



LEEDS
BECKETT
UNIVERSITY

Citation:

Grishina, A and Chinnici, M and Kor, AL and Rondeau, E and Georges, JP (2020) A machine learning solution for data center thermal characteristics analysis. *Energies*, 13 (17). ISSN 1996-1073
DOI: <https://doi.org/10.3390/en13174378>

Link to Leeds Beckett Repository record:

<https://eprints.leedsbeckett.ac.uk/id/eprint/7292/>

Document Version:

Article (Published Version)

Creative Commons: Attribution 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.



LEEDS
BECKETT
UNIVERSITY

Citation:

Grishina, A and Chinnici, M and Kor, AL and Rondeau, E and Georges, JP (2020) A machine learning solution for data center thermal characteristics analysis. *Energies*, 13 (17). ISSN 1996-1073
DOI: <https://doi.org/10.3390/en13174378>

Link to Leeds Beckett Repository record:

<http://eprints.leedsbeckett.ac.uk/id/eprint/7292/>

Document Version:

Article

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.

Article

A Machine Learning Solution for Data Center Thermal Characteristics Analysis

Anastasiia Grishina ¹, Marta Chinnici ^{2,*}, Ah-Lian Kor ³, Eric Rondeau ⁴ and Jean-Philippe Georges ⁴

¹ Department of Mathematics and Computer Science, Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands; a.grishina@tue.nl

² Department of Energy Technologies and Renewable Sources, ICT Division, ENEA Casaccia Research Center, 00123 Rome, Italy

³ School of Built Environment, Engineering and Computing, Leeds Beckett University, Leeds LS1 3HE, UK; a.kor@leedsbeckett.ac.uk

⁴ Université de Lorraine, Centre National de la Recherche Scientifique (CNRS), Centre de Recherche en Automatique de Nancy (CRAN), F-54000 Nancy, France; eric.rondeau@univ-lorraine.fr (E.R.); jean-philippe.georges@univ-lorraine.fr (J.-P.G.)

* Correspondence: marta.chinnici@enea.it; Tel.: +39-349-3553790

Received: 13 July 2020; Accepted: 20 August 2020; Published: 25 August 2020



Abstract: The energy efficiency of Data Center (DC) operations heavily relies on a DC ambient temperature as well as its IT and cooling systems performance. A reliable and efficient cooling system is necessary to produce a persistent flow of cold air to cool servers that are subjected to constantly increasing computational load due to the advent of smart cloud-based applications. Consequently, the increased demand for computing power will inadvertently increase server waste heat creation in data centers. To improve a DC thermal profile which could undeniably influence energy efficiency and reliability of IT equipment, it is imperative to explore the thermal characteristics analysis of an IT room. This work encompasses the employment of an unsupervised machine learning technique for uncovering weaknesses of a DC cooling system based on real DC monitoring thermal data. The findings of the analysis result in the identification of areas for thermal management and cooling improvement that further feeds into DC recommendations. With the aim to identify overheated zones in a DC IT room and corresponding servers, we analyzed thermal characteristics of the IT room. Experimental dataset includes measurements of ambient air temperature in the hot aisle of the IT room in ENEA Portici research center hosting the CRESCO6 computing cluster. We use machine learning clustering techniques to identify overheated locations and categorize computing nodes based on surrounding air temperature ranges abstracted from the data. This work employs the principles and approaches replicable for the analysis of thermal characteristics of any DC, thereby fostering transferability. This paper demonstrates how best practices and guidelines could be applied for thermal analysis and profiling of a commercial DC based on real thermal monitoring data.

Keywords: data center; thermal characteristics analysis; machine learning; energy efficiency; clustering; unsupervised learning

1. Introduction

Considerable efforts have been made by Data Centers in terms of their energy efficiency, reliability and sustainable operation over the past decade. A continuous increase in computing and power demands has spurred DCs to respond and upgrade their facilities in terms of size and stability [1,2]. A rapid growth of the information technology (IT) industry, advent of IoT, and AI technologies requires

an exponentially expanding amount of data to be stored and processed. Consequently, smart DC management is on the rise to meet this demand. If a data center experiences a system failure or outage, it becomes challenging to ensure a stable and continuous provision of IT services, particularly for smart businesses, social media, etc. If such a situation occurs on a large scale, it could be detrimental to the businesses and public sectors that rely on DC services, for example, health systems, manufacturing, entertainment, etc. In other words, a data center has emerged as a mission-critical infrastructure [3] to the survival of public and business sectors enabled by smart technologies. Therefore, it warrants an exceptional necessity for the backup system management and uninterruptible power supply (UPS) systems so that computing system stability can be maintained even in emergency situations. Overall, thermal management involves the reduction of excess energy consumption by cooling systems, servers, and their internal fans. This encompasses the compliance of the IT room environment to the requirements stipulated in IT equipment specifications and standards that ensure better reliability, availability, and overall improved IT equipment performance.

The mission-critical facility management for the stable operation of a DC leads to huge cost increases, and careful reviews must be performed starting from the initial planning stage [4,5]. Moreover, IT servers require uninterruptible supplies of not only power but also cooling [6,7]. Undeniably, a significant increase in power density has led to a greater cooling challenge [8]. For this purpose, the design of a central cooling system in a liquid cooling architecture includes cooling buffer tanks. This design provides chilled water supply in emergency situations. During cooling system outages, the interruption of chillers triggers activation of emergency power and cooling supplies to restore cooling services. However, such emergency situations are infrequent on the scale of a DC life-cycle. In addition, recent the development of IT equipment has increased servers' tolerance to various operational challenging conditions compared to that in the past. Consequently, the operating times and capacities of chilled water storage tanks could be considerably reduced. The same is true for energy and water consumption of the liquid cooling system. Undeniably, it is imperative to explore the thermal specifications of the DC IT equipment. They are expressed as (but not limited to) different admissible ranges for temperature, humidity, periods of overheating before automatic power off, etc.

