The TUCool Project - Low-cost, Energy-efficient Cooling for Conventional Data Centres

Matthias Vodel* and Marc Ritter*
*Chemnitz University of Technology
*Chemnitz, Germany

email: ^{1}matthias.vodel@hrz.tu-chemmnitz.de,
email: ^{2}marc.ritter@cs.tu-chemmnitz.de

***Abstract*—Air-conditioned cooling concepts still represent the usual cooling solution for thousands of mid-sized data centres. These data centres consists of different types of IT-infrastructure components, as well as different hardware generations. During the last decade, the optimisation regarding energy-efficiency for such central IT locations becomes one of the most important challenges. Green-IT improvements for existing data centres means an adaptive and safe parametrisation of the air-conditioning-system and its control mechanisms. But in order to handle these issues, the available amount of sensor data is critical. A large diversity of distributed sensor devices allows a more precise system management. In this context, the TU Chemnitz develops a cost-efficient and smart solution to improve the sensor knowledge base as well as the control mechanisms. We are using local sensor capabilities within the hardware components and combine these information with actual system loads to create an extended knowledge base, which also provides adaptive learning features. In a first research stage, we analyse the actual cooling environment and measure several operational scenarios for creating a detailed simulation model. The respective results demonstrates a huge optimisation potential. Accordingly, an optimised trade-off between power consumption and cooling capacity may result in significant cost savings.**

Keywords* - *air-conditioning; data centre; optimisation; energy-efficiency; adaptive; sensor fusion; control loop; control system.

I. INTRODUCTION

Traditional data centres are characterised by heterogeneous hardware components and generations. Multiple hardware generations over several decades are running side-by-side. Such locations include all kinds of IT-infrastructure like storage systems, network core components, and server systems. Due to this mixed environment, compromises regarding the cooling capabilities are necessary. Due to physical limitations regarding cooling power and energy density per rack, a large amount of space capacity inside each air-cooled server rack is wasted [1][2]. Accordingly, the optimisation of such traditional air-cooled data centre environments regarding energy- and cost-efficiency is one of the central challenges for hundreds of institutions in the public and educational domain [3].

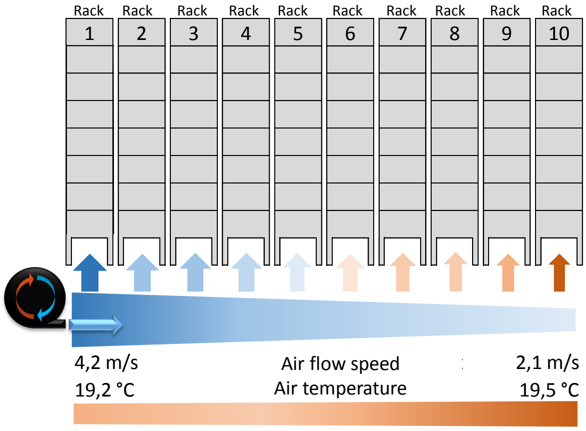

Figure 1. Key problems for traditional, air-cooled data centre environments. In-homogeneous air temperature and air flow speed dependent on the positioning of the server rack. Starting from the air intake on the left side, the cooling capacity shrinks from rack to rack [4].

II. PROBLEM DESCRIPTION

There are two major problems for usual air-cooled data centres: *Inhomogeneous air temperature* and the *inhomogeneous air flow* inside the data centre. These parameters are strongly dependent on the server rack location within the room and even on the position of each individual server component inside the rack. These two challenges are shown in Figure 1 based on measurements in our TU Chemnitz data centre.

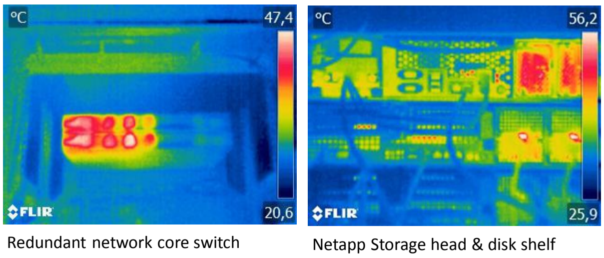

Redundant network core switch

Netapp Storage head & disk shelf

Figure 2. Detailed heat analysis for the air offtake behaviour in different hardware components.

