
Citation:

Grishina, A and Chinnici, M and Kor, A-L and De Chiara, D and Guarnieri, G and Rondeau, E and Georges, JP (2022) Thermal Awareness to Enhance Data Center Energy Efficiency. *Cleaner Engineering and Technology*, 6. ISSN 2666-7908 DOI: <https://doi.org/10.1016/j.clet.2022.100409>

Link to Leeds Beckett Repository record:

<https://eprints.leedsbeckett.ac.uk/id/eprint/7316/>

Document Version:

Article (Accepted Version)

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

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

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

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

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

Highlights

Thermal Awareness to Enhance Data Center Energy Efficiency

A. Grishina, M. Chinnici, A.-L. Kor, D. De Chiara, G. Guarnieri, E. Rondeau, J.-P. Georges

- Demonstrate the implementation of Sharma et al. [14] recommended framework for thermal management at varying levels of granularity within a data center (i.e. thermal data is collected for server node level to IT room);
- Large commercial data center (DC) thermal assessment through the analysis of thermal and energy-related big real monitoring dataset data (over a substantial period of time – 1 year's worth of data) for a commercial data center and generation of heat maps;
- Evidence-based recommendations for the improvement of DC thermal and energy efficiency policies based on ENEA DC thermal big data analysis and generated heat maps, evaluation of thermal metrics, and thermal characteristics of the DC IT room environment;
- Highlight the impact of the ineffective air-cooling system which results in dangerous hotspots that could reduce IT equipment reliability (due to hardware degradation) and lifespan;
- Recommendations in the form of a list of thermal management and monitoring improvements for the DC cluster (which could be transferrable to DCs with air-cooling systems);

Thermal Awareness to Enhance Data Center Energy Efficiency^{*}

A. Grishina^a, M. Chinnici^{b,*}, A.-L. Kor^c, D. De Chiara^d, G. Guarnieri^d, E. Rondeau^e and J.-P. Georges^e

^b *Eindhoven University of Technology, Department of
Mathematics and Computer Science, Netherlands*

^b *ENEA-R.C. Casaccia, 00196, Rome, Italy* ^b *Leeds Beckett
University, LS1 3HE, Leeds, UK* ^c *ENEA-R.C. Portici, 80055, Portici
(Naples), Italy* ^d *Université de Lorraine, CNRS, CRAN, F-54000,
Nancy, France*

ARTICLE INFO

:
data center
thermal metrics
energy efficiency
recirculation
bypass
hotspots
ASHRAE guidelines

ABSTRACT

Data centers aim at provisioning on-demand processing, storage and networking capabilities in a reliable and scalable way. In this context, proper maintenance of IT equipment within DC premises is crucial as it ensures prolonged lifetime of servers and uninterrupted availability of resources. DC management teams’ sustainable operation effort comprises various approaches to directly and indirectly reduce DC energy consumption. Thermal management aims to reduce excess energy consumption by air cooling and compute systems. This paper focuses on the analysis of the exact temperatures in a real DC cluster rather than considering device setpoints or guidelines. An extensive statistical analysis of available thermal data collected by server-level sensors, global and local thermal metrics evaluation is conducted. It enables isolating possible risks engendered by potential negative covert cooling-related factors. The ultimate outcome of this research is to bring about improvement of DC thermal management for sustainable operations.

1. Introduction

A plethora of data-driven enterprises, government, educational bodies adopt smart systems that generate an increasing amount of data. They require extensive as well as intensive on-demand processing, storage, and networking capabilities [2, 3]. Data centers provide the facilities for processing these data to ensure reliable and scalable resource provisioning and meet requirements of Service Level Agreements (SLAs), high level of Quality of Service (QoS) and Quality of Experience (QoE). In this context, proper IT equipment maintenance within DC premises is crucial as it ensures prolonged lifetime of servers and uninterrupted availability of resources. Specific focus of DC management is thermal characteristics of IT rooms. IT equipment within a DC encompasses servers, uninterruptible power supply (UPS) units, switches, and other networking as well as processing components. Humidity, temperature and overheating periods (not resulting in automatic shutdown) are assumed varying in predetermined and acceptable (regarding specifications) scales [4]. Additionally, maintaining healthy operational conditions is a complex task because IT devices might have different recommended ranges of operation. Undeniably, covert factors such as bypass, recirculation, hotspots and partial rack overheating may impact IT and power devices, devices critical for efficient DC operations. For instance, considering partitioning of a room in two aisles (one for cold, the other for hot), improper isolation may result in recirculation of hot air or cold air bypass [5]. Consequently, such emerging challenges call for the need for optimized thermal conditions within a DC facility.

Decreasing energy consumption is part of thermal management. It deals with servers (more specifically, their internal fans and their load processing) and cooling systems as well. While such management needs to satisfy temperature requirements and standards, it will inevitably impact on reliability, availability, and overall improved server performance. Current research activities regarding thermal management for data centers lead to: highlighting the main cooling issues for high power density data centers [6]; recommendation of a list of thermal management strategies based[7]; experimentation of the effect of a PUE cooling approach on PUE based on air spraying equipment [8]; investigation on thermal performances of air-cooled (with raised and non-raised floor setups) data centers [9] and quantification of thermos-fluid processes through performance metrics [10]; proposal of a thermal model for joint cooling and workload management [11]; exploration of thermal-aware job scheduling, dynamic resource provisioning and cooling [12]; utilization of real thermal information about servers, inlet/outlet air temperature, air mover speed to create thermal and power maps to monitor the real time status of a data center [13]. Most of these works focus on numerical modelling or simulations [6, 9, 10,

11, 12] or empirical research involving R&D on small-scale data center [8, 13] such that large data centers need more attention (with empirical research based on real relevant data). Doubtless, an efficient management would require a methodology for isolating air dynamics and hotspots and determining negative effects. Equipped with such objective information, DC operators will be able to improve their data center thermal design and ensure uninterrupted steady compute system operations. Additionally, it will be an added value if thermal management-related research adheres to [14] recommended thermal management framework at varying granularity of data centers.

To facilitate DC thermal management, several metrics have been defined by industry and academia [5]. Metrics are the fundamental components of the DC IT room thermal assessment leading to the identification of the root causes of the thermal issues when appropriately deployed. Thus, this research work aims at analyzing room thermal characteristics and determining both the potential solutions to improve the cooling system and to effect even distribution of server waste heat within a DC.

This paper focuses on the assessment of thermal conditions in an IT room of a DC cluster against thermal equipment setpoints and guidelines. It extends previous works [5, 15, 16, 17, 18] in terms of exploring the intricacies of deploying the theoretical application framework in a real data center. Data analysis is achieved from real sensors (close to servers) data to determine negative covert factors of the cooling policy. This paper particularly relies on extensive statistical analysis of available thermal data, global, and local thermal metrics evaluation. The ultimate goal of the current work is improvement of DC thermal management to foster sustainable operations. Identification of appropriate measures to address covert thermal factors would undeniably decrease energy consumption of the DC thermal equipment. This is facilitated through optimized design of IT rooms and improved set up of thermal system configurations.

This paper relies on an evaluation of IT room thermal features and IT devices energy consumption of ENEA Portici CRESCO6 cluster. Additionally, the aim of this paper is to formulate recommendations on how to manage thermal issues in DC. This work exploits extensive statistical analysis and addresses the following research objectives:

- To reveal the temperature ranges of the exhaust, inlet and inner air of the nodes (or servers) during the cluster testing and end-user utilization periods;
- To evaluate the variation of temperature of the air surrounding the computing nodes between two consecutive measurements;
- To identify pitfalls of the cluster's current cooling system design in terms of hotspots, bypass, recirculation;

- To apply global (room-level) and local (node-level) thermal metrics to the available dataset to identify possible existence of bypass and recirculation;
- To provide recommendations related to the thermal management in the IT room.

