A Survey of Energy Measurement Methodologies for Computer Systems

David Freina
Vrije Universiteit Amsterdam

Matthijs Jansen
Vrije Universiteit Amsterdam

Animesh Trivedi
Vrije Universiteit Amsterdam

Abstract

Energy measurement systems for computer systems have emerged as a promising approach to analyze the energy consumption of one or more computer systems. This paper presents a comprehensive review of current state-of-the-art energy measurement solutions for computing systems, both hardware-based and software-based. Furthermore, it presents two parallel timelines: the evolution of computer systems and the advancements in energy measurement technologies within the same domain. We present a detailed taxonomy of the solutions suggested in the existing literature on energy measurement systems. We also present the limitations and challenges of current energy measurement methodologies.

1 Introduction

With regard to the energy usage of computer systems, one of the most significant changes is the increasing number of data centers. In their Energy Efficiency Directive, the European Union (EU) published that data centers consumed 76.8TWh of energy in 2018. This equals 2.7% of the total energy consumption within the EU. Furthermore, they predict an increase in energy consumption of 28% (98.6TWh) until 2030 which would equal 3.21% of total usage. [12]

However, data centers are only one step in the evolution of computer systems that is characterized by a series of significant technological breakthroughs and architectural changes. From the onset of multi-core processors, the landscape of computing has undergone transformations with far-reaching consequences. These advances have not only enhanced the performance and computational efficiency, but also increased the energy usage of computer systems that are in use today.

To combat the energy consumption trends detailed in the preceding paragraph, there is a growing interest in enhancing the energy efficiency of computing systems. To accurately examine the efficiency of two systems, comparing the energy consumed to complete an identical task allows for a direct evaluation, with the system requiring less energy being deemed more efficient. Therefore, we will analyze state-of-the-art energy measurement systems in Section 4. They promise accurate and fine-granular measurements of a computer system's energy usage, which are essential to reliably ensure the efficiency of a system. Furthermore, we present an overview of the limitations and challenges that these systems face today in Section 9.

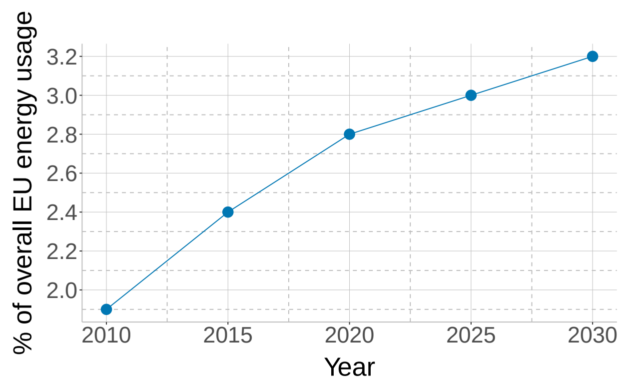

Figure 1: Energy usage of data centers in EU proportional to overall electrical energy demand in EU (predicted data from 2020-2030) [11]

The demand for reliable, accurate, and fine-grained energy measurements is further inflated under increasing pressure from legislative bodies. From May 2024 onward, data center owners and operators within an EU member state are required to report their annual energy usage according to the EU Energy Efficiency Directive. Consequently, the development of precise methodologies for energy measurement has become a highly relevant area of research and development in both academia and industry.

To gain a comprehensive understanding of the functioning of modern energy measurement systems and their focus on areas for improvement, it is crucial to examine the historical development of computer systems and the concurrent advancement of energy measurement technology. Initially, the main approach for measuring computer systems involved coarse-grained measurements that evaluated entire systems, without differentiating between individual components or specific processes [13,25,35]. However, Section 3 details how these approaches have evolved significantly, so that nowadays it is possible to measure the energy usage of most components of a system, such as GPUs or hard drives, at much higher frequencies (e.g. [34]). This progression was facilitated by the increased utilization of specialized hardware and enhanced vendor support for direct measurement gathering from system components.

Contributions: The contributions of this survey revolve mainly around answering the following research questions:

RQ-1: How did the evolution of computer systems influence the development of modern energy measurement solutions?
RQ-2: What are the current state-of-the-art energy measurement systems?
RQ-3: What are their current challenges or limitations?
RQ-4: How does the field prepare for the future (IoT, Edge, Large-Scale computing)?

Therefore, we contribute the following points through this work:

- We offer a thorough examination of the interconnected development of computer systems and energy measurement infrastructure, with the aim of assisting the reader to comprehend current research trends and areas of focus in energy measurement.

- We present a detailed taxonomy of the solutions suggested in the existing literature on energy measurement systems for computer systems, both hardware-based and software-based.

- Using this taxonomy allows us to discuss the differences in quality between different energy measurement systems that aim to provide the same type of measurement. Furthermore, it allows us to analyze the limitations and challenges that different systems may face.

- In the final part of the survey (Section 10) we want to provide an outlook for the future by analyzing energy measurement trends for emerging computer systems such as Exascale computing infrastructure.

2 Related Work

There is existing research on energy measurement systems that overlaps with this survey. However, most of the existing work focuses on specific areas of the field, while our objective is to provide a comprehensive view of the whole field.

Software-based energy measurement The most notable and closest related work on software-based energy measurement is performed by Jay et al. (2023) [22]. They conduct a comparative analysis of software-based power measurement solutions with a focus on CPU and GPU energy consumption. Their classification of these solutions includes three categories: *Energy Calculators*, *Energy Measurement Software*, and *Power Profiling Software*. We diverge from their classification and classify software-based power meters as either *based on specification* or *based on performance counters (PMCs)*.

