
Citation:

Grishina, A and Chinnici, M and De Chiara, D and Rondeau, E and Kor, A (2019) Energy-oriented analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center. Lecture Notes in Electrical Engineering, 574. pp. 286-300. ISSN 1876-1100 DOI: https://doi.org/10.1007/978-3-030-21507-1_41

Link to Leeds Beckett Repository record:

<https://eprints.leedsbeckett.ac.uk/id/eprint/5364/>

Document Version:

Article (Accepted Version)

Creative Commons: Attribution-Noncommercial 4.0

The aim of the Leeds Beckett Repository is to provide open access to our research, as required by funder policies and permitted by publishers and copyright law.

The Leeds Beckett repository holds a wide range of publications, each of which has been checked for copyright and the relevant embargo period has been applied by the Research Services team.

We operate on a standard take-down policy. If you are the author or publisher of an output and you would like it removed from the repository, please [contact us](#) and we will investigate on a case-by-case basis.

Each thesis in the repository has been cleared where necessary by the author for third party copyright. If you would like a thesis to be removed from the repository or believe there is an issue with copyright, please contact us on openaccess@leedsbeckett.ac.uk and we will investigate on a case-by-case basis.

Energy-oriented analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center

Anastasiia Grishina¹, Marta Chinnici², Davide De Chiara³, Eric Rondeau¹,
Ah Lian Kor⁴

¹ Université de Lorraine, CNRS-CRAN, Vandoeuvre-lès-Nancy, F-54000, France
{anastasiia.grishina2, eric.rondeau}@univ-lorraine.fr

² ENEA-ICT Division, C.R Casaccia Via Anguillarese 301, ROMA 00123, Italy
{marta.chinnici}@enea.it

³ ENEA-ICT Division, C.R. Portici Piazzale Enrico Fermi, 1, Portici NA 80055, Italy
{davide.dechiara}@enea.it

⁴ Leeds Beckett University, School of Computing, Leeds, UK
{A.Kor}@leedsbeckett.ac.uk

Abstract. This work analyzes a very subtle kind of energy metrics for Data Centers (DCs), namely productivity metrics which affect the global energy efficiency assessment in DC since they focus on the energy used for processing computing operations. By exploiting the available set of energy consumption data of operating systems in ENEA-DC, HPC-Cluster, the authors evaluated the energy consumed by different queues with several running applications. The queues energy waste has been calculated to provide an assessment for the ineffective use of computation-related energy load within the Cluster. This work shows an increment innovation beyond state-of-the-art for productivity metrics (e.g. useful work), and it will also help provide an invaluable insight into useful energy use and the use of enhanced sustainability metrics with the goal of driving a more sustainable DC. Additionally, sustainability concept in DC operations is driven by estimation of its indirect carbon emissions, which is shown in this work.

Keywords: Data Center, Energy Efficiency, Energy Consumption, Metrics, Cluster, Data Analysis, Workload Management, HPC, Sustainability.

1 Introduction

The critical driver of a sustainable Data Center (DC) is embodied within its energy efficiency strategy. The baseline for the development of this strategy is a structured measurement and control framework that can be used to quantify DC energy efficiency and provide insights into ways of improving it. Currently, a comprehensive list of energy metrics [1, 2, 19, 20] has been presented in existing literature, and progress has been made with energy efficiency metrics that measure computing and non-computing energy used in a DC [17]. This category of parameters, namely

productivity metrics [16, 18], measure useful work within a DC and hence, energy efficiency at various levels of granularity starting from the operating systems to the entire DC in a holistic manner. However, it is difficult to measure the energy consumption at the operating system level. This is due to the fact that energy consumption data disaggregated by DC sub-components may not be often available. It is also challenging to exactly know the number of operating systems and virtual machines running in a DC at a certain point of time. For this reason, measuring the energy consumption at the sub-component level of a DC is an ongoing challenge. Therefore, recently, the authors in [3, 4, 5] propose a novel methodology that addresses problems relating to measurement, calculation, and evaluation of energy productivity in a DC. It entails the collection of data associated with the proportion of useful energy effectively employed by the computing systems and energy wasted during computational work. This is followed by conducting relevant statistical analysis of energy data using bespoke software supported by appropriate mathematical formulation, and interpretation.