The remainder of the paper is organized as follows: Section 2 is dedicated to discussion of Related work; Section 3 introduces the paper methodology; Section 4 contains Results and discussion; while Section 5 concludes the paper.

2. Metrics for DC Assessment

Optimal thermal management of the IT room requires a holistic approach of mapping various types of IT equipment to corresponding ambient air conditions as well as energy saving strategies such as free cooling. These in combination with application of thermal metrics to a DC IT room monitoring data, provide an invaluable insight into the condition of the IT room environment. Although several works argue that metrics merely highlight distinct features of the IT room thermal conditions and are not suitable for creating a complete picture of the thermal environment [19, 20, 21, 22], the insight provided by metrics utilization is a step forward to better understand possible areas of improvement in the thermal conditions. Here, we would like to highlight the fact that theoretical advancement of thermal metrics research outweighs practical metrics evaluation form of research. Consequently, this research aims to bridge the gap between theoretical and practical employment of the metrics by applying global and local-scale metrics to real IT room monitoring data. The theoretical aspect of this research addresses thermal guidelines by international bodies and thermal metrics (discussed in the ensuing section).

Research centres, companies and voluntary programs (e.g. Code of Conduct for Energy Efficiency in Data Centers [23, 24]) introduce guidelines to improve sustainability in DC. This involves the promotion of renewable energy sources, adapted hardware and software to contribute to better power efficiency (whether for calculation or cooling), energy consumption.

Among these initiatives is the Energy Star programme, which sets out a number of requirements (particularly in terms of energy usage) for IT equipment to qualify for an eco-label [25]. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [4, 26] maintain its own sustainable rules related to operational devices and power supply to DC. To determine the level of integration in sustainability of a given site (DC), it is also possible to use the Code of Conduct for Energy Efficiency in Data Centers defined by the Joint Research Centre (JRC) [24]. This assessment is based on a holistic framework considering as a whole, DC exploitation and monitoring, power, cooling, building and how general policies and good practices are

implemented. The possible outcome for the deployment of this methodology is a score between 1 and 5 (max) for each DC service. Such scores are important since it enables comparative analyses of DC performances. For a given operator, it is then possible to know if an intervention intended to improve the sustainability is (or not) a good practice. See subsequent section for a more detailed review of such practices.

The ASHRAE organization has been dedicating its remarkable effort into the unification of guidelines for efficient thermal management of DCs since 2004. In the past, DCs rooms were generally cooled down to around a target of 20-21°C corresponding to the harsh thermal operational limit of the devices [26]. Currently, ASHRAE Thermal Guidelines for Data Processing Environments provide one of the most comprehensive list of recommendations. ASHRAE TC9.9 [26], resulting from a broad consensus among major IT manufacturers, introduces an approach to work with recommended envelope for allowed thermals working points. Thermal ranges are hence defined in such a way as to combine both reliability and sustainability, with a particular consideration of energy efficiency. It permits the limits to be temporarily exceeded, while remaining acceptable and for a limited period of time, so that there is no reliability deterioration of the IT equipment.

[26] suggests several optimization procedures presented in a flowchart to guide DC operators. It is organized based on three steps: i) evaluate existing DC according to best practices; ii) classify and determine allowable limits; iii) customize and optimize the DC using different options suggested by the ASHRAE committee. The organization recommendations are related to the extent by which violation of the envelope is permitted. Options are related to temperature and humidity (indoor or outdoor) monitoring. It also encompasses the introduction of new components, upgrades of existing components or retrofit.

The ASHRAE guidelines [26] also consider the side effects introduced by such modifications of the DC thermal setpoints or introduction of economizers and chiller-free techniques. Indeed, optimizations may increase cost, noise, power use by servers, air flow speed and also altered reliability. Next, we would like to highlight a drawback of the ASHRAE guidelines since it specifies a flat rate of 20°C into server inlet as a failure whereas DCs may support different setpoint levels and also has analysis has to be conducted in the context of active and non-active servers.

Differences between DCs stem from variations in terms of dimensions, operational purposes or accessibility (public vs private, confederated) [20]. Requirements in terms of confidentiality, availability, reliability or redundancy may also vary. In all cases, it has been evidenced that DCs have an environmental impact [27, 28], especially in terms of energy/power consumption and (indirect) carbon footprints. These findings (despite differences in terms of

accuracy and representativeness as shown in [20]), are steps forward for promoting awareness of DC environmental impact (amongst research and industry communities) due to increased DC energy demand and the urgent need for appropriate improvements to reduce such impact. One of the pressing challenges that we have identified is the absence of indicators (metrics or measurements) that could provide evidence of the causal relation between management strategy (policies or actions) and environment impact. On the other hand, there are many existing metrics (as defined in [29]) for each DC characteristics. All those metrics have been examined and critically reviewed in literature and the evaluation of DC.

Our focus is on thermal metrics that evaluate the efficiency and effectiveness of the cooling equipment and IT room design [30, 31, 32]. They could uncover global thermal and airflow phenomena (e.g., Return Temperature Index (RTI), Supply Heat Index (SHI) and Return Heat Index (RHI) and local (e.g., Rack Cooling Index (RCI)). Thermal metrics could be exploited to reveal unfavorable air flow caused by physical infrastructure compromise. Widely accepted local metrics are Recirculation (R), By-Pass (BP), Balance (BAL), employed for evaluating if air distribution within the IT room satisfies server requirements. The index β could be used to determine the control performance of the cold aisle temperature since it reveals presence of self-heating due to recirculation while RCI (%) measures of how effectively equipment racks are cooled. Generally, global metrics used are RTI (%) and RHI. RTI evidences whether bypass or recirculation is globally present. RHI has different interpretations: coldness of air cooling IT devices; unexpected sources of cold air in the hot aisle, or more generally, cold air mixes in the underfloor plenum.

All the thermal metrics are specialized, dedicated to a single characteristic. The only solution to provide a methodology that addresses a holistic DC evaluation through the combination of a repertoire of metrics. In this paper, we are going to propose such methodology, with a special focus on relevant thermal awareness metrics.

3. Methodology

This section aims to describe the facility involved for the metrics evaluation (i.e. the IT room with thermal data monitoring devices). Only a limited number of sensors have been installed within the data center under study, and thus, some approximations have been used for data preparation as well as analysis purposes and metrics assessment.

To reiterate the research objectives, we aim to explore thermal phenomena in proximity to the cluster

nodes and their correlation with servers' power consumption. The underlying paradigm for improving DC energy efficiency as an indirect benefit of optimal thermal conditions is the moto of this paper. Hence, detecting hotspots and therefore, identification of hotspots and undesirable effects such as recirculation or bypass could provide evidence-based information to data center owners for policies prioritization as well as enhancement of their facility thermal design to ensure uninterrupted steady operations.

3.1. Facility and Dataset Description

This paper is based on data collected for the new cluster CRESCO6 in ENEA Portici Research Center premises (up and running since summer 2018). It was set up to cope with the increasing demand for research center computational and analytic activities as well as the general motivation to keep abreast with emerging technologies (e.g. big data and IoT data streaming).

The High-Performance Computing cluster CRESCO6 has Nominal computing power of around 500 TFLOPS – 700 TFLOPS, the result obtained on High Performance Computing Linpack Benchmark, a computational power test that performs parallel calculations on dense linear systems with 64 bit precision). It complements the CRESCO4 HPC system, already installed and still operating in the Portici Research Center, with nominal computational power of 100 TFLOPS. CRESCO6 alone provides a multiplication increase factor of

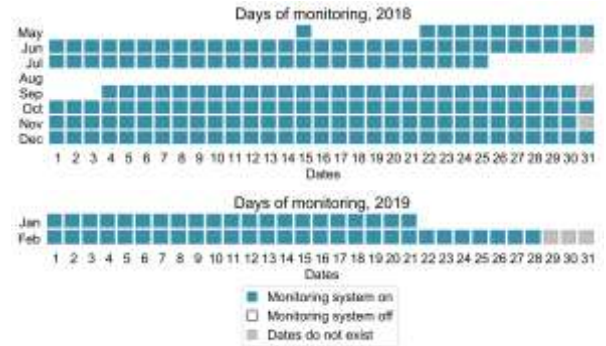


Figure 1: Period of available measurements data in May-December 2018 and January-February 2019.

x7 of the entire computing capability currently available for computational activities in the ENEA research center.

