16 Years of SPEC Power:
An Analysis of x86 Energy Efficiency Trends

Hannes Tröpgen, Robert Schöne, Thomas Ilsche, Daniel Hackenberg

ZIH, CIDS, TU Dresden, 01062 Dresden, Germany

{hannes.troepgen, robert.schoene, thomas.ilsche, daniel.hackenberg}@tu-dresden.de

Abstract—The SPEC Power benchmark offers valuable insights into the energy efficiency of server systems, allowing comparisons across various hardware and software configurations. Benchmark results are publicly available for hundreds of systems from different vendors, published since 2007. We leverage this data to perform an analysis of trends in x86 server systems, focusing on power consumption, energy efficiency, energy proportionality and idle power consumption. Through this analysis, we aim to provide a clearer understanding of how server energy efficiency has evolved and the factors influencing these changes.

Index Terms—Computer architecture, Performance analysis, High performance computing, Processor energy efficiency

I. INTRODUCTION

SPECpower_ssj 2008[1] is the most prominent server energy efficiency benchmark. Its rigorous methodology and healthy benchmark submission review process have led to 16 years of continuous benchmark submissions and corresponding published data. These results allow hardware vendors to rank and promote their systems with respect to energy efficiency, measured in ssj_ops/W. This metric gives customers an idea of how much computing they get for each invested Joule of energy, where a lower power consumption can increase it as well as a higher processing performance. Figure 1 illustrates some strengths of the benchmark: Due to the simplicity and scalability of the benchmark, server systems with multiple sockets and/or nodes can be measured. The workload can also be executed on different operating systems (OS) and different hardware, even though non-x86 processors are rare, and up to 2017, more than 97% of results use Windows.

Based on the benchmark results available via the SPEC Power website, we track the performance and power efficiency of x86 processors over the previous 16 years. We analyze data for different load levels to evaluate energy proportionality, as well as active idle power consumption trends.

II. BACKGROUND AND RELATED WORK

SPECpower_ssj 2008 [1], [2] is designed to measure "the performance and power consumption of servers". It consists of an integer-heavy transactional Java-based client/server workload with six differently weighted transaction types. A calibration phase is used to determine the maximum throughput of the system under test (SUT), which runs the server side. Partial

---

[1]*SPECpower_ssj 2008* is the first and so far only release of the *SPEC Power benchmark suite* released by the *Standard Performance Evaluation Corporation* (SPEC).

loads of $10\%, 20\%, \ldots, 90\%$ are created by scaling down the number of transactions proportionally. This allows the SUT to apply power-saving mechanisms [3, Section Power-Saving Techniques] and can be used to analyze energy proportionality [4], [5]. The test regime also includes a 0% load point, which greatly helps to record, track, and optimize *active idle* power consumption.

Based on the SPEC Power methodology, vendors and performance engineers generated hundreds of reports with 1017 being published on the SPEC website[2] at the time of writing. This long history of vendor-submitted data distinguishes it from other energy efficiency benchmarks, e.g., Green500 [6] or SPEC OMP 2012 [7]. The SPECpower committee[3] is not only responsible for the SPECpower_ssj 2008 benchmark, but also for the definitions and tool infrastructures for power measurements [2], e.g., the ptdaemon interface, SERT suite, and the Chauffeur