Hackenberg et al. (2013) [19] provide a comprehensive overview of various methodologies applicable for power consumption measurements. Their experimental validation confirms the simplicity and accuracy of common methodologies for both AC and DC measurements. Furthermore, their analysis of Intel RAPL and AMD APM reveals certain limitations: providing energy data rather than power data limits the applicability of RAPL for detailed analysis, while AMD APM demonstrates issues with *systematic inaccuracies*. We compare their findings to more recent ones to answer the question of whether Intel RAPL is a usable tool for energy measurements.

Hardware-based energy measurement We compare hardware-based energy measurement systems [2,25,26,34] with regard to their measurement frequency, accuracy, and applicability.

Measurement quality We have identified three primary factors that impact the quality of measurement. They are derived from the collective analysis of various other studies.

Da Costa et al. (2017) [8] investigate the challenges and biases associated with power measurements using power meters and performance counters. This study provides significant information regarding the overhead of power measurement and its impact on the measurements themselves, the precision discrepancies between system-level and hardware-level sensors, and synchronization issues in large-scale systems. The accuracy of hardware-based and software-based energy measurement systems is reported by [2,19,20,22,25,26,34]. Furthermore, our discussion about the importance of high temporal granularity of measurements is motivated by [17,20,22].

Limitations & Challenges We present the limitations and challenges of state-of-the-art energy measurement systems which are covered by other work but not presented in a combined manner.

We have analyzed the existing work on standardization of the energy measurement infrastructure [3,5,7,17]. Furthermore, we extend [38] by identifying more possible points of heterogeneity in modern computer systems.

3 Parallel evolution: Computing and energy measurement

This section presents a comparative analysis of two parallel timelines: the evolution of computer systems and the advancements in energy measurement solutions within the same domain. The goal is to clear up how the rapid technological progress in computing, from the inception of multi-core processing to the latest trends in hardware specialization, has influenced and been mirrored by the development of energy measurement methodologies.

This comparison is motivated by the growing demand for high-performance computing systems and the corresponding need for more sophisticated energy measurement solutions, as discussed in section 1. The energy measurement systems mentioned in this section will be examined in more detail in the following sections.

Figure 2 presents a joint timeline of advancements in commonly used computer systems that create a challenge for energy measurement systems and the corresponding energy measurement solutions that solve these challenges.

Clock Speed: Between 1993 and 2003, there was a notable trend in processor technology where clock frequencies doubled approximately every 18 to 24 months [30]. This increase in clock speed, allowing processors to perform more operations within the same timeframe, necessitated an enhancement in the resolution of energy measurement systems to accurately capture the power usage of computer systems. This need for higher resolution marked a significant advancement from earlier efforts in system power measurement, such as those documented by Lapujade et al. [25]. In 2003, Isci et al. [21] are able to achieved much higher resolution in measuring the power usage of CPUs.

Multi-Core: The mid-2000s marked the onset of a transition to multi-core processors in mainstream systems, driven by the thermal and physical constraints of increasing clock speeds and reducing die sizes in single-core systems [24]. This shift significantly improved parallel processing capabilities and overall computational efficiency [1, 14]. This migration necessitated an adaptation in energy measurement systems to maintain their capability to measure energy information for individual processes because these processes could now run concurrently on one or more cores. One notable example that allows for per-process energy measurements is the Energy Consumption Library [7].

RAM: Concurrently with the advent of multi-core architectures, RAM technology also underwent several revisions. These incremental updates led to substantial increases in clock rate and bus rate (for instance, the bus rate from the original DDR in 1998 to DDR3 in 2007 increased by up to 10 times). Subsequent generations, such as DDR4 and DDR5, achieved even higher clock and bus speeds. In contrast to the previous example, advancements in RAM technology have not prompted substantial changes in energy measurement systems. The relatively low power consumption of RAM and the enhanced efficiency of newer generations [10] have resulted in limited focus on developing RAM-specific energy measurement solutions. However, two notable exceptions are Intel's RAPL [9] and WattProf [34], which do address energy measurement for RAM.

Cloud Computing: The mid-2000s introduced a pivotal architectural shift in computer systems, particularly from an energy usage perspective, with the emergence of cloud computing. This period saw major industry players like Amazon, Google, and Microsoft launching cloud computing-focused products between 2006 and 2008, marking a significant transition in the landscape of computing infrastructure:

- Elastic Compute Cloud (EC2) released in 2006/2007
- Google App Engine (GAE) released in 2008
- Microsoft Azure released in 2008

[33] Today data centers account for 2.7% of the total energy consumption within the EU, with a predicted increase of up to 3.21% by 2030 [12]. To effectively analyze the energy usage of facilities housing numerous computer systems, energy measurement systems have had to evolve to support distributed data collection. Early publications such as [37] already present a measurement system that is scalable to multi-node clusters, however their data collection rate is limited to 4 samples per second per node. Solutions such as PowerInsight [26] improve upon this, by providing per-node measurement infrastructure that offers out-of-band collection of measurement data to allow for large scalability while still retaining high measurement frequency and granularity. This helps to accommodate the scale and complexity of data center environments.

Hardware Specialization: A more recent trend in computing hardware is the move towards hardware specialization for specific tasks. An early instance of this trend can be traced back to 2006 with NVIDIA's release of CUDA, which facilitated the use of GPUs for general-purpose computing (GPGPU) [28]. Furthermore, the development of more specialized hardware accelerators, such as TPUs (from 2015 onwards [23]) for specific applications like machine learning, exemplifies the ongoing trend towards task-specific optimization in computing hardware.