In summary, this work aims to investigate productivity metrics based on the analysis of available disaggregated energy consumption data. This set of data is related to the energy consumption of operating systems of HPC Clusters within ENEA DC. It aims to estimate the energy use by different queues where different (regarding the kind - parallel/serial - and services area) sets of applications are running. The following objectives will support this aim:

- a) Analyze the energy consumed by individual applications and to extrapolate the energy consumed at queues level during a period T ;
- b) Calculate the energy used by the queues in terms of useful computing (useful work) and non-useful computing (waste energy) work by employing the Energy Waste Ratio (EWR) metric during the period T ;
- c) Evaluate energy consumption for running parallel versus serial jobs (which comprise different sets of applications) during the period T ;
- d) Assess the load of the energy consumed by a DC Cluster at computing level based on (a-c) during the period T ;
- e) Assess sustainability in the DC: Translate (d) into sustainability metrics - outcomes in terms of energy use are transformed into carbon emissions related metrics (i.e. Carbon Usage Effectiveness (CUE)) to provide insight into the CUE of a real physical DC.

The paper is organized as follows: Section I – Introduction; Section II – Related Work; Section III – Methodology; Section IV – Results and Discussion; Section V – Conclusions and Future Work.

2 Related Work

In recent years, significant research effort and technological development have been devoted to DCs targeting energy efficiency and sustainability-related issues. Undeniably, eco-friendly DCs prioritizes sustainability relating to energy consumption, carbon emissions, and water usage, thus, exerting more sustainability-

driven control over their decisions on growth, geolocation, and outsourcing strategies. The authors in [6] defined *Sustainability* as an e-infrastructure strategy that posits:

- a) Any energy consumption level should be kept as low as possible;
- b) Any resource should be consumed as effectively and efficiently as possible. In other words, wastage should be minimized;
- c) Timely and accurate information should be made accessible for the assessment of energy usage, efficiencies and resource use (wastage) to guide and implement processes or policy improvement;
- d) A complete environmental and social impact of activities should be considered;
- e) The level of IT resource provision should be appropriate to the task being undertaken.

By extending the *Sustainability* definition based on the above five guiding principles to the specifics of sustainable large-scale infrastructure (DCs), it implies that energy consumption ought to be kept at a minimum level as far as possible with the available technology. As mentioned in the Introduction Section, the critical driver of a sustainable DC is embodied within its energy efficiency strategy, which is based on a measurement and control metrics framework. Even though many metrics have been proposed, the debate on a set of globally accepted metrics is still an ongoing challenge, particularly, in the areas of:

- *Productivity metrics*: yet to be explored comprehensively, and there is no existing proposed metric related to a direct measurement of useful work in a DC;
- *Environmental impact*: yet to conduct a complete assessment within a real physical context. Even though there are some carbon and hydro-based metrics such as Carbon Usage Effectiveness (CUE) and Water Usage Effectiveness (WUE), the deployment of these metrics in a real context is yet to be made.

Regarding the productivity metrics (e.g., useful work), the works [3, 4, 5, 7] have shown advance beyond state-of-the-art. On the other hand, the authors in [7] propose a methodology that addresses the problem of measurement, calculation, and evaluation of energy productivity assessment in a DC, which encompasses both the proportion of energy employed for computing and energy wasted during computational processing work. Environmental impact parameters comprise sustainability metrics which are also called “green” metrics. They measure the environmental impact of a DC and its components, highlight the importance of green energy as well as measure the efficiency of recycling and reuse in a DC [2]. One of the most important green metrics, Carbon Usage Effectiveness (CUE), has been proposed by The Green Grid [9] to quantify the annual CO₂ footprint of a DC. For DCs using grid power, the CUE is defined as the ratio of total CO₂ emissions from total energy consumption over the total IT equipment energy:

$$CUE = \frac{\text{Total CO}_2 \text{ Emissions Caused by the Total DC Energy}}{\text{Total IT Equipment Energy}}. \quad (1)$$