With focus on an entire data centre with multiple server racks and hundreds of server systems, an additional issue becomes critical: Turbulences and interferences between different air flows around the individual racks. These effects have a huge impact on the cooling efficiency. In order to quantify this impact, we are analysing detailed heat images for each hardware component. The heat distribution and model-specific air flow characteristic results in individual heat patterns for each component class. Figure 2 visualises such a component-based heat analysis for two different air offtakes.

Facing these efficiency challenges from an administrative perspective, the monitoring and measuring of the respective values appears in a very basic manner [5][6]. Usual data centre environments only provide a few global temperature sensors for the entire room. Accordingly, the control loop for the air conditioning is very simple. Besides the global room temperature, no further information are available.

III. RELATED WORK

Due to these issues, several professional solutions try optimise this situation regarding monitoring capabilities, sensor data sources, management & control processes as well as cost- and energy savings.

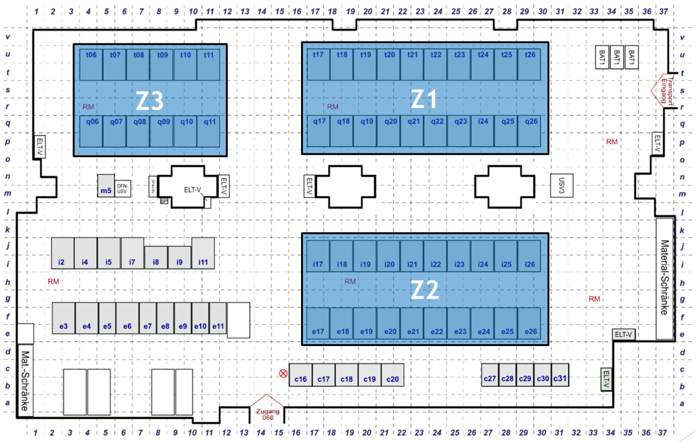

Figure 3. TU Chemnitz data centre with three cold aisle containments, which represent the operational zones Z1, Z2, and Z3.

A. Cold Aisle Containment & Air Boosters

One of the most efficient optimisation steps for traditional air-cooled data centres represents *cold aisle containments*, which allows us to concentrate the cooling capacity directly to the server hot spots within the room. Accordingly, we reduce the effective volume from the entire room space to single enclosures with a significant smaller capacity. Figure 3 shows the three realised cold aisle containments of the TU Chemnitz data centre.

Each containment provides individual temperature sensors and is equipped with optional *booster* elements. The booster technology is shown in Figure 4. As one can see, the boosters allow us to modify the air flow individually for each zone. In order to establish such cold aisle containments, each hardware component has to be re-organised regarding the direction of the air flow. Air intakes have to be located inside the containment, air offtakes outside the enclosures. Accordingly, the installation of these containments is very time-consuming, requires a detailed timeline and is critical with respect to system downtimes or failures.

But anyway, the control cycles as well as the information database for adapting the boosters and the air conditioning system are still the same. The control loops only operate in a static, reactive approach, based on single temperature measurements inside the containment. No further information are available.

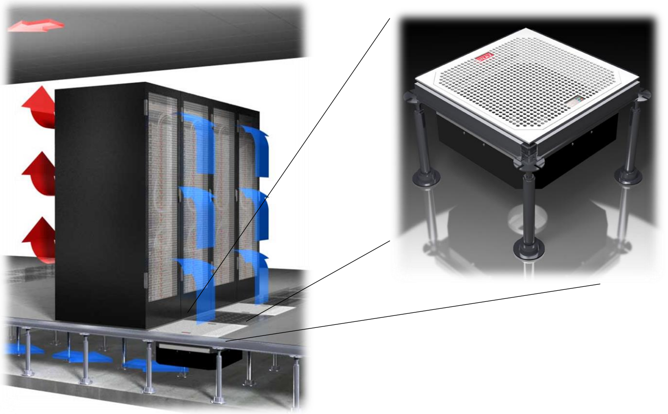

Figure 4. Booster components for dynamic adaptation of the air flow in different, individual housing areas [4].

