

Thermal Management in Large Data Centres: Security Threats and Mitigation

Betty Saridou¹[0000–0001–6939–9900], Gueltoum Bendiab²,
Stavros N. Shiaeles², and Basil K. Papadopoulos¹

¹ Democritus University of Thrace, Xanthi, GR
dsaridou@civil.duth.gr, papadob@civil.duth.gr

² Cyber Security Research Group, University of Portsmouth, Portsmouth, UK
gueltoum.bendiab@port.ac.uk, sshiaeles@ieee.org

Abstract. Data centres are experiencing significant growth in their scale, especially, with the ever-increasing demand for cloud and IoT services. However, this rapid growth has raised numerous security issues and vulnerabilities; new types of strategic cyber-attacks are aimed at specific physical components of data centres that keep them operating. Attacks against temperature monitoring and cooling systems of data centres, also known as thermal attacks, can cause a complete meltdown and are generally considered difficult to address. In this paper, we focus on this issue by analysing the potential security threats to these systems and their impact on the overall data centre safety and performance. We also present current thermal anomaly detection methods and their limitations. Finally, we propose a hybrid method to prevent thermal attacks, that uses multi-variant anomaly detection and a fuzzy-based health factor to enhance data centre thermal awareness and security.

Keywords: Anomaly Detection · Security · Data Centre · Thermal Sensors · Cooling System · Thermal attacks

1 Introduction

Data Centres are experiencing unprecedented growth and will continue to scale operations to meet the ever-increasing service demands. A recent study by Gartner³ estimated that about 425 million new servers might be needed by 2020 to support 30 billion connected IoT devices around the world, while the average data centre will support more than 100,000 servers [30]. The average data warehouse size usually ranges from 100ft² to 400,000ft² [35] and operates on a high-power consumption. A plethora of security issues has arisen due to this rapid growth, causing data centres to become vulnerable to strategic cyber-attacks on the physical infrastructure vital to maintaining the uninterrupted operation of data centres [12, 13, 19, 38]. Power supply, cooling, temperature monitoring and even security systems can serve as entry points for attacks against data centre operators or companies using data centre services. Security experts warn that new

³ <https://www.gartner.com/en>

sophisticated malware such as Triton and Trisis [28] are very effective against power and HVAC (Heating, Ventilating, and Air Conditioning) systems, and therefore, put data centres' safety at risk. Moreover, reported cyber-incidents showed that cyber-attacks on such connected facilities can destroy thousands of servers by overheating the environmental atmosphere, or manipulating energy and temperature settings to cause a fire or an explosion incident [12, 13, 38].

Attacks against temperature monitoring and cooling systems, also known as thermal attacks, are considered dangerous and difficult to tackle. In this paper, we aim to address this issue by exploring the most important security threats against thermal management systems, as well as their overall impact on data centre security and performance. In this context, we describe a number of vulnerabilities that may be exploited by attackers who aim to alter environmental conditions inside a data centre. Based on a combination of thermal hardware requirements, thermal anomaly detection is considered a relevant approach for detecting abnormal behaviour of data centre indoor temperature. In this context, several approaches have been proposed for temperature measurement analysis and abnormal behaviour definition on various environments, including smart homes, health, and data centres. In this paper, we analyse the effectiveness and limitations of these approaches with respect to the protection against thermal attacks in data centres. Additionally, we denote the importance of intelligent solutions and present initial work on a hybrid framework that uses multi-variant anomaly detection as a preventive measure. Finally, to further enhance thermal awareness and security of data centres, we introduce the idea of a fuzzy-based health factor.

The remainder of this paper is organised as follows. In section 2, a high-level overview of temperature monitoring in large data centres is presented, along with their importance and benefits. Section 3 discusses the potential security threats on temperature monitoring systems and their impact on the overall data centre security and performance. In Section 4, we provide an overview of the existing solutions and their limitations. In Section 5, we present a conceptual model that combines multi-variate anomaly detection methods and fuzzy logic for protecting data centres against thermal attacks. Finally, Section 6 concludes this paper and outlines future work.

2 Thermal management in data centers

Data centres are large infrastructures composed of a building or a group of buildings dedicated to house computer systems and their services [13]. They usually host a large number of high-density servers, which run simultaneously and generate extensive amounts of heat per second. Thus, existing data centres are equipped with multiple cooling technologies to cool down servers, including chilled water systems, cold aisle/hot aisle design, Computer Room Air Conditioner (CRAC), liquid cooling and free air cooling [12, 32]. For efficient thermal management, they also use robust thermal design and temperature regulation and monitoring as heat countermeasures. In modern data centres, HVAC and




Fig. 1. Typical cooling system in data centers [11].