---

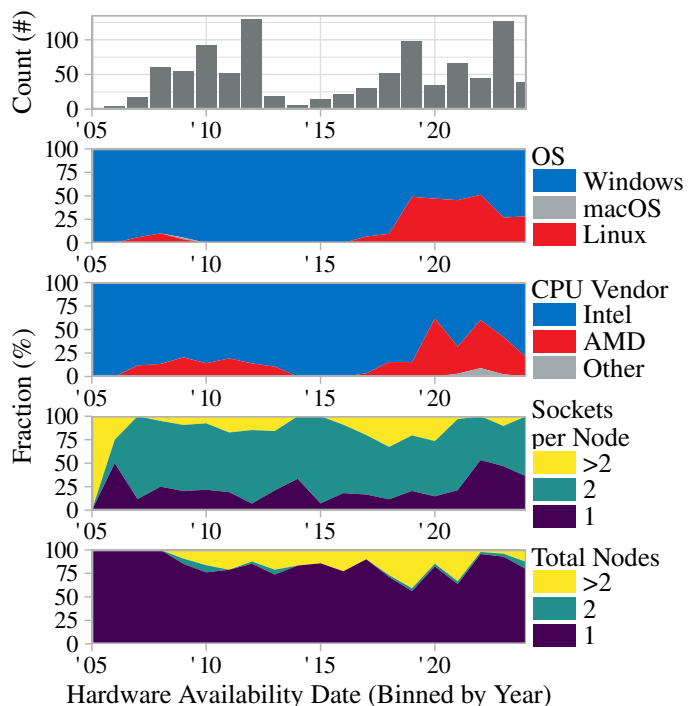[2]https://www.spec.org/power_ssj2008/results/

[3]https://www.spec.org/power/

Figure 1: Share of features on all 960 successfully parsed (unfiltered) SPECpower_ssj 2008 results (as of June 2024)
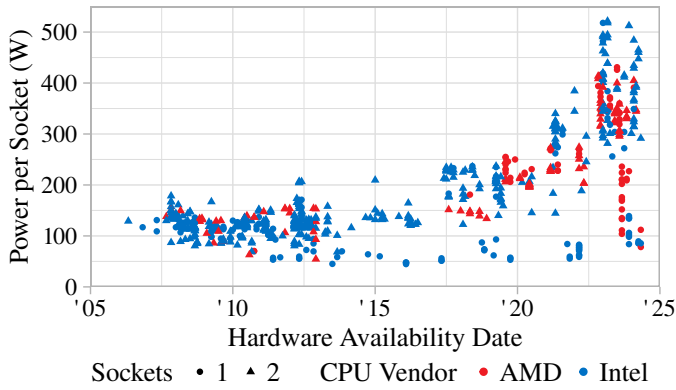
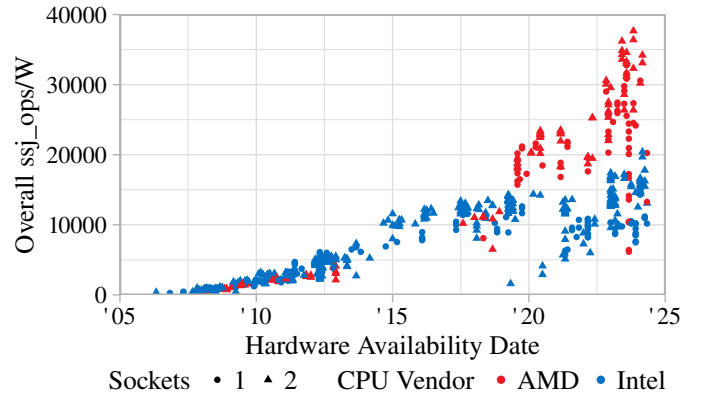Figure 2: Power consumption (per socket) at full load trend

Figure 3: Overall efficiency trend

Worklet Development Kit, which are also used for other benchmarks and certifications [8]. Recently, members of SPEC concluded that the workload of SPECpower_ssj 2008 does not represent the current demands and described the next version of the benchmark, the still unreleased SPECpowerNext [9], [10]. It will use different technologies for interfaces and measurement handling, but it will also use different workloads targeted at accelerators and CPUs.

As dataset for this paper, we download all 1017 `.txt` result files,[4] and extract information for hardware and software stack, as well as performance and power measurement results. We check the consistency, filtering runs that have not been accepted by SPEC (40), runs with ambiguous (3) or implausible (4) dates, ambiguous CPU names (3), or missing node count (1), as well as submissions where reported core/thread counts are inconsistent (5) or implausible (1). This leaves a dataset of 960 successfully parsed runs. Each run has four associated dates: (a) The test date, (b) the submission date, as well as (c) hardware and (d) software availability dates. As we discuss trends in hardware, we use the hardware availability date, which indicates the month at which the system became "generally available" [11]. Hence, even though the earliest results were published in 2007, some runs are associated with earlier dates.

