LUT University
School of Engineering Science
Erasmus Mundus Master's Program in Pervasive Computing & Communications for Sustainable Development (PERCCOM)

Anastasiia Grishina

DATA CENTER ENERGY EFFICIENCY ASSESSMENT BASED ON REAL DATA ANALYSIS

Supervisors: Doctor Marta Chinnici (ENEA Casaccia Research Center)
Doctor Ah-Lian Kor (Leeds Beckett University)
Professor Eric Rondeau (University of Lorraine)
Professor Jean-Philippe Georges (University of Lorraine)

Examiners: Prof. Eric Rondeau (University of Lorraine)
Prof. Jari Porras (LUT University)
Prof. Karl Andersson (Luleå University of Technology)

This thesis is prepared as part of an European Erasmus Mundus Programme PERCCOM - PERvasive Computing & COMmunications for sustainable development.

Co-funded by the Erasmus+ Programme of the European Union

This thesis has been accepted by partner institutions of the consortium (cf. UDL-DAJ, n°1524, 2012 PERCCOM agreement).

Successful defense of this thesis is obligatory for graduation with the following national diplomas:

- Master in Complex Systems Engineering (University of Lorraine)
- Master of Science in Technology (LUT University)
- Master of Science in Computer Science and Engineering, specialization in Pervasive Computing and Communications for Sustainable Development (Luleå University of Technology)

ABSTRACT

Lappeenranta University of Technology
School of Engineering Science
Erasmus Mundus Master's Program in Pervasive Computing & Communications for Sustainable Development (PERCCOM)

Anastasiia Grishina

Data Center Energy Efficiency Assessment Based on Real Data Analysis

Master's Thesis

105 pages, 13 figures, 10 tables, 4 appendices

Examiners: Professor Eric Rondeau (University of Lorraine)
Professor Jari Porras (Lappeenranta University of Technology)
Assoc. Professor Karl Andersson (Luleå University of Technology)

Keywords: Data Center, Energy Efficiency, Smart City, Productivity Metrics, Thermal Metrics, Thermal Management, Machine Learning, Real Data Analysis

This work covers energy efficiency analysis of Data Center (DC) operations. DCs empower a wide variety of applications and enhance decision making processes. Having such a crucial role in the modern life, DCs remain large power consumers due to their IT and cooling systems' demand for electricity. Since sustainability has become one of the main global goals, DCs should incorporate eco-friendly strategies to continue their operations without violating sustainability requirements. For a DC, sustainable goals could be interpreted as pursuing energy efficiency of all the operations. Therefore, energy efficiency has been addressed in this work from the point of IT equipment energy productivity and thermal characteristics of an IT room. Mathematical modelling, statistical analysis, productivity and thermal metrics evaluation and a Machine Learning (ML) technique applied to monitoring data collected in a real DC have resulted in a set of recommendations for DC energy efficiency improvement.

PUBLICATIONS

A. Grishina, M. Chinnici, D. De Chiara, G. Guarnieri, A. Kor, E. Rondeau, and J.-P. Georges, “DC Energy Data Measurement and Analysis for Productivity and Waste Energy Assessment”, in *2018 Proceedings - 21st IEEE International Conference on Computational Science and Engineering, CSE 2018*, pp.1-11. IEEE. ISBN 978-1-5386-7649-3 DOI: https://doi.org/10.1109/CSE.2018.00008.
[Online]: https://ieeexplore.ieee.org/document/8588212

M. Chinnici, A.-L. Kor, E. Rondeau, and A. Grishina, “Sustainable data center in smart cities: the role of sustainability-related metrics,” Abstract in *Proceedings of the First Sustainable Solutions for Growth SSG 2018*, 2018, p. 35. [Online]: http://ssg.budzianowski.eu/wp-content/uploads/sites/9/2017/11/Proceedings-1st-SSG-2018.pdf

A. Grishina, M. Chinnici, D. De Chiara, E. Rondeau, and A.-L. Kor, “Energy-oriented analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center,” in *Lecture Notes in Electrical Engineering, 3rd International Conference on: Applied Physics, System Science and Computers*, Dubrovnik, Croatia, 2018. Ed. Springer, 2019. [Online]: https://link.springer.com/chapter/10.1007/978-3-030-21507-1_41

A. Grishina, M. Chinnici, A.-L. Kor, E. Rondeau, J.-P. Georges, and D. De Chiara, “Data Center for Smart Cities: Energy and Sustainability Issue,” Chapter in *Big Data Platforms and Applications - Case Studies, Methods, Techniques, and Performance Evaluation (In-Press)*, F. Pop, Ed. Springer, 2019. **In-Press**

ACKNOWLEDGEMENT

Firstly, I would like to express my sincere gratitude for my main supervisor Dr. Marta Chinnici at ENEA R.C. Casaccia for her helpful and not limiting expert advice, for her motivation to articulate our findings in publications as well as encouragement towards excellence. Special thanks to her for being my best guide to Italy and a soulmate mathematician. My thanks would never be enough for such a productive and pleasant time that I had during my last semester with her.

I am grateful for Prof. Ah-Lian at Leeds Beckett University for fruitful discussions and immense contribution throughout all the research process as well as for extending the journey of my thesis studies to a month of Machine Learning training in Leeds.

With special mention to Dr. Davide de Chiara who has been a supportive research collaborator and monitoring data provider on the CRESCO4 and CRESCO6 technical side of the ENEA R.C. Portici data center.

A very special gratitude goes out to all my supervisors and collaborators Prof. Eric Rondeau and Prof. Jean-Philippe Georges at University of Lorraine, Dr. Guido Guarnieri and Dr. Fiorenzo Ambrosino at ENEA R.C. Portici who have contributed with fresh ideas for a research direction and technical explanations whenever requested.

I would like to thank the European Commission and the PERCCOM consortium for supporting this research [1]. Special gratitude to Prof. Eric Rondeau, Prof. Jari Porras and Prof. Karl Andersson, the program coordinators and wonderful Professors. To professors from ITMO University for the unforgettable summer school in Saint-Petersburg.

To my outstanding PERCCOM classmates for sharing the hard time of studying and the joy of travelling, for all the lifelong memories which I sincerely hope would grow during the following years. I am grateful to the destiny for having met you.

To my beloved family and dearest friends back home whose support and warm wishes have reached me regardless of the distance between us. You strengthen my belief in love.

Anastasiia Grishina

3 June 2019

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF SYMBOLS AND ABBREVIATIONS

ASHRAE	The American Society of Heating, Refrigerating and Air-Conditioning Engineers
BAL	Balance
BP	Bypass
CI	Confidence Interval
CoC	Code of Conduct
CRAC	Computer Room Air Conditioning
CRAH	Computing Room Air Handler
CRESCO	Centro computazionale di RicErca sui Sistemi COmplessi
CUE	Carbon Usage Effectiveness
CWR	Carbon Waste Ratio
DB	Database
DC	Data Center
DCeP	Data Center Energy Productivity
DVFS	Dynamic Voltage and Frequency Scaling
EE	Energy Efficiency
EWR	Energy Waste Ratio
FCFS	First Come, First Served
GHG	Greenhouse Gas
HPC	High Performance Computing
IaaS	Infrastructure as a Service
IoT	Internet of Things
ITE	Information Technology Equipment
ML	Machine Learning
PaaS	Platform as a Service
PDU	Power Distribution Unit
PUE	Power Usage Effectiveness
QoS	Quality of Service
R	Recirculation
RCI	Rack Cooling Index

RHI	Return Heat Index
RTI	Return Temperature Index
SaaS	Software as a Service
SC	Smart City
SHI	Supply Heat Index
UPS	Uninterruptible Power Supply
WCSS	Within Cluster Sum of Squares

1 INTRODUCTION

An estimation made by United Nations states that 66% of the world's population will live in cities by 2050 [2]. To address current and foreseen environmental and social challenges, cities tend to exploit Information and Communication Technologies (ICT). This helps optimise urban management and marks a process of their phasing into smart cities [3], [4]. ICT involvement fosters numerous applications to emerge in the cities and contributes to the quality of everyday life through enhancement of transportation, facilitating medical and governmental services as well as leveraging e-commerce and other spheres [5], [6]. For their effective work, applications require collection and processing of massive quantities of data (i.e. Big Data) related to urban living from objects (e.g., IoT), systems (e.g., energy infrastructure) and society (e.g., city residents as applications users). These diverse big data create useful content for various stakeholders, including citizens, visitors, the local government, and companies. In this scenario, the Data Centers (DCs) play a fundamental role, since they satisfy the demand to process a vast amount of urban big data which comes from interconnected systems operating in the cities. DCs, as High-Performance Computing (HPC) facilities that process urban applications, could be used to foster smart city sustainability providing computational resources for smart technologies. However, these processing demands have led to a tremendous increase in energy consumption, and undeniably, electricity usage contributes to the highest portion of expenditure in DCs [7]. High energy consumption leads to extensive use of energy resources and affects the environment by indirect carbon emissions as well as resources exhaustion. This implies that DC sustainability and in particular, its energy efficiency are crucial goals to be achieved by current emerging computational technologies.

1.1 Background

The context of this work is defined by three partially intersecting notions of a smart city, a data center and sustainability as shown in Fig. 1. To approach the central intersection of all three notions depicted in Fig. 1 which is a focus of this work, the following paired review is conducted: a data center and sustainability, a smart city and a data center, a smart city and sustainability, and finally, a sustainable data center in a (sustainable) smart city.

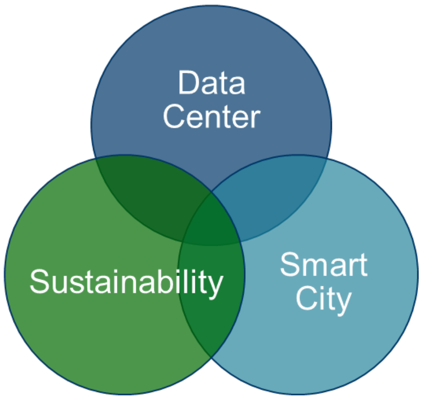

Figure 1. Intersection of smart city, data center and sustainability represents a contextual ground for this work

Data Center and Sustainability

Sustainability of a DC is most frequently regarded to as its energy efficiency and adoption of best practices for optimal DC infrastructure management [8]. In the context of sustainable DC operations, energy efficiency comprises cooling and IT equipment utilisation optimised to maintain recommendable IT room conditions and to satisfy service level agreements with minimal energy consumption. Moreover, sustainable DC practices include integration of renewable energy as a resource produced with minimal carbon emissions, heat recovery as mentioned before as well as regular evaluation of DC productivity and sustainability indices with a set of pre-defined metrics [9].

Pursuing DC sustainability is a challenging task due to a large number of factors affecting DC productivity and energy efficiency. For example, a trade-off between colder locations for the free air-cooling and sunny places for solar power plants is an issue yet to be analysed [10]. Another challenge concerns thermal equipment: raising the setpoint of cooling equipment or lowering the speed of CRAC (Computer Room Air Conditioning) fans to save energy used by thermal equipment may in the long-term decrease the IT systems' reliability, thus, a balance is yet to be found [10], [11]. Furthermore, an ongoing challenge of power overprovisioning and causing energy waste for idle servers has brought about research works on energy storage in UPS (Uninterruptible Power Supply), optimal allocation of PDUs (Power Distribution Units) with respect to servers, and multi-step algorithms for power monitoring and on-demand provisioning reviewed in [10]. Other challenges encompass workload management, network-level issues as optimal routing, VM allocation, balance between power savings and network QoS (Quality of Service) parameters as well as choice of appropriate metrics for DC evaluation.

One standard metric used by a majority of industrial DCs is Power Usage Effectiveness (PUE) proposed by Green Grid Consortium [9]. It shows the ratio of total DC energy utilisation with respect to the energy consumed solely by IT equipment. A plethora of metrics currently under development evaluates thermal characteristics, a ratio of renewable energy use, energy productivity of various components and other parameters. DCs experience an urgent need for a holistic framework that would thoroughly characterise them with a fixed set of metrics and find potential pitfalls in their operation. Although such attempts have been found in the research work, no framework has been standardised so far [12]–[15].

Smart City and Data Center

ICT is assumed to be a characterising attribute of a smart city as technologies help decision-makers optimise urban management and automate it [16], [17]. The very concept of a smart city stems from definitions of information cities, digital cities, intelligent cities, and only since 2010 a smart city notion have appeared more frequently in literature than its predecessors. This development of the notion of a digitally enhanced city emphasizes the important role of ICT in smart cities [18].

DCs with their large consolidated computing power enable storage and processing of big data coming from interconnected urban systems and residents. They can provide elastic on-demand virtualised resources for smart city computational needs [19]. Indeed, exponential demand for big data processing creates a need for scalable Data Analytics applications. Moreover, tremendous growth of data is predicted to reach 35 trillion gigabytes by 2020 [20]. To cope with such amounts of data, DCs provide Infrastructure, Platform and Software as a Service (IaaS, PaaS, SaaS respectively) for developers to implement new smart city solutions that help businesses and governmental organisations in decision making process [16], [19], [20].

Alternatively, some research work develops a concept of DCs as individual smart city players. For example, exploiting flexibility of energy consumption by IT equipment of DC for delay-tolerant workload provides a DC with a potential to play an active role in smart city power grids [21]. In this scenario, a DC with a local renewable power generation plant should use advanced task scheduling. It is supposed to reshape the load minimising the power purchased from the grid and maximising the generated power offered to the grid. Additionally, heat recovery technologies could be utilised to supplement existing heating solutions in a city [22], [23].

In essence, a DC provides computing resources to smart city stakeholders and could be regarded to as enabler of urban big data applications as well as power grid player and supplementary source of heat. Meanwhile, if DC energy consumption is not optimised, it contributes to indirect carbon-related emissions. It also needs frequent retrofits because of hardware exhaustion and violates sustainable environmental requirements of a smart city.

Smart City and Sustainability

Population growth and high urbanisation rate create a number of social, environmental, technological and other challenges for cities. Complex by their nature, cities comprise advanced systems that provide transportation, governmental and medical services, places of living and leisure. These systems might be undermined as many cities have not been created to support current or future estimated number of residents [3]. Therefore, social, economic and environmental sustainability are key factors that would ensure cities' steady operation under the circumstances of growing urban population [18].

Initiatives of the cities to pursue various sustainability goals differ in their nature: some cities solely invest in technologies while other cities rely on future human capital, foster innovation and entrepreneurship. These initiatives may comprise energy, water and waste management, transportation and medical services enhancement, e-government and other improvements in different aspects of city life. A city that incorporates one or more of advanced technological solutions could be called a *smart city*. Degrees of "smartness" of a city may be measured by a number of technological or other initiatives developed in the city and their integration into the city infrastructure [24], [25]. However, there is lack of consensus about the definition of a smart city and ambiguities still persist [3], [17], [18].

Similar to the definition of a smart city, *sustainability* is a term that is widely discussed and interpreted in various ways. A definition suitable in the scope of this work has been proposed in [17]. It emphasizes the need for a balance amongst "economic development and prosperity with environmental protection and integrity and social equity and justice". Within the context of a smart city it would imply adequate growth of a city and development of urban applications as well as improvement of quality of life while minimising its influence on the environment and combatting social inequality.

Sustainability goals and requirements tend to become essential parts of smart cities' development, although some research work still argues that a smart city and a sustainable city are not interchangeable notions [17]. According to [17], smart cities prioritise modern technologies and efficient solutions for everyday life, while sustainable cities focus on sustainability goals and design in the first place. Smart cities need to incorporate smart solutions under environmentally friendly and sustainable frameworks, and sustainable cities should exploit advanced technological solutions for their goals. In this way stronger connection between smart and sustainable concepts could be achieved.

To clarify, as an example of discrepancies between smart and sustainable cities, a waste management system could be considered. In a hyperbolised scenario of a smart but not sustainable city, waste collection trucks might have optimised routes and empty garbage bins on time to effectively avoid street pollution, but the litter is solely disposed in dumps where it decomposes for years and has negative environmental effect. If a sustainable but not smart city scenario is considered, there might be opportunities for waste recycling, but the waste collection system is not organised well, so waste recycling plants do not contribute to the city's cleanliness.

A recent tendency to include sustainability as one of the necessary requirements for a smart city has fortunately narrowed the gap between smart and sustainable cities [18]. Cities use IoT for sensing air quality, e-health for providing accessible medicine to patients, smart and sustainable waste management, develop renewable energy and smart grids, which tends to improve cities from both technological and environmental aspects [5], [26]–[28].

Sustainable Data Center for a Smart (and Sustainable) City

To consolidate the three notions based on descriptions above, sustainability of a smart city could be fostered by ICT, including DCs that process big data coming from urban applications. Smart city applications should be designed in a way that they aim for a balance between high quality of life and resource utilisation, not undermining environmental and social sustainability goals. For a DC as a smart city actor, a critical driver of sustainability is embodied within its energy efficiency strategy. This strategy is based on a structured measurement and control framework that could evaluate DC energy efficiency and provide insights into ways of its improvement. Since the thermal and IT equipment are the major energy consumers within a DC, it should be the primary focus of the energy efficiency framework. Finally, if sustainability requirements are met by smart city with the help of ICT (and DCs, in particular), eco-friendly policies become essential for DCs to follow.

1.2 Motivation

Mankind approaches climate change with various environmental targets as well as estimation of global energy consumption and carbon emissions caused by industrial activities [29], [30]. According to Gartner (2007), ICT accounts for 2% of global carbon emissions with 23% DC share in total ICT emissions. The DC electricity consumption increased twice from 2000 to 2005 and slowed from 2005 to 2010 partially due to the economic crisis of 2008-2009 and energy efficiency orientation starting from 2005. Total electricity use by DCs counted for 1.3% in 2010 [31]. DC carbon emissions are predicted to grow at 7% rate and reach 0.29 and 0.36 $GtCO_2$ by 2020 and 2030 respectively [29], [30]. Accounting for continuous growth of ICT electricity consumption and its environmental impact, energy efficiency strategy, as a part of EU 2020 and 2030 energy climate targets, comprises important measures to mitigate carbon emissions, improve the security of energy supplies and the business competitiveness compared to “business as usual” [32].

Several studies have investigated the use of metrics for DC assessment and identified the relevant set of parameters to assess the energy consumption and evaluate the benefits of energy and sustainability strategies [13], [33]. Additionally, some improvement is proposed by authors [34] in terms of a more comprehensive metrics framework and, above all, parameters for direct evaluation of energy used for productive computing operations, or *useful work*, in a DC [35]–[37]. The concepts of energy efficiency and sustainability represent future challenges in smart cities that depend on urban applications empowered by DCs, and ICT in general. In the meantime, complex issues in DCs from the design to utilisation stages should be addressed.

1.3 Problem Definition

Despite the emergence of studies and analysis in the corresponding fields, understanding the energy efficiency and sustainability concerns of DCs as well as their environmental assessment remain limited in practice [17]. Specifically, the following challenges persist:

1. A common regulatory framework encompassing explanatory metrics and methodologies for DC sustainability assessment is still unavailable [38], [39].

2. Due to its ease of use, a current standard industrial metric for measuring DC energy efficiency is de facto PUE. However, it does not fully reveal the real energy performance of DCs, e.g. IT equipment efficiency [40], [41]. Specifically, *energy waste* generated by inefficient use of computing resources is not widely investigated.

3. Limited attention has been devoted to evaluation of IT room thermal characteristics in real DCs. Although some frameworks are suggested in this area by the research work [38], [42], case studies are still infrequent.

4. Airflow efficiency of a DC is most commonly modelled with Computational Fluid Dynamics (CFD), a fluid mechanics approach. Systems that realise this approach in practice frequently have high computing resources and memory requirements which makes repeated evaluation of DC efficiency expensive from sustainability point of view [42]. While these models are beneficial for theoretical investigation, practical real-time analysis could be facilitated using other less resource-consuming approaches.

Identified research gaps are addressed in this work through data analysis of real DC power utilisation and thermal characteristics.

1.4 Research Aims and Objectives

General aim

Motivated by the mutual dependency between DC energy consumption and sustainable requirements for "smartness" of modern technologies and cities, the aim is to explore different facets of DC energy efficiency: computing systems energy productivity and thermal management.

To achieve this aim, we divide the work into three phases. One of these phases is covered in detail in this thesis while the detailed discussion the other two phases are found in appendices (Appendix 2 and Appendix 3 and will be submitted to scientific journals.

Phase 1. Energy Efficiency Analysis of IT Processes

Aim

The aim is to improve computing processes energy efficiency assessment methods through the investigation of productive energy consumption of ENEA Portici CRESCO4 cluster IT equipment using dataset 1 (see Appendix 1). In this phase, the following research objectives are addressed.

Research objectives

RO1.1. Evaluate energy utilisation by productive computing processes and energy waste within a DC cluster through the employment of appropriate metrics.

RO1.2. Propose metrics for the evaluation of carbon emissions associated with energy waste caused by premature abortion of computational jobs to improve the DC sustainability.

RO1.3. Provide recommendations for the improvement of IT-related energy productivity within the computing cluster under consideration.

Phase 2. Analysis of Data Center Thermal Characteristics (see Appendix 3 Phase 2 for methodology, results and discussion)

Aim

The aim is to increase DC thermal awareness and provide recommendations for effective thermal management based on the study of thermal characteristics of the DC IT room environment and IT equipment energy consumption of ENEA Portici CRESCO6 cluster using dataset 2. Phase 2 targets the statistical analysis of IT room thermal characteristics and thermal metrics evaluation. To achieve this aim, the following research objectives are addressed.

Research objectives

RO2.1.1. Investigate on typical temperature ranges within a cluster IT room.

RO2.1.2. Apply macro (room-level) and micro (node-level) thermal metrics as well as statistical methods to reveal possible existence of cooling system design pitfalls (e.g. hotspots, bypass, recirculation).

RO2.1.3. Formulate recommendations to improve thermal management in the IT room of the cluster in consideration.

Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis

Aim

The aim is to identify individual servers that frequently occur in the hotspot zones by applying a clustering algorithm to available dataset 2 with thermal characteristics of ENEA Portici CRESCO6 computing cluster. The following research objectives will facilitate the achievement of this aim.

Research objectives

RO3.1. Apply an appropriate clustering algorithm to a chosen subset of available data concerning IT room thermal characteristics to determine servers (with IDs) rate of incidence in the following categories: high, moderate, or low incidence in hot, moderate or cold zones within the cluster.

RO3.2. Provide a list of recommendations for thermal design to address the issue of local hotspots.

1.5 Delimitations

This work aims to create a methodology for holistic evaluation of DC characteristics and enhancement of its operation. However, the goal is not to create any automated measurement and evaluation system, but rather to provide a proof of concept how energy consumption, thermal characteristics and environmental effects could be estimated based on raw data from the monitoring system, what problems have to be addressed during data analysis and what assumptions are suitable for a similar DC case study.

