Airflow Management Using Active Air Dampers in Presence of a Dynamic Workload in Data Centers

Sadegh Khalili [1], Ghazal Mohsenian [1], Anuroop Desu [2], Kanad Ghose [2], Bahgat Sammakia [1]
[1] Departments of Mechanical Engineering, Binghamton University-SUNY, NY, USA
[2] Departments of Mechanical Engineering, Binghamton University-SUNY, NY, USA
E-mail: skhalil6@binghamton.edu

Abstract

The dynamic nature of today's data centers requires active monitoring and holistic management of all aspects of the facility, from the applications to the air conditioning. The most significant aspect of implementing a dynamic data center is the requirement to actively monitor and manage the infrastructure assets. It is vital to ensure information technology (IT) equipment has access to sufficient air (provisioned) at a proper temperature to assure their optimal and continues operation. Hot air recirculation, elevated fan speed, and hot spots are known consequences of an under-provisioned cold aisle. On the other hand, over-provisioning a cold aisle can lead to a significant loss in energy due to bypass of cooling air and leakages. Besides, the number of active servers in an aisle may be varied by load balancers due to short or long-term IT load changes. This demonstrates the need for an active airflow management scheme that is able to respond to airflow demand in different aisles of a data center. In this study, remotely controllable air dampers are implemented to regulate airflow delivery to a cold aisle containment (CAC) during workload changes in a data center. The energy saving opportunities are investigated and practical considerations are discussed.

Keywords

Airflow Management, Active Control, Air Damper, Dynamic Workload, CRAH, Blower, Chiller, PUE, MLC.

Nomenclature

CAC	Cold aisle containment
COP	Coefficient of performance
CRAH	Computer room air handler
IT	Information technology
ITE	Information technology equipment
MLC	Mechanical load component
OAR	Open area ratio
SLB	Smart Load Balancer

1. Introduction

Data centers play a vital role in many aspects of modern life such as banking, entertainment, business, communications, government, security, healthcare, airlines, education, etc. These mission-critical facilities are responsible for storing, processing, and managing huge amounts of data, and have been growing fast to propel the expanding demand for online services. Consequently, the share of data centers energy consumption in the world has grown significantly over the past decades. One of the most common concerns in the data center industry is cooling efficiency. Currently, air-cooling is the most popular technology in data centers due to its proven high reliability, compatibility with the environment and lower initial and maintenance costs. In a large portion of air-cooled data centers, the space underneath a raised floor is pressurized by cold air from computer room air handlers (CRAH) or computer room air conditioners (CRAC) and the supplied cold air is delivered to server racks via perforated floor tiles. In such data centers, cold air can be delivered to where it is needed by simply placing a perforated tile at that location. The fans of the information technology equipment (ITE) are responsible for drawing cold air from the adjacent cold aisles to cool their internal components. The airflow demand of ITE depends on various parameters such as vendor, model, configuration, environmental conditions (e.g. temperature and differential pressure across the ITE) and the workload.

In recent years, implementation of containment systems (either hot or cold aisle containment) has become an important energy-saving strategy to minimize the mixing of hot and cold air. By minimizing the mix of cold and hot air streams, the cooling capacity and thermodynamic efficiency increases. Impacts of hot and cold aisle containment on data center efficiency are discussed extensively in the literature [1–9]. Minimizing the mix of cold and hot air streams allow for a higher temperature setpoint of cooling units which leads to significant savings in chiller plant energy. In addition, elimination of cold air bypass allows for further savings by lowering blower speeds in the cooling units. Based on affinity laws, blower power is proportional to the cube of airflow rate. Therefore, optimizing the airflow rate can significantly decrease fan energy consumption. Ideally, the supplied airflow rate to an aisle should match the ITE airflow demand precisely. However, from a practical perspective, a perfect airflow match may not guarantee the elimination of recirculations. The distribution of supplied air across the containment is often not uniform [8,10,11]. Although the containment self-balances to a large extent, gaps in cabinets and the containment along with local negative/positive pressures in the containment allow for local recirculation. If the supplied airflow rate is too close to the ITE demand in a CAC with significant leakage paths, a portion of cold air would bypass the ITE and leak to the hot aisle (desired leaks), but there could also be leakage into the containment [7,12]. For instance, Khalili et al. [11] showed that the racks farther from the cooling unit in a CAC system may experience recirculation through the under-cabinet gap while bypass of air was observed through the gap below the racks that are closer to the cooling unit. Therefore, slight over-provisioning is recommended in addition to a proper sealing of the containment. The amount of recommended excess air is specific to the installed ITE and the quality of sealing in the containment.

Heterogeneous data centers have become popular as they offer efficient use of available resources, elevated reliability, higher security, and lower cost. These data centers utilize a wide range of networking, servers, and storage purchased from various vendors. Hence, the airflow requirement of the equipment and its response to environmental conditions and workload can be considerably different. Also, the distribution of the IT equipment between aisles can be significantly different in these data centers. In addition, most data center

workload demands exhibit daily or weekly periodic patterns [13–18] in which the demand for services can vary significantly. Furthermore, upgrading ITE may alter the airflow demand of the corresponding aisle. Therefore, airflow demand in different aisles can vary significantly and be a function of time. Heterogeneous data centers usually utilize a common open supply airflow path (raised floor) to provide flexibility to various clients. A disadvantage of traditional raised floors is that cooling units have no control over the destination of the supplied cold air. As a result, some of the aisles with a low airflow demand may receive excess airflow (i.e. is overprovisioned). This is while there are limited permanent ways to balance the airflow rate between aisles such as implementing perforated tiles with different open area ratio (OAR) or adding physical obstructions to the raised floor. Clearly, the above solutions require significant physical work and downtime. Besides, these solutions are semi-permanent and cannot address variation in airflow demand due to a transient workload. Therefore, an active approach for managing and balancing airflow is not only a way for operating such data centers efficiently but eases and reduces the cost of upgrading a data center by eliminating the need for replacing perforated floor tiles.

