Cleaner Engineering and Technology 6 (2022) 100409

Contents lists available at ScienceDirect

Cleaner Engineering and Technology

journal homepage: www.sciencedirect.com/journal/cleaner-engineering-and-technology

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], J.-P. Georges [e]

[a] *Simula Research Laboratory and the University of Oslo, Kristian Augusts gate 23, 0164, Oslo, Norway*
[b] *ENEA-R.C. Casaccia, 00 196, Rome, Italy*
[c] *Leeds Beckett University, LS1 3HE, Leeds, UK*
[d] *ENEA-R.C. Portici, 80 055, Portici, Naples, Italy*
[e] *Université de Lorraine, CNRS, CRAN, F-54 000, Nancy, France*

ARTICLE INFO

Keywords:
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 data center (DC) premises is crucial as it ensures prolonged lifetime of servers and uninterrupted availability of resources. This paper focuses on the analysis how the exact temperatures in a real big DC cluster of 10 368 cores correspond to thermal guidelines for IT devices. 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. The ultimate outcome of this research is the improvement of DC thermal management for sustainable operations via a set of recommendations.

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 (Hashem et al., 2016, Cloud Council, 2014). Data centers provide the facilities for processing these data to ensure reliable and scalable resource provisioning. Proper thermal characteristics within DC premises are crucial as they ensure prolonged lifetime of servers and uninterrupted availability of resources.

Specific focus of DC management is dedicated to the physical characteristics of IT rooms. Humidity, temperature and overheating periods (not resulting in automatic shutdown) are assumed varying in predetermined and acceptable scales regarding specifications (ASHRAE, 2016). In reality, maintaining healthy operational conditions is a complex task, because IT devices, such as servers, uninterruptible power supply (UPS) units, and networking 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. For instance, if a room is partitioned in hot and cold aisles, improper isolation may result in recirculation of hot air or cold air bypass (Capozzoli et al., 2014; Chinnici, 2020a,b). Consequently, such emerging challenges call for optimized thermal conditions within a DC.

Current research activities regarding thermal management for DCs lead to: highlighting the main cooling issues for high power density data centers (Bash et al., 2003); recommendation of a list of thermal management strategies (Zhang et al., 2018); studying the effects of a cooling approach based on air spraying equipment on PUE metric (Fredriksson et al., 2019); investigation on thermal performances of air-cooled with raised and non-raised floor setups DC (Srinarayana et al., 2014) and quantification of thermo-fluid processes through performance metrics (Schmidt et al., 2005); proposal of a thermal model for joint cooling and workload management (Mirhoseini Nejad et al., 2020); exploration of thermal-aware job scheduling, dynamic resource provisioning and cooling (Fang et al., 2017); utilization of real thermal information about servers, inlet/outlet air temperature, air moving speed to create thermal and power maps to monitor the real time status of a DC (Zhang et al., 2014).

Most of thermal management approaches indicated in the literature focus on numerical modeling or simulations (Bash et al., 2003, Srinarayana et al., 2014, Schmidt et al., 2005, MirhoseiniNejad et al., 2020,

---

* Corresponding author.

E-mail addresses: a.a.grishina17@gmail.com (A. Grishina), marta.chinnici@enea.it (M. Chinnici), A.Kor@leedsbeckett.ac.uk (A.-L. Kor), davide.dechiara@enea.it (D. De Chiara), guido.guarnieri@enea.it (G. Guarnieri), eric.rondeau@univ-loraine.fr (E. Rondeau), jean-philippe.georges@univ-loraine.fr (J.-P. Georges).

https://doi.org/10.1016/j.clet.2022.100409
Received 28 October 2020; Received in revised form 11 January 2022; Accepted 11 January 2022

Fang et al., 2017), and empirical research involving R&D on small-scale data center (Fredriksson et al., 2019, Zhang et al., 2014). Therefore, large DCs need more attention in terms of empirical research methods based on real relevant data. An efficient management would require a methodology for isolating air dynamics and hotspots and determining negative effects. So far everal metrics have been defined by industry and academia to facilitate DC thermal management (Capozzoli et al., 2014). Moreover, it will be an added value if thermal management-related research adheres to recommended thermal management framework (Sharma et al., 2005) at varying level of granularity of data centers. Thus, this research work aims at analyzing IT room thermal characteristics and determining both potential solutions to improve the cooling system and ways to smooth 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 set-points and guidelines. The contribution of the work is determining negative covert factors of the current DC cooling policy. Specifically, the factors include the presence of bypass and recirculation, dangerous temperature variations around the nodes and the degree to which the nodes are heated on average during observed months. It extends previous works (Capozzoli et al., 2014, Capozzoli et al., 2015, Grishina et al., 2018, Grishina et al., 2019a, Grishina et al., 2019b) in terms of exploring the intricacies of deploying the theoretical application framework in a real DC. Data is gathered by real sensors placed close to the servers. 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. We aim at highlighting and quantifying the dangerous thermal factors that persist over several months.

This paper relies on an evaluation of IT room thermal features and IT devices energy consumption of ENEA Portici CRESCO6 cluster. Having identified covert negative thermal factors, we aim to formulate recommendations on how to manage thermal issues in DC. This work addresses the following research objectives:

- To reveal the temperature ranges of the exhaust, inlet and inner air of the servers (we refer to servers as *nodes* in this paper) during the cluster stress 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 transferable to other air-cooled DCs.

In summary, the novel contribution of this paper is conducting thermal characteristics analysis based on real thermal related data collected for a big data center cluster of 10 368 cores. Additionally, analysis has been conducted at varying levels, that is from node to IT room level.