As the diversity of hardware in computing systems expands, there is a growing demand for standardized, multichannel measurement infrastructure to accommodate this variety. WattProf [34] addresses this need by offering a solution

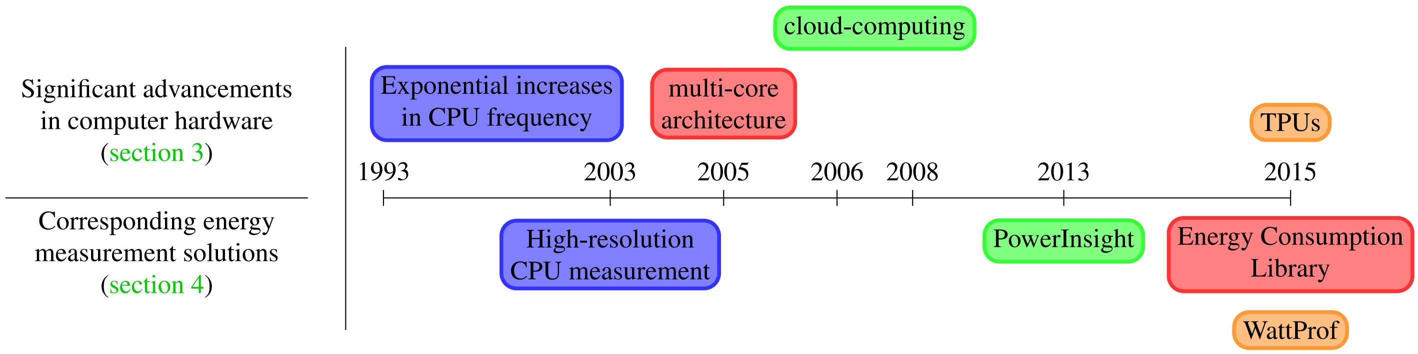

Figure 2: Timeline of computer system and energy measurement system evolution

equipped with sensor boards that are compatible with various connectors and capable of supporting up to 128 measurement channels.

Observation 1 (O-1): Over the past 30 years, the field of computer system energy measurement has transitioned from low resolution, system-wide measurements to high resolution, component-specific methodologies. This evolution has been driven by advances in (commodity) computer systems through the development of sophisticated hardware and software tools for high-resolution energy monitoring in real time.

O-2: Energy measurement in computer systems has become increasingly important, primarily due to growing environmental concerns and the need for energy-efficient technology in response to escalating energy costs and regulatory pressures.

4 Taxonomy of state-of-the-art energy measurement systems

To effectively organize and provide a comprehensive overview of the field of energy measurement systems, we present a structured taxonomy in Figure 3. The primary focus of this survey is on energy measurement systems; however, a thorough understanding of these systems requires an examination of the metrics used in measurement. Furthermore, to allow for a structured discussion about the topic of energy measurement, the exploration of the survey is divided into three major categories. We selected these categories because the majority of energy measurement tools we examined included a description of their function within these three categories.

- Domain: The specific areas or components within computer systems where energy measurement is applied.
- Methodologies: The various techniques and approaches used to measure energy in these systems.
- Quality Criteria: The standards that define the quality of energy measurement methods and measured data.

The subsequent sections will dive into each category of energy measurements, along with a discussion of the associated metrics.

5 Taxonomy: Energy measurement metrics

Understanding the metrics used to present measured data is crucial when discussing any type of measurement. In the context of this survey, when we refer to energy, we specifically mean electrical energy, which is typically quantified in kilowatt-hours (kWh). However, it is worth noting that some publications may use joules (J) as the unit of measurement because joules represent the general energy that can be converted into electrical energy. The conversion between these two units is given by the equation:

$$3.6MJ = 1kWh$$

Apart from electrical energy, there is another significant metric for energy measurement known as (electric) power, which is measured in watts (W). Power is also related to both metrics mentioned above, and the relationship can be expressed as:

$$1W = 1J/s$$

Consequently, 1 kWh represents the amount of energy transferred over the course of one hour with a power output of 1000W.

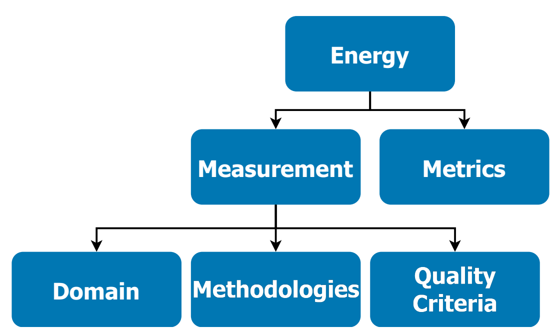

Figure 3: The main focus points of this survey

6 Taxonomy: Energy measurement domain

We are able to identify four different major domains that can categorize current and past energy measurement systems for the taxonomy in Figure 4. These classes are chosen based on the measurement domains that we have observed in various projects that were analyzed.

The evolution of measurement technologies has shown (Section 3), that a significant evolution is evident in the measurement domains of energy measurement solutions in various projects. Initially, up until 2003, the focus of energy measurement was predominantly at the node-level. Those measurements can range from simple approaches, such as connecting a wattmeter to each node's PSU, to more advanced methods like employing Smart PDUs that are able to measure per-node data in serverracks in data centers. Subsequently, there was a shift towards component-level measurement, where most projects focus on CPU measurements with options to extend the measurement coverage to some or all I/O components like HDDs, GPUs, or networking hardware. Ultimately, the scope expanded to encompass infrastructure-level measurements, reflecting the broadening scale of computer systems in data centers. At the infrastructure-level, energy consumption is typically monitored through electric meters or similar devices.

However, it is important to recognize that the measurement domain of an energy measurement system is not restricted to a single domain. Certain projects, such as the Energy Consumption Library [7], demonstrate this by offering the capability to implement a sensor at the node-level and the application-level. This flexibility is important to allow for more comprehensive energy measurement approaches that can accommodate a wide range of monitoring needs across different levels of computing systems.