CRAC systems are the most widely used cooling solution for managing temperature, airflow and humidity. At the same time, continuous monitoring with Wireless Sensor Networks (WSNs) helps to create and maintain the energy-efficient cooling environment.

As shown in Figure 1, computer rooms are arranged into hot and cold aisles; server rack fronts face each other to create cold aisles due to the front-to-back heat produced by servers [32]. CRAC units, which are usually positioned at the end of hot-aisles or around the room perimeter (see Figure 1), push the cold air under the raised floor and through the cold aisle to absorb the heat produced by servers. The resulting hot air enters the hot aisle from the front side of the server racks and is returned to the CRAC devices to cool down by the chillers once more [12, 13, 32]. In order to avoid overheating, the inlet temperatures of servers are continuously monitored using a network of thermal sensors. If the inlet temperature exceeds a predefined threshold, parts of servers or even racks are forced to shut down to avoid permanent hardware damage [12].

In general, control and monitoring of heat diffusion inside a data center is a complicated process that involves rack positioning design, air flow simulations and monitoring devices distribution. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) supplies detailed suggestions for temperature readings and defines three models regarding placing temperature sensors: a) managing the space, b) setting up the space, and c) troubleshooting the space. More specifically, it suggests that temperature should be measured at every fourth rack position at the center of cold aisles. Additionally, temperature measurement should be taken at the center of the air intakes of the top, middle, and bottom equipment mounted on a particular rack [33].

3 Security threats and their impact

Data centres play an increasingly important role in modern society and digital economy. However, their security is still a challenging area with many open problems. This growing reliance on digital infrastructures to accommodate large

volumes of valuable information make them prone to strategic attacks. In this context, many studies showed that most data centres and their associated building automation systems have little or no protection policies in place [8, 12, 13, 19]. In addition, a recent report by CyrusOne⁴ confirmed that its New York data centre was attacked by the REvil ransomware, also known as Sodinokibi, affecting more than six customer companies of CyrusOne [20]. Another report by Verizon found that DDoS attacks and privilege misuse were the most common attack vectors in last year’s cyber attacks against data centres, mainly due to the proliferation of poorly secured connected IoT devices [21].

The temperature monitoring system, which is one of the most critical components of a data centre, is vulnerable to cyber-attacks too [12, 19]. Temperature monitoring is usually guaranteed through hundreds or even thousands of temperature sensors (i.e.thermal sensors) that form a WSN, which can then be remotely controlled and configured. Thus, malicious cyber actors can easily take control of the network and manipulate the thermal conditions of a specific server or rack; for example, maintaining servers in a relatively hot environment [12, 13]. To achieve their goals, attackers will run thermal-intensive workloads into victim servers or VMs (Virtual Machines) to rapidly generate large amounts of heat. This heat will consequently worsen the thermal condition of the peripheral environment, and therefore raise the inlet temperature of other servers [12, 14]. What is more, these networks are vulnerable to a wide variety of attacks due to several restrictions, such as low computational capability, small memory size, limited energy resources and unreliable communication channels [31]. Overall, these insecure devices can be easily compromised in order to remotely launch attacks against the temperature monitoring system.

The adversary who launches a thermal attack can be an individual hacker, a competing provider, or a specialised organisation for committing cyber-crime. A study by [34] showed that thermal attacks can also come from insider malicious tenants, especially in the case of multi-tenant data centres. In this case, malicious tenants can inject additional thermal loads, exceeding the shared cooling system capacity, which can lead to overheating and possible system downtime. It is clear that the impact of thermal attacks is not limited to victim servers, but it can also affect surrounding environmental temperature, which, in turn, can impact thermal conditions of other servers and the entire data centre [14]. In the same context, a study by [12] found that a successful thermal attack targeting only 2% of the servers can dramatically affect the thermal conditions of the entire data centre. Similarly, affecting thermal conditions of adjacent servers and causing local hotspots would significantly raise cooling costs, or lead to a cooling breakdown.

One of the main reasons that renders data centres vulnerable to thermal attacks is the extensive use of aggressive cooling and power management policies, such as power oversubscription [12, 38]. According to this phenomenon, a single data centre hosts a larger number of servers, exceeding its power capacity [12]. Implementations of power oversubscription can lead to severe security gaps; for

⁴ <https://cyrusone.com/>

instance, malicious attackers can manipulate multiple servers to create a local hotspot or generate simultaneous power peaks [38]. Another reason is that thermal attacks try to intentionally keep the outlet temperature of victim servers at a high level preventing chip-level temperature sensors from monitoring the temperature at server- or data centre-level. Consequently, these sensors are unable to prevent local hotspots since they only provide information about the server’s core temperature, which is different from the the inlet and outlet measurements [12].