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. In combination with application of thermal metrics to a DC IT room monitoring data, such a holistic approach provides a valuable 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 (Daim et al., 2009, Whitehead et al., 2014, Chinnici et al., 2016, Quintiliani and Chinnici, 2016), the insight provided by metrics utilization is a step forward to better understand possible areas of improvement in the thermal management. 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 bodies, companies, and voluntary programs (e.g. Code of Conduct for Energy Efficiency in Data Centers (Acton et al., 2018b; Acton et al., 2018a)) 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 computing or cooling resources, energy consumption and the withdrawal of electronic devices.

Among these initiatives is the Energy Star program, which sets out a number of requirements, particularly in terms of energy usage, for IT equipment to qualify for an eco-label (Energy Star, 2019 (accessed 2019-04-30). The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) (ASHRAE, 2016, ASHRAE Technical Committee 9.9, 2011) maintains 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) (Acton et al., 2018a). This assessment is based on a framework that considers the following aspects of a DC as a whole: data center exploitation and monitoring, power, cooling, building configurations, and how general policies and good practices are implemented. A possible outcome for the deployment of this methodology is a score between 1 and 5 (max) for each DC aspect. Such unified ranking is important since it enables comparative analyses of different DCs performance. For a given operator, it is then possible to know if an intervention intended to improve the sustainability is (or not) a good option. 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, DC rooms were generally cooled down to a target of around 20–21°C corresponding to the harsh thermal operational limit of hosted devices (ASHRAE Technical Committee 9.9, 2011). Currently, ASHRAE Thermal Guidelines for Data Processing Environments provide one of the most comprehensive list of recommendations. ASHRAE TC9.9 (ASHRAE Technical Committee 9.9, 2011), resulting from a broad consensus among major IT manufacturers, introduces an approach to work with a recommended envelope for allowed thermal 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, so that there is no reliability deterioration of the IT equipment.

The ASHRAE guidelines (ASHRAE Technical Committee 9.9, 2011) also consider the side effects introduced by modifications of the DC thermal set-points 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 set-point levels and also has analysis has to be conducted in the context of active and non-active servers.

Differences between DCs stem from variations of their size, operational purposes or accessibility (public, private, confederated) (Whitehead et al., 2014). In all cases, it has been evidenced that DCs have an environmental impact (Neves et al., 2012; Global e-Sustainability Initiative GeSI, 2015), especially by means of excessive energy/power consumption and (indirect) carbon footprints. Despite differences in terms of accuracy and representativeness shown in (Whitehead et al., 2014), these findings are steps forward for promoting awareness of DC environmental impact among research and industry communities. 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, expressed as either policies or actions, and environment impact. On the other hand, there are many existing metrics (as defined in (Ferreira and Pernici, 2016)) 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 (Quintiliani et al., 2016, Shehabi et al., 2016, Herrlin, 2008). 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. RCI (%) measures how effectively equipment racks are cooled. Commonly used global metrics 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, i.e. dedicated to a single characteristic. The only solution to combine their output is 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 a methodology, with a special focus on relevant thermal awareness metrics.

3. Methodology

This section aims to describe the facility involved for the metrics evaluation, it is an IT room with thermal data monitoring devices. Only a limited number of sensors have been installed within the DC 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 servers. Based on revealed factors, DC operators can prioritize policies and enhance the thermal design of their facility to ensure uninterrupted steady operations.

3.1. Facility and dataset description

This paper is based on data collected for the cluster named CRESCO6 in ENEA Portici Research Center premises. The cluster has been operating 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 processing and IoT data streaming.

During the preceding decade, CRESCO High Performance Computing (HPC) system has enabled and supported ENEA participation in national and international projects in various technological sectors which range from bio-informatics 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.

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, installed beforehand and still operating in Portici Research Center, with nominal computational power of 100 TFLOPS. CRESCO6 provides a multiplication increase factor of 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. Each node houses two Intel® 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 10 368 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 μs. CRESCO6 could satisfy high scalability needs via the execution of parallel codes.

Each computing node of CRESCO6 is equipped with sensors installed directly on the motherboard. The sensors on board could read 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 27/3/2 020] 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 that always remains active in the background, a reliable automatic procedure, calls up Confluent instructions at a 1-min interval cycle. The procedure 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-min 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.

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 used in the current work 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 node with thermal sensors installed inside the servers, inlet and exhaust in cold and hot aisle respectively placed in the front and rear parts of every node. In summary, the 3 different categories of sensors deployed for this research are: (i) Category 1 - embedded sensor inside the refrigerating machines Emerson Liebert PDX-PCW; (ii) Category 2 - embedded sensors in the compute/calculation nodes; (iii) Category 3 - external environmental sensors that detect humidity and temperatures outside the nodes and air conditioners (Liebert SN-T rack temperature/humidity sensor). A summary of the wireless sensor infrastructure is shown in Table 1. On the cold side, 4 sensors have been placed respectively, one for each rack. On the hot side 4 sensors, one for each rack and in addition a fifth thermal sensor adjacent to the hot side of the network switch section and a sixth sensor placed at the top of the hot aisle, to understand the temperature at the ceiling, in the halfway point between racks and air conditioners. Fig. 1 show the installed thermal sensors in both the hot as well as cold aisles. These are IP-based

Table 1
Installed wireless thermal sensor infrastructure.