Available datasets that are composed of measurements of real DC facilities characteristics provides unexhaustive ground for DC metrics evaluation and assessment of its characteristics. Methods used to address issues of missing values or incomplete data are described in subsequent sections. For example, the total energy consumption of the DC was not available in any of the datasets. This either results in the approximation of some values for computed metrics or impedes the evaluation of some other metrics.

Transferability of the work depends on monitoring systems used in DCs, the quality and coverage of measurements data as well as individual DC characteristics. For example, DC providers might define and assess useful work of computational processes in a way most suitable for their infrastructure and purpose of DC operation. This current work shows a use-case of metrics that include estimation of useful work and motivates DC operators to closely investigate portions of IT equipment energy consumption used for jobs with different status of fulfilment but does not limit them in defining the types of jobs exit status and other inferences.

As a remark on terms used throughout this work to facilitate comprehension, the word "cluster" is dedicated to a set of servers connected in a separate infrastructure with its own network, load scheduling, and central management system. Several clusters in the use-cases are not interconnected and should be seen as independent structures both physically and logically. One cluster should be regarded as a small independent data center. For that reason, we do not evaluate characteristics which would cover several clusters within one DC but study them individually.

1.6 Novel Contributions

Overall contribution of this work is the identification of mutual interconnections between the smartness of the city, sustainability concept and DC involvement into urban operations. A three-phased methodology is proposed to assess DC energy efficiency as a main sustainability requirement imposed on the DC to provide benefits for smart cities while minimising negative environmental impact of large electricity consumption. A set of recommendations is formed based on unravelled pitfalls of the real DC clusters work.

The degree of transferability of applied methods depends on the monitored data of a DC willing to integrate proposed methods, however, the concepts covered in this work are useful for energy efficiency evaluation of any DC. This work showcases applicability of best practices and guidelines to a real DC and goes beyond the set of existing metrics for DC sustainability assessment. Contributions of three distinct phases are displayed below.

Phase 1 Contributions

C1.1. Assessment of the IT productivity metrics and waste energy evaluation based on collected real data over a period of 12 months to address the gap between metrics definition and their exploitation in a real DC context;

C1.2. Suggestions on energy waste and productivity metrics utility, namely Energy Waste Ratio and Data Center Energy Productivity, within the general methodology of energy efficiency assessment and overall DC sustainability framework;

C1.3. Proposal of a new metric, Carbon Waste Ratio based on Energy Waste Ratio, that links useful computing work, energy waste and its associated carbon emissions;

C1.4. A set of recommendations is proposed to enhance a DC cluster IT equipment energy productivity.

Phase 2 Contributions

C2.1. Thermal and energy efficiency policies for the DC are improved through real data center thermal data analysis, evaluation of thermal metrics and characteristics of DC IT room environment;

C2.2. Conducted analysis has increased operators' general awareness of possible thermal related weak points in DC thermal management;

C2.3. The problem of the air-cooling system which results in dangerous hotspots that could reduce IT equipment reliability and lifetime is highlighted;

C2.4. A list of thermal management and monitoring improvements is proposed for a DC cluster analysed in the Phase 2.

Phase 3 Contributions

C3.1. The hotspots are localised around individual cluster servers with the help of K-Means clustering algorithm applied to time series of IT room thermal characteristics;

C3.2. A set of measures is suggested to overcome an issue of hotspots in the DC cluster under consideration.

1.7 Structure of the Thesis

This thesis work is structured as follows:

Chapter 1: Introduction provides the background, motivation, research goals and objectives, key contributions and delimitations of this work;

Chapter 2: Related Work gives an overview of recent approaches towards smart cities and their sustainable requirements for DCs;

Chapter 3: Research Methodology introduces methods and techniques used in phases 1-3 of this work;

Chapter 4: Phase 1. Energy Efficiency Analysis of IT Processes covers power consumption and load scheduling data analysis of a real DC cluster through mathematical modelling to assess energy productivity of the cluster IT part;

Chapter 5: Sustainability Analysis of the Work outlines main sustainability contributions of this thesis;

Chapter 6: Conclusion provides a summary of findings and possible future work.

Appendices: description of datasets, **Phase 2, Phase 3**, list of recommendations for a DC in question.

2 RELATED WORK

The concepts of a data center, sustainability and a smart city are presented interconnected in the existing literature as smart cities rely on the development of ICT and DCs and pursues sustainability goals, DCs partially operate to satisfy smart city needs and tend to reduce energy consumption and thus environmental footprint and sustainability embraces a number of practices and approaches for both DC and smart city. Mutual interconnections between these three notions that appear in literature are studied in the current part of this work to further strengthen them with obtained results of three-phased data analysis.

2.1 Smart Cities Improved by ICT

The notion of smart cities appears in the 21st century with emerging ICT capabilities and rising environmental awareness as a trade-off with improved quality of life. The city is recognised as "smart" if it integrates enhanced technologies in one or several of the following sectors: education, governmental support, healthcare, transportation, safety, clean energy production and other industrial spheres [5]. Solutions deployed in a smart city aim to reduce negative environmental impact and increase the comfort of everyday life. Smart cities' solutions are empowered by technologies that typically rely on interconnected monitoring and reactive components, as well as large quantities of data generated by IoT and other involved systems [4], [16], [20], [27], [43]. Aggregation of historical data and data generated by societal use of applications contributes to the Big Data (BD) phenomenon with characteristics that match at least 3 to 7 V's versions of a BD definition [44], [45].

Smart cities are still in their early years of development, so notions and definitions regarding the concept of a smart city are being discussed in the literature. As outlined in the review paper [18], smart city is now a term that has recently outperformed digital city, information city and sustainable city in the number of citations and thus now is most widely used. The majority of works cited in the review paper include environmental awareness as a necessary point of smart city development. This point can be interpreted in a variety of ways, from achieving a balance between resource utilisation for urban needs and protection of the environment to energy-related savings, to overall thoughtful resource exploitation. However, the authors emphasise that growing cities that attract more people by good living conditions generate environmental outcomes that should be tackled within an umbrella of measures so that smart city and sustainable city would become interchangeable notions.

Following the discussion of a smart city as an ICT-enabled urban area, in the work [5] a variety of definitions of smart city and big data are shown as well as benefits of combining these two emerging principles in healthcare, transportation system, governmental use, etc. Authors propose a set of big data application requirements suitable for any smart city project, for example, security enhancement, governmental and citizen involvement, smart network, specialised platforms, enhanced algorithms, etc. The paper [5] is concluded with challenges concerning smart cities on a global scale, mostly from ethical point of view:

- Seamless data sharing between urban departments with varying privacy policies;
- Data format unification;
- Creating a knowledge base for a smart city with high interoperability between devices and platforms;
- Data quality enhancement, especially when collected from humans (tackling objectiveness) or from sensors of a third party;
- Data security improvement while it is being transferred via the network to different applications and actuators and identification of privacy rights of data owners;
- Decreasing the cost of smart projects and raising governmental and societal willingness to launch them;
- Development of smooth deployment and testing procedures so that new systems do not result in temporary problems of the integration stage in the sector that they are destined to improve
- Scalability of applications, especially under the circumstances of growing population that is prone to create increasing amounts of data in a smart city
- Reduced response time and enhanced reliability of real-time applications

The way these and other challenges are met by a certain city allows to place it on the scale of smart city maturity model described, for example, in [24], [25]. According to IDC

Energy Insights Smart Cities maturity model, each city can be placed on a specific level depending on the city's components and their performance: scattered (several smart projects are being developed, but not interconnected), integrated (initiatives are combined together and first positive results are achieved), connected (all projects coexist together and are managed by one committee) [24]. EUP maturity levels differ from the IDC levels in the sense that they are applicable to separate initiatives or projects and not to the whole city [25].

Several examples of smart cities are discussed in the literature. For instance, the paper [6] focuses on the integration of Big Data analytics in the smart cities, and through the case studies shows that Big Data analytics potentially can play an important role in the smart city environment and gives tools for business and research bodies to address the upcoming challenges of a smart city. It discusses some North European cities which incorporated several urban automated systems: waste management & inner city traffic are enhanced through smart applications in Stockholm. The city of Helsinki provides open public data stored in databases including transport, economics, conditions, well-being. Copenhagen aims to become the first carbon neutral capital by 2025, it introduces smart technologies to transportation, waste, water, heating systems and develops alternative energy sources.

Among the big data challenges reported in the paper [6], the authors identify business and technological concerns. The business issues consist in cost of essential devices, their scarcity, difficulties in planning an efficient solution, sustainable and secure use of stakeholders` information, and integration of cloud computing which may require data centers collocation for easier user access in various geographical areas. Technological challenges consolidate confidentiality of private data, efficient GIS-based 3D visualisation, support of a certain level of quality of service and enhancing computational intelligence algorithms for datasets of a smart city scale. Results of data analytics applied to big urban datasets are suggested to provide authorities a clear vision of current urban environment and become the basis for new legislation. For our study the paper gives insight into the data center role in a smart city and its place in future business models that involve big data processing and cloud computing.

2.2 Data Center Sustainability as a Smart City Requirement

We emphasise that a city is known as a smart urban environment if it has reached a level of environmental sustainability [3], [18], [24], [25]. DCs are pivotal actors of a technologically advanced smart city, must not disregard their role and responsibilities of maintaining a healthy environment and effective use of resources. It is, therefore, important to provide insight into the origins of DC environmental influence and explore the best practices proposed by international bodies (e.g. EU Code of Conduct for Data Center Energy Efficiency [8]) to address sustainability from the DC point of view.

2.2.1 Role of Data Center in Smart Cities

Smart cities extensively rely on big data processing thus far primarily provided by cloud technologies, and, therefore, DCs. Characteristics of computational, storage and network resources such as their reliability, availability and accessibility, security, and optimal power management are crucial for smart cities and their associated applications which can impact humans' life and safety [6], [19]. Overall, as an enabler of smart city services, DC's positive impact on the quality of life should outperform the negative environmental impact caused by indirect carbon emissions from electricity production, heat and material waste, as well as noise pollution that is expected to increase with the growth of DCs. However, limited attention has been accorded to the actual DC operation in the context of smart city and, a DC is often viewed as a separate area of study. This current study focuses on DC sustainability, energy and thermal efficiency in the context of smart cities.

2.2.2 DC Energy Efficiency

For the DC sector to continue its seamless integration in the smart city, pursuing energy efficiency is mandatory for a number of reasons listed further and explained by examples in the literature review afterwards. Firstly, DC is an integral part of a smart city as an enabler of city services, but at the same time, a huge consumer of energy. Secondly, energy efficient strategies can contribute to prolonged lifetime of the IT equipment through optimisation of its utilisation and decrease or slow down the amount of material waste generated by DCs. Thirdly, energy efficiency could also be interpreted as optimal thermal management of IT rooms and other places in the DC, which will positively impact the overall DC energy consumption and decrease heat waste. Moreover, integration of renewable energy is a plus to every DC site, as it allows to approach a problem of high carbon emissions caused by traditional energy production process through low-emissions procedure of energy generation.

Incessantly increasing demand in High Performance Computing (HPC) Data Centers require growing energy consumption, due to both data processing and cooling activities. For this reason, Data Center must be seen as a Cyber-Physical system, taking into account the thermal and computational resources [36]. This view of a DC coincides with the aim of the thesis investigation on energy efficiency of cooling system and IT Equipment (ITE) as of major energy consumers within the DC, where cooling refers to the physical part and performance of ITE concerns the cyber part of the notion. The findings of the authors in [36] contribute to the problem of estimation of "useful work" in terms of IT applications performed by DC and confirms the hypothesis about non-zero idle mode power consumption within the DC. The authors outline two problems to be analysed and solved: not-uniform DC`s workload overtime that results in fluctuations of power consumption, and not ideally proportional performance, i.e. non-zero idle power rates and non-linear power utilisation by DCs, which are shown on a case study example. Phase 1 of this current work can be seen as a continuation of the study in [36].

The study [31] reports changes in electricity use by data centers in the USA and worldwide through 2000 to 2005 and 2010 based on the data from International Data Corporation (IDC) on installed base of servers. It helps define four scenarios of growth in electricity consumption and identify existing challenges:

- Server peak power is different from server annual electricity use, which affects the trends and electrical network load.
- Network and data storage equipment electricity consumption should be also measured. Power needed for storage devices is defined by spindle movements, which differs with the growing density and capacity of storage facilities.
- Cloud computations decrease the need for installation of new servers and thus positively affect the electricity consumption. Nevertheless, there is too little data on the ratio of cloud computing servers within DCs.

2.2.3 Direct and Indirect Waste Created by DC

When a DC is not optimised, it contributes to different types of wastage. A DC generates physical waste during refurbishment and upgrade, heat waste as a result of servers processing IT jobs, and energy waste due to low computational productivity in comparison to energy resources used as discussed above. Moreover, these types of waste are interconnected: the IT equipment lifetime may be directly impacted by the temperature inside the IT room, unoptimised resource allocation, poor energy and cooling management. Reduced have an incidence about the rate of DC electronic waste generation.

LCA analysis and eco-labeling could be applied to tackle physical waste. Furthermore, thermal energy waste could be reused in the process called heat recovery, when heated water or air in the DC is directed to a heating system (within the DC or nearby buildings) that supplements existing heating processes [21]–[23], [46], [47]. Unfortunately, energy waste caused by inefficient use of electricity for cooling or computation cannot be reused.

Energy waste assessment has been addressed in academia and industry both qualitatively and quantitatively. Inefficient energy use causes increased electricity demand and also negative environmental impact if non-renewable energy is used in the DC. Some research work explores VM allocation-related *energy waste* that is particularly crucial for cloud paradigm in DCs which provide computing resources to users in the forms of infrastructure, platform and software as a service. Such work proposes Virtual Machine (VM) allocation strategies and algorithms which increase the performance and QoS characteristics of DCs [48]–[50]. In other research work, *energy waste* is discussed in terms of heat generation and in such cases, thermal energy reuse is suggested as a potential solution. For example, the heat recovery in smart cities can be used for heating (sometimes partially) the nearby buildings, or even the premises of the same DC to provide good working conditions for offices within DC premises [22], [46], [47].

Useful work, as opposed to *energy waste,* refers to the useful outcome of DC activity in terms of IT jobs processing. The definition is ambiguous, because useful results of data processing depend on application type and cannot be uniformly measured. Thus identified on the application level, *useful work* varies from the number of floating-point operations, number of service invocations, number of transactions, or another essence related to the individual application [9], [51]. In [52] the authors classify tasks failures-based causes such as server or software failure, scheduler issue and evaluates energy spent on such tasks.

2.2.4 Integration of Renewable Energy Sources

A part of sustainable DC strategies, reducing the carbon footprint of DC worldwide is a considerable challenge under the pressure of big data deluge and smart city-related processing [53]. A series of projects has been created under the paradigm of sustainable DC [54]. For example, DC4Cities focuses on creation of energy adaptive eco-friendly DCs that operate to support smart city applications. Thus, DC involvement in a Smart City life is defined by storing and processing the data coming from smart sensors and administration procedures. This data modification and knowledge extraction may simplify decision-making process. The authors in [27] mention Data-Information-Knowledge-Wisdom (DIKW) pyramid which has raw data as the basis, contextualised data or information on the second level, actioned or processed data at the knowledge level and automated data representing the wisdom level in the smart city context, because it helps increase effectiveness and add value to the decision-making process.

The project DC4Cities assumes that smart city is focused on increasing the share of renewable energy sources in their energy supply, which is aligned with the citizens involvement in sustainability goals and active use of smart home systems. Energy mix within smart cities is thoroughly studied during the DC4Cities project realisation [55], as well as possible evolution of electricity grids components to smart grids with extensive share of distributed energy sources. Trials of the developed methodology are made on the sites of Barcelona and Trentino DCs.

As aforementioned, in recent years serious effort has been made by consortia involving the industry, academia and public authorities to address the increasing energy demand challenge of the DC sector. Although such effort does provide valuable tools and practices towards reducing energy consumption, they should be merely considered as the beginning of a journey towards environmental targets. In a smart city context, past energy inefficient practices, such as ignoring the potential use of waste heat or renewable sources, are not sustainable. Now, the research work proposes to plan DC activities according to forecasted availability of renewable power sources and clean energy from the grid to minimise associated carbon and equivalent emissions [21]. The Real time workload and Delay Tolerant workload developed in [21] could be used with two advantages: (1) better management of task scheduling, (2) better adaptation between DC activities and green energy produced locally (solar panel on DC roof, for example) for reducing carbon emission.

2.2.5 Sustainable DC Guidelines and Best Practices

A lot of industrial and research effort has been dedicated to defining a sustainable DC and, more importantly, to providing suggestions on the incorporation of sustainability goals and practices. They cover all aspects of DC energy efficiency mentioned before and go beyond them. The sustainability-related practices and standards encompass Life-Cycle Assessment (LCA) of DC operations that include equipment, energy, and other resources use throughout the DC lifecycle, including its expansion, and upgrade of hardware as well as software components. LCA is a methodology that could assess interlinked environmental impacts of a DC while single-issue metrics do not provide a holistic overview [56].

Several guidelines for sustainable DC operations have been developed by different research and industrial bodies, as well as voluntary programs (e.g. Code of Conduct for Energy Efficiency in Data Centers [8], [57]). They cover renewable energy use, power efficiency in computational and cooling processes, recommendations for appropriate hardware, software, reduced energy consumption, and electronic equipment disposal.

Specifically, Energy Star programme has developed a set of requirements concerning energy use and optimised operations that should be satisfied by IT equipment and its manufacturers to be assigned an eco-label [58]. ASHRAE has developed several guidelines concerning power equipment and DC operational requirements for in the pursuit of sustainability [59], [60]. JRC Commission has proposed a holistic framework for assessment of the level of sustainability practices integration in a specific site in its Code of Conduct for Energy Efficiency in Data Centers [8].

The Code of Conduct provides a methodology for DC operators to assess their sites in terms of general policies adoption, IT, power use and cooling efficiency, building exploitation, and monitoring. Application of this methodology results in a DC evaluation on the scale from 1 to 5 (best score) in all DC areas that the methodology encompasses. This evaluation also allows DC operators to compare their DC's performance and metrics indices before and after some sustainability-related actions are undertaken. A more detailed overview of the practices and guidelines is displayed below.

2.2.5.1 EU Code of Conduct Guidelines

The practices concerning the entire data center, comprise, for example, forming an approval group for important decisions to regulate them in accordance with energy efficiency strategy, auditing the equipment to measure and optimise its usage, prepare plans for environmental and energy management. Air quality monitoring is a suggestion after ASHRAE 2011 white paper results (2011 Gaseous and Particulate Contamination Guidelines for Data Centers') which brings focus to dangerous corrosive elements in the air that can influence the equipment quality and lifetime.

Guidelines on provisioning and resilience level of data center operation highlight that infrastructures should be built as needed for business requirements and adjusted to maximise energy efficiency under conditions of partial and growing load of the facility. The latter adjustments are possible when power and cooling systems have several levels of resilience and when the whole DC is planned to be modularly scaled in the future.

The best practices also cover the process of choosing appropriate IT equipment with the help of customised or standardised metrics, for instance, making use of Energy Star, SERT or SPECPower. Not only will these measures improve energy efficiency, but also bring reduction in average utilisation cost. When purchasing a set of new equipment, it is crucial to verify temperature and humidity operation levels. Operators should set them carefully to consider the designed power capacity. During the selection process, equipment benchmarking should be verified to conform to the full allowable temperature ranges. Equipment with energy efficiency labelling and energy-aware design should be matched with the infrastructure and room configuration. Once cooling is concerned, air flow of new devices is required to match existing air flow schema, as well as the added IT equipment should comply to temperature and humidity levels typical or adjustable on the site.

As suggested in guidelines, equipment acquisition and deployment should be adjusted to business requirements and avoid overprovisioning. Same practices concern software selection and development. All the existing equipment should be carefully audited and analysed to remove or power off idle and standby components.

Major data management issues are concerned with unnecessary data duplication or heavy protection and archivation. Thus, data storage policies should be developed by the organisation and characterise data to preserve time limits and protection levels. Optional measures involve efficient snapshots to be used and data cleaning days to be organised, which in the long term should lead to overall storage volume reduction.

Guidelines on effective cooling include containment and separation of hot and cold air flows, positioning of blanking panels to eliminate air recirculation where the space is not occupied by any equipment and maintaining raised floors without apertures or obstructions. Recirculation should be minimised through tuning the pressure of air stream slightly higher than that of IT equipment air flow. Equipment requiring different environmental conditions should be separated and in case of colocation data centers charged with respect to the strictness of SLA in order to incentivise energy efficiency concept through billing policies. Cooling equipment settings should be reviewed upon every alteration of the facilities and IT equipment placement, for example, cooling system should be turned off in empty rooms, cooling units should be calibrated not to work against each other, they should also be properly maintained and cleaned. Temperature and humidity settings are deeply connected, since with overcooling comes increased humidifier energy consumption. Thus, raising intake air temperature, widening humidifier range and optimising water temperature to set it to the optimal level. Free cooling could facilitate easy energy conservation by allowing fresh cold air to cool the air or water used in DC cooling systems. It is also a good practice to use centralised humidity controller that would eventually benefit to potential use of adiabatic humidification and free cooling.

When cooling system requires refrigeration, the following aspects are important to consider: high Coefficient of Performance of chillers, decreased difference between cooling system temperatures, adjusting the cooling system to expected continuous partial load, including speed drives for cooling system elements and possibility of "free cooling" is also an important characteristic in the areas where this type of temperature management is possible.

Furthermore, cooling systems aimed at IT equipment appropriate conditions should not be affected by temperature management for other purposes. Computer Room Air Conditioners should allow variable speed of fans and configured to control on supply temperature, in order to handle varying conditions and loads. At the same time, operators should avoid multiple humidifier controllers. Waste Heat reuse options are discussed in the report, listing direct reuse of warm air in offices adjacent to the DC, introducing additional heat pumps to warm the water and heat nearby buildings or districts. Efficiency of such undertakings is proposed to be measured with Energy Reuse Factor and Energy Reuse Effectiveness metrics from The Green Grid.

Power equipment guidelines comprise suggestions on modular scalable power supply units, which comply with EU Code of Conduct requirements and perform efficiently when partially loaded. Existing power equipment should be audited and adjusted to the frequency of their usage. Their power factor should be high enough to guarantee less negative side effects such as electrical inefficiency and cable losses.

Energy use for the overall non data floor areas should be as well optimised according to building standards. Simple practices of switching off the lights when they are not needed, using energy efficient bulbs and providing energy reports from the hardware installed in the offices could improve sustainability on the site of a DC. DC should be located and engineered so as to benefit from all the natural conditions, facilitating free cooling, avoid high humidity areas, possibly collocate with the power source and capture rainwater.

