

# Adaptive Sensor Data Fusion for Efficient Climate Control Systems

Matthias Vodel<sup>1</sup>✉, Marc Ritter<sup>2</sup>, and Wolfram Hardt<sup>1</sup>

<sup>1</sup> University Computer Centre, Chemnitz, Germany  
marc.ritter@informatik.tu-chemnitz.de


<sup>2</sup> Junior Professorship Media Computing, Technische Universität Chemnitz,  
09107 Chemnitz, Germany  
{matthias.vodel,wolfram.hardt}@hrz.tu-chemnitz.de

**Abstract.** Thousands of data centres are using traditional air-conditioned cooling concepts for the entire payload. Most of these data centres include multiple hardware generations and different types of IT-infrastructure components, i.e. storage, compute, and network devices. In the context of Green-IT, an efficient and safe parameterization of the air-conditioning-system is essential - keep the temperature as low as necessary, but not too low. Usually, only a few amount of temperature sensors are available to handle these important control cycles. But in order to optimise the cooling capacity, several scenario-specific parameters have to be considered, including the shape of the room, air flow, or component placements. In this context, the TU Chemnitz develops novel concepts to improve this process. 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. First measurement scenarios show huge optimisation potential. The respective trade-off between power consumption and cooling capacity results in significant cost savings.

**Keywords:** Data centre · Air conditioning · Adaptive · Sensor · Data fusion

## 1 Introduction

The optimisation of 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. Multiple hardware generations over several decades are running side-by-side. New hardware components provide a significantly higher energy density and accordingly, the respective cooling capacity becomes a critical issue. Due to physical limitations regarding cooling power and energy density per rack, a large amount of space capacity inside the air-cooled server racks is wasted (see Fig. 1). In order to improve this situation, we have to analyse the key parameters, which have a direct impact on the cooling efficiency.



**Fig. 1.** Key issue within traditional, air-cooled data centre environments – optimising the rack efficiency by maximising the filling level.

## 2 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 Fig. 2 based on measurements in our TU Chemnitz data centre.

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.

Facing these efficiency challenges from an administrative perspective, the monitoring and measuring of the respective values appears in a very basic manner [1, 2]. 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.

