5-1-2017

Thermal Fingerprinting—Multi-Dimensional Analysis of Computational Loads

Matthias Vodel
Chemnitz University of Technology, vodel@hrz.tu-chemnitz.de

Marc Ritter
Mittweida University of Applied Sciences, ritter@hs-mittweida.de

THERMAL FINGERPRINTING—MULTI-DIMENSIONAL ANALYSIS OF COMPUTATIONAL LOADS

Matthias Vodel
Chemnitz University of Technology
vodel@hrz.tu-chemnitz.de

Marc Ritter
Mittweida University of Applied Sciences
ritter@hs-mittweida.de

Abstract:

Digital fingerprinting is used in several domains to identify and track variable activities and processes. In this paper, we propose a novel approach to categorize and recognize computational tasks based on thermal system information. The concept focuses on all kinds of data center environments to control required cooling capacity dynamically. The concept monitors basic thermal sensor data from each server and chassis entity. The respective, characteristic curves are merged with additional general system information, such as CPU load behavior, memory usage, and I/O characteristics. This results in two-dimensional thermal fingerprints, which are unique and achievable. The fingerprints are used as input for an adaptive, pre-active air-conditioning control system. This allows a precise estimation of the data center health status. First test cases and reference scenarios clarify a huge potential for energy savings without any negative aspects regarding health status or durability. In consequence, we provide a cost-efficient, light-weight, and flexible solution to optimize the energy-efficiency for a huge number of existing, conventional data center environments.

Keywords:
fingerprints; monitoring; control system; air-conditioning; data center; optimization; energy-efficiency; adaptive; sensor fusion

1. Introduction

Conventional data centers are wasting a massive amount of energy for cooling installed, heterogeneous hardware components. The cooling capacity is calculated as a static value, which represents a worst case cooling scenario with maximum thermal load. In order to optimize the air-conditioning baseline, several best practice approaches are used. Some of them are focusing on individual adaptations based on practical administrator experiences, i.e., static / periodical schedules of different cooling capacities or physical rearrangements of the given hardware components (Patel et al. (2002), Greenberg et al. (2006), Lent (2016)). Also DCIM solutions for an optimized data center management are well-known toolkits (Gilbert et al. (2013)).

Many other approaches deal with technical solutions, for example cold aisle containments or additional air boosters inside the double-floor. Another well-known approach deals with the migration from conventional CRAC concepts (computer room air conditioning) to more efficient, local cooling systems for each rack. And of course very familiar solutions represent the modernization of core units inside the air-conditioning system (optimized supercharger, DX evaporative coolers, radiators, EC brushless fans etc.) (Vogel et al. (2007), Pakbaznia & Pedram (2009)). Further, more innovative approaches are introducing additional sensor units inside the data centers (Liu & Terzis (2012), Wang et al. (2014)). Due to an easy installation and fast maintenance routine, wireless solutions are preferred (Jia et al. (2011), Zanatta et al. (2014)).

The key issues for all of these solutions are very simple. Either the approach is not flexible and provides only a reduced safety buffer, or, on the other hand, the optimization of the energy-efficiency requires massive hardware modifications which costs a huge amount of money and engineering efforts (Buyya et al. (2010)).

In contrast, our approach focuses on the usage of given sensor data sources which are merged into an adaptive software framework, which is able to control the cooling capacity of the available air-conditioning systems (Vodel et al. (2015)). The core feature represents thermal fingerprints—a promising and innovative instrument to classify computational tasks regarding its impact on the thermal behavior inside the data center.

Accordingly, the following paper is structured as follows. First of all, section 2 summarizes related approaches in the domain of data center monitoring, adaptive control loops, and sensor extensions for data center environments. After that, section 3 presents the key concept of thermal fingerprints, their specific parameters, and the use cases for optimizing the cooling capacity. The following section 4 deals with the reference scenarios, detailed environmental conditions, and the setup for the adaptive learning algorithm. The respective results are presented in section 5, including a critical discussion on system behavior and worst case scenarios. The final section summarizes the paper and clarifies the practical benefits of the approach.

2. Related Work

During the last decade, digital fingerprints were started to be used in multiple application domains. One important aspect deals with conventional security scenarios, e.g. digital forensics (Swaminathan et al. (2008)) or authentication/authorization concepts based on biometric watermarking (Noore et al. (2007)). The authors are using specific parameters from inside hardware & operating system, the software applications, and the characteristic content representation of the peripheral components.

Further application scenarios for digital fingerprinting deal with tracking and profiling of user or process behavior. This area includes OS fingerprinting to identify host operating systems without direct access to the hardware (Fifield et al. (2015)) as well as browser and canvas fingerprinting for the analysis of user behavior during web-sessions without using dedicated cookie techniques (Takei et al. (2015)).