The assumption made is that the DC is not a direct producer (i.e. not Greenhouse Gas Protocol Scope 1 emission [21]) but an indirect producer (i.e. Scope 2 emission

[21]) of CO₂ emissions due to consumption of energy provided via the national energy grid. However, the type of external energy source will determine the carbon footprint of 1 kWh consumed by the DC. DC electric energy could be sourced from a broad range of generators: hydro, sun, wind, coal, gas, and nuclear. During an average day, the proportion of each source and hence, the carbon footprint per kWh in the grid could vary. Also, other GHGs such as methane may be produced. Consequently, real-time carbon emission data obtainable from electricity suppliers should be used and this data ought to take into account other GHG emissions that have been converted into *carbon equivalents* (CO₂eq). CUE has the metric, kilograms of CO₂ (KgCO₂eq) per kilowatt-hour (kWh). Its ideal value is 0.0, indicating that no carbon use is associated with the DC operations. However, even if a CUE is the result of the operations within a DC, this formula does not include computing operations and hence, the useful work made within DC.

3 Methodology

This work focuses on energy-oriented analysis based on productive energy consumption in DC [10, 11, 12], particularly, on the evaluation of the energy consumed by queues and energy wasted during job execution. To address this challenge, we analyze the data coming from the experiments conducted in the ENEA HPC DC, on the Cluster named CRESCO4 (hosted by ENEA-R.C. Portici). The main goal is to calculate the energy consumption of applications running on the cluster and estimate energy performance characteristics of queues as well as parallel and serial jobs. Useful and waste energy consumption have been further translated into CO₂ or equivalent greenhouse gases (GHG) emissions to show the environmental impact of the Cluster's work. In this work, improvement has been made concerning the energy measurement and control with respect to the previous studies conducted on available data on CRESCO4 [3, 4, 5]. The authors particularly provide an assessment of the load of the energy consumed by DC Cluster at computing level and also a more in-depth knowledge of the useful work done by the Cluster. The results are also expressed in terms of carbon emissions to improve the understanding of the Carbon Usage Effectiveness (CUE) metric applied in a real physical DC context.

3.1 Facility and Dataset Description

The work focuses on assessment of energy consumption and calculation of energy waste occurring within the CRESCO4 cluster of ENEA DC during the period from February 2017 to January 2018. Briefly, the cluster CRESCO4 consists of 38 Supermicro F617R3-FT chassis, each hosting 8 dual CPU nodes. Each CPU, (specifically an Intel E5-2670) hosts 8 cores, and the total number is 4864 cores. These operate at a clock frequency of 2.6 GHz. The system is provided with a RAM memory of 4 GB per core. Computing nodes access a DDN storage system, for a total storage amount of 1 Pbyte. Computing nodes are interconnected via an Infiniband 4xQDR QLogic/Intel12800-180 switch (432 ports, 40Gbps).

The analysis is based on two available datasets: Platform LSF (Load Sharing Facility) report for one year and Zabbix yearly records of power usage retrieved from installed PDUs. Their intersection has provided information for 11 months, which defines the scope of the current study. The LSF dataset includes information about the number of cores and queue name assigned to every process (process, job, and application are synonyms in this context), start and end time of the application activity, names of executable file and directory, and the marker of whether the process has finished successfully (“done”) or with an error (“exit”). The job is recognized as a serial one if it uses only one core and parallel otherwise. The dataset retrieved from Zabbix tool contains hourly timestamps and values of average power usage observed during every hourly interval. The datasets have covered over 8700 rows in Zabbix and 530668 rows in LSF datasets.