Current anomaly detection approaches in data centre temperature monitoring exhibit several weaknesses in preventing thermal attacks [12, 25]. For instance, most of these approaches are not capable of distinguishing between benign and malicious thermal-intensive workloads [12]. Furthermore, a study by [13] demonstrated that existing anomaly detection approaches cannot defend against thermal attacks at data centre-level, where a large number of accounts run different workloads simultaneously.

4 Thermal anomaly detection in data centers

Anomaly detection is a relevant approach in the problem of temperature monitoring inside data centres, because it allows data centre managers to identify abnormal system behaviour in a proactive manner. In general, thermal anomaly detection is based on a combination of a) thermal hardware requirements that have been applied to cooling systems, and b) historical observations that are recorded and/or streamed by environmental sensors inside the critical infrastructure.

Several researchers have suggested anomaly detection methods for temperature data in a variety of environments. In an attempt to detect possible risk events inside a household, an anomaly detection method was initially applied to smart home data by [18] through temporal data mining. For their investigation, researchers used sensors to measure temperature, humidity, light, and other environmental factors. The model looked for temporal interactions among frequent activities to learn “resident” behaviour, while events with low probability were reported as “unusual”. Despite using a combination of real and synthetic data, the method proved effective, but was not used to perform predictions [18].

Deviations from historical patterns were also examined on environmental data by [15]. The study presented an anomaly detection method using a data-driven uni-variate auto-regressive model and a moving window to predict the next measurement based on historical data. The method, which employed sensors reporting through telemetry, performed well without requiring previous classification of data, and proved applicable to larger data sets as well [15]. In their extensive study, [6] proposed a spatio-temporal correlation to detect anomalies from temperature, humidity and light sensor nodes in wireless sensor networks. Even though this method aimed to reduce energy and spectrum consumption, the information-gain from aggregate neighbourhood sensor data showed that it

can prove critical in high-risk cases where temperature variations need to be addressed immediately.

In an attempt to ensure healthy operating conditions and reliability of High-Performance Computing systems (HPC) and data centres, researchers in [5] used autoencoders to train data collected from monitoring devices that were mounted on computing nodes. Their method used several metrics, such as core load, power consumption, room temperature, GPU usage, and cooling fans speed, and was able to catch anomalies when tested on real tier-1 HPC systems.

Despite the fact that temperature was not one of the monitoring values, research conducted by [1, 16] demonstrated the use of anomaly detection in critical infrastructure. On one hand, [16] presented an analysis of multi-variate anomaly detection methods to assist in condition-based maintenance. For their application, they used real aircraft turbofan engine data, where their model was able to successfully detect anomalies and recognise failing factors. On the other hand, [1] proposed a multi-agent swarm system to detect abnormal behaviour in cloud data centres. However, neither of the approaches considered collecting data from the surrounding environment for the evaluation, but rather focused on the system’s parameters only.

Following the rapid development of data warehouses around the globe, thermal anomaly detection approaches based on data centres only started to be explored during the last decade. Studies published by Marwah et al. [25, 26, 27] in 2009-10 were the first to address the issue of autonomous anomaly detection of temperature values in data centres using machine learning on sensor data sets. In their first paper on the topic, [27] underlined the importance for an autonomous system that is able to catch temperature values beyond a certain threshold. They also listed common technical reasons, as well as their symptoms, that lead to anomalies. As they stated, most of these anomalies can go undetected by traditional systems, while they demonstrated how Principal Component Analysis (PCA) can assist as the primary detection mechanism. In [25], they extended their initial hypothesis by introducing four detection mechanisms. More specifically, they compared a simple threshold method, a moving average method, an Exponentially Weighed Moving Average (EWMA) and Naive Bayes to predict thermal outliers on a three-month data set. Naive Bayes outperformed the other models and was able to predict 18% of the anomalies at an average of 12 minutes before occurrence. Next, in [26], they used hierarchical PCA for real-time detection, which resulted in a 98% accuracy of predicting anomalous cases. What is more, the method was applied to data coming from a part of the data centre where traditional methods would not be able to raise an alarm.