7 Taxonomy: Energy measurement methodologies

Through combining the classifications of Jay et al. [22] and Qi et al. [32] which focus primarily on software-based energy measurement systems and the methodologies observed in various other projects, we create our taxonomy. This taxonomy presents two methodologies that are employed today in energy measurement systems: software-based energy measurement and hardware-based energy measurement. Both methodologies are further segregated in their respective subsections.

7.1 Software-based energy measurement systems

For the scope of this survey, energy measurement systems that do not depend on the usage of additional hardware being present in a system are classified as software-based energy measurement systems. They obtain their measurement data through estimations based on hardware specifications ([1, 7, 21]) or through the energy measurement infrastructure already available built into the hardware ([5, 9]).

Estimation based on specification Estimation processes can be based on the thermal design power (TDP) of the hardware provided by the manufacturer (e.g., CodeCarbon). The TDP specifies the maximum heat generated by a component in a system. Although the TDP does not exactly give an estimate on the CPU power usage, it can be used in conjunction with information about the average CPU usage and total execution time to obtain an estimation of the total CPU power usage. [22]

Estimation based on performance counters (PMCs) Hardware PMCs offer an alternative method to estimate the energy usage of a system. These counters are specialized registers that record data on hardware events, such as cache accesses and misses, or CPU clock cycles and instructions. [7]

A vendor-agnostic example of an energy measurement project that implements an estimator would be the Energy Consumption Library [7] which is designed to estimate the energy use of applications and / or machines. The sensors used by the library are mainly based on hardware performance counters mixed with operating system data. It allows the estimation of energy usage on an application or machine basis by either assigning sensors to specific PID's or the whole machine. The library already includes one linear estimator and two dynamic estimators. The former estimator is based on the CPU usage and minimum and maximum power consumption of the machine and uses this information to estimate a power usage. However, rather than being limited to these estimators, the library is created to serve as a baseline with

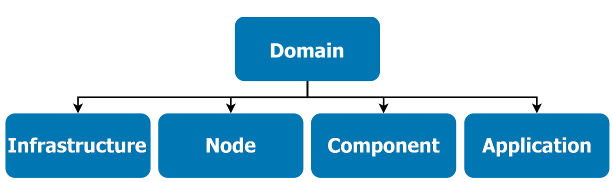

Figure 4: Subdivisions of energy measurement domains

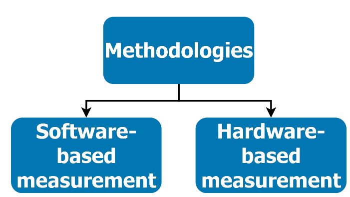

Figure 5: Subdivisions of energy measurement methodologies

easy extension options to allow for more sophisticated power estimators and custom power sensors.

Measurement using model-specific registers (MSRs) It is important to note that the support of manufacturers for integrated energy measurement circuits has improved in 2011 when both Intel and AMD include new measurement capabilities in their respective CPU generations [19]. Those measurement circuits are available to the kernel through model-specific registers (MSRs) that provide detailed information on the energy usage of the processor. Intel's Running Average Power Limit (RAPL) [9] and AMD's Average Power Management (APM) are two vendor-specific tools that estimate energy usage. These two tools derive their estimates by accessing the data stored in MSRs. [19] A third manufacturer, NVIDIA, also introduces API commands one year later in their NVIDIA Management Library (NVML) [29] that allow one to obtain energy and power measurement data.

A different example using the energy measurement infrastructure that is already available in modern hardware is the Energy Measurement Library [5]. This library leverages hardware-specific interfaces such as Intel RAPL or NVIDIA NVML to obtain power readings directly from the hardware. It is developed to act as a middle layer between the hardware and the software to allow architectural and experimental independence when measuring the energy consumption of a system.

Most other software-based energy measurement projects that have been introduced or updated after MSRs are available include support for Intel RAPL (e.g., PowerAPI [3], Scaphandre [31], Energy Scope [27], and Perf [36].

O-3: Support for MSRs is added to most software-based energy measurement systems making it the state-of-the-art methodology to obtain energy data from a computer system.

7.2 Hardware-based energy measurement systems

Hardware-based energy measurement systems, in contrast to software-based approaches, require the use of additional hardware to gather measurement data from computer systems. Within this category, a distinction can be made between systems that utilize commodity hardware and those that require specialized hardware, such as purpose-built measurement boards or FPGA-based technologies.

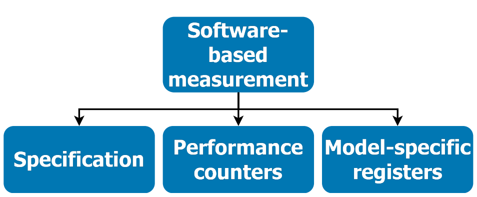

Figure 6: Subdivisions of software-based energy measurement

7.2.1 Commodity Hardware

Commodity hardware in this context typically includes external power meters and (smart) Power Distribution Units (PDUs). Examples of such hardware include OmegaWatt meters, Eaton power meters, and APC units. These devices are generally used to measure the total energy consumption of a whole system. However, they mostly lack the capability to provide fine-grained, component-specific data. The performance of these devices in terms of accuracy and granularity can vary significantly depending on the specific model and manufacturer. While they offer a more accessible and cost-effective solution for energy measurement, their utility can be limited by their inability to provide insights into both the energy usage of individual system components and the total system usage.

However, there is research, notably in 2003, that marked a substantial improvement in measurement frequency compared to the earlier coarse-grained measurements presented in 1994 and 1996 ([13, 25]). By employing a clamp ammeter connected to a digital multimeter, this study achieved a granular measurement rate of 1000 power readings per second with a resolution of 0.12W. However, this approach was limited to measuring a single system component, specifically the CPU. [21]

Furthermore, WattScope [18] presents an ML-based solution that utilizes node-level and rack-level power measurement infrastructure that is widely available in data centers. The authors claim that WattScope can provide application-level power usage data by disaggregating node-level and rack-level power measurements. However, their normalized mean absolute error across 1,100 servers running on average 40 jobs is almost 10%. [18]

7.2.2 Specialized Hardware

Unlike using commodity hardware, some projects have employed specialized hardware for energy measurement. This specialized hardware typically offers superior measurement capabilities, targeting one or more of the different measurement quality criteria detailed in Section 8.