Another well-known domain for digital fingerprinting are services for music, video or image identification, e.g. the Shazam app or the Amazon Firefly service (Han et al. (2015)). Here, complex multimedia content is reduced to simple, categorisable hashes, allowing the user to recognise contents within almost any playback environment. Therefore, only some media snippets are required, which may only contain poor quality samples. With respect to the application domain for energy-efficient cooling concepts for data centres, other and yet complex approaches have been explored. For example, by Li et al., who proposed a forecasting model for the health status inside data centers (Liu et al. (2011)). The model allows multiple data sources as well as self-learning capabilities. The key problem of this approach correlates with the training efforts as well as missing features for classification and event recognition. Furthermore, Liang et al. (2009) developed RACNet—a high-fidelity data center sensing network based on several sensor units. This system also tries to enable a more precise adjustment of the given cooling capacities. Here, the aggregated sensor data is processed with static rule sets in order to calculate a single health status level.

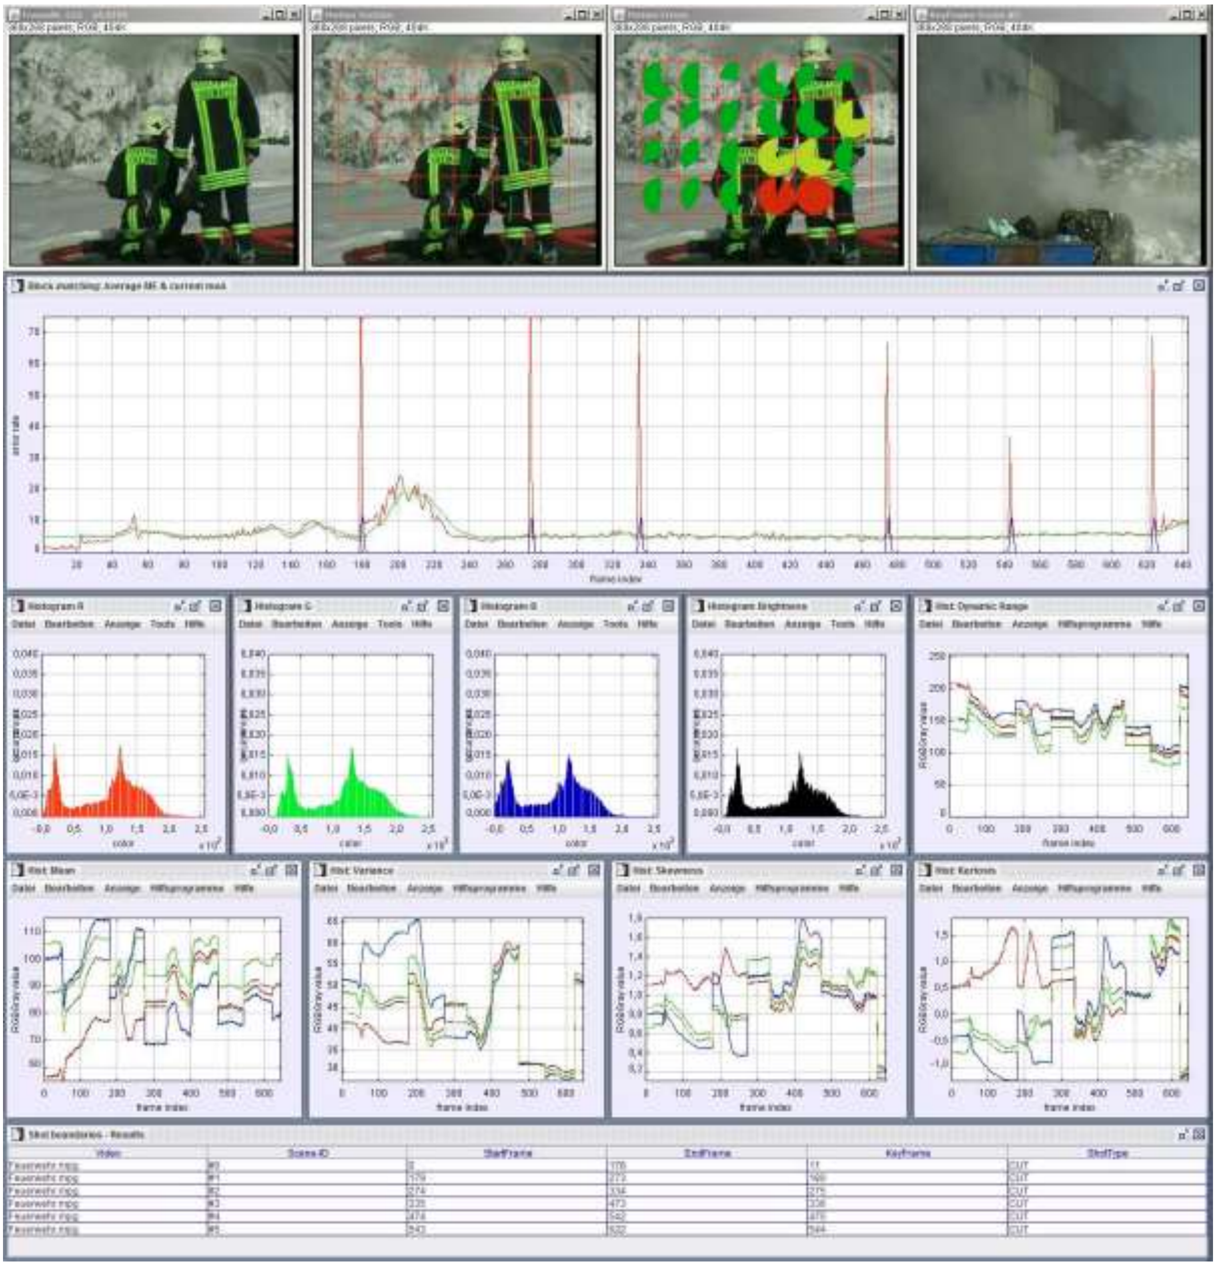

Figure 1. Visualization of two different methods for shot detection available in the AMOPA framework for automated video analysis. Top row: Original frame (left), motion estimation (middle) and representative key-frame of the last detected sequence (right). 2nd row: Course of motion-compensated error functions; spikes indicate a detected shot boundary. 3rd and 4th row: Conventional histograms of the current frame and their statistical measures in the course of time that are harder to trace and to predict than the method shown in the top rows. Bottom: Table of the detected sequences. (From: (Ritter, p.268))

We are now focusing on the adaptation of AMOPA (Automated MOving Picture Annotator) (Ritter (2014), Ritter & Eibl (2011)), a powerful educational and research software framework, that was originally build for the analysis of multimedia content. AMOPA was developed at Chemnitz University of Technology in the Professorships of Media Informatics and Media Computing. The framework automatically segments the structure of videos by applying stateof-the-art shot detection approaches (cf. to Fig. 1) Another layer introduces concepts and algorithms to identify and track predefined sequences or objects in unknown video content. Here, the operator is able to annotate samples, train classifiers and rapidly evaluate the classification outcomes. The application for the processing and detection of object instances in the big data context has been proven within the successful participation in the international TRECVid evaluation campaign (Ritter et al. (2015), Ritter et al. (2014)).