The cluster comprises 216 Lenovo nodes with FatTwin™ 2U form factor, housed in a total of 5 racks.* Xeon® Platinum 8160 CPUs, where each node houses two Intel 24 cores and operates with a frequency clock of 2.1 GHz, for a total of 10368 cores. Each node also houses an overall RAM of 192 GB, equivalent to 4 GB/core. Finally, the nodes are interconnected by an Intel® Omni-Path network with 15

switches of 48 ports each, bandwidth equal to 100 Gb/s, latency equal to 1 μ s. CRESCO6 could satisfy the high scalability needs via the execution of parallel codes. This resource is aimed to support Research and Development activities in ENEA Research Center.

During the last decade, CRESCO HPC system has enabled and supported ENEA participation in national and international projects in various technological sectors which range from bioinformatics to structural biology with impacts on medical and environmental fields such as design of new materials to fluid dynamics for different energy sectors (e.g. photovoltaic, nuclear, energy from the sea, combustion). Furthermore, thanks to the availability of the CRESCO infrastructure, ENEA is a partner of the European Center of Excellence EoCoE (Energy oriented Center of Excellence), Focus CoE (Center of Excellence) projects: one of eight Centers for HPC applications financed by the Horizon2020 program. EoCoE aims to contribute to accelerating the transition to a carbon-free economy by exploiting the growing computational power of HPC infrastructures.

Each computing node of CRESCO6 is equipped with sensors installed directly on the motherboard. Lenovo nodes comprise FatTwin™ 2U form factor. Each node houses two Intel® Xeon® Platinum 8160 CPUs, each with 24 cores and operating with a frequency clock of 2.1 GHz. The sensors on board could vital and non-vital parameters of the hardware for the entire calculation node. These sensors detect various temperatures at different points of the calculation node (particularly CPU and RAM), cooling fans rotation speeds, volume of air that passes through the node and an energy meter that provides the state of energy consumption each time it is invoked. Confluent platform [Confluent site, <https://docs.confluent.io/platform.html>, last accessed 27/3/2020] provides access to various power and cooling data on the monitored hardware. This is possible due to two general data access strategies, that is, from a shell such as a bash, or using an API (over the web, via python, or using the confetty CLI API browser). By means of a bash script (always remain active in the background), a reliable automatic procedure (at a 1-minute interval cycle), calls up Confluent instructions that facilitates reading of values for all installed sensors (e.g. for telemetry, the nodesensors command provides access to available power and cooling related data). Thus, for each 1-minute interval, we have the reading of all the sensors in the block. The data is then inputted into a relational MySQL database, specially designed to store this data in tables based on months and years.

Apart from enhanced hardware, the monitoring system of the cluster has also been improved. It comprises energy and power meters, temperature and air flow sensors, as well as fans speed registration. The data were collected during the initialization and the tuning of the cluster (May-July

2018) and its usage (September2018-February2019), i.e. during roughly 9 months as represented in Fig. 1.

The measurement system covered all 216 nodes, out of which 214-215 nodes were consistently monitored and other 1-2 nodes had missing values or were turned off. The monitoring system consisted of energy meter, power meter of CPU, RAM and the entire IT system utilization of every node, CPU temperature for both processing units of each

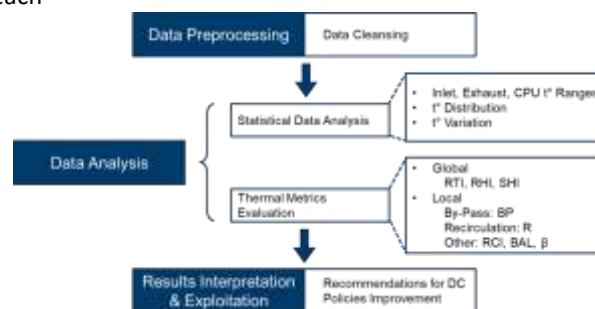


Figure 2: Methodology and Data Lifecycle.

node with thermal sensors installed inside the servers, inlet and exhaust in cold and hot aisles respectively placed in the front and rear parts of every node.

3.2. Data Analysis

The obtained measurements facilitate investigation on thermal characteristics evolution and evaluation of thermal metrics at several locations within the DC. As depicted in Fig. 2, an adapted data lifecycle methodology has been employed for the purposes of this work. The methodology comprises stages of data preprocessing, data analysis, results interpretation and exploitation in the form of recommendations. Fig. 2 also clarifies substages of the work: data analysis involves statistical analysis of thermal data and evaluation of thermal metrics. Available readings of servers' exhaust air temperature, inlet air temperature, as well as CPUs temperature have been analyzed to uncover general statistical properties which are subsequently aggregated into several descriptive metrics that reveal global and local trends. Fig. 2 provides an overview of the methodology.

The data cleansing step includes extracting important thermal data features and removing incomplete or erroneous data. In addition, data preprocessing involves timestamps and user information formatting for further analysis.

Data analysis stage includes several substages. Firstly, temperature measurements are averaged over each month. The objective is to identify periodic variations of the cluster temperature in the cold and hot aisles as well as within the nodes. These four thermal

sensors' locations are specified throughout the entire measurement period. This analysis stage aims to answer research objective 1, i.e. to discover temperature ranges of the air surrounding the computing nodes.

To obtain a more detailed picture of cluster thermal conditions, we explore the distribution of observed air temperature levels every month and throughout the entire observation period (research objective 2). Firstly, we need to explore the air temperature levels in the dataset. In this step, the distinction amongst the nodes is not our priority because our focus is to investigate on the global temperature distribution. The percentage of occurrences for every temperature level in the dataset is computed. Visualizations of the results help reveal the most frequently observed temperature in the proximity of the nodes.

In addition to global thermal characteristics evaluation, the variation of temperature between all pairs of sequential measurements for each node is explored. Pairs of sequential measurements for every node are chosen if the time difference between them is less than a certain interval threshold. This interval threshold is chosen based on the frequency of measurements and the goals of the research. This step provides an overview of how drastic the temperature changed around every node within the limits of a specified interval. The term variation in the context of this paper refers to temperature changes, such as rises and drops, rather than statistical variance (i.e. square of standard deviation). Other types of temperature variation explored at a node level are: difference between the inlet air and CPU temperatures; difference between CPU and exhaust air temperatures. This will provide an insight into how the cold aisle air was heated by the node and effectiveness of the servers' internal fans in lowering the air temperature heated by CPU dissipated waste heat. These findings will contribute to the identification of local pitfalls of the cooling system operation (research objective 3).

The next stage of data analysis is devoted to thermal metrics selection and evaluation (researchobjective4). The metrics were extracted from the literature [5, 21, 32, 33, 34] with the aim to comprehensively characterize the IT room environment, constrained by data obtained from only four thermal sensors. Following globally recognized procedures for metrics evaluation, the relevant metrics employed for this research are: Recirculation (R), ByPass (BP), Balance (BAL), Return Temperature Index (RTI), and Return Heat Index (RHI).

Extra measurements were taken by collaborators working in the DC to better estimate cold air temperature setpoints of the cooling system. These measurements have led to three distinct scenarios used for several metrics evaluation: low, medium, and high load of the servers that correspond to low, medium, and high CRAC output cold temperature setpoints respectively. This step

is essential due to the fact that the CRAC unit is configured to self-adapt to the changes in the inlet air temperature (the one in the hot aisle temperature) and thus, the setpoint for the CRAC output air temperature needed for Rack Cooling Index (RCI) metric evaluation, for example, has frequently changed.