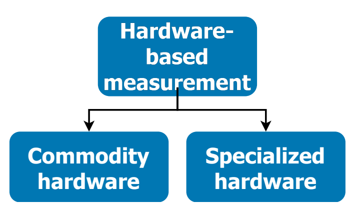

Figure 7: Subdivisions of hardware-based energy measurement

Section 3 comprehensively outlines the significant advancements in methodologies for measuring energy and power consumption in computer systems over time. This evolution is evident in the development of specialized hardware designed for energy measurement in these systems. Research for this survey identifies the first example of custom hardware tailored for energy measurements in host systems, dating back to 1994. In this initial study, the energy usage of entire computer systems and their peripherals, such as printers and monitors, was assessed using power meters combined with data loggers. In particular, these power meters were developed in a unique way and called "multiple outlet monitors (MOMs)". The MOM was essentially an enhanced power strip equipped with measurement capabilities that allowed monitoring of up to seven devices simultaneously. [35]

Further advances in this area can be observed with the publication of PowerMon and PowerMon2[1]. Similarly to the measurements using MOM from the previous paragraph, PowerMon/PowerMon2 supports 6/8 measurement channels, respectively. However, the transition to in-band monitoring represents a significant shift in the methodology of energy measurement. Unlike the MOMs, which are out-of-band measurement devices positioned between the power source and the Power Supply Unit (PSU) of the computer system, in-band measurement involves placing the measurement hardware between the computer system's PSU and its individual components. This change in hardware placement enables monitoring at the component level, allowing the collection of separate measurement data for CPUs, hard drives, and GPUs, and thus providing a more detailed and component-specific analysis of energy consumption. Furthermore, PowerMon provides measurements at a rate of up to 50Hz per channel, while PowerMon2 enhances this capability, achieving up to 1024Hz on a single channel or a combined 3072Hz across aggregated channels. [2]

During the research for this survey, PowerInsight [26] is identified as the next step in hardware-based energy measurement systems, building on the foundations laid by PowerMon/PowerMon2. It notably expands the number of measurement channels, increasing from the 6/8 channels offered by PowerMon/PowerMon2 to 15 channels. A key feature of PowerInsight is its ability to collect out-of-band data, enabling centralized data aggregation from multiple PowerInsight units through a control node. Although the exact measurement frequency is not explicitly detailed, the authors of PowerInsight assert a rate "faster than 1KHz" [26] indicating a similar level of measurement granularity to that of PowerMon2.

Improving on the capabilities of PowerMon/PowerMon2 and PowerInsight, which are standalone devices with PowerMon2 and PowerInsight fitting in a 3.5" drive bay, WattProf [34] introduces a distinct approach with its main monitoring component. This component, called the "monitoring board," is designed as a PCIe expansion card. Like PowerMon/PowerMon2 and PowerInsight, WattProf is primarily focused on delivering a system capable of fine-grained and highly granular measurements. However, the WattProf monitoring board further expanded the measurement domain and frequency, allowing up to 128 channels of sensor measurements with a resolution of 12 kHz per channel. This design significantly exceeds PowerMon/PowerMon2 and PowerInsight in terms of data collection capacity, offering enhanced granularity and more detailed energy measurement capabilities. Like PowerMon, PowerMon2, and PowerInsight, WattProf achieved this through the use of highly specialized hardware. However, the significant expansion in measurement domain and frequency achieved by WattProf, compared to PowerMon / PowerMon2 and PowerInsight, is attributed to its use of a Xilinx Spartan Field-Programmable Gate Array (FPGA) for power monitoring. Unlike the microcontroller-based PowerMon/PowerMon2 and the ARM-based PowerInsight, the FPGA in WattProf provides a highly customizable platform. This flexibility allows authors to tailor the central controller programming specifically to their measurement needs, enabling more precise and extensive monitoring capabilities.

O-4: Over the last three decades, there have been significant advances in specialized hardware for energy measurement systems. These advances have resulted in notable improvements in both measurement frequency, which has increased from 0.2Hz to 12kHz, and applicability, with systems now capable of supporting 128 measurement channels for one node, enabling measurement of multiple components simultaneously.

8 Taxonomy: Energy measurement quality criteria

In Figure 8, the quality criteria for energy measurements are categorized into three distinct sections, each addressing a crucial aspect of measurement quality:

- Accuracy: The reliability of an energy measurement system is highly dependent on the accuracy of the data it captures. Accuracy is primarily influenced by the quality and capabilities of the sensors used in the system, as well as the frequency at which these sensors can collect data. High accuracy ensures that measurements are reliable and reflect the actual energy consumption of the system.

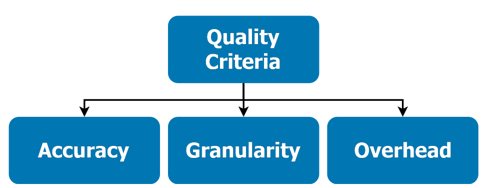

Figure 8: Subdivisions of energy measurement quality

[1]The capabilities of PowerMon and PowerMon2 differ slightly; hence, they will be referred to as x/y in the following text, where x represents PowerMon and y represents PowerMon2

- Granularity: Measurement granularity is divided into two subcategories: spatial and temporal granularity.
  - Spatial Granularity: This refers to the system's ability to obtain measurements from specific components of the system; it was previously defined as the "Measurement Domain" in Section 6.
  - Temporal Granularity: This aspect concerns the frequency with which the measurement system can provide data. Temporal granularity, like accuracy, is dependent on the capabilities of the sensor.
- Overhead: Software-based energy measurement systems are normally run locally on the machine being measured. For reliable measurements, the overhead of the software needs to be accounted for if it influences the measurements. In contrast to this, hardware-based energy measurement systems with out-of-band measurement capabilities do not incur any additional load on the system that is measured.