These basic concepts and the data analysis workflows of AMOPA are used for multimedia content identification. For the proposed approach, we adapt the workflows to the demands of computational software tasks. Accordingly, the system must be able to handle the respective thermal sensor data originating from different kinds of sensor classes.

3. Thermal Fingerprints

The proposed thermal fingerprint concept represents a feasible toolkit for software-based health-status estimation in heterogeneous data center environments. We try to solve the already mentioned challenges. At the same time, we avoid critical issues of related research projects, i.e. no hardware efforts, scalable usage for different environmental conditions and a short learning stage/initial setup. Based on the individual fingerprint patterns and the respective classification features, the expected re-use factor of the generated knowledge appears to be very high.

3.1 AMOPA Framework

A very important part for this approach is a sufficient software integration and the respective UI. Our goal is to integrate the entire data processing chain into the introduced AMOPA framework. In this context, we have to modify several input modules in order to support the following types of sensor sources:

- CPU temperature
- Memory temperature
- Chassis temperature
- GPU temperature
- Power supply temperature
- CPU load
- I/O load
- Network load
- S.M.A.R.T. information (work in progress)

Figure 2. Scheme for the generic creation of a simplified linear processing chain in AMOPA (Ritter (2014)).

The basic operational workflow (single process instance) is shown in Figure 2. Each process is derived from a single thread class that works on the input data and stores the resulting data for further processing steps. The chain can be flexibly concatenated as well as branched by using XML patterns. The data is automatically transmitted to the next process by the AMOPA framework. The adapted input process aggregates the already mentioned sensor types as well as predefined, environmental parameters. One key result of the workflow represents the calculation of two-dimensional thermal fingerprints.

In a further step, these fingerprints can be used to derive and assess multiple risk levels shown as different types of blue colored bars on the left hand side of Figure 3. Here, an adaptive Gaussian distribution curve allows a domain-specific, individual balancing of the threshold values. These levels are not static and represent an adaptive perspective to the health monitoring of the data center. Based on these levels, the system grants the operators or administrators to manage generic notification events and alert traps.

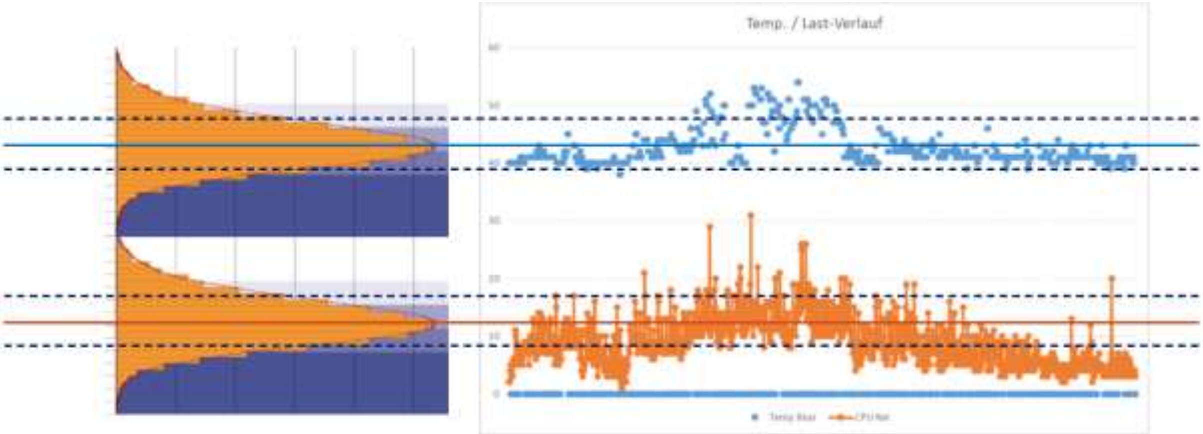

Figure 3. Data input for the calculation of two-dimensional, thermal fingerprints.