Device/Equipment	Specification	Private IP Address
Hot side thermal sensors	IBM temperature/humidity external sensor RJ45	172.17.1.101
		172.17.1.103
		172.17.1.109
		172.17.1.111
		172.17.1.117
		172.17.1.119
Cold side thermal sensors	IBM temperature/humidity external sensor RJ45	172.17.1.105
		172.17.1.107
		172.17.1.113
		172.17.1.115
Embedded internal sensors of CRAC	Refrigerating Machines Emerson Liebert PDX-PCW for thermal management	172.18.1.216
		172.18.1.217
		172.18.1.218
Embedded internal sensors of compute/ calculation nodes	Internal sensors of Lenovo Nodes FatTwin^{TM} 2U form factor	Internal LAN address of 216 Nodes
Humidity sensors' room	IBM temperature/humidity external sensor RJ45	172.17.1.101
		172.17.1.103
		172.17.1.109
		172.17.1.111
		172.17.1.117
		172.17.1.119
Dedicated Server	MySQL oracle RDBMS 5.7 version	192.107.70.130

wireless sensors that transmit data every 5 min to a dedicated server.

The data were collected during the initialization and the tuning of the cluster (May–July 2018) and its usage (September 2018–February 2019), i.e. during roughly 9 months as represented in Fig. 2.

3.2. Data analysis

The obtained measurements facilitate investigation on evaluation of thermal metrics at several locations within the DC. As depicted in Fig. 3, 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. 3 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

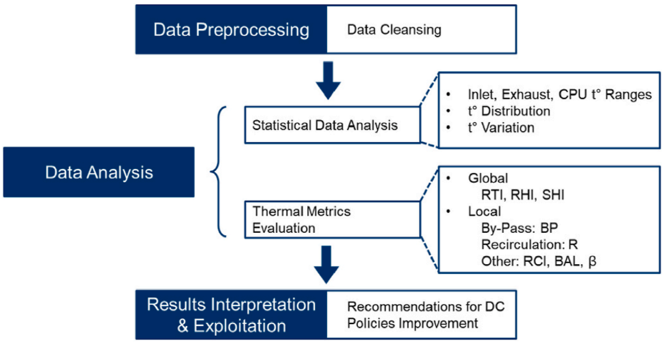

Fig. 3. Methodology and data lifecycle.

(a) Hot side of CRESCO6 and thermal sensors.

(b) Cold side of CRESCO6 and thermal sensors.

(c) Hot side of CRESCO6 and thermal sensors at a height of 2m and 3m.

(d) Hot side of CRESCO6 and thermal sensors at a height of 3m.

Fig. 1. Distribution and variation of monitored temperature values taken for all nodes and months.

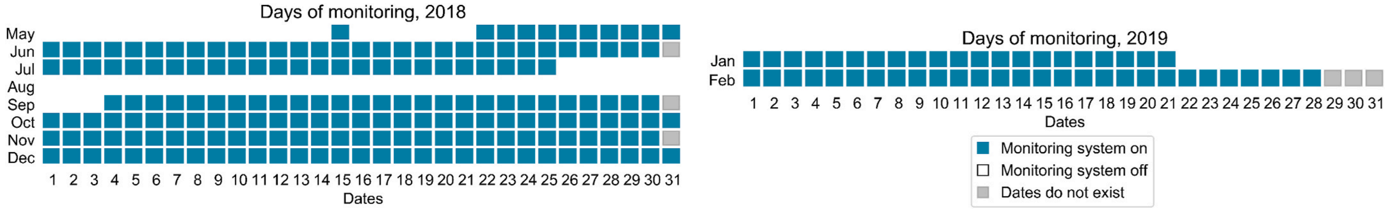

Fig. 2. Period of available measurements data in May–December 2018 and January–February 2019.

analyzed to uncover general statistical properties which are subsequently aggregated into several descriptive metrics that reveal global and local trends. Fig. 3 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 next 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 the first research objective, 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 (the second research objective). Firstly, we need to explore the air temperature levels in the dataset. In this step, the distinction among the nodes is not our priority because our focus is to investigate 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 drastically 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 (the third research objective).

The next stage of data analysis is devoted to thermal metrics selection and evaluation (the forth research objective). The metrics were extracted from the literature (Capozzoli et al., 2014, Chinnici et al., 2016, Herrlin, 2008, Herrlin, 2005, (Cupertino et al., 2015) with the aim to comprehensively characterize the IT room environment, constrained by data obtained from only four types of thermal sensors. Following globally recognized procedures for metrics evaluation, the relevant metrics employed for this research are: Recirculation, ByPass, Balance, Return Temperature Index, and Return Heat Index.

Extra measurements were taken by collaborators working in the DC to better estimate cold air temperature set-points 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 Computer Room Air Conditioning (CRAC) output cold temperature set-points 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 and thus, the set-point for the CRAC output air temperature (the one in the hot aisle) needed for Rack Cooling Index 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 DC facility.

4. Results and discussion

The 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.

4.1. Thermal ranges

Fig. 4 shows how the temperature varies at the inlet (cold aisle) and at the exhaust (rear side in hot aisle) points of the servers. Temperature measurements next to two CPUs of every node were also taken. The set-points of the cooling system were fixed approximately for the output at 18°C (blue vertical line) and for the input at 24°C (red line). In practice, it has been detected that set-points are variable such that they are included in 15–18°C and 24–26°C ranges.

Fig. 4 shows that the cooling equipment is able to achieve the temperature target at the inlet of the node. It is a proof that the design based on existing plastic panels which isolate the cold aisle from other spaces in the IT room of the data center of the cold aisle is sufficient. On the contrary, the exhaust temperature is above the set-point. The temperature overshot is around 10°C at the hot aisle. In fact, sensors are here located at the rear of the node such that measurements are much more representative of the hottest points of the aisle. Indeed, the cooling reaches the target of 24–26°C at the CRAC intake due to air circulation and air mix in the hot aisle. In this bar plot, please, consider the average values of inlet, CPU 1, CPU 2, and exhaust temperatures as the right-most ends of the bars, not the full range from the start to the end of the tape. Thus, the average inlet temperature over all nodes and all measurements in May is around 18°C; the average exhaust temperature is around 39°C; CPU 1 and CPU 2 correspond to the average 48°C and 52°C, correspondingly. We chose a horizontal barplot type of graph to highlight the deviation from the air conditioning set up (vertical lines). We believe that the colored volume between the temperature averages and vertical lines is easier to process for the eye, but it could also have been replaced by dots indicating the average temperatures that we wanted to show for every month per measurement type.