[20]

We derived these quality criteria by looking at the evaluation sections of various energy measurement projects and surveys that compared them. [22, 26, 34] However, we found that the overhead aspect was often overlooked when evaluating projects.

8.1 Accuracy

The accuracy of energy measurement systems varies between different projects, each with its unique specifications. PowerMon and PowerMon2, for example, claim to exhibit a worst-case accuracy range of -6.6% to +6.8% [2]. PowerInsight, on the other hand, presents an accuracy of $\pm$ 5% compared to precision voltmeters and ammeters, and achieved a coefficient of variation of 2.45% in an experiment involving a single node High Performance Linpack (HPL) benchmark. [26] For WattProf the authors claim an approximate total expected error of $\pm$5.65% in the worst-case scenario, while most measurements should fall within a more accurate range of $\pm$3.5%. [34]

Hackenberg et al. [19] conducted a comprehensive comparison involving Intel's Running Average Power Limit (RAPL), AMD's Average Power Management (APM), their own DC power measurement solution, and AC power measurement. They found that both Intel RAPL and AMD's APM experience systematic issues under certain workloads that negatively affect measurement accuracy. For Intel RAPL, they noted that the measurements remain relatively reliable, with systematic inaccuracies staying below 5%, as long as the energy data polling frequency is kept below 20 samples per second.

In recent years, however, it appears that either the systematic issues with Intel RAPL have been addressed, or that tools implementing Intel RAPL for energy measurement are not significantly affected by these issues. Jay et al. [22] demonstrated last year that various software-based energy measurement solutions (such as PowerAPI, Scaphandre, Energy Scope, and Perf) that utilize Intel RAPL can accurately measure the energy consumption of different benchmarks, closely aligned with readings from external power meters. They also observed that, for CPU-based benchmarks, a higher sampling rate tends to result in a stronger correlation with the values from external power meters. Even in the worst-case scenario, Perf, with a sampling frequency of 10Hz, achieved a high correlation of 0.93 with the readings of the external power meter.

8.2 Temporal Granularity

Modern energy measurement systems exhibit varying levels of temporal granularity, influenced by their architecture and the aggregation of data across different layers. At the lowest level, hardware measurements might be captured at frequencies as high as 10kHz, but these measurements are often aggregated at the node level and may be further aggregated at the infrastructure level, resulting in a lower frequency of measurements available to the user. The importance of the temporal granularity exposed to the user varies depending on the specific application of the energy measurement system and the characteristics of the tasks being monitored. Tasks that frequently change or have high variability require finer temporal granularity to capture accurate energy consumption data, whereas long-running tasks with stable system usage may be adequately monitored with coarser granularity. [17]

The significance of temporal granularity on measurement accuracy is underscored by a study from Jay et al. [22], as discussed in the previous Section 8.1. This study illustrates how lower temporal granularity can negatively impact measurement precision. Hackenberg et al. [20] further support this finding through their work with a SLURM plugin based on IPMI for power measurements in a High-Performance Computing (HPC) cluster. The standard IPMI power measurement plugin, which is limited to 1 sample per second with a 3W granularity for node-level measurements, was found to be insufficient for accurate power monitoring. By developing a custom IPMI OEM extension capable of 4 samples per second and utilizing the maximum granularity offered by the sensors, they demonstrated that increased temporal granularity and measurement detail lead to more precise power consumption assessments.

8.3 Overhead

Software-based measurement systems that are run locally on the same hardware that is getting measured inevitably influence the outcome of the measurement. However, Jay et al. [22] report an average energy overhead of sub 1% for their tested software-based energy measurement systems which they conclude to be insignificant to the overall measurements. Additionally to purely software-based systems, even hardware-based ones that do not utilize out-of-band data collection normally require a software package that is run locally to collect the measurement data. However, newer projects like PowerInsight [26] or WattProf [34] allow the collection of measured data on dedicated remote monitoring machines which will remove the overhead from the measured host system.

Furthermore, monitoring the PCs is not free and also lead to some measurable overhead. Depending on the amount and frequency of the monitored PCs the overhead can be significant to the measurements accuary. With 240 measured PCs with a sampling rate of 1 sample per second the worst case increase in power consumption can be as high as 4% according to [8]. However, the authors note that this increased is measured on a system that is otherwise idle which boosts the relative increase significantly. The absolute increase is measured with 0.87W to 1.31W which will result in a lower relative increase when the system is under load.

9 Limitations & Challenges

Accuracy In Section 8, we have already examined the factors that restrict the accuracy of measurements, such as the quality and frequency of the data obtained from the sensors. Additional inaccuracies in measurements may arise during the conversion of analog to digital values and in subsequent data processing procedures. To minimize these errors, the most effective approach is to employ state-of-the-art measurement hardware and analog-to-digital converters (ADCs) that provide the most precise measurements and conversion capabilities. Additionally, it must be ensured that possible data processing steps are clearly understood in relation to their implications on the accuracy of measured data. [20]

Standardization Furthermore, the standardization of energy measurements to enable interconnected use of multiple libraries remains an open area for research. We have identified the following approaches that aim to provide a unified interface to users while supporting multiple libraries to obtain data from the system:

- Energy Measurement Library (EML) [5]
- Power API [17]
- PowerAPI [3]
- Energy Consumption Library [7]

We found that the first two projects have not gained widespread adoption and seem abandoned at the time of writing. According to Google Scholar, [5] has been cited 32 times, and the repository [6] seems to have been inactive since the last commit 4 years ago.