3.2 Thermal Pattern Calculation

In order to calculate two-dimensional, individual thermal patterns, a dedicated data analysis sequence has to be processed. For this purpose, we are working with dynamic sliding window approaches in the time domain, as often used for distributed network simulators (Vodel et al. (2008), Vodel (2014)). Accordingly, the given sensor data will be merged and analyzed in several time spans in parallel by using different window sizes. Starting from the last 60 seconds and up to 7200 seconds of logged sensor data, the signal curves are processed. This allows both, the analysis of short term computational load behaviour as well as long term thermal sequences. The workflow is structured as follows being supplemented by Figure 4.

At first, we define a specific window size. Let the size be for example 120 units, what can be either seconds or minutes. Then we process each data channel of input sensors separately and on each specific window with overlap (e.g. at step size 10) in the top row (left). Here, the input sensors show values of *CPU load* (green), *Temperature* (red), and *CPU relative load*. The current classification of the original computation pattern is illustrated on the right being a set of *Computing* (red), *Web services* (yellow), and *Backup & Maintenance* (green). We used the traffic scheme in order to indicate the computational load of the machine. For each of the obtained windowed data distributions, we apply the following procedure:

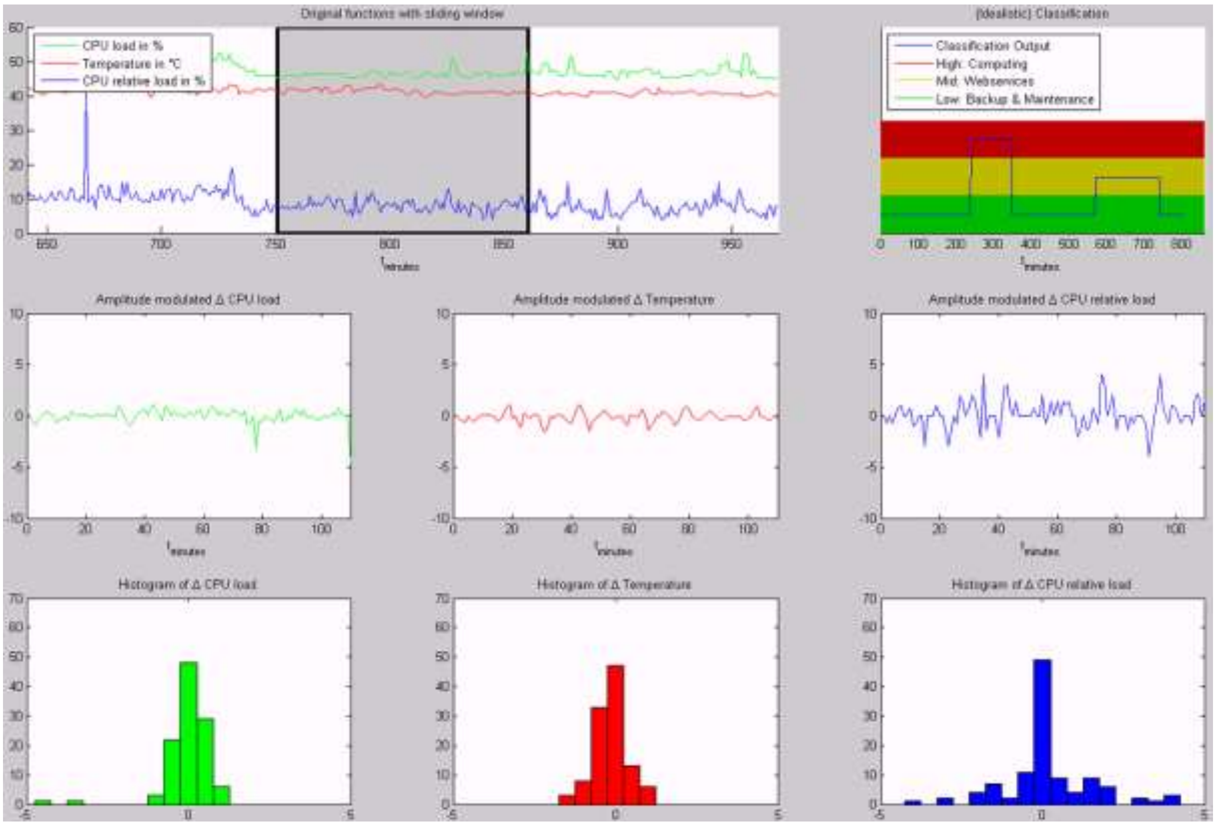

Figure 4. Example set for thermal fingerprint of backup & maintenance tasks in data center environments.

- Subtract the mean from the current data distribution window.
- Compose (multiply) the data with a sinusoidal wave that works as a carrier function ranging from 0 to 86 in quarter steps of π (middle row).
- Calculate the histogram from the composed data (bottom row).