The study [32] reports DCs participation in the EU Code of Conduct (CoC) initiative by 2016 and reveals that CoC Energy efficiency voluntary initiative is widely supported by

DCs across the EU. By December 2016, 325 DCs have applied for the CoC Participant status and 289 of them have been approved with average PUE value of all the latter sites of 1.8. The majority of approved DCs have applied 26-50 best practices while the number of mandatory practices has been 81. The results of the study confirm that DCs are heading toward sustainability and energy efficiency practices, but it s challenging to comply with all mandatory guidelines, especially owing to the fact that both energy requirements and CoC set of practices are updated every year while the latency of DC retrofit and upgrades is still high.

Overall, effective environmental management requires energy use monitoring, specifically of incoming energy consumption, IT energy consumption, room-level metering of supply air temperature and humidity, CRAC/CRAH unit level metering of air temperature as well as more granular metering. Undertaken measurements should be further analysed and reported to preserve statistics of energy use and economisation levels and use it for improvement of DC sustainability level. Usually, a set of metrics is exploited to provide final step of DC assessment, after all the data is gathered. Discussion of DC efficiency metrics is placed after the following part that particularly focuses on thermal guidelines from ASHRAE.

2.2.5.2 ASHRAE Thermal Management Guidelines

ASHRAE started unification of the environmental parameters which affected DC computing efficiency, performance, availability and reliability in 2004, and created their first set of thermal guidelines. In response to metrics development, namely, to the wide use of PUE metric, the organisation has created additional environmental equipment classes and guidance on their usage. The major achievement of the TC9.9 ASHRAE committee described in the whitepaper [59] is that ITE manufacturers agreed on recommended and allowable ranges for operational environment, which the committee summarised in the guidelines. Furthermore, the guidelines are formulated in terms of recommended and allowable envelopes, i.e. suggested sets of limits for thermal characteristics, for DC operators and aimed at two main factors: high reliability and energy efficiency. DCs are proposed a methodology to create their own envelopes with more suitable standards tuned for a specific site, if there is such need. It is emphasised in the cases when DC operators have an intention of changing these ranges, a set of possible side effects such as noise, variative speed of chemical reactions should be considered and tackled with, and the climate conditions should be primarily investigated.

Before the first edition of ASHRAE Thermal Guidelines for Data Processing Environments, DCs typically cooled their IT rooms down to 20-21 °C using the most stringent thermal requirement of all the equipment present in the rooms and a safety factor. One accomplishment of ASHRAE TC9.9 is the agreement between different major IT manufacturers on what thermal levels are regarded to as recommendable and allowable, so that DC operators are permitted to manage IT room with diverse components with a single agreed set of temperature and humidity ranges. The thermal ranges are expanded according to global desire to minimise the total cost of ownership and improve energy efficiency.
The guidelines delineate six types of IT equipment, out of which four classes refer to DC premises and two types to individual or point-of-sale use. The classes are attributed with thermal characteristics according to the mode of use of the equipment comprising them and statistical investigation with the help of industry representatives. The stringency of the temperature, humidity and other thermal characteristics is divided into two levels: the one called recommended envelope and the other being allowable limits. Recommended envelope in a set of thermal ranges which ensures the most reliable and reasonable operating environment, from energy efficiency point of view. In the meantime, if the ranges are exceeded within the allowable limits for short periods of time, it does not lead to deterioration of reliability.

The organisation suggests a clear set of steps for DC optimisation projects, where operators are proposed to first check their DC compliance with general best practices, then to apply ASHRAE classification and determine allowable limits suitable for a particular DC to further perceive to what extent the ASHRAE recommended operating envelope can be violated, and finish the optimisation project with optional customisation of recommended envelope of characteristics. The final optional step importantly includes consideration of a list of options for advanced optimisation brought together by ASHRAE committee. The options cover internal and outdoor temperature and humidity rating that includes monitoring of all types of IT, network, power and cooling equipment, airflow management and cooling architecture, requirements for economizer and chiller as well as type of DC affecting the reliability and availability requirements, and overall project type, covering introduction of new components, upgrades of existing components or retrofit.

Thus, optimisation projects should start with monitoring and statistical analysis of available data on climate, indoors thermal characteristics and architecture of cooling IT systems, with a goal to identify possible limitations, for example, some cooling equipment only allow maximum inlet temperature of 30°C. In addition, it is necessary to choose power equipment considering its mutual placement with IT equipment. For example, rack PDUs are usually placed in the back of IT equipment racks or at their sides, which makes it difficult for rack PDUs to access the cooling air from the cold aisle and requires power equipment to be designed to meet or exceed the IT equipment exhaust air. Having obtained the full picture of the DC operation and climate features, DC operators are suggested to start customisation of recommended envelope. All the optimisation steps are summarised in a flowchart as a guideline for DC operators for ease of use (Appendix F in [59]).

IT equipment, which operate properly and reliably at higher temperatures and relative humidity levels, pertain to ASHRAE classes A2-A4. Important to emphasise that A2-A4 ITE classes do operate reliably within allowable limits proposed by ASHRAE, but they still experience some increase in failure rate with changes of the working environment. While limited data is available for ITE of such classes, it is assumed that in the aftermath of adjusting temperatures to these classes DC will necessitate dealing with increased noise levels. Noise level augmentation by 3-5 DB caused by 2°C rise of inlet server temperature is referred to as a reasonable forecast. The problem of noise pollution produced within DCs may become an extreme challenge, if estimated noise level elevation reaches 97.9 dB(A) at 45°C ambient temperature and require action according to countries legislation in healthy working environment. This fact might inspire the idea of sound energy conversion in a DC into electrical energy useful for the data center, although this area requires further research.

Changes in DC thermal setpoints, introduction of economizers and chiller-free techniques, with the intention of increasing energy efficiency and improving the PUE value, are prone to side effects studied in ASHRAE guidelines [59]. The side effects include enlarged server power use, higher system air flow speed, increased noise levels, transformed reliability, performance and cost trends. One downside of the ASHRAE guidelines is that the base level of reliability is defined as failure rate at 20°C server inlet temperature, which may vary for different DCs and be incorrectly interpreted in terms of proportion of active servers.

Adoption of chiller-less economizer technologies requires prior analysis of time-at-temperature climate data to assess applicability of such changes. Thus, dry bulb, wet bulb temperature and humidity measurements should be obtained for the DC location as well as equipment failure rate in the projected resulting indoor conditions, i.e. a certain temperature rise from outdoor air to server inlet diapason is equally required. A similar research was conducted for some cities spread across the world in the same ASHRAE whitepaper. Although failure factors for locations with high annual temperatures increase in case of some economizer configurations, in general, the failure rate is predicted to be lower than those expected under continuously high temperature steady state operation conditions. In the view of the fact that IT equipment specifications and climate have undergone some changes since the release of ASHRAE guidelines, the exact numbers are exempt from the current work.

DCs, especially those which use economizers, need additional dehumidification, particulate and gas filtration in geographical locations with high pollution, temperature and humidity levels to mitigate condensation and pollution risks for ITE. High humidity levels are dangerous for dielectrics and may cause delamination of fine materials. The risks get higher when negative factors are combined: dust, gaseous or particulate pollutants together with high relative humidity facilitate corrosion of copper and silver miniature components. Therefore, it is recommended to maintain relative humidity less than 60% and apply filtration.

ITE performance might suffer from loose thermal limits, if the equipment lacks power management capabilities which would respond to variation in environmental conditions in the room. Therefore, DC operators are suggested to consult ITE manufactures to understand how the product will react to limits of allowable thermal envelopes. Added expenditure is also expected in DCs which aim to upgrade IT equipment to classes with higher allowable thermal limits and benefit the environment. However, improved materials for ITE, enhanced cooling systems with potentially lower energy consumption, or in general, new equipment adds to the cost of maintenance of a DC, which should be recalculated considering trade-offs between cooling solution, temperature ratings of ITE and its performance capabilities.

2.2.6 Metrics for DC Assessment

Data centers vary in their size, operational purposes and level of confidentiality of processed data [13]. They can be public, private and confederated and display different levels of reliability (probability of failure over a time period), availability (average time per time period of working without any downtime) and redundancy (availability in the event of failure). Several reports of GeSI, DatacenterDynamics and others on DC environmental impact and power usage focus on environmental impact of DCs, specifically their energy consumption and indirect carbon emissions. Although they show different levels of precision and representativeness [13], these steps forward to the prediction of DC environmental impact display rising consciousness of research and industry communities as well as DC providers about the DC energy demand and required improvements. Among current challenges authors identify the absence of indices which can show interrelations of management actions on different environmental consequences. However, a plethora of metrics exist for evaluation of distinct DC features, for example, its indirect carbon emissions, air/water/energy usage effectiveness, levels of systems' optimisation, and other specific characteristics. In general, a metric can be defined as an empirical, objective assignment of numbers, according to a rule derived from a model or theory, to attributes of objects or events with the intent of describing them [61].

A growing body of literature has proposed, examined, and critiqued the metrics for DC assessment [9], [34], [37], [51], [62], [63]. For example, in [9], a taxonomy of the state-of-art DC efficiency metrics is presented for further use by DC providers and researchers. A plethora of metrics is categorised (by their DC core dimensions) into groups: energy efficiency (e.g., DCeP, PUE), "greenness" (e.g., CUE, WUE), cooling systems (e.g., HVAC System Effectiveness, Recirculation Index), thermal and air management (e.g., Rack Cooling Index, Return Heat Index, Recirculation Ratio), performance or productivity (e.g., Idle-to-peak Power Ratio, Data Center Performance), security (e.g., Accessibility Surface, Defense Depth), network (e.g., Network Power Usage Effectiveness), storage (e.g., Response Time, throughput) and financial impact (e.g., CapEx, OpEx, ROI) metrics. Exhaustive information is given on each metric, including its expressivity, advantages and limitations, interrelationship and concepts which lie in the basis of the metric. The main advantage of the work is that it provides a quick access to a needed metric or group of metrics by addressing the category. The authors outline the current challenge in the area of metrics: no metric describes all the data center components at once, neither normalisation strategy nor metrics exist to compare different data centers, dependence of metrics on sites individual features, difficulties on metrics application to co-location DCs and overall complexity and unpredictable nature of data centers. While some work highlights new emerging metrics like Datacenter Performance per Energy (DPPE), ASHRAE Performance Index (PI) [64], others review and discuss existing metrics, such as internationally recognised but not exhaustive PUE metric [40], [41], [65], [66].

Special attention is dedicated to thermal metrics that identify efficiency of the cooling equipment and IT room design [14], [34], [38], [63], [67]. These metrics are divided into two categories that can identify global (e.g., RTI, SHI and RHI) and local (e.g., RCI) thermal and air-flow phenomena. They allow to discover infrastructurally caused disadvantageous processes of air bypass and air recirculation. It remains true for all the metrics that they are only single-purposed and do not assess DCs holistically unless they are used under a methodology that would encompass all available measurements and apply a set of metrics to evaluate the DC from different points of view. Therefore, this current work will only expand relevant metrics that are essential for the undertaken analysis.

2.2.7 Use Cases. How Do Real DC Providers Approach Sustainability?

Cloud Data Center providers claim that use of their services is more optimal and sustainable in terms of carbon emissions. We chose three cloud service providers, Amazon Web Services, Microsoft Azure and Google Cloud, which appear in Top-5 lists according to TechRadar reviews resource and Gartner Magic Quadrant 2018 [68]–[70], to make a brief overview of the companies' efforts to improve energy efficiency of their services.

The cloud actors mentioned above were also reviewed in 2017 by Greenpeace report on renewable energy use in data centers [71]. The report ranked Google Cloud, Microsoft and Amazon with A, B, and C grades correspondingly. The ranking criteria were based on ratio of clean energy sources in the whole range of sources used, transparency of activities and availability of information on sustainable efforts, adoption of renewable energy strategies and investments in this area, energy efficiency and level of attracted support in sustainability initiatives. It is also remarkable that the main focus of sustainable strategies varies for these three providers, being on-site renewable energy farms for AWS, carbon-neutrality through renewable energy purchases and grid interactions for Google and Microsoft, also thermal management for Google and IT operational energy efficiency and life-cycle carbon emissions assessment for Microsoft, as observed from the companies' websites and sustainability reports.

AWS based its estimations on NRDC report and deduced that, in the cloud, customers use only ¼ of the number of servers they would use on-premises [72], [73], thus leading to less power consumption rates and carbon emissions. AWS also claims that with the use of carbon-intense power mix together with optimised servers involvement, carbon emission reduction potential is 88%. This is the main driver for the Amazon company renewable energy sources development. With 3 wind farms and 6 solar farms, AWS reached an intermediary result of 50% renewable energy usage in January 2018 on its way to pursue the goal of being 100%-powered by renewable energy [73], [74]. The future goals include solar systems installation on in 50 facility rooftops by 2020 [75]. Focusing on sustainability together with employees, AWS encourages them to commit to work avoiding the use of personal cars and reduce packaging. One of the reasons for getting a low grade for renewable energy use from Greenpeace was the lack of transparency. Indeed, sustainability timeline published on AWS website starts from 2014, while, for example, Google is known to be carbon-neutral from 2007. So, AWS has made a leap to achieve energy efficiency and sustainability adoption on a high level since 2014.

Google Cloud purchases renewable energy and balances electricity use with sustainable resources provisioning, thus the company's renewable energy purchases zero out the entire carbon footprint of their electricity use [76]. Therefore, the company has improved since

2017 Greenpeace report on data center sustainability assessment [77]. The company assures that it provides energy efficient storage with fast data access [78]. Google Cloud shows remarkable trailing twelve-months PUE of 1.1 as a result of a set of measures undertaken by the company, which comprise high granularity continuous monitoring of energy use on all sites, customised efficient servers and power paths, airflow management, including blanking panels and plastic curtains, setpoint adjustment to 26.7°C, application of evaporation and sea water free cooling, recycle and reuse strategies only, since the company diverted 100% from the landfill [79], [80]. A worth mentioning fact is that while 100% carbon neutrality level is reached by Google data centers, carbon emissions are still present in their reports, although leveraged by renewable energy purchases.

In addition, Google reported that the company successfully applied Neural Networks, a Machine Learning (ML) technique, for predicting the PUE, as a proof-of-the-art, in which standard predictive modelling with fixed formulae resulted in large errors, because they failed to express complicated interdependencies within a dataset [81]. Automatic prediction of PUE can leverage performance alerting, real-time efficiency targets adjustment and troubleshooting, also DC management planning without physical changes and reduce uncertainty of future changes. The case study has shown that increasing the IT load in the range of 0-70% yields large efficiency gains which decline only slightly with IT load increasing even further after 70%, due to the cooling plant higher useful work percentage. It is also concluded that increasing the number of operational chillers, process water pumps, outdoor wet bulb temperature, outside air enthalpy (total outdoor air energy content), or any similar physical or management changes that increase the load on thermal control system, result in the rise in the PUE. The work has demonstrated that given a sufficient set of DC parameters, it is possible to predict PUE variation with increase of Process Water Supply Temperature, prove errors in new meters installed on site, optimise operational parameters in case of DC plant reconfiguration in the real use-cases.

Microsoft reported that running applications in Microsoft Azure cloud is up to 98% more carbon friendly and uses 22-93% less energy than when they reside on enterprise data centers [82], [83]. Being a cloud provider for 140 countries [84], Microsoft has created a set of policies to promote Cloud for Good initiative encompassing internet security, rural broadband gap reduction, realisation of the AI for Earth educational initiative which raises availability of cloud resources for research, clean energy promotion and other social and technological aspects of sustainability in the global view. In the Microsoft blog, the company's president reported that they have operated in a carbon neutral manner since 2012. The company used internal carbon fee model, charging units for their resultant carbon emissions to incentivise development of carbon-saving technologies. In Microsoft report [83] a comparative analysis of energy use and carbon emissions of Microsoft IT applications, compute and storage resources was conducted as opposed to equivalent on-premises deployments. The study showed that overall energy and carbon emissions could be reduced by switching from on-premises infrastructure to Microsoft cloud provider, mainly caused by decreased electricity use per useful output. However, the term of useful output is not explicitly defined, and the wide ranges of potential improvements predicted, such as 22-93% energy savings mentioned above, lack granularity and detailed explanation.

In essence, major cloud services providers mainly focus on carbon neutrality through renewable energy purchases and/or solar and wind farms installation directly on their sites. Some energy saving and waste reducing guidelines are provided to DC employees to further reduce environmental impact of the companies' environmental footprint. Cloud providers are motivated to be transparent in their environmental initiatives and DCs' assessment through globally accepted metrics, out of which only PUE is widely used so far.

3 RESEARCH METHODOLOGY

This current work encompasses several phases of the exploration of DC energy efficiency. In this section, the overall methodology provides an idea of research stages. Furthermore, every phase of the work is summarised. However, detailed methodology for every phase is given in the dedicated sections.

As discussed in the Background Section 1.1, sustainability of a data center is fostered through the pursuit of energy efficiency in all possible operational aspects, including IT jobs processing efficiency, optimal power, materials resources allocation, appropriate cooling technologies. Therefore, energy efficiency investigation is the main focus of this current work and the basis for data analysis throughout all phases of the work.

Fig. 2 shows an approach of this work to assessment of DC energy efficiency in the context of a smart city. The diagram in Fig. 2 should be read from left to right. As shown in Fig. 2, analysis is performed in three phases and is based on available monitoring data concerning IT and thermal parts of DC. It is noteworthy that the DC under consideration processes jobs that have been designed for a broad range of scientific research computations and, among others, for smart city applications such as smart homes, air quality modelling and monitoring. We closely analyse IT room energy efficiency of DC clusters that have monitoring devices installed to measure power consumption, jobs scheduling data, and some thermal characteristics. These data are analysed within three phases. The phases are reported in chronological order of the work: Phases 1 and 2 bring about the necessary breadth of the work concerning energy efficiency analysis of thermal and IT dimensions of the DC; Phase 3 extends analysis in Phase 2 through the identification of thermal management pitfalls. The current work is organised in such a way that Phase 1 research work is discussed in the main body of the thesis while Phases 2 and 3 are summarised in the Appendices 2 and 3. However, it ought to be noted that the focus of this research work is DC energy efficiency assessment, relevant metrics, and evidence-based recommendations to DC owners. As Fig. 2 depicts, the analysis results will provide evidence-based energy efficiency related recommendations to the DC owners.

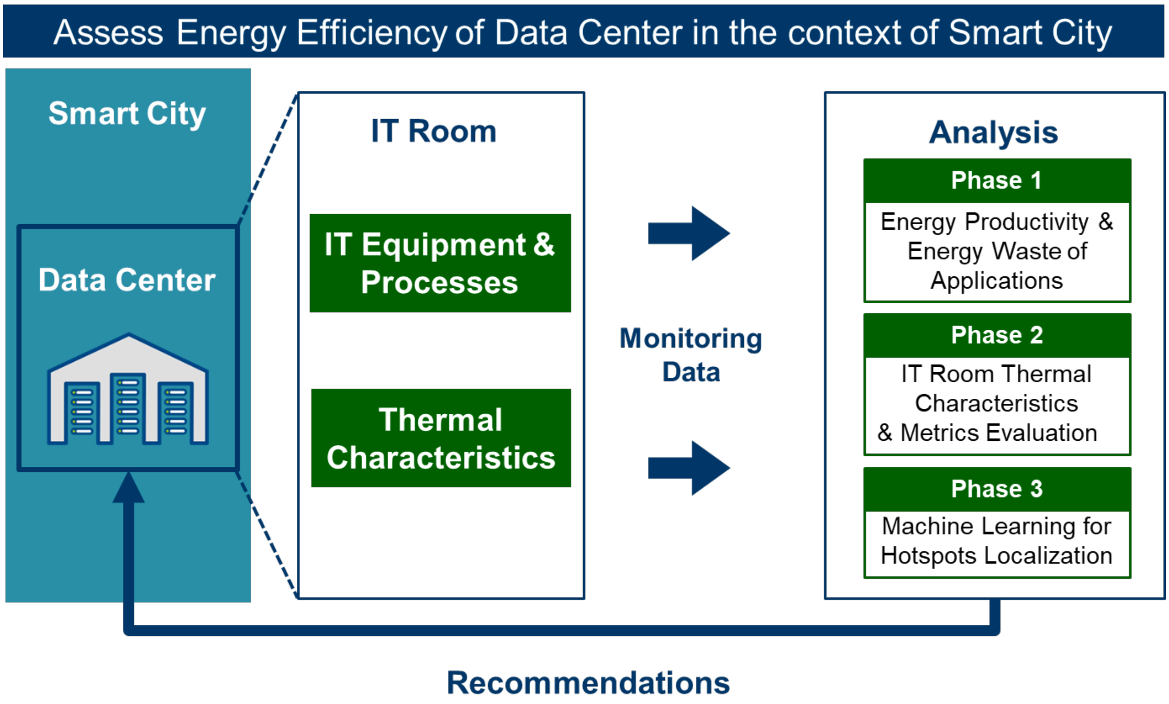

Figure 2. Overall thesis methodology comprising three phases of the work

Each phase has a separate methodology; however, all the phases assume the same structure based on an adapted data lifecycle methodology. While data lifecycle methodology exists in different forms, it typically includes stages of data collection, preprocessing, analysis, results exploitation and the cycle repeats. The scope of this current work does not cover data collection as the data has been provided before the commencement of this work. Thus, only stages of data preprocessing, analysis and results exploitation are implemented. The cycle is not iterated due to time constraints and lack of direct access to the monitoring systems settings. For example, commercial platform used for load scheduling and workload monitoring has a fixed algorithm of data acquisition and is not subject to changes. This commercial platform is called Load Sharing Facility (LSF), an IBM product, that will be discussed in Phase 1. Finally, adapted data lifecycle methodology of each phase is realised with Python programming language suitable for big datasets, computation of statistical characteristics, mathematical modelling and visualisation. Initially, all datasets have been received in .csv or .xlsx format.

Phase 1. Energy Efficiency Analysis of IT Processes

With the purpose of calculating energy waste and useful energy consumption of IT

Equipment, a mathematical modelling approach has been used. It consists in applying the law of energy conservation to available data, integrating time series measurements of power consumption and evaluating the hourly amount energy required by each process in one core. Further, categorisation of jobs processed by the cluster is performed to make a distinction between jobs which cause energy waste for different reasons. The Phase 1 section also includes evaluation of carbon emissions associated with *useful work* and wasted energy consumed by the cluster and a new sustainability-related metric, Carbon Waste Ratio, is proposed.