In [39], researchers introduced a sophisticated two-tier hierarchical neural network framework that detected server-level as well as data centre-level thermal anomalies. The method was able to achieve this by studying the relationships of heterogeneous sensors and consequently outperform other machine learning models. After extensive research on thermal maps for data centres, [22, 23] proposed a novel anomaly detection method that compares and maximizes the accuracy of constructed and observed maps to detect Regions of Interest (RoIs). The

method, which was notably based on thermal cameras and traditional temperature monitoring devices, demonstrated sufficient accuracy in anomalous cases. Later, and responding to the growing need for fast online detection of fault cases, [3, 40] introduced the use of Self-Organizing Maps and reputation systems.

In another study published in 2016, [4] presented a sophisticated four-step tool based on density estimation for anomaly detection and data exploration that was specifically designed for HPC sensors. Even though this method was tested and performed well on a real HPC environment, it was not designed for real-time detection, an approach that recent studies tend to follow. [24] focused exclusively on temperature data, when they used naive ensembles to detect anomalies in the cooling system of data centres. Nevertheless, because of its theoretical approach, their method needs to be extended to demonstrate sufficient performance on non-simulated scenarios. During the same year, another study published by [9] proposed an architecture sensing scheme. According to it, heat sensors collect server temperature data and transmit them to the cloud for further analysis. Later in 2018, [7] tested several machine learning models on server temperature data to examine cooling reliability of data centres. By constructing workload-independent cooling profiles servers, they were also able to detect both transient and lasting cooling failures of servers. Two extensive studies published by [17, 41] in 2019, deployed machine learning to detect fault operations of HVAC systems. In [41], recognising the importance of air cooling systems in data centres, researchers introduced a hybrid approach for four-fault decoupling features; their study included compressor valve leakage, condenser fouling, evaporator airflow reduction, and liquid line restriction. For data manipulation, they used the random forest algorithm, which proved rather effective in fault events diagnosis. Lastly and for the same reason, in [17], researchers considered three different types of anomaly detection methodologies, namely naïve point anomalies, contextual point anomalies, and level shifts. Then, they used several machine learning methods on a real data set, achieving good precision rates.

Taking everything into consideration, there have been several approaches to analyse temperature measurements and define abnormal behaviour of the respective environments. Recent studies have started to include critical infrastructure in their analyses as well, but their approaches are mainly focused on optimal system management, energy efficiency, and reliability of services. Given the recent increase of cyber-attacks in industrial systems, we believe that real-time temperature monitoring in data centres should start to be examined by the security community as an attack mitigation measure. Moreover, temperature monitoring should not be analysed separately, but rather in conjunction with network and security monitoring measures. Additionally, there is a clear need for efficient time response in anomaly detection methods of operational processes applied to data centres; this approach can be observed across the field of financial streaming analytics [2], where having the lead in information exchange provides an advantage in transactions. Finally, we should note that the type and quality of data used in model testing is of paramount importance. In the next section, we address these

issues by proposing a theoretical framework based on a combination of recent methods and tools published by [10, 29, 36, 37].

5 Multi-Variant Anomaly Detection & Fuzzy Logic

In an attempt to address the specific technical problem, we propose a hybrid method that merges the strongest elements of the methodologies found in [10, 29, 37]. A study by [10] highlighted the need for performing anomaly detection in multi-variate time sensing environments and proposed RADM, a real-time anomaly detection algorithm based on Hierarchical Temporal Memory (HTM) and Bayesian Network (BN). Their methodology outperformed anomaly detection in uni-variate sensing time-series, when tested on CPU, network and memory sensors. Additionally, [10] introduced the health factor α , which measures the overall health of the system by taking into account the individual anomaly scores calculated by HTM for each parameter (eg. CPU). We should note that HTM is a far more accurate representation of the neural structures and mechanisms of the brain than the widely spread neural networks.

In their study, [29] demonstrated a sophisticated decentralised scheme for fault detection and classification where wireless sensor networks are prevalent. According to their methodology, neighbour sensors were grouped by conducting calculations among sensor readings. For sensor measurement, their choice of time-series model was the Auto-Regressive Moving Average model (ARMA). Fault detection was performed when reading was compared to a certain threshold and then the fault classification algorithm was initiated to determine the fault type. The fault reading was checked for the frequency and continuity, as well as the presence of an observable pattern. According to these criteria it was then classified as a) *random*, if discontinuous and appearing randomly, b) *malfunction*, if discontinuous and appearing frequently, c) *bias*, if continuous and exhibiting no pattern, d) *drift*, if continuous and following a pattern. For their investigation, they used an outdoor temperature dataset, achieving accuracy rates between 85% and 95%.

The ever-increasing need for intelligent systems in the security domain, also led to non-traditional approaches, such as Fuzzy Logic and Fuzzy Inference Systems (FIS). Researchers in [37] proposed a fuzzy-based approach to overcome complexities in Building Energy Management Systems (BEMSs). More specifically, they were able to use the nearest neighbour and fuzzy rules to extract normal building behaviour. Then, they used fuzzy linguistic descriptors for a variety of parameters, such as zone temperature, exhaust fan load, and supply fan current.