As an intermediary result, we receive specific histograms for different time slots, which represent deviations from the given baseline. The threshold value between the different deviation zones represent the already mentioned risk levels.

In order to stabilize those results, we apply the window data distributions from step 2) to another method. The analysis of spectrograms is very common in the audio domain. The composed/transformed sinusoidal waveform allows us to use such methods in this domain of data. In order to yield reasonable spectrograms, we first have to transform the signal with a Butterworth filter that was designed by using the MATLAB DSP Toolbox for high-pass filters with the coefficients of filter order 10, a cutoff frequency for the point 3 dB point below the passband value of 25, and a sampling frequency of 200 Hz. For the already mentioned example above (120 units), we generate a spectrogram over the whole window with MATLAB standard parameters[1]. By default, those parameters divide the window data distribution into 8 segments with 50 % overlap using a Hamming window. The number of frequencies points is limited to 128.

[1] cmp. to http://de.mathworks.com/help/signal/ref/spectrogram.html, 2016-06-09

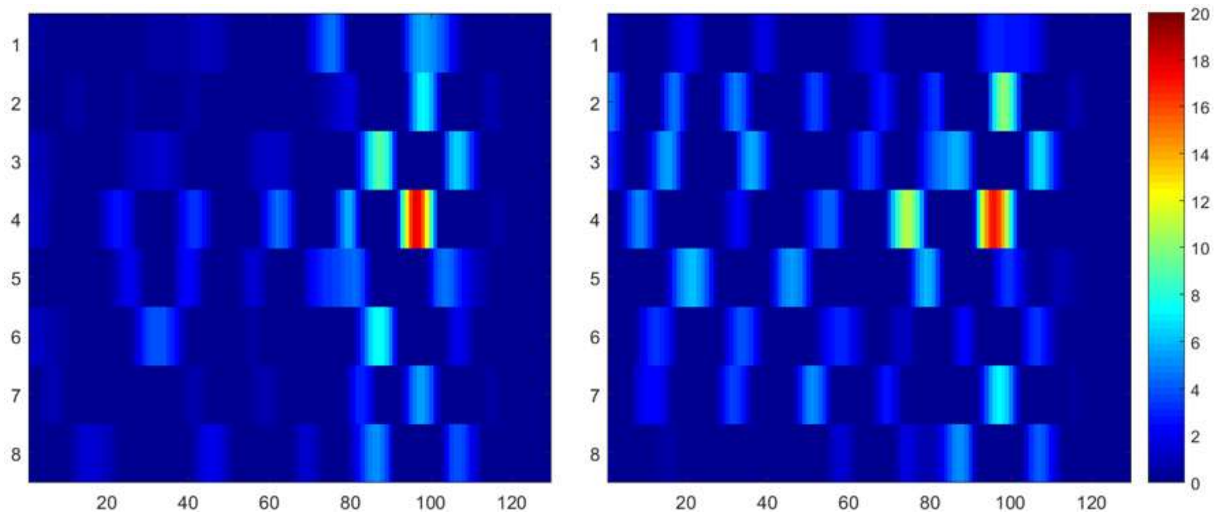

Figure 5. Example set for thermal fingerprint of the computational task in our data center environment for the CPU load (left) and the CPU relative load (right).

The results are two-dimensional, colored patterns, which represent unique fingerprints. Examples of the acquired spectrogram fingerprints are shown in Figure 5, providing several kinds of information. On the X axis, the time domain in the sample window sequence of 120 units is given whereas the Y axis contains the corresponding segments. The individual colored cells describe the signals intensities in pseudocolors. The figure clarifies higher frequencies (red) in the patterns for CPU load (left) and CPU relative load (right) at around 100 units within this window.

3.3 Fingerprints Storage & Discussion

The pattern can be stored in a database as a simple concatenation of the hexadecimal color values of each sector. The pattern resolution is predefined with a static number of X columns and Y rows. The fingerprint pattern can be used as a general knowledge base to identify and recognize repeating tasks.

Actual and future research work investigate topics for an efficient classification or categorization of these patterns. Here, the key challenge is represented by the different variations of similar tasks. These challenges are comparable to the already mentioned music and multimedia databases like Shazam or Amazon Firefly. The system must be able to detect similar computational tasks with modified environmental conditions or different initial system states. E.g., there is a huge thermal impact if several parallel tasks are executed at the same time. Also, the thermal pattern varies if the actual cooling capacity is different to the reference measurement. Accordingly, the patterns provide significant, domain-specific characteristics with scenario-specific, adapted shapes. Our system must be able to detect these constant characteristics in a generic way and with minimal training efforts.

The self-learning capabilities as well as the pattern classification are not in the focus of this paper and represent early work in progress.

4. Reference Scenarios & Test Bench

In order to test the proposed approach under real-world conditions, we've done some proof-of-concept. The test environment represents the Chemnitz University of Technology Computing Centre with its central server locations. The data centre consists of more than 200 server, storage, and network components. The hardware is cooled by a conventional air-conditioning system. The server racks are organised as cold aisle containment groups to shrink the air volume for cooling. Additional booster components inside the double floor allow the dynamic adaptation of the airflow and the respective local cooling capacity (see Figure 6).

