

Comparison and Analysis of Software and Hardware Energy Measurement Methods for a CPU+GPU System and Selected Parallel Applications

Grzegorz Koszczał, Mariusz Matuszek, and Paweł Czarnul

Faculty of Electronics, Telecommunications and Informatics
Gdańsk University of Technology
Narutowicza 11/12
80-233 Gdańsk, Poland
pczarnul@eti.pg.edu.pl

Abstract. In this paper authors extend upon their previous research on power-capped optimization of performance-energy metrics of deep neural networks training workloads. A professional power meter Yokogawa WT-310E is used, as well as Intel RAPL and Nvidia NVML interfaces, to examine power consumption of a much more comprehensive set of multi-GPU and multi-CPU workloads, including: selected kernels from NAS Parallel Benchmarks for CPUs and GPUs as well as Horovod-Python Xception deep neural network training using several GPUs. A comparison and discussion of results obtained by both power measurement methods has been performed using 2 systems, one with 2 Intel Xeon CPUs and 8 Nvidia Quadro RTX 6000 GPUs and the second 2 Intel Xeon CPUs and 4 Nvidia Quadro RTX 5000 GPUs. We compared power consumption between hardware and software interfaces for CPU, GPU and mixed CPU+GPU workload configurations, using 1-40 threads for the CPUs and 1-8 GPUs.

Keywords: high performance computing, performance-energy optimization, energy measurements, measurement accuracy, parallel benchmarks.

1. Introduction

In recent years an increased awareness of need for energy conservation can be observed, together with ever-increasing demand for high performance computing (HPC), either in a traditional or cloud-presented form. These two factors combined led to a series of research on performance-energy optimizations in high-performance computing environments, as well as introducing dedicated power measurement interfaces by Intel (RAPL) and Nvidia (NVML). In this paper we use both a Yokogawa WT-310E professional power meter [41] and software interfaces by Intel and Nvidia to examine reported power draw under a comprehensive set of multi-GPU, multi-CPU and mixed workloads, in order to gain a better insight into capabilities and limitations of those software interfaces.

This paper is a very significantly extended version of conference paper [25]. While the original work focused on power-capped optimization of performance-energy metrics of multi-GPU training of deep neural networks using a hardware power meter, this paper focuses on a thorough comparison of software and hardware power/energy measurement methods using the aforementioned application (Xception DNN training) and a set of selected NAS Parallel Benchmarks, run on CPUs, GPUs and CPU+GPU configurations.

1 The paper is organized in a standard pattern. We begin by highlighting relevant re-
2 search papers in this field. Next we present our motivation for research and list our main
3 contributions. Then a description of our measurement methodology is given, followed by
4 reasoning behind our choice of testbed applications. A large part of this paper is then
5 dedicated to a detailed description of our experiments and discussion of results. We fin-
6 ish the paper with a summary of results and short discussion on possible future work
7 and an appendix, containing detailed measurements including power values from individ-
8 ual computing devices and entire nodes, execution times and standard deviations for the
9 measurements.

10 2. Related Work

11 Intel Running Average Power Limit (RAPL) is a set of interfaces for reporting the accu-
12 mulated energy consumption of various power domains [17].

13 NVIDIA Management Library (NVML) is a C-based API that allows monitoring and
14 managing various states of the NVIDIA GPU devices, including their power draw [32].

15 The related work can be broadly assigned to two research categories. One is exam-
16 ination and validation of RAPL and NVML power measurement interfaces themselves.
17 The second one is using measurements obtained from said interfaces as inputs for some
18 energy/performance optimization scheme. Additionally, there are papers dedicated to re-
19 search on the accuracy and quality of standard benchmark patterns. Papers [30] and [6]
20 present NAS Parallel Benchmarks and their implementations, while papers [3] and [4]
21 discuss efficient NAS Parallel benchmarks on GPUs.

22 Hähnel et al. [14] report on using Intel RAPL energy sensors to measure energy con-
23 sumption of short code paths, with execution times substantially shorter than RAPL up-
24 date interval. The authors describe operating system level modifications made to ensure
25 code execution synchronous with RAPL register updates.

26 Lang and Runger [28] propose a statistical method for generating high resolution
27 power profiles on Nvidia GPUs using standard (lower) resolution power measurements
28 returned by NVML.

29 A power consumption model focused on rendering graphics of varying complexity on
30 a mobile GPU is proposed by Vatjus-Anttila et al. [39]. The model is based on graphics
31 primitives, such as triangles, render batches and texels, and thus is intended to be hardware
32 agnostic. The proposed approach allows to predict the complexity of any given 3D scene
33 at a content production phase. The authors verify their model with measurements on a
34 real-world content and hardware and report prediction errors in range of 0.3% to 3.2%.

35 A quantitative evaluation of the Intel’s RAPL power control system is presented by
36 Zhang and Hoffmann [42]. The authors evaluate the RAPL system by setting target power
37 limits and running a set of benchmarks, quantifying the results using metrics of accuracy,
38 stability, settling time, overshoot and efficiency.

39 Desrochers et al. [11] conducted detailed analysis of RAPL measurements (using the
40 perf event interface) concerning the CPU and DRAM versus system wide numbers from
41 a WattsUpPro? device. The testbed system included 3 machines with 2 desktop and 1
42 server Haswell CPUs. Three different types of DDR3 DRAM as well as two types of
43 DDR4 DRAM were tested. The authors concluded that, in general, usually RAPL mea-
44 surements follow the total power value trends, typically by a constant offset. In general

1 measurements matched within 20% with following relative phase behavior. Additionally,
2 in terms of memory, results matched best when the DRAM was heavily utilized, but were
3 not accurate where the system was idle or memory was used by an integrated GPU. The
4 authors also noted that the Haswell server machines returned more accurate results from
5 actual power measurements.

6 Lucas and Juurlink [29] examine power consumption dependency of arithmetic-logic
7 units (ALU's) in modern GPUs on data values being processed. They show large power
8 consumption differences between the same kernel processing different data. Use of mi-
9 crobenchmarks to characterize power consumption of GPU functional units is discussed
10 and obtained data is used to propose a simple and fast model of power consumption of
11 functional units, in order to improve accuracy of power consumption models.

12 Ferro et al. [13] use internal GPU power sensors to examine power profiles of two
13 Nvidia GPUs under both artificial (benchmark) and real application loads. Cards from two
14 different generations are measured and a comparison of accuracy between generations is
15 made. Power consumption of a whole node is also monitored using IPMI interface and
16 tool.

17 Ikram et al. [15] propose an experimental methodology for evaluating power and en-
18 ergy consumption of programs executing on Nvidia Kepler GPUs. The methodology is
19 applied to two common HPC programs: a matrix multiplication and a parallel sorting.
20 The authors conclude, that their methodology can be applied to any program executing
21 on Nvidia Kepler GPU, to obtain measurements of peak and average power, energy and
22 kernel runtime.

23 Khan et al. [23] performed detailed analysis of RAPL accuracy and usefulness for
24 power measurements. They used both customized benchmarks as well as Stream, Stress-
25 ing and ParFullCMS applications as well as two power measurement datasets from the
26 Taito – a supercomputing system of the Finnish Center of Scientific Computing. They
27 determined that measurements using RAPL are very highly highly correlated (≈ 0.99)
28 with AC power, with a high sampling rate ≈ 1 ms that even allows to distinguish various
29 phases of an application and negligible overhead less than 1% for standalone systems and
30 less than 2.5% for Amazon EC2. They concluded that RAPL can be used for measur-
31 ing energy consumption of servers without power meters. Selected, desired features that
32 would be useful were identified, including: reduction of the delay of reading RAPL MSRs
33 through hypervisors and more resolution, i.e., per core readings, not only per package.

34 Sen et al. [35] present a quality assessment methodology of GPU power profiling
35 mechanisms. Using the proposed methodology the authors assess the quality of four pro-
36 filing techniques: NVML via Allinea MAP, NVML via direct reads, PowerInsight (PI)
37 via vendor-provided drivers and PI via Allinea MAP. In addition the authors discuss the
38 influence of moving-average filters on the slow variation of some of measured power
39 profiles.

40 Paniego et al. [33] focus on monitoring and analyzing energy consumption in HPC
41 environment for a given application and architecture. A matrix multiplication application
42 on a shared-memory architecture is used to compare values reported by Intel RAPL with
43 physical measurements obtained through the processor power source. The authors report
44 that, for the application considered, there is an error of up to 22% between the average
45 CPU power and values predicted by RAPL.