In this paper, our goal is to design an integrated anomaly detection framework that addresses the issues of existing methodologies discussed in Sections 3 and 4. While our proposal method focuses on temperature, it considers various measured phenomena as inputs too, reflecting the highly monitored environment of data centres today. Even though our proposal is purely theoretical, it is based

on well-examined methods, which were tested individually and proved effective in their respective field.

The proposed methodology combines elements from the three works described previously and splits them into four main steps. First, we perform a grouping of sensors as per the type of measured phenomenon or the spatial distribution of sensors measuring the same reading. Second, we perform fault detection on the grouped time series models individually. In the case where a faulty reading has been detected, we perform fault classification to define its type according to continuity, frequency and pattern criteria. Third, the grouped time series are inserted into the RADM framework, where the global anomaly region is defined. If once again, the system detects an anomaly on the overall system, we perform fault detection to define the type. Finally, we present a permutation of the health factor α , based on Fuzzy Logic and Fuzzy Rules. More details about each step are provided in the following subsections.

5.1 Step 1: Thermal Sensors Grouping

In this step, we acknowledge that the ambient temperature of data centres is a product of more than one factor. For instance, new or advanced cooling systems, such as rack cooling, have been employed in recent years to aid in reserving the servers' ideal state. Consequently, there is a clear need to consider more than one input when building time series models for anomaly detection. For this reason, we propose a sensor grouping technique that groups same phenomenon sensor readings, e.g. grouping of all ambient air temperature sensors in a room, fan motor speed, etc. Similarly, sensors can be also grouped spatially, for example aggregating high-mounted ambient temperature sensors of a specific aisle in a single group, low-mounted ambient temperature sensors in a different one, creating a dedicated group for high-mounted rack temperature sensors, and so on. For this method, we employ the Neighborhood Voting scheme proposed by [29]. According to them, *the system does not require a priori knowledge about the environment. Instead, it takes advantage of the redundancy in measurements of sensor readings.* By changing the notation to keep the consistency of our framework, we present the Neighbourhood Voting algorithm as it appears in their study:

1. Collect the set of readings $R = r[1 \dots |Neighbour(Y_i)|]$ from all neighbours, excluding its own reading r_i .
2. Calculate the median of the group, $\mu = \{R\}_{\frac{1}{2}} = \bar{r}$.
3. Calculate the difference between r_i and \bar{r} , $D_{r_i\bar{r}} = |r_i - \bar{r}|$.
4. Compare the difference $D_{r_i\bar{r}}$ with a threshold τ , that can be adjusted.
 - If $D_{r_i\bar{r}} < \tau$ then r_i is a good reading.
 - If $D_{r_i\bar{r}} \geq \tau$ then r_i is a faulty reading. We define faulty readings as $D_{r_i\bar{r}} = \varepsilon_i$.

where Y_i is a node measuring a specific phenomenon, r_i is its reading. The set of its neighbours is denoted by $Neighbour(Y_i)$ and the number of neighbours by $|Neighbour(Y_i)|$.

5.2 Step 2: Fault Detection

In this step, we propose fault detection on the aforementioned grouped measurements. Because of the many advantages it exhibits for anomaly detection over traditional time-series, grouped readings are modelled with HTM time-series as suggested by [10], and not the ARMA model preferred in [29]. If a faulty reading is detected, the Fault Classification Algorithm proposed in [29] is initiated. The algorithm is run parallel to the HMT modelling to classify the faulty measurement as either *malfunction*, *random*, *bias*, or *drift*, according to the Fault Classification Algorithm (Algorithm 1), where T represents time intervals, $|\varepsilon_i|$ the numbers of faulty occurrences, and θ the desired threshold.

Algorithm 1: Fault Classification [29].

Input:

1. $R[1..T]$: vector of T sensor readings
2. $E[1..T]$ vector of the faulty state of $R[1..T]$

Output: C : fault type of sensor node in the interval

- 1 compute the occurrences of faults $|\varepsilon_i|$ in R
- 2 check the continuity
- 3 **if** ε_i *is discrete* **then**
- 4 Check the frequency
- 5 **if** $|\varepsilon_i| \geq \theta$ *i* **then**
- 6 $C = Malfunction$
- 7 **else**
- 8 $C = Random$
- 9 **if** ε_i *is continuous* **then**
- 10 Check the fault function ε_i
- 11 **if** $\varepsilon_i = const$ **then**
- 12 $C = Bias$
- 13 **else**
- 14 $C = Drift$
- 15 **return** C

5.3 Step 3: Anomaly Region Detection

During the third step of the process, the HTM time-series of the grouped variables are combined with the Bayesian Network to detect the anomaly region. The time-series follow the RADM framework proposed by [10]. As previously discussed, we define the temperature T of the data centre as a product of factors X , Y , and Z , according to Equation 1:

$$T(t) = (X(t), Y(t), Z(t)) \quad (1)$$

Then, the overall process, including steps 1, 2, and 3, is described in Figure 2.