One approach to manage airflow locally is the deployment of active tiles. An active tile utilizes multiple fans mounted on the anterior side to boost the airflow by increasing air draw from the plenum. Arghode et al. [19] studied the benefits of using active fan tiles for increasing the efficiency in open and enclosed cold aisles. They concluded that implementation of active tiles can improve the uniformity of temperature field in open and contained cold aisles compared to an under-provisioned cold aisle with passive tiles. However, a higher hot air entrainment was observed in presence of active tiles. This higher entrainment can be due to jets of air with high velocities resulting from fans. Active tiles can also be utilized locally to increase the flow rate of tiles that are affected by an underfloor vortex. High-velocity jets of air introduced by air handler units (AHU) or presence of underfloor obstructions can introduce vortices and stagnant regions in a plenum [20,21]. Besides all the benefits, installing active tiles may affect the flow rate of the adjacent tiles due to the suction created by fans of the active tile. In addition, the fans in active tiles are relatively small and consume power which increases the operating costs in data centers. In addition, a large portion of the tile surface is blocked for decreasing fans recirculation and installing large batteries that can provide power during power outages. This increases the resistance on the flow through the tiles when fans are not operated. An alternative approach for local airflow management is deploying air dampers. Remotely controllable air dampers are installed below directional tiles and allow adjusting the flow rate by controlling the angle of the dampers' vanes (see Fig. 1). A directional tile direct airflow toward the face of the adjacent racks and increases rack capture index. In the normal operation mode, the required power for operating the damper is supplied by a power over Ethernet (POE) switch via a network cable which eliminates the need for a separate power cabling. In addition, the electric motors in these dampers only operate for adjusting the angle of vanes when needed. This increases motors' lifetime and allows for smaller batteries. Also, the batteries can last longer in case of a power outage. However, unlike active tiles, the flow rate of individual tiles cannot be boosted by using active dampers.

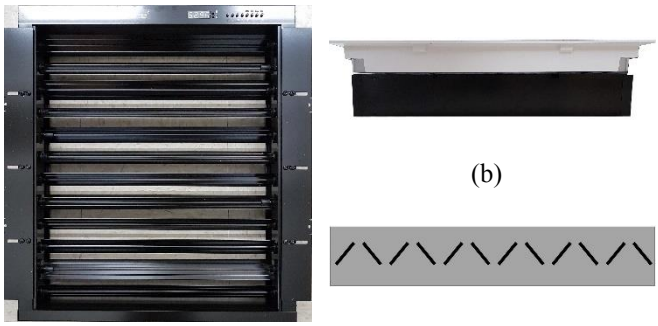

Fig. 1: Air damper: a) top view of the damper, b) side view of the damper mounted under a perforated floor tile, c) schematic of the damper's vanes.

In this study, the impacts of employing remotely controllable dampers on energy efficiency in presence of a dynamic load are investigated in aisle level. Consequent energy savings in blower and chiller powers are calculated. Also investigated are various approaches for turning off non-utilized servers, namely turning off all the servers in preselected racks and turning off multiple servers at the bottom of each rack. In addition, thermal imaging and smoke test are utilized to detect potential locations prone to a recirculation.

2. Modeling of IT Load Variation in Data Centers

Dynamic server provisioning can be employed to lower power consumption and airflow demand, increase system utilization. The basic idea is to consolidate the workload from several servers with low utilization onto fewer servers. This allows saving energy by turning off or enabling sleep state in the servers that are not needed for satisfying the workload requirements. Additional servers can be turned back on when the demand increases [15,16]. In the case of virtualized environments, the unneeded virtual machines can simply be released back to the cloud to save on rental costs. However, the sleep and off states can affect the response time in data centers without a sophisticated workload management system [15,16]. Stachecki and Ghose [15] proposed a Smart Load Balancer (SLB), a fully automated solution, for provisioning IT capacity to match the instantaneous offered workload that results in no performance loss against a baseline design where all servers remain active but operate at a lower CPU utilization. Figure 2 represents the variation of actual utilization trace of a data center operator from low IT load hours to the pick load period. In this trace, the utilization of the system stays below 20% for two-thirds of the duration of the trace. In order to consolidate IT load into a fewer number of serves, the SLB can be used to schedule the jobs. There are two variants of SLB depending on the configuration parameters, namely SLB-LT and SLB. In SLB-LT lower utilization thresholds (40% - 60%) are used to reduce the impact on the tail latency. During the dormant period of the trace, SLB-LT schedules the incoming workload onto 41% of the initial working set of servers and remaining 59% of servers are powered down. The second variant, SLB, uses higher CPU utilization thresholds (50% - 60%) than SLB-LT and schedules the incoming workload to a fewer number of

servers compared to SLB-LT. With the second variant, the results showed that 66% of the servers were not utilized and therefore, could be turned off during the dormant period of the trace.

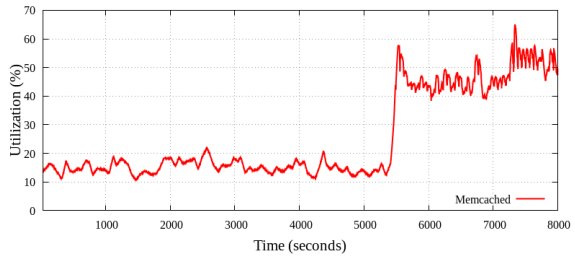

Fig. 2: Experimental CPU utilization as a function of time extracted based on a real data center trace.

A challenge in dealing with the dynamic workload is the possibility of abrupt changes in load. Several instances of load spikes have been documented during important events, such as elections, sporting events, natural disasters, Black Friday shopping, slashdot effects, etc. To handle sudden spikes as well as short-term fluctuations in the load, SLB keeps a fixed number of servers powered on while non-utilized. In the above experimental runs, 14% of the servers had been kept online to mimic avoiding a potential the setup time and violation of Service Level Agreement (SLA). In this study, a more conservative approach is chosen by turning on/off 33% and 47% (instead of 59% and 65%) of the servers. This increases the available capacity further for handling large unpredicted spikes in the load and minimizes a risk of SLA violation. It should be noted that higher energy savings can be achieved by utilizing SLB or SLB-LT in which a larger portion of servers can be powered off.