1 Fahad et al. [12] investigated accuracy of on-chip power sensors and predictive models
2 versus an external hardware power meter for measurement of dynamic energy of appli-
3 cations. They ran two scientific applications, matrix-matrix multiplication and 2D fast
4 Fourier transform, for various data sizes, on three Intel multi-core CPUs, two Nvidia
5 GPUs and an Intel Xeon Phi device and concluded that the average error using on-chip
6 power sensors can be as high as 73% and using predictive models with performance moni-
7 toring counters can be as high as 32% for Haswell and Skylake CPUs, for dynamic energy
8 profiles.

9 Kavanagh and Djemame [20] use RAPL and IPMI interfaces to come up with tuned
10 models for estimation of host-level power consumption, for use in Cloud and High Per-
11 formance Computing platforms. The methodology of models tuning and calibration, as
12 well as mitigating errors present in both interfaces is discussed in detail.

13 Ilsche [16] analyzed, in particular, the accuracy of RAPL measurements, for vari-
14 ous architectures of Intel CPUs, specifically: Sandy Bridge, Haswell and Skylake. For
15 SandyBridge, after testing several workloads, for the package domain correlation between
16 RAPL and reference power has been found to be weak with RAPL values being higher
17 or lower, depending on the workload. For DRAM, close correlation was observed for
18 higher power consumption while visible discrepancy was observed for smaller values.
19 This suggests that usage of RAPL is limited in that generation as a replacement for real
20 measurements, according to the author. In that generation RAPL used an internal, archi-
21 tectural model with a set of events and consideration of weights for prediction of power
22 consumption. From Haswell, an implementation based on physical measurements was in-
23 troduced. The author also evaluated RAPL for the Skylake generation. The accuracy of
24 measurements for that generation was shown to be considerably better, with relative dis-
25 crepancy for package + DRAM max. at 3.8% and 3.3% without idle measurement point
26 (but with much higher discrepancies for separate package and DRAM components).

27 Arafa et al. [2] present a detailed measurement of energy consumption of different
28 instructions that can be executed on Nvidia GPGPUs. Three different techniques to read
29 on-chip power sensors are being used and a comparison of these techniques is made.
30 Additionally, obtained measurements are verified against an external, custom-designed,
31 hardware power measuring device.

32 Aslan and Yilmazer-Metin [5] present a study on power and energy measurement in
33 Nvidia Jetson embedded GPUs, by validating and extending the methodology presented
34 by Burtscher et al. [8], originally developed for measurement of GPU power using the
35 K20 built-in sensor.

36 Shahid et al. [36] investigated various energy predictive models for multi-core CPUs,
37 performing tests for two systems with Haswell and Skylake CPUs and a variety of bench-
38 marks including: NPB kernels, MKL FFT, HPCG, MKL DGEMM, stress and naive matrix-
39 matrix and matrix-vector multiplication. They have determined, having analyzed the pre-
40 diction accuracy of linear energy models based on utilization variables only, PMCs only
41 and both utilization variables and PMCs, that the best models with both utilization vari-
42 ables and PMCs resulted in 3.6 and 2.6 times better average prediction accuracy compared
43 to the models based on utilization variables only and PMCs only.

44 Krzywaniak et al. [26,27] analyzed application of power capping for multi-core CPUs
45 and GPUs, in the context of performance-energy optimization considering metrics such as
46 Energy Delay Product (EDP), Energy Delay Sum (EDS). For a given parallel application

1 and a computing device, the goal was to find such a power cap so that the value of a
2 given metric was optimized, compared to the default power cap. Since this process was
3 performed dynamically at runtime by the proposed DEPO tool, available in versions for
4 the CPU and the GPU, both performance and power/energy were measured automatically,
5 even without the need for code instrumentation. For a given power cap, in the case of the
6 CPU, application progress is measured using an instruction counter and energy using
7 Intel RAPL. For a code running on a GPU, application progress is measured using a
8 kernel invocation counter while power is measured using Nvidia NVML which allows
9 computation of energy. Measurements are taken in an adjustable time window that should
10 be short enough for quick browsing of a range of available power caps and sufficiently
11 long for the measurements to be stable and representative. DEPO can browse the available
12 range of adjustable power caps using one of two algorithms: Linear Search (LS) and
13 Golden Section Search (GSS). As a result, the authors concluded that for the purpose of
14 the stated optimization, the accuracy of Intel RAPL and Nvidia NVML was sufficient,
15 as demonstrated by Krzywaniak et al. [26] (RAPL compared to HPE Metered 3Ph 22
16 kVA/60309 5-wire 32 A/230 V Outlets (30) C13 (3) C19/Vertical INTL PDU (D9N56A)
17 with $\pm 1\%$ or better accuracy in power monitoring) and [27] (NVML compared to power
18 meter WT310) respectively.

19 A very low-level study of energy cost associated with dynamic branch prediction in an
20 Intel CPU is given by Alqurashi and Al-Hashimi [1]. The authors use the RAPL interface
21 for reporting the CPU power and energy, while using known micro-benchmarks under
22 various run conditions to explore potential pitfalls in the measurement interface.

23 Departing from the prevailing Intel RAPL and Nvidia NVML tune, Tröpgen et al.
24 [38] show an evaluation of the energy measurement present in the IBM POWER9 on-chip
25 controller (OCC). The authors provide a detailed description and in-depth evaluation of
26 OCC-provided power measurements for several power domains, confronting the results
27 with externally measured data.

28 Yang et al. [40] investigated details of power/energy measurement using Nvidia-smi.
29 They proposed a suite of micro-benchmarks to benchmark profiles using Nvidia-smi for
30 power readings and have evaluated for over 70 different GPUs from several generations.
31 For the steady state error evaluation, in the majority of the cases the error was within
32 $\pm 5\%$. Additionally, they determined that for H100 and A100, Nvidia-smi only reports
33 the average of the past 25ms every 100ms, while for the previous Volta and Pascal
34 generations 10/20ms periods were reported. The authors proposed to incorporate controlled
35 delays between repetitions to exposing various segments of an application to the identified
36 measurement window.

37 Shalavi et al. [37] study the accurate calibration of power measurements from Nvidia
38 Jetson devices. The authors propose and utilise a regression model to map sensor mea-
39 surements to real accurate power readings obtained from external hardware. The authors
40 found, that internal sensors in Jetson devices underestimated the power draw by up to
41 50% and, by applying calibration, were able to reduce the error to within $\pm 3\%$.

42 A very interesting and thorough experimental comparison of software-based power
43 meters was presented by Jay et al. [19], focusing on both CPU and GPU workloads. The
44 authors discussed various methods to measure energy in computer systems, including:
45 power meters, intra-node devices, hardware sensors (with respective software interfaces
46 such as Intel RAPL and Nvidia NVML), as well as power and energy modeling (including

1 using e.g. a combination of RAPL and hardware performance counter events). Several
2 software-based power meters are discussed and compared, including: Carbon Tracker,
3 Code Carbon, Energy Scope, Experiment Impact Tracker, Perf, Power API, Green Al-
4 gorithms, ML CO2 Impact, and Scaphandre. Furthermore, experimental evaluation was
5 conducted using a cluster with Nvidia DGX-1 nodes, each with 2 Intel Xeon (Broadwell)
6 CPUs, 8 Nvidia Tesla V100 GPUs and 512 GB RAM, with an Omegawatt external power
7 meter (sampling frequency of 1Hz) and node's BMC with a sampling frequency of 0.2Hz.
8 Selected NAS benchmark kernels were used: EP, LU and MG (different sizes for CPUs
9 and GPUs). Power profiles over time were analyzed from Energy Scope, Scaphandre,
10 Perf, PowerAPI, BMC, and the power meter for both the CPU and GPU. In general, the
11 Pearson correlation coefficient between the tools and the power meter was approx. 0.95
12 with the highest value for EnergyScope 0.972 (sampling frequency of 2Hz). What is very
13 interesting in the context of this work is that the offset between the software power meters'
14 and the external power meter' values is not constant and, according to the authors, shall
15 be studied for various architectures/computing node. Energy overheads of the tools were
16 evaluated as of approx. 1%. Additionally, the authors tested several benchmarks running
17 in parallel, also benchmarks of various types, i.e., compute and memory intensive, with
18 PowerAPI and Scaphandre giving mostly coherent results but some differences occurred
19 as well. One interesting observation from the results was that after GPU workloads, GPUs
20 would still take some 20-30 seconds to return to idle load, apparently due to engaged fans
21 etc.

22 Raffin and Trystram [34] [perform](#) a critical analysis of Intel RAPL energy measure-
23 ment operations in order to provide RAPL users with best access practices. Qualitative
24 highlights of differences between measurement mechanisms are given, including evalua-
25 tion of overheads for various mechanisms present. The paper is focused on a comparative
26 analysis of the measurement mechanisms and on experimental evaluation of performance
27 and energy measurements.

