 v2 (released in 2014), and the CPU of the new system is an Intel Xeon Platinum 8480CL (released in 2023). In contrast to the scenario above, we assume the workload to be sorting, which is more representative of database workloads than SPECrate performance. We rely on our measurements of the parallel `std::sort` workload from Section 2.3 to estimate the performance difference between the current and the new system. Additionally, we assume the servers to be 60% utilized and to be placed in Sweden, leading to a GCI of 25g of CO₂/kWh.

Figure 6 shows that the new system’s manufacturing CO₂ footprint is very large in comparison to the saved CO₂ from operational improvement due to performance increase. Due to the low carbon intensity from Sweden (GCI of 25g of CO₂/kWh) and the limited performance improvements from the new hardware, replacing the system only reduces the accumulated CO₂ emissions after a duration of approximately 17 years.

Validating our Model. To validate our findings and the estimation of the operational carbon footprint, we compare the accumulated CO₂ emissions as reported by our model to those reported by HPE’s Power Advisor [18].⁶ Since HPE’s Power Advisor only supports a limited set of hardware, we keep the DRAM and SSD configurations from Scenario 1 and 2 but select two subsequent generations of servers with AMD CPUs that are used to report TPC-H benchmark numbers. The current system is an HPE ProLiant DL325 Gen10 server from 2019 with an AMD 7502P 2.5 GHz 32 Core CPU and we consider to replace it with the three years newer successor: an HPE ProLiant DL325 Gen11 server with an AMD 9334 2.7 GHz 32 Core CPU. For our comparison, we chose the CPU’s SPECrate numbers as a performance indicator, a 30% utilization, and Sweden as the location. Comparing Figure 7a showing the numbers generated by our model and Figure 7b showing the numbers from the HPE Power Advisor, we see that the absolute numbers are similar while the break-even points are five years apart. Several factors account for the differences between the models. The HPE Power Advisor considers the power usage utilization of the data center, networking, and power supply inefficiencies in addition to the components considered in our model. However, both analyses report a break-even point of more than ten years and strongly advise against an immediate hardware upgrade to the latest generation. This result is even more evident when we use the manufacturing carbon emissions for both systems as reported by HPE [19, 20]. Compared to our model, HPE estimates an approximately three times higher CO₂ manufacturing footprint, shifting the break-even point far beyond 17 years. This underlines our general observation that short replacement cycles are not recommended.

Comparison with non-DB workloads. Li et al. assessed the carbon footprint of various deep neural network (DNN) workloads on different high-performance computing (HPC) systems [30]. We compare our results with their findings. Similarly to our proposed scenarios, Li et al. analyze the impact of replacing a current system with more energy-efficient hardware on accumulated carbon emissions over time. They report a break-even point of less than five years for low (20g CO₂/kWh), medium (200g CO₂/kWh), and high

⁶We actively do not consider the CO₂ footprint for cooling from the HPE Power Advisor to be consistent across both models.