3. Experimental Setup

The experiments are conducted in ES2-Data Center Laboratory in Binghamton University. The lab is a 215 m^2 (2,315 ft^2) space with a 0.91 m (3 ft) deep raised floor and is equipped with two chilled water-based CRAH units, which are rated at 114 kW (32 tons) of cooling capacity. The nominal maximum airflow rates of the cooling units are 16,500 CFM and 17,500 CFM for CRAH 1 and CRAH 2, respectively. The CRAHs are equipped with variable frequency drives so that supplied airflow rates can be modulated. The layout of the lab is shown in Fig. 3. The IT equipment rows are arranged in an alternating hot aisle/cold aisle arrangement. The cold aisles in Aisles C and D are contained. In this study, air dampers are installed bellow directional tiles in Aisle C. Sixteen racks containing 272 2-RU servers from different vendors and generations are installed in Aisle C. Table 1 shows the list of installed ITE in this aisle and their corresponding airflow demand for an inlet air temperature of 20 °C. The airflow demand of each model was measured by mounting a server of the corresponding model on a flow bench designed in accordance with AMCA 210-99/ASHRAE 51-1999. The details of test procedures are similar to the characterization process described in [22]. It should be mentioned that the number of installed servers in the racks varies from 14 to 20 servers per rack (equipment density at the beginning of the aisle is higher).

Table 1: List of ITE and their airflow demand in Aisle C.

IT Equipment	Require Airflow Rate (cfm)	Total Number	Total Airflow Demand (cfm)
Dell PE R530 - Conf.1	35	79	2765
Dell PE R530 - Conf.2	37	32	1184
Dell PE R520	33	61	2013
Dell PE R730	33	6	198
HP PL DL380/385 G6	34	64	2176
HP PL DL380/385 G7	35	13	455
Dell PE 2950	44	3	132
Dell PE C2100	34	14	476
Network Switch	29	16	464
	Sum:	**286**	**9620**

The tiles in Aisle B and D are covered in this study while the tiles of aisles A and E were left open to provide air for some of the essential networking equipment for operating the data center lab. Three Bapi ZPT-LR pressure sensors with a measurement range of 1″ wc (248.84 Pa) and accuracy of ±0.25% of the range are installed at the top of the racks C1-1, C1-4 and C1-8 (see aisle C in Fig. 3) for monitoring pressure inside the CAC. A flow hood (CFM-850L) is used to measure airflow rate of the tiles. The accuracy of flow hood measurements is analyzed by mounting it on the flow bench and comparing its data with flow rates of the flow bench. It is found that the flow hood shows between 7% to 12% higher flow rates. To minimize the error in airflow measurements, the flow hood data is corrected based on flow bench readings. The maximum discrepancy between the flow bench and flow hood data reduced to 2% via the correction. This allows a more accurate assessment of provisioning level based on the data presented in Table 1 that was measured using the flow bench.

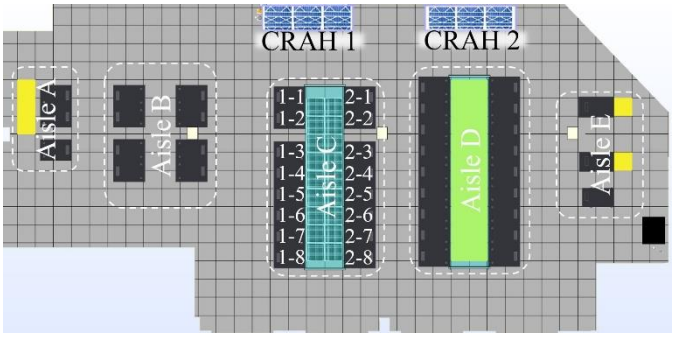

Fig. 3: Layout of the ES2 data center lab in Binghamton University.

A Fuzzy feedback controller is designed to adjust the OAR of the dampers based on pressure measurements in aisle C. Generally, fuzzy controllers are built from a set of if-then rules. An important advantage of fuzzy controllers is that a mathematical model is not necessary [23]. Hence, a fuzzy controller is an ideal choice for controlling complex systems in which a mathematical model of the system is not available or is computationally prohibitive. A CFD modeling of a data center is computationally expensive and a real-time simulation is not feasible with current resources. Also, characterization of

leaks through gaps in the cabinets, containment, and raised floor is cumbersome, therefore, a fuzzy feedback controller well suites this application. The designed controller adjusts the openness of the dampers via Simple Network Management Protocol (SNMP) with various rates based on the error signal. The error signal is defined as the difference between measured pressures and ideal value. For small error signals, the controller changes the openness of the damper with fixed small steps to avoid an overshoot in pressure. For larger error signals, the controller measures the response of the pressure in the CAC by making an arbitrary change in the OAR of the dampers. Subsequently, the controller estimates the next OAR by a linear extrapolation based on the latest response of CAC pressure to the OAR. In this study, the goal is to reach and maintain a neutral provisioning state while 1 Pa over-provisioning of the CAC is allowed (0 Pa $\leq p_{ideal} \leq$ 1 Pa).

4. Test Procedures

The developed fuzzy control system is tested in decreasing and increasing IT load cases in which some of the servers were powered off and powered on, respectively. Based on an analysis described in section 2, implementing SLB-LT and SLB allows turning off 59% and 65% of the servers during the dormant period of the trace. However, a more conservative approach is considered in this paper in which 33% and 47% (instead of 59% and 65%) of the servers are switched off/on. In the first scenario, the impact of under-provisioning Aisle C is visited briefly. Smoke tests and thermal imaging is used for detecting potential areas prone to a recirculation. Consequences of under-provisioning an aisle are explored in numerous studies, so is not discussed in this paper extensively. Interested readers may refer to [10,11,24–27]. In scenario 2, 6 racks (C1-6 to C1-8 and C2-6 to C2-8) are turned on/off to mimic variation of IT airflow demand due to a sudden change in IT load. The response of control system and pressures in the CAC are monitored from the initial steady state (when 33% of the servers were powered off) to the second steady state after powering on all the servers. Next, a decrease in IT load is mimicked by powering off the 33% of the servers and transient response of the system is monitored. A similar test procedure is used in scenario 3 except that 8 racks are powered on/off to mimic the status change in 47% of all the servers. After testing the performance of the control system in scenarios 1 and 2, 8 servers in the bottom half of all the racks are turned on/off to investigate the effect of location of off servers in scenario 4. The results are compared with scenario 3.

Authors noticed a drift in the differential pressure sensors in some of the tests. Thus, frequent calibration may be required for proper operation of the controller. In addition, it should be noted that the response time of the system is also a function of servers' fan speed. However, no significant change in fans speed of active servers is observed in the studied scenarios.