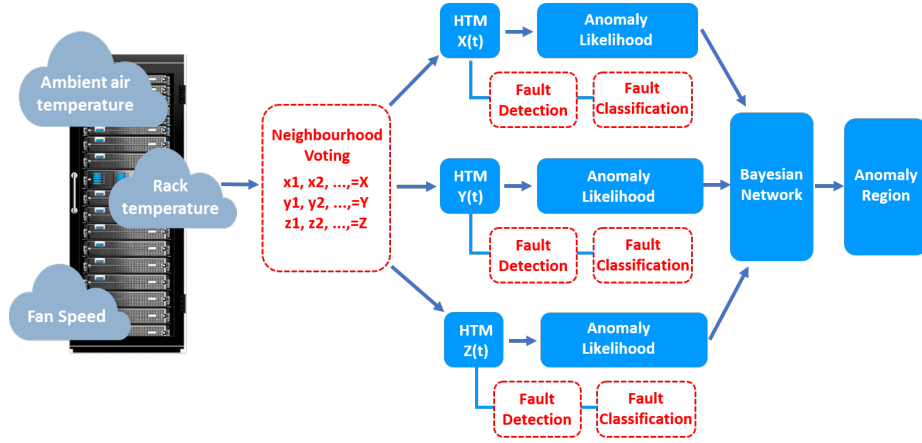


Fig. 2. Steps 1, 2, and 3 of the combined methods. Boxes in red represent methods proposed by [29], while boxes in blue comprise the RADM framework proposed by [10].

5.4 Step 4: Fuzzy sets

In the fourth step of our methodology, we propose a permutation of the health factor α proposed by [29], in order to assess the overall well being of our framework. Our method uses linguistic descriptors to describe the concept of the system's anomaly state in an aggregated and intelligent manner. The method uses fuzzy rules similar to the analysis by [37], except the input variables consist of the anomaly scores S_i calculated by the HTM model for the grouped variables. More specifically, the input variables are represented using five fuzzy sets, namely *very low*, *low*, *medium*, *high*, and *very high*. The output variable *Health* is represented again by five fuzzy sets, that describe the state of the system as: *verybad*, *bad*, *average*, *good*, and *very good*. Input and output linguistic descriptors are shown in Figure 3. Fuzzy rules are then constructed according to experience and describe the contribution of individual anomalies to the general state. The linguistic descriptors for the healthiest state with the three inputs used in our example would be:

**IF $S_{X(t)}$ IS Very Low AND $S_{Y(t)}$ IS Very Low AND $S_{Z(t)}$ IS Very Low
THEN Health IS Very Good**

By using this fuzzy-based system, the proposed methodology is able to use its own metrics -the anomaly scores- to assess the overall state of the system. FISs are able to support a large number of inputs and rules to achieve the desired accuracy. This approach offers enormous scalability for a critical environment with multiple sensors, as it is able to include many factors affecting server temperature. Finally, for a data set that could fit into a future application of this framework, we propose the DAD data set published by [36]. The DAD data set is a labelled IoT data set, which contains real-world behaviours of a data centre as

seen from the network perspective. The network and environmental values were obtained from a physical data centre, combined with temperature measurements transmitted by NFC smart passive sensor technology.




Fig. 3. Linguistic descriptors for (α) individual Anomaly Scores and (β) overall Health.

6 Conclusions and Future Work

In this paper, we examined how easily temperature monitoring systems can be used by attackers to manipulate heat distribution in such environments. We also presented and reviewed existing model-based anomaly detection methods that have been focused on measuring temperature deviations inside critical infrastructures. According to our findings, thermal attacks are fundamentally not

addressed as potential cybersecurity attacks for data centres, and consequently, there is a serious lack of frameworks proposed to eliminate this threat. We believe that future studies on the current topic should be focused on constructing intelligent solutions. In our research, we found several approaches to tackle this issue and addressed the most important aspects for a thermal attack mitigation methodology. As part of our initial work, we have designed a multi-variate anomaly detection method for fault classification, and a fuzzy-based health factor to assess the overall state of the system. Finally, we propose a curated data set that could assist in the exploration of the aforementioned models.