28 3. Motivations and Contribution

29 As energy-aware high performance computing has gained much attention [10,24], there
30 is a constant need for assessment of accuracy of various power and energy measurement
31 methods. This stems from the fact that methods such as Intel RAPL and Nvidia NVML
32 might be expected to be improved over time and also because of the fact that many system
33 configurations differ considerably in terms of setups such as numbers of CPUs, GPUs but
34 also density and additional devices that might be measured by hardware but not grasped
35 by software APIs, notably fans etc.

36 Our direct motivation for this research is following up on our previous research on
37 power-capped optimization of performance-energy metrics of multi-GPU training of deep
38 neural networks [25]. While in the latter work, we took power measurements under power
39 capping conditions using a professional hardware power meter Yokogawa WT-310E [41]
40 and used DNN training workloads only, in this paper we focus on comparing power
41 readings from Yokogawa and Intel RAPL and Nvidia NVML. For the comparison to be
42 meaningful, we broaden the scope of workloads with NPB Suite to better reflect the di-
43 versity of contemporary applications. Additionally, we want to explore potential impact
44 of possible power supply configurations on such measurements.

1 Our contribution is a comparison of power consumption readings between software
 2 and hardware measurement methods, for a set of workloads representative of HPC appli-
 3 cations, for a range of CPU and GPU loads, up to 2 multi-core CPUs and 8 GPUs. We also
 4 expose and discuss power consumption overheads which are not visible in the readings
 5 from software interfaces like Intel RAPL and Nvidia NVML.

6 4. Measurement Methodology

7 The methodology adopted in this paper assumes that we compare average power taken by
 8 a system running selected, representative benchmarks, using both hardware (power me-
 9 ter) and software based measurement methods (through Intel RAPL and Nvidia NVML)
 10 which are available and widely used by other researchers [19].

11 The ground truth for our measurements is provided by Yokogawa WT310E – an ex-
 12 ternal power meter that is a digital power analyzer that provides extremely low current
 13 measurement capability down to 50 micro-Amps, and a maximum of up to 26 Amps
 14 RMS. This device follows standards and certificates such as Energy Star, SPECpower and
 15 IEC62301/EN50564 testing. This model belongs to WT300E’s family of devices that of-
 16 fer a wide range of functions and enhanced specifications, allowing handling of all the
 17 measurement applications from low frequency to high frequency inverters using a single
 18 power meter. The WT300E series with a fast update rate of 100ms. The basic accuracy for
 19 all input ranges is 0.1% rdg + 0.05% rng (50/60Hz) and DC 0.1% rdg + 0.2% rng. In order
 20 to obtain readings from the Yokogawa WT310E power meter connected to a machine, a
 21 special software has been written – Yokotool [7]. Yokotool is a command-line tool for
 22 controlling Yokogawa power meters in Linux. The tool is written in Python and comes
 23 with the ‘yokolibs.PowerMeter’ module which can be used from Python scripts.

24 For software-based measurement APIs, we used Intel RAPL for obtaining energy and
 25 then power from the CPU(s) and the DRAM domain as well as Nvidia NVML for obtain-
 26 ing power taken by the GPU(s). [Energy measurements from RAPL are accessible by read-
 27 ing proper files exposed in a file system. Power values are reported by NVML by either us-
 28 ing the `nvidia-smi` command-line tool or calling the `nvmlDeviceGetPowerUsage`
 29 library function.](#)

30 While Yokogawa allows us to measure the total power of the whole system, the com-
 31 bined readings from RAPL and NVML are not able to account for additional system
 32 devices, including disks and especially cooling, including fans. By analyzing the differ-
 33 ence between the Yokogawa and combined RAPL+NVML measurements, under various
 34 configurations (CPU cores and GPUs used), for various applications, we can conclude
 35 how the difference changes with the configurations, applications and systems.

36 As the measurement frequency of the Yokogawa device is 10Hz, it was the frequency
 37 that we set for all the APIs, i.e., Yokogawa WT301E, Intel RAPL and Nvidia NVML.

38 Specifically, the following results will include computation of the offset (δ) as the
 39 difference between the subtraction of active node power draw (Yokogawa) (θ_Y) and the
 40 sum of active CPUs (RAPL) power draw (θ_R) and GPUs (NVML) power draw (θ_N) [W],
 41 and subtraction of idle node (Yokogawa) power draw (θ_{Y_i}) and the sum of idle CPUs and
 42 GPUs (RAPL+NVML) power draw [W] (θ_{R_i} and θ_{N_i} , respectively), computed for each
 43 configuration:

$$\delta = (\theta_Y - (\theta_R + \theta_N)) - (\theta_{Y_i} - (\theta_{R_i} + \theta_{N_i}))$$

$$= (\theta_Y - \theta_{Y_i}) - ((\theta_R + \theta_N) - (\theta_{R_i} + \theta_{N_i})) \quad (1)$$

1 5. Testbed Applications

2 In order to obtain results that are representative of a wide range of applications, it is best
 3 to rely on performing experiments using well established benchmarks. In our case, it is
 4 especially important that we target CPU, GPU and [mixed CPU+GPU workloads](#), for the
 5 purpose of comparison of power/energy measurement methods. Specific requirements in-
 6 clude: ability to run in parallel on various numbers of logical processors, run on one or
 7 more GPUs and being able to execute for various input data sizes so that configurations
 8 with reasonable compute to communication/synchronization time ratios can be selected so
 9 that parallelization is efficient, actual speed-ups are obtained [9]. We have decided to use
 10 the well-established NAS Parallel Benchmark Suite [as well as loads generated by Xception](#)
 11 [DNN training](#). The NPB Suite contains a series of kernels including: IS – Integer Sort
 12 (random memory access), EP – Embarassingly Parallel, CG – Conjugate Gradient (irreg-
 13 ular memory access and communication), MG – Multi-Grid on a sequence of meshes, FT
 14 – Discrete 3D fast Fourier Transform (all-to-all communication) as well as applications:
 15 BT – Block Tri-diagonal solver, SP – Scalar Penta-diagonal solver, LU – Lower-Upper
 16 Gauss-Seidel solver. These benchmarks are available in various sizes, in particular: classes
 17 A, B, C – standard test problems (roughly 4 times size increase from each of the previous
 18 classes) as well as classes D, E, F – large test problems (roughly 16 times size increase
 19 from each of the previous classes). In terms of the implementations tested, we used the
 20 following that satisfied the aforementioned criteria to put load on the CPU(s) and GPU(s):
 21 NAS Parallel Benchmarks (C++ with OMP) [30], NAS Parallel Benchmarks (Fortran with
 22 MPI) [6], NAS Parallel Benchmarks (CUDA) [3,4]. [In addition to the NPB Suite, a cus-](#)
 23 [tom deep learning model based on Xceptionnet with MPI communication was used. This](#)
 24 [model was tested previously with power capping for performance-energy optimization](#)
 25 [under power capping \[25\]](#). The rationale and methodology for selection of particular con-
 26 figurations is discussed in detail in Section 6.2.

27 It should be noted how sizes of the benchmarks are set in the following experiments,
 28 which stems from actual implementations for CPUs and GPUs. Specifically, the bench-
 29 marks for the CPU scale with a specific number of threads set which results in reduced
 30 execution times for larger numbers of threads. On the other hand, the compiled bench-
 31 marks for the GPU are run, by default, on a single GPU. For runs that use more than one
 32 GPU, we run the same benchmark independently on a given number of GPUs, which re-
 33 sults in essentially the same execution time of the test, regardless of the number of GPUs
 34 used. In the case of [mixed CPU+GPU benchmarks](#), we conducted a given test until the
 35 end of the shortest (out of the CPU and GPU) benchmarks.

36 6. Experiments

37 6.1. Testbed Systems

38 For the following tests, we used two systems with multicore CPUs and GPUs, with the
 39 following specifications:

1 **vinnana** : 2 x Intel(R) Xeon(R) Silver 4210 CPU (TDP 85W), 2.20GHz, for a total of
2 20 physical cores and 40 logical processors; 384 GB RAM, DDR4, 2400 MHz; 4 x
3 Nvidia Quadro RTX 5000 16GB (TDP 230W);
4 **sanna** : 2 x Intel(R) Xeon(R) Silver 4210 CPU (TDP 85W), 2.20GHz, for a total of
5 20 physical cores and 40 logical processors; 384 GB RAM, DDR4, 2400 MHz; 8 x
6 Nvidia Quadro RTX 6000 24GB (TDP 260W).

