

Politechnika Gdańska

Wydział Elektrotechniki i Automatyki

Gdańsk 2015

ACTIVE MANAGEMENT OF EQUIPMENT COOLING IN HOTELING DATA CENTERS

Artur OPALINSKI

Politechnika Gdańska, Wydział Elektrotechniki i Automatyki
tel.: 58 347 1347 e-mail: Artur.Opalinski@pg.gda.pl

Abstract: Hoteling data centers are designated for housing computing and storage units of many, usually small customers, as opposed to traditional data centers supporting own computing and storage resources of a bigger company. One of the services to be provided to consumer's equipment is cooling. Cooling in data centers is prevalently achieved by circulating air in computer room. Efficient cooling requires delivering cold air from central cooling units directly to the equipment that needs it, in appropriate quantity, and at the right time. But with centralized cooling, hot spots and cold spots arise in computer room, caused by uneven, uncontrolled heat generation. The proposed active cold air flow management is to periodically input data from sensors into a model which calculates parameters at every floor localization based on measurements taken at the most reliable measurement locations. Such model has been created and verified by implementation in an industrial data center.

Keywords: air flow, modeling, management, energy efficiency.

1. INTRODUCTION

Hoteling data centers are designated for housing computing and storage units of many, usually small customers, as opposed to traditional data centers supporting own computing and storage resources of a bigger company. Hoteling data centers only provide infrastructure, including physical space, networking connectivity, power and cooling. Customers provide their computing and storage resources to be installed in the physical cabinets provided by hoteling data center.

One of the services to be provided to consumer's equipment is cooling. Cooling in data centers is prevalently achieved by circulating air in computer room. Air is warmed by working computing and storage nodes and exhausted through their fans. This hot air is cooled by computer room air conditioner (CRAC) units which return cool air back into the plenum, to be re-used for cooling.

Efficient cooling requires delivering cold air from CRACs directly to the equipment that needs it, in appropriate quantity, and at the right time. Hoteling data centers are more troublesome in this regard, as they leave little freedom both in physical placement of heat-generating equipment, as well as in balancing processing between nodes. Physical placement is mostly influenced by grouping consumers physical assets, while processing load is totally dependent on customer activity, i.e.

uncontrollable and unpredictable from data center point of view. This makes cooling in hoteling data centers more challenging than in typical data centers.

Additionally, cooling requirements cannot easily be specified by customers; the factor most used is inlet temperature of the cooling air. But this factor is not enough to forecast the amount of heat energy dissipated by equipment in time, i.e. not enough to forecast the right quantity and timing of cool air delivery.

2. STATE OF THE ART

Energy consumption in data centers worldwide is approximated at $3 \cdot 10^{11}$ kWh [1], while mean power consumed by a datacenter is approximated at 2.6 MW [2]. This power is consumed on a relatively small area, which creates challenges not only to produce but also to transfer energy, and to remove the resulting heat [3]. Removing heat, which is generated in significant amounts on the small area has a key meaning, as computing and storage equipment should be operated in precisely prescribed conditions. Not adhering to conditions of ambient air temperature and relative humidity results in increased fault ratio and may void product guarantees [4].

So both economical as well as ecological factors influence the desire to reduce electrical energy consumption in data centers. Hoteling data centers have less possibilities to reduce energy consumption, as only efficiency and optimization of infrastructure can be influenced by managers. The biggest share of infrastructure energy is consumed by cooling systems [5] in such data centers..

Modern data center cooling is typically done in the following way: air is pre-conditioned by CRACs to the given values of temperature and humidity. The air is then pushed under raised floor, which constitutes a transport channel to deliver cold air to all parts of the computer room. The air escapes to the room through perforated tiles, which should preferably be installed near the air intake of the equipment. The more heat generated by equipment, the more perforated tiles may be needed. Cooling air is then sucked by equipment's fans and flows through equipment's chassis, usually from its front to its rear. At this stage heat from the equipment is absorbed by air. Hot air escapes through equipment rear to the open

plenum. Hot air from the open plenum is taken again by CRACs, and recirculated. To avoid mixing of cold air coming from under the floor with the hot air already in plenum, plenum gets sometimes divided into open or closed aisles, which separate equipment air intakes from equipment air outlets. Equipment is installed in industrial racks, which stay in rows alongside aisles. In older solutions perforated tiles were located under racks [6].