Figure 1 shows general trends over time. For the whole duration from 2005 to 2023, an average of 44.2 runs were submitted per year. Between 2013 and 2017, this drops to 15.2 runs per year. The increased number of submissions from 2018 onward coincides with an increase in submissions using Linux (from 2.2 % before 2018 to 36.3 % after 2018), and an increase in submissions using AMD processors (from 13.0 % to 31.3 %). The latter observation can be explained by AMD's introduction of its EPYC server CPUs in 2017.

To keep the systems within the dataset comparable, we exclude uncommon configurations: Runs with CPUs made by neither Intel nor AMD (9), and all runs not on server or workstation CPUs[5] (6) are filtered. Finally, we remove runs with more than one node or more than two sockets (269). After all filtering, 676 runs remain as the base for all further analysis.

III. PERFORMANCE AND EFFICIENCY TRENDS

Across all runs, the maximum power consumed rapidly increases over the years, as Figure 2 depicts. This trend is consistent across both AMD and Intel, although the spread increased in recent years. The dataset shows a broad insight into both low and high Thermal Design Power (TDP) processors. Overall, this upward trend of TDP cannot continue indefinitely, as cooling infrastructure has to be scaled with the power consumed, with air cooling becoming unfeasible at around 400 W TDP [12].

The increase in power consumption per socket is most pronounced at full load (100 %), where the mean since 2022 increased ∼2.5x compared to runs up to 2010 (119.0 W to 303.3 W). However, power consumption at all other load levels increases as well, e.g., by ∼1.8x at 20 % or by ∼2.2x at 70 % load using means across the same time spans.

Dividing the achieved performance (ssj_ops rate) by the mean power consumption results in the energy efficiency in ssj_ops/W. Across all runs and load levels, this energy efficiency improved over the years, as Figure 3 shows.[6] Here, AMD emerges as the driver of the upward energy efficiency trend, in particular from ∼2018 onward. Even though Intel's efficiency is also growing, out of the 100 most efficient runs 98 use AMD processors. Similarly to the overall score, the runs record the achieved ssj_ops and average power at every load level, which enables us to compute the efficiency per load level. We then scale this to the efficiency at full load, yielding the *relative efficiency* per load level. A relative efficiency <1 is less, >1 more efficient than full load; a relative efficiency of 1 at all load levels essentially corresponds to *energy proportionality*. We summarize these relative efficiencies from 60 % to 90 % load, binned by CPU vendor and year in Figure 4.

In the early years, lower load is consistently less efficient compared to full load. Over time, the relative efficiency approaches 1 also for lower load levels. Since 2012 Intel systems have a mean relative efficiency >1 with all load levels ≥70 %, but from 2017 on, we observe a regression

---

[4]https://www.spec.org/power_ssj2008/results/power_ssj2008.html
[5]I.e. CPUs marketed neither as *Xeon*, *Opteron*, nor *EPYC*.

[6]The overall efficiency (overall ssj_ops/W) for SPEC Power is defined as $\sum \text{ssj\_ops} / \sum P$ across all load levels including active idle [11].
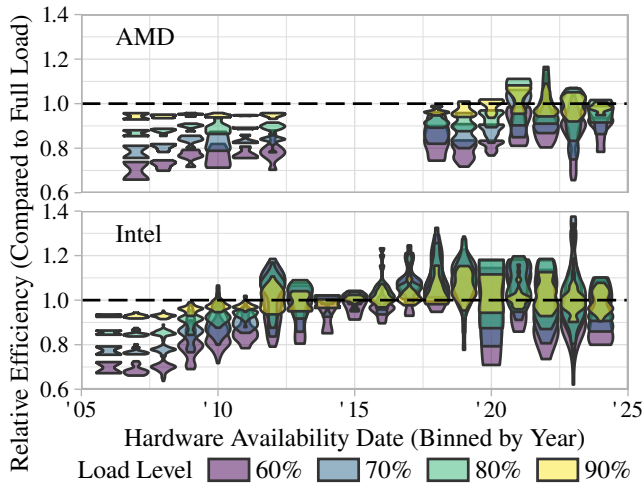
Figure 4: Distribution of relative efficiency at 60% to 90%, binned by year and CPU vendor.