Nevertheless, the significant difference from the set-point remains a weakness of the cooling. It implies that the cooling is not able to tackle the hotspots, which could be a serious challenge for the reliability of the servers. Thus, DC operators would have to address this issue, e.g. directed cooling at hotspots.

Remarkably, although the hotspots are present at the rear of the nodes, the cooling system does influence temperature around the nodes. Table 2 depicts air temperature variations averaged over all nodes and binned based on months. Cold air flows through the node and is

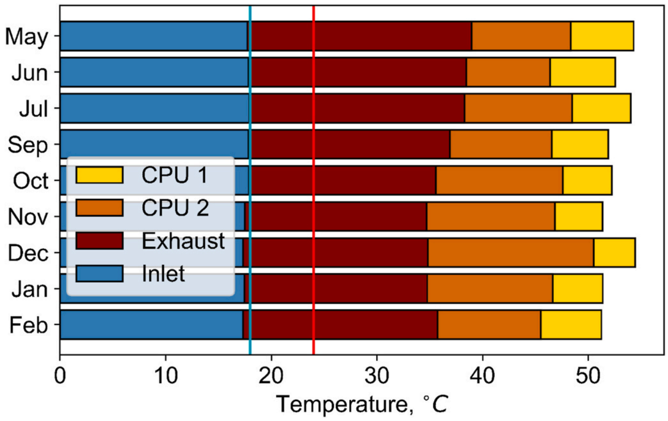

Fig. 4. Average temperatures and set-points for the cold and hot aisles (vertical lines).

Table 2

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	Between CPU Rise	CPU to Exhaust Drop
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

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 2. 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).

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. 4. 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. 5a. The boxes indicate the temperature occurrences within the inter-quartile 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 0.000 1% of cases). The inference drawn is that cold aisle temperature sits around the 15–18°C set-point 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. 5b 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. 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 favourable and safe operating conditions.

4.3. Evaluation of the metrics

Thermal metrics for DC and their formulae are given in (Capozzoli et al., 2014, Chinnici et al., 2016, Herrlin, 2008, Herrlin, 2005, Cupertino et al., 2015). We propose to review the list of sensors according to the notation provided in (Capozzoli et al., 2014). Sensors are also employed to make inferences based on the metrics values.

The DC cluster under consideration is equipped with air cooling (Fig. 6, Table 3). The cold air exits the CRAC unit with supply temperature T^C_{sup}. Next, air traverses underfloor plenum and obtains possibly higher supply temperature, T^{uf}_{sup}, and further exits the underfloor space to enter the cold aisle with supply temperature T^{CA}_{sup}. The cold aisle air reaches the nodes in the rack bringing individual inlet temperature T^r_{in} to every node. Having passed the server rack, the air is heated up to rack output temperature T^r_{out} and it finally returns to the CRAC unit with T^C_{ret}.

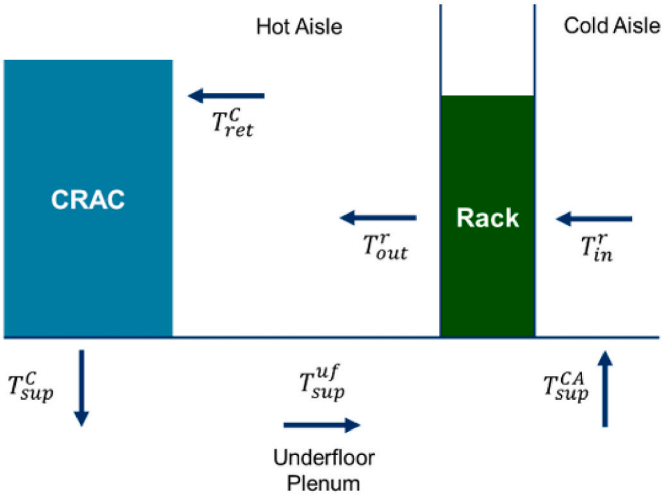

Fig. 6. Air flow in a DC.

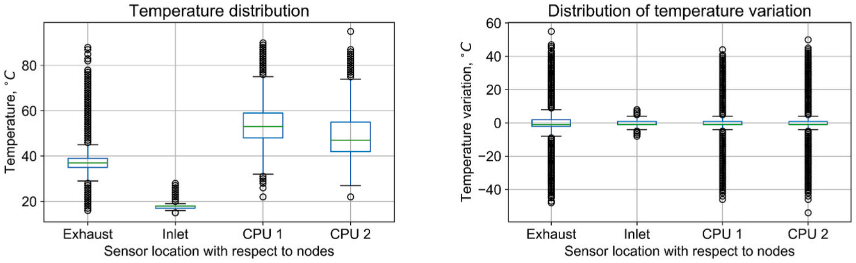

(a) Distribution of air temperature for sensed locations in the IT room.

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

Fig. 5. Distribution and variation of monitored temperature values taken for all nodes and months.

Table 3
IT room air temperature parameters.