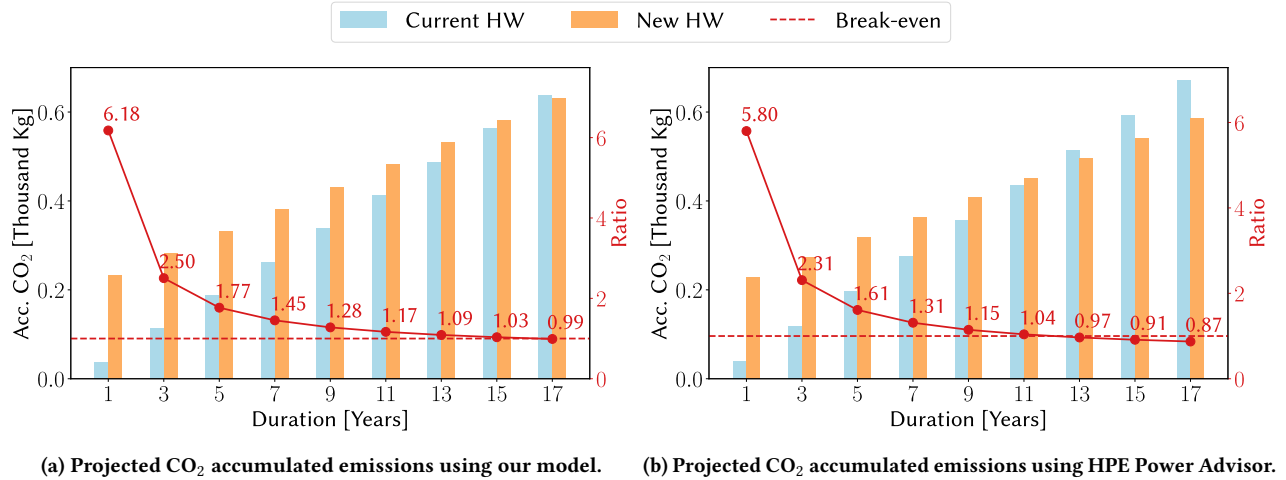


Figure 7: Comparison between the projected CO₂ emissions of our model and the HPE Power Advisor based on SPECrate numbers for 30% utilization and a Swedish energy mix.

(400g CO₂/kWh) carbon intensity scenarios as well as break-even points of less than two years for low (20%), medium (40%), and high (60%) utilization levels at a fixed medium carbon intensity (200g CO₂/kWh). Our results show a significant difference in the break-even point projection for the low carbon intensity scenario (20g CO₂/kWh). We identify four reasons for this difference. First, we consider the embodied carbon of the whole system (CPU, DRAM, HDD/SSD), while they only consider the embodied and operational carbon of the GPUs. Second, they utilize Carbontracker [2], a tool that gathers statistics during system operation to estimate the grams of CO₂ from the operation. Third, the workloads considered in their work are much more compute-heavy, which produces faster carbon savings given the better energy efficiency of the new hardware. Finally, the jumps between the performance improvements of new GPU generations considered in their analysis are higher than those we considered between CPU generations.

4 DISCUSSION

Continuously improving performance of state-of-the-art systems often comes with the price of upgrading system components to newer generations at an unsustainable rate. In this paper, we argue that this is not reasonable for database deployments. As our results show, replacing hardware frequently does not improve performance significantly for typical database workloads, but it leads to significantly higher carbon emissions. We propose that database systems research should actively target hardware-friendly algorithms to support extended hardware lifetimes. We envision future algorithms to incorporate hardware maintenance with the explicit goal of increasing hardware lifetime.

Previous research efforts focus on thermal management and update cycles to extend the lifespan of hardware. Fluctuating temperatures affect the lifespan of a CPU and increase cooling costs. For every 10°C temperature increase, the chip’s failure rate doubles [25]. Current solutions propose software-based thermal management techniques to extend the lifespan of chips [34, 40]. They slow the

major heat-generating processes in CPUs to keep temperatures below a certain threshold. NVM devices, such as SSDs, suffer from low endurance problems because the number of erase operations for memory cells is limited. The state-of-the-art solutions reduce SSD write traffic by proposing a comprehensive write buffer cache mechanism [12, 22]. Additionally, recent efforts regarding the lifespan of DRAM are focusing on investigating the aging behavior of DRAM cells, which leads to performance degradation [4]. We believe that the database community is positioned well to contribute to this line of research.

The lifetime of hardware, however, is limited, especially with high-intensity computing, so replacing it is inevitable. Hardware vendors recommend a product lifetime of four years [19], while cloud vendors target a six-year replacement rate in their fleets [1]. In Figures 5 and 6, we show that the break-even point of the systems’ carbon footprint can take up much longer. We observe a significant difference in the break-even points when utilizing energy with higher carbon intensity. Utilizing a carbon-intensive energy mix makes upgrading hardware to a new generation economical and ecological within less than one year. Transitioning towards low-carbon energy extends the efficient usage of systems past the recommended lifetime potentially by years.

To facilitate this, hardware vendors need to enable access and ease of maintenance of hardware components affected by high system utilization, such as SSDs and DRAM DIMMs. One notable initiative in this direction is the Right to Repair legislation proposed by the European Commission [11]. The legislation emphasizes the priority of repair over replacement, as well as integrating an ease of reparability standard. In combination, the utilization of low-carbon energy and the reparability of components contribute to more environmentally friendly data processing.

Another dimension we have not included in this paper is the effect of increasing application complexity on carbon emission and the ecological efficiency of hardware upgrades. The rise of complex machine learning models in mainstream software applications

drives hardware spending and deployment and the carbon footprint of the IT sector. Large companies are currently missing their carbon reduction targets mainly due to generative AI.⁷ Database research has the opportunity to set a positive example and reduce carbon emissions rather than increase them.