B. Genome Project

In order to provide a better sensor data knowledge base, Microsoft Research starts the *Genome project* [7][8], which adds dedicated wireless sensor nodes to each server rack. These nodes (called *Genomotes*, see Figure 5) are organised in a master-slave chained sensor network design (*RACNet*), based on the IEEE 802.15.4 low-power, low-data rate communication stack [9]. The RACNet infrastructure provides several information sets about the environmental status, including heat distribution, hot spots, and facility layout. Each node sends its data to a predefined data sink, which creates a global view regarding the health status. The entire raw data is merged together for different data representation tasks (analysis, prediction, optimisation, and scheduling).

C. SynapSense

Another company, which also uses such kinds of sensor nodes is *SynapSense* [10]. Here, several node classes with different types of information are available, e.g., *Therma Nodes* (see Figure 5), *Pressure Nodes*, or *Constellation Nodes*. The data sets from the nodes are processed in a centralised manner by a special software tool, which is able to adapt and to steer the air conditioning system.

All of these approaches possess two critical disadvantages. The first one deals with additional hardware costs for the different sensors. This includes costs for installation, configuration, operation, and maintenance. For large-scaled data centre environments, the required financial resources are very high [11]. The second disadvantage represents the type of data. All of these systems are measuring external parameters from the current point in time, thus providing no learning capability from the past. In addition, there are no server-internal data sources like the system load or any kind of hardware health status as well.

Figure 5. Sensor nodes for data centre monitoring. Left: *Genomotes* from Microsoft Research; Right: *ThermaNode* from SynapSense

IV. SENSOR DATA & SYSTEM MANAGEMENT

Our idea for a more efficient solution is very simple but quite efficient. With our monitoring and control approach, we include different software-based sensor plugins. Each plugin represents a class module for a specific kind of sensor data. In contrast to other approaches, we are using local sensor capabilities of each hardware components inside the data centre. E.g., a given server system typically provides several temperature sensors, located at the mainboard, the CPUs, and the cooling fans (illustrated in Figure 6).

Further information modules are monitoring and learning the system load values of physical/virtual server entities and the respective impact on the data centre temperature behaviour. We are mapping all of these temperature and system load information into one extended knowledge base. Furthermore, different sensor data sources are merged together to more abstract information sets. The fusion results indicate the actual health status of the data centre as well as a prediction trend for the future. Past monitoring data represents a continuous input for machine learning capabilities. In order to save energy and costs, a feasible prediction model [12] is necessary for adapting the cooling power.

Temperature peaks for short term loads and local hot spots are handled with an increased air flow, which means local air booster elements. Such short term situations include hundreds of boot processes of virtual desktops in the morning or backup tasks in the night. Also, small- and mid-size compute jobs for cluster installations may result in such short term temperature peaks.