7 Each system uses an Inspur YZMB01130107 motherboard as well as 4 x Delta Elec-
8 tronics DPS-2200AB-2 PSUs and runs Linux Ubuntu 22.04 LTS operating system. Use
9 of two very similar systems allowed us to achieve two objectives:

- 10 1. validate our findings (in the load ranges common to the two systems),
- 11 2. investigate the power consumption overhead patterns between less and more densely
12 GPU equipped systems, which was made possible by otherwise identical configu-
13 rations of the two systems. The only difference between the two systems was the
14 number and models of the installed GPUs.

15 6.2. Tested Configurations

16 Tested configurations correspond to various parallelization levels of application runs, in-
17 teresting from the point of view of taking advantages of the CPUs' numbers of cores
18 and numbers of GPUs, installed in the two testbed systems. Consequently, we varied the
19 number of CPU threads between 2 and 32 threads, considering powers of 2 for selected
20 simulations which is a typical approach in parallel computing benchmarking. For others,
21 we set the number of threads between 1 and 20 for 1 CPU as well as between 2 and 40
22 for 2 physical CPUs which is adjusted to the actual CPU models and core and logical
23 processor counts described in Section 6.1. For the GPU tests, we varied the number of
24 GPUs used, as available in a given system: 1, 2 and 4 GPUs for vinnana and 1, 2, 4 and 8
25 GPUs in sanna. For **mixed CPU+GPU tests**, we combined increasing numbers of threads
26 and GPUs from the aforementioned configurations.

27 In order to provide a larger variety of tests on similar machines (different in the GPU
28 configuration) we have decided to conduct tests of various benchmarks/implementations
29 on the two machines, as follows:

30 **vinnana** : CPUs – MPI-Fortran¹, GPUs – Horovod-Python, mixed – MPI-Fortran +
31 Horovod-Python;
32 **sanna** : CPUs – OMP-CPP, GPUs – OMP-CUDA, mixed – OMP-CPP + OMP-CUDA.

33 In order to select configurations for testing, we adopted the following assumptions.
34 Considering the sampling frequency of Yokogawa WT-310E of 10Hz, we assumed that
35 each test should last at least 20 seconds (in fact guaranteeing 199 probes). We performed
36 initial tests in order to select benchmarks and corresponding classes (sizes) to satisfy this
37 criterion. The following benchmarks and configurations were run in this initial phase on
38 particular systems, each of which was run 3 times and average values were recorded:

¹ note that throughout the paper, in the case of MPI-Fortran workloads, we refer to computational threads.

- 1 Implementation: OMP-CPP: benchmarks available: IS, FT, EP, CG, MG, LU, SP and
 2 BT, class sizes to choose from: B, C and D, 24 different combinations of benchmarks
 3 and class sizes in total, all tests were run on 2 CPUs and 40 logical processors in
 4 parallel;
- 5 Implementation: OMP-CUDA: benchmarks available: IS, FT, EP, CG, MG, LU, SP and
 6 BT, class sizes to choose from: C, D and E, 24 different combinations of benchmarks
 7 and class sizes in total, all tests were run on a single GPU with the same grid sizes
 8 each time;
- 9 Implementation: MPI-Fortran: benchmarks available: IS, FT, EP, CG, MG and LU, class
 10 sizes to choose from: B, C and D, 18 different combinations of benchmarks and class
 11 sizes in total, all tests were run on 2 CPUs and 32 logical processors in parallel;
- 12 Implementation: Horovod-Python: benchmark available: Xception, number of training
 13 epochs tested: 1, 3 and 5, 3 different combinations of model training parameters to
 14 choose from, all tests were run in configurations with 1, 2 and 4 GPUs, to check
 15 training behavior.

16 Based on these tests and satisfying the aforementioned requirements, we selected the
 17 following configurations for subsequent power/energy investigation for the following test
 18 types:

- 19 **CPU on vinnana** : ep.D.x – Embarassingly Parallel, class size ‘D’, lu.C.x – Lower-
 20 Upper Gauss-Seidel solver, class size ‘C’, is.D.x – Integer Sort, class size ‘D’;
- 21 **CPU on sanna** : bt.C – Block Tri-diagonal solver, class size ‘C’, is.D – Integer Sort, class
 22 size ‘D’, lu.C – Lower-Upper Gauss-Seidel solver, class size ‘C’;
- 23 **GPUs on vinnana** : number of epochs for the training of the Xception model was se-
 24 lected to be 1 as the test accuracy of model is relatively high (86.8%) and the execu-
 25 tion time is satisfactory (48.5s using 4 GPUs).
- 26 **GPUs on sanna** : lu.D – Lower-Upper Gauss-Seidel solver, class size ‘D’, sp.D – Scalar
 27 Penta-diagonal solver, class size ‘D’, ep.D – Embarassingly Parallel, class size ‘D’;

28 Then **mixed** versions were selected as follows:

- 29 **mixed on vinnana** : ep.D.x + Xception, is.D.x + Xception, lu.C.x + Xception;
 30 **mixed on sanna** : bt.C + lu.D, is.D + sp.D, lu.C + ep.D.

31 6.3. Results

32 All the following results are averages from 10 runs. We have also recorded standard de-
 33 viations, which are reported in tables, for clarity. Before we proceed with presentation of
 34 results of the selected configuration runs, we demonstrate values of the measured com-
 35 ponents, for selected runs on vinnana. Specifically, Table 1 presents power values for
 36 vinnana, for Xception run on 1, 2 and 4 GPUs, measured using Yokogawa as well as
 37 RAPL and NVML, for CPUs and GPUs, along with execution times, the latter showing
 38 scalability of the simulation. Standard deviation values are provided as well showing very
 39 good stability of the results.

40 Similarly, Table 2 presents values of analogous variables for a **mixed** run of ep.D.x+
 41 Xception using a respective number of 1, 2 and 4 GPUs and 8, 16 and 32 processes for
 42 computations. Table 3 presents idle power values for vinnana, measured using Yokogawa

Table 1: Execution times and power draws; server: **vinnana**, device: **GPUs**, implementation: **Horovod-Python**, benchmark: **Xception**

	GPUs		
Results from 10 runs	1 GPU	2 GPUs	4 GPUs
Avg. Exec. time [s]	133.363	80.633	56.574
Std. dev. of time [s]	0.403	0.360	0.350
(Yokogawa) Avg. power draw [W]	511.645	665.957	905.241
(Yokogawa) Std. dev. of avg. power draw [W]	1.120	2.813	2.909
(CPUs) Avg. power draw [W]	55.354	67.363	81.370
(CPUs) Std. dev. of avg. power draw [W]	0.074	0.203	0.283
(GPU: 0) Avg. power draw [W]	169.563	156.733	135.377
(GPU: 0) Std. dev. of avg. power draw [W]	0.878	0.832	1.089
(GPU: 1) Avg. power draw [W]	14.416	159.318	138.517
(GPU: 1) Std. dev. of avg. power draw [W]	0.097	1.436	1.123
(GPU: 2) Avg. power draw [W]	11.060	10.967	133.840
(GPU: 2) Std. dev. of avg. power draw [W]	0.054	0.120	1.084
(GPU: 3) Avg. power draw [W]	14.637	14.873	137.662
(GPU: 3) Std. dev. of avg. power draw [W]	0.046	0.091	1.053

Table 2: Execution times and power draws; server: **vinnana**, device: **Hybrid**, implementation: **MPI-Fortran+Horovod-Python**, benchmark: **ep.D.x+Xception**

	Hybrid		
Results from 10 runs (GPU(s) + Processes)	1 + 8	2 + 16	4 + 32
Avg. Exec. time [s]	131.454	80.068	73.330
Std. dev. of time [s]	0.552	1.590	0.355
(Yokogawa) Avg. power draw [W]	561.049	741.187	915.957
(Yokogawa) Std. dev. of avg. power draw [W]	1.145	6.528	2.154
(CPUs) Avg. power draw [W]	99.713	135.506	163.343
(CPUs) Std. dev. of avg. power draw [W]	0.123	0.217	0.221
(GPU: 0) Avg. power draw [W]	171.163	156.899	112.283
(GPU: 0) Std. dev. of avg. power draw [W]	0.676	2.686	0.845
(GPU: 1) Avg. power draw [W]	14.324	159.284	115.304
(GPU: 1) Std. dev. of avg. power draw [W]	0.087	2.510	0.371
(GPU: 2) Avg. power draw [W]	11.058	11.096	111.088
(GPU: 2) Std. dev. of avg. power draw [W]	0.052	0.068	0.945
(GPU: 3) Avg. power draw [W]	14.670	14.960	115.475
(GPU: 3) Std. dev. of avg. power draw [W]	0.034	0.064	1.035

1 for the whole node, using RAPL for CPUs, NVML for GPUs as well as presents the
 2 difference between the Yokogawa value and sum of the two latter.