Given that IT devices might have different recommended specifications for their operation, maintaining healthy operational conditions is a complex task. Undoubtedly, covert thermal factors tend to affect the health of IT and power equipment systems in a negative way. Such factors comprise recirculation, bypass, and (partial) server rack overheating and stable operation is critical for a DC [9]. For example, in the case where an IT room is divided into cold and hot aisles, ineffective partitioning of the aisles (e.g., poor physical separation of the aisles) may result in leaked air causing recirculation of hot air (note: this is the mixing of hot and cold air) or cold air bypass (note: this happens if the cold air passes by the server and does not cool it properly - this occurs when the speed of air-flow is too high or the cold air is unevenly directed at the hot servers) [10]. Consequently, such emerging challenges have to be addressed to effect thermal conditions optimization within a DC facility. Undeniably, an increase in ambient temperature could lead to an increase in power usage of IT equipment that could lead to hardware degradation [8]. Thus, it is necessary to address the issue of server waste heat dissipation with the ultimate goal to attain an even thermal distribution within a premise. This is possible by taking appropriate measures to prevent the emergence of heat islands that result in individual server overheating [11].

This work explores IT room thermal characteristics using data mining techniques for the purpose of relevant and essential knowledge discovery. The primary goal is to use an unsupervised machine learning technique to uncover inadequacies in the DC cooling system based on real monitored thermal data. Analysis in this research leads to the identification of areas for improved thermal management and cooling that will feed into DC recommendations. The proposed methodology includes statistical analysis of IT room thermal characteristics and identification of individual servers that are contributors to hotspots (i.e., overheated areas) [12]. These areas emerge when individual servers do not receive

adequate cooling. The reliability of the analysis has been enhanced due to the availability of a dataset of ambient air temperature in the hot aisle of an ENEA Portici CRESCO6 computing cluster [9].

In brief, clustering techniques have been used for hotspots localization as well as server categorization based on surrounding air temperature ranges. The principles and approaches employed in this work are replicable for thermal analysis of DC servers and thus, foster transferability. This work showcases the applicability of best practices and guidelines in the context of a real commercial DC that transcends the typical set of existing metrics for DC thermal characteristics assessment. Overall, this paper aims to raise DC thermal awareness and formulate recommendations for enhanced thermal management. This aim is supported by the following list of research objectives [9]:

- RO.1.To identify a clustering (grouping) algorithm that is appropriate for the purpose of this research;
- RO.2.To determine the criteria for feature selection in the analysis of DC IT room thermal characteristics;
- RO.3.To determine the optimal number of clusters for the analysis of thermal characteristics;
- RO.4.To perform sequential clustering and interpretation of results for repeated time series of air temperature measurements;
- RO.5.To identify servers that most frequently occur in cold or hot air temperature ranges (and clusters);
- RO.6.To provide recommendations related to IT room thermal management with the aim of appropriately addressing servers overheating issue.

In summary, typical identification of hotspots (and rack or cluster failure) in data centers is through the deployment of heatmaps (e.g., by Facebook in [13]) which comprise discrete scene snapshots of the premise (and not individual compute node) under study. However, our novel contribution is the employment of an Artificial Intelligence approach (i.e., Machine Learning clustering technique) based on ‘continuous’ data center environmental monitoring data, which will provide insights into the qualitative measure of the ‘duration’ a particular compute node is in a particular temperature range (i.e., low, medium, hot). Additionally, our proposed approach has an edge over typical heatmaps due to the low level granularity (i.e., with more details) information it provides for each node rather than aggregated information of the nodes (in clusters), so that targeted as well as effective corrective action or interventions can be appropriately taken by data center owners.

The remainder of the paper is organized as follows: Section 2 focuses on the background of the problem and related work; Section 3 presents the methodology used in this work; Section 4 provides discussion of experimental results and analysis; Section 5 concludes the paper with a summary and recommendation for future work.

2. Background and Related Work

In recent years, a number of theoretical and practical studies have been conducted on DC thermal management to better understand ways to mitigate inefficiencies of the cooling systems. This includes DC thermal and energy performance evaluation and load distribution optimization. Ineffective thermal management could be the primary contributor to DC IT infrastructure unreliability due to hardware degradation.

Existing DC-related thermal management research highlights the primary challenges of cooling systems in high power density DCs [14]; recommends a list of thermal management strategies based on energy consumption awareness [2,15]; explores the effect of different cooling approaches on power usage effectiveness (PUE) using direct air with a spray system that evaporates water to cool and humidify incoming air [16]; investigates the thermal performance of air-cooled data centers with raised and non-raised floor configurations [17]; studies various thermofluid mechanisms using cooling performance metrics [18]; proposes thermal models for joint cooling and workload management [19], while other strains of research explore thermal-aware job scheduling, dynamic resource provisioning, and cooling [20]. In addition, server-related thermal information, such as inlet/outlet air temperature and air mover speed, is utilized to create thermal and power maps with the ultimate goal to monitor the real-time status of a DC [21].