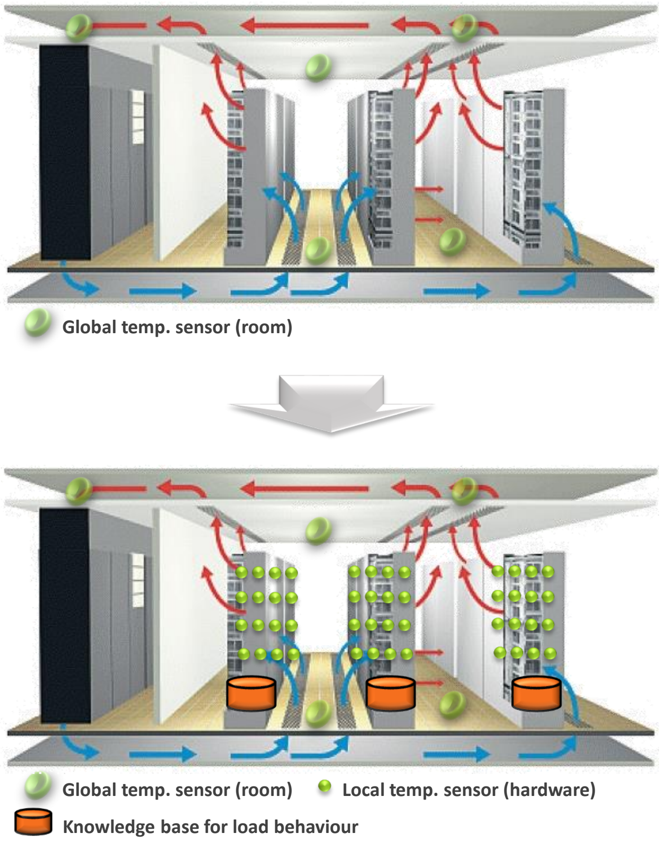

Figure 6. Extension of the knowledge base by using additional sensors and load data from the individual server systems [4].

On the other side, the control system must handle the long term temperature behaviour inside the data centre, e.g., the differences between working days and weekends as well as day & night periods. For such scenarios, the entire air conditioning system with its specific cooling capacity has to be adapted periodically.

V. ANALYSIS & RESULTS

In this paper, we analysed three different impact parameters, based on our data centre environment with three cold and one warm zone.

First of all, we measured the temperature differences between each zone and we identified noticeable influences between the encapsulated cold aisles. Dependent on the positioning of each server system within the zones, nearby area are heated up significantly (see Figure 7). A very simple but efficient solution represents any kind of obstacle between the cold zones. Accordingly, the air interferences are minimised and the temperature impact decreases.

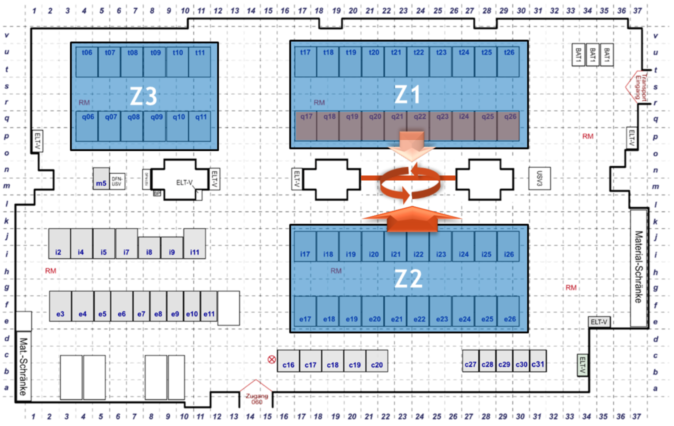

Figure 7. Impact of nearby cold aisles dependent on the hardware location. Indirect temperature increase for isolated cold aisle containments. The red line represents any kind of separating obstacle between these zones.

A second result deals with the hardware diversity and the hardware distribution within the data centre. We found out that a mixed positioning of different hardware types (storage, compute, network) in the available zones is much more energy-efficient than an organised positioning. If we put all the compute and server systems in one zone, we create hot spots regarding the temperature. Due to the fact that we only have one global air conditioning system, we have to increase the basic cooling capacity (and the respective power consumption) for cooling down only one critical zone.