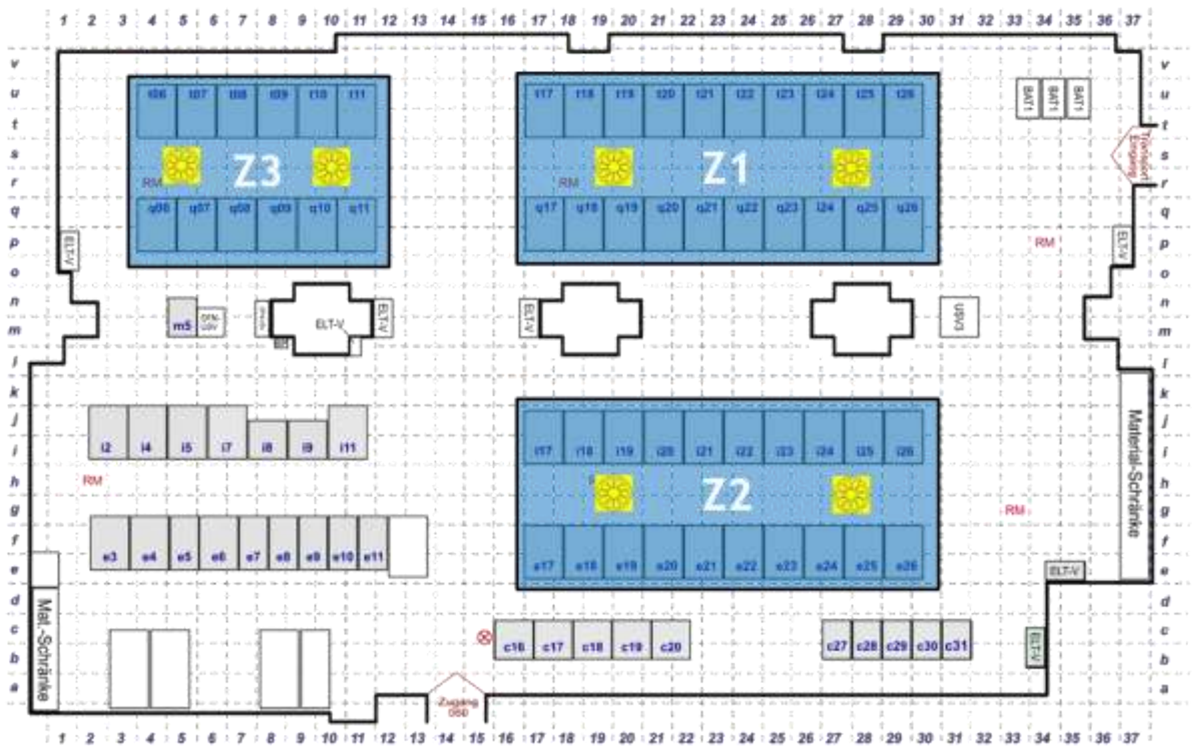

Figure 6. Schematic visualization of the TU Chemnitz data center. Cold aisle zones Z1 to Z3 (blue), booster fans (yellow).

4.1 Setup & Environmental Conditions

In order to provide sufficient sensor data, we measured all available temperature sensor inside the individual hardware nodes (CPU, chassis, memory, power supply, and GPU if available). The sampling rate is static and predefined with one measurement per second. Furthermore, the logging system grabs the CPU load, I/O load and network load with the same sampling frequency. The data was stored on a dedicated logging server for the entire post-analysis routines.

4.2 Limitations & Constraints

One limitation deals with the given sampling rate of one measurement per second. Here, a higher data resolution allows more precise data analysis.

Another limitation represents the missing reference values for the temperature measurements. We are using standard system interfaces for extracting the sensor data from the hardware components. These values are not referenced and may differ between similar systems. Only a few global *reference* temperature sensors inside the data center location are available, which are used as conventional control input for the air-conditioning system.

Further constraints address the mapping between sensor data and software tasks. In order to define the relation between measurement sets and the respective computational tasks, we are using the log files of the systems. Minimal time drifts between different systems are possible but uncritical. The system log correlates with the time stamps of the measurement sets and enables a direct mapping of a system load to the thermal load as well as the respective estimation for the thermal impact.

For this contribution, we analyzed different hardware from our IT infrastructure. This includes a Cisco UCS chassis for our desktop virtualization, a central storage system (NetApp FAS series) as well as Dell PowerEdge servers as dedicated compute nodes as service-providing systems. The proof-of-concept consists of 48 hours' log data for each system.