Phase 2. Analysis of DC Thermal Characteristics (see Appendix 2 for details)

Phase 2 is dedicated to the assessment of real thermal conditions in the IT room of one DC cluster as compared with thermal equipment setpoints and guidelines. Hidden factors such as bypass, recirculation, hotspots and partial rack overheating can negatively affect the health of IT and power equipment that is critical for the DC. Based on real data from server-level sensors, data analysis is conducted with the aim to identify potential risks caused by the possible presence of aforementioned hidden factors. Specifically, this phase involves extensive statistical analysis of available thermal data, global and local thermal metrics evaluation, investigation of possible correlation between power consumption of server components and temperature variation.

Phase 3. Machine Learning for DC Thermal Characteristics Analysis (see Appendix 3 for details)

To provide suggestions for hotspots localisation as well as categorisation of nodes based on surrounding air temperature ranges, Machine Learning techniques are used in Phase 3. Variability of thermal data and uncertainties in defining temperature thresholds for hotspots identified in Phase 2 have invoked a need for unsupervised learning. Therefore, a clustering algorithm is applied in Phase 3 to address the challenges that are beyond typical statistical techniques. In this phase, the number of clusters is determined using two indices (Silhouette metric and Within-Cluster Sum of Squares), and available thermal characteristics (i.e. exhaust temperature, CPUs temperatures) are inputs to a clustering algorithm. Subsequently, a series of clustering results are intersected to unravel nodes (identified by IDs) that frequently fall into high temperature areas of the cluster racks.

4 PHASE 1. ENERGY EFFICIENCY ANALYSIS OF IT PROCESSES

In this part of the work the concept of energy efficiency is approached through definition and assessment of *useful work* and *energy waste* of IT jobs processing based on real data related to the power consumption and load scheduling within HPC cluster CRESCO4 hosted by DC within the ENEA Portici Research Center. To evaluate energy consumption of IT processes, energy conservation law is employed in this part of the work to build an approximation of every process energy consumption from one core during each hour of monitored period. Several productivity metrics are then applied to evaluate *useful work* and *energy waste* on different levels of analysis granularity: these of individual jobs, cluster queues, groups of parallel and serial jobs and the whole cluster. Special attention is given to the interpretation of energy consumption profiles from sustainability point of view, i.e. in terms of associated carbon emissions. Additional point in favor of the current work is that the real data from a working DC is used for analysis, thus, in comparison with simulations of a DC operation, it shows real issues which should be addressed both by DC operators to optimally manage DC processes and by the user side to improve their applications performance. The research contained in this section synthesizes and develops the results published in [85]–[88] during the thesis work.

4.1 Data Center Facility and Datasets Description

The cluster CRESCO4 consists of 38 Supermicro F617R3-FT chassis, with 8 dual CPU nodes each. Each CPU is of the type Intel E5-2670 and hosts in its turn 8 cores, which results in a total number of 4864 cores. The CPUs operate at a clock frequency of 2.6 GHz. Furthermore, each core of the system is provided with a RAM memory of 4 GB. Computing nodes access a DDN storage system, constituting a total storage amount of 1 Pbyte. Computing nodes are interconnected via an Infiniband 4xQDR QLogic/Intel12800-180 switch (432 ports, 40Gbps).

This section exploits available data gathered on CRESCO4 cluster of ENEA DC during the period from February 2017, to January 2018. Datasets have been obtained from (1) Platform LSF (Load Sharing Facility) job scheduler and (2) Zabbix power consumption monitoring tool. To clarify, the words "process", "job", and "application" are used as synonyms in this section.

Briefly, LSF is a workload management platform and job scheduler for distributed HPC systems. This platform is concerned with deciding which process is to be run and is designed to keep CPUs as busy as possible. The LSF dataset covers details about the number of cores assigned by the scheduler for every process, 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"). Zabbix dataset contains average level of power consumption, minimum and maximum registered power consumption for each hour. These datasets are intersected based on monitoring time period, and the resulting dataset covers 11 months from 12:00, 19th of February 2017, to 12:00, 25th of January 2018, divided by 19th day, 12:00, of each consecutive month except January 2018.

The task scheduling of the cluster is based on First Come First Served algorithm. The queues characteristics are reported in Table I in [85]. The cluster processes approximately 40 types of applications and benefits several fields of research, such as climate modelling, renewable energy, environmental issues, materials science, efficient combustion, nuclear technology, plasma physics, biotechnology, aerospace, complex systems physics, HPC technology.

4.2 Methodology

Initial goal of this thesis phase is to obtain energy consumption data from available datasets. Once energy use is evaluated, analysis proceeds with energy efficiency investigation and finding answers to other research questions including recommendations on enhancing DC sustainability. Thus, datasets processing has been done according to adapted data lifecycle methodology as depicted in Fig. 3. It shows that data preprocessing stage covers formatting available datasets, their intersection and conversion of all the fields into the formats that would allow further processing (e.g. timestamps de-/encoding). Further, data analysis is divided into two major steps of mathematical modelling that would make a foundation for the second step, energy efficiency analysis. Energy efficiency analysis stage includes IT jobs categorisation, evaluation of productivity metrics and indirect carbon emissions. The phase results in recommendations for improving energy efficiency of IT equipment and jobs processing to feed into overall DC sustainability. Data analysis steps are explained in the following parts of methodology subsection.

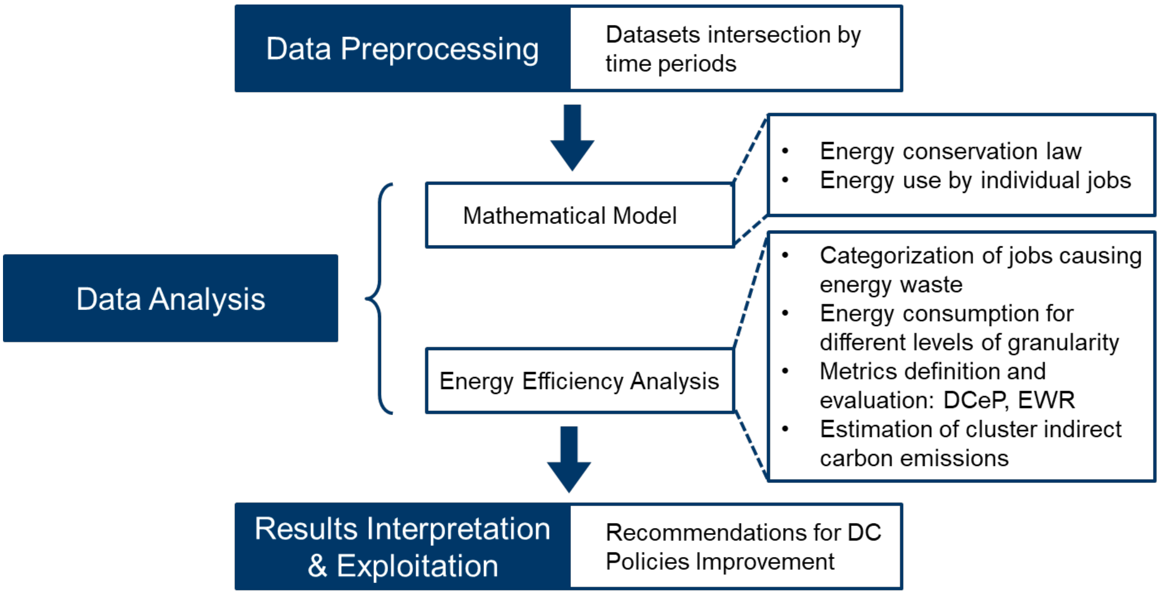

Figure 3. Phase 1. Data Lifecycle methodology adapted to mathematical modelling and energy efficiency evaluation of DC IT jobs processing, including metrics evaluation and estimation of indirect carbon emissions.

4.2.1 Mathematical Modelling for Estimation of Energy Consumption IT jobs

To estimate effective energy consumption by the cluster IT jobs and energy waste created by incorrectly finished jobs, energy conservation law has been applied to the combination of available data to find a set of introduced unknown variables and further estimate energy consumption of every process. With this intention equations of the energy conservation law have been expressed in terms of available characteristics of the DC cluster, see Equation (1):

$$\sum_{i=1}^{K} \int_{t_{i,j}^{0}}^{t_{i,j}^{1}} c_{i,j} \cdot x_j \, dt = E_j, \quad j = 1, \ldots, N, \tag{1}$$

for each hour j of every month, $j = 1, \ldots, N$, where N stands for the number of hours within one month. The right part of the equality represents the value of energy consumed by the cluster. The left part of the equation has been formed by the cluster power load generated by jobs, which had been processed by a certain number of cores (nodes), integrated over a period of time obtained from LSF dataset. The process of creation of these equations as well as notations are explained below:

1. Power used by the cluster will be integrated over the time of monitoring reported in the Zabbix dataset, i.e. over approximately one hour, and extrapolated to exactly fit one-hour period, resulting in the variable $\mathrm{E_j}$, denoting the number of watt-hours of energy consumed by the cluster during the hour j.
    a. LSF data allows to estimate the power used by an arbitrary process within every hour of the cluster activity. Consider the variables $c_{i,j}$ standing for the number of cores required to work on application i during the hour j, and $t^0_{i,j}$, $t^1_{i,j}$ – start and end time of the process i activity during the hour j. These variables can be devised from the monitored data.
    b. For the purpose of obtaining power value in the workload part of the equation, an unknown set of variables is introduced: x_j stands for the power required by arbitrary application every second from one core during the hour j.
    c. Multiplication of $(c_{i,j} \cdot x_j)$ produces the amount of power consumed by application i during the hour j, the integral of this product over the period of this application activity results in energy consumption registered for the application i during the hour j. Finally, summation over the number K of applications which had been active during the hour j provides the estimation of energy required for processing the applications during the hour under consideration.
2. Equation (1) is then transformed discrete format to avoid integration over non-continuous variable and is rewritten as follows:

$$\sum_{i=1}^{K} c_{i,j} \cdot \frac{t_{i,j}}{3600} \cdot x_j = E_j, \quad j = 1, \ldots, N, \tag{2}$$

   where, $t_{i,j}$ represents the duration of job i processing in seconds and is divided by the number of seconds in one hour. Equation (2) forms a sequence of linear equations with one unknown variable x_j for each equation from the sequence. The equations can be resolved by simple division of the right part over the sum from the left part.

3. Having resolved unknown variable x_j, one can estimate monthly energy consumption of the cluster using summation of E_j over j from 1 to N, or energy consumption by individual processes and that of other different granularity levels. Once calculated, these results will contribute to achieving the research objective **RO1.1**.

4.2.2 Quantitative Analysis of IT Jobs Energy Efficiency

1. Once the power required by any process from one core within each hour x_j is obtained, the processes can be categorised by their productivity for the end-users as described in Section III.D [85] and summarised below in subsection 4.2.2.1.
2. Using markers for processes that resulted in energy waste w, and Kronecker delta $\delta_i^w = 1,\ if\ \ i \in w; 0\ -\ otherwise$, where w – ID of processes that finished with errors, it is possible to estimate E_w – monthly energy waste as a sum of hourly energy waste E_{wj}:

$$E_w = \sum_{j=1}^{N} E_{wj} \tag{3}$$

$$E_{wj} = \sum_{i=1}^{K} c_{i,j} \cdot t_{i,j} \cdot x_j \cdot \delta_i^w\ ,\ j = 1,\ \dots, N \tag{4}$$

3. Monthly *energy waste* evaluation allows to estimate DCeP and EWR, energy productivity metrics, the overview and discussion on which is included in subsection 4.2.2.2.
4. Quantitative analysis concludes with interpretation of energy use in terms of indirect carbon emissions using carbon factor for Italy, c = 0.343 tCO_2/MWh [89]:

$$CO_2e = E \cdot c \tag{5}$$

   where E is energy in MWh and CO_2e stands for the amount of carbon emissions (tCO_2). Application of this formula will contribute to investigation on research objective **RO1.2**.

In addition, energy efficiency analysis of cluster queues and groups of parallel and serial jobs is explained in [86], [88].

4.2.2.1 Categorisation of jobs by productivity or energy waste type

All jobs that are marked by LSF as correctly finished processes are assigned to *useful work*. Furthermore, three categories of jobs that have caused energy waste are distinguished: (I) jobs that run for too short time (≤ 30 seconds); (II) jobs that exceed their queues' maximum allowed running time; (III) jobs that end with errors for other unknown reasons.

In case (I), the time is so short that it can only cover job scheduler activities and such jobs cannot bring about useful results to end-users. Even if the jobs are reported to have finished without errors, they should be marked with *energy waste* sign, therefore, this category reveals additional energy waste that is not registered as such by LSF. The category (II) comprises jobs that continue running after the queue maximum time is exceeded and are by default marked as erroneous jobs. However, within the queue maximum time limit these jobs have fed into useful processing results for the end-users and only the part that exceeded the queue time limit has caused *energy waste*. Hence, the second category splits jobs that automatically got an "ending with error" status to *useful work* and *energy waste* depending on the jobs' processing time. The third (III) category dumps all other jobs that ended with errors for unknown reasons and should be analysed by DC operators more thoroughly.

4.2.2.2 Energy Productivity Metrics, DCeP and EWR

As mentioned in previous sections, the DC energy consumption has increased dramatically over the last decade, and this situation has determined the quest for metrics that evaluate DC energy efficiency. Despite a great interest, traditional metrics for measuring energy efficiency in DC (e.g., PUE) are limited to calculating the energy required for the major IT components of the DC plus the energy for supporting infrastructure. In contrast, the present part aims to compute energy efficiency metrics based on a clear definition of the *useful work* which is a parameter intended to gauge the real computing carried out by a DC (**RO1.1**).

In the research that approaches *energy waste* with the use of metrics similarly to this work, useful energy, as the opposite to the *energy waste,* might have an ambiguous definition. *Useful work* depends on the type of application that it characterised and vary from the number of floating-point operations to number of transactions, network traffic or other measurable output [33], [51]. Therefore, unified approach to *useful work* assessment is highly needed.

The analysis on productivity metrics related to the *useful work*, despite the ambiguity of this term, is necessary to achieve sustainability goals. Generally, *useful work* of a DC can be represented by overall computing activity of the IT Equipment. The ITE activity comprises computing, storing and transferring data and is referred to as IT services. Appropriate productivity metrics are used to measure and assess such activity's characteristics [34], [37]. Nevertheless, productivity metrics differ in their approach to assess useful work. As a consequence, none of the metrics has provided a practical way to exactly calculate the work done or *useful work*, even though several attempts have been made to define the productivity metrics for DCs. Among all the productivity metrics, DCeP (Data Center Energy Productivity) is the most significant one [9] and is calculated as follows:

$$DCeP = \frac{Useful\ Work\ Produced}{Total\ DC\ Energy\ Consumed\ over\ Time} \tag{6}$$

The present stage of work is devoted to calculating it based on the DC operation data. DCeP metric evaluation is facilitated by the consideration of each IT job power consumption per core during each second obtained from Eq. 2 and the information about the fulfilment status of jobs. As described in step 2 of subsection 4.2.2, jobs that have caused energy waste are marked with Kronecker delta, all the other jobs have finished successfully and are assumed to contribute to the *useful work*. Thus, monthly energy consumption with separation on the energy for *useful work* and *energy waste* is obtained for each month of the investigated period. Further, DCeP is evaluated as the ratio of energy for *useful work* over the total energy consumption as mentioned in the Eq. 6. As *Total DC Energy Consumed Over Time* is not available in this work, the denominator contains total IT energy used by DC, which will indicate more precisely which part of IT energy is consumed by useful work.

Energy waste assessment has been addressed in academia and industry both qualitatively and quantitatively, for the reason that inefficient energy use causes increased electricity cost and negative environmental impact if the extra energy used is produced from non-renewable resources. Unlike *useful work*, *energy waste* is related to the energy that has been used for computing activities but has not produced results for the end user. Therefore, *energy waste* and energy spent on *useful* computational *work* are two supplementary portions together forming the total cluster ITE energy consumption. For this reason, EWR (Energy Waste Ratio) metric [35], [90] which is equal to (1 – DCeP) or can be otherwise expressed as in Eq. 7, is studied for individual applications and energy waste categories.

$$EWR = \frac{Energy\ Wasted\ for\ not\ Useful\ Work}{Total\ DC\ Energy\ Consumed\ over\ Time} \tag{7}$$

Similar to DCeP, the metric EWR shows what portion of energy has been wasted on jobs that have not resulted in any useful work of cluster processing activities. Since this metric estimates how much energy the cluster uses in vain, measures taken to minimise this value should result in rise of DC IT productivity, therefore, once the metric components are clearly defined, it is useful for DC energy efficiency analysis.

4.3 Results and Discussion

Energy distribution between tasks occupying the reported DC cluster is illustrated in Fig. 1 in section IV of [85]. In detail, the figure shows the proportion in which energy is consumed by different processes over the overall period of monitoring. In the meantime, it indicates the purposes of cluster computations: the variety of applications observed to reside on the cluster is typical for a data center which is adapted for smart city purposes. The variety spans from air quality monitoring, climate modelling, initial versions of smart home and other urban applications to Monte Carlo algorithms for particle physics simulations.

Out of all processes, statistical Monte Carlo methods for particle detection, transport and nuclear fusion are registered to have the highest energy demand and consume 35% of energy over the whole observed period of 11 months. The second group of applications is responsible for 23% of the cluster energy consumption includes air quality simulation and forecast. Other applications individually do not require more than 6% of cluster resource use, while the smallest considerable portion of energy is dedicated to genetic analysis and mathematical algorithm for turbulent flows simulation. Applications with less than 1% energy use over the period under consideration (11 months) have been combined into one group. Given that this group necessitates 16% of total cluster energy, the cluster utilisation pattern is visible: it processes a large number of applications with low energy demands.

Monthly energy consumption has been calculated as described in subsection 4.2.1. The results are represented in Fig. 4 along with DCeP metric evaluation. The largest portion of energy use is observed during the month from 19th of March to 19th of April reaching the point of 35.6 MWh, whereas the smallest portion of energy consumption is reported in the months from 19th of July to 19th of September. DCeP varies from the minimum of 0.61 in the last reported month to 0.84 in the June-July period. In the sense of sustainable resource utilisation, these findings bring evidence that around 60-80% of all energy is consumed by IT equipment to produce *useful work*, a ratio that could be improved with some practices that will be discussed in conclusion and concern users alerts and better load scheduling.

As a note on data analysis strategy, in the case when jobs are taken directly from the LSF data, without categorisation and identification of additional categories (I) and (II) from subsection 4.2.2.1, DCeP is reported to stay at a lower level than after preprocessing the LSF dataset and extracting categories. DCeP differences can be observed in Fig. 4 *(b)*, where no categorisation has been done, versus Fig. 4 *(a)* depicting values when the categorisation has been considered. The reason for such differences stem from the fact that in the raw LSF dataset *useful work* performed by jobs that exceeded queue maximum time have been hidden by the marker of erroneous job for the full period of jobs execution. In addition. However, as described previously, the energy used within the queue time had been spent on *useful work* and only the remaining part of processing period caused energy waste. Also, some short jobs have been marked as *useful work* which does not agree with our assumptions. Henceforward, the categorised dataset is used, i.e. the one corresponding to Fig. 4 *(a)*.

In addition, energy consumption of the processes is found to have been unevenly distributed. The majority of the processes consume less than 100 kWh per month. A more

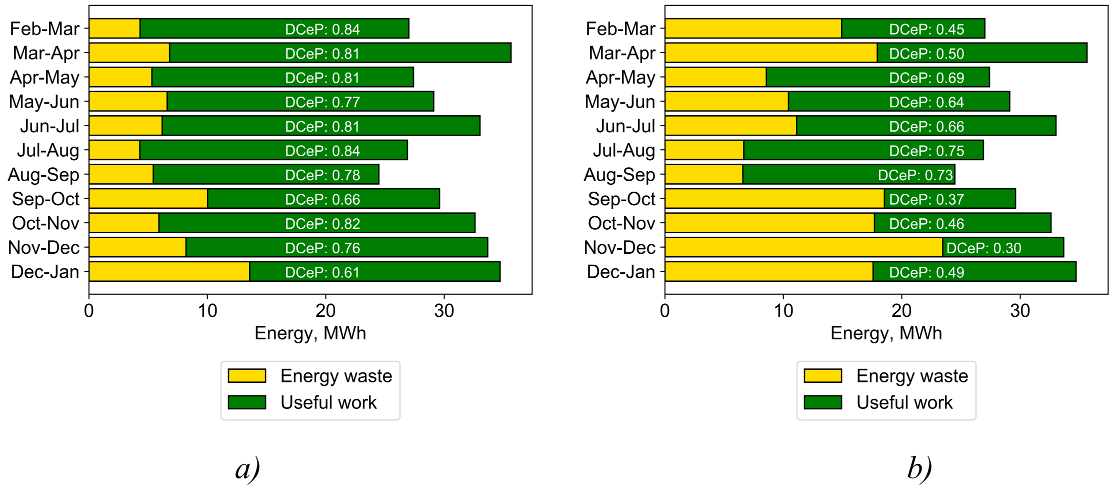

Figure 4. Monthly analysis of energy consumed by correctly finished jobs (useful work) and jobs which exited a queue with an error status (energy waste), and DCeP
a) Energy waste categories are considered;
b) Jobs are not categorised by causes of energy waste, data on jobs status is taken directly from LSF.

granular analysis showed that from 62% to 93% of the overall number of the cluster jobs consume less than 10 kWh per month as shown in Fig. 4 section IV.B in [85].

Energy use by queues and by groups of serial and parallel jobs is studied and reported in in-press works [86], [88] that are available upon request. The main findings are, however, the following. Energy consumption of all 18 queues ranges from 1 kWh to 207 MWh over the total period of monitoring. Number of jobs allocated to each queue reveals no correlation with energy consumed by the queue: for example, the queue with the second smallest energy consumption over the total period and EWR of 16% is reported to have had the most significant number of job allocations. The ratio of 99% of energy is consumed by 9.5% separate submissions, while there is no correlation between energy consumption of a queue and number of jobs submissions. Second smallest energy consumption has been detected in the queue with the highest number of job allocations.

The Energy consumption and EWR of parallel jobs generally prevail over serial jobs, while the number of serial jobs submissions is observed to have been higher than parallel jobs submissions in 10 out of 11 months. An even pattern of parallel jobs EWR has a mean value of 22%, whereas the same metric for serial jobs fluctuate between 0.025% – 4%. It is noted that the monitored cluster parallel jobs consume more energy and, if such a job fails, then the energy required for computations until to the failure point is largely wasted in comparison with serial jobs. In addition, serial jobs consume around two times less energy throughout the studied period, although they are submitted 200 times more frequently on average, the value having been dispersed throughout the months from 10 to 1000 times.