Open cold or hot aisles are sufficient up to approximate load of 4 kW per rack [7], but currently rack load can exceed 8 kW [3], therefore closed aisles are increasingly used.

But with centralized cooling, hot spots and cold spots arise in computer room, caused by uneven, uncontrolled heat generation. This results in inefficient cooling: to ensure environmental conditions at hot spots, cold spots get unnecessarily over-cooled. Two basic methods to counter this in traditional data center are:

- Distributing equipment among racks so that heat is generated evenly in computer room space [9]. This is not possible in hoteling data centers, where consumer decides on amount of equipment placed in the racks assigned to him.
- Matching the amount of cold air with local heat generation, by using tiles with different perforation or tiles with regulated vents [10][11]. This is not possible in hoteling data centers, where changes in processing load and thus in heat generation are dependent on consumer decisions and very variable.

Due to the limits of the above methods, only very approximate reduction of temperature differences in computer room is possible. Additionally none of them addresses the dynamics of consuming electrical energy by customer equipment, and thus the space-distribution and dynamics of heat generation. Therefore energy flow in data centers is currently mainly managed statically. Racks are sub-optimally equipped from heat-generation point of view, as the equipment is bound by other requirements and limitations [12].

Providers of cooling systems and providers of computing equipment recommend only very general guidelines to balance electrical and heat energy (kW to BTU/h). These general recommendation additionally assume installed power, and not the level of power actually used. They also neglect the specifics of cooling medium flow, which features varying parameters, despite common space under floor, due to the area of the floor, under-floor obstacles, and local air flow intensity. So adhering just to these general recommendations, while ensuring a total, lump balance of energy, and ensuring environmental conditions required by equipment, does not provide for optimal consumption of electrical energy by cooling systems [13].

Currently advised and used solutions, aiming at reduction of electrical energy used by cooling systems, encompass:

- IT services virtualization and consolidation. This trend does not pertain to hoteling data centers, however.
- Using static air directors which either separate cold and hot air, or direct air flows accordingly
- Free cooling, i.e. use of outer air, or natural water reservoirs to cool air in the computer room.
- Changing power supply components of IT equipment to reduce losses during local power

conversion, e.g. by using low-voltage alternating current mains, increasing frequency of mains, or using high-voltage direct current [14].

3. SOLUTION

The general solution to the above issues lies in active management and optimization of air flow, so as to direct cooling air according to electrical energy consumption at every computer room localization. But data from temperature, humidity, and flow sensors cannot be reliably used for decision taking, as air flow under floor is very perturbed. Therefore the complete solution is to periodically input data from sensors into a model which calculates parameters at every floor localization based on measurements taken at the most reliable measurement locations. The most reliable measurements describing computer room state are:

- electric energy consumption for each rack,
- localization and state (on or off) of CRACs,
- air down-flow forced by each CRACs,
- the localization and opening rate of floor vents.

After updating with current state, the model offers support for operator decisions to change cold flow pattern, mainly by manipulating opening state of floor vents. Implications of different decisions can be easily modeled by introducing new data into the model. This also includes switching CRACs on or off, or changing the level of air flow forced by CRACs.

Optimal management of environmental conditions in data center is difficult because cooling systems are centralized, while the prescribed conditions must be generally ensured in each point of the computer room, even with unequal heat generation from different equipment or different customers. Heat generated by equipment with its dynamics is a disturbance from the control system point of view (Fig. 1). It is usually assumed that all electrical energy (E_{in}) provided to equipment is eventually dissipated as heat, i.e. that energy escaping in the form of electromagnetic emissions is negligible. Therefore, with assigning E_{out} to the amount of heat energy removed, change in ambient temperature ΔT can be described as:

$$\Delta T(t) = f(E_{in}(t), E_{out}(t)) \quad (1)$$

In hoteling data centers, electrical energy (E_{in}) is out of control as it must be provided at the time and in the amount instantaneously consumed by customer's equipment. Therefore only amount of removed heat (E_{out}) is the component under control.