Finally, results of statistical analysis and metrics evaluation highlight pitfalls and provide evidence-based recommendations for thermal management improvement within the data center facility.

4. Results and Discussion

Number of features was reduced after data cleansing. Several measurements have been unavailable due to sensor failures, although the dataset contains partial values for these features. Data concerning 10 different fans' speed is excluded from analysis because the locations of the fans are not specified and thus, remains outside the scope of this study. However, cluster cooling system could be characterized by temperatures in the cold and hot aisles as well as CPU temperature measurements discussed in the subsequent section.

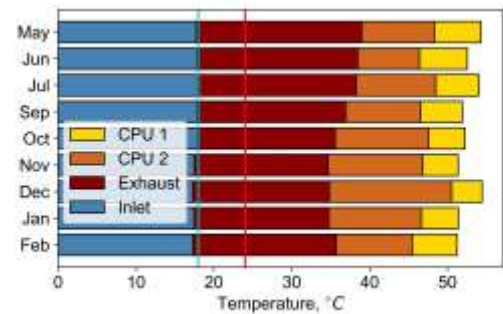


Figure 3: Average temperatures and setpoints for the cold and hot aisles (vertical lines).

4.1. Thermal Ranges

Fig. 3 shows how the temperature varies at the inlet (cold aisle) and at the exhaust (rear side in hot aisle). Temperature measurements for locations next to two CPUs of every node were also taken. The setpoints of the cooling system were fixed approximately at 18°C (blue vertical line) for the output, and for the input, at 24°C (red line). In practice, it has been detected that setpoints are variable such that they are sit in 15-18°C and 24-26°C ranges. Fig. 3 shows that the cooling machine is able to achieve the node's inlet air temperature target. It evidences that the cold aisle design (based on existing plastic panels which isolate the cold aisle from other spaces in the IT room of the data center) is sufficient. On the contrary, the exhaust air temperature is above the setpoint. The temperature overshoot is around 10°C at the hot aisle. In fact, sensors here are located at the rear of

the node such that measurements are much more representative of the hottest points of the aisle. Indeed, cooling reaches the target of 24-26°C at the CRAC intake due to air circulation and air mix in the hot aisle.

Nevertheless, the significant difference from the setpoint remains a weakness of the cooling. It implies that the cooling is not able to tackle the hotspots (since it mainly focuses on the ambient temperature), which could be a serious challenge for the reliability of the servers. They could support short overshoots but not feasible for longer periods. Thus, DC operators would have to address this issue (e.g. directed cooling at hot spots).

Remarkably, although the hotspots are present at the rear of the nodes, the cooling system does influence temperature around the nodes. Table 1 depicts air temperature variations averaged over all nodes and binned based on months. Cold air flows through the node and is measured at the inlet, followed by measurements at CPU 2 and CPU 1 locations and finally, at the exhaust point of the server. The differences between the observed temperature ranges in these locations are averaged for all the nodes and represented in Table 1. Following the air flow, inlet air temperature is heated by 30°C inside the servers until it reaches CPU sensors. It continues to increase by 4-6°C while moving from CPU 2 to CPU 1 sensors and, due to internal server fans, drops by 15-20°C by the moment it reaches the rear of the nodes.

An overview of the temperature distribution obtained with an infrared (IR) thermography camera is also included in the Appendix. The quantitative analysis is thus confirmed by the IR visual images. It should be noted that the IR camera captured air temperature outside the nodes while the sensor readings include only air temperature inside the nodes (servers).

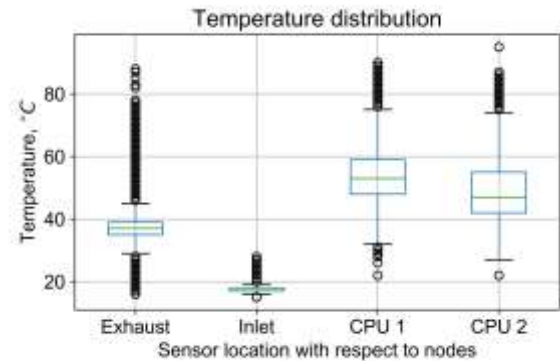
Table 1
Variation of air flow temperature in the immediate proximity of the nodes, averaged over all nodes taken for every month.

Month	Inlet to CPU Rise °C	Between CPU Rise °C	CPU to Exhaust Drop °C
May 2018	30.6	6.0	-15.3
Jun 2018	28.6	6.2	-14.1
Jul 2018	30.6	5.5	-15.7
Sep 2018	28.7	5.4	-15.0
Oct 2018	29.7	4.7	-16.7
Nov 2018	29.3	4.5	-16.7
Dec 2018	33.2	3.9	-19.6
Jan 2019	29.2	4.7	-16.6
Feb 2019	28.2	5.7	-15.5

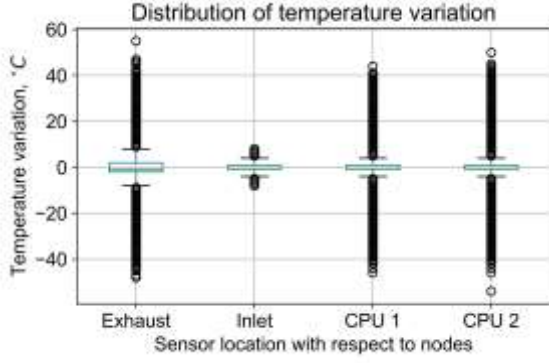
4.2. Distribution and Variation of Monitored Air Temperature Values

The study of observed temperature distribution contributes to the overall understanding of the DC thermal characteristics, as it provides a more detailed overview of prevailing temperature shown in Fig. 3. For every type of thermal sensor, the temperature values are recorded as an integer number, and the percentage of occurrences of each value is calculated and depicted in Fig. 4a. The boxes indicate the temperature occurrences within the interquartile range, a horizontal line inside the box shows the median value and dots outside the boxes depict inlet temperature values. A majority of the inlet temperature values are around 18°C while some have risen up to 28°C (i.e. account for approximately 0.0001% of cases). The inference drawn is that cold aisle temperature sits around the 15-18°C setpoint for most part of the monitored period. Ranges of the exhaust temperature and those of CPUs 1 and 2 are 15-90°C with most frequently monitored values in the intervals of 35-40°C, 50-60°C and 42-55°C respectively. Although these observations might contain measurement errors, they reveal the possibility of servers' risks as they are frequently found to be overheated.

Further, the study focuses on variation between subsequent thermal measurements with the aim to explore stability of the temperature around the nodes. Fig. 4b represents the percentage of temperature variations for every node's sequential measurements observed over the whole period of 9 months. All temperature types have distinct peaks of zero variation which decreases symmetrically and assumes a Gaussian distribution. It could be concluded that temperature tends to be stable in the majority of monitored cases.



(a) Distribution of air temperature for sensed locations in the IT room.



(b) Distribution of air temperature changes between consequent measurements in sensed spots of the IT room.

Figure 4: Distribution and variation of monitored temperature values taken for all nodes and months.

However, the graphs for exhaust temperature, CPUs 1 and 2 temperature variation reveal that less than 0.001% of the recorded measurements show an amplitude of air temperature changes of 40°C. Sudden infrequent temperature fluctuations are less dangerous than prolonged periods of high constant temperature. Nevertheless, further investigation is needed to identify causes of abrupt temperature changes so that measures could be undertaken by the DC operator to maintain longer periods of constant favorable and safe operating conditions.

4.3. Evaluation of the metrics

Thermal metrics for DC and their formula are given in

[5, 21, 32, 33, 34]. We propose to review the list of sensors with regards to the computation and according to the notation provided in [5]. Sensors are also employed to make inferences based on the metrics values.