5. Results and Discussions

5.1. Scenario 1

The goal of this scenario is to determine potential leak paths in Aisle C implementing thermal imaging via an infrared (IR) camera and smoke visualization. The IR camera shows the apparent surface temperature of the objects but air temperature. However, in a steady state, it can be assumed that the surface temperatures represent airflow temperature passing over them when heat generation in the surface is not significant. In this scenario, the aisle is slightly under-provisioned to ease detection of recirculation paths. Initially, racks #5 to #8 in rows C1 and C2 (four racks from in each row) were turned off and the aisle was provisioned. After assuring a steady state, racks #5 and #6 in both rows is turned on while the OAR of dampers was fixed. This dropped the CAC pressure to -2 Pa, -1 Pa and -1 Pa at the beginning, middle and end of the aisle, respectively. Figure 4a and 4b show thermal images captured from the end of the aisle after the initial steady state (neutrally provisioned aisle). Slight recirculation is observed in servers of rack C1-8 while the temperatures in rest of the aisle are cool. It should be mentioned that warmer temperatures corresponding to blanking panels are due to the contact of blanking panels with hot air in the hot aisle.

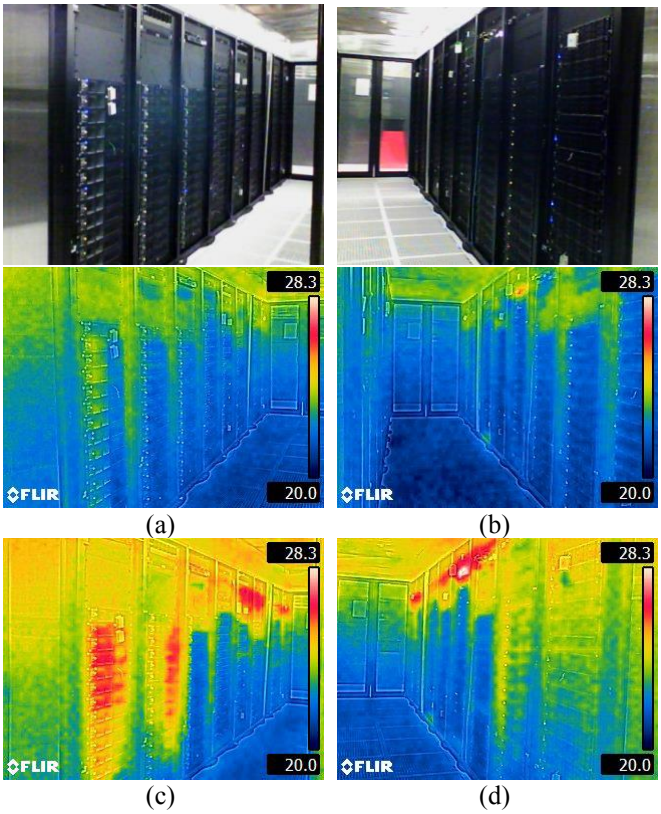

Fig. 4: Thermal images in the neutrally and under-provisioned CAC in scenario 1.

The thermal images in Figs. 4c and 4d are when Aisle C is under-provisioned. High-temperature areas in front of racks # 7 and #8 demonstrate that the powered-off servers are clearly a significant recirculation path although the magnitude of differential pressure across the containment was small. The leaked hot air mixes with cold air supplied by the tiles and slightly elevates the overall CAC temperature. Also, it is seen that surface temperatures increases across the racks' height. The colder surface temperatures at the bottom of the racks are due to the type of the floor tiles which direct airflow towards the face of the racks. The impingement of air to the front panel of servers cause a higher pressure at the front panel of servers that are closer to the tiles. This higher pressure creates a resistance on the recirculation paths through these servers. Figures 4c and 4d also reveal another recirculation path at the

top of the racks. This is where intake ducts connect the intake of rear mounted network switches to the CAC. Further investigation reveals that hot air leaks into the CAC through the brush grommet used between the intake duct and the switch due to the negative pressure in the CAC as shown in Fig. 5a. The total flow rate through the tiles was measured as 5880 cfm and 5986 cfm in the first and second steady states which shows a slight increase while the OAR of the dampers was fixed. The higher tiles flow rate in the second steady state is due to the increased airflow demand in the aisle after turning on the servers and the consequent negative pressure inside the CAC.

(a) (b) (c)

Fig. 5: Smoke visualization reveals the potential recirculation paths: a) recirculation through the intake duct of network switches, b & c) hot air sucked into neighboring servers.

To further investigate the impact of powering off the servers in a rack, another case is studied in which servers at the bottom of all the racks are turned off instead of powering off all the servers in some of the racks. Similar to the previous test in this scenario, it is seen that hot air recirculates through the servers that are powered off is observed. However, smoke tests showed that the leaked hot air was sucked into the active servers above the leak path (see Fig. 5b). This concentrates the impact of hot air recirculation on a few servers that are near the leak path and creates hotspots in a specific height of the racks. Similar behavior was observed when servers at the top of the rack were turned off as shown in Fig. 5c. Overall, the results demonstrate that a slight under-provisioning can cause recirculation through the servers that are powered off and create hot spots inside the containment.

5.2. Scenario 2

In this scenario, the performance of the controller is tested in presence of IT load variation. Figure 6 presents the variation of dampers' OAR and pressures at the beginning, middle and end of Aisle C. Between $0 < t < 10$ minutes, the servers in 6 racks (C1-6 to C1-8 and C2-6 to C2-8) were off and the pressures were within the ideal range. It should be noted that the step changes in pressures are due to the resolution of pressure measurements which is 1 Pa in this study. At $t = 10$ minutes, these racks were turned back on and the controller started to adjust dampers OAR to respond to the consequent negative pressures in the CAC. The controller response varies with respect to the magnitude of under/over-provisioning as seen in **Fig. 6**. For example, when the magnitude of the error signal is equal to or more than 3 Pa, the controller decreases the OAR with larger steps calculated based on the previous response of the system to smaller steps. Figure 7 demonstrates that the recirculation through servers shown in **Fig. 4**c is eliminated by properly provisioning the CAC. The total response time of the system is less than 8 minutes in this part. It should be mentioned that the controller monitors the pressures every 10 seconds in normal operation but sleeps 90 seconds after each change in the OAR of dampers to allow sufficient time for the system to reach to a quasi-steady state before assessing the pressure conditions again. This delay increases the stability of the control system and minimizes the chance of an overshoot due to fan speed changes during set up time of the servers when there is a temporary spike in servers' fan speed.