In future work, we intend to extend our research by conducting experiments on large data sets from real settings to further test the effectiveness of this approach. We also intend to compare it against previous studies to evaluate its performance and further explore aspects of development.

Acknowledgement



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 786698. This work reflects authors' view and Agency is not responsible for any use that may be made of the information it contains.

Bibliography

- [1] Adhikari, A.T.B., Karunananda, A.S.: Real time anomaly detection of cloud data centers with multi agents swarm. SRI LANKA ASSOCIATION FOR ARTIFICIAL INTELLIGENCE p. 5 (2015)
- [2] Ahmad, S., Purdy, S.: Real-time anomaly detection for streaming analytics. arXiv preprint arXiv:1607.02480 (2016)
- [3] Aransay, I., Sancho, M., Garcia, P., Fernández, J.: Self-organizing maps for detecting abnormal thermal behavior in data centers. In: Proceedings of the 8th IEEE International Conference on Cloud Computing (CLOUD). pp. 138–145 (2015)
- [4] Baseman, E., Blanchard, S., DeBardeleben, N., Bonnie, A., Morrow, A.: Interpretable anomaly detection for monitoring of high performance computing systems. In: Outlier Definition, Detection, and Description on Demand Workshop at ACM SIGKDD. San Francisco (Aug 2016) (2016)
- [5] Borghesi, A., Libri, A., Benini, L., Bartolini, A.: Online anomaly detection in hpc systems. In: 2019 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS). pp. 229–233. IEEE (2019)
- [6] Bosman, H.H., Iacca, G., Tejada, A., Wörtche, H.J., Liotta, A.: Spatial anomaly detection in sensor networks using neighborhood information. *Information Fusion* **33**, 41–56 (2017)
- [7] Chen, C., Wang, G., Sun, J., Xu, W.: Detecting data center cooling problems using a data-driven approach. In: Proceedings of the 9th Asia-Pacific Workshop on Systems. pp. 1–8 (2018)
- [8] Chou, J.S., Telaga, A.S.: Real-time detection of anomalous power consumption. *Renewable and Sustainable Energy Reviews* **33**, 400–411 (2014)
- [9] Das, N., Kundu, A.: Cloud based thermal management system design and its analysis. In: Information Systems Design and Intelligent Applications, pp. 519–528. Springer (2016)
- [10] Ding, N., Gao, H., Bu, H., Ma, H., Si, H.: Multivariate-time-series-driven real-time anomaly detection based on bayesian network. *Sensors* **18**(10), 3367 (2018)
- [11] ENERGY.GOV: Management technology for energy efficiency in data centers and telecommunications facilities. <https://www.energy.gov/eere/amo/management-technology-energy-efficiency-data-centers-and-telecommunications-facilities>, accessed: 05-08-2020
- [12] Gao, X., Liu, G., Xu, Z., Wang, H., Li, L., Wang, X.: Investigating security vulnerabilities in a hot data center with reduced cooling redundancy. *IEEE Transactions on Dependable and Secure Computing* (2020)
- [13] Gao, X., Xu, Z., Wang, H., Li, L., Wang, X.: Why” some” like it hot too: Thermal attack on data centers. In: Proceedings of the 2017 ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems. pp. 23–24 (2017)