back to ~1. This likely stems from overlapping related effects of increased power reduction at low lower load, but also the use of inefficient turbo states at full load, which were particularly pronounced around 2017. For AMD, the relative efficiency approaches 1 around 2021. Even though there are still visible differences between AMD and Intel results from 2021 onward, both have a large spread: No CPU vendor has a universally better relative efficiency at any given load point – their energy proportionality is diverse. When comparing absolute instead of relative efficiency however, AMD systems still clearly outperforms Intel systems (cf. Figure 3).

IV. IDLE POWER TREND ANALYSIS

In addition to the full and partial load, the power is measured for an active idle interval. During active idle, the SUT is ready to perform work, but no transactions are being processed. Partial load configurations already allow for significant power reduction by leveraging techniques such as Dynamic Voltage and Frequency Scaling (DVFS) and core C-states [13]. In active idle, the power can be reduced further by powering down shared components, e.g., implemented by package C-states (see [3, Chapter 2]).
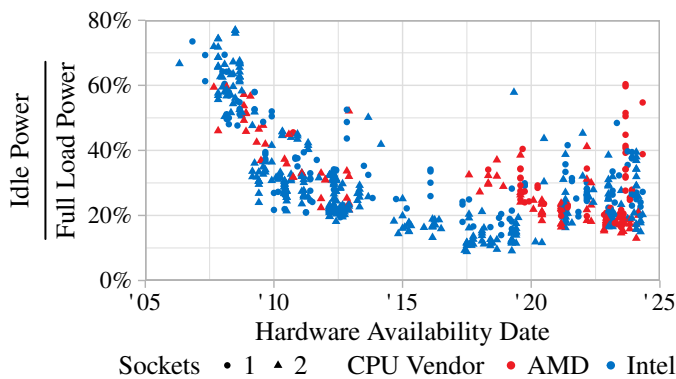
The active idle power consumption is particularly relevant for High-Performance Computing (HPC) systems. HPC systems strive to maximize utilization and may, on average, have a higher utilization than other data center applications. However, if no batch job is executed on an HPC node, its load level is truly 0%. In contrast, e.g., a web service during off-hours has a low load level that is typically still $>0\%$.

The early runs of SPEC Power in Figure 5 show the widespread introduction of power-saving mechanisms targeting active idle: From the earliest runs in 2006, where idle consumes a mean 70.1% power compared to full load (the *idle fraction*), the yearly mean drops to its minimum of 15.7% in 2017. Since then, the yearly mean idle fraction has increased again to 25.7% in 2024, marking a regression in idle-specific power optimizations. Intel seems more affected: In Figure 5 Intel's runs follow an upward trend, whereas AMD has a slightly falling trend – although there are both low and high idle fraction systems from both CPU vendors.

In an attempt to explain some of this recent development, we explored possible correlations between various run features, including the idle fraction. This exploration of runs since 2021 showed that the CPU vendor lineups, as well as submitted runs affect many features, confounding possible correlations. Most prominently, the core count of AMD (mean 85.8) is greater than that of Intel (mean 39.5). A more subtle example is the nominal frequency, where AMD and Intel share the same mean ($\sim2.3\,\text{GHz}$) but differ by spread (standard deviation $0.3\,\text{GHz}$ vs $0.5\,\text{GHz}$). Our correlation analysis to explain the recent development of the idle fraction remains inconclusive.