Statistical characteristics are taken from the monthly samples of data and are shown in Table 1. The table includes the minimum, maximum, mean value and standard deviation of the ratios of energy used by jobs from each category related to the general energy use. As might be observed from Table 1, processes with short running time consume the least share of energy (i.e. approximately 0.03%), whereas jobs which exceed the queue time used around 0.2% of total energy consumption. A considerable amount of jobs which are only processed by the scheduler and have a maximum running time of 30 seconds (category I)) represent from 14 to 56% of all submitted jobs throughout the whole period of investigation, Table 1. On the contrary, jobs, which exceed the queue time limit, form less than 1% during the majority of reported period.

Table 1. Energy Waste Ratio by Job Categories with Relation to Overall Energy, %

Statistical Characteristics	a) Running time ≤ 30 sec	b) Running time > queue time	c) Other reasons
Min	0.007	0.004	16
Max	0.06	0.3	39
Mean	0.03	0.2	23
Standard Deviation	0.01	0.09	7

To summarize, results obtained through the assessment of useful work and energy waste reveal the energy consumption patterns within the cluster. Firstly, the least energy is consumed during the summer months of annual vacations, whereas the most significant amount of wasted energy is observable in December-January when users might have worked remotely during the Christmas holidays. Secondly, a high percentage of jobs consume less than 10 kWh per month, which result in the energy spent on minor jobs rather than resource-hungry processes. Also, the cluster wastes most of the energy for jobs which end with errors for unknown reasons that require further examination. Regarding the energy waste categories, some jobs that are only preprocessed by the scheduler and do not provide any results, is considerably higher than the number of jobs removed from the queue because of the time limit conflicts.

The results in terms of energy consumption required for useful work and processes without positive results for the end user have been translated into CO_2 or equivalent greenhouse gases (GHG) emissions to show the environmental impact of the cluster's processing work. Fig. 5 shows the amount of carbon emissions produced by the computation facilities during the cluster processing. As a basis for this figure, the monthly energy consumption is calculated for all the jobs, which are successfully completed, and the jobs, which end up with errors. The values are converted to MWh and then multiplied by the carbon factor. Evaluation of CO_2 or equivalent GHG emitted only by IT equipment does not facilitate the assessment of CUE (Carbon Usage Effectiveness) metric for the cluster, because it requires data on total emissions caused both by IT equipment and supportive infrastructure of the DC. Thus, by analogy with EWR, we propose to use Carbon Waste Ratio, CWR, to express the same value in terms of CO_2 emissions as in Eq. 8. CWR metric and

$$CWR = \frac{Carbon\ Emissions\ Introduced\ by\ Unproductive\ Computations}{Total\ DC\ Carbon\ Emissions\ over\ Time}. \tag{8}$$

The overall CO_2 emissions fluctuate between 8 and 12.2 tonnes CO_2 per month. The proportion of CO_2 emissions caused by energy waste ranges from 16% to 40% of monthly emissions (CWR value in %). Fig. 5 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. These results meet the target **RO1.2**.

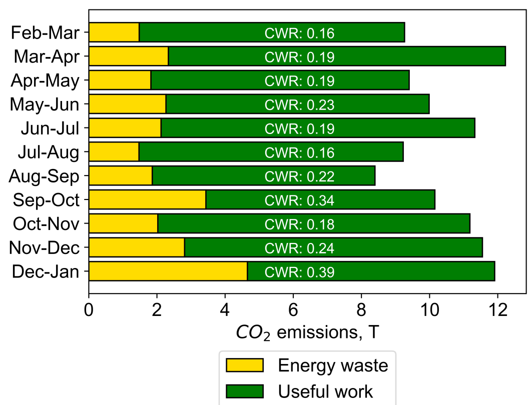

Figure 5. Monthly CO_2 (or equivalent) emissions caused by jobs which ended with errors and correctly finished jobs, CWR.

The conducted analysis provides a more in-depth insight into the useful work performed by the cluster IT equipment and, at the same time, waste energy. An incremental contribution towards a better understanding of DC sustainability has been presented in terms of carbon emissions for useful work and jobs associated with energy waste. However, the data available in the case study is not sufficient for the evaluation of any carbon or sustainability metric, therefore a new metric has been proposed.

4.4 Phase 1 Conclusion

Assessment of IT equipment energy efficiency has been addressed in this part of the work through evaluation of useful work and energy waste generated by IT jobs processing. Ambiguity of useful work definition has been overcome with the help of IT jobs markers of successful or faulty completion of every job assigned by Load Sharing Facility and enhanced by further categorisation of jobs that caused energy waste based on jobs individual characteristics. Investigation on energy efficiency pursues the goal of improving cluster's sustainability, since it helps identify weak points in cluster jobs scheduling processing so that DC operators and users are motivated to act appropriately and increase energy efficiency, reaching higher productivity and larger amounts of useful work with less energy consumption.

The study goes beyond energy efficiency and translates energy utilisation values to associated carbon emissions of the cluster. With the aim to quantify energy use in terms of environmental burden imposed by the DC, these results increase awareness of DC stakeholders about the severity of IT equipment ecological impact and the urgency of the need to improve energy efficiency.

Raising a question of transferability of the current findings, methodology proposed in this part of the work applicable for evaluation of energy use of the same DC cluster during its future operation. Moreover, general technique to apply energy conservation law for the study of energy distribution between applications is transferable to other systems with different patterns of monitored data. Use of carbon factor proposed to assess carbon emissions caused by the cluster is an approach that gives a rough estimate of associated carbon footprint but is widely applicable thanks to its simplicity. The carbon metric CWR is based on analogous EWR index and can complement the latter for a more rigorous assessment of energy-related impact. They both require distinction between energy waste and useful work, which has been discussed in this section. Since the final goal of separating jobs in these two categories, useful work and Not Useful Work, has the final goal of evaluation of energy by these groups, it is sufficient to assume some jobs provide computational results to the end user and then directly evaluate energy consumption of such jobs rather than, for example, interpreting useful work in terms of floating-point operations and translating it into the energy characteristics. Therefore, suggested methodology allows to avoid one extra step in useful work-related energy consumption in the settings where resultant jobs can be marked as such.

This part of the work concludes with energy efficiency-oriented recommendations based on the findings concerning the cluster IT equipment operation to address the objective **RO1.3**.

4.4.1 Recommendations for DC IT Jobs Energy Efficiency Enhancement

REC 1. Improve scheduling policies

REC 1.1. Currently utilised FCFS queuing algorithm could be replaced by other algorithms (appropriate under different circumstances), for example, Largest Job First, to optimise the system load, or Smallest Job First, to optimise the throughput, or other algorithms.

Introduce/enhance priorities and add backfilling approaches if necessary: reorder jobs to match the availability of resources and tasks priority.

REC 2. Enhance task resource allocation strategy

REC 2.1. Consider energy usage optimisation queuing strategy rather than to the currently employed chronological order-based strategy. According to the chronological order-based queuing the first job that enters the system is allocated the first available queue with required characteristics. By contrast, energy usage optimisation could foster the choice of the queue with minimal energy consumption for a specific job.

REC 3. Apply best practices for general energy efficiency

REC 3.1. Avoid overprovisioning: provision only the required IT power usage (guideline 4.1.9 from [8]), shut down and remove idle equipment (guideline 4.3.6 from

[8]), apply Dynamic Voltage and Frequency Scaling whenever possible; consolidate servers when needed (guideline 4.3.4 from [8]).

REC 3.2. Improve monitoring system to separately take power measurements from IT components such as different-purpose servers and PDUs.

REC 3.3. Review load characteristics and monitoring system.

REC 4. Alert and inform end users about optimal utilisation of the cluster.

REC 4.1. Jobs not optimised for parallelism must not be submitted to parallel queues as they can cause large energy waste. For example, submission of a job requiring only one core to a 24-core queue will cause idle power consumption of 23 remaining cores.

REC 4.2. Jobs must be well-designed and tested prior to their submission to the cluster queues; resubmission of faulty jobs should be avoided to minimise energy waste.

REC 5. Raise environmental awareness of the DC by auditing the energy consumption of existing equipment.

REC 5.1. Identify the degree of DC IT equipment compliance to Energy Star specifications.

REC 5.2. Compare monitored power or energy consumption values with technical specifications to determine if any equipment consumes extra power and investigate its underlying reasons.

REC 5.3. Regularly evaluate a cluster energy consumption and apply performance and productivity metrics for cluster energy efficiency assessment. Include an analysis of carbon emissions into a regular cluster evaluation to determine its environmental impact.

REC 5.4. Consider an integration of free cooling in cold months and renewable energy use.

5 SUSTAINABILITY ANALYSIS

This work covers energy efficiency analysis of data center operations. Energy efficiency is a driver for DC sustainability that has been covered from the point of IT equipment energy productivity, associated carbon emissions and thermal characteristics of an IT room to identify possible areas for improvement in terms of recommendations for a real DC under consideration. Optimised and reduced energy consumption could contribute to the global environment as it would imply emitting less indirect carbon emissions.

For a more detailed analysis of this work, following the model proposed in [91] with guidelines for its use in [92], sustainability contributions are studied from five perspectives: technical, economic, environmental, social, and individual. These five dimensions are further sorted by immediate, enabling (long-term) and structural (global cumulative) orders of effects that the work has on each of the dimensions. Key effects are depicted in Fig. 6.

As shown in Fig. 6, immediate effects originate from improved energy efficiency of a DC and comprise energy consumption (and thus, reduced electricity expenditure) and physical waste reduction as a result of quantitative analysis of IT jobs and thermal operation efficiency. The proposed framework for DC energy efficiency analysis is flexible and provides environmental benefits of reduced carbon emissions and fewer risks for air pollution-related health diseases as indicated in Fig. 6. Every dimension is considered in detail in the remainder of this chapter.

Environmental dimension

Suggested sets of recommendations have a principal goal of ***reducing energy waste*** generated by IT systems and thermal equipment. Analysis and recommendations targeting at energy efficiency reveal possibilities for ***better resource utilisation***. Moreover, optimal thermal conditions slow down deterioration of physical equipment and thus ***reduce the rate of physical waste generation***.

Since the DC under consideration does not provide information on utilisation of renewable

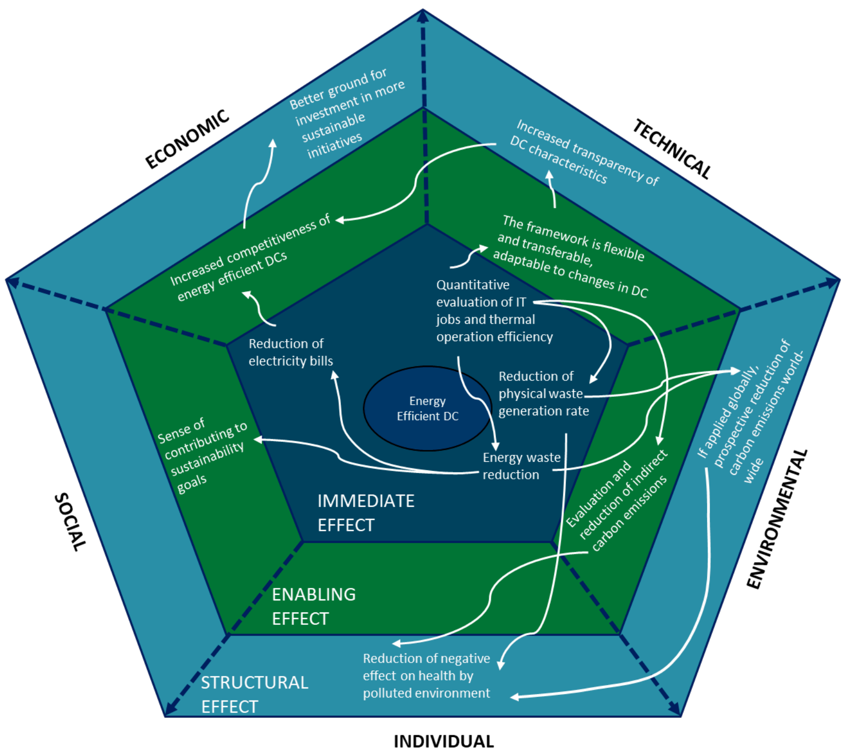

Figure 6. Sustainability analysis of the work

energy, ***carbon emissions*** of large-scale energy consumption might exacerbate the environmental situation to a smaller or larger extent. Moving towards energy efficient DCs could be regarded to as one measure in favour of their sustainable operation. In this regard, a ***Carbon Waste Ratio metric is proposed*** as an index that estimates the ratio of carbon emissions generated by futile IT work. Quantitative evaluation of carbon emissions could facilitate their reduction and, if applied globally, might ***reduce indirect emissions of DCs on a global scale***.

Addressing the issue of ***environmental cost of this work*** would entail the management and control of costs related to environmental impact of the DC operations. This would encompass changes made to the existing DC monitoring system installation (that has been set up prior to the start of this work) and its deployment. Therefore, the only direct environmental effect that the methodology proposed in this work would be the ecological footprint of data analysis performed on a PC which is negligible compared to the scale of the DC cluster considered here. Quantitative evaluation of the environmental cost of this work requires an in-depth and rigorous audit of the energy consumed by the PC used for this research work and also its life cycle accounting which is beyond the scope of this research work. However, it could be regarded as one of directions for future work.

Technical dimension

The conceptual framework of energy efficiency assessment presented in this work is mainly targeted at DC operators and engineers, i.e. ***prospective users*** who are specialised in DC maintenance and supervision. The framework is ***flexible*** to the changes of DC characteristics and adding/replacing features. For example, if a new metric is assumed beneficial, it can be easily added to a list of discussed indices. Alternatively, with extension of the monitoring system, new features available for supervision and control could be added to the analysis.

Accuracy of the ***results might be affected*** by errors of the monitoring system. Some monitoring inaccuracies have been found and avoided in Phase 2 during evaluation of energy consumption (see Appendix 2). However, some errors might have remained hidden, for example, for thermometers that provided extremely high temperature readings (up to 80°C). This implies that a monitoring system as a basis of this analytical work requires constant maintenance and tuning.

The methodology and results of this work are ***transferable*** to DCs other than the one considered here. For example, mathematical modelling approach described in Phase 1 is specifically developed for the type of CRESCO4 cluster monitoring system output and can be reused as long as this system is utilised. The concepts enabling this approach, i.e. energy conservation law applied to find power consumption of individual processes on a core level, have a potential for adaptability to other DCs. In addition, this work showcases applicability of thermal and productivity metrics for a real DC. A technical possibility of metrics evaluation could motivate DC operators to apply advanced monitoring and to report the state of their facilities for better control and management that could ***contribute to transparency of DC*** characteristics.

Economic dimension

The findings of this work contribute to DC operators in a way that they could potentially ***reduce electricity bills*** as well as prolong devices lifetime compared to business-as-usual strategies. In addition, a part of DC recommendations suggests notifying users in case of repetitive failures of their jobs so that they are motivated to submit good-quality code. ***Users' activity*** could help further reduce energy waste and extra electricity spending.

Immediate economic, technical and environmental effects of reducing energy and physical waste within a DC as well as electricity consumption make a DC ***more competitive*** on a global market of cloud and high-performance computations. Increased competitiveness could help ***attract more investment*** for further sustainable initiatives as a cumulative effect.

Social dimension

Users and employees of an energy efficient DC could benefit from a ***sense of contributing to global sustainability goals***. Inclusion in the EU CoC initiatives of applying best practices and guidelines for DC energy efficiency is a rewarding idea that could ***unite people*** concerned about the environment.

Individual dimension

Reduced carbon emissions as a result of lower electricity demand by an optimised DC improves regional environmental conditions and positively affects humans' ***health*** that leads to increased quality of life.

6 CONCLUSION

This chapter provides a summary of findings for this work, clarifies a role of a DC in smart cities as it has been considered in this work, and mentions suggestions for the future work.

6.1 Summary of Findings

Data centers play a crucial role in smart cities as enablers of urban applications, as summarised in Chapters 1 and 2. They provide scalable on-demand computing and storage resources to varying load of smart city transportation, governmental services, ecological monitoring, smart home and other applications that enhance residents' quality of life. Meanwhile, DCs are large electricity power consumers which impose certain challenges for a DC and a smart city that tends to accomplish sustainability goals. Indirect carbon emissions, generated through the process of electrical energy production that is later utilised by DCs, should be outweighed by positive environmental effects of smart city applications empowered by DCs. In this scenario, sustainability of a DC represents an essential goal for a large high-performance computing facility to operate within a smart city. In terms of DC operations, sustainability could be interpreted in a number of ways, and this work primarily focuses on energy efficiency.

Two major DC aspects have been considered in this work: IT equipment energy productivity and thermal characteristics of an IT room. The findings of this work are based on analysis of available monitoring data characterising two clusters of ENEA Portici DC, CRESCO4 and CRECSO6. Three phases of analysis have unravelled possible improvements for thermal design and load management with overall methodology of this work covered in Chapter 3.

In the first phase (Chapter 4), a question of IT jobs energy productivity has been raised. Mathematical modelling as well as energy efficiency metrics evaluation have been employed to investigate how effectively energy is used by IT equipment of CRESCO4 cluster to produce useful processing work. Not useful work or energy waste has been categorised based on the reasons for which IT jobs failed to produce results to the end users. The phase covered Energy Waste Ratio and Data Center energy Productivity metrics assessment and a proposal of a new metric Carbon Waste Ratio to better assess the portion of indirect carbon emissions generated during cluster processing activities that resulted in no useful work. Recommendations for the cluster IT equipment energy efficiency improvement encompassed better queue policies and resource allocation, users' notification about repeated failures of their jobs' submissions and auditing the IT equipment for better assessment of its environmental impact.

The second phase (Appendix 2) has been dedicated to statistical analysis and metrics evaluation of the CRESCO6 cluster IT room thermal conditions. Analysis of available real data obtained from temperature sensors has provided insights into thermal design pitfalls such as bypass and hotspots observed in the IT room. Statistical assessment of temperature ranges has shown that servers are overheated inside the rack. To address these issues a set of recommendations has been provided including air flow speed adjustment, better isolation of underfloor plenum, extension of the monitoring system to humidity and temperature sensors in various locations as well as regular maintenance and calibration of the monitoring equipment.

In the third phase (Appendix 3), a clustering technique has been employed to localise the hotspots identified in the second phase. Using the same dataset concerning CRECO6 IT room temperature measurements, sequential clustering has been performed to group nodes by thermal ranges in which they have resided most frequently during the period of observations. The ratio of 8% of all servers has been most frequently observed in the hot temperature range. Several measures to combat an issue of hotspots have been recommended concerning directional cooling, load management, and continuous monitoring of the IT room thermal conditions.

6.2 Emerging Challenges

The outcome of this work could not have been possible without having overcome some data analytics challenges in each phase. Firstly, the choice of granularity for the analysis being performed is a separate task. For a global level of a DC, it has been considered that monthly statistics would be sufficiently representative and useful for DC providers.

Furthermore, the datasets provided for Phase 1 have been challenging in two ways. They had uneven timestamps, therefore, additional methods were necessary to work with both datasets simultaneously. The most important challenge in Phase 1 was that the load scheduling dataset did not contain detailed information about which core IDs were used for every job execution. Instead, a number of all active cores was mentioned. Thus, the main result of mathematical modelling used in Phase 1 was to obtain average estimation of power use of every job per any arbitrary core. In the meantime, some cores were loaded more than other but this information could not be retrieved due to the mentioned challenge of Phase 1.

Phase 2 was initially provided with two datasets, one of which was not suitable for computations as more than 40% of data was missing. This challenge shaped the work into analysis of thermal characteristics exclusively. In addition, Phase 3 required sequential clustering as opposed to regular clustering or other techniques. For the type of data provided in Phase 3, a usual research objective would be prediction based on time series analysis. By contrast, the objective of Phase 3 was to obtain a group of frequently overheated nodes that required clustering, and since the measurements were repetitive, the problem transformed into sequential clustering which is rarely used and scarcely described in the literature. In this part of the work the most challenging task was to combine a sequence of results and reach a final outcome.

6.3 Future Work

Some suggestions for the future work could be made based on this thesis:

- A tool for automation of metrics evaluation based on the methods provided in this research could be useful for a DC. Automation of monitoring and immediate control of DC processes in case of undesirable thermal conditions or IT jobs failure could decrease the latency of DC operators' reaction to phenomena that reduce energy efficiency.
- For a better overview of the DC under consideration, energy efficiency analysis of the thermal equipment would be fruitful. So far, thermal analysis considered only the IT room characteristics, i.e. the output of the thermal equipment. However, the CRAC unit is a separate energy consumer and its efficiency could influence total DC energy consumption. In addition, monitoring of the CRAC power use as well as

other sources of the DC electrical consumption could enable PUE and other metrics evaluation and provide an opportunity to participate in EU CoC assessment.

- Distribution of the temperature ranges within the DC is a topic for further research both in the field of statistical analysis and fluid dynamics. A model of the temperature distribution could empower DCs to prevent and mitigate overheating in the areas suggested by such a model.
- To date, a limited number of approaches exist to address an issue of DC hotspots. Future research could be dedicated to preventing and cooling the hotspots;
- A further investigation into Machine Learning techniques for thermal data analysis could be fruitful. In particular, results of different ML methods applied to the thermal data could be compared and constitute a theoretical study of the effectiveness of ML techniques for the given problem. Cross-validation of the range of generated machine learning models will be deemed useful;
- Assessment of the environmental cost of this work could be important for DCs to evaluate added carbon emissions by monitoring systems' use and analysis of their data. A methodology could be proposed to evaluate environmental effects of changes brought to a DC during refurbishing and upgrades aimed at compliance with sustainability best practices;
- This work has shown that data centers and smart cities are closely related. This work could form the basis of a new research question "Why do smart cities need smarter and more sustainable data centers?", which could call for further consideration.

REFERENCES

[1] A. L. Kor, E. Rondeau, K. Andersson, J. Porras, and J.-P. Georges, “Education in Green ICT and Control of Smart Systems : A First Hand Experience from the International PERCCOM Masters Programme,” in *12th Symposium on Advances in Control Education (ACE 2019)*.

[2] D. of E. and S. A. United Nations, *World Urbanization Prospects: The 2018 Revision*. New York, New York, USA: United Nations, 2018.

[3] P. Neirotti, A. De Marco, A. C. Cagliano, G. Mangano, and F. Scorrano, “Current trends in Smart City initiatives: Some stylised facts,” *Cities*, vol. 38, pp. 25–36, Jun. 2014.

[4] C. Lim, K.-J. Kim, and P. P. Maglio, “Smart cities with big data: Reference models, challenges, and considerations,” *Cities*, vol. 82, pp. 86–99, Dec. 2018.

[5] E. Al Nuaimi, H. Al Neyadi, N. Mohamed, and J. Al-Jaroodi, “Applications of big data to smart cities,” *J. Internet Serv. Appl.*, vol. 6, no. 1, pp. 1–15, 2015.