3 We now proceed with presentation of the measurements for the applications and con-
 4 figurations derived in Section 6.2. Each of the following figures presents values for one
 5 computing device type (CPUs, GPUs, **mixed** CPUs+GPUs) for one system: vinnana or

Table 3: Power draw measurements of idle **vinnana** server

	Idle power draw
Measured components	Power draw readings [W]
Sum of entire node idle power draw (Yokogawa)	314.693
Sum of CPUs idle power draw (Linux Perf / Intel RAPL)	46.967
Sum of GPUs idle power draw (NVML)	50.255
Difference between idle node power draw and sum of CPUs and GPUs power draw	217.471

1 sanna, for a total of 6 figures: vinnana and applications running on CPUs – Figure 1,
 2 vinnana and applications running on GPUs and on CPUs+GPUs – Figure 2, sanna and
 3 applications running on 1 CPU – Figure 3, sanna and applications running on 2 CPUs
 4 – Figure 4, sanna and applications running on GPUs – Figure 5, sanna and applications
 5 running on CPUs+GPUs – Figure 6.

6 Each figure, for each application presents three graphs: average power measured for
 7 the whole node by Yokogawa, sum of average powers obtained from RAPL and NVML
 8 (in case of sanna measurements include the DRAM component) as well as the offset
 9 computed using Equation 1.

10 The factory hardware configuration of both sanna and vinnana includes 4 power sup-
 11 plies for each server, for required redundancy. In order to find out the power cost of re-
 12 dundancy we removed 2 power supplies from each server and rerun our test suite for
 13 CPU, GPU and **mixed** CPU+GPU workloads. We observed an average offset of approx.
 14 40 Watts in the overall power consumption between the two configurations, as reported
 15 by the Yokogawa power meter.

16 6.4. Discussion

17 From all the charts, looking at average power values gathered from application runs, for
 18 vinnana, for a given number of threads and/or number of GPUs we can see slightly dif-
 19 ferent results for various workloads, with the exception of very similar power values for
 20 ep.D.x+Xception and lu.C.x+Xception for 1 and 2 GPUs (for 4 GPUs values are visibly
 21 different), as shown in Figure 2. For sanna, very similar power values can be observed for
 22 GPU tests using sp.D and lu.D (Figure 5) and **mixed** bt.C+lu.D and is.D+sp.D (Figure 6).
 23 Some differences can also be observed in the speed of growth of the power (angle) for
 24 various applications, for instance between lu.C.x and is.D.x for vinnana (Figure 1). Inter-
 25 estingly, **mixed** is.D.x+Xception is not able to put much larger load onto 4 versus 2 GPUs
 26 and results from Yokogawa and RAPL+NVML are very consistent here.

27 The most important observations, in the context of the research goal of this paper, is
 28 the observation of the offsets between the hardware (whole node) power measurements
 29 from Yokogawa and the sum of measurements from the software APIs i.e. RAPL and
 30 NVML. The offsets, considering also differences in idle power values measured using
 31 Equation 1, are presented in all Figures 1-6. We can conclude that for all applications
 32 running on the CPUs, the offset is growing with an increasing number of threads but is
 33 relatively small and reaches only up to approx. 30W for a single CPU using 20 threads
 34 (Figure 3), up to approx. 40W using 32 threads on 2 CPUs on vinnana (Figure 1) and




Fig. 1: Measurements using CPU benchmarks on vinnana

1 40 threads on sanna (Figure 4). These constitute roughly up to 8.7% of the total power
 2 (Yokogawa) for largest numbers of threads in these tests. Those differences include other
 3 system components such as fans, disks and potential inaccuracies of RAPL and NVML
 4 as the measurements from Yokogawa are considered ground truth. We could also observe
 5 that the growth of the offsets is linear up to the number of available physical cores in the
 6 system (20) and then slightly dropping when Hyperthreading is engaged. **The similarity of**
 7 **observed patterns and values for offsets on the two similar systems allows us to conclude**
 8 **that the measurements are valid.**

9 For GPU workload (Xception) as well as mixed workloads (also including Xception
 10 on the GPU(s)) on vinnana, shown in Figure 2, we see linear growth of the offsets up
 11 to approx. 80W reaching the largest percentage of the total power (from Yokogawa) of
 12 approx. 11%. We can see that the offsets for all the applications are very similar.

13 For GPU based workloads as well as CPU+GPU mixed workloads run on sanna, we
 14 see a different behaviour. Firstly, we notice that the workloads run on sanna put much
 15 more load on the same number of GPUs (4) on both systems, which results from different




Fig. 2: Measurements using GPU and mixed CPU+GPU benchmarks on vinnana, number of threads and GPUs used is marked on the X axis




Fig. 3: Measurements using CPU benchmarks on 1 multi-core CPU on sanna

1 workloads and different GPU models – Quadro RTX 6000 with a higher TDP in sanna
 2 compared to Quadro RTX 5000 in vinnana. Additionally, there are 8 GPUs in sanna, in
 3 the same system (motherboard and chassis). This evidently shows in the offsets. We first
 4 notice that for these workloads, both GPU and **mixed** shown in Figure 5 and Figure 6
 5 respectively, RAPL+NVML report average powers growing linearly with an increasing
 6 number of GPUs. Additionally, starting with 4 GPUs, both for GPU and **mixed** work-
 7 loads, for all applications we see a large consistent jump in the offsets. These then either
 8 stay at this level or increase only very slightly for 8 GPUs. These increases are clearly
 9 visible in the Yokogawa readings. The values of the offsets for the largest 8 GPU cases
 10 range up to approx. 19% of the total power consumed by the node, and up to 22-26% for
 11 4 GPUs, as reported by Yokogawa. These ratios are much lower for 2 GPUs, e.g. approx.
 12 6% for **mixed** configurations on sanna. We shall note that while performing tests using the
 13 Xception application for paper [25], we also compared Yokogawa measurements versus
 14 the sum of Intel RAPL and Nvidia NVML and observed the same pattern of power differ-
 15 ence increase for 4-8 GPUs. This suggests that the issue is hardware rather than software




Fig. 4: Measurements using CPU benchmarks on 2 multi-core CPUs on sanna

1 related. We have compared our results with analysis of the requirements of cooling in high
 2 performance computer systems from the literature. Firstly, we note that in the sanna and
 3 vinnana systems, there are two rows of carriers that are four columns across, each carrier
 4 with two fans, for a total of 16 large chassis fans [21].

5 Itoh et al. [18] performed detailed analysis of power consumption of fans in an HP
 6 blade system. They modeled the speed of fan versus the temperature as well as power
 7 consumed versus fan speed (rpm) presenting formula $power = 22.1(rpm/10000)^3 + 8.2$.
 8 They concluded that, in their case, fans consumed approx 0.8~1kW which amounted to
 9 16~20% of the system. Similarly, Kennedy [22] analyzed the 1U and 2U system fan
 10 power consumption as a percentage of total power consumption, across 9 workloads. The
 11 1U fan power consumption was approx. 1% higher and amounted to 15~16% of the total
 12 power. Neudorfer [31] stated that server fans can consume 10% to 15% or more of the
 13 total power drawn by the server. Based on those findings, taking into account the our




Fig. 5: Measurements using GPU benchmarks on sanna




Fig. 6: Measurements using mixed CPU+GPU benchmarks on sanna, number of threads and GPUs used is marked on the X axis

1 offsets include also other system components such as disks, we believe these are in line
 2 with the aforementioned observations.

3 **7. Summary and Future Work**

4 In the paper, we performed detailed comparison of hardware and software based power/
 5 energy measurement methods using the hardware power meter Yokogawa WT-310E as
 6 well as Intel RAPL and Nvidia NVML interfaces respectively. We performed tests using
 7 2 systems, one with 2 Intel Xeon CPUs and 8 Nvidia Quadro RTX 6000 GPUs and the
 8 second 2 Intel Xeon CPUs and 4 Nvidia Quadro RTX 5000 GPUs. For thorough compar-
 9 ison we used selected kernels from NAS Parallel Benchmarks for CPUs and GPUs
 10 and [Xception](#) deep neural network training using several GPUs. Tests were conducted for
 11 CPU, GPU as well as [mixed](#) CPU+GPU configurations, using 1-40 threads for the CPUs
 12 and 1-8 GPUs.