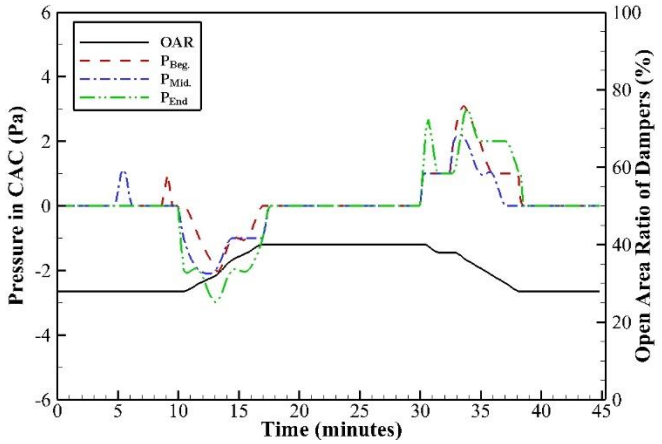

Fig. 6: Variation of pressure and OAR in scenario 2.

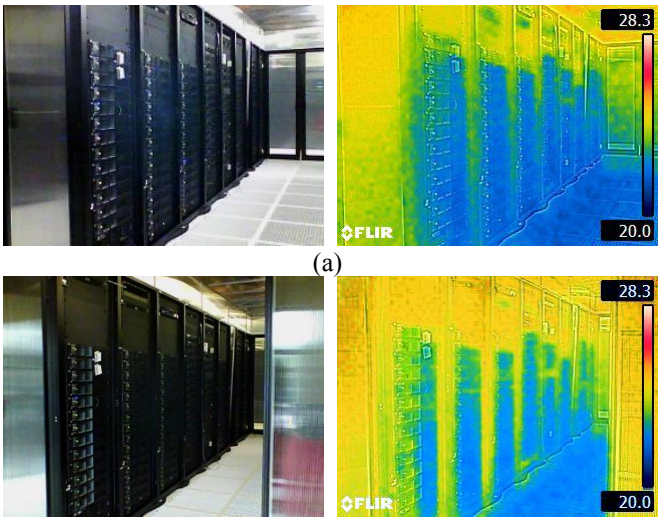

Fig. 7: Thermal image of row C1: a) before and b) after turning on the servers in scenario 2.

After reaching to the second steady state, at $t = 30$ minutes the same 6 racks were turned off to mimic the response of a load balancer to a drop in IT load. **Fig. 6** shows an initial steep change in the pressure by 2 Pa at the end of the aisle (where the racks were turned off). As the pressure exits the ideal range, the controller responds by decreasing the OAR of dampers by 2%. This slight decrease along with the diffusion of excess air in the CAC brings the pressures within the ideal range temporarily. After some time, the excess air builds up in the CAC and

increases the pressure which is followed by the controller response until the system reaches the final steady state of this scenario. Similar to the increasing load case, the overall response time of the system is less than 8 minutes. It should be mentioned that no significant change in the fan speed of active servers was observed in this scenario.

5.3. Scenario 3

The test procedure in this scenario is similar to scenario 2 except 8 racks (C1-5 to C1-8 and C2-5 to C2-8) were turned on/off to mimic a more aggressive load balancing due to increasing and decreasing IT loads. The transient response of the controller and pressures at the beginning, middle and end of the aisle is presented in Fig. 8. As expected, the magnitudes of pressure changes and the OAR of dampers are higher than scenario 2, which is due to the larger the number of servers that participate in handling an IT load. In addition, the system's overall response time is increased to 10 minutes. It is noted that the pressures at the end of the test are at the upper limit of the ideal pressure range which is due to the higher OAR of the dampers at the end of the test (23% compared to 21% at the beginning of the test). **Table 2** compares the measured tile airflow delivery and estimated total airflow demand in the aisle. The estimated airflow demands are calculated based on data presented in Table 1. The difference between the supplied flow rate through the tiles and ITE demand is presented in the "Difference" row of this table. The differences can be due to the bypass of airflow through containment gaps and the servers that are powered off, as well as uncertainty in the pressure and flow rate measurements. The overall uncertainty of the measurements does not exceed 6% of the measurements. This means the primary contributor to the "Difference" row in **Table 2** is the bypass of airflow through the servers. Also, this difference is higher in the final steady state which is due to the higher pressure in the CAC in the final steady state of this scenario. Figure 9 shows that that aisle is properly provisioned and no recirculation is detected.

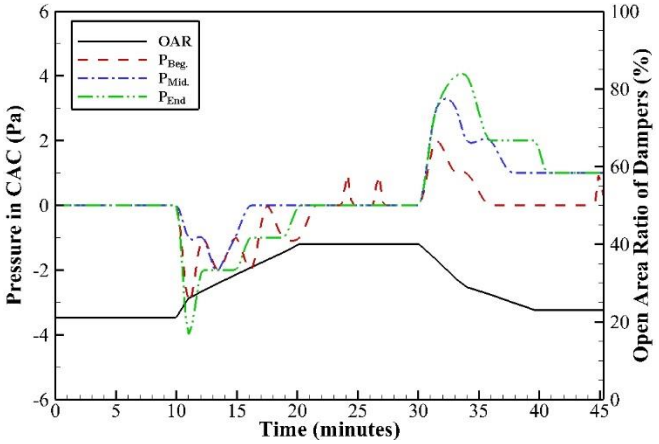

Fig. 8: Variation of pressure and OAR in scenario 3.

Table 2: Estimated and Measured flow rates in scenario 3

	Initial S.S.	Second S.S.	Final S.S.
Tiles' Flow rate	5935 cfm	9927 cfm	6164 cfm
Est. airflow demand	5430 cfm	9620 cfm	5430 cfm
Difference	505 cfm	307 cfm	734 cfm
% Overprovisioning	9%	3%	13%

(a) (b)

Fig. 9: Thermal image of row C1 after turning on the servers (the second steady state) in scenario 3.