The DC cluster under consideration is equipped with air cooling (Fig. 5, Table 2). The cold air exits the CRAC unit with supply temperature T^r . Next, air traverses underfloor plenum and obtains possibly higher supply temperature, T_s , and further exits the underfloor space to enter the cold aisle with supply temperature T_{sup}^{CA} . The cold aisle air reaches the nodes in the rack bringing individual inlet temperature

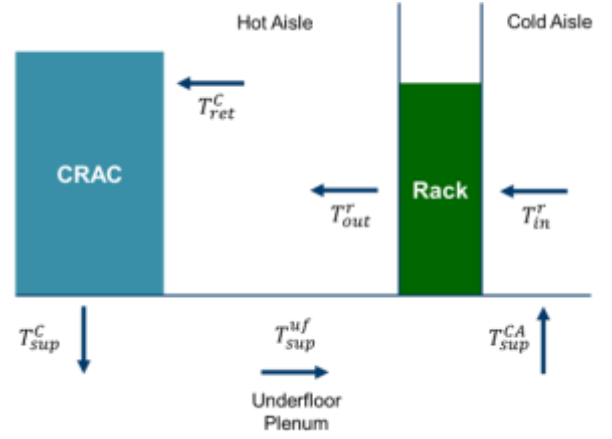


Figure 5: Air flow in a DC.

T_{in}^r to every node. Having passed the server rack, the air is heated up to rack output temperature T_{out}^r and it finally returns to the CRAC unit with T^r . Drawing a parallel between these widely accepted abbreviations for the temperature and available data, T_{in}^r is inlet air, T_{out}^r corresponds to exhaust air, T_{ret}^C is cooling system inlet air setpoint of 24-26°C, and T_{sup}^C is cooling system output setpoint that varies from 15-18°C. The setpoints are preconfigured by the system and the mentioned temperature ranges are reconfirmed by setting up thermal sensors at the locations of T_{sup}^{CA} and T^r .

It is assumed that the difference between T_{sup}^C and T_{sup}^{CA} is negligible. During a manual check, the most frequently observed temperature levels have been 15.0°C, 15.6°C, 16.5°C for T_{sup}^{CA} and 24.0°C, 25.0°C, 26.0°C for T_{ret}^C . Finally, T_{sup}^{uf} is unavailable, and following the previous assumption, it is taken to be equal to T_{sup}^C and T_{sup}^{CA} .

A classification in three scenarios is now introduced, relying on low, medium and high levels, in terms of computation and cooling load, or high T_{sup}^{CA} and low T^r , medium T_{sup}^{CA} and medium T_{ret}^C , low T_{sup}^{CA} and high T_{ret}^C . If values of T_{sup}^{CA} and T_{ret}^C are needed for a metric evaluation, they are calculated for three scenarios, low, medium, and high cooling system load. Available datasets also provide other measures several scenarios since we have noted that setpoints may vary just like T_{in}^r and T_{out}^r . It is really important to consider emergence of estimation uncertainties as in the case of a simple selection of a couple of inlet and outlet CRAC unit setpoints. After having considered these three scenarios, and computed all the values, general trends undeniably remain stable and are not significantly challenged by moderate evolution of metrics. The metrics evaluated for every month for medium cooling equipment load are consolidated in Tables 3 and 4. Other scenarios have not yielded any major findings. Consequently, they are excluded here to avoid repetition.

Table 2

Table 2 IT room air temperature parameters.

Parameter	Description
$T_{C_{sup}}$	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

According to RTI metric ($< 100\%$) and BP (> 0.5), the racks experience air bypass, low values of R, β , RHI (or high SHI=1-RHI) and RTI ($< 100\%$) are evidences of low (or absence of) recirculation. In all scenarios, BAL metric is equal to 1.5-3.0. It differs from its benchmark value of 1 and shows that server requirement for cooling air is BAL times overprovisioned. However, high BP values signify that cold air passes through the servers too fast and does not cool them to a desired value. Thus, even under the conditions of excess of cold air, servers are still subject to overheating, which supports previously discussed findings and local servers' hotspots. RCI metric does not depend on thermal scenario and depicts how far the cold air (from CRAC unit) deviates from the allowable and recommendable ASHRAE temperature ranges for DC IT equipment. This metric is evaluated to show compliance with A1 and A2 ASHRAE classes, to which cluster IT equipment could subsume. The good design of the cold aisle is established by the RCI large values. This also applies to the choice of the cold setpoints of the CRAC unit. These findings nonetheless are constrained because: they strictly correspond to assessment of rack inlet air compliance to the ASHRAE guidelines and they do not highlight phenomena within or at the rear of the node. Naturally, identified bypass points to evidence of cooling leaks, high fees, misleading metrics (e.g. BAL and RCI), and hotspots. Note that from the low to high temperature rise scenario, the metrics values change in such way that depict slightly higher possibility of recirculation, but they are too negligible compared to bypass.

Table 3
Medium cooling system load scenario.

Month	RTI	RHI	SHI	β	BP	R	BAL
May 2018	40.2	0.9	0.1	0.1	0.7	0.1	2.5
Jun 2018	41.4	0.9	0.1	0.1	0.7	0.2	2.4
Jul 2018	41.8	0.9	0.1	0.1	0.7	0.2	2.4
Sep 2018	44.8	0.9	0.1	0.1	0.6	0.2	2.2
Oct 2018	48.2	0.9	0.1	0.1	0.6	0.2	2.1
Nov 2018	49.9	0.9	0.1	0.1	0.6	0.1	2.0
Dec 2018	49.1	1.0	0.0	0.1	0.6	0.1	2.0
Jan 2019	49.7	0.9	0.1	0.1	0.6	0.1	2.0
Feb 2019	47.0	0.9	0.1	0.1	0.6	0.1	2.1

5. Discussion and Recommendations

Effective thermal management is crucial for any DC as it maintains recommended healthy conditions for the IT

Table 4

Metrics not related on a scenario.

Month	RCI_{hA1}	RCI_{hA1}	RCI_{hA2}	RCI_{hA2}
May 2018	100.0	66.3	100.0	87.4
Jun 2018	99.0	66.1	99.9	87.3
Jul 2018	100.0	66.8	100.0	87.6
Sep 2018	100.0	66.5	100.0	87.5
Oct 2018	100.0	66.6	100.0	87.5
Nov 2018	100.0	64.8	100.0	86.8
Dec 2018	100.0	65.1	100.0	86.9
Jan 2019	100.0	64.1	100.0	86.5
Feb 2019	100.0	61.9	100.0	85.7

critical equipment and helps prolong its lifespan. Another essential factor following effective thermal operation is that energy consumed by cooling equipment is optimized while providing sufficient cooling to the servers to prevent overloading in servers' internal fans. Optimal thermal management, achieved through continuous and rigorous monitoring as well as tuning of cooling equipment brings about positive effects towards DC sustainability goals of efficient energy use, reduced physical and heat waste and increased lifespan of the critical equipment. Also, metrics evaluation is a step forward to equip a DC with a competitive edge in the global market of HPC facilities. Regular assessment and optimal operation of the DC render it an attractive computing facility option for processing data-intensive applications. Its attractiveness is derived from the exploitation of Artificial Intelligence and Machine Learning to bring about sustainability, and corporate social responsibility for the general good of societies and businesses, for example, in smart cities.

In this case study, the CRAC unit cold aisle set points comply with guidelines for A1 and A2 ASHRAE ITE (IT Equipment) classes. However, the air temperature changes from around 18°C to 45-50°C on average within the nodes and to around 35°C at the rear of the servers. These observations are detrimental to the health of ITE, as they evidence the presence of hotspots which are caused by high CPU power consumption and thus, overheating. RTI values indicate the presence of bypass, and albeit BAL metric shows that server cooling requirements are overprovisioned 1.5-3 times. Additionally, bypass drastically decreases efficiency and fails to cool the nodes and effectively prevent hotspots. A positive aspect of the cluster IT room thermal design is that recirculation is mitigated by blanking panels and effective air isolation between cold and hot aisles.