## 3 Related Work

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

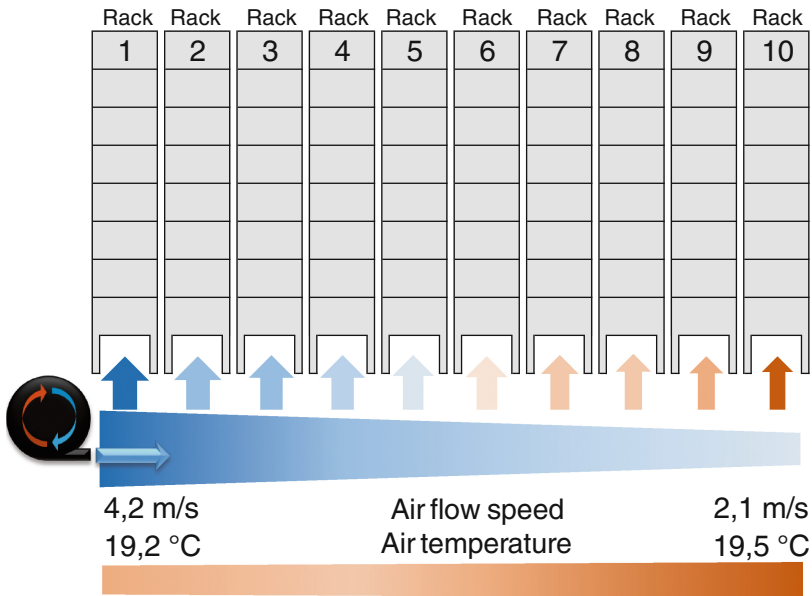


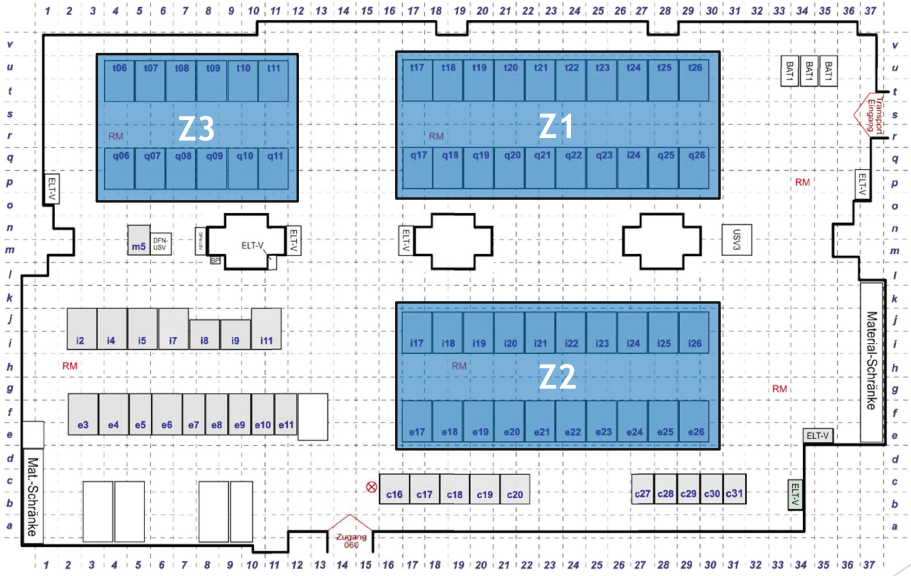
Fig. 2. Key problems for traditional, air-cooled data centre environments. Inhomogeneous air temperature and air flow speed dependent on the positioning of the server rack. Starting from the air intake on the left side, the cooling capacity shrinks from rack to rack

### 3.1 Cold Aisle Containmentment and 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 Fig. 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.



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


### 3.2 Genome Project

In order to provide a better sensor data knowledge base, Microsoft Research starts the *Genome project* [3,4], which adds dedicated wireless sensor nodes to each server rack. These nodes (called *Genomotes*) 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 [5]. 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).

### 3.3 SynapSense

Another company, which also uses such kinds of sensor nodes is *SynapSense* [6]. Here, several node classes with different types of information are available, e.g. *Therma Nodes*, *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



**Fig. 4.** Booster components for dynamic adaptation of the air flow in different, individual housing areas.

centre environments, the required financial resources are very high [7]. 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.

Nevertheless, all of these solution offer the same benefits, which are equal to the objectives of our research work:


- Enabling real-time monitoring & control
- Optimised change management
- Optimised capacity planning
- Optimised server positioning and provisioning
- Optimised fault diagnostics
- Optimised energy- & cost-efficiency (TCO)

## 4 TU Chemnitz Adaptive Cooling Approach

Based on the related research projects and products, we developed a more flexible, more cost-efficient and smarter solution for heterogeneous, air-cooled data centre environments. *TUCool* denotes our adaptive cooling approach at TU Chemnitz. Instead of using dedicated measurement hardware, we decided to use the already available hardware components inside the data centre. Accordingly, each single server system, each network core switch, each storage system becomes an additional sensor source for environmental data.

#### 4.1 Knowledge Base Extension with Sensor Data Fusion

The idea is simple but quite efficient. With the TUCool monitoring and control approach, we include different sensor plugins. Each plugin represents a class module for a specific kind of sensor data. A given server system typically provides several temperature sensors, located at the mainboard, the CPUs, and the cooling fans (illustrated in Fig. 5).



**Fig. 5.** Extension of the knowledge base by using additional sensors and load data from the individual server systems.

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. Accordingly, TUCool is able to map temperature and

system load information for an efficient adaptation of the cooling capacity. 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 the machine learning capabilities.

## 4.2 Adaptive Control Loop

The core control mechanism for the air conditioning is operating like a PID element (*Proportional plus Integral plus Derivative action*). In order to save energy and costs, a feasible prediction model [8] is necessary for adapting the cooling power. The TUCool system has to handle two control parameters for different cooling scenarios.

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.


On the other side, the TUCool 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.

In general, TUCool with its extended knowledge base is able to differentiate between short term and long term adjustments. From the physical perspective, we are able to balance short term temperature peaks with an increased airflow. In consequence, one key benefit of such a system is the possibility to increase the local cooling capacity without adapting the main air conditioning system. With these control features and this sensor knowledge base, we are now able to reduce the basic cooling level for saving massive amounts of energy. The prediction system avoids short term temperature peaks without any disadvantages for the hardware or the data centre health status.