Regarding the scheduling of tasks in the DC Cluster, the jobs are stored in the workload queues where the priority of each job and recorders are based on First Come First Served (FCFS) policy. The cluster processes jobs from 18 different queues, 11 of which are composed of jobs to be performed with parallelization, 3 queues are considered with exclusively serial jobs and the remaining – with both parallel and serial jobs. Approximately 92% of all submitted jobs have been processed in a serial mode, which leaves room for only 8% of jobs being processed with parallelization techniques. Several applications (approximately 40 types) run on the CRESCO4 and cover different fields of research.

The following subsection explains how this information is employed for the energy consumption estimations.

3.2 Data Analysis

To achieve the goal of estimating effective energy consumption and its waste with further extrapolation of results on the queue level, the following system of equations is used for each month:

$$\sum_{i=1}^K \int_{t_{i,j}^0}^{t_{i,j}^1} c_{i,j} \cdot x_j dt = E_j, j = 1, \dots, N. \quad (2)$$

Here x_j is the only unknown set of variables which stands for the power per second and per core, required by any application during the hour interval j . Id of a process is expressed by i with K being the number of active processes for the considered month; $c_{i,j}$ refers to the number of cores required to work on application i during the hour interval j . The process starts at time $t_{i,j}^0$ and ends at $t_{i,j}^1$. The cluster consumes E_j watt-hours of energy during the hour interval j ; and N stands for the number of hours in the extracted month. Thus, for each hour interval j , a separate equation is obtained by a summation of integrals. Every integral represents energy consumption of an application i during the hour interval j . Multiplication of unknown variable x_j and

calculable coefficient $c_{i,j}$ under the integral yields the power consumed by an application. These equations are rewritten as follows:

$$\sum_{i=1}^K c_{i,j} \cdot \frac{t_{i,j}}{3600} \cdot x_j = E_j, j = 1, \dots, N, \quad (3)$$

where, $t_{i,j} = (t_{i,j}^1 - t_{i,j}^0)$ and stands for duration of a process in seconds within the hour interval j . Therefore, N linear equations with one unknown set of variables x_j are formed for each month.

The following steps are undertaken to obtain the results shown in Section 4:

1. Calculate:
 - a. values x_j for each month from Equation (3);
 - b. $(c_{i,j} \cdot x_j)$ and sum these values for all the hour intervals j when the process i is active. For this step, energy consumption of every process within a month is obtained.
2. Group processes by queues using data from LSF, evaluate number of job submissions, energy consumption and apply Energy Waste Ratio (EWR) metric to queues.
3. Group processes by parallel/serial modes of execution and evaluate number of parallel/serial jobs submissions, their energy consumption and EWR.
4. Estimate:
 - a. Power consumption of two categories of jobs for each month: with the use of markers taken from LSF for successfully done jobs producing “useful work” and those ended with a type of an error causing “energy waste”;
 - b. CO₂ emissions for month caused by the cluster “useful work” and “energy waste” using a carbon factor for Italy.

The EWR metric is used as in [13, 14] and expressed as follows:

$$EWR = \frac{\text{Energy Wasted for Not Useful Work}}{\text{Total DC Energy Consumed over Time}} \quad (4)$$

where, the term *Not Useful Work* or “energy waste” refers to the jobs which end with an error for one of the following reasons: a) their running time is too short and represents the work of the scheduler but not of the jobs themselves. Their empirically obtained running time is less than 30 seconds for the CRESCO4 cluster; b) their running time exceeds the queue limit time; in this case the running time is divided into two parts: the “useful work” is done within the queue maximum time and the excess time is regarded to as causing “energy waste”; c) jobs which quit the queue with an error for any other reason. Here, *Total DC Energy Consumed over Time* stands for the energy used for all jobs (both successful and aborted ones).

4 Results and Discussion

The available datasets are used in the previously mentioned steps to analyze the energy consumption of individual applications and further extrapolated to obtain data at the queues level and parallel/serial groups of jobs. EWR metric is applied to assess what percentage of energy has been spent on applications, which has not produced any “useful work”. Moreover, results expressed in terms of energy use are converted into carbon emissions.