A majority of previously listed research work focuses on simulations or numerical modeling [2,16–20] as well as on empirical studies involving R&D or small-scale data centers [16,21]. Thus, there is a need for more empirical research involving real thermal-related data for large scale data centers. Undeniably, it is tremendously beneficial to identify hotspots and the air dynamics (particularly its negative effects) within a DC IT room. Such useful evidence-based information will help DC operators improve DC thermal management and ensure uninterrupted steady computing system operations. This is made possible when affected servers continue to perform graceful degradation-related computations or enter the ‘switch off’ mode once the temperature threshold is breached. Thermal management could be improved in a number of ways based on evidence-based analysis. For example, some corrective actions could be: identify cold air leakages and erect isolating panes; adjust speed, volume, and direction of the cold air stream; apply free cooling wherever possible or adjust the humidity levels. An exhaustive guideline for DC thermal management improvement can be found in [7].

A crucial step forward in DC thermal management related research could be adherence to the recommended thermal management framework [22] at varying DC granularity levels. As a part of the framework, thermal metrics have been created by research and enterprise DC communities [10]. Employment of the metrics aims to reveal the underlying causes of thermal-related issues within a DC IT room and to assess the overall thermal conditions of the room. A recently proposed holistic DC assessment method is based on biomimicry [23]. This integrates data on energy consumption for powering and cooling ICT equipment.

This paper is an extension of the previous authors’ work [10,11,24–28], which focus on real DC thermal monitoring data. In detail, this current research focuses on the analysis of DC IT room thermal characteristics to uncover ways to render a more effective cooling system as well as explore possibilities to employ machine learning techniques to address this issue. Appropriate data analytics techniques have been applied on real server-level sensor data to identify potential risks caused by the possible existence of negative covert physical processes related to the cooling strategy [2]. In summary, this work is based on the analysis of DC thermal characteristics using machine learning (ML) techniques. ML has been generally employed for virtual machines allocation, global infrastructure management, prediction of electricity consumption, and availability of renewable energy [29]. Thus far, there is work on ML for thermal characteristics assessment and weather conditions prediction, but only limited available work on thermal management. Typically, Computational Fluid Dynamics (CFD) techniques have been employed for the exploration of DC thermal management. Their drawbacks are high computational power and memory requirements. Therefore, the added value of this research is an evidence-based recommendation for a cooling system for more targeted temperature management through thermal characteristics analysis for localization of overheated areas in the DC IT room.

3. Methodology

This section discusses the thermal characteristics analysis of an ENEA R.C. Portici cluster CRESCO 6. An ML clustering technique was chosen for a more in-depth analysis of overheated servers’ localization based on an available dataset of CRESCO6 server temperature measurements. The terms “server” and “node” are used interchangeably in this work, while “hotspot” is utilized to indicate an overheated area next to a server in the IT room and results in an overheated or undercooled server. The drawback of a typical analysis of temperature measurements is that it could not locate the specific nodes which cause rack overheating. Hence, to address this issue, we have applied node clustering to localize potentially harmful hotspots. To identify overheated areas in the CRESCO6 group of nodes, we sequentially grouped the nodes into clusters characterized by higher or lower surrounding air temperature [9]. The term “group of nodes” stands for the DC “cluster” (note that this term is not used to avoid its confusion with the term “cluster”, which is the outcome of running an ML clustering algorithm).

3.1. Cluster and Dataset Description

Thermal analysis was based on monitoring data of the CRESCO6 cluster in the premises of ENEA-Portici Research Center (R.C.). Data collected were cluster power consumption of IT equipment (servers) and measurements of ambient air temperature. This cluster has been up and running since May 2018. It is used to augment the computing resources of the CRESCO4 system, already installed and still operating in the Portici Research Center. The reason for the augmentation is due to the rise in demand for analytic activities. Thus, with the addition of the cluster CRESCO6, the overall computing capability of ENEA R.C. has increased up to seven-fold. The cluster comprises 418 Lenovo nodes housed in a total of 5 racks. Each node includes two CPUs, each with 24 cores (with a total of 20,064 cores). This pool of resources is aimed to support Research and Development activities in ENEA Research Center [9].

Measurements of ambient air temperature were recorded during two phases. The first phase was during cluster set up and performance tuning (subject to thermal control strategies) and other indicator specifications during the months of May–July 2018. Subsequently, end-users were allowed to submit and run their jobs and relevant parameters had been monitored and measured for approximately 9 months (September 2018–February 2019). The measurements were paused in August 2018 as shown in Figure 1 [9]. Data collected encompassed all 216 nodes, out of which 214–215 nodes were consistently monitored, and the other 1–2 nodes had missing values or were turned off. The monitoring system consisted of an energy meter, a power meter of CPU, RAM and computing power utilization of every node, and CPU temperature for both processing units of each node with thermal sensors installed inside the servers, at the inlet and exhaust air locations in cold and hot aisles respectively (i.e., placed in the front and rear parts of every node).




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

3.2. Data Analytics

Variation of the air temperature was captured and analyzed for different areas of the IT room. The variability of thermal data and uncertainty in defining temperature thresholds for overheated areas has provided a justification for the use of an unsupervised learning technique. Hence, a k-means algorithm was employed to address the limitations of typical statistical techniques and cluster the servers according to their surrounding air temperature. Silhouette metric and within-cluster sum of squares were used to first determine the number of clusters. Available thermal characteristics (i.e., exhaust air temperature, readings of CPU temperature) served as inputs to the k-means algorithm. A set of surrounding air temperature measurements for the nodes was clustered the same number of times as the number of measurements taken for the batch of all the nodes. Subsequently, the resulting series of cluster labels were intersected to unravel nodes (distinguishable by their IDs) that frequently occurred in the high-temperature cluster.