5.4. Scenario 4

In this scenario, 8 servers at the bottom of all the racks of Aisle C are turned off (47% of all the servers). The controller was able to manage airflow delivery to the CAC successfully. The results were similar to scenario 3 except that OAR of dampers in the initial and final steady states was slightly higher. Graphs are not presented here for the sake of space. **Table 3** presents the tiles flow rate in the initial and the second steady states. By comparing tiles' flow rate between the initial steady states in tables 2 and 3, it is perceived that turning off the servers at the bottom of the racks increases air bypass through these servers. As mentioned earlier, the installed tiles direct airflow toward the face of the racks which creates a higher pressure at the inlet of the servers. This effect is stronger at lower elevations where servers are closer to the tiles, i.e. servers at the bottom of the racks. Therefore, cold air bypass through the powered-off servers in this scenario is more than scenario 3. servers.

Table 3: Estimated and Measured flow rates in scenario 4

	Initial S.S.	Second S.S.
Tiles' Flow rate	6424 cfm	9950 cfm
Est. airflow demand	5495 cfm	9620 cfm
Difference	929 cfm	330 cfm
% Overprovisioning	16%	3%

6. Energy Saving Opportunities

Iyengar and Schmidt [28] presented an analytical model for thermodynamic characterization of cooling systems in data centers using experimental data and empirical equations. In a case study, they found the chilled water air handler as the second largest energy drain on the cooling system after the chiller (see **Fig. 10**). Patterson et al. [12] reported a similar observation.

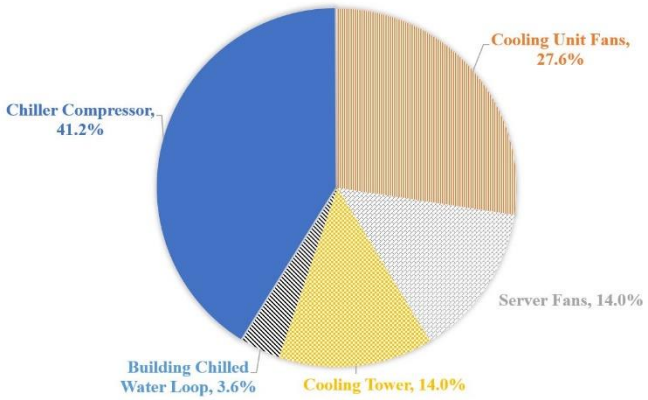

Fig. 10: Cooling energy breakdown for the case study example in [28].

Comparing the tiles' flow rates in **Table 2** shows that air delivery to Aisle C can be decreased by approximately 40% via implementing the controller. The most significant energy saving associated with a decrease in airflow delivery to an aisle comes from adjusting the fans speed of cooling units. The energy required for moving air by a fan is proportional to the cube of the airflow rate.

$$\frac{P_2}{P_1} = \left(\frac{Q_2}{Q_1}\right)^3 \quad (1)$$

According to affinity laws (1), cutting airflow rate by 40% will reduce CRAH's blower power by 80% approximately. However, many data centers do not dedicate a cooling unit per aisle. The cooling units are usually oversized to serve multiple aisles and benefit from higher efficiency larger blowers or provide cooling redundancy in data centers.in addition, a portion of the supplied air bypasses through raised floor gaps and do not reach the aisles. In practice, the percent of the cut in a CRAH's flow rate is smaller than the cut in an aisle's airflow demand. Therefore, the actual power saving is expected to be less than 80%. The actual saving can be estimated by making some assumptions. To have an accurate estimate, raised floor leaks should also be taken into account. Amount of air bypass through raised floor gaps depends mainly on the pressure of the plenum. A reasonable assumption is that this leakage remains fixed if the pressure in the plenum is controlled via adjusting blower speed of the cooling units. As the OAR of tiles in aisle A and E are fixed, it can be assumed that flow rates of these tiles do not change if the plenum pressure is fixed. Maintaining a fixed pressure in the plenum also minimize the impact of changing OAR of dampers in an aisle on the airflow delivery to other aisles. However, maintaining a fixed pressure in the plenum require a separate controller that controls the fan speed of cooling units, and is out of the scope of this paper. The required CRAH's flow rate can be estimated by

$$Q_{CRAH} = Q_{Aisle\ C} + Q_{Aisle\ A} + Q_{Aisle\ E} + Q_{R.F.\ Leaks} \quad (2)$$

in which $Q_{R.F.\ Leaks}$, $Q_{Aisle\ A}$ and $Q_{Aisle\ E}$ remains constant when plenum pressure is fixed. Using (2), the flow rate of leaks through raised flow ($Q_{R.F.\ Leaks}$) can be calculated. Following above assumptions, required supply airflow after turning off servers is calculated which shows a 25% decrease compared to when all the servers were powered on. This decrease in CRAH's airflow translates to a 42% drop in CRAH's power.

Another significant benefit of airflow management is in the chiller plant. In ES2 Data Center Lab, the chilled water is supplied from the building's chiller. Thus, a direct measurement of savings was not possible. Instead, typical coefficient of performance (COP) values can be used to estimate the compressor work. The chiller power can be calculated via (3) [29]:

$$P_{Chiller} = \frac{\text{Removed Heat}}{COP_{Chiller}} \quad (3)$$

Turning the servers off and consolidating workload by the load balancer decreases the number of idling servers but increases the power of active servers. However, there is a logarithmic relation between CPU utilization and a server power. Therefore, operating in a higher CPU utilization increases the efficiency of a server. The power consumption measurements in different scenarios of the study showed the overall IT power decreases by consolidating workload and turning off excess servers. The rise in air temperature supplied to aisle C due to IT heat dissipation can be calculated using heat balance:

$$\Delta T_{\text{air}} = \frac{\text{IT Heat Dissipation}}{\dot{m}_{air} \cdot c_p} \quad (4)$$

Assuming negligible change in air properties, equation (4) gives an estimate on the return air temperature to the CRAH unit for a fixed supply air temperature of the CRAH unit. Powering off the non-needed servers decreases the total IT power and airflow demand in the aisle. Two cases are compared in Table 4. In case 1, all the servers in Aisle C are active and operate at 15% CPU utilization. Case 2 represents Aisle C with active airflow management and a smart load balancer. Table 4 shows that despite the decrease in the total heat dissipated by ITE, the overall temperature rise is higher in case 2. This is due to the cut in supplied airflow rate to the aisle and higher workload on the active servers. A relatively high rack air temperature rise leads to a more efficient heat transfer in the CRAH's coils and higher fluid temperature returning to the chiller evaporator which improves overall cooling performance and increases the COP of the chiller. Breen et al. [30] concluded that COP gain due to an increase rack air temperature rise is considerable when this rise is below 12.5 ℃ which is the case in Table 4. In this study, a COP of 3.0 is assumed for the base case and an increase of 0.2 in COP per ℃ increase in rack air temperature rise. Implementing above assumptions, the chiller power can be calculated using (3). The results show a 36% saving in the chiller power consumption.