4.1 Energy Use by Queues

The queues analysis comprises the estimation of their energy use, EWR, and number of submitted applications. Fig. 1 represents energy consumption levels of different queues over the total period of 11 months. The range of energy load covers values from 207 MWh to 1 kWh. However, EWR metric occurs to be non-proportional to energy use as shown in Fig. 2. EWR of every queue is defined as energy wasted on jobs processed by a queue which end with any type of error divided by the total energy used by the queue in question. Remarkably, one queue resident jobs end only with errors and thus, yielding an EWR value of 100%. This observation as well as of other queues, high EWR values may be used to notify users of repetitive negative jobs statuses and motivate them to check their algorithms to prevent energy waste. The queues analysis in Fig. 3, depicts the number of job submissions over the global period of investigation. Since one job may have been submitted several times, the term “submission” is preferred. The term refers to every submission (even from the same applications) which is counted separately. The most significant number of submissions is reported so that they could be directed to the queue with the second smallest energy use and EWR of 16%, which is twice below the average level for all queues (32.5%). Nevertheless, 99% of energy is consumed by 9.5% of submissions over the observed period. Aggregated values are obtained by grouping consumption and submission values of the first 14 out of 18 queues ordered as in Fig. 1. Essentially, the queue analysis has shown that the scheduling technology may be further improved via user notification in the event of repeated job failures. Typically, when users find their jobs are not running properly, they will terminate and resubmit them. However, this will incur a lot of computer resources and therefore, energy wastage. A good practice is to test and debug programs and ensure they can work properly and produce the expected results. Additionally, it is crucial to understand when it is better or worse to implement a serial or parallel job. When a “parallel” modality is chosen, it is recommended to optimize the algorithms for the use of all available resources.




Fig. 1. Energy consumption (kWh) of queues estimated for the total period of 11 months.




Fig. 2. Energy Waste Ratio (EWR) metric applied to queues, expressed in percent of energy use during the total period of 11 months.




Fig. 3. Number of applications submitted to queues during the total period of 11 months.

4.2 Energy Use by Parallel and Serial Jobs

Jobs submitted to different queues are further partitioned by their mode of execution, namely parallel and serial, with the help of information on core usage available in LSF dataset. Analogous to the queue analysis, parallel and serial jobs groups are investigated in terms of energy consumption, EWR and number of submissions (depicted in Fig. 4, Fig. 5 and Fig. 6 respectively). Consideration of these two categories of execution modes fosters a higher granular analysis of parallel and serial jobs characteristics for every month. Energy consumption and EWR of parallel jobs prevail over serial jobs, while the number of serial jobs submissions is observed to be averagely higher than parallel jobs submissions. An even pattern of parallel jobs EWR is depicted in Fig. 5 with a mean value of 22%, whereas the same metric for serial jobs fluctuates between 0.025% - 4%.

Typically, running parallel jobs can speed up computation or simulation. Resources could also be used more efficiently. However, some jobs submitted to the parallel queues do not actually run in a parallel mode nor make efficient use of the resources. The authors list many factors that cause “no-work” (such as, malfunctioning of the systems or algorithms errors) when dealing with serial and/or parallel jobs. In general, different LSF queues are configured for different kinds of jobs. (Serial queues are for running jobs in serial mode on only a single core. Parallel queues are designed for running jobs in a parallel mode where jobs would run with multithreads on all processor cores). When jobs not optimized for parallelism are submitted to the parallel queues, they will cause big resource wastage. Generally, parallel jobs are poorly coded (programmed) because they have a programming model more suited for serial jobs. For example, when a job is submitted to parallel queues with 24-processor

core and a specification that the job can only run using one processor core, consequently, the other 23 cores will be left idle. This means more than 95% of resources are wasted. The waste will be significant if the job runs for an extended time. Also, when more of such jobs are running in parallel queues, more parallel jobs could be held back in the queue and remain in pending status for an extended time.