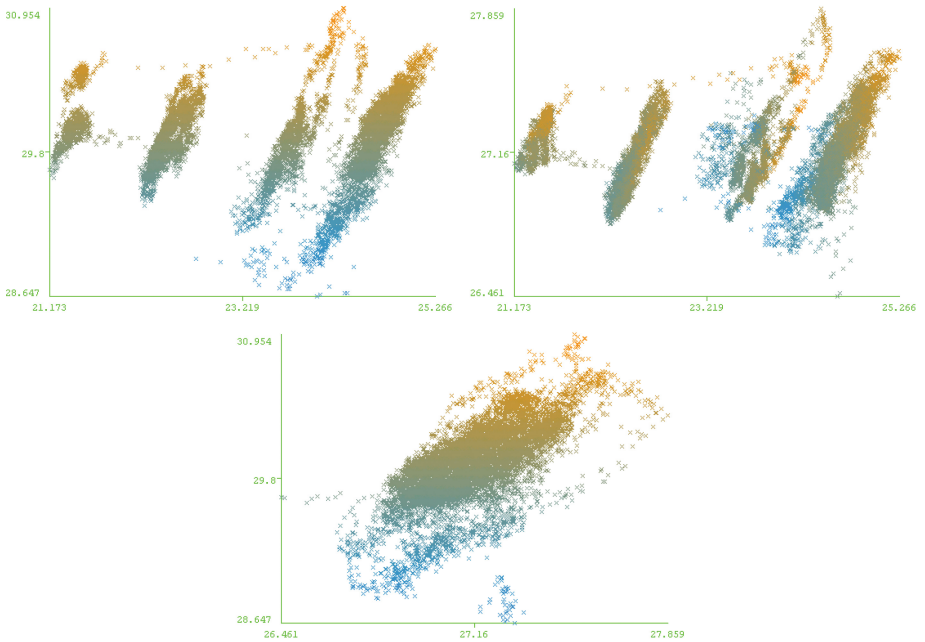
**Static Constraints.** In order to control the cooling system, respective policies or rule sets are necessary. For defining these rule sets, two approaches are possible. The first one deals with static definitions, which are situation-specific predefined by the administrator. The different policy classes can be structured as follows:

- *Temperature hot spot* (local short term thermal peak) → increase booster level
- *System load peak* (local behaviour of a server bay or rack) → increase booster level
- *Average zone temperature* (cooling zone hits a predefined thermal value) → increase cooling capacity





**Fig. 6.** TU Chemnitz data center heat map. Hot spots without cold aisle containment in the bottom left corner are clearly visible.



**Fig. 7.** Co-occurrences of temperatures from three different and spatially distributed sensor groups: q10 vs. i05 (top left), q10 vs. t21 (top right), and t21 vs. i05 (bottom).



- *Time slot entered* (predefined, time-specific behaviour) → increase/decrease cooling capacity / booster level

These static rules represent a basic set of control mechanisms for a given cooling system. In contrast to related research projects, we are focusing on both internal and external sensor data for adapting the cooling behaviour.

**Machine Learning Capabilities.** For further improvements, our future research work deals with automated processes for a continuous optimisation of the entire cooling system. This represents the second control approach. Starting from a static rule set, the system has to provide self-learning features. Accordingly, such a control system is switching from a *re-active* adaptation of the cooling capacity to a *pro-active* adaptation of the respective rule sets. The input for the machine learning features consists of different types of data as well as different time periods. Especially in the context of data centre environments, the knowledge about frequently repeating event in time is very helpful for optimising the energy-efficiency of the cooling system.

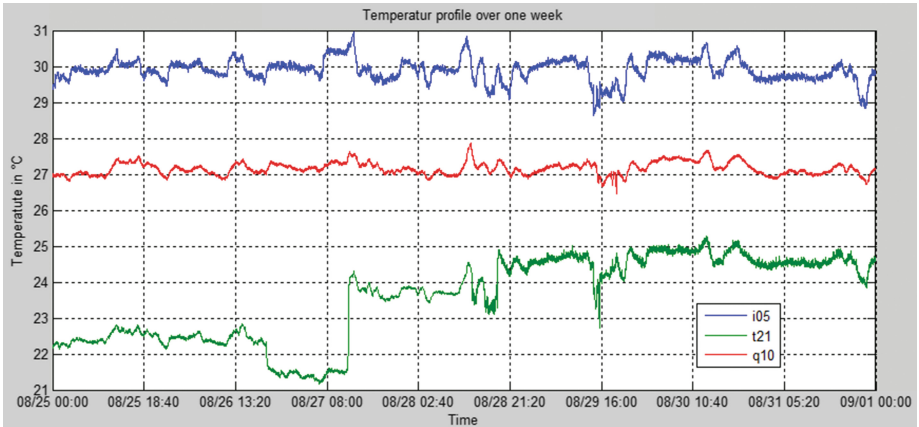
Another key benefit deals with the efforts for maintenance and the adaptation processes. Time resources of IT administrators are limited and accordingly, adaptations and optimisations for the cooling system are very cost-intensive. For future research work in the TUCool project, we only want to define one initial rule set as well as some safety limits for the entire cooling system. The continuous monitoring and control process will be executed by the management software without further manual efforts.