We see the need for new efficiency metrics to facilitate research in environmentally friendly database research. Current database metrics are typically measured in queries or transactions per second for transactional systems or query execution time for analytical systems. Efficiency metrics used in benchmarks such as those from the Transaction Processing Performance Council also consider cost or power. While cost is calculated based on a three-year deployment – too short, as we have seen in our analysis – power is only measured during the actual run. We see the need for an attribution of the lifecycle carbon footprint to each performance run.

Energy-efficient and green database systems design receives attention from industry and academia. Barroso and Hölzle introduced the notion of *energy-proportional* machines after observing that the energy efficiency—defined as the ratio of performance to power—of thousands of Google servers dropped by more than 50% at average utilization levels (20%-30%) compared to the energy efficiency at peak performance [3]. In an *energy-proportional* server, power consumption is proportional to the delivered performance, meaning that when the server is idle, it does not draw any power and only draws the highest power at peak performance. Designing *energy-proportional* systems is still a significant challenge, primarily due to 1) the nonlinear relationship between current hardware power consumption and performance, 2) the lack of broader dynamic power range modes of hardware, and 3) the poor interaction between software and the power management features offered by hardware [13]. Research on *energy-proportional servers* is essential for reducing the carbon footprint of database systems and extending the operational lifespan of the hardware behind them. For this to succeed, hardware vendors must offer a broader range of power features, e.g., inactive power modes or a wide dynamic power range. At the same time, database system designers must prioritize the tight coupling of software and hardware’s power management features to develop systems that consume power in proportion to their work. More efficient hardware and hardware operation will ultimately also increase the ecological server lifetime, further mandating a focus on research in the area of lifetime extension of database servers.

5 MEASURING EFFICIENCY

Current database benchmark metrics are typically measured in queries or transactions per second for transactional systems or query execution time for analytical systems. With a narrow focus on peak performance, database research fails to incentivize efficiency of hardware and software.

Established efficiency metrics used in benchmarks such as those defined by the TPC consider cost or power. While cost is calculated based on a three-year deployment, power consumption is measured during benchmark execution. Both metrics are relevant but hardly ever used in research. Additionally, current benchmarks fail to

⁷For example, Microsoft reports a 29.1% increase in carbon emissions since 2020 [31] despite their targets to net carbon neutrality in 2030, other companies have similar trends.

represent the ever-increasing software complexity and application bloat. While transactional throughput and query latency improve over time due to hardware and database architecture improvements, these are quickly eaten up by applications adding additional queries and transactions per user. Applications and complete deployments should, therefore, report resources consumed per user or end-to-end workload unit rather than per simple database interaction.

With our model for carbon footprints, we hope to lay a foundation for reporting ecological efficiency of database systems and applications. In our current model, we focus on static workloads. This model can also be used to measure the increasing efficiency or inefficiency of database applications by normalizing the carbon footprint of different deployments or application versions per user.

6 RELATED WORK

Increased carbon emissions are widely discussed and impact many decisions made by commercial companies [39]. Gupta et al. [14] developed a framework to measure the carbon footprint of hardware. Their framework involves three cases: the trade-off between general-purpose and specialized hardware, minimizing hardware resources, and recycling hardware. The MIT Materials Systems Laboratory designs PAIA [32] to provide an efficient and cost-effective estimate of the carbon footprint of products. This tool is used by many commercial companies, e.g., Intel, Dell, Cisco, and IBM. Li et al. [30] analyze the carbon footprint of high-performance computing (HPC) systems. They consider the carbon footprint of the hardware components and the geographic carbon intensity to characterize the carbon footprint of a system across its lifecycle. The authors conclude that hardware upgrades often come with performance improvements at the price of increased embodied carbon depending on the carbon intensity of the energy source.

Nambiar and Poess make an early analysis of the performance improvements of DBMS in comparison to Moore’s law [33]. Their results from 2010 show a close resemblance of TPC-C improvements to the general hardware improvements. Koomey et al. [24] provide an overview of trends in performance, costs, and energy use for servers. The performance trends for servers track Moore’s law well.

Our research extends previous work to database workloads and shows that current hardware improvements do not translate to database performance.

7 CONCLUSION

In this paper, we analyze the efficiency of database hardware upgrades and their impact on lifecycle carbon emissions. We show that current hardware improvements do not fully translate to database performance and servers require very long lifetimes to break even with previous generations in regions with low carbon intensity power generation. With this work, we hope to show the necessity for research on environmentally conscious research in database systems including extending hardware lifetime and support for older hardware.

ACKNOWLEDGMENTS

This work was partially funded by the German Research Foundation (ref. 414984028), the European Union’s Horizon 2020 research and innovation programme (ref. 957407), and by SAP.