A further problem could be related to the jobs that fail to produce results after being dispatched. Some jobs can be dispatched and run, but no results are generated due to problems with the executable, parameters or any options used. Usually, when users find their jobs are not running correctly, they terminate and resubmit them. However, to reiterate, a lot of computer resources and users' time have already been wasted. A good strategy would be to test and debug programs and ensure they can run successfully and produce expected results.

Finally, it will be greatly beneficial if appropriate codes have been written and implemented for parallel jobs to exploit real parallelization computer processing. On the contrary, jobs that are not well parallelized will waste resources. It will take longer to yield results and has adverse effects on other processes. Some factors can be considered to evaluate a proper parallelization process. Generally, the formula considered is CPU Time Consumed divided by Turnaround Time minus Waiting Time in the queue (e.g $[CPU\ Time / (Turnaround\ Time - Waiting\ Time)]$). When the value obtained via the equation approaches the number of cores used, then the parallel algorithm can be considered appropriate.




Fig. 4. Monthly energy use by parallel (grey) and serial (dotted) jobs submitted to all the queues.




Fig. 5. Energy Waste Ratio (EWR) metric applied to parallel (*grey*) and serial (*dotted*) jobs, expressed in percent of overall monthly energy use.




Fig. 6. Number of monthly submissions of parallel (*grey*) and serial (*dotted*) jobs (to all the queues).

4.3 Sustainability

DCs have progressively become subject to green audits which encompasses environmental impact audits and assessment of all their IT and supporting facilities. In the context of the reported ENEA setting, a carbon factor is used to translate energy use of IT facilities (i.e. operating under a specific workload) to carbon (and its equivalent) emissions. Note that the carbon factor in this paper is obtained from [15] for Italy. It converts electrical energy usage in MWh to tonnes of CO₂ or equivalent GHG and is equal to 0.343 tCO₂/MWh.

Fig. 7 depicts the amount of carbon emissions produced by the computation facilities for processing work within the Cluster. As a basis for this figure, the monthly energy consumption is calculated for all the jobs, which have successfully completed, and the jobs, which end up with errors. The values are converted to MWh and then multiplied by the carbon factor. The overall CO₂ emissions fluctuate between 8 and 12.2 tonnes CO₂ per month. The proportion of CO₂ emissions caused by “not useful work” ranges from 16% to 40% of monthly emissions. Fig. 7 shows a low granularity analysis of the approach described in Section 3 and is used here to highlight the importance of identifying jobs, which do not produce any “useful work”, but negatively impact on the energy consumption and environment.




Fig. 7. Monthly CO₂ (or equivalent) emissions caused by jobs which ended with errors (*dotted parts*) and correctly finished jobs (*inclined lines*).

The conducted analysis provides a more in-depth insight into the useful work and at the same time, waste energy in the Cluster. For this reason, an incremental contribution towards a better understanding of DC sustainability issues, using an improved Carbon usage effectiveness (CUE) metric. Undeniably, for this metric, it is

mandatory to consider the energy consumed by computing processes and not only the energy consumed for IT equipment usage. As a result of this, this unravels a novel and innovative perspective for DC sustainability metrics and greenness related policies.

5 Conclusion and Future Work

Measuring energy at various levels in a DC is a technique, which can enforce and validate the merits and advantages of any green, sustainable or even a simple energy reduction policy, delivered through many sustainable strategies. In this paper, we have investigated a specific class of energy metrics (i.e. productivity and waste energy metrics) for DCs. This paper has presented the assessment results of ENEA-HPC-CLUSTER based on energy-oriented analysis related to productive energy consumption and waste energy.