[17] has been cited 47 times according to Google Scholar, but the only mention of a publicly available repository that we could find is in a presentation from 2020 [15] and it appears that the repository [16] is no longer available.

Furthermore, we are unable to find any available source code or download for the Energy Consumption Library [7].

However, the third PowerAPI [3] project still has an active and available repository [4]. It allows the usage of various sensors as sources for energy measurements. They can range from hardware power meters (e.g., Smart PDUs) to performance counters (e.g., RAPL).

Heterogeneity As compute platforms become more diverse and include additional hardware, such as accelerators, the traditional model of a computer system consisting of a CPU, GPU, hard drives, RAM, and expansion cards such as network cards may no longer be sufficient. In the future, systems could be much more heterogeneous, with different hardware accelerators, possibly even on a single chip with large FPGAs, combined with multiple general-purpose processors. This trend towards heterogeneity is further exemplified by modern CPU architectures like ARM's big.LITTLE or Intel's use of P-cores and E-cores. Consequently, there is a need for specialized measurement infrastructure that can accurately capture the unique power profiles exhibited by this type of hardware. [38] Projects like WattProf [34] have already made progress in this area by utilizing specialized hardware built on FPGAs to enable per-component monitoring. However, further research is encouraged by us to develop support for energy measurement on these highly heterogeneous systems.

O-4: We can observe that the challenges and limitations that we identified are already being addressed by various projects to varying degrees. Changes in computer systems have created opportunities for research in the past (Section 3) and continue to do so.

10 Conclusion and Future Outlook

Investigation into energy measurement systems for computer systems has underscored the capability of achieving detailed and precise energy measurements. However, significant gaps remain, particularly in the realm of cloud environments housed within large data centers. Despite scalable hardware-based measurement solutions such as PowerInsight [26], granular energy measurement in cloud computing infrastructure has not seen widespread implementation, possibly due to economic factors and the intrusive nature of such measurement systems. Projects like WattScope [18] aim to address these challenges by leveraging existing infrastructure and advanced software techniques, although they face limitations in accuracy.

The rapidly evolving hardware landscape, marked by the introduction of TPUs, multiarchitecture CPUs, and the expansion into edge computing, IoT devices, and exascale data centers, presents additional challenges to the applicability of current energy measurement methodologies. These developments require innovative approaches to capture reliable energy data accurately.

The future advancement of energy measurement technologies should focus on creating non-intrusive, scalable, and hardware-agnostic methods. There is promising potential in software-based energy estimation models that utilize machine learning and statistical analysis to bridge the current gaps. Furthermore, the establishment of standardized energy measurement protocols on a variety of computing platforms, including next-generation architectures such as ARM and RISC-V, is crucial.

References

[1] Basmadjian, R., and De Meer, H. Evaluating and modeling power consumption of multi-core processors. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet* (Madrid Spain, May 2012), ACM, pp. 1–10.

[2] Bedard, D., Lim, M. Y., Fowler, R., and Porterfield, A. PowerMon: Fine-grained and integrated power monitoring for commodity computer systems. In *Proceedings of the IEEE SoutheastCon 2010 (SoutheastCon)* (Concord, NC, USA, Mar. 2010), IEEE, pp. 479–484.

[3] Bourdon, A., Noureddine, A., Rouvoy, R., and Seinturier, L. PowerAPI: A Software Library to Monitor the Energy Consumed at the Process-Level. *ERCIM NEWS* (2013), 43.

[4] Bourdon, A., Noureddine, A., Rouvoy, R., and Seinturier, L. PowerAPI. https://github.com/powerapi-ng/powerapi, 2024.

[5] Cabrera, A., Almeida, F., Arteaga, J., and Blanco, V. Measuring energy consumption using EML (energy measurement library). *Computer Science - Research and Development 30*, 2 (May 2015), 135–143.

[6] Cabrera, A., Almeida, F., Arteaga, J., and Blanco, V. Energy Measurement Library. https://github.com/HPC-ULL/eml, 2018.

[7] Cupertino, L. F., Da Costa, G., Sayah, A., and Pierson, J.-M. Energy Consumption Library. In *Energy Efficiency in Large Scale Distributed Systems*, J.-M. Pierson, G. Da Costa, and L. Dittmann, Eds., vol. 8046. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 51–57. Series Title: Lecture Notes in Computer Science.

[8] Da Costa, G., Pierson, J.-M., and Fontoura-Cupertino, L. Mastering system and power measures for servers in datacenter. *Sustainable Computing: Informatics and Systems 15* (Sept. 2017), 28–38.

[9] David, H., Gorbatov, E., Hanebutte, U. R., Khanna, R., and Le, C. RAPL: memory power estimation and capping. In *Proceedings of the 16th ACM/IEEE international symposium on Low power electronics and design* (Austin Texas USA, Aug. 2010), ACM, pp. 189–194.

[10] Desrochers, S., Paradis, C., and Weaver, V. M. A Validation of DRAM RAPL Power Measurements. In *Proceedings of the Second International Symposium on Memory Systems* (Alexandria VA USA, Oct. 2016), ACM, pp. 455–470.

[11] European Commission. Directorate General for Communications Networks, Content and Technology., Environment Agency Austria., Borderstep Institute, and Umweltbundesamt. *Energy-efficient cloud computing technologies and policies for an eco-friendly cloud market: final study report.* Publications Office, LU, 2020.

[12] European Parliament. Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 (recast).

[13] Floyd, D. B., and Parker, D. S. Measured Performance of Energy-Efficient Computer Systems. *H&H - Symposium on Improving Building Systems in Hot and Humid Climates* (1996).