REFERENCES

- [1] Amazon Web Services. 2024. AWS Cloud - Amazon Sustainability. Online. Retrieved 2024-11-15 from <https://sustainability.aboutamazon.com/products-services/aws-cloud>
- [2] Lasse F. Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. 2020. Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models. *CoRR* abs/2007.03051 (2020). arXiv:2007.03051 <https://arxiv.org/abs/2007.03051>
- [3] Luiz André Barroso and Urs Hölzle. 2007. The Case for Energy-Proportional Computing. *Computer* 40, 12 (2007), 33–37. <https://doi.org/10.1109/MC.2007.443>
- [4] Md Kawser Bepary et al. 2022. DRAM retention behavior with accelerated aging in commercial chips. *Applied Sciences* 12, 9 (2022), 4332.
- [5] James Bucek et al. 2018. SPEC CPU2017: Next-Generation Compute Benchmark. In *ACM/SPEC ICPE*. ACM, 41–42.
- [6] Intel Corporation. 2023. Thermal Design Power (TDP) in Intel Processors. Online. Retrieved 2024-11-15 from <https://www.intel.com/content/www/us/en/support/articles/000055611/processors.html>
- [7] Solid State Storage Technology Corporation. 2024. SSDs vs. HDDs: The Green Power Consumption Advantage with CVB SATA SSD. Online. Retrieved 2024-11-15 from <https://www.ssstc.com/knowledge-detail/ssd-vs-hdd-power-efficiency>
- [8] Markus Dreseler et al. 2019. Hyrise Re-engineered: An Extensible Database System for Research in Relational In-Memory Data Management. In *Extending Database Technology (EDBT)*. OpenProceedings.org, 313–324.
- [9] Electricity Maps. 2024. Live 24/7 CO2 emissions of electricity consumption. Online. Retrieved 2024-11-15 from <https://app.electricitymaps.com>
- [10] Hadi Esmailzadeh et al. 2011. Dark Silicon and the End of Multicore Scaling. In *International Symposium on Computer Architecture (ISCA)*. 365–376.
- [11] European Commission. 2024. Right to repair: Making repair easier and more appealing to consumers. Online. Retrieved 2024-11-15 from <https://www.europarl.europa.eu/news/en/press-room/20240419IPR20590/right-to-repair-making-repair-easier-and-more-appealing-to-consumers>
- [12] Ziqi Fan and Dongchul Park. 2019. Extending SSD Lifespan with Comprehensive Non-Volatile Memory-Based Write Buffers. *Journal of Computer Science and Technology* 34, 1 (2019), 113–132.
- [13] Binglei Guo et al. 2022. Energy-Efficient Database Systems: A Systematic Survey. *Computing Surveys* 55, 6, Article 111 (Dec. 2022), 53 pages. <https://doi.org/10.1145/3538225>
- [14] Udit Gupta et al. 2022. ACT: designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the Annual International Symposium on Computer Architecture (ISCA)*. 784–799.
- [15] Udit Gupta et al. 2022. Chasing Carbon: The Elusive Environmental Footprint of Computing. *IEEE Micro* 42, 4 (July 2022), 37–47.
- [16] Gabriel Haas and Viktor Leis. 2023. What Modern NVMe Storage Can Do, And How To Exploit It: High-Performance I/O for High-Performance Storage Engines. *Proc. VLDB Endow.* 16, 9 (2023), 2090–2102. <https://doi.org/10.14778/3598581.3598584>
- [17] John L. Henning. 2006. SPEC CPU2006 benchmark descriptions. *SIGARCH Computer Architecture News* 34, 4 (Sep 2006), 1–17. <https://doi.org/10.1145/1186736.1186737>
- [18] Hewlett Packard Enterprise. 2024. HPE Power Advisor. Online. Retrieved 2024-11-15 from <https://poweradvisorext.it.hpe.com/?Page=Index>
- [19] Hewlett Packard Enterprise. 2024. HPE Product Carbon Footprint – HPE ProLiant DL325 Gen10 Server. Online. Retrieved 2024-11-15 from <https://www.hpe.com/psnow/doc/a50004544enw>
- [20] Hewlett Packard Enterprise. 2024. HPE Product Carbon Footprint – HPE ProLiant DL325 Gen11 Server. Online. Retrieved 2024-11-15 from <https://www.hpe.com/psnow/doc/a00140770enw>
- [21] Intel Corporation. 2024. Intel Product Specifications. Online. Retrieved 2024-11-15 from <https://ark.intel.com/content/www/us/en/ark.html>