In this paper, the proposed methodology for IT room thermal characteristics assessment of the air-cooled DC cluster (located in the region where free air cooling is unavailable) comprises (but is not limited to) the following:

- Low-level granularity analysis of temperature ranges, their mean values and variation between consecutive measurements over all nodes ought to be computed for each month (or any other chosen period);
- To evaluate a comprehensive set of global and local thermal metrics that would reveal occurrences of covert aerial phenomena within an IT room. Several examples are bypass and recirculation when coupled with the metrics' testing, could assess the degree of compliance to IT equipment thermal management guidelines (due to effectiveness of the cooling system operation and adopted set points).

The main contribution of this paper is the analysis of a use case Data Center Thermal Characteristics that encompasses statistical analysis and thermal metrics evaluation of its cluster-based on monitoring data. A list of the contribution of this research work is as follows:

- Evidence-based recommendations for the improvement of the DC thermal and energy efficiency policies based on ENEA real data center thermal big data analysis, testing of thermal metrics, and thermal characteristics of DC IT room environment;
- Conducted data analysis and its associated findings help increase DC operators' general awareness of possible thermal related weaknesses in DC thermal management;
- Highlight the impact of ineffectiveness air-cooling system which results in dangerous hotspots that could reduce IT equipment reliability (due to hardware degradation) and lifespan;
- Recommendations of thermal management and monitoring improvements for the DC cluster – which could be transferrable to DCs with air-cooling systems (see below).

The following section includes a set of recommendations for the IT room thermal characteristics management based on revealed pitfalls of the real DC cluster analysis work:

- R.1. Improve the cooling system efficiency reducing bypass phenomena addressing the issue of hotspots;
 - R.1.1 Optimize the velocity of air injected to the cold aisle using floor grilles to ensure that the air reaches the elevated servers of the rack as evenly as possible, i.e. it neither overshoots the top nor is seized on the low levels of the rack;
 - R.1.2 In order to ensure the air supply temperature independent from the load on the CRAC unit,

the switch control of the cooling system setpoints from CRAC return temperature to supply temperature is necessary (as suggested in [35])

- R.3.3 To investigate the operating cooling unit fans to improve a slight oversupply of air compared to IT equipment flow demand to avoid a superfluous oversupply of air volume and minimization of recirculation in the room. In contained air systems (partitioned in hot and cold aisles), a slightly positive pressure should be maintained in the cold air stream with respect to the hot air stream;
 - R.4.4 Once the bypass problem is overcome, temperature and humidity ranges must be reviewed for potential lowering load and widening its distribution on the cooling system.
- R.2. Improve the design of IT room;
 - R.2.1. To provide a new design of floor tiles and remove any obstacle from above the tiles;
 - R.2.2. Separate and isolate areas with components that run with hotter ambient temperature (e.g. PDUs) than servers (note: they are more sensitive to temperature changes);
 - R.2.3. Seal air gaps in the raised floor using: improvement of floor tiles; use of cable brushes to isolate underfloor cold air passages and block its diffusion to the cold aisle; use of foam pillows;
 - R.2.4. Once the bypass problem is overcome, temperature and humidity ranges must be reviewed for potential lowering load and widening its distribution on the cooling system.
 - R.3. Redesign the load distribution in order to redistribute the load and allow more time for their cooling e.g: if some nodes are constantly overloaded;
 - R.4. To adopt and/or improve the monitoring system;
 - R.4.1. Measure the Negative Pressure (NP) to benefit from the full set of mutually linked thermal metrics;
 - R.4.2. Review (Periodically) CRAC set-point calibration and properly maintain the cooling unit;
 - R.4.3. Use the monitoring system to ensure high accuracy and uninterrupted measurements.

The results of the current work have contributed to the evaluation of thermal management of a real DC cluster use case at the early stages of its lifecycle. This research work findings form the basis of future regular assessment of the cluster thermal effectiveness and it is expected to grow whilst meeting the demands of a smart city as well as HPC-related computing processing power. Moreover, the methodology proposed for the use case DC is transferable to other DCs with the condition that the DC is equipped with a monitoring system that yields a set of measurements that

are comparable (in terms of expressivity) to that of CRESCO6.

6. Conclusion

This work has included analysis of the thermal characteristics of ENEA HPC DC CRESCO6 cluster with the aim to expose covert effects related to air cooling. To improve thermal awareness, statistical analysis has been conducted, focusing on sensors in proximity to cluster servers. Analyses included in this research are: estimation of inlet air, exhaust air, and internal server temperature ranges, temperature variation and its correlation with servers power use. Investigation conducted also encompasses a review of the main thermal metrics as an evaluation of the cold aisle design, CRAC unit setpoint efficiency in combination with effects of bypass and recirculation that typically occur around servers.

The proposed methodology in this research is applicable to any DC that would benefit from a monitoring system installed in the IT room and collects measurements data of thermal characteristics. It is recommended that sensors be installed in close proximity to the nodes so as to better identify local phenomena and take appropriate actions to mitigate them locally instead of tuning global setpoints (guideline 9.1.8 in [24]). It is, however, essential to have thermal sensors at the room level as well as at the CRAC unit supply and return air levels (guidelines 9.1.3, 9.1.4 in [24]). Thus, if the monitoring system is sufficiently complex to cover all levels of granularity, such a DC may apply the entire methodology outlined in this research. Otherwise, parts of the analysis could be replicated given only partial measurements in comparison to CRESCO6 dataset available for this research work.

The analysis presented in this research work could be enhanced in future work by taking into account precise hotspots localization. This paper introduces a list of recommendations to be combined with applicable best practices and other research work [23, 24, 36]. To reiterate, DC thermal awareness ought to be prioritized viewing the fact that it is closely associated to DC energy efficiency and environmental impact.

A. IR Thermography on CRESCO 6 Cluster

Thermographic investigations (IR thermography) have been performed on the CRESCO6 cluster on the 25th of July 2019 between 15:00 and 17:00 CEST. This type of investigative examination has also allowed us to detect problems of thermal hotspots and possible leaks of cold air flows from the air conditioners. The pictures were taken with FLIR E40 Infrared Thermal Imaging Camera 64501-0101 enabled with 0.07°C sensitivity.

In order to obtain representative images, the cluster was thermo-photographed in its entirety, both in the cold aisle, where cold air is pumped into the room closed by a partition comprising VC panels that prevents the escape of fresh air both from the side of the hot aisle where hot air leaving the calculation nodes is released.

Electromagnetic radiation of an object is measured through the detection of emitted electromagnetic radiation (IR). We had to record values of basic parameters of the camera. There are several parameters that must be taken into consideration in order to make an accurate temperature measurement.

Of greatest importance are the evaluation of the emissivity of the framed subject and the apparent reflected temperature of the surrounding environment. Emissivity represents the material's ability to emit electromagnetic radiation in certain wavelengths. In our case, the emissivity indicates (in percentage terms) how much thermal radiation in the thermal imager's wavelength is actually emitted by the object and how much is reflected. With emissivity 1, an ideal black body is represented, that is all the radiation is emitted by the subject without any reflection at all. With lower emissivity levels, it will also be necessary to enter the reflected apparent temperature, the value which, processed by the thermal imager, allows clean signal from the energy not directly emitted by the subject. In our case, given that the framed material is an anodized black body with a temperature up to 50%, it was considered appropriate to set the emissivity between 0.95 and 0.97.