- [14] Gao, X., Xu, Z., Wang, H., Li, L., Wang, X.: Reduced cooling redundancy: A new security vulnerability in a hot data center. In: NDSS (2018)
- [15] Hill, D.J., Minsker, B.S.: Anomaly detection in streaming environmental sensor data: A data-driven modeling approach. *Environmental Modelling & Software* **25**(9), 1014–1022 (2010)
- [16] Hu, X., Subbu, R., Bonissone, P., Qiu, H., Iyer, N.: Multivariate anomaly detection in real-world industrial systems. In: 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence). pp. 2766–2771. IEEE (2008)
- [17] Huang, J., Chai, Z., Zhu, H.: Detecting anomalies in data center physical infrastructures using statistical approaches. In: *Journal of Physics: Conference Series*. vol. 1176, p. 022056. IOP Publishing (2019)
- [18] Jakkula, V., Cook, D.J., et al.: Anomaly detection using temporal data mining in a smart home environment. *Methods of information in medicine* **47**(1), 70–75 (2008)
- [19] Knapp, K.J., Denney, G.D., Barner, M.E.: Key issues in data center security: An investigation of government audit reports. *Government Information Quarterly* **28**(4), 533–541 (2011)
- [20] Korolov, M.: Cyrusone confirms ransomware attack, says six customers affected. <https://www.datacenterknowledge.com/cyrusone/cyrusone-confirms-ransomware-attack-says-six-customers-affected>, accessed: 04-08-2020
- [21] Korolov, M.: The four main types of cyberattack that affect data center uptime. <https://www.datacenterknowledge.com/security/four-main-types-cyberattack-affect-data-center-uptime>, accessed: 04-08-2020
- [22] Lee, E.K., Viswanathan, H., Pompili, D.: Model-based thermal anomaly detection in cloud datacenters. In: 2013 IEEE International Conference on Distributed Computing in Sensor Systems. pp. 191–198. IEEE (2013)
- [23] Lee, E.K., Viswanathan, H., Pompili, D.: Model-based thermal anomaly detection in cloud datacenters using thermal imaging. *IEEE Transactions on Cloud Computing* **6**(2), 330–343 (2015)
- [24] Li, C.: Cooling anomaly detection for servers and datacenters with naive ensemble. In: 2016 32nd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM). pp. 157–162. IEEE (2016)
- [25] Marwah, M., Sharma, R., Bash, C.: Thermal anomaly prediction in data centers. In: 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. pp. 1–7. IEEE (2010)
- [26] Marwah, M., Sharma, R., Lugo, W., Bautista, L.: Anomalous thermal behavior detection in data centers using hierarchical pca. *Knowledge Discovery from Sensor Data (SensorKDD’10)* p. 78 (2010)
- [27] Marwah, M., Sharma, R.K., Lugo, W.: Autonomous detection of thermal anomalies in data centers. In: *International Electronic Packaging Technical Conference and Exhibition*. vol. 43604, pp. 777–783 (2009)
- [28] Mehrfeld, J.: Cyber security threats and incidents in industrial control systems. In: *International Conference on Human-Computer Interaction*. pp. 599–608. Springer (2020)

- [29] Nguyen, T.A., Bucur, D., Aiello, M., Tei, K.: Applying time series analysis and neighbourhood voting in a decentralised approach for fault detection and classification in wsns. In: Proceedings of the Fourth Symposium on Information and Communication Technology. pp. 234–241 (2013)
- [30] RACK-SOLUTION: 400 million new servers might be needed by 2020. shorturl.at/cqzCY, accessed: 04-08-2020
- [31] Salehi, S.A., Razzaque, M.A., Naraei, P., Farrokhtala, A.: Security in wireless sensor networks: Issues and challenges. In: 2013 IEEE International Conference on Space Science and Communication (IconSpace). pp. 356–360. IEEE (2013)
- [32] Schmidt, R.R., Cruz, E.E., Iyengar, M.: Challenges of data center thermal management. *IBM Journal of Research and Development* **49**(4.5), 709–723 (2005)
- [33] Seaton, I.: Temperature Sensor Location - Data Center Efficiency | Upsite (Oct 2018), <https://www.upsite.com/blog/data-center-temperature-sensor-location/>, library Catalog: www.upsite.com Section: Blog
- [34] Shao, Z., Islam, M.A., Ren, S.: A first look at thermal attacks in multi-tenant data centers. *ACM SIGMETRICS Performance Evaluation Review* **46**(2), 93–94 (2019)
- [35] Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., Lintner, W.: United states data center energy usage report. Tech. rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States) (2016)
- [36] Vigoya, L., Fernandez, D., Carneiro, V., Cacheda, F.: Annotated dataset for anomaly detection in a data center with iot sensors. *Sensors* **20**(13), 3745 (2020)
- [37] Wijayasekara, D., Linda, O., Manic, M., Rieger, C.: Mining building energy management system data using fuzzy anomaly detection and linguistic descriptions. *IEEE Transactions on Industrial Informatics* **10**(3), 1829–1840 (2014)
- [38] Xu, Z., Wang, H., Xu, Z., Wang, X.: Power attack: An increasing threat to data centers. In: NDSS (2014)
- [39] Yuan, Y., Lee, E.K., Pompili, D., Liao, J.: Thermal anomaly detection in datacenters. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* **226**(8), 2104–2117 (2012)
- [40] Zapater, M., Fraga, D., Malagón, P., Banković, Z., Moya, J.M.: Self-organizing maps versus growing neural gas in detecting anomalies in data centres. *Logic Journal of the IGPL* **23**(3), 495–505 (2015)
- [41] Zhu, X., Du, Z., Jin, X., Chen, Z.: Fault diagnosis based operation risk evaluation for air conditioning systems in data centers. *Building and Environment* **163**, 106319 (2019)