An adapted data lifecycle methodology was employed for this work, as depicted in Figure 2. The methodology comprises several data preprocessing steps, data analysis, followed by interpretation of the results and their exploitation in the form of recommendations for the DC under consideration [9]. A detailed discussion of all data analytics stages represented in Figure 2 are found in the ensuing section.

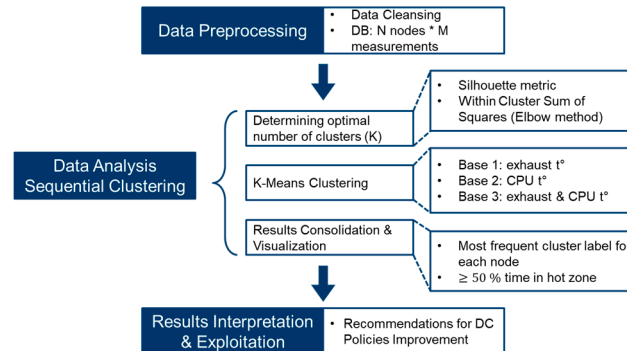


Figure 2. Data analysis lifecycle methodology adapted to sequential clustering of DC servers based on their thermal characteristics.

The data preprocessing step consisted of data cleansing of zero and missing values and formatting. The dataset was organized as shown in Table 1. This table summarizes the results of monitoring of the overall number of nodes in a computing cluster, N . In addition, data preprocessing involved timestamps formatting for further exploitation. In detail, the system was configured so that the monitoring system recorded the thermal data for every node with an interval of around 15 min, including a slight latency between each pair of consecutive readings of temperature sensors around the nodes. The readings resulted in a set of N rows with the information for every node ID.

Table 1. Dataset using for clustering analysis.

Time Label	Real Time of Measurement	Node ID	Inlet T (°C)	Exhaust T (°C)	CPU 1 T (°C)	CPU 2 T (°C)	Cluster Label
t_1	$t_1 + t_{1n_1}$	n_1	T_{in11}	T_{exh11}	T_{CPU111}	T_{CPU211}	$C_1^{base_{range}}$
	$t_1 + t_{1n_N}$	n_N	T_{in1N}	T_{exh1N}	T_{CPU11N}	$T_{CPU2s1N}$	$C_1^{base_{range}}$
t_2	$t_2 + t_{2n_1}$	n_1	T_{in21}	T_{exh21}	T_{CPU121}	T_{CPU221}	$C_2^{base_{range}}$
	$t_2 + t_{2n_N}$	n_N	T_{in2N}	T_{exh2N}	T_{CPU12N}	$T_{CPU2s2N}$	$C_2^{base_{range}}$

The data analysis step included several substages. The sequential clustering substage encompassed the investigation into the optimal number of clusters followed by server clustering into groups (with three possible levels: low, medium, and high) of the surrounding air temperature. The results were further consolidated to ascribe final cluster labels for each server (i.e., low, medium, or high temperature) based on the frequency of occurrences for each node label in the sequence of results [9]. Clustering was performed M times, where M is the overall number of time labels at which measurements were taken for all cluster nodes. Each new set of monitoring system readings was labeled with a time label t_1, \dots, t_M . The exact timestamp for the extracted information was marked with $t_i + t_{in_j}$ for every node j . Depending on the available dataset, a number of relevant features described the thermal state of every node and their different combinations could be used as a basis for clustering (RO.2 will be more considered in detail in Section 4). Thus, we introduce *base* in the last column of Table 1 which denotes the basis for clustering, i.e., a combination of temperature measurements taken as k-means input (see Results and Discussions for details). The indicator *base* also corresponds to the temperature of the cluster centroid [9].

In this work, k-means was chosen as a clustering algorithm due to the following reasons (RO.1):

1. The number of features used for clustering was small. Therefore, the formulated clustering problem was simple and did not require complex algorithms;
2. K-means has linear computational complexity and is fast to use for the problem in question. While the formulation of the problem is simple, it requires several thousands of repetitions of clustering for each set of N nodes. From this point of view, the speed of the algorithm becomes an influential factor;
3. K-means has a weak point, namely the random choice of initial centroids, which could lead to different results when different random generators are used. This does not pose any issue in this use case since the nodes are clustered several times based on sets of measurements taken at different timestamps and minor differences brought by the randomness are mitigated by the repetition of the clustering procedure.

The number of clusters, K , is an unknown parameter that also defines the number of ranges for $C_{i \text{ base range}}$. It is estimated for each of the three combinations or *bases* using two metrics separately: the average silhouette coefficient and within cluster sum of squares (WCSS) metric [9,30,31] (RO.3). The application of these two indices to derive the suitable number of thermal ranges or clusters is shown in Appendix A. In brief, the silhouette coefficient was computed for each clustered sample of size N and showed the degree of isolation for the clusters, thus, indicating the quality of clustering. The +1 value of silhouette index for a specific number of clusters, K , indicated the high density of clusters, -1 showed incorrect clustering, and 0 stood for overlapping clusters. Therefore, we focused on local maxima of this coefficient. WCSS was used in the Elbow method of determining the number of clusters and was used here to support the decision obtained from the silhouette coefficient estimation. It measured the compactness of clusters, and the optimal value of K was the one that resulted in the “turning point” or the “elbow” of the WCSS (K) graph. In other words, if we increase the number of clusters after reaching the elbow point, it does not result in significant improvement of clusters compactness. Although it could be argued that other features could be additionally used for determining the number of clusters, the combination of the two aforementioned methods had converged on the same values of K for our chosen *bases*, which was assumed to be sufficient for this current research.