Parameter	Description
T^C_{sup}	CRAC unit supply air temperature, measured by internal sensors of CRACs (Category 1)
T^{uf}_{sup}	Underfloor plenum supply air temperature, measured by external sensors placed under floating floor (Category 3), see Fig. 1b
T^{CA}_{sup}	Cold aisle supply air temperature, measured by a group of external sensors placed under racks (Category 3), see Fig. 1b
T^r_{in}	Rack inlet air temperature, measured by internal sensors of each single node (Category 2)
T^r_{out}	Rack output air temperature, measured by internal sensors of each single node (Category 2) and external sensors placed at height 2m and 3m (Category 3), see Fig. 1c
T^C_{ret}	CRAC return air temperature, measured by external sensors at height 2m and 3m (Category 3) see 1c,1d and internal sensors of CRACs (Category 1)

Drawing a parallel between these widely accepted abbreviations for the temperature and available data, T^r_{in} is inlet air, T^r_{out} corresponds to exhaust air, T^C_{ret} is cooling system inlet air set-point of 24–26°C, and T^C_{sup} is cooling system output set-point that varies from 15 to 18°C. The set-points are pre-configured by the system and the mentioned temperature ranges are reconfirmed by setting up thermal sensors at the locations of T^{CA}_{sup} and T^C_{ret}. It is assumed that the difference between T^C_{sup} and T^{CA}_{sup} 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^{CA}_{sup} and 24.0°C, 25.0°C, 26.0°C for T^C_{ret}. Finally, T^{uf}_{sup} is unavailable, and following the previous assumption, it is taken to be equal to T^C_{sup} and T^{CA}_{sup}.

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^{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 three scenarios, low, medium, and high cooling system load. Available datasets also provide other several scenarios since we have noted that set-points may vary just like T^r_{in} and T^r_{out}. It is important to consider estimation uncertainties as in the case of a simple selection of a couple of inlet and outlet CRAC unit set-points. After having considered these three scenarios and computed all the values, general trends 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 4 and 5. Other scenarios have not yielded any major differences from the ones presented here and are excluded to avoid repetition.

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.

Table 4
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

Table 5
Metrics not related on a scenario.

Month	RCI_h^{A1}	RCI_l^{A1}	RCI_h^{A2}	RCI_l^{A2}
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

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 set-points 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.

5. Discussion and recommendations

Optimal thermal management, achieved through continuous and rigorous monitoring as well as tuning of cooling equipment introduces a positive trend 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 (Grishina et al., 2020, Kubler et al., 2019) to ensure sustainability and corporate social responsibility for the general good of societies and businesses.

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. In addition, the work (Grishina et al., 2020) identifies the IDs of frequently overheated servers using sequential clustering of temperature measurements. RTI values indicate the presence of bypass, and BAL metric shows that server cooling requirements are overprovisioned 1.5–3 times. Additionally, bypass drastically decreases efficiency and the thermal equipment 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);
- A comprehensive evaluation of a set of global and local thermal metrics that would reveal occurrences of covert aerial phenomena within an IT room such as bypass and recirculation. In general, metrics 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 value of this paper is that it performs statistical analysis and thermal metrics evaluation for Data Center monitored data compared to thermal guidelines. A list of the contributions of this research work is as follows:

- Evidence-based recommendations are summarized to help improve the DC thermal efficiency policies based on ENEA (real) data center thermal big data analysis, evaluation 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 weaknesses in thermal management of individual DCs;
- Recommendations of thermal management and monitoring improvements for the DC cluster could be transferred to DCs with air-cooling systems (see below).

The following part includes a set of recommendations for ENEA-DC but applicable to any DC.

R.1. Improve the cooling system efficiency reducing bypass phenomena addressing the issue of hotspots identified in this research (look at infrared thermal images in Appendix);

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 set-points from CRAC return temperature to supply temperature is necessary (as suggested in (Chinnici et al., 2017))

R.1.3. Investigate the operating of cooling unit fans to improve a slight oversupply of air compared to IT equipment flow demand. In this way it is possible to avoid a superfluous oversupply of air volume and minimize air 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.1.4. 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 due to hot air recirculation and cold air bypass effects (based on BP value for bypass and RCI value for recirculation as presented in Tables 3 and 4);

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. power distribution units) than servers (note: servers are more sensitive to temperature changes);

R.2.3. Seal air gaps in the raised floor using: floor tiles adjustment; cable brushes to isolate underfloor cold air passages and block its diffusion to the cold aisle; foam pillows.

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 system because the current system only monitors air temperature and velocity;

R.4.1. Measure the Negative Pressure to benefit from the full set of mutually linked thermal metrics;

R.4.2. Periodically review CRAC calibration and properly maintain the cooling unit;

R.4.3. Use the monitoring system to ensure high accuracy and uninterruptible measurements.

In particular, the exploration of the ENEA-Cluster with thermal characteristic analysis, 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 CRAH unit setpoints could have resulted in poor estimation with large uncertainties. However, 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.

The results of the current work constitute therefore the evaluation of thermal management of a real DC cluster use case at the early stages of its life-cycle. This research work findings form the basis of future regular assessment of the cluster thermal effectiveness. 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 expressiveness to that of the studied CRESCO6 cluster.

6. Conclusion

In this work we have analyzed thermal characteristics of data center cluster ENEA HPC CRESCO6 with the aim to expose covert effects related to air cooling. To improve thermal awareness, statistical analysis has been performed on the data gathered by thermal sensors placed in the proximity to cluster servers. Analyses included in this research are: estimation of inlet air, exhaust air, and internal server temperature ranges, as well as temperature variation. The current study also encompasses a review of the main thermal metrics, an evaluation of the cold aisle design, CRAC unit set-point 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 of thermal characteristics. It is recommended to install sensors in close proximity to the computing nodes so as to better identify local phenomena and take appropriate actions to mitigate them locally instead of tuning global set-points (guideline 9.1.8 in (Acton et al., 2018a)). 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 (Acton et al., 2018a)). Thus, if the monitoring system is sufficient 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 (Acton et al., 2018b, Acton et al., 2018a, Tozer and Salim, 2010). To reiterate, DC thermal awareness ought to be prioritized viewing the fact that it is closely associated to DC energy efficiency and environmental impact.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The research reported here was supported and funded by the

PERCCOM Erasmus Mundus Program of the European Union (PERCCOM—FPA 2013–0231). The authors would like to express their gratitude to all the partner institutions, sponsors, and researchers involved in the PERCCOM program (Kor et al., 2019).