To better understand idle power optimizations, we introduce the *extrapolated* active idle power consumption: We extrapolate the power consumed at active idle through linear regression from the power consumed at 20% and 10% load. The result represents the power consumption during active idle if there would be no specific optimizations for full idle (rather than just individual idle cores). We then divide this extrapolated by the measured active idle power consumption and plot it over time in Figure 6. We refer to this quotient as *extrapolated idle quotient*. Higher values correspond to more effective idle-specific power optimization, 1 corresponds to none at all. However, higher values might also indicate a worse energy proportionality at

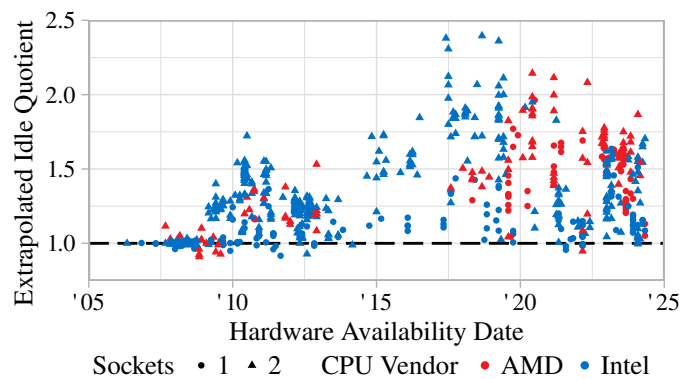
Figure 5: Idle power consumption trend

Figure 6: Trend of extrapolated vs measured active idle power

Table I: Comparison of two dual processor Lenovo systems, for the benchmarks SPECpower_ssj 2008, SPEC CPU Floating Point Rate Base, and SPEC CPU Integer Rate Base; *Factor* refers to the relative AMD/Intel performance difference

Benchmark	Result	Factor	System	CPU	TDP	Date	OS	RAM
power_ssj 2008	15112	1		Intel Xeon			Windows Server 2019 Datacenter	256
CPU 2017 FP	926	1	SR650 V3	Platinum 8490H	350 W	Feb 23	SUSE Linux Enterprise Server 15 SP4	512
CPU 2017 Int	902	1		1.90 GHz			Red Hat Enterprise Linux release 9.0 (Plow)	512
power_ssj 2008	31634	2.09		AMD EPYC 9754			Windows Server 2022 Datacenter	384
CPU 2017 FP	1420	1.53	SR645 V3	2.25 GHz	360 W	Aug 23	SUSE Linux Enterprise Server 15 SP4	1536
CPU 2017 Int	1830	2.03					SUSE Linux Enterprise Server 15 SP4	1536

low loads. Although Figure 6 has an upward trend overall, there is a large spread, in particular in newer runs. Idle-specific optimizations are not universally effective in recent runs.

The reasons behind this trend are obscured by two indistinguishable mechanisms. On the one hand, we suspect that processor architectures have an increasingly large share of power being used by shared resources, such as caches and on-chip communication, but also corresponding idle power-saving implementations for these resources. This effect could increase the extrapolated idle quotient if the latter effectively benefits from the energy-saving techniques. On the other hand, we speculate that it is becoming more difficult to effectively leverage idle power-saving mechanisms.

Consider, for instance, background tasks that are replicated for each logical CPU. Their activity prevents the system from fully utilizing idle states for short times each. With increasing core counts in recent processor generations, more of those tasks are running, reducing the relative time spent in the most efficient idle states. The first effect – architecturally low measured active idle – can explain an increased ceiling of the extrapolated idle quotient. The second effect – more difficult effective idle – can explain the large variation. From the quantitative data alone, we cannot fully distinguish the compound effects.