Once we obtained the optimal parameter K , we performed the clustering procedure for the chosen *bases*. For a fixed *base*, the sequence of cluster labels was examined for every node. Based on the frequency of occurrences for each cluster label (low, medium, or hot air temperature labels), the node was ascribed the final most frequently occurring cluster label $C_{\text{base range}}$ in the sequence and was assigned to the corresponding set of nodes as $N_{\text{base range}}$. Furthermore, we took the intersection of the sets of nodes in the hot air range for every *base*. Thus, we obtained the IDs of the nodes that were most frequently labeled as the nodes in “danger” or hot areas of the IT rack by three clustering algorithms: $N_{\text{hot}} = \cap_{\text{bases}} \{N_{\text{base hot range}}\}$ (RO.4). In the following section, we will discuss the results of k-means sequential clustering and identify the nodes that occurred in the overheated areas more frequently than others.

4. Results and Discussions

For every combination of measured thermal data, the results of servers clustering into cold, medium, and hot temperature ranges had been further analyzed to calculate the frequency of occurrences of each node in each cluster and determined their final frequency label (i.e., cluster label or temperature range). These labels were further intersected with labels obtained for different bases. Each set of $N = 216$ nodes was clustered at once, followed by temperature-based clustering for the same set of nodes using measurements taken at the next timestamp. This process is referred to here as sequential clustering. The indicator *base* of the input temperature data was used in three possible combinations of available thermal data: exhaust air (*base 1*), CPU (*base 2*), and exhaust air and CPU temperature measurements (*base 3*) (RO.2).

The dataset contained $M = 15,569$ sets of temperature measurements. Each M sets consisted of 216 node-level temperature measurements of the sensor data. The sensors were installed in different

locations with respect to the node: in the front (inlet), rear (exhaust) of every node, and two sensors inside each node (CPU temperature). The optimal number of clusters K was influenced by the *base* chosen for clustering. Using the silhouette metric and WCSS, we obtained the optimal number of K for *bases* 1–3 (exhaust, CPU, and exhaust and CPU measurements) and the K values equaled to three, five, and three clusters, respectively [9].

During sequential clustering, each node was labeled with a particular temperature range cluster. Since clustering was repeated for each set of measurements grouped by time label, every node was repeatedly clustered several times and tagged with different labels while the algorithm was in progress (RO.4). Figure 3a–c shows the results of sequential clustering, where one set of three or five vertically aligned graphs represents the result for sequential clustering using one input *base*. In detail, Figure 2 compares how frequently each of the 216 nodes is labeled with low, medium, or hot temperature range. Every vertical graph corresponds to the proportion of occurrences (in %) of low, medium, or hot temperature labels for one node. This figure indirectly implies the incidence rate, “duration”, or tendency of a particular node experiencing a certain temperature range (see legend in Figure 3a–c). Here, the medium temperature range label most frequently occurred for the majority of the nodes for all cluster *bases*. We also observe that some nodes remain in the hot range for more than 50% of clustering iterations. This information could alternatively be obtained if the temperature levels (lower and upper bounds of low, medium, and hot temperature ranges) were preset by an expert, but for this research, this estimation is not available. Therefore, the consideration of the clustering results is crucial for the DC in a situation when temperature ranges are unavailable. Figure 2 provides a means for asynchronous assessment of the thermal state of the IT room and unravels relevant thermal trends.




Figure 3. The ratio of nodes labeled by different temperature ranges (low, medium, and hot) based on different thermal data input (*base*): (a) Exhaust temperature, (b) Exhaust and CPU temperature, (c) CPU temperature. The legend includes coordinates of cluster centroids.

In cases where nodes remain in the hot range for a prolonged period or frequently fall in this range, it implies that they are overheated (and the cooling appears to be ineffective). Consequently, this could cause hardware degradation where the nodes have reduced reliability and accessibility as they automatically switch to lower power mode when overheated. Therefore, we continued with the analysis to identify the actual node IDs that had most frequently been clustered within the hot air temperature ranges. Table 2 provides an insight into the ratio of nodes with the highest frequency of occurrences in cold, medium, or hot air temperature range (RO.5). Depending on the cluster base, 50% to 86% of all nodes had the highest frequency of occurrence in the medium range. The hot range encompassed 11%–37% of all nodes, and only 0.5%–4% had been clustered within the cold range.

Table 2. Ratio of cluster sizes and intersection of node labels from three thermal data combinations (*bases*).

Cluster Type Ratio%	$N_{exh_{cold}}$ 2.8	$N_{exh_{med}}$ 86.0	$N_{exh_{hot}}$ 11.2		
Cluster Type Ratio%	$N_{CPU_{cold}}$ 4.2	$N_{CPU_{med1}}$ 20.0	$N_{CPU_{med2}}$ 28.4	$N_{exh_{hot1}}$ 31.2	$N_{exh_{hot2}}$ 16.2
Cluster Type Ratio%	$N_{exh_{CPU_{cold}}}$ 2.0	$N_{exh_{CPU_{med}}}$ 63.0		$N_{exh_{CPU_{hot}}}$ 35.0	
Cluster Type Ratio%	N_{cold} 0.5	N_{med} 40.0		N_{hot} 8.0	
Hot Range Node ID	-	-		30, 31, 32, 45, 46, 48, 68, 79, 94, 96, 105, 117, 118, 120, 182, 183, 189, 198	

Finally, the results of sequential clustering with three bases were intersected for cross-validation purposes. These results in the intersection of nodes were labeled as cold, medium, and hot surrounding air range. One node (equal to 0.5% of nodes) was labeled as the cold air temperature range for all three *bases*. The cluster characterized by the medium air temperature range had the largest intersection among the *bases* (i.e., more than 90%), while 8% (or 18 nodes) were binned in the hot air temperature range most frequently using all three *bases*. The highlight of this exploration is that we were able to identify the IDs of the most frequently overheated nodes. DC operators could further exploit this evidence-based information to improve thermal conditions in the cluster IT room. Possible corrective actions to mitigate overheating of the localized nodes could be to improve the existing natural convection cooling by directing the cold air to the hottest nodes. In addition, DC operators could update the load scheduling and decrease the workload of the identified nodes to indirectly prevent overheating (RO.6). In summary, DC operators could improve resource allocation policies and cooling strategies to effectively address this issue.