## 5 Measurement Scenario and Results

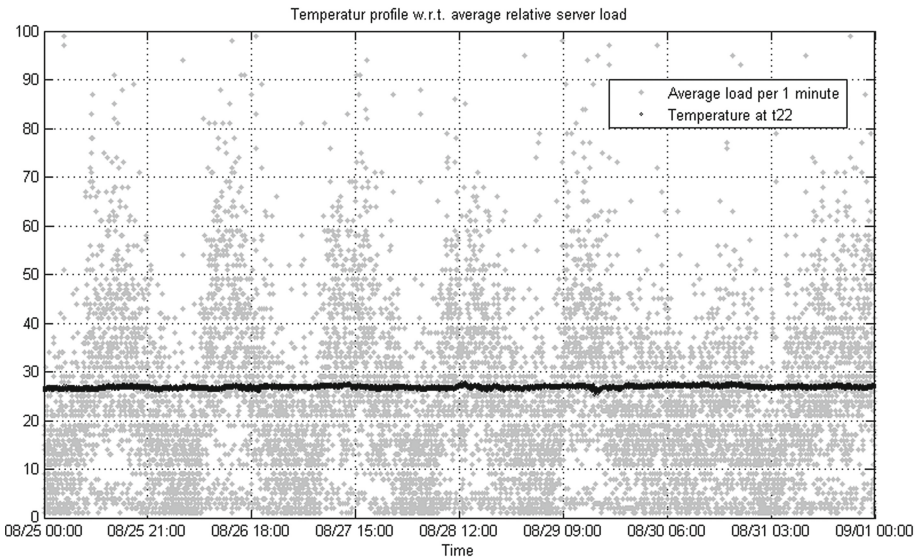
In this paper, we analysed several server systems and time periods in our data centre. We focused on the efficiency of the the current control loops. Therefore, we installed sensors to measure the room temperature at different locations as well as within the hardware components. We found correlations between neighbored sensors areas. We focus on multiple spatially distributed locations: q10 represents the temperature output in the area Z1, t21 in Z2, and i05 in the remaining warm aisle. All locations are mapped in Fig. 3.

In order to verify our approach, we measured the sensor profiles of these components over a time period of one week, starting at Monday, 25th of August 2014 at 00:00 A.M. while ending at Sunday, 31st of August at 11:59 P.M. We subsampled all measurements to 1 sample per minute leading to 10,080 data points per sensor. The resulting heat map of our data centre environment is shown in Fig. 6.

In addition, correlations of the occurring temperatures between the three different sensors are illustrated in Fig. 7, yielding to three different temperature distributions. The combination with the time series plot of these sensor values in Fig. 8 confirms a certain degree of dependency between those areas leading



**Fig. 8.** The one week temperature profile shows three sensors in three locations from above yielding to distinct temperature zones. These were controlled by using the proposed adaptation scheme thus being capable of stabilizing the interconnections of the varying temperature gradients.



**Fig. 9.** The relative average CPU load per minute of a DeLL 2960 server located at position t22 (grey dots) does not affect the overall cooling temperature in the region at t21 (black line).

to similar gradients. A globally small but locally noticeable change in the temperature (represented by some spikes) is visible. This results from periodically executed tasks like server maintenance, software distribution processes, virtual desktop management, and storage deduplications. Further relevant processes

include virtualization cluster boot-up tasks each morning as well as backup task for all critical services at night.

As mentioned before, we also recorded the inputs and outputs of temperature sensors within the hardware components. In this context, Fig. 9 visualises the relative CPU load of a server system with relation to its output temperature. Despite different, recurring and intense shifts in work loads, the local temperature can be kept at a stable level.

Finally, we illustrate the temperature ranges for hardware components in operation during the tests in Fig. 10 including minimum, maximum, and average values. These different profiles are consistent with the different hardware types that range from diverse server systems over large storage devices to network switches.