Fig. 1. Ambient temperature as result of imbalance between electrical energy (E_{in}) provided for equipment and heat energy (E_{out}) taken from the equipment. Data center designed as P , control system designed as C

To model flow of air under the floor, a grid approach is used. Each tile is a grid node. Flows under each tile are considered in four basic horizontal directions: North, East, South, West, designated as: F_N , F_E , F_S , F_W , respectively. Additionally for each node vertical direction is considered: up, i.e. air escape to the computer room (F_U) and down, i.e. air moved under the floor (F_D). Flows to the East, to the South and flow down are considered positive. For each node of the grid a single equation is created based on the law of conservation:

$$F_N + F_E + F_S + F_W + F_U = F_D \quad (2)$$

For a data center floor with R tile rows and C tile columns applying (2) to each tile results in $R \times C$ equations, which can be represented in matrix notation as:

$$F_{R \times C} \cdot R_{C \times I} = \Theta_{R \times I} \quad (3)$$

Equation matrix is sparse, and can be numerically solved by LU (lower upper) decomposition. Model parameters are: horizontal air flow resistance between neighboring tiles, vertical air flow resistance through vent tiles (considering its opening), and velocity of air forced by each CRAC. As a result, matrix $P_{R \times C}$ is obtained, where cell values represent pressures under each tile.

4. VERIFICATION

The model has been implemented in software as a C++ program with LAPACK/BLAS (Basic Linear

Algebra Subprograms) [15] library for algebraic matrix computations, and with wxWidgets library [16] for visualization. Results of solving (3) are represented by program (see Figure 2) as a map of data center floor, with each tile colored accordingly to pressure under this tile. Additionally resultant horizontal flow vector for each tile is calculated and presented as arrow on the tile. This feature significantly improves understanding of reasons of observed air flow phenomena.

All model parameters are easily editable in graphic window. Recalculation of matrix $P_{R \times C}$ by LU decomposition takes less than one second on a 16 GB, 2.2 GHz Intel i7 computer under 64-bit Windows (program compiled for 32-bit mode). This supports user in taking decisions, as various combinations of CRAC states and various degrees of floor vents opening can be checked without labor on the real plant. Model results are stable, as opposed to volatile air flow measurements, which gives a clear picture of data center cooling state.

Model parameters in regard to flow resistances, tile vent characteristics, and CRAC air flows have been validated by conducting measurements of flows under floor with a commercial anemometer at 127 tiles in all parts of data center floors. Model-calculated values differed from the measurements by less than 20 % for 101 tiles, i.e. in case of 79 % of all validated tiles, located in all parts of computer room. Obtaining a better result was not possible due to high variability and volatility of air flow measurement results.




Fig. 2. Results obtained from selected flow model, visualized in computer program written by the Author

5. CONCLUSION

The approach with modeling cold air flows in addition to taking measurements of most reliable computer room parameters proved correct. A decision support system has been created, which allows to model, visualize, and check possible scenarios for different

configurations of cooling system. Decision support system response time is acceptable low.

The model has been validated and the system has been implemented in commercial data center of Pomorskie Centrum Przetwarzania Danych.

6. ACKNOWLEDGEMENT

The research described in this paper has been conducted based on data center of Pomorskie Centrum Przetwarzania Danych, as part of project Developing active heat transfer management system for data centers, financed from European Union Funds by Ministry of Infrastructure and Development, Operational Programme Innovative Economy, Activity 1.4. Support for Special Purpose Projects.