- [22] Saeed Kargar and Faisal Nawab. 2023. Challenges and future directions for energy, latency, and lifetime improvements in NVMs. *Distributed and Parallel Databases* 41, 3 (2023), 163–189.
- [23] Kashif Nizam Khan et al. 2018. RAPL in Action: Experiences in Using RAPL for Power Measurements. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* 3, 2 (2018), 9:1–9:26.
- [24] Jonathan G Koomey et al. 2009. Assessing trends over time in performance, costs, and energy use for servers. *Lawrence Berkeley National Laboratory, Stanford University, Microsoft Corporation, and Intel Corporation, Tech. Rep* (2009).
- [25] Clemens JM Lasance. 2003. Thermally driven reliability issues in microelectronic systems: status-quo and challenges. *Microelectronics Reliability* 43, 12 (2003), 1969–1974.
- [26] Patrick Lathan. 2019. AMD Ryzen TDP Explained: Deep-Dive on TDP Definitions & What Cooler Manufacturers Think. Online. Retrieved 2024-11-15 from <https://gamersnexus.net/guides/3525-amd-ryzen-tdp-explained-deep-dive-cooler-manufacturer-opinions>
- [27] Seunghak Lee et al. 2021. GreenDIMM: OS-assisted DRAM Power Management for DRAM with a Sub-array Granularity Power-Down State. In *International Symposium on Microarchitecture (MICRO)*. ACM, 131–142. <https://doi.org/10.1145/3466752.3480089>
- [28] Charles E Leiserson et al. 2020. There’s plenty of room at the Top: What will drive computer performance after Moore’s law? *Science* 368, 6495 (2020), eaam9744.
- [29] Alberto Lerner and Gustavo Alonso. 2024. Data Flow Architectures for Data Processing on Modern Hardware. In *International Conference on Data Engineering (ICDE)*. IEEE, 5511–5522. <https://doi.org/10.1109/ICDE60146.2024.00439>
- [30] Baolin Li et al. 2023. Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. ACM, 19:1–19:15.
- [31] Microsoft. 2024. Our 2024 Environmental Sustainability Report. Online. Retrieved 2024-11-15 from <https://blogs.microsoft.com/on-the-issues/2024/05/15/microsoft-environmental-sustainability-report-2024>
- [32] MIT Materials Systems Laboratory. 2024. Product Attribute to Impact Algorithm - PAIA. Online. Retrieved 2024-11-15 from <https://msl.mit.edu/projects/paia/main.html>
- [33] Raghunath Othayoth Nambiar and Meikel Poess. 2010. Transaction Performance vs. Moore’s Law: A Trend Analysis. In *TPC Technology Conference (TPCTC) (Lecture Notes in Computer Science, Vol. 6417)*. Springer, 110–120.
- [34] Shehenaz Shaik and Sanjeev Baskiyar. 2017. Proactive thermal aware scheduling. In *International Green and Sustainable Computing Conference (IGSC)*. IEEE, 1–6.
- [35] Anil Shanbhag, Samuel Madden, and Xiangyao Yu. 2020. A Study of the Fundamental Performance Characteristics of GPUs and CPUs for Database Analytics. In *Proceedings of the International Conference on Management of Data (SIGMOD)*. ACM, 1617–1632. <https://doi.org/10.1145/3318464.3380595>
- [36] Standard Performance Evaluation Corporation. 2024. SPEC Benchmarks and Tools. Online. Retrieved 2024-11-15 from <https://www.spec.org/benchmarks.html>
- [37] Statista. 2024. Intel/AMD x86 computer CPU market share 2024. Online. Retrieved 2024-11-15 from <https://www.statista.com/statistics/735904/worldwide-x86-intel-amd-market-share/>
- [38] Elias Stehle and Hans-Arno Jacobsen. 2017. A Memory Bandwidth-Efficient Hybrid Radix Sort on GPUs. In *Proceedings of the International Conference on Management of Data (SIGMOD)*. ACM, 417–432. <https://doi.org/10.1145/3035918.3064043>
- [39] Rida Waheed, Sahar Sarwar, and Chen Wei. 2019. The survey of economic growth, energy consumption and carbon emission. *Energy Reports* 5 (2019), 1103–1115.
- [40] Shikang Xu, Israel Koren, and C Mani Krishna. 2019. Thermal aware task scheduling for enhanced cyber-physical systems sustainability. *IEEE Transactions on Sustainable Computing* 5, 4 (2019), 581–593.