As shown in this current work, measurements at different granularity levels (i.e. job or queue), and execution modes (i.e. parallel or serial) provide invaluable statistical data. Rigorous analysis of such collated data highlights areas for improvement and could further inform appropriate sustainable operations and process related policies. Additionally, this research has demonstrated the need for super optimized codes in order to effect a higher level of energy efficiency. Also, carbon emissions estimation (using CUE) provides an insight into the environmental impact of working IT facilities and this poses as a step forward for more environmentally related metrics. Finally, the findings in this research foster a better understanding and enhancement of the existing pool of DC sustainability- related metrics. This research provides a new perspective for green DC sustainability metrics and policies. Further research could be a rigorous analysis of the bottom-up cascading effect of energy savings and carbon (or greenhouse gases) emissions reduction at processor level to the DC and Grid level. It is reported by Gartner (Datacenter Knowledge, 2017) that Google has 2.5 millions servers in 2016. We could further create a *what if* scenario which scales ENEA datacenter up to Google's level. For instance, the cluster for our research has 608 servers with 0.343 tCO₂/MWh with greenhouse gas emissions. If the number of servers is 2.5 millions then the total emissions will be 1410 tCO₂/MWh. According to IEA (2017), the total emissions for the year 2016 is 32.5 GtCO₂. This means the percentage is almost 4.33×10^{-7} . The point raised is *what if* we consider the emissions of waste energy for all the other datacenters in the world. Undeniably, this would have a great impact of waste energy on the environment in a *global* scale.

For future work, there is a need to extend the energy consumption of CPUs to different components of a computer system (e.g. memory, disk, network, etc...). A more holistic energy-assessment ought to be conducted. This will take into consideration the implication of cooling (e.g. fan speed, etc...) and relevant quality of thermal environment (including incoming and outgoing air temperature of each computing node) on DC energy consumption as well as communication units' impact. Additionally, direct measures of DC carbon footprint and sustainability metrics will provide a more comprehensive energy/sustainable DC assessment as well as we will take a time of a day into consideration to foster new strategies for jobs execution within the cluster. There is a need to invest in instrumentation for accurate measurement of parameters that will be relevant for our research on performance and energy efficiency metrics for a DC. Finally, it is necessary to consider the deployment of machine learning algorithms for an AI-defined datacenter. This would involve the analysis of the impact of individual vectors and their complex interactions on a datacenter energy efficiency.

Acknowledgments. The research work has been supported and funded by the PERCCOM Erasmus Mundus Program of the European Union [8]. Moreover, the authors would like to express their gratitude to the research HPC group at the ENEA-R.C. Portici for the useful advice on modelling and control of ENEA-Data Center.

References

1. M. Chinnici, A. Capozzoli, and G. Serale, "Measuring energy efficiency in data centers", in *Pervasive Computing Next Generation Platforms for Intelligent Data Collection*, Chapter 10, pp. 299-351 (2016)
2. V. D. Reddy, et al., "Metrics for Sustainable Data Centers", *IEEE Transaction on sustainable computing*, vol. 2, No. 3 (2017)
3. M. Chinnici, D. De Chiara, and A. Quintiliani, "Data Center, a Cyber-Physical System: Improving Energy Efficiency Through the Power Management", in *Proceedings of 2017 IEEE 15th International Conference on Dependable, Autonomic and Secure Computing, 2017 IEEE 15th International Conference on Pervasive Intelligence and Computing, 2017 IEEE 3rd International Conference on Big Data Intelligence and Computing and 2017 IEEE Cyber Science and Technology Congress, DASC-PICom-DataCom-CyberSciTec 2017*, vol. 2018-January, pp. 269-272 (29 March 2018)
4. A. Quintiliani, M. Chinnici, and D. De Chiara, "Understanding 'workload-related' metrics for energy efficiency in Data Center", in *20th International Conference on System Theory, Control and Computing, ICSTCC*, pp. 830-837 (2016)
5. M. Chinnici, D. De Chiara, and A. Quintiliani, "An HPC-Data Center Case Study on the Power Consumption of Workload", in *Applied Physics, System Science and Computers II*, Springer, APSAC 2017, Lecture Notes in Electrical Engineering, vol. 489, pp. 183-192 (2017)
6. C. Pattinson et al., "Green Sustainable Data Centres, Measurement and Control", online course materials, e-Infranet, <http://e-infranet.eu/>.