7. REFERENCES

1. Koomey, Jonathan, Worldwide electricity used in data centers. *Environmental Research Letters*. vol. 3, no. 034008. Available online, URL: <http://stacks.iop.org/1748-9326/3/034008> (DOA:2015.09.01). 2008.
2. Koomey, J.: Growth in Data center electricity use 2005 to 2010. Oakland, CA: Analytics Press. August 1, 2011.
3. Hachman M.: Data Centers Expanding, Along with Power Budgets: Survey, available online, URL:<http://slashdot.org/topic/datacenter/data-centers-expanding-along-with-power-budgets-survey/>, April 2013.
4. Alfonso Capozzolia, Gianluca Seralea, Lucia Liuzzoa, Marta Chinnici, Thermal Metrics for Data Centers: A Critical Review, 6th International Conference on Sustainability in Energy and Buildings, SEB-14, Energy Procedia, Vol.62, pp.391–400, 2014.
5. Neudorfer J.: DataCenter Knowledge, Data Center Energy Efficiency, Executive Guide Series, Part3, 2012.
6. Van Geet O.: Trends in Data Center Design - ASHRAE Leads the Way to Large Energy Savings, ASHRAE Conference backed by NREL (national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy), Denver, June 24, available online, URL: <http://www.nrel.gov/docs/fy13osti/58902.pdf> (DOA 2015.09.01), 2014.
7. Data Center Cooling Efficiency & Containment, DataCenter Experts, available online, URL: <http://www.datacenterexperts.com/products/data-center-cooling-efficiency-and-containment.html> (DOA: 2015,09.01).
8. Ana Maria Juan Ferrer, Jérôme Brun, Mathieu Peyral, Mick Symonds, Chee Tan, Command and control for Data Centers, ATOS, 2014.
9. Samadiani E.: Energy efficient thermal management of data centers via open multi-scale design, PhD dissertation, G.W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, 2009.
10. Tang Q., Mukherjee T., Gupta S. K. S, Cayton P.: Sensor-Based Fast Thermal Evaluation Model For Energy Efficient High-Performance Datacenters, Intelligent Sensing and Information Processing, ICISIP 2006. Fourth International Conference on, pp.203-208, 2006.
11. Bash, C.E., Patel, C.D., Sharma, R.K.: Dynamic thermal management of air cooled data centers, Thermal and Thermomechanical Phenomena in Electronics Systems, 2006. ITherm '06. The Tenth Intersociety Conference on, pp.446-452, 2006.
12. Yao J., Guan H., Luo J., Rao L., Liu X.: Adaptive Power Management through Thermal Aware Workload Balancing in Internet Data Centers, Parallel and Distributed Systems, IEEE Transactions on, pp.2400-2409 Vol.26, Issue:9, 2015.
13. Wang X., Wang X., Xing G., Chen J., Lin C.-X., Chen Y.: Intelligent Sensor Placement for Hot Server Detection in Data Centers, Parallel and Distributed Systems, IEEE Transactions on, pp.1577-1588 Vol.24, Issue: 8, 2013.
14. Huang W., Allen-Ware M., Carter J.B., Elnozahy E., Hamann H., Keller T., Lefurgy C., Jian Li, Rajamani K., Rubio J.: TAPO: Thermal-aware power optimization techniques for servers and data centers, 2011 International Green Computing Conference and Workshops (IGCC), Proceeding of, 2011.
15. LAPACK — Linear Algebra PACKage, Project Home Page, available online URL: <http://www.netlib.org/lapack/> (DOA: 2015.09.01).
16. wxWidgets Cross-Platform GUI library, Project Home Page, available online URL: <https://www.wxwidgets.org/> (DOA: 2015.09.01).

AKTYWNE ZARZĄDZANIE PRZEPLYWEM POWIETRZA W CENTRACH DANYCH ŚWIADCZĄCYCH USŁUGI KOŁOKACJI

Centra danych świadczące usługi kolokacji oferują przestrzeń fizyczną i infrastrukturę do działania sprzętu obliczeniowego małych klientów, w odróżnieniu od tradycyjnych centrów danych obsługujących zasoby własnej, większej firmy. Jedną z usług świadczonych klientom jest usługa chłodzenia ich sprzętu. Najczęstszą metodą chłodzenia w centrach danych jest obieg powietrza chłodzącego w serwerowni. Efektywne chłodzenie polega jednak na tym, by doprowadzać chłodne powietrze z centralnych jednostek chłodzących do odpowiedniej lokalizacji w serwerowni, we właściwej ilości, i we właściwym czasie. Stosowanie centralnych jednostek chłodzących bez dodatkowych zabiegów skutkuje jednak powstawaniem w serwerowni miejsc przegrzanych i miejsc przechłodzonych. Proponowane rozwiązanie aktywnego zarządzania powietrzem chłodzącym polega na modelowaniu parametrów w każdej lokalizacji serwerowni na podstawie pomiarów dokonywanych periodycznie w miejscach serwerowni oferujących najbardziej niezawodne wyniki. Stosowany model został opracowany i zweryfikowany poprzez wdrożenie w jednym z przemysłowych centrów danych.

Słowa kluczowe: przepływ powietrza, modelowanie, zarządzanie, efektywność energetyczna.