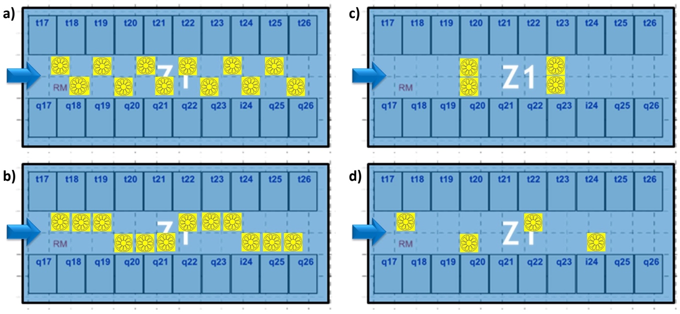

Figure 8. Different booster configurations in a given cold aisle containment.

A third optimisation result focuses on the booster technology. In order to find an optimal trade-off between hardware efforts and cooling efficiency, the amount of booster elements is critical. But there is no direct linear relation between the amount of boosters and the cooling capacity. Furthermore, the booster locations are relevant. Figure 8 visualises this topic.

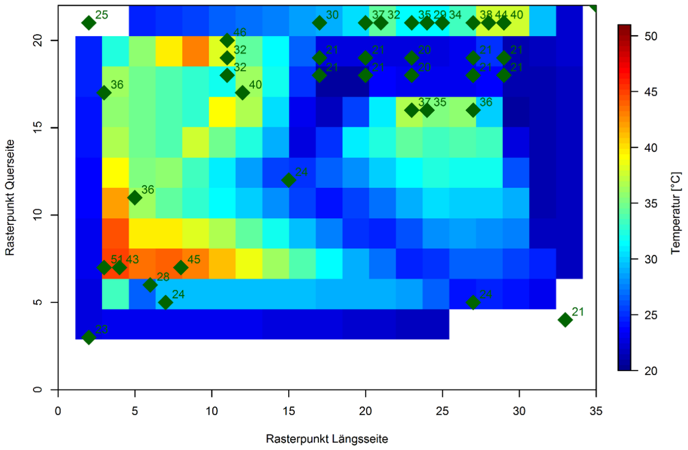

Figure 9. TU Chemnitz data center heat map [4]. Hot spots without cold aisle containment in the bottom left corner are clearly visible.

Type a) and b) is very hardware-extensive but the benefits in comparison to type d) are minimal. In contrast, the cooling efficiency of type d) is significant better than type c) while using the same amount of boosters. Accordingly, the well-organised booster positioning allows excellent cooling capabilities without massive hardware efforts.

All measurements in our data centre are monitored over a time period of three months. We visualise the results in complex heat maps as shown in Figure 9. The sample rate for all temperature measurements is 1 data set per minute, leading to more than 100000 data points per sensor. Server-internal system load measurements are monitored with 1 data set per second to calculate a feasible average value per minute. Accordingly, we map global and local temperature values together with local system loads and fan speed information. Hence, the merged data is available in one central knowledge base to allow a detailed and situation-specific adaptation of the cooling capacity.

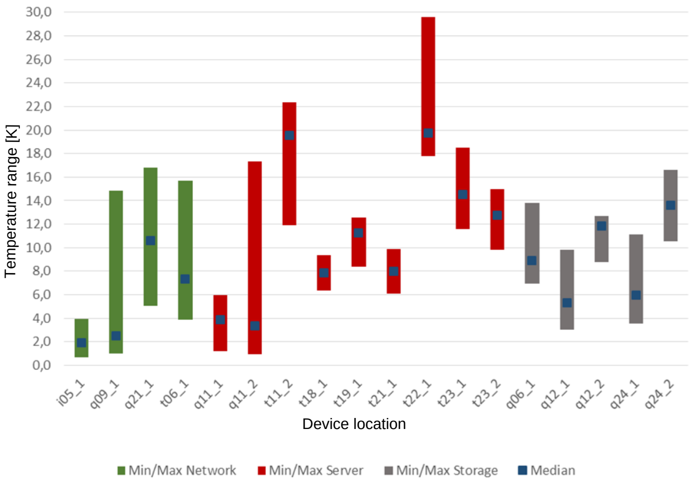

Figure 10. Measured temperature ranges for several hardware components at different locations (coordinates referenced in Figure 3).