Appendix A. IR Thermography on CRESCO 6 Cluster

Thermographic investigations (IR thermography) have been performed on the CRESCO6 cluster on the 25^{th} 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 64 501–0101 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 Figs. 7–9. 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. 7. The central area of the connected vertical racks experiences cold air leakage visible in Fig. 8 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. 9).

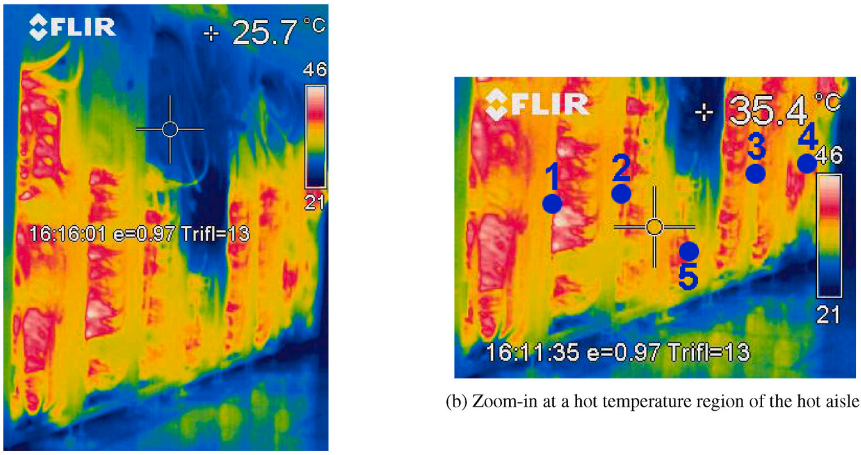

(a) Zoom-in at a low temperature region of the hot aisle.

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

Fig. 7. Hot aisle of the cluster, at the rear of the nodes.

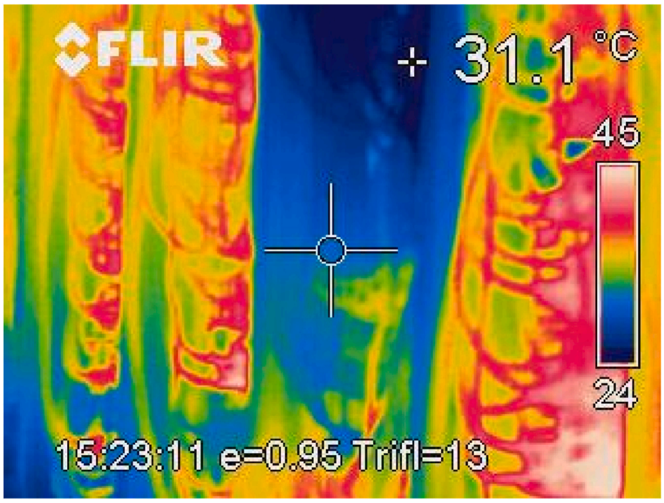

Fig. 8. Hotspots (red) and loss of cold air (blue) in the central area.

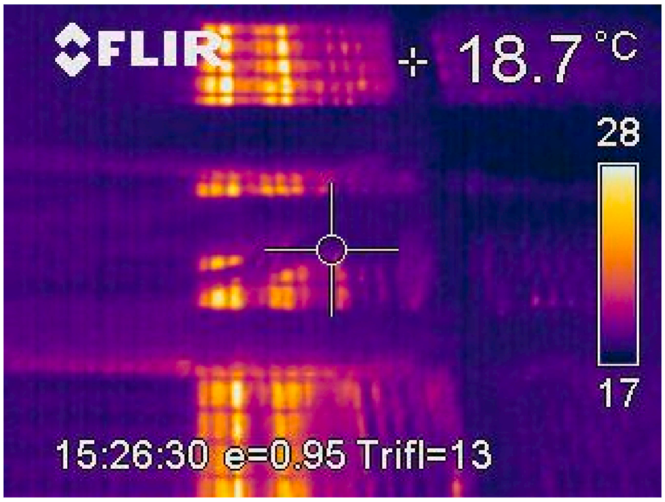

Fig. 9. Cold aisle floor temperature.

Appendix B. Thermal Metrics Calculation (Capozzoli et al., 2014)

$$RTI = \frac{T^{C}_{ret} - T^{C}_{sup}}{T^{r}_{out} - T^{r}_{in}} \cdot 100\%;$$

$$RHI = 1 - SHI;$$

$$SHI = \sum_{i}\sum_{j} \frac{{T_{in\,i,j}}^{r} - T^{C}_{sup}}{{T_{in\,i,j}}^{r} - T^{C}_{sup}};$$

$$\beta = \frac{T^{r}_{in} - T^{C}_{sup}}{T^{r}_{out} - T^{r}_{in}};$$

$$BP = \frac{T^{S}_{out} - T^{C}_{ret}}{T^{S}_{out} - T^{uf}_{sup}};$$

$$R = \frac{T^{S}_{in} - T^{uf}_{sup}}{T^{s}_{out} - T^{uf}_{sup}};$$

$$BAL = \frac{T^{S}_{out} - T^{S}_{in}}{T^{C}_{ret} - T^{C}_{sup}};$$

$$RCI_{l} = \frac{1 - \sum_{i}\left(T_{low\,rec} - T_{in\,(i)}\right)}{n\cdot\left(T_{low\,rec} - T_{low\,allow}\right)}, \ if\ T_{in\,(i)} < T_{low\,rec};$$

$$RCI_{h} = \frac{1 - \sum_{i}\left(T_{in\,(i)} - T_{high_{rec}}\right)}{n\cdot\left(T_{high_{allow}} - T_{high_{rec}}\right)}, \ if\ T_{in\,(i)} > T_{high_{rec}}$$

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

References

Acton, M., Bertoldi, P., Booth, J., Newcombe, L., Rouyer, A., Tozer, R., 2018a. 2018 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency volume EUR 29103. URL: http://publications.jrc.ec.europa.eu/repository/bitstream/JRC110666/kjna29103enn.pdf.