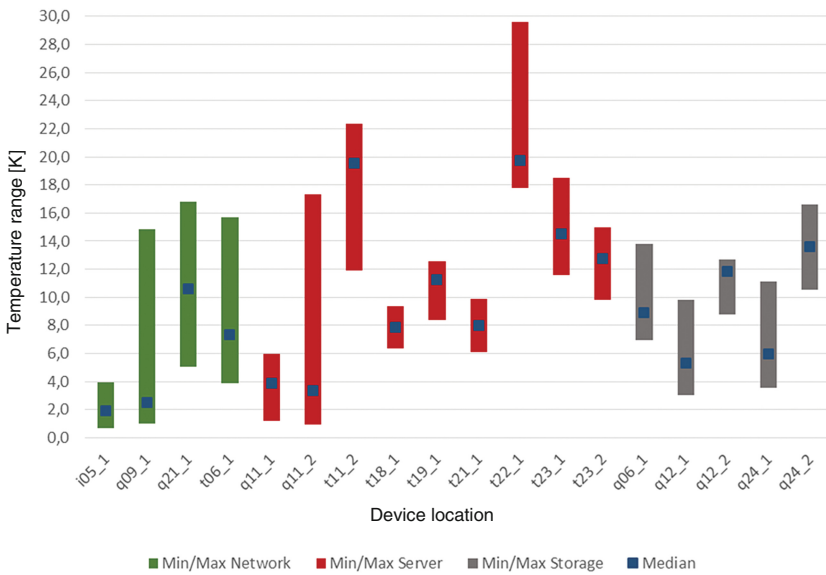


Fig. 10. Measured temperature ranges for several hardware components at different locations.

## 6 Conclusion and Future Work

In this paper, we presented TUCool, an innovative approach for optimising heterogeneous data centre environments with traditional air cooling systems. In contrast to other professional solutions and research projects, TUCool 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. Accordingly, based on this sensor data fusion approach, TUCool is capable of controlling and optimising the entire cooling system automatically and continuously over the time. This results in massive energy and cost savings.

In the next project steps, we want to use these measurements and results to develop a detailed simulation model for heterogeneous data centre environments. The research goal deals with the vision of optimising an entire data centre environment based on extensive simulation processes without trial-and-error approaches using real hardware. Critical optimisation parameters might include energy- and cost-efficiency as well maintenance efforts and load balancing.



**Acknowledgement.** This research work is supported by the *RenewIT* project, which is co-financed by the 7<sup>th</sup> Framework programme of the European community (FP7) under grant agreement number 608679. Furthermore, the work was partially accomplished within the InnoProfile-Transfer-Initiative *localizeIT* under grant number 03IPT608X, funded by the *Federal Ministry of Education and Research* (BMBF, Germany) in the program of *Entrepreneurial Regions*. All measurements were done within the data centre environment of the *Technische Universität Chemnitz*.

## References

1. Liu, J., Terzis, A.: Sensing data centers for energy efficiency. *Philos. Trans. R. Soc. A* **370**, 136–157 (2012)
2. Vodel, M.: Energy-efficient communication in distributed, embedded systems. Ph.D thesis, TU Chemnitz, Chemnitz, February 2014. ISBN 978-3-941003-18-7
3. Microsoft Research. DC Genome Project (2009). <http://research.microsoft.com/en-us/projects/dcgenome/>. [Accessed: 2015/02/18]
4. US Department of Energy. Wireless sensors improve data center energy efficiency. Technical report CSO 20029, USA, September 2010
5. Liang, C-J.M., Liu, J., Luo, L., Terzis, A., Zhao, F.: RACNet: a high-fidelity data center sensing network. In: *Proceedings of the SenSys (ACM Conference on Embedded Network Sensor Systems)*. ACM, November 2009
6. SynapSense. Active Control (2015). <http://www.synapsense.com> [Accessed: 2015/02/18]
7. Rodriguez, M.G., Ortiz Uriarte, L.E., Jia, Y., Yoshii, K., Ross, R., Beckman, P.H.: Wireless sensor network for data-center environmental monitoring. In: *Proceedings of the 5th International Conference on Sensing Technology (ICST)*, pp. 533–537, November 2011
8. Li, L., Liang, C-J.M., Liu, J., Nath, S., Terzis, A., Faloutsos, C.: Thermocast: a cyber-physical forecasting model for data centers. In: *Proceedings of the SIGKDD (Conference on Knowledge Discovery and Data Mining)*. ACM, August 2011