[6] I. A. T. Hashem *et al.*, “The role of big data in smart city,” *Int. J. Inf. Manage.*, vol. 36, no. 5, pp. 748–758, Oct. 2016.

[7] M. Dayarathna, Y. Wen, and R. Fan, “Data Center Energy Consumption Modeling: A Survey,” *IEEE Commun. Surv. Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.

[8] M. Acton, P. Bertoldi, J. Booth, L. Newcombe, A. Rouyer, and R. Tozer, “2019 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency,” no. v 10.1.0, pp. 1–48, 2018.

[9] V. D. Reddy, B. Setz, G. S. V. R. K. Rao, G. R. Gangadharan, and M. Aiello, “Metrics for Sustainable Data Centers,” *IEEE Trans. Sustain. Comput.*, vol. 2, no. 3, pp. 290–303, Jul. 2017.

[10] X. Jin, F. Zhang, A. V. Vasilakos, and Z. Liu, “Green Data Centers: A Survey, Perspectives, and Future Directions,” *arXiv*, vol. 1608, no. 00687, 2016.

[11] L. Cupertino *et al.*, “Energy-efficient, thermal-aware modeling and simulation of data centers: The CoolEmAll approach and evaluation results,” *Ad Hoc Networks*, vol. 25, pp. 535–553, Feb. 2015.

[12] T. Daim, J. Justice, M. Krampits, M. Letts, G. Subramanian, and M. Thirumalai, “Data center metrics: An energy efficiency model for information technology

managers," *Manag. Environ. Qual. An Int. J.*, vol. 20, no. 6, pp. 712–731, 2009.

[13] B. Whitehead, D. Andrews, A. Shah, and G. Maidment, "Assessing the environmental impact of data centres part 1: Background, energy use and metrics," *Build. Environ.*, vol. 82, pp. 151–159, Dec. 2014.

[14] M. Chinnici, A. Capozzoli, and G. Serale, "Chapter 10 - Measuring energy efficiency in data centers," in *Pervasive Computing*, 2016, pp. 299–351.

[15] A. Quintiliani and M. Chinnici, "DC4Cities. An environmentally sustainable DC for Smart Cities. Final DC4Cities standardization framework and results description of the European Cluster," Rome, Italy, 2016.

[16] I. A. T. Hashem *et al.*, "The role of big data in smart city," *Int. J. Inf. Manage.*, vol. 36, no. 5, pp. 748–758, Oct. 2016.

[17] S. E. Bibri and J. Krogstie, "Smart sustainable cities of the future: An extensive interdisciplinary literature review," *Sustainable Cities and Society*. 2017.

[18] E. P. Trindade, M. P. F. Hinnig, E. M. da Costa, J. S. Marques, R. C. Bastos, and T. Yigitcanlar, "Sustainable development of smart cities: a systematic review of the literature," *J. Open Innov. Technol. Mark. Complex.*, vol. 3, no. 1, p. 11, Dec. 2017.

[19] Cloud Council, "Deploying Big Data Analytics Applications to the Cloud : Roadmap for Success," *Cloud Stand. Cust. Counc.*, 2014.

[20] Z. Allam and Z. A. Dhunny, "On big data, artificial intelligence and smart cities," *Cities*, vol. 89, pp. 80–91, Jun. 2019.

[21] T. Cioara *et al.*, "Exploiting data centres energy flexibility in smart cities: Business scenarios," *Inf. Sci. (Ny).*, vol. 476, pp. 392–412, Feb. 2019.

[22] M. Wahlroos, M. Pärssinen, J. Manner, and S. Syri, "Utilizing data center waste heat in district heating – Impacts on energy efficiency and prospects for low-temperature district heating networks," *Energy*, vol. 140, pp. 1228–1238, Dec. 2017.

[23] M. Antal, T. Cioara, I. Anghel, C. Pop, and I. Salomie, "Transforming Data Centers in Active Thermal Energy Players in Nearby Neighborhoods," *Sustainability*, vol. 10, no. 4, p. 939, Mar. 2018.

[24] R. Yesner and F. Ozdemir, "Understanding Smart City Transformation with Best Practices," *Paper, White*, no. November. pp. 1–24, 2017.

[25] C. Manville, R. Europe, J. Millard, D. Technological Institute, A. Liebe, and R. Massink, "Mapping Smart cities in the EU," 2017.

[26] N. S. Godbole and J. Lamb, “Using data science & big data analytics to make healthcare green,” *2015 12th Int. Conf. Expo Emerg. Technol. a Smarter World*, pp. 1–6, 2015.

[27] R. A. Alshawish, S. A. M. Alfagih, and M. S. Musbah, “Big data applications in smart cities,” in *2016 International Conference on Engineering & MIS (ICEMIS)*, 2016, pp. 1–7.

[28] S. E. Bibri, “The IoT for smart sustainable cities of the future: An analytical framework for sensor-based big data applications for environmental sustainability,” *Sustain. Cities Soc.*, vol. 38, pp. 230–253, Apr. 2018.

[29] Global eSustainability Initiative (GeSI), “SMARTer 2020: the role of ICT in driving a sustainable future,” pp. 1–243, 2012.

[30] Global e-Sustainability Initiative GeSI, “#SMARTer2030-ICT Solutions for 21st Century Challenges,” pp. 1–8, 2015.

[31] Koomey, “Growth in Data Center Electricity use 2005 to 2010,” *Oakland, CA Anal. Press. August*, vol. 034008, pp. 1–24, 2011.

[32] M. Avgerinou, P. Bertoldi, and L. Castellazzi, “Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency,” *Energies*, vol. 10, no. 10, 2017.

[33] V. D. Reddy, B. Setz, G. S. V. R. K. Rao, G. R. Gangadharan, and M. Aiello, “Metrics for Sustainable Data Centers,” *IEEE Trans. Sustain. Comput.*, vol. 2, no. 3, pp. 290–303, Jul. 2017.

[34] M. Chinnici, A. Capozzoli, and G. Serale, “Measuring energy efficiency in data centers,” in *Pervasive Computing: Next Generation Platforms for Intelligent Data Collection*, 2016, pp. 299–351.

[35] M. Chinnici and A. Quintiliani, “An Example of Methodology to Assess Energy Efficiency Improvements in Datacenters,” in *2013 International Conference on Cloud and Green Computing*, 2013, pp. 459–463.

[36] M. Chinnici, D. De Chiara, and A. Quintiliani, “Data Center, a Cyber-Physical System: Improving Energy Efficiency Through the Power Management,” in *2017 IEEE 15th Intl Conf on Dependable, Autonomic and Secure Computing, 15th Intl Conf on Pervasive Intelligence and Computing, 3rd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology*

Congress(DASC/PiCom/DataCom/CyberSciTech), 2017, pp. 269–272.

[37] A. Quintiliani, M. Chinnici, and D. De Chiara, "Understanding 'workload-related' metrics for energy efficiency in Data Center," in *2016 20th International Conference on System Theory, Control and Computing (ICSTCC)*, 2016, pp. 830–837.

[38] A. Capozzoli, G. Serale, L. Liuzzo, and M. Chinnici, "Thermal Metrics for Data Centers: A Critical Review," *Energy Procedia*, vol. 62, pp. 391–400, Jan. 2014.

[39] A. Capozzoli, M. Chinnici, M. Perino, and G. Serale, "Review on performance metrics for energy efficiency in data center: The role of thermal management," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 8945, pp. 135–151, 2015.

[40] J. Yuventi and R. Mehdizadeh, "A critical analysis of Power Usage Effectiveness and its use in communicating data center energy consumption," *Energy Build.*, vol. 64, pp. 90–94, 2013.

[41] N. Horner and I. Azevedo, "Power usage effectiveness in data centers: Overloaded and underachieving," *Electr. J.*, 2016.

[42] K. Zhang, Y. Zhang, J. Liu, and X. Niu, "Recent advancements on thermal management and evaluation for data centers," *Appl. Therm. Eng.*, vol. 142, pp. 215–231, Sep. 2018.

[43] A. M. S. Osman, "A novel big data analytics framework for smart cities," *Futur. Gener. Comput. Syst.*, vol. 91, pp. 620–633, Feb. 2019.

[44] M. A. Khan, M. F. Uddin, and N. Gupta, "Seven V's of Big Data understanding Big Data to extract value," in *Proceedings of the 2014 Zone 1 Conference of the American Society for Engineering Education*, 2014, pp. 1–5.

[45] A. Gandomi and M. Haider, "Beyond the hype: Big data concepts, methods, and analytics," *Int. J. Inf. Manage.*, vol. 35, no. 2, pp. 137–144, Apr. 2015.

[46] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renew. Sustain. Energy Rev.*, vol. 31, pp. 622–638, 2014.

[47] G. F. Davies, G. G. Maidment, and R. M. Tozer, "Using data centres for combined heating and cooling: An investigation for London," *Appl. Therm. Eng.*, vol. 94, pp. 296–304, Feb. 2016.

[48] N. Beldiceanu *et al.*, “Towards energy-proportional clouds partially powered by renewable energy,” *Computing*, 2017.

[49] A. Khosravi, S. K. Garg, and R. Buyya, “Energy and carbon-efficient placement of virtual machines in distributed cloud data centers,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2013, vol. 8097 LNCS, pp. 317–328.

[50] Z. Royaee and M. Mohammadi, “Energy aware Virtual Machine Allocation Algorithm in Cloud network,” in *2013 Smart Grid Conference (SGC)*, 2013, pp. 259–263.

[51] L. Cupertino *et al.*, “Energy-efficient, thermal-aware modeling and simulation of data centers: The CoolEmAll approach and evaluation results,” *Ad Hoc Networks*, vol. 25, pp. 535–553, Feb. 2015.

[52] P. Garraghan, I. S. Moreno, P. Townend, and Jie Xu, “An Analysis of Failure-Related Energy Waste in a Large-Scale Cloud Environment,” *IEEE Trans. Emerg. Top. Comput.*, vol. 2, no. 2, pp. 166–180, Jun. 2014.

[53] S. Klingert *et al.*, “DC4Cities. An environmentally sustainable data centre for Smart Cities. First Market Analysis,” Rome, Italy, 2014.

[54] European Commission Information Society, “Smart Data Centers Projects.” [Online]. Available: http://ec.europa.eu/information_society/activities/sustainable_growth/funding/prj_data/index_en.htm. [Accessed: 10-Jan-2019].

[55] T. Schulze *et al.*, “DC4Cities. An environmentally sustainable DC for Smart Cities. Final report on Trials Requirements and Metrics,” 2013.

[56] B. Whitehead, D. Andrews, A. Shah, and G. Maidment, “Assessing the environmental impact of data centres part 2: Building environmental assessment methods and life cycle assessment,” *Build. Environ.*, vol. 93, pp. 395–405, Nov. 2015.

[57] M. Acton, P. Bertoldi, J. Booth, L. Newcombe, A. Rouyer, and R. Tozer, *2018 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency*, vol. EUR 29103. 2018.

[58] Energy Star, “Data Center Equipment.” [Online]. Available: https://www.energystar.gov/products/data_center_equipment. [Accessed: 30-Apr-

2019].

[59] ASHRAE Technical Committee 9.9, "Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance," 2011.

[60] ASHRAE, "Data Center Power Equipment Thermal Guidelines and Best Practices," 2016.

[61] A. M. Ferreira and B. Pernici, "Managing the complex data center environment: an Integrated Energy-aware Framework," *Computing*, vol. 98, no. 7, pp. 709–749, Jul. 2016.

[62] M. K. Herrlin, "Rack cooling effectiveness in data centers and telecom central offices: The Rack Cooling Index (RCI)," *ASHRAE Trans.*, vol. 111, no. Part 2, 2005.

[63] M. K. (Ashrae) Herrlin, "Airflow and Cooling Performance of Data Centers : Two Performance Metrics," *ASHRAE Trans.*, vol. 114, no. Part 2, 2008.

[64] D. M. Rădulescu, C. Z. Rădulescu, and G. Lazaroiu, "New Metrics for the Measurement of Energy Efficiency in Data Centers," in *International Multidisciplinary Scientific GeoConference : SGEM : Surveying Geology & mining Ecology Management*, 2018, vol. 18, no. 2.1, pp. 603–610.

[65] R. Zhou, Y. Shi, and C. Zhu, "AxPUE: Application level metrics for power usage effectiveness in data centers," in *Proceedings - 2013 IEEE International Conference on Big Data, Big Data 2013*, 2013, pp. 110–117.

[66] R. C. Zoie, R. D. Mihaela, and S. Alexandru, "An analysis of the power usage effectiveness metric in data centers," in *Proceedings - 2017 5th International Symposium on Electrical and Electronics Engineering, ISEEE 2017*, 2017, vol. 2017-Decem, pp. 1–6.

[67] A. Shehabi *et al.*, "United States Data Center Energy Usage Report," *Lawrence Berkeley Natl. Lab. Berkeley, CA, Tech. Rep.*, pp. 1–66, 2016.

[68] Gartner, "Best Cloud Infrastructure as a Service, Worldwide of 2018 as Reviewed by Customers," 2018. [Online]. Available: https://www.gartner.com/reviews/customers-choice/public-cloud-iaas. [Accessed: 20-Feb-2019].

[69] Gartner, L. Leong, R. Bala, and D. Smith, "Magic Quadrant for Cloud Infrastructure as a Service, Worldwide," 2018. [Online]. Available:

https://www.gartner.com/en/documents/3875999. [Accessed: 20-Feb-2019].

[70] TechRadar, N. Drake, and B. Turner, "Best cloud computing services of 2019," 2019. [Online]. Available: https://www.techradar.com/news/best-cloud-computing-service. [Accessed: 30-Apr-2019].

[71] G. Cook and E. Jardim, "Guide to Greener Electronics 2 0 1 7," 2017.

[72] The Natural Resources Defense Council (NRDC) and Anthesis, "Data Center Efficiency Assessment," 2014. [Online]. Available: https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf.

[73] A. W. Services, "AWS & Sustainability." [Online]. Available: https://aws.amazon.com/about-aws/sustainability/. [Accessed: 22-Feb-2019].

[74] J. Barr, "Cloud Computing, Server Utilization, & the Environment," 2015. [Online]. Available: https://aws.amazon.com/blogs/aws/cloud-computing-server-utilization-the-environment/. [Accessed: 20-Feb-2019].

[75] Amazon and Day One Staff, "Sustainability by the numbers," 2018. [Online]. Available: https://blog.aboutamazon.com/sustainability/sustainability-by-the-numbers. [Accessed: 20-Feb-2019].

[76] Google, "Google Cloud Renewable Energy." [Online]. Available: https://cloud.google.com/renewable-energy/. [Accessed: 20-Feb-2019].

[77] GreenPeace *et al.*, "Clicking clean: who is winning the race to build a Green Internet?," 2017.

[78] Google, "Google Cloud Sustainability." [Online]. Available: https://cloud.google.com/sustainability/. [Accessed: 22-Feb-2019].

[79] Google, "Efficiency: How we do it," 2019. [Online]. Available: https://www.google.com/about/datacenters/efficiency/internal/#measuring-efficiency. [Accessed: 30-Apr-2019].

[80] Google, "Google 2018 Environmental Report," p. 53, 2018.

[81] J. Gao, "Machine Learning Applications for Data Center Optimization," *Google White Pap.*, pp. 1–13, 2014.

[82] Microsoft and M. Branscombe, "How Microsoft is Keeping Its Cloud More Efficient Than Your Data Center," 2018. [Online]. Available: https://www.datacenterknowledge.com/microsoft/how-microsoft-keeping-its-cloud-more-efficient-your-data-center. [Accessed: 22-Feb-2019].

[83] Microsoft Corporation, “The Carbon Benefits of Cloud Computing,” *Int. J. Technol. Diffus.*, p. 26, 2018.

[84] Microsoft Azure, “Azure Regions.” [Online]. Available: https://azure.microsoft.com/en-us/global-infrastructure/regions/. [Accessed: 30-Apr-2019].

[85] A. Grishina *et al.*, “DC Energy Data Measurement and Analysis for Productivity and Waste Energy Assessment,” in *2018 IEEE International Conference on Computational Science and Engineering (CSE)*, 2018, pp. 1–11.

[86] A. Grishina, M. Chinnici, D. De Chiara, E. Rondeau, and A. L. Kor, “Energy-Oriented Analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center,” Applied Ph., K. Ntalianis, G. Vachtsevanos, P. Borne, and A. Croitoru, Eds. Dubrovnik, Croatia: Springer, 2019, pp. 286–300.

[87] M. Chinnici, A.-L. Kor, E. Rondeau, and A. Grishina, “Sustainable data center in smart cities: the role of sustainability-related metrics,” in *First Sustainable Solutions for Growth SSG 2018*, 2018.

[88] A. Grishina, M. Chinnici, A.-L. Kor, E. Rondeau, J.-P. Georges, and D. De Chiara, “Data Center for Smart Cities: Energy and Sustainability Issue,” in *Big Data Platforms and Applications - Case Studies, Methods, Techniques, and Performance Evaluation (In-Press)*, F. Pop, Ed. Springer, 2019.

[89] B. Koffi, A. Cerutti, M. Duerr, A. Iancu, A. Kona, and G. Janssens-Maenhout, *Covenant of Mayors for Climate and Energy: Default emission factors for local emission inventories - Version 2017*. 2017.

[90] I. Munteanu, V. Debusschere, S. Bergeon, and S. Bacha, “Efficiency metrics for qualification of datacenters in terms of useful workload,” in *2013 IEEE Grenoble Conference PowerTech, POWERTECH 2013*, 2013.

[91] C. Becker *et al.*, “Requirements: The Key to Sustainability,” *IEEE Softw.*, vol. 33, no. 1, pp. 56–65, Jan. 2016.

[92] L. Duboc *et al.*, “Do we really know what we are building? Raising awareness of potential Sustainability Effects of Software Systems in Requirements Engineering,” in *27th IEEE International Requirements Engineering Conference*, 2019.

[93] R. Tozer and M. Salim, “Data center air management metrics-practical approach,” in *2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical*

Phenomena in Electronic Systems, 2010, pp. 1–8.

[94] L. Kaufman and P. J. Rousseeuw, *Finding Groups in Data: An Introduction to Cluster Analysis*. Hoboken, NJ, USA: John Wiley & Sons, Inc., 1990.

[95] A. Kassambara, “Determining The Optimal Number Of Clusters: 3 Must Know Methods.” [Online]. Available: https://www.datanovia.com/en/lessons/determining-the-optimal-number-of-clusters-3-must-know-methods/. [Accessed: 06-May-2019].

APPENDIX 1. FACILITY AND DATASET DESCRIPTION

CRESCO4 datasets format

Table 2. Zabbix dataset (power consumption of servers)

Start Timestamp (1h interval)	Min Power	Avg Power	Max Power	Measurement Duration (sec)
1487505600	38600	40637	41500	60
1487509200	40010	40863	41530	60

8673 rows, 4 columns

Table 3. Load Sharing Facility (jobs running on the servers)

Job ID	Start (UTC)	End (UTC)	Duration (sec)	Num Cores Used	User	Queue	File Directory	Executable File Name	Job Status
1	<utc start 1>	<utc stop 1>	102383	256	guarnier	system		wrapper.sh	DONE
2	<utc start 2>	<utc stop 2>	103760	256	guarnier	system		wrapper.sh	DONE

571471 rows, 10 columns

CRESCO6 IT Room

Cold aisle containment

Hot aisle

The CRESCO6 cluster is equipped with a cooling machine Vertiv PX054DD connected with two HCR51 condenser units.

(continues)

APPENDIX 1. Facility and dataset description (continues)

CRESCO6 dataset format

Table 4. Thermal dataset – description of features

Node Name	server ID, integer from 1 to 216;
Timestamp	timestamp of a measurement;
System, CPU, Memory Power	one server instantaneous system, memory, CPU power use in three corresponding columns, W;
Fan 1a, Fan1b, …, Fan 5a, Fan 5b	speed of a cooling fan installed in the node, RPM;
System, CPU, Memory, I/O utilisation	ratio of component utilisation, %, missing data;
Inlet, CPU1, CPU2, Exhaust temperature	temperature at the front, inside (CPU1 and CPU2) and at the rear of every node;
SysAirFlow	speed of air traversing the node, CFM;
DC Energy	total energy that the server has used by the corresponding timestamp, kWh

24 columns, different number of measurements for every month

3347335 rows in total for months May 2018 – February 2019

(continues)

APPENDIX 1. Facility and dataset description (continues)

CRESCO6 dataset snapshot

nodename	tempo	sys_power	cpu_powe	mem_pow	fan1a	fan1b	fan2a	fan2b	fan3a	fan3b	fan4a	fan4b	fan5a	fan5b	sys_util	cpu_util	mem_util	io_util	amb_temp	cpu1_tem	cpu2_tem	exh_temp	sysairflow	dcenergy
cresco6x001	Tue 15 May 16:05:	150	110	13	3584	3712	4736	4992	4736	4992	4864	4992	3584	3712	0	0	0	0	19	45	40	38	13	260.19502
cresco6x002	Tue 15 May 16:05:	150	110	12	3584	3712	4736	4992	4736	4992	4864	4992	3584	3712	0	0	0	0	18	45	40	38	13	251.28987
cresco6x003	Tue 15 May 16:06:	150	110	12	3584	3712	4736	4992	4736	4992	4864	4992	3584	3712	0	0	0	0	20	47	39	40	13	253.87433
cresco6x004	Tue 15 May 16:06:	150	110	10	3584	3712	4736	4992	4736	4992	4864	4992	3584	3712	0	0	0	0	17	44	37	37	13	232.74909
cresco6x005	Tue 15 May 16:06:	150	110	13	3712	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	19	43	37	38	13	243.38527
cresco6x006	Tue 15 May 16:06:	150	110	13	3712	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	18	44	37	38	13	258.80151
cresco6x007	Tue 15 May 16:06:	140	110	14	3712	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	19	44	38	38	13	249.74964
cresco6x008	Tue 15 May 16:06:	150	110	12	3712	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	17	45	38	37	13	249.53803
cresco6x009	Tue 15 May 16:07:	150	110	10	3584	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	18	43	40	38	13	76.54134
cresco6x010	Tue 15 May 16:07:	150	110	14	3584	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	17	43	39	37	13	167.99527
cresco6x011	Tue 15 May 16:07:	150	110	14	3584	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	18	44	39	38	13	262.81563
cresco6x012	Tue 15 May 16:07:	150	110	13	3584	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	18	44	38	38	13	104.02956
cresco6x013	Tue 15 May 16:07:	150	110	11	3712	3712	4736	4992	5120	4992	5120	4992	3584	3712	0	0	0	0	18	45	37	38	13	29.29375
cresco6x014	Tue 15 May 16:07:	160	110	14	3712	3712	4736	4992	5120	4992	4992	4992	3584	3712	0	0	0	0	18	45	41	39	13	273.19437
cresco6x015	Tue 15 May 16:08:	150	110	12	3712	3712	4736	4992	5120	4992	5120	4992	3584	3712	0	0	0	0	18	43	39	37	13	245.60738
cresco6x016	Tue 15 May 16:08:	150	110	14	3712	3712	4736	4992	5120	4992	5120	4992	3584	3712	0	0	0	0	18	44	38	38	13	140.46233
cresco6x017	Tue 15 May 16:08:	150	110	13	3584	3712	4736	4992	4736	4992	4864	4992	3712	3712	0	0	0	0	18	44	41	38	13	257.32675
cresco6x018	Tue 15 May 16:08:	150	110	11	3584	3712	4736	4992	4736	4992	4864	4992	3712	3712	0	0	0	0	18	43	38	38	13	28.72178
cresco6x019	Tue 15 May 16:08:	150	110	14	3584	3712	4736	4992	4736	4992	4864	4992	3712	3712	0	0	0	0	18	43	37	38	13	271.46561
cresco6x020	Tue 15 May 16:08:	150	110	10	3584	3712	4736	4992	4736	4992	4864	4992	3712	3712	0	0	0	0	18	44	38	38	13	254.41421
cresco6x021	Tue 15 May 16:08:	150	110	15	3584	3712	4736	4992	4736	4992	4736	4992	3712	3712	0	0	0	0	18	45	38	38	13	251.13167
cresco6x022	Tue 15 May 16:09:	150	110	13	3584	3712	4736	4992	4736	4992	4736	4992	3712	3712	0	0	0	0	18	44	38	38	13	256.38365
cresco6x023	Tue 15 May 16:09:	150	110	14	3584	3712	4736	4992	4736	4992	4736	4992	3712	3712	0	0	0	0	18	44	40	38	13	256.26152
cresco6x024	Tue 15 May 16:09:	160	120	13	3584	3712	4736	4992	4736	4992	4736	4992	3712	3712	0	0	0	0	18	45	39	38	13	266.06671
cresco6x025	Tue 15 May 16:09:	160	120	13	3584	3712	4736	4992	4736	4992	4736	4992	3584	3712	0	0	0	0	18	46	39	38	13	212.98285

APPENDIX 2. PHASE 2. ANALYSIS OF DATA CENTER THERMAL CHARACTERISTICS

In this section we shift the focus from IT jobs productivity to thermal characteristics of IT room. Here, we rigorously explore the monitored thermal data in a new cluster of ENEA DC that has been assembled and set up in ENEA Portici Research Center. The cluster started processing end user tasks since September 2018 but collected dataset of available thermal and power measurements also covers a period of the cluster stress-testing in May-July 2018.