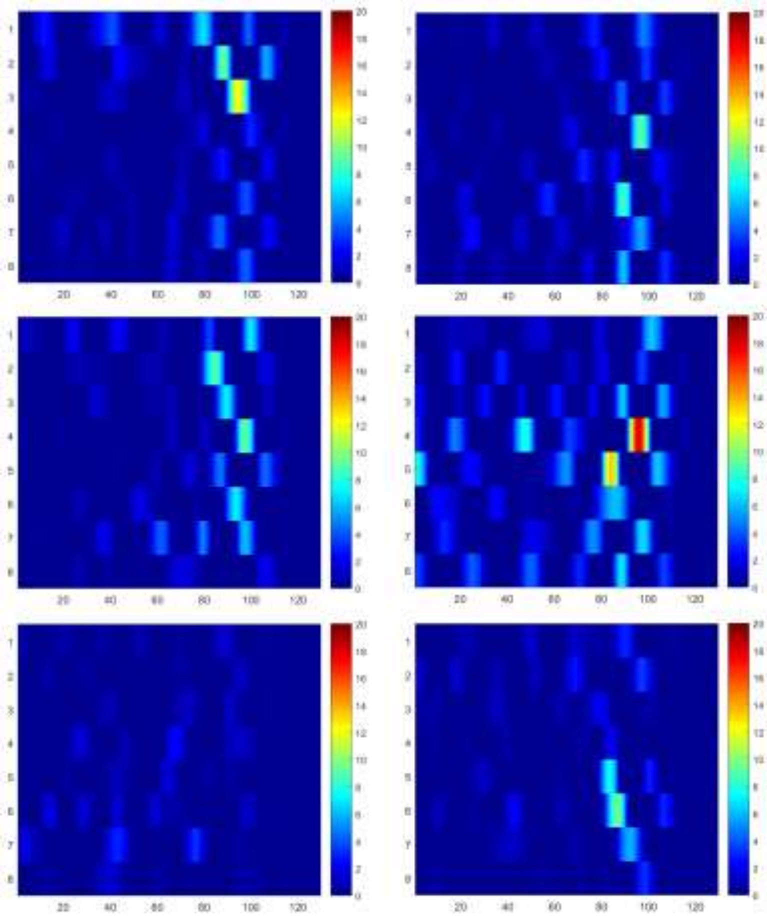

Figure 7. Six different thermal fingerprints for different types of computational tasks.

5. Result Analysis & Scale-Out

Based on the given data center infrastructure, several systems are monitored and processed. The proposed approach generated individual fingerprints for any type of system (both physical or virtual). Figure 7 illustrates a small number of reference patterns.

The figure includes the following tasks:

1. Top left: Cisco UCS virtualization cluster → cluster-internal VM migration task for cluster maintenance preparation.
2. Top right: Cisco UCS virtualization cluster → boot up sequence for a virtual PC pool of 20 hosts.
3. Center left: DELL PowerEdge 710 server → short term batch $<$5 minutes (user-specific computational task)
4. Center right: DELL PowerEdge 710 server → parallelized long term batch (about 60 minutes' user-specific task with I/O and computational load)
5. Bottom left: NetApp FAS 3240C → copy process for 25 GB of user data
6. Bottom right: NetApp FAS 3240C → data integrity test for 25 GB of user data

The measured thermal fingerprints are stored in a database knowledge base. This data is used to trigger an adaptive control loop for the air-conditioning system of our data center. This specific control system is called "TUCool" (TU Chemnitz Cooling Approach) (Vodel et al. (2015)). The system is able to adapt the cooling capacity dynamically, dependent on the system behavior and given temperature data. Basically, TUCool is working with a predefined or self-learned rule set. The different parameters are used to calculate possible future environmental conditions inside the data center.

The proposed thermal patterns represent the second stage for this adaptive software-based optimization approach. The fingerprints allow more accurate, faster prediction for the required (near-future) cooling capacity.

If we scale-out the approach to the usage at several locations or institutions, the system allows an open exchange of the fingerprint data. In consequence, an open access database for thermal fingerprints provides the opportunity to minimize efforts for learning the initial setups and schedules. At this point the authors would like to clarify, that such an open access database will be very hard to manage. Due to endless different environmental conditions and operational scenarios, the classification and recognition of known thermal loads at different locations is not that simple.

In discussion of possible worst case scenarios, the following example may summarize the critical points: Here, a similar computational task would be executed periodically on several comparable hardware platforms but with changing environmental conditions. In consequence, a suboptimal data handling results in numberless variations of thermal fingerprints for an identical task. Accordingly, such a scenario generates massive amounts of useless data and makes the identification / recognition process much more complicated. Thus, a robust implementation of the proposed core features is essential for a feasible and efficient solution.

6. Conclusion

The paper has introduced a very promising thermal fingerprint approach for the classification and recognition of thermal loads in heterogeneous data center environments. Administrators are able to optimize the cooling baseline for the given air-conditioning system without additional hardware efforts. This results in significant savings with regards to energy and money. At the same time, we are able to ensure a constant health status for any hardware component within the entire data center.

Furthermore, the proposed solution does not require an initial invest for additional monitoring hardware or system upgrades. TUCool as well as the thermal fingerprint patterns are software-only approaches, which are using the given hardware environment and the given sensor data. In a further step, the approach also allows the integration of different rooms and locations into one central database.

Accordingly, our research goal can be summarized as follows: provide an economic, easy-to-use toolkit to minimize the power consumption as well as the carbon footprint for hundreds of existing, conventional data center.

Based on the results from this proof-of-concept, current and future research work focuses on robust classification and recognition capabilities for this approach.