7. A. Grishina et al, "DC Energy Data Measurement and Analysis for Productivity and Waste Energy Assessment", *accepted* paper in CSE 2018: IEEE International Conference on Computational Science and Engineering, Bucharest (Romania) (October 2018)
8. A. Klimova, E. Rondeau, K. Andersson, J. Porras, A. Rybin, A. Zaslavsky, "An international Master's program in green ICT as a contribution to sustainable development", *Journal of Cleaner Production*, vol. 135, pp. 223-239 (2016)
9. C. Belady, D. Azevedo, M. Patterson, J. Pouchet, and R. Tipley, "Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric", The Green Grid, White Paper No.32, 2010, <https://www.thegreengrid.org/en/resources/library-and-tools/241-Carbon-Usage-Effectiveness-%28CUE%29%3A-A-Green-Grid-Data-Center-Sustainability-Metric>.
10. M. Chinnici, A. Quintiliani, "An example of methodology to assess energy efficiency in Data Centers", in IEEE Third Int. Conf. Cloud Karlsruhe, pp. 459-463 (2013)
11. A. Capozzoli, M. Chinnici, M. Perino, and G. Serale, "Review on Performance Metrics for Energy Efficiency in Data Center: The Role of Thermal Management", in Springer E2DC 2014, Cambridge, vol. 8945, pp. 135-151 (2015)
12. A. Capozzoli, G. Serale, L. Liuzzo, and M. Chinnici, "Thermal Metrics for Data Centers: A Critical Review", in ENERGY PROCEDIA Elsevier, vol. 62, pp. 391-400 (2014)
13. L. Cupertino et al., "Energy-efficiency, thermal aware modeling and simulation of data centres: the CoolEmAll approach and evaluation results", *Ad Hoc Netw.*, vol. 25, pp. 535-553 (2015)
14. I. Munteanu, V. Debusschere, S. Bergeon, and S. Bacha, "Efficiency metrics for qualification of datacenters in terms of useful workload", in IEEE Grenoble Conference PowerTech, POWERTECH 2013 (2013)
15. B. Koffi et al., "Covenant of Mayors for Climate and Energy: Default emission factors for local emission inventories– Version 2017", EUR 28718 EN, JRC107518, Publications Office of the European Union, Luxembourg (2017)
16. The Green Grid, "A Framework for Data Center Energy Productivity", White Paper #13 (2008), <https://www.greenbiz.com/sites/default/files/document/GreenGrid-Framework-Data-Center-Energy-Productivity.pdf>.
17. Horner, N., and Azevedo, I., "Power usage effectiveness in data centers: Overloaded and Underachieving," *The Electricity Journal*, Vol. 29, pp. 61–69, (2016)
18. The Green Grid, *The Green Grid Data Center Power Efficiency Metrics: PUE And DCiE*, 2007. [Online]. Available: Premier Solutions Co., http://www.premiersolutionsco.com/wp-content/uploads/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf.
19. BCS, "Data centre energy efficiency metrics: Existing and proposed metrics to provide effective understanding and reporting of data centre energy," (nd.a), <https://bcs.org/upload/pdf/data-centre-energy.pdf>
20. The Green Grid, "The Green Grid Data Center Power Efficiency Metrics: PUE And DCiE," (2007), http://www.premiersolutionsco.com/wp-content/uploads/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf
21. The Greenhous Gas Protocol, <https://ghgprotocol.org/sites/default/files/standards/ghg-protocol-revised.pdf>
22. Datacenter Knowledge, (2017) url: <https://www.datacenterknowledge.com/archives/2017/03/16/google-data-center-faq>, accessed date: 28/9/2018
23. IEA. (2017). *Global Energy & CO₂ Status Report 2017*, url: <https://www.iea.org/publications/freepublications/publication/GECO2017.pdf>, accessed date: 28/9/2018