Referring to research objectives for Phase 2, we explore temperature ranges around the cluster nodes and possible pitfalls of thermal design of an IT room in question. The underlying paradigm of improving DC energy efficiency remains the dominant direction of this work. Optimised thermal management reduces excess energy consumption by conditioning units from one hand and servers that require less energy for internal fans from the other hand. Moreover, compliance of IT room environment with recommended temperature ranges contributes to steady reliability, availability and overall server performance without breakdowns. Therefore, identification of hotspots and negative effects of air dynamics such as bypass or recirculation are useful for DC operators who could improve thermal design and ensure uninterrupted steady operations within their facilities.

Data Center Facility and Datasets Description

Analysis described in this section is founded on server power and surrounding air temperature monitoring data of the new cluster CRESCO6 in ENEA Portici Research Center premises introduced in summer 2018. The new cluster was created due to the growing demand for research center computational and analytic activities as well as the general motivation to keep abreast with current modern technologies.

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

The High-Performance Computing cluster CRESCO6 has nominal computing power of around 700 TFLOPS (500 TFLOPS the result obtained on High Performance Computing Linpack Benchmark, a computational power test that performs parallel calculations on dense linear systems with 64bit precision). It complements the CRESCO4 HPC system, already installed and still operating in the Portici Research Center, with nominal calculation powers of 100 TFLOPS. CRESCO6, on its own, provides increase equal to a factor x7 of the entire computing capability currently available for computational activities in the ENEA research center.

Apart from enhanced hardware, improvement has also been made to the monitoring system of the new cluster. It comprises energy and power meters, temperature and air flow sensors and fans speed registration. Measurements were taken throughout the period from cluster initialisation and performance tuning in the months of May-July to the months of cluster utilisation by end users in September 2018-February 2019 for approximately 9 months in total with a break in the month of August 2018. The measured characteristics are represented in Table 4 of Appendix 1.

Phase 2 Methodology

The nature of measurements does not facilitate the evaluation of energy consumed to produce useful work and energy waste of the new cluster CRESCO6 as it has been done for CRESCO4. Instead, it facilitates the investigation on temperature variation in different parts of the IT room and evaluate thermal metrics. Additional investigation on cluster energy use and idle mode power threshold is shown in Appendix 3. As depicted in Fig. 7, adapted data lifecycle methodology employed for Phase 2 comprises stages of data preprocessing, analysis as well as results interpretation and exploitation in the form of

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

recommendations for the DC. Fig. 7 also clarifies substages of the work: data analysis comprises statistical analysis of thermal data and evaluation of thermal metrics. Available readings of servers' exhaust, inlet, CPUs temperature have been investigated to find general statistical properties and then aggregated into several descriptive metrics that

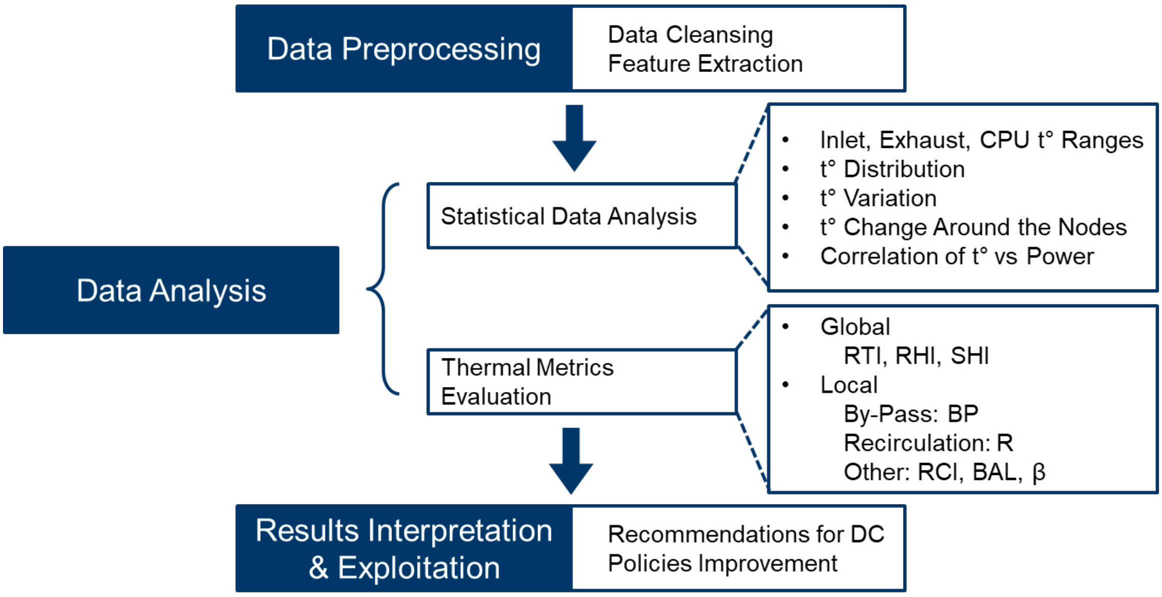

Figure 7. Phase 2. Data Analytics methodology adapted to statistical analysis and metrics evaluation of DC thermal characteristics.

reveal global and local phenomena within the IT room. All stages represented in the Fig. 7 are described in detail below.

Data cleansing step includes extracting valuable features of the thermal data and removing direly incomplete or erroneous data. For example, zero or negative values of temperature measurements should be marked with NaN as not a number will be automatically omitted by the statistical software used the analysis. Such selective marking of missing or erroneous values helps maintain a sizable cleaned dataset. Additionally, it is required to convert all the timestamp fields into the datetime format to be able to perform mathematical operations on them.

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

Data analysis stage includes several substages. Firstly, observed temperature ranges are consolidated and averaged for every month to investigate on periodical fluctuations of the overall cluster air temperature in the cold, hot aisle and inside the nodes. These four thermal sensors' locations are fixed and used for air temperature assessment throughout the entire phase. This stage of analysis aims to meet the research objective **RO2.1.1**.

The next stage of data analysis is devoted to thermal metrics choice and evaluation (**RO2.1.2**). Following globally recognised procedures for metrics evaluation [14], [38], [51], [62], [63], we investigate the efficiency of IT room design, focusing on possible bypass, recirculation, temperature increase within a rack and other factors. They can be categorised into two groups: local and global thermal metrics. Most widespread local thermal metrics comprise Recirculation (R), ByPass (BP), Balance (BAL), shows how well server requirements are met in terms of air distribution in the IT room. The index β indicates presence of self-heating due to recirculation while Rack Cooling Index (RCI, %) shows how effectively the cold aisle temperature is maintained. A list of most frequently discussed global thermal metrics includes Return Temperature Index (RTI, %) that identifies if bypass or recirculation is present globally. It also encompasses Return Heat Index (RHI) that indicates how much the air is mixed in the hot aisle with some unwanted sources of the cold air how effectively the cold air is used to cool the IT equipment or if it there are any air mixes in the underfloor plenum or the hot aisle.

Finally, results of statistical analysis and metrics evaluation have been visualised, interpreted and exploited to provide a list of observed pitfalls and recommendations for the DC operator to improve thermal management (**RO2.1.3**).

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

Results and Discussion

The data cleansing step has reduced the number of features in the resulting dataset as several measurements such as CPU, memory and overall system utilisation are unavailable in reality, although the dataset contains some values for these features. Data concerning 10 fans' speed is excluded from analysis because it is not clear where exactly these fans is beyond the scope of this work. Nevertheless, thermal operation of the cluster cooling system could be characterised by temperature in the hot and cold aisles and CPU temperature measurements as described below.

Thermal Ranges

Average temperature observed at the inlet of the nodes in the cold aisle and exhaust temperature at their rear side in the hot aisle, is represented in Fig. 8. The temperature measurements are also taken next to two CPUs of every node. The setpoints of the cooling system were fixed approximately on 18°C at the output and 24°C at the input of the cooling machine which are represented in Fig. 8 as blue and red vertical lines respectively. It is subsequently discovered that the lower setpoint is variable and provides supply air at 15-18°C as well as high setpoint varies between 24-26°C.

As observed from the graph, cold aisle preserves the setpoint temperature at the inlet of the node, which affirms the efficient design of the cold aisle (i.e. supported by existing plastic panels isolating cold aisle from other spaces in the IT room of the data center). However, exhaust temperature is registered on average at 10°C higher level than the hot aisle setpoint. Notably, exhaust temperature sensors are directly located at the rear of the node (i.e. in the hottest parts of the hot aisle). Therefore, the air in the hot aisle is distributed in such a way that the hotspots are immediately located at the back of server racks and the hot

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

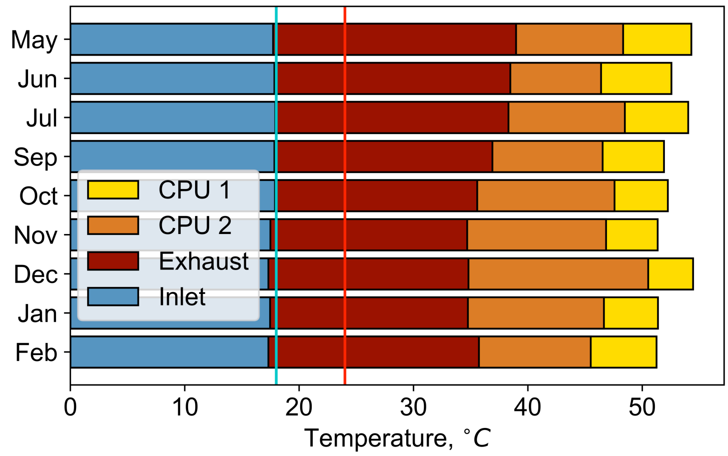

Figure 8. Temperature observed on average in all nodes during consecutive months with vertical lines corresponding to cold and hot aisle setpoints.

aisle air is cooled down to the 24-26°C input level of the cooling system at the CRAC intake due to air circulation and mix in the hot aisle.

Meanwhile, the previously mentioned difference of 10°C between the hotspots and the ambient temperature unravels the cooling system weak points, since it does not account for hotspots directional cooling. In the long term, constant presence of the hot spots might affect the servers' performance which should be carefully addressed by the DC operator.

Thermal Metrics Evaluation

Further assessment of IT room environment will be done through evaluation DC thermal metrics. The formulae for these metrics can be found in literature [14], [38], [51], [62], [63]. Following the notations of [38], we explain which sensors delivered specific information for the metrics calculation and make inferences based on the metrics values. The DC cluster under consideration is equipped with air cooling which operates as depicted in Fig. 9 with all the notations corresponding to the ones in Table 5 (as in [38]).

Based on results of manual sensing of the temperature in cold and hot aisles, three thermal scenarios are developed. They are assumed to correspond to potentially low, medium and high processing loads and account for low, medium and high cooling system load respectively, or high T^{CA}_{sup} and low T^{C}_{ret}, medium T^{CA}_{sup} and medium T^{C}_{ret}, low T^{CA}_{sup} and high T^{C}_{ret}. If values of T^{CA}_{sup} and T^{C}_{ret} are needed for a metric evaluation, they are calculated for

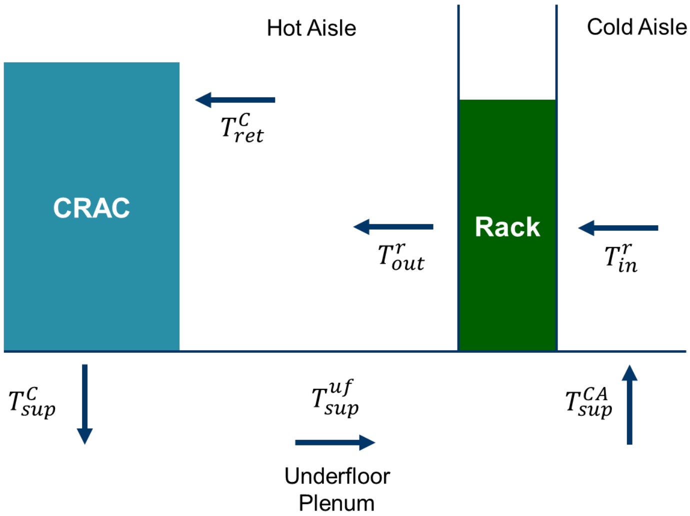

Figure 9. Layout of air distribution in an air-cooled DC.

Table 5. IT room air temperature nomenclature

T^{C}_{sup} – CRAC unit supply air temperature
T^{uf}_{sup} – underfloor plenum supply air temperature
T^{CA}_{sup} – cold aisle supply air temperature
T^{r}_{in} – rack inlet air temperature
T^{r}_{out} – rack output air temperature
T^{C}_{ret} – CRAC return air temperature

three scenarios, low, medium and high cooling system load. Other temperature measurements, T^{r}_{in} and T^{r}_{out}, are taken from available dataset.

The metrics evaluated for every month are consolidated in Tables 6-9. Thermal metrics are evaluated according to three scenarios defined through manual temperature sensing to overcome uncertainties of CRAC unit setpoints.

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

1. Low ITE temperature rise – T_{sup}^{C}=18, T_{ret}^{C}=24.

Table 6. Thermal metrics evaluation in low ITE temperature rise scenario

	RTI	*RHI*	*SHI*	*β*	*BP*	*R*	*BAL*
May 2018	31,42	0,98	0,02	0,02	0,7	0,06	3,18
Jun 2018	31,93	0,98	0,02	0,02	0,7	0,06	3,13
Jul 2018	31,95	0,98	0,02	0,02	0,7	0,07	3,13
Sep 2018	34,46	0,98	0,02	0,02	0,68	0,06	2,9
Oct 2018	36,94	0,98	0,02	0,02	0,65	0,06	2,71
Nov 2018	40,33	0,98	0,02	0,02	0,62	0,06	2,48
Dec 2018	40,64	0,98	0,02	0,02	0,62	0,06	2,46
Jan 2019	40,4	0,97	0,03	0,03	0,62	0,06	2,48
Feb 2019	38,9	0,97	0,03	0,02	0,64	0,06	2,58

2. Medium ITE temperature rise – T_{sup}^{C}=16.5, T_{ret}^{C}=25.

Table 7. Thermal metrics evaluation in medium ITE temperature rise scenario

	RTI	*RHI*	*SHI*	*β*	*BP*	*R*	*BAL*
May 2018	40,21	0,94	0,06	0,06	0,66	0,15	2,49
Jun 2018	41,37	0,94	0,06	0,07	0,65	0,16	2,42
Jul 2018	41,84	0,93	0,07	0,07	0,65	0,17	2,39
Sep 2018	44,78	0,93	0,07	0,07	0,62	0,16	2,23
Oct 2018	48,24	0,93	0,07	0,08	0,6	0,16	2,07
Nov 2018	49,87	0,94	0,06	0,06	0,56	0,12	2,01
Dec 2018	49,06	0,95	0,05	0,05	0,56	0,1	2,04
Jan 2019	49,67	0,94	0,06	0,06	0,56	0,12	2,01
Feb 2019	47,02	0,95	0,05	0,05	0,58	0,11	2,13

(continues)

3. High ITE temperature rise – T^C_{sup} =15, T^C_{ret} =26.

Table 8. Thermal metrics evaluation in high ITE temperature rise scenario

	RTI	RHI	SHI	β	BP	R	BAL
May 2018	51,78	0,89	0,11	0,13	0,61	0,25	1,93
Jun 2018	53,33	0,88	0,12	0,14	0,6	0,26	1,88
Jul 2018	53,95	0,87	0,13	0,14	0,6	0,26	1,85
Sep 2018	57,74	0,87	0,13	0,15	0,57	0,26	1,73
Oct 2018	62,2	0,86	0,14	0,16	0,54	0,26	1,61
Nov 2018	63,94	0,87	0,13	0,14	0,5	0,23	1,56
Dec 2018	62,88	0,88	0,12	0,13	0,5	0,21	1,59
Jan 2019	63,64	0,87	0,13	0,14	0,51	0,22	1,57
Feb 2019	59,94	0,87	0,13	0,13	0,53	0,21	1,67

Table 9. Evaluation of thermal metrics that do not depend on scenario type

	RCI_h^{A1}	RCI_l^{A1}	RCI_h^{A2}	RCI_l^{A2}
May 2018	100	66,33	100	87,37
Jun 2018	99,91	66,05	99,94	87,27
Jul 2018	100	66,81	100	87,55
Sep 2018	100	66,54	100	87,45
Oct 2018	100	66,56	100	87,46
Nov 2018	100	64,82	100	86,81
Dec 2018	100	65,13	100	86,93
Jan 2019	100	64,1	100	86,54
Feb 2019	100	61,91	100	85,72

Three scenarios have a similar general pattern and the findings comprise a dangerous and inefficient combination of overprovisioning of the cooling air and bypass, and a very low possibility of recirculation. High values of RCI metric give evidence of good cold aisle structure and appropriate low setpoints of the CRAC unit. However, RCI is only limited to assessment of rack intake air compliance to the ASHRAE guidelines (A1 and A2) and does not reveal issues that occur within or at the rear of the node. In essence, identified bypass results in lost cooling capacity, higher cooling costs, misleading metrics as in the case of BAL and RCI, and hotspots.

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

Exploration of several scenarios has been an essential step from a theoretical point of view, because the setpoints of the systems are variable and picking only one pair on inlet and output CRAC unit setpoints could have resulted in poor estimation with large uncertainties. However, once the values are computed for all three scenarios, it is clear that general trends stay the same and slight variation of metrics values do not bring about remarkably new results. From the low to high temperature rise scenario, the metrics' values change in a way to depict slightly higher possibility of recirculation, but they are too negligible to warrant superiority of recirculation over bypass.

Phase 2 Conclusion

Phase 2 has covered analysis of ENEA DC CRESCO6 cluster thermal characteristics to unravel hidden effects that occur during IT room air-cooling. Thermal characteristics have been studied through statistical analysis of sensors data installed around all cluster servers. Analysis included estimation of inlet, exhaust, and internal server temperature ranges. The investigation also comprises an evaluation of a set of main thermal metrics that provide an overview of cold aisle design, CRAC unit setpoints efficiency in combination with effects like bypass and recirculation that happen around servers.

Results of the current phase have contributed to evaluation of thermal management of the given DC cluster on the early stages of its operation. They have provided a basis for regular future estimation of the cluster thermal effectiveness as it grows and operates for the smart city and research purposes. Moreover, the methodology proposed for the DC under consideration is applicable to other HPC cites given that their monitoring system provides a set of measurements comparable by their expressivity with the available dataset concerning CRESCO6.

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

The analysis presented in Phase 2 can be further improved by hotspots localisation within CRESCO6 cluster and is formulated as a separate goal of the following phase (Phase 3). Phase 2 concludes hereby with recommendations for IT room thermal management consolidated below based on applicable best practices and other research work [8], [57], [93].

Recommendations for DC IT Room Thermal Management

REC 1. Improve efficiency of the cooling system and reduce bypass to address the issue of hotspots

REC 1.1. Optimise velocity of air injected to the cold aisle through the floor grilles to ensure that the air reaches all the elevated servers of the rack as evenly as possible, i.e. it neither overshoots the top nor is seised on the low levels of the rack;

REC 1.2. Switch control of cooling system setpoints from CRAC return temperature to supply temperature as suggested in [8] to ensures an even supply air temperature independent on the load on CRAC unit;

REC 1.3. Investigate operating cooling unit fans to ensure a slight oversupply of air compared to IT equipment flow demand so that oversupply of air volume is avoided as well as recirculation is minimised. In contained air systems with separate hot and cold aisle, a slightly positive pressure should be maintained in the cold air stream with respect to the hot air stream;

REC 1.4. Once an issue of bypass is overcome, temperature and humidity ranges must be reviewed for potential widening and lowering load on the cooling system.

REC 2. Improve IT room design

REC 2.1. Review the positioning of floor tiles and remove any obstructions from above the tiles;

(continues)

APPENDIX 2. Phase 2. Analysis of Data Center Thermal Characteristics (continues)

REC 2.2. Separate and isolate areas with components that run hotter, such as PDUs, with plastic curtains, which currently are placed on the vertical sides of the racks beside the servers and touch them while servers are more sensible to temperature changes;

REC 2.3. Seal air gaps in the raised floor: improve floor tiles, use foam pillows, cable brushes to isolate underfloor cold air passages and block ways of its dispersion on the way to the cold aisle.