Table 4: Comparison of IT power in without and with implementing controller and load balancing for the trace presented in Fig. 2.

Case #	CPU Util.	# of Active Servers	IT power (kW)	Airflow Rate (cfm)	ΔT_{air} (℃)	COP_R	Chiller Power (kW)
1	15%	271	44.8	9927	8.1	3	14.9
2	65%	161	30.4	5935	9.2	3.2	9.5

The power usage effectiveness (PUE) corresponding to cases 1 and 2 in Table 4 can be calculated via (5)

$$\text{PUE} = \frac{1}{\text{DCiE}} = \frac{\text{Total Facility Energy}}{\text{ITE Energy}} \quad (5)$$

where total facility energy includes power used by chiller, pumps, blowers, fans, and lighting. A direct measurement of pumps power associated with ES2 data center lab was not feasible because the chiller pump is shared with the building. However, the power consumed by pumps is negligible compared to chiller and blower powers, and therefore, is not considered in the calculations. Table 5 shows a 6% improvement in PUE in case 2 compared to case 1. Recently, mechanical load component (MLC) is introduced by ASHRAE as an improved measure for the mechanical efficiency in data centers. The MLC is defined as the sum of all cooling, fan, pump, and heat rejection power divided by the data center ITE power. Following the above assumptions, MLC is calculated for both cases which shows a 17% improvement in case 2.

Table 5: PUE analysis.

Case #	Blower Power (kW)	Chiller Power (kW)	IT power (kW)	Lighting (kW)	PUE	MLC
1	8.2	14.9	44.8	0.2	1.52	0.52
2	3.45	9.5	30.4	0.2	1.43	0.43

7. Conclusions and Next Steps

In recent years, airflow management has gained significant interests in the data center industry to mitigate increasing energy costs. According to various surveys, IT equipment is the most significant power consumer in data centers followed by chiller plants and air handling units. A wise and CPU aware load balancing scheme can consolidate IT load into a fewer number of servers and put the non-utilized servers into sleep mode or powered them off. This can significantly decrease IT power consumption and chiller power. However, airflow delivery to the aisle is often not adjusted in the data centers which causes over/under-provisioning in some of the aisles.

In this paper, remotely controllable air dampers are utilized to ensure just enough air is supplied to an aisle at different airflow demands due to variation in IT load. The differential pressure between the CAC and room is considered as the control parameter. This allows designing a control system which is independent of the model, generation, layout, workload and the number of installed ITE and also eliminates the need for characterization of the servers. A fuzzy controller is designed to regulate the OAR of dampers based on the pressure measurements in the CAC. This controller ensures just enough airflow is delivered to the aisle with a small over-provisioning margin. The performance of the controller is tested for various load balancing schemes. The designed controller successfully regulated CAC pressure in various cases by adjusting supplied air flow to the aisle via controlling OAR of dampers, and thus ensured proper provisioning of ITE in both decreasing and increasing IT load scenarios. It is shown that powering off the servers near the directional tiles (at the bottom half of the racks) can increase cold air bypass. It is shown that implementing the proposed airflow management scheme along with a smart load balancer in Aisle C of ES2 Data Center Lab can save 42% and 36% in the CRAH blower and the chiller powers, respectively. Also, it is shown that utilizing a load balancer without assuring a proper provisioning of an aisle can lead to recirculation through the servers and create hot spots.

In this study, the controller is tested in Aisle C of ES2 Data Center Lab successfully. However, a potential challenge can be the implementation of the controller in multiple aisles which are connected through a common plenum. In a multi-aisle data center with a common raised floor, adjusting OAR of dampers in one aisle can affect air delivery to other aisles. As a result, OAR of dampers in other aisles can vary to respond to this change. This can cause a cascade change in OARs which increases the response time of the control system. An important parameter in such airflow management systems is the ability to maintain a fixed plenum pressure by adjusting blower speed of cooling units and allowing sufficient tolerance for the ideal pressure range. The step would be designing a holistic control approach which manages blower speed of cooling units based on various pertinent parameters including IT load, number of active servers, plenum pressure, etc.

Acknowledgments

The authors would like to thank Russ Tipton and Arash Golafshan from Vertiv and Mark Seymour from Future Facilities for their support and advice. This work is supported by NSF IUCRC Award No. IIP-1738793 and MRI Award No. CNS1040666.

References

[1] Sundaralingam, V., Arghode, V. K., Joshi, Y., and Phelps, W., 2014, "Experimental Characterization of Various Cold Aisle Containment Configurations for Data Centers," J. Electron. Packag., **137**(1), p. 11007.

[2] Nemati, K., Alissa, H. A., Murray, B. T., and Sammakia, B., 2016, "Steady-State and Transient Comparison of Cold and Hot Aisle Containment and Chimney," *In Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 15th IEEE Intersociety Conference On*, pp. 1435–1443.

[3] Bharath, M., Shrivastava, S. K., Ibrahim, M., Alkharabsheh, S. A., and Sammakia, B. G., 2013, "Impact of Cold Aisle Containment on Thermal Performance of Data Center," *InterPACK2013-73201*, pp. 1–5.

[4] Niemann, J., Brown, K., and Avelar, V., 2011, "Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency," Schneider Electr. Data Cent. Sci. Center, White Pap., **135**, pp. 1–14.

[5] Alissa, H. A., Nemati, K., Sammakia, B. G., Schneebeli, K., Schmidt, R. R., and Seymour, M. J., 2016, "Chip to Facility Ramifications of Containment Solution on IT Airflow and Uptime," IEEE Trans. Components, Packag. Manuf. Technol., **6**(1), pp. 67–78.

[6] Khalili, S., Alissa, H., Desu, A., Sammakia, B., and Ghose, K., 2018, "An Experimental Analysis of Hot Aisle Containment Systems," *Proceedings of the 17th InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, ITherm 2018*, San Diego, CA USA, pp. 748–760.