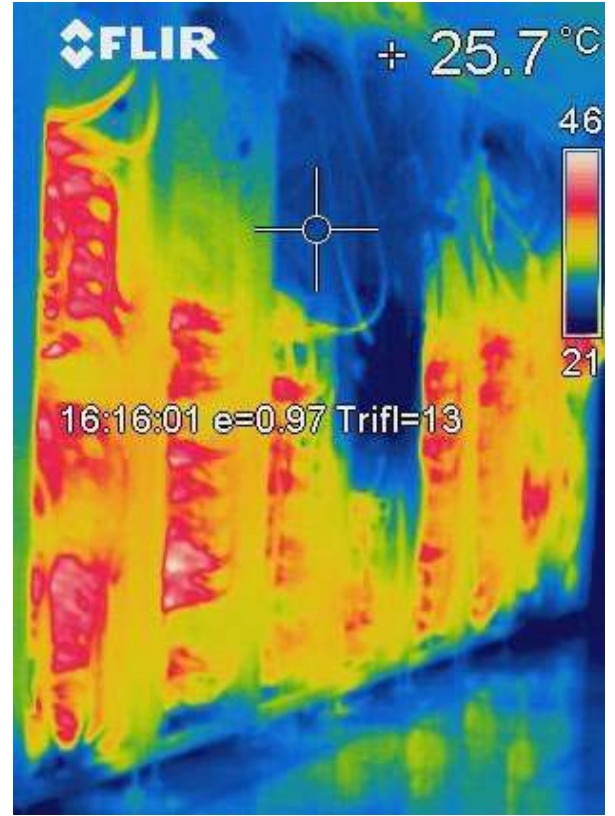
Resulting air temperature heat maps are shown in Fig. 6, 7, 8. The figures include information on the captured momentary temperature distribution, where colder air is depicted in blue and hotter air in red. The photos include parameter settings: in the lower part of the image, time of the photography, emissivity level e , *Trifl* parameter indicating that background temperature is captured. *Flir* stands for the camera model; circle surrounded by four straight lines is the pinpoint of the picture for which the temperature is indicated in the upper right corner of the figure. The hot aisle experiences the highest air temperatures from around 25.7-35.2°C and up to 46°C as shown in Fig. 6. The central area of the connected vertical racks experiences cold air leakage visible in Fig. 7 due to a specific design of a rack: the column combines nodes (hotter areas) and switches (colder area). The cold aisle supply air temperature varies between 17°C and 27°C (Fig. 8).

B. Thermal Metrics Calculation [5]

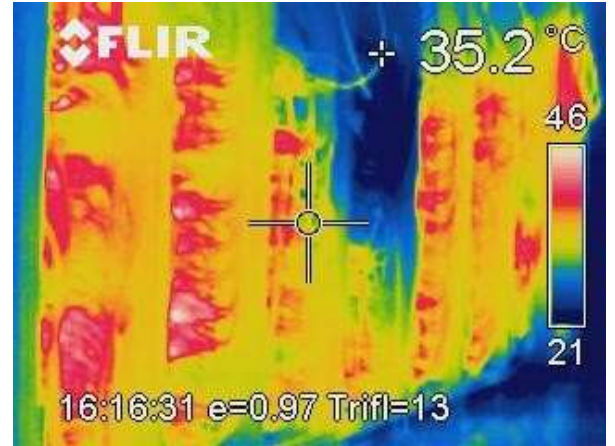
$$RTI = \frac{T_{retC} - T_{supC}}{T_{retC} - T_{inr}} \cdot 100\%;$$

$$\begin{aligned}
& T_{out} \\
RHI &= 1 - SHI; \\
SHI &= \frac{\sum_i \sum_j T_{ini,jr} - T_{supC}}{T_{ini,jr} - T_{supC}}; \\
& T_{inr} - T_{supC} \\
\beta &= \frac{T_{outS} - T_{inr}}{T_{outS} - T_{supC}}; \\
BP &= \frac{T_{outS} - T_{retC}}{T_{outS} - T_{supC}}; \\
R &= \frac{T_{inS} - T_{supC}}{T_{outS} - T_{supC}}; \\
BAL &= \frac{T_{outS} - T_{inS}}{T_{outS} - T_{supC}}; \\
&= \frac{\sum_i (T_{lowrec} - T_{in(i)})}{T_{retC} - T_{supC}} \\
RCH &= \frac{1 - \sum_i (T_{in(i)} - T_{lowrec})}{n \cdot T_{lowrec} - T_{lowallow}} \quad \text{if } T_{in(i)} < T_{lowrec}; \\
RCH &= \frac{1 - \sum_i (T_{in(i)} - T_{highrec})}{n \cdot T_{highallow} - T_{highrec}} \quad \text{if } T_{in(i)} > T_{highrec}
\end{aligned}$$

where i represents number a rack and j — number of a row,
 n — total number of racks.



(a) Zoom-in at a low temperature region of the hot aisle.



(b) Zoom-in at a hot temperature region of the hot aisle.

Figure 6: Hot aisle of the cluster, at the rear of the nodes.

References

- [1] Klimova, A., Rondeau, E., Andersson, K., Porras, J., Rybin, A., Zaslavsky, A.. An international Master's program in green ICT as a contribution to sustainable development. *Journal of Cleaner Production* 2016;135:223–239. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0959652616307077>. doi:10.1016/j.jclepro.2016.06.032.
- [2] Hashem, I.A.T., Chang, V., Anuar, N.B., Adewole, K., Yaqoob, I., Gani, A., et al. The role of big data in smart city. *International Journal of Information Management* 2016;36(5):748–758. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0268401216302778>. doi:10.1016/j.ijinfomgt.2016.05.002. arXiv:arXiv:1011.1669v3.

- [3] Cloud Council, . Deploying Big Data Analytics Applications to the Cloud : Roadmap for Success. Cloud Standards Customer Council



Figure 7: Hotspots (red) and loss of cold air (blue) in the central area.

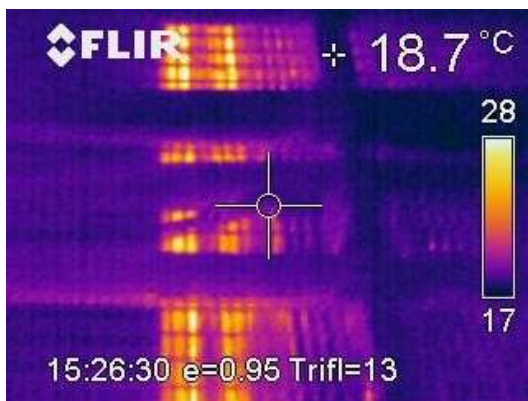


Figure 8: Cold aisle floor temperature.