13 We shall note that the software power measurements from Intel RAPL and Nvidia
 14 NVML only capture the CPU, GPU and DRAM components without e.g. cooling and
 15 disks. We have determined that the offset, computed using formula 1 in Section 4, grows
 16 with an increasing load (numbers of CPU threads and GPUs used) and amounts to roughly
 17 8.7% of the total system power for CPU workloads. For GPU and mixed workloads on
 18 the system with 4 GPUs the offset reaches up to 11% of the total system power. At the
 19 extreme end, under the highest 2 CPU + 4-8 GPU load on the system with 8 Quadro RTX
 20 6000, it rises up to roughly 19-26% of the total system power. In the high-load range
 21 (multi-GPU workloads on sanna) obtained offset patterns clearly indicate node cooling
 22 system saturation point.

23 In the future, we plan to extend the scope of tests in terms of applications and systems.
 24 Additionally, we plan to perform a similar comparison with the power measurements ob-
 25 tained from the system platform as well, when provided. We also plan to correlate values
 26 of system metrics corresponding to the loads of CPUs and GPUs with power measurement
 27 offsets. That would allow runtime estimation of whole node power draw using previously
 28 obtained node characteristics.

29 **A. Detailed Results from sanna**

30 For clarity, this appendix section contains detailed results of measurements of selected
 31 benchmarks, including execution times, power measured for individual computing de-
 32 vices and the whole sanna node, as well as standard deviation values. Table 4 contains
 33 measurements for the lu.C benchmark executed in parallel using various numbers of
 34 threads on a single multi-core CPU, Table 5 presents analogous results using 2 multi-core
 35 CPUs, Table 6 details measurements from the ep.D benchmarks using GPUs and finally
 36 Table 7 contains values corresponding to [mixed](#) runs of lu.C and ep.D on both CPUs and
 37 GPUs.

Table 4: Execution times and power draws; server: **sanna**, device: **1 CPU**, implementation: **OMP-CPP**, benchmark: **lu.C**

	1 CPU			
Results from 10 runs	1 Thread	5 Threads	10 Threads	20 Threads
Avg. Exec. time [s]	709.838	146.984	76.672	64.195
Std. dev. of time [s]	1.198	0.124	0.201	0.196
(Yokogawa) Avg. power draw [W]	349.991	379.51	411.501	432.449
(Yokogawa) Std. dev. of avg. power draw [W]	0.085	0.212	0.829	1.553
(CPU: 0) Avg. power draw [W]	34.419	54.615	74.217	82.381
(CPU: 0) Std. dev. of avg. power draw [W]	0.221	0.053	0.079	0.055
(CPU: 1) Avg. power draw [W]	28.919	29.018	28.887	28.803
(CPU: 1) Std. dev. of avg. power draw [W]	0.297	0.075	0.042	0.030
(RAM) Avg. power draw [W]	23.574	27.971	30.029	31.413
(RAM) Std. dev. of avg. power draw [W]	0.093	0.030	0.053	0.046
(GPU: 0) Avg. power draw [W]	5.384	5.414	5.378	5.387
(GPU: 0) Std. dev. of avg. power draw [W]	0.077	0.101	0.073	0.056
(GPU: 1) Avg. power draw [W]	9.691	10.083	10.046	9.759
(GPU: 1) Std. dev. of avg. power draw [W]	0.180	0.139	0.185	0.186
(GPU: 2) Avg. power draw [W]	13.812	13.757	13.811	13.945
(GPU: 2) Std. dev. of avg. power draw [W]	0.164	0.148	0.041	0.190
(GPU: 3) Avg. power draw [W]	17.234	17.248	17.240	17.233
(GPU: 3) Std. dev. of avg. power draw [W]	0.014	0.017	0.012	0.011
(GPU: 4) Avg. power draw [W]	4.045	4.037	4.017	4.025
(GPU: 4) Std. dev. of avg. power draw [W]	0.011	0.017	0.006	0.011
(GPU: 5) Avg. power draw [W]	18.450	18.661	18.709	18.585
(GPU: 5) Std. dev. of avg. power draw [W]	0.111	0.136	0.164	0.117
(GPU: 6) Avg. power draw [W]	6.469	6.393	6.593	6.621
(GPU: 6) Std. dev. of avg. power draw [W]	0.077	0.082	0.108	0.079
(GPU: 7) Avg. power draw [W]	11.371	11.745	12.035	11.495
(GPU: 7) Std. dev. of avg. power draw [W]	0.280	0.249	0.191	0.147

Table 5: Execution times and power draws; server: **sanna**, device: **2 CPUs**, implementation: **OMP-CPP**, benchmark: **lu.C**

	1 CPU			
Results from 10 runs	2 Threads	10 Threads	20 Threads	40 Threads
Avg. Exec. time [s]	518.970	75.024	46.207	44.716
Std. dev. of time [s]	26.472	0.910	0.385	1.309
(Yokogawa) Avg. power draw [W]	354.402	405.909	452.840	466.669
(Yokogawa) Std. dev. of avg. power draw [W]	0.550	0.808	1.639	4.044
(CPU: 0) Avg. power draw [W]	33.168	52.455	67.115	70.499
(CPU: 0) Std. dev. of avg. power draw [W]	0.155	0.204	0.163	0.524
(CPU: 1) Avg. power draw [W]	33.589	51.560	67.221	70.781
(CPU: 1) Std. dev. of avg. power draw [W]	0.307	0.247	0.215	0.605
(RAM) Avg. power draw [W]	23.927	28.316	30.873	31.369
(RAM) Std. dev. of avg. power draw [W]	0.096	0.132	0.117	0.234
(GPU: 0) Avg. power draw [W]	5.559	5.514	5.508	5.538
(GPU: 0) Std. dev. of avg. power draw [W]	0.204	0.114	0.093	0.061
(GPU: 1) Avg. power draw [W]	10.276	10.210	10.480	10.347
(GPU: 1) Std. dev. of avg. power draw [W]	0.298	0.185	0.174	0.064
(GPU: 2) Avg. power draw [W]	13.689	13.966	13.833	13.863
(GPU: 2) Std. dev. of avg. power draw [W]	0.125	0.126	0.166	0.070
(GPU: 3) Avg. power draw [W]	17.229	17.238	17.226	17.256
(GPU: 3) Std. dev. of avg. power draw [W]	0.034	0.017	0.014	0.011
(GPU: 4) Avg. power draw [W]	4.004	4.012	4.039	4.061
(GPU: 4) Std. dev. of avg. power draw [W]	0.043	0.016	0.013	0.023
(GPU: 5) Avg. power draw [W]	18.408	18.451	18.554	18.537
(GPU: 5) Std. dev. of avg. power draw [W]	0.183	0.121	0.097	0.097
(GPU: 6) Avg. power draw [W]	6.278	6.446	6.314	6.491
(GPU: 6) Std. dev. of avg. power draw [W]	0.381	0.087	0.113	0.097
(GPU: 7) Avg. power draw [W]	11.684	12.216	12.250	12.131
(GPU: 7) Std. dev. of avg. power draw [W]	0.367	0.140	0.144	0.212

Table 6: Execution times and power draws; server: **sanna**, device: **GPUs**, implementation: **OMP-CUDA**, benchmark: **ep.D**

	1 CPU			
Results from 10 runs	1 GPU	2 GPUs	4 GPUs	8 GPUs
Avg. Exec. time [s]	28.056	28.212	28.350	28.975
Std. dev. of time [s]	0.080	0.031	0.051	0.133
(Yokogawa) Avg. power draw [W]	539.594	703.358	1306.636	1949.073
(Yokogawa) Std. dev. of avg. power draw [W]	8.704	6.079	15.716	11.492
(CPU: 0) Avg. power draw [W]	33.286	37.151	37.555	45.274
(CPU: 0) Std. dev. of avg. power draw [W]	0.065	0.639	0.223	0.092
(CPU: 1) Avg. power draw [W]	28.970	28.776	37.179	44.509
(CPU: 1) Std. dev. of avg. power draw [W]	0.152	0.231	0.095	0.073
(RAM) Avg. power draw [W]	21.765	21.899	22.030	21.994
(RAM) Std. dev. of avg. power draw [W]	0.077	0.094	0.106	0.074
(GPU: 0) Avg. power draw [W]	156.550	156.842	153.124	150.186
(GPU: 0) Std. dev. of avg. power draw [W]	1.426	0.428	0.299	0.892
(GPU: 1) Avg. power draw [W]	10.206	160.103	158.563	155.258
(GPU: 1) Std. dev. of avg. power draw [W]	0.151	1.263	0.448	0.692
(GPU: 2) Avg. power draw [W]	13.757	13.844	161.103	158.964
(GPU: 2) Std. dev. of avg. power draw [W]	0.138	0.106	1.131	1.452
(GPU: 3) Avg. power draw [W]	17.208	17.205	174.830	172.491
(GPU: 3) Std. dev. of avg. power draw [W]	0.019	0.023	1.426	0.342
(GPU: 4) Avg. power draw [W]	4.055	4.047	4.107	153.803
(GPU: 4) Std. dev. of avg. power draw [W]	0.013	0.011	0.012	1.957
(GPU: 5) Avg. power draw [W]	18.628	18.739	19.708	168.175
(GPU: 5) Std. dev. of avg. power draw [W]	0.145	0.152	0.135	0.994
(GPU: 6) Avg. power draw [W]	5.995	5.979	6.795	158.945
(GPU: 6) Std. dev. of avg. power draw [W]	0.081	0.114	0.152	1.391
(GPU: 7) Avg. power draw [W]	12.572	12.626	13.555	172.951
(GPU: 7) Std. dev. of avg. power draw [W]	0.215	0.201	0.163	1.657