V. Limitations and Generalization

Since we only evaluate data from one benchmark, there is an argument to be made regarding the possibility of generalizing our observed energy efficiency trends. We evaluated other benchmarks that include power/efficiency: The TOP500/Green500 [6] are of limited value with respect to single-node performance due to their additional complexity (e.g., scalability challenges, inter-node networks). SPEC OMP 2012 [7] and SPEC ACCEL [14] perform floating-point heavy parallel workloads, representing typical parallel scientific workloads. However, their power measurement support never gained track with submitters; only 8 and 27 submissions include them, respectively. Other benchmarks, such as SPEC CPU [15], include much more general workloads than SPEC Power, with equally rich and well-reviewed publicly available datasets, but lack the power measurements required for efficiency analyses.

To assess the similarity to floating-point workloads, we resorted to screening the SPEC CPU results for recent runs of similar class CPUs with similar TDP from the same vendor that are also available in our SPEC Power dataset. Table I shows one example of two Lenovo nodes with Intel/AMD CPUs, both powered by 1100 W power supply units, and compares the results of SPEC Power and SPEC CPU Rate Base (throughput). In accordance with our expectations, the relative performance difference between the two systems is similar for SPEC Power and SPEC CPU integer, while AMD's performance advantage in the SPEC CPU floating-point benchmark suite is less pronounced. The integer-heavy SPEC Power workload favors AMD CPUs, while Intel's 2x advantage in AVX register width reduces the performance gap for floating-point calculations on wide vectors. Therefore, the observed energy efficiency trends can not be generalized to floating-point workloads.

Our results can also not be generalized to accelerator-based systems. While alternative benchmarks can make use of such devices (Green500, SPEC ACCEL), we did not consider them for the reasons mentioned previously.

VI. Conclusion and Future Work

Our analysis of 16 years of SPEC Power benchmark results shows continuous increases in *power consumption* of x86 processors. While there are certainly physical limitations to this growth, they are not yet visible in the data.

Energy efficiency also increases continuously and substantially, with AMD clearly providing superior efficiency in recent years. Due to the integer-heavy properties of the SPEC Power benchmark, this observation may not be generalized to more floating-point-intensive workloads. We also observe a positive trend towards better *energy proportionality* for both CPU vendors; although this trend is not universal.

Our analysis of *active idle power* data shows a conclusive trend towards lower consumption between 2006 and 2017, driven by the introduction of successively more effective sleep state mechanisms. Since then, a substantial share of runs show a regression in idle-specific power optimizations. We believe that the high variation in the results serves as an indication that particular attention should be paid to practical active idle power in the hardware selection and system operation. Especially for systems that may spend substantial time in active idle, such as HPC systems, idle power optimizations can improve economical and ecological performance.

The major *limitations* of this analysis are a lack of data for floating-point workloads, for other processor architectures such as ARM, and for accelerators such as AMD or NVIDIA graphics processing units (GPUs). The latter will hopefully be addressed by the SPECpowerNext benchmark [9]. This would be the industry-standard, vendor-driven benchmark, filling an important gap and enabling *future work* on GPU energy efficiency trend analysis.

Acknowledgments

Availability

We provide all scripts to reproduce this paper (parsing, analysis, and plotting), together with all data (raw and processed) online [16].

References

[1] SPEC, "Design document ssj workload specpower_ssj2008," 2012, https://www.spec.org/power/docs/SPECpower_ssj2008-Design_ssj.pdf.

[2] J. von Kistowski, K.-D. Lange, J. A. Arnold, S. Sharma, J. Pais, and H. Block, "Measuring and benchmarking power consumption and energy efficiency," in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 57–65, DOI: 10.1145/3185768.3185775.

[3] C. Gough, I. Steiner, and W. Saunders, *Energy Efficient Servers: Blueprints for Data Center Optimization*, 2015, DOI: 10.1007/978-1-4302-6638-9.

[4] C.-H. Hsu and S. W. Poole, "Power signature analysis of the specpower_ssj2008 benchmark," in *(IEEE ISPASS) IEEE International Symposium on Performance Analysis of Systems and Software*, 2011, pp. 227–236, DOI: 10.1109/ISPASS.2011.5762739.

[5] ——, "Revisiting server energy proportionality," in *2013 42nd International Conference on Parallel Processing*, 2013, pp. 834–840, DOI: 10.1109/ICPP.2013.99.