REC 3. Review the load distribution: if some nodes are constantly overloaded, redistribute the load, allow more time for their cooling

REC 4. Improve the monitoring system

REC 4.1. Measure the NP (negative pressure) to benefit from the full set of interrelated thermal metrics;

REC 4.2. As far as this cluster design is not finalised, availability of cooling must be reviewed prior to any ITE changes to correspond to rising ITE cooling demand;

REC 4.3. Periodically review CRAC setpoints calibration and properly maintain the cooling unit; maintain the monitoring system to ensure high accuracy and uninterruptible measurements.

APPENDIX 3. PHASE 3. MACHINE LEARNING FOR DATA CENTER THERMAL CHARACTERISTICS ANALYSIS

This section is devoted to the localisation of hotspots during further CRESCO6 cluster thermal characteristics analysis. Statistical analysis of temperature measurements described in Phase 2 could not pinpoint specific nodes which caused rack hotspots. Therefore, here, we aim to apply Machine Learning techniques for node clustering to identify the incidence of hotspots. The results of this phase are expected to help the DC maintenance team mitigate negative effects of the hotspots (note: this was impossible having only statistical estimations from Phase 2).

Methodology

To reiterate, a Machine Learning clustering technique is chosen for deeper analysis of hotspots location and applied to the dataset of CRESCO6 nodes temperature measurements described in Phase 2. Locating hotspots in the CRESCO6 group of nodes* is achieved through clustering of sequential sets of nodes into clusters with higher or lower hot aisle and internal server temperature.

The steps of data analysis are presented in Fig. 10. They involve data preprocessing (cleansing and dataset organisation), three data analysis substages which lead to results exploitation in the form of recommendations for the DC. On the data analysis step sequential clustering requires determining the optimal number of clusters (done with the use of two indices), actual clustering of servers into groups with low, medium and high surrounding air temperature ranges and consolidation of results to obtain the most frequently occurred cluster label for each server.

To elaborate, on the data preprocessing stage, the dataset is cleansed zero and missing

*Here the term "group of nodes" stands for the data center "cluster", but the latter is not used to avoid its confusion with clusters of data which will be introduced further.

(continues)

APPENDIX 3. Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis (continues)

values, and is organised as shown in Table 10. In Table 10 $base$ is an indicator of one of the three combinations of measurements used as the basis for clustering and $range \in \{cold, medium, hot\}$ corresponds to the temperature of the cluster centroid.

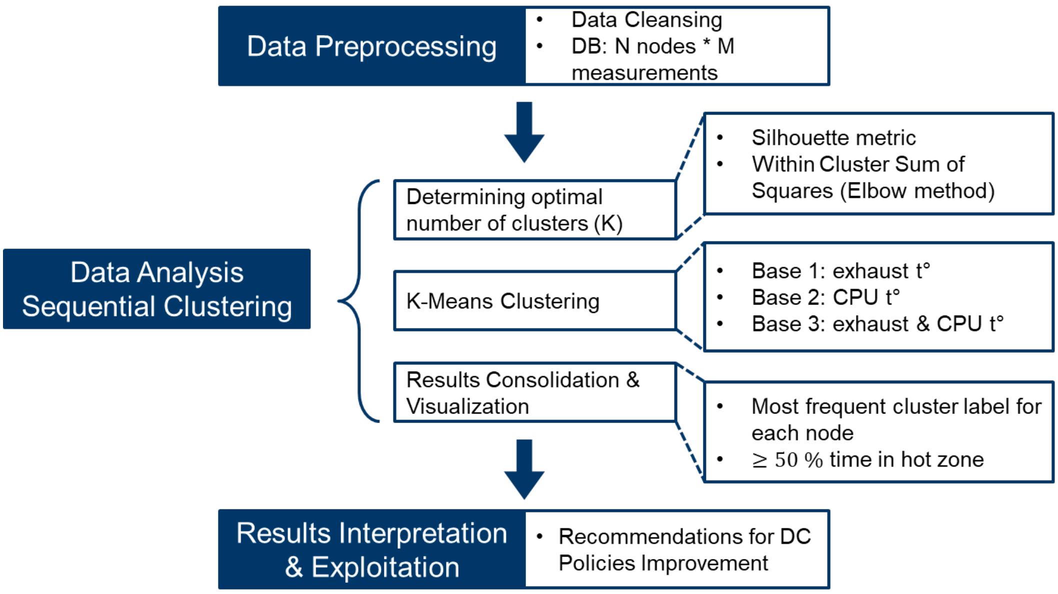

Figure 10. Phase 3. Data Analytics methodology adapted to sequential clustering based on DC thermal characteristics.

Table 10. Dataset for clustering

Time label	Real time	Node ID	Inlet T°	Exhaust T°	CPU 1 T°	CPU 2 T°	Cluster label
t_1	$t_1 + t_{1n_1}$ … $t_1 + t_{1n_N}$	n_1 … n_N	$T_{in_{11}}$ … $T_{in_{1N}}$	$T_{exh_{11}}$ … $T_{exh_{1N}}$	$T_{CPU1_{11}}$ … $T_{CPU1_{1N}}$	$T_{CPU2_{11}}$ … $T_{CPU2s_{1N}}$	$C_1\ base_{range}$ … $C_1\ base_{range}$
t_2	$t_2 + t_{2n_1}$ … $t_2 + t_{2n_N}$	n_1 … n_N	$T_{in_{21}}$ … $T_{in_{2N}}$	$T_{exh_{21}}$ … $T_{exh_{2N}}$	$T_{CPU1_{21}}$ … $T_{CPU1_{2N}}$	$T_{CPU2_{21}}$ … $T_{CPU2s_{2N}}$	$C_2\ base_{range}$ … $C_2\ base_{range}$

(continues)

APPENDIX 3. Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis (continues)

In this work, K-Means algorithm is chosen for clustering the nodes for the reasons that it is fast and suitable for repetitive computations required for sequential clustering based on a relatively small number features (1-3). In addition, the drawback of the algorithm, namely difference in final results due to the random choice of initial centroids, is made infinitesimal by repetition of the clustering procedure which could be seen as cross-validation in the current problem **(RO3.1)**.

The number of clusters K, i.e. number of ranges for $C_{i\,base_{range}}$, is an unknown parameter which was estimated for each of three combinations separately using two metrics: average Silhouette Coefficient and Within Cluster Sum of Squares (WCSS) metric [94], [95]. These two indices are shown in practice in the Results subsection.

Once the optimal number of clusters is obtained, actual clustering is performed for the chosen bases. For every cluster base we further examine how frequently every node is assigned to each cluster and deduce the final cluster label as one of $C_{base_{range}}$ and corresponding sets of nodes as $N_{base_{range}}$. Subsequently, sets of nodes in the hot range for every cluster base are intersected to unravel nodes that are clustered to be in "danger" or hot zone with the highest frequency by three clustering algorithms: $N_{hot} = \bigcap_{bases}\left\{N_{base_{hot\ range}}\right\}$ (**RO3.1**). The next section will discuss results of this clustering procedure and list the nodes that fall in the hot zone.

Results and Discussion

Sequential clustering is further performed for each set of $N{=}216$ samples based on three combinations of available thermal data: exhaust (base 1), CPU (base 2), exhaust and CPU

(continues)

temperature measurements (base 3). The full dataset consists of $M=15569$ sets of temperature monitoring data where each set consists of 216 node samples with data from sensors installed in different locations: in the front (inlet), rear (exhaust) of every node and two sensors inside each node (CPU temperature).

The optimal number of clusters depends on the base chosen for clustering. Two metrics are computed for random sets to be clustered. Their visualisation can be found in Appendix 4. Identified optimal number of clusters K is 3, 5 and 3 for bases 1-3 (exhaust, CPU, exhaust & CPU measurements), which will be used in sequential clustering with these bases.

As a remark, a number of clusters could be determined using several approaches that are currently widespread among data scientists. However, none of them is considered accurate as they all provide an approximate value. This work utilizes two methods: Within Cluster Sum of Squares (WCSS) or an elbow method, and average Silhouette Index. These indices are computed for a range of cluster numbers K and an optimal value is then chosen based on the indices' values. WCSS is a measure of the cluster's compactness and is calculates as follows:

$$WCSS(K) = \sum_{j=1}^{K} \sum_{x \in C_j} \left\| x - \mu_j \right\|^2, \tag{8}$$

where C_j is cluster j, K stands for the total number of clusters and μ_j denotes cluster sample mean. WCSS should be minimised and, in practice, an optimal value of K that is a turning point of the graph where the rate of WCSS decrease slows down, or an elbow of the graph. The method is based on the idea that increasing the number of clusters after the turning point or an elbow is not meaningful, since WCSS decreases only slightly and the positive impact of every next K is low.

(continues)

APPENDIX 3. Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis (continues)

In the Average Silhouette method, a Silhouette index is computed for every data point (or every member of every cluster) and then is averaged over all data points. It estimates consistency of the data within clusters and should be maximised for better separation of the clusters. The Silhouette index is calculated for every data point as follows:

$$s(i) = \frac{b(i) - a(i)}{\max\{b(i), a(i)\}}, \tag{9}$$

where $a(i)$ is the mean distance between a data point and all other points in the same class, $b(i)$ is the mean distance between a sample and all other points in the next nearest cluster. Example of these indices utilisation is shown in Fig. 11. 12 for one step of sequential clustering for exhaust temperature basis. The optimal elbow point of WCSS is K=3 and same for Silhouette index local maximum.

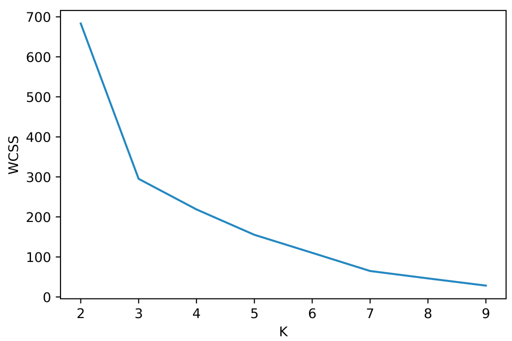

Figure 11. WCSS estimation for clustering based on exhaust temperature

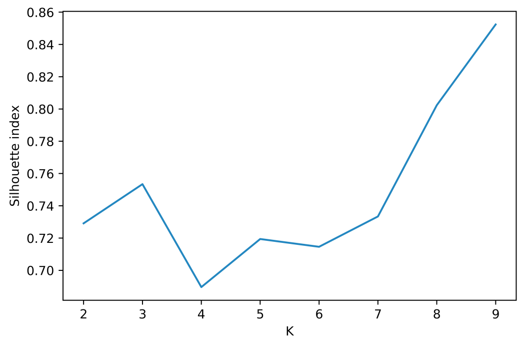

Figure 12. Average Silhouette Index estimation for clustering based on exhaust temperature

Fig. 13 *(a-c)* shows the frequency of occurrence of every node in a particular cluster based on available measurements and clustering algorithm. This information indirectly implies “duration” that particular node resides in a certain temperature range (see legend in Fig. 13 *(a-c)*). Here, the nodes most frequently occur in the medium temperature range for all cluster bases. However, some nodes remain in the hot range for more than 50% of clustering cases.

(continues)

Finally, to cross-validate the clustering we have taken the intersection of nodes clustered into cold, medium and hot ranges. Only one node (or 0.5% of all nodes) has been clustered in the cold range for all three bases algorithms, the medium range has the highest intersection range while 8% (or 18 nodes) are captured in the hot range.

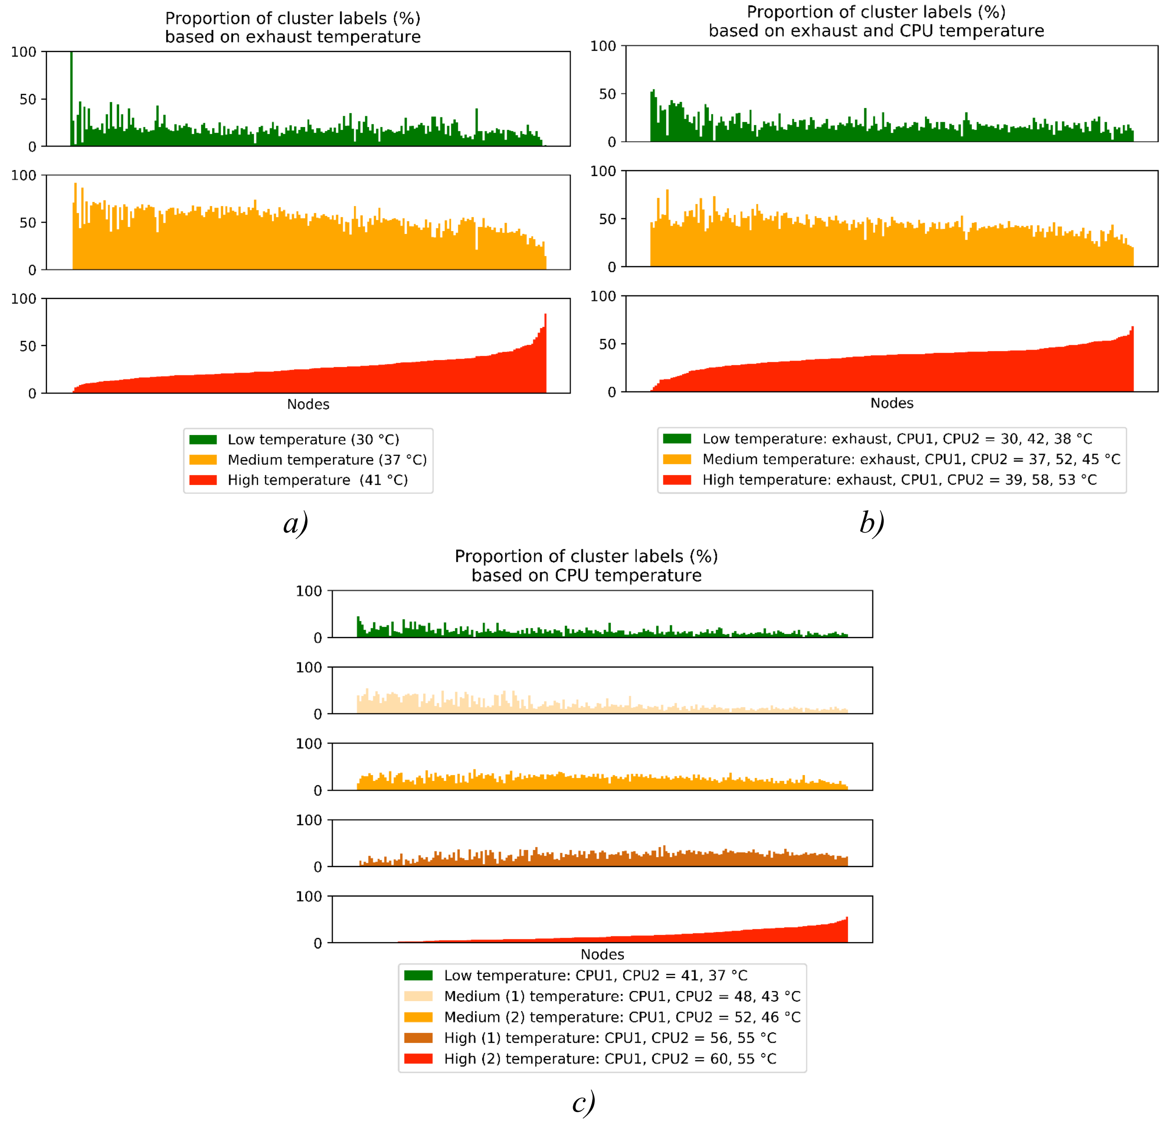

Figure 13. Proportion of nodes clustered into different temperature ranges based on *(a)* Exhaust temperature, *(b)* Exhaust and CPU temperature, *(c)* CPU temperature

(continues)

APPENDIX 3. Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis (continues)

The principle result of Phase 3 analysis is identification of the hot range node IDs and this could be exploited by DC operators to improve thermal conditions in the cluster IT room. Possible solutions could comprise nodes localisation in the room, upgrading cooling system to directional cooling with pumps that could push cold air to the hottest nodes in addition to existing natural convection approach, and improving load scheduling to avoid overloading and overheating of identified nodes.

Phase 3 Conclusion

High-granularity analysis of this section has considered temperature ranges of the air around individual servers to identify and localise frequently overheated servers. A machine learning technique, K-Means clustering, has been applied to sequential sets of thermal measurements for all the cluster servers. The results are further intersected to obtain IDs of the servers that most frequently fall into the hot temperature range.

This part of the work has contributed to thermal characteristics analysis of the DC cluster addressing an issue of hotspots. Being a thermal design pitfall, hotspots impose a risk of local overheating and deterioration of servers exposed to high temperature for prolonged periods of time. In this regard, localisation of hotspots is crucial for better overview and control of the IT room temperature distribution. It provides a direction of future thermal management improvements that would mitigate the mentioned risk.

Finally, the results infer that the majority of the servers operated in the medium and hot temperature ranges. Given that 8% of all cluster servers have been most frequently labelled as hot range nodes, a list of recommendations is suggested below to address the issue of hotspots (**RO3.2**).

(continues)

APPENDIX 3. Phase 3. Machine Learning for Data Center Thermal Characteristics Analysis (continues)

Recommendations for DC IT Room hotspots mitigation

REC 1. Locate nodes by identified hot range IDs and find possible patterns in overheated nodes, for example, position in the rack, and proximity to the PDUs;

REC 2. Tune load sharing so that these nodes are not overloaded in the future;

REC 3. Add directional cooling, for example, spot cooling;

REC 4. Continue monitoring IT room thermal conditions in the immediate proximity of the nodes to evaluate in what way recommendations would affect the IT room temperature.

APPENDIX 4. FULL LIST OF RECOMMENDATIONS

Recommendations for DC IT Jobs Energy Efficiency Enhancement

REC 6. Improve scheduling policies

REC 6.1. Currently utilised FCFS queuing algorithm could be replaced by other algorithms (appropriate under different circumstances), for example, Largest Job First, to optimise the system load, or Smallest Job First, to optimise the throughput, or other algorithms.

Introduce/enhance priorities and add backfilling approaches if necessary: reorder jobs to match the availability of resources and tasks priority.

REC 7. Enhance task resource allocation strategy

REC 7.1. Consider energy usage optimisation queuing strategy rather than to the currently employed chronological order-based strategy. According to the chronological order-based queuing the first job that enters the system is allocated the first available queue with required characteristics. By contrast, energy usage optimisation could foster the choice of the queue with minimal energy consumption for a specific job.

REC 8. Apply best practices for general energy efficiency

REC 8.1. Avoid overprovisioning: provision only the required IT power usage (guideline 4.1.9 from [8]), shut down and remove idle equipment (guideline 4.3.6 from [8]), apply Dynamic Voltage and Frequency Scaling whenever possible; consolidate servers when needed (guideline 4.3.4 from [8]).

REC 8.2. Improve monitoring system to separately take power measurements from IT components such as different-purpose servers and PDUs.

REC 8.3. Review load characteristics and monitoring system.

REC 9. Alert and inform end users about optimal utilisation of the cluster.

REC 9.1. Jobs not optimised for parallelism must not be submitted to parallel queues as they can cause large energy waste. For example, submission of a job requiring only one core to a 24-core queue will cause idle power consumption of 23 remaining cores.

(continues)

APPENDIX 4. Conclusion. Full list of recommendations (continues)

REC 9.2. Jobs must be well-designed and tested prior to their submission to the cluster queues; resubmission of faulty jobs should be avoided to minimise energy waste.

REC 10. Raise environmental awareness of the DC by auditing the energy consumption of existing equipment.

REC 10.1. Identify the degree of DC IT equipment compliance to Energy Star specifications.

REC 10.2. Compare monitored power or energy consumption values with technical specifications to determine if any equipment consumes extra power and investigate its underlying reasons.

REC 10.3. Regularly evaluate a cluster energy consumption and apply performance and productivity metrics for cluster energy efficiency assessment. Include an analysis of carbon emissions into a regular cluster evaluation to determine its environmental impact.

REC 10.4. Consider an integration of free cooling in cold months and renewable energy use.

Recommendations for DC IT Room Thermal Management

REC 1. Improve efficiency of the cooling system and reduce bypass to address the issue of hotspots

REC 1.1. Optimise velocity of air injected to the cold aisle through the floor grilles to ensure that the air reaches all the elevated servers of the rack as evenly as possible, i.e. it neither overshoots the top nor is seised on the low levels of the rack;

REC 1.2. Switch control of cooling system setpoints from CRAC return temperature to supply temperature as suggested in [8] to ensures an even supply air temperature independent on the load on CRAC unit;

REC 1.3. Investigate operating cooling unit fans to ensure a slight oversupply of air compared to IT equipment flow demand so that oversupply of air volume is

(continues)

APPENDIX 4. Conclusion. Full list of recommendations (continues)

avoided as well as recirculation is minimised. In contained air systems with separate hot and cold aisle, a slightly positive pressure should be maintained in the cold air stream with respect to the hot air stream;

REC 1.4. Once an issue of bypass is overcome, temperature and humidity ranges must be reviewed for potential widening and lowering load on the cooling system.

REC 2. Improve IT room design

REC 2.1. Review the positioning of floor tiles and remove any obstructions from above the tiles;

REC 2.2. Separate and isolate areas with components that run hotter, such as PDUs, with plastic curtains, which currently are placed on the vertical sides of the racks beside the servers and touch them while servers are more sensible to temperature changes;

REC 2.3. Seal air gaps in the raised floor: improve floor tiles, use foam pillows, cable brushes to isolate underfloor cold air passages and block ways of its dispersion on the way to the cold aisle.

REC 3. Review the load distribution: if some nodes are constantly overloaded, redistribute the load, allow more time for their cooling

REC 4. Improve the monitoring system

REC 4.1. Measure the NP (negative pressure) to benefit from the full set of interrelated thermal metrics;

REC 4.2. As far as this cluster design is not finalised, availability of cooling must be reviewed prior to any ITE changes to correspond to rising ITE cooling demand;

REC 4.3. Periodically review CRAC setpoints calibration and properly maintain the cooling unit; maintain the monitoring system to ensure high accuracy and uninterruptible measurements.

Recommendations for DC IT Room hotspots mitigation

REC 1. Locate nodes by identified hot range IDs and find possible patterns in overheated nodes, for example, position in the rack, and proximity to the PDUs;

(continues)

APPENDIX 4. Conclusion. Full list of recommendations (continues)

REC 2. Tune load sharing so that these nodes are not overloaded in the future;

REC 3. Add directional cooling, for example, spot cooling;

REC 4. Continue monitoring IT room thermal conditions in the immediate proximity of the nodes to evaluate in what way recommendations would affect the IT room temperature.