Table 7: Execution times and power draws; server: **sanna**, device: **Hybrid**, implementation: **OMP-CPP+OMP-CUDA**, benchmark: **lu.C+ep.D**

	1 CPU			
Results from 10 runs	1 + 4	2 + 8	4 + 16	8 + 32
Avg. Exec. time [s]	27.974	28.101	28.314	29.211
Std. dev. of time [s]	0.082	0.072	0.076	0.186
(Yokogawa) Avg. power draw [W]	570.861	748.939	1385.889	2024.047
(Yokogawa) Std. dev. of avg. power draw [W]	8.010	9.171	26.194	11.160
(CPU: 0) Avg. power draw [W]	53.715	72.342	65.580	71.738
(CPU: 0) Std. dev. of avg. power draw [W]	0.106	5.246	5.428	0.738
(CPU: 1) Avg. power draw [W]	29.490	29.483	65.091	72.040
(CPU: 1) Std. dev. of avg. power draw [W]	0.108	0.069	7.089	0.824
(RAM) Avg. power draw [W]	27.094	29.173	27.993	29.224
(RAM) Std. dev. of avg. power draw [W]	0.086	2.041	2.760	0.385
(GPU: 0) Avg. power draw [W]	157.522	157.090	153.753	150.166
(GPU: 0) Std. dev. of avg. power draw [W]	1.623	1.619	1.082	0.598
(GPU: 1) Avg. power draw [W]	10.251	159.154	157.782	154.078
(GPU: 1) Std. dev. of avg. power draw [W]	0.181	0.775	1.027	0.878
(GPU: 2) Avg. power draw [W]	14.145	14.103	161.310	158.264
(GPU: 2) Std. dev. of avg. power draw [W]	0.191	0.182	1.294	0.991
(GPU: 3) Avg. power draw [W]	17.251	17.239	173.964	171.584
(GPU: 3) Std. dev. of avg. power draw [W]	0.026	0.031	1.835	1.215
(GPU: 4) Avg. power draw [W]	3.966	3.973	3.996	154.742
(GPU: 4) Std. dev. of avg. power draw [W]	0.005	0.013	0.012	0.731
(GPU: 5) Avg. power draw [W]	18.808	18.940	19.680	167.668
(GPU: 5) Std. dev. of avg. power draw [W]	0.164	0.194	0.344	0.620
(GPU: 6) Avg. power draw [W]	6.049	6.233	6.476	156.762
(GPU: 6) Std. dev. of avg. power draw [W]	0.081	0.346	0.107	0.808
(GPU: 7) Avg. power draw [W]	10.969	11.067	12.323	171.879
(GPU: 7) Std. dev. of avg. power draw [W]	0.157	0.220	0.261	1.756

1 Acknowledgements

2 This work is partially supported by CERCIRAS COST Action CA19135 funded by COST.

3 References

- 4 1. Alqurashi, F.S., Al-Hashimi, M.: An experimental approach to estimation of the energy cost
5 of dynamic branch prediction in an intel high-performance processor. *Computers* 12(7) (2023),
6 <https://www.mdpi.com/2073-431X/12/7/139>
- 7 2. Arafa, Y., ElWazir, A., Elkanishy, A., Aly, Y., Elsayed, A., Badawy, A.H., Chennupati, G.,
8 Eidenbenz, S., Santhi, N.: Nvidia gpgpus instructions energy consumption. In: 2020 IEEE
9 International Symposium on Performance Analysis of Systems and Software (ISPASS). pp.
10 110–112 (2020)
- 11 3. Araujo, G., Griebler, D., Rockenbach, D.A., Danelutto, M., Fernandes, L.G.: Nas parallel
12 benchmarks with cuda and beyond. *Software: Practice and Experience* 53(1), 53–80 (2023),
13 <https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.3056>
- 14 4. Araujo, G.A.d., Griebler, D., Danelutto, M., Fernandes, L.G.: Efficient nas parallel benchmark
15 kernels with cuda. In: 2020 28th Euromicro International Conference on Parallel, Distributed
16 and Network-Based Processing (PDP). pp. 9–16 (2020)
- 17 5. Aslan, B., Yilmazer-Metin, A.: A study on power and energy measurement of nvidia jetson em-
18 bedded gpu using built-in sensor. In: 2022 7th International Conference on Computer Science
19 and Engineering (UBMK). pp. 1–6 (2022)
- 20 6. Bailey, D.H.: *NAS Parallel Benchmarks*, pp. 1254–1259. Springer US, Boston, MA (2011),
21 https://doi.org/10.1007/978-0-387-09766-4_133
- 22 7. Bityutskiy, A., Laakso, A., Correia, H.: Yokotool, <https://github.com/intel/yoko-tool>, accessed
23 on March 5th 2024
- 24 8. Burtscher, M., Zecena, I., Zong, Z.: Measuring gpu power with the k20 built-in sensor. In:
25 *Proceedings of Workshop on General Purpose Processing Using GPUs*. p. 28–36. GPGPU-7,
26 Association for Computing Machinery, New York, NY, USA (2014), [https://doi.org/](https://doi.org/10.1145/2588768.2576783)
27 [10.1145/2588768.2576783](https://doi.org/10.1145/2588768.2576783)
- 28 9. Czarnul, P.: *Parallel Programming for Modern High Performance Computing Systems*. CRC
29 Press, Taylor & Francis (2018), ISBN 9781138305953
- 30 10. Czarnul, P., Proficz, J., Krzywaniak, A.: Energy-aware high-performance computing: Sur-
31 vey of state-of-the-art tools, techniques, and environments. *Sci. Program.* 2019, 8348791:1–
32 8348791:19 (2019), <https://doi.org/10.1155/2019/8348791>
- 33 11. Desrochers, S., Paradis, C., Weaver, V.M.: A validation of dram rapl power measurements.
34 In: *Proceedings of the Second International Symposium on Memory Systems*. p. 455–470.
35 MEMSYS '16, Association for Computing Machinery, New York, NY, USA (2016), [https://doi.org/](https://doi.org/10.1145/2989081.2989088)
36 [10.1145/2989081.2989088](https://doi.org/10.1145/2989081.2989088)
- 37 12. Fahad, M., Shahid, A., Manumachu, R.R., Lastovetsky, A.: A comparative study of methods
38 for measurement of energy of computing. *Energies* 12(11) (2019), [https://www.mdpi.](https://www.mdpi.com/1996-1073/12/11/2204)
39 [com/1996-1073/12/11/2204](https://www.mdpi.com/1996-1073/12/11/2204)
- 40 13. Ferro, M., Yokoyama, A., Klöh, V., Silva, G., Gandra, R., Bragança, R., Bulcão, A., Schulze,
41 B.: Analysis of gpu power consumption using internal sensors. In: *Anais do XVI Work-*
42 *shop em Desempenho de Sistemas Computacionais e de Comunicação*. SBC, Porto Ale-
43 gre, RS, Brasil (2017), [https://sol.sbc.org.br/index.php/wperformance/](https://sol.sbc.org.br/index.php/wperformance/article/view/3360)
44 [article/view/3360](https://sol.sbc.org.br/index.php/wperformance/article/view/3360)
- 45 14. Hähnel, M., Döbel, B., Völp, M., Härtig, H.: Measuring energy consumption for short code
46 paths using rapl. *SIGMETRICS Perform. Eval. Rev.* 40(3), 13–17 (jan 2012), [https://](https://doi.org/10.1145/2425248.2425252)
47 doi.org/10.1145/2425248.2425252