Acknowledgement: *We gratefully acknowledge the cooperation of our project partners and the financial support of the DFG (Deutsche Forschungsgemeinschaft) within the* Federal Cluster of Excellence EXC 1075 *"MERGE".*

References

Buyya, R., Beloglazov, A. & Abawajy, J. (2010). Energy-efficient management of data center resources for cloud computing: a vision, architectural elements, and open challenges. arXiv preprint arXiv:1006.0308.

Fifield, D., Geana, A., Garcia, M.L. et al. (2015). Remote operating system classification over IPv6. In: Proc. of the 8th ACM Workshop on Artificial Intelligence and Security, 57–67.

Gilbert, P., Ramakrishnan, K. & Diersen, R (2013). From data center metrics to data center analytics: How to unlock the full business value of DCIM. CA technologies: white paper.

Greenberg, S., Mills, E., Tschudi, B. et al. (2006). Best practices for data centers: Lessons learned from benchmarking 22 data centers. In: Proc. of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar. 76–87.

Han, B., Hou, Y. Zhao, L. et al. (2015). A filtering method for audio fingerprint based on multiple measurements. In: Proc. of the Int. Conf. on Information Technology and Computer Application Engineering, Hong Kong, CRC Press. 377-382.

Jia, Y., Yoshii, K., Ross, R. et al. (2011). Wireless sensor network for data-center environmental monitoring. In: Proc. of the 5th Int. Conf. on Sensing Technology, 533–537.

Lent, R (2016). Evaluating the cooling and computing energy demand of a datacentre with optimal server provisioning. In: Future Generation Computer Systems 57, 1–12.

Liang, C.J.M., Liu, J., Luo L. et al. (2009). RACNet: A High-Fidelity Data Center Sensing Network. In: Proc. of the ACM Conference on Embedded Network Sensor Systems, 9–12.

Liu, J., Nath S. Terzis, A. et al. (2011). ThermoCast: A Cyber-Physical Forecasting Model for Data Centers. In: Proc. of the SIGKDD (Conference on Knowledge Discovery and Data Mining), 13–15.

Liu, J. & Terzis, A. (2012). Sensing data centers for energy efficiency. In: Philosophical Transactions of the Royal Society A, 136–157.

Noore, A., Singh, R. Vatsa, M. et al. (2007). Enhancing security of fingerprints through contextual biometric watermarking. In: Forensic Science International, 169(2):188–194.

Pakbaznia, E. & Pedram, M. (2009). Minimizing data center cooling and server power costs. In: Proc. of the International Symposium on Low Power Electronics and Design. 145–150.

Patel, C.D., Sharma, R., Bash, C.E. et al. (2002). Thermal considerations in cooling large-scale high compute density data centers. In: The Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. ITHERM. 767–776.

Ritter, M. (2014). Optimization of algorithms for video analysis: A framework to fit the demands of local television stations. In: Wissenschaftliche Schriftenreihe Dissertationen der Medieninformatik (3). PhD Thesis, TU Chemnitz, Germany, 1–336.

Ritter, M. & Eibl, M. (2011). An Extensible Tool for the Annotation of Videos Using Segmentation and Tracking. In: Design, User Experience, and Usability. Theory, Methods, Tools and Practice, volume 6769 of LNCS, 295–304.

Ritter, M., Heinzig, M., Herms, R. et al. (2014). Technische Universität Chemnitz at TRECVID Instance Search 2014. In: Proc. of TRECVID Workshop, Orlando, FL, USA, 1–8.

Ritter, M., Rickert, M., Juturu Chenchu, L. et al. (2015). Technische Universität Chemnitz at TRECVID Instance Search 2015. In: Proc. of TRECVID Workshop, Gaithersburg, ML, USA, 1–12.

Swaminathan, A., Wu, M. & Liu, K.J. (2008). Digital image forensics via intrinsic fingerprints. In: IEEE Transactions on Information Forensics and Security, 3(1):101–117.

Takei, N., Saito, T., Takasu, K. et al. (2015). Web browser fingerprinting using only cascading style sheets. In: 10th Int. Conf. on Broadband and Wireless Computing, Communication and Applications, 57–63.

Vodel, M. (2014). Energy-efficient communication in distributed, embedded systems. Habilitation thesis, TU Chemnitz, ISBN 978-3-941003-18-7.

Vodel, M., Ritter, M. & Hardt, W. (2015). Adaptive sensor data fusion for efficient climate control systems. In: 9th Int. Conf. on Universal Access in Human-Computer Interaction, 582–593.

Vodel, M., Sauppe, M., Caspar, M. et al. (2008). Simanet–a large scalable, distributed simulation framework for ambient networks. In: Journal of Communications, 3(7):11–19.

Vogel, M.R. Novotny, S. & Lee, M. (2007). Data center room cold aisle deflector. US Patent 7,266,964.

Wang, Y., Xing, G. Wang, X. et al. (2014). Sensor system for high-fidelity temperature distribution forecasting in data centers. In: ACM Trans. Sen. Netw. 11, 2, Article 30, 25ff.

Zanatta, G., Bottari, G. D., Guerra, R. et al. (2014). Building a WSN infrastructure with COTS components for the thermal monitoring of datacenters. In: Proc. of the 29th ACM Symposium on Applied Computing, 1443–1448.