Acton, M., Bertoldi, P., Booth, J., Newcombe, L., Rouyer, A., Tozer, R., 2018b. 2019 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency, pp. 1–48. http://publications.jrc.ec.europa.eu/repository/bitstream/JRC110666/kjna29103enn.pdf.

ASHRAE, 2016. Data Center Power Equipment Thermal Guidelines and Best Practices. Technical Report.

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

Bash, C.E., Patel, C.D., Sharma, R.K., 2003. Efficient thermal management of data centers—immediate and long-term research needs. HVAC R Res. 9, 137–152. https://doi.org/10.1080/10789669.2003.10391061.

Capozzoli, A., Chinnici, M., Perino, M., Serale, G., 2015. Review on performance metrics for energy efficiency in data center: the role of thermal management. Lect. Notes Comput. Sci. 8945, 135–151. https://doi.org/10.1007/978-3-319-15786-3_9.

Capozzoli, A., Serale, G., Liuzzo, L., Chinnici, M., 2014. Thermal metrics for data centers: a critical review. Energy Proc. 62, 391–400. https://doi.org/10.1016/j.egypro.2014.12.401. URL: https://www.sciencedirect.com/science/article/pii/S1876610214034328 https://linkinghub.elsevier.com/retrieve/pii/S1876610214034328.

Chinnici, M., et al., 2020. A machine learning solution for data center thermal characteristics analysis. Energies 13 (17), 4378. https://doi.org/10.3390/en13174378. In this issue.

Chinnici, M., et al., 2020. Data mining for big dataset-related thermal analysis of high performance computing (hpc) data center. LNCS 12143, 367–381. https://doi.org/10.1007/978-3-030-50436-6_27. In this issue.

Chinnici, M., Capozzoli, A., Serale, G., 2016. Measuring energy efficiency in data centers. In: Pervasive Computing: Next Generation Platforms for Intelligent Data Collection, pp. 299–351. https://doi.org/10.1016/B978-0-12-803663-1.00010-3 (chapter 10).

Chinnici, M., De Chiara, D., Quintiliani, A., 2017. 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). , IEEE, Orlando, FL, USA, pp. 269–272. https://doi.org/10.1109/DASC-PICom-DataCom-CyberSciTec.2017.56. URL: http://ieeexplore.ieee.org/document/8328399/.

Cloud Council, 2014. Deploying Big Data Analytics Applications to the Cloud : Roadmap for Success. Cloud Standards Customer Council.

Cupertino, L., Da Costa, G., Oleksiak, A., Piątek, W., Pierson, J.M., Salom, J., Sisó, L., Stolf, P., Sun, H., Zilio, T., 2015. Energy-efficient, thermal-aware modeling and simulation of data centers: the CoolEmAll approach and evaluation results. Ad Hoc Netw. 25, 535–553. https://doi.org/10.1016/j.adhoc.2014.11.002. URL: https://linkinghub.elsevier.com/retrieve/pii/S1570870514002364.

Daim, T., Justice, J., Krampits, M., Letts, M., Subramanian, G., Thirumalai, M., 2009. Data center metrics: an energy efficiency model for information technology managers. Manag. Environ. Qual. Int. J. 20, 712–731. https://doi.org/10.1108/14777830910990870.

Energy Star, 2019 (accessed 2019-04-30). Data Center Equipment. URL: https://www.energystar.gov/products/data_center_equipment.

Fang, Q., Wang, J., Gong, Q., Song, M., 2017. Thermal-aware energy management of an hpc data center via two-time-scale control. IEEE Trans. Ind. Inf. 13, 2260–2269. https://doi.org/10.1109/TII.2017.2698603.

Ferreira, A.M., Pernici, B., 2016. Managing the complex data center environment: an Integrated Energy-aware Framework. Computing 98, 709–749. https://doi.org/10.1007/s00607-014-0405-x. URL: http://link.springer.com/10.1007/s00607-014-0405-x.

Fredriksson, S., Gustafsson, J., Olsson, D., Sarkinen, J., Beresford, A., Käufeler, M., Minde, T.B., Summers, J., 2019. Integrated thermal management of a 150kw pilot open compute project style data center. In: 2019 IEEE 17th International Conference on Industrial Informatics. INDIN, pp. 1443–1450. https://doi.org/10.1109/INDIN41052.2019.8972145.

Global e-Sustainability Initiative GeSI, 2015. #SMARTer2030-ICT Solutions for 21st Century Challenges, pp. 1–8. http://smarter2030.gesi.org.

Grishina, A., Chinnici, M., De Chiara, D., Guarnieri, G., Kor, A.L., Rondeau, E., Georges, J.P., 2018. DC energy data measurement and analysis for productivity and waste energy assessment. In: 2018 IEEE International Conference on Computational Science and Engineering (CSE). IEEE, Bucharest, Romania, pp. 1–11. https://doi.org/10.1109/CSE.2018.00008. URL: https://ieeexplore.ieee.org/document/8588212/.