- 1 15. Ikram, M.J., Abulnaja, O.A., Saleh, M.E., Al-Hashimi, M.A.: Measuring power and energy
2 consumption of programs running on kepler gpus. In: 2017 Intl Conf on Advanced Control
3 Circuits Systems (ACCS) Systems & 2017 Intl Conf on New Paradigms in Electronics & In-
4 formation Technology (PEIT). pp. 18–25 (2017)
- 5 16. Ilsche, T.: Energy Measurements of High Performance Computing Systems: From In-
6 strumentation to Analysis. Ph.D. thesis, Technischen Universität Dresden (March 2020),
7 <https://tud.qucosa.de/api/qucosa%3A71600/attachment/ATT-0/>
- 8 17. Intel Corporation: Running average power limit energy reporting / cve-
9 2020-8694 , cve-2020-8695 / intel-sa-00389 (February 2022), [https://www.intel.com/content/www/us/en/developer/articles/
10 technical/software-security-guidance/advisory-guidance/
11 running-average-power-limit-energy-reporting.html](https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/advisory-guidance/running-average-power-limit-energy-reporting.html)
- 12 18. Itoh, S., Kodama, Y., Shimizu, T., Sekiguchi, S., Nakamura, H., Mori, N.: Power consumption
13 and efficiency of cooling in a data center. In: 2010 11th IEEE/ACM International Conference
14 on Grid Computing. pp. 305–312 (2010)
- 15 19. Jay, M., Ostapenco, V., Lefevre, L., Trystram, D., Orgerie, A.C., Fichel, B.: An experimental
16 comparison of software-based power meters: focus on cpu and gpu. In: 2023 IEEE/ACM 23rd
17 International Symposium on Cluster, Cloud and Internet Computing (CCGrid). pp. 106–118
18 (2023)
- 19 20. Kavanagh, R., Djemame, K.: Rapid and accurate energy models through calibration
20 with ipmi and rapl. *Concurrency and Computation: Practice and Experience* 31(13),
21 e5124 (2019), [https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.
22 5124](https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5124), e5124 cpe.5124
- 23 21. Kennedy, P.: Inspur systems nf5468m5 review 4u 8x gpu server (March 2019), [serveTheHome](https://www.servethehome.com/inspur-systems-nf5468m5-review-4u-8x-gpu-server/),
24 <https://www.servethehome.com/inspur-systems-nf5468m5-review-4u-8x-gpu-server/>
- 25 22. Kennedy, P.: Deep dive into lowering server power consumption (February 2022), [serveThe-
26 Home](https://www.servethehome.com/deep-dive-into-lowering-server-power-consumption-intel-inspur-hpe-dell-emc/), [https://www.servethehome.com/deep-dive-into-lowering-server-power-consumption-
27 intel-inspur-hpe-dell-emc/](https://www.servethehome.com/deep-dive-into-lowering-server-power-consumption-intel-inspur-hpe-dell-emc/)
- 28 23. Khan, K.N., Hirki, M., Niemi, T., Nurminen, J.K., Ou, Z.: Rapl in action: Experiences in using
29 rapl for power measurements. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 3(2) (mar
30 2018), <https://doi.org/10.1145/3177754>
- 31 24. Kocot, B., Czarnul, P., Proficz, J.: Energy-aware scheduling for high-performance computing
32 systems: A survey. *Energies* 16(2) (2023), [https://www.mdpi.com/1996-1073/16/
33 2/890](https://www.mdpi.com/1996-1073/16/2/890)
- 34 25. Koszczal, G., Dobrosolski, J., Matuszek, M., Czarnul, P.: Performance and energy aware train-
35 ing of a deep neural network in a multi-gpu environment with power capping. In: Zeinalipour,
36 D., Blanco Heras, D., Pallis, G., Herodotou, H., Trihinas, D., Balouek, D., Diehl, P., Cojean,
37 T., Furlinger, K., Kirkeby, M.H., Nardelli, M., Di Sanzo, P. (eds.) Euro-Par 2023: Parallel
38 Processing Workshops. pp. 5–16. Springer Nature Switzerland, Cham (2024)
- 39 26. Krzywaniak, A., Czarnul, P., Proficz, J.: Depo: A dynamic energy-performance optimizer
40 tool for automatic power capping for energy efficient high-performance computing. *Software:
41 Practice and Experience* 52(12), 2598–2634 (2022), [https://onlinelibrary.wiley.
42 com/doi/abs/10.1002/spe.3139](https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.3139)
- 43 27. Krzywaniak, A., Czarnul, P., Proficz, J.: Dynamic gpu power capping with online perfor-
44 mance tracing for energy efficient gpu computing using depo tool. *Future Generation Com-
45 puter Systems* 145, 396–414 (2023), [https://www.sciencedirect.com/science/
46 article/pii/S0167739X23001267](https://www.sciencedirect.com/science/article/pii/S0167739X23001267)
- 47 28. Lang, J., Runger, G.: High-resolution power profiling of gpu functions using low-resolution
48 measurement. In: Wolf, F., Mohr, B., an Mey, D. (eds.) Euro-Par 2013 Parallel Processing. pp.
49 801–812. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
- 50

- 1 29. Lucas, J., Juurlink, B.: Alupower: Data dependent power consumption in gpus. In: 2016
2 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and
3 Telecommunication Systems (MASCOTS). pp. 95–104 (2016)
- 4 30. Löff, J., Griebler, D., Mencagli, G., Araujo, G., Torquati, M., Danelutto, M., Fernandes, L.G.:
5 The nas parallel benchmarks for evaluating c++ parallel programming frameworks on shared-
6 memory architectures. *Future Generation Computer Systems* 125, 743–757 (2021), <https://www.sciencedirect.com/science/article/pii/S0167739X21002831>
- 7 31. Neudorfer, J.: Optimizing server energy efficiency (February 2009), techTarget,
8 <https://www.techtarget.com/searchdatacenter/tip/Optimizing-server-energy-efficiency>
- 9 32. NVIDIA Corporation: Nvidia management library (nvml), <https://developer.nvidia.com/management-library-nvml>, accessed Sep 2024
- 10 33. Paniego, J.M., Gallo, S., Pi Puig, M., Chichizola, F., De Giusti, L., Ballardini, J.: Analysis of
11 rapl energy prediction accuracy in a matrix multiplication application on shared memory. In:
12 De Giusti, A.E. (ed.) *Computer Science – CACIC 2017*. pp. 37–46. Springer International
13 Publishing, Cham (2018)
- 14 34. Raffin, G., Trystram, D.: Dissecting the software-based measurement of cpu energy consumption: a comparative analysis (2024)
- 15 35. Sen, S., Imam, N., Hsu, C.H.: Quality assessment of gpu power profiling mechanisms. In: 2018
16 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). pp.
17 702–711 (2018)
- 18 36. Shahid, A., Fahad, M., Manumachu, R.R., Lastovetsky, A.: Improving the accuracy of energy
19 predictive models for multicore cpus by combining utilization and performance events model
20 variables. *Journal of Parallel and Distributed Computing* 151, 38–51 (2021), <https://www.sciencedirect.com/science/article/pii/S0743731521000137>
- 21 37. Shalavi, N., Khoshshirat, A., Stellini, M., Zanella, A., Rossi, M.: Accurate calibration of power
22 measurements from internal power sensors on nvidia jetson devices. In: 2023 IEEE International
23 Conference on Edge Computing and Communications (EDGE). pp. 166–170 (2023)
- 24 38. Tröpgen, H., Bielert, M., Ilsche, T.: Evaluating the energy measurements of the ibm power9
25 on-chip controller. In: *Proceedings of the 2023 ACM/SPEC International Conference on Performance Engineering*. p. 67–76. ICPE '23, Association for Computing Machinery, New York, NY, USA (2023), <https://doi.org/10.1145/3578244.3583729>
- 26 39. Vajtas-Anttila, J.M., Koskela, T., Hickey, S.: Power consumption model of a mobile gpu based
27 on rendering complexity. In: 2013 Seventh International Conference on Next Generation Mobile Apps, Services and Technologies. pp. 210–215 (2013)
- 28 40. Yang, Z., Adamek, K., Armour, W.: Part-time power measurements: nvidia-smi's lack of attention (2023)
- 29 41. Yokogawa: WT310E/WT310EH/WT332E/WT333E Digital Power Meter. User's Manual (October 2017), 2nd edition, IM WT310E-01EN, <https://cdn.tmi.yokogawa.com/IMWT310E-01EN.pdf>
- 30 42. Zhang, H., Hoffmann, H.: A quantitative evaluation of the rapl power control system. In: *Proceedings of the 10th International Workshop on Feedback Computing* (2015), <https://api.semanticscholar.org/CorpusID:13950838>
- 31
32
33
34
35
36
37
38
39
40
41
42