[6] W.-c. Feng and K. Cameron, "The green500 list: Encouraging sustainable supercomputing," *Computer*, vol. 40, no. 12, pp. 50–55, 2007, DOI: 10.1109/MC.2007.445.

[7] M. S. Müller, J. Baron, W. C. Brantley, H. Feng, D. Hackenberg, R. Henschel, G. Jost, D. Molka, C. Parrott, J. Robichaux, P. Shelepugin, M. van Waveren, B. Whitney, and K. Kumaran, "Spec omp2012 — an application benchmark suite for parallel systems using openmp," in *OpenMP in a Heterogeneous World*, B. M. Chapman, F. Massaioli, M. S. Müller, and M. Rorro, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 223–236, DOI: 10.1007/978-3-642-30961-8_17.

[8] U.S. Environmental Protection Agency, "Energy star® program requirements product specification for computer servers final draft test method," 2023, https://www.energystar.gov/sites/default/files/asset/document/Computer%20Servers%20Version%204.0%20Final%20Draft%20Test%20Method.pdf.

[9] N. Schmitt, K.-D. Lange, S. Sharma, N. Rawtani, C. Ponder, and S. Kounev, "The specpowernext benchmark suite, its implementation and new workloads from a developer's perspective," in *Proceedings of the ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 225–232, DOI: 10.1145/3427921.3450239.

[10] M. Meissner, K.-D. Lange, S. Sharma, J. Arnold, A. Cragin, P. Galizia, M. Petrich, B. Zhang, and S. Kounev, "Challenges and future directions in efficiency benchmarking (vision paper)," in *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '23 Companion. New York, NY, USA: Association for Computing Machinery, 2023, p. 51–55, DOI: 10.1145/3578245.3585034.

[11] SPEC, "SPECpower_ssj2008 Result File Fields," 2018, https://www.spec.org/power/docs/SPECpower_ssj2008-Result_File_Fields.html.

[12] American Society of Heating Refrigerating and Air-Conditioning Engineers, "Emergence and Expansion of Liquid Cooling in Mainstream Data Centers," Tech. Rep., May 2021. [Online]. Available: https://www.ashrae.org/file%20library/technical%20resources/bookstore/emergence-and-expansion-of-liquid-cooling-in-mainstream-data-centers_wp.pdf

[13] D. Hackenberg, R. Schöne, T. Ilsche, D. Molka, J. Schuchart, and R. Geyer, "An energy efficiency feature survey of the intel haswell processor," in *2015 IEEE International Parallel and Distributed Processing Symposium Workshop*, 2015, pp. 896–904, DOI: 10.1109/IPDPSW.2015.70.

[14] G. Juckeland, W. Brantley, S. Chandrasekaran, B. Chapman, S. Che, M. Colgrove, H. Feng, A. Grund, R. Henschel, W.-M. W. Hwu, H. Li, M. S. Müller, W. E. Nagel, M. Perminov, P. Shelepugin, K. Skadron, J. Stratton, A. Titov, K. Wang, M. van Waveren, B. Whitney, S. Wienke, R. Xu, and K. Kumaran, "Spec accel: A standard application suite for measuring hardware accelerator performance," in *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation*, S. A. Jarvis, S. A. Wright, and S. D. Hammond, Eds. Cham: Springer International Publishing, 2015, pp. 46–67, DOI: 10.1007/978-3-319-17248-4_3.

[15] J. Bucek, K.-D. Lange, and J. v. Kistowski, "Spec cpu2017: Next-generation compute benchmark," in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 41–42, DOI: 10.1145/3185768.3185771.

[16] H. Tröpgen, R. Schöne, T. Ilsche, and D. Hackenberg, "Artifacts to Reproduce "16 Years of SPEC Power: An Analysis of x86 Energy Efficiency Trends"," Jul. 2024. [Online]. Available: https://doi.org/10.5281/zenodo.12656360