Grishina, A., Chinnici, M., De Chiara, D., Rondeau, E., Kor, A.L., 2019a. Energy-Oriented Analysis of HPC Cluster Queues: Emerging Metrics for Sustainable Data Center, Applied ph ed. Springer, Dubrovnik, Croatia, pp. 286–300. https://doi.org/10.1007/978-3-030-21507-1_41. URL: http://link.springer.com/10.1007/978-3-030-21507-1_41.

Grishina, A., Chinnici, M., Kor, A.L., Rondeau, E., Georges, J.P., 2020. A machine learning solution for data center thermal characteristics analysis. Energies 13. https://doi.org/10.3390/en13174378.

Grishina, A., Chinnici, M., Kor, A.L., Rondeau, E., Georges, J.P., De Chiara, D., 2019b. Data center for smart cities: energy and sustainability issue. In: Pop, F. (Ed.), Big Data Platforms and Applications - Case Studies, Methods, Techniques, and

Performance Evaluation. (In-Press). Springer chapter Data Center for Smart Cities: Energy and Sustainability Issue.

Hashem, I.A.T., Chang, V., Anuar, N.B., Adewole, K., Yaqoob, I., Gani, A., Ahmed, E., Chiroma, H., 2016. The role of big data in smart city. Int. J. Inf. Manag. 36, 748–758. https://doi.org/10.1016/j.ijinfomgt.2016.05.002. URL: https://linkinghub.elsevier.com/retrieve/pii/S0268401216302778. arXiv:arXiv:1011.1669v3.

Herrlin, M.K., 2005. Rack cooling effectiveness in data centers and telecom central offices: the Rack Cooling Index (RCI). Build. Eng. 111.

Herrlin, M.K.A., 2008. Airflow and cooling performance of data centers : two performance metrics. Build. Eng. 114.

Kor, A.L., Rondeau, E., Andersson, K., Porras, J., Georges, J.P., 2019. Education in green ICT and control of smart systems : a first hand experience from the international PERCCOM masters programme. In: 12th IFAC Symposium on Advances in Control Education ACE. Philadelphia, USA.

Kubler, S., Rondeau, E., Georges, J.P., Mutua, P., Chinnici, M., 2019. Benefit-cost model for comparing data center performance from a biomimicry perspective. J. Clean. Prod. 231, 817–834. https://doi.org/10.1016/j.jclepro.2019.05.183.

MirhoseiniNejad, S., Moazamigoodarzi, H., Badawy, G., Down, D.G., 2020. Joint data center cooling and workload management: a thermal-aware approach. Future Generat. Comput. Syst. 104, 174–186. https://doi.org/10.1016/j.future.2019.10.040.

Neves, L., Krajewski, J., Jung, P., Bockemuehl, M., 2012. 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). https://doi.org/10.1111/j.2006.0906-7590.04873.x.

Quintiliani, A., Chinnici, M., 2016. DC4Cities. An Environmentally Sustainable DC for Smart Cities. Final DC4Cities Standardization Framework and Results Description of the European Cluster. Technical Report D7.3. Rome, Italy.

Quintiliani, A., Chinnici, M., De Chiara, D., 2016. Understanding "workload-related" metrics for energy efficiency in Data Center. In: 2016 20th International Conference on System Theory, Control and Computing (ICSTCC). IEEE, Sinaia, Romania, pp. 830–837. https://doi.org/10.1109/ICSTCC.2016.7790771. URL: http://ieeexplore.ieee.org/document/7790771/.

Schmidt, R.R., Cruz, E.E., Iyengar, M., 2005. Challenges of data center thermal management. IBM J. Res. Dev. 49, 709–723. https://doi.org/10.1147/rd.494.0709.

Sharma, R.K., Bash, C.E., Patel, C.D., Friedrich, R.J., Chase, J.S., 2005. Balance of power: dynamic thermal management for internet data centers. IEEE Internet Comput. 9, 42–49. https://doi.org/10.1109/MIC.2005.10.

Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., Lintner, W., 2016. United States data center energy usage report. Lawrence Berkeley National Laboratory, Berkeley, CA, pp. 1–66. Tech. Rep. http://eta.lbl.gov/publications/united-states-data-center-energy-usag. LBNL-1005775.

Srinarayana, N., Fakhim, B., Behnia, M., Armfield, S.W., 2014. Thermal performance of an air-cooled data center with raised-floor and non-raised-floor configurations. Heat Tran. Eng. 35, 384–397. https://doi.org/10.1080/01457632.2013.828559 arXiv: https://doi.org/10.1080/01457632.2013.828559.

Tozer, R., Salim, M., 2010. Data center air management metrics-practical approach. In: 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. IEEE, pp. 1–8. https://doi.org/10.1109/ITHERM.2010.5501366. URL: http://ieeexplore.ieee.org/document/5501366/.

Whitehead, B., Andrews, D., Shah, A., Maidment, G., 2014. Assessing the environmental impact of data centres part 1: background, energy use and metrics. Build. Environ. 82, 151–159. https://doi.org/10.1016/j.buildenv.2014.08.021. URL: https://linkinghub.elsevier.com/retrieve/pii/S036013231400273X.

Zhang, K., Zhang, Y., Liu, J., Niu, X., 2018. Recent advancements on thermal management and evaluation for data centers. Appl. Therm. Eng. 142, 215–231. https://doi.org/10.1016/j.applthermaleng.2018.07.004.

Zhang, S., Zhou, T., Ahuja, N., Refai-Ahmed, G., Zhu, Y., Chen, G., Wang, Z., Song, W., 2014. Real time thermal management controller for data center. In: Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. ITherm), pp. 1346–1353. https://doi.org/10.1109/ITHERM.2014.6892436.