- 2014;.
- [4] ASHRAE, . Data Center Power Equipment Thermal Guidelines and Best Practices. Tech. Rep.; 2016.
- [5] Capozzoli, A., Serale, G., Liuzzo, L., Chinnici, M.. Thermal Metrics for Data Centers: A Critical Review. *Energy Procedia* 2014;62:391–400. URL: <https://www.sciencedirect.com/science/article/pii/S1876610214034328><https://linkinghub.elsevier.com/retrieve/pii/S1876610214034328>. doi:10.1016/j.egypro.2014.12.40/ 1.
- [6] Bash, C.E., Patel, C.D., Sharma, R.K.. Efficient thermal management of data centers—immediate and long-term research needs. *HVAC&R Research* 2003;9(2):137–152. doi:10.1080/10789669.2003.10391061. [7] Zhang, K., Zhang, Y., Liu, J., Niu, X.. Recent advancements on thermal management and evaluation for data centers. *Applied Thermal Engineering* 2018;142:215 – 231. doi:<https://doi.org/10.1016/j.applthermaleng.2018.07.004>. [8] Fredriksson, S., Gustafsson, J., Olsson, D., Sarkinen, J., Beresford, A., Käufeler, M., et al. Integrated thermal management of a 150kw pilot open compute project style data center. In: 2019 IEEE 17th International Conference on Industrial Informatics (INDIN); vol. 1. 2019, p. 1443–1450. doi:10.1109/INDIN41052.2019.8972145.
- [9] Srinarayana, N., Fakhim, B., Behnia, M., Armfield, S.W.. Thermal performance of an air-cooled data center with raised floor and non-raised-floor configurations. *Heat Transfer Engineering* 2014;35(4):384–397. doi:10.1080/01457632.2013.828559. arXiv:<https://doi.org/10.1080/01457632.2013.828559>.
- [10] Schmidt, R.R., Cruz, E.E., Iyengar, M.. Challenges of data center thermal management. *IBM Journal of Research and Development* 2005;49(4.5):709–723. doi:10.1147/rd.494.0709.
- [11] MirhoseiniNejad, S., Moazamigoodarzi, H., Badawy, G., Down, D.G.. Joint data center cooling and workload management: A thermal-aware approach. *Future Generation Computer Systems* 2020;104:174 – 186. doi:<https://doi.org/10.1016/j.future.2019.10.040>.
- [12] Fang, Q., Wang, J., Gong, Q., Song, M.. Thermal-aware energy management of an hpc data center via two-time-scale control. *IEEE Transactions on Industrial Informatics* 2017;13(5):2260–2269. doi:10.1109/TII.2017.2698603.
- [13] Zhang, S., , Zhou, T., , Ahuja, N., , Refai-Ahmed, G., , Zhu, Y., , Chen, G., , et al. Real time thermal management controller for data center. In: Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm). 2014, p. 1346–1353. doi:10.1109/ITHERM.2014.6892436.
- [14] Sharma, R.K., Bash, C.E., Patel, C.D., Friedrich, R.J., Chase, J.S.. Balance of power: dynamic thermal management for internet data centers. *IEEE Internet Computing* 2005;9(1):42–49. doi:10.1109/MIC.2005.10.
- [15] Capozzoli, A., Chinnici, M., Perino, M., Serale, G.. Review on performance metrics for energy efficiency in data center: The role of thermal management. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2015;8945:135–151. doi:10.1007/978-3-319-15786-3_9. [16] Grishina, A., Chinnici, M., De Chiara, D., Guarnieri, G., Kor, A.L., Rondeau, E., 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). Bucharest, Romania: IEEE. ISBN 978-1-5386-7649-3; 2018, p. 1– 11. URL: <https://ieeexplore.ieee.org/document/8588212/>. doi:10.1109/CSE.2018.00008. [17] Grishina, A., Chinnici, M., De Chiara, D., Rondeau, E., Kor, A.L.. Energy-Oriented Analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center. *Applied ph ed. Dubrovnik, Croatia: Springer; 2019, p. 286–300. URL: http://link.springer.com/10.1007/978-3-030-21507-1_41. doi:10.1007/978-3-030-21507-1_41.*
- [18] Grishina, A., Chinnici, M., Kor, A.L., Rondeau, E., Georges, J.P., De Chiara, D.. Data Center for Smart Cities: Energy and Sustainability Issue. In: Pop, F., editor. *Big Data Platforms and Applications Case Studies, Methods, Techniques, and Performance Evaluation (InPress)*; chap. Data Center for Smart Cities: Energy and Sustainability Issue. Springer; 2019,.
- [19] Daim, T., Justice, J., Krampits, M., Letts, M., Subramanian, G., Thirumalai, M.. Data center metrics: An energy efficiency model for information technology managers. *Management of Environmental Quality: An International Journal* 2009;20(6):712–731. doi:10.1108/14777830910990870.
- [20] Whitehead, B., Andrews, D., Shah, A., Maidment, G.. Assessing the environmental impact of data centres part 1: Background, energy use and metrics. *Building and Environment*

2014;82:151–159. URL: <https://linkinghub.elsevier.com/retrieve/>

pii/S036013231400273X. doi:10.1016/j.buildenv.2014.08.021.

- [21] Chinnici, M., Capozzoli, A., Serale, G.. Measuring energy efficiency in data centers. In: *Pervasive Computing: Next Generation Platforms for Intelligent Data Collection*; chap. 10. ISBN 9780128037027; 2016, p. 299–351. doi:10.1016/B978-0-12-803663-1.00010-3.
- [22] Quintiliani, A., Chinnici, M.. DC4Cities. An environmentally sustainable DC for Smart Cities. Final DC4Cities standardization framework and results description of the European Cluster. Tech. Rep. D7.3; Rome, Italy; 2016.
- [23] Acton, M., Bertoldi, P., Booth, J., Newcombe, L., Rouyer, A., Tozer, R.. 2019 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency 2018;(v 10.1.0):1–48.

URL: <http://publications.jrc.ec.europa.eu/repository/bitstream/>

JRC110666/kjna29103enn.pdf. [24] Acton, M., Bertoldi, P., Booth, J., Newcombe, L., Rouyer, A., Tozer, R.. 2018 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency; vol. EUR 29103. 2018.

URL: <http://publications.jrc.ec.europa.eu/repository/bitstream/>

- [25] Energy Star, . Data Center Equipment. 2019 (accessed 2019-04-30). URL: https://www.energystar.gov/products/data_center_equipment.
- [26] ASHRAE Technical Committee 9.9, . Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance. Tech. Rep.; 2011.
- [27] Neves, L., Krajewski, J., Jung, P., Bockemuehl, M.. SMARTer 2020: The role of ICT in driving a sustainable future. Technical report, A Report by The Climate Group on Behalf of the Global e-Sustainability Initiative. Global eSustainability Initiative (GeSI) 2012;doi:10.1111/j.2006.0906-7590.04873.x.
- [28] Global e-Sustainability Initiative GeSI, . #SMARTer2030-ICT Solutions for 21st Century Challenges 2015;;1–8URL: <http://smarter2030.gesi.org>.
- [29] Ferreira, A.M., Pernici, B.. Managing the complex data center environment: an Integrated Energy-aware Framework. Computing 2016;98(7):709–749. URL: <http://link.springer.com/10.1007/s00607-014-0405-x>. doi:10.1007/s00607-014-0405-x.
- [30] Quintiliani, A., Chinnici, M., De Chiara, D.. Understanding “workload-related” metrics for energy efficiency in Data Center. In: 2016 20th International Conference on System Theory, Control and Computing (ICSTCC). Sinaia, Romania: IEEE. ISBN 978-15090-2720-0; 2016, p. 830–837. URL: <http://ieeexplore.ieee.org/document/7790771/>. doi:10.1109/ICSTCC.2016.7790771.
- [31] Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., et al. United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, Berkeley, CA, Tech Rep 2016;;1–66URL: <http://eta.lbl.gov/publications/united-states-data-center-energy-usag>. doi:10.1005775.
- [32] Herrlin, M.K.A.. Airflow and Cooling Performance of Data Centers : Two Performance Metrics. ASHRAE Transactions 2008;114(Part 2).
- [33] Herrlin, M.K.. Rack cooling effectiveness in data centers and telecom central offices: The Rack Cooling Index (RCI). ASHRAE Transactions 2005;111(Part 2).
- [34] Cupertino, L., Da Costa, G., Oleksiak, A., Piotek, W., Pierson, J.M., Salom, J., et al. Energy-efficient, thermal-aware modeling and simulation of data centers: The CoolEmAll approach and evaluation results. Ad Hoc Networks 2015;25:535–553. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1570870514002364>. doi:10.1016/j.adhoc.2014.11.002.
- [35] Chinnici, M., De Chiara, D., Quintiliani, A.. 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). Orlando, FL, USA: IEEE. ISBN 978-1-5386-1956-8; 2017, p. 269–272. URL: <http://ieeexplore.ieee.org/document/8328399/>. doi:10.

1109/DASC-PICom-DataCom-CyberSciTec.2017.56.

- [36] Tozer, R., Salim, M.. Data center air management metrics-practical approach. In: 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. IEEE. ISBN 978-1-4244-5342-9; 2010, p. 1–8. URL: <http://ieeexplore.ieee.org/document/5501366/>. doi:10.1109/ITHERM.2010.5501366.