Finally, we illustrate the operational temperature ranges for different hardware components during the tests. Figure 10 shows minimum, maximum, and average values.

VI. CONCLUSION

In this paper, we analysed the energy-efficiency of heterogeneous data centre environments with traditional air cooling systems. We introduced a cost-efficient and smart approach for improving the sensor data in order to implement an adaptive management system for the cooling capacity. The proposed concept does not require any further hardware components or installation efforts. The system utilizes given sensor sources from each hardware system and aggregates these data sets into one single knowledge base. Furthermore, we also discuss the impact of weak parameters like the hardware positioning inside the data centre and its cold aisles. Also, the booster configuration, as well as influences between different cold and warm zones were analysed. As already mentioned, improved control cycles for the cooling systems allow us to reduce the basic cooling level. Short term temperature peaks can be avoided with a local adaptation of the air flow using booster components. Accordingly, we are able to save massive energy and costs without any increased risk level for the hardware [4]. In this context, the presented results offer high potential for optimisations. The research goal deals with the vision of optimising an entire data centre environment based on extensive simulation processes without any trial-and-error approaches on the real hardware.

ACKNOWLEDGEMENT

This research work is supported by the *RenewIT* project, which is co-financed by the 7^{th} Framework programme of the European community (FP7) under grant agreement number 608679. All measurements were done within the data centre environment of the *Technische Universität Chemnitz*.

REFERENCES

[1] R. Lent. Evaluating the cooling and computing energy demand of a datacentre with optimal server provisioning. *Future Generation Computer Systems 57*, pages 1–12, April 2016.

[2] Y. Wang G. Xing X. Wang B. Punch J. Chen, R. Tan and D. Colbry. Sensor system for high-fidelity temperature distribution forecasting in data centers. *ACM Trans. Sen. Netw. 11, 2, Article 30*, page 25 ff., December 2014.

[3] R. Guerra G. Zanatta, G. D. Bottari and J. Leite. Building a wsn infrastructure with cots components for the thermal monitoring of datacenters. In *In Proceedings of the 29th Annual ACM Symposium on Applied Computing (SAC '14)*, pages 1443–1448, 2014.

[4] M. Vodel and M. Ritter. Adaptive Sensor Data Fusion for Efficient Climate Control Systems. In *Proc. of HCII 2015 - International Conference on Human-Computer Interaction*, pages 1–11. Springer, 2015.

[5] J. Liu and A. Terzis. Sensing data centers for energy efficiency. *Philosophical Transactions of The Royal Society A*, pages 136–157, January 2012.

[6] M. Vodel. *Energy-efficient communication in distributed, embedded systems*. PhD thesis, TU Chemnitz, Chemnitz, February 2014. ISBN 978-3-941003-18-7.

[7] Microsoft Research. DC Genome Project, 2009. [Online; last access May 2016].

[8] US Dept. of Energy. Wireless sensors improve data center energy efficiency. Technical Report CSO 20029, USA, September 2010.

[9] L. Luo A. Terzis C.-J. Mike Liang, J. Liu and F. Zhao. RACNet: A High-Fidelity Data Center Sensing Network. In *Proc. of the SenSys (ACM Conference on Embedded Network Sensor Systems)*, pages 9–12. ACM, November 2009.

[10] SynapSense. Active Control. http://www.synapsense.com, 2015. [Online; last access May 2016].

[11] Y. Jia K. Yoshii R. Ross M. G. Rodriguez, L. E. Ortiz Uriarte and P.H. Beckman. Wireless sensor network for data-center environmental monitoring. In *Proc. of the 5th International Conference on Sensing Technology (ICST)*, pages 533–537, Nov 2011.

[12] J. Liu S. Nath A. Terzis L. Li, C-J. Mike Liang and C. Faloutsos. ThermoCast: A Cyber-Physical Forecasting Model for Data Centers. In *Proc. of the SIGKDD (Conference on Knowledge Discovery and Data Mining)*, pages 13–15. ACM, August 2011.