[7] Makwana, Y. U., Calder, A. R., and Shrivastava, S. K., 2014, "Benefits of Properly Sealing a Cold Aisle Containment System," Thermomechanical Phenom. Electron. Syst. - Proceedings Intersoc. Conf. ITherm2014, pp. 793–797.

[8] Patterson, M. K., Weidmann, R., Leberecht, M., Mair, M., and Libby, R. M., 2011, "An Investigation Into Cooling System Control Strategies for Data Center Airflow Containment Architectures," pp. 479–488.

[9] Shrivastava, S. K., and Ibrahim, M., 2013, "Benefit of Cold Aisle Containment During Cooling Failure," ASME 2013 Int. Tech. Conf. Exhib. Packag. Integr. Electron. Photonic Microsystems, (55768), p. V002T09A021.

[10] Alissa, H. A., Nemati, K., Sammakia, B. G., Seymour, M. J., Tipton, R., Mendo, D., Demetriou, D. W., Schneebeli, K., Tipton, R., Mendo, D., Demetriou, D. W., and Schneebeli, K., 2016, "Chip to Chiller Experimental Cooling Failure Analysis of Data Centers: The Interaction Between IT and Facility," IEEE Trans. Components, Packag. Manuf. Technol., **6**(9), pp. 1361–1378.

[11] Khalili, S., Tradat, M., Nemati, K., Seymour, M., and Sammakia, B., 2018, "Impact of Tile Design on the Thermal Performance of Open and Enclosed Aisles," J. Electron. Packag., **140(1)**, pp. 010907-010907-12.

[12] Patterson, M. K., Weidmann, R., Leberecht, M., Mair, M., and Libby, R. M., 2011, "An Investigation Into Cooling System Control Strategies for Data Center Airflow Containment Architectures," (44625), pp. 479–488.

[13] Gmach, D., Rolia, J., Cherkasova, L., and Kemper, A., 2009, "Resource Pool Management: Reactive versus Proactive or Let's Be Friends," Comput. Networks, **53**(17), pp. 2905–2922.

[14] Arlitt, M., and Jin, T., 2000, "A Workload Characterization Study of the 1998 World Cup Web Site," IEEE Netw., **14**(3), pp. 30–37.

[15] Stachecki, T. J., and Ghose, K., 2015, "Short-Term Load Prediction and Energy-Aware Load Balancing for Data Centers Serving Online Requests*."

[16] Gandhi, A., 2013, "Dynamic Server Provisioning for Data Center Power Management," PhD diss., Intel, (June), pp. 1–174.

[17] Atikoglu, B., Xu, Y., Frachtenberg, E., Jiang, S., and Paleczny, M., 2012, "Workload Analysis of a Large-Scale Key-Value Store," *ACM SIGMETRICS Performance Evaluation Review*, ACM, pp. 53–64.

[18] Hazelwood, K., Bird, S., Brooks, D., Chintala, S., Diril, U., Dzhulgakov, D., Fawzy, M., Jia, B., Jia, Y., Kalro, A., Law, J., Lee, K., Lu, J., Noordhuis, P., Smelyanskiy, M., Xiong, L., and Wang, X., 2018, "Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective," *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 620–629.

[19] Arghode, V. K., Sundaralingam, V., and Joshi, Y., 2016, "Airflow Management in a Contained Cold Aisle Using Active Fan Tiles for Energy Efficient Data- Center Operation Airflow Management in a Contained Cold Aisle Using Active Fan Tiles for Energy Efficient Data-Center," **7632**(October 2017).

[20] Alissa, H. A., Nemati, K., Sammakia, B., Ghose, K., Seymour, M., and Schmidt, R., 2015, "Innovative Approaches of Experimentally Guided CFD Modeling for Data Centers," Therm. Meas. Model. Manag. Symp. (SEMI-THERM), 2015 31st, pp. 176–184.

[21] Chen, K., Federspiel, C. C., Auslander, D. M., Bash, C. E., and Patel, C. D., 2006, "Control Strategies for Plenum Optimization in Raised Floor Data Centers," HP Lab. White Pap.

[22] Khalili, S., Alissa, H., Nemati, K., Seymour, M., Curtis, R., Moss, D., and Sammakia, B., 2018, "Impact of Internal Design on the Efficiency Of IT Equipment In A Hot Aisle Containment System - An Experimental Study," *ASME 2018 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems Collocated with the ASME 2018 Conference on Information Storage and Processing Systems*, San Francisco, CA, p. In press.

[23] Zadeh, L. A., 1973, "Outline of a New Approach to the Analysis of Complex Systems and Decision Processes," Syst. Man Cybern. IEEE Trans., (1), pp. 28–44.

[24] Athavale, J., Joshi, Y., Yoda, M., and Phelps, W., "Impact of Active Tiles on Data Center Flow and Temperature Distribution."

[25] Alissa, H. A., Nemati, K., Puvvadi, U. L. N., Sammakia, B. G., Schneebeli, K., Seymour, M., and Gregory, T., 2016, "Analysis of Airflow Imbalances in an Open Compute High Density Storage Data Center," Appl. Therm. Eng., **108**, pp. 937–950.

[26] Alkharabsheh, S. A., Sammakia, B. G., and Shrivastava, S. K., 2015, "Experimentally Validated Computational Fluid Dynamics Model for a Data Center With Cold Aisle Containment," J. Electron. Packag., **137**(2), pp. 21010–21019.

[27] Breen, T. J., Walsh, E. J., Bash, C. E., Punch, J., Shah, A. J., Bash, C. E., Rubenstein, B., Heath, S., and Kumari, N., 2011, "From Chip to Cooling Tower Data Center Modeling: Influence of Air-Stream Containment on Operating Efficiency," *ASME/JSME Thermal Engineering Joint Conference*, pp. T10074-T10074-10.

[28] Iyengar, M., and Schmidt, R., 2009, "Analytical Modeling for Thermodynamic Characterization of Data Center Cooling Systems," J. Electron. Packag., **131**(2), p. 021009.

[29] Cengel, Y. A., and Boles, M. A., 2002, "Thermodynamics: An Engineering Approach," Sea, **1000**, p. 8862.

[30] Breen, T. J., Walsh, E. J., Punch, J., Shah, A. J., and Bash, C. E., 2010, "From Chip to Cooling Tower Data Center Modeling: Part I Influence of Server Inlet Temperature and Temperature Rise across Cabinet," *2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, pp. 1–10.