This current paper has contributed to thermal characteristics awareness for a real DC cluster and addressed the issue of servers overheating. This has two positive effects in terms of sustainability. Firstly, local overheating could be considered as an IT room thermal design pitfall. It leads to a high risk of hardware degradation for servers that are frequently and/or for long time exposed to high surrounding air temperature. From this perspective, the localization of hot regions of the IT room performed in this study (via a ML technique) is crucial for providing a better overview of thermal distribution around the servers which could be fed into better thermal control and management strategies. In other words, future thermal management improvements could be aligned to the direction provided in this study with the aim of mitigating the abovementioned risk. Secondly, a clustering technique used in this phase requires less computational resources compared to heatmaps, computational fluid dynamics modeling and/or simulations performed on existing simulation packages [9]. Therefore, this work evidences the benefit of less computationally intensive analytical techniques (in yielding sufficient information) to incentivize improvement of thermal conditions (through even thermal distribution) in data centers. In summary, conclusions that could be drawn from this research are that a majority of the nodes were

located in medium and hot air temperature ranges. Joint results of three clustering algorithms had shown that that 8% of cluster servers were most frequently characterized as having hot surrounding air temperature. Based on this evidence, we have formulated a list of recommendations (see subsequent section) to address the problem of repeated or prolonged overheating of servers (RO.6).

5. Conclusions and Future Work

Analysis of IT and cooling systems is necessary for the investigation of DC operations-related energy efficiency. A reliable cooling system is essential to produce a persistent flow of cold air to cool servers that could be overheated due to an increasing demand in computation-intensive applications. To reiterate, Patterson [8] has maintained there is an impact of DC ambient temperature on energy consumption of IT equipment and systems. However, in this paper, the focus is on thermal characteristics analysis of an IT room. The research methodology discussed in this paper includes statistical analysis of IT room thermal characteristics and the identification of individual servers that frequently occur in the overheated areas of the IT room (using a machine learning algorithm). Clustering techniques are used for hotspot localization as well as categorization of nodes (i.e., servers) based on surrounding air temperature ranges. This methodology has been applied to an available dataset with thermal characteristics of an ENEA Portici CRESCO6 computing cluster. In summary, this paper has presented a proposed methodology for IT room thermal characteristics assessment of an air-cooled DC cluster located in a geographical region where free air cooling is unavailable. The steps involved for evidence-based targeted temperature management (to be recommended for air-cooled DCs) are as follows:

1. Explore the effectiveness of the cooling system by firstly uncovering nodes with hot range IDs (e.g., change direction, volume, speed of cooling air). Additionally, directional cooling could be recommended (e.g., spot cooling to cool overheated nodes). Next, unravel covert factors that lead to nodes' repetitive overheating (e.g., location next to the PDUs that have higher allowable temperature ranges);
2. Revise cluster load scheduling so that these frequently overheated servers are not overloaded in the future (note: this is to enable an even thermal distribution within the IT room. See [11] for details). In other words, it is recommended to formulate a resource allocation policy for the purpose of a more even thermal distribution of ambient air temperature;
3. Perform continuous environmental monitoring of the IT room and evaluate the effectiveness of recommended actions and their influence on the ambient temperature.

To reiterate, the approaches covered in this work are transferrable for thermal characteristics analysis in any air-cooled DC context enabled with a thermal monitoring system. This study illustrates the applicability of the best practices and guidelines to a real DC and uses an ML approach to perform IT room thermal characteristics assessment. This work could be extended by incorporating an integrated thermal management with existing energy efficiency policies-related research (e.g., energy awareness [15]; job scheduling using AI [32], temporal-based job scheduling [33], work-load aware scheduling [34], and queue theory [35]; resource utilization [36] of multiple applications using annealing and particle swarm optimization [37]). Another direction that could be taken could be the energy efficiency policies and waste heat utilization [26].

Author Contributions: Conceptualization A.G., M.C., A.-L.K., E.R. and J.-P.G.; methodology, A.G., M.C., A.-L.K.; data analysis, A.G.; data curation, M.C.; writing—original draft preparation, A.G., M.C., A.-L.K. and J.-P.G.; writing—review and editing, A.G., M.C., A.-L.K.; project administration, E.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by EMJMD PERCCOM Project [38].