[14] Geer, D. Chip makers turn to multicore processors. *Computer 38*, 5 (May 2005), 11–13.

[15] Grant, R. E. PowerAPI: A Standardized Interface to Power/Energy Monitoring and Control. https://www.osti.gov/servlets/purl/1798421, 2020. Accessed: 2024-02-11.

[16] Grant, R. E., Levenhagen, M., Olivier, S. L., DeBonis, D., Pedretti, K. T., and Laros Iii, J. H. Power API. https://github.com/pwrapi/puwerapi, 2016.

[17] Grant, R. E., Levenhagen, M., Olivier, S. L., DeBonis, D., Pedretti, K. T., and Laros Iii, J. H. Standardizing Power Monitoring and Control at Exascale. *Computer 49*, 10 (Oct. 2016), 38–46.

[18] Guan, X., Bashir, N., Irwin, D., and Shenoy, P. WattScope: Non-intrusive Application-level Power Disaggregation in Datacenters. Publisher: arXiv Version Number: 1.

[19] Hackenberg, D., Ilsche, T., Schone, R., Molka, D., Schmidt, M., and Nagel, W. E. Power measurement techniques on standard compute nodes: A quantitative comparison. In *2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)* (Austin, TX, USA, Apr. 2013), IEEE, pp. 194–204.

[20] Hackenberg, D., Ilsche, T., Schuchart, J., Schone, R., Nagel, W. E., Simon, M., and Georgiou, Y. HDEEM: High Definition Energy Efficiency Monitoring. In *2014 Energy Efficient Supercomputing Workshop* (New Orleans, LA, USA, Nov. 2014), IEEE, pp. 1–10.

[21] Isci, C., and Martonosi, M. Runtime power monitoring in high-end processors: methodology and empirical data. In *22nd Digital Avionics Systems Conference. Proceedings (Cat. No.03CH37449)* (San Diego, CA, USA, 2003), IEEE Comput. Soc, pp. 93–104.

[22] Jay, M., Ostapenco, V., Lefevre, L., Trystram, D., Orgerie, A.-C., and Fichel, B. An experimental comparison of software-based power meters: focus on CPU and GPU. In *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)* (Bangalore, India, May 2023), IEEE, pp. 106–118.

[23] Jouppi, N., Young, C., Patil, N., and Patterson, D. Motivation for and Evaluation of the First Tensor Processing Unit. *IEEE Micro 38*, 3 (May 2018), 10–19.

[24] Knight, W. Two heads are better than one [dual-core processors]. *IEE Review 51*, 9 (Sept. 2005), 32–35.

[25] Lapujade, P., Parker, D. S., and Center, F. S. E. Measured Energy Savings of an Energy-Efficient Office Computer System. *ACEEE Sumer Study on Energy Efficiency in Buildings 9* (1994).

[26] Laros, J. H., Pokorny, P., and DeBonis, D. PowerInsight - A commodity power measurement capability. In *2013 International Green Computing Conference Proceedings* (Arlington, VA, USA, June 2013), IEEE, pp. 1–6.

[27] Limpens, G., Moret, S., Jeanmart, H., and Maréchal, F. EnergyScope TD: A novel open-source model for regional energy systems. *Applied Energy 255* (Dec. 2019), 113729.

[28] NVIDIA Corporation. About CUDA. https://developer.nvidia.com/about-cuda. Accessed: 2024-01-24.

[29] NVIDIA Corporation. NVML API REFERENCE MANUAL. https://developer.download.nvidia.com/assets/cuda/files/CUDADownloads/NVML/nvml.pdf, 2012. Accessed: 2024-01-24.

[30] Parkhurst, J., Darringer, J., and Grundmann, B. From Single Core to Multi-Core: Preparing for a new exponential. In *2006 IEEE/ACM International Conference on Computer Aided Design* (Double Tree Hotel, San Jose, CA,USA, Nov. 2006), IEEE, pp. 67–72. ISSN: 1092-3152.

[31] Petit, B. Scaphandre. https://github.com/hubblo-org/scaphandre, 2021.

[32] Qi, X., Chen, J., Dong, Y., Yuan, Y., Xu, T., Deng, R., Li, Z., Zhou, K., and Wang, Z. HighRPM: Combining Integrated Measurement and Sofware Power Modeling for High-Resolution Power Monitoring. In *Proceedings of the 52nd International Conference on Parallel Processing* (Salt Lake City UT USA, Aug. 2023), ACM, pp. 369–379.

[33] Qian, L., Luo, Z., Du, Y., and Guo, L. Cloud Computing: An Overview. In *Cloud Computing*, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, M. G. Jaatun, G. Zhao, and C. Rong, Eds., vol. 5931. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 626–631. Series Title: Lecture Notes in Computer Science.

[34] Rashti, M., Sabin, G., Vansickle, D., and Norris, B. WattProf: A Flexible Platform for Fine-Grained HPC Power Profiling. In *2015 IEEE International Conference on Cluster Computing* (Chicago, IL, USA, Sept. 2015), IEEE, pp. 698–705.

[35] Szydlowski, R., and Chvala, W. J. Energy consumption of personal computer workstations. *ACEEE Sumer Study on Energy Efficiency in Buildings 2* (Feb. 1994).

[36] Torvalds, L. perf.c. https://github.com/torvalds/linux/blob/master/tools/perf/perf.c, 2024.

[37] Xizhou Feng, Rong Ge, and Cameron, K. Power and Energy Profiling of Scientific Applications on Distributed Systems. In *19th IEEE International Parallel and Distributed Processing Symposium* (Denver, CO, USA, 2005), IEEE, pp. 34–34.

[38] Zoni, D., Galimberti, A., and Fornaciari, W. A Survey on Run-time Power Monitors at the Edge. *ACM Computing Surveys 55*, 14s (Dec. 2023), 1–33.