Acknowledgments: This work was supported by EMJMD PERCCOM Project [38]. During the work on the project, the author A.G. was affiliated with ENEA-R.C. Casaccia, Rome, Italy. Moreover, the authors extend their gratitude to the research HPC group at the ENEA-R.C. Portici for the discussion of ENEA-Data Center control and modeling.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

Several approaches are widely used by data scientists to identify the optimal number of clusters. However, it is worth noting that none of the approaches is considered accurate, instead they provide a suggestion for the number of clusters. The current study applies two indices: WCSS, also known as an elbow method, and average Silhouette Index [32]. We apply a systematic approach and calculate these metrics for different values of cluster numbers, K . Subsequently, an optimal value is chosen based on the values of the indices. WCSS characterizes the compactness of the cluster and is defined as follows:

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

where K is the number of clusters, C is a set of clusters (C_1, C_2, \dots, C_j), and μ_j represents a certain cluster sample mean. The target value of this metric should be a possible minimized value. Practically, this refers to a point where the value of the metric continues to decrease but at a significantly slower rate in comparison to a smaller number of clusters and is considered an optimal one. It is visually associated with an “elbow” of the graph. The justification for choosing an elbow point is that with the increase in the number of clusters, the metric decreases only slightly, while computations become increasingly more intensive.

The second method applied in this paper is the average silhouette. It is computed using average silhouette index over all data points (or cluster members). It estimates within-cluster consistency and should be maximized to achieve the effective cluster split. The formula for the Silhouette index is as follows:

$$s(i) = \frac{b(i) - a(i)}{\max\{b(i), a(i)\}}, \quad (2)$$

where, a is the mean distance between one cluster member and all other members of the same cluster. Parameter b is the distance between a cluster member and all other points in the nearest cluster. We depict the results of this metrics calculation in Figure A1. It shows the metrics for one-step of sequential clustering based on server exhaust air temperature. At $K = 3$, we observe a local maximum for silhouette index and an elbow point of WCSS graph.




Figure A1. (a) WCSS; (b) Average Silhouette Index. Both indices are computed for one step of the sequential clustering procedure based on the exhaust air temperature.

References

1. Hashem, I.A.T.; Chang, V.; Anuar, N.B.; Adewole, K.S.; Yaqoob, I.; Gani, A.; Ahmed, E.; Chiroma, H. The role of big data in smart city. *Int. J. Inf. Manag.* **2016**, *36*, 748–758. [CrossRef]
2. Zhang, K.; Zhang, Y.; Liu, J.; Niu, X. Recent advancements on thermal management and evaluation for data centers. *Appl. Therm. Eng.* **2018**, *142*, 215–231. [CrossRef]
3. Datacenter Knowledge. A Critical Look at Mission-Critical Infrastructure. 2018. Available online: <https://www.datacenterknowledge.com/industry-perspectives/critical-look-mission-critical-infrastructure> (accessed on 26 June 2020).
4. Hartmann, B.; Farkas, C. Energy efficient data centre infrastructure—Development of a power loss model. *Energy Build.* **2016**, *127*, 692–699. [CrossRef]
5. He, Z.; Ding, T.; Liu, Y.; Li, Z. Analysis of a district heating system using waste heat in a distributed cooling data center. *Appl. Therm. Eng.* **2018**, *141*, 1131–1140. [CrossRef]
6. Nadjahi, C.; Louahlia, H.; Lemasson, S. A review of thermal management and innovative cooling strategies for data center. *Sustain. Comput. Inform. Syst.* **2018**, *19*, 14–28. [CrossRef]
7. AT Committee. Data Center Power Equipment Thermal Guidelines and Best Practices Whitepaper. ASHRAE, Tech. Rep., 2016. Available online: https://tc0909.ashraetcs.org/documents/ASHRAE_TC0909_Power_White_Paper_22_June_2016_REVISED.pdf (accessed on 6 June 2019).
8. Patterson, M.K. The effect of data center temperature on energy efficiency. In Proceedings of the 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Orlando, FL, USA, 28–31 May 2008; pp. 1167–1174. [CrossRef]
9. Grishina, A. Data Center Energy Efficiency Assessment Based on Real Data Analysis. Unpublished PERCCOM Masters Dissertation. 2019.
10. Capozzoli, A.; Serale, G.; Liuzzo, L.; Chinnici, M. Thermal metrics for data centers: A critical review. *Energy Procedia* **2014**, *62*, 391–400. [CrossRef]
11. De Chiara, D.; Chinnici, M.; Kor, A.-L. Data mining for big dataset-related thermal analysis of high performance (HPC) data center. In *International Conference on Computational Science*; Springer: Cham, NY, USA, 2020; pp. 367–381.
12. Chinnici, M.; Capozzoli, A.; Serale, G. Measuring energy efficiency in data centers. In *Pervasive Computing: Next Generation Platforms for Intelligent Data Collection*; Dobre, C., Xhafa, F., Eds.; Morgan Kaufmann: Burlington, MA, USA, 2016; Chapter 10; pp. 299–351. ISBN 9780128037027.
13. Infoworld. Facebook Heat Maps Pinpoint Data Center Trouble Spots. 2012. Available online: <https://www.infoworld.com/article/2615039/facebook-heat-maps-pinpoint-data-center-trouble-spots.html> (accessed on 20 June 2020).
14. Bash, C.E.; Patel, C.D.; Sharma, R. Efficient thermal management of data centers—Immediate and long-term research needs. *HVAC&R Res.* **2003**, *9*, 137–152. [CrossRef]
15. Fernández-Cerero, D.; Fernández-Montes, A.; Velasco, F.P. Productive Efficiency of Energy-Aware Data Centers. *Energies* **2018**, *11*, 2053. [CrossRef]
16. Fredriksson, S.; Gustafsson, J.; Olsson, D.; Sarkinen, J.; Beresford, A.; Kaufeler, M.; Minde, T.B.; Summers, J. Integrated thermal management of a 150 kW pilot Open Compute Project style data center. In Proceedings of the 2019 IEEE 17th International Conference on Industrial Informatics (INDIN), Helsinki, Finland, 22–25 July 2019; Volume 1, pp. 1443–1450.
17. 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 Transf. Eng.* **2013**, *35*, 384–397. [CrossRef]
18. Schmidt, R.R.; Cruz, E.E.; Iyengar, M. Challenges of data center thermal management. *IBM J. Res. Dev.* **2005**, *49*, 709–723. [CrossRef]
19. MirhoseiniNejad, S.; Moazamigoodarzi, H.; Badawy, G.; Down, D.G. Joint data center cooling and workload management: A thermal-aware approach. *Future Gener. Comput. Syst.* **2020**, *104*, 174–186. [CrossRef]
20. Fang, Q.; Wang, J.; Gong, Q.; Song, M.-X. Thermal-aware energy management of an HPC data center via two-time-scale control. *IEEE Trans. Ind. Inform.* **2017**, *13*, 2260–2269. [CrossRef]
21. Zhang, S.; Zhou, T.; Ahuja, N.; Refai-Ahmed, G.; Zhu, Y.; Chen, G.; Wang, Z.; Song, W.; Ahuja, N. Real time thermal management controller for data center. In Proceedings of the Fourteenth Intersociety Conference on

- Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 27 May 2014; pp. 1346–1353.
22. Sharma, R.; Bash, C.; Patel, C.; Friedrich, R.; Chase, J.S. Balance of Power: Dynamic Thermal Management for Internet Data Centers. *IEEE Internet Comput.* **2005**, *9*, 42–49. [[CrossRef](#)]
 23. Kubler, S.; Rondeau, E.; Georges, J.P.; Mutua, P.L.; Chinnici, M. Benefit-cost model for comparing data center performance from a biomimicry perspective. *J. Clean. Prod.* **2019**, *231*, 817–834. [[CrossRef](#)]
 24. Capozzoli, A.; Chinnici, M.; Perino, M.; Serale, G. Review on performance metrics for energy efficiency in data center: The role of thermal management. *Lect. Notes Comput. Sci.* **2015**, *8945*, 135–151.
 25. Grishina, A.; Chinnici, M.; De Chiara, D.; Guarnieri, G.; Kor, A.-L.; Rondeau, E.; Georges, J.-P. DC Energy Data Measurement and Analysis for Productivity and Waste Energy Assessment. In Proceedings of the 2018 IEEE International Conference on Computational Science and Engineering (CSE), Bucharest, Romania, 29–31 October 2018; IEEE: Piscataway, NJ, USA, 2018; pp. 1–11, ISBN 978-1-5386-7649-3.
 26. Koronen, C.; Åhman, M.; Nilsson, L.J. Data centres in future European energy systems—Energy efficiency, integration and policy. *Energy Effic.* **2019**, *13*, 129–144. [[CrossRef](#)]
 27. 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*; Springer: Dubrovnik, Croatia, 2019; pp. 286–300.
 28. 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 *Big Data Platforms and Applications—Case Studies, Methods, Techniques, and Performance Evaluation*; Pop, F., Ed.; Springer: Berlin, Germany, 2020.
 29. Athavale, J.; Yoda, M.; Joshi, Y.K. Comparison of data driven modeling approaches for temperature prediction in data centers. *Int. J. Heat Mass Transf.* **2019**, *135*, 1039–1052. [[CrossRef](#)]
 30. Kaufman, L.; Rousseeuw, P.J. *Finding Groups in Data: An Introduction to Cluster Analysis*; John Wiley & Sons: Hoboken, NJ, USA, 2009.
 31. Kassambara, A. (Ed.) Determining the Optimal Number of Clusters: 3 Must Know Methods. Available online: <https://www.datanovia.com/en/lessons/determining-the-optimal-number-of-clusters-3-must-know-methods/> (accessed on 6 May 2019).
 32. Fernández-Cerero, D.; Fernández-Montes, A.; Ortega, J.A. Energy policies for data-center monolithic schedulers. *Expert Syst. Appl.* **2018**, *110*, 170–181. [[CrossRef](#)]
 33. Yuan, H.; Bi, J.; Tan, W.; Zhou, M.; Li, B.H.; Li, J. TTSA: An Effective Scheduling Approach for Delay Bounded Tasks in Hybrid Clouds. *IEEE Trans. Cybern.* **2017**, *47*, 3658–3668. [[CrossRef](#)]
 34. Yuan, H.; Bi, J.; Zhou, M.; Sedraoui, K. WARM: Workload-Aware Multi-Application Task Scheduling for Revenue Maximization in SDN-Based Cloud Data Center. *IEEE Access* **2018**, *6*, 645–657. [[CrossRef](#)]
 35. Fernández-Cerero, D.; Irizo, F.J.O.; Fernández-Montes, A.; Velasco, F.P. Bullfighting extreme scenarios in efficient hyper-scale cluster computing. In *Cluster Computing*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 1–17. [[CrossRef](#)]
 36. Fernández-Cerero, D.; Fernández-Montes, A.; Jakobik, A.; Kołodziej, J.; Toro, M. SCORE: Simulator for cloud optimization of resources and energy consumption. *Simul. Model. Pract. Theory* **2018**, *82*, 160–173. [[CrossRef](#)]
 37. Bi, J.; Yuan, H.; Tan, W.; Zhou, M.; Fan, Y.; Zhang, J.; Li, J. Application-Aware Dynamic Fine-Grained Resource Provisioning in a Virtualized Cloud Data Center. *IEEE Trans. Autom. Sci. Eng.* **2015**, *14*, 1172–1184. [[CrossRef](#)]
 38. 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. *J. Clean. Prod.* **2016**, *135*, 223–239. [[CrossRef](#)]

